Radar & Applications Course (RAC)



Presented by

The Warning Decision Training Division

TAB

Radar & Applications Course

Topic:

Principles of Radar

Table of Contents
Topic: Principles of Radar

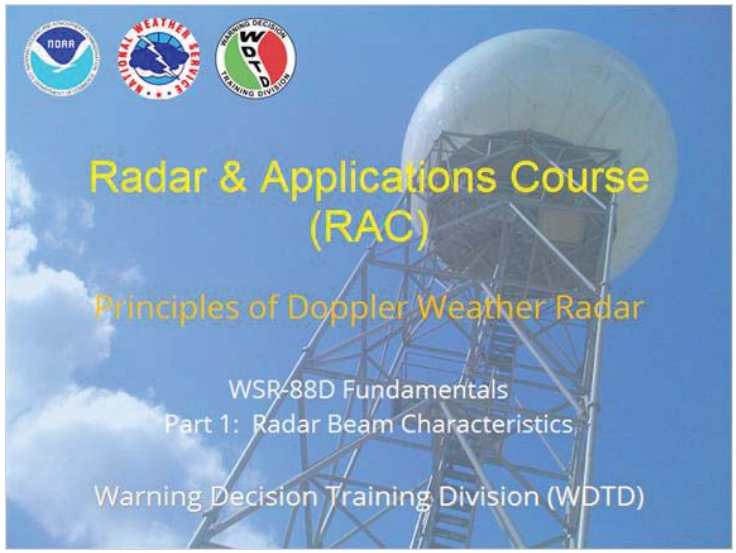
Click to jump to lesson

Lesson 1	**WSR-88D Fundamentals Part 1: Radar Fundamentals**
Lesson 2	**WSR-88D Fundamentals Part 2: Weather Radar Equation**
Lesson 3	**WSR-88D Fundamentals Part 3: Transmitting & Receiving Characteristics**
Lesson 4	**WSR-88D Fundamentals Part 4: Non-Standard Beam Consequences**
Lesson 5	**WSR-88D Fundamentals Part 5: Data Collection**
Lesson 6	**Radar Sampling Issues**
Lesson 7	**VCP Selection**
Lesson 8	**Dynamic Sampling**
Lesson 9	**RPG HCI Controls**
Lesson 10	**RPG HCI Functions**
Lesson 11	**Legacy Base Data Generation**
Lesson 12	**Super Resolution Base Data Generation**
Lesson 13	**Dual-Pol Base Data Generation**
Lesson 14	**Clutter Filtering**
Lesson 15	**Range Unfolding**
Lesson 16	**Velocity Dealiasing**
Lesson 17	**Recombination at the RPG**
Lesson 18	**Dual-Pol Base Data Preprocessing at the RPG**
Lesson 19	**WSR-88D Base Data Quality**
Lesson 20	**Radar Rainfall Estimation Issues**
Lesson 21	**The Legacy Precipitation Processing Subsystem (PPS) Algorithm**
Lesson 22	**The Dual-Pol Quantitative Precipitation Estimation (QPE) Algorithm**
Lesson 23	**PPS and QPE Comparison**
Lesson 24	**Snow Accumulation Algorithm**

WSR-88D Fundamentals Part 1: Radar Beam Characteristics

1. Intro to Radar Beam Characteristics

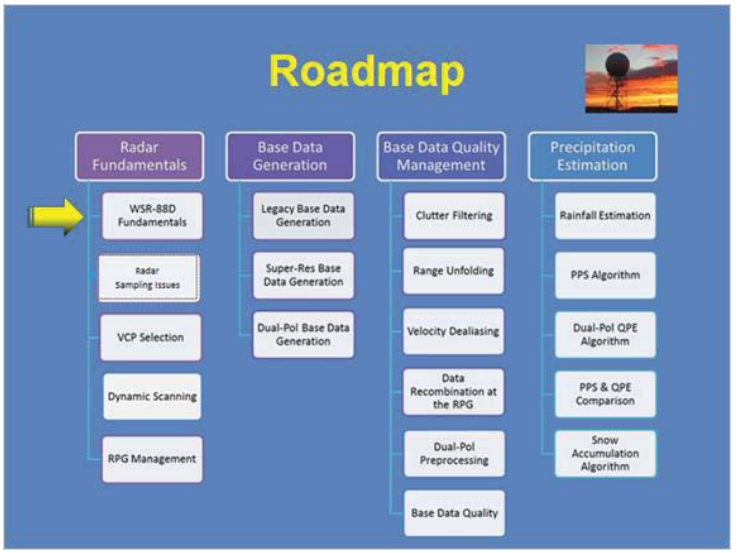
1.1 Welcome



Notes:

Welcome to the Radar & Applications Course (RAC) Principles of Doppler Weather Radar. This lesson is on WSR-88D Fundamentals Part 1: Radar Beam Characteristics. It is presented by the Warning Decision Training Division (WDTD). Let's get started!

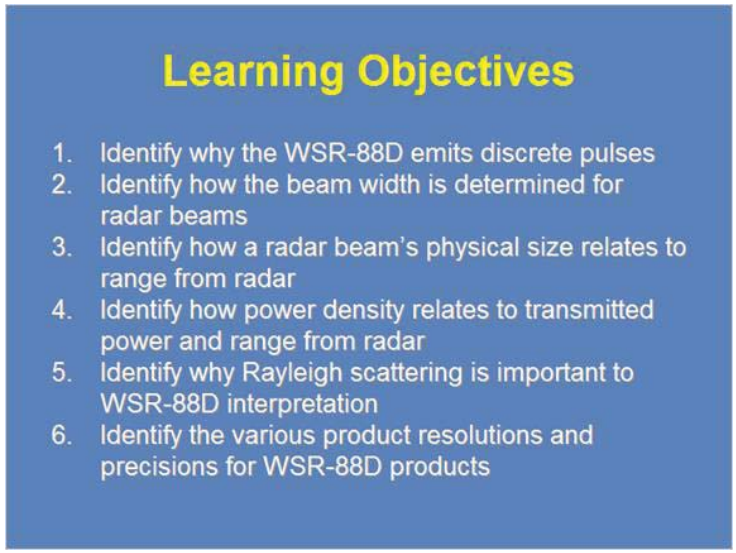
1.2 Roadmap



Notes:

Here is the complete roadmap for the entire "Principles of Doppler Weather Radar" section of RAC. You are currently in the WSR-88D Fundamentals portion of this section, and this portion consists of 5 lessons. Let's keep going!

1.3 Learning Objectives

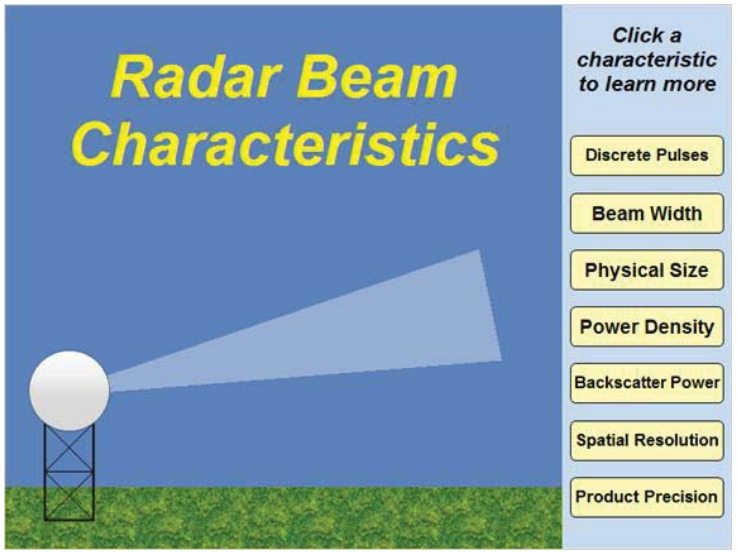


Notes:

Here are the learning objectives for this lesson. Please take a moment to review these objectives, as the quiz at the end of this lesson is based on these objectives.

2. Main Radar Graphics

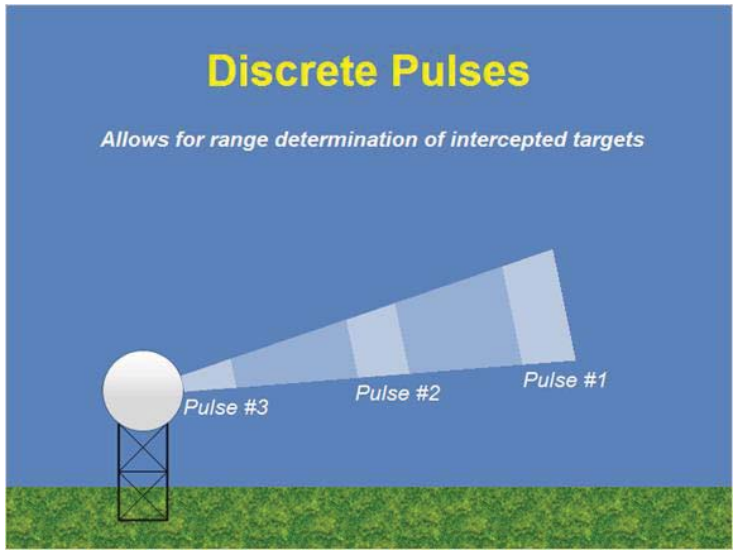
2.1 Radar Beam Characteristics HOME



Notes:

The radar beam has many different characteristics which help determine the information you will see as a warning forecaster. These various characteristics include: discrete pulses, beam width, physical size, power density, backscatter power, spatial resolution, and product resolution. Click on each of these characteristics to learn more. Once you have completed viewing each of these characteristics, you will be ready for the quiz to test your understanding.

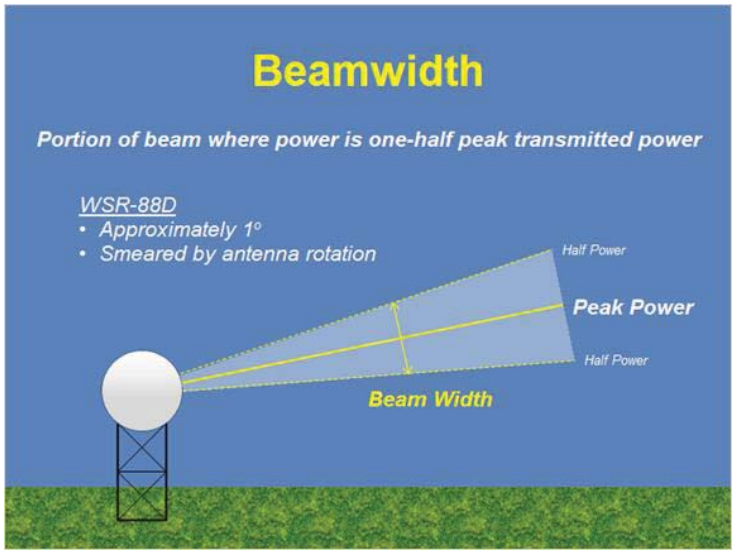
2.2 Discrete Pulses



Notes:

The WSR-88D emits pulses of energy into the atmosphere at pre-defined intervals. This discrete pulse mode allows for the radar signal processor to determine the range of the intercepted targets.

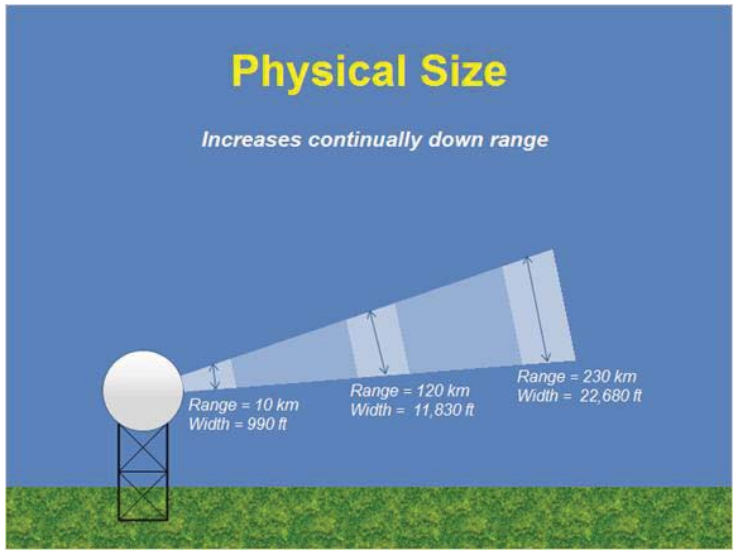
2.3 Beamwidth



Notes:

The radar generates the electromagnetic radiation at the transmitter and the antenna focuses this radiation into a beam that is then reflected into the atmosphere toward the intended targets. Since the radiation doesn't magically have boundaries, the width of the beam is defined as the point at which the power along the beam reaches one-half the peak transmitted power. For the WSR-88D, this width is approximately 1 degree, but is somewhat broader due to the rotation of the antenna while the beam is being transmitted which is referred to as the effective beam width.

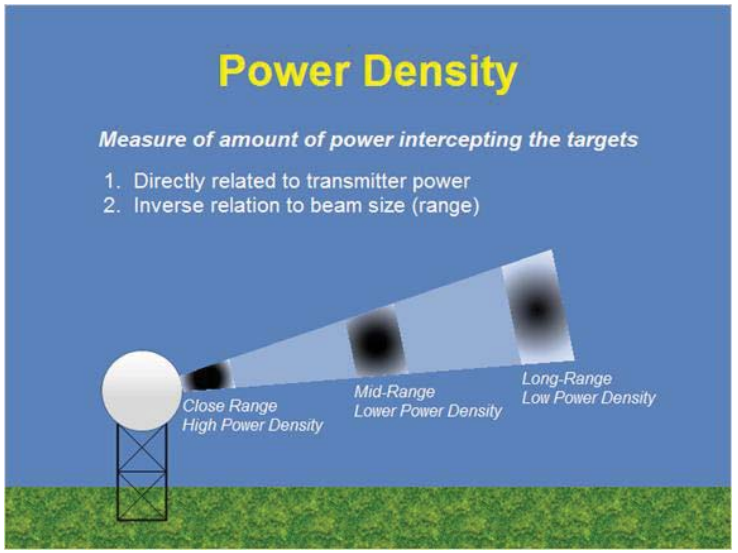
2.4 Physical Size



Notes:

While the beam width does not change down range from the radar, the physical size does increase quite dramatically down range. Let's quickly look at how big the WSR-88D pulse gets as we approach the far reaches of the WSR-88D range. At 10 km from the radar, the width of the beam is already 990 feet. That's almost 3 football fields! When the beam gets to 120 km range, the width is over 2 miles! And, when the beam gets to the outer edges of the first trip (which we'll define later), the beam is over 4 miles. So, targets within one radar beam can be as far as 4 miles apart!

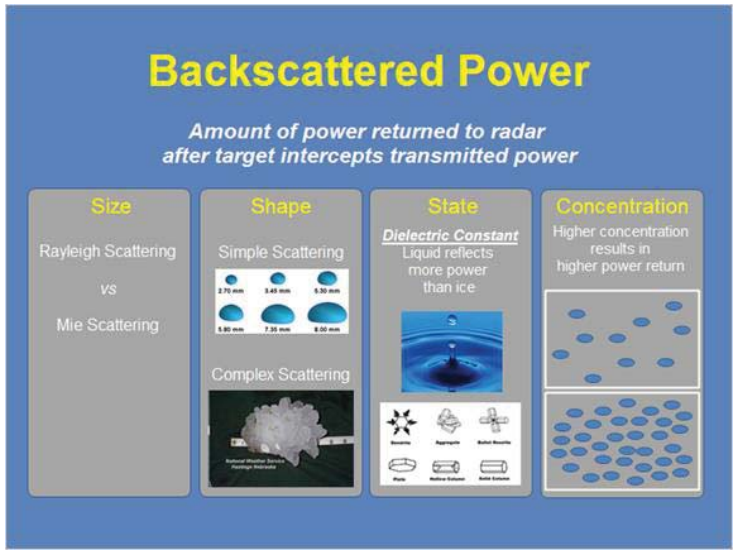
2.5 Power Density



Notes:

The purpose of sending out a pulse of radiation is to have weather targets intercept this energy and reflect some of it back to the radar so we can determine the intensity of the weather targets. The amount of energy that intercepts these weather targets is called the power density, and it is dependent upon two factors: transmitter power and beam size. First power density will increase with increasing transmitter power. However, as the beam goes down range, remember it increases in size, so that same amount of power is spread across a larger area, so the power density actually decreases as the beam goes down range, but it's not equally distributed across the beam. The beam centerline contains most of that power.

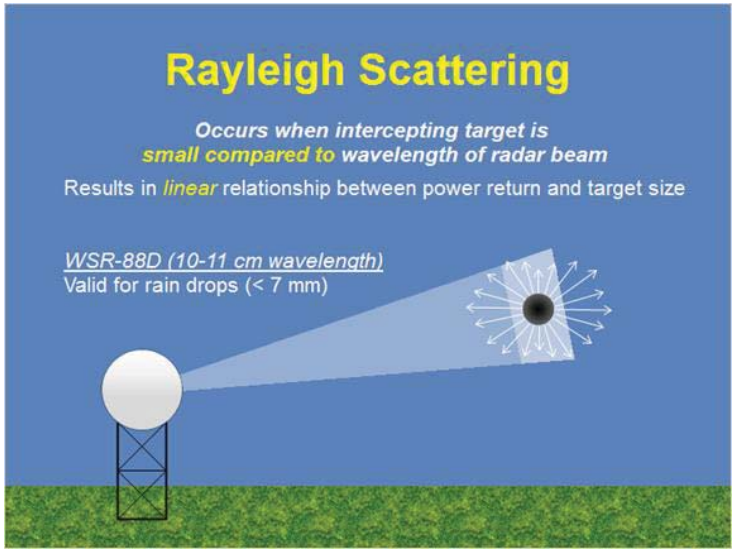
2.6 Backscattered Power



Notes:

Once the radar intercepts some of the transmitted power, it reflects some of the power back to the radar. The amount of power returned back to the radar is referred to as "backscattered power" and is dependent upon 4 characteristics of the weather target. This includes size (which we'll discuss shortly). The next is shape. Fairly smooth-shaped objects like rain drops will have simple scattering properties, but spiky hail stones will have complex scattering properties. The next is the state of the weather target. We'll also refer to this as the dielectric constant which basically means "how reflective is the weather target". Liquid reflects radar energy much more effectively than does ice. Therefore, for the same exact size/shape, water will return a significantly higher amount of power than will the ice particle. Finally, we have concentration. Within a given beam, if there only exists a few rain drops, for example, the power returned will be fairly low. However, if you take those same rain drops and increase the amount by, let's say, triple…the amount of power returned will increase. So, these are the four factors affecting backscattered power, but let's dive a little deeper into the relationship between size and power return.

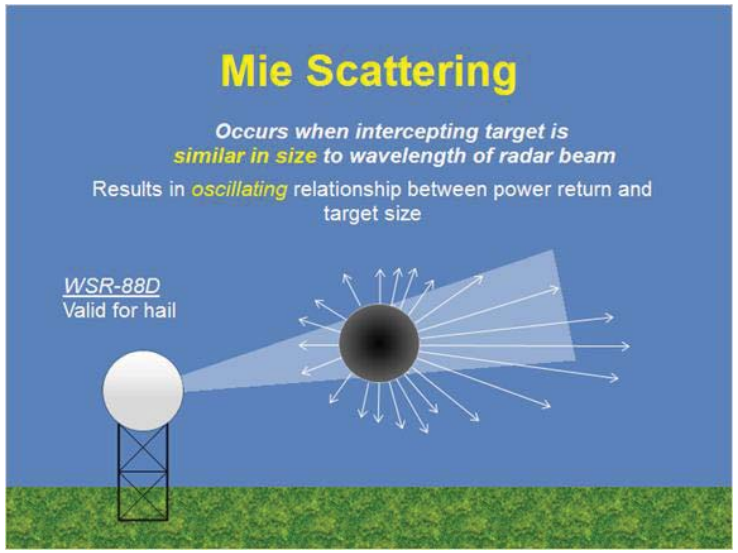
2.7 Rayleigh Scattering



Notes:

The size of the weather target compared to the wavelength of the radar determines the scattering regime. For targets small compared to the wavelength of the radar beam, the scattering is fairly uniform in all directions and the amount of power reflected increases linearly with increasing size. This type of scattering is called Rayleigh scattering and it is preferred because the linear relationship between power return and target size is very helpful (as we'll find out later). The wavelength of the WSR-88D is approximately 10-11 cm. Therefore, Rayleigh scattering is a good assumption for all weather targets smaller than 7 mm which includes practically all rain, but does not include hail. This basically means for most weather objects, the power returned is linearly related to the size of the weather target.

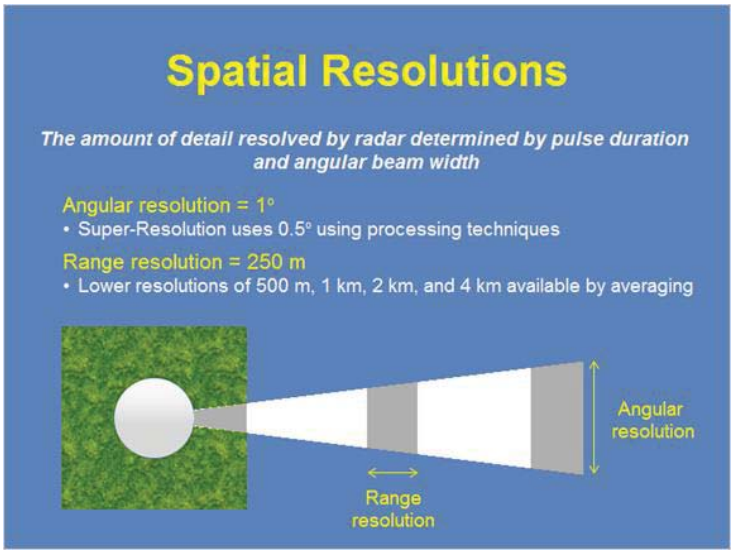
2.8 Mie Scattering



Notes:

When weather targets become roughly similar in size to the wavelength of the radar beam or larger, the scattering properties are no longer linearly related, and most energy is forward scattered. The energy that is reflected back to the radar does not have a linear relationship to size, but rather an oscillating relationship. This is called Mie scattering. So, relating size to power return is not so straightforward. For the WSR-88D, hail is the primary target that falls within the Mie scattering regime.

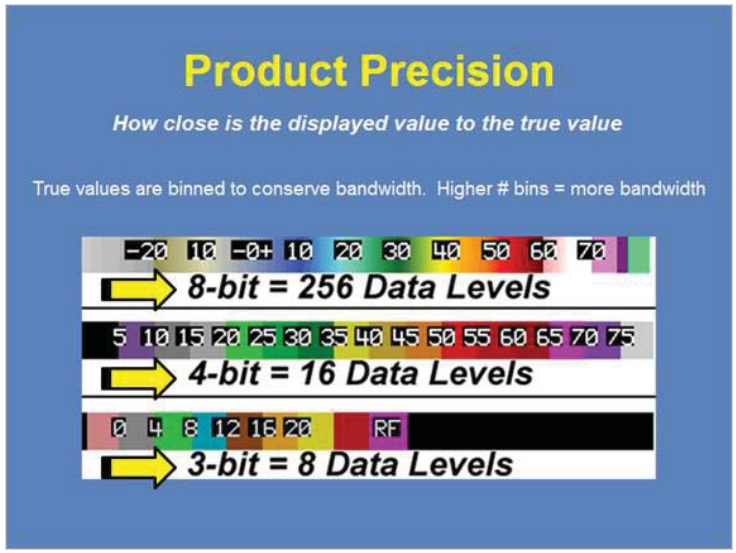
2.9 Spatial Resolutions



Notes:

So, how much detail can we see with data collected by the WSR-88D? That is all dependent upon the pulse duration and angular beam width. Since we already know the beam width is 1 degree, that is the best angular resolution we can get without any processing of the data. I say "without any processing" because recently, radar engineers came up with a processing technique to display 0.5 degree angular resolution data, which is called "super-resolution". This technique will be covered in a later lesson in this section, but for now, just know it is possible to get 0.5 degree angular resolution with the WSR-88D. As for the range resolution, this is determined by the pulse duration (or how long does the pulse transmit). The longer the pulse duration, the coarser the range resolution. For the WSR-88D, the best range resolution possible is 250 m. Other coarser resolutions are available by averaging the individual 250-m bins.

2.10 Product Precision



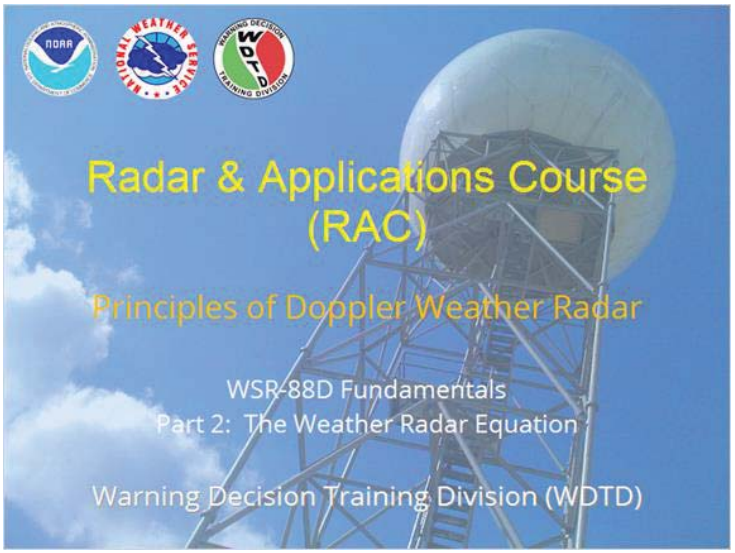
Notes:

Once the information is received at the radar, it has to be quantified and binned to be displayed. The more bins you have to fit the data into, the more precise the value displayed is to the real value. For example, if you have a bin that contains values from 5 to 10, then any real value from 5 to 10 will be displayed as 5. However, if you have a bin that contains only values from 5 to 7, then only values from 5 to 7 will be displayed as 5. In other words, the second example has higher precision. The reason for binning the data is to conserve bandwidth. For the WSR-88D, there are three basic binning levels… these are 3-bit, 4-bit, and 8-bit. 3-bit has 8 data levels, 4-bit has 16 data levels, and 8-bit has 256 data levels. So, 8-bit data is more precise than the 3-bit or 4-bit counterparts, but takes up more bandwidth. More of this will be discussed in the base and derived products section of RAC, but for now, just know the three different levels of binning.

WSR-88D Fundamentals Part 2: Weather Radar Equation

1. Intro to Radar Beam Characteristics

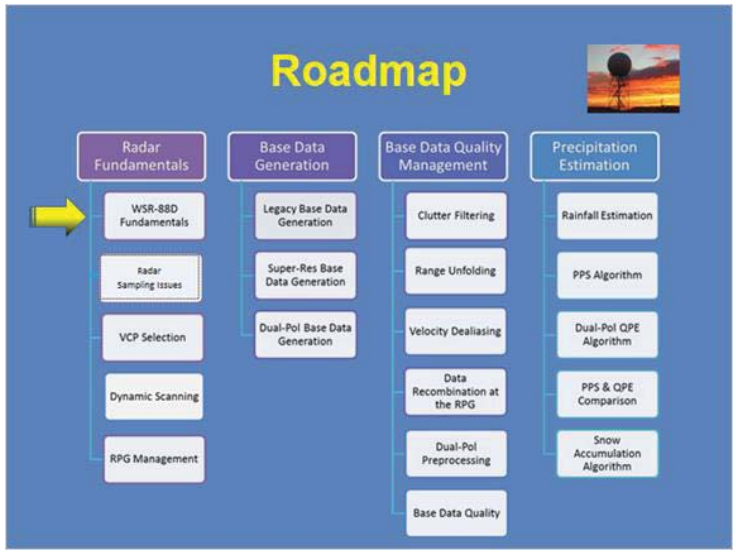
1.1 Welcome



Notes:

Welcome to the Radar & Applications Course (RAC) Principles of Doppler Weather Radar. This lesson is on WSR-88D Fundamentals Part 2: The Weather Radar Equation. It is presented by the Warning Decision Training Division (WDTD). Let's get started!

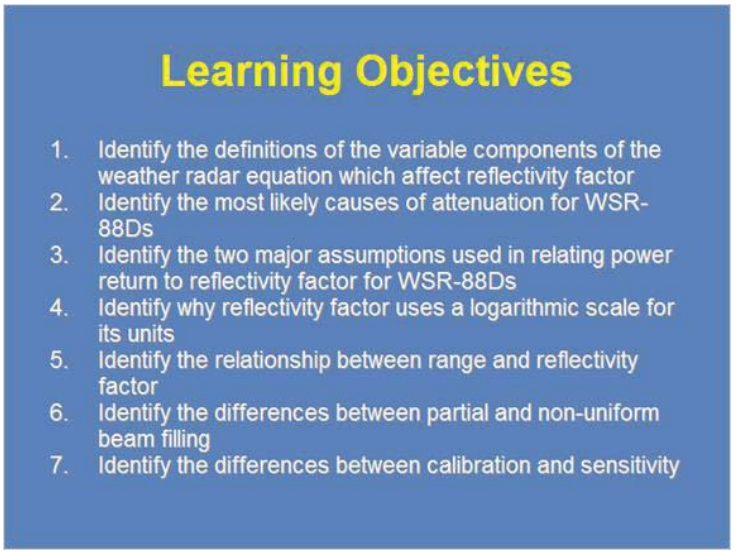
1.2 Roadmap



Notes:

Here is the complete roadmap for the entire "Principles of Doppler Weather Radar" section of RAC. You are currently in the WSR-88D Fundamentals portion of this section, and this portion consists of 5 lessons. Let's keep going!

1.3 Learning Objectives

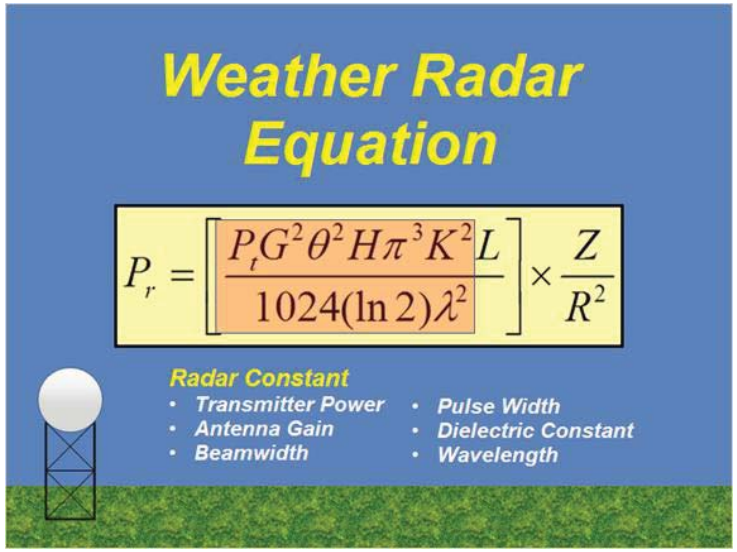


Notes:

Here are the learning objectives for this lesson. Please take a moment to review these objectives, as the quiz at the end of this lesson is based on these objectives.

2. Main Radar Equation

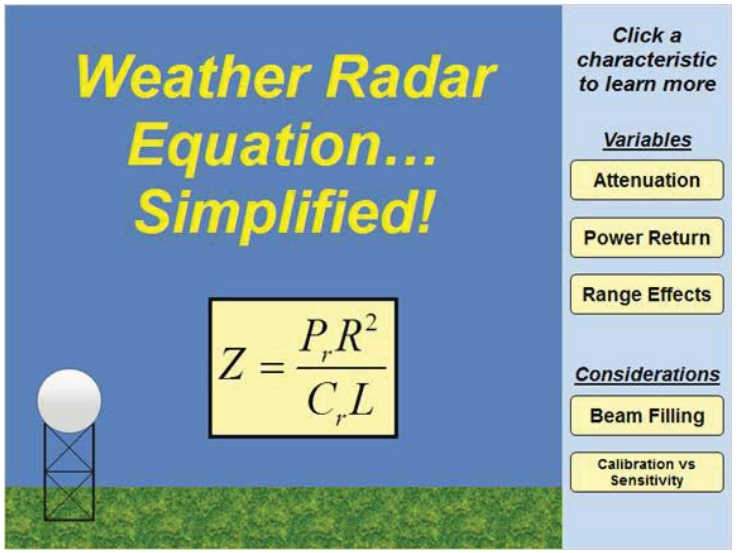
2.1 Weather Radar Equation



Notes:

Here is the full weather radar equation which is full of constants, variables, and equations…OH MY! But don't worry, we can drastically simplify this equation by combining many of these terms into what is called the "radar constant" which is just various aspects of the radar system that usually remain constant or are assumed to be constant with the WSR-88D. These constants are… 1) transmitter power, antenna gain, beamwidth, pulse width, dielectric constant, and wavelength. The entire goal of the weather radar equation is to take the power returned from weather objects, and convert that value into something useful, and this something useful is called reflectivity factor, or Z. So, let's keep moving…

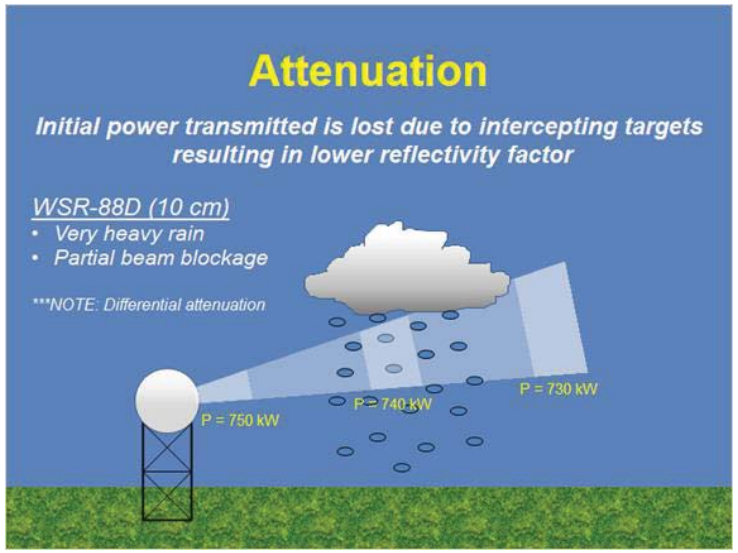
2.2 Weather Radar Equation HOME



Notes:

By combining all those terms into the radar constant, and solving for Z, we get this simplified radar equation (I promise…this is my last attempt at making this sound like a math class). But now that we have this simplified radar equation solved for reflectivity factor, let's examine each component (or variable) of the equation and see how reflectivity factor is affected. We'll also look at some considerations that are important to keep in mind when interpreting reflectivity factor. So, click on each of the variables and considerations listed on the left to learn more about the weather radar equation. Once you have viewed all items, a quiz button will appear and you will be ready to test your knowledge! One last note, because the WSR-88D is dual-polarized, this means it transmits both a horizontal and vertically polarized pulse. How we solve for both polarization is identical, so when referring to Z, just know it applies to both the horizontal and vertical channels.

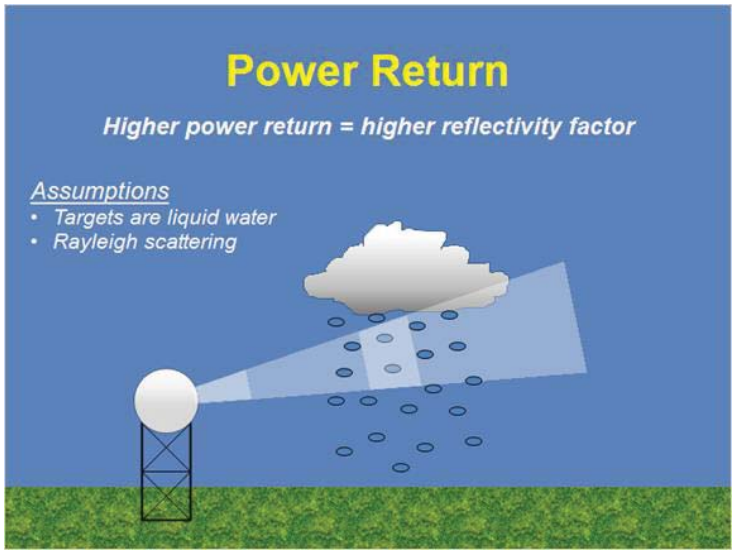
2.3 Attenuation



Notes:

The first term of the radar equation we'll look at is attenuation. It is defined as the loss in initial power transmitted due to intercepting targets. An increase in attenuation results in a lower reflectivity factor. For the WSR-88D, which is a 10-cm wavelength radar, attenuation is often negligible. A couple instances where attenuation may occur is in very heavy rainfall, or in areas of partial beam blockage. One other important consideration related to dual-polarization is differential attenuation. This is an instance where the horizontal channel attenuates slightly more than the vertical channel, or vice versa, and this can result in artifacts seen in the differential reflectivity product. More on this phenomenon will be covered later in this section.

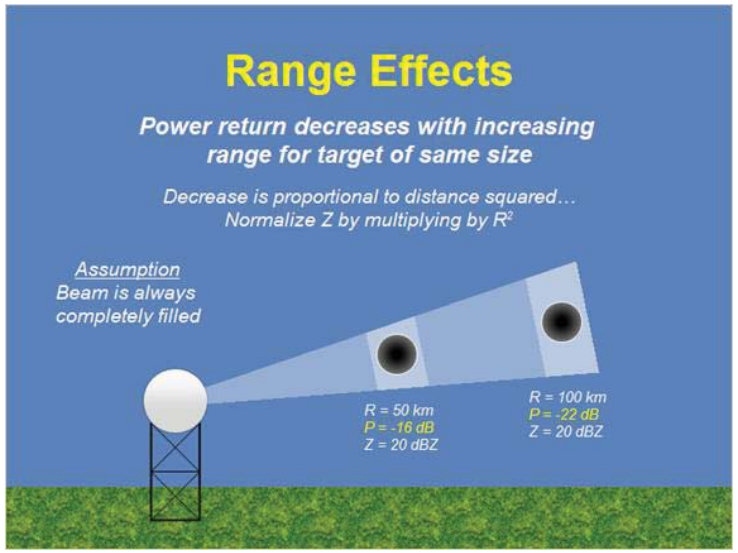
2.4 Power Return



Notes:

The next variable affecting reflectivity factor is the power return. This is pretty simple…higher power return equals higher reflectivity factor. However, our equation makes two very big assumptions when it comes to the power return. First, it assumes the power return is from liquid water targets (remember that K squared term that we rolled into the radar constant?). Ice particles have much lower dielectric constants (or K squared), so the radar equation is going to underestimate the reflectivity factor for ice. Therefore, ice will generally have lower reflectivity factor values than liquid water. The other assumption is that all scattering is Rayleigh in nature. We saw in the previous lesson that this is mostly true for WSR-88Ds, but not always. So, determining reflectivity factor from Mie scattering objects like hail using this equation is not entirely valid, but will get you close.

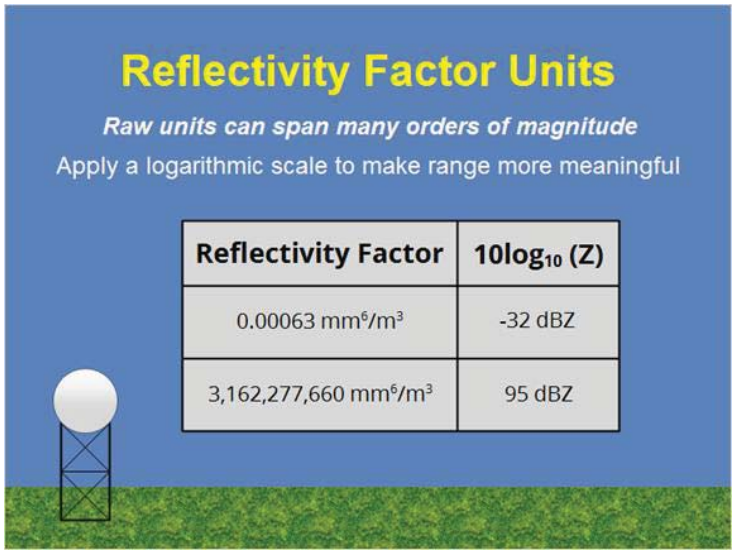
2.5 Range Effects



Notes:

The final variable affecting reflectivity factor is range. Recall that power return decreases with increase range for a target of the same size. This decrease in power is proportional to the distance squared. Therefore, the equation attempts to normalize the reflectivity factor by this distance squared. Look at the example… at 50 km, the target is producing a power return of -16 dB resulting in a Z = 20 dBZ. This same target at 100 km only produces a power return of -22 dB, but because of the range normalization, Z still equals 20 dBZ. One major assumption here is that the radar beam is always completely filled, but we'll find out soon this assumption is rarely valid.

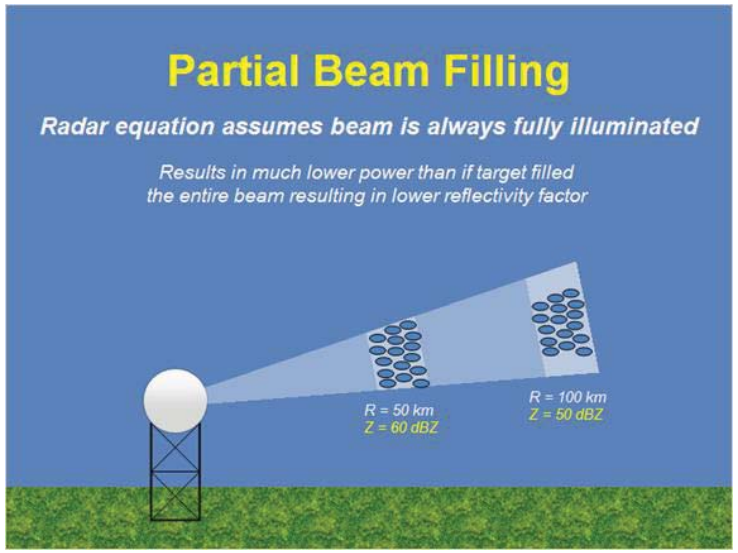
2.6 Reflectivity Factor Units



Notes:

Before we look at the considerations to the assumptions made in the radar equation, I want to quickly take a look at the units of reflectivity factor. The raw units are mm6/m3, and these values can span many orders of magnitude. So, a logarithmic scale is applied to the raw reflectivity factor to compress these values into a more meaningful range. For the WSR-88D, this range is from -32 dBZ to 95 dBZ, and you can see the raw reflectivity values to the left. The low end has a value of much less than 1, but the high end has raw values in the billions!

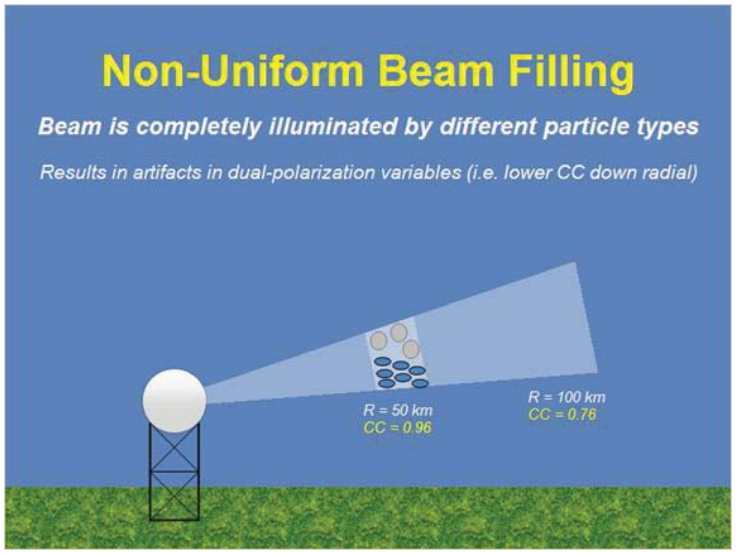
2.7 Partial Beam Filling



Notes:

We'll actually discuss two types of beam filling issues with the WSR-88D, but the first is partial beam filling. This results when the beam is not completely illuminated by weather targets. You might have rain in one portion of the beam and nothing in another portion. This partial beam filling will result in lower power return than if the target completely filled the beam which results in a lower reflectivity factor. In this example, a 60 dBZ echo results when these rain drops completely fill the beam at 50 km, but at 100 km, these same drops do not fill the beam completely, and the reflectivity factor is actually around 50 dBZ.

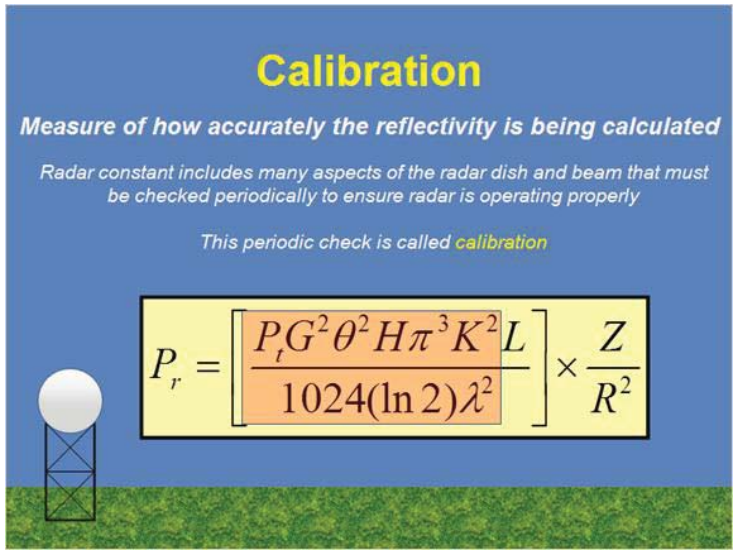
2.8 Non-Uniform Beam Filling



Notes:

Another type of beam filling issue is non-uniform beam filling, and we'll discuss it more later on in this section, but for now just know that this phenomenon occurs when weather targets of varying type (typcially hail and rain) exist in different portions of the beam. This causes varying propagation effects on the radar beam, and adversely affects the dual-polarization variables (especially CC). It will primarily show up as a valley of reduced CC along the affected radials.

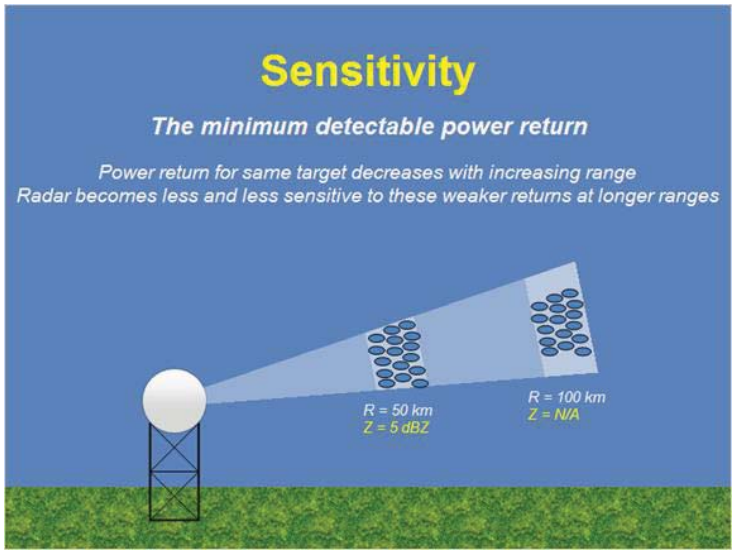
2.9 Calibration



Notes:

Oh no! It's that dreaded full equation again! Stop the insanity! Okay, don't freak out just yet… we need to see all this to better understand calibration of the radar. The radar constant consists of many physical aspects of the radar and radar beam. Because these can actually change with time, the radar operators will periodically check the radar system to see how these "constants" are behaving and make sure they are residing within acceptable limits. This periodic check is called calibration and it helps keep the measurement of reflectivity factor as accurate as possible.

2.10 Sensitivity



Notes:

One other consideration is to look at the sensitivity of the radar itself. This is basically defined as the minimum detectable power return. In other words, if the power return is too weak, the radar can't pick it up. For example, we have this group of weather targets at 50 km which produces a 5 dBZ reflectivity factor. However, at 100 km, this same group of targets will produce a weaker power return which is not detectable by the WSR-88D, and therefore will not assign a reflectivity factor. This ability to detect the power return is the sensitivity.

WSR-88D Fundamentals Part 3: Transmitting & Receiving Characteristics

1. Transmitter & Receiver Characteristics

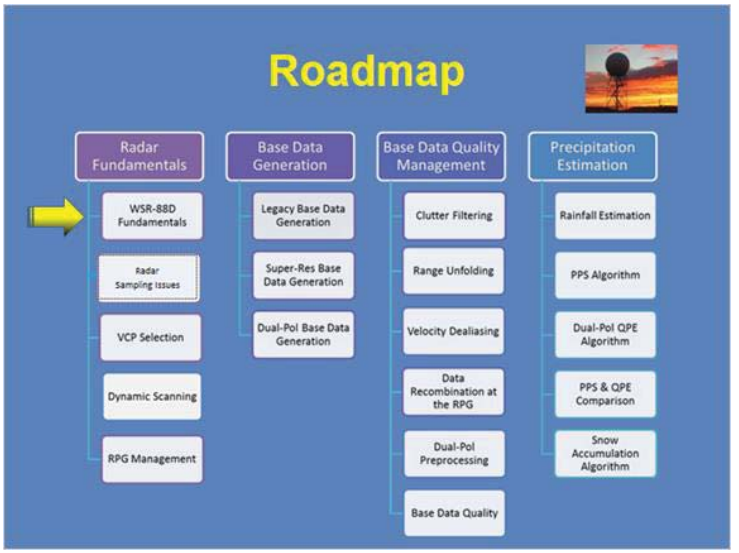
1.1 Welcome



Notes:

Welcome to the Radar & Applications Course (RAC) Principles of Doppler Weather Radar. This lesson is on WSR-88D Fundamentals Part 3: Transmitting and Receiving Characteristics. It is presented by the Warning Decision Training Division (WDTD). Let's get started!

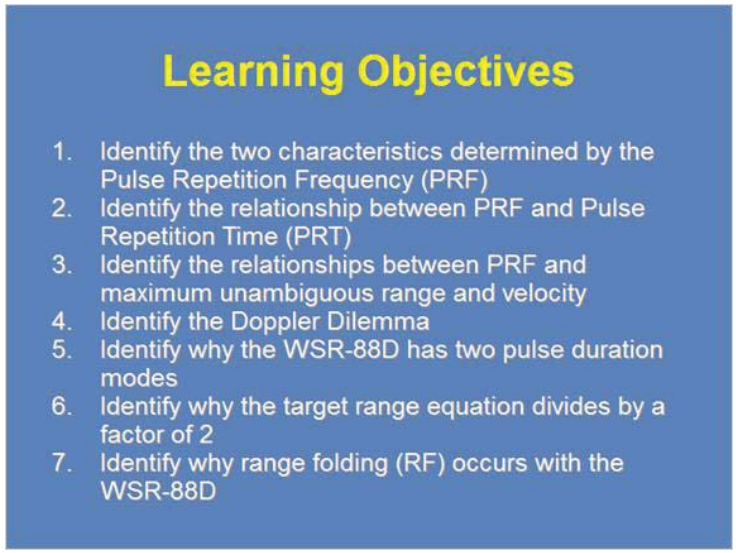
1.2 Roadmap



Notes:

Here is the complete roadmap for the entire "Principles of Doppler Weather Radar" section of RAC. You are currently in the WSR-88D Fundamentals portion of this section, and this portion consists of 5 lessons. Let's keep going!

1.3 Learning Objectives

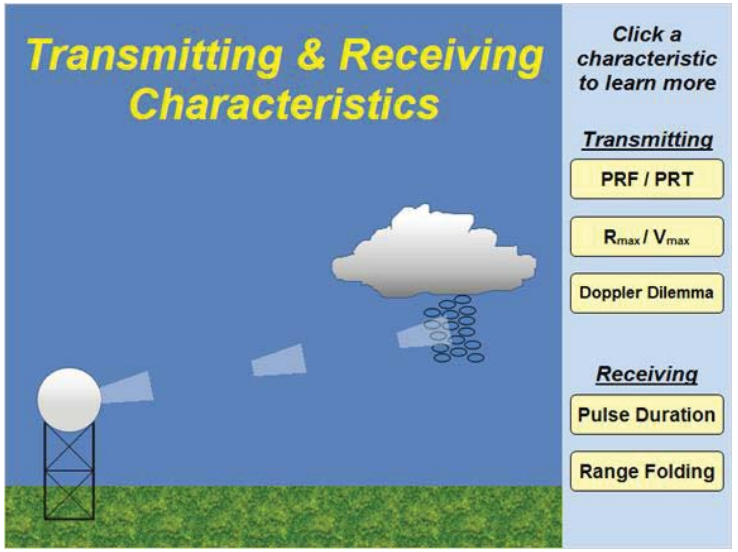


Notes:

Here are the learning objectives for this lesson. Please take a moment to review these objectives, as the quiz at the end of this lesson is based on these objectives.

2. Transmitting & Receiving Characteristics

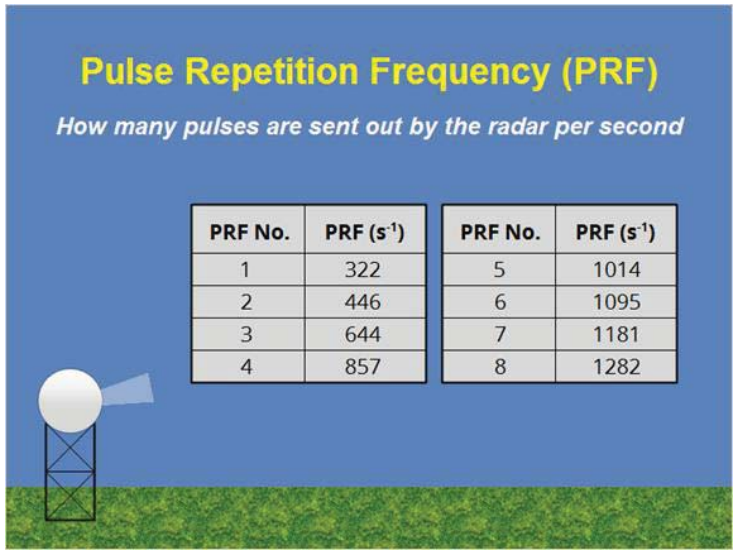
2.1 Transmitting & Receiving Characteristics HOME



Notes:

In the first lesson of this section, we learned that the WSR-88D emits pulsed radiation so that it can determine range to target. Well, there is a limit to the range it can accurately detect, and the WSR-88D also measures velocity which is dependent upon this pulsed transmission. In this lesson, we will take a look the transmitting and receiving characteristics that will affect the range and velocity values we can measure. Click on each of these characteristics on the right to learn more.

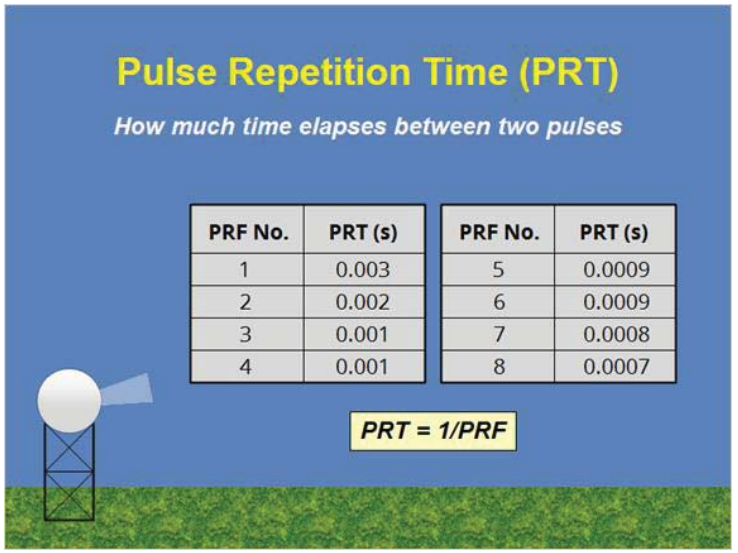
2.2 Pulse Repetition Frequency (PRF)



Notes:

The first term we'll introduce is pulse repetition frequency (PRF) which is nothing more than how many pulses are sent out by the radar every second. For the WSR-88D, there are 8 different PRF modes which range anywhere from 322 pulses per second up to 1,282 pulses per second. That's a lot of pulses!

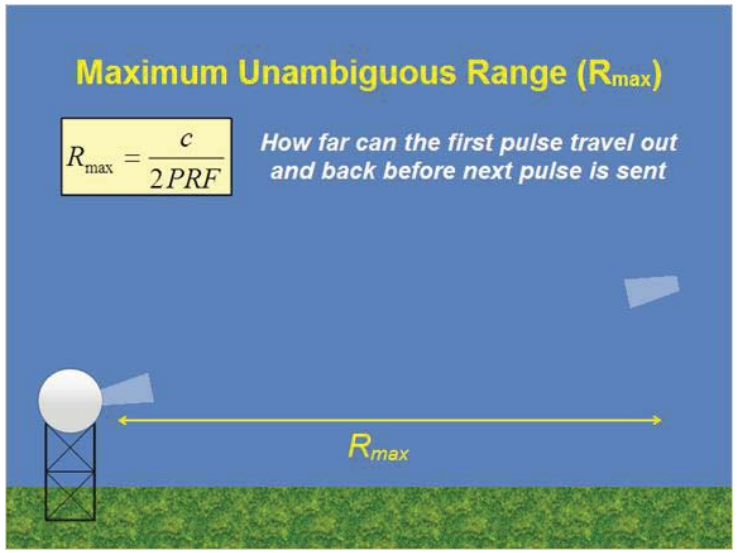
2.3 Pulse Repetition Time (PRT)



Notes:

And, if we just take the reciprocal of the PRF, we get the Pulse Repetition Time (PRT). This just tells you how much time elapses between two consecutive pulses. For the 8 WSR-88D PRF modes, the associated PRTs are listed here in this table. There isn't a lot of time in between pulses.

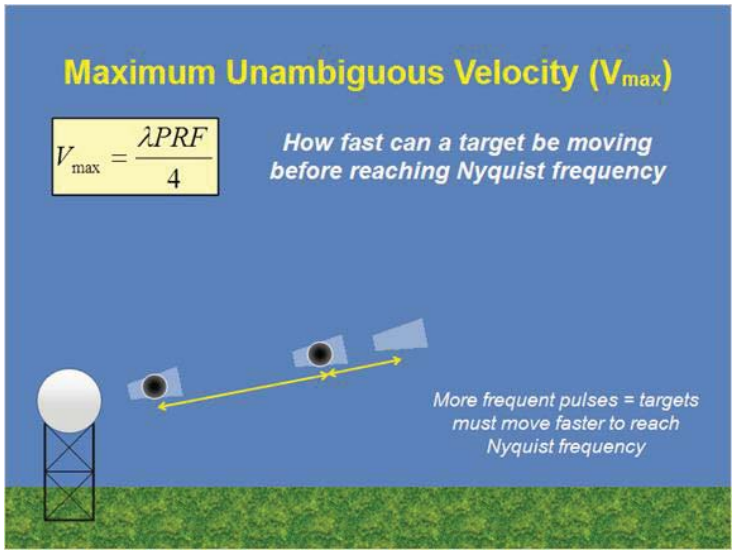
2.4 Maximum Unambiguous Range (Rmax)



Notes:

So we know that the WSR-88D emits pulsed radiation to determine range to targets. But, is there a limit to this range? Yes! This is called the maximum unambiguous range (Rmax) and it tells you how far the first pulse can travel out and back before the next pulse is transmitted. The equation for this is pretty simple… take the speed of light and divide by 2 times the PRF. What this tells us is the maximum unambiguous range increases with decreasing PRF. This makes sense because the fewer pulses we send out, the further it can travel out and back before the next pulse is transmitted.

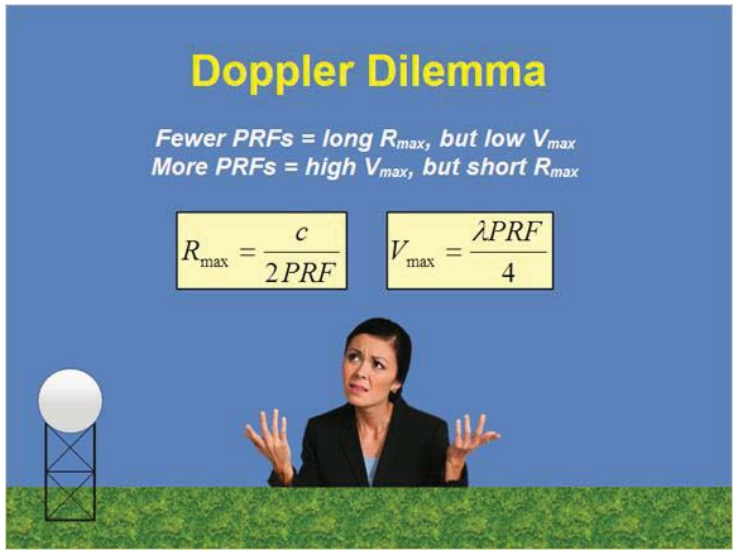
2.5 Maximum Unambiguous Velocity (Vmax)



Notes:

When it comes to measuring velocity, the radar does it by measuring the frequency shift between pulses. It's called the Doppler effect, and hence why we call the WSR-88D, Doppler radars. Well, if only a few pulses are sent out, the radar can only reconstruct a few pieces of the frequency waveform of the intended target. If the target is moving too fast, then the reconstructed waveform from the multiple pulses will have multiple solutions (or ambiguous velocities). The frequency at which we reach these ambiguous velocities is called the Nyquist frequency. Since faster particles have higher frequency shifts, we need to send out more pulses to better reconstruct the waveform. Therefore, higher PRFs lead to higher unambiguous velocities. To put it another way, the more frequently we send pulses, the faster the targets have to move in order to reach the Nyquist frequency.

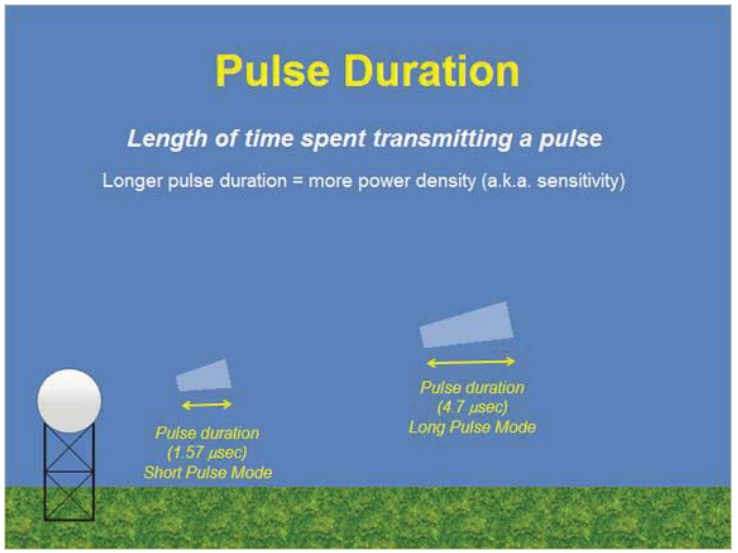
2.6 Doppler Dilemma



Notes:

So, you might be thinking there's a problem here… and you are absolutely right! If lower PRFs give me better range detection, but reduce the velocities that can be unambiguously measured, and vice versa, then we have a dilemma! This is called the Doppler Dilemma. By choosing one PRF, we have to sacrifice either Rmax or Vmax. But don't fret, the radar engineers at the Radar Operations Center have come up with scanning strategies where we can do multiple scanning strategies to maximize both values which we'll discuss later, but just know, Rmax and Vmax both depend on PRF, but in different ways and this can cause a problem.

2.7 Pulse Duration



Notes:

The length of time the transmitter remains on while transmitting a pulse is called the pulse duration. The longer the pulse duration, the more power density resides inside the pulse which can increase its sensitivity. For the WSR-88D, there are two pulse duration modes: short pulse mode and long pulse mode. The short pulse mode is 1.57 microseconds, and the long pulse mode is 4.7 microseconds. All of the scanning strategies employ the short pulse mode except one which uses the long pulse mode. We'll talk more about scanning strategies in a later lesson, but for now we'll look at pulse durations for another reason on the next slide…

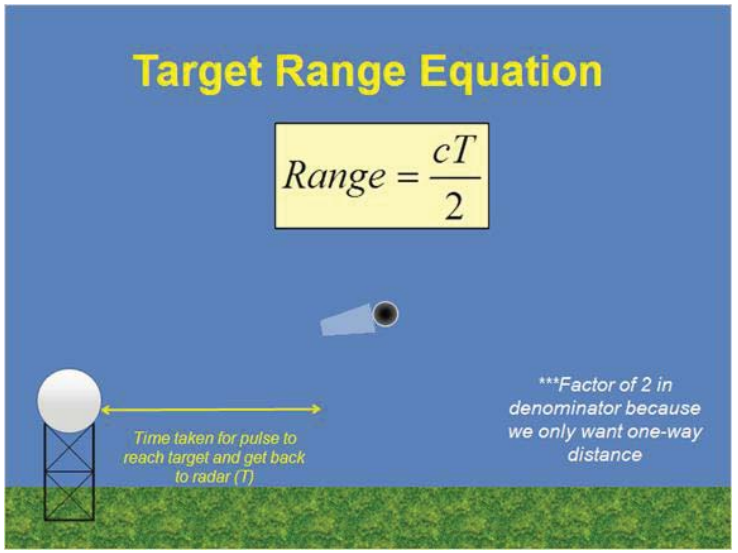
2.8 Listening Period



Notes:

When the radar is transmitting, it can't be listening for radar echo returns. Therefore, longer pulse durations means less time listening for radar echo returns. This listening period is defined as the time between the end of the first pulse and the beginning of the second pulse. For the WSR-88D, the radar is in listening mode while in short pulse mode for 99.8% of the time. In long pulse mode, it is in listening mode for approximately 99.5% of the time. So, you can see, the radar is primarily listening to what's going on in the atmosphere.

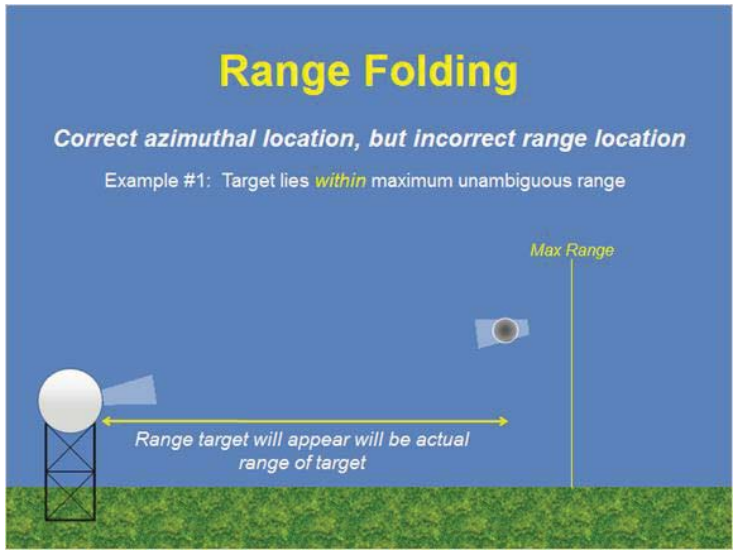
2.9 Target Range Equation



Notes:

When a radar echo return arrives at the radar, it is nice to know how far away the radar echo resides. Since we know how fast the pulse is traveling, and how long it has been since we transmitted the pulse, we just simply multiply the speed of light by the time it took the pulse to be transmitted and then received back at the radar. We have to then divide this value by 2 since the time between transmission and reception is a round-trip value, and we only care about the one-way distance.

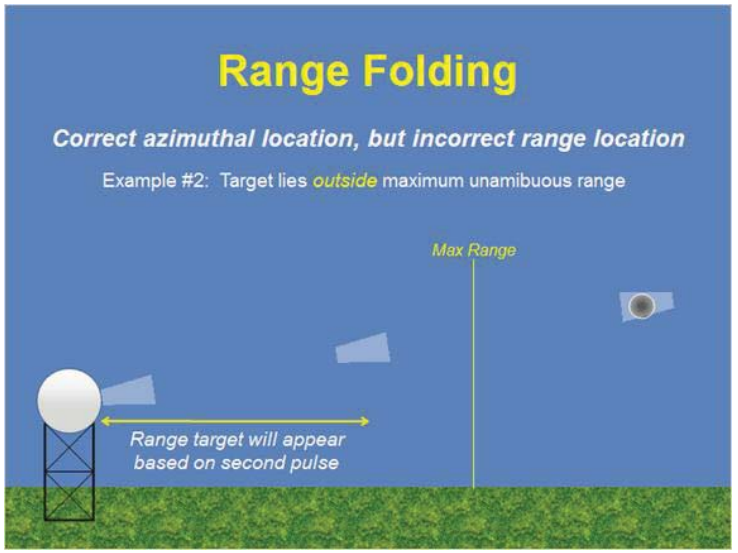
2.10 Range Folding



Notes:

The range equation leads us nicely into the next topic…range folding. As long as the first pulse makes it back to the radar before the second pulse is transmitted, the radar will correctly measure the range to the target. However, if the first pulse makes it back to the radar after the second pulse has been transmitted, then the time variable will be incorrect because it will be based on when the second pulse was transmitted, not the first pulse. This will incorrectly place the return echo at a closer range than where it actually occurred. This phenomenon is called range folding, and let's look at two examples. This first example shows the normal case where no range folding will occur. The pulse is sent out, it encounters a target within the maximum unambiguous range, and returns an echo to the radar. It arrives at the radar before the second pulse is transmitted and therefore gets the correct range applied to it.

2.11 Range Folding



Notes:

However, in this second example, the target lies outside the maximum unambiguous range. Thus, when the first pulse encounters the target, the target returns some of the power back towards the radar. However, before the first pulse return power makes it to the radar, a second pulse is transmitted. Then, when the first pulse finally does make it back to the radar, the radar thinks this power return is from the second pulse, not the first pulse, and therefore thinks the return power is from a target at a range closer to the radar than where it actually resides. It is range folded.

WSR-88D Fundamentals Part 4: Non-Standard Beam Consequences

1. Intro to Radar Beam Characteristics

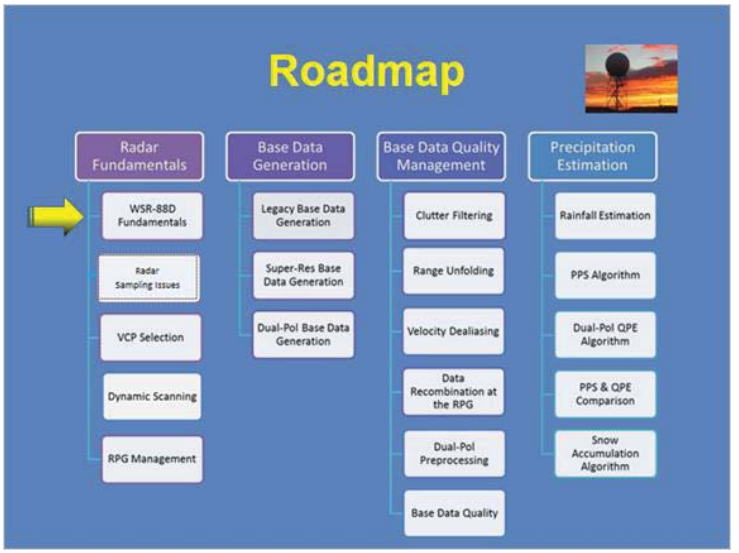
1.1 Welcome



Notes:

Welcome to the Radar & Applications Course (RAC) Principles of Doppler Weather Radar. This lesson is on WSR-88D Fundamentals Part 4: Non-Standard Beam Consequences. It is presented by the Warning Decision Training Division (WDTD). Let's get started!

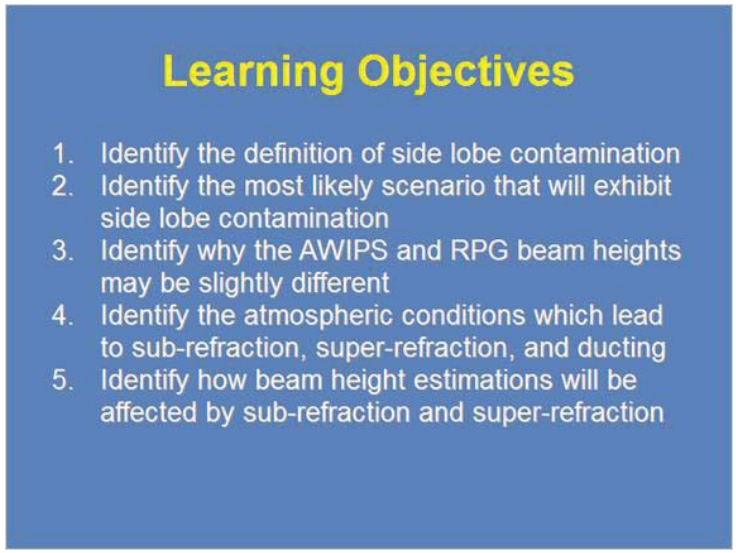
1.2 Roadmap



Notes:

Here is the complete roadmap for the entire "Principles of Doppler Weather Radar" section of RAC. You are currently in the WSR-88D Fundamentals portion of this section, and this portion consists of 5 lessons. Let's keep going!

1.3 Learning Objectives

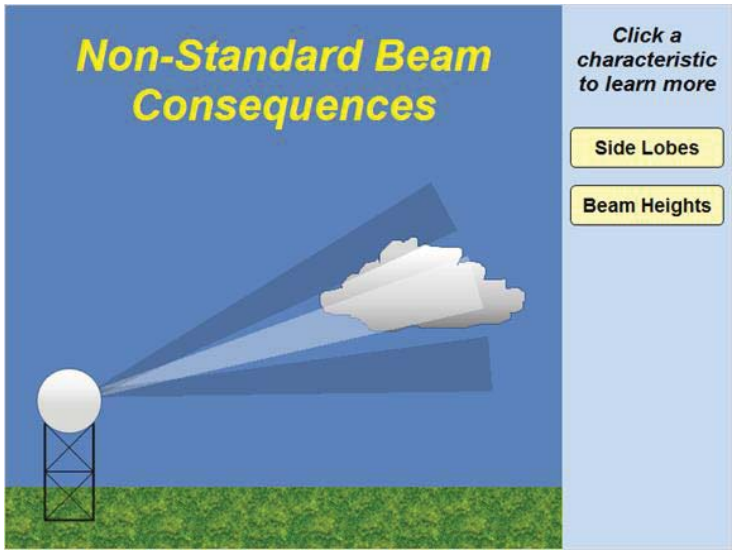


Notes:

Here are the learning objectives for this lesson. Please take a moment to review these objectives, as the quiz at the end of this lesson is based on these objectives.

2. Non-Standard Beam Consequences

2.1 Non-Standard Beam Consequences HOME



Notes:

Many of the general concepts of beam propagation are assumed to occur in a "standard" atmosphere, or we assume the beam is a rigid object. Well, the atmosphere is rarely, if ever, "standard" and the beam is not a rigid object. Click on the buttons to the right to learn more about the non-standard beam consequences known as side lobes and beam height estimations.

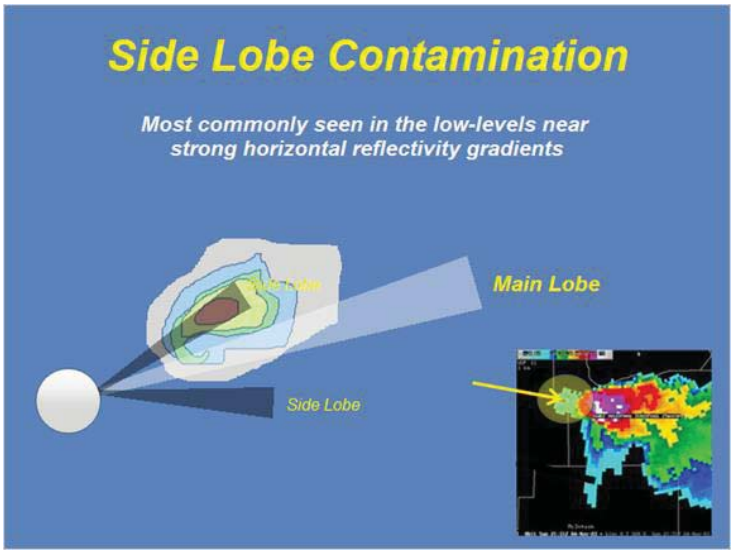
2.2 Side Lobe Definition



Notes:

Recall that the initial energy for the beam is generated by the transmitter and is isotropic in nature (radiates in all directions equally). Well, that's where the radar antenna comes into play. It focuses this energy into the 1 degree beam which is what is sent out into the atmosphere to detect the weather objects. However, the antenna doesn't focus all of the energy into this 1 degree beam. Some of it is focused into regions just outside the main lobe called side lobes. These side lobes contain a very small fraction of the total energy transmitted, but can intercept weather targets and produce returns at the radar which are strong enough to be seen on the radar display. These side lobes can occur either in the vertical or horizontal, and we'll discuss those next.

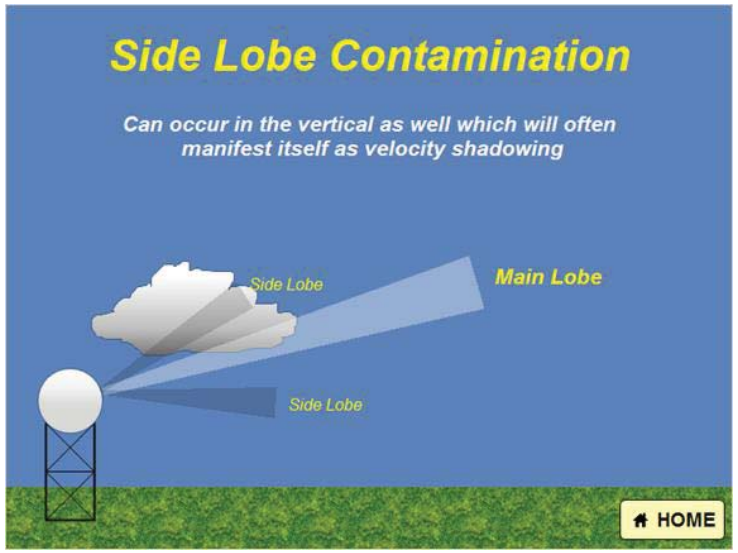
2.3 Side Lobe Horizontal



Notes:

Effects from side lobes are most commonly seen in the horizontal. For example, here is a typical conceptual model of a supercell thunderstorm. Once the main lobe passes by the core of the storm, the side lobe samples the core while the main lobe is sampling very low returns. Because the signal returned to the radar is dominated by the side lobe returns, that is the signal processed by the radar. However, because the radar thinks the return came from the main lobe, it places this return where the main lobe is sampling, which is just off to the side of the core in the clockwise direction. So, side lobe contamination will show up as weak reflectivity just to the side of a core. Here is an example of side lobe contamination.

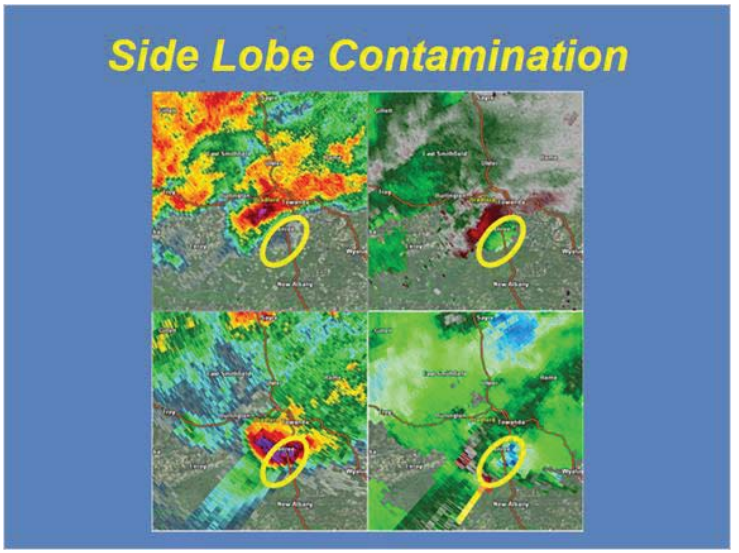
2.4 Side Lobe Vertical



Notes:

The last example of side lobe contamination was in the horizontal. However, side lobes exist in the vertical as well. The most common scenario where you will see the effects of side lobe contamination is when the main lobe is sampling the low levels, but the side lobe is sampling a fairly intense overhang. However, the reflectivity will not be the prominent feature, but rather the velocity signature. This is often called the velocity shadow. What happens is the velocity signature from the overhang is pretty much superimposed in the low levels. This can lead to spurious velocity couplets.

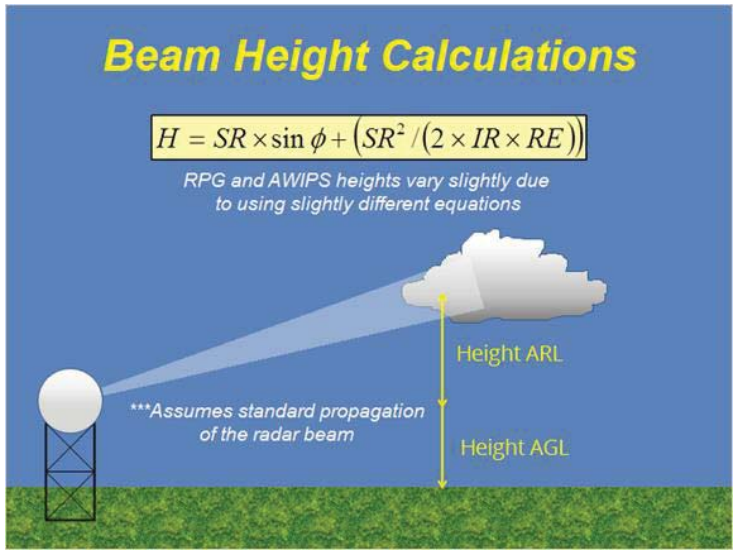
2.5 Side Lobe Vertical



Notes:

Here is an example. The top two images are the low level reflectivity (left) and velocity (right). Notices the intense inbounds well away from the core and near the core there are moderate outbounds which make it appear as though there is rotation. However, stepping up in elevation, there is as strong core directly above with very strong inbound velocity. These strong inbounds aloft are basically being superimposed below because of side lobe contamination. Therefore, be aware of this limitation when viewing velocity values in weak signal areas.

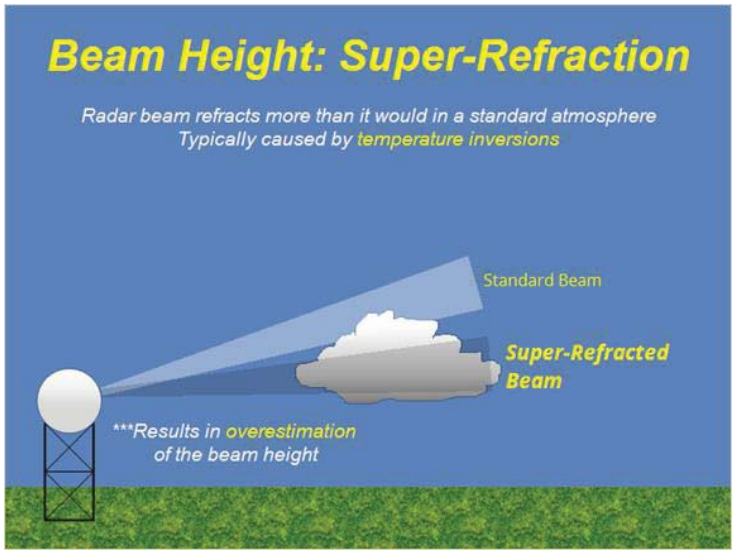
2.6 Beam Height Calcs



Notes:

So, here's another one of them funny equations, but don't worry, you won't be tested on it… okay, seriously, contain your enthusiasm :) The main thing to note here is this equation assumes a standard atmosphere for which the beam propagates through. As we all know, the atmosphere is rarely standard, so the beam will always be somewhere slightly different than the equation suggests. Also, this equation here is used by the RPG, but AWIPS uses a slightly different equation, so RPG heights and AWIPS heights may be off just slightly.

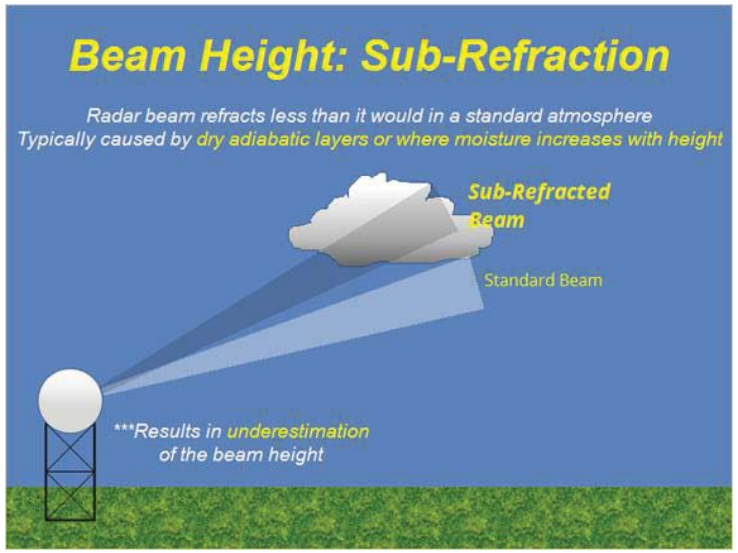
2.7 Super-Refraction



Notes:

Let's look at these different propagation anomalies. The first case we'll examine is super-refraction. Super-refraction occurs when the beam bends more than normal towards the ground. This phenomenon typically occurs when there is a temperature inversion near the ground. Because the height of the beam is lower than the equation suggests, then the reported beam height is overestimated. Let's look at the next case…sub-refraction.

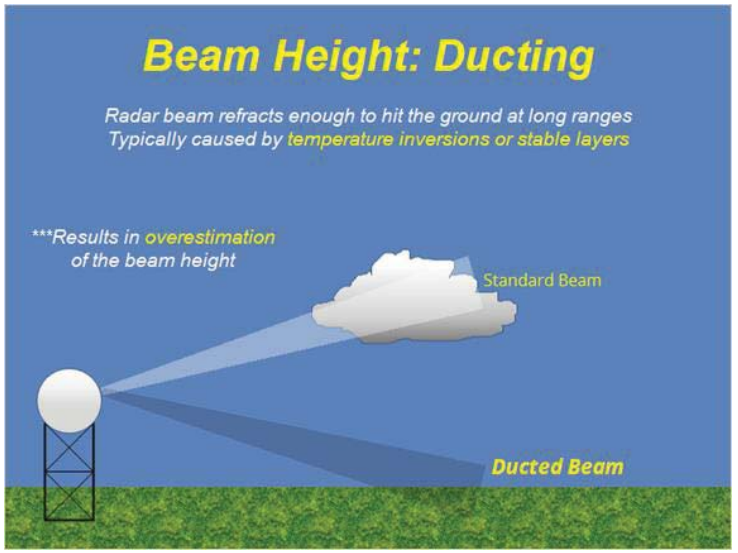
2.8 Sub-Refraction



Notes:

When the beam is sub-refracted, it bends upward a little more than normal, or is refracted a little less than normal. This type of atypical refraction occurs when there are dry adiabatic layers or areas where moisture increases with height. Because the beam is actually higher than the equation suggests, the reported height of the beam is an underestimation. The last case we'll look at is called ducting.

2.9 Ducting



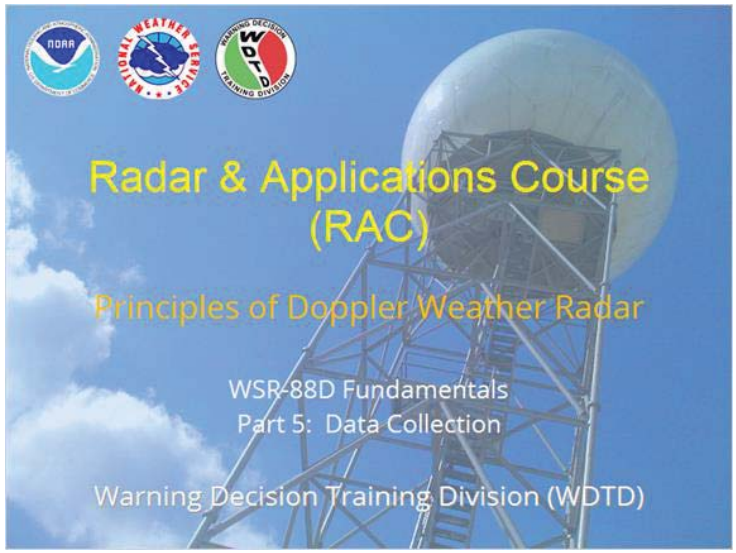
Notes:

A radar beam can sometimes get trapped in a layer and actually bent downward enough that it hits the ground at long ranges. This type of atypical propagation is called ducting. Ducting usually occurs when there are stable layers in the atmosphere or sharp temperature inversions. Because the actual beam height is lower than the equation suggests, the equation is actually overestimating the beam height. This concludes our section on beam height estimations and anomalies.

WSR-88D Fundamentals Part 5: Data Collection

1. Data Collection

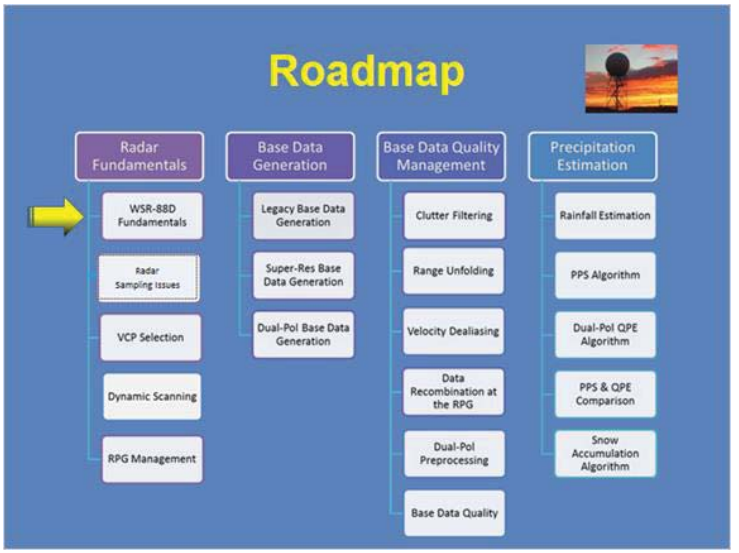
1.1 Welcome



Notes:

Welcome to the Radar & Applications Course (RAC) Principles of Doppler Weather Radar. This lesson is on WSR-88D Fundamentals Part 5: Data Collection. It is presented by the Warning Decision Training Division (WDTD). Let's get started!

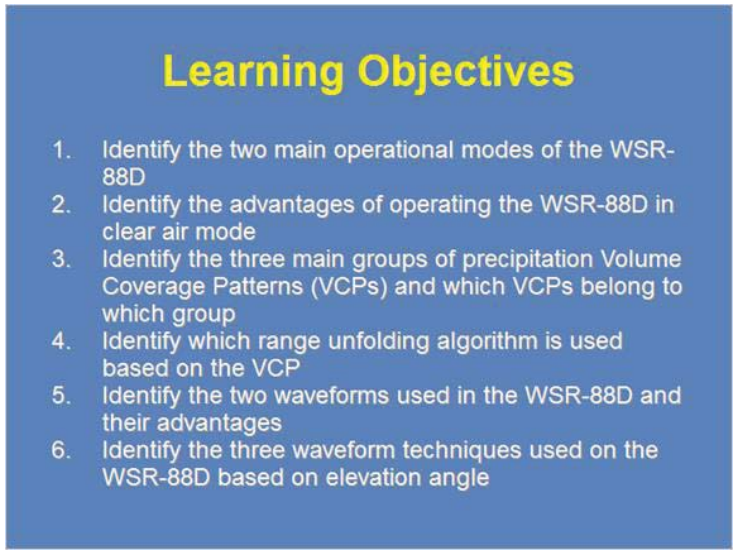
1.2 Roadmap



Notes:

Here is the complete roadmap for the entire "Principles of Doppler Weather Radar" section of RAC. You are currently in the WSR-88D Fundamentals portion of this section, and this portion consists of 5 lessons. Let's keep going!

1.3 Learning Objectives

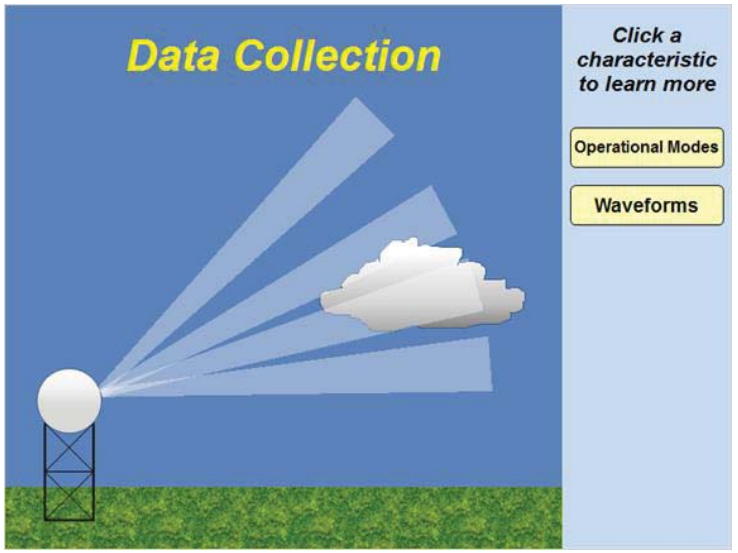


Notes:

Here are the learning objectives for this lesson. Please take a moment to review these objectives, as the quiz at the end of this lesson is based on these objectives.

2. Transmitting & Receiving Characteristics

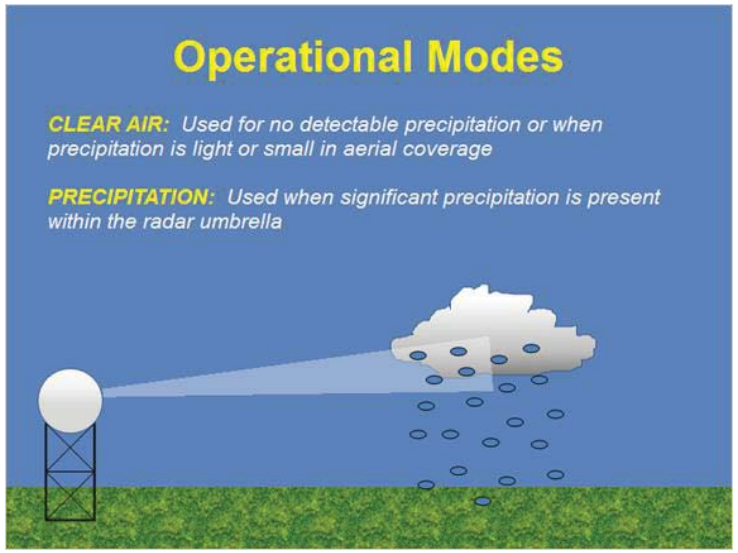
2.1 Data Collection HOME



Notes:

With a solid understanding of how the radar transmits and receives energy in order to detect meteorological targets, we can now take a look at how the radar is operated in order to collect data in the most efficient way possible. Every radar has a number of preset scanning strategies called operational modes and Volume Coverage Patterns (VCPs), and each VCP has various waveform techniques it employs to collect the data in the most efficient manner. More details on VCPs and scanning strategies will be covered in a later lesson. For now, we'll just look at the basics.

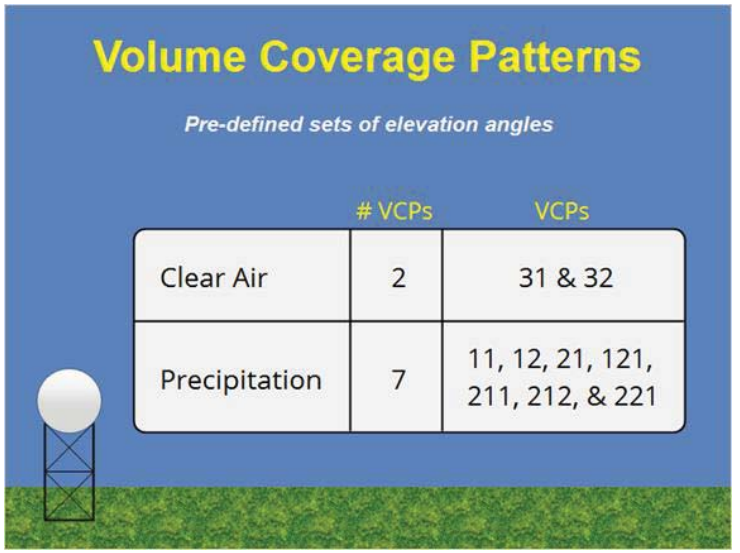
2.2 Operational Modes



Notes:

The whole idea of the WSR-88D is to detect precipitation echoes, but the radars run 24/7, so when weather is not present, the radar is running. So, it naturally follows there are just two operational modes which are Clear Air and Precipitation. Clear Air mode is primarily used when there are no detectable precipitation echoes within range, or when there is light precipitation or aerial coverage is small. When significant precipitation is present, Precipitation mode is enacted. Within each of these operational modes, there is a subset of different scanning strategies, and we'll look at those next…

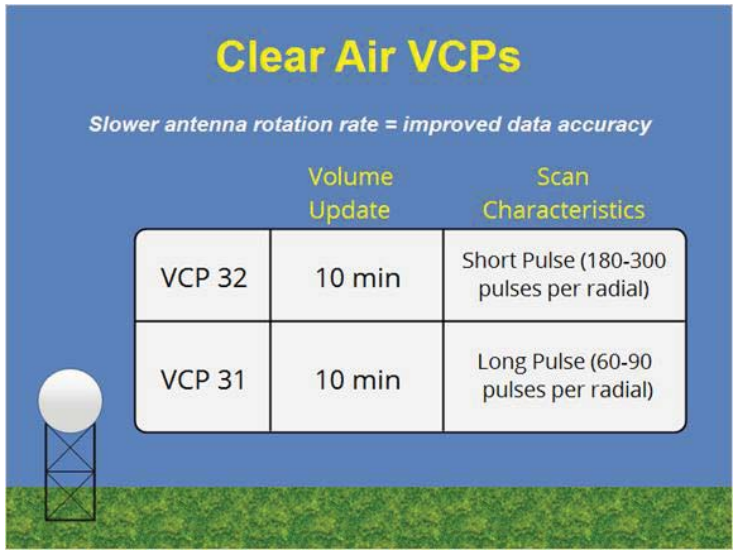
2.3 VCPs



Notes:

The pre-defined set of elevation angles run for each operational mode is called a Volume Coverage Pattern, or VCP. For Clear Air mode, there are two VCPs to choose from which are VCP 31 and 32. For the Precipitation mode, there are 7 different VCPs to choose from which are 11, 12, 21, 121, 211, 212, and 221. In the next couple slides, we'll take a little closer look at the general characteristics of these VCPs…

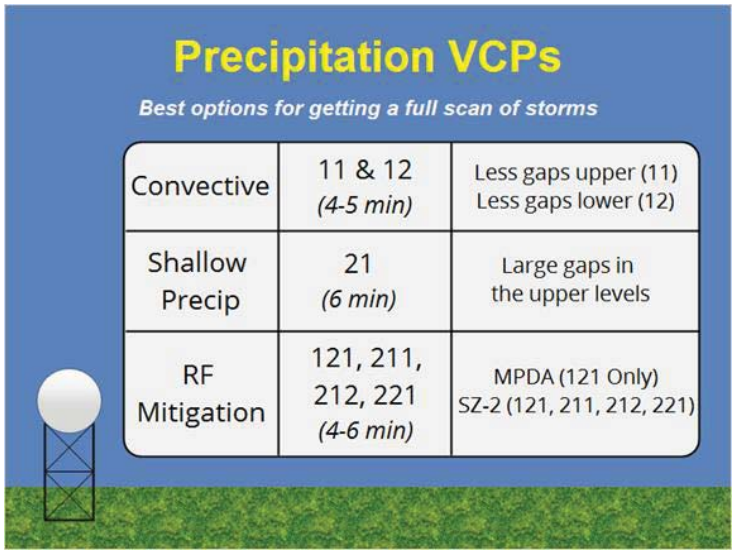
2.4 Clear Air VCPs



Notes:

The Clear Air mode VCPs operate using a much slower antenna rotation rate than any of the other VCPs. This is because the need for rapid updates is minimal. This slower antenna rotation, however, allows for more pulses per radial to be transmitted which allows for improved data accuracy. The total volume update time for the two VCPs is around 10 minutes. One last thing to note is VCP 31 is the only VCP to use long pulse mode for transmission. This long pulse mode allows for greater power density within the beam and therefore increases the sensitivity of the radar by roughly 3 dB. However, the VCP primarily uses low PRF pulses which leads to higher velocity dealiasing failures. Let's take a look at the Precipitation VCPs next…

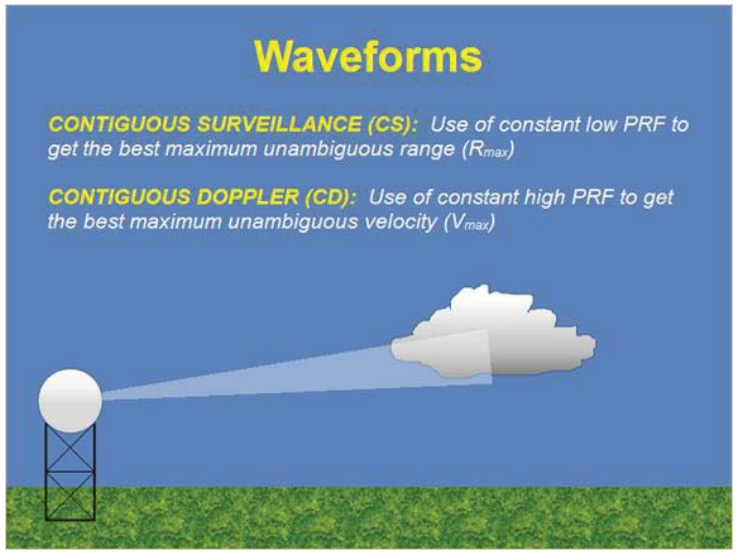
2.5 Precip VCPs



Notes:

The Precipitation VCPs are the best options to choose during significant precipitation events because they scan to higher elevations, therefore giving you a more complete picture of the storms. The seven VCPs can be divided up into 3 main groupings: Convective, Shallow Precipitation, and RF Mitigation. The Convective grouping consists of VCP 11 and 12. VCP 11 was the original convective VCP and it does adequate low level scanning with few gaps aloft. VCP 12 came about to fill in gaps in the lower level but reduced the number of elevations in the upper part of the scan. VCP 21 was developed for those very shallow and slow moving stratiform events, so there is decent low level sampling, but very few scans in the upper levels. Finally, recent upgrades to the WSR-88D have come with new techniques to improve range folding mitigation which are included with VCPs 121, 211, 212, and 221. VCP 121 uses an algorithm called the Multiple PRF Dealiasing Algorithm (MPDA) and the 200 series VCPs use the SZ-2 algorithm to unfold range ambiguities and recover velocity estimates in weaker signals. Alright, that's the high level overview of the VCPs themselves, let's go back home and take a look at how scanning strategies change from elevation to elevation.

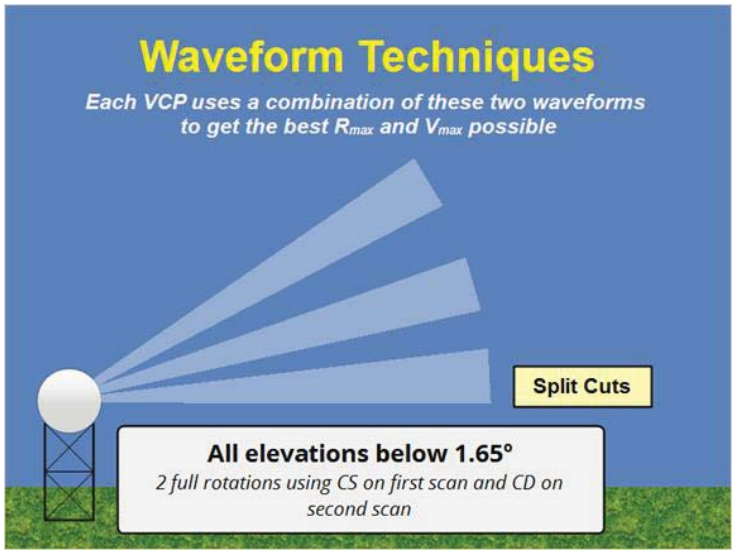
2.6 Waveforms



Notes:

Within each elevation angle, the WSR-88D employs two scanning waveforms. These are either contiguous surveillance (CS) or contiguous Doppler (CD). CS uses a constant low PRF to get the best maximum unambiguous range, while the CD uses a constant high PRF to get the best maximum unambiguous velocity. How these two waveforms are used depends on which elevation the radar is scanning. We'll take a look at this next…

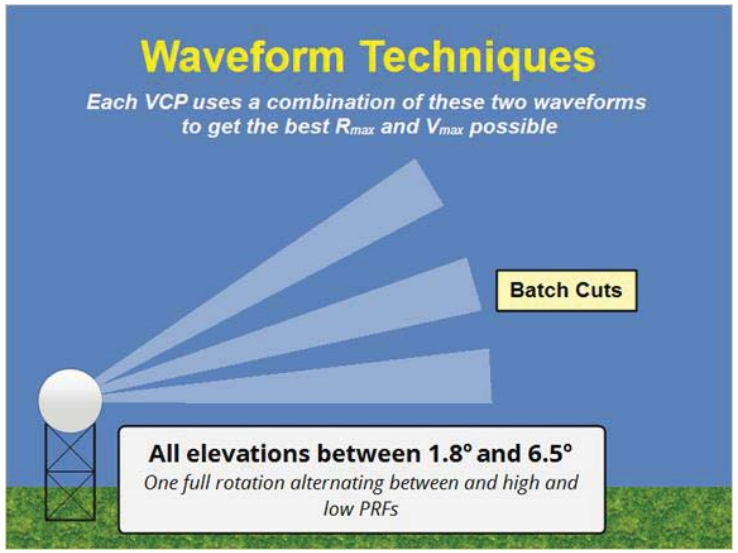
2.7 Waveform Techniques



Notes:

Split cut elevations include all elevations below 1.65 degrees. Each elevation consists of two full scans with the first one being in CS mode and the second scan being in CD mode.

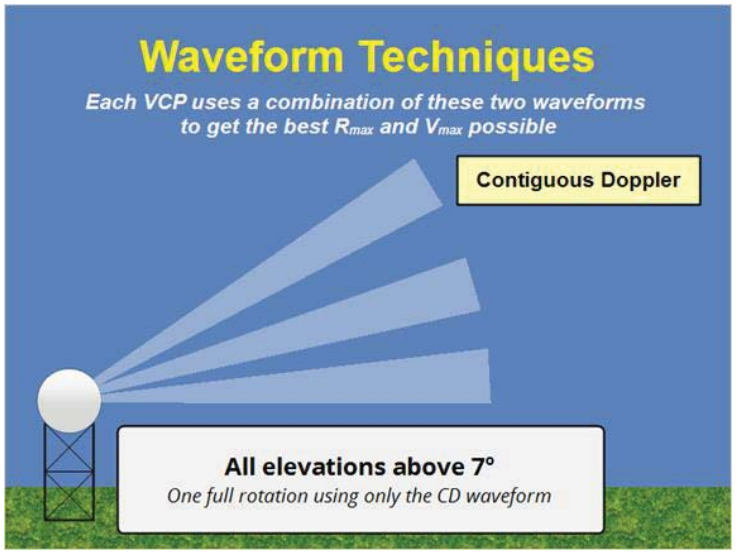
2.8 Waveform Techniques



Notes:

In the Batch elevations, which include all elevations between 1.8 degrees and 6.5 degrees, each elevation does one full scan where the pulses alternate between high and low PRF.

2.9 Waveform Techniques



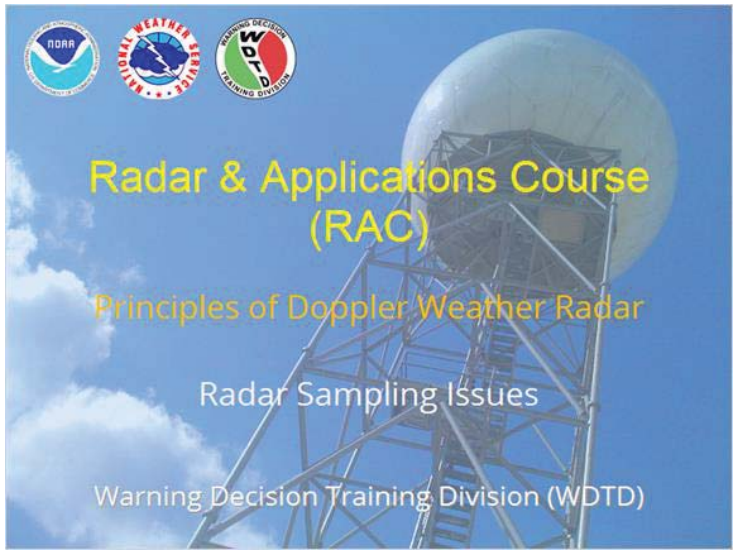
Notes:

All elevations above 7 degrees perform one scan using the CD mode only.

Radar Sampling Issues

1. Intro to Radar Beam Characteristics

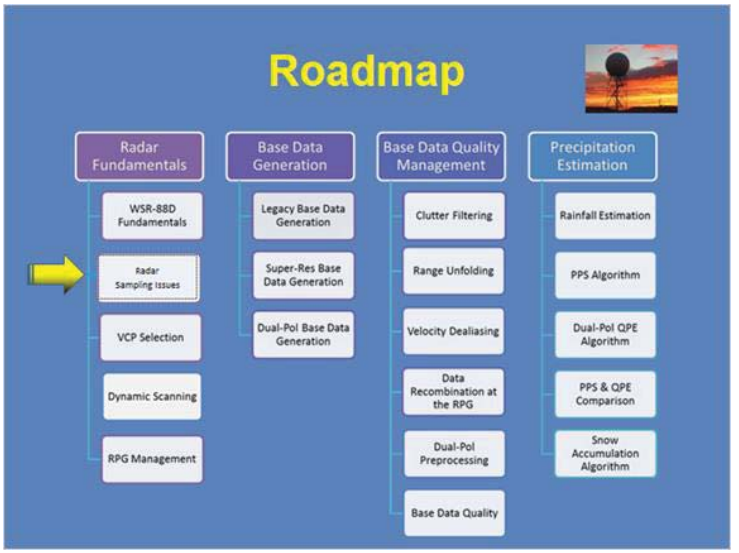
1.1 Welcome



Notes:

Welcome to the Radar & Applications Course (RAC) Principles of Doppler Weather Radar. This lesson is on Radar Sampling Issues. It is presented by the Warning Decision Training Division (WDTD). Let's get started!

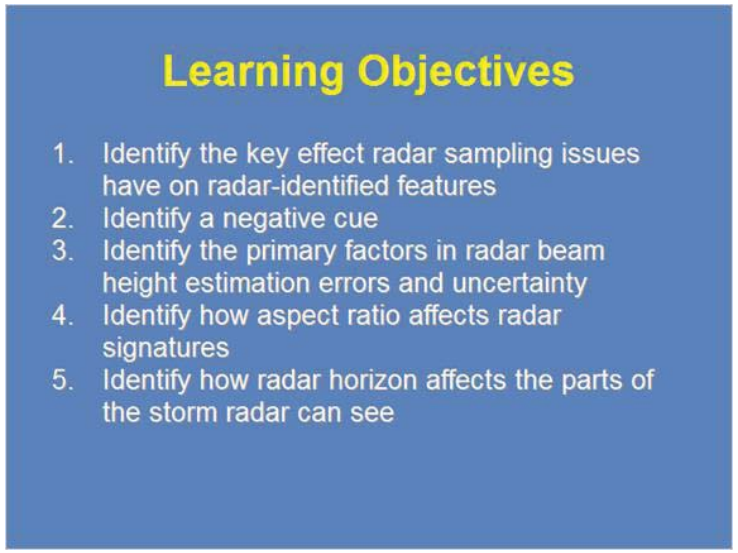
1.2 Roadmap



Notes:

Here is the complete roadmap for the entire "Principles of Doppler Weather Radar" section of RAC. You are currently in the Radar Sampling Issues portion of this section, and this portion consists of 1 lesson. Let's keep going!

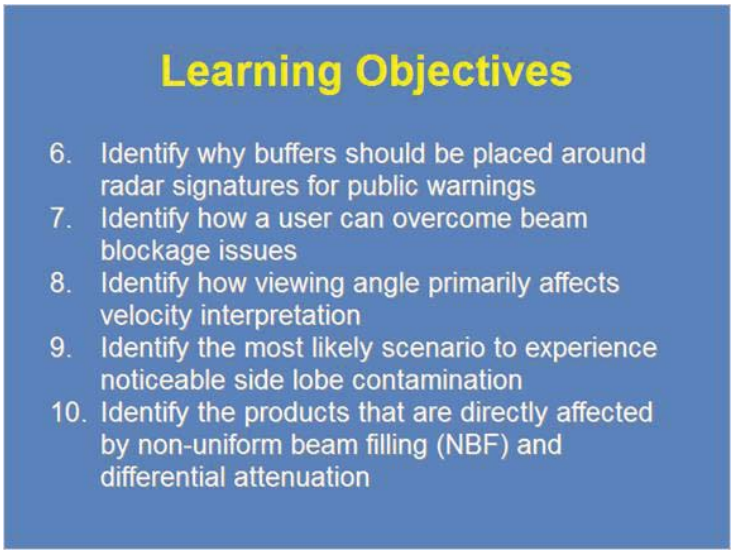
1.3 Learning Objectives



Notes:

There are ten learning objectives for this lesson. Here are the first five, and the last five will be on the next slide. Please take a moment to review these objectives, as the quiz at the end of this lesson is based on these objectives.

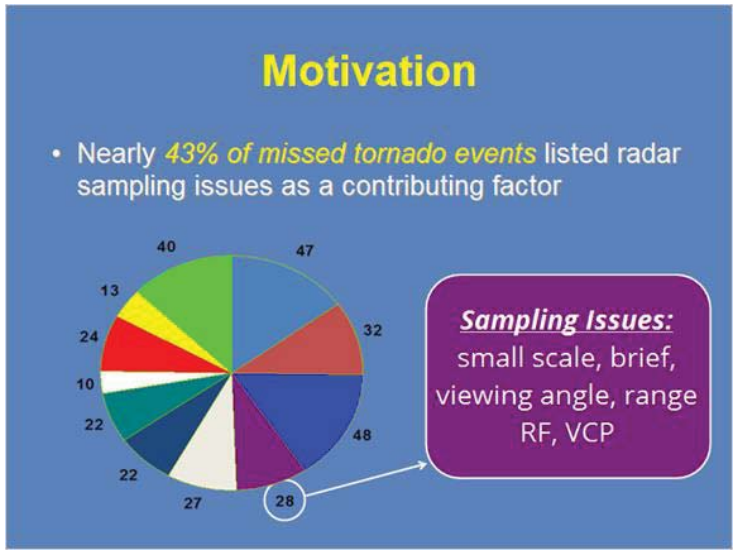
1.4 Learning Objectives



Notes:

Here are the last five learning objectives. Feel free to go back and forth between this slide and the previous to jot down these objectives.

1.5 Motivation

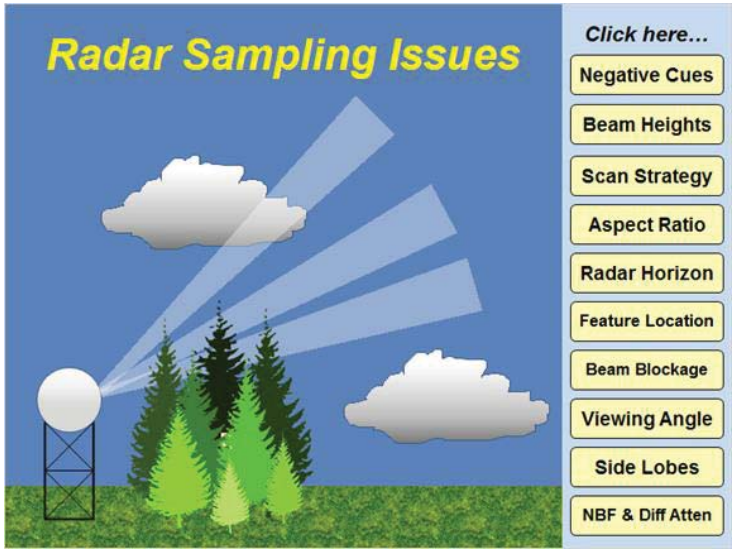


Notes:

A root cause analysis study revealed that in 65 missed tornado events, 28 of those missed events listed radar sampling issues as a contributing factor. Therefore, it is imperative we examine radar sampling issues and try to better understand them so as to avoid future missed events. Let's get started!

2. Transmitting & Receiving Characteristics

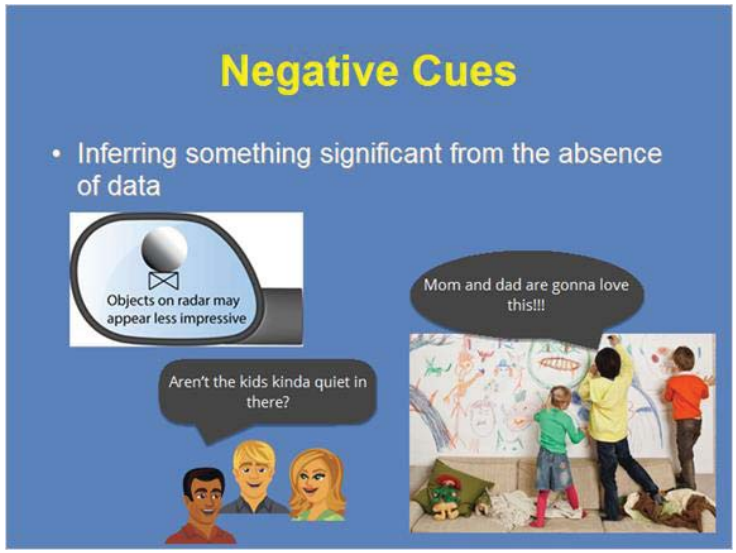
2.1 Sampling Issues HOME



Notes:

Radar sampling issues come in all shapes and sizes. Some can make interpretation almost impossible, and other are just annoyances. However, if you understand the situations in which these sampling issues can occur, there are ways to mitigate them. This lesson will show you common examples where these various sampling issues occur, and discuss how to mitigate them where possible. Let's get started!

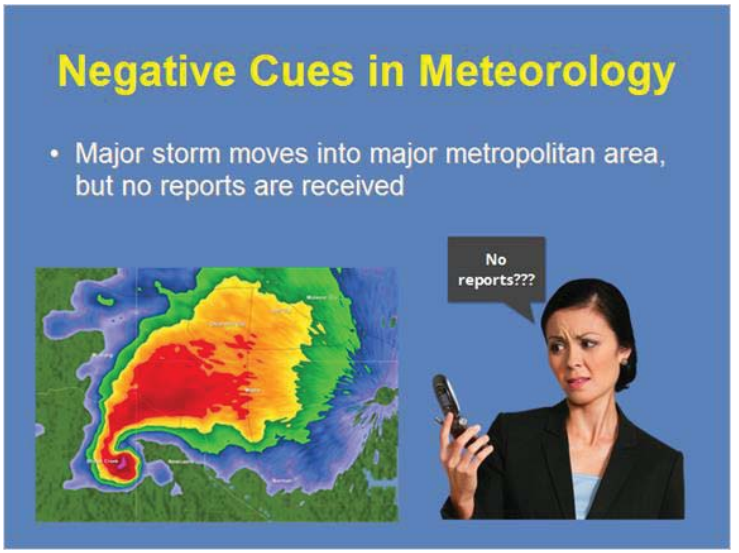
2.2 Negative Cues



Notes:

Before we actually jump into the actual sampling limitations, we need to discuss the concept of negative cues. Many of the sampling limitations we'll discuss impacts radar signatures usually by making the features appear LESS IMPRESSIVE than the features appear in actuality. And sometimes, these limitations mask the feature altogether, and we have to infer something significant by the absence of data. This concept is called negative cues. In every day life, let's say a group of parents are at a house talking in one room, and the kids are playing in another room. All of sudden, the kids' room gets quite. The parents notice there is NO noise coming from the kids' room and therefore gets their attention and they go check on the kids. The absence of noise from the kids' room was a negative cue to the parents that something significant was occurring in the kids' room.

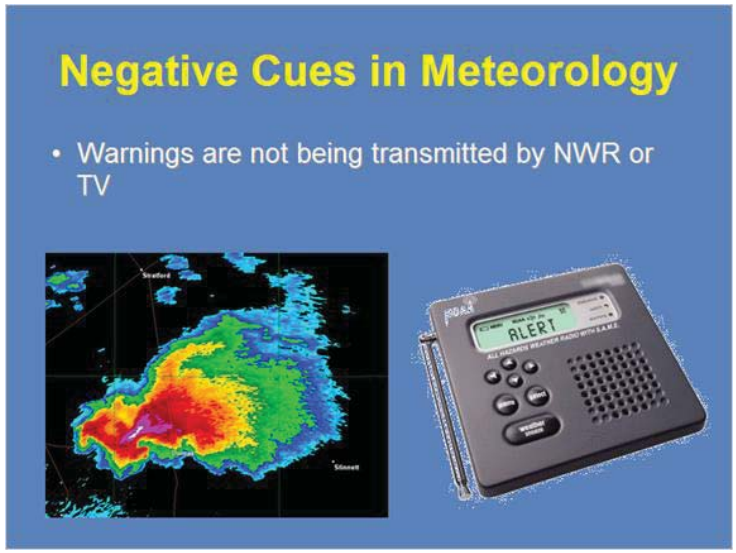
2.3 Negative Cues in Meteorology



Notes:

Let's take a look at some examples of negative cues in the warning world. First example here shows a very classic supercell with a hook echo moving into a major metropolitan area. With such a large population being affected by this storm, a forecaster would certainly expect reports coming into the office, right? What if there were no reports coming in? Wouldn't that seem suspicious? The lack of reports would signal to you, the forecaster, something isn't right. Maybe the storm is causing massive destruction and no one is thinking to call into the office. Maybe your phone lines are down and you can't receive incoming calls. Whatever it is, the lack of reports coming into the office is a negative cue that something significant is happening.

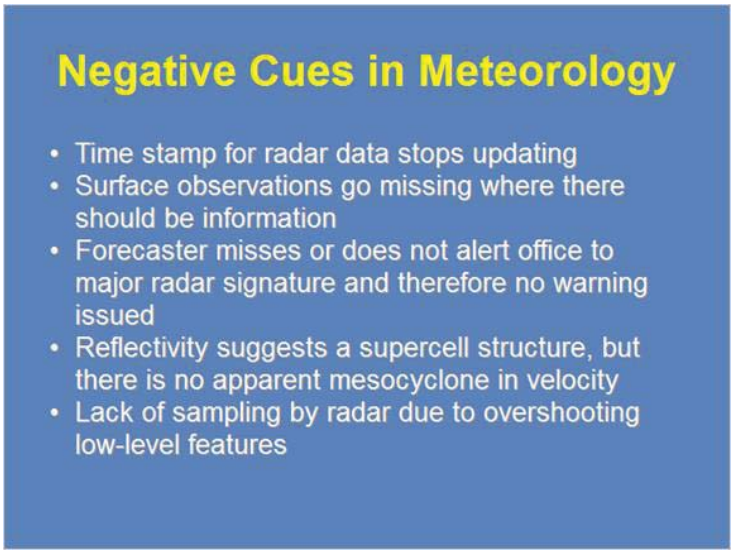
2.4 Negative Cues in Meteorology



Notes:

Another example might be you see this storm and decide to issue a tornado warning (which is a good decision). However, you keep getting calls and tweets from the public and media partners saying, "Where is that tornado warning?" They proceed to tell you that their weather radios aren't sounding for this storm even though family and friends are saying there's a tornado. What does this tell you? Did you issue for the right storm? Did you not click submit on the WarnGen interface? Is the NOAA Weather Radio system down? This lack of information to the public is the negative cue to you that something isn't right.

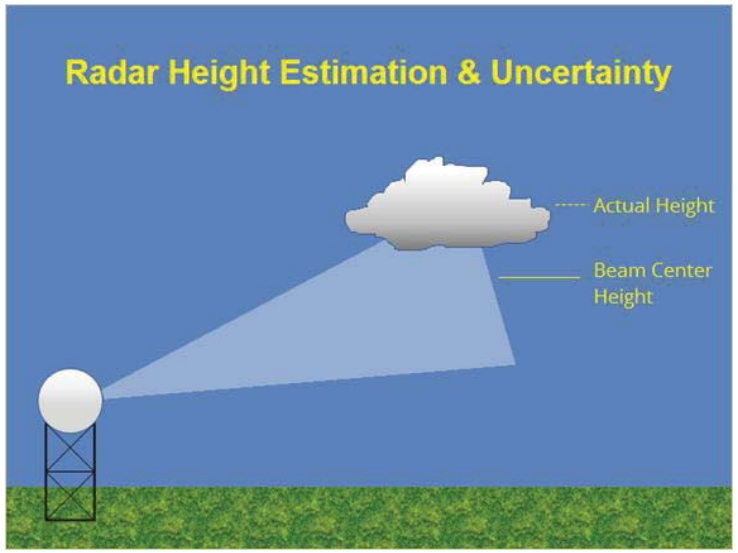
2.5 Negative Cues in Meteorology



Notes:

We could go on and on with great examples of negative cues in warning operations, but we want to learn about radar sampling issues, right? So here are some other examples of how negative cues occur in warning operations. Notice many of them deal with radar interpretation. That's the point of this lesson… radar sampling issues often lead to negative cues in the warning environment. Therefore, having a good grasp on radar sampling issues will help you better spot these negative cues in warning operations and make you a better forecaster. Let's go look at these sampling issues now.

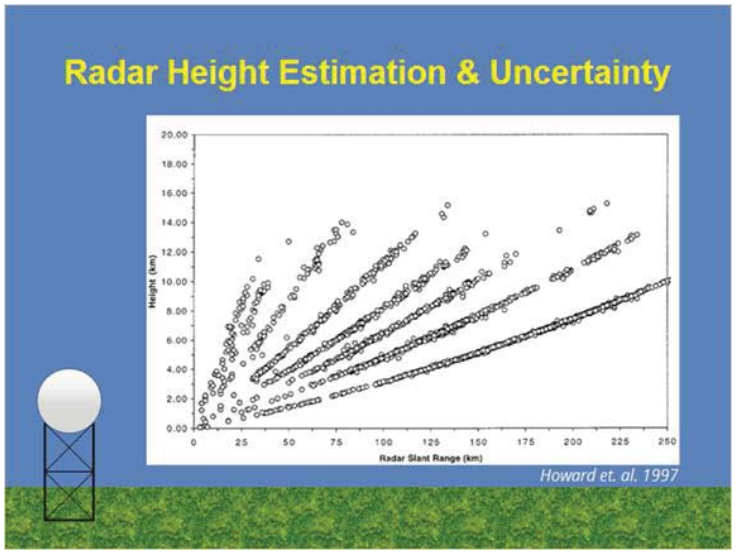
2.6 Beam Heights & Uncertainty



Notes:

Recall the radar beam spreads out as it goes down range making the beam width wider and wider. At distances as close as 60-70 miles, the beam width is over 1000s of feet. Therefore, an object detected at the edge of the beam, like this cloud here, might have an actual height of 22,000 feet, but the beam center is located at 18,000 feet. So, the radar will think the cloud echo is at 18,000 feet, whereas it is really at 22,000 feet. Just keep this fact in mind when interpreting radar echo heights.

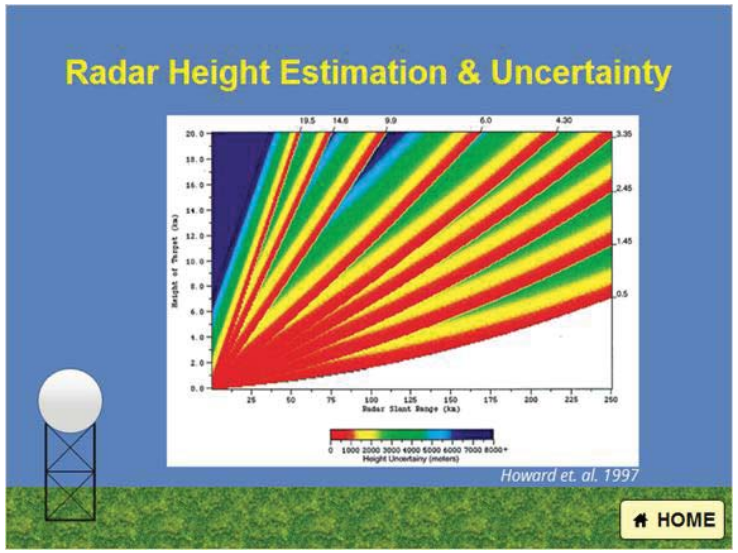
2.7 Radar Height Estimates



Notes:

A study done by Howard et. al. (1997) actually plotted the echo top heights as a function of range from radar. Look at how all those radar echo top heights fall nicely along straight lines. Is this the reality of storms? Absolutely not! This graphic illustrates how assigning all echoes to beam center results in echo heights being constrained to the heights based on the elevation angle scanned. So, storms exist in a continuous space, but the radar can't determine that. So, what is our uncertainty in these heights? We'll look at that next.

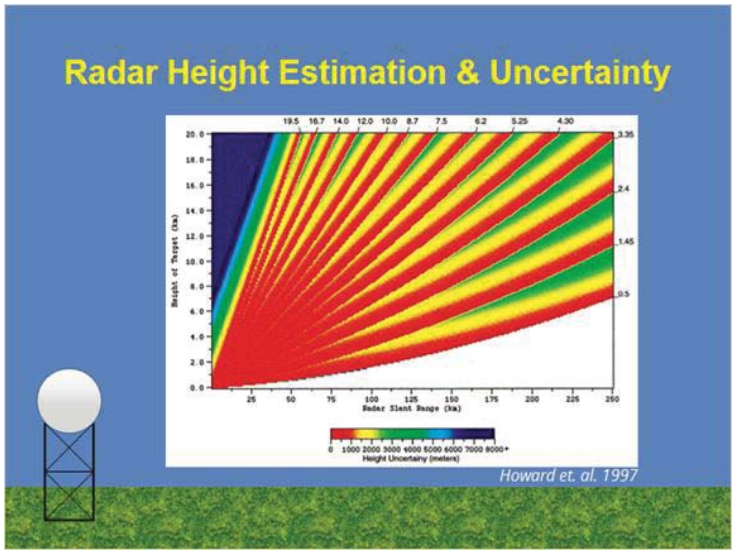
2.8 Radar Height Uncertainties



Notes:

This first graphic here shows the uncertainty (in meters) of the height estimations for each elevation angle for VCP 21. Red is good, blue is bad! Notice how within 50 km range and below 4 km height, the estimates are pretty accurate (within 1000 m). However, as you get higher is height, or farther in range, the uncertainty can be as much as 3 to 7 km because there are gaps in the scanning strategy.

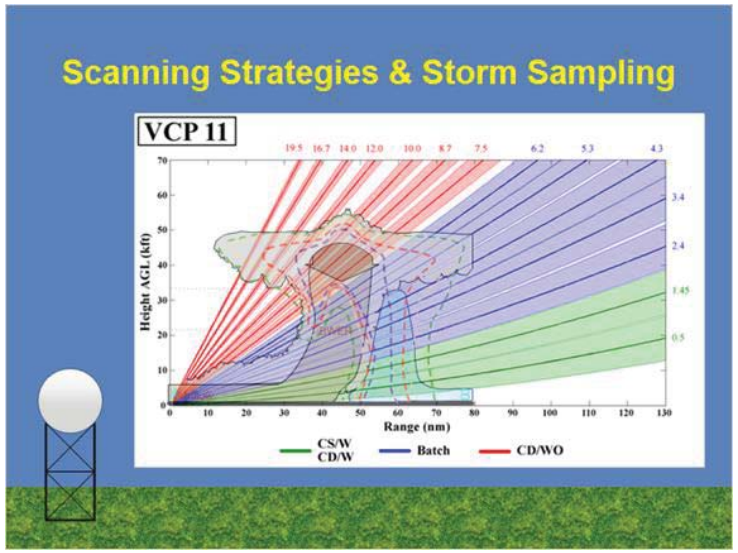
2.9 Radar Height Uncertainties



Notes:

However, look at VCP 11… The uncertainty doesn't really become a factor until past a range of 75 km and a height of 8 km. So, choosing a scanning strategy with more elevation scans can decrease the likelihood of increased height uncertainty in the height estimations.

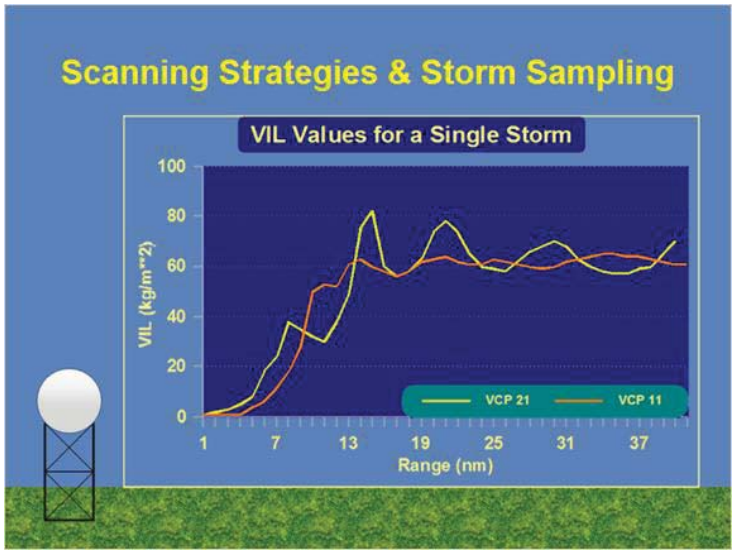
2.10 Scanning Strategies



Notes:

However, if we switch over to VCP 11, notice how those gaps in the upper levels are filled in with more elevation scans. So, different scanning strategies can lead to better or worse sampling. Be mindful of the scanning strategy you are choosing.

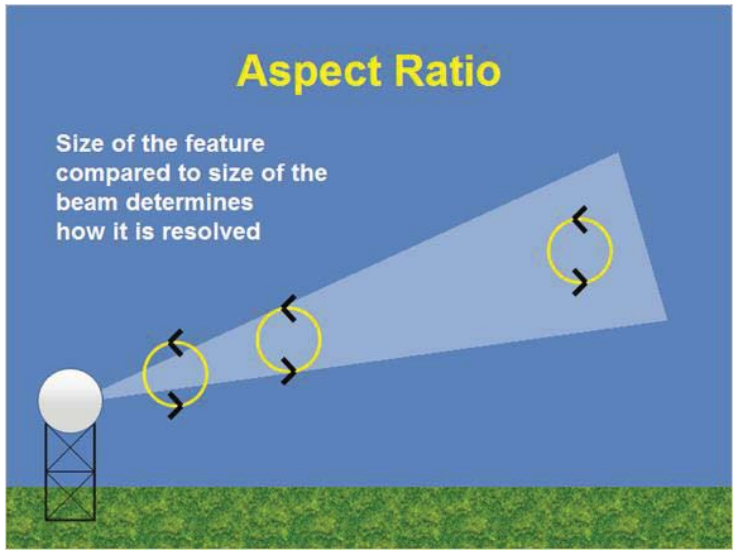
2.11 VIL Differences between Strategies



Notes:

Because sampling is affected by the scanning strategy you choose, so do the algorithm outputs. For example, here is a range plot of the VIL values for a single storm based on VCP 21 and VCP 11 scanning strategies. Notice how they are not the same. In fact, at some ranges, the difference between the two is quite large! Does this mean the storm is less intense because you chose a different VCP? Not at all! It's just that the storm is not being sampled properly.

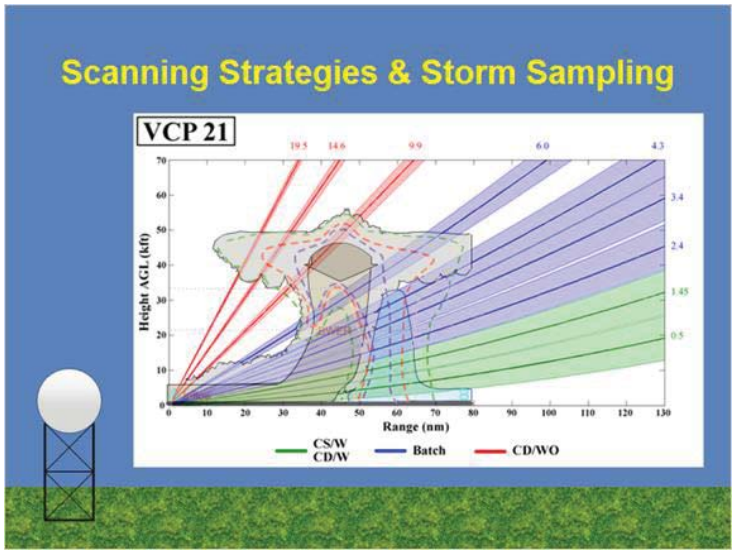
2.12 Aspect Ratio



Notes:

Because the beam spreads out as it propagates down range, a feature of constant size will appear differently at different ranges. The size of the feature compared to the size of the beam is called aspect ratio. How does it affect radar sampling issues? Take for example a circulation feature. At close range, multiple beams will sample the circulation and you'll get a detailed view. At medium ranges, the circulation will be roughly the size of the beam, and so you'll still get some detail, but not as much as if multiple beams are sampling it. At far ranges, only a portion of the beam is sampling the feature, and here you'll get very little detail and, in fact, the feature may not be resolvable at all. Basically, the farther away a feature is, the less impressive a feature will appear, and it will have less detail.

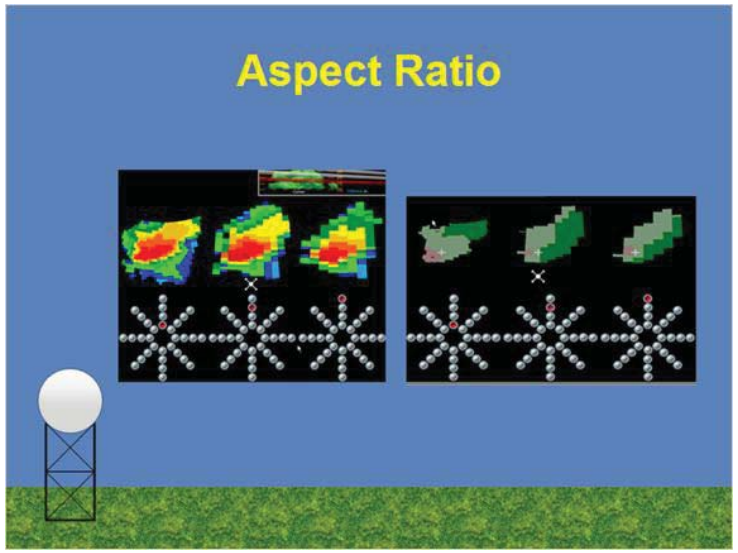
2.13 Scanning Strategies



Notes:

Suppose we have this supercell located down range of the radar. If we were to sample this storm using VCP 21, this is how the sampling would look like. Notice how near the low levels, sampling is such that there aren't many gaps. But, when we get to the higher elevations, there are significant gaps. In fact, some of the gaps are large enough that we possibly could miss some important features like the Bounded Weak Echo Region, etc. However, if we switch over to VCP 11, notice how those gaps in the upper levels are filled in with more elevation scans. So, different scanning strategies can lead to better or worse sampling. Be mindful of the scanning strategy you are choosing.

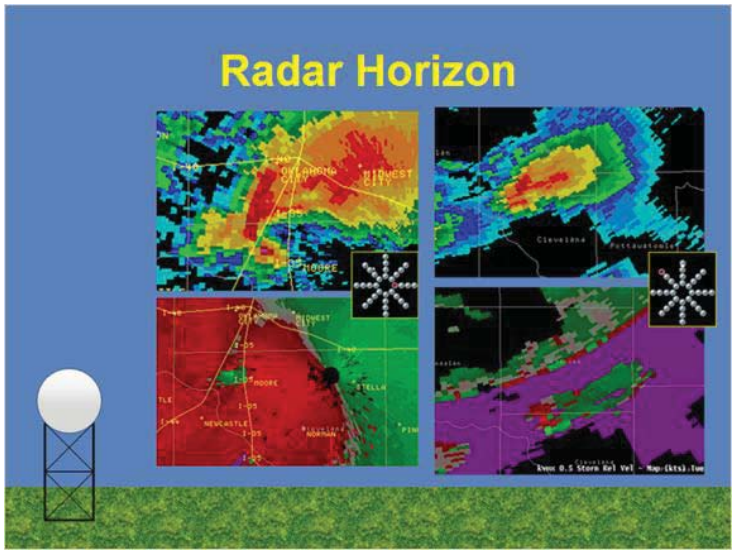
2.14 Aspect Ratio EXAMPLE



Notes:

Here is an example of a storm that is located south of 3 different radars. We'll take a look at the reflectivity and velocity images. The left image is from a radar located very close to the storm. The middle image is from a radar located at mid-ranges from the storm. And, the right image is from a radar located at far ranges from the storm. In reflectivity, notice the detail you see in the image on the left… a very nice reflectivity gradient to the southeast and pixel sizes are small. At mid and long ranges, the reflectivity appears a little more washed out and blocky. Granted, the radar beam is at different heights for each case, and we'll get to that next. But, for now, just pay attention to how the storm just looks different and maybe a little less intimidating at the mid and far ranges, than the close range. For the same storm, toggling over to velocity shows the same pattern. There is a good circulation noted in the radar closest to the storm, but the farther ranges either detect much weaker circulation.

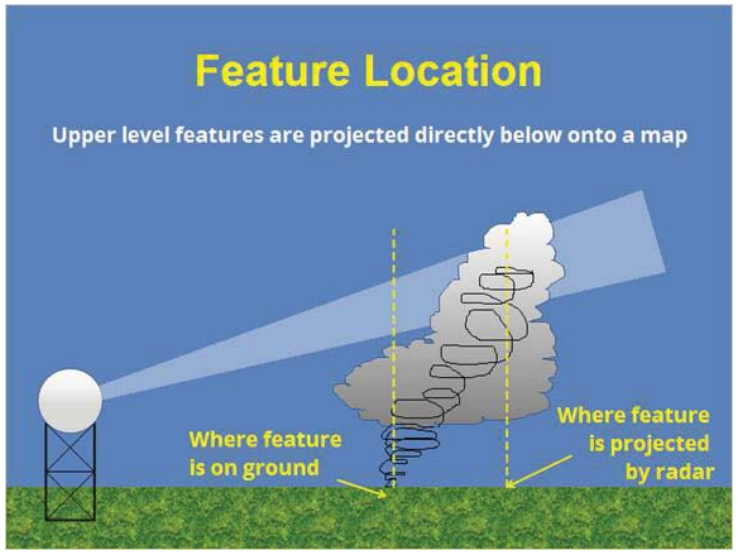
2.15 Radar Horizon EXAMPLE



Notes:

Here is a supercell thunderstorm producing a tornado moving through a major metro area. The two images on the left are from a radar to the east that is very close to the storm. The image on the right is from a radar farther away to the northwest. Notice how on the left, you can see fine-scale features in the reflectivity and velocity, but from the radar to the northwest, you aren't seeing these features. This is because the beam from the radar to the northwest is overshooting these low-level features. Therefore, if low-level features are your primary concern, then make sure you keep this limitation in mind and look for the closest radar.

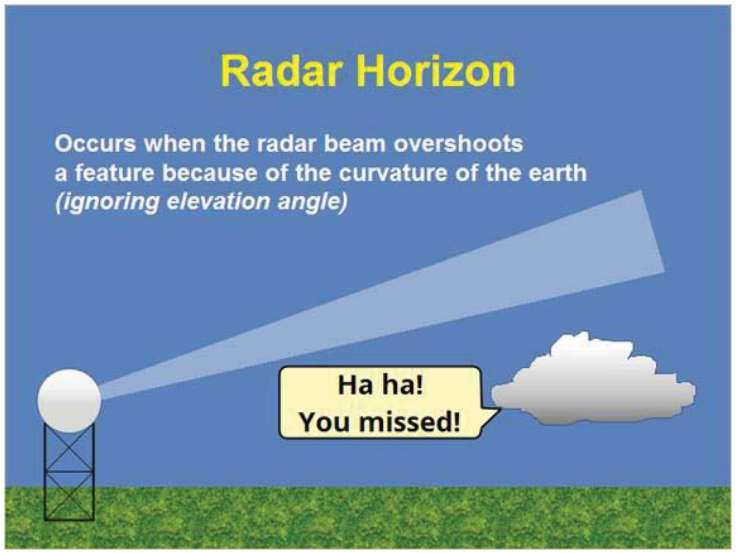
2.16 Feature Location



Notes:

As we've mentioned before, the radar beam progressively gets higher and higher as the range from the radar increases. However, when we look at radar images on a screen, we are projecting these features onto the ground directly below. Since there is often a tilt to features such as circulations associated with tornadoes, or the decent of hail cores translates horizontally, when we detect features aloft, this usually means it is not going to happen at the surface in the exact location we have projected it on a map. Look at this illustration of a tornadic circulation. The beam intercepts this circulation near the top of the storm, so when it is projected onto a map, it shows this circulation happening over this area on the ground. However, the actual circulation is occurring below the radar beam on the ground in this location. Therefore, be careful with how much preciseness you infer from the circulation you see on radar.

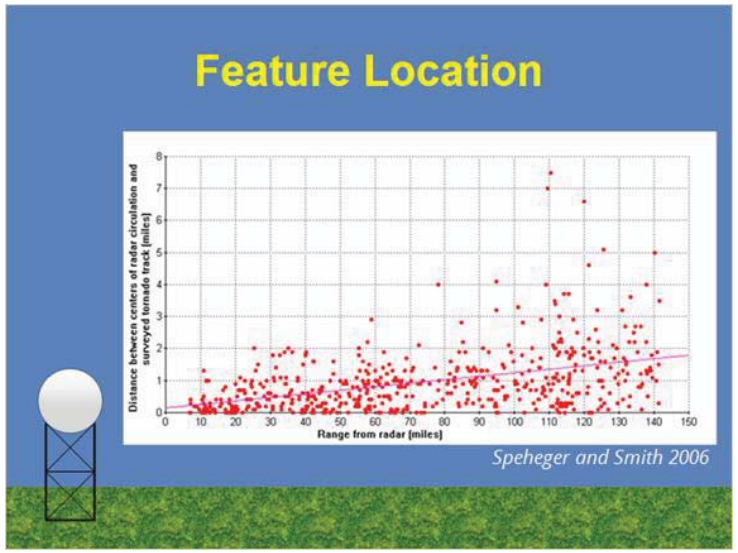
2.17 Radar Horizon



Notes:

Recall as the beam propagates away from the radar, the earth's curvature causes the beam to become higher and higher off the ground (beyond the rise due to elevation angle). This basically means low level features far from the radar will not be seen by radar. Therefore, you will either have to switch to a closer radar, or draw conclusions based on the information you do have. Let's look at an example.

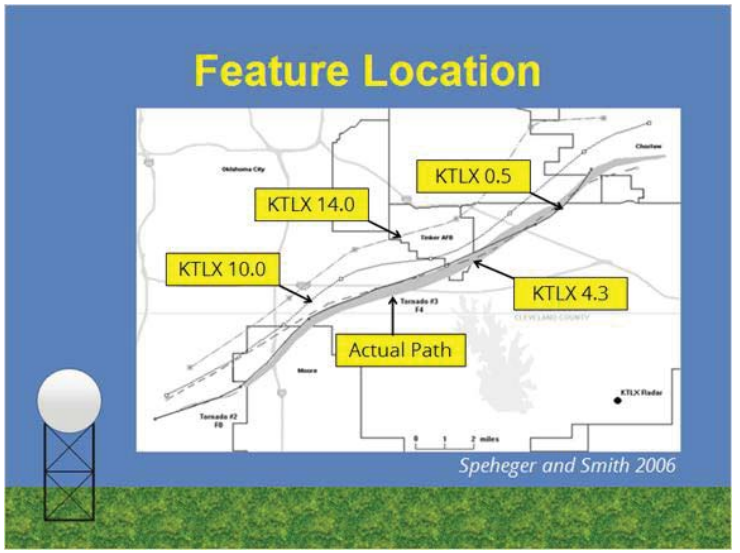
2.18 Feature Location EXAMPLE



Notes:

This study by Speheger and Smith (2006) shows just how far a radar identified circulation exists relative to the actual circulation on the ground. Near the radar, roughly within 50 miles, the error will be small (< 2 miles). However, once you get further away from the radar, these errors can become as large as 7-8 miles! Therefore, be careful with how exact you portray features on radar with where they are occurring on the ground.

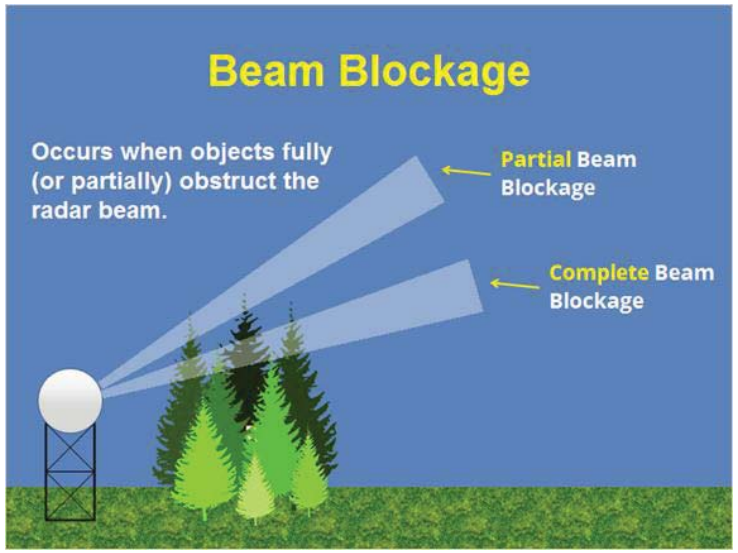
2.19 Feature Location EXAMPLE 2



Notes:

Here's an example of how elevation angle can also affect feature location. This is also from Speheger and Smith (2006). Notice how the radar identified circulations as you go higher and higher up do not coincide exactly with the actual path. Therefore, always be careful assuming a ground location when looking at radar data, especially in the upper levels.

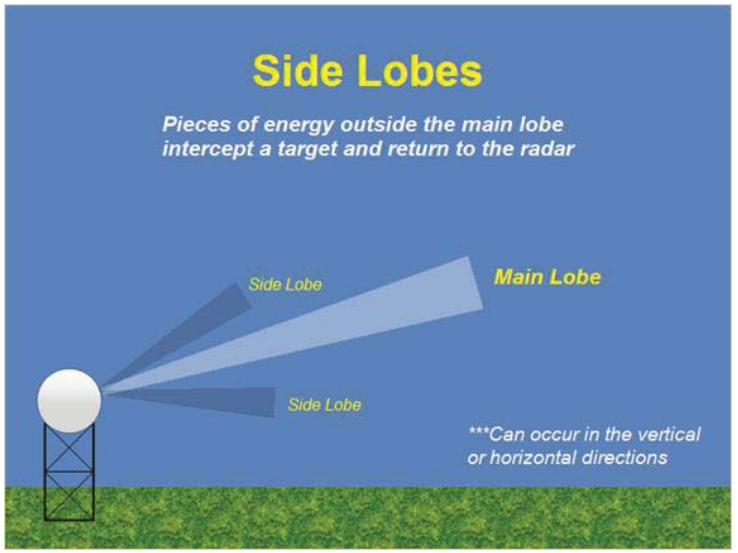
2.20 Beam Blockage



Notes:

When tall objects reside near the radar, they can obstruct the beam and cause the power loss to either be great enough that no features in that line of sight will be visible on radar, or the power returned by the features in that line of sight are much weaker than they would be in the obstructing objects were not there. These objects can be anywhere from trees across the street to giant mountain ranges. In this example, the lower beam is experiencing full beam blockage where the upper beam is experiencing partial beam blockage.

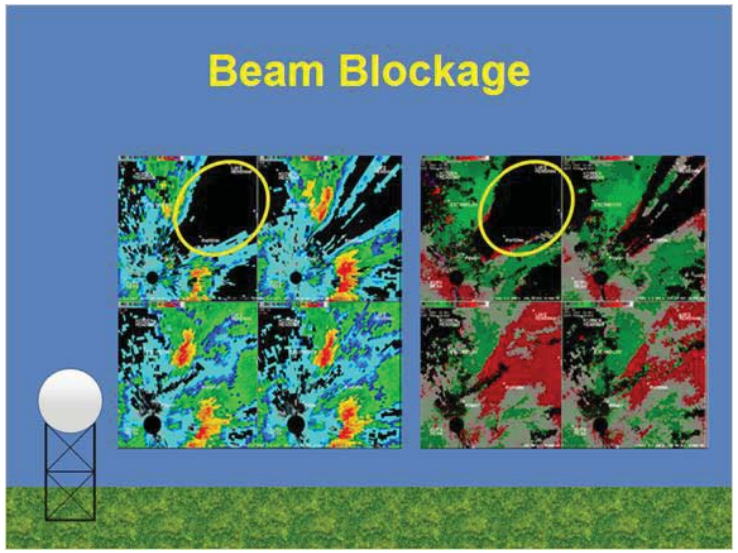
2.21 Side Lobes



Notes:

Recall from the previous section of this topic that there are pieces of energy from the transmitter that exist outside the main lobe that can intercept targets and return power to the radar and show up on your display. These side lobes can occur in the vertical and horizontal. Most of the time, the main lobe is sampling something at the same time the side lobe is sampling another target and the power return from the main lobe dominates, but in some instances, the main lobe will be sampling nothing and the side lobe returns dominate. Let's briefly look at some examples…

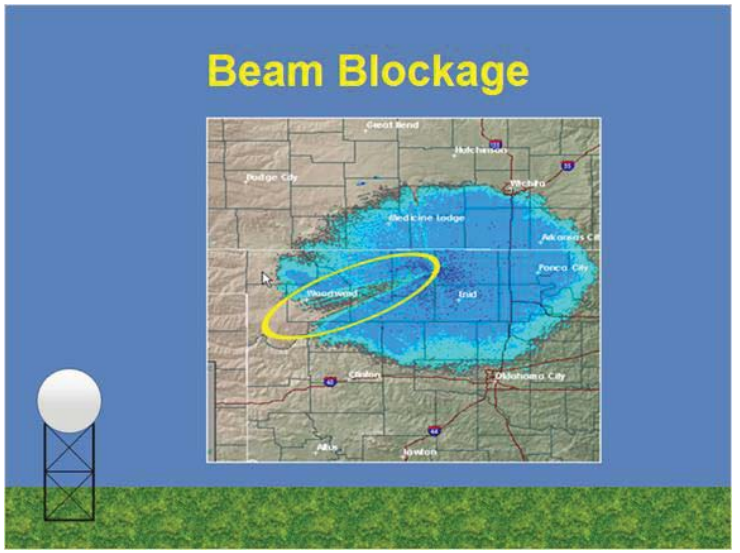
2.22 Beam Blockage EXAMPLES



Notes:

Here is an example of beam blockage from a nearby radar. The lowest elevation angle (top left of both images) shows a data gap to the northeast of the radar. Does this mean there are no storms in this area? No! This area is blocked by a nearby obstacle. And the obstruction is powerful enough to completely block the beam. In this particular example, the beam blockage was blocking a tornadic signature in the lower levels. Aloft (lower two panels), you can see there is a storm there with rotation.

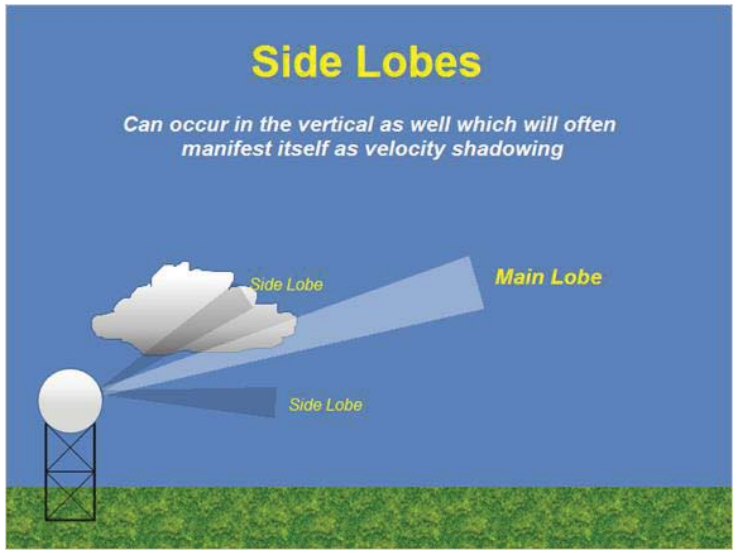
2.23 Beam Blockage EXAMPLES



Notes:

In this example, there is no complete beam blockage at the lowest level, but there is obviously some partial beam blockage. In instances like this, it is good to remember this because in these areas, you will get returns, but they will be weaker than reality. Therefore, you will need to mentally adjust your expectations.

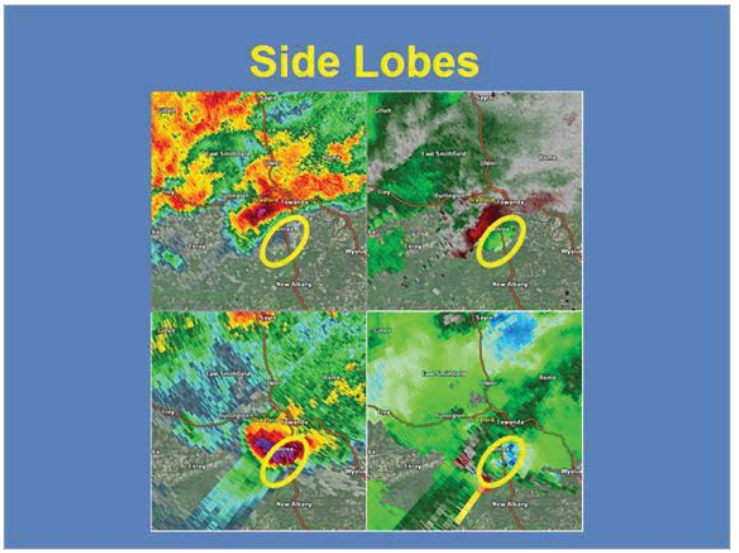
2.24 Side Lobes EXAMPLE 2



Notes:

The last example of side lobe contamination was in the horizontal. However, side lobes exist in the vertical as well. The most common scenario where you will see the effects of side lobe contamination is when the main lobe is sampling the low levels, but the side lobe is sampling a fairly intense overhang. However, the reflectivity will not be the prominent feature, but rather the velocity signature. This is often called the velocity shadow. What happens is the velocity signature from the overhang is pretty much superimposed in the low levels. This can lead to spurious velocity couplets.

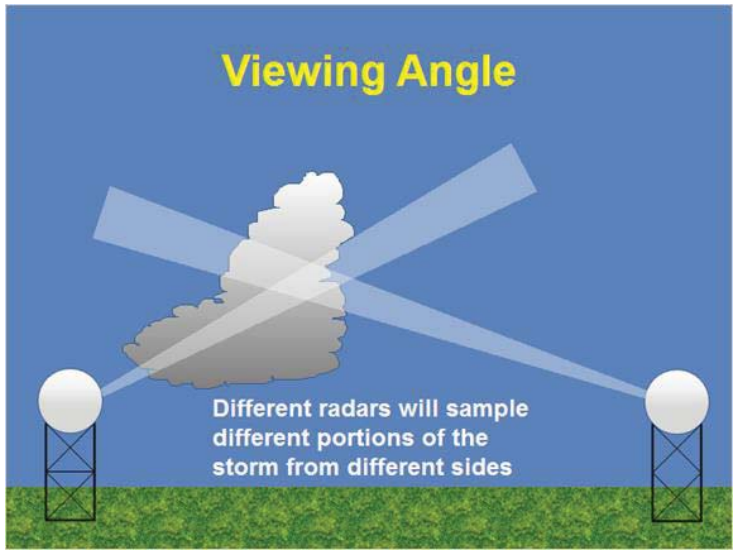
2.25 Side Lobes EXAMPLE 2



Notes:

Here is an example. The top two images are the low level reflectivity (left) and velocity (right). Notice the intense inbounds well away from the core and near the core there are moderate outbounds which make it appear as though there is rotation. However, stepping up in elevation, there is as strong core directly above with very strong inbound velocity. These strong inbounds aloft are basically being superimposed below because of side lobe contamination. Therefore, be aware of this limitation when viewing velocity values in weak signal areas.

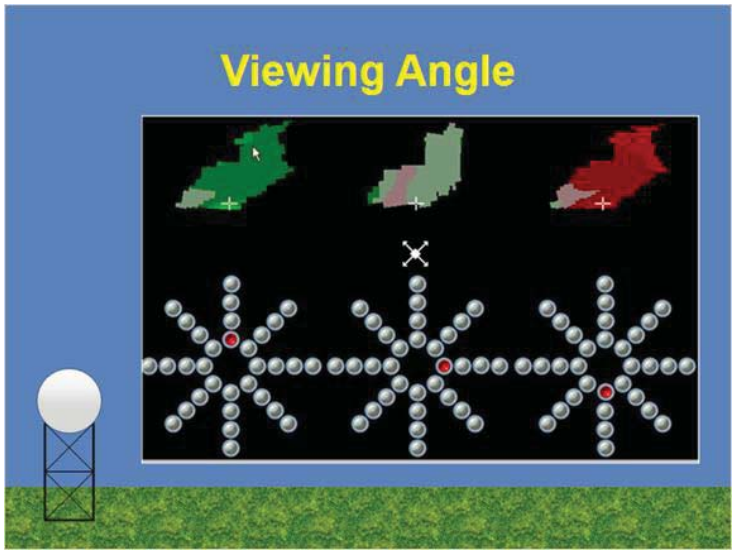
2.26 Viewing Angle



Notes:

Since the radars are somewhat close together throughout most of the US, there are some areas where storms will be sampled by multiple radars. Because storms are not symmetrical in the least bit, getting different views of the storms may be helpful, but this viewing angle can cause some problems such as masking circulations or missing three-body scatter spikes. Let's look at a few examples…

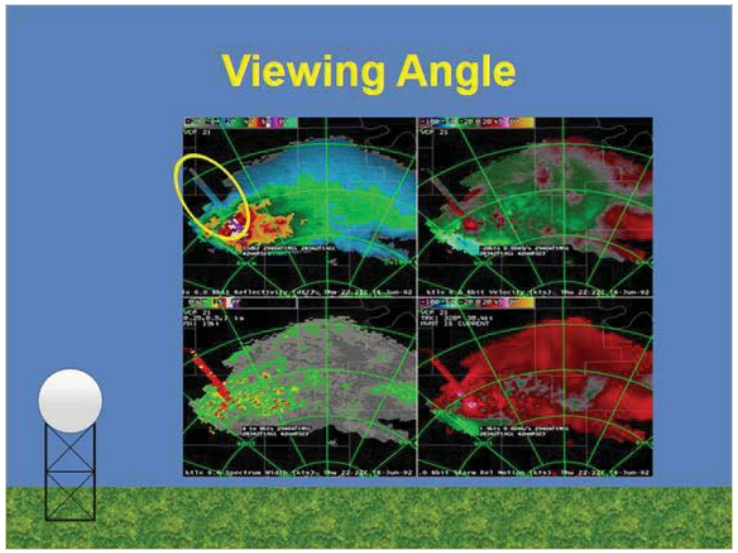
2.27 Viewing Angle EXAMPLE



Notes:

Recall that if the wind is blowing down the radial, you measure all of it. If it's blowing perpendicular to the radial, you measure none of it. Usually it's somewhere in between. This image shows the same radial velocity data when viewed from 3 radars, one north of the echo, one east of the echo, and one south of the echo (the echo is assumed to be in the middle of the "grid" of radars). Notice how the interpretation of the wind field changes with the radar you choose while the actual wind field is not changing. Therefore, keep in mind where your radar is located, and switch radars if you can to get a better look, especially when looking at velocity signatures.

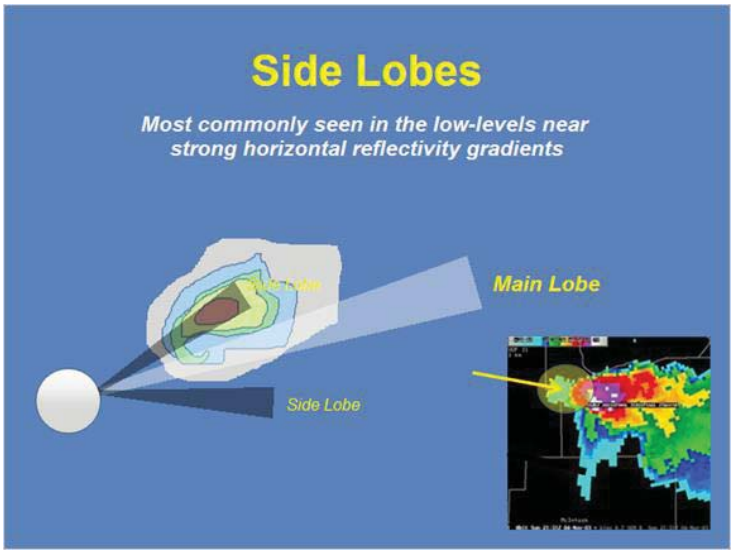
2.28 Viewing Angle EXAMPLE



Notes:

While the velocity is often the victim of viewing angle issues, reflectivity is not immune. Recall the three-body scatter spike (TBSS) is a feature seen in reflectivity that is often down radial of significant hail cores. Here is an example. Notice the TBSS to the northwest of the storm. If another storm was located in this area, the TBSS may be obscured by that storm, but could be visible if the storm was sampled by a radar in a different location. Also, TBSS tend to mess up velocity estimates and TBSS can sometimes occur within the mesocyclonic region, so if a radar beam passes through the hail core first and then into the mesocyclone area, then there could be some potential issues. Therefore, keep viewing angle in mind when looking at velocity and reflectivity signatures.

2.29 Side Lobes EXAMPLE

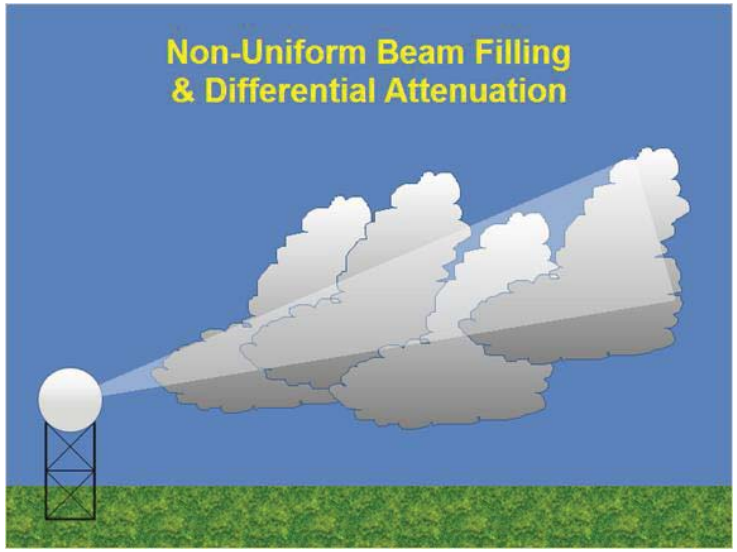


Notes:

Here is a typical conceptual model of a supercell thunderstorm. Once the main lobe passes by the core of the storm, the side lobe samples the core while the main lobe is sampling very low returns. Because the signal returned to the radar is dominated by the side lobe returns, that is the signal processed by the radar. However, because the radar thinks the return came from the main lobe, it places this return where the main lobe is sampling, which is just off to the side of the core in the clockwise direction. So, side lobe contamination will show up as weak reflectivity just to the side of a core. Here is an example of side lobe contamination. The main lobe has passed by the main core and is sampling clear air, but the side lobes are still picking up the main core. So, there is return power at the radar from the side lobes that is dominating the signal but because the main lobe is off to the side, that is where the radar is thinking the power is coming from.

2.30 Non-Uniform Beam Filling

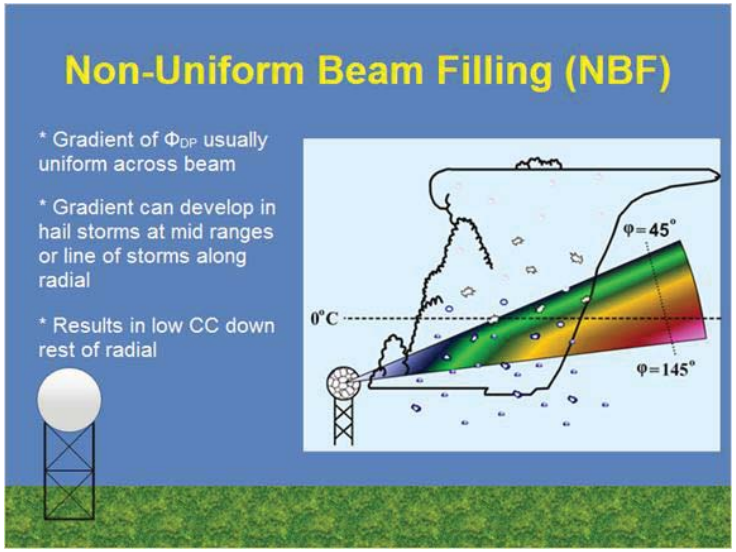
& Differential Attenuation



Notes:

The last two sampling issues we'll look at primarily impact the dual-pol variables of correlation coefficient and differential reflectivity. These two issues are non-uniform beam filling and differential attenuation. As you'll learn soon, these two sampling issues often occur together, but are caused by completely different processes. The most common situation where both of these occur are in very heavy cores with hail present, and when storms align along a radial. So, let's take a look at both of these issues a little more…

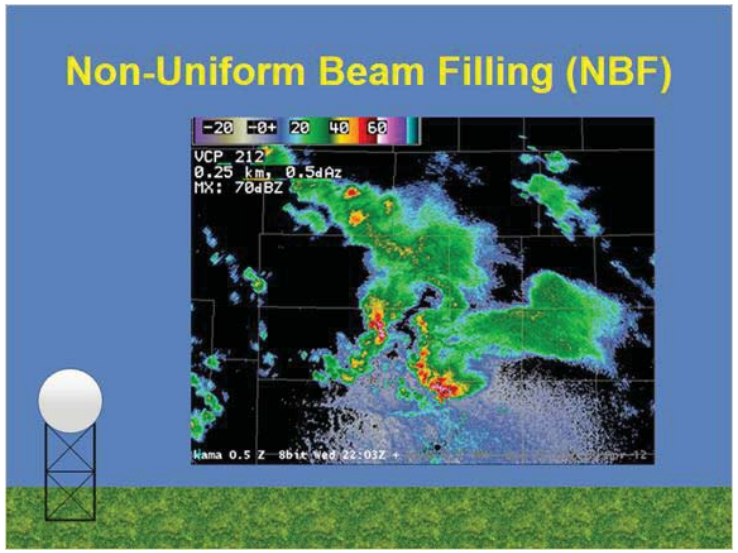
2.31 Non-uniform Beam Filling



Notes:

Most of the time, the gradient of precipitation across the beam is such that differential phase shifting is fairly uniform across the beam. However, with hail storms at mid ranges from the radar, the beam can experience little phase shifting at the top of the beam where it is sampling hail, but have large phase shifting at the bottom of the beam where it is sampling rain. You'll learn more about phase shifting later in this topic, but for now this will suffice. This large gradient of phase shifting across the beam causes the correlation coefficient to be reduced and this effect is seen down the rest of the radial. So, NBF is often noted where low CC is noted from the source of the NBF all the way down the radial. Another common place to see the NBF effect is when a long line of storms align along a radial. Let's look at a couple examples…

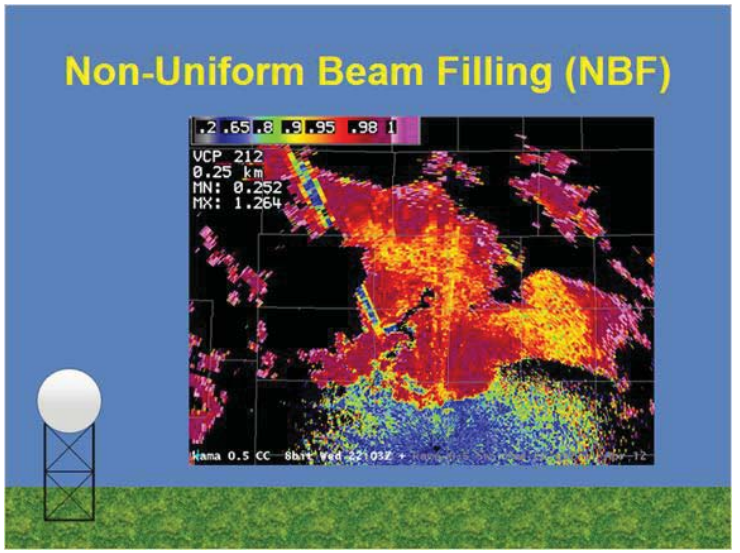
2.32 NBF Example



Notes:

Here is an example of some storms that are causing non-uniform beam filling (NBF) issues. Can you determine which ones are causing issues?

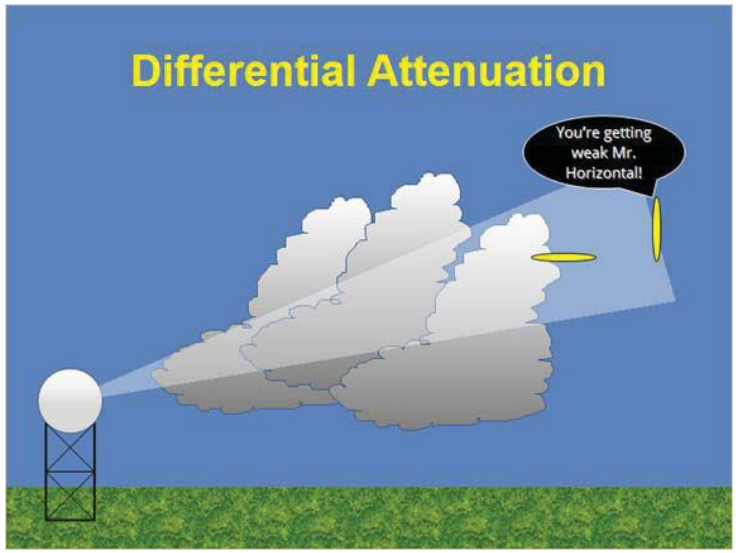
2.33 NBF Example



Notes:

Let's look at the CC product and see if that helps… Better, worse, about the same? Hopefully you said better :) The storm to the northwest is experiencing significant NBF issues and that is evident by the significant drop in CC from where the hail storm is residing all the way down the rest of the radial. The other storm causing issues is just to the north of radar. While the CC drop is not as prominent, there are noticeable drops.

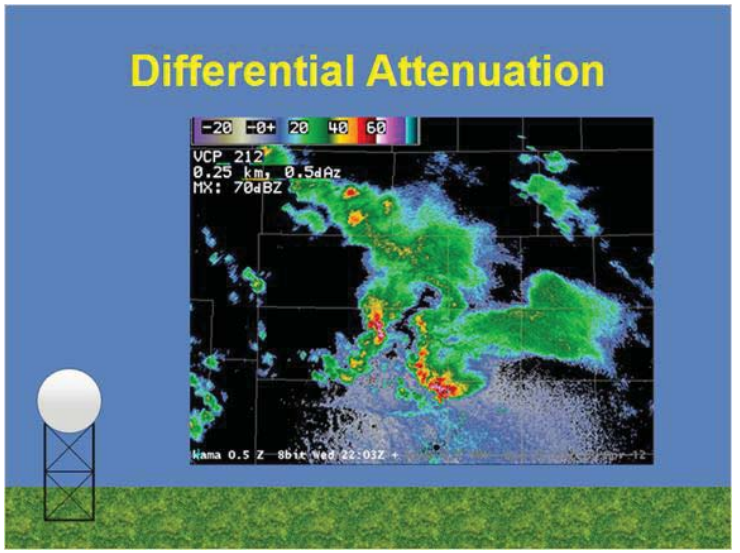
2.34 Differential Attenuation



Notes:

As the two polarizations propagate through the atmosphere, they usually attenuate (or slow down/get weaker) at the same rate because S-band radiation usually does not attenuate much in weather. However, there are some instances where horizontal pulses will attenuate quite a bit, but the vertical pulses will not. This is usually near storms with very heavy rain/hail cores, or when many storms align along a radial. In these instances, the vertical pulse will not slow down/get weaker nearly as much as the horizontal pulse. As a result, when differential reflectivity is computed in areas of differential attenuation (which is the horizontal power return minus the vertical power return), you get a slightly lower value because the vertical is stronger than the horizontal. And many time, these values go negative. Let's look at an example…

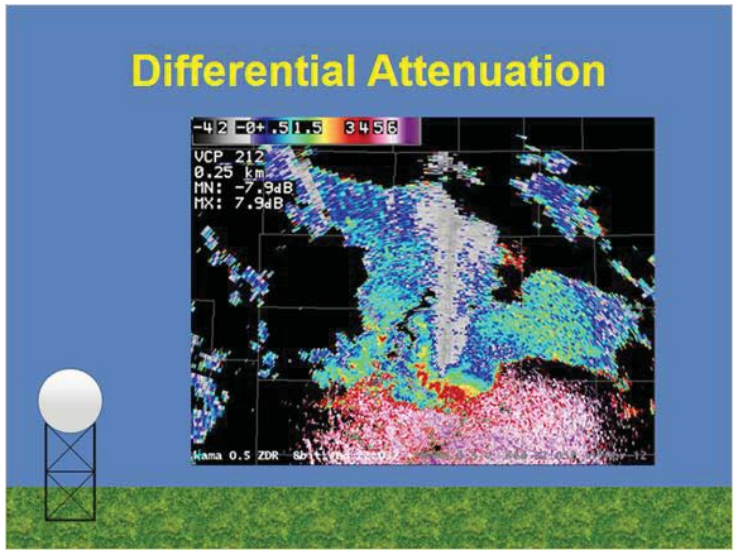
2.35 DA Example



Notes:

This example ought to look familiar :) This is the NBF case, so why I am I showing it here? Because most NBF situations also cause differential attenuation. Look downstream of the two storms we noted before. We don't see a noticeable change in the horizontal reflectivity even though there is probably some attenuation happening in the horizontal channel.

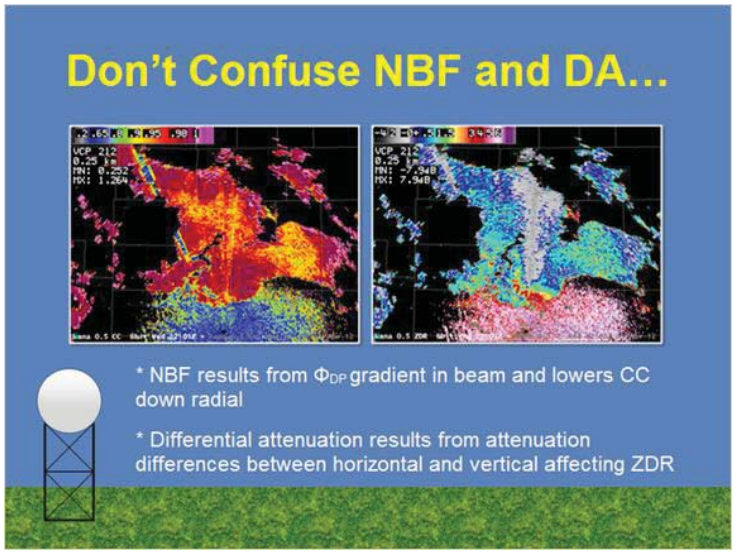
2.36 DA Example



Notes:

However, look at the ZDR product… Behind both of these heavy rain/hail cores, we see significantly lower ZDR values down the rest of the radial. This is the result of differential attenuation.

2.37 NBF/DA Differences

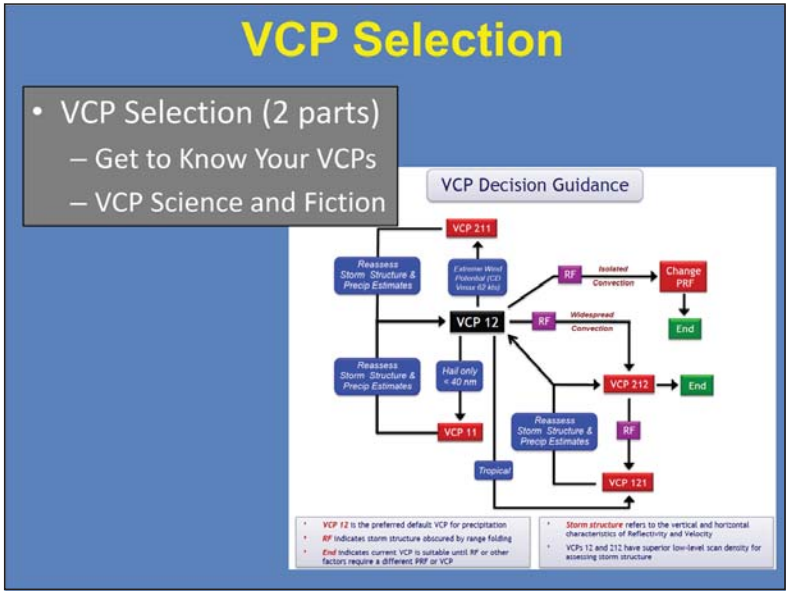


Notes:

So I alluded to this before, and I'll state it here again. NBF and differential attenuation often occur in the same situations (heavy hail/rain cores or storms along a radial), but they affect different variables for different reasons. NBF results from a PhiDP gradient within the beam and lowers CC the rest of the radial, and differential attenuation results from attenuation differences in the horizontal and vertical causing ZDR to become lower than expected.



Welcome to Topic 2's lesson on VCP Selection training, with a little bit of VCP fact or fiction built-in, but this is get you familiar with all the VCPs the Doppler Radar has to offer, so you'll be equipped when you need to make some choices.



In this lesson on VCP Selection, we will cover the two main parts we want to address in the VCP decision-making process. First, you have to know your VCPs, what their strengths and limitations are – and how they compare to one another. Then, we'll spend some time debunking some VCP myths or fiction in this case, with science. Armed with those two, you should be able to choose the right VCP for the right situation.

Get to Know Your VCPs Objectives

1. Identify distinguishing characteristics of each VCP, such as vertical sampling, editable Doppler PRFs, or range unfolding techniques.

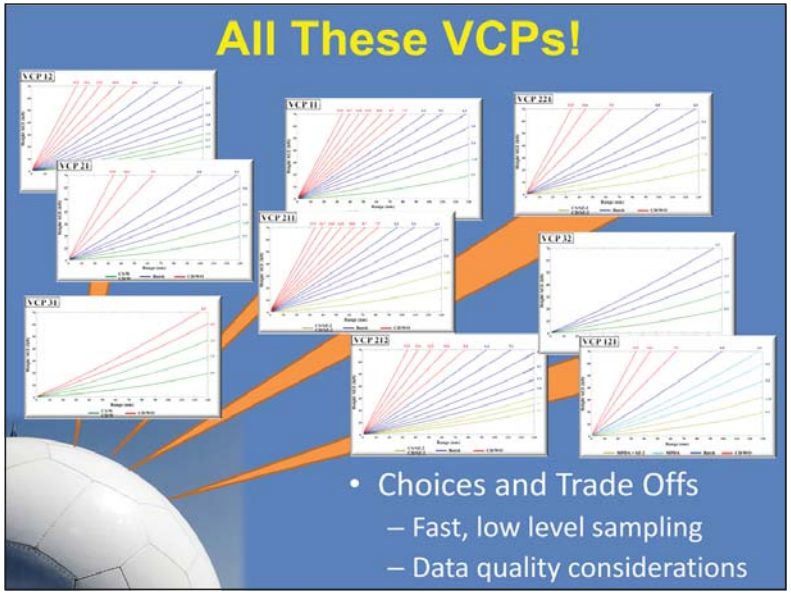
2. Identify the newer software-based features that enhance current VCPs function better for real-time meteorological situations.

For this lesson's objectives, we'll split it into two slides, with this one showing the learning objectives for getting to know the individual VCPs, their strengths and limitations, and how they compare to one another.

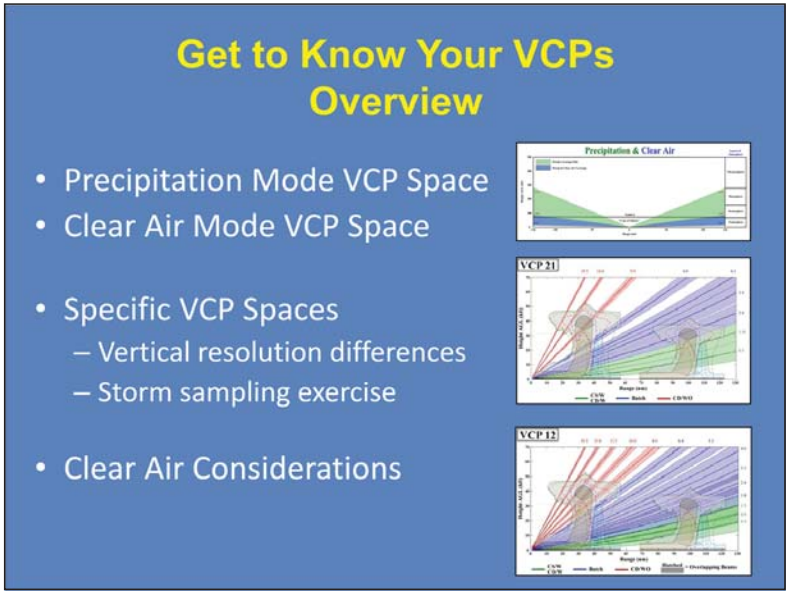
VCP Science and Fiction Objectives

1. Identify the difference between normal wear and tear vs. excessive wear and tear.
2. Identify the impact of the NEXRAD Technical Requirements on antenna system wear and tear.
3. Identify the two characteristics of antenna motion that contribute to wear and tear.
4. Identify best practices that enable use of all the VCPs with minimum stress to the system.
5. Identify the benefit of VCP 31 in clear air or light precipitation events.
6. Identify the impact of faster VCPs on rainfall estimation.

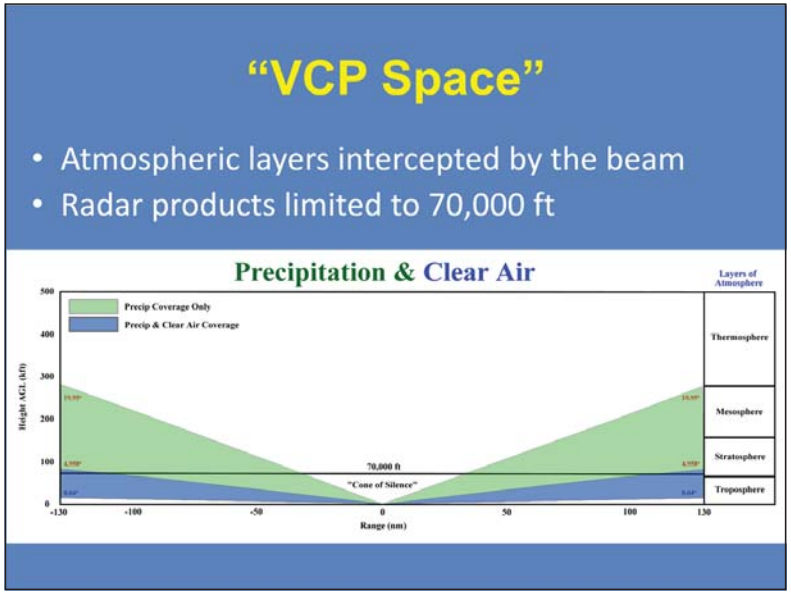
Here are the learning objectives for VCP Science and Fiction. Notice that the first four focus on VCP choice and how it relates to wear and tear on the WSR-88D hardware, which has been a significant concern with the introduction of the faster VCPs. The remaining objectives address the use of VCP 31 and the impacts of faster VCPs on rainfall estimation.



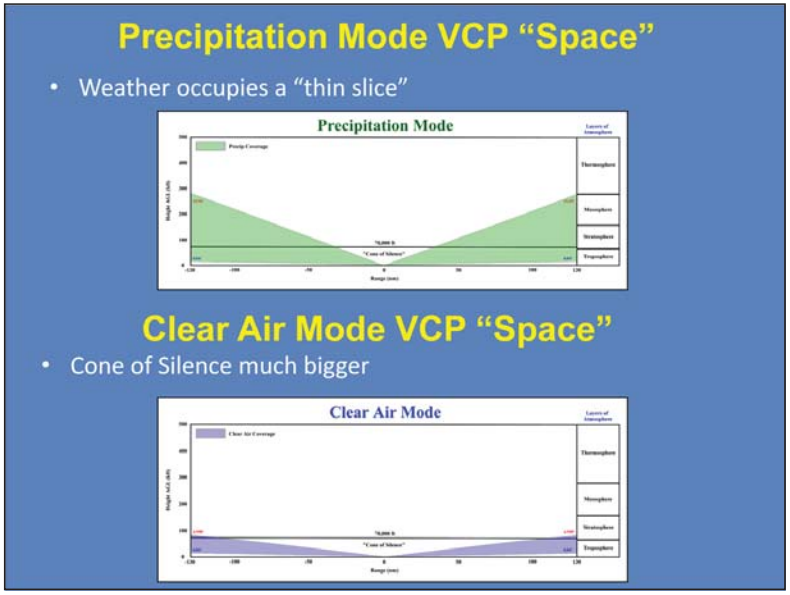
Currently, there are nine different VCPs to choose from, which has evolved over the past couple of decades with from changing needs and derivative capabilities of older VCPs. Your choices involve trade offs, such as the benefits of fast, low level sampling vs. the impact on data quality of fewer pulses per radial. We'll try to demystify the options you see here, showing some advantages and disadvantages of each.



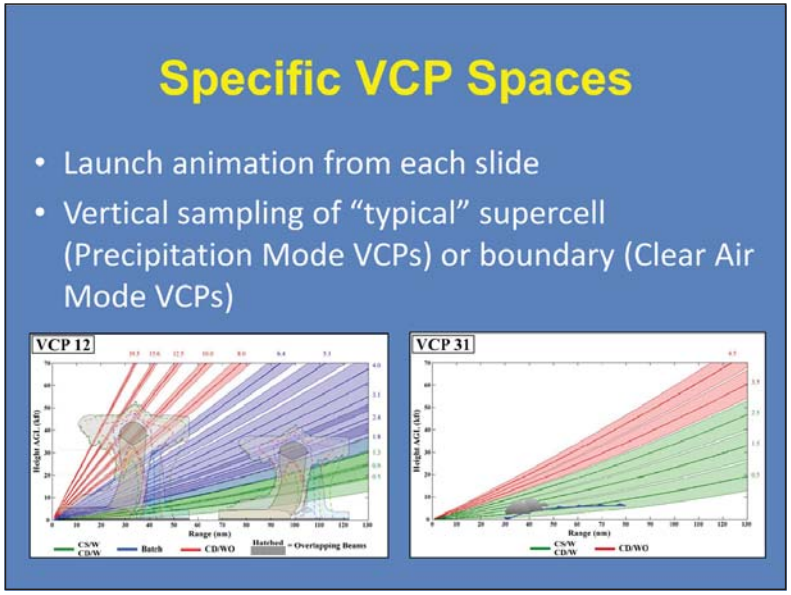
To set the foundation for getting to know your VCPs, we'll start with the atmospheric space that Precipitation Mode and Clear Air mode sample. Then the space sampled for each of the VCPs is presented, along with animations of the angles used during a volume scan. These animations will help you to visualize the elevation angle spacing for the different VCPs. Just from the two still images on this slide, look at the differences between VCPs 12 and 21 with respect to mid and upper level sampling of a storm at close range. Then, we'll close out with some outside the box thinking for our two Clear-Air mode VCPs.



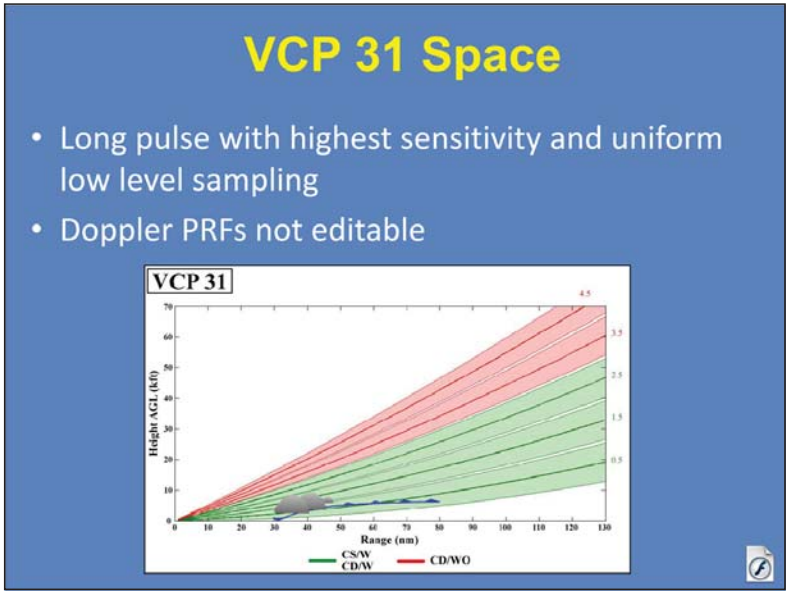
Let's look at a quick cross-section of the quote VCP space, or the amount of atmosphere that's actually being sample by a doppler radar, both in clear air and precip modes – and showing just how vast the cone of silence is above a particular radar. This image is mostly a reminder of how thin the Tropospheric layer is and how remarkable it is that such dramatic weather occurs in this layer. Why is there a line at 70,000 ft? It's because radar data are not assigned to products any higher than 70,000 ft!



Here's a look at the both modes separately. The cone of silence is always something to remember and hopefully you have sufficient adjacent radars to mitigate some of the data loss. Compare the size of the cone of silence for Clear Air mode. By design, it covers much less atmosphere because nothing is going on, at least, if you are in Clear Air Mode, there should be nothing going on...with a few exceptions, which we will discuss later.



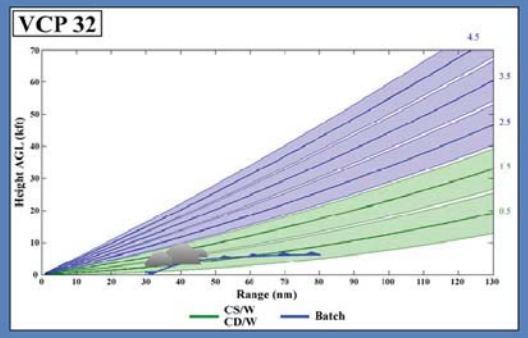
Now for a series of animations of the spaces that each VCP occupies. The following slides will each present a particular VCP, and you can launch the animation from that slide. For the Precipitation mode VCPs, you'll see a couple of "typical" (not too big, not too small) supercells, for you to see the sampling strengths and limitations of each VCP. For the Clear Air mode VCPs, you'll see a frontal boundary with small cumulus forming, since convective initiation is one of the applications of the Clear Air VCPs.



We'll start off with the Clear Air Mode VCPs. VCP 31 is the only VCP that uses long pulse, offering the highest sensitivity with uniform low level sampling. The Doppler PRFs are held constant and are not editable.

VCP 32 Space

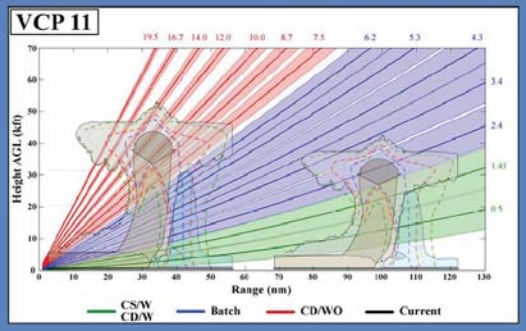
- Uniform low level sampling
- Doppler PRFs editable Split Cuts and Batch



VCP 32 is the short pulse Clear Air mode VCP, with uniform low level sampling. The Doppler PRF is editable for the Split Cut and the Batch elevations.

VCP 11 Space

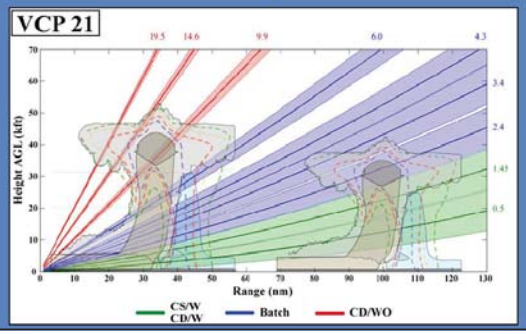
- Generally uniform vertical sampling
- Doppler PRFs editable Split Cuts and Batch



Now to the Precip Mode VCPs. We'll start off with VCP 11, which was actually the original convection VCP with relatively uniform vertical sampling. The Doppler PRF is editable for the Split Cut and the Batch elevations.

VCP 21 Space

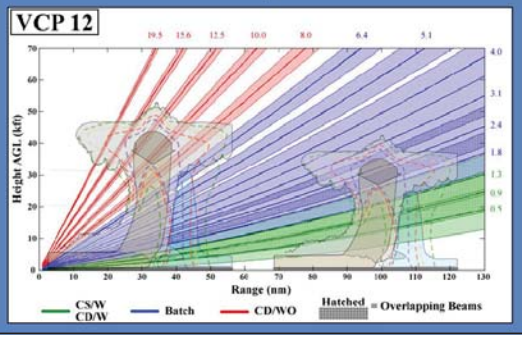
- Uniform low level sampling with significant gaps aloft
- Doppler PRFs editable Split Cuts and Batch



The widely used and widely defaulted, VCP 21 was the original stratiform precipitation VCP with uniform low level sampling, but significant gaps aloft. The Doppler PRF is editable for the Split Cut and the Batch elevations.

VCP 12 Space

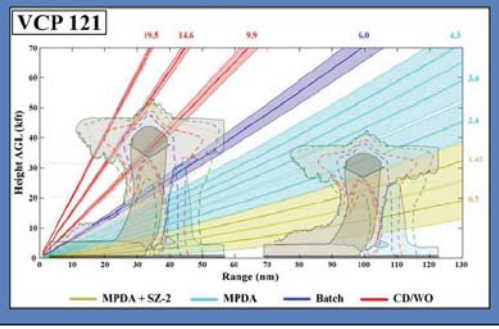
- Best low level vertical sampling with uniform gaps aloft
- Doppler PRFs editable Split Cuts and Batch



VCP 12 has the best low level vertical sampling with uniform gaps aloft. The Doppler PRF is editable for the Split Cut and the Batch elevations.

VCP 121 Space

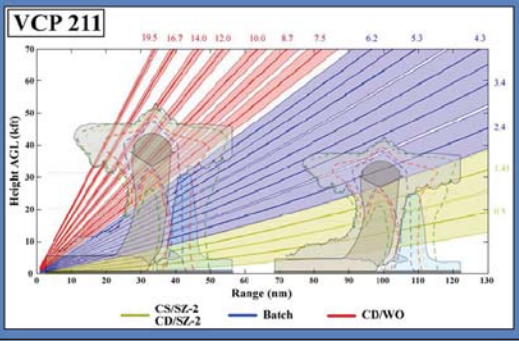
- Uniform low level sampling, significant gaps aloft
- Almost no RF data at low levels (MPDA & SZ-2)
- Doppler PRFs not editable



VCP 121 has uniform low level sampling with significant gaps aloft (just like VCP 21!). It's strength lies in the processing of the lower elevations, combining the SZ-2 and MPDA techniques. The result is almost no RF data for the lowest two elevation angles. The Doppler PRFs are held constant and are not editable.

VCP 211 Space

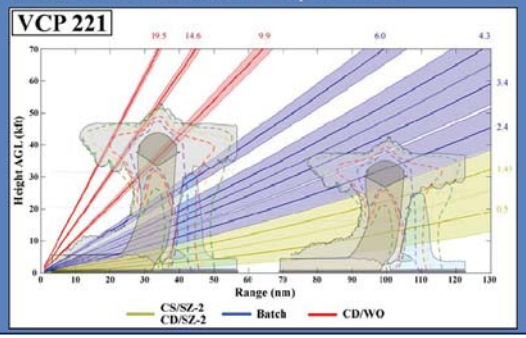
- Generally uniform vertical sampling
- Reduced RF data at low levels (SZ-2)
- Doppler PRFs not editable on Split Cuts



VCP 211 has relatively uniform vertical sampling, with reduced RF data due to SZ-2 processing on the lowest 2 elevations. The Doppler PRFs are not editable on the Split Cuts, but can be edited on the Batch elevations.

VCP 221 Space

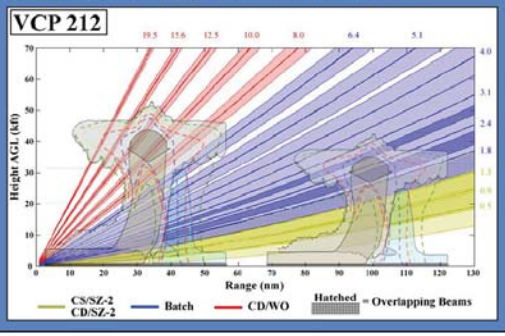
- Uniform low level sampling, significant gaps aloft
- Reduced RF data at low levels (SZ-2)
- Doppler PRFs not editable on Split Cuts



Like VCP 21, VCP 221 has uniform low level sampling with significant gaps aloft. Due to SZ-2 algorithm processing on the lowest 2 elevations, RF data is reduced. The Doppler PRFs are not editable on the Split Cuts, but can be edited on the Batch elevations.

VCP 212 Space

- Best low level vertical sampling, uniform gaps aloft
- Reduced RF data at low levels (SZ-2)
- Doppler PRFs not editable on Split Cuts



VCP 212 has the best low level vertical sampling with uniform gaps aloft. Due to SZ-2 processing on the lowest 3 elevations, RF data is reduced. The Doppler PRFs are not editable on the Split Cuts, but can be edited on the Batch elevations.

VCP Usage:
Radar Operations Center survey

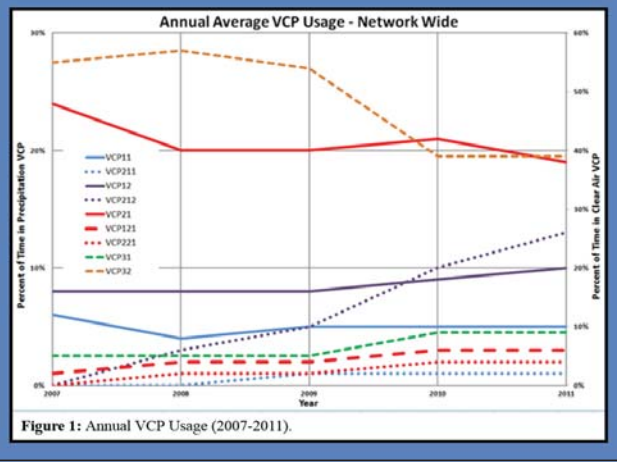


Figure 1: Annual VCP Usage (2007-2011).

The most recent figures from the Radar Operations Center were placed into this graphic, a survey of VCP usage nationwide, averaged out and graphed from 2007 to 2011. This period of time was occurring when offices were getting used to the newer SZ-2 VCPs, the 200-series VCPs. Though we don't have updated values since then, you can see the general trends here during this period. For one, notice the slight increase in VCP 31 and the larger drop in 32 for clear air modes. Then, we have most popular precip mode, 21, showing a flat-lining for a while but a bit of a recent decrease, which may be continuing, while the 200-series VCPs are starting to become more of a part of WFO operations, especially 212.

The Future of VCPs

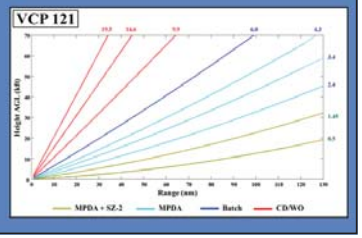
- Dynamic scanning has changed the way VCPs operate (AVSET, SAILS, MESOSAILS)

VCP	Without AVSET	Shortest Update with AVSET
11	~5 min	3 min, 12 sec (up to 6.2°)
12	~4.5 min	3 min, 10 sec (up to 6.4°)
212	~4.5 min	3 min, 30 sec (up to 6.4°)
21	~6 min	4 min, 55 sec (up to 9.9°)

One of the major changes to VCPs over the past few years is not even related to VCPs themselves – as in, making new ones or combining them, it is software changes that enhance the current ones. You all now know them as AVSET, SAILS, and the newer MESO-SAILS. AVSET terminates a volume scan early if there is no longer relevant data at the higher elevations. This means that products will update faster with AVSET, depending on the location and height of echoes. If storms are at close range, the entire volume scan up to 19.5 degrees may be needed. If storms are more distant, AVSET can terminate the volume scan once the beam is above the storms, giving you faster product updates. The table shows the shortest possible product updates for some of the VCPs. On top of that, we now have SAILS and MESO-SAILS which both serve to add 0.5 degree scans into one volume scan, so instead of just one, you can get an additional one, two, or three lowest cuts, depending on the setting you choose.

The Future of VCPs

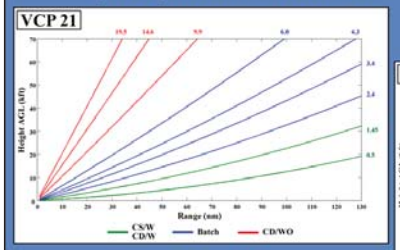
- Staggered Pulse Repetition Time (PRT) applied to Batch elevations
  - Reduces RF data and velocity dealiasing errors
  - Apply to VCP 121 Batch elevations
  - Increase product update rate to ~ 5 minutes!
- Build 17?

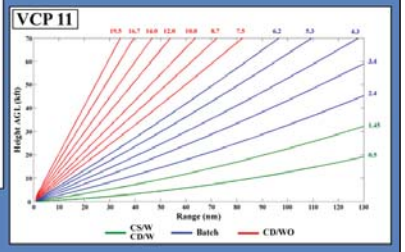


Staggered PRT is a technique that is currently being developed, reducing both RF data and velocity dealiasing errors in the Batch elevations. It is expected to be implemented with VCP 121, with an improved product update rate. This would result in very little RF data for both the Split Cut and Batch elevations for VCP 121.

The Future of VCPs

- Build 18 (not yet determined)
  - Replace VCPs 11 ***and*** 21 with a "general surveillance" precipitation mode VCP





As of this recording, it is already in the works to combine VCPs 11 and 21, creating a new general surveillance VCP. This combination takes the uniform angle spacing of VCP 11 and the slow antenna speeds of VCP 21 to generate a VCP with better vertical sampling and data quality for non-severe precipitation.

VCP Science and Fiction

- VCP ***design*** based on sampling the weather
- 9 VCPs to choose from! Oh my!
- VCP ***choice*** based on
  - Science?



  - Or Fiction?



Given all of these VCP choices, we'll now examine some of the factors affecting VCP selection in real time. Each VCP was designed to sample the atmosphere in a different way. What we want to address here is whether VCP choices are based on science or fiction.

VCP Science and Fiction

- VCPs were ***designed*** for weather
- VCP ***choices*** based on
  - Expected weather hazard(s)
  - Concerns about ***potential*** hardware impacts and unplanned outages
  - Are these sometimes in conflict?



VCP choice is certainly driven by the current or expected weather, especially hazardous weather. Another concern may be potential impacts on hardware due to use of the faster VCPs. These factors can be in conflict. A radar outage during a severe weather event is something we would all want to avoid.

Use of VCP 12, 212 and 121

- Do VCPs 12, 212 and 121 cause ***too much*** wear and tear on the system?
- Antenna ***speed*** matters
- ***Change*** in antenna speed ***also*** matters



We want to look at the true impact of the use of VCPs 12, 212 and 121, because these VCPs do have the fastest antenna rotation rates. Antenna speed does matter. Along with antenna speed is how often the antenna speed changes during a VCP. Speeding up and slowing down such a large antenna also imposes hardware stress.

VCPs 12, 212, and 121

- Most challenging to the hardware? Yes!
  - Avoid them? No!
- Wear and tear is ***within*** system design
- Use VCP(s) that ***support the NWS mission***





It is true that VCPs 12, 212 and 121 are the most stressful to the hardware. That does not mean that they should be avoided. Any associated wear and tear is within system design, and this will be explored in subsequent slides. Given thoughtful overall VCP usage, at appropriate times, VCPs 12, 212 and 121 can be critical to supporting the NWS mission.

Wear and Tear Contributors

- VCP Choice
  - Antenna rotation rate
  - Changes in antenna rotation rate
- Other factors
  - RDA environment (temperature, dust, moisture)
  - Lack of regularly scheduled preventative maintenance
  - Not following established maintenance procedures



VCP choice does contribute to wear and tear, specifically fast antenna rotation rates and frequent changes in antenna rotation rates. There are other, equally important factors. Some RDAs are located in pretty tough environments with respect to extremes in temperature, dust, or moisture. The greater these extremes, the greater the need for regularly scheduled preventative maintenance that closely follows established procedures.

Pedestal Wear and Tear

- Antenna speed ***and*** changes in the speed ***both*** affect
  - Gear Box for azimuth (small gear + bull gear)
    - Mechanically turns the antenna
  - Servo motor
    - Drives the gear box
  - Encoder
    - Sensor for antenna position/tells motor what to do



There are many parts of the antenna pedestal that are affected by both VCP choices and frequency of preventative maintenance. The gears (small gear and bull gear) for the antenna azimuth, the servo motor, and the encoder are some examples. The bull gear is the larger one because it lies at the base of the antenna. The small gear and the bull gear together mechanically rotate the antenna. The servo motor drives the gear box, while the encoder tells the motor what to do.

NEXRAD Technical Requirements

- System is operating ***below*** design limitations
- VCPs 12, 212 and 121
  - Do ***not*** result in ***excessive*** wear and tear
- Increased use of VCPs 31 and 32 decreases normal wear and tear
- Regularly scheduled preventative maintenance following established procedures still critical



Given the NEXRAD Technical Requirements, even the fastest VCPs are operating well within the system design. VCPs 12, 212 and 121 do not result in excessive wear and tear. Also, the Mode Selection Function has resulted in an increased use of VCPs 31 and 32, which are the slowest, decreasing normal wear and tear overall. Regularly scheduled preventative maintenance following established procedures is still critical.

Regardless of the "wear & tear" NWS Mission Comes First!

- Faster VCPs: use 'em when you need 'em
  - Back to slower VCPs when you don't
  - Think: "What supports the mission?"
- Combine pro-active maintenance and thoughtful VCP usage
  - Protect the pedestal components and the public!

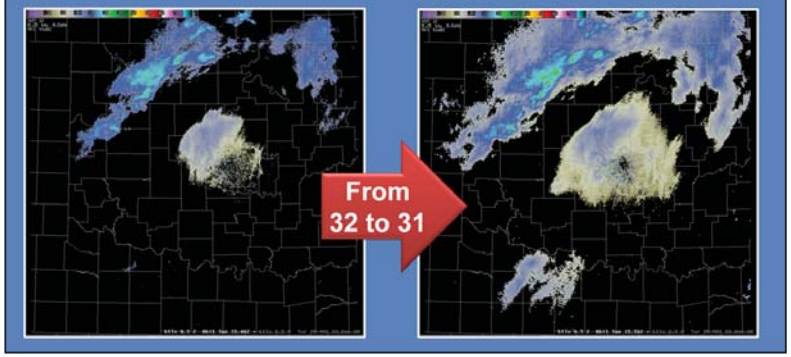




The bottom line is that the National Weather Service's mission of protecting lives and property comes first. If VCP 12 or 212 is the best option for your convective event or other hazardous weather event, use it and don't look back. If parts break, they will get fixed. So, use the faster VCPs when you need them, but DO be mindful to switch to slower VCPs when the event is over. Even the most chaotic weather events have a relatively short life, so use the faster VCPs for the time period, then downgrade to 21 or even clear air if the precip is gone. A combination of thoughtful VCP usage and pro-active maintenance can protect both the pedestal components and the public!

VCP 31 Considerations

- It is the **only** long pulse VCP
  - Improves sensitivity for very light precipitation, such as freezing drizzle and snow flurries



What is the operational advantage of VCP 31? Long pulse provides greater sensitivity, which means seeing returns from very light precipitation that you would not see with short pulse. In this freezing drizzle event, VCP 31 detects this very light, but high impact type of precipitation much more effectively than 32. Neither VCP will be producing precip accumulation products, so keep that in mind, but if you want to basically see the most plausible areas of very light falling precipitation, you won't do better than VCP 31. If you're getting reports of flurries and the radar is in 32 or even one of the precip modes, switch over to 31 and see if scope begins to fill up much more with echoes.

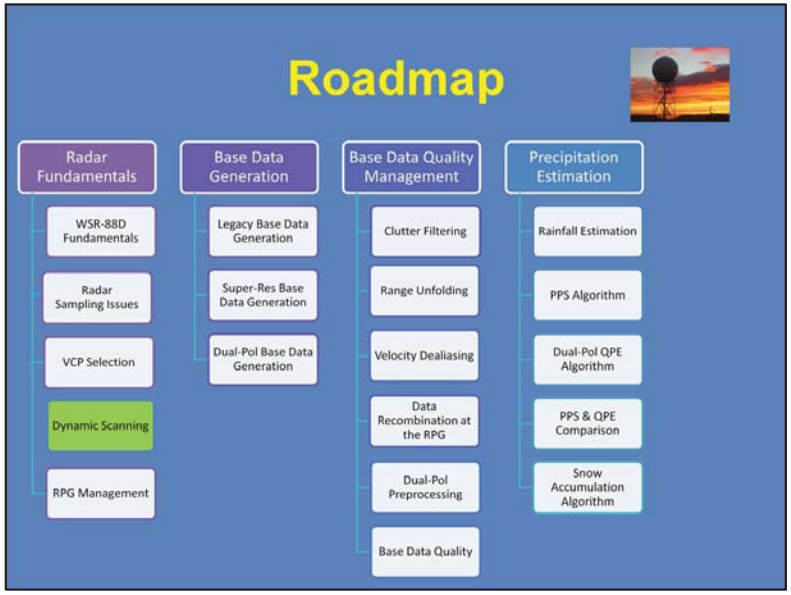
The End
VCP Selection

- VCP Selection
  - Get to Know Your VCPs
    - Old friends and new friends
- Jami.B.Boettcher@noaa.gov
  - 405-325-2986
- Gregory.M.Schoor@noaa.gov
  - 405-325-3004

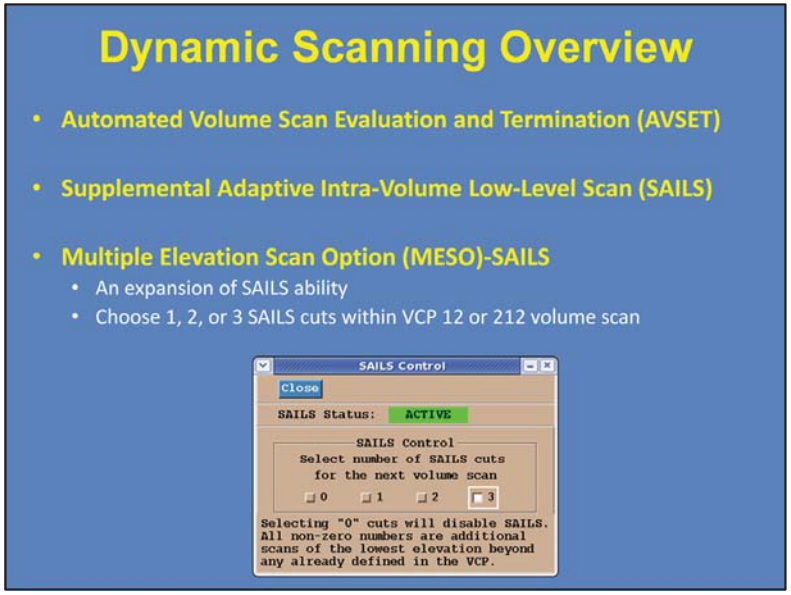
This concludes our lesson on VCP selection training, if you have any questions – here are the resources.



Welcome to this lesson on Dynamic Scanning. What is "dynamic scanning", you may ask? This is a term we will utilize to describe how the WSR-88D radars have, in recent years, attained a capacity for more robust and user-defined ways to scan the skies. In this lesson, we will take a closer look at the applications of AVSET, SAILS, and MESO-SAILS.

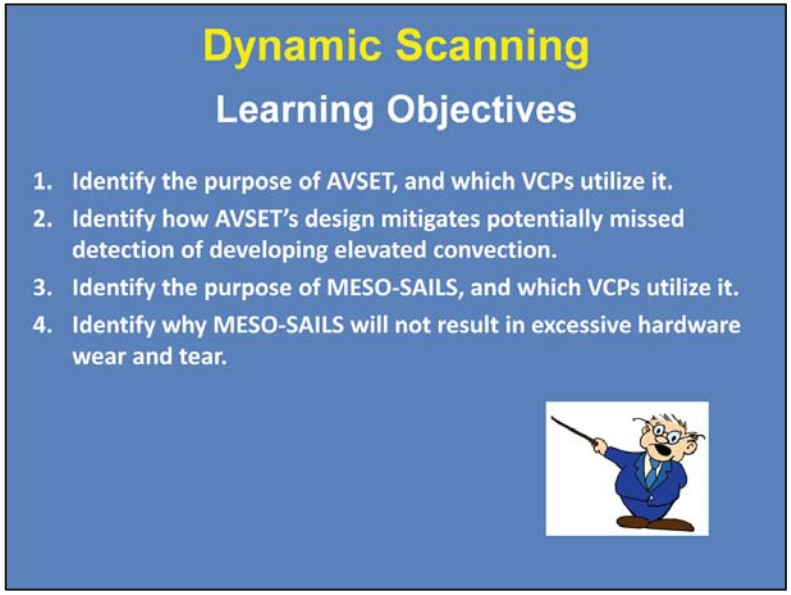


Here is the "roadmap" with your current location.

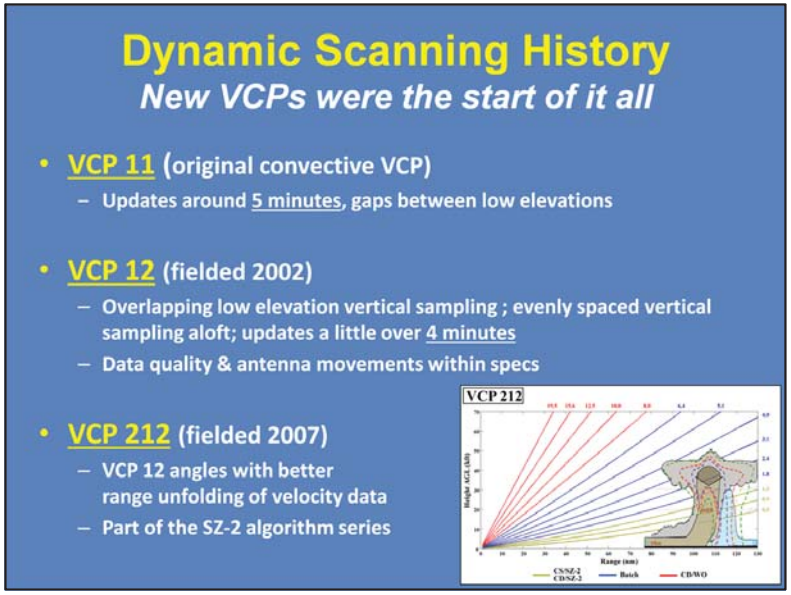


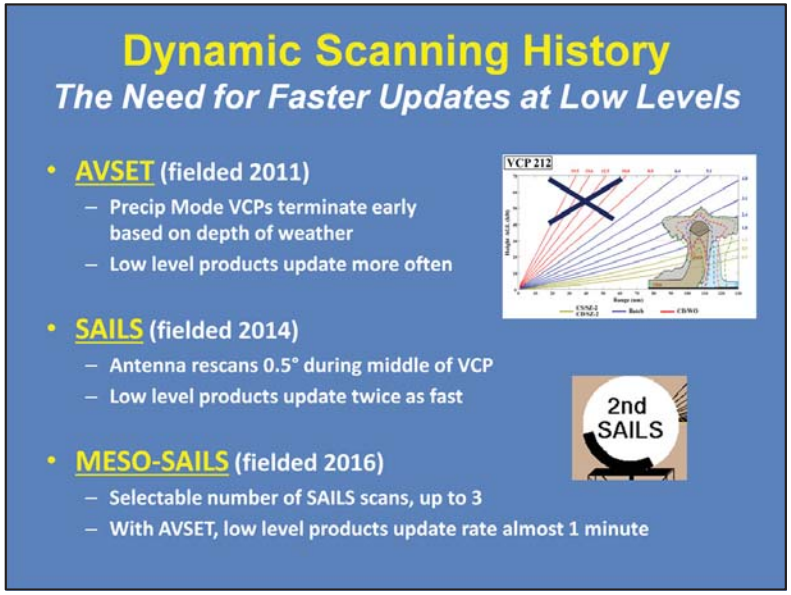
Before we define the lesson objectives, we will spell out the components of this Dynamic Scanning ability of the WSR-88Ds. The first application is called the Automated Volume Scan Evaluation and Termination, or AVSET. This allows any of the Precip mode VCPs to truncate the upper level tilts of a particular volume, if there is no detected echoes above a certain level. The fewer unnecessary tilts, the faster the VCP will be completed and the faster you'll get more data in the lower tilts.

Next is the SAILS application, which stands for Supplemental Adaptive Intra-Volume Low-Level Scan. For precip VCPs, this feature adds a 0.5° degree tilt to every volume right in the middle of the volume, giving you twice as many lowest degree products per volume. Then, taking that idea one, two, and even three steps further, MESO-SAILS to allow one, two, or three additions 0.5° cuts in one volume. With three SAILS cuts, you get almost one-minute updates on the 0.5° products.



Here are the learning objectives for this lesson, please read through them and keep them in mind as you go through the lesson.





Let's start off with a short background of how we even got to this era of Dynamic Scanning. Some of the VCPs currently in use did not exist from the beginning. They were developed over time and the Radar Operations Center folks looked for more effective ways to use the radars. VCP 11 was the original convective VCP which generated volumes in about 5 minutes and had a fairly uniform scan strategy from top to bottom. However, in real time, people were noticing that this created gap in the low elevations, which are usually the tilts you would use for much of your storm interrogation. So, back in the early 2000s, VCP 12 was created, which helped to fill in the gaps in the low elevations, the antenna was sped up a bit to account for the additional elevation angles with a little over 4-minute full volumes. The data quality and antenna movements for VCP 12 are within the NEXRAD specifications. In the mid to late 2000s, the SZ-2 algorithm was created, to better mitigate range folding with velocity data. VCP 212 is the result, with VCP 12 angles and better velocity data collection.

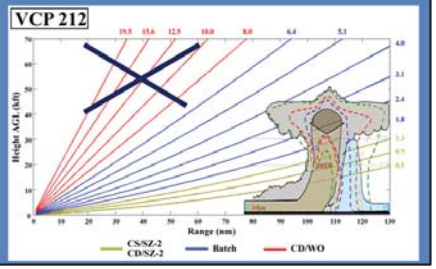
Even still, with all these enhancements, the fastest we could get new radar products was to 4 minutes.

A few years ago, knowing that even faster rotation rates would mean data quality below NEXRAD specifications, the attention turned to getting the most of each individual volume scan. First came AVSET, which began with the question – "why do we need to scan clear air at the higher tilts, above any sampled precipitation?" The response was to terminate the volume for situations like this, where the radar would be scanning above any sampled echoes, resulting in faster product updates by not scanning these upper tilts. AVSET brought the update rate to every 3 and a half to 4 minutes at best. Terminal Doppler Weather Radars, or TDWRs, have a capacity for one-minute lowest degree tilts with their scan strategies. So the next thought was, "why can't we meet that in the middle and have at least one additional lowest degree tilt in a volume?" So, SAILS was created to foster this ability. When SAILS was deemed a success, then MESO-SAILS was created and recently deployed, to allow an addition one, two, or three lowest degree tilts within one volume.

AVSET in a Nutshell



- Vision: Faster updates for Precip Mode VCPs
- Sampling above the weather?
  - Stop current volume scan & start a new one!
  - Only works for Precip VCPs (not VCP 31 or 32)



As we learned in the history of these applications, there has long been a desire for faster VCP updates. AVSET can meet this need when the weather return is limited to the lower elevation angles, though it is dependent on the depth of any detected precipitation and its range from the radar. This image illustrates how a convective storm is at a long enough range that the middle and upper elevations of the VCP are sampling nothing of significance. AVSET only runs on the Precip Mode VCPs, and can truncate different numbers of the upper elevations each volume scan because it dynamically works in each individual volume – which makes for maximum efficiency of the application.

How AVSET Works

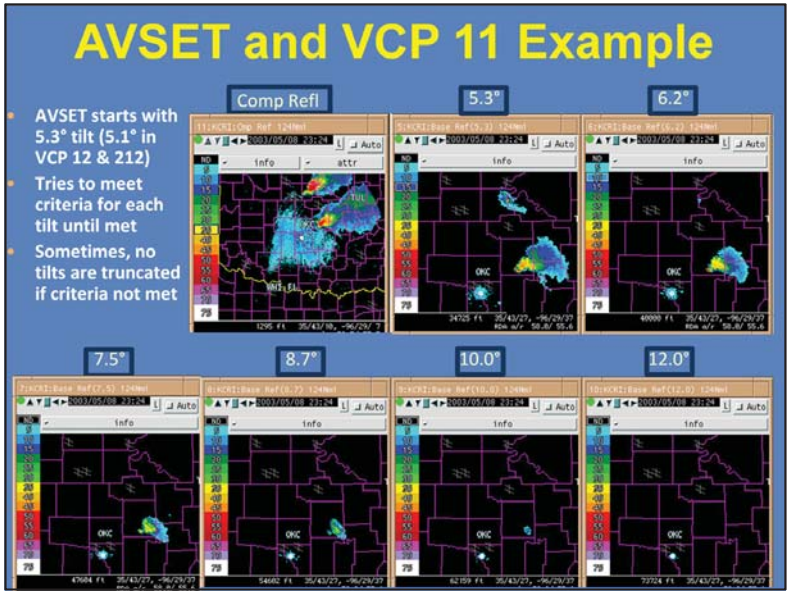


- Terminates volume scan once returns fall below thresholds of dBZ & areal coverage
- Only analyzes data *above* 5.1°
  - To determine when terminate remainder of the volume
- To terminate, each condition *must be met*:
  1. ≥18 dBZ over < 80 km²
  2. ≥ 30 dBZ over < 30 km²
  3. areal coverage ≥18 dBZ has *not* increased by 12 km² or more since the last volume scan
- AVSET terminates VCP after scanning one angle above scan that meets thresholds

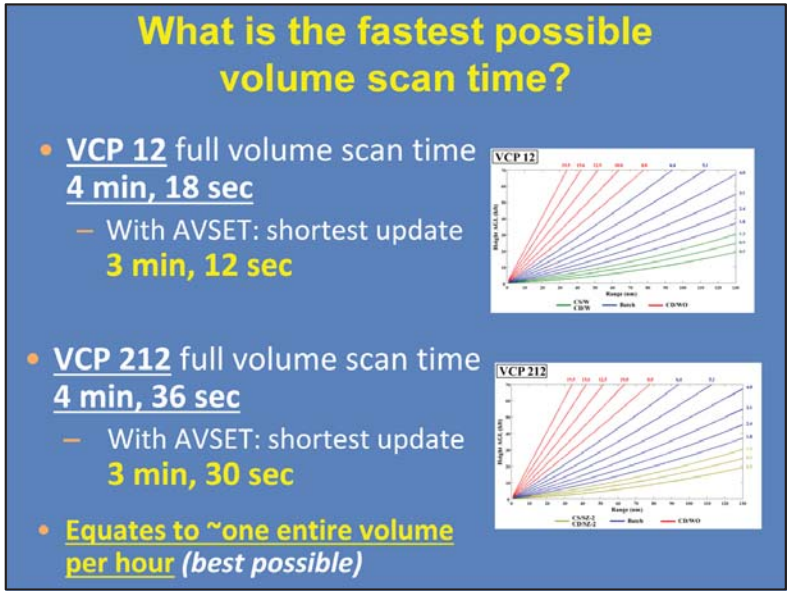
How does this process work? Once the radar gets above 5.1°, it starts to calculate whether the sampling of precipitation meets certain dBZ and areal coverage thresholds. AVSET checks the ***total*** areal coverage of returns above both 18 and 30 dBZ. To terminate the volume scan, there are three conditions and all three of them must be met.

The coverage at or above 18 dBZ must be less than 80 km².
The coverage at or above 30 dBZ must be less than 30 km².
The coverage at or above 18 dBZ has not increased by 12 km² or more since the last volume scan.

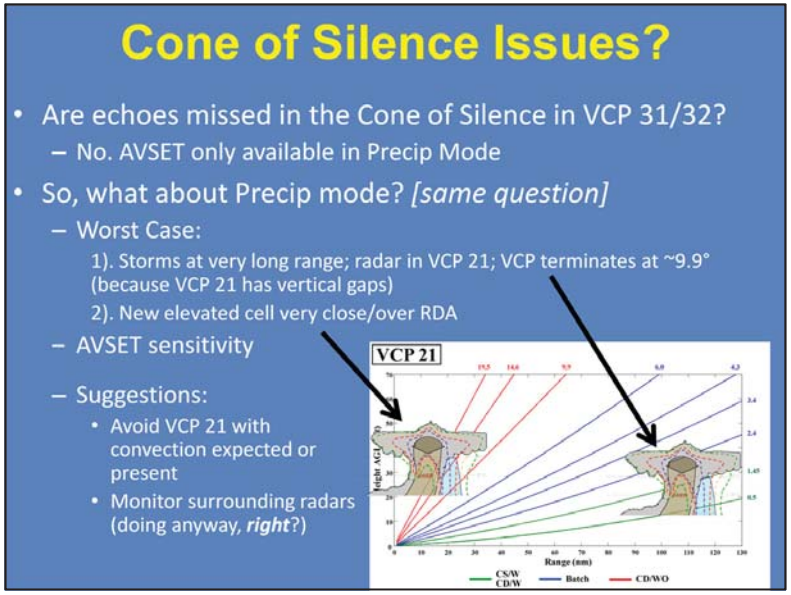
If these three conditions exist, AVSET then truncates the rest of the volume after sampling the next higher elevation. For the angle where returns fall below the threshold, AVSET samples one angle higher, then terminates.



Here is a quick example of how AVSET works with a couple of storms on the display. You can see in the Composite Reflectivity, the first image on the top-left, there are a couple of storms to the north and east of the radar. Starting with the 5.3° tilt, which is 5.1° degrees in VCPs 12 and 212, AVSET will calculate the area of returns above 18 and 30 dBZ throughout that scan. As we go up in elevation, you see the returns decreasing and if the thresholds from previous slide is met, AVSET will scan one more elevation above, then terminate the VCP. In this example, 10.0° is where AVSET is able to meet its criteria, it adds the 12.0° tilt, which is the final elevation for that volume scan. Thanks to AVSET, 3 unnecessary elevation angles were avoided and the volume time was shortened about 45 seconds.



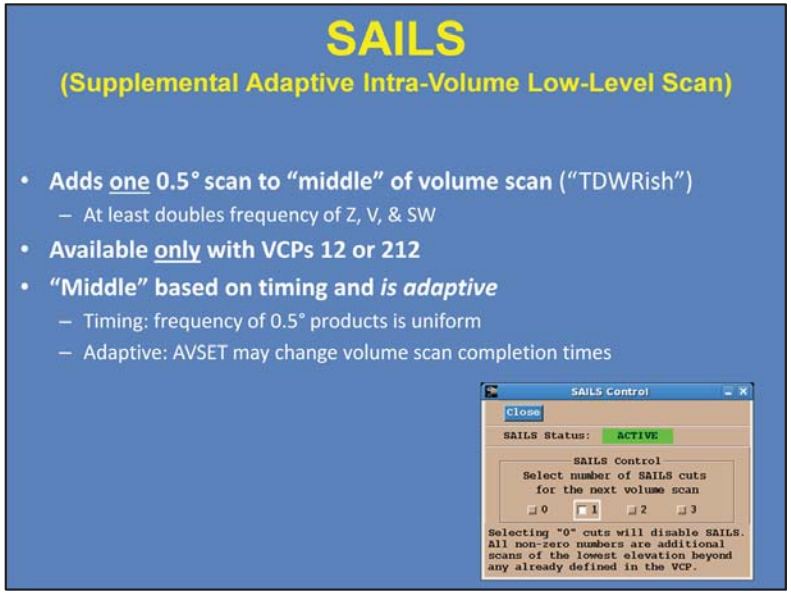
So, after all this talk of speeding up VCPs, what ***IS*** the fastest possible volume scan time with AVSET enabled? Here, we use the two fastest VCPs, 12 and 212. With AVSET disabled and performing a full volume scan, the times are 4 and 4.5 minutes, respectively. But with AVSET enabled and able to truncate all elevation tilts above 6.4°, the volume scan time drops by nearly one minute for each VCP. This means, over the course of one hour, you would receive one extra volume that you wouldn't have without AVSET.



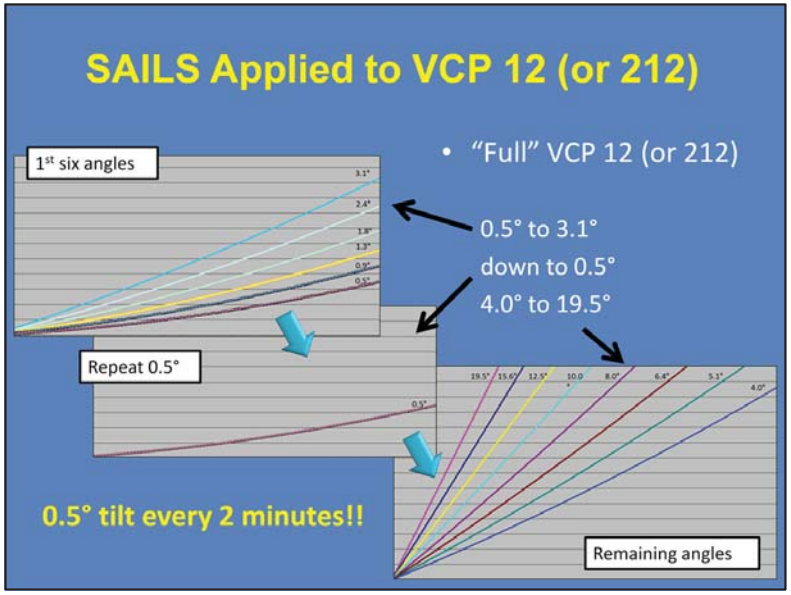
One more consideration for AVSET is the Cone of Silence. Since AVSET starts at 5.3°, it has no affect on VCPs 31 and 32. There is no change in detection of elevated convection in Clear Air Mode.

So, what about the Precip Mode? Could AVSET miss any cells aloft? The first possibility is with echoes that don't have much areal extent and are very far from the radar and you're in VCP 21, which has sizeable vertical scanning gaps at long ranges. The other is if there is a cell developing aloft very near or right over the radar, and VCP 21 again contributes to potentially missing the cell due to gaps aloft. Also, AVSET is very sensitive to developing returns (18 dBZ cannot increase more than 12 km^2).

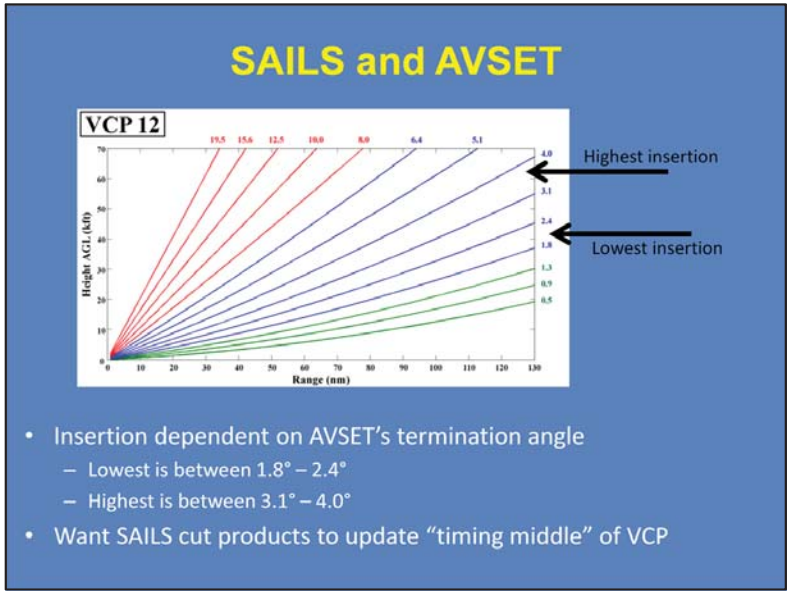
The suggestions for this are to avoid using VCP 21 when you expect convection or there is ongoing convection. It's also important to monitor surrounding radars (a practice you're already doing, right?). This includes the TDWRs, if you have them, to get a more complete regional picture.



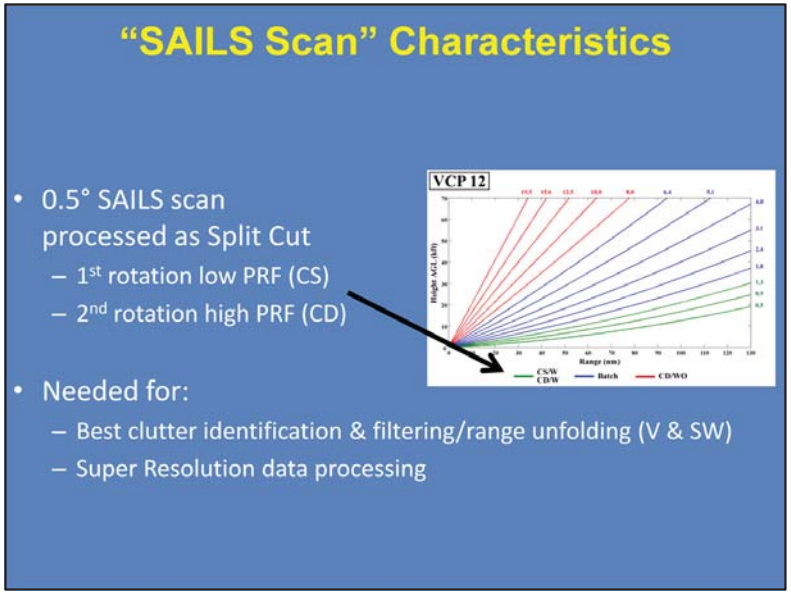
Next in our utility belt of Dynamic Scanning is SAILS, which stands for – Supplemental Adaptive Intra-Volume Low-Level Scan. The lowest degree tilt is by far the most widely used tilt, especially when scanning for low-level features that contribute to tornadogenesis. TDWRs are well known for their one-minute lowest tilt updates and SAILS was the first iteration of the WSR-88D to move toward that capability. SAILS is only an option for VCPs 12 or VCP 212 and it adds an extra 0.5° tilt in the "middle" of the volume scan. This "middle" of the volume scan is based on timing, as you'll see in the next slide. It is also adaptive when AVSET is on because the volume scan completion times will change.



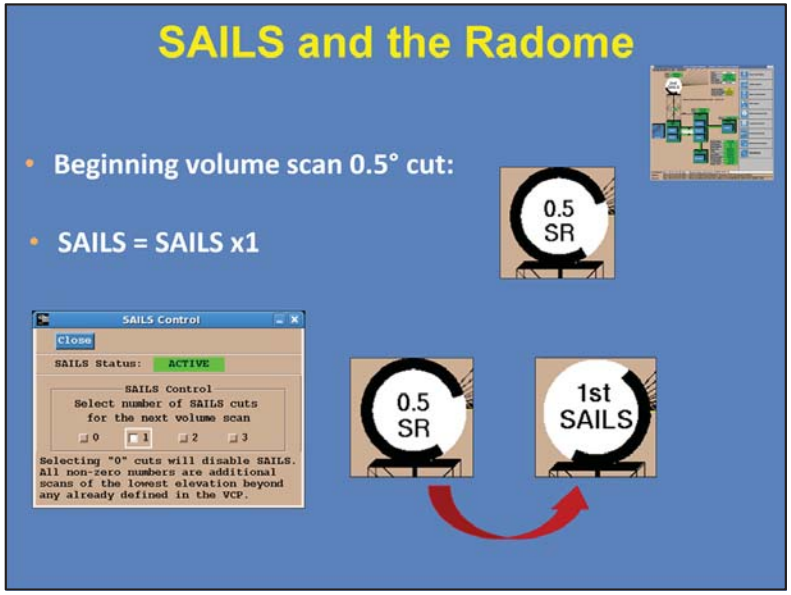
Here's an example of SAILS applied to VCP 12 – or 212, with all the angles sampled, so this would be either with AVSET disabled or there are enough storms that no upper tilts are truncated. The first six angles are samples as normal, elevations 0.5° through 3.1°. The antenna then drops down to sample 0.5°, which is the SAILS cut, then returns to complete the volume scan, elevations 4.0° through 19.5°. The total volume scan time for this example would be around 4 and a half minutes, so the additional SAILS cut results in low level products updating in just over 2 minutes.



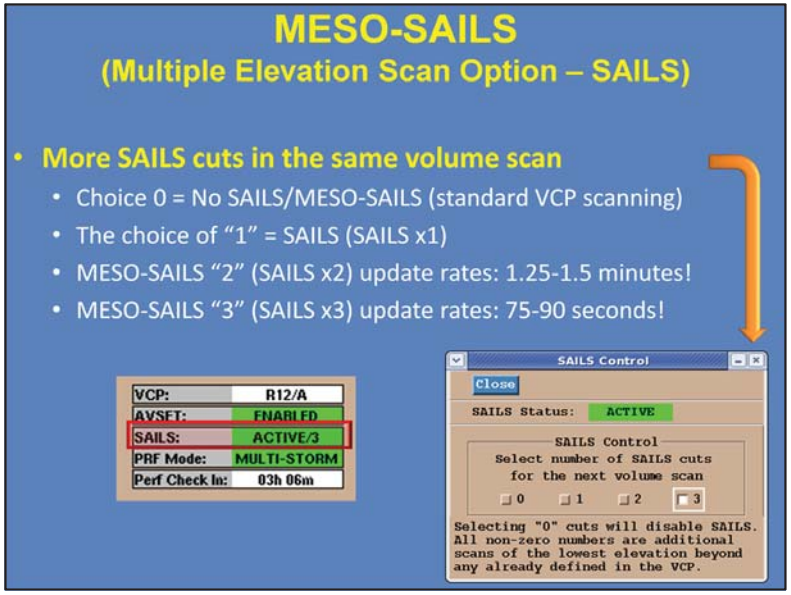
Speaking of AVSET… When it is active and the termination angle changes from volume to volume, the timing of the extra SAILS cut, or the insertion point, will also change. The lowest insertion is between 1.8° and 2.4°, while the highest is between 3.1° and 4.0°. This variation meets the goal of having the SAILS cut product in the middle of the volume scan with respect to timing.

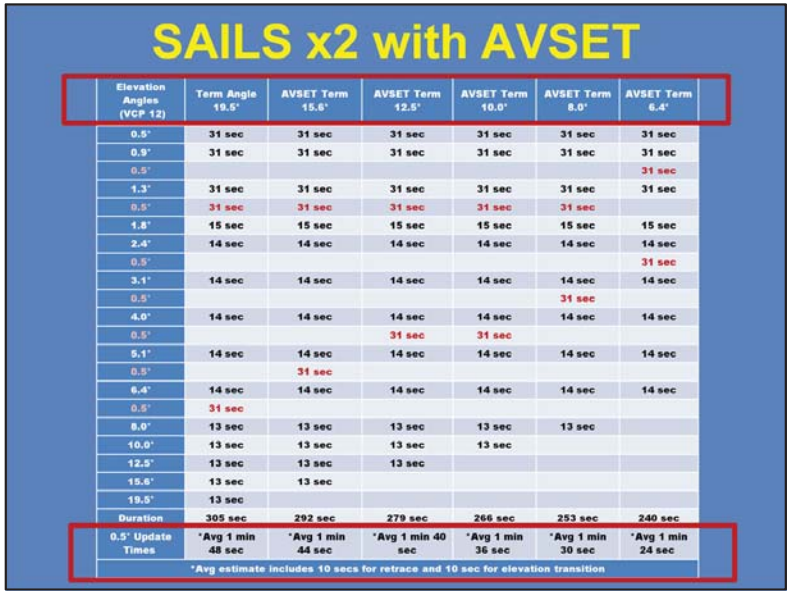


This extra 0.5° scan is known as the SAILS cut. It is processed as a Split Cut, which means the first rotation is at low PRF (Contiguous Surveillance), followed by a second rotation at high PRF (Contiguous Doppler). The Split Cut sampling is needed for the best data quality, better clutter identification and filtering, super resolution data processing, and better range unfolding of velocity and spectrum width.



With both SAILS and AVSET active, there is a need to know when the SAILS cut is executed. To support this, there are two features on the radome on the RPG Control/Status Window. At the beginning of the volume scan, the radome displays 0.5° and SR. In this example, SAILS x1 has been selected, which is the same as what started as just "SAILS". When the SAILS cut is executed, "1st SAILS" appears on the radome.





Elevation Angles (VCP 12)	Term Angle 19.5°	AVSET Term 15.6°	AVSET Term 12.5°	AVSET Term 10.0°	AVSET Term 8.0°	AVSET Term 6.4°
0.5°	31 sec	31 sec	31 sec	31 sec	31 sec	31 sec
0.9°	31 sec	31 sec	31 sec	31 sec	31 sec	31 sec
0.5°						31 sec
1.3°	31 sec	31 sec	31 sec	31 sec	31 sec	31 sec
0.5°	31 sec	31 sec	31 sec	31 sec	31 sec	
1.8°	15 sec	15 sec	15 sec	15 sec	15 sec	15 sec
2.4°	14 sec	14 sec	14 sec	14 sec	14 sec	14 sec
0.5°						31 sec
3.1°	14 sec	14 sec	14 sec	14 sec	14 sec	14 sec
0.5°					31 sec	
4.0°	14 sec	14 sec	14 sec	14 sec	14 sec	14 sec
0.5°			31 sec	31 sec		
5.1°	14 sec	14 sec	14 sec	14 sec	14 sec	14 sec
0.5°		31 sec				
6.4°	14 sec	14 sec	14 sec	14 sec	14 sec	14 sec
0.5°	31 sec					
8.0°	13 sec	13 sec	13 sec	13 sec	13 sec	
10.0°	13 sec	13 sec	13 sec	13 sec		
12.5°	13 sec	13 sec	13 sec			
15.6°	13 sec	13 sec				
19.5°	13 sec					
Duration	305 sec	292 sec	279 sec	266 sec	253 sec	240 sec
0.5° Update Times	*Avg 1 min 48 sec	*Avg 1 min 44 sec	*Avg 1 min 40 sec	*Avg 1 min 36 sec	*Avg 1 min 30 sec	*Avg 1 min 24 sec
*Avg estimate includes 10 secs for retrace and 10 sec for elevation transition						

While SAILS adds one additional 0.5° cut in the "middle" of VCP 12 or 212, MESO-SAILS takes this two and even three steps further, giving you the option of having an additional 2 or 3 0.5° cuts. From the image in the lower-right, you can see the options for 0, 1, 2, or 3. Zero, means that SAILS and MESO-SAILS are disabled, so no additional 0.5° cuts. Choosing "1" is basically SAILS, so you just get one additional 0.5° cut. So, the choice of either 2 or 3 additional 0.5° cuts is the function of MESO-SAILS. Choosing "2" gives an update rate of around a minute and a half, while MESO-SAILS "3" gives an 0.5° update rate of just over one-minute!

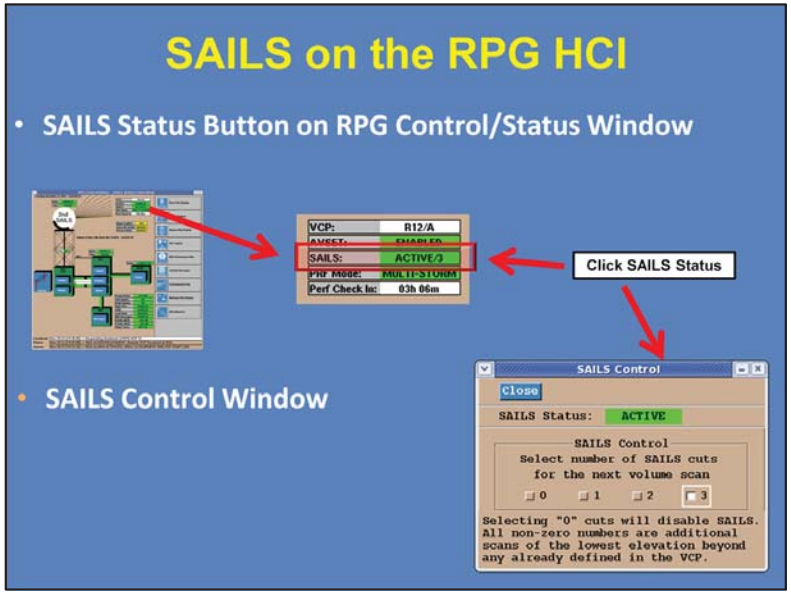
The next couple of slides show tables for the update rate times of all possible AVSET and MESO-SAILS interactions. This table is for SAILS x2. You can see that with AVSET enabled and the termination angles, or "term" in the header, the red values are where the SAILS 0.5° cuts appear in the scanning strategy for each AVSET possibility. Then, at the bottom, you can see the average update times for the 0.5° cuts and above it, the times for the full volume. With AVSET disabled or terminating at 19.5°, the volume will take just over 5 minutes to complete, but you also get two additional 0.5° cuts during that volume.

SAILS x3 with AVSET

Elevation Angles (VCP 12)	Term Angle 19.5	AVSET Term 15.6°	AVSET Term 12.5°	AVSET Term 10.0°	AVSET Term 8.0°	AVSET Term 6.4°
0.5°	31 sec	31 sec	31 sec	31 sec	31 sec	31 sec
0.9°	31 sec	31 sec	31 sec	31 sec	31 sec	31 sec
0.5°			31 sec	31 sec	31 sec	31 sec
1.3°	31 sec	31 sec	31 sec	31 sec	31 sec	31 sec
0.5°	31 sec	31 sec				
1.8°	15 sec	15 sec	15 sec	15 sec	15 sec	15 sec
0.5°			31 sec	31 sec	31 sec	31 sec
2.4°	14 sec	14 sec	14 sec	14 sec	14 sec	14 sec
3.1°	14 sec	14 sec	14 sec	14 sec	14 sec	14 sec
0.5°		31 sec				31 sec
4.0°	14 sec	14 sec	14 sec	14 sec	14 sec	14 sec
0.5°	31 sec				31 sec	
5.1°	14 sec	14 sec	14 sec	14 sec	14 sec	14 sec
0.5°			31 sec	31 sec		
6.4°	14 sec	14 sec	14 sec	14 sec	14 sec	14 sec
0.5°		31 sec				
8.0°	13 sec	13 sec	13 sec	13 sec	13 sec	
0.5°	31 sec					
10.0°	13 sec	13 sec	13 sec	13 sec		
12.5°	13 sec	13 sec	13 sec			
15.6°	13 sec	13 sec				
19.5°	13 sec					
Duration	336 sec	323 sec	310 sec	297 sec	284 sec	271 sec
0.5° Update Times	*Avg 1 min 29 sec	*Avg 1 min 26 sec	*Avg 1 min 23 sec	*Avg 1 min 19 sec	*Avg 1 min 16 sec	*Avg 1 min 13 sec
*Avg estimate includes 10 secs for retrace and 10 secs for elevation transition						

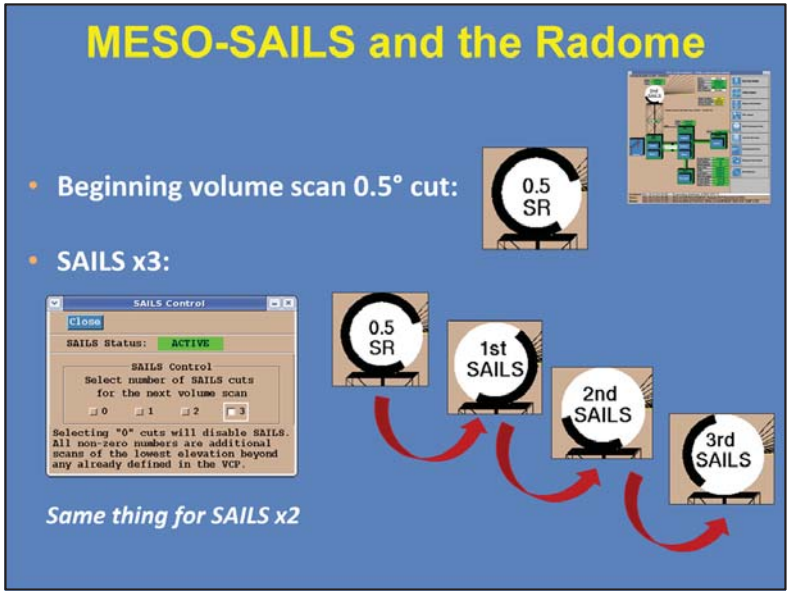
Here's the table for SAILS x3. When might you need nearly one-minute 0.5° products? Are you monitoring low-level circulations near the radar that are more important than monitoring the upper tilts? Possibly and maybe for only a brief period. You may have some QLCS or mesovortex features appearing that would be missed completely if you had standard VCP scanning with 4 to 5 minute updates. But, with SAILS x3 enabled, you wait longer for the upper tilts, sometimes almost 6 minutes if AVSET is not terminating early.

The key takeaway is: **there is always a trade off**. The MESO-SAILS benefit is not "set it and forget it". You may have to change strategies multiple times during an event, just depends on what is happening.

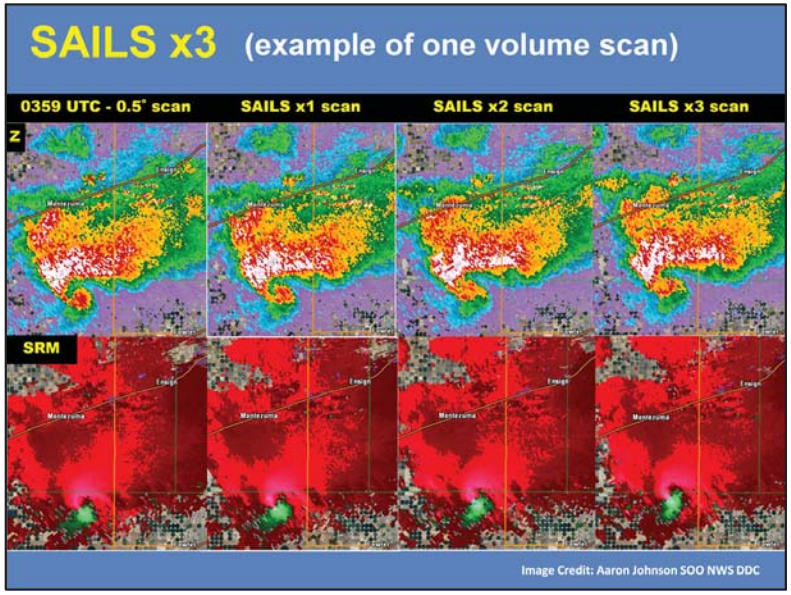


How do you invoke the number of SAILS cuts desired? Starting on the RPG Control/Status Window, the SAILS Status button tells you the current setting. In this example, it is SAILS x3, and since we're in VCP 12, SAILS is ACTIVE.

Click on the SAILS Status button and the SAILS Control Window appears. This window allows you to select the number of SAILS cuts you desire.

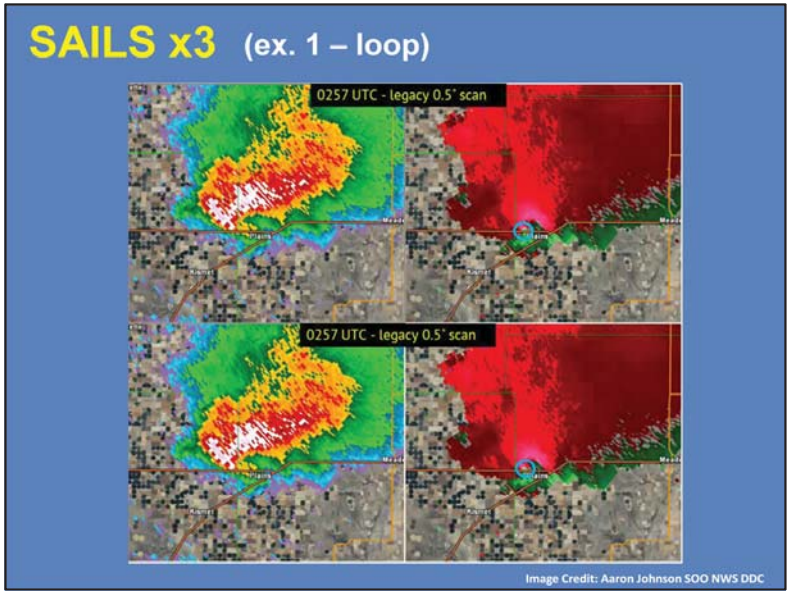


As with the SAILS x1 example, the radome tells you when the SAILS cuts are scanned, identified as 1st, 2nd, or 3rd, depending on the number of cuts selected. Here we have an example of each of the SAILS cuts on the radome when SAILS x3 has been selected.

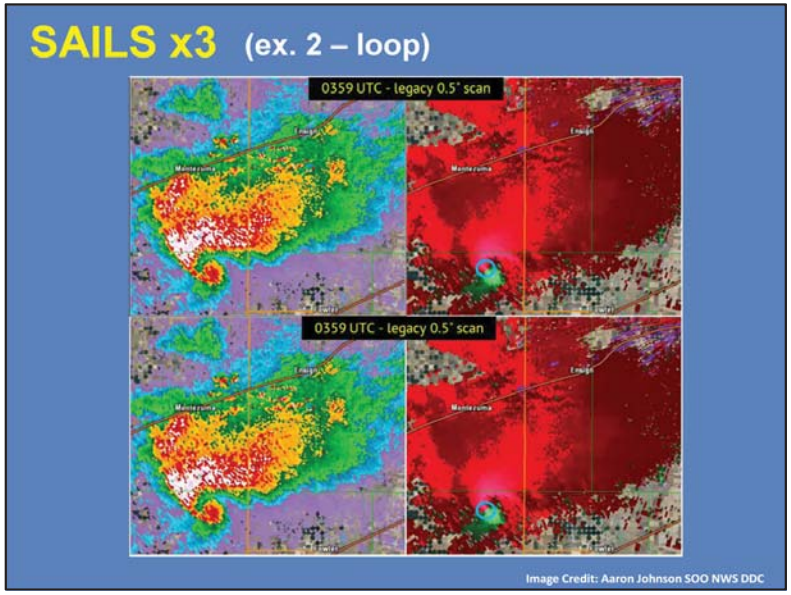


It is fairly intuitive what one additional 0.5° set of products would look like, updating at twice the "normal" rate. But what does almost one-minute 0.5° data look like and what could it mean in terms of watching the evolution of features?

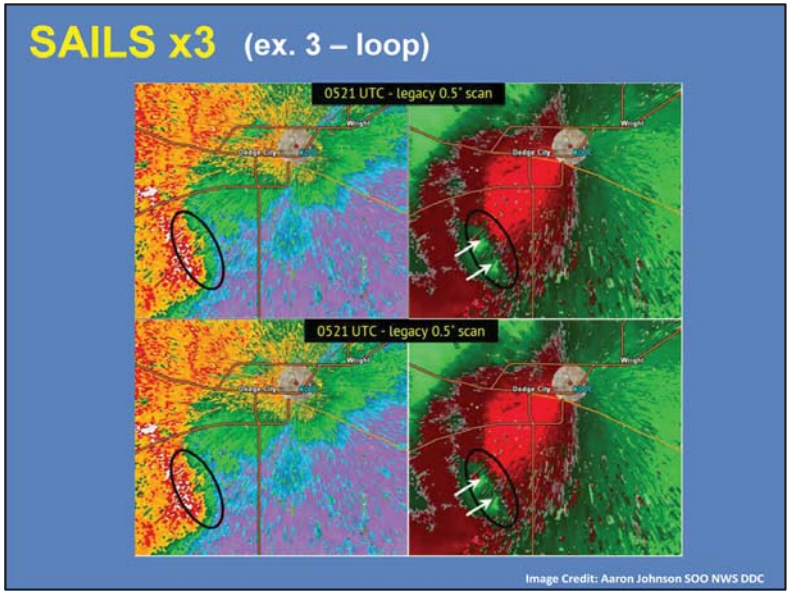
The following examples were provided by Aaron Johnson, SOO at the WFO in Dodge City, KS. This one shows a clear evolution of a tornadic supercell fairly close to the radar – for just one volume scan. We're at 03:59 Z starting on the left, with Base Reflectivity and Storm-Relative Velocity. Between these image on the left and on the right, about 3 to 4 minutes later, the velocity couplet in the cell's hook echo has moved at least a few miles. With no MESO-SAILS enabled, there would be no intermediate images, and you would have to wait 4 to 5 minutes to see where this strong tornadic signature has moved.



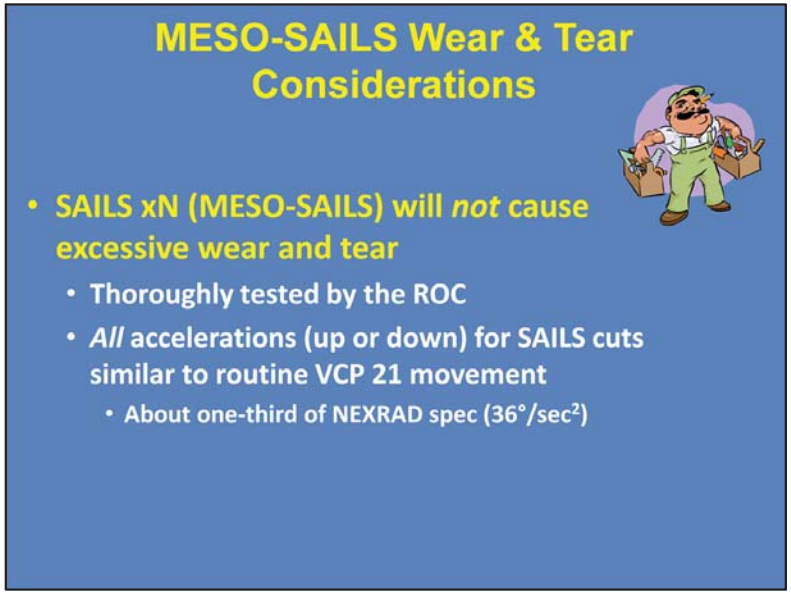
This is the first of 3 examples that are all short gif animations, that will keep replaying from start to finish. SAILS x3 is enabled, with Base Reflectivity on the left and SRM on the right. The top half of these loops are the SAILS x3 cuts and the bottom half show just the first 0.5° cut from each consecutive volume – so, standard VCP scanning. Notice that while the bottom half images are basically "standing still", the top half move consistently through scans that are coming in at intervals of just over one minute! Notice how the estimated track of the mesocyclone, added by Aaron, shows a more detailed and refined track then the legacy side which looks like it's just connecting the dots with much lower spatial accuracy.



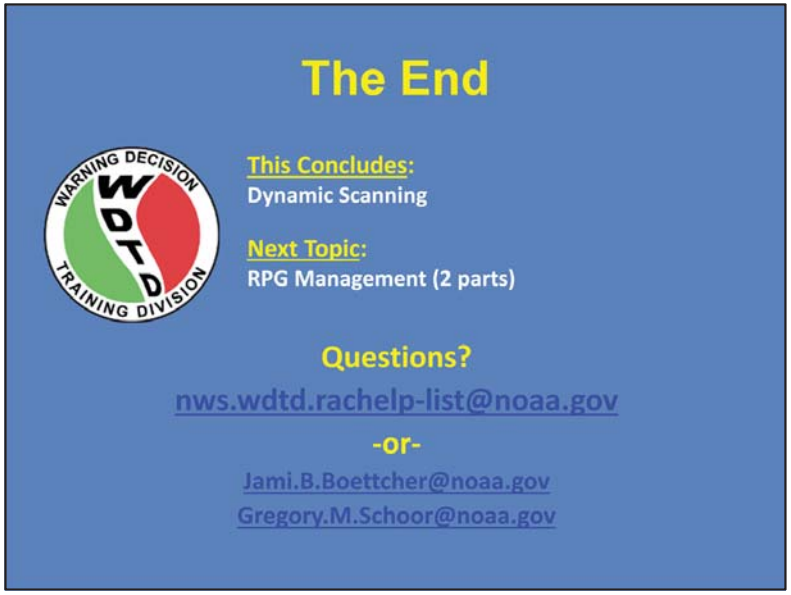
Now for the second example, again showing the tracks of multiple possible mesocyclones. Not only are we able to better track a more exact path of a potential tornado or at least the best area of rotational velocities, but we now have a bit more lead time on the precursors of these types of signals. As you watch, the second couplet doesn't quite develop and the third one does. The second, which may have been a brief satellite tornado, wasn't even detected with the legacy scan timing.



The last example is of an RFD surge that swept just south of Dodge City, KS – just a few miles away from the RDA. There are some brief and embedded couplets with this surge. Notice the difference in the evolution of this feature with SAILS x3 enabled and compare it to the legacy scan timing on the bottom. With a feature like this so close to the radar, you are probably not getting very good data from the upper level tilts anyway. It may be of more benefit to have SAILS x3 running in cases like this, so you can track the low-level progression of such features. If you are more concerned about a core height, you may not have good sampling of the upper tilts of the storm or the core aloft anyway and may have to utilize another nearby radar for that. All of this is a balancing act, where you need to determine, sometimes in real-time, which option is the best to see the features you are looking for.



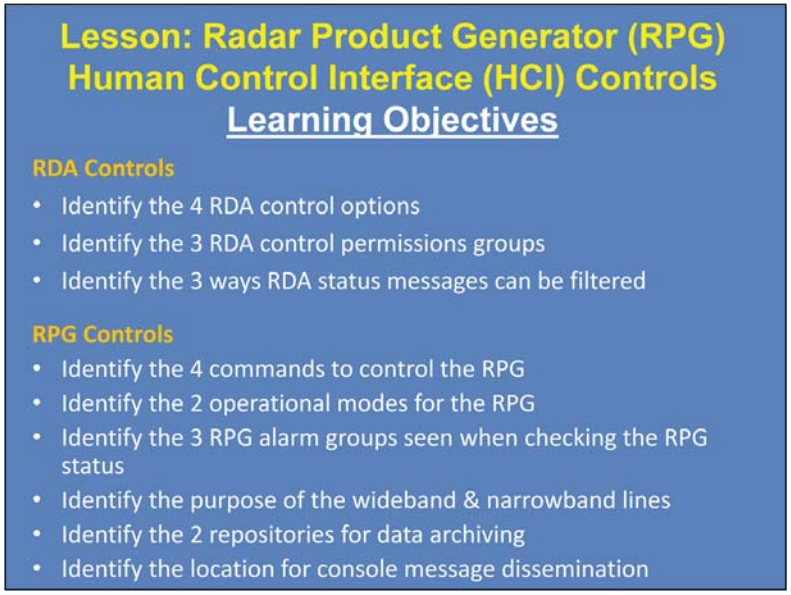
When the faster VCPs, 12 and 212, were fielded, there was initial concern about more wear and tear on the moving hardware of the WSR-88Ds. When SAILS and MESO-SAILS were being tested before deployment, it was understood that no feature was going to be deployed that resulted in excessive wear and tear. The engineers at the Radar Operations Center have evaluated these features and have found no need for concern. SAILS uses the ***same*** azimuthal rotation rates as VCPs 12 and 212 and the antenna acceleration and deceleration rates are comparable to VCP 21, with it's wide gaps between elevations. The antenna motions for both VCP 21 and MESO-SAILS are, at most, ***one third*** of the design specification.



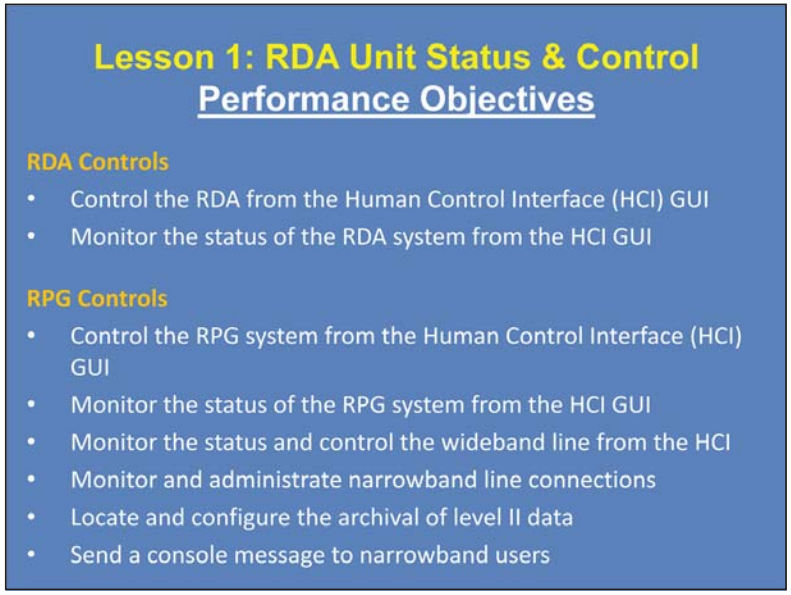
This concludes the lesson on Dynamic Scanning. The next lesson will dive into a 2-part lesson on working with understanding the RPG – or Radar Product Generator. If you have any questions, please direct them to any of the email addresses listed.



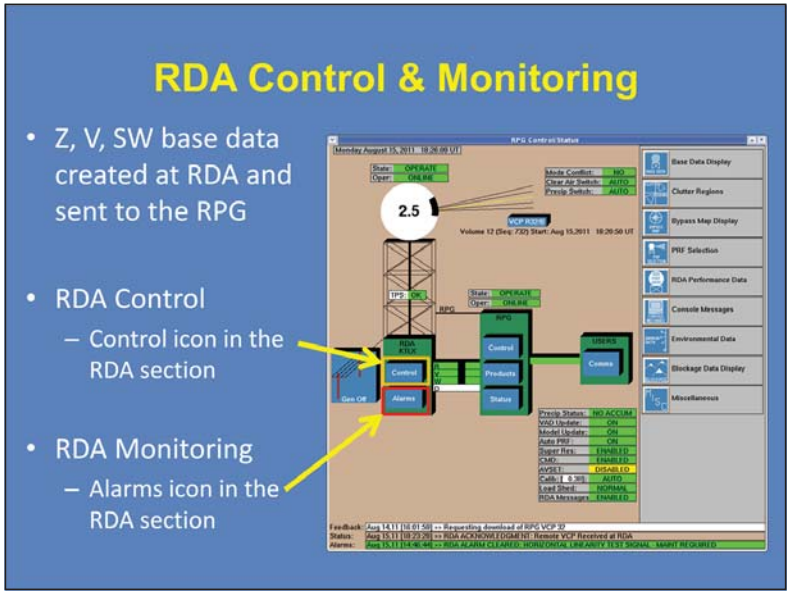
Welcome to Topic 2 and the first of two lessons for RPG Management. RPG stands for Radar Product Generator and is basically what creates the products you see that come from the RDA, the radar itself. This lesson will focus on the interface itself for the RPG, showing you the buttons and what they mean. The next lesson will be about the operational functionality of these buttons and features. So, let's get started!



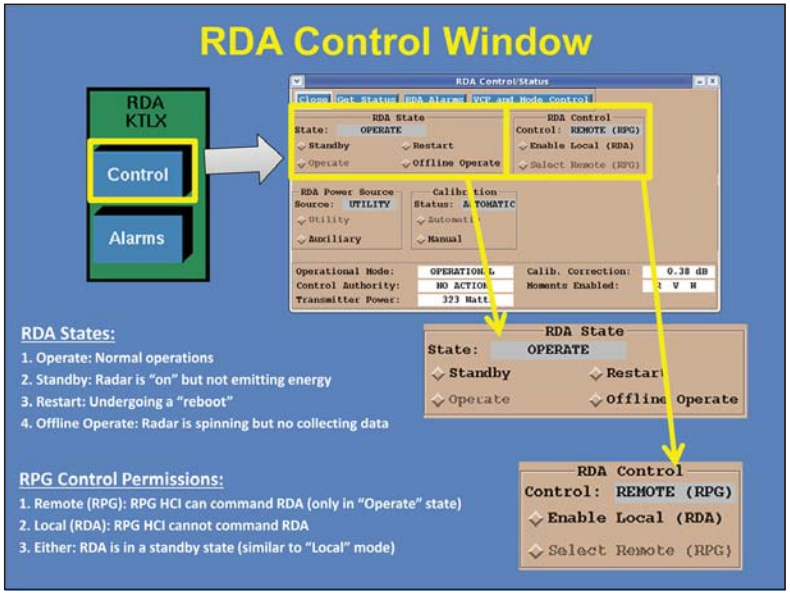
With so many buttons and features we will go over in this lesson, here are the objectives which are all essentially the identification of the various options with each button and feature. So, please read through these and keep them in mind as you go through the lesson. The next slide is the performance objectives which is a more hands-on version of these objectives.



As mentioned on the previous slide, these are the performance objectives for this lesson, please read through them and keep them in mind as you go through the lesson.

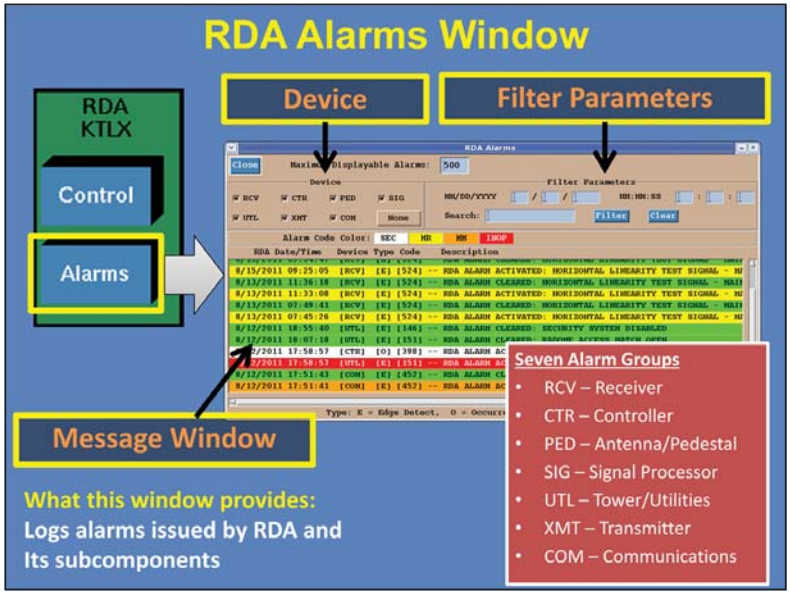


So this window you see is the RPG HCI. HCI stands for Human Control Interface. This is your control panel for the entire radar, if you will, the one module you will use to view, modify, and monitor settings of the radar itself and the products it generates. The RPG, or Radar Product Generator, basically converts base data from the RDA into other products that are derived and displayed in AWIPS and other display modules. This lesson is all about the three darker green modules in the middle of the image you see, which is an example of an RPG HCI window. These three modules hold clickable buttons, the blue buttons, where you'll find a number of other important settings. The next few slides will talk about how the RDA can be controlled and monitored.

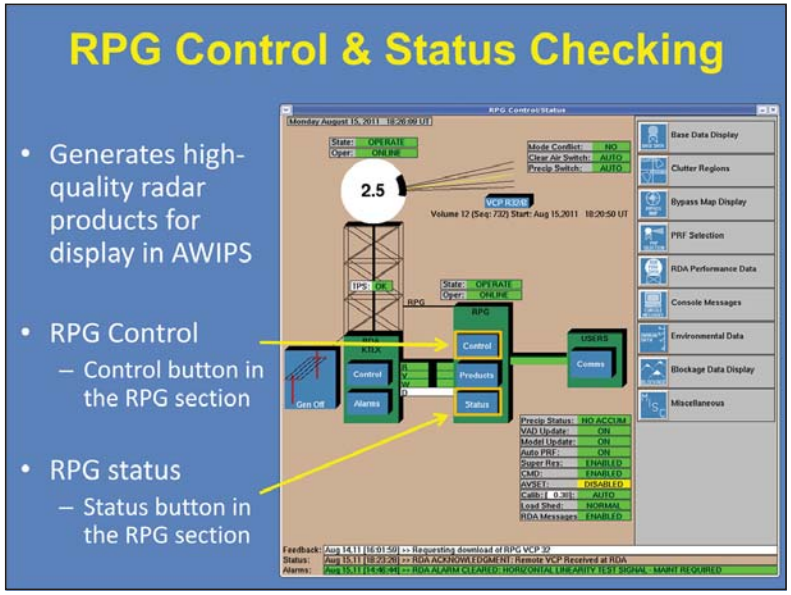


The RDA control window is just for that purpose – to control the functions of the RDA. First of all, if the RDA is running routinely, the state is called "Operate" which you'll see in the status window on the main GUI as well. When you click on the 'Control' icon in the RDA as shown here, the window on the right will appear, showing the current state of the RDA. This is also where you would change the operational state of the RDA if necessary. Since the RPG is also a part of this equation, the RDA control portion of the window allows the user to specify whether or not the RPG can send commands to the RDA.

Please take a moment to look through the different states and aspects of these two sections of the window, down on the left hand side and continue when you're ready.

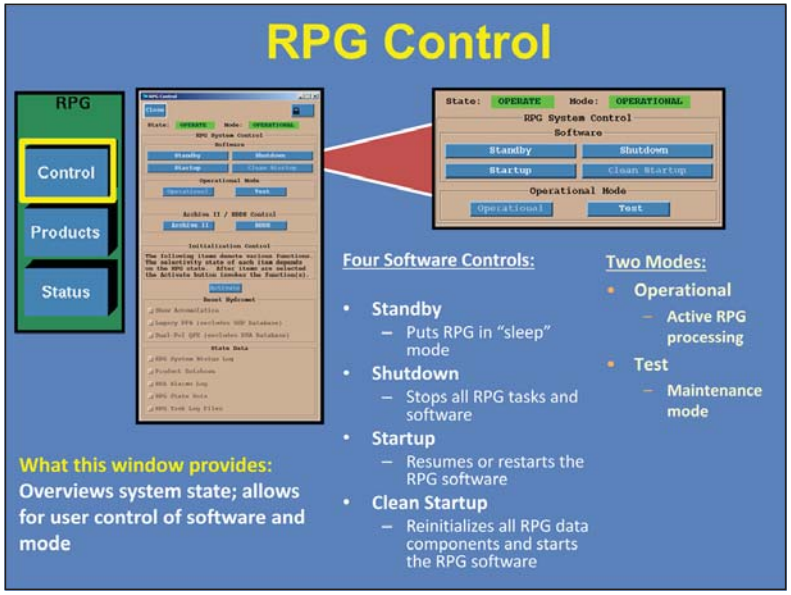


You can monitor the RDA, click on the 'Alarms' button. From this window, you can view and filter all current and previous alarms issued by the RDA. Messages are displayed as a scroll list in the message window. The Device Box allows you to sort and isolate messages by the RDA sub-component that triggered the alarm. There are seven alarm sub-groups from the RDA, which are spelled out in the red box on the lower right for reference. Unchecking a box will only hide the alarms from the message window not fix the alarm. As a result, you will have to be careful with keeping these boxes unchecked from long periods of time. In general, a vast majority of the yellow and green alarms can be ignored. Most of them as just status messages as opposed to actual alarms. There are so many different components and algorithms that go into a radar's operating system, that many of these are just internal checks on the system. However, when you see orange and especially red alarms that don't clear themselves within a least a couple of minutes, then you should alert the El Techs or shift supervisor so that troubleshooting techniques can be administered.

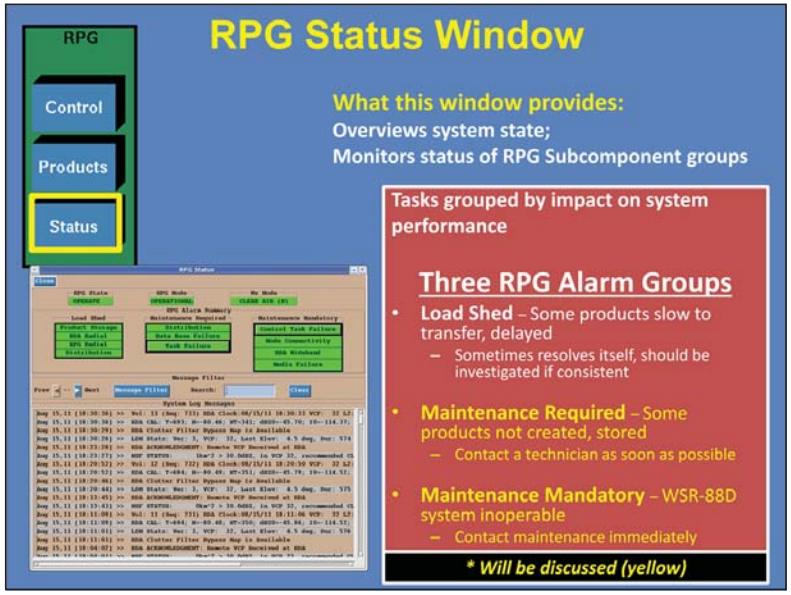


Now that we've covered the 2 big buttons on the RDA, let's go down the line to the RPG, which has 3 big blue buttons. Recall that the RDA is basically the radar itself and the various receiving hardware in the shelters underneath the radar. The RPG is the box that receives this data and converts it into what we know as products that can be displayed either in AWIPS or other platforms.

The RPG is controlled through the Control button you see here, and RPG status monitoring is accessed through the Status button. We will address each of these components on the slides ahead. The products button is geared toward the actual products generated, which you won't need to monitor or modify, so we will skip that for this lesson.

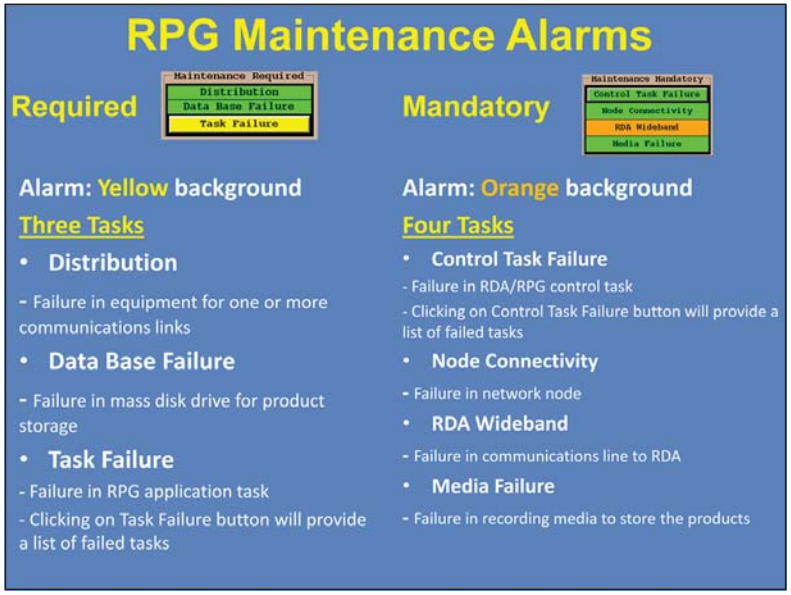


There are 4 software control commands that can be initiated on the RPG software, there in the lower middle – Standby, Shutdown, Startup, and Clean Startup. Read through each of the control functions and know that these are options for times when the radar needs to go into some other mode that is not fully operational. If the radar goes down unexpectedly or for routine maintenance, these are the different functions to react to in those cases, much like a regular PC. The RPG also has two similar modes but slightly different because of the configuration it has compared to the RDA. The RPG's two mode are Operational, which means the RPG is active and processing data and then you and Test it, which can be initiated for a number of reasons such as maintenance at either the RDA or the RPG. With an idea of the different control types available to the RPG, we will now focus our attention on monitoring the status of the RPG.

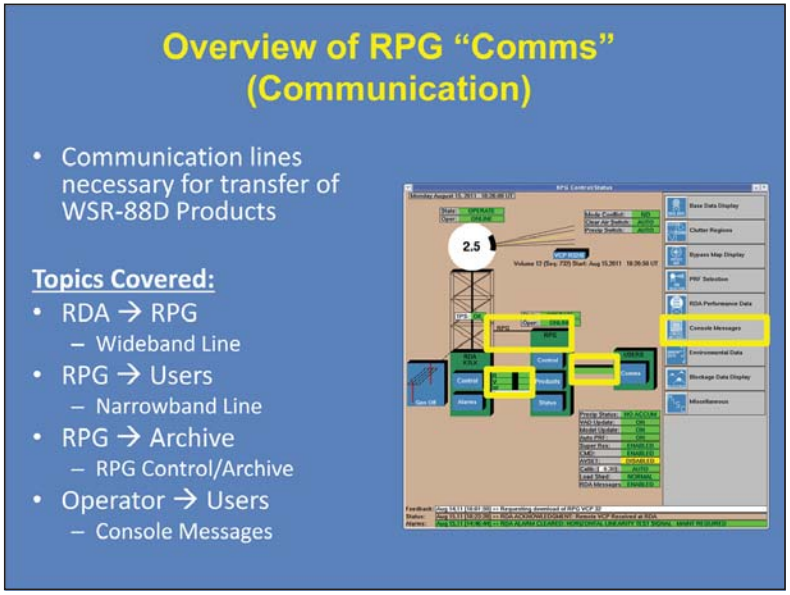


To launch the RPG status window, click on the 'Status' icon in the RPG section of the HCI. This window provides an overview of the RPG system state and current mode, along with the status of several RPG subcomponents. We will elaborate on the two types of maintenance alarms you see in the lower right window, the required and mandatory maintenance types. Load Shed type alarms sometimes appear but are less critical and mostly clear themselves up.

Consistent alarms should be investigated for a potential data bottlenecks.

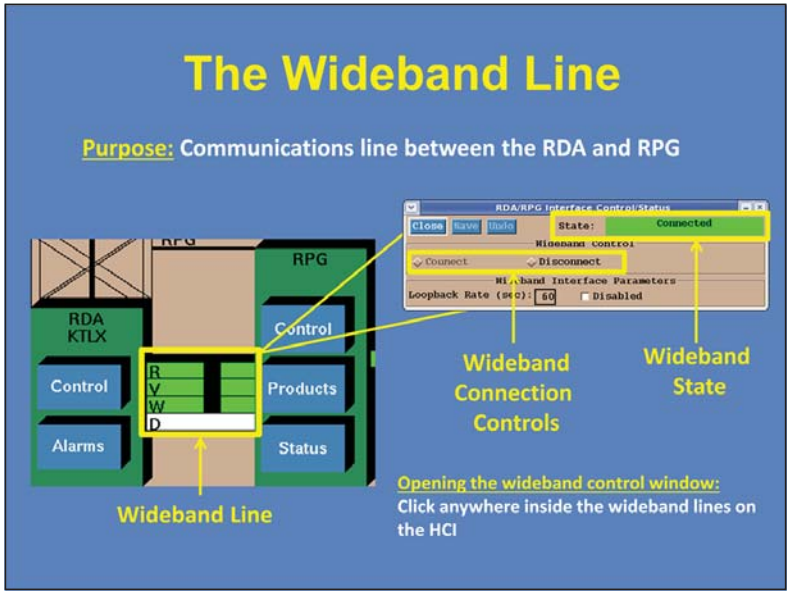


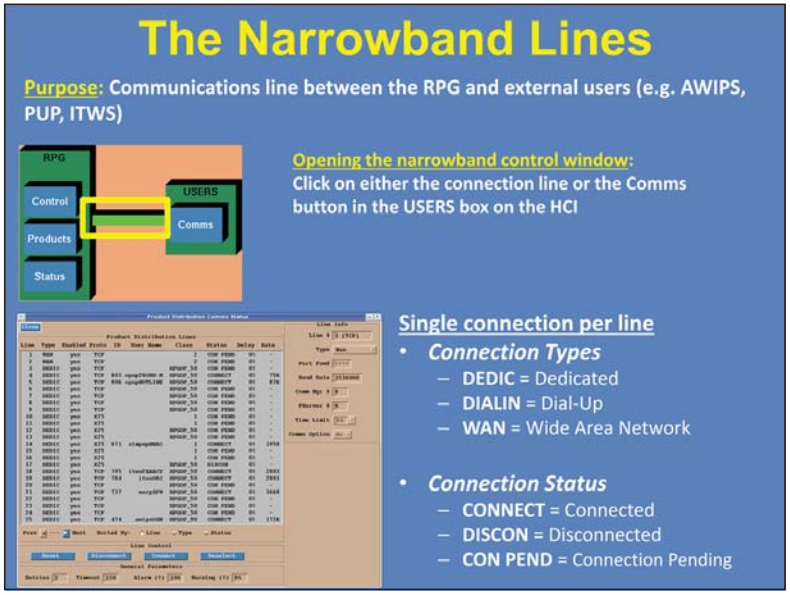
The other two types of alarms are the ones that will require action to be taken, if they don't clear themselves automatically or appear consistently. The RPG's alarms differ from the RDA in that they are more about the software side of things than the hardware and parts of the RDA – or the radar itself. So common issues for the RPG usually have to do with a backlog of products, improper generation of products for various reasons, and connectivity issues. On the left, the "Required" alarms are the slightly less serious but still important alerts, which appear in yellow. Then, on the right, you have the more time-sensitive alarms – which are called mandatory because the RPG cannot do it's job, relative to this alert, and needs immediate attention. Please also take a moment to read through the types of associated mandatory alerts.



In this last section, we will cover system communication. Communication lines are important for the transfer of data throughout the WSR-88D system and ultimately to it's end-location, such as your AWIPS environment. While there are numerous communication procedures employed by the WSR-88D system, we will focus our efforts to 4 communication topics:

1) First, how data are transferred between the RDA and the RPG
2) Second, how data flows between the RPG and the users
3) Third, how data are archived to use by external partners to the NWS
4) And finally, how you the radar operator can relay information and messages with other end-users

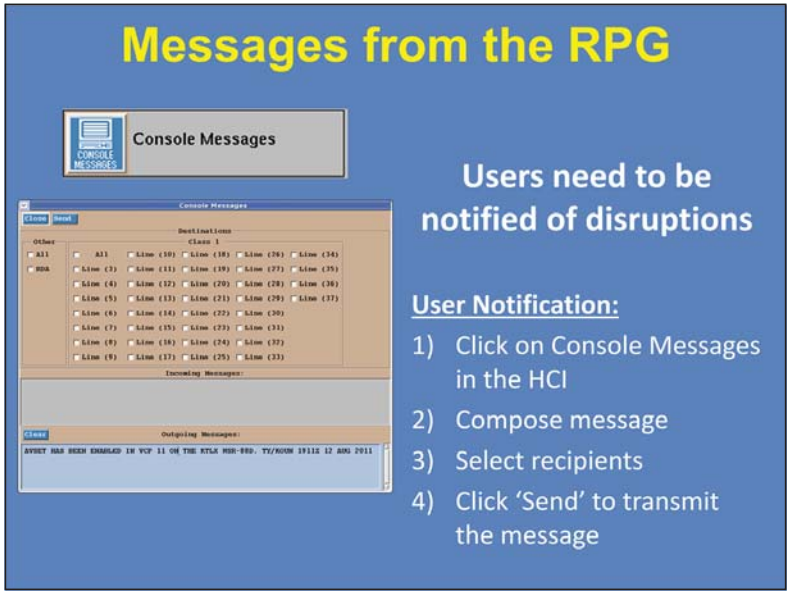




The wideband line is basically the main communications link between the RDA and the RPG. See the 4 lines there in the image on the left? This basically is showing the data stream of base products produced at the RDA and is sent directly through this line to the RPG, so that the RPG can generate the various derived products *from* the base products.

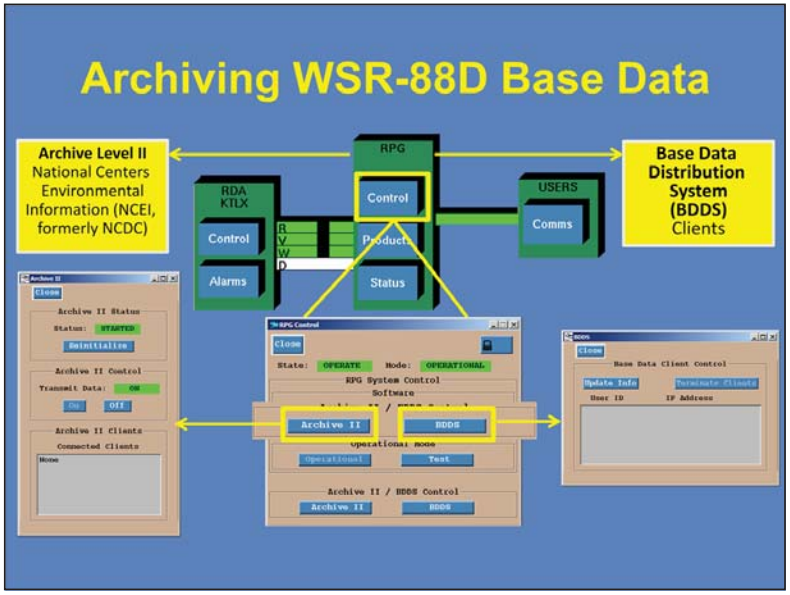
If there is ever an instance where it is necessary to disconnect, the Electronics Technicians or EI Techs should perform this task. But clicking on this zone will pop up the window you see on the right.

The other type of communications lines are the narrowband lines. They serve as the communication medium between the RPG and external display systems, such as your operational AWIPS. Users are typically entities with their own RPG and use it to receive products from this radar through this line. Think of it as a dial-up internet connection to a server, with differences in the technicalities. Your Radar Program Leader should have a list of what's called the "Dedicated Users" or entities that have one of these line feeds. The main thing to monitor is whether one of your users is having problems with their connection, which will either show Disconnect or Connect Pending. Not every entry that has this status has a problem though, most or all of these that you'll see are set up a certain way for specific purposes – but something to keep in mind when doing a routine checklist check.



Messages can be sent from the RPG which, in AWIPS 2 for instance, will show up as an FTM product – or "Free Text Message", which will go out to the users. The main purpose of this feature is to notify end users of something having to do with the radar's operability. If the radar has failed for some reason yet unknown, or if a technician is taking it down for routine maintenance, this is the way to do the quick notification.

To generate a message, click on "Console Messages" on the right side of the HCI. Compose the message on the bottom of the window. All caps is default and the only way it will appear. Make sure that "All" and "RDA" and selected on the left-side – which will send it to everyone. Unclicking them will make it an internal message. Then click "Send". Nothing will happen on this end, but you will see the message if you type FTM and the 3-letter WFO identifier in an AWIPS text window. In fact, that's a good practice to make sure the message sent.



Finally, let's look at how WSR-88D data are archived. In addition to the base data being transferred to the narrowband lines, the base data can also be transferred to the following two places:

1) First, all the level II and level III data goes directly to the National Centers for Environmental Information, NCEI which is the former NCDC. They inventory the base data and make it available for public distribution.
2) Second, Level II base data will also be provided to specific clients through the Base Data Distribution System (BDDS).

Archive status can be viewed and modified by clicking on the RPG control button. Inside, you will find 2 buttons, "Archive II" and "BDDS". Clicking on each of these will reveal their individual configuration windows. Here, you can monitor the status of the archive and manage the sites receiving the data. All of this should already be set up for your radar with no real reason to make modifications.

RDA Control & Monitoring Summary

Four commands for RDA software

- **Operate** – the radar is scanning and processing data
- **Standby** – the radar is on but not emitting energy
- **Restart** – the RDA components will be restarted
- **Offline Operate** – the system is operational but not collecting data

Three control permissions

- **Remote (RPG)** – commands can be sent from the RPG
- **Local (RDA)** – commands can only be sent locally from the RDA
- **Either** – RDA is in a standby state ; similar to Local (RDA)

RDA Status Monitoring

- Alarms are separated by RDA component (7 total alarm groups)
  - receiver, controller, antenna/pedestal, signal processor, tower/utilities, transmitter, communications
- Messages can be filtered by
  - Date, search term, and/or RDA component

Since this lesson was essentially 3 separate topics rolled into one with a lot of options associated with each feature of the RDA and RPG, we'll have 3 separate summary slides for you to act as a reference for each. First off, this is a summary of the RDA control and monitoring window. Please read through them and go to the next slide when you're ready.

RPG Control & Monitoring Summary

Four commands for RPG software

- **Standby** – RPG tasks are in a suspended state
- **Shutdown** – RPG tasks and processes are terminated
- **Startup** – RPG is active from either standby or shutdown state
- **Clean Startup** – RPG restart with database/linear buffer purge

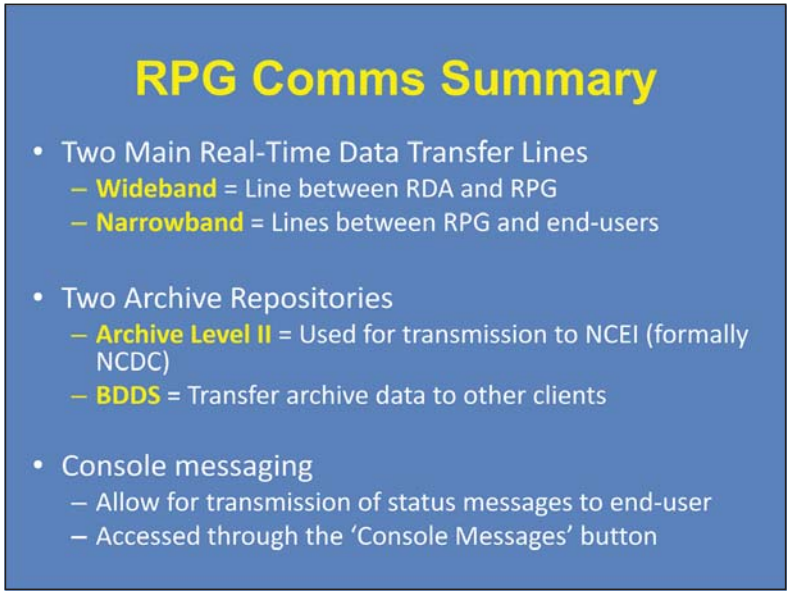
Two RPG System Modes

- **Operational** – the RPG is active and processing data
- **Test** – RDA or RPG is undergoing maintenance or testing

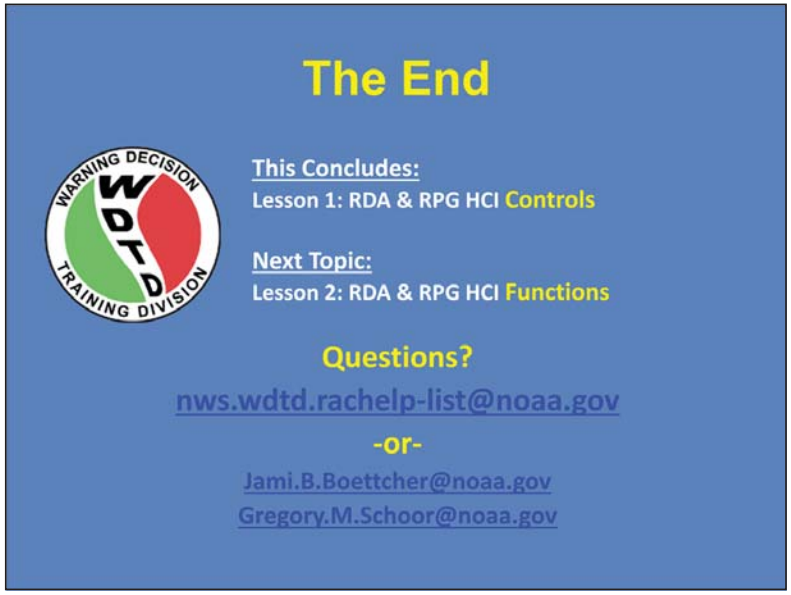
Three classes of RPG alarm groups

- **Load Shed** – Investigate if problems are consistent
- **Maintenance Required** – Contact technician ASAP
- **Maintenance Mandatory** – Contact technician immediately

Next, we jump over to the RPG, which initially looks like it has similar characteristics to the RDA's features, which it does, but the RPG, being mostly a software operation, has some very different status options and associated alerts. Please read through this summary and go to the next slide when you're ready.



Lastly, the 88D comms must also be taken into account. Since there is a radar or radars that are based out of your office, it is up to you to monitor where the data is going and if it is getting there, through the communication lines and making sure the data and products go to internal and external users. Please read through this summary and go to the next slide when you're ready.



This concludes the lesson on the RDA and RPG interface management. The next lesson will dive into the operational usage and functions of these windows. If you have any questions, please direct them to any of the email addresses listed.



Welcome to Topic 2 and the second of two lessons for RPG Management. After learning about the different buttons and windows in the RPG HCI for controlling and monitoring the system, we will now focus on the functionality of some of these buttons and windows, to see what they can do for you in an operational sense. So, let's get started!

Learning Objectives

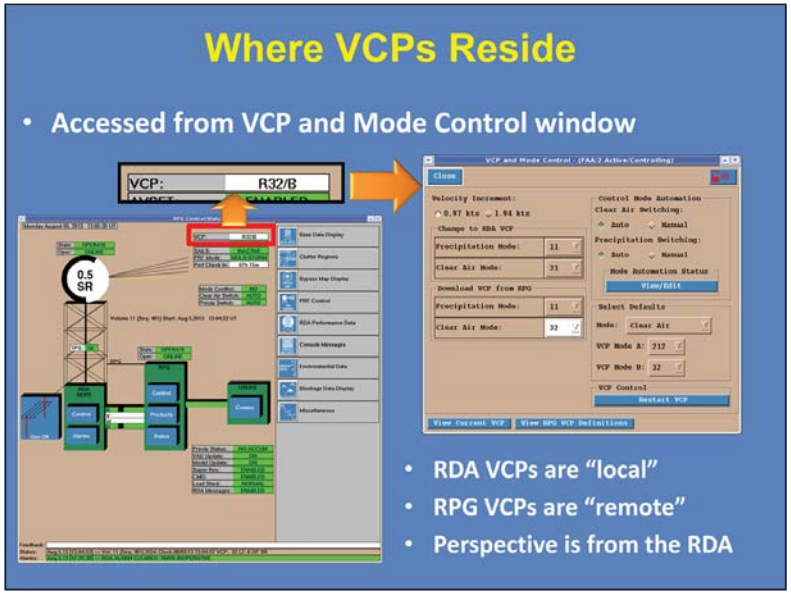
Through the RPG HCI:

- Define the three VCP groups in the system
- Identify the 3 PRF Control options
- Identify the 2 mode types for switching between VCP groups
- Identify and define the 3 environmental data parameters configured in the HCI

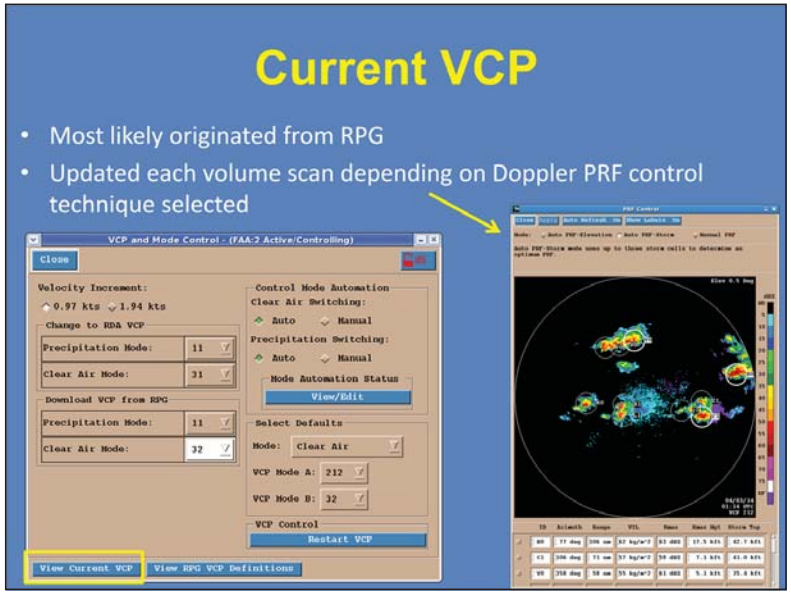
Performance Objective

- View and edit the 3 environmental data parameters from the HCI

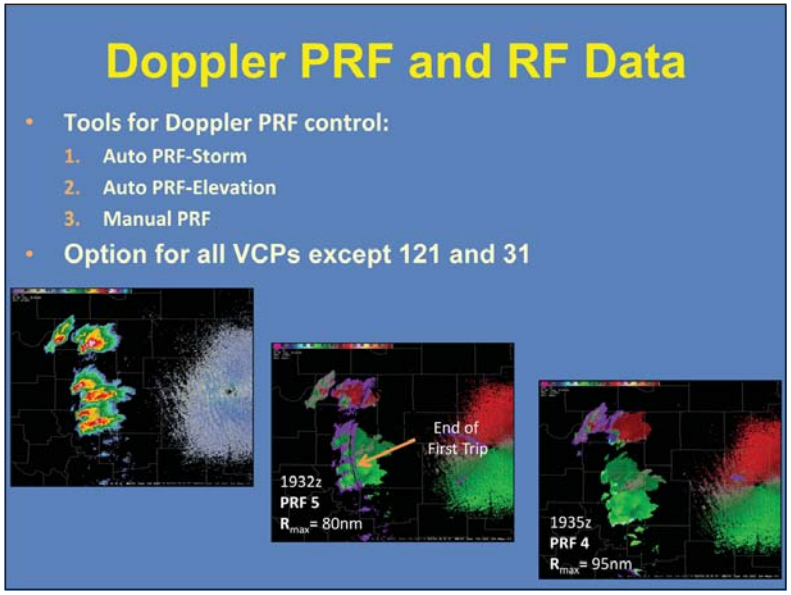
There are four learning objectives with this lesson and one performance objective. Please take a moment to read through these and move to the next slide when you're ready.



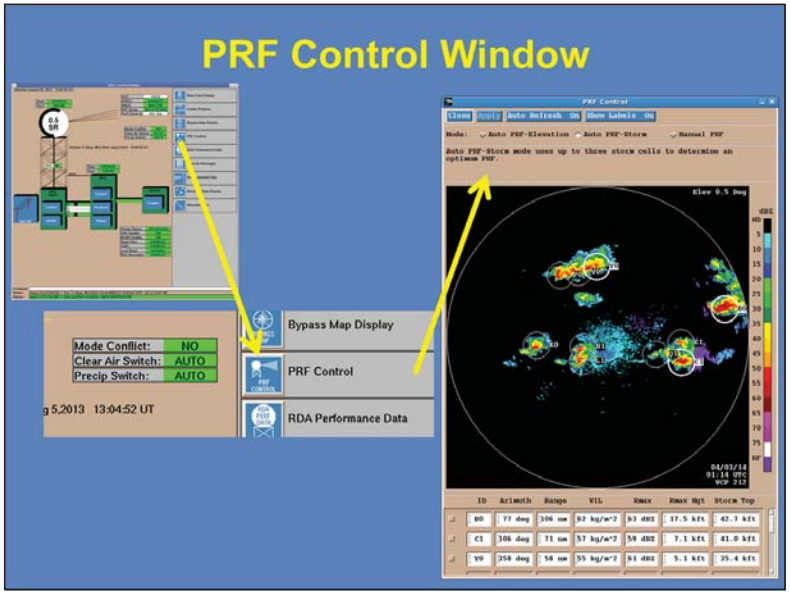
All of the VCPs are accessed from the VCP and Mode Control window. From the RPG HCI main page, clicking on the VCP button reveals the VCP and Mode Control window. There are VCPs stored at both the RDA and the RPG. The RDA VCPs, 11, 21, 31, and 32, have been stored at the RDA since the original WSR-88D deployment. At the RPG, the original 4 are available, as well the newer VCPs, such as 12 and 212. The VCPs stored at the RDA referred to as Local VCPs. The VCPs stored at the RPG are referred to as the Remote VCPs. The concept of Local vs. Remote is from the perspective of the RDA. If you look closely at the current VCP as it is listed on the front page of the RPG HCI, there are some extra letters. In this case, the R that precedes VCP 32 means that it's origin was Remote. Remote means it was downloaded from the RPG to the RDA. The /B that follows VCP 32 is a reminder that 32 is a Clear Air Mode VCP.



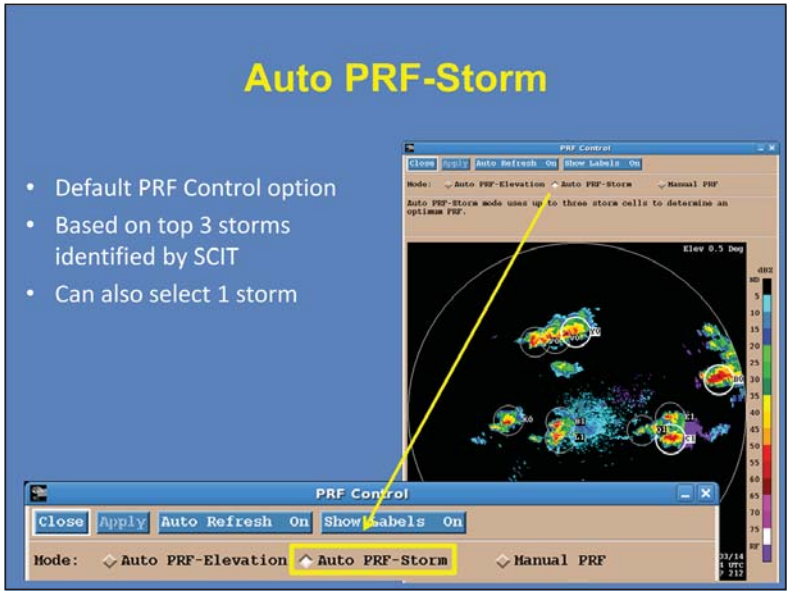
On the previous slide, VCP 32 was the current VCP. There is a "Current VCP" every volume scan. Though it may have the same VCP number, some characteristics of the Current VCP can change as often as every volume scan. For example, the Doppler PRF for the lower elevations can change, depending on the option chosen for PRF control – which is what we will take a closer look in the next slide.



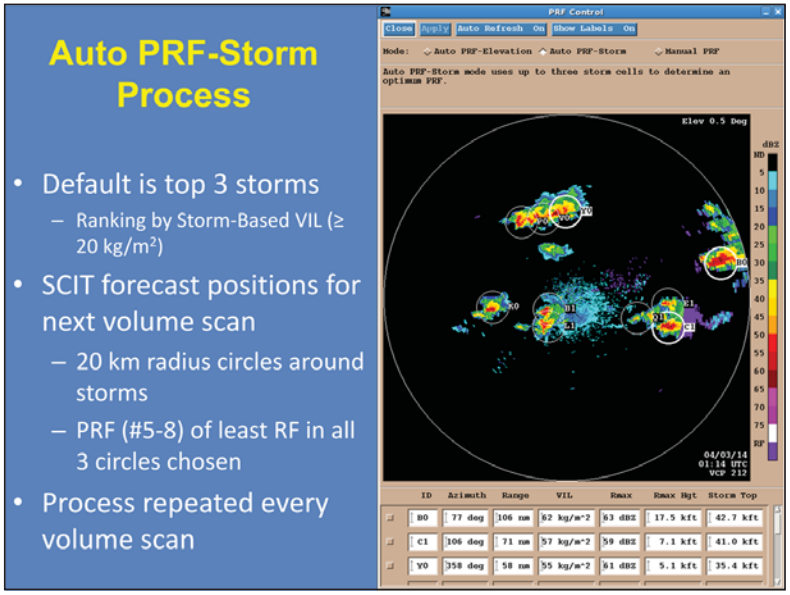
The WSR-88D routinely displays velocity data beyond the first trip, and there are five different Doppler PRFs. Each of these PRFs has a different resulting R_{max}. The end of the first trip can range from ~65 - 95 nm, depending on the Doppler PRF used. There are three different tools for Doppler PRF control, and they will be discussed in the order presented on this slide. None of these options apply to VCPs 121 or 31. The reflectivity and velocity images here show the power of applying the appropriate PRF Control technique.



All of the PRF Control techniques are selected from the PRF Control window. This window is accessed from the RPG HCI Main Page by clicking on the PRF Control applications button. The graphical portion of this window shows the 0.5° Reflectivity and the RF from the Velocity data only. Details of the graphical and table portions of the display change dependent on the PRF Control method chosen.

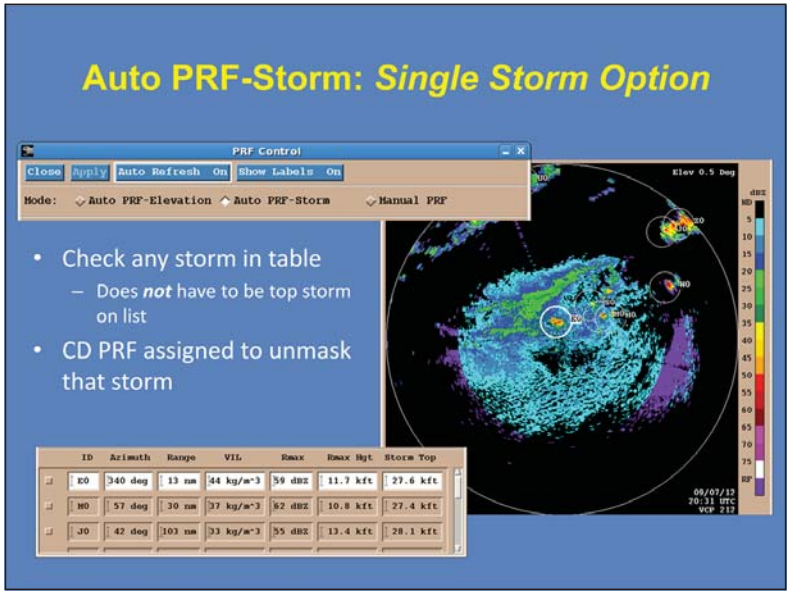


The default option is “Auto PRF-Storm”. It chooses the Doppler PRF that has the least amount of RF data for the top 3 storms. Auto PRF-Storm also has the option of protecting a single storm from RF data.

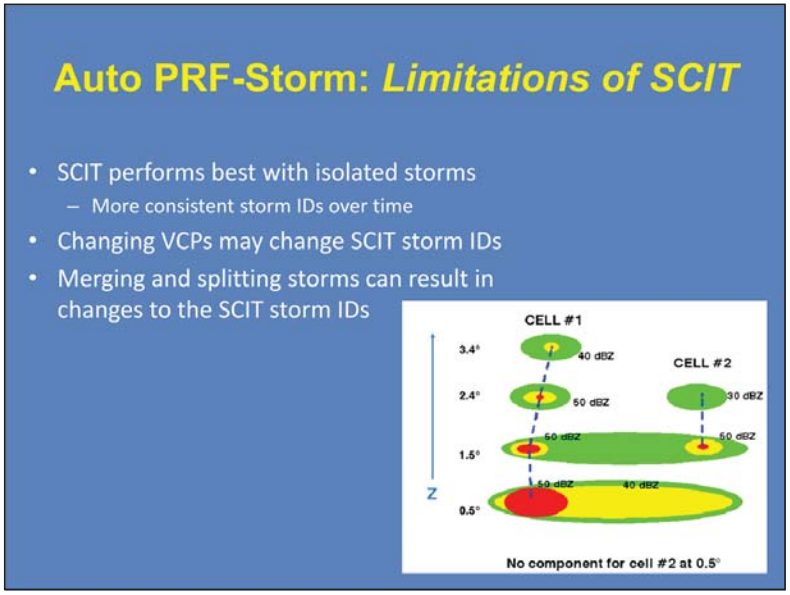


Auto PRF-Storm is reliant on storms identified by the Storm Cell Identification and Tracking (SCIT) algorithm, and the storm-based VIL of these storms. The default setting for Auto PRF-Storm selects the 3 most significant storms, based on storm-based VIL. The IDs and attributes (such as VIL) of the top three storms are listed in the table at the bottom of the PRF Control window. Those storms also have wide circles around them on the graphical area of the PRF Control window.

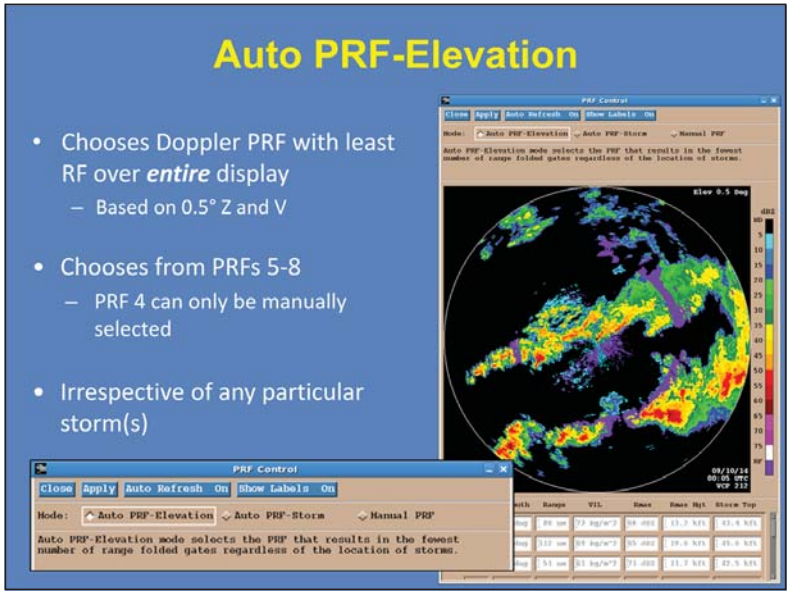
For these storms, the SCIT forecast positions for the next volume are used to create 20 km radius circles around each storm. The Doppler PRF (#5-8) with the least RF in these 3 circles is downloaded to the RDA for the next volume scan. This process is repeated every volume scan.



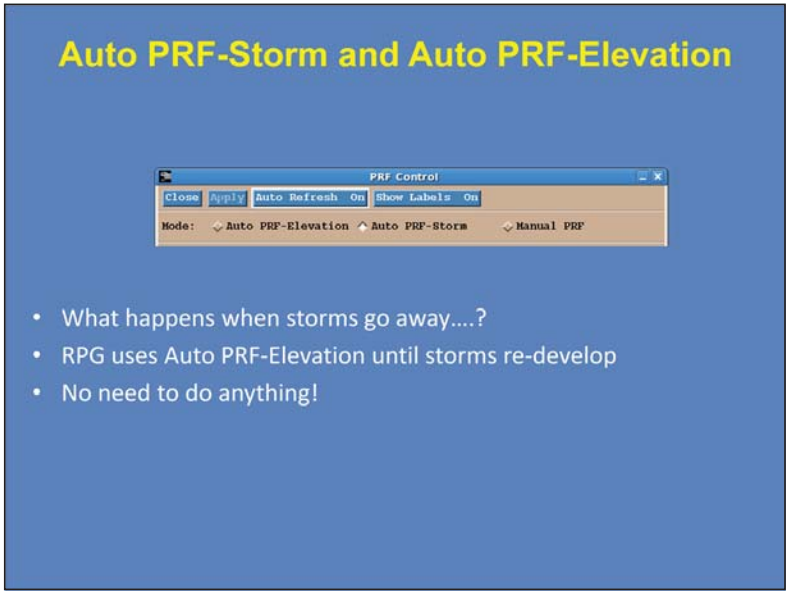
Auto PRF-Storm also has a single storm option, which can be invoked by clicking on the checkbox associated with a single storm in the table at the base of the PRF Control window. The chosen storm does not have to be the one with the highest VIL.



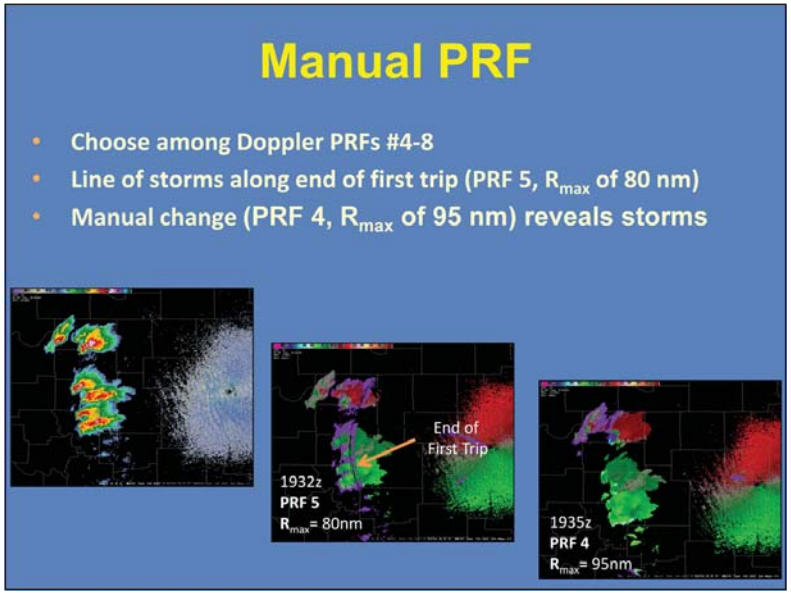
Since “storms” as defined by Auto PRF-Storm are those that are identified by the Storm Identification and Tracking Algorithm (SCIT), the limitations of SCIT must be kept in mind when using Auto PRF-Storm. SCIT performs best with isolated storms, which retain the most consistent storm IDs over time. There are many things that can change the ID for a given storm, such as changing the VCP or storms that merge or split. The diagram below shows how the SCIT algorithm loses cell #2 at the lowest degree because it is elevated and cell #1’s echoes bleed over into cell #2’s area. So SCIT would have trouble putting a track on cell 2 with this type of issue.



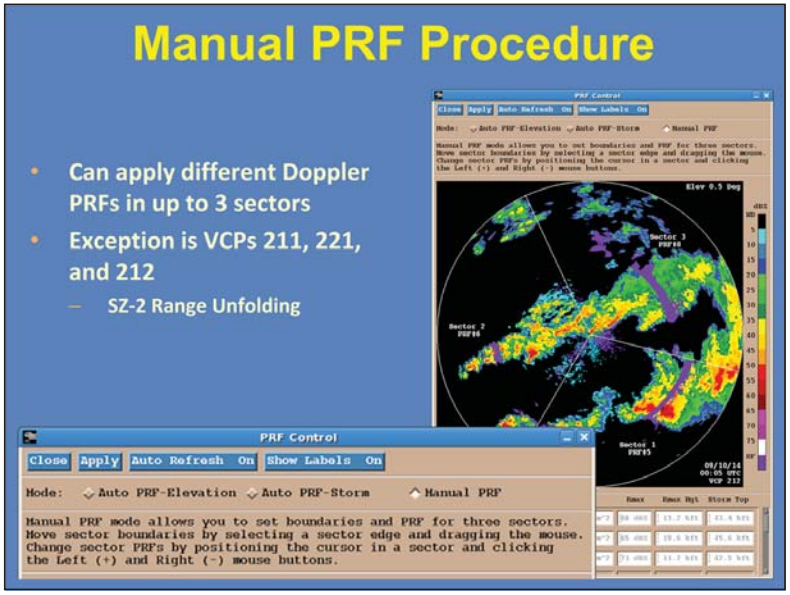
Auto PRF-Elevation uses the 0.5° reflectivity and velocity data, and examines the total areal coverage over the display for Doppler PRFs 5-8. The PRF with the lowest areal coverage of RF over the entire display is chosen and downloaded to the RDA for the next volume scan. Auto PRF-Elevation pays no attention to actual storms or intensities on the scope, only the area of RF coverage, so keep that in mind when choosing this option.



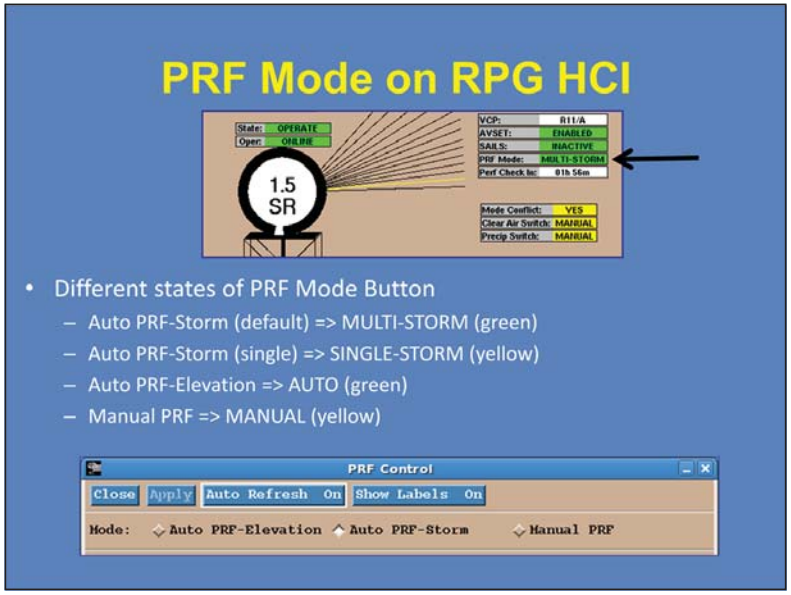
Now that the basics of Auto PRF-Storm and Auto PRF-Elevation have been presented, the RPG has an additional feature that is a good thing to know. When storms dissipate, Auto-PRF Elevation will automatically control the Doppler PRF. Once storms redevelop, Auto PRF-Storm will automatically control the Doppler PRF selection. There is no need to do anything!



The last option for PRF Control is Manual PRF. This allows for Doppler PRF selection directly. In this example, a line of storms is approaching the end of the first trip with PRF 5 chosen, which has an R_{max} of 80 nm. By choosing PRF 4, the R_{max} moves to 95 nm, and the velocity data for this line is no longer obscured by RF data.



The Manual PRF option allows for different Doppler PRFs in up to 3 different sectors. In the window on the right, probably the easiest thing to do is to take the mouse cursor, left-click once on the display portion of the window until you get the most appealing option for PRF sectors, with the least amount of purple haze and especially the least amount within the strongest storms on the scope. This sectorizing of Doppler PRFs is not available for VCPs 211, 221, and 212.



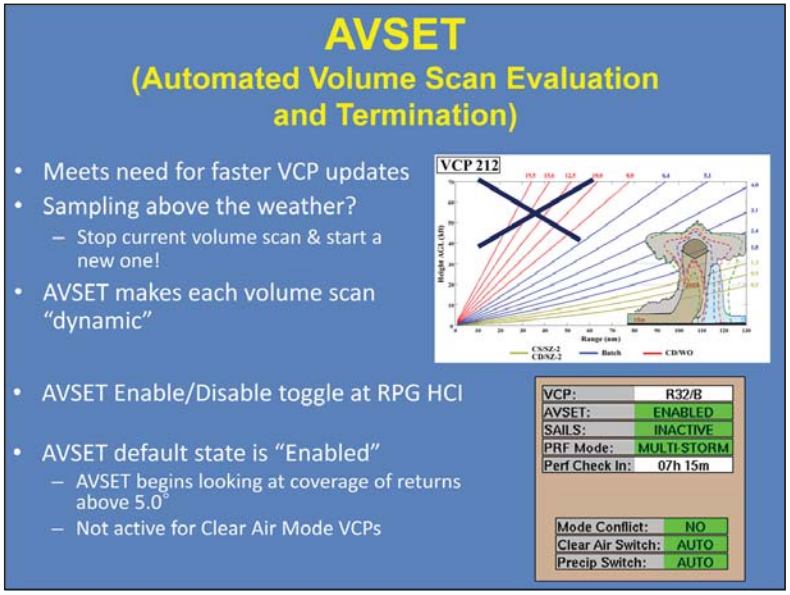
Depending on the PRF control option selected, there are four possible states for the PRF Mode button on the RPG HCI main page.

For Auto PRF-Storm (default), the button reads MULTI-STORM with a green background

For Auto PRF-Storm (single), the button reads SINGLE-STORM with a yellow background

For Auto PRF-Elevation, the button reads AUTO with a green background.

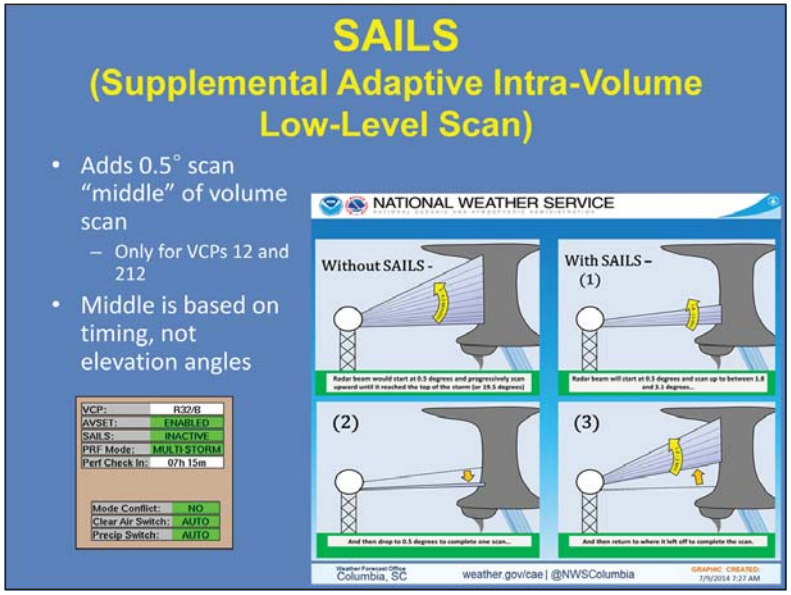
For Manual PRF, the button reads MANUAL with a yellow background



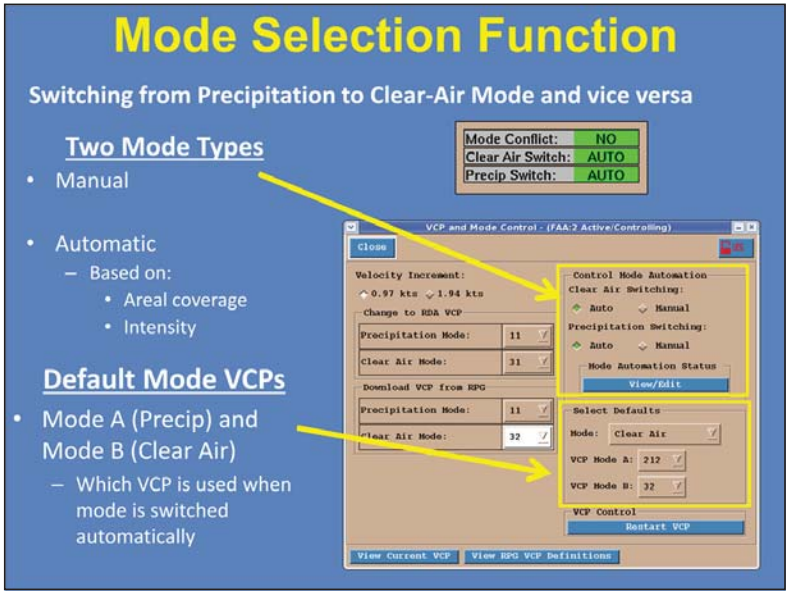
AVSET stands for Automated Volume Scan Evaluation and Termination, and is active for the Precipitation Mode VCPs. AVSET meets the need for faster VCP updates when the weather return is limited to the lower elevation angles. This is dependent on both the depth of any weather and it's range from the radar.

In this graphic, a convective storm is at a long enough range that the middle and upper elevations of the VCP are sampling nothing of significance for the NWS mission. AVSET makes each volume scan "dynamic". Once we are sampling "above the weather", AVSET stops the current volume scan and starts a new one. AVSET runs at the RDA and is commanded from the RPG.

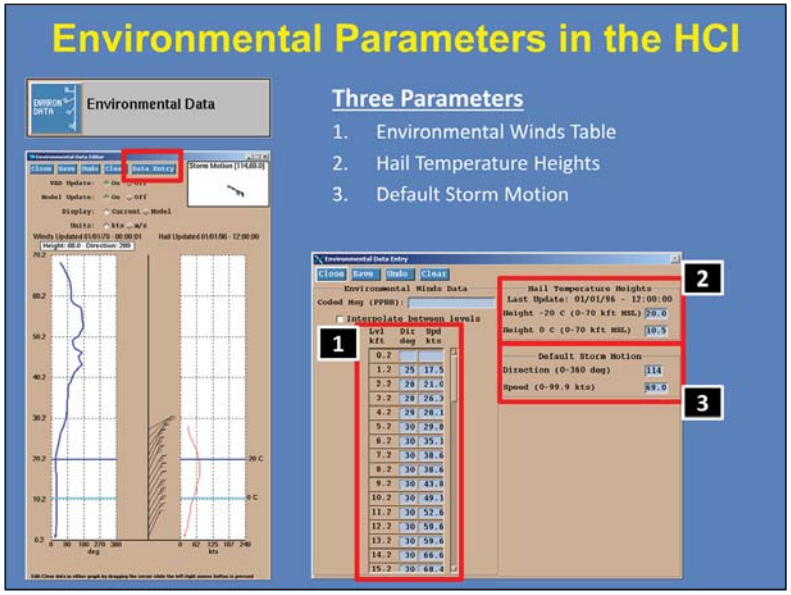
AVSET is controlled by a toggle at the RPG HCI, and it is Enabled by default. Since AVSET begins checking the areal coverage of returns above 5.0°, it is Enabled, but not active for the Clear Air Mode VCPs.



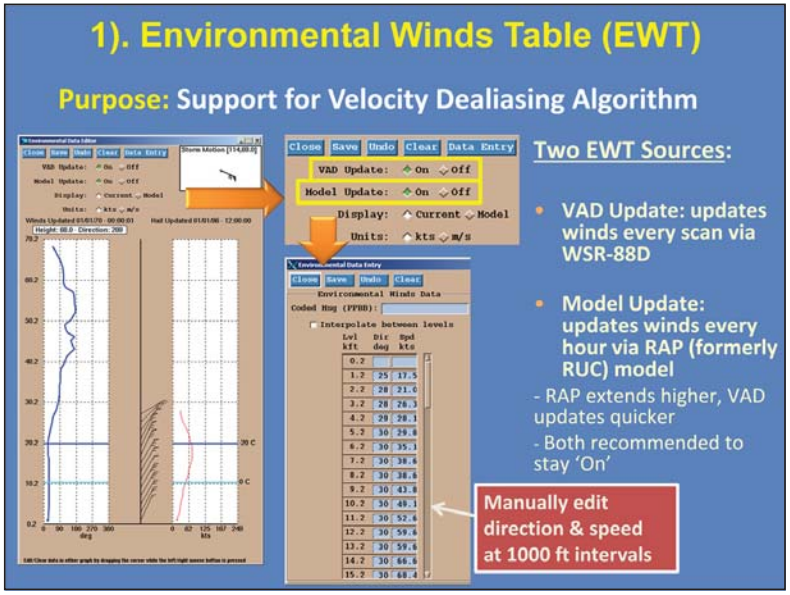
Another significant feature that makes VCPs dynamic is the Supplemental Adaptive Intra-Volume Low-Level Scan (SAILS). It is available only for VCPs 12 and 212. SAILS adds a 0.5° scan in the "middle" of the volume scan. The middle is not defined by the elevation angle, but by timing. The goal is for the 0.5° products to arrive with uniform timing. With AVSET on, the volume scan time changes, thus the timing of the SAILS scan changes. Many thanks to WFO Columbia, SC, for this great SAILS graphic posted to their Facebook page!



On to the Mode Selection Function, which controls how the WSR-88D transitions to and from Clear Air and Precipitation Mode. Notice that we are back to the VCP and Mode Control window. The two methods are Manual and Automatic. The Manual method requires more oversight and communication from shift to shift. The Automatic method is based on the areal coverage of returns with a given intensity. For Automatic Mode switching, there is the need for a Default VCP for each of the modes. Once the URC password is provided, these Default VCPs can be selected for each mode. This choice may be made seasonally. For example, for Precipitation Mode, VCP 212 may be preferred during the warm/convective season, while VCP 11 or 21 may be preferred during the cool season.



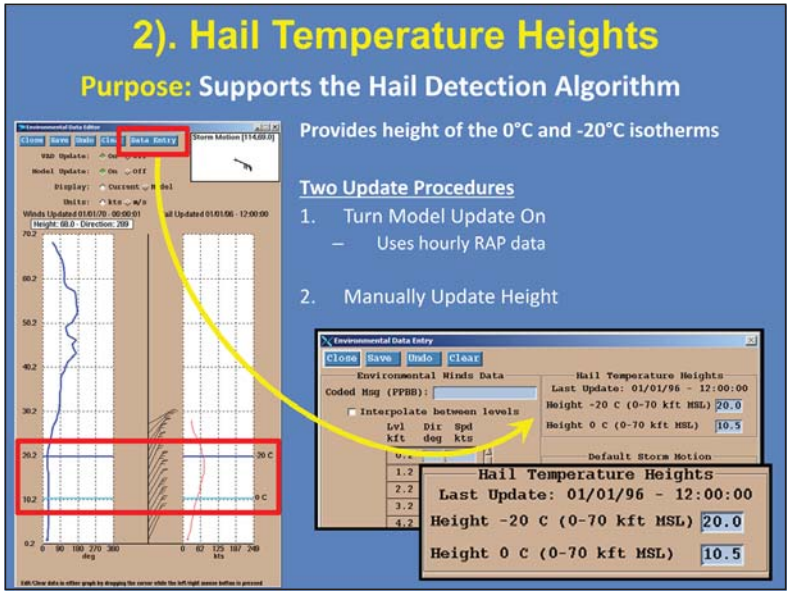
We'll close out this lesson with the overview of the features and functions on the Environmental Data collected and used by the radar. The main purpose of the Environmental Data is that it is utilized to support several RPG algorithms that go on to create output for AWIPS. On the HCI, you can monitor and configure the data or just let the data ride that already exists and is updated automatically. The three main environmental data parameters configured from this window are: the environmental winds table, the hail temperature heights, and default storm motion. All of which, we are going to delve into over the next 3 slides.



The environmental winds table parameter is very important to the quality of the base data as it supports the velocity dealiasing algorithm. In the Environmental Data Editor, the EWT components are displayed graphically on the window and are controlled through radio buttons at the top of the window. With the EWT supporting the velocity dealiasing algorithm, the EWT must be representative of winds aloft in real-time. To do this, the EWT can be updated from two sources.

1) When Velocity Azimuth Display (or VAD) Update is set to On, the VAD algorithm will update the EWT with WSR-88D generated winds aloft every volume scan.
2) When Model Update is set to On, RAP (formerly RUC) model output will send its gridded u/v components to the EWT table on an hourly basis.

Under most conditions, the RAP data will extend higher than the VAD, but the VAD will provide more frequent updates to the lower-level winds. As a result, it is recommended that VAD Update and Model Update are set to On most of the time. The current environmental winds can be interrogated in the graphical display below the command buttons. Here, wind direction and speeds are plotted as a function of height. Sometimes, you may have to manually modify the environmental winds table. This can be done in one of two ways. First, graphically by drawing your environment wind speed and direction inside the environmental data editor window. Second, you can insert wind speeds and directions into a table launched from the 'Data Entry' button at the top of the environmental data editor window. Inside, you will find heights separated out by default in 1000 ft. intervals.

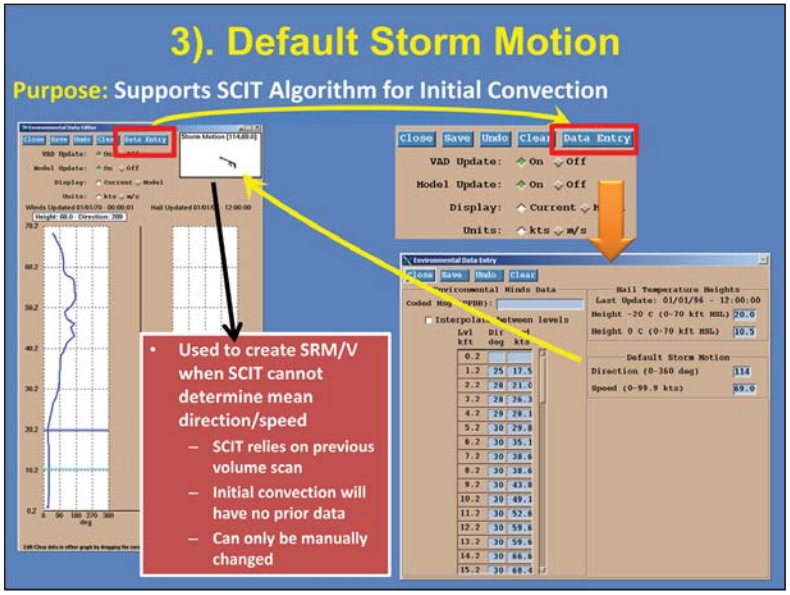


The hail temperature heights parameter supports the hail detection algorithm by providing the altitude of the 0 and -20 degrees Celsius isotherms.

There are two ways these isoterm heights are generated:

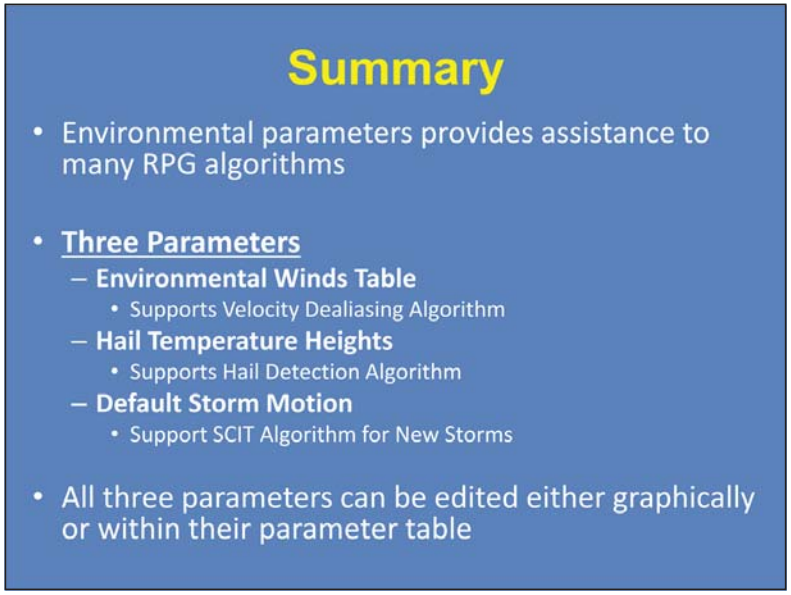
1) First, if Model Update is turned On. The RAP temperature analysis will be ingested hourly and this will set the hail temperature heights.
2) Second, you the operator can manually set these values based on another ingest source, such as current sounding data.

You can edit the heights graphically by clicking and dragging the 0 and -20 degree isotherm lines on the environmental winds table but that is much less precise than just inputting the values in the Environmental Data Entry window, as shown on the lower right.

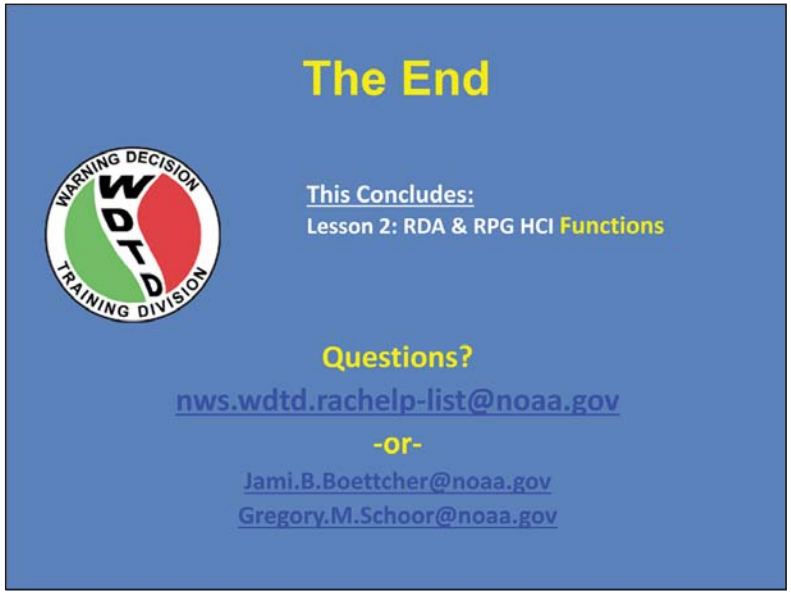


The default storm motion parameter supports the Storm Cell Identification and Tracking (or SCIT) Algorithm, specifically by using the default speed and direction for the creation of Storm Relative Mean Radial Velocity Products. In this example, though it's small, the storm motion is set to 114 degrees at 69 knots. To change this, if your hodograph analysis says different, click on the "Data Entry" button which will pop up this larger window on the bottom-right. You'll see the Default Storm Motion section with the 114 and 69 values, which you would then change and click "Save".

In most cases, this will be used only on initial cell development since the SCIT relates storms found in the current volume scan to storms in the previous volume scan. With new convection, no prior data exists so default storm motion will be used. This environmental parameter does not rely on any other ingested meteorological data and as a result can only be manually changed.



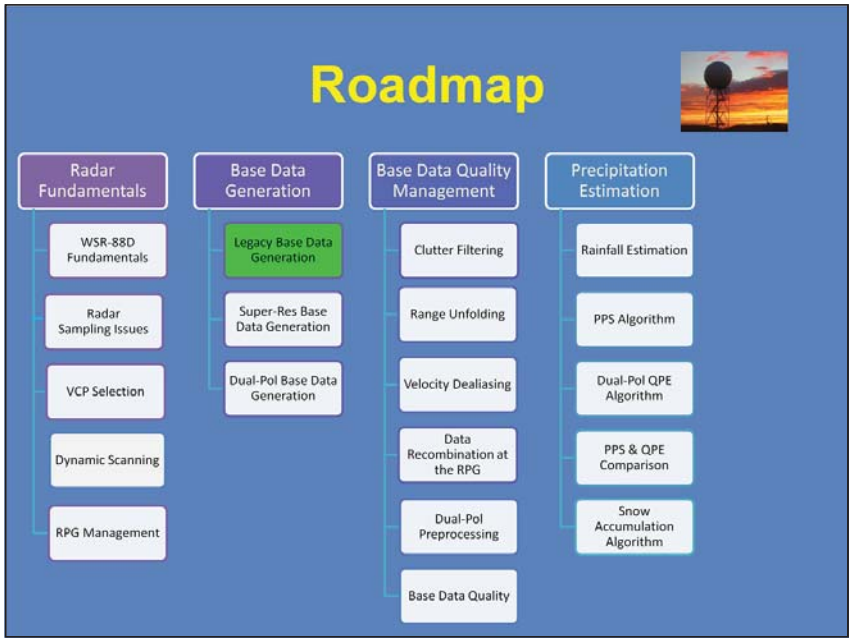
This is a summary of the features and functions of the most widely used and necessary buttons and windows from the RPG HCI. Please take a moment to review these and head out to the closing slide when you're ready.



If you have any questions, please direct them to any of the email addresses listed.



Welcome to Legacy Base Data Generation.



Here is the “roadmap” with your current location.

Base Data Generation
Objectives



1. Identify how Doppler information is obtained by the WSR-88D to determine atmospheric motion
2. Identify the relationship between V_{max} and the interval of first guess velocities
3. Identify how the returned signal is used to generate:
   a) Reflectivity (Z)
   b) Radial Velocity (V)
   c) Spectrum Width (SW)

There are 3 objectives in Legacy Base Data Generation, and these objectives will be taught in sequence during this module.

Doppler Effect



- "The change in frequency with which energy reaches a receiver when the receiver and the energy source are in motion relative to each other."

- **What matters: Frequency shift proportional to target motion**

You're likely familiar with the definition of the Doppler Effect. Since the radar location is fixed, any relative motion comes from the target's motion. The good news is the any Doppler frequency change is directly proportional to the target's motion.

Doppler Equation

$c = f \lambda$

$V_r = - (f_{dop} \lambda)/2$

c = speed of light
f = frequency
λ = wavelength

V_r = radial velocity
f_{dop} = Doppler shift
λ = WSR-88D wavelength

- Minus sign for inbound vs. outbound
- Factor of 2 for initial target illumination + backscatter

On the left is the basic relationship of frequency and wavelength to the speed of light. Based on that equation, on the right, is the relationship of the Doppler shift to radial velocity for Doppler weather radar (derivation not required!).

The factor of two is there because the signal is transmitted and interacts with the target, then is reflected back. The minus sign is there for target direction. By convention, inbound velocities are negative and outbound velocities are positive. For example, an inbound target produces a positive Doppler shift, making the velocity negative. An outbound target produces a negative Doppler shift, making the velocity positive.

Sound Waves & Doppler Shift

- Sound source moving 50 kts toward or away from receiver
  - Sound frequency = 10,000 Hz
  - Doppler shift = ± 800 Hz
    - +800 Hz inbound, -800 Hz outbound
- Doppler shift ~8% of original frequency
- Detectable by the human ear



A common way to demonstrate the Doppler effect is with the change in pitch of the sound of a train or ambulance as it first moves toward you, then moves away. In this example with a speed of 50 kts, the frequency shift is 800 Hz, +800 when the sound source is moving toward you and -800 Hz when the sound source is moving away from you. This Doppler shift is then 8% of the original frequency. That is why this type of Doppler shift is detectible by the human ear. We next look at the WSR-88D's listening ability, which is much more precise compared to the human ear!

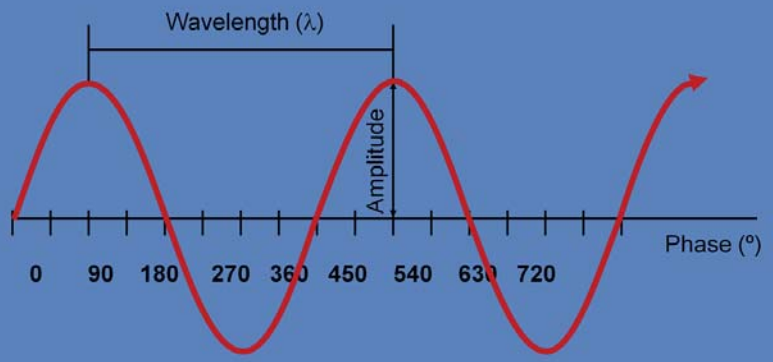
WSR-88D & Doppler Shift

- Target moving 50 kts toward or away from WSR-88D
  - Transmitted frequency = 2850 MHz (2,850,000,000 Hz)
  - Doppler shift = ± 487 Hz
- Doppler shift ~.00002% of original frequency
- WSR-88D does not directly measure such small frequency changes

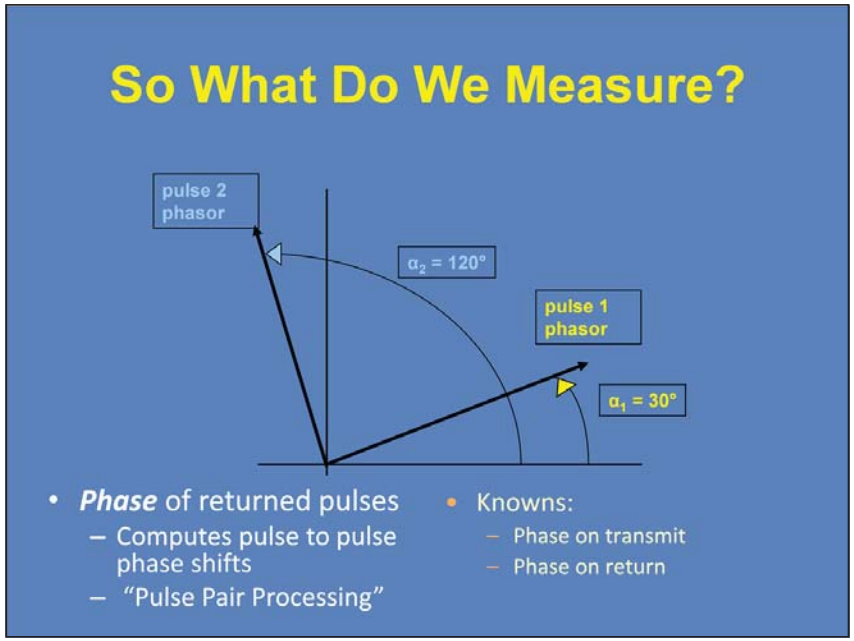


Now for the same target motion of 50 kts, but the listening device is the WSR-88D. We are no longer dealing with sound waves. Note that the frequency of transmission is very high compared to the frequency of a sound wave. The frequency shift is tiny compared to the original frequency…too small to be measured. Thus the shift in frequency is not what is used to determine target motion. We need something else.

So What Do We Measure?

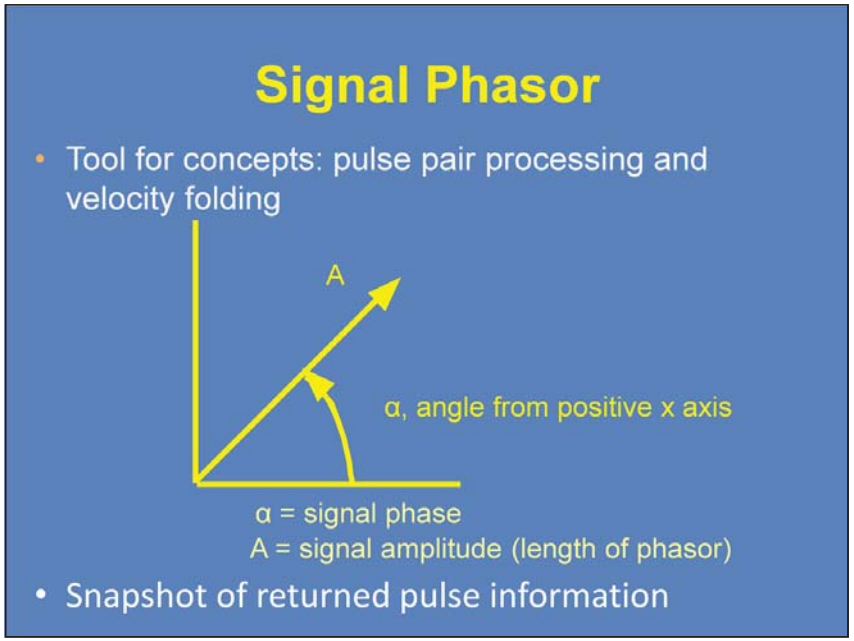


Since we don't measure the frequency shift, what else is there? Here is a reminder of the various characteristics of wave energy. The wavelength is the distance for one complete cycle, which is about 10 cm with the WSR-88D. The amplitude is the signal strength, which is directly related to reflectivity. The phase is a particular point along the wave, which can be used to determine velocity information.

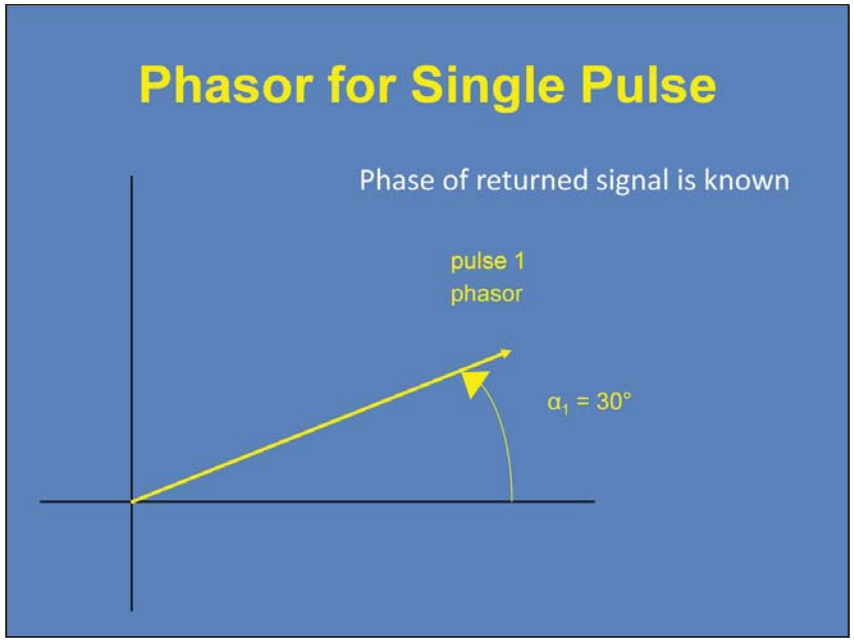


It turns out that the WSR-88D measures the phase of each returned pulse and is able to compare the phase values from one pulse to the next. The phase shift from one pulse to the next is directly related to the radial velocity. This technique is called Pulse Pair Processing.

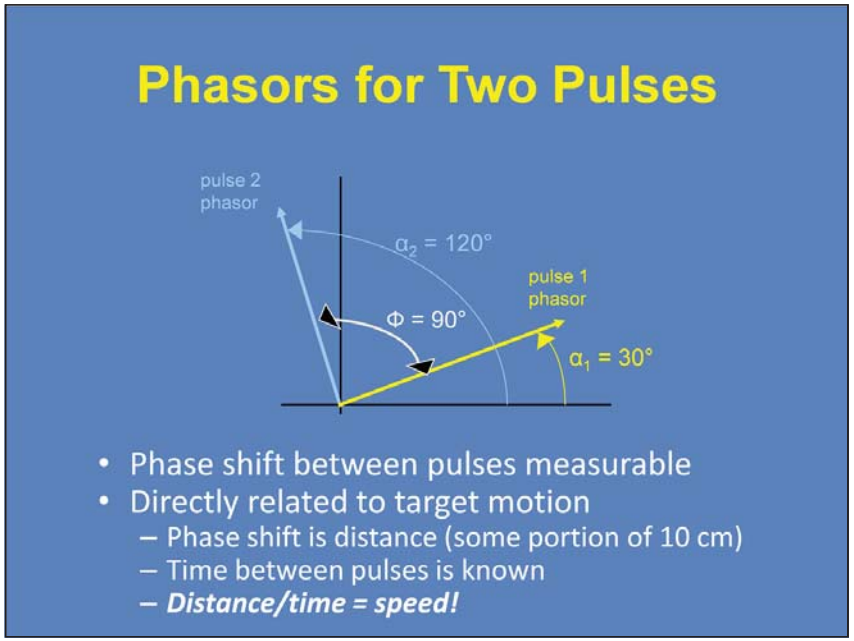
Pulse Pair Processing is possible because the initial phase is known when each pulse is transmitted. The initial phase for each returned pulse is also known. In a nutshell, a phase value is assigned to each pulse, then compared from one pulse to the next.



One way to represent the concept of a pulse pair phase shift is to use phasors. A phasor is a tool for temporary use only to support your understanding of one of the fundamental ambiguities with Doppler weather radar: velocity folding or aliasing. A phasor represents the necessary information from each returned pulse. The phase of that pulse is the angle of the phasor from the positive x axis. The length of the phasor is the signal amplitude. If the WSR-88D were continuously transmitting and receiving, the phasor would be rotating. However, pulses are needed for target range, and each phasor is a snapshot of information for each returned pulse.

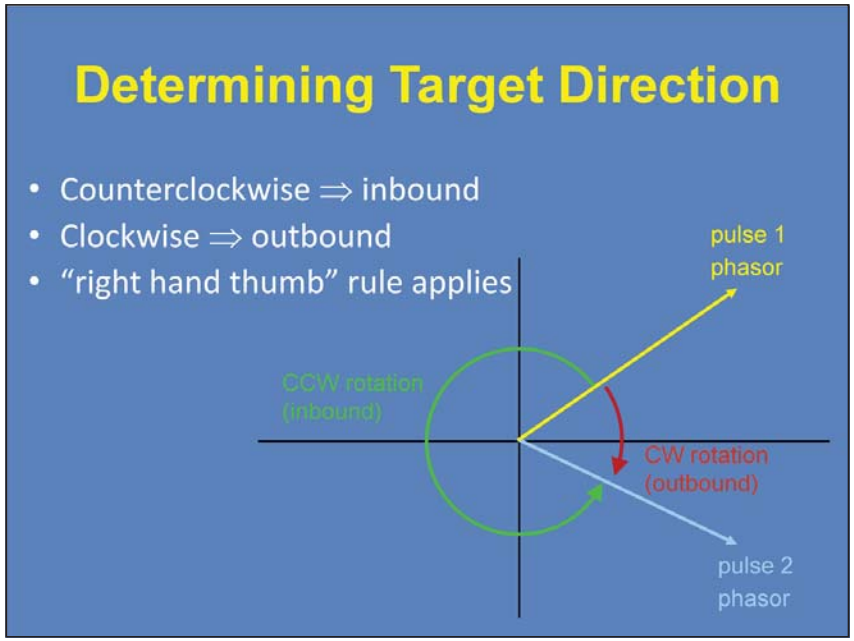


In this example, pulse 1 has been transmitted, has interacted with a target, and the returned signal has been processed. The phase value for pulse 1 is 30°.

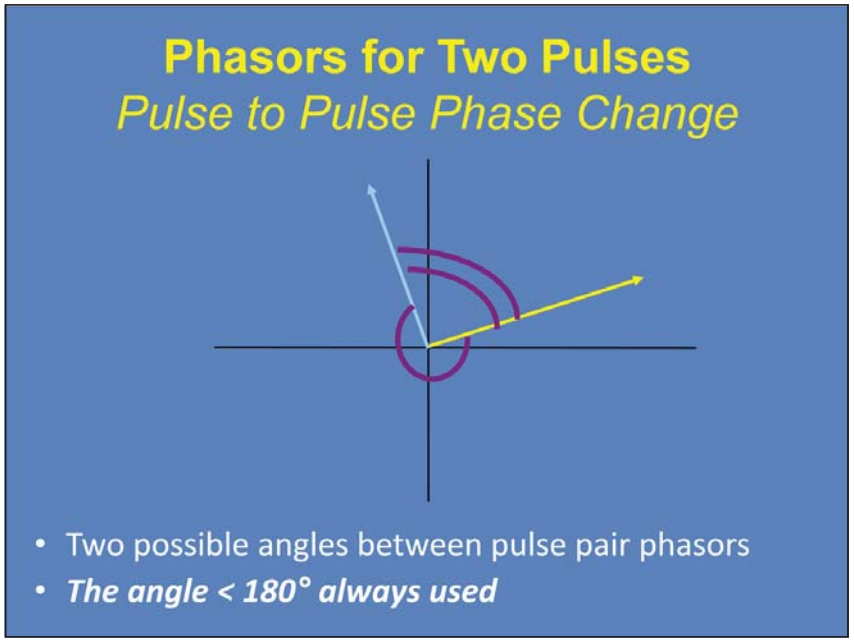


Using the pulse 1 phase of 30°, assume that the target is in motion and the phase value for pulse 2 is 120°. The angle between the two phasors (90°) is called the pulse pair phase shift.

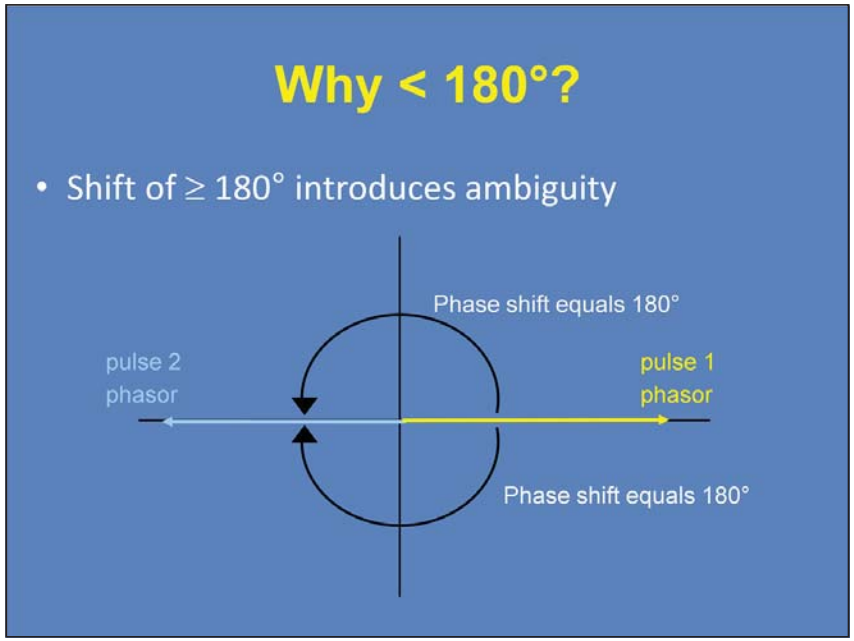
The key here is that the phase shift between pulses is directly related to target motion. Since the wavelength is 10 cm, the phase shift is distance, i.e. some portion of 10 cm. Since the PRF is known, the time between pulses is known. We then have both ingredients for target speed: distance and time.



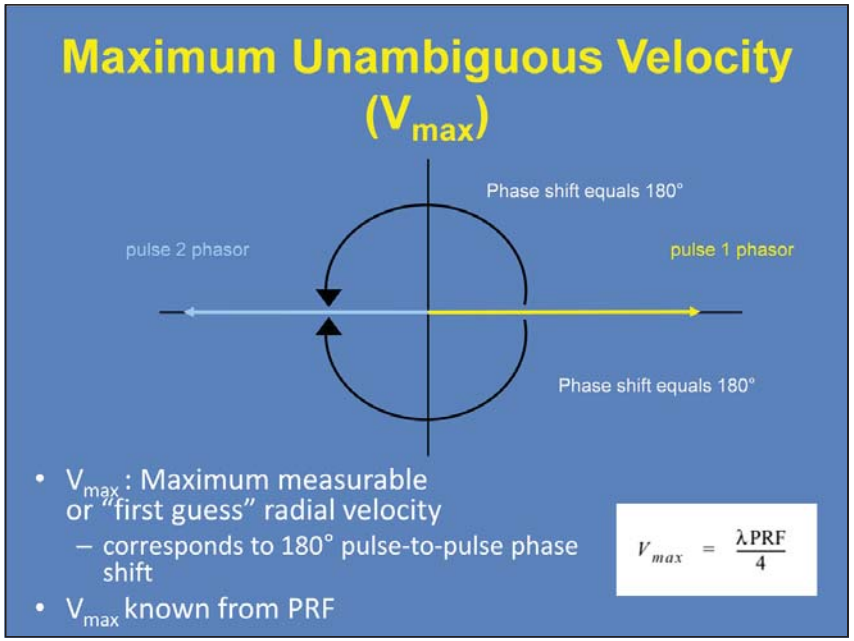
The target direction, inbound vs. outbound, is determined by the phasor rotation from pulse 1 to pulse 2. If clockwise, the direction is outbound. If the rotation is counterclockwise, the direction is inbound. You can also use the right hand thumb rule for the cross product of two vectors. Using the angle <180°, the result is clockwise rotation and outbound motion. If you use the right hand rule, your thumb would be pointing away from you.



There are two phasors representing the information for two different pulses. Since these phasors are snapshots, some assumptions have to be made about what happened in between the two pulses. There are two possible angles between the phasors that represent these pulses. Which one is used? It is always the angle < 180°.



Why use the angle < 180°? A phase shift of exactly 180° introduces ambiguity…it is unknown which direction the phasor rotated to get from pulse 1 to pulse 2. If the target moves so much between pulses that the true phase shift ≥ 180°, there is ambiguity in determining the velocity. Stay tuned for how we deal with that ambiguity.



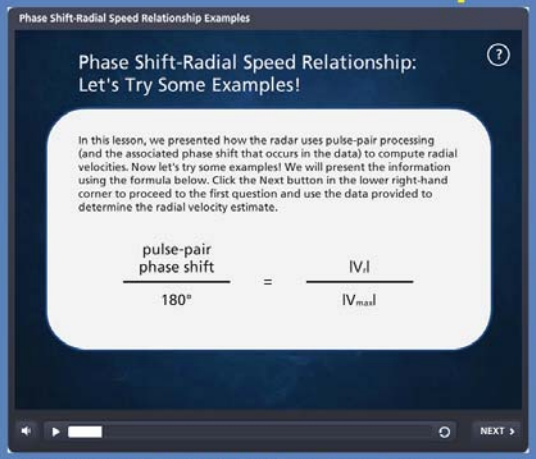
The maximum velocity that can be measured is called the maximum unambiguous velocity. It corresponds to a pulse pair phase shift of 180° (actually 179.99999…°), and is dependent on the pulse repetition frequency (PRF). With the WSR-88D, V_{max} values range from about 16 to about 64 kts.

Phase Shift-Radial Speed Relationship

$$\frac{\text{pulse-pair phase shift}}{180^{\circ}} = \frac{|V_r|}{|V_{max}|}$$

- $|V_r|$ = radial speed
- $|V_{max}|$ = maximum unambiguous speed

Phase Shift - Radial Speed

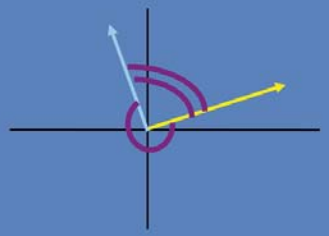


If no pop-up window appears that looks like the above, open a browser and go to:
http://www.wdtd.noaa.gov/courses/rac/principles/interactions/phaseshift-radialspeed/

Once the pulse-pair phase shift and the V_{max} are known, computing the first guess radial speed is straightforward. That's because the pulse-pair phase shift is some portion of the maximum shift of 180°, and the radial speed is that same portion of the maximum speed, or absolute value of V_{max}.

V_{max} and First Guess Velocities

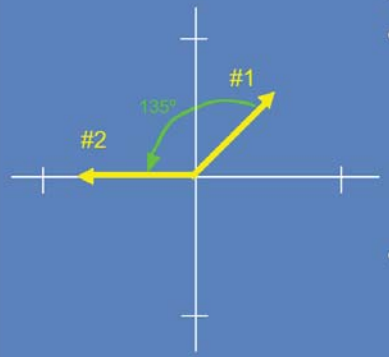
- First guess velocities from phase shift <180°
- V_{max} defines interval of first guess velocities
  - V_{max} = 60 kts; first guesses within ±60 kts
- Every first guess velocity has a set of known possible velocities, or aliases



The previous examples were all based on the pulse pair phase shift that is <180°. A velocity that is based on this assumption is called the "first guess velocity". Since V_{max} is associated with 180°, V_{max} then defines an interval of first guess velocities. For example, when V_{max} = 60 kts, the first guess velocities will be from -60 kts to +60 kts when V_{max} = 54 kts, the first guess velocities will be from -54 kts to +54 kts, etc.

Sometimes the first guess velocity is not the correct one, but the good news is that the other possible velocities are known and can be used if the first guess is incorrect.

Putting it Together: *First Guess Correct*



- V_{max} = 60 kts
- First guess phase shift 135°

$$\frac{135°}{180°} = \frac{|V_r|}{60}$$

60 (3/4) = $|V_r|$ = 45 kts

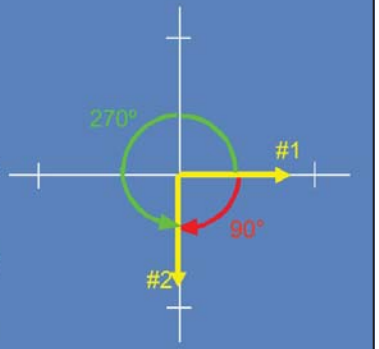
- Counterclockwise rotation
  - First guess radial velocity -45 kts

This example combines the concepts of pulse pair shift plus V_{max}, which gives us the first guess speed, along with phasor rotation which gives us target direction. In this case, we'll assume that the true phase angle between pulses is the one that is <180°, and that it represents the true target motion.

The true phase shift is 135°, and V_{max} = 60 kts. Since 135° is three fourths of 180°, the first guess speed is 45 kts (three fourths of 60 kts). In this case, using the angle <180°, the phasor rotation is counterclockwise. So the first guess velocity is -45 kts, and in this case, it is the correct radial velocity.

Putting it Together: *First Guess Incorrect*

- V_{max} = 60 kts
- First guess phase shift 90° with clockwise rotation
  - First guess radial velocity +30 kts
- Actual phase shift 270° with counterclockwise rotation
  - Actual radial velocity -90 kts



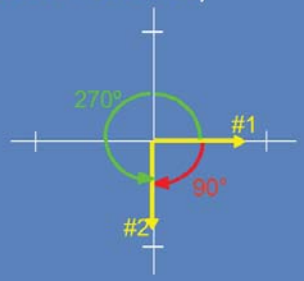
So what happens when the first guess velocity is not correct, i.e. the actual phase shift is >180°? The good news is that for every first guess velocity, there are other possible velocities which are known.

In this example, the first guess velocity is based on the phase shift of 90° in the clockwise direction. With a V_{max} of 60 kts, the first guess is then +30 kts. The true radial velocity, based on the phase shift of 270° in the counterclockwise direction, is -90 kts. Though +30 kts is incorrect, -90 kts is computed as a possibility velocity or alias. How these aliases are used to find the true radial velocity is discussed in a later lesson.

First Guess and Alias Velocities

- V_{max} defines interval of first guess velocities
  - V_{max} = 60 kts; first guesses within ±60 kts
  - V_{max} = 54 kts; first guesses within ±54 kts
  - Etc.
- Each first guess has aliases (possible velocities)

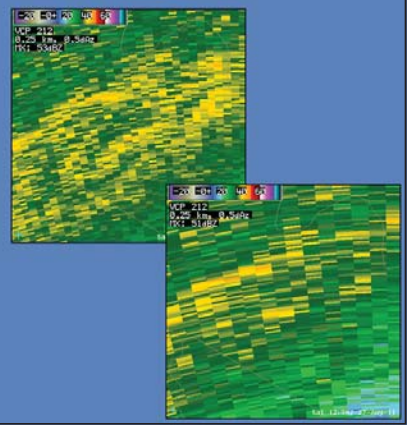
First guess radial velocity +30 kts
Possible velocities: -90 kts,



First guess velocities are based on the phase shift <180°, and V_{max} is the maximum unambiguous velocity, associated with a phase shift of 180. Each V_{max} thus defines an interval of first guess velocities. For example, for V_{max} = 60 kts, first guess velocities range from -60 kts to +60 kts, for V_{max} = 54 kts, first guess velocities range from -54 kts to +54 kts, etc. Since we cannot be certain that any first guess is correct, the good news is that the other possible velocities are known and can be used instead (more about that later).

Base Reflectivity (Z) Generation

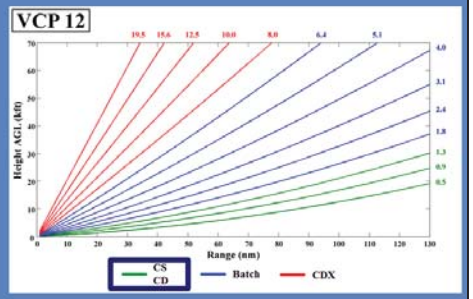
- Pulses/radial vary (6 to 64) for each .25 km range bin
- Average P_r converted to Z
  - Z converted to dBZ
- Super Res on Split Cuts
  - 0.5° azimuth x .25 km
- Legacy Res Batch & higher
  - 1.0° azimuth x .25 km



For each range and azimuth, there are multiple pulses used to generate base reflectivity. The returned power for these pulses is averaged for each .25 km range bin, then converted to Z using the Probert-Jones radar equation. The Z value is next converted to dBZ for product generation. For the Split Cuts, the best resolution Z product (what you use most of the time) is 0.5° azimuth by .25 km. For the Batch or higher elevation Z products, the best resolution to 1.0° azimuth by .25 km.

Base Reflectivity (Z) Generation

- Low PRF => long R_{max}
- Split Cut
  - 1st rotation CS/low PRF
    - Z & Dual-pol
  - 2nd rotation CD/high PRF
    - V & SW



A low PRF provides a long R_{max} and is used for Reflectivity data. Split Cut mode is used for the lowest 2 or 3 elevations for all the VCPs except VCP 121. Split Cut first uses one rotation in Contiguous Surveillance (CS), which is a low PRF mode. Base reflectivity and the dual-pol data are generated from the CS rotation. Then there is a second rotation at the same elevation in Contiguous Doppler (CD), which is a high PRF mode, used for base velocity and spectrum width. The trade off is that the R_{max} for CD mode is short, and multiple trip, range folded echoes are common.

The data collected from these two rotations are used together to “range unfold” velocity and spectrum width. The range unfolding techniques are presented in a later lessons.

Z Generated from P_{rH}

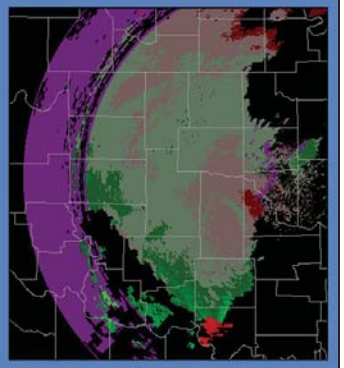
$$P_r = \left[\frac{P_t G^2 \theta^2 H \pi^3 K^2 L}{1024(\ln 2)\lambda^2}\right] \times \frac{Z}{R^2}$$

- P-J radar equation converts P_r to Z for both horizontal and vertical channels
- Base Reflectivity (Z) calculated from P_{rH} only

Base reflectivity is calculated from the average returned power, that is then converted to reflectivity from the Probert-Jones radar equation. Now that the WSR-88D has been upgraded to dual-polarization, this conversion from returned power to reflectivity is performed on both the horizontal and the vertical channels. However, the Base Reflectivity product is built from the horizontal channel information only.

Base Velocity (V) Generation

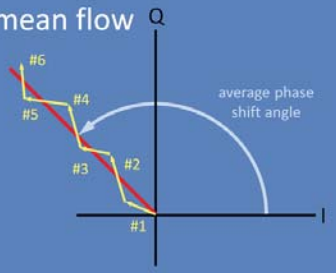
- Super Res on Split Cuts
  - 0.5° azimuth x .25 km
- Legacy Res Batch & higher
  - 1.0° azimuth x .25 km
- Maximum display range 162 nm
- Pulse pair processing
  - Phase changes between successive returned pulses averaged



Velocity also has a .25 km range resolution. Just as with reflectivity, super resolution is defined as 0.5° azimuth and is available only for the Split Cut elevations. For the Batch and higher elevations, velocity data has an azimuthal resolution of 1.0°. The maximum display range for velocity is 162 nm. Pulse pair processing refers to averaging the phase changes between a series of returned pulses to achieve a velocity estimate. However, this is not a linear average.

Base Velocity (V) Generation

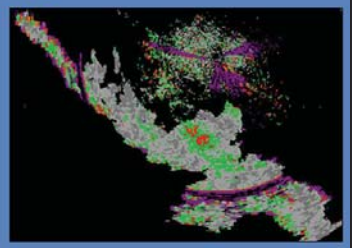
- Pulse pair processing
  - Averaging of pulse pair phase changes ***not*** linear
  - Average is ***power weighted***
  - Larger scatterers affect V average
  - Larger scatterers move with mean flow



Instead of a linear average of pulse pair phase shifts, the average is weighted toward those pulses that return higher power. This means that the larger scatterers in the volume will have a greater influence on the velocity estimate, and the larger scatterers are more likely to move with the mean flow.

For the calculation, each phasor in this graphic represents the information for one pulse pair. The phasor's angle from the positive x axis is the pulse pair phase shift, while differences in phasor length relate to the returned powers of the two pulses. It turns out that a vector sum of these phasors results in a power weighted average pulse pair phase shift.

Base Spectrum Width (SW) Generation



- Super Res on Split Cuts
  - 0.5° azimuth x .25 km
- Legacy Res Batch & higher
  - 1.0° azimuth x .25 km
- Maximum display range 162 nm
- Measure of velocity dispersion
  - Proportional to variation in wind speed/direction
- SWs typically high with
  - Boundaries, thunderstorms, high shear

Spectrum width also has a .25 km range resolution. Just as with reflectivity, super resolution is defined as 0.5° azimuth and is available only for the Split Cut elevations. For the Batch and higher elevations, spectrum width has an azimuthal resolution of 1.0°. The maximum display range for spectrum width is 162 nm.

Spectrum width is a measurement of the velocity dispersion or variability within a range bin. It is proportional to the variability of wind speed and direction. Spectrum widths can be expected to be high in areas such as boundaries, thunderstorms or any high shear environment.

Base Spectrum Width (SW) Generation

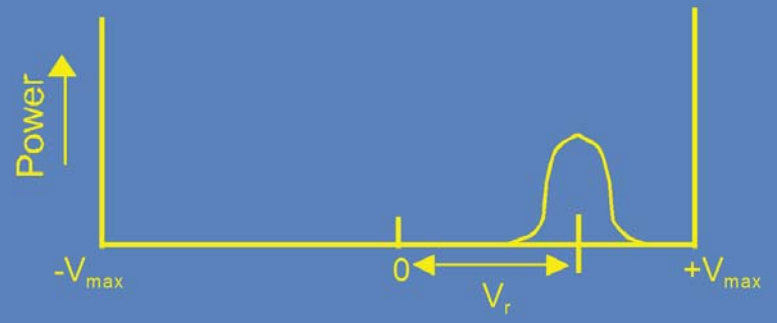
- SW technique "Autocorrelation"
  - How successive pulse pair shifts correlate to one another
- High variation in phase shifts =>high SW
- Low variation in phase shifts =>low SW
- Visual tool: Doppler power spectrum
  - "weather" well approximated by Gaussian curve



The technique used to calculate Spectrum Width is called Autocorrelation. What's being correlated? The series of phase shifts from one pulse to the next. If there is a lot of variation in the pulse pair phase shifts, the spectrum width will be high. If there is little variation, the spectrum width will be low. The best way to visualize spectrum width is through the "Doppler Power Spectrum". It turns out that "weather" can be well approximated by a Gaussian curve. The Doppler Power Spectrum is a representation of the base data analysis process.

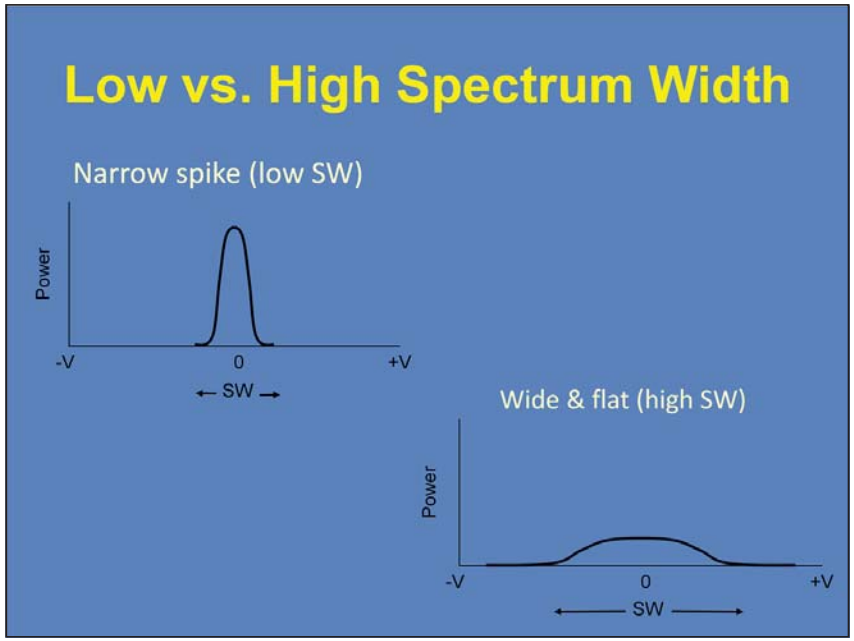
Doppler Power Spectrum

- Base data assignment for a single range bin



First Guess Radial Velocity

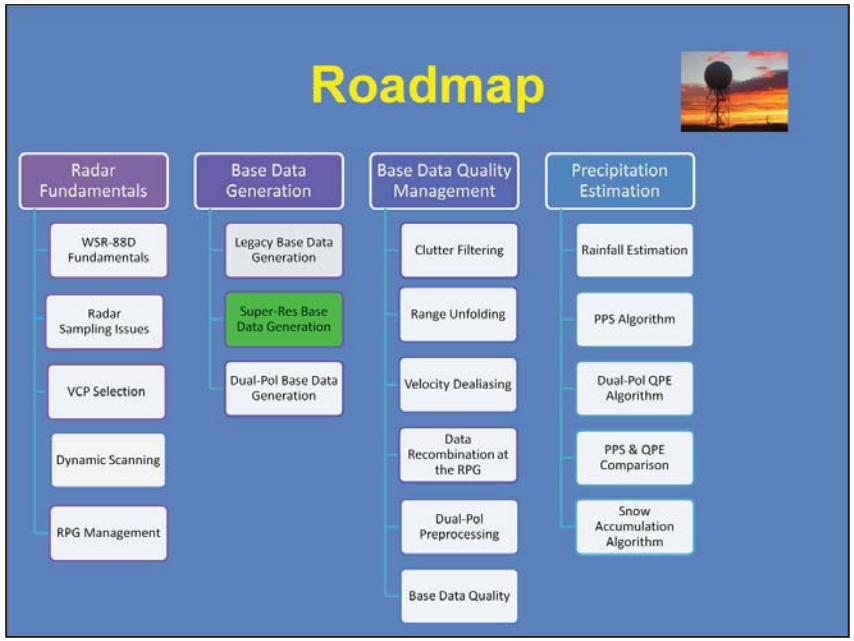
The Doppler Power spectrum represents the base data analysis process for a single range bin. The power and velocity information from a series of pulses are converted to points known at "spectral coefficients". A bell curve is fit to these coefficients. The average returned power (thus reflectivity) is the area under the curve. The mean radial velocity is where the midpoint of the curve falls along the horizontal axis. The width of the curve is proportional to the magnitude of spectrum width.



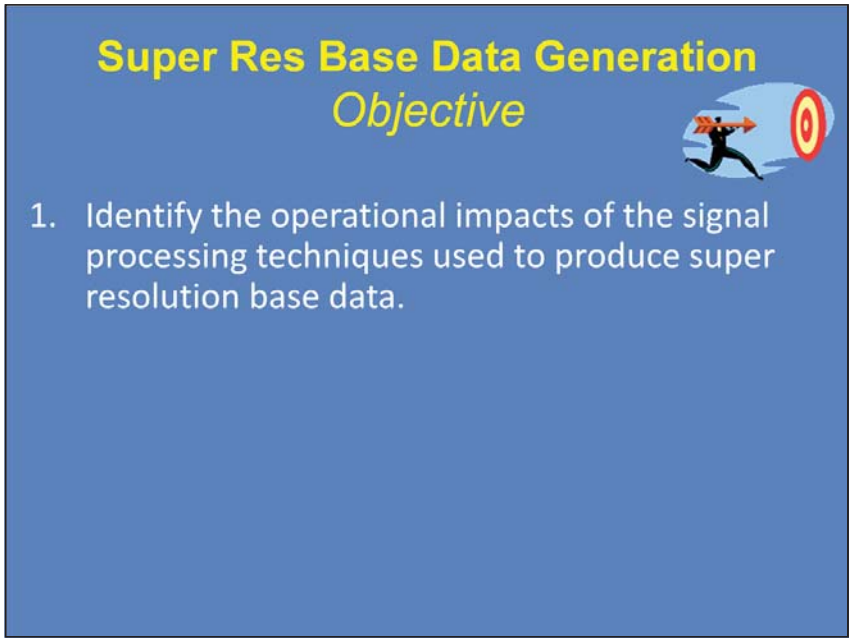
The magnitude of spectrum width will vary depending on the shape of the power spectrum. Returned pulses from ground clutter will likely have strong power and near zero velocity (upper left image). There is minimal variation in pulse pair phase shifts and the Doppler Power Spectrum curve is narrow and centered near zero velocity. This also results in a low spectrum width. On the other hand, some type of weather is returning low power, but a wide variety of velocity values (lower right image). The average velocity is near zero, but the width of the Doppler Power Spectrum is much greater than with the clutter example, and the associated spectrum width value would be high.



Welcome to Super Resolution Base Data Generation.

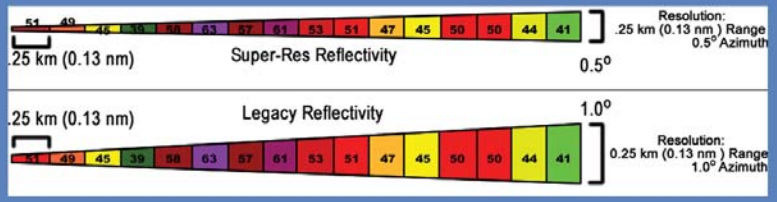


Here is the “roadmap” with your current location.



There is one objective in Super Resolution Base Data Generation.

Super Resolution Signal Processing



- Super Resolution: 0.5° azimuth for legacy base data on Split Cuts
- How to narrow azimuthal resolution to 0.5°?
  - Overlapping radials
  - Data windowing
- "Effective" beamwidth
  - Physical beamwidth (single pulse) plus the antenna is moving

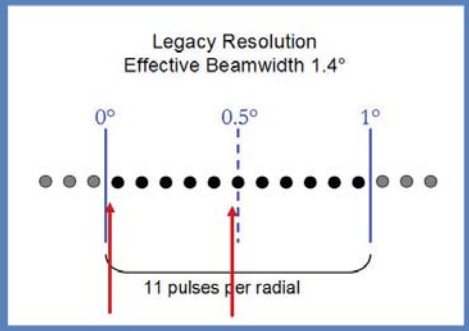
Super resolution is defined as a 0.5° azimuth. It is available only for legacy base data on the Split Cut elevations.

The upgrade to super resolution was based on signal processing techniques, not on new hardware. There are two signal processing techniques used to narrow the azimuthal resolution from 1.0° degree to 0.5°, overlapping radials and data windowing.

In order to understand this approach, we need to start with the concept of effective beamwidth. The "beamwidth" of ~1.0° presented in WSR-88D Fundamentals is based on the antenna being stationary. Antenna motion produces what is know as the effective beamwidth.

Effective Beamwidth

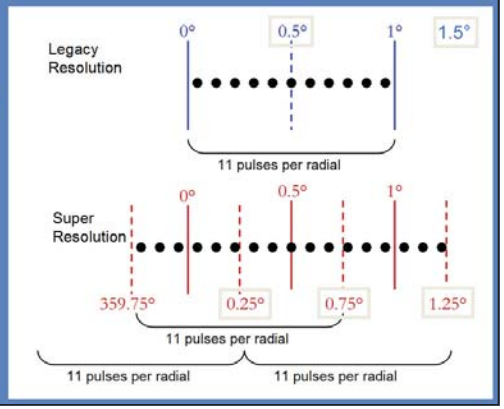
- *Effective* beamwidth is ~1.4°
  - Beamwidth ~1° for a single pulse



In order for a pulse to be used for the base data estimate for a radial, the beam centerline must be somewhere within that radial. In this example, each dot represents the location of the beam centerline for a single pulse. There are 11 pulses that fall within this simplified 1° radial. As the antenna rotates, the pulse with a centerline that is just on the inside edge of this radial still has a physical beamwidth of 1°, so the beam is sampling a volume that is both inside and outside the radial. Only when the beam is centered on this radial is the associated pulse contained within the radial. This process of capturing volumes outside a radial as base data are estimated is what is known as the effective beamwidth, which for the WSR-88D, is about 1.4°.

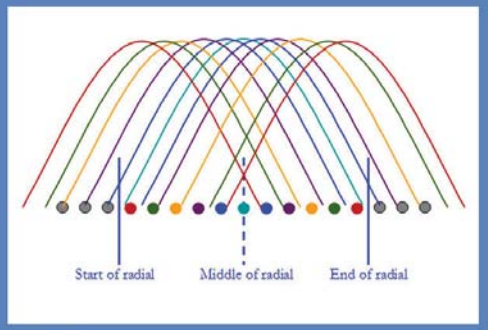
Overlapping Radials

- Overlapping radials: change defined center of each radial
- Effective beamwidth is ***still*** 1.4°



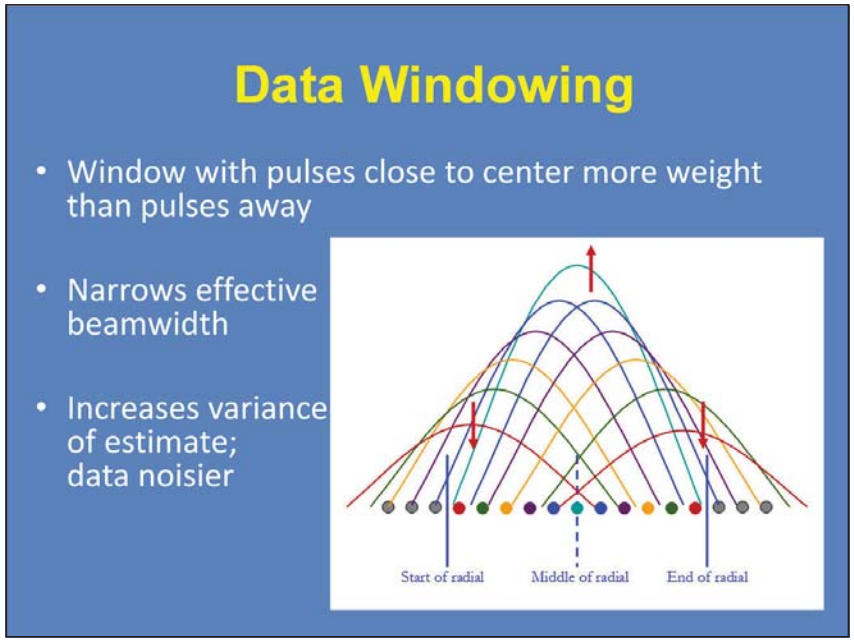
The process of overlapping radials is simply changing the definition of the center of each radial. The radial centers for legacy resolution are 0.5°, 1.5°, 2.5°, etc. The radial centers for super resolution are 0.25°, 0.75°, 1.25°, 1.75°, etc. However, since the number of pulses per radial cannot decrease and the effective beamwidth is still 1.4°, the volume that is sampled outside of each 0.5° radial is too large. Simply choosing new radial centers is not sufficient. The next step is needed to narrow the effective beamwidth.

Data Windowing

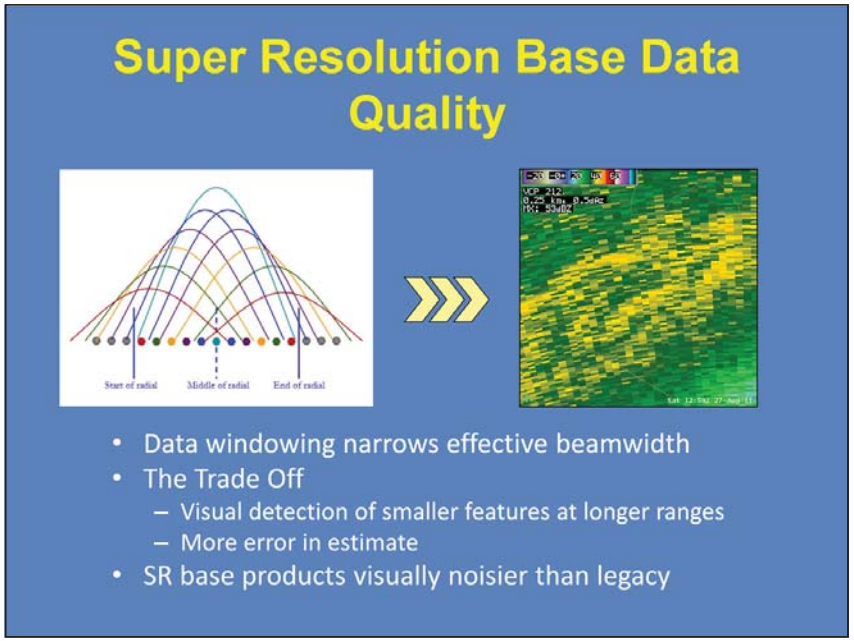


- Data Windowing
  - Apply weighting function to pulses that comprise base data for a radial
- Rectangular window: all pulses have equal weight

Data windowing is the next step required to achieve super resolution by narrowing the effective beamwidth. Windowing is a signal processing technique that applies a weighting function to the pulses that are used to generate base data for a particular radial. A technique that applies equal weight to all of the pulses is known as the Rectangular window.



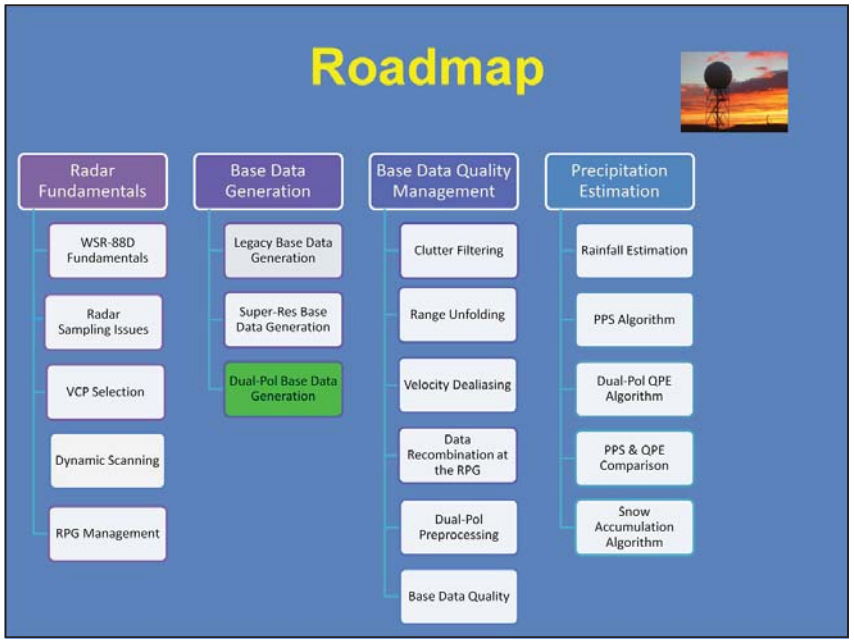
What is needed for super resolution is a window that narrows the effective beamwidth. This is accomplished by giving pulses near the center of the radial more weight, progressing to less weight applied to pulses away from the center. This technique meets the need of narrowing the effective beamwidth, but it does introduce more variance or error in the estimate. This increase in error occurs because some of the pulses are overemphasized while others are underemphasized. The result is that the data on super resolution products is noisier than the legacy resolution.



The 0.5° azimuthal resolution is obtained through signal processing techniques, specifically data windowing. This technique narrows the effective beamwidth, but there is a trade off. Smaller scale features are visually detectable at longer ranges with super resolution, but there is also more error in the base data estimate. In general, super resolution base products are visually noisier than corresponding legacy resolution base products. Many of the RPG algorithms were not designed to ingest super resolution base data.



Welcome to Dual Polarization Base Data Generation



Here is the “roadmap” with your current location.

Dual Pol Base Data Generation
Objectives

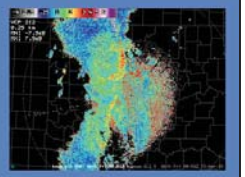


1. Identify how the returned signal is used to generate
   a) Differential Reflectivity (ZDR)
   b) Correlation Coefficient (CC)
   c) Differential Phase (Φ_{DP}), then Specific Differential Phase (KDP)
2. Identify the similarities and the differences between SW and CC
3. Identify the radar volume characteristic that has the greatest impact on the magnitude of KDP

There are 3 objectives in Dual Pol Base Data Generation. These objectives will be taught in sequence during this lesson.

RDA Generation of ZDR

- ZDR calculated from P_{rH} and P_{rV}
  - $Z_{rH} = P_H C_H r^2$
  - $Z_{rV} = P_V C_V r^2$



$$ZDR = 10\log_{10}\left(\frac{Z_H}{Z_V}\right)$$

$$= 10\log_{10}\left(\frac{P_h}{P_v}\right) + 10\log_{10}\left(\frac{C_h}{C_v}\right)$$

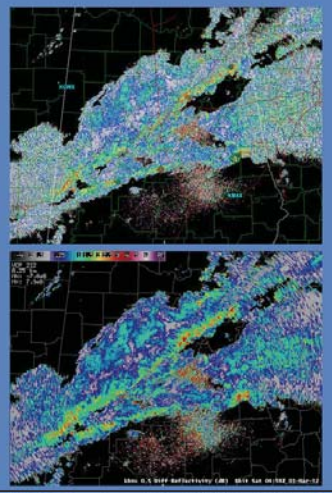
⟸ Calibration of both channels matters!

From the Probert-Jones radar equation, the Z value for each channel is equal to the returned power, times the range squared times a constant that is based on the radar's calibration.

Just as with Base Reflectivity, ZDR is calculated from returned power, but ZDR uses returned power from the horizontal and vertical channels. The ZDR equation can be written as ZDR = 10log (Z_H/Z_V). When the Zs are substituted with returned power, the radar constants and the range, the range terms cancel. I'm showing the logarithmic separation of returned power and the radar constant terms to underscore the importance of calibration of both channels for an accurate ZDR. There are operational implications of the need to calibrate both channels that will be presented later in the course.

Dual Pol on Split Cuts

Dual Pol from CS pulses:
- Avoids range folding
- At RDA, 0.5° azimuth (noisy)

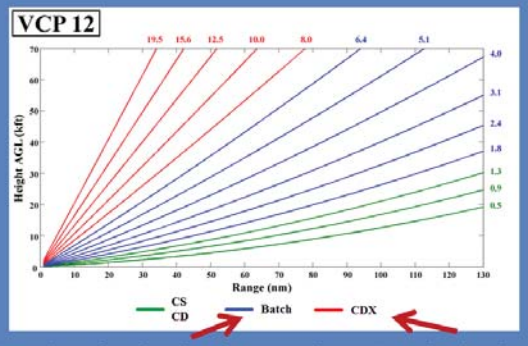


- Recombined at RPG to 1.0° azimuth

For the Split Cut elevations, ZDR and the other Dual Pol data are built from the low PRF, Contiguous Surveillance pulses on the first rotation. This is to avoid multiple tripping or range folding, and thus the need to "range unfold" the ZDR data. As with the legacy base moments (Z, V, & SW), the azimuthal resolution for ZDR on the Split Cuts is 0.5°. With this resolution, ZDR, and all the other Dual Pol data are even noisier in appearance than Z, V & SW. On the upper right is a ZDR image built from the 0.5° azimuthal resolution Level II data. On the lower right is that same data that has been recombined to 1.0° azimuth and smoothed. This is how ZDR appears on the AWIPS display.

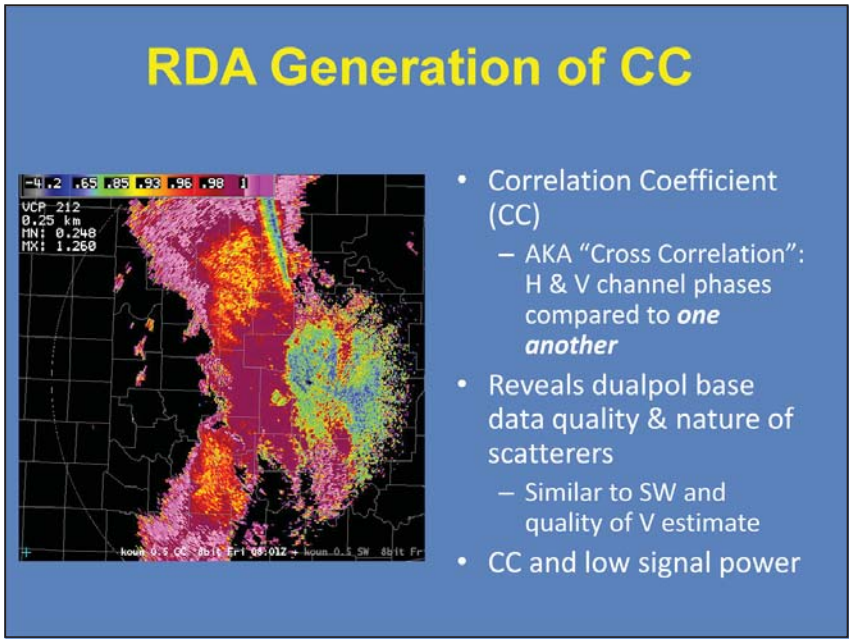
All of the Dual Pol base data are recombined at the RPG to 1.0° azimuth and smoothed, before the products are built, and there will be more about this "preprocessing" in a later lesson.

Dual Pol on Batch & Above



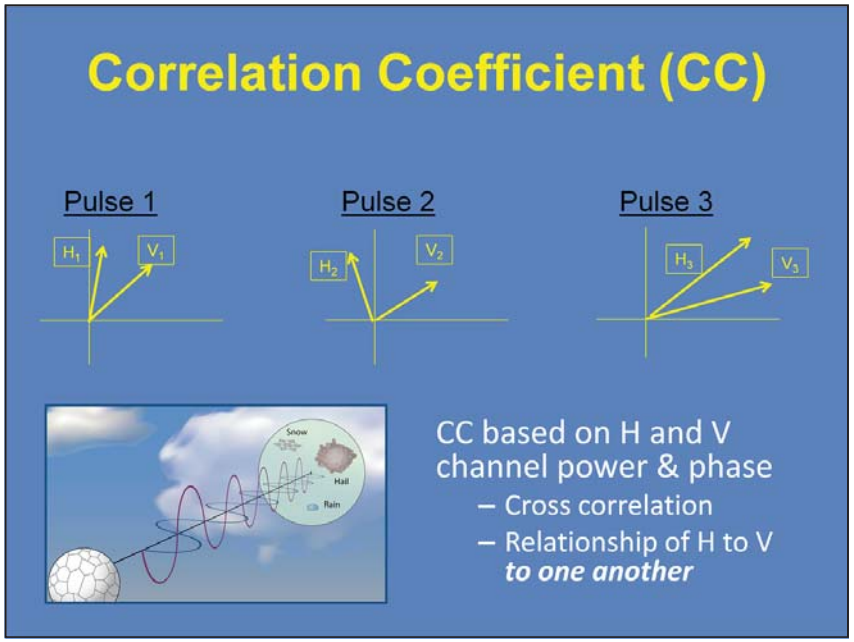
- Batch: CD pulses for dual pol
  - More CD pulses than CS
  - 1.0° azimuth
- Above Batch: dual pol built from CDX pulses
  - That's all there is!

For the Batch elevations, the antenna makes a single rotation alternating between low PRF, Contiguous Surveillance, and high PRF, Contiguous Doppler, mode. ZDR, and the other Dual Pol data, are built from the Contiguous Doppler pulses for each radial, since there are more of them. It is possible to see range folded, RF, data on the Dual Pol products at these elevations since there is no range unfolding applied. The azimuthal resolution is 1.0°, so no recombination is required. For the elevations above Batch, ZDR and the other Dual Pol data are built from the CDX pulses because that is all that is used. Since multiple trip echoes are so unlikely at these higher elevations, the CDX means Contiguous Doppler with no range unfolding.

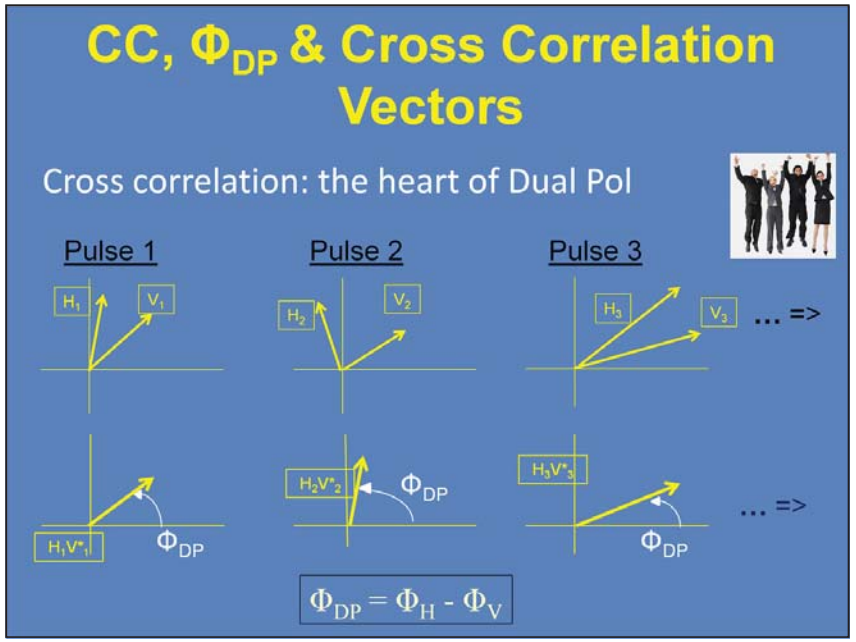


Correlation Coefficient (CC) measures the consistency of the H and V returned power and phase with one another for each pulse. This "cross correlation" looks at how the returned power and phase of one channel compares to the other channel. If the consistency is high (for example, stratiform light rain is being sampled), the phase change with one channel is similar to the phase change with the other channel. CC's measure of consistency reveals information on the nature of the scatterers. For example, uniform hydrometeors are much more consistent than ground clutter or smoke.

CC also provides information on the quality of the Dual Pol base data estimate, in some ways similar to the relationship between spectrum width and velocity. Spectrum width measures the consistency of the phase shifts from one pulse to next, which then relates to the reliability of the associated velocity value. There are both similarities and differences between CC and SW coming up soon, along with what happens to the validity of CC values in areas of very weak returned signal.



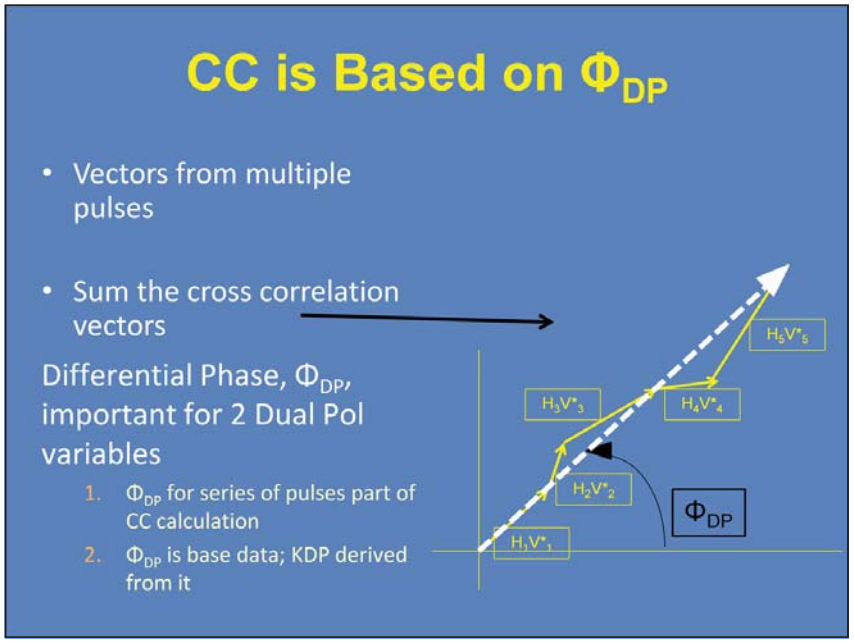
For each pulse, the returned power and phase from the H and V channels is known and can then be compared to one another. This type of comparison is know as a cross correlation. Of interest here is the magnitude of and the angle between the H and V vectors, which can be determined by vector multiplication. This "cross correlation" vector (next slide) of H and V is checked for each pulse.



Cross correlation vectors are at the heart of Dual Pol base data. Here are a series of pulses, with the individual phasors for the H and V channels of each pulse. The cross correlation vector for each pulse captures how the horizontal and vertical information relate to one anther.

The cross correlation vector for each pulse is computed by multiplying the H vector by the complex conjugate of the V vector. This multiplication creates a new vector whose phase is the angle between H and V. This angle is known as Φ_{DP}. It is the horizontal phase minus the vertical phase: $\Phi_{DP} = \Phi_H - \Phi_V$

There is a Φ_{DP} for each individual pulse, as well as an average Φ_{DP} value assigned to each range bin, known as the Differential Phase.



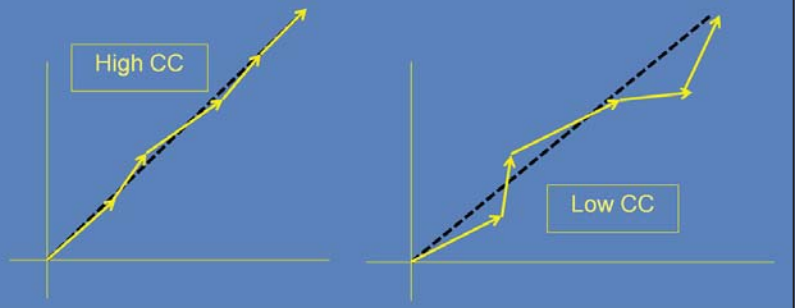
Since we don't assign any type of base data with just one pulse, the cross correlation vectors for a series of pulses are summed. This vector sum (white dashed arrow) is what's needed for the remaining two Dual Pol variables. Differential Phase, Φ_{DP}, is the angle of this vector sum, and it is part of the base data generated at the RDA for each range bin.

It turns out that Differential Phase, Φ_{DP}, is important for two of our Dual Pol variables:

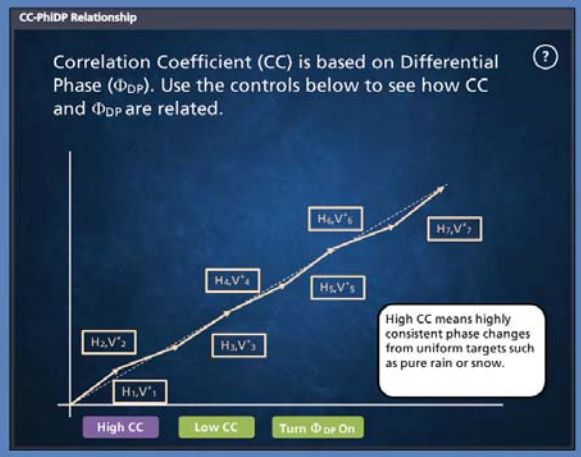
1. Φ_{DP} for a series of pulses is part of the calculation of CC.
2. Φ_{DP} is base data generated at the RDA and sent to the RPG. Specific Differential Phase, or KDP, is derived from it.

CC is Based on Φ_{DP}

- CC = length of cc vector ÷ averaged H & V powers (huh?)
- 0 < CC < 1
  - fraction of "perfect" consistency



CC-PhiDP Relationship



If no pop-up window appears that looks like the above, open a browser and go to: http://www.wdtd.noaa.gov/courses/rac/principles/interactions/cc-phidp/

For a little more detail, CC is calculated by taking the length (or amplitude) of the vector sum of the cross correlation vectors and dividing it by the averaged H and V powers. This calculation captures the variation of the individual cross correlation vectors that contributed to the sum.

CC is a unitless variable, a number between 0 and 1. It is a fraction of "perfect" consistency of scatterers.

If pure rain is being sampled, there is minimal variation between the H and V channels, the cross correlation vectors line up nicely, and CC is close to 1.

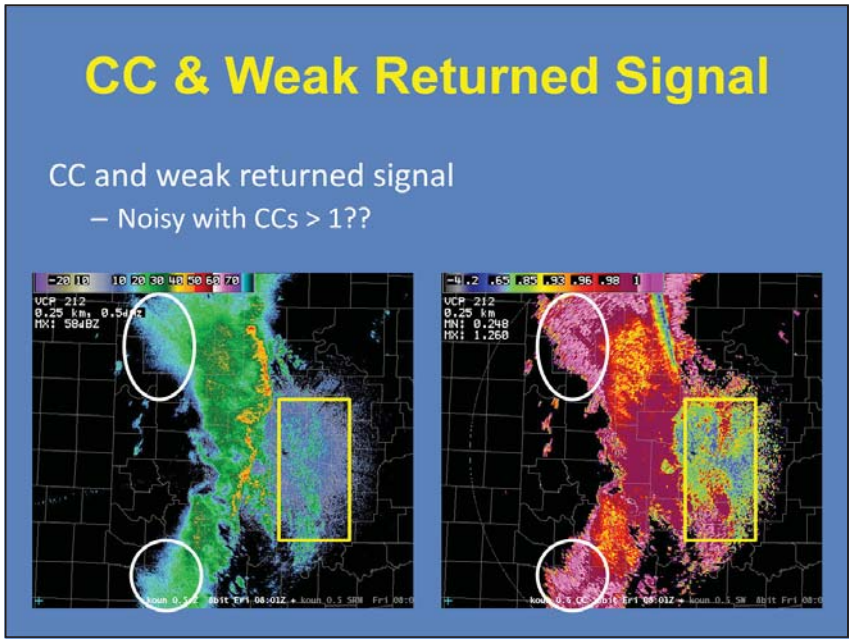
The more diverse the scatterers, the more variation with the cross correlation vectors, and CC gets closer to 0.

I have the word perfect in quotes, because CC is never exactly equal to 0 or 1.



What does CC tell us? Low CC (<0.80) implies low consistency between H and V in the estimate and lot of diversity of the scatterers. In fact, CC <0.80 is so diverse the scatterers are unlikely to be meteorological, such as birds or insects. This distinction between biological and meteorological targets is one of the great benefits of Dual Pol.

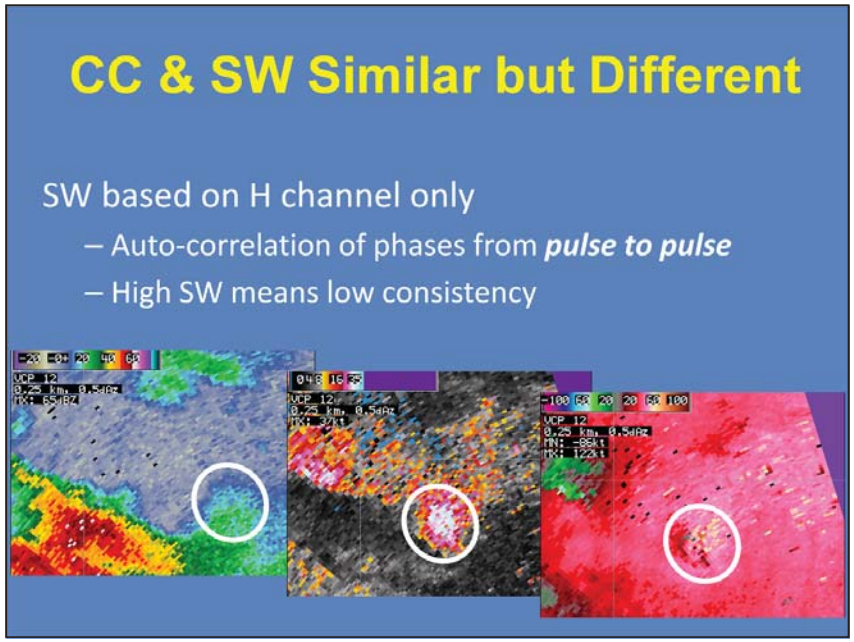
On the other hand, a high CC (>0.97) tells us that the Dual Pol base data estimate is high in consistency between H and V. The scatterers are very uniform in size and shape, such as pure rain or snow.



Correlation Coefficient is an important indicator of Dual Pol data quality. In general, the Dual Pol data will be noisier, and less reliable in weak signal areas than the legacy data. The good news is that CC can help to identify where the Dual Pol data are least reliable.

In areas of weak signal, CC is often noisy in appearance and the magnitudes can vary. Near the radar (boxed area) the CC values are generally low, and there are likely to be non-hydrometeors present. At longer range on the fringe areas of the precipitation (circled areas), the CC values are noisy with values greater than 1. CC > 1 is an estimation artifact, meaning that the estimate is unreliable at that location. It would be misleading to truncate these values at 1, so they are intentionally displayed as > 1.

Though the Z values near the radar and at long range are similar, the CC values tells us more. Z is range normalized, while returned power is not. The weak signal areas at longer ranges and the associated CC values are less reliable than weak signal areas at closer range.

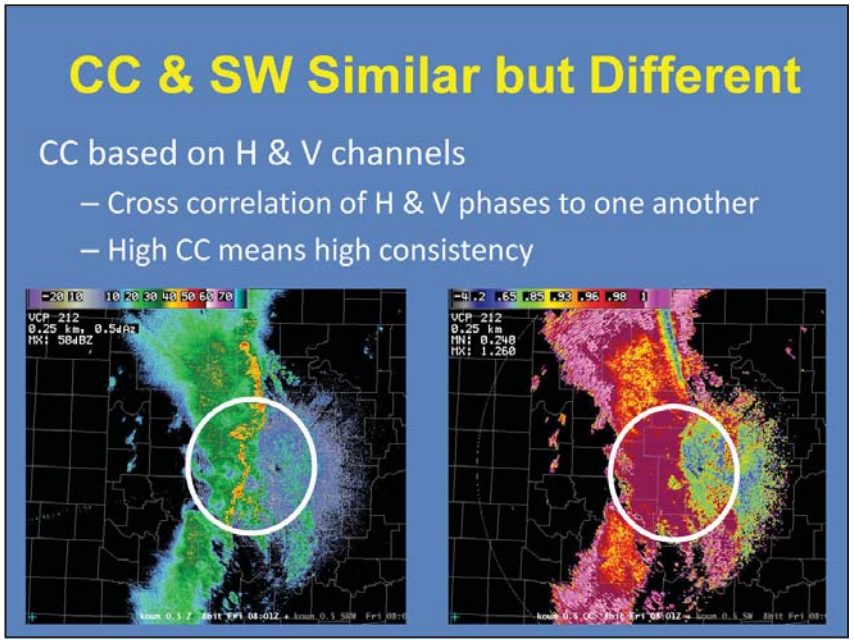


Correlation coefficient and spectrum width are analogous, but there are some important differences. Spectrum width is calculated from the horizontal channel only, and tells us something about the nature of the data from which the base velocity was calculated. Spectrum width is derived from auto-correlation, which compares horizontal channel phase shifts from one pulse to the next.

The greater the variation of these phase shifts, the greater the spectrum width. This means that a high spectrum width implies low consistency as the phase shifts are processed.

The circled area is weak signal close to an intense supercell. There is high spectrum width (middle image) due to both the weak signal and the turbulence. Notice that the velocity field in this same area (right image) is noisy.

A high spectrum width implies a low consistency of pulse to pulse phase shifts. It is an inverse relationship.

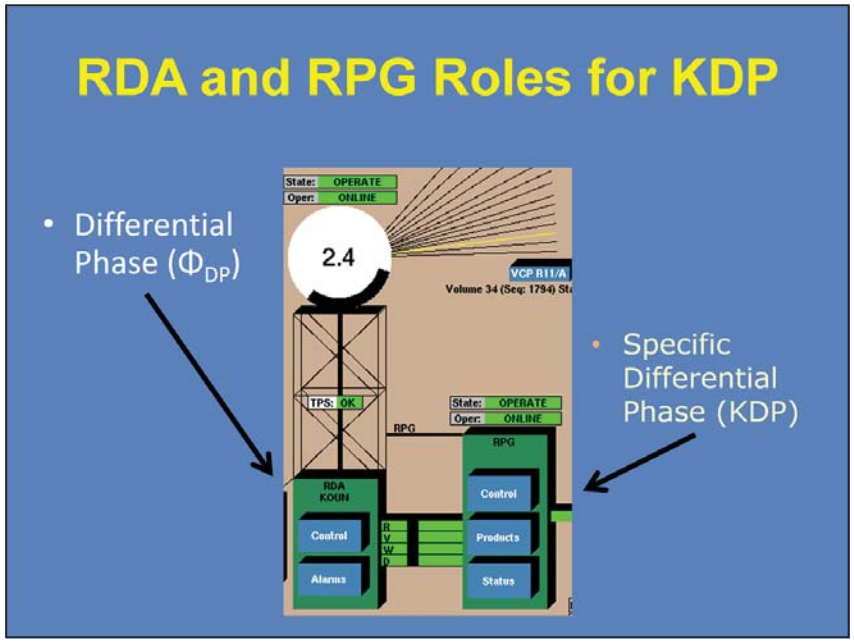


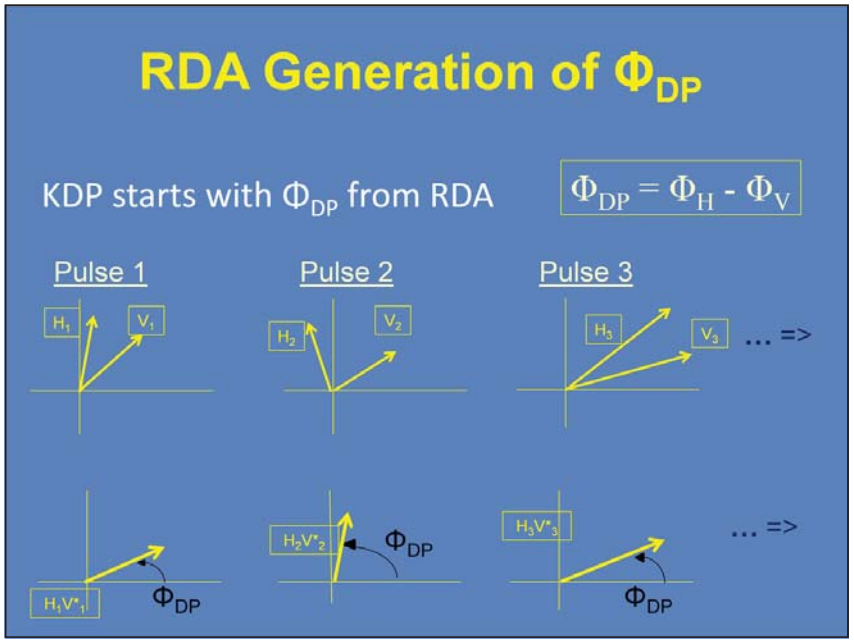
Unlike spectrum width, correlation coefficient is calculated from both the horizontal and vertical channels, and tells us something about the nature of the scatterers that were sampled. CC is derived from cross-correlation, which compares phases from the horizontal and vertical channel pulses to one another.

The greater the variation between the horizontal and vertical channels, the lower the cross-correlation, or CC value. On the other hand, a high correlation coefficient implies lower variation.

The circled area captures both precipitation and clutter near the radar. The correlation coefficient associated with the precipitation is high compared to the clutter. CC is an excellent discriminator between precipitation and non-precipitation echo. You'll see a lot more related to Correlation Coefficient interpretation later in the course.

A high Correlation Coefficient implies a high consistency between the horizontal and vertical channels. It is a direct relationship.



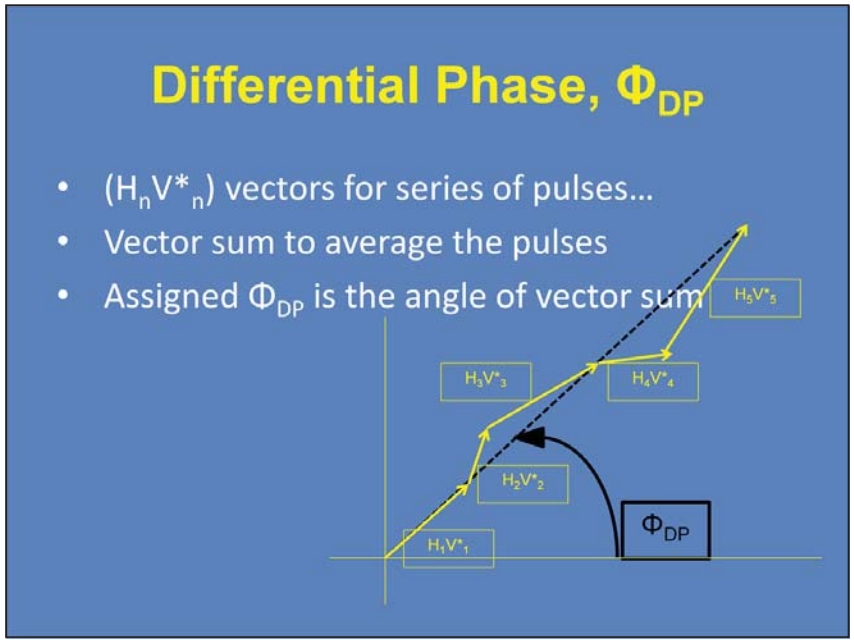


Another Dual Pol base product is Specific Differential Phase, or KDP, though it is technically a derived product generated at the RPG.

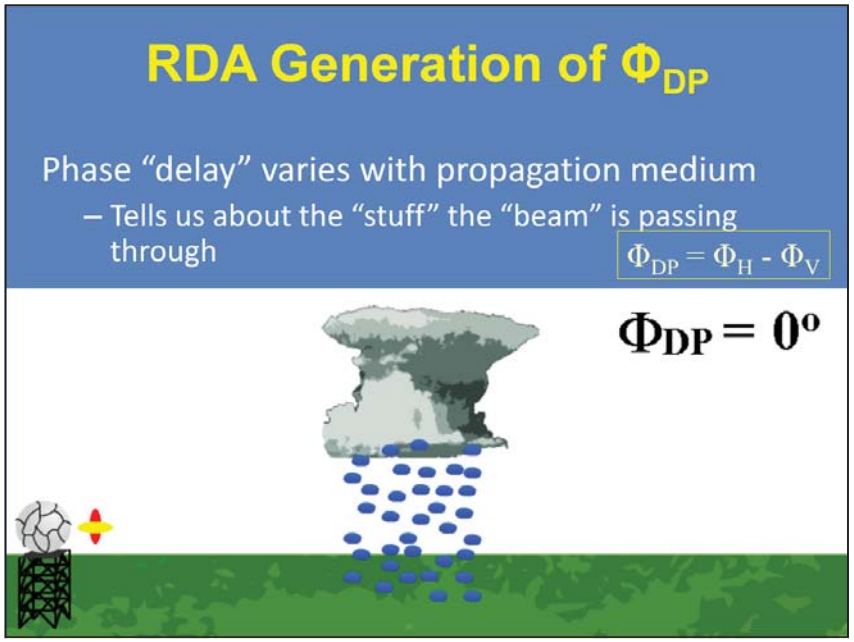
KDP is built from the Differential Phase, or Φ_{DP} base data, which are generated at the RDA and sent to the RPG.

The benefit of KDP is that it tells you something about the type of medium (light rain? heavy rain?) that the beam has passed through.

In order to understand KDP, we must go back to differential phase, Φ_{DP}, which is generated by the RDA signal processor. Recall that for a single pulse, Φ_{DP} is the angle of the cross correlation vector, or the horizontal phase minus the vertical phase. For a series of pulses, Φ_{DP} is the angle of the vector sum of the cross correlation vectors (next slide).

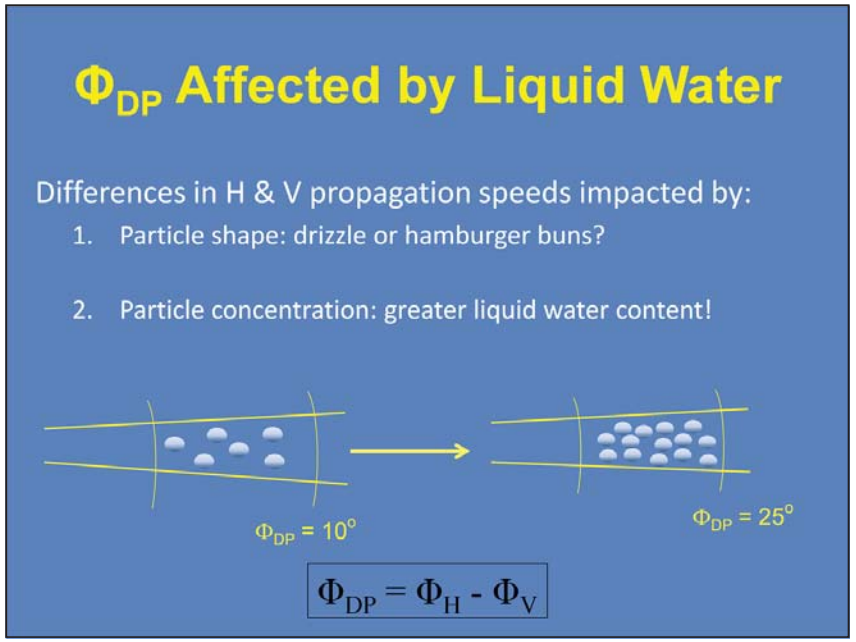


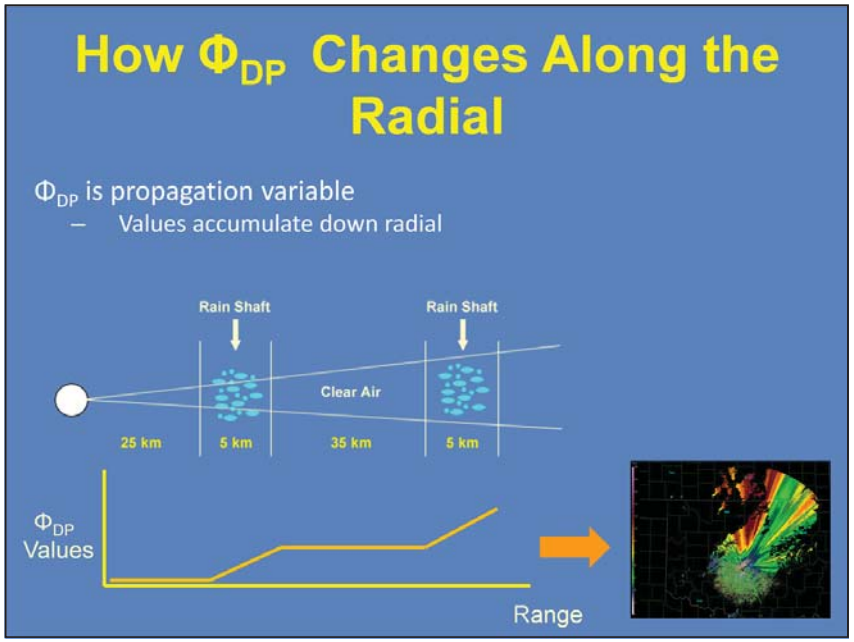
The differential phase, or Φ_{DP} value assigned as base data comes from the vector sum of the single pulse cross correlation vectors. The angle of this vector sum is the assigned Φ_{DP} for that range bin.



As the pulse propagates through different media (light rain, heavy rain, etc.), there is a delay that is apparent in the phase of the returned pulse. Since we have both horizontal and vertical phase, we can compare how the "H delay" differs from the "V delay". This gives us valuable information on the nature of the "stuff" that the radar pulse is passing through.

Liquid water provides "resistance" to the outgoing pulse. In this animation, the pulse is passing through raindrops, which have a larger horizontal extent than vertical. There is more resistance in the horizontal direction compared to vertical, creating a longer delay in the returned H phase compared to the V phase. The returned phase value for H will be greater than for V. This means that Φ_{DP} for that range bin will be positive.





For any given atmospheric volume, the value of Φ_{DP} is affected by differences in propagation speeds of the horizontal and vertical waves. Propagation is slowed by both particle shape and/or by particle concentration.

Here are 2 examples that would result in a slower horizontal propagation compared to the vertical, and thus a higher Φ_{DP}:

1. If the beam is passing through large raindrops (think hamburger buns!), there is more propagation delay in the horizontal direction than in the vertical. That was the example on the previous slide.

2. Assume the same size and shape raindrops in each of two volumes. However, there is a greater concentration of them on the right. This greater concentration, which means more liquid water content, also creates a greater propagation delay in the horizontal direction than in the vertical.

This direct relationship between Φ_{DP} and liquid water content is what makes this Dual Pol variable so valuable.

The value of Φ_{DP} propagates down radial, accumulating with range. There is no way to "reset" Φ_{DP} as the pulse travels outbound, encountering one or more areas of precipitation.

In this super simple example, the beam first passes through clear air, which leaves the Φ_{DP} values unchanged. Then a rain shaft is encountered, which means $\Phi_{DP} > 0$ for a series of range bins, and increases with each bin in the rain shaft. The beam then progresses to another patch of clear air and the Φ_{DP} value stays constant. Finally, another rain shaft is encountered, again increasing the Φ_{DP} value down radial. Throughout this process, the Φ_{DP} value does not "reset" to 0.

Since Φ_{DP} is an angle, the units are degrees, and given enough liquid water, Φ_{DP} will "fold" back to 0°. Φ_{DP} is particularly noisy and subject to data quality problems with the small number of pulses used within many of the WSR-88D VCPs. The possibility of folding, and the general noisiness, make interpretation of Φ_{DP} as a base product challenging.

Why KDP?

- Specific Differential Phase, KDP, easier to interpret
- Range derivative of Φ_{DP}
  - Φ_{DP} change in small chunks along radial
  - KDP units ° per km

$$KDP = \frac{\phi_{DP}(r_2) - \phi_{DP}(r_1)}{2(r_2 - r_1)}$$

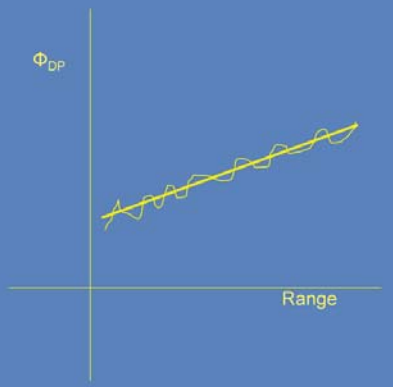




Differential phase, Φ_{DP}, is sent as base data from the RDA. Though both KDP and Φ_{DP} are available in AWIPS, KDP is easier to interpret. KDP is defined as the range derivative of Φ_{DP}. KDP is a way to capture how Φ_{DP} changes over very short ranges, which gives us more useful information. Thus the units for KDP are ° per km.

This equation does not represent the actual calculation of specific differential phase, or KDP. It is used to represent the concept of subtracting differential phase, or Φ_{DP} , over a range interval. The actual calculation involves a least squares fit of multiple differences along the radial, centered at the range bin.

KDP and Z



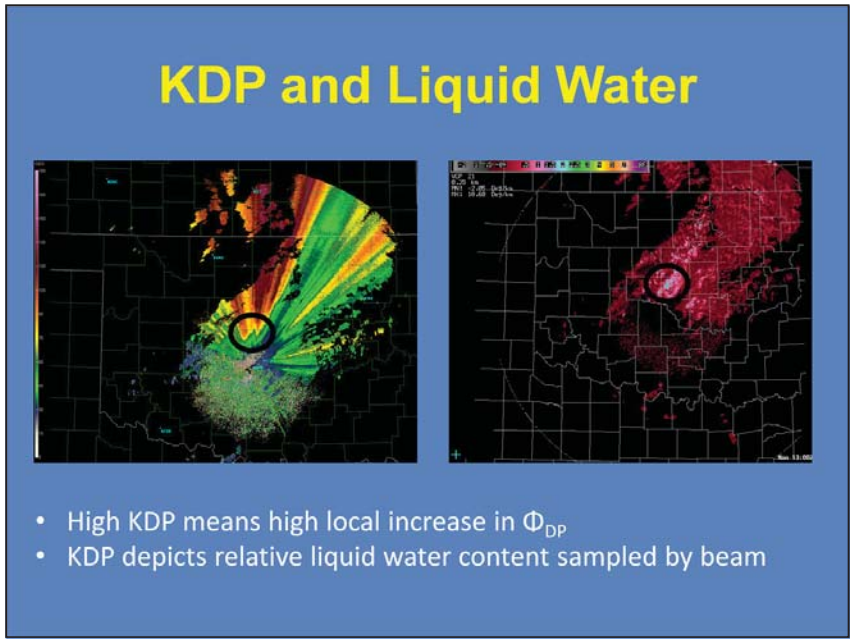
- "Range" depends on Z
  - For Z ≥ 40 dBZ, KDP based on 9 bins
  - For Z < 40 dBZ, KDP based on 25 bins
- KDP captures increase of Φ_{DP} down radial
- Magnitude of KDP directly related to liquid water content

The span of range bins used for the KDP calculation is dependent on the Z value. Higher reflectivity generally corresponds to better data quality, and less noisy Φ_{DP} values. Thus for higher Z, KDP is calculated along a shorter interval of the radial, as compared to lower Z values.

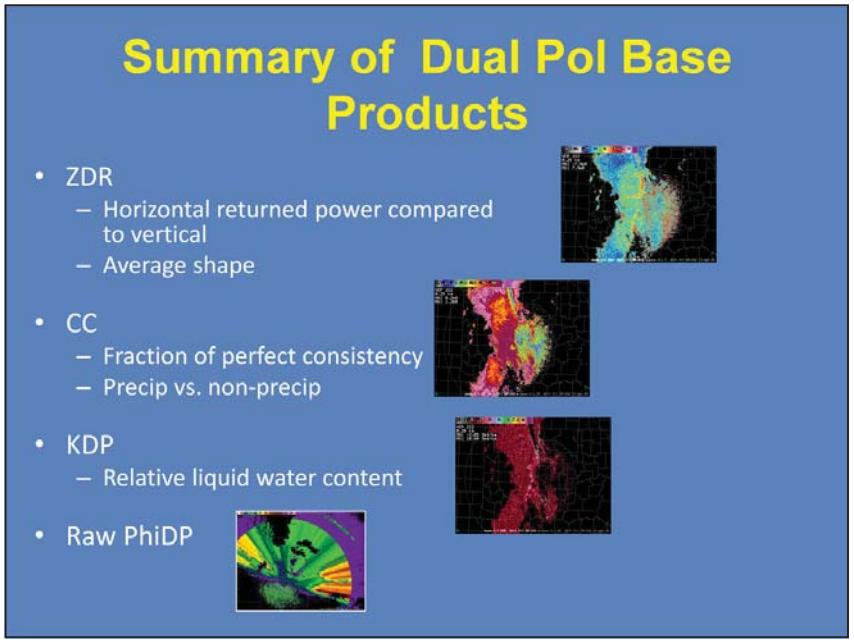
For Z ≥ 40 dBZ, KDP is based on an integration of 9 bins (4 bins back and 4 bins forward along the radial). For these higher Z values, there is less smoothing required, and fewer bins are used.

For Z < 40 dBZ, KDP is based on an integration of 25 bins (12 bins back and 12 bins forward along the radial). For lower Z values, there is the potential for more noise in the data, thus more bins are used for greater smoothing.

Specific differential phase, KDP, is capturing the magnitude of any increase in differential phase, Φ_{DP}, along a radial. The greater the increase, the greater the liquid water content present.



KDP is the last of the Dual Pol variables. It is of value because of its direct relationship to the amount of liquid water content sampled by the beam. On the left is the Φ_{DP} base data displayed in a Level II viewer (GR Analyst). On the right is the associated KDP product displayed on AWIPS. The higher KDP values correspond to the areas of highest Φ_{DP} gradients along the radial.



Differential Reflectivity, ZDR, is computed from the returned power in the horizontal and vertical channels. It does not use returned phase information. ZDR tells us about the average shape of the scatterers.

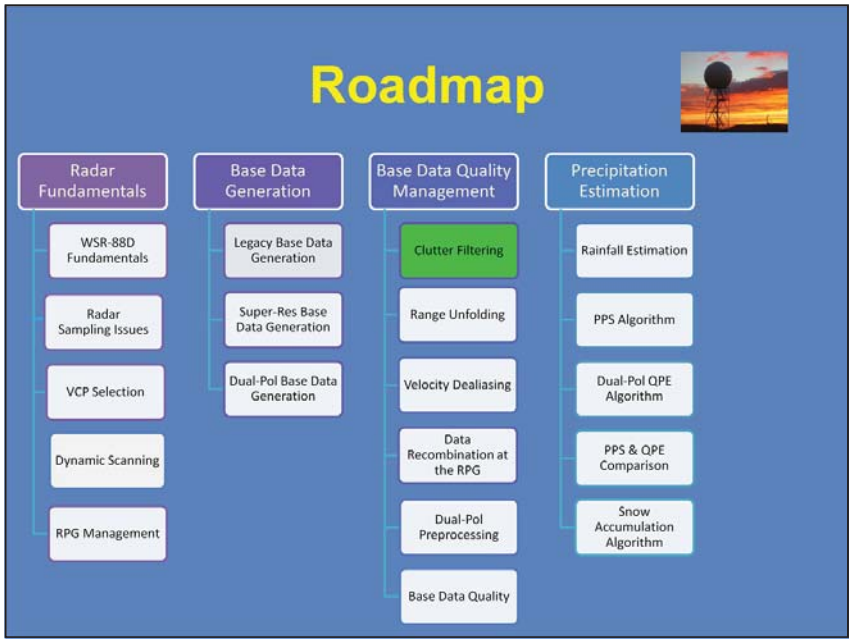
Correlation Coefficient, CC, is based on the returned power and phase, especially how the horizontal and vertical returned phase values compare to one another. CC is a fraction of perfect consistency of these returned phases, which reveals the consistency of the scatterers sampled by the radar. Among many other applications, CC is a highly effective discriminator for precipitation vs. non-precipitation echoes.

Specific Differential Phase, KDP, depicts the relative liquid water content of the volume sampled by the beam.

There is a Φ_{DP} product available, Raw PhiDP. It is the unprocessed, level II Φ_{DP} base data from the RDA. It's use is limited, and it will be discussed in the Products section of RAC.



Welcome to Clutter Filtering.



Here is the “roadmap” with your current location.

Clutter Filtering Objectives

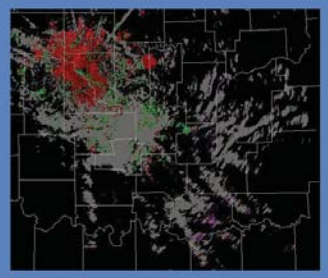


1. Identify the purpose, strengths and limitations of the following clutter suppression algorithms:
   a) Clutter Mitigation Decision (CMD)
   b) Gaussian Model Adaptive Processing (GMAP)
2. Identify examples of moving ground-based targets that cannot be identified by CMD.

Here are the two objectives for Clutter Filtering, which will be taught in sequence during this module.

Ground Clutter Contamination

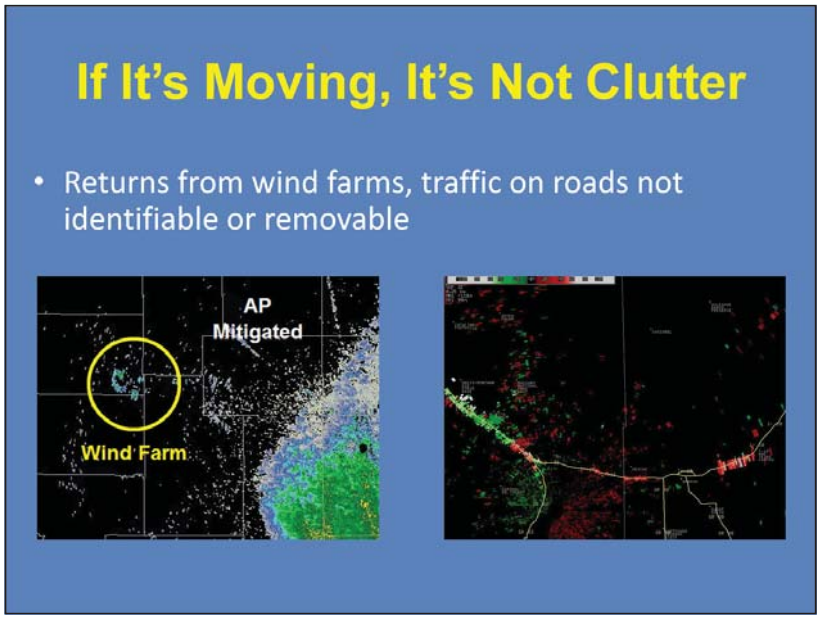




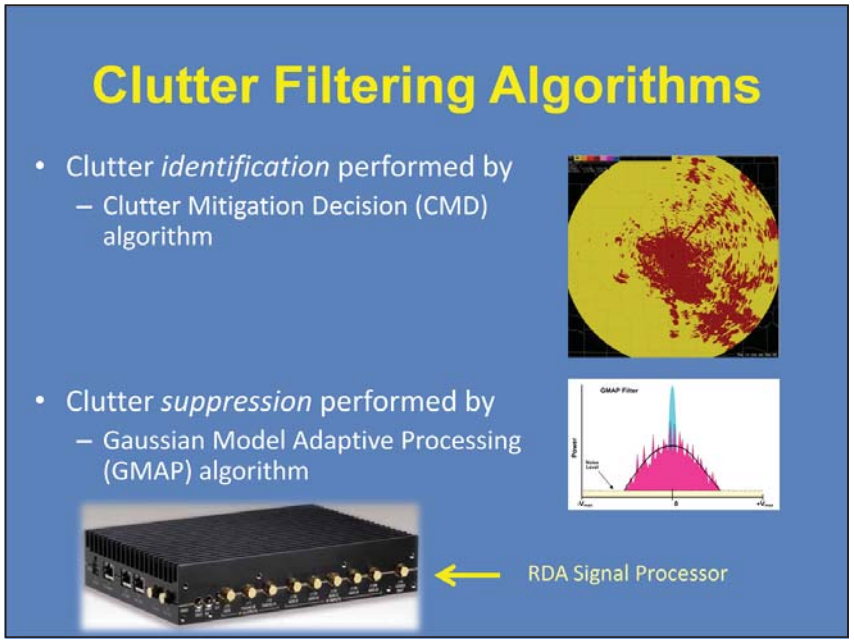
- Returns from non-moving ground targets
  - Affects ***all*** radar products
- Two types of clutter contamination
  - Normal ground clutter
  - Anomalous propagation clutter

In general, ground clutter on WSR-88D radar products is return from stationary or nearly stationary ground targets that has not been filtered. Clutter suppression is applied at the signal processor just before the base data are built. Since all the base and derived products are built from the base data, unfiltered clutter will negatively impact all the products. There are two types of clutter contamination. The first is normal ground clutter, meaning features that are present all the time, such as terrain, buildings, etc. The second type is transient and dependent on beam propagation, known as anomalous propagation clutter.

For the reflectivity (left) and velocity (right) images, clutter filtering is only applied very close to the radar. There is extensive Anomalous Propagation (AP) clutter contamination to the east through the south of the radar. Note the near zero velocities throughout the AP clutter areas.



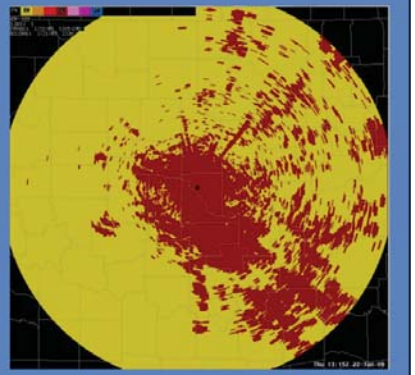
The WSR-88D clutter algorithms are designed to detect ground clutter, which means near zero velocity and spectrum width. There will still be contamination from moving ground targets such as wind farms and traffic on highways.



There are two different algorithms in the clutter filtering process. The data are first processed by the Clutter Mitigation Decision (CMD) algorithm, which does the job of identifying clutter on a bin by bin basis. For each bin identified by CMD, the Gaussian Model Adaptive Processing (GMAP) algorithm then applies the signal reduction, or suppression, of the clutter. Both CMD and GMAP are run at the RDA signal processor (the black box!)

CMD in a Nutshell

- CMD identifies both normal & AP clutter every volume scan
  - Builds "Dynamic Bypass Map" every rotation
- CMD maps used for bin by bin signal removal

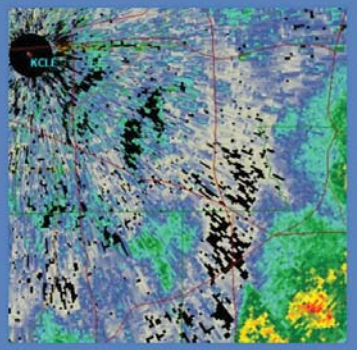


Clutter Mitigation Decision (CMD) offers an automated approach to the management of clutter filtering. CMD can identify both normal and AP clutter every volume scan, which eliminates the need for manually defining and downloading regions files to address AP clutter. CMD builds a dynamic Bypass Map that shows the bins that contain clutter. These maps are then used for a bin by bin signal removal. Suppression of clutter is only performed on the those bins identified by CMD.

The image on the right is a visualization of a CMD generated Bypass Map. It is called the Clutter Filter Control (CFC) product. Each of the red bins has been identified by CMD as containing clutter. The yellow bins do not contain clutter, based on CMD's analysis.

CMD Inputs

- Fuzzy logic with many inputs
- Z texture
  - Smooth (low) = weather
  - Rough (high) = clutter
- Z spin
  - Z gradient sign changes
- CPA (phase consistency)
  - High CPA = clutter
  - Low CPA = weather



CMD is a fuzzy logic algorithm with multiple inputs. It uses nearly all the base data in its most primitive form to determine where clutter exists.

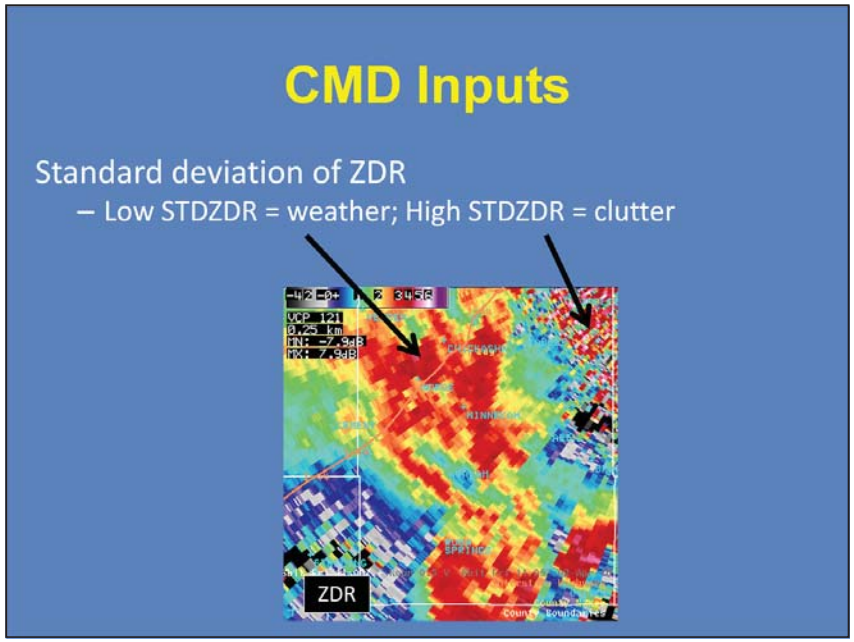
We start with the inputs based on Reflectivity, Z.

1. Z texture, which is smooth (a low value) for a meteorological target and rough (a high value) for a clutter target

2. Z spin, which looks at the number of times that the Z gradient changes sign along a portion of the radial. The higher the number of sign changes, the higher the likelihood of clutter.

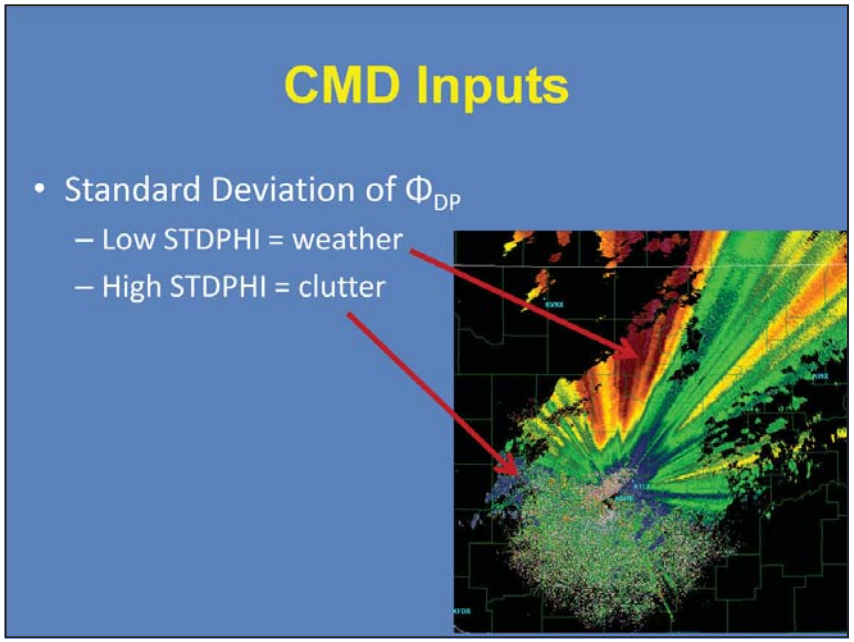
In the associated Z image, the weak signal areas close to the radar have a visually rougher texture and one would find more frequent changes in sign of the Z gradient.

3. Clutter Phase Alignment (CPA) captures the variance of pulse to pulse phase changes. Since CMD is looking for clutter targets that don't move, there is low variation in returned pulse phase changes with clutter. A high CPA means good alignment of phasors, a low variation of returned phase, and a higher likelihood of clutter.



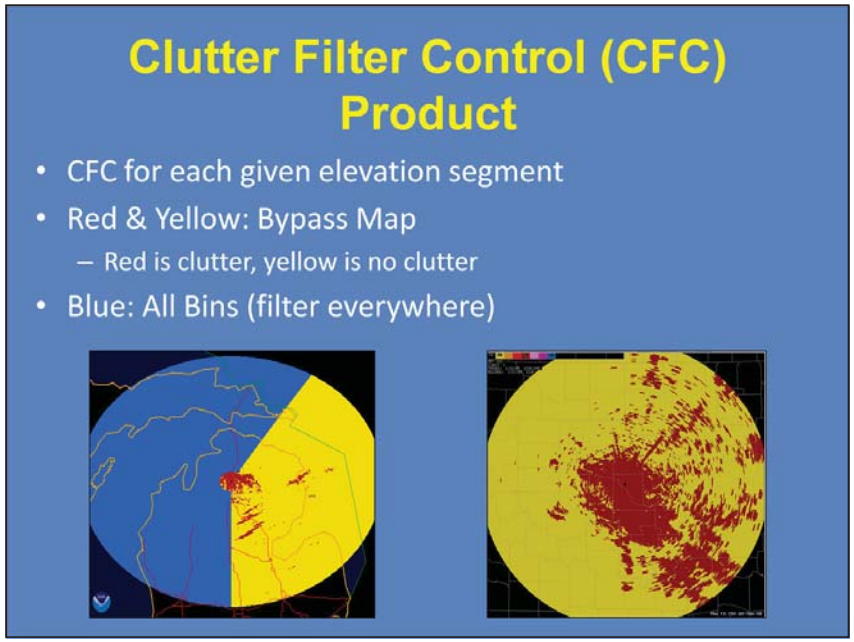
There are two inputs based on Dual-Pol base data. They are standard deviations of ZDR and Φ_{DP}.

The higher the standard deviation of ZDR, the higher the likelihood that the range bin contains clutter. In this example, you can likely tell by looking at the smoothness of the ZDR data where the clutter vs. the weather is located. That variation (or lack of it) is captured in the standard deviation data. The higher the standard deviation of ZDR, the more likely it is that the bin contains clutter.



The second CMD input based on Dual-Pol base data is the standard deviation of Φ_{DP}.

The higher the standard deviation of Φ_{DP}, the higher the likelihood that the range bin contains clutter. In this example, compare the noisiness of the Φ_{DP} data surrounding the radar to the areas of precipitation to the north through the east (trust me on the precipitation part). That variation (or lack of it) is captured in the standard deviation data. The higher the standard deviation of Φ_{DP}, the more likely it is that the bin contains clutter.

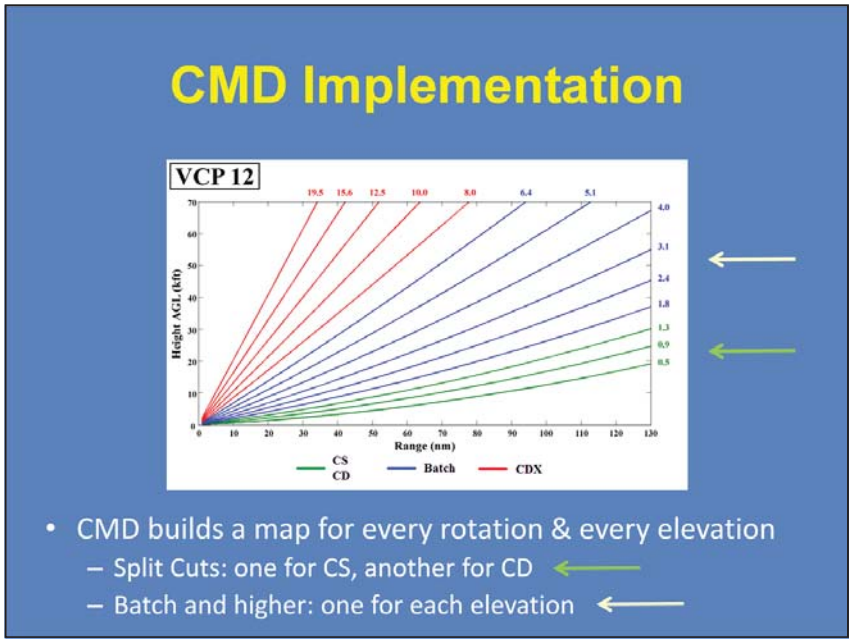


For each elevation segment, the CFC displays the type of clutter filtering that has been invoked for the lowest elevation in that segment. The area with a red/yellow combination are where the Bypass Map is in control. That means there is bin by bin identification of clutter, with red for clutter (to be filtered) and yellow for no clutter (to be left alone). The blue area defines All Bins filtering, which means every bin is going to filtered.

The CFC on the left hand side shows a clutter filtering scheme that has All Bins (blue) in control over a large area from south through northeast. Outside of the All Bins area, the Bypass Map is on control, with red bins identified as containing clutter and yellow bins identified as not containing clutter.

The CFC on the right hand side shows a clutter filtering scheme that has the Bypass Map in control everywhere.

These varying clutter filtering schemes are implemented manually by downloading clutter regions files from the RPG to the RDA.



CMD is active for every rotation and every elevation, building a bypass map each time. For the Split Cuts, there are two rotations at the same elevation, Contiguous Surveillance (CS), then Contiguous Doppler (CD). CMD builds a different map for each of these two rotations. For the remaining elevations in the volume scan, CMD builds a new map each time.

There is not a Clutter Filter Control (CFC) product build for every elevation, thus you cannot see all of the maps built by CMD. This may be challenging for CMD troubleshooting efforts.

CMD Implementation

- Keep Bypass Map in control all the time
  - Default clutter regions file
- All Bins suppression rarely needed
  - Does not increase amount of power removed by GMAP

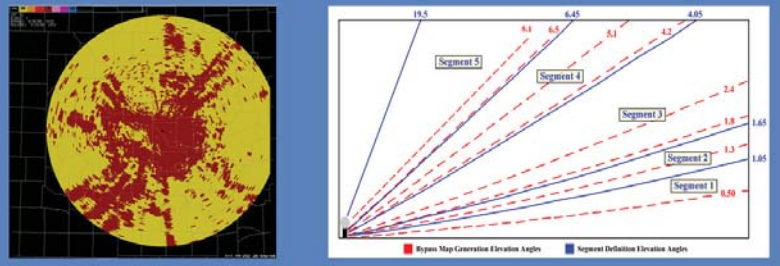


CMD offers “hands off” clutter suppression, though some data quality monitoring may be needed from time to time. In order for CMD to build maps for every elevation, the Bypass Map must be in control for the entire display for all elevations. The “Default” clutter regions file is designed to do just that. It may be necessary (as in this example) to create a local version of this file, “SUPPRESSION ONLY W/ BYPASS MAP”. Whether it is the Default file, or a locally defined version, the key point is to have the Bypass Map in control for the entire display for all elevations.

With CMD active, All Bins suppressions is rarely needed. All Bins only defines where suppression occurs. It has no affect on the amount of power removed by GMAP when the suppression is performed.

CMD Summary

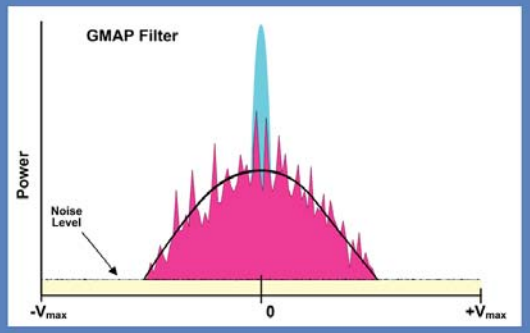
- CMD identifies clutter
  - Builds map for every rotation & every elevation
  - Performance affected by pulses/radial
- One CFC per elevation segment



In summary, CMD’s job is to identify the bins that contain clutter, both normal and AP clutter. CMD performs this identification and builds a clutter map for every rotation and every elevation of every VCP (that’s a lot of maps). The Clutter Filter Control product is available to help visualize where CMD has identified clutter. However, CFC products are limited to one map for each of the elevation segments on the graphic on the right, only showing a subset of the maps actually built by CMD.

The number of pulses per radial impacts CMD’s performance, especially with respect to discriminating clutter from weather with little movement, like stratiform rain. The faster VCPs, 12, 212, and 121, have the lowest number of pulses per radial. The impact of faster VCPs on data quality will be explored later.

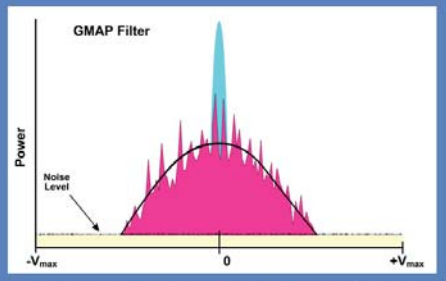
Clutter Filtering Algorithms



- CMD has ***identified*** clutter
- For the bins identified by CMD, filtering performed by Gaussian Model Adaptive Processing (GMAP) algorithm

Now that clutter have been identified by CMD, it is time to perform the filtering, or removal of the power from the clutter portion of the returned signal. The Gaussian Model Adaptive Processing (GMAP) algorithm applies filtering only to those bins identified by CMD.

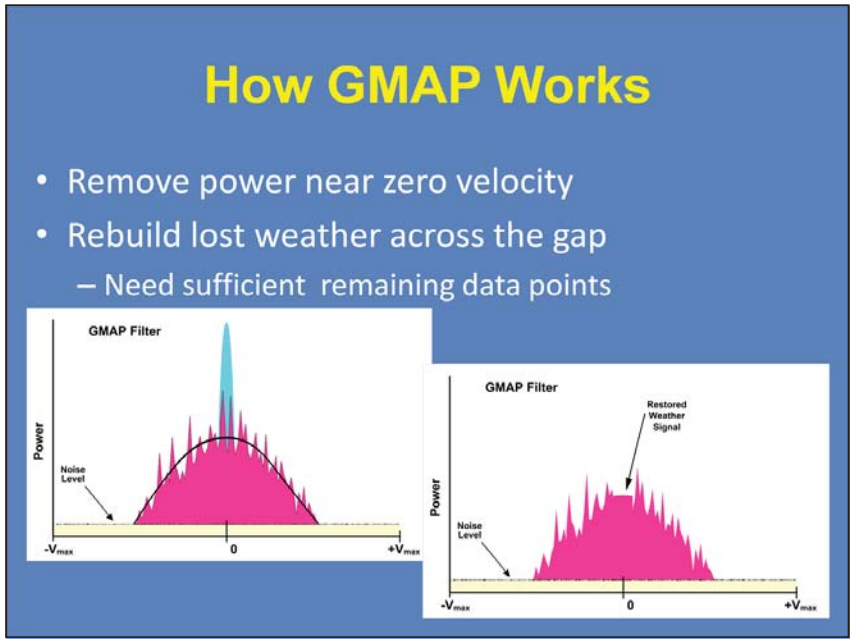
How GMAP Works



- Weather and clutter signals look different
  - Can be approximated by Gaussian curve
- Remove power from narrow spike near zero velocity
- Once power removed, GMAP attempts to rebuild lost weather signal

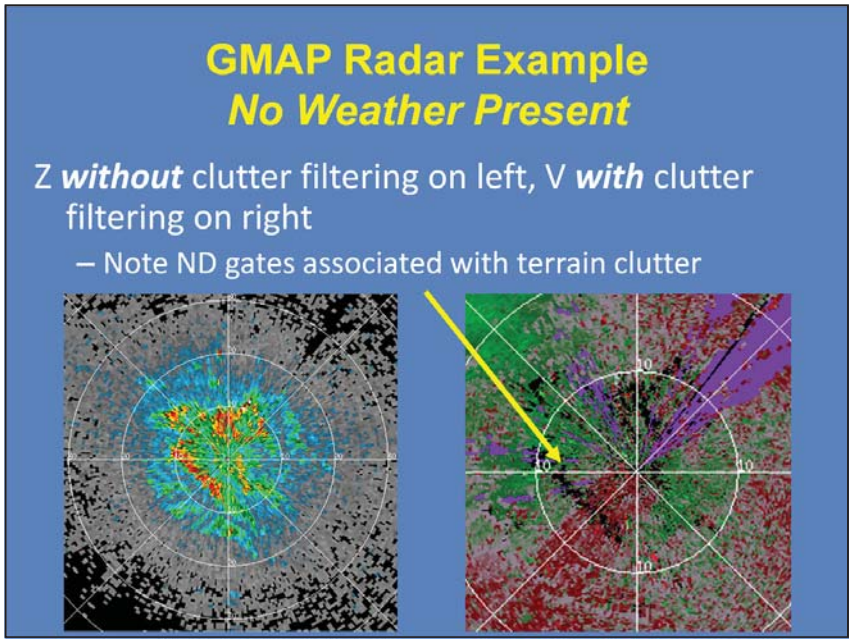
Weather and clutter signals have different characteristics. A clutter signal (blue green spike) has high power, is centered at zero velocity and has a narrow spectrum width. A weather signal (broader pink blob) will have varying power, velocity and spectrum width. This difference between clutter signals and weather signals can be used to filter out the clutter signal with minimal damage to the weather signal. Another aspect that is part of GMAP's design is that both clutter and weather signals can be well represented by Gaussian curves. This is a super simple graphic. For example, weather signals are not usually centered at zero velocity.

GMAP first removes power from the signal near zero, hopefully as much of the "spike" as possible. If there is enough of the weather signal remaining, GMAP can rebuild it.

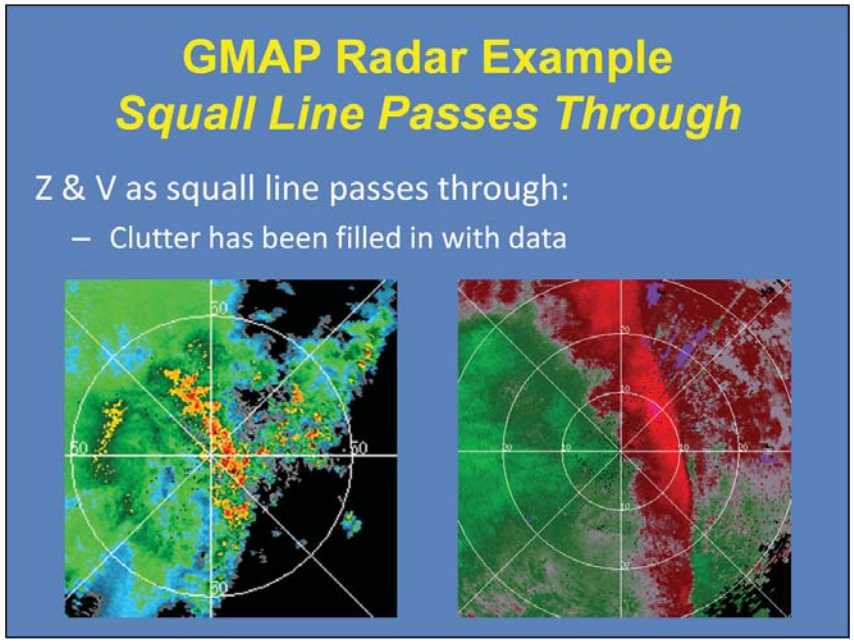


Once the width for signal removal is determined, filtering is applied to all of the signal within that width. In this case, both the clutter and weather signal within this interval will be removed. However, GMAP has the ability to rebuild the weather signal that was lost. This is dependent on the availability of data points outside of the gap. If there is sufficient weather signal data outside the gap to be represented as a Gaussian curve, GMAP can rebuild the weather signal across the gap using the Gaussian estimate. To see an animation of this process, click on the beginning image on the left hand side.

The number of pulses per radial impacts GMAP's performance, especially with respect to the rebuilding of the signal. The faster VCPs, 12, 212, and 121, have the lowest number of pulses per radial. The impact of faster VCPs on data quality will be explored later.

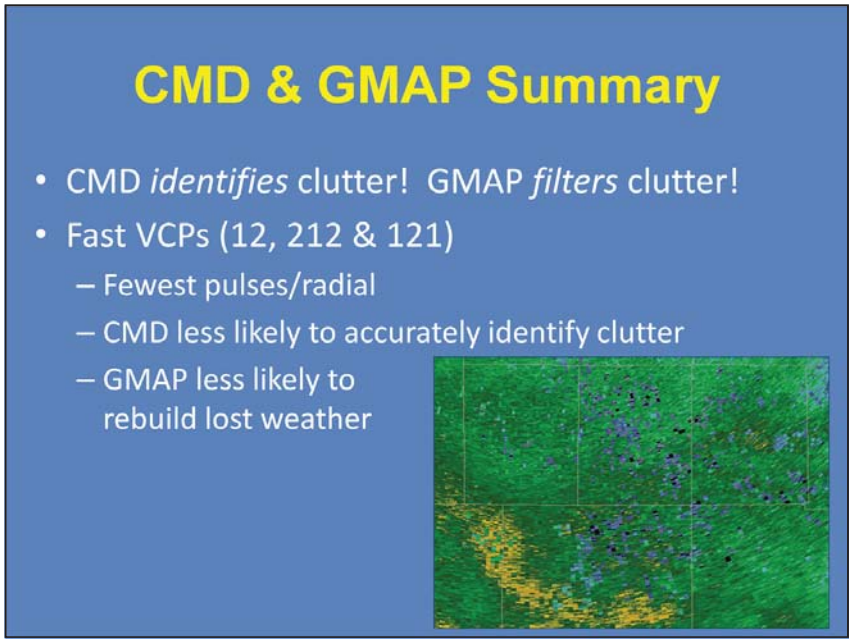


As an example of GMAP performance, we start with no weather with either of these images. On the reflectivity on the left, clutter filtering has been turned off to identify the local terrain clutter. Of particular importance is the ridge line to the southwest. On the velocity on the right, clutter filtering has been applied. There is also a second step known as clutter censoring, which attempts to remove additional signal for bins with only clutter in them. Clutter filtering and censoring have produced the bins with no data on the velocity product. Again, the ridge to the southwest is apparent.



At a later time, a squall line has passed through the area. The reflectivity on the left and the velocity product on the right show the squall line. Compared to the reflectivity, the velocity product has been zoomed in.

Note that in both products the ridge and other terrain clutter areas are no longer apparent. The weather signal is strong enough, and there are enough pulses available, for GMAP to rebuild the weather for the bins that contain clutter.



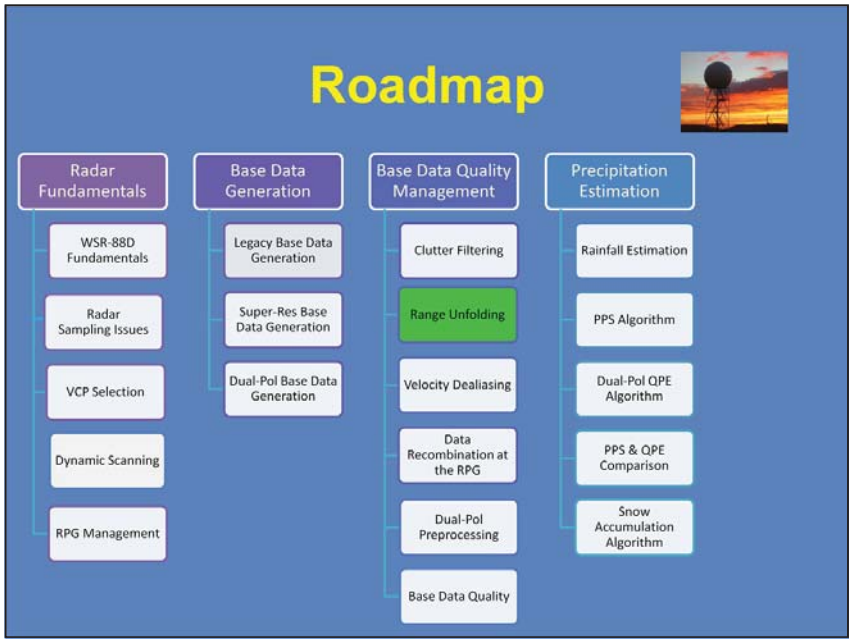
In summary, Clutter Mitigation Decision (CMD) and Gaussian Model Adaptive Processing (GMAP), work together to suppress clutter. CMD first identifies clutter, both normal and AP, on a bin by bin basis. CMD builds maps for each elevation and rotation showing which bins contain clutter and which do not. GMAP then applies filtering only to the bins identified by CMD. GMAP removes power from the signal near zero velocity, then rebuilds the weather signal “across the gap” if there are sufficient pulses remaining.

The performance of both CMD and GMAP are impacted by the number of pulses per radial. The faster VCPs, 12, 212, and 121, have the lowest number of pulses per radial. For these VCPs, it is more difficult for CMD to discriminate clutter from weather with little movement, like stratiform rain. For these VCPs, it is more difficult for GMAP to rebuild the weather signal after the clutter has been removed.

This reflectivity product is a stratiform rain event with VCP 12 invoked, even though VCP 12 is designed for more intense convective precipitation. The low number of pulses per radial with VCP 12 in stratiform rain makes it harder for CMD and GMAP, resulting in numerous gates of “lost” data.



Welcome to Range Unfolding.



Here is the “roadmap” with your current location.

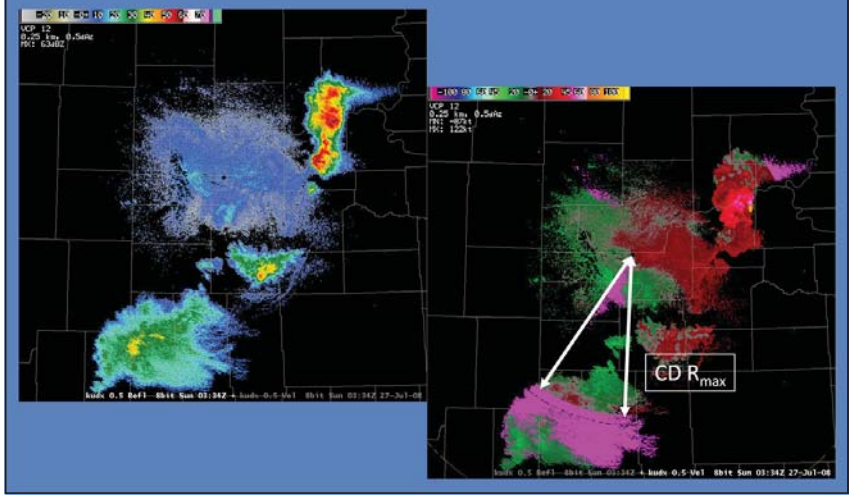
Objectives



1. Identify the purpose, strengths and limitations of the following range unfolding algorithms:
   a) Legacy Range Unfolding
   b) SZ-2 Range Unfolding

There is one objective for Range Unfolding, which will be taught in sequence during this module.

Now For RF on V & SW



Now for range unfolding of velocity and spectrum width base data. Recall that high PRFs are needed for good velocity estimates, but a high PRF also means a short R_{max}. It is necessary to “unfold” the velocity data to its appropriate range.

The next several slides will help explain how the purple “RF” data area assigned on the velocity and spectrum width products. Notice the discontinuity in the velocity data at a fixed range. You will often see this on velocity and spectrum width as long as there is sufficient areal coverage of echo. This range is the R_{max} for the Contiguous Doppler/high PRF mode that is used for velocity data collection. You are seeing the end of the first trip and the beginning of the second trip. The question is “how did any velocity data end up being assigned beyond the first trip?”

Range Folding Rarely on Reflectivity

- Multiple tripping rare on Z products
  - Low PRFs; R_{max} typically ~250 nm
- No RF assigned to Z products
  - Purple means very high dBZ!



It's important to remember that when the reflectivity data are collected, low PRFs are used, resulting in long R_{max} values. For the lowest elevation angles, the reflectivity data were collected with an Rmax ≈ 250 nm, so all the data are within the first trip.

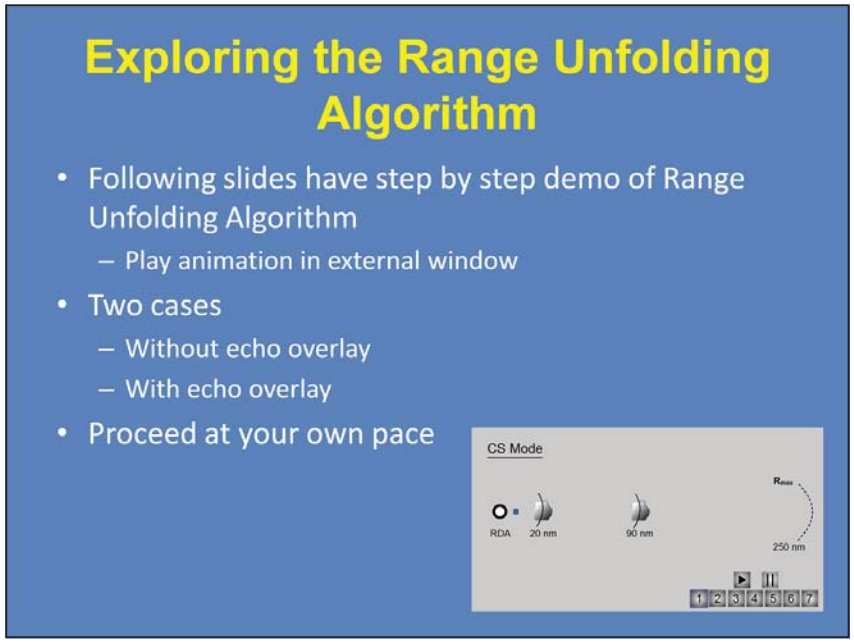
The color purple on reflectivity products indicates very high dBZ values, not RF.

Legacy Range Unfolding Algorithm

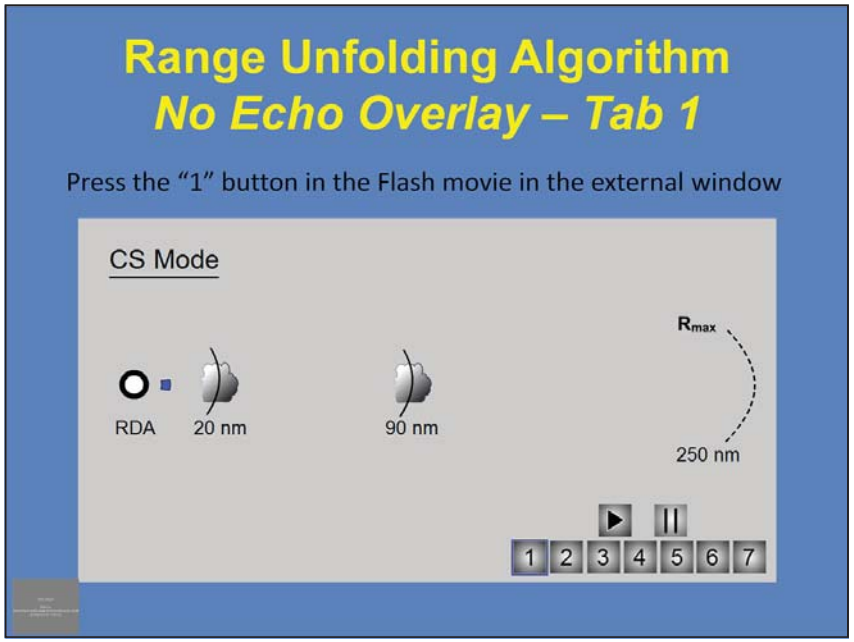
- In action since WSR-88D original deployment
- Assigns V & SW to proper range
  - May be beyond CD R_{max}
- Run at RDA signal processor for all VCPs except 121



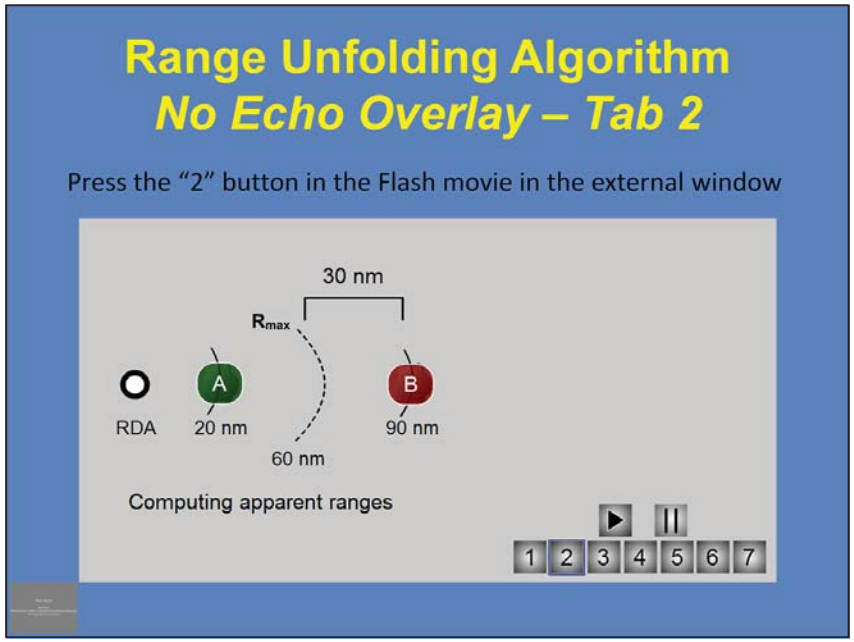
The first Range Unfolding algorithm to present will be known as the legacy Range Unfolding Algorithm. That is because this technique has been in use since the WSR-88D was originally deployed. The goal of the Range Unfolding algorithm is to assign velocity and spectrum width data to its appropriate range. This may be in the first trip, the second trip, and in rare cases, the third trip. The Legacy Range Unfolding algorithm is run at the RDA signal processor, which is literally a black box!



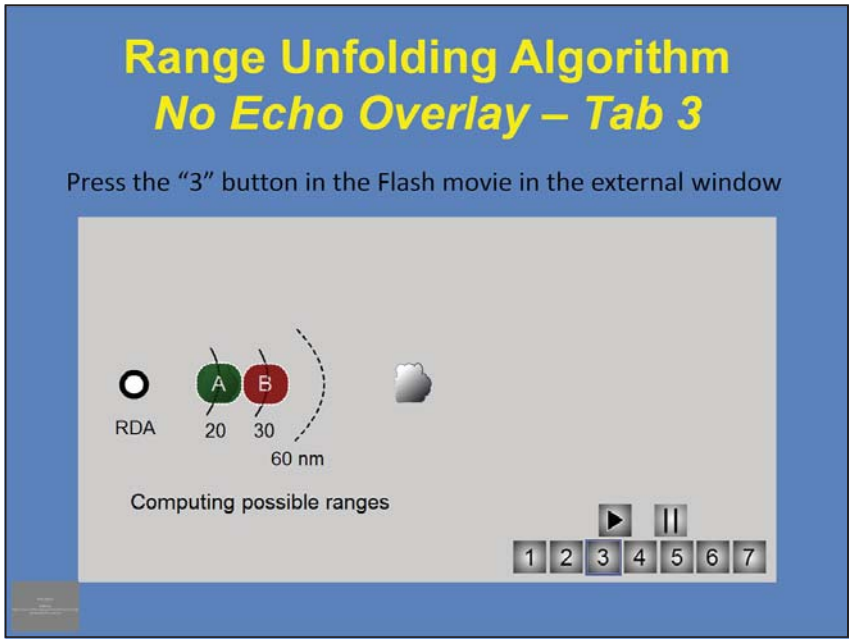
The following slides give you a step by step demonstration of how the Range Unfolding algorithm works. Each slide has an animation in an external window that you can play over and over again, at your own pace. You'll see two cases. The first one has no echo overlay. The second case involves no echo overlay.



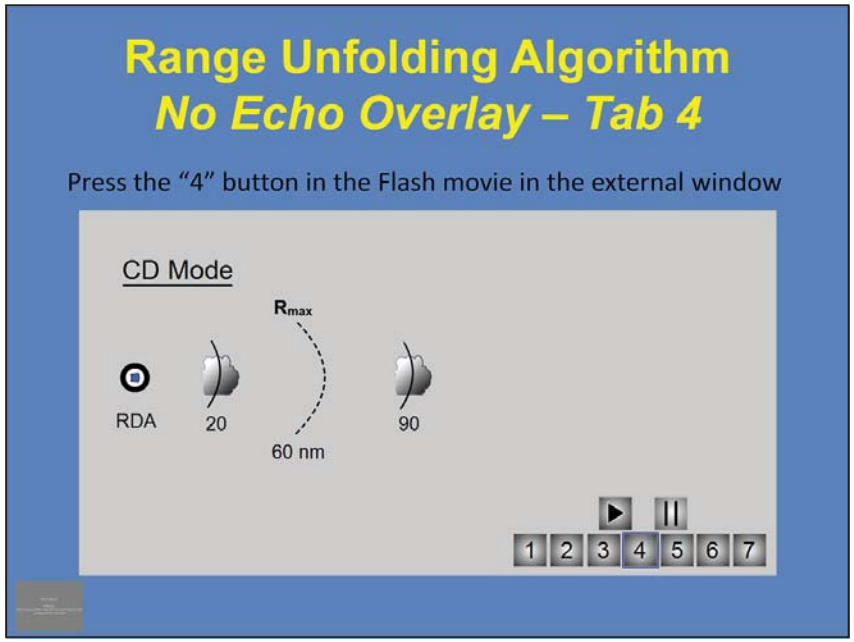
Tab 1. We are looking down a single radial in Contiguous Surveillance mode. That means we have a nice long unambiguous range. We have two targets. The first one is at 20 nm, the second one is at 90 nm. Since both of these targets are within the first trip, we know their true range and we know their returned power.



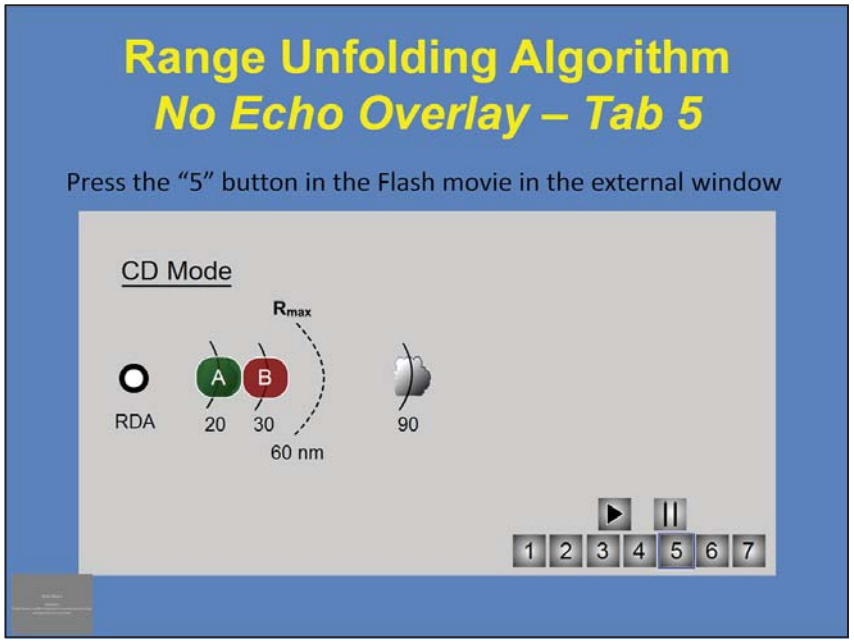
Tab 2. In anticipation of going to Doppler mode in the next rotation, the algorithm is able to compute what the apparent ranges of these two targets will be, and in this case, we're going to have an Rmax of 60 nm when we go into Doppler mode. So that means target A will stay at 20 nm, target B will be folded into an apparent range of 30 nm.



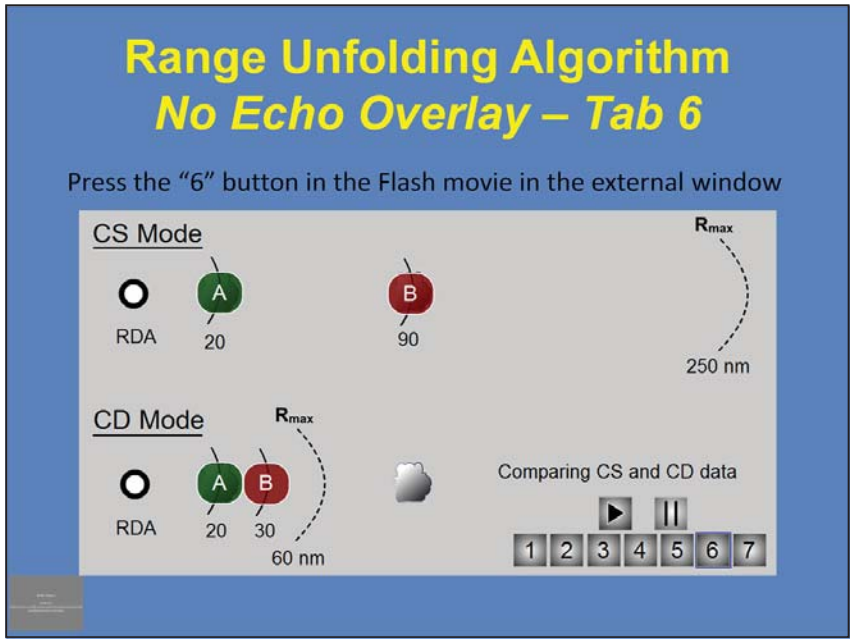
Tab 3. We know that target B is going to have an apparent range of 30 nm. At this step, we're going to compute all the possible ranges for all of the targets that show up in the first trip, and these possible range go out into the second, third, and even fourth trips.



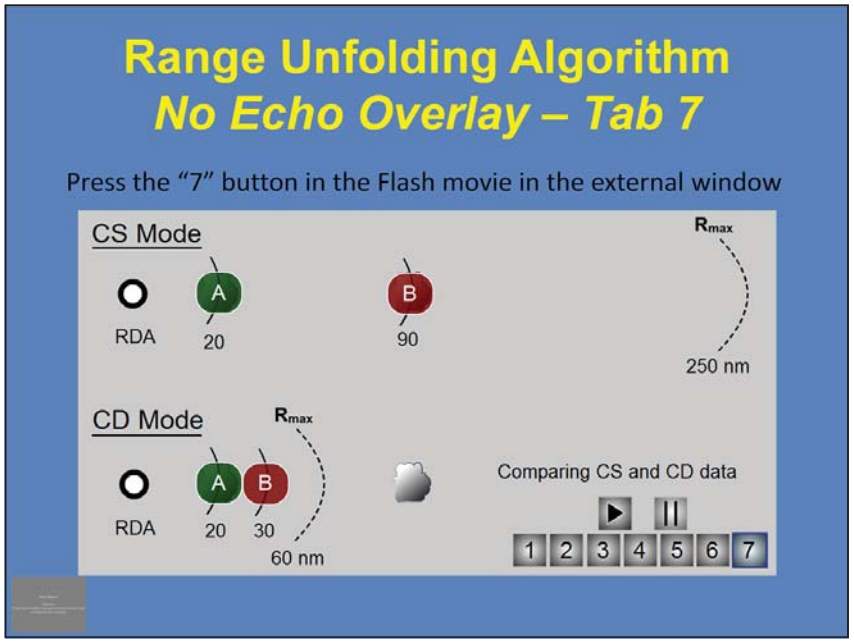
Tab 4. Now we're actually collecting data in Doppler mode, and the energy from pulse 1 for target B which was at 90 nm is on its way back, but it does not arrive before the energy from pulse 2 is transmitted. That is how we get an apparent range for target B at 30 nm.



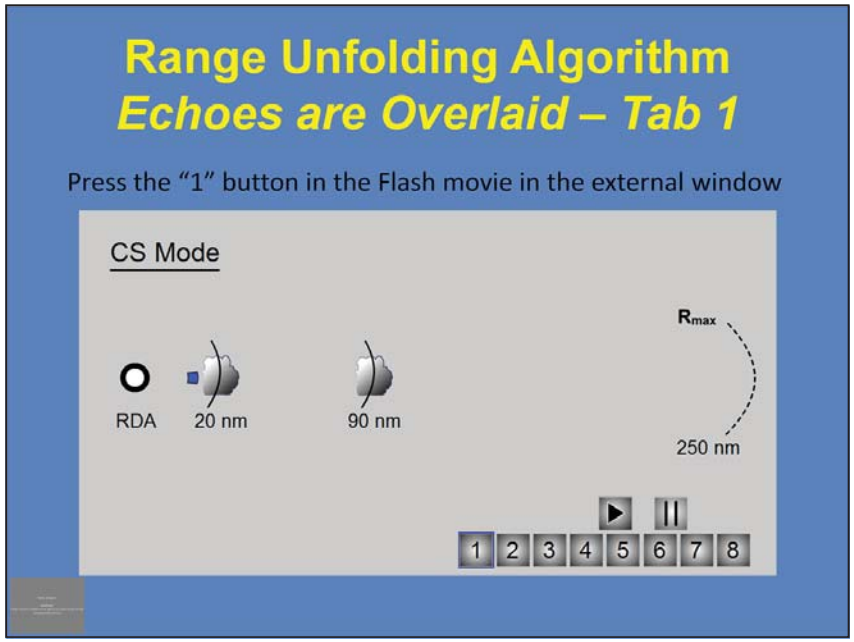
Tab 5. Now we begin the actual comparison down radial of the Surveillance data vs. the Doppler data, checking to see along the radial wherever there is a target to see if there was something at that same range in Surveillance mode. So when we get to target A at a range of 20 nm, we ask the question "was the something at 20 nm in the Surveillance data?" And the answer is Yes, so we know the velocity for target A belongs at the range of 20 nm.



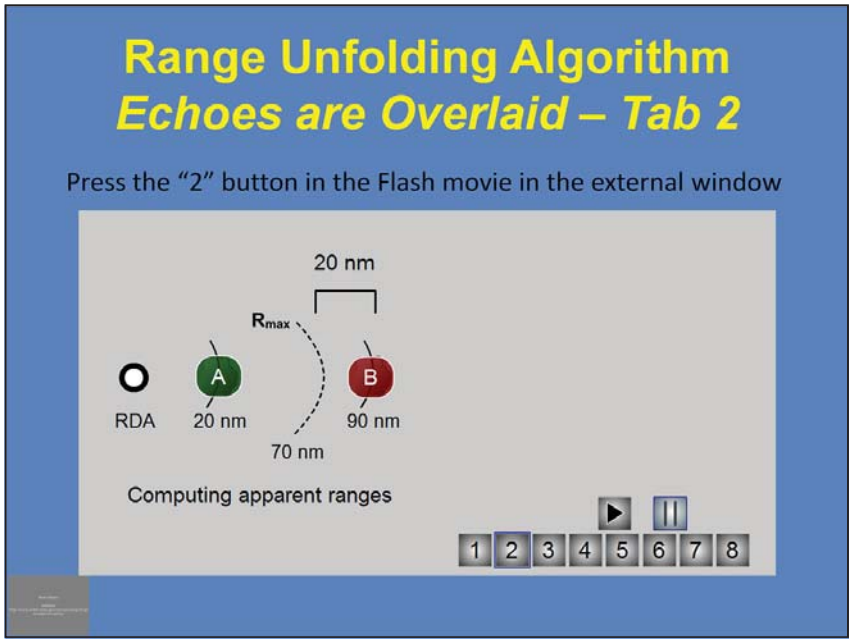
Tab 6. Now we find a target at an apparent range of 30 nm, compare that to the Surveillance data to ask the question "was there something at 30 nm?" No, there wasn't.



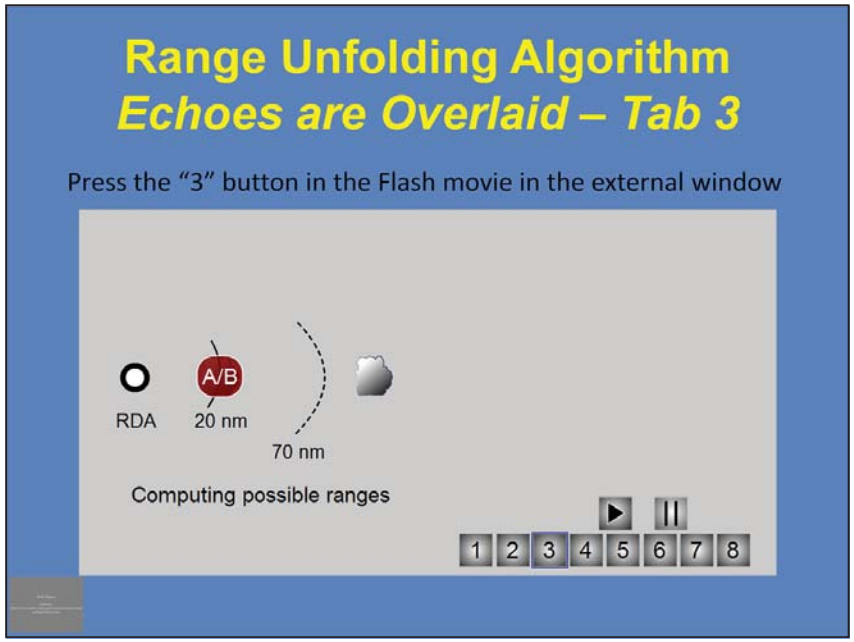
Tab 7. Since we don't have anything at 30 nm, the next step is to check the next possible range which is 90 nm, and yes, there was something in the Surveillance data at that range. That's how we know that the velocity value that appears to be coming from 30 nm actually belongs at 90 nm.



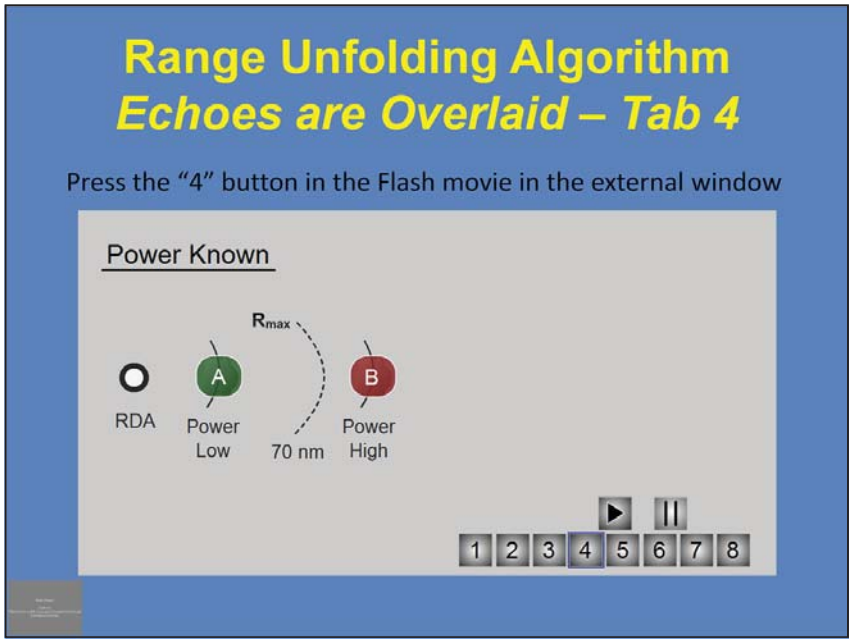
Tab 1. We start off with the same initial conditions. We're in Surveillance mode, we have a target at 20 nm and another target at 90 nm, maximum unambiguous range is 250 nm, so we have returned power and we have range for each of these returned targets.



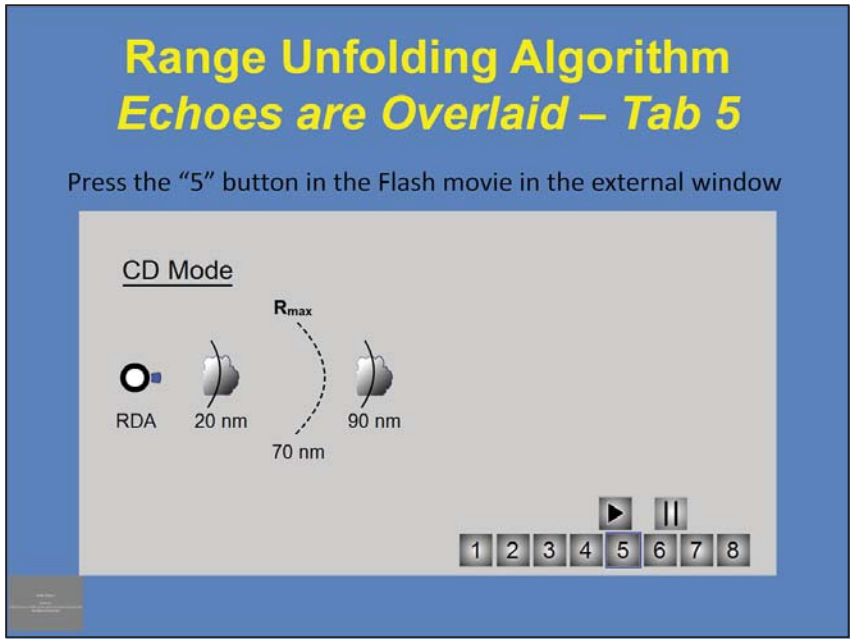
Tab 2. This time, when we go to Doppler mode, the maximum unambiguous range is going to be 70 nm. That means that both our targets will have an apparent range of 20 nm. It also means that we'll be getting pulses back from both of these targets at the same time, so this is an overlay situation.



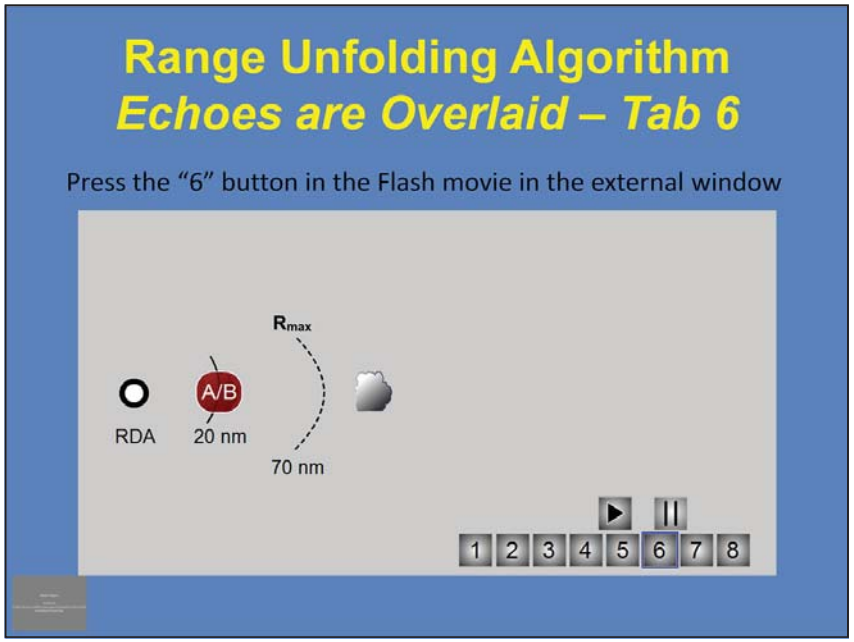
Tab 3. Just as before, not only does the algorithm account for the apparent ranges of the targets, it accounts for the possible ranges out into the second, third and even fourth trips.



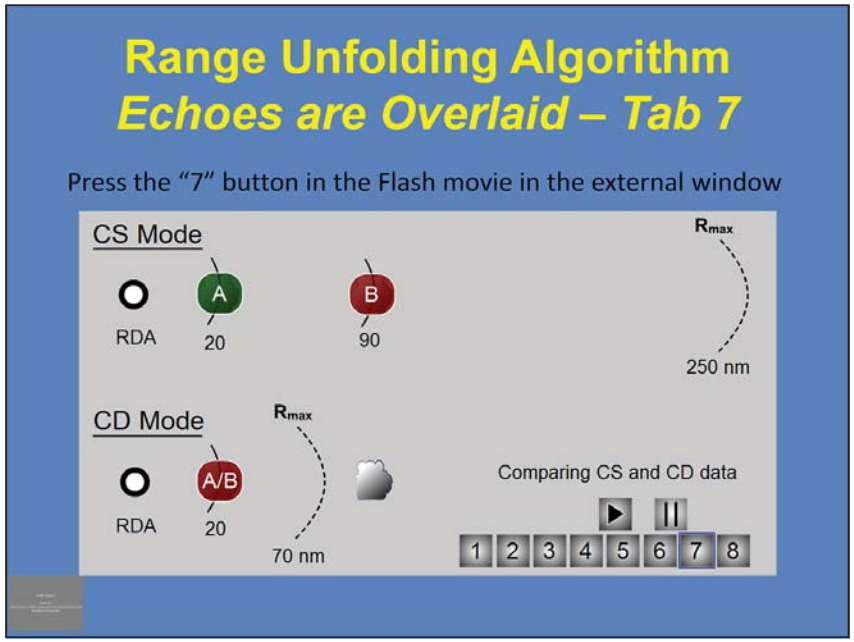
Tab 4. Here's where we take advantage of the fact that we know the returned power for each of these two targets and their original range. Those returned powers can be compared to one another, and a threshold, called TOVER, is used to determine whether we are going to assign velocity and spectrum width to one of these two targets. So the higher power is compared to the lower power, and if we reach this threashold, we have the ability to assign velocity and spectrum width to one of these overlaid targets.



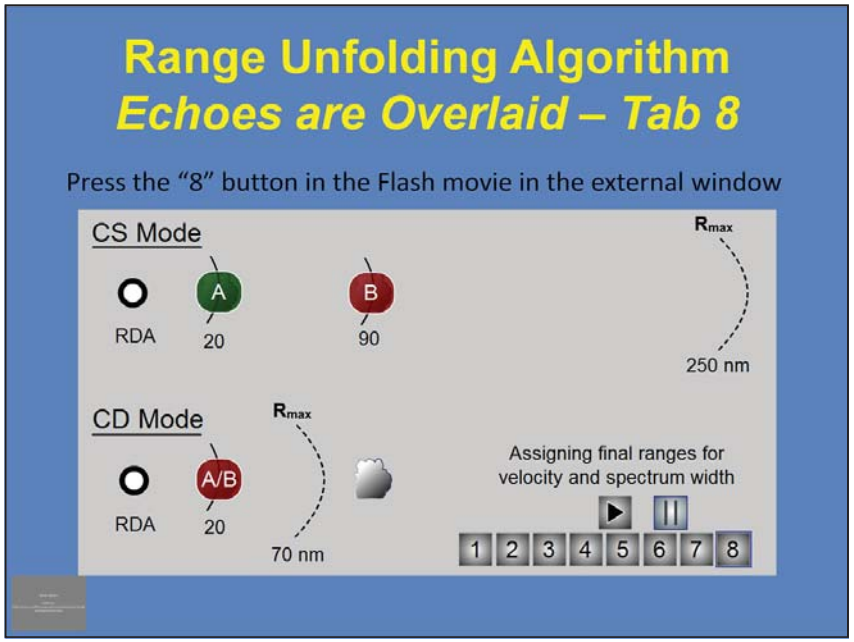
Tab 5. Now we actually do the Doppler data collection, and notice that information from pulses 1 and 2 come back the radar at the same time. So we have a velocity value with an apparent range of 20 nm, that is composed of pulses from both targets A and B.



Tab 6. Now we begin the comparison of the Doppler data to the Surveillance data, asking the question of whether a target was at a specific range, but actually including the question "was that the target that had the higher power?" And in this particular example, Target B had the lower power, so the answer is No at a range of 20 nm.



Tab 7. Now asking the same question, looking for a possible range for this target, "was there something at 90 nm?" Yes, and "was this the higher power target?" and the answer is Yes.

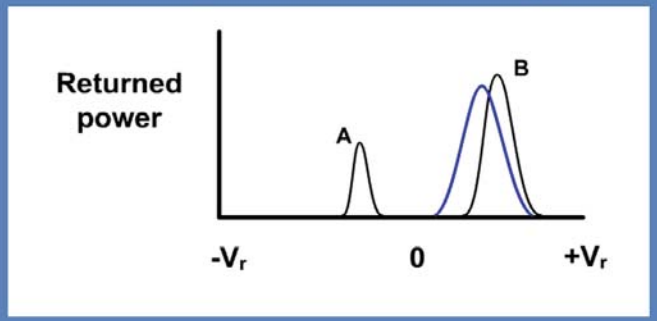


Tab 8. Now that we know that this velocity value that appears to be coming from a range of 20 nm, is more representative of the target that was at 90 nm. So we assign velocity at 90 nm, and range folding, or purple at 20 nm.

Range Unfolding Algorithm
Echoes are Overlaid

A and B will have same apparent range

- Pulses returning from both targets at same time
- Is base data representative of either target?



This graphic represents the base data estimation process for a range bin when there is an echo overlay case, such as 20 nm from our previous example. Pulses from two different targets are received and processed at the same time, apparently from 20 nm. If the pulses from these two different targets could be analyzed separately, the result would be the black bell curves for targets A and B. However, they can only be analyzed as a single target, which is represented by the blue bell curve.

Which one of the original targets does the blue curve better represent? Target B, and that is because B is returning significantly more power than target A. Recall that velocity estimates are power weighted, so this "hybrid" velocity estimate is going to be much closer to target B than target A, since B is returning much more power.

Now we need to determine if the power returned by B is sufficiently greater than A for the "hybrid" velocity to be assigned to target B.

Range Unfolding Algorithm
Echoes are Overlaid

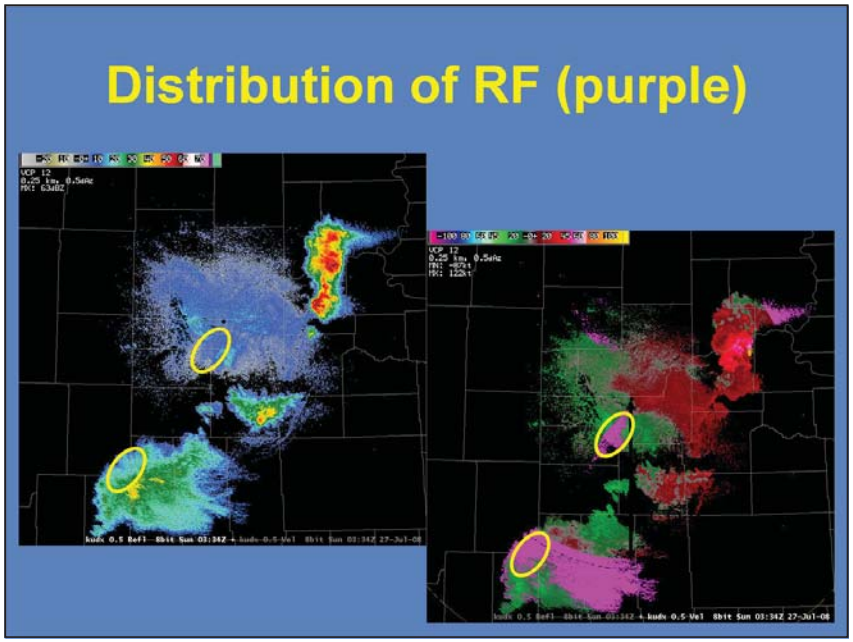


- If power ratio exceeds TOVER (5 dB)
  - V & SW assigned to echo with higher power
  - Other echo assigned RF
- If power ratio does not exceed TOVER
  - Both echoes assigned RF

The Range Unfolding algorithm uses a parameter known as "threshold over" or TOVER, to determine if overlaid echoes have a sufficiently higher returned power difference to assign the hybrid velocity to one of them.

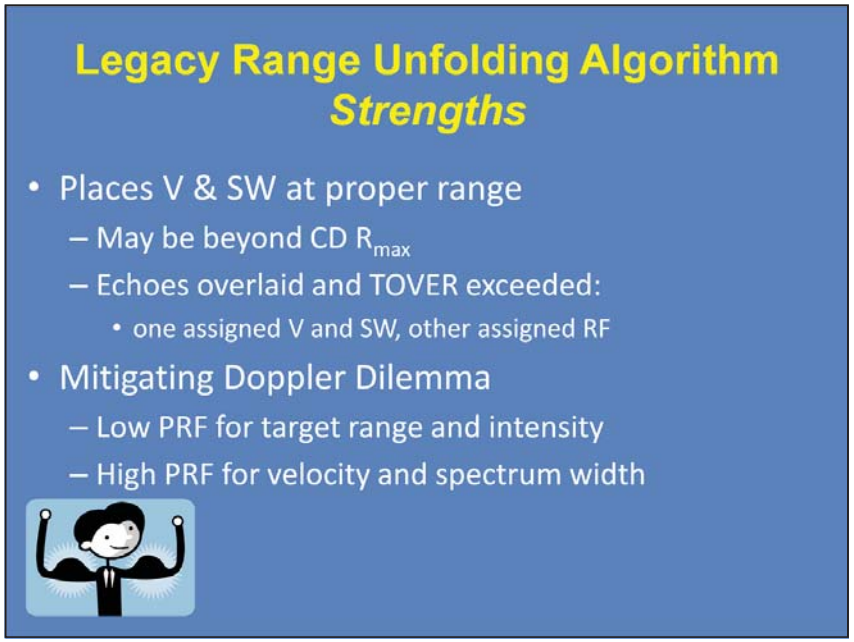
The current setting for TOVER is 5 dB, and if the power ratio exceeds TOVER, velocity and spectrum width are assigned to the echo which returned higher power. The other echo is assigned RF.

If the power ratio does not exceed TOVER, both of the overlaid echoes are assigned RF.



Now for another look at the distribution of RF on the velocity product, given the location of the echoes on the reflectivity. The strongest storms to the northeast are within the first trip, so there is no issue with the availability of the associated velocity data. There is a “blob” of weaker echo to the south and southwest of the radar, some of which falls within the first trip of the velocity data, some of which falls within the second trip.

The yellow ovals are meant to capture a group of gates in the same relative positions in the first and second trips, with purple assigned to both of the bins. The overlaid echoes here did not return enough of a power difference to accurately assign velocity data to either one, so purple (RF) was assigned to both.



Now for the strengths of the Legacy Range Unfolding algorithm. Most of the time, it is able to achieve its objective of assigning velocity and spectrum width data to the appropriate range, which may be beyond the R_{max} for Contiguous Doppler (CD) mode. When echoes are overlaid, the Legacy Range Unfolding algorithm can, at best, assign velocity and spectrum width to one of the overlaid echoes. This assignment is based on the returned power of the separate echoes that contribute to the overlay. If TOVER is exceeded, that means one of the echoes has returned power that is sufficiently high to be assigned the velocity and spectrum width. The other echoes are assigned RF.

The Legacy Range Unfolding algorithm is designed to mitigate the Doppler Dilemma. For every radial, the low PRF, Contiguous Surveillance (CS) pulses provide returned power and target range. For that same radial, the high PRF, CD pulses provide accurate velocity and spectrum width data. Comparing the information from CS and CD along the radial allows for “unfolding” the velocity and spectrum width to its appropriate range.

Legacy Range Unfolding Algorithm *Limitations*

- Extensive echo coverage with echoes extended along radials
  - Echo overlay maximized
- V & SW unavailable for overlaid echoes if power ratio does not exceed TOVER



The Legacy Range Unfolding algorithm is going to have the greatest limitation on its performance when there is extensive echo coverage aligned along radials, through both the first and second trips in the CD data. This situation maximizes the echo overlay along the radial. Since the Legacy Range Unfolding algorithm can assign velocity and spectrum width to, at best, one of the overlaid echoes, the result is a lot of RF data. In this example, the echo coverage is extensive for both the first and second trips to the north and the south of the radar. There is extensive RF coverage in the second trip because the TOVER requirement was met by the echoes in the first trip.

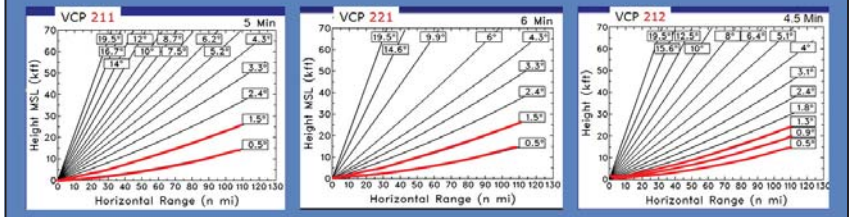
In some echo overlay cases, RF will be assigned to both trips if TOVER is not exceeded.

SZ-2 Range Unfolding Algorithm

- Run at RDA signal processor

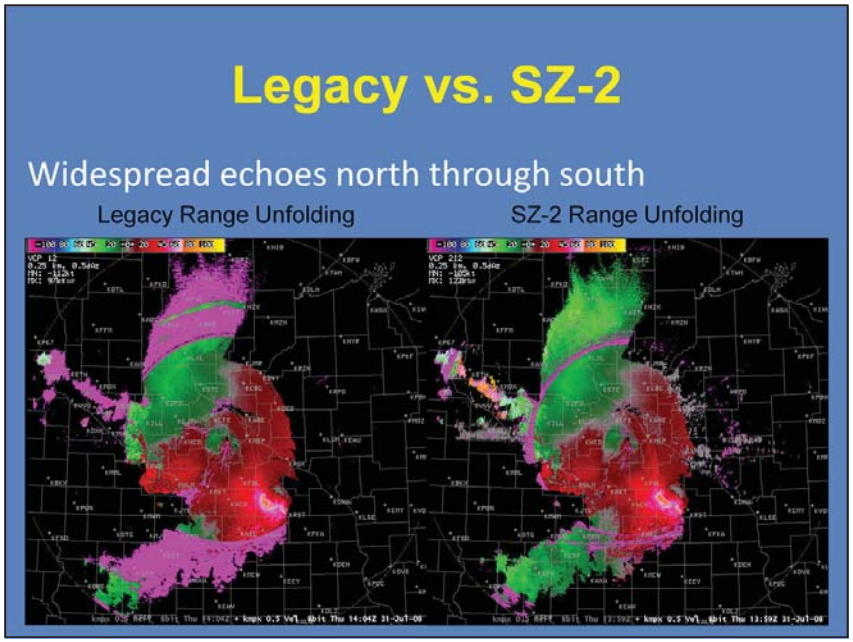


- Applied to Split Cut elevations only on VCPs 211, 221, and 212

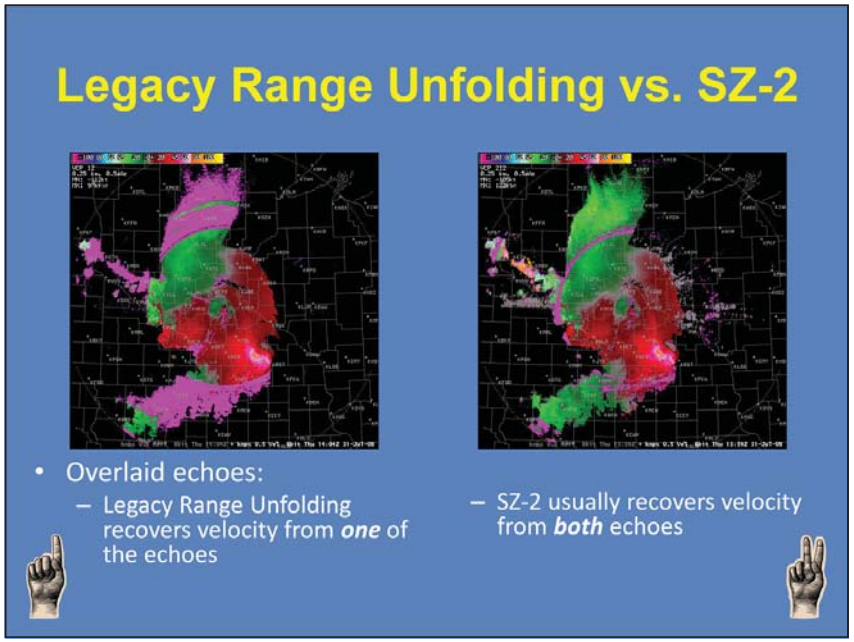


The next technique for range unfolding the velocity and spectrum width data was fielded in 2007, and is know as SZ-2 Range Unfolding. This technique is named for the two research scientists who developed it: Mangalore Sachidananda and Dusan Zrnic. It is also run at the RDA signal processor.

SZ-2 is available only for the Split Cut elevations of the three SZ-2 VCPs, 211, 221, and 212. The “2” as the first digit in the VCP name indicates that it is an SZ-2 VCP. The second two digits in the name tells you the structure of the VCP with respect to the angles used and the update rate. For example, VCP 212 is just like VCP 12 expect that SZ-2 is used to range unfold velocity and spectrum width on the Split Cuts.



The SZ-2 Range Unfolding algorithm is particularly effective with widespread echo coverage. In this case, the Legacy Range Unfolding algorithm (VCP 12) is used for the data on the left, with SZ-2 (VCP 212) on the right. There is a dramatic improvement in the availability of velocity and spectrum width when VCP 212 is used.



This example also demonstrates the key difference between the two range unfolding techniques currently available when there is echo overlay, i.e. returns from multiple trips in the CD data. When there is echo overlay, the Legacy Range Unfolding algorithm can recover, at best, one of the echoes. On the other hand, SZ-2 can usually recover both overlaid echoes.

SZ-2 Strengths

- SZ-2 Strengths
  - Significant increase in availability of velocity data
  - Best results with widespread returns
  - VCPs 212 and 211 designed for widespread rapidly evolving severe convective storms



The strength of the SZ-2 Range Unfolding algorithm is the increased availability of velocity data for multiple trips. The best results come with events with widespread echo coverage. VCPs 212 and 211 are designed for widespread rapidly evolving severe convective storms. However, VCP 212 has superior low level vertical resolution, just as VCP 12 does.

SZ-2 Limitations

- All Bins degrades SZ-2 velocity
  - CMD enabled: Default clutter file downloaded with SZ-2 VCP
  - CMD disabled: CMD enabled and Default clutter file downloaded with SZ-2 VCP



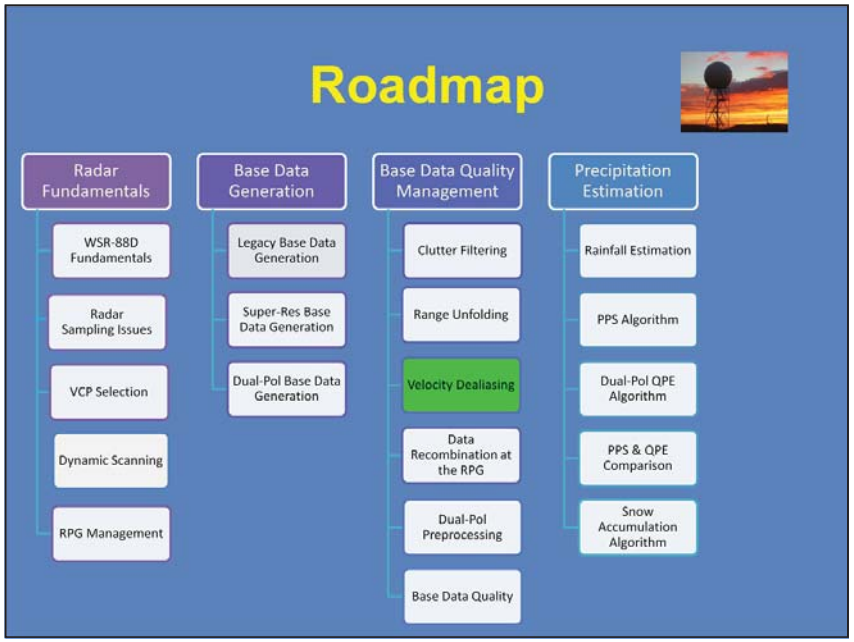
With CMD active, the need for All Bins clutter suppression is nearly zero (never say never?). SZ-2 is another reason to avoid All Bins suppression, as it can degrade the velocity data when SZ-2 is applied. The RPG software has some steps in place to avoid having All Bins and SZ-2 active at the same time.

If CMD is enabled, and an SZ-2 VCP has been downloaded, the Default clutter regions file will be downloaded along with it. This ensures that the Bypass Map is in control everywhere, overwriting any All Bins filtering that may already be in place.

If CMD has been disabled, and an SZ-2 VCP has been downloaded, CMD will be enabled, and the Default clutter regions file will be downloaded along with the SZ-2 VCP.



Welcome to this lesson on Velocity Dealiasing.



Here is the “roadmap” with your current location.

Objectives

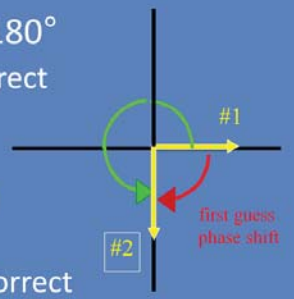


1. Identify the purpose, strengths and limitations of the following techniques to dealiase velocity data
   a) Legacy Velocity Dealiasing Algorithm (VDA)
   b) 2 Dimensional Velocity Dealiasing Algorithm (2D-VDA)
   c) VCP 121: Multiple PRF Dealiasing Algorithm

There is one objective for Velocity Dealiasing, and the different algorithms will be taught in sequence in this module.

Improperly Dealiased Velocities

- V estimated from pulse-to-pulse phase shifts
- First guess V based on shift <180°
  – True shift <180°, first guess correct
- Possible Vs (aliases) based on shift >180°
  – True shift >180°, first guess incorrect



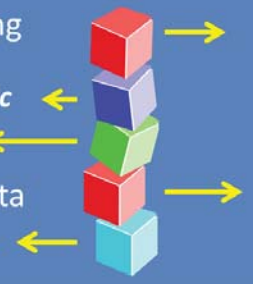
The previous lesson presented the process for assigning velocity and spectrum width to its appropriate range. The radial velocity value itself is the first guess velocity, which may not be correct. Perhaps one of the aliases of that first guess is the correct velocity. The example used in Base Data Generation had a first guess velocity of +30 kts, while the correct velocity was -90 kts.

Doppler velocity is estimated based on pulse-to-pulse phase shifts, with the first guess velocity calculated from the phase shift < 180°. For each first guess velocity, there are possible velocities, or aliases, based on phase shifts > 180°.

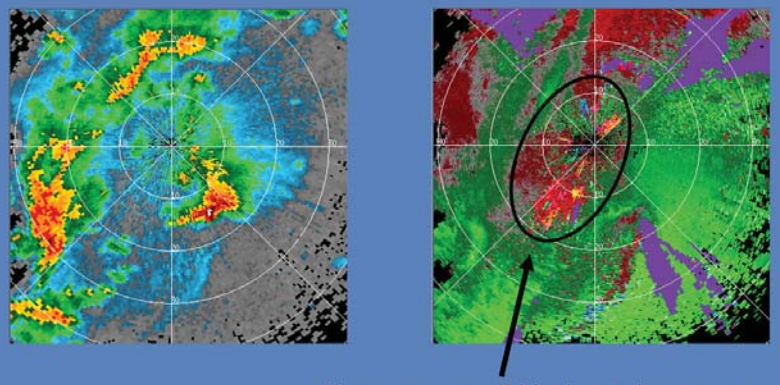
Improperly Dealiased Velocities
Product Characteristics

1. Small blocks of Vs in direction opposite from surrounding data
   - Usually at close range
2. Larger blocks or wedges of Vs in direction opposite from surrounding data
   - No zero velocity boundary; ***unrealistic shears***

Most likely in areas lacking velocity data continuity

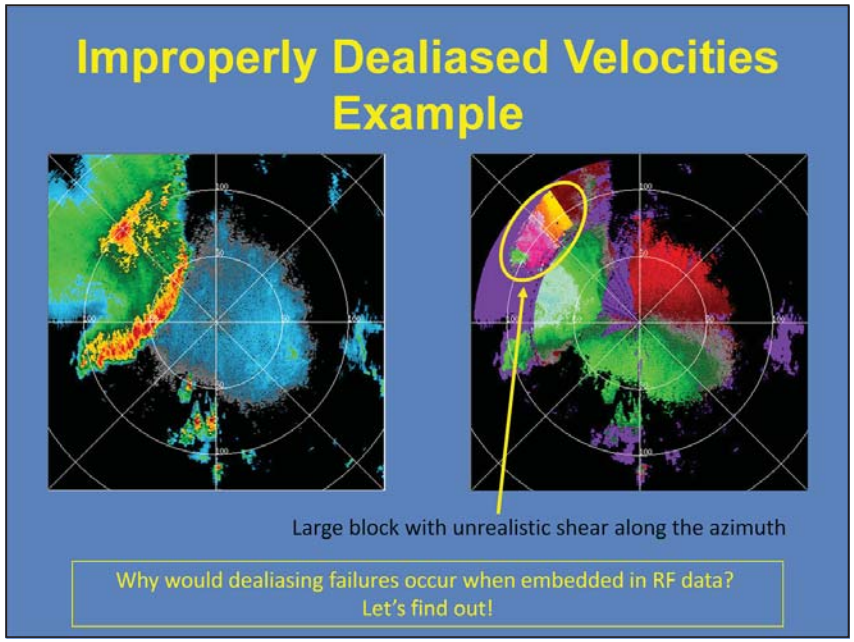


Improperly Dealiased Velocities
Example

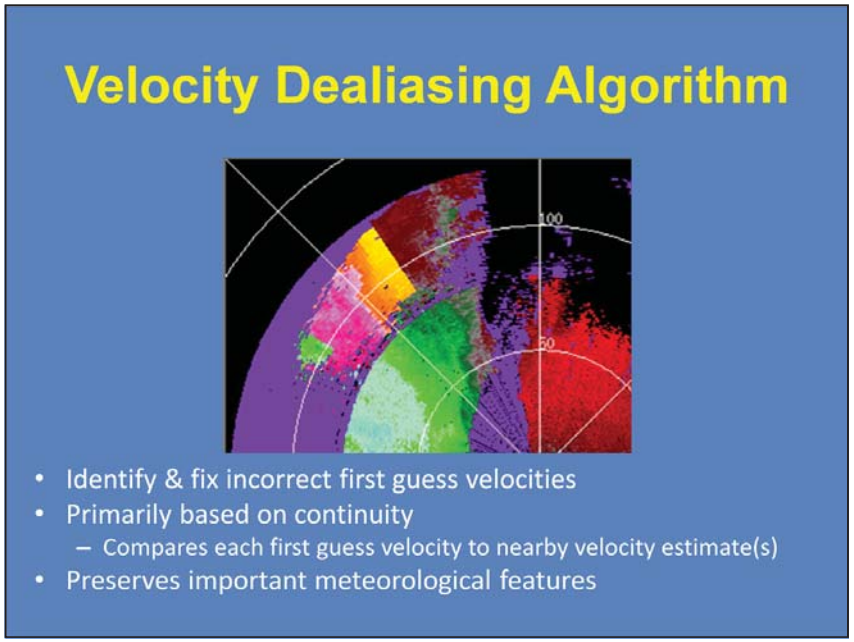


Before looking at the algorithm that "de-aliases" velocities, we first look at the impact of improperly dealiased velocities on the radar products. There are two types of improperly dealiased velocities. At close range, especially in residual ground clutter, small blocks of velocity values opposite in direction from the surronuding data often occur. Typically away from very close range, is another type of improperly dealiased velocity. These blocks or wedges of values opposite in direction from the surrounding data are generally larger. Sometime more challenging are the shears that appear along an azimuth. In some cases, it is difficult to determine if these azimuthal shears are meteorological, or the result of improperly dealiased velocities.

Here is an example of the type of improperly dealiased velocities frequently seen near the radar in the residual clutter region. They are usually not operationally significant.



This is an example of improperly dealiased velocities resulting in two azimuthal shears that are clearly not meteorological. Another term for this data artifact is dealiasing failures. The fact that these failures occur while embedded in an area where RF has been assigned is related to how the dealiasing algorithm works.



The goal of the Velocity Dealiasing algorithm is to assign the correct radial velocity to each range bin. It first identifies any first guess velocities that are "suspect", then attempts to assign one of the aliases if that is more appropriate. The decision making for what is "appropriate" comes from comparing each first guess velocity and its aliases to neighboring velocity values.

There are additional steps in this algorithm that are designed to preserve real meteorological shears, both from one azimuth to the next (for example, a circulation), and along a radial (for example, storm top divergence).

Velocity Dealiasing Algorithm

- First 3 steps search for V close to each 1st guess
  - Compare 1st guess against V neighbor
  - Compare 1st guess aliases against V neighbor
- Does +59 make sense?

-68	-56
-74	+59
-68	+50
-67	+53
-64	+53
-62	+59
-15	-18
-3	-14
0	-8
+9	-4
+9	0

The first three steps of the Velocity Dealiasing algorithm are similar in function, and are transparent to the user. Once a first guess velocity is identified along a radial, it and its aliases are compared to a "velocity neighbor". What constitutes a "velocity neighbor" varies with each step, getting a little further away from the first guess velocity. The essence of steps one through three is to assess whether than first guess velocity makes sense given the surrounding velocity field.

Velocity Dealiasing Algorithm
Significance of Steps 1 through 3

- Common goal of steps 1, 2 & 3:
  - Compare each 1st guess V to nearby, already dealiased V
- Each step looks a little further away for V to compare to 1st guess
- Step 4 most important since ***you*** have ***input and oversight***!



The first three steps of the Velocity Dealiasing algorithm have the common goal of comparing each first guess velocity to a "velocity neighbor". That "velocity neighbor" varies with each step, getting a little further away from the first guess velocity. With respect to everyday operations, the most important step with the Velocity Dealiasing algorithm is step four. This final step is the most important, because this is where users have both input and oversight, meaning here's where dealiasing failures can be mitigated.

Velocity Dealiasing Algorithm
Step 4: Environmental Winds Table

- 1st guess and aliases compared to EWT
  - Must be within threshold of EWT velocity
- *Important to have representative EWT to support Velocity Dealiasing Algorithm*
- EWT interface
  - Environmental Data Editor – graphical
  - Environmental Data Entry - tabular

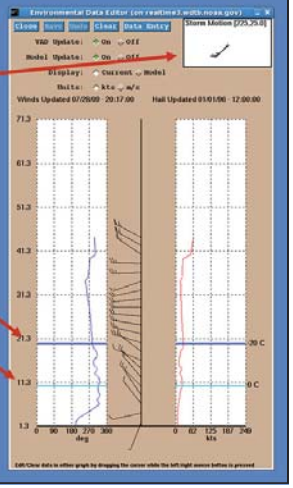


When the previous steps do not resolve a possible incorrect first guess velocity, the Velocity Dealiasing algorithm accesses the Environmental Winds Table to find a representative velocity for the given range and height. It is important that the Environmental Winds Table, which is stored and updated at the RPG, has a realistic picture of the state of the winds aloft.

There are two interfaces related to the Environmental Winds Table, one which is graphical, and another which is a table. The titles are actually Environmental Data, because environmental information other than winds aloft are stored, such as the height of 0° C.

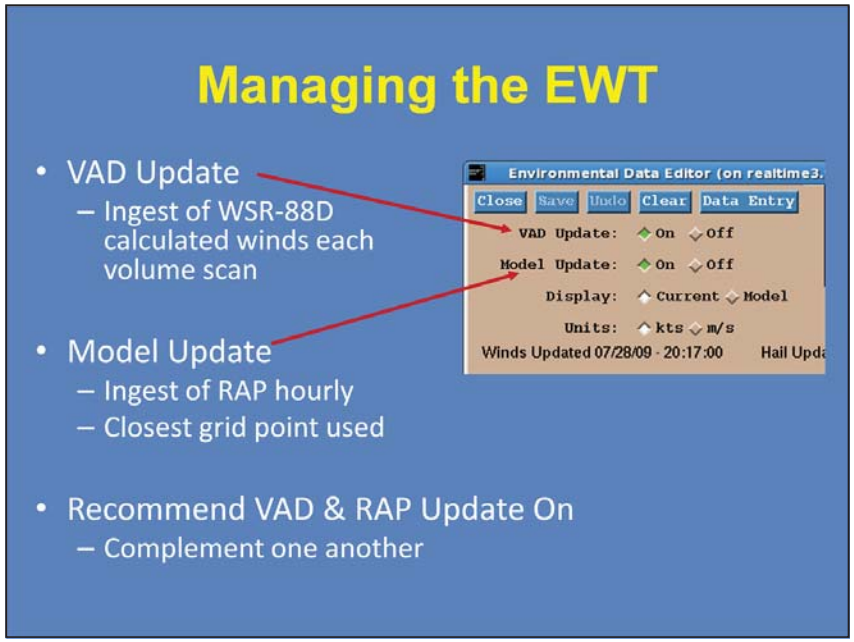
Environmental Data Editor Window

- EWT: Winds aloft from VAD and/or RAP
- Default storm motion and hail temperature heights (coming up)



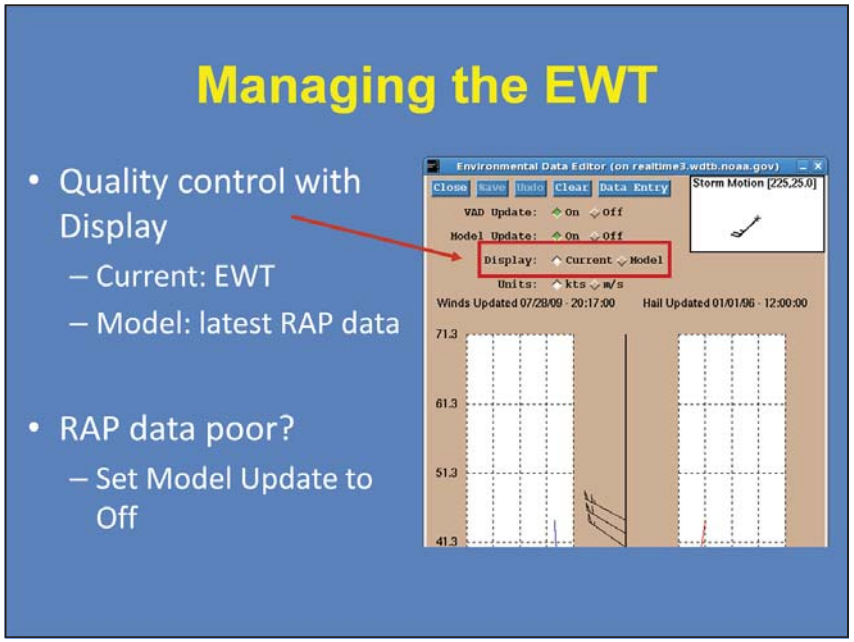
When initially accessing Environmental Data, this window provides multiple types of data. The local winds aloft is presented with the wind barbs in the center, as well as the wind direction and speed represented on the graphs. The input for the wind information is usually a combination of input from the Velocity Azimuth Display (VAD) and the Rapid Refresh (RAP) model. The VAD is an RPG algorithm that uses WSR-88D data to generate winds at a series of heights, while “Model Update” on this RPG window is referring to the RAP.

The Environmental Data Editor window also has the default storm motion and the hail temperature heights.



Here's a closer look at the buttons that control inputs to the Environmental Winds Table. When VAD Update is set to On, the EWT is updated every volume scan by VAD-generated winds aloft. When Model Update is set to On, the RAP data for the closest grid point are used hourly to update the EWT.

It is recommended that both of these updates be set to On, unless there is some kind of problem. These two data sources complement one another. The VAD provides updates every volume scan, but the data are limited to available scatterers to generate radar detected winds. The RAP is available hourly, with wind data available throughout the column at the grid point closest to the radar.



Managing the EWT is a task of monitoring it for relevance, and the Display button can help. When Current is selected, the window is displaying the current state of the EWT. Based on knowledge of winds aloft from other sources, you can verify if the EWT is representative. If there is a need to check the quality of the RAP data, then select Model after Display. The window will display the last hourly model input. If these data are poor, Model Update can be set to Off.

Legacy VDA
Strengths

- Best possible velocity data for algorithms
  - MDA, TDA, SCIT….
- Provides velocity estimates > V_{max}
- Preserves significant meteorological features
  - Gust fronts
  - Storm top divergence
  - Mesocyclones
  - TVS



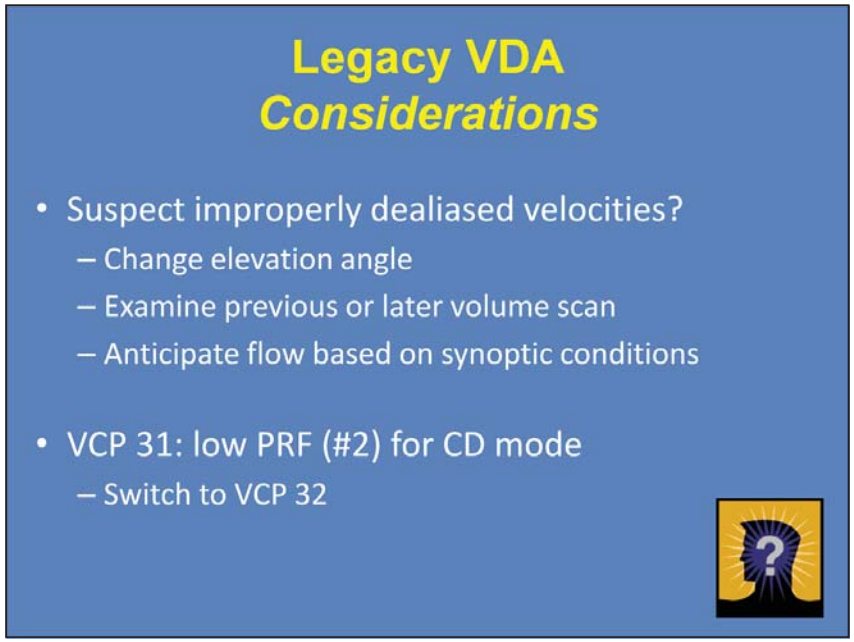
The Velocity Dealiasing algorithm is designed to provide the best possible velocity data, primary to support the RPG algorithms that look for significant features such as circulations. The Velocity Dealiasing algorithm attempts to assign the true radial velocity, even when it exceeds V_{max}. There are quality control steps with the Velocity Dealiasing algorithm that are designed to preserve significant meteorological shears, such as gust fronts, storm top divergence, mesocyclones and tornadic vortex signatures.

Legacy VDA
Limitations

- Performance degraded by
  - unfiltered clutter
  - weak returned power
  - limited pulses per radial required for faster VCPs
- VDA failures
  - can mask real shears
  - can contaminate algorithms (false MDA detections)
  - most likely in sparse data (leading edge of storms)

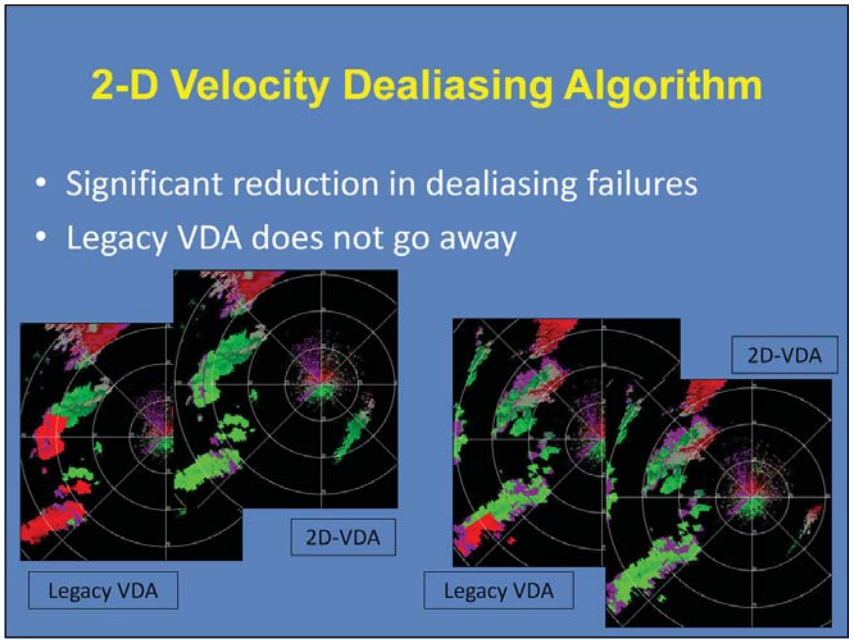


The performance of the Velocity Dealiasing algorithm is degraded by data quality problems that reduce the reliability of the first guess velocity. This includes unfiltered clutter, weak returned power, and in some cases, the low number of pulses per radial with the faster VCPs. When dealiasing failures occur, the result can mask real shears that you would want to see. Dealiasing failures can also contaminate the results of the RPG algorithms, such as false circulations from the Mesocyclone Detection Algorithm (MDA). Dealiasing failures are mostly likely to occur in areas of sparse data, which means there are no neighboring velocity values for comparison with a first guess velocity.



Sometimes dealiasing failures are subtle, and it can be hard to determine if you are seeing a valid shear. It is helpful to be aware that improperly dealiased velocities are not usually preserved from one elevation angle to the next, or from one volume scan to the next. For synoptic or mesoscale flow, it can be helpful to have an expectation based on other data sources.

For Clear Air mode operations, VCP 31 uses long pulse, which provides the best sensitivity, and is good for detecting light precipitation such as snow. However, VCP 31 uses a low PRF for velocity data, and dealiasing failures are more likely. For Clear Air mode operations, if velocity detection is the highest priority, VCP 32 is recommended.

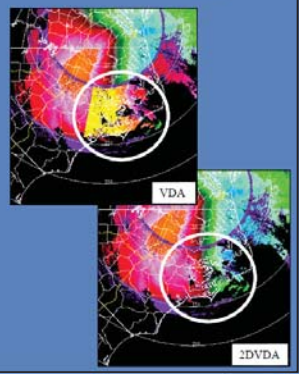


The Two Dimensional Velocity Dealiasing Algorithm (2D-VDA) was deployed in 2013, and offers a significant reduction in dealiasing failures. These examples come from an event with storms that were moving so fast that the legacy Velocity Dealiasing Algorithm had numerous dealiasing failures, while the 2D-VDA did much better.

Though the 2D-VDA is the default velocity dealiasing algorithm, it does not completely replace the legacy Velocity Dealiasing Algorithm.

2-D Velocity Dealiasing Algorithm

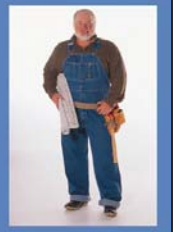
- Dealias entire elevation of V data
  - Azimuth/radial grid built with median V for each grid center
  - V field partitioned to dealias small features
  - Weighting factors reduce data noise
    - Low weighting where spectrum width is high



The 2D-VDA is much more robust than the Legacy VDA, and this slide provides a very brief overview. The overall approach is to use a least squares method to minimize errors in the velocity. For each elevation, 2D-VDA first builds a 2 dimensional grid (azimuth and radial) of the velocity data, with a median velocity value for each grid center point. This serves as a large scale dealiasing step. The velocity field is then partitioned in order to dealias small scale features such as mesocyclones and tornadic vortex signatures. There are also steps that involve applying weighting factors, primarily to reduce noisiness in the velocity data. For example, bins with a high spectrum width would have lower weighting, because velocity estimates are usually less reliable where spectrum width is high.

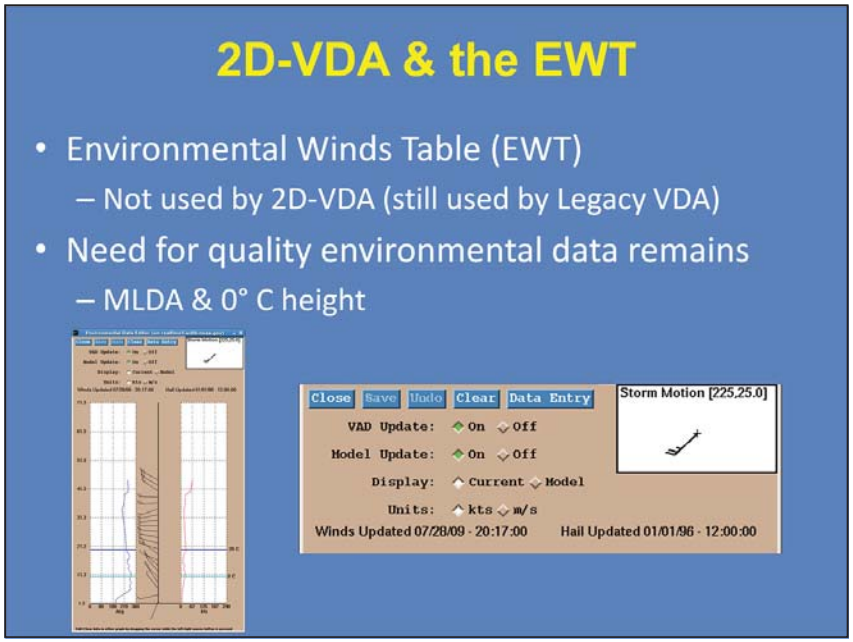
2D-VDA Implementation

- 2D-VDA is default dealiasing algorithm
- RPG software reverts back to legacy VDA:
  - Differing Doppler PRFs in sectors
  - VMI set to 1.94 kts
- All other conditions: RPG uses 2D-VDA
- Switching to/from 2D-VDA automated



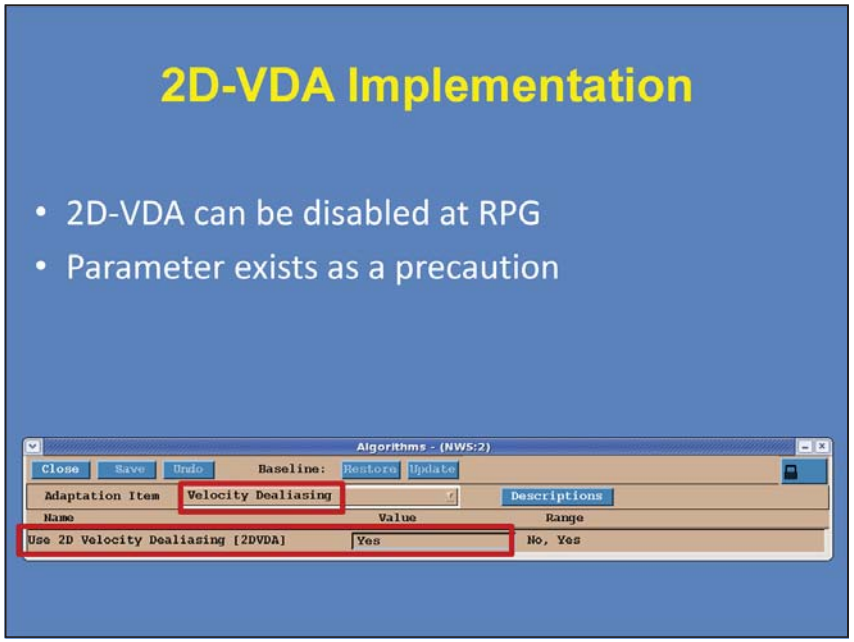
It is important to be aware of how the 2D-VDA has been implemented. Unless it is turned off, the 2D-VDA is the default velocity dealiasing algorithm. There are two conditions where the RPG software will automatically revert back to the Legacy VDA. The first is when a VCP with different Doppler PRFs in the three sectors has been downloaded and is active. The second condition is when the Velocity Measurement Increment (VMI) is set to 1.94 kts (the default VMI is 0.97 kts).

Outside of these two conditions, the 2D-VDA will be active unless it is manually turned off in the Algorithms window. The RPG will automatically switch to and from the 2D-VDA as needed.



The Environmental Winds Table (EWT) supports the performance of the Legacy VDA, and will continue to do so whenever the Legacy VDA is running. The 2D-VDA does not rely on the EWT.

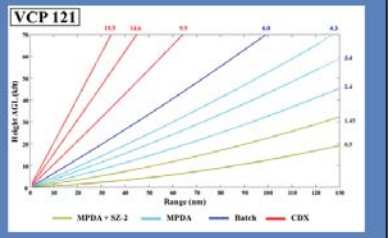
Though the environmental winds will likely be used less often for dealiasing velocity data, maintaining the validity of all the environmental data, such as the 0° and -20° C heights, remains important. For example, the Melting Layer Detection Algorithm (MLDA) will not always have sufficient radar detections to identify a melting layer. When that occurs, the MLDA relies on the RPG 0° C height, which was either manually entered or from the model data.



There is a new entry at the Algorithms window at the RPG, called “Velocity Dealiasing”. The parameter, “Use 2D Velocity Dealiasing”, controls whether the 2D-VDA is used, and the default setting is Yes. Setting this parameter to No is not expected to be needed, but is available as a precaution.

Multiple PRF Dealiasing Algorithm (MPDA)

- Designed to mitigate range folding ***and*** improperly dealiased velocities



- VCP 121 used solely for MPDA
  - VCP 121 is MPDA version of VCP 21
  - Additional CD rotations at lower elevations
    - At 0.5°, 1 CS and 3 CD rotations

The Multiple PRF Dealiasing Algorithm (MPDA) is a special application uniquely used for VCP 121. It is designed to mitigate both range folding and improperly dealiased velocities, with range folding mitigation the most apparent. VCP 121 samples the same elevations as VCP 21, and thus has the same limitations. VCP 121 has additional Doppler rotations for the lower elevations. For example, at 0.5, there is one CS rotation followed by 3 CD rotations. Each of these CD rotations uses a different Doppler PRF.

VCP 121

- SZ-2 + Legacy Range Unfolding + extra CD rotations

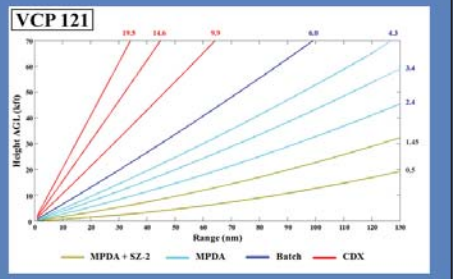
VOLUME COVERAGE PATTERN 121
SCAN STRATEGY MPDA/SZ2 SHORT PULSE

Scan				Surveillance		Doppler PRF No.				
Elevation (deg)	AZ Rate (deg/sec)	Period (sec)	WF Type	PRF No.	No Pulses	4 No. Pulses	5 No. Pulses	6 No. Pulses	7 No. Pulses	8 No. Pulses
0.5	18.677	19.28	SZCS	1	17	-	-	-	-	-
0.5	19.754	18.22	SZCD	8	-	43	51	55	59	64
0.5	27.400	13.14	CD	6	-	31	37	40	43	46
0.5	21.401	16.82	CD	4	-	40	47	51	55	59
1.45	19.842	18.14	SZCS	1	16	-	-	-	-	-
1.45	19.754	18.22	SZCD	8	-	43	51	55	59	64
1.45	27.400	13.14	CD	6	-	31	37	40	43	46
1.45	21.401	16.82	CD	4	-	40	47	51	55	59
2.4	19.205	18.75	B	1.8	6	27	32	34	37	40
2.4	27.400	13.14	CD	6	-	31	37	40	43	46
2.4	21.401	16.82	CD	4	-	40	47	51	55	59
3.35	21.599	16.67	B	2.8	6	28	33	35	38	40
3.35	27.400	13.14	CD	6	-	31	37	40	43	46
3.35	21.401	16.82	CD	4	-	40	47	51	55	59
4.3	16.304	22.08	B	2.4	6	40	48	52	56	61
4.3	29.498	12.20	CD	7		29	34	37	40	43
6.0	20.204	17.82	B	3.5	6	34	40	43	47	51
9.9	29.498	12.20	CD	7	-	28	34	37	40	43
14.6	29.795	12.08	CD	8	-	28	33	36	39	43
19.5	29.795	12.08	CD	8	-	28	33	36	39	43

Here's a snapshot of the design of VCP 121. For the lowest two elevation angles, there are three CD rotations with different Doppler PRFs, and one of these CD rotations uses SZ-2 to range unfold the velocity data.

VCP 121 Considerations

- VCP 21 angles; ***fastest antenna rotations***
  - 20 rotations in 5 mins 45 secs

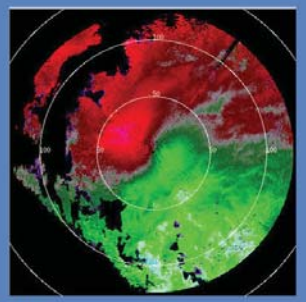


- VCP 121 ***not*** appropriate for fast moving or rapidly evolving storms
  - Use VCP 12 or 212
  - Better vertical sampling & faster updates

Since VCP 121 has additional rotations for several elevations, it has the greatest number of total rotations of any VCP. With an update rate of just under 6 minutes, VCP 121 has the fastest antenna rotation rates of any VCP, and a low numbers of pulses per radial. VCP 121 is not appropriate for fast moving and/or rapidly evolving storms. VCPs 12 or 212 are the appropriate choices, given their better low level vertical sampling and their faster updates.

Strengths of MPDA/VCP 121

- For lowest 2 elevations, VCP 121 recovers nearly ***all*** velocity data



- Designed for
  - Hurricanes
    - Not tornadic storms within rainbands
  - Events with widespread echo coverage



MPDA, aka VCP 121 is capable of recovering nearly all velocity data for the lowest two elevations. It is designed for sampling hurricanes while still offshore when large scale velocity structure is the priority. Once operations shift to looking for potentially tornadic storms within the rainbands, VCP 121 is no longer appropriate (VCPs 12 or 212 are better choices). VCP 121 is designed for events with widespread echo coverage, provided there are no severe convective cells to interrogate.

Limitations of MPDA/VCP 121

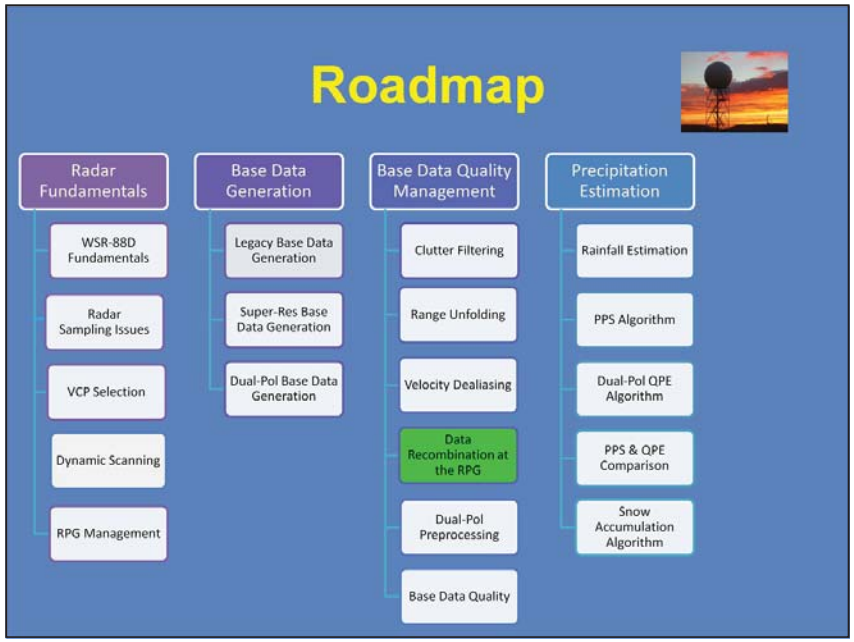
- VCP 121 *not* an appropriate choice for
  - Tornadic storms close to the RDA
  - Any situation where fast updates from low elevation base products are a priority
- High antenna rotations + CMD + GMAP + Super Res processing can degrade data quality



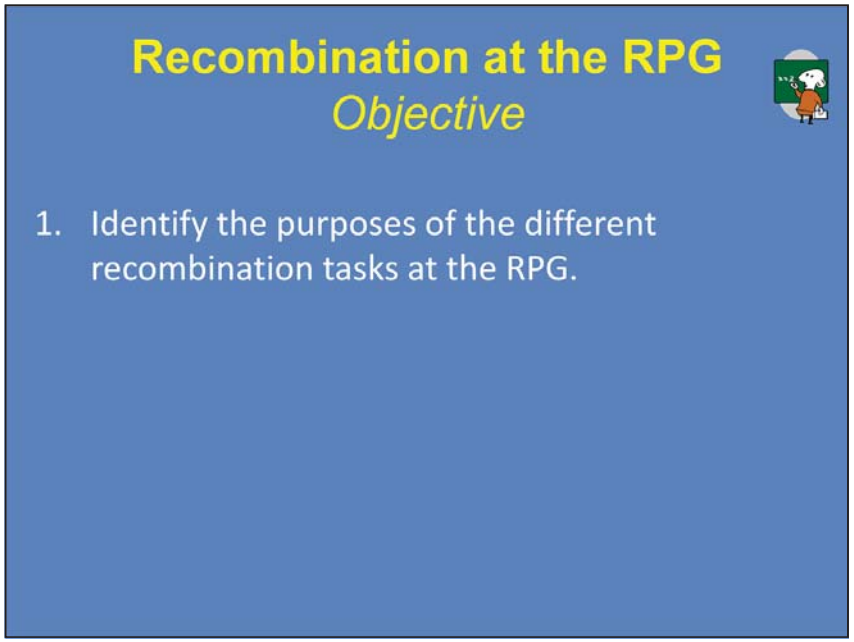
VCP 121 is not an appropriate choice for tornadic storms close to the radar, or any situation where fast updates from low elevation base products are a priority. VCP 121 has the highest antenna rotation rates, with low numbers of pulses per radial. With the application of CMD, GMAP, and super resolution processing, VCP 121 is more vulnerable to data quality problems.



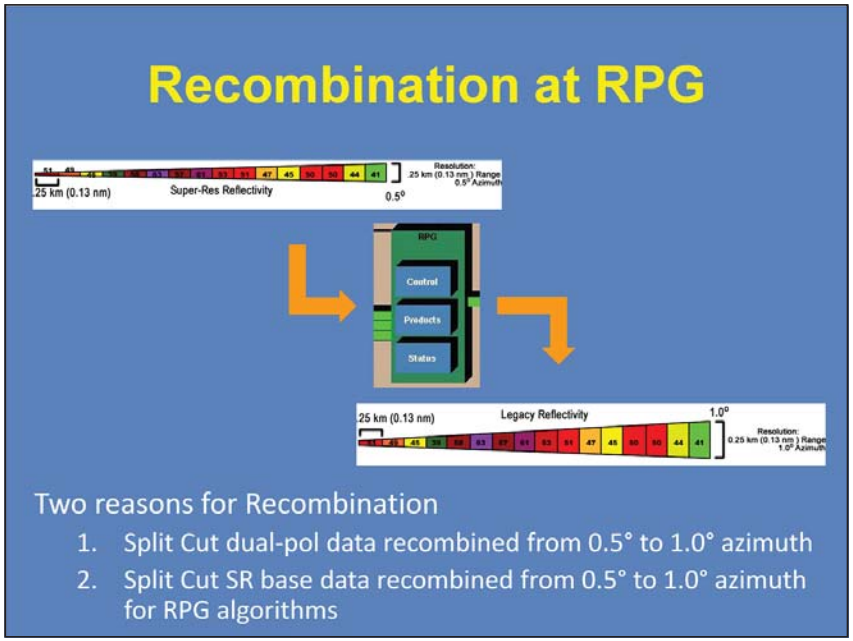
Welcome to Recombination at the RPG.

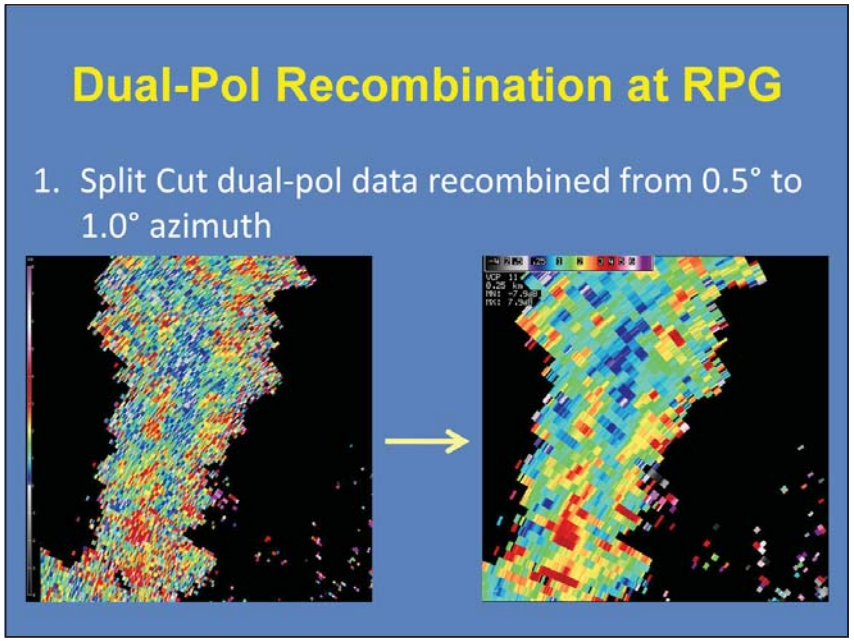


Here is the “roadmap” with your current location.



There is one objective in Recombination at the RPG.





Recombination is a process run at the RPG before many of the products are built. There are two cases where the base data that comes directly from the RDA needs adjustment before it can be used to build products or to be used for algorithm input.

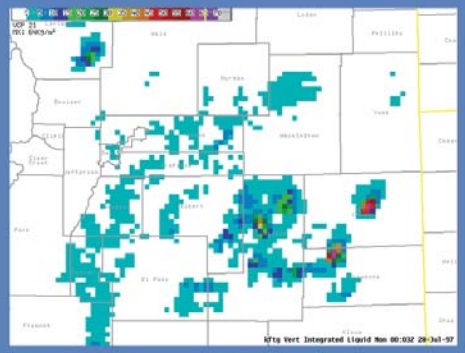
1. For the Split Cut elevations, the dual-pol base data are processed at the RDA with an azimuthal resolution of 0.5°, then sent via the wideband to the RPG. This includes Differential Reflectivity, ZDR, Correlation Coefficient, CC, and Differential Phase, Φ_{DP} , all of which are too visually noisy for direct product generation.

2. For the Split Cut elevations, the legacy base data (reflectivity, velocity and spectrum width) are super resolution. Many of the RPG algorithms cannot ingest data with 0.5° azimuth, so it must first be recombined to an azimuth of 1.0°.

The dual-pol base data arriving from the RDA are ZDR, CC, and Φ_{DP}. For the Split Cuts, all of these data have an azimuthal resolution of 0.5°, and are too visually noisy for direct product generation. These data are recombined to an azimuth of 1.0°.

The dual-pol base data are also “preprocessed”, which includes smoothing and converting Differential Phase, Φ_{DP}, into Specific Differential Phase, KDP. The Dual-Pol RPG Preprocessor algorithm will be discussed in a later lesson. On the left is the ZDR base data displayed in a Level II viewer (GR Analyst). On the right is the associated ZDR product displayed on AWIPS. The data have been recombined to 1.0°, as well as preprocessed.

Recombination at RPG

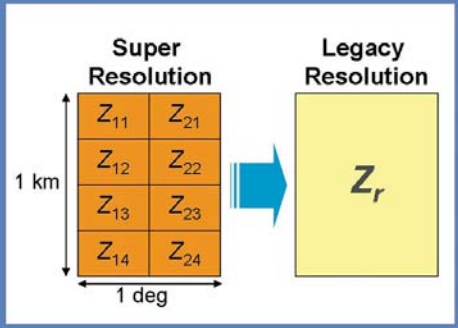


2. Many RPG algorithms (VIL, PPS, TDA,…) cannot ingest SR base data
   – MDA is an exception

Most of the RPG algorithms that rely on legacy base data cannot accept super resolution (0.5° azimuth) for input. Many also were designed for reflectivity base data with 1 km range resolution. Examples include the tornado detection algorithm, precipitation processing, vertically integrated liquid, and many others. One exception is the Mesocyclone Detection Algorithm (MDA). You will learn much more about the derived products later in this course.

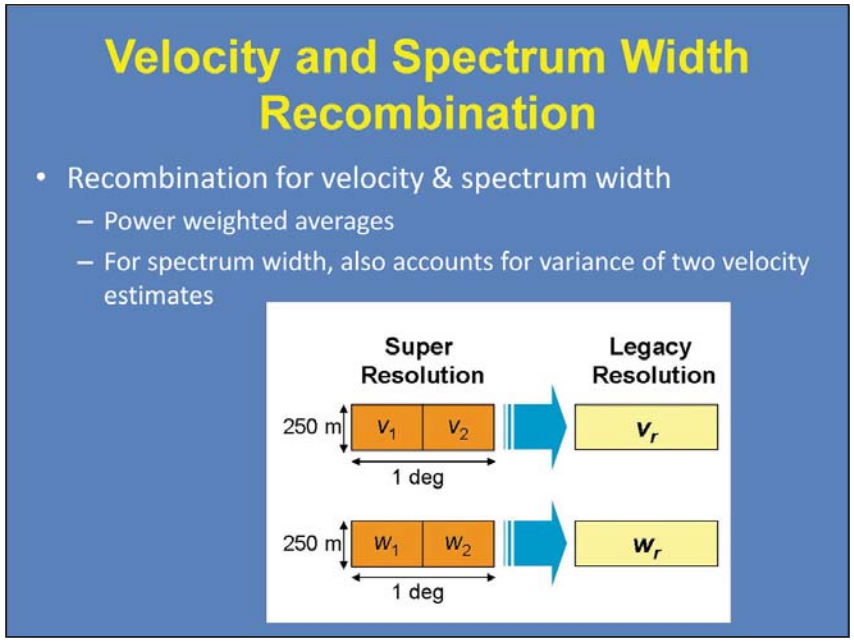
Reflectivity Recombination

- Linear average of 8 bins
- For bins with No Data, power estimated & associated Z included in average



This step is to support RPG algorithms that require 1.0° azimuth by 1 km range resolution for reflectivity input.

For reflectivity, the recombination process is a linear average of the 8 super resolution bins into the corresponding 1 legacy resolution bin. For bins that are assigned No Data, the associated power is estimated, converted to Z and included in the average.

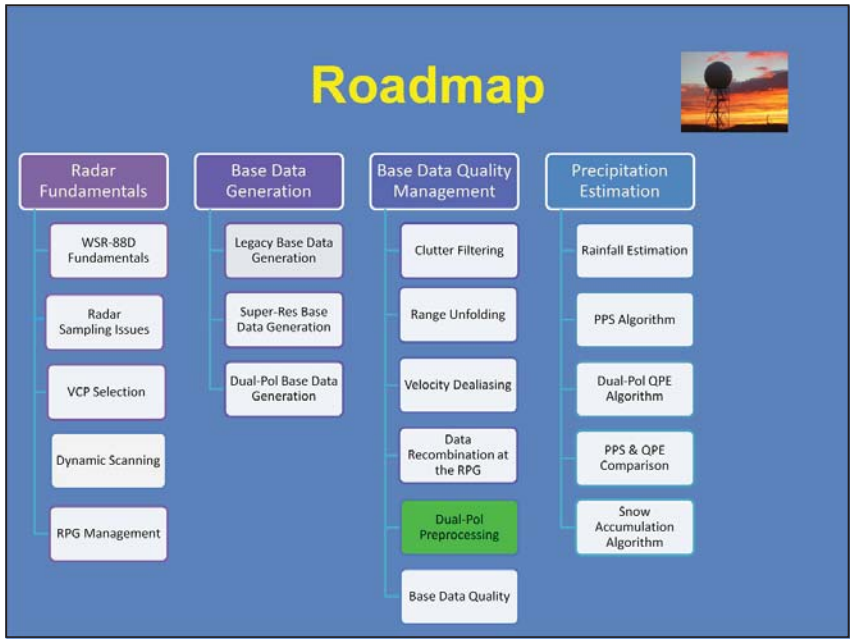


This step is to support RPG algorithms that require 1.0° azimuth by .25 km range resolution for velocity and spectrum width input.

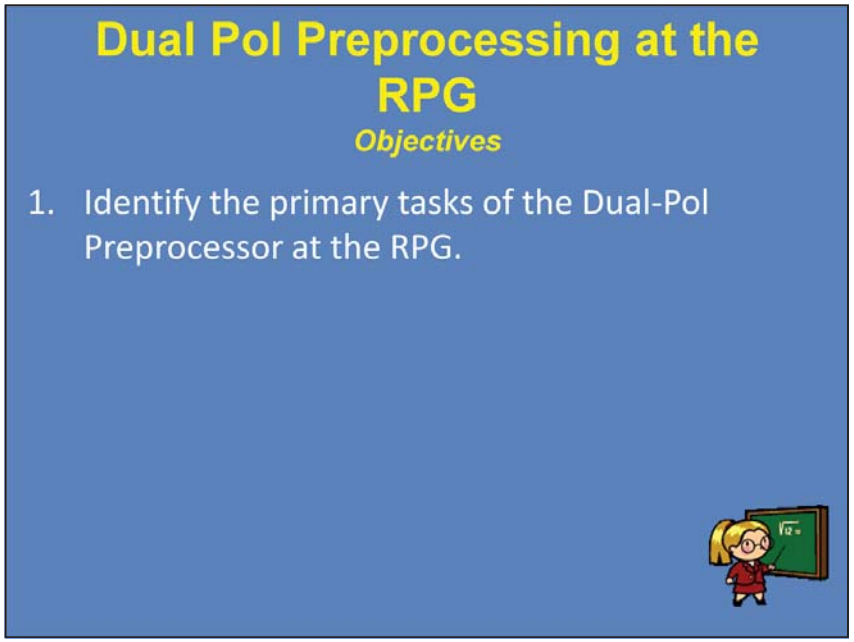
For velocity and spectrum width, the recombination process is a power weighted average of the 2 super resolution bins into the corresponding 1 legacy resolution bin. There is an additional step in the spectrum width processing to account for the variance of the two corresponding velocity values. If both bins are assigned No Data or RF, then the legacy bin is also assigned No Data or RF. If one of the two super resolution bins is assigned No Data or RF, then the remaining valid bin is assigned as the legacy resolution value.



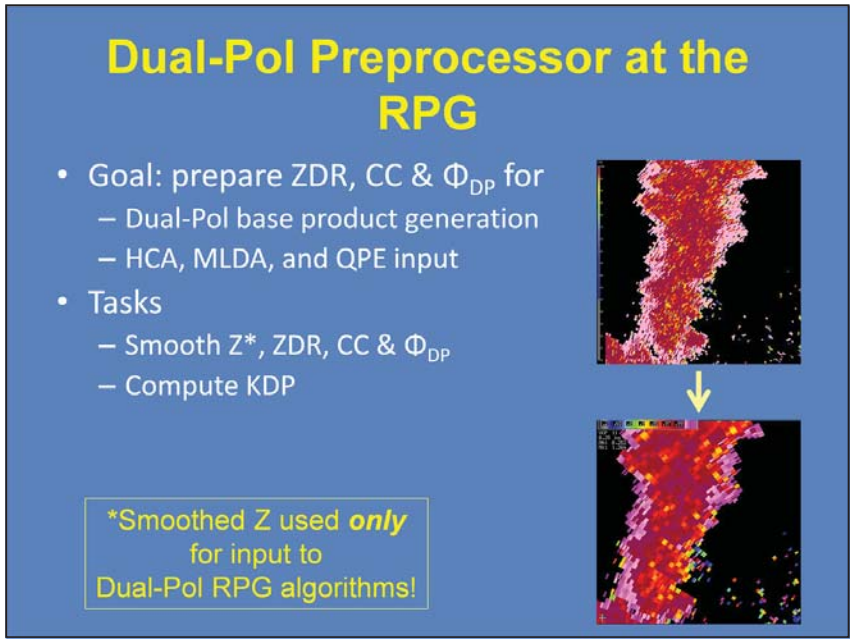
Welcome to Dual-Pol Base Data Preprocessing at the RPG

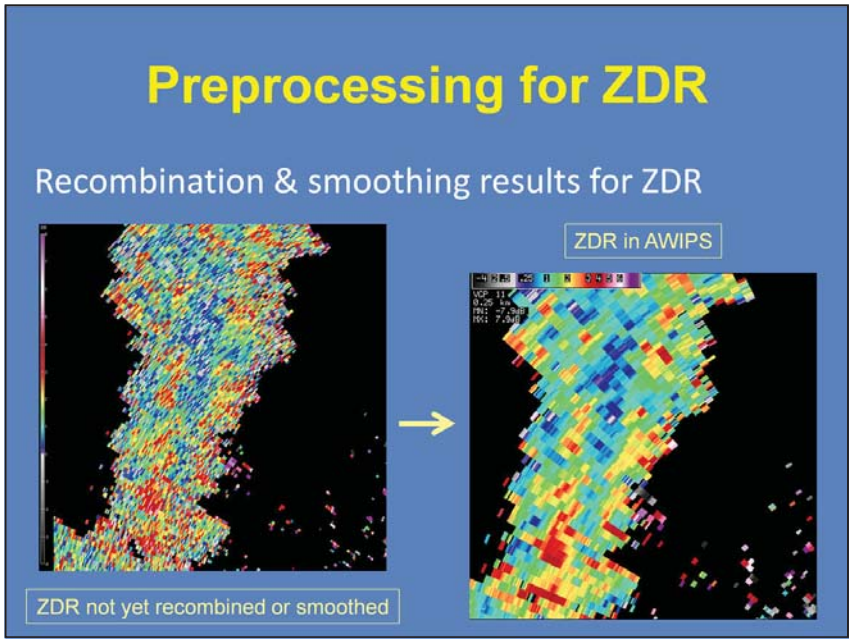


Here is the “roadmap” with your current location.



There is one objective for Dual Pol Preprocessing.





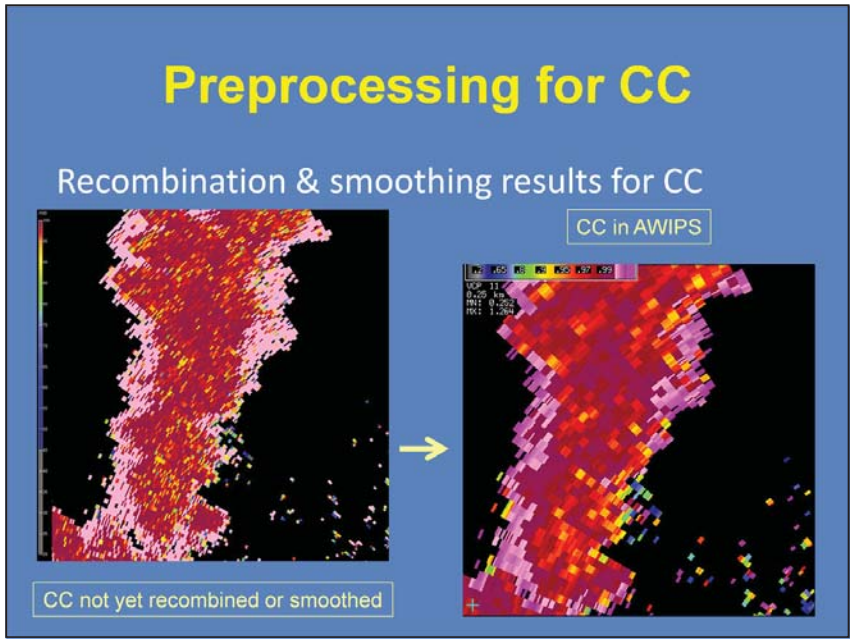
The RPG Dual-Pol Preprocessor is an RPG algorithm. It's purpose is to prepare the Dual-Pol base data for base product generation and for input into the RPG Dual-Pol algorithms: the Hydrometeor Classification Algorithm (HCA), the Melting Layer Detection Algorithm (MLDA), and the Quantitative Precipitation Estimation Algorithm (QPE).

The Dual-Pol base data sent from the RDA have a 0.5° azimuthal resolution, and are too noisy for human interpretation and algorithm input. For each range bin, the Dual-Pol base data are first recombined to 1.0° azimuth.

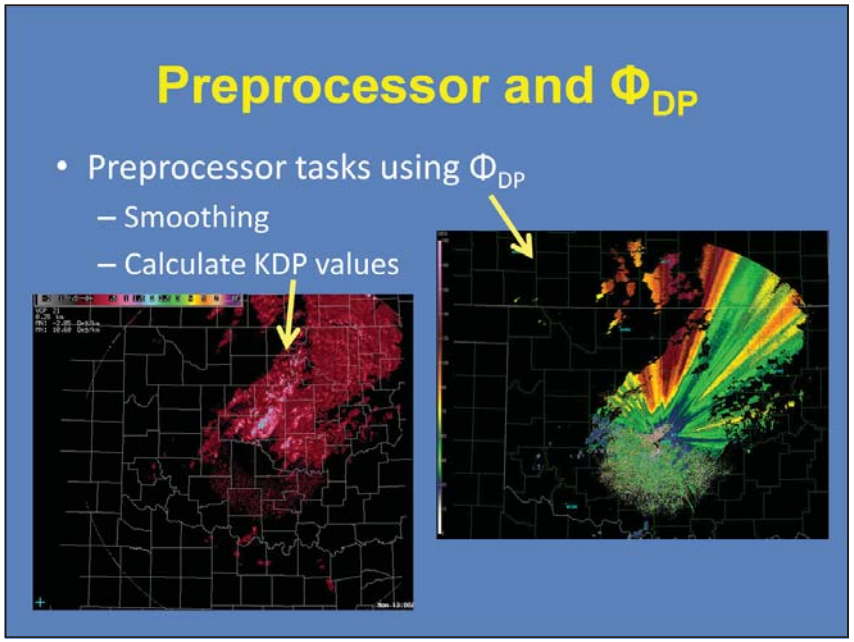
The Preprocessor smoothes Z, ZDR, CC & Φ_{DP} data along each radial. These smoothed Z data are only used for input to the Dual-Pol RPG algorithms. There is no change to the Z values used to generate all the reflectivity-based products that you are familiar with.

The remaining task for the Preprocessor is computing the Specific Differential Phase (KDP) value.

The example on the left is raw Differential Reflectivity, ZDR, from the RDA, at 0.5° azimuthal resolution and .25 km range resolution. It has not yet been recombined or smoothed. It is pretty noisy for even human interpretation. The image on the right is the same data displayed in AWIPS after recombination and Preprocessor smoothing. The Preprocessor smoothing technique applies a linear average to a segment (of varying length) of data along the radial. This average value is then assigned to the original range bin, which is at the center of the segment.



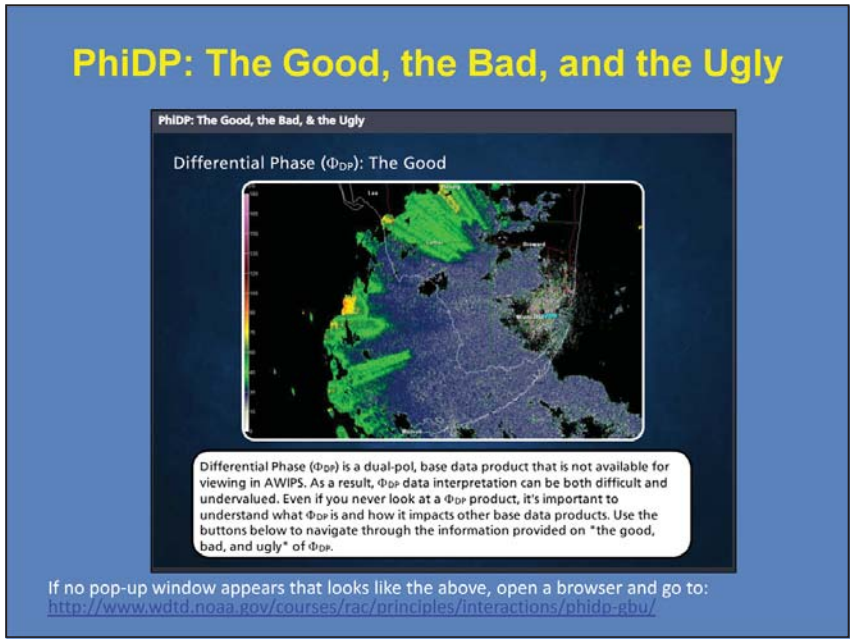
Here's a similar comparison for Correlation Coefficient, CC. We have raw CC on the left and the recombined and smoothed CC that becomes a product in AWIPS on the right. As with ZDR, the Preprocessor smoothing technique applies a linear average to a segment (of varying length) of data along the radial. This average value is then assigned to the original range bin, which is at the center of the segment.

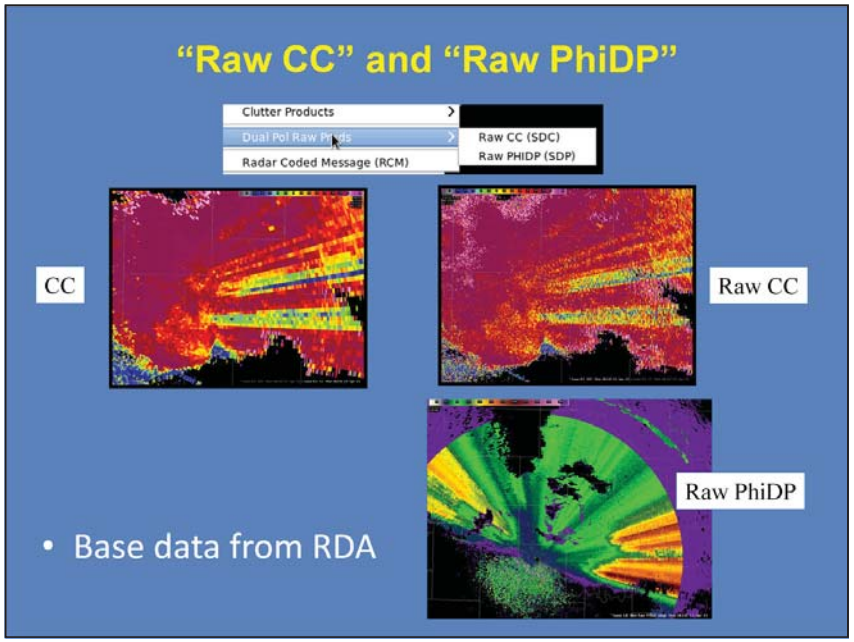


As with ZDR and CC, the Differential Phase, Φ_{DP}, base data are first recombined, then smoothed. On the right is an example of Φ_{DP} base data, not yet recombined or smoothed. This image is from GR Analyst, showing the raw Level II data.

Once the Φ_{DP} data have been smoothed, the Preprocessor then calculates Specific Differential Phase, or KDP. The KDP values are then available for generation of the KDP product (image on the left) and for input to the Dual-Pol algorithms.

These two images are a good example of why Φ_{DP} can be more difficult to interpret than KDP.





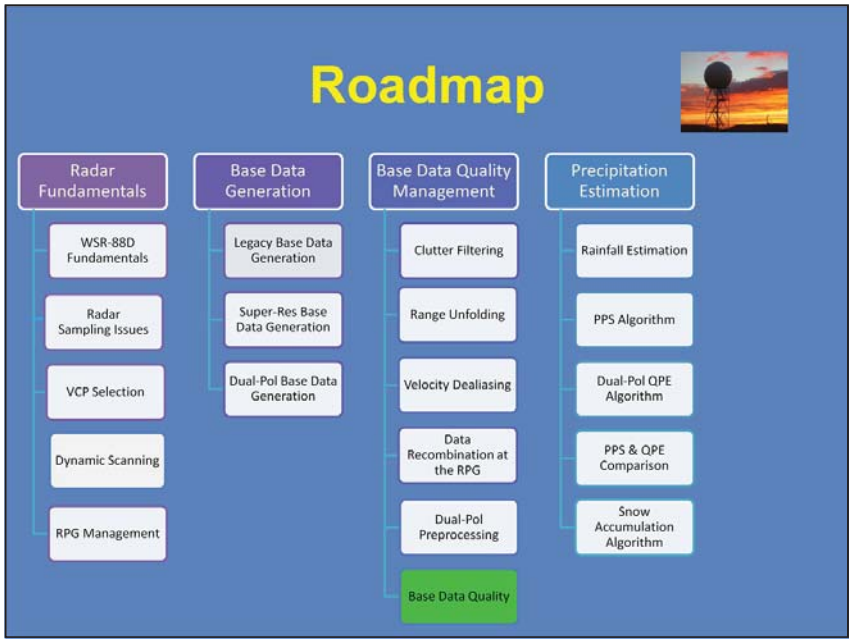
There are two dual pol products in AWIPS that ***not*** Preprocessed. They are both titled "Raw" to indicate that you are seeing ***only*** the base data sent from the RDA.

The Raw CC has higher azimuthal resolution (0.5º), however it is not a substitute for CC. The dual pol base data are noisier than the legacy base data, and the use of Raw CC is limited to (perhaps) earlier detection of a Tornadic Debris Signature (TDS).

The Raw PhiDP may be helpful for diagnosing dual pol base data quality issues. You will see both of these products described in the Products set of the course.



Welcome to WSR-88D Base Data Quality



Here is the “roadmap” with your current location.

Objectives



1. Identify areas of CMD false detections, and the trade off that can contribute to these false detections
2. Identify the strengths and limitations of the VCPs that are designed to mitigate RF data
3. Identify the impact of differential attenuation, non-uniform beam filling, and depolarization on the Dual-Pol products
4. Identify the "trade offs" involved with producing high quality base data vs. meeting operational constraints

Here are the 4 objectives for WSR-88D Base Data Quality, which will be taught in sequence during this lesson.

…And Now for the Really Cool Stuff about Doppler Weather Radar!



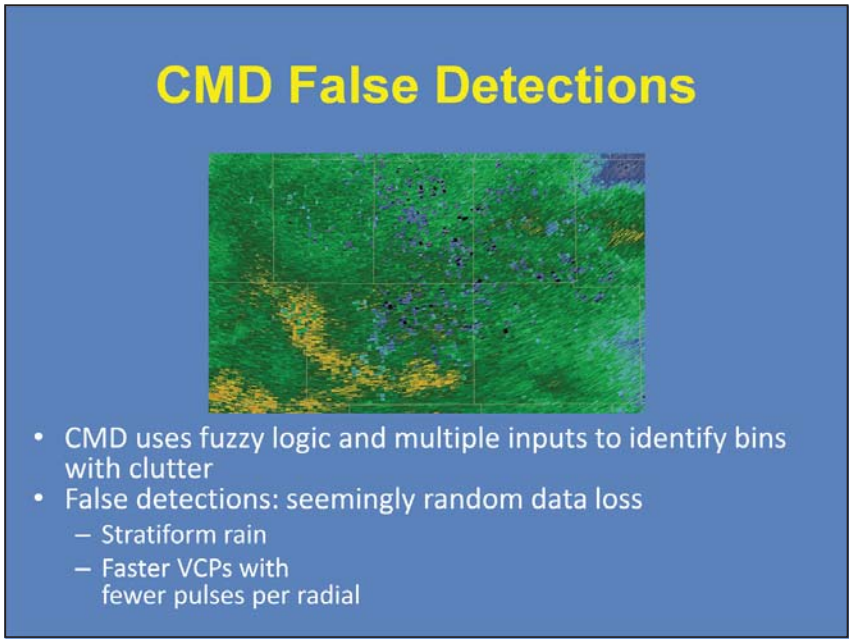
WSR-88D Data Quality



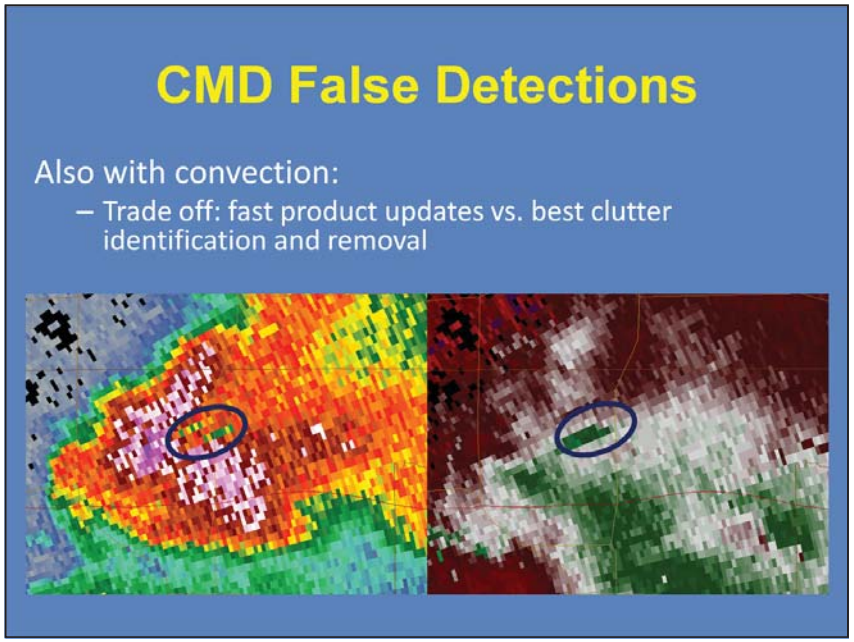


"If your base data ain't any good, nothin else is gonna be"

This lesson brings together all the previous Radar Principles concepts, exploring how you can optimize your base data quality, as well as recognizing the trade offs between optimal base data vs. operational needs for fast updates.

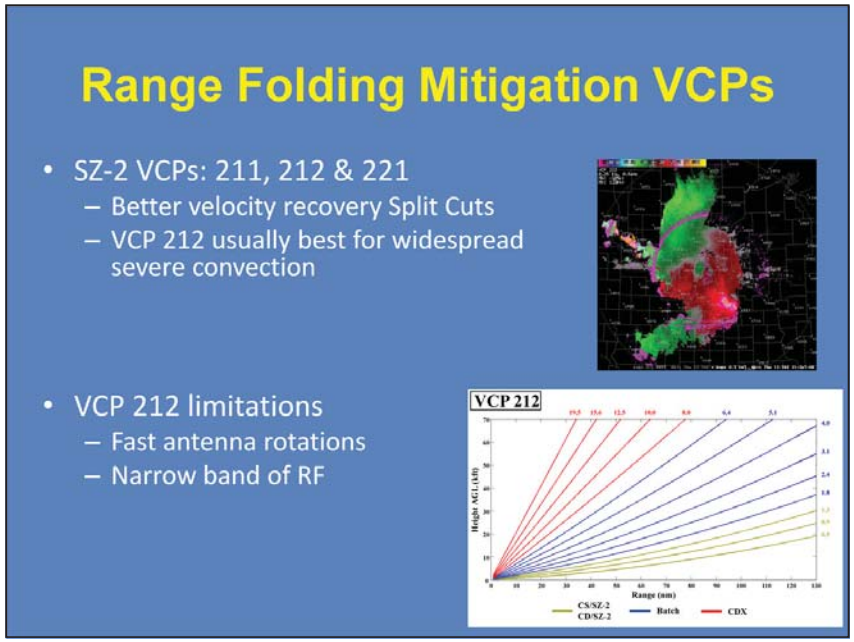


CMD is a complex algorithm with multiple inputs, and performs best with strong returned signal and lots of pulses per radial. CMD performance is most challenged with weak stratiform precipitation, especially when one of the faster VCPs is also being used. With these conditions, CMD is more likely to falsely identify bins without clutter. These false detections can result in noisy data with sporadic gates of signal removed that are not clutter.



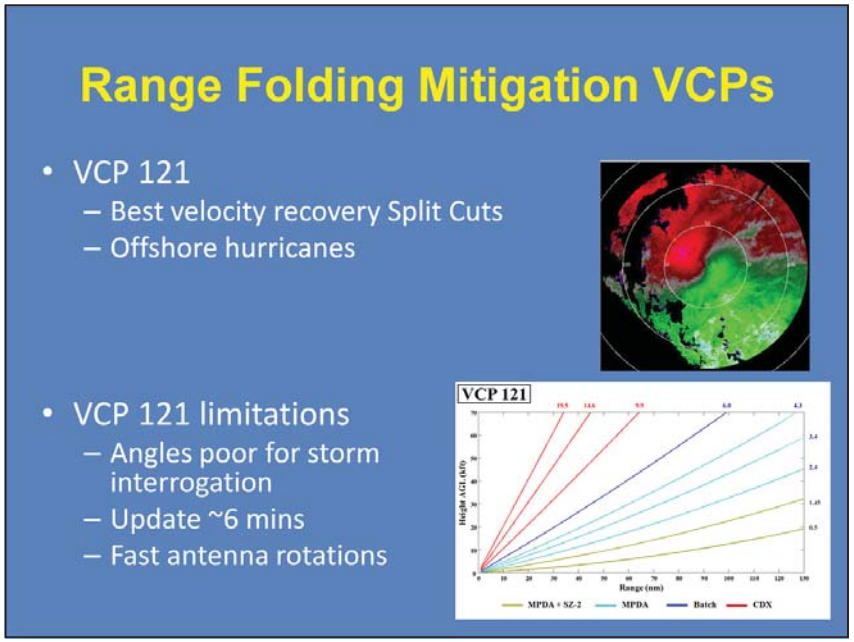
CMD false detections can even occur in or near convective storms. The result is seemingly random blocks of data loss that do not persist in space or time. The trade off at work here is the need for fast product updates with convective events, which means fewer pulses per radial, vs. the best performance of the clutter suppression algorithms: CMD for identification and GMAP for removal.

For severe convection, the need for VCP 12 or 212 overrides the need for perfect clutter suppression. For stratiform rain, VCP 21 is the better choice, providing more pulses per radial. Events between these two extremes are where the trade off can make the VCP decision tougher, though it is usually best to choose the VCP that is designed for the threat.



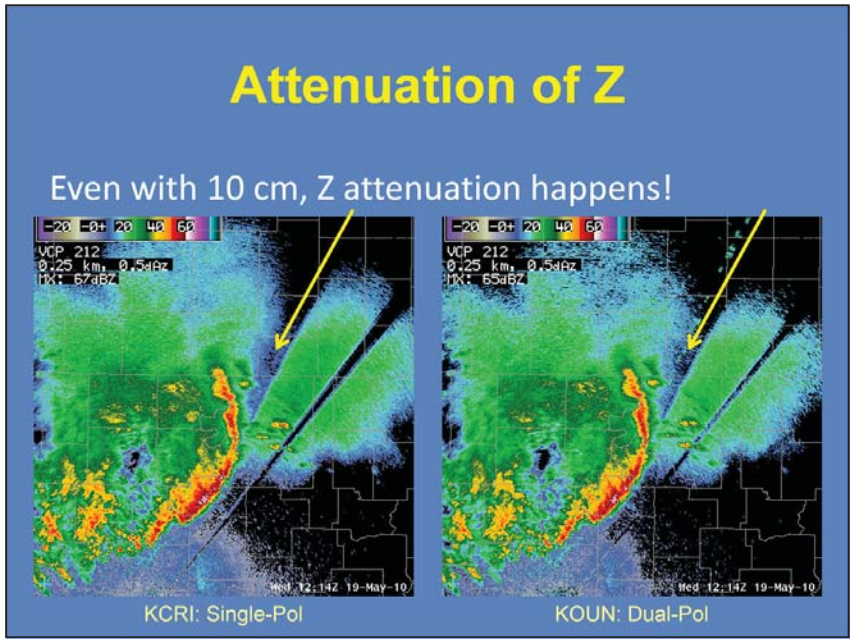
There is a class of VCPs that are designed to mitigate or minimize range folding in the velocity data. The first three in this group, VCPs 211, 212, and 221, share the fact that SZ-2 Range Unfolding is applied on the Split Cut elevations. The advantage of SZ-2 is much greater availability of velocity data, even with echo overlay conditions.

VCP 212 is the most frequently used VCP of this group, being a good choice for widespread severe convection. VCP 212 provides an update rate of about 4.5 minutes, with good vertical resolution, especially for the lower elevations. VCP 212 has fast antenna rotation rates, which can degrade data quality, especially when used for events other than severe convection.



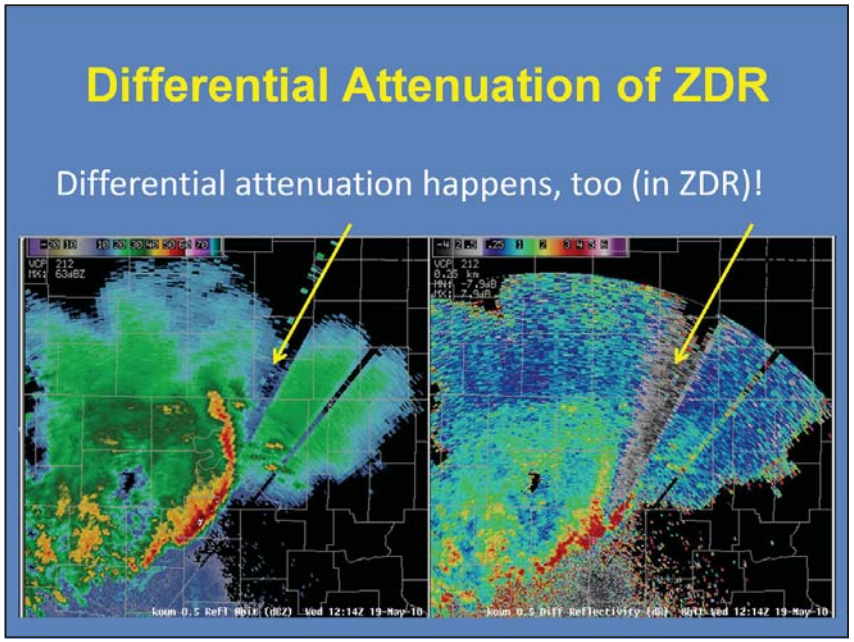
The other VCP in the group of range folding mitigation VCPs is 121. The primary benefit of VCP 121 is that for the lowest two elevations, nearly all the velocity data are recovered. VCP 121 is usually the best choice for offshore hurricanes or widespread non-severe convection.

VCP 121 is not a good choice for severe convection, with respect to both sampling and data quality. The elevation angles used by VCP 121 are not optimized for storm interrogation, especially at the lower levels. The update rate of almost 6 minutes is slow for severe convection. Since VCP 121 has multiple rotations at the same elevation, it has the fastest antenna rotation rates of any VCP, which can degrade data quality.



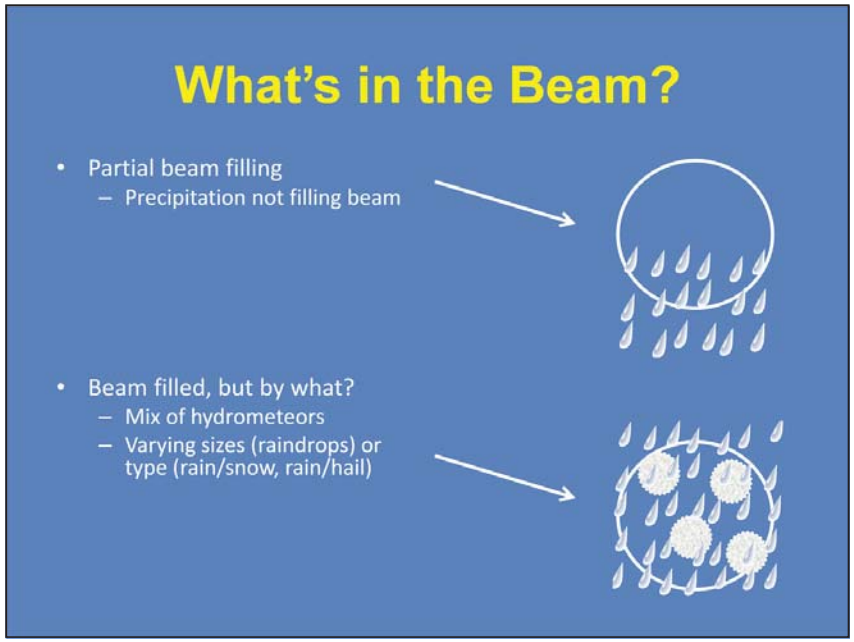
Attenuation of Z has always been with us, and will continue to be with Dual-Pol. We are very fortunate that the WSR-88D is a 10 cm radar, which attenuates much less than 5 cm radars. Of course, attenuation still happens and we need to take a look at how the Dual-Pol variables are impacted.

Here is a squall line sampled by two nearby WSR-88D radars, a Single Pol on the left and a Dual-Pol on the right. The squall line parallel is parallel to several radials and you can see the attenuation down radial in both of the Z products. Once the signal is attenuated, that loss cannot be recovered and propagates down radial.



With this same squall line case, there are very low ZDR values down radial (right image) that visually correlate with the Z attenuation. With ZDR, it is possible to have “differential attenuation”.

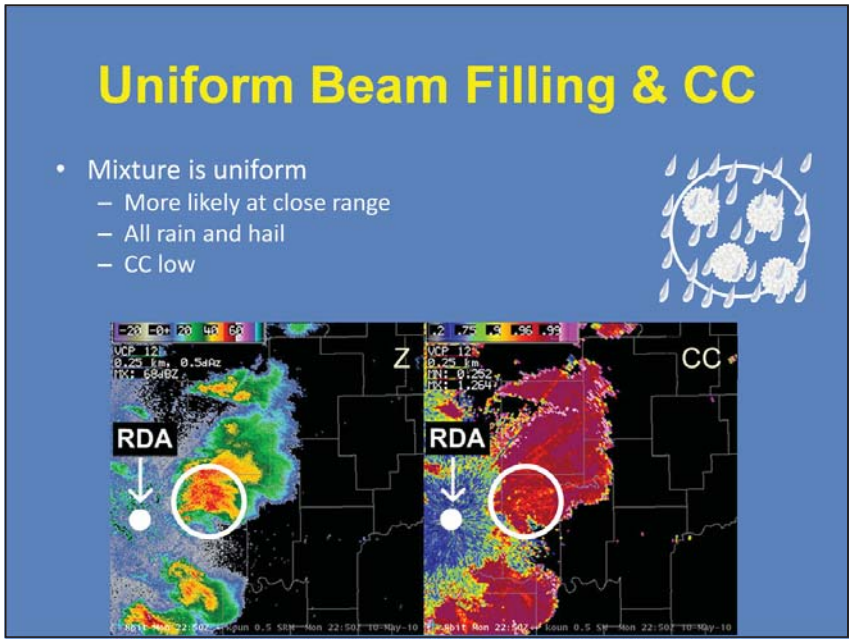
In this case, the beam encounters heavy rain with large drops. These large liquid drops results in more attenuation in the H direction compared to the V. With more signal loss in the H direction than the V direction, the ZDR is much lower that expected. For example, in areas where we know large to medium sized raindrops exist, large positive ZDR values are expected. Instead, ZDR values are generally negative, and extend down radial from the storm cores. Once the signal is attenuated, the loss in ZDR cannot be recovered and propagates down radial.



Non-uniform beam filling often occurs, and has always had implications with weather radar data quality, especially as range increases. However, the impacts have not been as apparent as it can sometimes be with Dual-Pol data. Even with Dual-Pol data, a specific type of non-uniform beam filling is required for the base data quality to be compromised. In these graphics the white circle represents the radar beam as if you were standing at the RDA looking outbound along a radial.

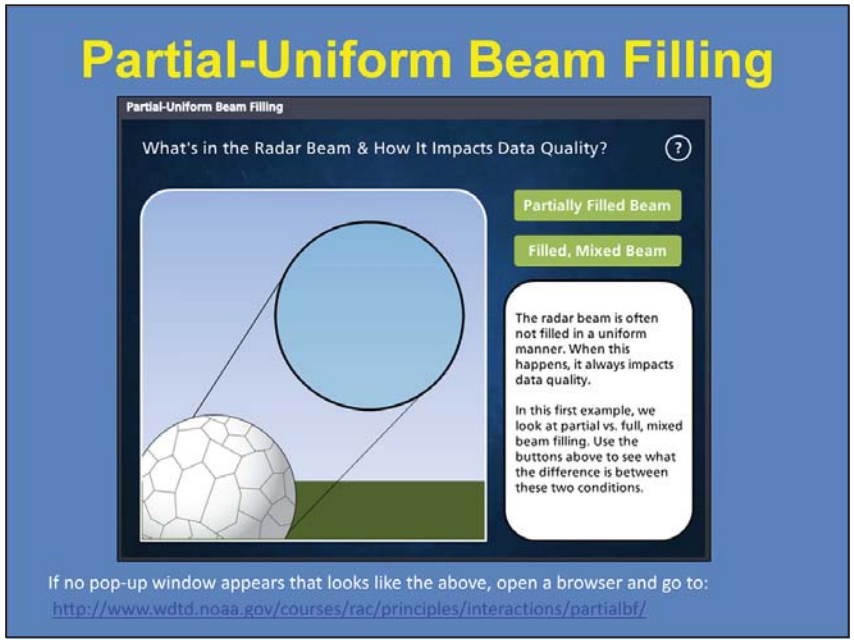
The top image represents partial beam filling, which is familiar, resulting in underestimated Z values.

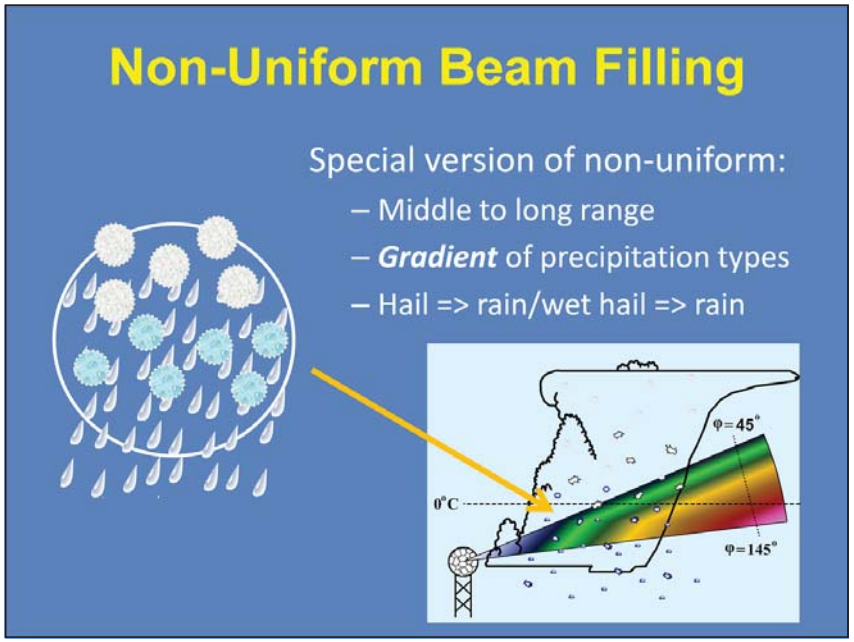
On the lower image, the beam is filled, but by a mix of precipitation sizes and types. The mix may be varying sizes of raindrops or hail stones or it could be varying precipitation types such as a rain/snow mix or a rain/hail mix. The nature of this mix and its distribution within the beam is relevant for Dual-Pol data quality.



It turns out that Dual-Pol products are negatively impacted by what is called Non-Uniform Beam Filling (NBF), and there are examples coming up. Though non-uniform beam filling in the literal sense occurs frequently, we also use NBF to describe a specific type of signature on Dual-Pol products that results from a specific type of non-uniform beam filling.

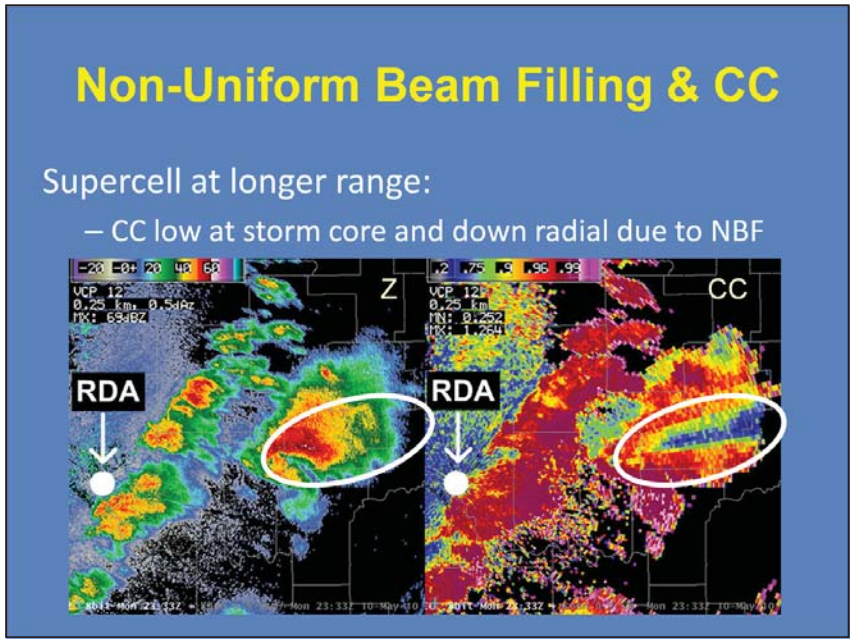
In this image, there is a supercell close to the radar and the associated CC product is on the right. In the circled area, the radar is sampling a mixture of rain and hail. Note that the CC values are lower within the core areas of the storm. This is expected when the radar samples a mixture of rain and hail that is relatively uniformly distributed across the radar beam cross section.

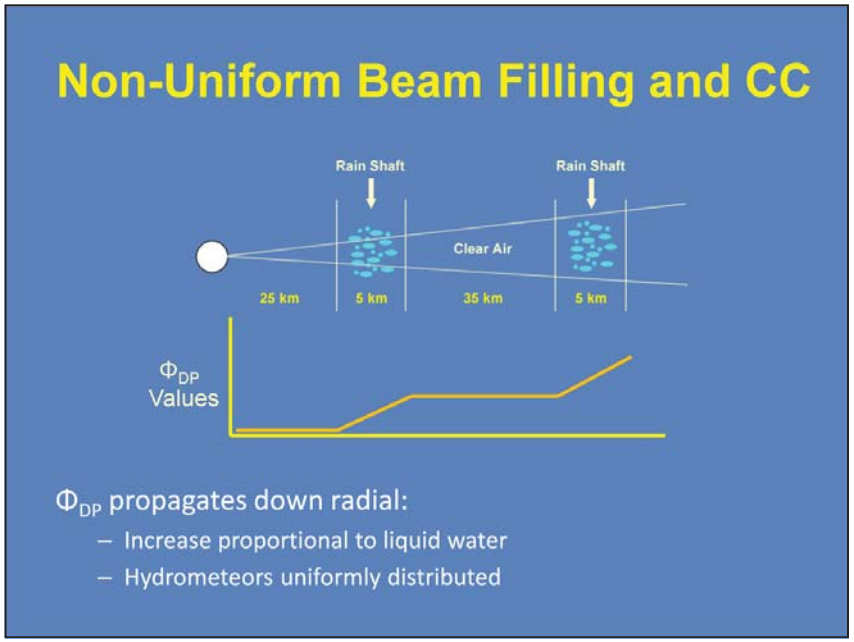




A non-uniform mixture can produce a gradient of precipitation types within the beam. This is more likely to occur at middle to long range. For example, the top of the beam may be sampling mostly hail, the middle sampling rain and wet hail, and the bottom of the beam sampling rain only. This gradient of precipitation types produces the version of non-uniform beam filling that is most likely to result in the Dual-Pol data artifact that we call Non-uniform Beam Filling (NBF).

Recall that Φ_{DP} contributes to both CC and KDP. With this gradient of precipitation type, this graphic represents the associated gradient of Φ_{DP} from the top to the bottom of the beam, if we had the vertical resolution to measure it. The gradient of precipitation types and the associated gradient of Φ_{DP} is the bottom line for low CC values locally and down radial.



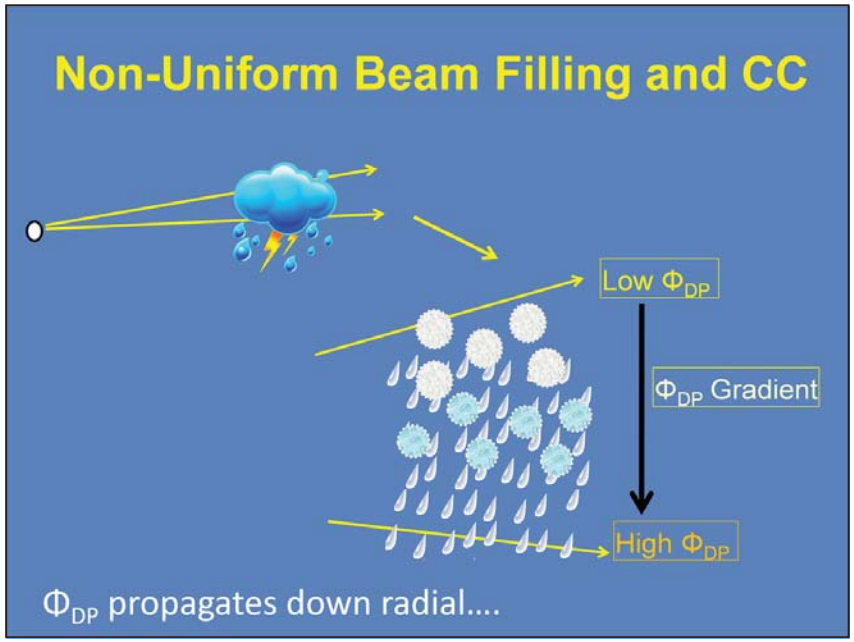


In the radar example, the supercell has moved to the east and is at a longer range, with the beam sampling a larger volume of the storm. There are radial swaths of low CC that originate from the storm core. This is an example of non-uniform beam filling and its impact on the CC product. This has impacts on other Dual-Pol products, with examples coming up.

By now you've probably seen a new window pop up with an animation of this event. You see Z and CC every other volume scan as the storm moves away from the radar. Once the storm is at a longer range, the non-uniform beam filling results in low CC values over a large wedge. This wedge persists even after the last frame of this loop. We know from the associated Z product that these low CC values do not make sense.

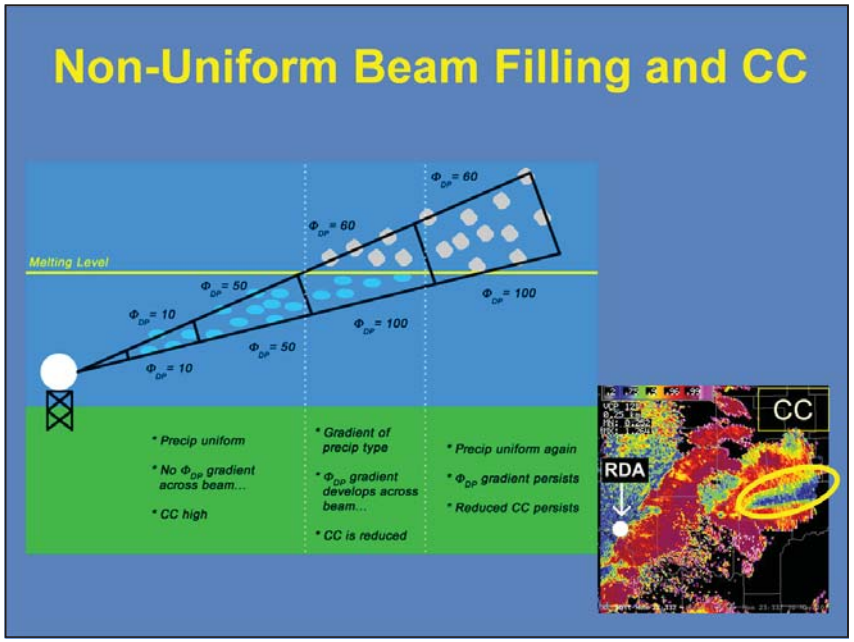
It is important to be aware of the potential for NBF on the CC products, because it has consequences for other Dual-Pol products.

First recall that Φ_{DP} values propagate down radial. When the hydrometeors are uniformly distributed, life is good. Φ_{DP} increases down radial as the beam passes through areas of pure rain. Since Φ_{DP} does not reset, the values are cumulative down radial.



When sampling a convective storm at longer range or a squall line along a radial, there is an increasing chance of capturing a gradient of precipitation types within the beam. At the top can be hail and/or graupel, while the bottom of the beam is sampling liquid drops.

This matters with Dual-Pol base data because the Φ_{DP} values are significantly different for ice than for liquid water. This is because Φ_{DP} responds to the amount of liquid water content. Though we cannot measure it, there is a significant gradient of Φ_{DP} within the beam. Since Φ_{DP} propagates down radial, this gradient does not "reset" down the radial.

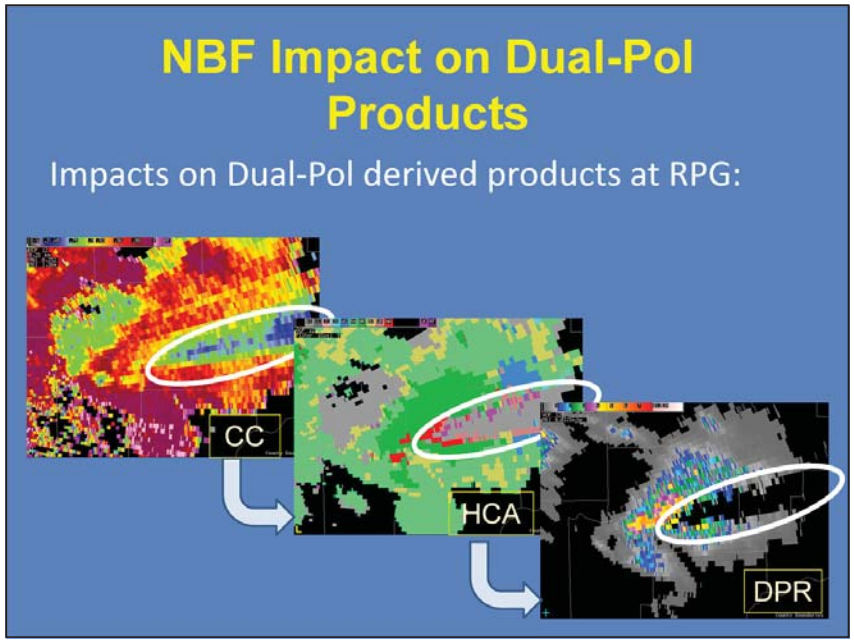


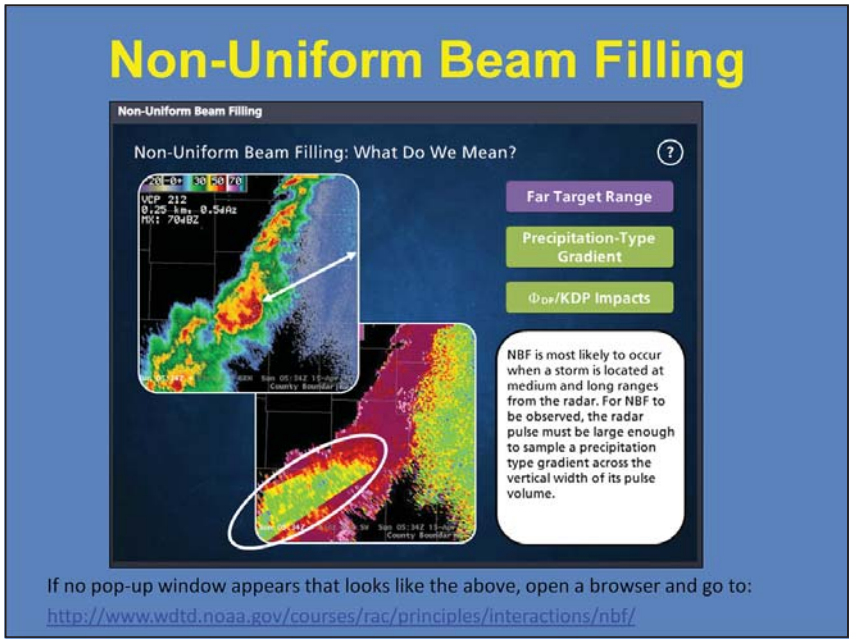
Here's a super simple example of what happens to the CC down radial with only four range bins. Note the Φ_{DP} values at the top of the beam and at the bottom of the beam for each of these bins.

The first two bins closest to the radar are sampling pure rain. Since the beam is uniformly filled, there is no Φ_{DP} gradient across the beam. CC values would be high for these bins.

The next bin encompasses the melting layer, with frozen hydrometeors at the top of the beam and liquid at the bottom of the beam. For this range bin, the CC value is low, there is a gradient of precipitation type across the beam, and thus a significant Φ_{DP} gradient across the beam.

Since Φ_{DP} does not reset down radial, this Φ_{DP} gradient will persist even as the beam is sampling uniform hydrometeors above the melting layer. This also means that the lowered CC will persist down radial.



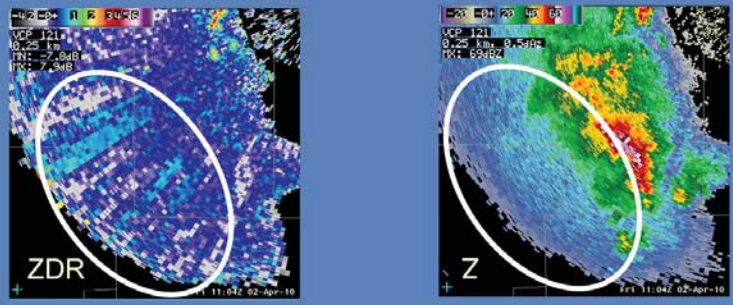


The artifact of a swath of low CC values due to NBF can be either easy to spot or subtle. By comparing it to other radar data and understanding the environment, you can ask yourself if the CC values make sense.

It is important to be mindful of this artifact because of the potential impact on the RPG algorithms that use CC as input. For example, CC affects the Hydroclass value that gets assigned, which then affects whether or nor rainfall is accumulated.

ZDR and De-Polarization

- Down radial from ice crystal regions
- Transient for any given radial
  - Canting of needles due to electrification
- Usually low operational significance

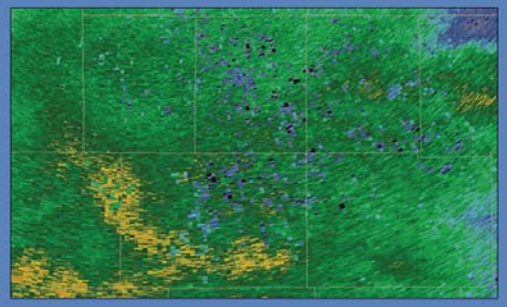


Depolarization will sometimes be apparent on the ZDR product. Depolarization means that the reflected energy from a particle switches polarization, from horizontal to vertical, vertical to horizontal, are maybe both.

Depolarization only affects the ZDR product. It appears as radial spikes which are transient with time. Though it may rarely occur in hail, depolarization is far more likely to happen in the upper regions of thunderstorms when the electrification causes canting of the ice crystals. Since the electrification varies with time, so does the impact of depolarization.

Fortunately, regions that are down radial from thunderstorm tops are usually of low operational significance. Be aware that this is a known ZDR data artifact, and is not a cause for concern.

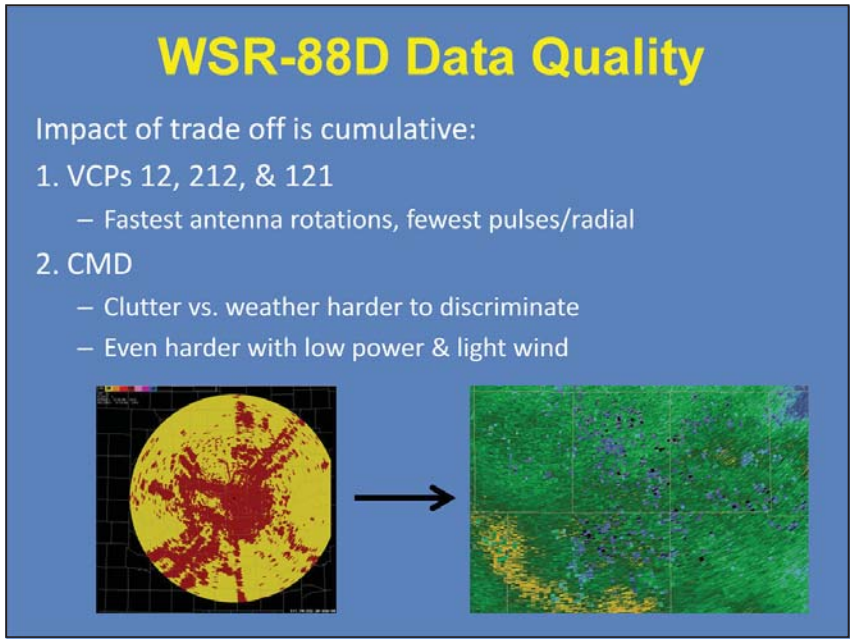
WSR-88D Data Quality



- Trade off:
  - high quality base data vs.
  - low level sampling and fast product updates
- Benefits of VCPs 12, 212, & 121 are obvious
  - Use 'em when you need 'em!

The WSR-88D is the most robust Dual-Pol Doppler radar fleet in the world. There is a big difference between the needs of operating a weather radar for research vs. operating one to meet the NWS mission.

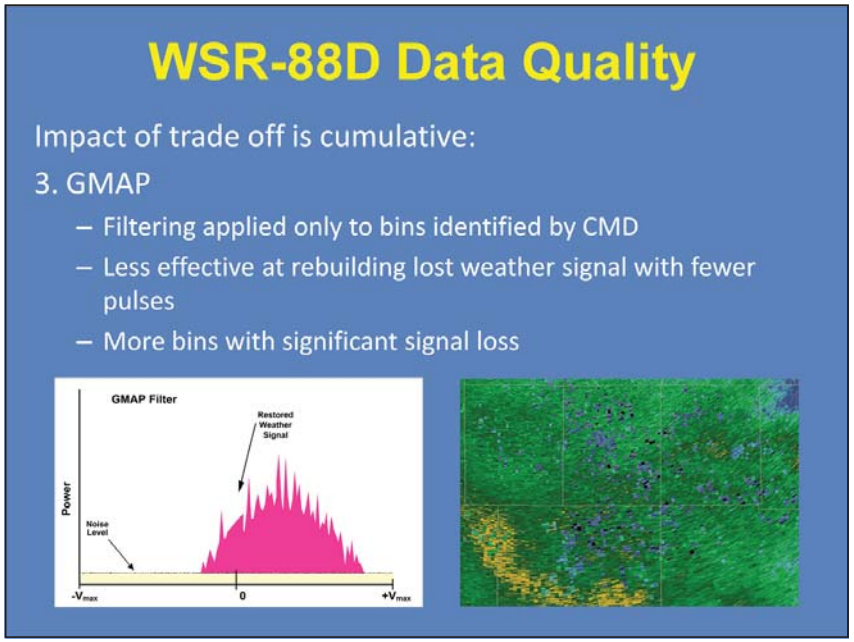
There is an inherent trade off between having the best quality base data, and meeting operational goals such as fast product updates and sufficient vertical resolution of elevation angles, especially at the lower levels. Our fastest VCPs, in terms of antenna rotation rates, have very obvious benefits for sampling severe convection (12 and 212) or offshore hurricanes (121). Do not be reluctant to use these VCPs when it is appropriate, though the fast antenna rotation rates push the limit of base data quality.



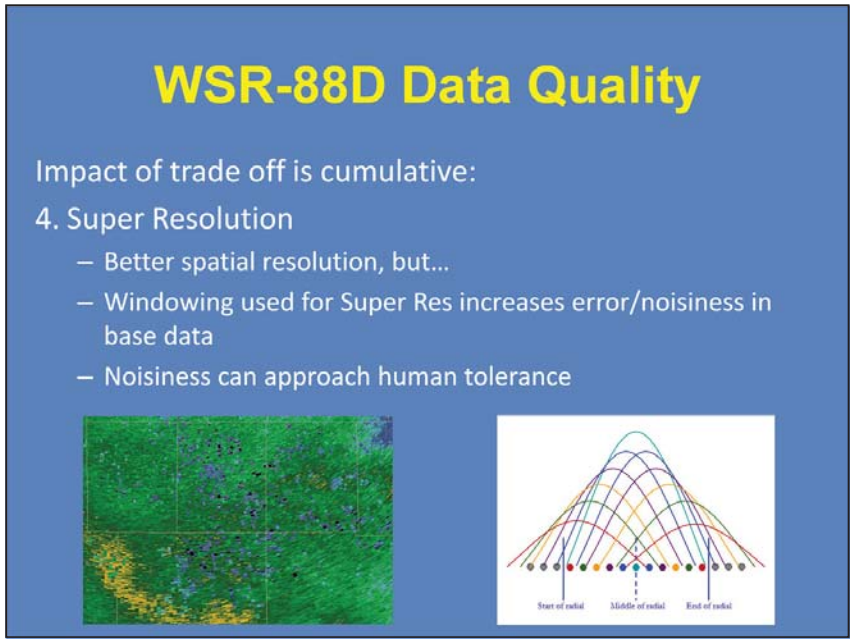
The impact of this trade off is cumulative, and I focus on VCPs 12, 212, and 121, since they have the fastest antenna rotation rates and thus the fewest number of pulses per radial.

For this group of VCPs, CMD can be less effective at discriminating clutter from weather. CMD false detections are also typically higher in areas of lower returned power and light winds. For example, using VCP 12 or 212 for a stratiform rain event is likely to result in more CMD false detections than when using VCP 12 for severe convection.

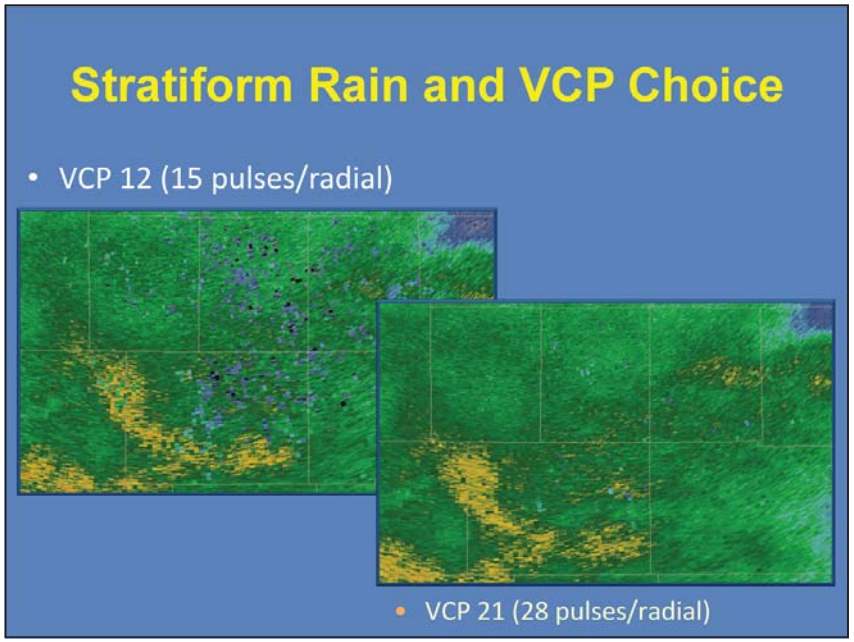
In this example, VCP 12 is active with no severe convection, only stratiform rain. Notice the seemingly random distribution of gates with reduced power (or data missing entirely) throughout the image. CMD is one of the contributors, by falsely identifying bins that contain clutter.



Once CMD identify which bins need to have clutter suppression applied, GMAP does the actual signal removal. It first isolates the clutter signal near zero velocity, then removes power just from the notch, or interval, around zero velocity. One of the strengths of GMAP is its ability to rebuild a lost weather signal across the zero velocity “gap”. However, this rebuilding is dependent on having a sufficient number of pulses remaining after the clutter portion has been removed. For the faster VCPs, there are fewer pulses per radial to work with, increasing the chance that GMAP will not be able to rebuild the weather portion of the signal. The result on the products is that more of the bins identified by CMD have data loss because GMAP cannot rebuild the weather signal.



Super resolution processing is another trade off. The benefit of better spatial resolution is obvious, but super resolution processing includes a windowing technique that introduces some error in the base data estimate. The cost is an increase in noisiness in the base data. The fewer the pulses per radial, the greater this noisiness can be.

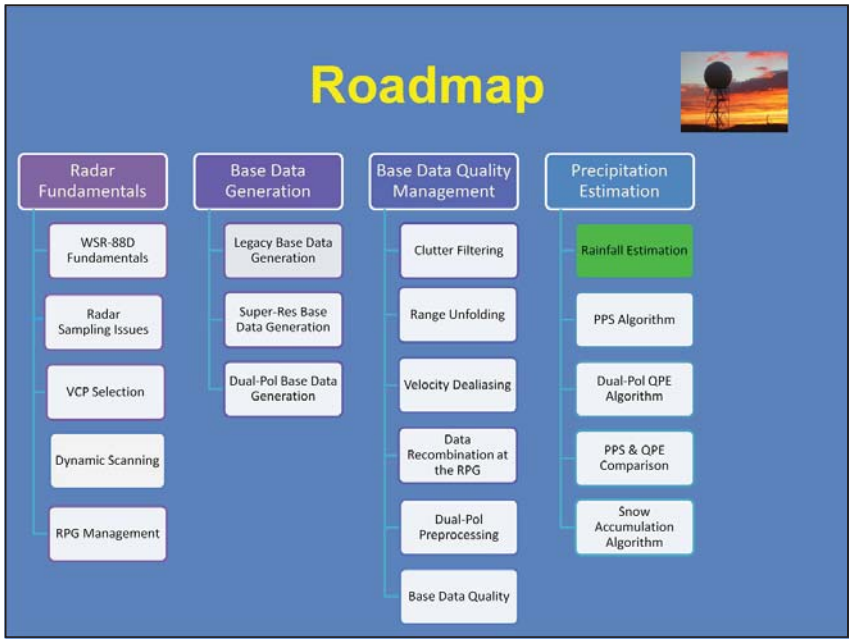


Now for the grand finale on the impact of this trade off between the need for high quality base data, along with fast product updates and low level sampling. This is a stratiform rain event, initially with VCP 12, which is the image on the left hand side. The staff noticed numerous gates of data loss over the rain area. These gates varied in space and time, but were numerous enough to cause concern. They decided to switch to VCP 21, which for the lowest elevation, has 28 pulses per radial. VCP 12 has 15 pulses per radial, almost half the number of VCP 21.

For severe convection, do not hesitate to use VCP 12 (or 212). That is what they are designed for, and the higher power returns will mitigate most of these errors. For stratiform rain, VCP 21 is recommended, as the larger number of pulses per radial will mitigate the errors due to weak signal plus light winds.



Welcome to Radar Rainfall Estimation Errors.



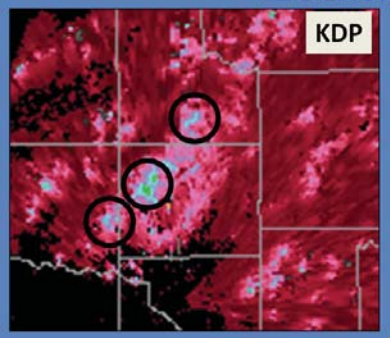
Here is the “roadmap” with your current location.

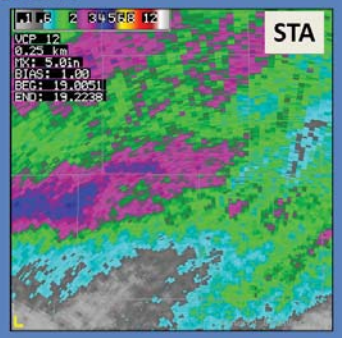
Radar Rainfall Estimation Errors
Objectives

1. Identify the impact of the given potential error sources associated with using a radar to estimate rainfall.
2. Identify the available radar base data inputs for generating a rainfall rate, and their limitations.

There are two objectives for Radar Rainfall Estimation Errors.

Heavy Rain Detection vs. Rainfall Estimation





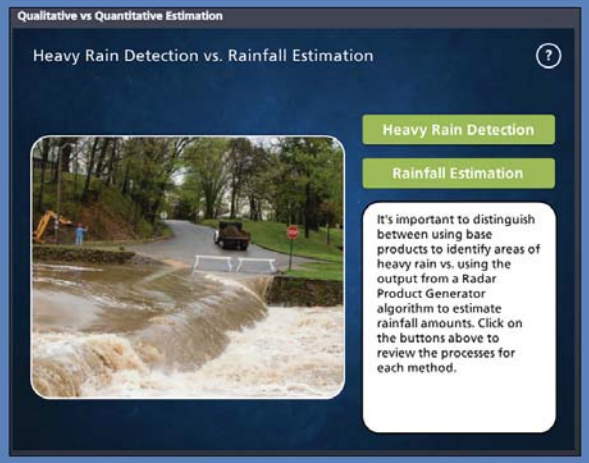
- Base Data (qualitative)
  - Patterns of heavy rain
- RPG Algorithms (quantitative)
  - Rainfall amounts over time

It's important to make a distinction between using the base products to identify areas of heavy rainfall vs. using the output from an RPG algorithm to estimate rainfall amounts.

Assume that we are looking at radar data that is below the melting layer where hydrometeors should be liquid. Among the available base products, Specific Differential Phase, KDP, is the best indicator of relative liquid water content. That means that observing KDP over time can identify areas that are most vulnerable to significant rainfall. Observing areal coverage patterns over time can reveal areas of potentially significant rainfall, growth and movement of precipitation, as well as linear vs. circular patterns.

On the other hand, using the output from the RPG algorithms to estimate rainfall amounts over specific durations first requires some situation awareness based on a thorough analysis of the base products. For example, do the locations of greatest rainfall make sense? Next, an understanding of the strengths and limitations of these algorithms is needed to use these products effectively.

Qualitative vs Quantitative Estimation



If no pop-up window appears that looks like the above, open a browser and go to:
http://www.wdtd.noaa.gov/courses/rac/principles/interactions/qual-vs-quant/

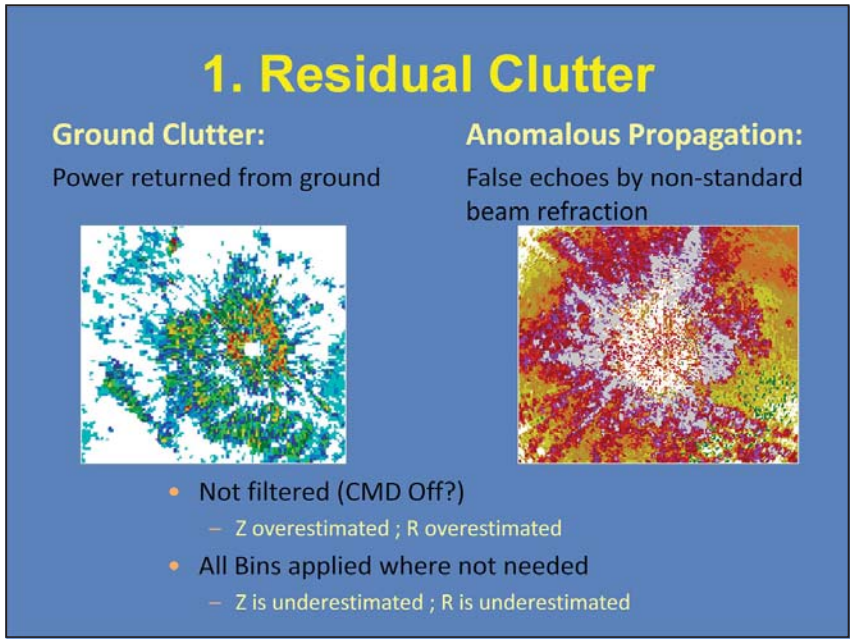
Issues of Using Radar to Estimate Rainfall

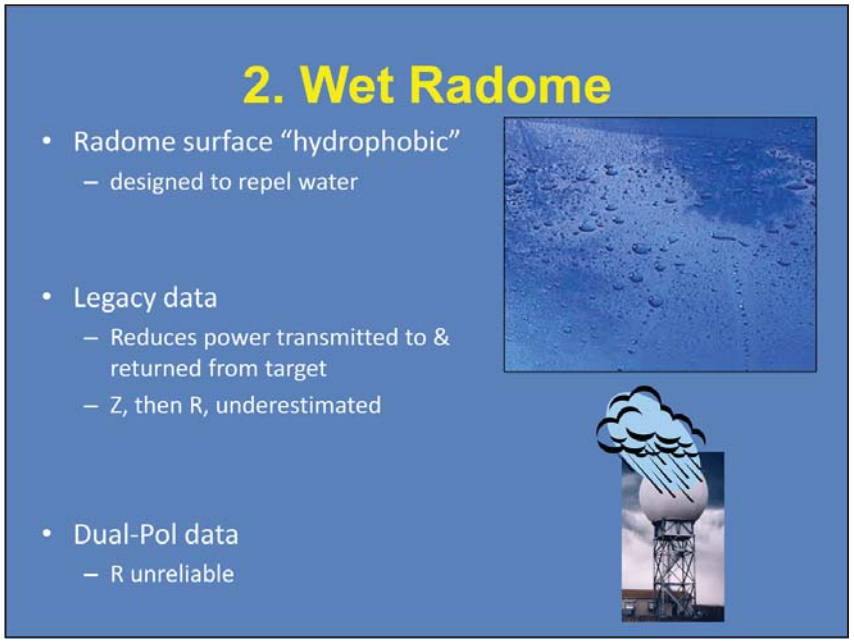
1. Residual clutter
2. Wet radome
3. Incorrect calibration
4. Below beam effects
5. Beam is in or above the melting layer
   - Sampling freezing or frozen precip
6. Partial or non-uniform beam filling
7. Coefficients & exponents for any rain rate equation vary for different areas





There are many issues that make using a radar to estimate rainfall amounts extraordinarily challenging, even before we discuss the algorithm design! Each of these items will be explored on the remaining slides.





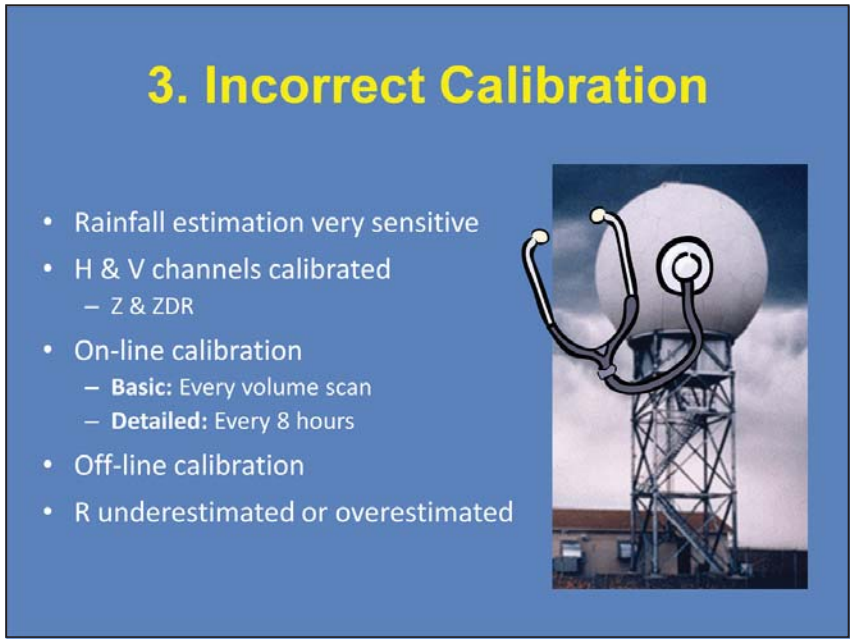
Residual Clutter describes unfiltered clutter, and can occur when clutter filtering has been turned off or when clutter remains after filtering is applied. Ground clutter indicates returns from certain ground based targets are always present, such as buildings and terrain. Technically, Anomalous Propagation (AP) clutter refers to any returns due to a superrefracting beam. AP often describes any clutter contamination that results when the beam is striking ground targets at varying ranges due to superrefraction. Unlike ground clutter, AP clutter is transient in space and time.

If no filtering is applied to either normal ground or AP clutter (CMD set to off?), Reflectivity values will be overestimated. Overestimated Z values will result in an overestimate of rainfall rate (R). Though the Dual-Pol RPG algorithms should better identify clutter and prevent it from being converted to rainfall, the potential for overestimation still exists.

It is also possible to apply clutter suppression where it is not needed through the use of All Bins suppression. Applying All Bins suppression can result in underestimated Reflectivity values, that then result in an underestimate of rain rate. The data quality impact of applying All Bins suppression unnecessarily to the Dual-Pol base data has the potential to negatively impact the quality of the dual-pol rainfall estimates.

A wet radome can cause multiple issues with base data. The radome surfaces are designed to be "hydrophobic", repelling water like wax on a car, in order to prevent water coating the surface. It still happens, of course, but the condition is usually transient.

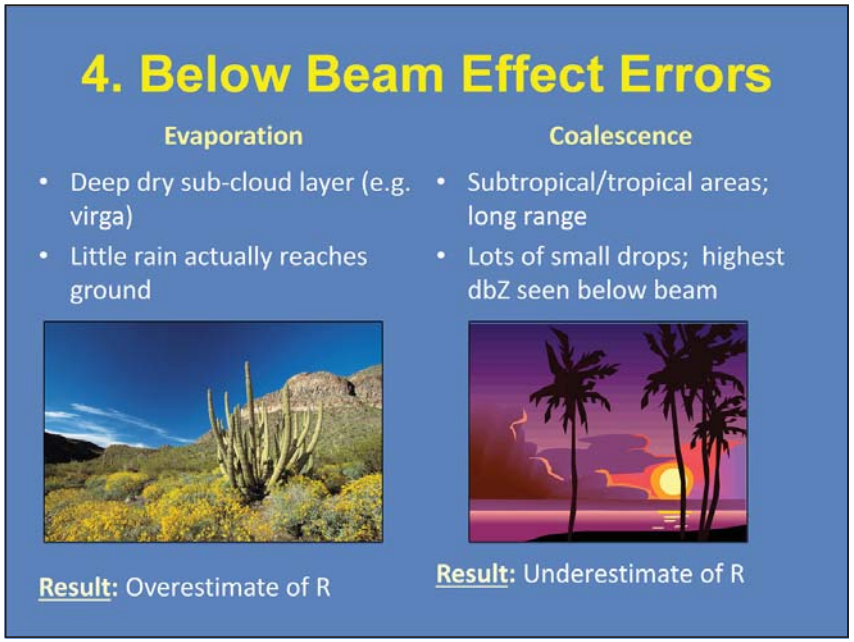
A water coating on the radome reduces the amount of transmitted power. For the legacy base data (i.e., the horizontal channel), the reduction in transmitted power results in a reduction in returned power, which leads to an underestimate of both reflectivity and rain rate. For the dual-pol data, which is based on both horizontal and vertical channels, the impact of a wet radome on R can be either underestimated or overestimated.



Rainfall estimation is particularly sensitive to calibration errors. A valid Z value is dependent on a well calibrated horizontal channel, while a valid ZDR value depends on each of the horizontal and vertical channels being well calibrated.

There are two different types of calibration: on-line and off-line. One on-line calibration is performed at the end of every volume scan, as the antenna is moving back to 0.5° to begin the next volume scan. A second on-line calibration is performed every 8 hours. This 8 hour "Performance Check" also includes multiple tests to assess the "health" of the radar. Off-line calibration requires the technicians to have control of the radar for a more lengthy process.

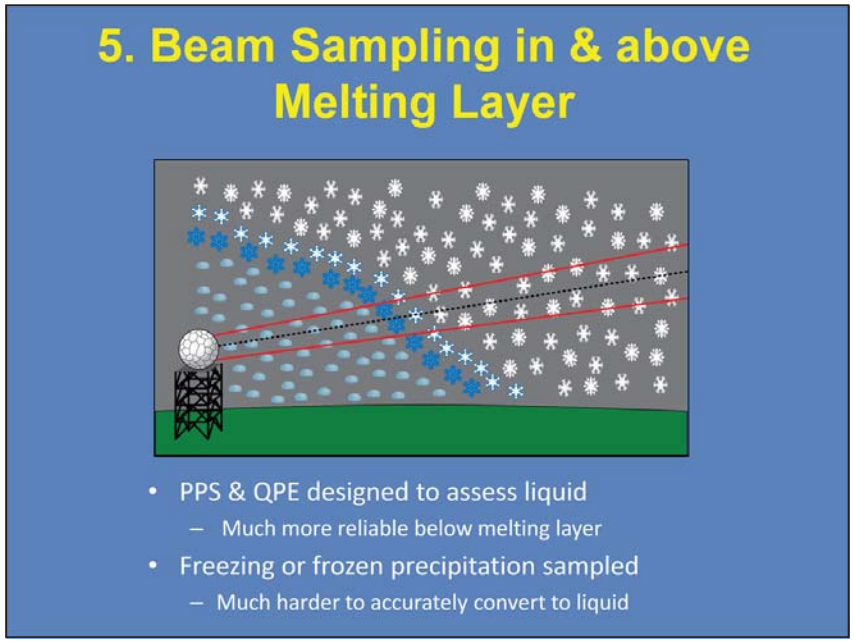
With respect to rainfall estimation, calibration errors can result in either an underestimate or an overestimate.



Below beam effects are a fundamental challenge when using a radar to estimate rainfall. We must remember that rainfall estimates are based on the hydrometeors that were sampled within the radar beam. Depending on range, it may be a very long way down to reach the ground! Consideration of the layer below the beam is very important.

Evaporation below the beam can occur in the presence of a deep, dry sub-cloud layer. The desert southwest often has convection with significant below beam evaporation. This layer will evaporate some (or even all) of the rain leaving the cloud, causing a smaller amount rain to reach the ground. As a result, overestimates of rainfall by the radar are likely when sub-beam evaporation is possible.

Coalescence below the beam occurs primarily in subtropical or tropical areas and at long distances from the radar. Where warm rain processes are dominant, a large number of small drops collide and lead to raindrop growth. Except for precipitation at short ranges, the largest drops are often too low to be sampled by the radar beam. As a result, underestimates by the radar are likely when coalescence occurs beneath the radar beam.



It is important to remember that both the RPG rainfall algorithms are designed to estimate liquid rainfall. Algorithm performance is much more reliable at locations where the radar beam is intercepting liquid hydrometeors below the melting layer. Within the melting layer and above, hydrometeors are frozen (or freezing), and converting to liquid rainfall on the ground really requires a snow conversion algorithm, which will be discussed in a later lesson in this topic.



The Probert-Jones radar equation converts returned power to reflectivity with the assumption that the beam is uniformly filled with scatterers. You can probably guess how hard that condition is to meet, especially as the beam increases in size with range.

Here we look at two specific cases that do not meet that condition. The first is partial beam filling, where only a portion of the beam is sampling precipitation. The radar beam spreads with range, increasing the chance that targets may only partially fill the beam. When partial beam filling happens, the reflectivity and rainfall rates are both underestimated. Since the beam volume is greater than the actual precipitation volume, the areal coverage is overestimated, also.

Literally speaking, non-uniform beam filling is likely a common phenomenon. However, a particular type of non-uniform beam filling has a significant impact on the dual-pol products. When a gradient of precipitation types exists across the beam, CC can be underestimated down the radial from that gradient. Active precipitation areas down radial from the non-uniform beam filling may not be converted to rainfall. There was an example of this problem in an earlier lesson in this topic.

7. Coefficients & Exponents Vary

$$Z = 300R^{1.4} \qquad R(Z) = (0.017)Z^{0.714}$$

$$R(Z, ZDR) = (0.0067)Z^{0.927} ZDR^{-3.43}$$

$$R(Z, ZDR) = (0.0142)Z^{0.77} ZDR^{-1.67}$$

$$R(KDP) = 44.0|KDP|^{0.822} sign(KDP)$$

Equations that convert to rainrate are empirical
- Vary with dropsize distribution

You will see these equations again later in the lessons that discuss the specific rainfall estimation algorithms. For now, just understand that these rain rate equation use empirical coefficients and exponents. The applicability of these equations varies when dropsize distributions differ significantly from those used to determine these values.

RPG Algorithms to Estimate Rainfall



- Legacy PPS
  - Based on Z,V,SW
  - Rainfall rates based on Z
- QPE
  - Based on R, V, & Dual Pol
  - Rainfall rates based Z, Z&ZDR, or KDP

Note: ***either*** algorithm overestimates or underestimates for ***multiple*** reasons
Underestimate ≠ "cold" radar & Overestimate ≠ "hot" radar

There are two RPG algorithms designed to estimate liquid rainfall at the surface. The first is called the Legacy Precipitation Processing SubSystem (PPS), and it has been in place since the original deployment of the WSR-88D. The PPS has seen substantial design changes over the years, as with many of the RPG algorithms. It's inputs are reflectivity, velocity and spectrum width. The PPS computes rainfall rates based on Z, from a choice of Z-R relationships.

The second rainfall algorithm is the Quantitative Precipitation Estimation (QPE) algorithm. The QPE algorithm was recently fielded as part of the Dual-Pol upgrade. Along with reflectivity, and velocity, QPE uses the dual-pol base data and related algorithms as inputs.

There is a misconception that rainfall estimate errors from either of these algorithms are directly related to radar calibration. Given the number of potential errors possible with either algorithm, forecasters should not assume underestimates imply a "cold" radar, or that overestimates imply a "hot" radar.

Z is Estimated





- **Unknown**: Dropsize Distribution
- **Known**: Returned Power

Z ***estimated*** from P-J radar equation

$$P_r = \left[\frac{P_t G^2 \theta^2 \, \mathrm{H} \pi^3 \mathrm{K}^2 \, L}{1024(\ln 2)\lambda^2}\right] \times \frac{Z}{R^2}$$

Reflectivity, Z, is an obvious input for estimating rainfall. It is important to remember how Z itself is estimated, and its associated limitations.

If any given dropsize distribution could be measured, Z (as well as the rainfall rate, R) could be computed directly. However, dropsize distribution cannot be measured directly and is unknown. What is known is the power that is returned to the radar. Based on that returned power, reflectivity is an estimate that comes from the Probert-Jones radar equation.

Rain Rate Equation Input: Z

- Z is estimated from P_r
- Confidence in Z for rainfall estimation dependent on
  - Range
  - Dropsize distribution
  - Calibration
  - Non-uniform or partial beam filling
  - Attenuation



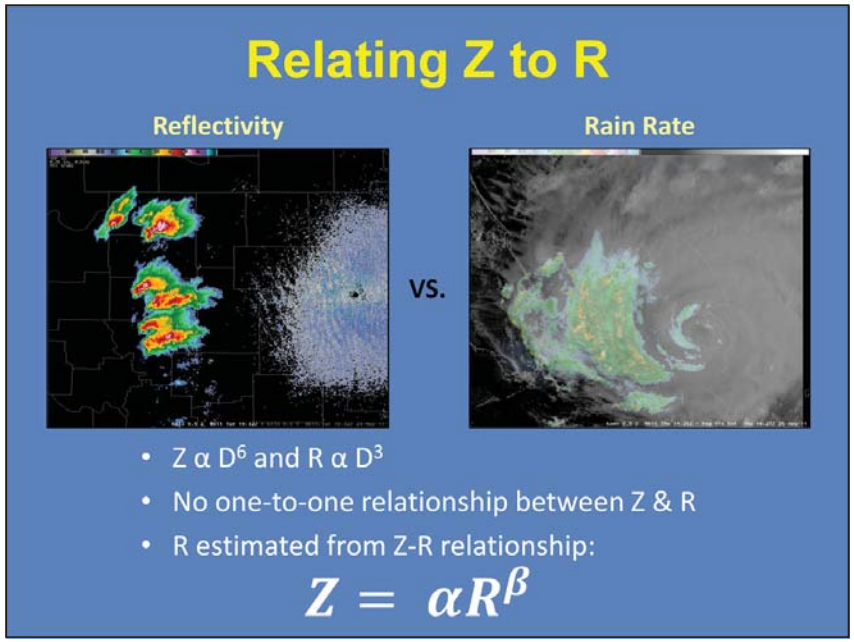
Estimating rainfall from radar is a very complicated business. It's amazing it works as well as it does! Remember that there are potential errors in the Z values even before a rain rate is calculated. Here are the limitations on the validity of the Z value itself.

As range increases, so does the size of the volume that is sampled by the beam. At far ranges, the chance that the radar beam is sampling above the melting layer (i.e., all snow and/or ice crystals) increases dramatically. The Probert-Jones radar equation assumes liquid (not frozen) hydrometeros, affecting the accuracy of Z.

Variations in dropsize distribution can occur at multiple spatial scales, from within the radar umbrella to within a sample volume.

A Z value from a poorly calibrated radar can introduce significant errors in the rainfall estimate.

When the beam is partially or non-uniformly filled by hydrometeors (which is typical at longer ranges), the Z value may not be representative. Though the WSR-88D is a 10 cm radar (yeah!), signal attenuation still happens with heavy rain, resulting in significant underestimates of Z down radial.



Both Z and R are dependent on the dropsize distribution, which is unknown, but with different dependencies. Z is proportional to the drop diameter to the sixth power, while the rainfall rate is proportional to the drop diameter to the third power. There is no one-to-one relationship between Z and R. As a result, R is estimated though a Z-R relationship, expressed as a power law equation. Here, alpha and beta are empirical constants that change depending on the meteorological event being analyzed.

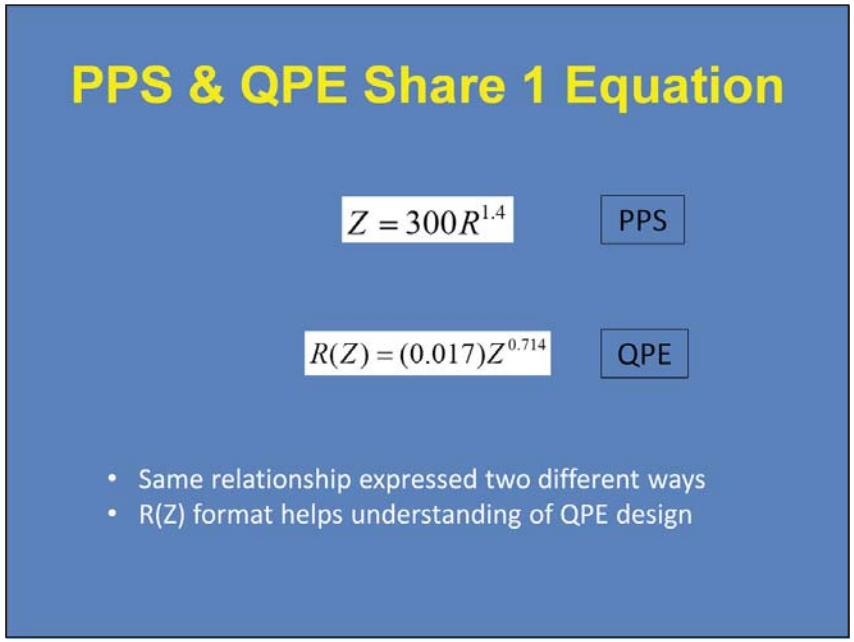
PPS: Z-R Relationships are Editable

Relationship	Optimum for:	Also for:
Convective $Z = 300\ R^{1.4}$	Deep convection	non-tropical convection
Tropical $Z = 250\ R^{1.2}$	Tropical convective systems	
Marshall-Palmer $Z = 200\ R^{1.6}$	General stratiform	
East-Cool Stratiform $Z = 130\ R^{2.0}$	Winter *rain* east of continental divide	Orographic rain east
West-Cool Stratiform $Z = 75\ R^{2.0}$	Winter *rain* west of continental divide	Orographic rain west

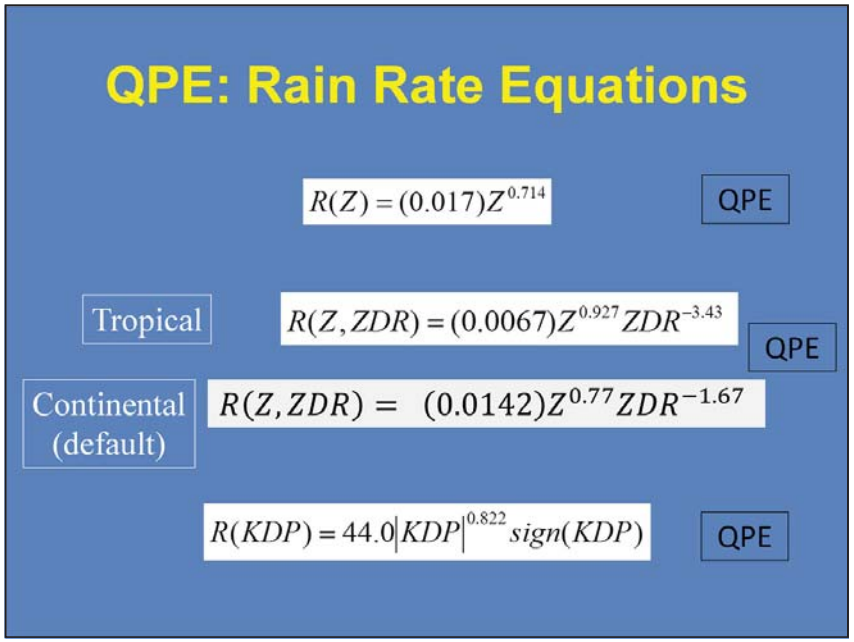
All used for estimating *liquid* precipitation

The Precipitation Processing Subsystem (PPS) relies solely on Z for rainfall estimation, by applying a Z-R relationship. There are five Z-R relationships that have been developed over the years, and they can be selected by editing the multiplier and the coefficient at the RPG. This table summarizes the available WSR-88D Z-R relationships, along with their optimal environments.

First, notice the differing empirical constants used in each relationship. Second, even though two of these relationships are optimized for wintertime, the goal is still to estimate liquid precipitation on the ground.



The legacy PPS and the Dual-Pol Quantitative Precipitation Estimation (QPE) algorithm share one equation for converting Z to rainfall rate. The first equation shows the more familiar format, solved for Z. The second equation is the same as the first, just rewritten in a format that is solved for R. We introduce the R(Z) format here to help your understanding of the QPE design (coming up!). Though you've just seen five different Z-R relationships available for use with the PPS, the only Z-R relationship used by QPE is Z=300R**1.4.

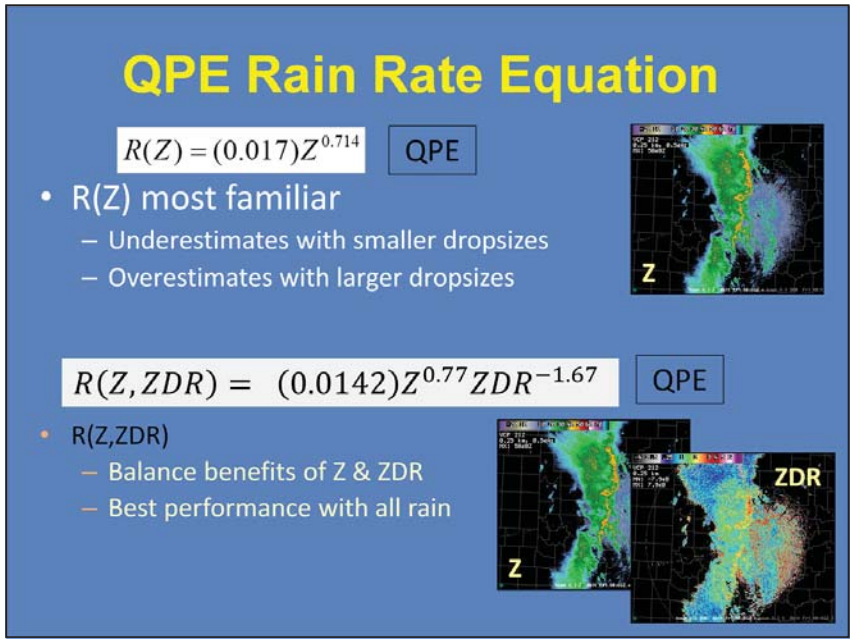


There are four equations for computing rain rate used by QPE. How QPE determines which of these equations to use will be explored in a later lesson.

The first equation is the Z=300R**1.4 relationship solved for R, as you saw on the previous slide.

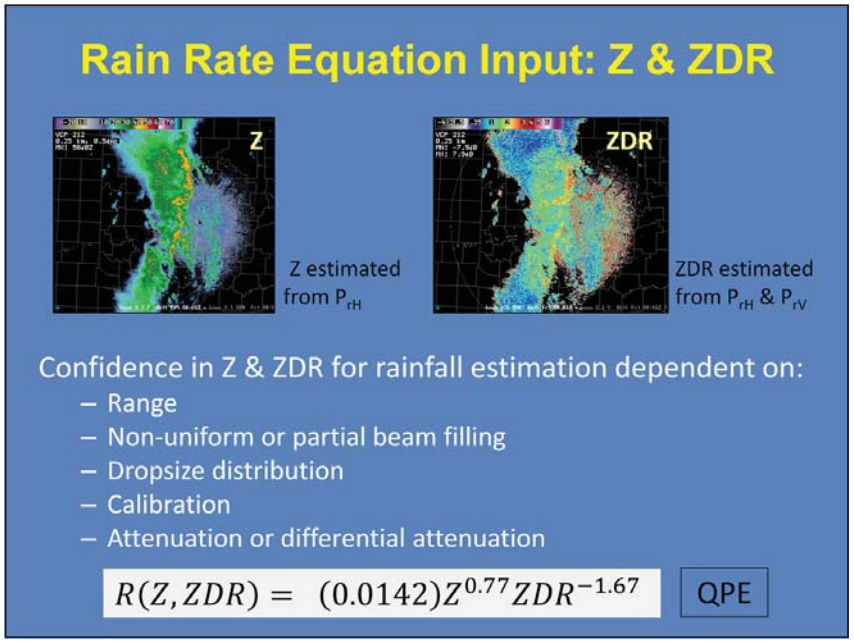
The second and third equations combine Z and ZDR as input, in linear units (mm^6/m^3). The R(Z,ZDR) equation labeled Continental is the default, is recommended for the cool season and for warm season deep convection. The Tropical R(Z,ZDR) is recommended for the warm season and for events that are dominated by warm rain processes, especially hurricanes and tropical systems. Southern and coastal locations may find the Tropical R(Z,ZDR) sufficient for a majority of events.

The R(KDP) is used where hail has been identified. You may notice that the R(KDP) equation includes the possibility of a negative rain rate. The QPE algorithm logic rejects R(KDP) when the rate is negative, so don't worry about that possibility.



The limitations of $Z = 300R^{1.4}$, now presented in its R(Z) format, are probably familiar. It often underestimates rainfall in a warm rain dominant event, while overestimates the water at the surface from mixed precipitation types, such as wet snow and hail.

R(Z,ZDR) combines input from both Z and ZDR. The goal is to balance the benefits of each of these inputs, and this combination is generally best where the beam is sampling all rain.



Using Z and ZDR as inputs for computing rain rate have similar dependencies that affect their reliability for rainfall estimation. As a reminder, Z is estimated from the returned power in the horizontal channel, while ZDR is estimated from the returned power in both the horizontal and vertical channels.

As range increases, the volume sampled also increases. The larger radar volume makes it harder for the beam to be uniformly filled with the same hydrometeors. Also at longer ranges, the beam can be filled with by all snow and/or ice crystals. Frozen hydrometeors differ from those assumed in the Probert-Jones radar equation, which affects the accuracy of both Z and ZDR.

Variations in dropsize distribution, especially the presence of hail, affect both Z and ZDR values.

Both Z and ZDR are sensitive to calibration errors.

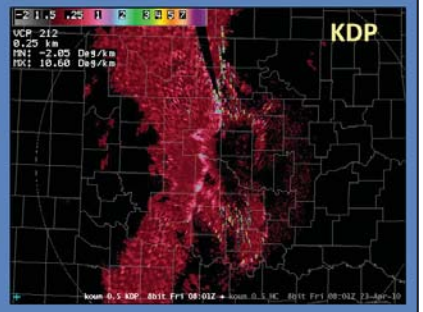
Even at 10 cm, attenuation (or differential attenuation) happens. Significant attenuation of either kind can make the Z or ZDR estimates less reliable.

QPE Rain Rate Equation

$$R(KDP) = 44.0|KDP|^{0.822} sign(KDP)$$ QPE

R(KDP):

- KDP mostly immune to partial beam blockage & hail
- Helps QPE mitigate hail contamination
- Sign(KDP) is a quality control



R(KDP) has the advantage that KDP is mostly immune to both partial beam blockage and hail contamination. The magnitude of KDP is related to the liquid water content in the volume, but not to ice. For example, if a volume contains mostly hail, KDP values can be quite low. On the other hand, a large quantity of small melting hail can result in KDP being very large. Since KDP is mostly immune to hail, it can be used to estimate rainfall where hail is present. This equation includes the term "sign(KDP)". Negative KDP values are possible, and this term is included as a quality control. If R(KDP) < 0, it is not used to estimate rainfall.

Rainrate Equation Input: KDP

$$R(KDP) = 44.0|KDP|^{0.822} sign(KDP)$$ QPE

- KDP
  - Less dependent on dropsize distribution
  - Mostly immune to partial beam blockage & hail
- Confidence in KDP for rainfall estimation dependent on
  - Returned signal strength
  - Number of pulses per radial
  - Associated CC value

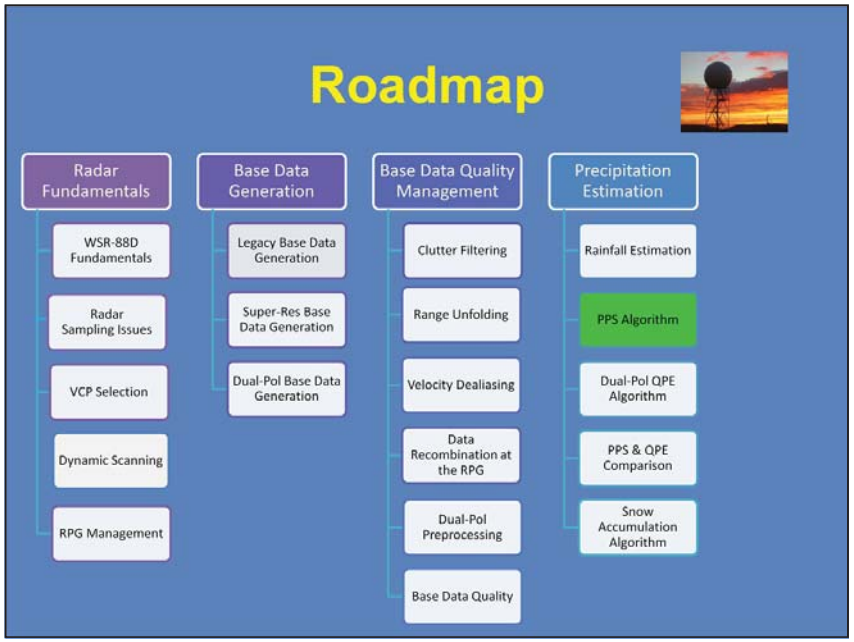


KDP as input for computing rainrate is less dependent on dropsize distribution than Z and ZDR. KDP can be useful where hail is present or there is partial beam blockage.

Compared to Z and ZDR, KDP is noisier and more dependent on sufficient signal strength and pulses per radial. Where the returned power is low (or fast VCPs are used), KDP is noisy and thus less reliable for conversion to rainfall. However, the noisiness is usually not a problem for qualitative human interpretation. So, KDP can still be very useful for identifying areas of heavy rain even when noisy. The use of KDP for rainfall estimation is also dependent on the associated CC value. The details of this will be presented in the QPE lesson, but for now, if CC is low, R(KDP) is not used for rainfall estimation.



Welcome to Legacy Precipitation Processing Subsystem (PPS) Algorithm.



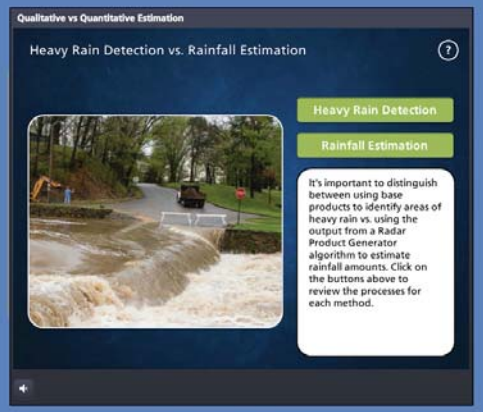
Here is the “roadmap” with your current location.

Precipitation Estimation
Objectives



1. Identify strengths & limitations of the Legacy Precipitation Processing Subsystem (PPS).

Review: Heavy Rain Detection vs. Rainfall Estimation



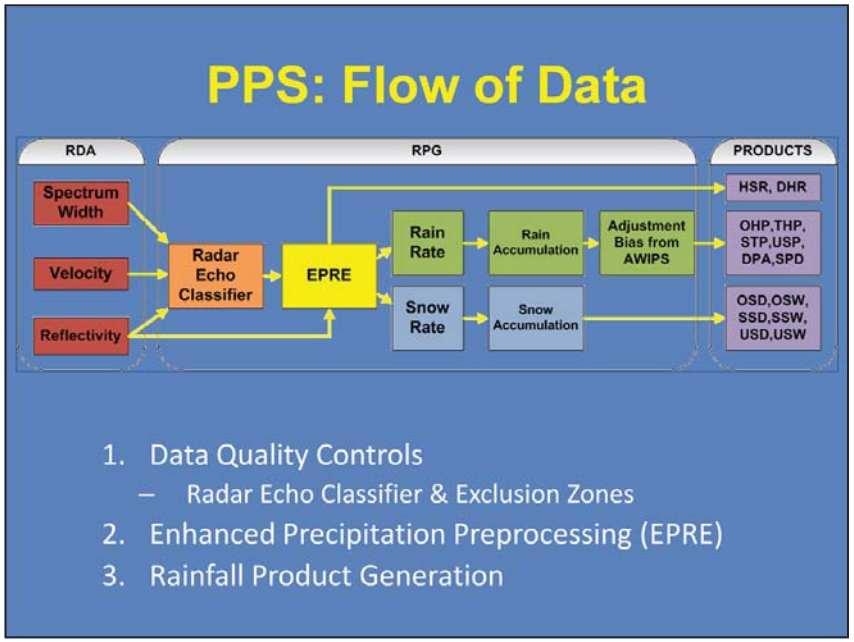
If no pop-up window appears that looks like above, open a browser and go to:
http://www.wdtd.noaa.gov/courses/rac/principles/interactions/qual-vs-quant/

There is one objective for this lesson, focusing on the design of the PPS.



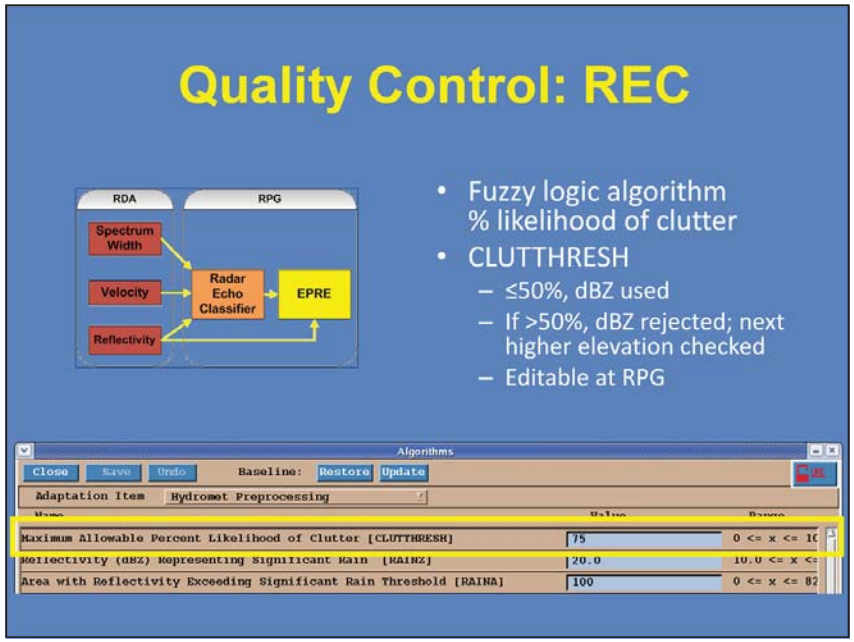
Now for a “tour” of the Legacy (aka been around awhile) Precipitation Processing Subsystem (PPS). The PPS is a series of algorithms that use base reflectivity, velocity and spectrum width data as input, then estimate rainfall and generate rainfall products. This process contains several quality control steps and the rainfall estimates are provided out to a range of 124 nm. In addition, there are several adaptable parameters utilized that provide configuration flexibility.

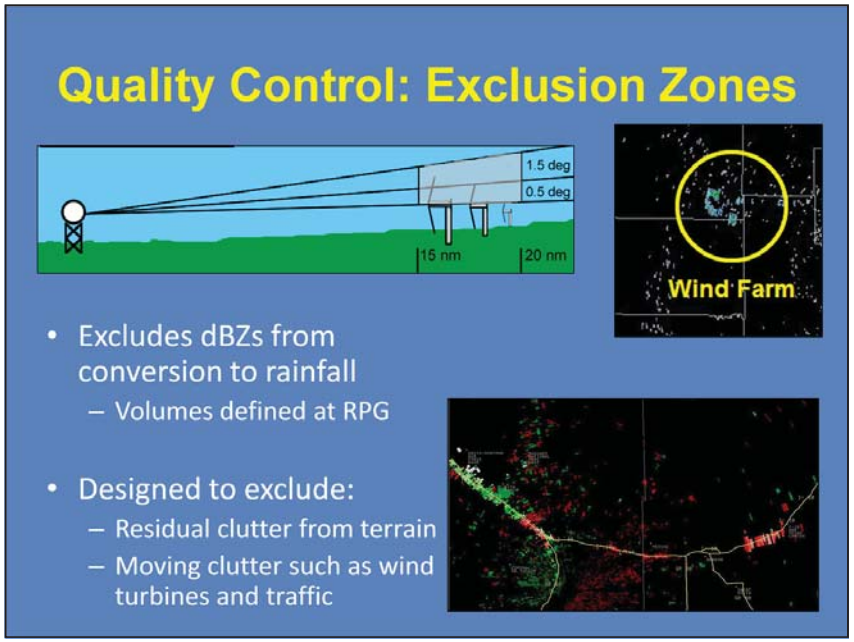
This is a complex algorithm with multiple steps. The focus of this “tour” will be the portions of the PPS where there is some operator control.



This graphic provides a high-level look at the flow of information through the Precipitation Processing Subsystem (PPS) and will serve as an outline for the PPS components. There are multiple quality control steps that will be highlighted, including a look at how the base data are “prepared” by the Enhanced Precipitation Preprocessing (EPRE) algorithm. Some important steps that affect product generation will also be provided.

The two boxes that have the word “snow” on them are not related to the PPS. They will be discussed in a later lesson, which presents the Snow Accumulation Algorithm.





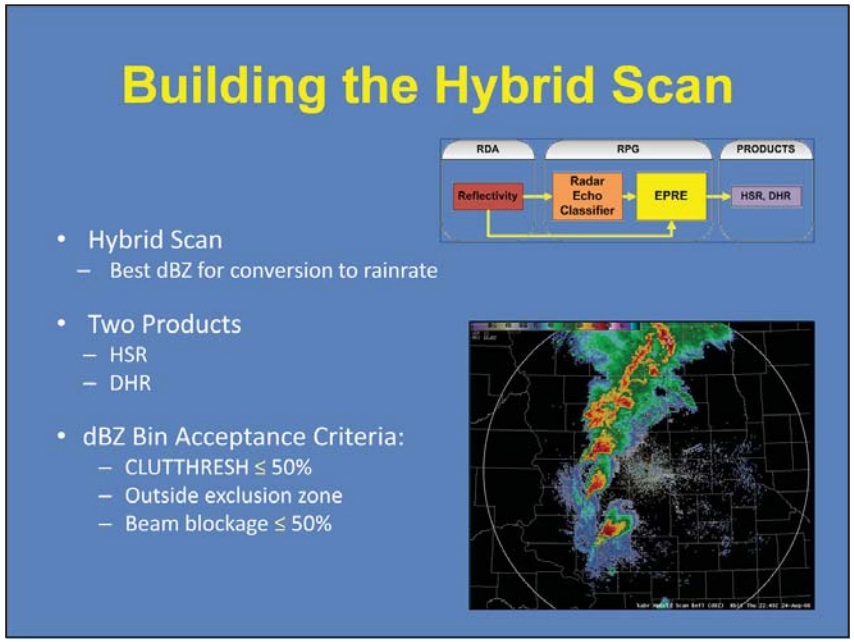
The Radar Echo Classifier is a fuzzy logic algorithm that uses reflectivity, velocity, and spectrum width to assign a likelihood that a particular bin contains clutter.

EPRE ingests this guidance for comparison against a parameter called CLUTTHRESH which determines whether or not that dBZ is used in rainfall product generation. The default setting for CLUTTHRESH is 50%, which means that if a bin is assigned less than or equal to 50% by the REC, the dBZ for that bin is used for conversion to rainfall. If the REC has assigned a % greater than the CLUTTHRESH value, the bin is rejected. For rejected bins, the next higher elevation bin is then checked.

The CLUTTHRESH parameter is editable at the RPG. Here's an example where CLUTTRESH was increased to 75% during the warm season for a location with very little terrain clutter. The intent was to use the lowest elevations possible for rainfall estimation to better avoid hail contamination.

Another quality control option that is part of EPRE is the application of Exclusion Zones, which are defined locally. Exclusion zones are used to prevent reflectivity from specific areas and elevations from being converted to rainfall. An exclusion zone is actually a volume, defined from azimuth to azimuth, range to range, and up to a maximum elevation. Exclusion zone definition is done at the RPG.

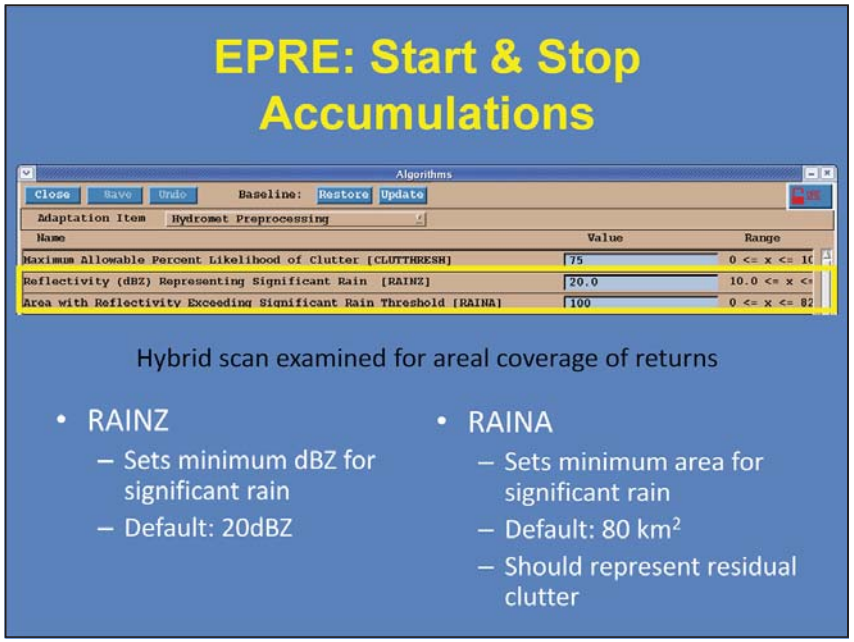
Optimally, these zones will be used to prevent residual clutter from terrain and moving clutter originating from sources like wind turbines and traffic from being converted to rainfall. These areas, especially at close ranges to the radar will cause high reflectivity that cannot be removed by the clutter filters.



The “grand finale” of the EPRE algorithm is the building of the hybrid scan. The idea of a hybrid scan is to find an optimal dBZ value at each range bin only for the purpose of converting to rainrate. There are two Hybrid Scan products that represent the reflectivity field that was used by the PPS for that volume scan. Specific products generated are the Hybrid Scan Reflectivity and Digital Hybrid Scan Reflectivity. More information on these products will be presented in products portion of this course.

In order for the EPRE to accept a dBZ bin into the hybrid scan, it must meet the following criteria:

1) Must have a clutter likelihood of less than the CLUTTRESH setting (50% by default)
2) Must fall outside of a defined EPRE exclusion zone.
3) Beam blockage must be no more than 50%



Based on the hybrid scan, EPRE determines when accumulations begin and when they end. The idea of whether or not it is “raining” is based on the areal coverage of returns above a certain dBZ value. There are two EPRE adaptable parameters that govern the start and stop of rainfall, RAINZ and RAINA. RAINZ is the minimum dBZ that “counts” as rain. The default value for RAINZ is 20 dBZ, which is generally considered to be the minimum dBZ for precipitable returns.

RAINA is the minimum areal coverage of returns at or above RAINZ for accumulation to either begin or to continue. The default value for RAINA is 80 km^2, which is often too small. RAINA is meant to represent the average areal coverage of residual clutter for each radar. If RAINA is smaller than the residual clutter area, the PPS may be accumulating clutter instead of precipitation.

With the default settings for RAINZ and RAINA above, if 80 km^2 of returns at or above 20 dBZ are detected, the PPS processes for accumulating rainfall begin, and will continue each volume scan that the thresholds are met. The rainfall accumulations are automatically reset to zero once conditions fall below the RAINZ or RAINA thresholds for one hour.

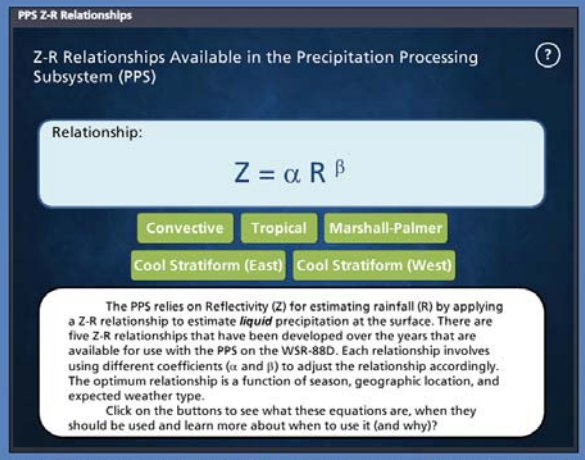
Resetting Accumulations



Automatic reset to zero after 1 hour below RAINZ or RAINA

Manual reset at RPG
- Legacy PPS

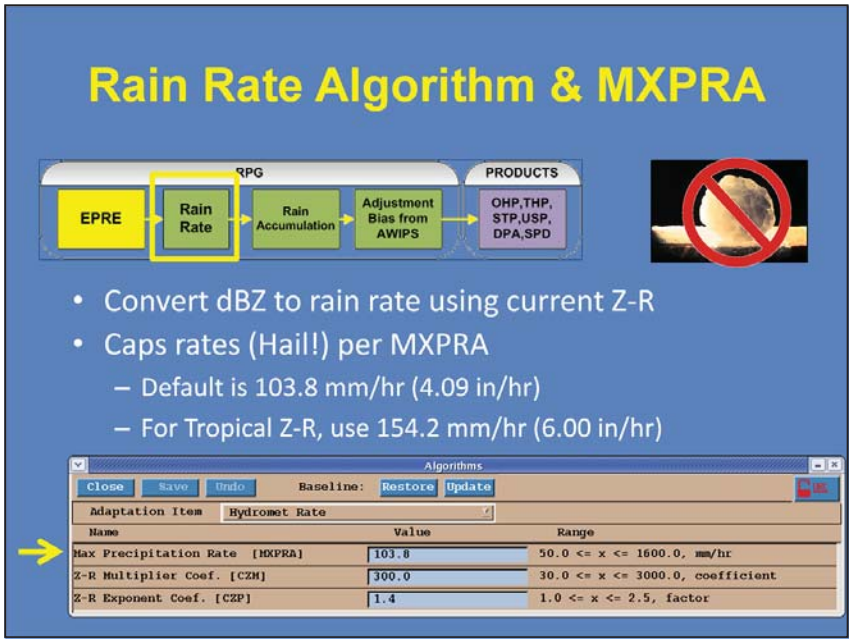
PPS Z-R Relationships



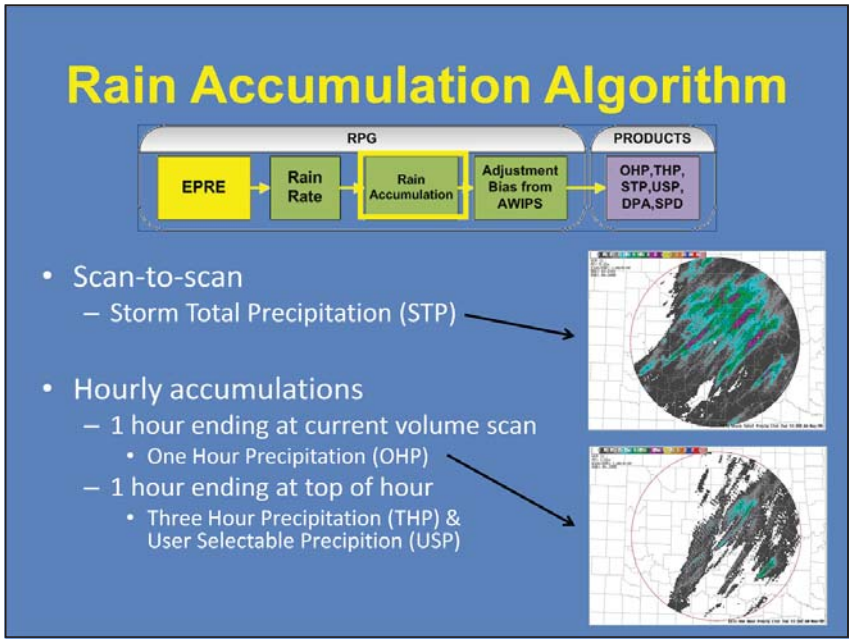
If no pop-up window appears that looks like above, open a browser and go to:
http://www.wdtd.noaa.gov/courses/rac/principles/interactions/pps-zr/

There are two approaches to resetting rainfall accumulations to zero. The first is the automatic reset when the conditions fall below RAINZ and RAINA for one hour. The storm total is a type of accumulation that continues as long as RAINZ and RAINA are exceeded. For some locations, this can be too long!

There is a manual reset of the storm total accumulation available at the RPG, specifically the RPG Control window. Manual resets are actually available for both the rainfall algorithms, as well the snowfall algorithm, but for now, we focus on the “Legacy PPS”.



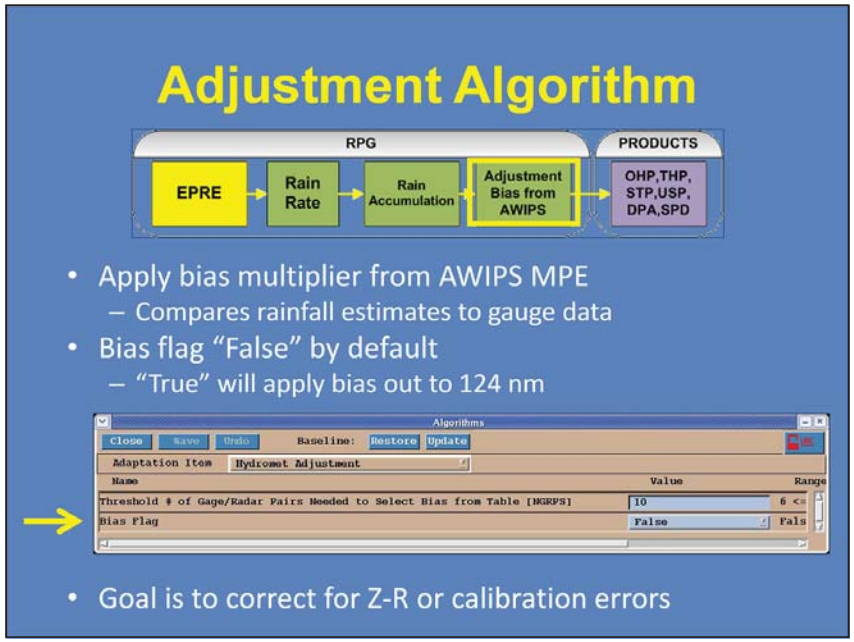
The Rain Rate algorithm converts the dBZ values from the hybrid scan to rain rate using the current Z-R relationship applied at the RPG. The Rate Algorithm also applies a parameter called the Max Precipitation Rate (MXPRA). MXPRA works as a cap to prevent hail contamination. The default setting for MXPRA is 103.8 mm/hr, which is 4.09 in/hr. This means that any rain rates that exceed 4.09 in/hr will be capped at this value. If the Tropical Z-R is used, it is recommended that MXPRA also be adjusted to allow for higher rain rates. The recommended setting is 154.2 mm/hr, or 6.00 in/hr.



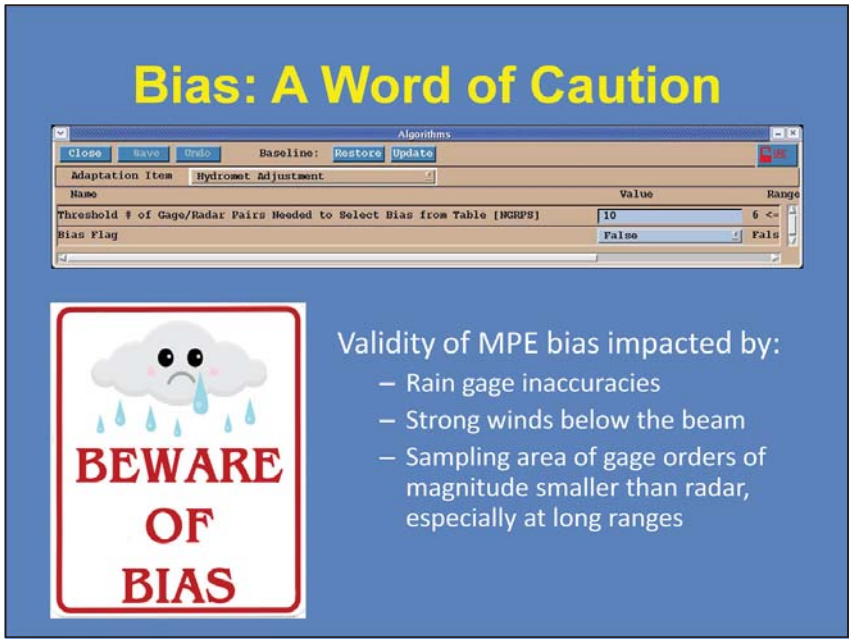
The Accumulation Algorithm, uses the calculated rain rates and differing durations to accumulate rainfall. There are two different types of accumulations.

Scan to scan accumulations continue every volume scan as long as RAINZ and RAINA are exceeded. Scan to scan is the accumulation displayed on the storm total products, the one most commonly used is the Storm Total Precipitation (STP) product.

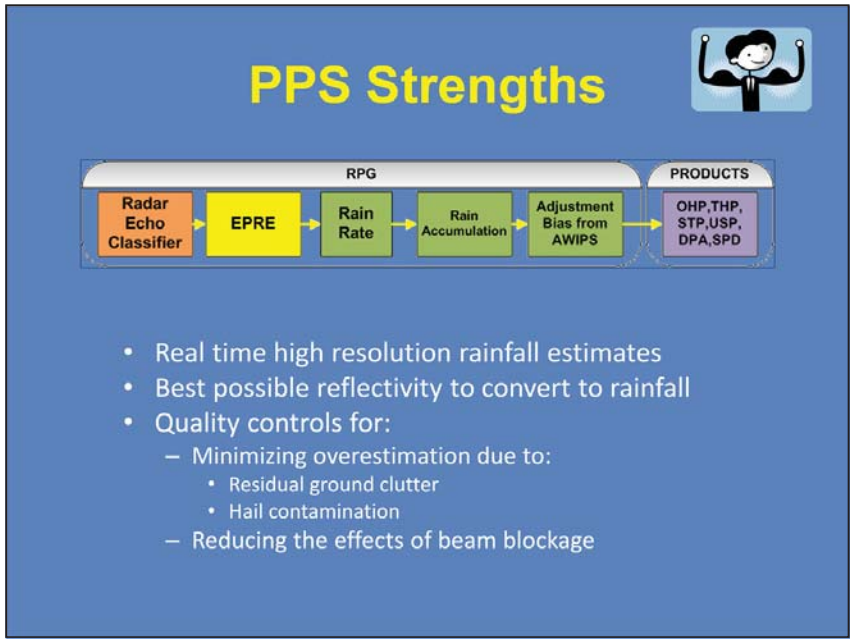
The second type of accumulation is hourly. There is a one hour accumulation ending at the current volume scan time. This is the accumulation displayed on the one hour products, the one most commonly used is the One Hour Precipitation (OHP) product. There is also a one hour accumulation that ends at the top of each hour. These hourly ending at the top of the hour accumulations are used to build the Three Hour Precipitation (THP) and User Selectable Precipitation (USP) products.



The adjustment algorithm is the last of the PPS algorithms, providing the option of applying a bias multiplier to the rainfall accumulations. The AWIPS Multi-Sensor Precipitation Estimator (or MPE) compares radar rainfall estimates to gauge data and sends a bias table to the RPG once an hour. Applying the bias is controlled by a parameter known as the Bias Flag. It is set to false by default. Setting the Bias Flag to true will apply the best bias generated by the MPE out to 124 nm. The goal of the bias adjustment is to correct for a non-representative Z-R relationship or calibration errors.

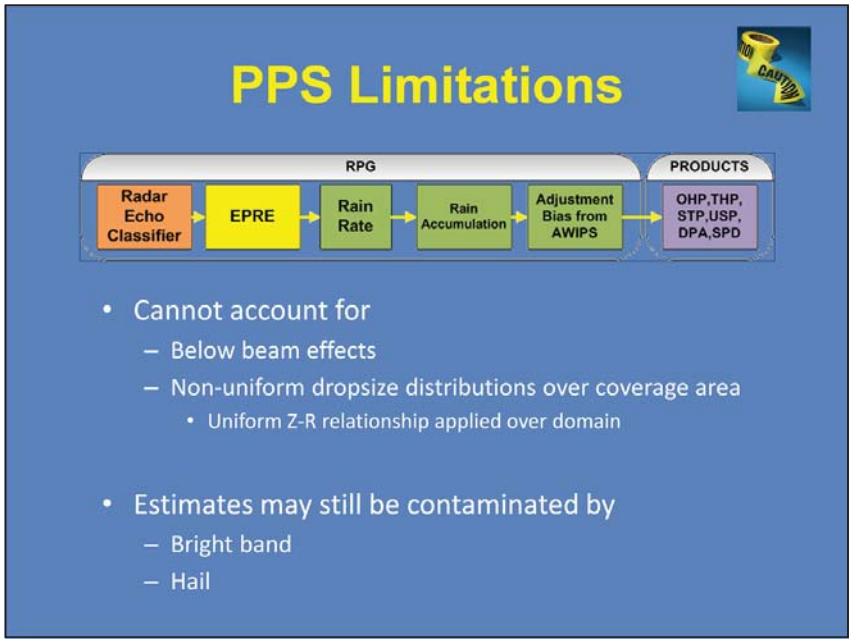


The MPE bias output isn't perfect and there are some scenarios to consider that could affect the validity of this bias. Rain gage values can be inaccurate for a variety of reasons. There may be strong winds below the beam. The most important consideration is that the rain gage sampling area is orders of magnitude smaller that the radar.



Here are the strengths of the PPS algorithm. This is the only source of real time high resolution rainfall estimates. It is important to remember that the qualitative spatial information of the rainfall pattern can be valuable.

The PPS uses the reflectivity value closest to the ground that is not contaminated with clutter. There are quality control steps to minimize hail contamination and to avoid radials with beam blockage.

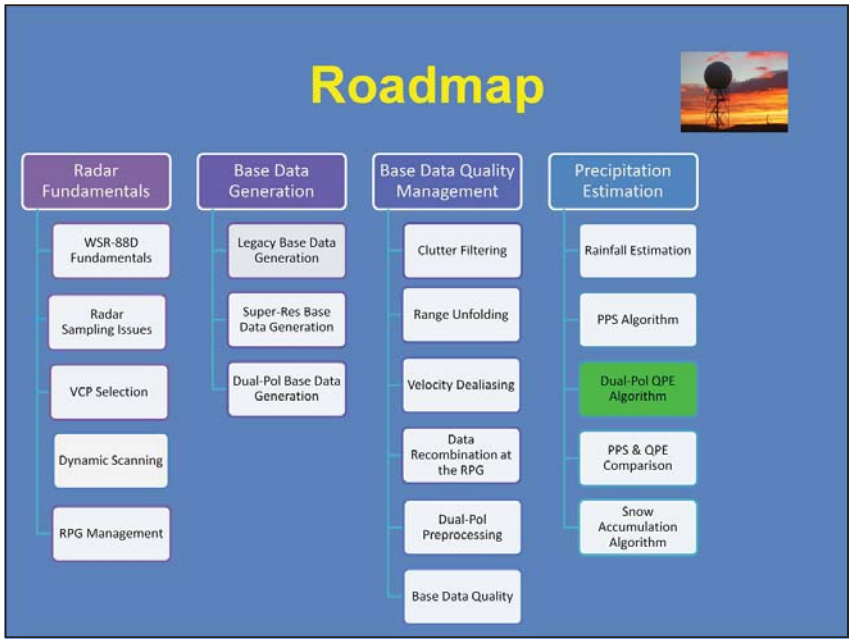


There are also some limitations to the PPS. The PPS cannot account for below beam effects such as strong winds, evaporation, or coalescence. Since a single Z-R relationship is applied over the entire domain, the PPS estimates can be compromised by non-uniform dropsize distributions in the radar area.

While the PPS does attempt to mitigate some errors, it may not always be 100% successful. Even though the PPS chooses the lowest viable reflectivity with sufficient coverage, the beam will strike the melting layer and bright band contamination can occur. Even with the application of the Max Precipitation Rate parameter, hail contamination may not be completely removed.



Welcome to the Dual-Pol Quantitative Precipitation Estimation (QPE) Algorithm



Here is the “roadmap” with your current location.

Learning Objectives



1. Identify strengths & limitations of the Quantitative Precipitation Estimation (QPE) algorithm

There is one objective for this lesson, focusing on the design of the QPE algorithm.

QPE Tour

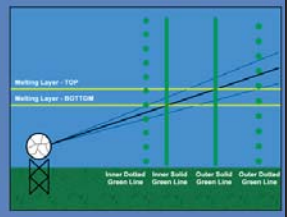
$$R(Z) = (0.017)Z^{0.714}$$

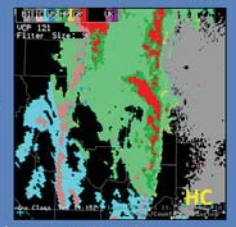
$$R(Z, ZDR) = (0.0067)Z^{0.927} ZDR^{-3.43}$$

or

$$R(Z, ZDR) = (0.0142)Z^{0.77} ZDR^{-1.67}$$

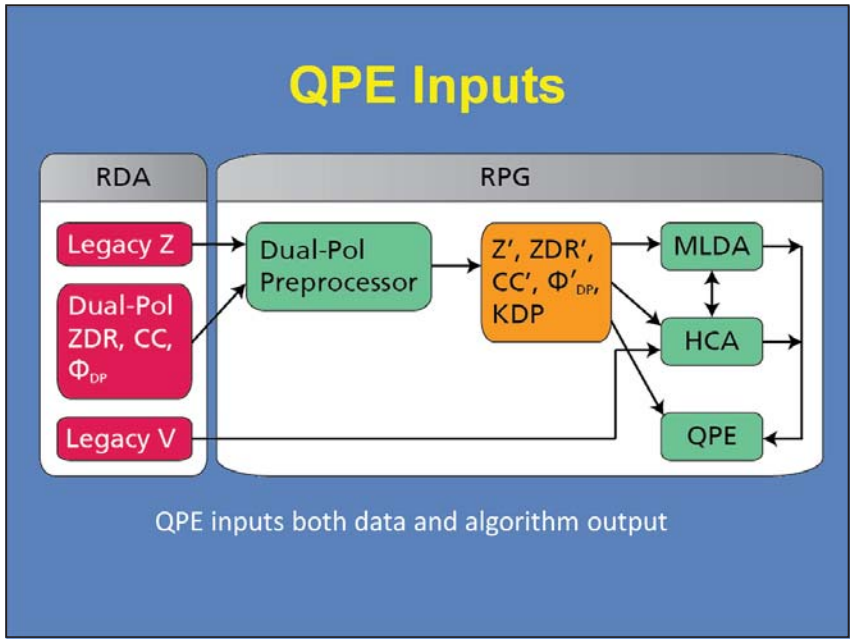
$$R(KDP) = 44.0|KDP|^{0.822} sign(KDP)$$

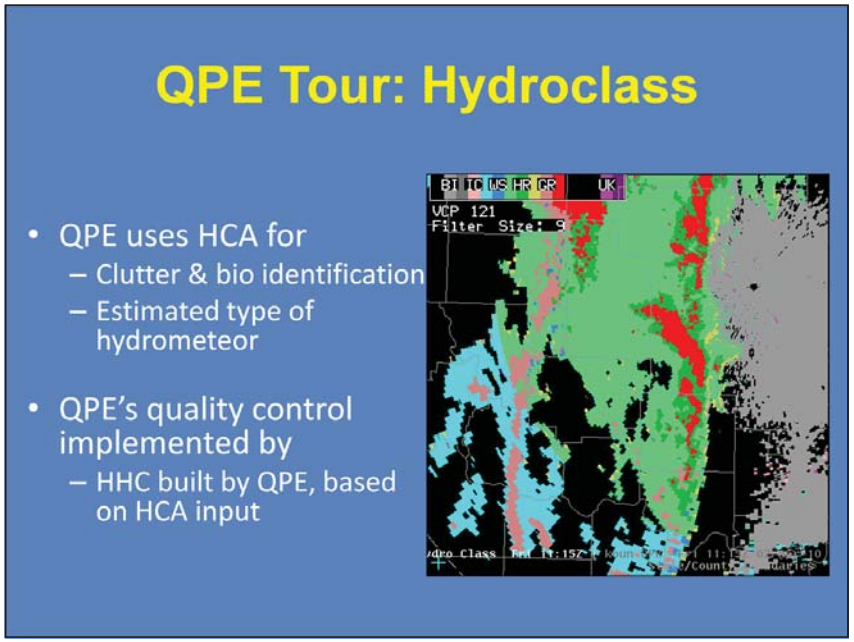




- PPS has sequential quality control steps
- QPE's quality control implemented by
  - Three different rain rate equations
  - Input from MLDA & HCA determines which equation is used

The QPE Tour will be somewhat different from the PPS Tour. The Legacy Precipitation Processing Subsystem (PPS) was presented as a sequence of quality control steps. The QPE Tour will be less sequential, but focused on key design elements of the QPE. This QPE Tour also provides more information on how QPE relies on the output from the Melting Layer Detection Algorithm (MLDA) and the Hydrometeor Classification Algorithm (HCA), which both determine which rain rate equation is used for any given range bin. The use of three different rain rate equations and the reliance on MLDA and HCA makes QPE significantly more complex than any other RPG algorithm.





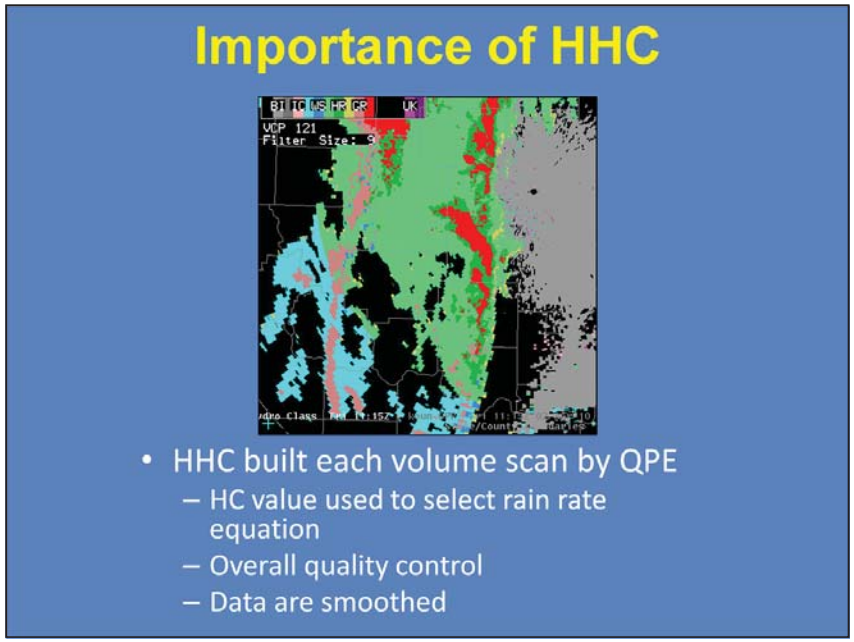
The Dual-Pol Quantitative Precipitation Estimation (QPE) algorithm is the most complex RPG algorithm yet implemented. QPE relies on multiple base data inputs, as well as the output from two new RPG algorithms. These algorithms are the Melting Layer Detection Algorithm (MLDA) and the Hydrometeor Classification Algorithm (HCA).

With the exception of velocity, the base data inputs are first passed through the Dual-Pol Preprocessor. Reflectivity (Z), Differential Reflectivity (ZDR), Correlation Coefficient (CC), and Differential Phase (Φ_{DP}) are smoothed and (in some cases) corrected for attenuation. The "prime" notation is being used to denote the data that have been preprocessed. Specific Differential Phase (KDP) is also generated by the Dual-Pol Preprocessor and input to all of the Dual-Pol RPG algorithms.

In addition to the base data, QPE is reliant on the assessed height of the melting layer from MLDA and the expected hydrometeor (or non-hydrometeor) type from the HCA.

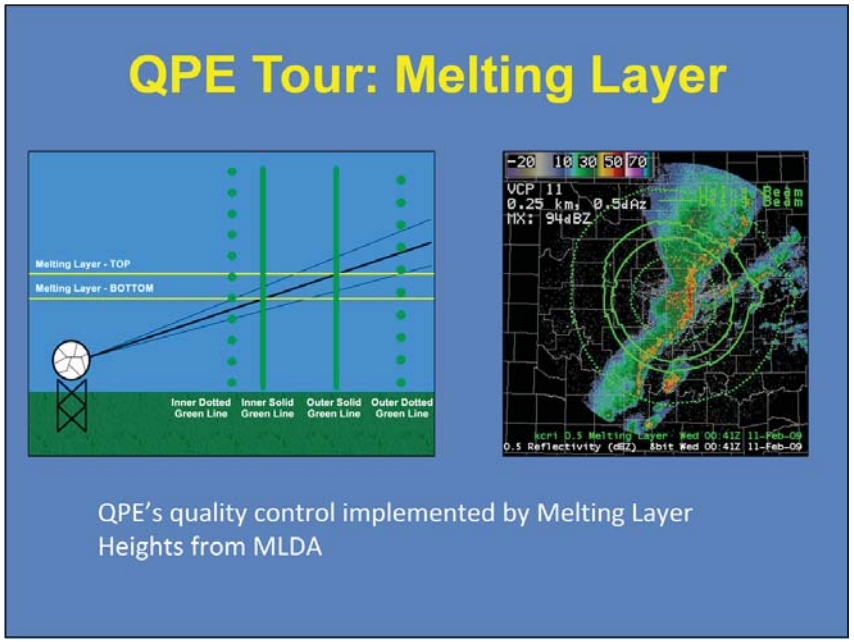
The QPE relies on the HCA in two ways. Bins that have clutter or biological targets identified by HCA are not converted to rainfall by the QPE. For bins that HCA identifies as having some type of precipitable hydrometeor, the QPE uses one of the different rain rate equations based on the type of hydrometeor (much more about that soon)!

The Hybrid Hydrometeor Classification (HHC) product is output from the QPE. It shows you, on a bin by bin basis, which HCA values were used by QPE to generate the rainfall accumulation products. The HHC can be very useful as a quality control check for the QPE rainfall products.

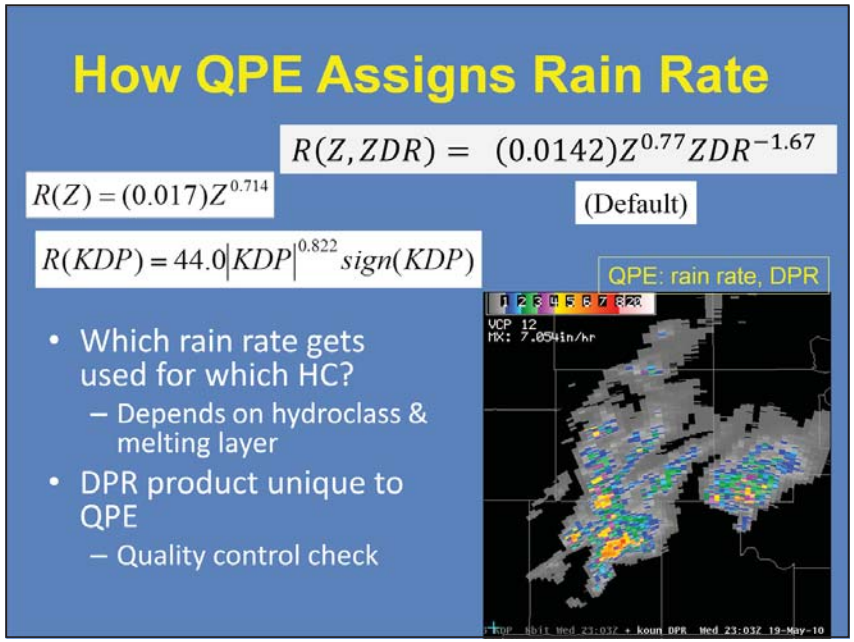


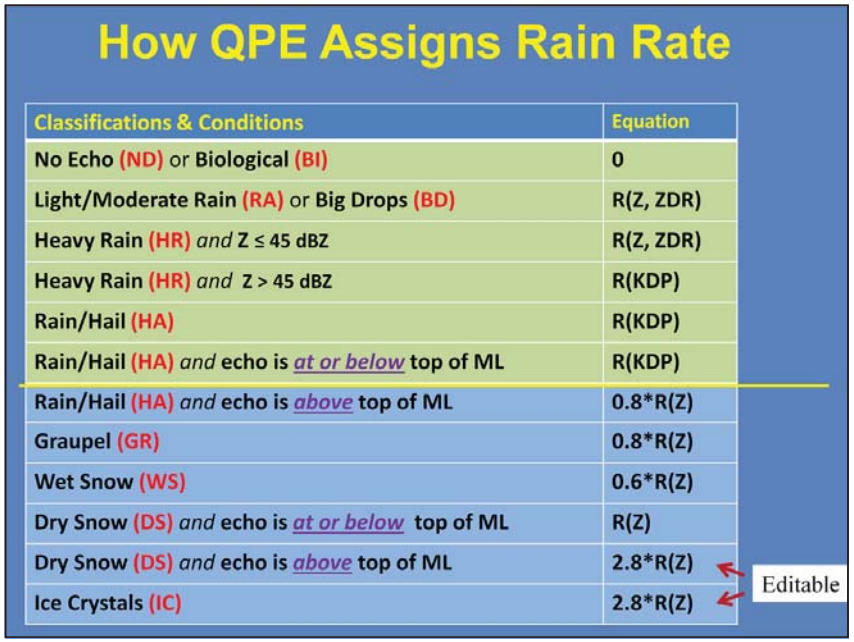
The Hybrid Hydroclass (HHC) product is built by the QPE and represents the hydrometeor classification values that were used to determine the rain rate equation applied, on a bin by bin basis, for each volume scan.

This product can be used for an overall quality control check, but be cautious about checking every single bin. The data on this product are smoothed. The technique is called a 9 bin filter. For each bin, the most common hydroclass value for the surrounding 9 bins is assigned. This has the effect of reducing speckling on the product.



The QPE also relies on output from the MLDA. For each elevation angle, the MLDA provides four different heights related to the melting layer, which are displayed on AWIPS as an overlay product. For the purpose of QPE quality control, the top and bottom of the melting layer are used along with the Hydrometeor Classification values. Here's an example: for the Hydroclass value of Dry Snow (DS), one rain rate equation is used where DS is above the top of melting layer, while something different is used where DS is in or below the melting layer.



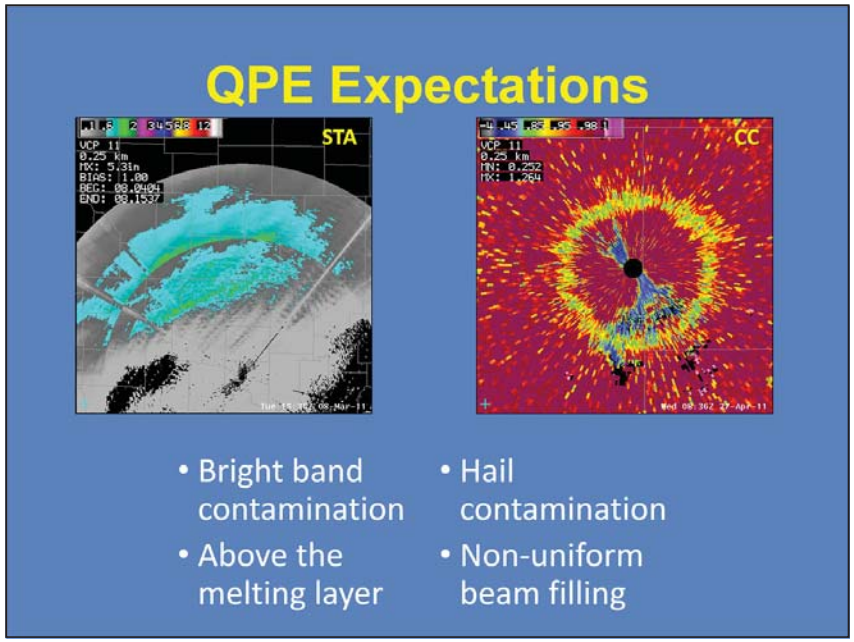


There are a number of steps involved in determining which rain rate equation is used, given the Hydroclass value and position of the range bin with respect to the melting layer. In some cases, a multiplier is used with the Z-R equation, such as 0.8* R(Z). Variations of the R(Z) are used in or above the melting layer. The R(Z,ZDR) and R(KDP) equations are used below the melting layer where the expectation is that the beam is sampling rain (or where hail possibly mixed with rain is identified).

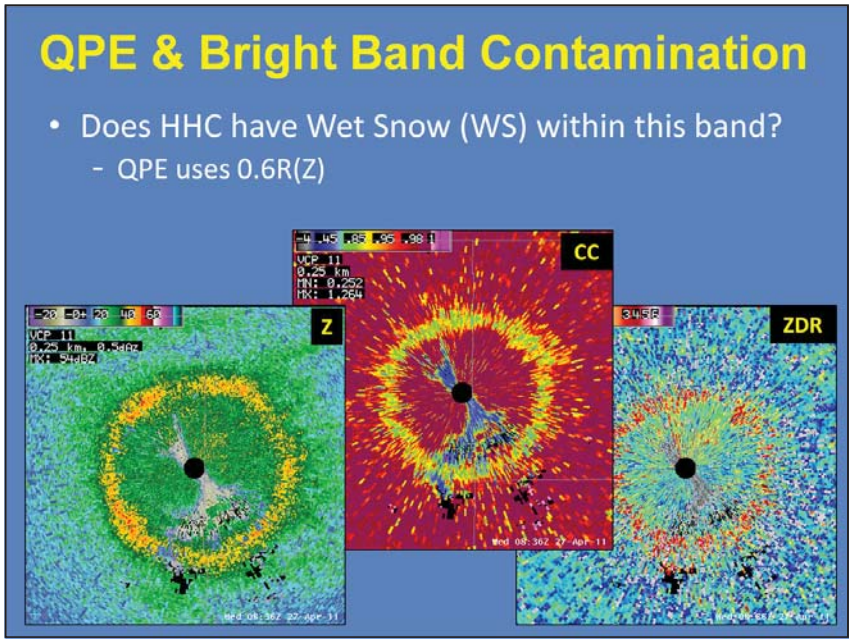
QPE has one product that has no PPS counterpart. It is the Digital Precipitation Rate (DPR) product. It is generated every volume scan and presents the rainfall rates in inches per hour that were used to generate the suite of QPE rainfall accumulation products. The DPR can be used as a quality control check to determine if the precipitation rates seem reasonable.

This table is a summary of the hydroclass values and position relative to the melting layer that determines which rain rate equation is used by the QPE. The primary idea to note is that rain rate equations that include the dual pol variables, R(Z,ZDR) and R(KDP) are used where the hydroclass values are mostly liquid and below the melting layer. The exception is Rain/Hail at or below the top of the melting layer. This group is in the top of the table.

At the bottom of the table, for hydroclass values that non-liquid, R(Z) is used with or without a multiplier. For Dry Snow (DS) above the top of the melting layer and Ice Crystals (IC), the 2.8 multiplier has resulted in significant overestimates in many cases. This multiplier is now editable, once sufficient local research has been done to determine a more appropriate value.

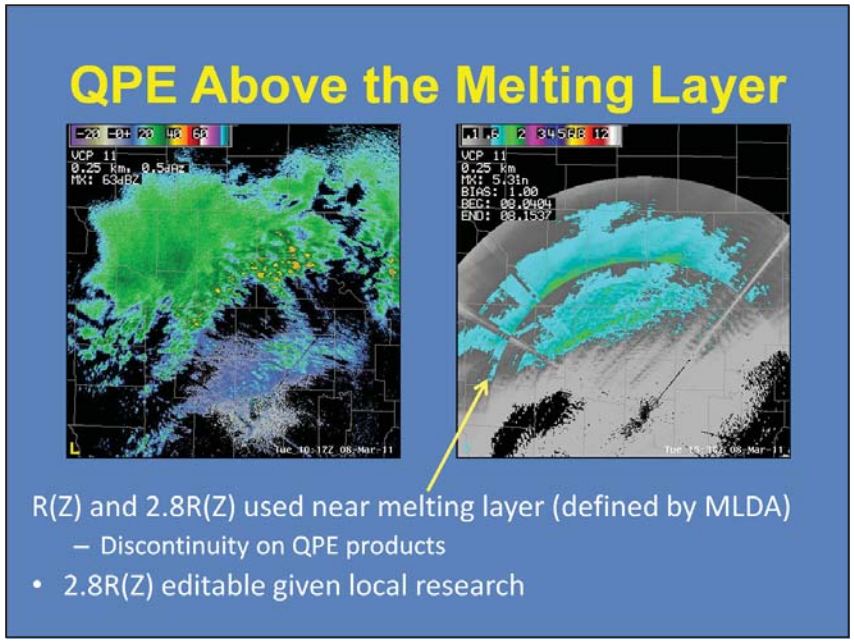


The next several slides summarize how QPE is designed to mitigate some of the most difficult problems with using a radar to estimate rainfall.



Bright band contamination has long been a challenge for using a radar to estimate rainfall. The QPE approach is to adjust the rain rate equation for water coated frozen hydrometeors such as wet snow that are typically located within a mesoscale melting layer (stratiform rain event). A multiplier of 0.6 is applied to the R(Z) relationship for the WS bins. There are some considerations to remember. The 0.6R(Z) equation is only applied to bins that fall within the melting layer (as defined by MLDA) and are identified as Wet Snow (WS) by the HCA.

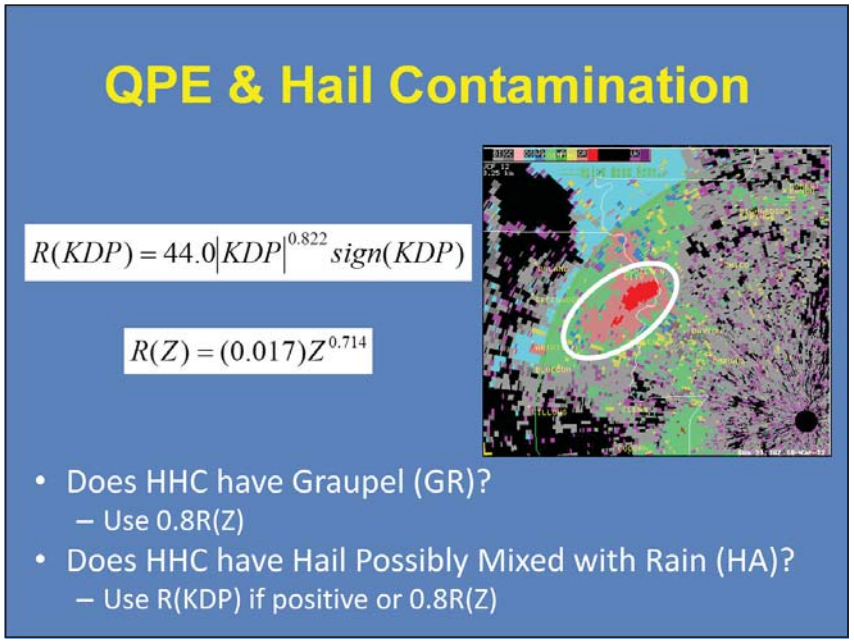
The bright band is often apparent on Z, or CC, and ZDR, allowing for a base data quality control check for these algorithms. It is important to remember that QPE's ability to mitigate the overestimate of rain within the bright band is totally dependent on a band of Wet Snow (WS) being properly identified by the HHC coincident with the bright band on the base data products.



When the HHC assigns IC (Ice Crystals) or DS (Dry Snow), and the DS is above the melting layer, 2.8R(Z) is used to generate a rainfall rate. The location of the melting layer is based on output from the MLDA. This multiplier is now editable, given sufficient local research.

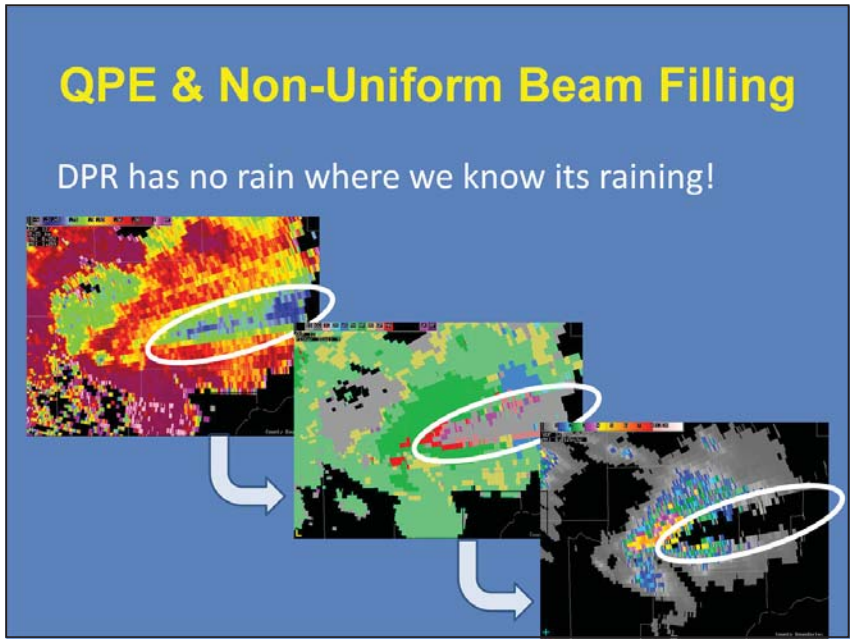
For DS in or below the melting layer, R(Z) is used. This use of R(Z) in or below the melting layer and 2.8R(Z) above can result in a discontinuity on QPE products near the melting layer for long term stratiform events. Here's an example Storm Total Accumulation (STA) product from an lengthy stratiform rain event.

This illustrates the difference between a human assessment of a transition vs. an algorithm. For an algorithm, the "top" of the melting layer is a sharp transition, while our human understanding is that the top is really a layer in itself.



QPE addresses hail contamination for range bins that are likely to contain hail or graupel, as identified by the HHC. The approach is to use 0.8R(Z), which lowers the rain rate by the 0.8 multiplier, or to use R(KDP). KDP is the Dual-Pol variable that is sensitive to the liquid water content in the volume. The presence of hail or the size of hail has little impact on the KDP value, so it is a good choice for conversion to rain rate. The R(KDP) equation includes the term sign(KDP) to check for the possibility that KDP is negative. Negative rain rates based on KDP are not used by QPE.

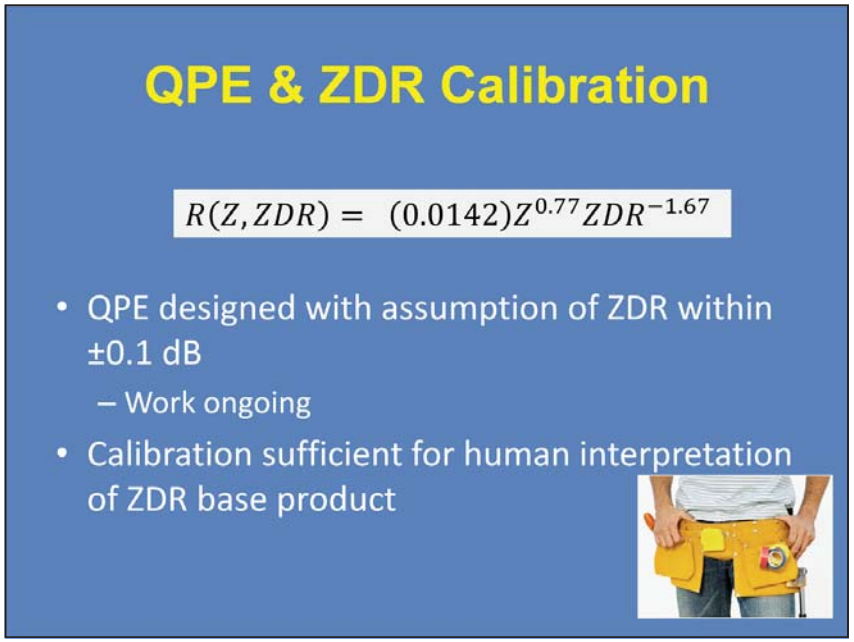
In this HHC example, most of the bins within the white circle are either Graupel (GR) or Hail.



The artifact of a swath of low CC values due to non-uniform beam filling (NBF) can be either easy to spot or subtle. By comparing it to other radar data and understanding the environment, you can ask yourself if the CC values make sense.

It is important to be mindful of this artifact because CC affects the Hydroclass value that gets assigned, which then affects whether or nor rainfall is accumulated. In this case, biological targets are identified by the Hydrometeor Classification Algorithm, and QPE assigns no rain rate to bins with the Biological Hydroclass value.

There are steps within the RPG Dual Pol Preprocessor to identify bins with NBF, though they may not always be successful.

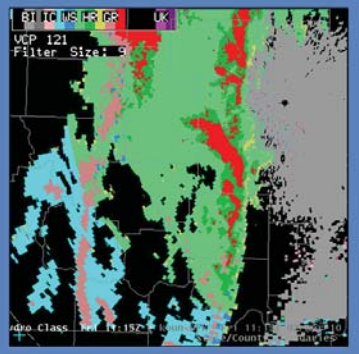


The QPE logic and parameter design was based on an assumption of ZDR values that are calibrated to within ±0.1 dB. As of this writing, it is not known if that level of accuracy has been achieved.

ZDR calibration is based on an accurate calibration of the horizontal and the vertical channels. With rare exceptions, ZDR values are sufficient for human interpretation. For example, many of the benefits of ZDR, such as updraft or hail detection, are based on relative values of ZDR. A relative minimum, rather than a specific value of ZDR adds confidence in the presence of hail.

The QPE, and the other RPG Dual-Pol algorithms have a higher sensitivity to ZDR calibration, and work is underway to develop techniques to refine calibration of the horizontal and vertical channels.

QPE "Ground Rules"



- Frequently investigate the:
  - ML product to verify that it makes sense
  - HHC product to verify that it makes sense
- Recognize situations when QPE performance is likely to be negatively impacted
  - Expect better performance below the melting layer

Since QPE relies so heavily on the output of the MLDA and the HCA, it is important to avoid using the QPE products "as is". The QPE "ground rules" include monitoring the output from the MLDA and looking at the HHC product to verify that the melting layer and the hydroclass values make sense.

It is also important to be mindful of situations where QPE performance is likely to be negatively impacted, such as above the melting layer. Better performance can be expected below the melting layer.

QPE Strengths

1. Prevent returns dominated by non-meteorological targets from conversion to rainfall
2. Better rain rates based on hydrometeor classification
3. Mitigate bright band overestimation
4. Mitigate hail contamination
5. Rain rate product every volume scan

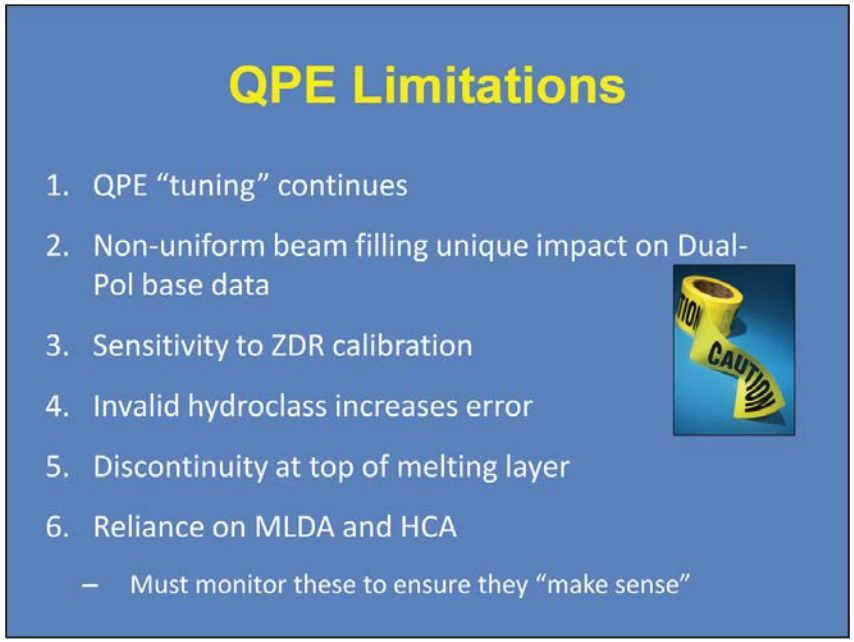


The strengths of QPE are mostly based on using the benefits of Dual-Pol to mitigate long standing challenges with using any radar to estimate rainfall.

Preventing returns from non-meteorological targets from conversion to rainfall is based on identification of Ground clutter and biological returns by the HCA.

For bins that are identified with precipitable returns, QPE has three different rain rate equations that are applied based on the hydroclass value and the height with respect to the melting layer. These choices help to mitigate both bright band and hail contamination.

QPE also has a rain rate product generated every volume scan, which can be an asset for quality control.



As with most new algorithms, there are a number of limitations to QPE that will likely improve over time. The original QPE design and associated parameters were based on research in Oklahoma, and regional "tuning" is ongoing.

Non-uniform beam filling has a unique impact on Dual-Pol base data, which translates to the QPE performance.

The performance of the R(Z,ZDR) equation is highly dependent on accurate ZDR values. The accuracy of ZDR is sufficient for human interpretation of the ZDR base product, but QPE performance is more sensitive to ZDR.

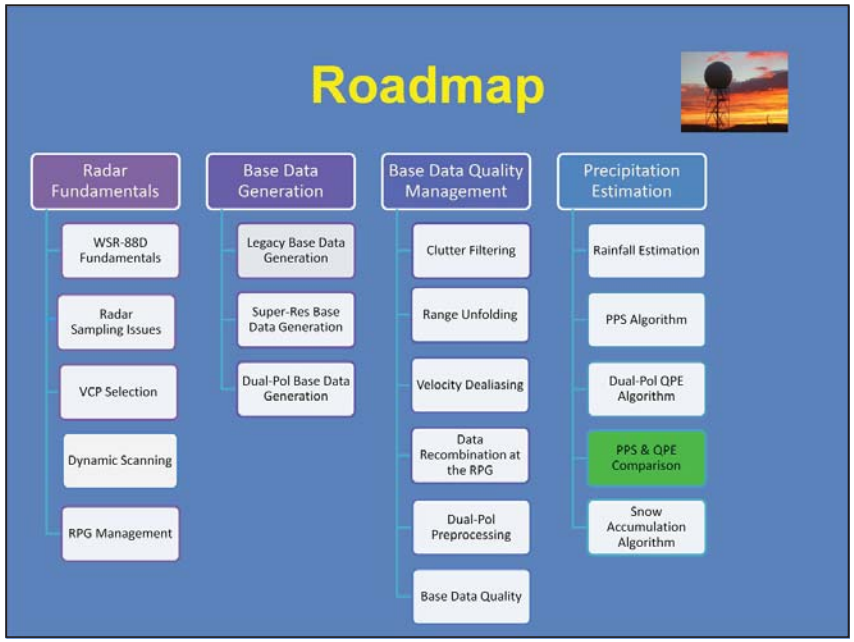
If an assigned hydroclass value is invalid, that will increase the error of the rainfall estimate.

For long term stratiform events, especially in winter, a discontinuity at the top of the melting layer is likely, unless the 2.8*R(Z) multiplier has been edited based on local research.

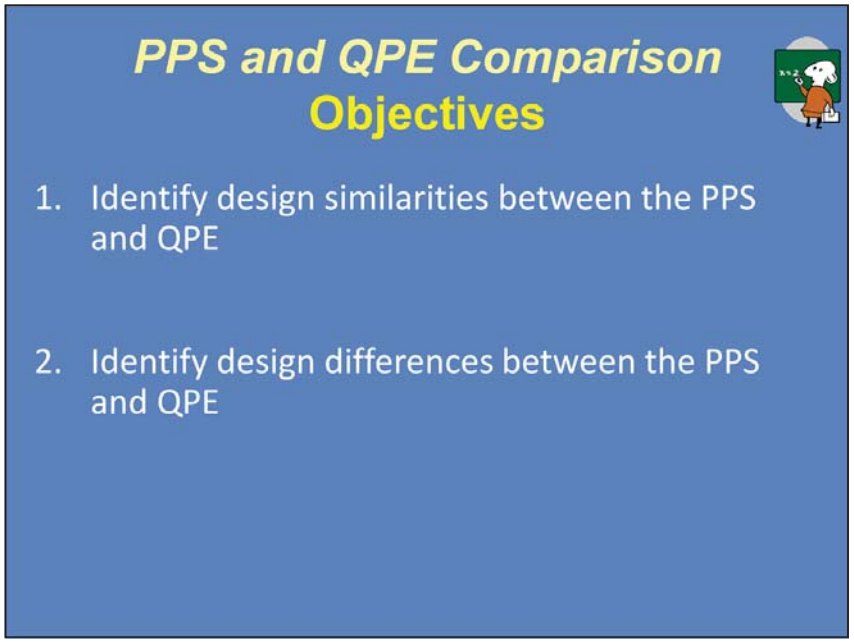
QPE is different than other RPG algorithms in terms of its reliance on the output from both the MLDA and the HCA. This means that using QPE in operations requires monitoring (at least) the HHC product to ensure that the outputs "make sense".



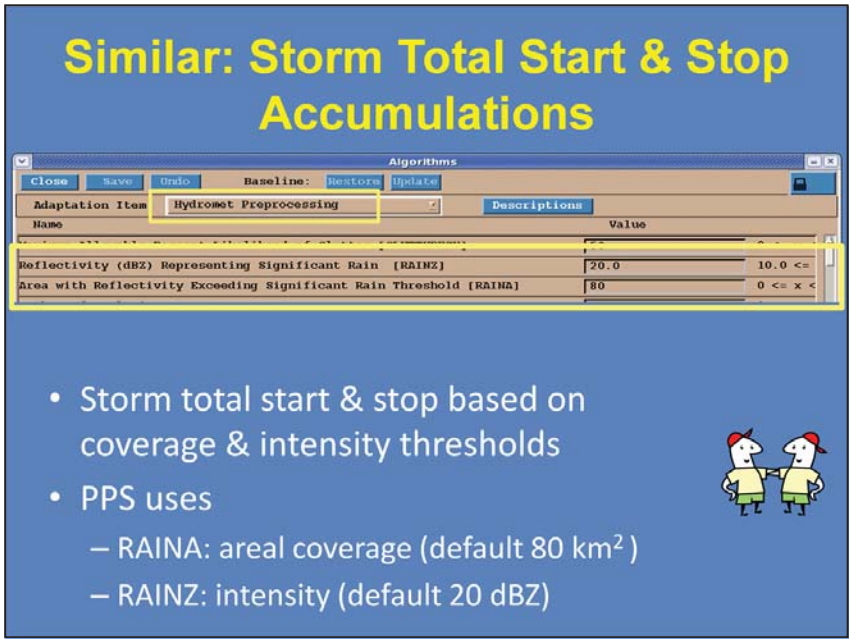
Welcome to the PPS and QPE Comparison lesson

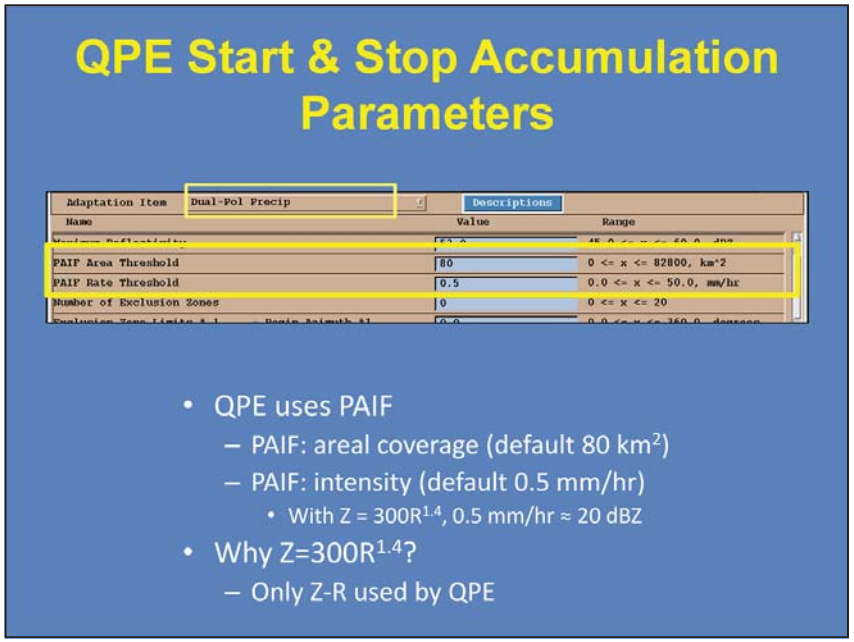


Here is the “roadmap” with your current location



There are two objectives for PPS and QPE Comparison, and they will be taught in sequence during this module.





The next several slides present the similarities and differences between the PPS and the QPE.

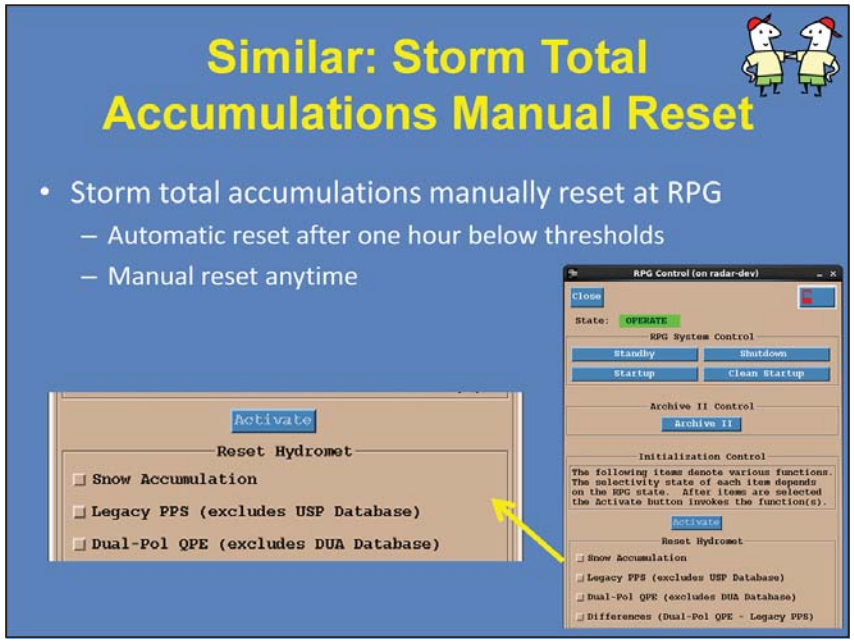
The first similarity between the PPS and QPE is how the storm total accumulations start and stop. They both use the same concept. If the radar returns exceed thresholds of areal coverage and intensity, accumulations begin. Once the returns fall below the thresholds for one hour, accumulations stop.

Here are the thresholds for the PPS, along with their default values. RAINA is the areal coverage, while RAINZ is the intensity, or dBZ, threshold. RAINA and RAINZ are accessed from the Algorithms window at the RPG under "Hydromet Preprocessing".

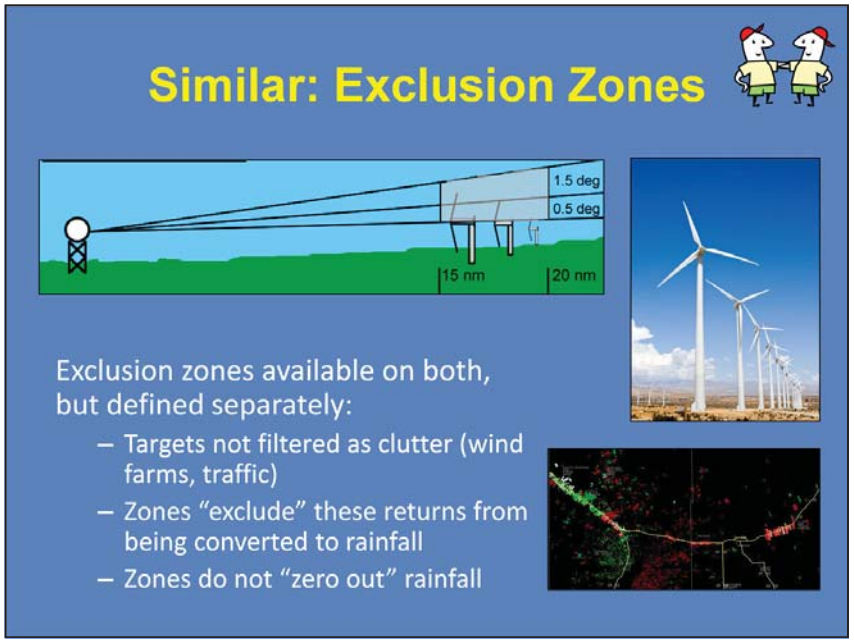
QPE uses the same concept for start and stop of storm total accumulations, it just uses different threshold names. The QPE thresholds start with the Precipitation Accumulation Initiation Function (PAIF). The PAIF Area Threshold is analogous to RAINA, while the PAIF Rate Threshold is analogous to RAINZ. Though the units differ, the default values are the same as for the PPS. The PAIF area threshold is 80 km^2, while the PAIF intensity threshold is 0.5 mm/hr. You probably don't spend your time converting dBZ to mm/hr, but for $Z=300R^{1.4}$, 0.5 mm/hr is equivalent to 20 dBZ. Why use $Z=300R^{1.4}$? In the QPE, this is the only direct Z-R relationship used.

The PAIF thresholds are accessed from the Algorithms window at the RPG under "Dual-Pol Precip".

Though the units differ, the PPS and QPE parameters for starting and stopping accumulations have the same default values.

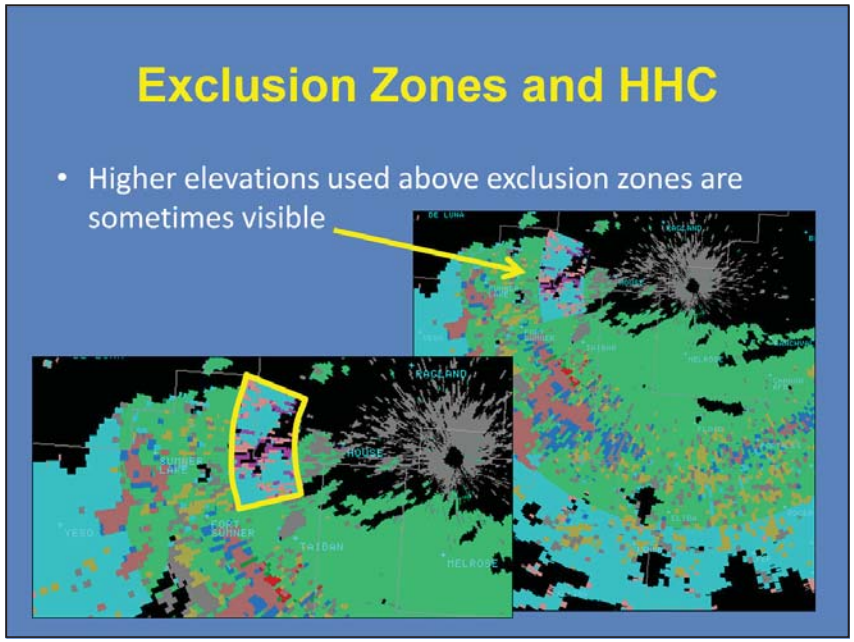


Both the PPS and the QPE have an automatic reset of the storm total accumulations after one hour of radar returns below their respective thresholds. Both the PPS and the QPE allow for a manual reset, which can be done from the RPG Control Window, selecting the “Legacy PPS” and the “Dual-Pol QPE”, respectively.

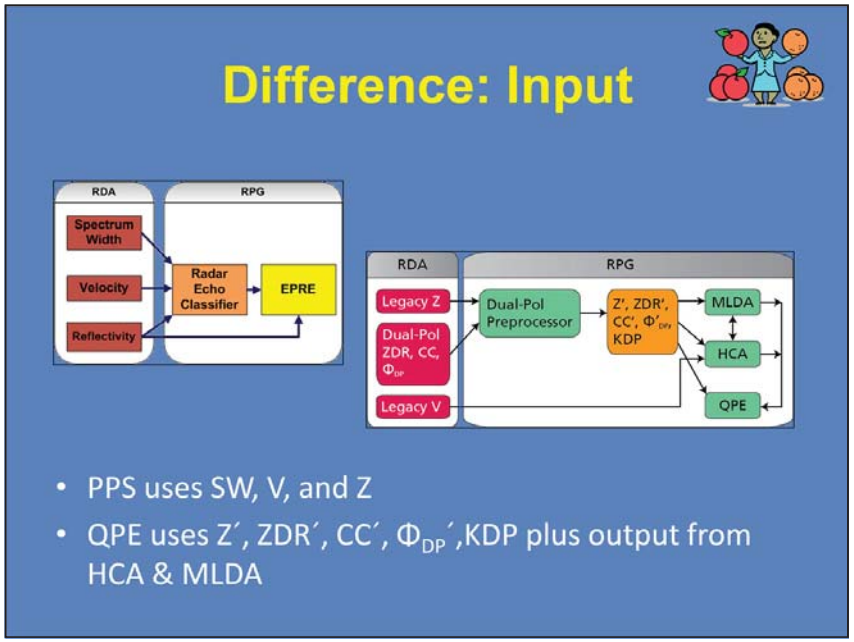


The next similarity between the PPS and the QPE is the use of exclusion zones. There are moving ground targets that are not filtered as clutter, such as rotating wind turbine blades and traffic on roads. Returns from targets like this usually exceed the intensity thresholds (RAINZ and PAIF Rate). Exclusion zones can be applied to both the PPS and QPE, preventing this type of radar return from being converted to rainfall.

Exclusion zones are an important tool, defining a volume from azimuth to azimuth, range to range, and up to (and including) a maximum elevation angle. There is a misconception that exclusion zones “zero out” rainfall estimates within the zone. If a range bin falls within an exclusion zone, PPS and QPE use the lowest elevation that is above the exclusion zone to estimate rainfall.



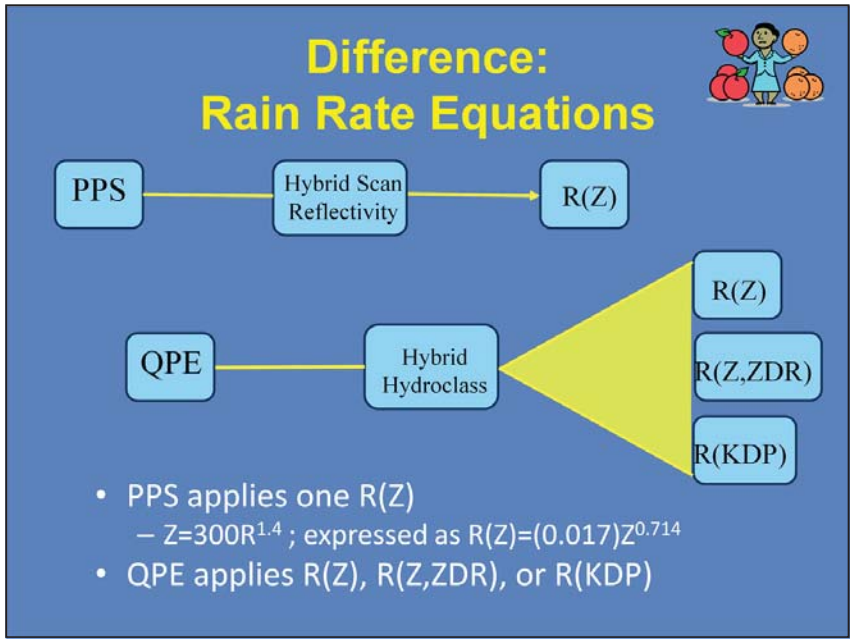
Here is an example where the use of exclusion zones for QPE becomes evident in the HHC product. Higher elevations are used by QPE for the azimuths and ranges within the exclusion zone, and that can sometimes be reflected in the hydroclass types on the HHC. In this case, you can see the transition from light to moderate rain (light green) to dry snow (light blue) at a consistent range from the radar. The exception is the block just to the west of the radar, which has mostly dry snow and some graupel. This is the result of the higher elevation above the exclusion zone intersecting the melting layer.



The Legacy PPS relies on the legacy base data: spectrum width, velocity and reflectivity. Though reflectivity is the only input for conversion to rainfall, spectrum width and velocity are used by the Radar Echo Classifier to identify bins that likely contain clutter.

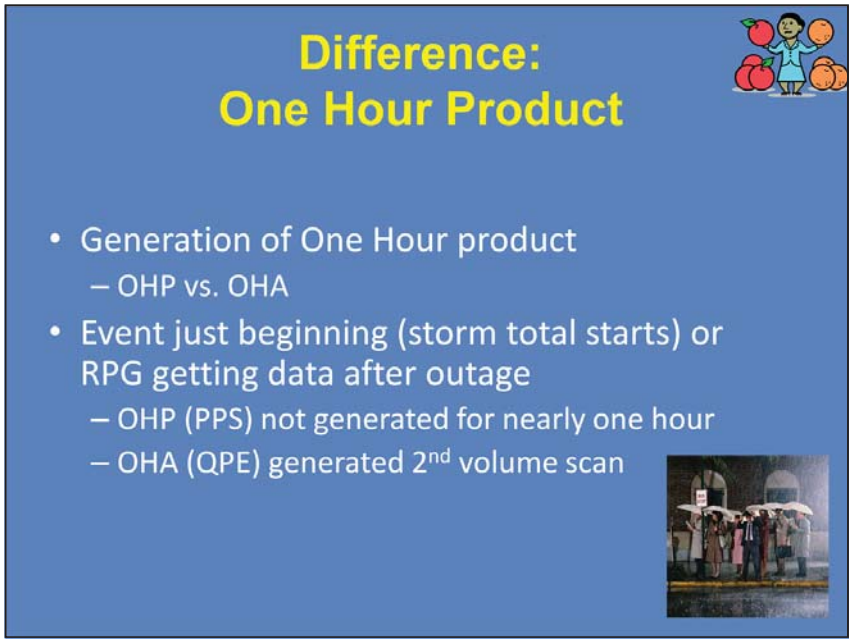
QPE relies on the reflectivity and Dual-Pol base data, but only after it has passed through the RPG Dual Pol Preprocessor algorithm. QPE also uses the output from the HCA and the MLDA to choose the rain rate equation for each range bin. HCA and MLDA help to prevent non-meteorological returns from being converted to rain rate and to determine the best rain rate equation for the particular range bin.

So there is a significant difference with the inputs to the PPS, compared with the QPE.



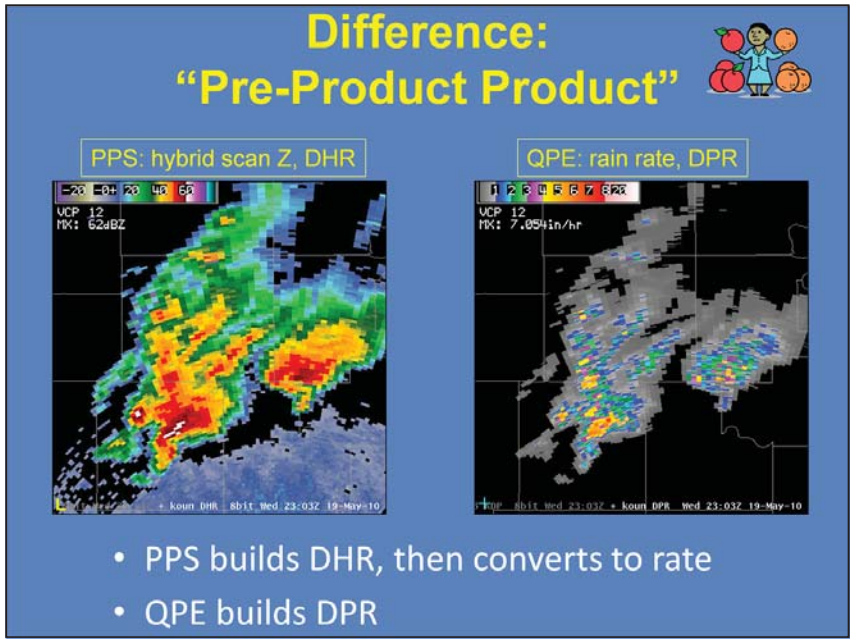
A key difference between the PPS and the QPE is the approach for converting base data to rainfall rate.

The PPS relies solely on a Z-R relationship, with editable Z-R relationship parameters. The familiar $Z=300R^{1.4}$ can be represented as $R(Z)=(0.017)Z^{0.714}$, which is a better representation to understand the QPE rainfall rate equations. The 3 different QPE rain rate calculation methods are each based on different inputs, and the notation tells you the input. These equations are selected based on the type of hydroclass assigned on the HHC product. For some hydroclass values, R(Z) is used with a multiplier.



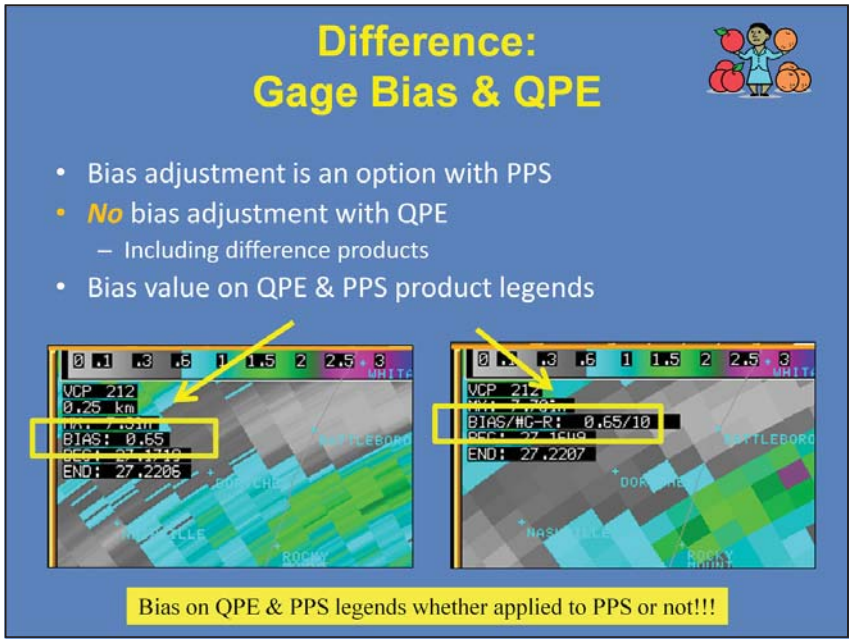
There is a significant difference between the PPS and QPE on the generation of their respective one hour products, the OHP and the OHA. This difference occurs for the beginning of an rainfall event or the return of base data to the RPG after an outage. The beginning of an event means that the storm total thresholds have been exceeded and accumulations have begun. The return of base data to the RPG after an outage means that there has been some kind of failure (wideband or RDA) that prevents base data from getting to the RPG.

In either case, the PPS will not generate an OHP for nearly one hour, while the OHA will be available beginning with the 2nd full volume scan.



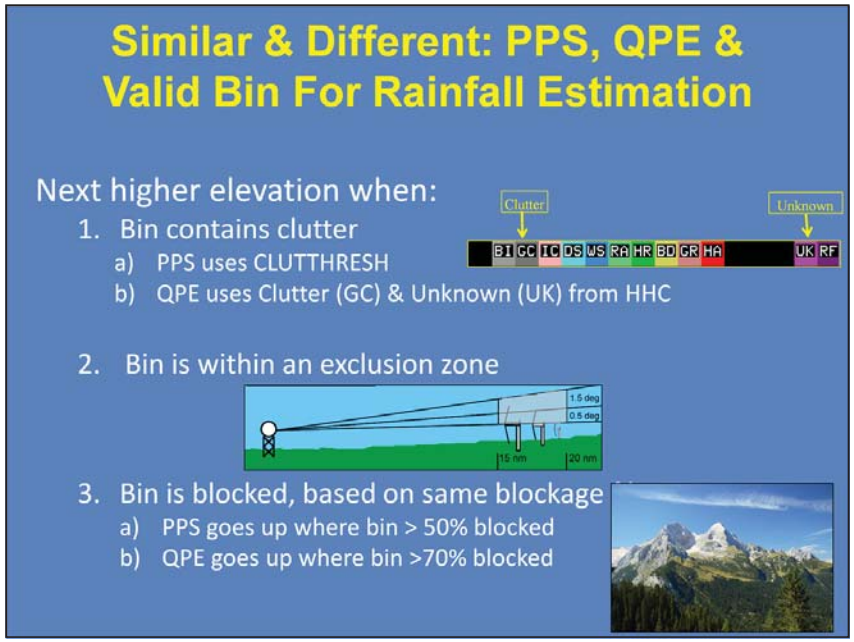
Both the PPS and QPE end up with rain rates assigned to every range bin before the products are built, but the respective methods are different, and the products available to see what was used to generate the accumulations differ.

In terms of a "pre-product" product, the PPS gives you the Digital Hybrid Scan Reflectivity (DHR), which is the dBZ value for each range bin before it is converted to rainfall rate (example on the left). The QPE gives you the rainfall rate directly, via the Digital Precipitation Rate (DPR), which is the instantaneous rate for each range bin that is used for the product accumulations (example on the right).



There is also a difference between the PPS and the QPE with respect to the application of a rain-gage bias. Bias adjustment is an option with the PPS, but there is no such option to apply a rain-gage bias to the QPE accumulations. All of the QPE-generated products are un-biased, including the QPE Difference products. You will learn more about the QPE products in the topic on base and derived products.

Unfortunately, a bias value is shown in the product legends of both the PPS and QPE products, irrespective of whether the bias has been applied to the PPS. For QPE products, always disregard the bias value. For the PPS products, if there is a decision to apply the bias, it must be communicated to the current staff and to forecasters on subsequent shifts.



In a general way, the approach for selecting a bin to be converted to rainfall is similar for the PPS and the QPE. There is some difference in the implementation. There are three requirements that can cause the PPS or the QPE to use a higher elevation angle.

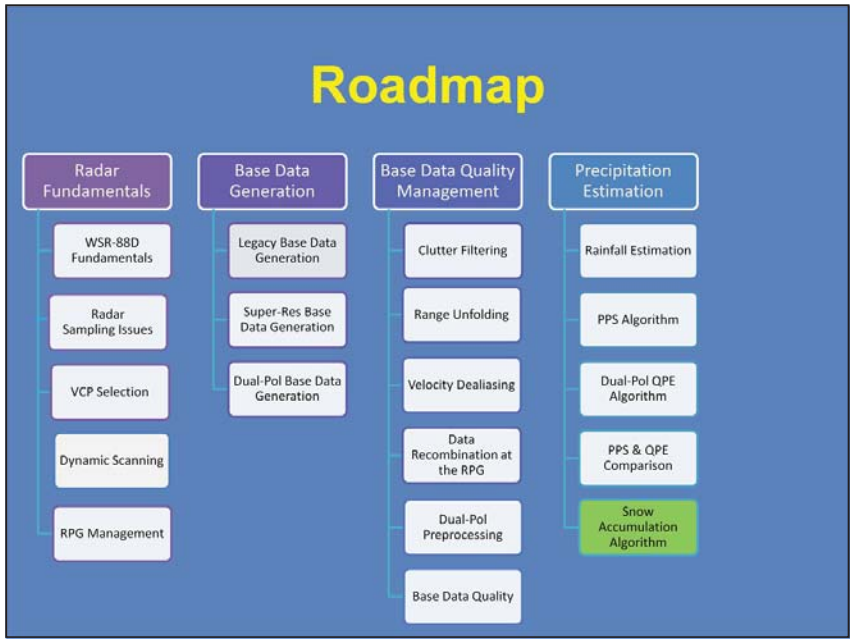
1. If the bin is suspected of containing clutter, the next higher elevation is used. For the PPS, this decision is based on whether the output of the Radar Echo Classifier exceeds the setting of CLUTTRESH. For the QPE, this decision is based on whether ground clutter (GC) or Unknown is assigned from the HHC.

2. If the bin falls within an exclusion zone, the next higher elevation is used. It is recommended that the same exclusion zones be defined for both the PPS and the QPE.

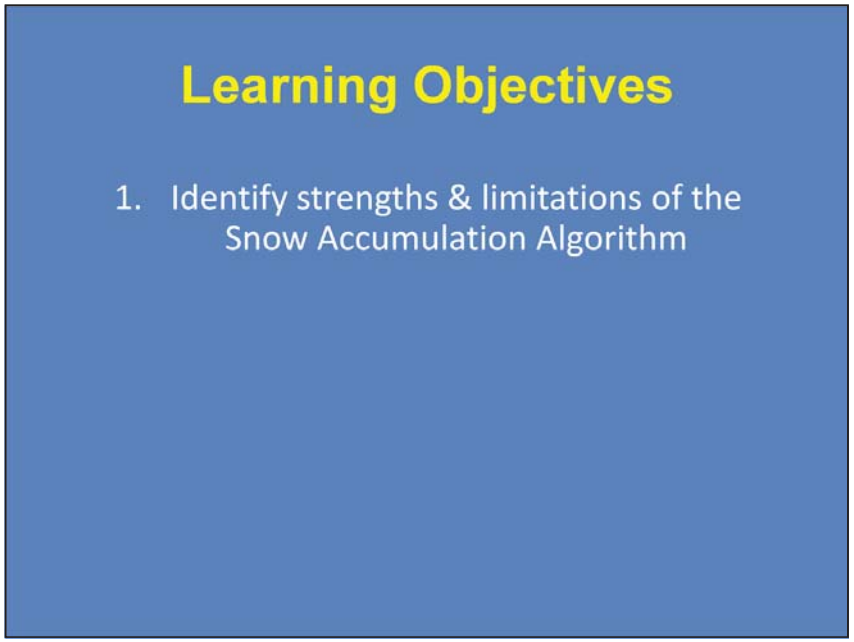
3. Based on the same blockage data file, if the bin is partially blocked beyond a threshold, the next higher elevation is used. For the PPS, the next higher elevation is used if the partial blockage exceeds 50%. For the QPE, the next higher elevation is used if the partial blockage exceeds 70%.



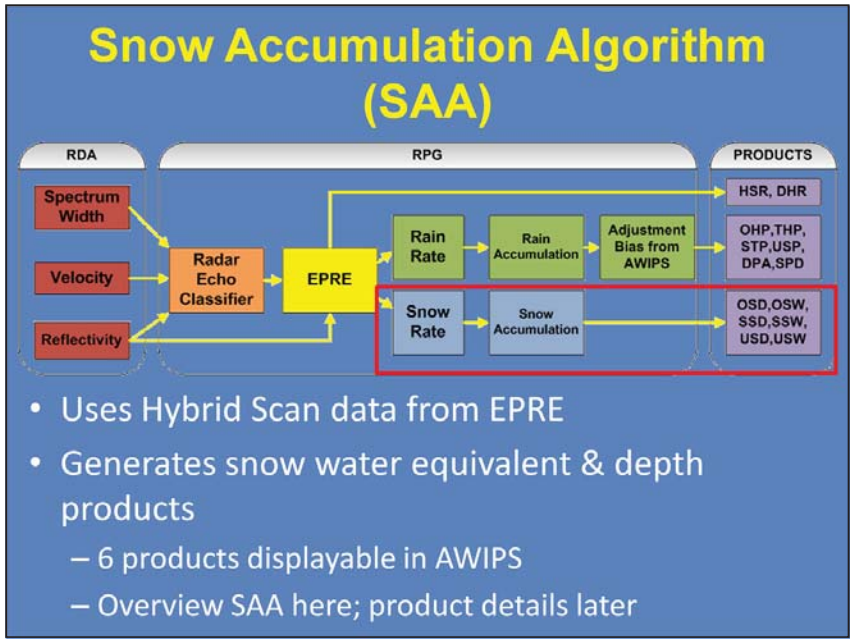
Welcome to the lesson on the Snow Accumulation Algorithm.

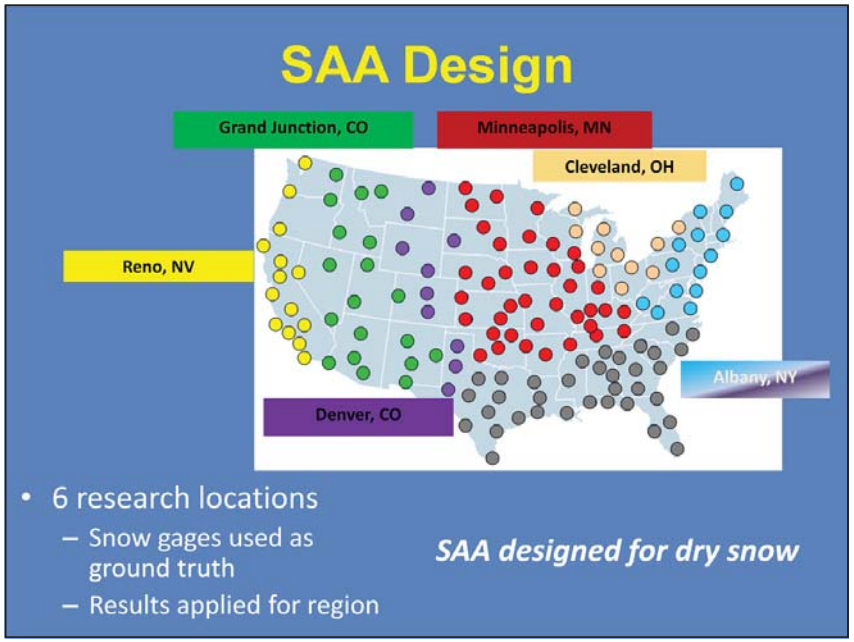


Here is the Roadmap with your current location.



There is one objective for this lesson.





We now focus on the generation of snow water equivalent and snow depth products from the Snow Accumulation Algorithm (or SAA).

The SAA uses hybrid scan data from the Enhance Preprocessing Algorithm (EPRE)

After data processing, this algorithm produces 6 total products in AWIPS that represents values for snow water equivalent and depth over various time ranges.

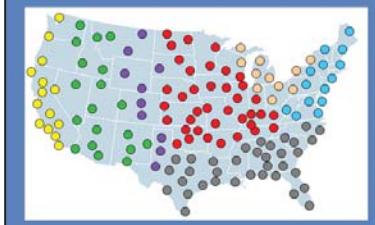
Currently 6 products in AWIPS and more information on the products will be presented in a later lesson.

The SAA was developed from data collected at 6 different research locations. At each of these locations, a network of high quality snow gages was used as ground truth against the radar snowfall estimates. The output of this research provided default adaptable parameters that are used at each of the regions on the map. Note that there was not a research site selected in the southern United States. The data for Albany was selected for use in both the Northeast and Southern United States.

One of the most important assumptions with the SAA is that it was designed for dry snow events.

Z-S Relationships

- Reflectivity (Z) to rate of snow water equivalent (S)
- Same default Z-S relationship for each region



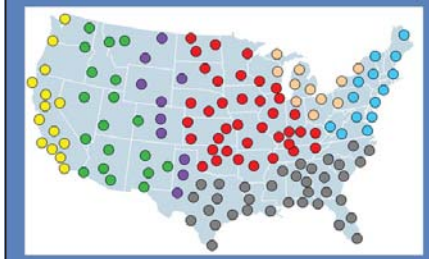
Research Location	Relationship
Albany, NY	$Z = 120\ S^2$
Cleveland, OH	$Z = 180\ S^2$
Minneapolis, MN	$Z = 180\ S^2$
Denver, CO	$Z = 130\ S^2$
Grand Junction, CO	$Z = 40\ S^2$
Reno, NV	$Z = 222\ S^2$

Similar to how there are Z-R relationships to estimate rainfall rates from reflectivity, there are Z-S relationships developed to estimate snow water equivalent from reflectivity. Using the EPRE Hybrid Scan data as input, the returned power is plugged into a Z-S relationship using coefficients developed from one of the regional research locations.

This table lists the default Z-S relationships for each research location.

SAA and Snow Ratio

- Snow ratio
  - Water equivalent to snow depth
- Same default Snow Ratio for each region



Research Location	Snow Ratio
Albany, NY	11.8
Cleveland, OH	16.7
Minneapolis, MN	11.8
Denver, CO	13.3
Grand Junction, CO	14.3
Reno, NV	8.0

The snow ratio used for converting snow water equivalent to snow depth is another adaptable parameter and the default values for each region are listed here. Even within a given region, it is expected that the appropriate ratio will vary from event to event.

SAA Adaptable Parameters



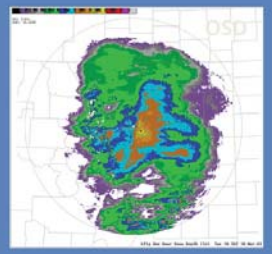
- Default values based on region
- Seven URC adaptable parameters
  - Z-S, Snow Ratio, Height Correction



All the sites within a given region have the same default values for the SAA adaptable parameters. There are seven SAA adaptable parameters that are editable under URC guidelines in the RPG. Here you can edit the coefficients in the Z-S relationship, and modify the snow-water ratio.

SAA Products

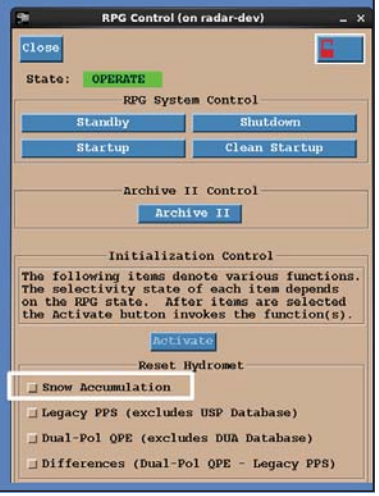
- 1km x 1°; 16 data levels; range of 124 nm
  - OSW: One Hour Snow Water Equivalent
  - OSD: One Hour Snow Depth
  - SSW: Storm Total Snow Water Equivalent
  - SSD: Storm Total Snow Depth
  - USW: User Selectable Snow Water Equivalent
  - USD: User Selectable Snow Depth



There are six snow products generated, all with 16 data levels, a resolution of 1km by 1 degree and a range of 124 nm. There are three durations: one hour, storm total and user selectable. For each of these durations, there is a snow water equivalent and a snow depth product. Examples of the products will be shown with more information in a later lesson.

Begin and End of Snowfall Accumulations

- ***No*** automatic reset of snow accumulations
- Must be ***manually*** reset prior to event



The SAA is designed to be event driven, and there is no automatic reset of the accumulations. This means the snow accumulations must be reset to zero at the beginning of an event. Resetting is done at the RPG Control window, just as with a reset of the PPS or QPE storm total accumulations.

SAA Strengths



- Only source of real time high resolution snowfall accumulations
- Uses best possible reflectivity (close to ground) to convert to snowfall
- Accumulations can be reset to zero as needed
- Available Z-S relationships and snow ratios are editable

The first strength of the SAA is that it is the only source for real time high resolution snowfall accumulations.

Since the SAA also uses EPRE as input, the SAA uses a reflectivity value closest to the ground that is not from clutter and is not blocked by the beam.

At the RPG, snow accumulations can be reset to zero as needed.

The Z-S relationships and snow ratios are editable.

SAA Limitations



- SAA designed for dry snow
- Ground truth likely needed
  - to verify precipitation type
  - to determine onset of accumulation
- No automatic reset of accumulations; must be done manually
- Available Z-S relationships and snow ratios may not be representative

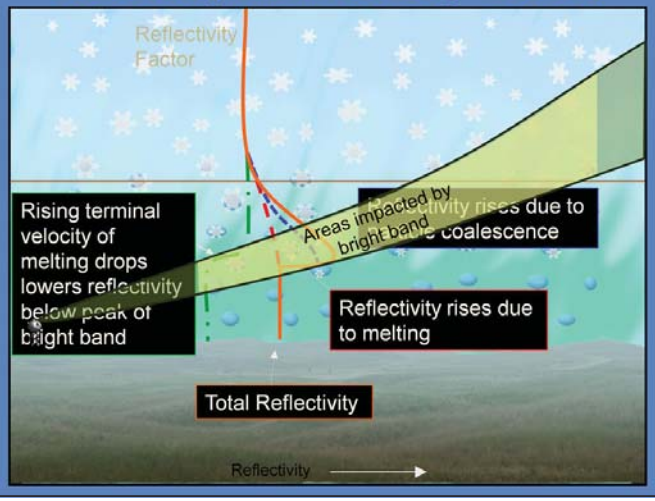
Perhaps the most important thing to remember about the SAA is that it designed for dry snow, snow that does not melt as it falls or when it hits the ground.

Ground truth will likely be needed to verify precipitation type and to determine the onset of snow accumulation.

The onset must be known in order to reset the snow accumulations, which must be done manually.

Finally, the default Z-S relationship were developed at specific locations and applied across an entire region. So both the Z-S relationship and the snow ratio may not be representative for your CWA and may require adjusting.

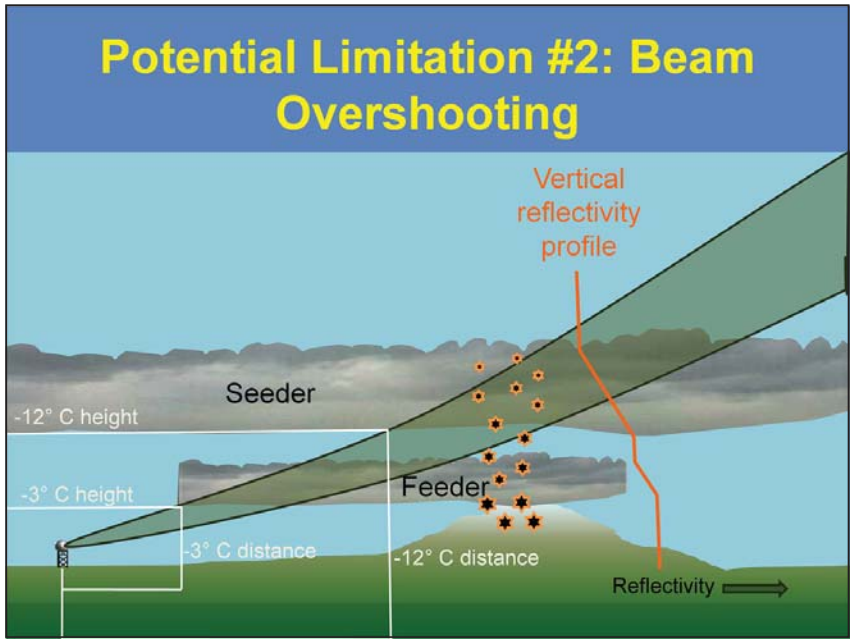
Potential Limitation #1: Assumption of Dry Snow



The Z-S equation in the Snow Accumulation Algorithm is based on the assumption that the precipitation at the surface is dry snow. We know this assumption isn't true in the bright band, and here is a closer look at what's going on.

As snow flakes approach the melting layer, liquid resides on their ice surfaces. The increased water coating helps colliding ice particles to stick together and snow flakes begin to increase in size. Larger particles form and the radar reflectivity increases. The liquid water coating itself also helps to increase radar reflectivity. An offset to the increasing reflectivity occurs when the terminal velocity of these particles increases with melting. Increasing terminal velocity increases the separation between hydrometeors and lowers the reflectivity.

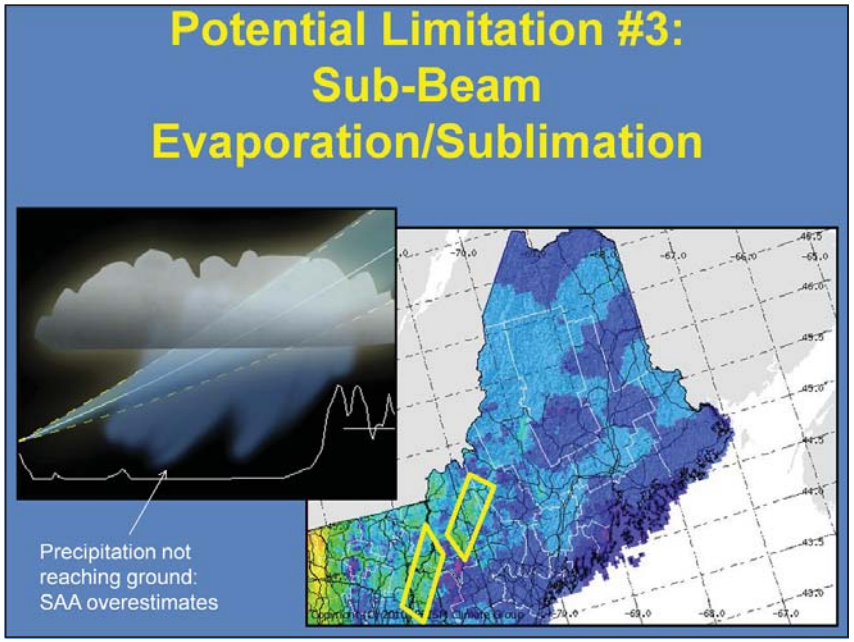
Due to these mechanisms, overestimation can occur in areas where the radar samples the bright band because wet snow has a higher reflectivity than dry snow. This overestimation can occur even when only portions of the beam (i.e., the top or bottom) are sampling the bright band.



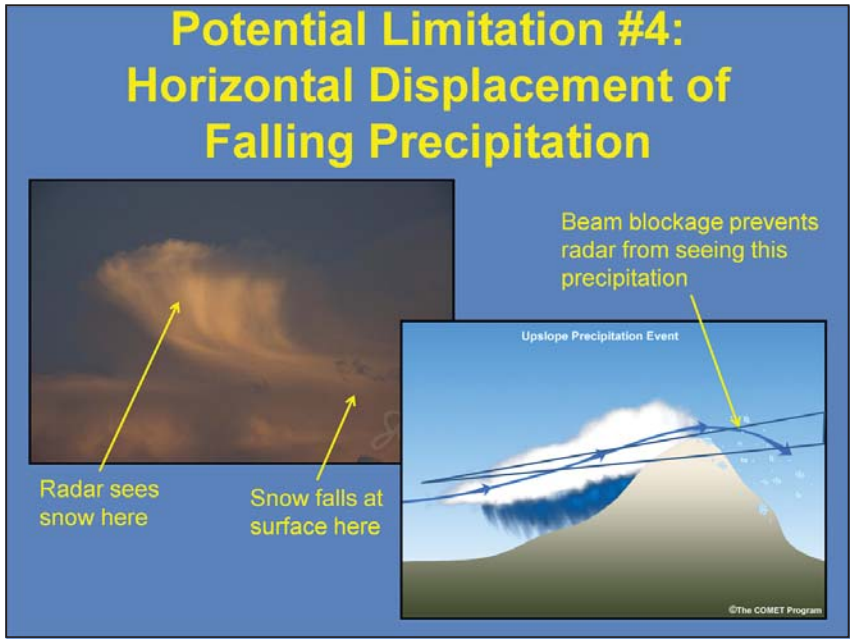
Assume we have a cloud pictured here actively generating precipitation so the intensity increases from top to bottom. Reflectivity begins to degrade once the top of the radar beam is above this precipitation production layer, and the signal is lost once the bottom of the beam is above the cloud.

The dendritic growth zone (i.e., -12 to -18 degrees Celsius) is usually the region of maximum growth. However, high cloud liquid water content between the dendritic growth zone and the bright band can contribute significant amounts of riming and needles. Collision-coalescence becomes more active in warmer, saturated clouds. Any of these precipitation production zones can be shallow, causing reflectivity to degrade quickly as range from the radar increases.

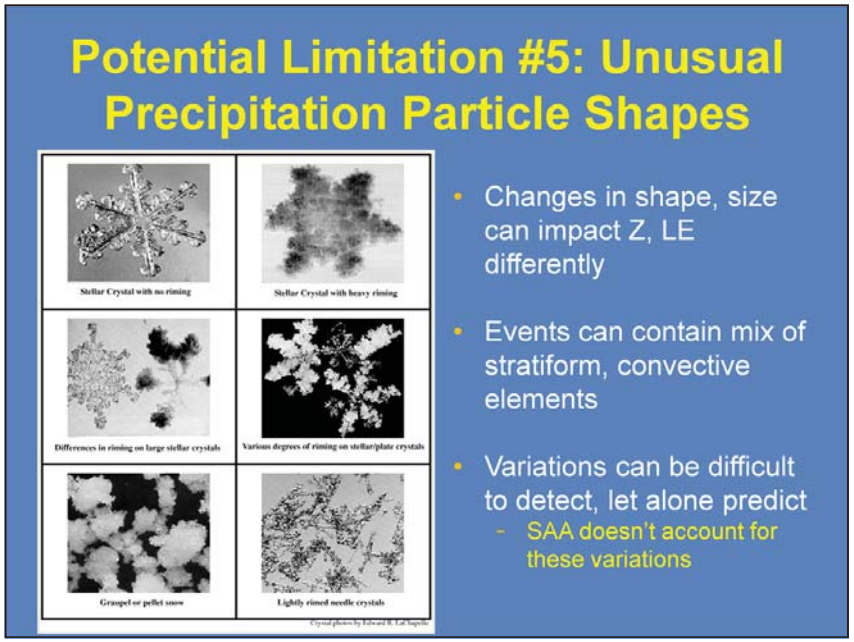
A particularly acute example of this problem is with orographic precipitation. With precipitation very close to upslope terrain, it's difficult to separate ground returns from real precipitation. So, even when the beam does sample the precipitation generation regions, reflectivity can be degraded.



Another data quality issue to consider is sub-beam evaporation (or sublimation). When sub-beam evaporation/sublimation occurs, expect the Snow Accumulation Algorithm to overestimate the liquid equivalent precipitation at the surface. The problem is most common when the near-surface airmass is dry (such as in areas of downslope winds or valleys/basins with locally low elevations). These areas may have lower annual precipitation totals such as areas north of the White Mountains in Northern New Hampshire and adjacent parts of Maine shown in the figure on the right.



Under strong, sub-beam horizontal wind conditions (or areas of strong sub-beam vertical wind shear), precipitation may drift horizontally a long distance before reaching the ground. In some cases, the radar may not observe the precipitation that reaches the ground, such as when there is lee side spillover of orographic precipitation.



Precipitation particle shape and size can significantly alter reflectivity without a corresponding change in the liquid equivalent precipitation rate. If you experience snow events, take a look at how the shapes and sizes of the snowflakes change over short periods of time.

Many precipitation systems can contain a mixture of stratiform and convective elements with variations in vertical velocity profiles with respect to the thermal profile. The result is a rapid change in particle shapes.

This error source can be very difficult to detect, let alone predict. You may notice areas of precipitation particles that vary significantly from the "average" for your region using dual-pol base data. While you could identify these areas as likely deviating from the Z-S equation being used, you can't edit for localized areas within the radar coverage area.

TAB

Radar & Applications Course

Topic:

Velocity Interpretation

Table of Contents
Topic: Velocity Interpretation

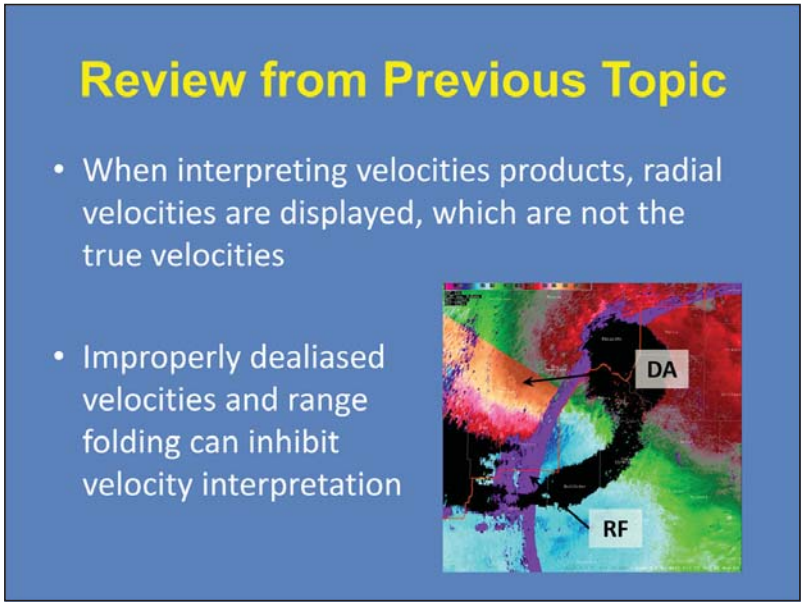
Click to jump to lesson

Lesson 1	**Large-Scale Doppler Velocity Patterns**
Lesson 2	**Storm-Scale Doppler Velocity Patterns**

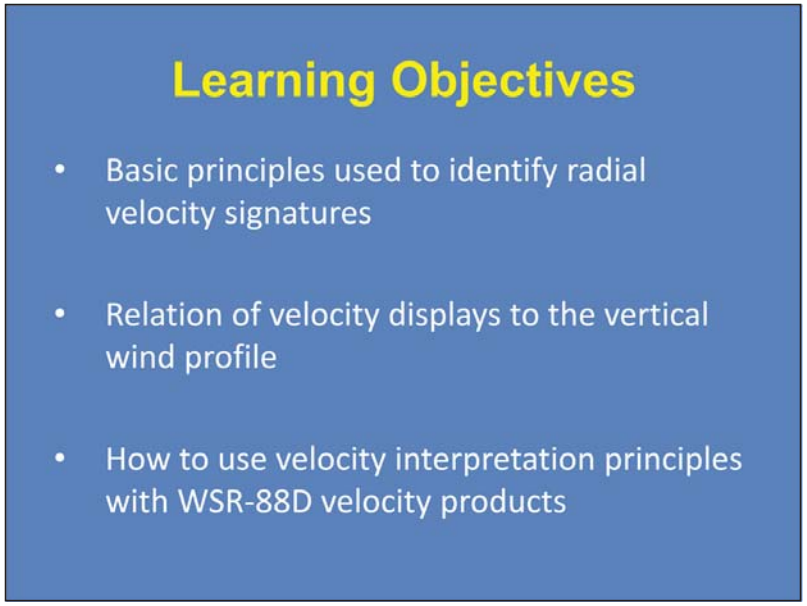


Hi, my name is Jill Hardy and welcome to this Topic on Velocity Interpretation. This lesson will cover large-scale Doppler velocity patterns, while the next lesson will focus on storm-scale patterns.

We have a guest speaker for this lesson: Steve Martinaitis of OU CIMMS at NSSL. But if you have any questions regarding the material, please feel free to contact me, or the RAC team. Our contact information will be on the next slide.



When interpreting velocity data, the radial velocities are displayed, which are not the true velocities. Also, the RPG dealiases velocities while the RDA performs range folding. These tasks are effective most of the time, but failures do occur, which can inhibit your ability to interpret velocity products.



There are three learning objectives for this lesson. By the end of the lesson, you should be able to understand the basic principles used to identify radial velocity signatures, relate the velocity display to the vertical wind profile, and apply velocity interpretation principles to the WSR-88D velocity products.



There are also three performance objectives with this lesson. The first objective is to be able to interpret Doppler velocity patterns under uniform, non-uniform, ambiguous, and meteorologically complex conditions. You will be identifying inbound vs. outbound velocities, constant wind speed and direction, wind speed and direction changing with height, velocity maxima, confluence and diffluence, vertical discontinuities, and boundaries.

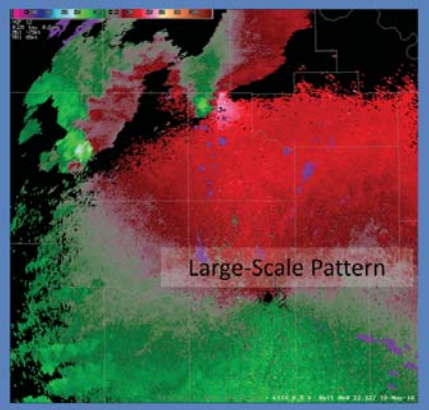
Performance Objectives

2. Construct vertical wind profiles for uniform and non-uniform horizontal wind conditions

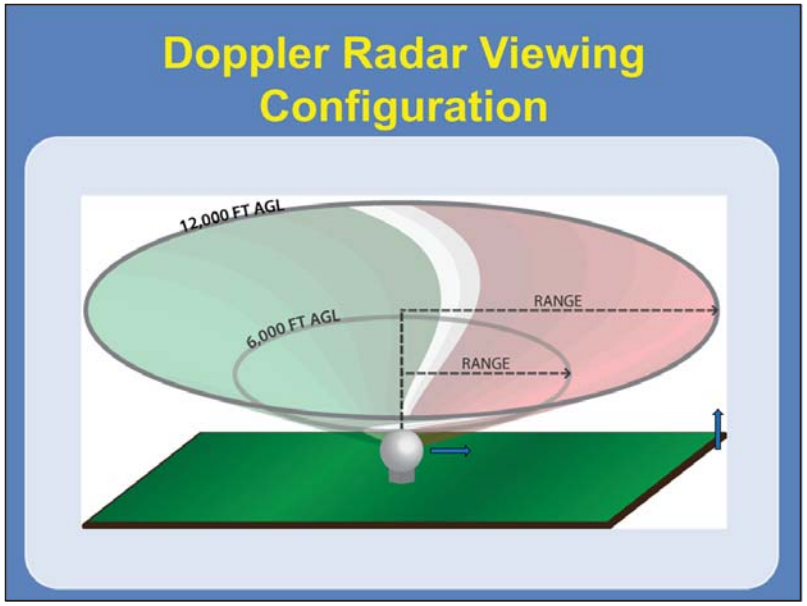
3. Assess the meteorological conditions associated with the identified velocity patterns

The second performance objective is the ability to construct vertical wind profiles for uniform and non-uniform horizontal wind conditions. And the final performance objective is to assess the meteorological conditions associated with the identified velocity patterns.

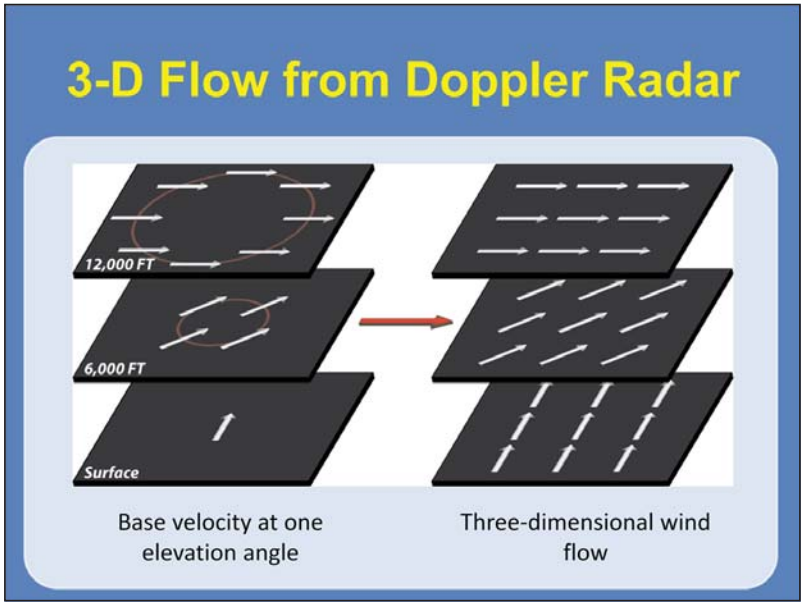
Doppler Velocity Patterns



In this lesson, we will be discussing large scale velocity patterns. The second lesson will then focus on storm-scale velocity signatures.



When you are looking at a velocity product, you are viewing the display from above looking into a cone. As you move farther away from the RDA, you are also increasing in height above the ground.

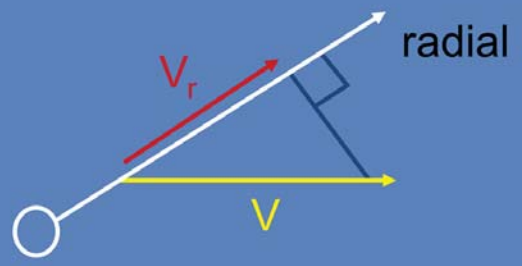


By using this concept, you can determine the wind flow at various levels and construct a three-dimensional wind profile of the atmosphere around the RDA. In this graphic here, you can see that with one elevation scan, the flow can be determined at different ranges, which are proportional to different heights. The flow at each range can then be assumed for a constant altitude.

At an AWIPS workstation, hold down the left mouse button at a selected cursor location. This will display the height in feet (AGL/MSL), azimuth, range (Statute Miles), and the radial velocity (knots).

Radial Velocities

Defined as the component of target motion parallel to the radar radial (azimuth)



Before we begin to look at conceptual models of velocity patterns, we need to discuss radial velocities. A radial velocity is defined as the component of target motion parallel to the radar radial, or azimuth. In this diagram, you can see the actual target motion (yellow arrow) and the radial velocity target motion (red arrow) along the radial (white arrow).

Radial Velocity Principles

1. Radial velocities will always be less than or equal to actual target velocities.
2. Radial velocity equals actual velocity only where target motion is directly towards or away from the radar.
3. Zero velocity is measured where target motion is perpendicular to a radial or where the target is stationary.

There are three basic principles with regards to radial velocities. The first principle is that radial velocities will always be less than or equal to actual target velocities. The second principle is that the radial velocity equals the actual velocity only when the target motion is directly towards or away from the radar. The third principle is that a radial velocity of zero is measured when the target motion is perpendicular to a radial or when the target is stationary.

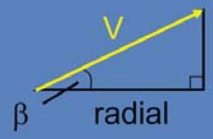
You will see why in the following slides.

Radial Velocity Equation

$$|V_r| = |V| \cos \beta$$

- V_r = radial velocity
- V = actual velocity
- β = smaller angle between V and radar radial

- When β = 0°, then $|V_r| = |V|$
- When β = 90°, then $|V_r| = 0$



The relationship between a target's actual velocity and the radial velocity depicted by the RDA can be described using the Radial Velocity Equation. Here, the absolute value of the radial velocity is equal to the actual velocity multiplied by the cosine of the angle β. The angle β represents the smaller angle between the actual velocity and the radar radial. When β equals 0°, then the radial velocity is equal to the actual velocity. When β is at 90°, then the radial velocity is zero.

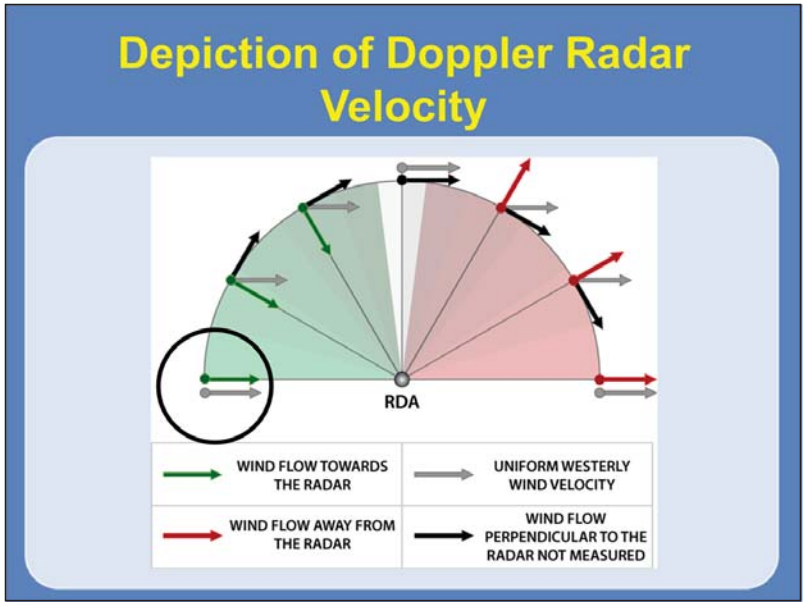
Percentage of Actual Velocity

$$|V_r| = |V| \cos \beta$$

β (degrees)	Cosine β	Percent Measured
0	1	100 %
5	0.996	99.6 %
10	0.985	98.5 %
15	0.966	96.6 %
30	0.866	86.6 %
45	0.707	70.7 %
60	0.500	50.0 %
75	0.259	25.9 %
90	0	0 %

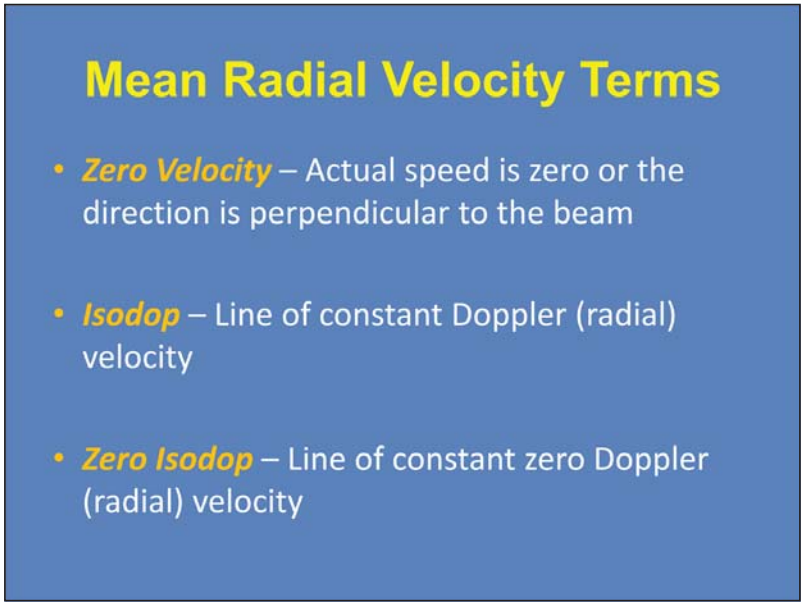
The table here compares various angles of β and the percentage of the actual target motion measured at each radial. The greater the angle between the target's velocity vector and the radar azimuth, the smaller the percentage of the target's actual velocity that will be measured and depicted on the velocity products.

Note that at a 45° angle, the radar is measuring approximately 70% of the motion, not 50%. This is because of the cosine function within the equation, which creates a non-linear relationship.

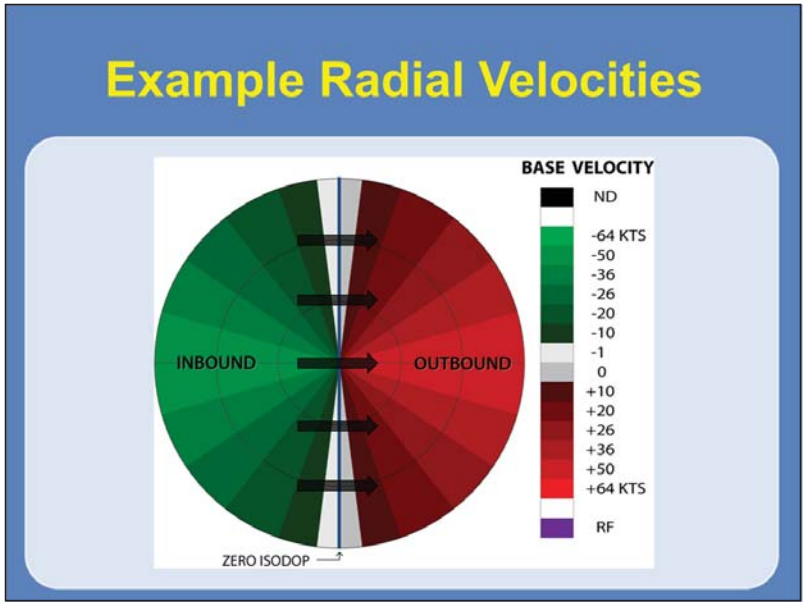


This graphic depicts the radar's ability to measure velocities and what the operator sees. When the wind is parallel to radial, the full component of the wind is measured. As the radial becomes more perpendicular to the actual wind, the radial component decreases. When the radial is perpendicular to the wind, the radar displays zero velocity. However, the actual velocity has not changed. This is the reason that the colors change or speed seems to decrease as you move away from the actual wind direction/speed.

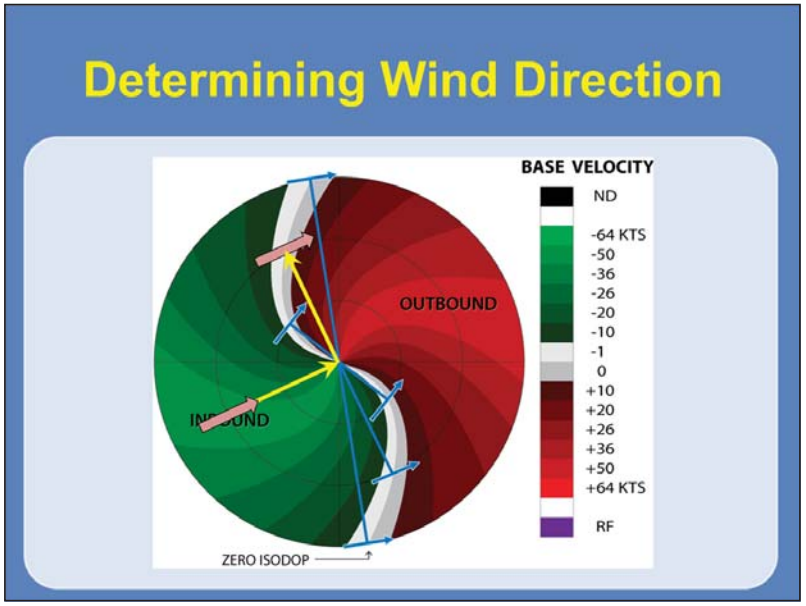
Here, you can also see that inbound velocities (positive) are depicted by cool colors (green) and outbound velocities (negative) are depicted by warm colors (red). The reason for these colors is that the first Doppler radar pointed straight up, so downdrafts (negative vertical motion) pointed towards the radar.



Before we go on, it is important that we define a few terms here. Zero velocity is when the actual speed is zero or the direction is perpendicular to the radar beam (which can also be described as zero radial velocity). An isodop is a line of constant Doppler (or radial) velocity. Finally, a zero isodop is a line of constant zero Doppler (radial) velocity.



Here is an example of radial velocities on a Plan Position Indicator (PPI) scope. The inbound velocities are in green while the outbound velocities are in red. The zero isodop represents the line of constant zero radial velocity. A straight zero isodop, as seen here, represents a uniformly directional flow at all levels. The wind direction is simply perpendicular to the radar beam at the zero isodop and is from the inbound to the outbound side. In this example, a homogeneous westerly flow exists over this area.

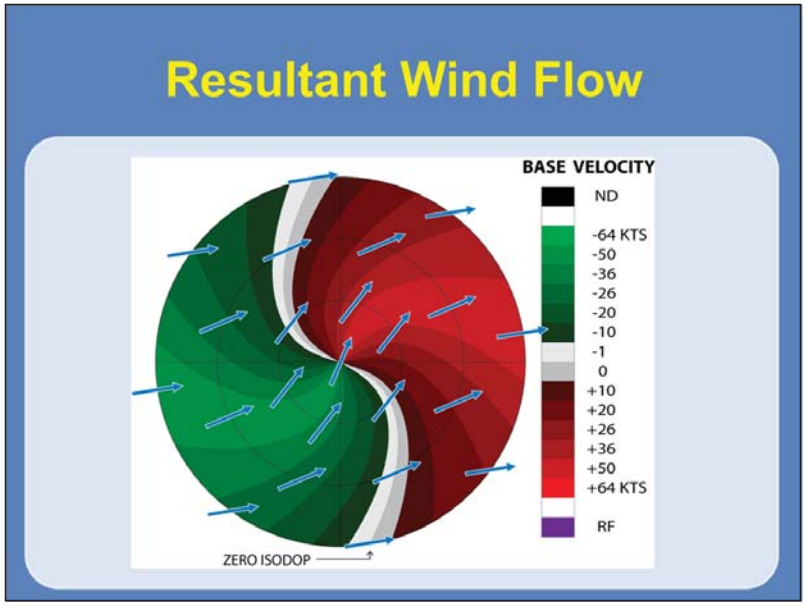


There are two methods for determining wind direction. However, both of these methods carry with it the assumption that the flow over the entire area is homogeneous for each level.

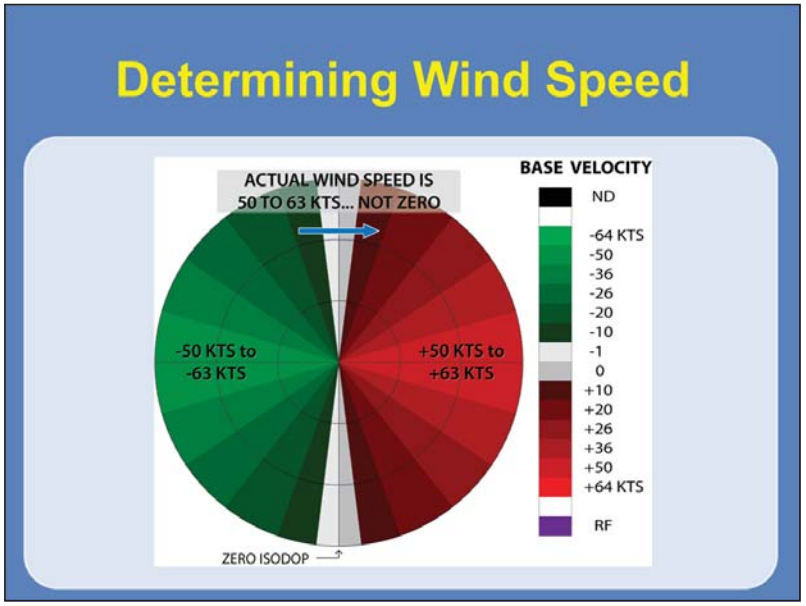
The first method uses the zero isodop to determine wind direction. First draw a line along a radial from the RDA to some point along the zero isodop. Next, draw an arrow perpendicular to the line along the radial. The arrow should be pointing from inbound to outbound velocities. Assuming homogeneous flow, the arrow represents the wind direction at that range (height).

The second method uses the direction of the maximum inbound and outbound velocities from the radar. At a certain range, draw an arrow pointing from the location of the inbound maximum velocity towards the RDA. For both methods, note that the wind direction changes at each level, so you'll need to repeat for multiple levels.

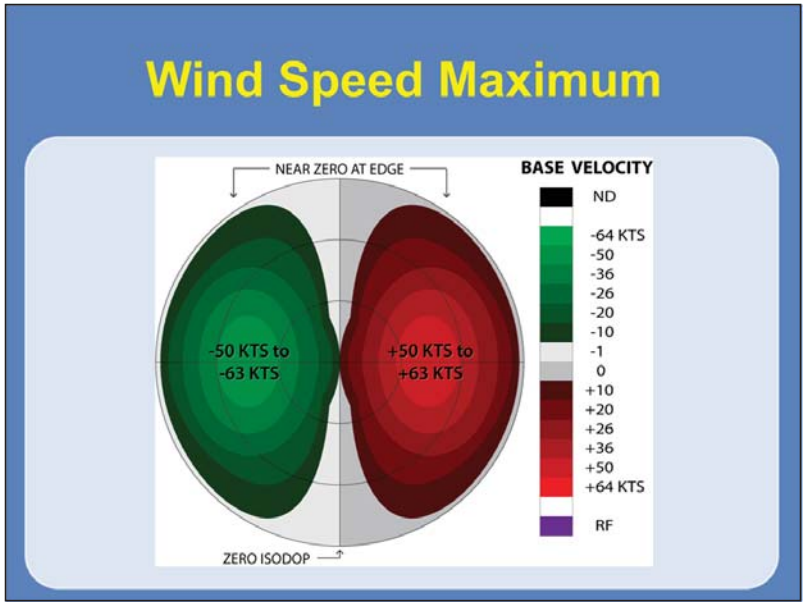
You may need to use either method to estimate the wind direction. The flow may be horizontally homogeneous over just a part of the radar, so you may need to determine the wind direction for than one region on your PPI. Also, there will be many cases when one of these methods will not work. At times, you may have to use a combination of the two methods.



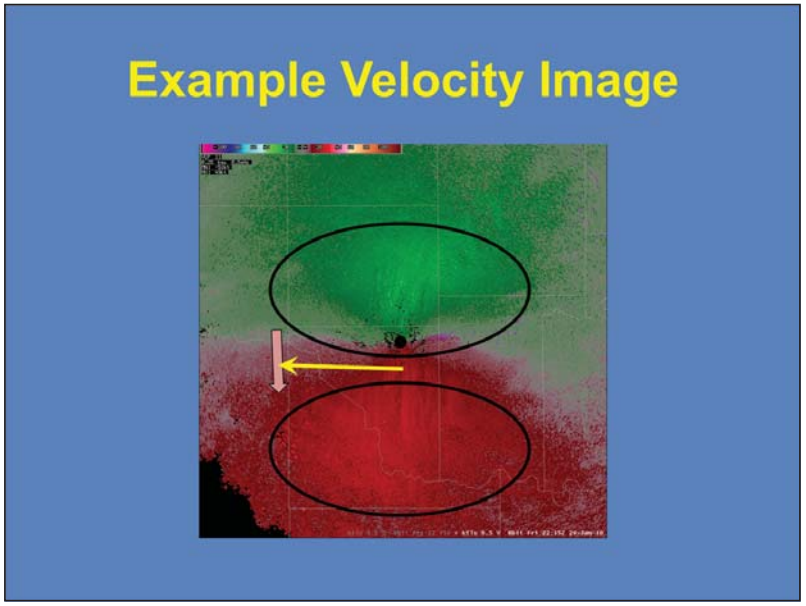
By using the methods described previously, we can determine the wind direction at any range (or height). It is important to draw the arrows perpendicular to the line from the RDA (radial), not the zero isodop itself.



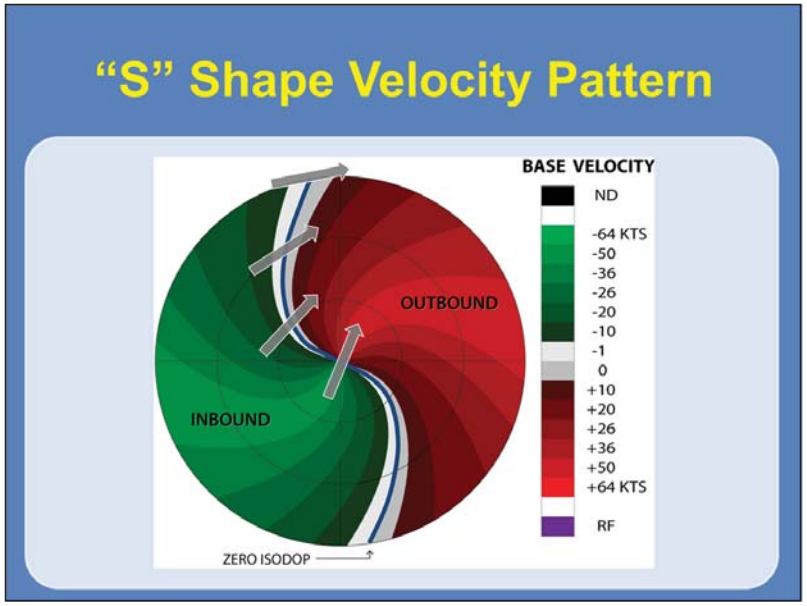
Now that we have seen how wind direction is determined, it is time to look at wind speed. Staying with the homogenous westerly flow example, we can easily tell the maximum wind speed from the radar either looking into or directly away from the wind. This example shows a constant wind speed with height of 50-63 kts. Even though the radial velocity is zero at the zero isodop, the actual velocity is still 50-63 kts.



A wind speed maximum is identified by closed isodops surrounding a maximum velocity value. Here is a depiction of a low-level wind speed maximum, as if we had a low-level jet over weak surface flow and underneath a layer of near zero winds at the height that is represented by the farthest extent of the PPI. Starting near the RDA, you can see winds increase up to a closed area of 50-63 kts between the first and second range rings. From there, they decrease to near zero at the edge of the display.

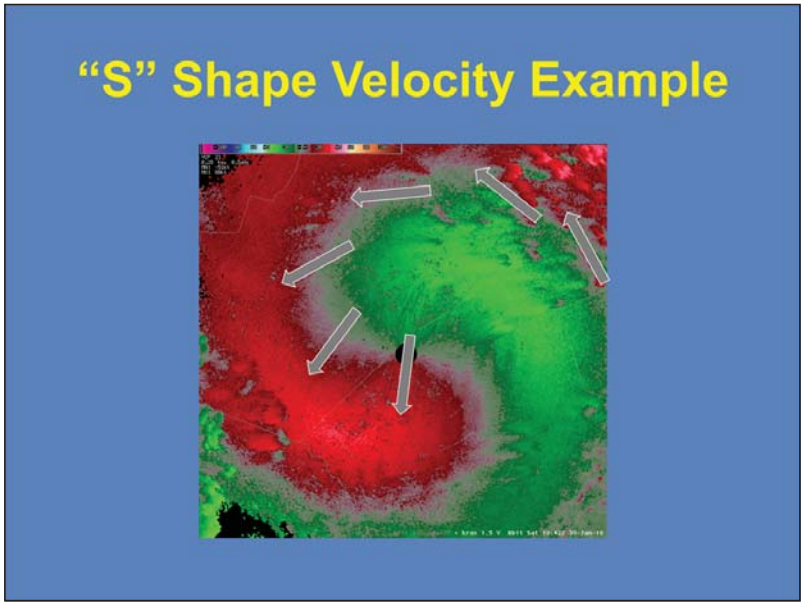


Here is a velocity example that represent some of the concepts learned in the previous slides. First, you can see a fairly straight zero isodop extend from east to west across the display. With the inbound velocities north of the RDA, the general flow is from north to south. You can also see that the wind speed maximum is close to the radar. As you move farther away from the radar in the N-S direction, the maximum velocity values decrease, signifying a low-level wind maximum.

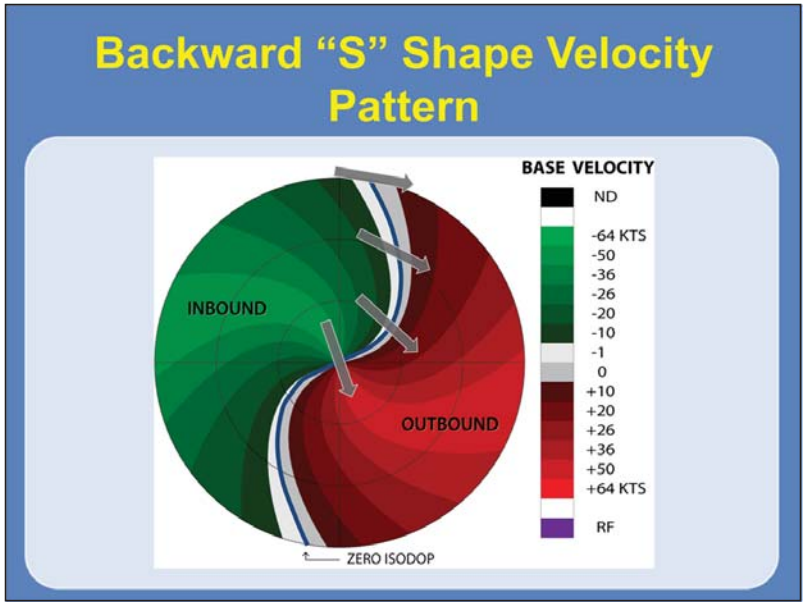


Now, we will begin looking at more complex velocity patterns.

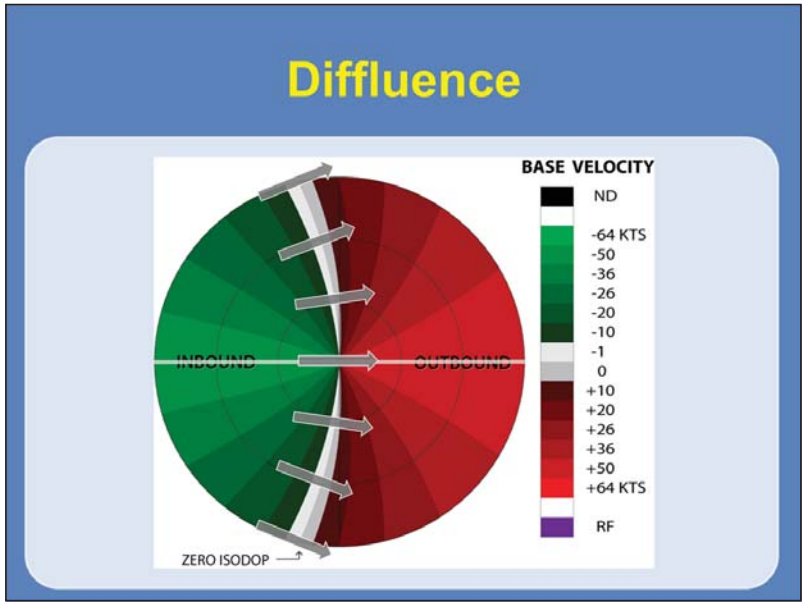
Curvature of the zero isodop represents changing wind direction with height. In this case here, the zero isodop is shaped like the letter "S." The wind direction near the RDA is from the south-southwest while the wind direction is from the west-southwest near the edge of the display. The associated vertical wind profile shown here indicates that the winds are turning clockwise with height. The meteorological term for this is veering. Veering generally indicates that warm air advection is occurring.



Here is a real-world example of veering winds with height. There is a "S" shape to the zero isodop, with winds from the north-northeast near the surface to winds from the south near the edge of the display. Thus, the winds are veering with height throughout the entire layer.

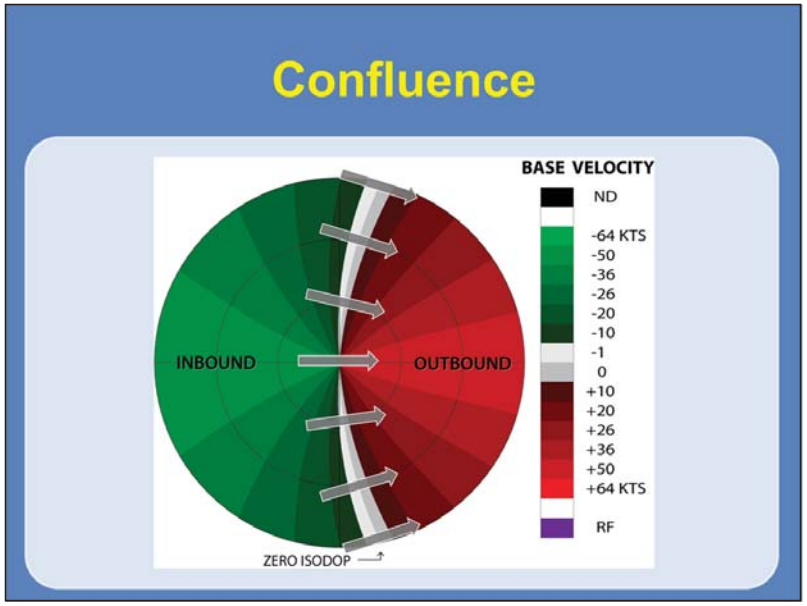


Some velocity patterns can exhibit a backward “S” shape pattern. In this case, the winds are turning counterclockwise with height, which is completely opposite of the standard “S” shape profile you saw earlier. In this graphic, the winds are from the north-northwest near the RDA and from the west-northwest near the edge of the display. When the vertical wind profile is turning counterclockwise with height, as seen here, it is referred to as backing. Backing winds are generally associated with cold air advection.

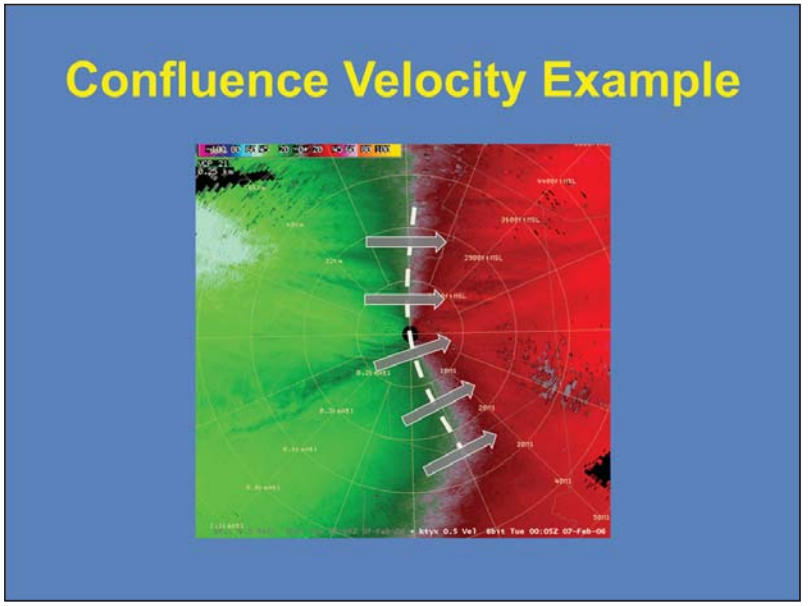


We will now diverge away from pure homogeneous wind fields and talk about diffluent and convergent wind fields. The best way to examine these kinds of wind fields is to split the display into two parts. Starting with the top half, the wind direction changes from the west near the surface to west-southwest at the edge of the display. Looking at the bottom half of the display, the wind direction is westerly near the RDA but changes to west-northwest at the southern edge of the display. This associated pattern shows the air spreading out as it passes over the RDA.

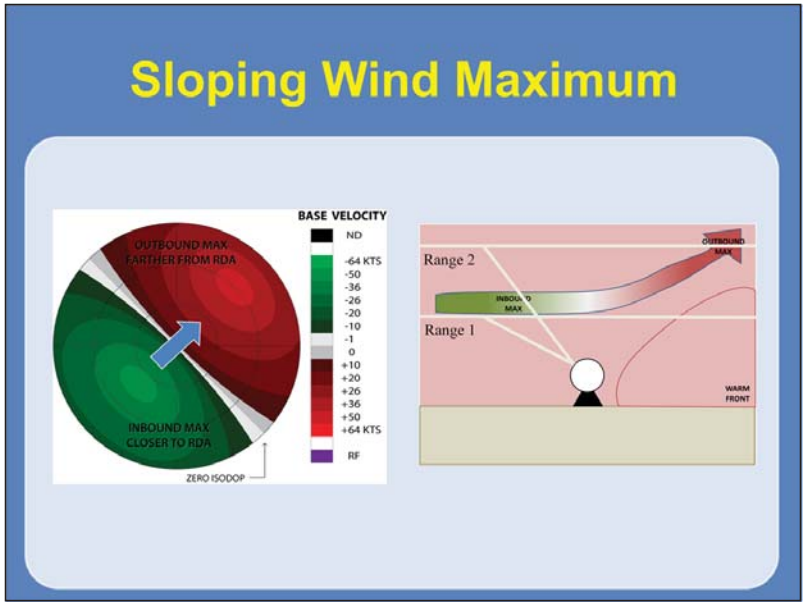
Note that the zero isodop has a bowing shape to it and that the inbound velocities are within the bow. Now, at any one point, the vertical wind profile is still going to be unidirectional; however, we now add that complexity of the wind field changing over horizontal distances (not horizontally homogeneous).



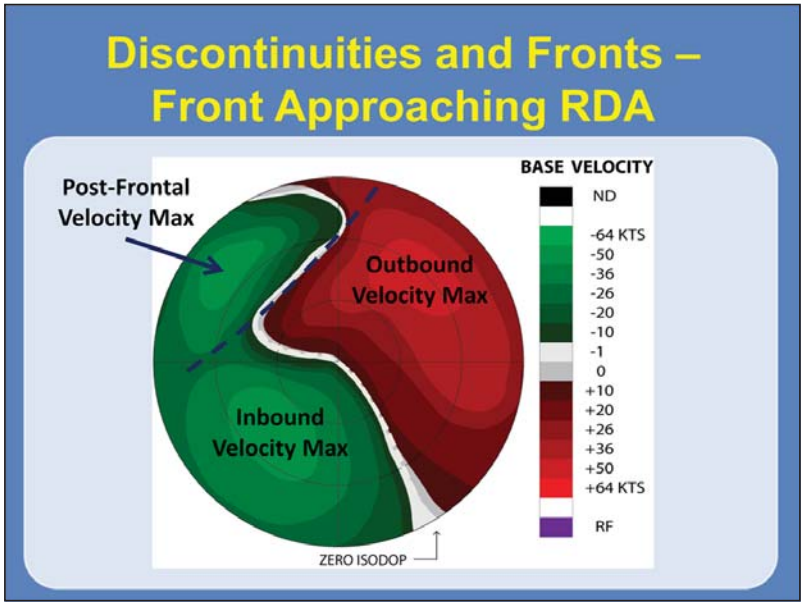
Looking at a diagram at what a confluent wind flow would be, you can see that the outbound velocities are now on the inside of the bow-shaped zero isodop. Here, the winds are coming together over the RDA.



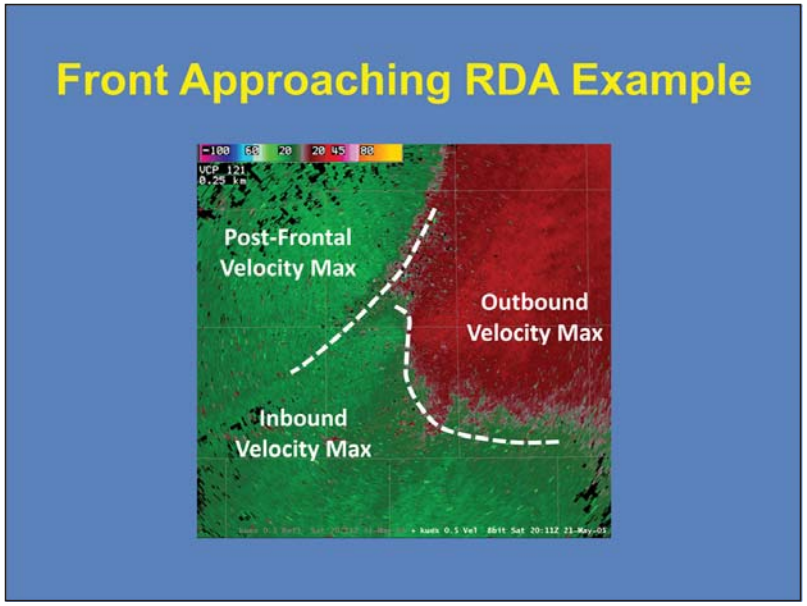
Here is a real-life example of a confluence zone as seen by the WSR-88D. A lake effect snow band is passing over this radar from southwest to northeast. You can see that the zero isodop bends in a way that the outbound velocities are on the inside of the bowing shape.



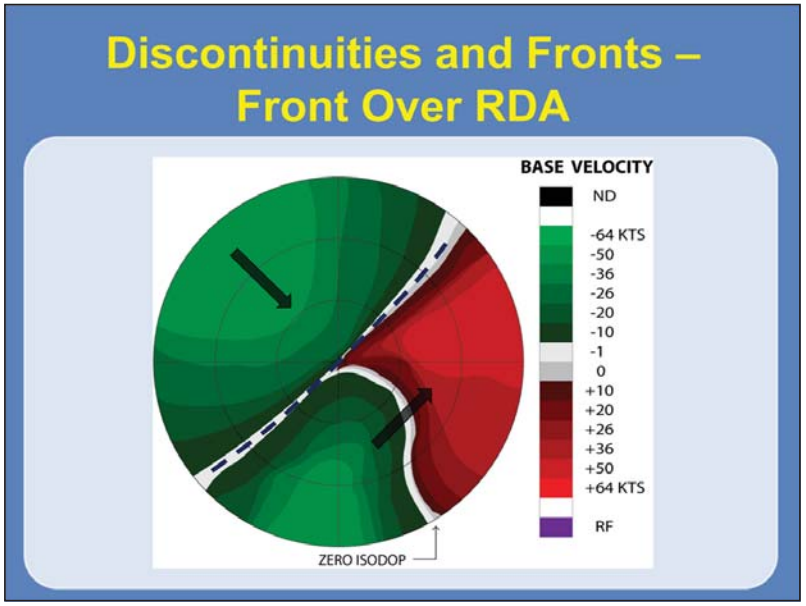
Sometimes, you might run into a situation where you have a sloping wind maximum, like a low-level jet moving over a warm front. In this graphic here, the general flow at all levels is from the southwest to the northeast. If you were to look at the location of the velocity maximum for the inbound winds, you can see it is just beyond the first range ring. Now looking at the velocity maximum on the outbound side, you can see it is over the second range ring. This indicates that the wind maximum is increasing with height as it moves across the display.



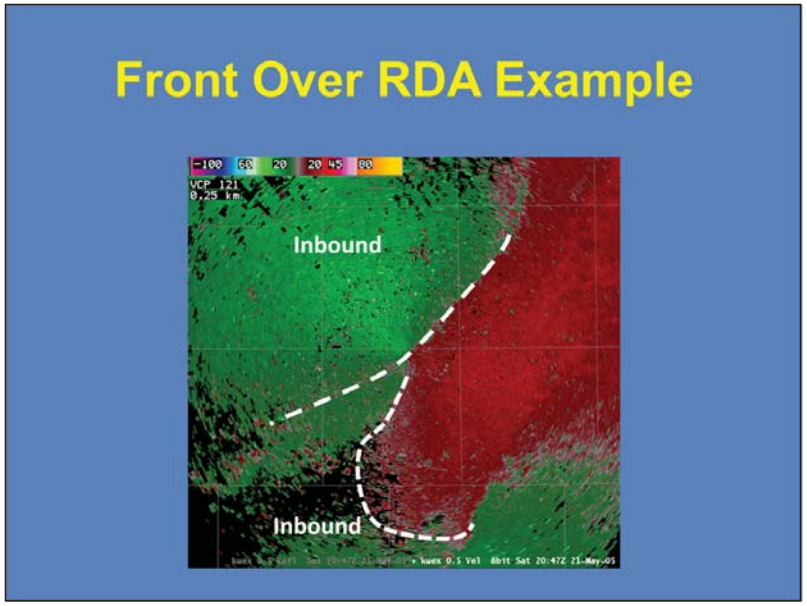
Now lets move on to an even more complicated wind pattern. The following examples will show a frontal boundary moving through the display area, and the expected wind patterns at various stages of the passage. First, we will start with the front approaching the RDA, in this case, from the northwest. Here, you can see that the southeastern 2/3 of the display has an “S” shaped pattern, with velocity maxima located to the northeast and southwest of the RDA. Behind the front, located here, you can see a secondary inbound wind maximum to the northwest, which is not “connected” with any of the other two maxima on this display.



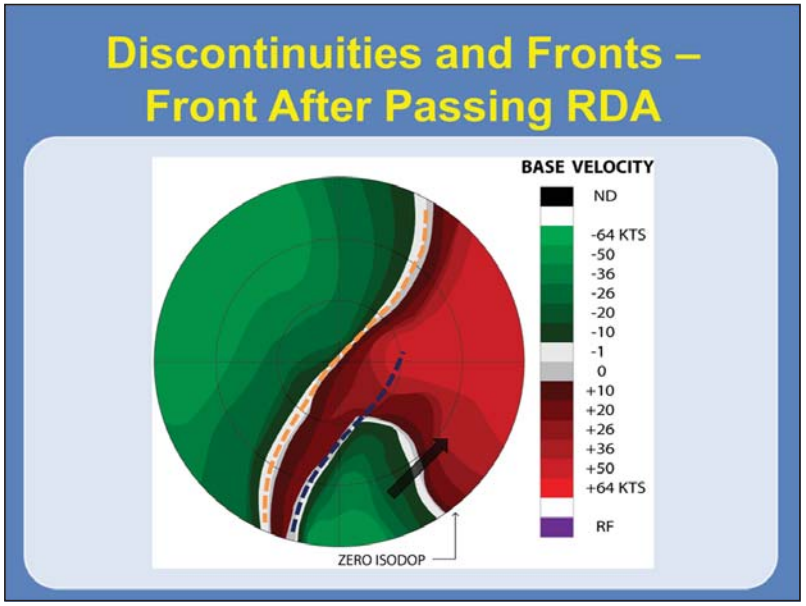
This real world example also has a front approaching the RDA, this time from the north-northwest. In the real world, it is not always easy to see frontal boundaries, especially in just one volume scan. Using the static image here, you can see the where the front is located via the cutoff between the inbounds and outbounds to the north and the velocity minimum to the west. A backing wind pattern is seen south of the boundary. The two connected wind maxima are located to the west-southwest and east-northeast of the RDA, while the disconnected post-frontal velocity maximum is located to the north-nortwest of the RDA. Note that this boundary is aloft in this display. That is, the radar beam is sampling the elevated portion of the boundary.



Moving forward in time, here is a graphical representation of a frontal boundary now located over the RDA. The boundary is still oriented from southwest to northeast, as seen by the zero isodop. The winds are from the northwest behind the boundary, while they are generally from the southwest in a veering pattern ahead of the front.

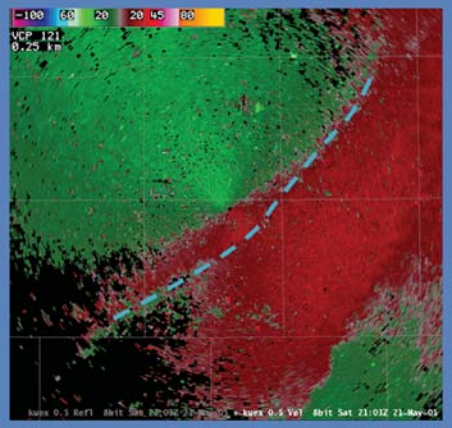


Now moving forward in time with the real-world example, the frontal boundary is now at the RDA. The inbound velocities are located to the south and west of the RDA ahead of the boundary and to the northwest of the RDA behind the boundary. A sharp change in speeds indicates the location of the boundary. Also note the backing wind profile ahead of the front.



The front has now passed over the RDA and is now located to its southeast. The winds ahead of the front (southeastern 1/3 of the display) are still from the southwest. The front is approximately located where the blue line is drawn. Behind the front, the winds are now backing with height.

Front After Passing RDA Example



Back to the real world example, the front is now located to the southeast of the RDA, which is denoted by the blue line here. Notice that this front is losing definition with time. Only to the southwest of the RDA is there a leading edge of inbound velocities. Otherwise, use the leading edge of the enhanced outbounds as the frontal interface.

Summary

- When looking at a display, you are looking into a cone with north at the top of the display.
- The full component of the wind will be measured only when it is parallel to the radial. When the wind is perpendicular to the radial, none of the wind is measured.

Let's review what has been discussed in this lesson. First, when you are looking at a velocity display, you are looking into a cone with north at the top of the velocity display. The full component of the wind will be measured only when it is parallel to the radial. When the wind is perpendicular to the radial, none of the wind is measured.

Summary Continued

- Inbound velocities are negative and are assigned cool colors. Outbound velocities are positive and are assigned warm colors.
- Wind speed at a particular range (height) is determined by the highest Doppler velocity at that range if in a homogeneous flow field.

Inbound velocities are negative and are assigned cool colors. Outbound velocities are positive and are assigned warm colors. Wind speed at a particular range (height) is determined by the highest Doppler velocity at that range if in a homogeneous flow field.

Summary Continued

- A normal “S” shape zero isodop produces a clockwise turning vertical wind profile (veering).
- A backward “S” shape zero isodop produces a counterclockwise turning vertical wind profile (backing).

Continuing on with the summary, a normal “S” shape zero isodop produces a clockwise turning vertical wind profile (veering), while a backward “S” shape zero isodop produces a counterclockwise turning vertical wind profile (backing).

Summary Continued

- A bow-shaped zero isodop with inbound velocities inside the curve represents diffluence.
- A bow-shaped zero isodop with outbound velocities inside the curve represents confluence.

Finally, a “bowed” shape zero isodop with inbound velocities inside the curve represents diffluence, while a “bowed” shape zero isodop with outbound velocities inside the curve represents confluence.



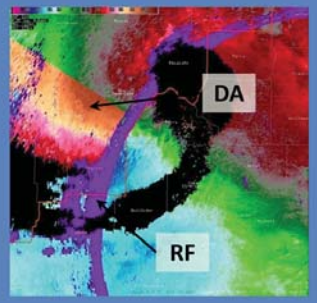
Hi again! It’s Jill Hardy and welcome to this lesson on storm-scale Doppler velocity patterns.

Again, we have a guest speaker: Steve Martinaitis of OU CIMMS at NSSL.

But if you have any questions regarding the material, please feel free to contact me, or the RAC team. Our contact information will be on the next slide.

Review from Previous Topic

- When interpreting velocities products, radial velocities are displayed, which are not the true velocities
- Improperly dealiased velocities and range folding can inhibit velocity interpretation

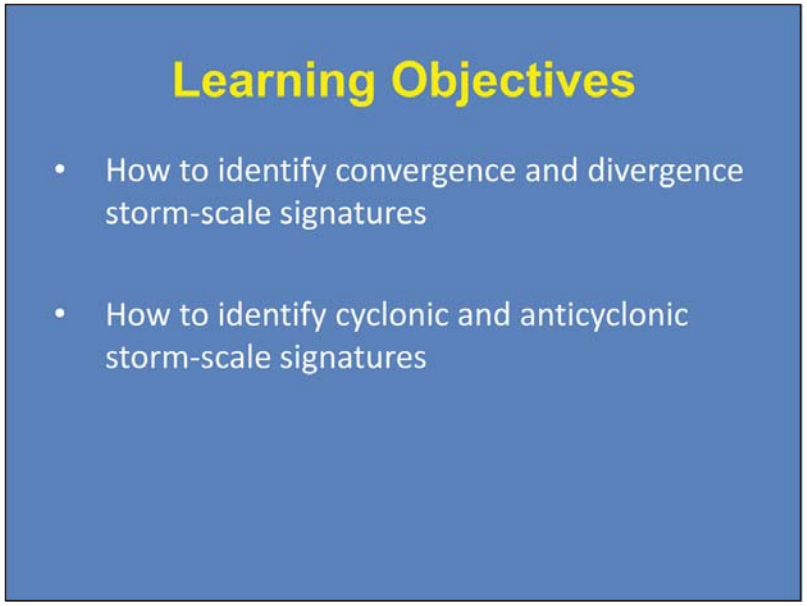


Just like Lesson 1, we will recall that when interpreting velocity data, the radial velocities are displayed, which are not the true velocities. Also, the RPG dealiases velocities while the RDA performs range folding. These tasks are effective most of the time, but failures do occur, which can inhibit your ability to interpret velocity products.

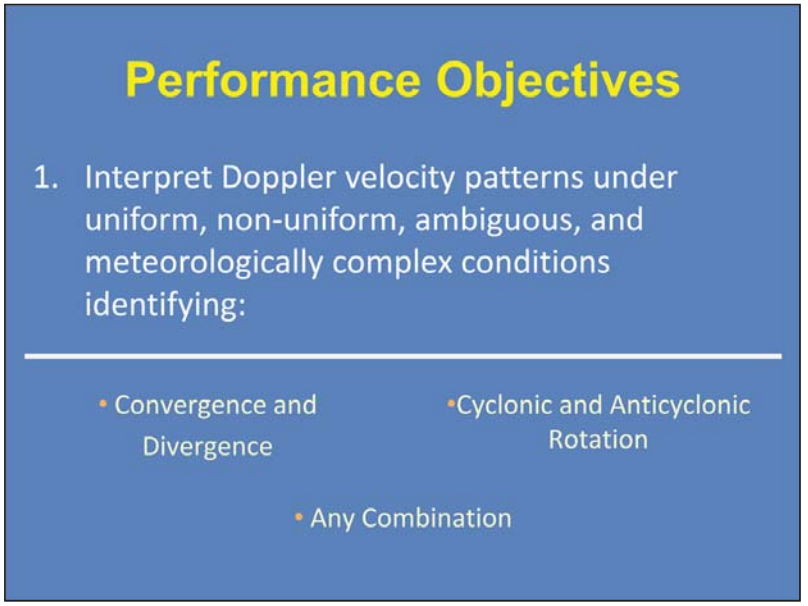
Review from Lesson 1

- When looking at a display, you are looking into a cone with north at the top of the display.
- The full component of the wind will be measured only when it is parallel to the radial. When the wind is perpendicular to the radial, none of the wind is measured.
- Wind speed at a particular range (height) is determined by the highest Doppler velocity at that range if in a homogeneous flow field.

From Lesson 1, recall that when you are looking at a velocity display, you are actually looking into a cone where the farther from the RDA you get, the greater the height above the ground the data is being measured. The full component of the wind will be measured only when it is parallel to the radial. When the wind is perpendicular to the radial, none of the wind is measured. You saw this through the Radial Velocity Equation. Also, wind speed at a particular range (height) is determined by the highest Doppler velocity at that range if in a homogeneous flow field. Of course, you saw how complex this became when the wind field became horizontally non-homogeneous.



There are two learning objectives to this lesson. By the end of this presentation, you should be able to understand how to identify storm-scale convergence and divergence signatures, and understand how to identify cyclonic and anticyclonic storm-scale signatures.



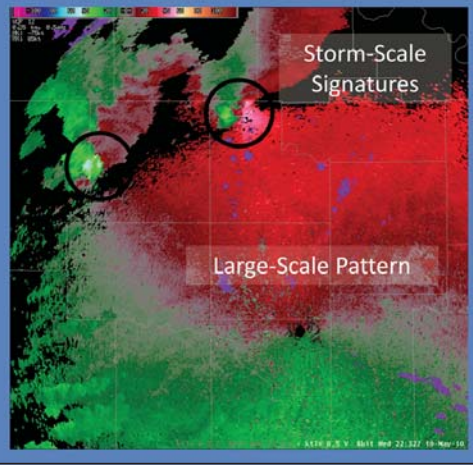
There are also two performance objectives with this lesson. The first objective is to be able to interpret Doppler velocity patterns under uniform, non-uniform, ambiguous, and meteorologically complex conditions. You will be identifying convergence and divergence, cyclonic and anticyclonic rotation, and any combination of the above.

Performance Objectives

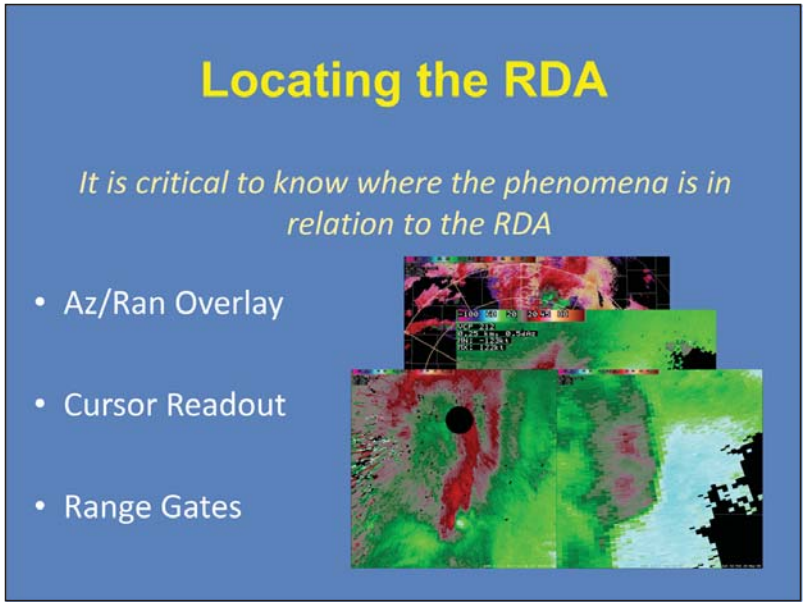
2. Assess the mesoscale meteorological conditions and threats associated with the identified velocity patterns

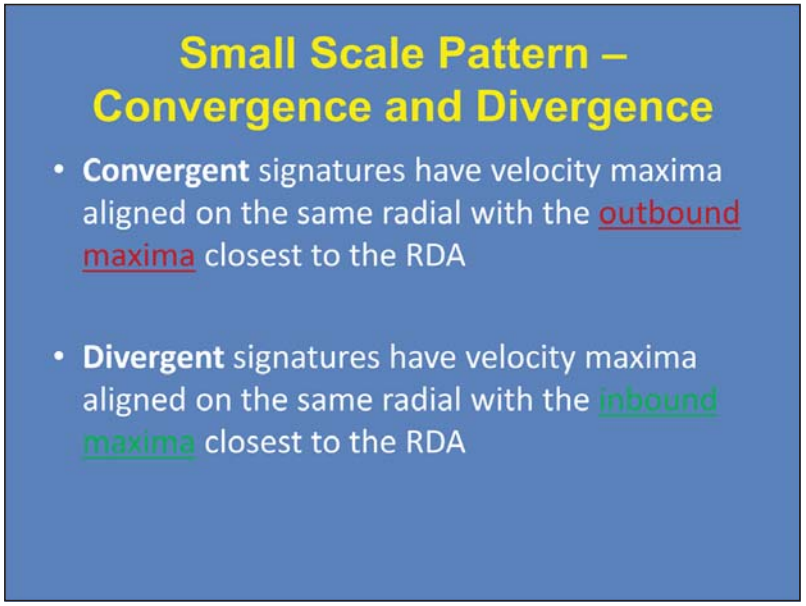
The second objective is to assess the mesoscale meteorological conditions and threats associated with the identified velocity patterns.

Doppler Velocity Patterns



In the first lesson, you saw a variety of factors and examples that influence the large-scale velocity field and its display in the AWIPS environment. Now, we will go ahead and focus on the small-scale phenomena, which cover only a few range gates, and therefore, have a relatively small change in elevation.



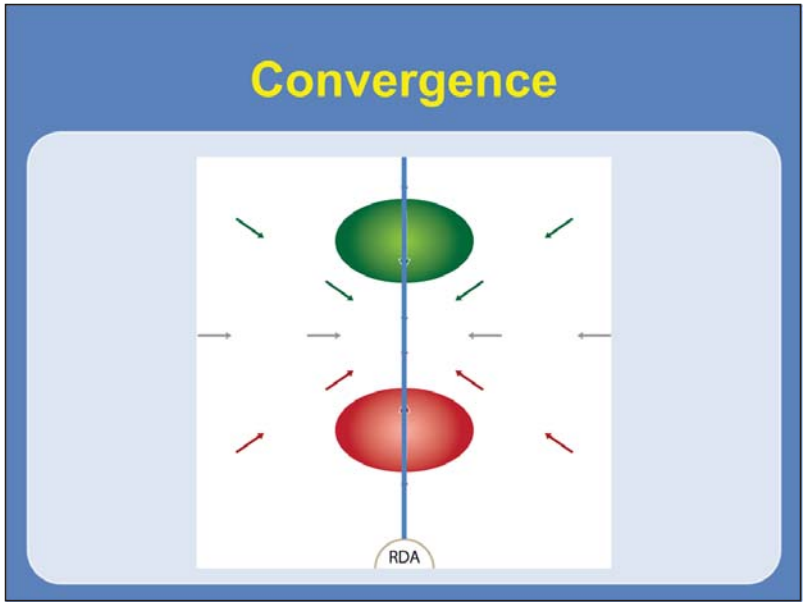


When examining data on this scale, you will be using the Zoom feature in order to see small scale rotation and/or convergence and divergence. It is critical to know where the phenomena is in relation to the RDA. Here, you can no longer assume that the RDA is in the center of the display, or on the display at all. The following three actions, either used separately or in combination, will help in locating the RDA.

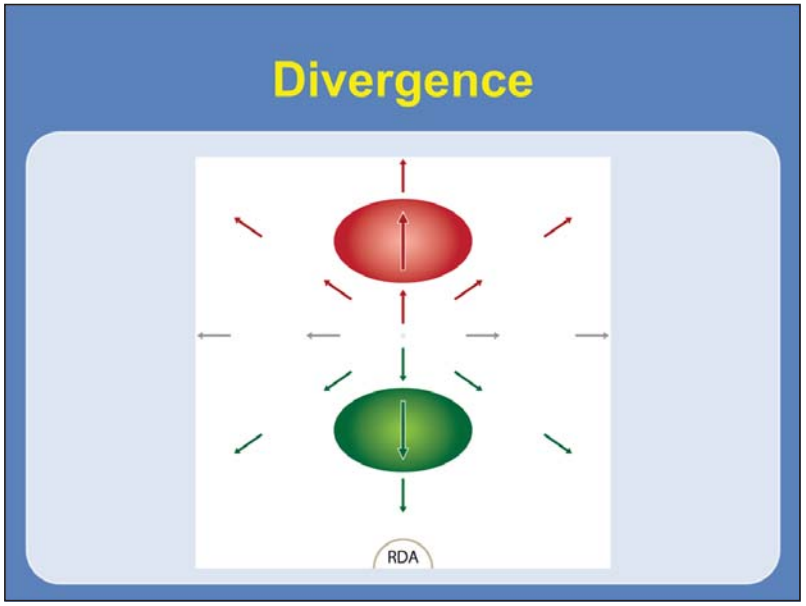
You can select the Azimuth and Range (or Az/Ran) Overlay from the Tools menu to help determine the location of the RDA by overlaying a polar grid centered on the RDA. Or you can place your cursor at the point of interest and hold down the left mouse button. The cursor readout will give the azimuth and range (in statute miles) from the RDA. Make sure that the Home Location tool is not turned on. Or you can visually analyze the range gates in your velocity display. Range gates increase in width along each radial as they increase in distance from the RDA. This is one advantage to an unsmoothed radar display.

When interpreting pure divergence or convergence patterns, the velocity maxima lie along the same radial. Whether the pattern is divergent or convergent is dependent upon which maximum is closest to the RDA. With a convergent signature, the outbound maxima is closest to the RDA. With a divergent signature, the inbound maxima is closest to the RDA.

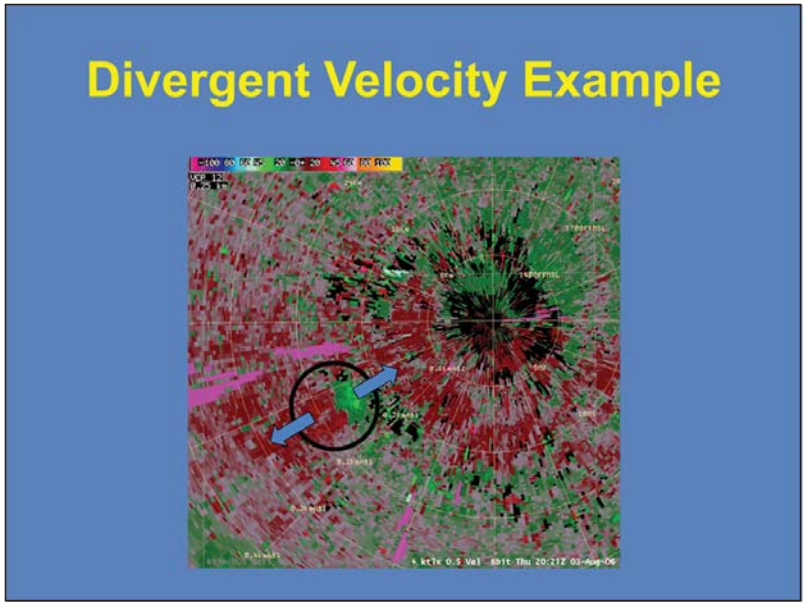
Note that in the following examples, the RDA is located to the south of the velocity signature.



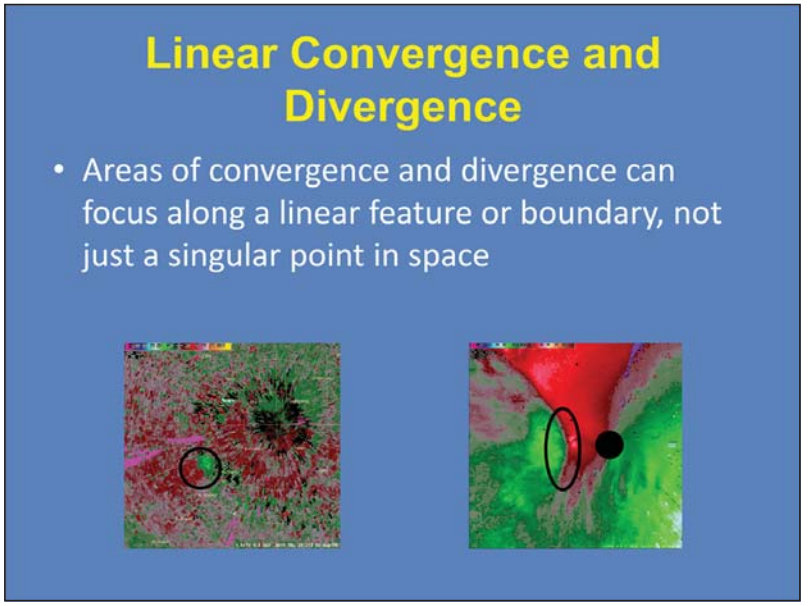
Here is a basic diagram of a convergent velocity signature. As you can see here, both the maxima lie along the same radial with the outbound velocity maximum closest to the RDA.



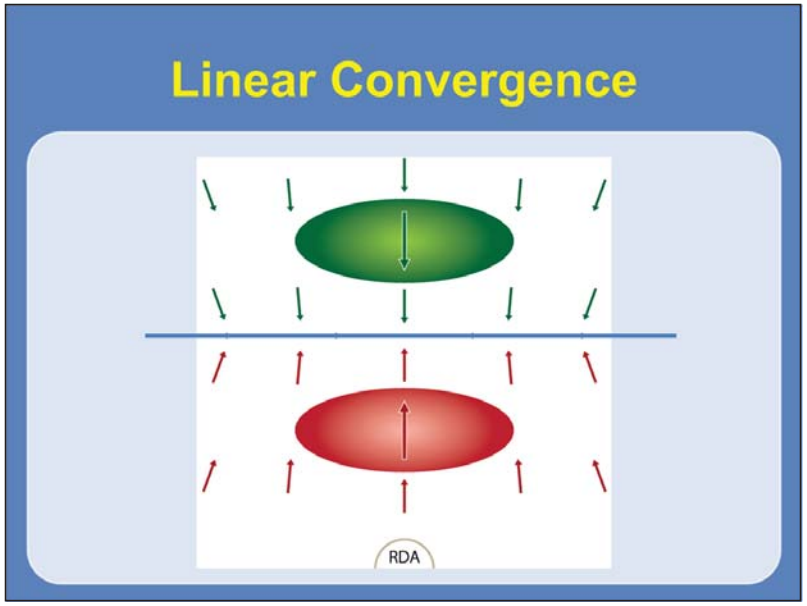
Same thing here for the divergent velocity signature, except that the inbound velocity maximum is now closest to the RDA.



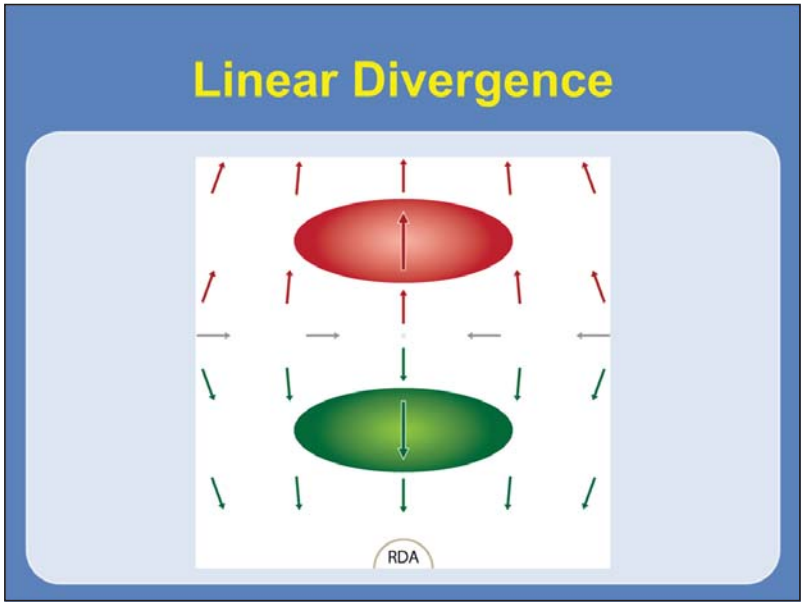
In this real-world example, there is a divergence signature located to the west-southwest of the RDA with a circular outflow pattern. This occurred just after a downburst from a thunderstorm. Note that the maximum inbound velocity is closer to the RDA than the maximum outbound velocity.



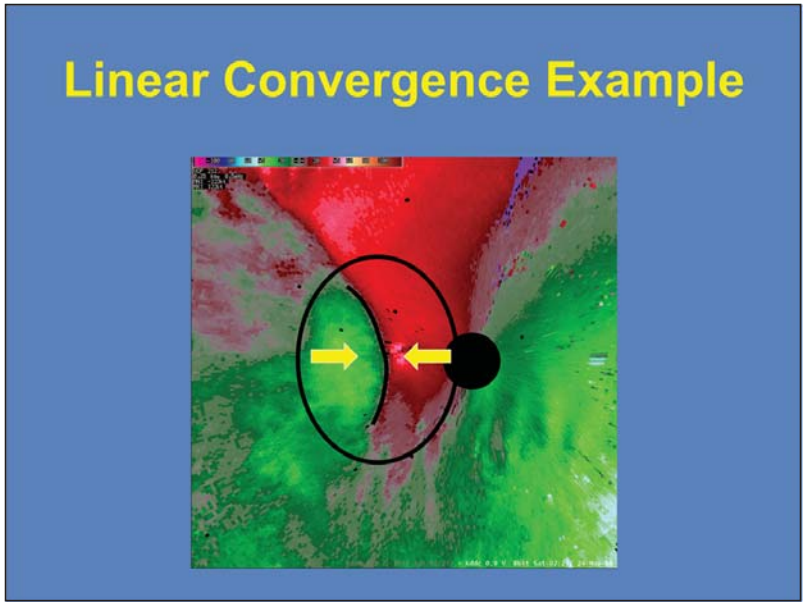
The examples shown in the last few slides were of conceptual models of pure convergence and divergence on a single point in space. This was also shown in the real-life downburst signature you just saw. However, areas of convergence and divergence can also focus along a linear feature.



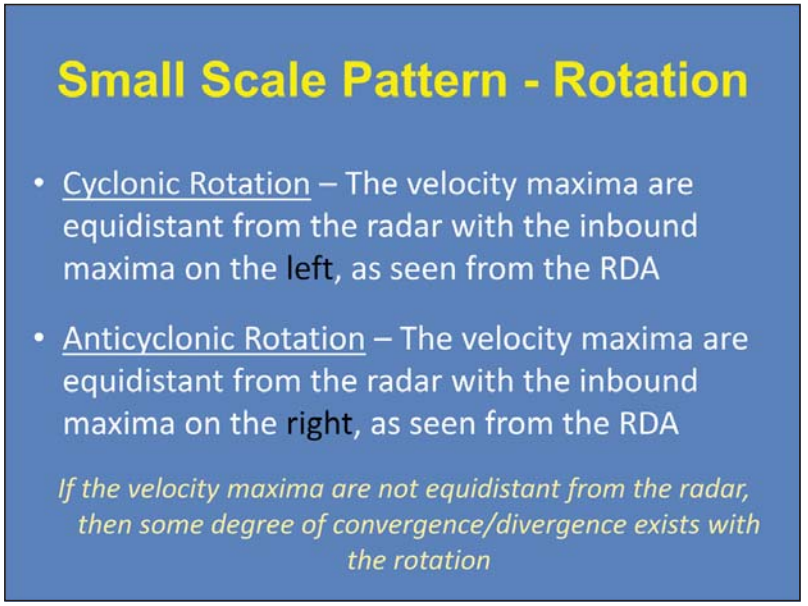
Here is a basic diagram of a convergent velocity signature focused along a linear feature. As you can see here, both the maxima still lie along the same radial with the outbound velocity maximum closest to the RDA. The exception is that these maxima are elongated across a number of radials at about the same range from the RDA.



Same thing here for the divergent velocity signature. Again, both maxima are elongated across a number of radials, but now the inbound maximum is closest to the RDA.

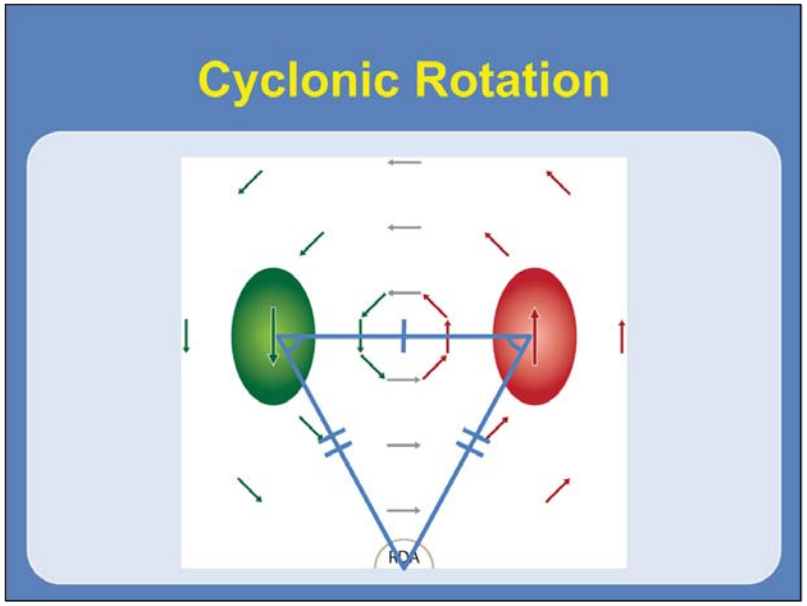


Here is a real-world example of a linear storm-scale convergence signature. In this case, a QLCS is approaching the RDA from the west, and a segment of the line is bowing out at this point, creating an enhanced convergence signature. Areas of enhanced convergence along a line segment could lead to mesovortex formation, which can enhance the wind threat and increase the probabilities of a tornado in areas of vorticity that are generated.

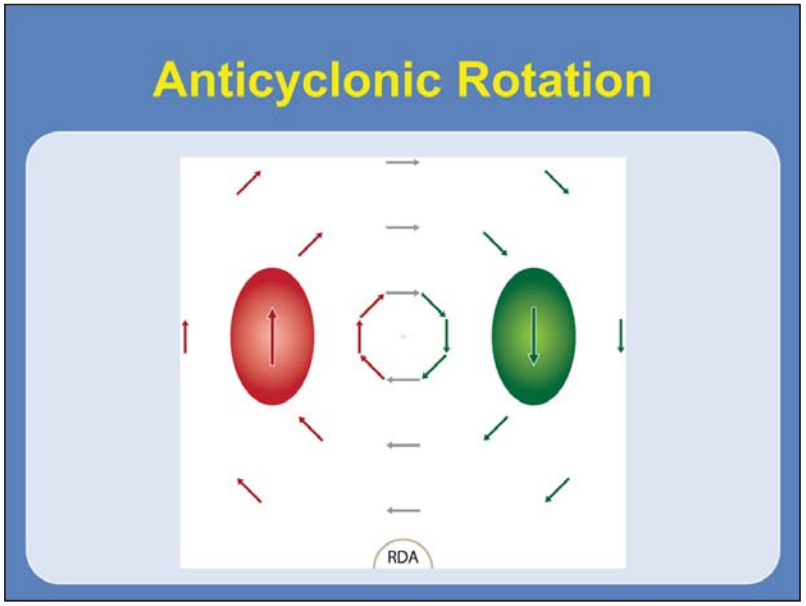


When examining pure rotational patterns, the velocity maxima are equidistant from the radar. Whether the pattern is cyclonic or anticyclonic is dependent upon whether the inbound maximum is on the left side or the right side of the signature, as seen by the RDA. With cyclonic rotation, the inbound maximum is on the left hand side, while the inbound maximum is on the right hand side with anticyclonic rotation. Velocity maxima oriented any other way means some combination of rotation and convergence or divergence is occurring.

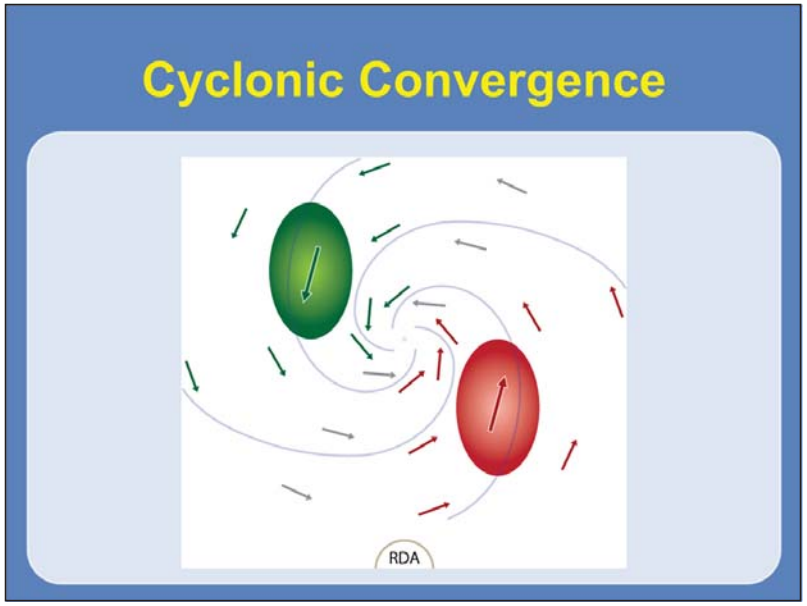
Again, for the following examples, the RDA is located to the south of the velocity signature.



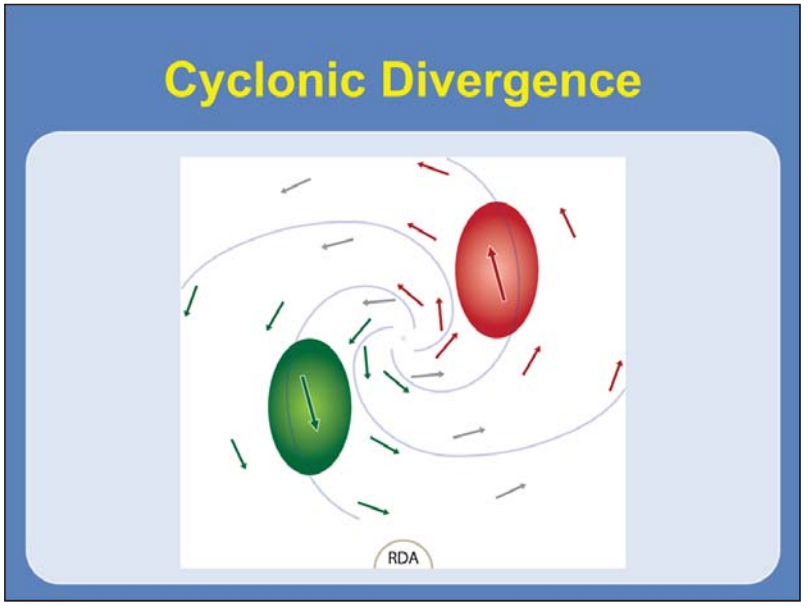
Here is a basic example of a pure cyclonic rotational signature. As you can see here, both of the velocity maxima are equidistant from the RDA with the maximum inbound velocities on the left side of the signature.



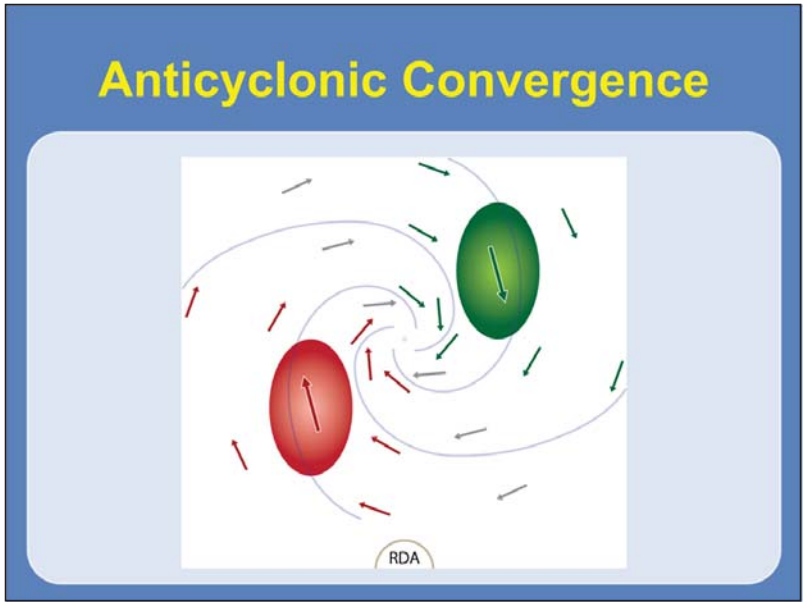
Now with the pure anticyclonic rotational signature, both velocity maxima are again equidistant from the radar with the inbound velocity maximum on the right side of the signature.



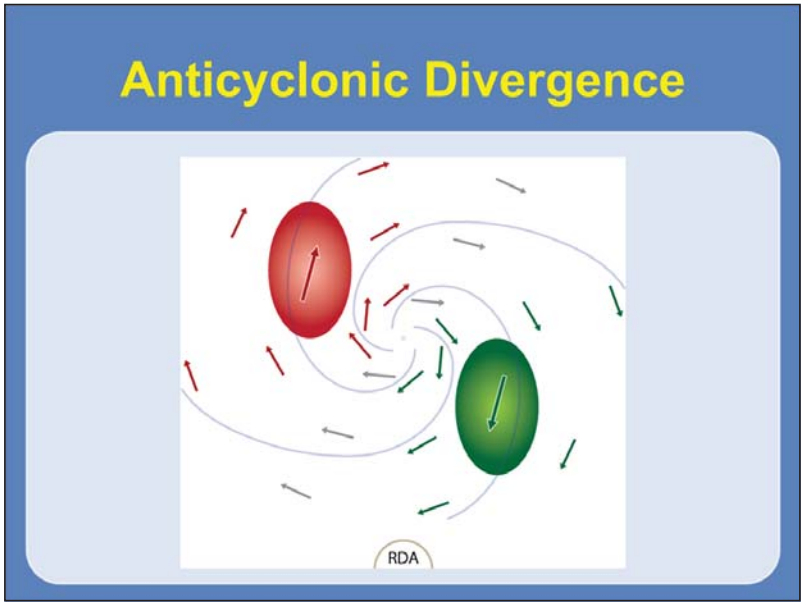
Now to make things more complicated, we will show combinations of both rotation and convergence or divergence. This is an example of cyclonic convergence. First note that both the maxima are not on the same radial and not equidistant from the radar. Here, the outbound maximum is closest to the RDA, signifying convergence, and the inbound maximum is to the left, signifying cyclonic rotation.



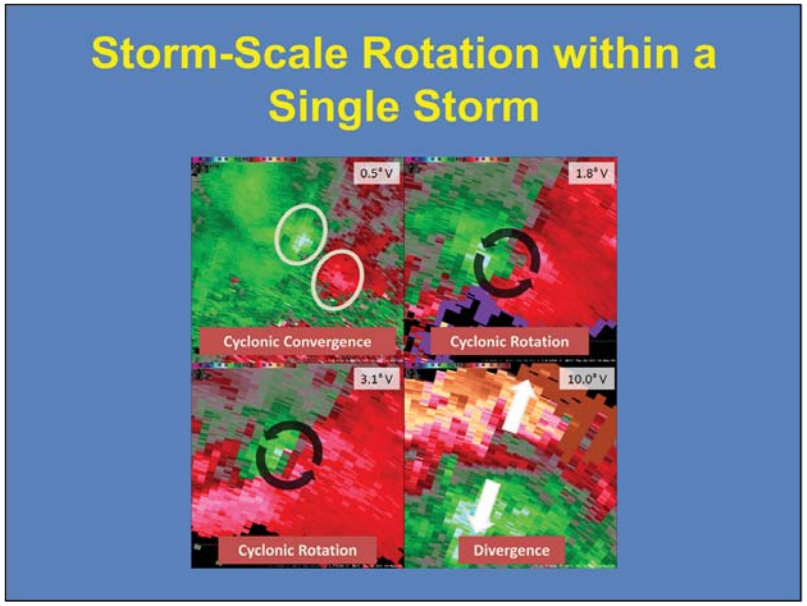
Here is an example of cyclonic divergence. The inbound maximum is closest to the RDA, signifying divergence, and the inbound maximum is to the left, signifying cyclonic rotation.



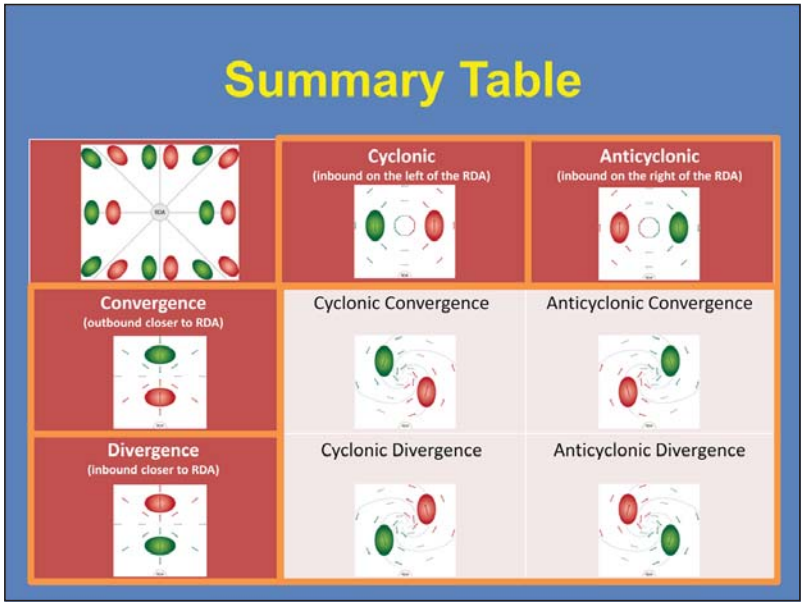
Here is an example of anticyclonic convergence. The outbound maximum is closest to the RDA, signifying convergence, and the inbound maximum is to the right, signifying anticyclonic rotation.



Finally, here is an example of anticyclonic divergence. The inbound maximum is closest to the RDA, signifying divergence, and the inbound maximum is to the right, signifying anticyclonic rotation.



Here is a real-world example of storm-scale signatures through various tilts of what would become a tornadic supercell. Note that the RDA is located to the south-southwest of the storm. In the upper-left panel (0.5° tilt), you see a cyclonic convergence signature with the storm. The next two elevation scans shown in this example (1.8° tilt in the upper-right panel and 3.1° tilt in the lower-left panel) are close to “pure” cyclonic rotation. The 10.0° tilt (lower-right panel) is an example of storm-top divergence.



Let’s go ahead and summarize storm-scale velocity signatures. Convergence signatures have the velocity maxima lie along the same radial with the outbound maximum closest to the radar. Divergence signatures have the velocity maxima lie along the same radial with the inbound maximum closest to the radar.

Cyclonic rotation signatures have the velocity maxima equidistant from the radar with the inbound maximum to the left, as seen from the radar. Anticyclonic rotation signatures have the velocity maxima equidistant from the radar with the inbound maximum to the right, as seen from the radar.

Velocity maxima oriented any other way means some combination of rotation and convergence or divergence is occurring.

TAB

Radar & Applications Course

Topic:

Base and Derived Products

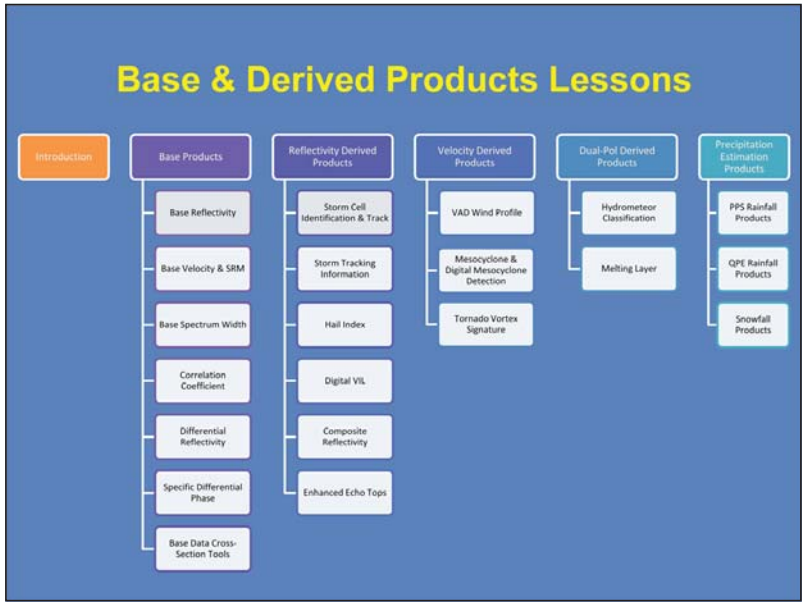
Table of Contents
Topic: Base and Derived Products

Click to jump to lesson

Lesson 1	**Introduction to Base and Derived Products**
Lesson 2	**Base Reflectivity (Z)**
Lesson 3	**Base Velocity (V) & Storm-Relative Velocity Mean Radial Map (SRM)**
Lesson 4	**Base Spectrum Width (SW)**
Lesson 5	**Correlation Coefficient (CC)**
Lesson 6	**Differential Reflectivity (ZDR)**
Lesson 7	**Specific Differential Phase (KDP)**
Lesson 8	**Base Data Cross-Section Tools (FSI, Volume Browser, & GR2Analyst)**
Lesson 9	**Storm Cell Identification and Tracking (SCIT) Algorithm**
Lesson 10	**Storm Tracking Information (STI)**
Lesson 11	**Hail Index (HI)**
Lesson 12	**Digital Vertically Integrated Liquid (DVIL)**
Lesson 13	**Composite Reflectivity (CZ)**
Lesson 14	**Enhanced Echo Tops (EET)**
Lesson 15	**Velocity Azimuth Display Wind Profile (VWP)**
Lesson 16	**Mesocyclone (MD) & Digital Mesocyclone (DMD)**
Lesson 17	**Tornadic Vortex Signature (TVS)**
Lesson 18	**Hydrometeor Classification (HC)**
Lesson 19	**Melting Layer (ML)**
Lesson 20	**PPS Rainfall Products**
Lesson 21	**QPE Rainfall Products**
Lesson 22	**Snowfall Accumulation (PPS)**



Welcome to the Radar and Applications Course topic on Base and Derived Products. This lesson introduces this topic where you will learn about the most operationally relevant products produced by the WSR-88D. You'll learn when, why, and how to use them (aka: their applications), as well as when to question or even ignore their output (aka: their limitations). This lesson will take approximately 15 minutes to complete.



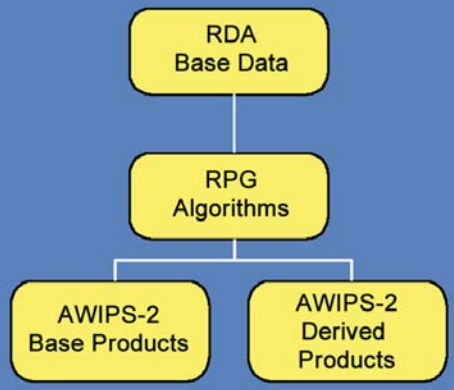
Here is a roadmap for the lessons in this topic. This lesson is the Introduction, shown shaded in orange, that provides an overview of the entire topic. As you take each lesson in this topic, we will highlight "where you are" with respect to the topic as a whole.

Learning Objectives

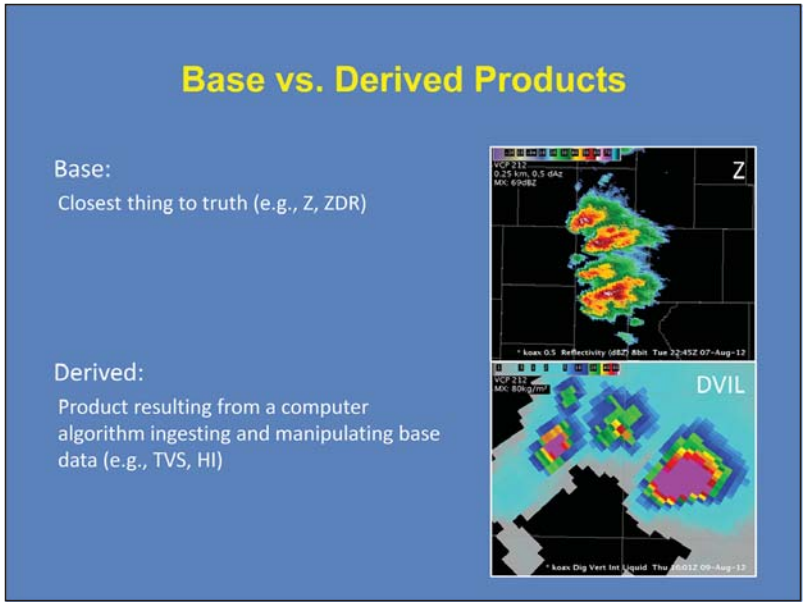
1. Identify the general applications and limitations of base and products
2. Identify the general applications and limitations of derived products
3. Identify the technical variables that impact your ability to view base and derived products in AWIPS, their definitions, and the best practices mentioned related to each variable

These are the learning objectives for this lesson. Please take a moment to review the objectives and click on the next button when you are ready to advance to the next slide.

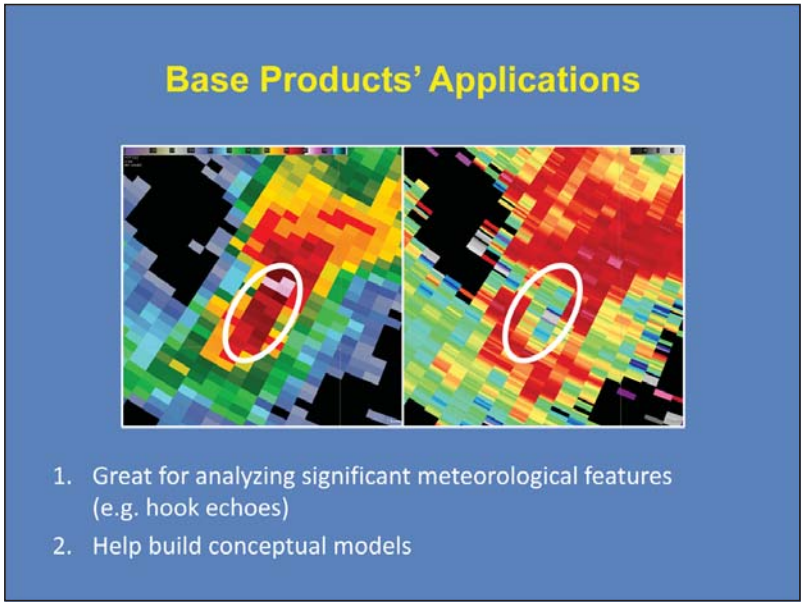
Review of WSR-88D Data Flow



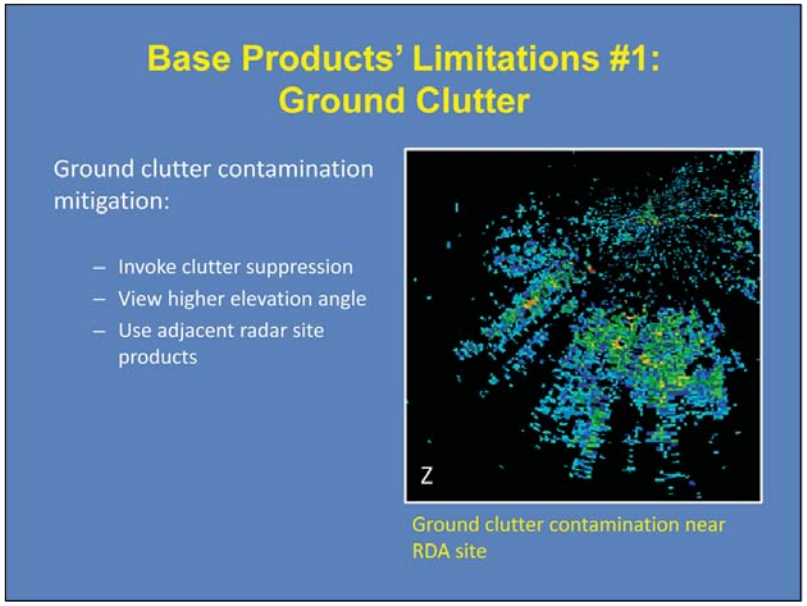
Recall that base data are collected by the RDA and sent to the RPG to be processed by the meteorological algorithms residing at the RPG. Ultimately the RPG will then produce and distribute both base and derived products to the user workstations.



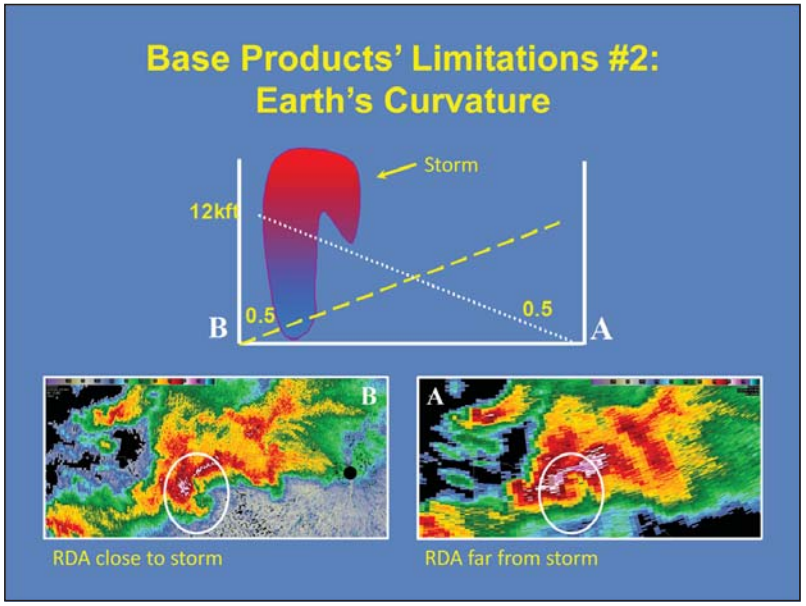
So, what are base and derived products? Base products are basically the closest thing you can get to the truth. It's Base Reflectivity, Differential Reflectivity, and similar products that visualize the base data streams. Derived products are products generated from a computer algorithm ingesting and manipulating the base data. Examples of derived products include the Tornado Vortex Signature or the Hail Index products.



In general, base products are great for recognizing and analyzing significant meteorological features such as hook echoes, hail cores, etc. Base products also help you develop and apply conceptual models that make storm analysis and warning issuance quicker and more efficient.



The next few slides will cover general limitations of all base products. The first limitation is ground clutter. Ground clutter is usually an issue near the radar where tall objects or beam bending can cause the radar returns from objects on the ground. Ground clutter often appears speckled in nature. To help mitigate ground clutter contamination, you can invoke clutter suppression at the RPG, view a higher elevation, or view data from a nearby radar.



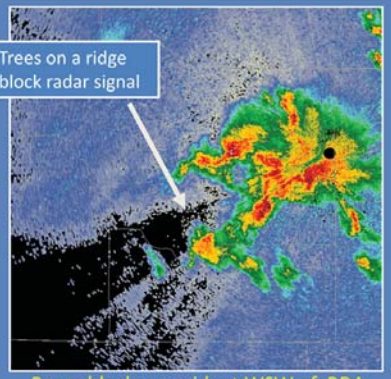
As you know…the earth is round! While this is hardly a shocking revelation, it does play an important role in our next limitation.

The radar beams propagate away from the radar at increasing altitudes with range due to the earth's curvature. In other words, the beam samples higher and higher into the atmosphere the further away the transmitted signal travels. When the radar beam is 10s of kilometers away, it may be sampling the midlevels of a storm. In the example provided, radar B is very near the storm and notice how you can see a hook echo, an inflow notch, and other details. The same storm viewed from radar A, which is farther away, appears different because it is sampling a different part of the storm, among other reasons. This is a classic example of the earth's curvature limitation.

Base Products' Limitations #3: Beam Blockage

Beam blockage mitigation:

- Choose higher slice
- Query other radars
- Ground truth from spotters



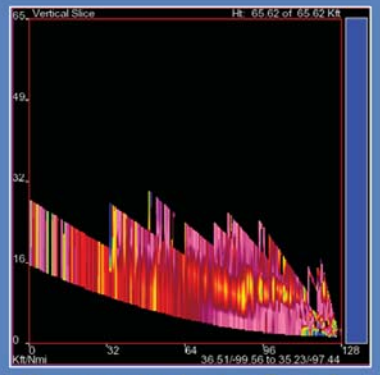
Beam blockage evident WSW of RDA site

Beam blockage occurs when tall objects near the radar completely, or partially, absorb (or reflect) the radiation from the radar pulse. Since much of the transmitted signal is blocked, targets down radial from such objects cannot be observed properly by the radar. Beam blockage manifests itself as a wedge of low-value, or missing, returns. Here is an example from the KVNX radar in NW Oklahoma. When you see beam blockage in radar data, you can mitigate these effects by either choosing a higher slice, looking at a nearby radar, or obtaining ground truth from spotters in those regions.

Base Products' Limitations #4: Beam Broadening

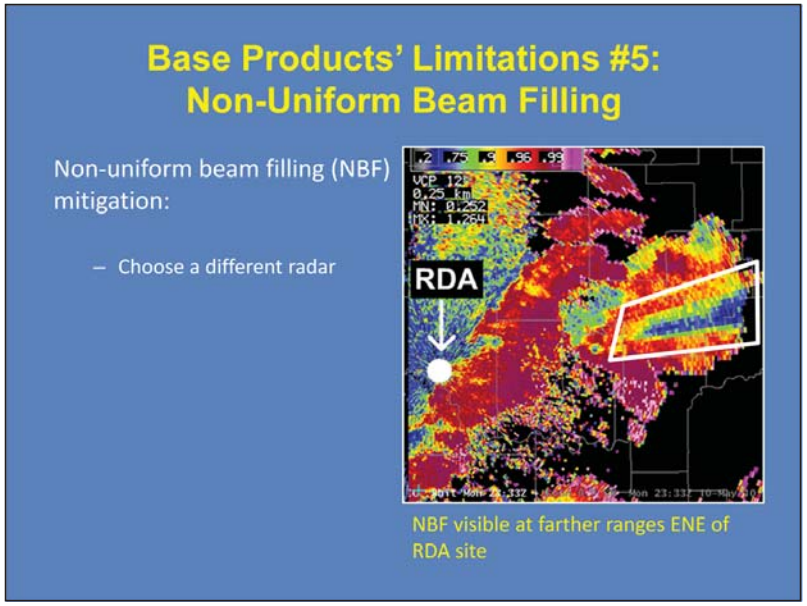
Beam broadening mitigation:

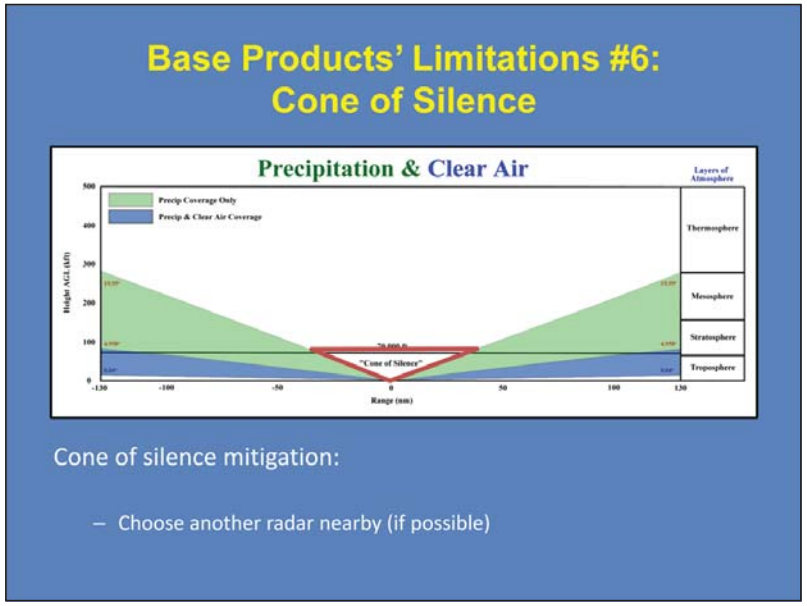
- Choose a closer radar



FSI cross-section showing effects of beam broadening at far range

As the radar beam propagates away from the radar, it broadens. In fact at 60 km from the radar, the beam is roughly 1 km wide! This broadening effectively smears all radar signatures. For example, the melting layer, shown by the yellowish/orange colors in the image, becomes broader the farther away from the radar it is sampled. To reduce beam broadening effects, you must view data from a closer radar, if one is available.

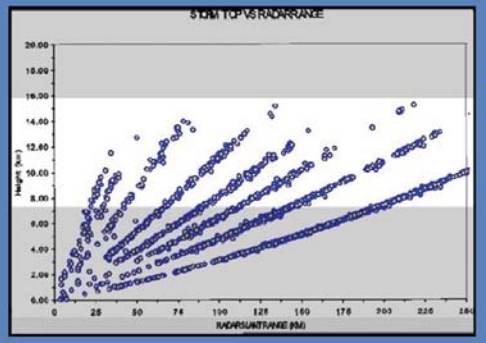




Non-Uniform beam filling (or NBF) has been around for as long as radar has, but has gained visibility recently with the WSR-88D upgrade to dual-polarization. While there are many degrees of non-uniform beam filling, it is only operationally significant when storms line up along a radial or the radar samples a significant hail core in a storm around the height of the environmental melting level. An example of NBF is noted in the image on the right by the radials of reduced Correlation Coefficient (CC; Ryzhkov and Zrnic, 1998; Gosset 2004). When NBF impacts your storm analysis, you need to use data from a different radar to mitigate the issue.

Recall that the highest elevation scanned by the WSR-88D is 19.5 degrees. So, this means there is a lot of the atmosphere that is not sampled! The unsampled region above the highest tilt is called the cone of silence. The cone of silence impacts your analysis only when storms are located close to the radar. If you experience cone of silence issues, choose a different radar where the storm of interest is not in its cone of silence.

Base Products' Limitations #7: Beam Height Estimations

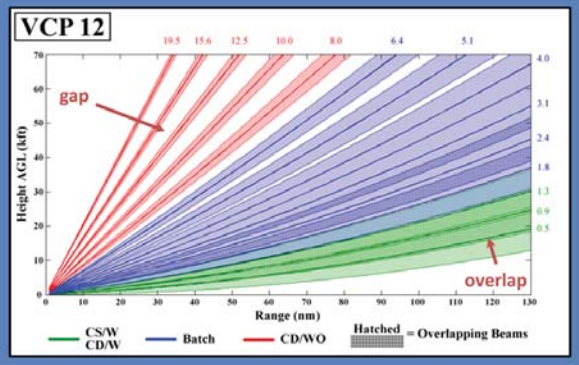


Plot of estimated storm top heights as a function of range

Remember that displayed beam heights are only estimates!

As the antenna rotates, it doesn't stay exactly on the elevation defined. For example, when we say the radar is at 0.5 degrees, it could be 0.4 or 0.6 degrees and still be in spec. However, a slight difference here can mean the height estimate of the beam may have significant errors at far ranges. Likewise, changes in the atmospheric temperature structure can lead to errors in beam height estimates. AWIPS-2 assumes a standard atmosphere for beam height estimates, which is usually not true. So, always keep in mind that the beam heights listed are only "estimates" and in actuality they could be as much as a few thousand feet off.

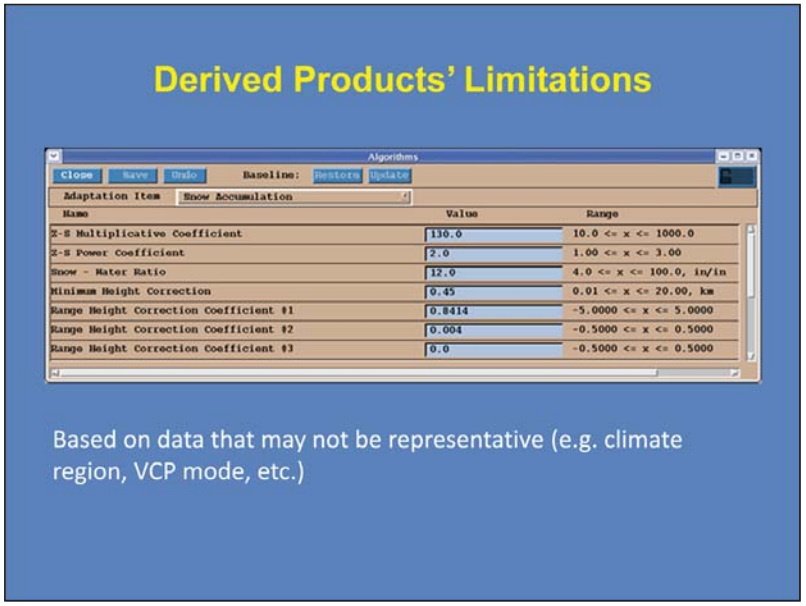
Base Products' Limitations #8: Discrete Sampling



Gaps between elevation angles affect forecaster interpretation & introduce algorithm artifacts

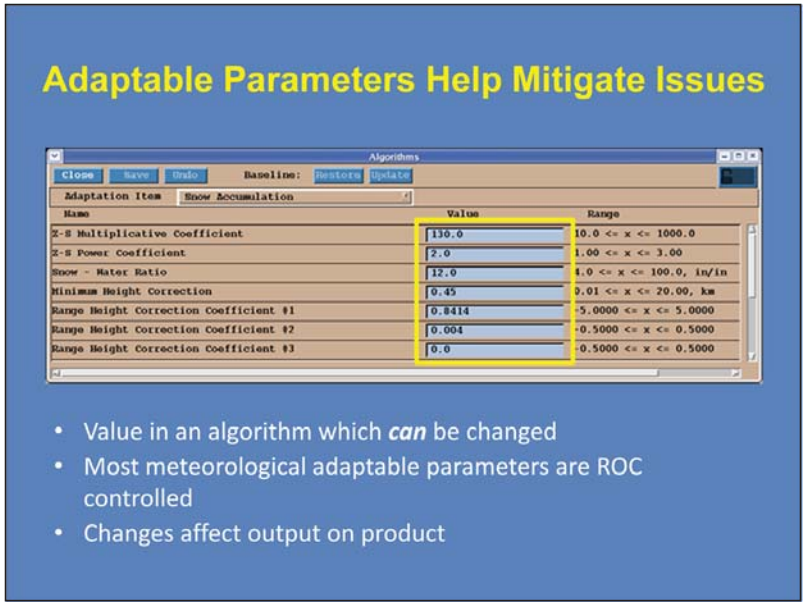
In the lowest elevations, there are no gaps between elevation slices. In fact, you can see from the figure shown here that there is even some overlap. On the other hand, the upper portions of each volume scan often have substantial gaps, especially above 8 degrees. These gaps affect forecaster interpretation as well as introduce artifacts in algorithm output.



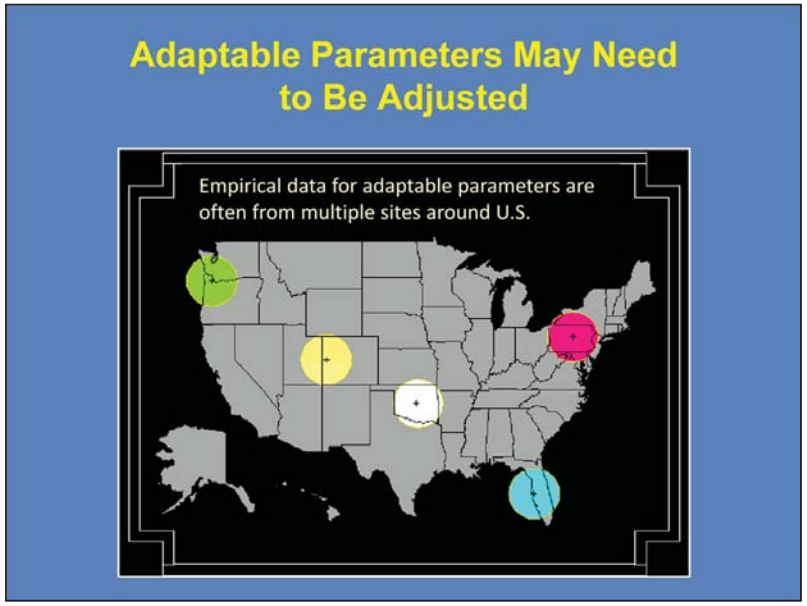


Now that we have covered base products, lets talk about the applications and limitations of derived products. The primary advantage to using derived products is they speed up analysis of large amounts of data. However, you should only use derived products as a safety net. Basically, if you see something in a derived product that grabs your attention, then you should verify that signature with your knowledge of the base data before making a warning decision. One thing you should never do is issue a warning based solely on derived products!

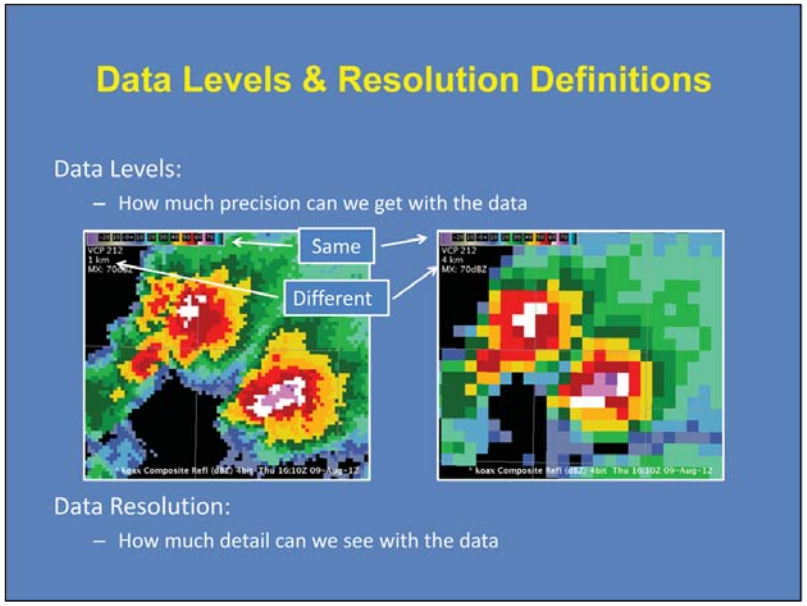
Recall that derived products result from a computer algorithm ingesting base data and analyzing it for significant features. Well, not every location experiences the same types of weather and the algorithms were developed using data from certain locations. So, the data used to develop the algorithm may not be representative of all areas.



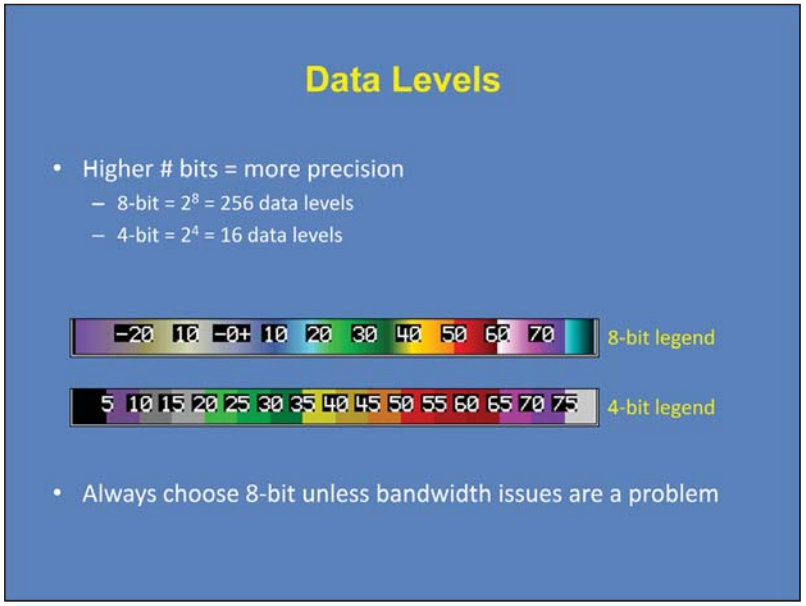
Fortunately, adaptable parameters help address problems related to RPG algorithms! Many products use adaptable parameters that allow for modification of algorithm behavior. Most of these parameters are password protected by the Radar Operations Center, but a few can be changed by the local office. Any changes made to the adaptable parameters will impact the look and feel of the product for ALL users. An example adaptable parameter window is shown here for the Snow Accumulation Algorithm.



Most algorithms were optimized empirically using a series of datasets. However, to perform best in your area at different times of the year, the adaptable parameters for a particular algorithm may need to be modified. This process takes both research and coordination with the Radar Operations Center.

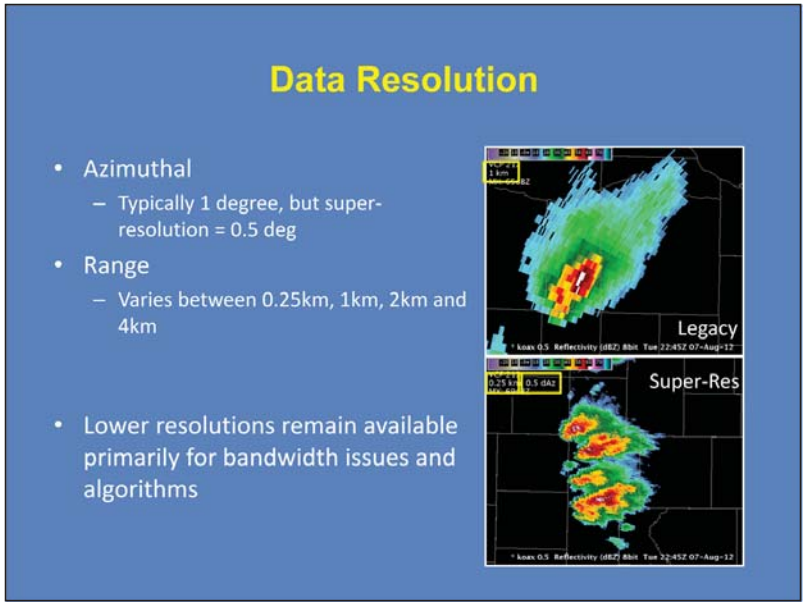


Two important definitions to remember here are data levels and resolutions. Data levels tell you how much precision we can get with the data. The higher the data levels, the more precision that is possible. Data resolution tells us how much detail can we see with the data. The higher the data resolution, the more detail we can see. The next few slides will go into a little more depth of these definitions.

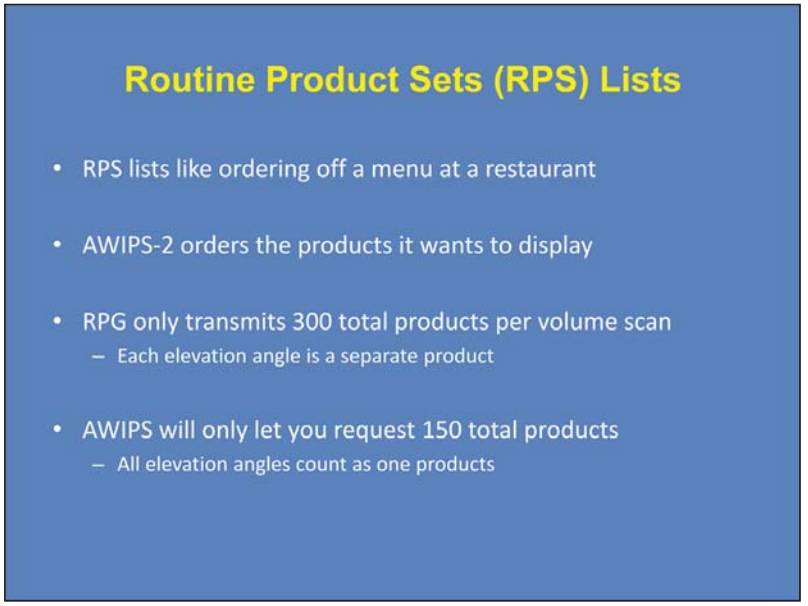


Data levels are characterized by how many bits are available to record them. The more bits available, the more precise the data are. For example, 8-bit products have 256 data levels while 4-bit products only have 16 data levels. To figure out exactly how many data levels are available, just raise 2 to the power of the product bit number.

The color bars at the bottom of this slide illustrate the difference very nicely. The top color bar is for 8-bit Base Reflectivity, while the bottom legend is for 4-bit Base Reflectivity. Notice how for 4-bit, a single color represents roughly a 5 dBZ range, whereas for 8-bit products each color represents a much smaller range. So with 8-bit we can delineate smaller increments in reflectivity changes than 4-bit can. Operationally, higher bit products are always preferable unless you are experiencing significant connectivity issues in your office.

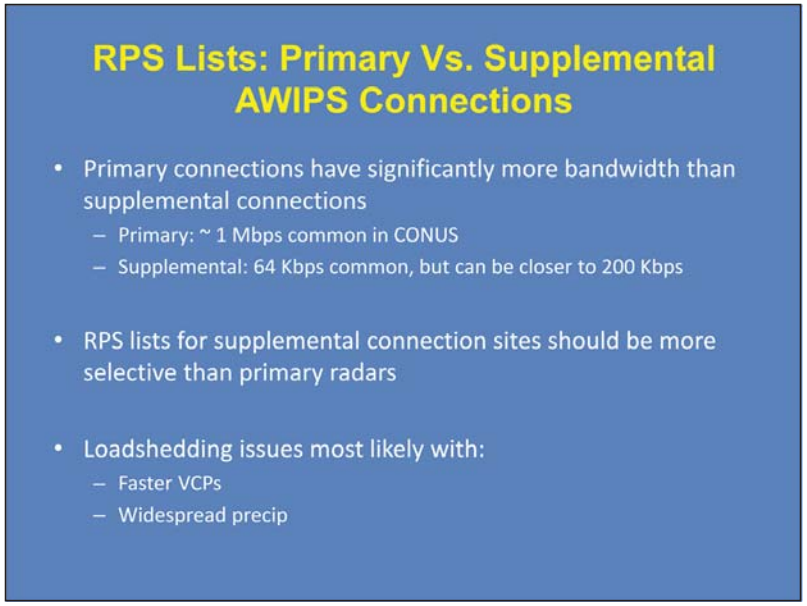


Radar data resolution can be broken out into two categories: azimuthal and range resolution. For WSR-88D products, most products will have an azimuthal resolution of 1 degree, except for super-resolution products which will be 0.5 degree azimuthal resolution. Most WSR-88D products have a range resolution of 0.25 km, but coarser resolution products are still available (i.e. 1, 2 and 4 km). It is highly recommended to use the highest resolution data available unless you experience performance issues with your workstation or your office experiences significant connectivity issues.



The last item to cover is the Routine Product Sets (or RPS) list. The RPS list is similar to an order off of a menu. Just like we can't eat everything off the menu (well at least I hope you can't), your local AWIPS can't download all products that are generated from your dedicated RPG. The RPS list allows you to select the products that are most relevant to your operations.

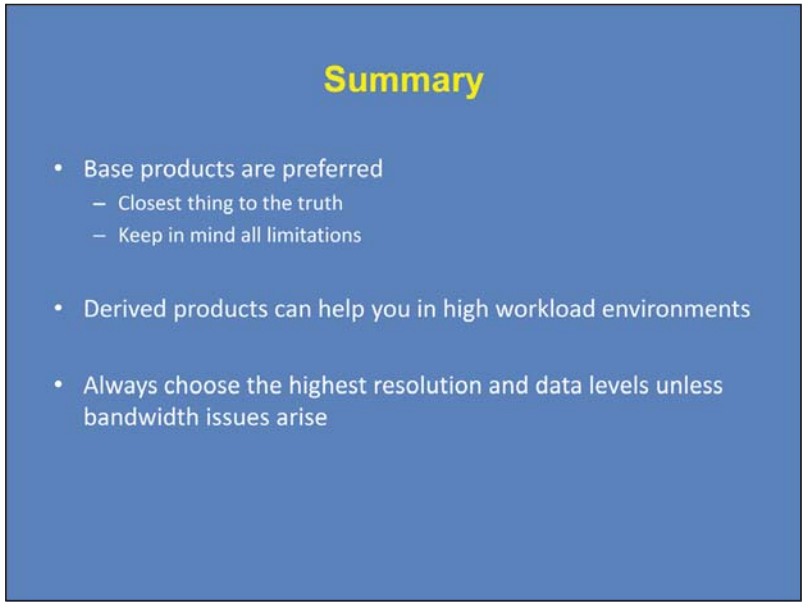
RPS lists are Volume Coverage Pattern (VCP) dependent, and the number of products you request varies between AWIPS and the RPG. The RPG counts every single product separately and will only transfer up to 300 total products. AWIPS groups separate elevation angles of the same product ID as a single product. As a result, AWIPS will only let you include 150 products (by its count) in a RPS list.



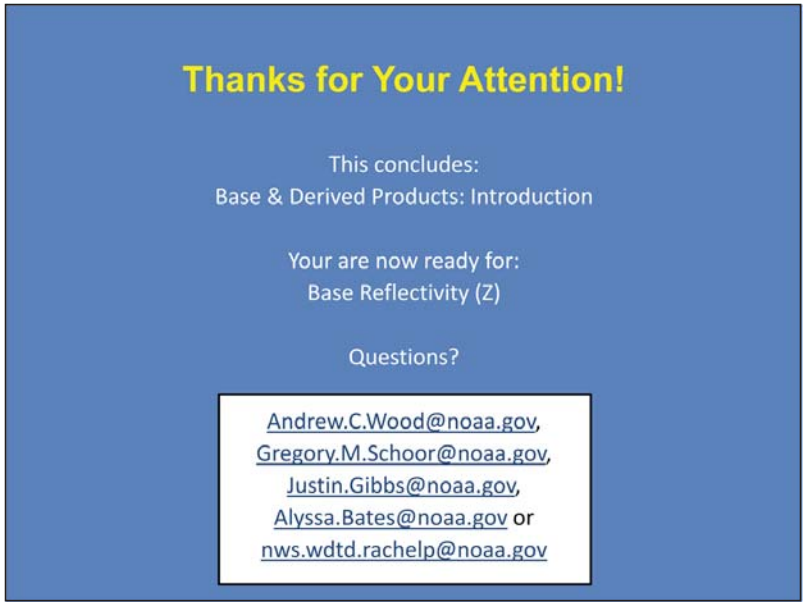
When building RPS Lists, it's important to know whether the products you request are from your dedicated radar (through a primary AWIPS connection) or a non-dedicated radar (through a supplemental AWIPS connection). Primary AWIPS connections have significantly more bandwidth than supplemental connections. And when I say significantly more, it's anywhere from 5-20 times the bandwidth, on average (Frashier, 2016).

When building a RPS List for a supplemental connection, forecasters should be selective. These sites are often radars located in surrounding areas that provide improve coverage of the fringes of your CWA. You probably don't need to include base data products for the entire volume from these sites. A little trial and error should help figure out what elevation(s) serve as a good cutoff.

The smaller bandwidth on supplemental connections may result in loadshedding in certain situations. Loadshedding occurs most frequently with faster VCPs (such as VCP 12 or 212) and in widespread precipitation as the product file sizes tend to increase with increasing storm coverage.



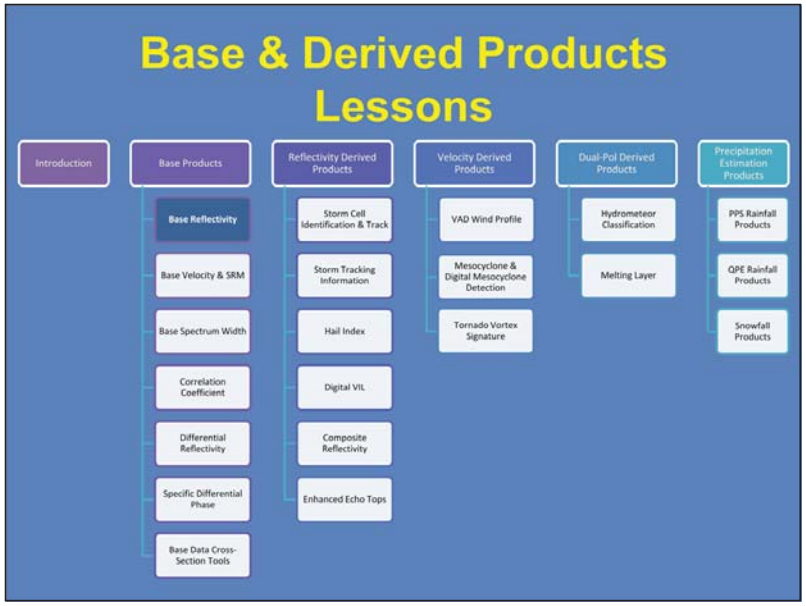
Here is a brief summary of this lesson. Remember that base products are the closest thing we have to meteorological truth, but these products have several important limitations. Derived products are great as a safety net, especially in high workload environments, but you should never issue a warning based solely on a derived product. Finally, always choose products with the largest number of data levels and highest resolution possible unless bandwidth issues arise. One such example would be requesting products using an RPS List on a supplemental AWIPS connection. Please proceed to the next slide to take the quiz.



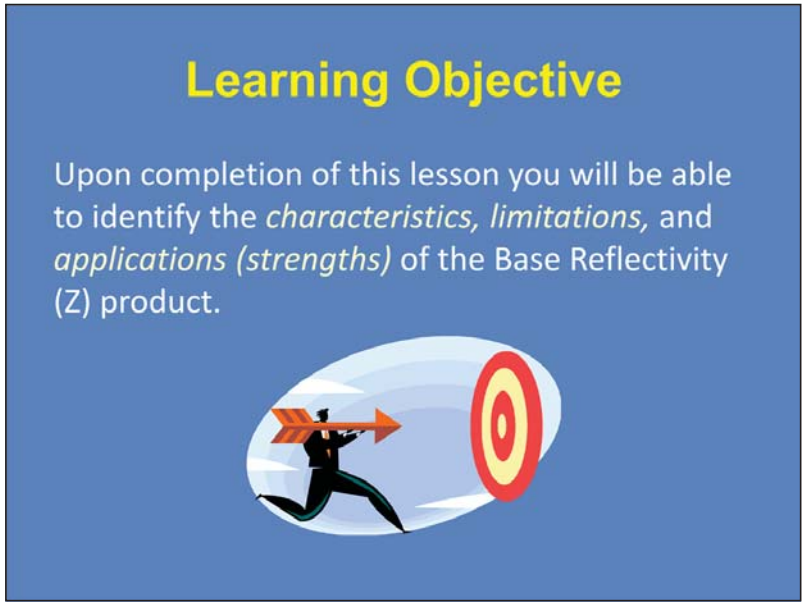
If you have passed the quiz, then you have successfully completed this lesson. You are now ready to move onto the next lesson, Base Reflectivity (Z). If you have any questions, please contact us using any of the e-mail addresses listed on the bottom of the slide.



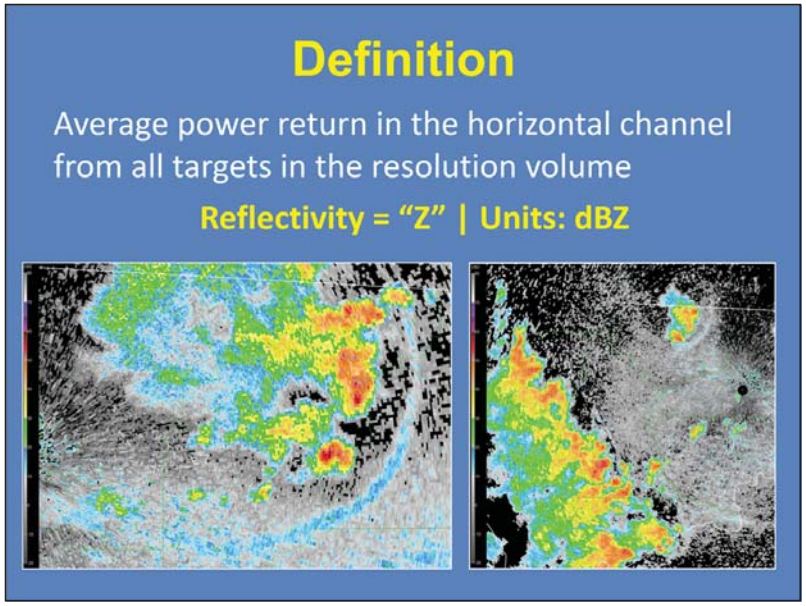
Welcome to Topic 4, Base and Derived Products. In this first lesson, we will start to build a foundation of understanding with the best known and most widely used base product – Base Reflectivity. So, let's get started!



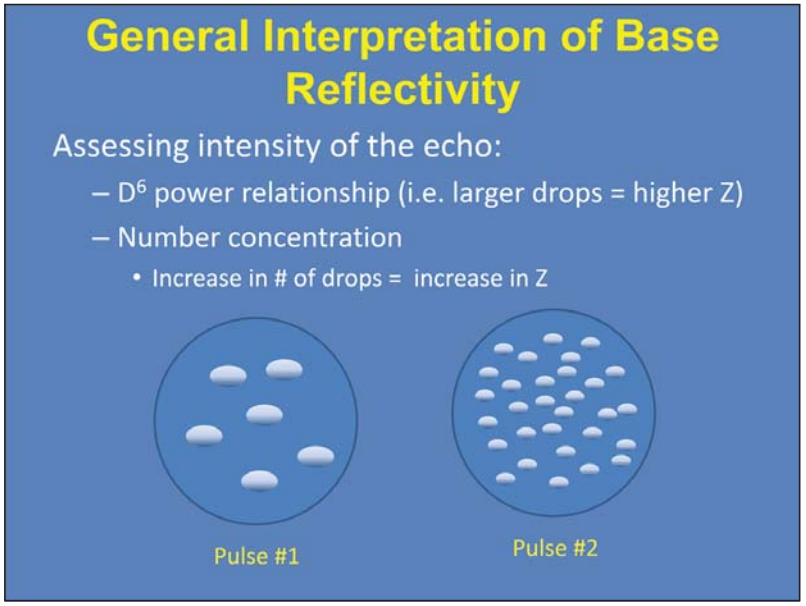
So this is the first lesson for the base products – along with the rest of the lessons in this topic. Also note the additional products to the right, which are derived from the base products coming up in following lessons.



Here is the learning objective for this lesson. Please advance the slide when you are ready.



Base Reflectivity is defined as the average power return in the horizontal channel from all targets within the resolution volume and is denoted with Z in terms of the product ID. So, the average power return for each pulse that is sent out comes back as to the RDA, is filtered and converted to a form of a decibels – like the amount of sound returned in an sound-driven echo, with the units of dB or decibels – of – Z, or reflectivity.



This power return, or reflectivity, is essentially the intensity of the echo – or as maybe you're used to hearing it – how reflective the targets are. It has a diameter to the sixth power relationship meaning larger drops result in substantially larger Z. Also, Z has a number concentration relationship. The more drops per unit volume that are present, the higher the Z. In the example shown, it is possible that even though pulse 1 has larger drops, pulse 2 may have higher Z because of the larger number of drops present.

Products Available from the RPG

Product Name	Units	AWIPS ID	RPG ID	RPG Code	Data Levels	Range Res	Az Res	Max Range
Base Reflectivity	dBZ	Z	R	19	4-bit (16)	1 km	1 deg	230 km
Base Reflectivity	dBZ	Z	R	20	4-bit (16)	2 km	1 deg	460 km
Base Reflectivity	dBZ	Z	R	21	4-bit (16)	4 km	1 deg	460 km
Base Reflectivity Data Array Product	*dBZ*	*Z*	*DR*	*94*	*8-bit (256)*	*1 km*	*1 deg*	*460 km*
Super Res Reflectivity Data Array Product	*dBZ*	*Z*	*SDR*	*153*	*8-bit (256)*	*0.25 km*	*0.5 deg*	*460 km*

**** Only the gray highlighted products will be discussed in this lesson*

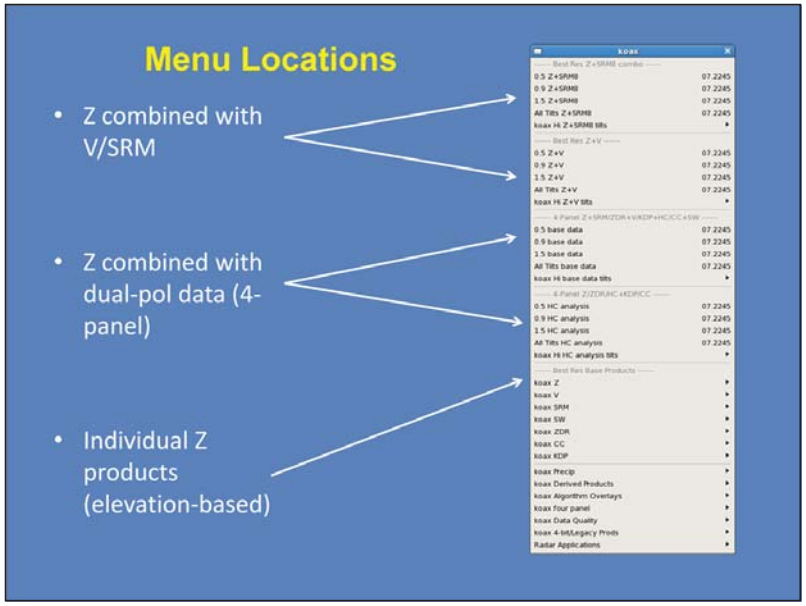
Here is a table listing all the Reflectivity products available from the RPG. In this lesson we will only focus on the ones highlighted in gray. The other products are lower resolution products that are primarily available for bandwidth reasons only.

General Product Characteristics

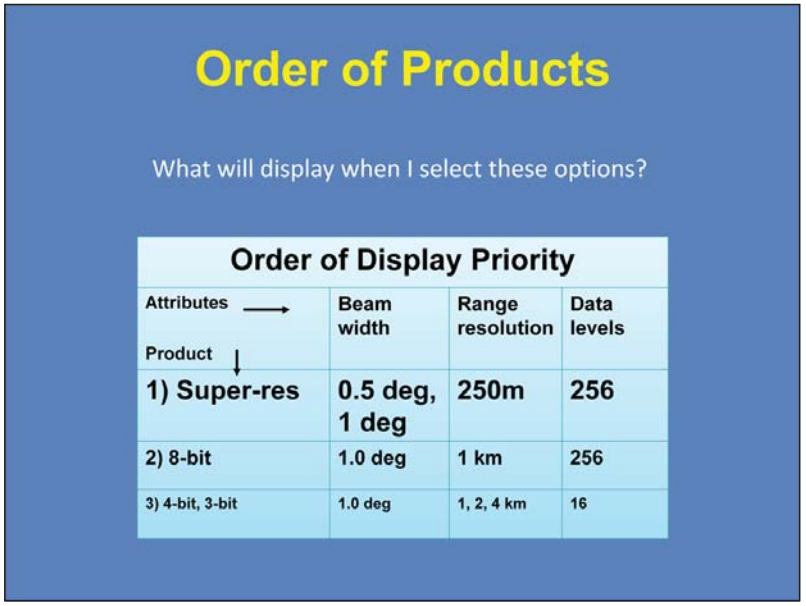
- Displayed as a polar coordinate with a 1 degree beam width
  - Super-res has 0.5 degree beam width
- Displayed relative to the RDA site
- Available for any elevation angle of the current VCP



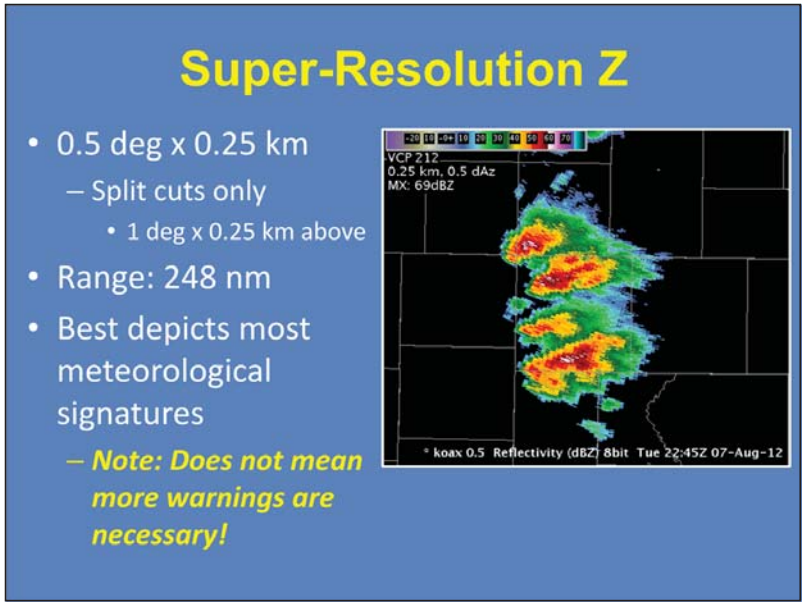
Here are the general characteristics for Reflectivity. It is displayed in polar coordinates with a 1 degree beam width, except for super-res products which have a 0.5 deg beam width. All pixels are relative to the RDA site and the data are available for all elevations for the current VCP.

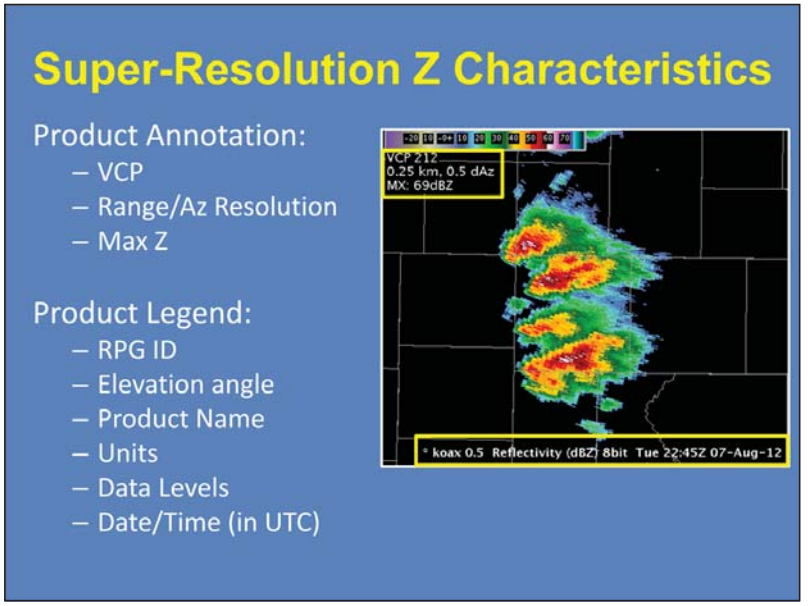


In your dedicated radar's drop-down menu, you can find Z products in these locations. The top four listed are Z combined with the various other base products, like velocity or SRM, so you can switch back and forth quickly, at different levels. If you'd like to just load Z by itself, it can be found toward the bottom of this list, shown here by the last white arrow.



Just a quick reference to the order of products, which is important with regard to how AWIPS-2 loads products. It starts by looking for the highest resolution possible, and then works its way down the list. Here is a table showing the order of loading for Z in AWIPS-2. The same will be true for the velocity products.

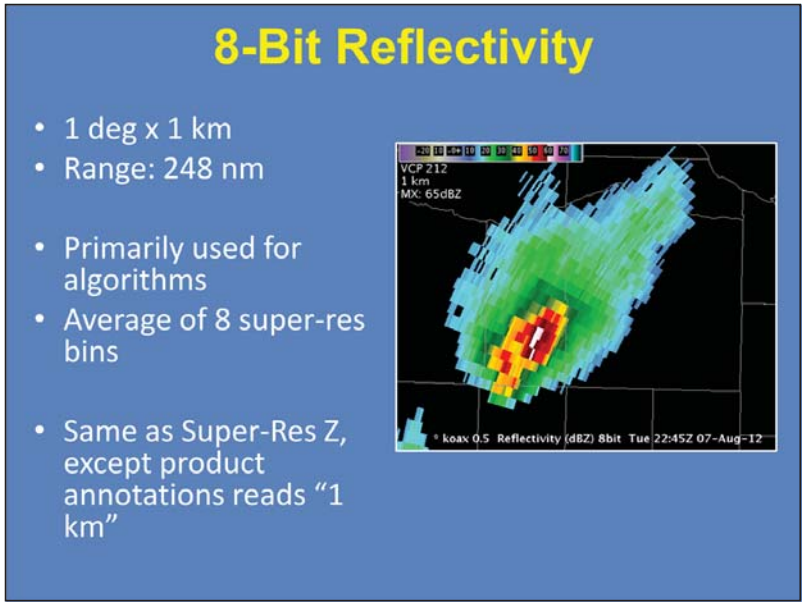


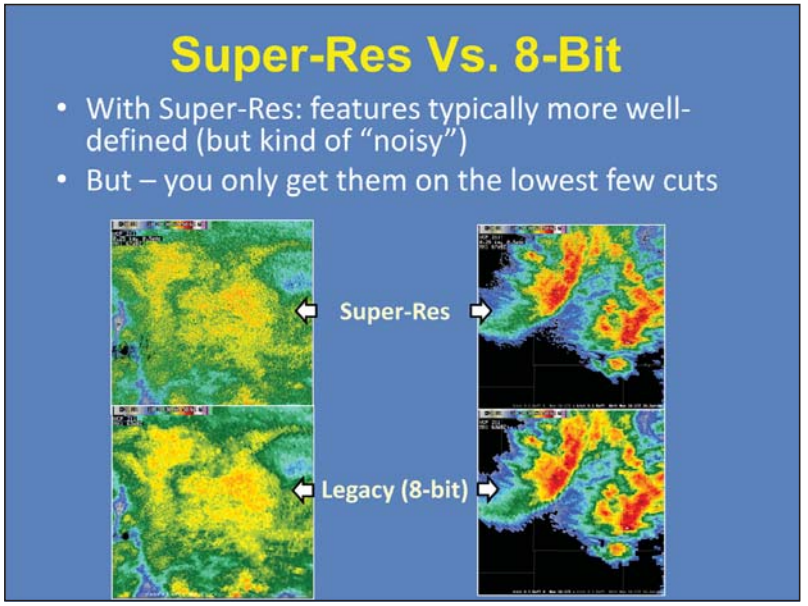


The first Reflectivity product we'll discuss is super-resolution reflectivity, which has been around since the late 2000s. Before then, 8-bit reflectivity was the highest resolution. It has a resolution of 0.5 deg by 0.25 km in the split cut elevations, and 1 deg by 0.25 km ABOVE the split cut elevations. It extends to a range of 248 nm (or 460 km) and is the best product to use for identifying most meteorological signatures. However, do not assume easier identification necessarily means more warnings.

In the upper left corner of the super-res Z product, you will find the product annotation information. Here you will find what VCP the radar is operating in, the range and azimuthal resolutions of the product (in the case of super-res, you'll see 0.25 km, 0.5 dAz), and the maximum Z anywhere in the radar umbrella. Note that it does not tell you where that max reflectivity occurs, which doesn't really matter unless you know generally where that max would be, because often times – especially if there is terrain within the radar's reach, it may flag a pixel or two of the top of a hill or mountain with a very high DBZ value, so keep that in mind.

On this example image, note the product legend on the bottom right. The legend tells what radar it is, the elevation angle being displayed, the product name along with the units, how many data levels are in the product and the date/time in UTC.

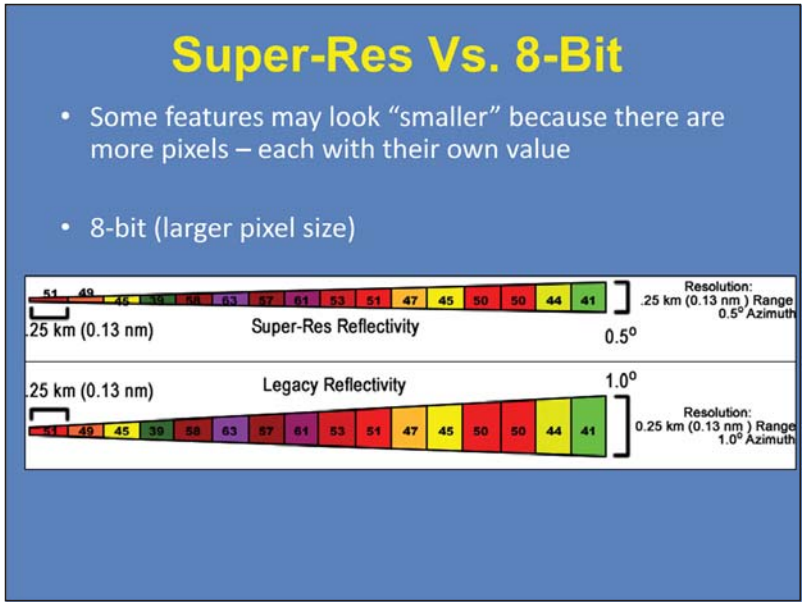




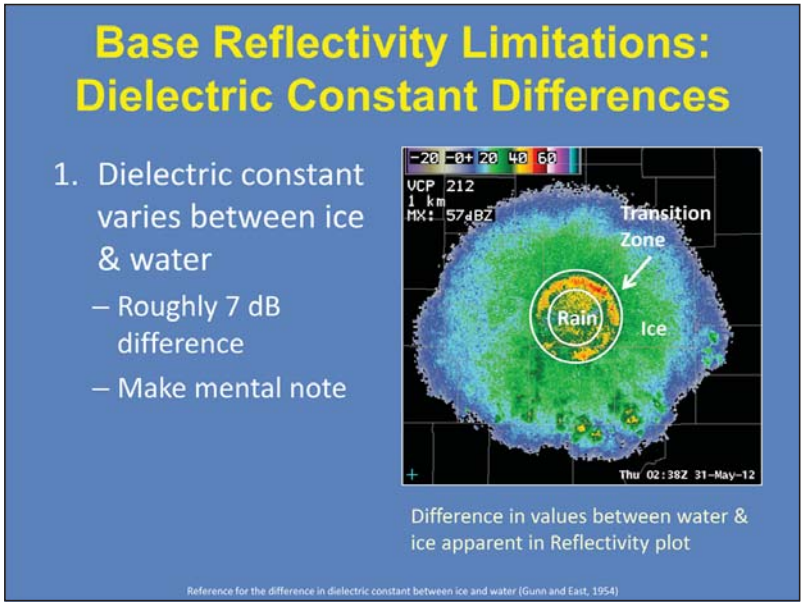
The 8-bit Reflectivity product is much like the super-res Z product except it is 1 deg by 1 km in resolution. But it does have it uses, mainly for some of the algorithms in derived products. When Super-Res came about algorithms were not been modified to ingest super-res Z data. It is generated at the RPG by taking the average power of the 8 super-res bins that comprise one 8-bit bin and then converts that average power to a Z value.

The product legend and annotations are identical to the super-res Z product except the resolution just reads 1 km.

In general, super-res Z is much preferred over the 8-bit counterpart because it tends to show features much more easily. However, one thing to note is that super-res Z is noisier than the 8-bit product due to the way super-res Z is computed. But this should not hinder interpretation.



Also note that higher resolution doesn't necessarily mean bigger, better features, most of the time, it means smaller features. The reality is – with super-res, hail cores and other signatures, within the split cuts, that have high DBZs will usually be smaller than what you would see in the 8-bit version. Remember that what is displayed in each pixel is the highest DBZ in that pixel, so if you have a larger pixel, like 8-bit typically does, then it will look like there is more area for that really high value, when in reality it could be a smaller area. This is especially true at higher cuts typically, because there, you won't have Super-Res, just 8-bit, so if the storm is far enough away from the radar, the size of, say a 70 DBZ echo will actually be smaller than if you would be able to view it with Super-Res.



So, believe it or not, there are some limitations to reflectivity, or maybe you could see it as – things to watch out for.

First, let's look at this aspect of the dielectric constant, which can and usually does affect the values of Z. Ice has a roughly 0.2 dielectric constant, whereas liquid water has nearly 1.0. This means for the same size particle, ice will be roughly 7 dB lower than the liquid particle. In the image on the right, near the radar and scanning at a relatively higher elevation, the beam is heading up into the cloud a steep angle, where liquid water, rain, is detected close to the radar but then the beam encounters the melting layer, inside the two white circles, where you see the ring of red and yellows which is the transition zone between liquid water and ice. Then, as the beam continues on up into the cloud, to the right of the outside white circle, it is all ice. Notice how the reflectivities are generally lower in ice than in liquid. This is due to the dielectric constant difference in ice and water. And even though this is technically a feature-detection – of the melting layer – it is a limitation for precipitation estimation. The radar tends to add up values pretty quickly over the area where the melting layer resides, which obviously is not reality at the surface, just a feature hanging several thousand feet or more up in the cloud, depending on the time of year. (Reference: Gunn and East, 1954)

Base Reflectivity Limitations: Chaff

2. Chaff looks similar to precipitation in Z
   - Military ops
   - High dBZ initially, with gradual decrease
   - Review all observation systems
   - Recurrence over same area



Stream of chaff visible in Base Reflectivity

Chaff, is pretty rare to see, but does happen on occasion and will appear on the scope as a long, thin stripe, like in the example image. Chaff is made up of very tiny long and thin fibers that are released by military aircraft for military training purposes, which is not an incredibly common thing but something to be aware of. It usually shows up as light to moderate Z during initial release and then decreases in intensity as the chaff disperses down wind. There are ways to figure out that this is not precipitation, mainly using Dual-Pol products, but also, you can tell with various satellite data whether this is actually a dense precipitation-bearing cloud there, and the military normally choses not to release this stuff when there is precipitation, because it would kind of defeat the purpose of detecting it.

Base Reflectivity Applications

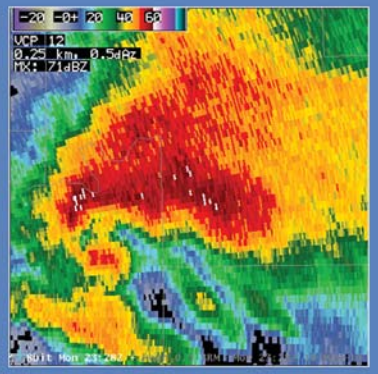
1. Observe precipitation intensity, movements & trends
   - Time lapse feature is valuable
   - Clear air mode for light precipitation



So, let's go over the various applications and benefits of reflectivity. The number one application of reflectivity is to observe precipitation intensity, movements, and trends. Setting a time lapse of reflectivity is invaluable. This allows you to see the morphology of an event and anticipate future development. Additionally, in light precipitation events, clear-air mode is preferred for the same reasons.

Base Reflectivity Applications

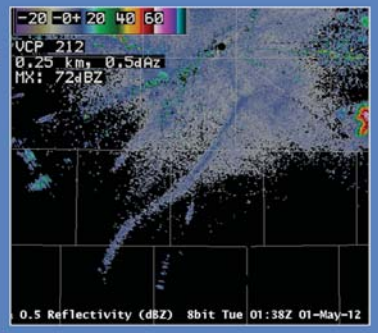
2. Determine significant storm structure features
   - WERs, BWERs, Hooks, Weak Echo Channels, TBSS, etc.
   - Use All-Tilts, 4-panels or FSI



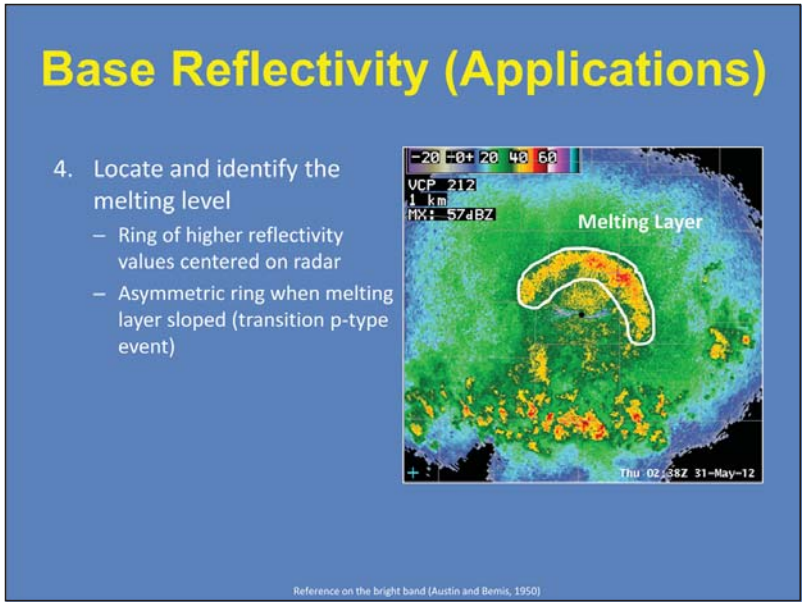
When interrogating storms, reflectivity can be used for identifying significant storm structure features that indicate potential for a severe thunderstorm. Some of these features include weak echo regions (WERs), bounded weak echo regions (BWERs), Three body scatter spikes (TBSS), etc. It is highly recommended that you display different levels of reflectivity in either 4-panel or All-Tilts views and utilize FSI – which is the 4-D Storm Interrogation tool, which is a separate window you can use in AWIPS-2 (which we'll discuss in a later lesson) to observe these features.

Base Reflectivity (Applications)

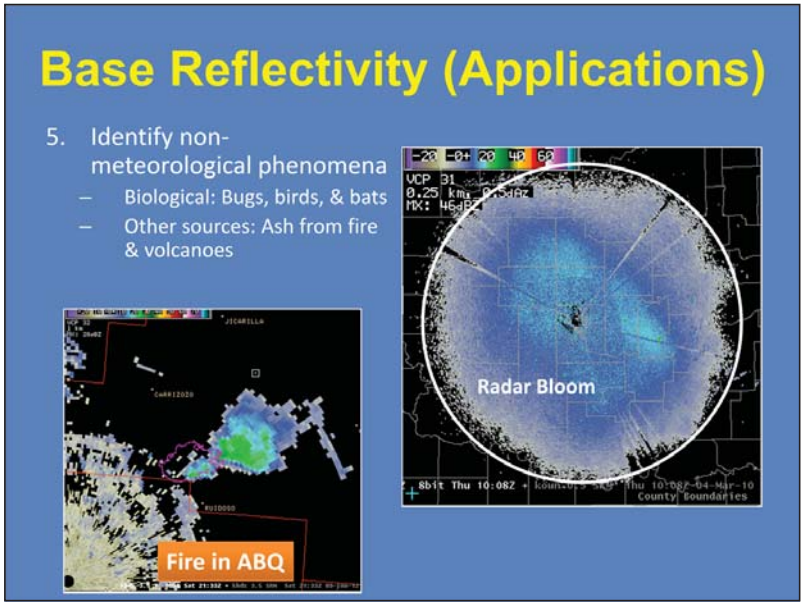
3. Determine location and motion of fronts and other boundaries



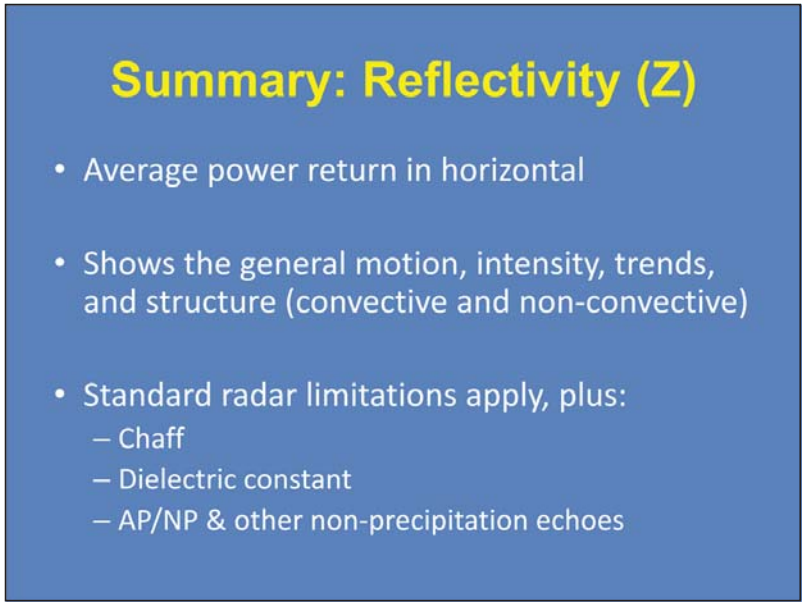
Along the lines of the first application, reflectivity can be useful in identifying the motion and trends of fronts and boundaries. The image on the right shows an example of a boundary approaching the radar from the southwest. Most of these detected boundaries show very low reflectivity values and appear because they are carrying non-meteorological particulates in the air as they progress forward.



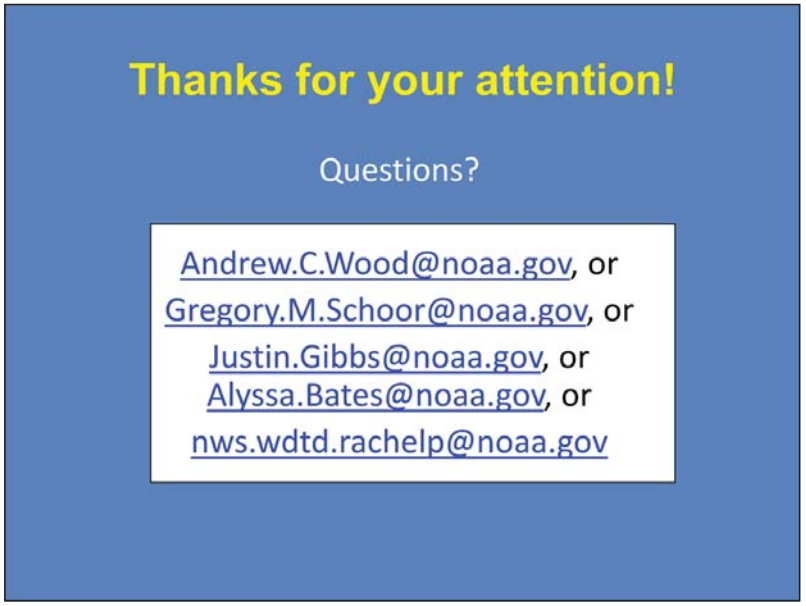
Back to the melting layer, which we noted for its limitations due to the overestimation of precipitation in that zone. However, this is also an application of reflectivity, because it does show the melting layer and when you have sampling turned-on, you can see both where it is above the surface, if there is a melting layer, and what the general thickness is of that layer. If the melting layer shows up as an arc – as you can see in this image, you most likely have a transition precipitation event. Reference: (Austin and Bemis, 1950)



Finally, reflectivity can identify non-meteorological echoes, and even though fronts and boundaries fall into this category, here are some additional features. Typically, during the evening and morning hours, you will note an increase in the coverage of reflectivity in clear-air situations. This is often called the “radar bloom”. This is prime time for bird migrations, insect feeding, etc. So, what you are actually seeing is biological patterns. Over time, this can be helpful in interpreting radar signatures during different times of the year. Also, sometimes you can detect fires or wildfires, and even volcanic ash.



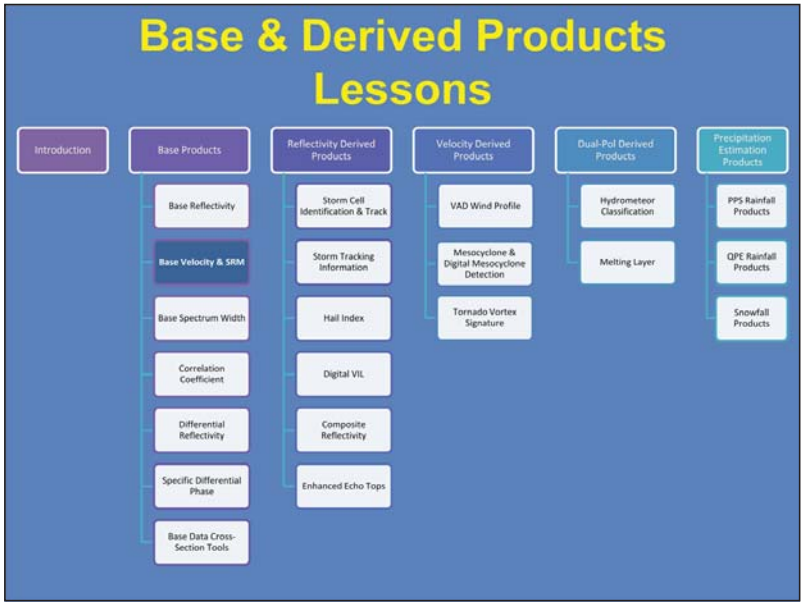
To summarize our lesson on reflectivity, recall that it is the average power return in the horizontal channel. The resultant power amount returned to the radar can show the general motion of echoes, intensity, trends, and structure of precipitation within a cloud, whether you're looking at convection or non-convective activity. As always, standard limitations apply in addition to chaff appearing as precipitation and the dielectric constant being different between ice and water, as well as other non-precipitation type features, such as terrain.



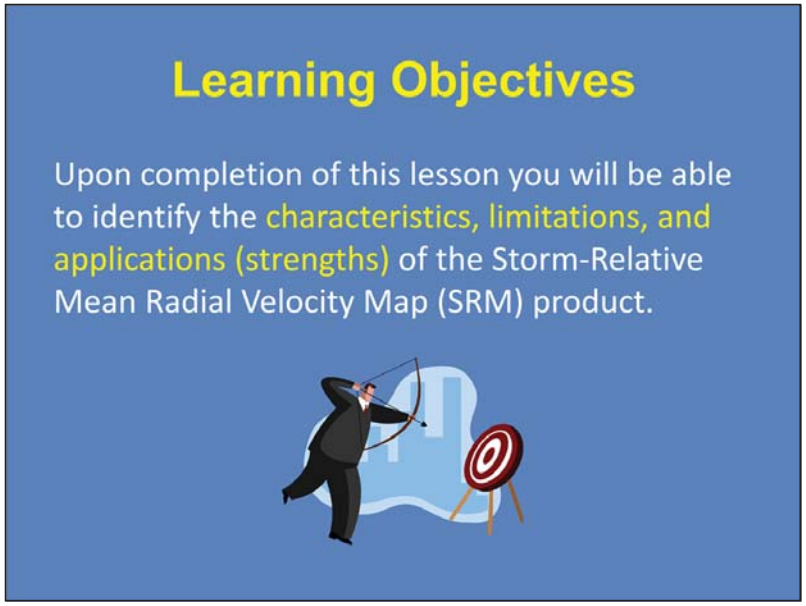
If you have any questions, feel free to email any of the contacts below.



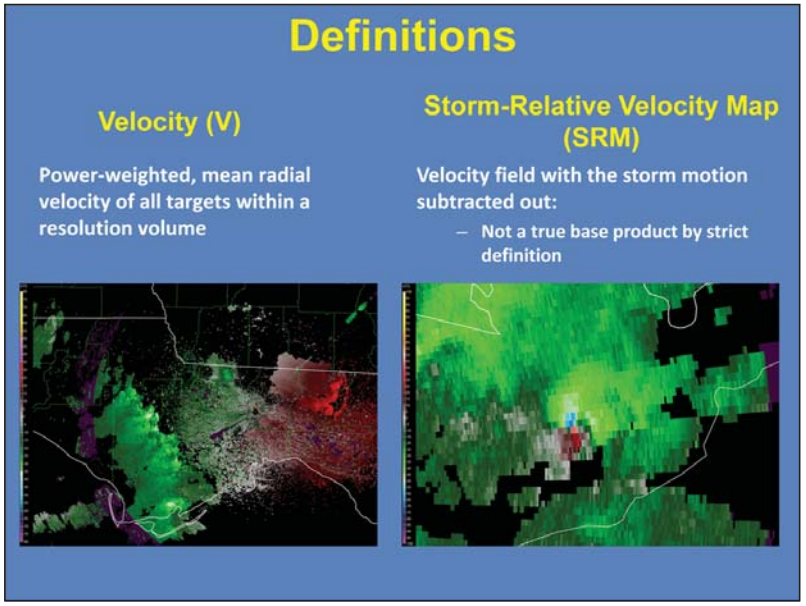
Welcome again to Topic 4, and in this lesson, we'll be talking about base velocity and Storm-Relative Velocity Map, what is commonly referred to as SRM. These lessons were previously separate and even though SRM is not a base product, but a mathematical derivative of velocity, they are essentially the same in terms of their fundamental specs, and are now combined in order to draw on their similarities and their differences. So, let's get started!



We're now into the second lesson on base products, with more to come in terms of the derived products after this series of lessons is complete.

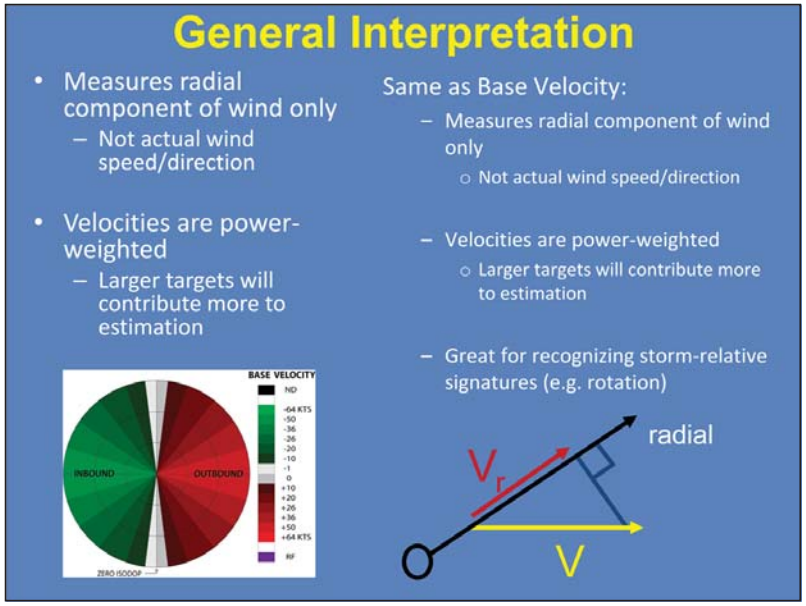


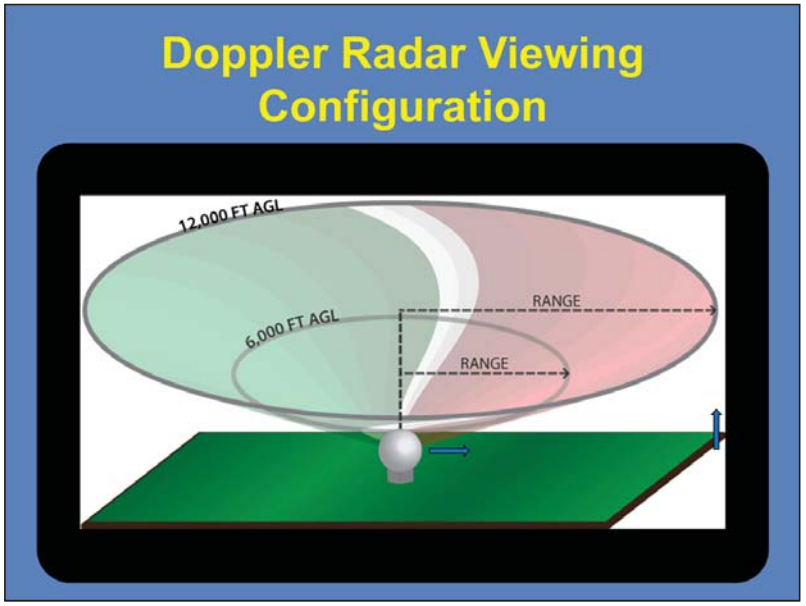
Here is the learning objective for this lesson. As mentioned, in the past, velocity and SRM were separate lessons. But even though SRM is a derived version of base velocity, they can both be compared and discussed at the same time. So, remainder of this lesson will be a kind of back and forth between the two along with some comparisons made. Advance to the slide when you are ready.



What is velocity? Well actually, it is technically called radial velocity, and it is the power-weighted, mean radial velocity of all targets within a resolution volume.

And Storm-Relative Velocity Map, or SRM, which is not a true base product in the strict sense, is really just velocity with a storm motion subtracted out. But as we'll see over the next few slides, the characteristics, strengths and limitations are very similar.

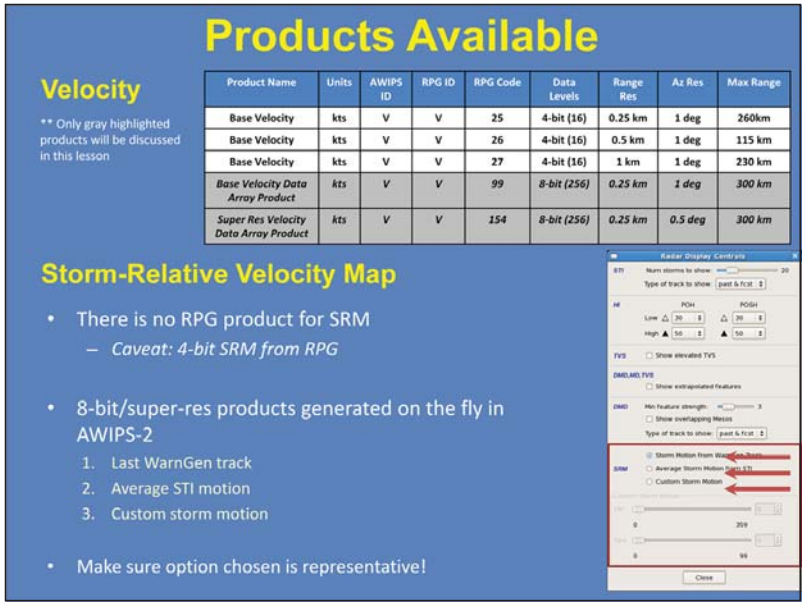




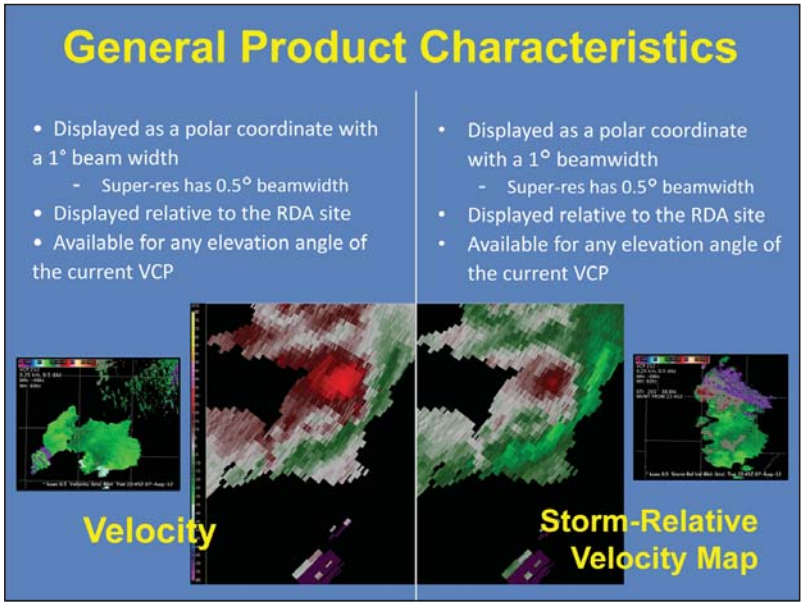
Recall from Topic 3 that radial velocity only measures the component of the wind toward and away from the radar. Any component perpendicular to the radar is not measured. So base velocity is not the measure of the true wind unless the true wind is oriented along a radial. Also, because velocities are power-weighted, larger targets will contribute more to the estimation than smaller targets.

Interpreting SRM is identical to V. It is a power-weighted measure of the radial component of the wind and velocities. The main difference between V and SRM is SRM is great for recognizing storm-relative signatures like rotation, whereas V is great for recognizing ground-relative signatures like microbursts, keeping in mind the sampling angle.

More about this sampling angle. When you are looking at a velocity product, you are viewing the display from above looking into a cone. As you move farther away from the RDA, you are also increasing in height above the ground. When you are looking at any plan-view product on the screen, so not just velocity, keep in mind how you're increasing in height further away from the radar. When the screen is filled with high in-bound and out-bound velocities, it doesn't necessarily mean that those speeds are translating to the ground, unless there's a tropical feature on the scope – you may have a jet max overhead at some level several thousand feet up. But notice how the velocities decrease on the top side of this cone and at the bottom but are maximized on the right and left sides. This does not mean that the winds necessarily drop off from one side to the other, based on the radial sampling however, the values of velocity do decrease as you go from left to right, until you cross the zero isodop – white, grayish line, then increase again toward the right.



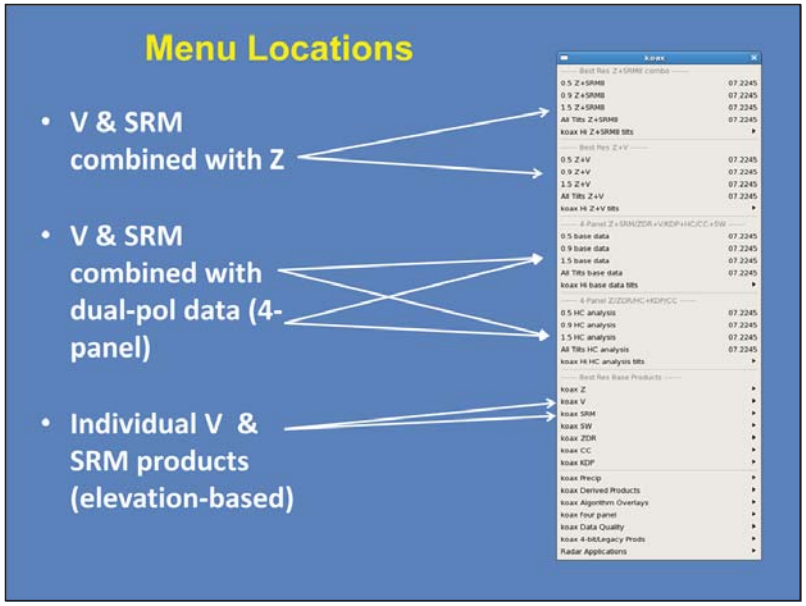
Product Name	Units	AWIPS ID	RPG ID	RPG Code	Data Levels	Range Res	Az Res	Max Range
Base Velocity	kts	V	V	25	4-bit (16)	0.25 km	1 deg	260km
Base Velocity	kts	V	V	26	4-bit (16)	0.5 km	1 deg	115 km
Base Velocity	kts	V	V	27	4-bit (16)	1 km	1 deg	230 km
Base Velocity Data Array Product	*kts*	*V*	*V*	*99*	*8-bit (256)*	*0.25 km*	*1 deg*	*300 km*
Super Res Velocity Data Array Product	*kts*	*V*	*V*	*154*	*8-bit (256)*	*0.25 km*	*0.5 deg*	*300 km*



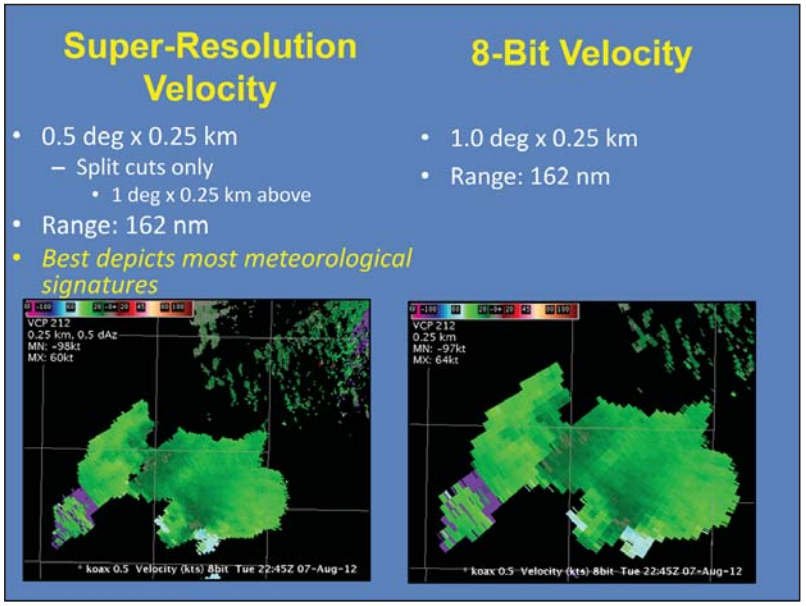
Here's a table of all radial velocity products available in the RPG. We will only focus on the ones highlighted in gray. The 4-bit products are only really needed if there are bandwidth issues.

There is no RPG product for SRM, except for the 4-bit RPG generated SRM – but that is probably not a widely used product. The 8-bit and super-res versions of SRM are actually generated on the fly in AWIPS-2 via one of the three options available from the Radar Display Controls window. Those three options are: the latest WarnGen storm motion, the average motion of all SCIT identified storms, or a custom storm motion. The first two options are nice because they automatically update when each of those products update. However, when the storm motions from those products become unrepresentative of the storm you are investigating, then the custom storm motion becomes handy, but that storm motion only updates when you update it. Which option you choose is totally up to you, but make sure whatever option you do choose best fits the storm motion of the storm you 're interrogating.

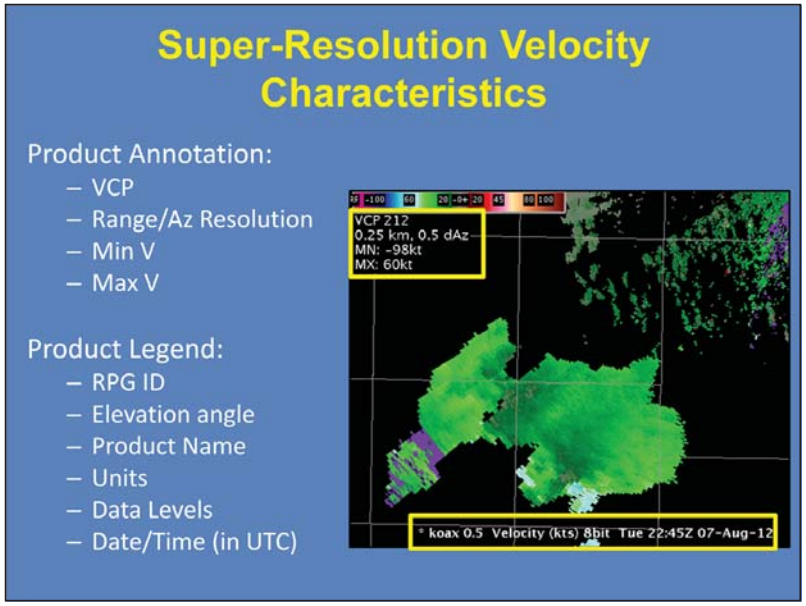
Velocity and SRM share generally the same characteristics because they are the same, SRM just has the storm motion factored in – as we mentioned. For all velocity products, the data are displayed in polar coordinates with a 1 degree beam width, the exception being super-res which has a 0.5 degree beam width. All data points are relative to the RDA site and are available for all elevations angles for the current VCP.



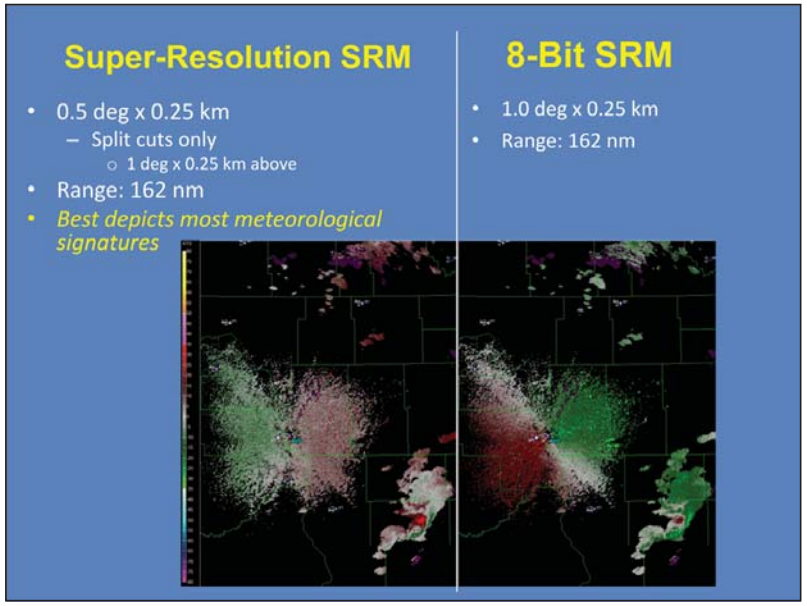
Take moment to see where in the menu the velocity and SRM products are located. Unlike most products, velocity and SRM are unique in-that they are combined with reflectivity in a number of angles – for good reason. So, velocity products are all over this menu, whether in their standalone versions or in combos of 4-panel and all-tilts.



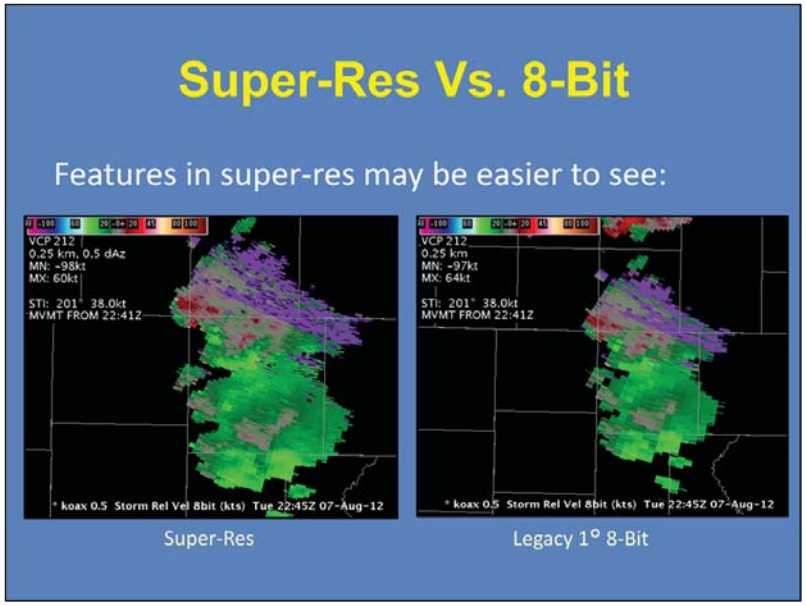
The characteristics of super-res velocity are listed here. In the split cuts, it is 0.5 deg x 0.25 km, and above the split cuts the azimuthal resolution goes to 1 deg. Like reflectivity, super-res velocity is best for identifying most meteorological signatures. As you can see from the two images, the one of the left, the super-res have much more detail for the same region of sampled velocities than the one on the right. However, super-res velocity can be noisier than 8-bit, but that should not hinder any interpretation. And just like super-res reflectivity, just because you have higher resolution, it doesn't necessarily mean more warnings should be issued.



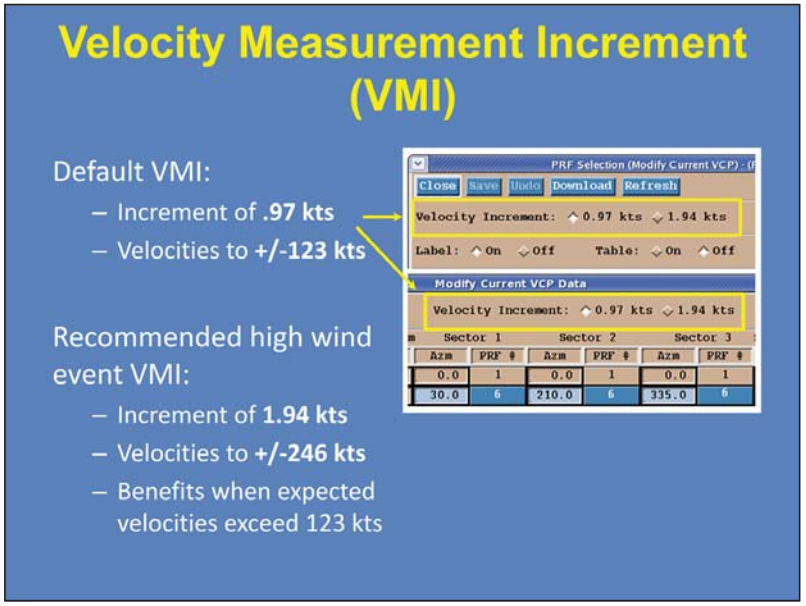
In the upper left corner of the velocity product is the product annotation information. This tells you the current VCP mode, the range and azimuthal resolution and the min/max velocity values anywhere in the radar domain. Note, that the location of these min/max values are not known. In the lower right is the product legend. This tells you what radar you are viewing, the elevation angle, product name (in this case Velocity), the units, how many data levels and the date in UTC.



Since SRM requires the velocity product to be available, the super-res SRM has the same characteristics as velocity. It is 0.5 deg by 0.25 km in the split cuts and 1 deg x 0.25 km above that. It goes out to a range of 162 nm, and is best for identifying most meteorological signatures. Again, don't assume it will mean more warnings.

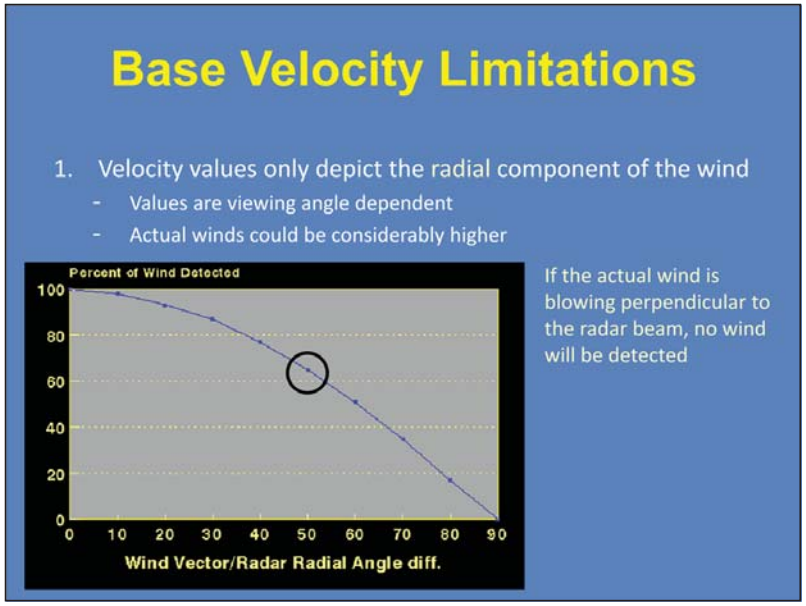


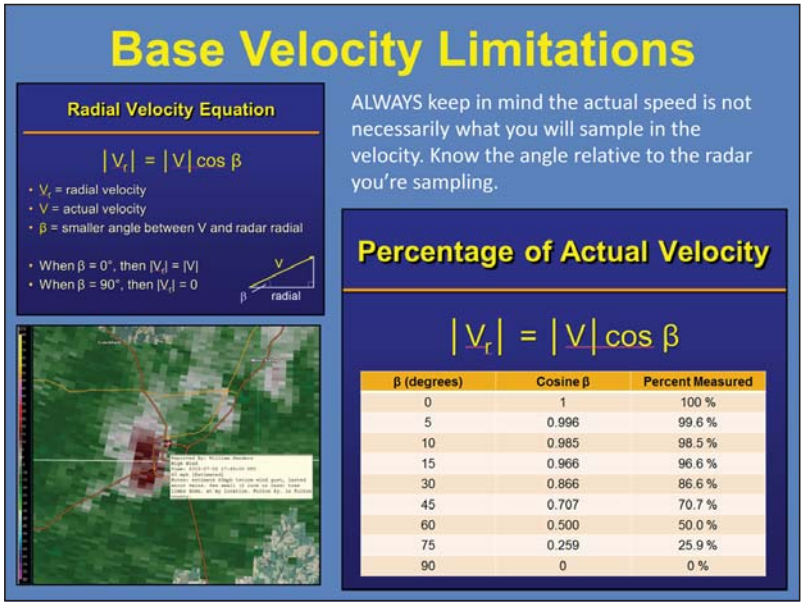
Again, just like velocity, the one thing you will notice with super-res SRM versus 8-bit SRM is that features in super-res may be easier to see than features in the 8-bit product. However, remember that super-res data are noisier than 8-bit data, but that should not affect your interpretation.



There is a unique element to velocity measurement that we need to briefly cover. Recall with data levels that for a given range of values, you can only get a certain precision. Well, with velocity, the default precision is 0.97 knots, which for 8-bit gives you a range of velocities of 123 knots inbound and outbound. In other words, even if there are winds over 123 knots, at that precision, the radar won't be able to sample it and will max-out at 123 knots. This may be an issue if there is a tropical system that is carrying winds above that level.

The only way to increase this range is to add more bits, or change the precision of the data. We can't change the bits, so this is where the Velocity Measurement Increment (VMI) comes into play. By changing the VMI (or precision) from 0.97 knots to 1.94 knots, the range of velocity values increases to 246 knots inbound and outbound (or doubles), which is basically double. The default VMI is 0.97 knots and for most meteorological events, that is sufficient, but in high wind events like hurricanes, changing this value is recommended. You can find the VMI settings in the RPG HCI via the PRF selection window.





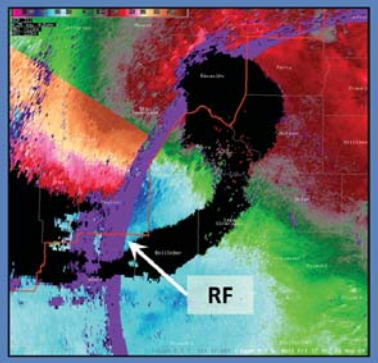
Now we'll look at some issues that may affect velocity interpretation. Velocity has a couple of potentially big limitations. It just depends on the scenario you're dealing with, whether they become an issue. The first is Range folding, or RF signals with no velocity assigned – purple haze as it's sometimes referred to, as the default color is purple. This is one of the most crucial elements to remember when you're looking at live radar data.

The other big thing to remember is that velocity only depicts the radial component of the wind. Therefore, if the actual wind is blowing perpendicular to the radial, no wind will be detected by the radar just based on the geometry. The graph below shows the percentage of wind that is measured based on the angle it makes to the radial. At an angle of 50 degrees, the radar is only seeing roughly 65% of the actual wind! Remember this and continue to remember it well, that when you are sampling velocity values, keep in mind that angle – relative to the radar – that the data is being sampled it, it may not be a true value to what the velocity actually is at that level and that angle.

Actually, I'd like to hammer this point home a bit more because velocity will be such a large part of your radar interrogation and warning strategy. When you're sampling areas, especially around 30 degrees on either side of the due north and due south portions of the radar's view, you're only sampling 50% or less of the true velocity of those particulates. The two blue background images are oldies but goodies, in terms of understanding that you can have a large reduction of values on the northern and southern peripheries of the radar. You obviously won't be able to do this equation for each pixel on the scope while you're interrogating radar data, but keep this in mind in case you run into a situation like the example on the lower left. It may be a little tough to see, but essentially there is a marginally severe downburst occurring near and within the darker red area. However, sampling this region, you're only going to get a range of 10 to 20kt outbound pixels at best, and we are around 190 degrees relative to the radar – or south-southwest. That's a potential reduction of half or more from the true value!

Base Velocity Limitations: Range Folding

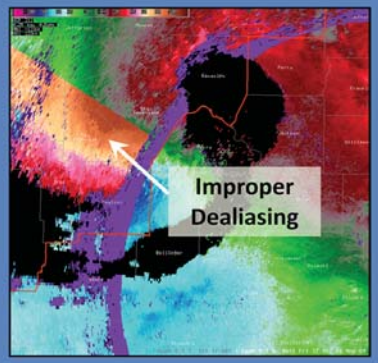
2. Mitigate range folding impacts by:
   - Changing PRF
   - Changing to VCP 121, 211, or 212
   - Invoking clutter suppression
   - View a higher slice
   - Select another radar



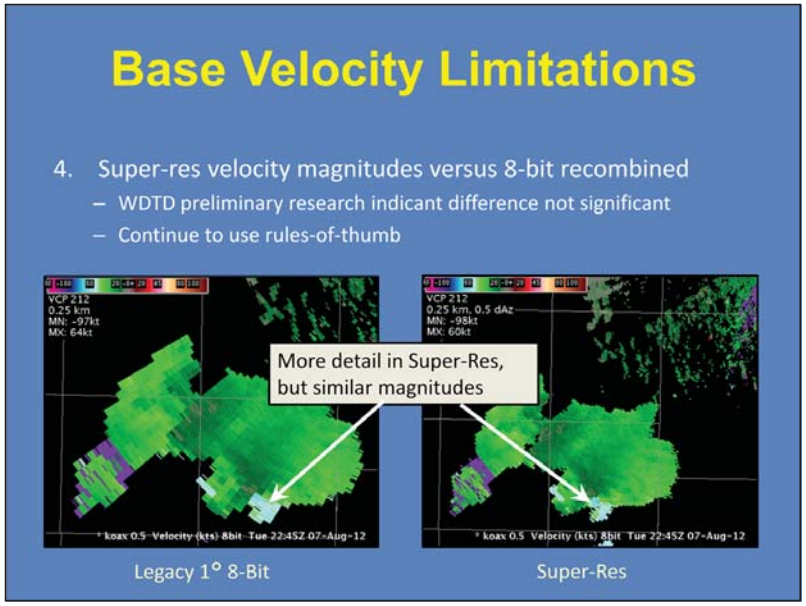
As mentioned in the previous slide, Range folding – or RF – is a major issue with velocity. It can mask features that are crucial to warning decisions (e.g. tornadoes). If RF is an issue, consider the following: Change the PRF, switch VCPs to 121, 211 or 212 since they are best at mitigating RF, invoke clutter suppression, change to a higher elevation, or choose a different radar.

Base Velocity Limitations: Improperly Dealiased Velocities

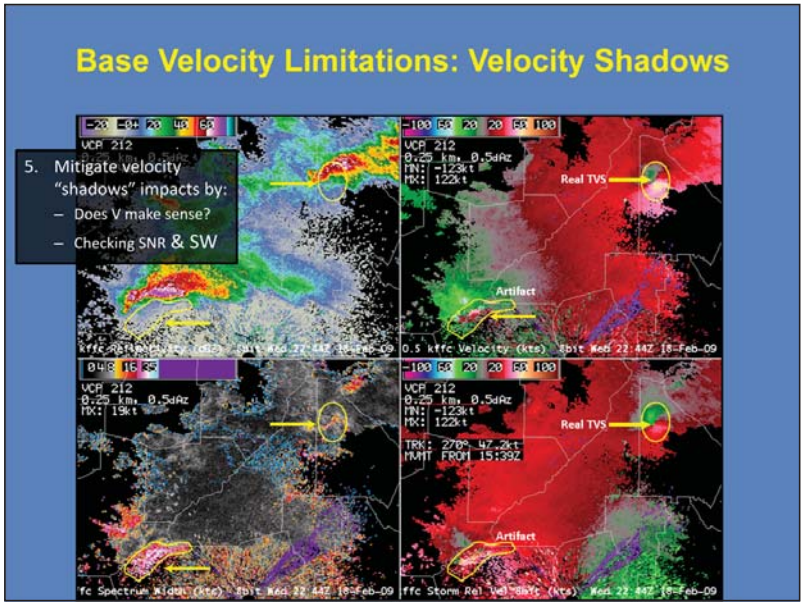
3. Mitigate improperly dealiased velocity impacts by:
   - Changing to VCP 121
   - Updating Environmental Winds Table
   - Changing PRF
   - Viewing different slice
   - Selecting another radar



Improper dealiasing can be an issue with velocity although it's mostly been minimized with some recent work to mitigate it through a number of different algorithms. The image on the right is a good example of dealiasing, which is a tough word to say and even more tough to properly extrapolate true values when you see it. In the image, there is a large wedge of dealiased data, mainly the bright orange colors indicated by the black arrow. Here are few things you can do: Change to VCP 121 (it has a special dealiasing algorithm in it), update the environmental winds table, change the PRF, select a different elevation angle, or change radars.



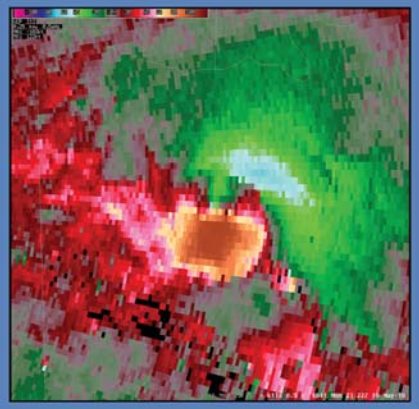
It was believed that with super-res the velocity magnitudes, or thresholds, of certain features, like the TVS, may need to be adjusted. Results from an in-house study at the Warning Decision Training Division have shown that there is not a significant difference in velocity magnitudes from super-res versus legacy for features such as the TVS. More studies need to be conducted to confirm this finding, but at this point rule-of-thumb thresholds should still be used for warning purposes. It should be noted again here that this does not mean features might not be easier to identify. Super-res data will provide more detail and might make recognizing the features easier, but don't confuse that with the magnitudes of the velocities.



Another velocity limitation loosely termed "velocity shadows" can significantly affect warning decisions, particularly for tornado warnings. Numerous examples exist from the field where tornado warnings were sent out because of this signature. Let's define what they are: Velocity shadows are found in very weak reflectivity returns, normally on the inflow side of storms. You'll typically see them on the lowest elevation angles, and they are likely due to side lobe from very strong echo overhang. Here is an example of two storms at the 0.5 degree elevation angle. Notice that the SW storm has very weak reflectivity returns on the inflow side, and there are very strong rotational signatures in this area. For comparison, the storm to the NE also has very strong rotational signatures, but is a real signature and is far less noisy. For any strong velocity signature, check to see if it makes sense in terms of where in the storm it is located, if there is sufficient reflectivity return associated with it, and if spectrum width and/or general data quality is not overly noisy.

Base Velocity Applications

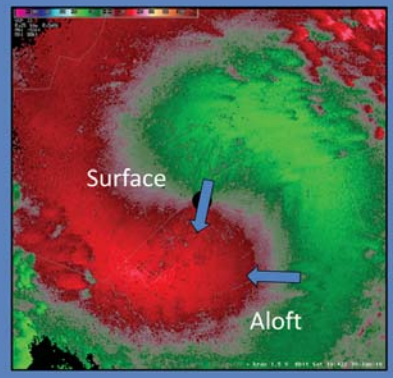
1. Analyze storm structure
   - Gust fronts
   - Microbursts
   - Straight-line winds, etc.



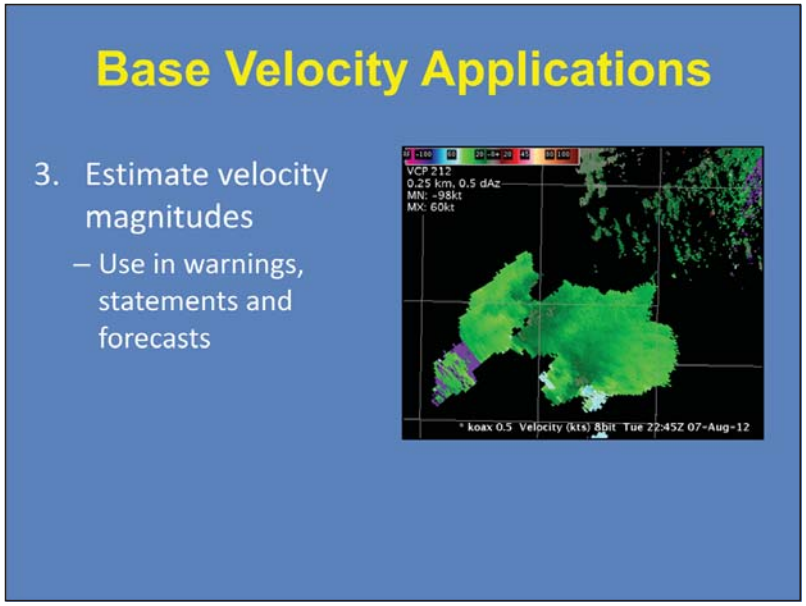
The primary application of velocity is to identify significant storm-scale structures such as gust fronts, microbursts, straight-line winds, and many other storm-scale signatures.

Base Velocity Applications

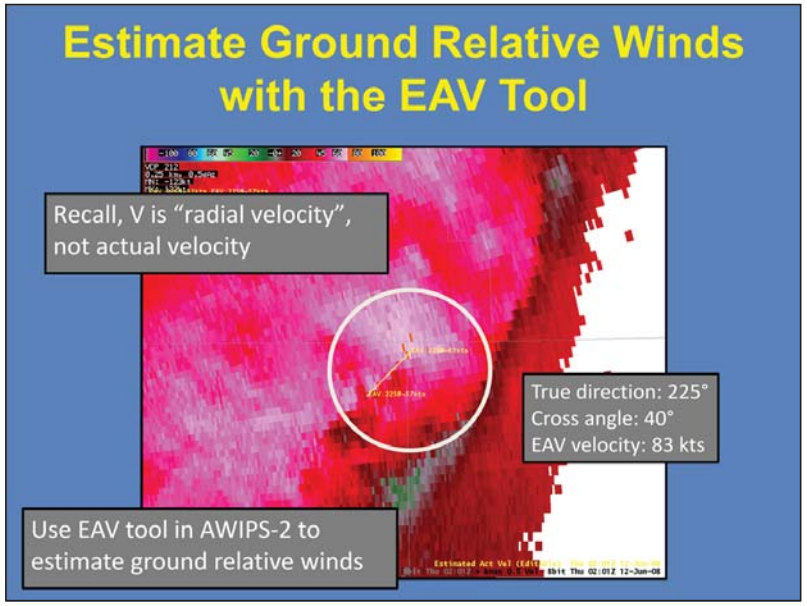
2. Determine atmospheric structure
   - Veering/backing winds
   - Low level jets, etc.



Velocity can also be used to determine synoptic scale wind patterns such as veering or backing winds, and low level jets. In the image on the right, the S-shaped pattern to the velocity indicates that the winds are turning clockwise with height meaning the winds are veering with height. This type of information can be useful for forecasts and anticipating storm mode when combined with other tools.



Velocity can also be used to estimate velocity magnitudes for possible use in warnings, statements and forecasts. But, remember that it is radial velocity and not actual velocity. More on that in the next slide.



AWIPS-2 has a tool to help you estimate the actual velocity from a base velocity image, known as the estimated actual velocity or EAV tool. It is loaded from the tools menu in CAVE. Here is an example of the EAV tool on a 0.5 degree super-res velocity image of a bow echo. The goal is to orient the line along the direction you think the fastest winds are blowing. Keep in mind that errors of even 10-20 degrees from the true wind direction can result in large errors in the estimated ground relative wind speeds. In this example the gust front was moving from about 225 deg, resulting in a crossing angle of 40 degrees at the northeast tip of the EAV tool. The radial velocity is 70 knots outbound, and the EAV tool estimates 83 knots winds.

Base Velocity (Applications)

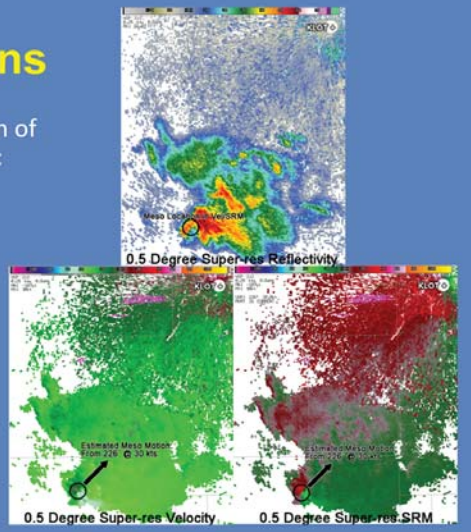
4. Update/adjust hodographs



Finally, velocity information can be used to update and adjust hodographs. This is especially useful since RAOB information is only updated twice, maybe 3 times, per day and radar updates every 5-10 minutes. So in fast evolving cases, this type of updating can be very helpful.

SRM Applications

1. Improved detection of velocity signatures:
   - TVSs
   - Mesocyclones
   - Microbursts
   - Boundaries



Now, on to some SRM applications or benefits, to wrap-up this section. The concept behind SRM is to help improve detection of certain velocity signatures. For example, look at the image on the right. In the lower right is SRM and in the lower left is velocity. Notice how the rotational signature shows up in both products, but it is easier to see in SRM. This is the beauty of SRM; improved detection of velocity signatures like rotation.

SRM Applications

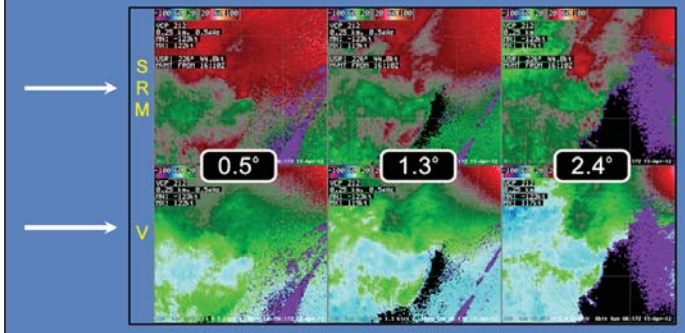
2. High SRM values (up to 248 kts) are displayable and viewable on cursor readout sampling



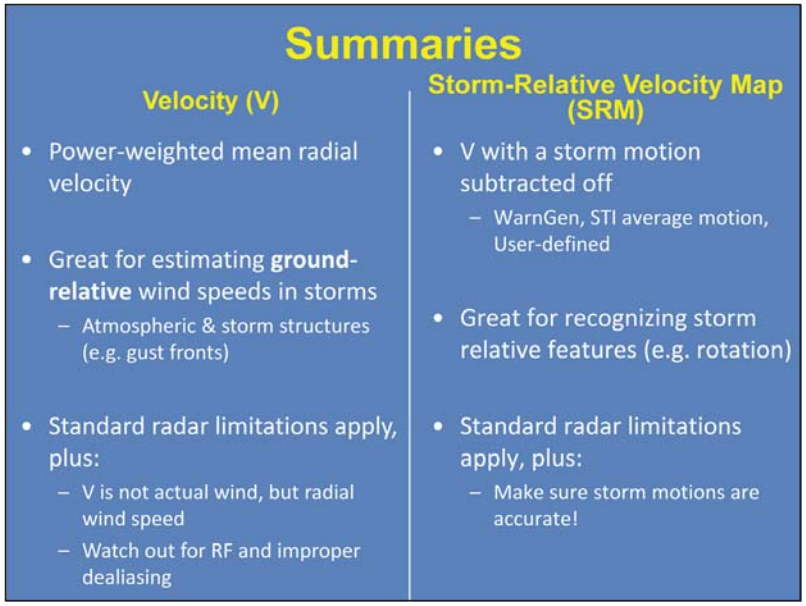
Using cursor readout, the values from the SRM product are displayable, so you can estimate various values such as rotational velocity or delta-V.

SRM Applications

3. Useful for examining the velocity structure of fast moving storms (>10 kts)



In fast moving storms, the velocity field can become saturated by the inbound or outbound colors, making it difficult to see certain signatures. In this example, the lower panel is V and the upper panel is SRM. Notice how both show a velocity couplet, but it is much easier to identify in SRM.

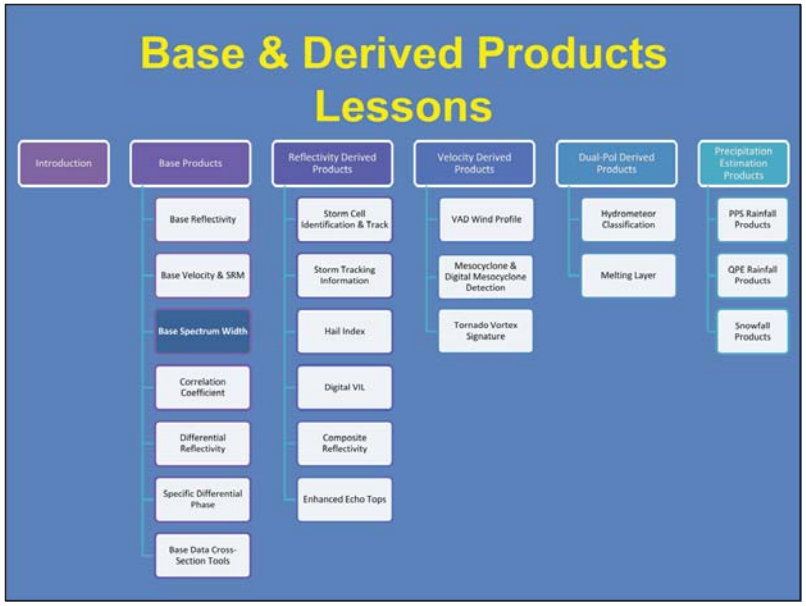


I know we kept bouncing back and forth between velocity and SRM but the purpose was to show how similar and how different they are from each other, at the same time. So, let's go ahead summarize this lesson. First, velocity is the power-weighted, mean radial velocity of all targets within the resolution volume. It is great for estimating ground-relative winds speeds in storms and determining atmospheric and storm-scale structures. Standard radar limitations apply, and remember that V is not actual wind, but radial wind speed – so just because you sample of pixel and get a number, like 42.7 kts – that does not necessarily mean that the true velocity or speed of the particulates in that pixel. Also, watch out for RF and improper dealiasing.

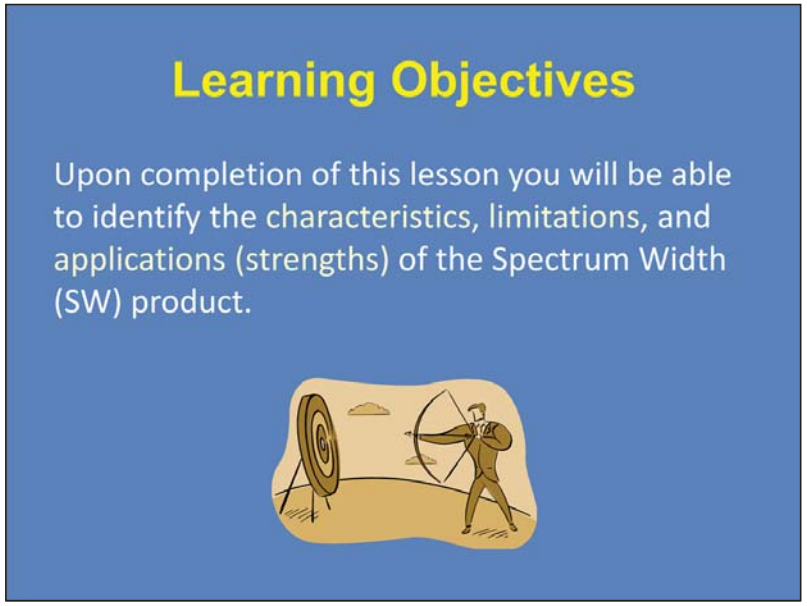
On to SRM, which is basically the velocity with storm motion subtracted off of it. This aids in recognizing storm-relative features much more easily. However, make sure your storm motions are accurate in addition to other limitations that apply, such as the angle relative to the radar and which you are sampling.



Welcome to the much underrated but no less important, Base Spectrum Width. Let's get started.



This is the third lesson for the series of Base Products, so you are progressing nicely. Head to the next slide, when you're ready.

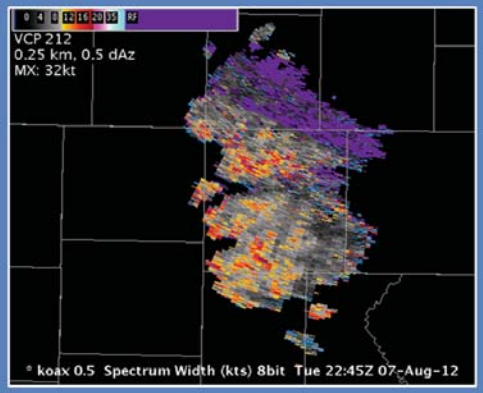


Here is the learning objective for this lesson. Please take a moment to review it, and advance the slide when you are ready.

Definition

Velocity dispersion within a resolution volume

- or…the amount of “chaos” occurring



Spectrum Width is defined as the velocity dispersion within a resolution volume. In other words, the product shows the variability of the individual target motions within a resolution volume. Another way to look at it, is spectrum is generally seen as the amount of chaos going on within a range bin. The more chaos or relative turbulence targets are experiencing within that area, the higher the value.

General Interpretation

- **Low SW** = low velocity dispersion (smooth flow)
  - Velocity estimate more accurate
- **High SW** = high velocity dispersion (chaotic motion)
  - Less trustworthy velocity estimate
- Meteorologically speaking:
  - SW high in regions of turbulence, wind shear, etc.
  - SW low in regions of quiescent precip (stratiform) and smooth flow (e.g. strong updrafts)

Low spectrum width implies low velocity dispersion, or smooth flow. The “mean radial velocity” values should be reliable in locations where SW is low. On the other hand, high SW implies high velocity dispersion, or chaotic motion. Forecasters can infer that velocity values associated with bins of higher SW values are less trustworthy. For meteorological purposes, SW is high in regions of turbulence, wind shear, etc. In regions of quiescent precipitation, like stratiform precipitation and smooth or near uniform flow, like in strong updrafts, SW is low.

Products Available

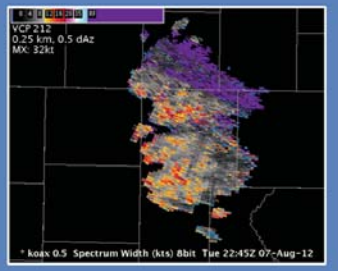
Product Name	Units	AWIPS ID	RPG ID	RPG Code	Data Levels	Range Res	Az Res	Max Range
Base Spectrum Width	kts	SW	SW	28	3-bit (8)	0.25 km	1 deg	60 km
Base Spectrum Width	kts	SW	SW	30	3-bit (8)	1 km	1 deg	230 km
Super Res Spectrum Width Data Array Product	*kts*	*SW*	*SDW*	*155*	*8-bit (256)*	*0.25 km*	*0.5 deg*	*300 km*

Only products highlighted in gray will be discussed in this lesson

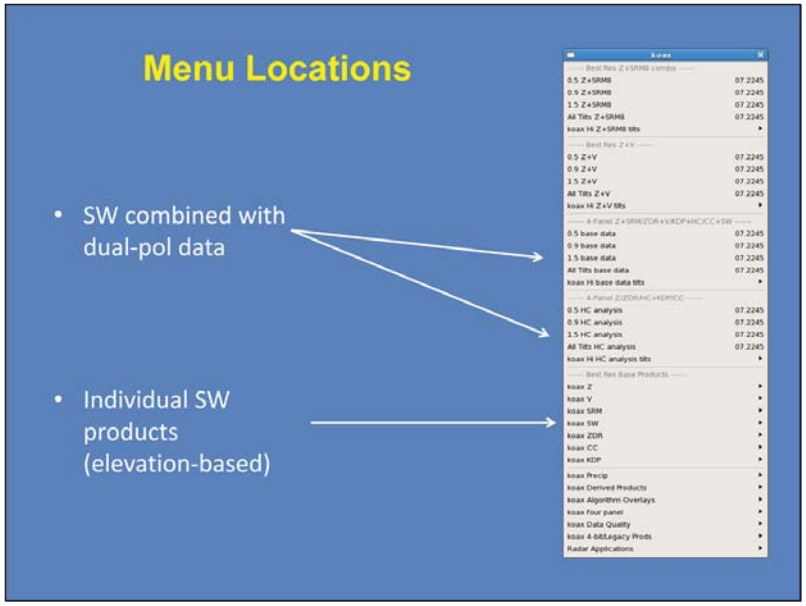
Here are the SW products available from the RPG. Note there is only one 8-bit, super-res product. That is all we will talk about in this lesson. The 3-bit products are available, but use these only if you are having bandwidth issues in your office.

General Product Characteristics

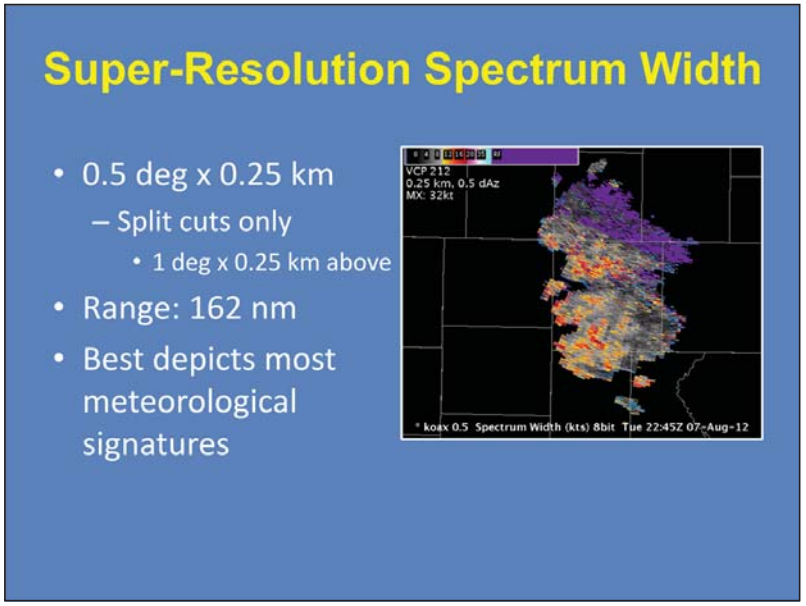
- Displayed as a polar coordinate with a 1° beamwidth
  - Super-res has 0.5° beamwidth
- Displayed relative to the RDA site
- Available for any elevation angle of the current VCP



The Spectrum Width product is displayed in polar coordinates with a 1 degree beamwidth except for the split cut elevations, where it has a 0.5 degree beam width. All data are displayed relative to the RDA site and are available for any elevation angle of the current VCP.



You can find Spectrum Width combined with other base data in the two 4-Panel areas highlighted in the dedicated radar's drop down menu. You can also find individual spectrum width elevation-based products in the Best Res Base Products section.

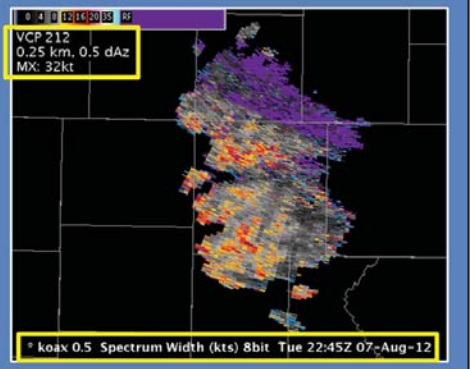


Super-res SW has 0.5 deg x 0.25 km resolution in the split cut elevations and 1 deg x 0.25 km resolution above that. The range resolution extends to 162 nm and is the preferred product to help identify most meteorological features.

Super-Resolution SW Characteristics

Product Annotation:
- VCP
- Range/Az Resolution
- Max SW

Product Legend:
- RPG ID
- Elevation angle
- Product Name
- Units
- Data Levels
- Date (in UTC)

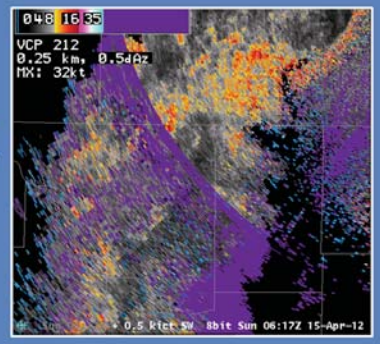


Here is the product annotation for SW. It shows you what VCP the radar is operating in, the range and azimuthal resolution, and the maximum spectrum width value anywhere in the product. Note, it does not show where in the product that maximum value occurred.

The product legend information can be found in the lower right. It has such information as radar ID, elevation angle, product name (i.e. Spectrum Width), the units of SW, the number of data levels and the date/time in UTC.

Spectrum Width Limitations

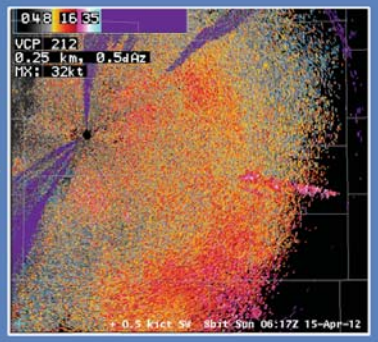
1. Range folding may obscure some data
   - Apply same solutions for V to mitigate RF in SW



So now, some limitations. Since Spectrum Width is just the measure of the velocity dispersion, range folding impacts Spectrum Width just as it does other velocity products. So apply the same solutions you learned previously for mitigating range folding to Spectrum Width.

Spectrum Width Limitations

2. System noise from weak power returns
   - Increases SW
   - Most notable in clear-air returns



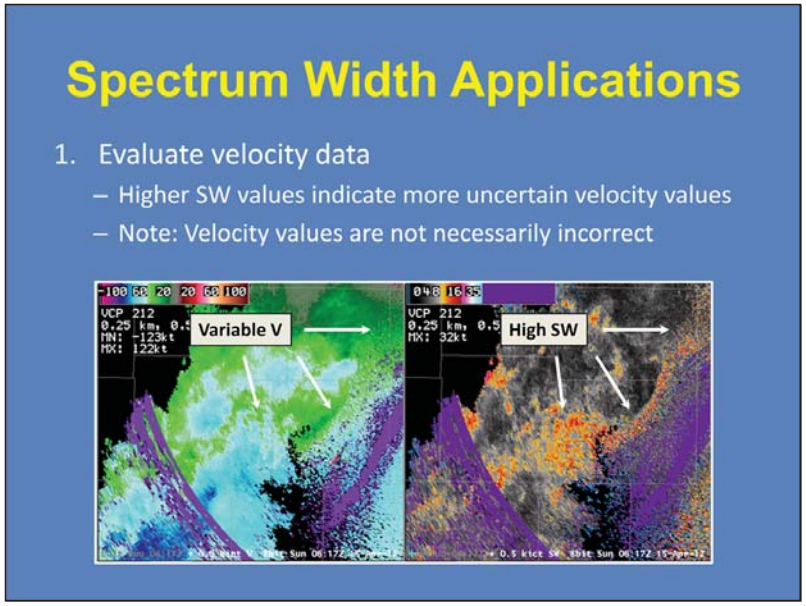
If the power return is very weak and the signal is dominated by system noise, then Spectrum Width values will increase substantially. The most likely region for this to occur is in clear-air.

Spectrum Width Limitations

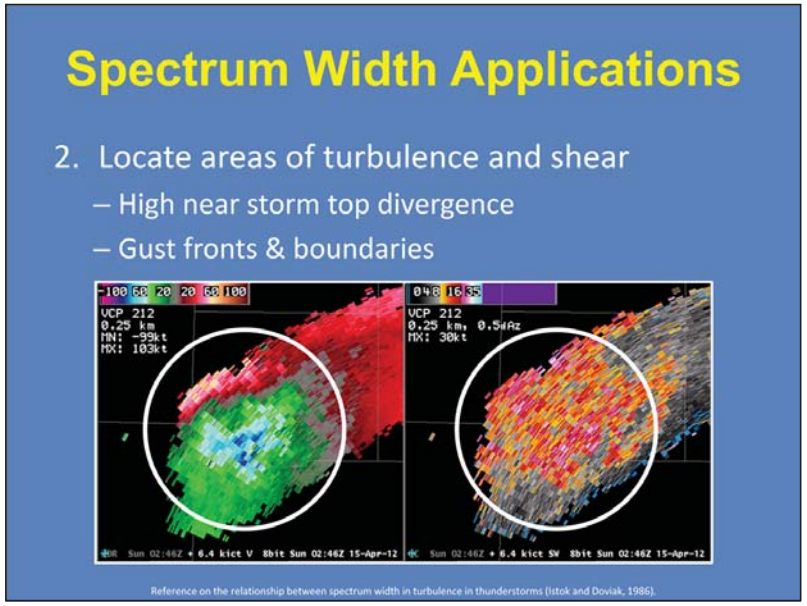
3. Long transmission times
   - Can bog down your narrowband



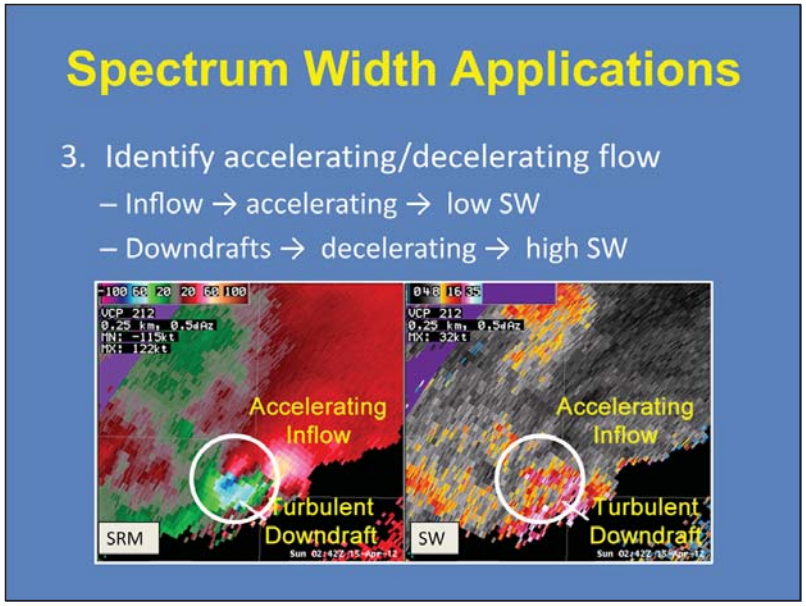
The 8-Bit Spectrum Width has a large file size and thus takes time to transmit over narrowband comms. Thus, it can bog down your system. If you are having issues with your bandwidth, limiting the amount of spectrum width data requested is an option.



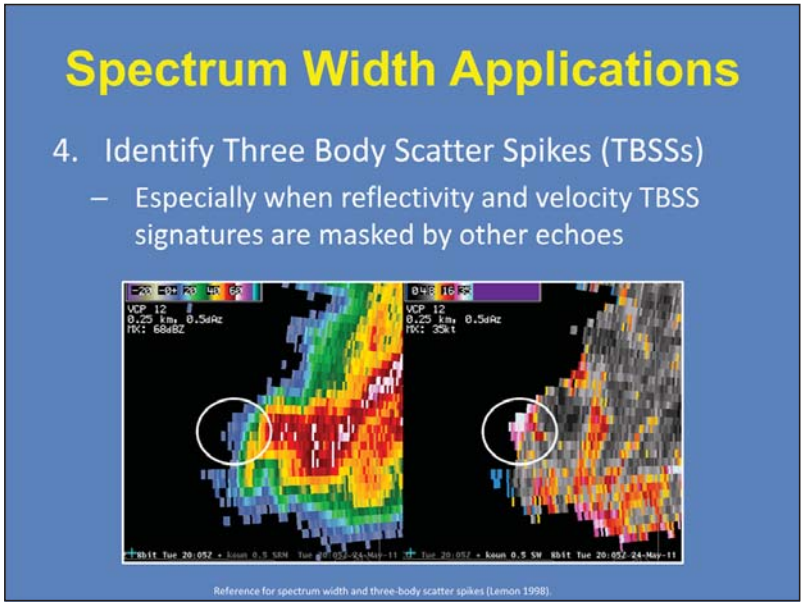
So now, on to some applications of spectrum width. Maybe by the end of this lesson, you'll be able to find some use for it, and not just use reflectivity and velocities for interrogation of structures. One application of Spectrum Width is its help in evaluating the quality of the velocity data. Broad spectrum width values can indicate a poor estimate of the mean radial velocity values due to turbulence, shear, or a low signal-to-noise ratio within the pulse volume. Recall, that a low signal-to-noise ratio can indicate that reflectivity values are probably too low to obtain good velocity estimates. However, if you check the corresponding reflectivity values and find them to be even modestly high (normally > 5 dBZ but dependent on range to the target) then poor SNR is probably not the reason for the broad spectrum width. Note that in this example the variability we see in the velocity field is consistent with the Spectrum Width product. Also, it is important to note here that high spectrum width does not imply incorrect velocity estimate, but rather that the estimate is less trustworthy.



Spectrum Width can be useful for identifying regions of intense turbulence. Some examples where turbulence and shear are high are storm-top divergence and gust fronts. An example of high spectrum width associated with storm-top divergence can be seen in the images shown. Reference on the relationship between spectrum width in turbulence in thunderstorms (Istok and Doviak, 1986).



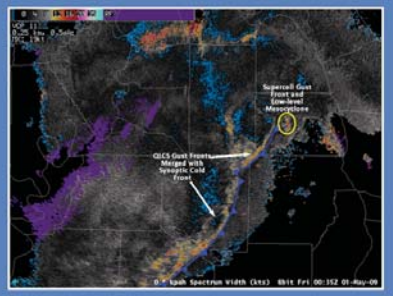
Spectrum Width can be used to locate areas of accelerating and decelerating flow. Low spectrum width values within the updraft are expected for intense updrafts or updrafts characterized by higher helicity (updraft correlated with high vertical vorticity) values. Thus, the low spectrum width values within the updraft are almost always seen with supercells or with the most intense convective storms. The correlation of low spectrum width with the more intense updrafts (especially in the lowest 1/3 of storm depth) indicates turbulence dampening when flow is undergoing acceleration, as you would expect within an intense updraft. Conversely, flow undergoing deceleration occurs within downdrafts descending near the ground. In this example, the very high spectrum width correlates very well with the presence of a strong mesocyclone with a downdraft at low levels, indicating very turbulent flow.



Spectrum Width can be useful in locating Three Body Scatter Spikes, a feature often associated with severe criteria hail. In fact, many times it is easier and faster (as compared to the use of reflectivity alone) to identify the TBSS by simply scanning two or three SW products that cut through storms in mid-levels where three-body scattering is often more apparent. In the example, the TBSS is not obvious in the reflectivity product, but stands out dramatically in the SW. Reference for spectrum width and three-body scatter spikes (Lemon 1998).

Spectrum Width Applications

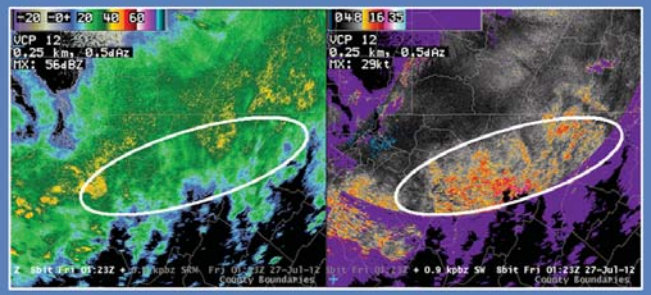
5. Identify important meteorological boundaries
   - Gust fronts
   - Synoptic fronts
   - Outflow boundaries



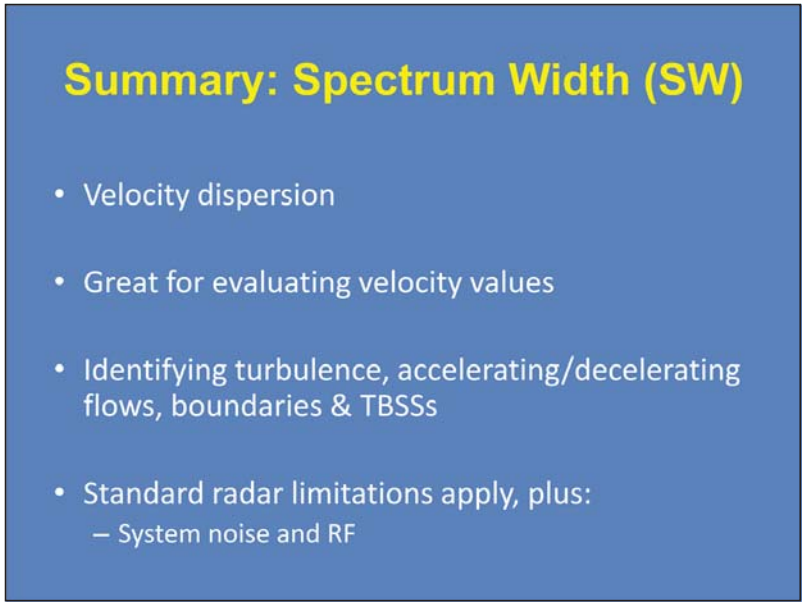
Because many meteorological boundaries create low-level turbulence, particularly boundaries that are moving rapidly, they show up nicely in Spectrum Width. Though Base Reflectivity and Velocity products are used to identify these types of boundaries, Spectrum Width can be used as well. In some cases, the boundary shows up best in Spectrum Width. This application should only be applied with the Super-Res Spectrum Width products available on the split cuts.

Spectrum Width Applications

6. Locate the melting layer
   - Differing fall velocities within layer result in broad spectrum widths



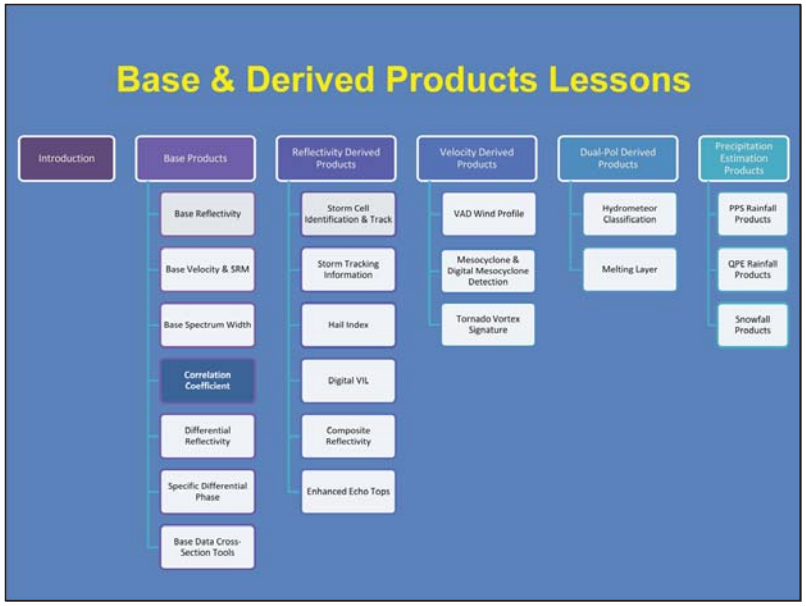
Within the melting layer there are complete ice-phase particles near the top and completely melted liquid-phase particles near the bottom. Liquid particles have faster fall speeds than ice particles, and the melting layer is usually thin compared to the width of the beam. These traits of the melting layer usually result in an observed velocity dispersion, which results in an increase in Spectrum Width values.



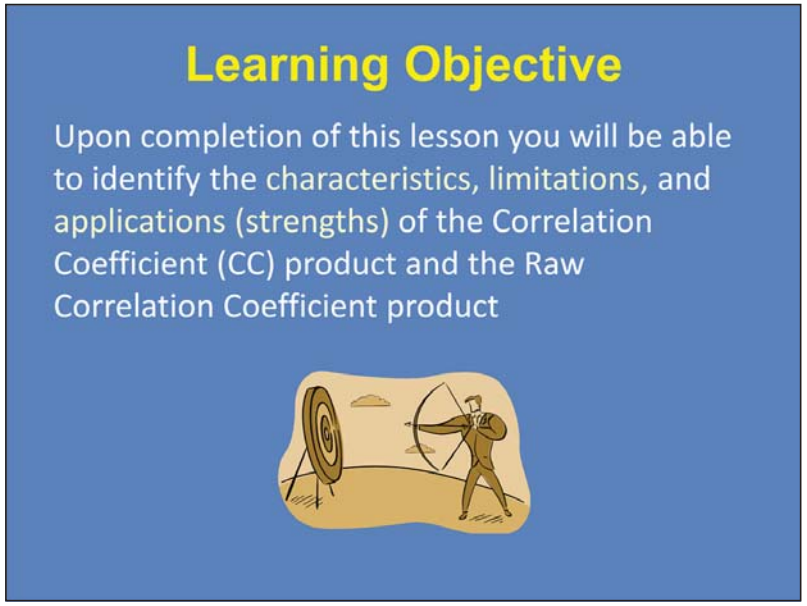
Summarizing this lesson, SW is a measure of the velocity dispersion within a resolution volume. It is great for evaluating velocity values and identifying turbulence, accelerating/decelerating flows, boundaries and TBSSs. Be careful of interpretation SW in regions of low SNR and range folding.



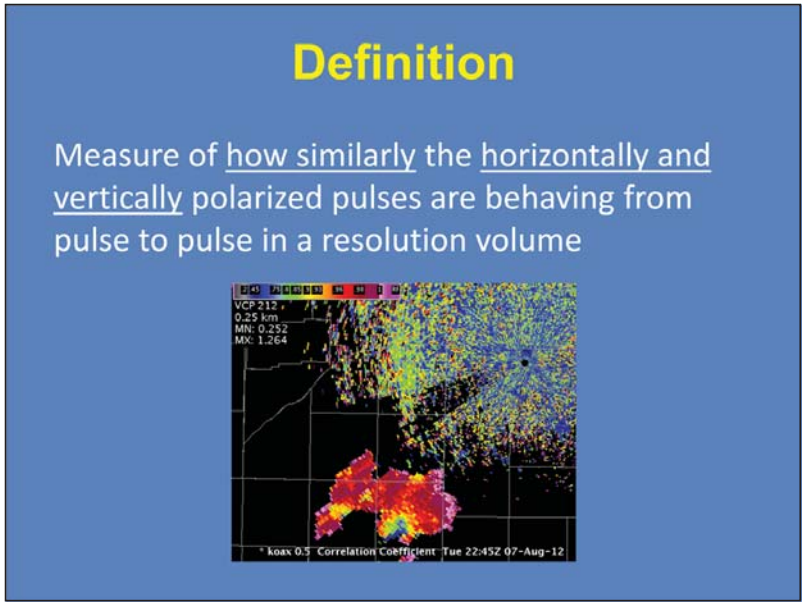
Welcome to Topic 4 Lesson 4, Correlation Coefficient. This is the first of the three dual-pol base products. So let's get started.



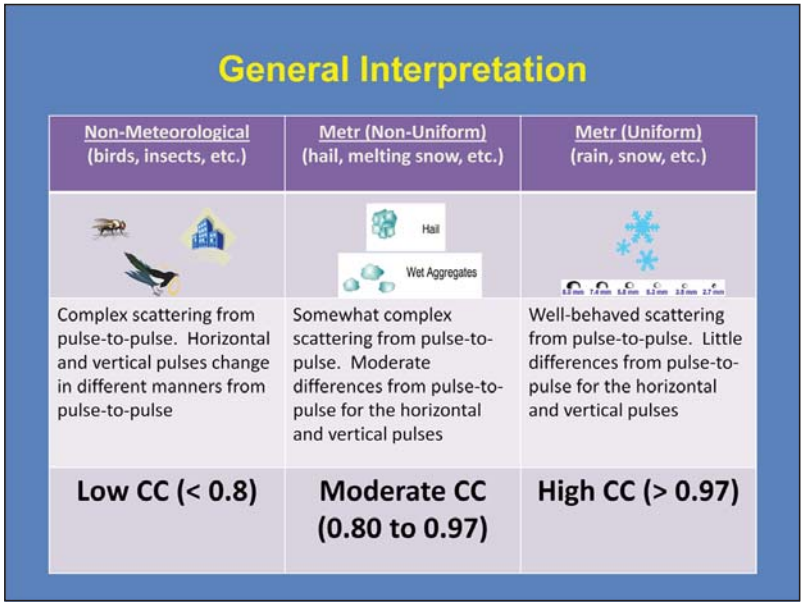
Here is the Topic 4 lessons roadmap, for reference. Advance to the next slide when you're ready.



Here is the learning objective for this lesson. Take a moment to read this slide and advance the slide when you are ready.

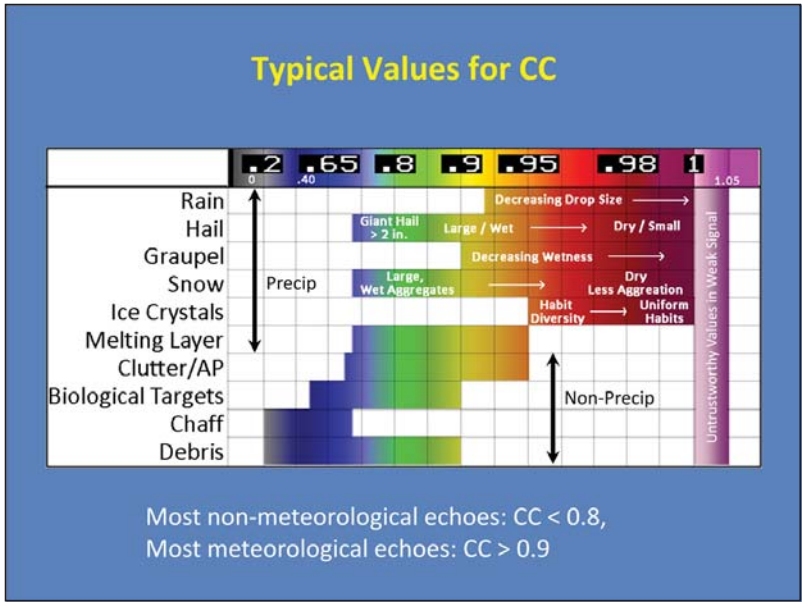


Correlation coefficient (CC) is defined as the measure of how similarly the horizontally and vertically polarized pulses are behaving from pulse to pulse within a resolution volume.



In terms of the Dual-Pol base products, CC will usually be your first stop on the journey to determining what you came for, which is to discriminating between non-meteorological and meteorological echoes, and p-types. The chart here illustrates this. For non-meteorological echoes such as birds, insects and ground clutter, the scattering of these objects can be quite complex. This causes the horizontal and vertical pulses to change in different manners from pulse-to-pulse resulting in CC typically below 0.8.

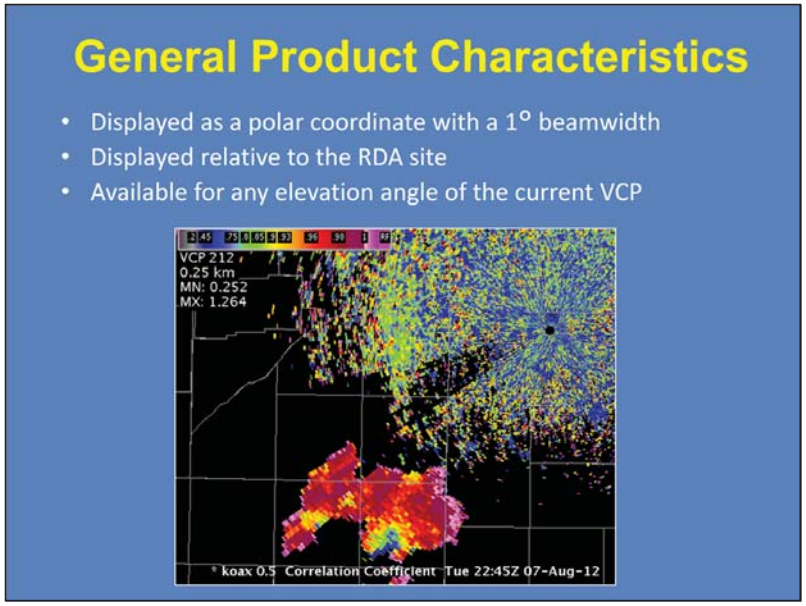
For meteorological echoes that are non-uniform in shape and size such as hail and melting snow, the scattering of these objects can be complex, but not nearly to the degree of complexity as the non-meteorological echoes. This causes the horizontal and vertical pulses to change in somewhat different manners from pulse-to-pulse typically resulting in CC between 0.8 and 0.97. Lastly, for meteorological echoes that are fairly uniform in shape and size such as rain and snow, the scattering off the objects is quite well-behaved. This causes the horizontal and vertical pulses to change in very similar manners from pulse-to-pulse resulting in CC greater than 0.97.



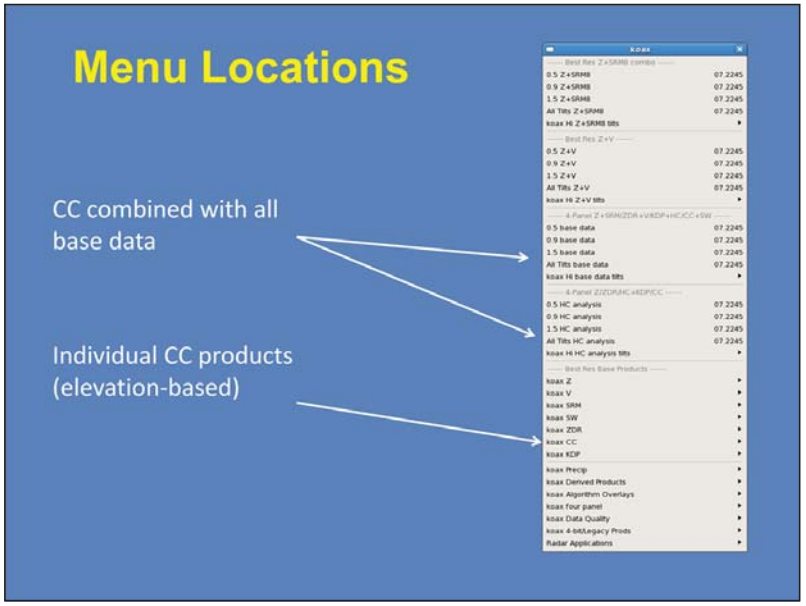
Now that we've seen a general idea of the CC values, here is a chart that dives into a little more detail. The thing to note with this chart is that there is pretty good break between CC values in non-meteorological echoes (bottom half of the chart), and meteorological echoes (top half of the chart). Basically, CC less than 0.8 is going to be non-meteorological, and CC greater than 0.9 is going to be meteorological.



We will discuss these two 8-bit CC products in this lesson, with the first one being the one you're used to, the processed CC data but now with 88D build 16.1 completed, we also have a super-resolution version of CC – what you're used to seeing in GR2 Analyst. There are some advantages to this product but also some key limitations, which we will discuss later in the lesson.



CC products are displayed in polar coordinates with a 1 degree beamwidth and data are displayed relative to the RDA site. CC is available for all elevation angles for the current VCP.

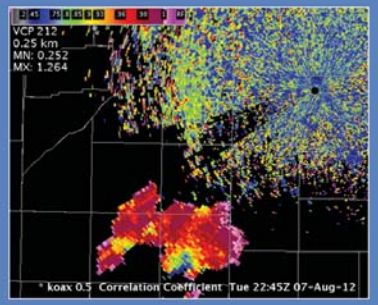


You can find the CC products combined with all base data in your dedicated radar drop-down menu in the first two options highlighted. And, the individual, elevation-based CC products can be found in the third option.

8-Bit Correlation Coefficient

- 1 deg x 0.25 km
- Range: 162 nm

Note: In the split cuts, there will be a difference between resolutions for Z, V, SRM and SW

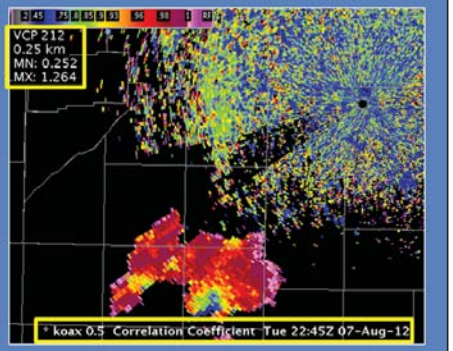


The 8-bit CC product is 1 deg x 0.25 km resolution and extends to a range of 162 nm. It is worth noting here when viewing products in the split cut elevations, toggling between CC and Z,V, SRM or SW there will be a difference in resolutions. This should not change your interpretation though.

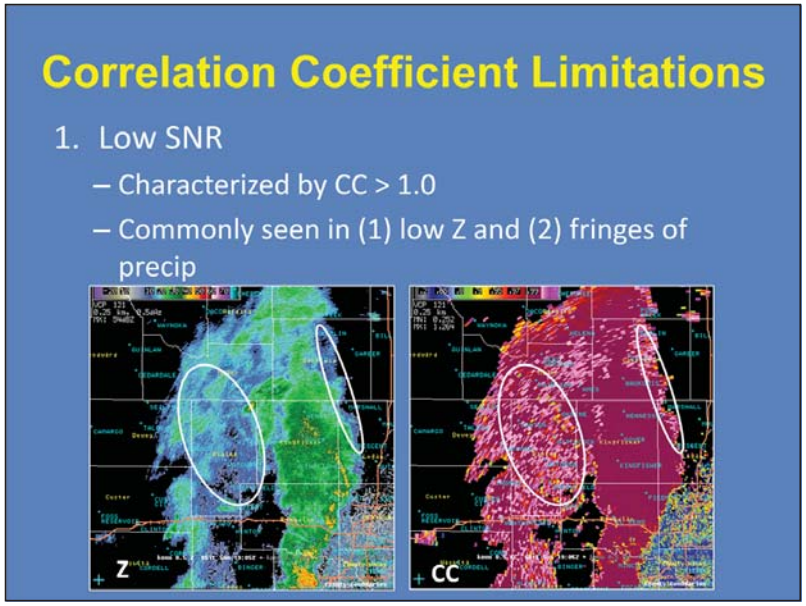
8-Bit CC Characteristics

Product Annotation:
- VCP
- Range Resolution
- Min CC
- Max CC

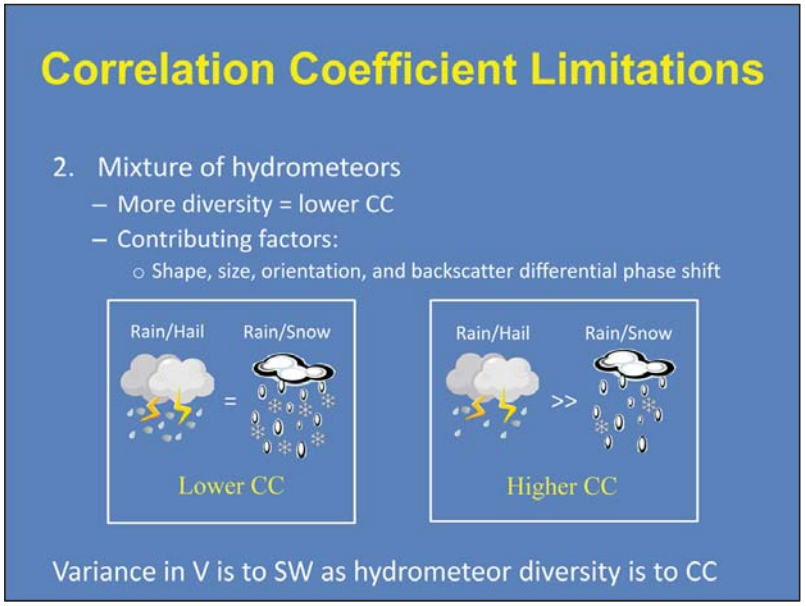
Product Legend:
- RPG ID
- Elevation angle
- Product Name
- Date/Time (in UTC)



In the upper left portion of the CC product is the product annotation. Here you can find information on the radar's current VCP mode, the range resolution and the min/max value anywhere in the CC field. Note, you will not know where in the CC product those values occur. In the lower right is the product legend information. There you can find the radar ID, elevation angle, product name (in this case CC) and the date/time in UTC.



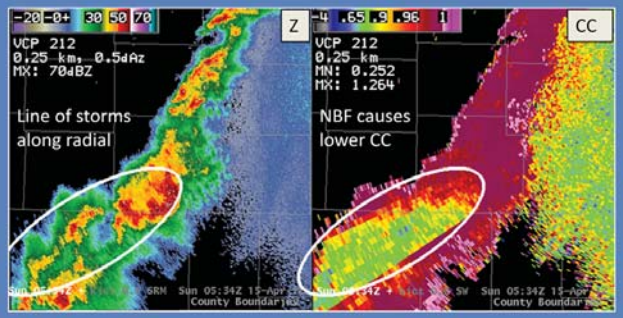
What are the limitations of CC? The first one we'll look at is regions of low signal-to-noise ratio (SNR). In the regions, CC becomes unreliable and can be noted by CC greater than 1.0 which shows up as a pink color in AWIPS-2. The most common regions of low SNR are in low Z regions (usually less than 20 dBZ), and along the fringes of precipitation.



When two or more types of hydrometeors are present within a pulse volume, the correlation coefficient will decrease. The amount of decrease in the correlation coefficient is dependent upon the relative contributions of each hydrometeor type present to the overall signal. The lowest correlation coefficient will occur when relatively equal contributions of each hydrometeor type to the signal are similar. For example, if snow and rain contribute relatively equally to the signal, the correlation coefficient will be lower than if the rain contributed more to the signal than did the snow. Another way to think about this is CC is like spectrum width is to velocity. The more random the velocities in a pulse volume, the higher the spectrum width just like the more random the particles in a pulse volume, the lower the CC.

Correlation Coefficient Limitations

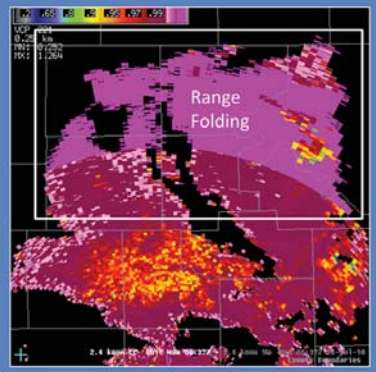
3. Non-uniform beam filling (NBF)
   – Radial "valleys" of lower CC



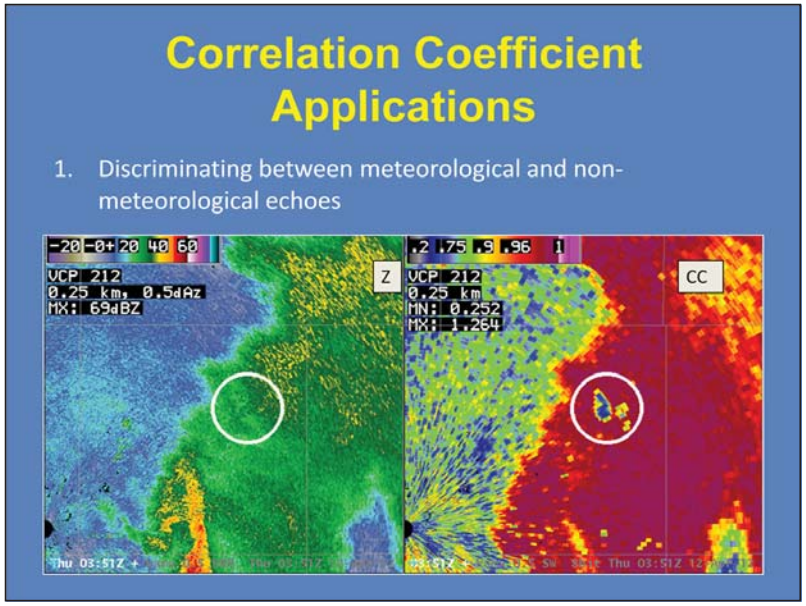
Non-uniform beam filling (NBF) can be detrimental to CC. NBF occurs when significant gradients of reflectivity and the polarimetric variables exist within the radar beam. This within-beam gradient can either be in the horizontal or the vertical. The result of these gradients is an increase in the bias of the polarimetric variables, especially CC. In CC, it is noted by a negative bias, so you will see a noticeable drop in CC in areas of NBF. The most common situations are when storms line up along a radial and areas downstream of a significant hail core. Here is an example where a line of storms became aligned along a radial. In the CC product, see how much the values decrease along these radials. This is directly caused by NBF. One last thing to note with NBF is that once significant gradients of reflectivity and the polarimetric variables occur within a beam, the effects will be seen the rest of the radial.

Correlation Coefficient Limitations

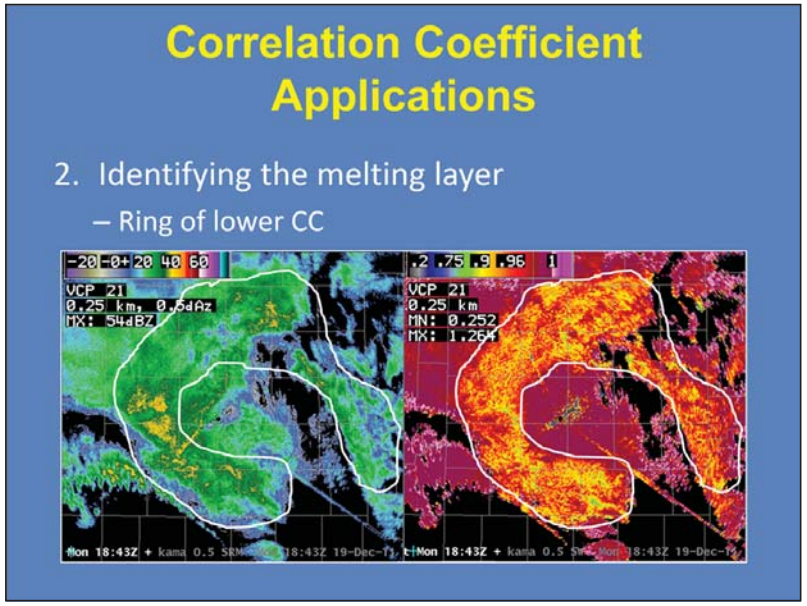
4. Range folding in the batch cuts
   – Between 1.65° & 6.5°



The last limitation with CC that we'll look at is in the batch cuts, range folding may obscure some signatures in CC. As a refresher, the batch cuts for the WSR-88D network are elevation angles between 1.65 degrees and 6.5 degrees. In these cuts each radial uses a series of alternating low and high PRF pulses and the dual-pol variables are only computed using the high PRF pulses which are subject to range folding. Here is an example of RF affecting CC.



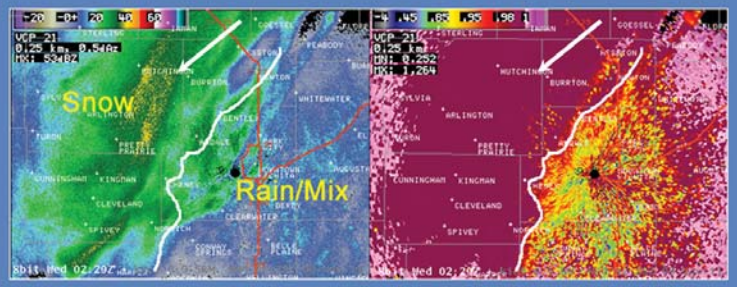
Now, on to the applications of CC. For one, it is best at discriminating between meteorological and non-meteorological echoes – which is why I mentioned that this will probably be – and should be – your first stop on the Dual Pol trail of products. Here is an example where in Z, it is very difficult, if not impossible, to see that the encircled region is mostly non-precipitation. However, switching over to CC, it becomes readily apparent by the very low values of CC, less than 0.5 that we have non-meteorological echoes.



The next biggest advantage of CC is identification of the melting layer. It almost always manifests itself as a ring of lower CC as long as there is sufficient precipitation present. In the example below, note how the melting layer does not show very clearly in Z, but in CC it stands out like a sore thumb.

Correlation Coefficient Applications

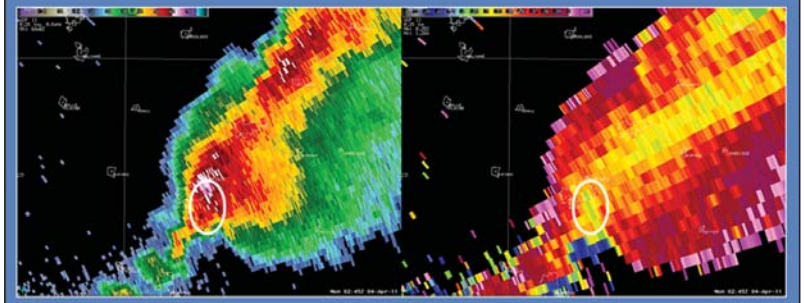
3. Identifying regions of rain vs. snow



Applying the melting layer signature concept to rain/snow transitions, CC is great at delineating rain/mix areas from all snow. Make sure you incorporate surface data into your analysis, but basically, the rain snow transition line is a manifestation of the melting layer near the ground. In this example, reflectivity shows what appears to be a bright band to the west, but when looking at CC we see that that is all snow (and based on surface temps below freezing). The true transition region is further east as seen in CC. So, without CC, the rain snow transition line would have been misplaced.

Correlation Coefficient Applications

4. Identify regions of significant severe hail
   - CC < 0.9
   - Diameters > 2 inches (i.e. golfballs)



When hail becomes significantly large, and we'll define this as hail larger than golf balls, CC begins to behave oddly. This is because at 2 inches or greater, Mie scattering becomes quite dominant which causes CC values to dip below 0.9 in hail. Therefore, when reflectivity is high, and you suspect large hail based on storm structure, if CC is less than 0.9, there is almost no doubt that hail larger than golf balls are present. Also, notice how the hail core messes with the CC values down those radials, which draws back to our discussion on limitations and something to keep in mind.

Correlation Coefficient Applications

5. Identify regions of irregularly shaped hydrometeors
   - Large, spiky hail

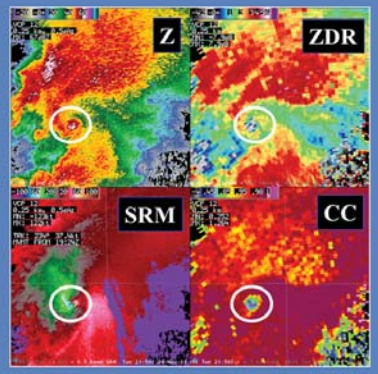


Hail like this stone can cause CC values of 0.75 or lower

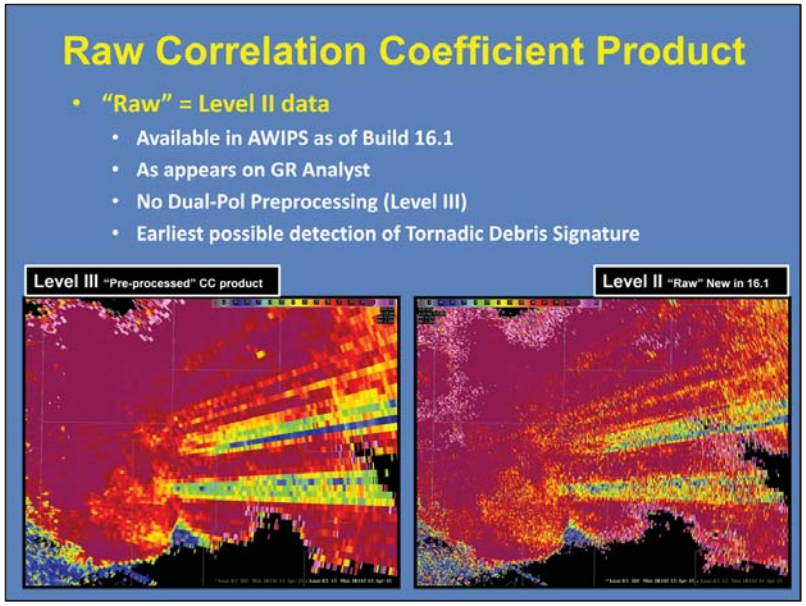
The same logic can be applied to spiky hail. Larger protuberances can cause CC to behave very oddly, reducing CC. There is not as much of a relationship here as with the previous slide. But, if CC are anomalously low, roughly less than 0.75, it is possible that irregularly shaped hail present. We can only imagine what CC would have been with this record hail stone in Nebraska!

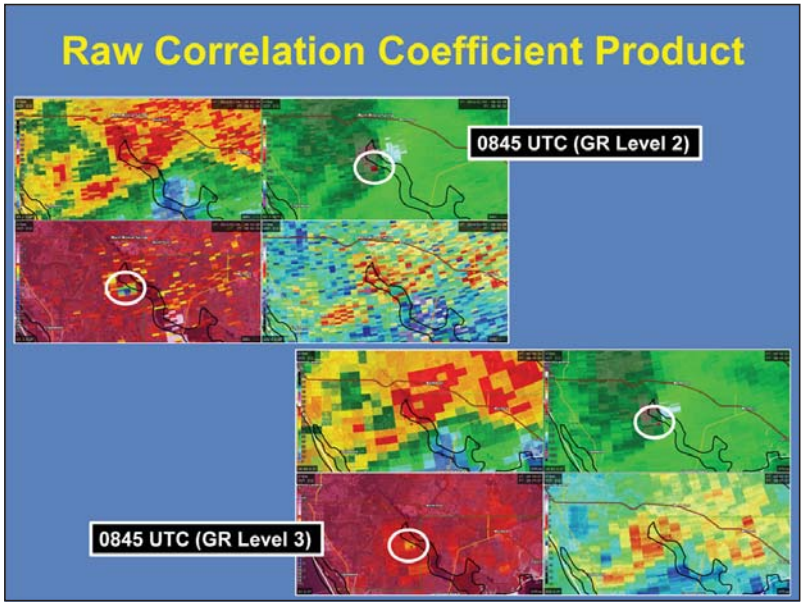
Correlation Coefficient Applications

6. Identify tornadic debris
   - Low CC (typically < 0.9)
   - Must be associated with velocity couplet!
   - Extent of damage unknown
   - Enhance wording in subsequent statements



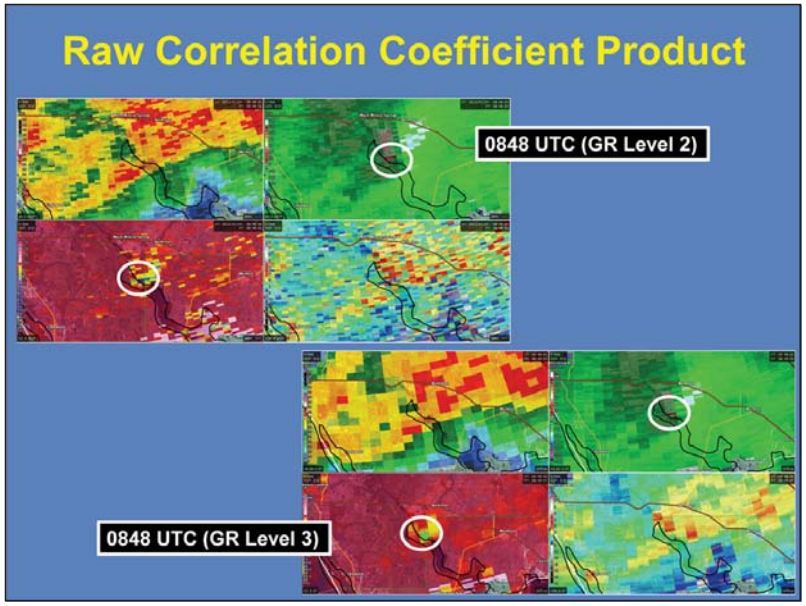
Probably one of the most well known and useful signatures with CC, the tornadic debris signature, or TDS. When you see a storm capable of producing a tornado by first looking at a combination of reflectivity, velocity, and SRM, CC can be used to help confirm the existence of a tornado. If you see CC's less than 0.9, co-located with the velocity couplet and within a region favorable for tornadoes, then you can confidently say that there is a tornado causing damage. The extent of the damage is unknown, although research suggests that this assessment may be possible but that is still to be determined. What you can do is use more confident wording in your subsequent warning statements to convey that a tornado is definitely occurring. You can also monitor trends in CC what stage in the potential life cycle of the tornado is occurring.



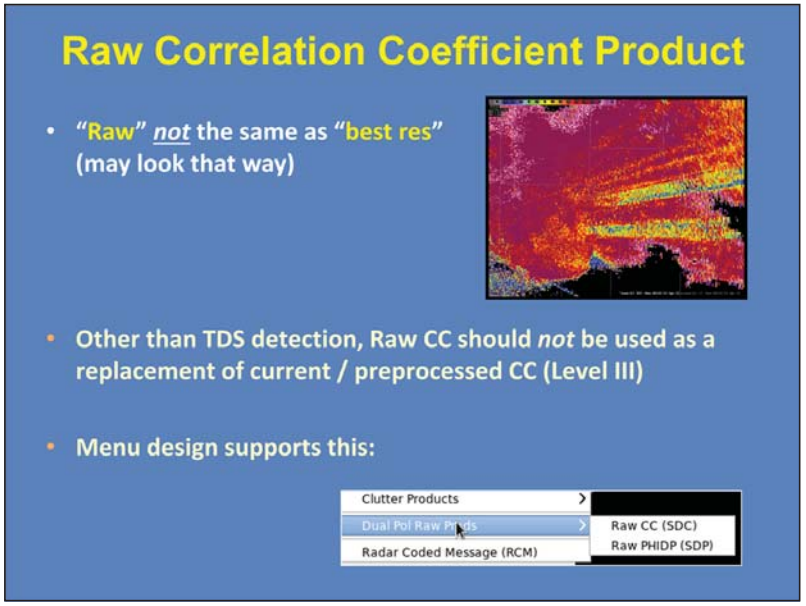


The second significant update with RDA/RPG Build 16.1 is the addition of the Raw Correlation Coefficient (CC) Product. This product has been added to RPG product suite at the request of numerous WFOs, particularly in the southeastern U.S. Comparisons of the Level II, Raw, CC data on GR Analyst with the Level III CC product on AWIPS during warning operations has shown that Tornadic Debris Signatures (TDS) are easier to see in the Level II data. This is of particular important where tornadoes can spin up along squall lines, and/or occur at night, conditions more likely in the southeastern U.S. Getting a Raw CC product into the AWIPS baseline supports the earliest possible detection of a TDS. The example images on the bottom, show the Level III data, on the left, is actually Pre-Processed and what you're used to seeing on AWIPS, while the Level II CC on the right, looks like it has finer detail, only because it is not Pre-Processed, which is why it is considered "Raw".

Here is an example at 08:45 UTC (definitely overnight), which is overnight and both of these example images are at the same time. The Level II four panel is on the upper left, with a corresponding Level III four panel on the lower right. Note that in both images there is a local minimum in CC collocated with the small velocity couplet. However, the CC minimum in the Level II image, compared to the minimum in the Level III image, is more visually apparent, and the minimum itself is lower. It looks more like a TDS in the Level II image than in the Level III image.



Stepping forward to the next volume scan, the visual difference between the Level II and the Level III CC images is still apparent. Switch back and forth between this slide and the previous one, if you want to get a sense of this progression – then move to the next slide, when you're ready.



It is important to remember that "Raw" for CC does ***not*** mean the same thing as best res or Super-Res for a product like reflectivity. The legacy base data (Z, V, and SW) are less noisy than the Dual Pol base data, and can be interpreted with 0.5° azimuthal resolution. Recall that the Dual Pol base data, which arrives from the RDA at 0.5° azimuthal resolution, are Preprocessed at the RPG before any of the Dual Pol products are generated. Preprocessing recombines the 0.5° azimuth to 1.0° azimuth, and smooths the Dual Pol base data. In some cases, details in a small TDS can be less apparent due to the Preprocessing algorithm.

The Raw CC is ***just*** the CC base data from the RDA generated as a product. Other than TDS detection, the Raw CC should not be your new "go-to" CC. For all other applications, the Level III, Preprocessed CC is still superior (think winter weather, hail, etc.). The design of menu supports this as well, keeping the Raw CC product in a menu stream separate from the Level III CC product.

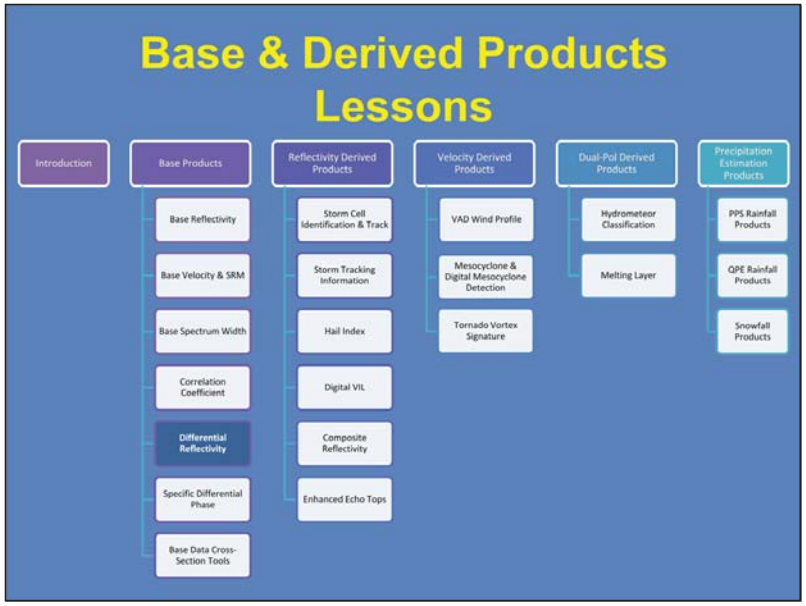
Summary: Correlation Coefficient (CC) & Raw CC

- Similarity of horizontal and vertical pulses in a resolution volume
- Meteorological vs. Non-meteorological echoes
- Melting layer, rain/snow, hail and TDS
- Standard radar limitations apply, plus:
  - Low SNR, NBF and RF in the batch cuts

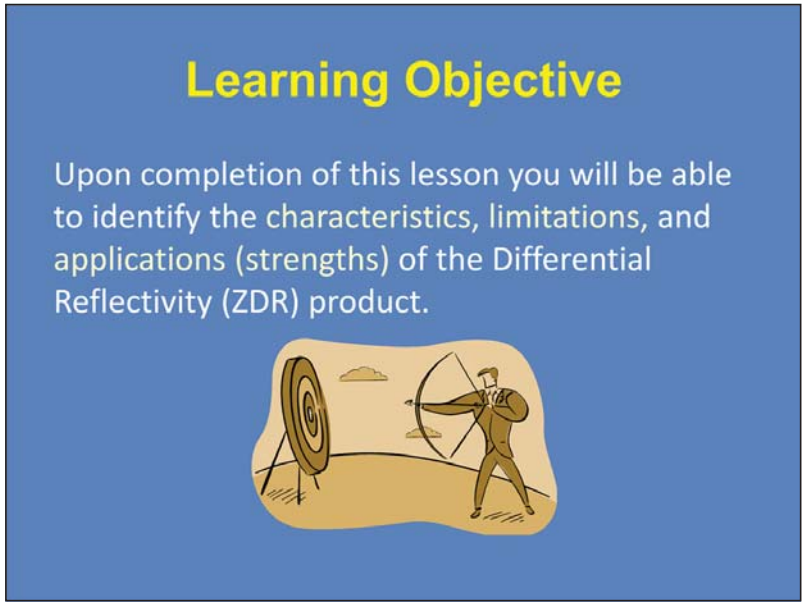
In summary, we have seen that CC is a measure of the similarity of the horizontal and vertical pulses in a resolution volume. It is best at discriminating meteorological and non-meteorological echoes and can definitively identify the melting layer, rain/snow transitions, large hail and tornadic debris. While standard limitations do apply, other considerations while looking at CC are low SNR, NBF and RF in the batch cuts.



Welcome to RAC Topic 4, on Differential Reflectivity (ZDR). Let's get started.



Now, we're on to the second of the Dual-Pol products lessons. Here is the roadmap for topic 4, for reference. Continue to the next slide when you're ready.



Here is the learning objective for this lesson. Take a moment to read through it and advance to the next the slide when you are ready.

Definition

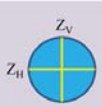
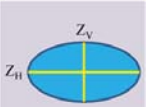
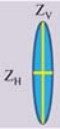
Difference between the horizontal and vertical reflectivity factors (in dBZ units)

$$ZDR = Z_H - Z_V$$

ZDR – or differential reflectivity is just that, it is literally the difference in reflectivity values of a particulate's horizontal and vertical components - in dBZ units, just like base reflectivity. The equation is seen here.

General Interpretation: Overall

Good indicator of dominant drop shape

Spherical (drizzle, small hail, etc.)	Horizontally Oriented (rain, melting hail, etc.)	Vertically Oriented (i.e. vertically oriented ice crystals)
		
$Z_H \sim Z_V$	$Z_H > Z_V$	$Z_H < Z_V$
$Z_H - Z_V \sim 0$	$Z_H - Z_V > 0$	$Z_H - Z_V < 0$
ZDR ~ 0 dB	ZDR > 0 dB	ZDR < 0 dB

In general, ZDR is great at indicating the dominant drop shape. For example, spherical objects will have near 0 dB because the horizontal and vertical reflectivity factors will be similar. For horizontally oriented objects, the horizontal reflectivity factor will be higher than the vertical reflectivity factor, therefore, ZDR will be positive. The opposite is true for ZDR with vertically oriented objects.

General Interpretation: Rain

- Strong relationship between raindrop diameter and shape
- Leads to strong relationship between raindrop size and ZDR

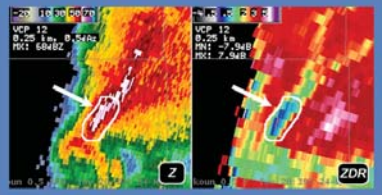
Major Axis Diameter (mm)	Image	ZDR (dB)
< 0.3 mm		~ 0.0 dB
1.35 mm		~ 1.3 dB
1.75 mm		~1.9 dB
2.65 mm		~2.8 dB
2.90 mm		~3.3 dB
3.68 mm		~4.1 dB
4.00 mm		~4.5 dB

Let's dive into ZDR a little deeper. First we'll look at rain. There is a strong relationship between raindrop diameter and shape. This leads to a strong relationship between raindrop size and ZDR. The table on the right illustrates this concept. Note how larger rain drops have larger ZDR values, keying in on the horizontal shape diagram from the previous slide, the positive ZDR values.

General Interpretation: Hail

- Tends to tumble
  - Local minimum
  - ZDR < 2 dB (with high Z)
- Small, melting hail
  - Appears as giant raindrop
  - High ZDR
    - As high as 5-6 dB





Hail, unlike rain, tends to tumble as it falls making it appear spherical to the radar. This causes a local minimum in ZDR values. The classic hail signature is where high Z and low ZDR co-exist. An example of this can be seen in the lower right. One caveat to this classic signature is small, melting hail. Small, melting hailstones can appear to the radar as large rain drops which can cause ZDR values to become extremely large (as high as 5-6 dB).

General Interpretation: Snow/Ice

Widely varying:

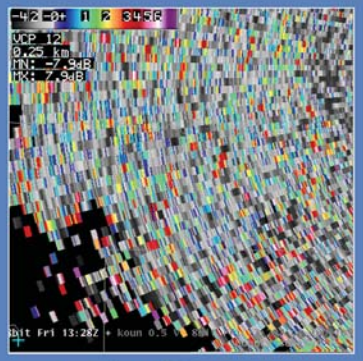
- Wet or dry?
- Low or high density?
- Preferred orientation?

Snow	ZDR
Dry / Aggregated	0.2 to 0.3 dB
Wet / Melting	2 to 3 dB
Ice	**ZDR**
Low-density / Random orientation	< 1 dB
High-density / Preferred Orientation (Horizontal)	As high as 4 to 5 dB
High-density / Preferred Orientation (Vertical)	- 2 to 0 dB

For snow and ice, interpreting ZDR values can be very tricky. You have to ask yourself a few questions. Is the snow wet or dry? Low or high density? Do the hydrometeors have any preferred orientation? All these factors play a role in the ZDR value that results. As a general guideline, the table on the right gives various expected values of ZDR for different ice crystal habits.

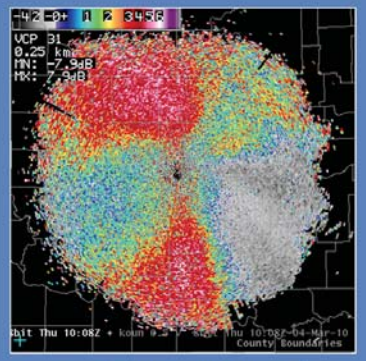
General Interpretation: Non-Met

Ground Clutter



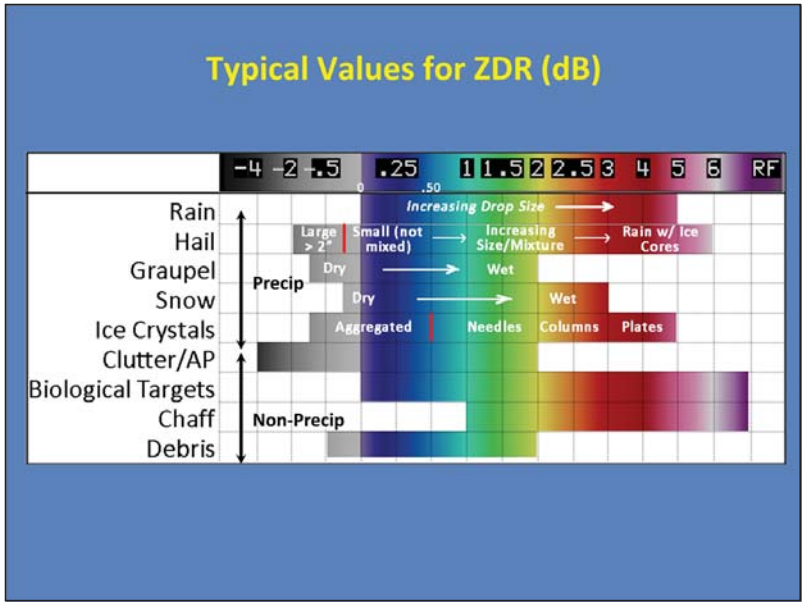
Lots of variability…very noisy

Biologicals



Less variability…orientation & flight direction

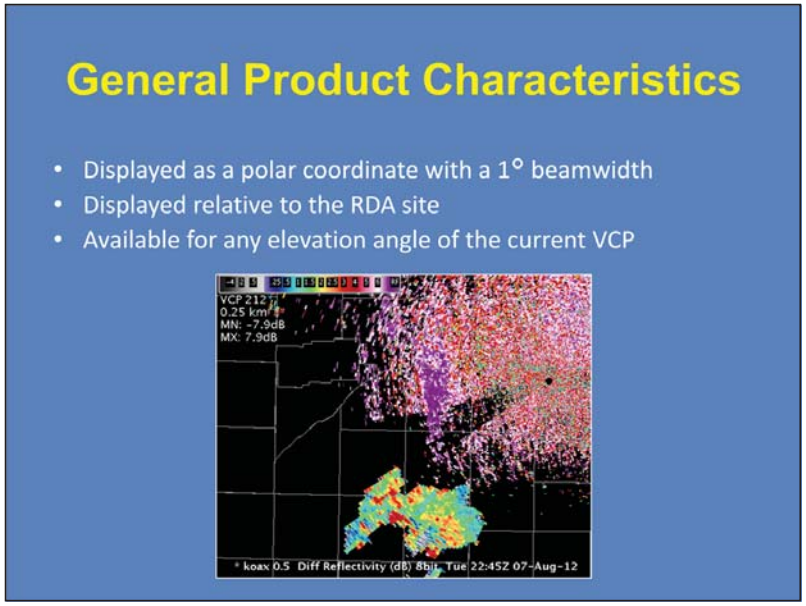
Non-meteorological echoes have more variability in shapes and sizes than do meteorological echoes, thus increasing the variability of expected ZDR values. The ZDR product is best at distinguishing non-meteorological echoes that are ground clutter (either normal or AP) and biological scatterers. For ground clutter, ZDR will typically appear very noisy and have both high and low values. Birds and insects, however, result in ZDR values that are more predictable and are dependent on the targets shape, orientation, and flight direction.



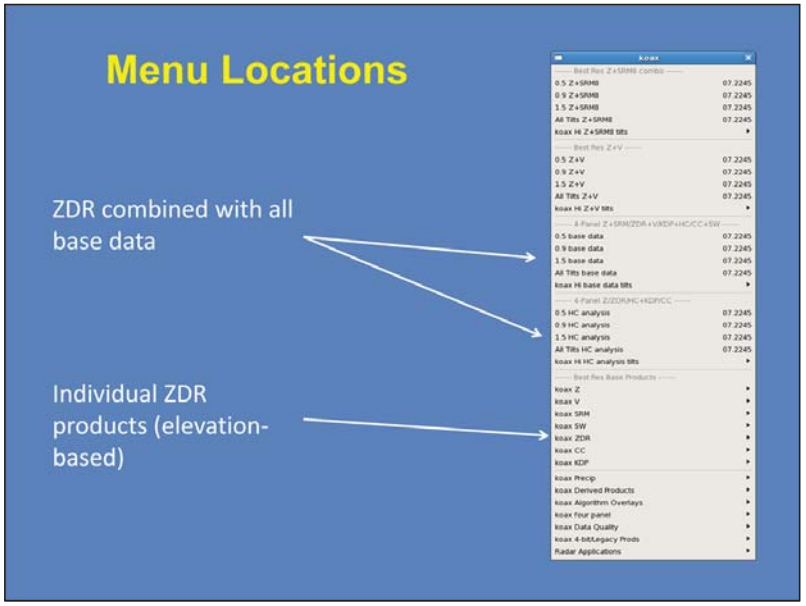
Here is a table to better illustrate and summarize a lot of the elements were just went over, showing the expected values of ZDR for various hydrometeor types. Notice how there is not as clear of break between non-meteorological and meteorological echoes as there was with CC.



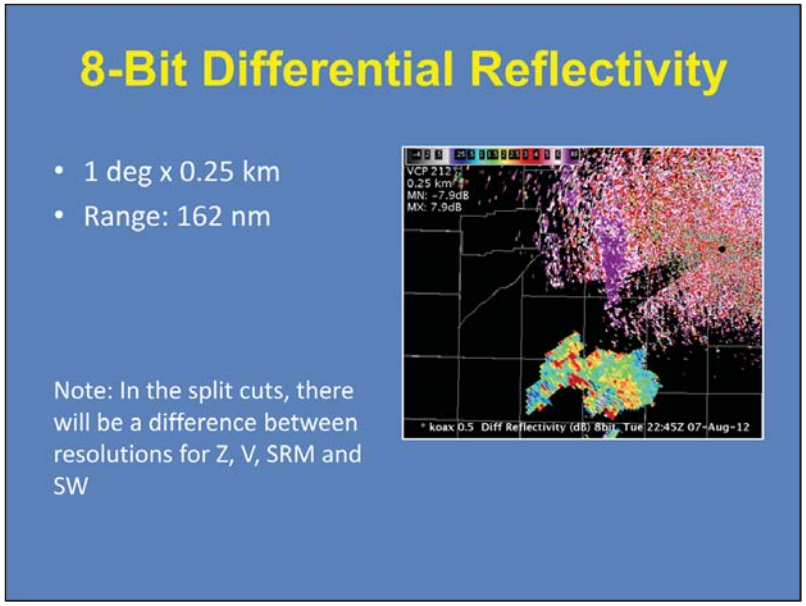
There is only one ZDR product that comes from the RPG, the 8-bit product ZDR and there is no super-res ZDR product available.

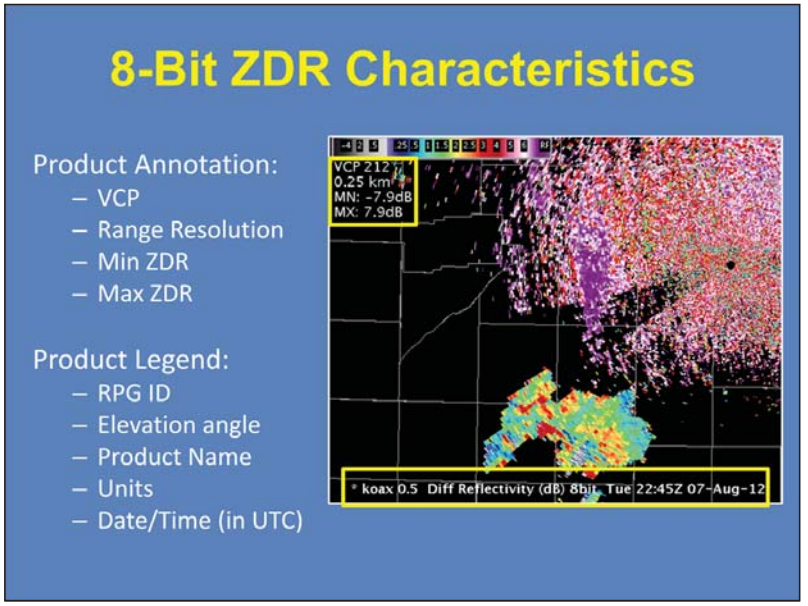


ZDR products are displayed in polar coordinates with a 1 degree beam width and data are displayed relative to the RDA site. ZDR is available for all elevation angles for the current VCP.



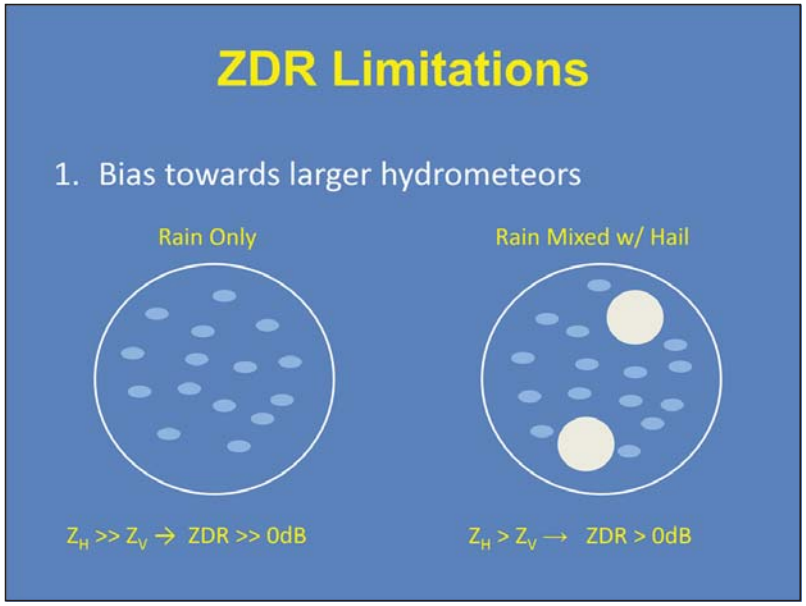
You can find the ZDR products combined with all base data in the two 4-Panel sections of your dedicated radar drop-down menu. There is also an individual, elevation-based ZDR product available in the Best Res Base Products section of the same menu.



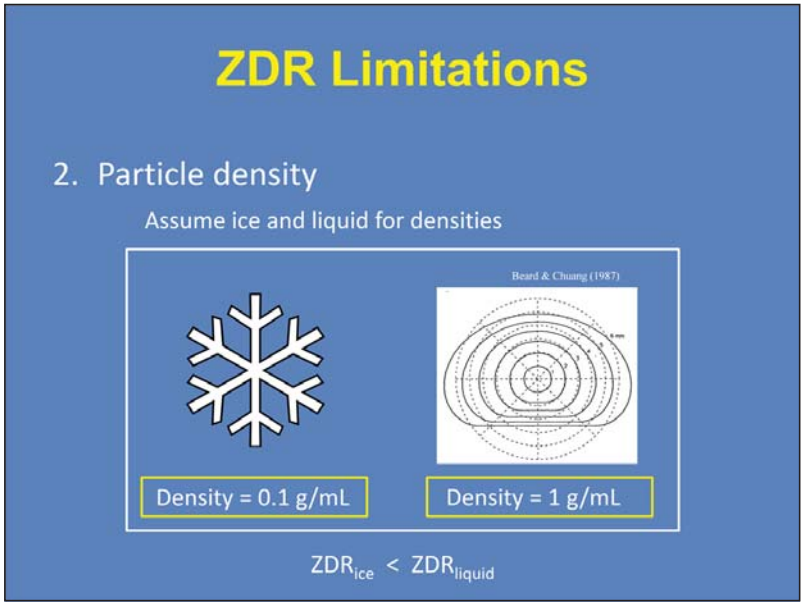


The 8-bit ZDR product is 1 deg x 0.25 km resolution and extends to a range of 162 nm. It is worth noting that while viewing all the base data products on the split cut elevations, you may notice a change in resolution when toggling between ZDR and non-dual-polarization products. While this resolution change may be distracting, it should not change your interpretation of ZDR features.

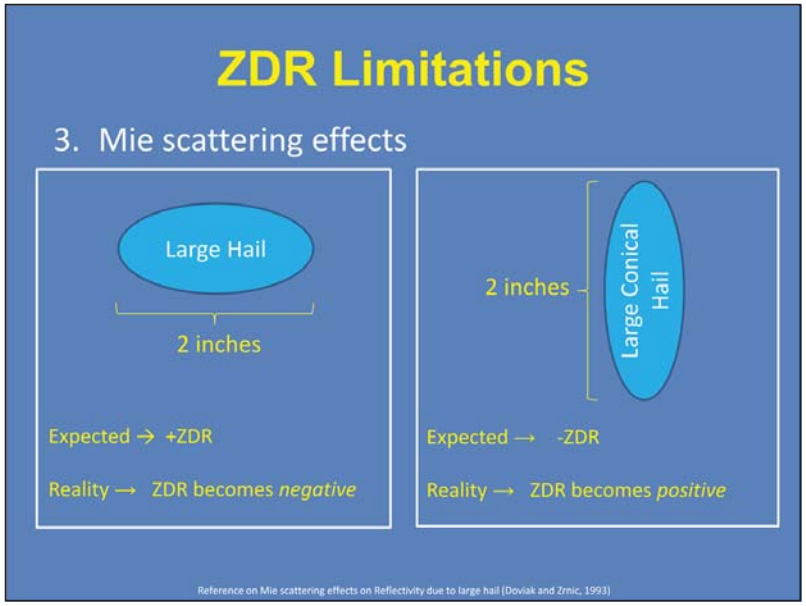
In the upper left portion of the ZDR display is the product annotation. Here you can find information on the radar's current VCP mode, the range resolution and the min/max value anywhere in the ZDR field – though it will not say where on the scope these values exist. In the lower right portion of the display is the product legend information. There you can find the radar ID, elevation angle, product name (in this case ZDR) and the date/time in UTC.

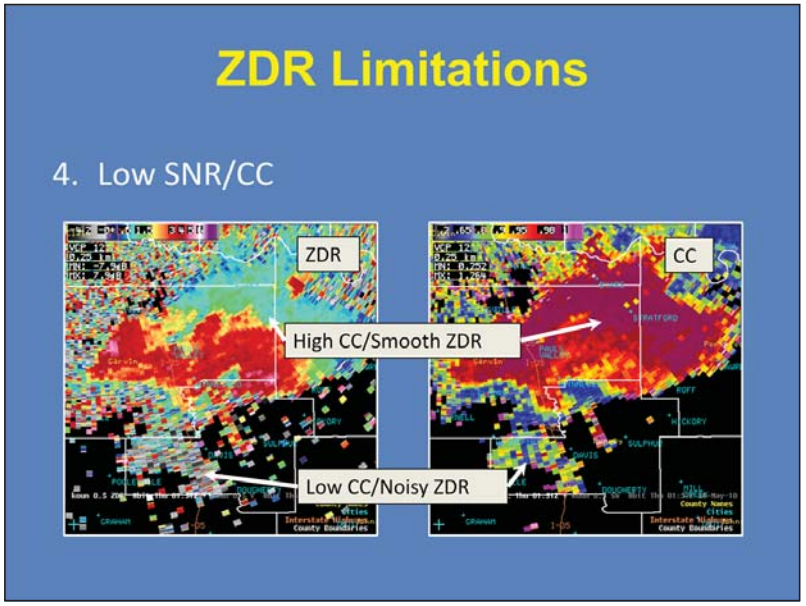


Now let's move on to ZDR's limitations. Since ZDR is just the difference between two reflectivity factors, it only makes sense that both Z and ZDR are biased towards larger targets. For example, in the image on the left, the radar samples pure rain. In this case, ZDR is going to be positive. The image on the right has exact same amount of rain and drop sizes, but with some hail mixed in the volume. The ZDR will be lower for this example because ZDR will be more influenced by the larger hailstones, which are characterized by lower ZDR values.



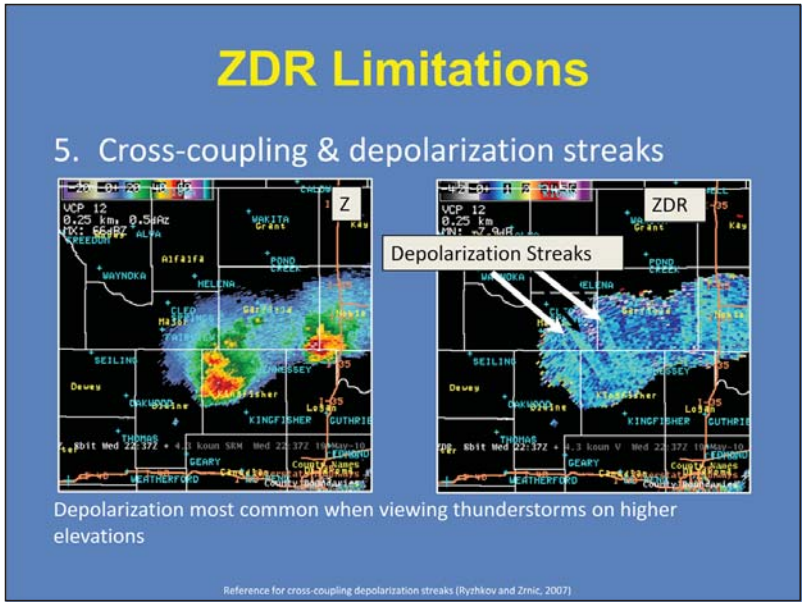
The next factor to consider when viewing ZDR is particle density. The lower the particle density, the lower the ZDR. For example, if you have one volume filled with ice (with a density of 0.1 g/mL) and another volume filled with raindrops of the same size and concentration (but a density of 1 g/mL), then the ZDR of the ice volume will be lower than the ZDR of the rain. The reason for this difference, despite the shapes being the same, is due to the difference in the dielectric constants of ice and liquid water. The higher the dielectric constant of a radar target, the more power that target returns to the radar, and the higher the reflectivity factor will be. This relationship is why reflectivity values for snow are often lower those for rain. NOTE: High density ice particles (such as needles) that are horizontally oriented can have significantly positive ZDR values even though Z values are generally low.





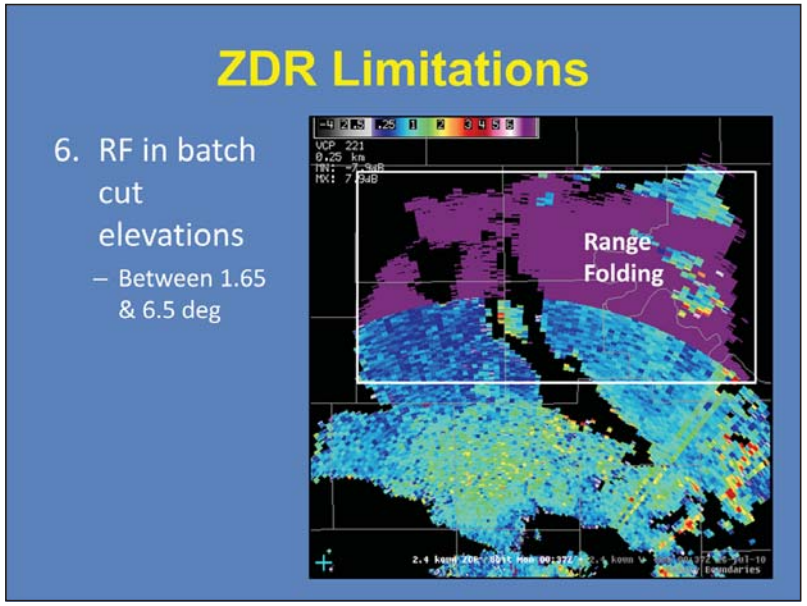
Recall from the previous lesson that Mie scattering can have some crazy effects on CC. The same issues can be true for ZDR, especially when targets in the radar volume are approximately 2 inches in diameter or larger. Research suggests that when targets are around 2 inches in diameter, Mie scattering effects can cause ZDR values to switch signs. In other words, for a horizontally oriented particle that is roughly 2 inches in diameter, you would expect ZDR values to be positive, but it may actually appear negative. The opposite would also be true for vertically oriented particles. Operationally, you will most likely see this impact with large hail. In regions of hail larger than golf balls, expect more negative ZDR values due to this Mie scattering effect. So, if you see high Z, low CC (< 0.9) and negative ZDR values, there is a good chance that golf ball size or larger hail is present. Reference on Mie scattering effects on Reflectivity due to large hail (Doviak and Zrnic, 1993).

When the correlation coefficient drops below 0.95, the error in ZDR becomes greater than 0.3 dB. When errors are this large, the ZDR product begins to appear noisy. The same result can happen in low signal-to-noise ratio, or SNR, regions because of similar drops in CC. In the example shown, ZDR is on the left and CC is on the right. Where the CCs are a maroon color, greater than 0.98, the ZDR field looks very smooth, but where CCs are yellow and blue in color, less than 0.95, the ZDR fields look much more noisy.

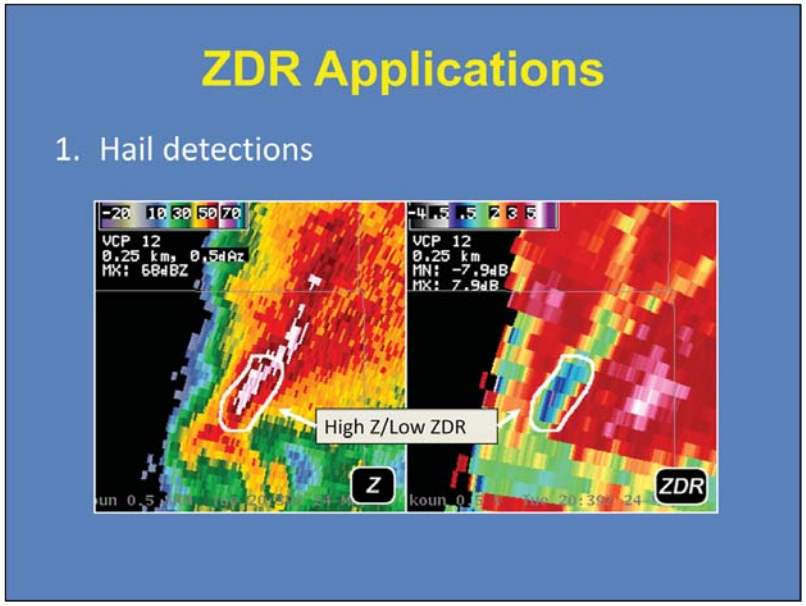


One of the issues that can occur with a radar that simultaneously transmits a horizontally and vertically polarized pulse is an effect called cross-coupling. Cross-coupling occurs when a portion of the transmitted power in one polarization is scattered (forwards and backwards) into the other polarization. Normally, cross-coupling is nominal. However, when a significant number of targets in the radar beam become canted (i.e., preferentially aligned) at an angle between 0 or 90 degrees, cross-coupling can significantly impact the quality of ZDR values down radial from where the cross-coupling began. The end result of this cross-coupling is a streak of enhanced positive or negative ZDR values called depolarization.

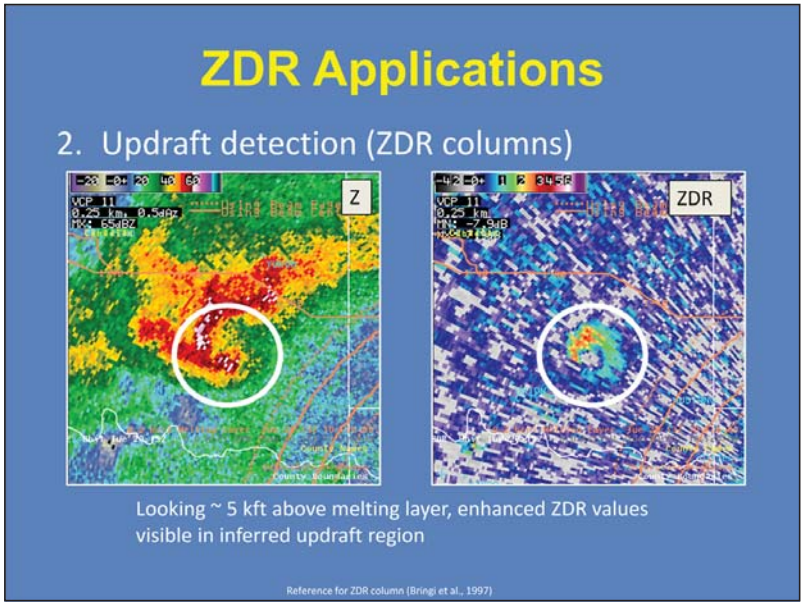
Depolarization is most common when ice crystals aloft become preferentially aligned in an electric field inside a thunderstorm. Depolarization streaks are best seen on higher elevation scans. Depolarization only affects ZDR and shows up as radial spikes of high and/or low ZDR based on the type of cross-coupling occurring in the region of depolarization. Reference for cross-coupling depolarization streaks (Ryzhkov and Zrnic, 2007).



The last limitation we will discuss with ZDR occurs in the batch cuts. On these tilts, range folding may obscure some signatures in ZDR. As a refresher, the batch cuts for the WSR-88D network are elevation angles between 1.65 degrees and 6.5 degrees. In these cuts each radial uses a series of alternating low and high PRF pulses and the dual-pol variables are only computed using the high PRF pulses. These pulses are subject to range folding. An example of RF affecting ZDR is shown.

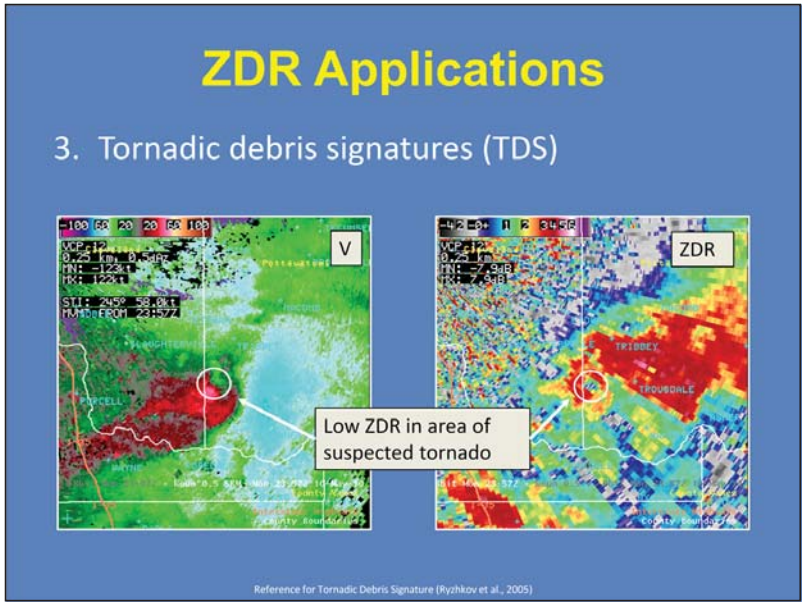


Now let's talk about some applications of ZDR. As mentioned on an earlier slide, hail detection is a major advantage with ZDR. In fact it was one of the first signatures noted with dual-pol data. The classic signature is characterized by high Z and near zero ZDR. An example of this signature is shown. It's important to note that small melting hail can cause ZDRs to be very large due to those hydrometeors appearing as giant rain drops to the radar.

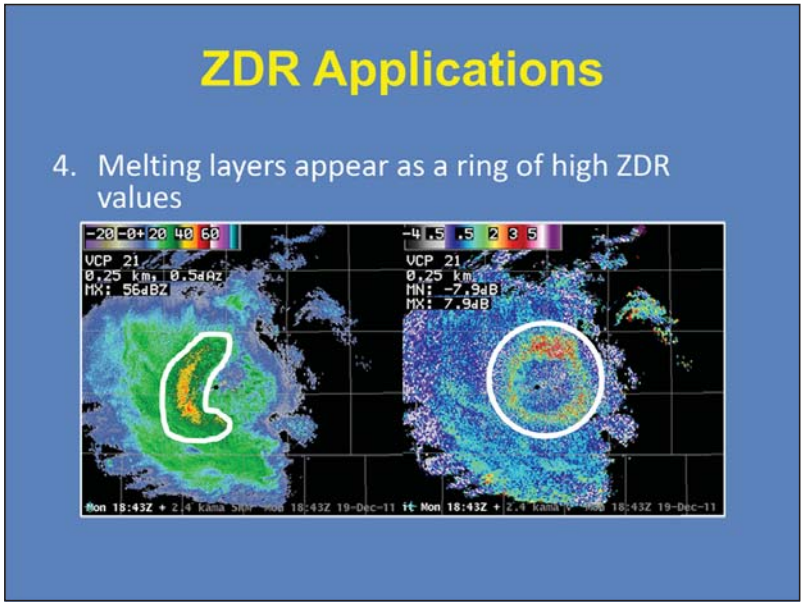


When intense updrafts develop, if enough liquid water is present within them, this liquid water will be lofted well above the environmental 0 degree Celsius level. This process will result in an area of locally enhanced ZDR within an updraft called a ZDR column. A good rule of thumb is look for ZDR > 1 dB above the melting layer height that is connected with higher ZDR below the melting layer. In other words, the column has to have vertical continuity.

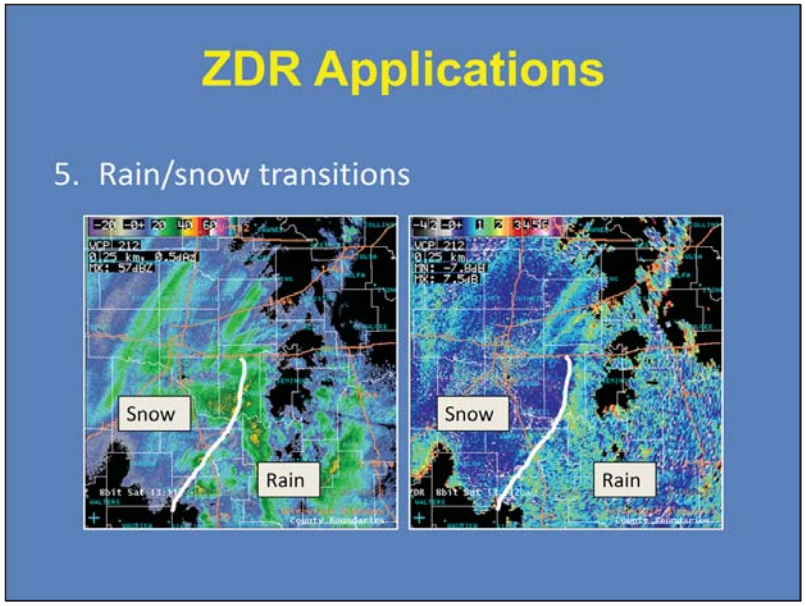
Here is an example of a ZDR column. The melting layer height in this example was roughly 10,500 feet and we are looking at 15,700 feet. Looking at reflectivity, we see a supercell with an inflow notch and hook echo. The location of the inflow notch should be roughly the location of the updraft. If we look at ZDR, we see a localized area of enhanced ZDR (> 2 dB) in the inferred updraft region. Due to the ZDR column, we can confidently say this region is where the storm updraft is located. Reference for ZDR column (Bringi et al., 1997).



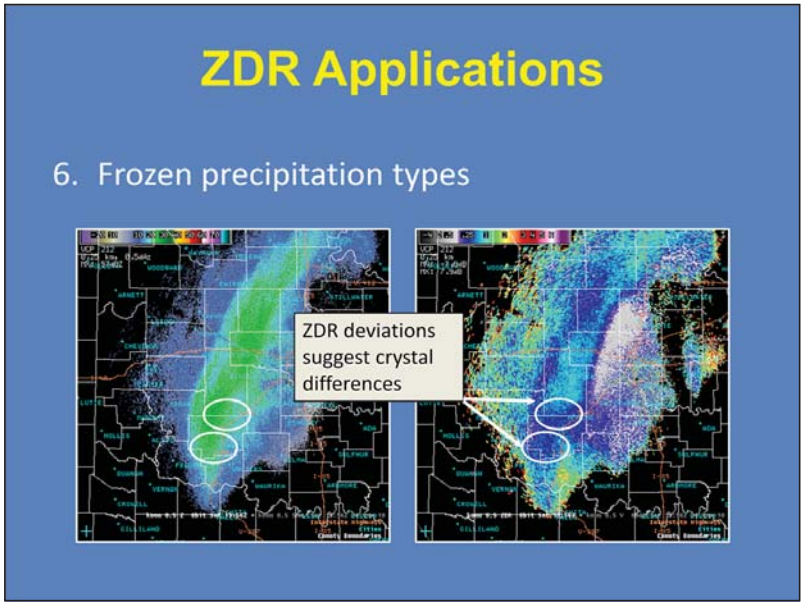
Like CC, ZDR can also be used to verify tornadic debris. Since tornadic debris has no preferred orientation, ZDR is going to be near 0 or slightly negative. So, when looking for tornadic debris in ZDR, first identify the tornadic signature using Z and SRM. Then toggle over to CC to verify it is debris. Then toggle to ZDR. If you see low, or slightly negative, ZDR then that gives you added confidence in a tornado causing damage. However, remember to look at Z, SRM, and CC before ZDR when identifying a TDS. Reference for Tornadic Debris Signature (Ryzhkov et al., 2005).



Within melting layers, ice becomes water coated and melts into a rain drop. During this transition, these particles essentially look like large rain drops. This causes ZDR to increase. Thus, the melting layer will manifest itself as a ring of high ZDR. Here is an example where you can see a partial bright band in Z, but the ring in ZDR is much more noticeable.



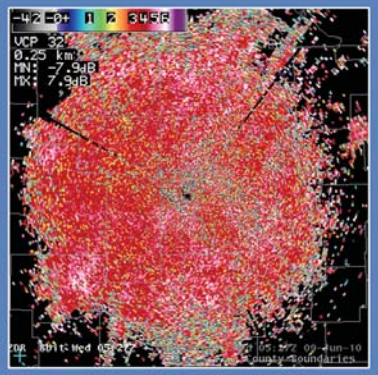
Along the same lines as melting layers, ZDR is great at delineating rain/snow transition regions. Basically, it is a manifestation of the melting layer near the ground. Again, always use surface observations to corroborate your evidence. Here is an example of ZDR delineating the rain/snow transition line (white line). To the east of the white line, ZDR values are generally higher indicating rain. To the west of the white line, ZDR are lower indicating snow. Surface observations confirmed this, but are not shown in the image.



Dual-pol is showing promise in delineating many types of precipitation. Here is an example where ZDR is clearly noting two different types of frozen precipitation. Despite Z being similar throughout the two white circles, ZDR in the circle to the north is slightly higher than ZDR to the circle to the south. This tells us something is different in these regions. Research is on-going to help determine operational applications for these differences. In the future, it may be possible to use ZDR to determine frozen precip crystal type and estimate snow-to-liquid ratios more accurately.

ZDR Applications

7. Non-meteorological echoes
   - Local variability exists
   - Knowledge of biological populations helps



Lastly, ZDR can help delineate different types of non-meteorological echoes, especially birds and insects. Migrating birds (and some insects like monarch butterflies) usually will appear with a distinct azimuthal pattern of higher and lower ZDR values based on their body shape and flight orientation. Other insects that do not migrate will generally have high ZDR values at all radials.

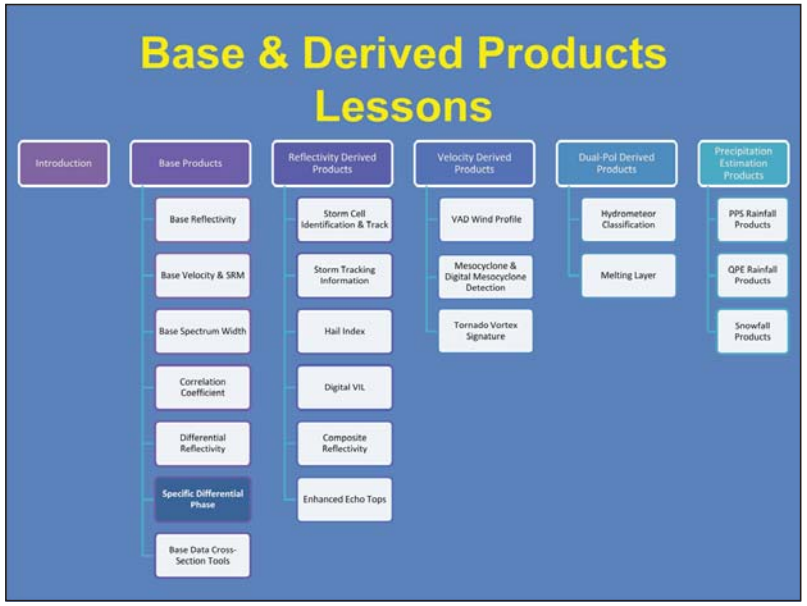
Summary: ZDR

- Difference between horizontal and vertical reflectivity factors
- Drop sizes in rain, hail detection & updrafts
- Melting layers, ice habits, rain/snow transitions
- Standard radar limitations apply:
  - Bias toward larger echo types, low SNR/CC, depolarization, RF in batch cuts

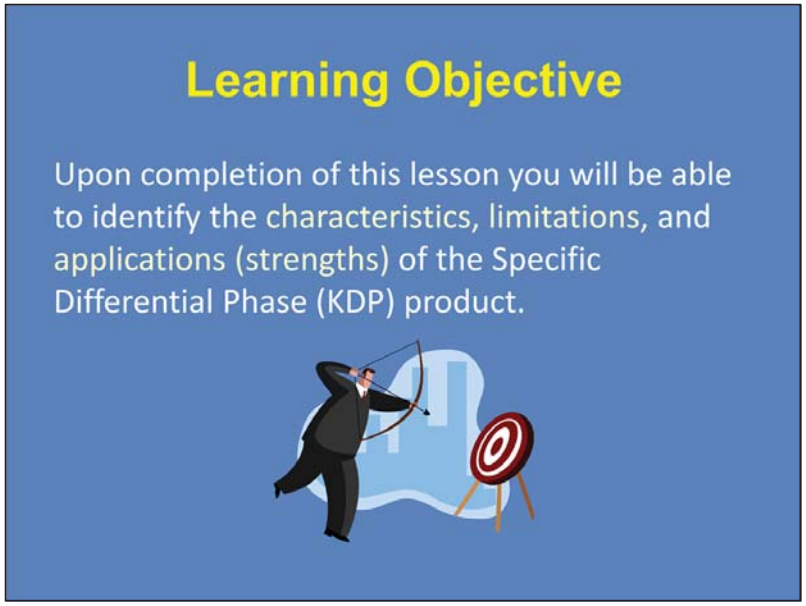
To summarize ZDR, it is the difference in the horizontal and vertical reflectivity factors. It is great for determining drop sizes in rain, detecting hail and updrafts, identifying melting layers, ice habits, and rain/snow transitions. However, ZDR is biased by larger hydrometeors, low SNR & CC, becomes “streaky” in depolarization regions and is masked by RF in the batch cuts.



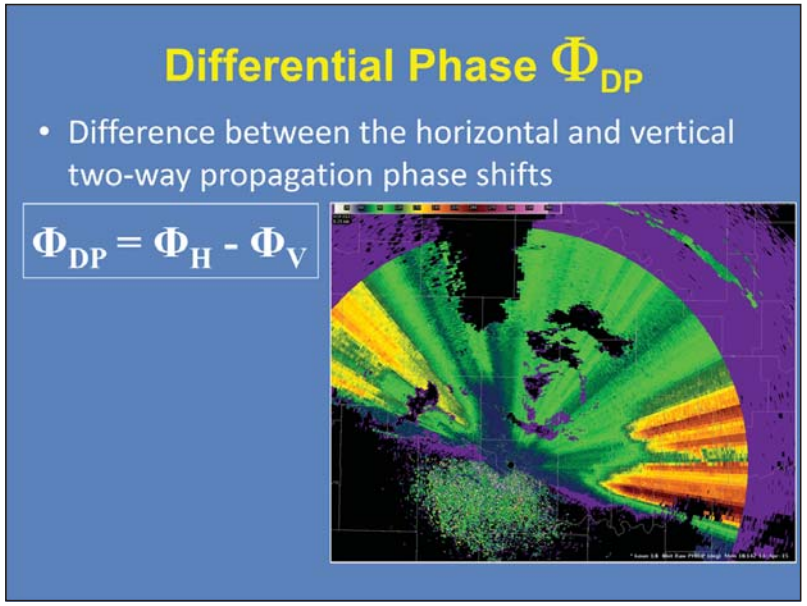
Welcome to Topic 4 and our lesson on Specific Differential Phase – or KDP, which is the third of the 3 dual-pol base products. Let's get started!



Almost finished with the base products lessons and we're at the end of the 3 dual-pol products, as mentioned. Head on over to the next slide when you're ready.

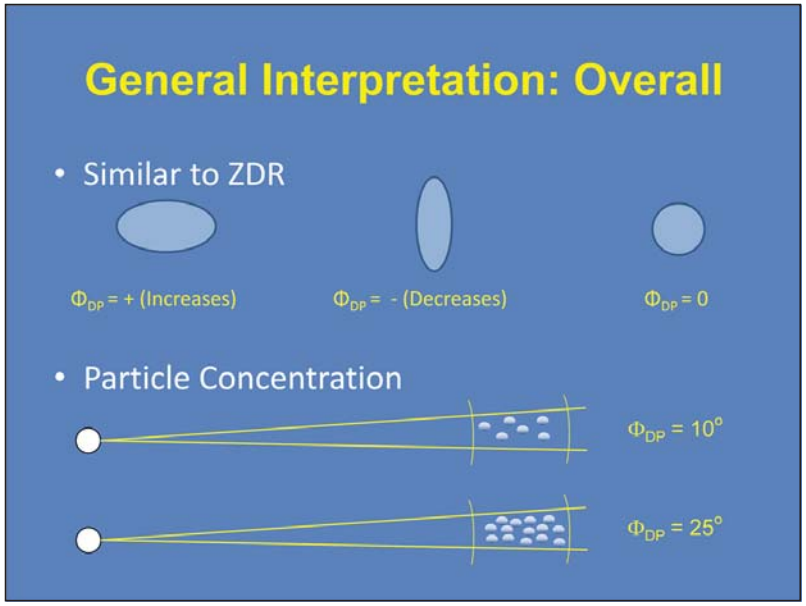


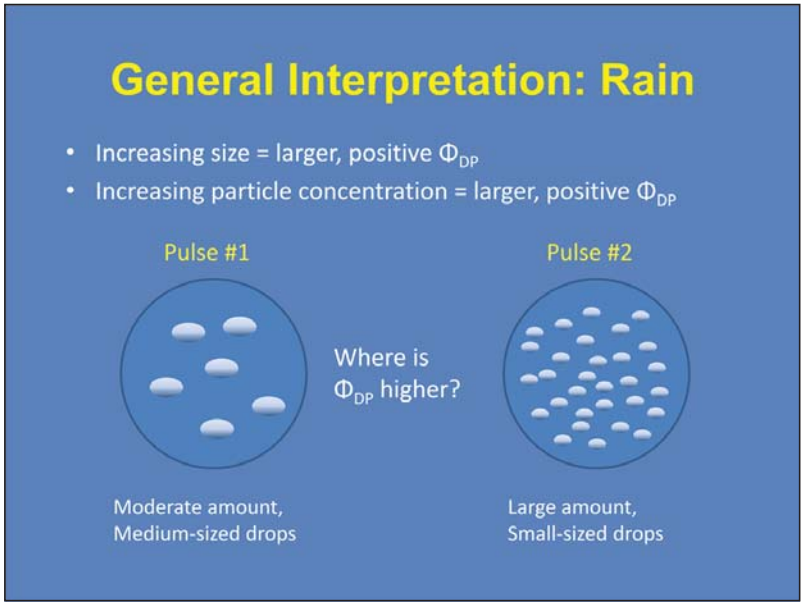
Here is the learning objective for this lesson. Take a moment to read this slide and advance the slide when you are ready.



Before diving into KDP, we need to learn about where it comes from. Back when Dual-Polarization was fielded on the 88Ds, the decision was made to make available the only the base Dual-Pol products and not immerse the field with additional products, including this product called Differential Phase. However, KDP, which is what much of this lesson is about, is actually derived from Differential Phase or Differential Phase Shift to be more precise – which is denoted as Phi DP. Phi DP is the difference between the horizontal and vertical two-way propagation phase shifts. What is a phase shift? That…you will learn in the next several slides on general interpretation. For now, just know that Phi DP is the difference between the horizontal channel – H in this case, and the vertical channel – or V. Phi DP can be used for data quality control and monitoring.

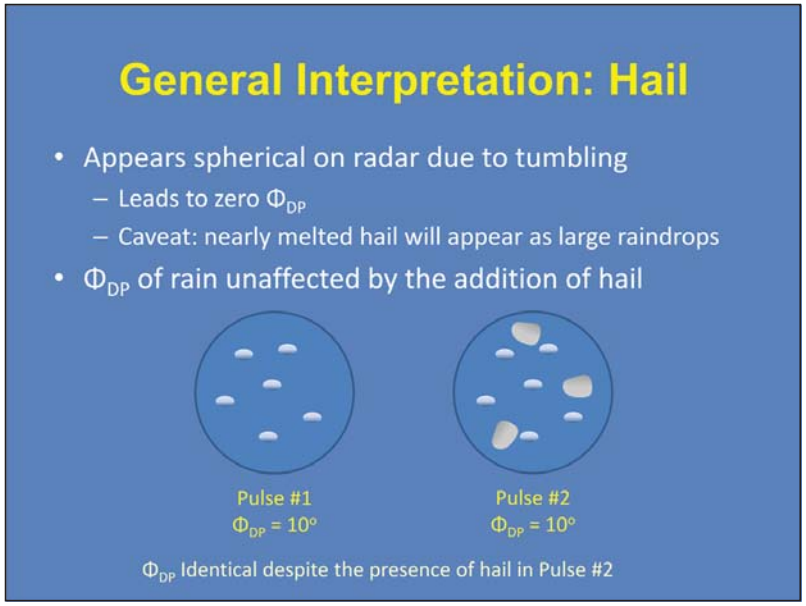
Why the strange-looking spikes and increases in the values down radial? Because Phi DP increases with liquid water content down the radial. In this image, gradients increase as the beam enters rainfall along the radial and elicits higher values with higher reflectivity. The product is also cumulative, so it is very difficult to tell where the important areas are, especially way down the radial once values are already high.



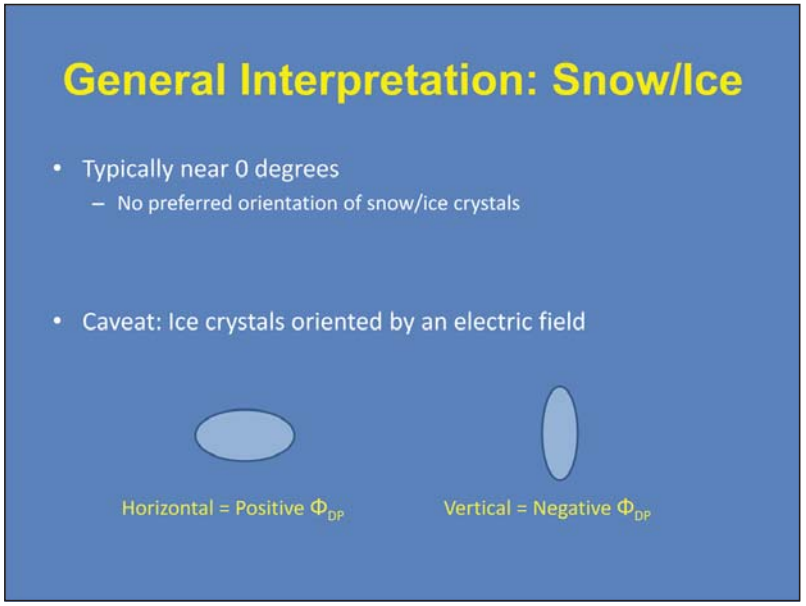


Much like differential reflectivity, the shape of the target affects the differential phase shift. Horizontally oriented targets will produce an increasing, positive differential phase shift with range. Vertically oriented targets will produce a decreasing, negative differential phase shift with range. And, spherical targets will produce near zero differential phase shift with range. Additionally, unlike ZDR, differential phase shift is dependent on particle concentration. The more particles there are in a pulse volume, the more differential phase shifting will occur. For example, the more horizontally oriented targets there are within a pulse volume, the higher the positive differential phase shifting.

Larger raindrops are more oblate than smaller raindrops. So, from the last slide we would expect small raindrops, such as drizzle, to have near zero differential phase shifting and larger raindrops to have larger, positive differential phase shifting. This is very analogous to the concept of ZDR, larger particles have larger, positive differential phase shifting. The caveat here is particle concentration. Let's say one pulse volume has a moderate amount of medium-sized raindrops and a second pulse volume has a large amount of smaller-sized raindrops. Even though the pulse volume with the medium-sized raindrops should have higher differential phase shift according to the size differences, the larger amount of smaller-sized particles could actually make the differential phase shifting in the second pulse volume higher than the first pulse volume.



Since hail tends to tumble as it falls, it appears on radar to be spherical. This means hail will typically have near 0 degrees of differential phase shifting. The one caveat to this is that when small, sub-severe hail is nearly melted, it appears on radar as large raindrops which can cause differential phase shifts to be very high. Another important point is the low dielectric constant of ice. Part of the reason that a melting hail stone produces unusually large KDP values is that the hail is wet. For the most part, differential phase shifting is unaffected by the presence of hail. This means that if we have two pulse volumes with the same amount and type of rain and one has hail and the other does not, the differential phase shifts will be identical despite the presence of hail.



Since most ice and snow crystals do not fall with a preferred orientation, differential phase shifts in snow and ice crystal regions are typically near 0 degrees. The one exception to this rule is when ice crystals become aligned due to some outside force, such as a strong electric field inside a thunderstorm. This external forcing leading to crystal alignment can cause significant non-zero differential phase shifting. Depending on the orientation of the alignment, the differential phase shifting can be positive or negative. It is positive when the alignment is in the horizontal, and negative when the alignment is in the vertical. The difference in dielectric constant between ice and snow is also an important factor. For example, you may see a preferential orientation of frozen hydrometeors in dendritic growth zones. Even when the hydrometeor concentration and shape are similar to heavy rain, the difference in Phi DP will be much smaller – but still greater than zero – with dendrites, because they are frozen.

General Interpretation (Non-Met)

- Appears very noisy
- Noise due to backscatter differential phase shifting
- Previous discussion: Propagation differential phase
  - Forward propagation effects only
- Backscatter differential phase:
  - Near zero in meteorological echoes
  - Highly variable in non-meteorological echoes, such as:


Birds


Bugs


Clutter

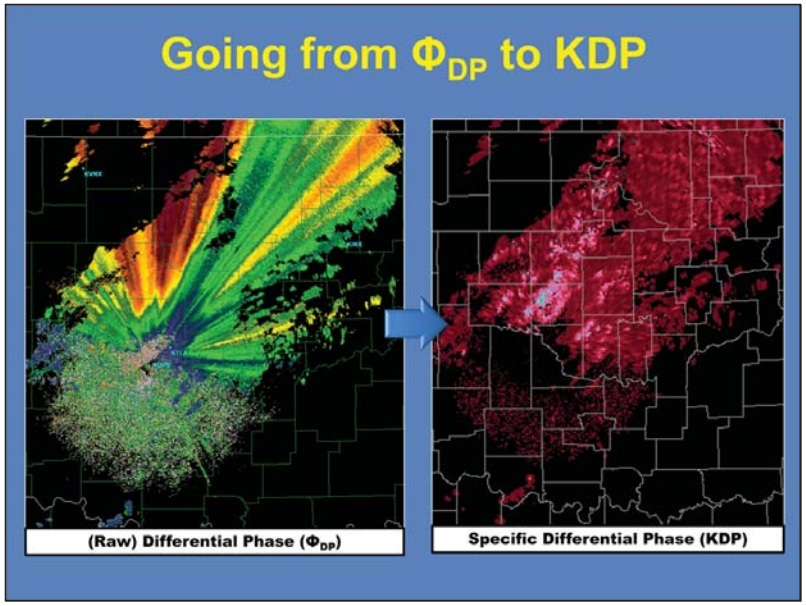
In non-meteorological echoes, differential phase shifting is very noisy. This is primarily due to the effects of backscatter differential phase shifting. This type of differential phase shifting is slightly different from what we have discussed so far. To this point we have only looked at what is called propagation differential phase shift. It is the amount of differential phase shifting that occurs solely as a result of forward propagation (i.e. how much differential phase shift has occurred as the pulse has gone forward in space). Backscatter differential phase shifting is the amount of differential phase shifting that occurs once the pulse scatters off of an object. In short, meteorological echoes typically experience no backscatter differential phase shifting, where as non-meteorological echoes such as birds and ground clutter experience significant backscatter differential phase shifting. The amount of backscatter differential phase shifting is highly variable depending on the object shape and size, hence why in non-meteorological echoes differential phase shifting appears very noisy.

Definition for KDP

Range derivative of the specific differential phase shift (Φ_{DP}):

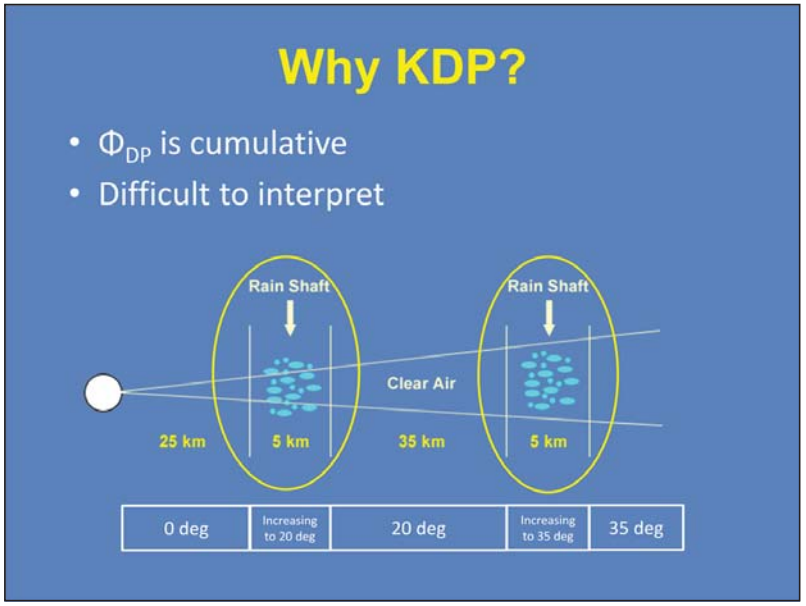
$$KDP = \frac{\Phi_{DP}(r_2) - \Phi_{DP}(r_1)}{2(r_2 - r_1)}$$

Alright, enough about Phi – which is not an operational Dual-Pol product, however, it does factor in to one that is, which is called Specific Differential Phase – or phase shift, and is denoted by (KDP). So what is KDP? It is the range derivative of the differential phase shift between the horizontal and vertical pulse phases. The next several slides will examine in detail the physical characteristics of KDP.

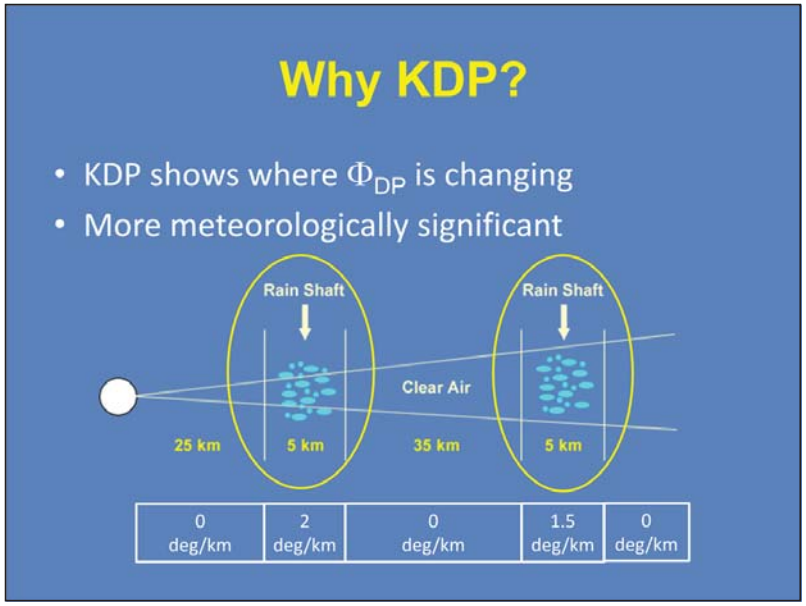


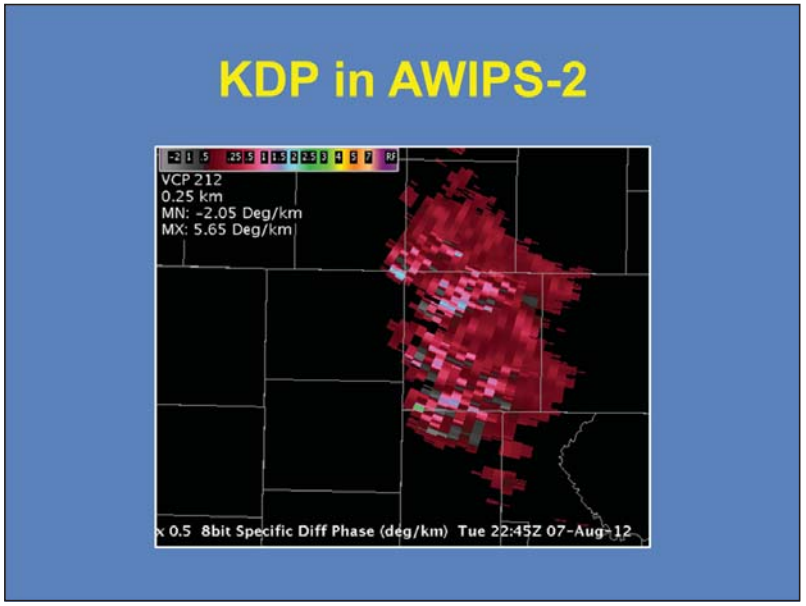
The Raw Differential Phase product, or Φ_{DP}, has been available in AWIPS since radar Build 17.0.

This product, which is called Differential Phase is the building block for KDP because it shows the difference in phases of the horizontal and vertical channels. So, just as ZDR shows the size difference in targets, this product determines the phase difference, which is a function of that target's movement within the space it's in. Is it a snowflake that is changing it's orientation as it falls? Or a tumbling hailstone? The movement or change in movement is relative to the horizontal and vertical planes. Think of it in a storm-relative or target-relative sense. On the left is an example Φ_{DP} product, and on the right is the associated KDP product.



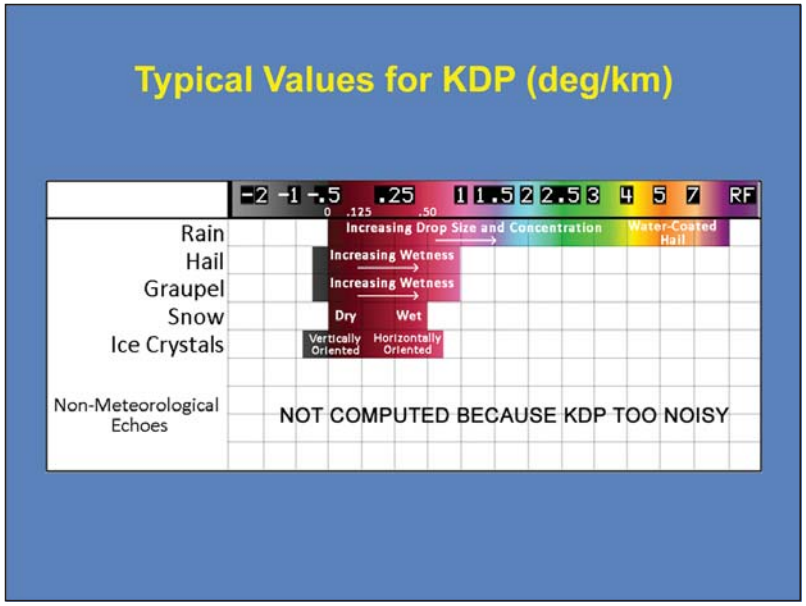
First of all, why KDP? Why can't we just used Phi DP? Phi DP certainly has discernable characteristics in rain versus hail versus snow/ice crystals. Well, let's look at the example here of Phi DP for two rain shafts at 30 and 70 km from the radar. As the horizontal and vertical pulses propagate toward the first rain shaft, they experience zero differential phase shifting, so Phi DP is zero. As the pulses go through the first rain shaft, the horizontal pulse slows down faster than the vertical pulse resulting in a positive differential phase shift. We'll say it is 20 degrees. After exiting the rain shaft, the pulses enter clear air and experience zero additional differential phase shifting. However those bins in the clear air will show a differential phase shift of 20 degrees because the differential phase shift cannot reset itself along a radial. In the second rain shaft, the differential phase shift will increase again, and let's say it increases by 15 degrees. In those bins and any bin further down range, the differential phase shift will be 35 degrees. As you can see, the differential phase shift is cumulative and the absolute value tells you nothing about what is going on in that particular bin, but rather all that has happened along the radial up to that point.





Now we'll look at the same example except using KDP, or the range derivative of Phi DP. Up until the rain shaft no differential phase shifting is occurring, so 0 degrees divided over any distance will give 0 degrees/distance. Therefore, anywhere within 25 km of the radar has a KDP of 0 deg/km. Inside the rain shaft, we said there was differential phase shifting of 20 degrees. If we divide this by twice the distance over which it occurred (10 km) we get a KDP of 2 deg/km. In the clear air past the first rain shaft, the differential phase shift remains 20 degrees but does not change over this distance. So, KDP will go back to 0 deg/km in this region because the difference between any two differential phase shifts over any distance in this region will be 0. In the second rain shaft, the differential phase shift increases from 20 degrees to 35 degrees, so it increases by 15 degrees. Dividing this value by twice the distance over which it occurred (10 km) gives a KDP of 1.5 deg/km. Past the second rain shaft, the KDP goes back to 0 deg/km for the same reasons it did in the clear air in between the two rain shafts. As you can see, KDP is much better at giving you information about what is happening at that particular bin than is Phi DP.

Here is an example of KDP in AWIPS-2. Notice how it is easier to discern regions of interest in KDP rather than Phi DP.



Here are the typical values for KDP given the various types of echoes listed on the left. Notice first how rain is the only meteorological echo that has a wide range of KDP values possible. All other meteorological targets are less than 1 deg/km. The one exception being small, melting hail, which we'll discuss later. The other note to make is for non-meteorological echoes, KDP is not displayed because KDP is too noisy due to the low CC.



Product Name	Units	AWIPS ID	RPG ID	RPG Code	Data Levels	Range Res	Az Res	Max Range
Specific Differential Phase Data Array Product	*deg/km*	*KDP*	*DKD*	*163*	*8-bit (256)*	*0.25 km*	*0.5 deg*	*300 km*
Super Res Specific Differential Phase Data Array Product	*deg/km*	*SDP*	*SDP*	*168*	*8-bit (256)*	*0.25 km*	*0.5 deg*	*300 km*

There are two 8-bit KDP products from the RPG. The second one is the Super Res product which is better known as Raw Phi DP.

General Product Characteristics

- Displayed as a polar coordinate with a 1° beamwidth
- Displayed relative to the RDA site
- Available for any elevation angle of the current VCP



KDP products are displayed in polar coordinates with a 1 degree beam width and data are displayed relative to the RDA site. KDP is available for all elevation angles for the current VCP.

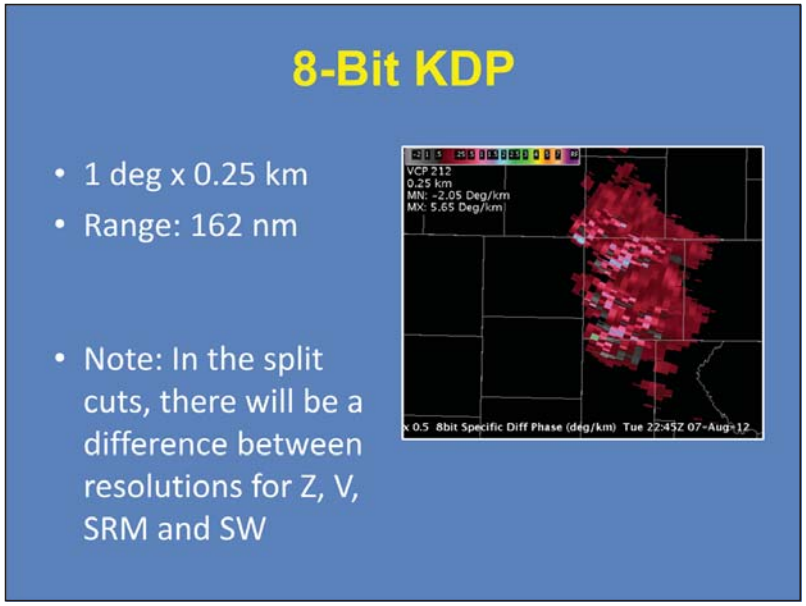
Menu Locations

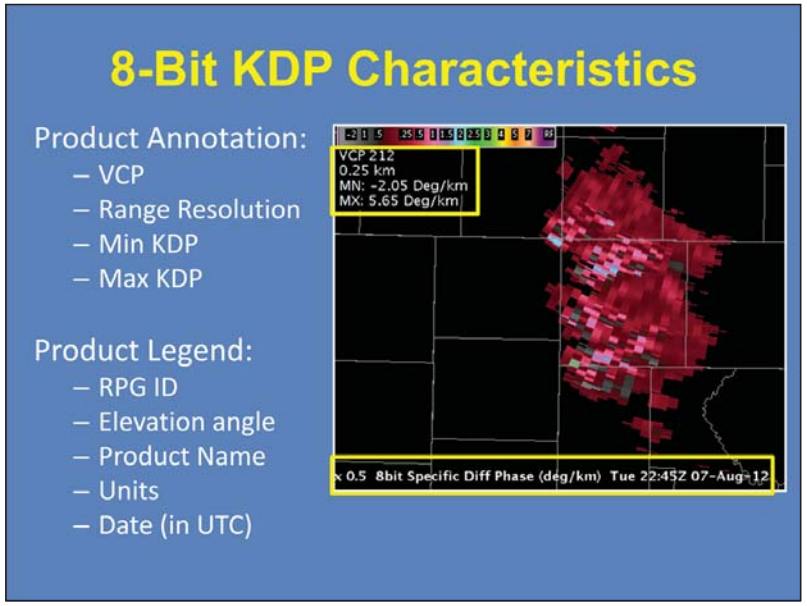
KDP combined with all base data

Individual KDP products (elevation-based)



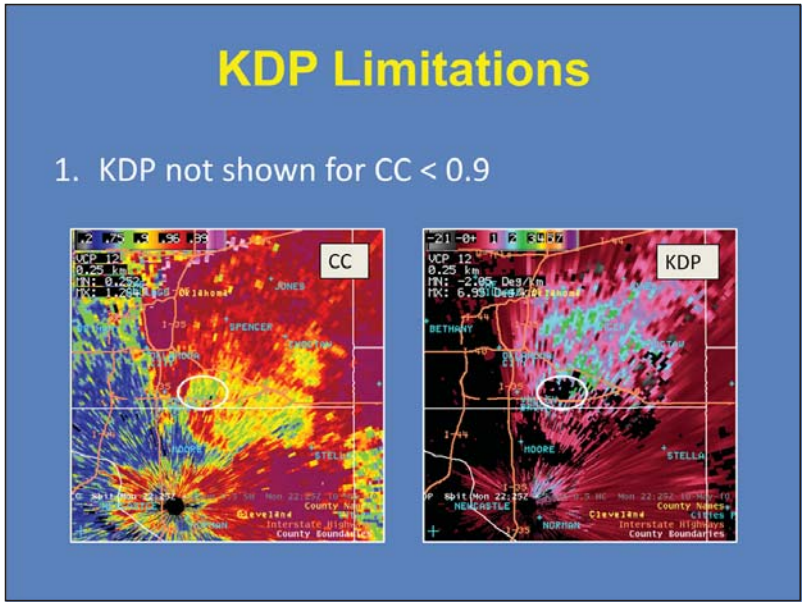
You can find the KDP products combined with all base data in your dedicated radar drop-down menu in the first two options highlighted. And, the individual, elevation-based KDP products can be found in the third option.



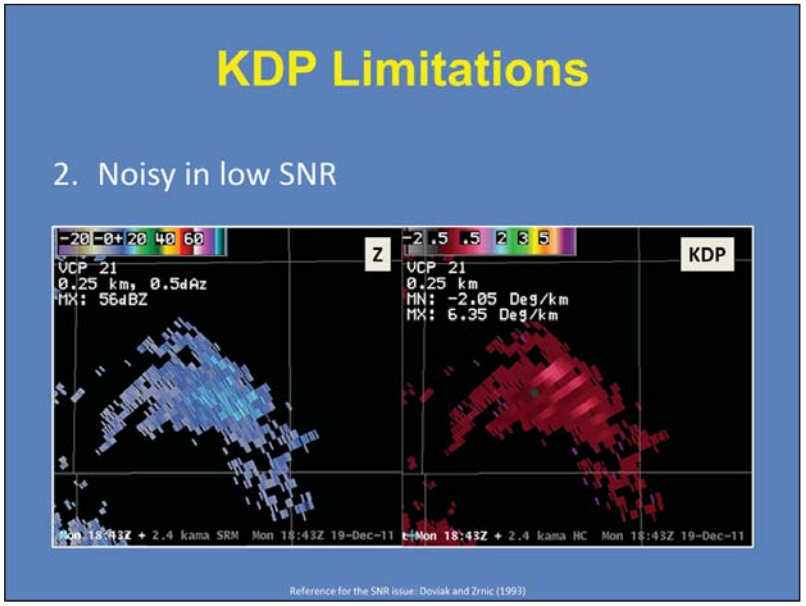


The 8-bit KDP product is 1 deg x 0.25 km resolution and extends to a range of 162 nm. It is worth noting here when viewing products in the split cut elevations, toggling between KDP and Z,V, SRM or SW there will be a difference in resolutions. This should not change your interpretation though.

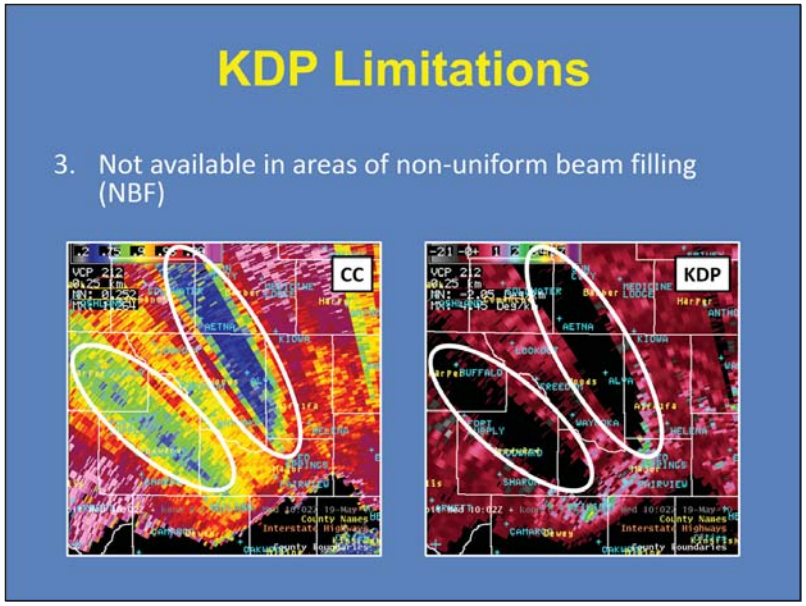
In the upper left portion of the KDP product is the product annotation. Here you can find information on the radar's current VCP mode, the range resolution and the min/max value anywhere in the KDP field. Note, you will not know where in the KDP product those values occur. In the lower right is the product legend information. There you can find the radar ID, elevation angle, product name (in this case KDP) and the date/time in UTC.



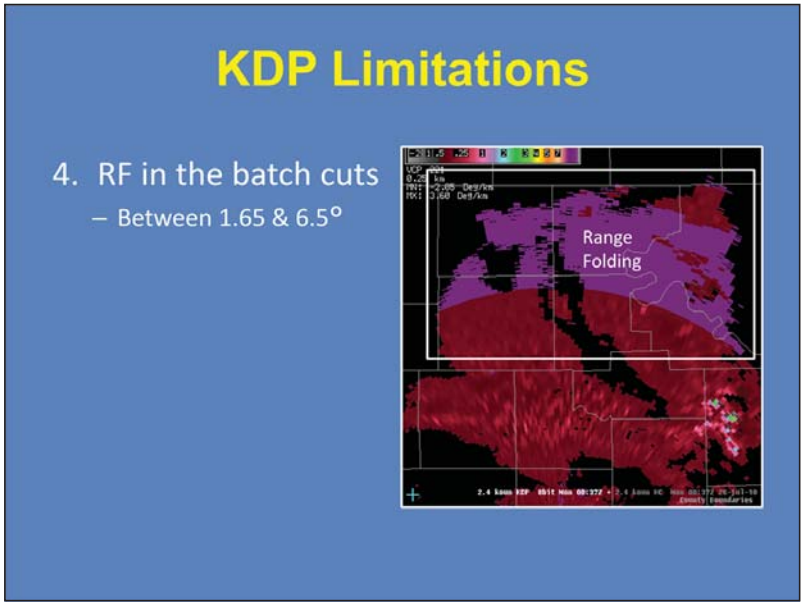
When the CC is below 0.9, Phi DP accumulates significant errors making it appear very noisy. This noisy Phi DP makes it useless to compute KDP. Therefore, in bins with CC less than 0.9, KDP will not be plotted, so it will appear as black holes in the data. Here is an example where KDP is “blacked-out” because CC was less than 0.9.



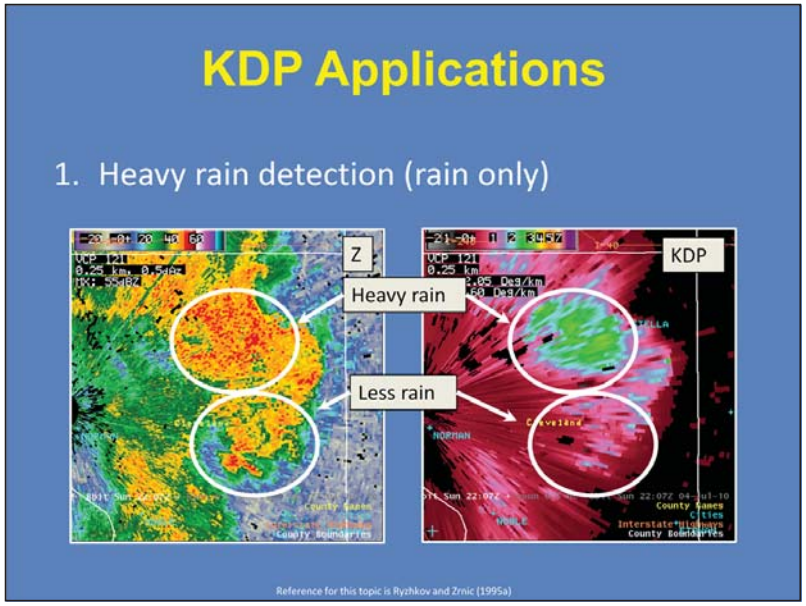
In regions of very low signal where the signal-to-noise ratio (SNR) is very low, KDP will look noisy. Note in the example how KDP is some what chaotic in the low Z regions far from the radar. Reference for the SNR issue: Doviak and Zrnic (1993).



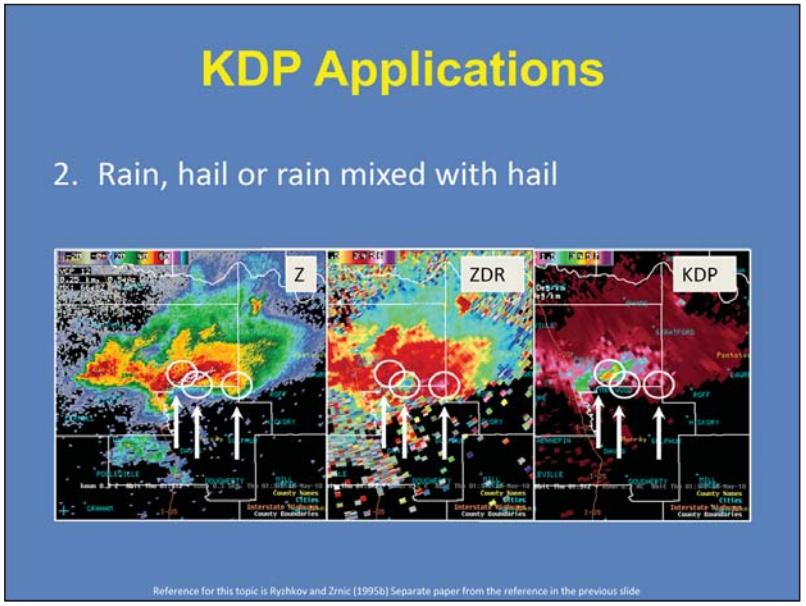
We saw in the CC lesson that non-uniform beam filling causes CC to dramatically decrease along the radials that experience NBF. And, as we just saw a couple of slides ago, if CC less than 0.9, then KDP does not get computed. Therefore, in regions of NBF, especially when CC less than 0.9, KDP will be “blacked-out”. Here is an example of extreme NBF that caused large areas of KDP not to be displayed.



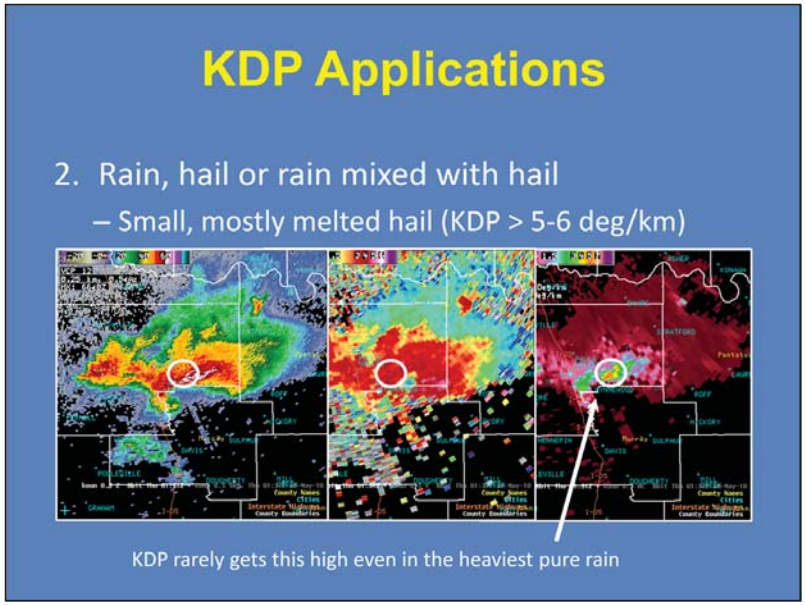
The last limitation with KDP that we’ll look at is in the batch cuts, range folding may obscure some signatures in KDP. As a refresher, the batch cuts for the WSR-88D network are elevation angles between 1.65 degrees and 6.5 degrees. In these cuts each radial uses a series of alternating low and high PRF pulses and the dual-pol variables are only computed using the high PRF pulses which are subject to range folding. Here is an example of RF affecting KDP.



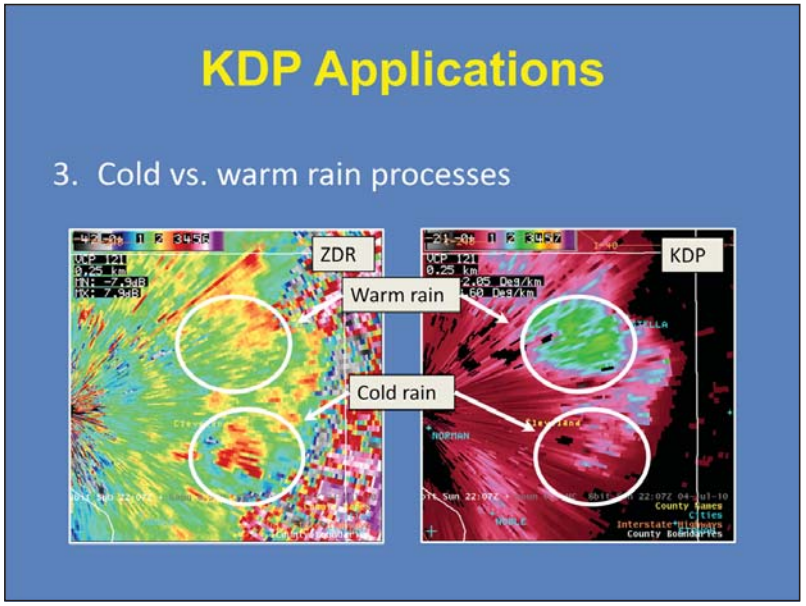
The primary advantage of KDP is its ability to detect heavy rain situations. Here is an example where we have a modest area of greater than 40 dBZ echoes that is fairly uniform in intensity. Looking at KDP, we see higher KDP values to the north, and lower KDP values to the south despite reflectivity values being almost identical. This tells us that there is heavy rain falling where the KDP values are higher and not as much rain where KDP is lower. We'll talk more about this example in a few slides. A reference for this topic is Ryzhkov and Zrnic (1995a), focusing more on R(KDP) vs. R(Z), but it highlights the utility of KDP as a useful tool for identifying areas of moderate to heavy rain.



KDP is mostly immune to hail contamination. Here is an example to illustrate. Looking at the white circle farthest to the east, we see Z is moderately high and ZDR is very high. However, KDP is very low, barely reaching 1 deg/km. This tells us we have large rain drops, but in low concentrations in this region. So, rainfall is going to be minimal. Nearer the core (middle white circle), we see that Z are higher, so we might be thinking increased rain. However, ZDR tells us that this is likely hail because ZDR is a local minima (less than 2 dB) and KDP is very near 0 deg/km. So, this area is mainly hail, with very little rain mixed in. The last area to investigate is the farthest west white circle. Z is roughly the same as the last example, but now ZDR is moderately large (around 4 dB) telling us we have good-sized drops in this region. Looking at KDP, values are approaching 5 deg/km. This tells us there is a lot of liquid water in this region, so flash flooding could be an issue.



One caveat to point out here is small, mostly melted hail. In these situations, flash flooding may not be a threat but KDP could become very large. This is because in regions of very small, mostly-melted hail, the radar sees these particles as very large rain drops, and if there are a lot of these particles present, KDP sky rockets to values as high as 10 deg/km! A good rule of thumb here is that in pure rain situations, KDP will rarely get above 4-5 deg/km even in the heaviest of rains. So, if KDP are greater than 5-6 deg/km, you can bet that there is some small, melting hail present and you will need to mentally adjust your estimates of rainfall based on KDP, or rely on other means of estimating rainfall.

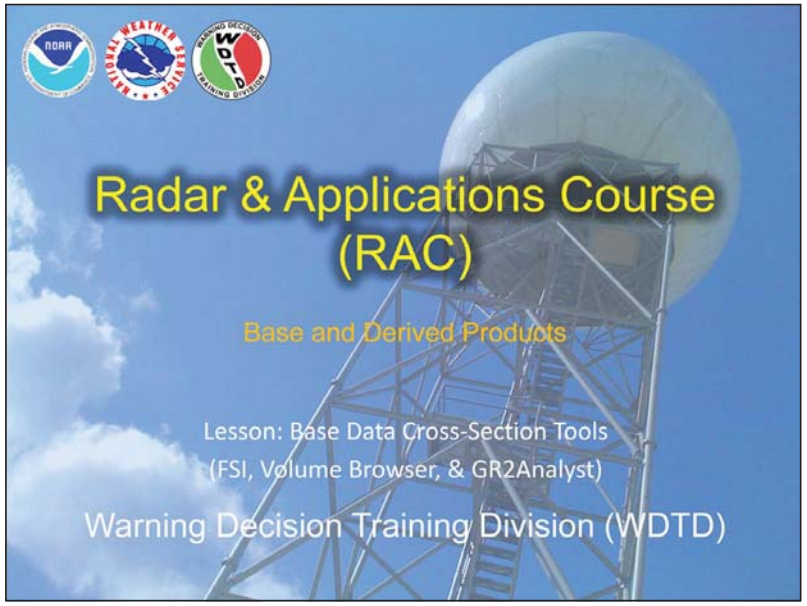


Here is the example from 2 slides ago, except now we are showing ZDR and KDP together. Recall that Z was similar in both areas highlighted here. Note, to the north, ZDR are moderately high but don't really exceed 3 dB. However, to the south, ZDR are approaching 4-5 dB. This tells us that larger drops are occurring to the south and smaller drops to the north. But, KDP is telling us heavier rain is occurring to the north. What gives? This is an example of how KDP, with ZDR, can reveal warm-rain vs cold-rain processes. And, if you see warm-rain processes, this indicates very efficient rainfall and might warrant a flash flood warning soon given all other factors are in place.

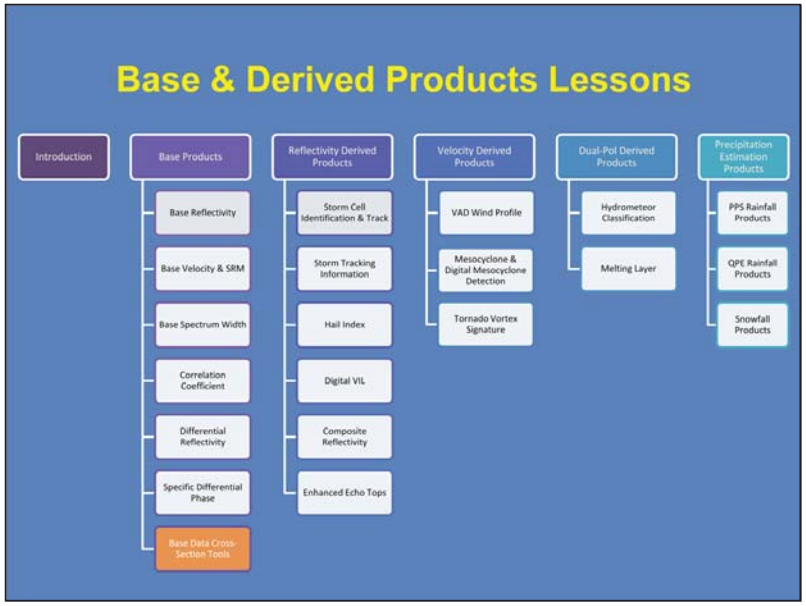
Summary: KDP

- Range derivative of differential phase shift
  - Great for identifying areas of high liquid water content
- Standard radar limitations apply, plus:
  - Data voids where CC < 0.9
    - Same for areas of NBF
  - Noisy in low SNR
  - RF in batch cuts

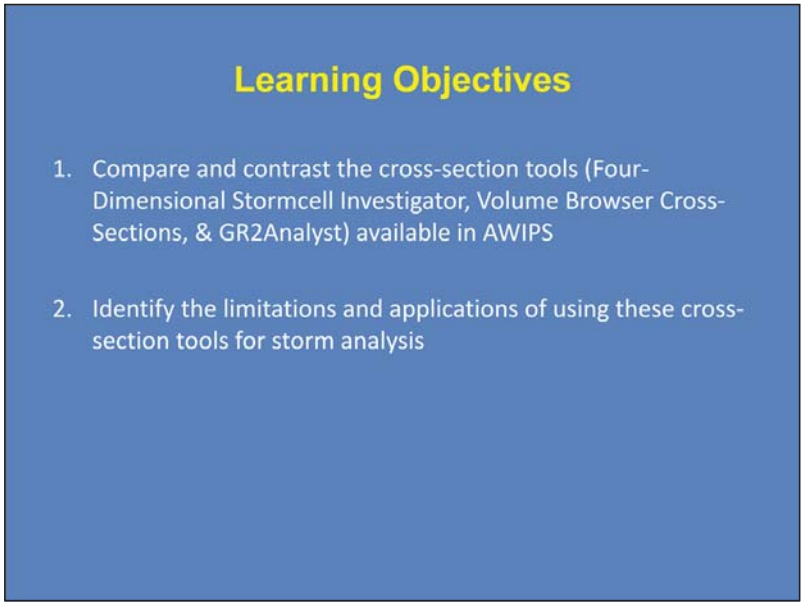
In summary, KDP is the range derivative of the differential phase shift, and is primarily good for identifying regions of heavy rain, even in hail contamination cases. However, KDP is not computed where CC less than 0.9, is noisy in low SNR regions and is RF in the batch cuts.



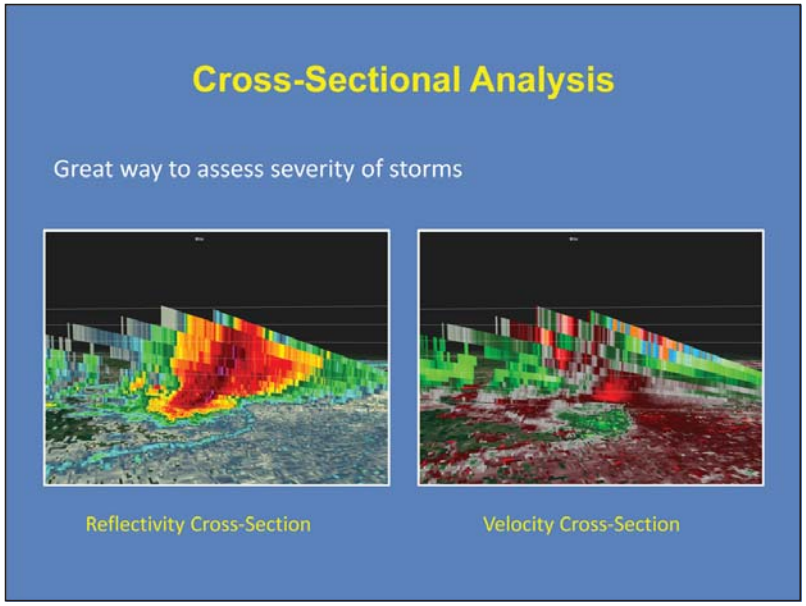
Welcome to this lesson on Base Data Cross-Section Tools. This training is part of the Base and Derived Products topic in the Radar and Applications Course. Let's get started.



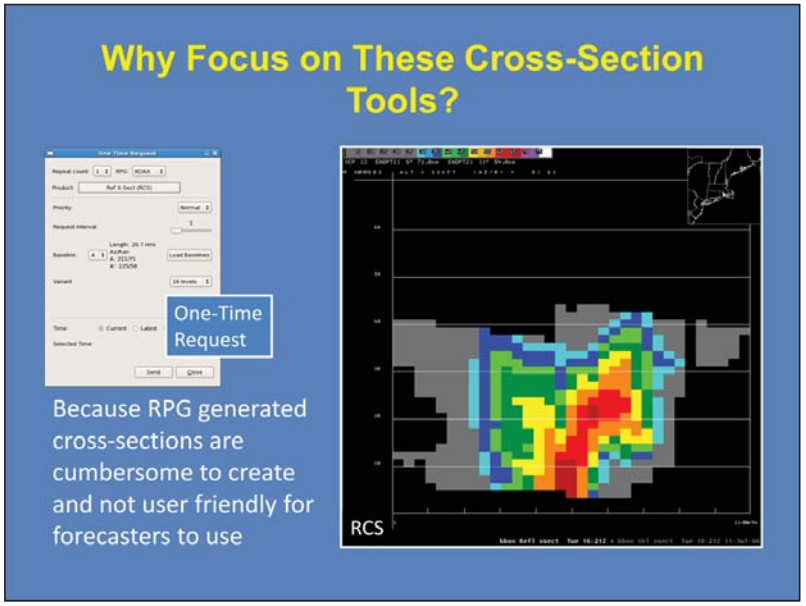
Here is a roadmap for the lessons in this topic. This lesson on the Base Data Cross-Section Tools, shaded in orange, is the last part of the Base Products Section of this topic.



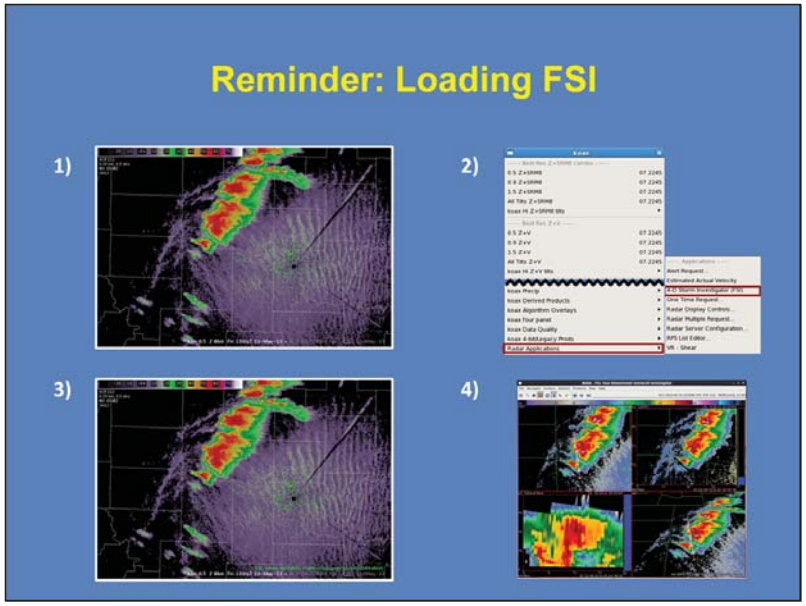
Here are the learning objectives for this lesson. Take a moment to read these objectives. Advance the slide when you are ready.



Cross-sections are a great way to assess storm severity, especially cross-sections of base data like Z, V, ZDR, CC, etc. This lesson introduces the three tools NWS forecasters can use to generate cross-sections. Interpretation of cross-section plots will be covered in the convective storms topic lessons found later in the course.

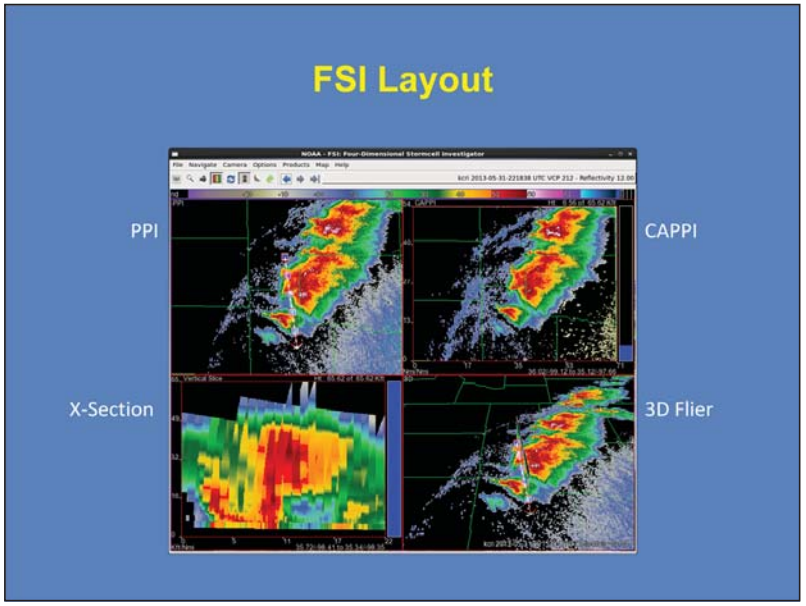


The Radar Product Generator can produce cross-section products available through One-Time Requests or Radar Multiple Requests. However, generating cross-section products in this way will feel antiquated and cumbersome to most forecasters. As a result, we will focus on the three cross-section tools available to NWS forecasters: The Four-Dimensional Storm Investigator, Volume Browser (both in AWIPS-2) and the standalone application GR2Analyst. These tools are more user friendly and offer other applications that the RPG cross-section products cannot match.

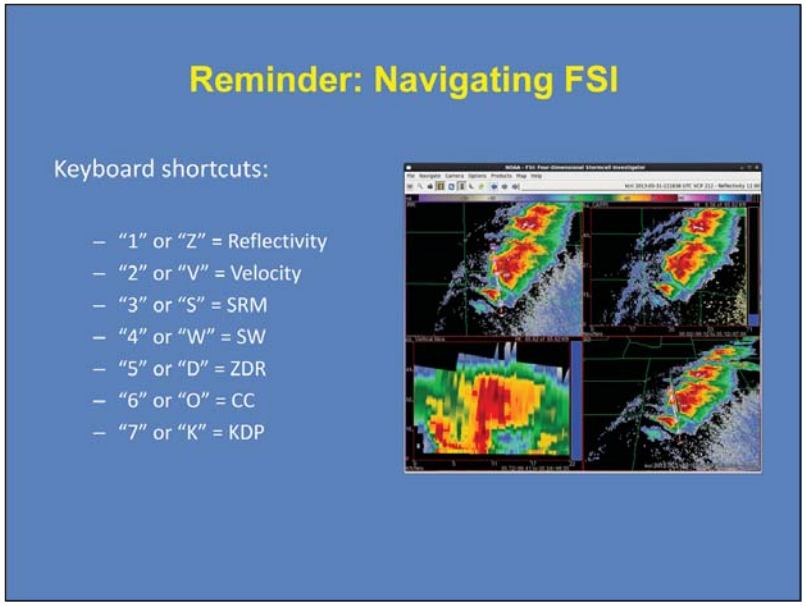


We will start with FSI. In case you need a reminder, here are the steps for loading FSI from the AWIPS Fundamentals training:

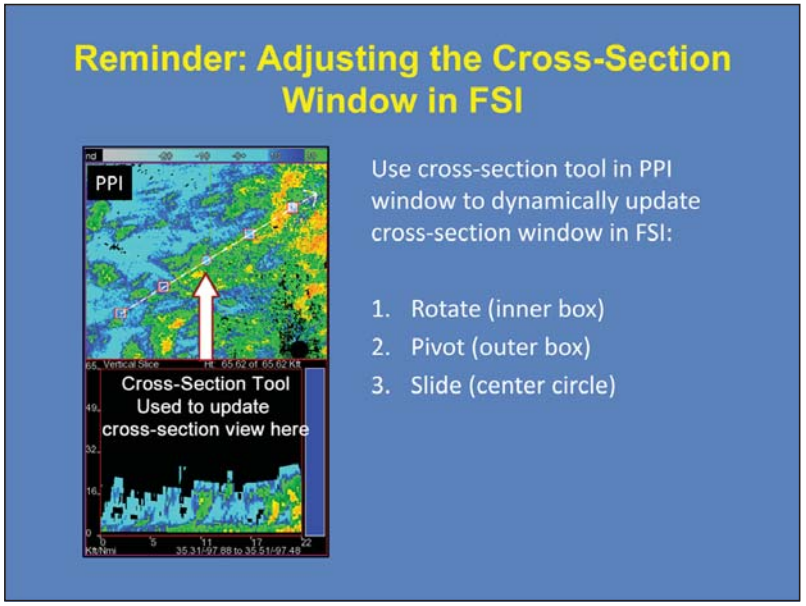
- First, load a 0.5 deg base data display with Reflectivity;
- Choose “4-D Storm Investigator” from either a specific radar menu or the Tools menu;
- Right click on the storm you want to investigate; and
- The FSI will appear momentarily in a separate window.



Here is an example of the FSI window. In the top left is the PPI, or what you normally see in AWIPS-2. The top right is the CAPPI panel, which stands for constant-altitude PPI. Here you can define a constant height, and see how the different base data look at that height. The bottom left is the vertical cross-section window. Here you can see vertical cross-sections of the base data, which will be the main focus of this lesson. Finally, the lower right is the 3D flier window. This is where you can get as close to a 3D look as possible by combining a perspective view with the CAPPI and vertical cross-section views.



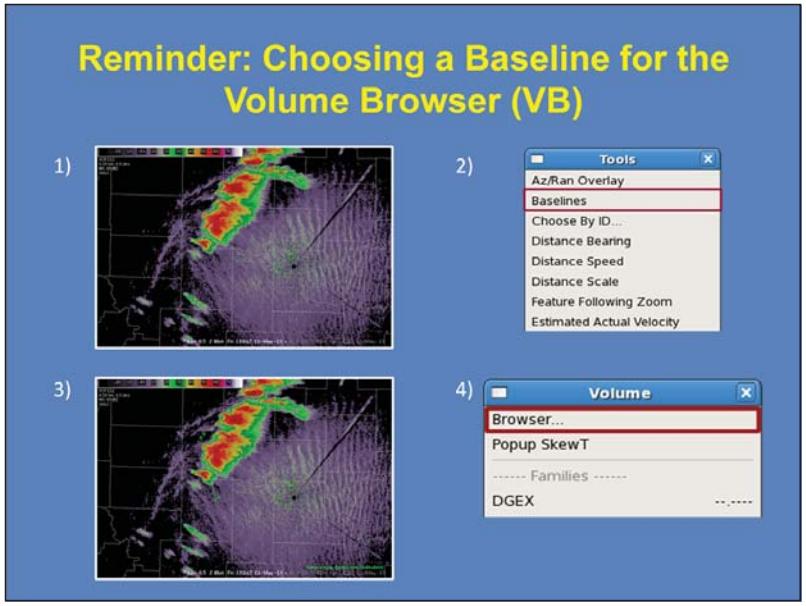
Just to remind you, all base data are available for viewing in FSI. Use the alphanumeric shortcut keys shown in order to toggle between each of the base products. Look over the display and, when you are ready to proceed, click "Next" to go to the next slide.



More than likely, you will want to adjust the cross-section baseline that FSI provides when it opens. There are three ways to manipulate the baseline provided by dragging different objects on the feature. You can:

1) Rotate the line (using the inner boxes);
2) Pivot the line (using the outer boxes); and
3) Slide the line (using the center circle).

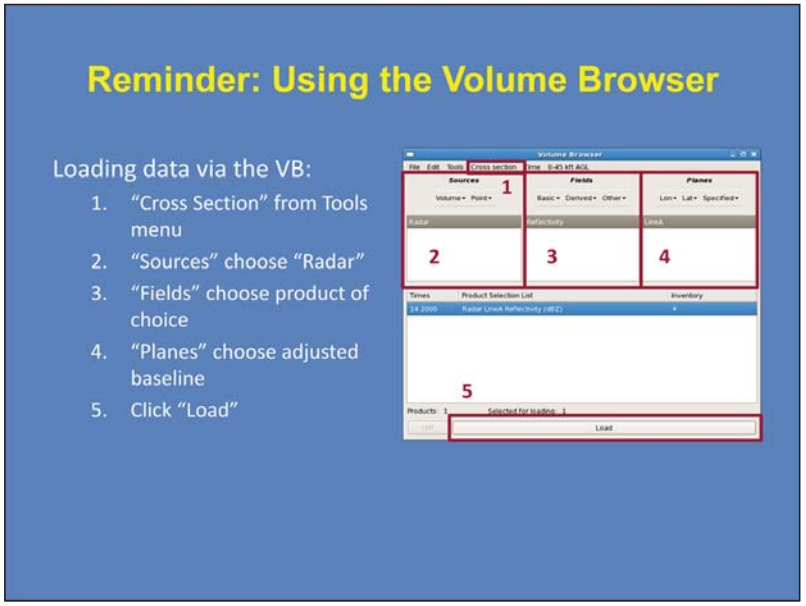
That covers the basics of FSI.



Now let's discuss cross-sections generated using the Volume Browser. Here are the steps for loading a cross-section using this tool:

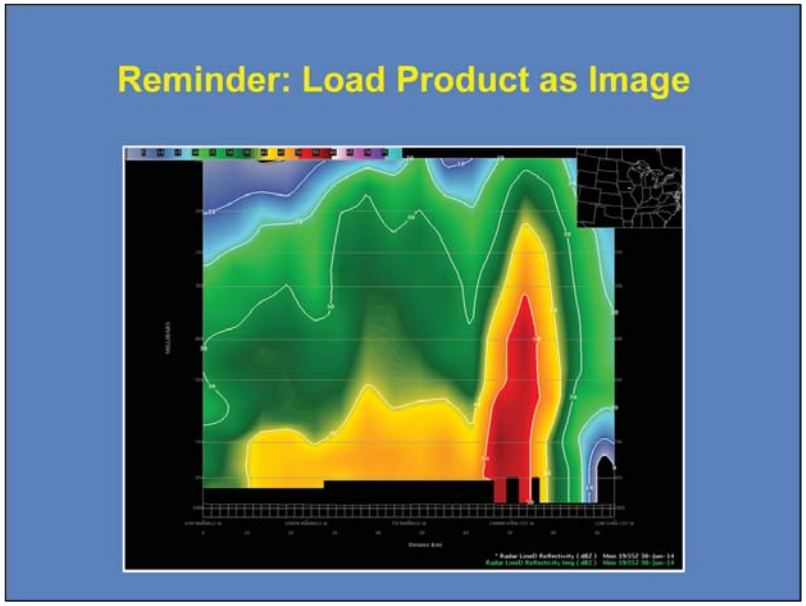
1) Load a 0.5 deg base data display with Reflectivity;
2) Load "Baselines" from the Tools menu;
3) Adjust one of the baseline's length and orientation to sample your storm of interest; and
4) Start the Volume Browser from the Volume menu.

We will cover the remaining steps on the next slide.



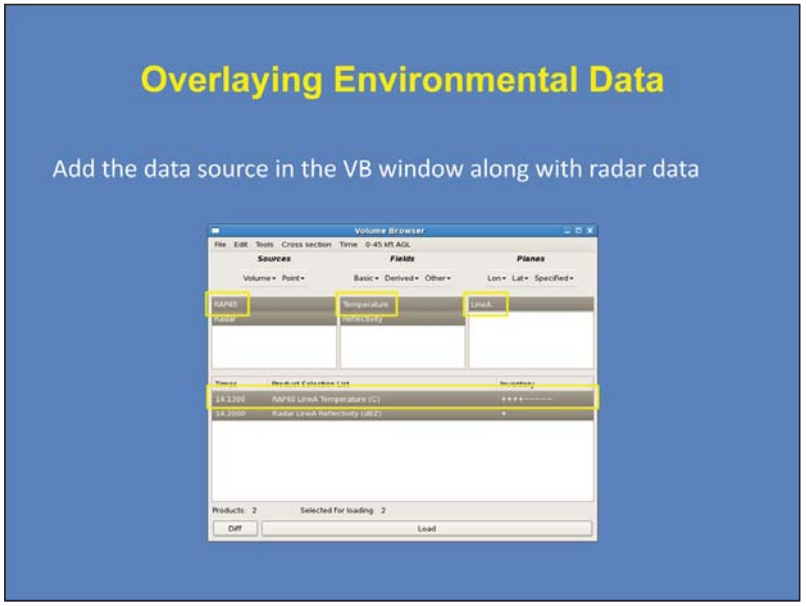
Once the Volume Browser window appears, follow these steps:

1) Choose "cross-section" from the tools menu;
2) Choose "Radar" for the Source;
3) Choose the base data product you want to view for the Fields;
4) Choose the baseline you edited for the Planes; and
5) Once the product you selected appears in the bottom panel, click "Load".

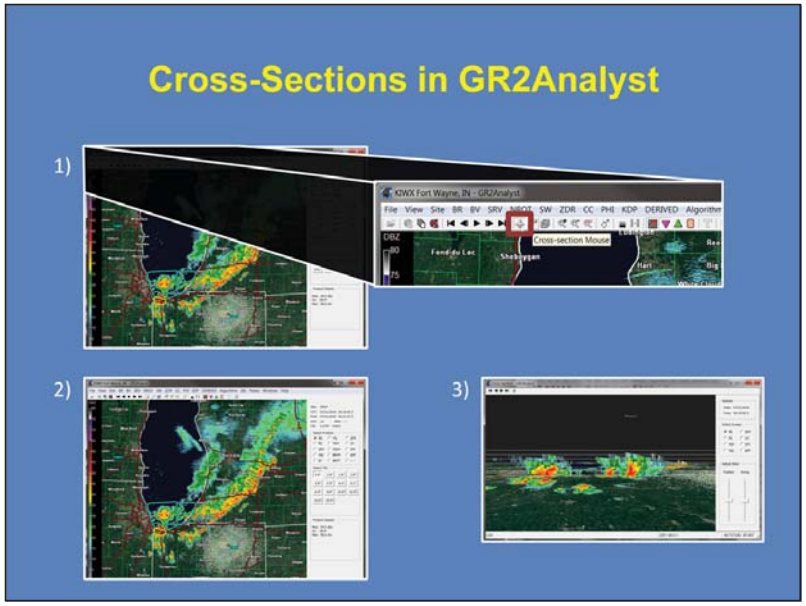


When your cross-section first loads, it appears initially as just contours. You will likely want to make your radar product cross-section an image product instead. To make that change, right-click over the product legend info and choose “Load as Image” from the menu that appears. Here’s an important note: You will want to create a new color table specifically for cross-sections if you continue to use them. The cross-section plots use a smaller range of values than standard PPI displays, so using the same color table in a volume browser cross-section will make storms appear visually weaker.

One other thing you should know. One major issue limits use of Volume Browser cross-sections at the moment. Only Reflectivity will display properly. The functionality should be addressed at some point in the future, but the fix has been kicked down the road for now.

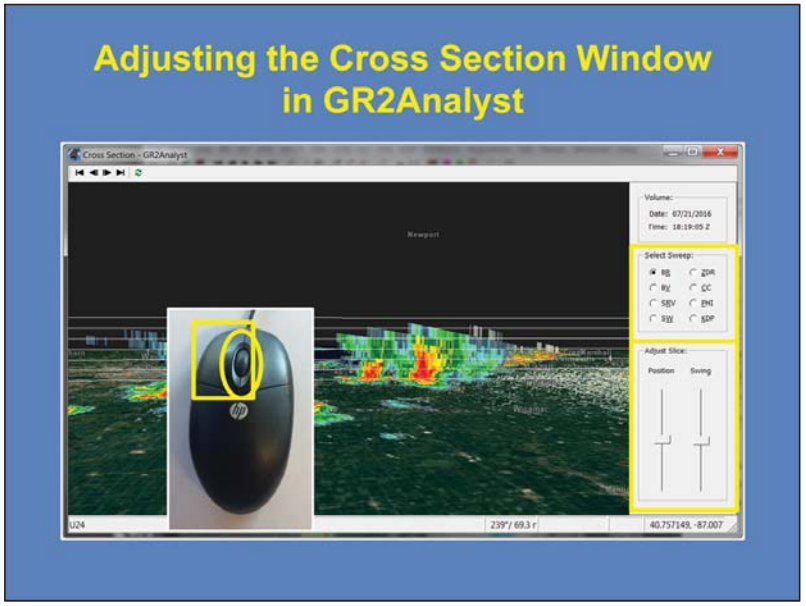


With Volume Browser product cross-sections, you can overlay environmental data such as temperature directly on your image. Just choose the data source and field for the parameter you want to overlay, then use same baseline for your Plane, and AWIPS will plot the parameter on top of the original cross-section.



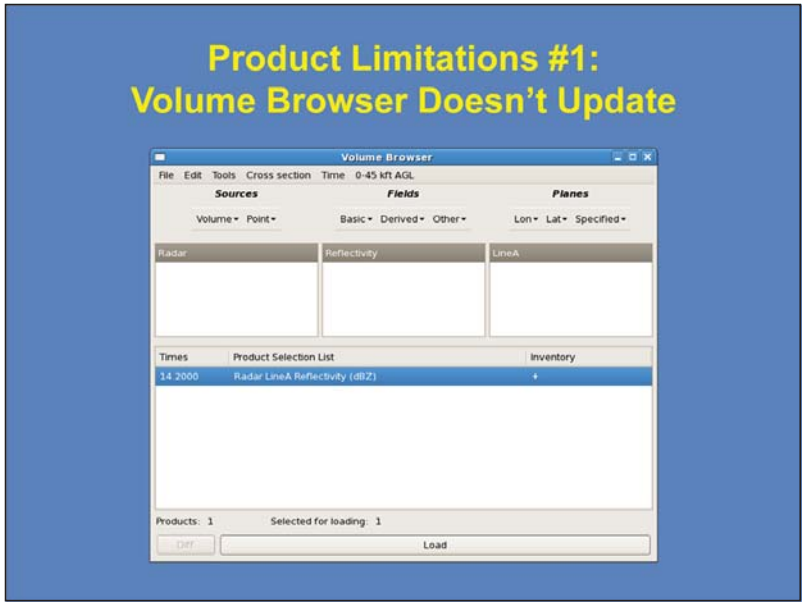
Lastly, let's talk about GR2Analyst. Unlike the previous two methods of generating cross-sections, we have not been training you on how to use this application in previous lessons. However, for the sake of time, I'm just going to assume that most NWS forecasters have an elementary knowledge of GR2Analyst as a third-party application. With that in mind, here are the steps for generating a vertical cross-section:

1) Select the Cross-section Mouse from the toolbar in the main GR2Analyst window;
2) Drag out a straight-line in the display window; and
3) The cross-section will appear in a new Cross Section window.

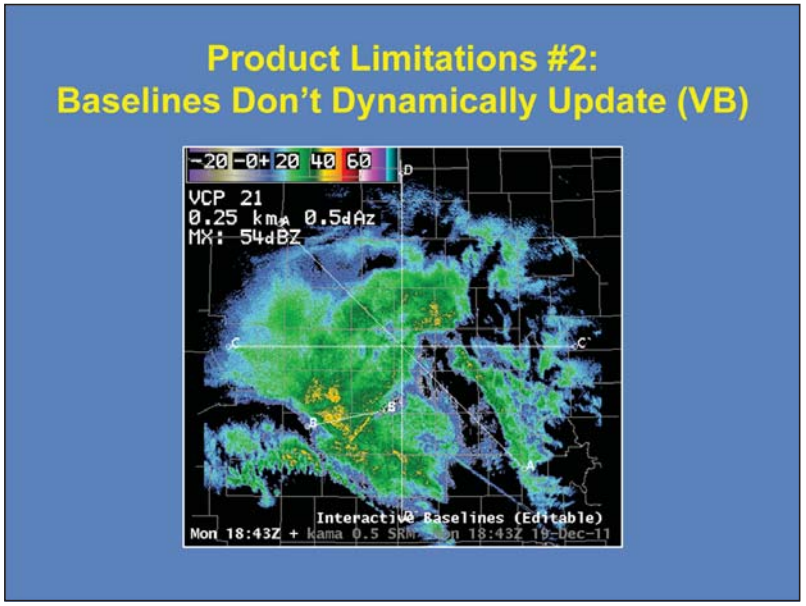


You have several options available to adjust the cross-section display in GR2Analyst:

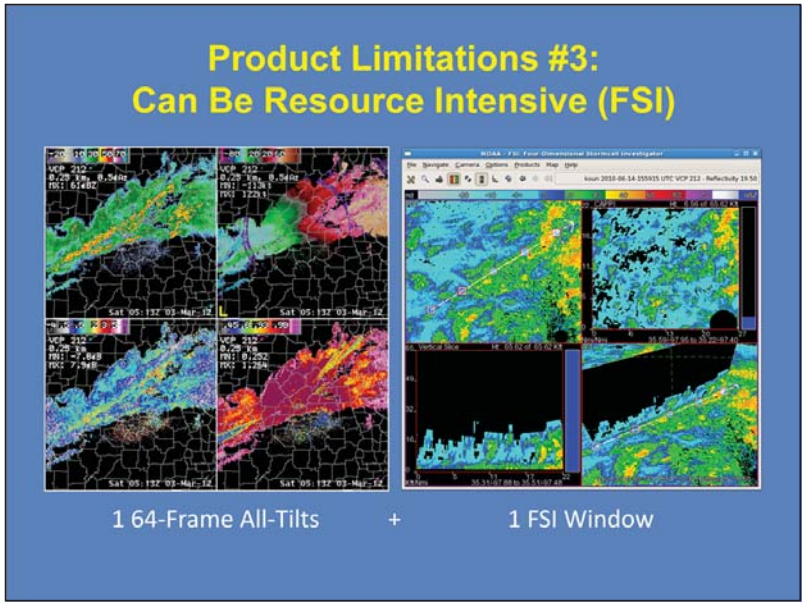
1) To zoom in and out, you can use the scroll wheel;
2) A left-mouse click and hold will allow you to drag or rotate your display in three-dimensions;
3) You can use the "Select Sweep" radio buttons on the right to change the product shown in the display; and
4) The "Adjust Slice" slider bars allow you to find tune the location of your cross-section baseline.



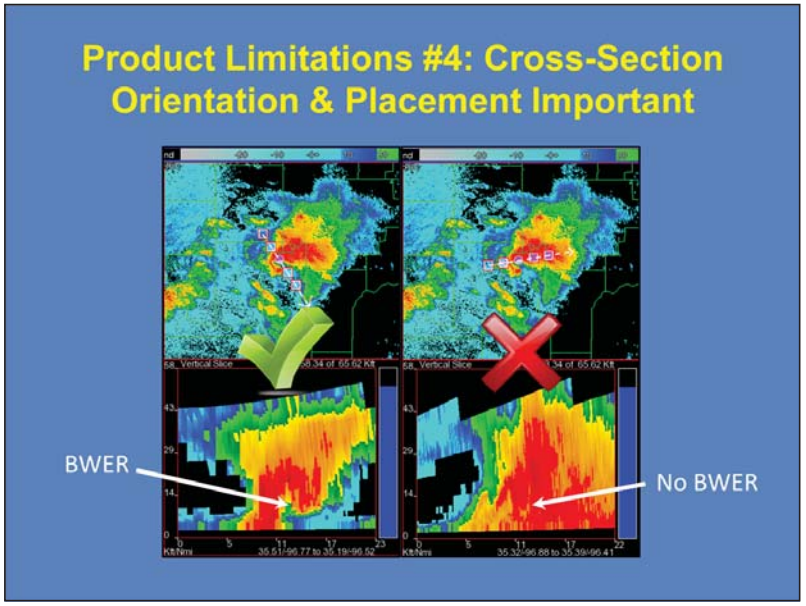
So let's start with the limitations of these cross-section tools. Much like the cross-sections from the RPG, Volume Browser cross-sections do not update as new data become available. This is a bug in AWIPS-2 and should be fixed soon, but until then, if you notice the cross-section not updating as new data come in, just re-load the product (adjusting the cross-section orientation if necessary).



Unlike in FSI and GR2Analyst, cross-sections generated by the Volume Browser do not dynamically update when you change the baseline. In other words, each time you change the baseline, you must re-load the product from within the Volume Browser to display the new cross-section. This workflow can be prohibitively cumbersome during high end events.

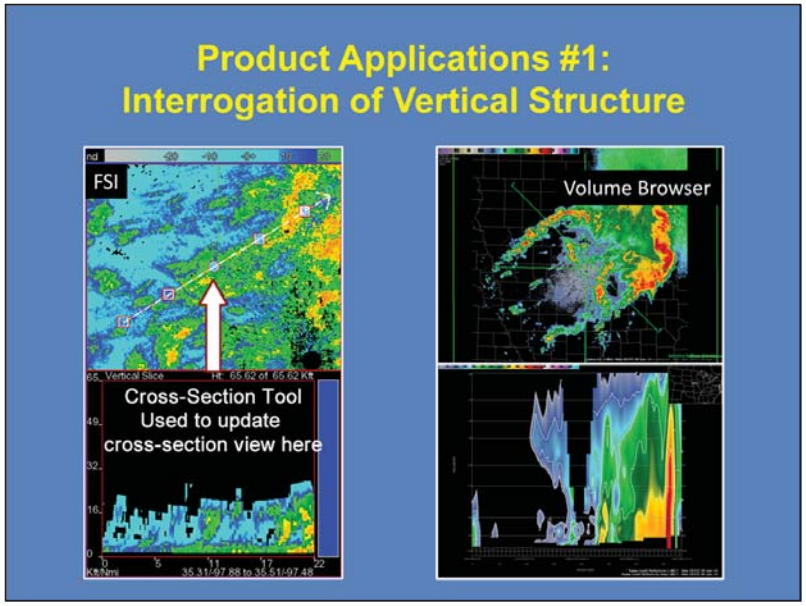


The downside to FSI's dynamic updates is that it requires a lot of computer resources to do it. Depending on what you have running, FSI can negatively impact performance by significantly slowing your system down. We recommended limiting yourself to only one 64-frame, radar all-tilts display and one FSI window running on the same monitor. If you notice slow system performance even with these restrictions, try closing your FSI window and see if any improvement occurs.

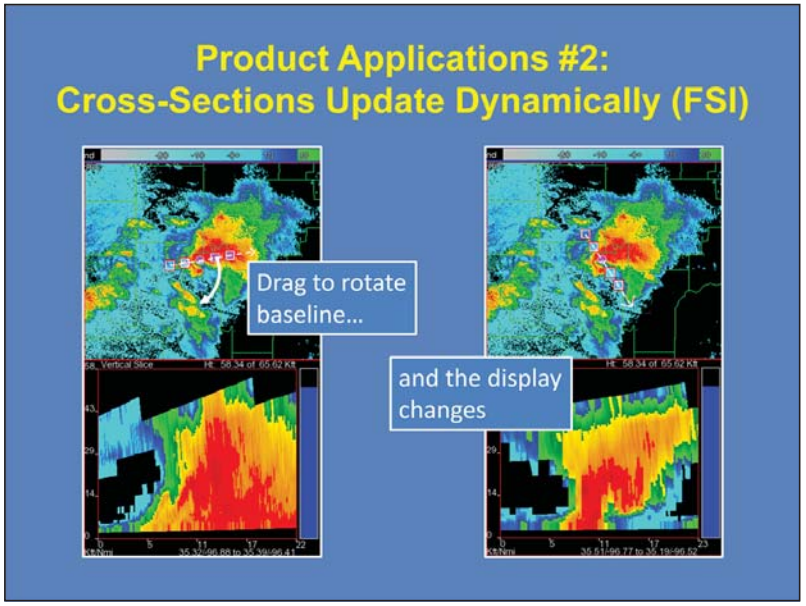


The last limitation we will mention regards the orientation of the cross-section. Improper orientation of your cross-section baseline can result in you missing important storm features. Let's use the example we show on the slide.

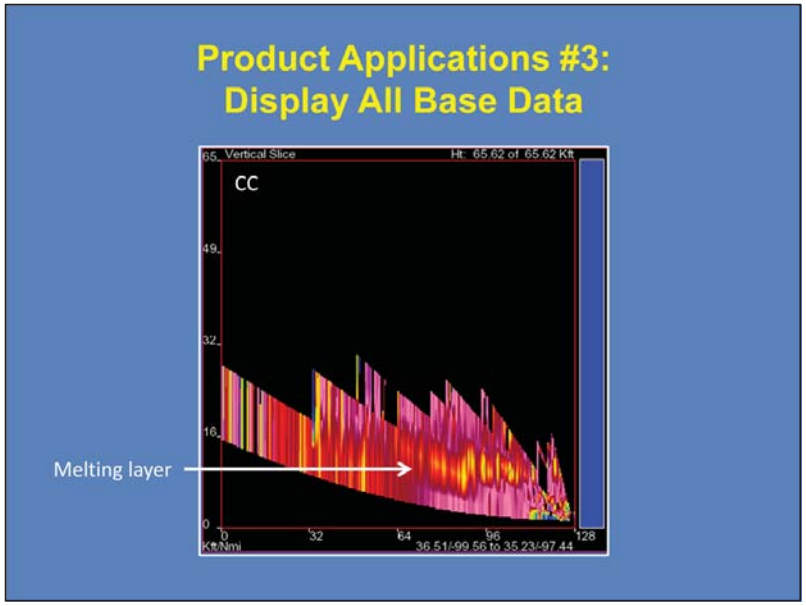
Here is a supercell thunderstorm that I suspect has a Bounded Weak Echo Region (BWER) that I want to analyze using a cross-section. The baseline on the left properly dissects the inflow notch, allowing me to clearly see the BWER and reflectivity overhang. The baseline on the right misses this critical part of the storm's structure and, instead, cuts through the storm's core. The end result: No BWER or nice updraft structure like I was suspecting in this second baseline. So, you can see how the orientation and placement of the baseline is crucial when searching for specific features.



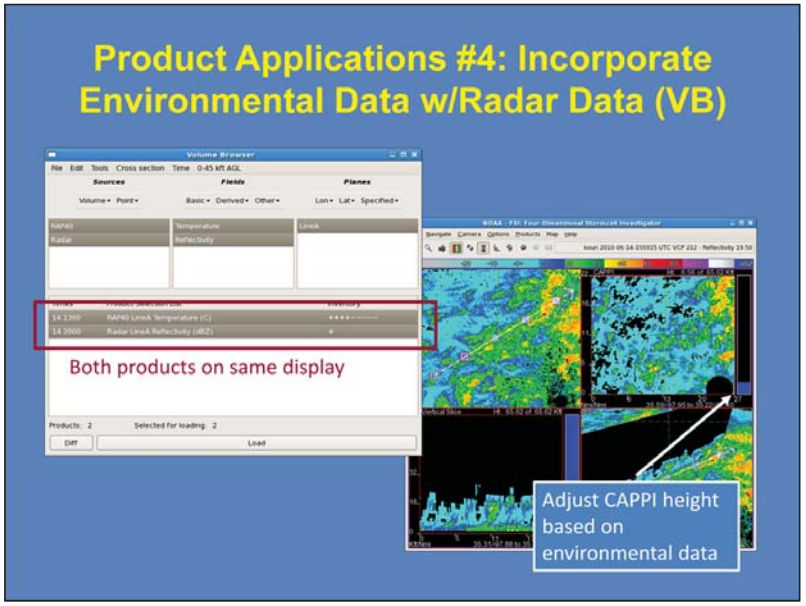
Now let's look at some applications with these products. All three types of cross-sections are much easier to create than the old RPG-based cross-sections. These solutions make it simpler to analyze the vertical structure of storms in real-time, which aides in warning operations.



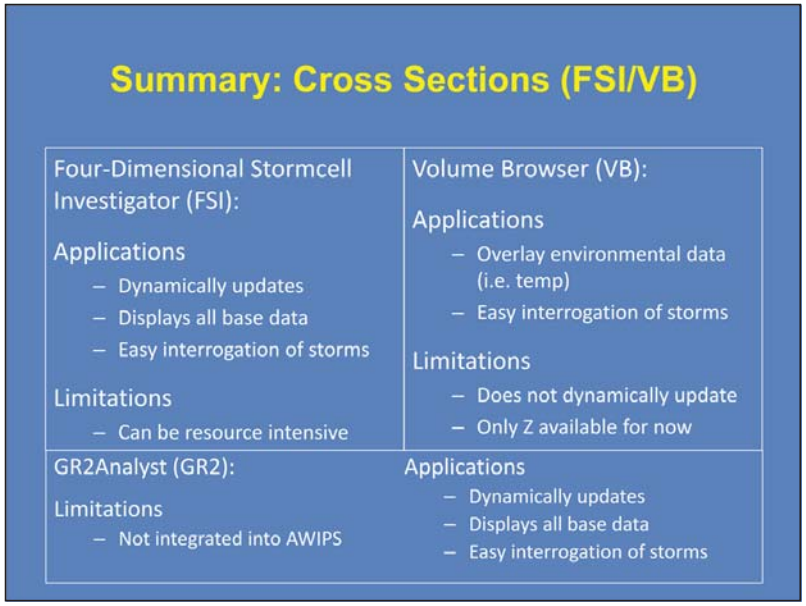
If you choose FSI or GR2Analyst to display cross-sections, the display will update dynamically as you edit your baseline. This feature allows you to quickly change from storm to storm or alter your view of a particular storm. This dynamic update feature can be highly beneficial during high end events where you need to interrogate multiple storms quickly.



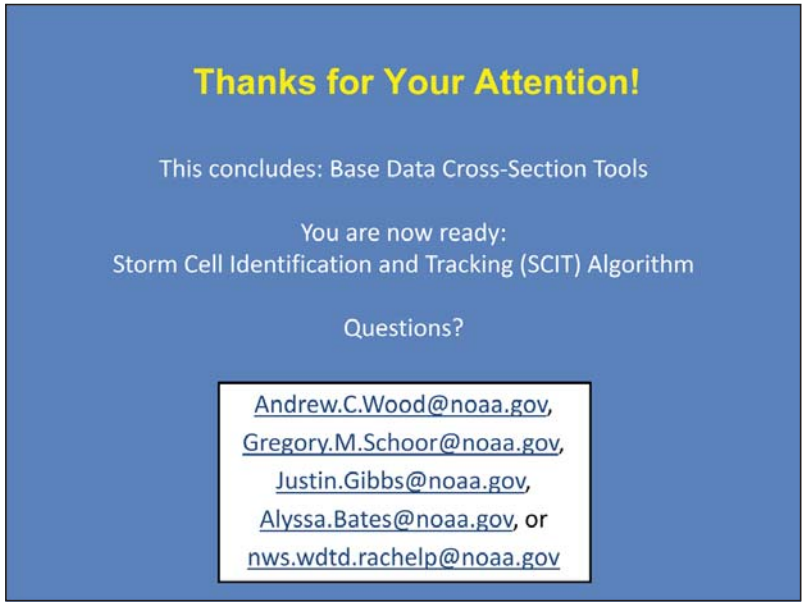
The RPG only generates cross-sections for Reflectivity and Velocity. FSI and GR2Analyst can display all base data in a cross-section, which can be very handy if you are trying to identify a feature, like the melting layer, as in the Correlation Coefficient image shown.



Lastly, the Volume Browser cross-sections allow you to overlay environmental data directly on the display. While FSI lacks this specific capability, you can adjust the CAPPI height to match heights of specific temperature levels and manipulate the display based on the environmental data at your disposal. GR2Analyst has the ability overlay surface data and other GIS shape files, too, but it doesn't provide the same capabilities as the volume browser cross-sections.



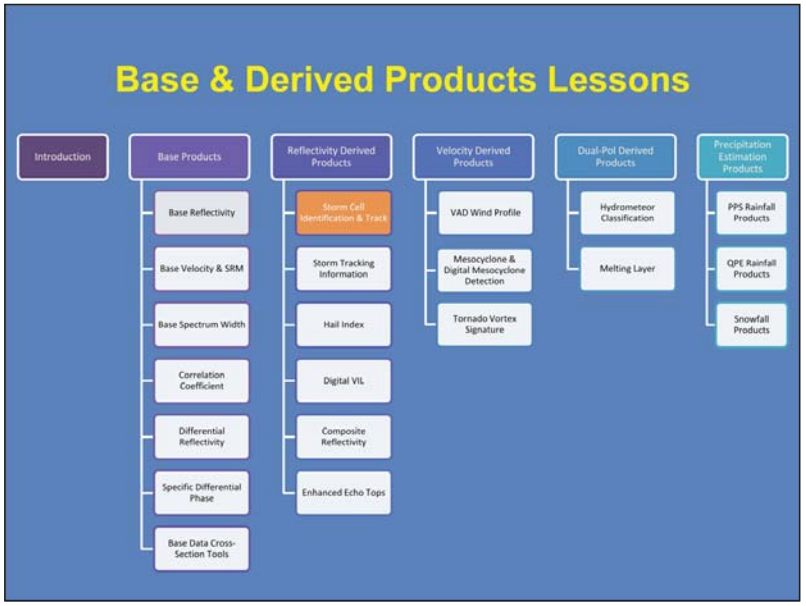
In summary, FSI, Volume Browser, and GR2Analyst are all useful tools for doing cross-sectional analysis. There are several advantages to these tools. First, all base data are available (at least with FSI and GR2 Analyst). It's easier to update the data when needed. You can also include environmental data into your analysis (at least with the AWIPS-2 tools). While there are limitations with these tools, most of the limitations exist with any cross-sectional tool you would use for radar analysis.



If you have passed the quiz, then you have successfully completed this lesson. You are now ready to move onto the next lesson (on the SCIT algorithm). If you have any questions, please contact us using any of the e-mail addresses listed below.



Welcome to this lesson on the Storm Cell Identification and Tracking (or SCIT) algorithm. This training is part of the Base and Derived Products topic in the Radar and Applications Course. Let's get started.



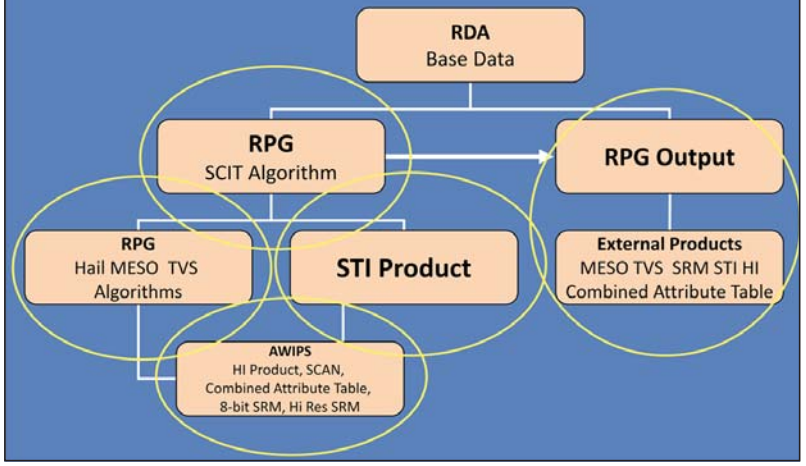
Here is a roadmap for the lessons in this topic. This lesson on the SCIT algorithm, shown shaded in orange, is the first in the Reflectivity Derived Products Section of this topic.

Learning Objectives

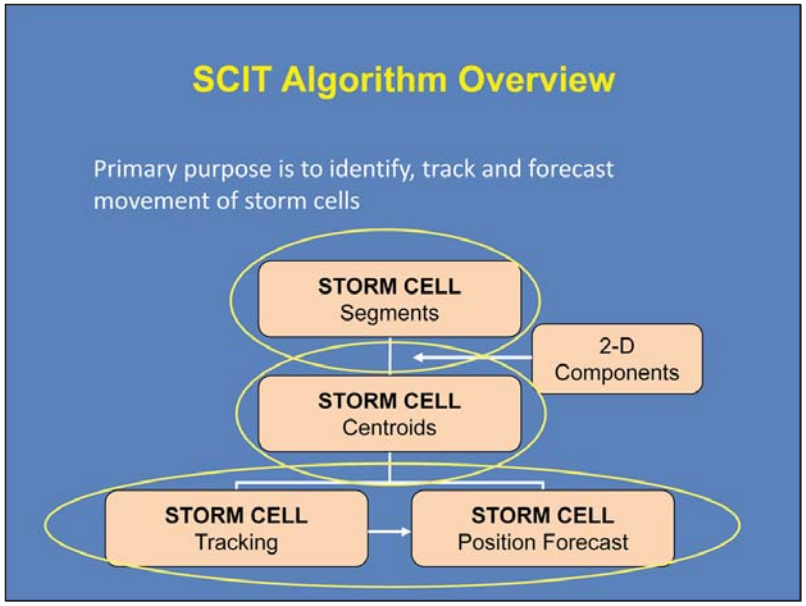
1. Identify the primary role of the SCIT algorithm and how it relates to other RPG algorithms
2. Identify the key storm traits the Storm Cell Identification and Tracking (SCIT) algorithm characterizes
3. Identify the outputs of the SCIT algorithm and the other RPG algorithms/products that use these outputs

Here are the learning objectives for this lesson. Take a moment to read these and advance to the next slide when you are ready to proceed.

SCIT Data Flow: It's Everywhere!

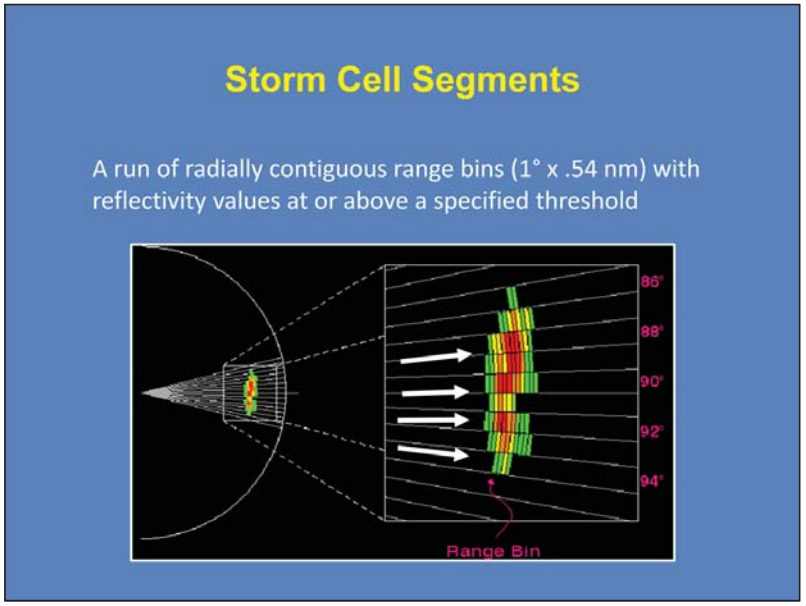


We start our discussion on algorithms with SCIT because it plays an important roll in so many other algorithms at the Radar Product Generator (RPG). As base data flows into the RPG, the SCIT algorithm (Johnson et al., 1998) processes the data and produces the Storm Track Information product (which will be discussed in detail in a later lesson). The information in this product is then used as input for other algorithms, such as the Hail, Mesocyclone, and TVS Detection Algorithms. The information are also displayable in overlay products such as Hail Index, SCAN, and Composite Reflectivity (in the Combined Attribute Table). So, SCIT is important because it impacts the performance of several other algorithms and the appearance of several products (including those that are accessible to our customers and partners).



The SCIT algorithm will input Reflectivity data and process it into building blocks called segments. Segments are one-dimensional reflectivity objects of comparable value. Segments, which are oriented radially, are then combined with comparable, adjacent segments into two-dimensional reflectivity objects called components. Once all of the components are defined for each elevation slice, the algorithm compares the components on each tilt to correlate the components with each other and calculate a cell's centroid location. The cell's centroid is key because it will be used to characterize storm traits, track the storm's motion, and even forecast future locations for the storm.

Let's discuss each step in a little more detail so you can better understand how this process works.



Storm cell segments are the fundamental traits that the algorithm identifies. A segment is defined as a run of contiguous range bins oriented down a specific radial. The SCIT algorithm uses legacy Reflectivity data resolutions of 1 deg by 0.5 nm for its range bins. As the algorithm processes the data on a radial, it looks for data greater than or equal to a specified threshold.

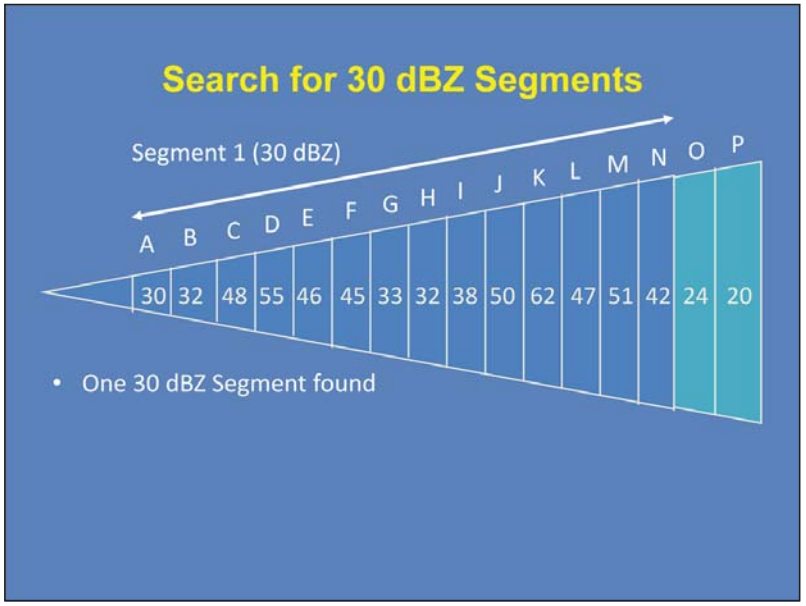
The graphic on the slide shows Reflectivity data plotted with the individual radials and range bins demarcated for emphasis. The arrows indicate how the algorithm processes the data radial by radial.

Storm Cell Segments: ROC Controlled Adaptable Parameters

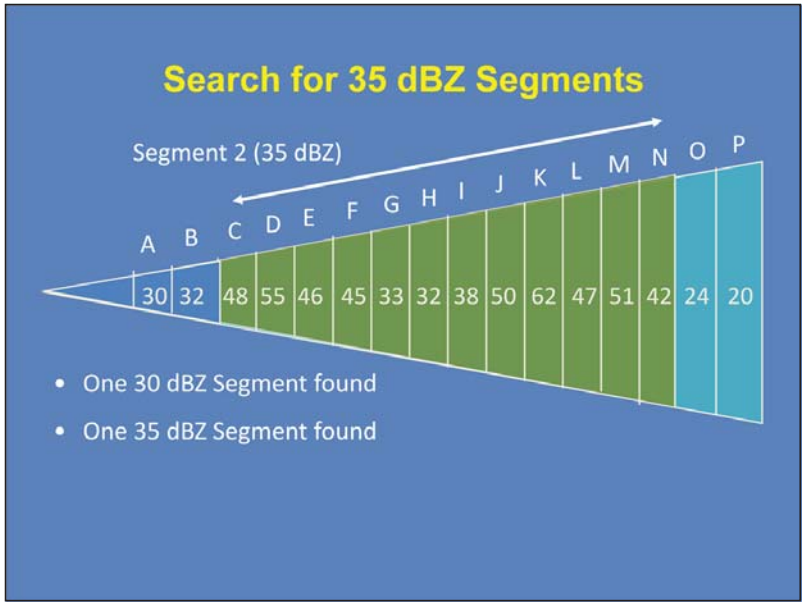
- Segments of up to 7 minimum reflectivity thresholds
  - 30, 35, 40, 45, 50, 55, 60 dBZ
- Minimum segment length
  - 1.9 km/2 bins
- Algorithm has a dropout forgiveness factor:
  - Up to 2 range bins
  - Reflectivity ≤ 5 dBZ below minimum reflectivity threshold

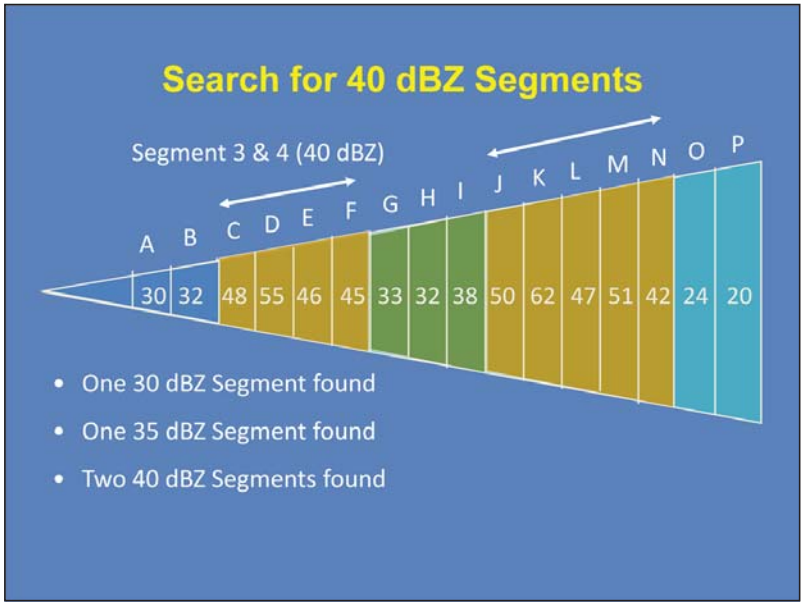
The SCIT algorithm relies on adaptable parameters that are controlled by the Radar Operations Center. These parameters are: minimum reflectivity thresholds and minimum segment length. The algorithm uses seven minimum reflectivity thresholds from 30 to 60 dBZ in 5 dBZ increments. This means the radar must detect 30 dBZ in a range bin before a segment can start. The minimum segment length is 2 range bins (which is at least 1.9 km). So, you would need two range bins next to each other of at least 30 dBZ for a segment to be identified.

The algorithm can be forgiving at times. In some circumstances, up to 2 contiguous range bins below the minimum threshold can be included in the segment if their value is within 5 dBZ of the threshold.



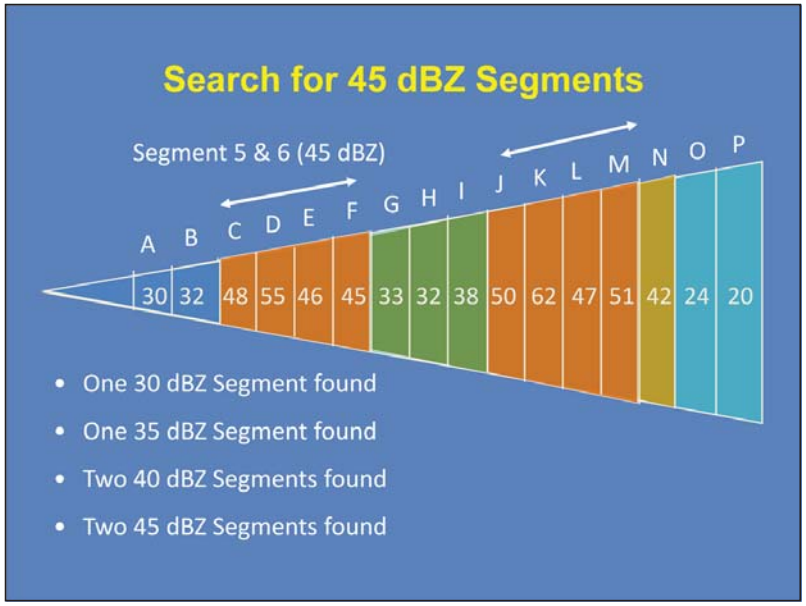
Let's walk through an example just to make sure I am making sense to you. The slide shows you a segment of a single radial of Reflectivity data. Once the algorithm finds a couple of bins at or over 30 dBZ, so it starts identifying a segment. The algorithm proceeds looking at each range bin down radial until it encounters an out of range bin (which in this case is the 24 dBZ value in the location labeled “O”). The initial segment, or 30 dBZ segment, then terminates at this point.



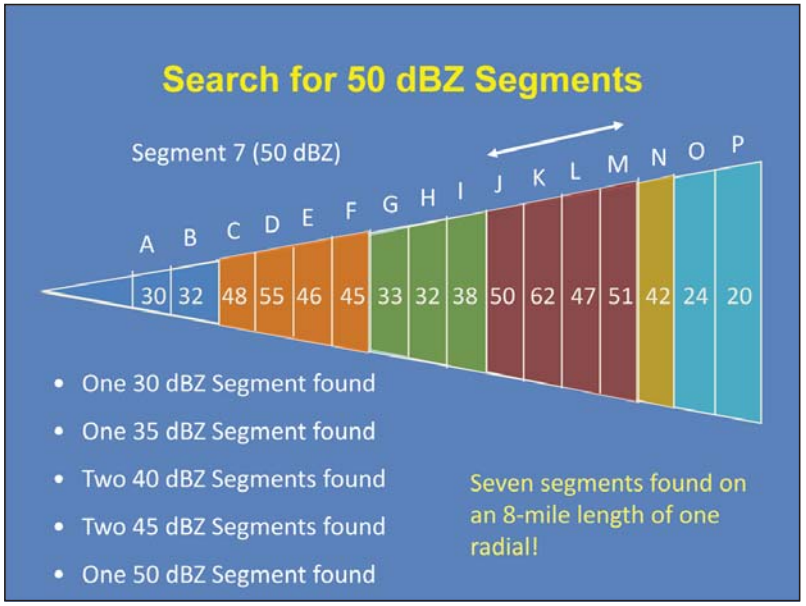


The algorithm proceeds to the next segment threshold by looking at each of the 30 dBZ segments. The algorithm attempts to isolate and identify the higher reflectivity cores in each segment, so it will continue to incremental increase the minimum reflectivity threshold with each pass. So, using the 35 dBZ threshold, the next segment begins in bin C where we have a 48 dBZ value and ends at the same point as the previous segment.

The process gets a little more interesting at the 40 dBZ threshold. A segment starts at bin C again, but now terminates at bin F because of the drop in values in the bin G. However, the algorithm continues down the radial and finds another segment starting in bin J and ending at bin N.

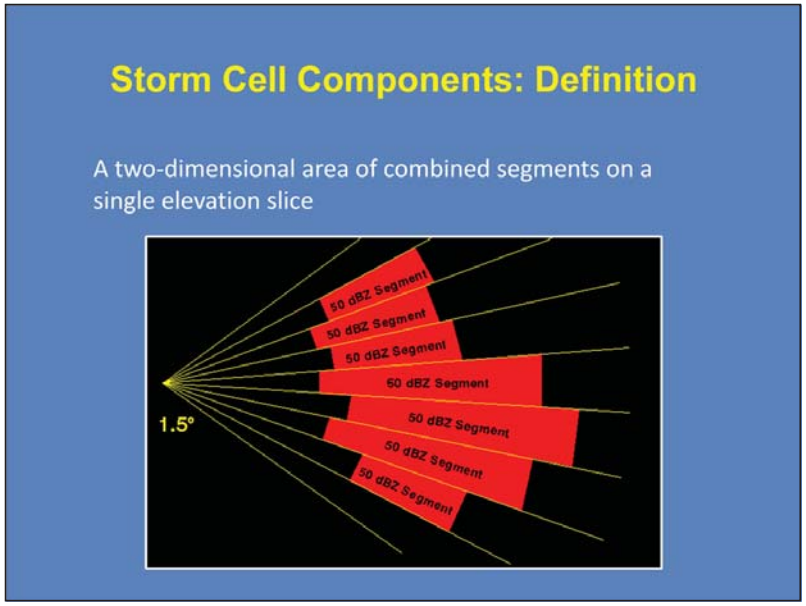


Continuing on at 45 dBZ, the algorithm furthers its attempt to isolate the higher reflectivity cores in these two segments. Note, the value in bin N is within the 5 dBZ dropoff range, but there are more than two range bins between locations of 45 dBZ values. Hence, this range bin is not included in the segment.

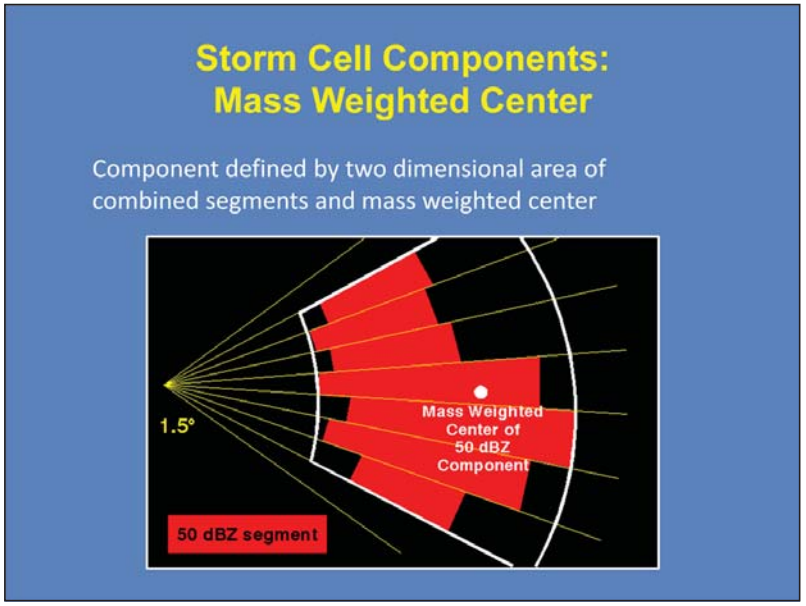


Finally, the algorithm identifies its last segment on this portion of the radial in bins J through M. Notice how bin L is included in the segment with a value 3 dBZ below the threshold because it is located between two 50+ dBZ bins.

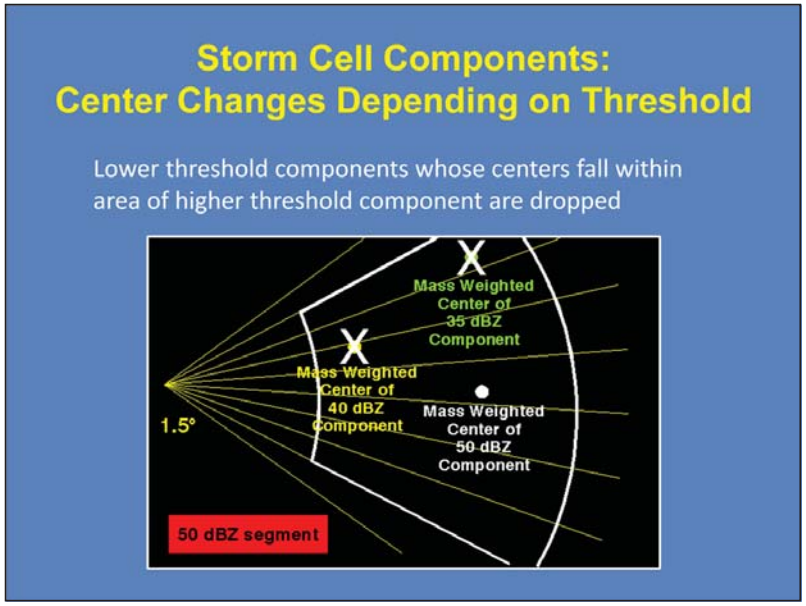
The end result is the algorithm found 7 segments on this small part of a single radial!



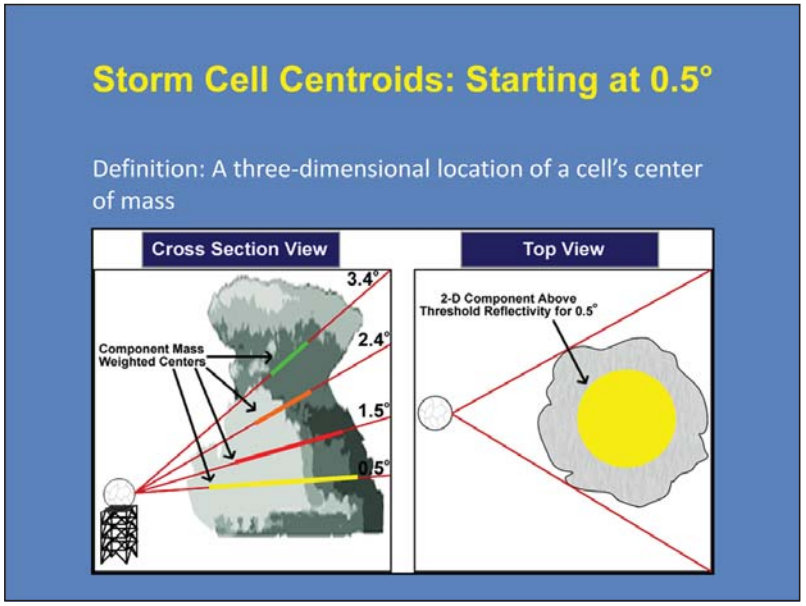
Once all the segments are identified on a particular elevation slice, the algorithm can start associating adjacent segments with the same threshold values into two-dimensional components. These components are traits that indicate the two-dimensional representation of a cell at a particular elevation angle.



As components are assembled, the algorithm looks for the mass-weighted center of the component. This image shows the 50 dBZ component with its mass weighted center identified. The term “mass-weighted” comes from the reflectivity values “pulling” the storm centroid in that location. This process prevents the algorithm from plotting the storm centroid simply at the geographic center of the component, which could be misleading to the forecaster looking at the algorithm output.

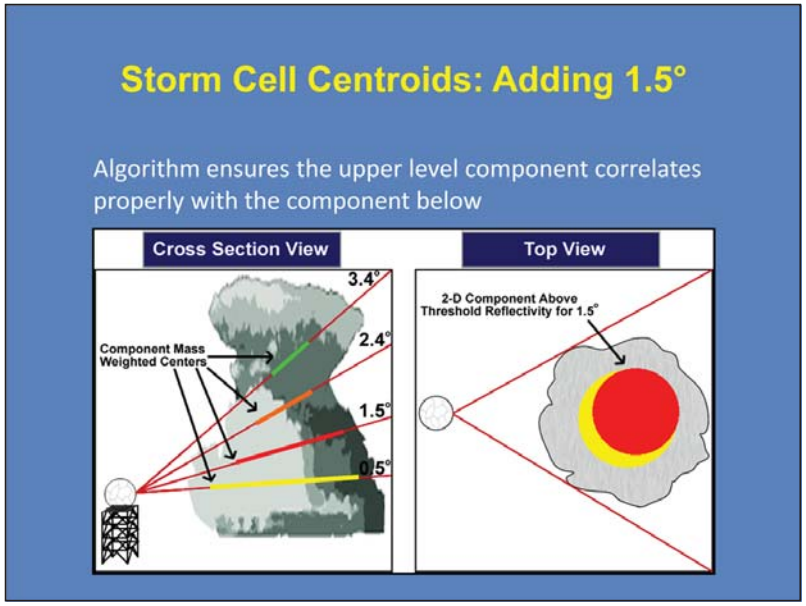


Many times, the mass-weighted centers will move as the reflectivity thresholds increase. The mass-weighted center that ultimately leads to identifying the storm centroid will use the location for the highest threshold value for the component. In this case, the component is built out of contiguous 50 dBZ segments. Hence, the 50 dBZ mass-weighted center is used as we proceed to the next step.

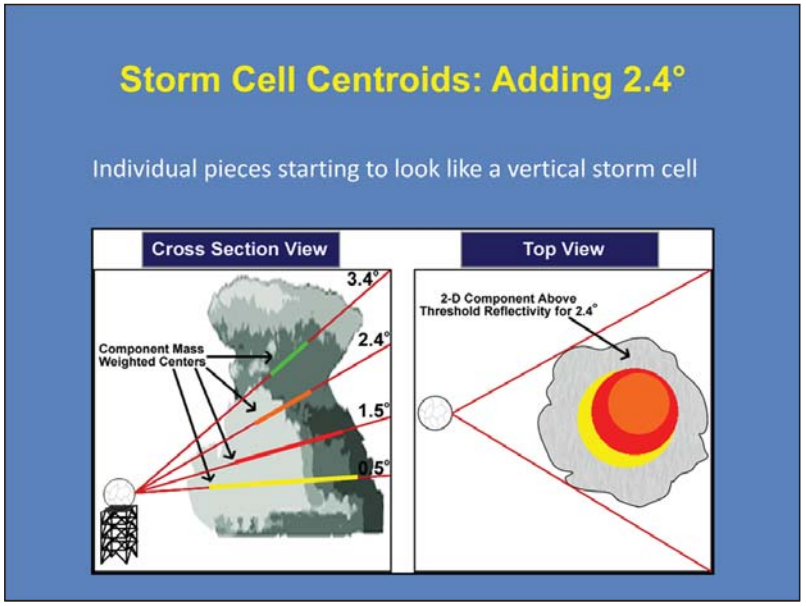


The last trait that the algorithm will identify are centroids. Centroids are identified by looking at the mass-weighted centers of components on each elevation slice and correlating them together vertically to create a three-dimensional object.

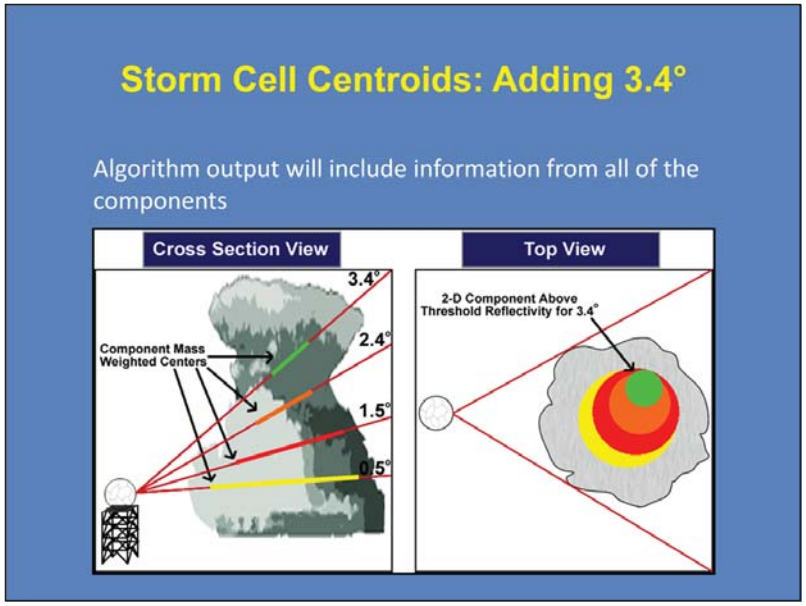
Let's walk through this process, by comparing mass-weighted centers from several elevation angles (shown in the graphic on the left). The image on the right shows the component from the 0.5 degree tilt. Let's add components from other slices and see what we come up with.



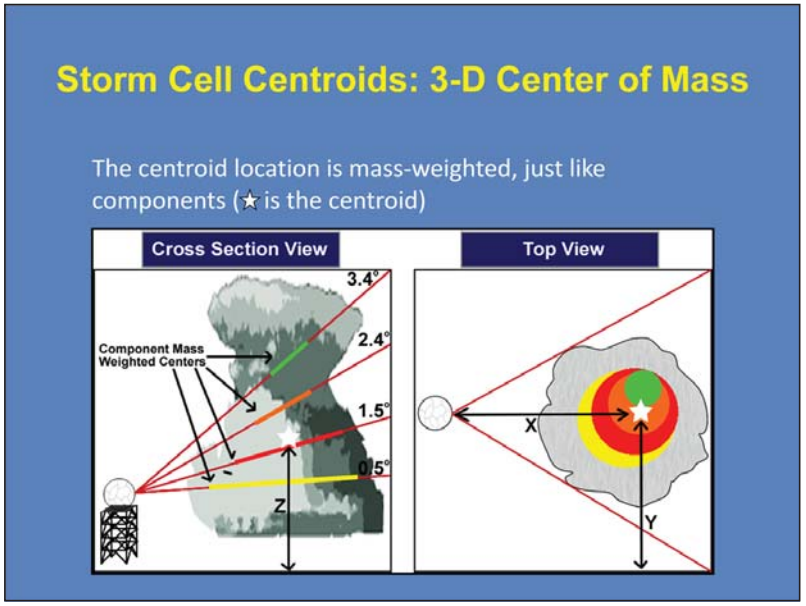
Adding the 1.5 degree component, there appears to be some agreement between the two. The algorithm will perform some checks to make sure that two components are found within a particular search radius of each other. Specifically, the algorithm ensures the upper level component correlates properly with the component(s) below it.



As we move up to 2.4 degrees, we add the next component. You can see how the individual pieces start looking more like a vertical storm cell, algorithmically speaking.



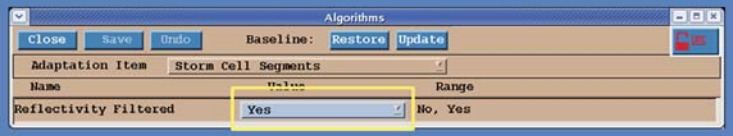
We include the last component, found on the 3.4 degree tilt. The algorithm output for this storm will include the information calculated from all four of these components. If the algorithm found components on higher tilts, it would include those, too.



Once the three-dimensional structure of the storm components has been analyzed, the storm centroid is calculated in the much the same way as each component's mass-centered weight. The storm centroid will be influenced by the location of higher reflectivity values. So, don't expect the centroid to be in the three-dimensional geographic center of the cell. The centroid will have an azimuth, range, and height associated with it since it's a three-dimensional feature. The RPG will use the storm centroid to plot the location of the cell in this and other algorithm's outputs you see in AWIPS.

Adaptable Parameters: Reflectivity Filter

- Recommend it be turned on ("yes")
- Does not impact high reflectivity values used in defining cell attributes

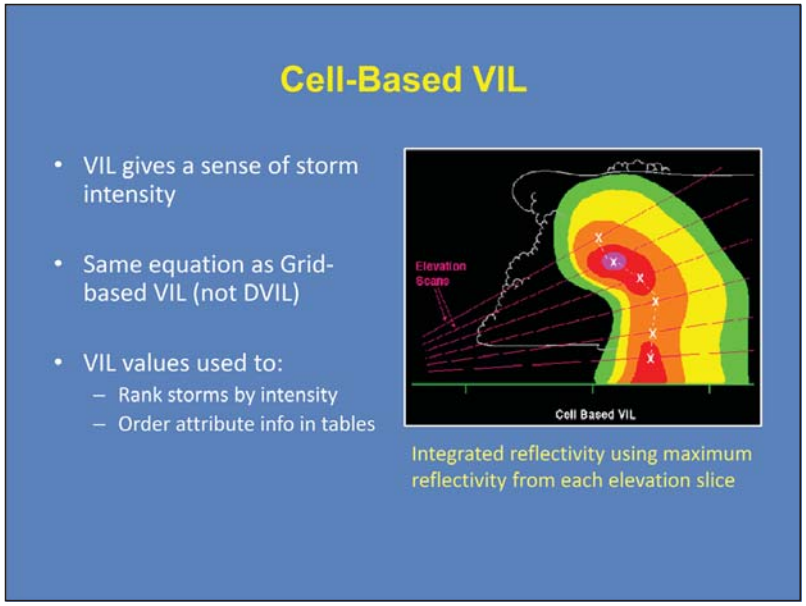


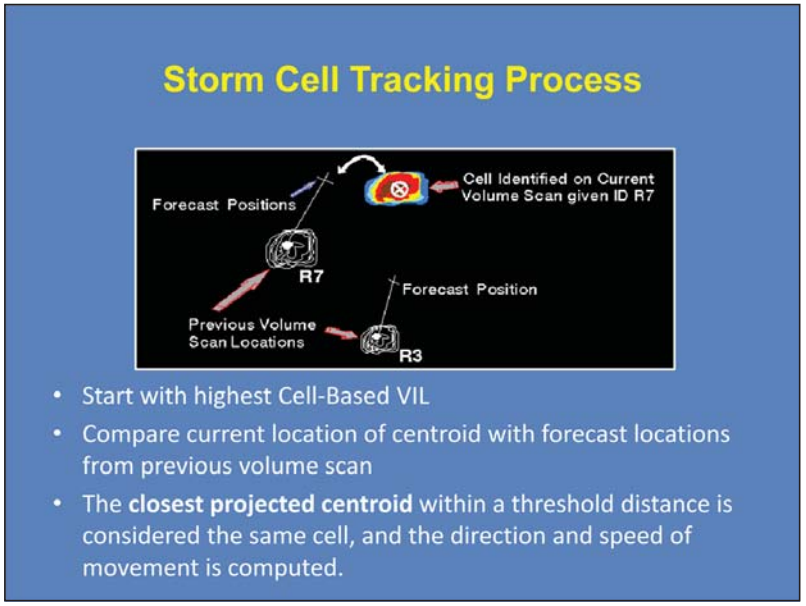
There's one adaptable parameter your office has control over under Storm Cell Segments called "Reflectivity Filtered" that we want to highlight. The purpose of this filter is to eliminate spurious noise in the Reflectivity that might mess up the SCIT algorithm. It only filters the values that go into the algorithm, not what you see on your Reflectivity displays in AWIPS. As a result, we recommend you set this parameter to "Yes" to aid the algorithm's performance.

Storm Cell Centroids Output: Available on Up to 100 Cells

- Centroid (in polar coordinates)
- Height of the centroid (ARL)
- Maximum (3-bin averaged) reflectivity
- Height of the maximum reflectivity (beam center point height - ARL)
- Cell Top and Base (beam center point height-ARL)
- Number of components
- Cell-based Vertically Integrated Liquid (VIL)

Once the algorithm has processed all of the data for the volume scan, it is ready to provide a list of all the storms it identified. The algorithm output can include up to 100 different cells. Each identified cell will have an azimuth, range, height (above radar level), the maximum reflectivity value and height, plus the cell base and top heights. The algorithm will also document the number of components identified in a cell, along with a cell-based Vertically Integrated Liquid, or VIL, value.



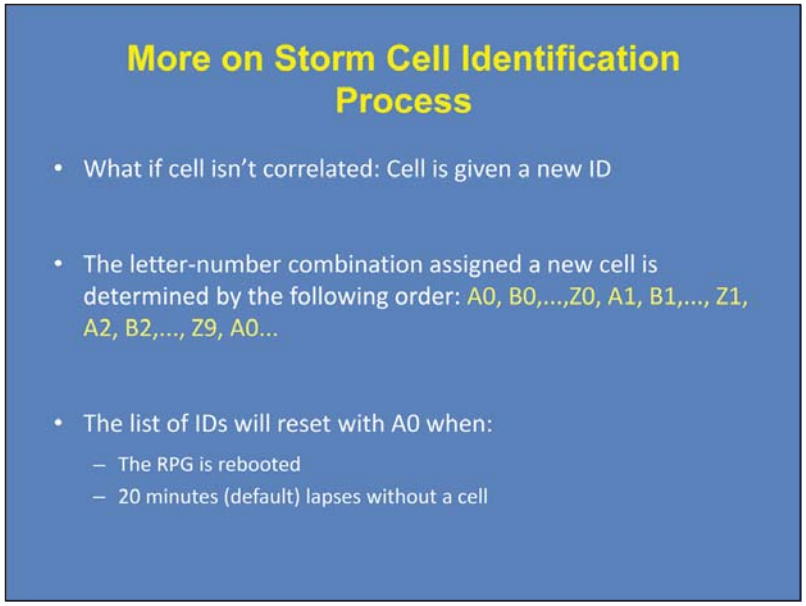


If you are taking the Base and Derived Products lessons in the order prescribed in the RAC curriculum, than you haven't seen a discussion of VIL (actually Digital VIL) yet. So, I want to take a moment to discuss it here.

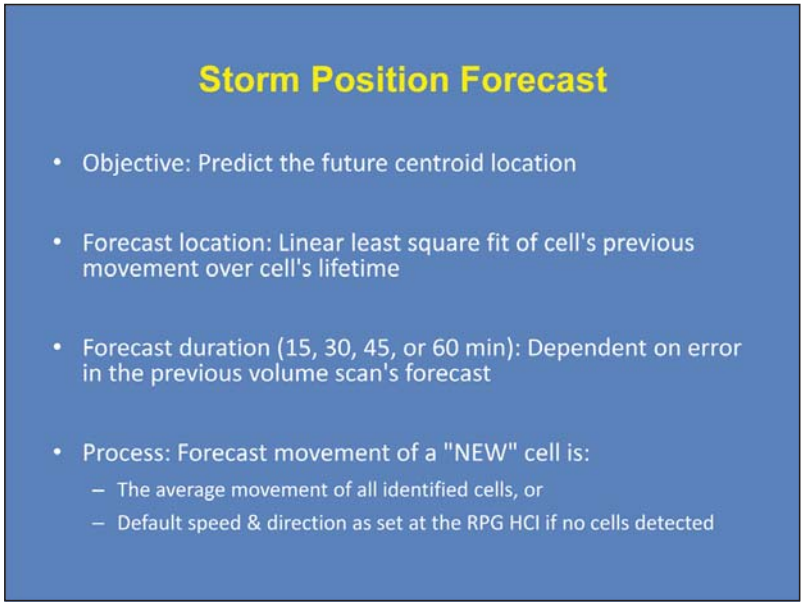
Cell-based VIL integrates the maximum reflectivity values throughout the vertical extent of the cell to give you a sense of the storm's intensity. It's important to note the SCIT algorithm uses the same equation as the original VIL product, not digital VIL. These VIL values are important because the algorithm will use VIL to rank the storms it identifies. Storms with higher VIL values will generally appear before lower VIL storms in the various tables where storm cell information is displayed.

Once all of the cells surrounding the radar have been identified, the SCIT algorithm attempts to correlate them with storms from the previous volume scan. This process is called storm cell tracking. The algorithm processes each cell in order by cell-based VIL values, starting with the highest.

The image on the slide shows an example of this process. We see one cell (shown in color) for the current volume scan, along with two nearby cells from the previous volume scan. The algorithm will attempt to correlate the current cell with one of the previous ones. How you might ask? Well, the algorithm will compare the cells current location with forecast positions (calculated during the previous volume scan) of the old storms. The two older storms have the IDs R7 and R3, respectively, and their forecast positions are indicated on the graphic. While the current storm is not at either of these locations, the algorithm will see if the distance between the forecast position and current location falls within a threshold value. In this case, R7 falls within that threshold. As a result, the current storm is given the ID R7 and the storm motion and forecast position are updated based on the new information.

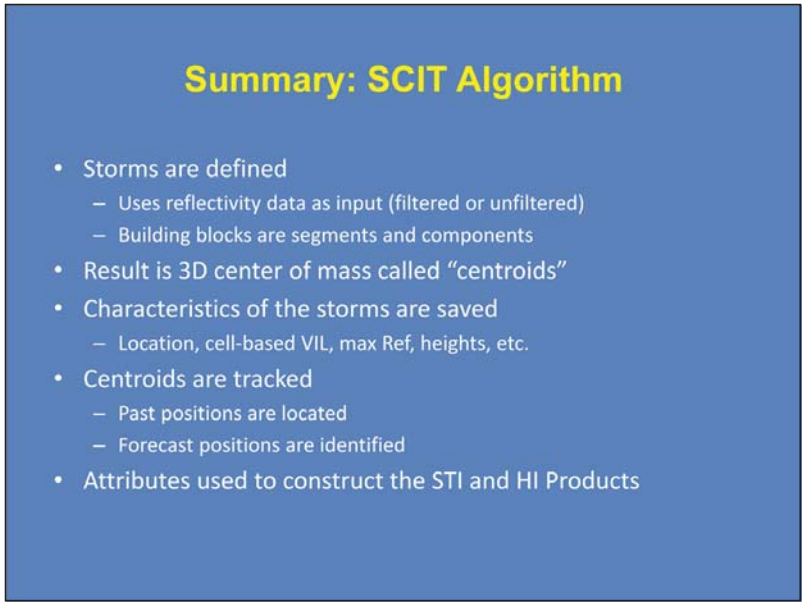


So, what happens if a storm doesn't correlate with any previously existing cells? Well, the cell will be considered a new storm and given a new ID based on a simple letter/number convention. The first identified storm is given the ID of A0. The next storm is given the ID B0, and so on, until the algorithm labels a storm Z0. From there, the IDs roll over to A1, B1, etc. During an active weather event, the storm IDs don't really tell you much about a storms strength, just how long a particular storm centroid has been tracked by the algorithm. The lower the number, the longer the storm has been around. The IDs are eventually reset when the RPG is either rebooted, or when there have been 20 minutes or more without an identified storm.



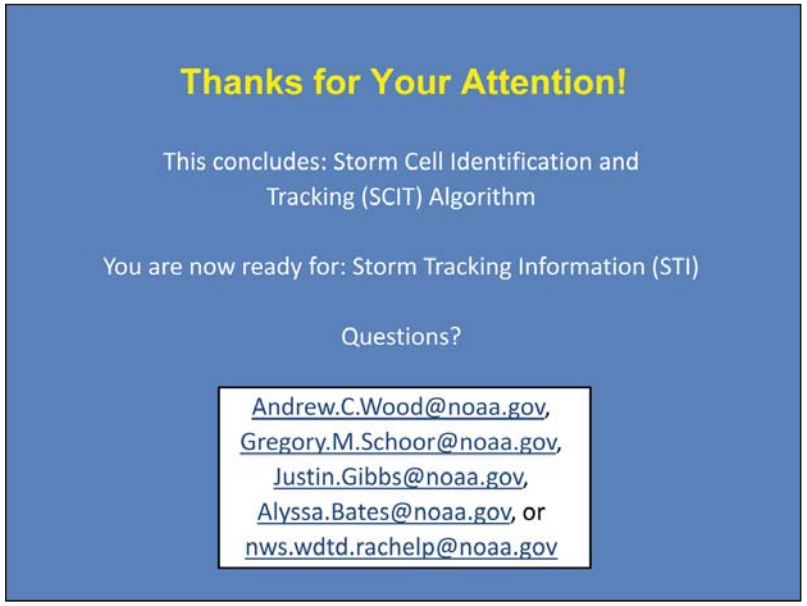
We talked previously about the storm position forecast. Let's look at this process in more detail. This part of the SCIT algorithm attempts to predict the future location of the storm centroid. This forecast is based on a linear least square fit of the storm's previous movement over its lifetime. If a storm has been tracked by the algorithm for several volume scans, this forecast track can be fairly good and can extend out for 15 to 60 minutes, depending upon the error calculated during the least square fit process. Basically, the higher the error, the shorter the forecast duration.

So, you might be asking yourself what happens with cells identified in the last volume scan? Well, new cells are given one of two forecast movements. Most of the time, it will use an average speed from all of the identified cells as a first guess. If the storm in question is A0 (or identified in the same volume scan as A0), then it will use the default speed and direction set at the RPG.



Let's summarize the SCIT algorithm and how it will impact the algorithms and derived products that will be discussed in later lessons. The purpose of this algorithm is to identify storms. It uses reflectivity (which we recommend filtering via your adaptable parameters) as its primary input. The reflectivity data are used to identify one-dimensional, radially aligned segments. Contiguous segments are used to build two-dimensional components that fall on a single elevation angle. Components are then stacked on top of each other to build a three-dimensional centroid. Centroids have a center of mass that is used to locate and track the storm's position. Various characteristics of the storm (including location, cell-based VIL, maximum reflectivity, and centroid height) are saved by the algorithm. Storm centroids are then tracked by comparing these data to data (specifically forecast position) from storms identified in the previous volume scan and correlating the storms when they appear to be the same. This information is then passed along to other algorithms and derived products (such as the Storm Tracking Information and Hail Index products).

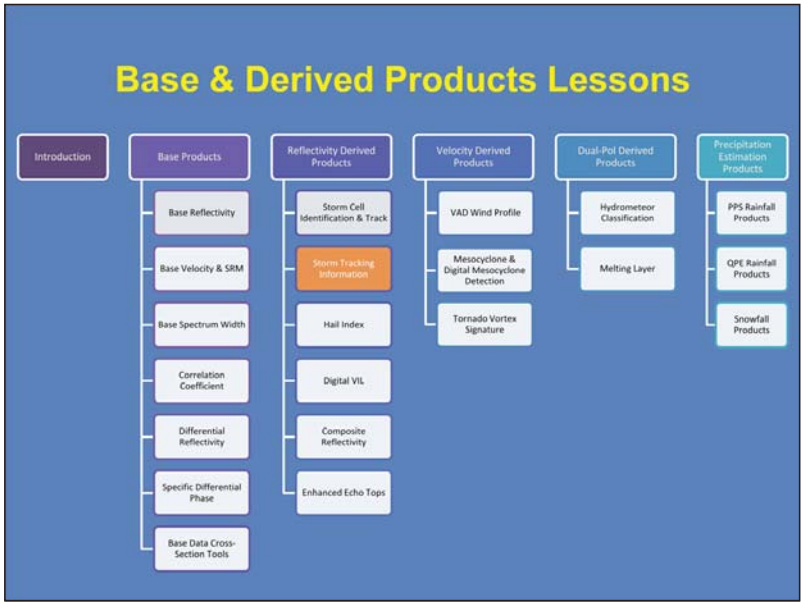
Please take a few moments to complete the quiz on the next slide.



If you have passed the quiz, then you have successfully completed this lesson. You are now ready to proceed to the next lesson on the Storm Tracking Information (STI) product. If you have any questions, feel free to e-mail the contacts listed below.



Welcome to this lesson on the Storm Tracking Information (STI) product. This training is part of the Base and Derived Products topic in the Radar and Applications Course. Let's get started.



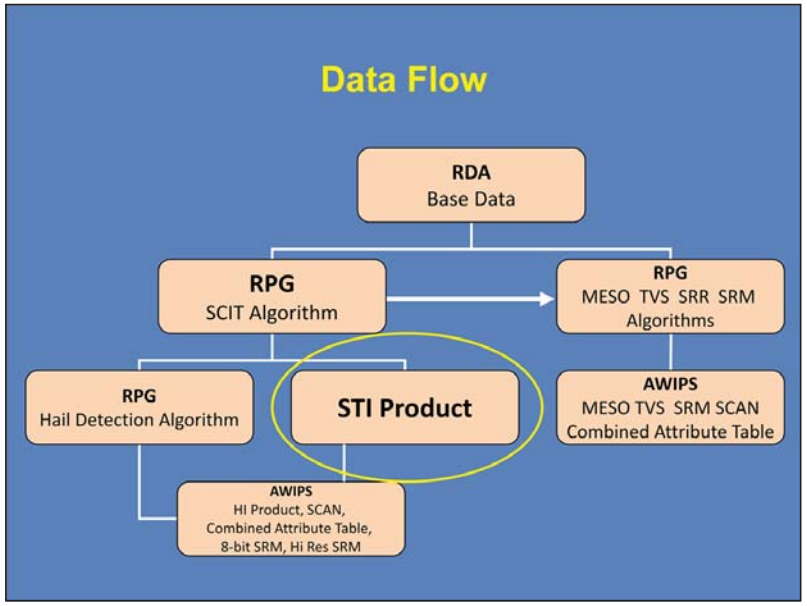
Here is a roadmap for the lessons in this topic. This lesson on the STI product, shaded in orange, is part of the Reflectivity Derived Products Section of this topic.

Learning Objectives

1. Identify the key attributes and display characteristics (including how the Radar Display Controls impacts them) of the Storm Tracking Information (STI) product in AWIPS
2. Identify the limitations & applications of the Storm Tracking Information product

Here are the learning objectives. Take a moment read them and advance the slide when you are ready.

Data Flow



The Storm Cell Identification and Track algorithm (or SCIT) lesson discussed how the WSR-88D RPG identifies storms as a three-dimensional object. This lesson will talk about one product that is produced from that algorithm's output: the Storm Tracking Information, or STI, product.

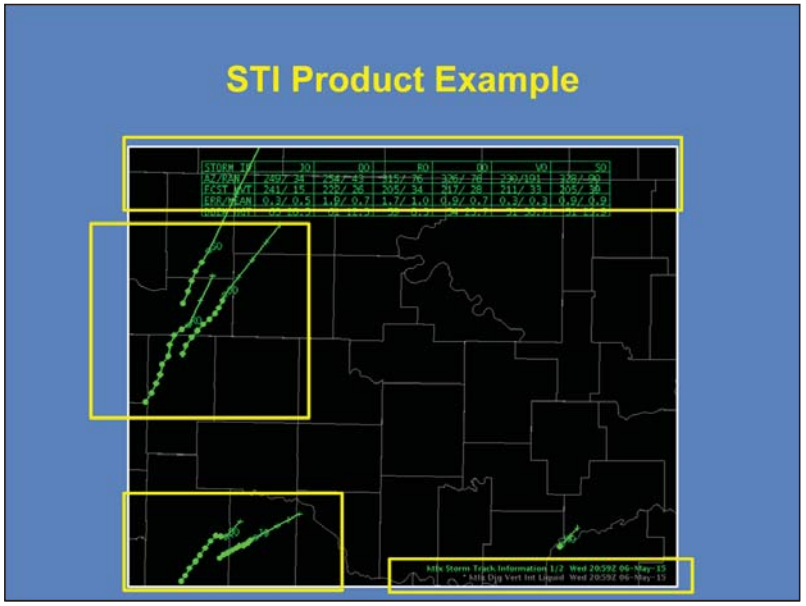
Product Description

Displays up to 100 cells identified by the SCIT Algorithm along with actual past positions and forecast positions with the following symbols:

- Centroid location
- Past position (volume scan increments with a line between each symbol)
- Forecast position (15 minute increments with a straight line connecting all forecast positions)
- Cell moving less than 5 knots

The STI product allows you to see the output of the work done by the SCIT algorithm. The product will display data on up to 100 cells identified by the SCIT algorithm, including the past, current, and forecast locations. Each of these items are identified using different symbols. The current centroid location is marked by a circle with an x in it. Positions from previous volume scans are marked by a filled circle connected by lines. Forecast positions will appear with a plus symbol (spaced out in 15 minute increments) with a straight line connecting locations. Cells moving at less than 5 knots will have an additional circle around the centroid location marker, similar to how a ASOS site would report calm winds on a surface map.

STI Product Example



Here is what the Storm Track Information product looks like. We see several storms with various amounts of past and forecast positions. There is also an attribute table which gives additional information about each cell. More on that info in a little bit. Lastly, you can see in the product legend that the STI product will have pagination info about the table. In this case, page 1 of a two-page table is visible.

STI Attribute Table

STORM ID	J0	00	R0	00	V0	S0
AZ/RAN	249/ 34	254/ 43	315/ 76	326/ 76	230/191	328/ 90
FCST MVT	241/ 15	222/ 36	205/ 34	217/ 28	211/ 33	205/ 30
ERR/MEAN	0.3/0.5	0.0/0.2	0.2/0.1	0.0/0.0	0.3/0.3	0.0/0.0
DBZM HGT	63 28.5	61 12.3	59 8.3	54 23.7	51 36.7	51 23.9

- STORM ID: Cells in order of Cell-based VIL
- "NEW": Appears on FCST MVT line for 1st scan of new cell
- ERR: Distance (nm) from forecast position to actual position
- MEAN: Mean error (ERR) during the cell's lifetime
- DBZM: Maximum reflectivity
- HGT: Height of the maximum reflectivity in kft ARL

Let's take a look at the Storm Track Information attributes table. Each storm that has been identified will be listed on one of the pages of this table. The Storm ID is the letter/number combination discussed in detail during the SCIT lesson. Cells will be listed in order of cell-based VIL. In this example, J0 has the largest cell-based VIL and gets listed first.

The azimuth and range listed is for the current centroid location. It's not the location of the largest hail, high winds or heavy rain, but rather the storm's 3-D center of mass. Forecast movement shows the expected direction and speed for a particular cell. If it's a brand new storm, you will see the word "NEW". The next line, "ERR/MEAN", is a measure of how well the algorithm thinks it's doing tracking that particular storm. The error number indicates the distance of current position from previous forecast positions. The mean number is the average of all the errors during the lifetime of the storm. The last line shows the maximum reflectivity for that storm centroid, and the height (in thousands of feet) where that value is located.

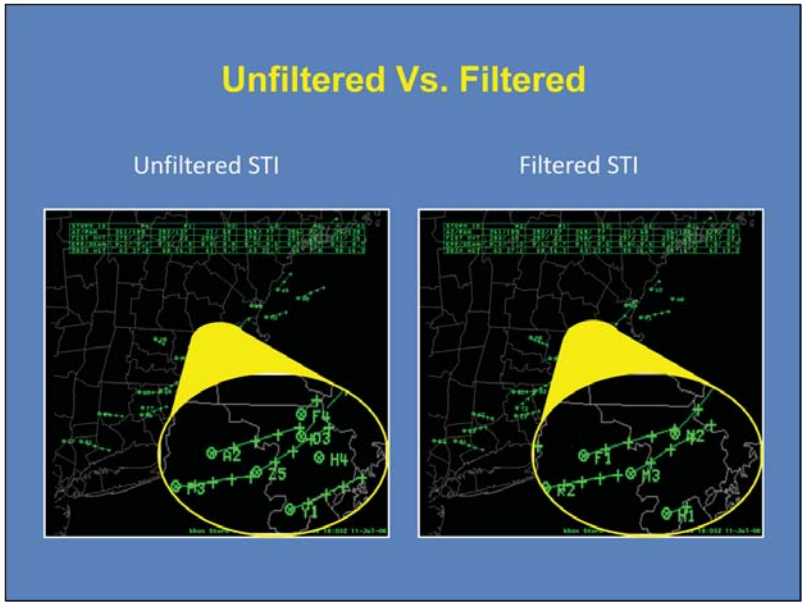
Radar Display Controls: Displays Can Be a Mess!

- Display can be cluttered if numerous cells present
- Limit number of cells displayed on product
  - Sorted by Cell-based VIL
- AWIPS doesn't annotate display with this setting!
- Can select to not display past tracks or future positions

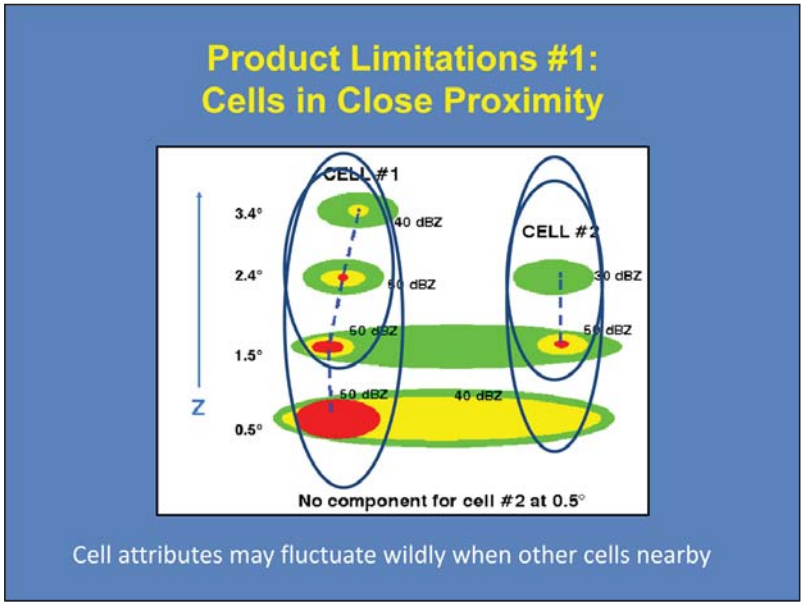


If the RPG identifies numerous, persistent storm centroids, the STI product can become very cluttered. Users have the ability to control the appearance of the STI product using the Radar Display Control GUI. At the top of the interface, a slider bar limits the number of storm centroids that are displayed. You can choose any number between 0 and 100. Remember that the storm centroids that are displayed will depend on their cell-based VIL. So, if you choose to display only 10 cells, then you will only see the ones with the highest cell-based VIL values. If storms are cycling in intensity on short-time scales, storm centroids could disappear and reappear from volume scan to volume scan if you have the display value set very low. Also, there is no way to tell in AWIPS how many cells are not being displayed because of this cut off value. So, if you set the value to 10, but the SCIT algorithm identified 30 storms, then AWIPS will not tell you that 20 storms are not being displayed.

You can also change how past and future cell tracks are displayed. You can choose to display past & future tracks, each individually, or neither. These settings can further help de-clutter your display.

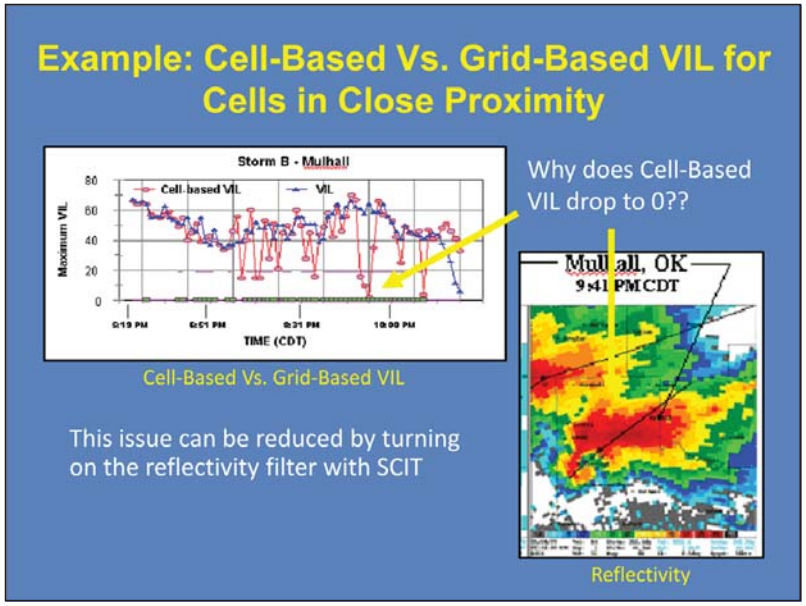


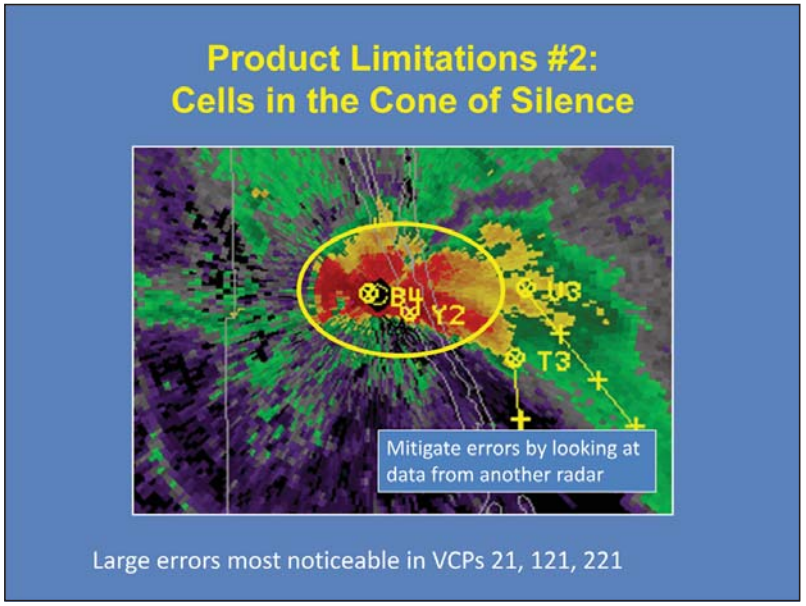
Recall from the SCIT lesson that users have the option to filter the reflectivity data input to that algorithm. The option that you choose will affect the appearance of the STI display. Let's look at an example with unfiltered data shown on the left and filtered data on the right. We will focus on the areas inside the yellow oval. You can see more centroids have been identified in the unfiltered data. While the unfiltered data gives you more storm centroids, the filtered STI likely shows you better centroid locations, making it a more reliable product.



Now that we have covered the basics of the Storm Track Information product, let's talk about its limitations. Errors can sometimes occur in storm centroid trait calculations, or storm centroids could not be identified altogether, if cells are in close proximity to each other. Let's use an example to illustrate this limitation.

We see two cells, Cell 1 and Cell 2, that are made up of storm components, some of which they have in common, on multiple elevation angles. The SCIT algorithm doesn't "share" components between cells in close proximity. So, the large component at 0.5 degree is assigned completely to Cell 1. Even though a human can clearly see an echo at 0.5 degree for Cell 2, the SCIT algorithm truncates the bottom of this cell at 1.5 degrees. This algorithmic truncation will impact the attributes for Cell 2 that you see in STI, such as cell-based VIL. If cell-based VIL is reduced significantly, it would reduce were Cell 2 should fall in the algorithm's ranking system. Worse yet, if the structure of the 0.5 degree component changes significantly in the next volume scan, the algorithm could attach this component to the Cell 2 centroid next. The end result is that storm centroids in close proximity can appear to oscillate abnormally in intensity (and rank) due to this artifact in the SCIT algorithm.

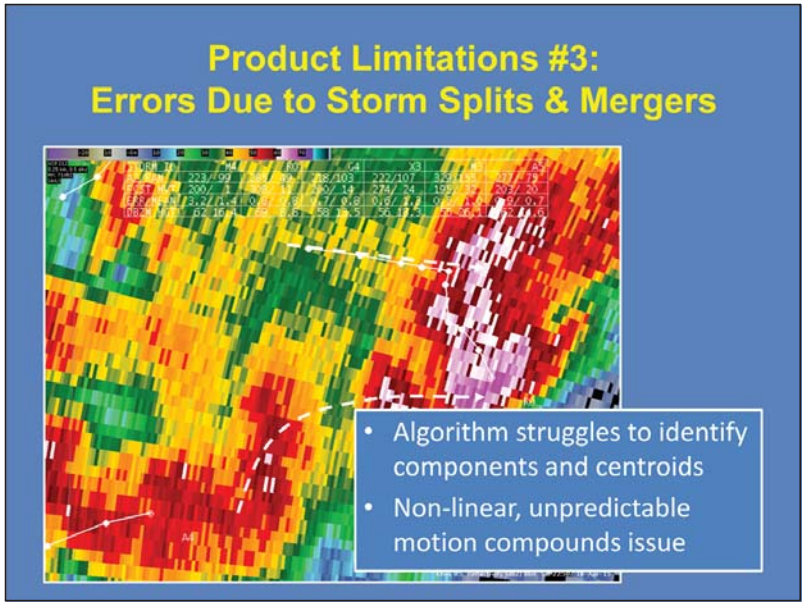




Here is another example of the errors that can occur when cells are in close proximity. This graph shows cell-based and grid-based VIL values for a particular storm over time (shown on the right). The cell-based VIL, shown in red, tends to vary much more than the grid-based VIL for this storm. It even drops all the way to 0 kg/m^2 a couple of times. One of those dropouts coincides closely with the reflectivity image shown at the right. Now, do you really think that storm should have a cell-based VIL of nearly 0? I don't think so. Clearly there is some storm segment shenanigans going on because two or more centroids are close together.

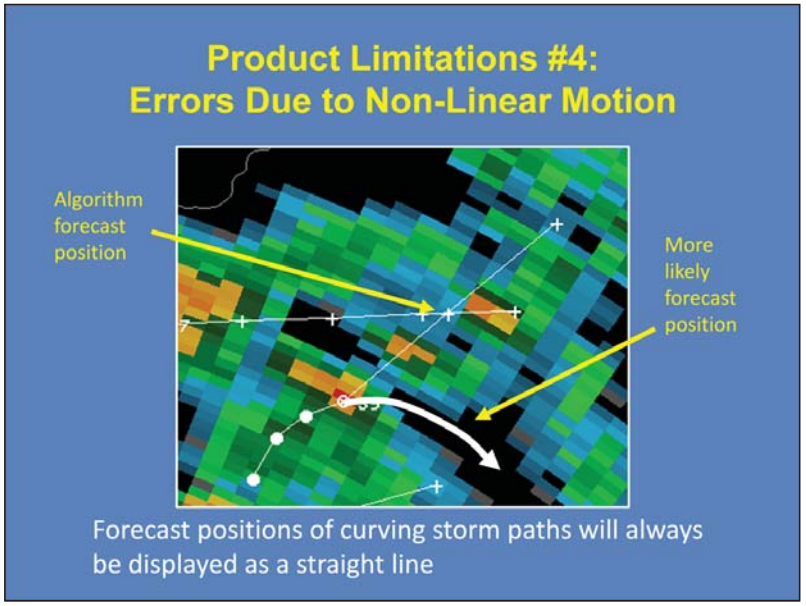
Some of these issues can be avoided by using filtered reflectivity data. However, using that adaptable parameter will not always prevent this issue from popping up.

A second limitation is large errors in cell attributes can occur for cell located in the cone of silence near the RDA. This limitation is most noticeable when using VCP 21, 121, or 221 because they use fewer elevation cuts, with larger gaps, than other precipitation mode VCPs. You can mitigate this issue by viewing data from a nearby radar (if one is available). That secondary radar should provide better cell attribute data, even though the storm is located further from the radar.

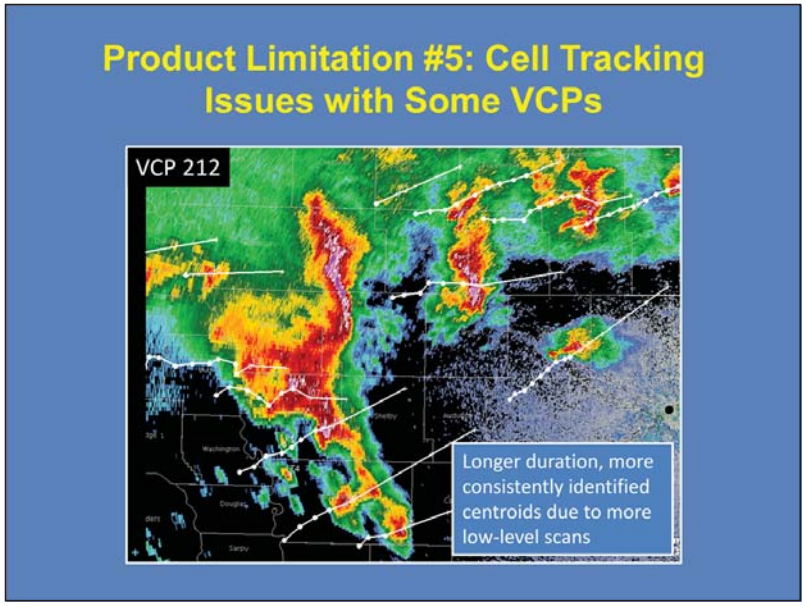


A third limitation occurs when you have a large storm that splits or smaller storms that merge. As these processes occur, the algorithm can struggle to identify the storm components and centroids in a manner that makes the most sense to a human forecaster. Combine this issue with the non-linear, and sometimes unpredictable, cell propagation and redevelopment that occurs during these transitions, and you may be left scratching your head at the results.

In the example on the slide, we have two cells (labeled R0 and B4) that are on a collision course. As we move forward in time, you can see that B4 is moving initially to the north, but curving to the right. SCIT starts getting confused as they near each other and treats them as one centroid. As the B4 ID disappears, R0 starts taking a hard jog to the south. The core for this cell still is moving to the east at 0.5 degrees. You can see as we reach the last of the frames that the cores, which started merging aloft several volume scans ago, has now more or less merged at 0.5. The frames during the pre-merger phase may confuse you if you don't realize what is happening. If you use the STI regularly, you will need to catch these spurious identifications as they happen or they can lead you down the wrong path.

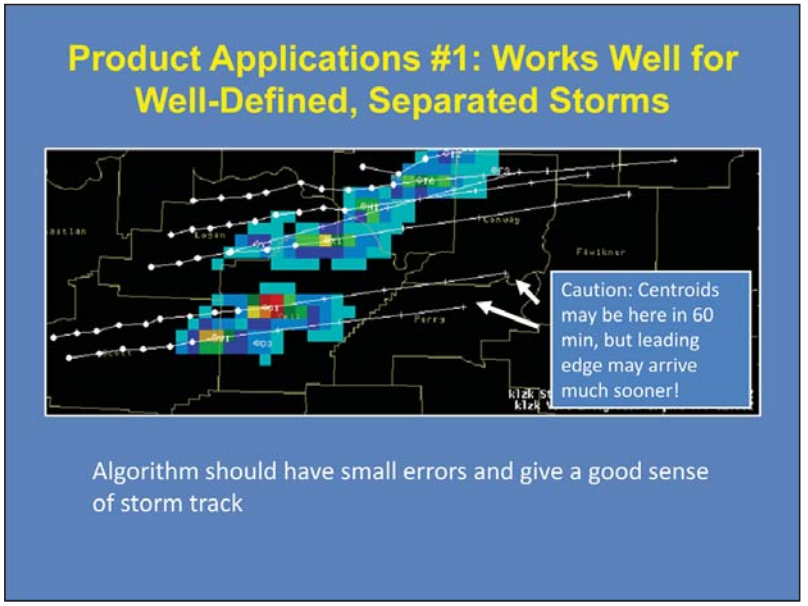


A fourth limitation occurs with cells experiencing non-linear motion. For each storm centroid, the SCIT algorithm generates a linear storm motion. So, the forecast positions shown in STI will always be a straight line. In the example on the slide, cell S9 clearly has been following a right-ward curving path. However, SCIT has assigned a linear forecast location for this cell. In reality, the curved arrow overlaid on the image gives a truer sense of the forecast path of this cell.



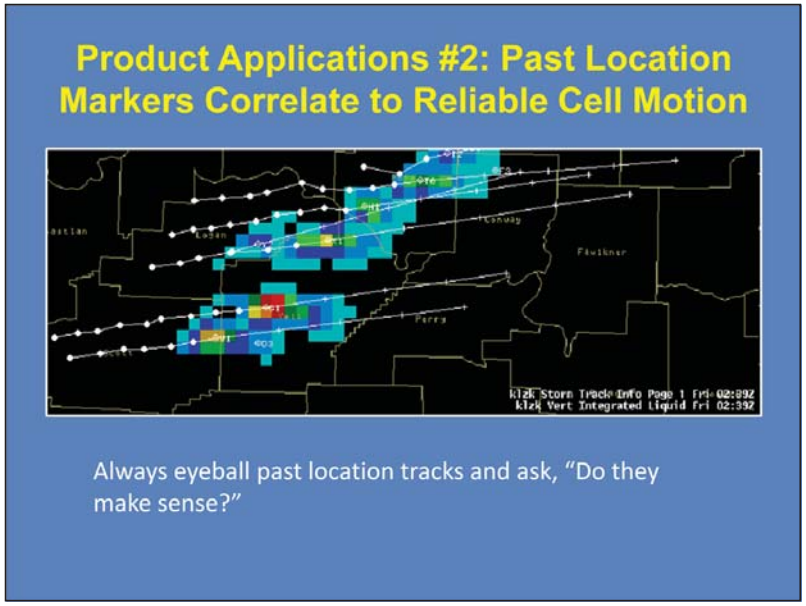
The last limitation presented here has to do with the consistency of cell tracks over time. Depending on which VCP you use, storm tracks can have more continuity or less continuity. In general, cells detected using VCPs 12 or 212 will be tracked with more consistency by the SCIT algorithm than cells detected using VCPs 11, 21, and other similar scanning strategies (Brown et al., 2005). The improved consistency results from VCPs 12/212 more thorough sampling at lower elevation angles.

In a nutshell, we have demonstrated why you can't blindly trust the STI displays. Now, let's talk about how this product can be usefully applied.

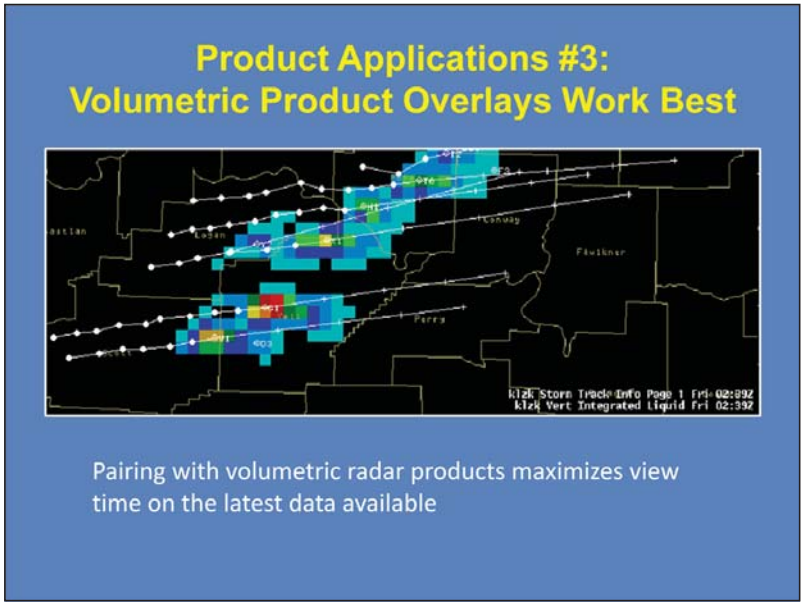


STI is most useful when you have well-defined, separated storms in your CWA. In the example shown, we have the STI overlaid on a gridded VIL product. With this display, we can see past storm locations, the current centroid location matching well with local maxima of VIL values, and future cell positions based on linear propagation. In situations like this one, the algorithm should have small errors and give forecasters a good sense of where these storms may track in the next hour.

If you are using STI to give you a sense of future storm positions, there's an important caveat to consider. The future positions just tell you where the algorithm thinks the storm centroid will be located. There's no indication where the leading edge of the cell is located. That location may not be the most important when it comes to tracking potential severe weather threats or, more importantly, drawing any warning polygons that are needed.



In most cases, seeing a large number of past centroid locations markers (and preferably 4 forecast position markers) signifies a more reliable storm movement. You should always eye-ball these tracks and see if they make sense. Looking at this example, the algorithm appears to have a good handle on things.



As I mentioned previously, this STI is overlaid on a VIL product. While STI can be matched with any radar graphics product, we think it works best to display with volumetric products for time-matching reasons. If you try to plot STI on 0.5 degree Reflectivity, it will take several minutes for STI to appear on your products since STI is created at the end of the volume scan. That gives you very little time to view both products together before the next reflectivity image comes. Volumetric products are generated at the end of the volume scan, so the STI product will reside on the display longer when paired with them.

Other Product Applications

4. Cell motion can be used in storm-relative velocity products



5. Attributes are input to the Hail Index product and SCAN display

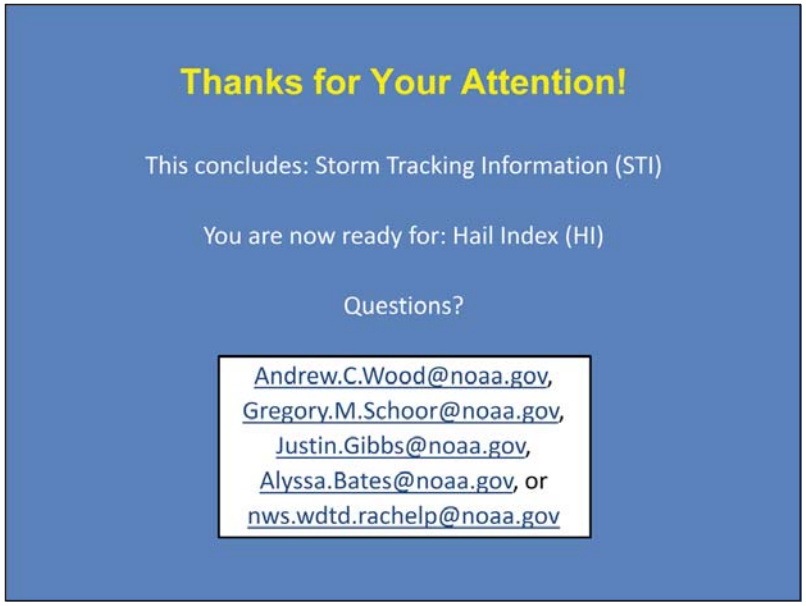


There are two other product applications you should know about with STI. STI cell-motions are used to determine the average cell motion velocity, which is one of three velocities used to calculate the Storm Relative Mean Radial Velocity Map (SRM) product. This average velocity is often a good first guess estimate for use with SRM. Second, the storm attributes shown in STI are the same critical inputs that go into the Hail Index product and SCAN displays. If you use those products, you may want to have STI available in an AWIPS display so that you look at the data going into those products.

Summary

- Errors occur in cell identification and tracking when cells are in close proximity
- Cell attributes are less reliable in VCP 21 (121,221) close to the RDA
- Cell identification and tracking works best when storms are separated and little development or dissipation is occurring
- A large number of past tracks and/or four forecast positions are indications of reliable tracking

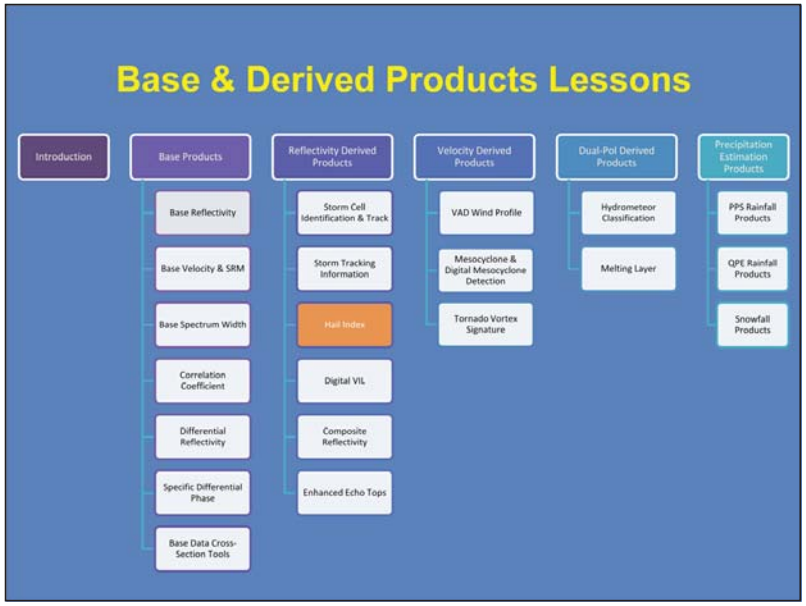
The Storm Tracking Information product provides a graphical display of the storm cell attributes generated by the SCIT algorithm. This products has several limitations. Errors can occur when storms are located close together, fall within the radar's cone of silence, or when they move (or propagate) in a non-linear fashion. Be very suspicious of the attributes in those instances. In particular, cell attributes are unreliable close to the radar when using VCPs 21, 121, and 221. In fact, they can be impacted out to 60 nm of the RDA. Likewise, the STI product works best when storms are widely separated, temporally persistent, and move in a linear fashion. When STI shows numerous past storm centroid locations and four or more forecast positions, the algorithm thinks that the storm motion is quite reliable.



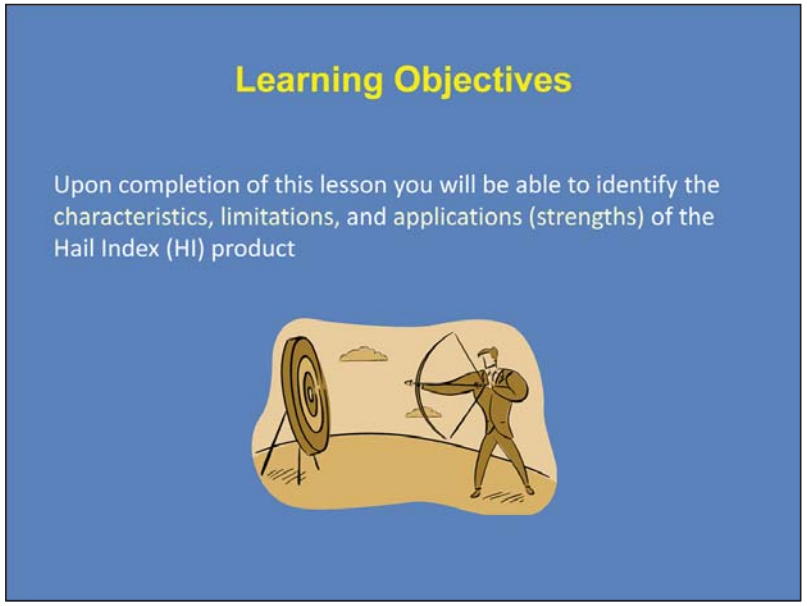
If you have passed the quiz, then you have successfully completed this lesson. You are now ready to move onto the next lesson, Hail Index (HI). If you have any questions, please contact us using any of the e-mail addresses listed on the bottom of the slide.



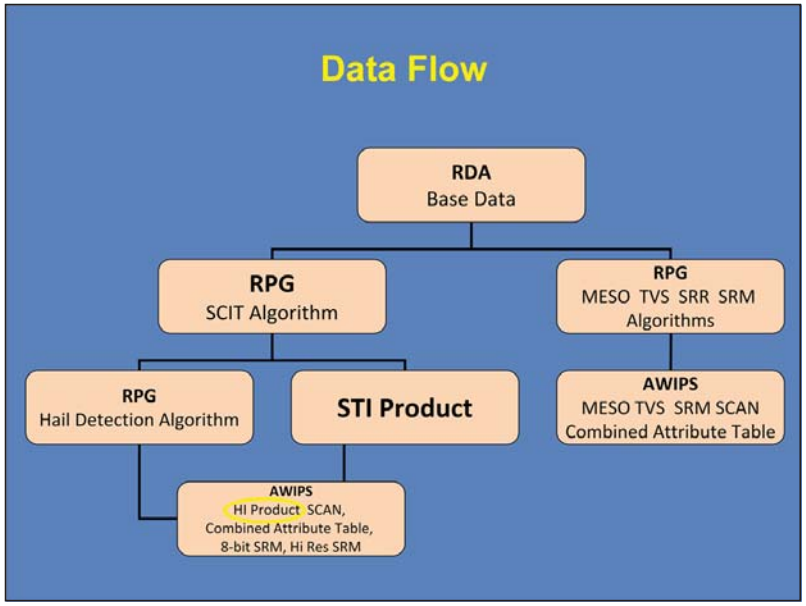
Welcome to the next lesson on Reflectivity-Derived Products. This lesson will cover the Hail Index algorithm.



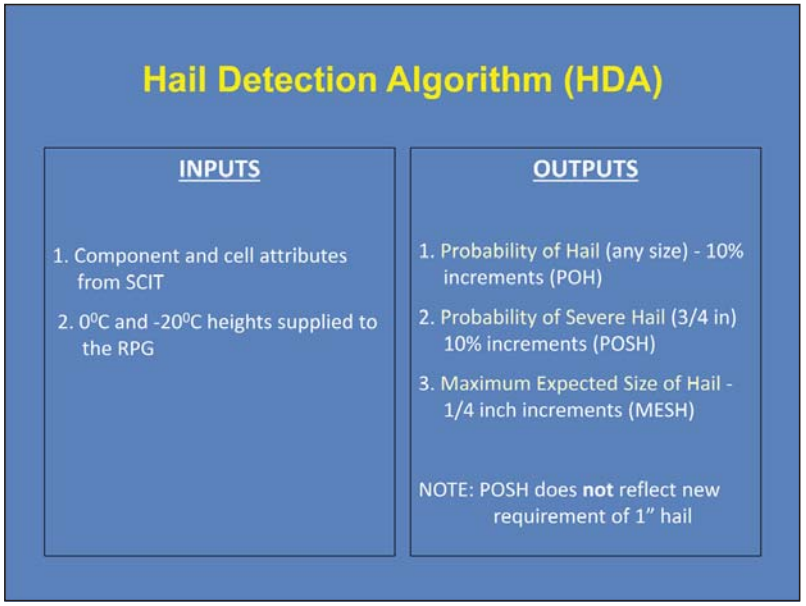
Here is a roadmap for the lessons in this topic. This lesson on the Hail Index algorithm, shaded in orange, is the third in the Reflectivity Derived Products Section of this topic.



Here are the learning objectives for this lesson. When you are ready click advance to start this lesson.



Here is the flow of data from the RDA to the RPG and ultimately to your AWIPS-2 workstation that you have seen before. On the right side is velocity derived products that we'll discuss soon, but for this lesson we can see that output from the SCIT algorithm is ingested by the Hail Detection Algorithm which then produces the hail index product. Let's dive deeper into this Hail Index product.



Now let's go over a summary of the inputs and outputs of the HDA. The Hail Detection Algorithm requires just 2 inputs: Component and cell attributes from SCIT, and the 0 and -20 degree Celsius heights supplied to the RPG. The HDA outputs 3 things that get put into the Hail Index…probability of hail which is any size in 10% increments…probability of severe hail (or POSH) which means hail greater than ¾" in 10% increments…and finally the maximum expected size of hail in ¼" increments which we call MESH. One note here…despite the new 1" criteria for severe hail, the POSH was never updated to reflect this.

Hail Detection Algorithm (HDA)

POH:

- Uses the altitude of the highest component above the 0°C level with a reflectivity > 45 dBZ
- The greater the height, the greater the POH

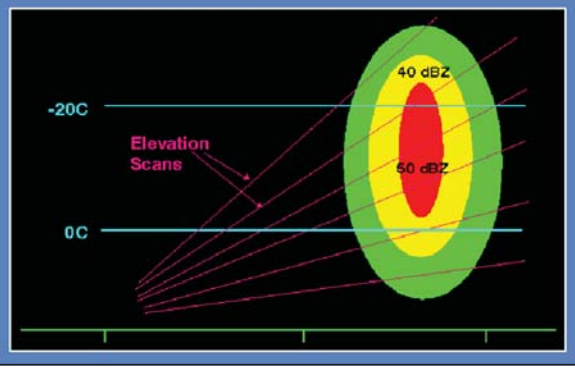
POSH:

- Uses all components above 0°C level that have maximum reflectivities > 40 dBZ
- The greater the height, the greater the weighting factor
- Reflectivity > 50 dBZ and above -20°C altitude get the highest weighting

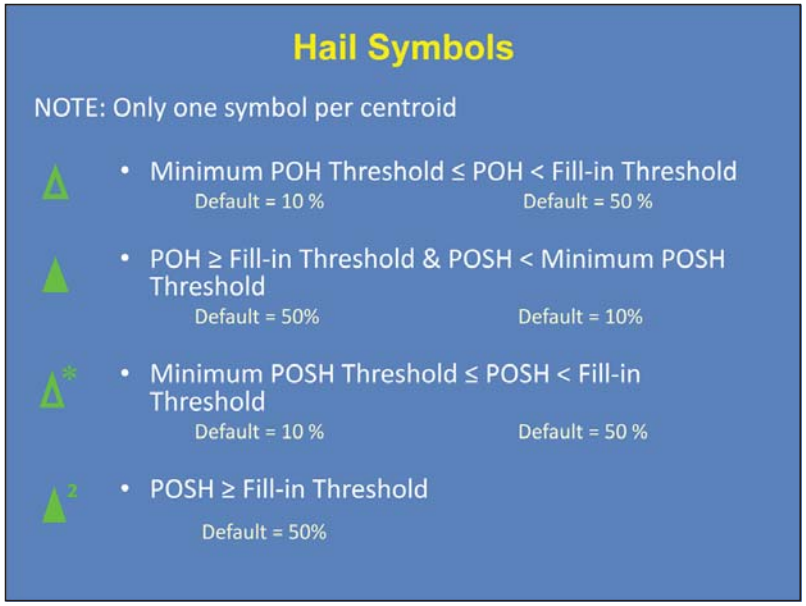
Let's examine the details of the HDA output, the probability of hail first. This uses the altitude of the highest component above the 0 degree Celsius level with a reflectivity greater than 45 dBZ. It is a linear relationship. The greater the height, the higher the probability of hail. POSH uses all components above the 0 degree Celsius level that have maximum reflectivity values greater than 40 dBZ. The greater the height, the greater the weighting factor. Reflectivity values greater than 50 dBZ and above the -20 degree Celsius altitude get the highest weighting. We'll explain this a little further on the next slide.

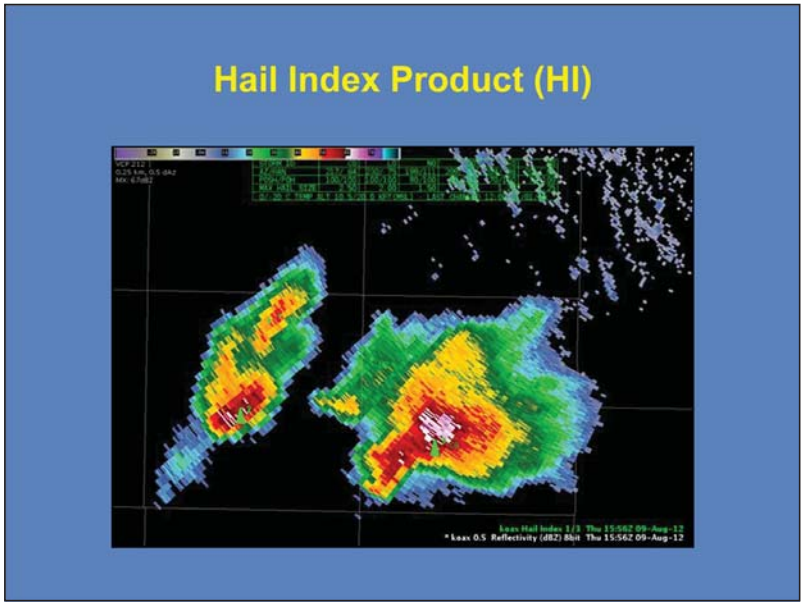
Hail Detection Algorithm

- > 50 dBZ, higher than -20°C → Full weight
- < 40 dBZ, lower than 0°C → Removed entirely



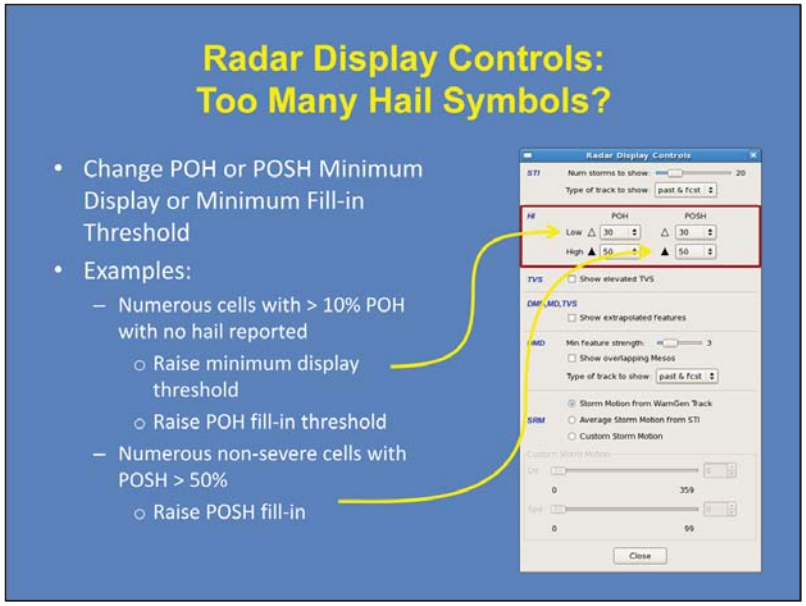
So to recap, the HDA uses 2 weighting function, one is for temperature and the other uses reflectivity values. Reflectivity values greater than 50 dBZ occurring at heights higher than the -20 C level receive full weight. On the flip side, reflectivity values less than 40 dBZ and below the 0 C level are removed entirely. The weighting functions really come into play when reflectivities are between 40 and 50 dBZ or is located between the 0 and -20 C heights.

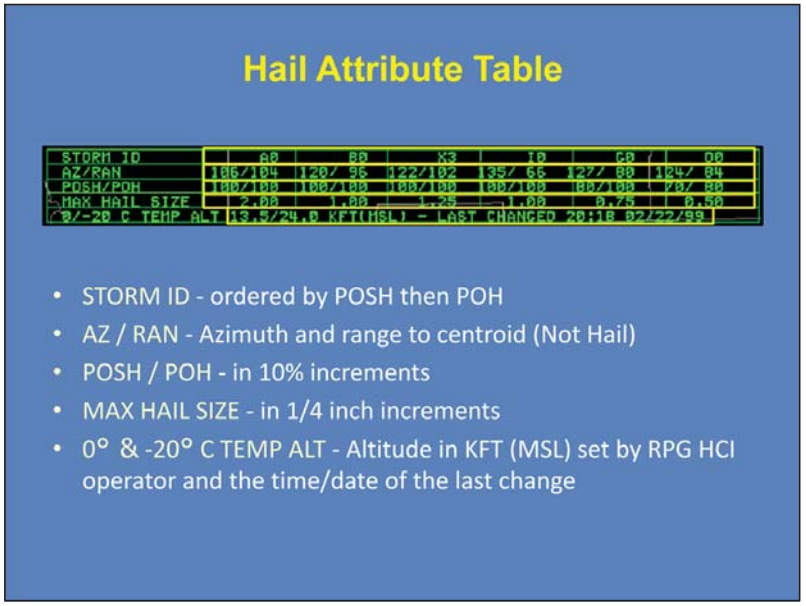




There are multiple symbols that can be assigned to a hail centroid, but only 1 is allowed. Let's look at these different symbols and see what they mean. The first is a small, open triangle shown at the top. This means that the POH value is between the minimum POH threshold (10% by default) and the minimum POH fill-in threshold (50% by default). The next symbol is a small, filled-in triangle which means that the minimum POH fill-in threshold (50% by default) has been met, but the minimum POSH threshold (10% by default) has not been met. A large, open triangle means that the minimum POSH threshold (10% by default) has been reached, but the minimum POSH fill-in threshold (50% by default) has not been met. Finally, a large, filled-in triangle means the minimum POSH fill-in threshold (50% by default) has been met. It is worth noting here that these thresholds are defined at your workstation through the Radar Display Controls GUI which we'll discuss in a few slides. And, the last thing here is the numbers listed in the upper right of some of the symbols. These represent the maximum expected hail sizes in inches. An asterisk just means that the maximum expected hail size is less than ¾ inch. In a warning environment, a forecaster pretty much will only concern themselves with the large, filled-in triangles.

Here is an example of the HI product overlaid on top of reflectivity as you will see it in AWIPS-2. We will take a closer look at the various characteristics of this product over the next few slides.



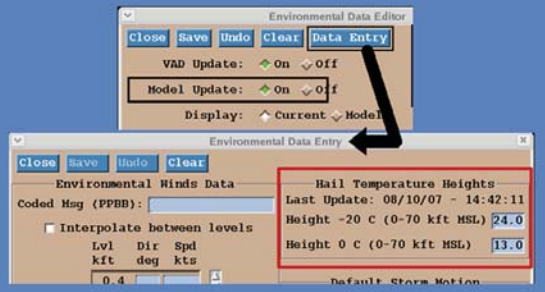


First, the thresholds for the triangles in the HI product are adjustable at your AWIPS-2 workstation. You most likely will never have to adjust these thresholds, but they are adjustable from the Radar Display Controls window found in your Tools menu. A couple of examples where you might want to change the thresholds are…If there are too many symbols and no hail is reported, you may want to raise the thresholds for POH. Another scenario would be if numerous non-severe cells are displaying POSH > 50%, or lots of filled-in, large triangles. In this instance you would want to raise the POSH fill-in threshold.

Here is the Hail Attribute Table that can be found toward the top of the HI product. Let's examine the various information available in this table. In the first row is Storm ID. The ordering of these IDs is by POSH, then POH. These are just the storm IDs that are assigned by the SCIT algorithm. The next row is the azimuth and range to the centroid of the storm identified which may or may not be the hail core. Remember the centroid is the 3D center of mass from the SCIT identified storm. The next row is the POSH and POH values for the identified storm in 10% increments. The fourth row is the Maximum Expected Hail Size (MESH) and is listed in ¼" increments. Finally, the bottom row lists the 0 and -20 C heights in kft MSL as defined in the RPG HCI environmental window. It also tells you the last time these values were last updated.

Hail Temperature Height Selection

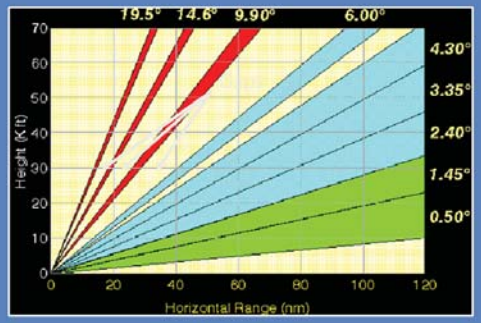
Environmental Data at RPG HCI:



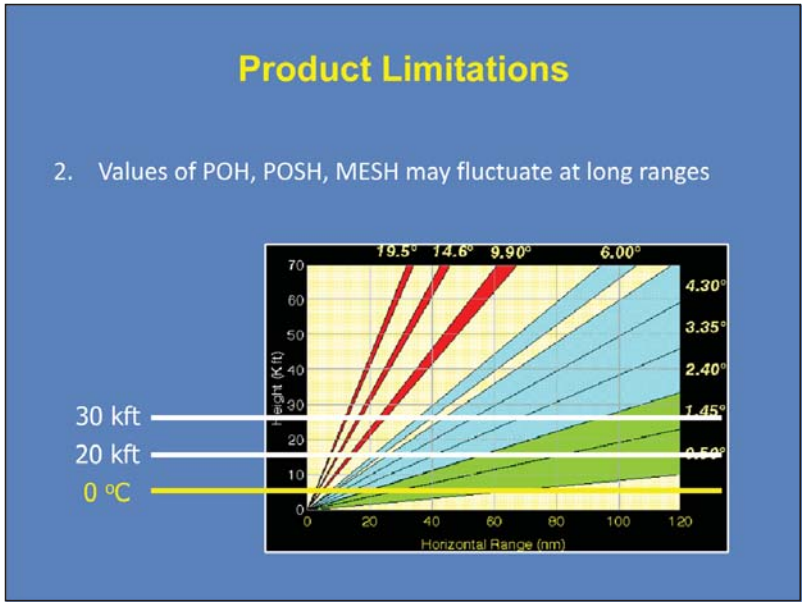
In case you are wondering where these 0 C and -20 C heights are defined, here is the environmental data window from the RPG HCI. The hail temperature heights are highlighted in red. You can access this window by clicking on the Environmental button on the main page of the RPG HCI. Then from the Environmental Data Editor, click “Data Entry”. The window that appears is where you can edit these heights. It is also worth noting here that manually updating these heights can be cumbersome in intense operations. That is why it is recommended that the Model Update be turned “On” in the Environmental Data Editor window at the top. This will use the latest model data to constantly update the 0 and -20 C heights automatically.

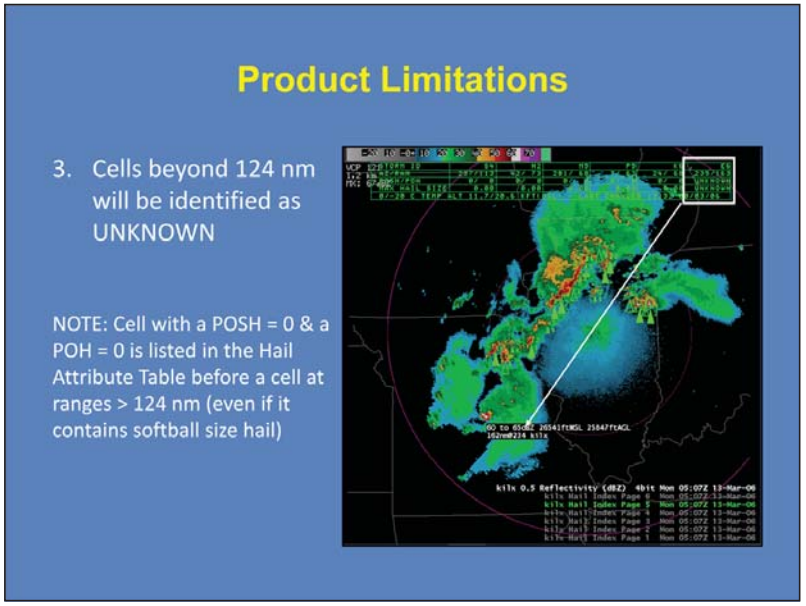
Product Limitations

1. Values of POH, POSH, and MESH will fluctuate at close ranges, specially in VCP 21,121, 221



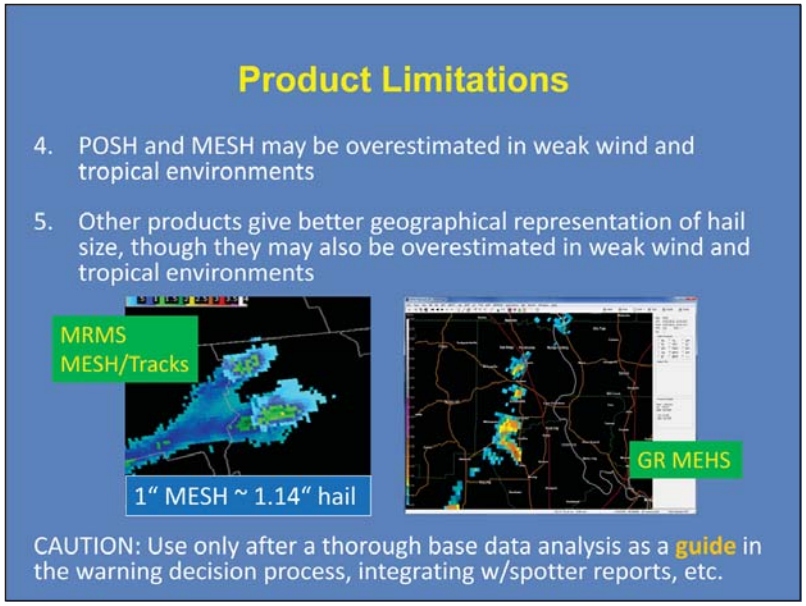
Now for the limitations. The first one is values of POSH/POH and MESH will fluctuate at close ranges to the radar, especially in VCPs 21, 121, and 221. The primary reason is the gaps in elevation which are more prominent when storms are close to the radar.

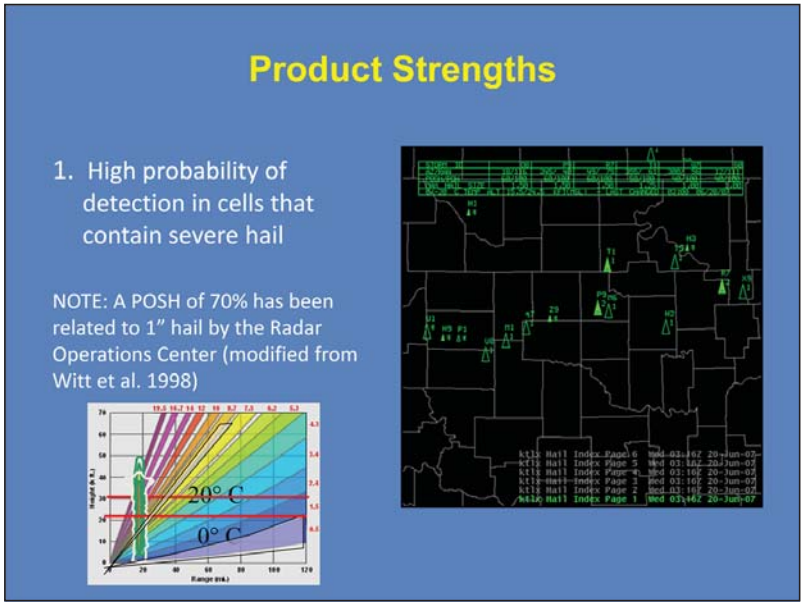




The opposite is also true. At further ranges the values of POSH, POH and MESH may fluctuate. Now, this is not so much tied to VCP choice as it is the earth curvature issue. Recall that at long ranges, the beam is sampling higher and higher into the atmosphere. So at long ranges, only a couple elevations may sample the storm, but as it moves closer more elevations sample the storm and the values of POH, POSH and MESH will fluctuate and this may or may not be representative of what is actually occurring.

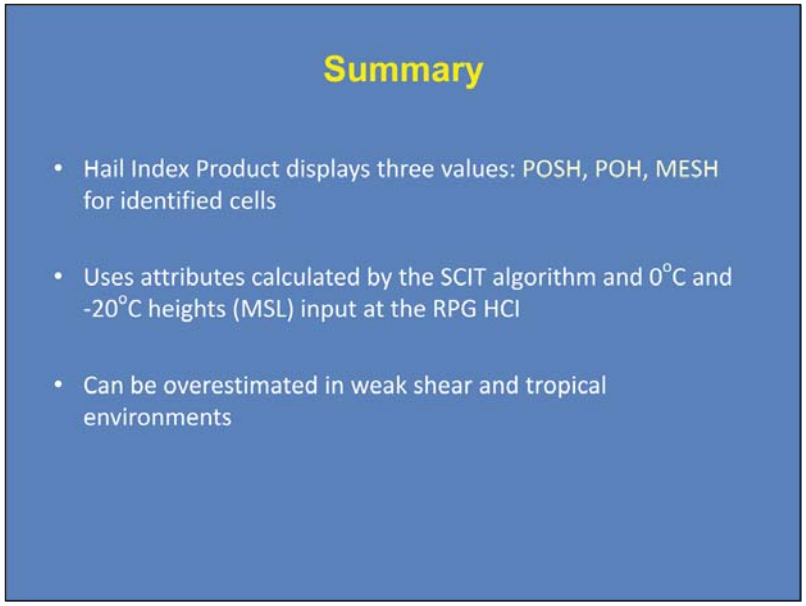
Storms can be identified out to ranges beyond 124 nm, but the HI product only goes out to a range of 124 nm. So storms identified at ranges beyond 124 nm will have HI characteristics listed as “Unknown”. Something interesting here is that cells with POH and POSH of zero will be listed before a storm at ranges further than 124 nm, even if that storm is severe and producing softball-sized hail! Notice in this example that you have to go all the way to page 5 to see this storm identified here which was at a range greater than 124 nm so its HI characteristics were listed as “Unknown”. But, you can clearly see it is severe in nature so you would want to pay attention to it.



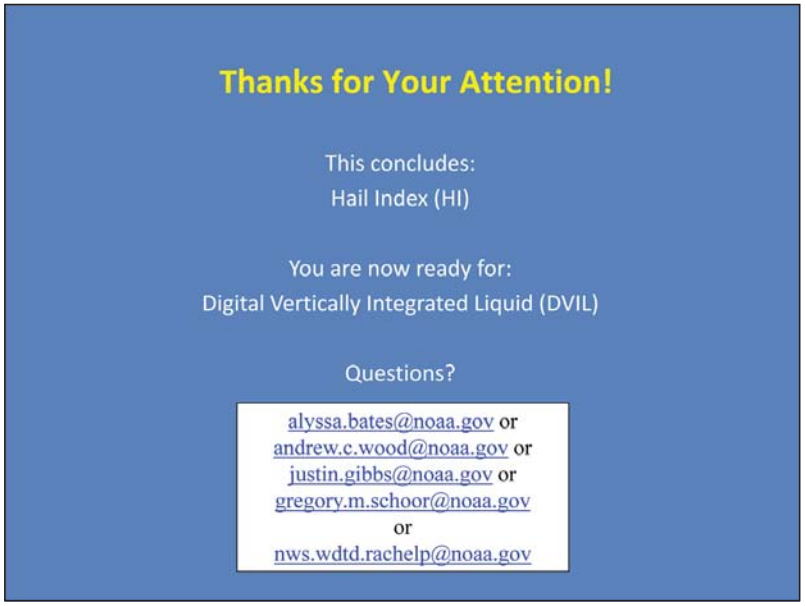


Additionally, several local studies have shown that POSH and MESH may be overestimated in weak wind and tropical environments. Finally, it is worth noting that other hail products can give better spatial representations of hail size and even trends, though they may also exhibit the same overestimation as the HI. Multi-Radar/Multi-Sensor (MRMS) Maximum Estimated Size of Hail (MESH) and its MESH Tracks pinpoint a storm's core and other features such as trends and deviant motion. Studies (Wilson et al. 2009, Cintineo et al. 2012) have shown MESH has utility in determining the presence of severe hail but not to discriminate between severe and giant hail. Use this with caution, but Cintineo's study found 1" MESH roughly correlates to 1.14" hail. If you're a GR user, Gibson Ridge's Maximum Expected Hail Size can also give you a hail size estimate and an inferred core location. GR uses the same hail size algorithm as MRMS MESH, one developed at the National Severe Storms Laboratory, but can still give different estimates. Taking into account all these limitations, it is highly recommended that the HI product only be used as a GUIDE in the warning decision process. Do a thorough base data analysis integrated with spotter reports, etc. and then use the HI product as a consistency check with regards to hail threat.

Let's examine how you'd apply hail index products. Despite the high false alarm rate, the hail index product scores a high probability of detection for severe hail, which is 1 inch in diameter or larger. What this means is that with Hail index it won't often miss severe hail events, but as discussed earlier it false alarms often. If using POSH, the Radar Operations Center found 1" hail corresponds 70% POSH.



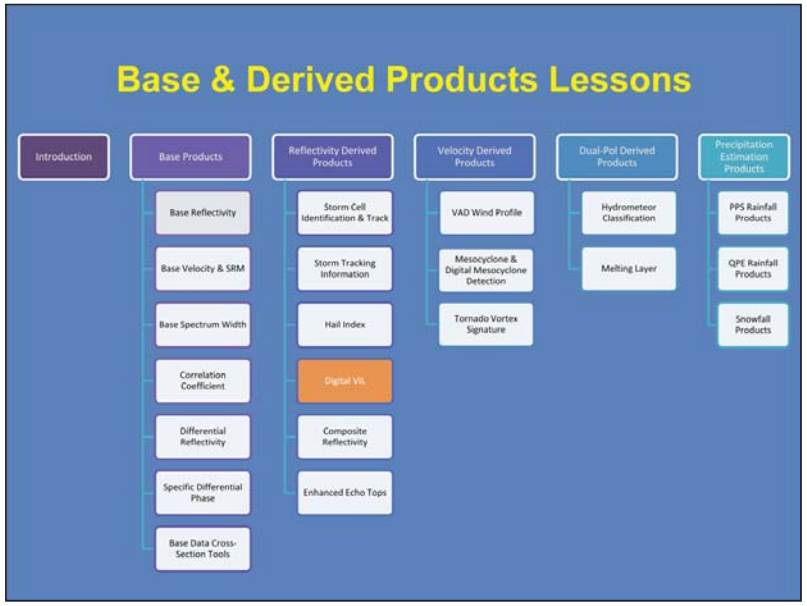
Summarizing this lesson, HI displays three different values; POH, POSH and MESH for all identified cells. It uses attributes calculated by the SCIT algorithm and the 0 and -20 C heights MSL input at the RPG HCI. Keep in mind it can be overestimated in weak wind and tropical environments. The next slide will be a short quiz on this lesson.



This concludes the lesson on the HI. You can now move onto the next lesson on Digital Vertically Integrated Liquid. If you have any questions, feel free to email the contacts listed below.



Welcome to the next lesson on Reflectivity-Derived Products. This lesson will cover the Digital Vertically Integrated Liquid (DVIL).



Here is a roadmap for the lessons in this topic. This lesson on the Digital Vertically Integrated Liquid (DVIL), shaded in orange, is the fourth in the Reflectivity Derived Products Section of this topic.

Learning Objectives

Upon completion of this lesson, you will be able to identify specific characteristics, limitations, and applications (strengths) of the Digital Vertically Integrated Liquid (DVIL) product



Here are the learning objectives for this lesson. Please advance the slide when you are ready to begin.

Digital Vertically Integrated Liquid (DVIL)

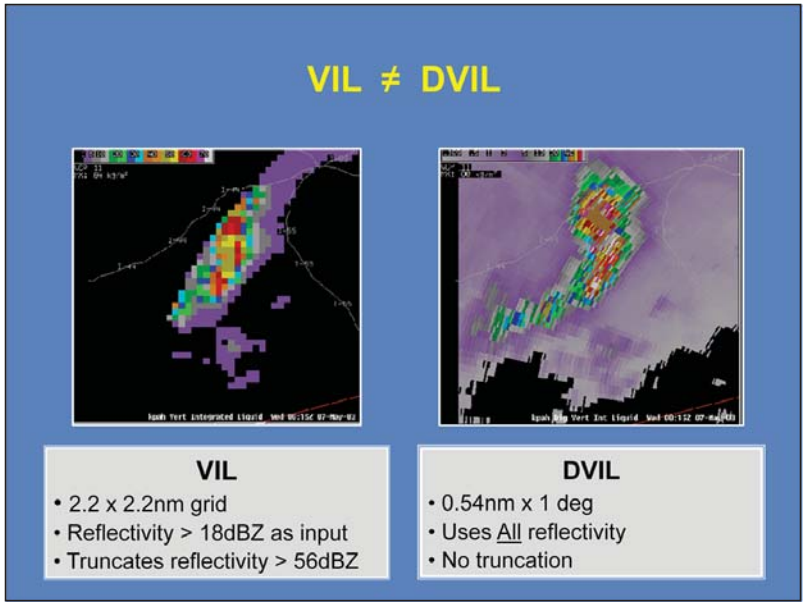
- Uses same equation as VIL for converting reflectivity values to a liquid water content

$$M = 3.44 \times 10^{-3} Z^{(4/7)}$$

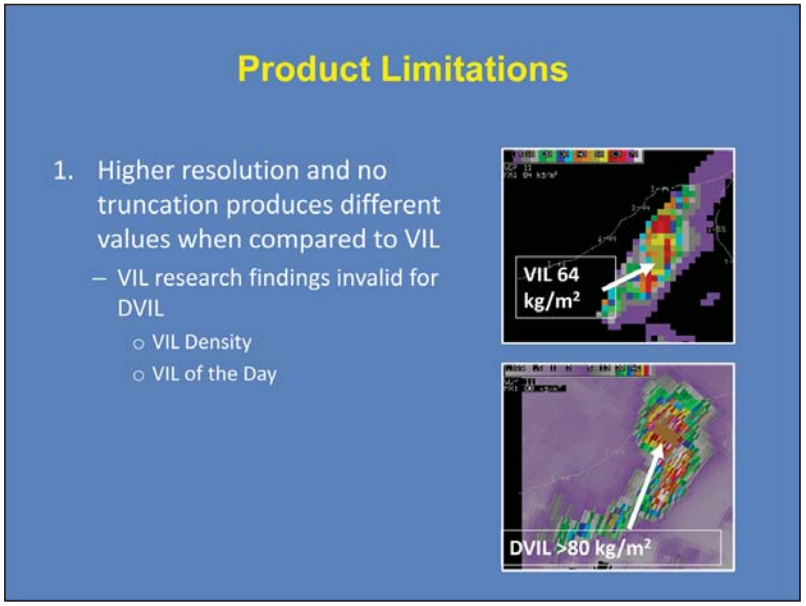
M = liquid water content (g m^{-3})
Z = radar reflectivity (mm^6 m^{-3})

See RAC References NOAA Virtual Laboratory (VLab) page for information about VIL

First, let's note that this lesson only covers the Digital VIL product. If you'd like to learn more about the VIL product, you can refer to the RAC References NOAA Virtual Laboratory (VLab) page for this lesson. What's important here is to note that DVIL uses the same equation as VIL which is shown here, where M = liquid water content (g m-3) and Z = radar reflectivity. The values are derived from each 1 km (0.54 nm) x 1 degree range bin, and then vertically integrated. DVIL values are output in kg/m2, and displayed in 256 data levels from 0 to 80 kg/m2 out to a range of 460 km (248 nm).However, DVIL and VIL differ in many ways. We'll see that in the next slide.



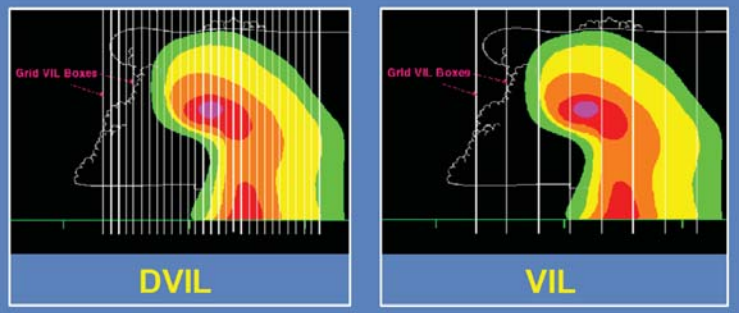
Here are the differences between VIL and DVIL. First VIL is a gridded product that has a resolution of 2.2 nm x 2.2 nm. DVIL is a polar coordinate product that is 0.54 nm by 1 deg. Secondly, VIL does not use reflectivity values below 18 dBZ whereas DVIL uses all reflectivity values. And on the other side, VIL truncates reflectivity values greater than 56 dBZ where as DVIL uses all reflectivity values as mentioned. So, you can see in the example shown that VIL and DVIL are definitely not apples to apples comparisons.



So looking at the limitations of DVIL. The higher resolution and no truncation produces different values when compared to VIL. This difference renders all VIL study findings invalid when using DVIL values. Therefore, you cannot use VIL diagrams when viewing DVIL. You can see this illustrated in the example shown. The VIL and DVIL values are vastly different, so trying to use your "VIL of the Day" while looking at DVIL would not be wise.

Product Limitations

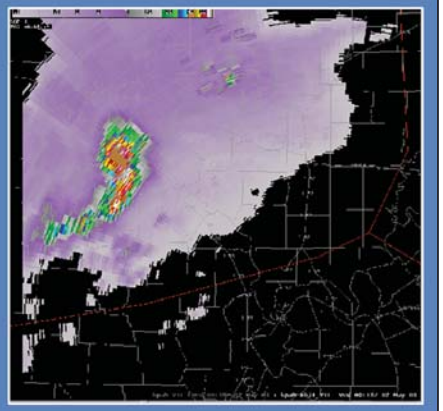
2. Greater impact of fast movement and tilt on DVIL than VIL



Because of the higher resolution of the DVIL product, there is greater impact of fast moving storms and tilt on the values of DVIL than you will see on VIL. So, watch for erroneous DVIL values when you have fast storm motion.

Product Limitations

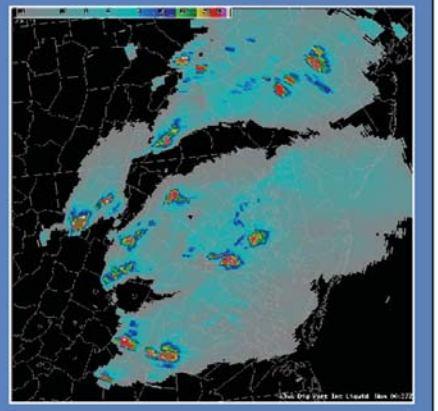
3. 80 kg/m^2 max value is commonly reached due to lack of truncation



Because we don't truncate reflectivities used in the DVIL algorithm, you tend to get larger areas of high VIL values. In this case, 80 kg/m^2, which is the maximum value the product will display, can be reached much more easily resulting in large areas of the maximum value. If you'd like to fix this issue, you can raise the maximum threshold of your color scale.

Product Applications

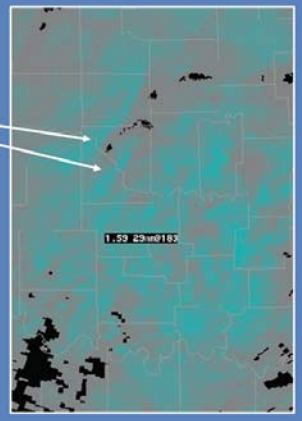
1. Helps locate areas of more significant Z cores



The primary application for the DVIL product is it can help you locate the areas of more significant Z more easily. Notice how in this image, you can quickly identify which areas are experiencing more significant storm. If you were to look at reflectivity, some of the weaker areas may be contaminated by high Z due to ground clutter, but DVIL has not been impacted by it, making it easier to see the strengths of the storms relative to each other in the DVIL product.

Product Applications

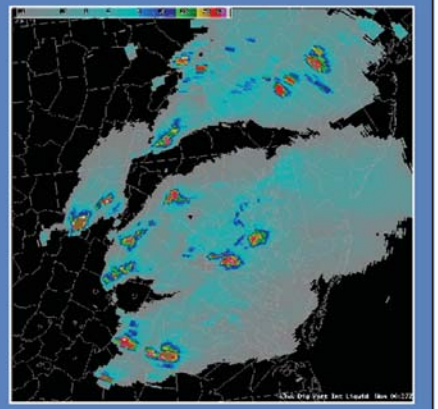
2. Displays low reflectivity features
   - Snow bands
   - Gust fronts
   - Smoke plumes



As is often the case with most of these products, sometimes its limitations in one area becomes a strength elsewhere. This is an example we have with DVIL. Because we are including all reflectivities, especially the lower reflectivities, we can see weak features such as snow bands, gust fronts or smoke plumes in the DVIL product. This was actually a surprise when we first fielded the product, but because of the inclusion of lower reflectivities in the integration, we can see these feature more frequently in the DVIL product.

Product Applications

3. Ground clutter has less impact on DVIL than other volume products (e.g. Composite Reflectivity)
   - Used by FAA

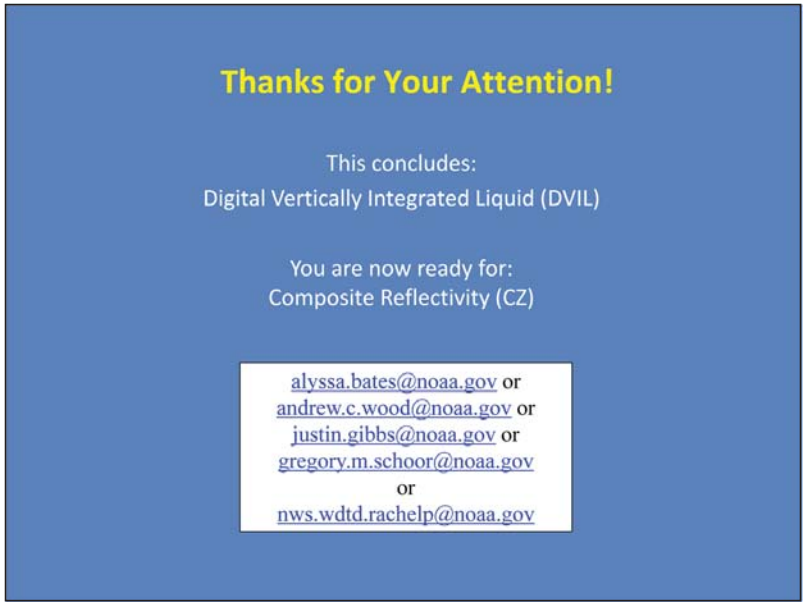


You might be able to better able to sort out ground clutter on DVIL than you would be able to on certain products such as CZ. When you are looking at high Z in the CZ product, and DVIL is low, you know that that region does not have much vertical extent so you can sort out what is AP more easily. This application is used by the FAA to detect ground clutter.

Summary

- Values differ from VIL due to different grid size and lack of truncation of high reflectivity values
- Weak reflectivity signatures displayed
- Impact of ground clutter less than on other volume products

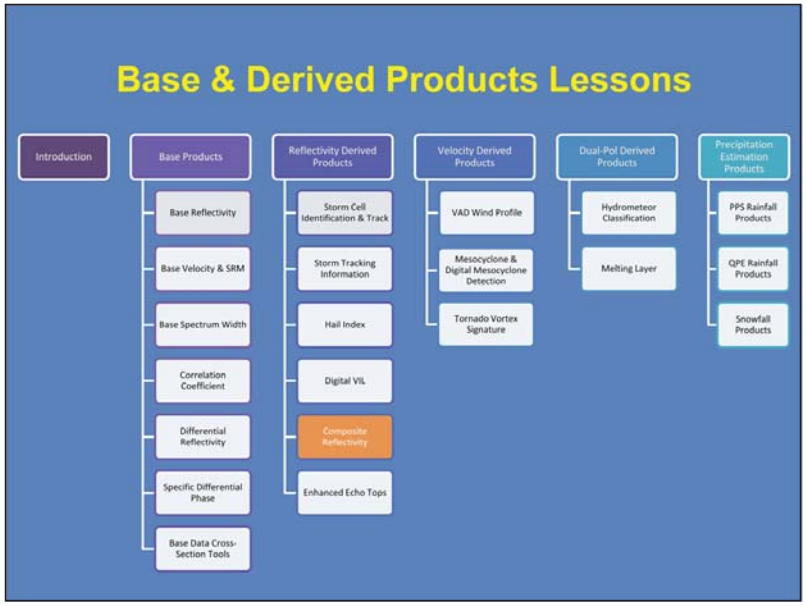
In summary of the digital VIL or DVIL product…the values differ from VIL due to different grid size and the lack of truncation of high reflectivities. You can see very weak reflectivity features displayed, and the impact of ground clutter is less than that of other volume products such as CZ. The next slide will be a short quiz over this lesson.



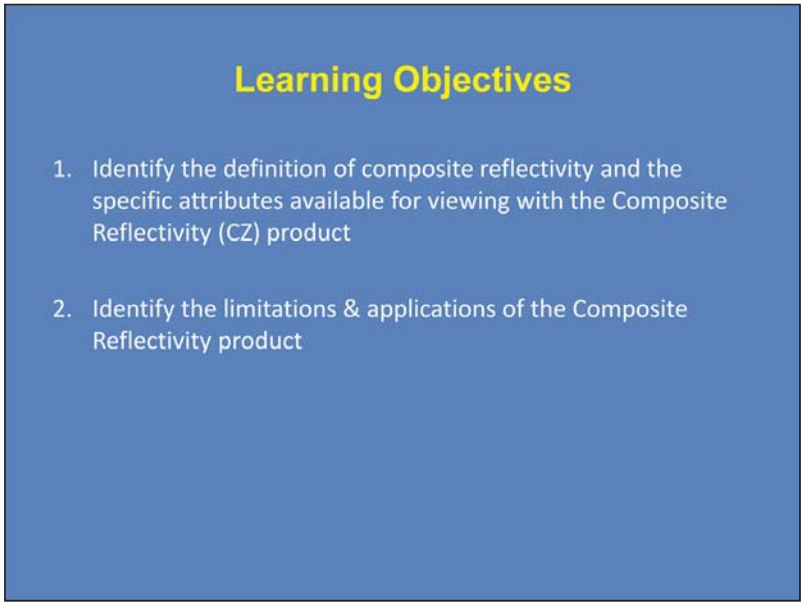
This concludes the lesson on the DVIL product. You are now ready to move onto on the composite reflectivity product. If you have any questions or comments, please email any of the contacts below.



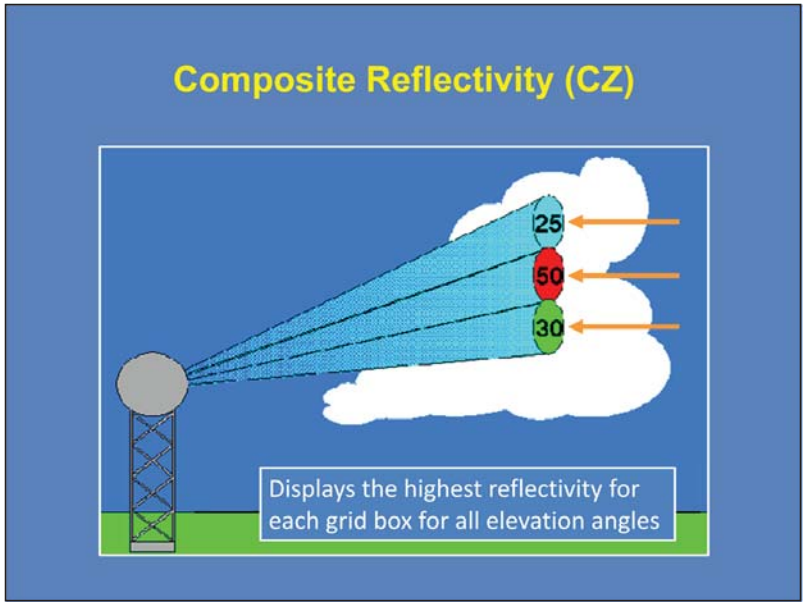
Welcome to this lesson on the Composite Reflectivity (CZ) product. This training is part of the Base and Derived Products topic in the Radar and Applications Course. Let's get started.

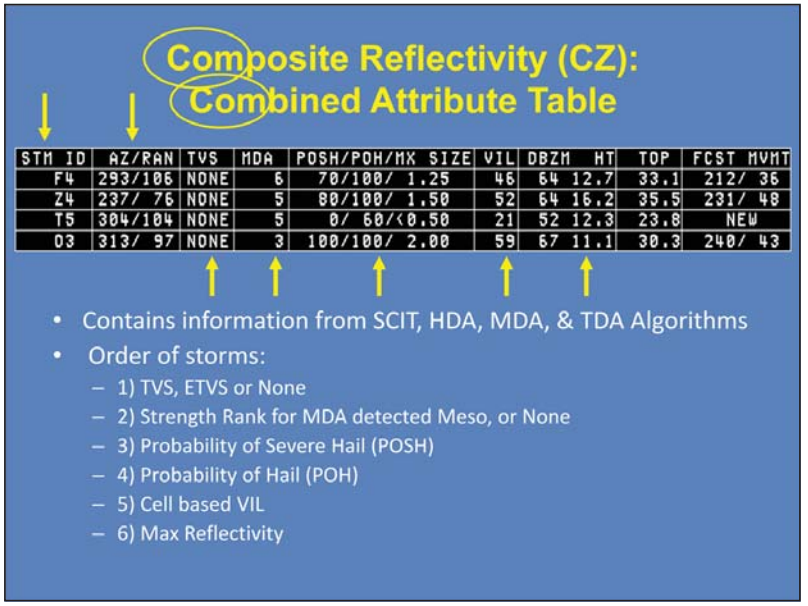


Here is a roadmap for the lessons in this topic. This lesson on the Composite Reflectivity product, shaded in orange, is part of the Reflectivity Derived Products Section of this topic.



Here are the learning objectives for this lesson. Take a moment to read these and advance the slide when you are ready.





The composite reflectivity product is based on a straight-forward concept. The purpose of CZ is to indicate the maximum reflectivity at a point for the entire volume scan. The RPG accomplishes this task by picking the highest reflectivity value for that point from each elevation angle scan from that volume.

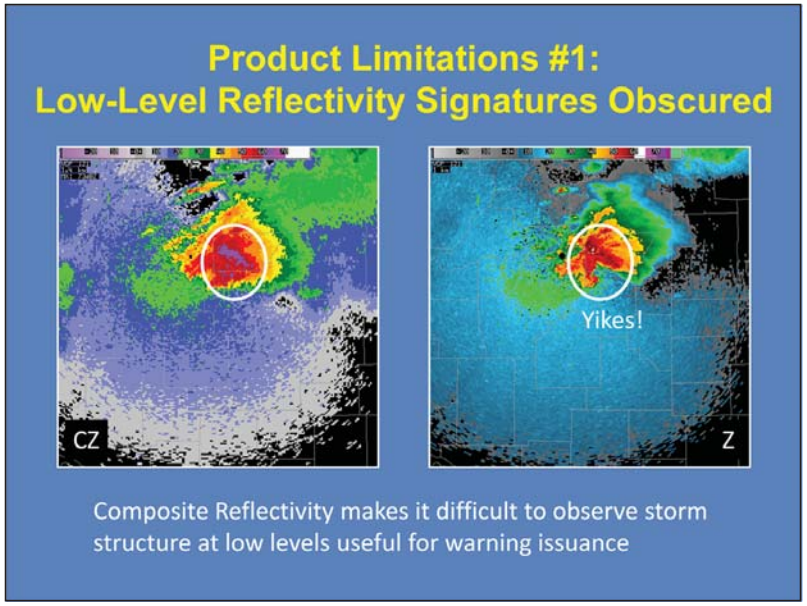
So, if we look at this example here, we are looking at the reflectivity values above a particular grid box. At 0.5 deg, the radar observes 30 dBZ at that location. At 1.5 deg, 50 dBZ is observed. And, at 2.4 deg, the radar detects 25 dBZ. Composite Reflectivity will pick the highest reflectivity value from these (and any other available) data and display that value for that location.

The key point is this: when we look at the CZ product, we don't know from which elevation angle the displayed value was observed. We just know that was the highest value in the volume above that grid box.

The Composite Reflectivity display is different from other reflectivity products because it displays with the Combined Attributes Table. This table displays information from the SCIT, HDA, MDA and TDA algorithms. The Combined Attribute Table gives you a dashboard view of severe weather likelihood. One way to remember which product this table comes with is Combined and Composite both begin with COM.

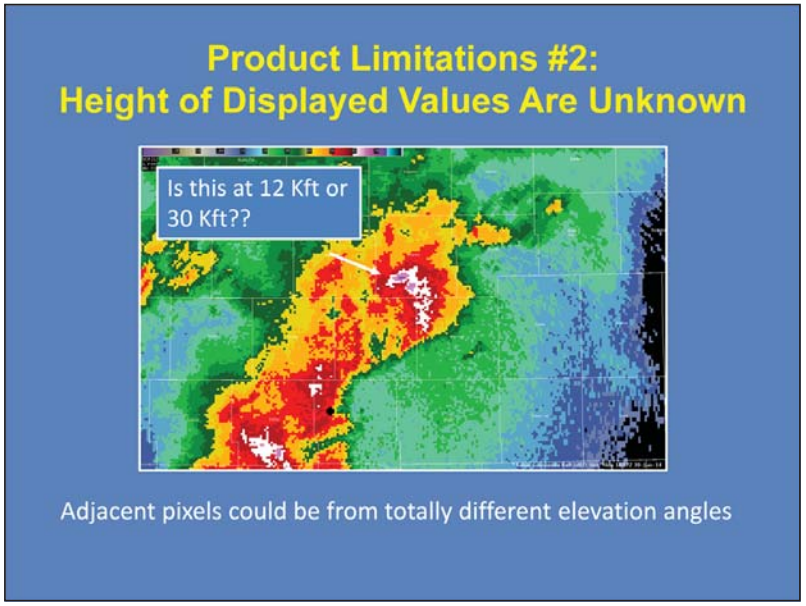
Notice that the table is organized with the Storm ID from SCIT on the left. The next column contains the azimuth and range where the SCIT defined storm centroid is the located (not necessarily the most significant portion of the storm weatherwise).

The remaining data help determine the order in which the storms are listed. The storm order gives precedence to the most "significant" storm first. Storms with a TVS detected are listed at the top. After that, storms with a MDA detection will be listed next in order of the circulation's strength rank. The next criteria used are the storms POSH and POH values. Then, the algorithm uses the cell-based VIL computed by the SCIT. Lastly, storms are sorted by the highest reflectivity value and the height of that value.

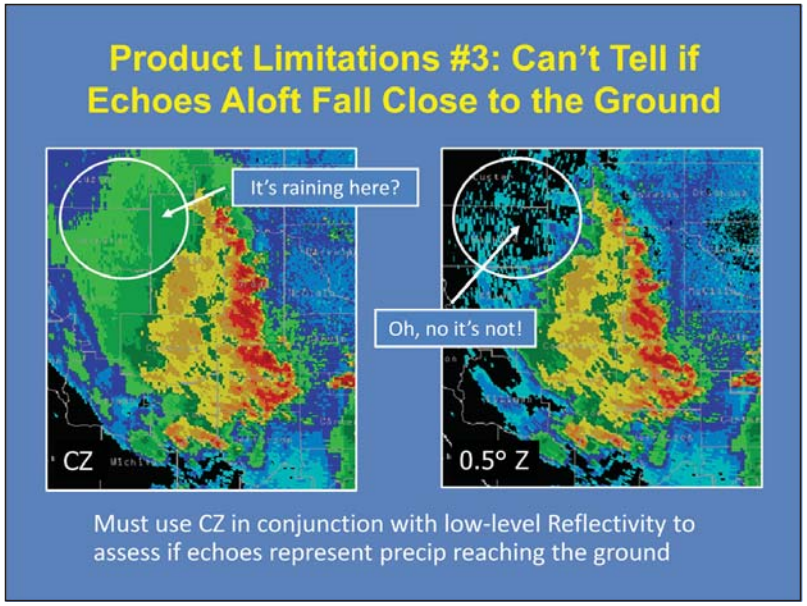


So, let's talk about the limitations of CZ. Arguably the biggest limitation is the tendency of Composite Reflectivity to obscure low- and midlevel signatures. CZ makes it difficult to observe storm structure at these levels useful for warning issuance.

Let's look at an example. The CZ image on the left shows a big blob of a storm with high Z, approaching 70 dBZ. If you were just using this product, then you would miss the well-defined hook echo seen in the image on the right (the 0.5 deg Z product). The lesson here is make sure you use Base Reflectivity along with CZ when analyzing storm structure.

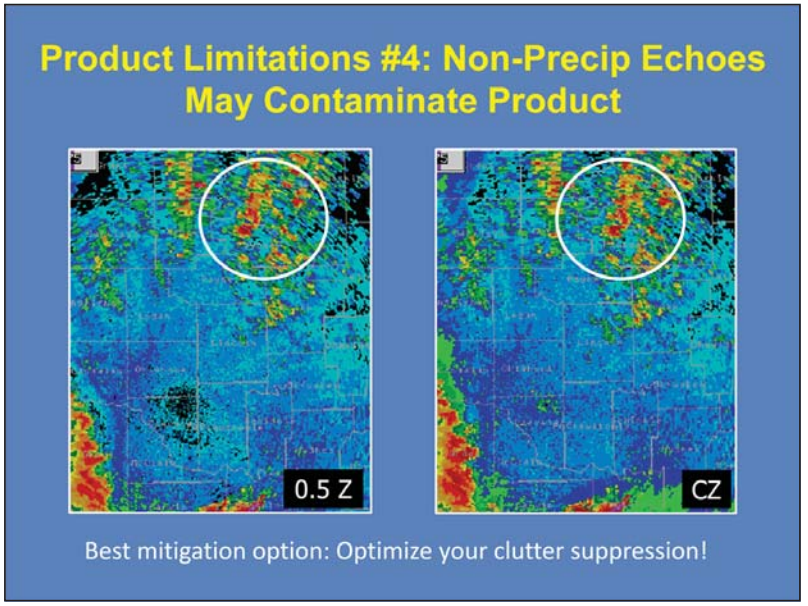


By design, the CZ product gives us the max reflectivity anywhere in the volume scan. The problem is we don't know where it came from. We have no idea if that max Z is from 0.5 deg or 19.5 deg. In this example, we don't know if the pixel being pointed out displays data collected at 12 kft, 30 kft, or anywhere else. All we know is that it is somewhere in that volume above that particular location. This limitation is another reason not to use CZ for identifying storm features because adjacent pixels could be from totally different elevation angles.



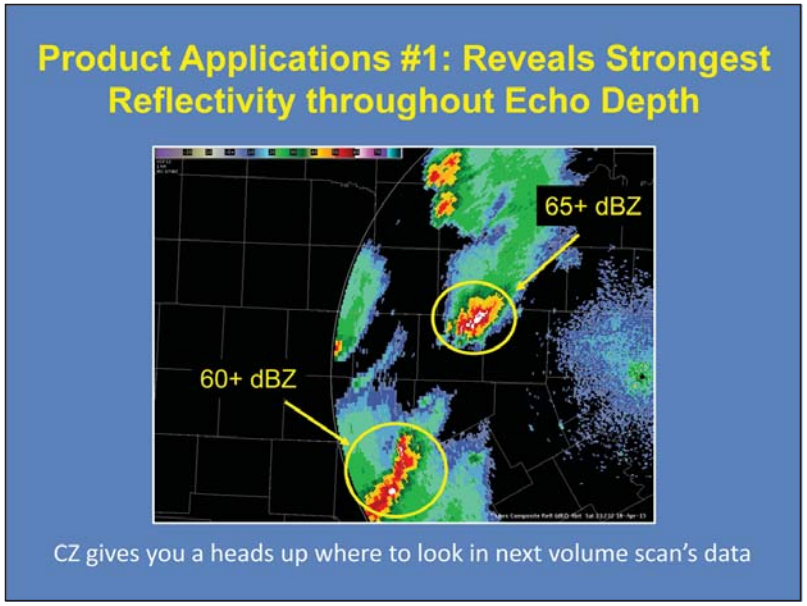
This next limitation is similar to the one we just talked about. Composite Reflectivity doesn't allow you to discriminate whether the reflectivity displayed is precip that actually reaches the ground (i.e., virga).

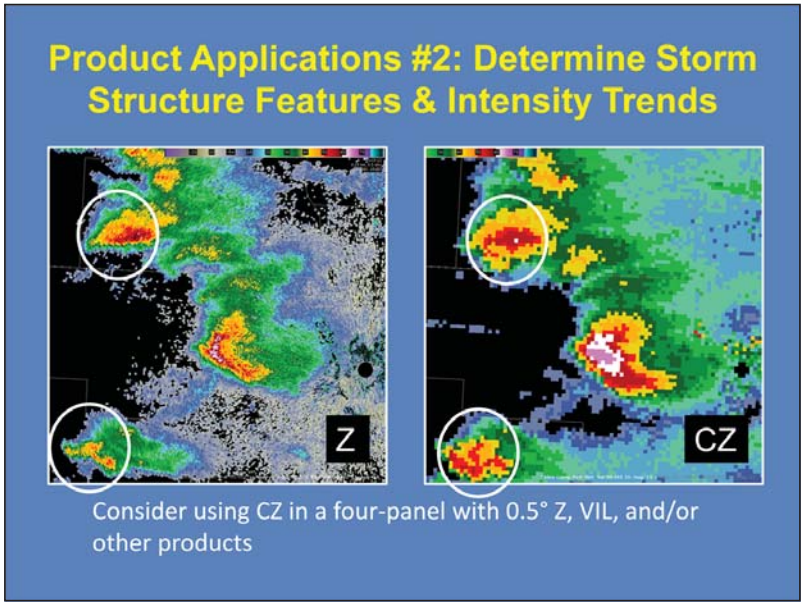
In this CZ product shown on the left, we have a lot of echo. If you just looked at this product, you might think it was raining at the surface in the circled area. However, one look at the 0.5 deg Z and you'll see that those echoes are not precipitation reaching the ground. So, I continue to beat the broken drum that you must use CZ in conjunction with Z to better assess if echoes represent precipitation reaching the ground.



The next limitation involves anomalous propagation and ground clutter. The images on this slide are an example of how these returns on low elevation angle scans can contaminate the CZ product. Remember, the values displayed in CZ are not required to have vertical extent to be included in the product.

The best option is to optimize your clutter suppression. However, you need to be aware that these returns may appear in your CZ product regardless of your office's best efforts to mitigate them.





Enough with the limitations, let's look at how CZ can be useful to you.

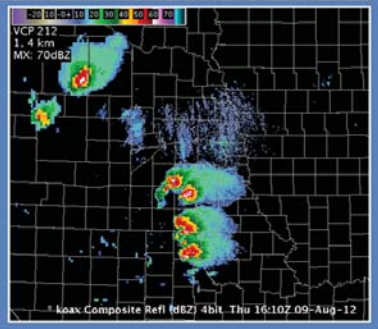
CZ provides a quick look at the highest reflectivity values without having to check each elevation product. I can see in this example that there are three areas with storms to the west of the RDA. The storm in the middle of the display grabs our attention because of its peak reflectivity values exceeding 65 dBZ. The storm to the south appears to be strong, but a little weaker than the first storm. The area of storms to the north appear to be smaller and weaker that the first two. This quick view helps me focus my attention on these areas for a more thorough, base data analysis using the elevation-based products.

I should note that CZ is one of the last products generated during the volume scan, so it will likely give you a heads up on where to look for the next volume scan as that data comes in.

You can get around the limitations we mentioned by using other products with CZ. Often, you will want to use CZ in conjunction with Z. In the example, some of the storms have noticeably strong low-level reflectivity values. However, the storm to the southwest is noticeably stronger aloft than it is at low levels. Likewise, the storm to the northwest has a strong echo aloft that suggests a possible BWER with this storm. When analyzing CZ, display the product in a four-panel with Z, VIL, and/or other products for comparison purposes.

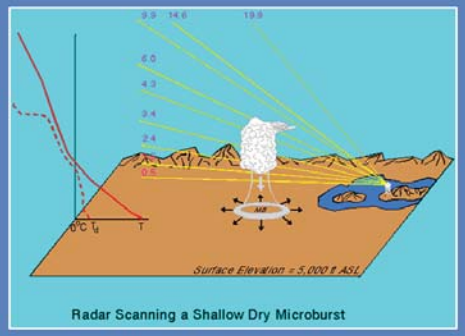
Product Applications #3: Combined Attribute Table

STM ID	AZ/RAN	TVS	MDA	POSH/POH/MX SIZE	VIL	DBZM HT	TOP	FCST MVMT
P4	146/ 31	TVS	NONE	10/100/ 0.75	39	57 12.5	50.1	244/ 58
W7	87/ 92	NONE	19	60/100/ 1.50	61	62 24.4	45.3	287/ 45
N6	75/ 97	NONE	12	60/100/ 1.50	42	59 26.4	50.2	240/ 35
V1	104/ 69	NONE	8	40/100/ 1.25	53	60 20.4	31.8	271/ 31



Another strength of the CZ product is that we can view the Combined Attributes Table. Much of this information also finds its way into SCAN which is available at your workstation and will be discussed at the RAC workshop.

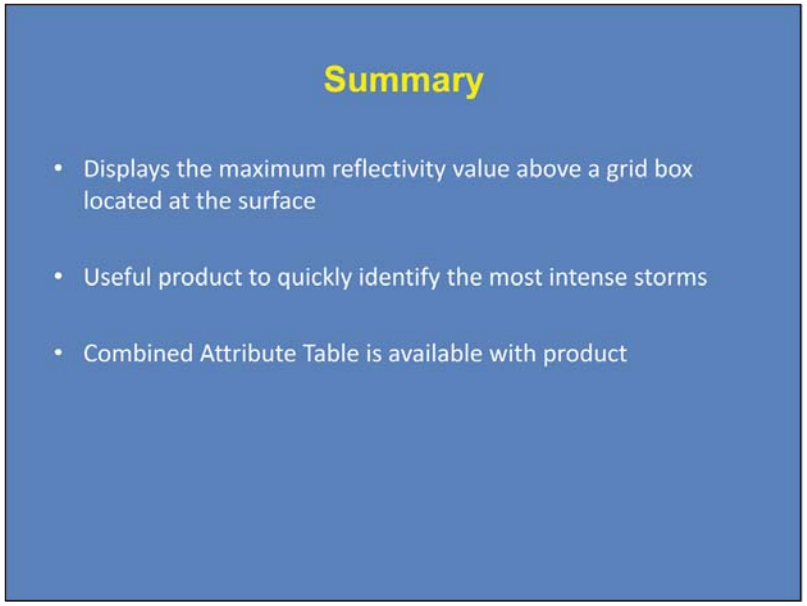
Product Applications #4: Detecting Elevated Storms



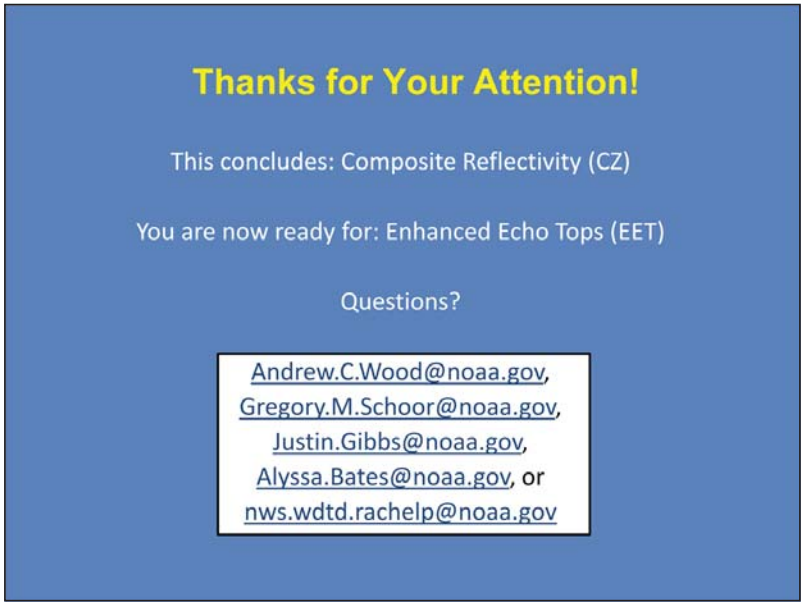
Radar Scanning a Shallow Dry Microburst

Composite Reflectivity can be helpful as pulse storms form aloft

In some regions, like the western U.S., elevated pulse storms may lead to significant downbursts (Wakimoto, 1985). Radar echoes at the lowest elevation tilt may not appear until the downburst is occurring. Composite Reflectivity can be helpful for spotting these storms as they form aloft.



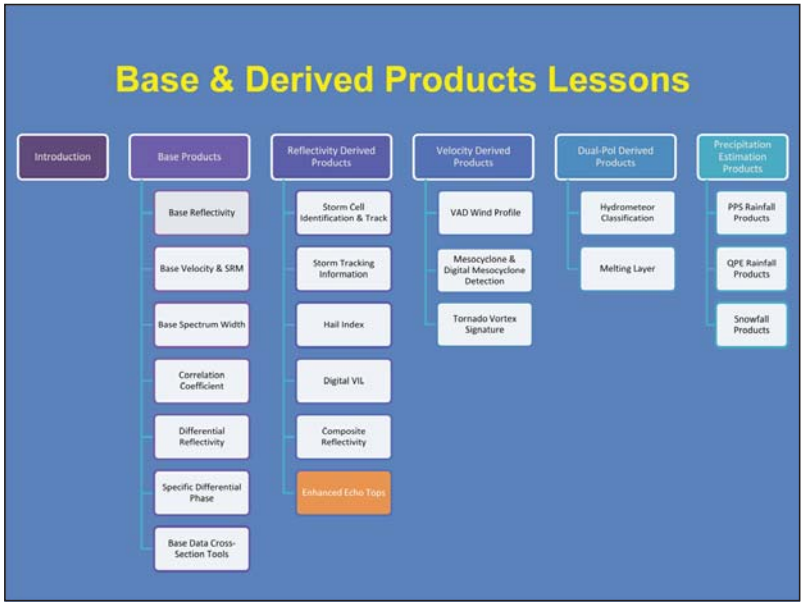
To summarize, the CZ product determines the highest reflectivity value above a grid box at the surface for you radar coverage area using all the elevation scan reflectivity data in the current volume coverage pattern. It is most useful to quickly identify the most intense storms in your local area. CZ also has the Combined Attributes Table.



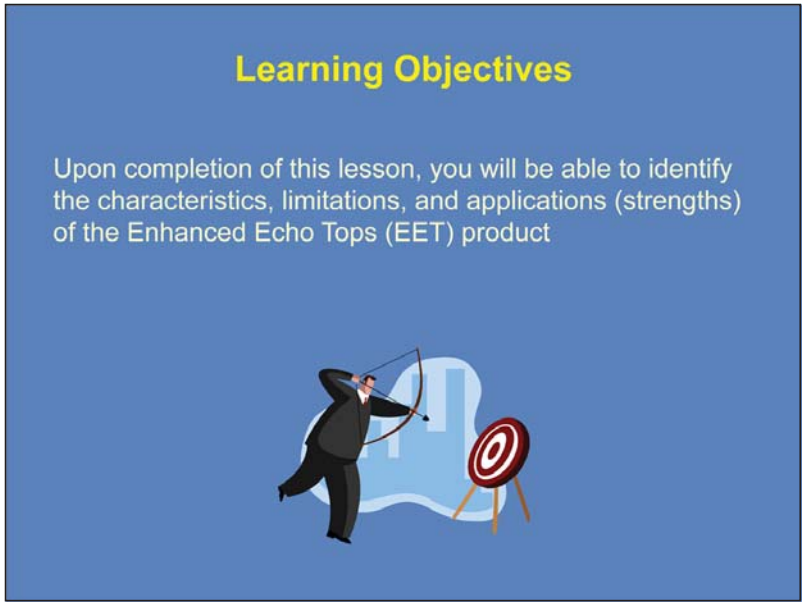
If you have passed the quiz, then you have successfully completed this lesson. You are now ready to move onto the next lesson, Enhanced Echo Top (EET). If you have any questions, please contact us using any of the e-mail addresses listed on the bottom of the slide.



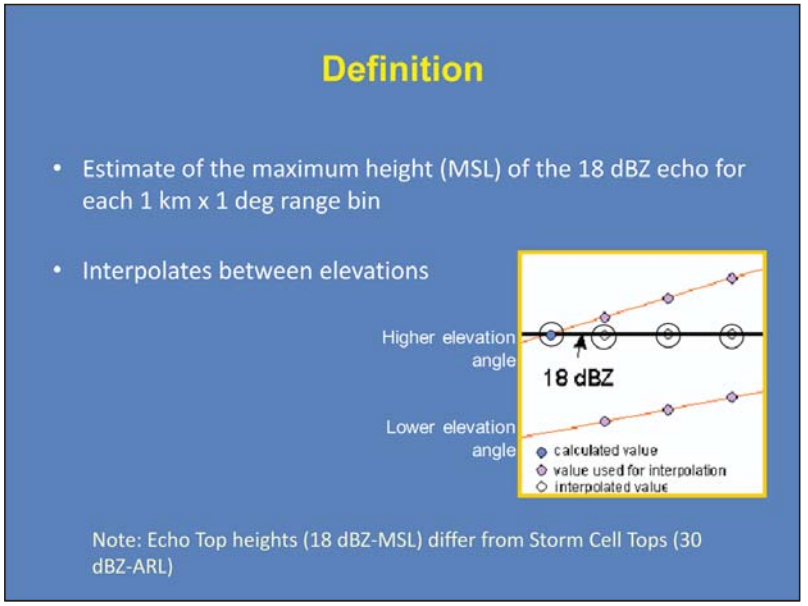
Welcome to the last lesson on Reflectivity-Derived Products. This lesson will cover the Enhanced Echo Tops (EETs).



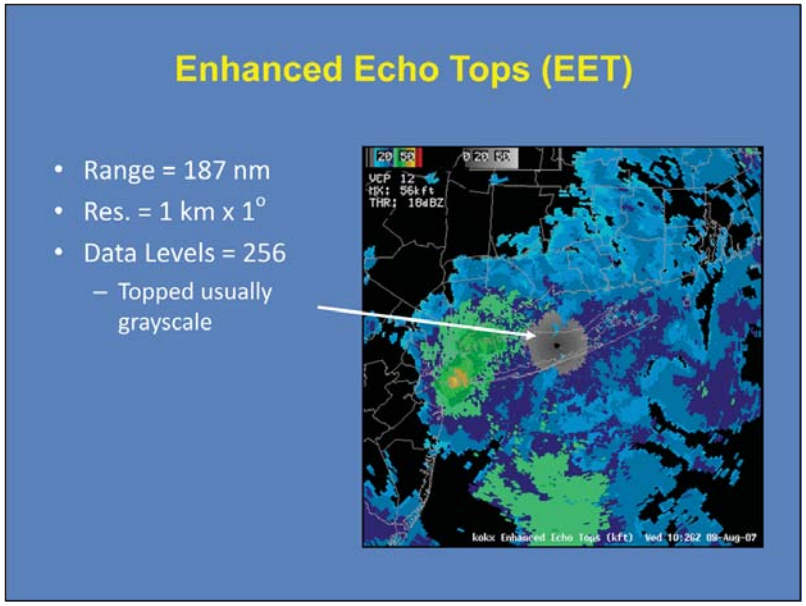
Here is a roadmap for the lessons in this topic. This lesson on the Enhanced Echo Tops (EETs), shaded in orange, is the sixth and last in the Reflectivity Derived Products Section of this topic.



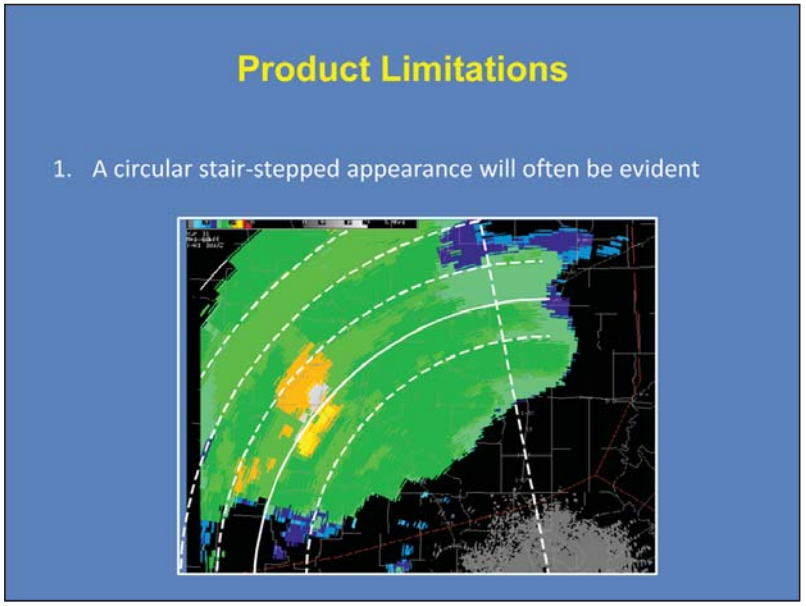
Here are the learning objectives for this lesson. Please advance the slide when you are ready to begin.



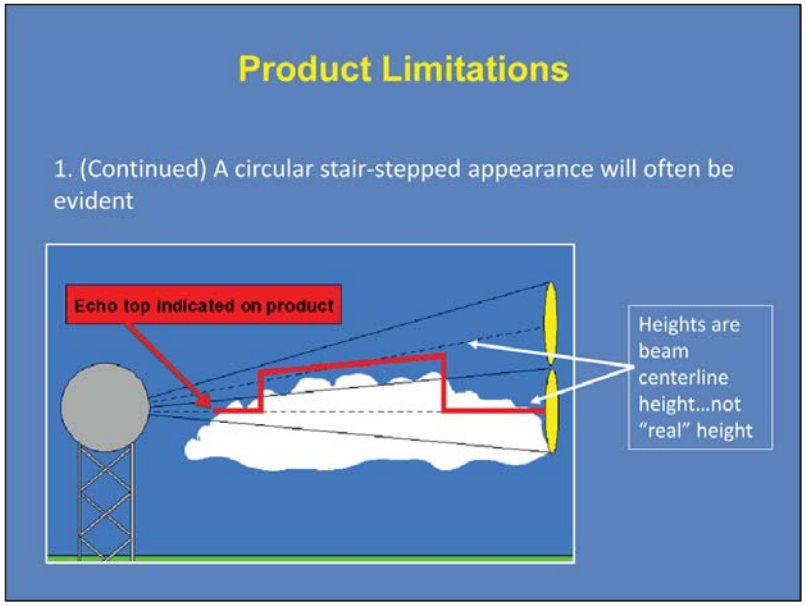
Let's start by defining enhanced echo tops. The echo top is defined as the estimate of the highest height of the 18 dBZ echo for each 1 km x 1 deg range bin. Heights are given in MSL and interpolation is used to determine the height of the 18 dBZ echo in between elevation angles. Look at the graphic on the right for an illustration. The slanted red lines are the lower and upper elevation angles. The light purple dots represent the reflectivity values at each range bin. Let's assume at the first range bin on the upper elevation, the bin is exactly 18 dBZ…no interpolation needed. But for the other bins, let's say that the upper elevation bins are less than 18 dBZ and the lower elevation bins are greater than 18 dBZ. It will use linear interpolation between these elevations to determine the height of the 18 dBZ interpolation. Recall from the SCIT algorithm that it tries to estimate storm tops. These are not the same thing. SCIT uses height of 30 dBZ ARL and EET uses height of the 18 dBZ MSL.



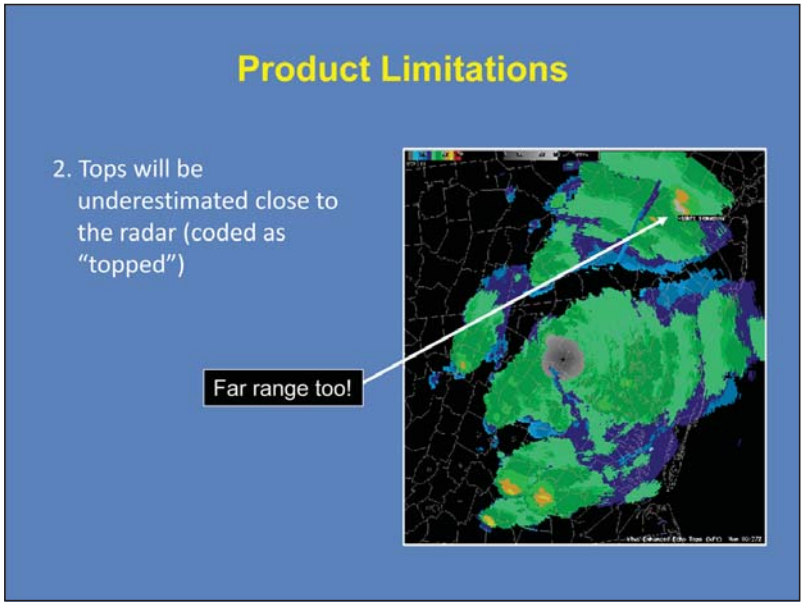
Here are the characteristics of the EET product. It goes out to a range of 187 nm. The resolution is 1 km x 1 deg and it has 256 data levels, or is 8-bit. It also has a gray scale for something it calls “topped”. Bins are tagged as “topped” for areas where the highest elevation angle has a value of 18 dBZ or higher. The algorithm cannot extrapolate the height in these situations, so it tags that bin as “topped”, or gives a basic estimate of what it thinks it could be. The last thing here is in the upper left is the product annotation. It tells you the VCP mode, the maximum EET height anywhere in the product and the threshold used to compute the echo top, in this case it is 18 dBZ which is also the default.



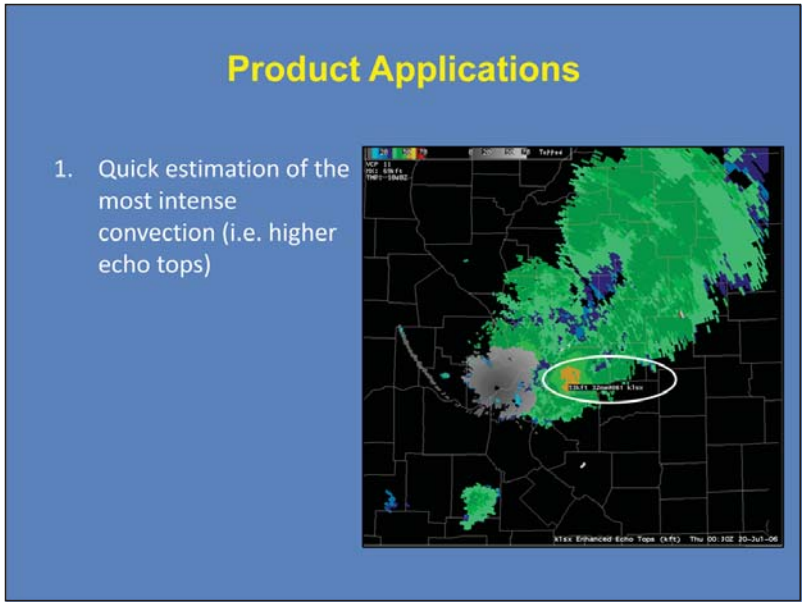
So, for limitations, a circular stair-stepped appearance will often be evident. Recall from other reflectivity derived products with stair-stepped appearances, this is due to the gaps in elevations and the fact that beam center point is being used in the calculation of the height. We’ll look at that on the next slide.



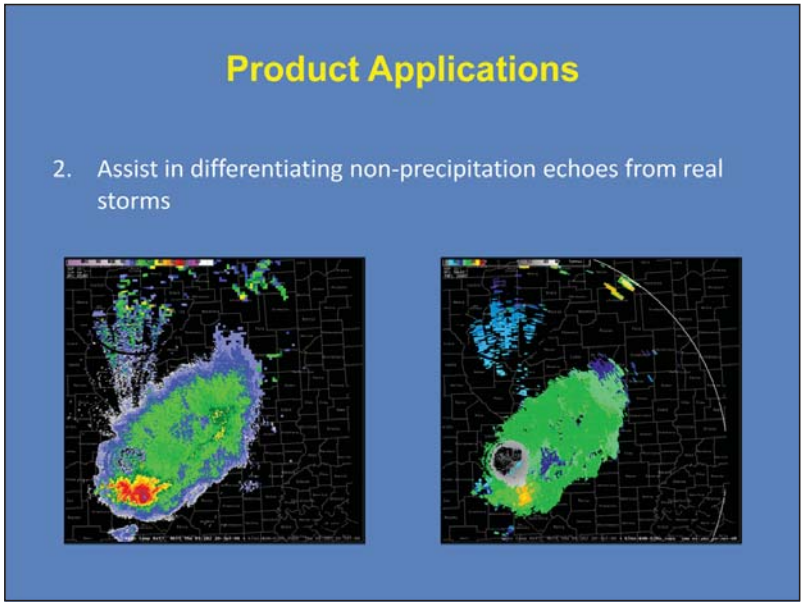
This graphic illustrates how a stair-stepped appearance can be evident. The red line is the echo top displayed on the product where the white cloud shows where the echo top should be located. Because the heights along each radar beam use the beam center point, not the actual height, the algorithm thinks that the heights are actually higher in the middle of this area than they actually are because the higher beam is picking up this reflectivity and it is higher than 18 dBZ. As soon as the higher elevation no longer samples the could top, the echo height goes back down the height of the lower beam.



A second limitation is tops will be underestimated close to the radar. Recall from a previous slide that these are coded as topped because the highest elevation angle is showing reflectivity higher than 18 dBZ. However, topped can happen at far ranges too, but for a different reason. You may recall that reflectivity data are cut off at 70 thousand feet. Thus with each successive elevation angle, the range for which data are shown gets shorter and shorter. So, at very far ranges from the radar, as in the example shown, may have data at the 0.5 and 1.5 deg elevations. However, the 2.5 deg elevation angle at that range is above 70 thousand feet, so it is displayed as no data. Since the 1.5 deg has data and the 2.5 deg has no data, the algorithm tags it as topped just like it would near the radar. This is pretty rare, and most likely will only show up in the most intense convective storms.



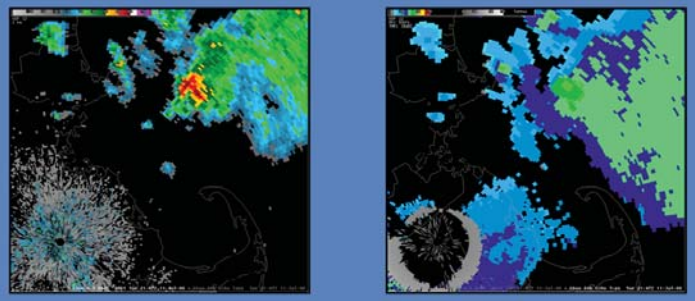
One of the applications of EET is to get a quick estimation of the most intense convection. Typically the higher the echo tops relative to surrounding storms, the more intense the updraft, and thus more potential for more severe weather. Using cursor readout with EET is particularly useful.



A second application is to assist in differentiating non-precipitation echoes from real storms. Recall that this application applies to many of the reflectivity derived products. In the example below, we have a lot of AP in the CZ product in the area to the north of the radar. Notice in the EET product that the tops are blue which are well below 20 kft which is a good thing since AP is usually very low. So, this product can help you differentiate echoes that are real from those that are artifacts.

Product Applications

3. May detect mid-level echoes before low-level echoes are detected



Finally, EET can be very useful in detecting mid-level echoes before low-level echoes are detected. In the example below, the 0.5 deg Z is shown on the left and the EET product is shown on the right. The two little blobs of reflectivity in the white circle eventually went on to be mini supercells. Notice that there are some tops to almost 20 kft with the weaker echoes at low levels, but it isn't the greatest example since there are already echoes in the lowest levels, but it does illustrate the point.

Summary

- Estimate of the maximum height (MSL) of the 18 dBZ echo
- Primary use of the product is to identify storms with greater vertical development
- Can be used in differentiating real echoes from non-precipitation echoes

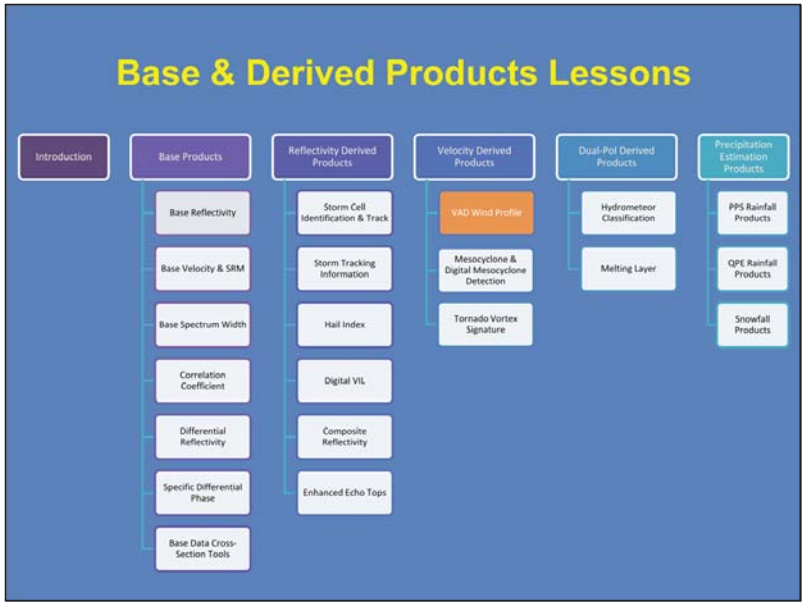
In summary, the EET product is an estimate of the maximum height, in MSL, of the 18 dBZ echo. It's primary application is to identify storms with greater vertical development and thus could be the most severe. Finally, it can be used in differentiating real echoes from non-precipitation echoes.



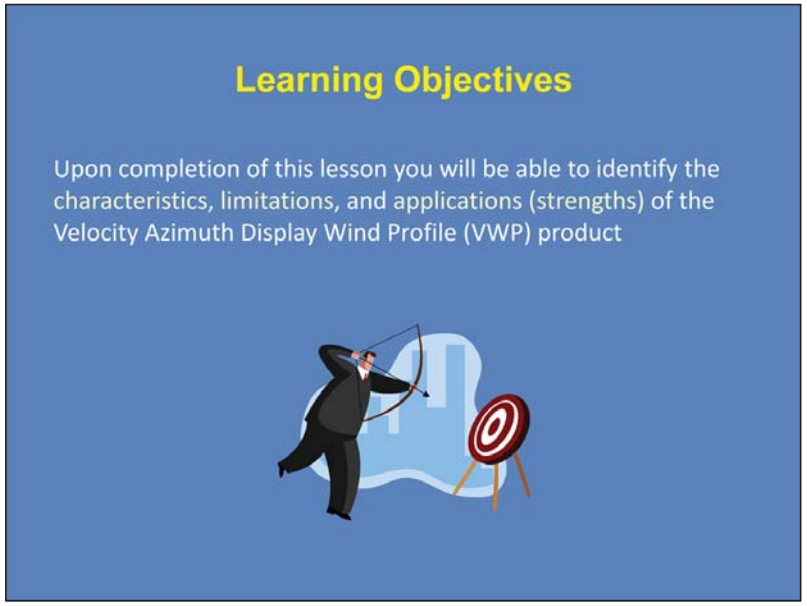
This concludes the lesson on the HI product. You are now ready to move onto lesson 14 on the DVIL product. If you have any questions or comments, please email any of the contacts below.



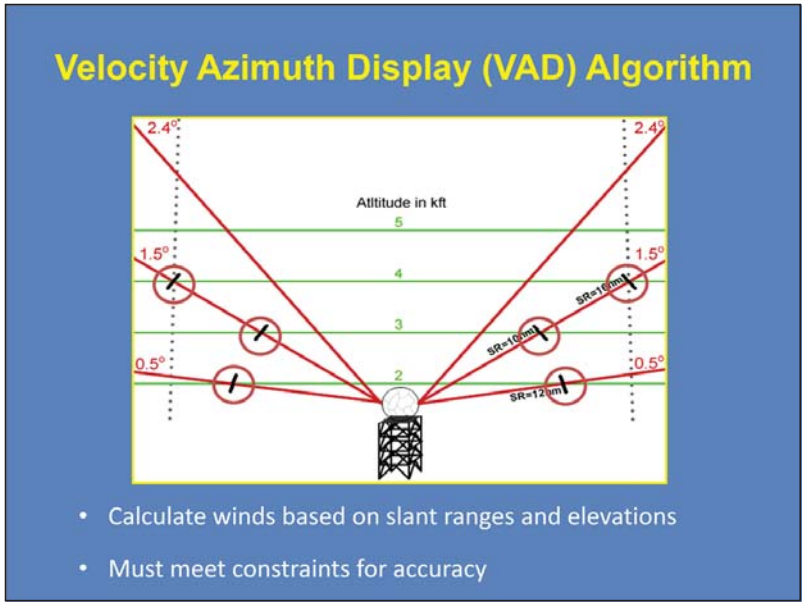
Welcome to the first lesson on velocity-derived products. This lesson will cover the Velocity Azimuth Display Wind Profile, or VWP.

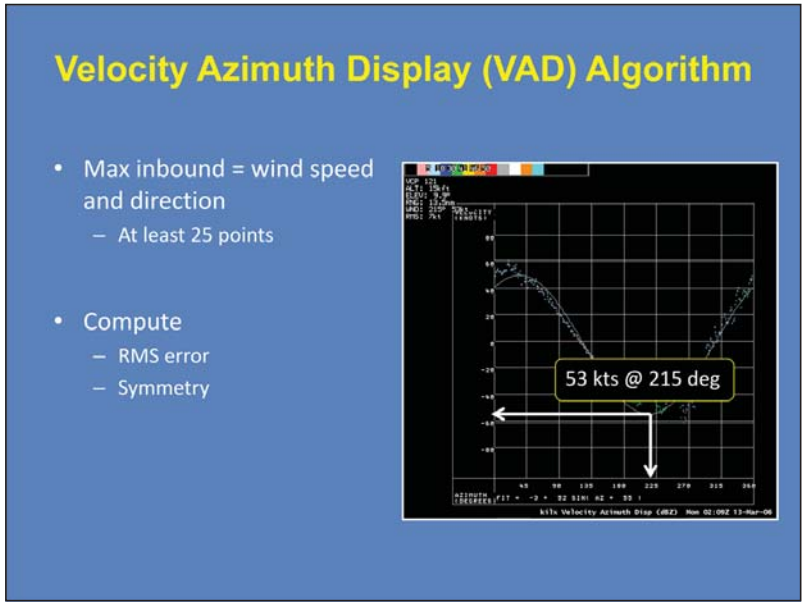


Here is a roadmap for the lessons in this topic. This lesson on the VWP profile, shaded in orange, is the first in the Velocity Derived Products Section of this topic.



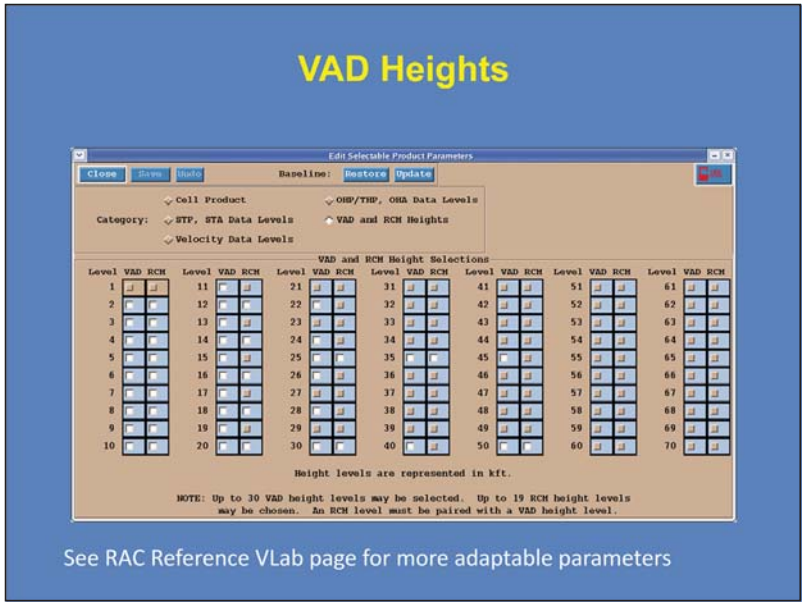
Here are the learning objectives for this lesson. When you have read them, click next to start.



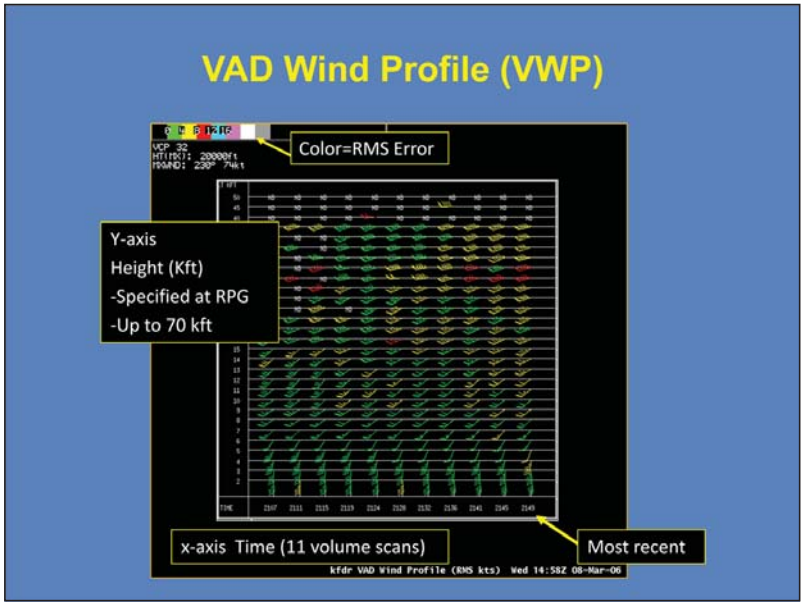


The VWP product is an output of the Velocity Azimuth Display (VAD) algorithm. So we'll quickly look at how this algorithm works, and then jump right into the VWP product. For each height requested, the VAD algorithm finds the elevation angle that closely intersects a slant range, at which it computes a wind speed and direction. It then uses multiple slant ranges and elevation angles at that height to compute more winds that are constrained for accuracy. This procedure is repeated for every requested height level. Note that the slant range increases with elevation angle and that the slant ranges are calculated between 10 and 120 km of a radar. In order for any wind from any elevation to be considered, there is an "error test" requirement that ensures that any error in the estimate is sufficiently low. Since VAD calculates more than one possible wind for any given height, a single wind is chosen using a validity weight or validity index to take into account the error, the wind speed, and the number of points used to generate the wind estimate..

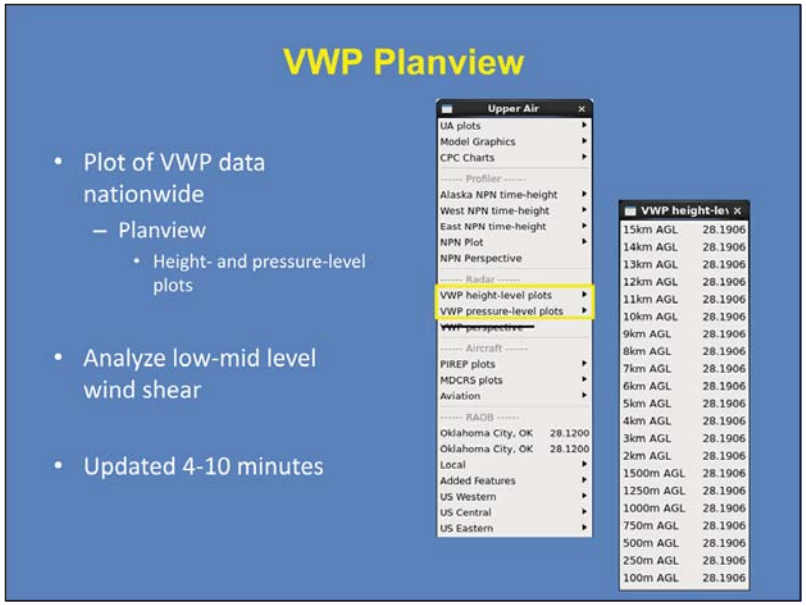
When all the points are plotted, it looks like this graph here. Notice that it looks like a sine wave, so the algorithm fits a sine wave to the plotted data. From the fitted sine wave, the maximum inbound velocity, or maximum negative velocity value, is recorded as the wind speed at that height. The direction of the wind is the azimuth that velocity value occurs. And, to fit a sine wave, there needs to be at least 25 points. So in this example, we see there are plenty of points and the wind at 15 kft from this fitted sine wave would be approximately 53 knots from 215 degrees. The algorithm then computes RMS error and symmetry basically telling us how close the data fit the sine wave telling us the reliability of that estimated wind speed and direction.



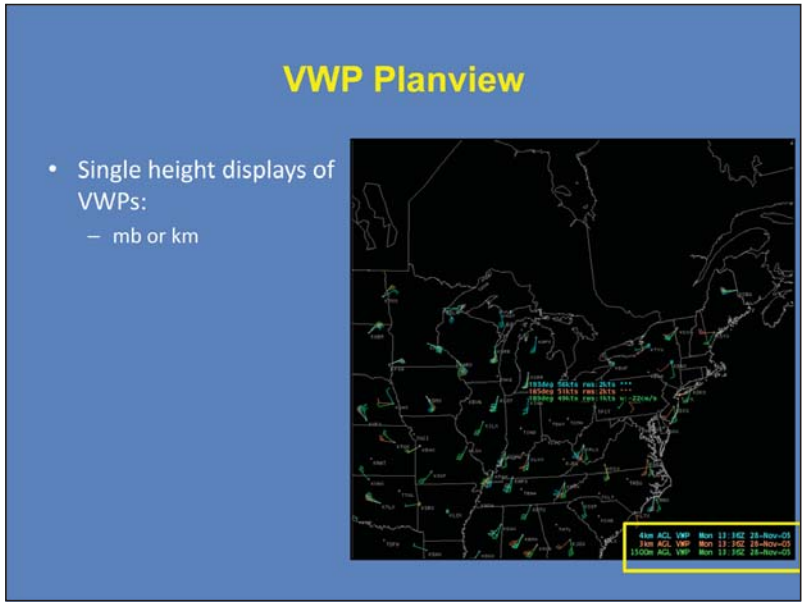
I'll let you explore the RPG at your leisure, but here is the RPG window where you can specify the heights that will be used to compute the VAD winds. You can select winds in 1,000 ft increments from 1 kft to 70 kft.



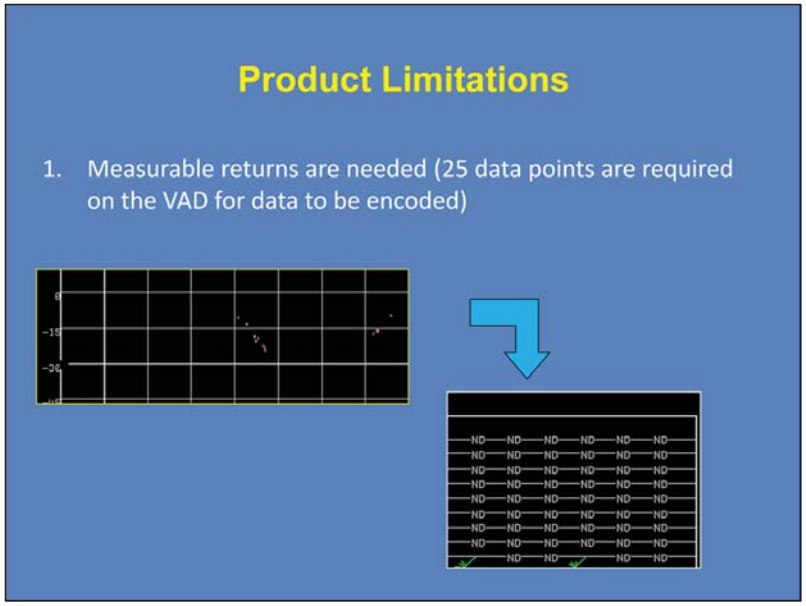
So, for each volume scan you can plot the wind at each height and construct a wind profile, which is the VAD Wind Profile product, or VWP. It's found in your local radar-derived products menu. The most recent profile is on the right and previous profiles are to the left. Time is on the x-axis and the heights requested are plotted on the y-axis in thousands of feet up to 70 kft. As shown in the previous slide, these heights are specified at the RPG. The color bar in the upper left corresponds to the color of the plotted winds and represents the RMS error. So, looking at this example, winds plotted in green have low RMS error, and winds plotted in red have high RMS error. Basically, this tells us the reliability of those plotted winds. You can trust the green plots, but you should be suspicious of the red winds, or at least check and make sure they can be verified.



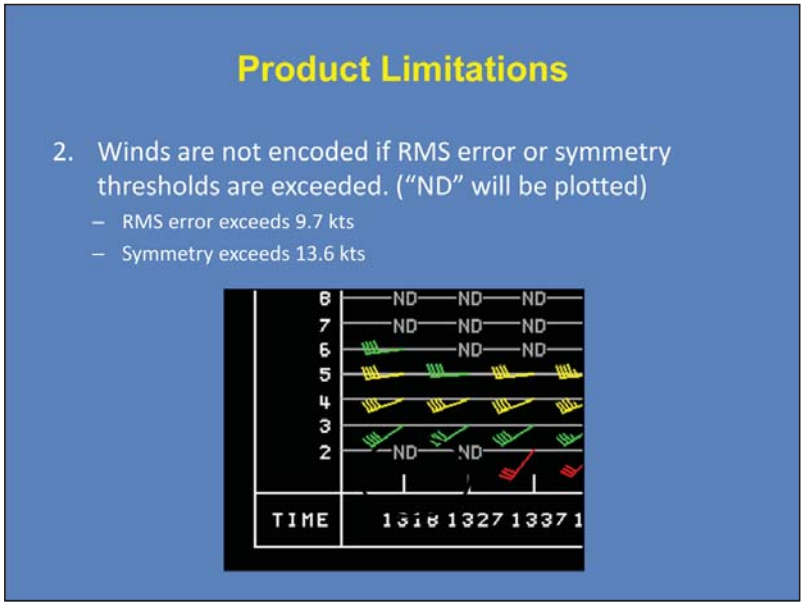
You can also plot the VWP winds on Plan views to get a spatial feel for the wind patterns when upper air plots from RAOBs are not available. You can find these plots in the Upper Air menu in CAVE as seen in the image. These winds are updated every 4-10 minutes, or whenever new volume scans are available.



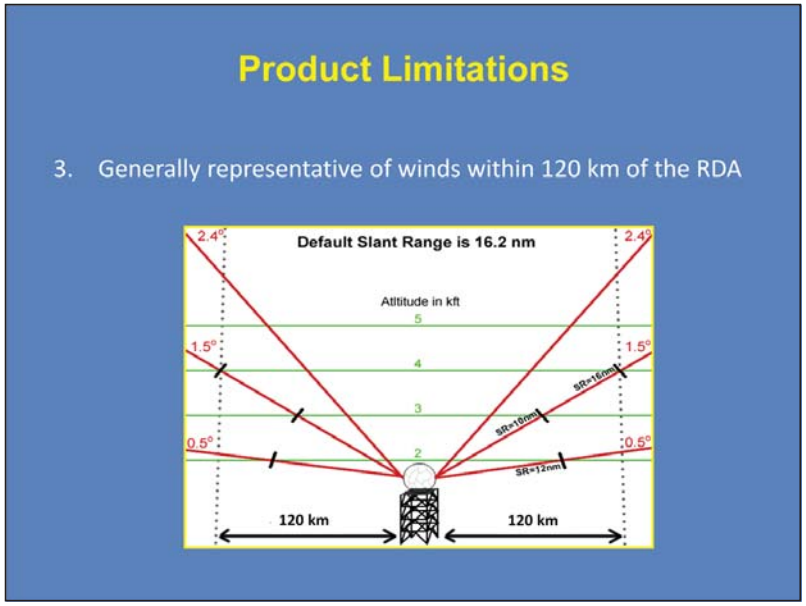
Here is an example of the Planview VWP product. Notice that you can plot not only one height, but multiple heights. This is great for getting a feel for the magnitude and direction of horizontal shear evolving throughout the day if you are anticipating severe weather. Obviously, the plots are dependent upon the availability of data to plot these wind speeds.



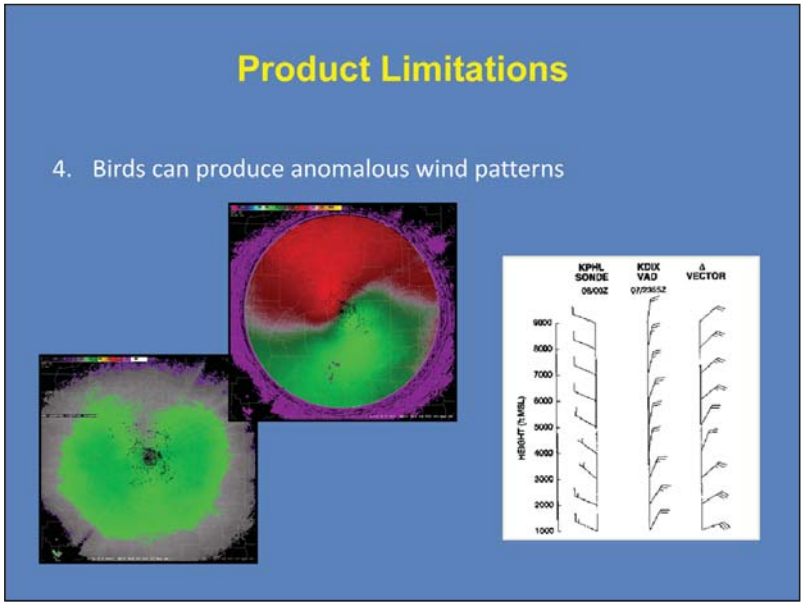
Now let's discuss the limitations of the VAD wind profile, we have to keep in mind the limitations of the VAD algorithm because that is what the product is based on. The first limitation is we need at least 25 points to produce a wind from the VAD algorithm to be plotted in the VWP. Here is an example of what happens when there are not enough data points. Since there are not enough points, it doesn't try to fit a sine curve and therefore it sends nothing to the VWP, and you'll see "ND" plotted where this occurs.



There is also another way ND is plotted, and that has to do with our second limitation. Winds will not be plotted if the RMS error or symmetry thresholds are exceeded. And, those values are 9.7 knots for RMS error, which is a measure of how well the plots fit the sine curve and 13.6 knots for symmetry, which is the measure of how homogeneous the wind flow is around the radar. So, if either of these thresholds are exceeded, then ND will be plotted as well.



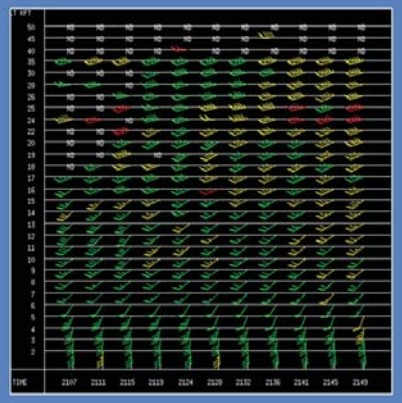
Our third limitation we have to keep in mind that the VAD winds plotted on the VWP are only generally representative of the winds within 120 km of the radar, because that's the limit of the slant range calculations. If you've got storms developing 150 km from the radar, you will want to use other sources of data to assess the wind profile in that area.



This is another limitation that ties back to the VAD algorithm. Flocks of migrating birds can affect the correctness of the VAD winds which are passed to the VWP. This can make you think you are developing a low level jet at night when in reality it is just the migrating birds. So, be mindful of this if you are in the path of migrating birds.

Product Applications

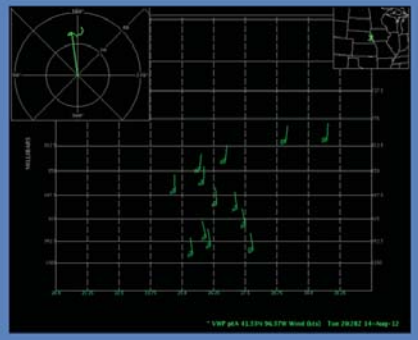
1. The VAD Wind Profile (VWP) will aid in the following operations:
   - Severe Weather
   - Aviation
   - Hydrology
   - Forecasting



The VWP product is useful in the severe weather, aviation, and hydrology sectors and forecasting in general. Obviously you can use the VWP to see how winds change with height in both direction and speed, plus you can analyze changes in time. You can also use it to estimate the depth of a cold air mass. For instance in the winter time, if you have a cold air mass moving in and you want to determine if it is dissipating or moving back north with time, you can use the VWP to animate and see that. That can help you decide what type of precip you are going to anticipate, either freezing rain, rain or snow.

Product Applications

2. The VAD Wind Profile (VWP) can be used to create or adjust hodographs and soundings



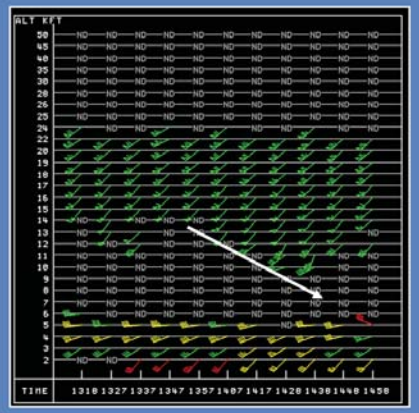
See RAC References VLab page for display instructions

The VWP can also be used to update hodographs for times in between upper air soundings.

Applications (Strengths)

3. Can indicate cloud depths and/or virga lowering with time

4. Can be used to estimate expected low-level trajectories for balloon launches

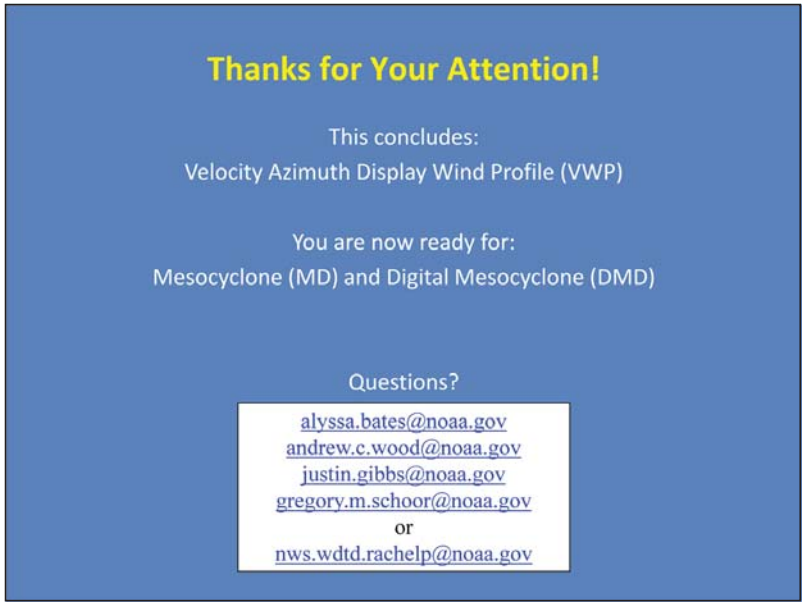


A surprising application of the VWP is the estimation of cloud depths or virga lowering with time. You can see that noted by the white arrow in the graphic. This is what virga descending will look like. As it descends it saturates the lower levels, in turn providing more data points, allowing the VWP to plot the data at those lower and lower heights. You can also use the VWP to estimate the expected low-level trajectories for balloon launches if you are trying to get a good track on a balloon as it's released.

Summary

- A composite vertical profile of VAD derived winds
- Excellent tool for forecasting, severe weather, hydrology, and aviation
- The VAD Wind Profile will display "ND" (no data) at a given height, if fewer than 25 data points exist, or if the symmetry or RMS error thresholds are exceeded

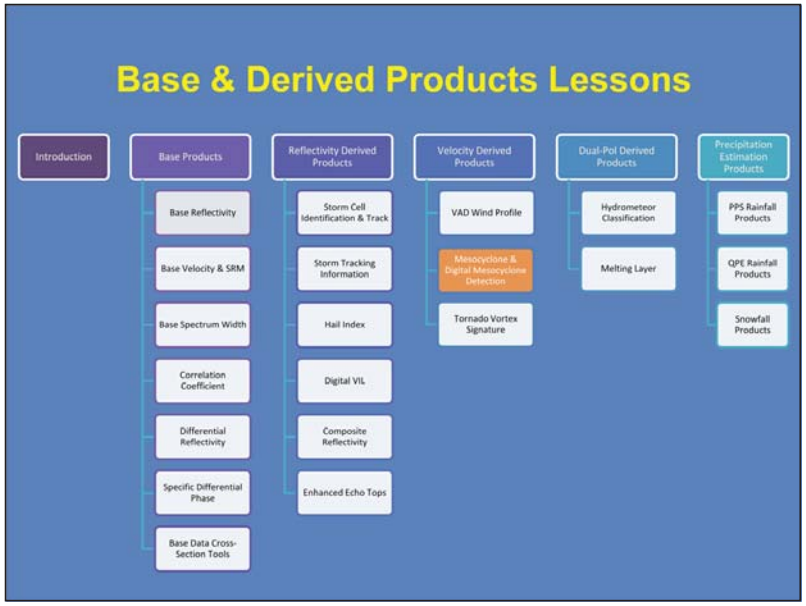
In summary, the VWP is a composite vertical profile of the VAD derived winds. It's a great tool for forecasting, severe weather, hydrology and aviation. The VWP will display ND at heights where there are fewer than 25 data points for the VAD algorithm or if symmetry or RMS error thresholds are exceeded. The next slide will be a short quiz on this lesson.



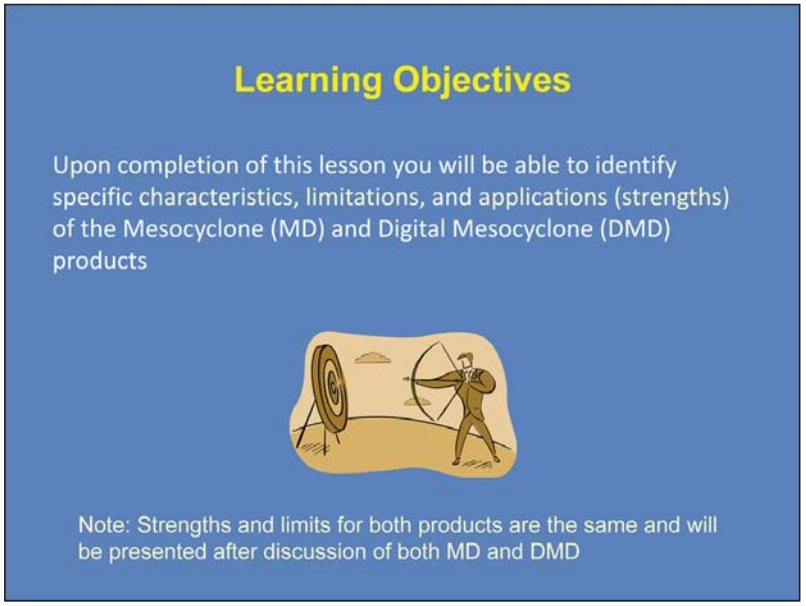
This concludes the lesson on the VWP. You can now move onto the next lesson on the Mesocyclone and Digital Mesocyclone products. If you have any questions, feel free to email the contacts listed below. See you in the next lesson!



Welcome to the next lesson on Velocity-Derived Products. This lesson will cover the Mesocyclone (MD) & Digital Mesocyclone (DMD) products.



Here is a roadmap for the lessons in this topic. This lesson on the Mesocyclone (MD) & Digital Mesocyclone (DMD) products, shaded in orange, is the second in the Velocity Derived Products Section of this topic.



Here are the learning objectives for this lesson. Notice the strengths and limitations are pretty much the same for both products. Advance the slide when you are ready.

Operator Defined Mesocyclone

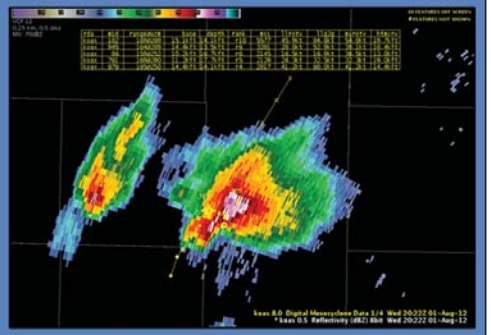
- Small scale rotation closely associated with a convective updraft that meets or exceeds established thresholds for:
  - **Persistence** (Minimum of two volume scans)
  - **Vertical extent** (Depth of at least 10,000 ft)
  - **Shear**
    - Core diameter < 5 nm
    - rotational velocity >= minimal mesocyclone strength in the table

Modified from Original NSSL Mesocyclone definition

Let's make a distinction real quick between what you as an operator would define as a mesocyclone versus what the algorithm will define as a mesocyclone. This definition comes from the original work done by NSSL done back in the 70s and 80s. An operator defined mesocyclone is defined as a small-scale rotation closely associated with a convective updraft that meets or exceeds established thresholds. One, it must be persistent, so it can't just show up and then disappear as will be the case often when viewing velocity data. Persistence in this case is going to be a minimum of 2 volume scans, but may need adjusting depending on the situation. Secondly, it needs to have vertical extent. By definition it needs to be at least 10,000 ft deep, but this again may be dependent upon the situation (like mini supercells). And it has to have a certain amount of strength, which is measured by shear. The core diameter must be less than 5 nm and the rotational velocity must equal or exceed a minimal mesocyclone strength that we'll see in a moment. So it has to be around for a while, have vertical extent and have some amount of strength.

Mesocyclone Detection Algorithm (MDA)

1. 1-D shear segments
2. 2-D features
3. 3-D features
4. Strength rank
5. Tracking
6. Feature parameters

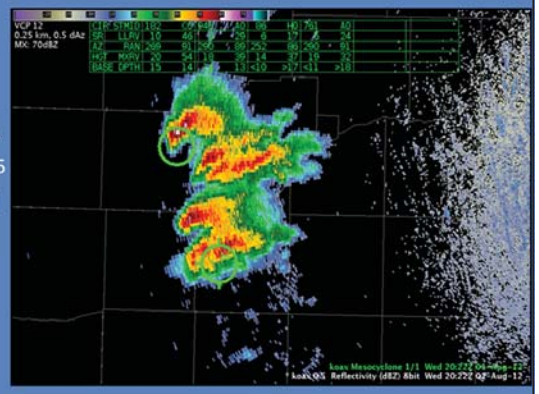


For more detailed information on how the algorithm works and adaptable parameters, see the RAC References VLab page

The mesocyclone detection algorithm will follow a similar path to the SCIT algorithm. Recall the SCIT used reflectivity as input and used building blocks to identify storms or centroids. The MDA will do a similar thing except use velocity as input. It will first identify what it calls 1-dimensional shear segments. It will combine those into 2-D features. It will then look vertically and build a 3-D feature we'll call a circulation. A strength rank will then be assigned to that circulation, and depending on the strength, you may call it a mesocyclone. It will then track the feature and give you a whole bunch of information on that circulation. For more detailed information on how the algorithm works, as well as information about the adaptable parameters associated with this algorithm, see the RAC References VLab page for this lesson.

Mesocyclone (MD) Characteristics

- Available at end of volume scan
- Icons
  - Thin circles = SR 1-4
  - Thick circles = SR ≥ 5
- Four spikes if found at lowest tilt or < 1 km ARL
- Attribute table available



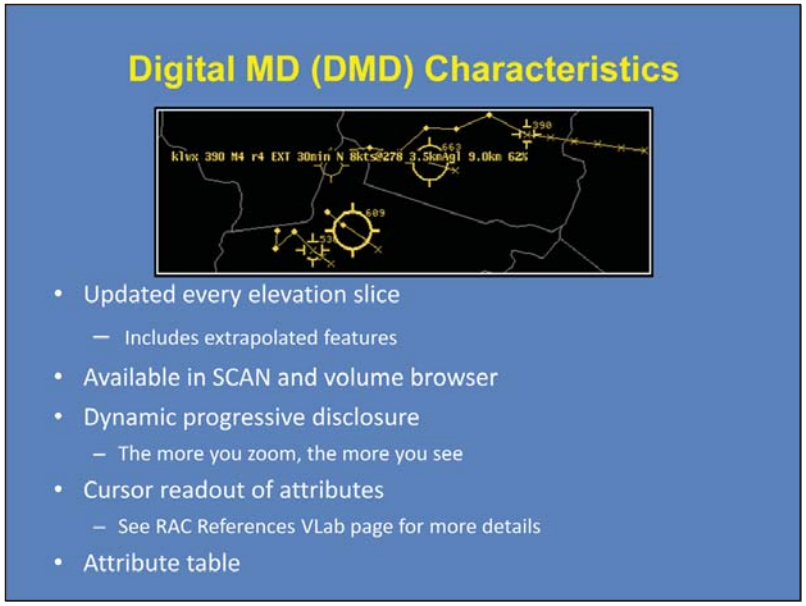
Enough about the algorithm. Now, let's get to the products. First, we'll talk about some of the characteristics of the mesocyclone detection product. It is available at the end of a volume scan and you'll see a few different icons with it. You'll see thin circles if the strength rank is 1 to 4, and thick circles for strength ranks greater than or equal to 5. If you see four spikes on these circles, then that tells you that the circulation was found on the lowest tilt or less than 1 km ARL. And, there's an attribute table that is available.

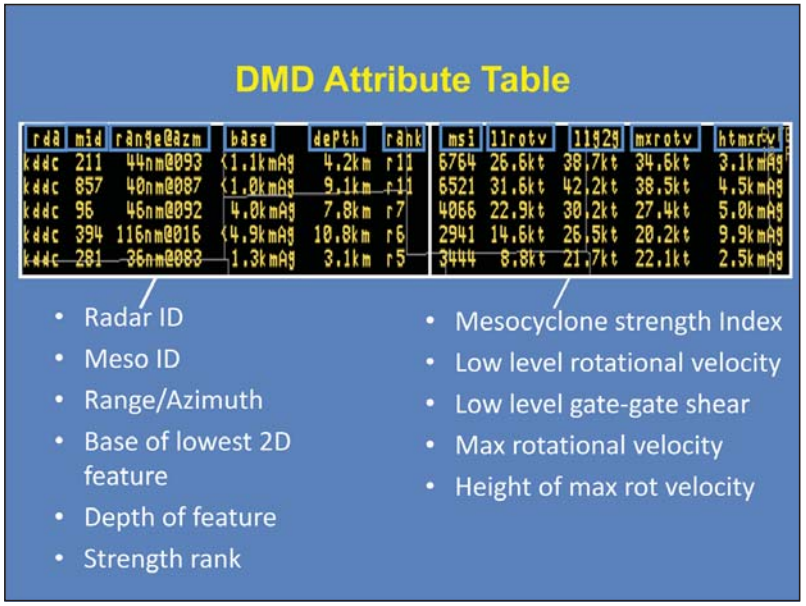
MD Attribute Table

CIR	STMID	353	I4	29	I4	644	W7
SR	LLRV	10	26	9L	53	9L	41
AZ	RAN	128	81	127	78	123	84
HGT	MXRV	14	68	16	61	10	41
BASE	DPTH	< 9	>17	< 9	>15	<10	>13

- CIRC STMID = Circulation ID Number and ID of closest SCIT storm
- SR = 3-D strength rank
- LLRV = Lowest level rotational velocity (kts)
- AZ = Azimuth of lowest 2D Feature
- RAN = Range of lowest 2D Feature
- HGT = Height of maximum rotational velocity (Kft)
- MXRV = Max Rotational Velocity in the feature (kts)
- BASE = Altitude of lowest 2D Feature (Kft) ("<" shown if Base on 0.5° slice)
- DPTH = Depth of 3D Feature (Kft) (">" shown if base on 0.5° slice)

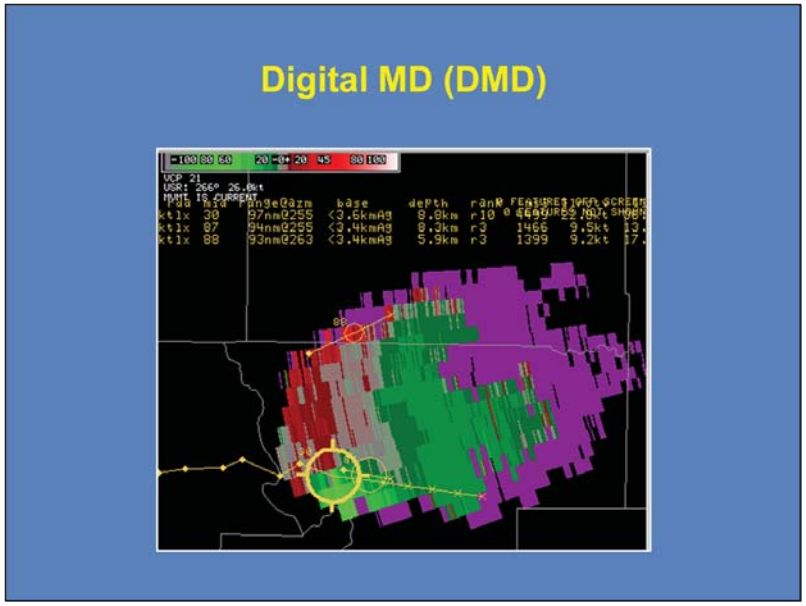
Let's take a moment to go through the MD attribute table. First you'll have a circulation identification number and this is an individual mesocyclone identification number, in this case 353. Next to that is the ID of the closest identified storm output by SCIT. Next, you have a strength rank, and it's a 3D strength rank, and then the low level rotation velocity associated with the feature. Next is the azimuth and range to the lowest 2D feature. The height is the height of the maximum rotational velocity. And finally you have the base of the feature which is the height of the lowest 2D feature, and you'll see the less than sign if the base is on the 0.5 deg slice indicating it could be lower than that. And, depth of the 3D feature with a greater than sign if the base is on the 0.5 deg slice since the depth could be greater than what is depicited.



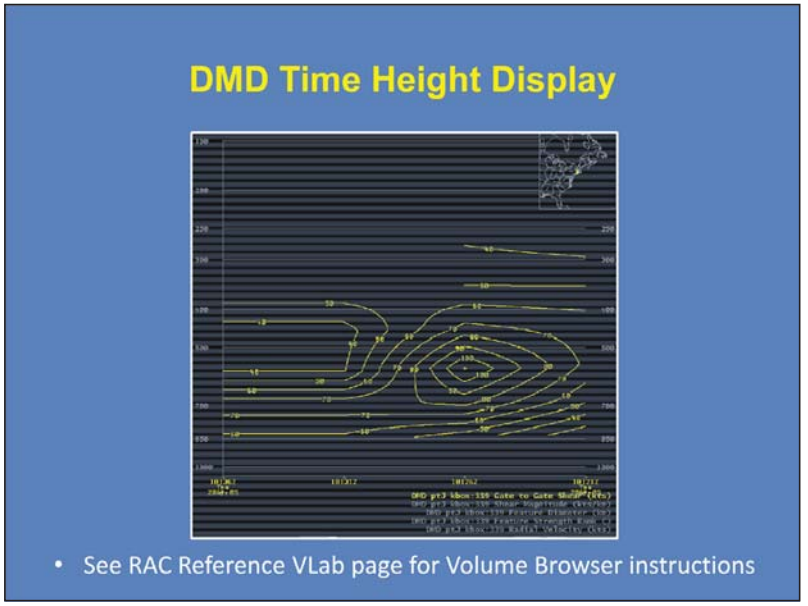


Here is the Digital Mesocyclone Detection product, or DMD. It is going to give you a little more information and done a little bit differently. First, it is updated every elevation slice. They wanted to give you some intermediate output instead of always having to wait until the end of the volume scan. Therefore, this product will have extrapolated features which appear as broken circles on the display. And those show you where the algorithm thinks the new update will be, and will become a full solid circle when the feature is correlated. Data are available in SCAN as well as the volume browser. There is dynamic disclosure with this product which means the more you zoom, the more detail you will see. And, cursor readout has a lot of the attributes available.

The DMD product has its own attribute table. On the left side we see information related to the radar ID, meso ID, azimuth and range, the base of the lowest 2D feature, the depth of that feature as well as the strength rank. On the right side we see a mesocyclone strength index, a low level rotational velocity, low level gate to gate shear, max rotational velocity and the height of the max rotational velocity.



This is what the DMD product looks like and we have overlaid on top of a base velocity image.



This is what you get when you send of a request for this data from the volume browser. You can see the time height series for the gate to gate shear and you can see the time characteristics and if it is strengthening with time. All of this is dependent of course on how well the algorithm has identified and maintained this feature.

Viewing DMD Output via SCAN

- More effective at viewing time/height series
- Easier to read
- Must start SCAN DMD Application
- More at workshop



Using the volume browser time-height option can be cumbersome. Another option you have is the DMD Icons and Table via the SCAN menu. This is a more effective way of viewing the time-height series and is easier to read. It does require you to start the SCAN application at your workstation and we'll cover that more at the workshop.

MD/DMD Display Control

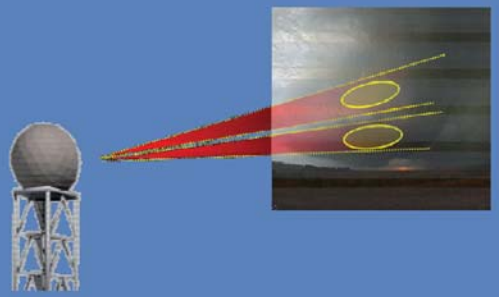
- Determine display attributes:
  - Extrapolated features
  - Minimum strength rank
  - Past & forecast tracks



As with many other algorithms, you can affect how the DMD data are displayed at your workstation by using the radar display controls sub menu that you are familiar with already. Here you can decide if you want to see extrapolated features. You can also decide what DMD features to see, what strength rank to filter, and whether to see past or forecast positions.

Products' Limitations

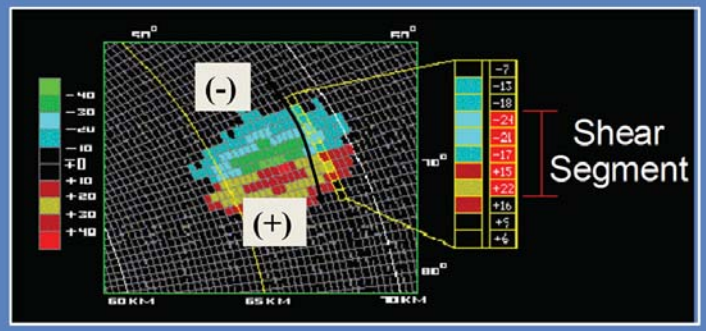
1. Detects numerous false detections and weak, shallow insignificant circulations
   - Doesn't require 10Kft depth



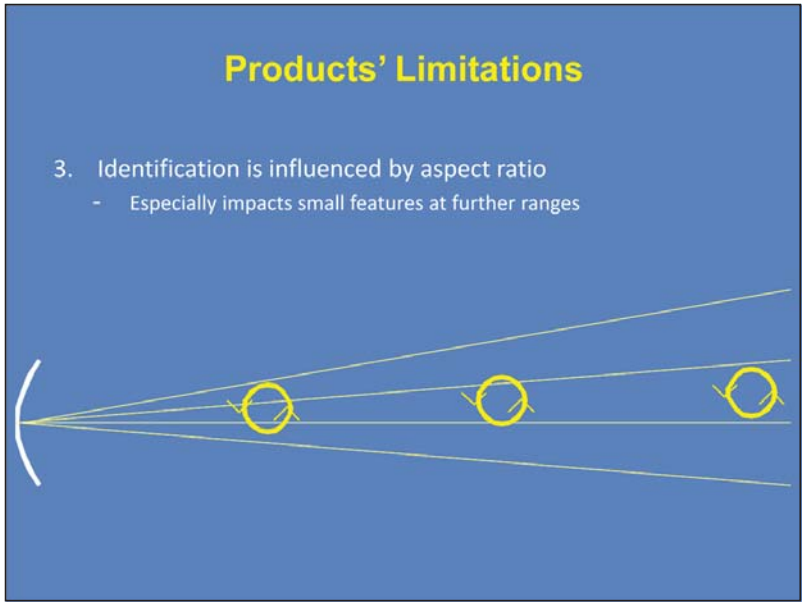
So let's look at the limitations of the MD/DMD products since they are pretty similar. The first is that we do not need a 10kft depth for the circulation but rather only 2 elevation cuts, which is different from the original operator defined mesocyclone. This means that you can get some weak, shallow features identified by the algorithm.

Products' Limitations

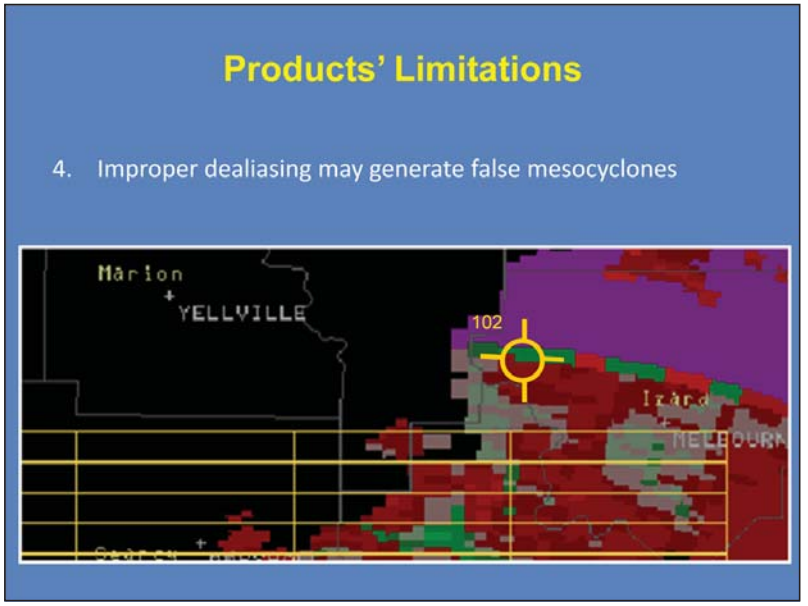
2. The algorithm only detects cyclonic rotations



Because we are searching for increasing velocities as the radar rotates clockwise, we are only searching for cyclonic rotations. We are not able to detect anticyclonic rotations.



Whether a feature is identified is influenced by aspect ratio. And by this, we mean the relationship between the size of the feature relative to the size of the beam. This primarily affects small features at small ranges where the feature can be contained solely within one beam.



If you get improper dealiasing that has vertical continuity, you can get some false detection. This can really be an issue at far ranges where you only need one elevation angle. That is why it is important that you always go back to the base data to verify features in the derived products.

Products' Limitations

5. Default settings are for large deep supercells

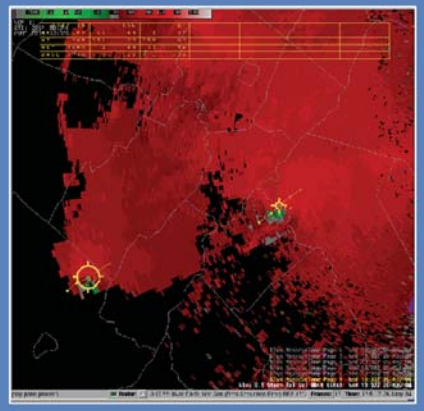
6. Numerous detections of circulations may require changes in adaptable parameters



Limitations 5 and 6 have something to do with the adaptable parameters for the algorithm. First, the default setting is for large, deep supercells and this is fine if that's what you have, but it may not be fine on other occasions. If you get a lot of detections, you may need to tweak these parameters such as increase min reflectivity, or change the values for what is displayed on your workstation.

Products' Applications

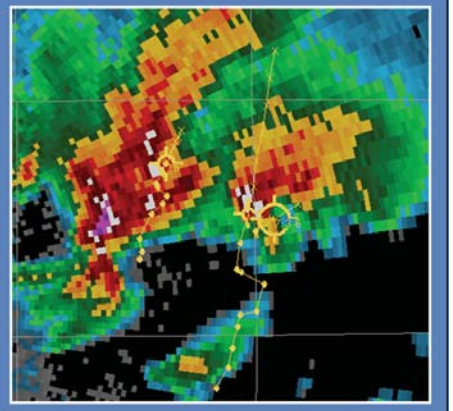
1. Identify mesocyclones, even those that are shallow



The strength of these products is that they can identify mesocyclones, even those that are shallow. It is good practice to always overlay the MD/DMD products on top of base velocity or SRM that way you can verify their existence. In this image, you can see there is good correlation between the algorithm and the base data.

Products' Applications

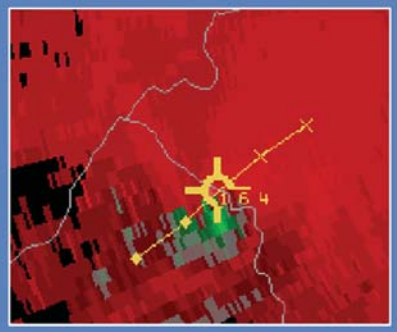
2. Weak circulations detected
   - Display by lowering minimum strength rank



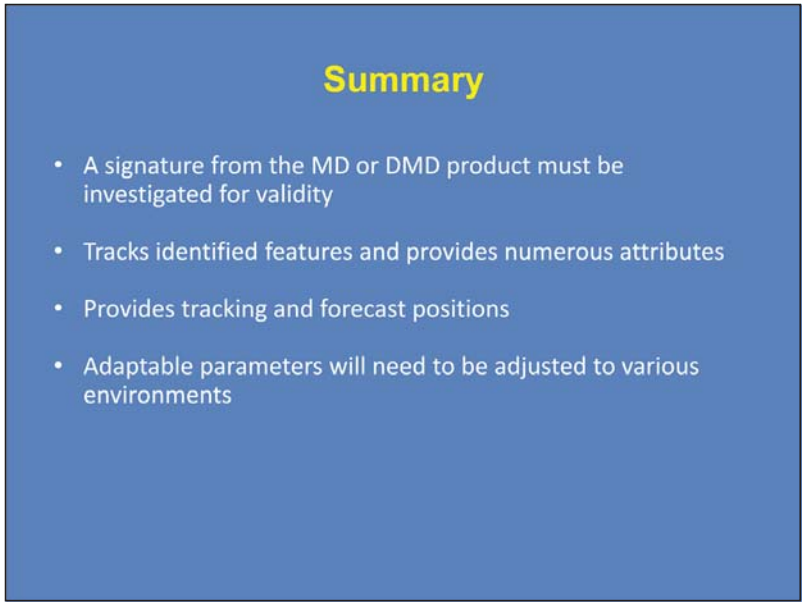
You can see weak circulations but only if you have the minimum display threshold set low enough. By default, 5 is the minimum strength rank displayable, so to see the thin circles, you have to set your display threshold to strength rank 1-4.

Products' Applications

3. Tracking attempts to account for time continuity



This product does attempt to give you time continuity by showing you the past and forecast positions.



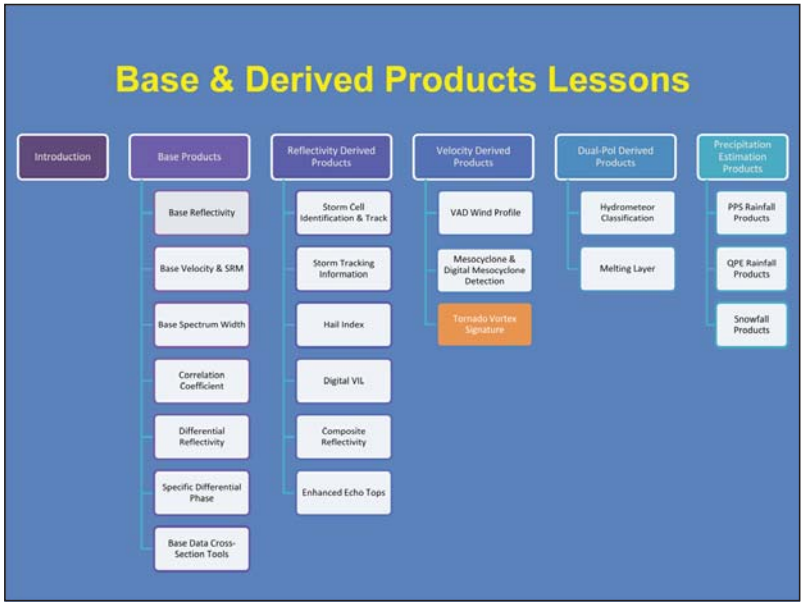
To summarize these products, the DMD or MD products must be validated by looking at velocity or SRM displays, correlate it with reflectivity and understand the environment the feature is located in. We track these identified features and provide numerous attributes as well as past and future positions. And, there are adaptable parameters that you can adjust, and probably should adjust, based on various environments. The next slide will be a short quiz on the DMD/MD products.



This concludes the lesson on the MD/DMD products. You can now move onto the next lesson on the Tornado Vortex Signature. If you have any questions, feel free to email the contacts listed below. See you in the next lesson!



Welcome to the next lesson on Velocity-Derived Products. This lesson will cover the Tornado Vortex Signature.



Here is a roadmap for the lessons in this topic. This lesson on the Tornado Vortex Signature, shaded in orange, is the third and last in the Velocity Derived Products Section of this topic.

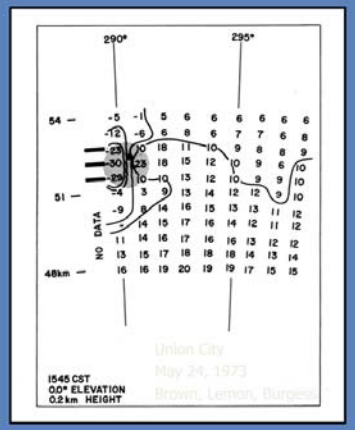
Learning Objectives

Upon completion of this lesson you will be able to identify specific characteristics, limitations, and applications (strengths) of the Tornadic Vortex Signature (TVS) product



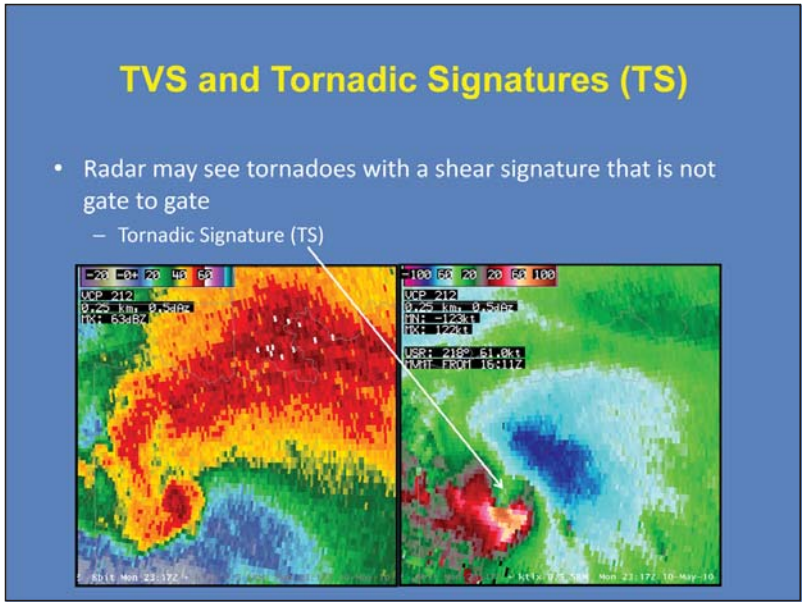
Here are the learning objectives for this lesson. The next slide will feature Les Lemon telling his account of the history of the TVS. Advance the slide when you are ready to get started.

TVS History

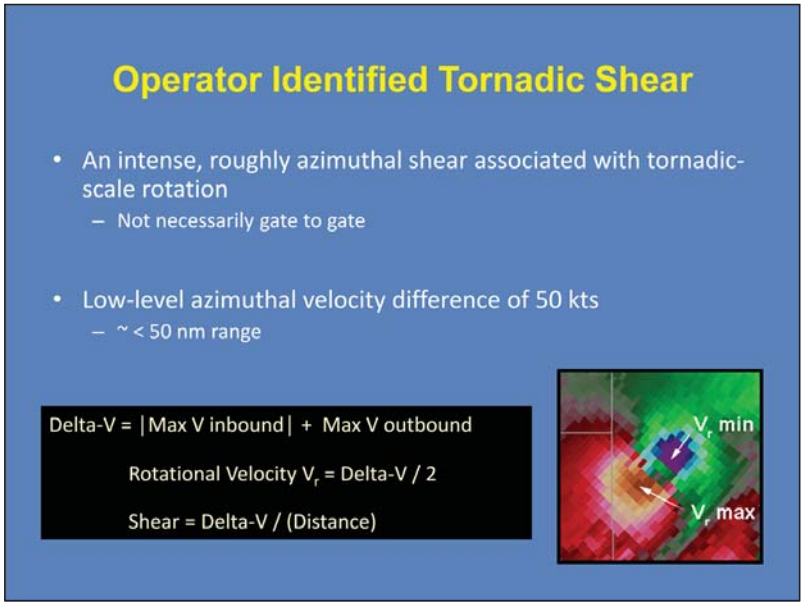


- Large NSSL Study (1973)
- Mesonet, aircraft, spotters
- Radar collected but not analyzed in real time
- Later noted a gate-to-gate shear (GTGS)
- Found to coincide with the Union City Tornado
- Renamed GTGS to TVS

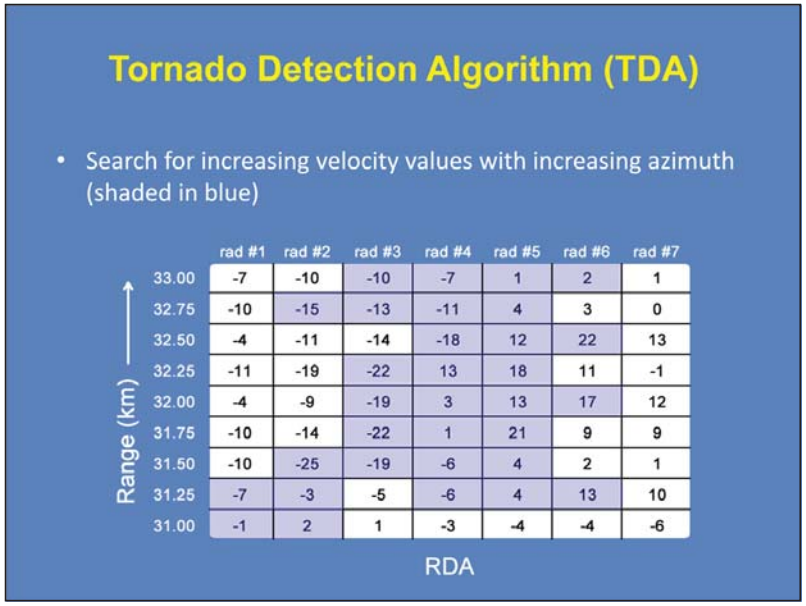
Back in the early days of NSSL, we had a spring data collection (1973). We were trying to sample storms with radar from inception to decay, and we were lucky to see the storm that passed over Union City, OK. We collected data on the storm and one thing to note is we could not see the data in real time. The processors were just not fast enough. So, after data collection, in fact during the winter after that event, we were analyzing the B-scans. In looking at them we started seeing what you can see in the image of intense gate-to-gate velocity differences. We thought it was a problem with the radar. Nobody anticipated seeing this signature. When we looked at higher elevations we saw something similar, so we knew this was something unique here and not something caused by an artifact or bad radar measurement. We called this a gate-to-gate shear (GTGS). We knew it was associated with a tornado when we launched a rocket from the damage path from a point where we collected data and where we saw the TVS on the ground and this rocket echo was in the exact same spot. We realized then this was a tornadic signature. But we only called it a gate-to-gate shear because we didn't know how prevalent it would be. Later we started seeing it in more situations but sometimes it was only aloft, never touching the ground. So, this is the origin of the TVS or tornadic vortex signature.



For many years the term tornadic vortex signature has been used to indicate gate to gate shear associated with tornadoes. Since 2008, WSR-88D's have been able to display super-resolution base products. As a result, the radar has been able to sample the actual tornado, and the shear signature associated with some tornadoes has not been gate-to-gate. Notice in this example in the lower right that there are a few very weak super-res velocity gates in between the maximum inbound and outbound velocities, yet that is an EF4 tornado from a range of 28 nm. These types of rotational signatures are called tornadic signatures and will be covered in depth in Topic 7. The tornado detection algorithm, which is what this lesson is based on, only searched for gate to gate signatures and did not detect a TVS with this tornado.

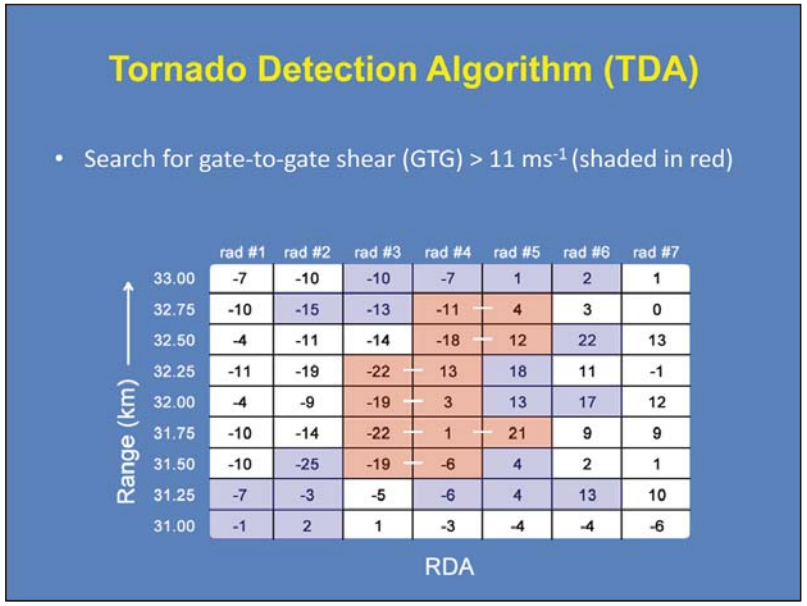


With the previous 2 slides in mind, I'd like to define operator identified tornadic shear…which is what you the forecaster will be looking for in the base products. Tornadic shear is an intense, roughly azimuthal shear associated with tornadic-scale rotation. It does not have to be gate-to-gate but should be at nearly the same range and the max inbound and outbound velocities shouldn't be separated by more than a radial or two in super-res velocity. The low-level azimuthal velocity difference should be at least 50 kts. Because of beam height increasing with range from radar, and beam width increasing, you normally cannot sample tornadic scale winds outside of about 50 nm. To avoid confusion with terms, Delta V is the velocity difference, and it just the addition of the absolute value of the maximum inbound and outbound velocities. Rotational velocity, Vr, is delta-V over 2. Finally, shear is delta-V divided by the distance between the maximum inbound and outbound velocities. Topic 7 will cover these topics at great length. The rest of this lesson focuses on how the algorithm identifies tornadic vortex signatures.



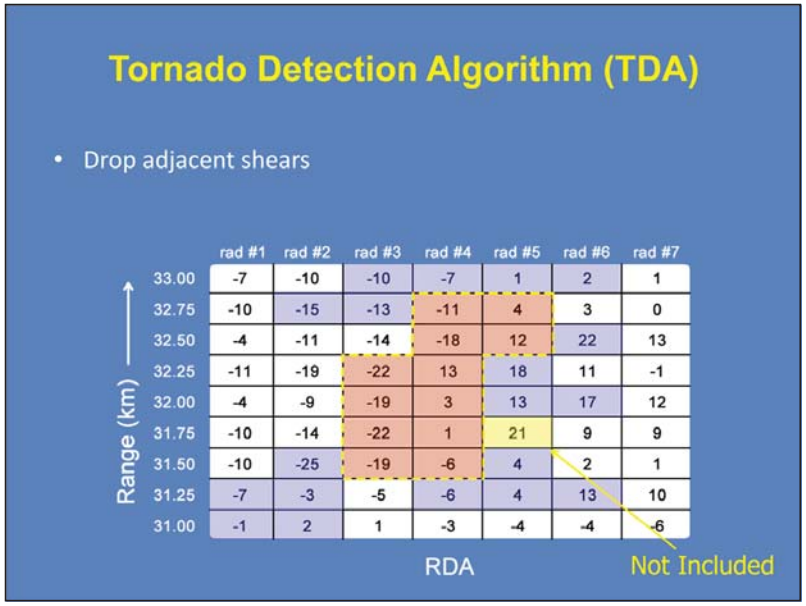
	rad #1	rad #2	rad #3	rad #4	rad #5	rad #6	rad #7
33.00	-7	-10	-10	-7	1	2	1
32.75	-10	-15	-13	-11	4	3	0
32.50	-4	-11	-14	-18	12	22	13
32.25	-11	-19	-22	13	18	11	-1
32.00	-4	-9	-19	3	13	17	12
31.75	-10	-14	-22	1	21	9	9
31.50	-10	-25	-19	-6	4	2	1
31.25	-7	-3	-5	-6	4	13	10
31.00	-1	2	1	-3	-4	-4	-6

We will now dig into the Tornado Detection Algorithm (TDA). Because the radar spins in a clockwise direction, it looks for increasing velocities along adjacent azimuths in a clockwise direction. The data in this table are velocity values and anything that is increasing from left to right, that's increasing azimuth, is shaded blue. Searching in a clockwise direction for increasing velocities means that the TDA only looks for cyclonic circulations which in the northern hemisphere is counterclockwise. The vast majority of tornadoes in the northern hemisphere spin counterclockwise.

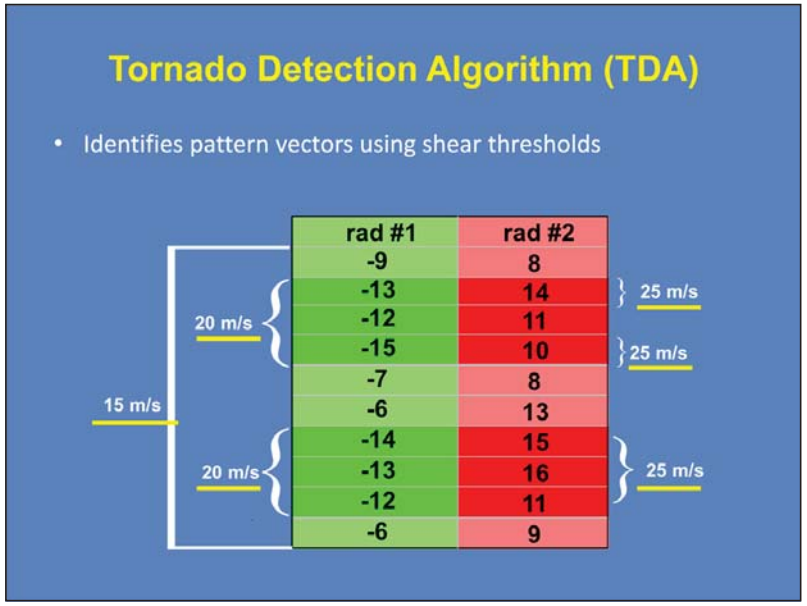


	rad #1	rad #2	rad #3	rad #4	rad #5	rad #6	rad #7
33.00	-7	-10	-10	-7	1	2	1
32.75	-10	-15	-13	-11	4	3	0
32.50	-4	-11	-14	-18	12	22	13
32.25	-11	-19	-22	13	18	11	-1
32.00	-4	-9	-19	3	13	17	12
31.75	-10	-14	-22	1	21	9	9
31.50	-10	-25	-19	-6	4	2	1
31.25	-7	-3	-5	-6	4	13	10
31.00	-1	2	1	-3	-4	-4	-6

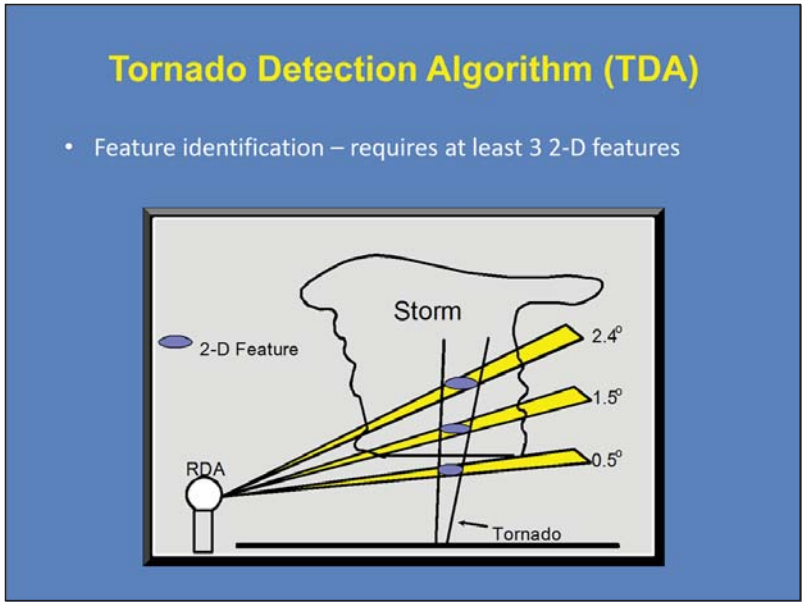
It then calculates the actual values of shear greater than 11 m/s, or 22 knots. The values of shear greater than 11 m/s, or 22 knots, are shaded in a reddish color.



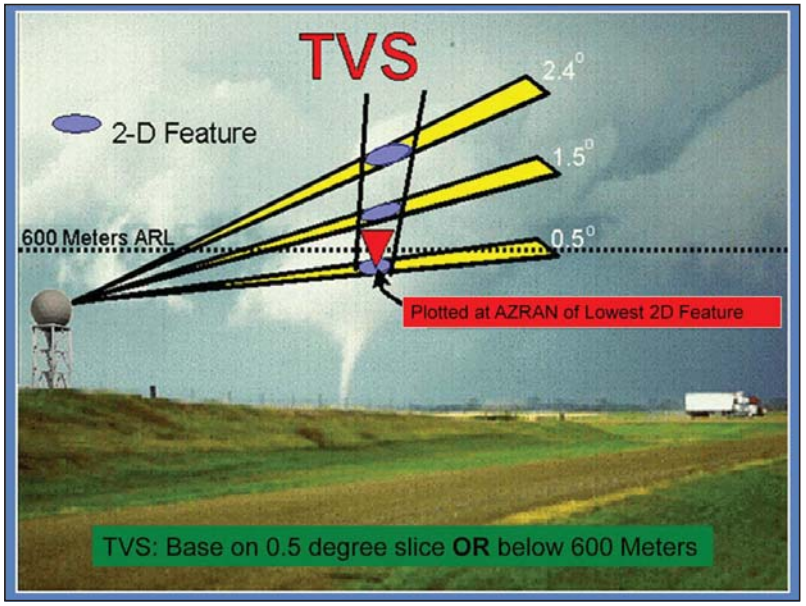
The TDA then uses a pattern matching scheme. Notice that everything colored red in this table has fairly strong negative, or inbound velocities on the left, and positive, or outbound velocities on the right. The 21 m/s outbound velocity which is colored yellow, is not included in this 2D feature because it does not fit the pattern of the rest of the signature.



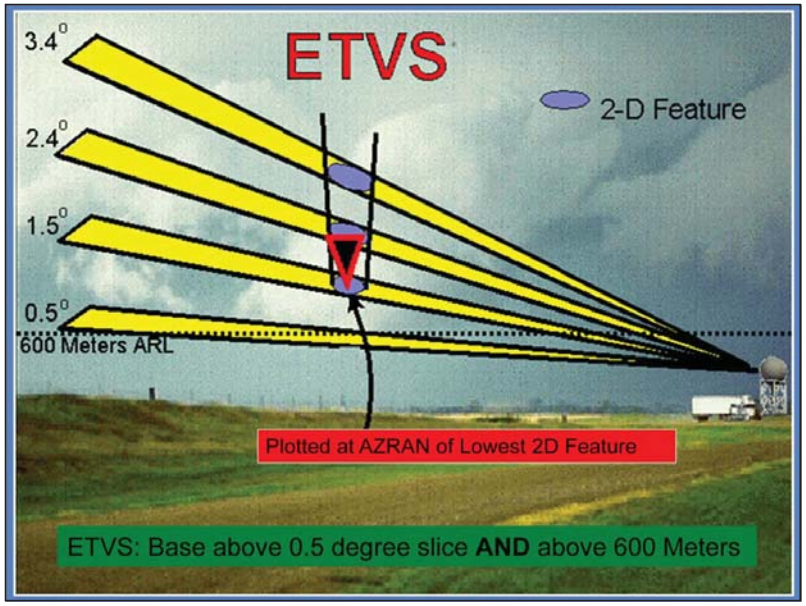
So, at this point, the algorithm has its 2D feature containing all gate to gate shears greater than 11 m/s, or 22 knots. The next step is to apply shear thresholds to all of the 2D features. Again, the velocities here are in meters per second. So, it first applies a 15 m/s threshold which includes the entire 2D feature. When a threshold of 20 m/s is applied, two features containing 3 gates each are identified. Finally, a 25 m/s threshold comes about, and three separate features are identified. The algorithm starts with the higher thresholds first in order to find a match with higher elevation angles first, which then helps to determine if it's a TVS signature or not. We'll talk about how it tries to correlate vertically next.



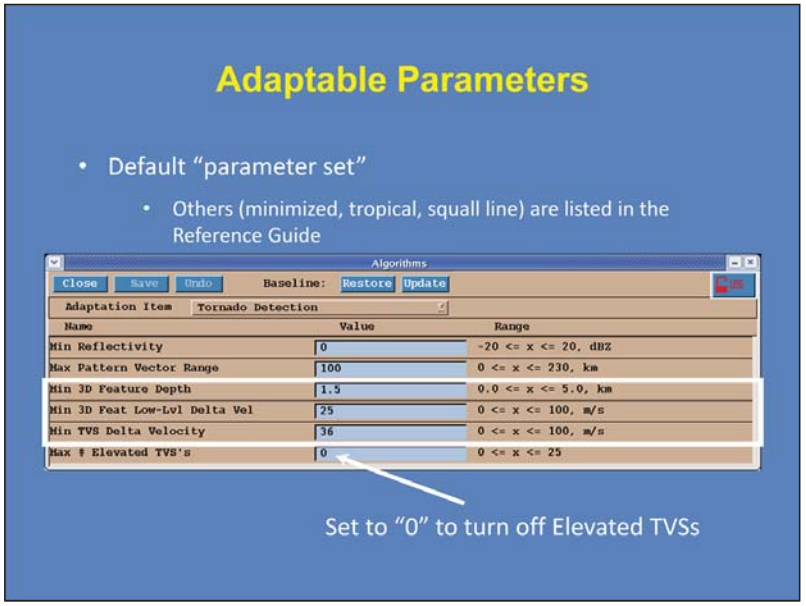
The final step in determining at TVS (starting with the strongest 2D circulations first) it tries to match at least 3 vertically correlated 2D circulations, and that would declare it a 3D circulation. Ideally, there will be no gaps in the elevation angles between the vertically correlated 2D circulations. However, a one elevation gap is allowed to account for base data issues such as range folding and velocity dealiasing. In the graphic shown here, 3 successive elevation angles have 2D features that match, therefore this is a TVS.



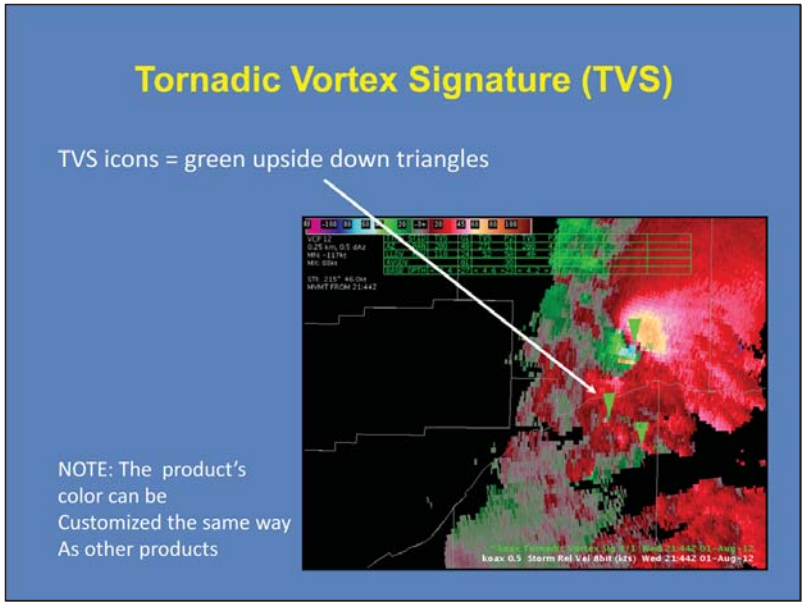
We saw just a few slides ago that for a TVS to be identified, there needs to be at least 3 vertically correlated 2D circulations identified and ideally these vertical correlations will have no gaps, but a one elevation gap is permitted for unforeseen data quality issues like range folding. Anyhow, one additional criteria is the lowest 2D feature must be on the 0.5 deg slice or below 600 meters ARL, which is around 2000 feet ARL. Provided that the TVS is on the 0.5 deg elevation angle or below 600 m, or less than 2000 feet ARL, the TVS symbol, which is a red, isocoleses triangle will be plotted at the azimuth and range of the lowest 2D feature.



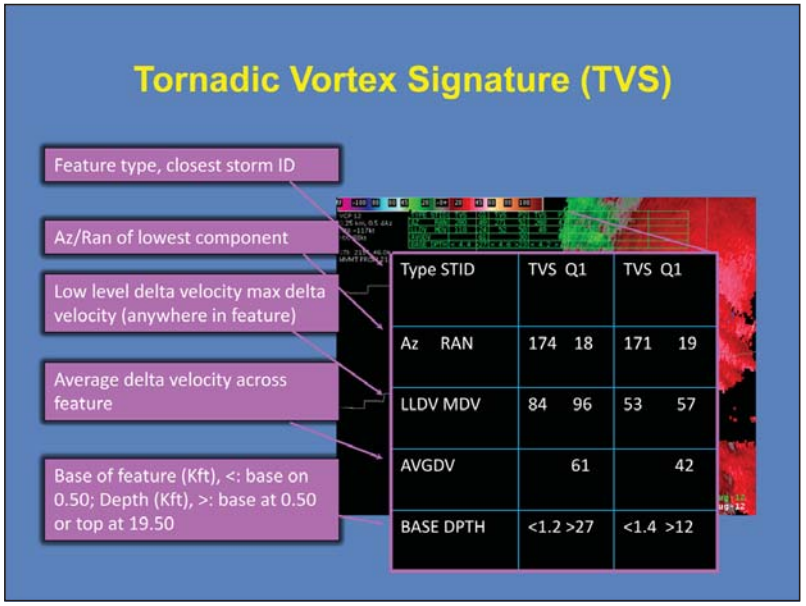
Even though these are turned off by default, we do need to examine how the ETVS is identified. As with the TVS, it needs at least 3 vertically correlated 2-D features. However, the constraint of this feature is the lowest 2D circulation needs to be above the 0.5 deg slice AND above 600 meters, 2000 feet ARL. An ETVs icon is plotted at the azimuth/range of the lowest 2D feature identified. The icon is an OPEN, red isocolese triangle.



Here are the adaptable parameters at the RPG which are listed in the Adaptable Parameters Handbook. The ones listed in the white box are for the default parameter set. These 3 can be changed as a set as outlined in the parameter handbook. For just about any meteorological situation that you might be expecting, for instance, mini-supercells, tornadoes associated with tropical systems, squall lines; those are all sets listed in the handbook that you can set at the RPG. The other 3 not listed in the white box…Min Reflectivity, Max Patter Vector Range and Max # of Elevated TVS's, can be set independently. As we'll talk about in just a little bit, elevated TVSs, by default are turned off due to high false alarm rates. And, any evidence for operational utility or research to back up ETVSs and tornado occurrence is lacking. Therefore, we recommend turning off ETVSs until research provides evidence to justify turning them on.



Here is what the TVS product looks like in AWIPS-2 CAVE. In this graphic I have the TVS overlaid on top the base SRM product. The closest storm ID from SCIT is plotted next to each TVS icon symbol.



Notice there is also an attributes table. The first row gives you the circulation type and the closest storm ID to that circulation feature. The second row is the azimuth and range of the lowest 2D feature. The third row gives you the low-level delta velocity in knots, and the max delta velocity in knots, anywhere associated with that feature. The fourth row gives you the average delta velocity of the feature. And, finally, the bottom row gives you the base of the feature in thousands of feet. Anytime you see a less than sign associated with the base, this tells you that the lowest 2D feature is on the 0.5 deg scan. The other part of the bottom row is depth, in thousands of feet. If you see a greater than symbol, that means that the base is at 0.5 deg or the highest 2D feature is on the 19.5 deg scan.

Radar Display Controls

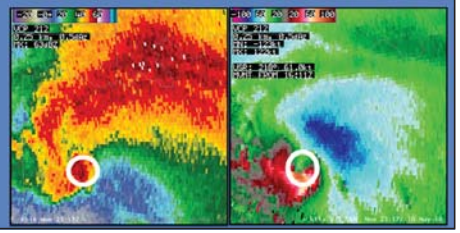
- Change AWIPS-2 workstation display to not view ETVS symbols
- CAUTION!
  - Only changes AWIPS-2 display!
  - You will not see ETVS
    - Outside users will see ETVS



Though by default ETVSs are turned off at the RPG, let's say you are using them. In AWIPS-2, you can control whether they are displayed or not using the radar display controls window. Partway down there is a little box you can check that will turn off ETVS icons when checked and turn them on when not checked in your CAVE window. However, this only changes the AWIPS-2 display, so outside users will see them despite you not seeing them. The only way to prevent outside uses from having there display cluttered with ETVSs is to have them turned off at the RPG.

TVS Operational Considerations

- TVS detection...consider:
  - Low-level wind shear and thermal profile
  - Signature's position in relation to reflectivity storm structure
  - Time continuity and range
- Beyond ~60 km (~45 nm), TVS most likely triggered by strong mesocyclone
  - TDA independent of MDA



The next several slides deal with operational considerations for incorporating TVS into the warning decision process. If you see a TVS consider the following: What is the low level wind shear and thermal profile in the near storm environment? Are they both favorable for tornadoes? Next, look at the TVS in relation to the reflectivity structure. Does it make sense where this TVS is located relative to the storm? Finally, is there time continuity and how far from the radar is this feature? Is the TVS associated with a mesocyclone that has been around a while? That is good time continuity. The problem with range, is that beyond 60 km, or 45 nm, the TVS is most likely triggered by strong mesocyclones, and not necessarily indicative of an impending tornado. Recall that the TDA is independent of the MDA, so at far ranges, they most likely will be identifying the same feature.

TVS Operational Considerations

- TDA shows time and space continuity
- TDA identifies TVSs nearly continuously on supercells
- TDA has poor performance in squall lines
- TDA tends to identify TVSs near the bend in a LEWP

Continuing on with operational considerations, the TDA does show some time and space continuity allowing you to track features. Also, the TDA nearly continuously identifies TVSs on supercells which could be a good or a bad thing. It's bad in that the false alarm is very high in these supercells. But it could be a good thing in maintaining your situational awareness that the storm you are looking at is still a supercell and still warrants a continuous and thorough base data analysis. TDA also performs very poorly in the squall lines depending on how the leading edge of the squall line oriented relative to the radar. You can have large areas of shear show up that aren't really tornadic by any means, but still they show up because they are areas of shear. This is especially true near the bends in LEWPs.

TVS Operational Considerations

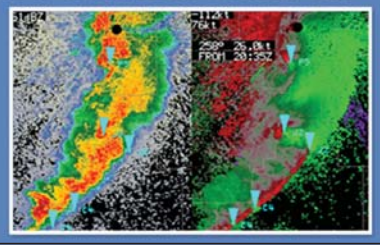
- Elevated TVSs correlation worse than TVSs
  - Recommend keeping them turned off at the RPG
- NEVER issue tornado warnings solely on TVS/ETVSs
  - Remember that algorithms serve to provide users with guidance



Some final operational considerations. The elevated TVS has worse correlations than TVSs relative to tornado occurrence. So, it is recommended that you keep the ETVS turned off at the RPG until evidence says otherwise. We would be doing our customers a great dis-service if we issued tornado warnings based solely on the TVS signature. Remember that these algorithm serve to provide users with guidance only. A thorough base data analysis coupled with near storm environmental analysis is critical in any tornado warning decision.

Product Limitations

1. Adaptable parameters need more research
2. High false alarm ratio, especially in squall lines and tropical cyclones
3. Little research has been done to date relating the occurrence of tornadoes to Elevated TVSs



Let's examine the limitations. First, adaptable parameters need more research. Second, high false alarm ratios occur especially with squall lines and tropical cyclones. And finally, there has been little research done to date, relating the occurrence of tornadoes to ETVs.

Product Applications

1. Searches for gate-to-gate shears
2. Multiple velocity difference thresholds
3. Provides information on the base and depth of circulations (attributes table)
4. Allows performance tuning through adaptable parameter changes

Here are the applications of the TVS. First, it searches for all gate-to-gate shears. So, after doing a thorough base data analysis, you can use the TVS product to see if you missed anything. Second, it does use multiple velocity difference thresholds. The product is displayed with an attributes table which provides information on the base and the depths of the circulations. And, finally, it allows for performance tuning through adaptable parameter changes based on the meteorological conditions you anticipate for an event.

Summary

- TVS product can alert operator of significant and possibly tornadic circulations
- TDA performance within squall lines and tropical cyclones is significantly worse than supercells





Image courtesy WFO ILX website

Summarizing this lesson, the TVS can alert the operator of significant and possibly tornadic circulations. Be careful using TVS products in squall line and tropical cyclone events as the TDA performance in these situations is significantly worse than in the supercells.

Thanks for Your Attention!

This concludes:
Tornado Vortex Signature (TVS)

You are now ready for:
Hydrometeor Classification (HC)

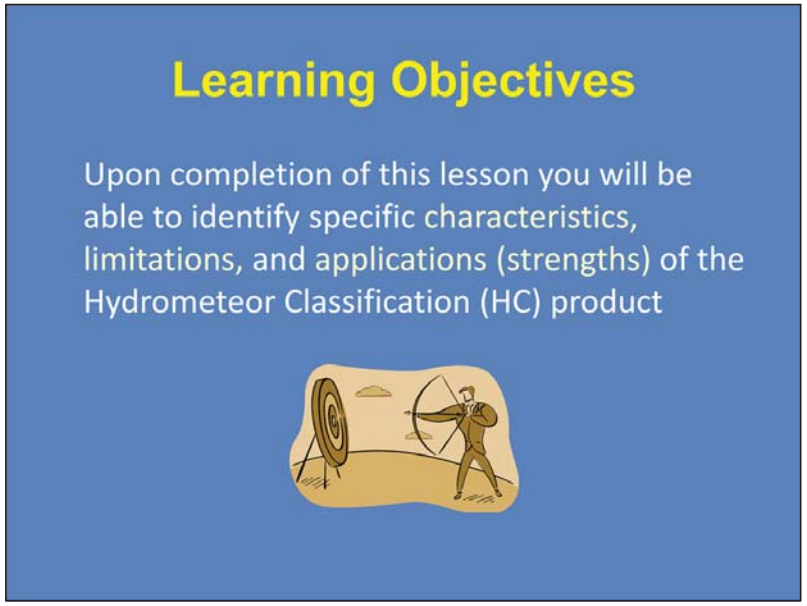
Questions?

alyssa.bates@noaa.gov or
andrew.c.wood@noaa.gov or
justin.gibbs@noaa.gov or
gregory.m.schoor@noaa.gov
or
nws.wdtd.rachelp@noaa.gov

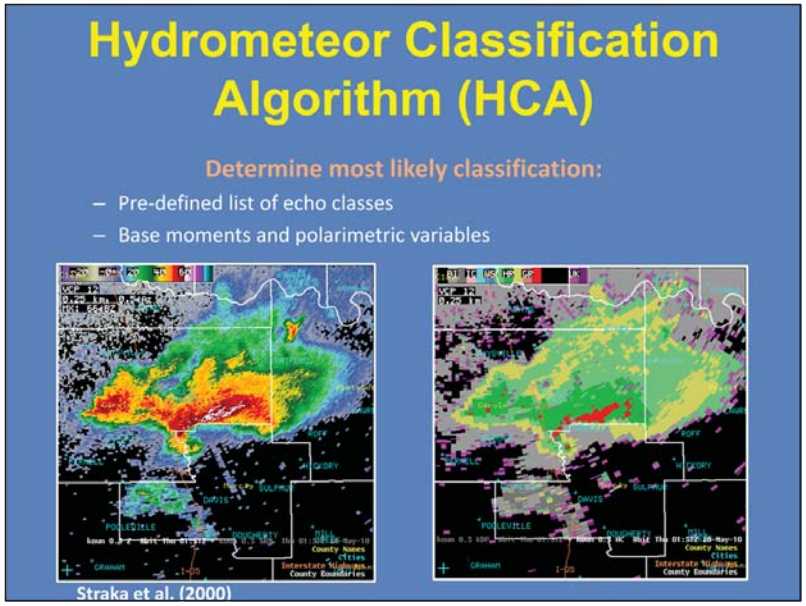
This concludes the lesson on the VWP. You can now move onto the next lesson on the Mesocyclone and Digital Mesocyclone products. If you have any questions, feel free to email the contacts listed below. See you in the next lesson!



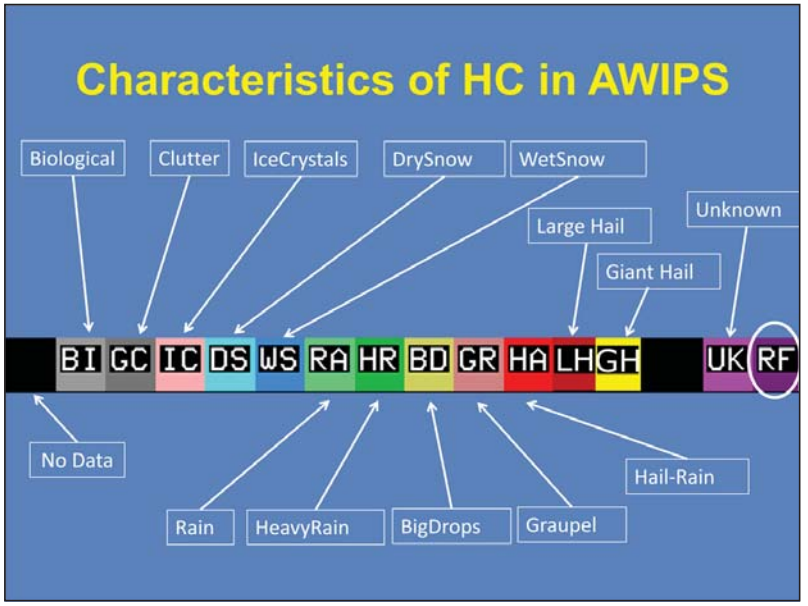
This Topic 4 lesson will be on Hydrometeor Classification. I'm Justin Gibbs of the Warning Decision Training Division.



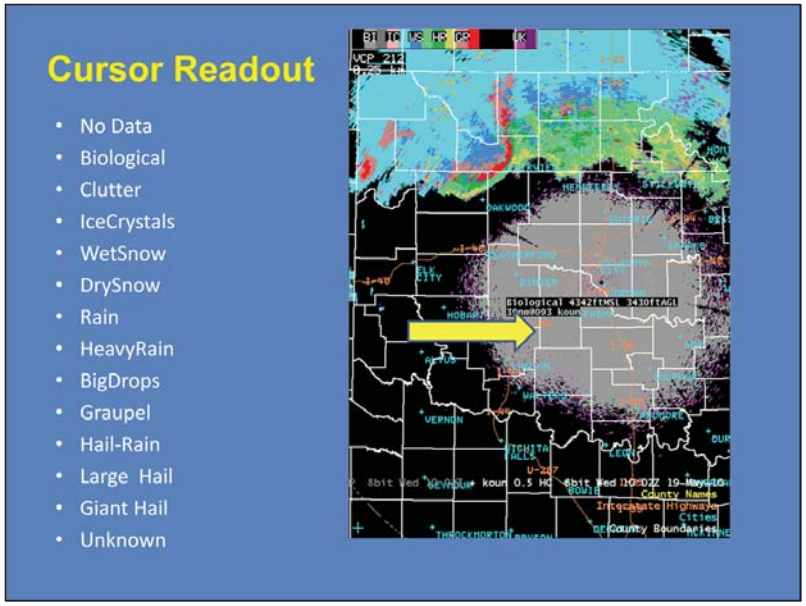
Here are the learning objectives for this lesson. Please advance the slide when you are ready to begin.



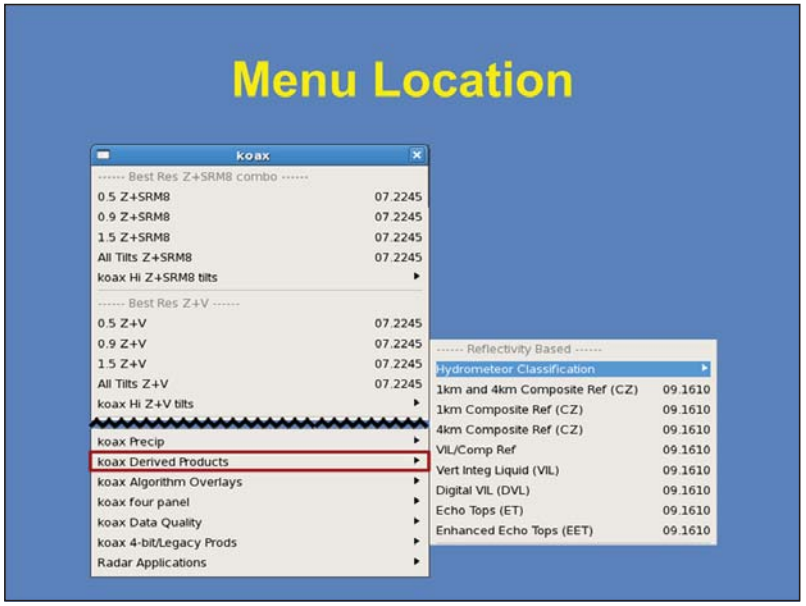
The purpose of the hydrometeor classification algorithm, or HCA, is to determine the most likely classification of all the range bins that the radar detects at all elevation angles. It attempts to put all the echoes into a predefined list of echo classes, and uses base data and polarimetric variables to reach its estimate.



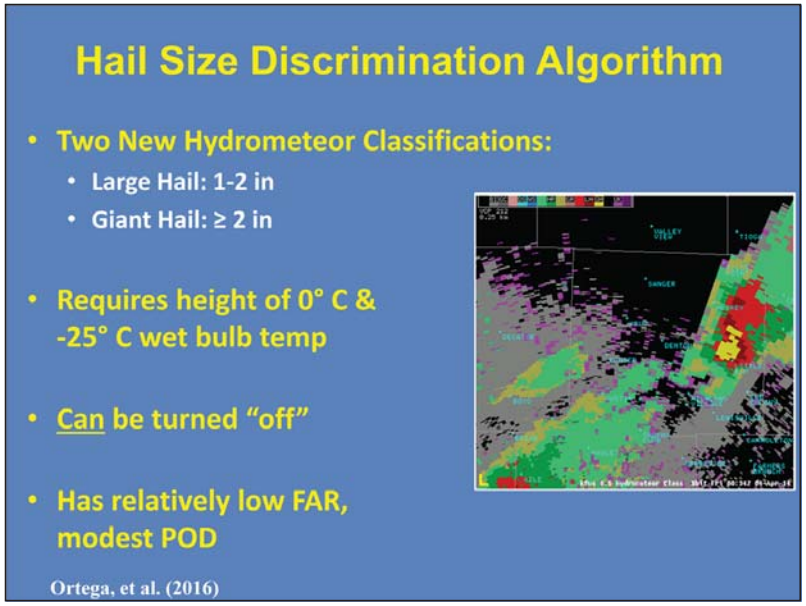
The bins that the HCA attempts to place echoes into are , No Data (ND), biological scatterers (BI), ground clutter / anomalous propagation (GC), ice crystals (IC), dry snow (DS), wet snow (WS), light/moderate rain (RA), heavy rain (HR), big drops (BD), graupel (GR) hail possibly mixed with rain (HA), Large Hail (LH), Giant Hail (GH) , and Unknown (UK). If new variables are developed they can be placed in the blank black area.



You can also enable sampling and do a mouse over of each target in plan view. In this case the target was biological scatterers at 4,342 MSL, and it will give you any of the returns we mentioned before.



To pull up the Hydrometeor classification algorithms output, go to your radars pulldown menu, select derived products, and Hydrometeor classification and you will get a pulldown menu that will allow you to select any individual, or all of the tilts of the HCA output.

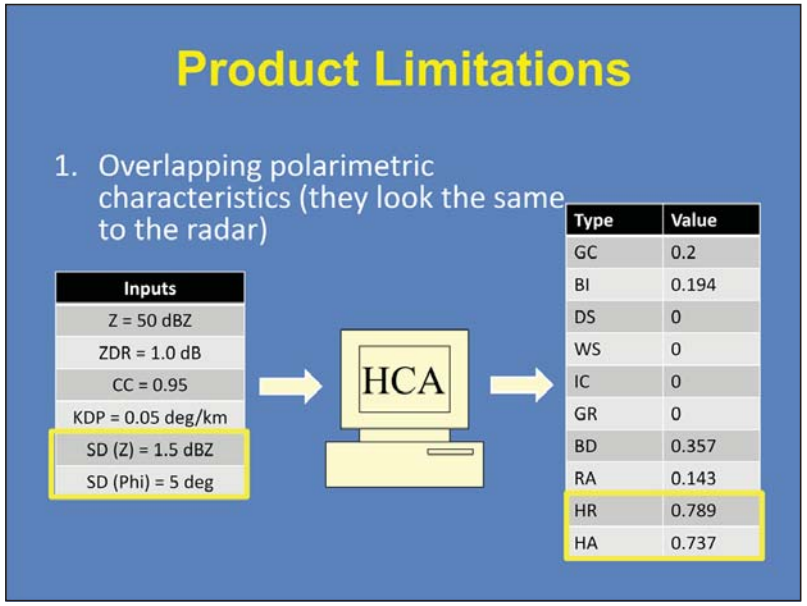


Brand new in 2016-17 is the hail size discrimination algorithm. This gives you your LH and GH outputs in the HCA. Which signals large hail, and giant hail. Large hail being severe hail between 1-2 inches, and giant hail, greater than 2 inches.

The HSDA requires a 0 and -25 degree C wet bulb temperature which should populate in the RPG.

It CAN be turned off if for some reason you need to, and then you would just get the standard rain hail mixed HCA output.

In validation testing it showed some promising skill, with a relatively low FAR, the probability of detection was a little low, so keep that in mind but it may have sufficient skill to be a useful alerting tool given its low FAR, meaning if its showing Giant hail and you don't think hail is falling, take another good look at your data because you may be missing something.



One of the key limitations of the hydrometeor classification algorithm are overlapping polarimeteric characteristics. Basically two different bins of possible answers that look the same to the radar. In this example the input has 50dbz a Zdr of 1db, CC of 0.95 an SD Z of 1.5 dBZ and an SD Phi of 5 degrees. When it runs through the algorithm the output says there's a 79% chance that its heavy rain and a 74% chance its large hail, so it can't tell if its hail mixed with rain, or just heavy rain.

Product Limitations

2. Uncertainty not portrayed

Example Inputs
- Z = 50 dBZ,
- ZDR = 1.0 dB,
- CC = 0.95,
- KDP = 0.05,
- SD (Z) = 1.5 dBZ,
- SD(Phi) = 5 deg

⇒ HR = 0.789
HA = 0.737

Heavy Rain (HR) will be displayed in AWIPS:
- Will not see that HA was very close in likelihood value

And its important to understand that uncertainty is not portrayed in the HCA output. It just takes the most likely answer and spits that out to you. So in our situation we had in the previous slide where its about 79% chance its heavy rain and 74% chance its hail, AWIPS is only going to display heavy rain. You wont see that hail was a very close in likelihood value, only about 5% difference. So when you have those overlapping characteristics where maybe its graupel or hail or rain, you wont know what those other options are other than what is displayed in the HCA.

Product Limitations

3. Available classifications are limited




No classification categories for scatterers such as ash or sleet

The available classifications are also limited, there's no classification categories for scatterers such as ash or sleet, or debris both from weather related phenomena or a rocket explosion, we've seen a couple of those in the last few years where debris shows up on the radar, its got low CC so humans can tell what it is, but the algorithm cant pick up on it.

Product Limitations

4. Thresholds are empirical and/or subjective

	Fuzzy-Logic Membership Function (Dry Snow)				Weighting (Dry Snow)
	X1	X2	X3	X4	Weight
Z	5	10	35	40	1.0
ZDR	-0.3	0.0	1.3	1.6	0.8
CC	0.95	0.98	1.0	1.01	0.6
KDP	0.001	0.003	10	100	1.0
SD(Z)	0	0.5	3.0	6.0	0.2
SD(Phi)	0	1	15	30	0.2

Values may not be representative for all regions of U.S.

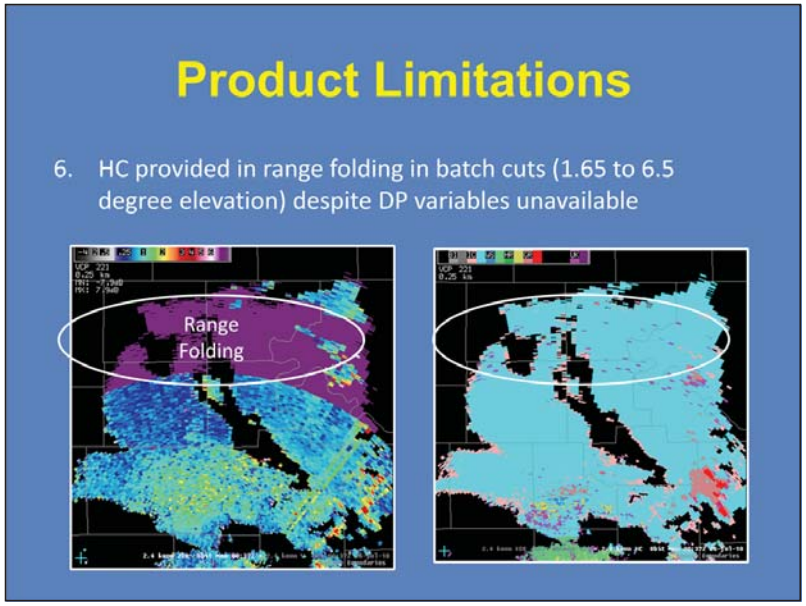
Another limitation is that the thresholds are empirical or subjective. What that means is a mathematical theory doesn't exist that says the Zdr of heavy rain should be say, 1.5 and the CC should be 0.95. There is no theory that we can go to say this is what we are expecting like there are with other variables in meteorology. So a fuzzy logic membership function was derived from a series of studies where they said "this is heavy rain falling, were going to point the radar at it" and these are the values that came out, so were going to say when we see these values in the future it must be heavy rain. The problem with that is that these values may not be representative for all regions of the US, so the algorithm will have that limitation as well.

Product Limitations

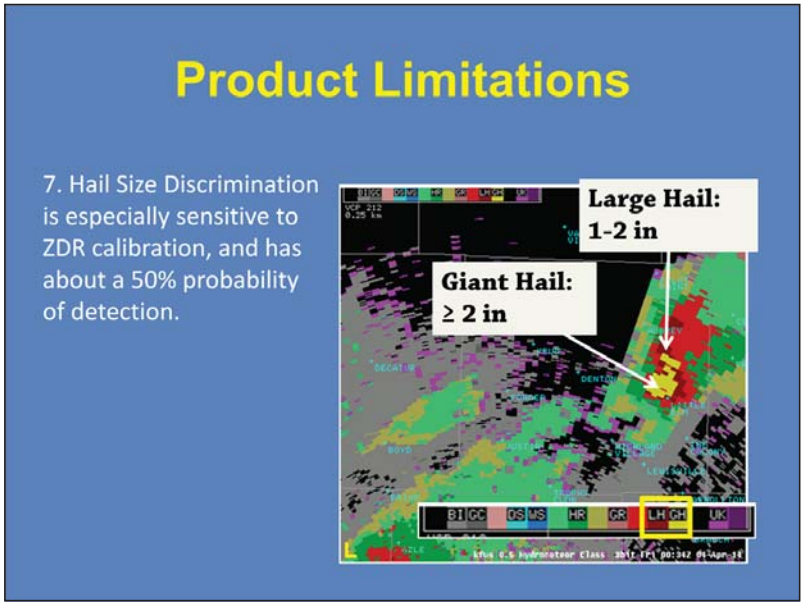
5. Be mindful of below beam effects!



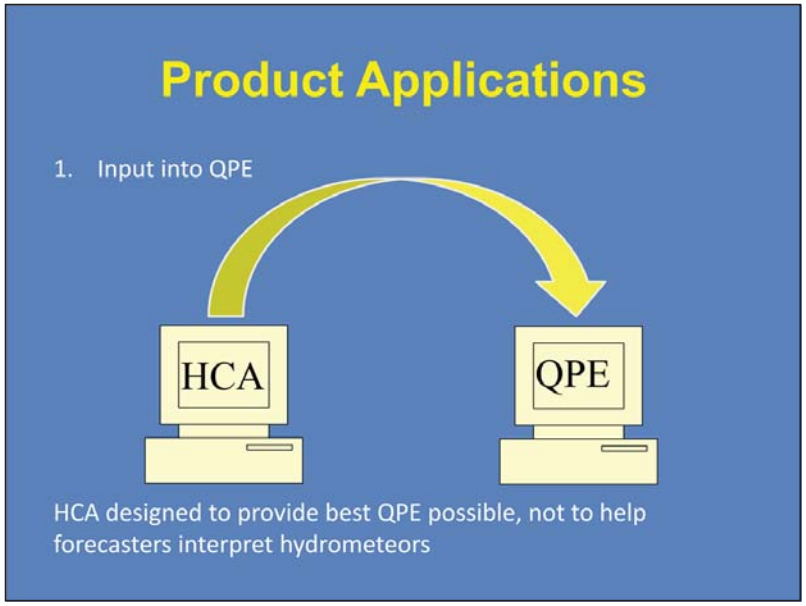
You also have to be mindful of below beam effects. The HCA is going to give you the output where the beam intersects the target, and whatever happens below the target the radar is not able to account for. So in this illustration the HCA is detecting hail and lets say 15,000 ft. Well that hail melts before it reaches the ground and produces rain at the ground. Its also a problem if its snow, heavy rain, hail you can imagine the possibilities where you could have a target at 8 or 10 thousand feet above the ground and how it could change before it reaches the ground.



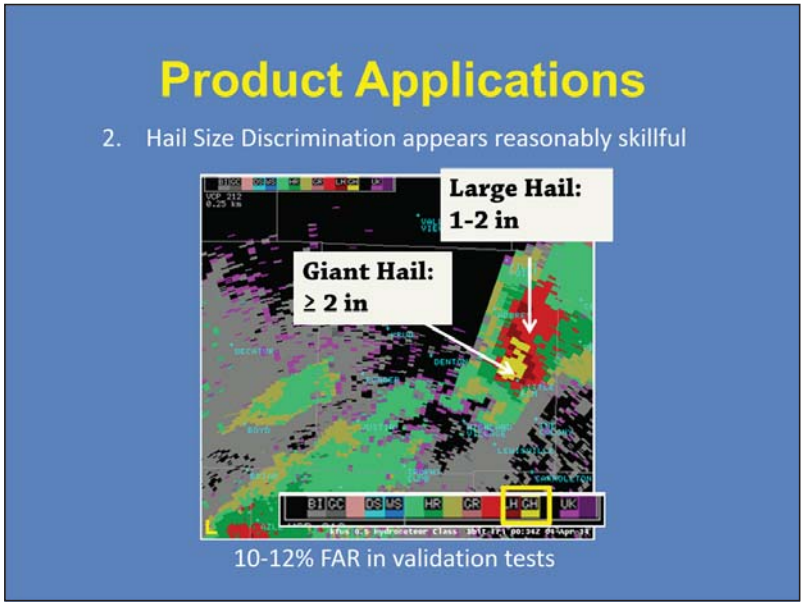
The sixth limitation, the HCA, part of what makes it work is the dual pol moments. The dual pol moments however aren't available in batch cuts, between 1.65 and 6.5 degrees elevation in range folding so this area of range folding that exists on the left hand image, there's no ZDR there's no KDP there's no CC but the HCA is going to make a guess anyway, and its not going to have nearly as much information as it would otherwise have to make a guess at what the p-type is and there's no indication in the HCA that, hey I don't have as much information as I normally have I don't have any ZDR or CC so the user has to be aware of that. Again that's where your meteorology comes in handy and your ability to know what is the most likely outcome based on the science that exists away from what the HCA is showing you.



The hail size discrimination algorithm is brand new, so we don't know too much about how it will perform in the field. It also has a relatively low POD, and is particularly sensitive to ZDR calibration issues.



One of the main applications of the HCA is that it is used to improve the QPE output. The HCA is really designed to provide the best QPE possible, and its not really there to help forecasters interpret hydrometeors. Its there to tell the radar what it is seeing, and its there to calculate rainfall rates and it just happens to be an output that's also available to you as you use it. So its going to provide you better rainfall rates, and if its not working correctly its going to introduce potential errors into your QPE and that's why you can see the output so you can be like "whoa, I've got a lot of hail I know there is so I'm going to be a little more questionable to my QPE for example" But it isn't there to help you interpret hydrometeors, that's why you're a meteorologist and you're a forecaster you make those determinations and provide that value added service to our customers.



The hail size discrimination algorithm shows promise, and appears reasonably skillful. It may be useful as a situational awareness too, and particularly as an alerting tool given its apparently low false alarm rate, meaning if its showing large or giant hail and you aren't anticipating large hail, to take a second look.

Product Applications

3. Provides quick look at regions of interest
   - Aviation nowcasting of hail





It also can be used to provide a quick look at regions of interest, perhaps aviation nowcasting of hail. You look at the image on the lower left and there's some 60 to 65 dBZ output and that correlates pretty well to where the algorithm is showing hail. The image on the right center looks a little more suspect but I don't know what's going on aloft in this image. So it gives you a quick first glance if you want to use it, but again base data analysis all the methods that you are being taught through this course and the courses to come are going to provide a superior service to the HCA.

Summary: HC

- Provides guess for echo type, HSDA appears skillful
- Input is used in QPE
- Thresholds/parameters are empirical and overlap
  - Uncertainty is not portrayed
- Below beam effects
- Not designed for deterministic operational use!

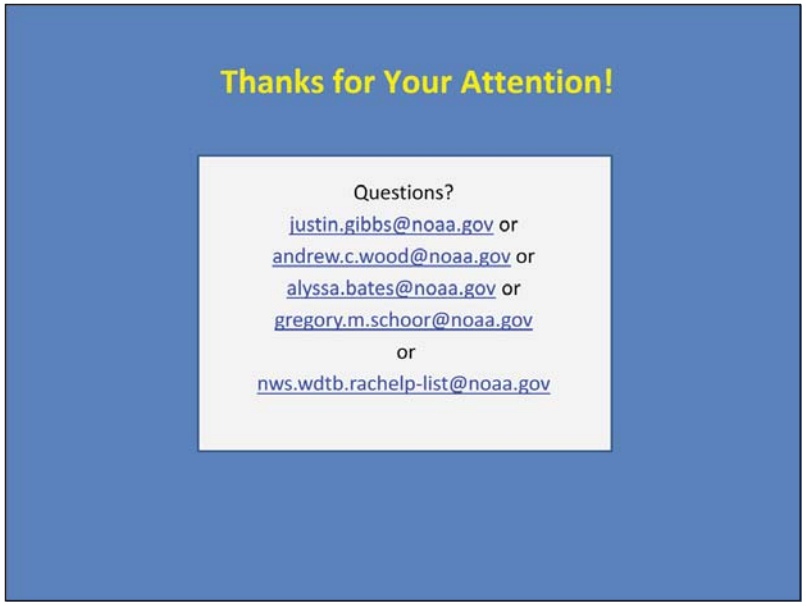
So in summary the HCA provides a guess for echo type, the new HSDA appears skillful but the algorithms

primarily used to improve QPE.

The thresholds are empirical they are not mathematically derived, and they can overlap and that uncertainty is not portrayed, there is a lot of fuzzy logic that goes into the developing the HCAs output.

Below beam effects what's going on at the ground may not be in any way representative as to what the radar beam is detecting.

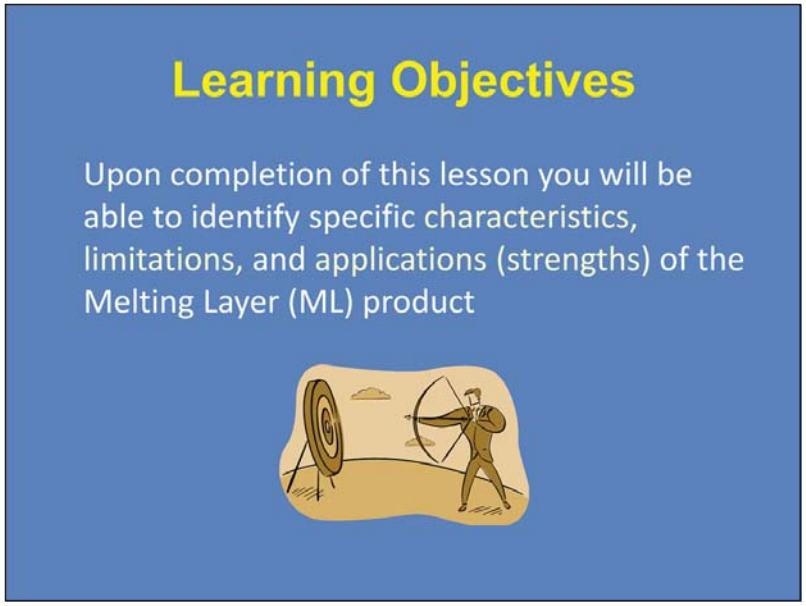
It is not designed for deterministic operational use and should not be used that way. If you're relying on your HCA output you are probably not going to be providing the best service, look at the other aspects of what the radar can do and meteorology to get that level of service up to the standards that we want it.



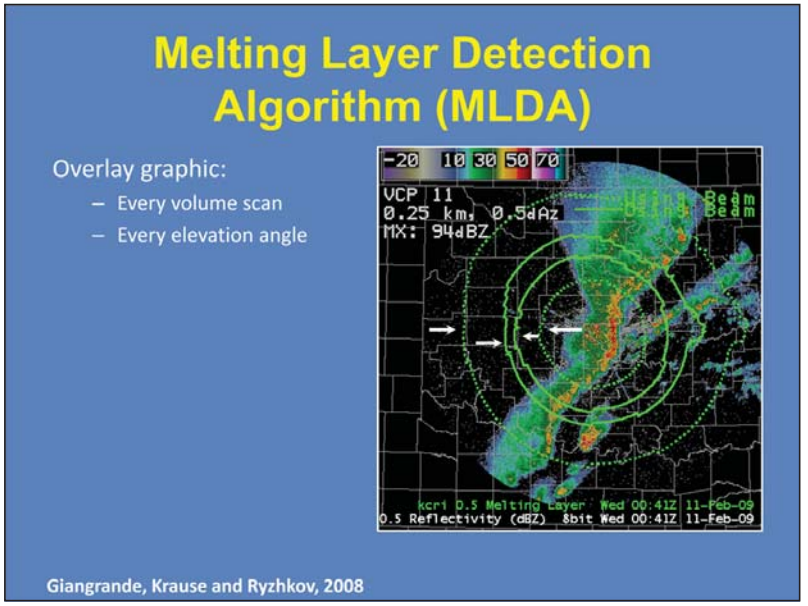
That concludes this lesson on hydrometeor classification If you have any questions email myself or any member of the team responsible for putting this course together. Thanks for listening!



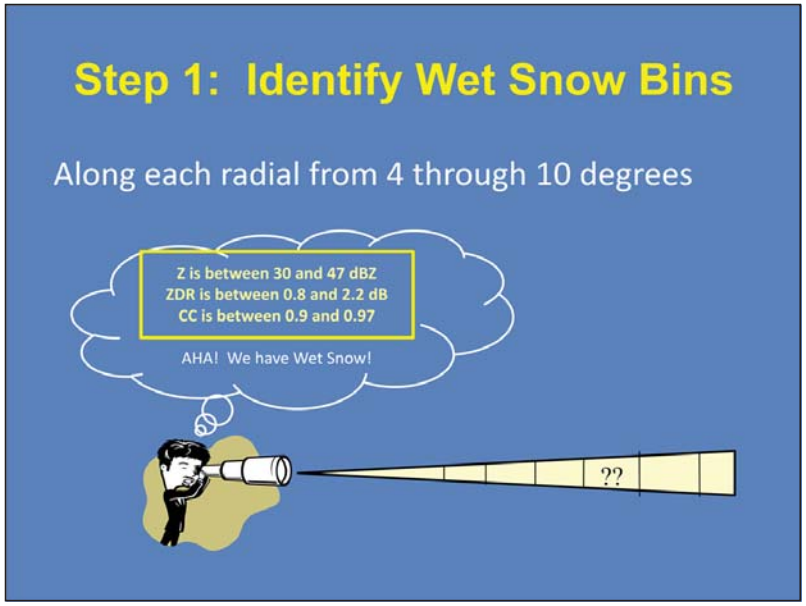
Welcome to this Topic 4 lesson on the Melting Layer.



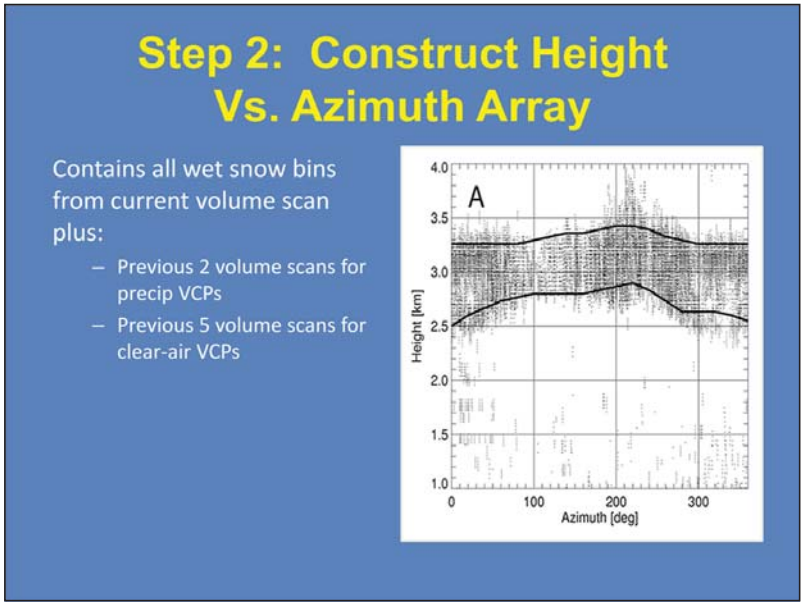
Here are the learning objectives for this lesson. Advance to the next slide when you are ready to begin.



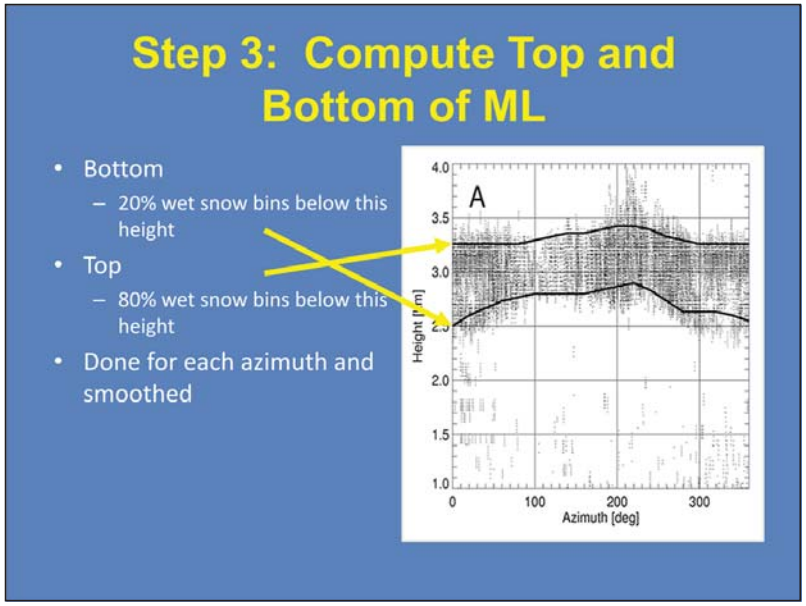
Recall that the ZDR and CC products are great at identifying the melting layer. In other words, the melting layer has a distinct signature in both ZDR and CC. This characteristic is used by the Melting Layer Detection Algorithm (MLDA) to automatically detect a melting layer from the radar data and then display it as an overlay on other radar products. The green solid and dotted lines in the graphic shown represent the melting layer product. This graphic is available every volume scan for every elevation angle. The next few slides will briefly describe how the MLDA works.



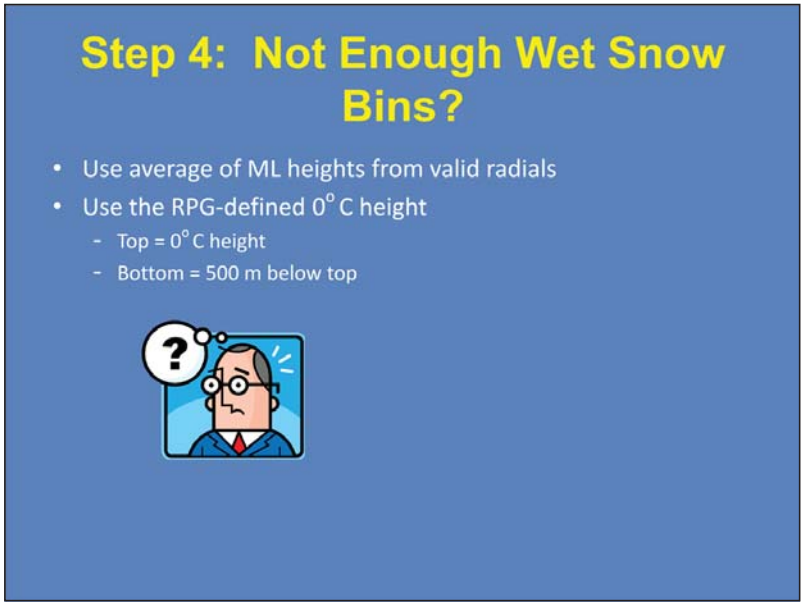
The algorithm starts by looking down each radial between the elevations of 4 and 10 degrees and identifies the total number of wet snow bins. A wet snow bin is identified in regions where there is a combination of high Z and ZDR and lower CC. For a bin to be “wet snow”, Z must be between 30 and 47 dBZ, ZDR between 0.8 and 2.2 dB, and CC between 0.9 and 0.97. It's important to note that “wet snow” in the MLDA is not the same as wet snow in the HC product.



The next step in the algorithm is to construct a height versus azimuth array (like the one shown here on the right) for all bins identified as wet snow from step 1. This array includes all wet snow bins from the current volume scan, plus the previous two volume scans in precip mode or the previous five volume scans for clear-air mode.

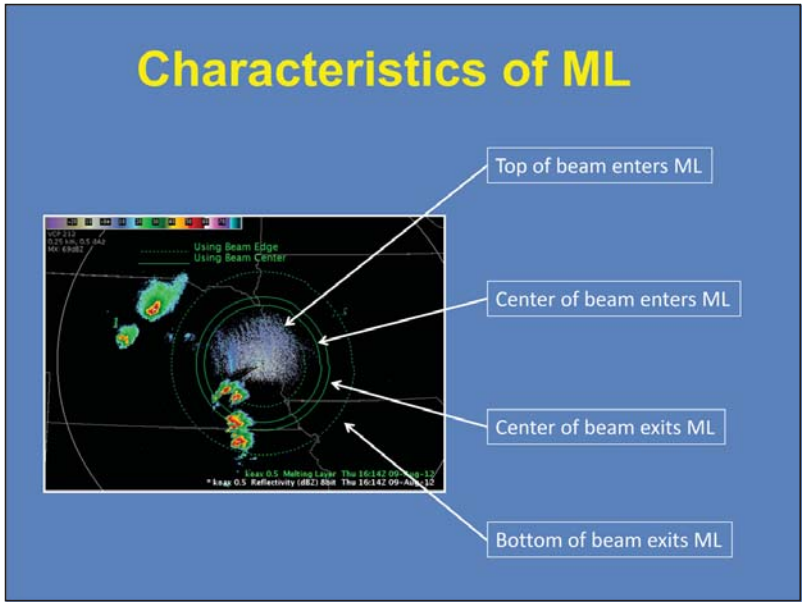


If a sufficient number of wet snow bins are identified, the top and the bottom of the melting layer are then computed using the height versus azimuth array in step 2. The top and bottom are determined using a percentage of wet snow bins along each azimuth of the array. The bottom of the melting layer is placed at the 20th percentile location (bottom black line) while the top of the melting layer is identified at the 80th percentile location (top black line). Once these height designations are identified for each azimuth, they are smoothed to maintain time and space consistency between radials.



So what happens if the MLDA doesn't identify enough wet snow bins to locate the melting layer? Well, the algorithm still identifies a melting layer height using one of two methods.

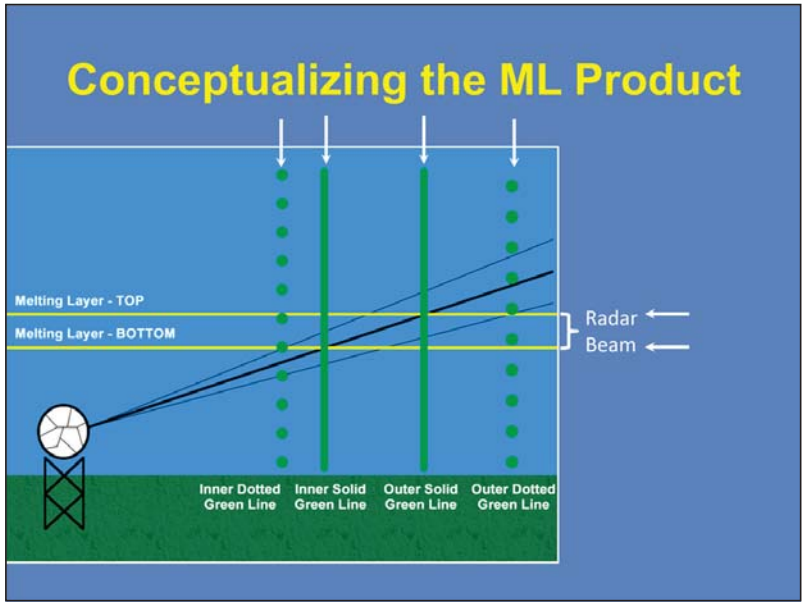
The first method interpolates the melting layer top and bottom from nearby radials in the current volume scan where there were valid detections. If the gap between radials where the melting layer was identified is too large, then the second method uses the 0 Celsius height in the gap as defined in the RPG. It will set the 0 Celsius height as the top of the melting layer, and the bottom of the melting layer will be 500-m below that height. This 0 Celsius height can be defined automatically if you have “Model Input” turned on at the RPG, or you can manually enter it if you think the model is out to lunch. Either way, just make sure you have an accurate 0 Celsius height defined in the RPG.



The heights determined by the MLDA are then used to construct the Melting Layer (ML) product. This example shows the 0.5 deg ML overlaid on top of the 0.5 deg Reflectivity. I should note that, although the algorithm only uses tilts between 4 and 10 degrees to identify the melting layer, the ML product is available for every elevation angle in the volume scan.

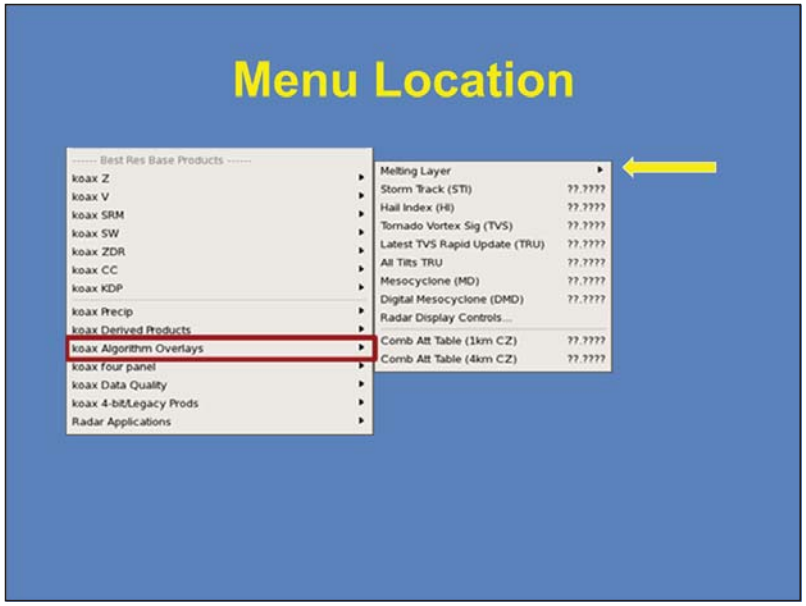
So let's look at this display and see what the ML product tells us. There are four different green lines displayed: Two are solid lines and two are dashed. We'll start closest to the radar and work our way down the radial. The innermost dotted line indicates where the top of the radar beam intersects the bottom of the melting layer. The innermost solid line indicates where the radar beam center intersects the bottom of the melting layer. The outermost solid line indicates where the center of the beam intersects the top of the melting layer. Lastly, the outermost dotted line indicates where the bottom of the beam intersects the top of the melting layer.

So, the region between the two solid lines is where the radar primarily samples the melting layer. However, the melting layer affects in some way all data collected between the two dashed lines.



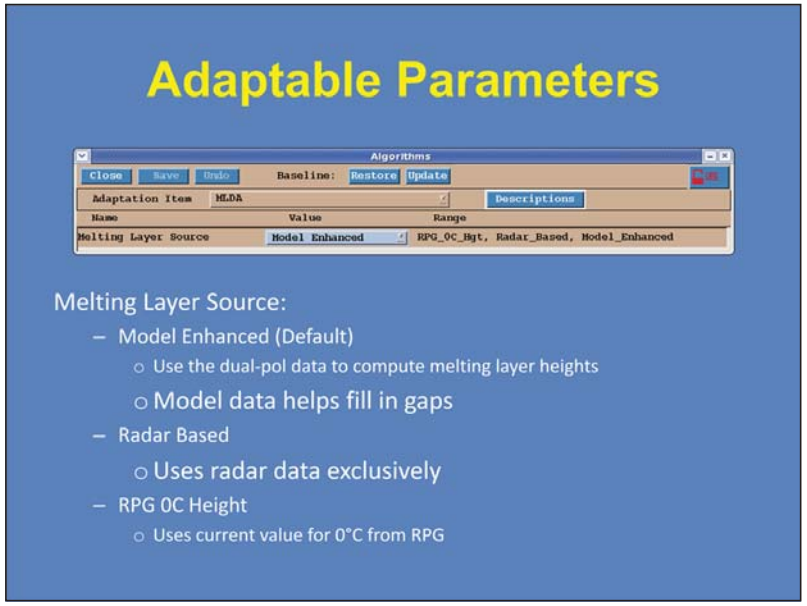
We provide this graphic to conceptualize the ML display further and aid your understanding further. In this graphic, the radar beam is illustrated with a thick black line for the center of the beam and two thin black lines for the radar beam edges. The two yellow lines represent the identified melting layer top and bottom.

Following the discussion from the previous slide, the first dashed line is displayed where the top of the beam intersects the bottom of the melting layer. The innermost solid line is plotted at the height where the center of the beam intersects the bottom of the melting layer. The outermost solid line represents where the beam center intersects the top of the melting layer. Finally, the outermost dashed green line is plotted at the height where the bottom of the beam intersects the top of the melting layer.



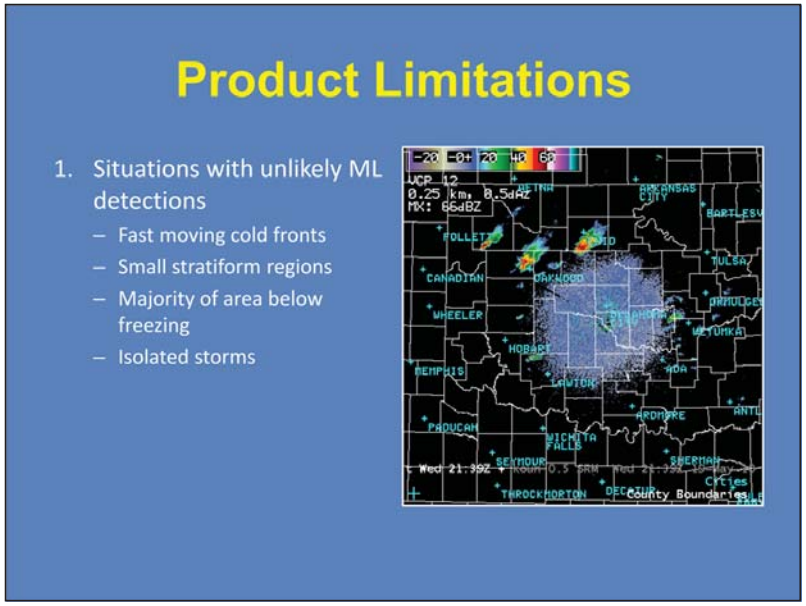
The ML product graphic can be found in two places in the dedicated radar's drop-down menu. You can load the ML product by itself for any elevation angle available using the Algorithm Overlays menu in your dedicated radar's menu. You can also load this product in conjunction with other dual-pol products in the four panel menu (not shown).

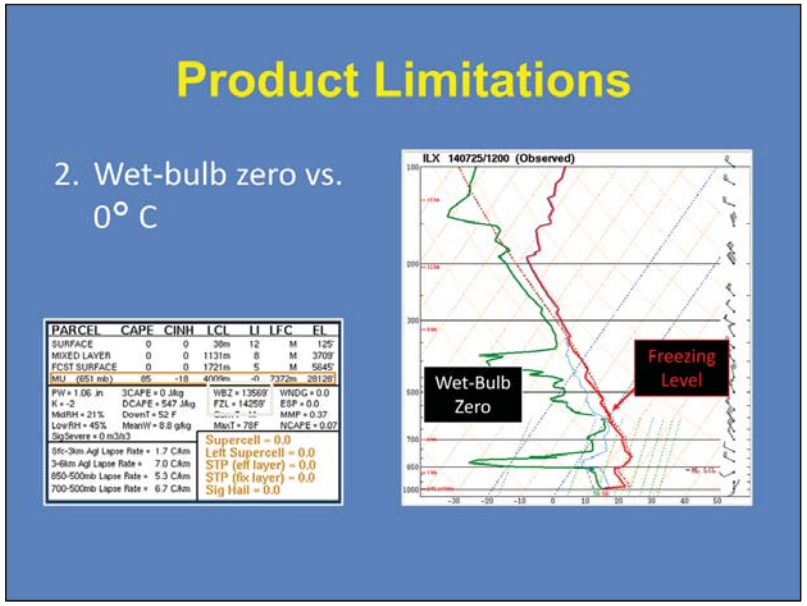
We recommend displaying the ML product only on other elevation-based radar products. And when you do, make sure the elevation angles for the ML and other products you have loaded match.



In this RPG window, you will see one adaptable parameter, entitled Melting Layer Source, with three possible options. "Model Enhanced" is the default option shown on the slide. When "Model Enhanced" is selected, the MLDA will look at the data available radial by radial. When sufficient data exists, the methodology discussed previously is used. Azimuths lacking sufficient coverage are addressed in two ways. If the data gap is less than 15 deg of azimuth, the RPG interpolates to fill the gap using radar-based heights only. If the data gap is greater than 15 deg of azimuth, the RPG interpolates using a blend of the radar-based heights and model data. If no radar-based estimate is possible at any azimuth, model data are used at all azimuths.

The other two options are "Radar Based" and "RPG 0C Height". When "Radar Based" is selected, the algorithm will use radar data exclusively. When "RPG 0C Height" is selected, a 500 m thick melting layer is assigned at all azimuths using the current value for 0 C in the RPG Environmental Data window.

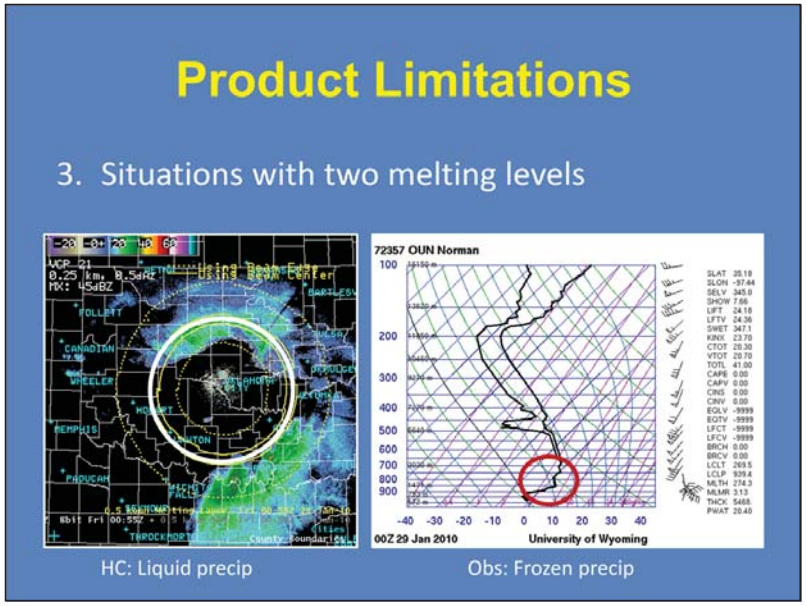




Let's talk about the limitations of the ML product. The algorithm requires the radar to detect the melting layer with a certain amount of spatial and temporal continuity. So, if an insufficient amount of wet snow detections are made in a 3-6 volume scan period, depending on VCP, then the algorithm will fail. So, when will failures most likely occur during non-quiescent weather?

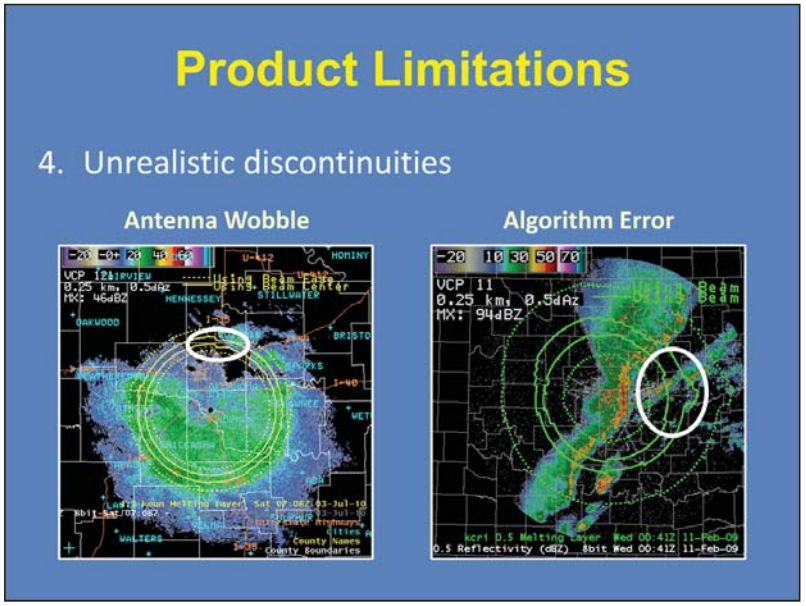
The first situation is when there are fast moving cold fronts. During quick frontal passages, the algorithm can have difficulty maintaining temporal continuity unless there is a lot of precipitation present. A second difficult situation occurs during small-scale stratiform rain events where there is not enough wet snow sampled to identify the melting layer heights. Problems can also occur when the majority of the local area is below freezing, especially if the 0 degree height listed in the RPG is incorrect. Lastly, issues can arise with the ML product when there are only isolated storms in the area, as in the example shown.

The top of the melting layer is considered synonymous to the 0º C height. While this is true in some cases, the top of the melting layer is actually the wet-bulb zero height. In some environments, there may be a significant difference (~1000 ft or more) between the two heights. If the algorithm has to switch between algorithm output and the environmental data values, noticeable differences may occur due to the difference in what is measured, not because the algorithm isn't working properly.



In the event there are multiple 0 degree Celsius levels, the algorithm will only detect one of them. An example of such a situation includes sleet events. You may see both melting layers in the base data itself under the right circumstances. Here is an example from 29 January 2010 around 1 Z. On the left is the Reflectivity with ML overlaid. On the right is the 00 Z sounding from OUN. Note that the sounding shows two melting layers present, but the ML product only indicates one.

One implication of this limitation is the HC product will indicate liquid precipitation in areas where frozen precip is observed at the ground. Therefore, you should rely on any surface-based spotter reports and sounding data that are available during these types of events.



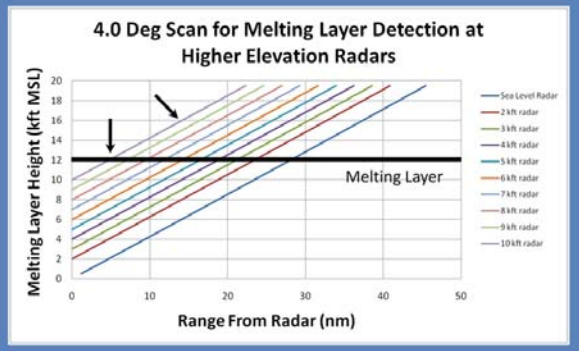
Another limitation to be aware of occurs when there are unrealistic discontinuities in the melting layer display. When this situation occurs, there are two likely causes: antenna wobble and algorithm error.

Antenna wobble typically occurs at the beginning of an elevation scan. When the antenna changes elevation angles, sometimes the dish has not finished moving vertically before data collection begins. As a result, you may see an odd, smooth discontinuity (such as in the example on the left) because the beam is under- or over-shooting at the beginning of the elevation scan. These discontinuities should only appear on elevation angles where the wobble occurs.

While antenna wobble discontinuities appear smooth, algorithm errors usually appear spiky and are visible on every elevation angle at the same azimuth. An example of algorithm error is shown on the right.

Product Limitations

5. Performance in mountain regions

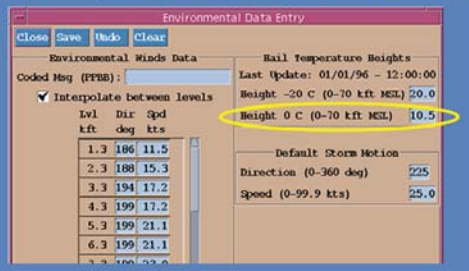


Recall that the MLDA computes the ML product using data between 4 and 10 degrees. At lower elevation regions, using these scans will rarely be a problem. However, radars in mountainous regions are another story. The higher the surface elevation of the radar, the greater the likelihood the radar beam will start in the melting layer or completely overshoot it. The graph shown on the slide illustrates this problem.

Say, for example, your radar is located at 10 kft above MSL. The purple line at the top of the graph represents the center beam height for your radar in a standard atmosphere. If the top of the melting layer is 12 kft MSL, then the center of the beam will already be sampling above the melting layer in the first 10 nm. Unless there is precipitation all around the RDA, it's unlikely the MLDA algorithm will work well in these conditions.

Product Applications

1. Updates every volume scan
   - Polarimetric data is not sufficient?
   - Uses 0° C height from RPG adaptation data -or- RAP



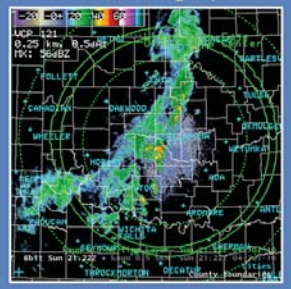
Enough with the limitations…let's move on to the ML applications.

With the ML product, you have the potential for measured melting layer heights updated every volume scan. Even if the radar cannot sample enough "wet snow" to compute the melting layer height, the algorithm can still provide melting layer heights (using the RPG 0 Celsius height from the RPG adaptation data) as an overlay for base data products.

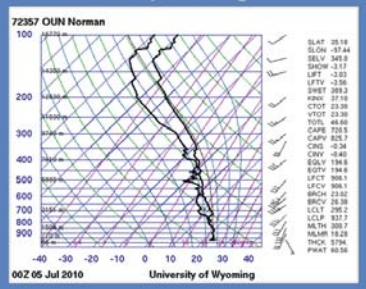
Product Applications

2. Aids in interpreting base data

0.5° Z w/Melting Layer



Nearby Sounding



ML = 15,500 ft MSL / Skew-T = 16,200 ft MSL

Summary: ML

- Provides ML height for each radial determined from base data
  - Updates every volume scan!
  - Uses model data if not enough data
- Many situations with unlikely ML detections
- Two melting level events are tricky
- Unrealistic discontinuities
- Poor performance during:
  - Fast moving cold fronts
  - Poor echo coverage
  - Mostly below freezing temperatures

By overlaying the ML product on base data products, you can more quickly tell where echoes are being sampled thermodynamically in the atmosphere, which can help you more quickly analyze precipitation type. This application can be very helpful because the melting level is not always uniform nor is the nearest sounding always representative of the area you are investigating.

Here is an example from 4 July 2010 in Central Oklahoma. On the left, I have overlaid the ML product on the 0.5 degree Z. With this product, now I can tell more quickly where the echoes are relative to the melting layer. Without the overlay, I would need to consult a sounding (like the one shown to the right) and then use cursor readout to determine the beam center height at a particular range.

Having said that, don't ignore sounding data. Use soundings for confirmation if the ML product look fishy. In this example, the algorithm appears to be working OK as the top of the radar-identified melting layer is around 15,500 ft MSL and the skew-T shows the 0 degree height around 16,200 ft MSL.

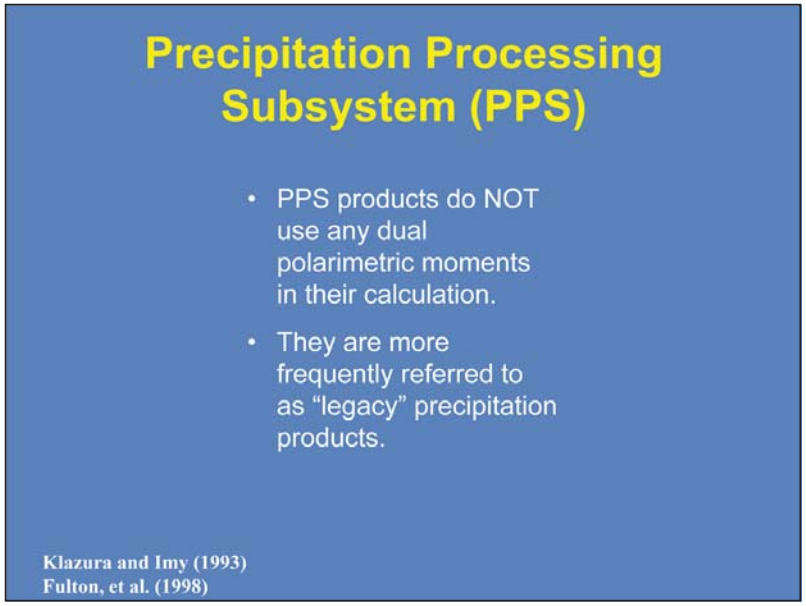
In summary, the ML product provides an updated melting layer height for every elevation each volume scan based on the base data and/or environmental input at the RPG! There will be many situations where the algorithm cannot identify the melting layer based on radar data alone. The algorithm can only identify one melting layer, so the algorithm doesn't handle refreezing events well. Also, watch out for unrealistic discontinuities due to antenna wobble or algorithm errors. Lastly, be cautious when using the melting layer product with fast-moving cold fronts, precipitation is occurring in only a small area, or below freezing temperatures are occurring throughout the air mass over most of your area.



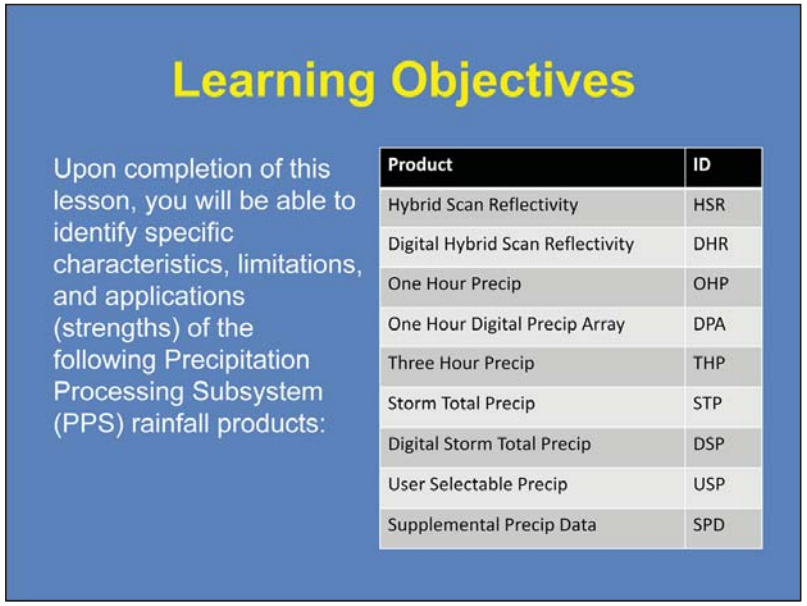
If you have any questions, feel free to contact us using any of the e-mails listed below.



Welcome to this Topic 4 lesson on PPS rainfall products.



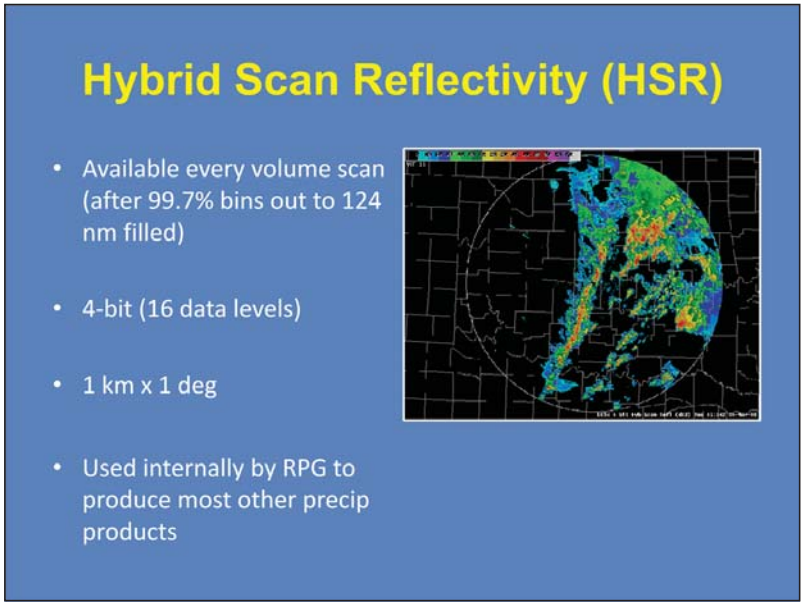
PPS stands for the precipitation processing subsystem. PPS products are the original estimation of precipitation algorithms that came with the WSR-88Dand we're last updated in 2004 . These predate dual pol and as a result do not have any dual polarimetric moments included in the calculations. Operationally, you most frequently hear these referred to as the legacy precipitation products. At your office when folks are referring to legacy products, legacy precip, 1 hour precip, that is the PPS products. We kept them because they're still somewhat skillful.



Here are the learning objectives for this lesson, and the chart on the right shows all the PPS products we'll discuss in this lesson. When you are ready, click next to continue.



Products are all located under your precipitation menu on your radar under the legacy precip accumulation products



First, we will discuss a hybrid scan reflectivity or HSR. That is available every volume scan after 99.7% of the beams are filled out to 124 nautical miles. Its called hybrid scan because it combines data from multiple elevation angles to produce the best reflectivity to estimate ground rainfall accumulations. It's trying to account for ground clutter, beam blockage and the data quality issues with this product. It is a 4-bit products with 16 different data levels, and it's available in one kilometer by 1 degree resolution. Its primarily used internally by the RPG to produce most of the other precip products.

Digital Hybrid Scan Reflectivity (DHR)

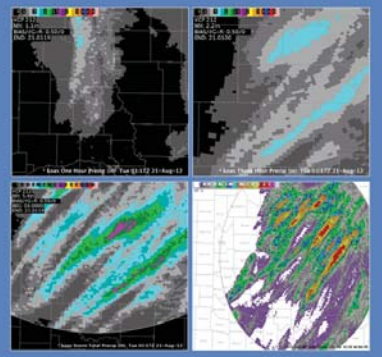
- 8-bit (256 data levels)
- 1 deg x 1 km
- 124 nm range
- Used primarily for FFMP and HPE
- Useful when working hydro events



Next the digital hybrid scan reflectivity or DHR. that's an 8-bit/256 data level product in 1 degree by 1 kilometer resolution and 124 nautical mile range. Its used primarily for input into FFMP, but I've also find it's got some operational use when you're working hydro, to determine intensity and location of rainfall. It is also a decent option when you're making Facebook or Twitter graphics or other public facing graphics as it gives a better idea for rainfall than just a 0.5 degree reflectivity image but its not as course it's a composite reflectivity which is in a much lower resolution.

Common Characteristics of All PPS Rainfall Accumulation Products

- Built from HSR
- 2 km (1.1 nm) x 1 deg
- 4-bit (one exception)
- 124 nm range



Characteristics of all the legacy rainfall accumulation products. They are built from hybrid scan reflectivity. They are in two kilometer by 1 degree resolution. They are 4-bit products with one exception, and have a range of 124 nautical miles. So if you get targets further away than 124 nautical miles you won't have any PPS rainfall accumulation products

One Hour Precipitation (OHP)

- Accumulations for past hour
- Use scan-to-scan accumulations
- Moving one hour window
- Available from the first volume scan with detected rainfall



The one hour precipitation product is one you will probably be accessing most frequently. It produces accumulations from the past hour although it is a one hour moving window so if the radar goes down or there's missing data, it could be more like a 1 hour 10 minute window. Normally though it will be be one hour. It uses scan to scan accumulations and is available from the first volume scan precip is detected.

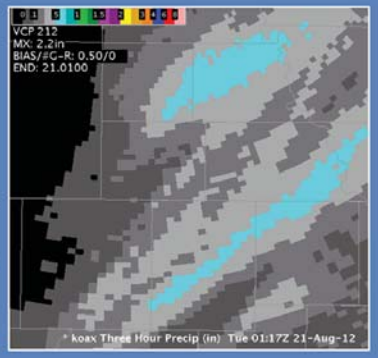
Digital Precipitation Array (DPA)

- Used by exterior systems
  - RFCs for NWSRFS
  - AWIPS Multisensor Precipitation Estimator (MPE) for bias calculation
  - FFMP - Note: DHR is preferable
- Moving one hour accumulation
  - Scan-to-scan accumulations
- 2.2nm x 2.2nm gridded product (easier to mosaic several radars)
- 8-bit (256 data levels)
- Not displayable

The digital precipitation array or DPA is mainly going to be used by external systems, like River forecast systems, the AWIPS multi-sensor precipitation estimator or MPE which is what goes into FFMP for bias calculation . It's a moving one hour accumulations calculator. DPA is a scan to scan accumulation similar to the one hour precip and is in 2.2 nautical mile x 2.2 nautical mile gridded product which is easier to mosaic several different radars, similar to concept that we see with the multi radar multi sensor project except at a much coarser resolution. Its an 8-bit 256 data level product and it is not displayed at your AWIPS workstations, it is sort of a behind-the-scenes product.

Three Hour Precipitation (THP)

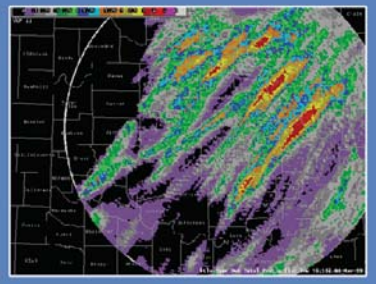
- Updated once per hour (after top of the hour)
- Uses top-of-the-hour accumulations
- Require 2 out of 3 top-of-the-hour accumulations (zero or nonzero)
- Not recommended for RPS list



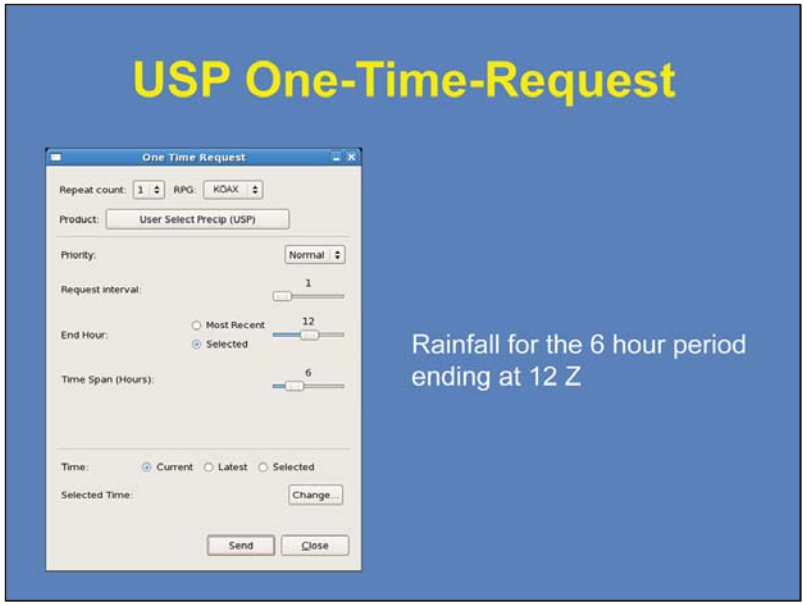
The three hour precipitation product or THP is updated once per hour at the top of the hour. It uses top of the hour accumulations and requires two out of the three top of the hour accumulations, zero or non-zero. Its not recommended for your RPS list because it only comes in at the top of the hour so it isn't as useful operationally as some of the other products.

User Selectable Precipitation (USP)

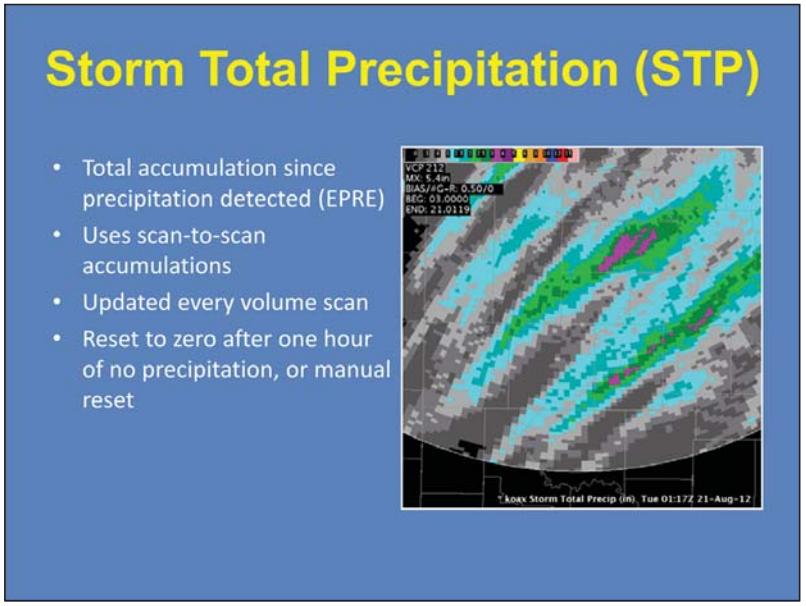
- Accumulations for user specified time
- Top-of-the-hour accumulations
- Past 30 hours are available
- User selects duration (up to 24 hours) & end time
- Default USP:
  - 24 hours ending 12Z



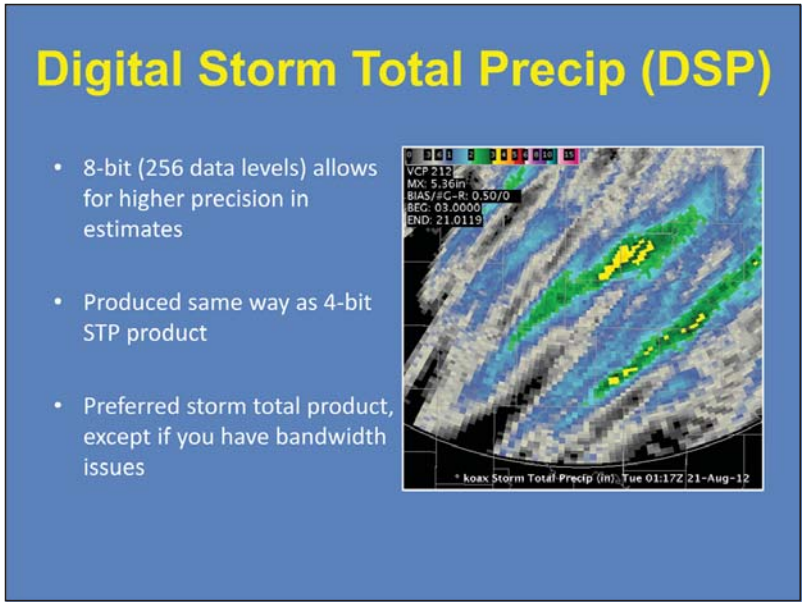
If you don't want a one or three hour term. You can use the user selectable precipitation product or USP. This gives accumulations for user-specified time and uses top of the hour accumulations. Up to the past 30 hours of accumulations are available. The user selects the duration this can be up to 24 hours then after 24 hours and choose the time you want the 24 hours ending at. The default is 12z.



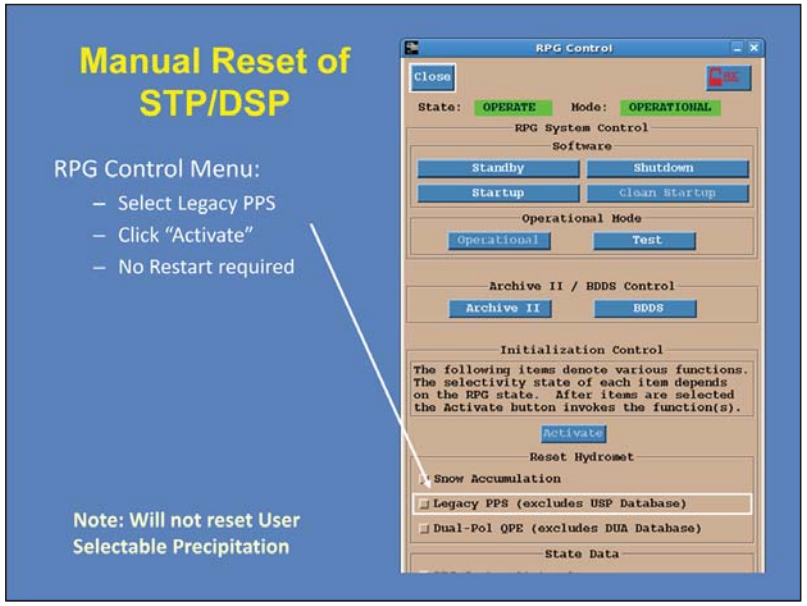
If you're interested in doing a USP one time request you can go to the one time request option on your radars menu, then select the end hour, you can choose the most recent or whichever hour that you would like. You can select a time span going back up to 24 hours. In this case you would get an accumulated rainfall for six hour period ending at 12Z.



Next product is the storm total precipitation product or STP. This is the total accumulation since precipitation was detected using the enhanced preprocessing algorithm which was last updated in 2004. It uses scan to scan accumulations and it's updated every volume scan. it's reset to 0 after 1 hour of no precipitation or can be manually reset which some offices do as part of their daily checklist.



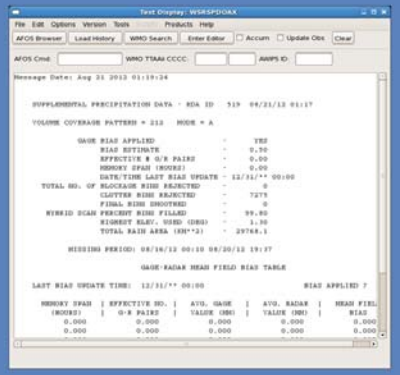
This is the one exception that I mentioned earlier with respect to 4-bit being the standard resolution. I know there was DPA, but it was not a displayable product. Anyhow, this product here is the high resolution storm total precipitation product, or DSP. It is 8-bit allowing for higher precision in estimates and is produced in the same way the 4-bit STP product is generated. By far, this is the preferred storm total product to request, provided you do not have bandwidth issues.



Manual reset of your STP and DSP can be conducted through your RPG control window on your HCI. It will not reset user selectable precipitations if you have a one-time request there that will continue to keep going. Just go to your RPG control, enter your URC password and then select legacy PPS and click Activate. A restart of your RPG is not required. This may be in in many cases part of your daily radar our check list. So you may have done this before this course but that is the way that you do it in case you're unfamiliar.

Supplemental Precipitation Data (SPD)

- Alphanumeric product
- Available every volume scan
- Bias computation
- PPS output
  - Gage-radar pairs
  - Missing data



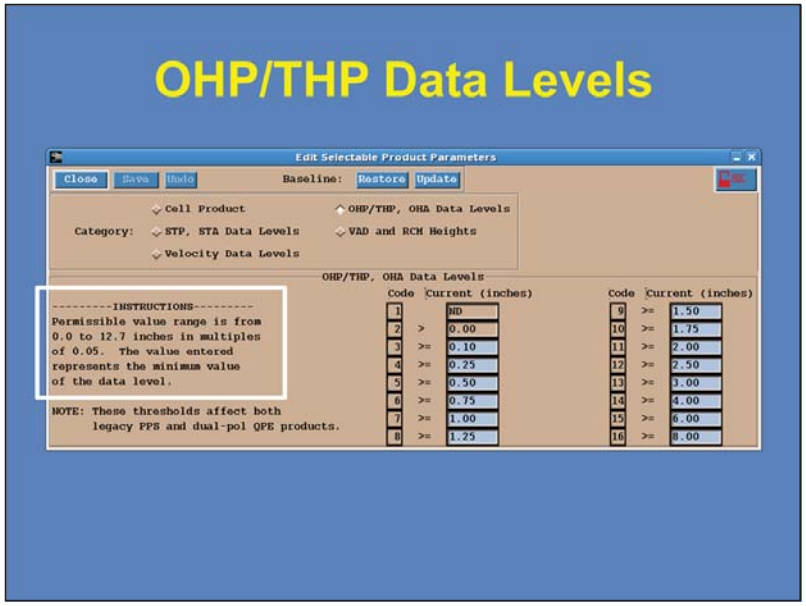
The final product we will talk about is the supplemental precipitation data product or SPD. Its an alpha numeric product and you can pull up the text workstations. Its available every volume scan and it contains a bias computation that was chosen for the other precipitation accumulation products. That also would typically show up on the product in your left hand corner, but this is another place you can see it. You should have access to that product again from any window on your text workstation and it will give you information about your gauge radar pairs and any missing data.

Adjusting Data Levels

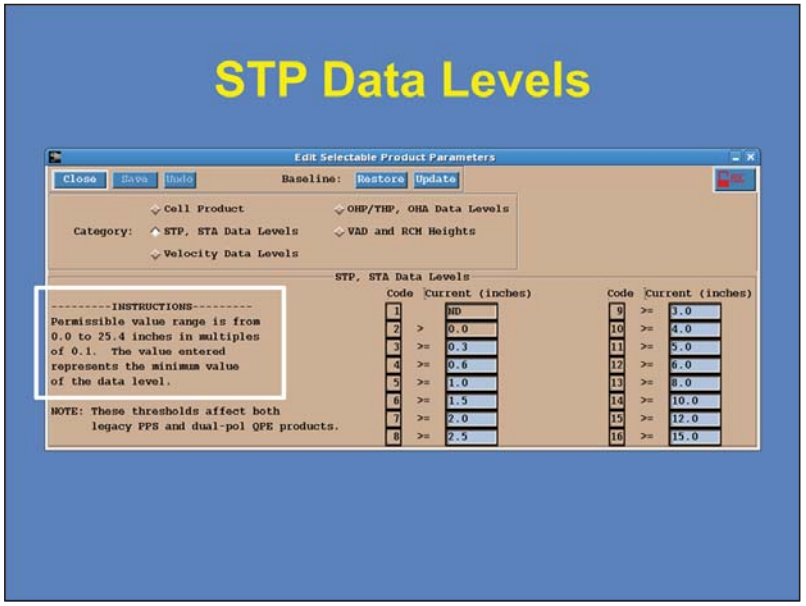
- Data levels on OHP, THP, and STP are editable (URC LOCA)
- Change for:
  - Climate
  - Topography
  - Expected weather



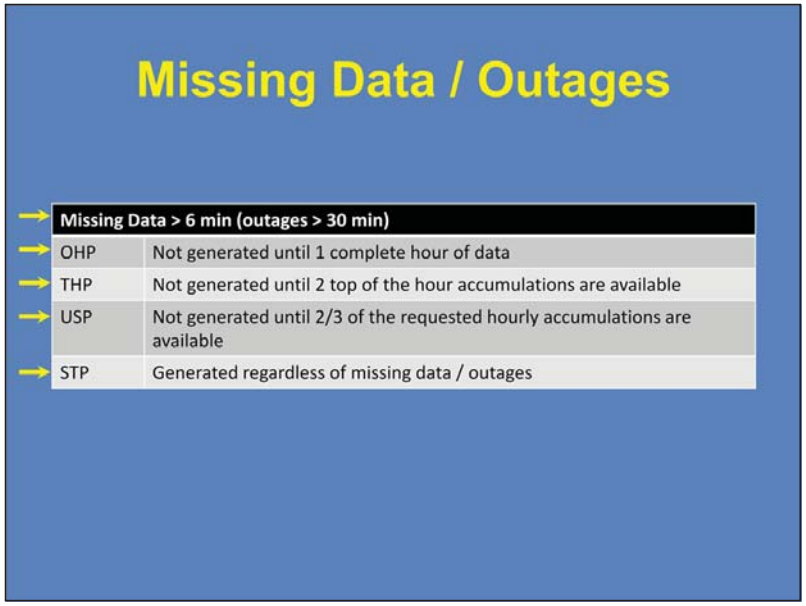
An important consideration of the precipitation accumulation product is at the data levels, like 1 inch 2 inch 3 inch half inch are editable, on the one hour, 3 hour, and storm total precipitation products. You need the URC password to go into the products in database display under your RPG products window. This might be useful depending on your climate where forecasting in Miami is going to be different than Bozeman far as what type of rainfall that you might be expecting. Your topography where say in Death Valley a quarter of an inch of rain can cause flash flooding versus New York City maybe an inch or two of rain but in the open rolling hills of Nebraska maybe a different level of rainfall and that way you can make calls on if you want to see a different level to account for those higher or lower values.



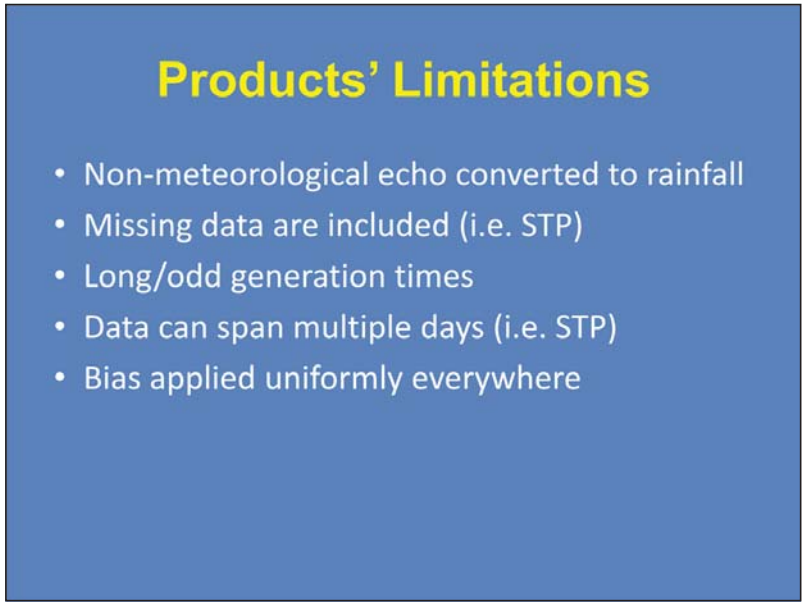
Here's what the edit window looks like in the RPG. Notice that you can go from 0 to 12.7 inches in multiples of 0.05 inches in the OHP in and THP products.



The storm total precipitation products are editable up to 25 point 4 inches in multiples of 0.1 But you notice in this case if you happen to have a 15 or 16 or 17 inch rainfall there wouldn't be any way to discriminate that so if you're expecting a tropical cyclone maybe you give up the first couple degrees of levels in order to discriminate a higher level rainfall.



What constitutes missing data in the accumulation products? Missing data implies greater than 6 minutes where an outage implies greater than 30 minutes of missing data. The 1 hour precip product is not generated until 1 complete hour is available. THP is not available until two top of the hour accumulations are available. The User-selectable precipitation product is not generated until at least two-thirds of the requested hourly accumulations are available. And storm total precip is generated regardless of missing data or outages, it just keeps trucking along.



Limitations of all of the precipitation products include non-meteorological echoes in the legacy products can be converted to rainfall. If you have a lot of biological scatterers, or wind farms or bird migrations or ground clutter that can be converted to rainfall. Missing data are included in the storm total accumulation product. It can have long or odd generation times if you're having trouble with your radar getting information. If everything is running smoothly though it shouldn't be a major problem. The data can span multiple days especially if its not reset. Like the rain event, you see this frequently, lets say you have flash flooding on a Tuesday and you come in on a Wednesday and you pull shift in your STP product is going to have that previous days rainfall rainfall is in there. Now that's useful because maybe you're looking at sore spots for additional flash flooding but if you want to see what's currently falling it could be contaminated by the previous days rainfall. And the bias is applied uniformly everywhere where in reality it is a little higher or lower depending on how the climate changes from one end of your CWA to the other.

Products' Applications

- Assessing flash flood potential (most commonly OHP and STP)
- Radar-gauge pair bias applied
- DHR used as input into FFMP/HPE
- Flexible time intervals (i.e. USP)
- Tracking rainfall paths
- Post storm analysis

Applications of the PPS or legacy products are primarily assessing flash flood potential, is the primary emergency reason you're going to be using these products like OHP and STP you could also use it for any other meteorological operational application like how much rain is fallen in the last hour, is it light rain or heavy rain. Like in a NOWcast you might say rainfall rate so "generally be less than a tenth of an inch." A radar gauge pair bias is applied so its including ground-truth gauge data in the calculation. Digital Hybrid Reflectivity is also an input into FFMP. Flexible time intervals are available using the user selectable Precip product. You can also track rainfall paths and conduct post storm analysis.

Summary: PPS Rainfall Products

- One Hour Precipitation (OHP)
  - Scan to scan accumulations
- Three Hour Precipitation (THP)
  - 3 top of the hour accumulations
- Storm Total Precipitation (STP)
  - From time precipitation detected
- User Selectable Precipitation (USP)
  - Top of the hour accumulations
- Digital Precipitation Array (DPA)
  - 256 data levels - used at RFC
- Supplemental Precipitation Data (SPD)
  - Alphanumeric - parameter settings

In summary, here are the products we just learned about and their primary uses. OHP tells you the past hour of precipitation that has fallen, and it is updated every volume scan. The THP tells you the past 3 hours of precip and it is generated using the past 3 top of the hour accumulations. STP tells you the amount of precip that has fallen since the first precip was detected. USP allows you to determine accumulations that are off the 1 and 3 hour time frames, using top of the hour accumulations. The DPA product is a 1-hour accumulation product that is 8-bit, but it is only used at the RFCs and is not displayable in AWIPS-2. Finally, the SPD is an alphanumeric product that gives you parameters setting used in EPRE.



That's it for lesson, Any questions email myself or anyone on this team responsible for putting everything together, and thank you for listening.



Welcome to the Topic 4 lesson on Dual-Pol QPE rainfall products. This will be the newest technology that exists as far as precipitation estimation that you will be using here in the NWS.

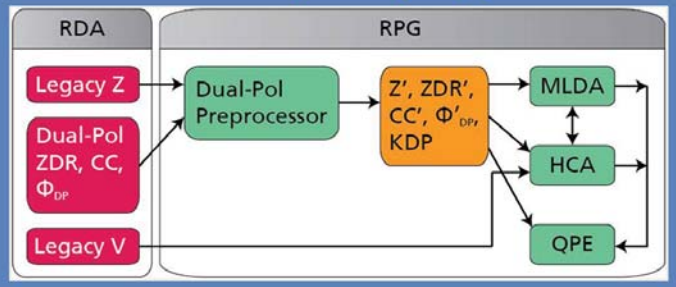
Learning Objectives

Upon completion of this lesson, you will be able to identify specific characteristics, limitations, and applications (strengths) of the following Quantitative Precipitation Estimation (QPE) rainfall products

Product	ID
Hybrid Hydrometeor Classification	HHC
Digital Precipitation Rate	DPR
One Hour Accumulation	OHA
Digital Accumulation Array	DAA
Storm Total Accumulation	STA
Digital Storm Total Accumulation	DSA
Digital User Selectable Accumulation	DUA
Digital One Hour Difference	DOD
Digital Storm Total Difference	DSD

Here are the products we'll discuss in this lesson, and what you should get out of learning about these products. Take a few moments to look at these products and the learning objectives for each. Advance the slide when you are ready.

Quantitative Precipitation Estimation (QPE) Algorithm

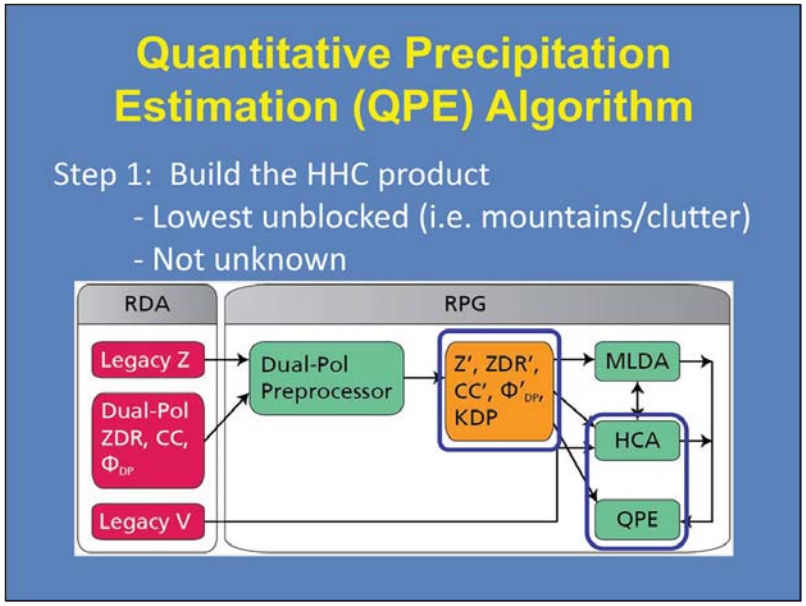


Very complex!

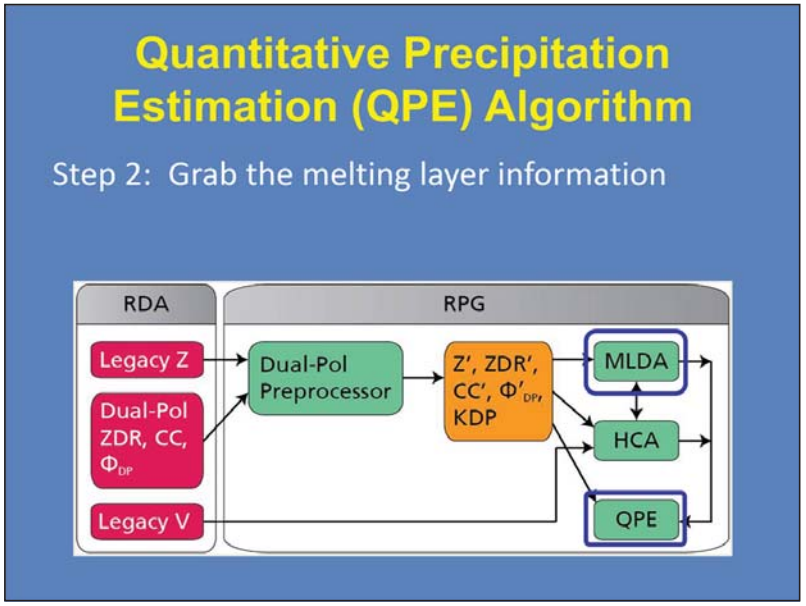
- Uses 3 different equations (instead of 1)
- Dependent upon HCA/MLDA output

Giangrande, et al. (2008)

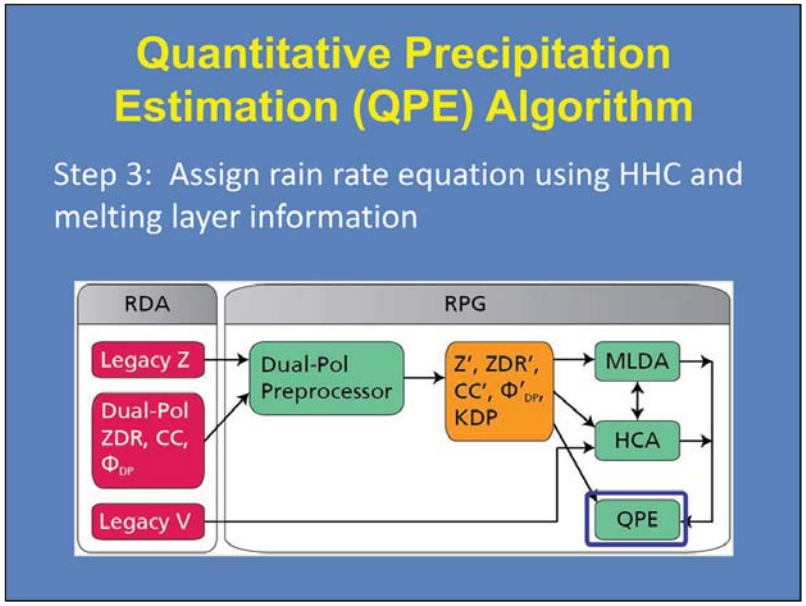
Like the PPS lesson, we'll quickly go over the algorithm that outputs the QPE rainfall products. As the name implies, the algorithm is the Quantitative Precipitation Estimation algorithm, or QPE. This graphic should look familiar from Topic 3. The QPE is much more complex than the PPS. It uses 3 different rainfall rate equations and that is dependent upon the hydrometeor classification for a given bin and where that bin lies relative to the melting layer. Comparing this to PPS, you can see how much more complex this algorithm is. I'll try to only hit the main points of this algorithm, so you can stay awake ☺



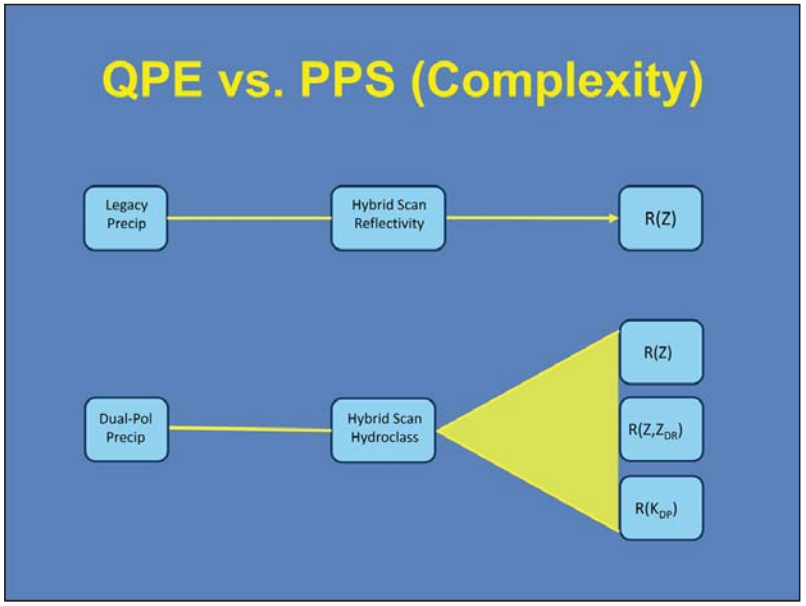
Step 1 in QPE is to build the hybrid hydroclass product, or HHC. It is the basis for what precipitation equation is used for each bin. This product is built by the base data first going into the HCA. From there, the HCA outputs the HC product, as we learned a few lessons ago. The HC product is input into the QPE algorithm where it is processed into the HHC product. Basically, the QPE smoothes the HC product to make the HHC. We'll look at this product shortly. Just know that the HHC in the QPE algorithm is analogous to the HSR/DHR product in the PPS. It is the building block for the QPE accumulation products.



Next, the QPE inputs the melting layer height information. This basically is used as a quality control check for what precipitation types are allowed in the HHC. For instance, rain above the melting layer does not make sense. Secondly, different rain rates are applied depending on the height of the bin relative to the melting layer.



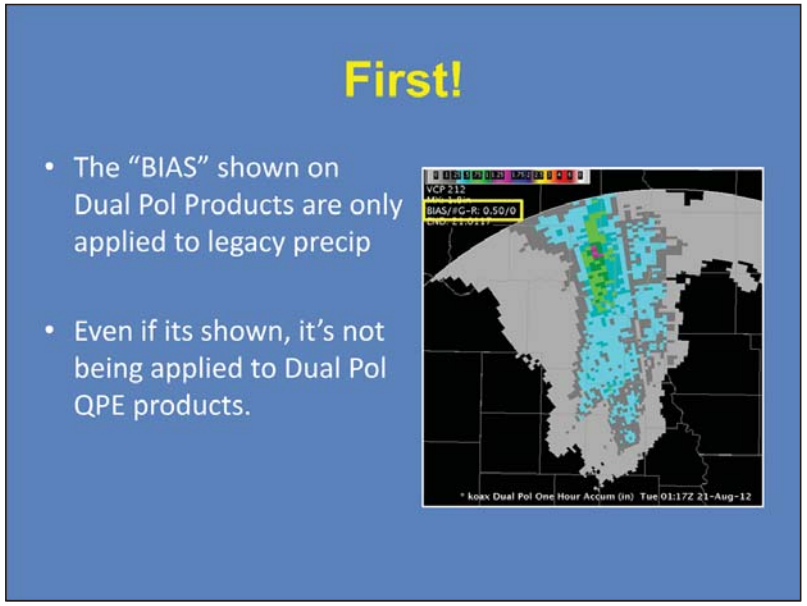
Finally, based on the HHC value for a bin, and its height relative to the melting layer, the QPE algorithm applies the appropriate rain rate equation to that bin and computes the accumulation products.



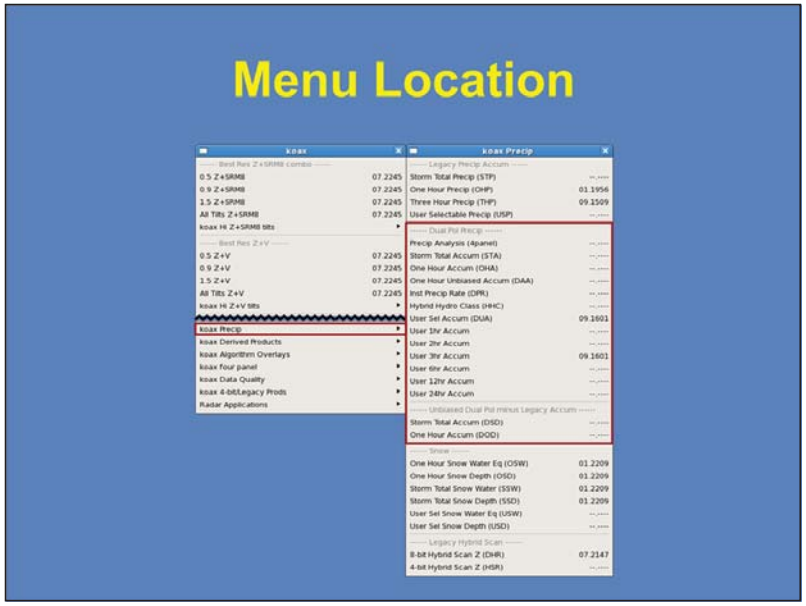
Here is just a graphic to show you the complexity of the QPE compared to the PPS. See how the PPS has only one rainfall accumulation equation, whereas the QPE has 3.

Conditions/Classifications	R method (mm/hr)
Ground Clutter (GC) or Unknown (UK)	Not computed
No Echo (ND) or Biological (BI)	0
Light/Moderate Rain (RA) or Big Drops (BD)	R(Z, Zdr)
Heavy Rain (HR) *and* blockage < 20% *and* ≤ 45 dBZ	R(Z, Zdr)
Heavy Rain (HR) *and* blockage ≥ 20% *or* Z > 45 dBZ	R(Kdp)
Rain/Hail (HA) *and* blockage ≥ 5%	R(Kdp)
Rain/Hail (HA) *and* echo is *at or below* the top of the melting layer (ML) *and* blockage < 5%	R(Kdp)
Rain/Hail (HA) *and* echo is *above* the top of the ML *and* blockage < 5%	0.8*R(Z)
Graupel (GR)	0.8*R(Z)
Wet Snow (WS)	0.6*R(Z)
Dry Snow (DS) *and* echo is *at or below* the top of the ML	R(Z)
Dry Snow (DS) *and* echo is *above* the top of the ML	2.8*R(Z)
Ice Crystals (IC)	2.8*R(Z)

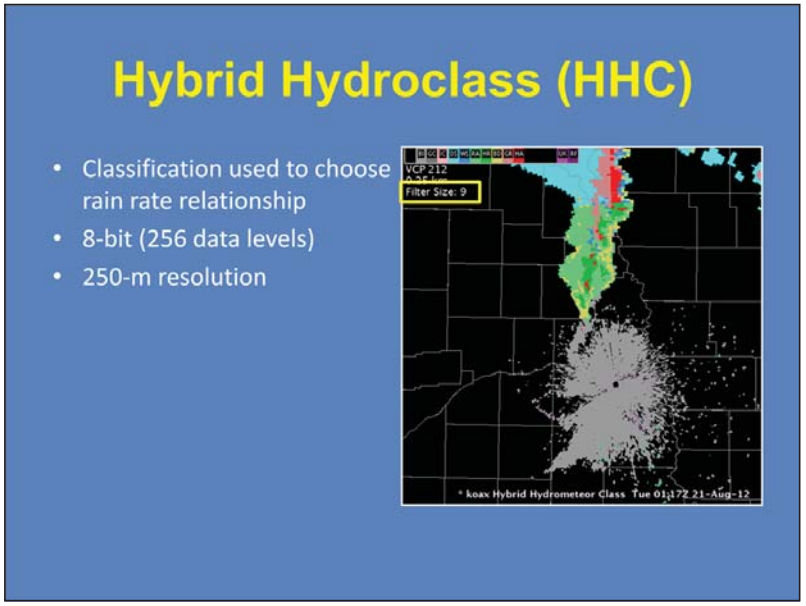
Just to give you some understanding of the complexity, and how everything is pieced together, when we talk about three different equations this is what we mean. Objects on the left are what the radar is seeing, dry snow, wet snow, heavy rain. Rain and hail, as the HCA identifies them and the HHC detects them, on the right is the R method the rainfall method, basically this is how hard its raining based on what the radar is returning. Some of the key differences, especially with regard to legacy PPS, and also some changes that have been made, keep in mind this is the latest technology. So for reflectivity greater than 45 dbz it uses KDP when calculating it with the value to get your QPE so higher KDP would be a higher QPE, where at lower dbzs you would just use Z and ZDR to give you say, 3/4 inch per hour rainfall rate, where with the KDP you might be looking at 3 inch per hour rainfall rates and the radar operations center has found that KDP works better at that level so they made that change, and then the same is true that dry snow and ice crystals use a standard R to Z, which means reflectivity means this amount of snowfall but thats customizable and the ROC suggests that you attempt to do that, to try to add some skill.



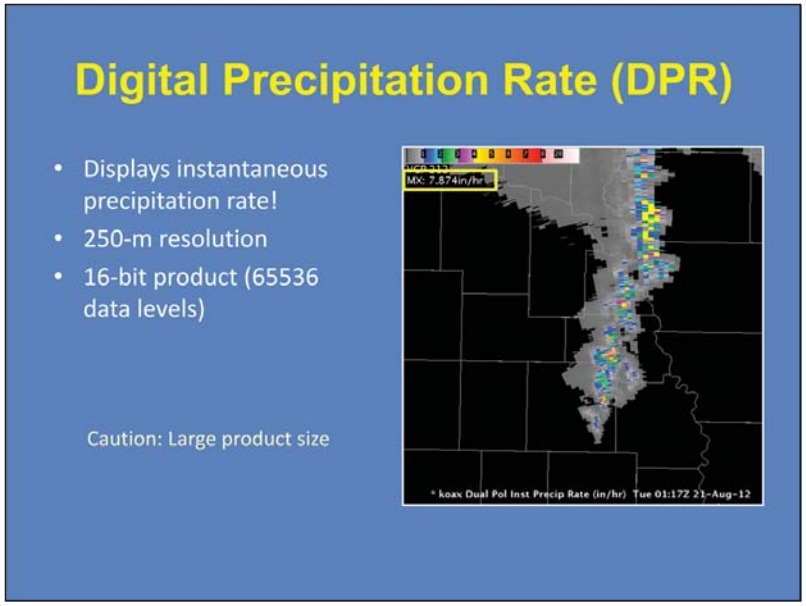
The QPE products can be found in the dedicated radar drop down menu inside the Precip submenu shown here.



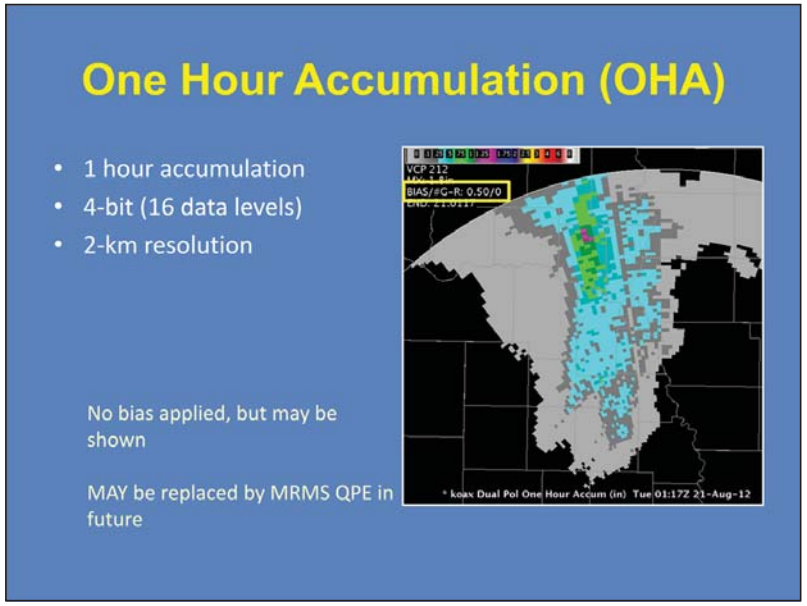
The QPE products can be found in the dedicated radar drop down menu inside the Precip submenu shown here.



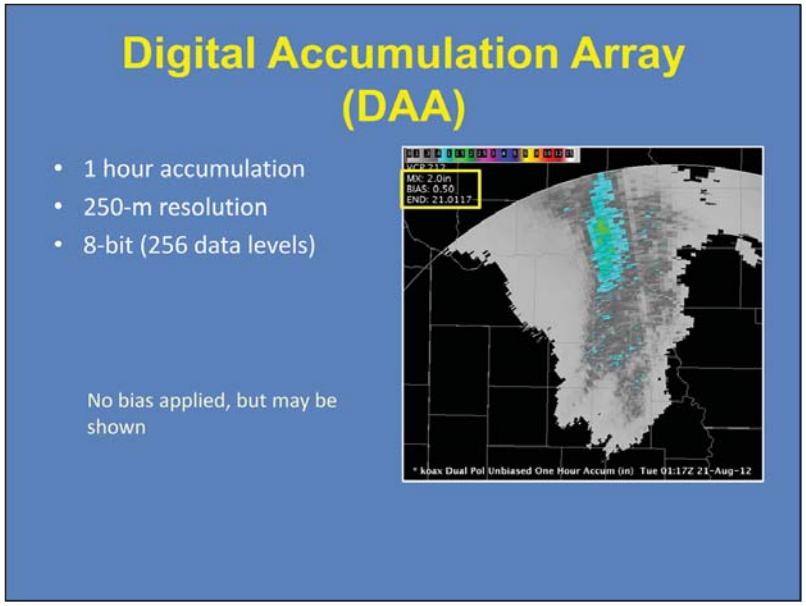
The first product, Hybrid Hydroclass (HHC), shows the hydrometeor classification that is used as an input to the dual-pol precipitation algorithm. The idea here is similar to that of the legacy DHR, and HHC has similar applications and limitations to DHR. The HHC product represents the radar's best guess of echo type to be used in the dual-pol QPE algorithm. It is an 8-bit product, meaning it has 256 data levels, but only 12 of these levels are used at this time. It has the highest available radial resolution of 250 m. The product legend here is indicating that the data have been filtered to remove speckling. The filter is not adaptable, so the filter size in the upper left of the product will always be 9. This is that smoothing referred to a few slides ago.



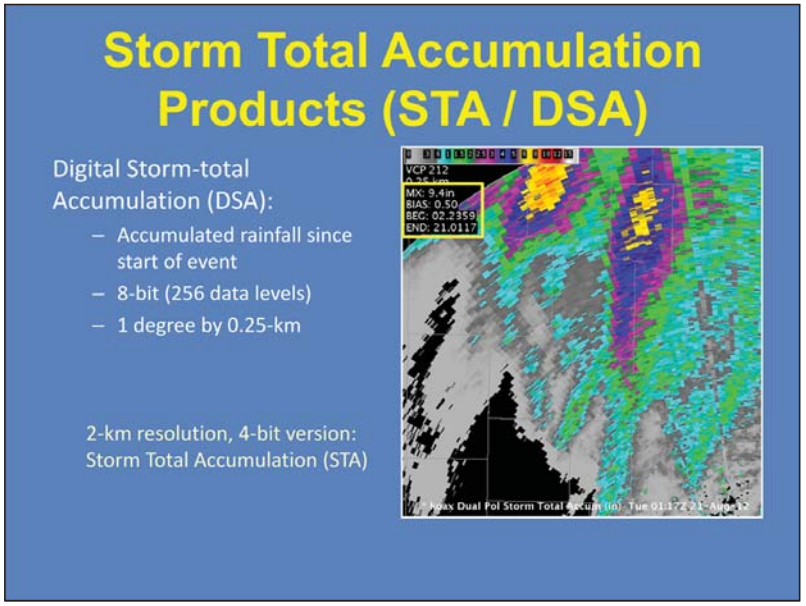
New with dual-pol is the ability to view the instantaneous precipitation rate using the Digital Precipitation Rate (DPR) product. One advantage of this product is that it can give you a good idea of where there may be heavy rain right now. DPR is also a 250 m resolution product, but is a **16 bit** product, meaning it has quite a few data levels: 65536 to be exact! The legend in the upper left lists the maximum precipitation rate anywhere in the product. One caution about DPR: because of file size, low bandwidth connections to some radars may not be able handle it.



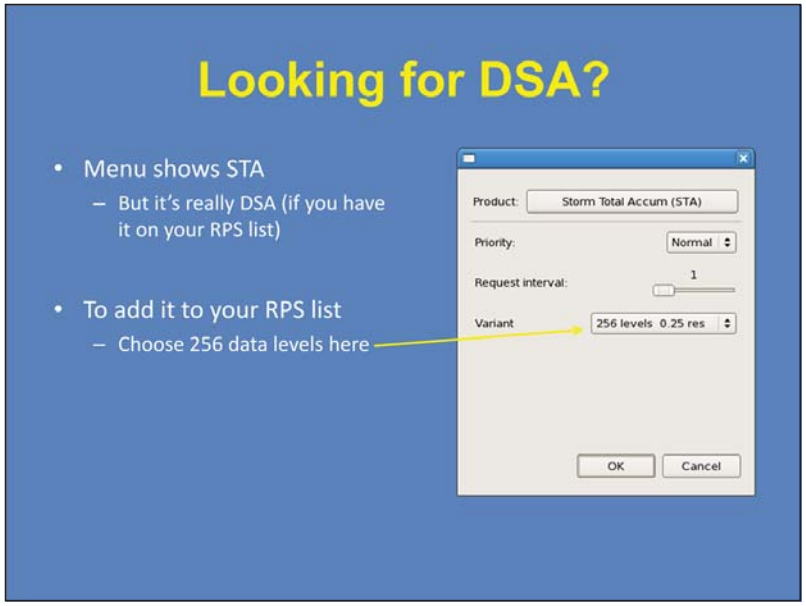
One Hour Accumulation, or OHA, is a 2 km, 4-bit, 1 hour accumulation product. This product is analogous to the legacy One Hour Precipitation (OHP) product. The legend in the upper left lists the maximum accumulation, the bias, and the product end time. Note, there is no bias applied to the QPE products at this time even though you see an entry in the product annotation. Also, if you do see a value other than 1 listed here, as is the case in this example, this means that the PPS is applying that bias to its products, and the QPE products will display that bias, but it is not applying it.



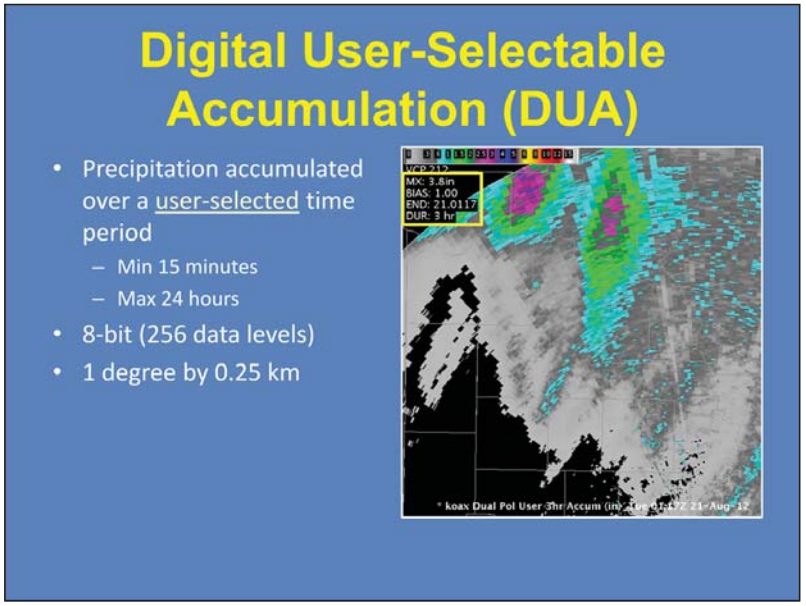
The Digital Accumulation Array (DAA) is also a 1 hour accumulation with no bias applied. However it is a 250 m, 8-bit (256 data level) product similar to the legacy Digital Precipitation Array (DPA). The legend in the corner lists the maximum accumulation in inches, the bias, and the product end time. The first 2 digits of the end time represent the day of the month and the last 4 digits are the time in UTC. Again, a bias may be shown, but it is not applied to this product.



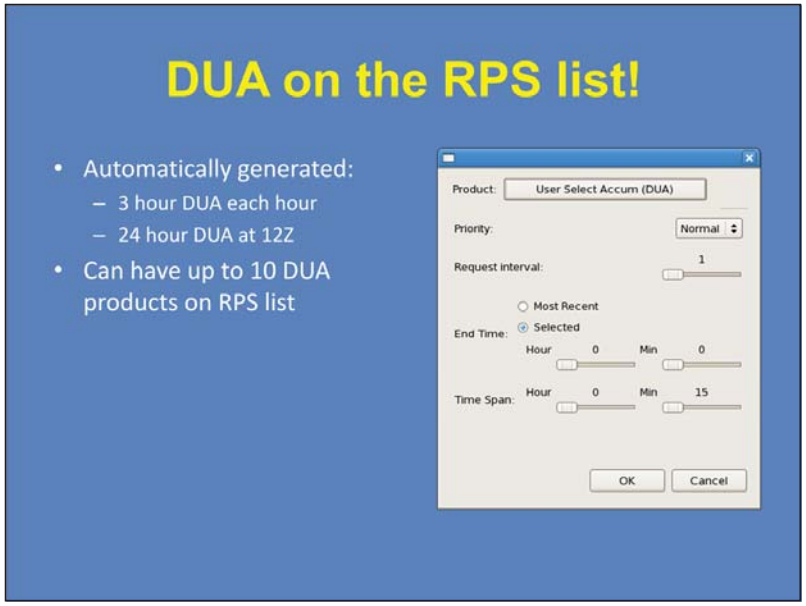
This is Digital Storm-Total Accumulation, or DSA. DSA, like the legacy Digital Storm Total Precipitation (DSP), shows total accumulated rainfall since the beginning of an event. This is an 8-bit product with a resolution of 1 degree by 0.25 km. In the legend you will find the value of the maximum accumulation (anywhere in the product), the bias (which is not applied even if it is not 1.0), and the beginning and ending times in the same format as in DAA. As with the legacy storm total products, one disadvantage is DSA may need to be reset manually at times. DSA may also include missing data. There is also a 2 km resolution, 4-bit version of this product, Storm Total Accumulation, which is not shown here but matches the resolution and data levels of the legacy Storm Total Precipitation (STP) product.



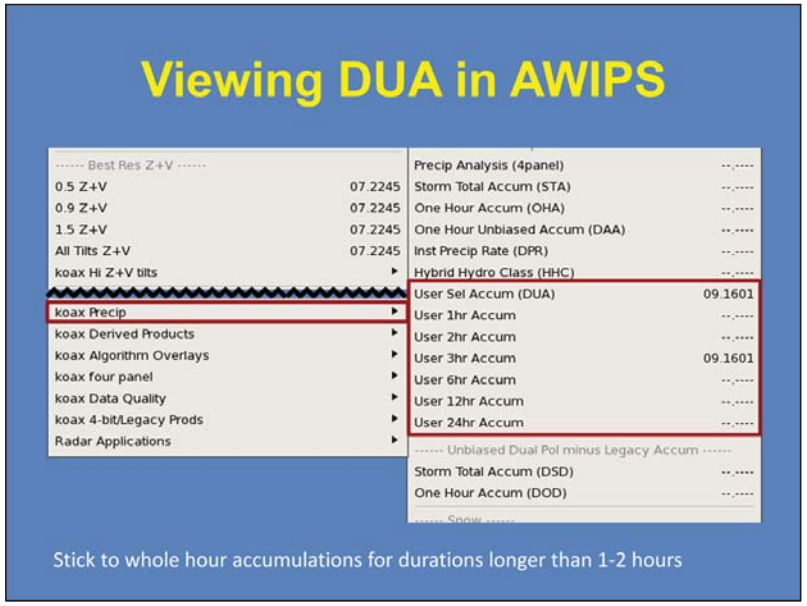
If you look in the menu for DSA, you'll never find it! Only STA is listed. So, where can you load it? If you have DSA in your RPS list and you choose STA from the menu, it will load the DSA product. Very confusing, I know! And, if you don't have DSA on your RPS list, when adding it, choose STA as the product, but then choose "256 levels, 0.25 res" to make it DSA. If you really want STA, choose "16 levels, 2 res".



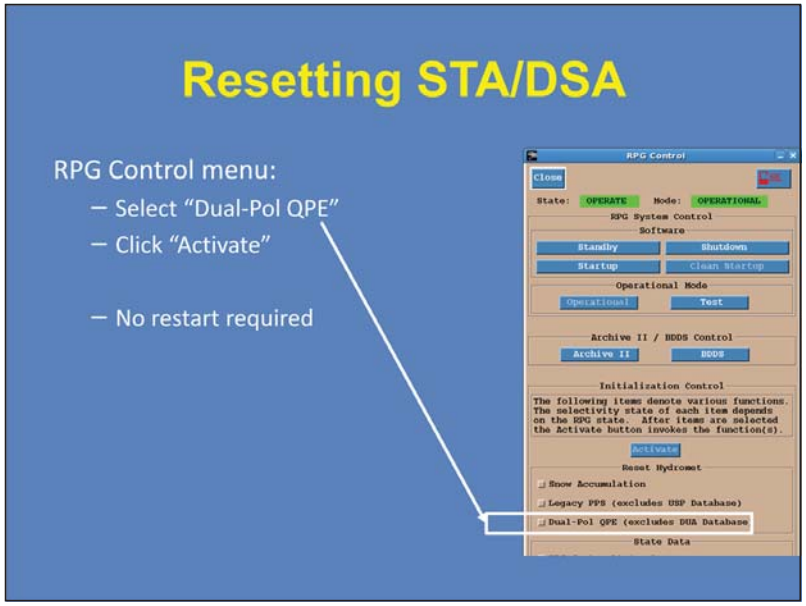
Digital User-selectable Accumulation (DUA) is accumulated precipitation over a time period chosen by the user. This time period can be as short as 15 minutes, or up to 24 hours. One advantage of this is that a precipitation product that most suits the situation can be requested, and in the following slides I'll show you how to set up customized DUA products for your office. This is an 8-bit product with a resolution of 1 deg x 0.25 km. The product legend in the upper left displays the maximum accumulation anywhere in the product, the bias (which again is not applied, even if it is not 1.0), the product end time, and the duration. As with the USP from the PPS, caution is advised because DUA may contain missing data.

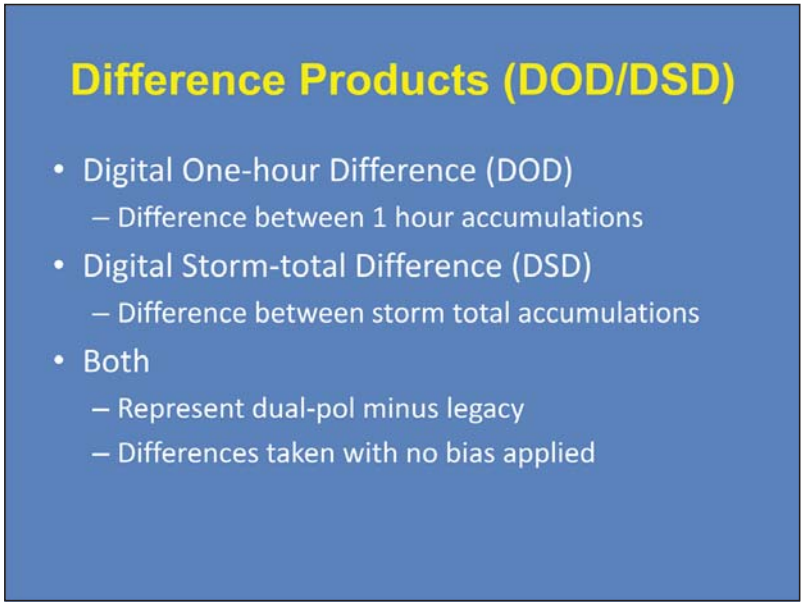


You can request DUA for other durations by adding them to the RPS list. Open the RPS list editor and choose “User Select Accum (DUA)” from the product dropdown menu. Use the two slider bars to choose the desired duration of the product in hours and minutes. Remember, the lowest duration you can choose is 15 minutes and the highest is 24 hours. Up to 10 additional DUA products can be on the RPS list at once.



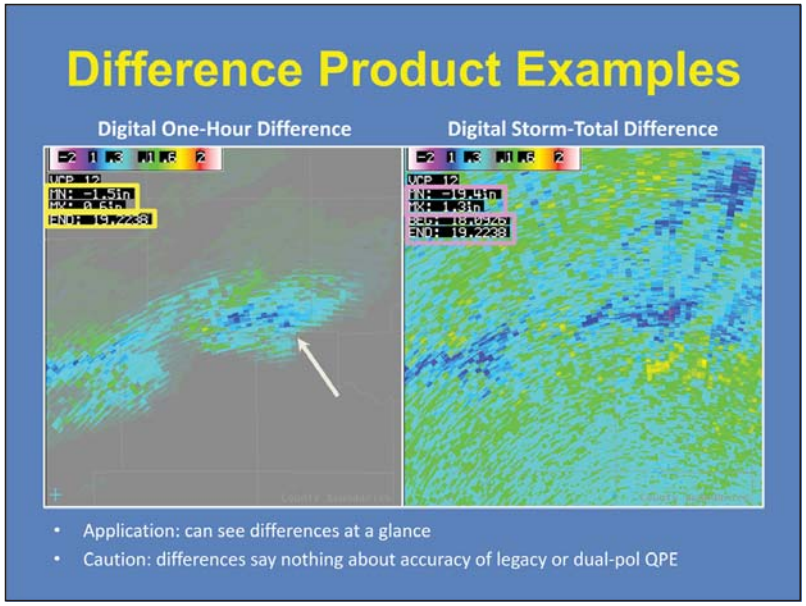
You will notice on the precip submenu that there are a number of entries for DUA. These 6 entries, seen inside the red box, will load any DUA product exactly matching the duration given in the menu, if that product is on the RPS list. Any off-hour products will be under the “User Sel Accum (DUA)” menu entry. A word of caution: if you request more than 1 off-hour product, for example, a 30 minute accumulation and a 90 minute accumulation, both products will be loaded from this menu entry. To view both of these products, use the arrow keys to step through them. It is recommended that for accumulations longer than 1 to 2 hours you request whole hour accumulations to make loading and viewing the products in AWIPS-2 straightforward.

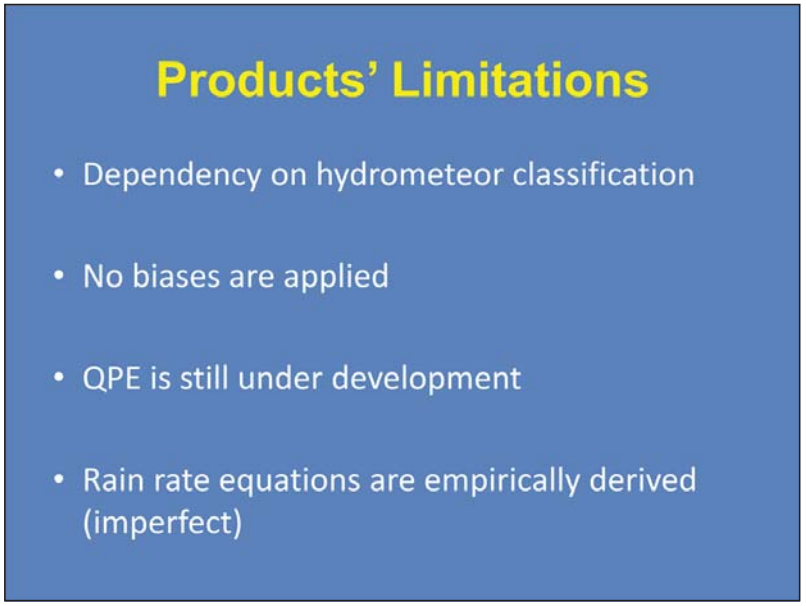




As mentioned, there may be times when the RPG does not automatically reset accumulations and you think it is warranted. Thus you are allowed to manually reset the storm total accumulation products at the RPG. Note that this will not reset the DUA products. From the main RPG HCI window, go to the "Control" button. Towards the bottom of the RPG control window are the options for resetting the Hydromet products. Look for the "Dual-Pol QPE (excludes DUA database)" option, and check that box. Now, just click the "Activate" button and this will not require an RPG restart.

Since PPS products are not going away with the arrival of QPE products, there are two difference products to help compare legacy precip products to dual-pol precip products. They are: Digital One-hour Difference (DOD) and Digital Storm-total Difference (DSD). DOD is the difference between dual-pol and legacy one hour accumulations while DSD is the difference between dual-pol and legacy storm total accumulations. Both difference products represent the dual-pol minus the legacy accumulations. These differences are taken with no bias applied. Also, note that this is not a straight subtraction because the resolutions between the analogous one hour and storm-total products used to compute these differences are not the same. The main point here is it is dual-pol minus legacy meaning positive values indicate dual-pol was higher, and negative values indicate PPS was higher, in general. Also, these products will not tell you which product is correct. It only gives you a relative measure of their differences.





Here is how the difference products look when displayed in AWIPS-2. Warmer colors, like the yellows and reds, indicate positive values, meaning the dual-pol precipitation estimate exceeds the legacy estimate, while cooler colors (blues) indicate negative values, meaning the dual-pol precipitation estimation is lower than the legacy precipitation estimation. In both DOD and DSD, maximum and minimum differences are listed in the legend in the upper left. DOD also lists the product end time, while DSD lists both the beginning and ending times. It is worth noting that while these products provide a way to see how the dual-pol and legacy precipitation estimates compare at a glance, they say nothing about which estimate is "right" or "wrong". Also, as with STA, DSD may need to be reset manually at times and may include missing data. Finally, both of these products will often have a "stair stepped" appearance, as seen here. This structure is due to the differences in range resolution between legacy and dual-pol.

With the QPE, a few unique limitations arise. First, the accumulations are only as good as the accuracy of the hydrometeor output, or HHC product. If that is off, then the QPE products will be off. Also, no bias is ever applied to any of the products at this point, even though a bias may be listed in the product. And, in general, remember that QPE is new to the WSR-88D and is still under development. As we begin to understand more, the QPE will begin to mature and become more robust. Finally, the new rain rate equations used are empirically derived, meaning these relationships may need tweaking in your local area, but that is part of the development process just mentioned.

Products' Limitations

- Standard radar limitations apply
- Missing data aren't highlighted (STA/DSA, and DUA)
- Data can span multiple days (STA/DSA)

Outside of the limitations unique to QPE, QPE is also subject to all standard radar limitations discussed to this point. Missing data affects STA/DSA and DUA. Plus, data can span multiple days for the STA and DSA products requiring a manual reset.

Products' Applications

- Assessing flash flooding events
- Tracking rainfall paths
- Flexible time intervals
- Post storm analysis

As for applications, these QPE products are useful for assessing flash flooding events…they have flexible time intervals with the DUA products. And, they can be used for tracking rainfall paths and for doing post storm analysis. These are applications identical to the PPS products.

Products' Applications

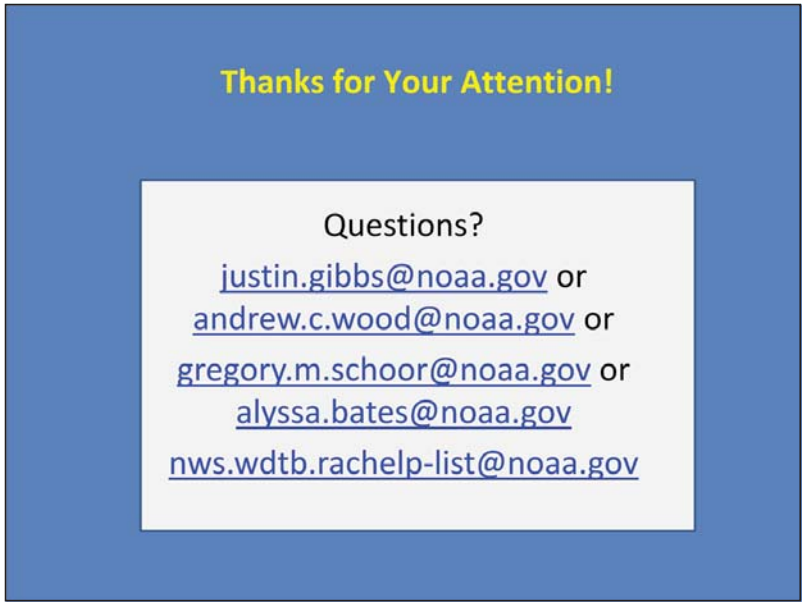
- Precipitation estimation is hydrometeor based
  - Less bright band sensitivity
  - Non-meteorological echoes don't contribute
- DPR is used as input into FFMP

Unique applications to the QPE products include precipitation is based on hydrometeor type which should result in better precipitation estimations, especially in regions of bright banding because the QPE is suppose to have less sensitivity to the bright band. Also, because the QPE is hydrometeor bases, non-precipitation echoes are not converted to rainfall! Finally, like DHR is used as input into the FFMP applications, so is DPR.

Summary

- All products are analogous to PPS
  - Except for new DPR!
  - No bias is applied
- Difference products aren't truth tellers
- QPE has shown to be better than PPS in certain situations
  - Still an algorithm under development

In summary, we saw that many of the new QPE products have an analagous PPS product except for the new DPR product. Also, no bias is applied to any of the QPE products despite it being shown in the product annotations. The difference products are not truth tellers. They only give you a sense of the relative differences between the PPS and QPE. Finally QPE has shown to be better than PPS in certain situations (like hail contamination areas) but it is still an algorithm under development. So, growing pains are inevitable. This is it for the QPE products…the next slide will be a short, ungraded quiz on this lesson. Please take it as it will help you with the end of section quiz after Lesson 24.



You are almost done with the online modules for Topic 5. This concludes Lesson 23 on the QPE rainfall products. The last online module will be on the PPS snowfall products. Go to this lesson when you are ready. If you have any questions, feel free to contact any of the emails listed below. See you in the last online lesson!



Welcome to this topic 4 lesson on PPS snowfall accumulation.

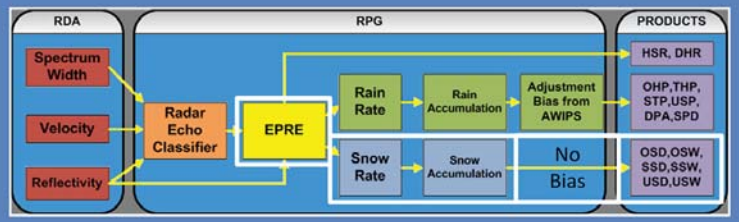
Learning Objectives

Upon completion of this lesson, you will be able to identify specific characteristics, limitations, and applications (strengths) of the following products:

Product	ID
One-Hour Snow Water Equivalent	OSW
One-Hour Snow Depth	OSD
Storm-Total Snow Water Equivalent	SSW
Storm-Total Snow Depth	SSD
User-Selectable Snow Water Equiv	USW
User-Selectable Snow Depth	USD

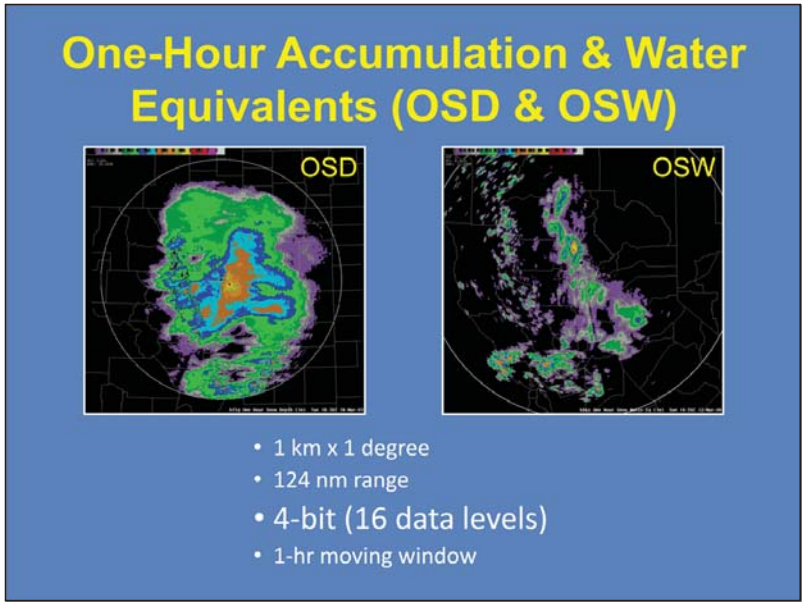
Here are the snowfall products we'll discuss. For each product you should be able to identify the characteristics, limitations and applications of each when you are finished with this lesson. Advance this slide when you are ready.

PPS Algorithm Flow for Snowfall

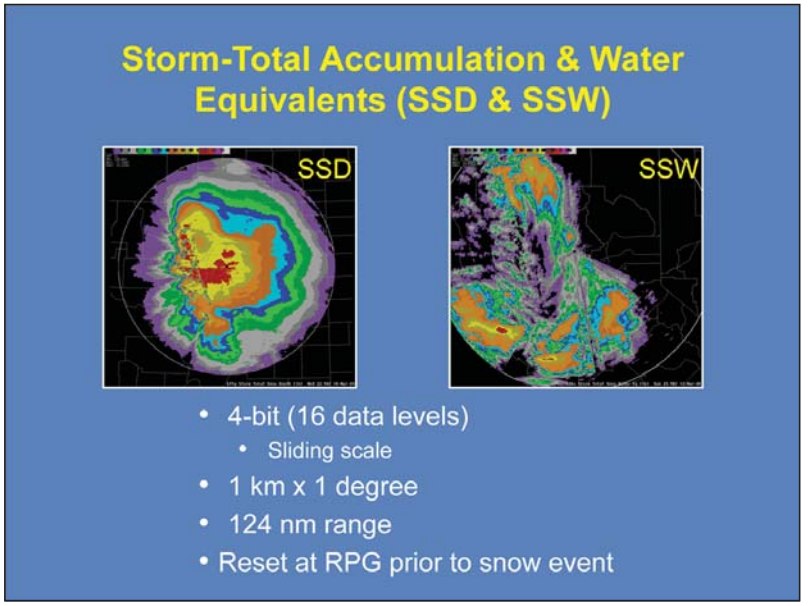


- Significantly limited in operational usefulness

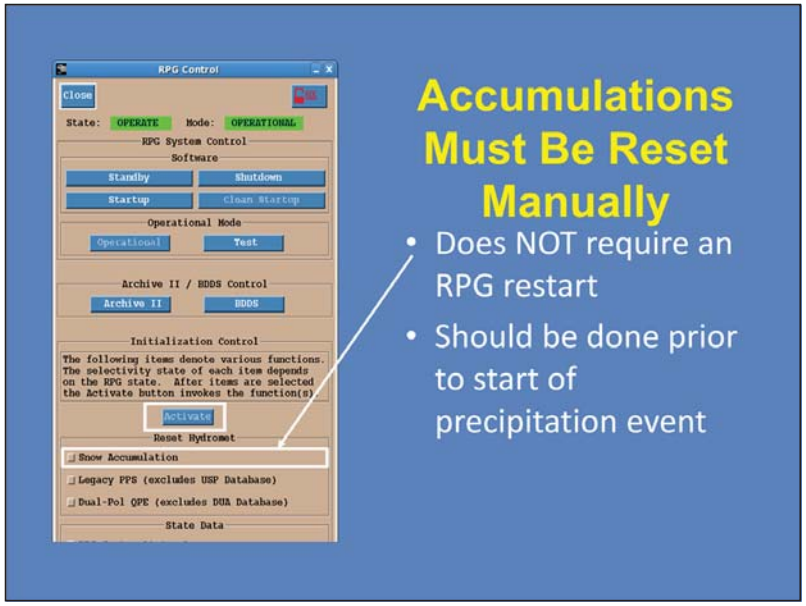
Recall this graphic from Topic 3 and the lesson on the PPS rainfall products. Starting with the EPRE function in the RPG, that computes a snow rate, which then is used to compute the 6 snowfall accumulation products we'll discuss in this lesson. Unlike the rainfall accumulation products, there are no bias adjustments with the snowfall products.



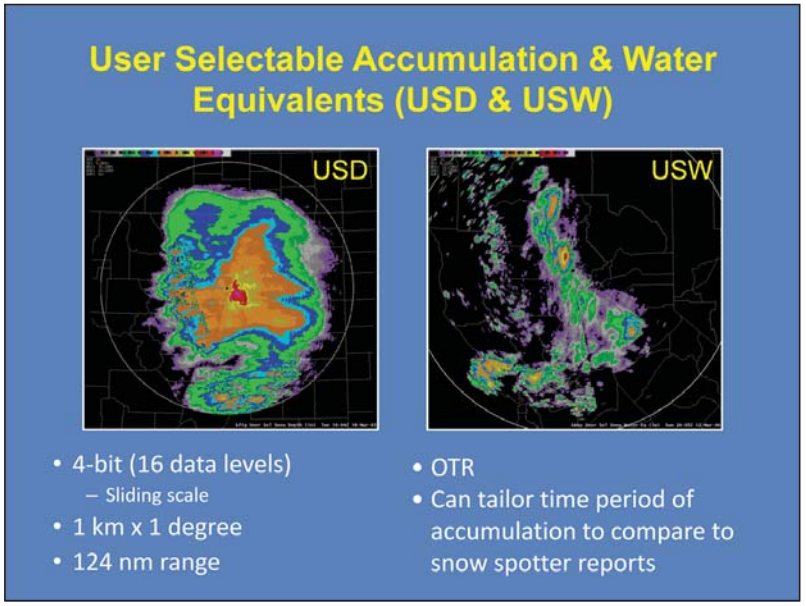
The first two products are the one hour snowfall accumulation with the product code OSD and the one hour snow water equivalent accumulation with the product code OSW. They are both 1 km by 1 degree range and azimuthal resolution out to a maximum range of 124 nm. Recall that ALL the rainfall accumulation products were 2 km by 1 deg. These have 16 data level and are accumulated over a one hour moving window.



The next two products are of the storm total persuasion: The storm total snowfall accumulation or SSD and the storm total snow water equivalent or SSW products. These are analogous to the 4-bit rainfall storm total accumulation or STP product. The one difference is that SSD and SSW are 1 km by 1 degree range and azimuthal resolution out to a maximum range of 124 nm. They have 16 data level product that use a “sliding scale” for assigning snowfall accumulations to each data level. For example, the storm total snowfall accumulation product begins with a range of data levels from 0 to 30 inches. If a sizable region of storm total snowfall accumulation exceeds 30 inches, all the data levels are doubled, increasing the max data level value to 60 inches. Both the SSD and SSW update every volume scan.



As opposed the rainfall accumulations, snowfall accumulations must be reset manually. In other words, there is no automatic reset like there is with the rainfall products. Luckily, this does not require an RPG restart. As a best practice, this should be done prior to the start of a precipitation event. To do that, go to the RPG control window…check the box for reset "Snow Accumulation" and click "Activate".



The final snowfall accumulation products are the user-selectable snowfall accumulation and snow water equivalent, product codes USD and USW respectively. The USD and USW have similar product characteristics as the storm total snowfall accumulation and snow water equivalent products. They are available as a one time request, and can be done for any duration up to 24 hrs for the previous 30 hours. They are produced as a top of the hour accumulation. Thus, you can tailor the time period of accumulation to compare to snow spotter reports.

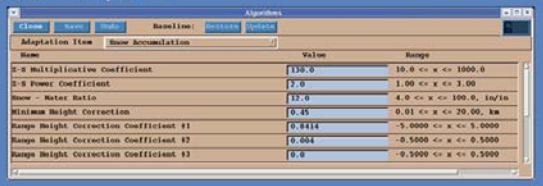
Products' Limitations

1. Assumes precip type at surface is dry snow
2. Rain falling anywhere in radar umbrella is treated as snow
3. Accumulation is not automatically reset
4. Z-S relationships and snow to liquid water ratios may not be representative
5. Melting snow, drifting snow, sublimating snow and beam overshooting may all impact amounts

Here are the limitations that can be applied to all the snowfall products. Many of these limitations are due to the fact that the radar cannot distinguish frozen precipitation types. Now, with dual-pol there is promise that this might be possible, but more research is needed to confirm this notion. Right now, snowfall products will only be available with the legacy algorithm, so we have to assume precip type is not known. Anyhow, the first limitation is PPS assumes the precip type at the surface is dry snow. Wet snow on the other hand is highly reflective, therefore estimates in these situations will be off. Also, the radar may be sampling dry snow, but below the beam it becomes wet snow due to melting and estimates will be off that way. Another limitations…rain falling anywhere in the radar umbrella is treated as snow. And, where this is the case, vast overestimation will result. Third, accumulations are never automatically reset. You must manually reset it before any event. Fourth, the Z-S relationships used and the snow-to-liquid ratios may not be representative of the event. Not only will these be different across events, but also across your CWA especially if you have a lot of terrain. Finally, melting snow, drifting snow, sublimating snow and beam overshooting may impact amounts. The radar just can't handle these types of things and it has nothing to do with the algorithm.

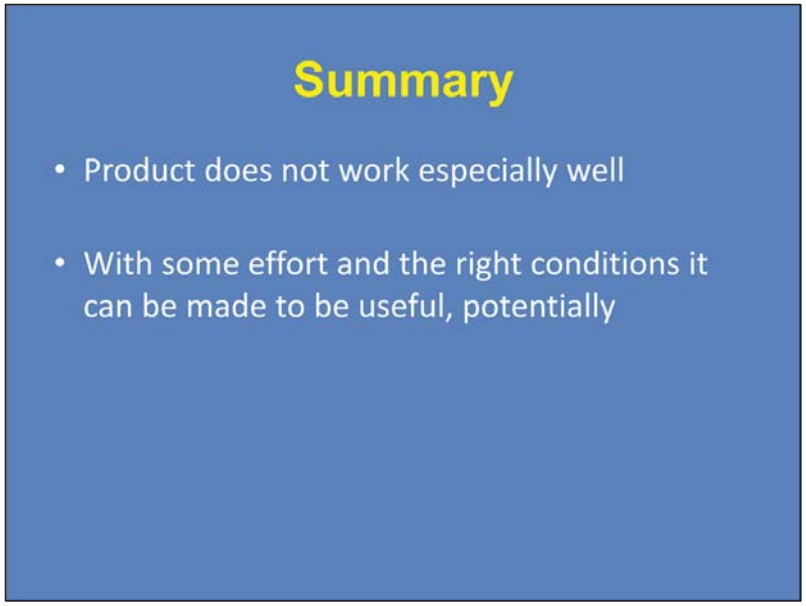
Products' Applications

1. Many Z-S relationships and snow to liquid water ratios available to tailor to local research or spotter reports



2. Updated each volume scan

Given the inherent difficulties the radar has in estimating snowfall amounts, applications for these products are certainly lacking. As mentioned there may be new hope with dual-pol, but it is too early to tell. Local studies will definitely be needed to help determine more applications in using radar-derived snowfall estimates. In the meantime, here are a few general applications. First, you can modify the Z-S coefficients and the snow-to-liquid ratio used in the algorithm at your RPG to tailor it to your local research or spotter reports. Another application is that these snowfall products are updated every volume scan.



In summary this product just does not work especially well, due to physical limitations and the application of the product. It was part of the legacy pps so its not being developed and updated, its just still in there. With some effort and the right conditions it can be made useful potentially. There are some customizations you can do.



That's it for lesson, Any questions email myself or anyone on this team responsible for putting everything together, and thank you for listening.

TAB

Radar & Applications Course

Topic:

Winter Weather Applications

Table of Contents
Topic: Winter Weather Applications

Click to jump to lesson

Lesson 1	**Winter Weather Precipitation Type Nowcasting**
Lesson 2	**Snowfall Nowcasting**



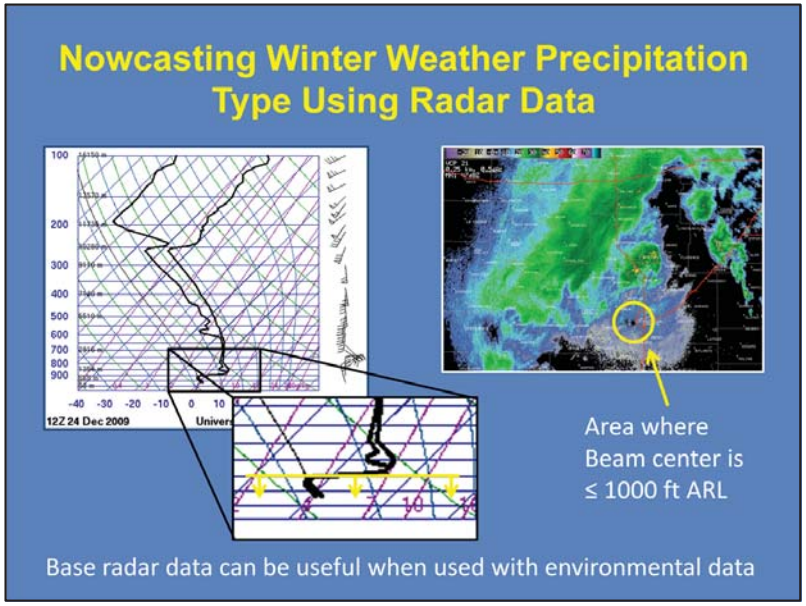
Welcome to this lesson on Winter Weather Precipitation Type Nowcasting. This training is part of the Winter Weather Applications topic in the Radar and Applications Course. Let's get started.

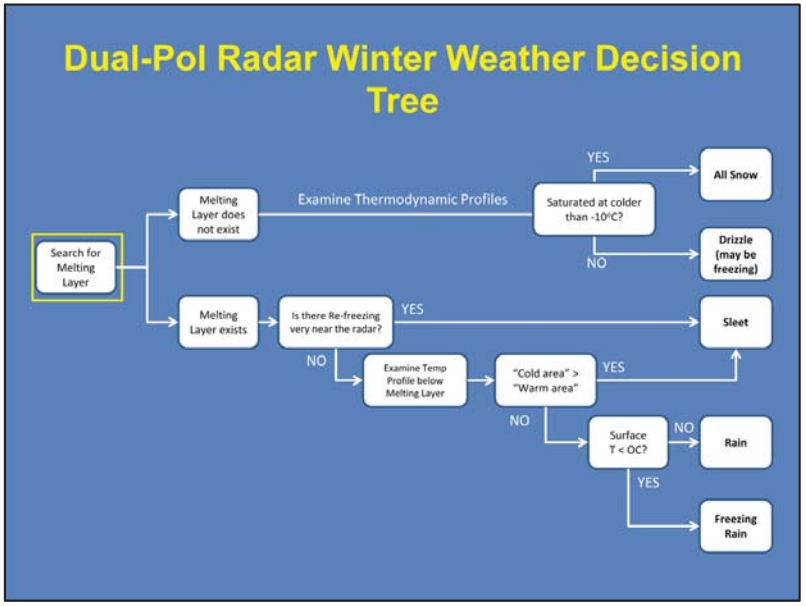
Learning Objectives

1. Identify how base radar data can indicate the presence of a melting layer and/or refreezing layer and their importance in precipitation type nowcasting

2. Identify the how the following meteorological concepts are important for determining winter precipitation type:
   - Condensation nuclei microphysics
   - Low tropospheric temperature profiles

There are two objectives for this lesson. Please take a few moments to read them over. When you are ready to proceed, advance to the next slide.





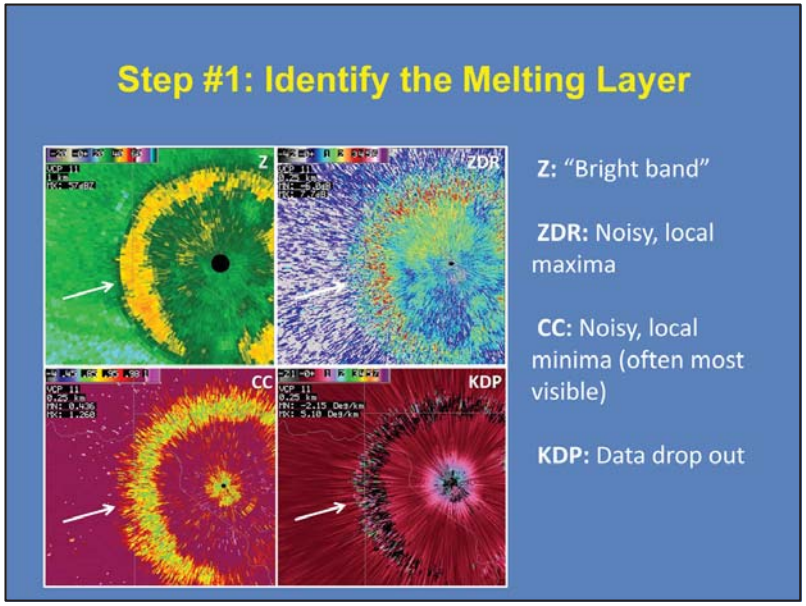
Observed winter weather precipitation type at the surface depends upon several factors, with the vertical temperature profile near the surface being particularly critical. As a result, using radar to determine precipitation type is very challenging. On the lowest radar tilt, 0.5 degrees, the center of the radar beam will already be 1000 feet above the radar height at a distance of approximately 10 nm from the radar. So, radar data alone is generally not sufficient to determine precipitation type at the surface.

That doesn't mean that base radar data isn't useful for determining precipitation type. Base radar products should be used in conjunction with other available environmental data. Using both radar and environmental data can greatly increase a forecaster's confidence in what precipitation type is occurring at the surface over large areas. Much of the discussion that follows focuses on using Reflectivity and dual-polarization base radar data in conjunction with environmental observations and model forecasts to determine precipitation type at the ground (WDTD, 2011).

To help you, the forecaster, with winter weather nowcasting, we are providing a basic methodology to determine the most likely hydrometeor occurring at the surface. This process involves radar and environmental data used in conjunction with conceptual models from winter weather forecasting.

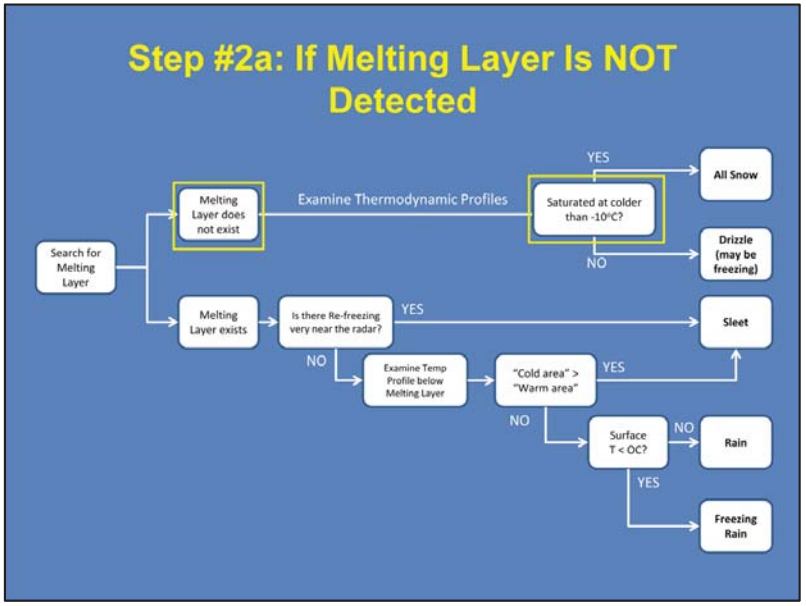
The bulk of this lesson will focus on this methodology. We will try to keep things as simple as possible. As you gain more operational winter weather experience, you will see that there is much more nuance in this process.

With all that said, let's dive in to the first step in the methodology: Searching for a melting layer!



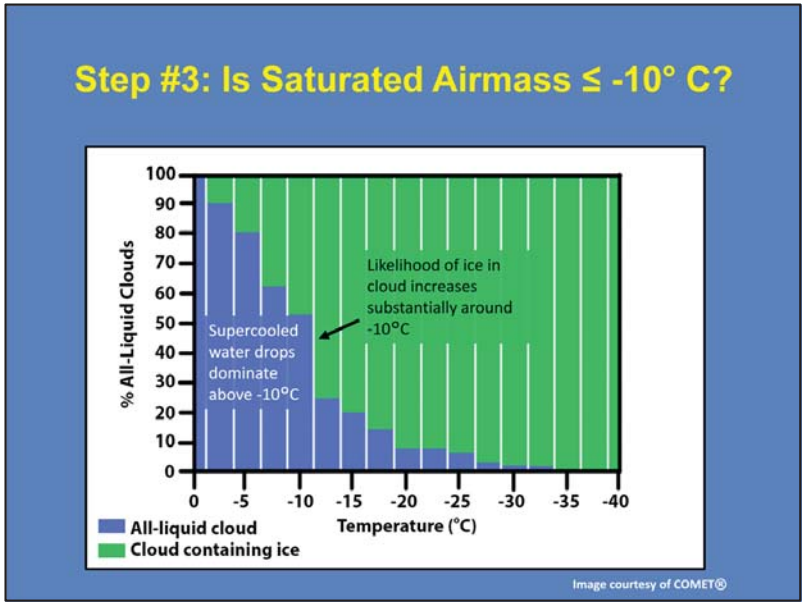
One of the most prominent applications of dual-pol base data in winter weather is the ability to detect a melting layer, when present. The melting layer appears as a "bright band" in Reflectivity (Doviak and Zrnic, 1993). However, this feature can be ambiguous, particularly on lower elevation tilts. The melting layer stands out much more prominently in the dual-pol base data (Sharfenberg and Manross, 2007).

The arrows in the graphics indicate where the melting layer is located in each product. In general, Differential Reflectivity will show a noisy, local maximum in the melting layer as frozen hydrometeors melt and appear to the radar as giant rain drops until they fully melt. Correlation Coefficient, which is often the best product for viewing the melting layer, will show a noisy, local minimum due to the increase in hydrometeor diversity. Lastly, Specific Differential Phase will usually show a data drop out in the melting layer as Correlation Coefficient values often drop below 0.9 in this region.



The next step in the process depends on whether or not you detect a melting layer in the base radar data. If you do not detect a melting layer, it likely means that the environmental temperature profile is completely below freezing. If that is the case, you are limited to two potential precipitation types at the surface: snow or drizzle.

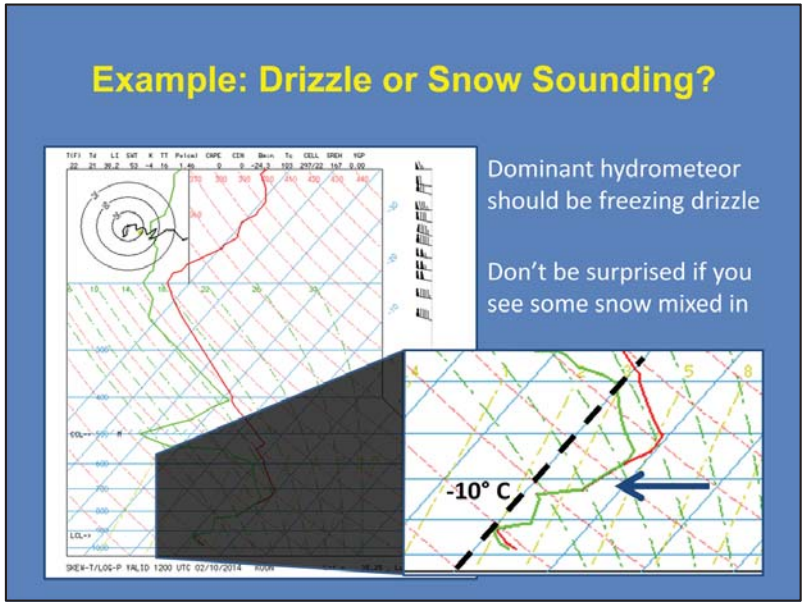
The next couple of slides will discuss how you can differentiate between the two. Before we do, I recommend you double-check your environmental data and make sure the lack of melting layer makes sense. Especially near the surface. The difference between freezing drizzle or liquid drizzle could be a shallow layer (say 100-200 m) of at or above freezing temperatures at the surface.



So why does the next step focus on a saturated air mass at, or colder than, -10 degrees Celsius? The answer involves cloud microphysics and the activation of cloud condensation nuclei (Baumgardt, 2001).

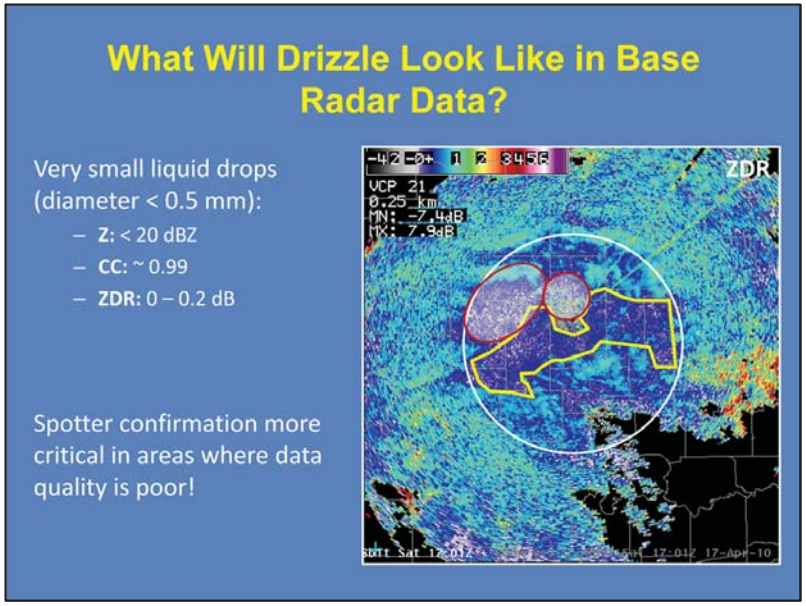
The chart on the slide illustrates it well. Cloud condensation nuclei tend to activate (create hydrometeors) as liquid or ice. Liquid CCN tend to dominate when the minimum saturated air temperatures are between 0 and -10 degrees Celsius. Once those temperatures get to -10 degrees Celsius, it's about a 50-50 split, with ice dominating when its colder than -10.

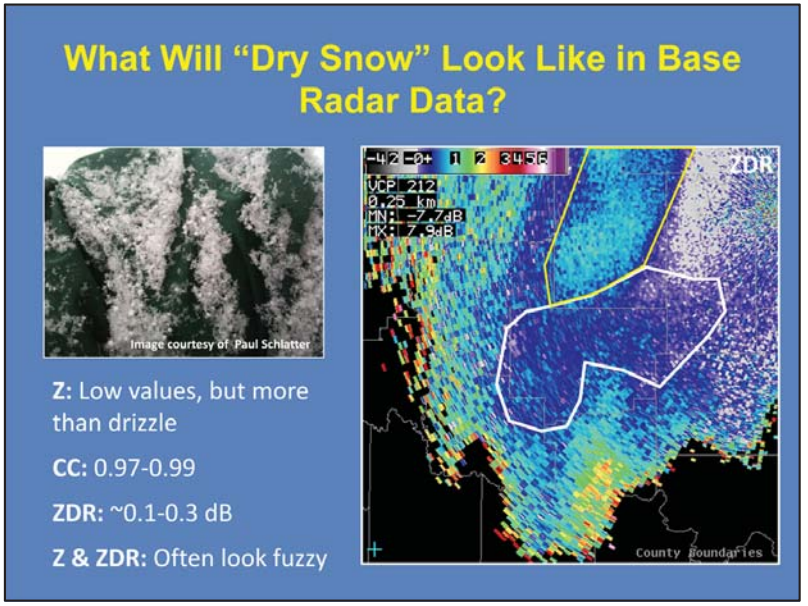
As a result, when the saturated air mass creating hydrometeors is warmer than -10 degrees Celsius, you will usually observe drizzle or freezing drizzle at the surface. Which one you get will depend on surface temperatures in that area. When colder, saturated temperatures are present, expect snow.



Here is an example sounding to illustrate the decision-making process discussed on the previous slide. This sounding represents a winter air mass that is saturated between 750 mb and the surface. I'll magnify the lower portion of the sounding so you can see the details a little better. The saturated portion of the sounding is between 0° C and -10° C. So, the dominant hydrometeor you would expect from this sounding is freezing drizzle.

Here's an important caveat to remember, especially in a sounding like this where part of the air mass is close to that -10° C threshold. What this process identifies is the dominant hydrometeor. You may see a mixture of snow and freezing drizzle because some ice cloud condensation nuclei have been activated. Those hydrometeors are just in the minority with regards to overall hydrometeor distribution.





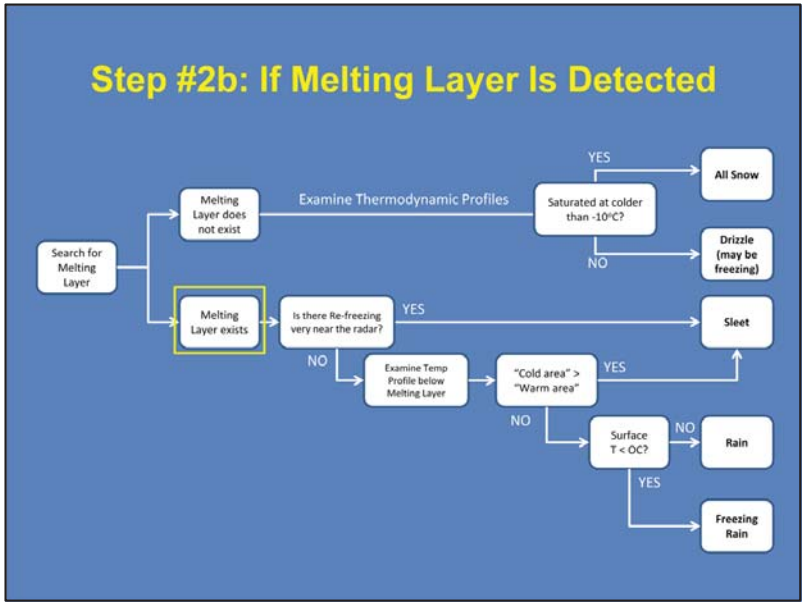
So what might drizzle look like in base radar data? Drizzle drops, by definition, are less than 0.5 mm in diameter. Drops at that diameter appear spherical and fall very slowly. They are associated with low Reflectivity values, generally lower than 20 dBZ. Since drizzle produces such a weak reflectivity signal, expect potentially poor data quality in the dual-pol products, especially at further ranges. If you trust the quality of these products, then expect Correlation Coefficient values to be high, around 0.99. Watchout for areas with CC values above 1.0 in suspected drizzle locations as data quality is poor there. Differential Reflectivity values should be low, between 0 and 0.2 dB. These values would support the generally spherical appearance of small hydrometeors

Let's look at ZDR a little more closely, focusing inside the white circle on the slide. Drizzle was observed at areas inside the yellow polygon at the surface with this event. The slightly higher ZDR values inside the white oval, but outside the yellow polygon were areas of light rain. The areas highlighted in red may have contained drizzle, too. However, the CC values were above 1.0 here. So, we would want some confirmation from observers in this area before we were confident that drizzle was present there.

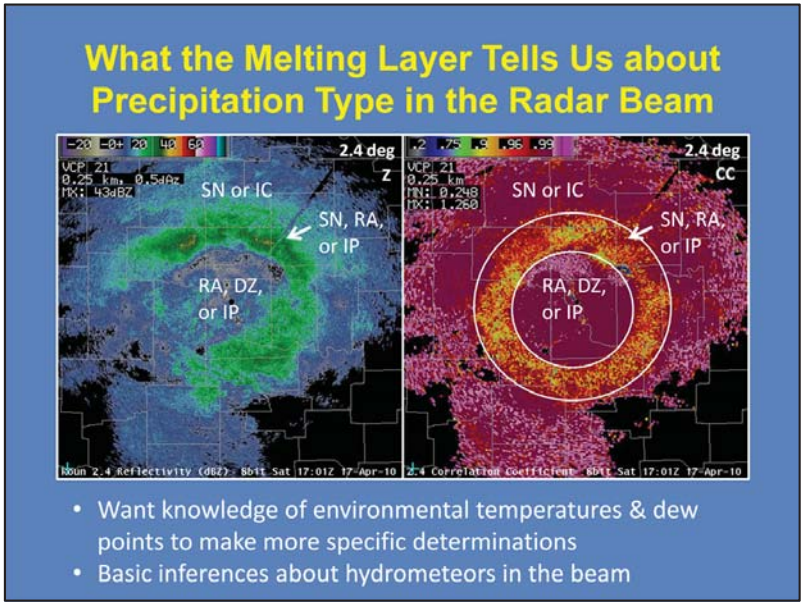
So how does "dry snow", or snow that's not melting, compare to drizzle? Well that will depend on several things, including snow structure, whether the snow crystals are aggregated, and the number of hydrometeors present. This picture shows a variety of pure and aggregated dendrites. You know, the stereotypical snow flake. Dry snow could also include plates, columns, needles and aggregated combinations depending on the environment where the crystals form and fall through.

Dry snow has generally low Reflectivity values, but it can be much higher than drizzle. As high as 40-50 dBZ in extreme cases. Correlation Coefficient values will generally be in the 0.97-0.99 range. So, similar to drizzle but noticeably lower. Likewise, Differential Reflectivity is low, but a touch higher than drizzle. In the 0.1 to 0.3 dB range. Pure crystals can be quite a bit higher than that, like the difference between the yellow polygon and the white polygon.

The last point I would make about snow is in regards to the appearance of the data. Both Reflectivity and Differential Reflectivity will often have a "fuzzy" appearance in areas of dry snow. I know that description is qualitative, but as you look at multiple cases you will get a better understanding of what I'm talking about.



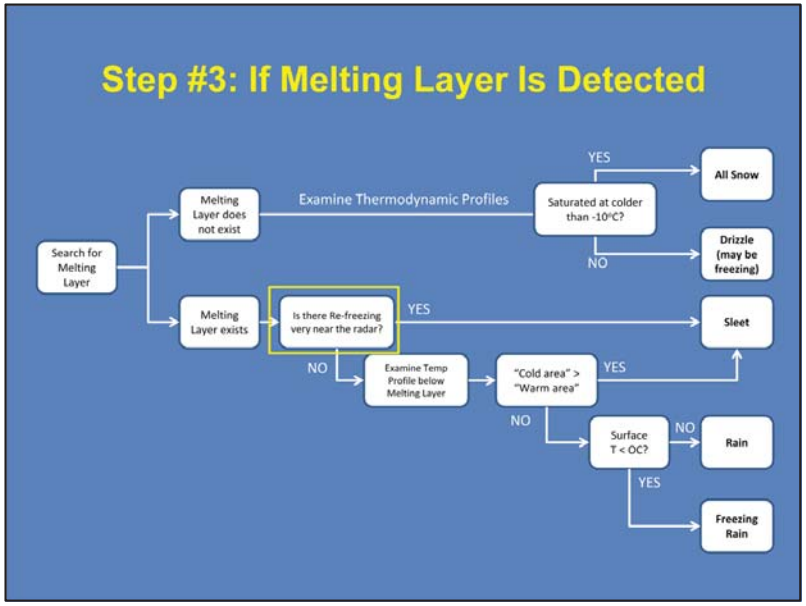
So, those steps cover situations where a melting layer isn't detected. Now let's cover what happens if a melting layer is detected. Before we move on to the next decision-making step, let's discuss what hydrometeors we are detecting in the radar beam (but not necessarily at the ground) when a melting layer is detected.



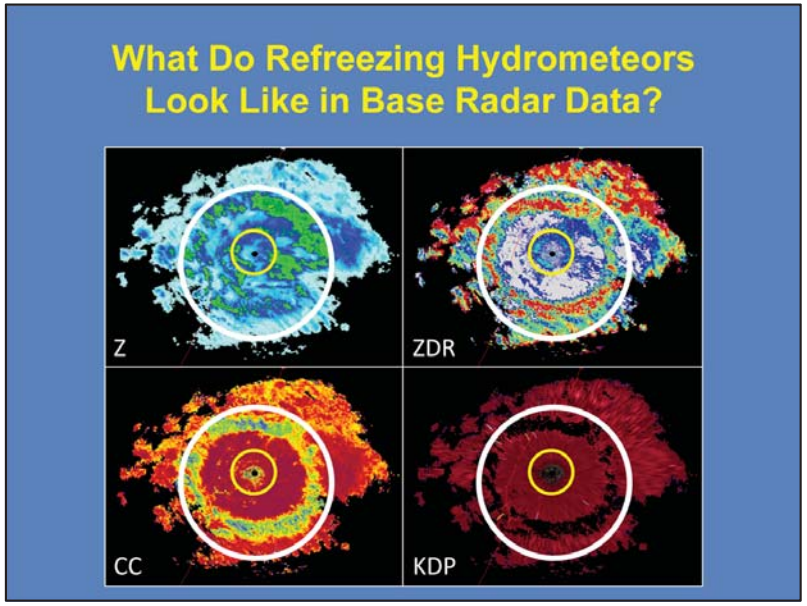
So, when we see a melting layer in the base radar data, what can we really say about the hydrometeors that the radar is detecting. While you always want to have knowledge of the temperature and dew point environmental data, there are some basic generalizations that we can realize from the base data alone.

When a single, well-defined melting layer is apparent in the base data (as in the example shown), we can infer the basic hydrometeor types that are present in the beam. Above the melting layer (or at ranges farther than the melting layer) you should expect some form of frozen hydrometeors. So, snow or ice crystals. Once you get into the melting layer, you will see some form of liquid-frozen mixture. Expect the percentage of liquid hydrometeors to increase as the radar samples towards the bottom of the melting layer. Once you get below the melting layer (or closer to the radar than the melting layer ring), the radar should be sampling some form of liquid hydrometeor. So, rain or drizzle. The one caveat is that you could also be sampling ice pellets or sleet in this area, too.

That sets us up for our next decision point: identifying a refreezing layer in the base data.

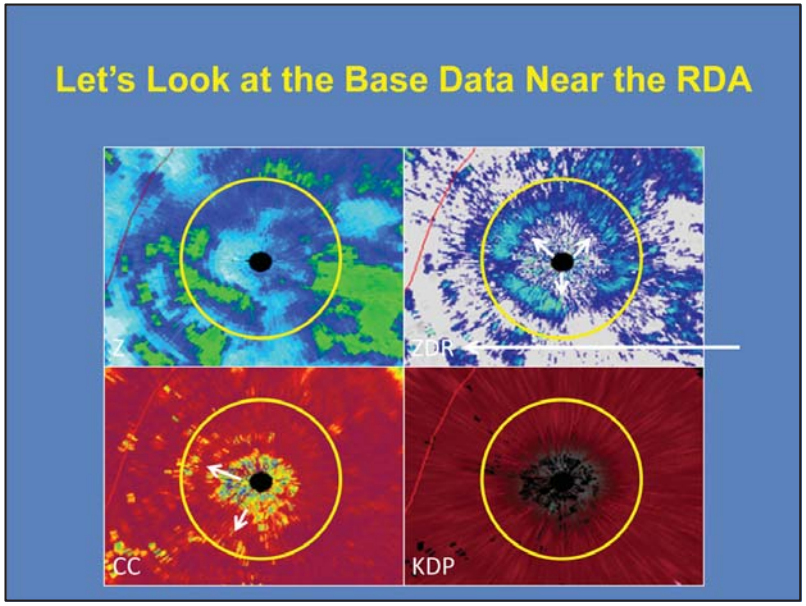


In some circumstances, the base radar data can detect the presence of hydrometeor refreezing near the earth's surface. The next step in our methodology is to check the base data and see if such a signature exists.



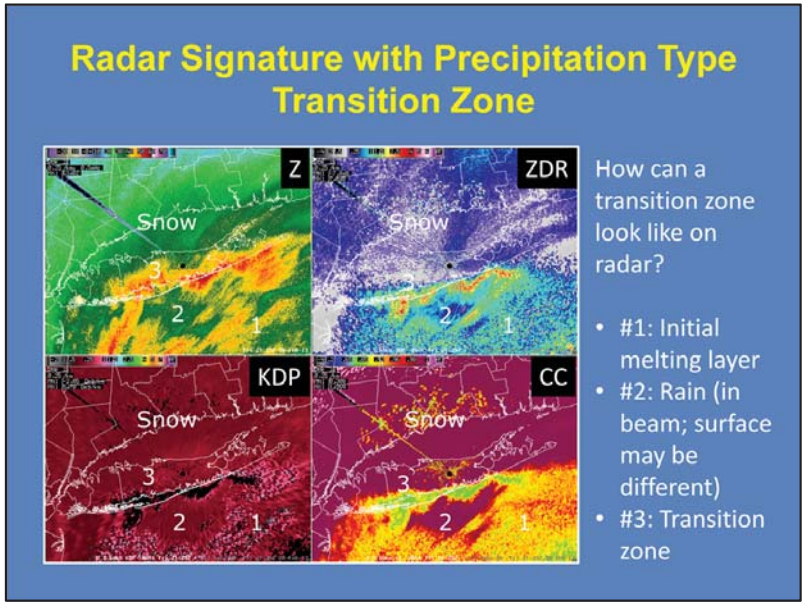
So what do refreezing hydrometeors look like in the base radar data? Well, the signature is often subtler than the primary melting layer. It will be easier to explain using an example.

I have a four-panel display on the slide with Reflectivity in the upper left, Differential Reflectivity in the upper right, Correlation Coefficient in the lower left, and Specific Differential Phase in the lower right. All of the products shown were collected at 3.4 degrees. Based on the training you've had so far, you can probably spot the melting layer in the white oval overlay. That signature is straight-forward. I now want to focus your attention on the area in the yellow circle near the radar.



Here are those same graphics, but zoomed in on the yellow circle. The key product to focus in on is Differential Reflectivity. You can see a ring of increased ZDR values almost completely around the radar. The current thinking about why this ring forms relates to how hydrometeors refreeze (Kumjian et al., 2013). The smaller, more circular drops are thought to freeze first. As they freeze, these hydrometeors reflect less energy because their dielectric constant is lower than liquid drops. Therefore, more of the returned power comes from larger drops that tend to have higher differential reflectivity values associated with them.

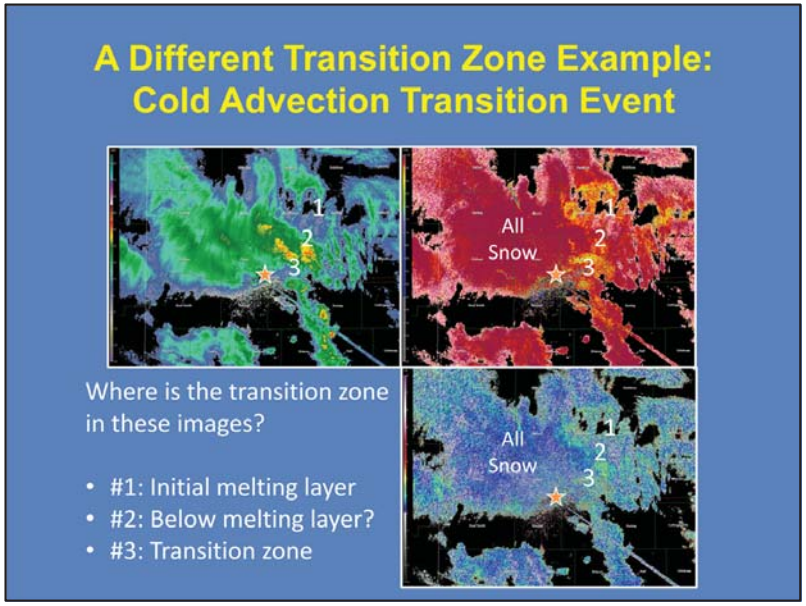
The other base data products are less functional at detecting a refreezing layer in most cases. Correlation Coefficient is the next best product, where you will usually see a noisy local minimum in values. However, the data field can look very spotty compared to ZDR. Reflectivity values tend to decrease once refreezing begins, leading to a donut-like appearance in values around the RDA. Lastly, KDP provides little assistance at all in identifying refreezing.



The previous slides present the "classic" refreezing signature. Now lets look at a common, but not so "classic" situation: When ice pellets, or even a narrow transition zone between snow and rain, are located at some distance from the RDA. Let's show you an example.

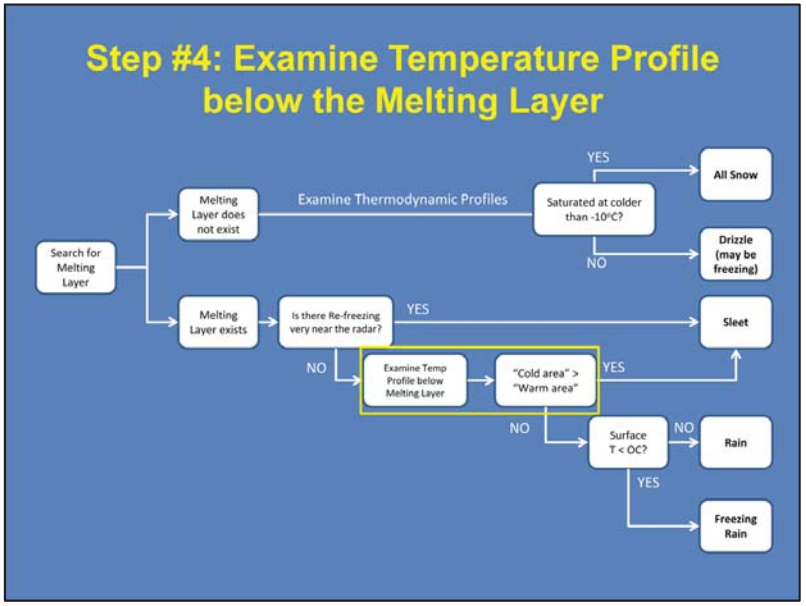
The images on the slide show the base data for a nor'easter impacting Long Island with a mixture of winter precipitation. At the RDA and to the north, no melting layer is visible and snow was observed there. Just south of the RDA is a completely different feature. If we look out over the ocean, we see a traditional melting layer aloft with higher Reflectivity, ZDR, and KDP, with lower CC values. Moving inside the melting layer, the parameters change as we would expect inside the melting layer. The hydrometeors sampled by the radar in this area are likely rain, but that may not be what is seen at the surface.

Moving closer to the radar leads to a completely new feature. We see what looks a bright band on steroids! The parameters change just like we would expect in the melting layer, but in a more intense fashion. The maxima/minima are much narrower, too. This feature corresponds to the transition zone between the snow to the north and rain further south. In this narrow transition zone, a mixture of snow and sleet was reported.

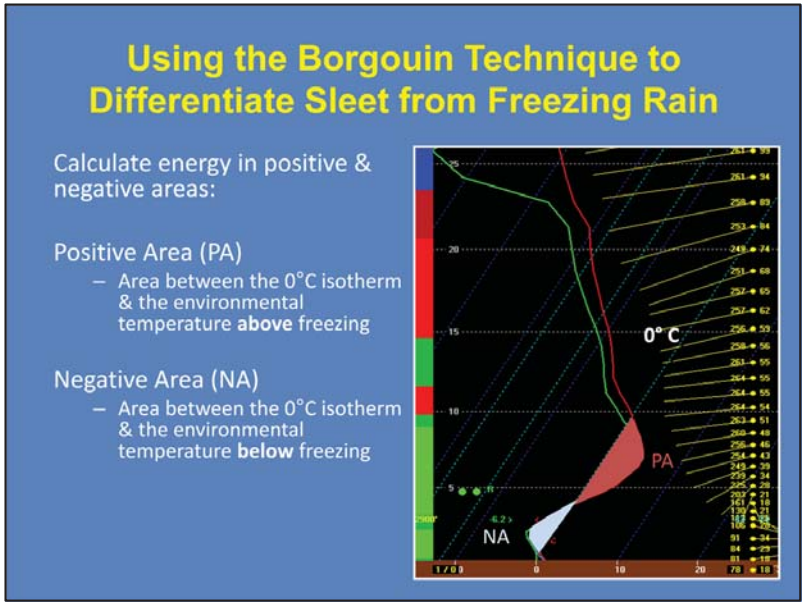


Just so you don't get the idea that identifying transition zones is always easy, here's a different example involving cold advection over the southern high plains. The star on the graphic indicates where the RDA is located. A front is bringing in cold air from left to the right, resulting in a change over from rain to snow from west to east, with a brief period of mixed precip in between.

I'll use the same labels as I did on the previous slide. "1" indicates where the melting layer is aloft. "2" is the area where we expect to see liquid precip in the radar beam. Notice in this example how the Correlation Coefficient and Differential Reflectivity values haven't fully rebounded to what you would normally expect below the melting layer. It's quite possible that the beam never fully samples below the melting layer. In the area of label "3" we see the high Reflectivity and ZDR in the same area as the very low CC values. The transition zone is likely in this area, but it's hard to say with certainty from just the radar data because the precipitation coverage isn't uniform.



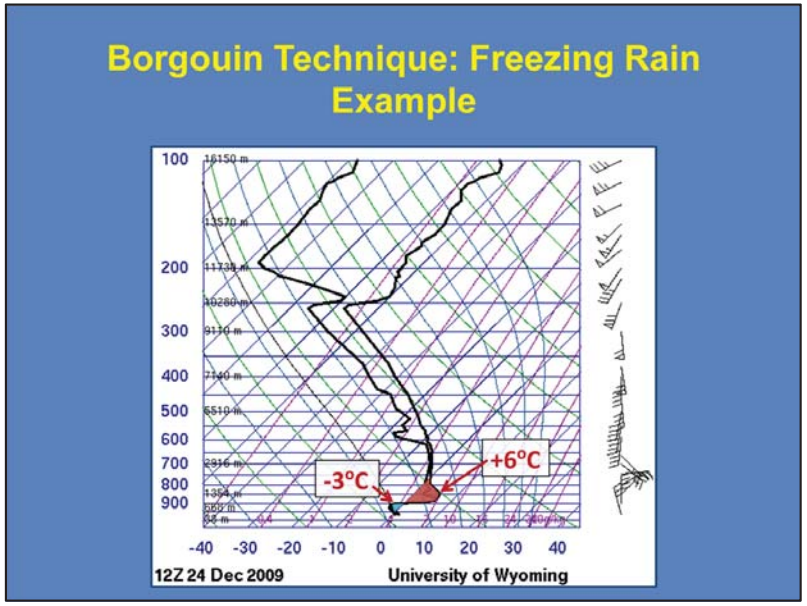
If you can't see a refreezing layer (or even a low-level transition zone) in the radar data, then you will need to look at low-level observed and model soundings to see if sleet is possible.



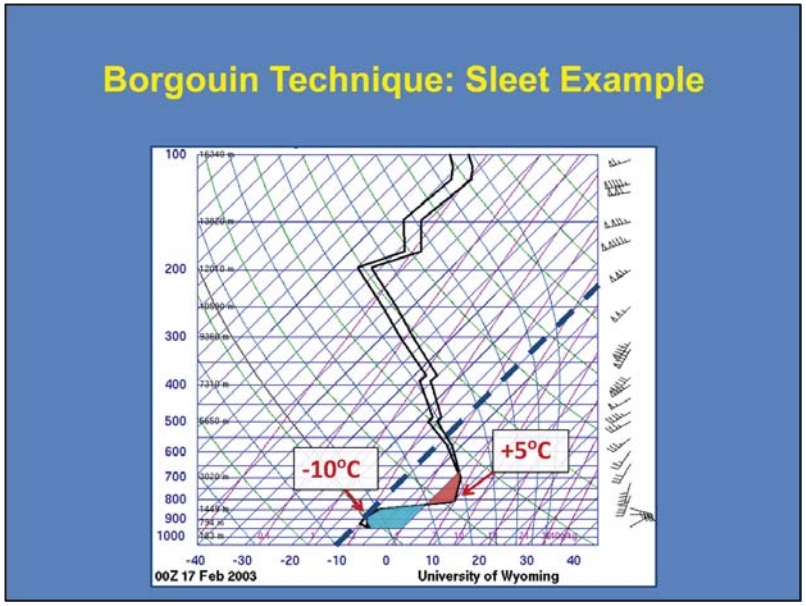
There are multiple techniques that have been published in scientific literature about determining precipitation type. For the purpose of this methodology, we will focus on the Borgouin Technique (Bourgouin, 2000) to identify whether we expect sleet or freezing rain to be observed.

The Borgouin Technique focuses on the amount of energy in positive and negative areas. What are these areas you ask? The positive area is the area between the 0 degree isotherm and the environmental temperature ABOVE freezing. Conversely, the negative area is the area between the same isotherm and the environmental temperature BELOW freezing. Which precipitation type you should expect will depending on the magnitude of these areas relative to each other.

Let's look at a couple of examples to see how this works.

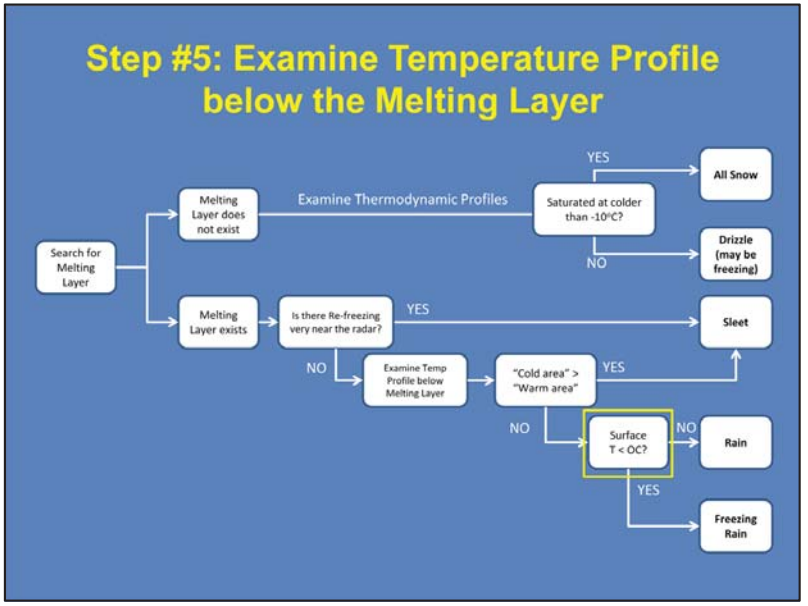


Here's the first example where we will apply the Borgouin Technique. We'll focus on the portion of the sounding below the highest melting layer. We see the positive area highlighted in red and the negative area highlighted in blue. The positive area is larger than the negative area, which indicates freezing rain should be expected with this sounding.

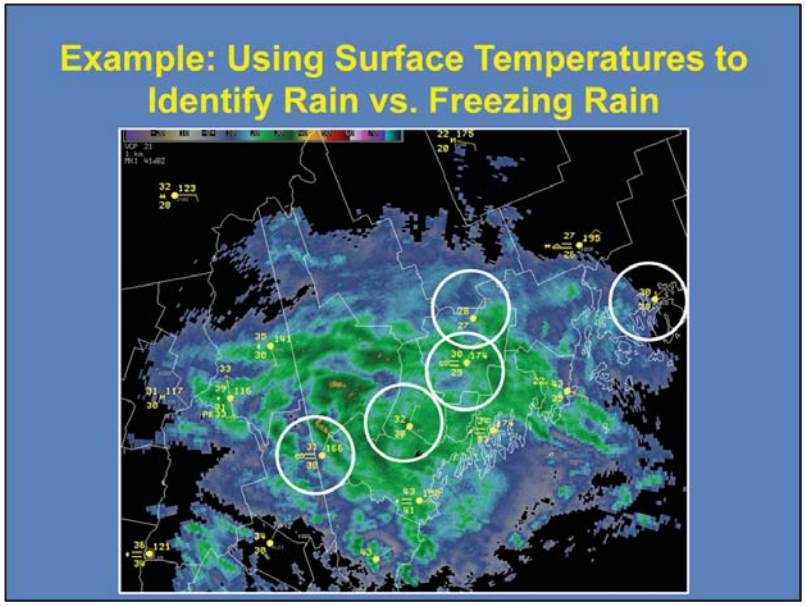


Here's a second example where will apply the Borgouin Technique. Unlike the previous example, you can see the negative area is much larger than the positive area. This suggests that sleet is likely with this sounding.

One important caveat to remember. Earlier we talked about the cloud condensation nuclei and the importance of -10 degree Celsius. If the saturated air mass in the sounding is completely at or above -10 degrees Celsius, than you should still expect freezing rain even if the negative area is larger than the positive area. If this is the case, then you probably didn't see a melting layer in the base data. Still, you should keep cloud microphysics in mind when analyzing soundings at this point of the methodology as well.

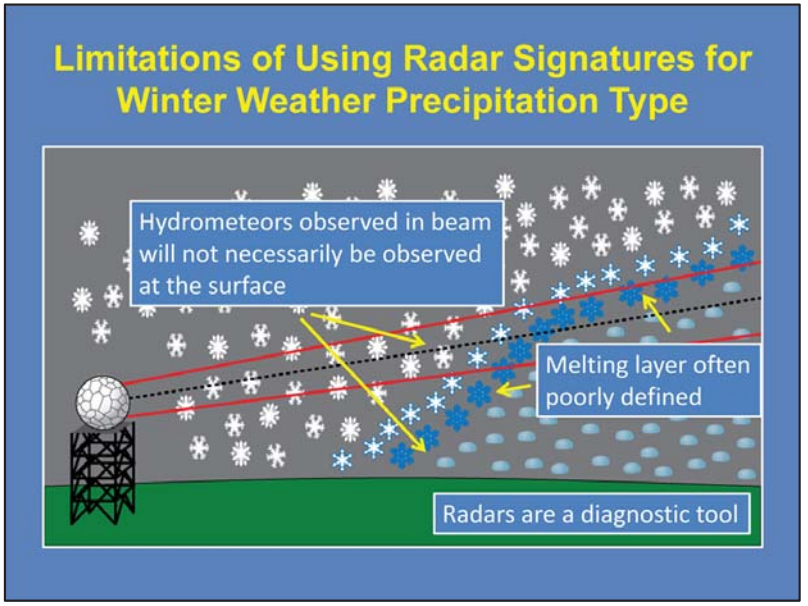


The final step in this methodology is to determine whether you have rain or freezing rain at the surface. To make this determination, you need to look at surface temperatures in the area where precipitation is falling.

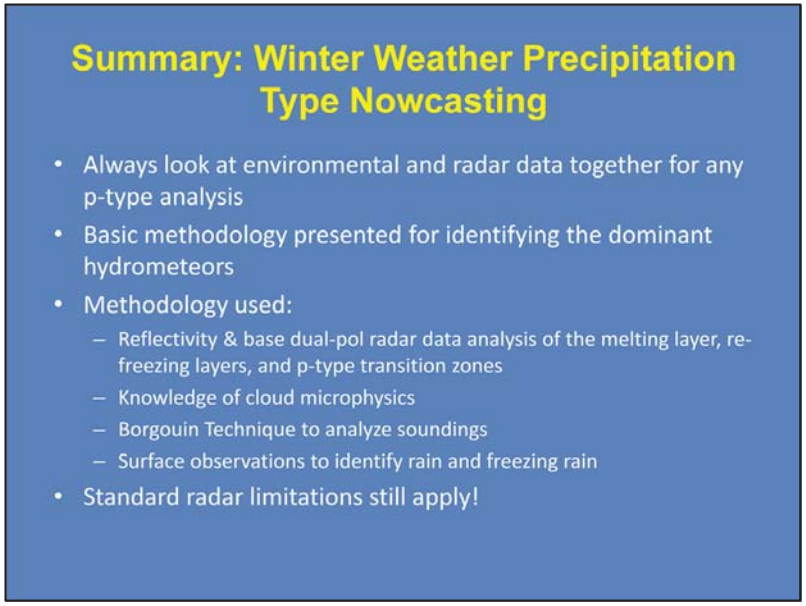


The only way to determine if you are getting freezing rain at the surface is by comparing areas of precipitation with observed surface temperatures. After all, freezing rain is still rain until it comes in contact with a freezing surface.

In the example shown on the slide, we have a broad area of precipitation visible on radar. The surface observations show a range of temperatures from the low 40s Fahrenheit (along the coast to the south) to the low 20s Fahrenheit to the north. In the center of the display, we see four observations in the precip that suggest a broad area of freezing rain is present. It helps that two of the surface observations are reporting freezing rain. However, there are two observations in this same area that are also below freezing. The observations along the coast are all above freezing except way to the east in Bar Harbor. These observations are consistent with shallow cold air damming along the southern side of the Longfellow Mountains in Maine, a classic set up for freezing rain in this area.



While we have covered our methodology to its logical completion, I want to remind you of some basic limitations for using radar data to diagnose precipitation type. The most obvious one for this topic is below beam effects. The hydrometeors detected in the radar beam will not necessarily be the same as what is observed at the surface due to below beam effects. Another limitation to remember is that the melting layer is often not well defined during winter weather events, especially when conditions are changing rapidly or when the 0 degree isotherm is strongly sloped. Lastly, remember that radars are a diagnostic tool. Radars alone cannot be a substitute for a detailed and thorough environmental analysis.





In summary, this lesson discussed how forecasters can use environmental data and radar data together to perform a rudimentary precipitation type analysis. It's important to use these data together as radar data alone will not provide the needed context for a thorough investigation.

I provided a simplistic methodology for determining the dominant hydrometeor in a given area. Remember that more than one hydrometeor type may be present in some circumstances. The methodology focused on four items. First, using Reflectivity and dual-pol base data to identify a melting layer and look for refreezing layers and precipitation type transition zones. Second, having a rudimentary understanding of cloud microphysics to know when liquid or frozen hydrometeors are more likely. Third, using the Borgouin Technique to analyze the portions of a sounding below the melting layer to determine if sleet or freezing rain is more likely at the surface. Lastly, we used surface observations to distinguish areas of rain from freezing rain. Since this methodology relies on radar data, it's important to remember that the standard radar data limitations apply.

The next slide contains the quiz for this lesson. You will need to get a score of 70% or higher to receive completion credit for the lesson.

If you have passed the quiz, then you have successfully completed this lesson. If you have any questions, please contact us using any of the e-mail addresses listed on the bottom of the slide.



Welcome to the winter weather applications lesson on snowfall nowcasting. This lesson should last about 35-40 minutes.

Learning Objectives

- Explain one major reason behind choosing the regional ZS coefficients in the Snow Accumulation Algorithm (SAA)
- Given an SAA output of snow water equivalent in a single WSR-88D domain, identify areas where precipitation rates are likely to be in error due to
  - Beam overshoot
  - Bright banding
  - Precipitation evaporation/sublimation
  - Beam blockage
  - Horizontal drift of falling snow
- Given a sounding, 1 km above ground level wind and dewpoint depression, and a surface observation network, determine the most likely sign of the SAA error given the potential error sources above.

Explain one major reason behind choosing the regional ZS coefficients in the Snow Accumulation Algorithm (SAA)

Given an SAA output of snow water equivalent in a single WSR-88D domain, identify areas where precipitation rates are likely to be in error due to

Beam overshoot

Bright banding

Precipitation evaporation/sublimation

Beam blockage

Horizontal drift of falling snow

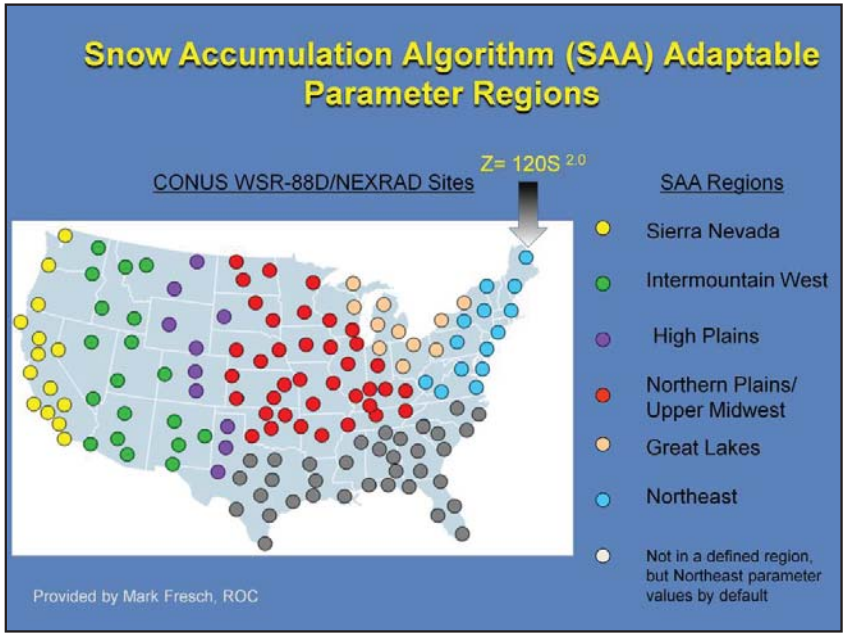
Given a sounding, 1 km above ground level wind and dewpoint depression, and a surface observation network, determine the most likely sign of the SAA error given the potential error sources above.

Recapping the source hybrid scan reflectivity product

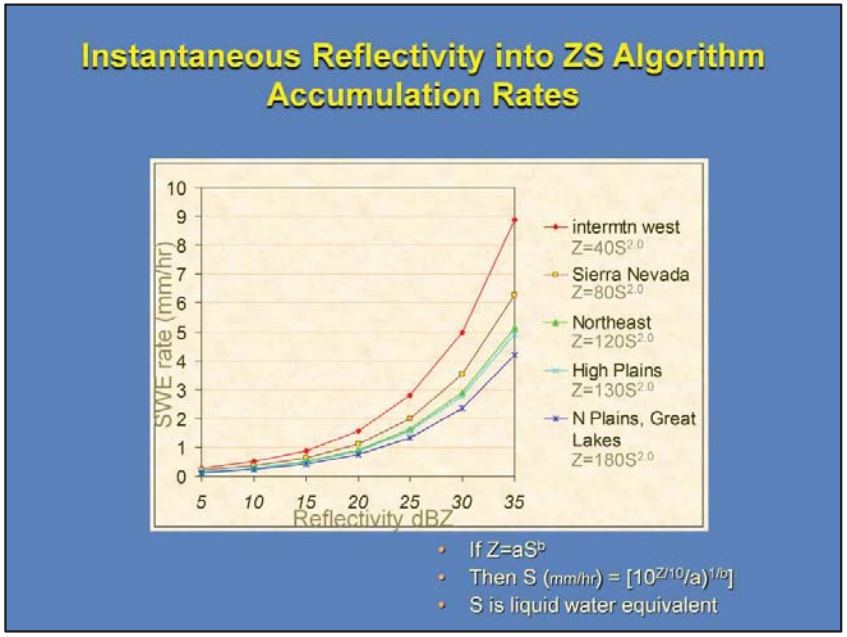
Lowest elevation scan reflectivity satisfying

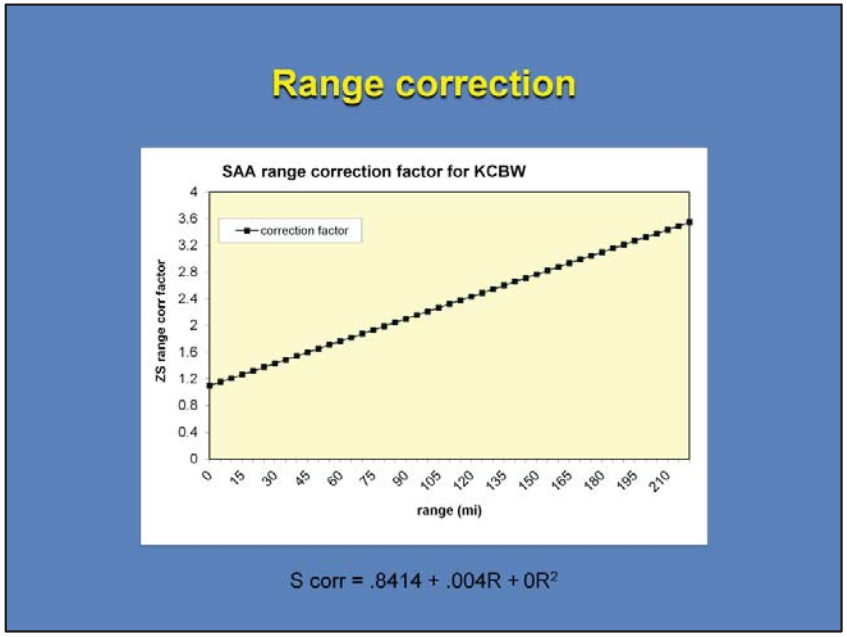
- Beam blockage ≤ 50%
- Outside an exclusion zone
- CLUTTHRESH ≤ 50%

As a review, the snow algorithm depends on the hybrid scan reflectivity product. Recall that this product is generated from the lowest elevation scan reflectivity that satisfies three criteria, less than or equal to 50% beam blockage, outside an exclusion zone, and the clutter threshold less than or equal to 50%.



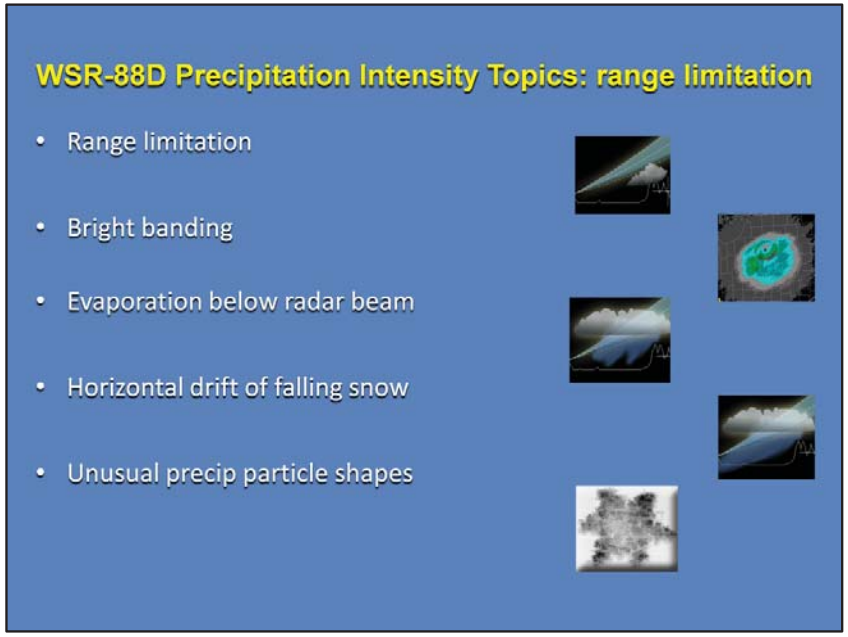
The ZS algorithm, was first deployed in the RPG in 2004. The Bureau of Land Management worked with the NWS to determine the most appropriate ZS algorithm for geographical regions. A representative office in each geographical region was the site of a one or more season's worth of high quality snow spotter data, where spotters not only sample snow depth but liquid equivalent too. After enough data has been collected, Super and Holroyd (1997) fixed the alpha and beta coefficients to one value that represents the minimum error between radar snowfall estimates and ground truth. Adjacent offices are also assigned these same coefficients based on the assumption that similar climatic conditions as the focus office exist, too.





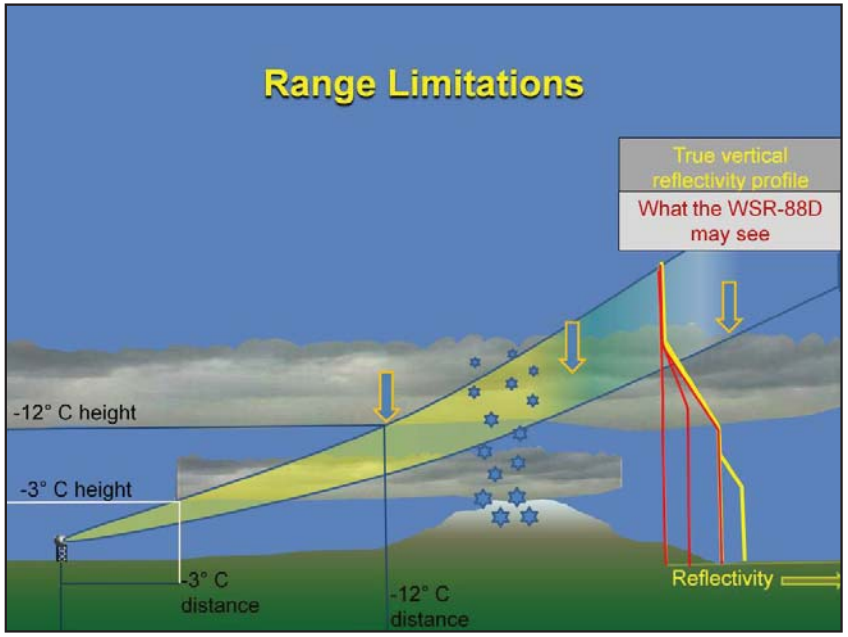
The default output of the ZS algorithm is liquid water equivalent rate in mm/hr. However, the ZS algorithm output will provide accumulated products of both liquid water equivalent and snow depth, both in English units. By the way, the snow depth will be derived using a fixed snow-to-liquid ratio. The ZS algorithm results show an exponential increase in snowfall rates with reflectivity for any region. The default coefficients are fixed over regions, however the weather is not. Variations in many factors can cause deviations from your default ZS algorithm settings.

The SAA applies a range correction to the estimates in order to compensate for the most obvious source of error of snow fall estimates. While the correction can use a second order polynomial, the last term is often neglected yielding a linear correction factor ranging from 1 at the radar and increasing to more than 3 at 200 mi. For all practical purposes, this correction is meaningless when the radar beam lifts above all precipitation echoes, which often occurs well short of the maximum theoretical range.



Five considerations adversely affect good precipitation estimates, and especially snowfall. Limited range becomes exacerbated in the winter when cold temperature microphysics routinely occur closer to ground. Bright banding becomes more of an issue as the freezing level is low enough to interfere with the hybrid scan selection. Evaporation below the lowest radar beam increases errors even at short range and magnifies as range increases.

Two more error sources include horizontal drift of falling snow and then unusual precipitation particle shapes. Both I'll explain later.

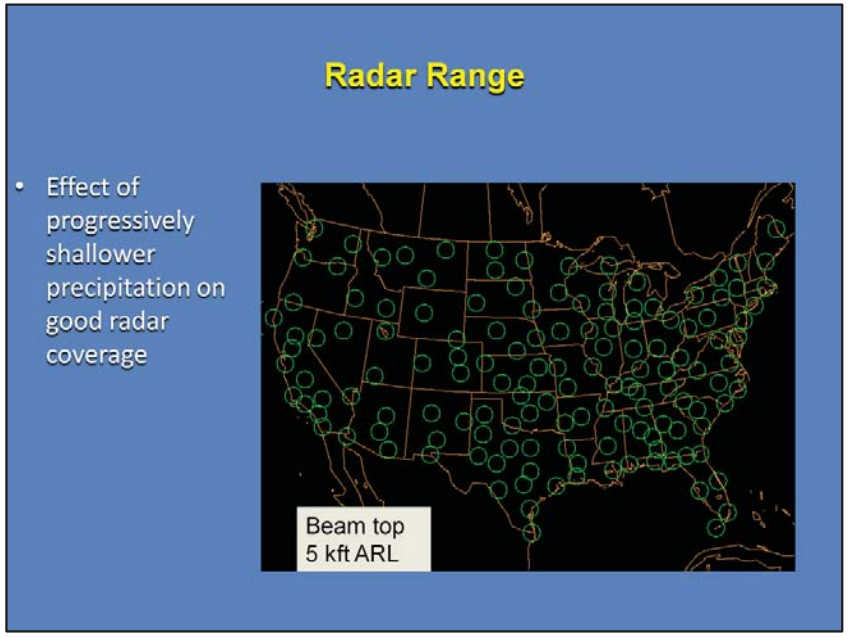


Assume we have a cloud pictured here actively generating precipitation so precipitation intensity increases going down through the cloud. Then if a radar was pointing up from underneath, the best vertical reflectivity profile shows increasing values as precipitation forms near cloud top and grows as the particles fall through the cloud followed by no growth as the precipitation falls out of the bottom of the cloud. Let's assume the reflectivity profile is the same everywhere under this cloud.

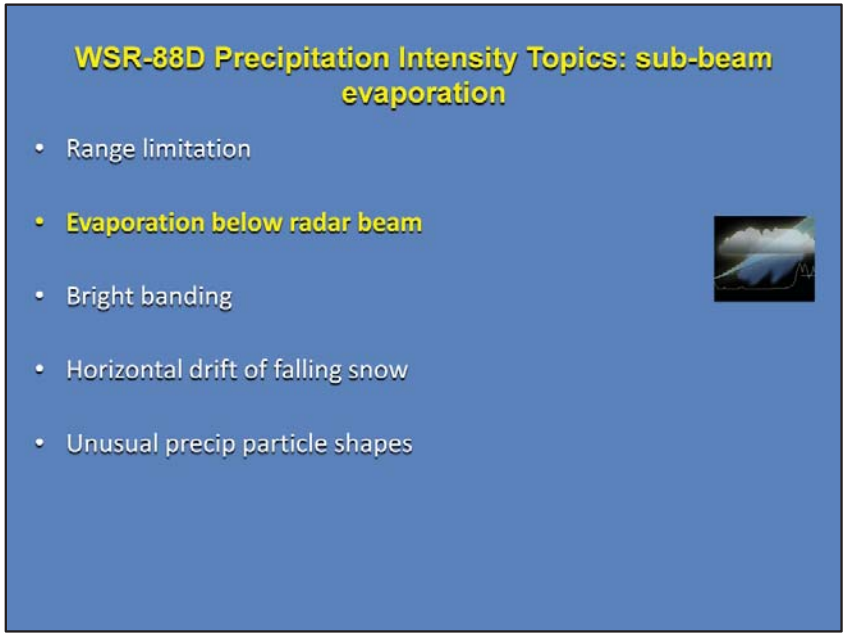
Now a WSR-88D from some distance away will detect all of the precipitation as long as the entire beam is below the precipitation production layer. Reflectivity begins to degrade once the top end of the radar beam climbs above the lowest part of the precipitation production layer because part of the beam is now sampling lower reflectivities. You're guaranteed to lose the signal once the bottom end of the beam departs the precipitation region.

The precipitation generation region is difficult to quantify. One definite zone is the maximum growth layer for dendrites (e.g., the -12 to -18 degrees C layer).

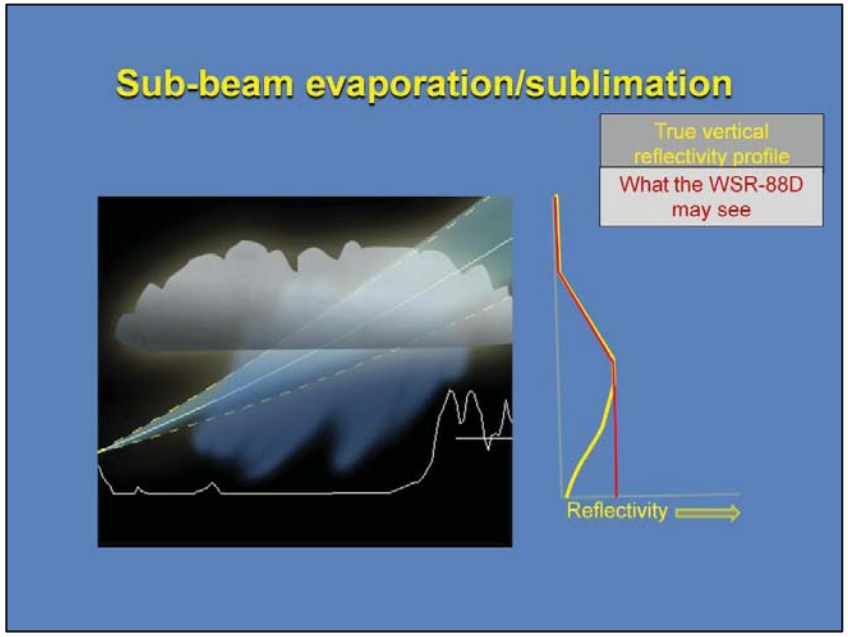
However, the presence of high cloud liquid water content in zones warmer than the dendritic growth layer but still below freezing can contribute significant amounts of riming and needles. In warmer saturated regimes, the collision-coalescence becomes active, too. Even the dendritic production layer can be fairly low in very cold weather, even near ground level. All of these precipitation production zones can be shallow and, therefore, cause reflectivity degradation at close ranges to the radar. As a side bar, orographic precipitation can occur very close to mountain sides. The WSR-88D has an exceedingly difficult time separating ground returns from real precipitation when both occur within a range gate. Clutter filtering can reduce or eliminate precipitation in range bins also containing clutter.



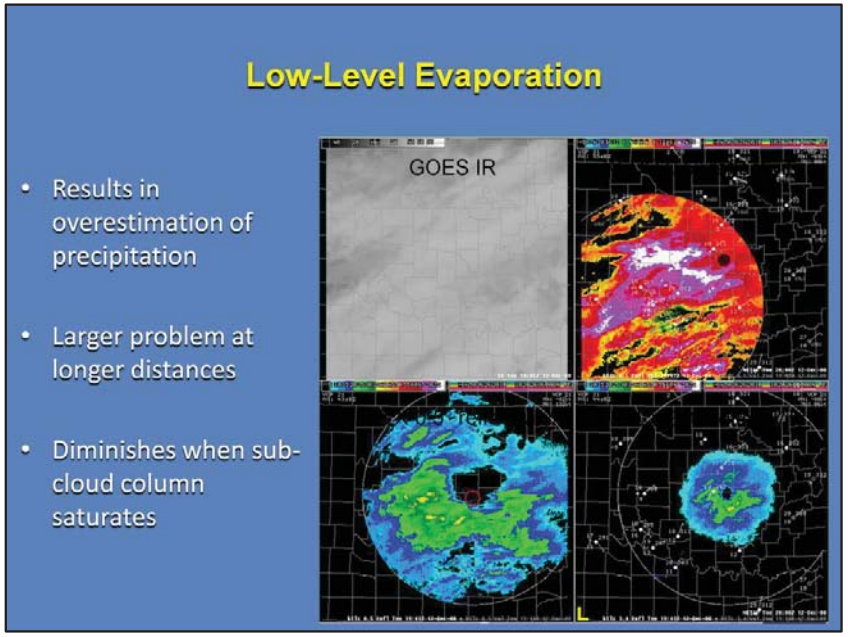
Here is the radar coverage for CONUS given that the dendrite production zone is relatively high, say > 22 kft. ARL. So for example, if there was a warm advection precipitation event where the dendrite production was the only layer producing precipitation, then expect good sampling by radar. Range limitations in this case are not a big problem. Now, the precipitating layer is lowering. Synoptic situations where this is common often occur in TROWALS, or along frontogenesis zones. Range degradation begins at a lower level and gaps in adequate coverage begins. Orographic clouds often hug the sides of mountains which can mean range is extremely limited. Also, in arctic outbreaks, shallow convection can result in significant snowfall rates, even with the precipitating layer at just 5 kft ARL. I have seen cases of power plant plumes and midwestern reservoirs producing significant snowfall whose cloud tops were only 1000'. Radar is extremely limited in its usefulness. Here is the coverage of typical lake effect snow, and precipitating orographic clouds where the precipitating layer is 8.5 kft ARL.



Let's talk about sub-beam evaporation of precipitation and where it affects radar precipitation estimates.



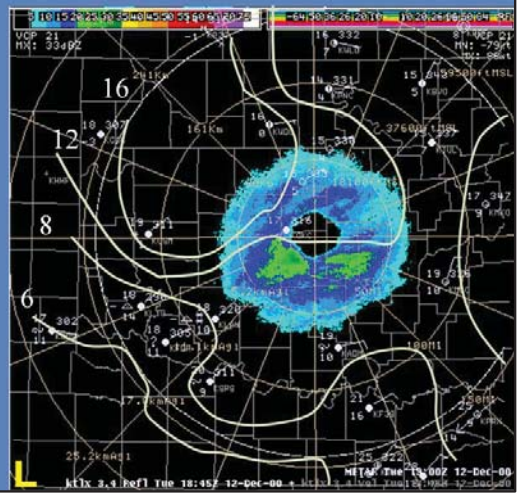
Sub-beam evaporation/sublimation presents the opposite problem for us as this process is often quite shallow and easily missed by radar resulting in precipitation overestimation.



The overestimation problem increases with increasing range from the source radar. The KFDR radar in the upper-right shows precipitation over an area where it is clearly virga, as viewed by the KTLX radar. If the environment saturates below the radar beam, this problem would diminish. The problem is if you don't have a nearby radar, how can you tell where the radar is overestimating precipitation rates. We'll look at reflectivity as a proxy for instantaneous precipitation rate and then look at some cases of hourly rates.

Low-Level Evaporation – dewpoint depressions

- Low-level dewpoint depression might give a clue to level of saturation
- However…



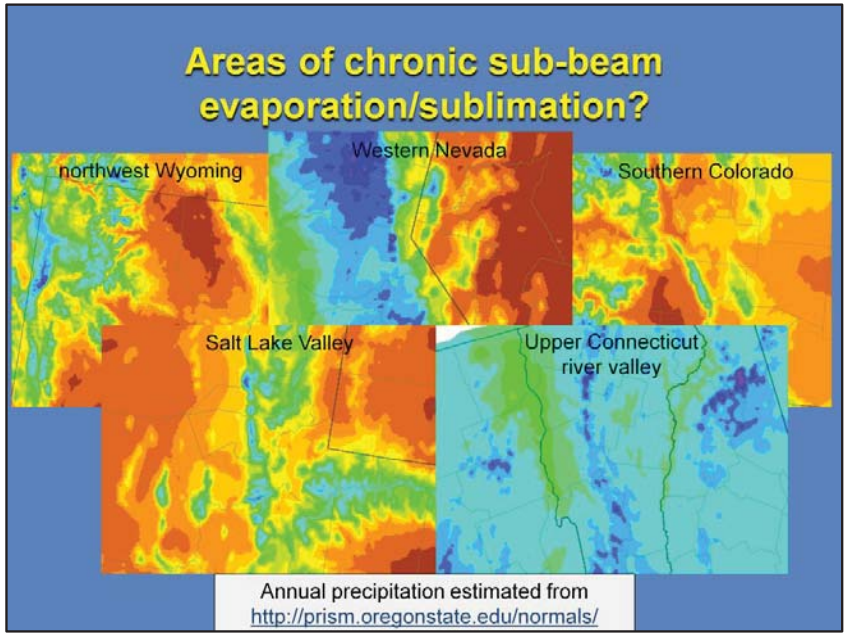
You can clue in on any subsaturated air by observing dewpoint depressions. Remember that these depressions are based on dewpoint, and not frost point. Thus if it's significantly below freezing, you probably will not see dewpoint depressions less than 5 degrees C.

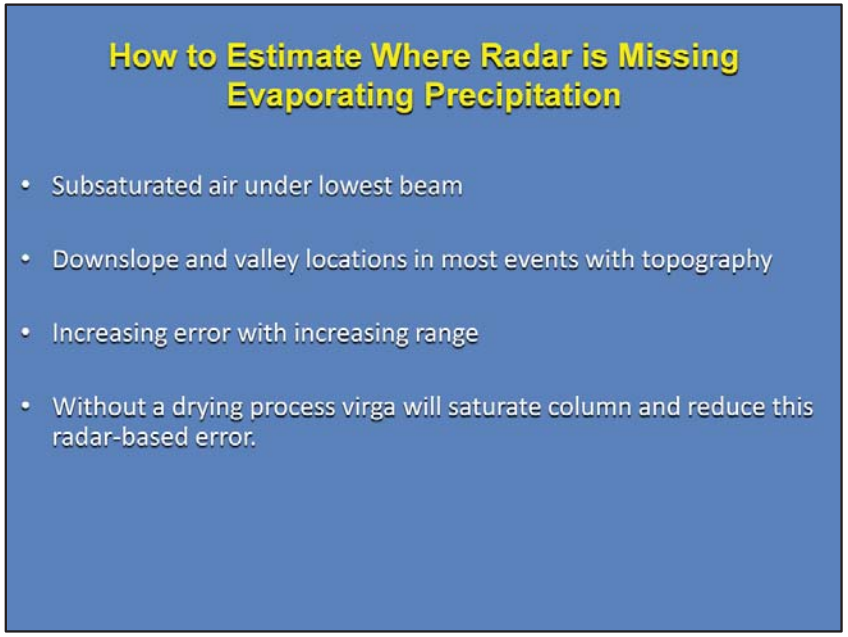
Low-Level Evaporation – elevated dry air

- Total evaporation is more dependent on integrated dry air exposure



Surface dewpoint depression in the last page was not entirely representative of the amount of total evaporation that a hydrometeor would experience on its way to the surface. It's also a function of the integrated dry air exposure a precipitation hydrometeor experiences along its downward path. We still know that precipitation rate will be overestimated by distant radars.





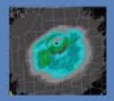
Those areas can quickly saturate as precipitation saturates the whole atmosphere.

Some geographical areas are quite likely to be subjected to chronic sub-beam evaporation, even with a radar nearby. Some areas listed above are examples, however I cannot possibly name all areas where this problem exists. Areas of downslope winds are primarily responsible for evaporating/sublimating precipitation. Since these areas change from one event to another, understanding your local climatology and how observed precipitation relates to your radar's estimates is of utmost importance.

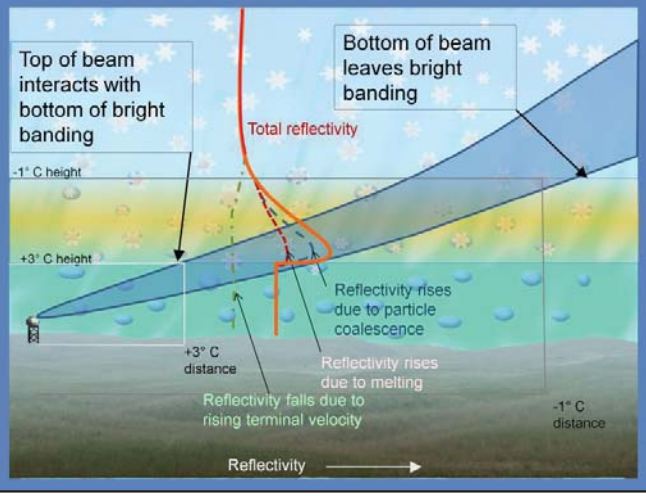
Anytime there's subsaturated air, precipitation evaporates. And as long as this is occurring underneath the radar beam, you're going to run into potential precipitation overestimates. Downslope and valley locations have the most likelihood for chronic sub-beam evaporation and such errors increase as range increases. Most areas subjected to hours of virga without any compensating process to dry the air will saturate fairly quickly and you'll watch your radar precipitation overestimates disappear fairly quickly from this error source.

WSR-88D Precip Intensity Topics: bright banding

- Range limitation
- Evaporation below radar beam
- **Bright banding**
- Horizontal drift of falling snow
- Unusual precip particle shapes



Bright Banding Reflectivity Factor



Five considerations adversely affect good precipitation estimates, especially snowfall. Limited range becomes exacerbated in the winter when cold temperature microphysics routinely occur closer to ground. Bright banding becomes more of an issue as the freezing level is low enough to interfere with the hybrid scan selection. Evaporation can also occur below the lowest radar beam.

Let's start with a stratiform precipitation region with a melting layer where the starting vertical profile of reflectivity in the snow appears as the red curve where reflectivity increases going to the right.

Breaking down the mechanisms behind the bright band, you see on radar because there is more than one. As snow flakes approach the melting layer, liquid resides longer on their ice surfaces before freezing. The increased water coating helps colliding ice particles to coalesce and snow flakes begin to increase in size. The larger particles increase the reflectivity.

The liquid water coating itself also helps to increase radar reflectivity because the dielectric constant increases as ice changes phase to liquid, especially when melting occurs.

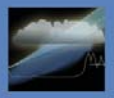
An offset to the increasing reflectivity occurs when the terminal velocity of the precipitation particles increases as melting accelerates. Increasing terminal velocity increases the separation between hydrometeors and lowers the reflectivity.

The combined result is an increase in reflectivity, maximized around +1 to +2 deg C.

Given that the center of a radar beam provides the most emitted energy, the strongest return from the bright-band effect is when the beam center is at these temperatures. But any part of the beam intersecting any part of the melting layer will also be affected.

WSR-88D Precipitation Intensity Topics: horizontal drift of falling snow

- Range limitation
- Evaporation below radar beam
- Bright banding
- **Horizontal drift of falling snow**
- Unusual precip particle shapes



We will discuss precipitation errors coming from horizontal drift of falling snow.

Horizontal displacement of falling precipitation?

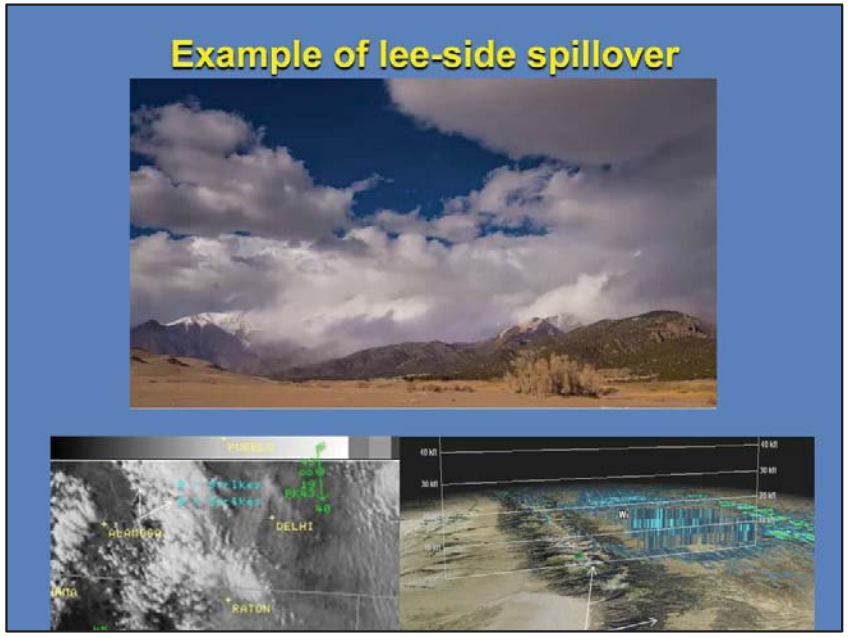




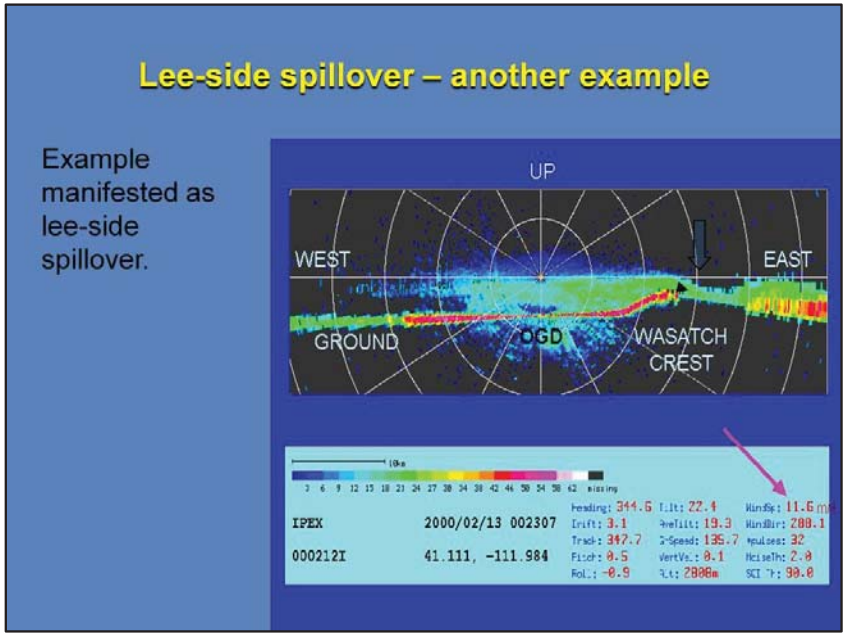
Not considered for today

Snow only falls at about 1 m/s, graupel, perhaps 3-4 m/s. The first case, the precipitation source is moving with the flow but the snow falling out of the base is subjected to horizontal displacement owing to vertical wind shear. Normally this may not be a problem for a radar beam to accurately locate estimates if the precipitation source is widespread. But perhaps some mesoscale snowbands would be displaced if the shear displaces the snow sideways under the beam and relative to the band.

The second case is where the precipitation source is anchored but horizontal winds displace the snow sideways such as with snow spilling over a mountain ridge from the orographic clouds on the windward side.



This time lapse from the Great Sand Dunes National Park shows a classic example of snow spilling over the Sangre De Cristo mountains into the San Luis valley as low-level upslope flow manages to make its way over the terrain into the lee-side. The time lapse is facing north with a field of view represented by the arrows on the satellite image and the reflectivity cross-section from the KPUX WSR-88D. While some of the upslope snow is visible from the KPUX WSR-88D east of the mountains, none of the lee-side spillover is visible to the radar. The spillover extends 2-3 nm downslope from the crest before sublimating.



What would lee-side spillover look like if you could see it from radar? illustrated nicely here in this cross section taken from the P3 aircraft tail radar just west of the Wasatch mountains. Snow forming in the upslope drifts down the lee side of the crest for several miles. In this part of the country, the 'spillover' effect is a blessing for ski areas on the east side of the crest.

In other parts of the country, this effect can occur if there is strong wind shear and rapidly moving transverse snowbands. If the location you're monitoring is far enough from the radar when one of these bands pass overhead, you may not have snow fall under the band for quite some time. Fortunately, this is one error source that can be mitigated by knowing the vertical wind profile and the height of the beam. The BLM attempted to implement a correction for horizontal drift of falling snow into the ZS algorithm. The results showed no significant improvement to the algorithm, possibly because other errors were so large.

Horizontal displacement of Falling Snow - recap

- High shear in synoptic systems
- Wind shear transverse to snow bands
- Over the crest of hills in topographic situations
- Other stationary sources of snow in high winds

- More important the further the event is from the radar

What you want to look for are situations where either you have a situation where rapidly moving bands transverse to the mean flow exist in regions of high vertical wind shear or a stationary snow production source embedded in strong winds such as orographics. The high shear forces the snow to horizontally drift relative to the source of the snow. Lake effect is a stationary source of snow by relative standards but the snow often falls within the axis of the band, just downwind. There are times though at the end of a lake effect band where snow production is nonexistent and all the snow that is falling is simply drifting there from upwind.

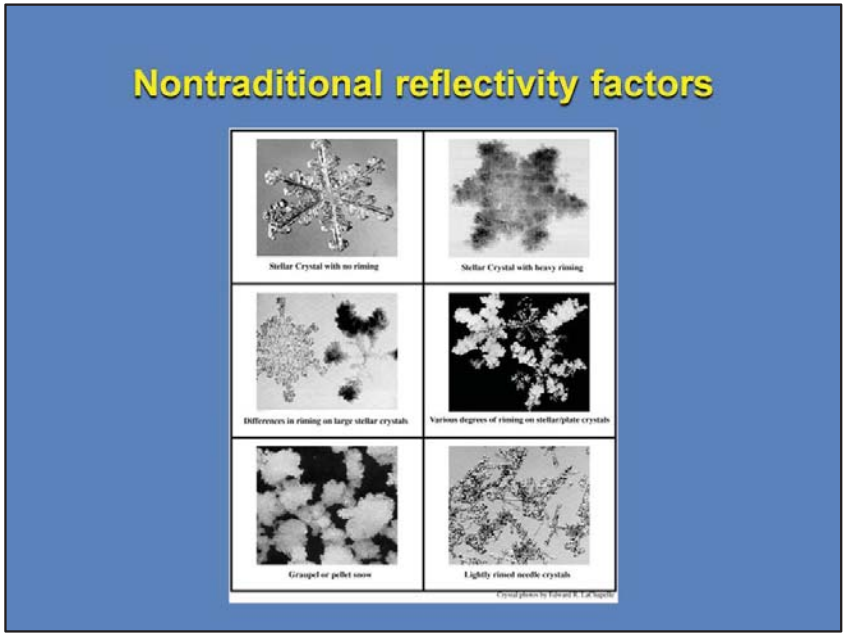
Finally, the further you are from the radar, the greater this source of error becomes.

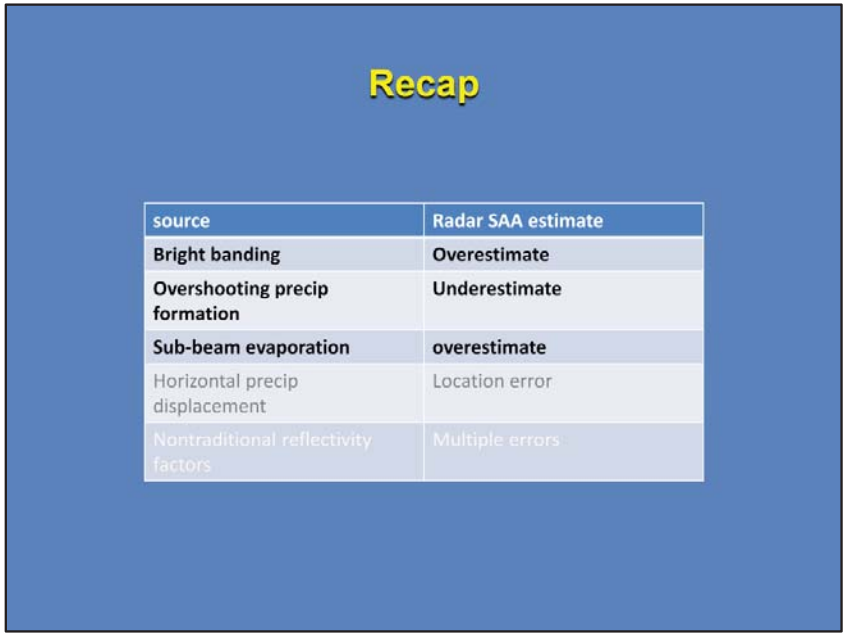
WSR-88D Precipitation Intensity Topics: precipitation particle shapes

- Range limitation
- Evaporation below radar beam
- Bright banding
- Horizontal drift of falling snow
- **Unusual precip particle shapes**



This final error source is probably the most intractable. Precipitation particle shape and size is something we'll talk about next.

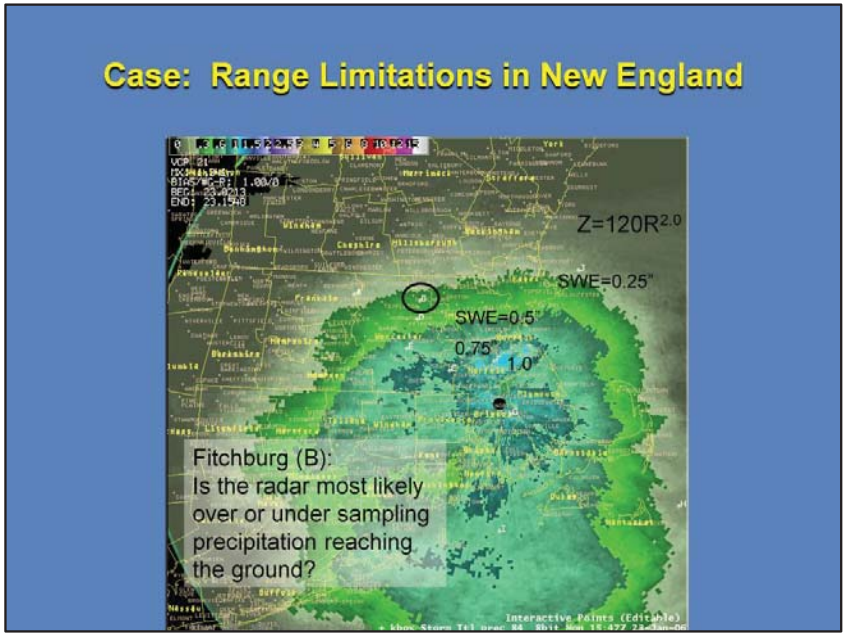


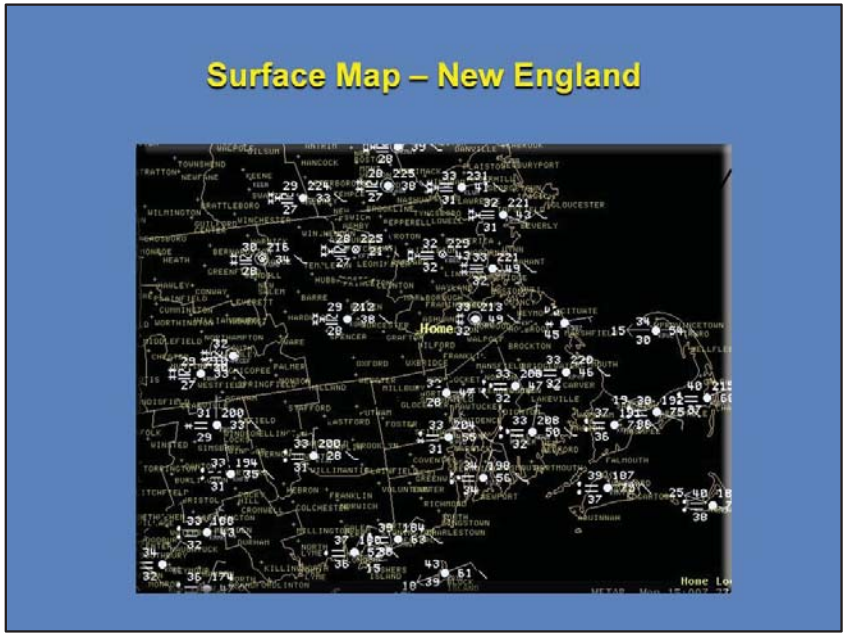


Snow particles come in all sizes, shapes, and densities, ranging from pure dendrites to needles, clusters of each, rimed clusters of each, and eventually, graupel. Given the same true precipitation rate, the diversity of precipitation shapes can cause changes in reflectivity. Some basic precipitation particle shapes can be gleaned from polarimetric radar, however at this time a polarimetric ZS algorithm is not imminent. We will also not consider this potential error sources in this course owing to the number of undetectable considerations that affect precipitation particle shapes beneath the lowest radar scan as well as the lack guidance available in dealing with the potential errors.

To recap, we have three sources of errors which can be assessed using operational data. Two of them induce precipitation overestimates, bright-banding and sub-beam evaporation/sublimation. Beam overshooting precipitating clouds is the major source of precipitation underestimations. Note that two of the three error sources magnify as range increases, however at a moderate range, errors from sub-beam evaporation can cancel out beam overshooting issues resulting in a fairly accurate ZS estimate – though for the wrong reason.

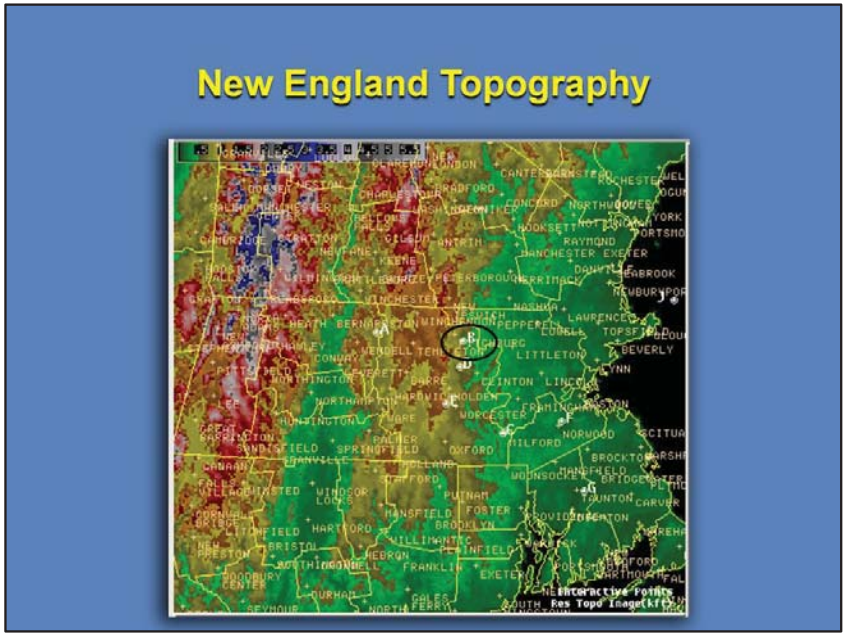
Horizontal snow displacement more or less affects the location of the precipitation reaching the ground relative to what the radar detects and can come from vertical wind shear below the precip source or precipitation falling laterally from an anchored precipitation source.

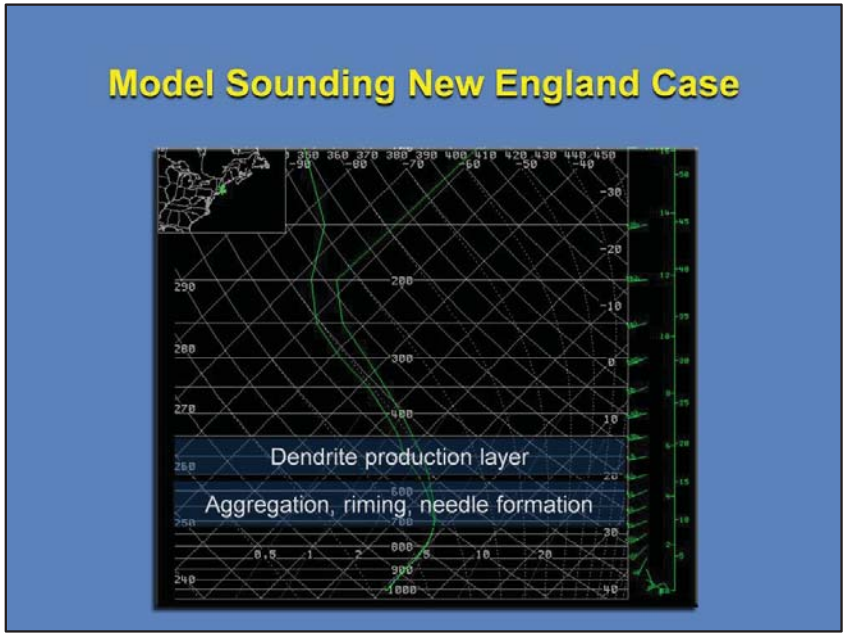




Let's try an example and take a look at how well the radar is estimating precipitation at Fitchburg, MA (in the circle). We ran the ZS algorithm using the default coefficients that exist for the Northeast U.S. and the minimum dBZ set to 10. There is no range correction applied here. Given the criteria I set, do you think the precipitation is likely to be over or underestimated for Fitchburg?

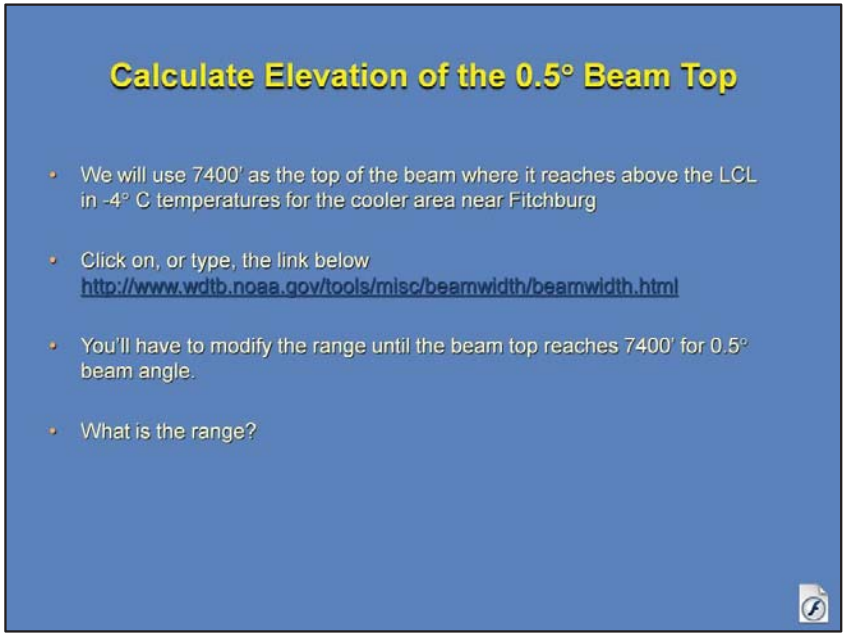
Fitchburg is well north of the rain snow line. The near saturation around the area also corresponds to relatively low cloud bases. Light winds in the area suggest that any bucket gauge should be able to be pretty efficient at capturing falling snow.

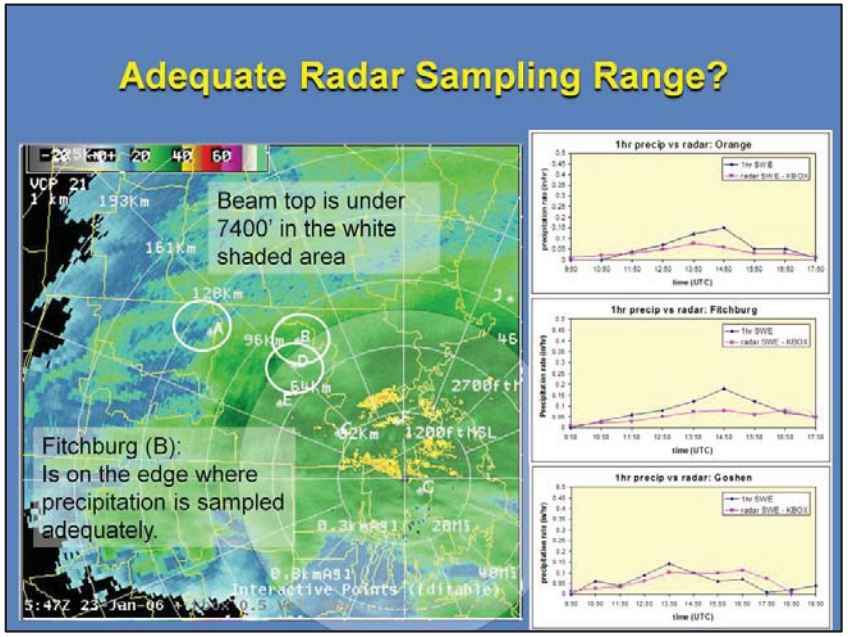




Fitchburg is in an area of relatively high terrain (600' MSL) and with the light easterly winds, there should be some upslope component to the flow. Upslope flow in subfreezing air means a greater possibility of low-level feeder clouds. Let's see what the morning sounding on Cape Cod shows. Let's see what you think about the sounding with the question coming up.

Let's assume that we need to have the lowest beam entirely below the dendrite production layer. Well, that's below about 560 mb layer or 15 kft. MSL. But it's not sufficient to be below just the dendrite production layer. We need to capture ice multiplication, aggregation, riming and needle formation. Most of these processes occur at temperatures colder than -4 degrees C. You'll need to have the beam top below 10 kft MSL to capture most of those processes. Given the deep saturated layer below, there may still be some additional aggregation as snow flakes become coated with thin films of water and get "sticky". So, reflectivity may go up a bit more. As a word of caution, thin films of water also increase the dielectric constant; therefore, reflectivity increases with no corresponding increase in actual liquid precipitation rate. Also, this sounding is not quite as far north as Fitchburg so the cold precipitation generation layer may be a little bit lower.





It is difficult to tell where in the sounding precipitation generation ceases and precipitation falls to the ground in its final form and intensity. We do know that below the dendrite production zone, frozen precipitation continues to develop. Even in warmer temperature you still have needle formation, riming and aggregation. We'll take temperatures warmer than -4 to -5 degrees C to be the point where most frozen precipitation growth will have already occurred. The soundings in the previous page show that temperature to be roughly 7400' MSL. We'll take that height and use the beamwidth tool.

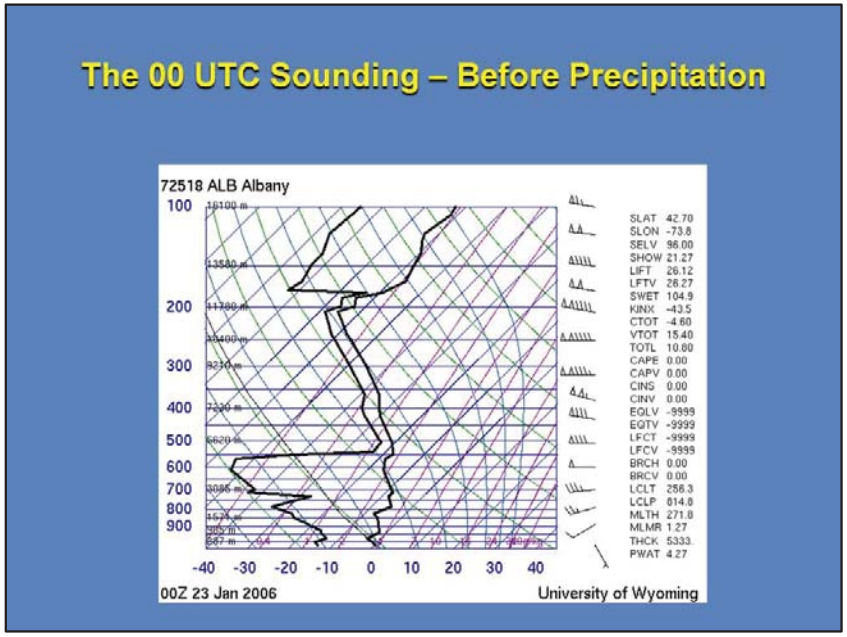
You're welcome to click on the beam width calculator below and modify the range until the beam top reaches 7400' for the 0.5 deg angle. What is the range that you find closest to 7400'?

It would be ideal if the beam was hugging the ground and there was no ground interference. Instead, it appears that perhaps we can adequately sample most precipitation if the entire beam remained below 7400' AGL. The white shaded circle satisfying this condition extends out about 53 nm from the radar. The Fitchburg observation is right on the outer edge of good radar sampling. Is 7400' enough or do we need to be lower? As it turns out, the Fitchburg ASOS is reporting higher hourly precipitation rates than the radar using the Z=120S2.0 relationship.

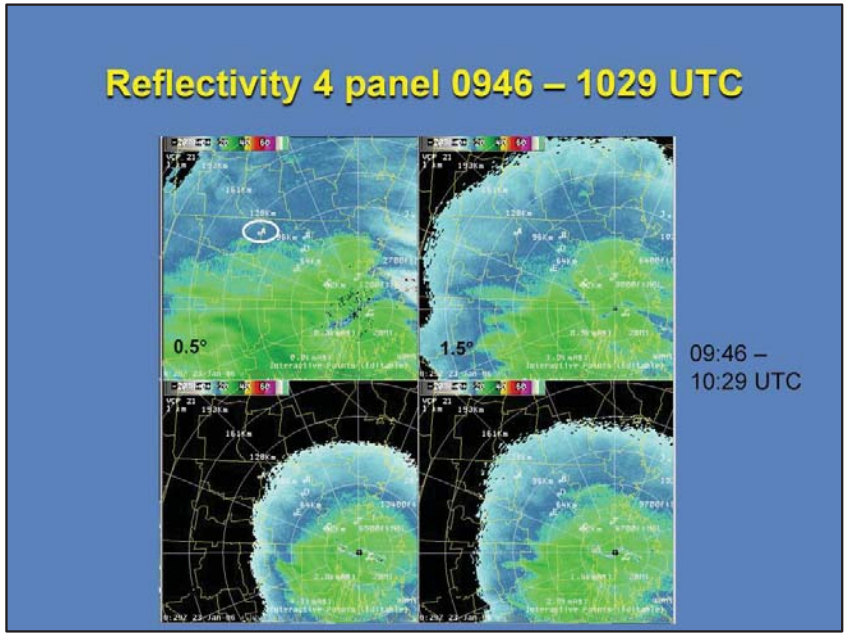
Going to the town of Orange, MA, which is further away, the same problem reveals itself. The radar is underestimating the hourly reports. Going a little closer to the town of Goshen, MA, the comparison is different. The radar is showing better agreement, perhaps even a bit of an overestimate.

If we believe the surface COOP station, then the radar beam is more accurately sampling the precipitation at Goshen than at Fitchburg. That is quite possibly because the beam is extending above some significant precipitation generation as one goes to Fitchburg and points beyond relative to the radar. I could make the argument that the COOP station is underestimating its precipitation and that would be a valid argument. However, I will show later in this lesson that there is nothing indicating that there is

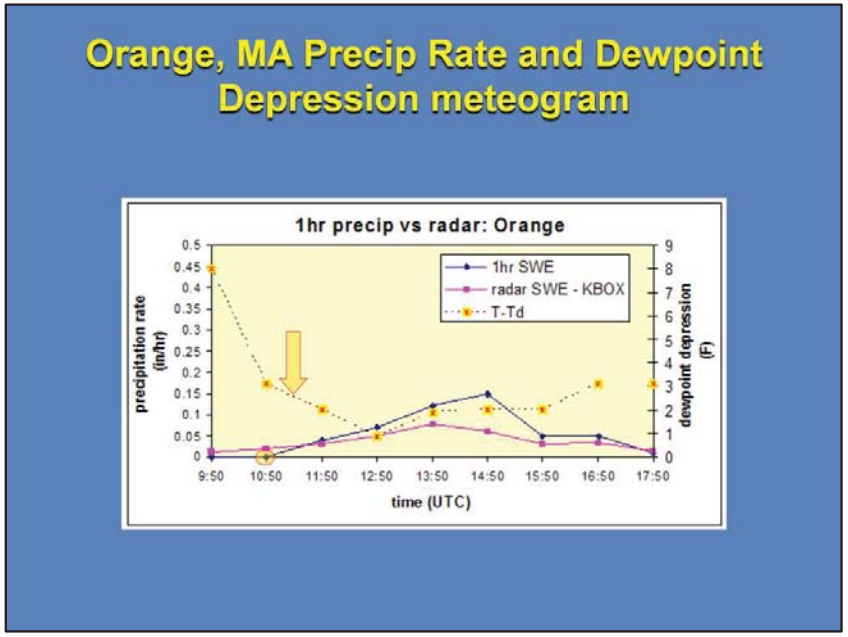
enough instrument error to change the conclusions that the radar beam over Fitchburg and beyond is overshooting generating precipitation.



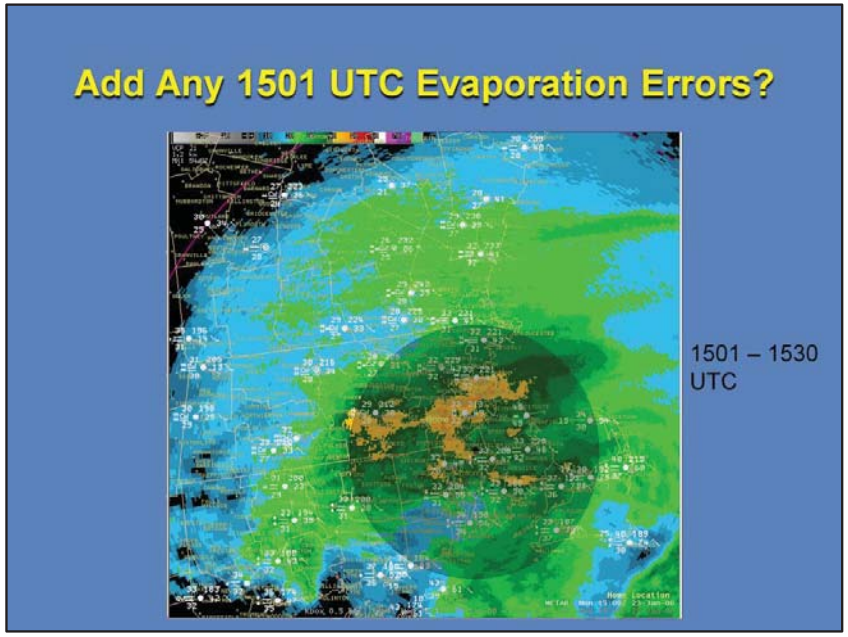
Now let's check out the sub-beam evaporation/sublimation potential. At 00 UTC on the 23rd, the Albany, NY sounding showed a very dry airmass below the midlevel moisture streaming in ahead of a short-wave trough. Let's take a look at the radar 4-panel image next.



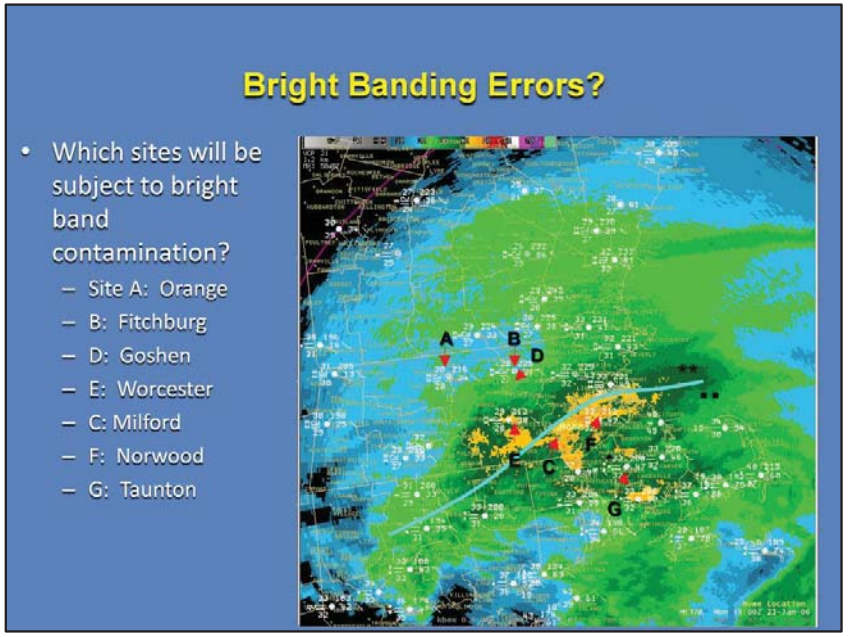
Take a look at this loop and given the previous sounding and the nature of the reflectivity echoes, determine the type of radar-based precipitation error you may observe here. Orange, in the white oval, is reporting a 9 degree dewpoint depression. I'll stop this from going on and allow you to come up with an answer before going on to the next slide.



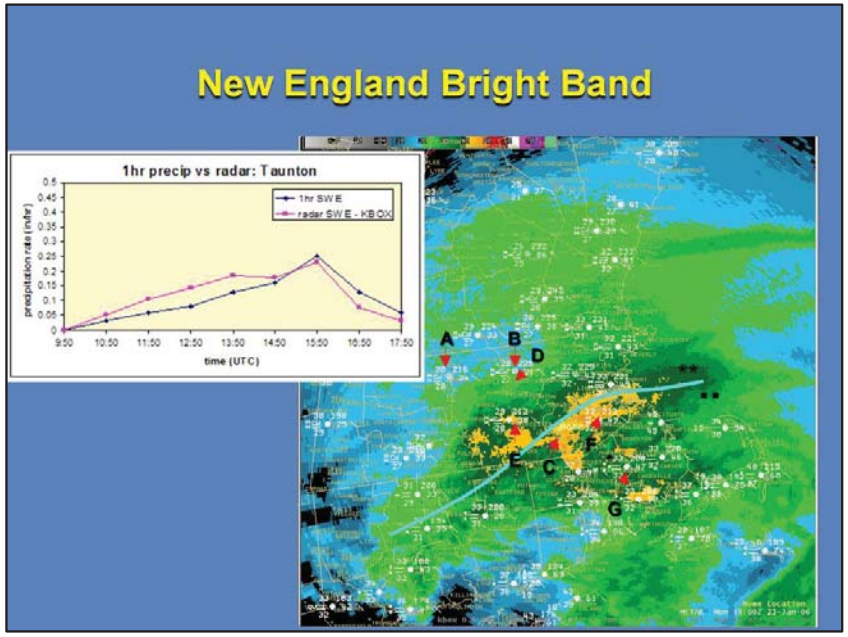
When you look at the dewpoint depressions, the dewpoint depressions correlate well with the propensity of the KBOX radar to overestimate precipitation rate at Orange. We're using Z=200S2, but we could raise the coefficient to 220 or 230 and I bet that wouldn't help because no precipitation is being recorded at Orange. Even if there was, note how the errors switch signs later on. You'd have to adapt by changing the ZS coefficients again. The error sign flipped as a large flux in moisture from above quickly saturates the air at Orange. I believe this is a common evolution for many sites that are located well away from the nearest radar.

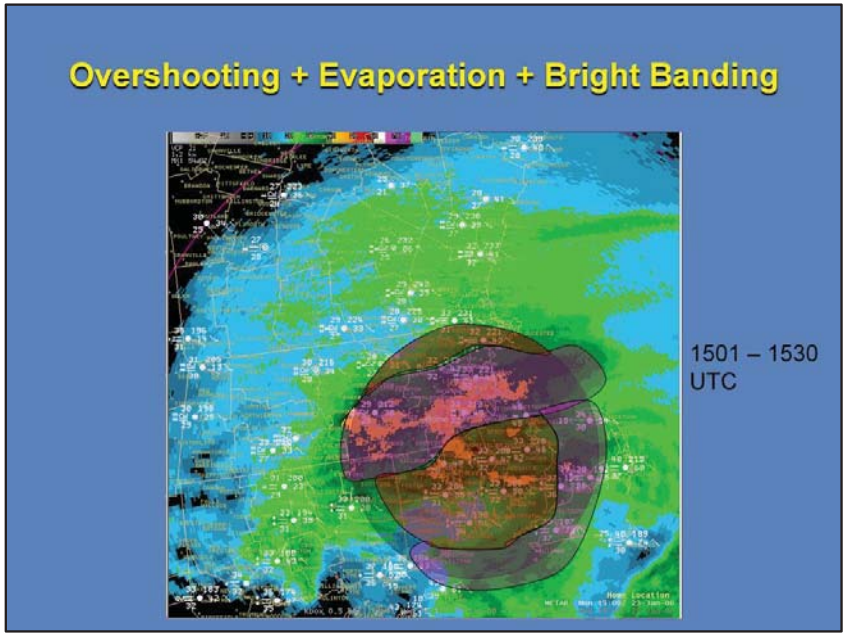


At this time, most everyone's saturated (at least with respect to ice). I doubt sub-beam evaporation is an issue now. Let's not include it in our considerations. So we still consider where the radar may overshoot precip outside the white shaded circle.



Let's try out your skill at locating bright band contamination. I have seven sites, some are ASOS, some are COOP sites. I am going to trust these sites as being relatively accurate. Understand that ASOS buckets are heated with poor shielding so some losses may occur due to evaporation and wind. Fortunately the wind is light and temperatures are fairly warm so it won't take much to melt snow into the bucket. The METARS tell me that the rain/snow line should be along the blue contour. The sites are, Orange, Fitchburg, Goshen, Worcester, Milford, Norwood, and Taunting. We'll go to a quiz.





Let's go through the answer to the last question. We start off with Orange (A) and see the problems with underestimation because of overshooting.

The same goes with Fitchburg (B).

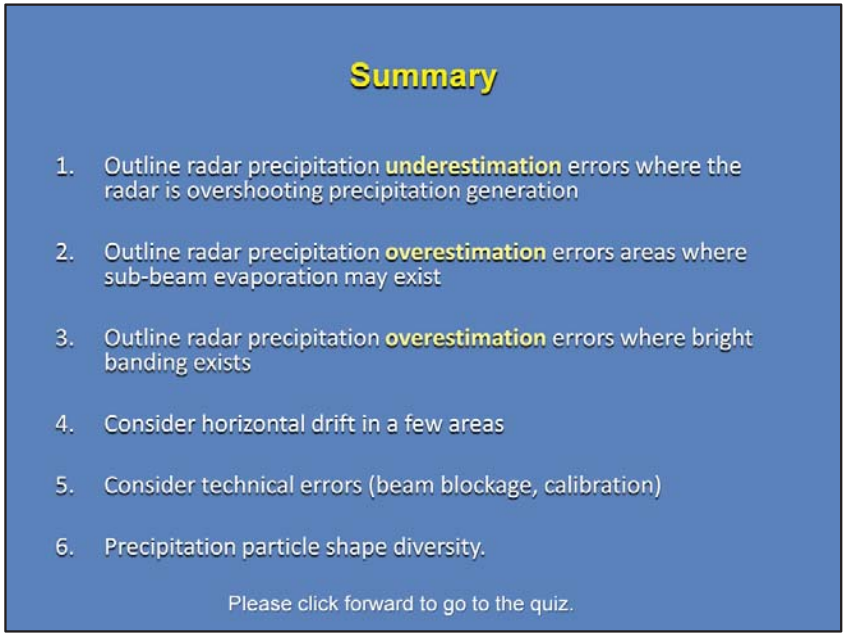
When we get to Goshen (D) we see the radar transition from underestimate to overestimation around 15 UTC. It could be either wet snow, sub-beam evaporation, change in precipitation particle shapes, or an error with the gauge. If we just talk about adequate precipitation coverage, it is doubtful there is sub-beam evaporation given the saturated conditions. I suggest that we are looking at wet snow aggregates with a larger than normal reflectivity factor than the liquid equivalent present. In other words, perhaps some bright banding is beginning even though it is all snow.

Going to Worcester (E), shows the same effect and here it is also snow.

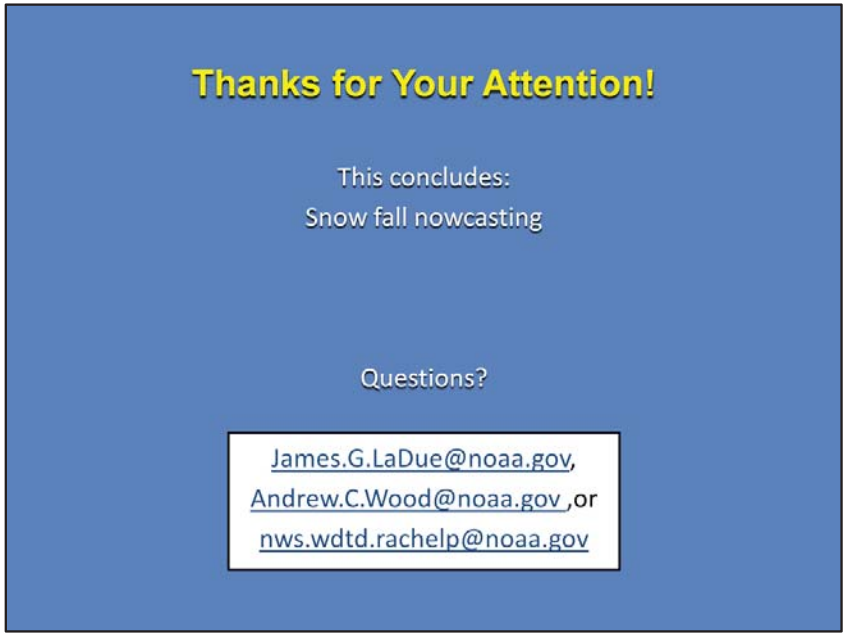
Going to Milford (C), and especially Norwood (F), we are looking at large waterlogged flakes with a huge reflectivity cross-section as the melting layer is fully sampled by KBOX.

Finally, the radar is sampling completely melted precipitation at Taunton (G) and there is better agreement there. Note, though, that earlier in the day there is some potential bright banding.

Beam overshooting problems exist outside the white shaded circle. There are likely no significant evaporation errors given the saturation across the radar domain. The most likely areas of bright banding appear in purple. The northern area is where the rain snow line exists at the surface but the errors leach into the wet snow areas to the north. Thus, the only areas of good precipitation sampling most likely exists in the orange areas. The northern orange area is the only area where snowfall is adequately sampled for the 15 UTC time frame.



Outline radar precipitation underestimation errors where the radar is overshooting precipitation generation. Outline radar precipitation overestimation errors areas where sub-beam evaporation may exist. Outline radar precipitation overestimation errors where bright banding exists. Consider horizontal drift in a few areas. Consider technical errors (beam blockage, calibration). Precipitation particle shape and density is the final error source that cannot be directly accounted for except when after considering the first 5 errors. If there is a consistent bias in the precipitation where sampling is good, you may consider the ZS algorithm coefficient to be in error because of precipitation particle shape.



If you have passed the quiz, then you have successfully completed this lesson. If you have any questions, please contact us using any of the e-mail addresses listed on the bottom of the slide.

TAB

Radar & Applications Course

Topic:

Convective Storm Structure and Evolution

Table of Contents
Topic: Convective Storm Structure and Evolution

Click to jump to lesson

Pre-requisite	**Hodograph Essentials for Convective Storms**
Lesson 1	**Fundamental Relationships Between Shear and Buoyancy on Convective Storm Structure and Type**
Lesson 2	**Ordinary Cell Convection**
Lesson 3	**Assessing Updraft Strength and Location**
Lesson 4	**Assessing Updraft Strength and Location with Polarimetric Radar Data**
Lesson 5	**Single Cell Downburst Detection**
Lesson 6	**Severe Hail**
Lesson 7	**Supercell Dynamics and Motion**
Lesson 8	**Supercell Morphology: Radar Reflectivity Signatures**
Lesson 9	**Supercell Velocity Signatures**
Lesson 10	**Supercell Morphology: Dual-Pol WSR-88D Signatures**
Lesson 11	**Supercell Archetypes**
Lesson 12	**Analyzing Tornadic Scale Signatures**
Lesson 13	**Tornado Hazards**
Lesson 14	**Multicell Archetypes**
Lesson 15	**Multicell Longevity and Severity**
Lesson 16	**Multicell Motion**
Lesson 17	**Rear-Inflow Jets in Multicells**
Lesson 18	**Line-End Vorticies and Bow Echoes**
Lesson 19	**Multicell Severe Wind Hazards**

Applied Performance Drills	
Drill 1	**Convective Initiation**
Drill 2	**Storm Type**
Drill 3	**Hazard Potential**
Drill 4	**Screen and Rank Storms**
Drill 5	**Convective Mode**
Drill 6	**Identifying the ZDR Column**
Drill 7	**Supercell Signatures**
Drill 8	**50 dBZ Height**
Drill 9	**Storm Top Divergence**
Drill 10	**Severe Hail**
Drill 11	**Measuring Rotational Velocity in Mesocyclones**
Drill 12	**The MARC Signature**
Drill 13	**Identify a TVS or TS**
Drill 14	**Tornado Debris Signature (TDS)**



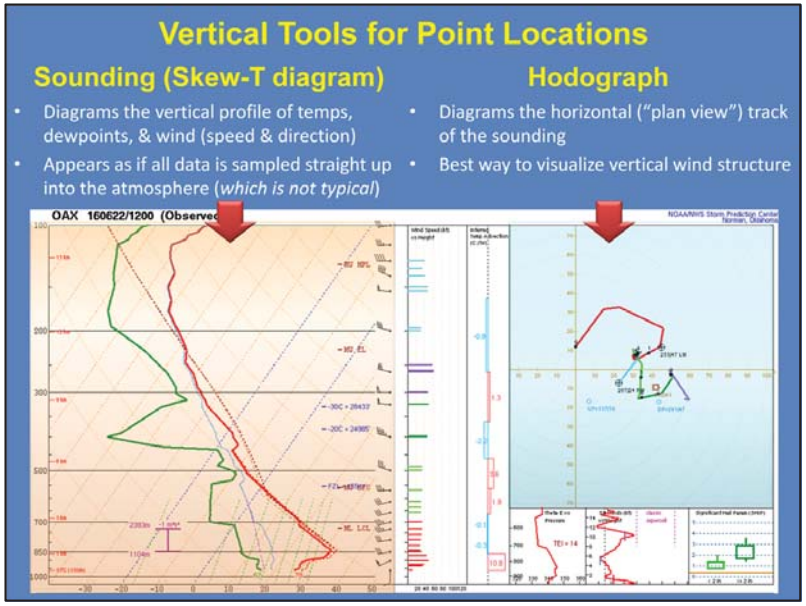
Welcome to the RAC Convective Storm Structure and Evolution lesson on Hodograph Essentials for Convective Storms. I'm Greg Schoor, WDTD Instructor. We will use this lesson to go over basics of hodographs and how to analyze them to get a better picture of the environment. This lesson is actually a prerequisite for the remainder of this topic, because it provides the foundation for the forecasting and understanding of different convective modes and the contribution of atmospheric shear. So, let's get started!

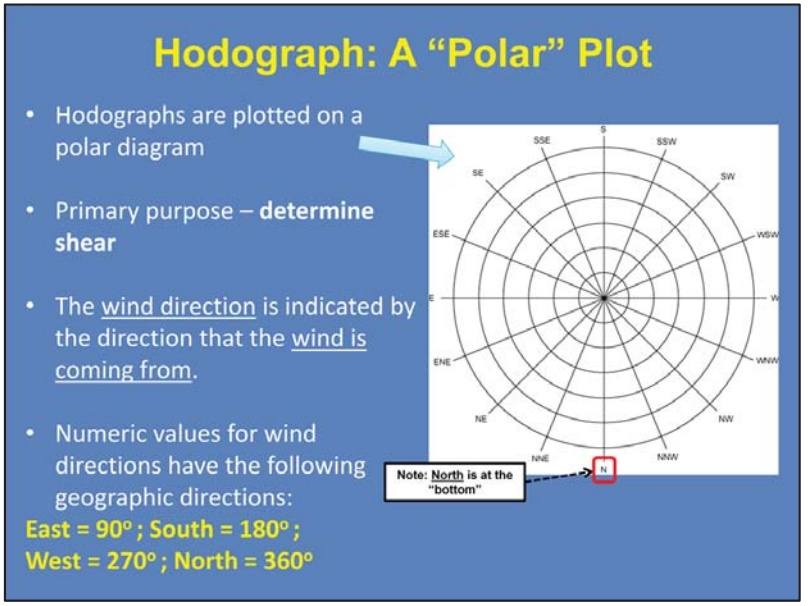
Learning Objectives

Given a list:

1) Identify the primary benefit of visualizing vertical winds on a hodograph versus a SKEW-T.
2) Identify the primary reason that hodographs are plotted on a polar diagram.
3) Define each of these types of shear parameters, vertical "bulk", total, and mean shear.
4) Identify the difference between ground-relative and storm-relative reference frame.
5) Identify the relevance of clockwise curved versus counter-clockwise curved hodograph layers.
6) Identify the difference between streamwise and crosswise vorticity.
7) Identify the two primary applications of using storm-relative helicity (SRH).

Here are the learning objectives. There are quite a few of them but this is a foundation lesson, so – a lot of important elements to it. Please take a moment to read through them – and refer back to them if need be.



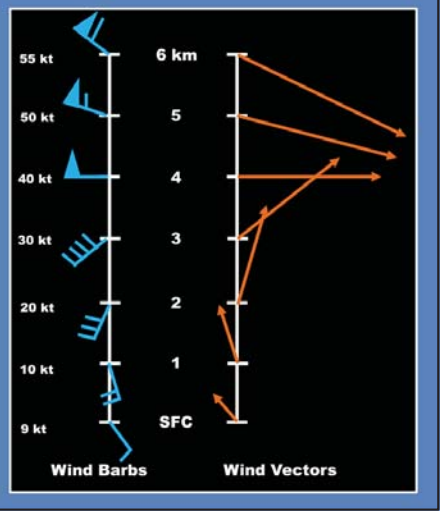


In terms of measuring the atmosphere, one radiosonde can serve two purposes – and really, both are necessary for proper analysis and forecasting of the atmosphere surrounding a point location. First, the sounding – on the left – also known as the Skew-T diagram, depicts the vertical temperature, dewpoint, and wind profile. But the way it is plotted on this chart, makes it seem like this information is being taken straight and directly up from the point where the weather balloon was released. In a controlled environment, with absolutely no vertical shear, that may be true, but realistically, it is not. But what we will focus on is the hodograph – on the right – which shows us in a horizontal, or "plan view" sense, where the radiosonde has travelled, start to finish. This gives the best visualization of the vertical wind structure and goes a long way to helping predict potential convective modes.

With Skew-T diagrams, three different elements have to be plotted on the same chart, but with hodographs, we are only concerned with wind. Plotting wind on a polar diagram allows us to determine various types of shear. Recall that shear is the change in wind speed and direction between 2 levels or with height, so we're talking about the gradient with the term, "shear". Also recall that wind direction is indicated by the direction that the wind is coming from, so if it's a "south wind", that means it is coming from the south. The hodograph's polar coordinates appear upside-down from what you're probably used to, with north being on the bottom, south on the top – for the Northern Hemisphere.

Winds: From Sounding to Hodograph

- Converts the wind barbs from the Skew-T profile into vectors
- The vectors point in the direction of the wind
- Vector length represents the magnitude (speed) of the wind



Next, we have to go back to the Skew-T to grab the wind information before it can be plotted onto a polar diagram. Wind barbs, in blue, look like they would on a Skew-T diagram, showing the magnitude – or speed – with either the tick marks or the flag symbols. The direction at each level goes from the tip of the barb to where the stem of the symbol meets the vertical line. For a hodograph, these values are converted into vectors, orange arrows on the right, in this example. Same information as the wind barbs, but these vectors now point in the direction the wind is going, instead of where it's coming from and they get longer with higher magnitudes. But what are the magnitudes? What does a 20 knot wind barb look like in vector form on a hodograph?

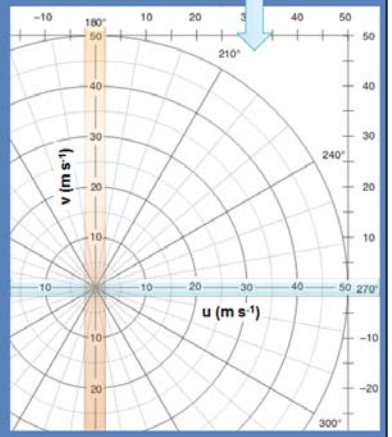
Hodograph: Magnitude Labels

- In the Northern Hemisphere, this quadrant is the most typical

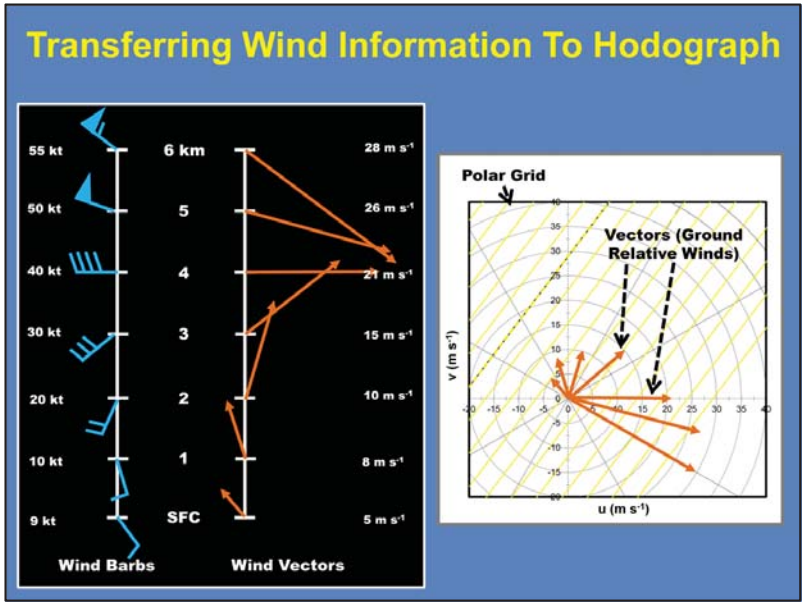
When plotting wind vectors…

- The wind vector is in knots (kts)
- The X & Y Axis are in meters per second ($m\ s^{-1}$)
- The rough conversion is: $m\ s^{-1} = ½(kts)$

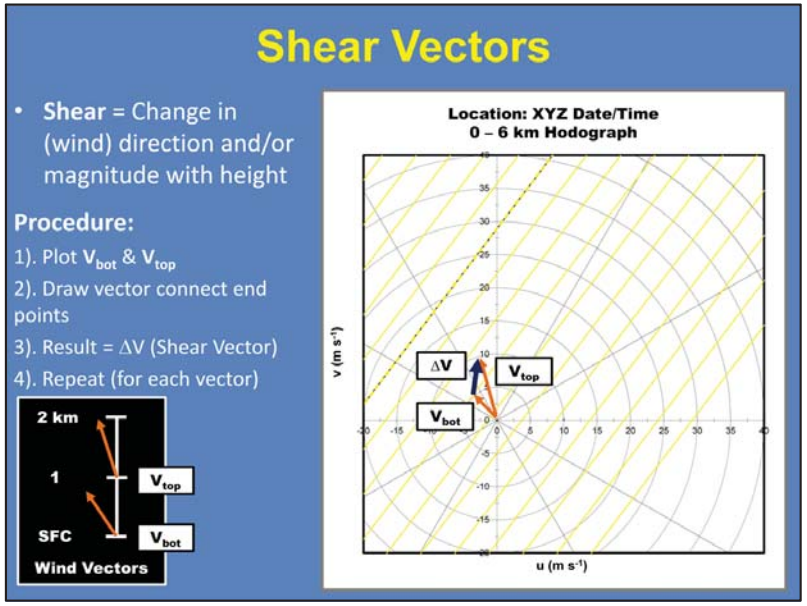
Examples:
10 kts = ~5 $m\ s^{-1}$ (5.144$\overline{4}$ $m\ s^{-1}$)
20 kts = ~10 $m\ s^{-1}$ (10.2889 $m\ s^{-1}$)
50 kts = ~25 $m\ s^{-1}$ (25.722$\overline{2}$ $m\ s^{-1}$)



A couple of slides ago, we introduced the polar diagram. So what about the X and Y axis and why are the units from 0 to 50? On a sounding, the wind information is in knots but on a hodograph, we convert that to meters per second. Fortunately, the conversion from knots to meters per second is just dividing the knots by 2…roughly. It's not an exact conversion but very close as you can see from the examples, all the way from 10 knots to 50 knots, the conversion to meters per second is just about half for each one.



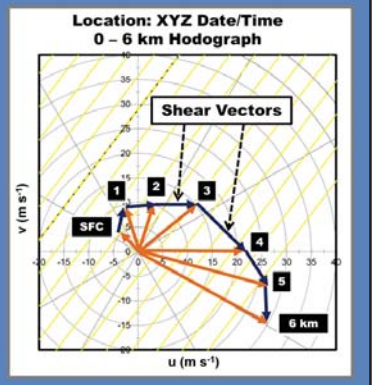
Now that we have converted wind barbs into vectors, we'll transfer them onto a polar diagram. So, we basically just pluck them at each level from the diagram on the left, taking the vector's origin and placing it on the center point of the hodograph, at 0, 0 – as you can see from this animation. Each vector points in the horizontal direction the wind is headed at that level. Next, we are going to dive into the concept of shear.



The foundation of the hodograph is this concept of shear…so what is it? Shear – and if we're talking about wind, it's wind shear – is the change in direction and/or speed, with height. From here on out, we're talking about the difference in the wind at two levels, every step of the way. Obviously, hodographs, like Skew-T's, are automatically generated but it is important to know how they are generated – so you can see where the various severe convective parameters come from. So, we'll start at the bottom, with the lowest two wind vectors from the earlier example. Plot them on the diagram and then place a connecting vector at the tips of the wind vectors, pointing the direction from bottom to top of the layer. The result is the Shear Vector, if it was just a line without an arrow at the end, we'd call it a Shear Segment – but a vector has direction. Then, you do this step for the remainder of your wind vectors through the column.

Tracing the Hodograph

- Once the vectors (ground-relative winds) are plotted – the "**Shear Vectors**" connect the end points
- When complete, this trace is the "hodograph"
  - Then the shear vectors disappear
  - Also, notice that the "hodograph" (blue line) doesn't start at the origin (unless it's calm at the surface)

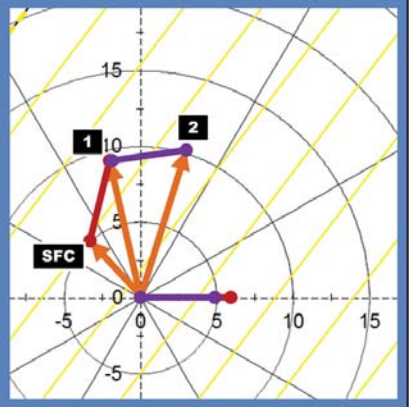


Now we have taken all of the wind vectors, plotted them, and generated Shear Vectors for each level. Remember, with this example, we had a vector at each 1 kilometer up, from the surface to 6 km. When all of the Shear Vectors are plotted through the column, that trace is referred to *as* the hodograph, the trace of all your wind shear values from the surface to the top of the sounding. Once this is plotted, then you can ignore the shear vectors. Also, you may be wondering why the hodograph line, which is the result of the blue arrows, doesn't start at the origin, or 0,0 on the chart. That would only happen if you have no wind at the surface and therefore your lowest point is a zero vector. Otherwise, you will see the line starting at some point away from the origin, depending on the strength and direction of the wind at that lowest elevation.

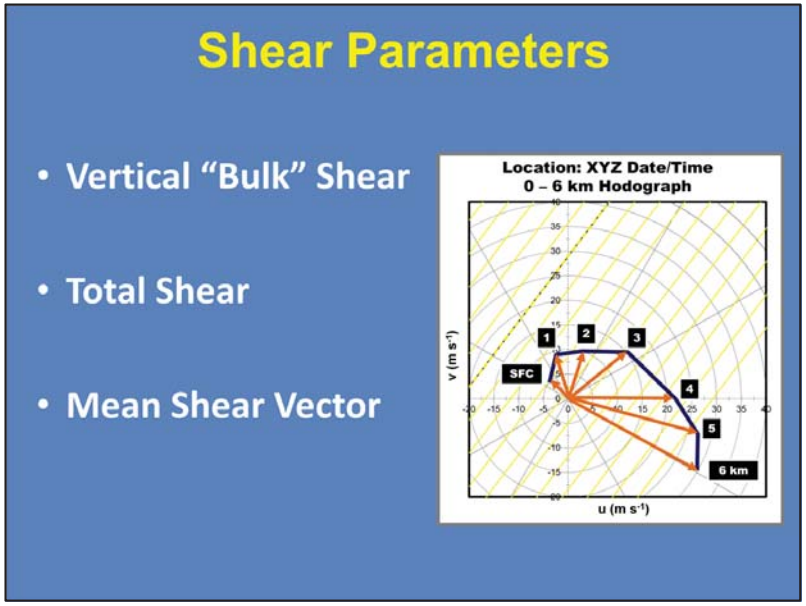
Shear Segments

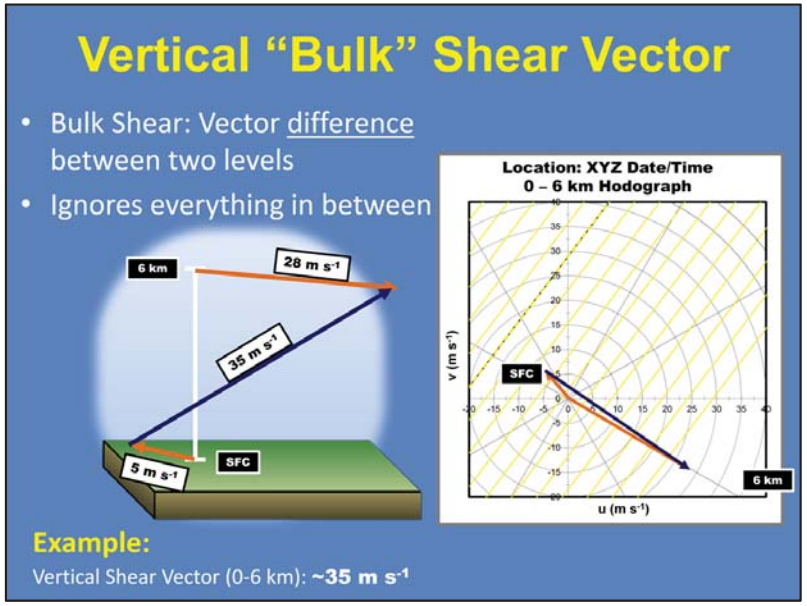
- **Shear Segment = Magnitude (length) between two levels**
  - Take first segment down to (0, 0) then determine x-axis length (estimate)
  - Same all additional segments
  - Then, add them all end-to-end for the "total" segment length

Segment	Length
SFC → 1	6 m s^{-1}
1 → 2	5
Total (SFC → 2)	11



The final important component of plotting the hodograph is determining the magnitude – or the length – of each shear segment. How do we do that? In this coordinate system, we can certainly complicate these calculations very quickly. However, we can actually just pluck the segments and transfer them down to either the X or Y axis and then estimate the length on that axis. In this example, we'll use the X-axis, starting with the surface to 1 km segment, which is the red line. Plopping that onto the X-axis, we come up with 6 m/s for the first segment and then 5 m/s for the 1 km to 2 km layer. Now, to add them up…we just add them up, end to end, which is 11 m/s. This is the total shear amount for the surface to 2 km layer. So now, let's build on that.



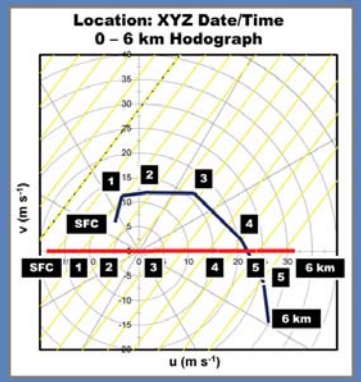


Once the hodograph is drawn, the real work starts because we need to make sense out of this trace. What does this shape mean in terms of the environmental wind and this time? Does this profile help or hinder convection? If we do have convection, what will become of it, what mode will it be? We can begin to answer these and other questions through variations of this concept of shear – which again, is what the hodograph is plotting, the differences in wind speed and direction with height. So, we're going to explore these parameters, Vertical or Bulk Shear, Total Shear, and the Mean Shear Vector listed in the next few slides.

First, we can determine the Vertical or Bulk Shear Vector which is just the vector difference between two given horizontal levels. In our example, we subtract the vectors from the surface to 6 km. This can be done for any two levels, 0 to 1, 3 to 6, and on and on. This is vector subtraction though and it is not as easy and just subtracting 28 minus 5 to get 35 m/s. So, where did 35 m/s come from then? We will find out how to determine magnitudes in the coming slides.

Total Shear

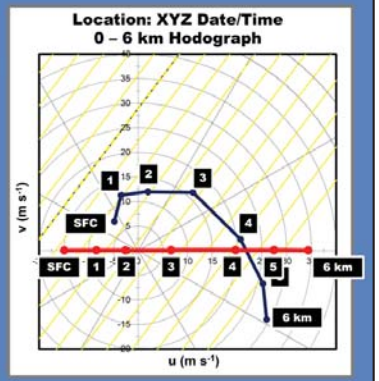
- **Total Shear**
  - "Net Length" of the hodograph
  - How is it derived?
  - Place the segments on the X-axis & add them up (the magnitudes)



Total Shear

- **Total Shear**
  - Determine the length of all segments
  - Put them end-to-end on the x (or y) axis

Segment	Length
SFC → 1	6 m s^{-1}
1 → 2	5
2 → 3	10
3 → 4	14
4 → 5	8
5 → 6	7
Total (SFC → 6)	50 m s^{-1}



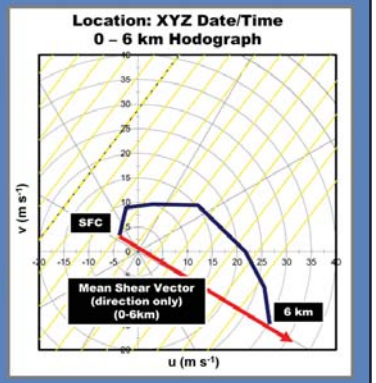
Remember how we came up with shear segment lengths? Now we're going to build on that and figure out the Total Shear for this entire plot. If you have a hodograph with sounding data that goes all the way to the troposphere, you will probably have data well past 6 km, maybe even between 10 and 15 km. Generally speaking though, most of the convective processes in terms of shear are found below 6 km, so that's where we will focus on. But if you do truncate levels higher than 6 km, then you have to denote this as the Total Shear from 0 to 6 km, for instance.

Since we've turned these vectors into segments, which is just magnitude, independent of direction, we can add them up end-to-end on the X-axis and get the total shear for that layer. On the next slide, we'll put them together.

Going through the same exercise as we did a few slides ago to determine the shear segment lengths, I've plotted the remainder of the magnitudes – or lengths – of each of the segments from 0 to 6 km from the same example. You can see the lengths there in the table. Then, we plop them all onto the x-axis, end-to-end, and normally, I would have started the end-to-end process at (0, 0) but as you can see with the total length of 50 m/s, the right side wouldn't have fit, but as you can see, on the X-axis it starts at -15 and goes to +35, which the absolute value is 50, so it works either way. So then, this is our Total Shear from 0 to 6 km for this example, a whopping 50 m/s.

Mean Shear Vector (Direction)

- **Mean Shear Vector can help determine supercell motion**
  - Total Shear is easier to deduce (just add up the magnitudes through a specified layer depth)
  - Mean Shear is a bit more labor intensive
  - (easy) First, determine ***JUST*** the Mean Shear vector **DIRECTION** by drawing a vector from beginning point to the end point (of the defined layer)



Example:

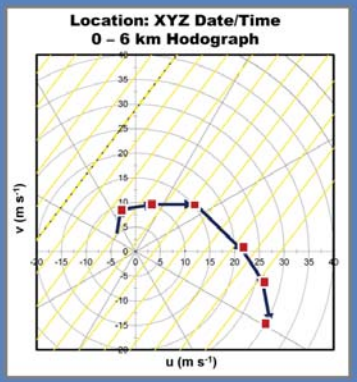
[red vector] Points in the direction ESE but is "West-northwesterly" (~300°)

Building on total shear, we delve into this idea of the Mean Shear Vector, which can help determine supercell motion – and again, this 0 – 6 km layer is a typical one to use, so we'll continue with this example. But the mean shear vector, hand-calculated, is not nearly as easy as finding the total shear – which we just took the shear segments end-to-end and added them up. Luckily, the direction of the Mean Shear Vector is pretty easy – just draw a line from the point where the SFC segment starts and connect it to the end of the 6 km segment and that direction this is your Mean Shear Vector for that layer. In this case, it's pointing toward around 300 degrees, which is West-Northwesterly, which means it is coming out of the West-Northwest but it is pointing in the direction of East-Southeast. Now, we have the task of determining its magnitude.

Mean Shear Vector (Magnitude): Part 1

- Convert segment points into the (x, y) coordinate system
- Treat each segment individually
- Act as if each segment's **starting** point is (0, 0) then find the (x, y) of the **end** point

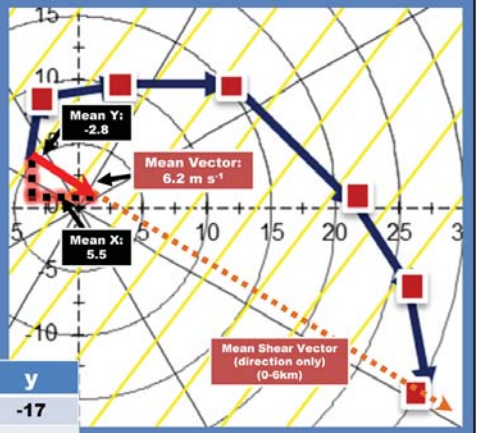
Segment	x	y
SFC → 1	+1	+6
1 → 2	+6	+1
2 → 3	+10	0
3 → 4	+10	-10
4 → 5	+5	-8
5 → 6	+1	-6
Total	+33 m s-1	-17 m s-1



Now, on to a bit more labor-intensive but critical step of finding Mean Shear magnitude. For each segment, we need to get an X and Y coordinate for the END point of that segment. How do you do that? By treating the START point as (0, 0) for each segment. Then, find the x-axis extent of that segment… for the first segment from 0 to 1 km, it only goes forward 1 in the positive x-direction, but it goes 6 in the positive y-direction. The next segment is 6 in the positive x-direction, then only up 1 in the positive y-direction… and so on until we reach 6 km. You see most of the y-directions being negative for these segments because we're going the opposite way on the chart from positive y. Add them all up for X and Y and you have 33 m/s total shear in the x-direction and -17 in the y-direction.

Mean Shear Vector (Magnitude): Part 2

- Divide the X and Y sides by 6 km
- Result is the **Mean Shear** (x, y)
- Plot this (x, y) back on the hodograph, starting at the SFC (0 km) point
- **Mean Shear: (5.5, -2.8)**
- **Vector: 6.2 m s^{-1} (ESE)**

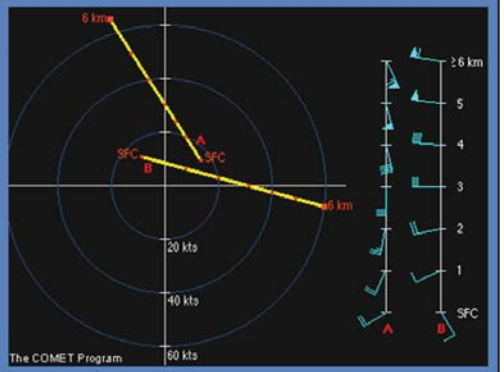


Segment	x	y
Total (0 → 6 km)	+33	-17
(divide each by 6 km)		
Mean / Avg.	+5.5 m s^{-1}	-2.8 m s^{-1}

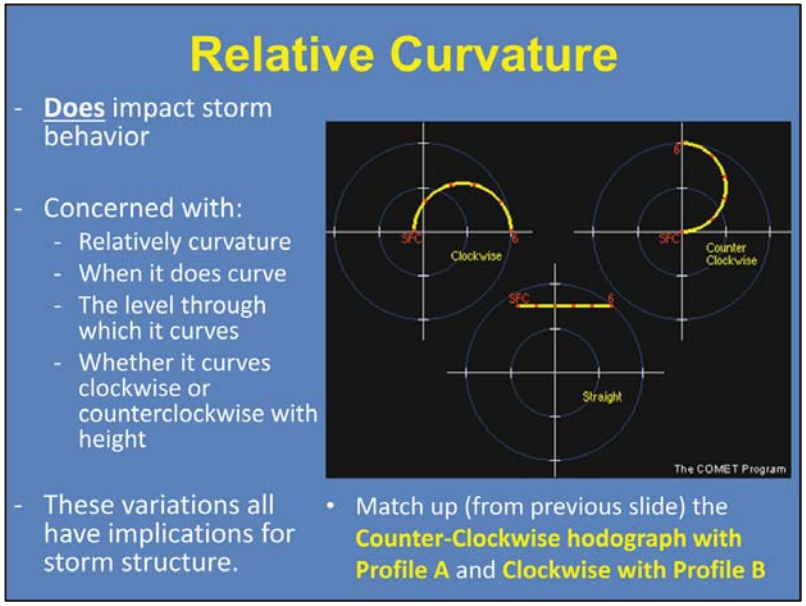
Almost there, but first we have to divide both the total length from X and Y and by 6 km and get the mean of each one. Then, we plot those back on the graph, starting with first hodograph point, down 2.8 on the y-axis and then over 5.5 on the x-axis, which equates to a 6.2 m/s magnitude line – the red line – that points to the ESE, which was the first thing we determined back a couple of slides ago.

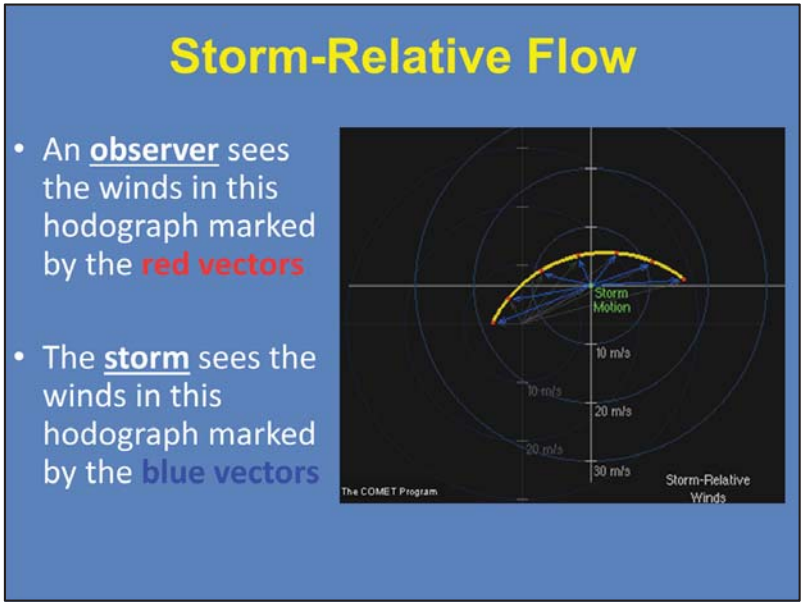
Shear Orientation

- **Does not** affect supercell behavior
- May be an indication of large scale synoptic conditions
- **Profile A** indicates backing winds (cold air advection; subsidence)
- **Profile B** reveals veering winds (warm air advection; rising motion)



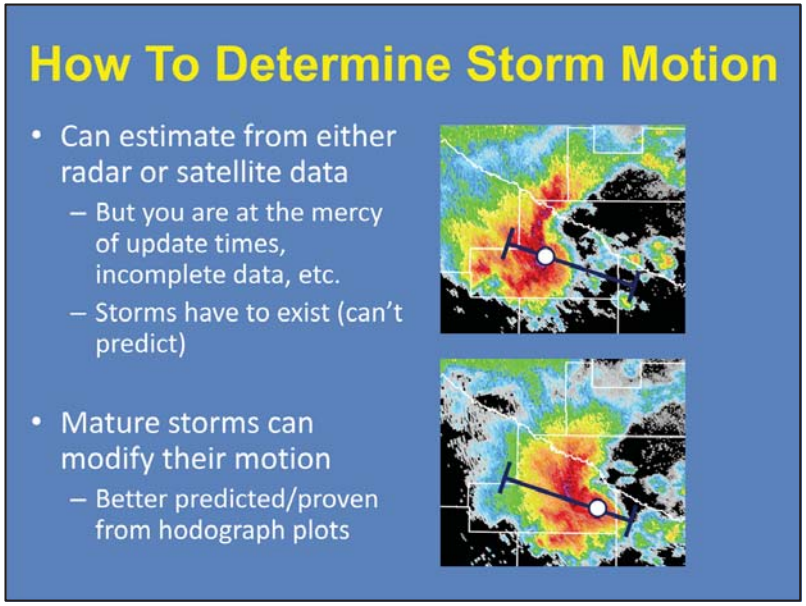
Let's go over some functions of hodograph shapes and different types of plots, which can tell us a lot about the environment. First, the Shear Orientation – what does mean? We'll take these two different wind profiles and just plot the direction or mean orientation of each profile. And so the orientation is more of a large-scale or synoptic behavior model as opposed to showing the behavior of individual storms or supercells. Profile A, which shows a backing orientation which is an indication of cold air advection – and therefore subsidence. Meanwhile, Profile B shows a veering winds, showing warm air advection in that layer and possible rising motion.

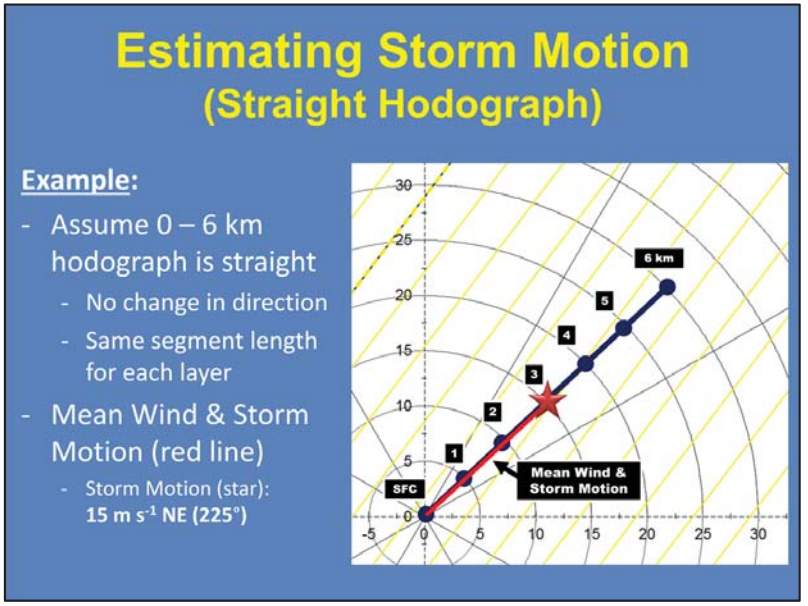




What does make a difference for individual storm behavior, especially supercellular ones, is the curvature of the shear on a hodograph. We are concerned about various aspects of curvature, including… the relative amount of curvature, i.e. how strong or intense, when it occurs – morning, noon, or night – where it occurs in the profile, and whether it is a counterclockwise or clockwise curvature. All of these aspects can be crucial in determining the type of environment you have for potential convection. In terms of what they mean, look at the wind profiles from the previous slide and match up the effects of the Counter-Clockwise hodograph with Profile A and Clockwise with Profile B.

After going through the Radar Products lessons, you should be familiar with the concept of "Storm-Relative", which has to do with the relative point of view difference between and observer at a fixed point and the storm itself. Similar to the SRM or Storm-Relative Map product, on the hodograph, the environmental winds are plotted from a fixed point, but sometimes we want to place that point on the storm itself and see how the environment affects the storm. This brings us to Storm-Relative Flow, starting off with the ground-relative location, which is the location of the observer and how the winds change from their perspective. Then, we transport the X and Y axis over to this Storm Motion point, to see how this environment is affecting the storm itself and thus, the storm-relative flow. But what is Storm Motion?





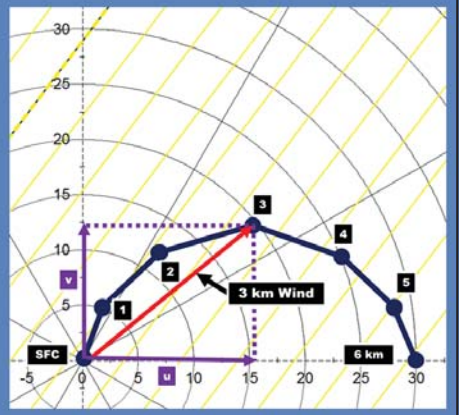
Next, we want to determine the storm motion, which will play into the relative effects that an environment can have on individual convective cells in a specific atmosphere. One can estimate storm motion, the direction and forward speed, based on either radar or satellite data. But, you have to have a storm to track, this won't do much good if you're trying to predict storm motion hours or even minutes in advance if there's nothing there. You are also at the mercy of how fast or slow the data updates and potential data quality issues. So, how do we do this on a hodograph?

How do we find Storm Motion on a hodograph? Let's start out with a very basic example with a straight hodograph, meaning, the winds throughout the column, 0 to 6 km in this case, are in the exact same direction, top to bottom. In this case, we have winds out of the SW that are also the same magnitude throughout the column, making it easy to average out both the mean wind speed and direction for this profile, which is 15 m/s and the vector would point toward the NE, or at 225 degrees, which is a Southwesterly wind. But in this case, we're not talking about vectors, we want a point which will be our new reference point, acting as the storm itself. So then, this red star is going to be our Storm Motion and you see where it is placed on the chart.

Estimating Storm Motion
(Curved Hodograph – Part 1)

Example:

- Curved 0 – 6 km
  - Change in speed & direction at all levels
- Find the Mean Wind (0 – 3 km)
- Create **u** and **v** components of that vector



Unfortunately for manual calculations, reality is usually much more complicated, so we'll add a little curve to the profile. You can see how the winds change in speed and direction as we go from the surface to 6 km. Recall a few slides back we determined the mean wind by making a vector go from the surface to the elevation of interest. In this case, we can just make a vector go from 0 to 3 km, with this red arrow. Then, we can break down the vector into its u and v components and plot them as well. To calculate the mean wind between 0 and 3 km, we'll average out u and v, like we did a few slides ago, then add these means together to get the total mean wind vector.

Estimating Storm Motion
(Curved Hodograph – Part 2)

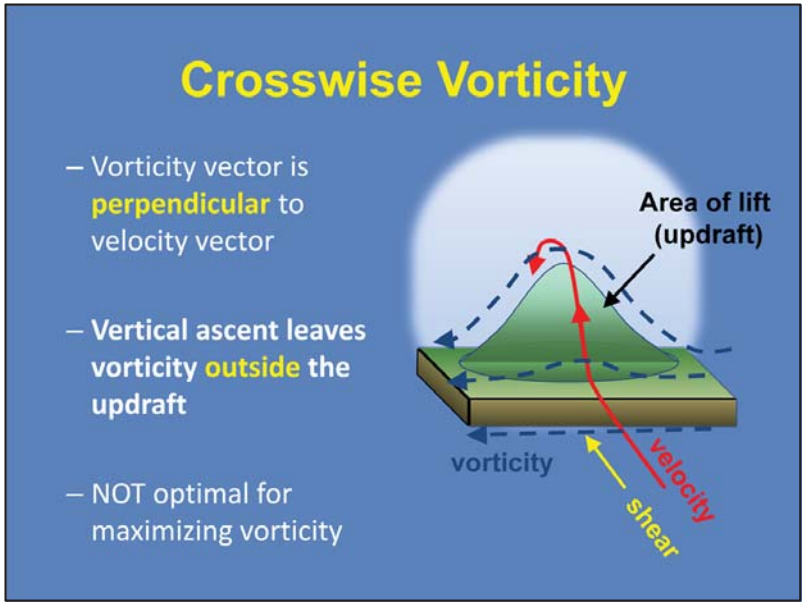
Example cont'd:

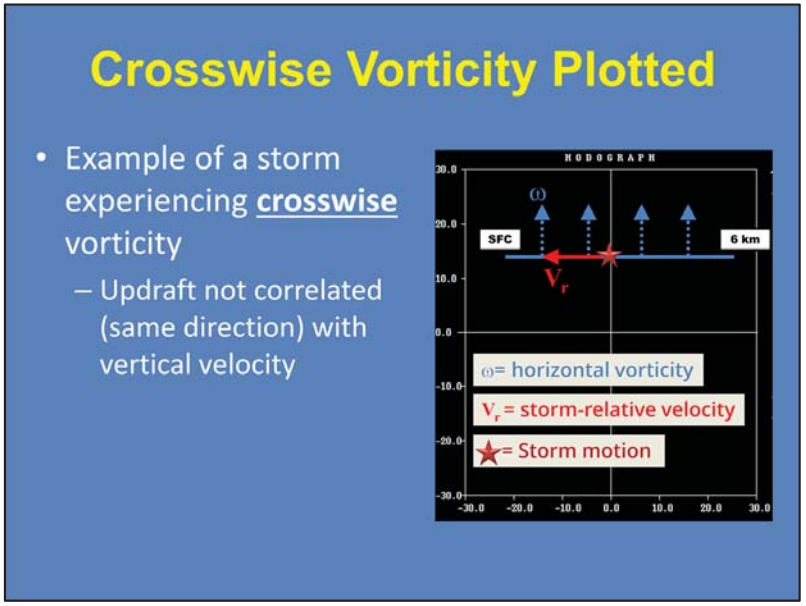
- First, estimate mean u for 0 & 6 km winds
  - (if you assume segments are relatively equal)
  - u (0 km) = 0 m s^{-1}
  - u (6 km) = 30 m s^{-1}
  - $\bar{u}$ (0 – 6 km) = 15 m s^{-1}
- Next, find the mean of all the v components
  - $\bar{v}$ (0 – 6 km) = 7 m s^{-1}
- Add u + v components (red line)
- Then, plot the Storm Motion (red star)



$\bar{u}$ = mean u
$\bar{v}$ = mean v

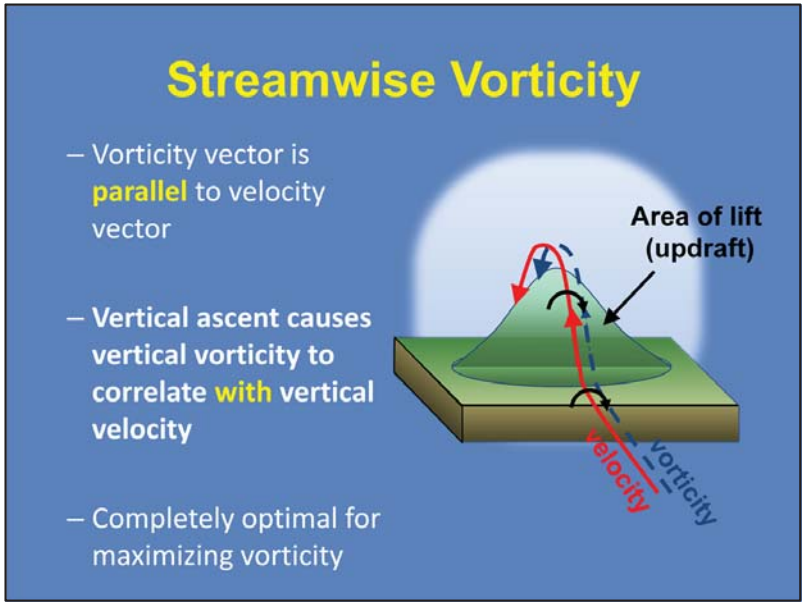
Continuing with this example, first we'll take the surface and 6 km wind and average them, to find the mean u. This is easy if all the segments are roughly equal, like in this example. It would be a bit more labor intensive if you had very different segment lengths. Then, do the same with the v components, which is about 7 m/s. Once the mean u and v are found, you connect them with the vector that is the mean wind, the red line. Then, the point with the star is the Storm Motion. Now…from here it gets more complicated with variations that can occur, even with this example, where this whole profile that, for instance, can be re-oriented onto a different axis – say, we re-orient this entire profile and all the vectors 45 degrees to the left – which would look a bit more realistic. Also, you can have multiple curves and make this profile look more like an S-shape which would be more difficult to compute but again, is not uncommon in the real world.

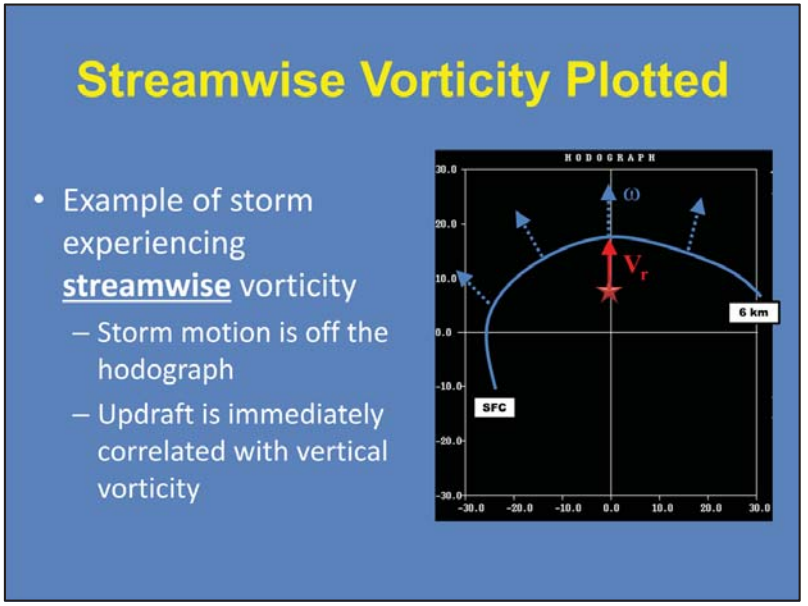




Now, we're going to add in this idea of vorticity, which – for hodographs – is broken down into two main types, Crosswise and Streamwise. Recall that vorticity is tendency for something to spin and a vorticity vector basically gives us the information about how that object will spin, so then it's up to the environment to effectively or ineffectively utilize the vorticity that is present. Crosswise, which is the example on the right, is the most absolutely inefficient method to utilize vorticity. In other words, if you need a certain amount of vorticity to go into an updraft and cause it to spin – you won't get it with crosswise vorticity because the shear and the mean flow are oriented perpendicular to the vorticity and passing by each other with no real interaction. If you have an updraft, the vertical ascent is not able to utilize the vorticity because of this skewed orientation.

Here is an idealized straight hodograph, from 0 to 6 km, with the Storm Motion plotted with the red star. The storm-relative velocity on this chart, the red V-sub R vector, represents the mean wind. You can see that this mean wind vector and the storm motion are essentially on top of the hodograph plot, which – also notice – is straight. The horizontal vorticity vectors are perpendicular to the hodograph which means that they are basically not affecting each other. They would need to be in the same direction, in order to maximize the potential of vorticity in rising motion. If you're trying to make sense of this in a conceptual format, head back to the previous slide and check out the diagram, to see how the vorticity lines just continue on their path, unaffected by the ambient wind, which is the point of crosswise vorticity.



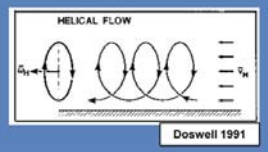


In direct contrast, there is streamwise vorticity, which means that the vorticity vectors are completely parallel to the velocity vector. Again, think of the velocity as the ambient wind and if you're trying to achieve a rotating updraft, you will efficiently utilize the vorticity into the updraft when both vectors are parallel, working in sync – in other words. Vorticity is your "spin" and when the velocity can efficiently tilt that spin into upward-moving air, then your potential for supercellular storms increases.

This is an example of what streamwise vorticity looks like on a hodograph. You have probably heard that for supercells and especially for possible tornadic potential, you would look for a "curved" hodograph, even if you aren't completely sure why that is. In this example, the storm motion is off of the hodograph, which allows the storm to utilize vorticity at all angles throughout the layer of interest. When you determine the mean wind, which again, is basically the red arrow, you see how it is in the same direction as the blue horizontal velocity vector. Again, if you need a more conceptual reference for this, go back to the previous slide and check out the diagram to see how both of these elements need to be parallel in order to maximize the available vorticity.

Helicity

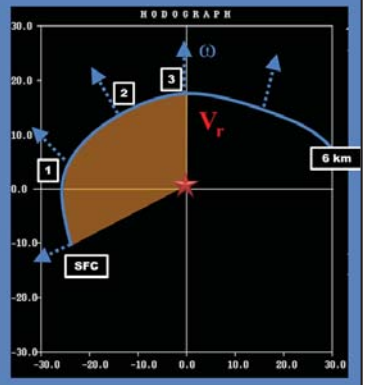
- **Helicity:** The property of a moving fluid to evolve into a helical flow
  - Corresponds to the transfer of vorticity from the environment to an air parcel (through convection)
- Proportional to the strength of the flow, the amount of vertical wind shear, and the amount of turning in the flow (i.e. vorticity)
- **Helical flow**: flow in the pattern of a corkscrew



Coming down the home stretch, we need to add in this last concept of helicity which builds into some of the most useful parameters that can be derived from hodographs. Helicity, in meteorology, refers to the property of a moving fluid to evolve into a helical flow and helical flow is basically in the pattern of a corkscrew. The transfer of vorticity from the environment to an air parcel occurs through the convective process and the amount of available helicity is a reflection how strong this convective process is.

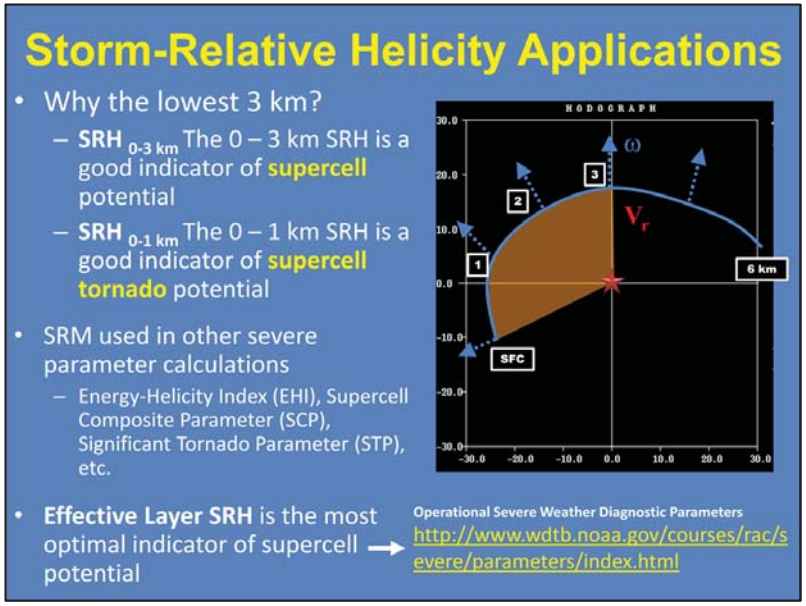
Storm-Relative Helicity (SRH)

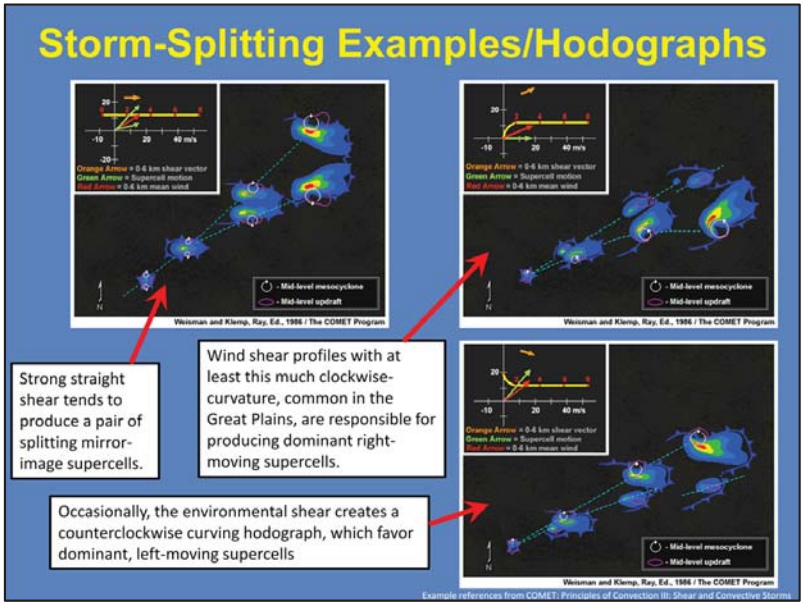
- **SRH** Definition: The potential for cyclonic updraft rotation in right-moving supercells
  - Integrated over a vertical layer
    - 0-1 km, 0-3 km
  - SRH is twice the area swept out between the hodograph and the Storm Motion between two levels *(0 – 3 km in the example)*



So why did we cover all these terms with helicity and vorticity? Why do they matter? Now, we can wrap together a number of these concepts into this term, Storm-Relative Helicity or SRH. Again, like the concept of Storm-Relative Velocity, as you learned about in the products lessons, we want to know how much relative streamwise vorticity is or could be ingested into a storm's updraft. Just like on a hodograph, when we experiment with the lifting parcel levels to get different values of CAPE for different situations, we use these concepts and calculations on a hodograph to find out a storm's potential for developing a mesocyclone (i.e. rotating updraft) which is the defining characteristic of a supercell thunderstorm.

SRH is integrated over a defined layer and is defined as basically the area under the curve between the hodograph and the Storm Motion, or the red star. We'll find out on the next slide why we use the 0 to 1 and the 0 to 3 km layers for this calculation.





Researchers using this parameter have found that the 0 to 3 km layer SRH is a good indicator of supercell storm potential, while the 0 to 1 km layer which is a more compact area near the surface, which is a good indicator of supercell tornado potential. The 0 to 3 km layer is more of a mean approximation of the storm's inflow and how the base portion of the storm is affected by the environmental shear, while the 0 to 1 km layer really takes into account the smaller-scale effects that normally determine whether a supercell tornado will form or not. And the story doesn't end there with SRM. It is used a number of other severe parameter calculations, each of them using this concept of storm-scale helicity usage in different ways. And then, Effective Layer shear is likely the most optimal of all parameters for indicating supercell potential. More on this can be found in the interactive module "Operational Severe Weather Diagnostic Parameters," just follow the link on the bottom right.

Lastly, one more application that groups together multiple concepts we learned about, here are 3 different idealized examples of a specific type of hodograph profile and what affect it has on storm-splitting. Look at the hodograph at the top-left of each image and see how each type of profile affects the production and maintenance of supercells. The terms left and right-moving refer to a storm's movement either to the left or right of a theoretical line in between two cells. These concepts will be covered in greater detail in some of your subsequent lessons in this Convective Storms Topic. But, at least you are getting a start on the basics and theory behind some of the more complex features and cases.

Takeaways (part 1)

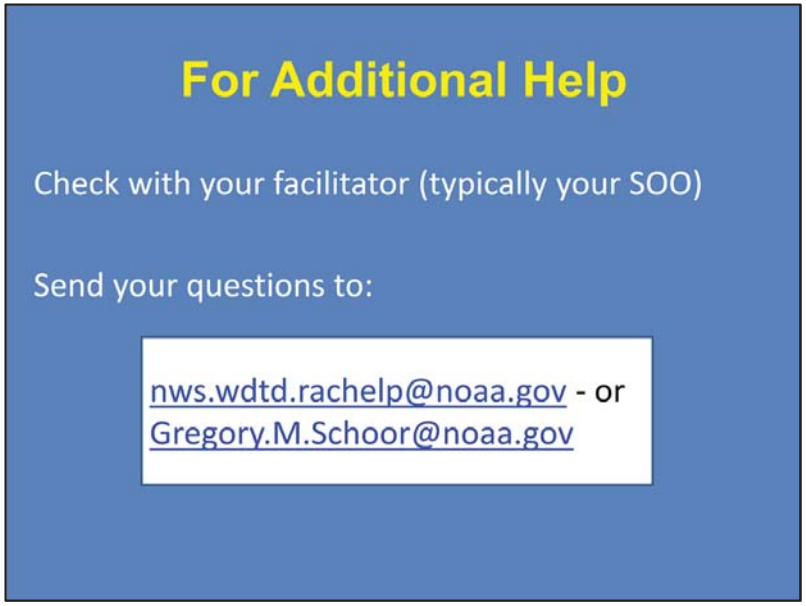
- **Hodograph is a plot of the vertical wind structure**
- **Shear**
  - Wind difference between two layers
  - Important for determining characteristics of convective storms
- **Certain specified layers of the atmosphere…**
  - are important for specific severe parameters
- **Bulk, Total, and Mean Shear**
  - Foundational for other severe parameters from the hodograph plot

Now, let's summarize what we covered. This lesson showed how hodographs are plots of the vertical atmospheric wind structure, using the wind information from a sounding. The hodograph is all about shear – which is the change or difference between two layers, any two layers – both speed and direction. From this information, we can determine what the possible characteristics, or you could even think of it as the side-effects of a storm being in a certain type of environment. We saw how specific layers of the atmosphere are interrogated for certain types of features, or effects, that could have implications on convective cells. And in order to know these things, we look at several different base parameters derived from the hodograph, which are the foundations for more specific calculations.

Takeaways (part 2)

- **Storm-Relative flow**
  - Can see how the environmental shear will affect a storm
- **Crosswise vs. Streamwise vorticity**
  - Crosswise is not optimal for utilizing vorticity
  - Streamwise is totally optimal for utilizing vorticity
- **Storm-Relative Helicity**
  - A measure of the potential cyclone updraft rotation in right-moving supercells

Then, we looked at this idea of storm-relative flow and how we need to see things from an individual storm's perspective. And then, getting into the heart of a storm's characteristics is how the two types of vorticity, crosswise and streamwise, will affect a thunderstorm in-particular, the updraft. And finally, the calculation for how vorticity and helicity measure-up for a thunderstorm with the Storm-Relative Helicity parameter.



For additional help, check with your facilitator (typically your SOO) or send your questions to the either of the e-mail address listed – and thanks for your time!



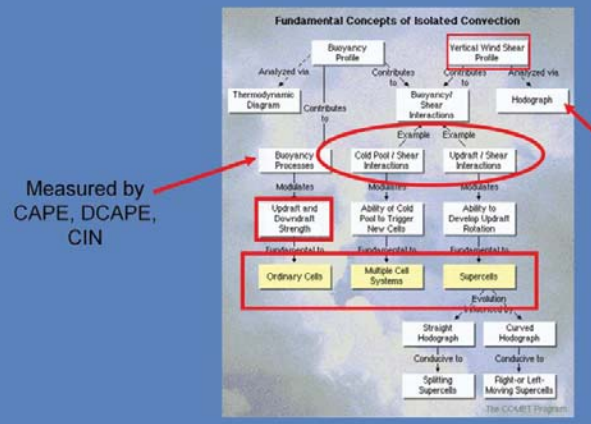
Welcome to fundamental relationships between shear and buoyancy on convective storm structure and type. This is a lesson in the RAC topic of convective storm structure and evolution. Hi, I'm Brad Grant, your instructor. The lesson is about 20 minutes in length.

Learning Objectives

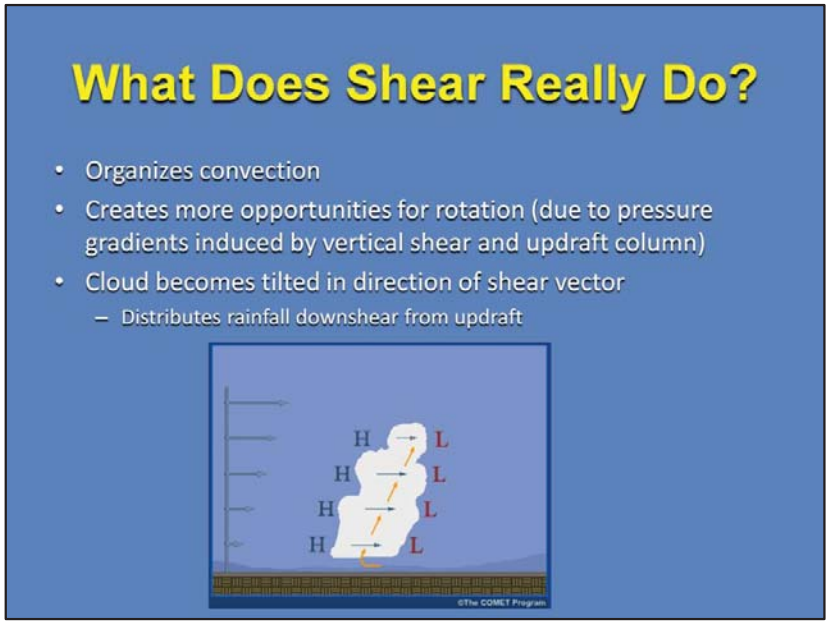
1. Determine influences of shear strength on overall storm structure and evolution.
2. Identify influences of the buoyancy profile on overall storm structure and evolution.
3. Identify influences of midlevel dry air on storm structure and evolution.
4. Identify the role of shear depth in controlling resulting storm structure and evolution.
5. Identify the role of hodograph curvature in controlling resulting storm structure and evolution for strongly sheared environments.

These are the learning objectives for lesson 1. Learn these items and you will be able to ace the test effortlessly.

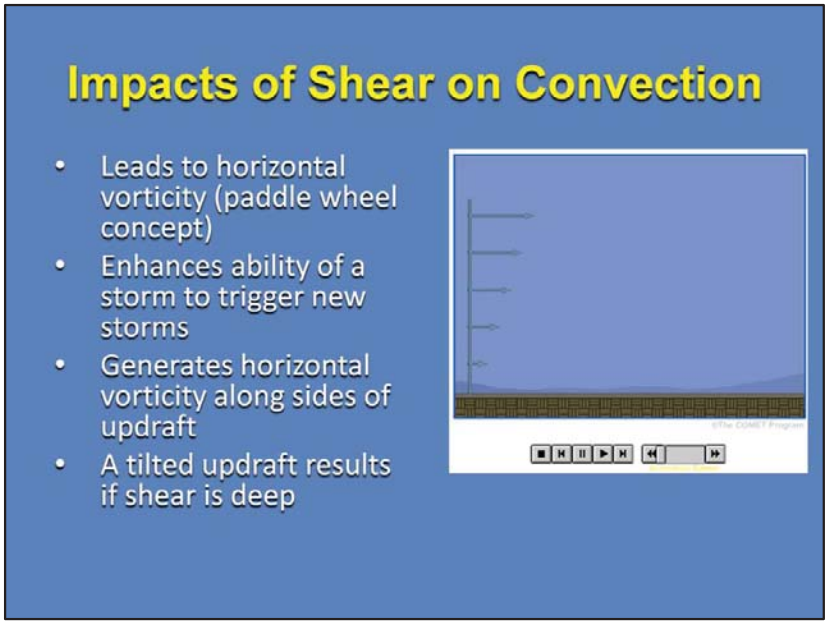
Spectrum of Convection



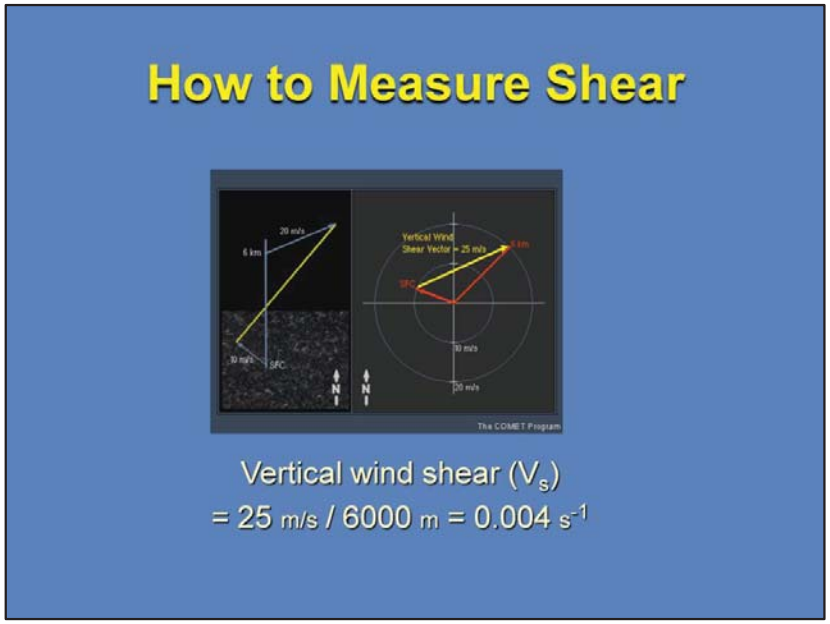
In the spectrum of convection, as seen here in this graphic from COMET, there are two primary mechanisms that control convective evolution, vertical shear and buoyancy. The vertical shear profile is analyzed by the hodograph and measured by shear and storm-relative helicity (SRH). We assess shear because it influences the updraft and can lead to internal dynamic processes important for supercells. Also, the interactions between the ambient shear and cold pool are fundamental to modulating multicell systems. The buoyancy profile is crucial in modulating updraft and downdraft processes which are fundamental for ordinary cells, or pulse storms, as they are often called. All three types of storms are directly affected by the distribution of buoyancy and shear, as well as other mechanisms such as forcing and triggers.

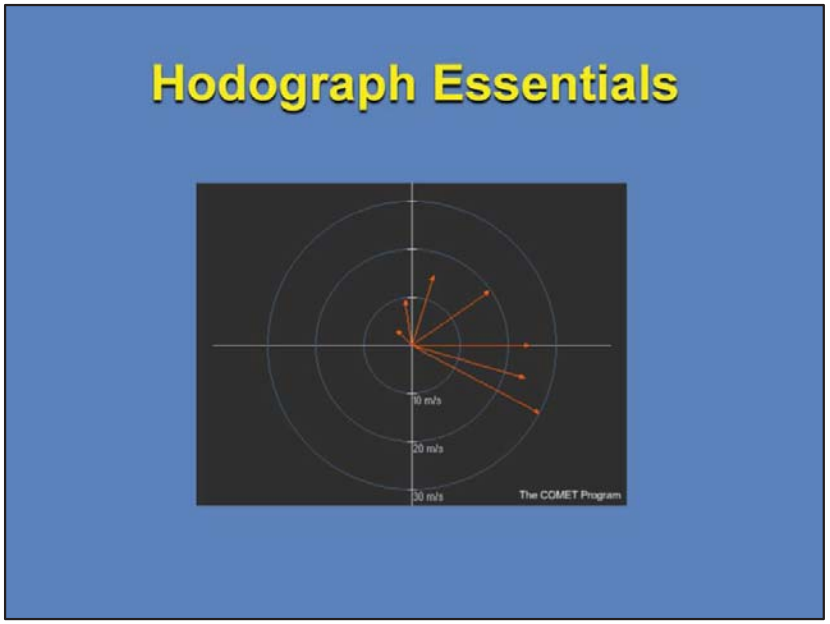


What does shear do in the environment for convective storms? Well, it organizes convection by creating more opportunities for rotation (due to pressure gradients induced by vertical shear and updraft column). And, it causes clouds to becomes tilted in direction on shear vector. By doing this it, distributes rainfall downshear from updraft.



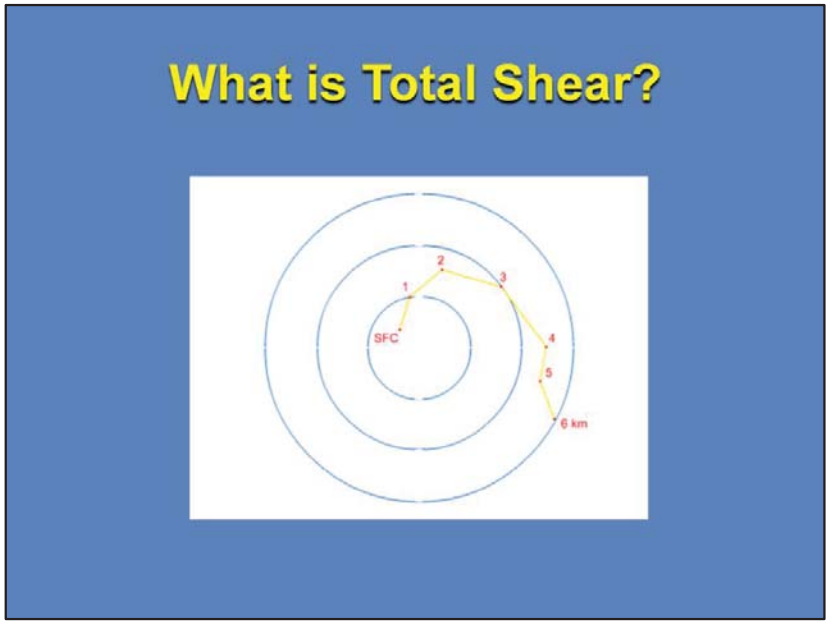
This graphic illustrates the impacts of shear on convection. There are four important processes that occur. 1) It leads to horizontal vorticity which is the paddle wheel concept, 2) It enhances the ability of a storm to trigger new storms. 3) It generates horizontal vorticity along sides of updraft due to the high to low pressure gradient, and 4) A tilted updraft results if shear is deep enough.



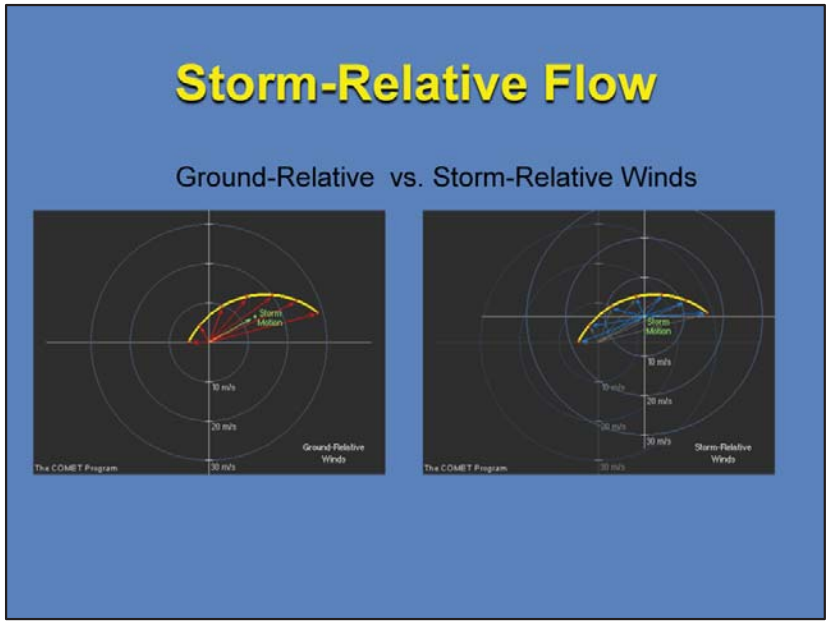


The hodograph shows how the wind changes with height, which we refer to as the vertical wind shear. Wind shear is a vector, defined as the vector difference between winds at different levels. This figure shows a vector difference between two points on the hodograph (the surface and 6 km). So, the resultant vertical wind shear is exactly 25 m/s over 6 km, which is equal to .004 per sec. By looking at consecutive segments of the hodograph, one can see how the shear vector changes with height. When the shear vector is veering with height, the hodograph is turning clockwise and when the shear vector is backing with height, the hodograph is turning counterclockwise with height.

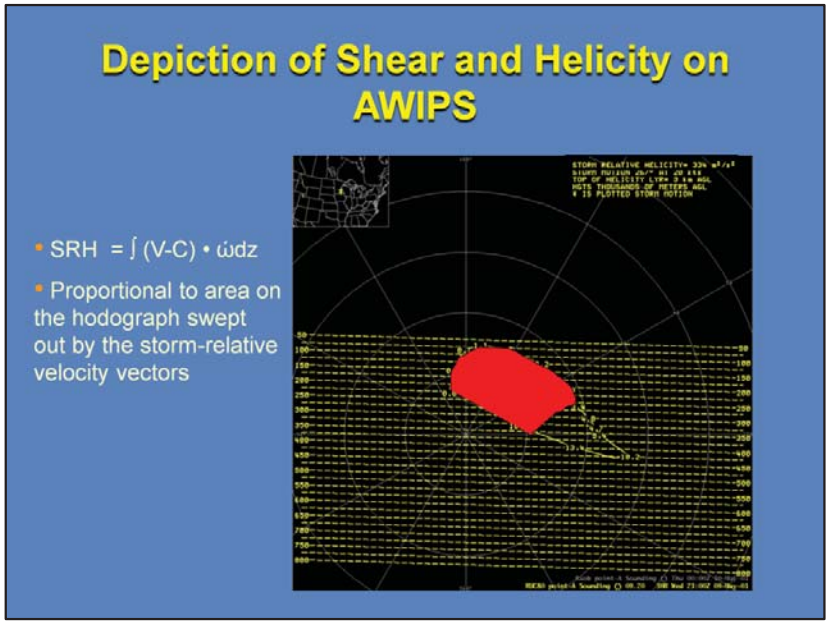
The hodograph is ideally suited for displaying vertical wind shear. Using a polar coordinate chart, shear is revealed by drawing shear vectors from the ends of each wind vector in sequence of increasing height. You can see that the line segments of the typical hodograph actually represent the vertical wind shear for each layer. If the wind vectors on a hodograph represent the winds at even intervals (usually each kilometer or 500 m), the shear vectors are equivalent in terms of the depth they represent. Their relative lengths then indicate the relative strength of the wind shear from layer to layer.



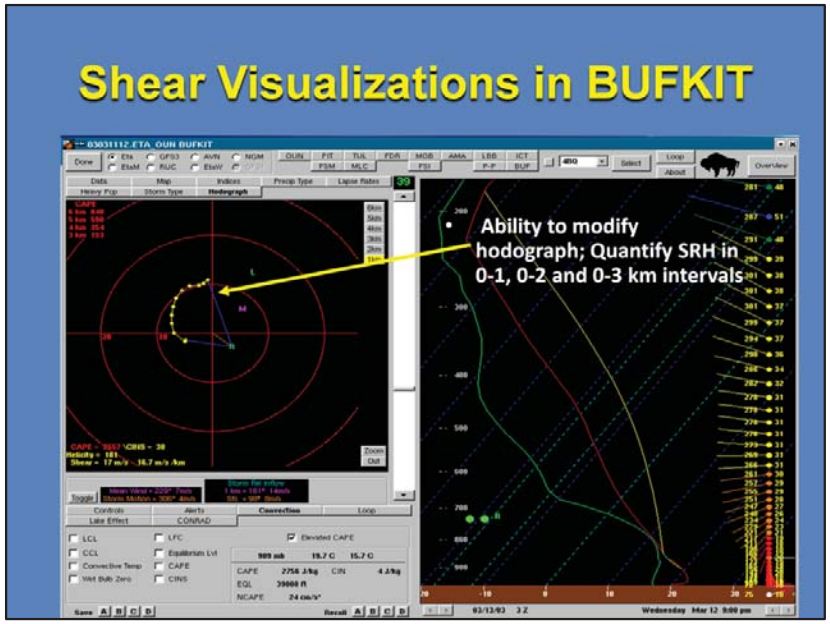
Estimating total vertical wind shear is done by combining the lengths of all the shear vectors over a particular depth (the net length of the hodograph). In this example, the total shear would be 60 m/s over 6 km.



In anticipating the evolution of a convective storm, it is important to consider both the environment that the storm is currently experiencing and the one into which it might move. For instance, the nature of the air making up the storm's inflow has implications for its evolution. To better understand this element of the storm's environment, you must first determine or estimate anticipated storm motion on the hodograph and then study the potential storm-relative winds. Because the storm moves through its environment, the wind it experiences is often very different from the ground-relative winds measured with the stationary sounding. This difference is apparent for the hodograph and storm motion shown on this page.



The area hatched under the hodograph is SRH from 0-3 km. Storm-Relative Helicity (SRH) is proportional to streamwise vorticity and storm-relative winds and takes into account storm motion. The mathematical expression for SRH, as defined Davies-Jones et al. (1990) is ∫ (V-C) • ώdz : where *V* is the horizontal velocity vector (ground-relative vector wind), *C* is the storm motion vector, and *w* is the horizontal vorticity vector. The integration is over the inflow layer of the storm from 0 km to some depth h (typically 1 to 3 km).



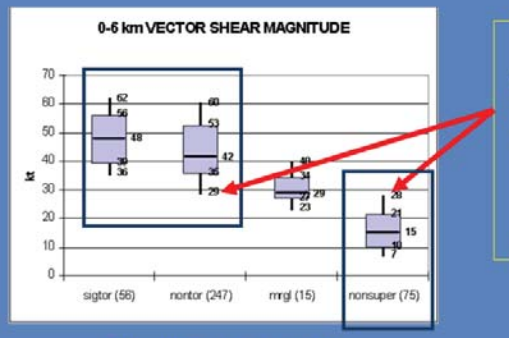
BUFKIT allows the user to modify the hodograph interactively to visualize changes in vertical wind shear, storm motion, and subsequent storm-relative helicity.

Vertical Shear Uses

- Important source of organization in severe thunderstorms (due to internal dynamic effects)
- Deep layer (0-6 km) vector shear magnitude of > 15 m/s indicates a potential supercell environment
- Low level shear (0-1 km AGL) of 10 m/s or greater indicates a potential for significant supercell tornadoes

Important source of organization in severe thunderstorms (due to internal dynamic effects)
Deep layer (0-6 km) vector shear magnitude of > 15 m/s indicates a potential supercell environment
Low level shear (0-1 km AGL) of 10 m/s or greater indicates a potential for significant supercell tornadoes

Some of Shear's Best Uses?

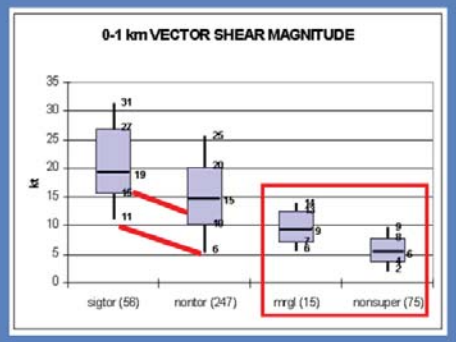


No overlap in values between 10th percentile for supercells and 90th percentile for nonsupercells

Thompson, et al. (2002)

The 0-6 km shear magnitude clearly discriminates between all supercells and non-supercells, with no overlap in values between the 10th percentile values for supercells (29-36 kt) and the 90th percentile for nonsupercells (28kt). So, Lower-bound thresholds of 0-6 km bulk shear of approx. 15-20 m/s (35-40kts) can be used as a first approximation for potential supercell environments. The top and bottom of each shaded box denotes the 75th and 25th percentiles, respectively, and the heavy horizontal line is the median value. The heavy vertical lines extend upward to the 90th percentile, and downward to the 10th percentile. Groups of supercells are labeled, with sample size in parentheses.

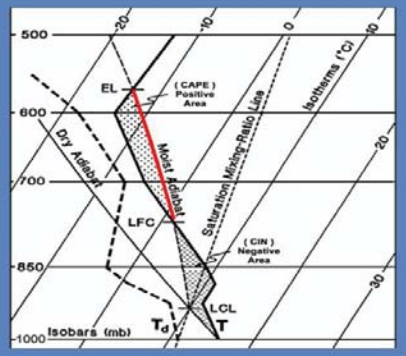
0-1 km Shear Use?



The amount of 0-1 km shear in the environment seems to differentiate between strong (or significant) tornadoes and weak or no tornadoes Vertical shear in the lowest km tends to be about 5 kt stronger for the significant tornadic supercells, and much weaker for the marginal and nonsupercell storms.

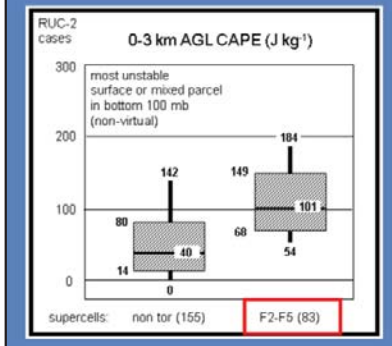
Buoyancy Influences

- If CAPE increases…
  - Size, depth, and strength of individual cell increases
- Must also consider forcing
- Must also consider convective inhibition (CIN)
- Must also consider water loading



The convective available potential energy (CAPE) is represented by the area on a skew-T diagram enclosed by the environmental temperature profile and the moist adiabat running from the LFC to the EL.(click) This area, depicted in this diagram, indicates the amount of buoyant energy available as the parcel is accelerated upward. CAPE is measured in units of joules per kilogram (J/kg). If CAPE increases… the size, depth, and strength of individual cell increases. But, there are lots of other factors beyond just CAPE that must be considered when assessing storm severity.

CAPE Uses



Davies (2002)

- Related to updraft velocity by this expression

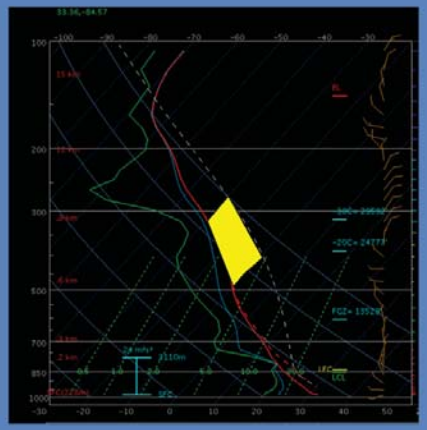
$W_{max} = \sqrt{(2 \cdot CAPE)}$

- Low-level CAPE (and LCL ht) strongly correlated to pre-storm "strong" tornado environments

CAPE is also related to updraft velocity by this expression Wmax= √ (2 • CAPE). Low-level CAPE (and LCL ht) strongly correlated to pre-storm "strong" tornado environments as you can see in this graph from Davies on 0-3 km MUCAPE and the associated correlation to supercell classes. Apparently, the vertical accelerations in the lowest levels of the updraft are much greater in significant tornado environments than in weaker or no-tornadic environments. See the Parameter web site for more details.

CAPE Uses (cont.)

- Large CAPE in hail growth zones (-10 to -30 deg C) is a necessary (but not sufficient) ingredient in severe hail production



Large CAPE in hail growth zones (-10 to -30 deg C) is a necessary (but not sufficient) ingredient for Hail production, assuming you have deep moist convection. Here we have the area bounded by the -10 to -30 deg C isotherms and the parcel ascent curve to view how the CAPE is distributed in this zone.

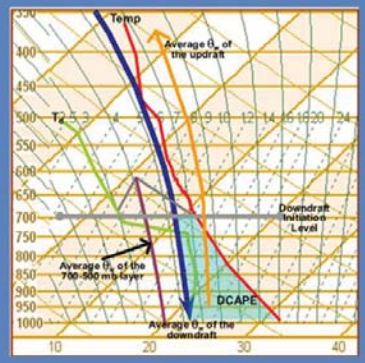
CAPE Limitations

- Assumes no mixing with surrounding environment
- Ignores effects of freezing and water loading
- Most calculations use ambient temperature instead of virtual temperature - lowers true CAPE
- Extremely sensitive to parcel origination and depth of integration

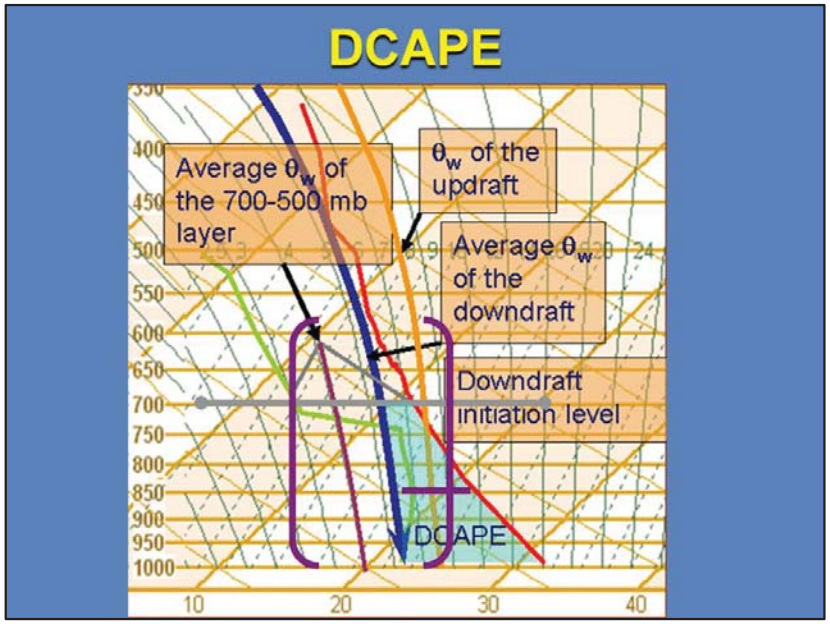
Tall-thin" CAPE is more susceptible to water loading than "short and fat" CAPE. For example, tropical storms, which develop in soundings characterized by high ELs, and tall-thin CAPEs, are not as likely to be as deep (in terms of convective growth) as shallow-topped cool season supercell storms , where representative soundings indicate short-fat CAPE.

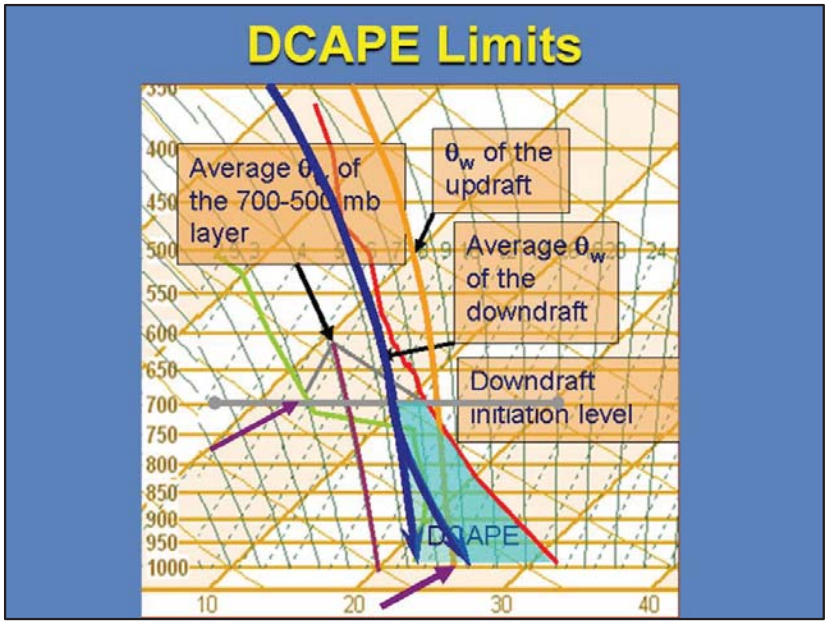
Dichotomous Effects of Dry Air

- Dichotomous effects
- Midlevel dry air can be entrained into both convective updrafts and downdrafts
- Decreasing positive buoyancy
- Increasing potential downdraft buoyancy



The reason for this dichotomous effect is that midlevel dry air can be entrained into both convective updrafts and downdrafts, decreasing potential updraft buoyancy, but increasing potential downdraft negative buoyancy.

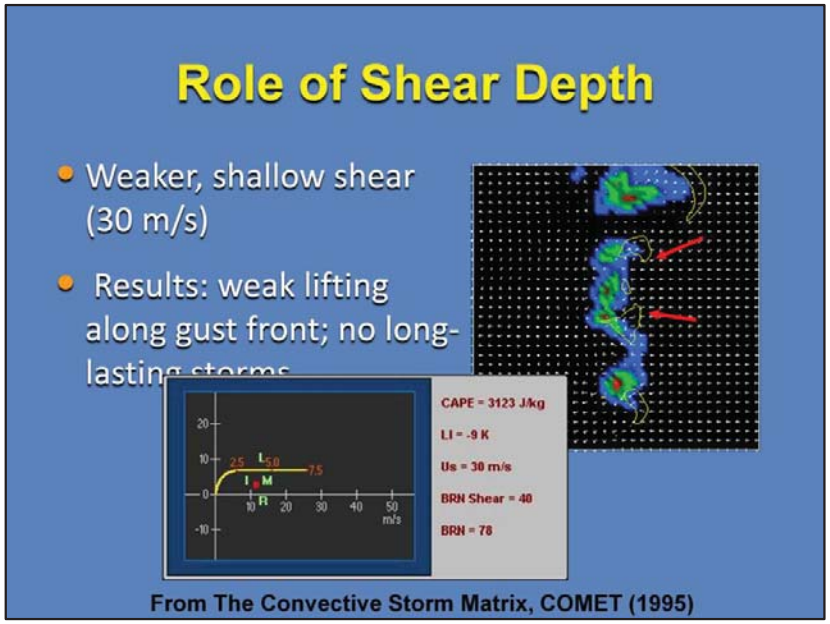


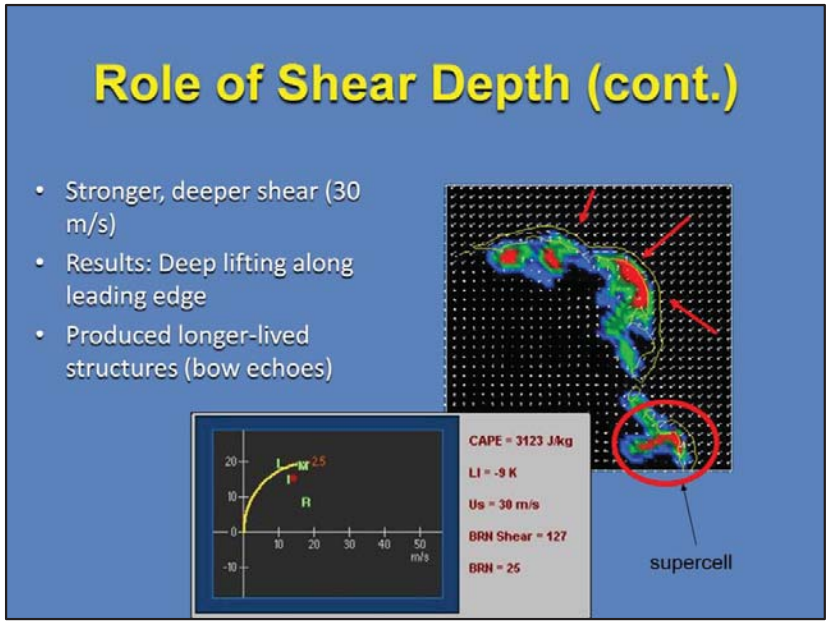


In this figure, we see the computation of DCAPE (green shaded region). Thick green (red) line is environmental dewpoint temperature. Thick blue line is the average Theta W, thick grey line is the downdraft initiation level situated at the bottom of the dry layer, and the thick orange line is the Theta W of the updraft.

If the updraft contains precipitation and environmental air mixes with ensuing evaporational cooling, it will create a negative thermal buoyancy over a vertical layer. The air cools to an average of the downdraft (brown curve) and the updraft (orange curve). A downdraft that is completely saturated would theoretically follow the mixed Theta W down to the ground (thick blue curve). We can then integrate the temperature difference between the environmental temperature profile (red) and the downdraft Theta w to calculate or DCAPE.

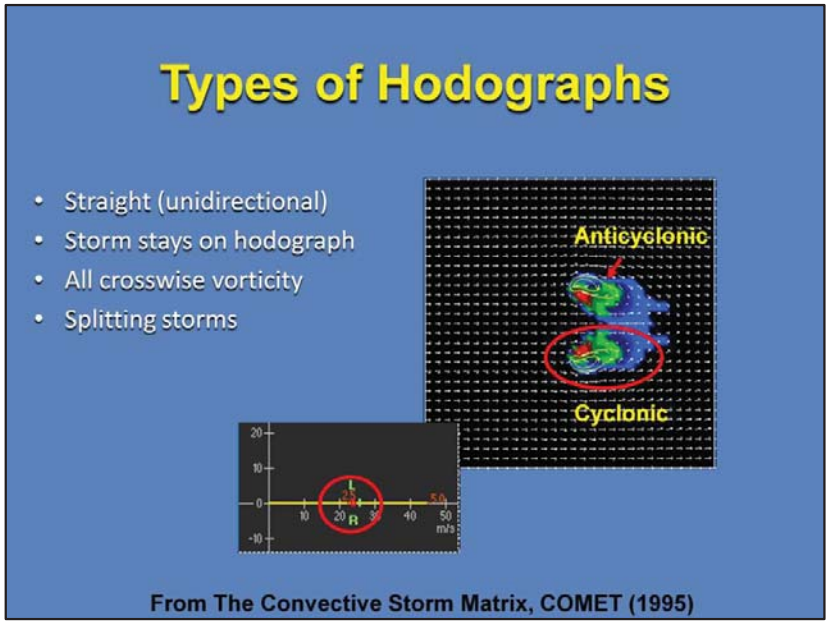
There are a few cautions to apply to properly interpret DCAPE. First, note that we started the integration at a specific level (700 mb) and called it the downdraft initiation level. In actuality, there is less certainty as to where the downdraft initiates than an updraft. Most downdrafts initiate over a layer rather than a level. That is why we used an average theta W in a layer, in this case, 700 - 500 mb. However, we could be equally justified to start the integration at a different level than 700mb. Picking a higher (lower) level most likely creates larger (smaller) DCAPE. A second caution of using DCAPE is that the downdraft most likely to be unsaturated as evaporation is never efficient enough to compensate for adiabatic compressional heating of dry air. The downdraft consequently never follows the theoretical Theta W curve and instead warms more quickly in reality. Most likely the realized DCAPE is much less than the theoretical leading to a weaker than expected downdraft. As a third caution, DCAPE does not account for negative buoyancy due to precipitation loading. A reflectivity core of 60 dBZ or greater may create enough precipitation loading to comparable to negative thermal buoyancy and therefore lead to a stronger downdraft than the DCAPE suggests. Finally, DCAPE does not account for nonhydrostatic downward directed pressure deficits that result from strong mesocyclogenesis or divergence beneath the level of interest. A Rear Flank Downdraft (RFDs) in a supercell (as you will see in one of the subsequent lessons) derives a significant portion of its forcing from a downward directed nonhydrostatic pressure deficit.

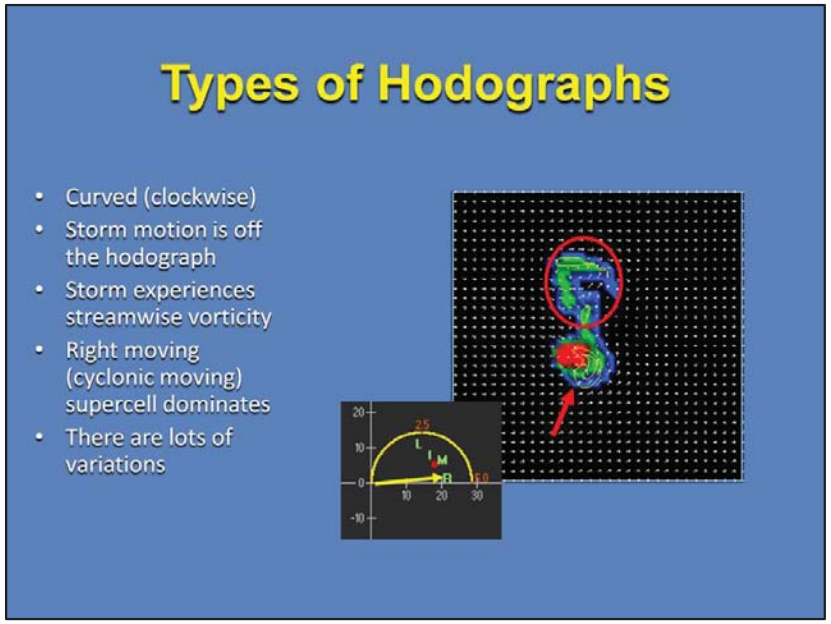




Vertical wind shear is crucial to the organization of convective systems. But just like CAPE, it's HOW the shear is distributed that is important in analyzing potential convective storm type. Here from the Convective Storm Matrix, we see a planview map at 4 km above surface of a model simulation of convection three hours after initiation. It uses the hodograph shown that possesses 30 m/s of shear over 7.5 km. The colored regions represent vertical velocity while thin yellow isohyets are vertical vorticity. The white vectors are system-relative winds. What happened in this simulation (and we observe this frequently in nature) is that in cases of weaker, shallow shear in a homogeneous environment of this amount of CAPE, it produces weak lifting along gust front; no long-lasting storms.

On the other hand, the results where strong, shallower shear environments exist (note the 30 m/s of shear is contained in the lowest 2-3 km) this leads to deeper lifting along the leading edge of convection, and longer-lived organized squall lines including bow echoes and occasional supercells. The simulations only tell part of the story. Environmental instability and system relative flow must also be considered when predicting eventual storm type.

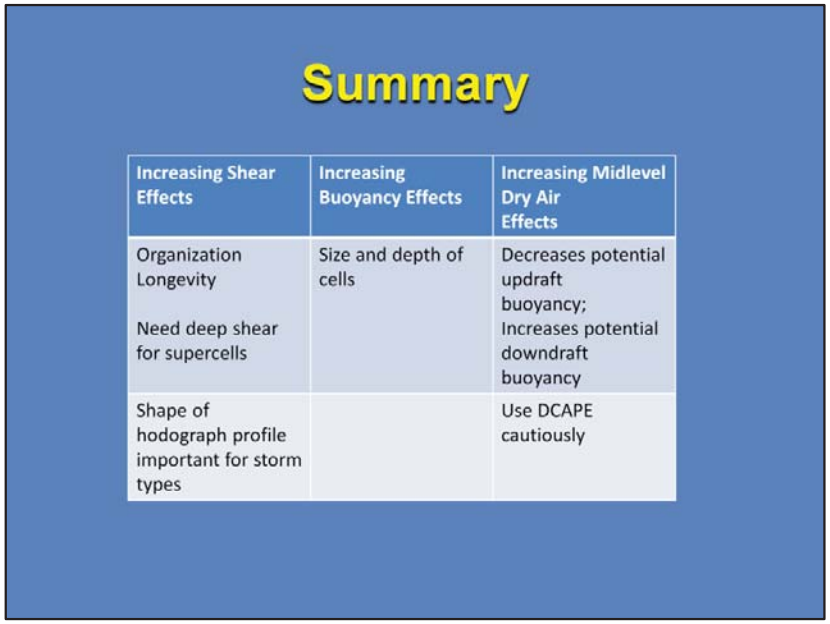




These next two slides illustrate two types of model simulated hodographs resulting from varying the shape of the hodograph. Essentially there are two basic types, straight and curved. These are the two most common types of variations seen in nature. What we have here are model simulations of updraft strength (shaded colors) and vorticity (yellow contours) at 4.6 km above surface and 1.5 hours after initiation for the straight hodograph shown in the inset. The hodograph has 46 m/s of shear over five kilometers.

Both straight and curved hodographs produce equally strong supercells given enough shear. But, straight hodographs allow both the right (cyclonic) and left (anticyclonic) moving supercells to be equally strong. Note the mirror image cyclonic and anticyclonic supercells in this simulation from an environment characterized by unidirectional shear . This is the straight hodograph example.

Clockwise (counterclockwise) turning hodographs favor the right-moving (left-moving) supercell and weakens the left-moving (right-moving) member. Remember, there are lots of variations.



In summary, I talked about the influences of shear on deep, moist convection. Essentially, the organization and longevity of convective storms and storm systems tend to increase with increasing magnitudes of vertical wind shear. Increasing the buoyancy in a convective storm or system tends to increase the size, depth, and strength of the individual convective cells, and the overall size and strength of the whole convective system. When assessing the likelihood of supercells, remember that shear and CAPE are often balanced. For example, in low CAPE environments (such as the cool seasons), stronger shear environments are often sufficient to produce tornadic storms. You need bulk shear on the order of 15-20 m/s through 6 km for supercell development. The shape of the hodograph is important in assessing helicity and resulting storm types, such as either a dominant right-mover, or splitting storms. In terms of midlevel dry air, we showed how dry air may be important for decreasing updrafts but increasing downdraft buoyancy. DCAPE can be used to assess downdraft strength potential but it needs to be used cautiously. More information about specific uses of thermodynamic and kinematic sounding parameters can be found in the pre-requisite lesson.



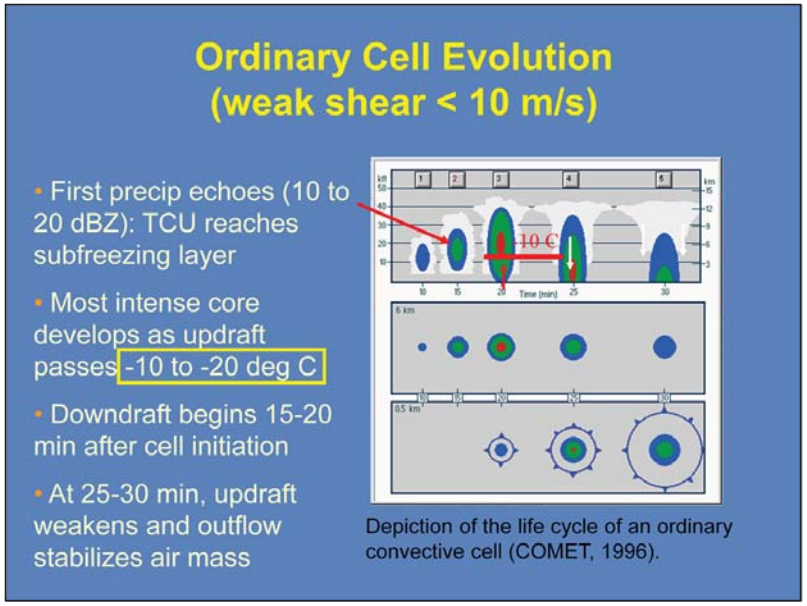
For additional help, check with your facilitator (typically your SOO) or send your questions to the listserv e-mail address here.

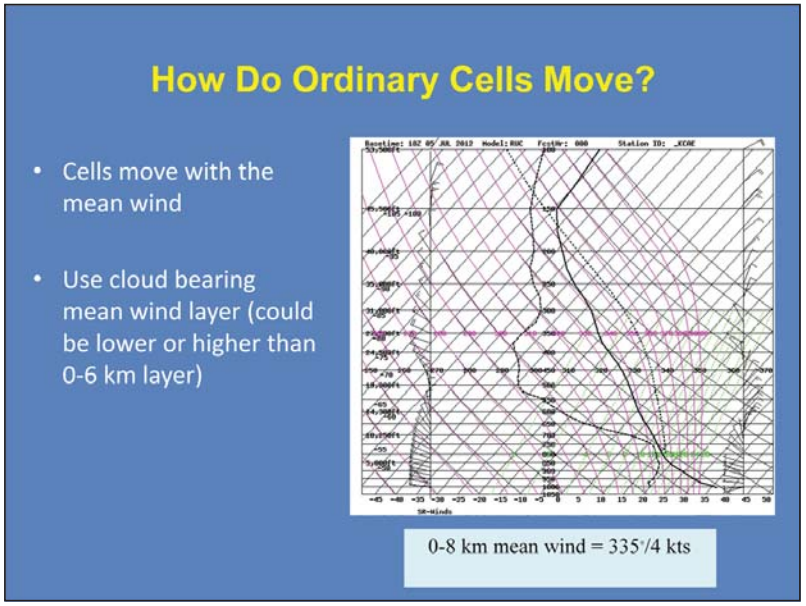


Learning Objectives

1. Identify the characteristics of convection associated with ordinary cells.
2. Identify how to anticipate the motion of ordinary cells.

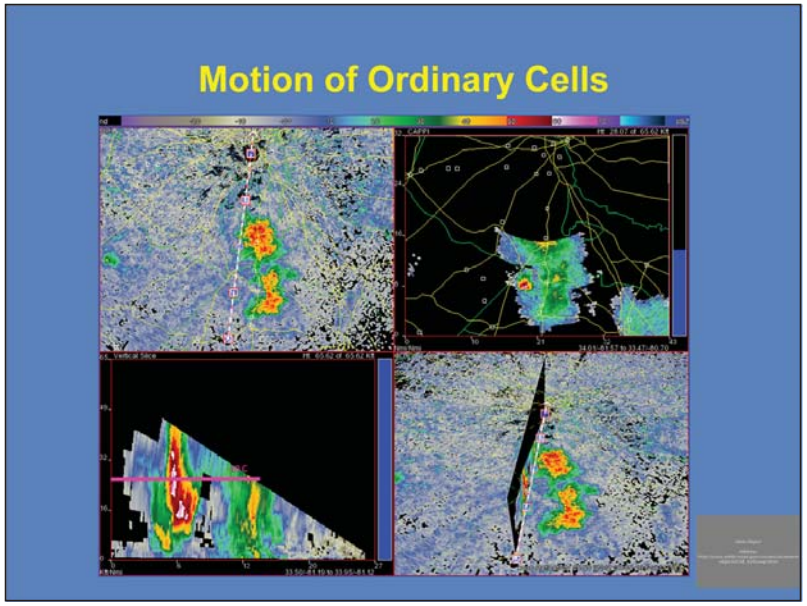
These are the learning objectives for this lesson. The test at the conclusion of this lesson includes questions based on these objectives.

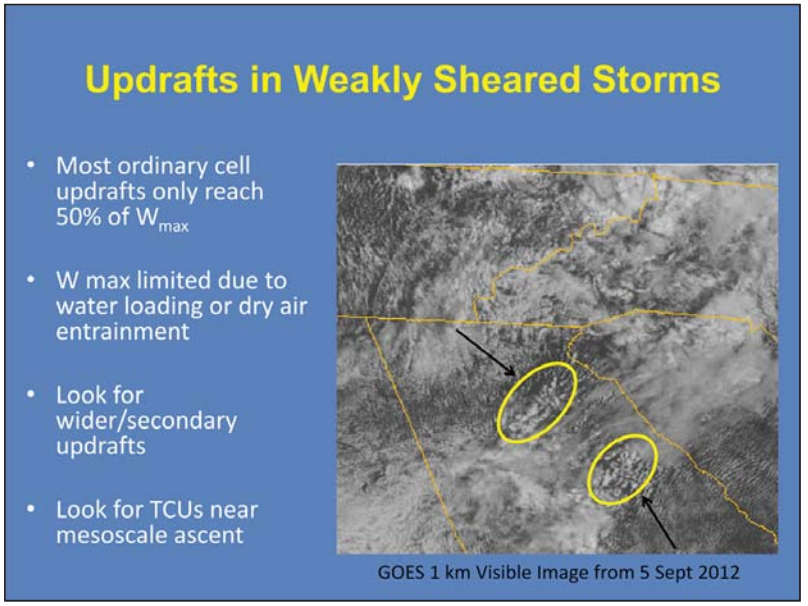




In this three panel graphic we see the cross section and plan views of an ordinary cell. The initial towering cumulus causes sharp gradients in the refractive index of the atmosphere along the cloud edges. These gradients scatter just enough of the incident WSR-88D energy back to result in -10 to 0 dBZ echoes just above the boundary layer. The first real precipitation echoes (10-20 dBZ) develop as the towering cumulus rises into the subfreezing layer. The most intense core develops as the updraft passes through the -10 to -20 deg C layer. The onset of downdraft is likely to occur as the precip core exceeds 45-50 dBZ. A downdraft usually begins between 15 and 20 minutes after cell initiation. The base of the descending precipitation core and the downdraft are typically, but not always, coincident. Therefore, when the core has reached the ground the downdraft begins to spread out into a cold pool. At this time, the updraft remains strong around a preferred side of the descending core. At 25-30 minutes after initiation, the updraft begins to weaken as the outflow stabilizes the low-level environment at its roots. Without a continuous feed of unstable low-level air in a weakly sheared environment, the updraft dies in the lowest several km above ground leaving an anvil behind.

Single cell storms in the absence of significant shear move with the flow at any level (which is not surprising since the flow at any one level is nearly the same as any other level). Based on Byers and Braham (1949) in the Thunderstorm Project, they found that the best estimate for steering layer flow was the mean wind in 0-6 km AGL layer. A common mean wind calculation is weighted by density giving more influence to the low-level flow. If the average uses a deeper layer (ex. 0-12km), then weighting the average may provide more accurate results. Also, low- (high-) topped thunderstorm motion may be influenced better by a shallower (deeper) mean wind. So, use the cloud-bearing mean wind layer, which could be lower or higher than 0-6 km layer. In this example, we can see from the sounding analysis output, the 0-8 km layer best represents the mean cloud-bearing layer is a mean wind vector from 335 deg at 4 kts, which would suggest isolated storms that develop would move very slowly SSE, which is exactly what they did.

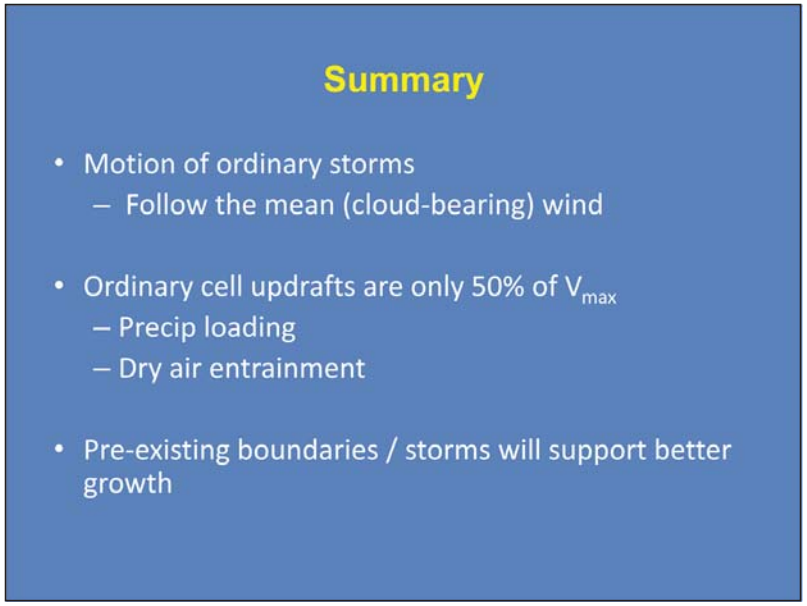




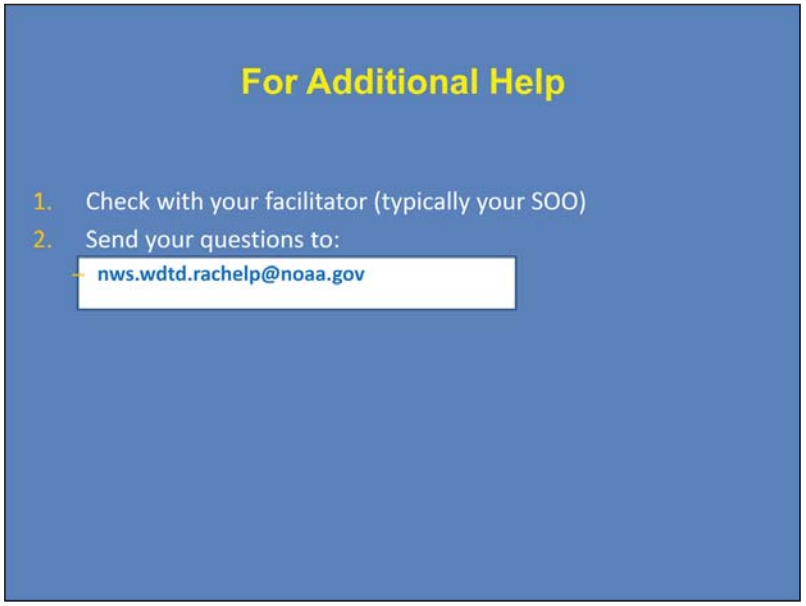
To illustrate an example of ordinary cells moving in weak shear, I am going to show an example from July 5, 2012 in Columbia, SC. There will be a flash animation (popping up in a new window) which illustrates a general southward drift of storms which developed south of the radar. In this animation, note the outflow boundary spreading outward from 20z to 21z. Pulse storms in this example tended to maximize their overall vertical extent after the leading edge of the outflow boundary and associated cold pool had passed. Outflow boundaries such as these typically spread out equally in all directions from a collapsing storm formed in an environment of weak ambient shear. The depth and orientation of the convergence in the boundary, plus the ambient air profile, were all factors in determining when and where storms would initiate. The still image shown on this slide is a screen capture at 1951z which shows a cross-section of a storm developing a 60 dBZ core above -20 deg C air. This pulse storm's updraft developed rapidly, then died off as the downdraft commenced within 45 minutes, but not before producing pea-size hail at 20z and a brief severe wind gust that downed trees at 2030z. A low-level divergence signature indicative of the small downburst associated with the downdraft stage of the storm is seen around 2034Z in the Velocity data.

The intensity of updrafts in ordinary cells is limited by their small size and speed. Most ordinary cell updrafts only reach 50% of their maximum updraft velocity due to water loading and/or dry air entrainment. A weaker updraft acceleration increases the chance that precipitation loading will diminish the strength of the updraft before it has a chance to reach the high theoretical speeds. Given the same CAPE, not all updrafts will be the same. Narrow updrafts are likely to be affected by lateral dry air entrainment so wider updrafts will allow the core to be protected. Use visible satellite imagery or radar representation of the midlevel core to estimate updrafts that will be less prone to entrainment. Secondary updrafts developing near a previous storm may grow in a more moist midlevel environment than what the models or RAOBs indicate.

A large area of towering cumulus growing in a region of mesoscale ascent (such as due to convergence from a sea breeze boundary or even differential heating) may provide clues that the environment is moister than expected. In this GOES-14 1 km visible image, you can see an updraft growing due to deeper moisture and stronger mesoscale ascent.



In summary, the motion of ordinary storms in the absence of even moderate shear, is dictated by the mean wind. Be cautious of using just 0-6 km layer alone to estimate the motion as the actual cloud bearing layer may vary based on the environment that the storms develop within. The actual updraft velocity of a storm is about 50% of its maximum updraft velocity due to water loading and/or dry air entrainment. A weaker updraft acceleration increases the chance that precipitation loading will diminish the strength of the updraft before it has a chance to reach the high theoretical speeds. To help compensate for these factors, look for updrafts developing near pre-existing boundaries or around previous storms as these will provide greater moisture and vertical growth.



For additional help, check with your facilitator (typically your SOO) or send your questions to the listserv e-mail address here.



Welcome to the Convective Storm Structure and Evolution lesson on assessing updraft strength and location.

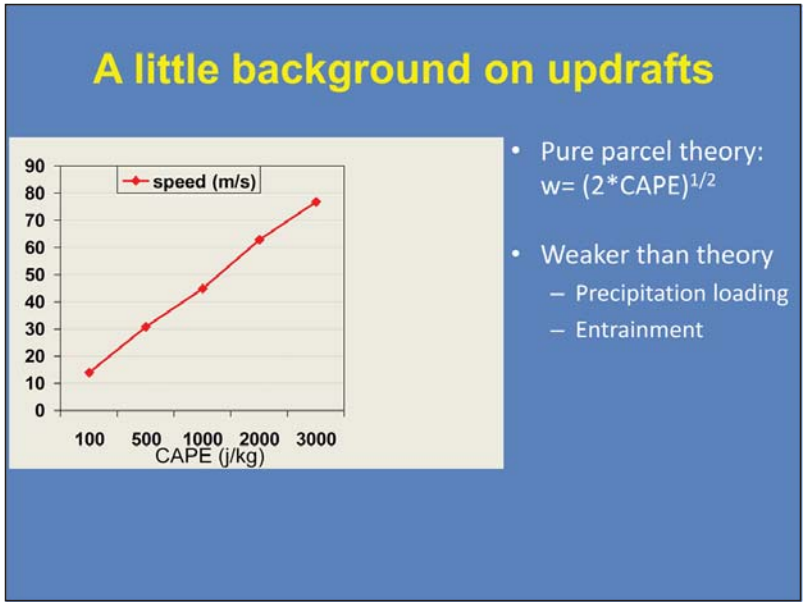
Learning Objectives

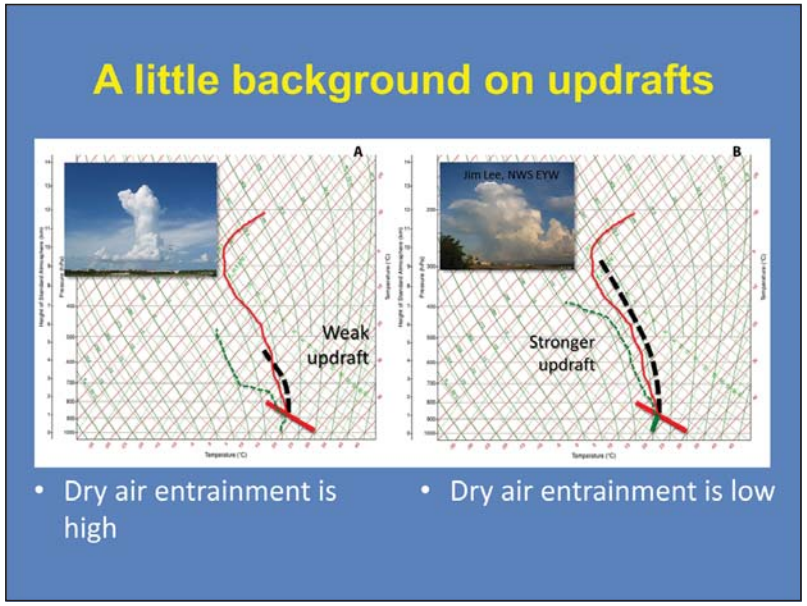
- Identify the guidance on assessing whether an updraft may lead to severe weather based on
  - the height and intensity of the upper-level reflectivity core,
  - low-, and upper-level convergence and divergence,
  - common updraft shape signatures
- Identify MRMS products, advantages and caveats with regards to identifying updraft locations.

The objective of this lesson is to provide you guidance on assessing whether an updraft may lead to severe weather based on:

the height and intensity of the upper-level reflectivity core,
low-, and upper-level convergence and divergence,
and common updraft shape signatures.

Identify MRMS products, advantages and caveats with regards to identifying updraft locations.





The estimation of the maximum updraft strength (Wmax) (based on the CAPE) does not take into account precipitation loading or dry air entrainment. Therefore, most ordinary cell updrafts reach only about 50% of Wmax due to these effects. That's the blue part of this domain where the red line is the pure parcel theory updraft speed.

Weaker updrafts can result from precipitation loading. A storm with 3000 J/kg of CAPE over 18 km of depth will have a weaker updraft acceleration than one with the same CAPE over 12 km. **A weaker updraft acceleration increases the chance that precipitation loading will diminish the strength of the updraft before it has a chance to reach the high theoretical speeds**. Stronger updraft accelerations advect cloud condensation nuclei upward so quickly that significant hydrometeor growth does not occur. Perhaps CAPE density can inform you about the updraft acceleration potential in a storm.

However, given the same CAPE and CAPE density, not all updrafts will be the same. Some storms remain weak regardless of the environmental CAPE. **Narrow updrafts are likely to entrain dry air to the core limiting updraft strength**. Also, significant midlevel dry air can increase the entrainment efficiency reducing the strength of an updraft even given large values of CAPE.

The impact of midlevel dry air is graphically represented by the more severe loss in parcel theta-E despite the same CAPE (or MLCAPE).

Updrafts can resist dry air entrainment if…

- If no dry air in the first place
- Wide updrafts
- Secondary updraft
- Strong mesoscale ascent

Given the effects of entrainment, look for these factors when considering the storm with the greatest updraft potential.

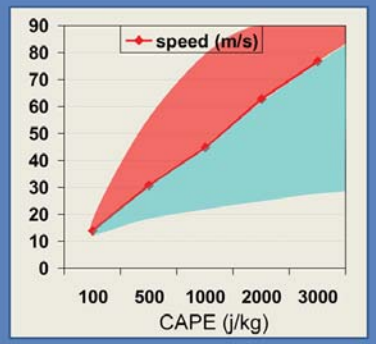
• **Look for the presence of dry air in a sounding** that could mix with the updraft air diminishing its buoyancy.

• **The widest updrafts allow the updraft core to be protected**. Satellite imagery of the width of the cumulus, or radar imagery of the midlevel precipitation core width, are two ways to estimate which storm will have the least entrainment potential. Large Bounded Weak Echo Regions (BWER)s can be used to infer updraft size. Wide updrafts may also manifest themselves as areas of low spectrum width.

• **Secondary updrafts developing near a previous storm may grow in a more moist mid level environment than what the models or RAOBs indicate**.

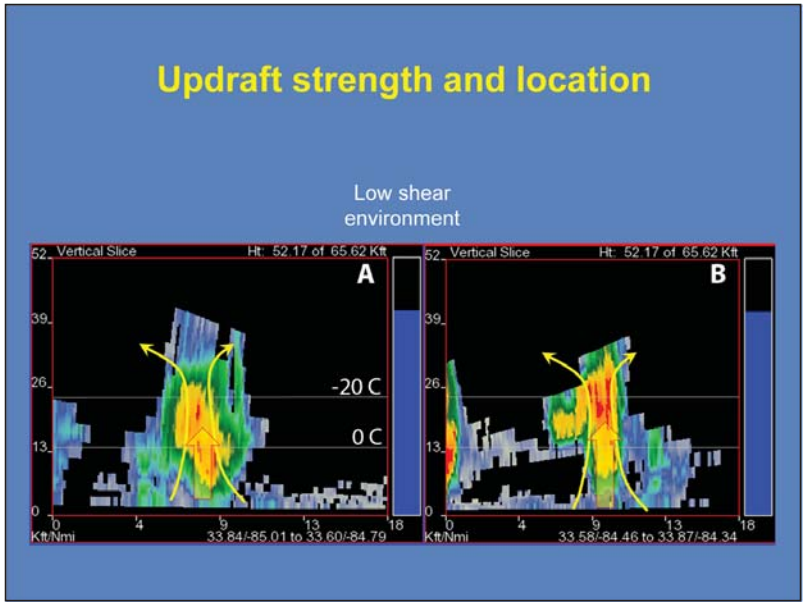
• **A large area of towering cumulus growing in a region of mesoscale ascent (e.g., a boundary) provides a clue that the environment will be more moist than analysis show**.

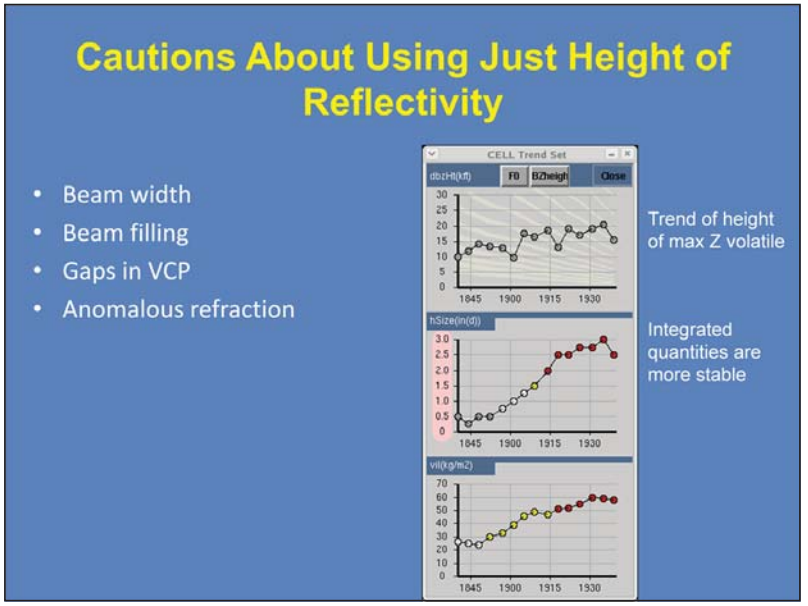
Can an updraft be stronger than accountable by parcel theory?



- Yes
- But it must contain significant vorticity

A very intense updraft can form in a relatively low updraft buoyancy environment if it is well correlated with significant vertical vorticity in mid-levels. As will be discussed in later lessons, a significant mid-level mesocyclone is occupied by a dynamic pressure perturbation pressure minimum that can significantly boost updraft strength. Some estimates based on numerical model studies suggest more than 50% of the updraft strength can be attributable to dynamic pressure forcing.





The most common technique for inferring an updraft's location is to observe the location of its upper-level reflectivity core as it reaches the maximum height in its lifecycle**.** Hydrometeor growth is maximized as the most intense part of the updraft passes through the -12° C to -20° C layer. Therefore, the highest reflectivity core in a layer centered just a bit higher should reveal the location of the strongest updraft.

This figure is a good example of a reflectivity signature of 40-50 dBZ surpassing the -20 ° C. Both storms were sampled by KFFC at the same stage in their development in the weakly sheared environment. However, note that the maximum reflectivity in the storm in B was at a higher altitude. It went on to produce a severe downburst in Atlanta while no severe reports were received from the storm in A.

So these examples point to what has been found before. **The intensity and altitude of the elevated core both increase as an updraft's intensity Increases**

As updraft intensity increases, the likelihood for intense downdrafts and large hail also increases given the same environment.

In fact, Cerniglia and Snyder (2002) noted that as the 55 dBZ reflectivity reaches higher altitudes, the False Alarm Ratio (FAR) decreases for some types of severe reports (wind or hail). Another study (Gerard, 1998) examined 64 storms from either the Jackson, MS CWA or the Cleveland, OH CWA, and found that those storms with 65 dBZ above the 0° C level were severe 96% of the time. At the time of these studies, the severe hail criterion was 0.75" in diameter.

The next slide is a video overview of the storm shown in the right panel.

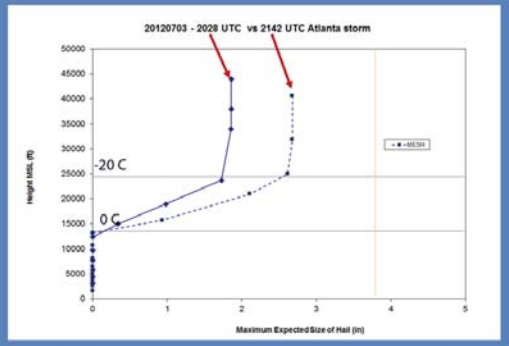
Though with any single height level radar observation, single height reflectivity guidance suffers from a rapid decrease in accuracy with increasing range owing to beam width, beam filling, gaps in VCP, and refractive index. Please be careful when you consider these radar-calculated heights in this lesson or in scientific literature. You will likely have more success comparing heights for storms at the similar ranges.

One way to reduce the impact of radar sampling limitations on estimating updraft intensity is to vertically integrate the vertical reflectivity profile of a storm from the freezing level to the top. Vertically integrated reflectivity is more resistant to changes in sampling from one volume scan to the next. Recall from the products topic that the Hail Detection Algorithm actually accomplishes such a task in its calculations (see Witt et al. 1998). The Severe Hail Index (SHI) is essentially a vertical integration of strong reflectivities above the freezing level.

This trend set shows how volatile the height of the maximum reflectivity in the storm can be relative to the VIL and maximum expected hail size (both integrated quantities). Take note that VIL and VIL density vertically integrates reflectivity through the whole storm depth, and so part of that reflectivity may be occupied by downdraft.

Young Severe Thunderstorm Reflectivity Profile

- View the whole Elevated Z profile gives an earlier heads-up to storm severity
- SHI profile reflects strength of updraft



MESH profile for the 2028 (2142) UTC cell shown as solid (dashed) curves respectively.

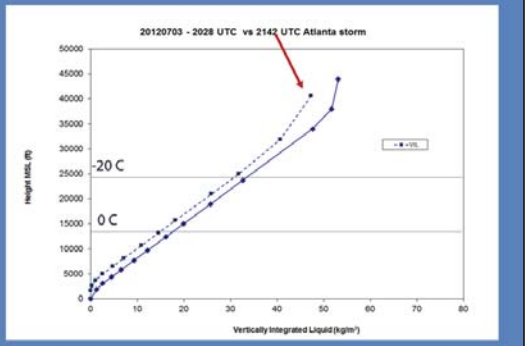
Reflectivity profile for the 2028 (2142) UTC cell shown as solid (dashed) curves respectively.

Contrast the reflectivity profile in this figure between two ordinary cells in their updraft-dominant phases. The dashed reflectivity profile of the cell that initiated at 2142 UTC showed higher values above the freezing level and lower values closer to the ground compared to the earlier storm at 2028 UTC.

The highest reflectivity for both storms was located above the freezing level. However the 2142 UTC cell had stronger reflectivities in the 0 to - 20° C layer. In fact, they helped to boost the Maximum Expected Size of Hail (MESH) much higher than the 2028 UTC cell. While these values were greatly overestimated for a weakly sheared and short-lived cell, the differences in the MESH helped to highlight the much stronger updraft in the 2142 UTC cell.

Young Severe Thunderstorm Reflectivity Profile

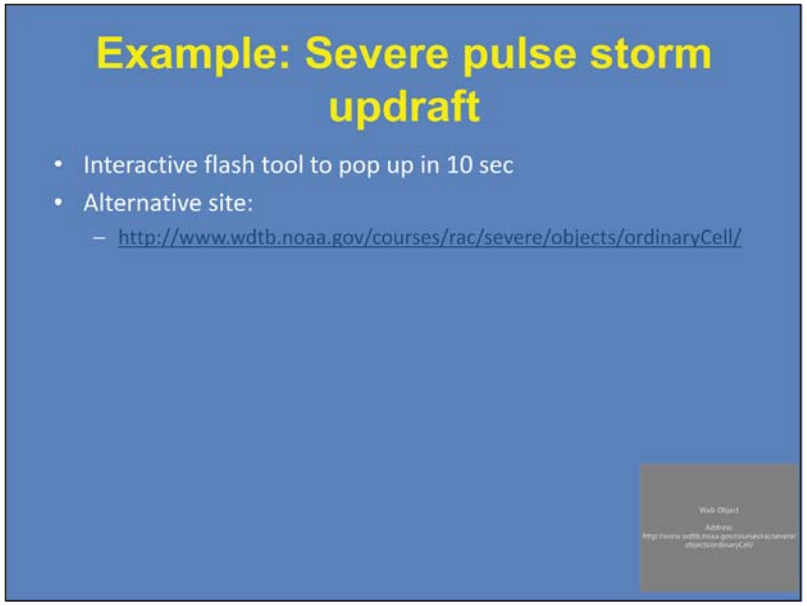
- VIL and VIL density <u>incorrectly assess</u> updraft strength



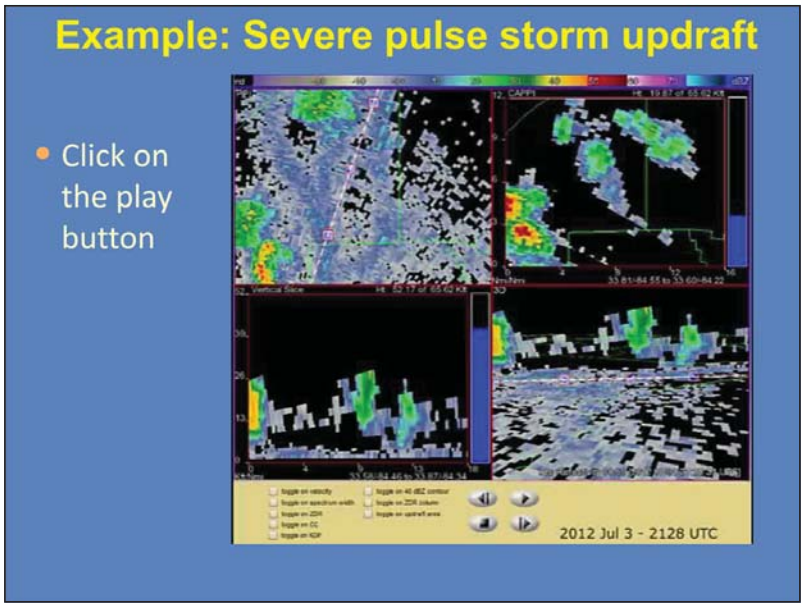
MESH profile for the 2028 (2142) UTC cell shown as solid (dashed) curves respectively.

Reflectivity profile for the 2028 (2142) UTC cell shown as solid (dashed) curves respectively.

The VIL failed to show a significant difference in value between these two cells in their updraft dominant phase. In fact, the cell with the weaker updraft had a slightly higher value of 51 kg/m^2. This was because the 2028 UTC scan of the cell showed more reflectivity at low levels. VIL density also showed higher values for the 2028 UTC cell. VIL is a great tool for showing the cells with the deepest and most intense reflectivities but it is not so great to evaluate which storm may have the strongest updraft.



In 10 seconds an interactive flash loop of the Atlanta, GA severe downburst producing pulse storm will appear. This is the same storm you saw in the last example but the loop starts at 2128 UTC.



I have an FSI loop of the Atlanta, GA storm where in the upper left panel lies the lowest elevation PPI, and the upper right shows the CAPPI at relevant altitudes. For reflectivity that's at about 20 kft or some just below -20 deg C. In the lower left is a cross-section and the lower right is the 3-D-like display with a cross section in the vertical and PPI in the horizontal.

Starting at 2128 UTC, the first sign of the significant updraft shows up as an area of 20-30 dBZ echoes starting around 13 kft (3 deg C) to above 26 kft (<-20 deg C). The CAPPI shows it nicely too. The lowest PPI shows two boundaries that just collided.

Clicking on velocity shows rather unimpressive convergence at the collision point but otherwise some faint hint of cloud top divergence in the cross-section.

If we were there we'd see an impressive towering cumulus cloud transitioning to a CB. The cell is all updraft.

Going to 2133 UTC, the reflectivity goes from around 30 dBZ to > 50 dBZ in just 5 minutes! The base of the core begins to descend to 13 kft and the top is brought upward to 35 kft. The PPI shows nothing yet but see how the CAPPI snags the core at near 20 kft. That's why the CAPPI is so useful.

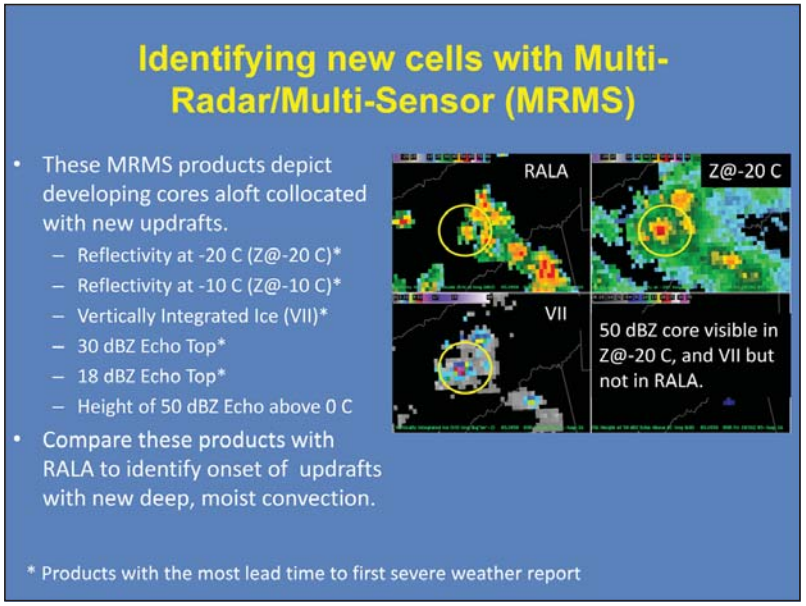
Velocity is again weak in the PPI and even in the CAPPI. I'm looking for signs of a downdraft to form. I don't quite see one yet as there is no elevated convergence. But look at the storm top divergence! We know the anvil is spreading out fast. And since there's no shear, the updraft is occupying the whole core. But things are about to change.

Five minutes later at 2138 UTC, the main body of the reflectivity core still hasn't reached the surface but the echoe has exceeded 65 dBZ up to 26 kft. It takes one intense updraft to develop such a strong echo. For a weakly sheared case, this

updraft is powerful. No doubt there's significant hail and the subsequent downdraft will likely be severe.

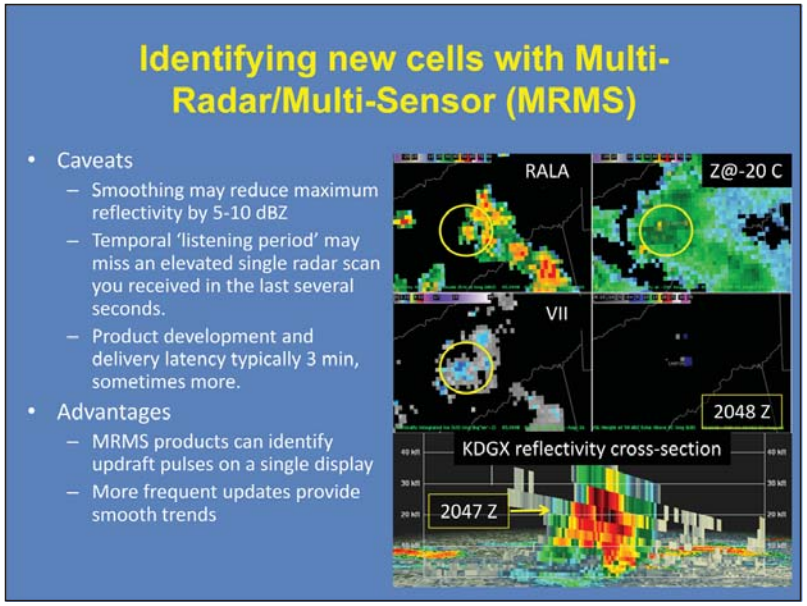
Velocity now begins to show an area of inbound on the far edge of the precip core centered around 14 kft, where the CAPPI is located. That's the beginning of the downdraft and now the updraft area is likely pushed to the south and up above the mid-altitude radial convergence, all the way to storm top.

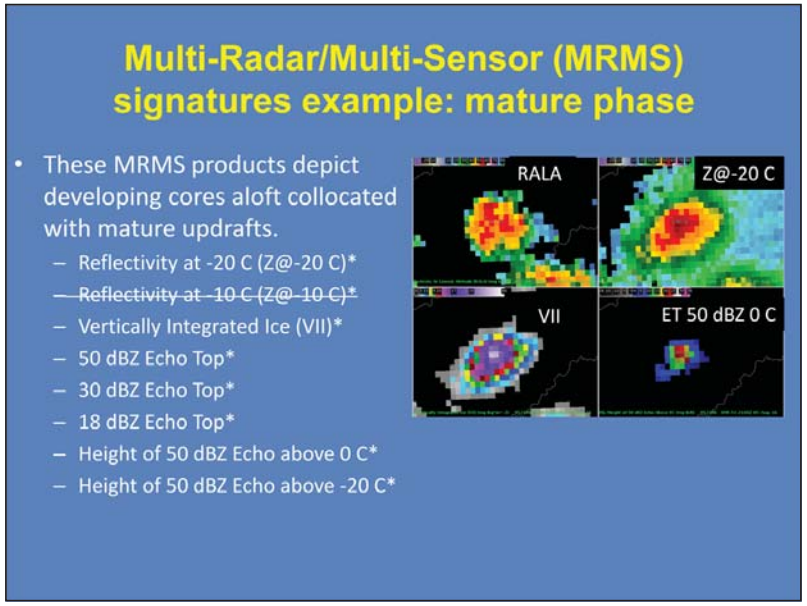
Going to 2142 -2147 UTC, the downdraft rushes to the ground but the updraft continues its ascent and diverges more quickly at anvil level. Notice the storm top divergence increase even more. But down below the -20 C level, the updraft is likely pushed further to the south or may even be dissipating.



Updrafts of new cells can be identified using the Multi-Radar/Multi-Sensor (MRMS) suite most readily using the isothermal reflectivity products, (e.g., Reflectivity at -20° and -10° C), Vertically Integrated Ice, the echo top products for 18 and 30 dBZ, and the height of the 50 dBZ echo top above 0° C. However waiting to identify a new updraft using the 50 dBZ echo top above 0° C may reduce your lead time to first lightning or severe weather report.

In the 4-panel display to the right, a new updraft created an elevated precipitation core intense enough to appear as a 50 dBZ echo in the Reflectivity at -20° C and 5 kg/m2 in VII before significant reflectivity appeared in the RALA product. Looping the MRMS helps a forecaster to quickly identify new updrafts using this display. Thus use these products with the RALA to identify new updrafts associated with deep, moist convection.

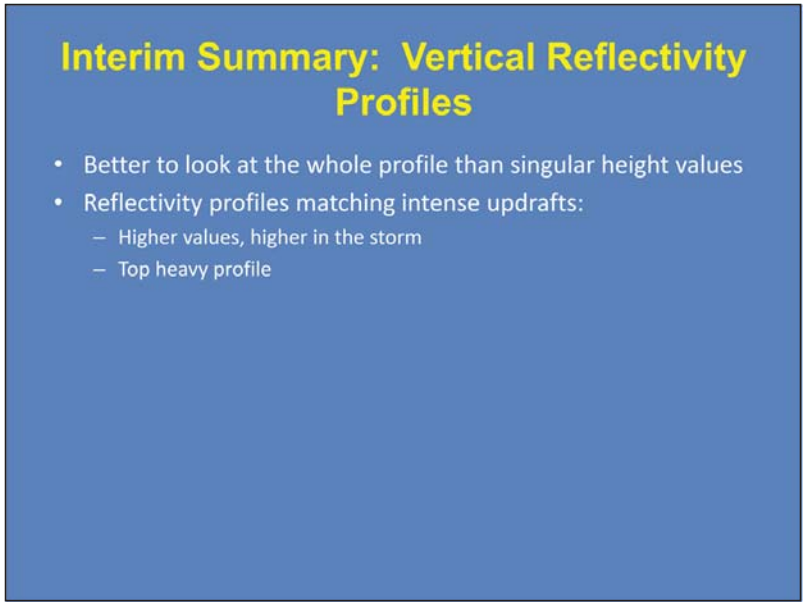




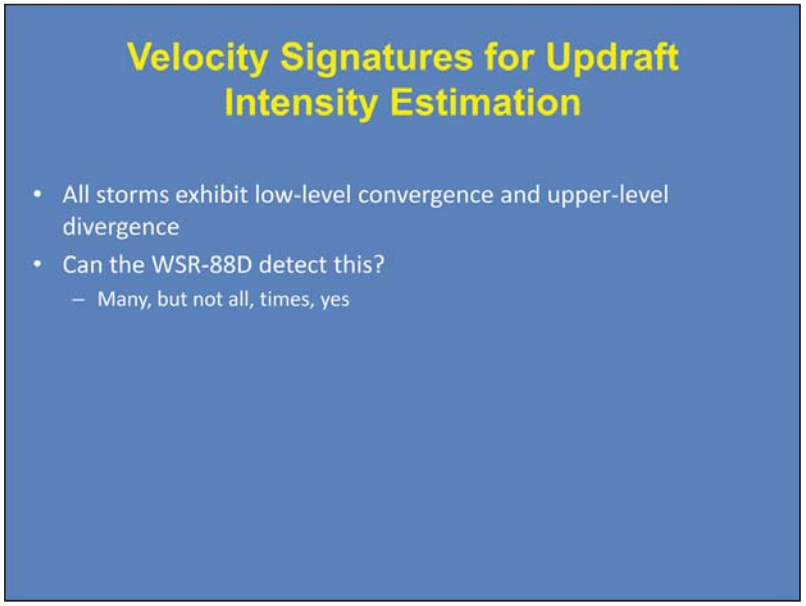
There are some caveats to using MRMS in identifying new cells. The reflectivity analysis smooths the highest peaks by sometimes up to 10 dBZ thus requiring a mental adjustments to thresholds you have set. Also, a new scan from a single radar may show you development of an elevated precipitation core but this new data may have missed the 'listening period' of the MRMS with the same time label as the single radar elevation scan. I illustrate this point here with a reflectivity cross-section from KDGX showing an elevated core where reflectivity exceeds 50 dBZ for the 2047 Z scan that cuts along the -20° C level. This data appeared in the 2050 Z MRMS reflectivity at -20° C but not at 2048 Z. Add a typical 3 minute latency, and it may be up to six minutes before you see the onset of this core. Thus there is one more reason to reduce your reflectivity threshold for identifying new updrafts.

Given these caveats, when you need a single display for identifying new thunderstorm updrafts, MRMS is the go-to-product suite. And the frequent updating means you can establish a trend more quickly, even if you have to reduce your thresholds.

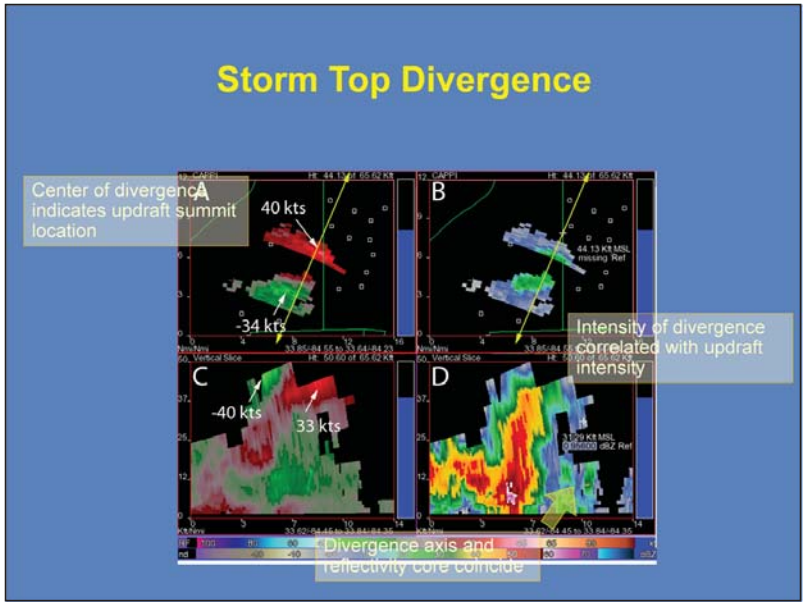
For mature thunderstorms, the split updraft/downdraft pattern means that the importance of the higher level products increases in locating updrafts in MRMS products. Now the reflectivity at -10° C may be occupied by downdraft so scratch this product. Add the echo top of the 50 dBZ echo above 0° and -20° C.

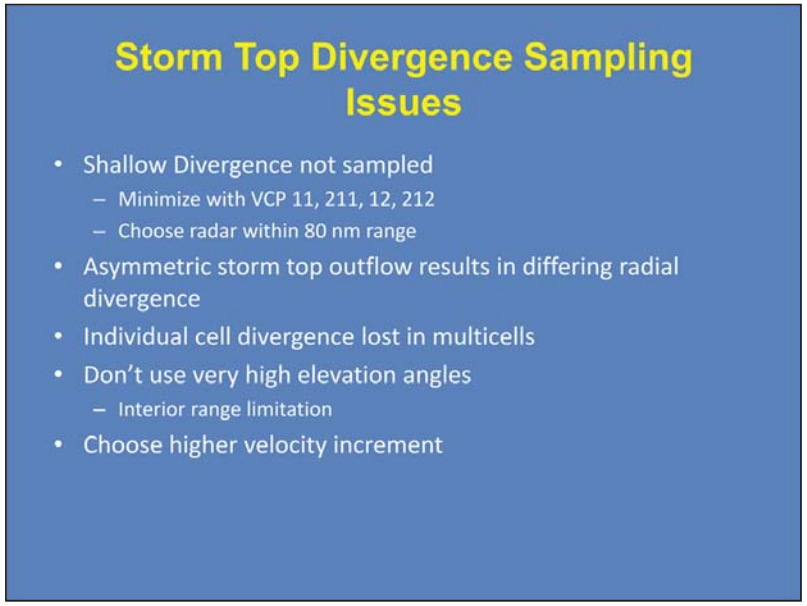


To summarize, the radar beam height uncertainties are just too great to allow you to hinge your warning decision on a single height threshold of reflectivity (e.g., height of the 55 dBZ echo). Instead, determine the shape and intensity of the reflectivity profile as it extends into the subfreezing layer in a storm as it's intensifying. Storms with a the most top heavy reflectivity profile (e.g., highest reflectivity at the highest levels) at this early stage are more likely to be severe than those with bottom heavy profiles. Additionally, storms with a top heavy reflectivity profile signifies the production of a core with large hail aloft before it descends to the surface.



All convective storms exhibit some amount of low-level convergence as air enters the updraft, and upper-level divergence near the equilibrium level as air diverges from the updraft. For severe storms, the updraft intensity is likely to be higher and so is the intensity of the convergence and divergence signatures. The ability of the WSR-88D to detect these signatures depends on how well the total convergence and divergence velocity patterns are reflected in the radial velocity-only patterns. In many cases, the WSR-88D is able to adequately sample convergence and divergence affording you another tool to evaluate updraft intensity.





Radar base velocity or storm-relative velocity data show a divergent flow pattern at the storm summit once the equilibrium level has been reached and an anvil begins to form. **The center of the divergence indicates the updraft summit location. The intensity of the divergence is positively correlated to the intensity of the updraft** (Witt and Nelson, 1990). The maximum inbounds and outbounds can be quite strong, exceeding 50 kts in both directions in the stronger storms. Note that the divergence axis and the reflectivity core roughly coincide. This storm top divergence (ΔV) was about 80 kts as determined from sampling the maximum and minimum velocities found around the overshooting top sampled by lifting and dropping the CAPPI and moving the vertical cross section east and west in FSI. What you see here is about 74 kts labeled.

There are sampling issues that may inhibit you from detecting the full divergence that a storm is generating:

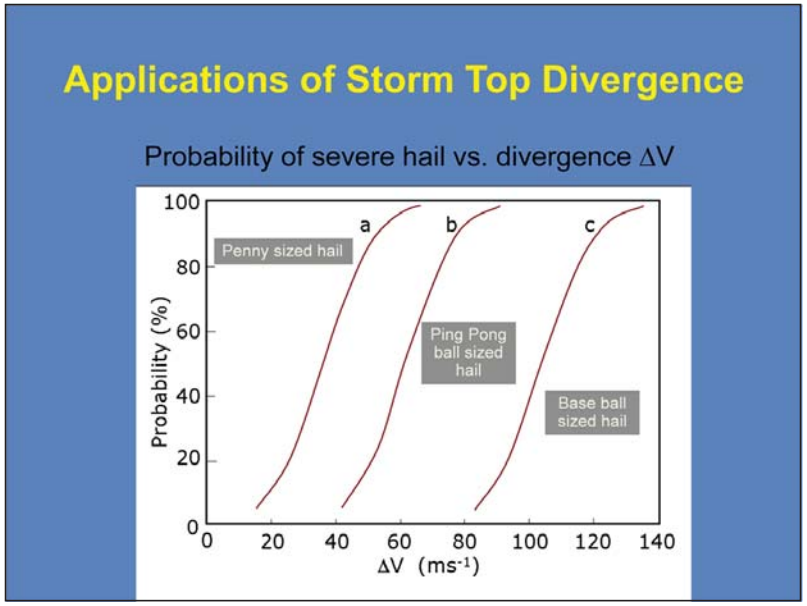
Shallow divergence may be missed by wide beams or VCP gaps. Beamwidths and/or gaps should be less than 2 km. To reduce its impact make sure you are using one of the 14 elevation angle VCPs (11, 211, 12, 212). Make sure you are sampling the storm inside of about 80 nm (120 km).

The storm top divergence, or outflow pattern may not be symmetric about the updraft summit. Thus radial divergence may change according to viewing angle. Try using an alternate radar for a better angle.

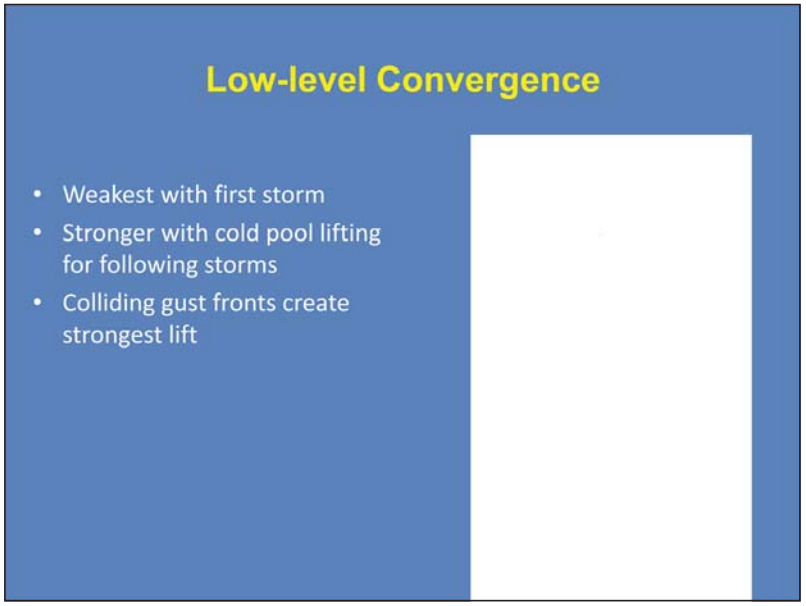
Individual storm top divergence signatures within large multicells may be difficult to detect.

Be careful about interpreting storm top divergence at very high elevation angles. Each individual beam may not sample both the minimum and maximum velocities. Try using the FSI CAPPI or a more distant radar.

Some anvils produce greater than 123 kts velocity. You will need to change the velocity increment to 2kts (1m/s) to adequately measure anvil divergence in these cases.

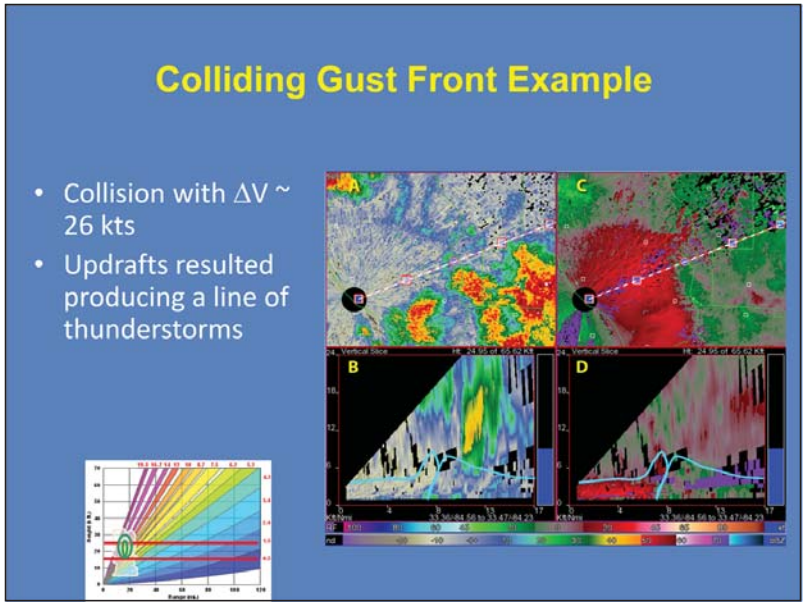


Updraft intensity estimations have been most closely tied to estimating hail size. Witt and Nelson, 1990 derived a useful correlation between maximum storm top divergence and probabilities of maximum hail size shown in this Figure . As a matter of caution, this graphic does not take into account the diversity of environments that you may face when estimating hail size.

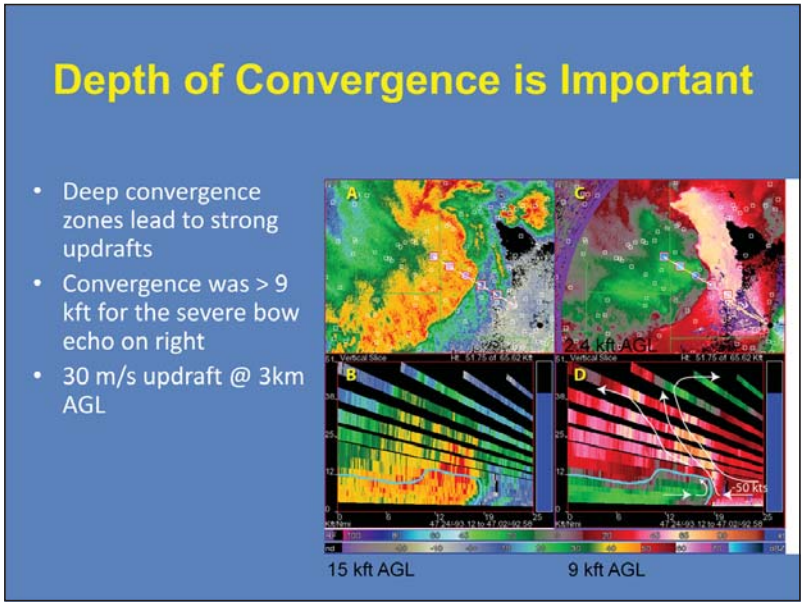


In addition to storm top divergence, you can attempt to detect the air converging into an updraft. During the initial stage of the first surface-based cell of the day, there may be a weak radial convergence feature within the lowest two kilometers of the ground as air flows in to feed the updraft. **Maximum radial velocities typically are very small for the first nonsevere cells of the day.** After cold pools develop, new surface-based updrafts receive much stronger initial baroclinic forcing along their edges leading to stronger updrafts.

Colliding cold pools, or cold pool interactions with other boundaries can generate vertical velocities exceeding 25 kts (12 m/s) within a few km above the surface. That's enough vertical velocity to generate graupel and split electric charges if the temperatures were cold enough.

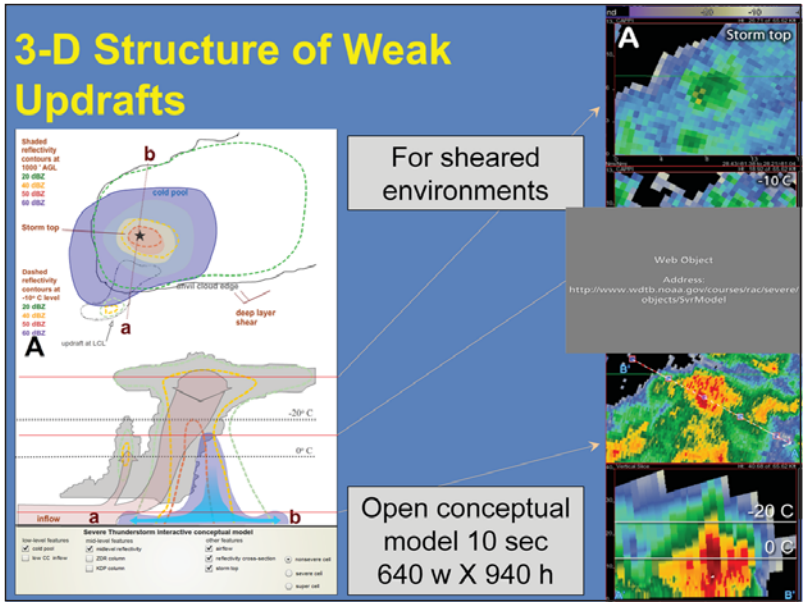


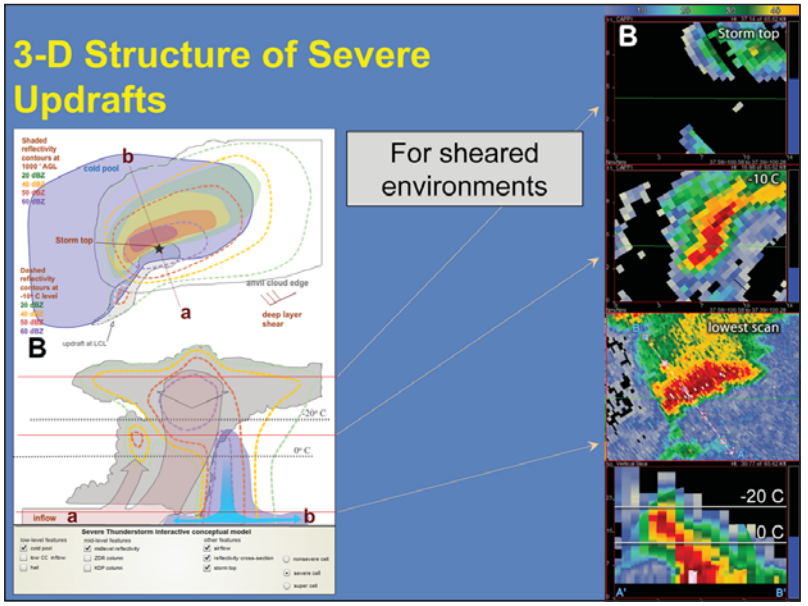
An example of colliding gust fronts initiating strong, but low-shear convection, is shown in this loop from just southeast of Atlanta GA. These gust fronts resulted in a rapid initiation of a line of broken ordinary cells, one of which we have seen before in this lesson.. The gust fronts collision resulted in approximately a ΔV ~ 20 kts over a 1 km distance.



For stronger storms, stronger low-level convergence is much more likely to develop as cold pool boundaries tend to be stronger and deeper. **It is not just the magnitude of the convergence but also the depth that is important to the strength of the updraft**, especially in the lower half of a storm. In the most severe multicell events, a cold pool leading edge may be up to 5 km (15kft) deep with more than 50 kts of velocity difference, leading to updraft strengths exceeding 70kts.

An example of a relatively deep convergence zone is highlighted in this image of a bow echo approaching Duluth. The gust front maintained strong convergence through a depth of at least 10 kft (3 km) with a DV ~ 20 kts across a boundary of 1 km wide (e.g., convergence > 0.01 s-1). By continuity, an updraft exceeding 30 m/s would occur at the top of the gust front's vertical interface.



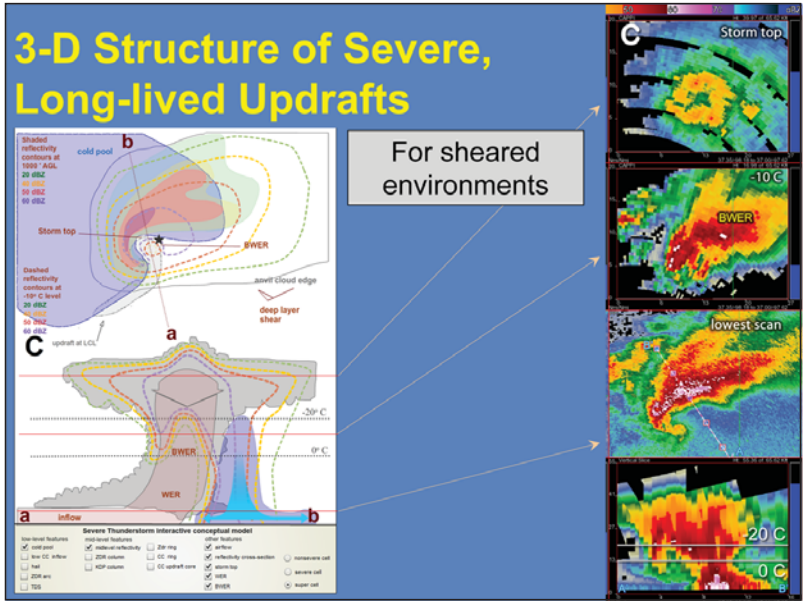


Before analyzing specific severe weather threats (e.g., hail and wind), it is important to take advantage of the volumetric radar data available to analyze the three-dimensional nature of storms. Lemon (1980), derived a methodology for volumetric discrimination between non-severe and severe convection. The technique focuses on the three-dimensional distribution of reflectivity through low-, mid- and upper-level indicators. Since updraft strength is an important controlling factor in severe weather, analyzing the 3-D shape of the reflectivity core is important. The following conceptual model is intended to represent convection as it evolves from nonsevere to severe modes in significant vertical shear environments. About a minute in, an external browser window will appear with an interactive version of this model. Follow along. This slide will stop at the end. The conceptual drawing on the left shows a horizontal planview on top and a vertical cross section on the bottom of a relatively weak updraft storm in a sheared environment. This storm can represent the onset of a severe storm to be, or be a mature storm that fails to utilize the instability and shear to its full potential. On the right, you are seeing a real example of a nonsevere storm as visualized in FSI. The cross section is displayed exactly as the conceptual drawing shows on the left with point A (B) facing toward (away) from the low-level inflow ingesting into the updraft. The shaded reflectivity shield in the conceptual drawing corresponds the low-level reflectivity shield and matches roughly to the scan second from the bottom on the right. The dashed reflectivity contours in the planview part of the schematic corresponds to the reflectivity map third from the bottom on the right and is roughly where the -20° C level should be located. The storm top reflectivity echo lies at the top of the figure on the right.

A weak updraft in a sheared environment slopes downwind and is typically unable to suspend any precipitation. Convergence at low-levels and the corresponding storm top divergence is relatively weak. Severe weather possibilities are relatively low with this kind of structure. An example of nonsevere convection in a sheared environment shows the least reflectivity overhang as the echo top lies mostly on top of the low-level reflectivity core. What overhang exists is mostly an artifact caused by the storm motion as the radar antenna ascended in elevation scans. Another example appears from KS on a high risk day in 2012 April 14 which shows similar structure.

As updraft strength increases, it becomes more erect and is able to suspend a heavy core of precipitation resulting in the Weak Echo Region (WER) (Fig. b). In this **Figure** , the reflectivity core shows obvious overhang above the WER in a direction facing the low-level storm-relative inflow and where the low-level core boundary exhibits concavity and a tight gradient. The echo top extends over the low-level reflectivity gradient next to the WER. The storm is more capable of producing severe weather. Large hail is quite likely when much of the echo overhang is above the freezing level.

Another example of a cell on 2012 April 14 in KS shows a similar structure of an echo overhanging the low-level WER on the side facing the low-level storm-relative inflow. The only difference is that the storm top significantly displaced downshear or into the plane of the cross section and so the echo appears somewhat shallow in the cross-section.



The most intense updrafts are erect and may exhibit a BWER . In this conceptual model, the BWER becomes more pronounced as an upward extension of the WER. The updraft in this storm is most likely stronger than in the non-supercellular severe storm model. The low-level reflectivity gradient facing the WER exhibits more curved concavity while the echo top may extend directly over the low-level WER and BWER. This reflectivity configuration is associated with the strongest updrafts, and the most severe weather reports.

The base of a severe updraft is typically located under the WER and BWER and next to strong reflectivity gradients and inflow notches. The WER is typically much larger than the saturated updraft; much of it exists because of intense reflectivities resident in the overlying anvil and the intense storm summit. You may use the low-level velocity to look for areas of strong convergence. The updraft base is most often rooted in the convergence.

The second example shows a more tornadic supercell where the BWER is a little more difficult to find at -10° C. The tornadogenesis may have helped weaken the updraft at the typical level of the BWER as the low-level rotation intensified and helped forced a rear flank downdraft.

Variations on 3-D Storm Structure

1. BWER in tropical cyclone supercells unlikely to infer large hail
2. Small BWERs more range limited
3. HPs and Bow echoes may exhibit front flank BWERs/WERs
4. Left moving storms exhibit WERs/BWERs on left side

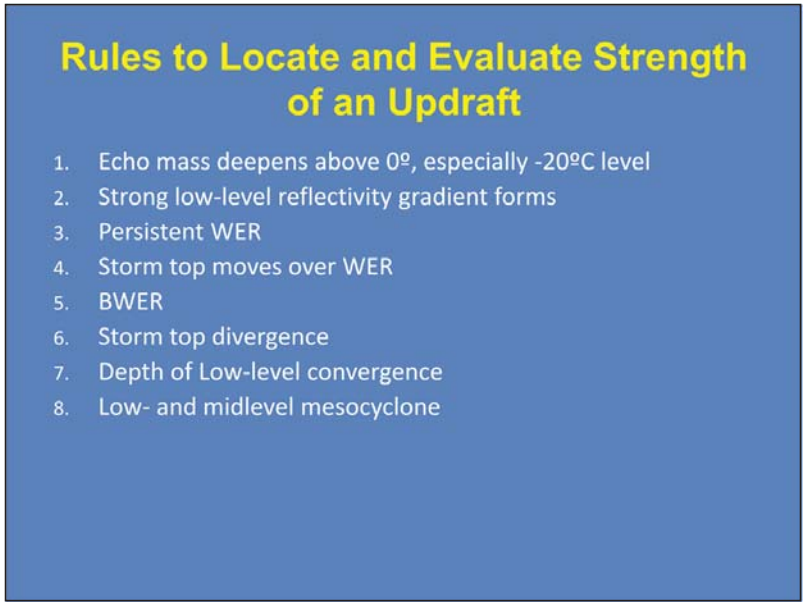
The updraft signatures discussed previously may result in different severe weather types depending on the storm environment. Here are a few examples:

A BWER in a nearly saturated, warm sounding (e.g., a tropical cyclone) is unlikely to infer the presence of large hail because the environment is too warm at the top of the BWER.

An environment with a low equilibrium level supporting mini supercells may indicate to the forecaster that BWERs may be too small to detect beyond a close range.

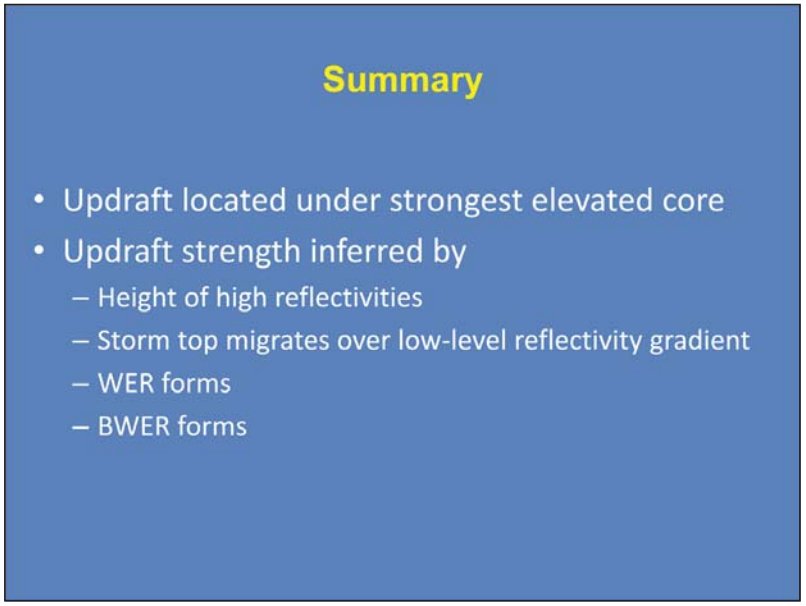
Many HP supercells and bow echoes may show WERs, and BWERs ahead of the main core with respect to the deep layer shear vector. In other words, these storms have front flank updrafts.

Straight or cyclonically curved hodograph environments favor left-moving storms with updraft signatures to the left front flank of the main core when you face its direction of motion.



Remember the following rules for locating and evaluating an intense updraft in a sheared environment:

1. Echo mass deepens above the freezing level, especially above the -20° C level.
2. Strong low-level reflectivity gradient develops.
3. A persistent strong echo overhang extends over the low-level, concave reflectivity gradient forming a WER.
4. The storm top moves over the lower level WER.
5. A BWER may form in the stronger updrafts as an upward extension of the WER.
6. Strong storm top divergence becomes strong.
7. Low-level convergence intensifies and becomes deep.
8. Low- to midlevel mesocyclone forms (not all of the mesocyclone is updraft).



For ordinary cell convection, the most generic proxy for locating an updraft is to locate the strongest elevated reflectivity core where it corresponds to temperatures less than -15° C. The higher you locate the core, the more likely it is dominated by updraft. The base of the updraft may begin with its roots in the boundary layer in the young phase of a cell. But over time, the base of the updraft becomes more elevated as intense core forces downdraft to form. As the storm nears its demise, the updraft may only be confined to the anvil layer.

The relative updraft strength can determine the maximum height achieved by the high reflectivities forming in the precipitation core. Generally the higher they form, the more likely there may be severe winds and/or hail.

Other reflectivity-based signatures come into play once the storm becomes more persistent. The more severe storm updrafts including severe multicells and individual supercells, exhibit a tendency for the high level echo core to migrate over a low level weak echo region in the vicinity of a tight reflectivity gradient. Sometimes the tight low-level reflectivity gradient becomes concave as low-level storm-relative inflow increases in response to the intensifying updraft. The size and extent of the Weak Echo Region (WER) increases with increasing severity of the updraft. BWERs form in the strongest updrafts.



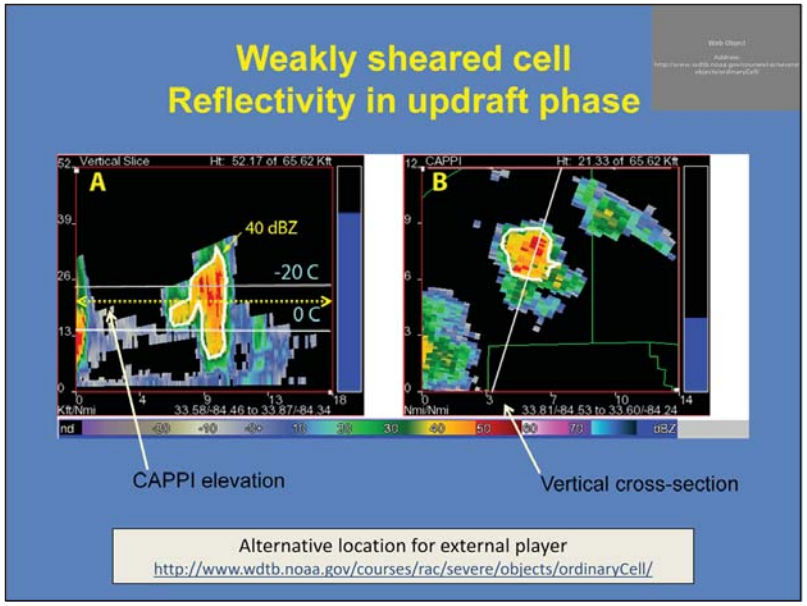
Welcome to the lesson on assessing updraft strength and location with polarimetric radar data. This lesson should last approximately 30 minutes.

Learning Objectives

- Analyze the location of a deep, moist convective updraft using polarimetric radar data using the following products:
  - ZDR, KDP, and CC
- Assess the relative strength of updrafts using all radar-based raw data

Analyze the location of a deep, moist convective updraft using polarimetric radar data using the following products: Differential Reflectivity (ZDR), Specific Differential Phase (KDP), and Correlation Coefficient (CC).

Assess the relative strength of updrafts using all radar-based raw data.

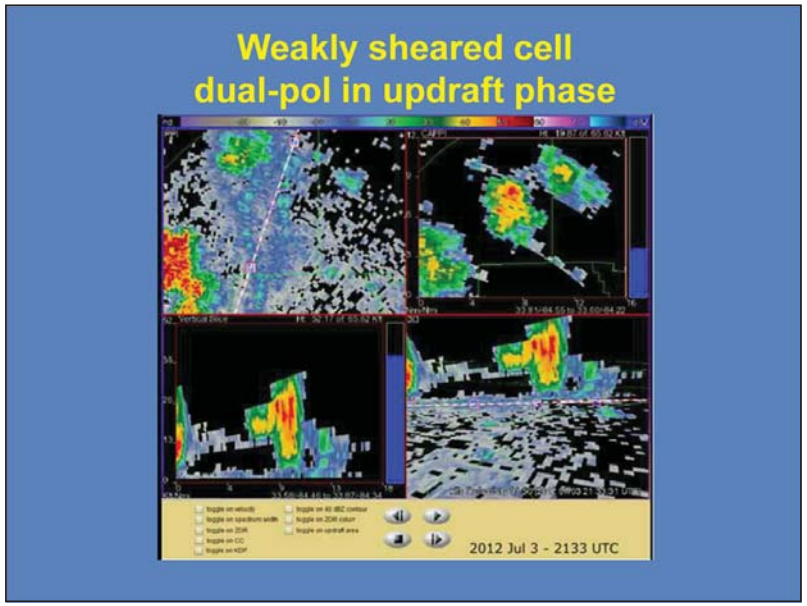


In a few seconds an interactive flash display will appear that you can use to follow along with my discussion on dual-pol updraft signatures of an ordinary cell. It is the same display you saw in an accompanying lesson on convective updraft identification, however there may be slight differences in the CAPPI elevation and vertical cross-section placement.

Capturing updraft signatures in weakly sheared ordinary cells depends on the timing of the radar's volume scans relative to the life cycle stage of the cell in question. The dual-pol signatures are transitory as the updraft bubble ascends toward its equilibrium level.

Let's start with viewing how these signatures **Updraft-dominant Phase** appear as the updraft pulse begins to produce precipitation. The case that we'll use is the 3 July 2012 Atlanta pulse severe thunderstorm case starting at 2133 UTC. This storm is quite typical in its appearance in the dual-pol products.

Keep the interactive display next to the articulate for the video walk-through of what's happening to the dual-pol products at this time in the next slide.



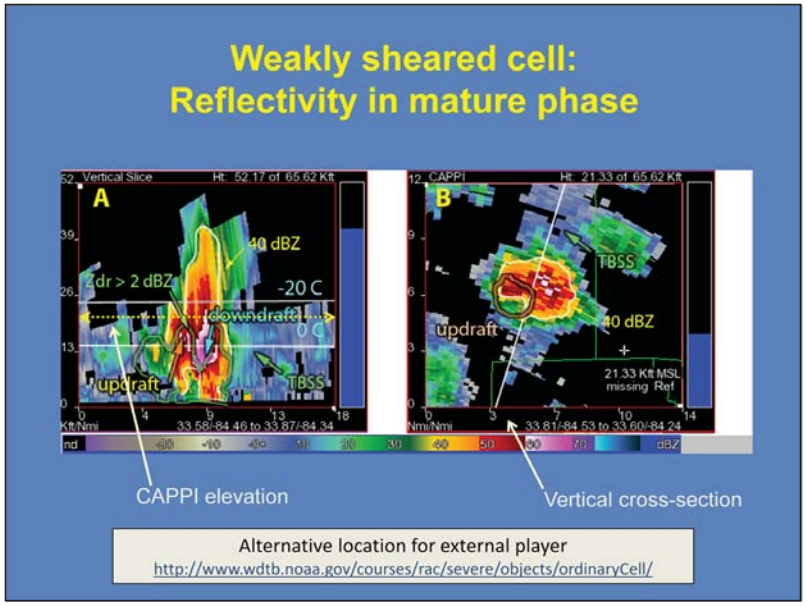
Click on the play button to start the video. You may follow along with the interactive viewer.

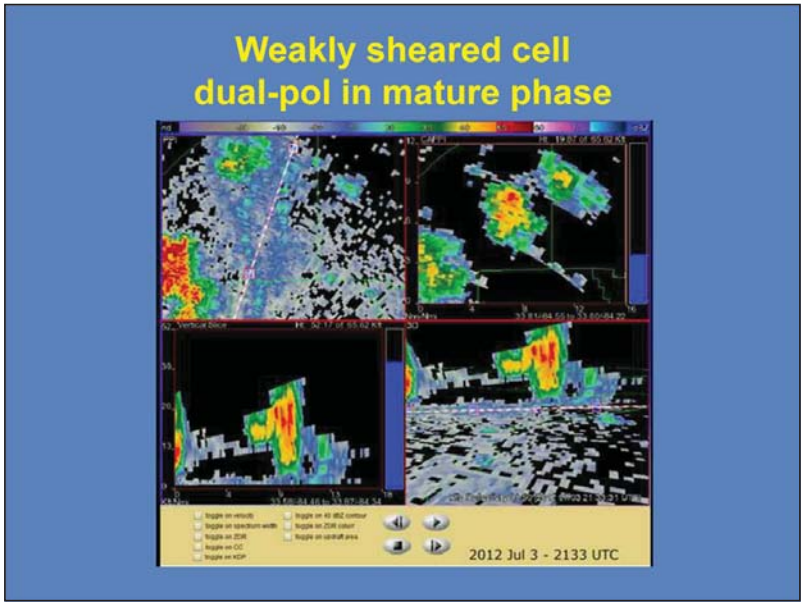
As the cell begins to produce precipitation, almost immediately, most of it is located above the freezing level (Figure 7-27A - Note from here on these are figure references in the student guide). Note that the ZDR shows large values in excess of 2 dB in the bottom portion of the reflectivity envelope that exceeds 40 dBZ. This is where large drops most likely dominate. Especially notable is the upward extension of likely large liquid drops above the freezing level (Figure 7-27C). The only way to get liquid drops above the environmental freezing level is through updraft. This is the area where we identify a ZDR column.

The vertical cross section shows a peak altitude of the ZDR column reaching about 18 kft above the radar (ARL), but the CAPPI (Figure 7-27D) shows the ZDR column extends up to its level at 21 kft ARL marked by a small area exceeding 2 dB. The vertical cross section missed the highest extent of the ZDR column. This is the area of strongest updraft within the broader updraft that is developing the echo. Note that the 2 dB value selected here is on the high side of a 1-2 dB threshold to consider the bounds of a ZDR column.

Where ZDR stands out with a large column exceeding 2 dBZ, the KDP remains small with only a small area exceeding 1deg/km near the freezing level (Figure 7-27E). This is because there is likely not much integrated water volume in the updraft. Only widely scattered large drops are lifting above the freezing level in the updraft.

The updraft within the ZDR column top exhibits slightly reduced CC values (Figure 7-27G and Figure 7-27H) that could be the result of a few ice particles forming amongst the large liquid drops. These values remain between 0.9 and 0.95.





Picking the location of an updraft through the depth of a mature ordinary cell is a bit more challenging. Now that the downdraft has begun, not all of the reflectivity core is occupied by updraft. Conventionally, we teach that the strong reflectivity core (> 40 dBZ) above the -20 deg C level is most likely occupied by updraft while that below this level is more than likely downdraft. Now with dual-pol data, we have the ability to better discriminate the location of the updraft.

Let's explore with another video tour in the next page. Meanwhile, keep your interactive graphic handy to follow along. When the next page loads, click the play button.

This video walks you through the dual-pol signatures of a weakly sheared cell in its mature phase. Hit the play button when you're ready.

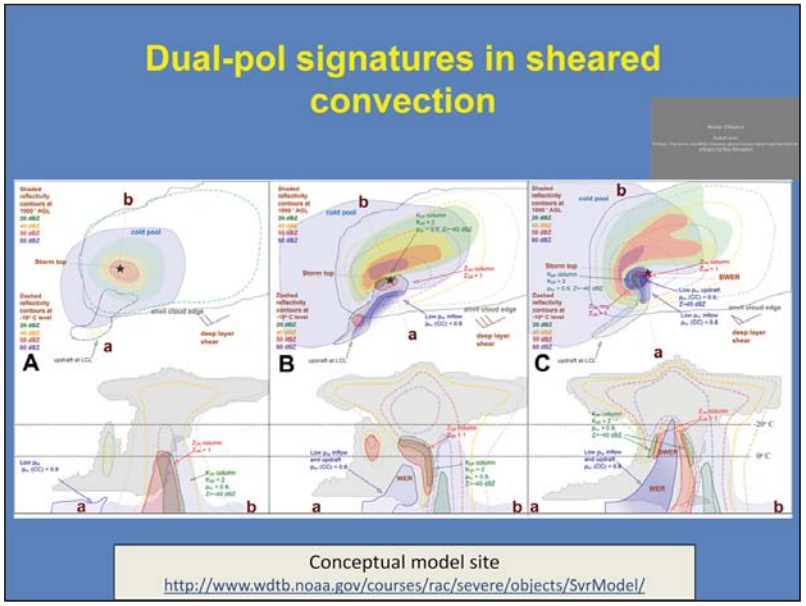
The onset of reflectivities > 55 dBZ above freezing level shown in Figure 7-28A is a strong signal that graupel and hail have formed 10 minutes after initiation (Figure 7-27). The ZDR images in Figure 7-28C and Figure 7-28D shows the depressed values < 1 DB that helps support the idea that the precipitation was dominated by ice. The downward plunge of the low ZDR precipitation core is quite likely associated with downdraft. At this time, the ZDR column has bifurcated somewhat with the primary updraft likely on the southwest flank of the precipitation core. The ZDR column there still reaches up to the 21 kft as the CAPPI in Figure 7-28D highlights.

Note that very high ZDR values exist down-radial from the precipitation core. This is a Three-Body Scatter Spike (TBSS) and not a ZDR column. Always suspect high ZDR values down-radial of a ZDR column when the reflectivity is low.

KDP indicates high concentrations of liquid water along the path of the radar beam. In Figure 7-28E the KDP is high (~2-3 deg/km) in the axis of low ZDR

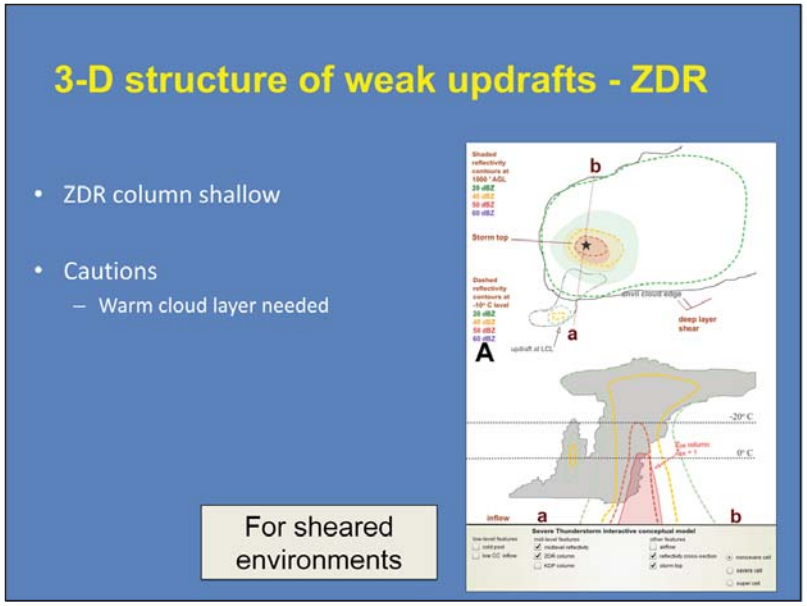
Values from straddling both sides of the freezing level. These high values mean there is a large amount of liquid water amidst the hail and graupel. How did so much liquid water wind up above the freezing level within what we believe is now downdraft? Two possibilities exist, the air may be downward moving but the temperature is still warm compared to the environment to melt hail or at least force drop shedding off of existing hail stones. Or perhaps the liquid water hasn't frozen from when they formed within the updraft in the previous 10 minutes. Most likely, the downdraft has just commenced, and the air is likely still
warmer than the environment. Thus, using the KDP column can give some clues that liquid water exists, but there is considerable question as to whether or not the air is still ascending. Note that the KDP is much lower closer to the 21 kft level and provides little information as to the location of the updraft (Figure 7-28F).

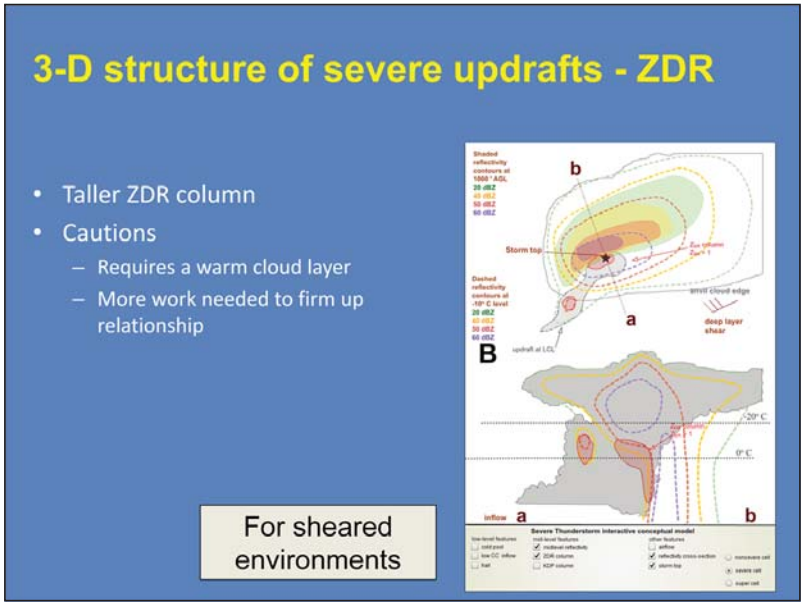
The CC is perhaps least associated with updraft. In Figure 7-28G, the ZDR column indicates somewhat depressed values (CC = 0.95-0.97) extending
above the freezing level, perhaps on two areas straddling the heavy precipitation core and downdraft. This may indicate some mixture of rain drop
sizes in these regions.



Instead of a short-lived pulse-like updraft, updrafts in sheared environments will exhibit a more steady state behavior. Indeed, there will always be multicellular behavior with updraft pulsing, but now at any one time in a storm's lifetime, updraft exists, either in a new cell or a mature one. You can use similar methods with dual-pol data to detect updraft location as with ordinary pulse cells. And there are some new signatures that appear in the strongest storms in a sheared environment. Now, let's add a dual-pol component to the reflectivity-based conceptual model of a non-severe, severe, and supercell thunderstorm in Figure 7-29.

The convective storm severity conceptual model will appear in a separate browser window. Use it for the up and coming discussion.





In the external conceptul model diagram, make sure you have the nonsevere cell radio button selected. Then select the following checkboxes: midlevel reflectivity, ZDR column, reflectivity cross-section, and storm top.

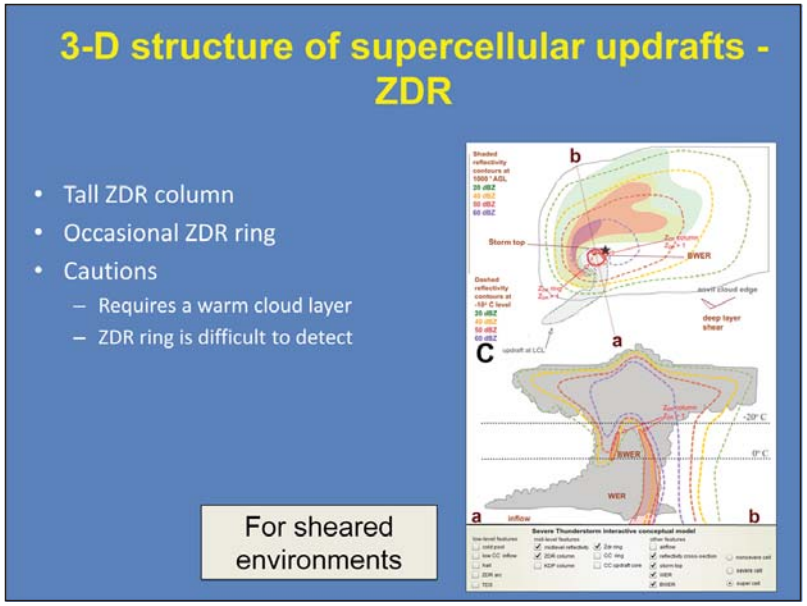
The ZDR column may appear for a longer duration than in an ordinary cell, yet the shape of the ZDR column may change. In an updraft that is weak, the ZDR column may only extend a few degrees above the freezing level (Figure 7-29A). It may also not exhibit any kind of overhang, just like the high reflectivity envelope.

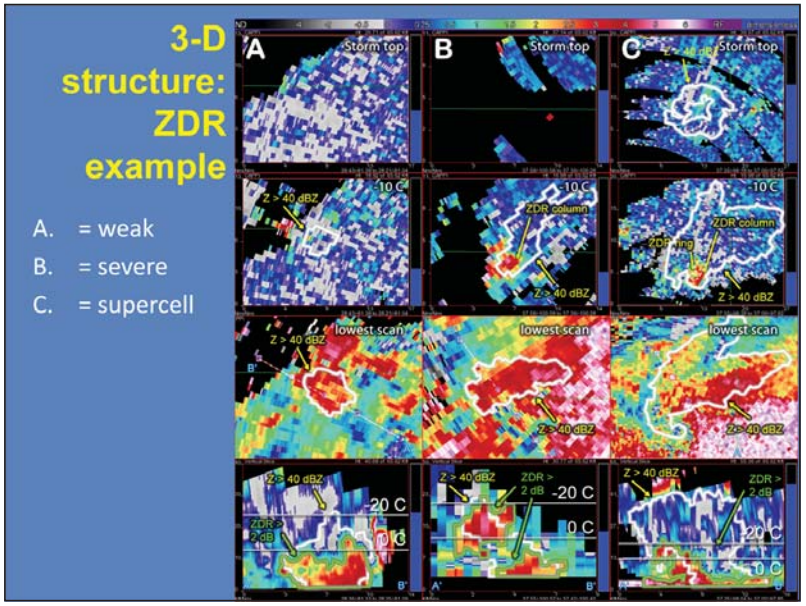
Remember that in order to get a ZDR column there must be a warm layer.

In the external conceptul model diagram, make sure you have the severe cell radio button selected. Then select the following checkboxes: midlevel reflectivity, ZDR column, reflectivity cross-section, and storm top.

As the updraft intensifies, the ZDR column expands upward and over the WER along the bottom of the strong reflectivity overhang (Figure 7-29B). The updraft and ZDR column are collocated.

Remember that in order to get a ZDR column there must be a warm layer. More work is needed to firm up guidance regarding the ZDR column and it's relationship to updraft severity.





Now the thunderstorm is strongly rotating and produces an updraft so strong that even large liquid rain drops don't have time to form by the time the air flows above the freezing level. Instead, what happens is that large drops reside along the periphery of the updraft into the sub-freezing air. This is why the ZDR column may lie on the periphery of the BWER (Figure 7-29C). The strong circulation may actually transport the large rain drops around the exterior of the updraft resulting in a ZDR ring. The ZDR ring, like the BWER itself, is often small and ephemeral meaning that poor radar sampling may prevent you from detecting many true events.

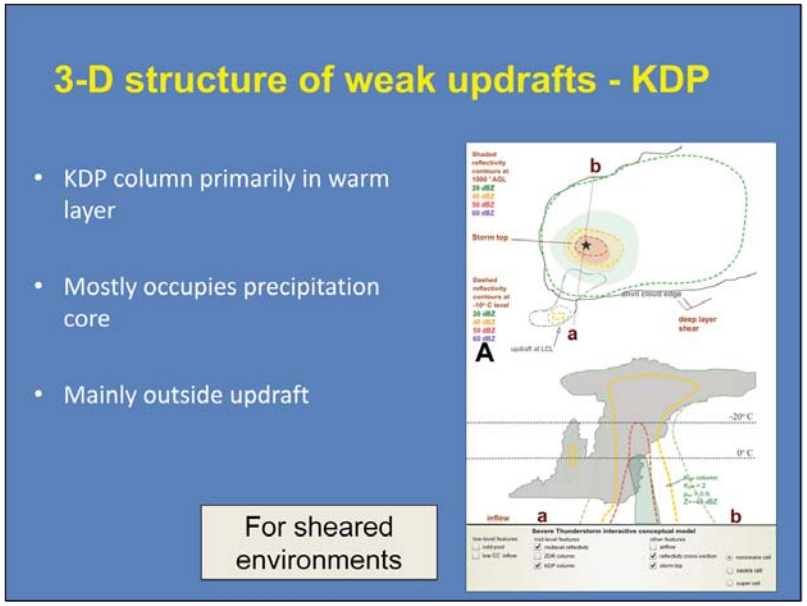
Remember that in order to get a ZDR column there must be a warm layer. More work is needed to firm up guidance regarding the ZDR column and it's relationship to updraft severity.

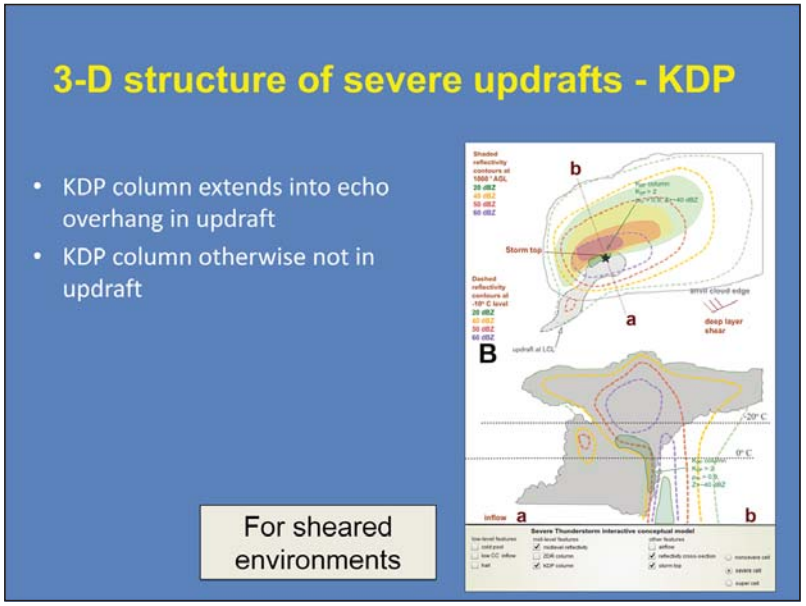
An example of how the conceptual model plays out with real storms appears in Figure 7-30.

In the non-severe storm case (Figure 7-30A), the ZDR column extends up close to the -10 deg C level on the western end of the updraft. At no place is there an overhang in the ZDR column.

The severe storm case in this column B (Figure 7-30B) shows a much taller ZDR column that extends above the -20 deg C level. The vertical cross-section reveals a substantial ZDR column overhang along the reflectivity extending over the low-level inflow and WER.

Finally, in the supercell example in Figure 7-30C, the ZDR column at -10 deg C is forced along the sides of the BWER. There is even indication of a ZDR ring around the lowest reflectivity portion of the BWER. This ZDR column extends upward to the -20 deg C level in the vertical cross-section and then descends down to the low-levels in the forward flank reflectivity gradient.



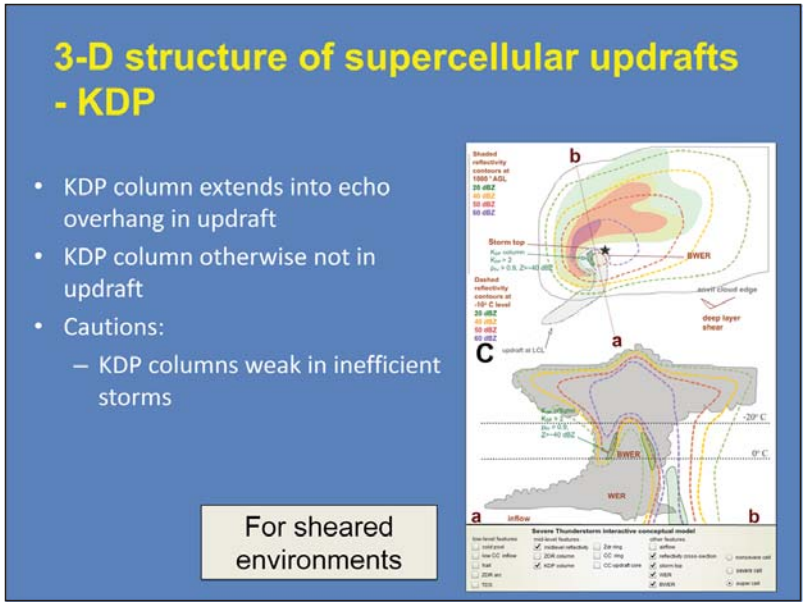


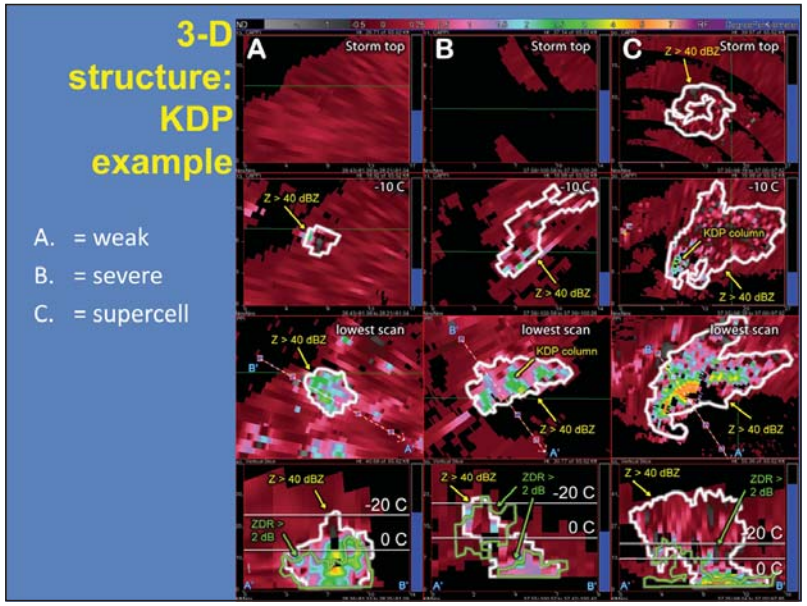
KDP columns do appear somewhat similar to ZDR columns, but there are differences that reflect the focus on integrated water content that KDP measures as opposed to the shape of large rain drops. In the non-severe storm conceptual model the KDP column extends up to just above the freezing level and it occupies more of the heavy precipitation where large volumes of rain and wet hail descend to the ground (Figure 7-30A).

Remember that in order to get a ZDR column there must be a warm layer.

In the severe thunderstorm case the KDP column extends upward and over the WER but at a slightly higher height than the ZDR column (Figure 7-30B).

In order for high KDP values to show, there needs to be larger quantities of liquid water. This area is likely to see that within the updraft, whereas the ZDR column may only be occupied by widely scattered large drops. This overhang of KDP often connects with the column of KDP extending to ground within the core of the storm. This cascade of high KDP falls outside of the main updraft. The higher extent of the KDP also can reflect the stronger updraft than with updraft in Figure 7-31A.

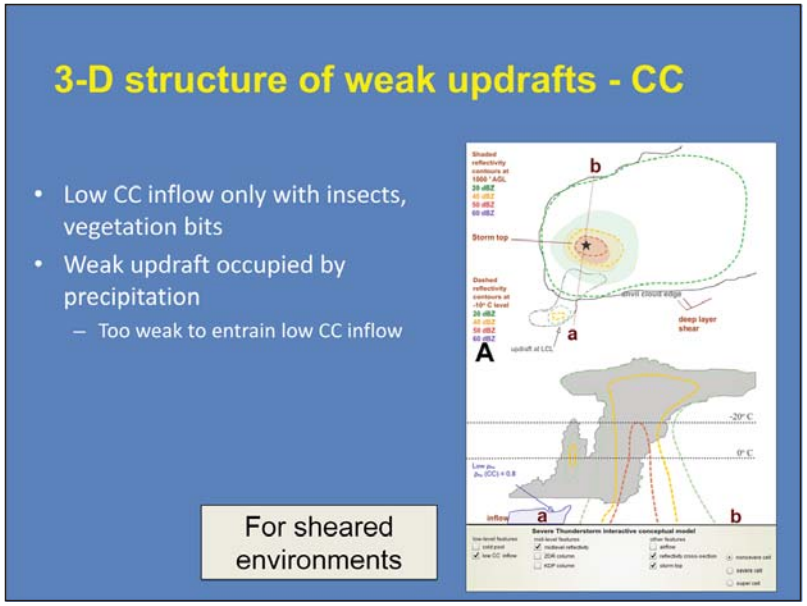


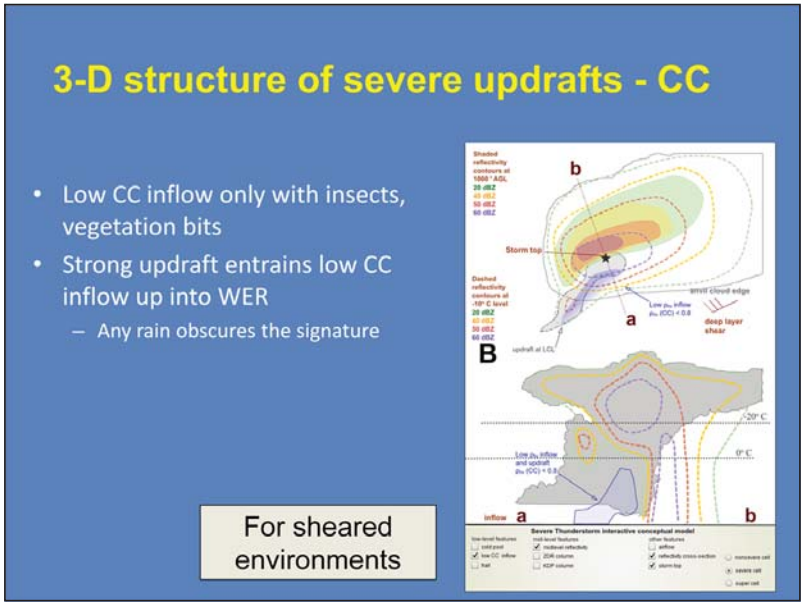


In the supercell, the KDP column also extends well above the freezing level. However, the highest values may lie on the upshear side of the BWER rather than the downshear location of the ZDR column. The KDP column extends into the core and descends to the ground in the heavy precipitation shield, often further into the core than the region of high ZDRs. Only in the highest parts of the KDP column would there be an association with the updraft. However much of the KDP, even at this altitude, may be further away from the updraft core than the ZDR column. Part of the reason why is that this region may be where the updraft is even weaker allowing more precipitation to fall out or be recycled. Yet, at an environmental temperature of -10 deg C, the presence of this much liquid water indicates that the updraft is still warming the area.

Sometimes, severe thunderstorms, especially supercells, may not exhibit much of a KDP column. This happens when severe storms produce a small amount of very large hail but not much volume of liquid or ice. Reflectivities may be high but KDPs stay low.

The same set of storms is now shown in KDP where the conceptual model is confirmed except for one case (Figure 7-32). Talking about where it is confirmed, note that the non-severe thunderstorm exhibits a substantial KDP column that rises almost to -10 deg C. Going to the severe and supercell thunderstorm cases, the KDP fields become more diminished aloft. Only down low do the KDP fields rise. Both of these storms appear to be dominated by drier hailstones. Nevertheless, at -10 deg C, there are regions of high KDP values, especially upshear (west) of the BWER in Figure 7-32B.

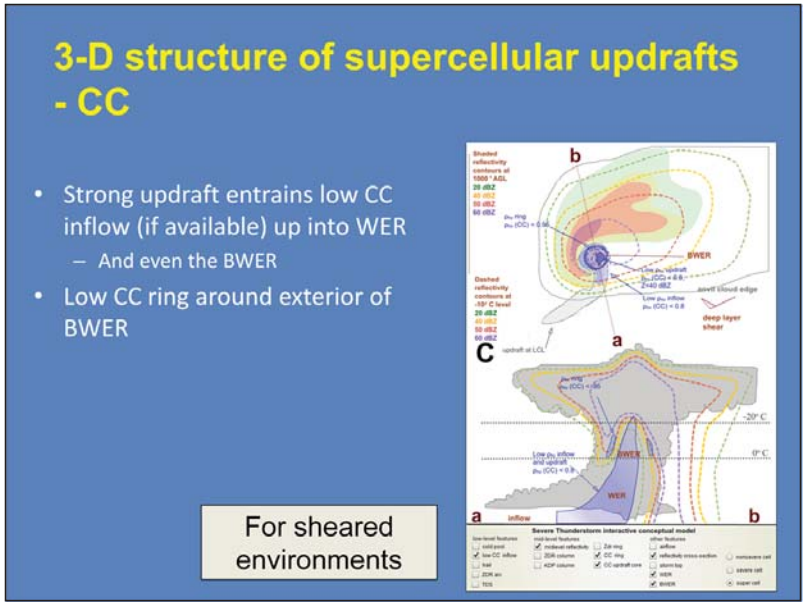




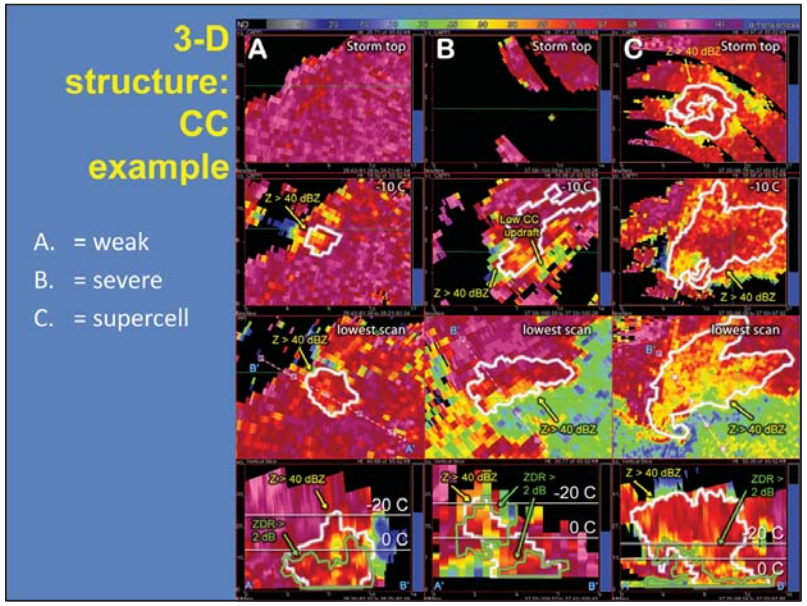
Correlation coefficient (CC) is perhaps one surprising tool to identify updrafts. Sometimes, thunderstorm updrafts show low values of CC (CC < 0.8) where reflectivity is very low. These low CC areas manifest themselves as an upward extension from the low CC clear echoes often found within a boundary layer occupied by flying insects or lofted light vegetative debris. Thus, the appearance of the low CC inflow and updraft is very dependent on the presence of these scatterers. It is also dependent on at least part of the updraft being free of any precipitation. Very small amounts of precipitation rapidly increase the CC and mask the detection of the non-meteorological scatterers.

In Figure 7-33, let's assume that the air is filled with insects or light vegetation and that the clear air boundary has a typical CC of less than 0.8. A non-severe thunderstorm typically sports an updraft too weak to result in an upward extension of the low CC air (Figure 7-33A). Most likely, the weak updraft is unable to produce a precipitation-free WER and the low CC signal is masked.

As the thunderstorm updraft intensifies, the likelihood of a precipitation-free updraft increases, allowing for the opportunity for the radar to detect the low CC echoes from the non-meteorological scatterers (Figure 7-33B) like insects and vegetation bits. Then the updraft can entrain this low CC inflow into the WER. Any rain in the updraft obscures this signal.



Strong supercells exhibiting a pronounced BWER present a challenge sometimes in that within the center of the BWER may show the weak CC signal from insects and/or debris as you can see from Figure 7-33C. But at the edges of the BWER, a low CC ring may appear just above the freezing level as precipitation encounters the melting between the environment and the updraft.



Finally, we have the appearance of CC with the same three storms shown in Figure 7-34. The non-severe case in Figure 7-34A did not have a low CC boundary layer owing to extensive light stratiform precipitation. As a side note, there was a TBSS confined to the subfreezing air. This was perhaps due to hail production after an updraft pulse. No severe hail was reported, however.

In Figure 7-34B there was a low CC boundary layer with values falling below 0.7 at times. Some of that low CC air was actually being entrained upward into the updraft within the WER high enough to be detected at the -10 deg C level adjacent to the precipitation core. In the supercell case in Figure 7-34C, there is substantially low CC in the low-level inflow, but the radar cannot detect the non-meteorological scatterers up to the -10 deg C due to light precipitation filling in the WER.

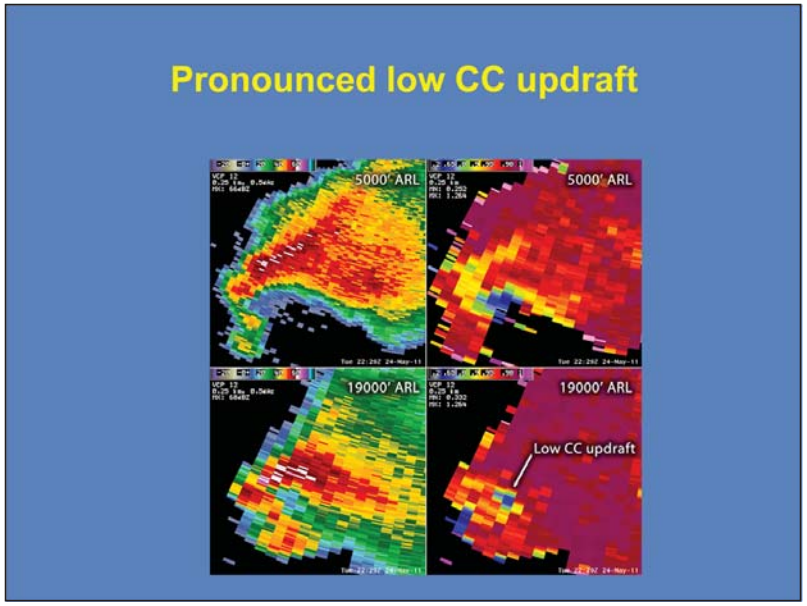
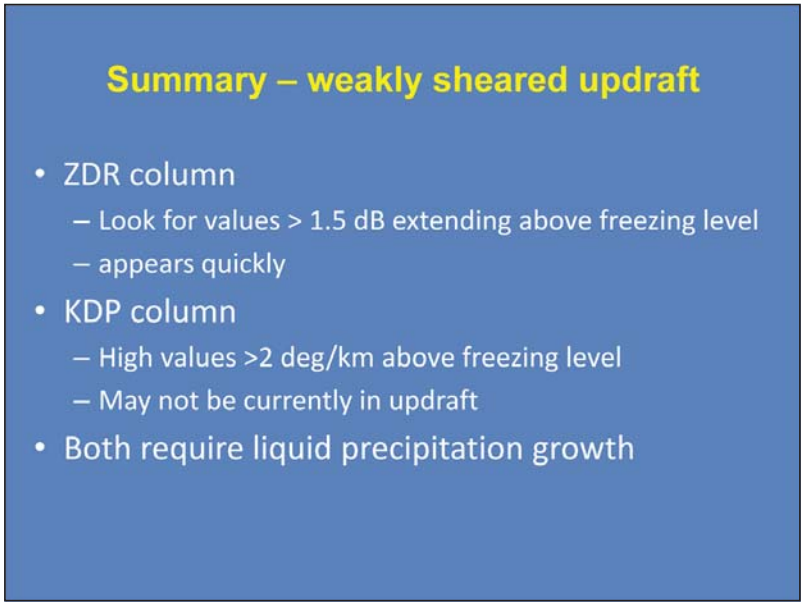


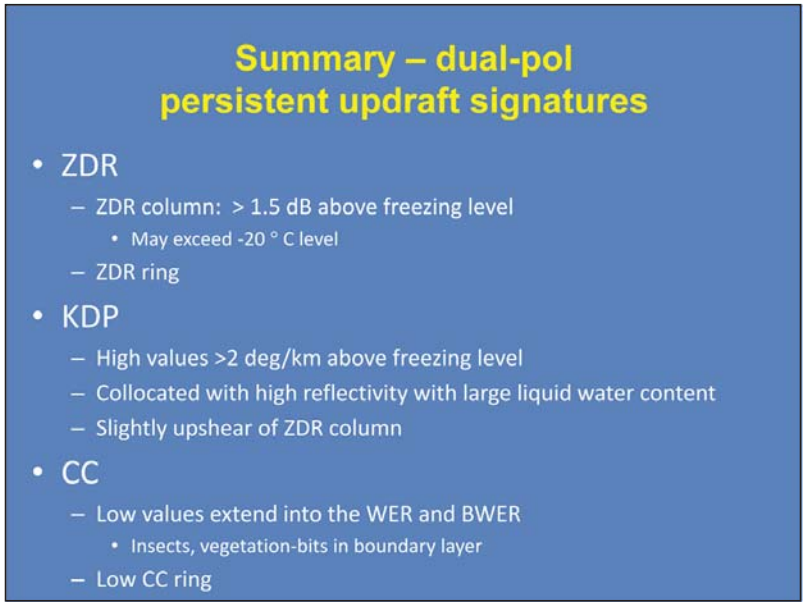
Figure 7-35 shows an example of a supercell with a prominent low CC updraft where the low values coincide well with the inflow notch at low levels extending into the BWER in the subfreezing air.

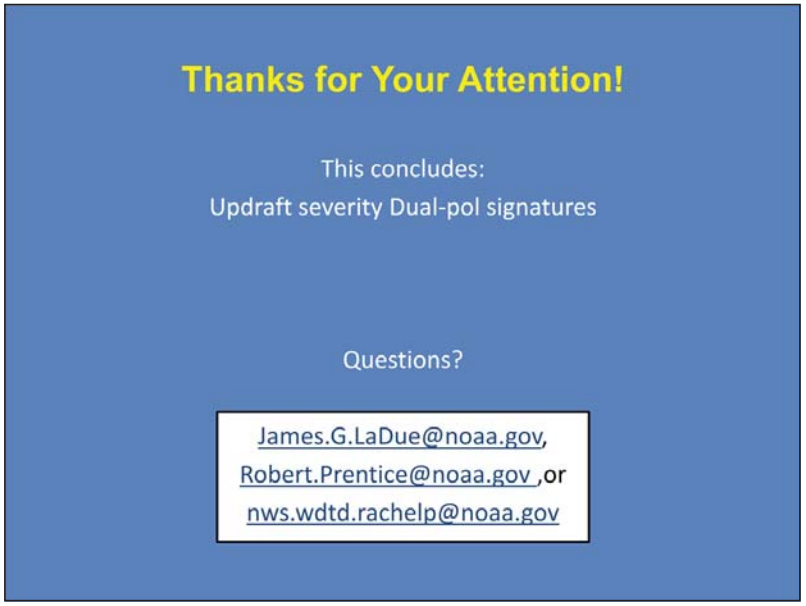


ZDR values in excess of 1.5 dB extending from low-levels up to and above the freezing level. The ZDR column will appear as soon as the precipitation core develops and is the best dual-pol updraft signature.

A KDP column appears after the ZDR column as the liquid water content increases in the updraft above the freezing level. However, by the time this happens, part of the column may be occupied by downdraft.

For both ZDR and KDP, there needs to be a warm cloud layer to enable liquid precipitation growth within the updraft.





ZDR

- Look for ZDR columns to extend into higher altitudes for the stronger updrafts in a given environment. Some ZDR columns may extend to -20 deg C.
- Strongly rotating updrafts may produce a high ZDR ring surrounding a BWER.

KDP

- The KDP column neighbors the ZDR column and will produce highest values where reflectivity is high.
- In supercells, the KDP column is often displaced slightly upshear of the ZDR column above the freezing level.
- KDP columns may not appear in some severe storm updrafts if they produce mostly large dry hail.

CC

- Low CC columns may appear in severe thunderstorm updrafts if the following conditions are met:
- Presence of low CC non-meteorological scatterers in the boundary layer (e.g., insects and/or light vegetation debris)
- Relatively weak reflectivities in the updraft.
- A low CC ring may appear in the periphery of the updraft just above the freezing level as a result of mixed phase precipitation.

If you have passed the quiz, then you have successfully completed this lesson. If you have any questions, please contact us using any of the e-mail addresses listed on the bottom of the slide.



Welcome to Single Cell Downburst Detection, a module in the Convective Storm Structure and Evolution topic in the Radar and Applications Course.

Learning Objectives

1. Identify the environmental and storm signatures favorable for dry and wet microbursts.
2. Identify favorable environmental and radar signatures of a high wind threat from supercells.

These are the learning objectives for this lesson. The end of lesson test includes some questions based on these objectives.

What is a Downdraft?

- Pulse downdraft spectrum
  - ➢Microbursts < 4 km < Downbursts
- Dry (<.01 inches, cell reflectivities < 35 dBZ)
- Wet (>.01 inches, cell reflectivities > 35 dBZ)



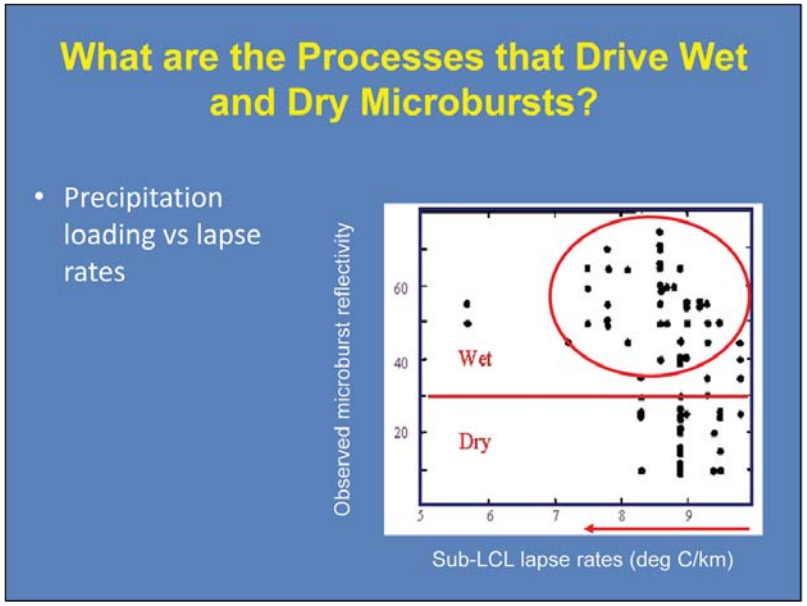
Wet downburst near Norman, OK on June 14, 2011

The downdrafts discussed in this section are typically on the same scale as the individual ordinary cell updraft. Downbursts and microbursts are outflows of an ordinary cell downdraft. The only difference is that downbursts are considered outflows larger than 4 km in diameter while microbursts refer to outflows less than 4 km in diameter.

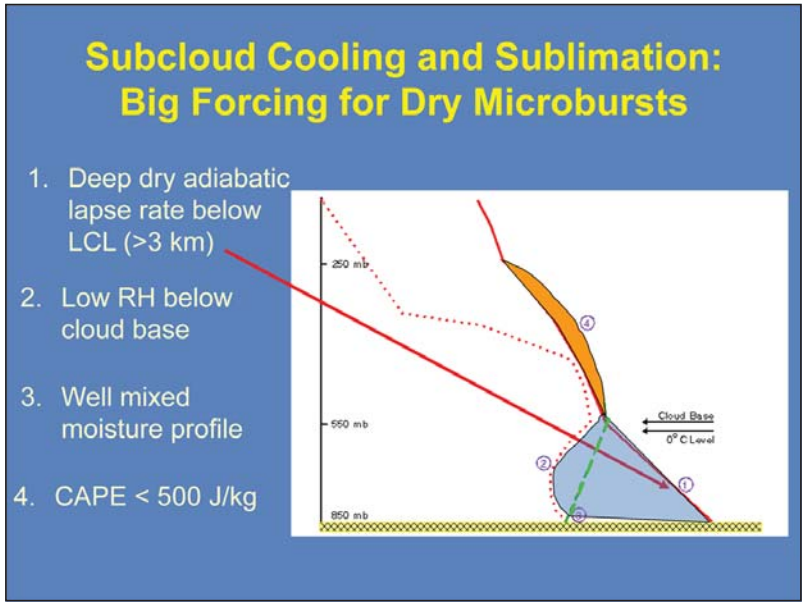
What Drives Downbursts?

1. Lateral dry air entrainment (lower Theta-E in mid layers)
2. Subcloud cooling (forcing for dry microbursts)
3. Sublimation (LCL < 0 deg)
4. Precip loading (Z > 45 dBZ)
5. Rear-Flank Downdraft (RFD) forcing

These are the primary factors that drive downburst processes: 1) Lateral dry air entrainment, which is measured by lower equivalent potential (theta-E) temperature in mid layers, 2) Subcloud cooling (which is the forcing for dry microbursts), 3) Sublimation, which occurs when the LCL is below freezing, and 4) Precipitation loading (which occur when lapse rates drop below 8 deg K/km). This factor (process) is observed with microbursts with a descending precipitation core of > 45 dBZ. Lastly, the fifth factor is Rear Flank Downdraft (non-hydrostatic) forcing in supercells. All of these factors will be covered in subsequent sections of this lesson.

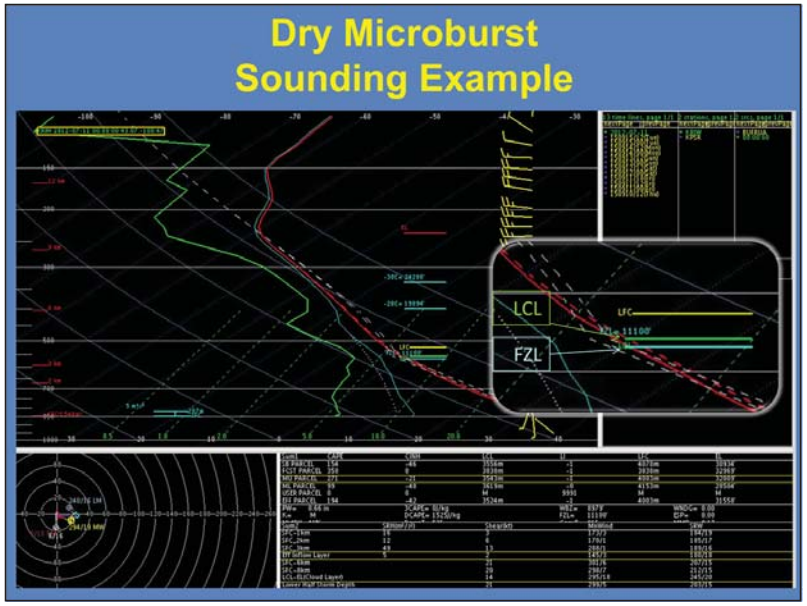


Taking a look at two of the primary processes that drive wet and dry microbursts, the primary factor is negative buoyancy and what forces it. Two forces at play are precipitation loading and lapse rates. They are both related. Let's look at this figure from Srivastava. As lapse rates decrease, downdrafts have an increasingly difficult time of maintaining their descent based on negative buoyancy effects alone. Heavy precipitation (especially those with reflectivity cores greater then 40 dBZ), can force the descent of a downdraft even if it loses negative buoyancy.



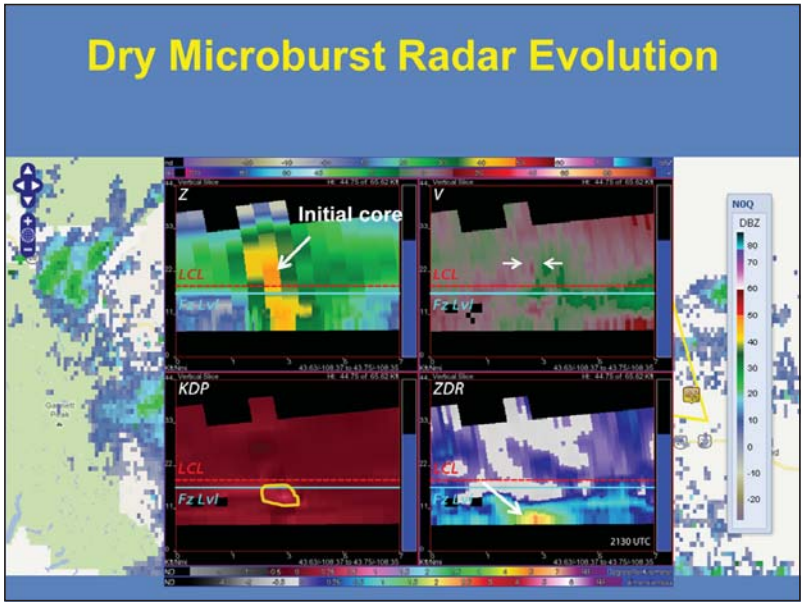
Dry microbursts are forced by evaporating precipitation below the LCL. These events are most common in the arid or semi-arid regions where LCLs are at least 3 km AGL. There are four basic characteristics of dry microbursts (kind of like clues in the environmental data: 1) A deep, dry adiabatic lapse rate below LCL, 2) Low relative humidity below the cloud base, 3) A well-mixed moisture profile (you can see the constant mixing ratio line from the surface to the LCL, and 4) Weak CAPE usually 500 j/kg or less with moist midlevels above the LFC. The typical dry microburst sounding is termed an "inverted V".

In these type of environments, the updrafts are relatively weak so precip loading and/or lateral entrainment are not major factors in contributing to downdraft strength. If the LCL is below freezing, the precipitation that does form initially cascades down as snowflakes which maximizes the surface area exposed to dry air.



We are going to show you an example of a typical dry microburst event. Here's the plotted RAP analysis sounding taken at Riverton, WY (KRIW) on 00 UTC July 11, 2012.

Note the LCL height is below freezing and there lies a deep, dry adiabatic layer extending to the ground with very little CAPE.



Next is the storm scale evolution of a microburst as seen by a multi-product cross section product from KRIW. There were two storms initially just northeast of the radar. The second one shown here is the one that became severe. You can see the warning polygon out of WFO Riverton. The microburst-producing storm initially produced an elevated core structure, as shown in the static 4-panel cross section at 2130 UTC (this is figure 7-40 in the Student Guide). Note the development of a 30-35 dBZ core 15-24 kft above the surface. This core of frozen precipitation eventually descended and sublimated just below the LCL. Radial convergence was weak (25-30 kts) as is typical for most dry microbursts. The ZDR column was apparent extending down to the surface and simultaneously weakening slightly prior to microburst impact. KDP showed larger values in the core indicative of a some intense rainfall. The loop which will pop up in a separate window when you advance to the next slide shows the core descending through the melting layer. Monitoring the descent of the reflectivity core helps provide some lead time of a dry microburst.

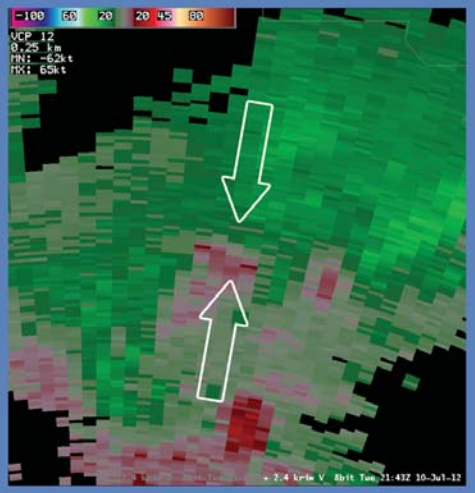
Flash Animation of a Dry Microburst

Web Object

Address:
http://www.wdtb.noaa.gov/courses/rac/severe/objects/RIW_DryMicro_FSI.html

This is a flash animation of an evolving dry microburst from KRIW.

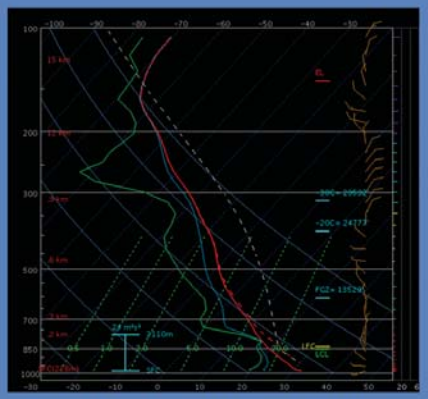
Weak Convergence



This is a velocity image of the dry microburst producing storm showing weak convergence near and just below cloud base (~ 20, 000 ft) indicating the downdraft has initiated. Radial convergence is weak (25-30 kts), as is typical for most dry microbursts, and more often than not, will be not be definable due to radar sampling limitations or flow tangential to the beam.

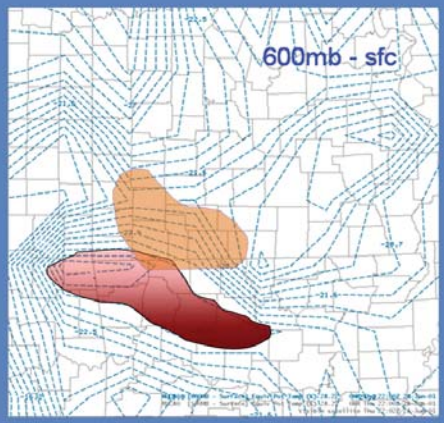
Forcing for Wet Microbursts: Lateral Dry Air Entrainment

- Midlevel dry entrainment and precip loading
- Large (> 25 deg K) theta-E differences from SFC to midlevels
- "Steep" lapse rates below LCL



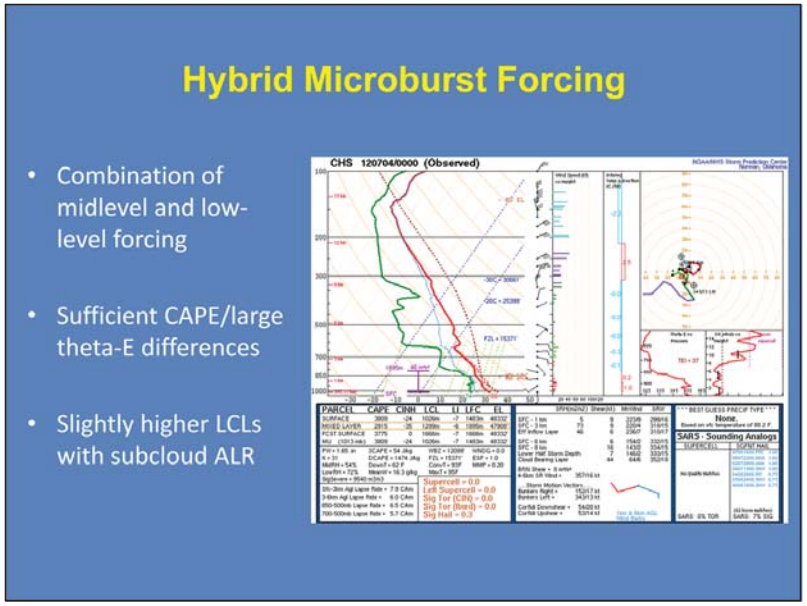
FFC RAP Sounding at 1800 UTC 07/03/12

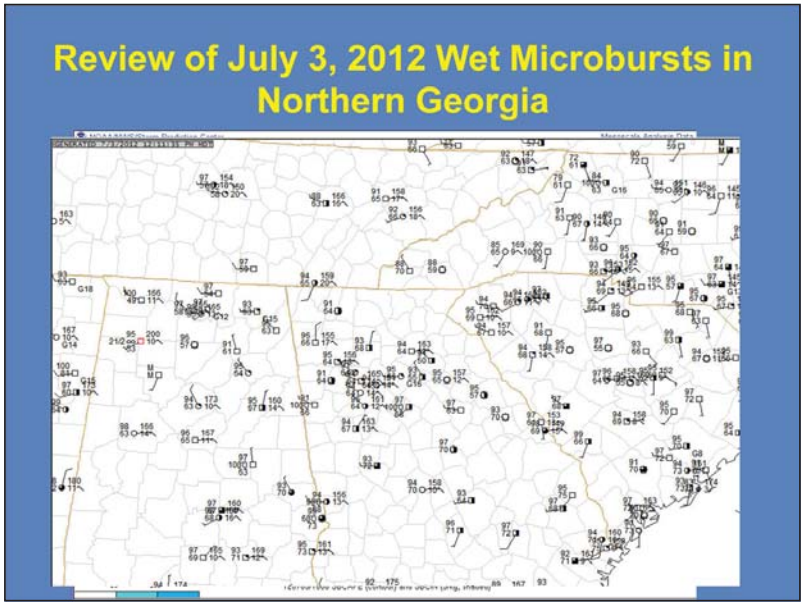
AWIPS Display of Theta-E Differences for Wet Microbursts



Wet microbursts are forced by midlevel entrainment and precip loading. Large (> 25 deg K) theta-E differences from the surface to midlevels are observed on days with wet microbursts. A "steep" lapse rate below the LCL is also important to allow the strength of the downdraft to reach the surface with strong outflow. However, the height of the LCL is not as important as with dry microburst cases. In this sounding from Atlanta from July of 2012, there was a SBCAPE of 2583 J/kg, no CINH, 4 kts of shear from 0-6 km, and substantial mid level dry air to support wet microburst generation. For reference, the avg. RH in midlevels of 63%. We are going to examine this case in more detail in a bit

The AWIPS volume browser allows you to subtract theta-E from two layers. In this figure, we see the isolines computed as theta-E difference between 600mb and the surface for a wet microburst event in Southwest Missouri. Values exceeding 25 deg K are shaded in red. The areas where significant microbursts occurred are shaded in blue. As is often the case, the risk area is not always confined to the maxima of Theta-E difference but throughout the gradient regions.

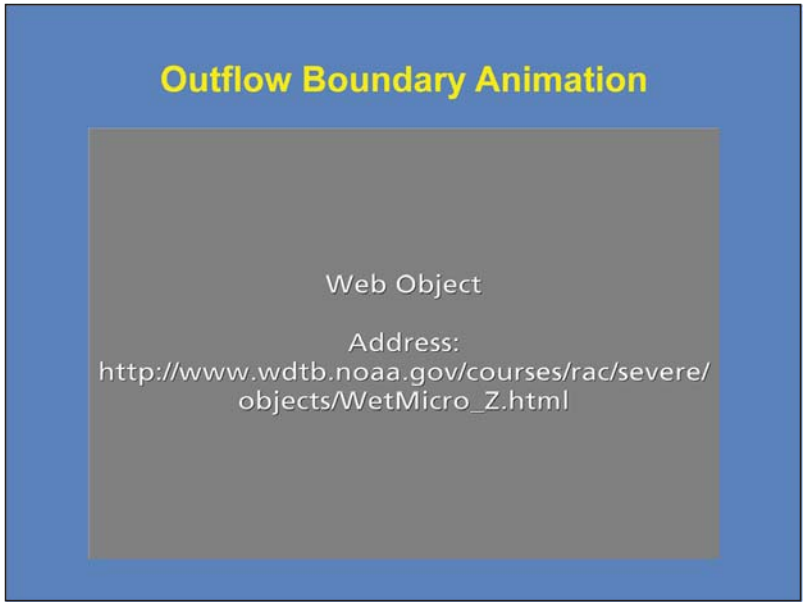




Hybrid microburst environments display characteristics of both dry and wet microbursts. Environments that support hybrid microbursts usually have sufficient CAPE, large theta-E differences from surface to mid-levels, and slightly higher than normal LCLs with adiabatic subcloud lapse rates. In a sense, hybrid microburst environments represent the middle of the microburst sounding spectrum. It could be said that most microburst soundings are hybrids.

This observed SPC NSHARP sounding from Charleston, SC was taken on the same day (July 3, 2012) that several microburst producing storms occurred all across GA, SC, NC, and VA. Note the theta-E vs Pressure plot right below the hodograph. There was an observed difference of 37 deg C from around 600 mb to the surface. This is a great environment for hybrid and wet microbursts.

A surface plot at 2011 UTC showed a large area of temperatures in the mid-90s and dew point temperatures in the mid-60s to around 70F in portions of central and northern Georgia. There was very weak surface convergence but with moderate instability increasing throughout the afternoon of 3 July 2012, forcing was accentuated by numerous outflow boundaries generated by thunderstorms across northern Georgia. You are going to see these outflows in motion when you advance to the next slide.



The window that has just popped up displays an animation of 0.5 deg Reflectivity from 2042Z to 2238z from KFFC (Peachtree City radar). As this was a weak shear environment, a number of thunderstorm outflow boundaries with large boundary motions were apparent, many of moved toward the radar. By 2105 UTC, which coincides with Figure 7-47 in the Student Guide, there was increasing thunderstorm development along and behind the outflow boundaries, especially to the east of KFFC. Some of the biggest storm intersections occurred along colliding outflow boundaries east of the radar from 2115z to 2205z. Next we will show a FSI cross section at 2142z through a storm that developed NE of the radar and produced severe winds.



The loop now playing shows a series of FSI cross-sections of Z, V, ZDR, and KDP from 2134 to 2210 UTC through a storm of interest northeast of the radar. Radar signatures show some of the signs of an impending microburst. First, you can see a large elevated core of high reflectivity (> 60dBZ) from 13-26 kft which had a TBSS. Then, the core begins to descend as the velocity shows some mid-level radial convergence as air flows into the top of the downdraft. About this same time (2142z) , there is a core of values of ZDR near zero indicating some hail in the downdraft column. If you toggle over to the KDP product, you can see high values of KDP (3-5 deg/km) in the downdraft core, indicating a high concentration of liquid water and potential hail accompanied the microburst. It is speculated that the large increase of KDP values prior to microburst may be due to the addition of hail meltwater to the rain already falling within the downdraft.

What are the Storm Precursors of a Wet/Hybrid Microburst?

- Stronger than normal initial updraft
  - Long lead time, relatively high FAR
- Descent of strong core/Mid-altitude radial convergence increase
  - Short lead time (< 10 min), moderate FAR
- Storm collapse
  - Short lead time (< 10 min), but high FAR!
- Decrease in ZDR (if hail is also occurring)

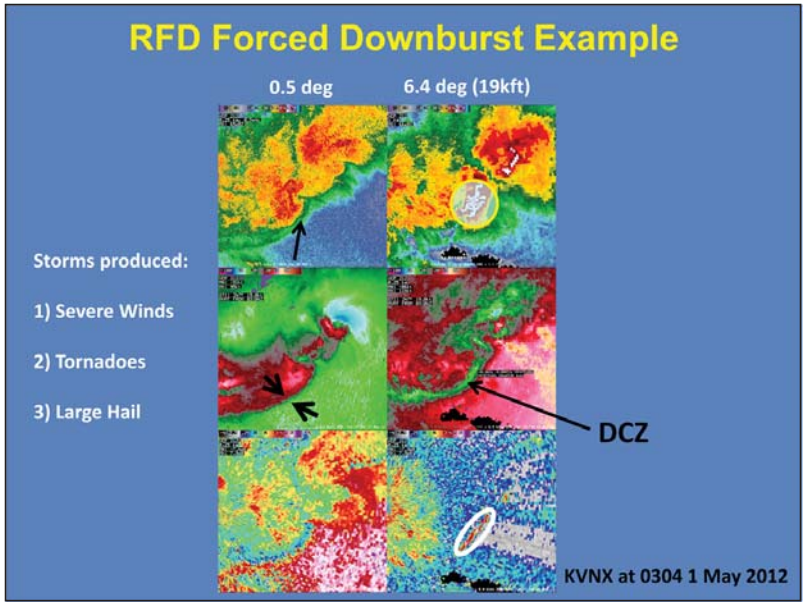
So, what are the signatures of an impending wet or hybrid microburst? Well, after the underlying threat assessment of the environment which can give clues to an increased possibility to an event, you should detect a stronger than normal initial updraft pulse. Now, in terms of verification of this signature, the stronger than normal initial updraft may provide a good lead time but also result in lots of false alarms, as the process is dependent on the dry air entrained to maintain or accelerate the downdraft. Secondly, as we saw in the Peachtree City microburst example, the decent of the strong core is a likely location for the base of the downdraft. This is often associated with the development of mid-altitude radial convergence as air flows into the downdraft source. However, this signature is often ill-defined and may yield a short lead time (on the order of 5 – 10 min) . And, it has a moderate False Alarm Rate (FAR).

Also, most severe pulse storms collapse as they are about to produce a downburst. This process will be evidenced by a simultaneous decrease in cell-based VIL and max reflectivity with time. But, not all weakening storms end in a downburst. And, since you have to wait till the end of volume scan for a time trend, the lead time is short. In addition, the lack of any signals of a storm collapse does not necessarily mean you won't get a downburst. Thus, very high FAR on this precursor. Finally, another microburst detection signal may be a simultaneous decrease in ZDR if hail is occurring in the microburst. But note that not all microbursts will have hail and again not all weakening storms end in a downburst.

Downbursts Forced from RFDs

- Low-level mesocyclogenesis
- Heavy precip (> 50 dBZ) within the hook
- Deep Convergence Zone (DCZ) on the backside of the meso

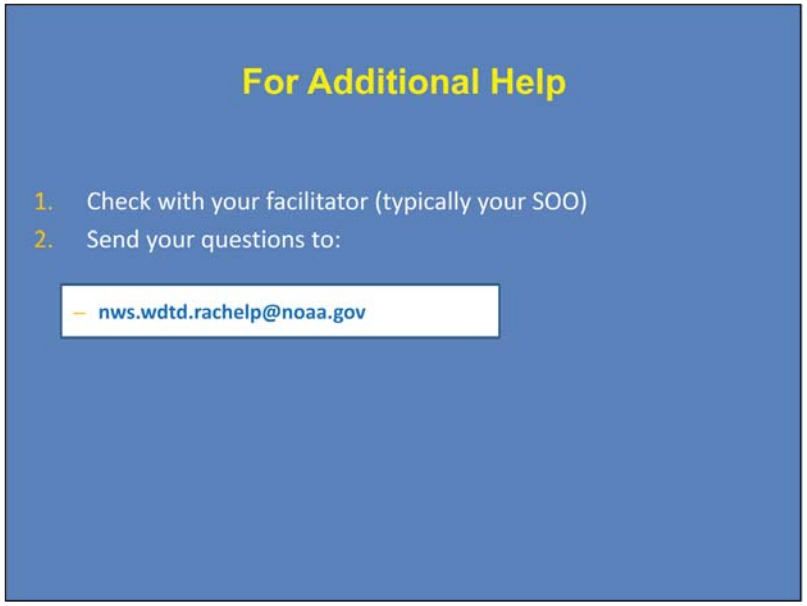
Now the same environmental parameters favoring hybrid or wet microbursts help to encourage a severe RFD**.** However, since we are talking about forcing from non-hydrostatic pressure deficits, they are not necessary if a strong, low-level mesocyclone is developing. Some indications in the supercell include heavy precipitation (> 50 dBZ) within the hook as this favors precipitation loading and evaporation cooling. Next, the indication of a deep convergence zone through 18-20 kft on the backside of the meso is a signal that strong damaging winds will occur usually just the right of the primary meso. Advance to the next slide to see an example of a northern OK HP supercell that produced a severe RFD.



Types of Downbursts	Dry Microbursts	Wet Microbursts	Hybrid Microbursts	Supercells
Forcing	Evaporative Cooling • Sub-Cloud (Low reflectivity) • Sublimation	Evaporative Cooling • Lateral dry air entrainment • Precip Loading • Boundary lifting	Same as Wet	Add Non-Hydrostatic Vertical Pressure Gradients
Environment	- Minimal CAPE - Very Deep Dry ALR - LCL Height Below Freezing - Weak Boundary Layer winds - Weak 0-6 km shear	- Steep midlevel Lapse rates - Moderate CAPE - Large Theta-E differences (> 25 to 30K) from surface to midlevel min) - Low cloud bases Dry sub-cloud ALR	- Sufficient CAPE (> 500 J/Kg) - Steep midlevel Lapse rates - Large Theta-E Differences - High LCLs - Deep, dry sub-cloud ALR	Supercell environment • 0-6 km shear > 15 m/s • Low LCL • Large CAPE • Steep sub-cloud ALR
Pre-cursor Radar Signatures	- Intense core above Freezing Level - Descent of max reflectivity core - Lowering ZDR (~0 DB); Midlevel radial convergence	- Descent of max reflectivity core - Mid-level radial convergence - Decrease in ZDR High KDP in descent	- Descent of max reflectivity core - Mid-level radial convergence - Decrease in ZDR	- Intensifying Mesocyclone with a RFD - Heavy precip (>50 dBZ) in the hook; DCZ

Here is an example of two supercell storms from the evening of April 30, 2012, in northern OK that illustrates most of the previously mentioned signatures of RFD forced downburst winds. With the storm closest to the radar, note a fat reflectivity hook at the 0.5 deg slice with huge inflow notch, a very high elevated Z core with a large 70 dBZ core at 6.4 deg slice (~ 19 kft MSL), elongated convergence zone with 70-80 kts of velocity difference extending up to 21 kft. The ZDR products (lowest two panels) display SW-NE oriented line of values > 1 dB on the 6.4 deg slice which indicates the updraft column extending above -20 deg C. Thus, the storm likely contains hail as well. It was just mentioned that the most damaging RFD winds likely occur with HP supercells, just to the right of the primary mesocyclone track. In this case, damaging winds and tornadoes occurred with both of the two distinct mesocyclones.

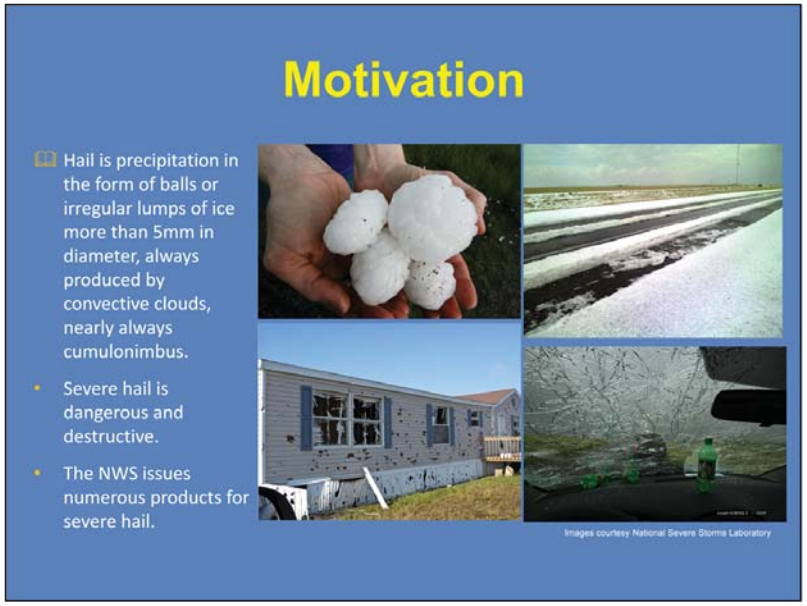
First, recognize clues in the environment. There will be parameters in proximity soundings suggesting an enhanced threat. Look for precursor signals based on the type of forcing. For dry microbursts, look for LCLs below 0 deg C, and an inverted V sounding. For wet microbursts, look for sources of dry air aloft and a subcloud ALR. In radar, you may detect a large, reflectivity core aloft with the initial pulse, then as the downdraft descends, a collapsing core with mid level convergence above the downdraft. For RFD forced events, watch for intensifying low-level mesos, large reflectivity in the hook, and a DCZ. These are some of the main detection signatures.



For additional help, check with your facilitator or send your questions to the listserv e-mail address here.

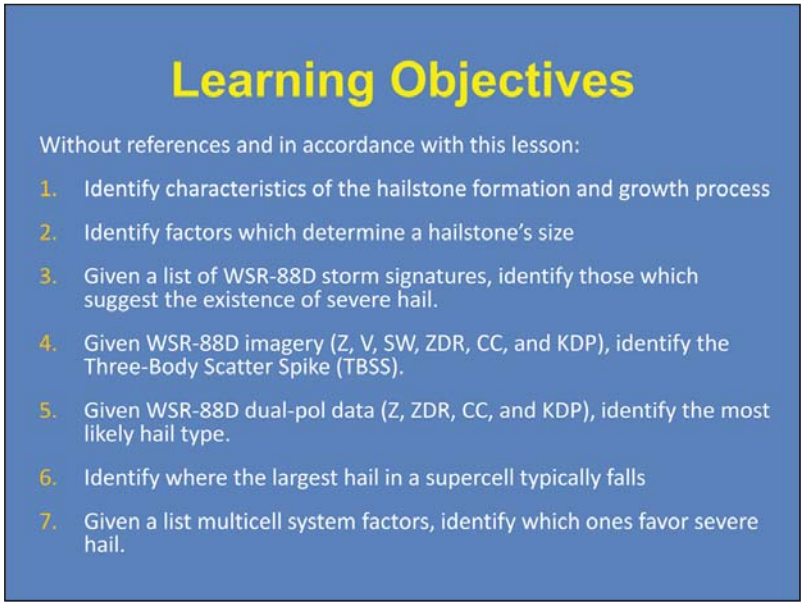


Welcome to the RAC Convective Storm Structure and Evolution lesson on Severe Hail.



Hail is precipitation in the form of balls or irregular lumps of ice more than 5mm in diameter, always produced by convective clouds, nearly always cumulonimbus. Hail can be deadly to people and livestock, and can damage property including homes, businesses, automobiles, and aircraft.

The National Weather Service issues numerous public and aviation products for severe hail. That's why it's important for you to learn about it.



Here are the learning objectives. Please take a moment to read them.

Supercooled Liquid Water (SLW)

- Supercooled Liquid Water (SLW) – Liquid water at a temperature below the freezing point (≈ 0°C)



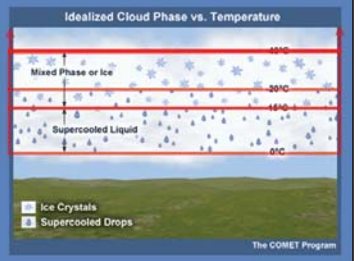
Courtesy Ohio State University

- Pure water suspended in the air does not freeze until it reaches a temperature of nearly -40°C

Let's begin by discussing the formation source of hail which is Supercooled liquid water (SLW). This is liquid water at a temperature below the freezing point (~0°C). It can exist at temperatures below zero because freezing is a complex process. Pure water suspended in the air does not freeze until it reaches a temperature of nearly -40°C.

Supercooled Liquid Water (SLW) in Clouds

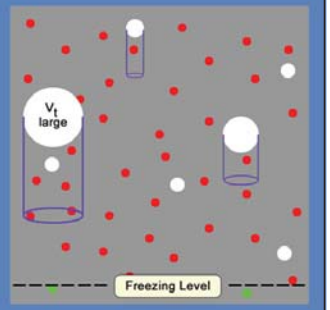
- Supercooled liquid water content of clouds varies with temperature
  - 0 to -15°C: SLW droplets dominate
  - -15 to -40°C: Both ice and SLW droplets coexist
    - At -20°C, the ratio is ≈50/50
    - Strong vertical currents (such as a thunderstorm) may carry SLW droplets to -40°C
  - < -40°C: All ice crystals



The supercooled liquid water (SLW) content of clouds varies with temperature. Between 0° to -15°C, supercooled liquid water droplets dominate. Between -15° to -40°C, both ice and supercooled liquid water droplets coexist. At -20°C, the ratio is about 50/50. However, strong vertical currents such as a thunderstorm may carry SLW droplets to -40°C. At temperatures colder than -40°C, clouds consist entirely of ice crystals.

Hail Formation & Growth Process

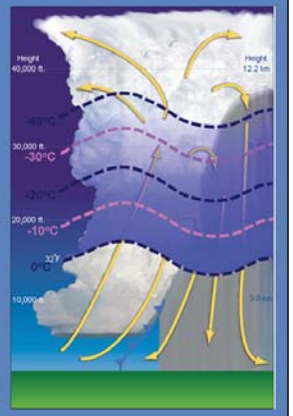
- Hail formation requires an embryo to accrete ice due to collisions with SLW droplets within a cumulonimbus cloud
  - As the embryo grows it, becomes larger and heavier than the surrounding SLW droplets and falls
  - The difference in fall velocities between the hailstone and the SLW droplets leads to growth as the stone "sweeps up" droplets along its path
  - Process takes at least 15-20 minutes
    - Longer for larger hailstones



Hail formation requires an embryo such as an ice crystal, frozen raindrop, dust or some other nuclei to accrete ice due to collisions with supercooled liquid water (SLW) droplets within a cumulonimbus cloud. As the embryo grows it, becomes larger and heavier than the surrounding supercooled liquid water droplets and falls (relative to the supercooled droplets). The difference in fall velocities between the hailstone and the SLW droplets leads to growth as the stone "sweeps up" droplets along its path. Process takes at least 15-20 minutes; longer for larger hailstones.

Hailstone Size Factors

- Size is largely dependent on its residence time in the preferred hail growth zone of supercooled liquid water (-10° to -30°C)
  - Strong, wide, persistent updrafts are most favorable
  - Growth continues until the hailstone's mass becomes large enough to overcome the updraft

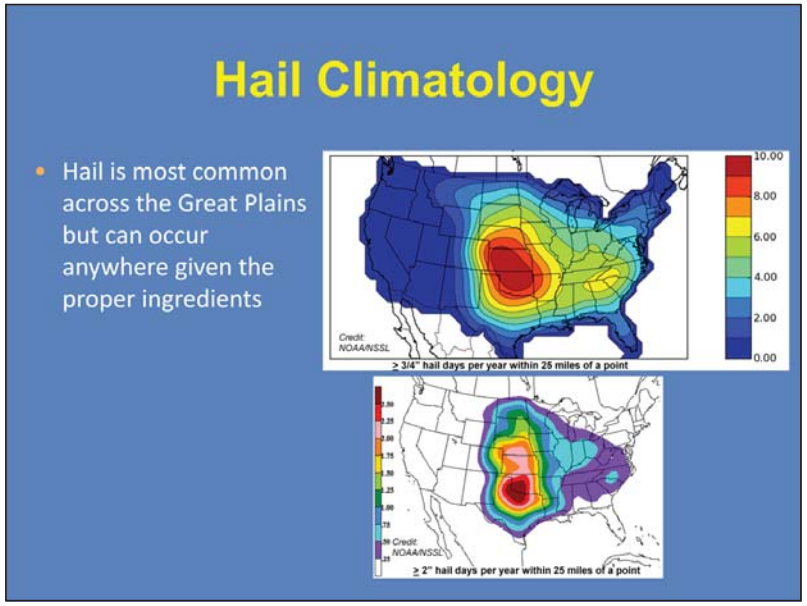


A hailstone's size is largely dependent on its residence time within the preferred hail growth zone of supercooled liquid water between -10° to -30°C. Strong, wide, persistent updrafts are most favorable. The growth rate is maximized near -13°C and rapidly diminishes at temperatures approaching -30°C as supercooled water droplets become rare at these colder temperatures. Growth continues until the hailstone's mass becomes large enough to overcome the attendant updraft.



Here is a table which equates hail description with its diameter and approximate updraft speed. The National Weather Service criteria for severe hail is one-inch diameter (or about quarter-size), which is based on research that indicates this is the threshold at which significant damage occurs.

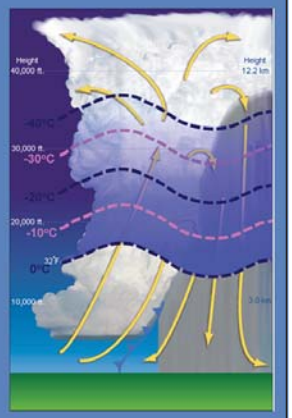
Storm spotters are taught to report the largest hailstone they observe. That's what WFOs log into their Local Storm Reports and is subsequently published in the National Climatic Data Center's official Storm Data publication.



Hail is most common across the Great Plains, but can occur anywhere given the proper ingredients. Please take a moment to view these two hail climatology graphics.

Warning Forecaster's Dilemma: Can't Measure Hailstone Residence Time

- A warning forecaster cannot measure a hailstone's residence time in the hail growth zone
  - Thus, WSR-88D proxy signatures must be used

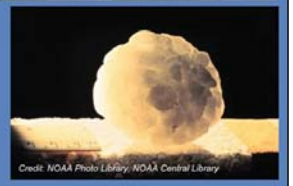


A warning forecaster cannot measure a hailstone's residence time in the hail growth zone. It's difficult to estimate updraft strength, let alone perform a real-time trajectory analysis. Thus, WSR-88D proxy signatures must be used to detect severe hail.

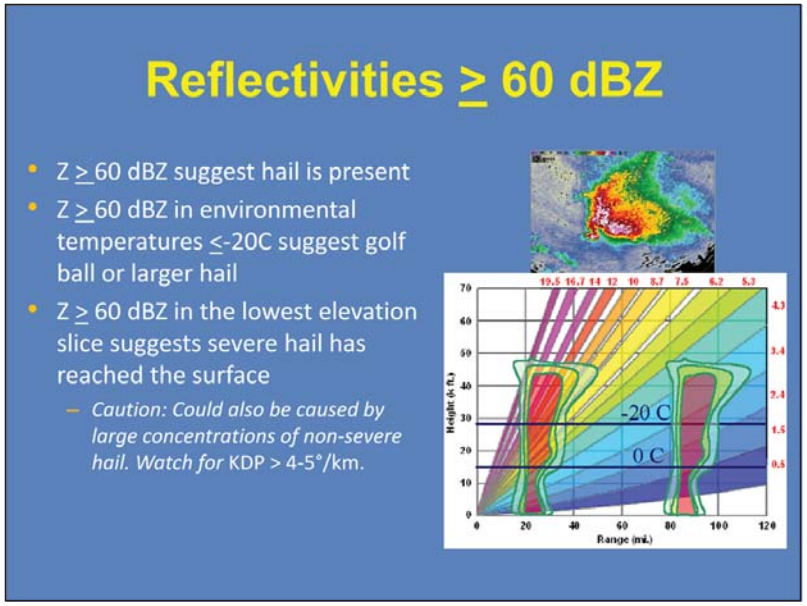
WSR-88D Detection of Severe Hail

- WSR-88D data be used can be used detect or infer the existence of severe hail via proxy signatures:
  - Z ≥ 60 dBZ
  - Three-Body Scatter Spike (TBSS)
  - Weak Echo Region (WER)
  - Bounded Weak Echo Region (BWER)
  - Mesocyclone
  - Storm-top divergence
  - Dual-polarization output





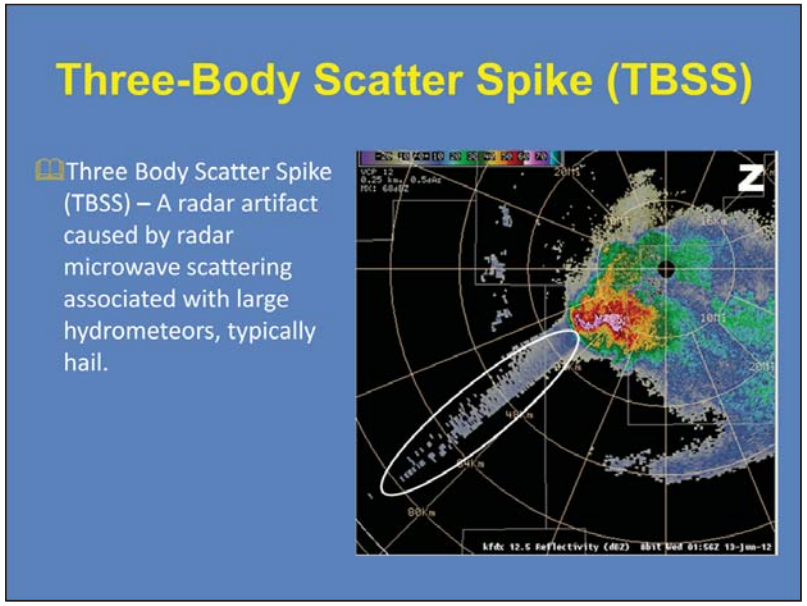
WSR-88D data can be used to detect or infer the existence of severe (≥ 1-inch) hail via the proxy signatures listed here. Please take a moment to view them before we discuss them in detail.



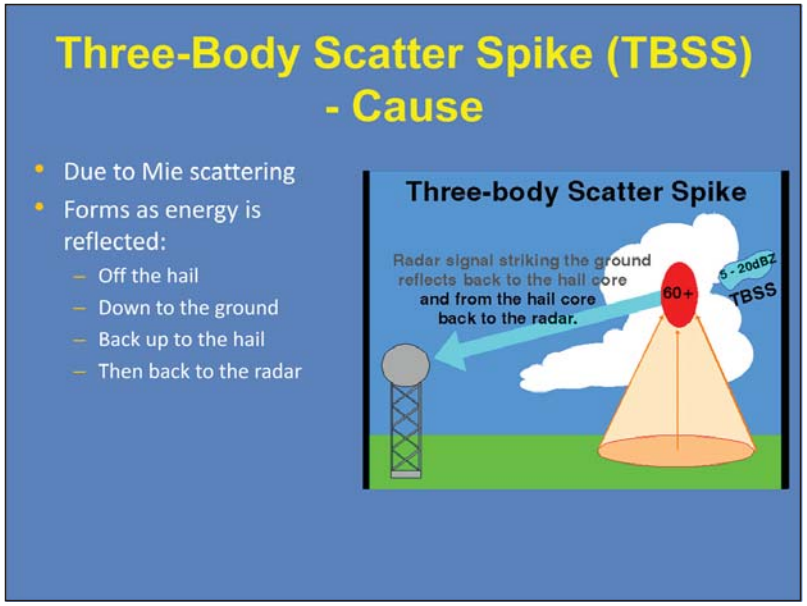
Reflectivities greater than or equal to 60 dBZ suggest hail is present. This is because pure water is almost impossible when reflectivity is greater than 55 dBZ. Thus, you can be confident that hail is present.

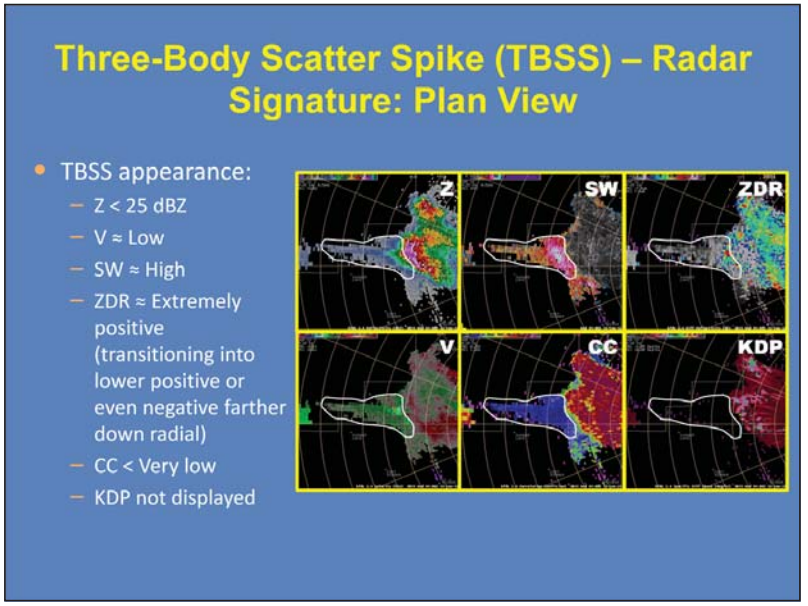
Reflectivities greater than or equal to 60 dBZ in environmental temperatures less than or equal to -20°C suggest hail greater than or equal to golf ball size (1.75”). Hail this large is unlikely to melt before it reaches the surface regardless of the melting level height. Increase reflectivities in this layer and the odds of severe hail increase dramatically. Reflectivity greater than or equal to 65 dBZ through the entire preferred hail growth zone (-10°C to -30°C) suggests the potential for giant (≥4”) hail (Blair et al., 2011). Thus, it’s important to know the height of the -20°C level.

Reflectivities greater than or equal to 60 dBZ in the lowest elevation slice strongly suggest the presence of hail that has reached the surface. Beware that is could also be caused by large concentrations of sub-severe melting hail so watch out for KDP values greater than 4-5°/km. Also, beware that giant, very dry hail in low quantities have been observed with reflectivities between 35 and 50 dBZ. When this occurs, the hail is usually falling in the strong reflectivity gradient immediately adjacent to an inflow notch and updraft with a supercell storm.



A Three-Body Scatter Spike (TBSS) (also known as a “hail spike”) is a radar artifact caused by radar microwave scattering associated with large hydrometeors, typically hail. The TBSS is characterized by extremely high ZDR and very low CC located (radially) just behind high reflectivity cores in hail-bearing storms.

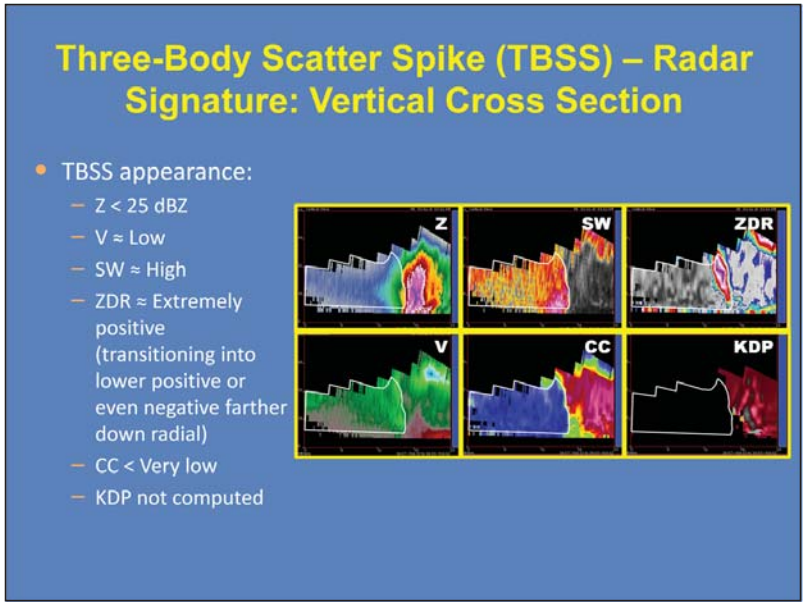




The TBSS is an artifact of the electromagnetic radar beam being subject to “Mie scattering” instead of the usual “Rayleigh scattering” process.

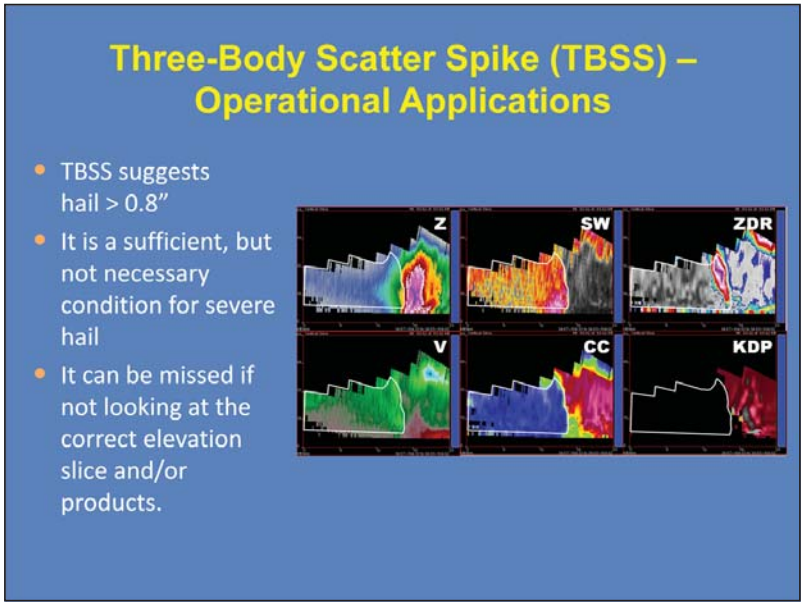
A TBSS forms as incident energy from the radar is reflected off the hail, down to the ground, then back up to the hail and back to the radar. Because of the delay in reception of the pulses, the radar circuitry displays the TBSS as downrange from the hail core.

The TBSS signature produces low reflectivities (generally less than 25 dBZ), low radial velocities (V), and high spectrum widths (SW). In ZDR, the TBSS appears as an area of extremely positive values just down-radial of the hail core (transitioning into lower positive or even negative values farther down-radial). In CC, the TBSS shows up very clearly as a spike of extremely low values (generally < 0.5), on the down-range side of the hail core. CC is especially useful in cases when the TBSS reflectivity signature is short or masked by precipitation echoes. The spike is not seen in KDP because of the 0.90 CC threshold filter discussed in previous lessons.



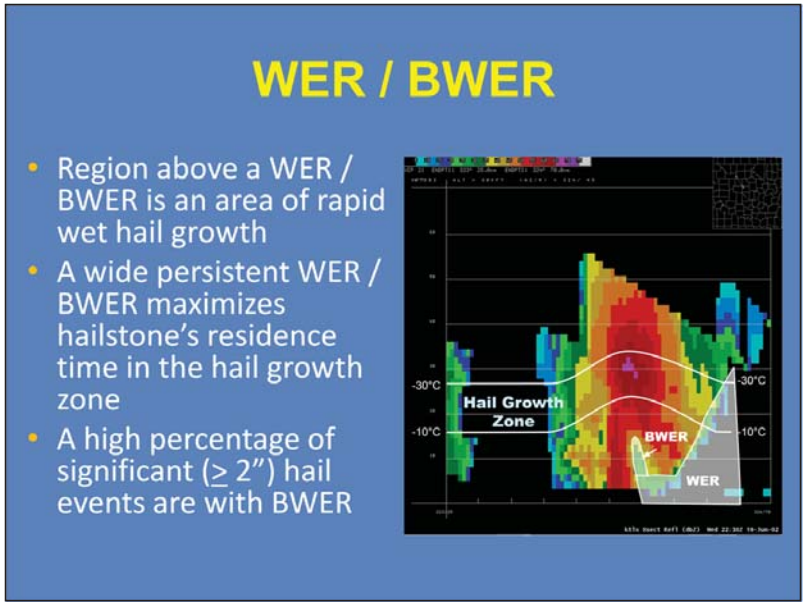
Taking a look at cross sections of the same storm, the TBSS is seen down-radial of the high reflectivity (Z) core. High spectrum widths (SW) and low radial velocities (V) exist down-radial from the hail core. In ZDR, the area of extremely high, positive values located immediately behind the hail core is usually somewhat wedge-shaped, while farther down-radial, there is a transition to negative values of ZDR. In CC, the TBSS is marked by very low values. The spike is not seen in KDP because of the 0.90 CC threshold.

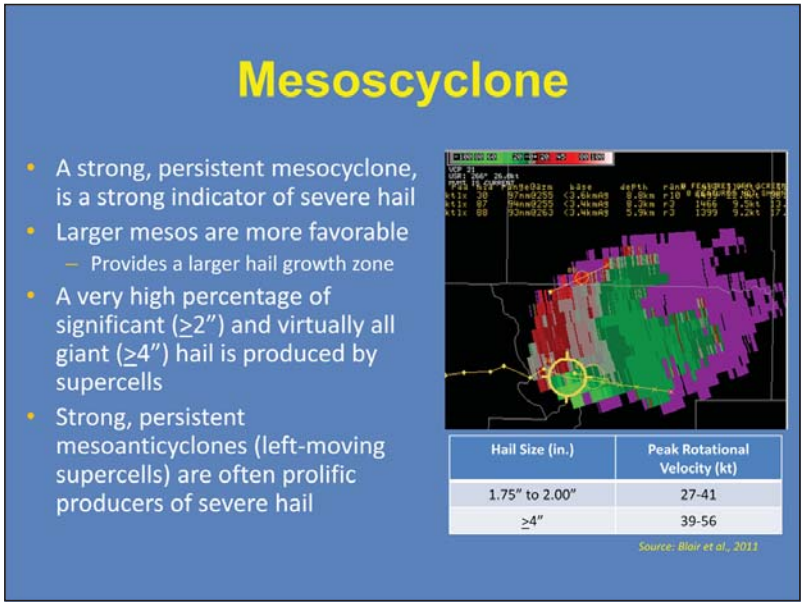
The strength and length of the TBSS is related to the intensity and vertical extent of the reflectivity core. Therefore, a TBSS should be easier to detect with a more intense and elevated reflectivity core. Also, the larger the highly reflective core area, the more extensive the TBSS.



The presence of a TBSS on an S-band (10 cm) radar such as the WSR-88D suggests the thunderstorm possesses hail larger than about 0.8” in diameter which is either reaching the surface or will do so within the next 10-30 minutes. Additionally, because storms with a TBSS often produce damaging surface winds, resulting warnings should contain mention of severe wind if the near storm environment is favorable.

Be aware of the limitations of using the TBSS signature. It is a generally sufficient, but not necessary condition for severe hail identification. It can be missed if not looking at the correct elevation slice and/or products. This signature is an artifact of Mie scattering and must not be construed as hail actually reaching the surface *beneath the echo spike itself*. The TBSS signature can only be applied on S-band (10 cm) radars, such as the WSR-88D. On C-band (5 cm) radars, the TBSS can be related to large raindrops rather than hail.



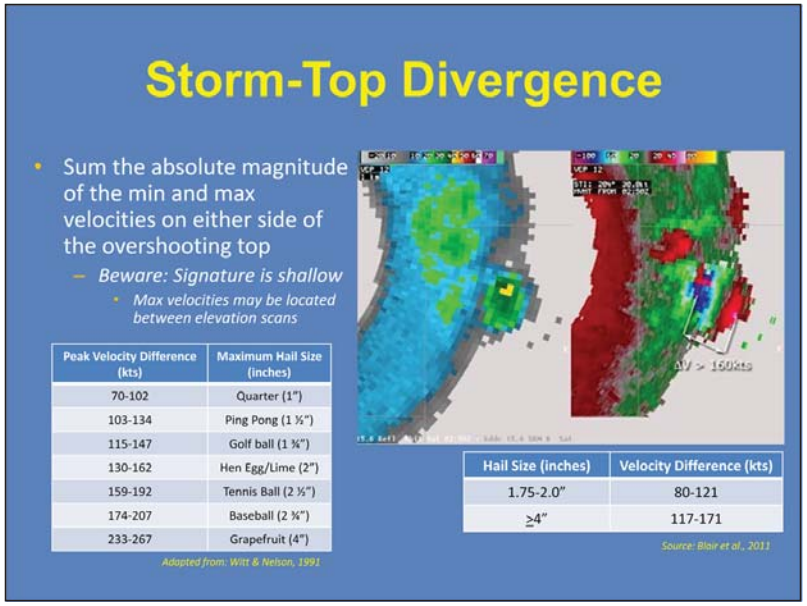


Studies have shown that the high reflectivity region above the top of the Weak Echo Region (WER) and Bounded Weak Echo Region (BWER) is an area where rapid wet hail growth is occurring in the core of an intense updraft. The existence of a WER/BWER suggests that these hailstones are subject to a massive influx of supercooled cloud water and growth, especially if in the favored hail growth zone (-10°C to -30°C). A wide, persistent WER/BWER helps to maximize a hailstone's residence time in the favored hail growth zone before it cascades into the core. A high percentage of significant (2-inch diameter and larger) hail events are associated with BWERs. Remember that a bona fide WER/BWER must be topped by intense reflectivities in order for it to be associated with updraft.

Severe hail can be inferred from velocity signatures as well.

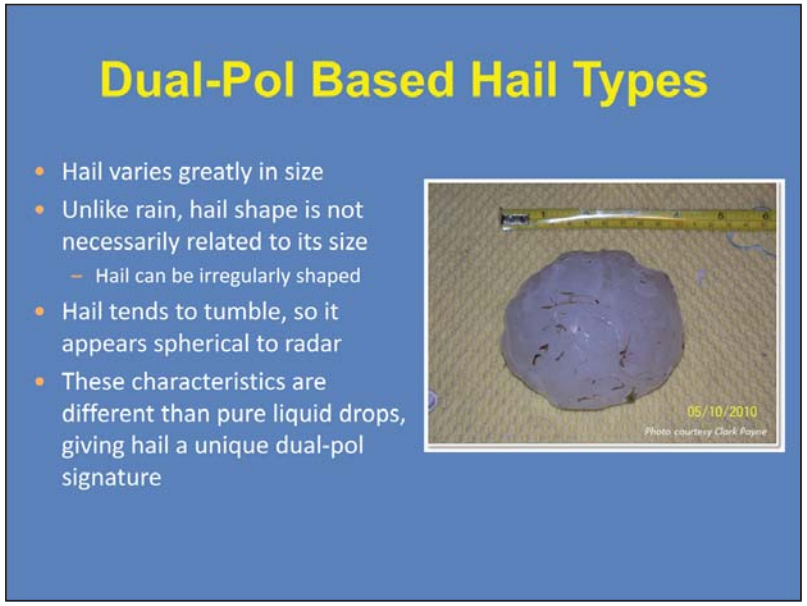
A strong, persistent mesocyclone, the defining characteristic of a supercell, is a strong indicator of severe hail. Dynamic pressure drops, especially in a strong mid-level mesocyclone, accelerate the updraft in the hail growth zone. It is hypothesized that mesocyclones produce favorable trajectories that lead to enhanced hail growth. Larger diameter mesocyclones are more favorable for severe hail because they provide a larger hail growth zone which also increases residence times and hail growth potential. Mesocyclone strength is especially important for the growth of large hail. A very high percentage of significant (≥2") and virtually all giant (≥4") hail is produced by supercells.

Strong, persistent mesoanticyclones (left-moving supercells) are often prolific producers of severe hail. This may be because they move faster and cover more area than right-movers.



Storm-top divergence can be used to assess a storm's maximum potential hail size. An example of this storm-top divergence technique can be seen here. First, use the proper elevation slice to sample the storm's overshooting top. Sum the absolute magnitude of the minimum and maximum velocities found on either side of the overshooting top and you'll get a representative sample of storm top velocity difference, in this case > 160 kts. If you cannot sample the overshooting top as shown here, then the technique may fail. There are other factors that may cause this technique to fail (e.g., poor data quality, mini-supercells, etc.), so use with caution.

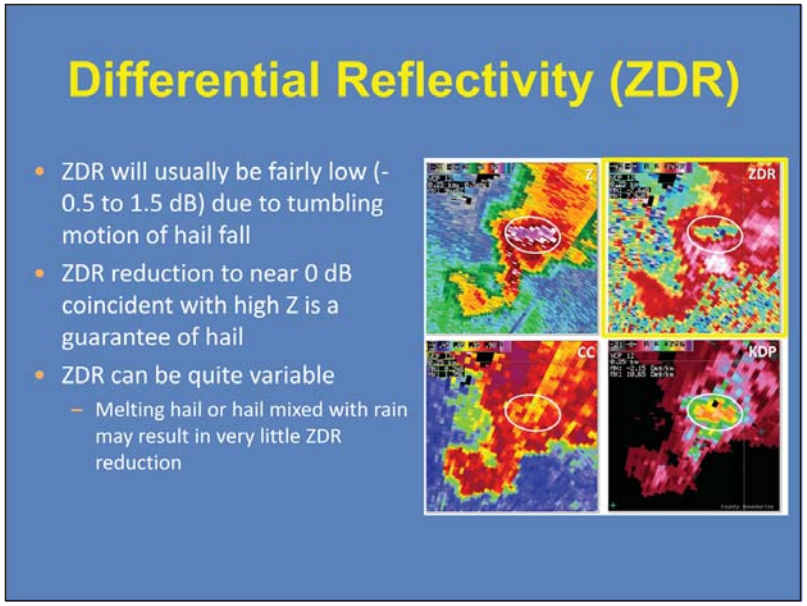
Two tables relating maximum storm top divergence velocity difference to hail size are shown here.

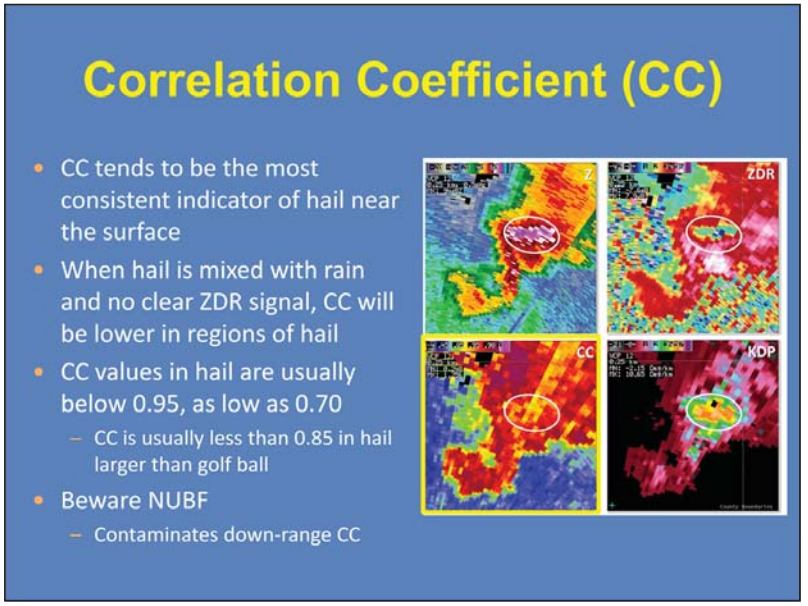


Dual-pol data is very useful for hail detection.

Hail varies greatly in size, from as little as a quarter of an inch up to 8 inches in diameter. Unlike rain, the shape of hail is not necessarily related to its size. Hail can be irregularly shaped, with some hailstones having large protuberances, and in some cases be elliptical with one particular dimension much larger than the other. Hail also has the tendency to tumble, so it tends to appear effectively spherical to the radar. These characteristics are different from pure liquid drops, giving hail a unique signature in dual-pol data.

Let's discuss the hail detection capability of dual-pol products individually and then show how these products can be used in combination with reflectivity (Z) to detect certain hail event types.

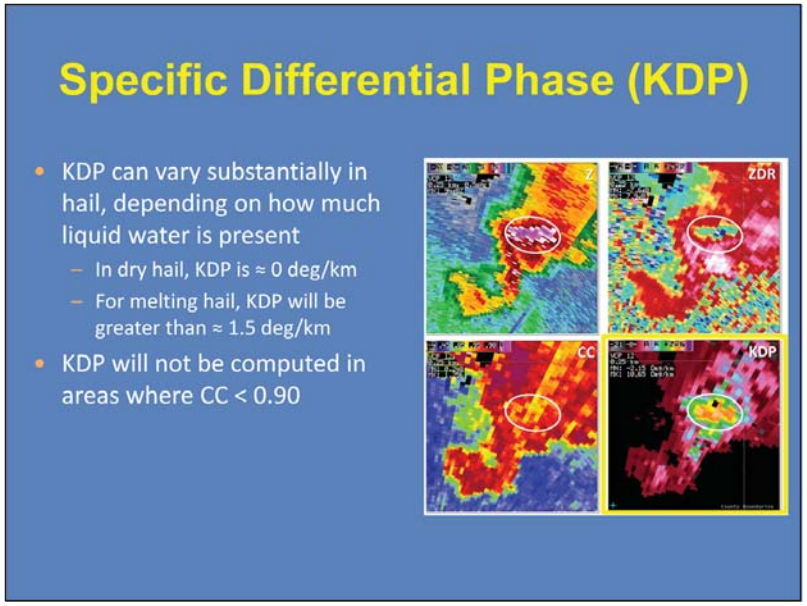




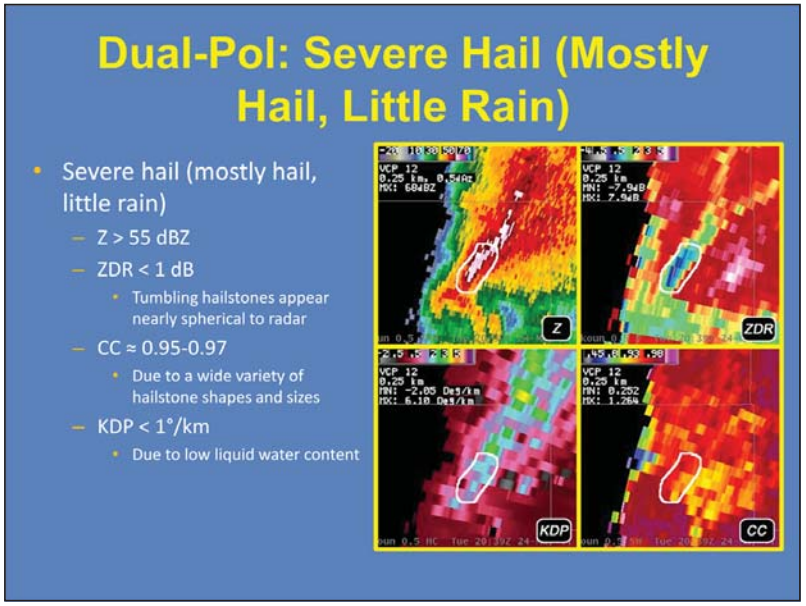
Differential Reflectivity (ZDR) will usually be fairly low, between -0.5 dB and 1.5 dB, due to the tumbling motion of the hail as it falls. A reduction in ZDR to near 0 dB coincident with high reflectivity is a guaranteed detection of hail. ZDR can be quite variable though, and in cases where the hail is melting or mixed with rain there may be very little reduction in ZDR.

Correlation Coefficient (CC) tends to be the most consistent indicator of hail near the surface. In cases when hail is mixed with rain and there is not a clear signal in ZDR, CC will be locally lower in the regions containing hail. Values of CC in hail are usually below 0.95 and can be as low as about 0.70. For hail larger than roughly golf balls (> 1.75-inches), CC is normally less than 0.85.

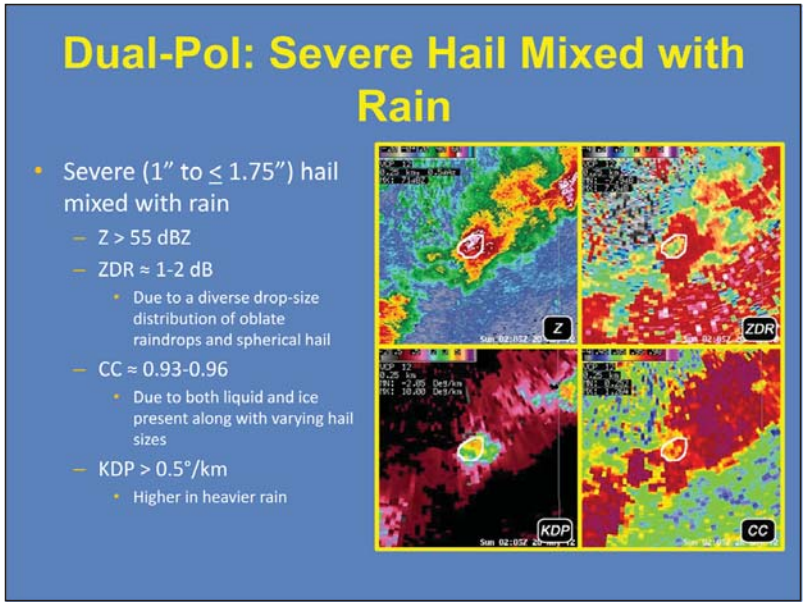
It is important to note that the CC values down range from the hail region in this figure are characterized by radial streaks of depressed CC values. This is known as non-uniform beam filling (NUBF) and is one of the limitations of the Correlation Coefficient product. It occurs when at least some of the radar pulse volumes are characterized by significant gradients of PhiDP across these pulse volumes. When this occurs, these down-range CC (and other dual-pol base data) values are “contaminated,” or compromised. This means that the CC values and all other dual-pol values along the affected radials cannot be trusted or used.



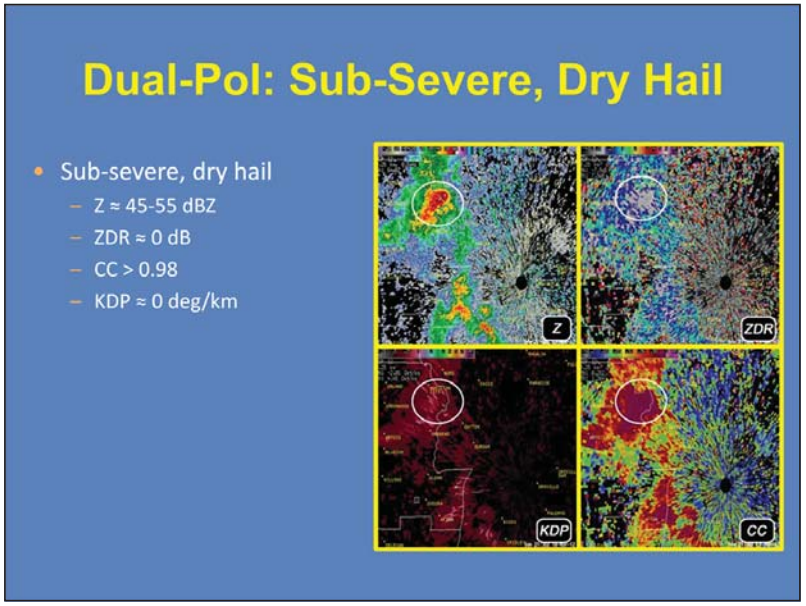
Specific Differential Phase (KDP) can vary substantially in hail, depending upon how much liquid water is present with the hail. In dry hail without much rain, KDP is near 0. For melting hail, KDP will be greater than about 1.5 deg/km. This assumes that KDP is computed in areas containing hail, which is not the case when CC is < 0.90 (notice the black range gates within circle).



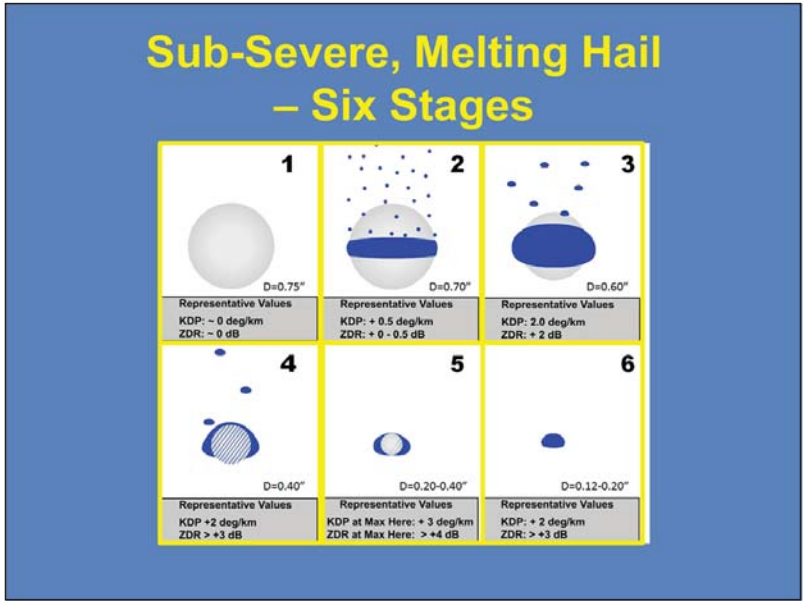
Many research papers refer to the “classic, severe hail signature.” Severe hail, by definition, has a diameter of at least 1 inch (although some research papers were written when severe hail was threshold was ¾-inch diameter). Reflectivity values for severe hail are usually larger than pure rain events (Z > 55 dBZ). Since hail often tumbles as it falls, hailstones appear nearly spherical to the radar (ZDR typically < 1 dB). A wide variety of hail shapes and sizes result in CC tending toward the 0.95-0.97 range. Finally, KDP values are lower than pure rain (KDP typically < 1 deg/km).



When rain is mixed with the "classic, severe hail signature," the dual-pol variables behave a little differently. Reflectivity values will still be very high (Z > 55 dBZ) because of the size dependence of Z. ZDR values will be more positive (ZDR ~ 1-2 dB) as the diverse drop-size distribution of oblate rain and spherical hail both contribute significant power returns. Likewise, CC will be slightly lower than the "classic, severe hail signature" (CC ~ 0.93-0.96) because there is now both liquid and ice present along with varying sizes of hail. KDP will increase (KDP > 0.5 deg/km) because it's not dependent upon drop size like Z and ZDR. KDP only depends upon drop shape and number concentration. As a result, tumbling hail doesn't contribute significantly to KDP, but oblate rain drops will. The heavier (and more concentrated) the rain, the higher the KDP values will be.



Sub-severe (< 1-inch), dry hail also has some unique dual-pol characteristics. Reflectivity will still be high, but not as high as other hail cores (Z ~ 45-55 dBZ). Small hail will tend to be smooth on the surface and appear spherical on radar (ZDR ~ 0 dB). CC should be near uniform (CC > 0.98) since the hail stones are similar in shape and size and little liquid water content is present. Likewise, KDP will be low (KDP ~ 0 deg/km).



Dual-polarization provides the capability to detect hail that has significantly melted, to the point that it's not likely to be severe. This signature is important because hail signatures can be found in just about any convective storm near and just above the freezing level. Therefore, observing a hail signature doesn't mean the hail will reach the surface. So let's discuss how hail melts and what it should look like on dual-pol products. A melting hailstone goes through six stages.

Stage 1: Begin with a solid ice sphere (D=0.75") falling outside of the updraft below the 0ºC level. Melting begins on the surface of the hailstone.

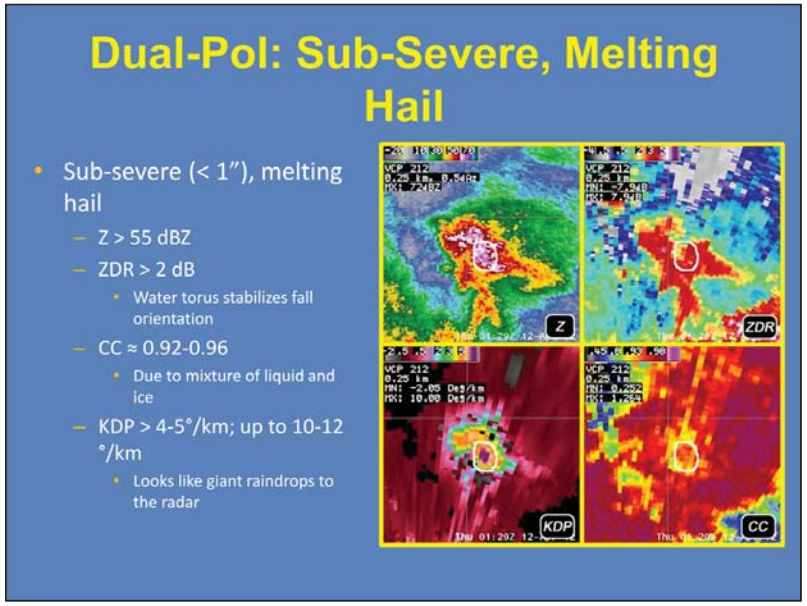
Stage 2: As the surface melts, the meltwater is advected into a torus (blue band around the equator of the hailstone) due to drag as it falls. Continuous shedding of small drops (~1 mm) occurs from the torus of water. Shed drops fall much slower. Hail diameter now 0.70."

Stage 3: Hail continues to melt and the torus moves upstream as the size of the ice particle decreases. Intermittent shedding of large drops (~3 mm) occurs from the unstable torus. Hail diameter now 0.60."

Stage 4: The torus loses its distinction, and a water cap forms around the top (lee side) of the ice core. Intermittent shedding of a few large drops (~3 mm). Ice core is now 0.40" diameter.

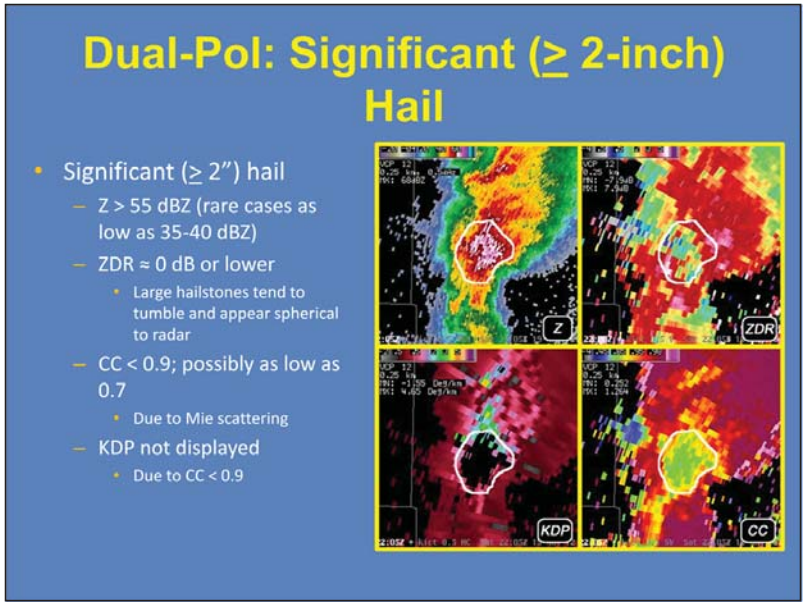
Stage 5: Meltwater forms a stable raindrop shape around the ice core. There is no drop shedding any longer. Horizontal axis diameter of ice and water coating ~0.2" to 0.4" (~5-9 mm).

Stage 6: Eccentric melting of ice core occurs until ice is completely melted. All that is left is a large, cold rain drop ~0.12" to 0.20" (3-5 mm) in diameter.



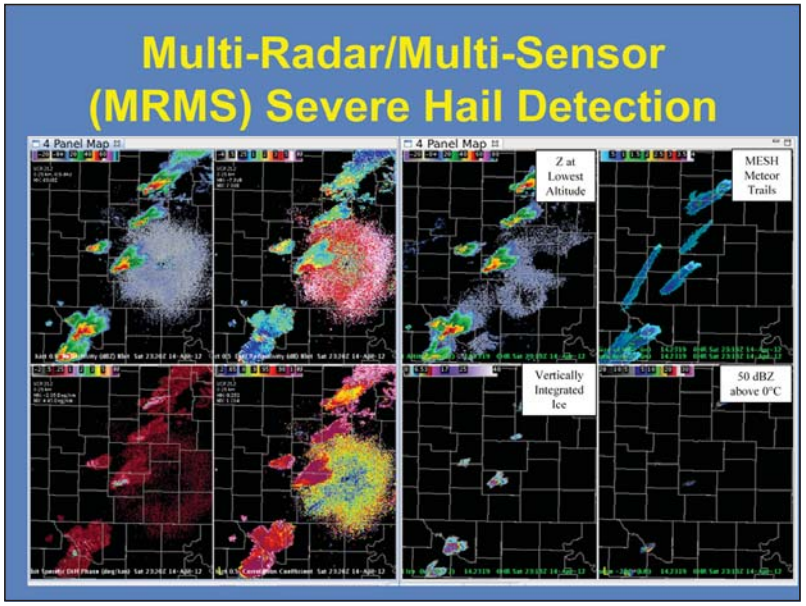
Now that you have seen how smaller hail melts, let's discuss how it appears to dual-pol radar. As you have seen, when sub-severe hail melts, it develops a water torus on the surface around its center. This water torus tends to stabilize its fall orientation, making it look like a giant rain drop to the radar. As a result, reflectivity values tend to remain high (Z > 55 dBZ). Likewise, ZDR increases (ZDR > 2 dB; possibly as high as 6 dB). CC decreases to around 0.92-0.96 due to the mixture of ice and liquid.

KDP is very revealing in this case. KDP in small, melting hail can become extremely large (KDP up to 10 deg/km!). Why is that? When there is a high concentration of these "giant raindrops" (that is sub-severe, melting hailstones with a water torus), then KDP values can become extremely large. In pure rain situations, KDP values will rarely go above 4-5 deg/km. When KDP is larger than those values, you can confidently assume that there is some small melting hail present.



When significant (2-inch diameter and larger) hail is present, the signature in the dual-pol products can be very pronounced. When hail gets to be larger than golf balls, Mie scattering effects begin to alter the way the dual-pol variables appear, and this signature is unique. Reflectivity (Z) will remain high (Z > 55 dBZ) except for rare cases when only a few, large hailstones fall in the Z gradient near the updraft/downdraft interface region (Z as low as 35-40 dBZ). Differential Reflectivity (ZDR) will still be near zero or even be mostly negative (ZDR ~ 0 dB or lower).

Correlation Coefficient (CC) is the most revealing product in this case. Mie scattering will cause CC values to drop significantly lower (CC < 0.9; possibly as low as 0.7!). Dropouts in the Specific Differential Phase (KDP) data will appear since gates where CC < 0.9 are filtered from the product. Therefore, if you see high reflectivity (or moderate reflectivity near a supercell updraft) and CC < 0.9, you can confidently say significant hail is present.

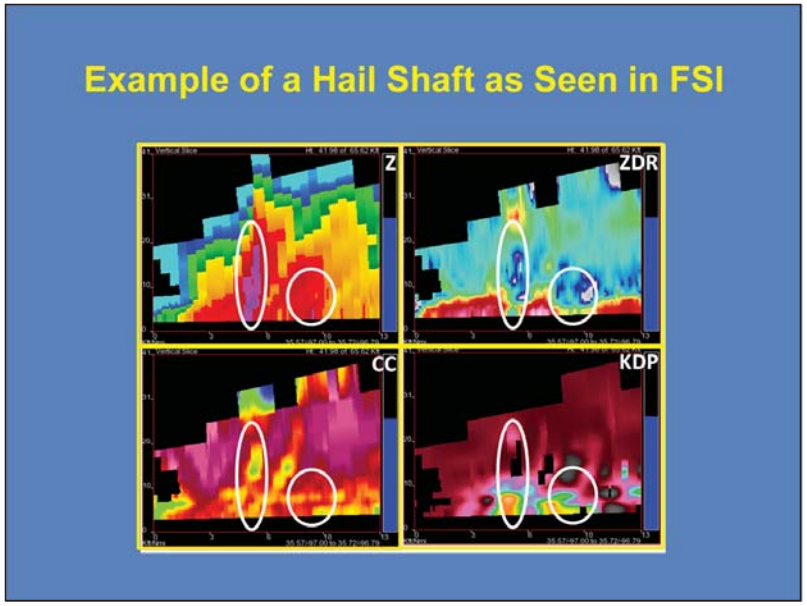


Multi-Radar/Multi-Sensor (MRMS) products are useful for severe hail detection. Perhaps the most useful is Maximum Estimated Size of Hail (MESH), and its associated MESH Tracks products. MESH has been shown to be very useful for assessing both the 2D distribution of hail and largest hailstone size associated with a storm. MESH Tracks products are useful for assessing both storm intensity trends and deviations in storm motion. Forecasters have discovered the benefits of image combining the MESH on top of the MESH Tracks products to create a MESH "Meteor Trails" display which is useful for the orientation of warning polygons.

Another group of MRMS products that can be used to detect severe hail are the reflectivity thickness products. In particular, the 50 dBZ Echo above 0°C product can help you to streamline the Donavon Technique.

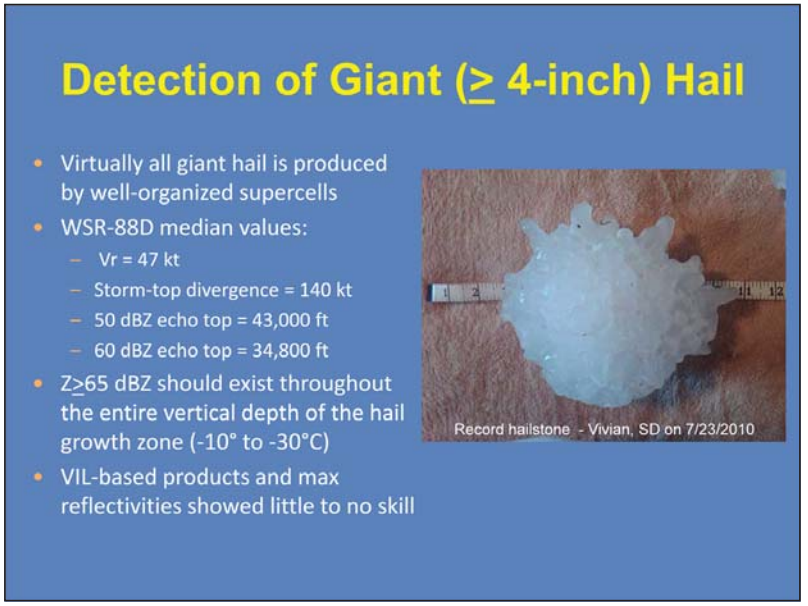
MRMS Vertically Integrated Ice (VII) can be used to assess storm severity including hail potential. It can also be used to assess changes in updraft intensity. Sudden increases (decreases) in VII often occur when an updraft is intensifying (weakening). It can be used to identify the region of a storm where new cell growth is occurring which is particularly useful for warning polygon orientation.

Several other MRMS products can be useful for severe hail detection, including the Reflectivity at x°C products. Make sure you keep a single-site radar display visible when using MRMS hail products during warning operations.



The AWIPS Four-Dimensional Stormcell Investigator (FSI) tool can be used for hail detection with dual-pol products. In the FSI cross section seen here, there are actually two hail shafts, marked by the white circles. As expected, both hail shafts have high reflectivity, although the one on the left, associated with a newer updraft, is deeper with higher values.

Taking a look at ZDR, both hail shafts are associated with fairly low positive to slightly negative values just above the melting layer, with the lower values extending higher aloft for the (younger) hail shaft on the left. Note also the depression in the transition to higher, positive values in each hail shaft. The signal in CC is not as easy to pick out as in the reflectivity or ZDR in this case, but notice the values of CC are generally low in both hail shafts. KDP is generally very high in both hail shafts reflecting increased liquid water content associated with melting hail and/or a rain/hail mixture.



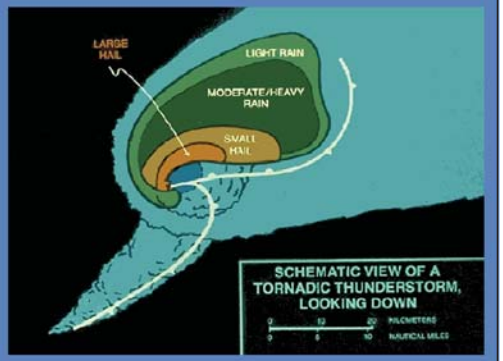
Giant (≥4-inch diameter) hail is a relatively rare phenomenon, accounting for less than 1% of all hail reports in the United States.

Blair et al. (2011) examined several radar signatures to assess their utility in discriminating storms most favorable for giant hail. It was found that virtually all giant hail-producing storms were supercells with a well-organized structure. They were characterized by median values of mesocyclone rotational velocitiy (Vr) of 47 kts (24 m s^{-1}), storm-top divergence value of 140 kt (72 m s^{-1}), 50-dBZ echo top of 43,000 ft (13,100 m), and 60-dBZ echo top of 34,800 ft. (10,600 m). Median reflectivies ≥65 dBZ were present through the entire vertical depth of the preferred hail growth zone (-10° to -30°C), suggesting that high reflectivity should reside throughout that layer with giant-hail producing supercells.

They also noted that: Vertically Integrated Liquid (VIL)-based products, maximum reflectivity within the storm, and reflectivity within the preferred hail-growth zone showed little to no skill in discriminating between giant and smaller hail sizes.

Storm-Relative Hail Location: Precipitation Size Sorting

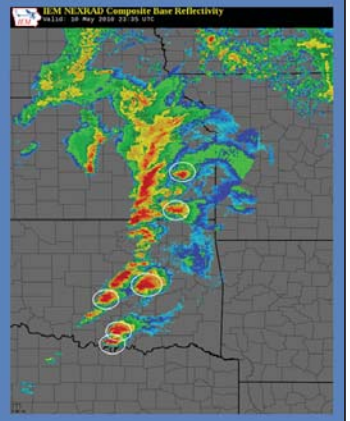
- Largest hail falls to the left and rear of the updraft along the updraft/downdraft interface region
- Progressively smaller hailstones fall at increasing distances to the left and left forward of the updraft
  - Due to precipitation size sorting.



Remember to always keep the radar products in context with the storm features, especially the updraft location as denoted by the Weak Echo Region (WER) and Bounded Weak Echo Region (BWER). The largest hailstones typically fall to the left and rear of a supercell updraft (relative to its movement), along the updraft/downdraft interface region denoted by the strong low-level reflectivity gradient, inflow notch, and sometimes hook echo. Progressively smaller hailstones fall at increasing distances to the left and left forward of the updraft due to precipitation size sorting.

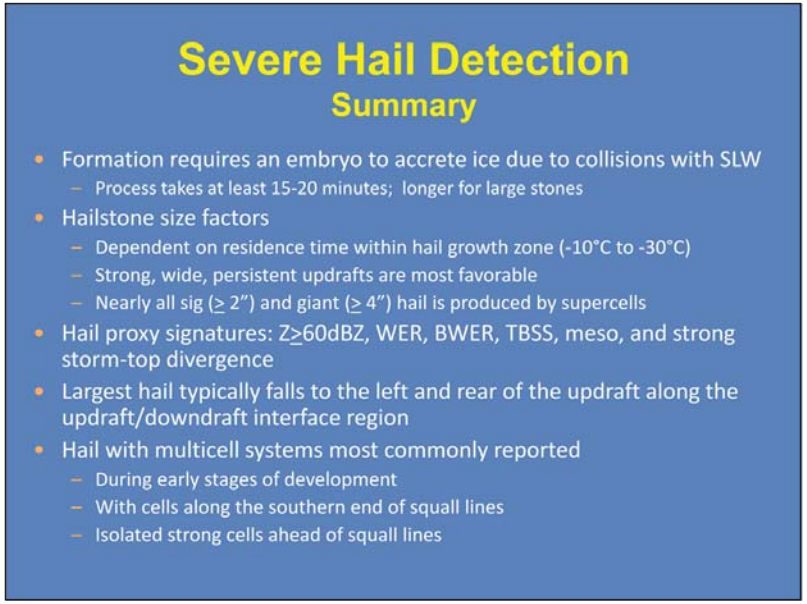
Factors Which Favor Severe Hail in Multicell Systems

- Tornado and hail reports are biased toward the early stages of multicell systems and are most frequently associated with:
  - Cells located along the southern end of squall lines
  - Isolated strong cells ahead of the squall lines
- Sig hail can form with quasi-stationary strong cells in a multicell complex, such as cells at a boundary intersection



Houze et al. (1990) documented severe weather locations for various mesoscale precipitation systems and found that tornado and hail reports were biased toward the early stages of multicell system development. They were most frequently associated with 1) cells located along the southern end of squall lines and 2) isolated strong cells ahead of the squall lines. This contrasts with high wind reports which are sometimes reported with isolated cells but are more numerous along well-developed convective lines. As multicell systems intensify, the effects of the cold pool and resulting increasing rear-to-front flow in the system tend to force an upright updraft along the leading edge. Any significant hail fall will likely occur in this region, not in the downdraft region or wake of the multicell system, which becomes dominated by cooler, saturated air.

Significant hail can occasionally form with quasi-stationary strong cells in a multicell complex, such as cells which form in the vicinity of a surface boundary where strong low-level convergence is focused near the updraft region of the complex.



Hail formation requires an embryo to accrete ice due to collisions with supercooled liquid water (SLW) droplets within a cumulonimbus cloud. The process takes at least 15-20 minutes; longer for larger hailstones.

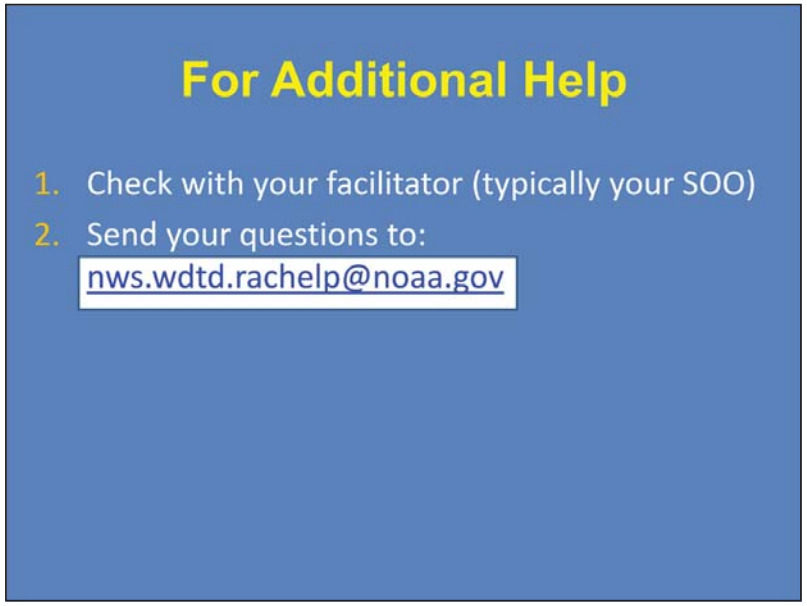
A hailstone's size is dependent on its residence time within the preferred hail growth zone of supercooled water between -10°C to -30°C. Strong, wide, persistent updrafts are most favorable. A very high percentage of significant (≥2-inch) and virtually all giant (≥4-inch) hail events are produced by supercells.

WSR-88D signatures can be used to detect or infer the existence of severe hail including: Reflectivities ≥60 dBZ, Weak Echo Region (WER), Bounded Weak Echo Region (BWER), Three-Body Scatter Spike (TBSS), Mesocyclone, and strong storm top divergence.

Hail detection is more robust with the inclusion of dual-pol data the inclusion of which can be used to detect hail type.

The largest hailstones typically fall to the left and rear of a supercell updraft, along the updraft/downdraft interface region denoted by the strong low-level reflectivity gradient, inflow notch, and sometimes hook echo.

Hail reports with multicell systems are biased toward the early stages of development and are most frequently associated with cells located along the southern end of squall lines and isolated strong cells ahead of the squall lines.



For additional help, check with your facilitator (typically your SOO) or send your questions to the listserv e-mail address here.



Welcome to the Convective Storm Structure and Evolution's lesson on supercell dynamics and motion. This lesson is about 30 min long.

Learning Objectives

- Identify the typical environment, storm structure, and evolution of supercells.
- Identify the effects of shear on storm propagation.
- Identify the technique to anticipate the motion of supercells.

These are the learning objectives that you need to learn in this lesson. Identify the typical environment, storm structure, and evolution of supercells.

Identify the effects of shear on storm propagation. Identify the technique to anticipate the motion of supercells.

Reasons for Storm Intensification in Significant Vertical Shear

- As deep vertical shear exceeds 20 m/s
  - Increased updraft/downdraft separation
  - Precipitation removal from updraft
  - Lower boundary-relative storm motion
  - Stronger storm-relative low-level inflow
  - Increased nonhydrostatic vertical pressure gradient
    - Due to updraft vorticity
    - Due to interaction of shear and updraft exterior

As updrafts encounter an increasingly sheared environment (e.g., 0-6 km shear >20 m/s), they become enhanced by: increased updraft/downdraft separation, precipitation removal from updraft, lower boundary-relative storm motion, stronger storm-relative low-level inflow, increased nonhydrostatic upward directed pressure forcing due to updraft vorticity, increased nonhydrostatic upward directed pressure forcing due to shear interacting with the updraft boundary.

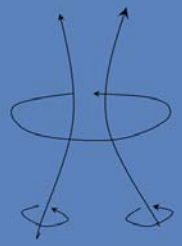
What Describes a Supercell?

- Updrafts last longer than individual air parcel transit time
- Updraft well correlated with vertical vorticity

If an updraft begins to persist for longer than an individual air parcel takes to traverse it, and it is well correlated with significant vorticity, the updraft is then called a **supercell**. Sometimes the effects of nonhydrostatic pressure forcing on updraft strength can exceed that of buoyancy.

Definition of Updraft Rotation and a Vortex

- Vortex is a local maximum collection of vortex lines
- Rotating Updraft is a local concentration of vortex lines with curvature vorticity

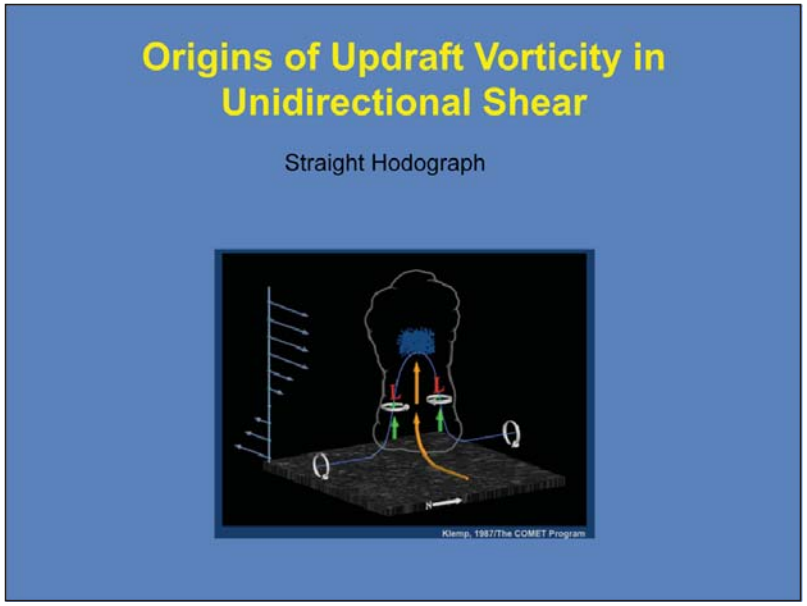


Before we proceed with the following discussion, there are some definitions of vortices, updraft rotation and updraft vorticity that you should know so as to avoid any misconceptions. When we discuss the term 'vortex', we refer to a local concentration of vortex lines. The vortex may or may not be rotational. In other words, a vortex could be a locally intense region of shear vorticity. Likewise a vortex could result from a concentration of vortex lines eminating from curvature vorticity or rotation. The term rotating updraft is a vortex with curvature vorticity. However air parcels within a rotating updraft may not complete a closed circuit as observations have shown. Supercell updraft air trajectories often show anticyclonic curvature even though the vertical vorticity is positive.

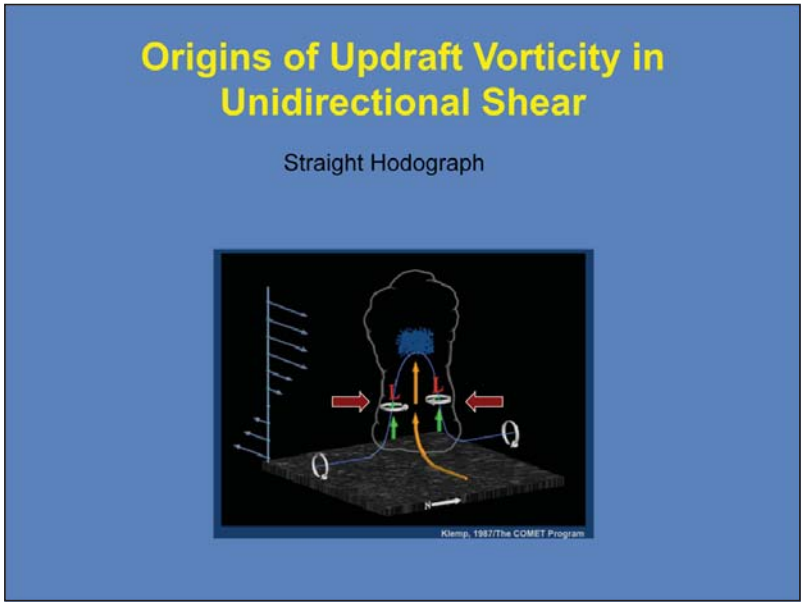
Origins of Updraft Vorticity

- Two models of getting updraft and vorticity correlated
  - Tilting of vorticity in unidirectional shear
  - Interaction of updraft edge to directional shear

The origins of updraft vorticity and storm motion deviant to the steering layer wind can both be explained by how the updraft is influenced by vertical wind shear. There can either be unidirectional or directional vertical shear in supercell environments. Fundamental origins of updraft vorticity and propagation are shared by both straight and curved sheared environments. However, there are important differences in the origins of updraft vorticity and propagation between unidirectional and directional vertical shear. These differences will be covered in this section.



We can visualize vertical shear as a continuous series of vortex lines oriented horizontally. A good analogy is a sheet of rolling logs. **As an updraft extends into a sheared environment, horizontal vorticity tilting acts to create two vertical vortices.** The strength of these vortices depends on the strength of the shear and the intensity of the updraft. Facing toward the direction of the shear, from left to right in Figure , on the right (left) side of the updraft lies a cyclonic (anticyclonic) vortex. Initially in this figure, the vortices lie along the periphery of the updraft, and thus contain no updraft within them. In other words, the updraft and vorticity are not correlated.



Both counter rotating vortices create a dynamic low. The stronger the vortex, the lower the pressure in its center. **Since tilting of the originally horizontal vorticity is most pronounced where the updraft is strongest (at midlevels), the vertical vortices are most intense there.** With the dynamic pressure at its lowest aloft, an enhanced upward directed pressure gradient force promotes the development of new updraft within their centers of rotation. The effect is a widening of the updraft and increasing correlation between updraft and vorticity on both flanks. Updraft strength is also augmented through this process.

Origins of Updraft Vorticity in Unidirectional Shear

Straight Hodograph

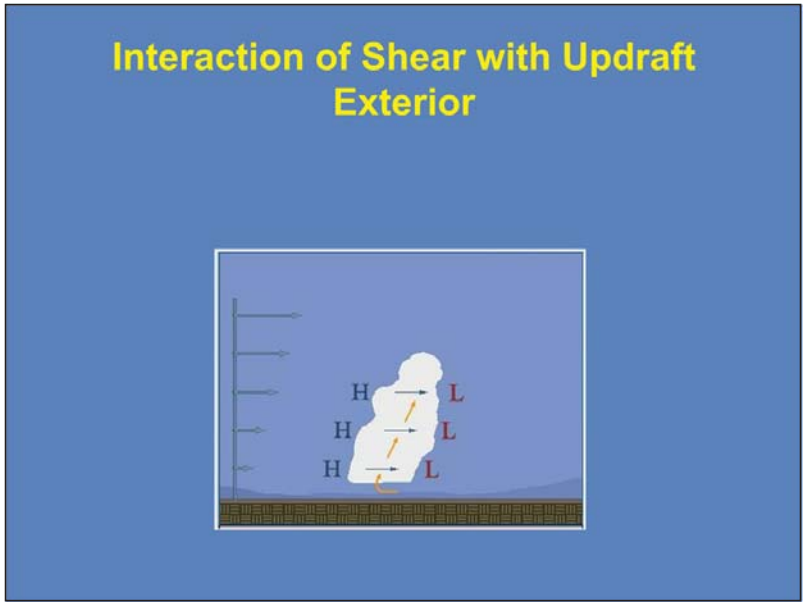


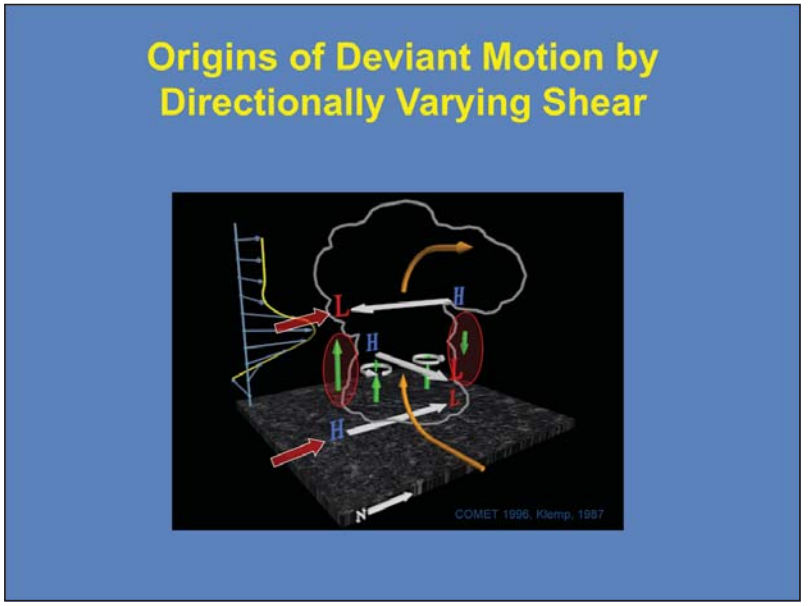
The greatest tilting of horizontal vorticity occurs right and left of the shear vector. **This means that the development of rotation and new updrafts also occur to the right and left of the shear vector.** Precipitation developing in the middle of the widening updraft acts to develop a downdraft which, in turn, helps to split the widening updraft into two parts. **The cyclonically (anticyclonically) rotating member moves to the right (left) of the shear vector.** Since both the cyclonic and anticyclonic updrafts experience similar upward dynamic pressure forcing, they are equally strong supercells in a straight hodograph environment. Once the supercell is deviating off the hodograph, it experiences streamwise vorticity, and storm-relative helicity in its inflow layer. Tilting of the streamwise vorticity into the updraft immediately produces vertical vorticity well correlated with updraft.

Origins of Rotation in Directionally Varying Shear

- Directional shear implies curved hodograph
- Curved hodograph suggests streamwise vorticity

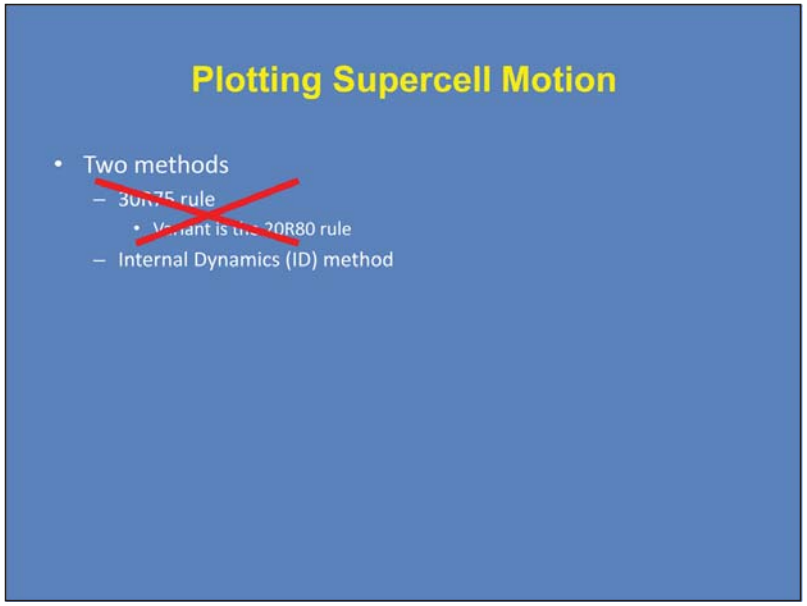
The processes that develop rotation in the unidirectional hodograph, also apply to curved hodographs. However, a curved hodograph implies that streamwise vorticity and helicity are available for the updraft to directly ingest upon its initial growth. Instead of the rolling log analogy to describe the vorticity in the environment, here the analogy is the thrown spinning football. This analogy represents the available streamwise vorticity that merely needs to be tilted into the vertical by the updraft in order for rotation and updraft to be well correlated. Therefore the evolution from ordinary cell to supercell is much quicker.





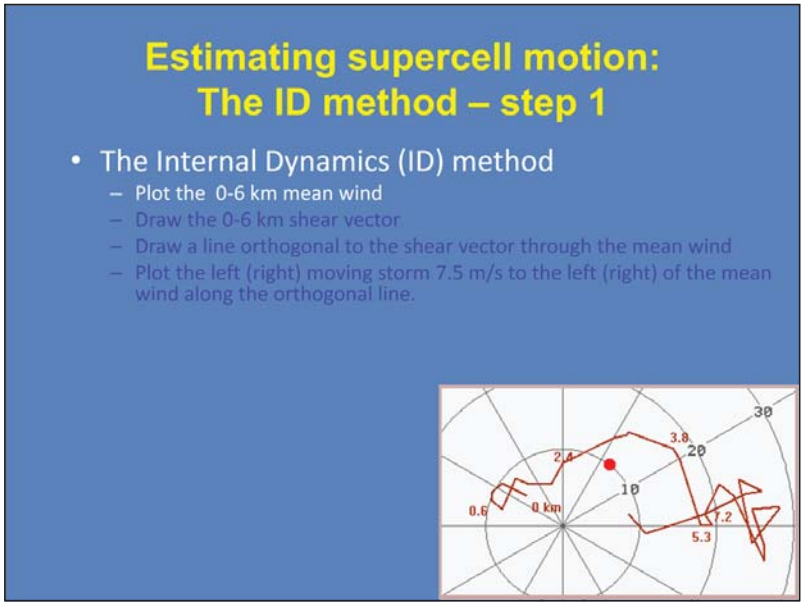
While the same processes that promote deviant motion in unidirectional hodographs will work in curved hodographs, the interaction of the changing shear vector with height will result in additional nonhydrostatic vertical pressure gradient forcing that promotes growth on only one flank of an updraft. This additional process is related to the same processes that force an updraft to tilt in the presence of vertical shear. On the upshear side of an updraft, high dynamic pressure forms as a result of partial flow blockage, while low pressure forms on the other side forcing the updraft to tilt. While this illustration deals with unidirection shear, we will next discuss how directional shear extends this concept to explain the origin updraft deviant motion and preference for the cyclonic member of a supercell to intensify.

When the shear direction changes with height so do the locations of the dynamic pressure maxima and minima. **We know that a dynamic high (low) forms on the up (down) shear side of an updraft.** In the example shown in this figure , the relative high is on the south side of the updraft at low levels. At higher levels, the shear vector pointing south would produce a relative low on the south side of the updraft. The result is an upward directed pressure gradient force that causes new updraft development and therefore, storm propagation to the right of its original motion. Meanwhile, the left side of the updraft would experience a downward directed dynamic pressure gradient force weakening, or even destroying, the side of the updraft containing the anticyclonic member of the rotational couplet. This is why a left-moving storm, given the hodograph in , would be suppressed.



There are two methods for estimating supercell motion for which to be aware. The "legacy Supercell method" and the ID method, which is the preferred technique, are presented next.

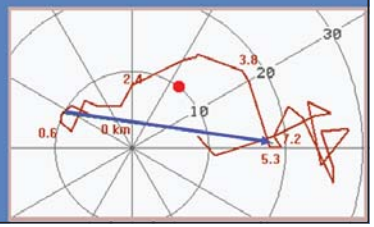
In the past, forecasters often based supercell motion on the 30R75 (Maddox, 1976) or 20R85 (Davies and Johns, 1993) rules. **The 30R75 rule estimates the cyclonically rotating supercell motion by adding 30° to the right of the 0-6 km steering layer flow direction and 75% of the speed**. The 20R85 rule was an adjustment for those supercells embedded in very strong flow. Unfortunately, these estimations are non-physically based and only apply in the Northern Hemisphere with the typical counterclockwise turning hodographs. The AWIPS skew-T program still uses this technique to estimate SRH.



Bunkers et al. (2000) developed a better method called the **ID method (Internal Dynamics)**, which uses the mechanisms by which updraft and shear interact to cause deviant motion. **This method can be used to calculate storm motion for both the cyclonically and anticyclonically rotating supercells resulting from a storm split**. The ID method is Galilean invariant allowing for its use in atypical hodographs (i.e., westerly shear with northerly mean winds). To estimate supercell motion using the ID method, the following steps work well: 1) Plot the 0-6 km non-pressure-weighted mean wind. An example in this Figure shows the mean wind as a red dot.

Estimating supercell motion

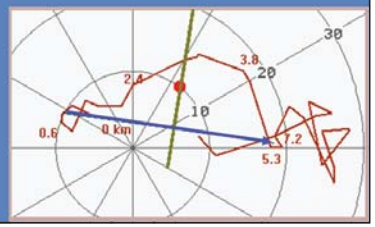
- The Internal Dynamics (ID) method
  - Plot the 0-6 km mean wind
  - Draw the 0-6 km shear vector
  - Draw a line orthogonal to the shear vector through the mean wind
  - Plot the left (right) moving storm 7.5 m/s to the left (right) of the mean wind along the orthogonal line.



Draw the shear vector from the mean wind in the lowest 0.5 km to the mean wind from 5.5-6 km.

Estimating supercell motion

- The Internal Dynamics (ID) method
  - Plot the 0-6 km mean wind
  - Draw the 0-6 km shear vector
  - Draw a line orthogonal to the shear vector through the mean wind
  - Plot the left (right) moving storm 7.5 m/s to the left (right) of the mean wind along the orthogonal line.

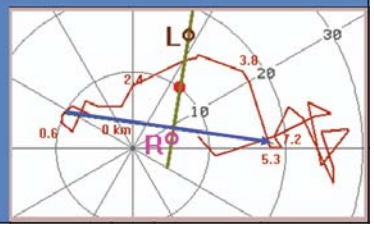


Draw a line orthogonal to the shear while passing through the mean wind. Note that the shear vector can be placed anywhere on the hodograph as long as it retains the same direction and magnitude.

Estimating supercell motion

- The Internal Dynamics (ID) method
  - Plot the 0-6 km mean wind
  - Draw the 0-6 km shear vector
  - Draw a line orthogonal to the shear vector through the mean wind
  - Plot the left (right) moving storm 7.5 m/s to the left (right) of the mean wind along the orthogonal line.

Bunkers et al. (2000)

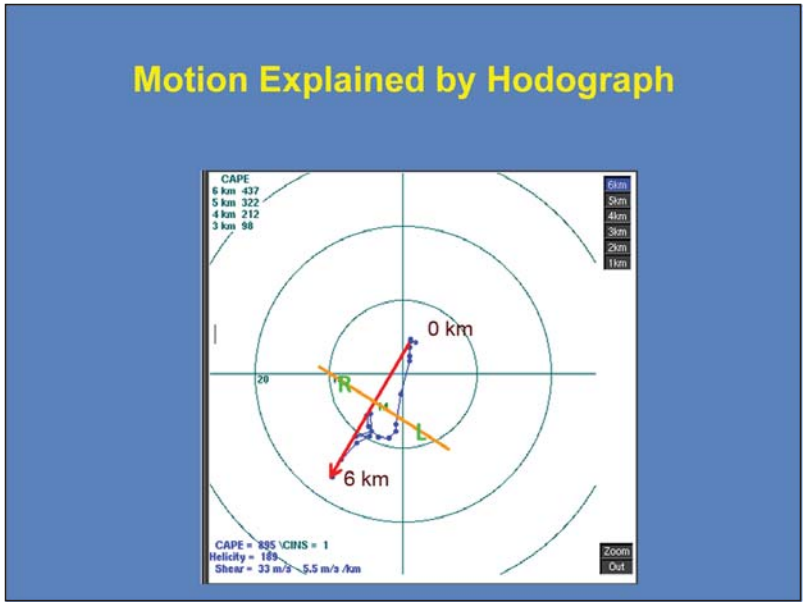


The right- (left-) moving supercell is drawn 7.5 m/s to the right (left) of the shear vector where shear vector intersects the shear-orthogonal line at the 0-6 km mean wind. Note that the storm motion remains on the shear-orthogonal line.

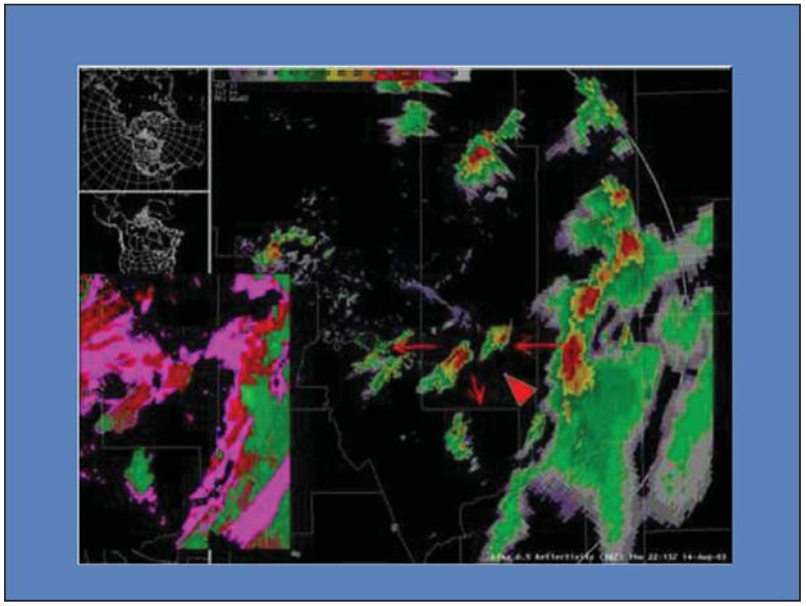
Limitations in ID Method

- ID Method contains uncertainties
- What is the "best" Deviant motion?
- How is it Modulated?
- Propagation Effects due to Boundaries

ID Method contains uncertainties. For example, what is the "best" Deviant motion? 7.5 m/s was chosen on a representative sample size. How is motion modulated? What about propagation effects due to boundaries?



This is the author's explanation of one of the previous quiz items.



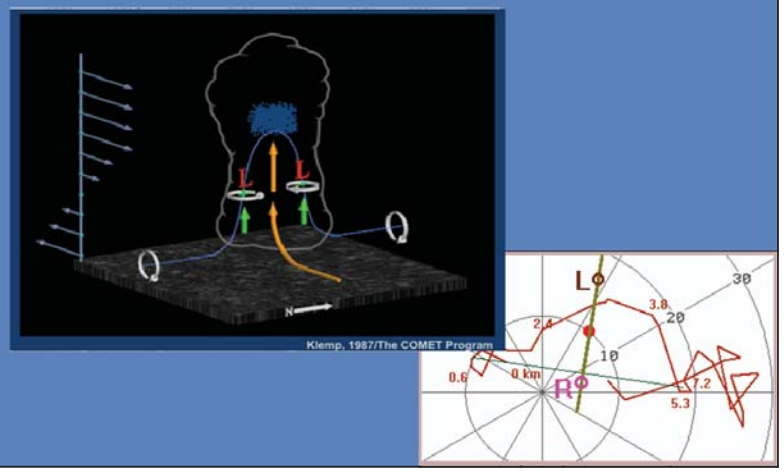
This is a radar loop (0.5 deg Z and 0.5 deg SRM in the lower left) from KFGZ of the resulting storm motions from the hodograph just analyzed.

Summary 1

- Strength of wind shear (> 20 m/s through 6 km) defines supercell longevity
  - Increased updraft/downdraft separation
  - Precipitation removal from updraft
  - Lower boundary-relative storm motion
  - Stronger storm-relative low-level inflow
  - Increased nonhydrostatic vertical pressure gradient
    - Due to updraft vorticity
    - Due to interaction of shear and updraft exterior
- Persistent updraft vorticity defines a supercell

This is the summary slide part 1.

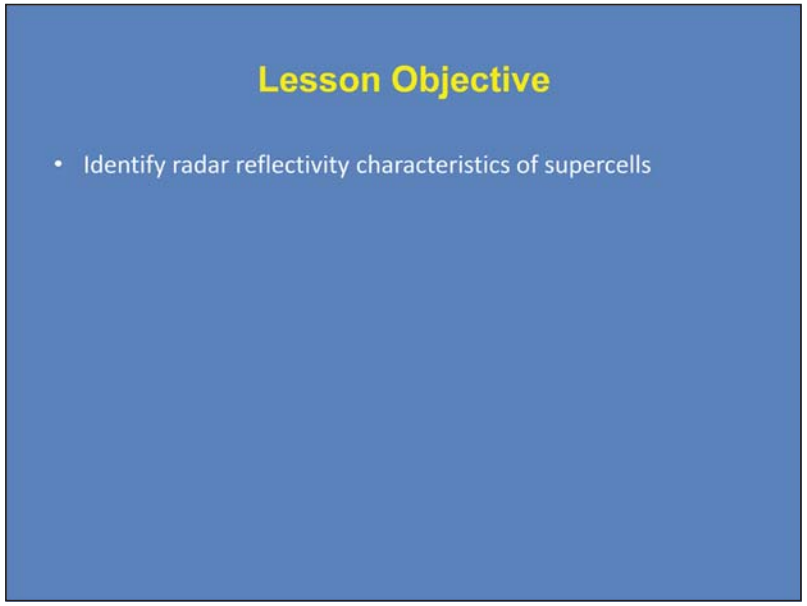
Summary 2



Summary slide part 2.

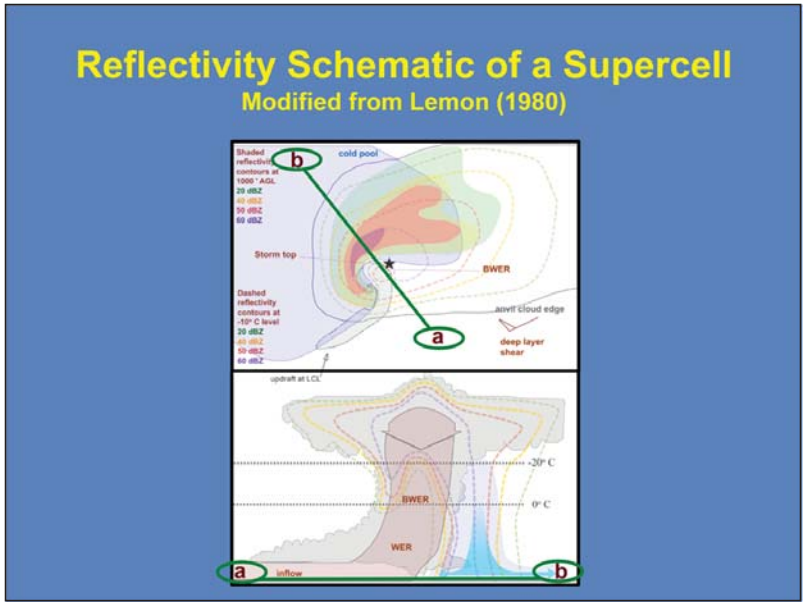


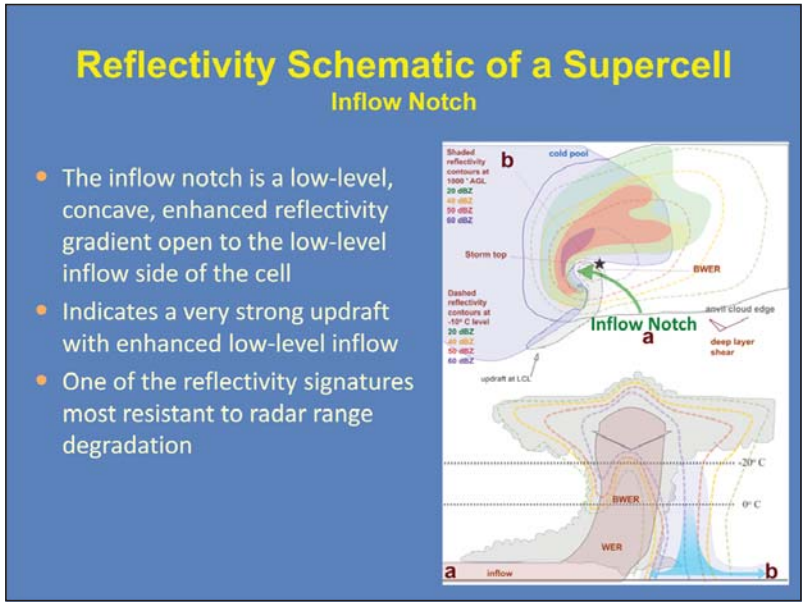
Welcome to the Convective Storm Structure and Evolution's lesson on Supercell Morphology: Radar Reflectivity Signatures.



The persistence and strength of a supercell thunderstorm updraft yields a distinctive appearance to its precipitation distribution. This lesson describes the common radar-based reflectivity characteristics associated with supercells.

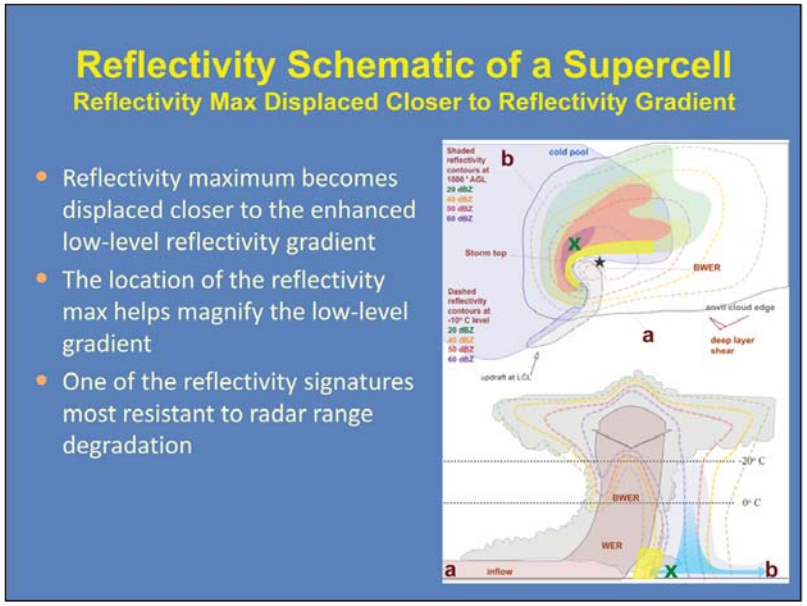
The objective is: Identify radar reflectivity characteristics of supercells.



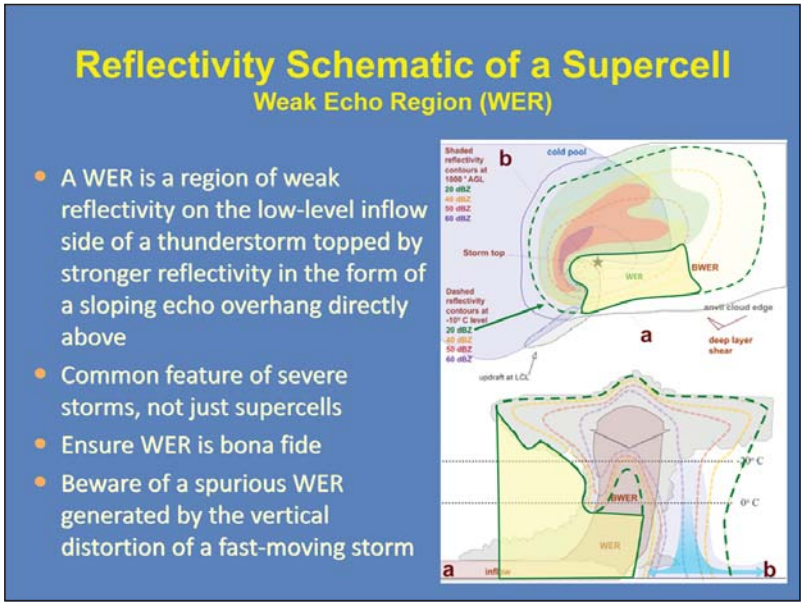


In 1980, Lemon identified radar reflectivity characteristics associated with supercells. The schematic shown here is a conceptual model of the reflectivity structure of a cyclonic, right-moving supercell in the northern hemisphere, modified from his paper. The horizontal plane is shown at the top, while a vertical cross-section shown at the bottom. The letters a and b denote the endpoints of the vertical cross section. Let's discuss those radar reflectivity characteristics using this schematic.

The most common radar reflectivity characteristic of supercells is the **"Inflow Notch"** which is a low-level, concave, enhanced reflectivity gradient open to the low-level inflow side of the cell. This signature indicates the presence of a very strong updraft with associated enhanced low-level inflow. If the storm is close to the radar, a surface trailing gust front may be seen wrapping into the region of the notch. The inflow notch is one of the reflectivity signatures most resistant to radar range degradation.



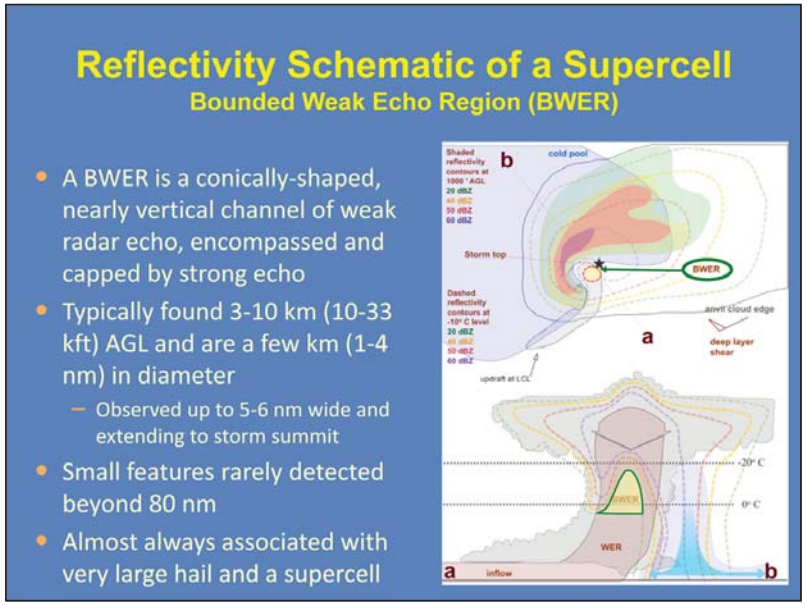
Another characteristic that indicates a storm has transitioned into a supercell is **the reflectivity maximum becomes displaced closer to the enhanced low-level reflectivity gradient.** The location of the reflectivity maximum helps magnify the low-level gradient. The enhanced reflectivity gradients is one of the features most resistant to radar range degradation.



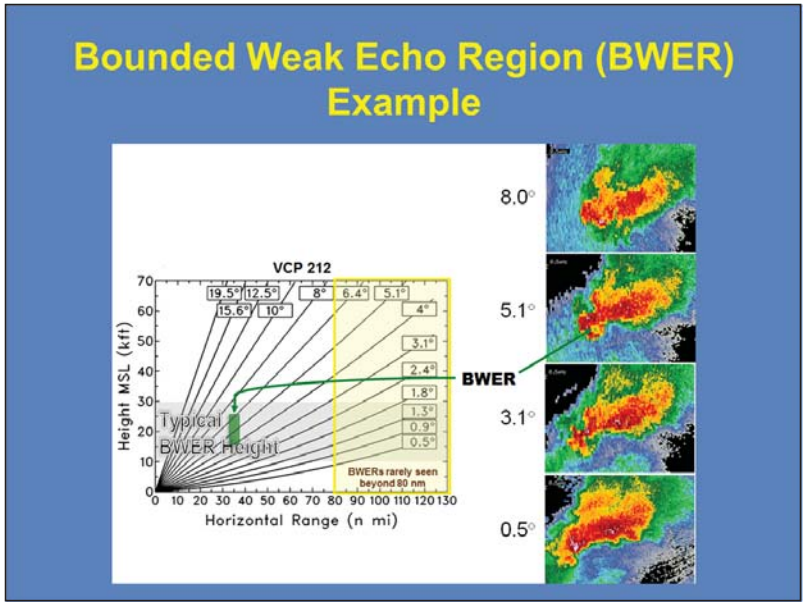
A **Weak Echo Region (WER)** is a region of weak reflectivity on the low-altitude inflow side of a thunderstorm topped by stronger reflectivity in the form of a sloping echo overhang directly above. The WER is produced by strong updraft and associated strong storm-summit divergence that carries large amounts of precipitation particles in all directions creating a high reflectivity echo-canopy (slopping echo overhang) over the low-level inflow of a strong or intense convective storm. The slopping nature of the overhang is created when precipitation begins to fall from the far edges of the overhang (visually, the edge of the thick anvil) and descends through the storm relative environmental winds finally reaching the ground in the strong low-level reflectivity echo. A WER is a common feature of severe storms in vertically sheared environments, not just with supercells. Note that the key ingredient that distinguishes storms is the strongly sheared environment. Therefore, features such as the WER, are not found with storms in a weakly sheared environment, such as the pulse storm.

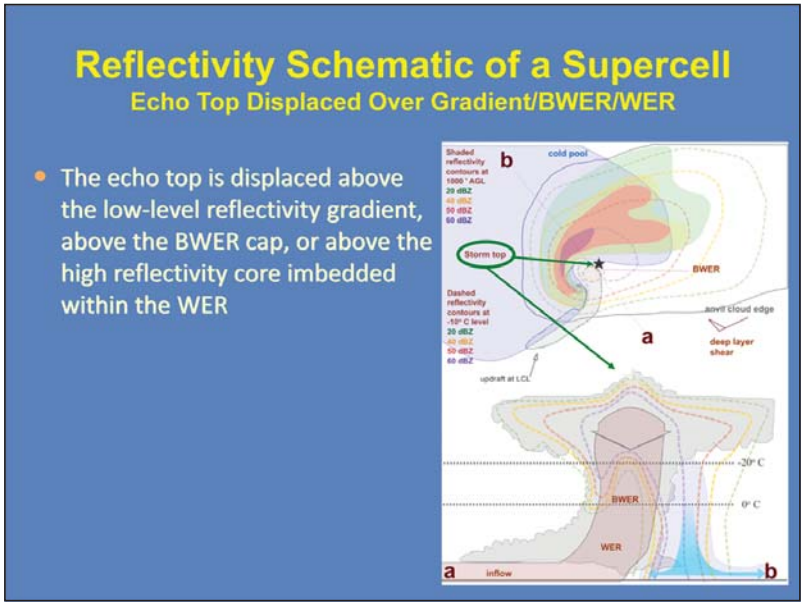
Care must be taken to ensure that a WER is on the updraft and inflow flank of the storm. A bona fide WER should be persistent (~ 10 minutes) and capped by high reflectivities (>45 dBZ) with the base of the slopping overhang beginning as high as the -20 to -30 degrees Celsius environmental temperature. False WERs not capped by strong reflectivity imply weak updraft, such as with an overspreading anvil layer. In addition, the WER should be found above the low-level inflow notch and strong reflectivity gradient.

Because a radar's volume coverage pattern (VCP) samples a storm from bottom to top, beware of a spurious WER oriented in the direction of storm motion generated by the vertical distortion of a fast-moving storm. For example, a storm moving at a speed of 60 kts can have its upper-level scans displaced up to 5 miles in the direction of storm motion.



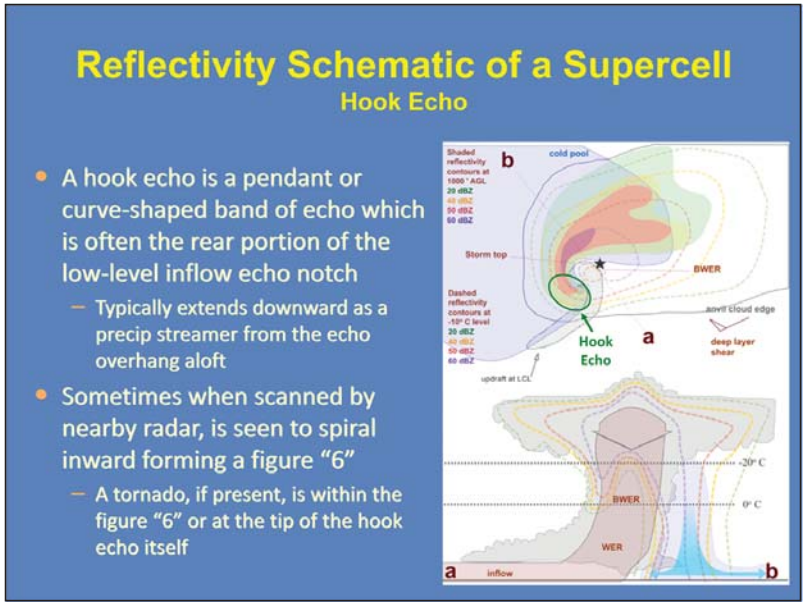
A **Bounded Weak Echo Region (BWER)** (also known as a "vault") is a conically-shaped, nearly vertical channel of weak radar echo, encompassed and capped by strong echo. The cap is composed of large concentrations of supercooled liquid water and rapidly growing hail. The BWER is the core of an intense updraft that carries newly formed hydrometeors to high levels before they can grow to radar-detectable sizes. BWERs are typically found imbedded in the slopping echo overhang and aloft above the apex of the low-level inflow notch. They are typically found 3-10 km (10,000 - 33,000 ft.) AGL and are a few kilometers (1-4 nm) in horizontal diameter. However, on rare occasions, they have been observed up to 5-6 nm wide and extending to storm summit. BWERs are small features rarely detected beyond 80 nm due to radar resolution limitations. The presence of a BWER is almost always associated with very large hail and is associated with a supercell. Note that the BWER is not associated with updraft rotation.



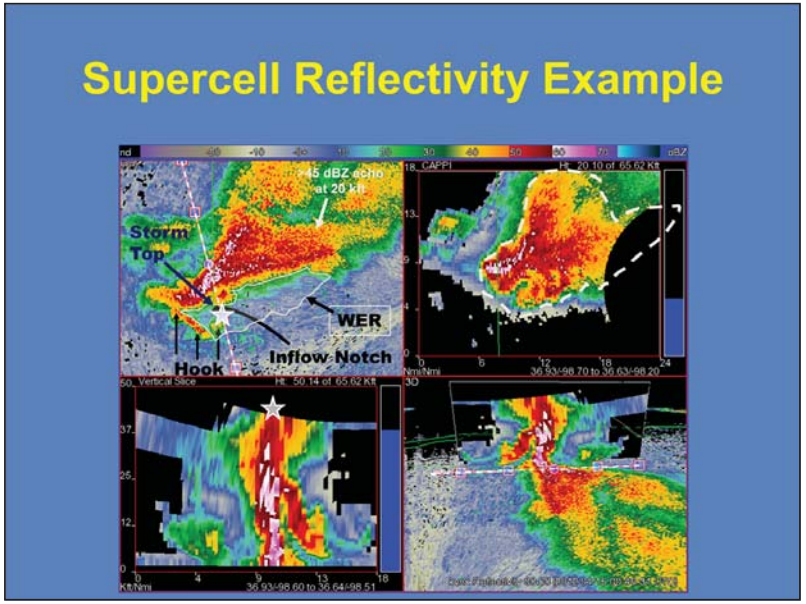


Here is an example of a BWER observed with a supercell 35 nm to the west-southwest of the radar. It is located between 15,000 to 26,000 feet above ground level which certainly falls within the typical BWER range of 10,000 to 33,000 feet above ground level. Remember that BWERs are rarely detected beyond 80 nm due to radar resolution limitations.

Another supercell reflectivity characteristic identified by Lemon is that the **echo top is displaced above the low-level reflectivity gradient, above the BWER cap, or above the high reflectivity core imbedded within the WER.**

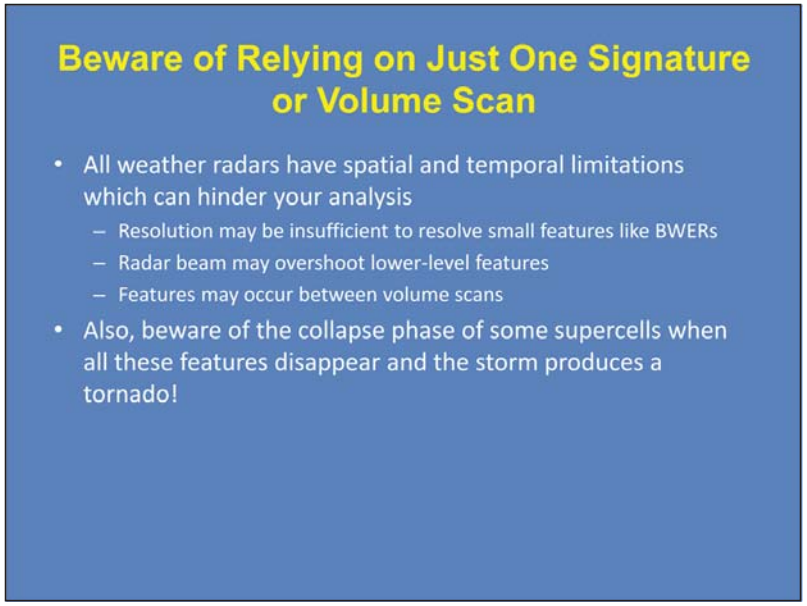


The final supercell reflectivity characteristic identified by Lemon is a **"hook echo"** which is a pendant or curve-shaped band of echo which is often the rear portion of the low-level inflow echo notch. It typically extends downward as a precipitation streamer from the echo overhang aloft. It is often a portion of echo bounding the BEWER on the rear. It may also be precipitation carried downward rapidly by the RFD or associated with the storm mesocyclone. Sometimes when scanned by nearby radar, is seen to spiral inward forming a sharply defined figure "6." A tornado, if present, is within the figure "6" or at the tip of the hook echo itself.

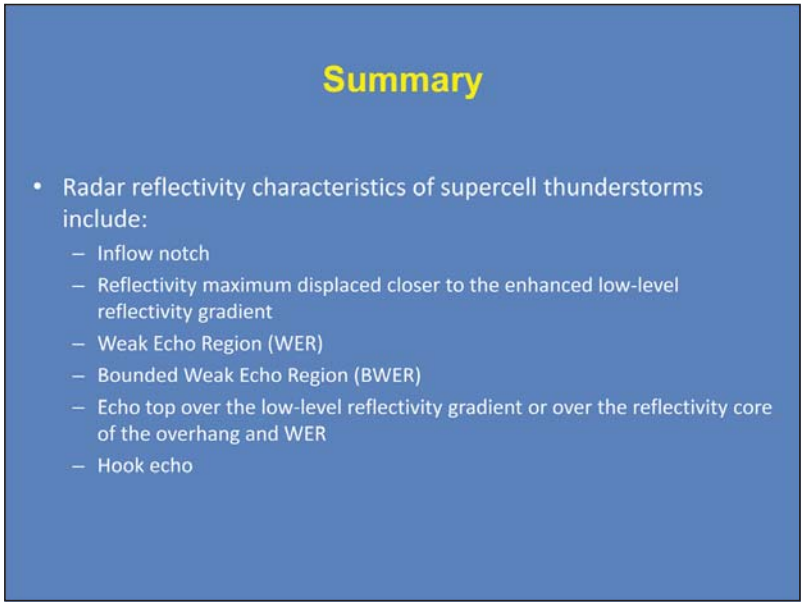


Here's a Base Reflectivity FSI example of the Cherokee, Oklahoma tornadic supercell of April 15, 2012 at 0046 UTC sampled from the Vance AFB WSR-88. It exhibits many of the characteristics we've discussed.

- An **"Inflow Notch"** was located on the low-level inflow side of the cell. This signature indicated the presence of a very strong updraft with associated enhanced low-level inflow.
- **The reflectivity maximum (in pink) displaced close to the enhanced low-level reflectivity gradient.**
- A **Weak Echo Region (WER)** was apparent as a region of weak reflectivity on the low-altitude inflow side of the storm topped by a sloping high-reflectivity echo-canopy directly above.
- The Cherokee supercell did not exhibit a well-defined, vertically-oriented, conically-shaped **Bounded Weak Echo Region (BWER)**, perhaps due to the rapid, cyclic tornado-genesis which occurred with this storm.
- However, the **echo top was displaced above the low-level reflectivity gradient, above the high reflectivity core imbedded within the WER.**
- And finally, a **"hook echo"** was evident as pendant shaped band of echo, spiraling inward to form a sharply defined figure "6." A tornado was located within the circular portion of the figure "6" on this volume scan.
- Bear in mind that both the reflectivity schematic and this example depict a cyclonic, right-moving supercell in the northern hemisphere. An anticyclonic, left-moving supercell in the northern hemisphere would appear as a mirror image of this.



Beware of relying on any one signature or volume scan in isolation when trying to identify a supercell. All weather radars have spatial and temporal limitations which can hinder your analysis of storm structure. Radar resolution may be insufficient to resolve smaller features at longer ranges such as BWERs or even hook echoes. The radar beam may overshoot lower-level features such as some hook echoes and WERs. Features may occur between volume scans. Plus, this lesson doesn't discuss deviant motion from the mean wind which is perhaps the most easily identifiable and reliable supercell characteristic. Finally, beware of the "collapse phase" of some supercells when all of these distinctive features disappear and the storm produces a tornado!



Radar reflectivity characteristics of supercell thunderstorms include:

- A strong reflectivity gradient bounding a concavity or "inflow notch"
- Reflectivity maximum displaced closer to the enhanced low-level reflectivity gradient
- Weak Echo Region (WER)
- Bounded Weak Echo Region (BWER)
- Echo top over the low-level reflectivity gradient or over the reflectivity core of the overhang and WER
- Hook echo

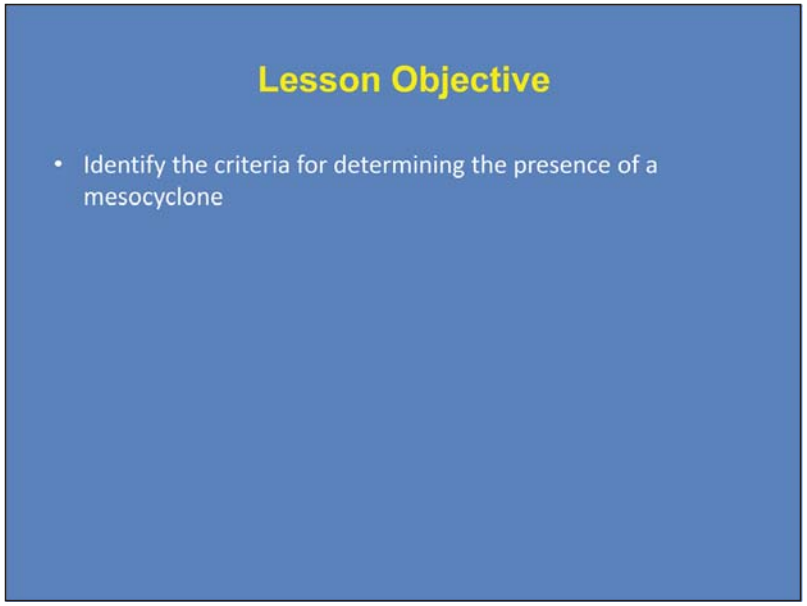
For Additional Help

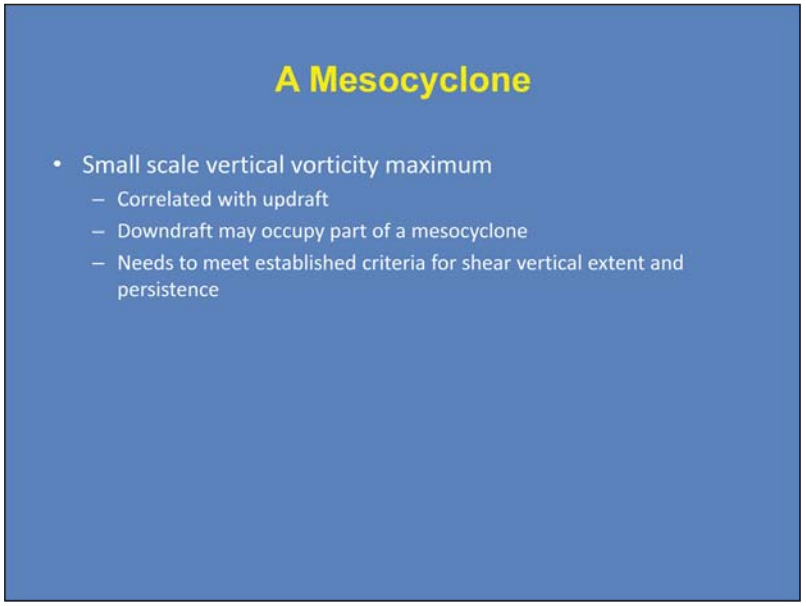
1. Check with your facilitator (typically your SOO)
2. Send your questions to:
   nws.wdtd.rachelp@noaa.gov

For additional help, check with your facilitator (typically your SOO) or send your questions to the listserv e-mail address here.



Welcome to the Convective Storm Structure and Evolution's lesson on supercell velocity signatures. This lesson is approximately 20 minutes long.

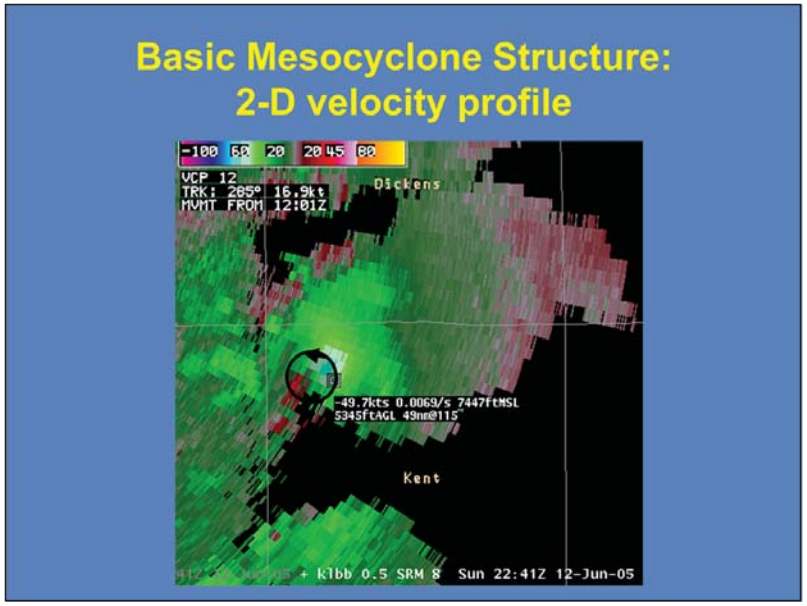




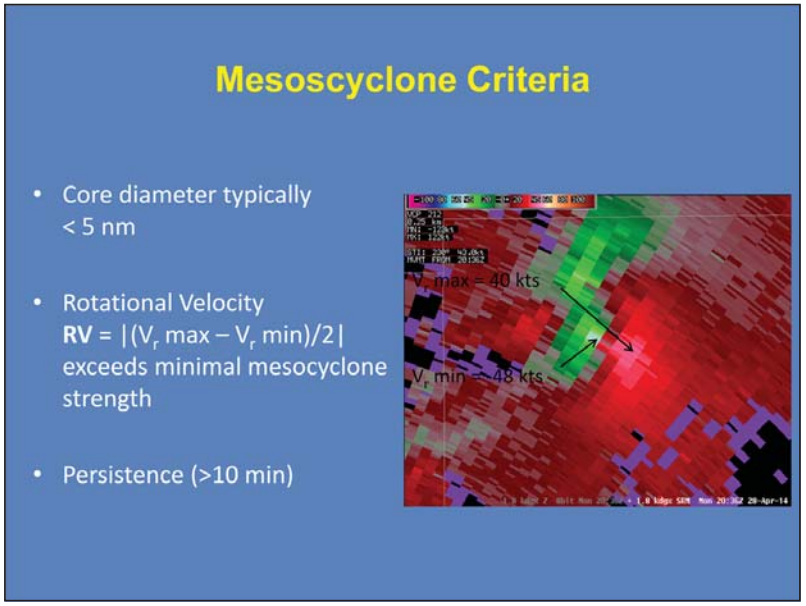
In many cases, a warning forecaster may have difficulty in distinguishing an ordinary from a supercell based solely on the reflectivity pattern. This is one of the major reasons that we have the WSR-88D network; to identify velocity patterns that compliment reflectivity by playing a critical role in identifying a supercell. This lesson describes the structure and morphology of supercell velociy signatures; focusing on aspects of the mesocyclone structure.

A localized region of vertical vorticity partially, or fully embedded within an updraft of a DMC (termed a mesocyclone) is one of the defining characteristics of a supercell. **By definition, a mesocyclone is a small-scale vertical vorticity maximum closely associated with the updraft, and downdraft, of a convective storm that meets or exceeds established criteria for shear, vertical extent and persistence**. Each of these criteria will be discussed.

Identify the criteria for determining the presence of a mesocyclone



Mesocyclone velocity structure is similar to that of a Rankine Combined Vortex. The core of the mesocyclone rotates as a solid body with the tangential velocity proportional to radius. Beyond this core, the velocity decreases exponentially with increasing radius from the mesocyclone center. Since only the radial velocity component is detectable from Doppler radar, only the radial components of the velocity can be detected. Therefore, the mesocyclone appears as a range adjacent couplet of inbound and outbound velocity.

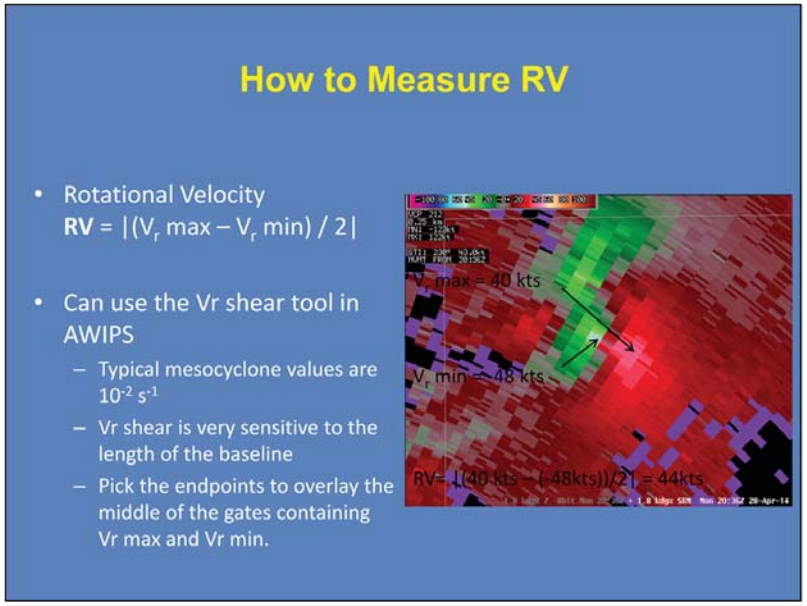


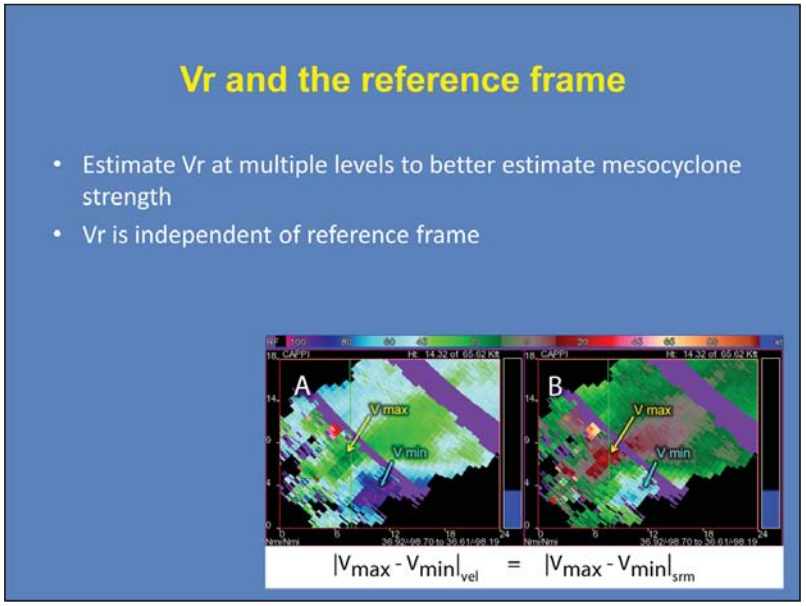
To establish the validity of a mesocyclone, we use a set of criteria for shear persistence and vertical depth. A circulation feature is labeled a mesocyclone when:

The core diameter (distance between the maximum and minimum velocities) is < 5 nm, and

The rotational velocity (RV = |(Vr max – Vr min|) / 2) equals or exceeds minimal mesocyclone strength. **Vr** min (Vr max) is the minimum (maximum) radial velocity in the circulation.

The feature persists for at least 10 minutes.





The inputs into calculating RV should represent the maximum and minimum velocities as illustrated by the inset in this figure. Note that the minimum and maximum velocities that contribute to the calculation of rotational velocity should be measured using representative peak values of the data levels in the velocity or SRM products.

In this example, the maximum Vr is 40 kts for a good representative value, and the minimum Vr is -48 kts. Taking the difference, dividing by two, and taking the absolute magnitude reveals a RV of 44 kts.

Vr shear is calculated by dividing RV by the distance between Vmin and Vmax. It can be easily calculated using AWIPS using the Vr shear tool. Values are on the order of 10**-2 s**-1 for mesocyclones. However, Vr shear can change by orders of magnitude just by changing the baseline distance without any actual increase in mesocyclone intensity. **Therefore, V**r shear should be calculated with great caution and consistency through successive volume scans. Pick the endpoints to overlay the middle of the gates containing Vr max and Vr min. Be aware of that you will need to adapt your baseline as the actual mesocyclone diameter changes. Estimating mesocyclone strength from RV alone is just as valid as that from Vr shear.

Note that estimating mesocyclone strength is more representative when assessing Vr from multiple levels rather than one alone.

A mesocyclone need not have in- and outbound velocities. The velocity difference, rotational velocity and shear across a mesocyclone are identical no matter the motion of the reference frame. Using the example in Figure 7-92, a forecaster may sample different velocity maxima and minima between the velocity and the SRM product for the mesocyclone moving toward the radar to the northeast. However, the Vr is identical.

Issues with Mesocyclone Criteria

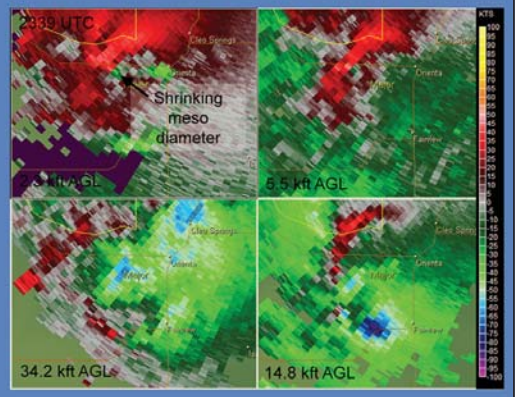
- Vr degrades when
  - Range increases
  - Mesocyclone size decreases
- Vertical criteria discriminates real mesocyclones from 2-D circulations

Establishing a minimal rotational velocity threshold requires knowledge of the distance of the feature, and the size of the supercell. **As radar sampling resolution degrades either by distance or by circulation size, the warning forecaster must reduce the minimal rotational velocity that discriminates mesocyclones from weaker circulations**.

The vertical criteria are required because of the number of shallow circulations uncorrelated with deeper vertical velocity features. Deep, vertically correlated circulations are most likely associated with updrafts and downdrafts because of vertical vortex stretching and advection of vorticity.

Mesocyclone Lifecycle

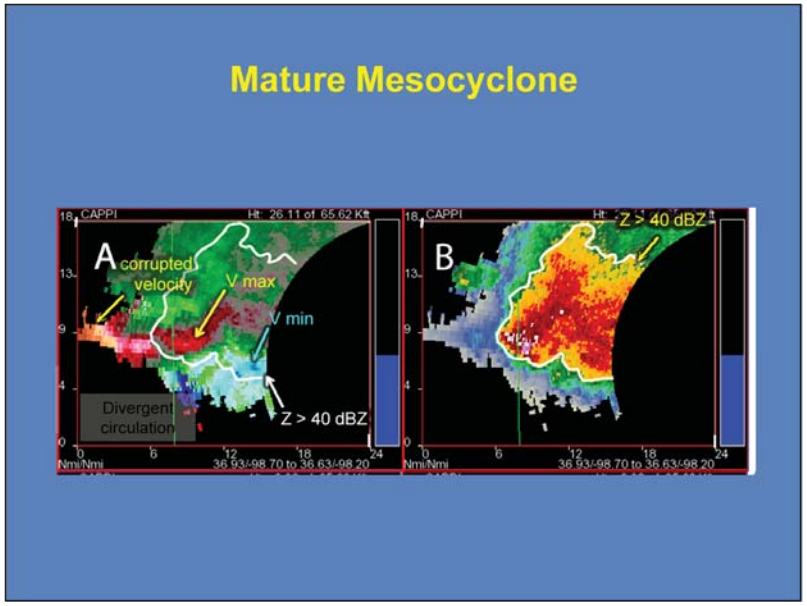
- Organizing
- Mature
- Dissipating



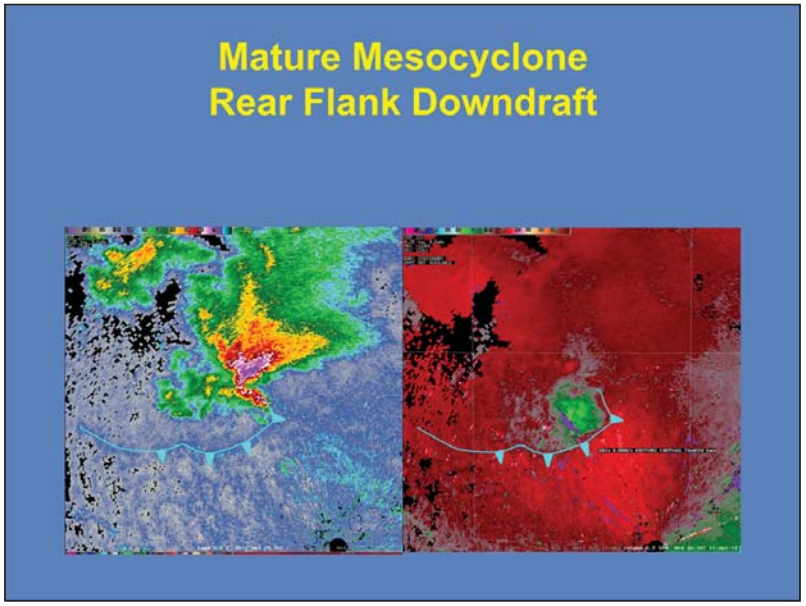
Mesocyclones typically undergo a life span where there is an organizing stage, mature stage and dissipating stage. The typical organizing mesocyclone begins at the level of maximum tilting or in the mid-levels of an updraft. The mesocyclone then begins to build downward and upward. The mid-level mesocyclone is dominated mostly by updraft. If the radar is close enough to the circulation, a convergent signature may be detected in association with the mesocyclone in the lowest slices.

An idealized mature mesocyclone has significant low-level convergence (panel 'G' in Figure 7-93), nearly pure rotation at mid-levels (panels 'C' and 'E' Figure 7-93), divergent rotation at upper-levels (panel 'A' in Figure 7-93).

In the decaying phase of a mesocyclone, the convergent rotational signature in the low-levels gradually transitions to that of divergent rotation as outflow begins to dominate. Mesocyclone depth decreases as does the maximum rotational velocity. As the mesocyclone weakens, it also broadens and becomes diffuse. **If the mesocyclone is tornadic and undergoes a dissipating stage, the tornado could persist for a period of time after all evidence of the parent mesocyclone has dissipated.**

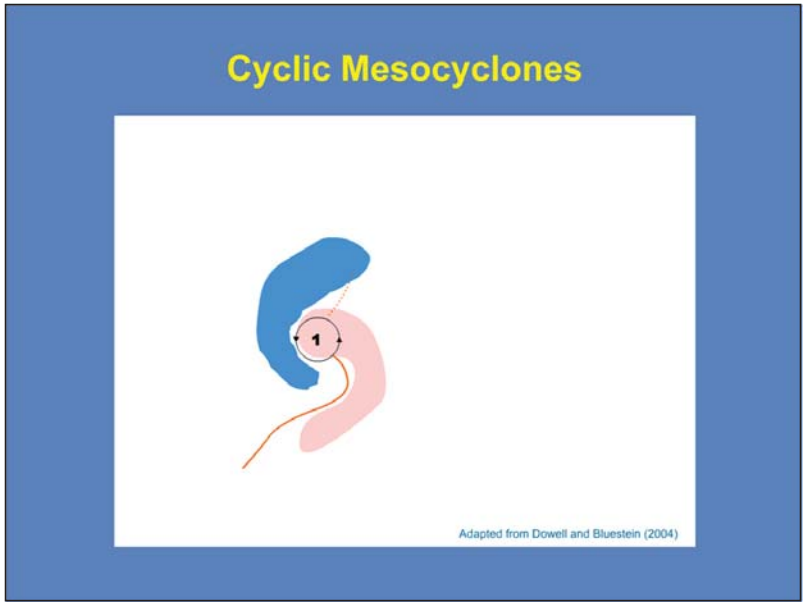


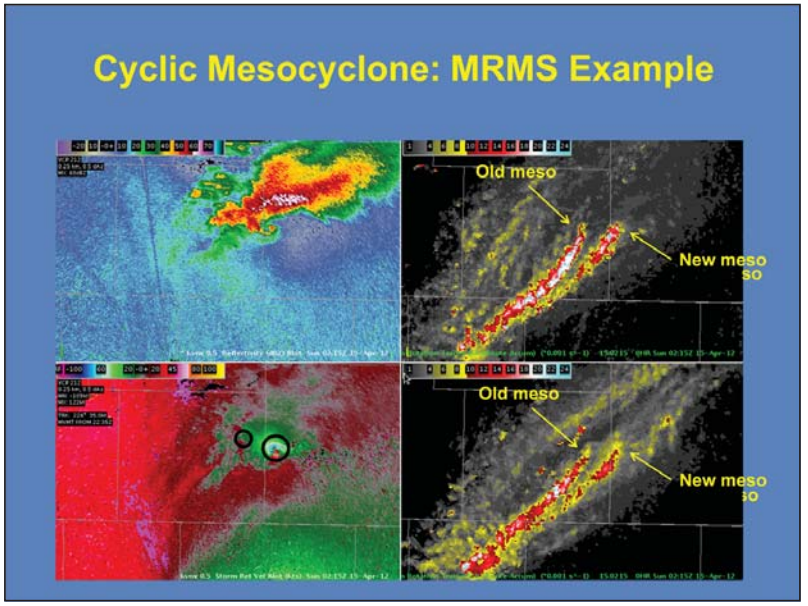
The example in Figure 93 shows a little more complexity than the ideal model. This is because the vorticity from the occluding low-level mesocyclone has been advected upward by the updraft within the larger mid-level mesocyclone producing an interior couplet of peak velocities.



The lower half of a mature mesocyclone is occupied by the rear flank downdraft, usually on its trailing side. The rear flank downdraft can be marked by the presence of strong localized convergence between the inbound to outbound velocities (Figure 7-94). Do not confuse the gust front with the RFD itself. The RFD is often associated with the hook, or pendant echo, and is a divergent outflow that creates the gust front. But this RFD divergence is often difficult to identify in contrast to itsassociated gust front.

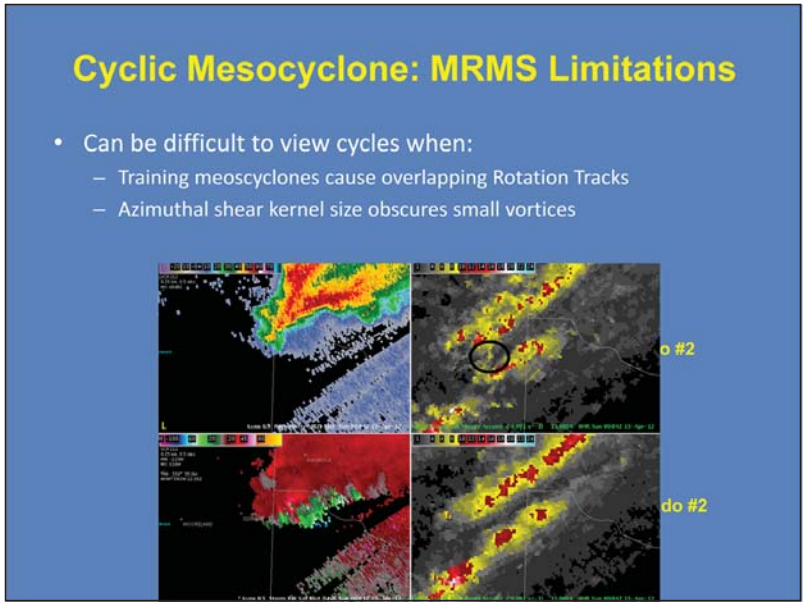
Additionally, the convergence along the RFD gust front should not be mistaken for the transition from in- to outbound velocities in a symmetric vortex.



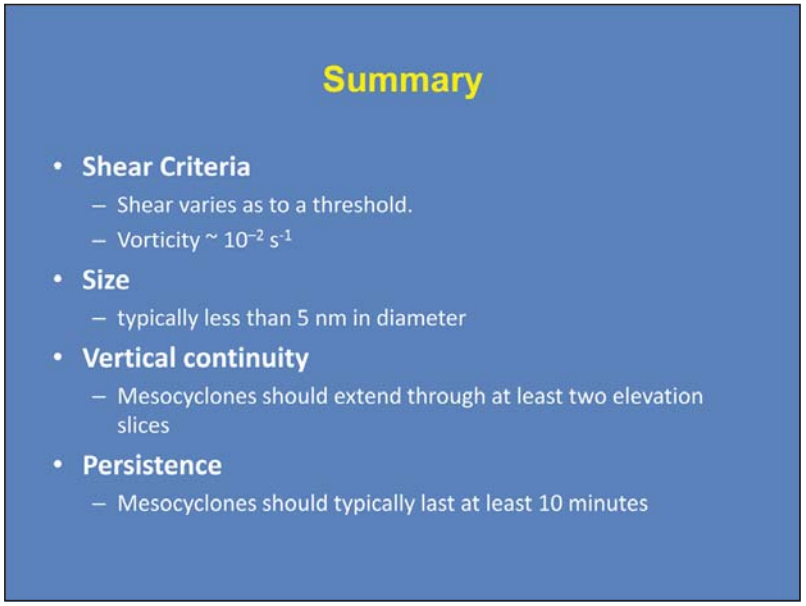


A supercell may produce more than one mesocyclone during its lifetime. In the flash movie coming up, you will see the first mesocyclone mostly occupied by updraft (denoted by the orange filled area) with downdraft on its backside (denoted by the blue area). The first mesocyclone typically takes the longest time to mature as the supercell remains outflow deficient. **Successive mesocyclones mature much more rapidly as they have the advantage of stronger lifting and vortex tilting from a stronger gust front** (denoted by the light brown border) . The life spans of successive mesocyclones may or may not be longer than the first one. The first mesocyclone extends to low-levels as the RFD reaches the ground. When the RFD matures, the outflow wraps cyclonically around the center of circulation, eventually closing it off from the inflow. **If the RFD is thermodynamically unstable, the primary mesocyclone can continue for an extended time**. However, the leading edge of the gust front associated with the RFD can quickly produce successive updrafts and mesocyclones owing to increased convergence and vertical low-level vorticity. In turn, the successive mesocyclones become wrapped by local RFD enhancement, and the process continues for possibly several hours. A family of tornadoes is produced in this fashion from one supercell.

Here's an example of a storm with a cyclic mesocyclone. We'll examine it using Multi-Radar/Multi-Sensor (MRMS) Rotation Tracks. At this time, a tornado is ongoing, as seen in the 0.5 degree Z panel. On the Low-Level Rotation Tracks product, this meso is exhibiting very high azimuthal shear levels near the surface. Strong rotation is also evident at mid-levels. At the next timestep, the old mesocyclone is still very strong and likely still producing a tornado. But now another Rotation Track is beginning to form south of the current one, indicating the growth of a new meso. Nine minutes later, the old meso has drastically weakened and arced to its left. It is probably not tornadic at this point. The new meso, however, as explained on the previous slide, has quickly strengthened into a large, tight circulation that could be tornadic. Finally, just five minutes later, the first meso has almost completely dissipated and the new one has taken over and is most likely producing a tornado.



MRMS Rotation Tracks are not able to help identify cyclic mesocyclones in every situation. It can be difficult to view cycles when training meoscyclones cause overlapping Rotation Tracks and when the MRMS azimuthal shear kernel size obscures small vortices. This is the same supercell we examined on the previous slide, but at this point in time it cycles through 2 tornadoes without obvious signatures. A brief tornado was reported at this time, but the velocity couplet is being partially obscured by beam blockage and MRMS Rotation Tracks don't show strong azimuthal shear. Five minutes later, the couplet is still obscured and the Rotation Tracks have belatedly noted some strong rotation. This delay could have been due to the relatively fast speed of the storm. Five minutes after that, another brief tornado was reported. Again, the vortex is so small the azimuthal shear kernel doesn't register it. The velocity couplet is visible but still partially obscured and another possible mesocyclone appears to its west. By our last timestep, the original couplet still isn't sampled well and there still appears to be another to its west. Also note that there's been some uncertainty of the northern edge of this storm's Rotation Track because of its close proximity to a previous one.



The criteria for determining a mesocyclone is:

Shear Criteria

Shear varies as to a threshold. Some mesocyclones may be poorly resolved and yet still carry considerable severe weather risk.

Mesocyclone vorticity lies on the order of 10–2 s-1

Size

Mesocyclones are typically less than 5 nm in diameter

Vertical continuity

Mesocyclones should extend through at least two elevation slices

Persistence

Mesocyclones should typically last at least 10 minutes



Lesson Objective

- Identify WSR-88D polarimetric characteristics of supercells including:
  - Mid- and upper-level signatures
    - ZDR ring, low CC ring, low CC updraft
  - Low-level signatures
    - ZDR arc, hail signature

As the supercell produces distinctive patterns in reflectivity, the WSR-88D polarimetric variables like ZDR, CC and KDP, also produce their own distinctive signatures.

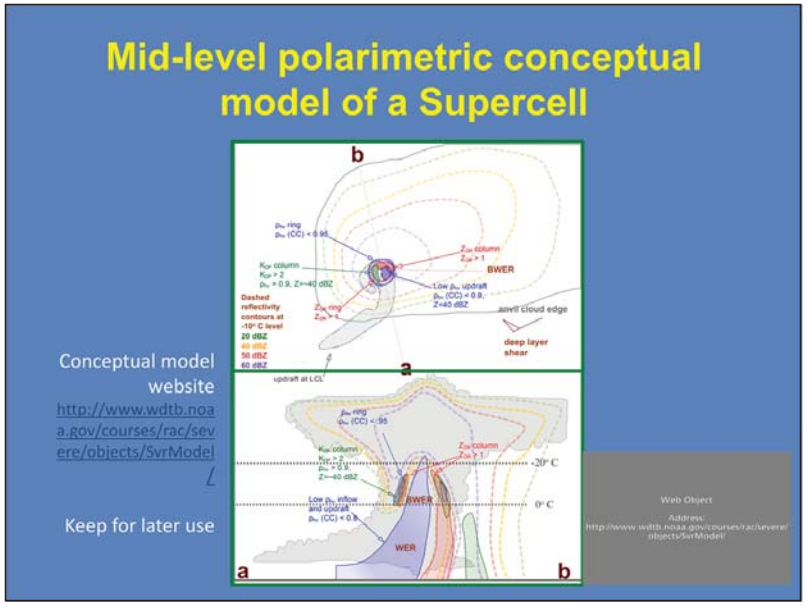
The objective is: Identify WSR-88D polarimetric characteristics of supercells.

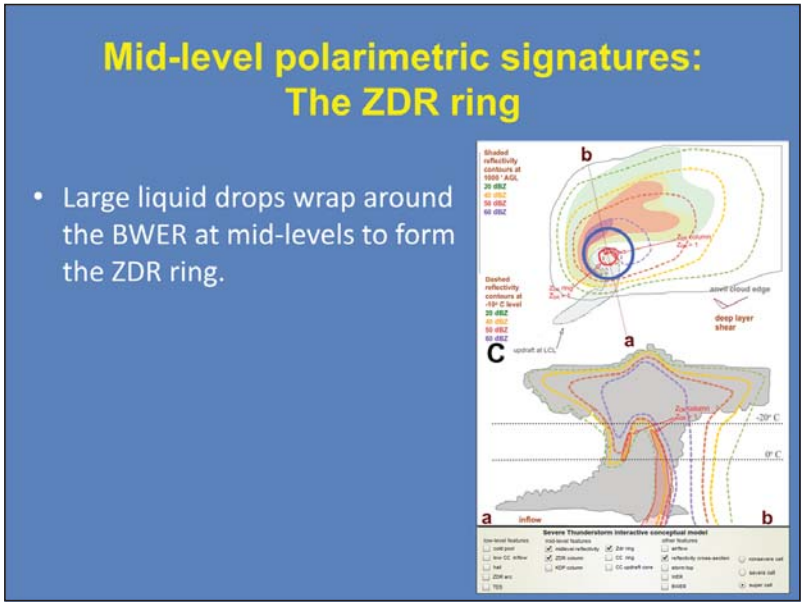
Mid- and upper-level signatures

ZDR ring, low CC ring, low CC updraft

Low-level signatures

ZDR arc, hail signature





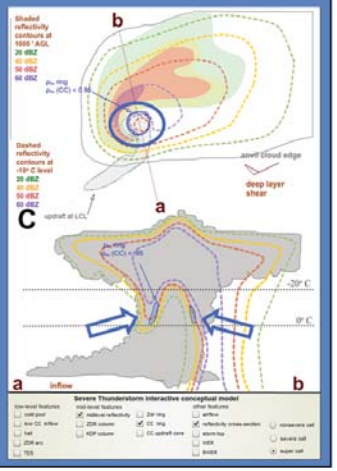
With dual-pol WSR-88D we now have the opportunity to observe important micro-physical structures that are unique to supercells. Detecting these structures is critical to successful warning decision making. This lesson overlaps a little with the lesson on dual-pol updraft signatures, and then it expands to describe other common radar-based dual-pol characteristics associated with supercells.

This conceptual model shows the dual-pol signatures that commonly accompany supercells at mid-levels within or close to the updraft. The flash version of this model will appear shortly and we'll frequently use it as we go on through this lesson. Since we've covered some common signatures with other lessons, we'll focus on the ones that most likely appear with supercells.

The ZDR ring often surrounds a BWER between the environmental freezing and -20 deg C levels. The reason that a ring forms has been attributed to the strong rotation within the updraft. Large water droplets rising within the updraft's outer edges advect around the center of the mesocyclone resulting in a ring-like structure. Not all supercells exhibit a ZDR ring and others exhibit a partial ring.

Mid-level polarimetric signatures: The low CC ring

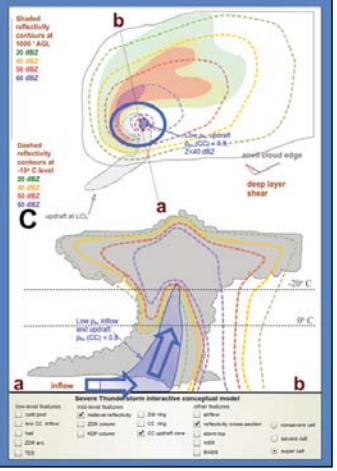
- CC ring (CC < 0.95)
- Mixed precipitation phase region around exterior of updraft just above environmental freezing level
- Ephemeral



Above the environmental freezing level and along the outermost perimeter of the updraft just outside the ZDR ring lies a ring of reduced CC values from 0.9 to 0.95. This ring forms as frozen particles from the main core interact with a raised melting layer of the updraft resulting in a region of mixed phased precipitation. Here, graupel, abundant liquid water, and growing hail are likely present. Because the perimeter of updraft where this occurs is very narrow, and the updraft edge changes quickly, this ring may not always be apparent on radar. This is especially true as the supercell undergoes mesocyclone occlusion processes when the updraft is partially disrupted.

Mid-level polarimetric signatures: The low CC updraft

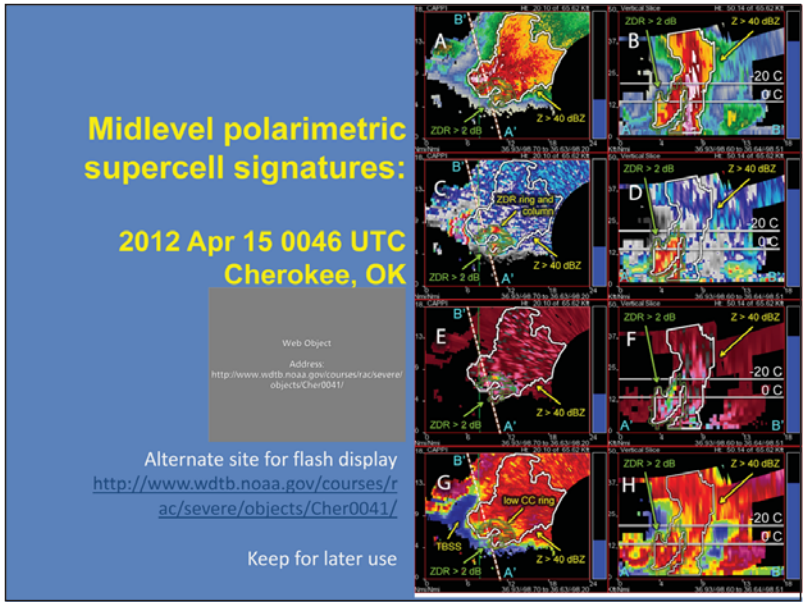
- Low CC updraft (CC< 0.8)
- Upward extension of low CC boundary layer air
- Insects and light vegetative debris
- extends as high as the BWER, and even more.
- Need clear air scatterers and no precipitation

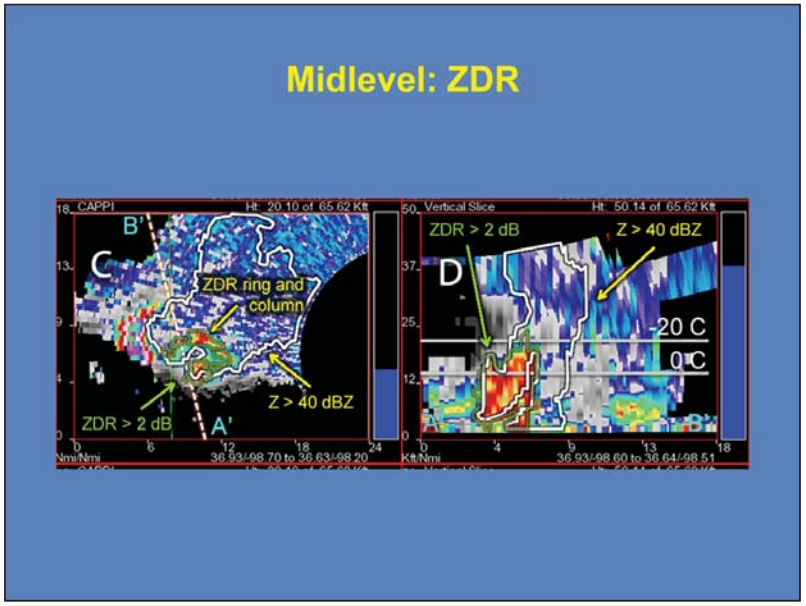


Sometimes, a core of very low CC air can be seen within the BWER at mid-levels.

This core appears to be an upward extension of the low CC boundary-layer inflow ahead of the storm. This signature is dependent on the low-level inflow exhibiting low CC.

This low CC inflow may be associated with non-meteorological scatterers, such as light vegetative debris or insects. In some stronger storms, this signature can even be observed at altitudes above the BWER. If the inflow is relatively clear of insects and light vegetation debris, then this signature may not appear. A low CC updraft column might also fail to appear if precipitation from this storm or an adjacent storm is entrained by the updraft. If so, then the precipitation signal will dominate.

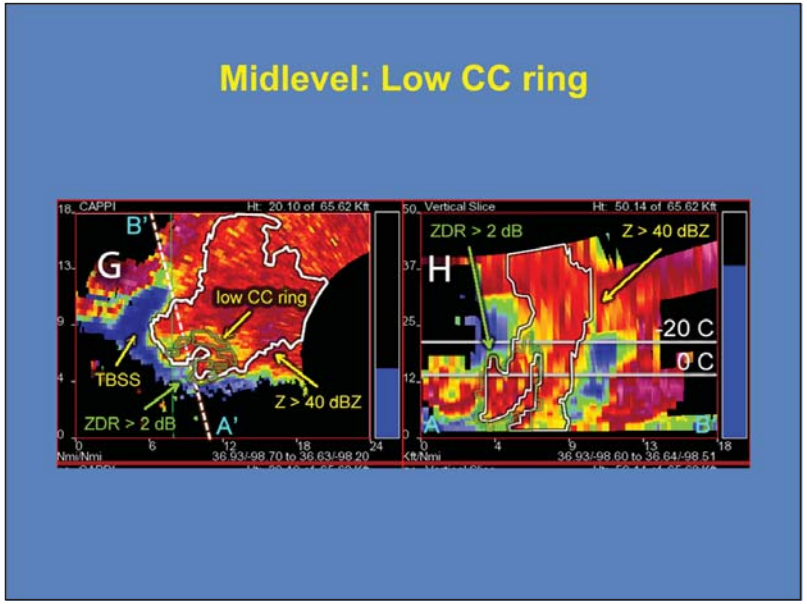




The example supercell thunderstorm, a cyclic tornado producer, was about to produce another tornado. Let's begin with a discussion of the features at the mid-levels and in the vertical cross-section.

An interactive flash display of this storm will appear with many CAPPI levels for this, and the previous time period. Keep it handy for the rest of the lesson.

Note that the vertical cross-section shows what appears to be a substantial BWER. In actuality, it was still a WER since there was an open end pointing into the cross section (Figure 7-97A, B). The reflectivity CAPPI was set at the -15 deg C level (20 kft ARL), high enough to isolate the ZDR column and nearly a ring (Figure 7-97C). Since the WER didn't really close off into a BWER, the ZDR ring also had an open end.



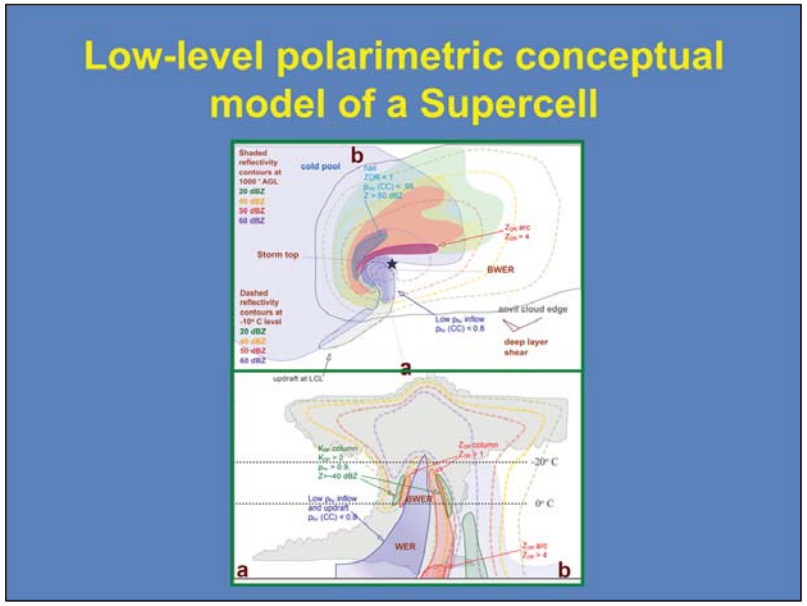
Likewise, the low CC region or arc in Figure 7-97G could also have been a ring had the high reflectivities closed off a BWER at -15 deg C. However, the same mechanism applies, and so, for convenience, we identify the low CC arc with the same name as the conceptual low CC ring.

One dual-pol signature that is notably absent is the low CC updraft signature. While the storm inflow CC was notably low (near the A' in Figure 7-97H), there was a substantial collapse of the high reflectivity and high ZDR echo overhang that resulted in some light precipitation entraining into the updraft, weakening the low CC signature within its core.

In Figure 7-97G, there is a field of very low CCs west (down-radial) of the precipitation core. This feature could be identified as the low CC updraft air since it appears to be in a notch surrounded by high CC echoes, but this is a Three-Body Scatter Spike (TBSS). Be skeptical of any adjacent low CC echo masquerading as a low CC updraft signature if it is:

1. Down-radial of intense reflectivities (> 60 dBZ) and, therefore, could be a TBSS, and

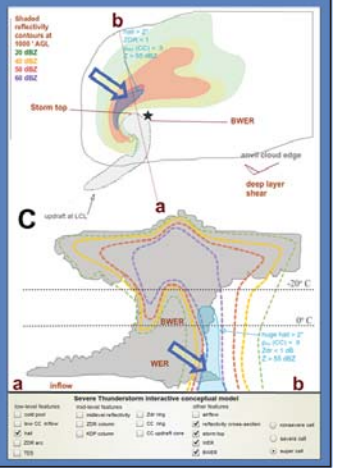
2. Contains a low signal-to-noise ratio (high spectrum width) on the edge of a reflectivity area.



At low-levels, the updraft signatures are replaced with others that are commonly associated with supercells. Low-levels are defined as typically from the ground to 3 km AGL for most surface-based supercells.

Low-level polarimetric signatures: significantly large hail

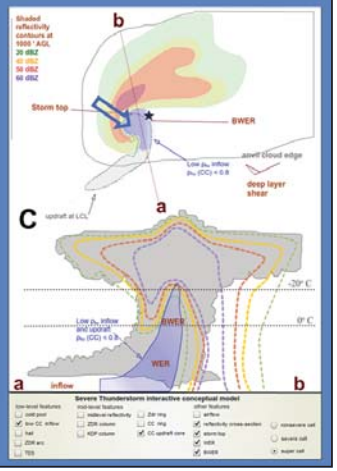
- Significant large hail (diameter ≥ 2")
  - ZDR < 1 dB
  - CC < 0.9
  - Z > 55 dBZ
    - (preferably > 60 dBZ)



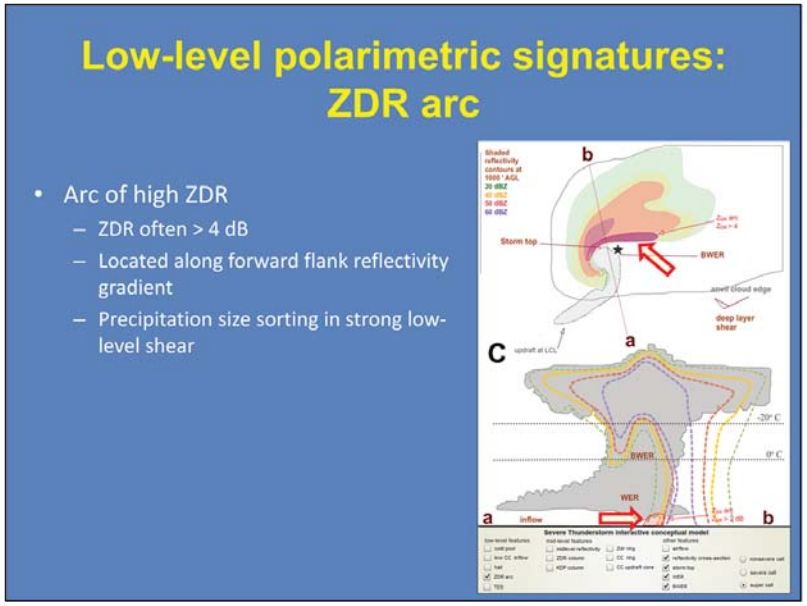
Supercells produce the vast majority of very large hail (diameter > 2"). Thus, the dual-pol signatures associated with this size hail commonly appear within the low-level cores of thunderstorms. These include CC roughly less than 0.9, ZDR less than 1 dB and preferentially near zero, and reflectivity greater than 60 dBZ (Figure 7-98). There is a severe hail identification lesson that explores the dual-pol hail signatures in greater detail.

Low-level polarimetric signatures: Low CC inflow

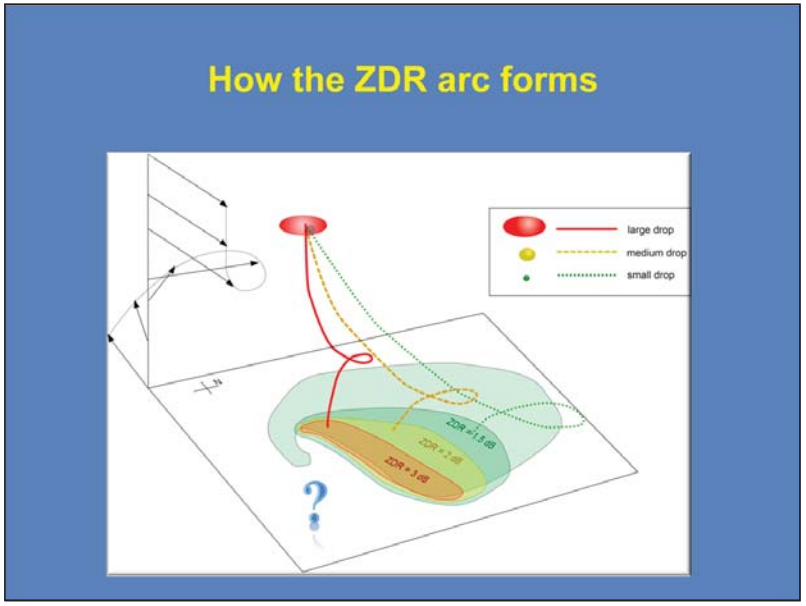
- Inflow accelerates into supercell updraft base
  - In addition to insects
  - May loft light vegetative debris



A little more about the low CC low-level inflow needs to be said. The lower extension of the low CC updraft core starts in the inflow layer. Usually this low CC inflow is an extension of a low CC precipitation-free boundary layer that is full of insects. However supercells have a tendency to accelerate air into the base of the updraft. If the inflow gets strong enough, then light vegetative debris also gets lofted and the CC may actually decrease in close proximity to the supercell. That's why the low CC inflow area in Figure 7-98 is shaded darker as the air flows into the base of the updraft.



One of the most intriguing signatures is the ZDR arc. This is a region of high ZDR precipitation echoes that lie along the sharp low-level reflectivity gradient facing the storm-relative inflow. Some of these hydrometeors are from the sloping echo overhang and others are from the edge of the precipitation cascade region. Recent research has theorized that the ZDR arc originates as the precipitation falling from aloft, is sorted by the vertical wind shear present in the environment, and enhanced along the forward flank outflow.

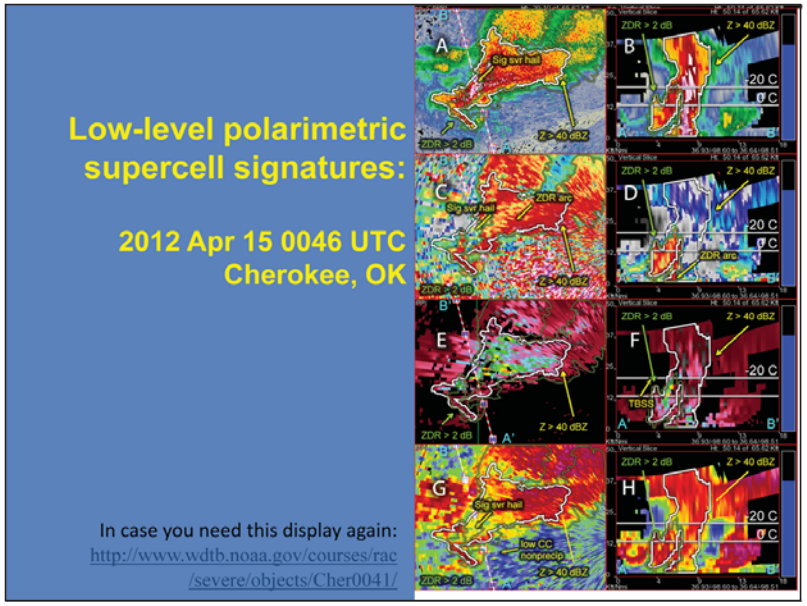


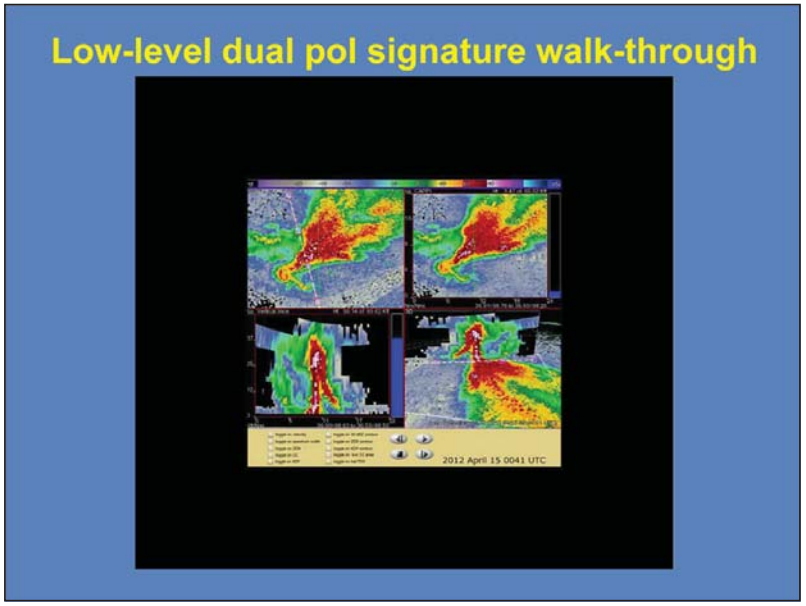
Imagine the wind profile with height changing in magnitude and direction as the vectors in Figure 7- 99 show. Then release a large, medium, and small size drops from the same position above the edge of the forward flank precipitation curtain. The larger droplets would respond less rapidly to the changing winds as they descend and therefore would fall closest to the precipitation-free low-level inflow. The smallest drops would respond most quickly to the changing winds and be carried away into the main core unless they evaporate first.

Because the size sorting continues to the ground, this feature is shallow, often below 6 kft above the ground. The fact is that strong vertical wind shear is required to produce the ZDR arc; therefore, this feature appears most commonly in supercells. In fact, Kumjian and Ryzhkov (2009) have suggested that the magnitude of the ZDR arc increases as the low-level storm-relative helicity increases.

It is enticing to think that monitoring the strength of the ZDR arc would give forecasters an assessment on the strength of the storm-relative helicity feeding the storm updraft, and that it could be used as a tornado precursor signature. However, there is no solid evidence yet and more research is needed on the ability of this signature to help in anticipating tornadogenesis.

In the meantime, the best way to detect this feature is to choose a radar close enough so that the lowest 6 kft can be sampled. Then choose either the lowest scan reflectivity image or drop the CAPPI into the lowest, clutter-free elevation possible.





Let's look at an example of the low-level dual-pol signatures in supercells by going back to the Cherokee, OK storm shown in the previous example.

An interactive flash display of FSI output will appear in a separate browser window. Use the right/left step buttons to go up in elevation in the CAPPI. As you continue to step to the right, you will go from the 0041 UTC scan to the 0046 UTC scan whereby continuing to step to the right will increase the CAPPI elevation.

Use the interactive flash display to follow along this video tour.

Okay in this video you may want to have your own display available and have adjacent to this.

Stop at any time

We have reflectivity in 4 panels. Again, it's the same display with PPI in upper left, CAPPI in upper right, then the cross-section lined up in the PPI.

We have a classic supercell with all the accoutrements of what you would expect in reflectivity including concave reflectivity hook e cho. Very high reflectivities exist just northwest of the concave gradient where values exceed 60 dBZ. Then in the cross-section you can see above the concave gradient we have a BWER.

So the highest reflectivities exceeding 60 dBZ indicate large hail. Toggle to Zdr and we see low values in these areas.

Going to CC and you can see values fall below 0.9.

Go to Kdp where cc < 0.9 and they drop out.

But I think you see enough of low Zdrs, right here, low CCs, and high reflectivities indicate very large hail close to the ground.

Let's look at ahead of the storm in CC and you see low values. This is

indicative of insects, chaff, dust in the strong inflow.

Take a look at velocity and indeed this is the area of Low CCs overlapping with strong inflow, which supports the idea of more dust and insects in the inflow.

Now going to Zdr, toggling on the 40 dBZ contour in white, and we'll take a look at this area of high values exceeding 4-5 dB here.

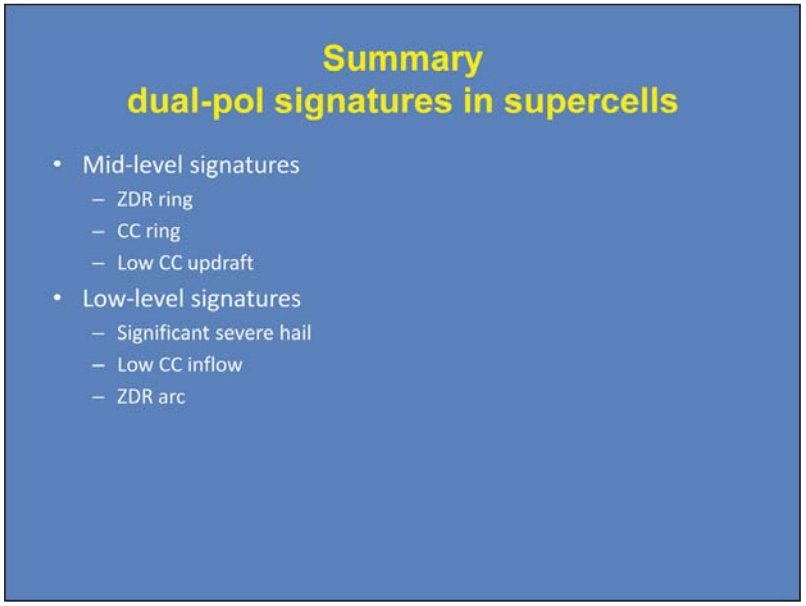
So this area represents the Zdr arc. If we look at the cross-section and we see those values and we'd expect to see a Zdr arc. This is the area we'd ex pect to see the Zdr arc due to precipitation size sorting.

The velocity here in the Zdr arc is where you can see the sorting as the storm accelerates the flow and the storm-relative flow helps in precipitation sorting.

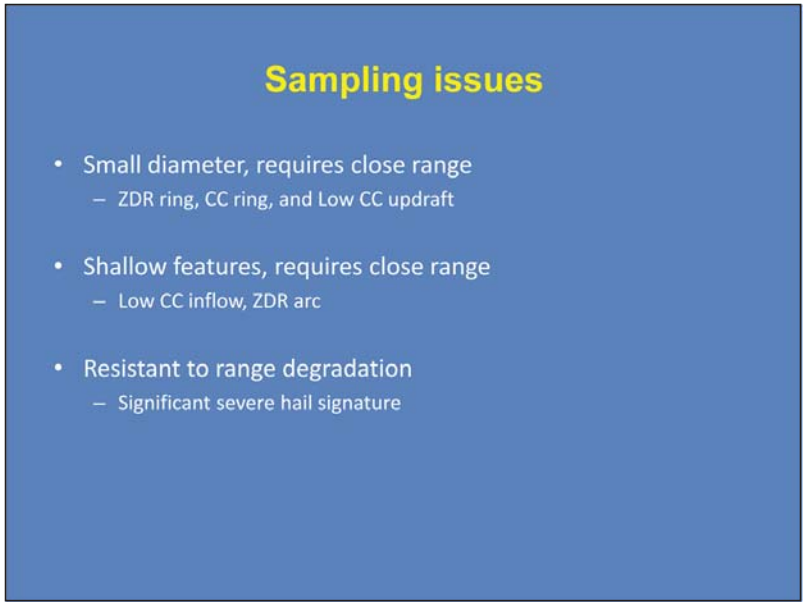
It's not just the low-level storm inflow but the environmental storm-relative flow too . You'd expect the storm-relative flow start from the southeast and then going up to the north at mid altitudes and then to the northeast at high altitudes.

All that indicates that small droplets are pushed to the north and large droplets fall quickly along the southern edge of the reflectivity forward flank.

You do note the high values of Zdr in the arc wraps around in the interior side of the hook.



We break down the dual-polarization signatures associated with supercells into two regions; the mid-levels between the freezing and -20 deg C level, and the low-levels focusing in the lowest 6 kft above ground.



Sampling some of these signatures is sometimes problematic given their small size or shallow nature. Here is a list of sampling issues with each signature:

• ZDR and CC rings are small, similar to the BWER in size. Radar needs to be close.

• The low CC updraft is small and the radar needs to be close.

• Significant severe hail signatures are large and relatively resistant to range degradation.

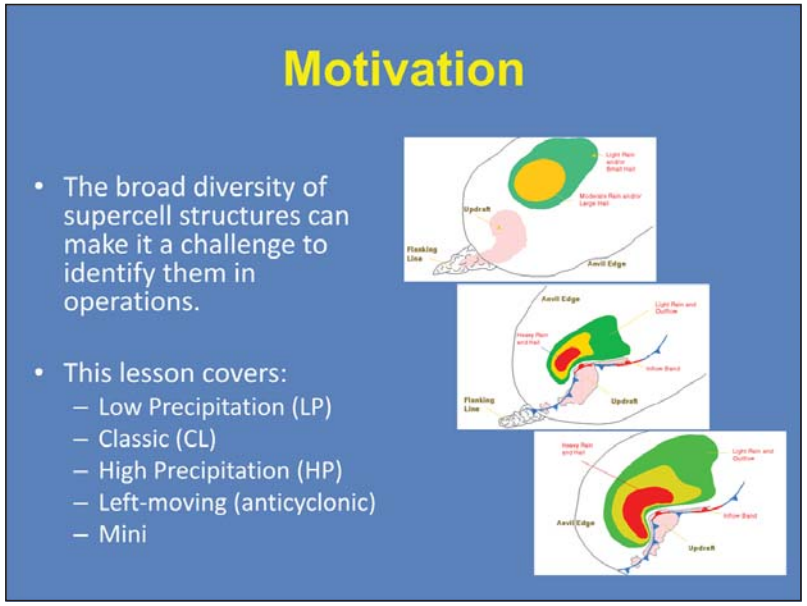
• The low CC inflow and ZDR arc are shallow and the radar needs to be close.



If you have passed the quiz, then you have successfully completed this lesson. If you have any questions, please contact us using any of the e-mail addresses listed on the bottom of the slide.



Welcome to the Radar & Applications Course, Convective Storm Structure and Evolution lesson on Supercell Archetypes.



Even though all supercells contain mesocyclones whose source of vertical vorticity is derived from vertical wind shear, the broad diversity of supercell structures can make it a challenge to identify them in operations. Thus, this lesson will describe some of the ways that supercells may appear through radar and other data.

This lesson will cover the characteristics of supercells with different precipitation distributions around their updraft which are: Low Precipitation (LP), Classic (CL), and High Precipitation (HP). We will also cover left-moving (anticyclonic) supercells and mini superecells, and how their characteristics influence warning decisions.

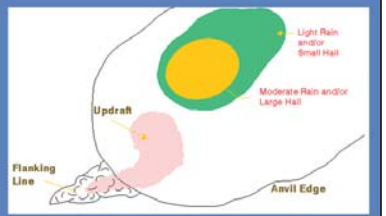
Learning Objective

- Describe the environmental, structural, and evolutionary differences that can produce low precipitation, high precipitation, classic, left-moving, and mini supercells.

This lesson has one objective: **Describe the environmental, structural and evolutionary differences that can produce low precipitation, high precipitation, classic, left moving and mini supercells.**

The Low Precipitation (LP) Supercell

- Dominated by updraft with little precipitation reaching the ground
- Low-level meso and hook echo are rare owing to the lack of an RFD
- Relatively dry boundary layer
  - Reduces available moisture
  - Adds to entrainment
- SR-anvil-layer winds > 58 kt
  - Little chance of hydrometeor recycling

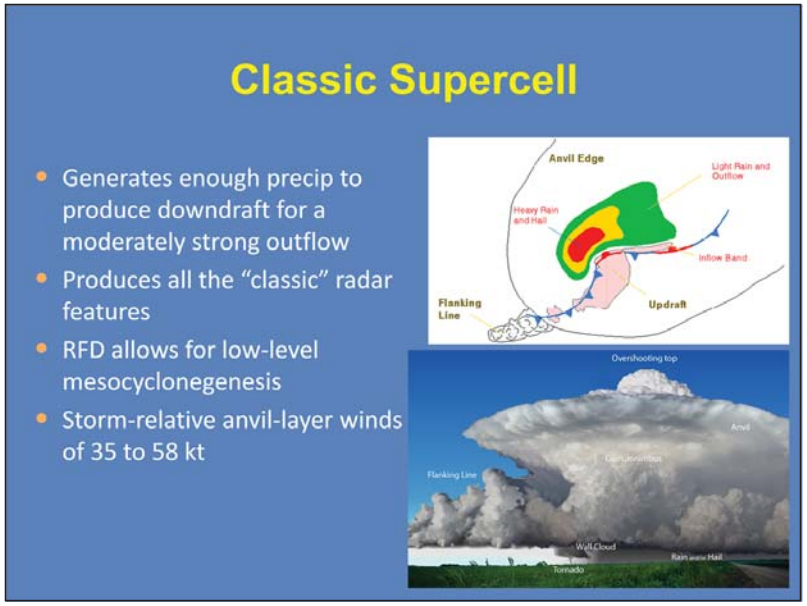




Supercells are grouped into three different structural classes depending on the amount of precipitation contained within the core, and where the mesocyclone is located with respect to the main core.

Low Precipitation (LP) supercells are generally dominated by updraft with little precipitation reaching the ground. These storms are visualized by exposed "skeletal" updrafts and translucent to nearly transparent precipitation cores. The relative lack of precipitation leads to weak downdraft formation and thus these storms could be said to be outflow deficient. LP supercell updrafts often show significantly strong midlevel mesocyclones. However, low-level mesocyclones are rare owing to the lack of a well defined Rear-Flank Downdraft (RFD). There is rarely a hook echo, and most of the precipitation is carried well downstream of the updraft by the storm-relative upper-level winds. Maximum reflectivities in LP supercells are relatively weak with the maximum values likely produced by a few, large, dry hailstones. True LP supercells represent the dry extreme of the supercell spectrum and are quite rare.

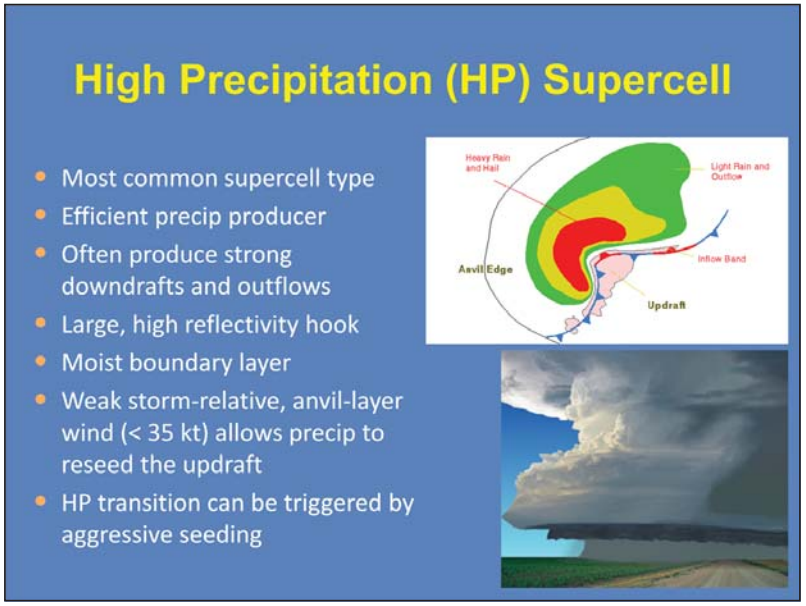
LP supercells require significant instability and shear, but other conditions help to reduce precipitation efficiency. Relatively dry boundary air reduces available moisture and adds to entrainment, but LP storms can also exist where boundary layer moisture is high. **Additionally, very high storm-relative anvil-layer winds (>30 m/s or >58 kts) transport rising hydrometeors well away from the updraft before they descend out of the anvil layer (Rasmussen and Straka, 1998).** Hydrometeors may have little chance of recycling back into the updraft, especially if the midlevels are dry.



Classic (CL) supercells generate enough precipitation to be able to produce enough downdraft for a moderately strong outflow. These storms are associated with all the classic radar features of a supercell including an inflow notch, WER, BWER, hook, and mesocyclone. The RFD is stronger with a classic supercell than with an LP supercell and therefore, low-level mesocyclogenesis is more likely. The result is a greater threat of severe weather from winds and tornadoes.

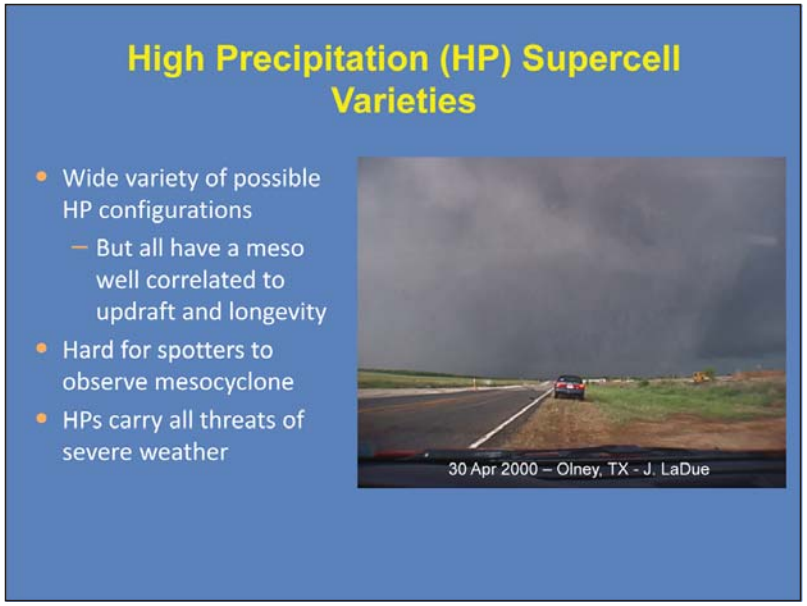
Classic supercells occur in moister environments than are typical for LP supercells. Storm-relative, anvil-layer winds are likely to be lower for classic supercells (mainly between 18-30 m/s or 35-58 kts).

These supercells produce the majority of long-lived tornadoes. They are also the common cyclic tornado producer.



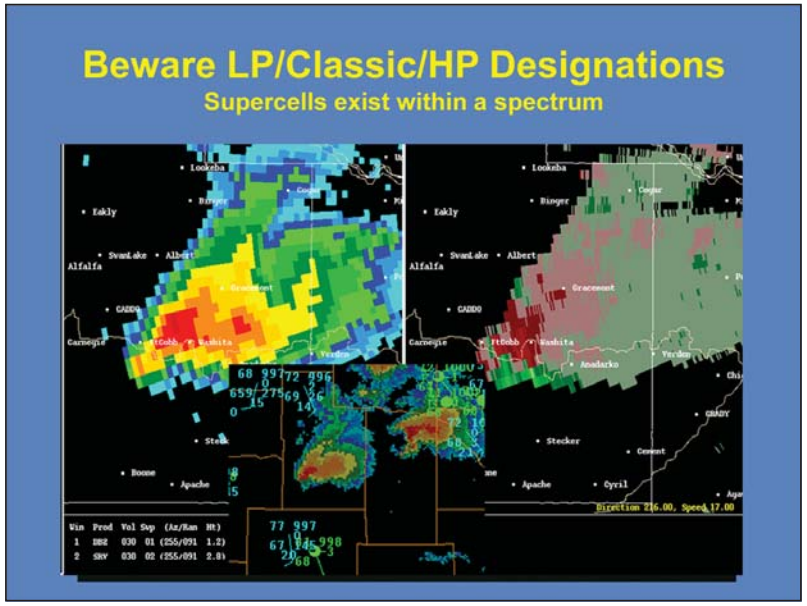
High Precipitation (HP) supercells are the most common of all supercells. They are highly efficient precipitation producers and often produce strong downdrafts and outflows. Large amounts of precipitation are available to wrap around the mesocyclone, producing a large, high reflectivity hook. Occasionally, the RFD gust front associated with the hook is intense enough to generate strong convection along its leading edge. The result is that the strongest core can be behind and to the right of the mesocyclone path. Occasionally, this process leads to HP supercells evolving into bow echoes.

HP environments typically show more boundary layer moisture than that of LP or even CL. However, high boundary layer moisture is not necessary for an HP. **Another possibility includes low anvil-level, storm-relative flow (<18 m/s or <35 kt) to allow precipitation to reseed the updraft improving precipitation efficiency.** A supercell can transition to HP if it is being seeded by aggressive cells on its flanking line or adjacent storms.



There is a wide variety of possible HP supercell configurations, however, they all share traits common to all supercells – a mesocyclone well correlated with an updraft and longevity. The mesocyclone is usually well sampled by radar owing to the high reflectivities in the hook. However, spotters in the field often have a difficult time observing the mesocyclone area most favorable for tornadogenesis.

HP supercells carry all threats of severe weather including: Large hail, damaging winds, tornadoes, and flash flooding.



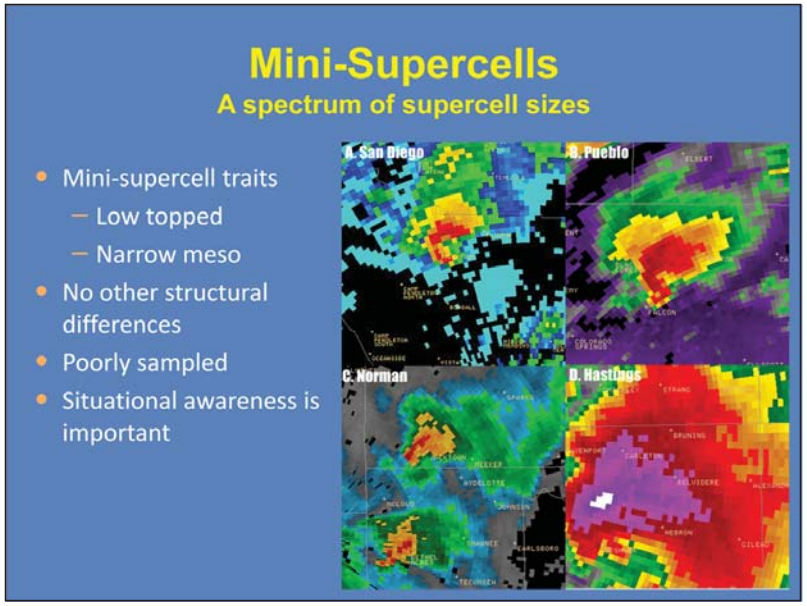
Having said all this, be cautious about spending lots of time trying to classify supercell type. There are no formal definitions, and many search papers refer to these supercell archetypes using different criteria. Supercells exist within a spectrum with no well-defined boundaries.

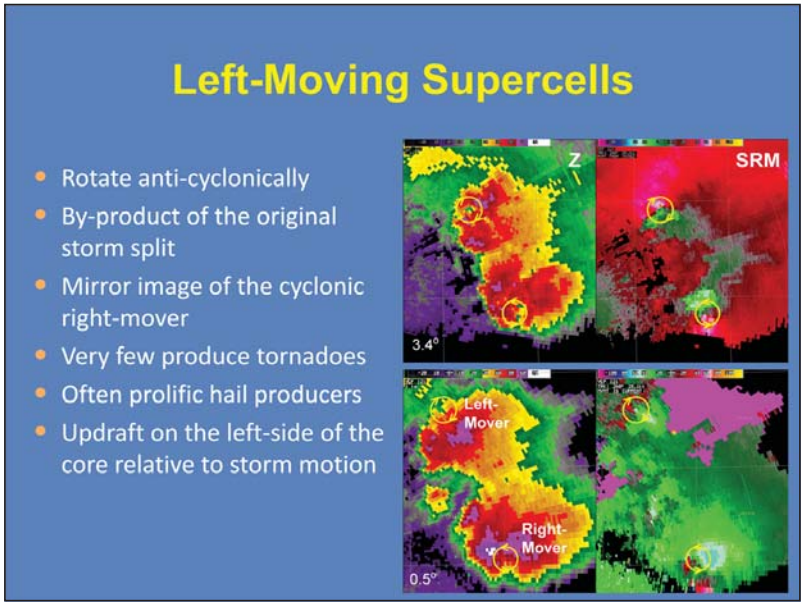
For example, the first designation of a Low Precipitation supercell was put forth in the 1970's (Davies Jones et al., 1976) as a name for storm they documented by radar with unusually low reflectivity. It was still a supercell and likely contained an active rotating updraft. Since their paper, the LP supercell has been photographically documented out in the field by others (Bluestein and Parks, 1983). Many storm spotters and storm chasers now label LP storms based on visual properties of a nearly transparent precipitation core and a fully exposed updraft tower. An example of what some spotters have designated as an LP storm is shown here.

The nearly transparent precipitation core in visual light can be deceiving. Much of the precipitation may be composed of large hail. In addition, there are precipitation shafts behind the lowered wall cloud under the right side of the updraft that may fall unnoticed by spotters.

The WSR-88D from 60 nm away showed that indeed the nearly transparent precipitation echoes were highly reflective. However, the radar was too far away to detect the hook echo. This supercell produced twenty tornadoes during its lifecycle, including ones shortly before and after these image were taken.

To summarize, when considering the potential hazards of a supercell, be careful not to base it too heavily on supercell classification. The storm in produced a tornado shortly after the image was taken.



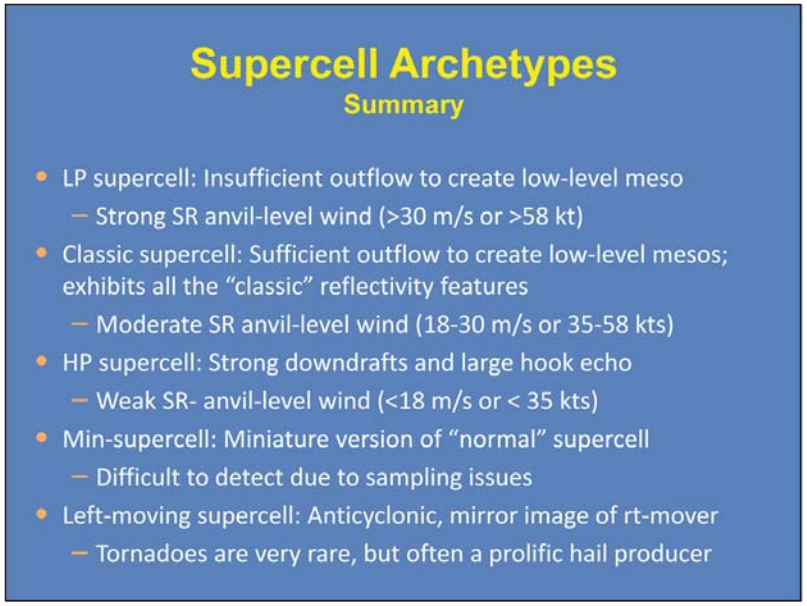


As supercells can vary in the amount of precipitation falling around their mesocyclone, they can also vary in size. There can be low-topped supercells with wide mesocyclones, or high-topped supercells with narrow mesocyclones. **A mini-supercell is one that is both low topped and contains a narrow mesocyclone.**

There are no structural differences between mini- and normal-sized supercells. However, there are differences in the expected severe weather. Giant hail (>2.5" in diameter) is rare because of limited extent of the updraft into the hail growth zone, and smaller horizontal dimensions of the updraft. **Poor radar sampling of small mesocyclones means that it is more difficult to measure high rotational velocities.** To illustrate this point, the supercells in panels A-C of this image were tornadic even though their associated mesocyclones were very small and/or apparently weak. Therefore, it is important to identify mini-supercells and be aware that their apparently weak circulations (V_r < 30 kts) can still carry a significant tornado risk as document by Prentice and Grant in 1996.

Left-moving supercells rotate anticyclonically (in the northern hemisphere) and are a by-product of the original storm split. They are structurally a mirror image of the cyclonic right-mover. Very few left-moving supercells produce tornadoes, and for reasons that are poorly understood, often produce long swaths of giant hail. As long as the hodograph is relatively straight, the left-mover can be as strong as the right-mover.

Left-moving supercells have their updraft of the left side of the reflectivity core relative to storm motion. Take a look at the left-mover in the 0.5 degree Base Reflectivity product. Notice the enhanced reflectivity gradient and concavity is located on the north side. The 0.5 degree SRM product shows anticyclonic shear associated with this updraft. Higher up, at 3.4 degrees, both the left- and right-movers have BWERs and rotation.



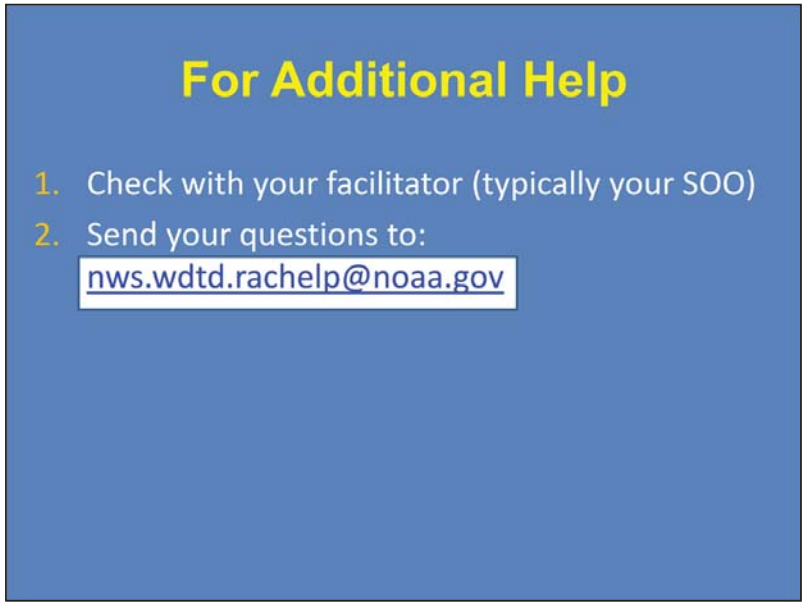
LP supercells exist with no real definition, and yet there is a consensus that LP's are unable to form a hook echo and also produce insufficient outflow to create low-level mesocyclones. They typically exist in dry boundary layers and/or strong anvil-level storm relative flow, however just as often, it is the way LPs initiate that provide a clue as to their existence.

Classic supercells exhibit a small hook relative to forward flank core which is accompanied by sufficient outflow to create low-level mesocyclone. They appear to form most often with moderate anvil-level SR flow (18–30 m/s or 35-58 kts) but not always.

HP supercells exhibit a large hook, sometimes with most of core following the mesocyclone. Intense RFD outflows often accompany HPs. They typically form with weaker anvil-level SR flow (<18 m/s or <35 kts), or perhaps through multistorm seeding.

Mini-supercells are structured very similarly to their larger counterparts, however, you suffer the disadvantage of not being able to detect their features as readily unless the storm is close to your radar.

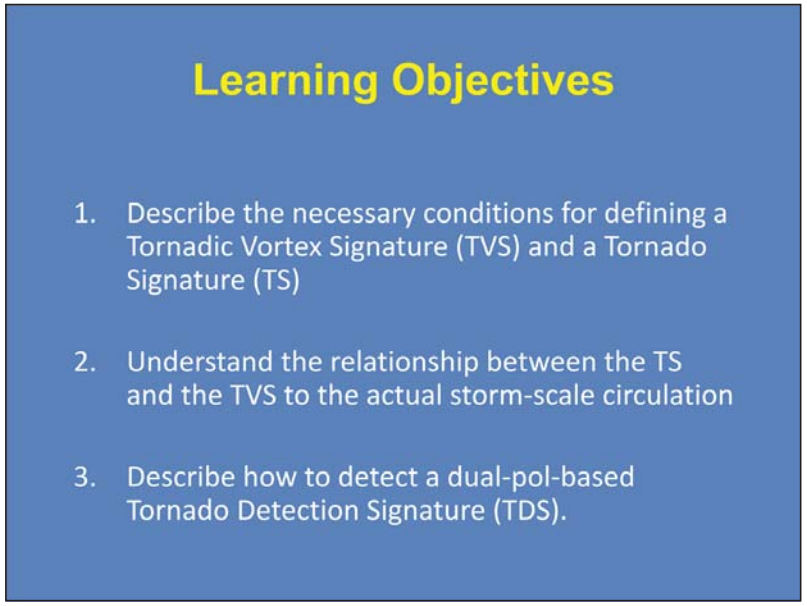
Left-moving, or anticyclonically-rotating supercells are structurally a mirror to their right-moving, or cyclonically-rotating counterparts. They rarely produce tornadoes, however, they are often prolific producers of hail. While they are rapid movers in most occasions in the Northern Hemisphere, some environments allow for the left-movers to be the slow movers. Their mesoanticyclones are currently undetectable by the MDA.



For additional help, check with your facilitator (typically your SOO) or send your questions to the listserv e-mail address here.



Welcome to the lesson entitled “Analyzing Tornadic Scale Signatures”. This lesson should last 30 minutes.

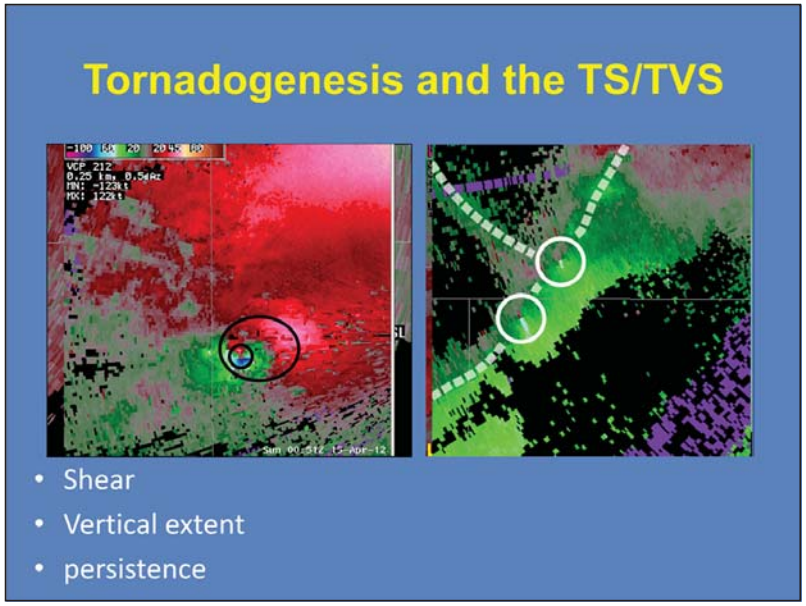


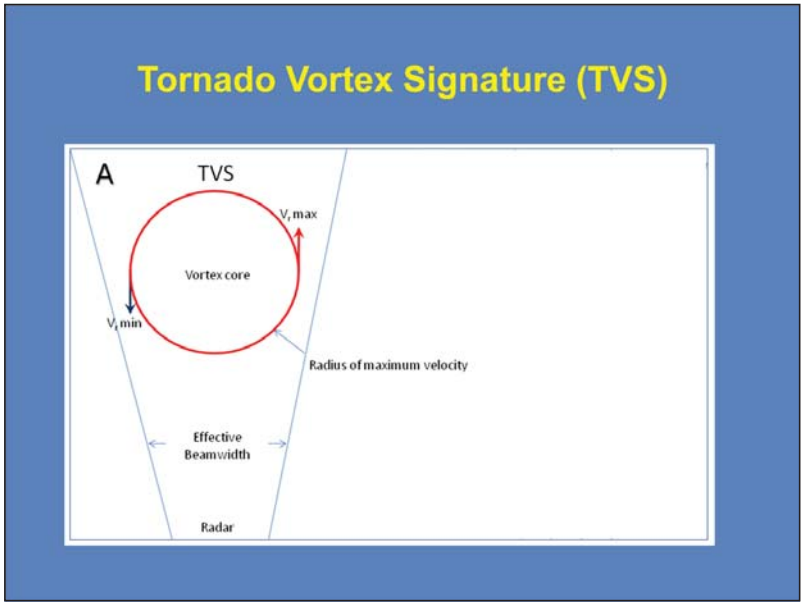
A WSR-88D with dual-polarization offers forecasters a broadened capability to infer ongoing or imminent tornadoes. Velocity data can detect a signature of a vortex that may be associated with a tornado. The signature may either be a Tornadic Vortex Signature (TVS) or a Tornado Signature (TS). Dual-pol data can detect lofted debris from columnar vortices connected to the ground. This detection is called a Tornado Debris Signature [TDS; Ryzhkov et. al (2005)].

This lesson is in two parts. The first describes how to identify a TVS *vs.* TS and how to assess the potential to identify an actual tornado. We also discuss the types of true circulations that may manifest themselves as a TS and TVS. The second addresses the TDS and how to identify it.

Objectives

- Describe the necessary conditions for defining a Tornadic Vortex Signature (TVS) and a Tornado Signature (TS).
- Understand the relationship between the TS and TVS to the actual storm-scale circulation.
- Describe how to detect a dual-pol-based Tornado Detection Signature (TDS).



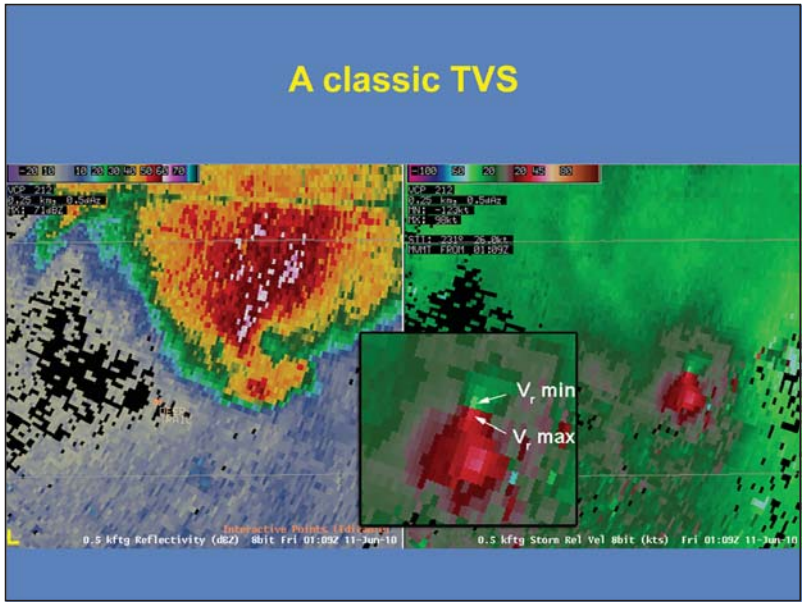


The tornadogenesis process traditionally manifests itself on radar as an increase in rotational velocity in the mid- and/or low-levels. The tightening of a region of circulation is common, which often leads to the development of a Tornado Signature (TS) and Tornadic Vortex Signature (TVS) by radar. On other occasions they appear to spring up out of nowhere. In reality, the range degradation of radar data prevents you from detecting the increased shear that exists just before the tightening process. Regardless of how it appears, the onset of a TS/TVS should be associated with the phasing of a strong updraft and increased low-level circulation. For the purposes of this lesson, we refer to the TS/TVS as one defined by the operator, not an algorithm.

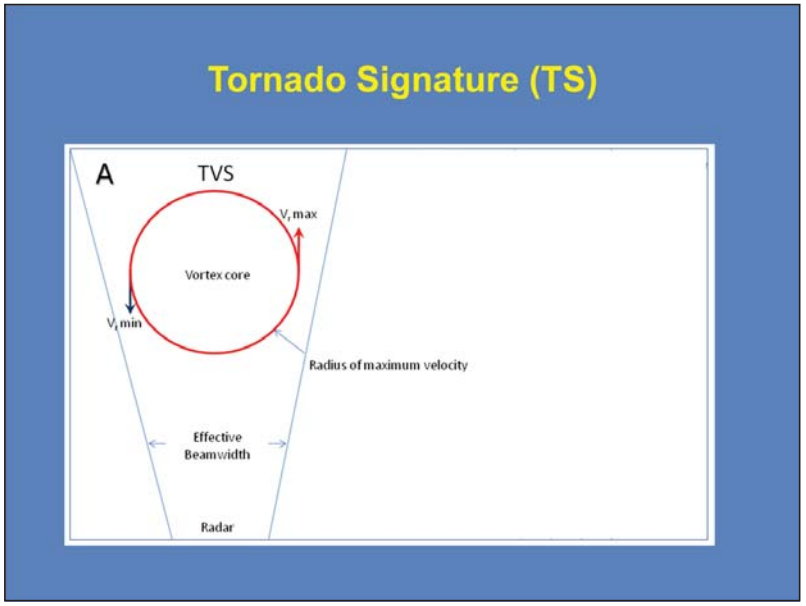
The type of circulations that satisfy this category are possibly associated with tornadic rotation that **meets or exceeds established criteria for shear, vertical extent and persistence. A TVS/TS can be described as a tornadic velocity profile superimposed on a larger mesocyclone. However, a larger parent circulation is not required and sometimes the TVS/TS is the mesocyclone. Let's elaborate more specifically on the TVS/TS.**

A TVS occurs when the core diameter of the tornado-scale circulation is smaller than the effective beam width of the radar. A TVS shows up as a signature where the radar detected maximum (Vr max) and minimum radial velocity (Vr min) are located on adjacent azimuths. That is unless the entire vortex core lies within a single beam. In such a case, the **Vr max, Vr min** would nearly cancel each other out leaving a nearly zero radial velocity and a very broad spectrum width.

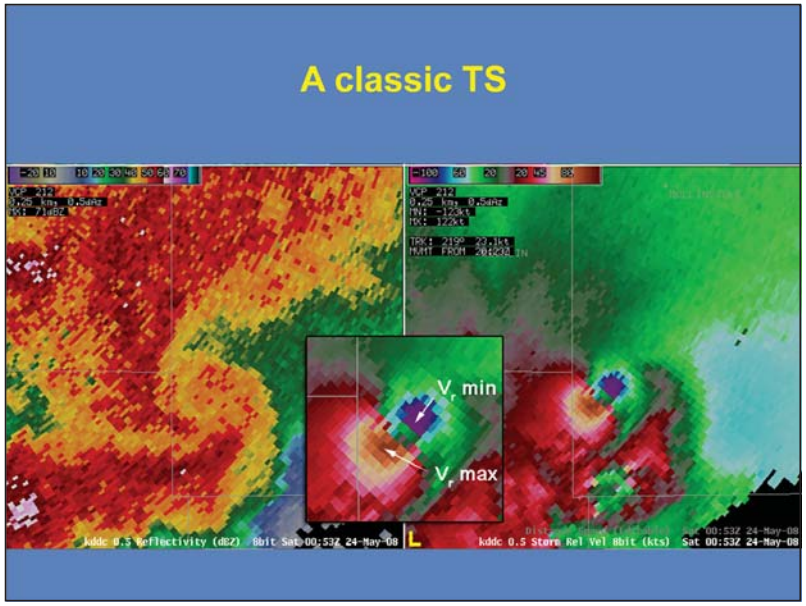
Also, when the azimuthal sampling interval is significantly less than the effective beam width, as is the case with super-resolution data, then there should be a transition zone between **Vr max, Vr min** as the beam sampling becomes less independent of one another. However, in reality, that's not frequent.



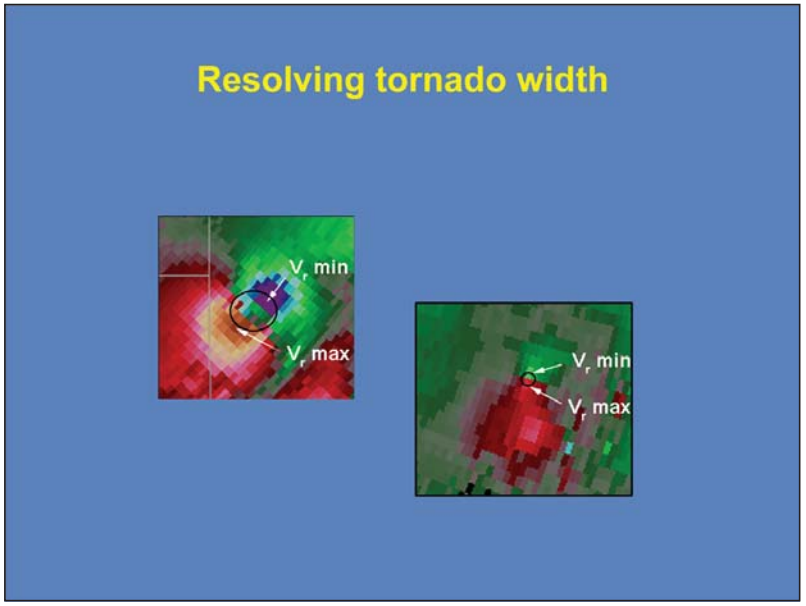
This image shows an example of a TVS sampled with the 0.5o elevation angle from the Deer Trail, CO tornadic supercell on 11 June 2010 at 0109 UTC. The maximum radial velocity was 35 kts (18 m/s). The TVS did not have a radar detectable Rankine combined vortex structure, because the effective beam width was too large to sample the inner core. Instead, the KFTG WSR-88D only detected the potential flow increasing in speed as the distance to the vortex center decreased until Vr max, and Vr min were found on adjacent azimuths. This is a classic TVS. However, there was also a separate V min without a corresponding V max that was probably associated with a rear flank downdraft.



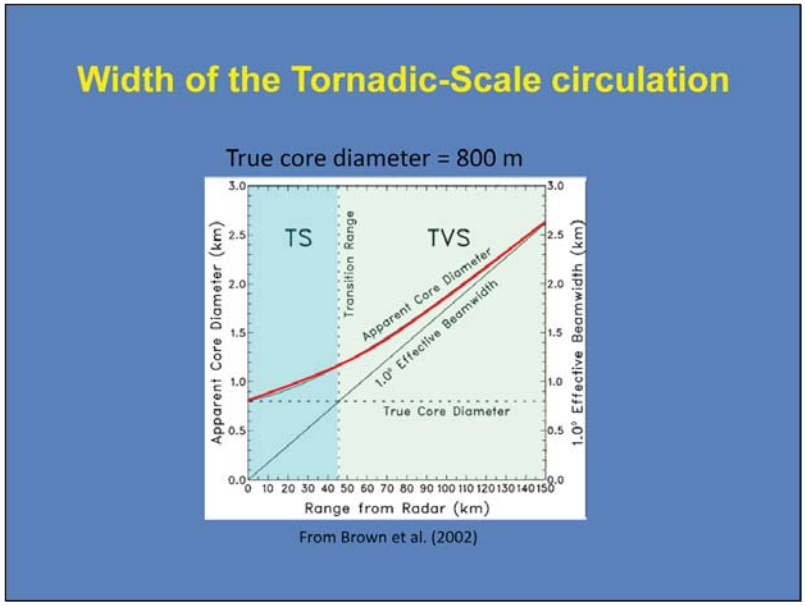
When the core diameter of a tornado-scale vortex is larger than the effective beamwidth, we call the vortex a TS. With more than one beam sampling the vortex core, Vr max and Vr min almost always appear separated by at least one radial.

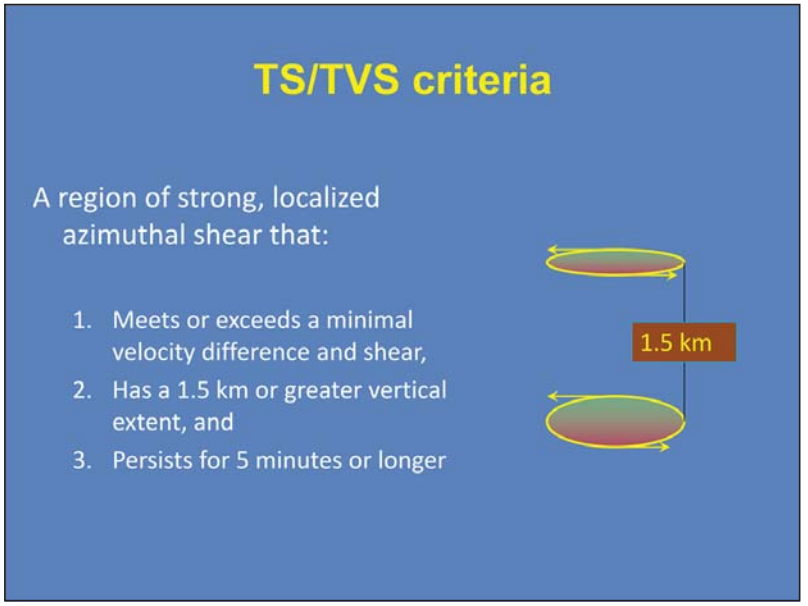


Here is an example of a TS where Vr max and Vr min are separated by four effective beam widths or 1.15 nm (1.85 km). The maximum inbound velocity was 124 kts (64 m/s). Given these velocity figures and the reports from the ground, this TS represented a high-end large tornado.



Since the radial velocity images show the same zoom magnitude, and each vortex was located 28 nm (51 km) in range, you may directly visualize the size differences between these two tornadoes. Even if they differ greatly in size and strength, both vortices exhibited an isolated velocity core leading up to a well defined Vr max, Vr min along with the maximum radial velocity gradient directed tangentially and counterclockwise. In other words, the detected velocity structure of both the TS and TVS was purely rotational.





Let me explain this graphic. Imagine a tornado whose true core diameter is 800 m and represented by the dotted line.

But note that as range increases from the radar represented by the X-axis, the 1.0 deg effective beamwidth increases in diameter.

Where the effective beamwidth is less than the true core diameter of the tornado, the radar should observe a TS. Beyond 45 km in range from radar, the effective beamwidth increases above the tornado diameter and this is where a TVS would be seen by radar. The transition range is that 'grey zone' between a TS and TVS. Again, this transition range is for this sized tornado only.

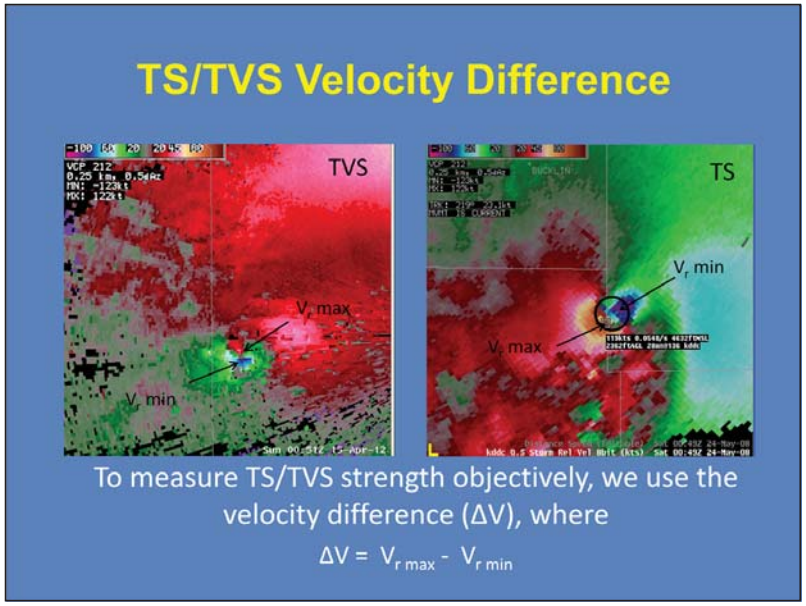
The existence of a TS does not guarantee that you are able to resolve the tornado width. The apparent core diameter of the TS immediately increases as range increases even if the true core diameter remains fixed. As this figure shows, there is no change in the how rapidly the diameter estimate increases with range even through the transition region from TS to TVS. Perhaps the apparent core diameter of the extremely large tornado would more closely match its true core diameter since its more than four radials wide, but even in this case, it is probably an overestimate.

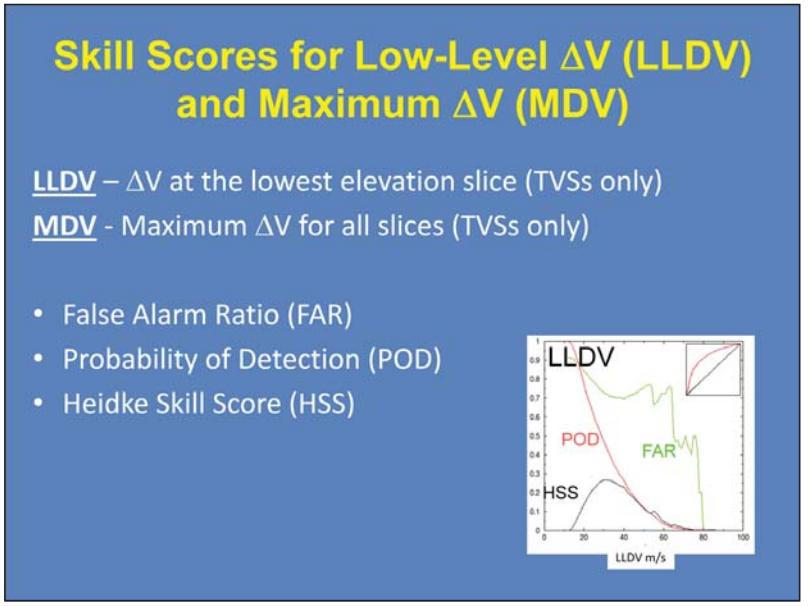
Some criteria are required to make sure that what circulation you are looking at is indeed a TVS. The three criteria are:

1. **A minimal shear:** There is no hard lower threshold in this criterion in the same way as an algorithm like the TDA would be assigned one. The minimal shear that an expert forecaster would define would depend on many things including the distance from the radar, the forecaster's assessment of the size of the vortex, near storm environment, and past experience. We will discuss what meaningful shears may be of significance.

2. **Vertical Extent:** At least some vertical continuity should be seen in a TS/TVS so that there is a high probability that an updraft is present in the circulation. For most events, the depth should be at least 1500 m (4900 ft.). Low topped supercells typically do not have deep TSs/TVSs, even if tornadic. Sometimes, and if detectable at all, only the lowest elevation angle contains a gate-to-gate rotational signature in tornadic low-topped supercells. Sometimes the vortex may appear as a TS and a TVS at adjacent elevation angles due to the vagaries of sampling and vortex structure. The vertical extent should include one or both manifestations as long as the true vortex appears to show vertical continuity.

3. **Persistence: In order to reduce the possibility of a circulation that randomly becomes vertically coordinated, you should ensure that the TS/TVS persists for at least five minutes**. However, mesocyclones can spin-up over a considerable depth in a very short time, and some legitimate TS/TVSs may become tornadic in less time. We suggest that if either signature forms in close proximity to a strong updraft signature, and a very supportive environment, persistence may not be a requirement to call it a TS/TVS and a tornado.





The maximum azimuthal shear found in a TS/TVS is perhaps the most consistent method for evaluating its strength; however, a forecaster needs a method that can be done quickly using base data. The fastest method is to simply take the radial velocity difference from where the maximum radial velocity is located as long as it represents the vortex core perimeter.

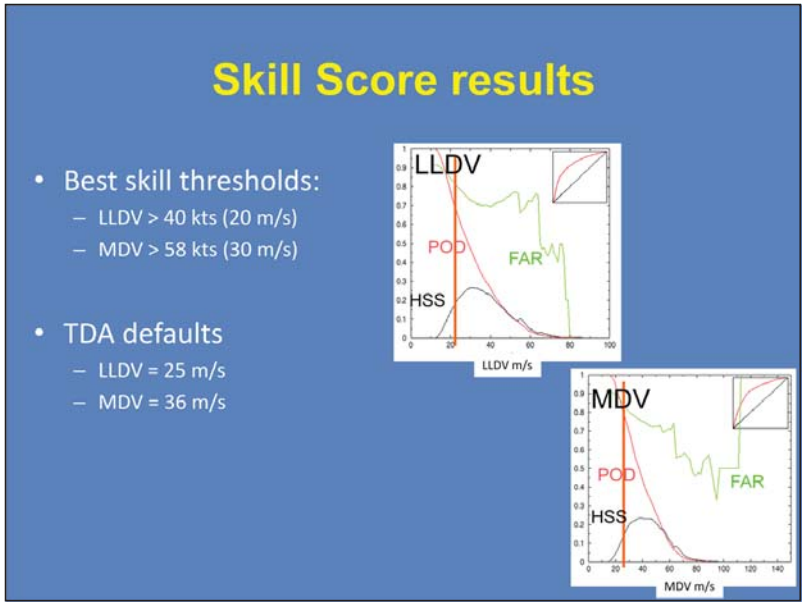
Therefore, the velocity difference (ΔV or DV) is shown as equal to the difference between the max and min Vr, where Vr max, and Vr min **are defined in the images**. Notice that we do not use inbound vs. outbound velocities because Vr min may still be the same sign as Vrmax. Because the distance between Vr max, and Vr min increases as the distance to the radar increases, DV is not really equivalent to shear. However, to simplify the process, we still use DV and account for how decreasing resolution can affect the relationship between DV and shear (more on this later).

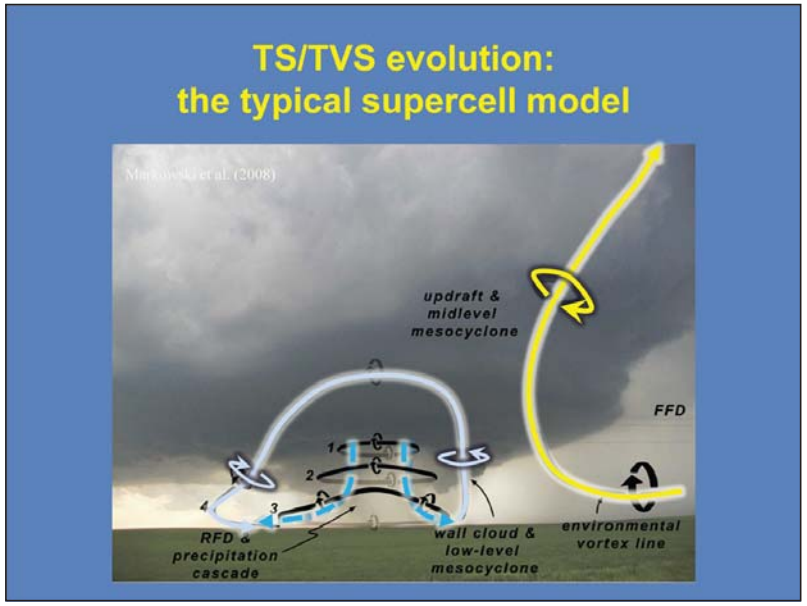
There are two DV calculations that are typically used:

- **The DV measured in the lowest slice is called the Low-Level Delta V, or LLDV**.
- **The maximum DV for all slices containing a TS/TVS is called Maximum Delta V (MDV)**.

To determine whether LLDV or MDV is large enough to satisfy part of the TS/TVS criteria depends on how useful it is to use these parameters in considering a tornado warning. Assuming a forecaster issues tornado warnings based solely on the presence of a TS/TVS, then the threshold LLDV and MDV are critically important to know. Unfortunately, there are many variables including storm type, environment, and distance to the radar, that impact and change these thresholds.

However, there is one way to provide some guidance to help comparing the likelihood that a certain LLDV and MDV is associated with a tornado. This guidance depends on incrementing the thresholds higher and higher and then look at how the False Alarm Ratio (FAR), Probability of Detection (POD), and Heidke Skill Score (HSS) change as the thresholds change using a large sample of TSs/TVSs of all storm types across the country. The HSS score compares FAR, POD, missed detections and correct nulls to show the best values for LLDV and MDV (TWG 2002). In other words, a forecaster's skill in issuing tornado warnings would peak when choosing the threshold values of LLDV and MDV where the HSS peaks.





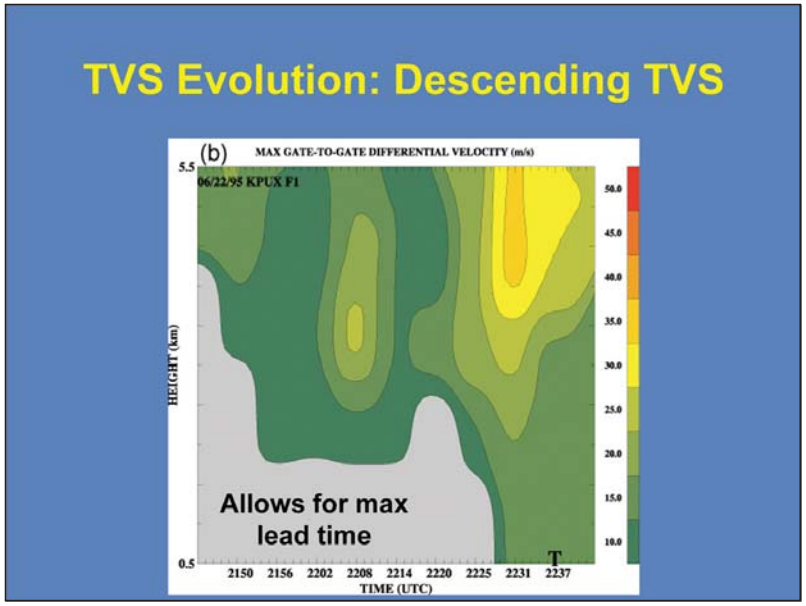
Results show significant skill score values are reached when LLDV exceed 20 m/s (40 kts) and MDV exceed 30 m/s (58 kts). The TDA default parameters are LLDV = 25 m/s and MDV = 36 m/s. As these values increase, the likelihood of a tornado also increases; however, a forecaster waiting for progressively higher values beyond where the HSS peaks before issuing a tornado warning suffers an increasing chance of missing a tornado.

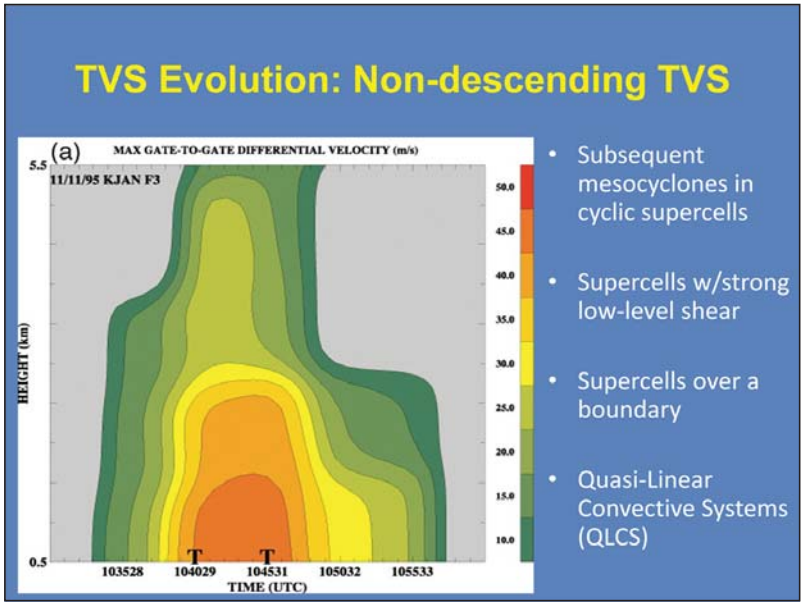
As a note, the data used in TWG 2002 was done with legacy resolution velocity data and using the TDA to collect only instances of TVSs. Current indications are that super-resolution velocity data involving TSs and TVSs offer similar skill in tornado discrimination, however the HSS peaks roughly 5 m/s higher. Given that super-resolution data is likely to detect higher peaks in velocity, this result is not surprising.

Traditional supercell mesocyclones often begin at midlevels as the updraft tilts environmental vorticity. As the updraft strengthens, the midlevel vortex may strengthen as well, possibly manifesting itself as a TS/TVS at those levels depending on the radar sampling.

At lower levels, the rear flank downdraft begins to generate horizontal vorticity around its exterior. As can be seen in this figure some of the vortex lines on the exterior of the RFD may get entrained into the main midlevel mesocyclone and updraft. As a result, a new low-level mesocyclone quickly develops at lower levels inside the wrapping RFD and under the updraft.

Since the low-level mesocyclone is feeding off of air of downdraft origins, it is often referred to as an occluded mesocyclone where the term occluded means the pre-storm air is no longer entraining directly into its base.



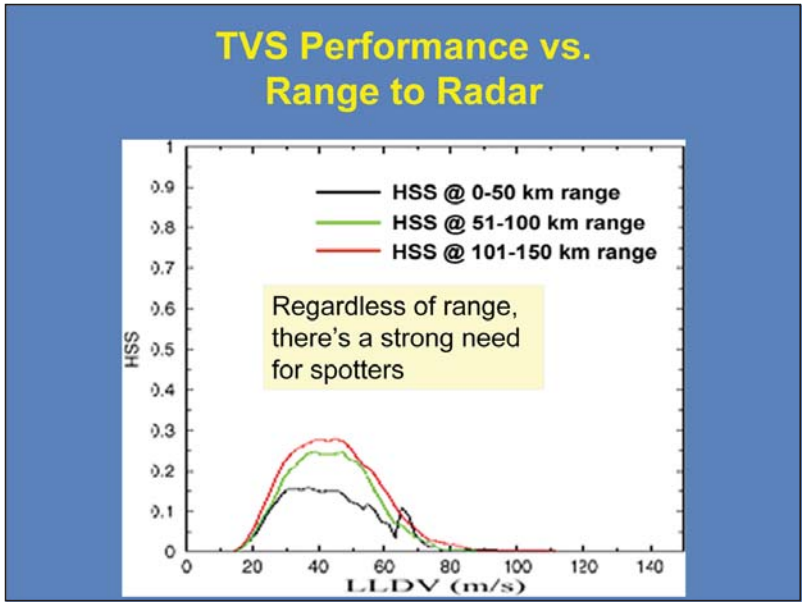


Following the mechanism of tornado formation in the last graphic, the WSR-88D may indicate that the mid-level TS/TVS is descending as the low-level mesocyclone strengthens. The process may or may not continue to intensify into a tornado, however to the warning forecaster, it may appear that the TS/TVS originates in the mid-levels and then descends to the ground over time. This process allows for the maximum lead time in a tornado warning.

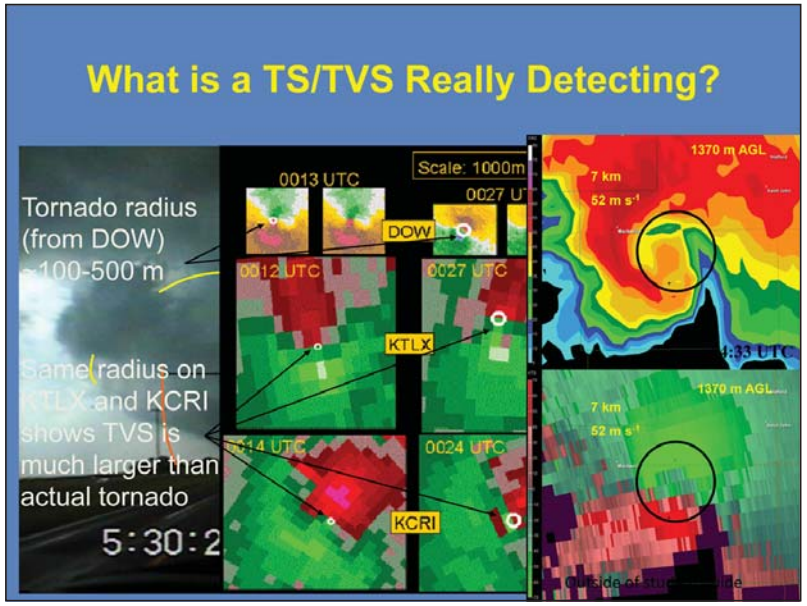
And you can see that in this graphic as the envelopes of MDV descend with time until a tornado was reported at 2237 UTC.

Non-descending TVS

Studies (Trapp et al., 1999; Wakimoto and Atkins, 1996) have indicated that not all mesocyclone induced TSs/TVSs descend from mid-levels to reach the ground. **About half originate at low-levels and then extend upward**. Often, this non-descending paradigm is associated with subsequent mesocyclones in cyclic supercells, or in supercells with very strong lowlevel shear, possibly from an outflow or other type of boundary. Non-descending TSs/TVSs occasionally originate within supercells above a boundary containing strong vertical vorticity (Wakimoto and Atkins, 1996). This is a critical observation, since low-level tornadogenesis can occur in moments. Warning lead time depends on monitoring the trend of the low-level TS/TVS shear, picking the right thresholds, and anticipating rapid tornadogenesis. Non-descending TSs/TVSs will be discussed further in the section on multicell squall lines.



TS/TVS detections are limited in range owing to degraded radar sampling with range. However, a comparison on the statistical performance of TSs/TVSs to detect tornadoes vs. range to radar indicates that there is little range degradation out to 150 km (~78 nm). These results show that other factors could be more important than radar range degradation - at least within the first 150 km. Therefore, there is a strong need for spotters regardless of range to the nearest radar.

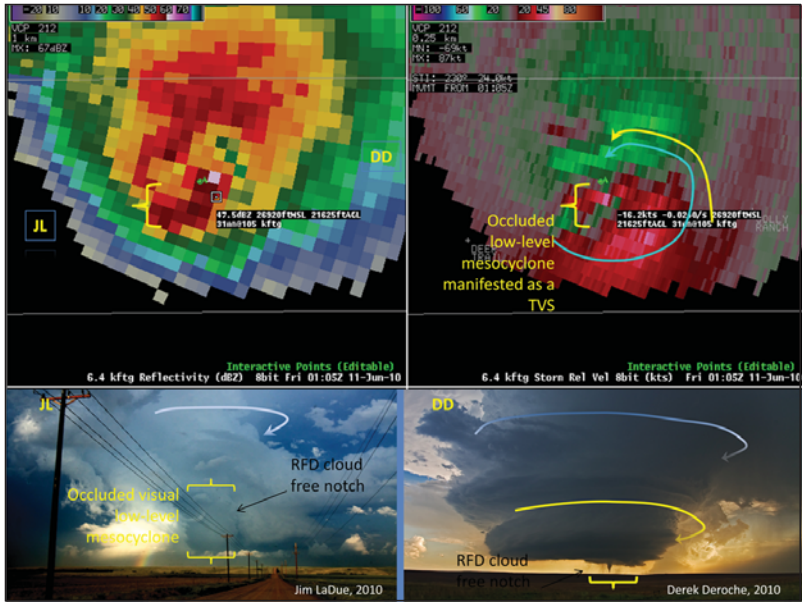


The TS/TVS in the image is showing a small circulation that is still much larger than a significant close-range tornado as depicted by the high resolution Doppler on Wheels (DOW) data. Note that the tornado widths represented by the white circles were sampled by the DOWs and then pasted upon the most recent scan from KTLX and then KCRI (a testbed radar). Now you can see how much smaller the actual tornado widths were than the distance between the velocity peaks from the WSR-88Ds.

At low-levels, the TS/TVS most likely represents part of the intensifying mesocyclone inside the wrapping RFD. The RFD axis is usually closely aligned with the axis of the wrapping hook echo. The low-level flow inside the hook/RFD gradually accelerates with decreasing distance to the circulation center.

At far ranges, the TS/TVS may be more appropriately called a non-divergent mid-level mesocyclone. Since they are relatively common, nondivergent mid-level mesocyclones often appear as TVSs at far ranges and that is why we often limit TS/TVS to ranges less than a range of 80 nm (150 km).

In a few rare cases, the WSR-88D can resolve the tornado where the vortex core diameter is four or more effective beam widths wide. Such tornadoes are essentially low-level mesocyclones whose strength reaches tornadic values. These features manifest themselves as TSs and it is possible to have subvortices reveal themselves as TVSs within the larger TS. The Greensburg, KS tornado of 4 May 2007 exhibited TS characteristics from the KDDC radar with legacy resolution data, and within the TS, there were asymmetries in the velocity data that may have suggested such an occurrence (Lemon and Umscheid, 2008).



I like to discuss how the WSR-88D views the difference between the mid- and low-level mesocyclone. Here is an example taken from one of the Project VORTEX2 supercells near Deer Trail, CO on 10 June, 2010. Starting at low-levels, there is a classic TVS where Vr max and Vr min are located on adjacent gates, even with super-resolution data. This indicates that the tornadic vortex core was too small to be resolved. However, the potential flow outside the velocity peaks is part of the circulation outside the core and it is easily resolvable. The gate-to-gate LLDV for this TVS is about 55 kts at the lowest scan. Here we are looking at 0.9 deg in elevation

There are two pictures taken from either side of the supercell that will help to define the physical nature of this TVS. The first picture was taken west of the supercell where we can see many important features from the back side of the storm. The second picture was taken from the more traditional front side of the supercell. Overall the two images depict quite different scenes. The backside shot depicts very convective looking towers right from the cloud base to near the anvil with no obvious features that imply a rotating updraft. The front side shot, however, depicts an updraft that we normally associate with a supercell, a smooth circular updraft with circular banding. These views are typical for a supercell. More importantly is the prominent cloud-free notch extending up from cloud base to 1/3 the way up the convective tower as viewed from the backside. From the front, that dry slot appears left of the developing tornado. The low-level mesocyclone denoted by the potential flow in the velocity image is most likely that area including the dry slot, the tornado, and all the way to the north side of the updraft wall (inside the brackets). The inner core of the low-level mesocyclone, the tornado, is too small to be resolved so the velocity signature in radar is that of a TVS.

Going up in altitude to 11 kft AGL, we see the upward extending low-level mesocyclone (manifested as a TVS), and then strong inbound velocity to the north. This level would be above the cloud base as seen from front side visual image where the inbounds would visually appear as strong horizontal flow going left to right and then around the northern edge, and then to the back side of the updraft. The northern side of the updraft in the backside image is obscured by the bright updraft.

The upward extent of the low-level mesocyclone is visible even to 22 kft AGL. We also still see the inbounds on the northern side of the updraft but there is now a strong outbound to the south of the updraft. The blue streamlines depict approximately the altitude this flow field exists.

Interim summary: TVS/TS

- A TVS is a tornadic vortex with a diameter ≤ 1 effective beamwidth
- A TS is a tornadic vortex with a diameter ≥ 1 effective beamwidth
- A TS/TVS must have
  - Minimal shear estimated by Delta-V
  - Vertical extent ≥ 4900’ (1500 m)
  - Minimal persistence (~ 5 min)
- A TS/TVS typically represents the scale between the tornado and the low-level mesocyclone

A TVS is a tornadic vortex with a diameter ≤ 1 effective beamwidth

A TS is a tornadic vortex with a diameter ≥ 1 effective beamwidth

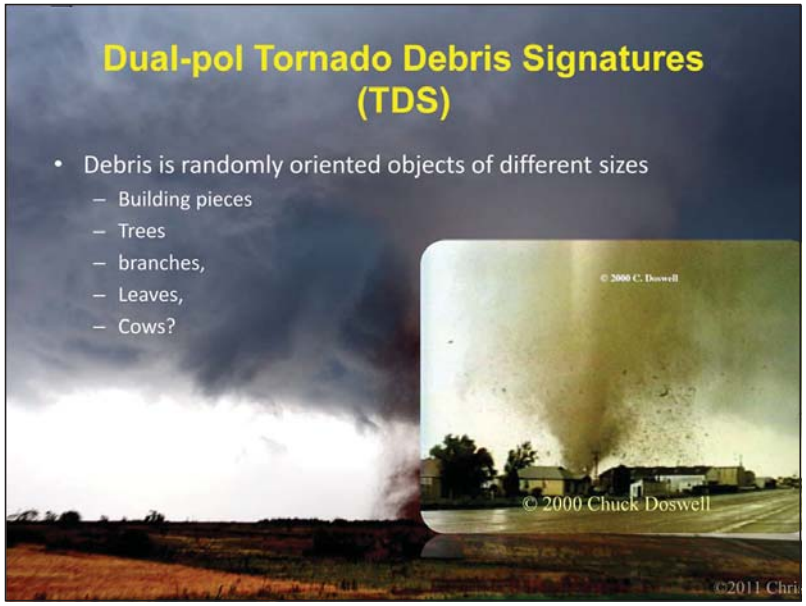
A TS/TVS must have

Minimal shear estimated by Delta-V

Vertical extent ≥ 4900’ (1500 m)

Minimal persistence (~ 5 min)

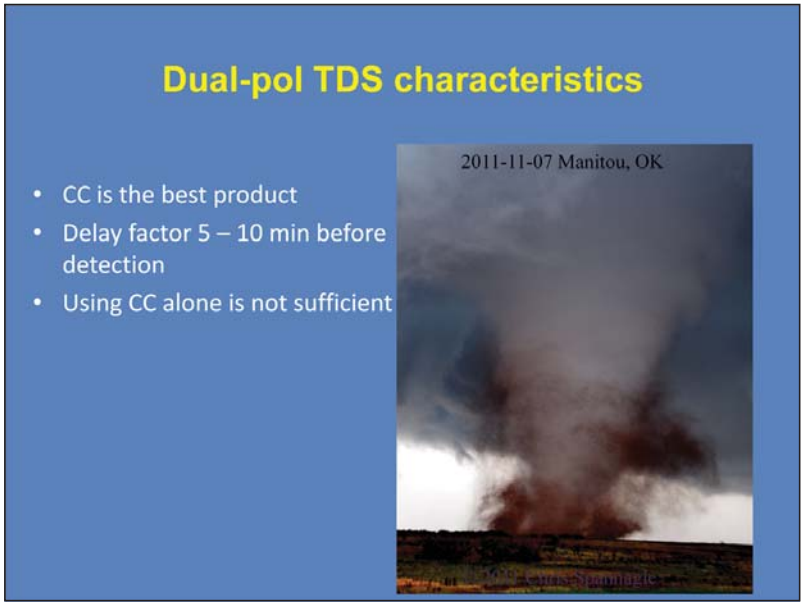
A TS/TVS typically represents the scale between the tornado and the low-level mesocyclone



A valid identification of a Tornado Debris Signature (TDS) helps a warning forecaster identify that a tornado is most likely occurring and is producing damage. A valid signature is likely to be considered as close to an actual tornado detection as a

spotter report. With that being said, the process of identifying a TDS must be done carefully to avoid an incorrect identification.

The radar is detecting tornado debris that is comprised of large, randomly oriented objects ranging from leaves to building fragments. This example shows debris as a tornado went through Pampa TX on 08 June 1995.



Because debris is randomly oriented, the dual-pol radar correlation coefficient product (CC) is by far the best product to discriminate debris from meteorological echoes. When analyzing a TDS, remember that debris was actually introduced to the circulation 5-10 minutes earlier. It takes time to loft and distribute the debris. And after the dissipation of the tornado, it takes time for the debris to settle out. However, identification of tornado debris with CC alone is not sufficient. Let's go through a method to make a good detection of a TDS.

Dual-pol TDS detection method

	product	strategy	Verdict?
			Yes or no
1	Velocity	Is azimuthal couplet present?	
2	CC	Are values < 0.8 in or near the velocity couplet?	
3	Reflectivity	Are values > 35 dBZ where CC < 0.8?	
4	ZDR	Values near zero help confirm above 3 criteria	
5	verdict		

1. First identify a storm-scale vortex such as a mesocyclone and/or a TVS (or TS) located in the vicinity of an updraft as per the instructions in identifying a mesocyclone, TS and TVS. There is no lower bound velocity threshold but the rotational couplet should be pronounced. In some cases, a vortex may be unresolvable and spectrum width may show a local and very high peak.

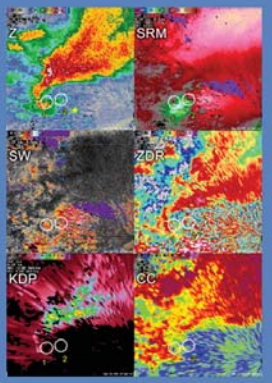
2. In the vicinity of a tornado vortex, look for a CC small minimum in CC. Typically, a value < 0.8 indicates a good potential for randomly oriented scatterers. This is not a hard threshold, however. Sometimes CC in a valid TDS may fall to 0.9 **only** when rain is mixed with debris. However, a TDS with a CC this high is rare.

3. If you have a localized CC minimum centered near a vortex, then check to see if the reflectivity is at least 35 dBZ. Lower reflectivities may result in untrustworthy CCs. In addition, the CC values may be the result of other non meteorological scatterers, such as insects or light suspended vegetation particles.

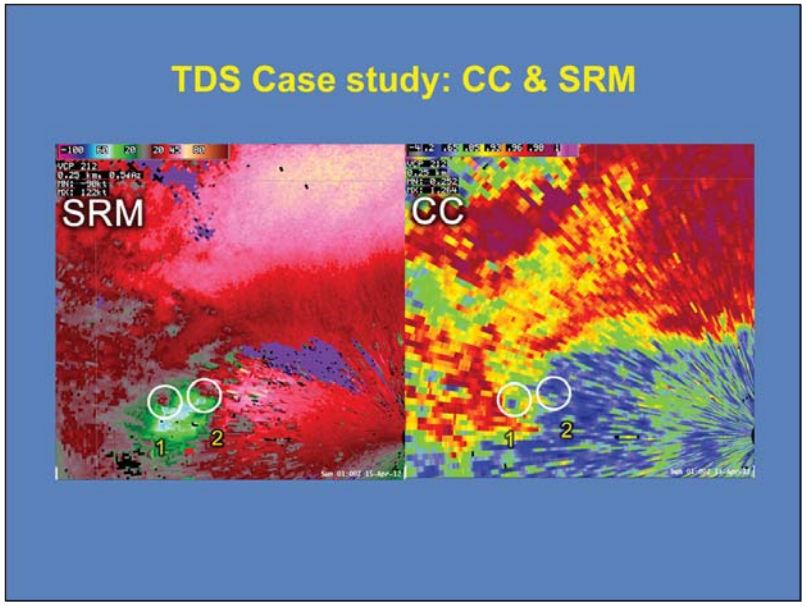
4. ZDR is typically near zero in valid tornado debris. However, the signature is not nearly as pronounced as CC. Nevertheless ZDR can be used as a confirmatory check.

Dual pol TDS case study

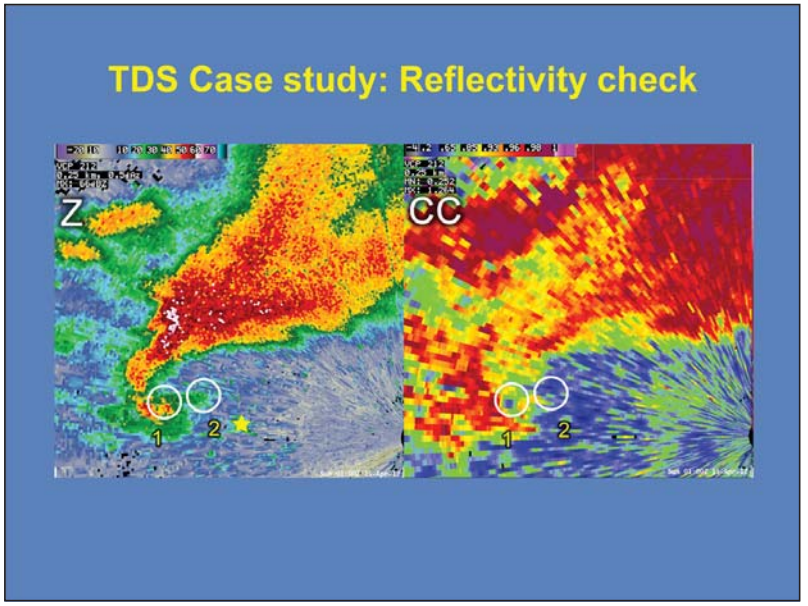
- 2012-April-15 0100 UTC
- Northwest of KVNX
- Two velocity couplets
  - Labeled 1 and 2
- Both in the same supercell



An example of two TDS candidates occurred with the 14 April 2012 Cherokee, OK supercell shown in this image. Follow along with the tornado identification methodology to find that two TVSs exist enclosed by Circles 1 and 2. Each vortex at the lowest scan has passed the vertical continuity check.



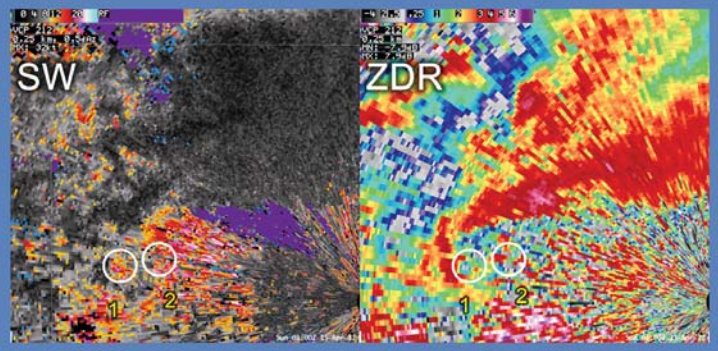
Both vortices exist along the edge of an RFD outflow and in the proper spot relative to the parent supercell. Going to the CC panel, note that there are low CC values within each circle. Circle 1 has a more pronounced CC minimum than Circle 2, while the low CC in Circle 2 is more in the low CC inflow.



The reflectivity panel shows that Circle 1 contains values above 20 dBZ and even some in excess of 40 dBZ. There may even be the suggestion of a debris ball, though not well defined. However, note that the lowest CC overlaps with 40 dBZ echoes in the southwest part of the circle. Circle 2 may have lower CC values but the reflectivity at the vortex center is just below 20 dBZ.

So, there is strong confidence that the low CC within Circle 1 is from tornado debris. However, confidence is low that the low CC values within Circle 2 are associated with any debris.

TDS Case study: ZDR and SW



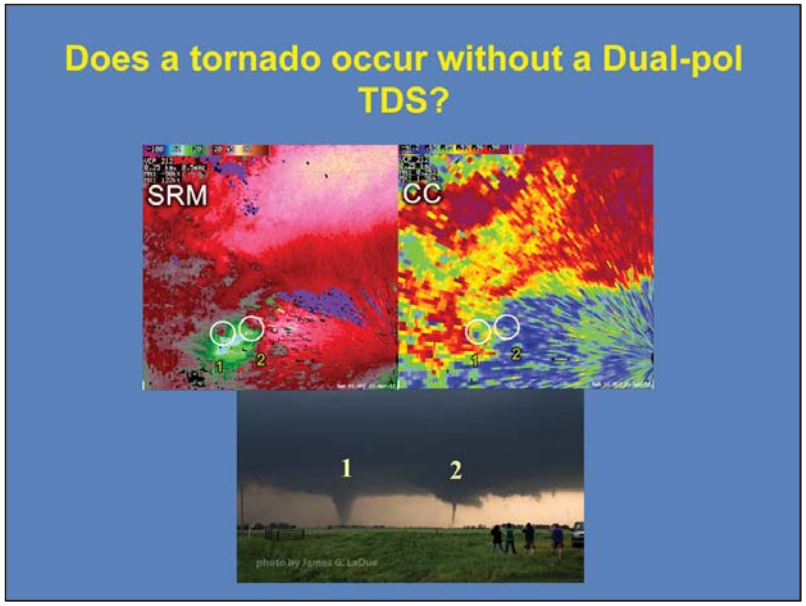
Within Circle 1 there is a well defined ZDR minimum at the same location as the CC minimum. This adds confidence that there is a TDS in Circle 1. The ZDR in Circle 2 is mottled with a mix of very low and high values, similar to the pre-storm air. Confidence remains low for a TDS in Circle 2.

Dual-pol TDS detection method

	product	strategy	Circ #1	Circ #2
			Yes/no	Yes/no
1	Velocity	Is azimuthal couplet present?		
2	CC	Are values < 0.8 in or near the velocity couplet?		
3	Reflectivity	Are values > 35 dBZ where CC < 0.8?		
4	ZDR	Values near zero help confirm above 3 criteria		
5	verdict			

Circle 1 contains a TDS. There is a TVS in the vicinity of a hook echo with a well defined CC minimum with sufficient reflectivity. Circle 2 shows no TDS signature. While the radar may be depicting debris, the signal cannot be separated out from the low CC non meteorological scatterers that exist around and within the inflow. This is the low CC inflow signature that is common in boundary layers with non-meteorological scatterers.

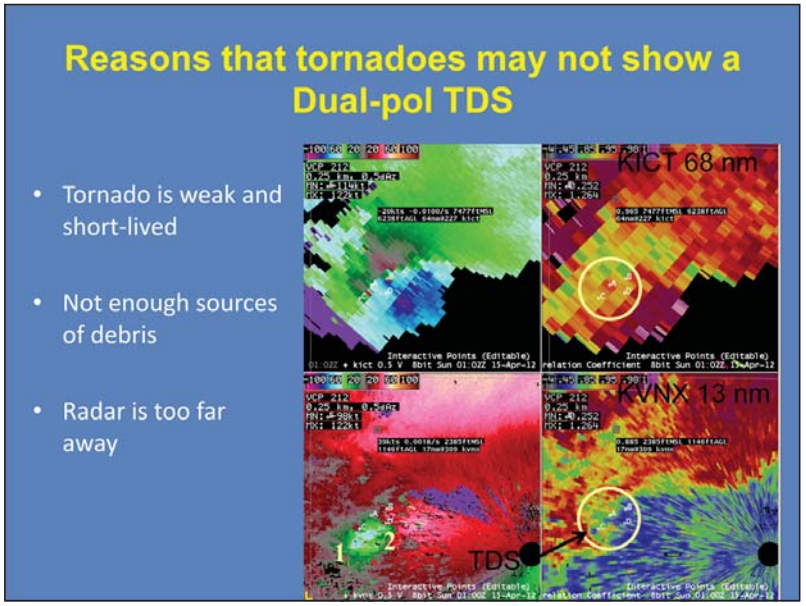
If there is a 'No' in any product in the first three rows (V/SRM, CC, Reflectivity) then you do not have confidence that you are seeing a Dual-pol TDS. The ZDR helps confirm the verdict but is not as strong of an influence.



Does this mean that only one tornado is in existence? In fact it doesn't. In the photo, two tornadoes are traveling across the landscape west of Cherokee, OK.

The tornado on the left corresponds to Circle 1 while the one on the right corresponds to Circle 2. At this time the left tornado is larger and has been in existence for longer. Both of those will result in more debris later. The one on the right has just formed from the new mesocyclone and has yet to loft enough debris to raise the reflectivity sufficiently from the KVNX radar to help discriminate insects from debris at this range and height.

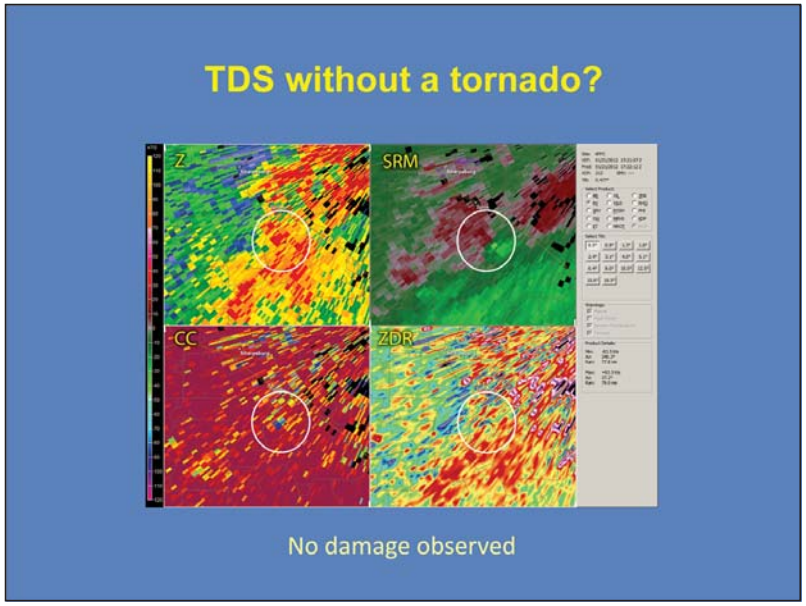
This example highlights one aspect of TDSs that is also common to many other radar signatures. That is the absence of a clear signature doesn't rule out the existence of the hazard for which the signature refers. But the presence of a TDS is as strong an indication of a tornado as a spotter report. The TDS should serve to raise confidence that tornado is or has been in progress. But warning issuance should never have to wait until a TDS occurs.



Three factors may lead to a tornado not exhibiting a Dual-pol TDS:

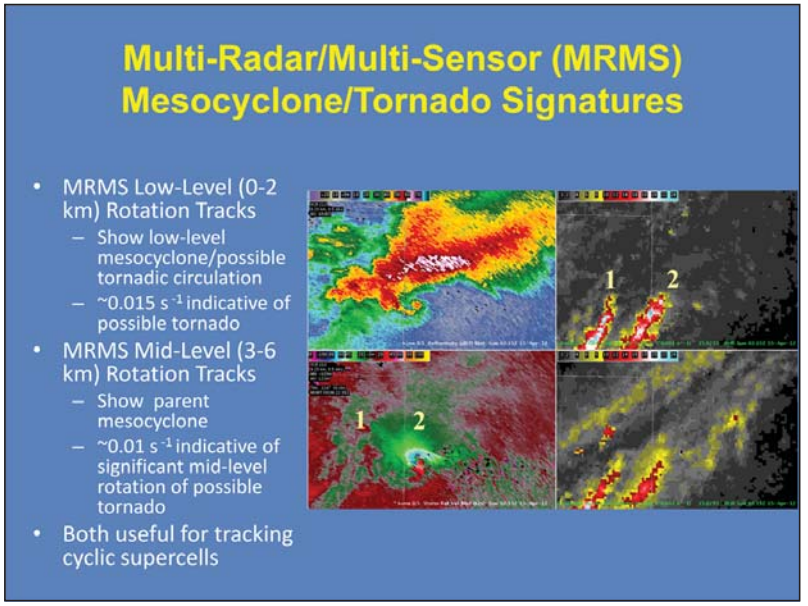
- The tornado is weak and short-lived. There is insufficient strength and time to loft detectable debris.
- There are not enough sources of debris. A tornado crossing an open dirt field is unlikely to generate as much detectable debris as a tornado of the same strength going through a town. Fine dust particles are too small for S-band radar detection in order to generate a TDS. A tornado must loft at least leaves, grass, and forest debris to generate a TDS.
- The radar is too far away. Range is everything. Weak tornadoes are unlikely to be detected more than 40 nm away. EF2 and greater tornadoes may generate a detectable TDS up to and possibly over 60 nm in range. The limiting factors are the height to which debris is lofted and the size of the debris footprint.

A great example of the distance problem comes up with the Cherokee storm when the same two tornadoes are viewed from KICT 68 nm away. The top two panels show the KICT radar where no CC < 0.8 was visible (top right panel) anywhere near the TDS that was visible from KVNX from 13 nm range. The letters represent the vortex locations from KICT (A and B), and KVNX (C and D, labeled 1 and 2).



Has a TDS been present without a tornado? The answer is possibly but very rarely. One such case occurred where a low CC bull's-eye was collocated with a rotational velocity couplet within high reflectivity in northern Georgia. NWS damage surveyors were unable to find significant tree damage. It is almost certain that a vortex was lofting light debris. The question is whether or not the vortex was strong enough to be defined as a tornado.

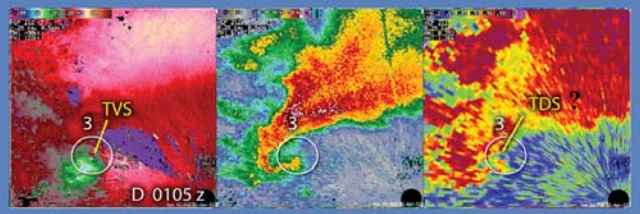
However, it is theorized that the main culprit in false detections seems to be when vortex signatures, low CC, and low ZDR values have been correlated within the weak reflectivity inflow notch ahead of a pendant, or hook echo. These associated vortex-like velocity signatures, in some instances, may have been side-lobe related, but in most instances, that is not the case. We simply do not know the origin of these “ghost-like” vortex signatures. However, the low values of CC and ZDR are easily explained in the updraft inflow notch of many supercells.



Another suite of radar products that can identify signatures near the tornadic-scale is Multi-Radar/Multi-Sensor (MRMS) Rotation Tracks products. Recall from the MRMS Products Course that these are time intervals of Azimuthal Shear and that they’re available at low- and mid-levels. The Low-Level Rotation Tracks can identify low-level rotation and possibly even tornadic circulations at its higher values. Little research has been conducted to test the utility of this product for tornado detection, but if you see values greater than .015 s^{-1}, you should check to see if your other tornado signatures correspond with the azimuthal shear magnitude.

Mid-level Rotation Tracks show the mid-level, parent mesocyclone. Values greater than .01 s^{-1} may indicate strong mid-level rotation. Additionally, both Tracks products make cyclic supercells evident. In the example on the right, our storm from the TDS example created another set of twin tornadoes during its lifecycle. The separate tracks are visible on both the low-and mid-level Rotation Tracks.

Tornado Lifecycle



- 0041 z - #2, convergent meso; #1 a TVS with TDS
- 0051 z - #2 a TVS with TDS; #1 dissipated, false TDS
- 0100 z - #3 convergent meso; #2 a TVS with TDS
- 0105 z - #3 a TVS with marginal TDS; #2 dissipated

http://www.wdtb.noaa.gov/courses/rac/severe/objects/CherMeso/

Now let's see how the tornado evolved over time for the same storm covered previously. In the 3 panel image a new rotational velocity couplet, mesocyclone (#2) begins with a strong convergent component as the old occluded rotational velocity couplet (#1) reveals itself as a TVS. The CC depicts a prominent TDS and in fact the reflectivity shows a donut hole at the tornado location.

Ten minutes later, the old TVS (#1) dissipated. What appears to be a continuing TDS is actually false. The reflectivity shows that the low CC bull's-eye is in very low reflectivity and there is no velocity couplet there. The new velocity couplet (#2) has become a TVS and is sporting a well defined TDS at the tip of the hook.

Ten minutes later, a new rotational velocity couplet (#3) has formed into a mesocyclone at the tip of the hook echo with a strong convergent couplet. Meanwhile rotational velocity couplet (#2) is still a TVS with an accompanying TDS.

Note that five minutes later, the rotational velocity couplet (#3) has become a TVS that is also accompanied by a TDS. This TDS is relatively marginal given the weak reflectivity but it has enough of a prominent CC minimum that debris could be the cause. Meanwhile, the rotational velocity couplet (#2) dissipated as it moved to the left of the track of the parent storm.

Summary

- A vortex is depicted as a
  - TVS when: effective beamwidth > vortex core diameter
  - TS when: effective beamwidth < vortex core diameter
- Three criteria satisfy a TS/TVS
  - Vertical continuity, time continuity, a localized peak in azimuthal shear with little convergent component, range within 150 km
- TS/TVS strength
  - $\Delta V = V_{r\,max} - V_{r\,min}$ whether or not the velocity peaks lie on adjacent azimuths.

Tornado-scale vortices manifest themselves as a nearly pure rotational velocity couplet with an isolated maxima in **Vr max, and Vr min**. If the Doppler radar's effective beam width is the same size or smaller than the vortex core diameter, the vortex manifests itself as a Tornado Signature (TS). If the effective beam width is larger than the vortex core diameter, then a Tornadic Vortex Signature is (TVS) is the result.

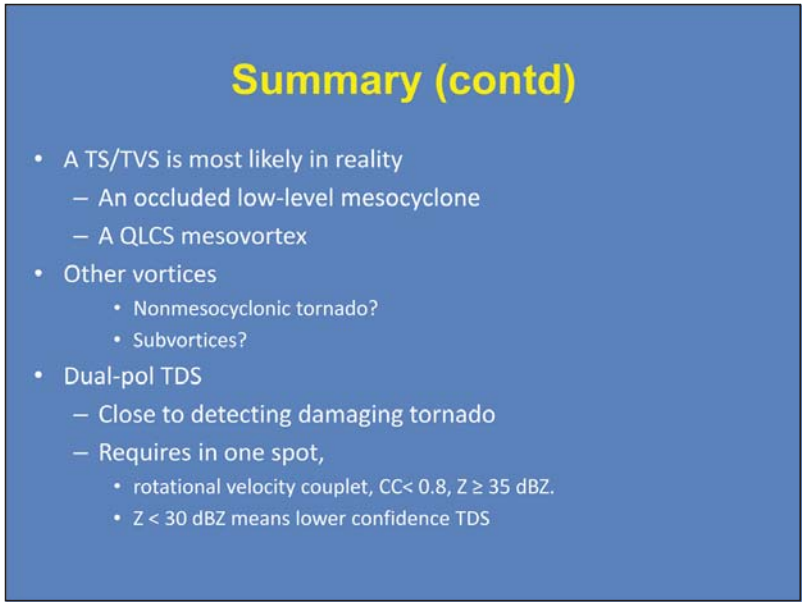
A TS exhibits both potential flow outside the vortex core, and solid body rotation within the core with Vr max, Vr min separated by at least one azimuth. This is the same as a Rankine combined vortex associated with nondivergent mesocyclones.

A TVS exhibits only potential flow and Vr max, Vr min on adjacent azimuths**.** Sometimes both Vr max, Vr min may be enveloped by the same radar beam resulting in a reduced mean velocity. In those cases, Vr max, Vr min may be separated by one effective beamwidth.

A valid TS/TVS needs exceedance thresholds in the velocity difference. However, the threshold is made to be flexible for you as a user, in order to accommodate variability in storm size and range degradation. Remember to discard hard and fixed rules on what constitutes a minimum velocity difference for an operator defined TVS.

A second criteria that is important is a TS/TVS should have some vertical continuity. We would like to see that continuity extend across at least two elevation scans. However, use 1500 m as a good starting point for vertical depth.

A third criteria is that the TS/TVS should persist for about five minutes. This criteria is not so hard since there are plenty of situations where a TVS may barely precede a tornado, if it does at all. Most often, when near the ground, a TVS or TS are ongoing tornadoes.





A TS/TVS may occasionally represent an occluded mesocyclone in a typical supercell at low levels where the radar is resolving a circulation somewhere in between the tornado scale and mesocyclone scale. Not all TVSs represent occluded mesocyclones, especially when considering non-supercell events.

Remember that you should not depend on a TVS as a primary consideration in a tornado warning. It is a signature that represents one of many cues in your decision making.

A TDS means that a tornado is almost certainly ongoing and capable of damaging any structures it impacts. However, **many tornadoes do not exhibit a TDS** due to range degradation, the strength and size of the tornado, and/or a lack of debris source.

1. Look for rotational velocity couplets in a velocity product

2. Identify an accompanying CC minimum, preferably less than 0.8

3. Check to see which CC minimum satisfies the minimum reflectivity threshold of 35 dBZ. Lower than 30 dBZ reflectivities mean a low confidence level of a TDS.

If you have passed the quiz, then you have successfully completed this lesson. If you have any questions, please contact us using any of the e-mail addresses listed on the bottom of the slide.



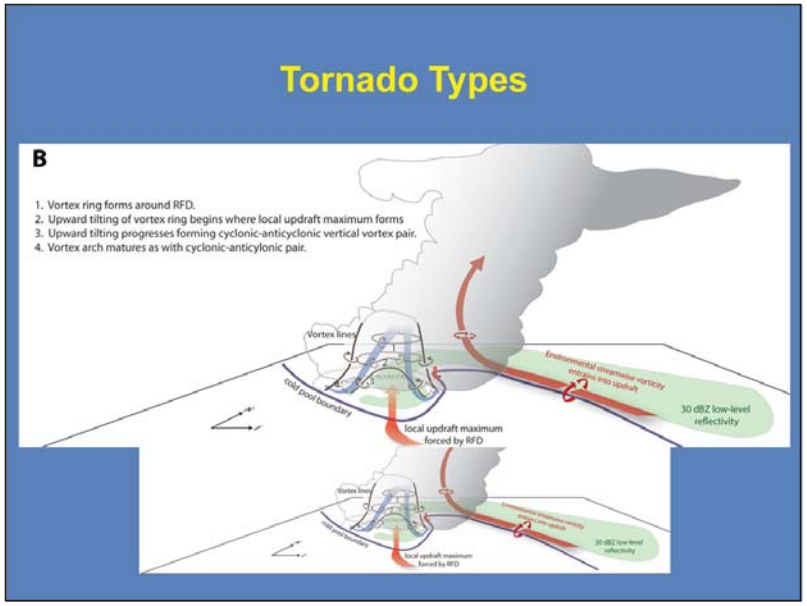
Lesson Objective

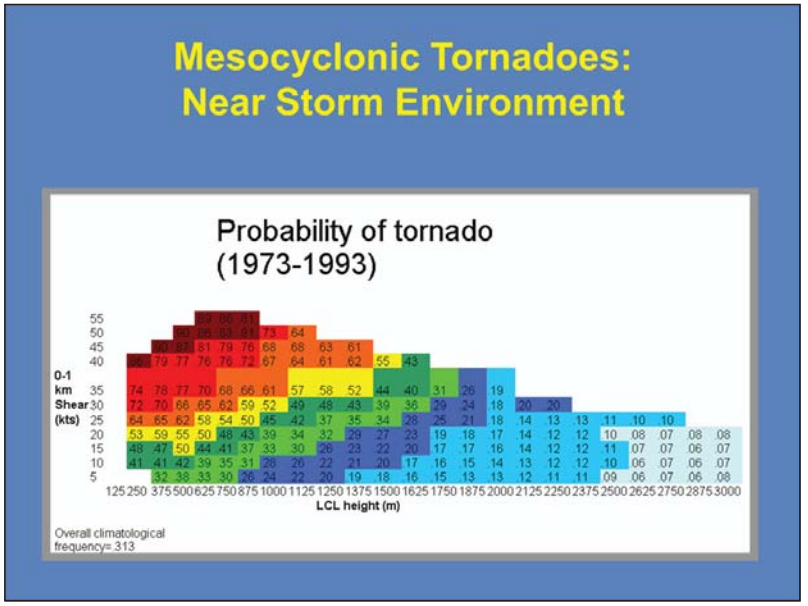
1. After reviewing the contents of this lesson,
   - Describe the quality of tornado precursor signatures associated with supercells, QLCSs and nonmesocyclonic tornadoes associated with low-shear convection.
   - Identify the typical considerations involved in the proper methodology for inferring a high threat from mesocyclonic and nonmesocyclonic tornadoes.
   - Determine the proper timing to end a tornado warning.
   - Describe the typical time-height evolution of strong azimuthal shear in a QLCS mesovortex signature
   - Identify the likelihood of a tornado based on the isolated vs multiple detections of a mid-, low-level mesocyclone signature and BWER signature.
   - Contrast the environments that favor nonmesocyclonic vs mesocyclonic tornadoes.
   - Describe how to orient a tornado warning polygon to account for the motion of mesocyclones in supercells
   - Discuss a tornado warning strategy for merging storms when one contains at least a moderate strength mesocyclone.

The problem of deciding whether or not you should issue a tornado warning creates arguably more consternation than any other warning type. Trying to anticipate a tornado is difficult since science does not clearly understand how all the ingredients for a tornado come together. In fact, we may yet not know all the ingredients for a tornado. But, we are not completely at a loss. There are many sources of information available that can supply us with evidence that we can use in making accurate tornado warning decisions. In this lesson, we will describe these sources and give some advice that you can use in your warning operations.

This lesson covers eight objectives in which you may review here before going on with the lesson.

You'll get the most of this lesson if you've had the lesson on Tornado signature identification.





Your strategy for assessing a tornado threat depends greatly on how the threat is evolving. At its most basic level, a tornado requires a high concentration of vertical vorticity embedded in converging and ascending flow into a convective cloud base. The question is, "where did the vertical vorticity originate?"

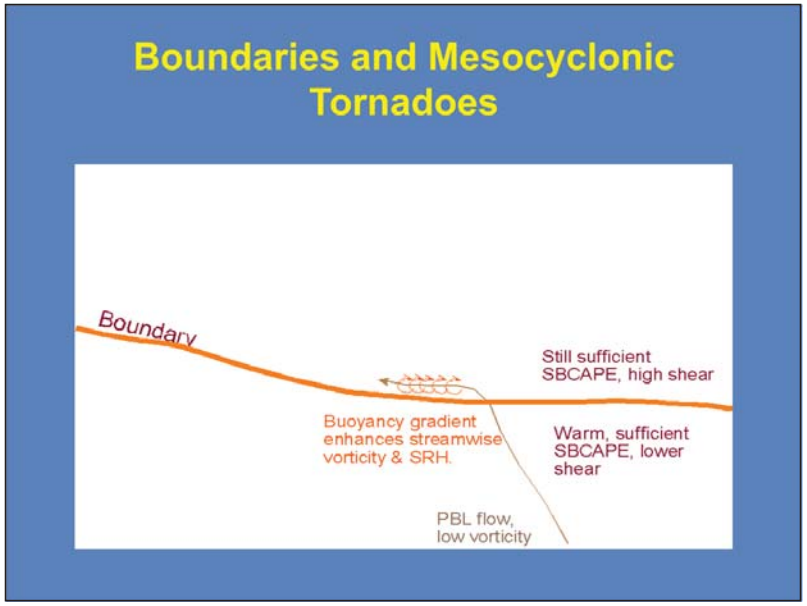
The most basic origination is that vertical vorticity was already embedded in a line of convergence and waiting for a locally intense updraft to intensify it. This type of tornado has been called a **landspout**, but is more appropriately termed a nonmesocyclonic tornado. This type of tornado is the only kind that can occur in association with updraft dominated ordinary cells. Note that at the time of tornadogenesis, the parent storm is in the towering cumulus phase where little to no downdraft is present in the cell. At this time, the boundary layer is often characterized by a steep or dry-adiabatic lapse rate offering little resistance to ascending motion. Subsequently, appreciable precipitation and a downdraft will develop, and this often leads to the demise of the tornado.

The second type of tornado occurs in an environment that had little or no vertical vorticity previously to the storm. In order to get a tornado, a negatively buoyant downdraft must be involved in order to generate horizontal vorticity

that can then be tilted upward through interaction with a positive gradient of vertical velocity from the rear to the front. An updraft cannot accomplish this by itself. We call this a **mesocyclonic tornado** that typically accompanies supercells. Many tornadoes may derive their vorticity from both sources. As a side note, Quasi-Linear Convective Systems also fit into 2nd tornado formation paradigm.

Most of the longest and most intense tornadoes accompany mesocyclones from supercells so it is important for development purposes that there is sufficient (> 15 m/s) deep layer shear and instability present in the environment. However, it is also true that most supercells are non-tornadic, so something more is needed to favor tornadogenesis.

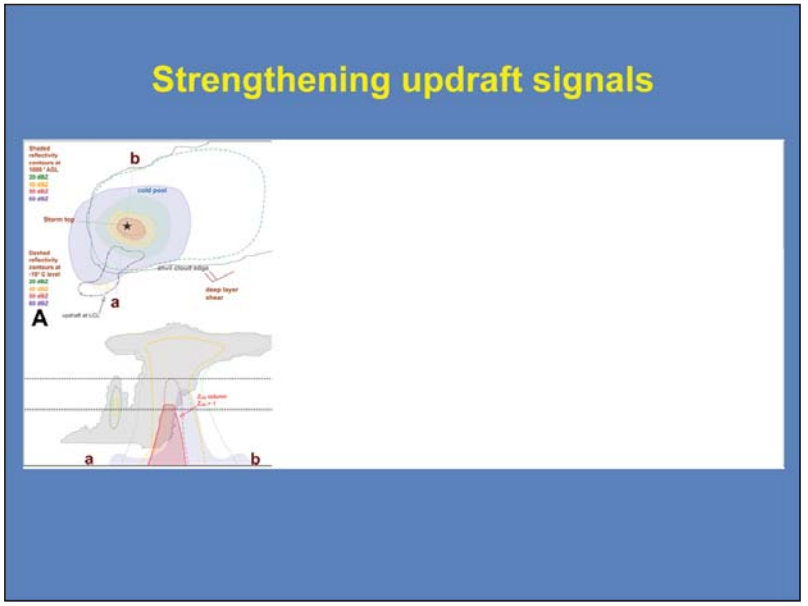
Warning decisions should involve the environment (and its changes) just as heavily as they do radar. Research has built up evidence that **mesocyclone-induced tornado environments favor strong 0-1 km shear and a low LCL.** Low-level shear helps to strengthen low-level mesocyclones. The low LCL is associated with buoyant rear flank downdrafts, allowing vertical vorticity to be easily stretched (Markowski et al., 2002). **The LCL and shear should be in an environment that promotes strong low-level updraft acceleration (i.e., low CIN and strong low-level convergence)**.



Boundaries are regions that can be locally favorable for tornadic supercells. Look for these types of boundaries:

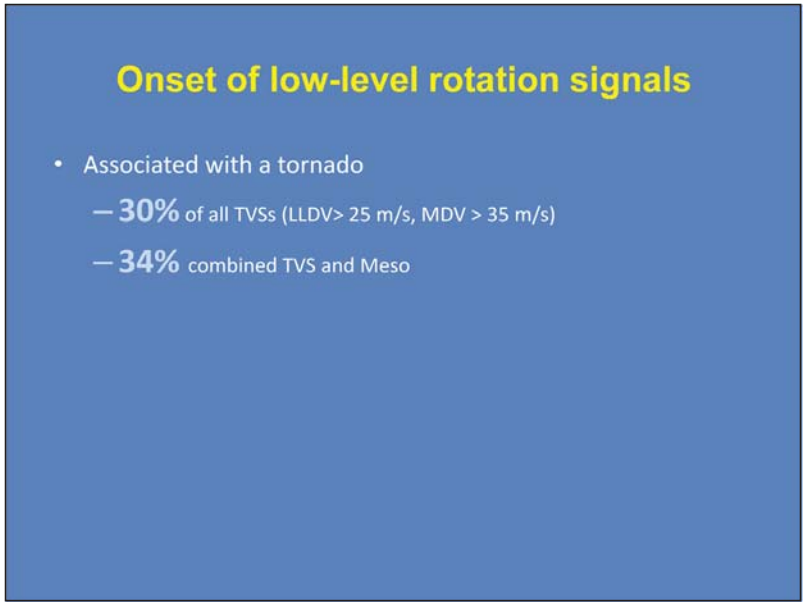
- **Subtle boundaries with backed winds and good SBCAPE providing a good clue of increased low-level shear, low LCL and little CIN , and**
- **Boundaries with strong vertical vorticity in supercell environments**.

Supercells can stretch environmental vertical vorticity along boundaries at least as effectively as pulse cells and provide little forewarning of tornadogenesis.

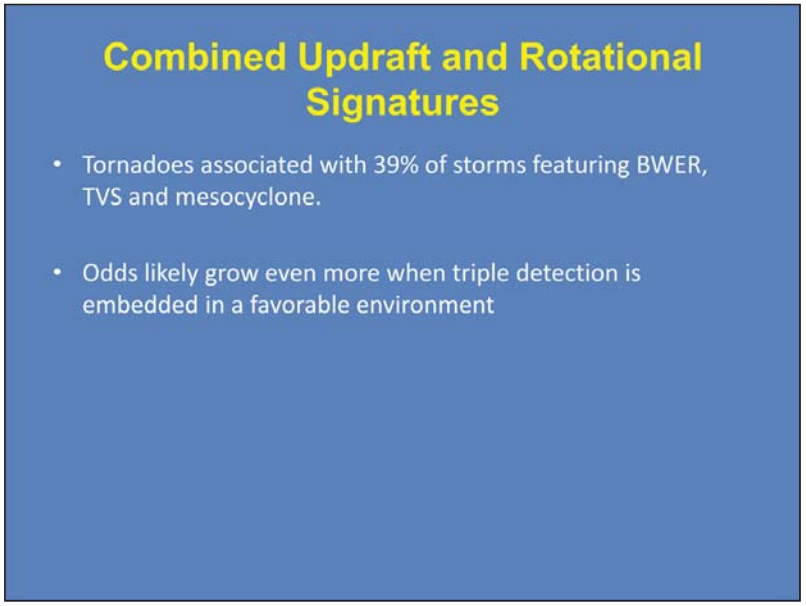


Evidence of a strong updraft in the lowest half of a storm provide even more support that a low-level circulation can be stretched into a tornado. Look for the classic reflectivity signatures, including the onset of a concavity in low-level reflectivity gradient, a strong echo overhang, displacement of the echo top over a WER, and evidence of change in storm motion.

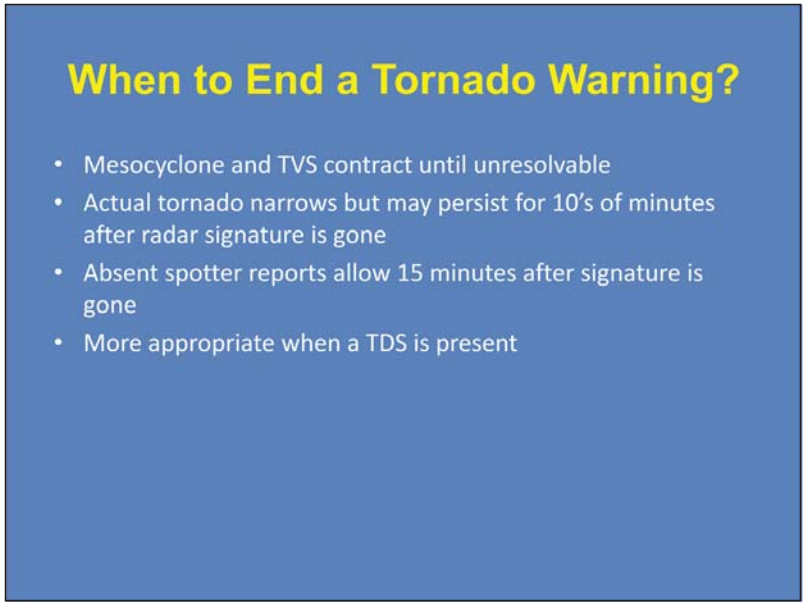
BWERs are rare, but if visible, there is an enhanced threat of a tornado when coupled with a strong mesocyclone and/or TVS/TDS. However, you should not depend on a BWER to consider a tornado warning.



Based on the National Severe Storms Laboratory TWG, 30% of all TVSs (LLDV > 25m/s; MDV > 35m/s) are associated with tornadoes (TWG, 2002). However, if the TVS is within a significant mesocyclone, the chances go up to 34% (TWG, 2002). **Look for improved odds of tornadic potential when there is a well developed and stronger mesocyclone in conjunction with a TVS**.



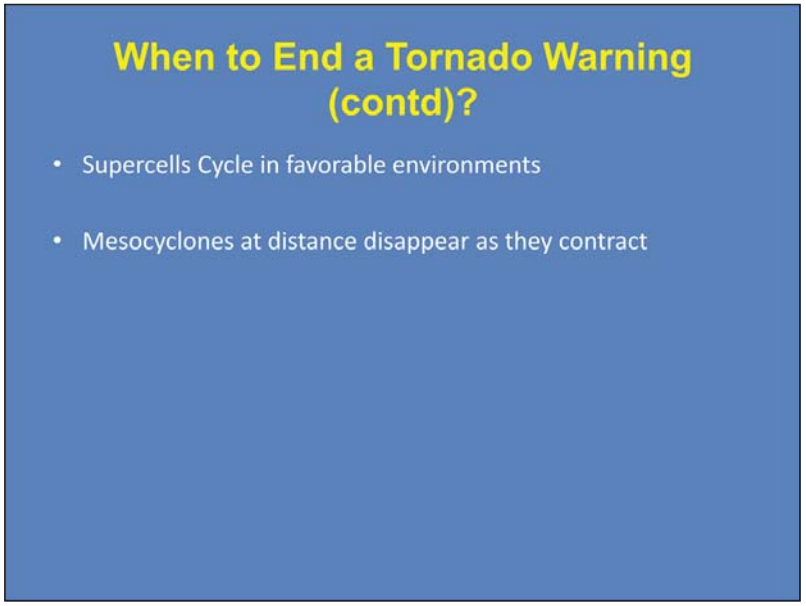
Recent research has shown that if storm has a BWER as well as both a TVS and a mesocyclone, the chances of a tornado jump to 39%. Remember that favorable parameters for tornadic supercells should be thought of as interrelated. For example, if there is a BWER, a TVS and a mesocyclone in an environment of low LCLs, steep low-level lapse rates, and no CIN, the odds for tornadogenesis are much higher than if the same three signatures were in a poor environment.



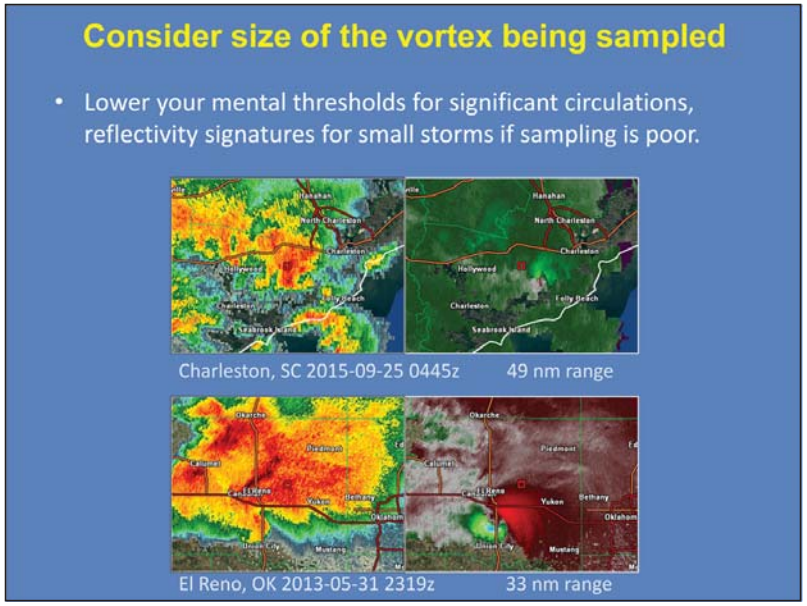
As a tornado nears the end of its life cycle, the parent mesocyclone, and even the TVS, contract to the point of becoming unresolvable. During the weakening stage, the vertical depth also typically decreases. Thus, at longer ranges, diminishing circulations, while still damaging, will likely become unresolvable earlier in the process.

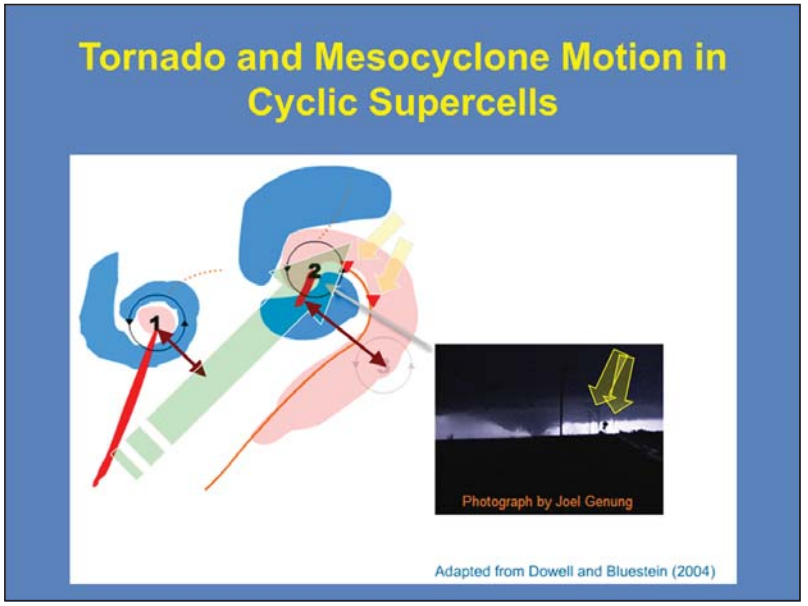
Do not terminate your warning prematurely. Allow 15 minutes for the end stage of a tornado.

This may be particular appropriate if a Tornado Debris Signature (TDS) using Dual Polarization data is apparent. Based on operational evidence, a TDS will start later and end later than the tornado (and likely any TVS signature). With high confidence, you can likely enhance wording in a warning or statement that debris is occurring with a valid TDS.



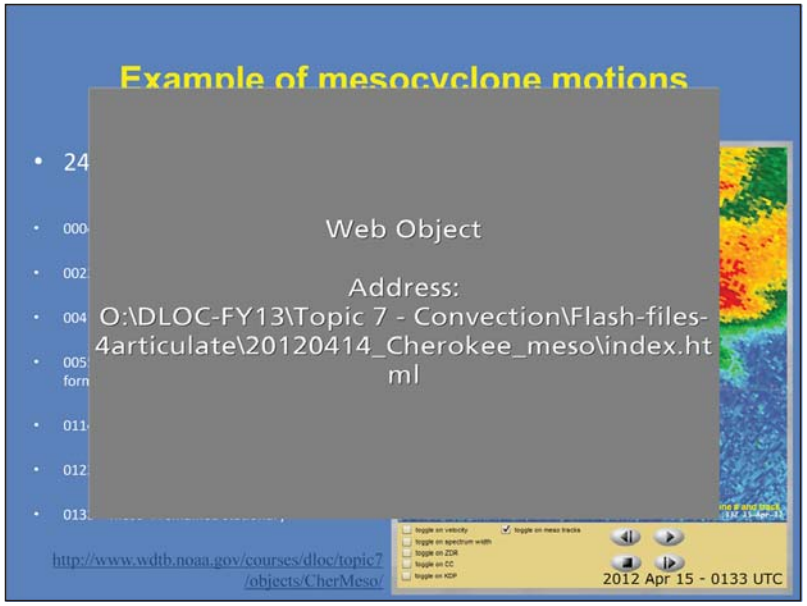
Be careful not to terminate a tornado warning for a supercell with a weakening mesocyclone/TVS signature that is rapidly approaching or already within a favorable tornado environment. The supercell may just be in the process of producing a new mesocyclone, or the mesocyclone may be contracting into a tornado. In some cases, the collapse of a BWER, lowering of the storm top, and shrinking of the mesocyclone diameter may mean tornadogenesis is underway (Burgess 1974).





The most important thing to remember is that ability for the WSR-88D to adequately sample a vortex is dependent on the ratio between the size of the vortex and the effective beamwidth. This means a well-sampled mini vortex is going to give you better details than a poorly sampled maxi supercell mesocyclone. Your thresholds for issuing a tornado should change accordingly. Lower them for poorly sampled storms and vortices, raise them when they're well sampled, given all else being equal. See the tornado signatures lesson on more details.

Account for differing tornado motions than storm motions, especially if a supercell is cycling. Recall from the lesson on supercells that older or decaying mesocyclones in cyclic supercells have significant motion to the left of the actual storm motion and newer mesocyclone motion. An old mesocyclone with a tornado can move up to 10 miles left of the main supercell path. Cyclic tornadogenesis appears to occur after the RFD to the right of a mature mesocyclone surges forward enhancing convergence at its leading edge of the cyclic mesocyclone section. The locally enhanced convergence helps initiate a local updraft which then tilts horizontal vorticity into the vertical. In a short time, a new mesocyclone forms on the head of the RFD surge while the old one continues motion to the left. The old mesocyclone often begins to move left of the main supercell either as a result of outflow from the main core pushing it rearward or from inflow locally pushing the RFD gust front backward. Meanwhile, more than one tornado may form on the RFD gust front, but outside the center of the low-level mesocyclone



This example is not in the student guide but it clearly shows how differently the individual mesocyclones move relative to the parent storm. A popup display will appear that allows you to explore in more detail what I'm about to describe here. The link is a backup to the popup. We'll take a tour of the Cherokee storm from KVNX as it moved northeast at 35 kts. I used the feature-following zoom to remove the parent storm motion from the loop. In the discussion I refer to these circulations as mesocyclones. They are low-level mesos that sometimes appear as TVSs. I refer to their motions in a storm-relative sense. The storm starts far southwest of the radar (35 mi) and passes within 10 mi west of the radar before retreating to the northeast at 25 mi at the end of the loop.

At 0004 UTC , mesocyclone 1 formed looking like a TVS. It moved to the southeast by 0023 UTC. A Dual-pol TDS developed by 0037 UTC.

At 0041 UTC, meso 1 backtracked and moved northwest as the Dual-pol TDS indicated a tornado. Meso 2 developed. Meso 2 was well centered under the parent storm updraft and produced a Dual-pol TDS by 0051 UTC as well as a strong TVS.

At 0055 UTC, meso 1 died while meso 2 was a strong TVS with a Dual-pol TDS while it moved more quickly to the west. The tornado with meso 2 was the one labeled circulation #1 in the lesson on Tornado scale signatures. Meso 3 formed well east of #2 and on an eastward extension of the RFD.

At 0114 UTC, meso 2 died while meso 3 also moved westward while quickly producing a TS and a Dual-pol TDS. The TDS was quite strong at this time.

At 0123 UTC, meso 3 moved quickly westward and is on the backside of the reflectivity envelope. The TS contracted into a TVS and the Dual-pol TDS continued. Meso 4 developed at this time.

At 0133 UTC, meso 3 died while meso 4 developed a TVS and the largest, most pronounced Dual-pol TDS of the series. This tornado was large.

Note out of all of this how deviant the tornadoes happened to be relative to the main reflectivity core of the supercell. Some of these tornadoes were up to 6 miles west the main track. This happened because the storm-relative inflow overpowered the relatively weak RFD surges sending the low-level mesos westward. However, the meso 4 remained under the parent storm for awhile likely because its RFD surge was stronger and well balanced with the inflow. The bigger tornadoes often accompany mesocyclones that can keep this balance.

Tornadoes from Storm Mergers

- RFD is enhanced to produce a tornado
- Storm ingests and tilts a high amount of streamwise vorticity to produce a tornado, or
- No tornado occurs at all.

- Key is to heighten awareness to possibilities

Be aware that mergers of a non-tornadic supercell with a gust front increases the possibility of a tornado. But a wide range of possibilities have been found to occur from this type of merger. Precisely what leads to tornadogenesis or failure is poorly understood. A non-tornadic supercell interacting with a neighboring storm may result in many evolutions:

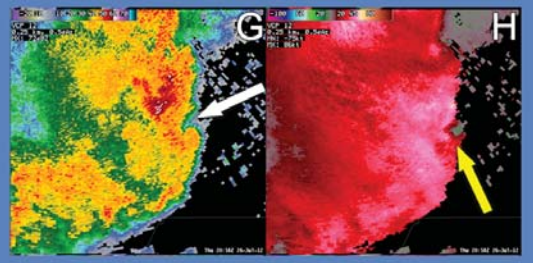
RFD is enhanced to produce a tornado

Storm ingests and tilts a high amount of streamwise vorticity to produce a tornado, or

No tornado occurs at all.

Additionally, collisions between left- and right-moving supercells may or may not assist tornadogenesis. **The key lies in attaining a heightened awareness of the possibility of a tornado when you see an impending storm merge with another storm or gust front.**

Squall Line Tornadoes

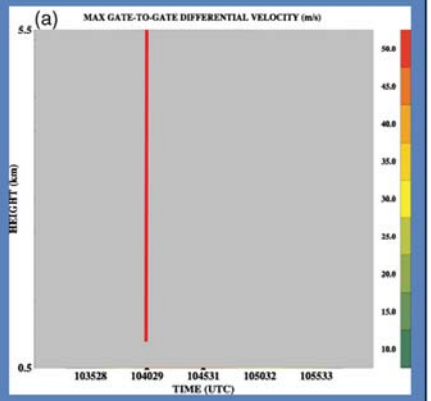


Because both bow echoes and supercells require strong vertical wind shear, supercells and severe bow echoes often occur in close proximity to one another, or evolve from one of these structures to the other during their lifetime. **Environments of bow echo tornadoes and supercell tornadoes are hard to distinguish.** Thus, it is important to examine storm structure of each individual cell within multicell structures, as supercell tendencies are frequently observed with well-organized multicell systems (i.e., those which develop in either sufficient deep shear and/or large CAPEs).

Well-defined front inflow notches often show up in reflectivity data to the north of a surging area of outflow prior to tornadogenesis. **In these cases, the vertical vorticity is enhanced by strong surface convergence and is at a maximum at low-levels. Thus, tornadoes in squall lines typically form much quicker (mean lead time of 5 minutes in the Trapp study) than with isolated supercells.**

Typical QLCS TVS Evolution

- QLCS = Quasi Linear Convective System
- Non descending 80% of the time



Sampling Considerations for QLCS Tornadoes



Poor sampling typically > 40 nm distant.

Tornadoes within quasi-linear convective systems tend to be associated with tornadic vortex signatures (TVS) that form from low-levels upward as opposed to some classic supercells which have midlevel circulations first and then build downward with time (Trapp et al, 1999). This non-descending paradigm for TVS evolution

Squall line tornadoes are quite often very difficult to detect at more than 20-40 nm away from the radar, as evidenced by the radar representation of this tornado-bearing squall line in southern New York on 26 July, 2012 . Notice the evolution of the front inflow notch in reflectivity from 2045 UTC to 2058 UTC. Tornadogenesis occurred around 2049 UTC scan. We are fortunate to be able to see the circulation here. Imagine if we were further away.

Low Shear Tornadoes: Radar Sampling Considerations

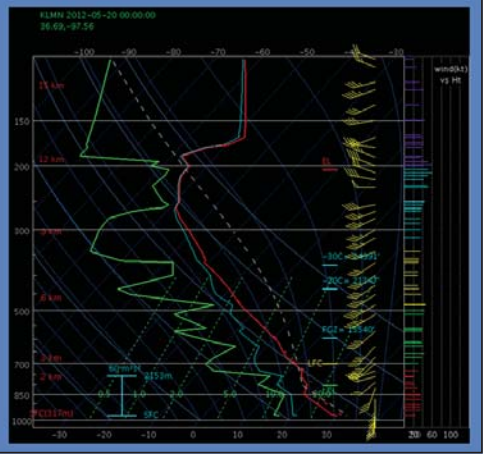
- Misocyclones typically < 2 km in diameter
- Are shallow
- Need to detect clear air boundary returns

- Absolutely depend on spotters

Because the origin of the rotation in these features is close to the ground, radar may easily overshoot the circulation unless it is within 50 km. The diameter and lifetime of these misocyclones are small (<2 km, 5-15 min) also limiting the ability of radar to resolve their velocity structures. This process creates a difficult job for Doppler Radar to provide adequate lead time of pulse storm tornadoes. Spotter reports of funnels in these situations should be taken very seriously (Lemon and Quoetone, 1995).

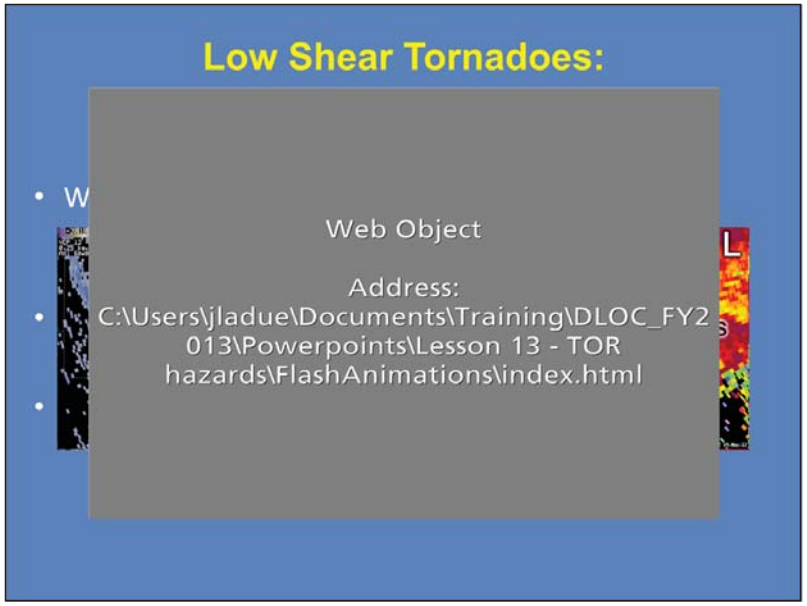
Low Shear Tornadoes: Sounding Considerations

- Steep lapse rates
- No CIN



Favorable conditions leading to the formation of a pulse storm tornado include:

An environment with steep lapse rates, strong surface heating and no CIN.

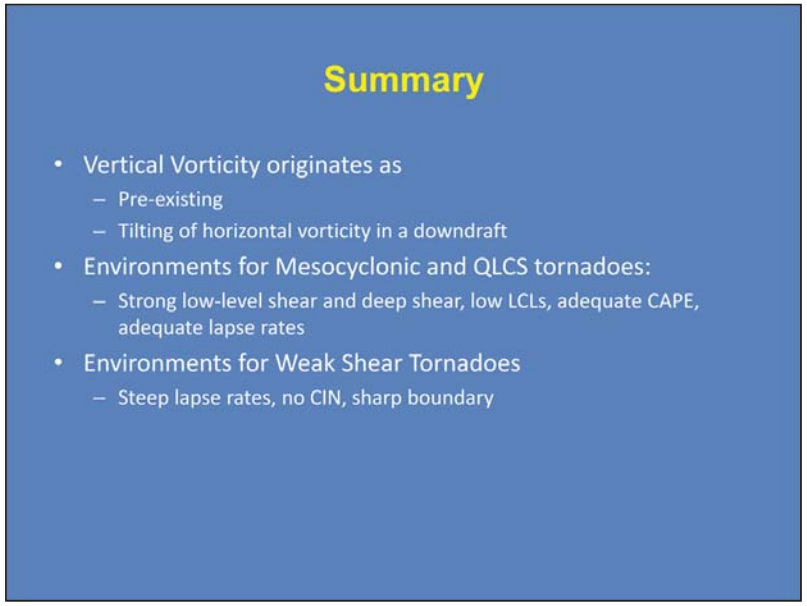


A popup window will display showing a case of a low-shear nonmesocyclonic tornado outbreak.

A well-defined boundary marked by a fine line in reflectivity, velocity discontinuity in base velocity, or a cumulus line visible from satellite. The boundary should have significant vertical vorticity. **Note that 10 m/s of shear across a 1 km wide boundary produces the same vorticity as a moderate mesocyclone (10-2 s-1).** Ideally, the boundary and cell motion should be nearly equivalent. Boundary collisions are also common regions of pulse storm tornadoes. An example of tornadoes forming along a boundary is shown in this image.

Colliding or intersecting boundaries with high potential for vertical vorticity production is an area to be closely monitored (Wakimoto and Wilson, 1989).

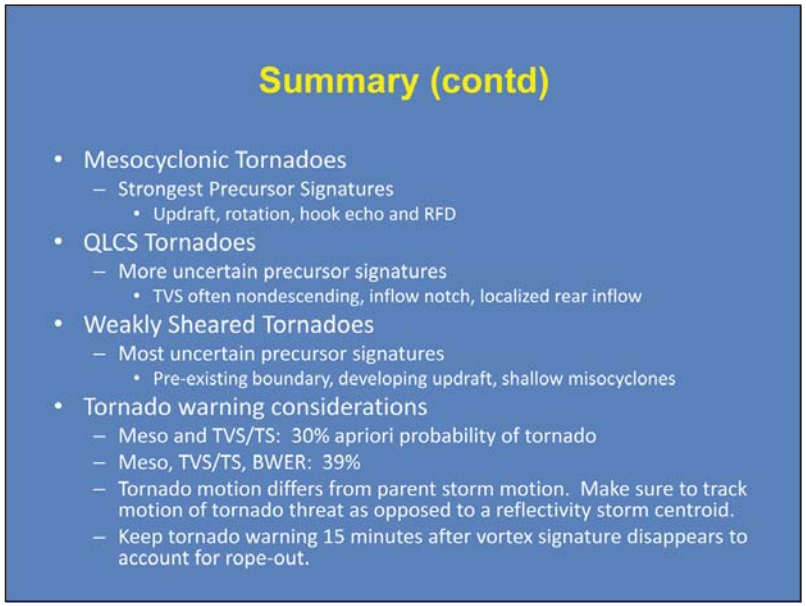
over the boundary is a strong clue. Note that many of these tornadoes are produced, at least initially, without nearby precipitation and may be hard to detect. At other times, thin lines close to the radar may attend the boundary. But in either case, visual observations can be critical in raising situational awareness. Be especially alert for tornadoes when the updraft forms an elevated reflectivity core.



Tornadoes derive their vorticity either from pre-existing vertical vorticity, or from tilting of horizontal vorticity into the vertical by a downdraft.

Three types of tornadoes we discussed in this lesson include mesocyclonic tornadoes, Quasi Linear Convective System (QLCS) tornadoes (also known as squall- line tornadoes), and nonmesocyclonic tornadoes.

Mesocyclonic and squall line tornadoes are favored in similar environments of strong low-level shear, strong deep layer shear, low LCLs, and sufficient CAPE. Nonmesocyclonic tornadoes require a sharp boundary with strong vertical vorticity, an uncapped atmosphere featuring significant CAPE, steep low level lapse rates, and a developing updraft riding the boundary.



Mesocyclonic tornadoes in conjunction with supercells are easiest to anticipate as they are typically preceded by strong low- and midlevel updraft signatures, strengthening circulation aloft, and an onset of a hook echo and RFD. Confidence of a successful warning rises with the simultaneous presence of multiple signatures including TVS, mesocyclone, and BWER.

Squall line tornadoes typically do not have deep updraft signatures, however low-level updrafts can be very strong. They typically occur on the leading edge of a bowing line segment and may be preceded by the onset of a front inflow notch. TVS signatures likely suddenly appear in a non descending fashion. A parent mesocyclone is not a requirement.

Nonmesocyclonic tornadoes are difficult to detect via radar given the lack of a preceding deep circulation. However, given close enough proximity, radar can infer the presence of low-level circulations. They are more likely to occur as a young updraft phases with a pre-existing low-level circulation.

Recognizing the different types of conceptual models for specific tornadic storm development is crucial in effective tornado warning decision making. A thorough analysis of the environment in which tornadoes occur, in addition to the best possible radar interrogation strategies, must be the prime considerations for an effective tornado warning methodology.

Following are considerations in tornado warnings.

Meso and TVS/TS: 30% apriori probability of tornado

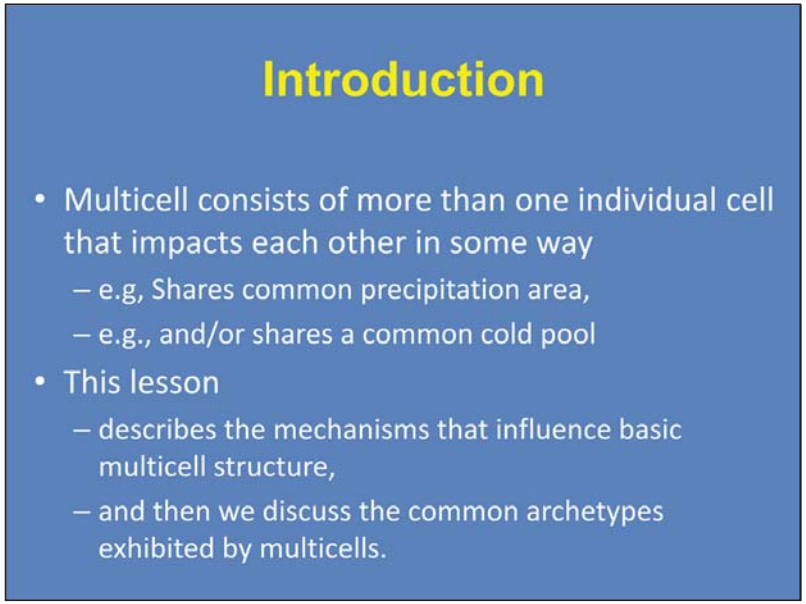
Meso, TVS/TS, BWER: 39%

Tornado motion differs from parent storm motion. Make sure to track motion of tornado threat as opposed to a reflectivity storm centroid.

Keep tornado warning 15 minutes after vortex signature disappears to account for rope-out.

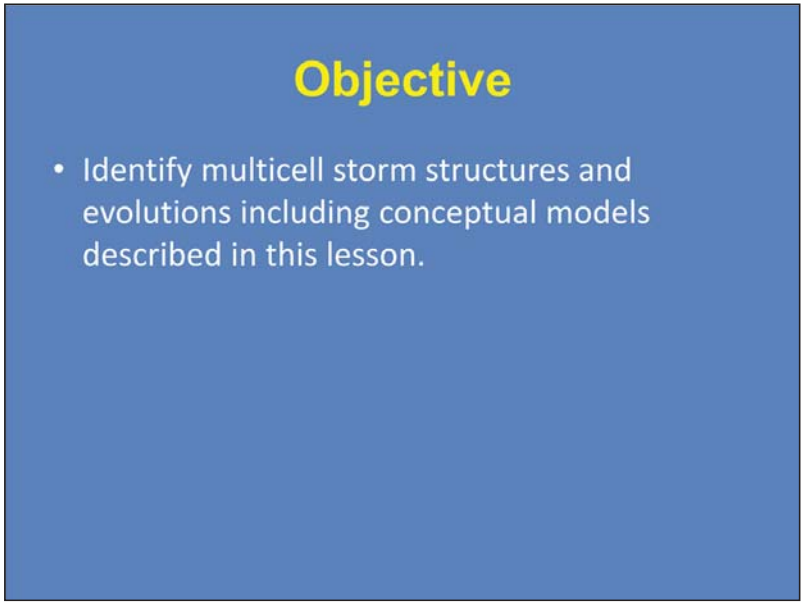


Welcome to the RAC Convective Storm Structure and Evolution lesson on Multicell Archetypes.

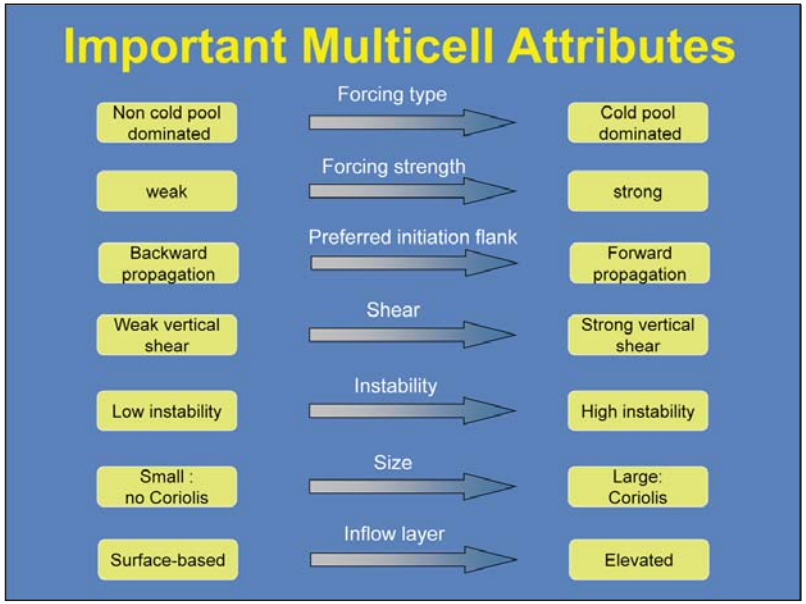


Multicell storms consist of individual cells, either ordinary or supercellular, in close enough proximity to affect each other in some way. For the purposes of this lesson, multicells are groups of two or more cells that at least share a common precipitation area and a cold pool. In nature, most Deep Moist Convection (DMC) becomes multicellular because there is typically more instability and forcing than one cell can alleviate. It is very rare for a single cell to be initiated in complete isolation from subsequent initiation; therefore, multicells are common in the broad parameter space of instability and vertical wind shear. However, the combination of forcing, vertical shear, and instability has an impact on the size and organization of multicell structures.

This lesson has two parts. The first part describes the mechanisms that influence basic multicell structure. The second part discusses the common archetypes exhibited by multicells.



The objective is to identify multicell storm structures and evolutions including conceptual models described in this lesson.



The categorization of multicells is quite complicated owing to the large variety of documented structures and forcing mechanisms that are immediately relevant to your severe weather forecasts. No one single multicell categorization scheme has been developed. Instead, let's describe the major characteristics which influence the type of multicell:

The type and orientation of forcing. Many multicells are driven primarily by their own cold pools. However, multicells can be dominated by non cold pool processes. The orientation of forcing may also influence the development of a multicell.

Strong forcing can certainly hasten the initiation of new cells more quickly than weak forcing (all other factors being equal).

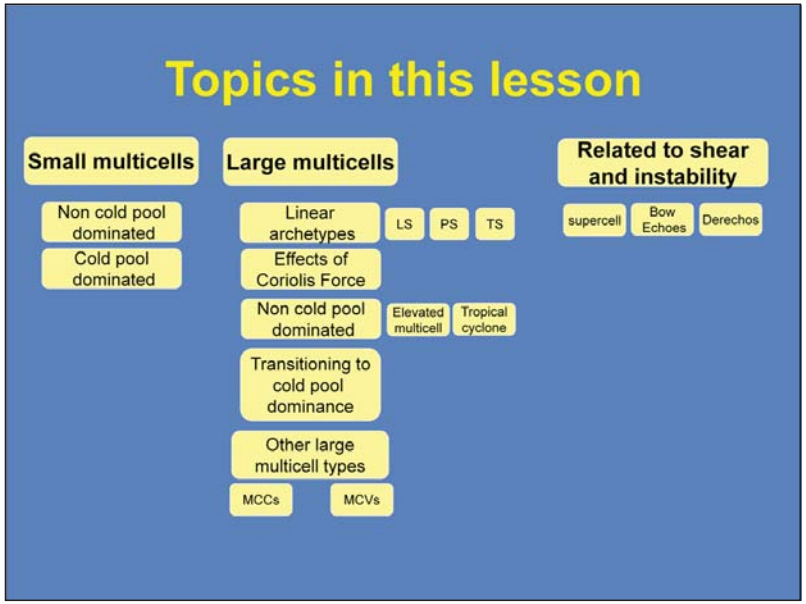
The location of forcing relative to the multicell influences its motion.

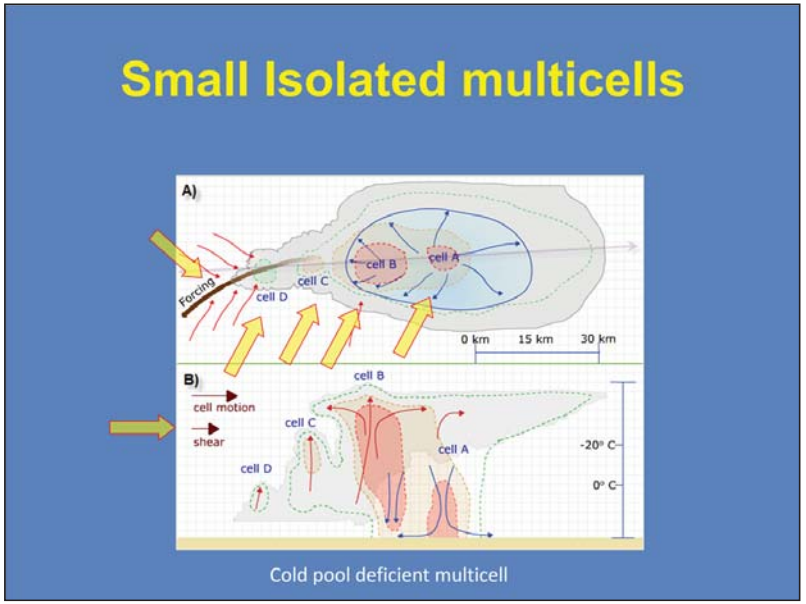
The vertical wind shear of the environment interacts with the cold pool and influences the ability for new cells to generate. We explain this more in the lesson on multicell motion. The interaction of the vertical shear with updrafts within the multicell also influences the potential for new cell generation. This is especially true for multicells with weak cold pools or for multicells with embedded supercells. We explained this in the lesson on supercell dynamics.

The vertical stability profile influences the ease at which forcing can initiate new convection and the ability for a multicell to create a significant cold pool.

Multicell size and the influence of the Coriolis force are positively related to each other. As multicell size increases, so does the likelihood of developing a dominant cold pool, developing other internal circulations (such as a rear inflow jet), and the likelihood that the Coriolis force is a significant influence.

And lastly, **the parcel inflow layer of a multicell** influences its behavior depending to what degree the updraft/downdraft base is elevated.





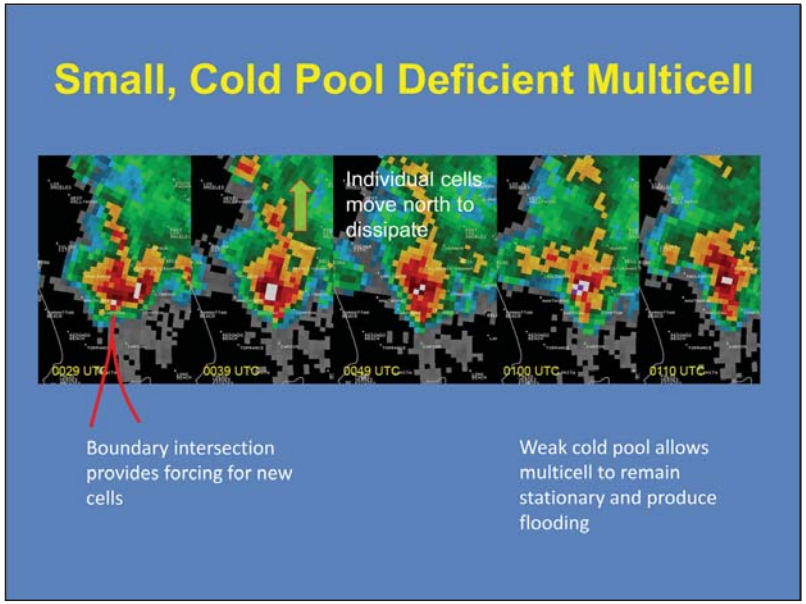
In this lesson we will talk about a variety of structures and archetypes. And we'll split them up according to size, because that's seems to have the greatest impact on how to discriminate multicell events. We'll start with small multicells and sub-categorize them by whether or not their evolution is dominated by their own cold pool.

Then we'll discuss large multicells. We'll talk about linear archetypes and a variety of sub-archetypes there. We'll also talk about the effect of Coriolis force on large multicells. Large multicells can also be non-cold pool dominated, so we'll talk about the elevated large multicells and even tropical cyclones. These large events can transition to cold-pool dominant events and we'll talk about the process that occurs there. And we'll also talk about other large multicell archetypes that have been defined using other instruments. Mesoscale Convective Complex's (MCCs) have been defined using satellite-based instruments. We'll also talk about Mesoscale Convective Vortices (MCVs).

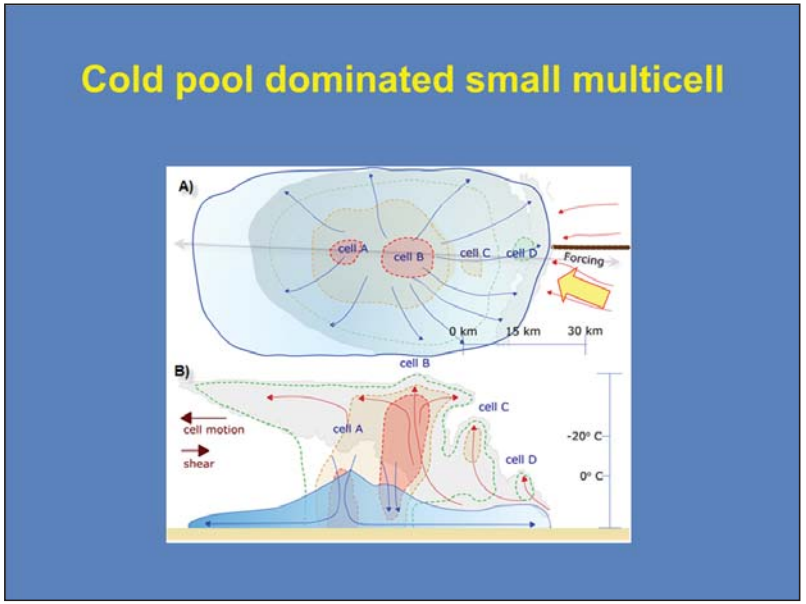
All of these events are greatly impacted by their environmental shear and instability. For example, we can get supercells out of small multicell events. We can also get bow echoes and derechos.

The smallest (meso-g scale) multicells typically feature several individual cells in different stages in development. You may have seen these also called isolated multicells. New cells form before old ones dissipate but in close enough proximity to share cloud material, and precipitation. Typically there are a small number of cells in any one stage in the lifecycle as depicted in this schematic. Cell A represents the dissipating stage where the precipitation core is raining and fully occupied by downdraft. Cell B is in the mature stage where the heavy precipitation core is descending with downdraft and the updraft bubble has overshot its equilibrium level. Cell C and D represent the newest members of this multicell and are primarily updraft dominant. This type of multicell is too small to create a linear organization.

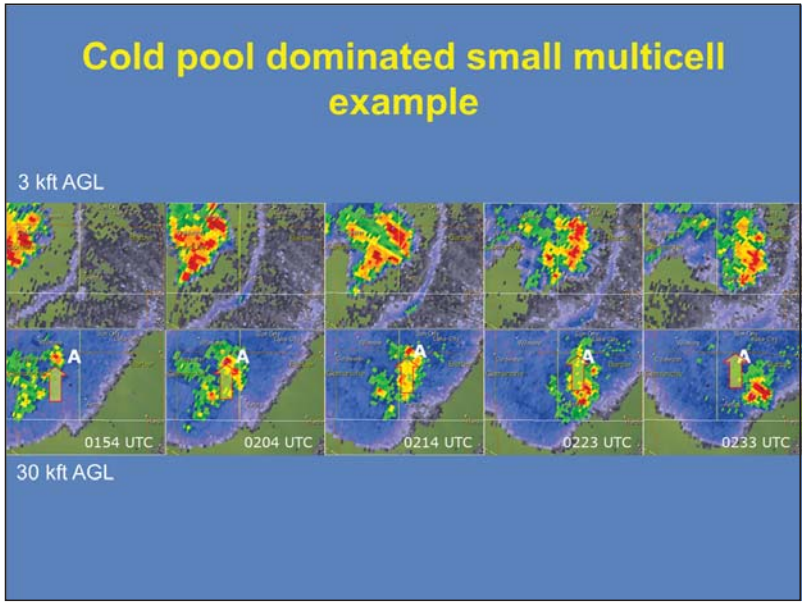
In the conceptual model shown in this figure, new cells are forming along an axis of forcing well outside of the cold pool boundary. The cold pool, while present, is too deficient to dominate the initiation of new cells. Small multicells are most likely to be cold pool deficient, and therefore, be dependent on external forcing and instability. While the forcing is depicted upstream of the steering layer flow, it does not have to be in order for this to be a non cold pool dominated small multicell.



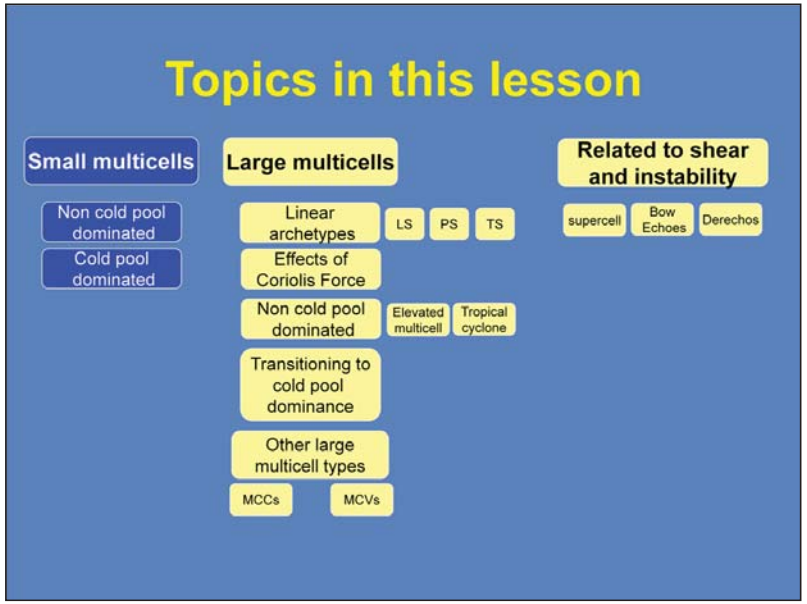
The multicell shown in this figure was primarily driven by a point source forcing manifesting as a nearly stationary interaction between a sea breeze front and a cumulus cloud roll. New cells formed at the intersection, then progressed north and dissipated on the north end of the multicell. The cold pool was too weak to force the multicell to move, and a backward propagating, flash flood producing multicell developed in south Los Angeles.



A small multicell could wind up being dominated by a significant cold pool in an environment conducive to its formation. Here is a conceptual model of a cold pool dominated small multicell. New cell initiation is forced primarily by ascent from the leading edge of the cold pool. The gust front typically moves faster than individual cells, and therefore, this type of multicell outpaces individual cell motion. We call this forward propagation. Note that the cell motion label here is relative to the whole multicell. Also note the line of "forcing. It's not necessary, but it does help to focus cell initiation in one spot.



Here is an example is of a small, cold-pool dominated multicell attempting to match speed with a surging gust front. Watch how cell A forms and quickly matures then dissipates as it moves to the east much more slowly than the gust front and the multicell.



Now, we'll talk about large multicells.

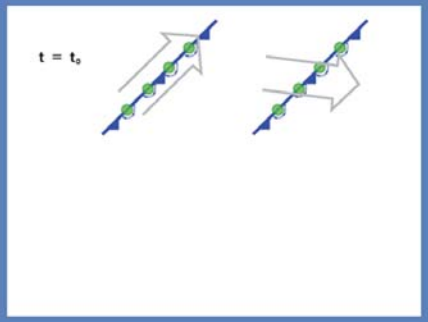
Large Multicells

- From meso-β to meso-α scales
- Many cells in similar stages of development
- Mesoscale Convective System (MCS)
  - Exceeds 100 km in any direction
- Impact of Coriolis force can be significant

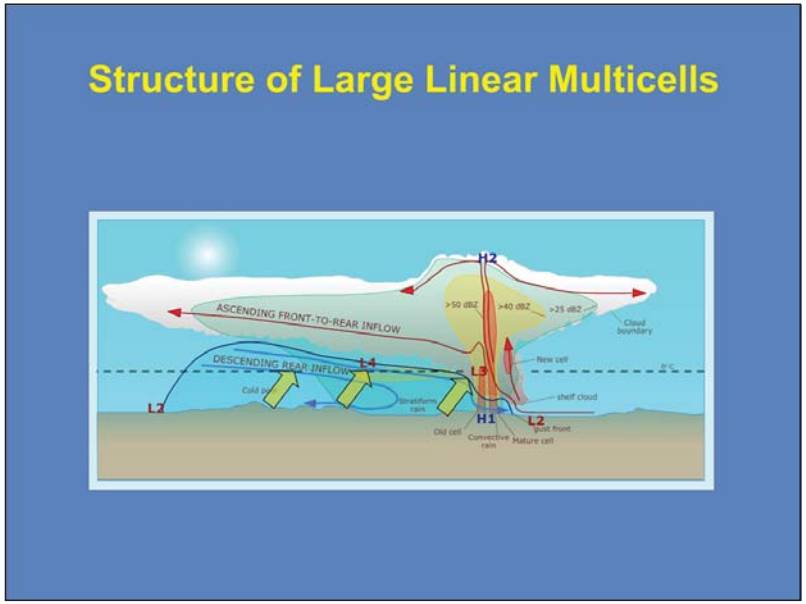
The dimensions of a large multicell can range into the meso-b and meso-a scales and persist for several or more hours. Large multicells typically contain many cells in similar stages of development.

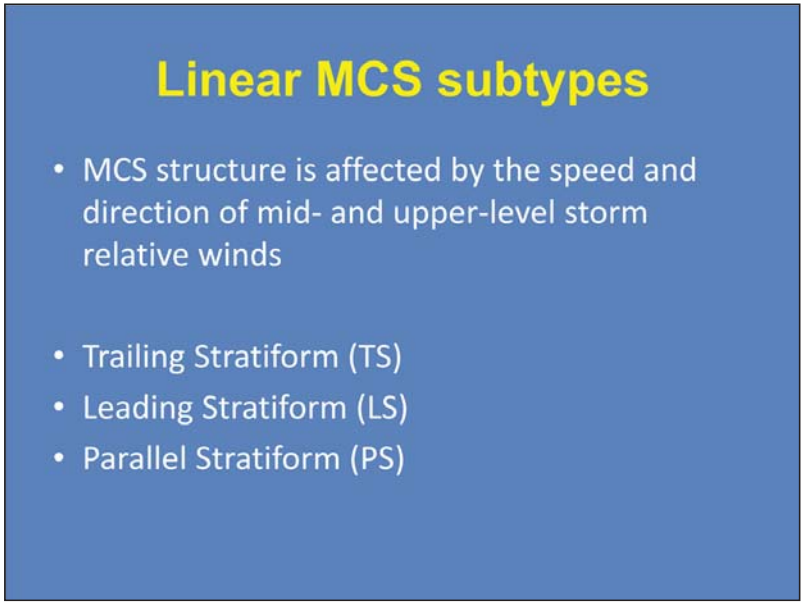
A **Mesoscale Convective System (MCS)** is a multicell whose contiguous precipitation area exceeds 100 km (54 nm) in any direction. However a large multicell need not satisfy such criteria. Unorganized MCSs may exhibit multiple flanks of sporadic, relatively infrequent, new cell initiation. Meanwhile, organized MCSs may exhibit relatively frequent new cell initiation on a preferred flank. MCSs produce large anvil shields and subsequent areas of stratiform precipitation in addition to strong system wide circulations that influence its structure and evolution. Only in persistent MCSs does the Coriolis force become a significant influence in its evolution.

The linear nature of large multicells



Large multicells are more apt to exhibit a linear nature to them reflecting the elongated lifting that commonly occurs along external forcing mechanisms (e.g., fronts), and internally generated cold pool boundaries. Fronts provide linear forcing, but multicells may not merge into a long line if the forcing is weak. However, if the deep layer shear is largely boundary-parallel, individual cold pools may more easily merge, reinforce the front, and enhance upscale growth into a long line.





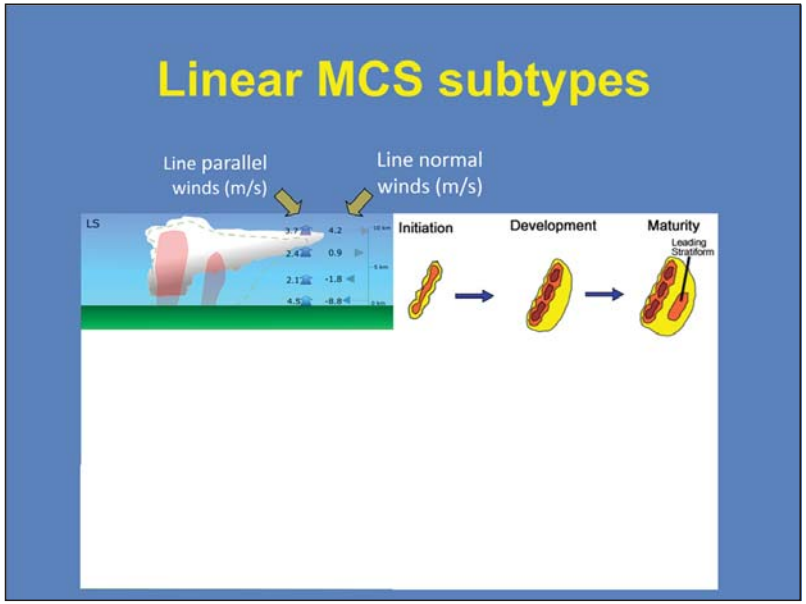
In addition to the propensity for linear development, MCSs also develop significant system wide circulations as the large anvil and cold pool modulate the pressure field from the surface to upper levels. Large areas of stratiform rainfall fall from underneath the anvil.

Two of the more widely accepted conceptual models of the complex flow structure are from Smull and Houze (1987) shown here.

A **Rear-inflow jet (RIJ)** is a mesoscale region of strong winds that originate in the trailing stratiform rainfall region of a squall line near the top of the cold pool and are directed toward the leading edge**.** In the Smull and House (1987) model of a mature MCS, development of the RIJ is attributed to mid-level, mesoscale areas of low pressure (labeled L3 & L4). The mesolow "L3", which forms immediately behind the leading line convection, is a hydrostatically-induced, negative pressure perturbation that develops under up-shear tilted warm, convective updrafts and above the evaporatively cooled downdrafts. Mid-level mesolow "L4" forms in the stratiform region in between the warm, buoyant air which gets pulled rearward past the cool, dry descending air flow. Note that the major difference between this figure and a small cold pool dominant multicell is the presence of a stratiform precipitation region, and attendant RIJ; both are features that more frequently accompany large MCSs.

Although MCSs develop a number of ways, typical mature systems contain both convective and stratiform precipitation regions. The eventual MCS type is determined to a large extent by the environmental conditions in which it develops and the strength of the system cold pool. Parker and Johnson (2000) studied numerous MCSs and determined the distribution of hydrometers and stratiform precipitation shapes were largely a result of mean storm-relative winds. **The speed and direction of the environmental mid- and upper-level winds relative to system motion affect the resulting evolution of the MCS.**

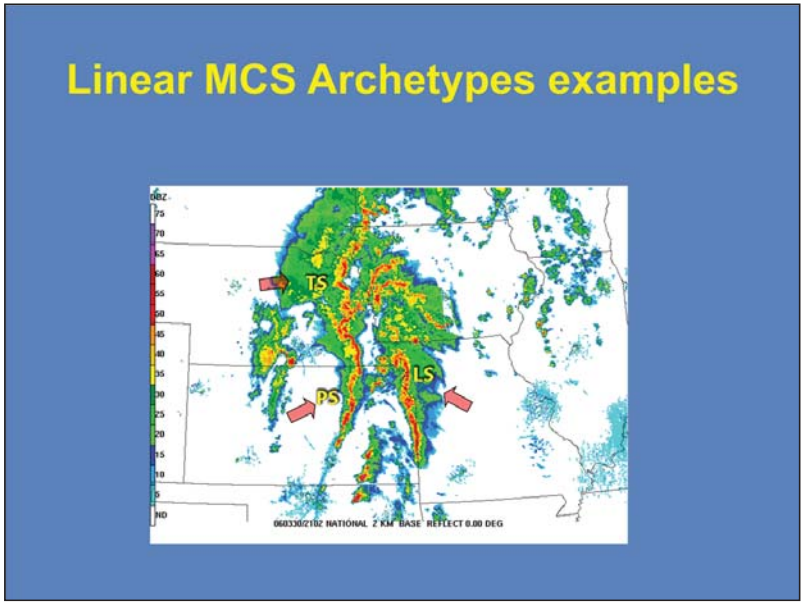
According to Parker and Johnson (2000), MCS squall lines evolve into three major archetypes: **1) trailing stratiform, 2) leading stratiform, and 3) parallel stratiform**. The main distinction arises from storm-relative flow fields.



The Leading Stratiform (LS) precipitation squall line archetype, which is typically slower-moving than trailing stratiform systems, is characterized by stronger mid- and upper-level storm-relative flow (often described as rear-to-front flow) than any of the other types. It's not so much the line parallel wind, but rather the line-normal wind that carries the anvil debris ahead of the main convective line. As a result, they tend to produce the weakest cold pools and may sometimes not be cold-pool dominated.

The Parallel Stratiform (PS) MCS structures represent a slightly less common form. The anvil and attendant stratiform precipitation expand outward, both in front of and behind the intense convective line. They occur in situations with strong along-line storm-relative flow, especially in mid- to upper-levels. They produce strong cold pools and therefore it is difficult to maintain this archetype for long time periods.

The Trailing Stratiform (TS) squall line type has a sloped front-to-rear flow produced by stronger system-relative flow in low-levels (and subsequent stronger low-level convergence along the leading edge). The line normal wind is from front to rear, which produces clouds and precipitation behind the line, which generates a strong cold pool with a rear inflow jet and provides most of the forcing for its maintenance. Trailing Stratiform squall lines tend to move quite quickly.



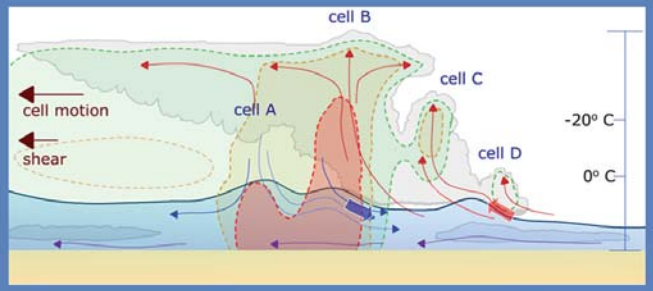
This figure shows an example of all three linear MCS archetypes occurring simultaneously ahead of an ejecting, strong, upper-level shortwave trough. Each MCS formed from a different boundary. The Trailing Stratiform formed on a stationary front north of the surface low, the Parallel Stratiform formed on a dryline, and the Leading Stratiform formed on what may have been a residual outflow boundary with the cooler air to the east. This case was document by French and Parker in 2006.

Effect of the Coriolis Force on Large Multicells



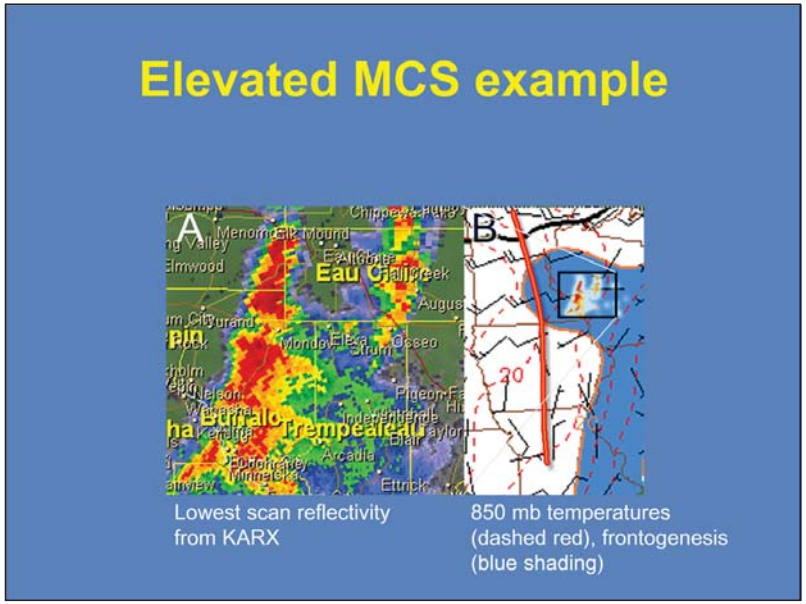
Large multicells, or Mesoscale Convective Systems (MCSs), are the convective events most subject to behavioral changes from the Coriolis force. As shown in figure, most linear MCSs develop book-end (or “line-end”) vortices through the tilting of horizontal vorticity, either from the environment or along the cold pool edge. Provided a sufficiently long lifespan, the cyclonic member of the system becomes reinforced by the Coriolis force while the anticyclonic member is weakened. The MCS becomes deformed and eventually exhibits a comma shaped configuration. A mid-level hydrostatic low under the anvil shield of the MCS also persists long enough to allow the Coriolis force to create a cyclonic circulation.

Large Non Cold Pool Driven Multicells: Elevated MCSs



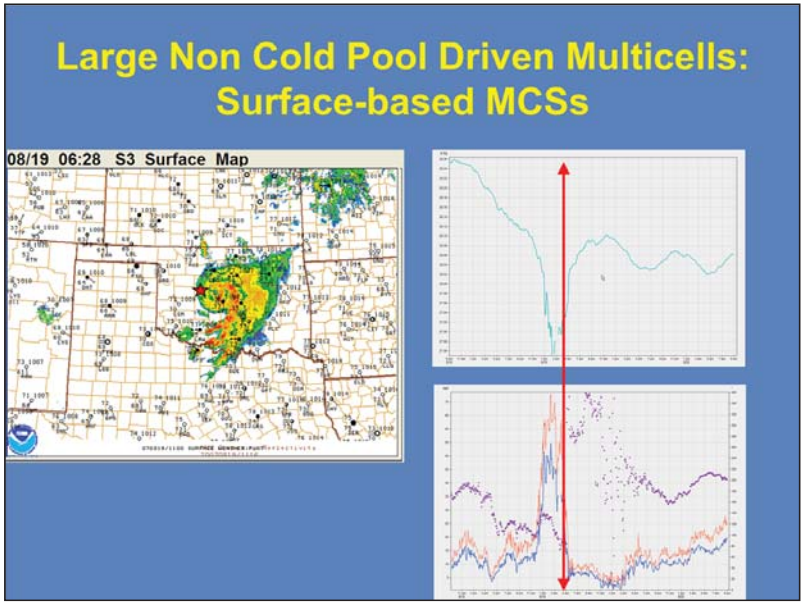
While most of the large multicell archetypes that we have discussed are dominated by cold pools, there are a host of large multicells embedded in environments that do not allow the cold pool to become significant, or are embedded within environments that do not allow the cold pool to become a significant driver.

The elevated MCS is one type of multicell likely not to be affected by a cold pool. Elevated multicell deep, moist, convection implies that the updraft parcel roots for each cell has a source above the ground, and that the air near the ground is conditionally stable. The morphology of elevated multicells depends even more on the shape and intensity of the original forcing within the context of the vertical stability profile. Forcing mechanisms are more predominantly associated with elevated lifting such as differential vorticity advection, localized warm advection, elevated frontogenesis, or gravity waves. Elevated multicells produce downdrafts; however, the resulting gust front is ineffective at creating new surface-based convection as long as the near surface air has zero Convective Available Potential Energy (CAPE). Downdrafts penetrating into the stable layer may not reach the ground to create a gust front, and instead may crated gravity waves that help initiation new convection as you can see in this figure.



Here is an elevated MCS example from July 2007 that shows a group of cells that initiated along a north-south axis which corresponds with a region of 850 mb warm advection and frontogenesis, and a region of elevated effective inflow base where buoyant parcel ascent was possible.

What happened to this MCS will be covered momentarily.



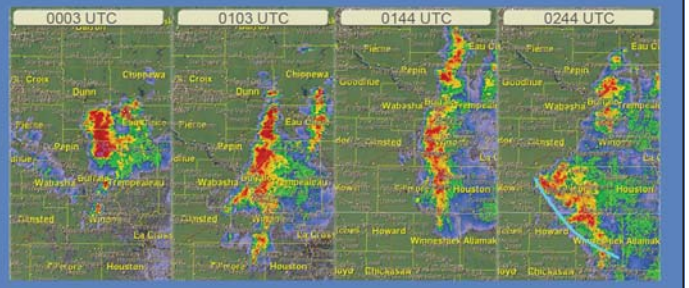
A surface-based MCS is also likely to be cold pool deficient in very moist environments with minimal DCAPE, or if there is significant DCAPE, a low-level heating source acts to modify any cooling.

An extreme example of this kind of multicell is a tropical cyclone. This is a warm core multicell whose pressure minimum under convective heating does not become concealed by a dense cold pool near ground, and the forcing is internally driven.

A tropical cyclone typically requires a constant heating source such as warm water in order to maintain itself and mitigate cold pools. However, similar structures have been found over land where cold pool production is weak. In this slide, the remnants of tropical storm Erin were reinvigorated as the circulation center redeveloped a new MCS. In turn, the MCS did not produce a cold pool, and a warm core low intensified right down to the surface.

Where the star is located, you can see the corresponding meteogram, with sea-level pressure in the top panel and winds in the bottom panel. Notice the wind maximum was co-located with the minimum of pressure, very typical of a tropical cyclone passage. The red arrow corresponds to the time of the radar image.

Transition from non cold pool, to cold pool driven multicells



Let's return to the elevated MCS from July 2007 and show what happened to it. Cold pools are prone to deepening as multicells grow upscale and/or persist. A multicell may transition from forcing-dominated to cold pool driven as lifting increases over the cold pool or by the gust front. This sequence shows our elevated multicell was dominated by frontogenesis and/or warm air advection forcing from 0003-0150 UTC. However, by 0244 UTC, a cold pool had formed which was large and powerful enough to force surface-based convection along the gust front boundary. As a result, the multicell propagated to the southwest.

Transition from Non Cold Pool to Cold Pool Driven Multicells contd



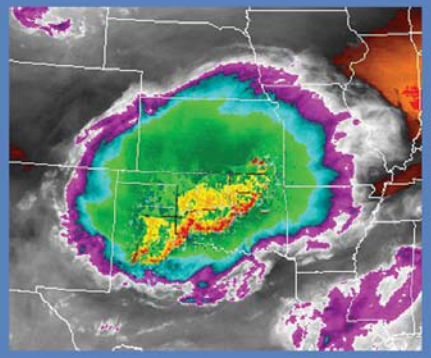
Initial convection and weak cold pool

Cold pool forms and balances with inflow shear

Cold pool grows to dominated the multicell

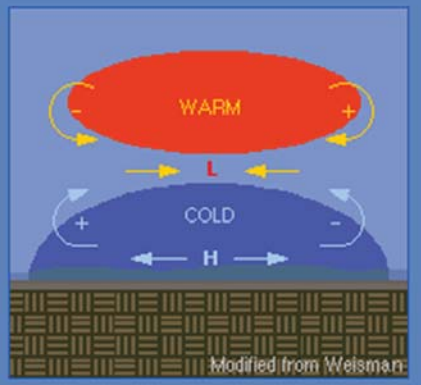
Another common example of this transition occurs when a line of discrete deep, moist convection transitions into a squall line. In this case, discrete multicellular convection initiates from an external forcing mechanism and generates a cold pool that strengthens over time. Early on, multicell forcing may be dominated by either external mechanisms or updraft-induced dynamic pressure gradients. Eventually, the cold pool deepens and becomes the dominant forcing mechanism. How long this process takes depends on the strength of the cold pool compared to the vertical wind shear and the strength and orientation of the initial forcing.

Mesoscale Convective Complexes (MCCs)

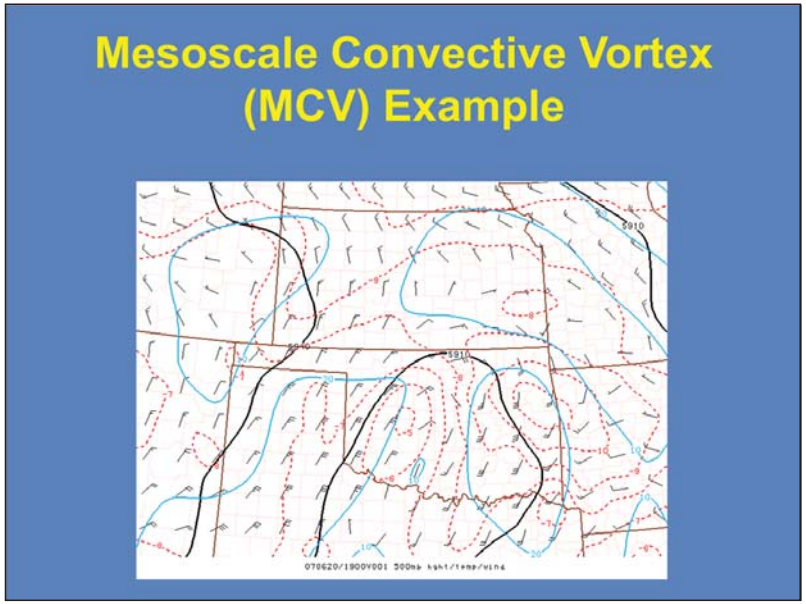


A Mesoscale Convective Complex (MCC) (defined by Maddox in 1980) is a subset of MCS that exhibits a large, circular (as observed by satellite) long-lived, cold cloud shield. The MCC's circular anvil shield masks the linear nature of the active convection underneath.

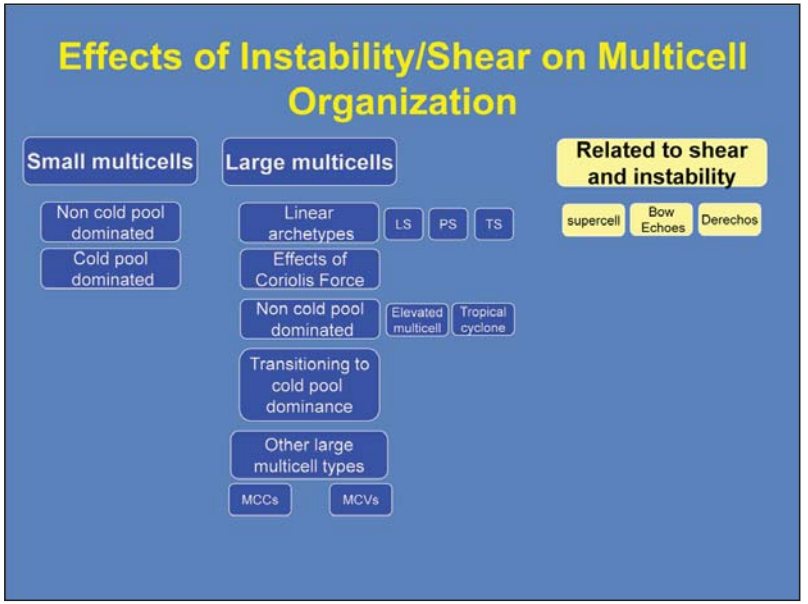
Mesoscale Convective Vortexes (MCVs)



Regardless of the underlying linear shape of the intense convection, the anvil's circular shape may allow for a coherent Mesoscale Convective Vorticity (MCV) maximum to form in the midlevels as anvil-layer heating produces a hydrostatically induced low. MCVs tend to persist more when the vertical shear is not strong. Perhaps that is why they tend to form most often on the equatorward edge of the midlatitude westerlies, and in the tropics. While many MCS structures produce a midlevel hydrostatically induced low, the shape and size of MCC anvils makes them most likely to produce long lasting MCVs which can, in turn, produce a variety of multicell structures the following day.



The MCC two slides ago produced a substantial MCV that became detached from the westerlies and persisted for days producing repeated episodes of storms and flash flooding over the southern Plains.



Now that we've covered both small and large multicells, let's now talk about how they are impacted by shear and instability.

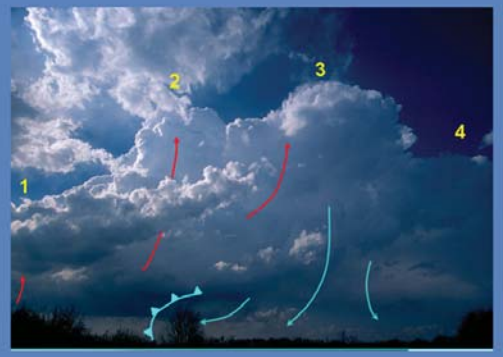
We first describe what it means for a multicell to be organized. More frequent and stronger initiation on a preferred flank of a multicell is one of the criteria for identifying a multicell as organized. In this section, we'll describe that process and what kind of structures result from that. We'll concentrate mostly on bow echoes and derechos because we discuss supercells in other lessons.

Weak Cold Pool Small Multicells – Weak Shear/Instability



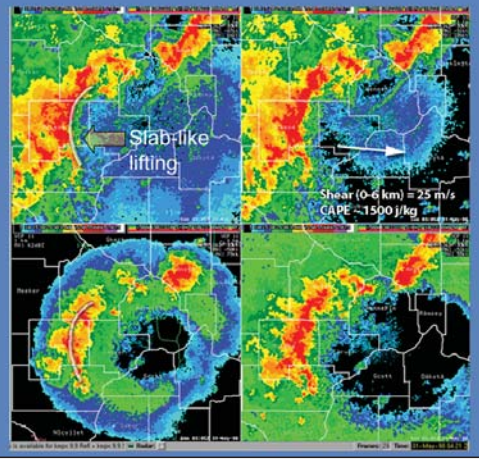
The weaker the instability and/or vertical shear, the more infrequent and weak cell regeneration. In this example, the individual cells are regenerating so infrequently that they are almost discrete.

Weak Cold Pool Small Multicells – Strong Shear/Instability



More organized small multicells exhibit an appearance of a more persistent, plume-like updraft, and adjacent heavy precipitation area as the rate of new cell initiation becomes more frequent. You may find that eventually the multicell acquires supercellular characteristics.

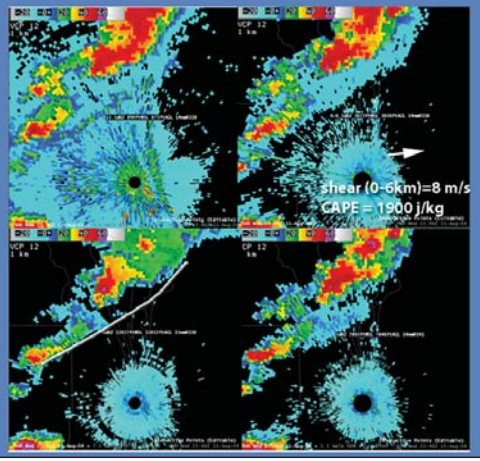
Large Cold Pool Dominant Multicells: Strong Shear/Instability



The impact of increasing vertical wind shear on a large, cold pool dominant MCS allows for the increased possibility of a continuous, stronger updraft along a preferred flank where a downshear component exists. An optimal state of shear would entail low-level shear well-balanced with the cold pool boundary-induced circulation, and/or sufficient deep-layer shear to allow the convective steering layer flow to match the gust front. Multicells satisfying these conditions exhibit slab-like lifting where you may have a difficult time discerning any discrete character to the cells. This image provides a good example of slab-like lifting along a gust front facing downshear. This event produced widespread significant wind damage.

We have a lesson on the environmental impacts on multicell organization in another lesson.

Large Cold Pool Dominant Multicells: Weak Shear/Instability



In this example, the difference between the shear and the cold pool strength has likely caused an imbalance which has limited lifting and the efficiency for creating new cells. You can see only scattered cells forming well behind the gust front. A shear/cold pool balance can be upset should the cold pool intensify such as changing of the instability while holding shear constant or changing the shear while holding instability constant.

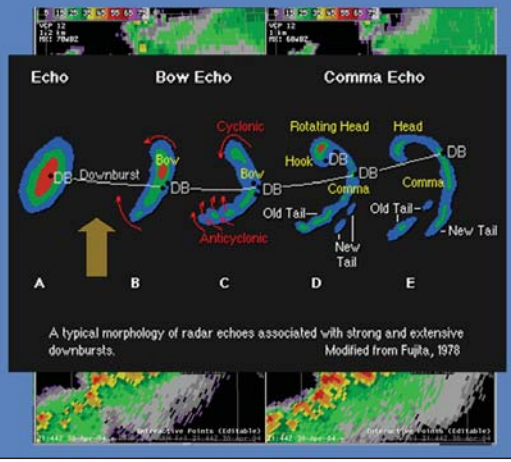
The most important component of the shear that impacts gust front lifting is that which is perpendicular to the gust front. And here you can see that the shear is mostly parallel to the gust front with almost no perpendicular component.

Derechos

- Concentrated of severe wind reports 400 km in length, some wind gusts > 65 kts, no gaps > 3 hours.
- Are severe weather events, not a form of multicell
- Produced by strong shear/instability
- Produced by MCSs often containing bow echoes.

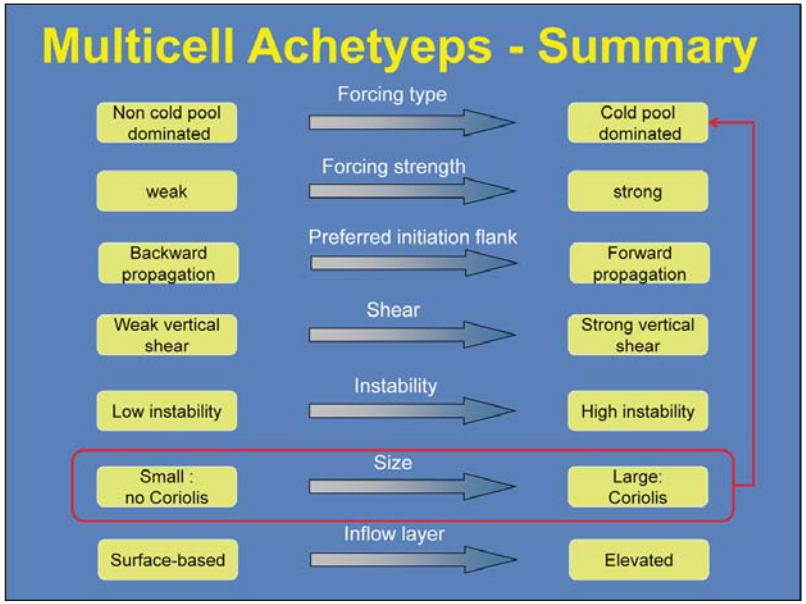
Strong shear interacting with a strong cold pool allows a cold pool forced multicell to become more severe and longer lasting, and can result in a **derecho**. A derecho is a widespread convectively induced straight-line windstorm that exhibits a concentrated area of damaging winds with a length of at least 400 km, shows an organized damage swath, contains gusts greater than 65 kt, and does not show gaps of more than three hours. These are severe weather events, not really a form of multicell, but they are produced by strong shear and instability combinations and are also produced by MCS's containing bow echoes.

Bow Echoes



A bow echo is a bow-shaped, multicell line of convective cells that is often associated with swaths of damaging straight-line winds and sometimes tornadoes. The bow appearance occurs because the precipitation has been deformed into the characteristic shape by a rear-inflow jet (RIJ). Bow echoes are meso-b scale features and often have lifetimes between 3-6 hours. Severe bow echoes are strongly favored in high CAPE/shear environments. They represent almost continuous, slab-like lifting along a deep gust front of a cold pool dominated multicell.

Here is a four panel reflectivity image of an organized cold pool dominated multicell where the center picture has been taken from point A in the upper left panel. In this figure, an organized, small multicell is taking on a bowing configuration as strong environmental vertical shear interacts with the deep cold pool. If you were to assign a stage in this bow echo life-cycle, it would likely fall between stage A and B in this schematic. Note that the most obvious manifestation of the bow occurs typically well after the onset of severe surface winds.



When you describe multicell archetypes, you may want to describe the multicell with fundamental attributes (much like adjectives) that help determine their behaviors. They include:

Small to large multicells, which represent a spectrum where large indicates an MCS potentially influenced by the Coriolis force. Large multicells are more likely to be cold pool dominant. Changing shear and instability impacts small and large multicells in different ways.

Cold pool dominant to non-cold pool dominant, multicells represent fundamentally different forcing mechanisms.

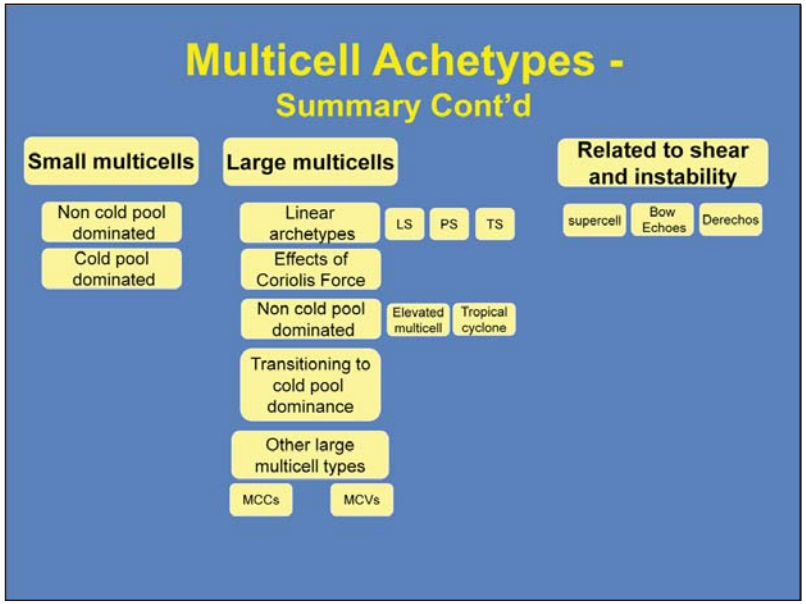
Non-cold pool dominant multicells are primarily forced by mechanisms external to the multicell (e.g., fronts, gravity waves, differential heating). Their shape, movement and size match closely with that of the forcing. Cold pools do not affect true elevated multicells or they are too weak to impact surface-based multicells.

The cold pool dominated multicells implies that they are forced by internally generated cold pools, though the strength of that forcing is modulated by the environmental vertical shear, stability, and the multicell itself. Multicells commonly evolve from non- cold pool to cold pool dominated modes as the multicell and cold pool grows with time.

The location of forcing relative to the multicell can determine its motion. If the forcing for new cells is to the rear (front) of the multicell then backward (forward) propagation is imminent.

The vertical wind shear and stability, assisted by strength of forcing, modulates the strength and frequency of new cell initiation. Small multicells may evolve into supercellular behavior with increasing shear. Large, cold pool dominated multicells may experience slab-like lifting with a balance between cold pool strength and shear. The severity of multicells is positively correlated with shear.

Finally, the behavior of the multicell is greatly influenced by its ability to tap the boundary layer.



The shape of the cloud and the precipitation distribution are other important descriptors of multicells that yield information valuable to a forecast process.

Most large multicells acquire some linear organization to them because of the shape of external forcing or by its own cold pool. Squall lines are an example of such accompanied by a brief burst of strong winds.

With linear large multicells, the stratiform precipitation distributions relative to the active convective line in large multicells lead to the following archetypes:

•Leading Stratiform precipitation (TS),
•Parallel Stratiform precipitation (PS),
•Trailing Stratiform precipitation (LS),

Trailing Stratiform MCSs tend to be cold pool forced, produce the most wind events of any MCS archetype and are the most common. Parallel Stratiform MCSs also produce intense cold pools, the updrafts tend to be more vertically oriented, and are responsible for both severe wind and heavy rainfall. Leading Stratiform MCSs are least likely to have strong cold pools.

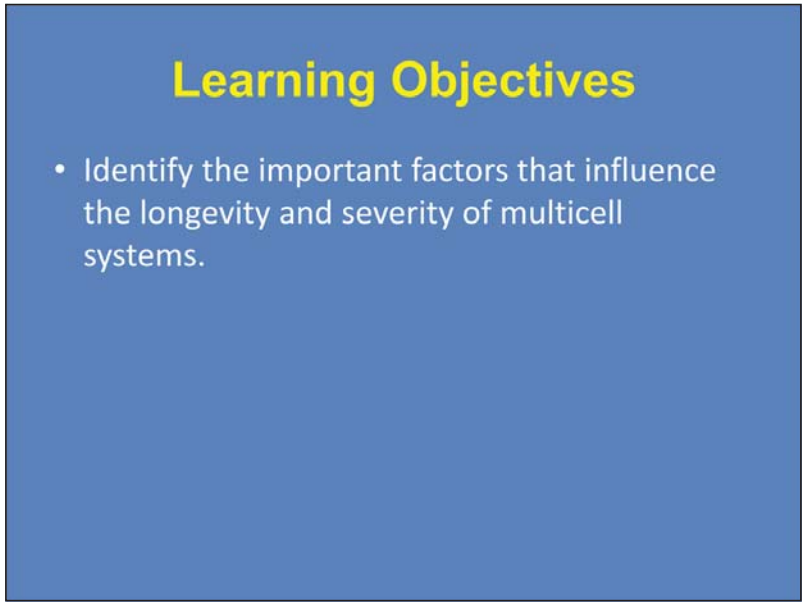
Bow echoes form a common precipitation shape that evolve either from organized small multicells/supercells, and from larger linear multicells of the Trailing Stratiform or Parallel Stratiform variety. They feature a forward bowing of an intense convective line ahead of a focused rear inflow jet or downdraft outflow and bracketed by at least one bookend vortex. They form in strongly sheared environments. Derechos represent long-lived severe wind episodes and often contain one or more bow echoes.

Multicells can be organized around other common precipitation shapes. A tropical cyclone is a large multicell circularly organized around a common low in the absence of a significant cold pool. Most small, isolated multicells are too small to acquire a linear organization and so we often call these clusters. Large clusters do occur, especially with a nonlinear shaped forcing mechanism.

A Mesoscale Convective Complex (MCC) is a version of an MCS archetype with a roughly circular, and large anvil top. The shape of their anvils, and the relatively light vertical shear environment makes them especially conducive to forming Mesoscale Convective Vortices (MCVs) in midlevels, and anticyclones in upper levels.



Welcome to the RAC Convective Storm Structure and Evolution lesson on Multicell Longevity and Severity.



Identify the important factors that influence the longevity and severity of multicell systems. Note that when we talk about multicells, there are lots of types of structures, as was mentioned in the previous lesson. We will use the terms: multicells, MCSs, squall lines, and even bow echoes throughout this lesson synonymously since we are focusing on the physical mechanisms.

Data Sets for Discerning Multicell Factors

- Proximity Soundings (observed and forecast)
- Numerical modeling simulations (RKW)

Bluestein and Jain (1985) → OK Squall Lines

Johns and Hirt (1987), Johns (1990) >>>> Derechos (Progressive, Serial)

Evans and Doswell (2001) >>>> Derechos

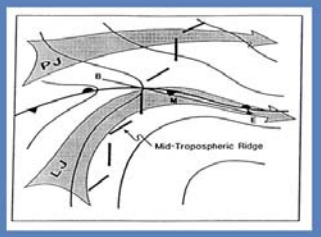
Weisman (1992, 1993)

Coniglio et al. (2004, 2006)

These are some of biggest sources for researched information for multicell factors. Over the years there have been some important differences in the results of the studies. We are going to try and focus on the most recent and consistent findings within both modeling and operational communities. In the student guide, you'll see numerous studies of squall lines, bow echoes, and the type of reports associated with derecho events. Bow echoes were studied exhaustively by Johns et al, and then by Evans and Doswell who enhanced the study to include all types of events by forcing. Weisman developed a lot of the RKW theory that is based on numerical weather simulations of multicells. Finally, we'll take a look at some of Coniglio's simulations which emphasize some of the impacts of mid and upper level shear on maintaining the severity of squall lines.

Previous Studies of Multicells

- Pattern forecastable but predictability limited
- Climatology biased toward warm season events
- Tried to find optimal values

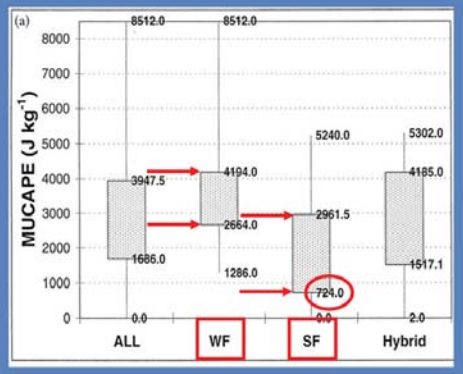


Warm Season
Progressive Derecho
after Johns (1993)

Previous studies of multicells focused on patterns that were forecastable but predictability was limited. Climatology of these types of researched events were biased toward warm season events. They tried to find optimal values that represented parameters and conditions favorable for warm season events.

What are the Factors that Influence Multicells?

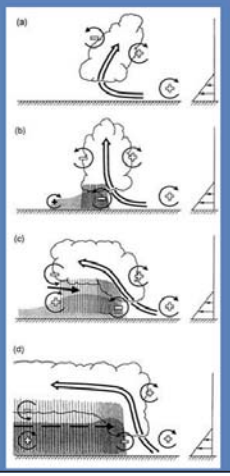
- CAPE Effects
- Wide range of values observed



Multicell systems can be long-lasting depending on the environmental conditions and resulting propagation. There are also modulating forcing effects and how the Rear-inflow Jet helps orient the deep lifting along the leading edge of convection. Needless to say, there are multiple mechanisms at work which influence longevity which makes multicell systems like Mesoscale Convective Systems very interesting. We are going to start with a look at the buoyancy or CAPE effects. In this study by Evans and Doswell, they analyzed proximity soundings of derechos, which are long-lasting multicell systems that can last from 2 to 24 hours and travel across multiple CWAs. They found a rather broad CAPE parameter space for derechos for various types of forcing. For weakly forced events (the WF group on the x-axis), the shaded boxes show the range of values between the 25th and 75th percentiles. The thin vertical lines extend to the max and min values. MUCAPE values were about 1000 j/kg lower for the Strongly forced event (labeled SF). Thus, long-lasting derechos (that are strongly forced) can tend to develop and persist in environments of rather low values of CAPE, even some that are conditionally stable.

What are the Factors that Influence Multicells?

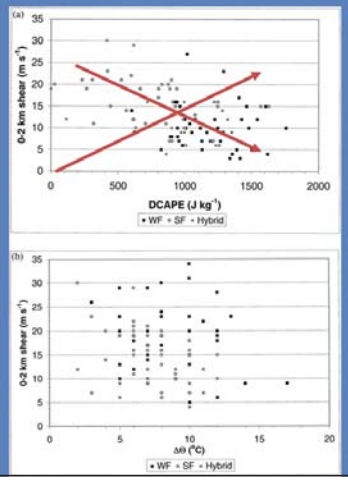
- RKW ('93)
- Focused on shear in lowest 2.5–5.0 km AGL
- "Optimal" conditions of Shear ≥ 20 ms^{-1}
- Is there a Balance?



Numerical cloud modeling studies starting in late 80s and continuing on through the 90s are termed **[box]** RKW for Rotunno, Klemp and Weisman. They sought to simulate long-lived severe squall line development. **[box]** These studies explored the storm-scale evolution in the development and maintenance of long-lived squall lines or bow echoes, with an emphasis on the effects of environmental shear. Results from RKW simulations indicated conditions become **[box]** "optimum" for sustained significant bow echo development when vertical wind shear of 20 m/s or greater occurs in the lowest 2.5–5 km above ground level (AGL). This led to a cold-pool circulation overwhelming the low-level shear and allowed the system to develop an upshear-tilted structure**[box]** . RKW is based on theory of vorticity balance between cold pool and low-level vertical wind shear. And the Rear-Inflow Jet, **[box]** if it stays elevated, helps to maintain a sloped front-to rear ascension of warm air developing over the updraft. Later research suggested that longevity and severity in multicells was not all about low-level shear. There may be some importance to deep shear above the cold pool as well.

The "Imbalance" Between Shear and Cold Pool

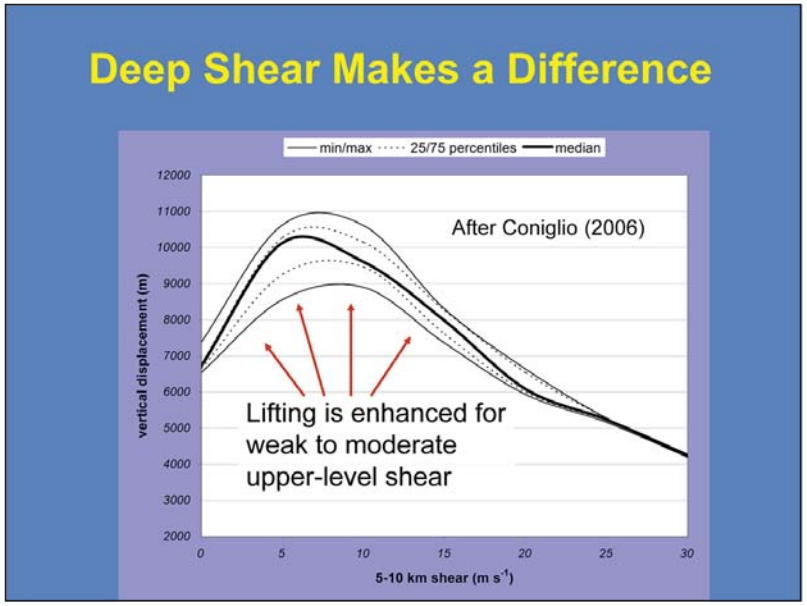
- No correlation in the data but unable to test RIJ's effects
- Is there something else to consider?



What About the Role of Deep Shear?

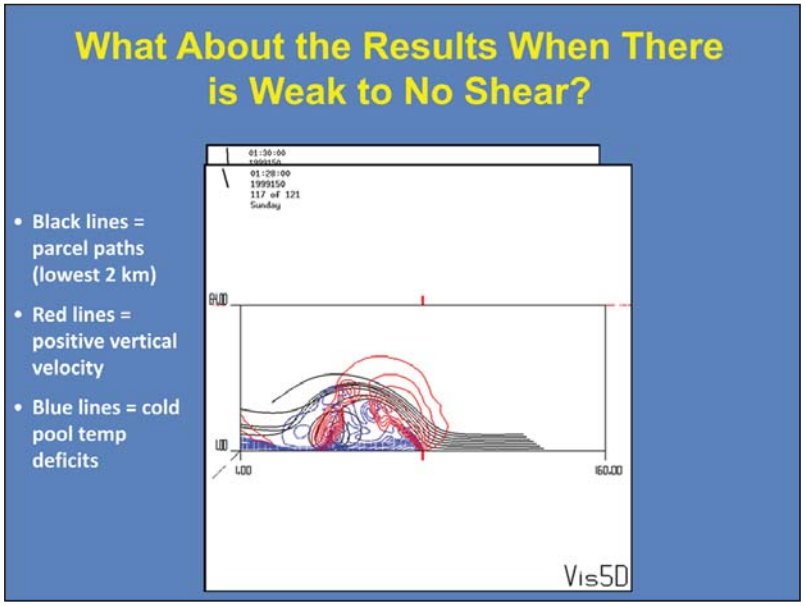
- Research suggests upper-level shear in mid-layers (5-10 km) increases vertical parcel displacements through an overturning updraft
- Leads to longer-lasting MCS archetypes
  - TS, PS types

The RKW theory found that strong low-level shear must be balanced by the cold pool to help support development of 3-D mesoscale features such as Elevated RIJs, and bookend vortices within the convective system. Evans and Doswell (2001) did not find a clear relationship between cold pool strength and low-level shear. Their data showed that DCAPE (used as a proxy for cold pool strength) and shear were not positively correlated. However, they were unable to test their observations against the assertion that the elevated RIJ was necessary to re-balance the cold pool/shear circulation (Weisman, 1993). But, the observations did provide a problem. It is true that cooler-season multicell events often do have strong low-level wind shear and weaker shear aloft, but the idealized numerical models probably most relevant to warm-season environments. Many warm-season cases have relatively weak low-level shear and significant shear above the cold pool. The RKW mechanism (cold pool/low-level shear balance) may not be the primary controlling mechanism.

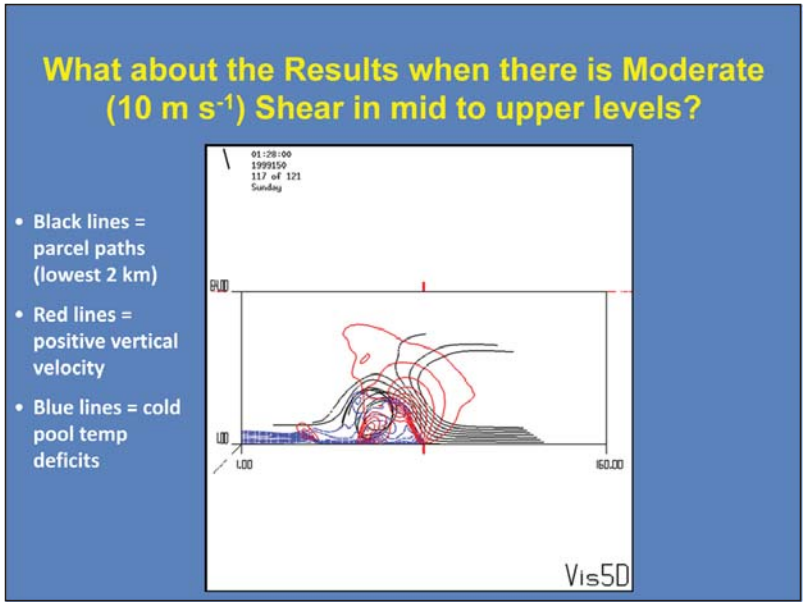


This is a graph showing numerical simulation results from Coniglio, who discovered that lifting above the cold pool in multicells was enhanced by

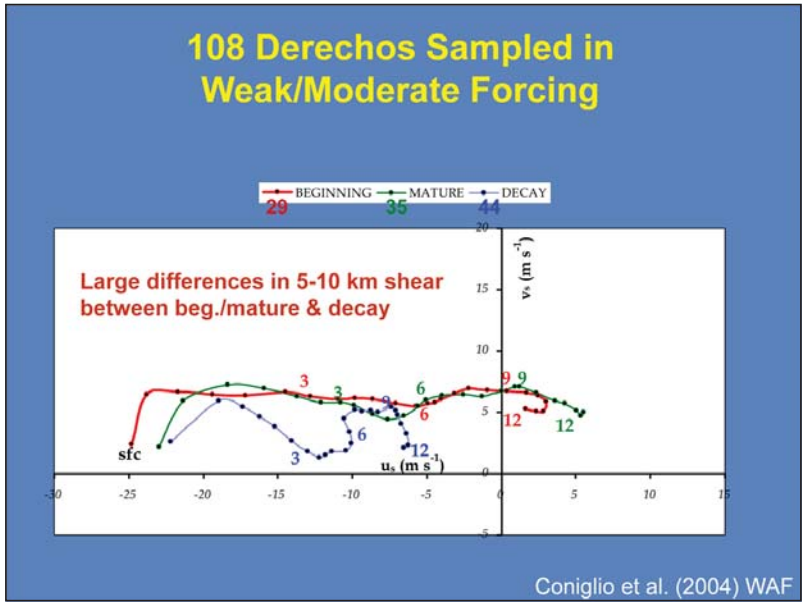
weak to moderate shear in the mid layers (5 to 10 km). This may help explain why RKW does not apply to all multicell cases.



In these Vis5D simulations from Coniglio, you can see the effects of weak midlevel shear. Very little vertical displacements occur (the black trajectory lines) results because the updraft becomes prematurely truncated as initially vertically ascending air parcels begins to experience a more slanted path rearward over the cold pool dome. The cold pool outruns the convective layer steering flow and initiating cells cannot maintain their proximity to the boundary induced lifting. The appearance on radar may be that of strong low-level reflectivity but weak reflectivity well below the equilibrium level.



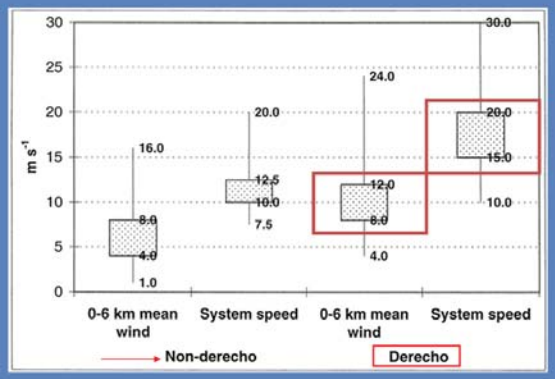
When the deep shear is moderate (on the order of 10 m/s or 20 kts), air parcels are forced up and begin to travel in the downshear direction as the interaction of the shear and updraft induce relative high pressure upshear. The airflow in this animation display slightly upshear trajectories then begin to tilt vertically with many parcels finally exiting the downdraft downshear (to the right). The updraft becomes much deeper and completes what is called an overturning cycle. The deep shear allows the updraft to remain on the boundary for longer periods of time. In addition, a larger percentage of air parcels are likely to be lofted toward the maximum available equilibrium level. This type of shear with continous overturning trajectories is found more often in TS and PS archetypes. In strong shear cases (animation not shown here) what happens is that most air parcels on average fail to reach the max potential height . However, some do make it and these may lead to isolated deep updrafts with supercell behaviors. The MCS as a whole may exhibit less slab-like lifting as the individual cells outrun the gust front more easily becase the steering layer flow is much faster and less stratiform precip is reinforcing the cold pool .



The results of the simulations were confirmed by observations. This is a composite of various hodographs at various stages of derecho development (beginning, mature, and decay). The hodographs indicate large differences in 5-10 km shear from beginning and mature stages to the decaying stage of observed derechos.

What Causes MCSs to Survive (or Die)?

- Mean winds

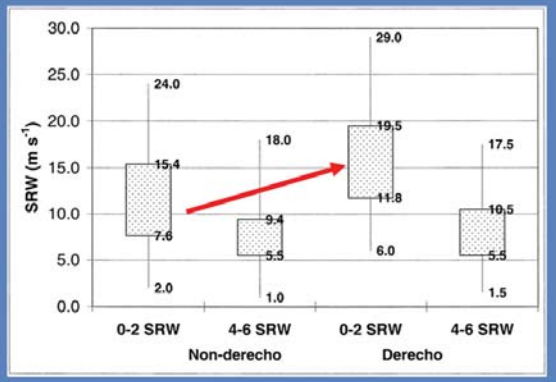


Evans and Doswell (2001)

Evans and Doswell (2001) results showed that the mean winds in the 0-6 km layer and the 0-2 km system-relative inflow were stronger because system speeds were faster for derecho events in Mesoscale Convective Systems (MCSs) as compared to MCSs that did not produce derechos.

What Causes MCSs to Survive (or Die)?

- Storm-relative flow

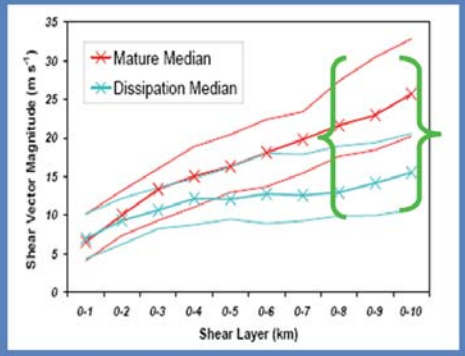


Evans and Doswell (2001)

Now here is the box and whisker plots showing the 25th to 75 percent quartiles, maxima and minima of observed values. The 0-2 km system-relative flow (labeled SRW) was enhanced in the strongest derechos likely due to faster forward speed and associated increased low-level convergence . The 4-6 km SRW did not show much difference between derecho and non-derecho events.

What Causes MCSs to Survive (or Die)?

- Upper-level shear

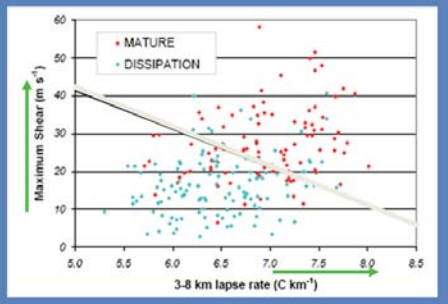


Observed derechos from Coniglio (2007)

The last thing that causes multicell systems like MCSs to survive or dissipate is upper-level shear. You can see in this graph from Coniglio that median shear vector magnitudes (lines marked by X) calculated for various depths among the mature (red) and dissipation (blue) sounding groups. The thin lines enclose the 25th and 75th percentiles of each distribution. The big spreads (click) in the median values occur as the shear layers grow deeper.

What Causes MCSs to Survive (or Die)?

- Mid-level lapse rates



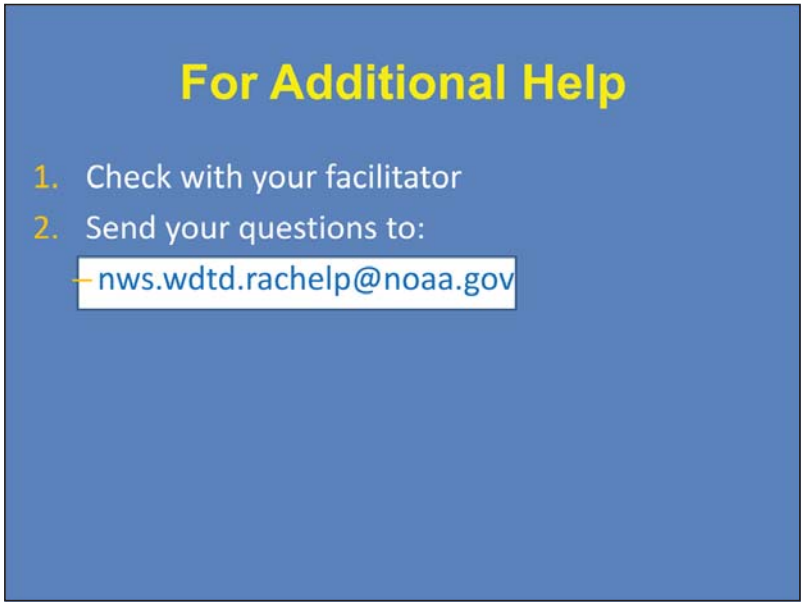
Observed derechos from Coniglio (2007)

One last influence for MCS survival: high midlevel lapse rates tend to be important to maintain buoyancy of the updrafts once they grow off the cold pool boundary. MCSs can survive in weaker lapse rates but only when there is minimal inhibition, deep moisture and stronger deep layer shear and you can see from this scatterplot of max deep shear vs 3-8 km lapse rates for the mature and dissipation soundings associated with observed derechos. The linear discrimination line separates 75% of the soundings correctly.



In summary, these are the environmental factors for influencing multicells severity and longevity:

It's still CAPE and Shear, but it's not just cold pool strength and low-level shear. It is deep shear and midlevel lapse rates that are also important for longevity of multicells. In the next lesson we'll see some of the mechanisms that affect multicell propagation.



For additional help, check with your facilitator or send your questions to the listserv e-mail address here.

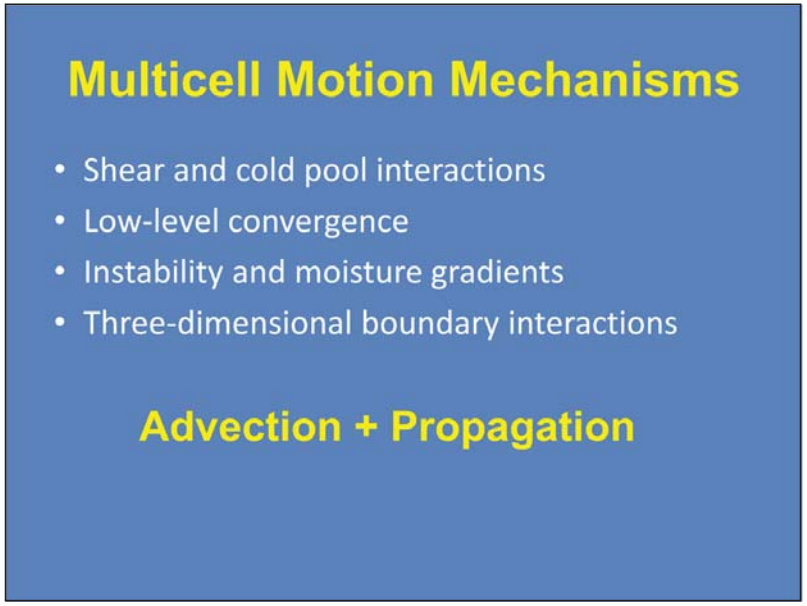


Welcome to the RAC Convective Storm Structure and Evolution lesson on Multicell Motion.

Learning Objectives

- Identify the primary vectors of motion of backbuilding and forward propagating multi-cell systems
- Identify what determines whether forward or backbuilding propagation is more likely
- Identify the role of boundary interactions on the motion of multi-cell systems
- Identify the role of instability gradients on the motion of multi-cell systems

This is the learning objective for Lesson 16.





Determining the motion of multicells is more challenging than for ordinary storms or supercells because there is usually quite a bit at work. Multicell motion is governed by shear and cold pool interactions, which

Influences the amount of low level convergence

Instability and moisture gradients can impact multicell storm motion

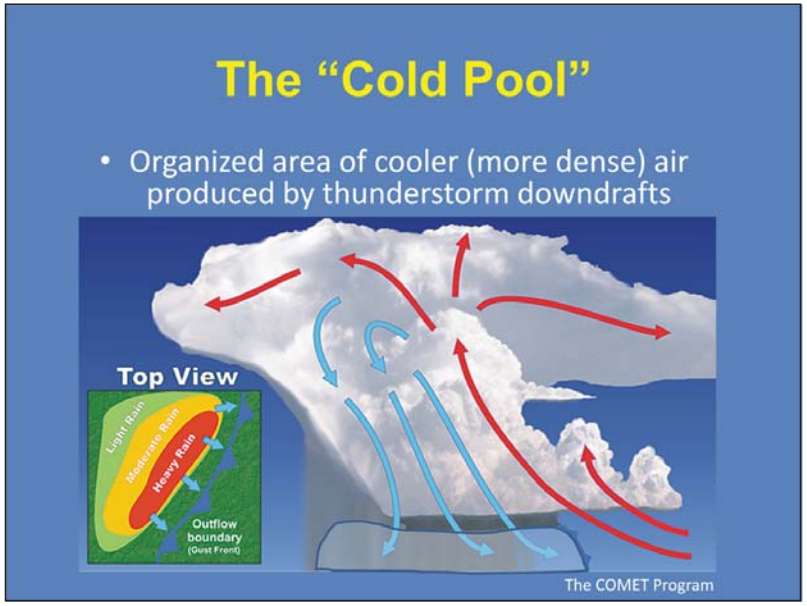
As can three-dimensional boundary interactions. The bottom line though multicell motion boils down to

Advection and propagation.

In the multicell or mesoscale convective system world Advection refers to shear and ambient winds moving and pushing the cells or MCS through the atmosphere

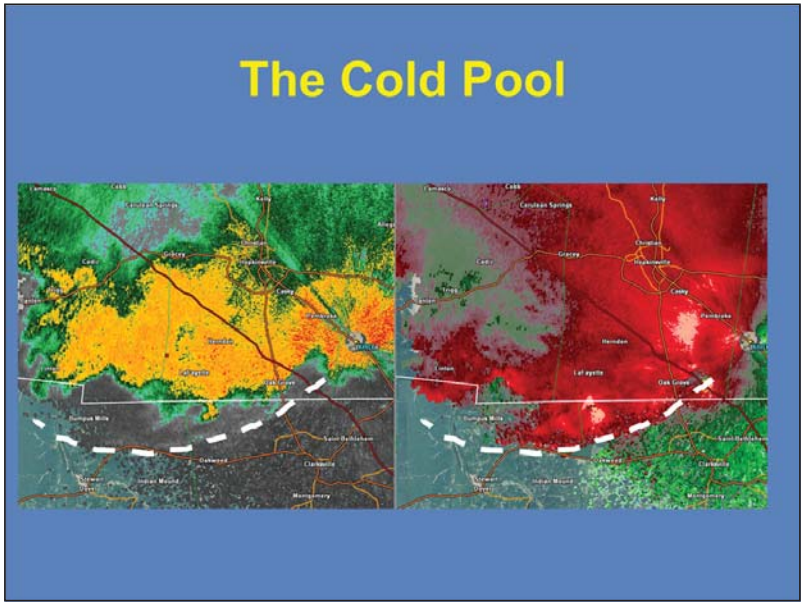
While propagation is the apparent movement of the system due to downdraft and gust front or more completely, the cold pool initiating new updrafts.

Corfidi, 2003



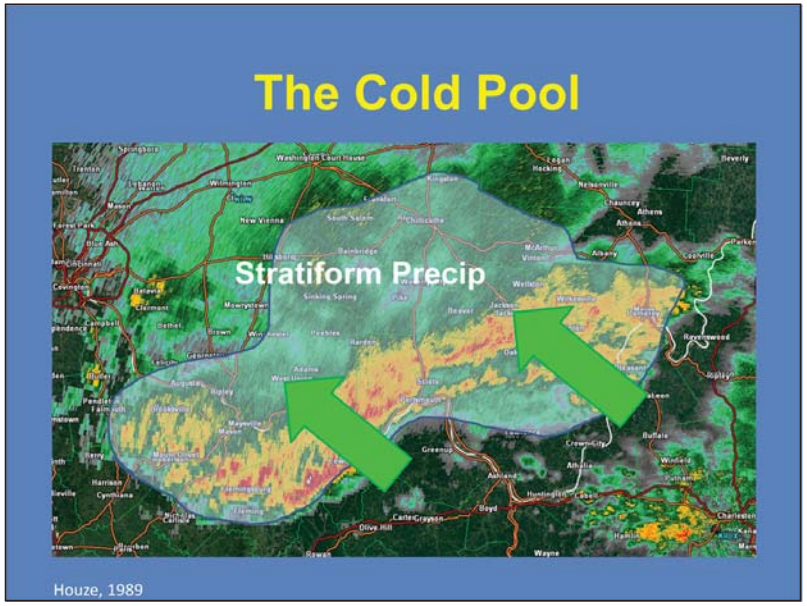
Just incase you are unfamiliar or forgot along the way, when we say “cold pool” we are just talking about the somewhat organized area of cooler, and more dense air produced by thunderstorm downdrafts.

Downdrafts from multiple cells can form a somewhat organized density current that helps drive storm and storm system motion.



On radar you can identify a cold pool, or outflow fairly easily. This image shows reflectivity on the left and storm-relative velocity on the right, so velocity from the perspective of the storm. Note that the inbounds are well ahead of the reflectivity/updraft region,

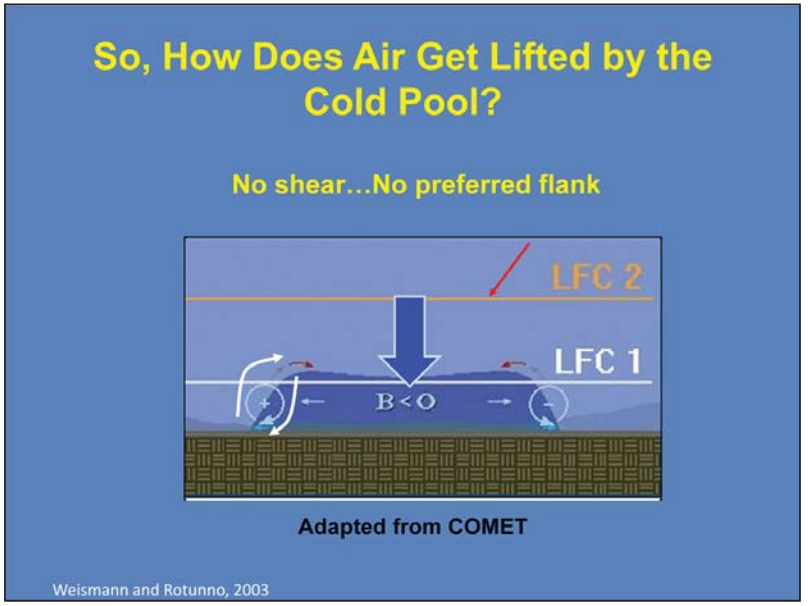
in this case the cold pool has moved out ahead of the convection.



Stratiform precipitation can also help identify the location of the cold pool.

Here the approximate location of a mesoscale cold pool is outlined

as ambient winds move warmer air over the more dense cold pool stratiform precip develops.



New storm development favors the side of a multicell storm where the ambient winds are positive relative to the orientation of the boundary.

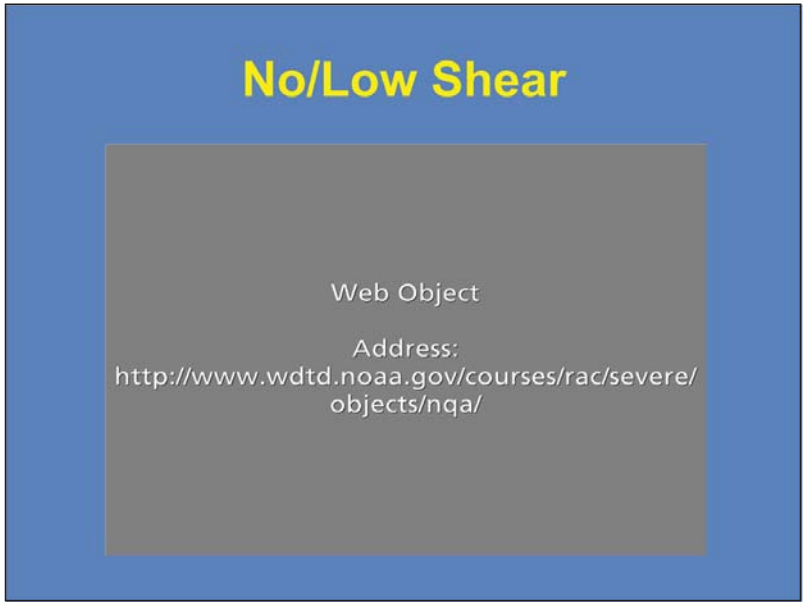
If no shear, or very minimal shear is present there will be no real preferred flank for new cell development.

A downdraft produces a new cold pool.

And the downdraft and cold pool produce low level convergence which can lead to new updrafts.

The height of the LFC, or height to which parcels will need to be lifted to develop new thunderstorms, will then determine whether or not new storm growth occurs. A lower LFC height, such as LFC1 labeled here would probably favor new storm growth with a cold pool at the depth shown.

But if it is at LFC2, new cell growth would appear unlikely in the absence of some other new forcing.



Here is an example of a low shear multi-cell system. Note how the outflow/cold pool expand almost perfectly symmetrically, in some areas new cells go up, in others they don't but there is no real preferred flank and this cluster never organizes in any appreciable way.



If we add shear things get a little more interesting. In this example we will see convection favored on the downshear side.

So in this image the cell on the right.

Cold pool winds, to the left of the cell with the negative sign, interact and converge with ambient winds to the right of the storm to favor development on that side.

On the left side the shear profiles are such that downward motion is favored on the left flank. Fast forward motion of the overall cold pool could help limit convergence on this flank, as could a convective feature like a rear-inflow-jet.

Nomenclature

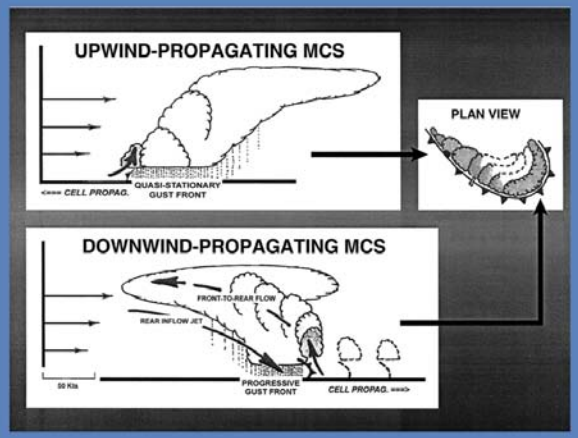
- Upwind/Upshear =
  "Backward Moving"
  Frequently Associated With Heavy Rain

- Downwind/Downshear
  "Forward Moving"
  Frequently Associated With Strong Winds

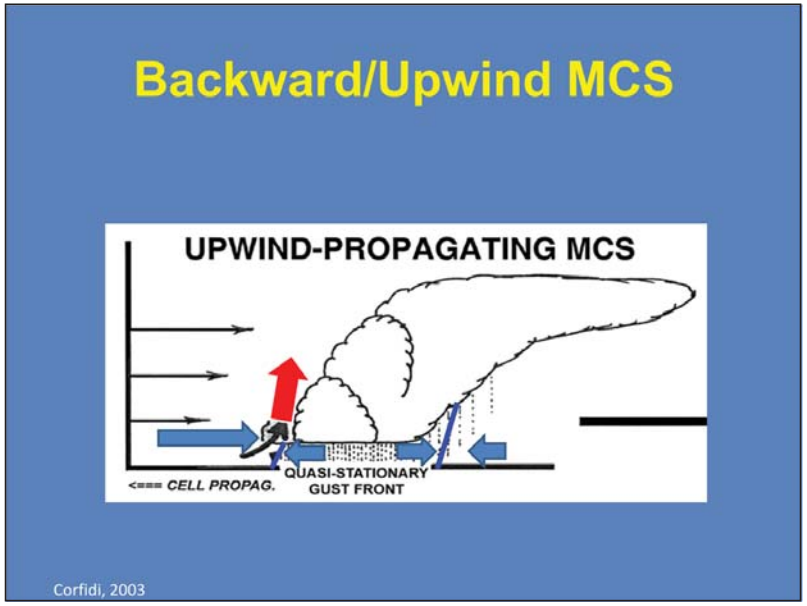
I always have a little trouble keeping MCS nomenclature straight. So just in case you do too. Upwind/Upshear systems are the type that appear "backward moving" or backward propagating. This results in slow ground-relative system motion and these systems are frequently associated with heavy rain, for more reasons than just the motion and we will get to that in a bit.

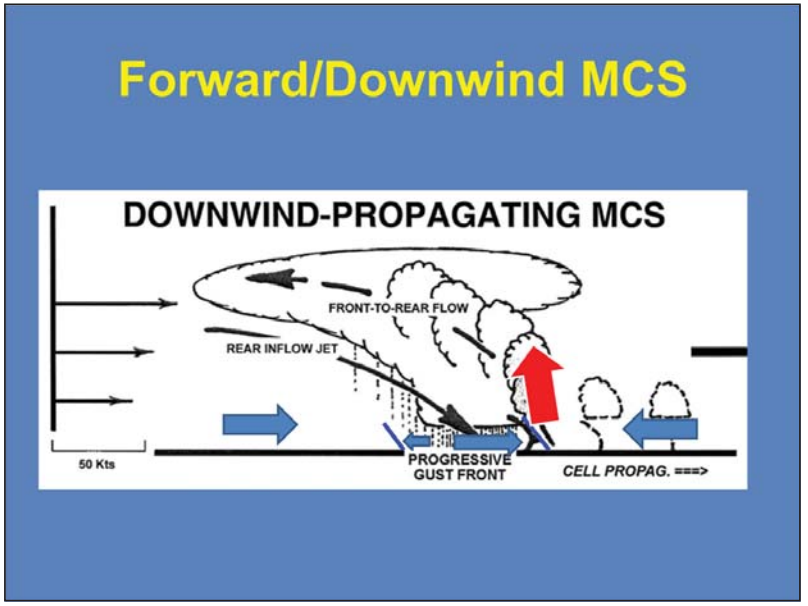
Downwind/downshear systems are forward moving, frequently associated with strong winds. The prototypical derecho or squall line would be a form of a downwind/downshear propagating multicell clusters.

Upwind/Downwind MCS Movement



Here is a side view schematic of upwind and down wind MCS movement. The upwind example our rainier, slow moving MCS, while the downwind system being our more classic squall line.





Looking more closely in an upwind propagating MCS. We have our cold pool

Roughly bounded here either side of the precipitation with outflow in both directions up against the cold pool boundary.

Shear to the left of this system might be relatively strong,

Compared with the ambient winds to the right, since its moving slowly on a quasi-staitonary cold pool. Its not surging forward against the headwinds.

This will lead to the strongest convergence on the back side of the storm.

So you have minimal advection, in a system largely moved by propagation.

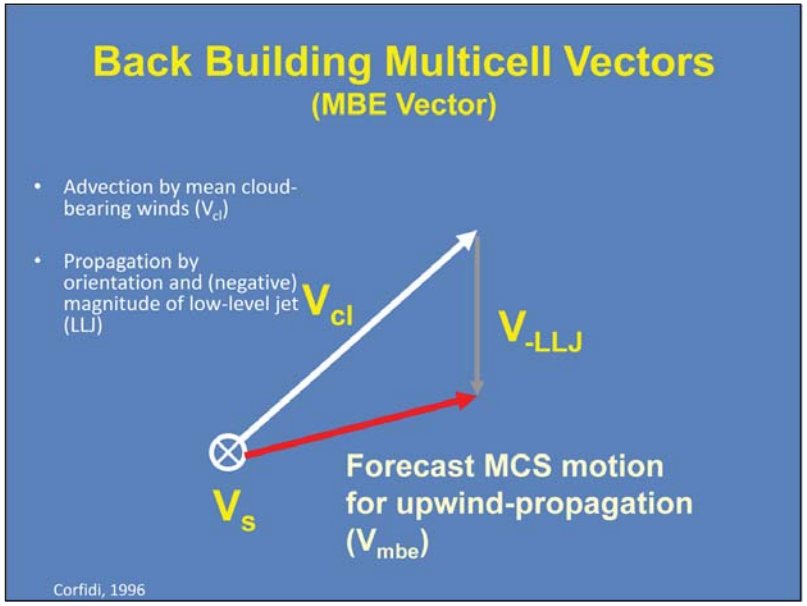
In the downwind example

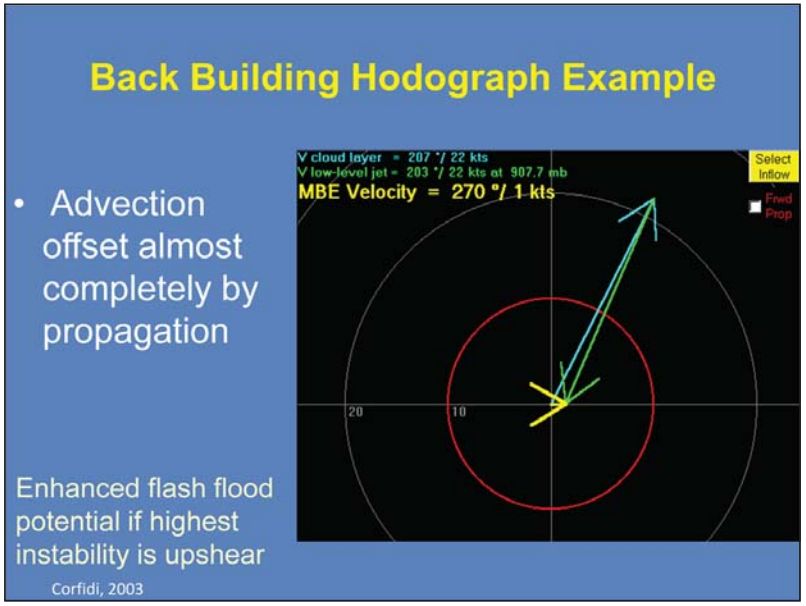
You have the cold pool boundary

And your downdrafts may be close to the same magnitude but the forward motion of the gust front gives the front facing winds more push

So the winds push harder against the oncoming winds

And new updrafts are favored on the right side. So you have the advection and propagation of this storm moving it forward.





We can estimate multicell storm motion through two different approaches. For back building multi-cell systems we can use the MBE vector, MBE stands for Meso-Beta scale Element, since that's the scale these systems usually operate on.

You first take the advection of the system by taking the speed of the the mean cloud bearing winds, usually 850 to 300 mb.

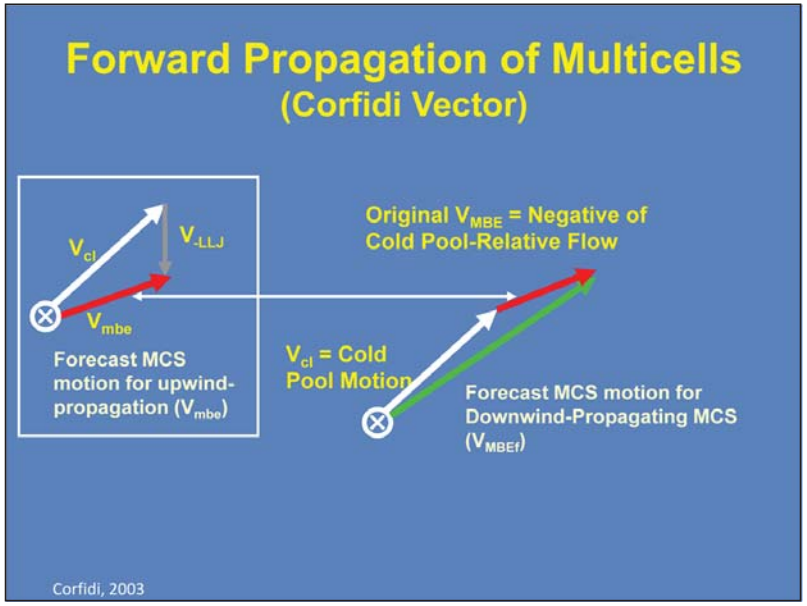
Then for your propagation you take the negative vector, or the opposite of the low level jet. Usually your 850 mb winds.

The resultant vector will give you're your forecast MCS motion for back building or upwind propagating Multicell systems. Normally this is available as either MBE Vector or Corfidi-Upwind in AWIPS.

Heres an MBE vector example, in this case the mean cloud layer winds are 207 at 22 knots, and the low level jet is 203 at 22 knots,

So the forward advection of the system is almost completely offset by propagation.

The flash flood potential is higher if the highest instablity is upshear in the direction the system is working towards.

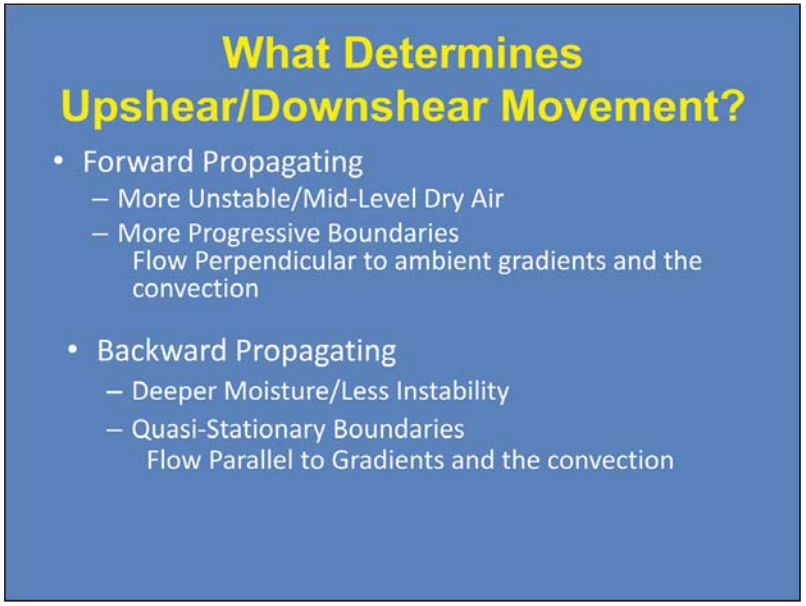


For forward moving, or downshear propagating systems you can use this technique.

We still use the original MBE vector

But we add it to the cold pool motion, which is roughly the mean winds of the cloud bearing layer. So it captures the same propagation as the original MBE vector but adds it to the larger scale advection, since propagation and advection are added in a forward moving MCS.

This vector is usually available as Corfidi-Downwind or Downwind MCS in AWIPS.



So what determines whether your system will follow upshear or downshear movement?

Forward propagating systems tend to occur in more unstable air, with drier mid level air that makes the cold pool stronger. They also tend to occur along more progressive boundaries which are a result of the fact the larger scale flow is perpendicular to ambient gradients and existing convection.

Backward propagating systems tend to occur in deeper moisture, with less overall instability, another reason they are a big flash flood concerns. They tend to be located along quasi-stationary boundaries due to the fact the flow is parallel to gradients and the convection.

So it's a combination of shear and instability, which of course are the processes producing advection and propagation.

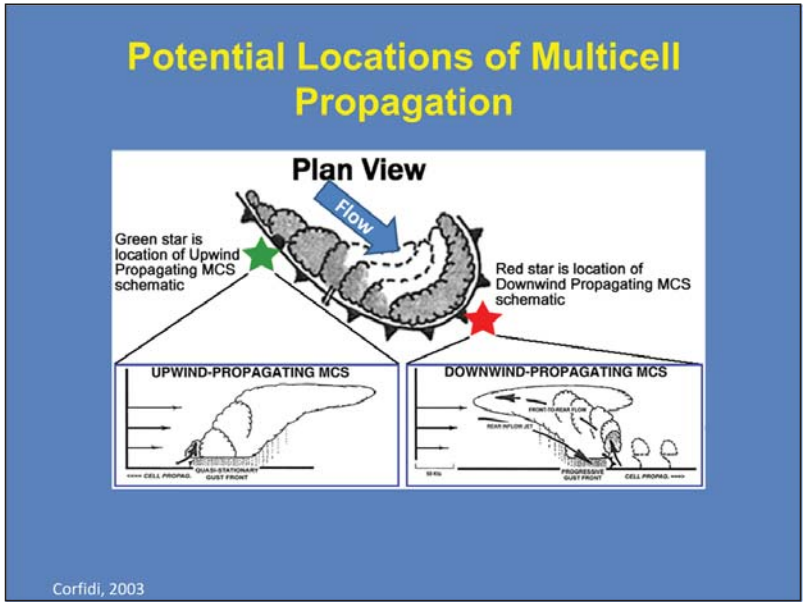


Heres an example of a pretty solid forward propagating MCS.

Notice the sharp reflectivity gradient along the leading edge and stratiform precipitation behind the system. This complex produced several wind damage reports as it swept across Kentucky and West Virginia.



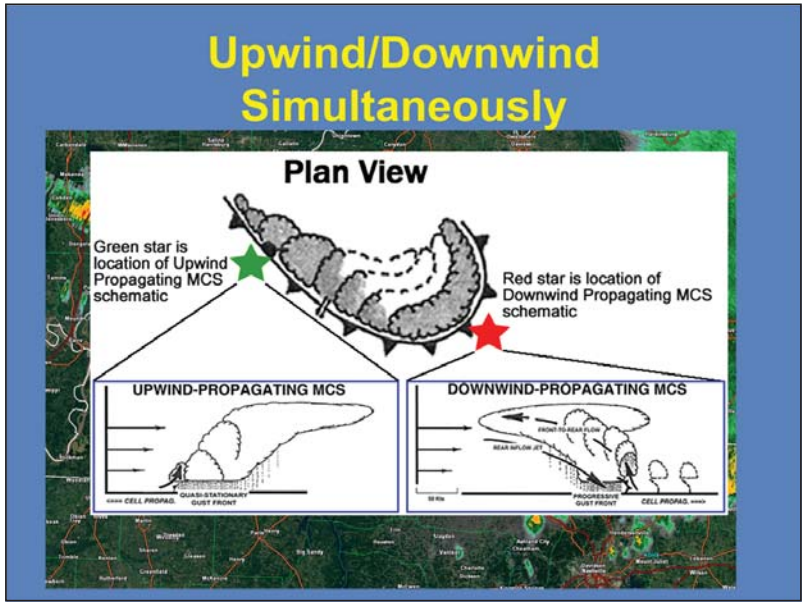
Here is a backward propagating MCS, the new cells developing on the left side of the complex. The mean flow is from the west to northwest and new cells are developing into the mean wind.



Can both occur at the same time?

Of course, it wouldn't be operational meteorology if it were cut and dried! In this case

With ambient flow out of the northwest the green star represents the location of upwind propagation, where the MCS and associated cold pool front is parallel to the flow and the Red star is where downwind propagation is occurring.



Here is an example of such.

The white dotted line represents the approximate location of the cold pool front.

The shear is from the northwest.

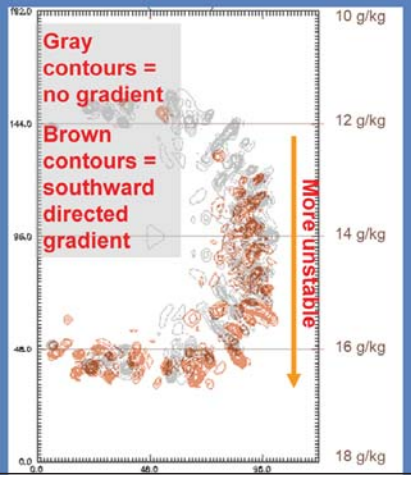
Convection here is forming downwind, or forward propagating

And here it is forming upwind, or backward propagating.

Looks a lot like this.

Theta-E Gradients

- Richardson's (1999) results showed instability gradients have an effect on the propagation of a multicell line



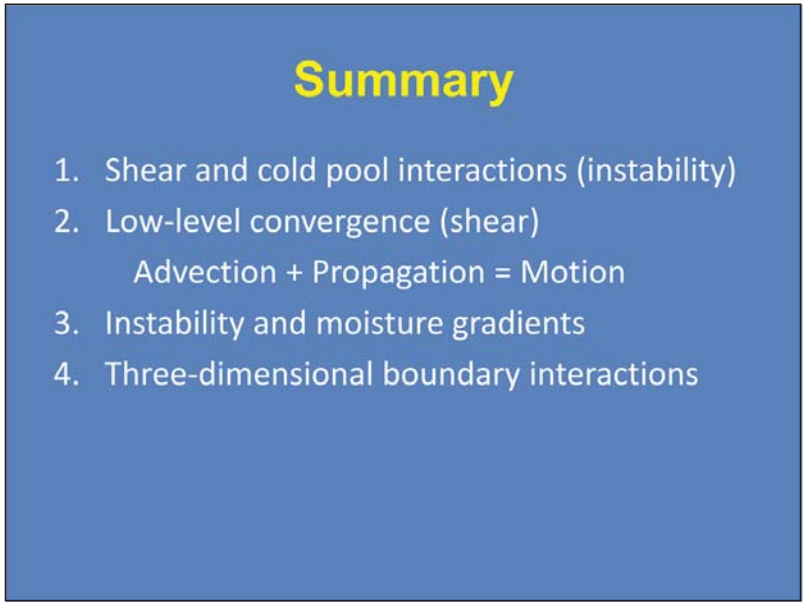
Theta-e gradients, thermal or moisture boundaries, or both can cause changes in local convergence that can influence propagation, and overall storm motion as well. Development tends to head towards more unstable regions due to lower LFCs and stronger resultant downdrafts reinforcing the cold pool.

Boundary Interactions

- Cold pool interacting with a boundary may initiate new convection



Local pre-existing boundaries can also impact storm motion with the propagation vector tending towards the location of the boundary.



To summarize there are four key drivers to multi-cell motion.

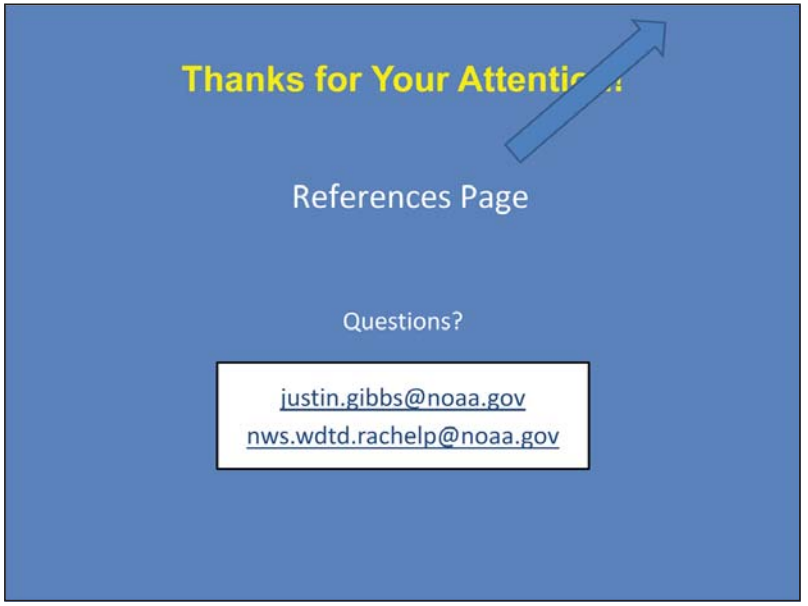
Shear and cold pool interactions, which can be estimated by available instability

And low level convergence, which can be estimated by ambient wind shear.

Advection and propagation combined form the overall motion

More difficult to predict, but still detectable through mesoscale analysis, instability and moisture gradients can influence multicell motion, with development favoring more unstable regions.

And pre-existing boundaries can also influence storm motion with cold pool interaction along that boundary helping initiate new convection.

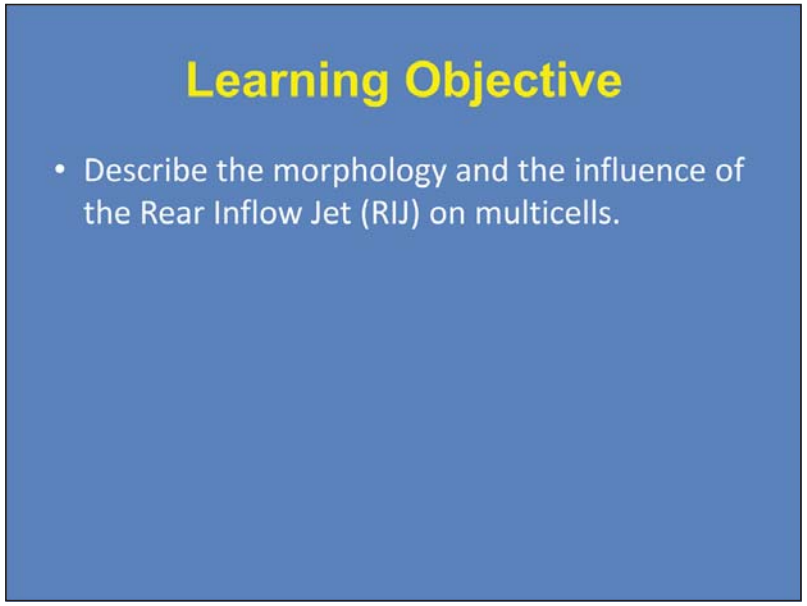


If you have passed the quiz, then you have successfully completed this lesson. If you have any questions, please contact us using any of the e-mail addresses listed on the bottom of the slide.

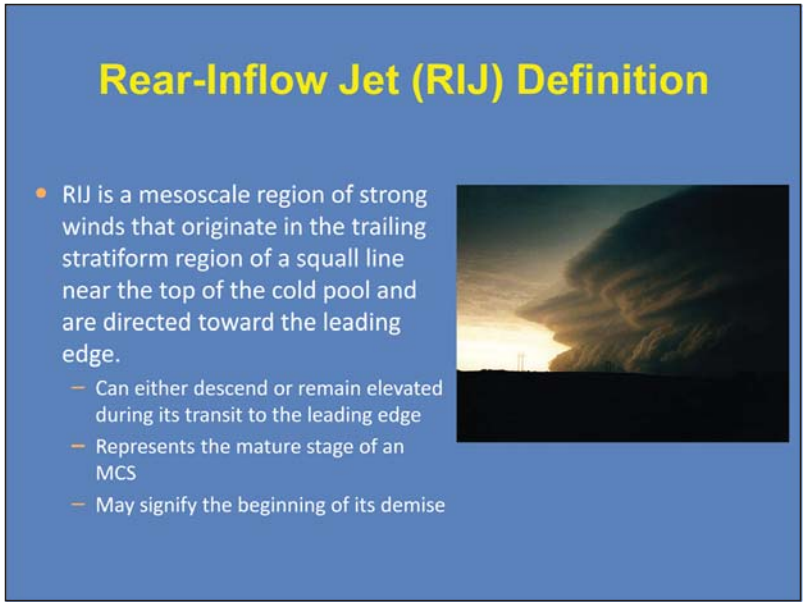
Check out the references page to links for papers on the topic that can deepen your understanding.

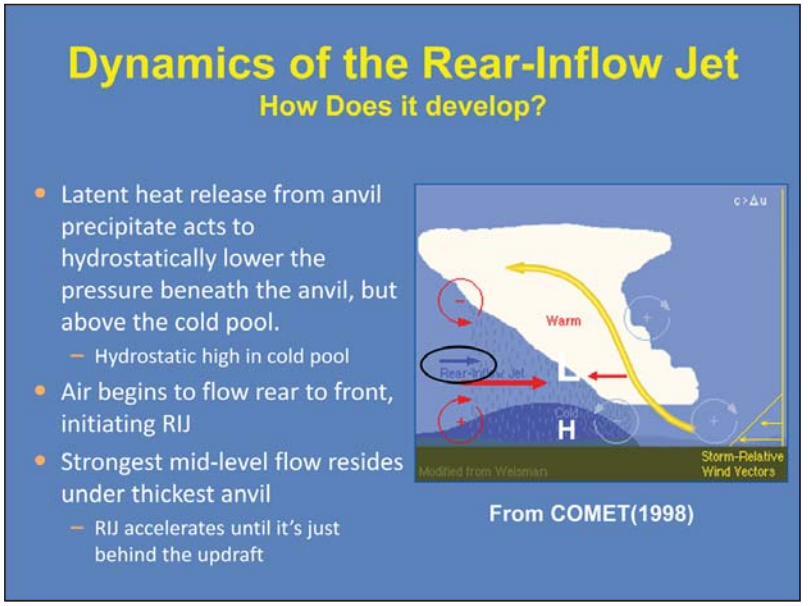


Welcome to the RAC Convective Storm Structure and Evolution lesson on Rear-Inflow Jets in Multicells. The morphology of the **rear-inflow jet (RIJ)** helps explain future evolution of a multicell and its potential to produce severe weather. This lesson describes the causes of a rear-inflow jet, its dynamics, favored environments, and the different manifestations of rear-inflow jets which impact the potential for severe weather.



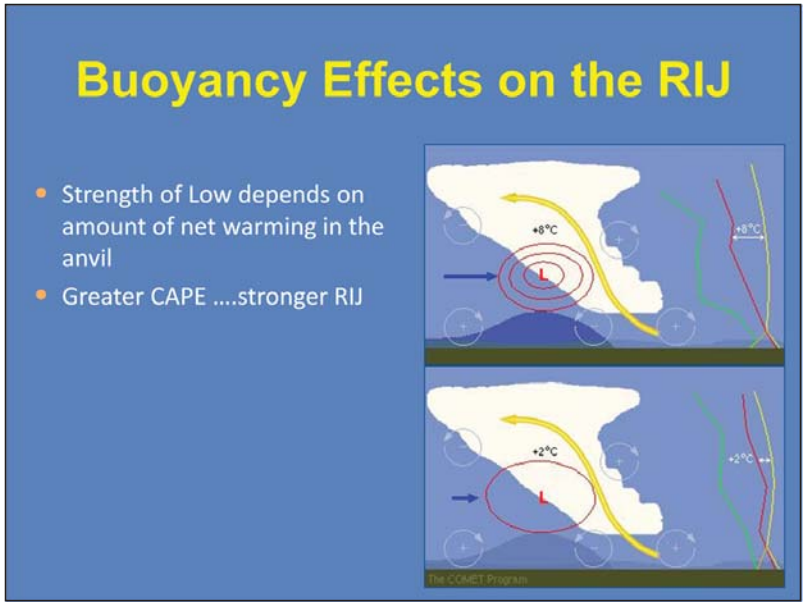
The learning objective for this less is describe the morphology and the influence of the Rear Inflow Jet (RIJ) on multicells.

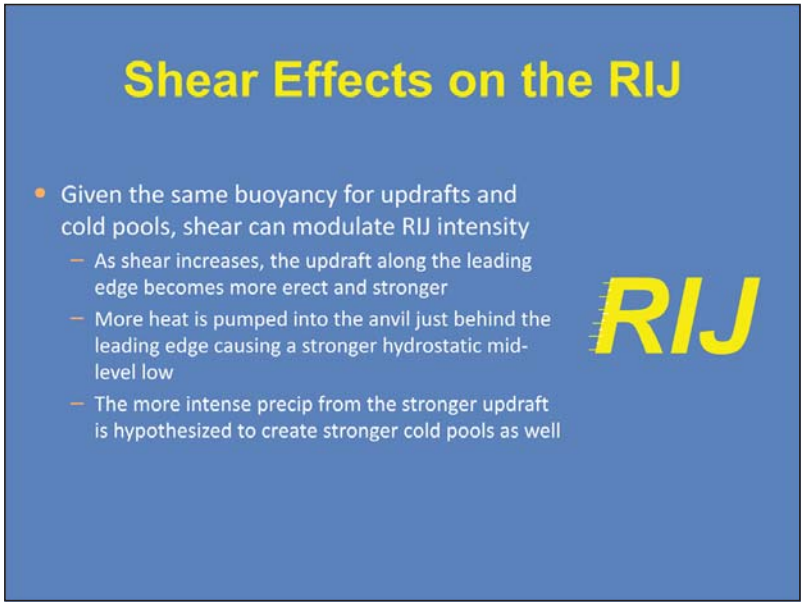




An important component of many **mesoscale convective systems (MCSs),** particularly the linear type, is a rear-inflow jet (RIJ). The RIJ is a mesoscale region of strong winds that originate in the trailing stratiform rainfall region of a squall line near the top of the cold pool and are directed toward the leading edge**. The RIJ can either descend or remain elevated during its transit to the leading edge.** It represents the mature stage of an MCS and may also signify the beginning of its demise. However, a large number of squall lines continue to show significant longevity and severity after the RIJ forms.

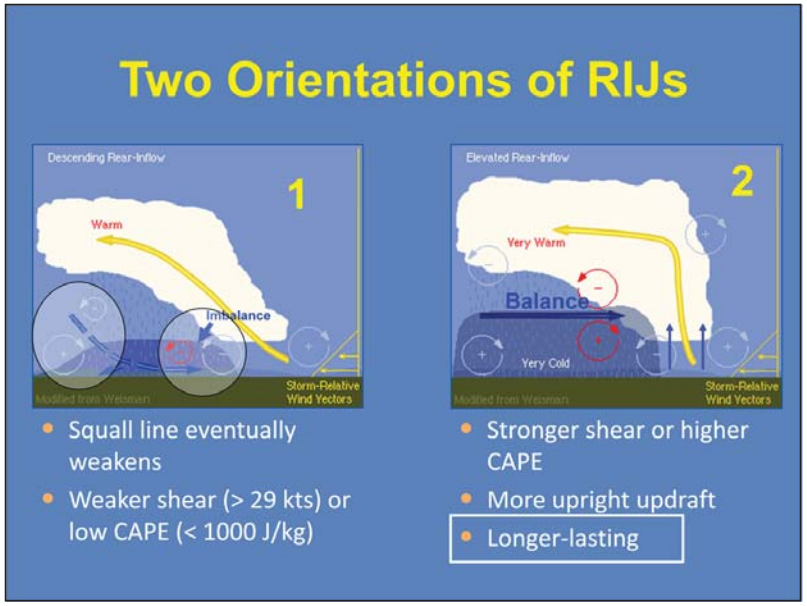
To explain the dynamics of the RIJ, we will start with the mature squall line schematic shown here. As a squall line matures, high-level anvil material begins to stream from the leading edge into the rear side of the squall line (represented by the yellow trajectory). Loaded with small- and medium-sized hydrometeors that have not fallen out of the leading edge, the anvil begins to precipitate, resulting in the region of trailing stratiform precipitation. The anvil material is also warming the upper-troposphere through latent heat released in the updraft along the leading edge. This heat acts to hydrostatically lower the pressure beneath the anvil but above the cold pool (marked by an 'L'). A hydrostatic high still exists near ground-level in the cold pool, dominating any pressure drop caused by anvil material aloft. In terms of pressure dynamics, the anvil-induced low induces air to laterally flow in from both the front and rear sides of the squall line (arrows to the right and left of 'L'). Air begins to flow rear to front, initiating the RIJ . Presumably, the strongest mid-level flow resides underneath the thickest part of the anvil just behind the deep updraft forced along the leading edge of the squall line. Therefore, the RIJ accelerates until it is just behind the updraft. We next discuss the strength of the acceleration and the factors that govern the slope of the RIJ.

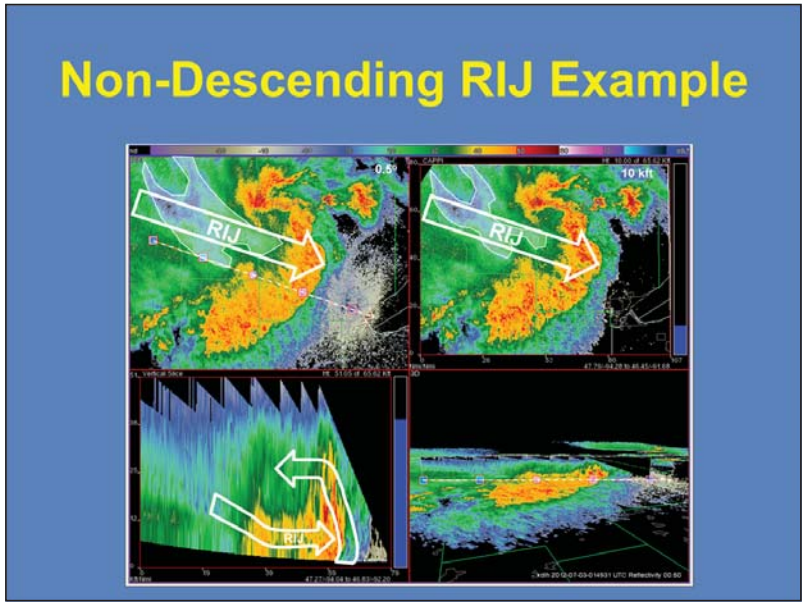




Let's talk about buoyancy effects on the RIJ. The strength of the low underneath the anvil depends on the intensity of the net warming in the anvil. Looking at this schematic, the squall line updraft with the greatest positive temperature excess is the one utilizing the greatest CAPE. Note the hypothetical sounding profile and the temperature excess of the updraft parcel to the environment. The result of higher CAPE is usually a stronger RIJ . Typically in large CAPE environments, lapse rates from the surface to mid-levels tend to be larger, promoting stronger cold pools. **From a vorticity argument, a stronger cold pool circulation to the rear of the squall line works with a more buoyant anvil aloft to generate strong mid-level horizontal inflow that forces the RIJ from the rear of the squall line.**

Given the same buoyancy for updrafts and cold pools, shear can modulate the intensity of the RIJ. According to numerical simulations, as shear increases, the updraft along the leading edge becomes more erect and stronger. More heat is pumped into the anvil just behind the leading edge causing a stronger hydrostatic low in the midlevels. The more intense precipitation from the stronger updraft is hypothesized to create a stronger cold pool as well.





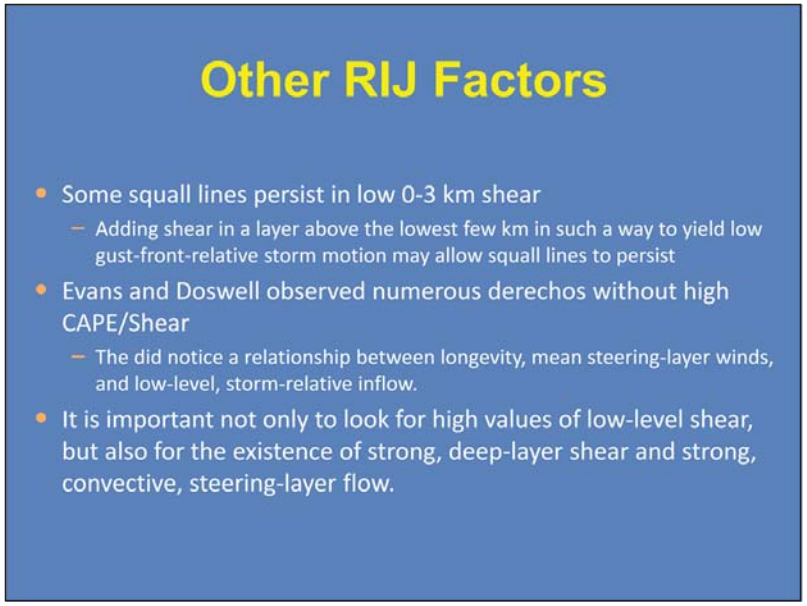
Weisman found that the longevity of the squall line may depend on the rear-to-front slope of the RIJ. Although there may be multiple slopes to the RIJs, there are two extremes: A descending RIJ and a non-descending RIJ.

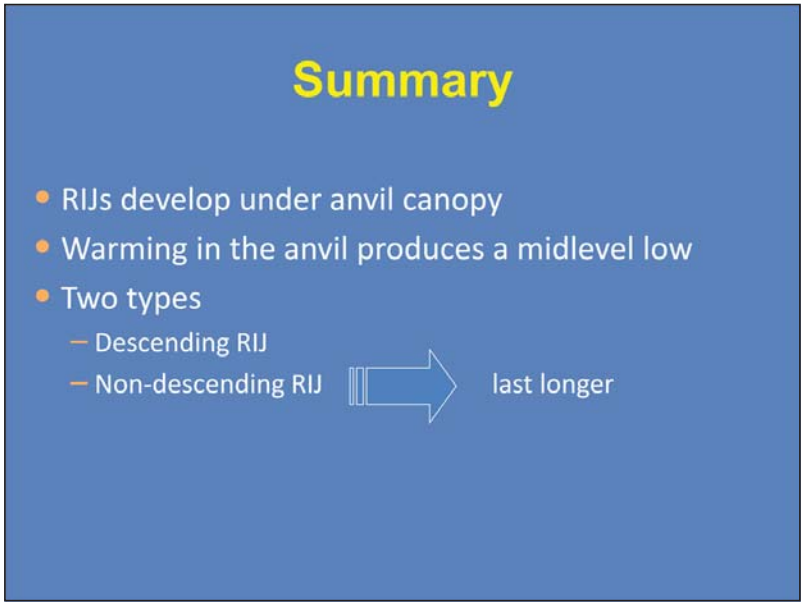
A descending RIJ occurs when the vorticity generated just underneath the ascending front-to-rear updraft is weaker than the vorticity generated of the opposite sign on the rear edge of the cold dome. The imbalance between the two circulations can be seen to help force the RIJ downward towards the ground prior to reaching the leading edge of the gust front. The RIJ then reinforces the vorticity along the leading edge increasing the imbalance between the cold pool and environmental vorticity. The squall line is theorized to become increasingly sloped rearward and thus weakening. According to simulations by Weisman (1992), this situation occurs with weakening shear (less than 15 m/s, or 29 kts, over the lowest several km) or if the environmental CAPE is less than 1000 J/kg.

For a non-descending RIJ in Figure 2, as CAPE and/or shear increases , the vorticity underneath the rearward expanding anvil becomes much larger due to increased buoyancy. The counter-rotating vorticity along the back edge of the cold dome does not increase as much. This situation results in the increased buoyancy-induced vorticity under the anvil matching the cold dome vorticity to invoke a more horizontally oriented RIJ .

This non-descending RIJ progresses towards the leading edge of the cold pool with a horizontal vorticity structure that interferes with the spreading cold pool vorticity near the gust front. Thus, the strength of the gust front vorticity decreases, becoming more balanced with the environment, and the updraft retains an upright nature. **Squall lines with a non-descending RIJ tended to live longer than their descending RIJ counterparts.**

Here is an example of a non-descending Rear Inflow Jet as seen using the AWIPS Four Dimensional Stormcell Investigator tool. In this reflectivity display, notice the weak echo channel in the top two panels which reveals the RIJ location, the upright nature of the updraft as revealed by the cross section, and how the gust front is not easily evident because it hugs close to the strong reflectivity gradient. In this velocity display, notice in the vertical cross section how the RIJ doesn't descend all the way to the ground.





The role of a non-descending RIJ in squall line longevity put forth by Weisman (1992) may not adequately explain the longevity of some severe squall lines in environments exhibiting low values of 0-3 km shear. Other numerical experiments provide evidence that adding shear in a layer above the lowest few km in such a way to yield low gust front-relative storm motion may allow squall lines to persist longer than predicted by shear/cold pool balance theory. In addition, strong synoptic-scale mid-level winds may boost the initiation time and strength of the RIJ. An example would be a cold-season, pre-frontal squall line in the warm sector of a surface extra-tropical cyclone.

Evans and Doswell (2001) observed numerous cases of derechos without high values of either shear or buoyancy. They did notice a relationship between longevity, mean steering-layer winds, and low-level, storm-relative inflow. The latter relationship is likely due to the fact that derechos move quickly. In addition, strong RIJs may be the result of dynamics beyond that of balancing anvil-level buoyancy with cold pool strength. For example, small amounts of CAPE are sufficient to vertically mix strong, synoptic-scale mid-level winds down to the surface yielding a strong RIJ-like structure.

Therefore, it is important not only to look for high values of low-level shear, but also for the existence of strong, deep-layer shear and strong, convective, steering-layer flow.

As is often the case, the parameter space in which long-lived multicell squall lines are observed is often much larger than simulations suggest. The RIJ strongly influences the longevity of MCSs, but there are several other very important environmental and storm scale features that modulate how long an MCS will survive, as well as how severe the MCS will become.

In summary, rear inflow jets develop underneath a large anvil canopy as a midlevel low forms in response to upper-level warming within the anvil. The more CAPE that gets pumped into the anvil, the stronger the midlevel low becomes.

Two types of RIJs have been documented: The descending RIJ, which typically evolves as the cold pool forcing exceeds that of the buoyant anvil. The descent of the RIJ enhances the forward motion of the cold pool forcing an increasing tilt to the convective updrafts along the leading edge. And, a non-descending RIJ, which forms in an environment of increasing vertical wind shear. The anvil is typically more buoyant due to stronger updrafts. The RIJ forcing is balanced between the cold pool and anvil. Non-descending RIJs help to restrain the advance of the cold pool relative to the motion of the system as a whole, and a deeper outflow boundary is the result. Squall lines associated with nondescending RIJs typically survive for longer periods.

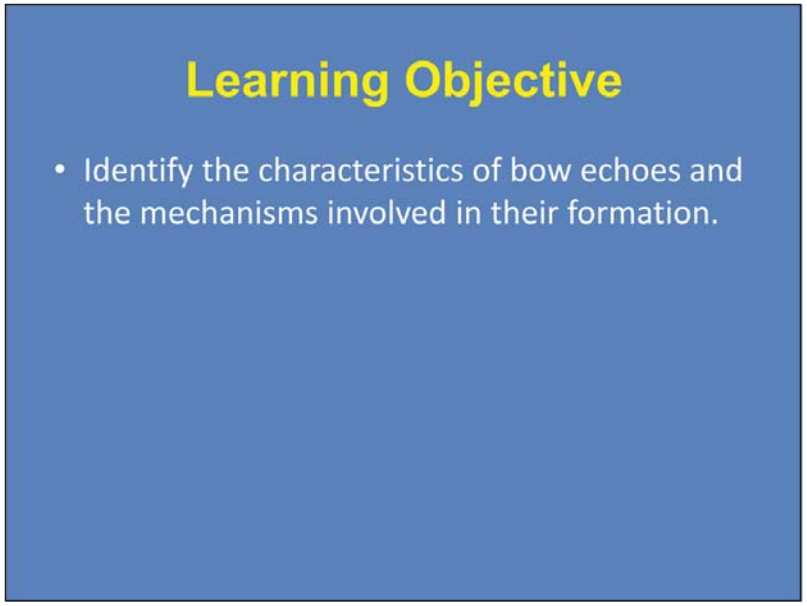
For Additional Help

1. Check with your facilitator (typically your SOO)
2. Send your questions to:
   nws.wdtd.rachelp@noaa.gov

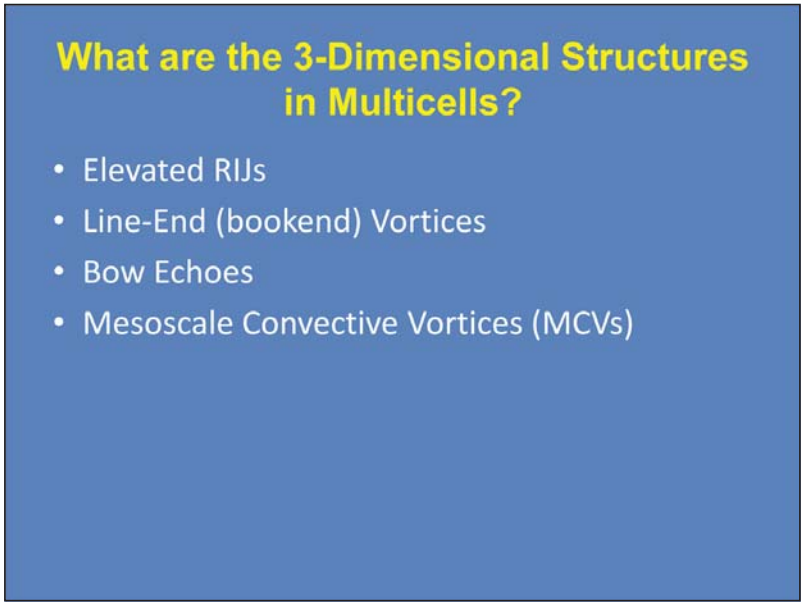
For additional help, check with your facilitator (typically your SOO) or send your questions to the listserv e-mail address here.



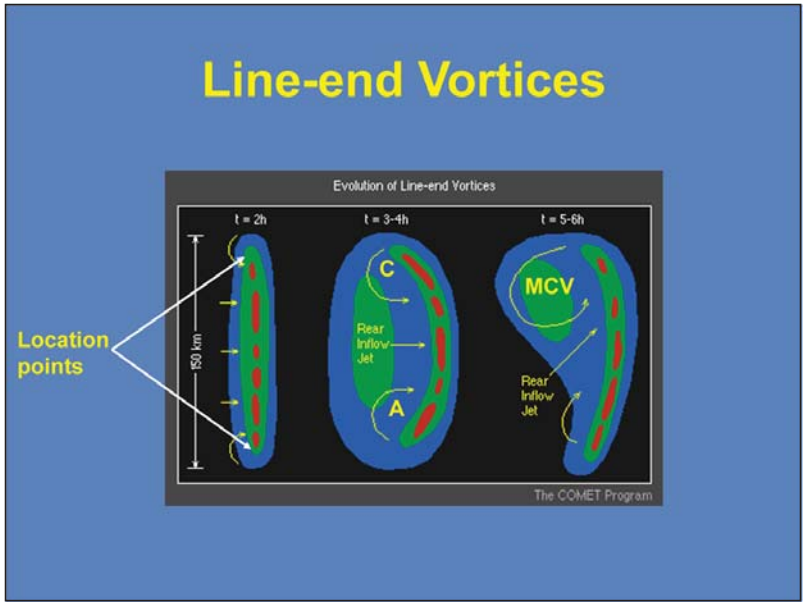
Welcome to the RAC Convective Storm Structure and Evolution lesson Line-End Vortices and Bow Echoes.

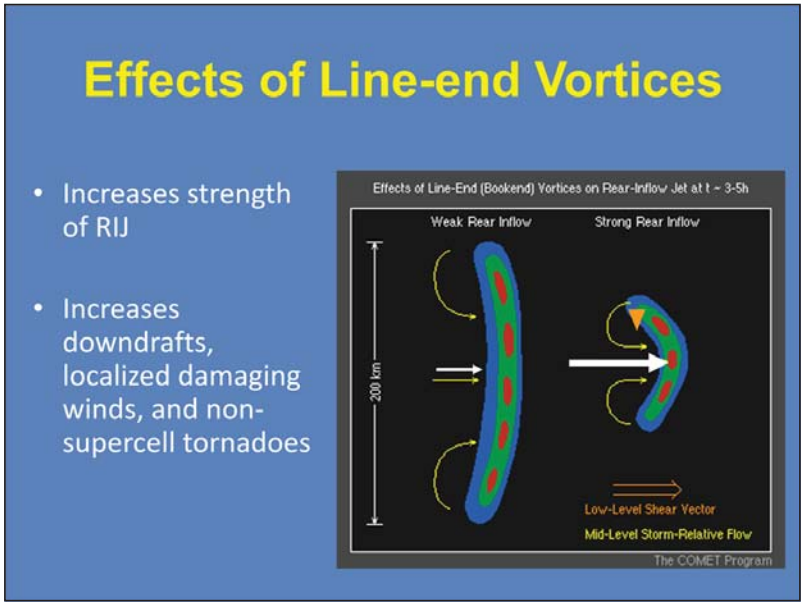


This is the learning objective for Lesson 18.



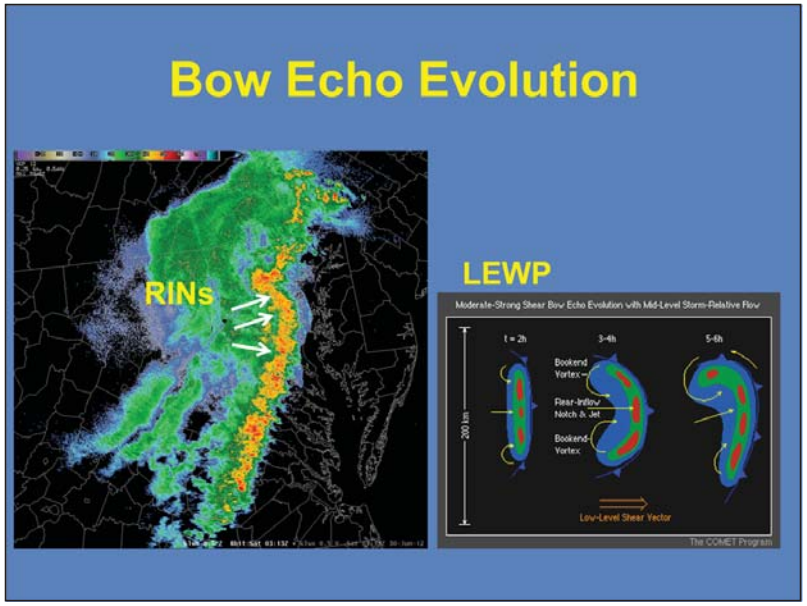
Well organized multicells exhibit numerous three-dimensional structures like elevated RIJs (previously discussed), line-end vortices (also called bookend vortices or mesovortices), bow echoes, and mesoscale convective vortices (known as MCVs). Line-end vortices and bow echoes are responsible for a disproportionately large amount of severe weather with multicells and therefore, you should know how to recognize them.





Line-end vortices by definition, evolve at the end of the line or at breaks within the line. The development of these features can alter the subsequent evolution of the system. In numerical simulations presented in the COMET *MCS module* (COMET, 1999), line-end vortices typically developed between 2-4 hrs into the lifetime of the convective system, just behind the zone of most active convection. The vortex at the northern end of the system had cyclonic rotation, while the vortex at the southern end of the system rotated anticyclonically. In the simulations, and in reality what happens after 5 or 6 hours in a mature squall line is that the asymmetrical shape develops because the Coriolis Force acts to increase convergence in the midlevels which helps to strengthen the northern cyclonic vortex and weaken the anticyclonic vortex. The dominant cyclonic vortex can last well beyond the lifetime of the originating convective system and grow upscale to become an MCV. These are the basics. Now, let's touch on some of the factors involved in line-end vortex intensity and associated severe weather development.

Since the line-end vortices develop within the downdraft portion of the squall line, they can enhance the strength of the RIJ between the vorticies, and produce severe weather such as tornadoes. Weisman (1992) found that as a squall line evolves to a bow echo, the smaller the distance between the line-end vortices produced greater enhancement to the midlevel flow and RIJ. Non-supercell tornadic development is also enhanced in the cyclonic vortex by downward tilting of positive vorticity by the downdraft in the back end of the bow echo. Tornadoes often form just to the left of the maximum wind of the RIJ. Note that bow echoes can also contribute to tornadic development by upward tilting of negative vorticity at the top of the cold pool by the storm's updraft. We'll see an example of that a bit later.

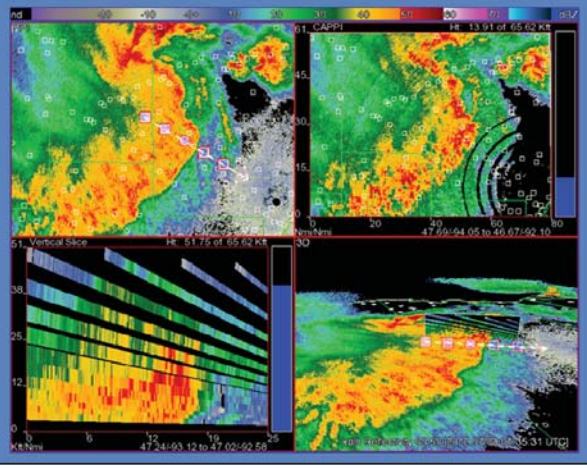




In terms of bow echoes, they certainly are a special structure unique to multicell systems and thus, their structures have been treated with a considerable amount of research and numerical modeling over the past 30 years. Fujita in 1978 was one of the first to depict a conceptual model of a bow echo morphology. In fact, Fujita first coined the phrase “bow echo” to describe the radar depiction of the accelerating concave-shaped squall line system of convective cells noted for producing long swaths of strong, surface winds. He found that the initial cells elongated symmetrically with cyclonic and anticyclonic rotating heads and then eventually became a comma-shaped echo after several hours. Bow echoes often occur from either isolated lines or as a series of bowing segments within larger convective systems. Often, you'll see a bulge in a thunderstorm which produces a wave-shaped “kink” in the line. This radar signature has been referred to as a LEWP. The potential for strong outflow and damaging straight-line winds increases near the bulge, which often resembles a bow echo. Another primary characteristic of bow echoes noted in simulations and in radar imagery is the Rear-Inflow Notch (RIN). The RIN is located well behind the core of the bow and it can signify the location of a strong RIJ. RINs have been frequently observed along the trailing edge of each individual bowing segment and signify the region of evaporatively-cooled lower Theta-E air being channeled toward the leading edge of the bow echo.

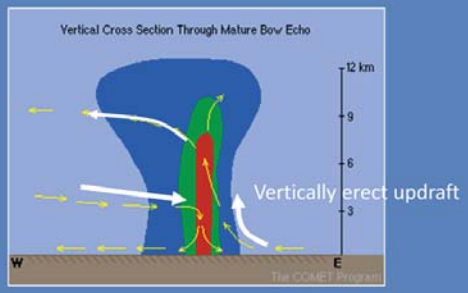
This is a loop of 0.5 deg reflectivity, SRM, ZDR, and KDP from KICT of a big squall line in southern KS/northern OK on May 30, 2012. Note there were several RINs associated with the larger RIJ structure just behind the apex of the bow that emerged as the multicell bowed out while approaching the radar. Toggle through the different products to gain a familiarization with storm structures of a Bow Echo.

Example of Rear Inflow Notches



Here's another example using FSI Z cross-section with great depiction of at least two distinct bow echo segments in the DLH CWA. Note multiple RINs observed within the larger RIJ structure.

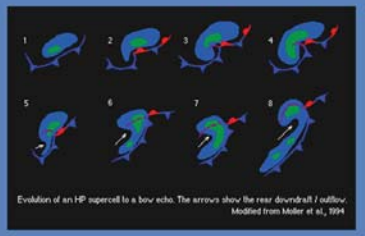
Bow Echo Cross Section



Vertical cross-sections in the core of mature bow echo simulations revealed a strong, vertically erect updraft at the leading edge of the system, a strong elevated RIJ impinging behind the updraft region before descending rapidly to the surface, and a front to rear flow near the top of the updraft which went back into the anvil and fed the trailing stratiform precip region. In reality, we often see slight modifications to this structure.

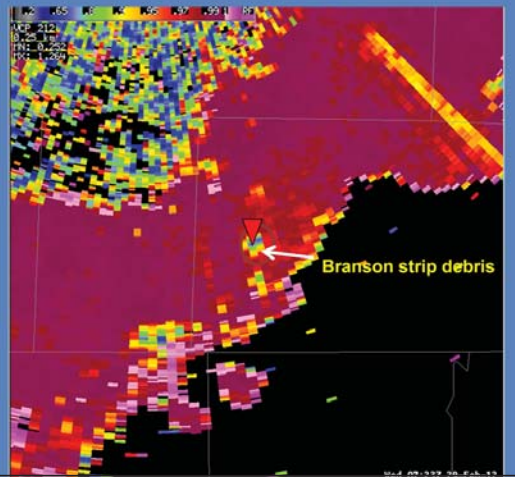
Bow Echo Transition

- Supercells to Bow Echoes
- Supercells within Bow Echoes

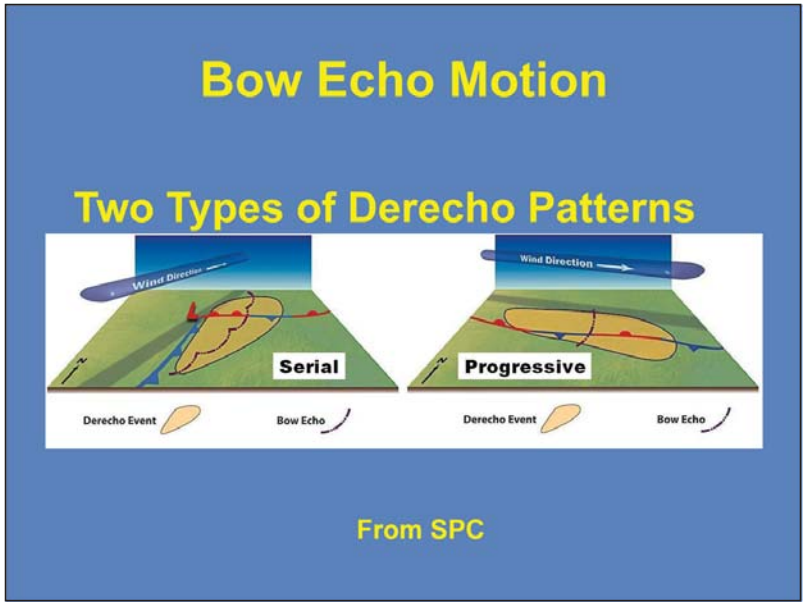


Supercells have been observed to evolve directly into a bow echo as the supercell decays. This type of evolution is typically seen with HP supercells. Supercells are also observed occasionally to exist within the larger bow echo structure. Next, I'll show you an example of this.

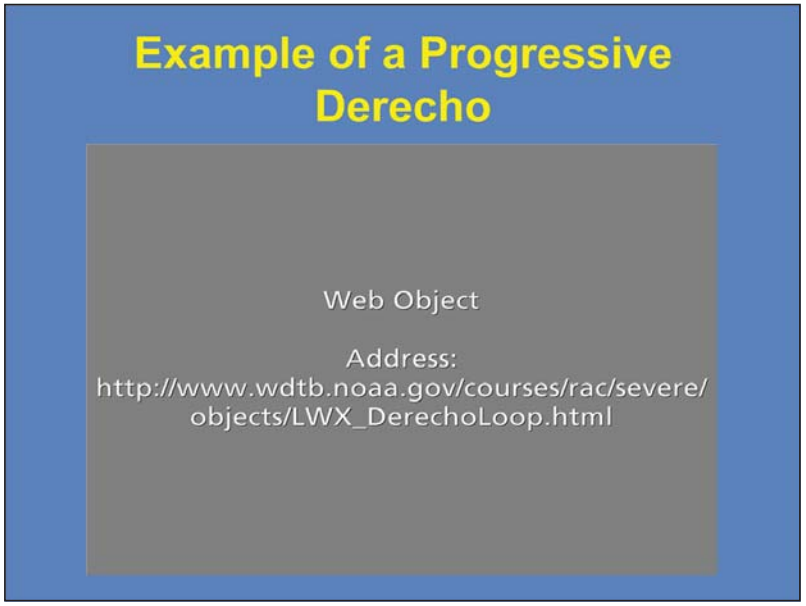
QLCS Mesovortex within Bow Echo



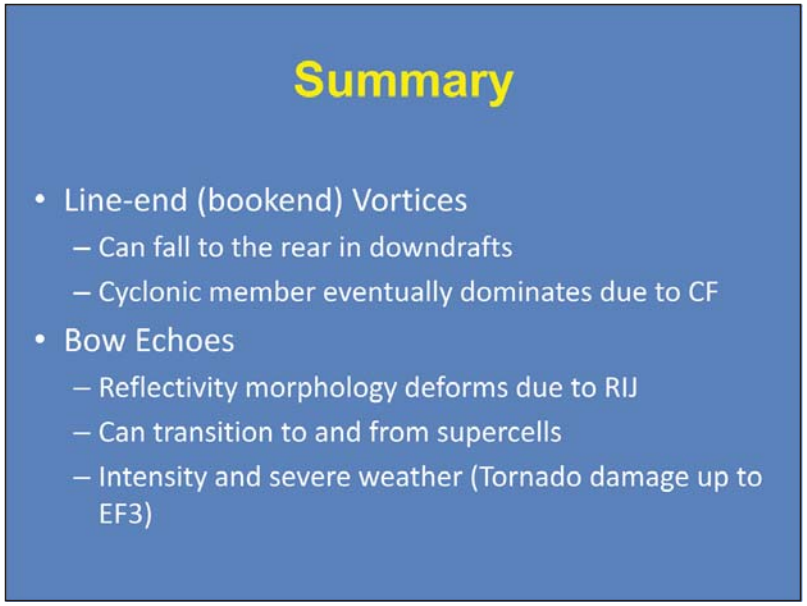
This is an image of 0.5 reflectivity from KSGF of a fast moving line of storms which produced a tornado of EF2 damage in Branson, MO at approximately 0723z 29 Feb 2012. The line contained an embedded QLCS mesovortex that formed a TVS and a TDS. As I fade to the SRM, you can easily detect the TVS, and the CC product helps to detect the TDS. As mentioned in the student guide, it is not uncommon to get up to EF3 scale damage from QLCS storms containing bow echo tornadoes.



Bows often move in the direction of the mean wind (or shear vector) at a speed related to its cold pool and system-relative flow. This is the case for both types of derecho events. Progressive derecho events usually occur along some sort of east-west oriented boundary and the bowing segments move parallel to the mean flow. Serial Derechos consist of an extensive squall line where the angle between the mean wind and squall line axis is relatively small (Johns and Hirt, 1987). The squall line typically moves normal to the mean wind at speeds of 30 kts or less, while the individual LEWPs and bow echoes move rapidly in the direction of the mean wind and tend to be most frequent near the northern end of the line. Serial derechos may contain supercells because the patterns which produce these types of derechos, the so called “dynamic pattern” is associated with strong shear.



The window that pops us displays a multi-product animation of a derecho event moving through the Washington D.C. area on evening of June 29, 2012. Note the motion was pretty much W-E with the exception of storms that moved northeast near the northern bookend and there were several embedded bow echoes with line end vortices and RINs especially observed as the line moved east of LWX radar. Take this opportunity to observe classic radar structures of a mature multicell producing derecho event, toggling through Z/V/ ZDR, and KDP products. When you are finished observing, advance to the next slide to view the summary and then, the Lesson Quiz.



This lesson covered a number of three dimensional multicell structures including line-end vortices, bow echoes and MCVs. Line-end vortices typically occur at the ends of breaks in convective lines with cyclonic and anticyclonic members on the northern and southern end of the breaks, respectively. These vortices fall behind the leading edge of a line segment if the shear is weaker. The Coriolis force tends to enhance the cyclonic member of the line-end vortices if they persist over a long enough period of time. Line-end vortices are typically occupied by downdrafts, however, tornadoes can spin up should any of the vorticity manages to phase with the cold pool boundary. Bow echoes are a special breed of mature multicells, usually when shear is moderate to strong, and are composed of an accelerating concave shape, rear inflow jet, deep lifting along the interface of updraft and downdraft. The reflectivity in the bow echo deforms from the enhanced channeling of the RIJ, which usually stays elevated and helps to make it long lasting and produce severe downdrafts and occasionally tornadoes.



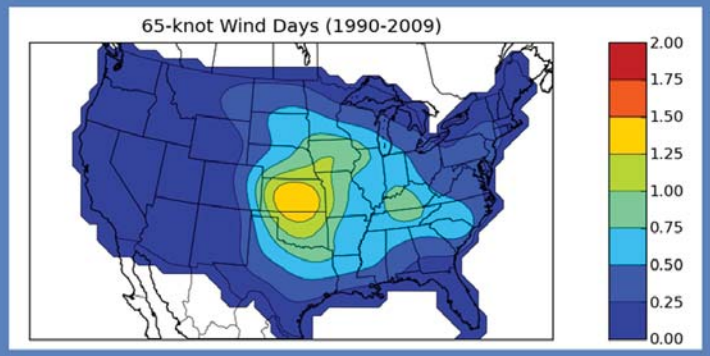
Welcome to the RAC Convective Storm Structure and Evolution lesson on Multicell Severe Wind Hazards.

Learning Objective

- Recognize multicell storm signatures for monitoring and anticipating damaging straight line winds.

This is the learning objective for Lesson 19.

Where Do Severe Winds Occur?



These are graphics from NSSL showing the mean number of severe thunderstorm wind days per year within 25 miles of a point. Note the three general frequency maxima: southern plains (OK/KS/MO/AR border region), Ohio /TN River Valley to the western Carolinas and one near DC and southern PA. There is a much weaker maxima in southern AZ due to the monsoons. When we fade to mean number of significant severe (> 65 kt) wind days per year, we see the maxima is centered squarely over KS and northern OK with a secondary maxima over TN/KY.

Multicell 2-D Structures

- 2-Dimensional (linear)
  - Leading edge high reflectivity cores
  - Trailing stratiform low reflectivity region
  - Movement essentially with mean wind
  - Most reflectivity gradients along the line do not extend well behind the updraft region
  - Gust front pushes well out ahead from location of leading edge convection as updraft weakens

Storm signatures associated with damaging winds from multicells are typically associated with squall lines containing supercells and/or bow echoes. Squall lines can form a variety of ways. The structure depends largely on the shear profile. With weak to moderate shear and subsequent slower system motion, the structure of the multicell complex is typically 2-dimensional with none of the characteristics discussed in the previous lesson. These structures would include:

• Leading edge possessing the high reflectivity convective cores.

• Trailing stratiform possessing the low reflectivity regions.

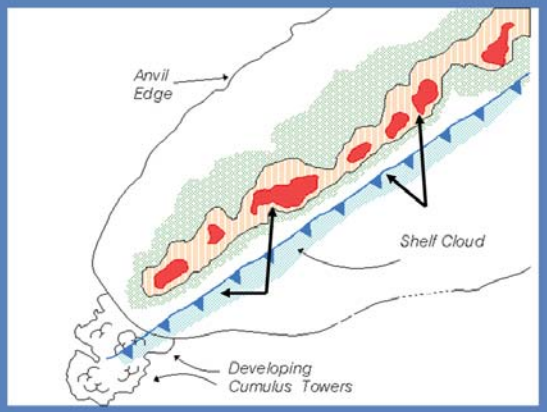
• Movement of the line with mean wind.

• No bowing, or in other words, points along the line do not extend beyond the leading edge.

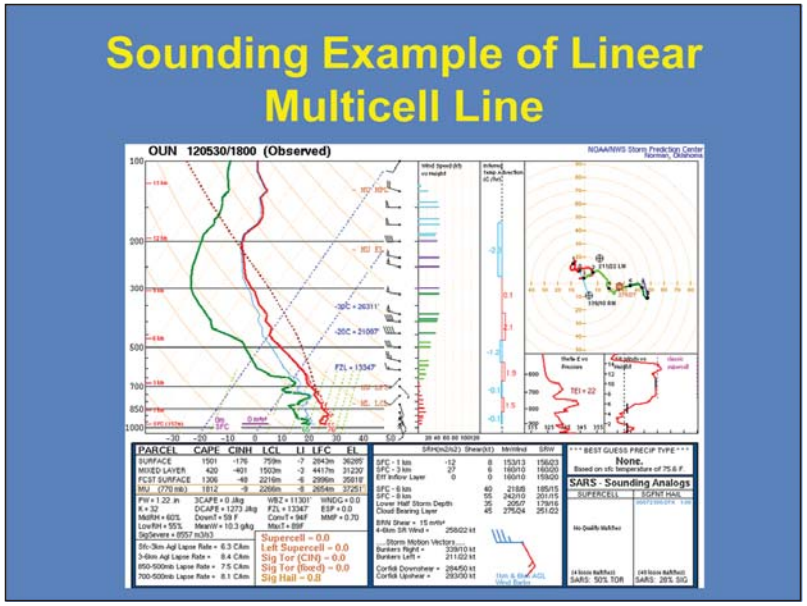
• Gust front pushes well out ahead of line.

I'll show a conceptual model of a plan view of this structure next.

Reflectivity 2-D Structure of a Multicell Line with Damaging Winds



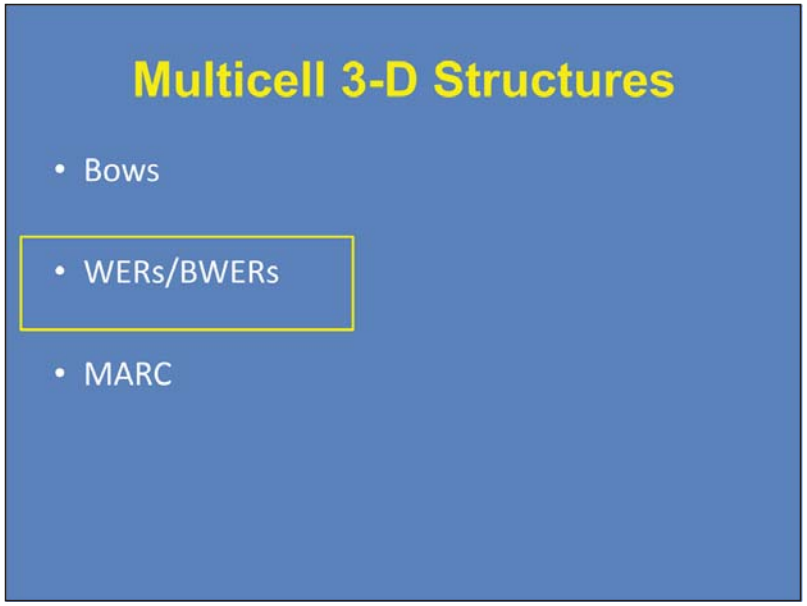
This is a conceptual model of a narrow multicell line possessing 2-D features. Note the position of the shelf cloud superimposed on the leading edge of the gust front, which, along with the regions around the stronger cores, is the most likely location for severe straight line winds. According to observations, flow is typically front to rear in cases such as these when the shear is not moderate to strong. Often this structure is observed in the early stages of squall line development.



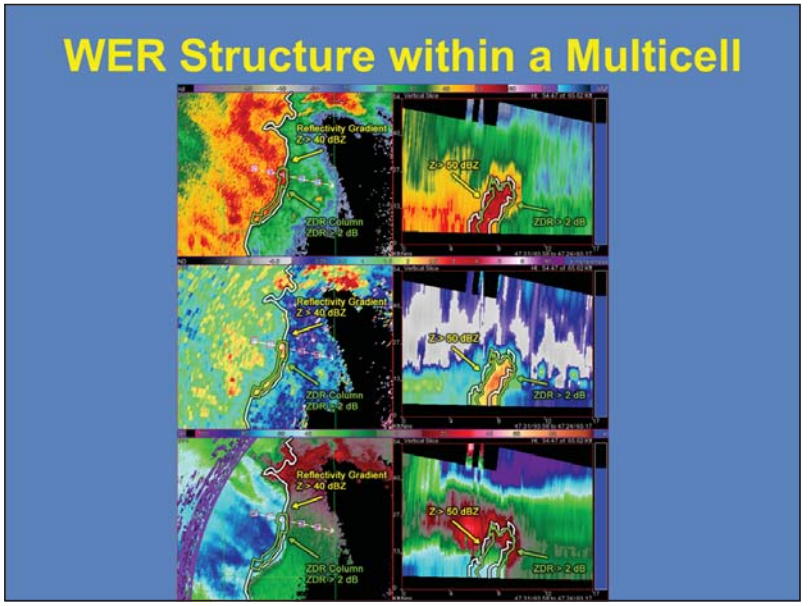
This is an upstream proximity sounding from Norman, OK at 1800 UTC on May 30, 2012. Note the weak flow in low to mid levels, but flow increases above 400 mb. The bulk shear from 0 to 6 km is 40 kts with a mean wind in the cloud bearing layer (from) 275 deg at 24 kts. Sounding has a MUCAPE of 1812 j/kg. There's still some surface based CINH, so weak winds in the diminished inflow layer. Above the LFC, the winds are unidirectional so most of the vorticity that can be tilted by the updraft will be crosswise. The upper level shear suggests that given this environment, you could get both linear 2-D and 3-D structures in subsequent multicell thunderstorm development.



This is a long loop of lowest scan reflectivity, velocity and ZDR from KICT from 2251z to 0224z showing the evolution of a multicell. Note how for roughly the first hour in the loop lifting with storms developing along the line were relatively discrete and motion dominated by individual cell propagation. Many storms in the multicell cluster were moving east while others were beginning to sag south as a result of different boundary relative flow. There was a large storm which had developed out ahead of the line just east of KICT and was moving away from the primary line to the north of the radar. There was an outflow boundary extending E-W from the storm to the middle of the line. When the gust front travels at the same speed as the multicell line, the boundary-relative flow maximizes potential for new cell growth along the leading edge. And the updraft becomes more erect, especially in the center portion of the line just north of the radar. 3D structures will evolve as the updraft becomes more upright and erect. Otherwise, when the gust front pushes well out , you can expect a more sloped front to rear 2D structure. If you toggle to the V product, you can see where the intense winds occurred just north and east of Wichita where trees and power lines were downed due to straight line winds up to 70 kts occurred. Intense updrafts due to the increasing shear and moderate buoyancy in the environment also produced some large hail up to 2" in diameter which is not uncommon in multicell structures especially in the southern Plains. The large trailing stratiform shield becomes more apparent after 0200z as the center part of the line weakens and more intense updrafts develop and accelerate along the line out to the west. Eventually the western part of the line increased and intersected the central part of the line.



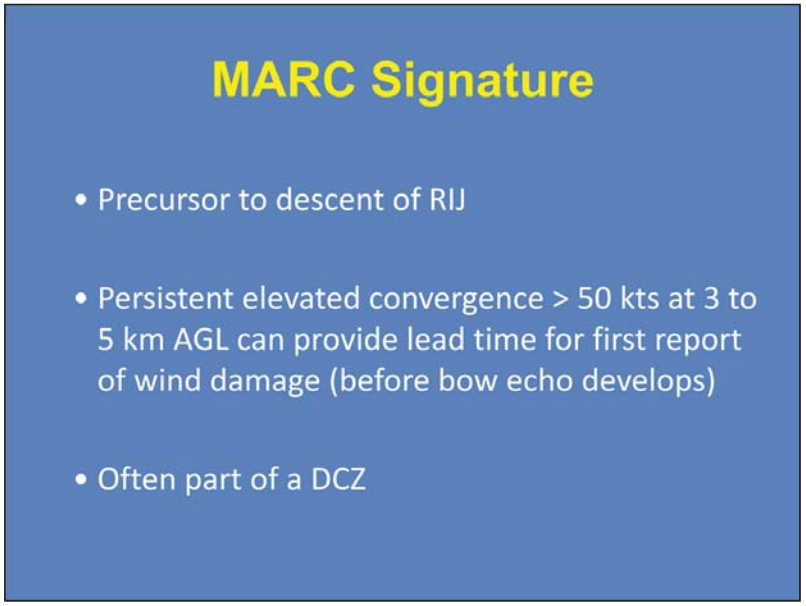
These are the 3-dimensional structures found in multicell squall lines and bows. Since we've already covered bow echo structures, I'm just going to show examples of a Weak Echo Region (WER) and a Mid-Altitude Radial Convergence (MARC) signature within a multicell.



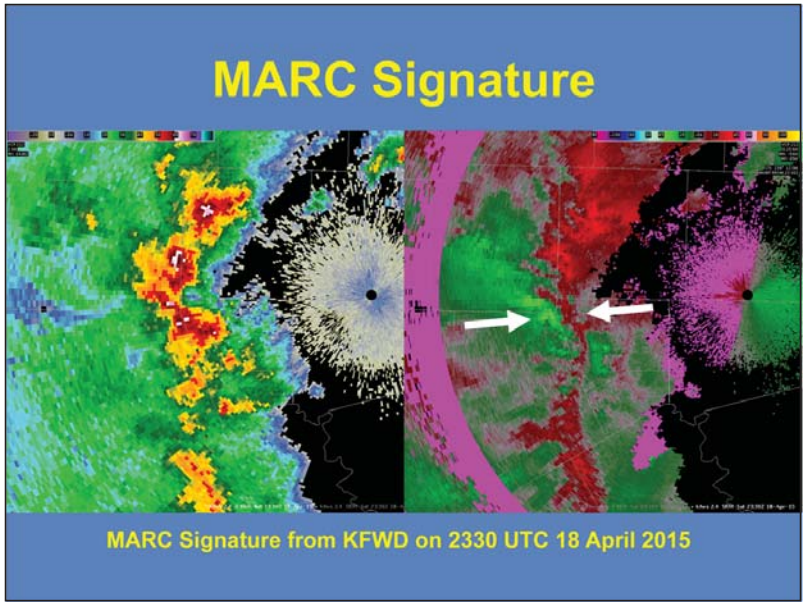
The existence of a WER (or BWER – "Bounded" WER) in the multicell structure of a QLCS most often indicates an enhanced potential for damaging surface winds. Remember that these features are associated with strong, deep layer wind shear. This figure (which is Figure 7-200 in the Student Guide) shows an example of a WER structure from a multicell from DLH on the evening of 2 July 2012. Note that the intense updraft is located along the leading edge of the line where the mid-level overhang and WER are located and where the strong low-level reflectivity gradient is located. Moreover, this line-segment is also bowing. The ZDR column where values are > 2DB are noted in the middle left panel. The cross-section views (right column panels) show the tilted updraft and ZDR column overhang due to shear and the deeper convergence extending up through the updraft. This case illustrates the characteristics of multicells structures in very strong shear. In terms of reports, numerous trees and power lines were down in Itasca County, MN.



This is a FSI screen capture at 0135 UTC from KDLH on July 3, 2012 of a highly sheared QLCS structure with a WER structure. You can toggle to V, ZDR, and KDP to see the DCZ and the vertically erect updraft in the cross-section portion of the 4-panel. Toggle to the velocity product and you can note the strong rear-to front flow with convergence all the way to the leading edge in association of a descending RIJ.



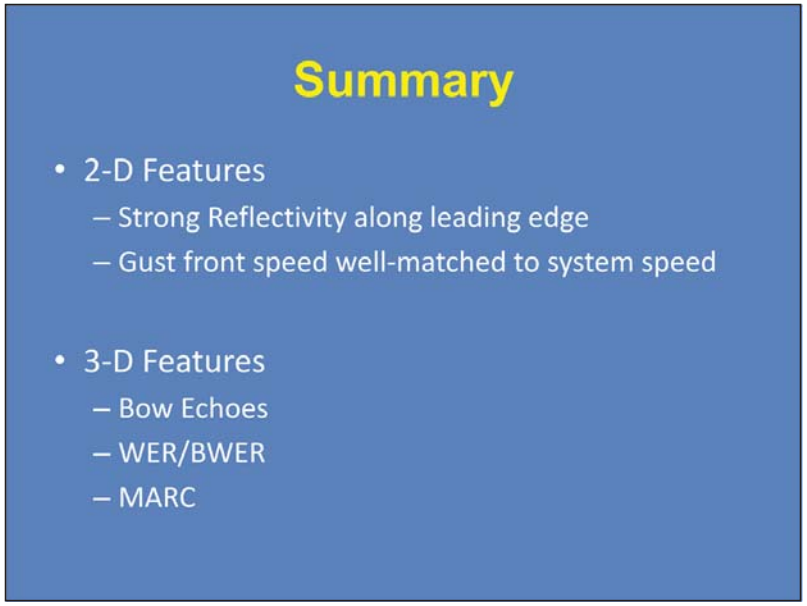
Another radar signature which is often associated with severe winds in multicells is the Mid-Altitude Radial Convergence Signature, abbreviated MARC. Observations of a MARC have been noted by Przybylinski and others as a precursor to the descent of the elevated RIJ. Enhanced velocity differentials which signify areas of strong convergence are often located just downwind of high reflectivity cores along the leading edge of the convective line. Persistent areas of MARC greater than 50 kts at 3-5 km AGL can sometimes provide lead time for the first report of wind damage (often before a well-defined bow echo with bookend vortex develops). MARC signatures are often part of a Deep Convergence Zone (DCZ) found in some intense updrafts where the Velocity differences of 30-55 m/s are found in both multicells and supercells. Let's look at an example of a MARC.



This is a 2.4 degree reflectivity and storm-relative velocity image of a Mid-Altitude Radial Convergence (MARC) signature. The white arrows indicate the location of the MARC signature. Remember, the MARC is detected via radial velocity data in the mid-levels (3-5 km AGL) within the intense reflectivity core of a squall line. Look for delta-V of at least 50 kts across the convergence axis. Advance to the next slide for a loop of this storm.



This is only a short loop but gives you the option to toggle back and forth from 2.4 deg Z and SRM, plus the 3 Dual-Pol products.



Multicell winds events arise from systems that occur on larger scales than individual downbursts. In the absence of mdt to strong shear, multicells usually possess 2-D characteristics such as strong reflectivities along the leading edge, and when the gust from remains close to the leading edge of the line.

When the multicell system is more mature, and when there is mdt to strong shear in the environment, 3-D features start to become more common such as bows, WERs/BWERs and MARCs. These all are classic signatures for severe winds in multicells (and in supercells). For a MARC signature, look for delta-Vs of at least 50 kts across the convergence axis through at least 3 km (~ 10 kft).

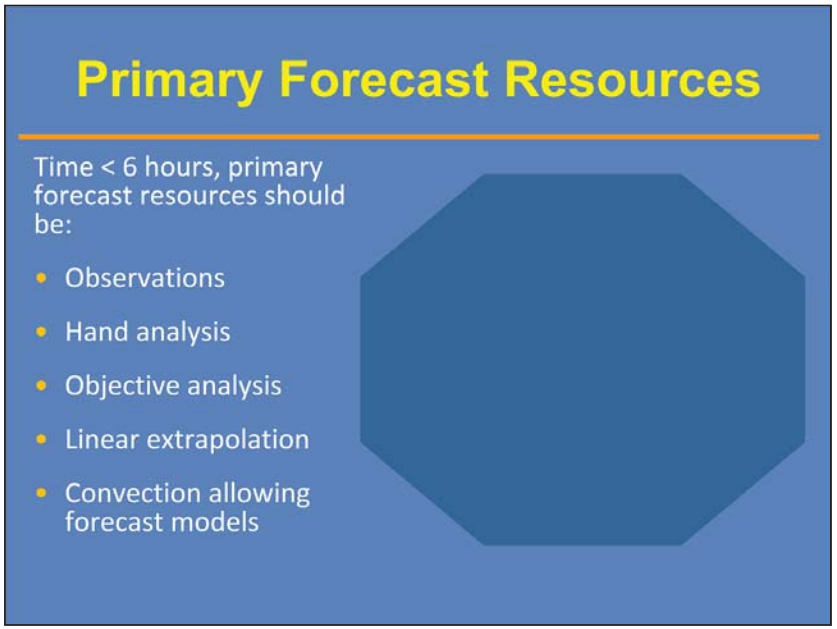
Applied Performance Drills	
Drill 1	**Convective Initiation**
Drill 2	**Storm Type**
Drill 3	**Hazard Potential**
Drill 4	**Screen and Rank Storms**
Drill 5	**Convective Mode**
Drill 6	**Identifying the ZDR Column**
Drill 7	**Supercell Signatures**
Drill 8	**50 dBZ Height**
Drill 9	**Storm Top Divergence**
Drill 10	**Severe Hail**
Drill 11	**Measuring Rotational Velocity in Mesocyclones**
Drill 12	**The MARC Signature**
Drill 13	**Identify a TVS or TS**
Drill 14	**Tornado Debris Signature (TDS)**



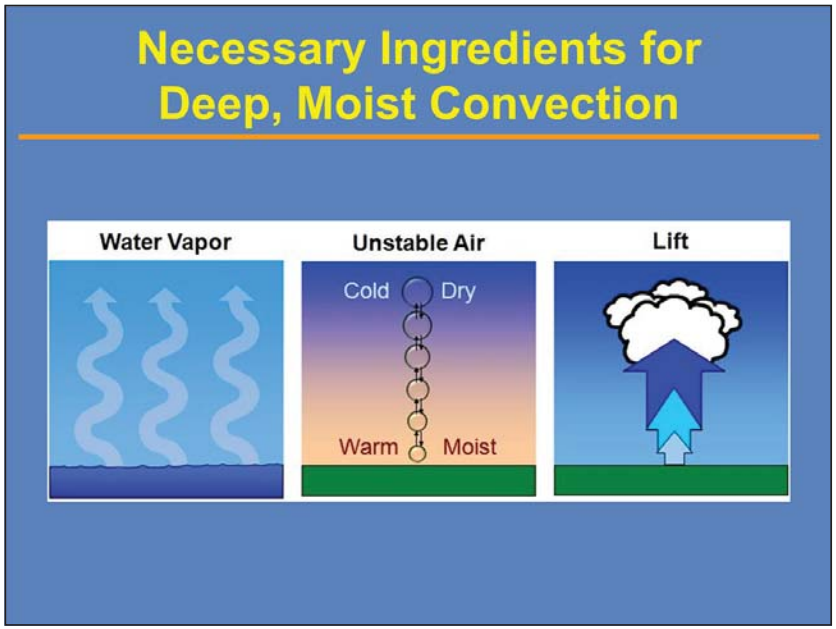
Welcome to the RAC Convective Storm Structure and Evolution, Applied Performance Drills: Convective Initiation.



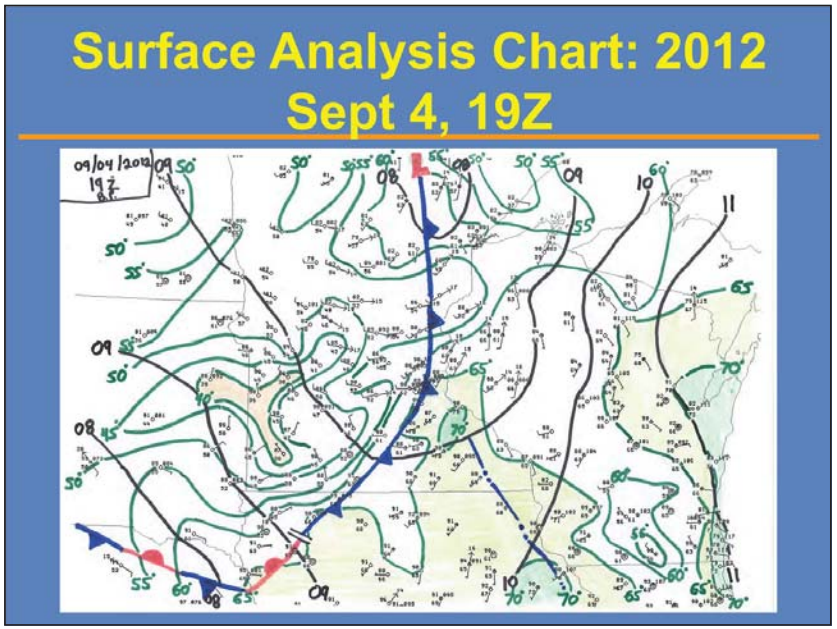
The performance objective for this lesson is: evaluate the pre-storm environment to assess its potential to produce deep, moist convection.



The potential for deep, moist convection can be forecast for most any time into the future, but for the operational warning environment, let's restrict ourselves to less than six hours. At this time scale, your primary forecast resources should be observations, hand analysis, objective analysis, linear extrapolation, and convection allowing forecast models.

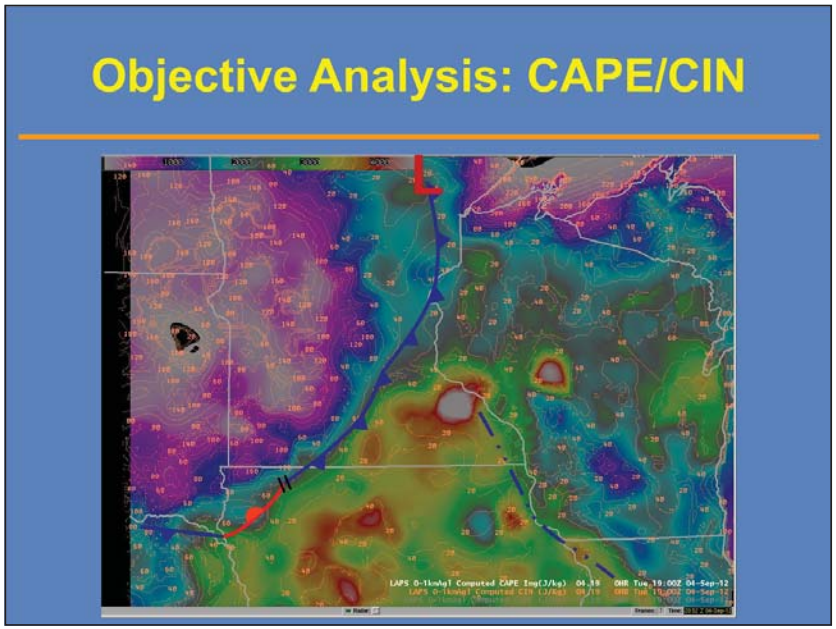


Deep, moist convection requires three ingredients: Sufficient water vapor, unstable air, and a lifting mechanism. Sufficient water vapor (commonly measured using dewpoint) must be present to produce unstable air. Virtually all showers and thunderstorms form in an airmass that is classified as "conditionally unstable." A conditionally unstable airmass requires a lifting mechanism strong enough to lift air parcels to their Level of Free Convection (LFC) to release the instability. Lifting mechanisms include: Converging winds around surface lows and troughs (via Ekman pumping), frontogenesis, upslope flow, drylines, outflow boundaries generated by prior storms, and local winds such as sea breeze, lake breeze, land breeze, and valley breeze circulations. Look for areas where all three necessary ingredients are co-located or soon will be.



Routine hand analysis of Surface Charts, augmented with high-resolution satellite and radar loops, is a critical step in assessing the atmosphere's potential to produce deep, moist convection. After all, how can we predict the future state of the atmosphere without starting with the best possible understanding of its present state? Hand analysis can reveal small, but critical features missed by objective analysis, particularly the location and nature of boundaries and lifting mechanisms. This is not to say objective analysis should be eschewed. Rather, it is to say that objective analysis should be trusted only upon its verification by hand analysis.

This 19Z Surface Analysis Chart revealed a low pressure center over northern Minnesota with a cold front extending south and southwestward into northwest Iowa. An outflow boundary from prior storms extended south-southeastward from southeast Minnesota across northeast Iowa.



Trends in mesoscale objective analysis can be used to anticipate convective development. Look for areas of decreasing Convective Inhibition (CIN) in regions where substantial Convective Available Potential Energy (CAPE) exists. Mixed Layer CAPE and CIN take into account the effects of parcel entrainment and, thus, typically do a better job than surface-based versions.

This 19Z objective analysis of 0-1km mixed layer CAPE and CIN revealed extreme instability in the warm sector with little or no CIN remaining, particularly over extreme westcentral Wisconsin and southeast Minnesota.



High-resolution (particularly visible) satellite loops can reveal where convective initiation is likely. It can reveal fine-scale details of lifting mechanisms including fronts and outflow boundaries. Cumulus fields reveal where surface-based convection is occurring within a shallow layer beneath the LFC. The development of towering cumulus indicates parcels have either reached their LFC or may do so soon, especially if there is strong, deep ascent associated with a lifting mechanism.

This visible satellite loop revealed cumulus and moderate cumulus along and ahead of the cold front and outflow boundary. It suggested storms were possible given further heating, destabilization, and lift.



Welcome to the RAC Convective Storm Structure and Evolution, Applied Performance Drills: Storm Type.

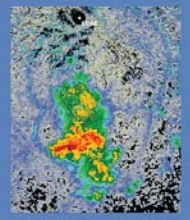
Performance Objective

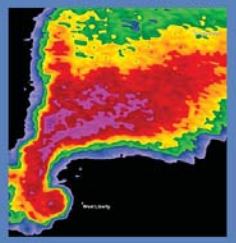
Evaluate the pre-storm environment and, given deep, moist convection, assess the expected storm type (ordinary cell vs supercell).

Our performance objective is: evaluate the pre-storm environment and, given deep, moist convection, assess the expected storm type. That is, ordinary cell versus supercell.

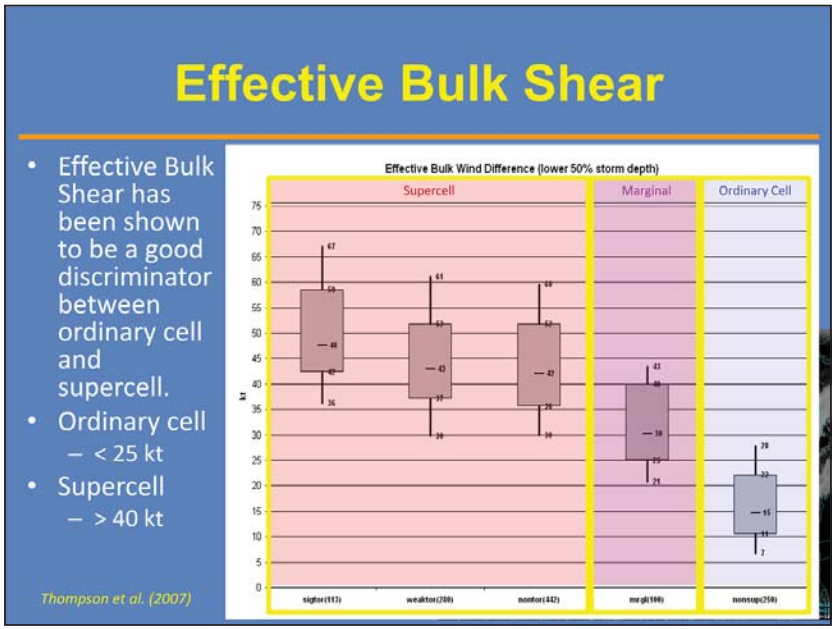
Supercell vs Ordinary Cell

- It's important to distinguish between supercell and ordinary cell environments
- Supercells produce:
  - Most ≥EF2 and virtually all ≥EF4 tornadoes
  - High percentage of ≥2" and virtually all ≥4" hail
  - Supercell Rear Flank Downdrafts (RFDs) often produce ≥65 kt wind
  - HP supercells can be prolific flash flood producers

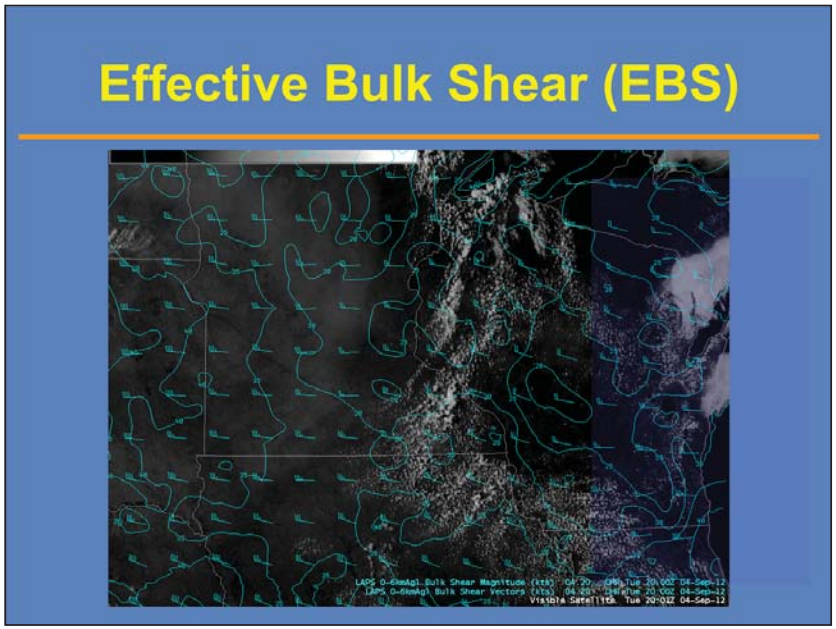




It's important to distinguish between environments which favor supercells versus those which only support ordinary cells. Supercells produce a disproportionate share of severe weather including most significant (>EF-2) and virtually all violent (>EF-4) tornadoes. Supercells produce a high percentage of significant (≥2" diameter) and virtually all giant (≥4" diameter) hail. Supercell rear flank downdrafts (RFDs) often produce significant (≥65 kt) wind as well. Finally, high precipitation (HP) supercells can be prolific flash flood producers.



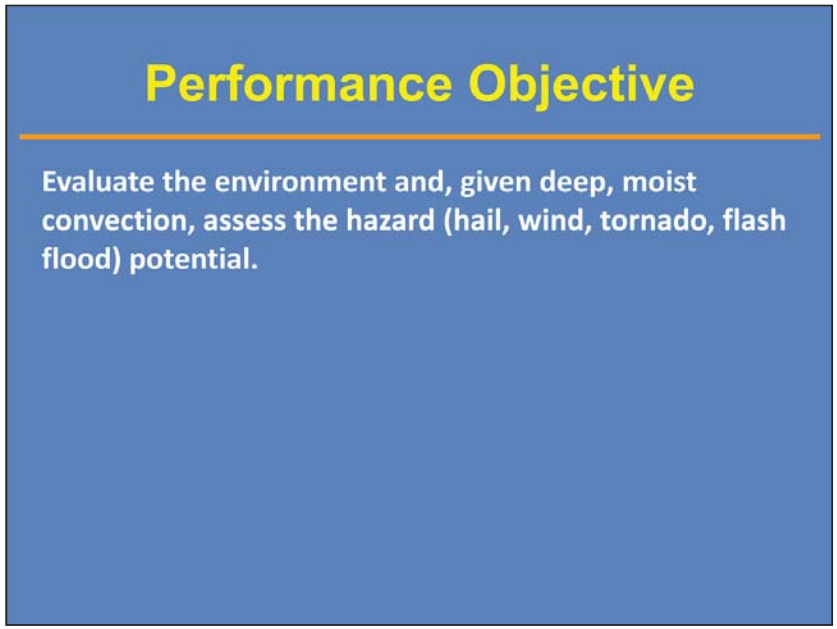
Effective Bulk shear (EBS) has been shown to be a good discriminator between an ordinary cell and a supercell. Supercell potential increases with effective bulk shear. In general, values less than 25 knots favor ordinary cells, while values greater than 40 knots favor supercells. Values between 25 to 40 knots represent a marginal area where supercells are possible, but not necessarily favored. You can find effective bulk shear on the SPC Hourly Mesoscale Analysis web page.



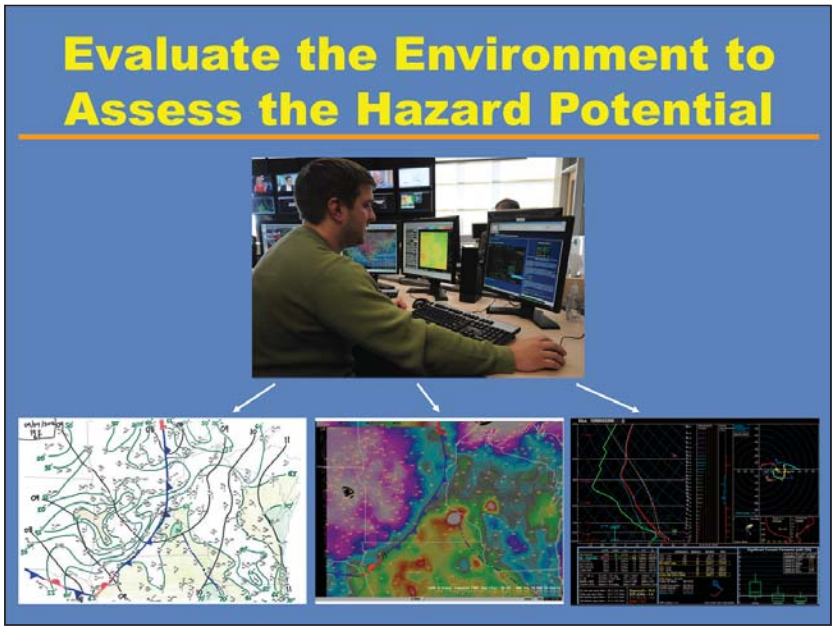
This 20Z visible satellite image is overlaid with the 20Z LAPS 0-6 km Bulk Shear Magnitude and Vectors. AWIPS does not calculate effective bulk wind shear, but this is a good alternative with environments supportive of relatively deep, surface-based convection. Values along the front and outflow boundary were in the marginal range with values generally between 25 to 35 knots.



Welcome to the RAC Convective Storm Structure and Evolution, Applied Performance Drills: Hazard Potential.

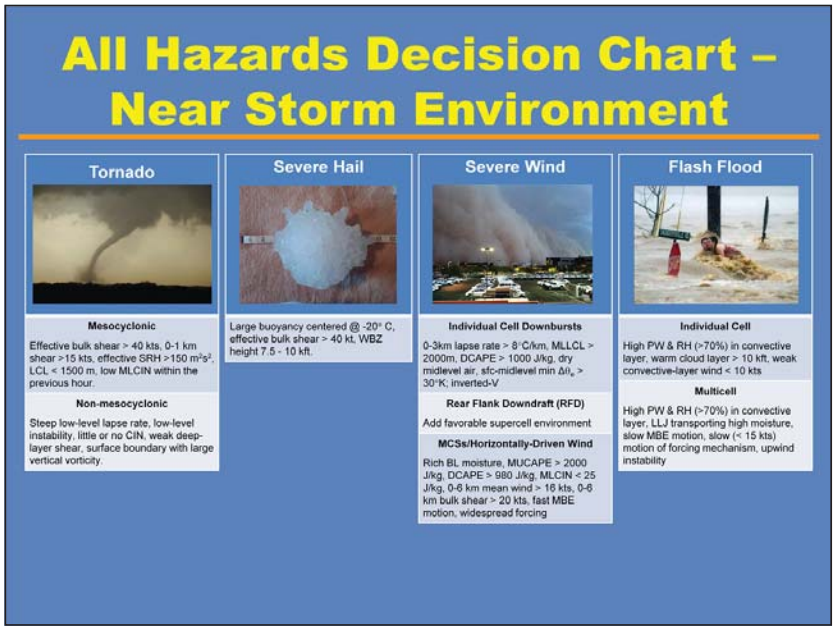


The performance objective for this module is: evaluate the environment and, given deep, moist convection, assess the hazard potential. That is, hail, wind, tornado, or flash flood.

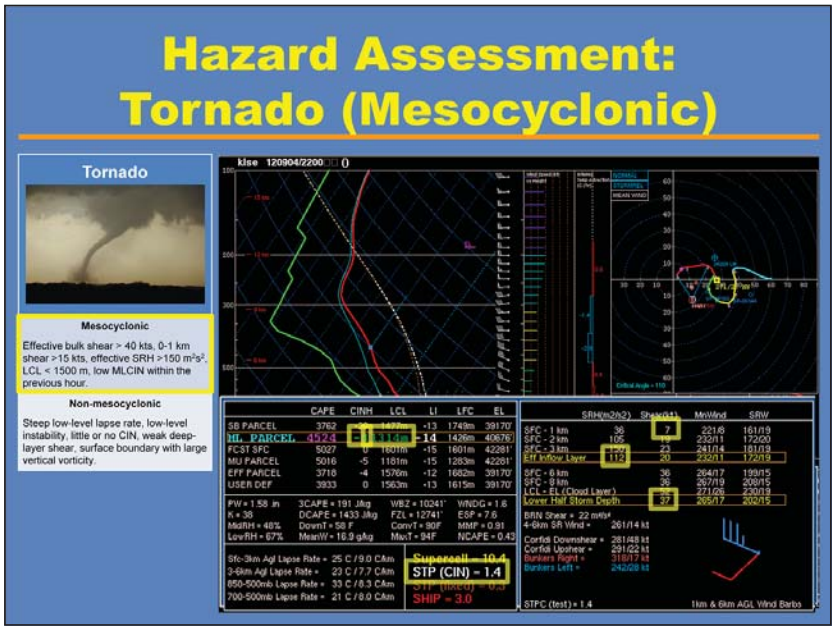


Warning forecasters should continually evaluate the environment to assess the severe weather hazard potential. This enables the warning operations team members to anticipate the type of severe weather that is expected (or at least possible) and bias their decisions appropriately. During active warning operations, this process is often best handled by a dedicated meso-analyst because forecasters in the "hot seat" may lose situational awareness of the near storm environment just when they need it the most.

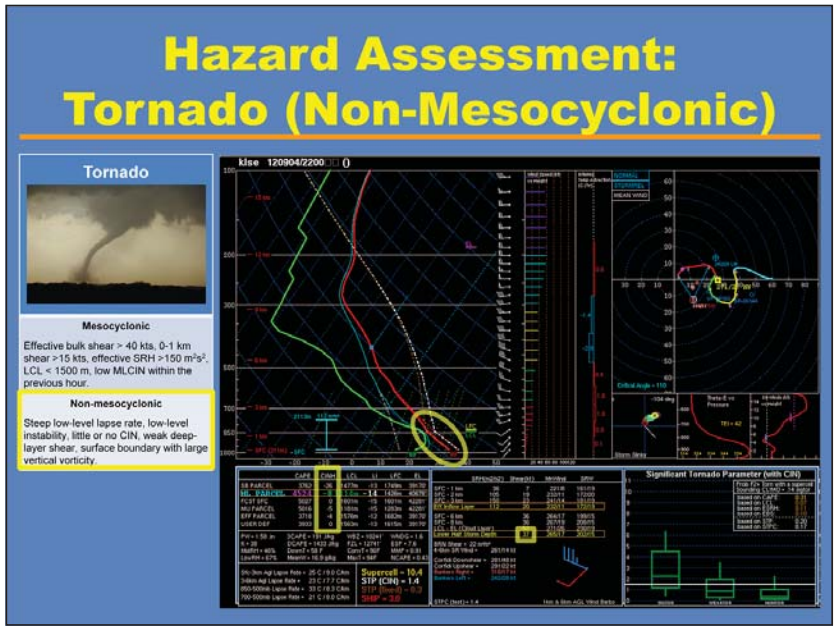
The environmental evaluation and hazard assessment process is best done using a combination of hourly hand analyzed surface charts, plan projection views of objective analysis and model forecast ingredients, and representative near-storm environmental soundings.



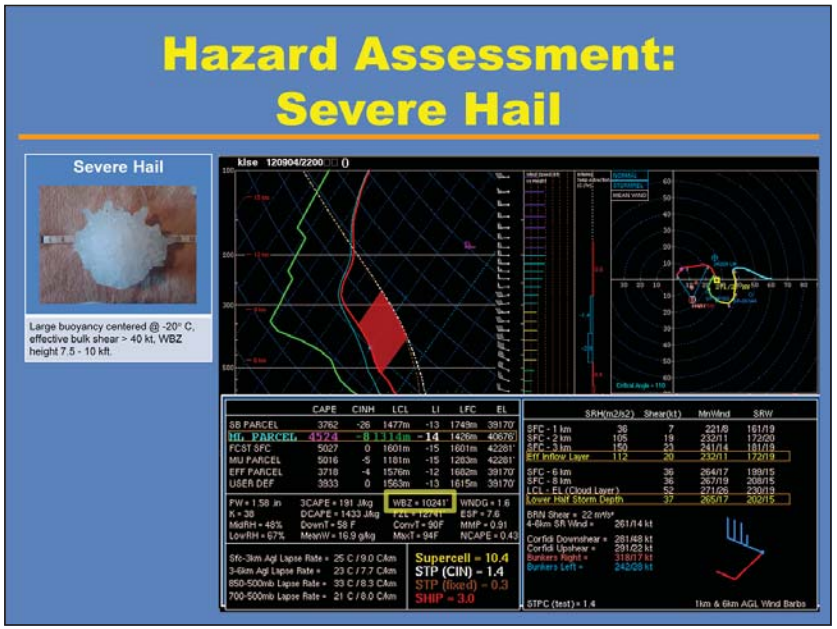
The WDTB "All Hazards Decision Chart" provides a summary of favorable near storm environmental characteristics associated with severe weather. Bear in mind, we base these ingredients and their values on the best available published scientific research whenever possible. But some are simply our best guess based on experience. Let's use this chart to evaluate the environment and assess the hazard potential for a severe weather event in the La Crosse, Wisconsin CWA on September 4, 2012.



Favorable ingredients for mesocyclonic tornado are highlighted in the yellow box. This RAP near-storm environmental sounding at La Crosse exhibited Effective Bulk Shear (denoted as "Lower Half Storm Depth) of 37 knots which was slightly below our threshold of 40 knots. 0-1 km shear was only 7 knots which was below our threshold of 15 knots. Effective ("Inflow Layer") storm-relative helicity was 112 m^2/s^2 which was also below our threshold of 150 m^2/s^2. The LCL height of a mixed-layer parcel was 1314m which was below our threshold of 1500m. MLCIN was only 8 J/kg which meets our criteria of low CIN. The combination of ingredients lead to an Effective Significant Tornado Parameter ("STP (CIN")) of 1.4 which suggested mesocyclonic tornado was possible.



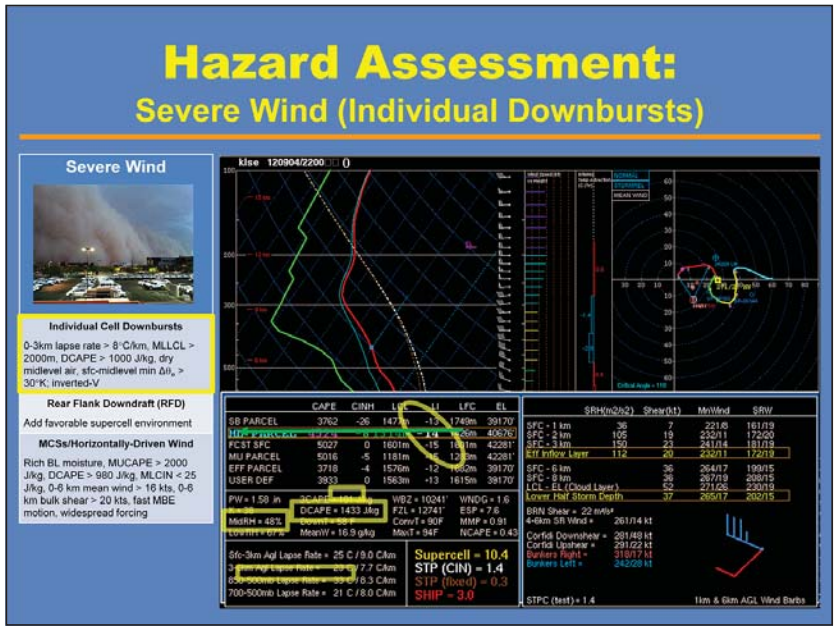
Favorable ingredients for non-mesosyclonic tornado are highlighted. The near-storm environmental sounding indicated a steep low-level lapse rate and little CIN, but the Effective Bulk Shear was 37 knots which tended to inhibit non-mesocyclonic tornado production. Plan projection views of surface vorticity objective analysis (not shown) indicated only low values along the front and outflow boundary. All these factors suggested the non-mesocyclonic tornado risk was low.



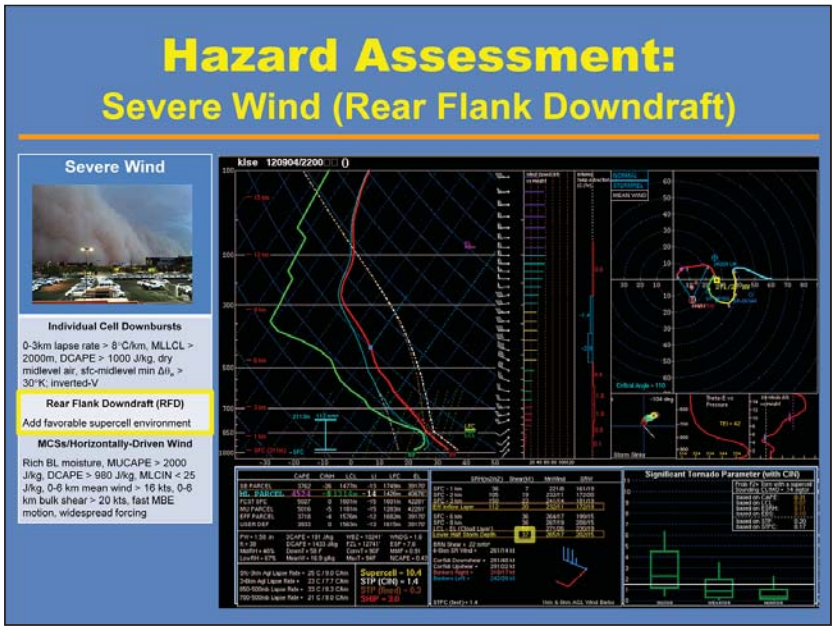
Hail forecasting is an immature science. On a nationwide basis, commonly used hail predictors show little or no skill in predicting hail size. However, supercells produce a very high percentage of significant ($\geq$ 2") and virtually all giant ($\geq$ 4") hail. Thus, given the potential for thunderstorms, every hail forecaster should evaluate supercell potential. The La Crosse sounding indicated an Effective Bulk Shear of 37 knots which was marginally supportive of supercells.

Large CAPE in the -10°C to -30°C layer suggests stronger updrafts in the supercooled liquid water region and the potential to suspend hailstones within this hail growth zone for a longer time, theoretically leading to larger hailstones. The La Crosse sounding indicated a large amount of CAPE in this hail growth zone (centered at -20°C) suggesting a high potential for severe hail.

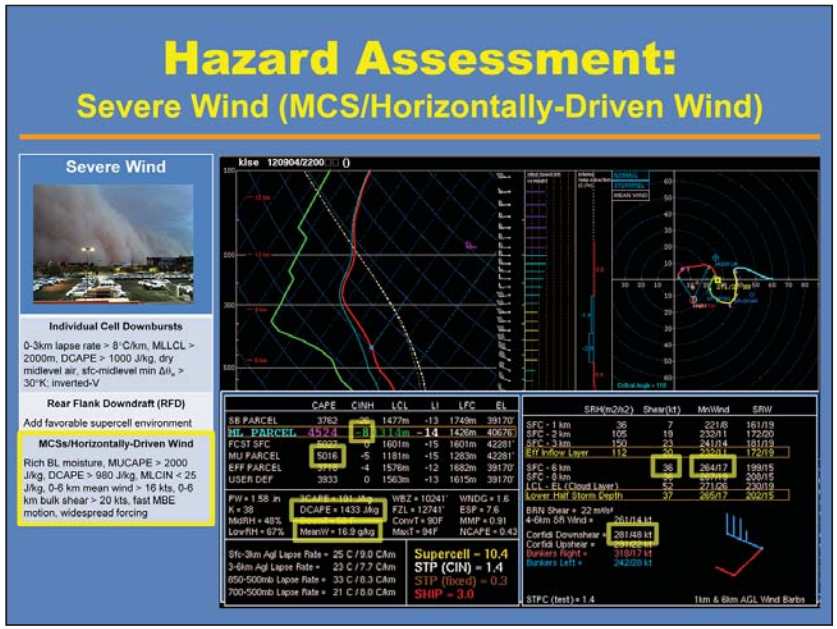
The wet-bulb zero height was 10,241 ft, slightly above our most favorable zone of 7,500 to 10,000 ft.



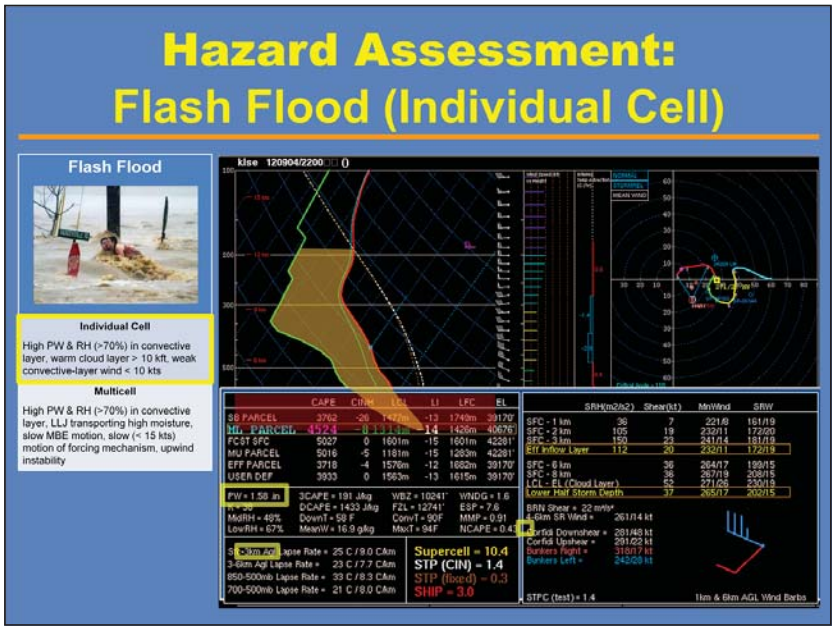
Favorable environmental ingredients for "Individual Cell Downbursts" is highlighted. The La Crosse sounding indicated a "Sfc-3km AGL Lapse Rate" of 9.0°C/km which exceeded our threshold of 8°C/km. Mixed-layer LCL was only 1314m, which was well below our threshold of 2000 m. But the DCAPE of 1433 J/kg easily exceeded our threshold of 1000 J/kg. Mid-level RH of 48% marginally met our criteria of "dry midlevel air." Surface to midlevel minimum theta-E is not denoted on this sounding so it's difficult to evaluate. The sounding did not exhibit an inverted-V profile common with dry microbursts which originate in midlevels. The upshot is severe downbursts are possible.



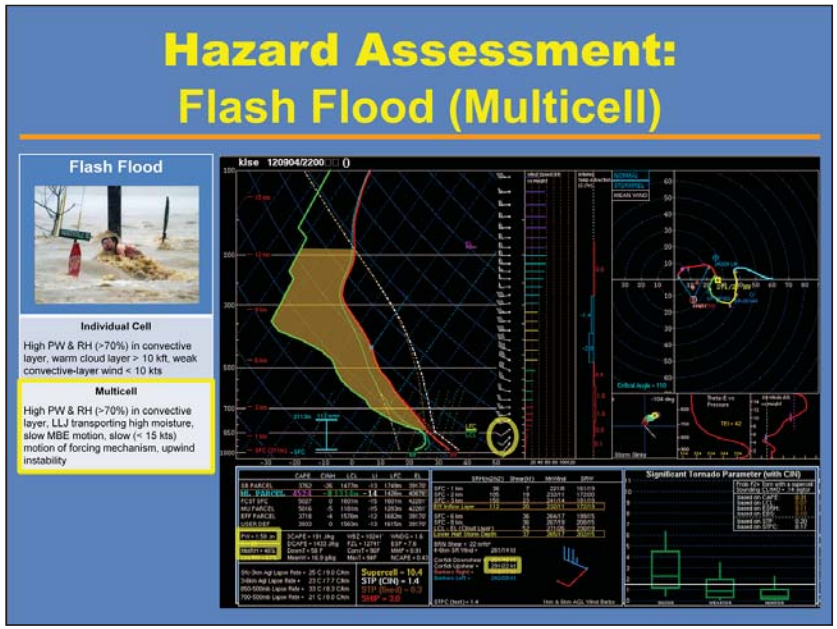
Favorable environmental ingredients for severe wind associated with supercell "Rear flank downdraft" is highlighted. The ingredients are the same as "individual cell downbursts" with the addition of favorable supercell environment characteristics. Effective Bulk Shear of 37 knots is marginally favorable for supercell.



The La Crosse sounding exhibited all the favorable characteristics for severe horizontally-driven wind associated with an MCS, including: Rich boundary-layer moisture (MeanW = 16.9 g/kg), MUCAPE greater than 2000 J/kg (5016 J/kg), DCAPE greater than 980 J/kg (1433 J/kg), MLCIN less than 25 J/kg, 0-6k mean wind greater than 16 knots, 0-6km bulk shear greater than 20 knots, fast Mesoscale Beta Elements (MBE) motion (denoted by Corfidi Downshear = 48 knots), and widespread forcing due to a cold front (not shown).



The La Crosse sounding was generally unfavorable for flash flooding associated with an individual cell. Precipitable water of 1.58 inches was between the 85-90th percentiles relative to climatology (not shown), but mean RH was well below 70% in the convective layer, the warm cloud layer depth from the LCL to the freezing level was only about 8,400 feet, and convective-layer wind of 52 knots was much stronger than the 10 knot (threshold). Remember that we should also consider antecedent soil moisture for the hydrological input into a flash flood potential assessment.



The La Crosse sounding was generally unfavorable for flash flooding with multicells as well. Precipitable water of 1.58 inches was between the 85-90th percentiles relative to climatology (not shown), but mean RH was well below 70% in the convective layer. There was no significant low-level jet to transport high moisture, and the mesoscale beta elements (MBE) motion ("Corfidi Upshear") was 22 knots. The outflow boundary was quasi-stationary and the cold front was moving east-southeast at around 10-15 knots, but there was no significantly increased upwind instability.



Welcome to the RAC Convective Storm Structure and Evolution, Applied Performance Drills: Screen and Rank Storms.

Performance Objective

Given a lowest tilt, Base Reflectivity, 60-minute time lapse loop with WSR-88D algorithm (TDA, MDA, HDA) overlays, rank the top three storms by their hazard (Hail, wind, tornado, flash flood) potential.

Our performance objective for this module is: Given a lowest tilt, Base Reflectivity, 60-minute time lapse loop with WSR-88D algorithm overlays, rank the top three storms by their hazard potential.

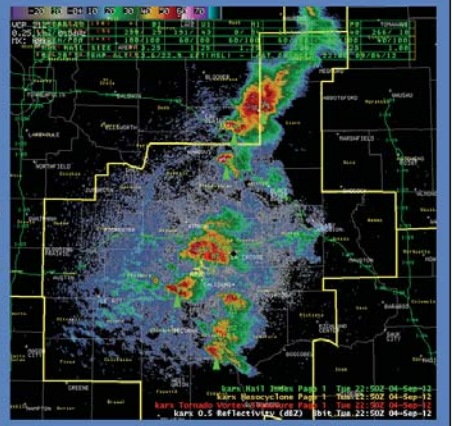
Suggested Warning Methodology
Screen, Rank, Analyze, Decision (SRAD)

- The WDTD Suggested Warning Methodology is a process by which NWS warning forecasters can efficiently evaluate and assess the severe weather potential of storms.
  - It's based on the SRAD process:
    1. **S**creen
    2. **R**ank
    3. **A**nalyze
    4. **D**ecision

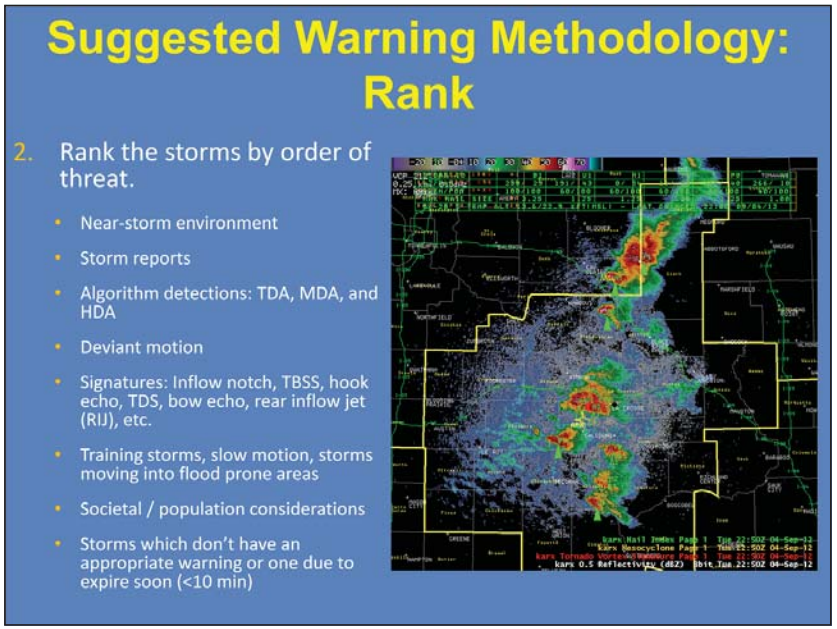
The WDTD "Suggested Warning Methodology" provides a structured process for making storm-based warning decisions. The first two steps in the process are "screen" and "rank."

Suggested Warning Methodology: Screen

1. Screen the storms that threaten lift and property over your CWA.
   - A lowest-tilt, Base Reflectivity, 60 minute time lapse loop with algorithm overlays works well.
   - *Note: An alternative could be a 4-panel: Z, SRM, STP, DVIL*



First, you should **Screen** the storms that threaten life and property over their County Warning Area (CWA). A lowest-tilt, Base Reflectivity, 60 minute time lapse loop with algorithm overlays works well. *(Note: An alternative could be a 4-panel: Z, SRM, STP, DVIL).*



Next, you should **Rank** the storms by order of threat. Factors to consider include:

Near-storm environment

Storm reports

Algorithm detections: TDA, MDA, and HDA

Deviant motion

Signatures: Inflow notch, three-body scatter spike (TBSS), hook echo, TDS, bow echo, rear inflow jet (RIJ) etc.

Training storms, slow motion, storms moving into flood prone areas

Societal / population considerations

Storms which don't have an appropriate warning or one that's due to expire soon (<10 min)



In this example from the La Crosse, Wisconsin CWA, let's rank the storms assuming none has a warning. I'd rank storm A1 the highest. Based on our previous assessment it's in the most favorable environment and along an old outflow boundary. It has supercell reflectivity characteristics including an inflow notch, hook echo, and deviant rightward motion. It has tornado vortex, mesocyclone, and (up to 3") hail algorithm detections. Plus, it's approaching the La Crosse metro area.

The second highest ranked storm is B1. It has supercell reflectivity characteristics as well including an inflow notch, hints of a hook echo, and deviant rightward motion. It also has algorithm indication of hail up to 3" in diameter.

The third highest ranked storm is debatable, but I'd go with storm F1. It's a physically large storm, appears to be right-moving, and produces periodic algorithm indications of a mesocyclone and hail up to 2" in diameter.

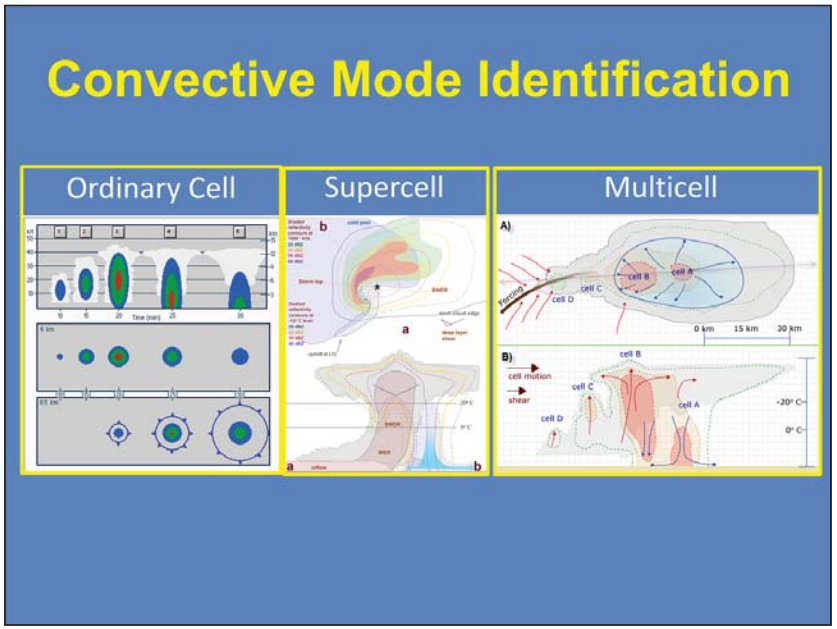


Welcome to the RAC Convective Storm Structure and Evolution, Applied Performance Drills: Convective Mode.

Performance Objective

- Given a specified storm, identify its convective mode (ordinary cell, supercell, or multicell).

The performance objective is: Given a specified storm, identify its convective mode. That is, if it is an ordinary cell, supercell, or multicell.



It's important for you as a warning forecaster to be able to identify a storm's convective mode: Ordinary cell, supercell, or multicell. That way, you get a good idea of how the storm is likely to behave and how to warn for it, including its expected severe weather hazards, motion, and evolution. Let's go back to basics and review a (truncated) definition of these three convective modes.

An **ordinary cell** the most basic component of a convective storm, consisting of a single main updraft that is usually quickly replaced by a downdraft once precipitation begins. Ordinary cells are especially observed in environments with weak vertical wind shear, and typically have lifetimes of 30-50 minutes. Ordinary cells are the primary component of multicell storms.

A **supercell** consists primarily of a single, quasi-steady rotating updraft (i.e., a mesocyclone).

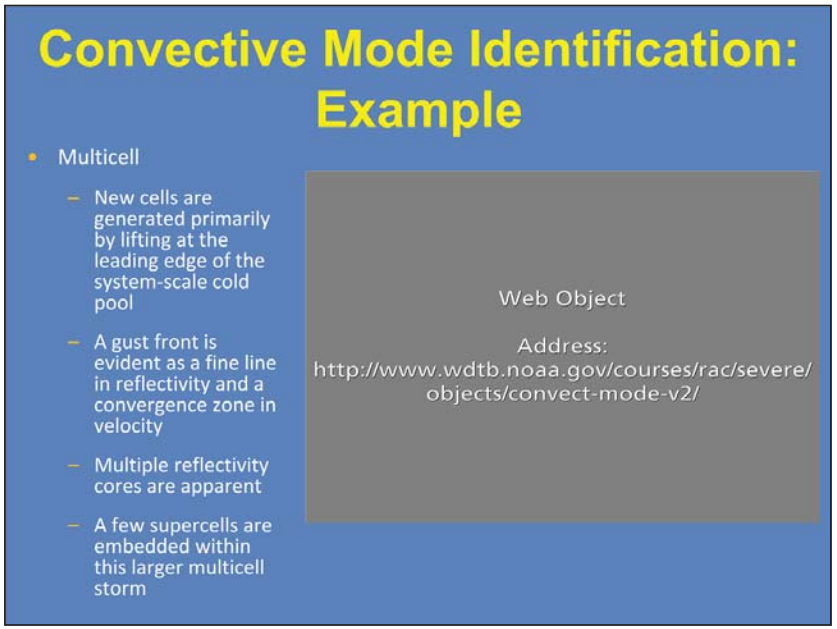
It typically has a very organized internal structure that enables it to propagate continuously. It may exist for several hours and usually forms in an environment with strong vertical wind shear. Supercells propagate in a direction and with a speed other than indicated by the mean wind in the environment. Such storms sometimes evolve through a splitting process, which produces a cyclonic, right-moving (with respect to the mean wind), and anticyclonic, left-moving, pair of supercells.

Supercells comprise a spectrum, but are often sub-categorized based on the extent to which their mesocyclone is wrapped in precipitation as revealed by their radar and/or visual appearance. The subcategories are: Low precipitation, high precipitation, and classic.

Supercell radar characteristics include: An inflow notch, reflectivity max displaced closer to the reflectivity gradient, weak echo region (WER), bounded weak echo (BWER), echo top displaced over the gradient/WER/BWER, mesocyclone, and hook echo.

A **multicell** is a cluster of ordinary cells and/or supercells at various stages of their life cycle in close enough proximity to at least share a common precipitation area and cold pool (gust front).

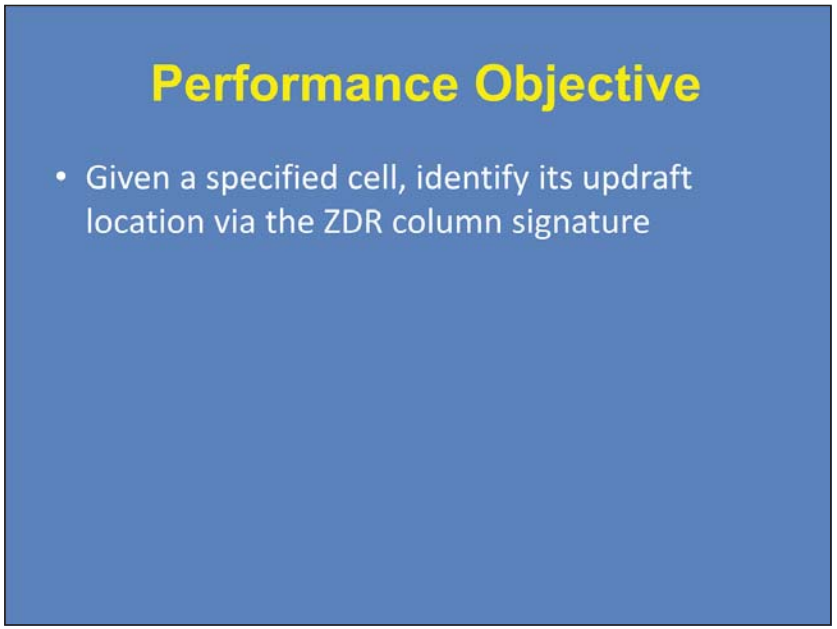
New cells are generated primarily by either low-level convergence along a preexisting boundary or by lifting at the leading edge of the system-scale cold pool that was produced by the previous cells. The cells move roughly with the mean wind. However, the storm motion usually deviates significantly from the mean wind due to discrete propagation (new cell development) along the gust front. The multicellular nature of the storm is usually apparent on radar with multiple reflectivity cores and maximum tops. Lifetime may be several hours.



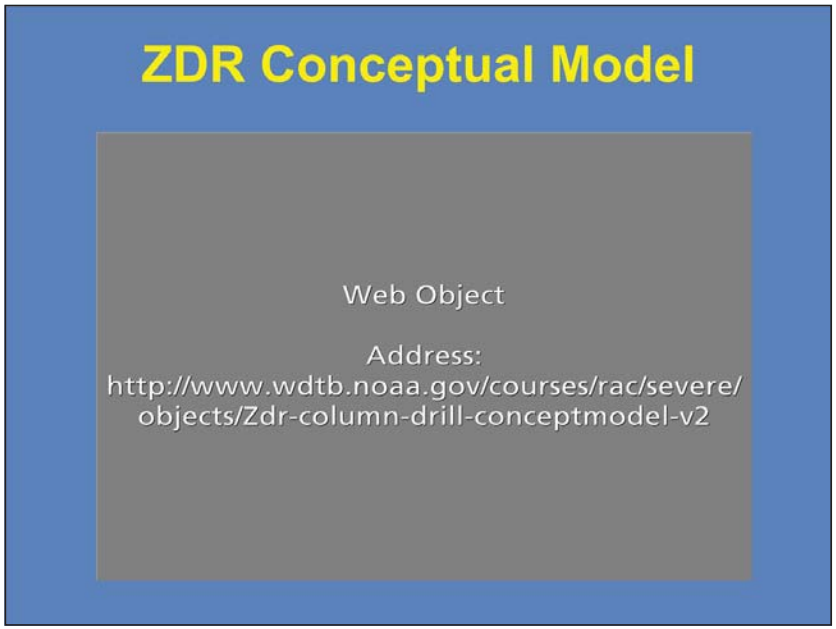


Let's evaluate the radar data from this storm near the La Crosse WSR-88D to identify its convective mode. I've got a 20-frame, 0.5 degree Reflectivity and Velocity loop to help us with the identification.

Radar assessment reveals this to be a multicell. New cells are generated primarily by lifting at the leading edge of the system-scale cold pool. A gust front is evident as a fine line in reflectivity and a convergence zone in velocity. The multicellular nature of the storm is apparent due to the multiple reflectivity cores. A few supercells are embedded within this larger multicell storm.



The objective for this drill is: Given a specified cell, identify its updraft location via the ZDR column signature.



Click the play button on this video to recap the conceptual model of a ZDR column. This video will last a minute.

Let's go ahead and do a drill on how to detect ZDR columns in AWIPS. But first look at this conceptual model on ZDR columns we presented in topic 7.

Remember that a ZDR column is an area of high ZDR of 1-2 dB that extend above the melting level in a region of fairly high reflectivities, as you can see here above 20 dBZ represented by the dashed green contour.

So it goes up into the midlevels of the storm in the cold regions where you can see the high reflectivities in the plan view closely associated with the updraft.

The updraft is transporting large drops from below the melting level upward into subfreezing air.

But remember one thing, the ZDR column may not form if the melting level is close to the LCL because there's not much room for warm rainfall processes or melting hail to recycle back into the updraft.



Play this video to get a first hand account of looking for a ZDR column with real radar data.

Let's take a look at the 3 hr NAM valid at 21Z to see where the ZDR column typically extends.

Now in the warm layers above freezing high ZDRs and heavy rainfall are commonplace in storms. But then those ZDRs tend to trail off once you go above the freezing level.

And only those areas of high ZDRs above this level you would call a ZDR column as long as they are co-located with high reflectivities and it makes sense that this is part of a thunderstorm updraft.

So a ZDR column starts at this level and extends into the subfreezing air, and in many cases extend to the -10 C level. In this case here, about 18 kft MSL.

Really strong updrafts will tend to push high ZDR values up to -20 C or about 23 kft MSL in this case.

Okay, now that we saw the sounding, let's take a look at the four panel display.

In the procedure group called radar, you'll see the radar name followed by all-tilts base data with standard environment lookup table. You'll see something that shows reflectivity (upper left), ZDR (upper right), CC (lower right) and KDP (lower left) in the front display, and going to the back display you'll see velocity images.

So what we want to do is to step up from the lowest levels that we can and note the high ZDR values > 1 dB where reflectivities are greater than 20 dBZ.

Look hard for those areas of high ZDR as you go above the melting level, and we go up one more scan to -2 C and we see high reflectivities and high ZDRs. This is the beginning of the ZDR column.

Note that this area extends upward contiguously to the next elevation scan where the temperature is -12 C.

There's one big area of about 2-3 dB at -12 C.

Then going up further, and we see the ZDR column extends upward and we have a few values > 1 dB in an area of high reflectivity on the west side of the storm.

This is probably the top of the ZDR column because it disappears above that.

So roughly -20 C is where we see the top of the ZDR column and that's indicating a strong updraft.

There's another area of high ZDRs further southwest but notice the CCs are very low, and this area is down radial from an intense reflectivity core and thus is a suspected three-body scatter spike, and not something real.

One thing to remember is that cannot explain every high ZDR value here. It's got to make sense. It should be associated with a storm like in the center part of this panel. But there are other areas of high ZDR that don't appear to be associated with a storm. There are reflectivities and high ZDRs going way above the melting level but it doesn't make sense.

A strong updraft should have a strong reflectivity core next to it, especially well above the melting layer, like the storm in the center.

So, again, this is not something (down here) I would call a ZDR column.

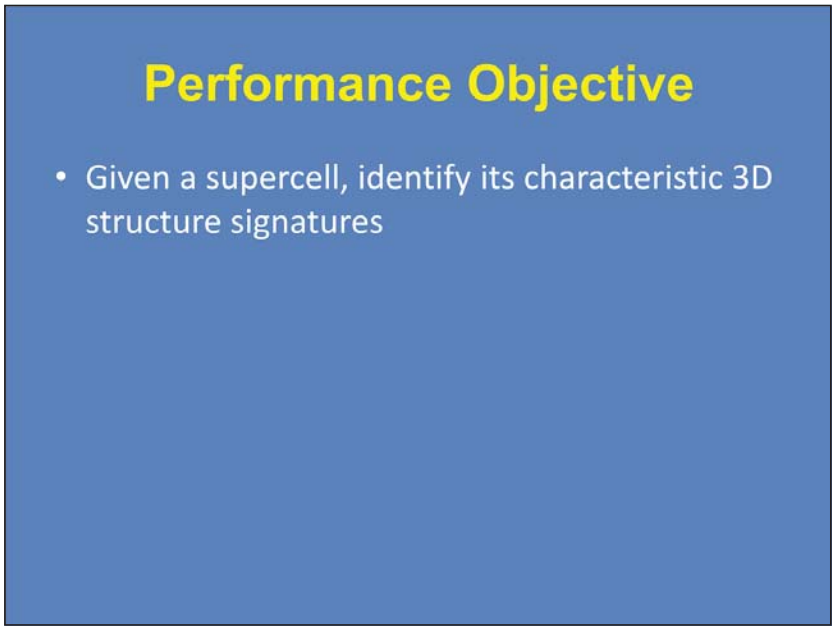
Summary

- ZDR column:
  - ZDR values > 1 dB in the context of a thunderstorm updraft above the environmental melting level
  - ZDR columns typically rise to the -10 C level
  - Strong updrafts push ZDR columns to -20 C
- Beware of ZDR column look-alikes
  - TBSS
  - Edges of precipitation areas



ZDR column: ZDR values > 1 dB in the context of a thunderstorm updraft above the environmental melting level, ZDR columns typically rise to the -10 C level, Strong updrafts push ZDR columns to -20 C

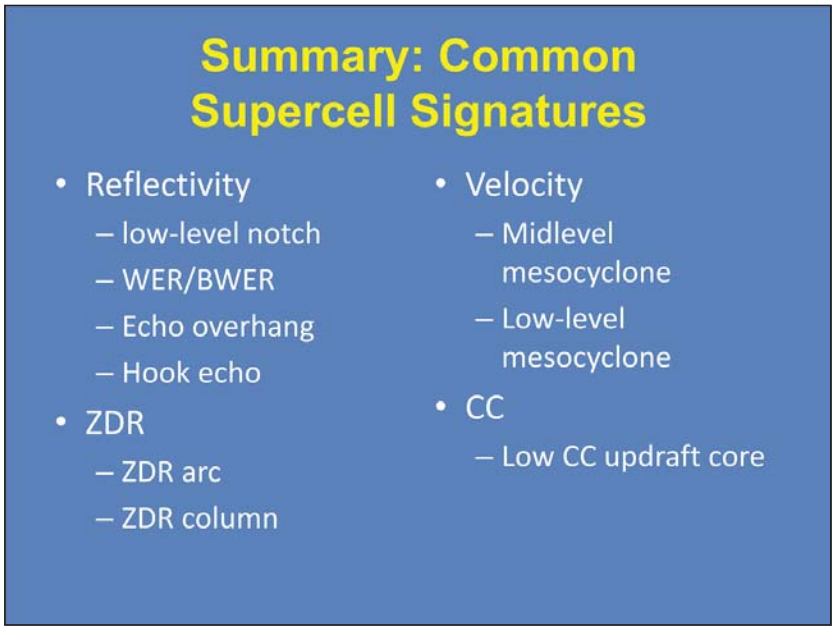
Beware of ZDR column look-alikes: TBSS, edges of precipitation areas



The objective of this drill: Given a supercell and WSR-88D data, identify its characteristic 3D structure signatures.



Please hit the play button to get a walk-through of identifying signatures often associated with a supercell.



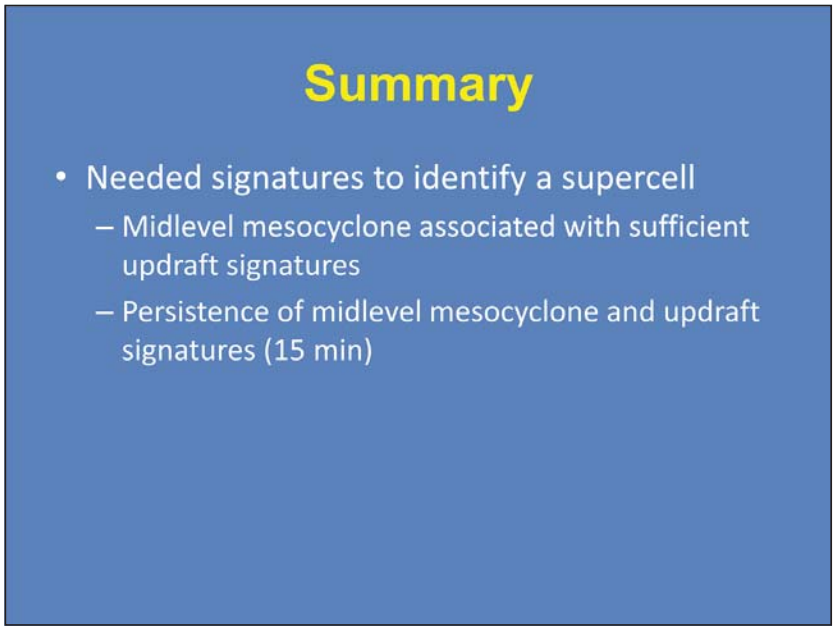
To summarize, I spent time identifying common and needed WSR-88D signatures to classify a supercell. Common supercell signatures include:

Reflectivity: low-level notch, WER/BWER, Echo overhang, Hook echo

ZDR: ZDR arc, ZDR column

Velocity: Midlevel mesocyclone, Low-level mesocyclone

CC: Low CC updraft core



But to really classify a supercell, you need to detect a midlevel mesocyclone that persists for at least 15 minutes.

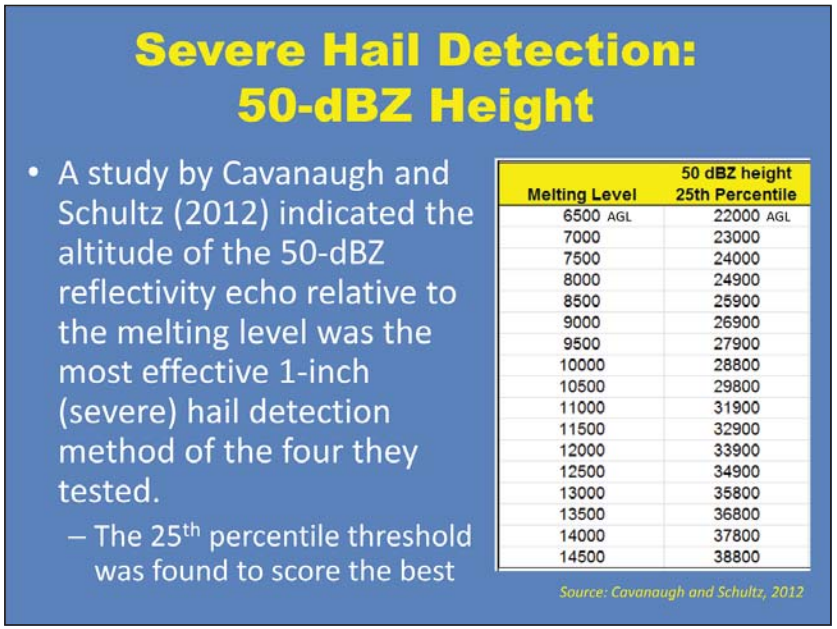


Welcome to the RAC Convective Storm Structure and Evolution, Applied Performance Drills: 50 dBZ Height.

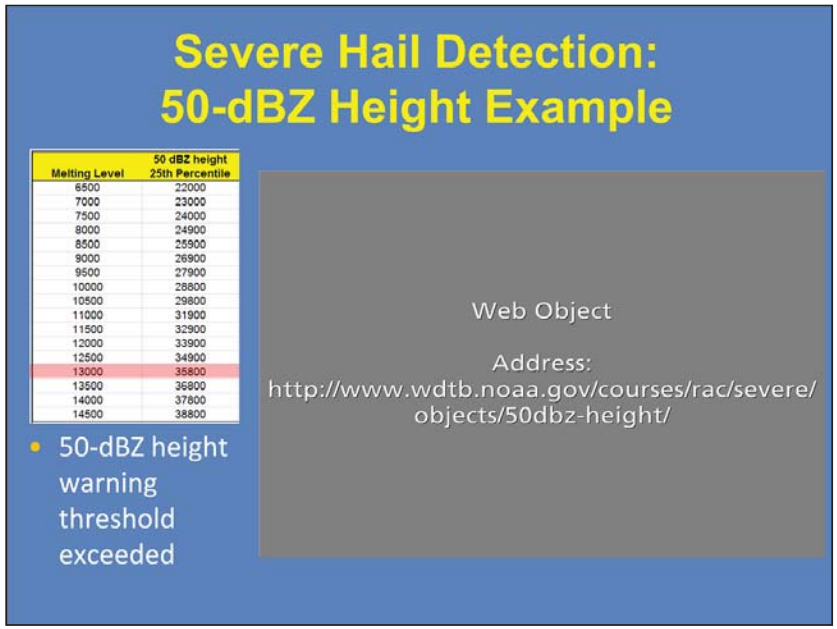
Performance Objective

Given a specified storm, identify the height of its 50 dBZ echo above the melting level.

Our performance objective for this module is: Given a specified storm, identify the height of its 50 dBZ echo above the melting level.



Building on previous research by Donavon (2007), a study by Cavanaugh and Schultz (2012) indicated the altitude of the 50-dBZ reflectivity echo relative to the melting level was the most effective 1-inch (severe) hail detection method of the four they tested. The 25th percentile threshold was found to score the best.



Let's use the results of the Cavanaugh and Schultz (2012) study to assess if this storm south of the La Crosse WSR-88D meets their criteria for a severe thunderstorm warning for hail. I've created a 14-frame, all-tilts reflectivity display with a LAPS freezing level height analysis overlay for our evaluation. The environmental melting level (denoted by the dashed purple isopleths) is 400 decameters or approximately 13,000 feet AGL. Our table on the left indicates that a 50-dBZ height of at least 35,800 feet AGL is necessary to meet their severe hail threshold.

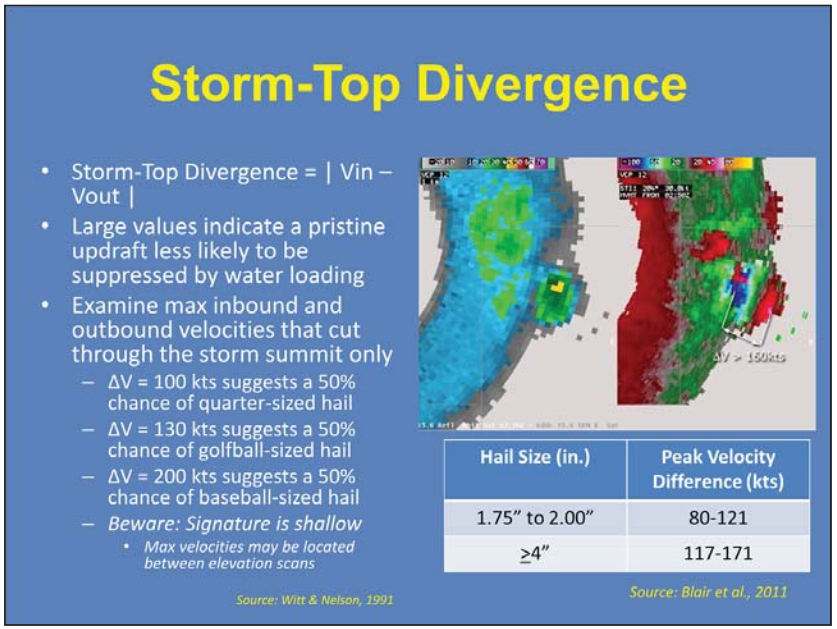
I turn sampling "on" and begin searching for 50-dBZ and greater echoes on the 0.5 degree elevation slice. I go to progressively higher tilts until I find the highest one with a 50-dBZ or greater echo. It's on the 6.4 degree tilt. Radar sampling indicates this is at an altitude of 40,093 feet which exceeds their suggested warning threshold by about 5,000 feet.



Welcome to the RAC Convective Storm Structure and Evolution, Applied Performance Drills: Storm Top Divergence.



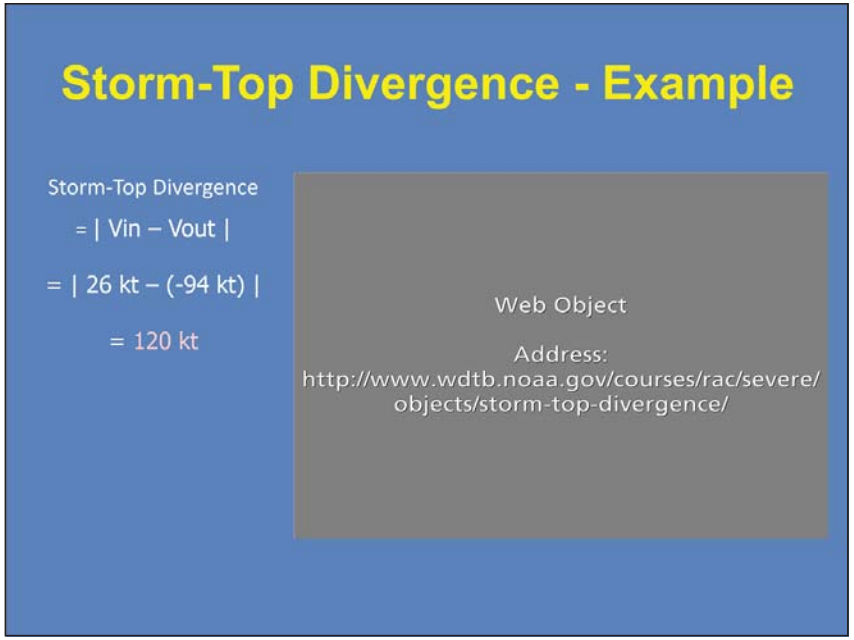
The performance objective for this lesson is: Given a specified storm, identify its storm-top divergence value.



Storm-Top Divergence is calculated as the absolute value of the maximum inbound velocity minus the maximum outbound velocity. Large values indicate a pristine updraft less likely to be suppressed by water loading. Check for scans that cut through the storm summit only, and examine the maximum inbound and outbound velocities. The presence of strongly divergent flow at storm top (|Vin - Vout| > 75 kts) indicates a strong updraft capable of producing at least marginally severe hail.

A study by Witt and Nelson in 1991 found that stop top divergence values equal to: 100 knots suggests a 50% chance of quarter-size hail, 130 knots suggests a 50% chance of golfball-sized hail, and 200 knots suggests a 50% chance of baseball-size hail. A study by Blair et al., 2011 indicated that golfball hail reports were generally associated with a delta-V of 80-121 knots, while giant hail (≥4") hail reports were generally associated with delta V of 117-171 knots.

Beware that the storm-top divergence signature is shallow and the true maximum velocity difference may be located between radar elevation slices.



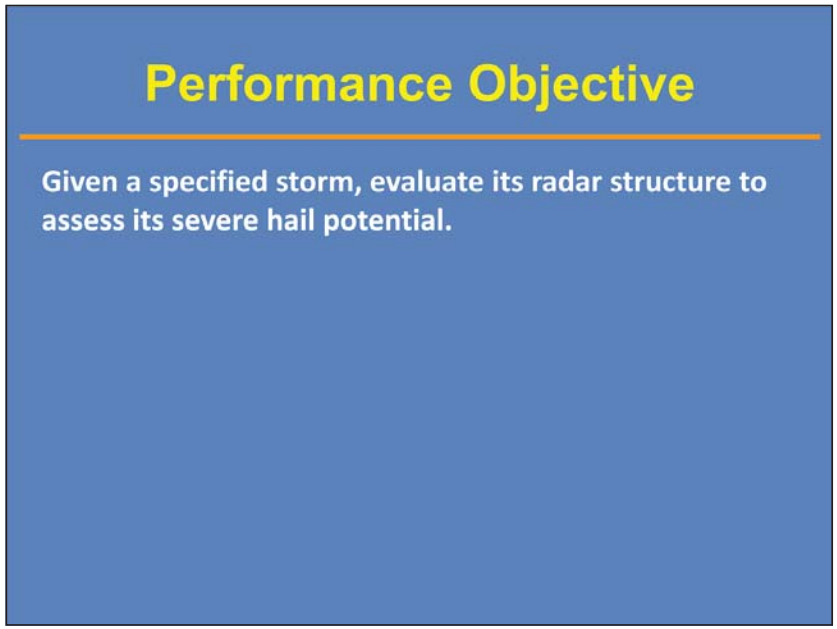
Let's identify the storm top divergence value of a storm southwest of the La Crosse, Wisconsin WSR-88D. I've created a 4-panel, all–tilts, base data display. I'll use the panel-combo rotate feature to evaluate the data.

I begin at the 0.5 degree Reflectivity product and move vertically upward until I'm at the updraft summit which in this case is between 12.5 and 15.6 degrees. I then switch over to velocity to find a suitable elevation slice that samples both the inbound and outbound values. Remember, the storm top divergence signature is shallow, so finding an elevation slice that both samples the storm summit and does not suffer from velocity dealising problems is often not possible. I find the 12.5 degree elevation slice meets these two requirements.

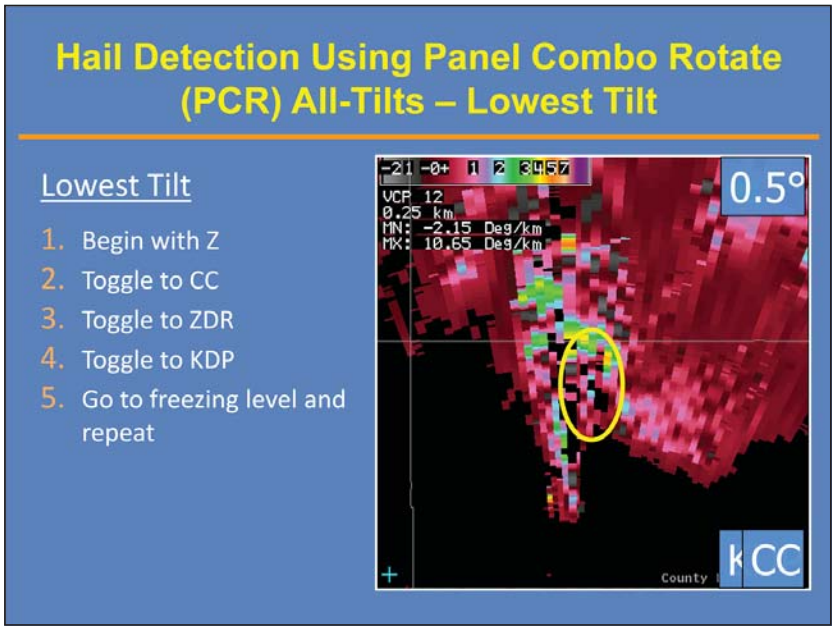
I turn sampling "on" and identify an outbound max velocity of 30 knots. The corresponding Reflectivity value is 28-dBZ which indicates we are not measuring clear air returns however, the corresponding Spectrum Width value of 16.5 knots is a bit high which places the validity of our velocity measurement at least somewhat in doubt. So let's search some more. I then find an outbound velocity value of 26 knots which is not much less than the 30 we measured before. The corresponding reflectivity value of 27.5 dBZ and a low spectrum width value of 6.8 knots gives me confidence this velocity value is very reliable, so I'll go with it. I then switch back to Velocity and identify a maximum inbound value of 94 knots. The corresponding reflectivity value of 54-dBZ and spectrum width value of 11.6 knots gives me confidence this velocity value is reliable. Summing these two velocities provides me with a storm-top divergence value of 120 knots, which suggests there is a greater than 50% chance the storm is producing severe (quarter-sized) hail.

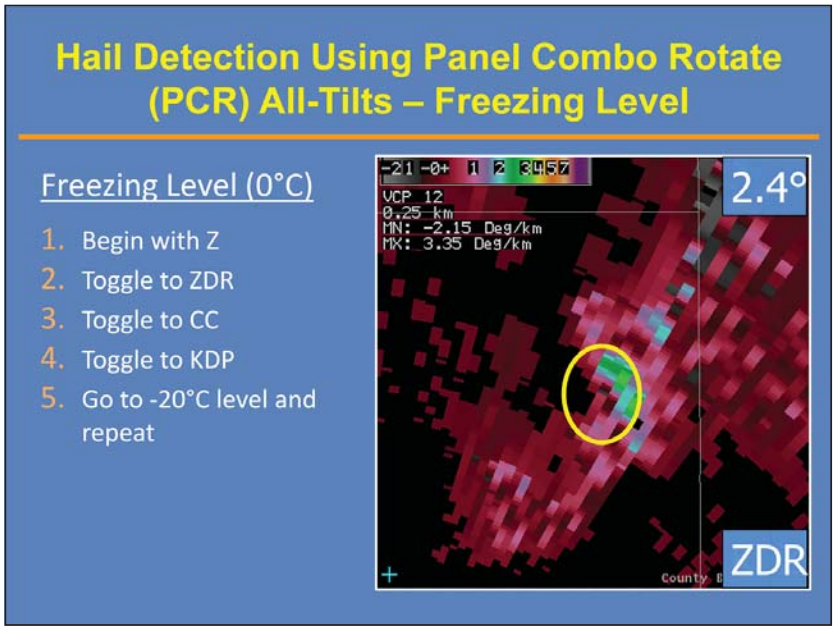


Welcome to the RAC Convective Storm Structure and Evolution, Applied Performance Drills: Severe Hail.



The objective for this lesson is: Given a specified storm, evaluate its radar structure to assess its severe hail potential.





The Advanced Weather Interactive Processing System (AWIPS) Panel Combo Rotate (PCR) All-Tilts is an efficient tool to examine a storm for hail. Start closest to the surface using the lowest tilt (0.5 degree slice in this example) because that's where we have information about the threat closest to the ground where the impacts are felt, or dented for that matter. But before using dual-pol products, start with what you know, reflectivity, and get an idea about storm structure and height of the 55 dBZ and 60 dBZ levels, and toggle back and forth to SRM to get a feel for any rotation and convergent signatures, both of which help point to a hail threat. Remember, strong deep shear helps to increase updraft strength, which is a precursor to hail growth. Try to use environmental sampling, in particular using temperature readout.

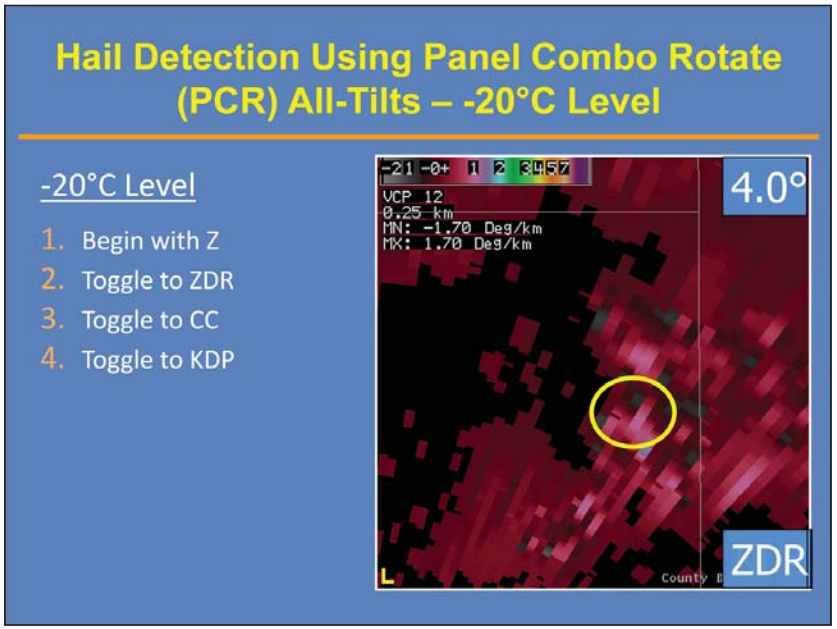
Once done with that, go to the 0.5 degree slice and toggle between Z and CC to find out where there is precip vs. non-precip. Where there is precip, find the highest reflectivity areas and look for CC less than 0.95 for hail. Then toggle to ZDR and look for local minima. If ZDR is low enough, less than 1.5 dB, you can say that hail is the dominant signal and that is the case here.

Finally, toggle to KDP to see if any rain is mixed in with the hail. Values greater than 0.5 deg/km indicate a good amount of liquid coincident with the hail. Keep in mind that radar resolution volumes where CC values drop below 0.90 are also areas where KDP values are not calculated. Thus, we see a number of those areas void of KDP values.

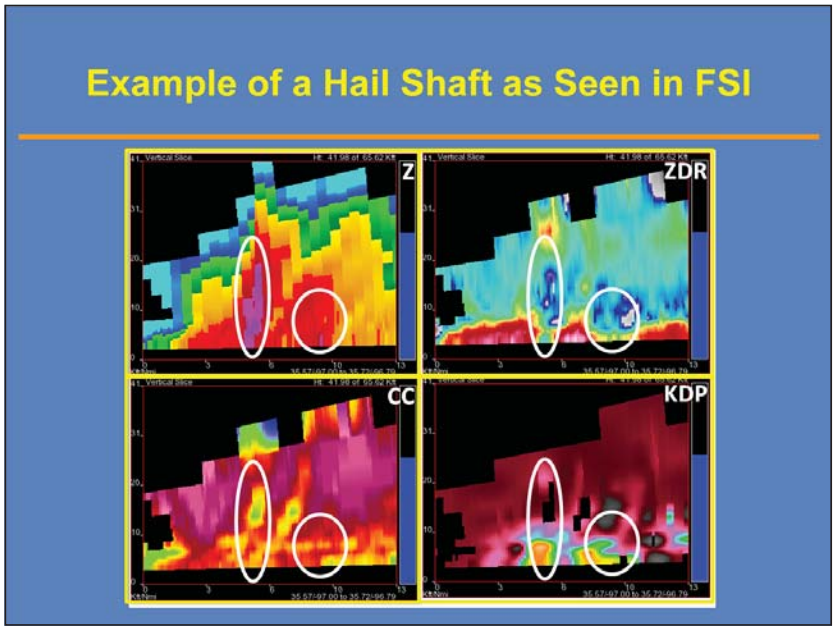
Next, go up to the freezing level and repeat.

The freezing level (0°C) was near the 2.4 degree slice in this example. High reflectivity is still present, and since you already know this has height continuity with the hail core at the 0.5 degree scan, you can toggle straight to ZDR to see what has changed at this height. Since this is right at the environmental melting layer, ZDR has dropped outside of the hail core to just over 0 dB, making it difficult to see the locally suppressed values of ZDR caused by the hail. Some slightly lower CC is still present, marking the location of the hail core. Looking at KDP, the eastern part of the hail core does contain fairly high amounts of liquid water.

Finally, go up to the -20°C level and repeat.



The -20ºC level was near the 4.0 degree slice in this example. There is still high reflectivity marking the location of the hail core, but ZDR is still low everywhere. CC has now become high everywhere due to the lack of mixed-phase hydrometeors at this elevation. KDP is pretty noisy and not greater than about 0.75 deg/km in the hail core. **Hail detection far aloft where no significant liquid is present becomes difficult because it looks the same as graupel to a dual-pol radar.** This is true for hail of any size unless you have a rare case of dry giant hail really high in a storm, which would reduce CC because of Mie scattering effects.



The AWIPS Four-Dimensional Stormcell Investigator (FSI) tool can be used for hail detection with dual-pol products. In the FSI cross section seen here, there are actually two hail shafts, marked by the white circles. As expected, both hail shafts have high reflectivity, although the one on the left, associated with a newer updraft, is deeper with higher values.

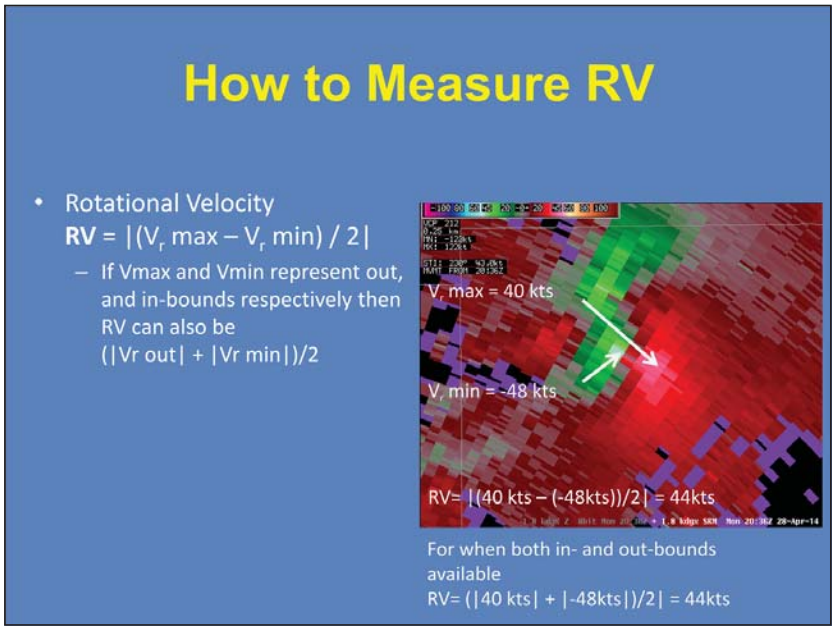
Taking a look at ZDR, both hail shafts are associated with fairly low positive to slightly negative values just above the melting layer, with the lower values extending higher aloft for the (younger) hail shaft on the left. Note also the depression in the transition to higher, positive values in each hail shaft. The signal in CC is not as easy to pick out as in the reflectivity or ZDR in this case, but notice the values of CC are generally low in both hail shafts. KDP is generally very high in both hail shafts reflecting increased liquid water content associated with melting hail and/or a rain/hail mixture.



Performance Objective

- Given a specified vortex signature and level, identify the rotational velocity (RV) including maximum RV and low-level Vr (LLRV).

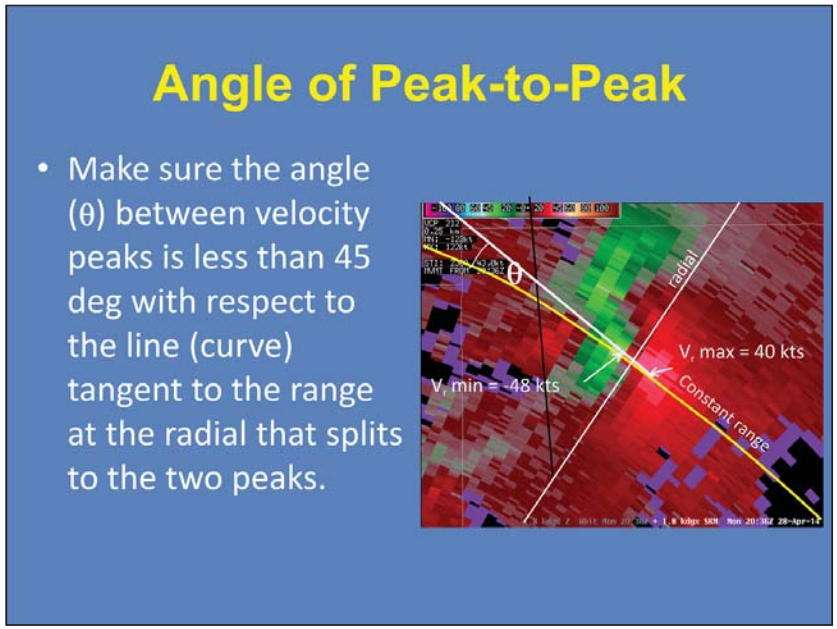
The objective for this lesson is: Given a specified vortex signature and level, identify the rotational velocity of (RV) including maximum RV and low-level RV (LLRV).



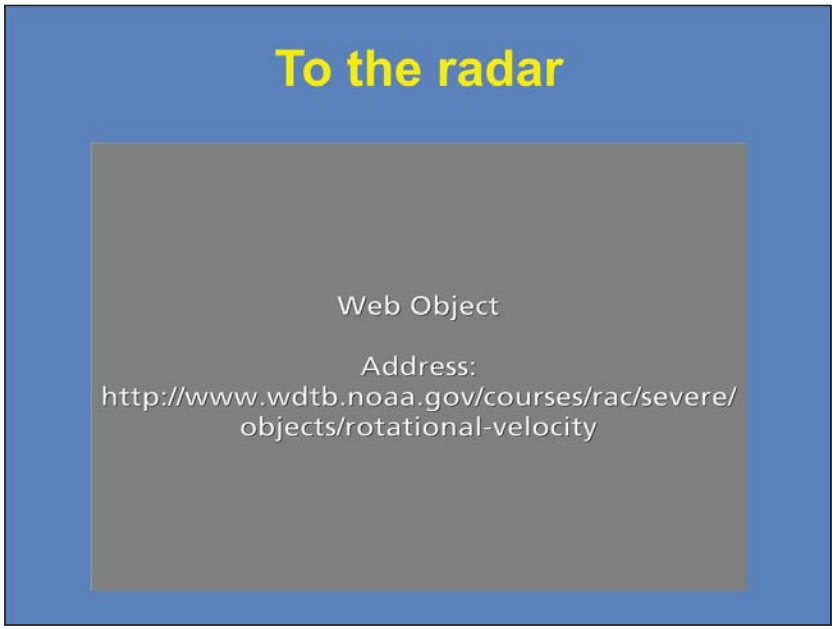
The inputs into calculating RV should represent the maximum and minimum velocities as illustrated by the inset in this figure. Note that the minimum and maximum velocities that contribute to the calculation of rotational velocity should be measured using representative peak values of the data levels in the velocity or SRM products.

In this example, the maximum Vr is 40 kts for a good representative value, and the minimum Vr is -48 kts. Taking the difference, dividing by two, and taking the absolute magnitude reveals a RV of 44 kts.

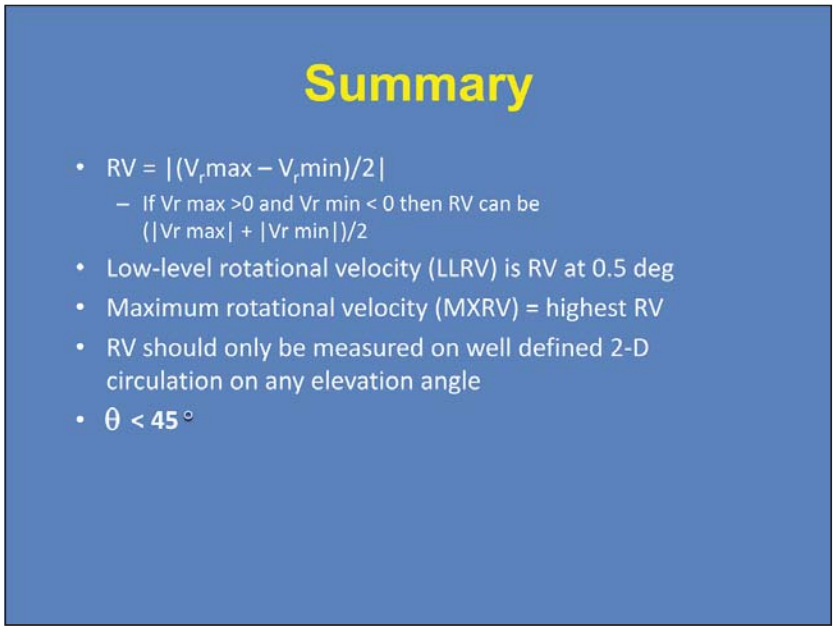
As it stands in this case, the Vr min is also an inbound velocity (Green) and Vr max is an outbound (red). Thus I can also take the average of the absolute magnitudes of these velocities and find RV. But as soon as Vr min and Vr max are both in-bound or out-bound then you have to go with the original formulation.



Make sure the angle between velocity peaks is less than 45 deg with respect to the line (curve) tangent to the range at the radial that splits to the two peaks.



Click the play button to get my walk-through of measuring RV using a real case.



To summarize, make sure RV is the absolute magnitude of the difference between the maximum and minimum Vr, divided by 2. If Vr max is > 0 and Vr min is < 0 then you can take the average of the absolute magnitudes of the velocities to get RV.

LLRV is the RV at the lowest elevation scan.

MXRV is the maximum RV in the mesocyclone

And make sure RV is measured with a well defined 2-D circulation on any elevation angle

Finally, Vr max and Vr min should be less than 45 deg from each other with respect to the line tangent to the constant range curve perpendicular to the radial between the two velocity peaks.

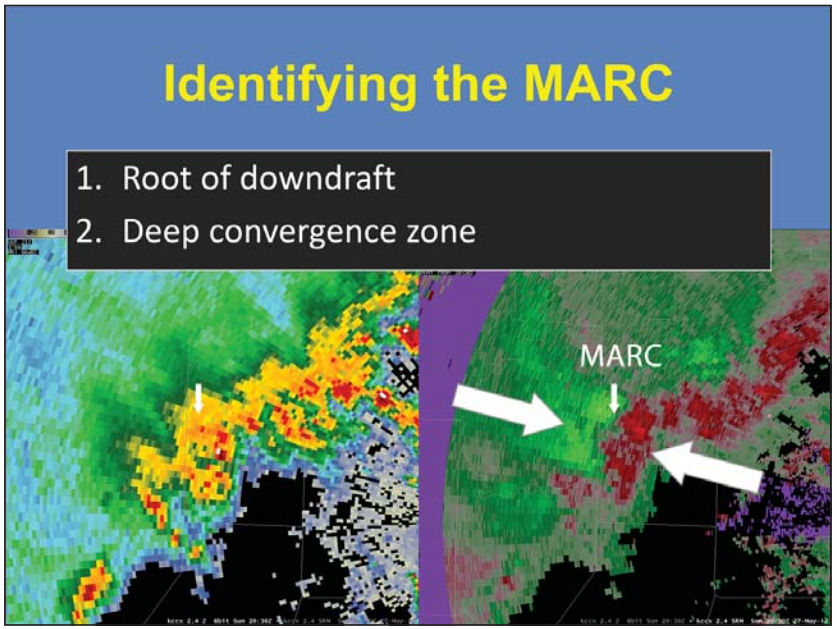


Welcome to the Applied Performance Drills lesson on identifying the MARC

Performance Objective

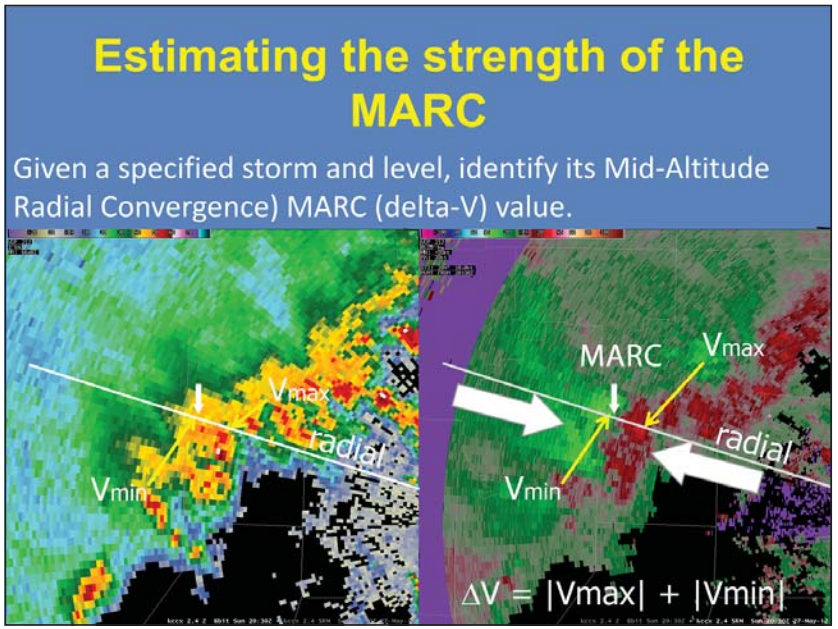
- Given a specified storm and level, identify its Mid-Altitude Radial Convergence) MARC (delta-V) value.

Given a specified storm and level, identify its Mid-Altitude Radial Convergence (MARC) (delta-V) value.



The MARC signature can be identified by radar as a localized area of elevated strong radial convergence collocated with strong reflectivity. There are two phenomena for which the MARC is associated, both are high wind precursor signatures:

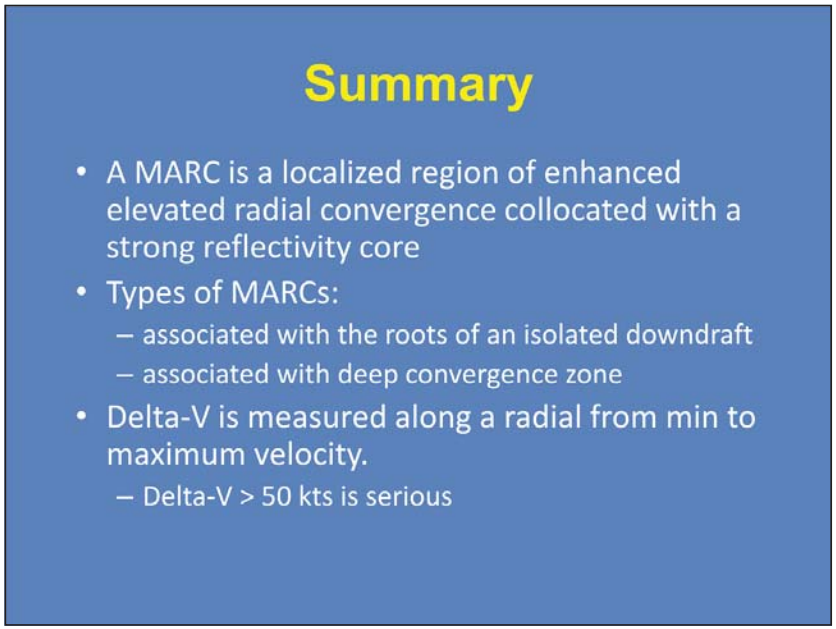
1. For an isolated cell downburst, the MARC is separated from the convergence along the gust front and represents the root of the downdraft.
2. But for QLCS, HP supercells, the MARC is part of, and connected to, a deep gust front and is a deep convergence zone.



Estimating the strength of a MARC signature can inform you of a storm's potential for producing severe winds. For a MARC, Vmin should be further away from Vmax relative to the RDA along the same radial. To calculate the strength of a MARC, find a relatively representative, isolated Vmax and Vmin oriented in a radial direction from each other. If no isolated velocity cores exist, pick a representative radial velocity that appear to be the largest values that appear involved in the convergence. The area between Vmax and Vmin should be in close proximity to the precipitation core, if not in it. Then add the absolute magnitudes of the peak velocities (Vmax and Vmin). That's your Delta-V.



Click play to watch me walk through identifying the MARC and estimating its strength.



--A MARC is a localized region of enhanced elevated radial convergence colocated with a strong reflectivity core

--Types: associated with the roots of an isolated downdraft, associated with deep convergence zone

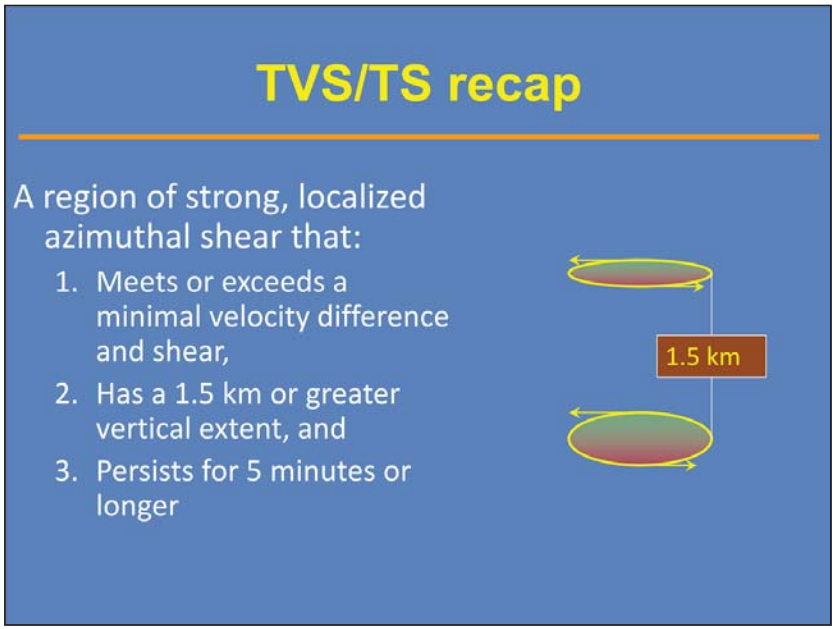
--Delta-V is measured along a radial from min to maximum velocity. Delta-V > 50 kts is serious.



Performance Objective

- Given a specified storm, identify its Tornado Vortex Signature (TVS) or Tornado Signature (TS)

Given a specified storm, identify its Tornado Vortex Signature (TVS) or Tornado Signature (TS).

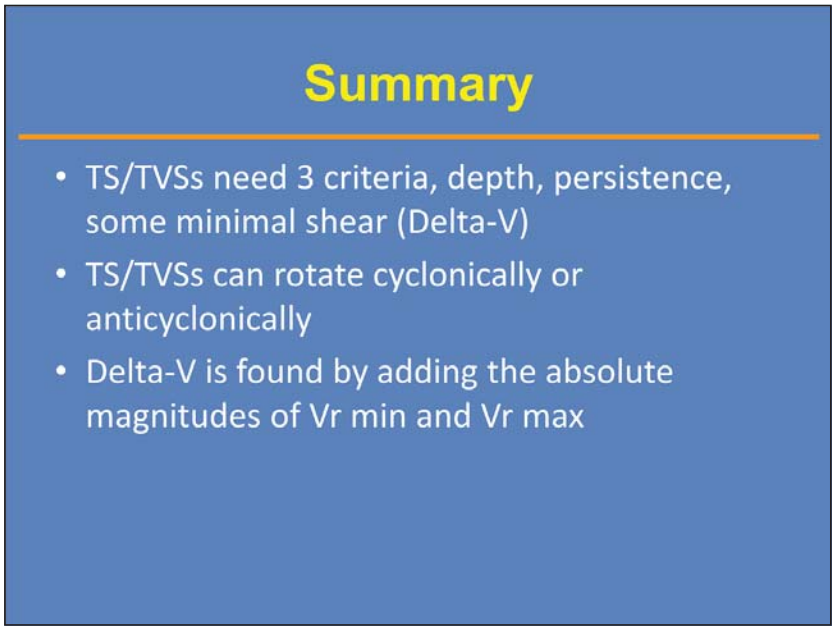


Some criteria are required to make sure that what circulation you are looking at is indeed a TVS. The three criteria are:

1. **A minimal shear:** There is no hard lower threshold in this criterion in the same way as an algorithm like the TDA would be assigned one. The minimal shear that an expert forecaster would define would depend on many things including the distance from the radar, the forecaster's assessment of the size of the vortex, near storm environment, and past experience. We will discuss what meaningful shears may be of significance.

2. **Vertical Extent:** At least some vertical continuity should be seen in a TS/TVS so that there is a high probability that an updraft is present in the circulation. For most events, the depth should be at least 1500 m (4900 ft.). Low topped supercells typically do not have deep TSs/TVSs, even if tornadic. Sometimes, and if detectable at all, only the lowest elevation angle contains a gate-to-gate rotational signature in tornadic low-topped supercells. Sometimes the vortex may appear as a TS and a TVS at adjacent elevation angles due to the vagaries of sampling and vortex structure. The vertical extent should include one or both manifestations as long as the true vortex appears to show vertical continuity.

3. **Persistence: In order to reduce the possibility of a circulation that randomly becomes vertically coordinated, you should ensure that the TS/TVS persists for at least five minutes**. However, mesocyclones can spin-up over a considerable depth in a very short time, and some legitimate TS/TVSs may become tornadic in less time. We suggest that if either signature forms in close proximity to a strong updraft signature, and a very supportive environment, persistence may not be a requirement to call it a TS/TVS and a tornado.



TS/TVSs need 3 criteria, depth, persistence, some minimal shear (Delta-V)
TS/TVSs can rotate cyclonically or anticyclonically
Delta-V is found by adding the absolute magnitudes of Vr min and Vr max



Welcome to the applied performance drills, identifying a Tornado Debris Signature. This drill lasts 8 minutes.

Performance Objective

- Given a specified storm, identify its Dual-Pol Tornado Debris Signature (TDS).

Given a specified storm, identify its dual-pol Tornado Debris Signature (TDS).

TDS Criteria

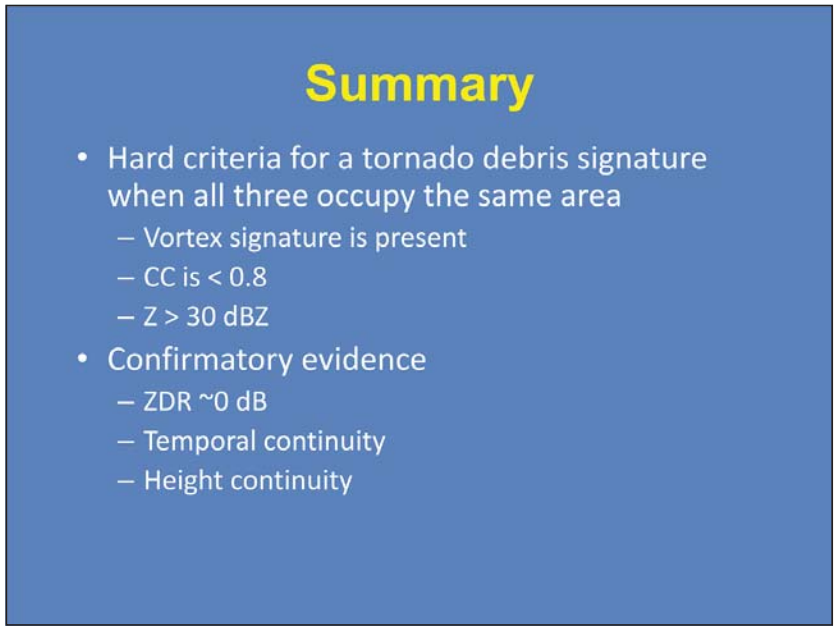
Given a specified storm, identify its dual-pol Tornado Debris Signature (TDS).

	product	strategy	Verdict?
			Yes or no
1	Velocity	Is azimuthal couplet present?	
2	CC	Are values < 0.8 in or near the velocity couplet?	
3	Reflectivity	Are values > 30 dBZ where CC < 0.8?	
4	ZDR	Values near zero help confirm above 3 criteria	
5	verdict		

Three signatures are absolutely needed for identifying a Dual-pol Tornado Debris Signature (TDS) because there are a lot of fakes out there. First, an azimuthal velocity couplet needs to be present – a vortex. It doesn't matter the sign of rotation, just that one is there. Within the bounds of the vortex, the CC product should show gates < 0.8. And for those gates, the reflectivity should be > 30 dBZ. To help confirm that these gates are candidates for a TDS, the ZDR should also be near zero dB.



Please click on this video to watch me evaluate two potential Tornado Debris Signatures.



Hard criteria for a tornado debris signature when all three occupy the same area: Vortex signature is present, CC is < 0.8, Z > 30 dBZ

Confirmatory evidence: ZDR ~0 dB, Temporal continuity, Height continuity

TAB

Radar & Applications Course

Topic:

Flash Floods

Table of Contents
Topic: Flash Floods

Click to jump to lesson

Lesson 1	**Flash Flood Meteorology**
Lesson 2	**Flash Flood Hydrology**
Lesson 3	**High-Resolution Precipitation Estimator (HPE) and Bias HPE**
Lesson 4	**Flash Flood Warning Decision Making (Part 1): Choosing Your Precip and Guidance Sources**
Lesson 5	**Flash Flood Warning Decision Making (Part 2): Using FFMP**



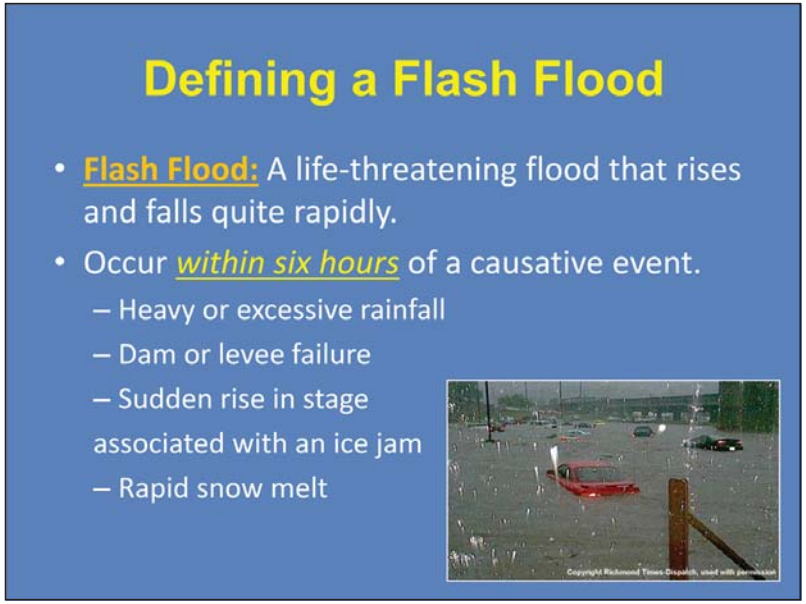
Hi, my name is Jill Hardy and welcome to this lesson on flash flood meteorology.

We have a guest speaker for this lesson: Steve Martinaitis of OU CIMMS at NSSL. But if you have any questions regarding the material, please feel free to contact me, or the RAC team. Our contact information will be on the next slide.

Lesson Objectives

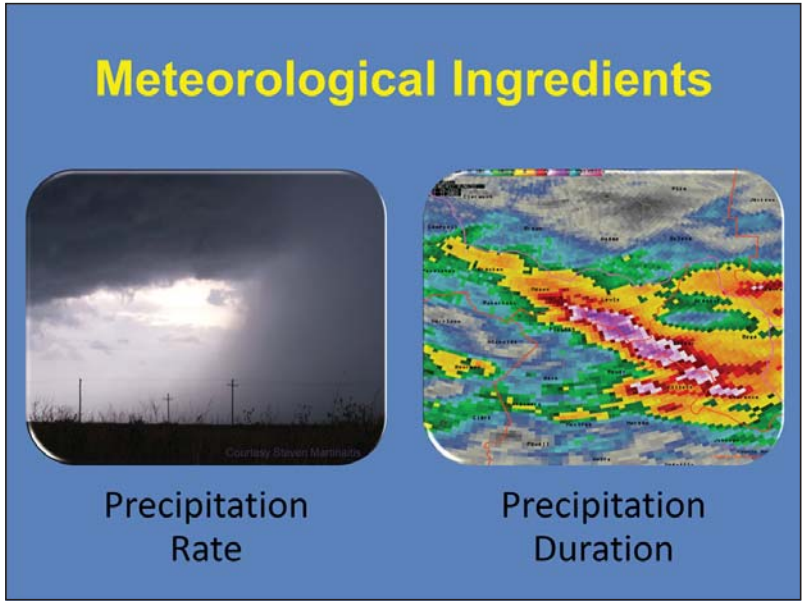
- Identify the mesoscale and storm-scale variables that contribute to the flash flood potential
  - Precipitation Rate/Efficiency
  - Precipitation Duration
- Identify heavy rainfall using WSR-88D and Dual-Polarization technology

There are two main objectives with this lesson. The first is to identify the variables related to precipitation rate and duration that contribute to the flash flood potential at a meso-scale and storm-scale levels. The second objective is to identify rainfall signatures using the WSR-88D and the new dual-polarization technology.



To guarantee that we are all on the same page, I want to make sure we understand how a flash flood is defined. Basically, it is a life-threatening flood that occurs quite quickly, i.e., within a six hour period. Flash floods can occur from a variety of events, such as heavy rainfall, dam failures, ice jams, or rapid snow melt.

For the purpose of these lessons, we will focus on flash flood events related to heavy rainfall.



When it comes to the meteorological aspects of flash flooding, the two most important things to consider are the precipitation rate and the precipitation duration. Let's focus on the factors that influence the rate first.

Precipitation Rate Factors

- Updraft Strength
- Liquid Water Content of Air Entering the Updraft

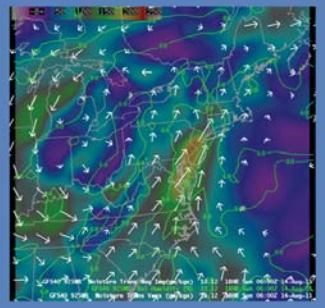
Upward Moisture Flux

- Precipitation Efficiency

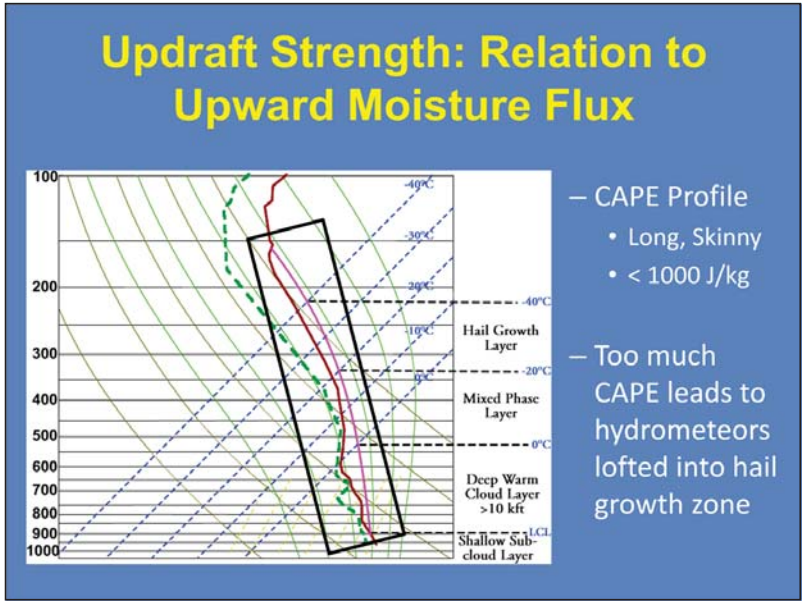
There are several factors that help determine the precipitation rate. The updraft strength and the liquid water content of the air that is entering the updraft contribute to the upward moisture flux into a storm. The percentage of that moisture flux that returns to earth as precipitation characterizes the precipitation efficiency of the storm.

Precipitation Efficiency

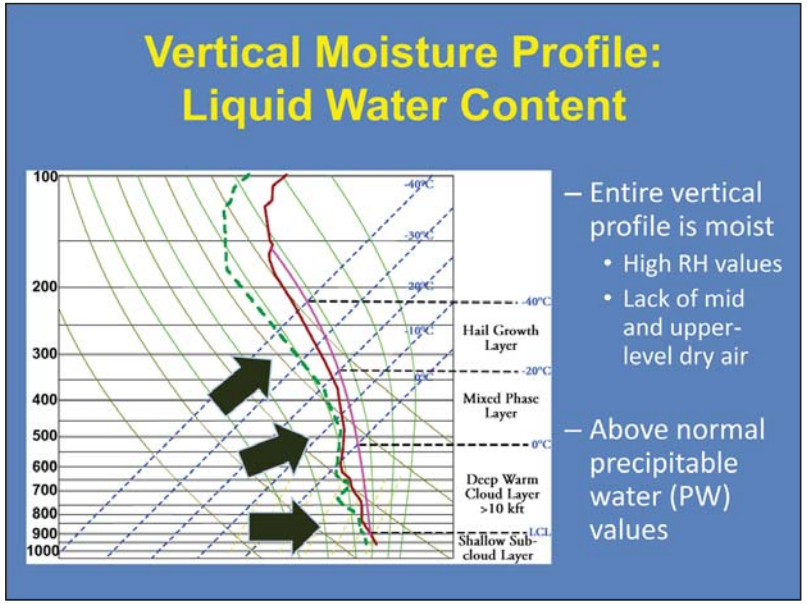
- Definition: Fraction of total moisture ingested by the updraft that falls back to the ground
- Dependent upon…
  - Updraft Strength
  - Moisture Profile
  - Warm Cloud Layer
  - Cloud Seeding



As you see here on the slide, precipitation efficiency is defined as the fraction of total moisture ingested by the updraft that returns as precipitation. Precipitation efficiency cannot be quantified in real time, so you will need to examine a number of factors to infer an efficiency. These factors include the updraft strength, the vertical moisture profile of the atmosphere, the depth of the warm cloud layer, and cloud seeding.

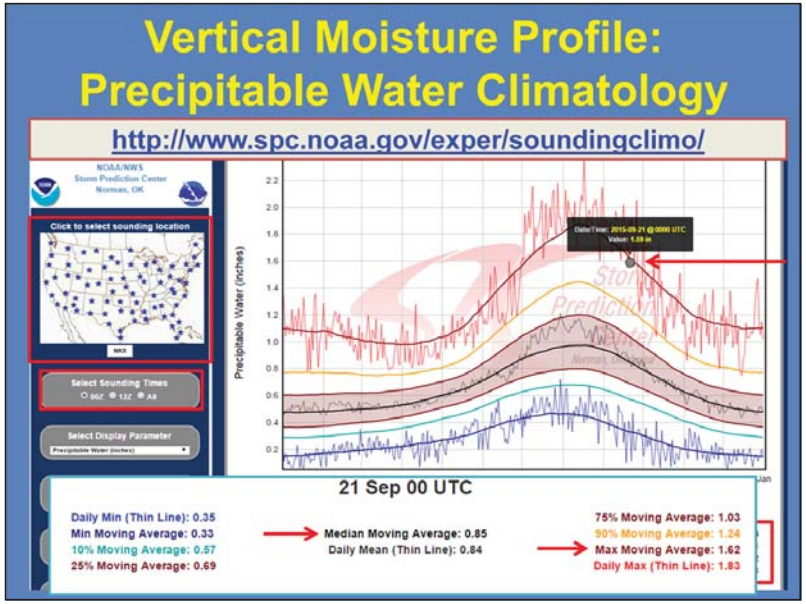


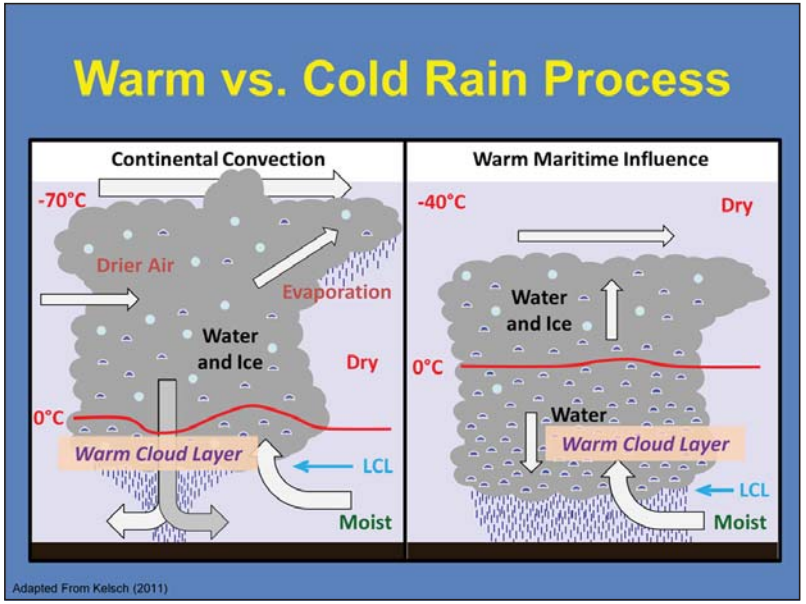
Let's start with the strength of the updraft. Shown here is the average atmospheric profile of flood and flash flood events. When considering the updraft strength of a convective storm, you would want to see a long and skinny CAPE profile. The amount of CAPE in the atmosphere should be under 1000 J/kg. Larger CAPE values will loft the hydrometeors ingested by the updraft into the hail growth zone.



Now looking at the temperature and dew point profile of this sounding, you can see that it is very moist at all levels. Notice how there is a lack of dry air at the mid and upper levels. This is important when you consider the depth of water within a column of the atmosphere if all the water were precipitated as rain, otherwise known as the precipitable water (PW) value. Seeing above normal PW values is a good indication of how moist the atmosphere is.

So how do you determine what is an above normal PW value?





I'm going to briefly hop in here, as there's been a change since this lesson was created. Many of you may have been familiar with Matt Bunkers' Precipitable Water Climatology page. However, the SPC now hosts the point sounding climatologies, similar to the previous website.

Using the SPC site, comparing model or observed precipitable water (PW) values is quite easy. Begin by navigating to the desired sounding location, and select a sounding time. The plot now shows the daily minimum, several moving percentile averages, as well as the daily maximum for each day of the year.

Let's use this 00Z sounding climatology plot for the KNKX radar near San Diego. When I overlay the latest sounding information, we see the current value is 1.59 inches. This is near the maximum moving average of 1.62 inches for this day. Historically on this day, the median PW is 0.85 inches, so we are quite a bit higher than that.

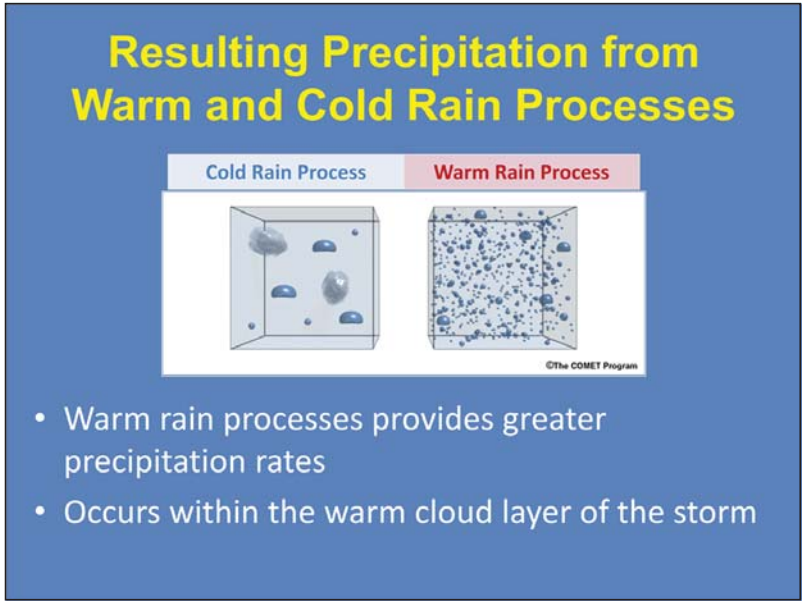
Heavy precipitation events that lead to flooding and flash flooding have values that are above the 75th percentile and usually approach the 90th or maximum moving averages. In fact, for this example, the San Diego WFO had a flash flood watch in effect for the majority of their CWA. Use the URL to access the PW climatology page.

Alright, back to Steve!

Now that we have analyzed the CAPE and moisture profile of the atmosphere, we can see how it helps determine whether a warm rain or a cold rain process is the predominant precipitation production method. Recall that precipitation forms through collisions and coalescence within a warm rain process while deposition and the Bergeron Process (the collision of ice crystals) define a cold rain process.

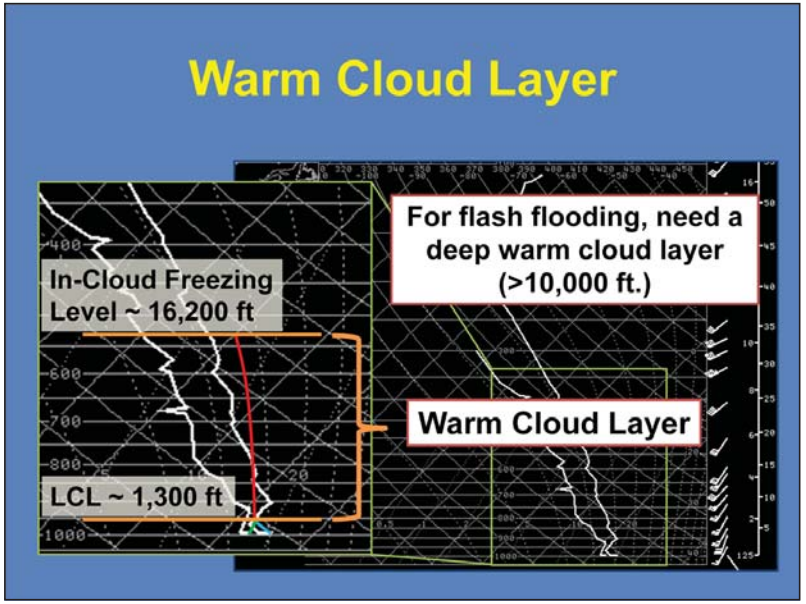
Looking at convection derived from a continental airmass, you can see that the LCL is relatively high while the in-cloud freezing level is quite shallow. The vertical separation between the LCL and freezing level is defined as the warm cloud layer. This is where warm rain processes occur. However, the warm cloud layer is generally not very deep with this type of convection. Within a strong CAPE environment, hydrometeors will be lofted beyond the warm cloud layer, where they will become frozen (resulting in the formation of hail) and become subjected to evaporation due to mid and upper level dry air entrainment. This region is where the majority of the hydrometeors undergo cold rain processes.

Now focusing on the convection influenced by a warm maritime airmass, you notice that the LCLs are relatively low, and the in-cloud freezing level is much higher. Therefore, you have a greater warm cloud layer. The weak CAPE profile allows for the majority of the hydrometeors to remain below the freezing level. The moist vertical profile also helps in diminishing the effects of evaporation and dry air entrainment. Here, warm rain processes will dominate precipitation production.



Comparing the resulting precipitation at the surface, you can see the dominant cold rain processes from the continental airmass yields a small quantity of rain drops that are generally large in size and can also include hail stones. The dominant warm rain processes in the maritime airmass has a substantial quantity of raindrops. So, the warm rain process results in a greater precipitation efficiency and greater precipitation rates.

In the example on the previous slide, you saw how the CAPE and moisture profiles influence the amount of hydrometeors that reside in the warm cloud layer, and thus, could undergo warm rain processes. Which leads to the next set of questions… How do we calculate the warm cloud layer? And how deep of a warm cloud layer do you need for precipitation rates that could potentially yield flooding?



To calculate the warm cloud layer, we will use this sounding from Grey, ME during a summer-time convective event that had some flash flooding. Focusing on the lower and mid levels of the atmosphere, you would start by finding the LCL. Follow the dry adiabatic lapse rate from the surface temperature and the saturation mixing ratio from the surface dewpoint until the two lines meet. The LCL for this case is about 1,300 feet.

From the LCL, follow the moist adiabat up to the freezing level. We choose the moist adiabat because that should be “in-cloud” and also where the warm rain process (collisions and coalesence) is occurring. In this case, the freezing level is around 16,200 feet. The difference in height between the LCL and in-cloud freezing level will be our warm cloud layer. Having a *deep* warm cloud layer is very important for flash flood forecasting. A warm cloud layer over 10,000 feet is considered deep.

For this example, our warm cloud layer is approximately 14,900 feet.

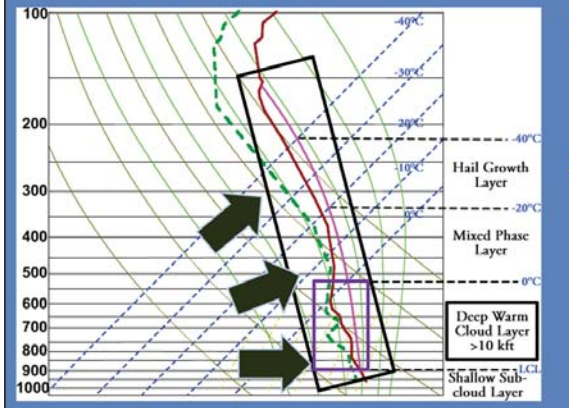
Cloud Seeding

- Definition: Jump start of precipitation product via the ingesting of hydrometeor embryos
- Inter-Storm Seeding
  - Increases upward Moisture flux
  - Increases environmental humidity



Another process that increases precipitation production is cloud seeding. We will focus on inter-cloud seeding here. This is the process where precipitation production is jump started by the updrafts ingesting hydrometeors from other storms. This will help increase the upward moisture flux and increase the local environmental humidity. In this example, an intense rain band forms with the remnants of Tropical Storm Hermine over central Texas. The combination of the tropical environment and the inter-cloud seeding enhanced rainfall production in an already efficient precipitation environment. Widespread rainfall totals of 6-10 inches were common with system.

Review: Precipitation Rate Variables



Upward Moisture Flux

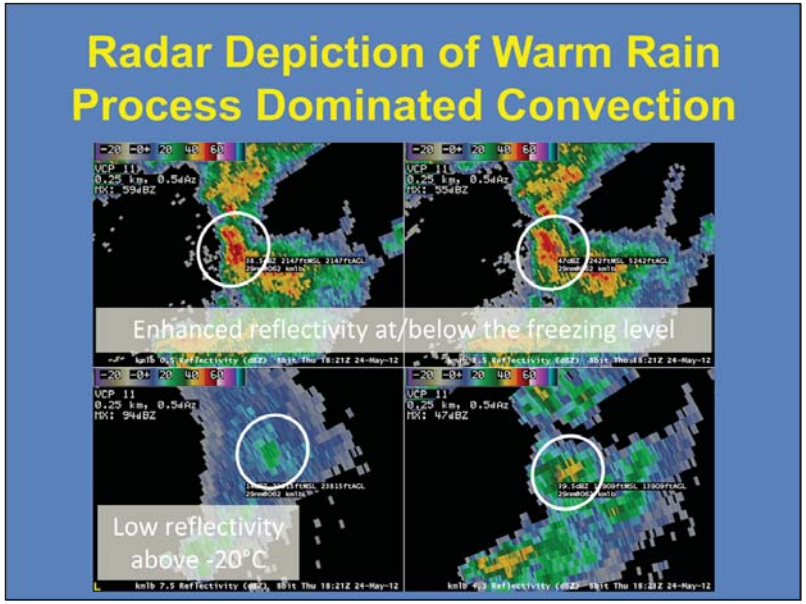
- Modest CAPE
- Moist vertical profile

Precipitation Efficiency

- Upward moisture flux
- Warm rain processes
- Deep warm cloud layer
- Cloud seeding

Let's quickly recap what influences the precipitation rate. First, you have to consider the variables that go into the upward moisture flux of a convective storm, such as a modest CAPE profile, generally under 1000 J/kg, and a moist vertical atmospheric profile.

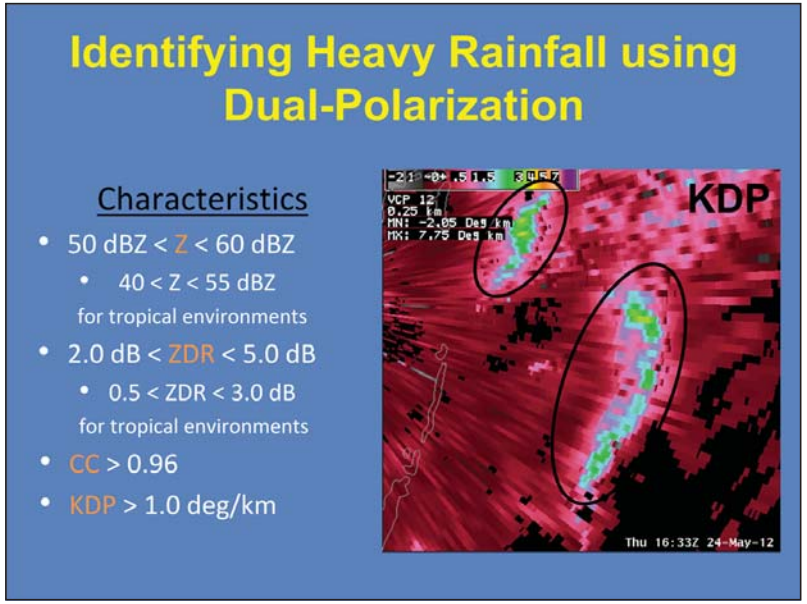
The fraction of the upward moisture flux that becomes precipitation defines the precipitation efficiency of the storm. Along with the upward moisture flux, recall that warm rain processes provide greater precipitation rates. Warm rain processes occur within the warm cloud layer. Remember that for a greater flash flood potential, you would like to a deep warm cloud layer of 10,000 feet or greater. You also have to consider inter-cloud seeding to increase precipitation production.



So what does convection dominated by the warm rain process look like on radar? Honestly, not much unless you know what you are looking for. The two main characteristics of convection dominated by the warm rain process are enhanced reflectivity values at or below the freezing level and low to non-existent reflectivity values above the -20°C level.

Using this example from the Melbourne, FL office, the top two images show reflectivity values between 50-60 dBZ below the freezing level. In the bottom-right panel, the 4.3° tilt scans the storm at 13,900 feet, just a few hundred feet below the freezing level. Here, there are very few pixels that meet or exceed 40 dBZ. The bottom-left panel shows the storm at the 7.5° tilt near the -20°C level. Reflectivity values here are below 25 dBZ. The storm does not exist on higher tilts.

This series of images shows what is called a low-echo centroid signature. This is where the majority of the precipitation core lies below the freezing level. The combination of this type of radar signature and a moist, slightly unstable environment should clue you in to warm rain processes being dominant here.



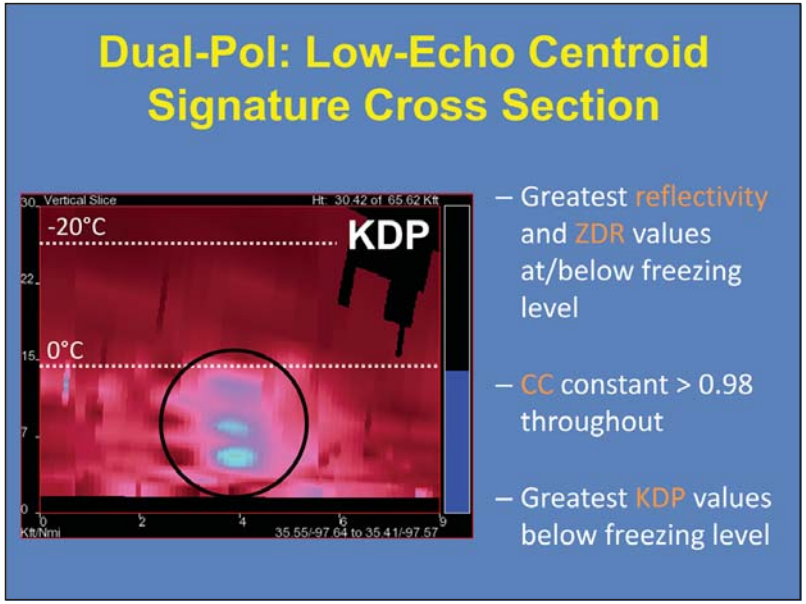
With the addition of dual-polarization technology, the new algorithms can help pinpoint areas of greater precipitation rates. This signature is from the Miami, FL radar and was related to a tropical disturbance that eventually became Tropical Storm Beryl.

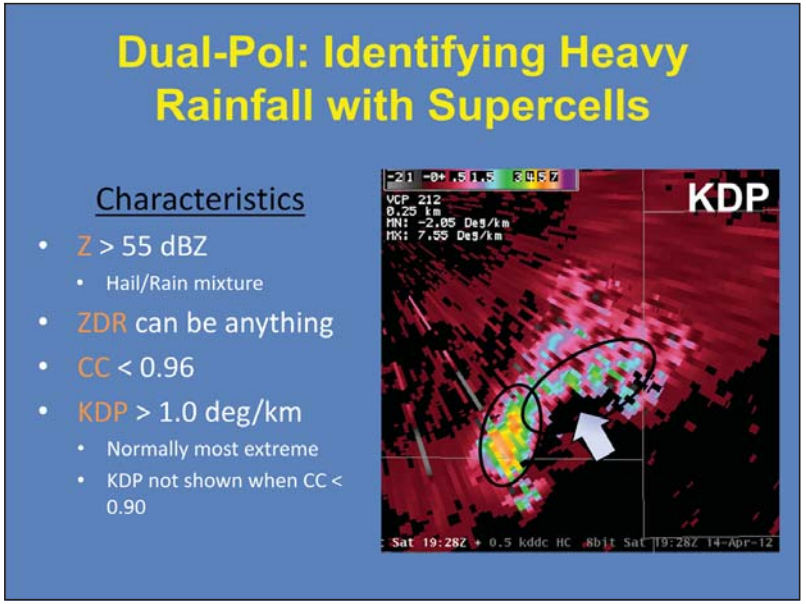
Starting with the reflectivity (Z), you would look for areas of enhanced values, generally in the 50-60 dBZ range; 40-55 dBZ for tropical environments. Here, we are highlighting two areas of enhanced values.

Now examining the differential reflectivity (ZDR), the difference between the horizontal and vertical reflectivity factors, you would look for ZDR values between 2.0 and 5.0 dB, and 0.5-3.0 dB for tropical environments. Remember, there is a strong relationship between the raindrop size and ZDR where the greater the ZDR values, the larger the raindrop diameter. Since we are dealing with a tropical environment in this case, the ZDR values are around 1.5 dB. Combine that with the high reflectivity values, you have a lot of small rain drops here.

Moving on to the correlation coefficient (CC), you should see very high values (above 0.96). This means that the type of precipitation that is being sampled is uniform. As you can see here, the areas that had the greater reflectivity have a CC of around 0.99, meaning all the precipitation here is rain.

Finally, values of the specific differential phase (KDP) should be above 1.0 deg/km. Higher KDP values can mean larger rain drops or a larger concentration of rain drops. Since we know this is a tropical environment and the ZDR values suggest small rain drops, then this means we are dealing with a larger concentration of rain drops, and thus, greater precipitation rates.





Now that you have seen what a warm rain process dominated storm looks like with base reflectivity and with the dual-pol products, we will now take a cross-section through a low-echo centroid signature. This example will look at a specific storm that was part of a system that produced significant flash flooding on the north side of Oklahoma City, OK.

As you saw in the four panel image earlier in the presentation, most of the enhanced reflectivity values lie at or below the freezing level and low reflectivity values exist near and above the -20°C level. The greater ZDR values, which represent rain drops, also lie below the freezing level. The very low ZDR values above the freezing level can represent very small water droplets, ice crystals, and/or hail.

The CC values are constant throughout the vertical profiles with them ranging from 0.98 to 0.995. The values closer to 0.98 (the darker purple shading with a slight orange tint) represent all rain with slightly larger drop sizes. Finally, the greater KDP values exist below the freezing level, showing where the greatest concentration of rain drops are occurring.

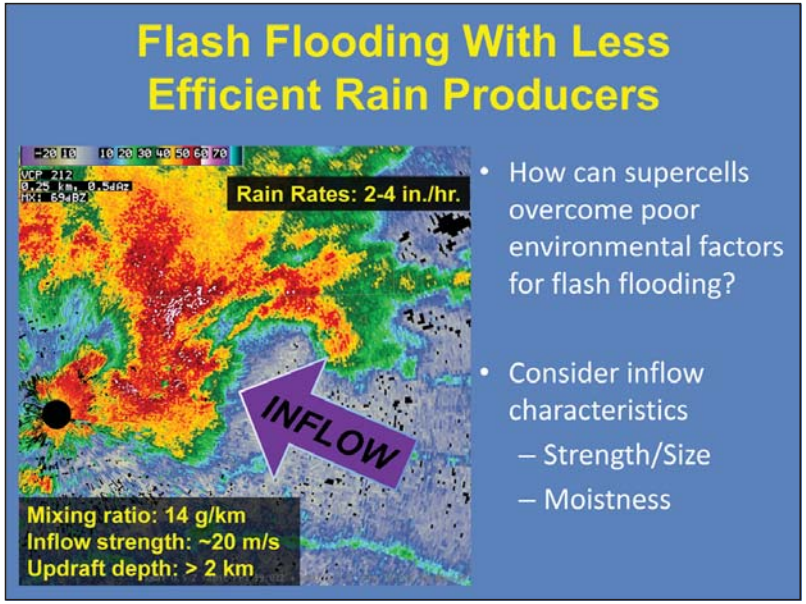
Since we have looked at what an efficient rainfall producer would look like with radar, let's take a look at what an inefficient storm would look like in dual-pol. For this example, we will use a supercell viewed from the Dodge City, KS office during the April 14, 2012 outbreak. Supercells can produce heavy rainfall, but you would need to examine the characteristics and motion of the storm to determine its flash flood potential.

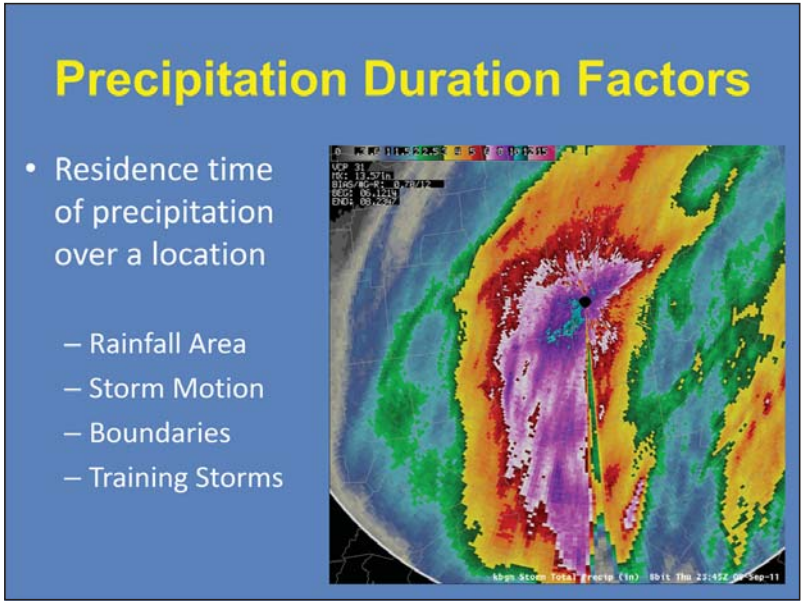
Starting again with reflectivity (Z), you would look for areas that are greater than 55 dBZ. Here, we have highlighted two separate areas within this supercell. These areas of enhanced values are probably areas of hail/rain mixture.

Now starting with the dual-pol products and differential reflectivity (ZDR), it should be noted that the ZDR values can be anything because of hail contamination. Severe hail can bring ZDR values to near 0 dB while water coated hail can have values up to 6 dB.

Since we are dealing with non-uniform precipitation types, correlation coefficient values will be below 0.96 in areas of rain/hail mix. Here, we see values ranging between 0.9 and 0.95, with some lower values within the forward flank downdraft. Now overlaying the hydrometeor classification algorithm (HCA), you can see where the radar is seeing the hail/rain mixtures in red.

Finally, looking at the specific differential phase, you would see values greater than 1.0 deg/km here, and you do in both of the highlighted areas. Some of the more extreme values, like the area of 4.0-7.0 deg/km near the rear flank downdraft, are where the greatest rainfall rates are occurring, but some values could be a result of water coated hail. It is important to note that KDP values will not display in areas of CC less than 0.90.



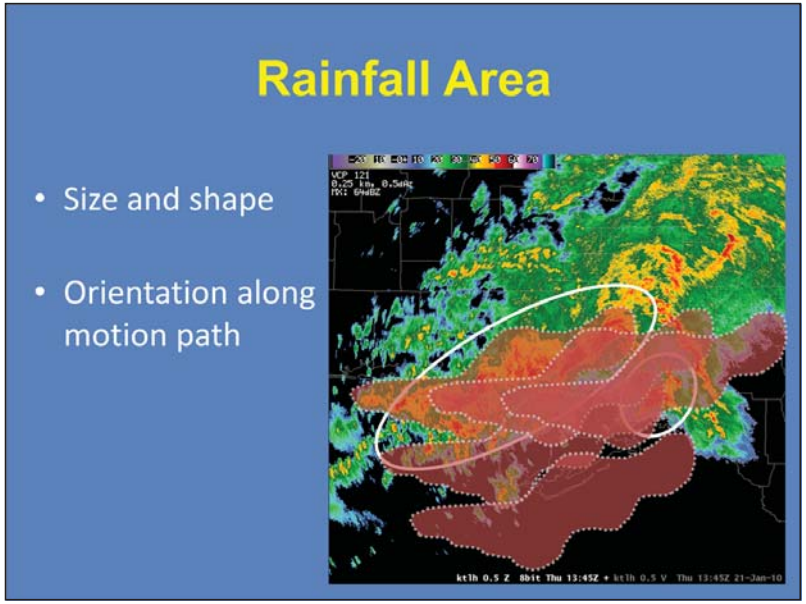


The key to high precipitation rates with supercells is understanding why supercells are such a threat despite poor precipitation efficiencies and at times swift movement. You just saw the dual-polarization characteristics of a supercell. However, there are some environmental factors to consider. In this case, we will focus on how much air it is ingesting and how moist is the inflow. Looking at this high-precipitation (HP) supercell near Midland, TX, this storm is in an environment where the profile is moist up to 700-mb and is considerably dry above it (The warm cloud layer is about 9,000 ft.). So, this storm is undergoing processes such as dry air entrainment and evaporation.

However, the storm has a very strong, moist updraft with it. It is ingesting very moist air (mixing ratio of 14 g/kg) at about 20 m/s. The updraft is wide and has a vertical depth of over 2 km. So, even though this supercell existed in an environment characterized by dry air above 700 mb, the storm produced rain rates of 2-4 inches per hour and fatal flash floods in the city of Midland. This shows that the factors that give HP and classic supercells a lower precipitation efficiency can be balanced by large values of moisture inflow and why supercells can produce high rain rates and flash flooding.

Now that we have examined the meteorological variables that influence the precipitation rate, let's examine the other meteorological factor that can influence the flash flood potential: precipitation duration.

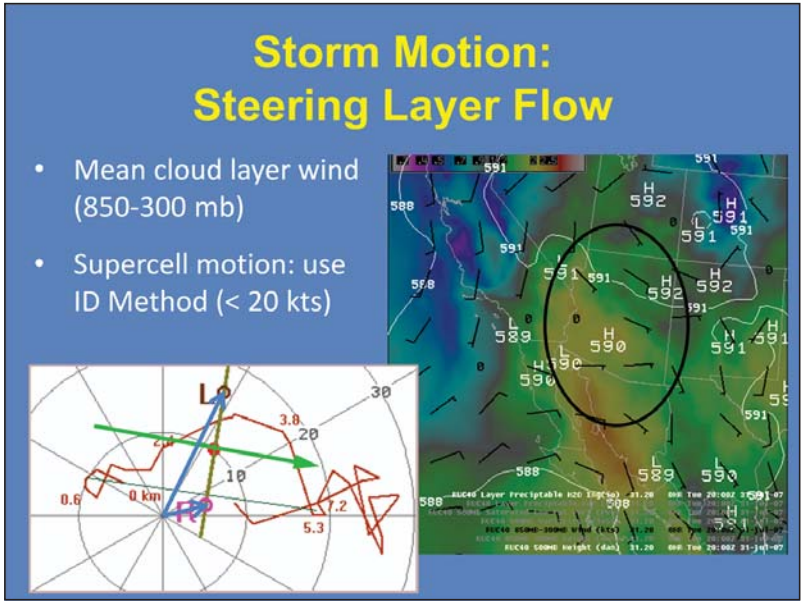
When we talk about duration, we are talking about the residence time of precipitation over a location. There are a number of things that affect the duration of rain over a specific area, such as rainfall area, storm motion, slow moving boundaries, and training storms. We will look at each one to see how they can increase precipitation duration.



The first thing to look at is the size and shape of the precipitation area. Using the Tallahassee, FL radar, you can see a supercell southeast of the radar. Supercells and pulse storms are small in size, and depending on movement, will generally have a small residence time over one location. Linear convection, like the complex to the west, cover a much greater area. Therefore, the residence time of rainfall over a specific point is increased.

One thing to look at is the orientation of the precipitation area with respect to the motion path. Let's assume that this convection highlighted here is moving towards the south at a constant speed of 40 mph. If we were to assume that the width of the area is approximately 20 miles, then the residence time of the moderate to heavy rain is about 30 minutes.

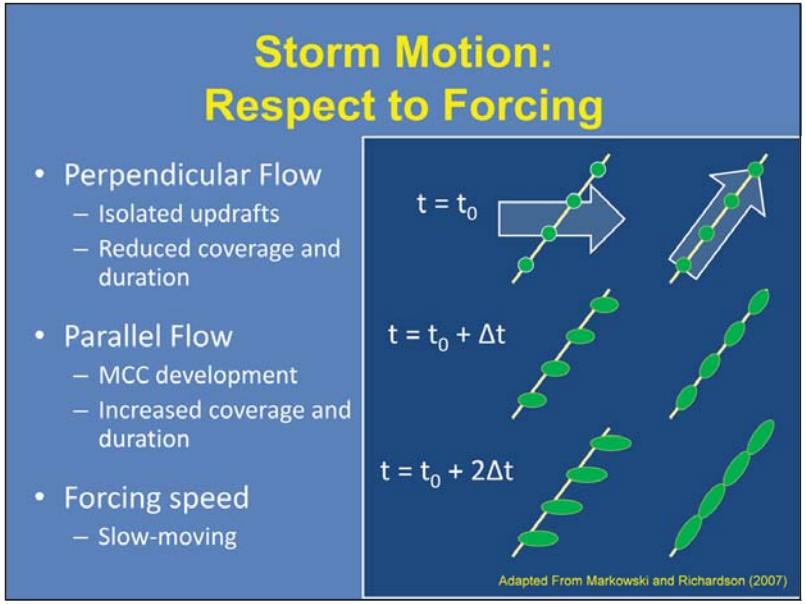
Now let's assume that this linear complex is moving to the east at 40 mph. If we were to assume that the length of the area is about 120 miles, then the residence time over this area is closer to three hours. With this event back in January 2010, the complex was moving towards the east and produced 4-7 inches of rain around the Tallahassee area.



Storm motion is a significant factor when it comes to precipitation residence time over an area. Obviously, slower storm motions lead to longer durations. But what would you look for to determine storm motion?

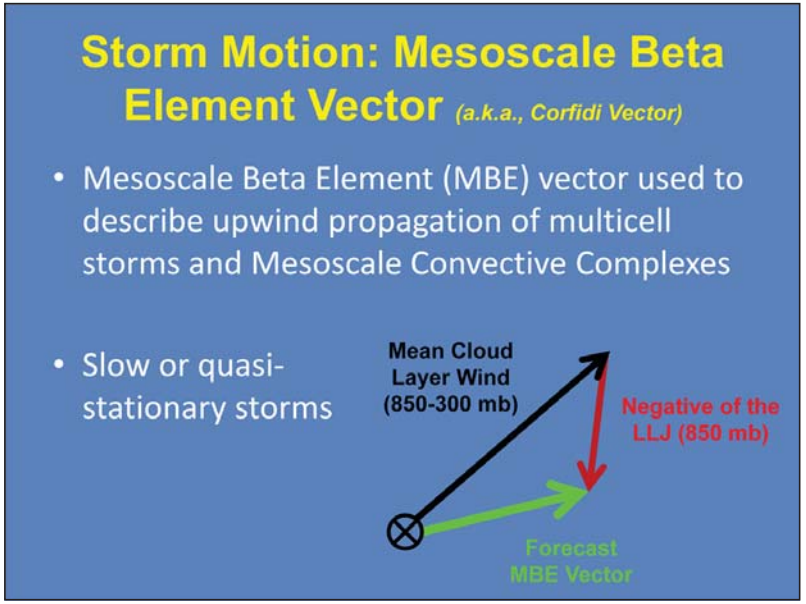
One factor to look at is the steering layer flow. You can use your volume browser in AWIPS to view the mean wind between 850 and 300 mb. Using the example over Arizona, you can see that the 850-300 mb winds over the state is generally from east to west at 5 kts. Very slow moving storms in this area did produce fatal flash floods in the Tuscon CWA.

For supercells, you can use the Internal Dynamics (ID) Method to calculate the motion of right and left moving supercells. A storm motion of under 20 kts is preferred. In the example shown here, you can see that right moving supercells with this hodograph would be moving just north of due east at about 5 kts. Recall how to use the ID Method with hodograph in the lesson on Supercell Dynamics and Motion.

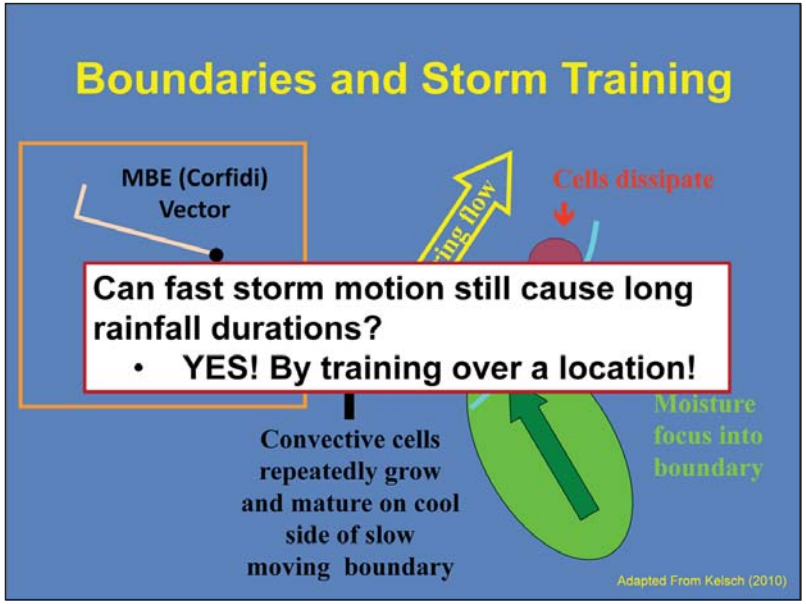


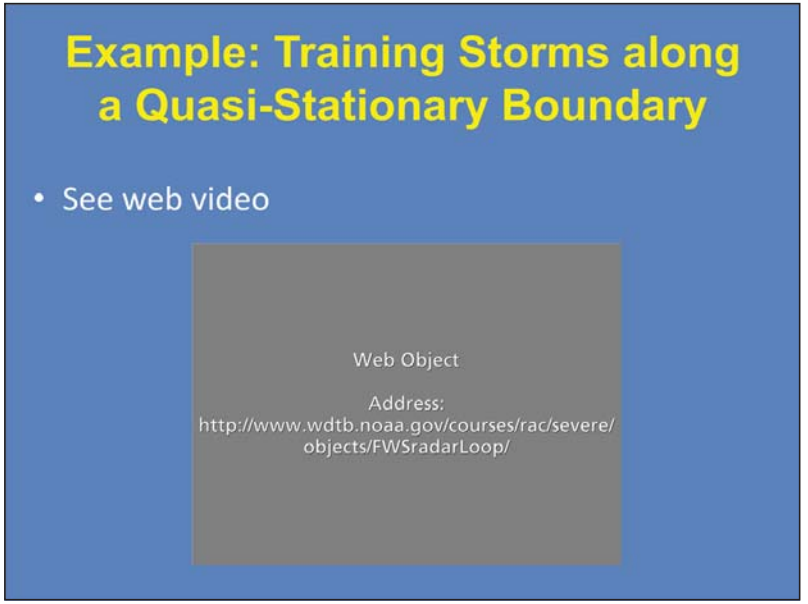
Forcing mechanisms play an integral role in the development and motion of convection. A forcing mechanism can range from fronts to outflow boundaries to topographic features. How storms form and move along a boundary can determine whether you have isolated updrafts or consolidated line segments and mesoscale convective complexes (MCCs).

Recall the work of Markowski and Richardson. Flow that is perpendicular to the forcing will lead to isolated updrafts, which in turn will have reduced areal precipitation coverage and smaller precipitation durations. Flow that is more parallel to the forcing will lead to linear convective formation. This will increase precipitation coverage and duration. Slow moving or quasi-stationary forcing mechanisms are best for increased precipitation residence time over an area.



When dealing with multicell storms and mesoscale convective complexes (MCCs), the Mesoscale Beta Element (MBE) vector can help describe the upwind propagation of multicells and MCCs. Recall from the lesson on Multicell Motion that the MBE vector is calculated from taking the mean cloud layer wind and adding the negative of the low-level jet (850 mb flow depending on the depth of the inflow layer). Small MBE vectors means that if there is upwind propagation, then the complex will be slow moving or even quasi-stationary.





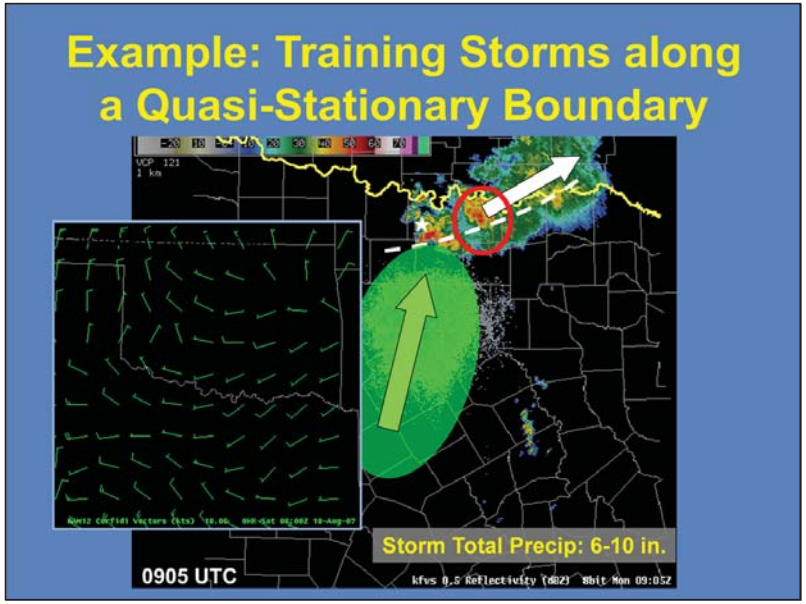
So far, we have talked about slow storm motions. What if storm motions are relatively fast? Can we still get large durations of rainfall? The answer is definitely yes.

If storms are training over the same location, it is easy to get the adequate duration for flash flooding to occur. One way is to have storms continuously propagate along a slow moving boundary. In this diagram, you have a SW-NE oriented boundary with an area of focused moisture transport. With enough lift and instability, convective cells will develop, move along the boundary, and dissipate. This cycle will continue so long as the boundary motion, moisture, instability, and trigger remain constant.

If you were to examine the vectors of this case, the mean flow parallels the boundary with expected storm motion of 25 kts. The MBE Vector shows that with backbuilding storms (upwind propagation), this system will move to the east at about 5 kts. This will allow for ample precipitation duration for flash flooding.

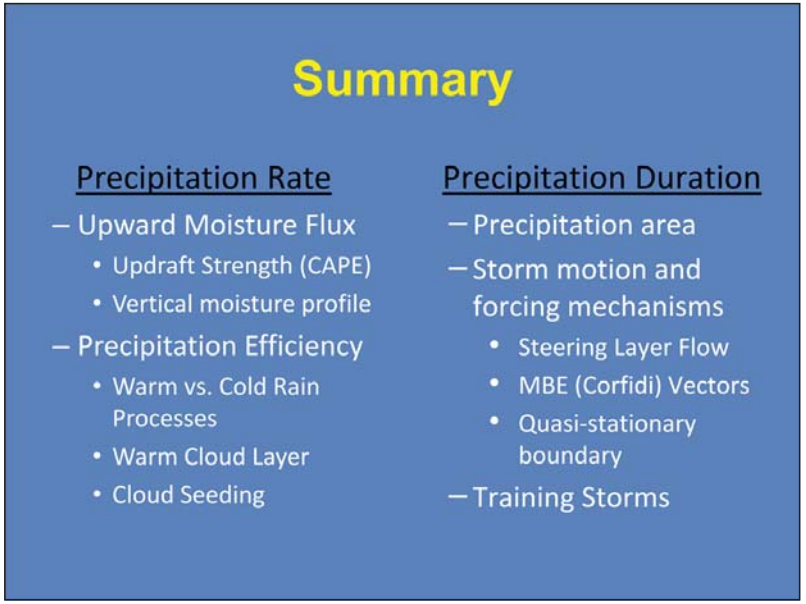
Here is an example of training storms that led to significant flash flooding. In this case from the Dallas/Fort Worth office, a series of storms train over the northern part of Texas near the Red River. This loop shows 5 ½ hours of radar data from KFWS. The star on the map shows the relative area of maximum focus and continuous development. Note how the storms train over the same area until a substantial cold pool develops for forward propagation.

Click next to advance to the analysis of this event when you are done viewing this loop.



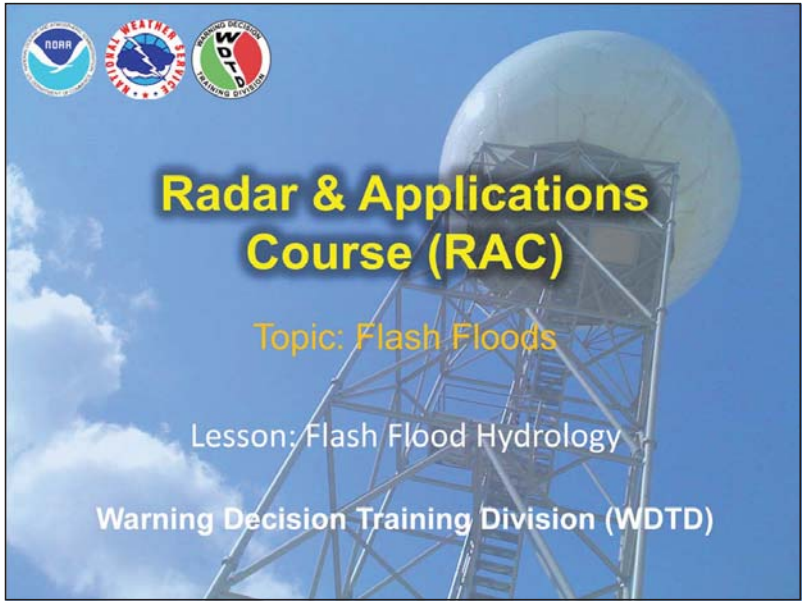
This event was created from a remnant mid-level circulation and boundary where a small vorticity maximum around the southern periphery is providing focus along the axis of forcing. The 1200 UTC sounding from Dallas/Fort Worth showed a very moist southerly 850 mb winds at 35 kts. You saw that storms initiated along the boundary where the forcing was maximized and then moved off to the ENE. However, the area of storms barely moved over a four hour time period. As you see here in the MBE, or Corfidi, Vectors, overall forecasted motion of the system is around 5 kts.

During this event, some areas received over four inches of rain in less than two hours, and storm total precipitation of 6-10 inches. There were six fatalities from these flash floods. Grayson County, which is circled in red here, had approximately 450 water rescues from vehicles and homes. There were hundreds of other water rescues in the surrounding counties.

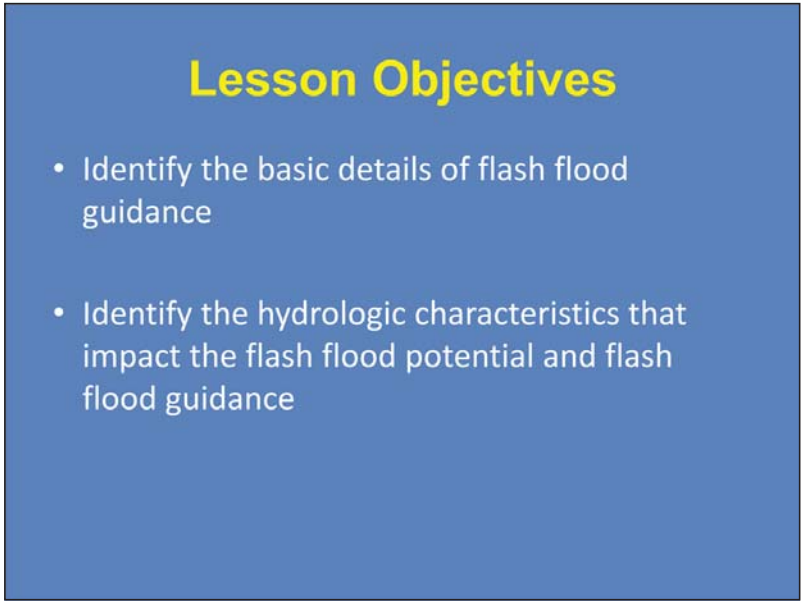


In summary, you saw that there were two primary meteorological factors regarding rainfall and flash flooding. With precipitation rate, you saw how the strength of the updraft and the overall vertical moisture profile played a role in the upward moisture flux into a storm. The fraction of that that is returned as precipitation is defined as the precipitation efficiency of a storm. This is dependent upon the type of rain processes that are dominant, the depth of the warm cloud layer, and cloud seeding.

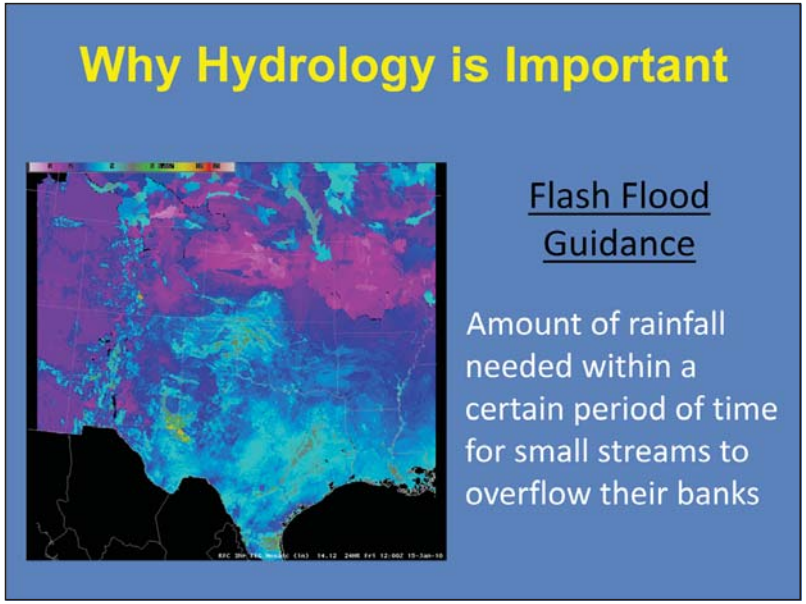
With precipitation duration, you have to consider the area and motion of the precipitation. Understanding storm motion and forcing characteristics, such as flow parallel to boundaries, weak steering layer flow, backbuilding complexes via slow MBE vectors, and slow or quasi-stationary boundaries, can help provide longer duration periods. Analyzing the mesoscale environment can help you determine the potential for training storms if storm motions are relatively fast.



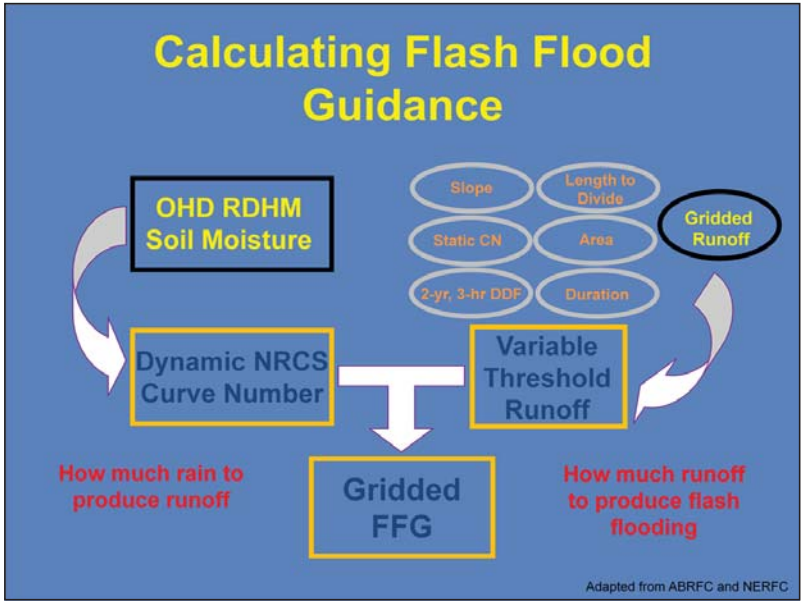
Jill Hardy again and welcome to this lesson on flash flood hydrology. Steve Martinaitis of OU CIMMS at NSSL will again be narrating. But if you have any questions regarding the material, please feel free to contact me, or the RAC team. Our contact information will be on the next slide.



There are two objectives to this lesson. At the end of this presentation, you should be able to identify the basic details regarding the creation of flash flood guidance and to identify the hydrologic characteristics that can impact the flash flood potential and the flash flood guidance product.

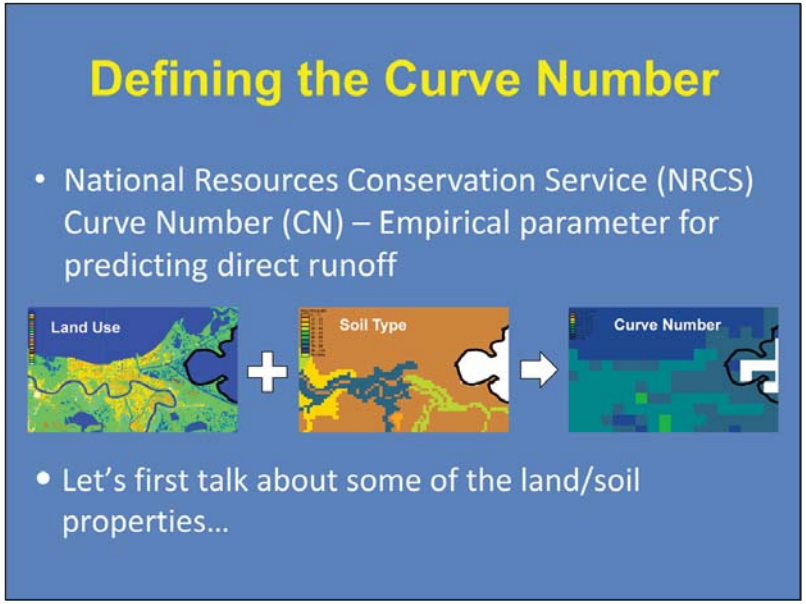


Understanding hydrologic properties, such as basin geometry, land use, and soil moisture, can help in determining what areas are prone to flash flooding and why. The flash flood guidance (FFG) product is created using these hydrologic properties to provide NWS forecasters a rainfall value needed within a certain temporal period for small streams and creeks to overflow their banks. And as you can see, these values can vary quite drastically across the country.

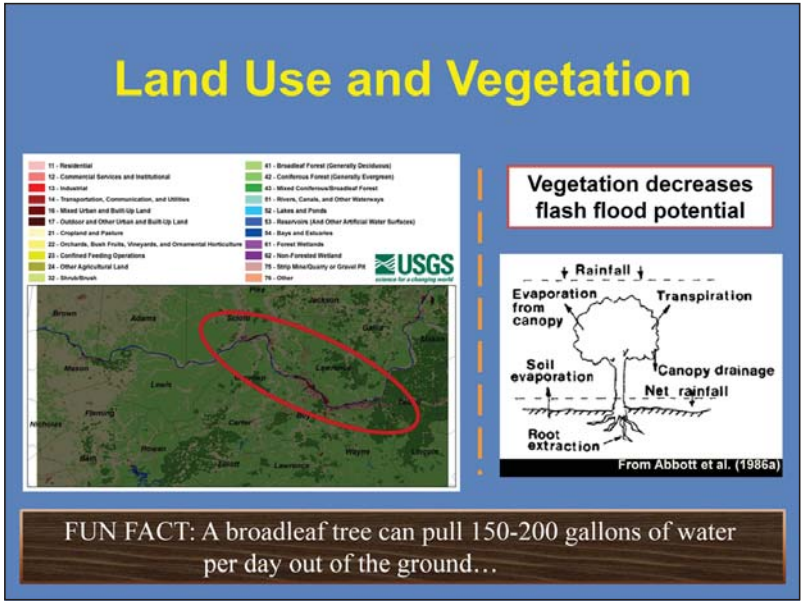


Most river forecast centers (RFCs) create a gridded flash flood guidance (FFG) product generally four or more times a day. FFG is derived from how much rainfall is needed to produce runoff via a dynamic National Resources Conservation Service, or NRCS, curve number (CN) and how much runoff is needed to produce flash flooding via a threshold runoff, or Thresh-R, value.

Both the NRCS curve numbers and Thresh-R values are calculated using different hydrologic properties. We will first focus on what hydrologic conditions influence the NRCS curve number.

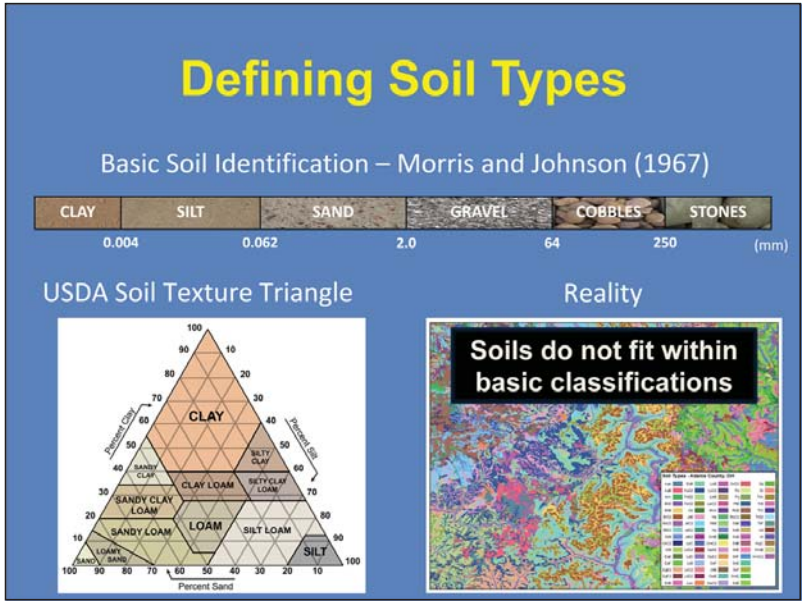


The National Resources Conservation Service, or NRCS, Curve Number (CN) is an empirical parameter for predicting direct runoff. The curve number is generally based on soil type and land use, which both impact the amount of rainfall that is intercepted and infiltrated. The higher the curve number, the less rain is needed to create runoff. So, before we go into the operational details of the curve number, let's focus on the characteristics of land usage and soil types.



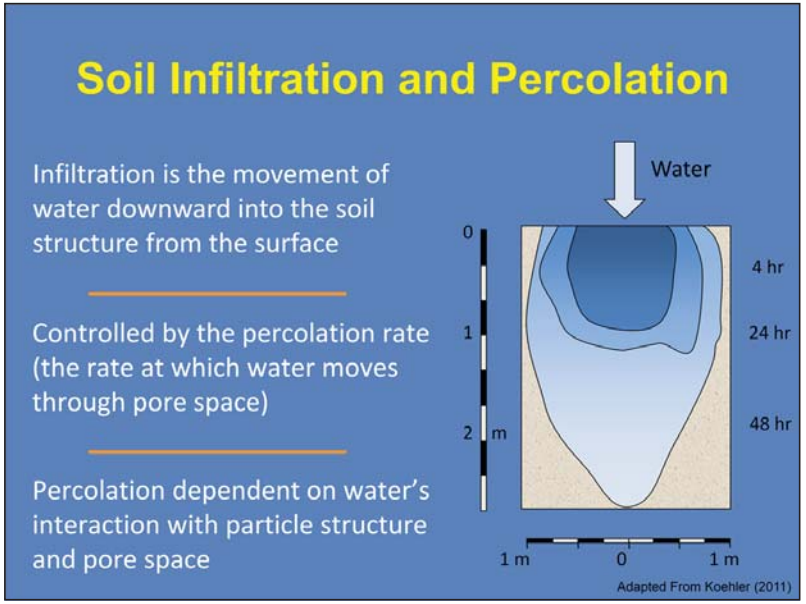
Land use can help determine how much water can be intercepted and how much water can be translated into runoff. Shown here is an example of a land use map for southern Ohio and northern Kentucky. The western part of this domain is dominated by croplands and pastures. The eastern two-thirds is a combination of broadleaf and coniferous forests. Urbanized areas are located along the Ohio River, and we will talk about the impacts of urbanization later in the module.

One the bigger influences of land use on the curve number is vegetation. Vegetation of all types help decrease the flash flood potential. Leaves can intercept rainfall before it reaches the surface. Water on and within the foliage undergo evapotranspiration. Roots help increase infiltration and can extract water from the surface and top layers of the soil. Did you know that one broadleaf tree can extract 150-200 gallons of water from the ground in one day? That's a lot of water. This is why areas devoid of vegetation or have been deforested have higher curve numbers and a higher potential to flash flood.



How soils are defined has evolved over the decades. The work by Morris and Johnson in 1967 defined the three basic soil types by the average size of their particles. Clay was defined by particles less than 0.004 mm while sand particles can approach 2 mm in size. The USDA then developed a soil texture triangle to better determine soil types based on the proportion of sand, silt, and clay after particles larger than sand have been removed. The USDA soil texture triangle defines 12 different soil types, including loam, which is a soil composed of sand, silt, and clay in relatively even proportion, and these types are commonly used in curve number and modeling calculations.

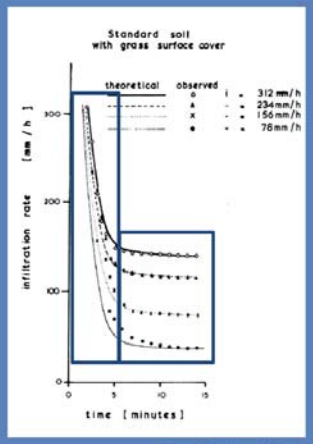
In reality, there are dozens upon dozens of soil types in the country. This image here shows 76 different soil types just within Adams County, OH. And in reality, soils do not fit within their basic classifications, such as those described by Morris and Johnson (1967).



One of the most important characteristics of soil that impact runoff potential is the ability of water to infiltrate the soil. Infiltration is the downward movement of water from the soil surface into the soil structure. The ability of water to infiltrate the soil structure is controlled by the percolation rate, which is defined as the rate at which water moves through pore space. The percolation rate is generally dependent upon how water interacts with the particle structure and volume of porous space. For example, the pore space of clay-type soils can vary, but how water interacts with the clay particle reduce the infiltration and percolation rates. Thus, it takes less water to generate runoff. There are a number of different equations that can be used to calculate soil infiltration, which will not be discussed here.

Decay of Infiltration During a Rain Event

- Infiltration capacity decreases exponentially to an equilibrium rate
- Influencing Factors
  - Percolation
  - Storage of water
  - Particle absorption and swelling



From Nassif and Wilson (1975)

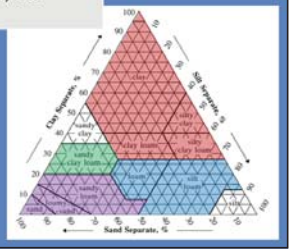
One important thing to know is that the maximum infiltration rate of the soil, more commonly referred to as the infiltration capacity, is *not constant*. A study by Nassif and Wilson (1975) compared rainfall rates with the infiltration properties of various soils. They tested different soil types, including some with grass surfaces, at different slopes using large soil trays and a sprinkler-type system to simulate instantaneous rain rates from 3-12 in./hr.

During their experiments, they found that within the first 5-10 minutes of applying the simulated rainfall, the infiltration rates of the soils decreased exponentially and began reaching an equilibrium rate. The experiment also demonstrated why a saturated ground cannot take in as much water as an unsaturated ground. The biggest factor that influences the ability of water to infiltrate the soil is percolation, while storage of water in the soil and particle absorption and swelling also influence this process.

Soil Infiltration and Curve Number

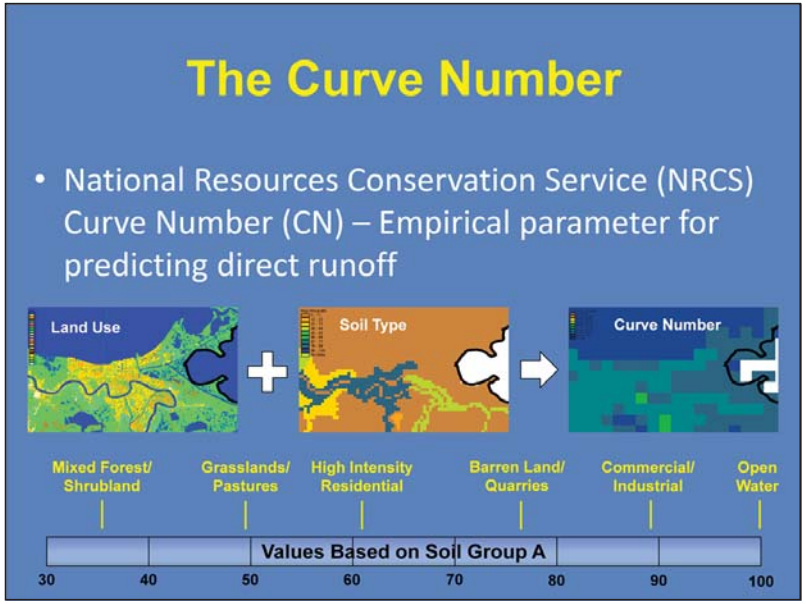
HSG	Soil Type	Characteristic
A	Sand, Loamy Sand, Sandy Loam	Low runoff potential, high infiltration rates
B	Silt Loam, Loam	Moderate infiltration rates
C	Sandy Clay Loam	Low infiltration rates
D	Clay Loam, Silty Clay Loam, Silty Clay, Clay	High runoff potential, low infiltration rates

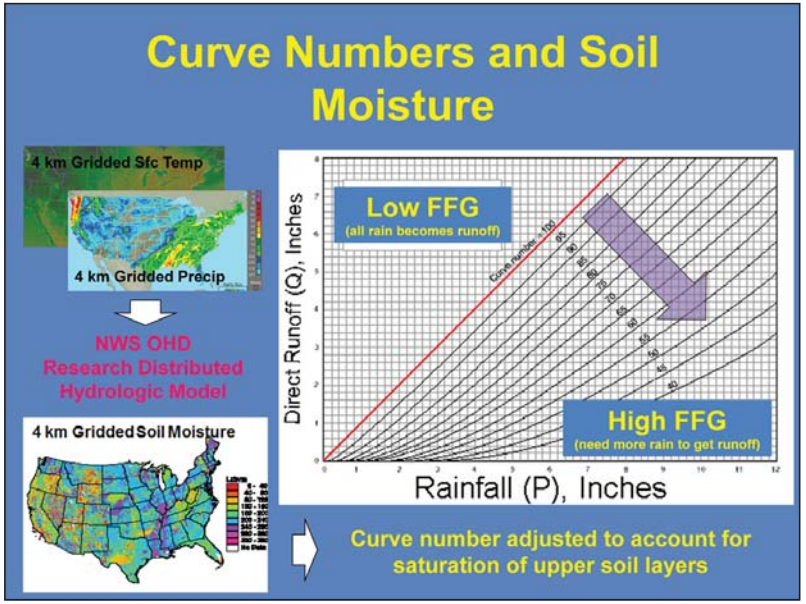
- Hydrologic Soil Group (HSG) classification used to determine NRCS Curve Number



Knowing all of this, we look back at how soil types and land use are operationally applied. NRCS soil scientists categorized the soil types defined by the USDA soil texture triangle into four Hydrologic Soil Groups (HSGs) based on infiltration, runoff characteristics, and texture. Group A consists of mostly sandy-based soils that have low runoff potential and high infiltration rates. Group B consists of silt loam and loam that have moderate infiltration rates and moderately coarse to fine particle textures. Group C consists of sandy clay loam, which is has low infiltration rates with moderately fine textures. And finally, Group D consists of mostly clay-based soils that have high runoff potential and low infiltration rates.

These four HSGs are then used to determine the NRCS curve number.

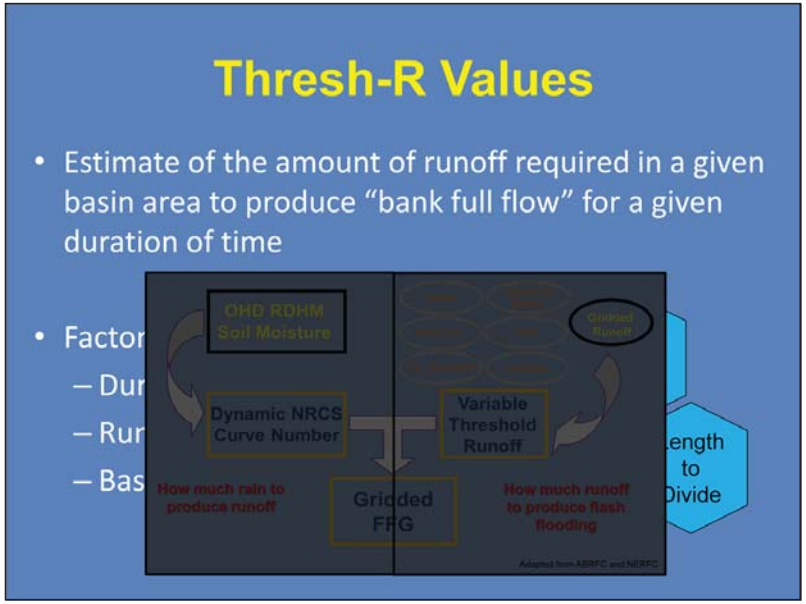


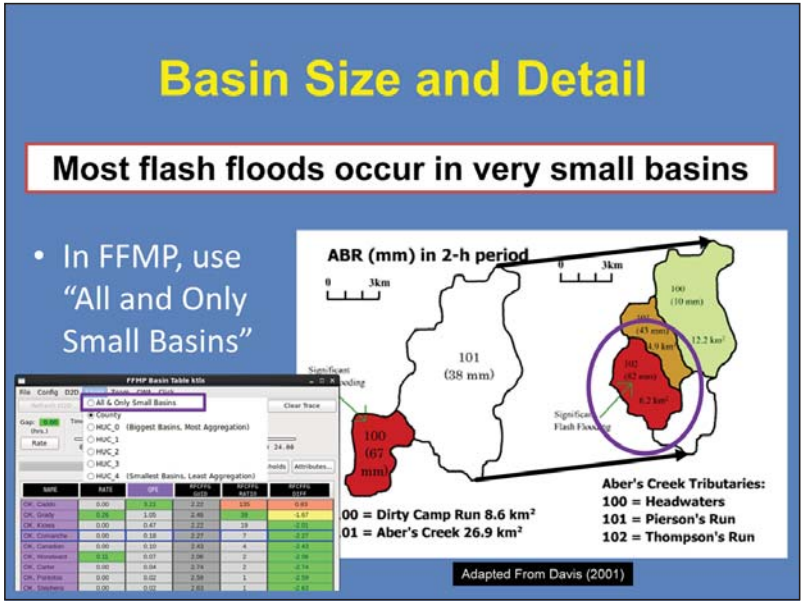


So let's see how the land use and soil types impact the NRCS curve number (CN). The CN is a unit-less number that ranges from 30 to 100. The higher the number, the less rain is needed to create runoff. Each type of land use is given four CNs, one for each Hydrologic Soil Group (HSG). In this example here, we will define the curve number for a variety of land uses for soil Group A, the sandy-based soils with low runoff potential and high infiltration rates.

Mixed forests, shrublands, grasslands, and pastures all have low curve numbers, which means these land uses for this soil group have low runoff potential and would need a lot of rain to generate any runoff. Here are where high intensity residential areas, barren lands, quarries, and commercial/industrial areas fall on the curve number spectrum. Notice how more urbanized land and land devoid of vegetation have much higher curve numbers. Open water has a curve number of 100, meaning all rainfall becomes runoff. Remember that these values are for Group A, the sandy-based soils. The curve number for each land use type is much higher with Groups B, C, and especially D.

For gridded FFG, the National Weather Service (NWS) Office of Hydrologic Development (OHD) uses a research distributed hydrologic model to create a dynamic curve number. The model uses CONUS-scale 4 km gridded surface temperatures and 4 km gridded precipitation to create a gridded soil moisture product. The CN is then adjusted based on the soil moisture calculated to account for saturation of the upper soil layers. In areas that have seen recent, heavy rainfall, the curve number can be increased to near 100, which means that the FFG is reduced and nearly all rainfall can potentially be translated into runoff. The lower the curve number, the greater the FFG and the more rain that is needed to start generating runoff.





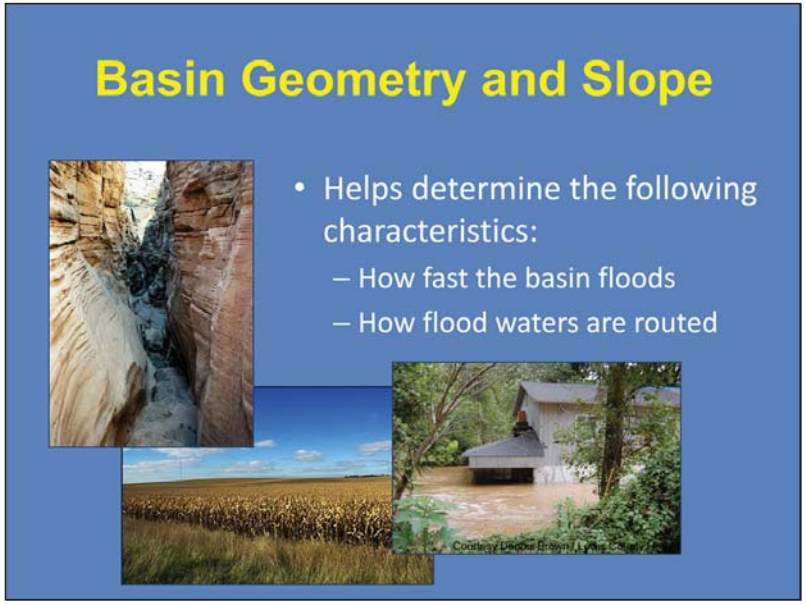
So we looked at the first half of the FFG calculation that asks how much rain does it take to generate runoff. Now we will focus on how much runoff will it take for small creeks and streams to reach "bank full flow" conditions, which is defined by the Thresh-R value.

The Thresh-R value considers a number of different factors, including duration of rainfall, runoff estimation, and a variety of basin characteristics. In this module, we will focus on what impacts basins have on the FFG.

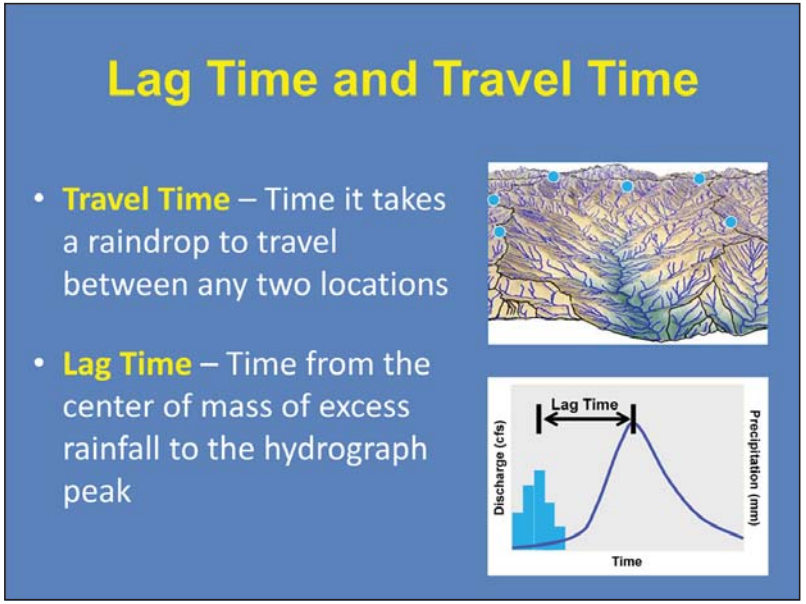
There are many levels of basin detail, ranging from the parent basin, such as the Mississippi River Basin or the Florida Watershed, to the smaller basins and sub-basins that compose them (a couple of square kilometers). The majority of flash floods occur in very small basins, mainly because the scale of the heaviest rainfall is also quite small.

Using this example from the Aberdeen, SD office, you can see two adjoining basins, Dirty Camp Run (#100) and Aber's Creek (#101). Significant flash flooding occurred in the Dirty Camp Run basin as a result of an average of over 2.5 inches of rain in two hours across the basin. In Aber's Creek, there was an average of 1.5 inches of rain across the basin. However, if you were to split the Aber's Creek basin into its sub-basins, you can see that the southwest part of the basin had almost 3.25 inches of rain in this two-hour period. Without looking at the smallest basins, it is possible that some areas would have went unwarned due to large scale basin averaging.

When using FFMP, you will get the greatest detail by going to the Layer menu in the FFMP Basin Table and selecting "All & Only Small Basins."



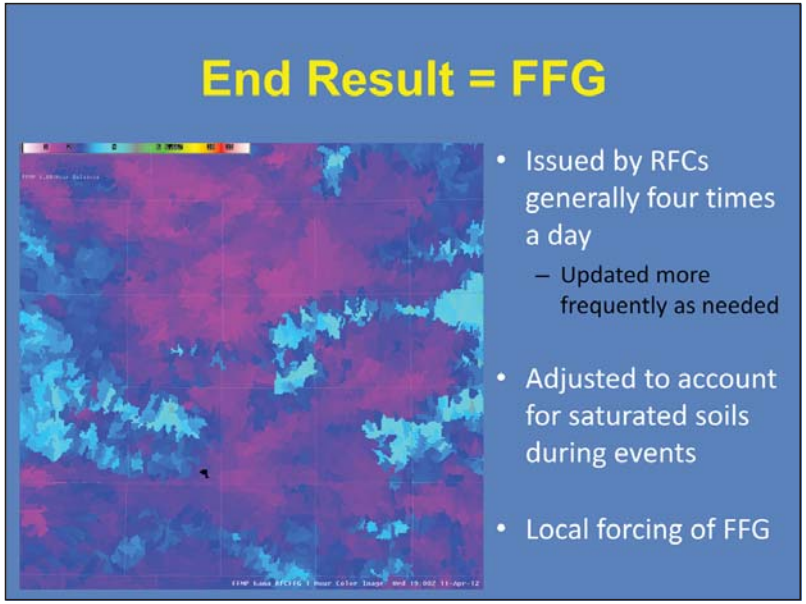
The three-dimensional geometric characteristics of a basin is important when determining how fast a basin can flood and how the flood waters are routed within the basin. A lot of this is dependent upon the slope of the terrain. If we are talking about the flat fields of the central Plains, then it could take awhile for runoff to move out of the area. In contrast, if you look at the mountainous terrain and slot canyons of the western U.S., then water can be quickly routed downstream. Think of a marble on a shelf. If the shelf is level, the marble is stationary. If the shelf sits at an angle, then marble rolls off. The same theory applies here.



Two terms that are used when talking about runoff in a basin is lag time and travel time. Travel time is the time that it takes a raindrop of water to travel between any two locations.

Lag time is the time difference between the center of mass of excess rainfall and the peak in a river hydrograph.

Both of these are dependent upon the size and shape of the basin, as well as the slope of the ground within the basin. Basins that are prone to flash flooding can have a lag time of just 15 minutes and the water can sometimes travel a considerable distance in a short amount of time. Basins not prone to flash flooding can have lag times that can exceed one hour.



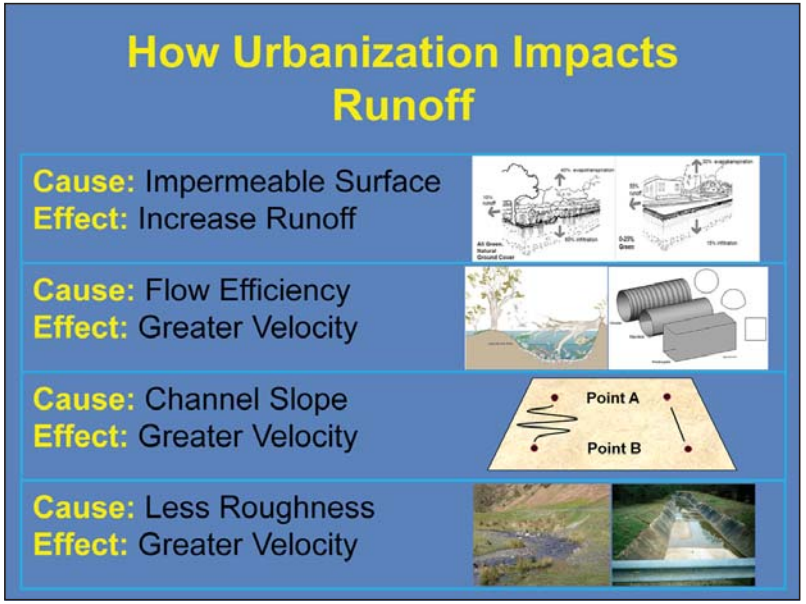
Through hydrologic modeling, the combination of dynamic NRCS curve numbers and Thresh-R values creates the operational FFG that we use today. FFG is issued by the RFCs generally four times a day, and they can be updated more frequently as needed. Since gridded FFG accounts for recent changes in soil moisture, they can be adjusted for saturated conditions during events. FFG can also be forced locally using a GUI in the AWIPS workstation.



Areas that have compromised FFG, and where FFG is generally forced, are urban areas and wildfire burn scars. Let's start with how urbanization impacts FFG.



Urbanization can create as much as five times more runoff than that of a completely forested area. This primarily has to do with areas being covered by concrete instead of soils and vegetation. The infiltration rate of concrete is near zero, so almost all rain that falls on to concrete surfaces is translated into runoff. In addition to the "concrete jungle," drainage systems within urban areas may not be adequate enough to handle very high quantities of water. But there is more to urbanization than just creating more precipitation runoff.



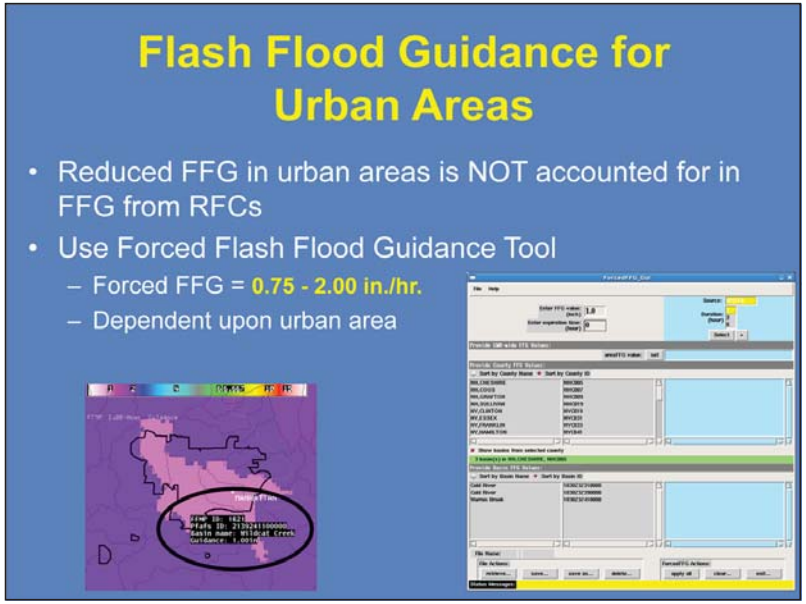
We know that when you have more impermeable surfaces (i.e., areas covered by concrete and asphalt), you are able to generate more runoff. But we must also look at what other factors urbanization can have on runoff.

Natural rivers and streams have very complex shapes, while urban drainage systems, culverts, pipes, etc. have simple shapes. This allows for greater flow efficiency of the runoff, and thus, gives the runoff a greater velocity.

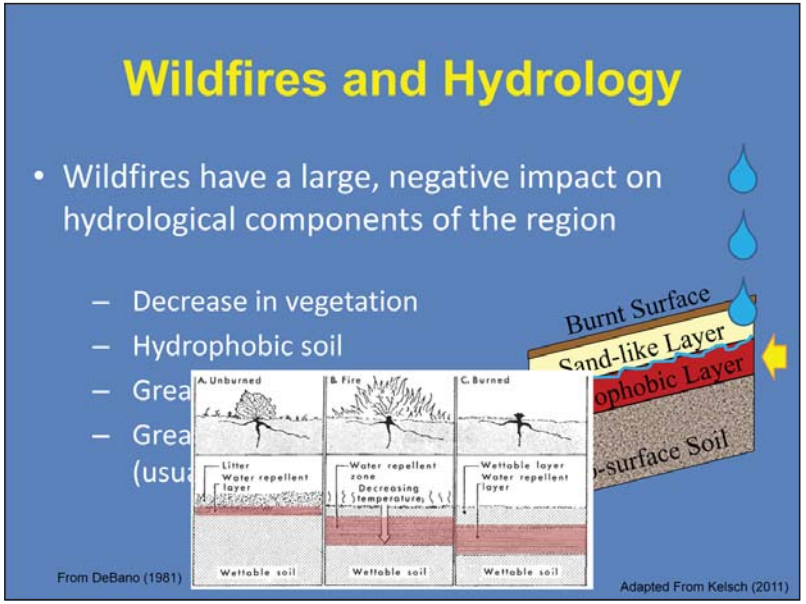
Streams and rivers have a tendency to meander and are never really straight. Urban areas tend to have very straight drainage paths. When placing a meandering path and straight path on a sloped surface, not only does the straight path take less time for runoff to go from Point A to Point B, the overall slope of the straight channel is much greater. Straighter channels in sloped urban areas will have a greater channel slope and a greater runoff velocity.

Also, natural streams and rivers have a lot of roughness that come from rocks, vegetation, and complex channel shapes. On the other hand, urban channels, especially those that are paved, tend to be quite smooth, and thus have less roughness. This again translates into greater runoff velocity.

So, not only does urbanization create more runoff, it can also provide a greater velocity to the runoff, which makes it that much more dangerous to life and property.



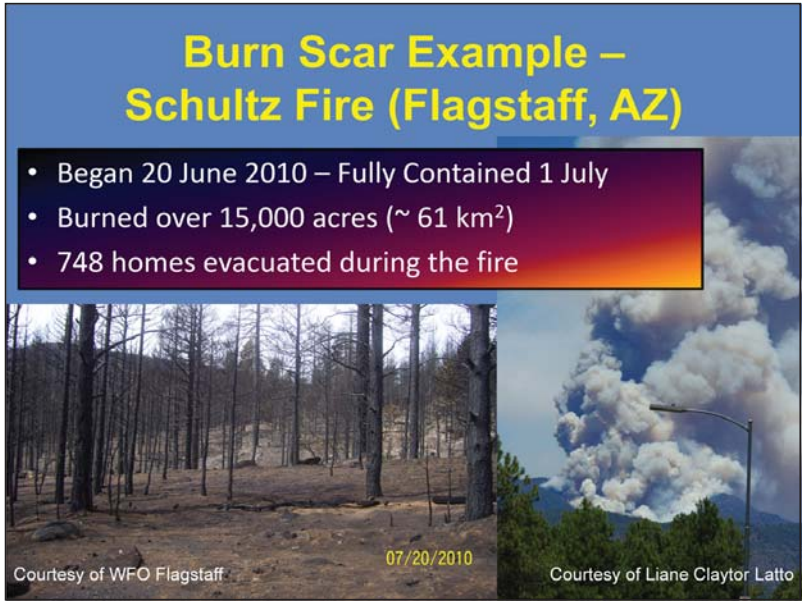
The reduced flash flood guidance (FFG) for urban areas are not accounted for in the FFG delivered by the RFCs. However, you can use AWIPS to force FFG for these areas. Depending upon the urban area, it can take anywhere from 0.75-2.00 in./hr. for flash flooding to occur. Ask your AWIPS focal point about using the Forced FFG interface.

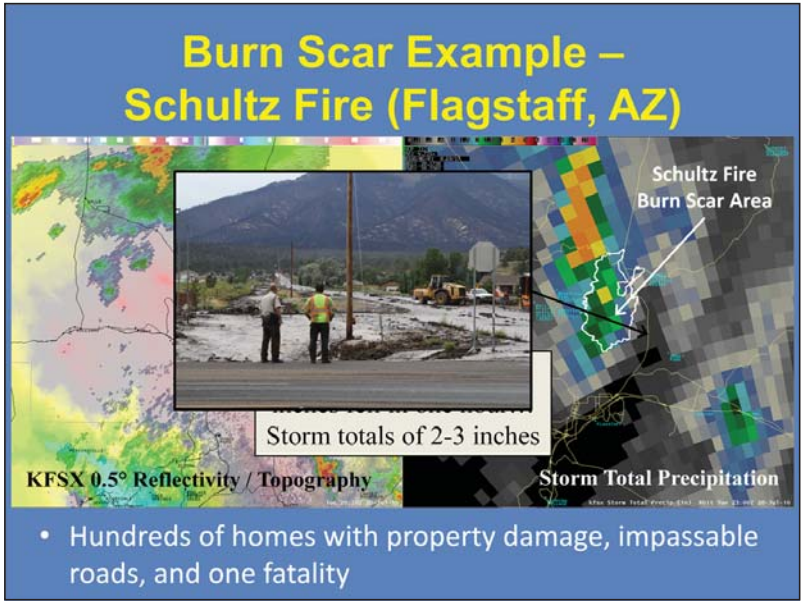


Now, I'm going to make a pretty obvious statement here and say that wildfires have a large and very negative impact on the hydrologic factors related to flash flooding. But, it is how the fires impact the area and for how long those impacts last that are critical to hydrology and flash flooding.

The part that we do see is the removal of vegetation. As we discussed earlier, vegetation helps capture water and increase soil infiltration, which reduces the potential for flash flooding. Without vegetation, there are no leaves to intercept rainfall or roots to extract water. One of the byproducts of combusting vegetation, especially in high-intensity forest fires, is the creation of a heavy gas that sinks and penetrates the soil profile. As this gas cools, it condenses and solidifies into a hydrophobic waxy coating around the soil particles. This is the part that we don't see and is very critical to flash flooding in burn scars. So, in the aftermath of high-intensity fires, a water-repellent sub-surface layer of soil is present. The greatest impacts from this occur within the first year after the fire, and these impacts remain for three to five years afterwards.

So, when it rains, the water will penetrate an initial layer of burnt soil and surface material and then a less-stable sand-like layer (all byproducts of the fire). When the water reaches the hydrophobic layer, it cannot penetrate any further. The water then becomes runoff with a high yield of sediment from the soils above this hydrophobic layer. Flash flood guidance is greatly compromised in burn scars, and it is usually around 0.50 in./hr.



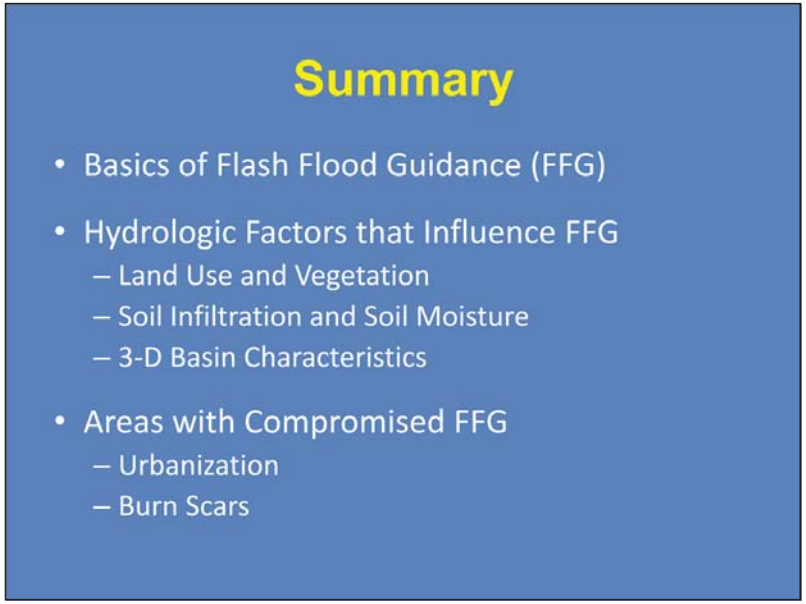


To show the impacts of a heavy rainfall event over a burn scar, we will look at the Schultz Fire of 2010. The fire began on the west side of the San Francisco Peaks, a volcanic mountain range north of Flagstaff, AZ, on June 20th, and the fire lasted about 12 days. The fire burned approximately 15,000 acres and resulted in the evacuations of nearly 750 homes.

The picture on the left shows what the terrain looks like one month after the fires. Note how the trees are devoid of any green vegetation and how the ground is burnt and lacks any underbrush or grass. Two hours after when this picture was taken, convection began to develop in the area.

On the left is the KFSX 0.5° reflectivity and topography image combination. The blue contour north of Flagstaff represents the burn scar area from the Schultz Fire. On this day, there were slow moving storms in and around the mountain range. Once convection developed within the burn scar, the office issued a Flash Flood Warning in anticipation of flash flooding based on very slow storm motions and high precipitation rates.

The image on the right shows the storm total precipitation from the event. Two to three inches of rain fell within the burn scar, most within a one hour time period. The runoff and sediment quickly moved down the mountain, especially with the aid of paved roads, including Highway 89. Downstream suburban areas were heavily impacted with home and property damage and impassable roads. A 12 year old girl was killed during the event when she was swept away by the flood waters in her neighborhood. This was a case where you had runoff easily generated in a burn scar area, which quickly moved down a steep mountainside and traveled into an urbanized area where roads and culverts quickly channeled the water in this dangerous situation.



In summary, you should have a basic understanding what Flash Flood Guidance (FFG) is and how it is calculated at the RFCs. We looked at a number of hydrologic factors that influence the values of FFG. We saw how land use, vegetation, and soil properties influence the dynamic curve number, which describes how much rain is needed to create runoff. We also saw how different basin characteristics help influence the Thresh-R value used to determine how much runoff is needed to create bank full conditions.

Finally, we took a look at two areas with greatly compromised flash flood guidance: Urbanized areas and burn scars from wildfires. The impacts of urbanization, such as large areas of concrete and hydraulically efficient channels, yield greater and faster moving water runoff. Meanwhile, the impacts of wildfire burn scars can last for years due to the removal of vegetation and the creation of a hydrophobic layer in the soil. Overall, both lead to situations where only a moderate amount of rain can produce dangerous flash flooding situations.



Hi, my name is Jill Hardy and welcome to this lesson on the High-Resolution Precipitation Estimator (or HPE) and Bias HPE products. Let's jump right in!

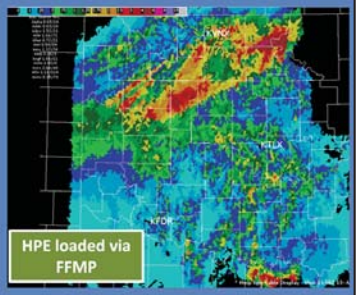
Learning Objectives

- By the end of this lesson, you will be able to:
  - Define what are the HPE, Bias HPE, and HPN products
  - Identify the default precipitation source for HPE, Bias HPE, and HPN
  - Interpret HPE and Bias HPE text overlays in the Volume Browser and FFMP
    - Precipitation source
    - Bias source
    - Bias factor and gauge-radar pairs
  - Identify how to determine areas of higher confidence in Dual-Pol HPE/Bias HPE precipitation estimates

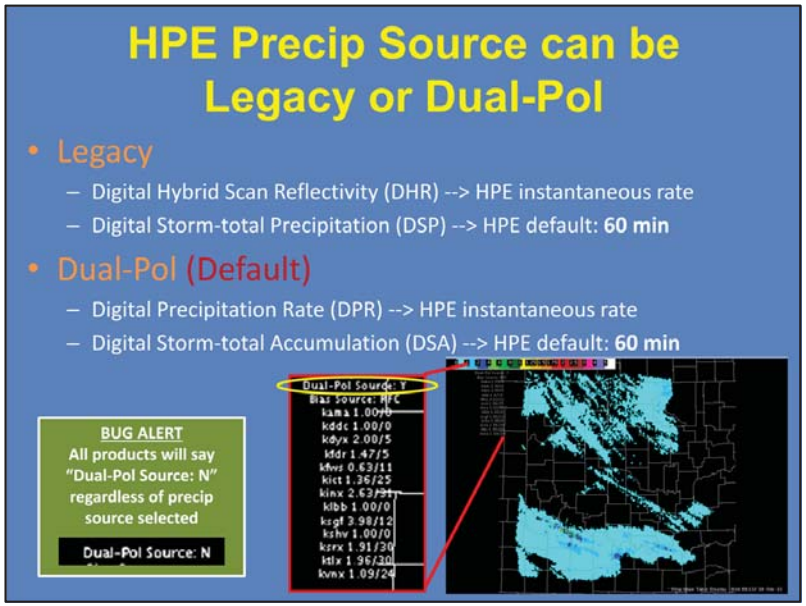
By the end of this lesson, you will be able to do the following…

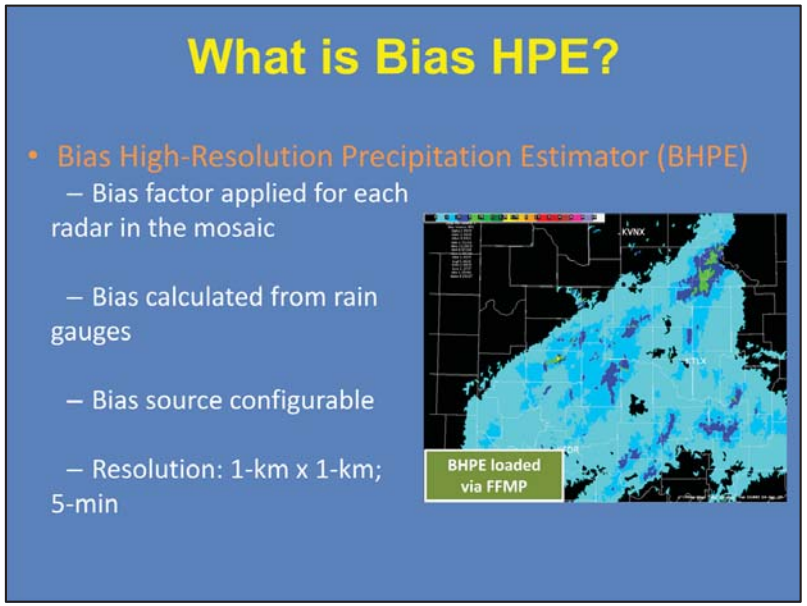
What is HPE?

- High-Resolution Precipitation Estimator (HPE)
  - Mosaic of single radar sources
    - Lowest altitude scan
      - Usually nearest radar unless beam blockage
  - Gridded rainfall fields
    - Instantaneous rain rate fields
    - Accumulated precipitation fields
      - Default = 60 min
  - Resolution: 1-km x 1-km; 5-min



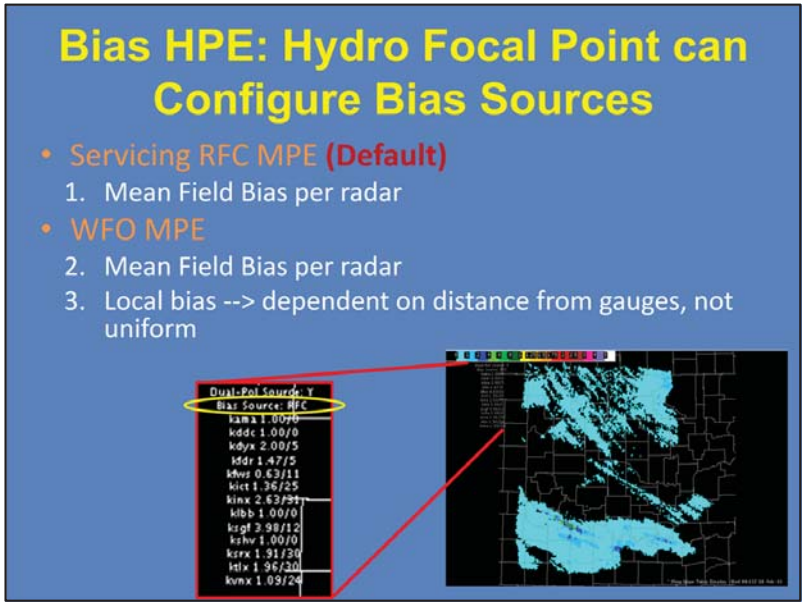
--The High-Resolution Precipitation Estimator (or HPE) was released in 2008, and was designed as a mosaic of rainfall data from all radars within your coverage area.
--One of the motivations behind HPE was to create a single precipitation mosaic that FFMP could display for your CWA and your back-up CWA, such as the image on the right that includes the radars overlaid.
--The mosaic is created from using the lowest altitude scan, which is usually the nearest radar. However, in areas of beam blockage, the mosaic may employ higher tilts.
--The precipitation source used in the mosaic is configurable, and the current source options will be discussed on the next slide.
--The two available HPE products are the instantaneous rain rate field and the accumulated precipitation field.
--A one-hour accumulation is the default in HPE, but this value can be altered by changing an HPE token with help from your hydro focal point.
--These products update every 5 minutes. However, the one-hour product needs at least an hour's worth of data before it begins updating every 5 min.
--HPE will stop collecting data once every contributing radar has not received rainfall within the last 20 min.
--In addition to FFMP, all HPE products are also available in the Volume Browser.



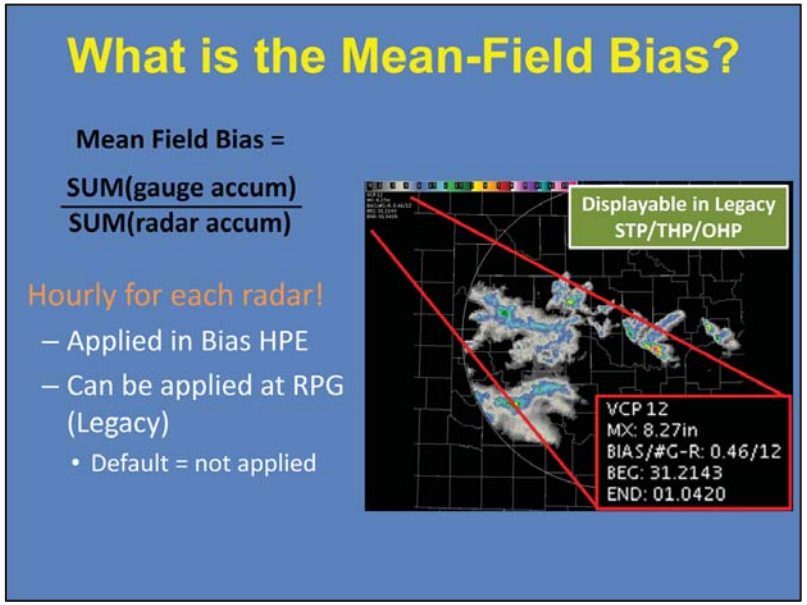


--Legacy OR Dual-Pol products can be configured by your hydro focal point to create the HPE and Bias HPE mosaics.
--When Legacy is the precip source, the DHR is used to create the HPE rain rates, and DSP to create the HPE accumulations.
--The default for HPE is to create a one hour accumulation.
--Similarly, the dual-pol counterparts are the DPR and the DSA.
--The Dual-Pol products are the default sources for creating the HPE and Bias HPE.
--The images below show an example of the HPE legend in FFMP. And you can see which precip source is being used. Here, it is Dual-Pol. If the source were changed to one of the Legacy products, the letter "N" would be displayed instead.
--However, there is currently a pretty bad bug which causes this label to **always** read "Dual-Pol Source: N", regardless of the precip source selected. Most offices will have the default configured (i.e., Dual-Pol), which unfortunately means…for most offices, this bug is valid until it is fixed in 16.4.1. But for the rest of this lesson, we use examples with the *correct* labels to avoid confusion.

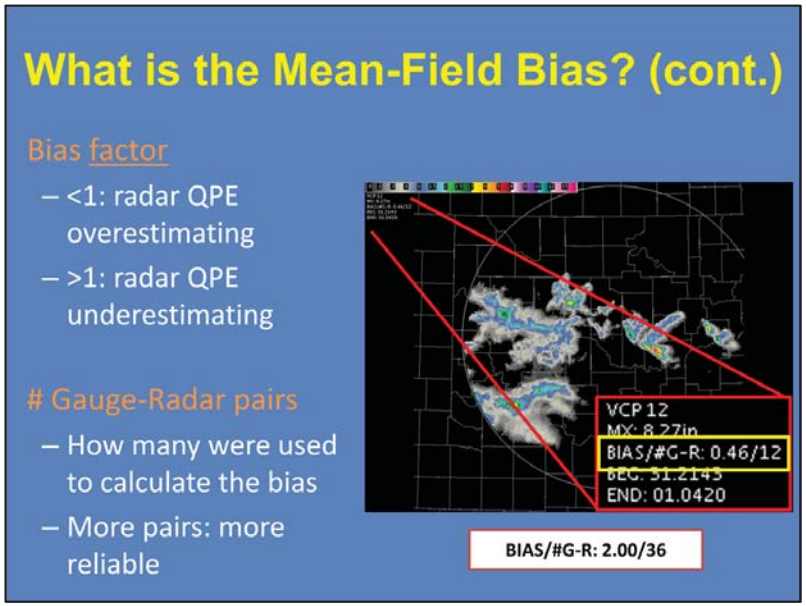
--Bias HPE is the same as HPE, except that a bias factor is applied for each radar in the mosaic, meant to help correct for radar uncertainties.
--In short, the bias is calculated by comparing rain gauge information with co-located radar data.
--The bias source is configurable, chosen by your local hydro focal point.
--Both have the same resolution.

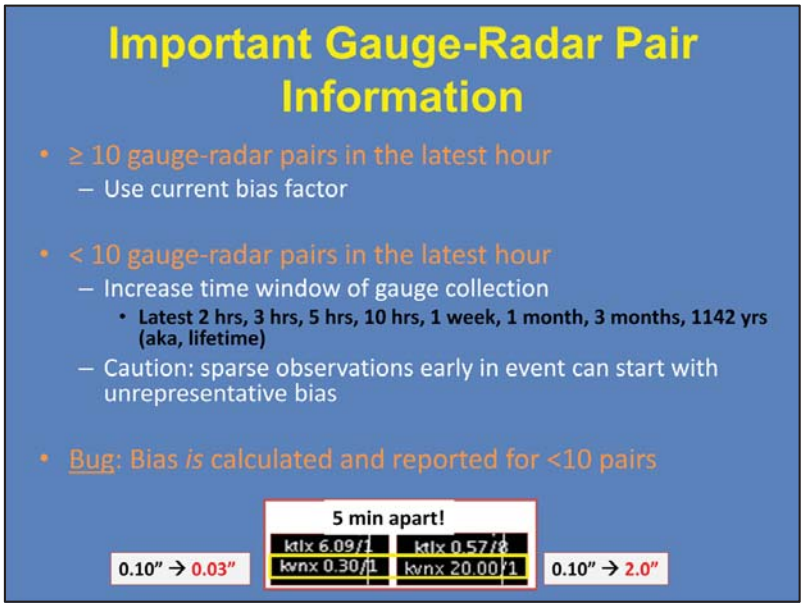


--There are three available bias sources for Bias HPE, and they are configurable by your local hydro focal point.
--First, you have the mean field bias for each radar, provided by the Mulitsensor Precipitation Estimator (or MPE) from your servicing RFC. This is the default bias source.
--The other option is to use your WFO's local MPE to choose from one of two biases.
--The first is a mean field bias per radar, similar to the RFC version.
--But the second option is what's called a “local bias”. Instead of one bias value being applied across the whole radar domain, this option applies different bias values across the domain. Each value is influenced by the nearby gauge-radar pairs, within a 40-km radius of influence. Some WFOs with more frequent flash flooding have found that the local bias has performed best.
--The image below shows an example where the bias source is the RFC mean field bias.



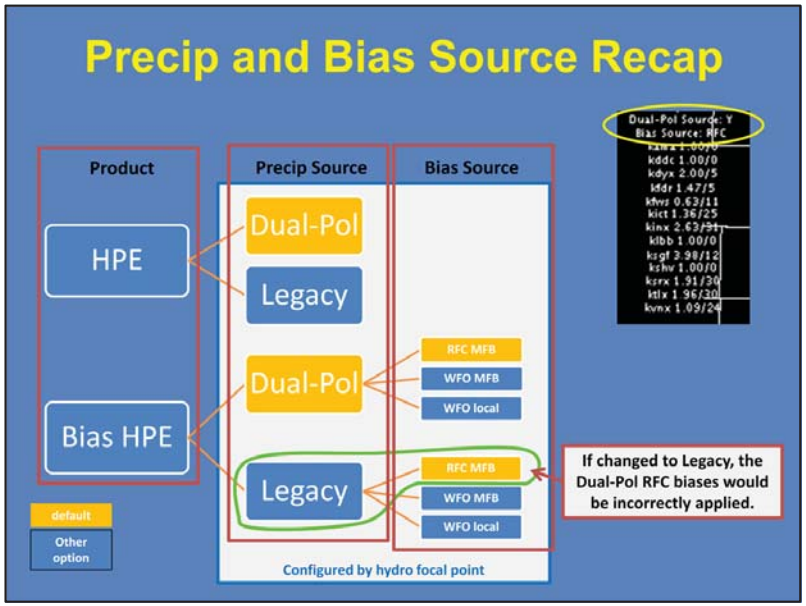
--In order to understand how to effectively use the new displays with Dual-Pol Bias HPE, you need to understand what the mean-field bias is. The mean-field bias is calculated using gauge-radar pairs of at least light precipitation. The sum of the gauge accumulations is divided by the sum of the radar bin accumulations. And outlier biases are removed.
--This bias information is calculated hourly for each radar that contributes to the Bias HPE mosaic.
--It can also be applied at the RPG and is displayable for your Legacy one-hour, three-hour, and storm-total precip estimates, but the default is that it is not applied. For more information on how to identify if the bias has been applied at the RPG, see the attached reference material on the VLab.

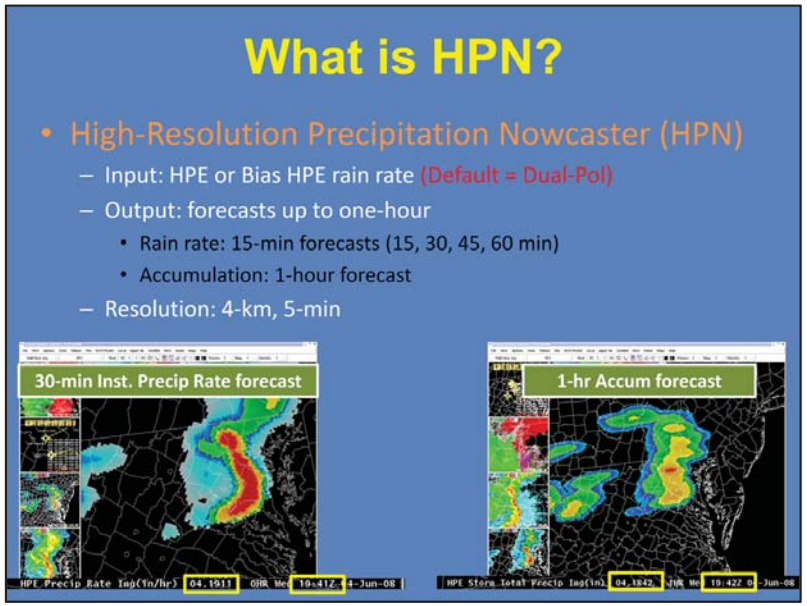




--Two of the most important components of the mean-field bias are, of course, the bias factor itself, and the gauge-radar pairs.
--The bias factor tells you if the radar's current precipitation source is under- or over-estimating rainfall, as compared to gauge measurements. It is a multiplicative factor, meaning this value is meant to be multiplied across the radar umbrella. For example, in this image, the bias factor is 0.46. Therefore, the DHR precipitation source is over-estimating precipitation, and the algorithm wants to multiply the QPEs by 0.46 to bring them down to where they should be. Similarly, a bias of 2.0 is saying that the QPEs are much lower than gauges measured, and therefore, should be doubled.
--The second component is the number of gauge-radar pairs used to calculate the bias. This is provided in order to add confidence to the bias factor. The more pairs used, theoretically, the more reliable the bias should be. In this example, there were 12 gauge-radar pairs used to get the 0.46 bias factor.

--There are a few important points to make about the gauge-radar pair information provided.
--First, how many gauge-radar pairs is significant to calculate a bias? The default number of pairs in MPE is 10.
--If there are 10 or more positive G/R pairs in the latest hour, then the current hour's bias factor is used.
--If there are less than 10 G/R pairs in the latest hour, then the algorithm increases the time window to collect more gauge information. If it reaches 10 or more gauges in the latest two hours, it creates a bias factor. If it is not reached by then, it will look at the latest 3 hours, and so on.
--A caution of this method is that if you are early in an event, and there hasn't been recent rainfall to trigger gauge accumulations, then you might be looking VERY far back in time to reach the 10 G/R pair threshold. Old G/R information, as much as 3 months or even the lifetime of MPE, may not be representative of the current event, thus creating an unrepresentative bias.
--A final note: with this methodology, bias information should never be reported using any less than 10 pairs. However, there is currently a bug that does calculate and report biases with less than 10 pairs, as shown below for KTLX and KVNX. These values were reported only 5 minutes apart. For KVNX for example, both scans only have one gauge reporting that affects the whole domain. That means, in the first scan, a radar estimate of 0.10" will be brought down to 0.03". But the very next scan, multiplied by 20, and brought up to 2". So be **very** careful interpreting bias information when there are less than 10 gauge-radar pairs reporting.





It can be a little confusing to keep track of all of the options between precip sources and bias sources. So here is a little flow chart to help summarize what is available. Default settings are in yellow.

--So you, as the forecaster, will have the option to choose between HPE or Bias HPE to load into FFMP, depending on if you want a bias applied to the mosaic field. From there, the precip source will either be Dual-Pol or Legacy, however, Dual-Pol is the default. If you choose Bias HPE, then one of three bias sources will be applied to the mosaic: either the RFC mean-field bias, WFO mean-field bias, or the WFO local bias. Here, RFC MFB is the default.

--All of these choices are configurable by your local hydro focal point. So as a forecaster, you will simply see what is configured in the product legend.

--If your local hydro focal point decides to change the default precip source from Dual-Pol to Legacy, it's important to understand that the biases being sent by the RFC are now dual-pol. Therefore, these Dual-Pol-derived mean field biases would be incorrectly applied to the DHR precip source. This will render the Bias HPE product unusable.

--The High-Resolution Precipitation Nowcaster (or HPN) is a QPF component of HPE and Bias HPE. Its intended use is to generate lead time in flash flooding for storms that are not rapidly changing.

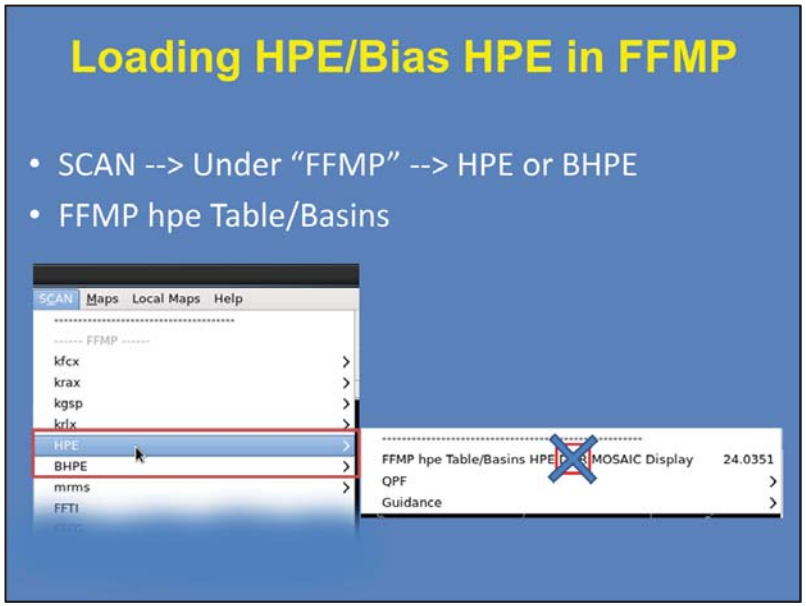
--This product uses its own feature tracking to extrapolate precipitation forecasts from an HPE or Bias HPE rain rate input. Therefore, the Dual-Pol default carries over into these HPN forecast products, as well.

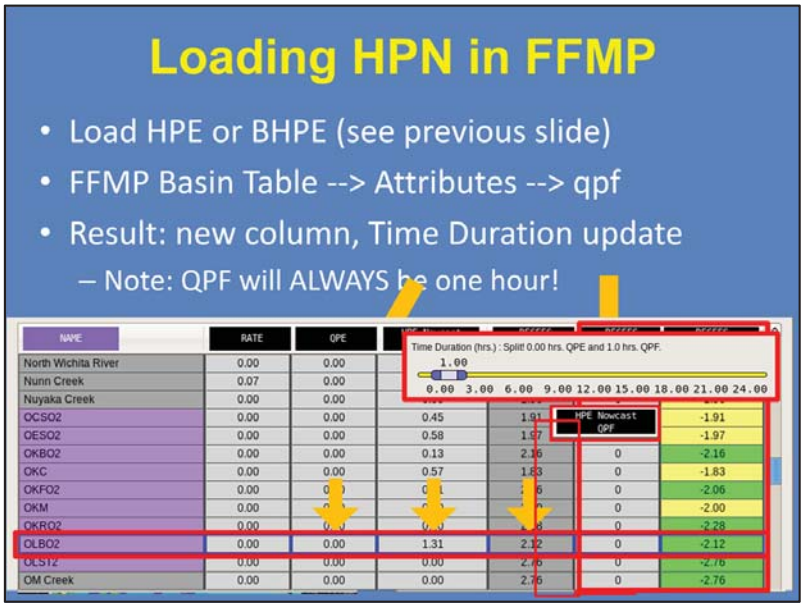
--The output is QPFs up to one hour in the future. For the rain rate products, four 15-min forecasts are created (at 15, 30, 45, and 60 minutes out). For the accumulated product, the rain rate products can be summed to create a single one-hour forecast.

--These products have a slightly lower spatial resolution, at 4km.

--Below is an example of each type of output from the Volume Browser. On the left, is a 30-min forecast of instantaneous precipitation rate for HPE. You can see the current time is 1911 and the product is valid at 1941Z. On the right is a one-hour accumulation, with the current time being 1842 and the forecast time as 1942Z.

--Both of these examples are forecasts of HPE, but remember, you can also look at forecasts of Bias HPE, as well.





--HPE was designed to make it easier to view multiple radar data in one product. This is especially useful in FFMP because you can load just one HPE mosaic, instead of multiple FFMPs for different radars.

--As a reminder, to load HPE via FFMP, go to the SCAN menu and click the HPE (or Bias HPE) sub-menu. From there, open the FFMP Table/Basins display.

--It is important to remember that the default HPE source is Dual-Pol, though, the label incorrectly says "DHR MOSAIC Display". To be correct, it should read "DPR MOSAIC Display".

--If you are interested in how to load these products, as well as the HPN grids, via the Volume Browser, please refer to the reference material.

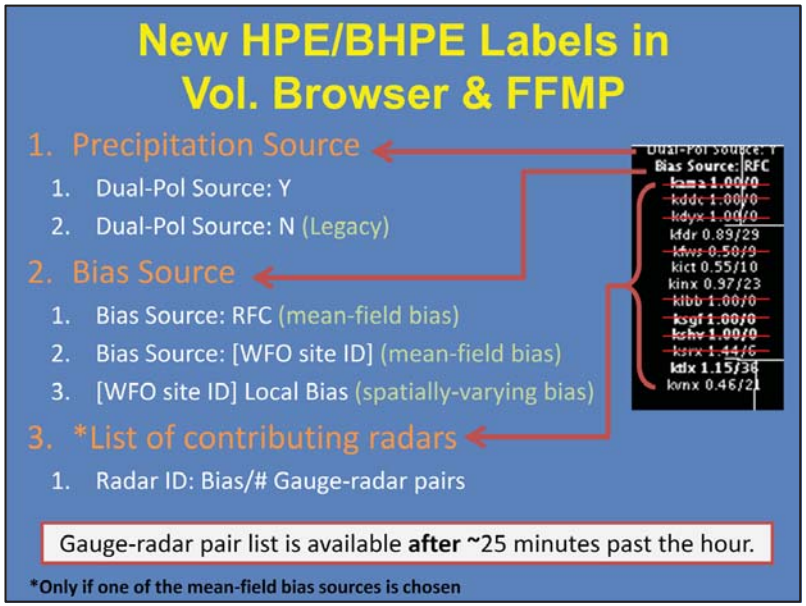
--So with the new default precip source being Dual Pol for both HPE and Bias HPE, the QPF extrapolation in the HPN forecast is also Dual Pol. Let's walk through how to load HPN QPFs in FFMP.

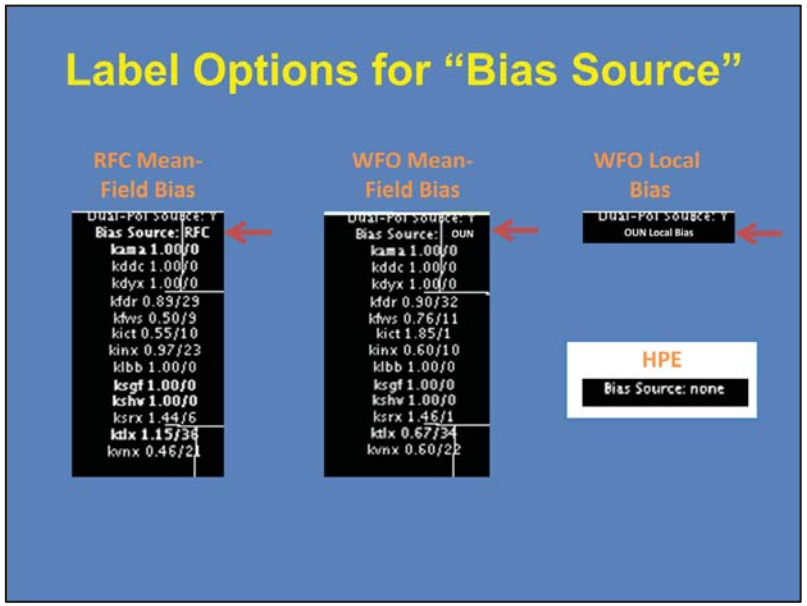
--Say we load HPE in FFMP (as described on the previous slide), and we click on the Attributes menu to turn on one of the QPF options. When the split window display comes up, we now have a 1 hour QPF combined with the QPE. Remember, the QPF extrapolation will ALWAYS be 1 hour, if this selection is made. It will never be more than 1 hour. We see the qpf column appear, and we can read off the forecasted precip one hour from now. Also, the Time Duration will tell you how many hours of QPE and QPF you are displaying, based on the slider bar total accumulation. In this example, since the slider bar is set for 1 hour, we have the 1-hour QPF in the "HPE Nowcast QPF" column, and there is no QPE displayed in the QPE column. If the slider bar was set to, let's say, 3 hours. Then the QPE column would have 2 hour accumulations, and the QPF column would remain the 1 hour forecast.

--In this example, the OLBO2 OKC west mesonet virtual gauge basin (highlighted with a purple name) had a QPF of 1.31", with no observed QPE due to storms just approaching the metro. The forecast was remarkably accurate (with 1.29" actually occurring), and it provided useful lead time information on some well-tracked storms. While the one-hour RFC FFG was only 2.12", this was enough precipitation to create urban flash flooding in the OKC metro.

--Unfortunately there is bug in FFMP of not incorporating the HPN QPF into the calculations for ratio and difference, but you can still get the QPF signal from the value in its column.

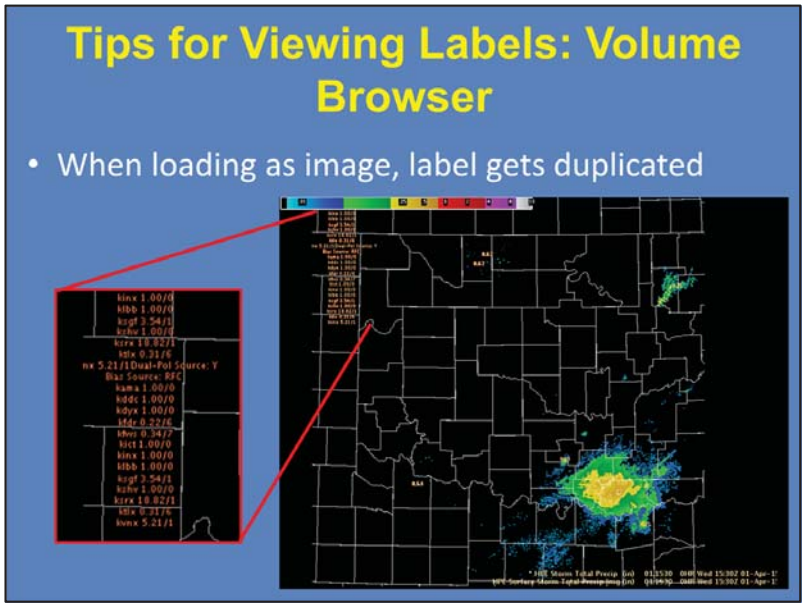
--For novice users, we recommend not using QPFs in FFMP because of the complexities discussed here.



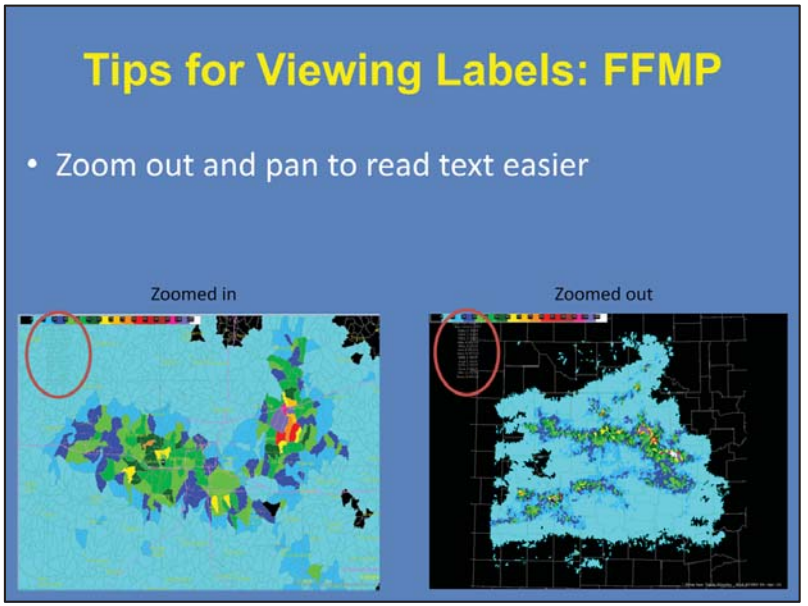


--With the option of Dual-Pol in HPE and Bias HPE, the labels on these products have changed to specify which precipitation source is being used, as well as adding more detailed bias information.
--The first line specifies the precipitation source. If it says "Dual-Pol: Y", then Dual-Pol is being used. If it is an "N", then Legacy is used.
--The second line specifies the bias source. Remember, there are three options here: if it is the RFC mean-field bias, it will say "Bias Source: RFC", like in this example. If it is the WFO's mean-field bias, then it will state the WFO's three-letter identifier. If it is the local spatially-varying bias, then it will specify the WFO's three-letter identifier with "Local Bias". The bias source is only relevant for the Bias HPE though it is displayed on both HPE and Bias HPE products.
--Finally, if one of the mean-field bias sources is configured, then the bias factor and number of gauge-radar pairs is listed. You can use this to understand how the bias factors are changing the Bias HPE precip estimates each hour throughout your mosaic.
--Remember, have higher confidence in precip estimates when you have more than 10 gauge-radar pairs, and if you have < 10 gauge-radar pairs you should not trust the precipitation estimates in those regions of your mosaic as they likely indicate a bug in the bias factor creation.
--An important note: Your current hour's gauge-radar pair list will NOT appear until after about 25 minutes past the hour. This is due to known latencies in gauge data collection. Therefore, after 25 minutes past the hour, the product overlay will automatically update on the current frame. But if you want to see the list for the earlier frames in the hour, you can clear and reload the HPE or Bias HPE product.

--So here are examples of how the three different bias sources display in Bias HPE.
--The first option is using the RFC MFB as the source. You will get "RFC" as the source, and a list of radars with their respective bias information.
--If the bias source is the WFO MFB, then the three-letter identifier is given (in this example, OUN), and the WFO MPE bias information is listed for each radar.
--Finally, if the bias source is the WFO's local bias, then it will say the WFO's identifier, followed by "Local Bias" and nothing will appear underneath. If you recall from the earlier slide, this is because this option does not offer a uniform bias for each radar.
--And since there are no biases applied for HPE, the bias source label simply says "none".



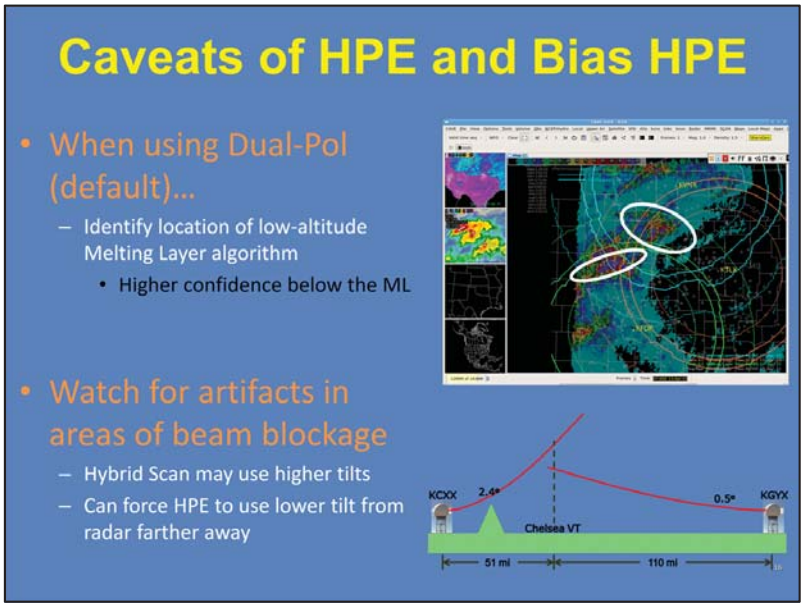
--There are a few things to be aware of when using this version of HPE and Bias HPE.
--When viewing in the Volume Browser, the label gets duplicated when you "Load as Image". This can make it a little difficult to read, but none of the text is overlapping, so the important information is still visible.



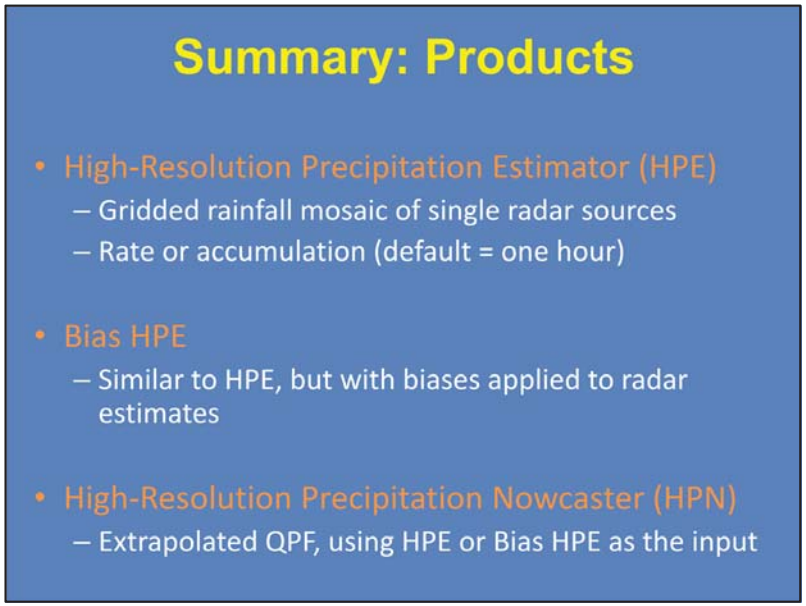
--Second, make sure to zoom out and pan to the left border in order to better read the label. When zoomed in, the label may be obscured by the basin fill. But if you zoom out, you can move it to a blank area to read easier.

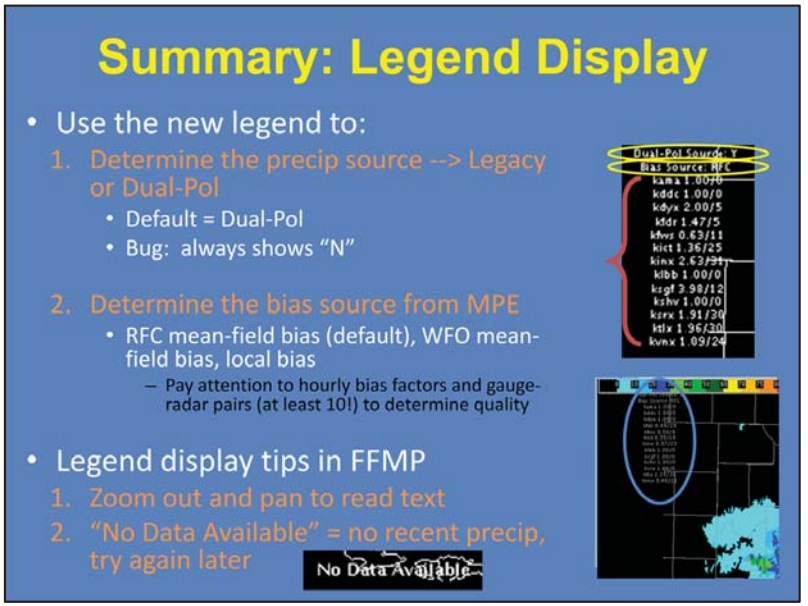


If you try to open HPE or Bias HPE when there has been no rainfall in your domain, you will get a blank pane with "No Data Available" in the text-legend in the bottom-right. This does not mean that HPE is broken! It simply means that the grids have not been created due to the fact that there has not been any recent rainfall. Just wait until the next rainfall occurs in the domain, and try opening again.



There are some caveats to using HPE and Bias HPE.
--For one, when using the dual-pol precipitation data in HPE, you should identify where the low-altitude melting layer algorithm begins to identify the melting layer. Have higher confidence in estimates below the melting layer. Within and above the melting layer, your estimates are going to change, and they won't be as good. This is due to the HCA assigning mixed phase or ice classifications at these heights, and thus, the Z-R relationships are being matched to those precip types, even if they are liquid when hitting the ground.
--On the right is an image of HPE. Let's consider the circled areas that may be experiencing flash flooding. By overlaying the melting layers for the three nearby radars, you can see the northeastern circled area is closest to the most northern radar, or the one with the blue ML circles. These values are located below the ML, so you can have higher confidence in them. However, for the southwestern circled area, it is within the green ML circles for the most southern radar. Thus, this area has more uncertainty. Checking the HHC to locate liquid precip types in the ML can further add to your confidence in this area.
--Secondly, watch for artifacts in areas of beam blockage. At times, your hybrid scan may be using elevations above the 0.5 degree slice in areas of beam blockage. But for QPE estimation, the lowest altitude estimate is the best. If there is a neighboring radar that uses the 0.5 degree tilt to cover the same area, that estimate may be better (even though the radar is farther way). You can force HPE to use the farther radar. So for more information on this topic, please see the Resources tab for the presentation from Greg Hanson, the service hydrologist at the Burlington, VT WFO.





So let's quickly summarize the products we discussed in this lesson.

1. The High Resolution Precipitation Estimator (or HPE) is a gridded rainfall mosaic, based on single radar sources. It can be created as an instantaneous rain rate or accumulation field, with the default accumulation being one hour.
2. The next product is the Bias HPE. This is similar to HPE, except that biases have been applied to the radar estimates to correct for uncertainties. These biases are calculated using available gauge data.
3. Finally, there is the High Resolution Precipitation Nowcaster (or HPN). This product is an extrapolated QPF, based on either the HPE or Bias HPE product. The forecast is available for 15-minute intervals up to one hour, or for a single one hour accumulation.

All three of these products update every 5 minutes.

The legends for HPE and Bias HPE can be used to interpret valuable information about the products.

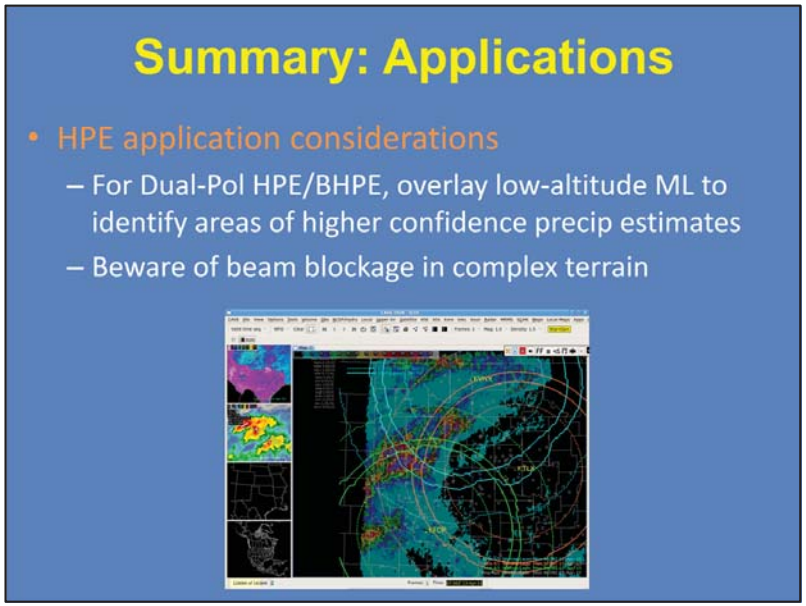
--For one, you can now see which precip source is being used, either Legacy or Dual-Pol. The default is to use Dual-Pol. Keep in mind there is currently a bug where this label will always say "N", even though most offices will have Dual-Pol configured.

--For Bias HPE, you can also determine the bias source from MPE. The default is the RFC mean-field bias, but you can also use your WFO mean-field bias or your WFO local bias.

--To get the most understanding out of these biases, make sure to pay attention to the hourly bias factors, as well as the number of gauge-radar pairs to determine the quality of the bias estimates.

--In FFMP, you can make viewing these legends easier by zooming out and panning to the edge to better read the text.

--Also, if you see a blank pane and "No Data Available" when loading in FFMP, it means there hasn't been enough recent precip for the HPE grids to be calculated. Simply try loading again once rainfall has occurred in the area.



In addition to the updates to the legend, there are some important takeaways to better apply HPE and Bias HPE.
--For Dual-Pol HPE or Bias HPE, you should overlay the low-altitude melting layer algorithm to identify areas of higher confidence precip estimates. Also, beware of beam blockage in complex terrain, as it may affect the QPE estimates.



Hi, my name is Jill Hardy and welcome to this lesson. This is Part 1 of 2 about the flash flood decision-making process, with this lesson focusing on the general approach, as well as how to best choose your precip sources.

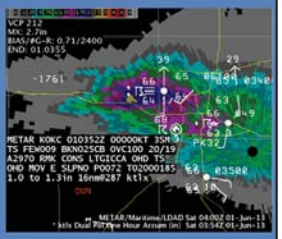
Learning Objectives

- By the end of this lesson, you will be able to:
  - Identify basic approach to flash flood decision making
    - Determining precip amounts, timing, hydro response, calibration
  - Identify how to choose a precipitation source

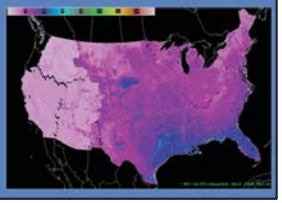
Here are the learning objectives for this lesson. When you have finished reading them, please continue to the next slide.

Flash Flood Decision Making Approach

- How much rain has fallen and when did it fall?
  - Best QPE
    - Compare to observations
    - Choose from multiple precip sources

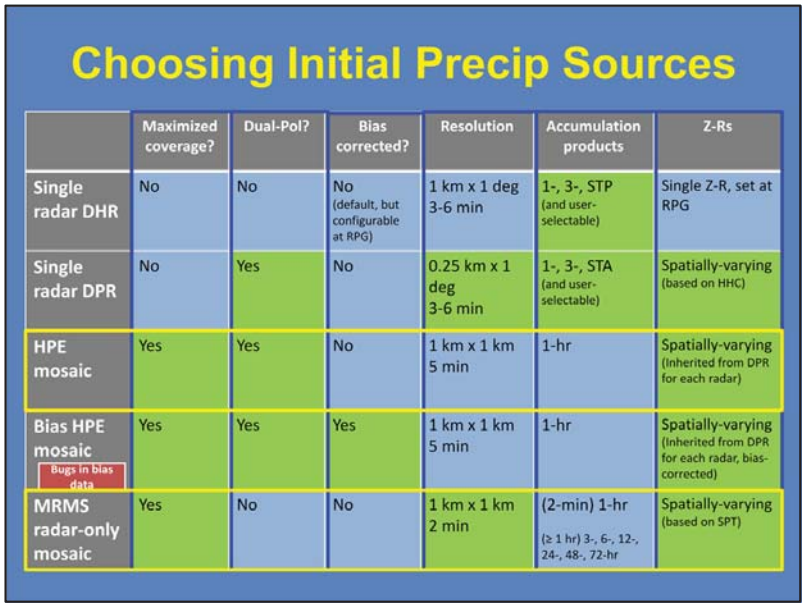


- Will runoff cause flash flooding?
  - Local knowledge
  - RFC FFG



The first question you should ask yourself is "How much rain has fallen and when did it fall?". To do this you will need to identify the QPE source that compares most favorably to surface observations and reports. There are multiple QPE options to choose from, so the next slide will step through the strengths and weaknesses of each to help you make an informed decision.

Next, you need to consider if runoff will cause flash flooding. Typically this starts with your local hydrological knowledge and is supplemented with River Forecast Center (RFC) Flash Flood Guidance, available at most CONUS WFOs.



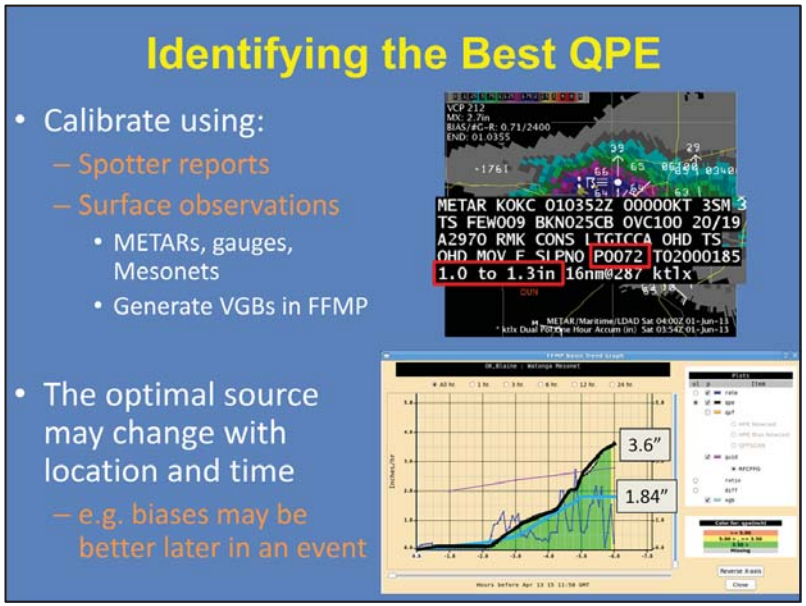
There are many factors to consider when choosing a precip source to use during an event, particularly early in an event before you have many observations.

1. For one, consider the coverage area. The mosaic HPE and MRMS products are good when you have multiple radars because you only have to open one grid or FFMP.
2. Next, is it Dual-Pol? Both HPE products default to using Dual-Pol precip sources. It is important to note that MRMS does NOT use Dual-Pol to calculate rain rates at this time, only for quality control purposes.
3. Do you want biases applied? By default, biases are only applied for Bias HPE. However, for DHR, the option is configurable by your local radar focal point. It is important to consider the quality of these biases. Only use these bias-adjusted products in areas where the biases make sense with reports you are getting from spotters or gauge reports you have confidence in. Note that MRMS does have bias-corrected products, but they are not timely enough for real-time flash flood use.
4. How do the resolutions compare? The single-radar DPR has the best spatial resolution at 250 m by 1 deg. This is useful for isolated events that are close to the radar. However, MRMS has the best temporal resolution, at 2 minutes. This is useful for convective events with high rain rates.
5. Consider what accumulation products are available outside of FFMP. The single radar sources offer 3-hour and storm total accumulations. Additionally, they offer the ability to do one-time requests of user-selectable accumulations. On the other hand, the HPE products only produce 1-hour accumulations by default. And, MRMS only produces the 1-hour accumulation every two minutes, with all of the other accumulations being one-hourly or longer.
6. Finally, how are rain rate relationships managed during product creation? DHR only has one Z-R, set at the RPG. Every other source uses spatially-varying Z-R relationships based on precip type.

While generally it may seem that mosaicked products are preferred for their spatial coverage, there are times when we recommend DPR or DHR. They are good choices if either HPE or MRMS have artifacts, like when multiple radars are covering complex terrain. When you simply want to view more than one source, like looking at both Legacy and Dual-Pol. Or if there is an isolated event close to a radar, which may provide better spatial resolution information.

So now, which precip sources should you use? First, there is one caveat worth noting. The Bias HPE product can be buggy due to problems with the biases being sent over from the RFCs, so only use Bias HPE if you carefully analyze and agree with the bias factors used.

Therefore, our current recommendations are: the HPE mosaic product, since it has great coverage, is Dual-Pol, and has spatially-varying Z-Rs. And, MRMS for its coverage, temporal resolution, and spatially-varying Z-Rs. From here, routinely check these sources against observations to see if you need to make any changes.

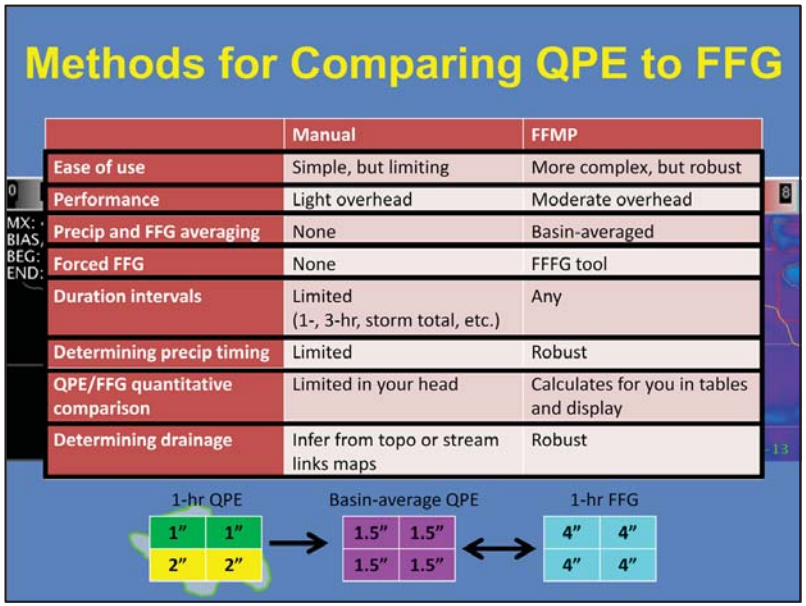


Choosing the right precip source doesn't end after your initial selection. You should routinely compare your precip sources to spotters and surface observations throughout the event, particularly in potential flash flood areas. You can sample surface observations by checking METARs, gauges, or Mesonet stations, as well as generate Virtual Gauge Basins in FFMP.

Here is an example that shows the one-hour accumulation product overlaid with METARs. We see the METAR estimates 0.72 inches, while the 1-hour DPR estimates 1.0 to 1.3 inches during the same time period. This shows that DPR is overestimating precip by a quarter to half an inch. You can use this information to self-calibrate the QPE estimate, and when considering other precip sources.

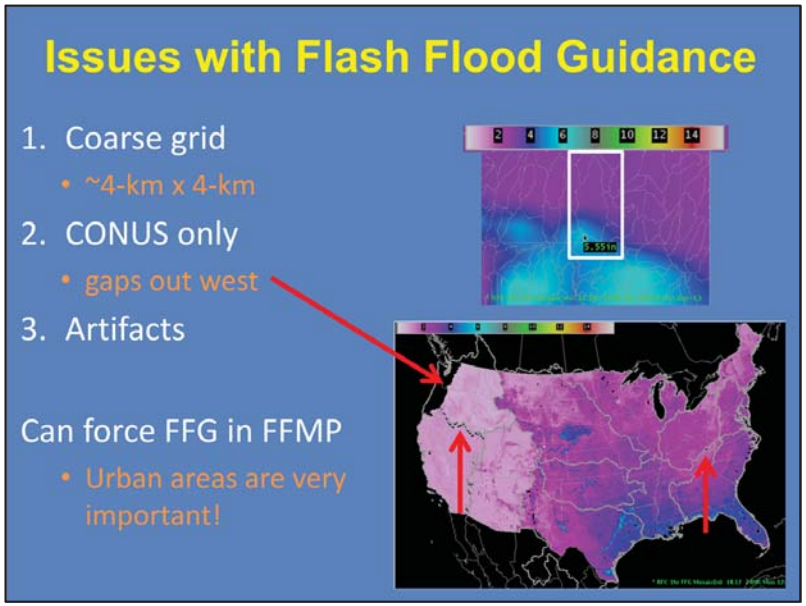
The Virtual Gauge Basins (or VGBs) in FFMP are a convenient way to compare a QPE source to METARs and mesonet observations, though the hourly updates of the METARS/mesonets can take a little getting used to. This image is for the Watonga Mesonet site in Oklahoma. The mesonet accumulation is indicated by the light blue line, and shows that the Mesonet accumulated 1.84 inches in the last 6 hours. The DHR VGB source, denoted by the black line, shows that radar estimated around 3.6 inches during the same time period at that location. In this case, the radar drastically overestimated the precip compared to the Mesonet observations, and it is easy to interpret using the VGB graph.

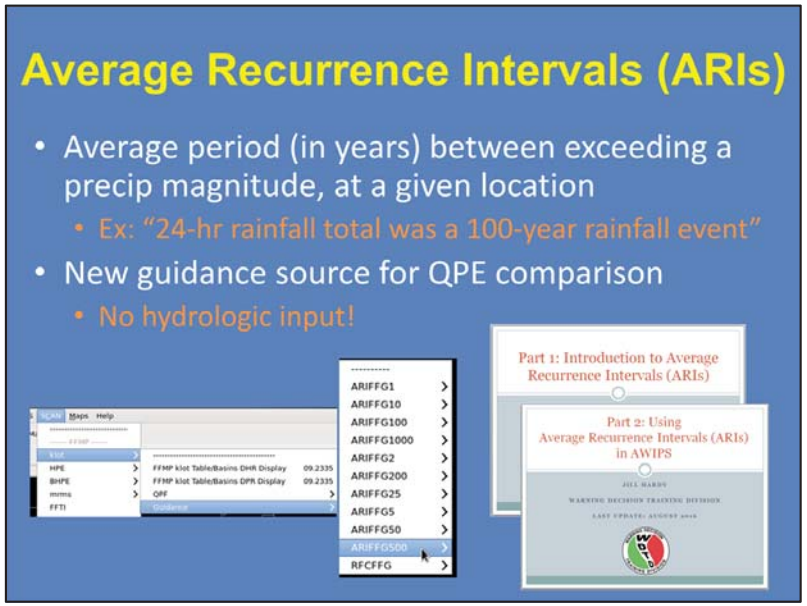
It is not unusual for the optimal precip source to change across your CWA and with time. For instance, once the bug is fixed in Bias HPE's raw data calculations, it typically improves over time as more gauge information is available. Routinely checking the precip estimates with gauges and reports is good best practice.



Okay, so you've identified the best QPE. The next step is to determine if the runoff will cause flash flooding. This is where you use your local rules of thumb for flashy basins or you compare QPE to the RFC Flash Flood Guidance. In D2D, you can interrogate QPE and FFG in two ways. For one, you can load the radar accumulations and RFC FFG manually from the radar and NCEP/Hydro menus, as shown here, and do a one-to-one comparison. The other option is to use FFMP. Here is a table outlining the main differences:

1. Loading products manually is simple, but has many limitations. FFMP is more complex, but it has more robust capabilities.
2. Performance overhead on the machine is light with manual loading, and moderate with FFMP.
3. One of the fundamental differences is that FFMP averages both the QPE and FFG over the basin area to provide more hydrologically-relevant calculations. To show this, let's start with a basin. Overlaying the corresponding QPE, we see there are spatial variations across the basin. FFMP averages those values, and then compares them to basin-average FFG of the same interval.
4. FFMP also has a tool to create your own FFG, but there is no way to do that in manual loading.
5. In manual loading you are limited to fixed intervals like 1-hr and storm-total precip which can make it hard to determine accumulations for durations like 2 or 4 hours. Some products like HPE or Bias HPE are limited to only 1 hour accumulations in their raw products by default, so it can be very difficult to just figure out a 3-hour accumulation for HPE and Bias HPE. Because FFMP accumulates precip every time a rate product is ingested, one of its greatest strengths is the ability to display any accumulation duration.
6. This accumulation-on-the-fly approach also allows FFMP to robustly display the precipitation timing information in the basin trends. When manually loading, it can be difficult to determine when precip occurred with training storms over long periods of time.
7. Another fundamental strength of FFMP is that it calculates QPE and FFG ratios and differences for you, whereas you have to manually calculate that for a limited number of points when you load the products and sample them.
8. Finally, when loading products manually it can be challenging to infer drainage from topo and stream links maps, but FFMP will show you downstream areas graphically.





If you have an idea of how much rain has fallen and you are going to use RFC flash flood guidance, it is important to note some of the challenges.

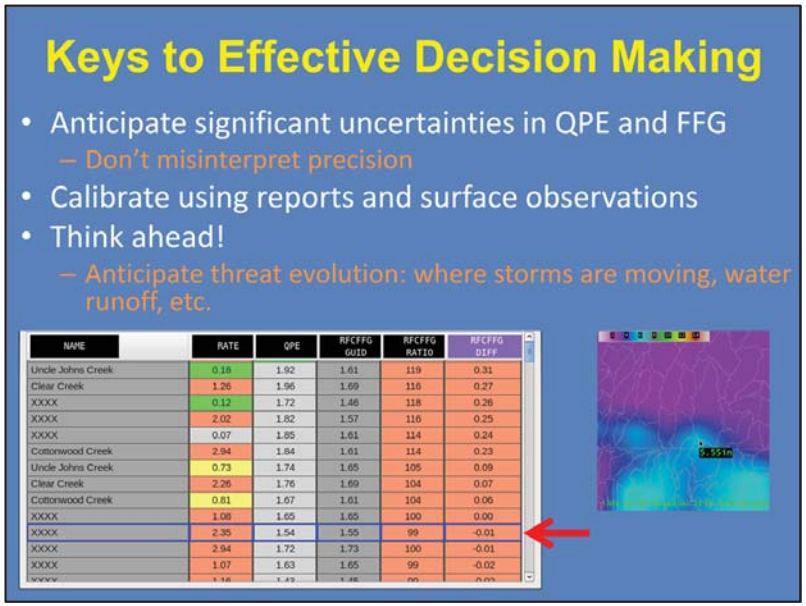
1. For one, RFC FFG has a rather coarse grid (approximately 4km by 4km) relative to the small size of many basins. Take this zoomed in 1-hr FFG product. Over this particular basin, FFG ranges from a little over 2 inches to 5.55 inches, likely due to two grid point observations that are not resolving the basin-scale detail. FFMP is going to average this out to one number, which may not adequately represent the hydrology over this basin.
2. Second, it's coverage is only over the CONUS, with some gaps out west, as shown here.
3. Finally, there are artifacts along some RFC boundaries where different methods of calculating guidance result in non-realistic sharp gradients, which we can see when overlaying the boundaries.

An advantage of using FFMP over manual comparison is: if you do not have RFC FFG or it isn't optimal, you can force it in FFMP. This is particularly useful in important urban areas which are not represented in RFC FFG. Overlaying your local "Urban Bounds" maps is a good practice to help locate urban flash flood-prone areas.

New to AWIPS, as of 16.2.1, are Average Recurrence Intervals (or ARIs). An ARI is defined as the average period (in years) between exceeding a precip magnitude, in a given location. You are probably more familiar hearing ARIs used like "Yesterday's 24-hour rainfall total was a 100-year rainfall event".

We bring ARIs up here because they are a new guidance source for QPE comparison, and you will see them when maneuvering around AWIPS and FFMP. For instance, here is a SCAN FFMP menu. Under the Guidance submenu, it's no longer just RFC FFG, but rather a plethora of ARI data, as well. One of the most important things to remember when using this dataset is that it was created solely as a measure of precip rarity. It does not include any hydrologic inputs. Therefore, it is fundamentally different than Flash Flood Guidance, even though they appear in a similar fashion throughout AWIPS and FFMP.

Before you begin using this dataset, we highly recommend taking the two ARI lessons released by WDTD in August 2016. There are a lot of details covered in these lessons that are crucial to effectively interpret this information. These lessons are included in the Warning Operations Course (or WOC) Flash Flood Track, if you plan to enroll in this course after RAC.

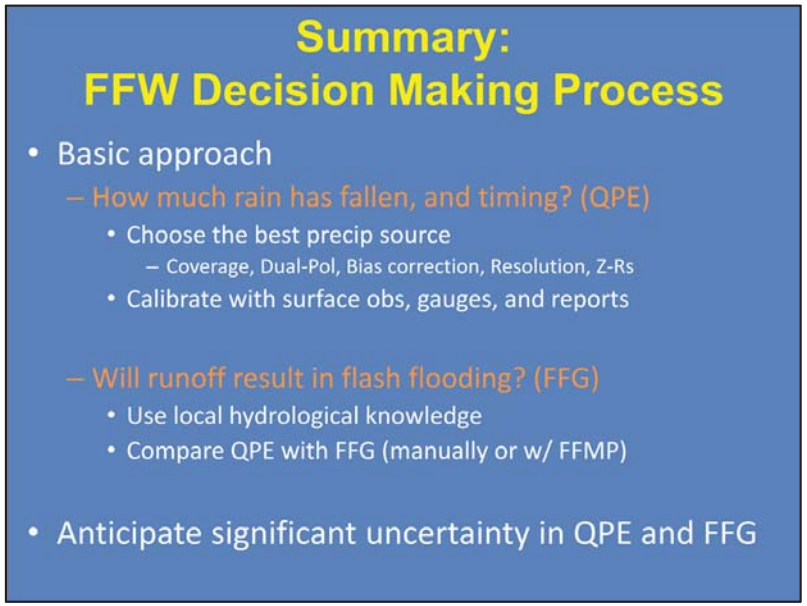


Now that you have been exposed to the general approach to flash flood decision making, there are several keys to making it as effective as possible.

For one, anticipate significant uncertainties in QPE and FFG. It is not uncommon to encounter QPE and FFG uncertainties on the order of 25%, or even 50% at times. However, when doing manual comparison or using FFMP, the values are shown with two decimal places. Don't misinterpret this precision. For instance while the selected basin may be 0.01" below FFG this could easily be a quarter inch above or below FFG due to uncertainties in the raw QPE or FFG data.

Therefore it is important for you to be routinely calibrating QPE using reports and surface observations, keeping in mind that surface obs can also have uncertainty.

Another key to effective decision making is to always think ahead. It is easy to become fixated on the complexities of what is going on now with tools like FFMP. Anticipate threat evolution by considering where the storms are moving and what the hydrological conditions will be in those areas. This will give you important lead time when drawing your FFW polygons.

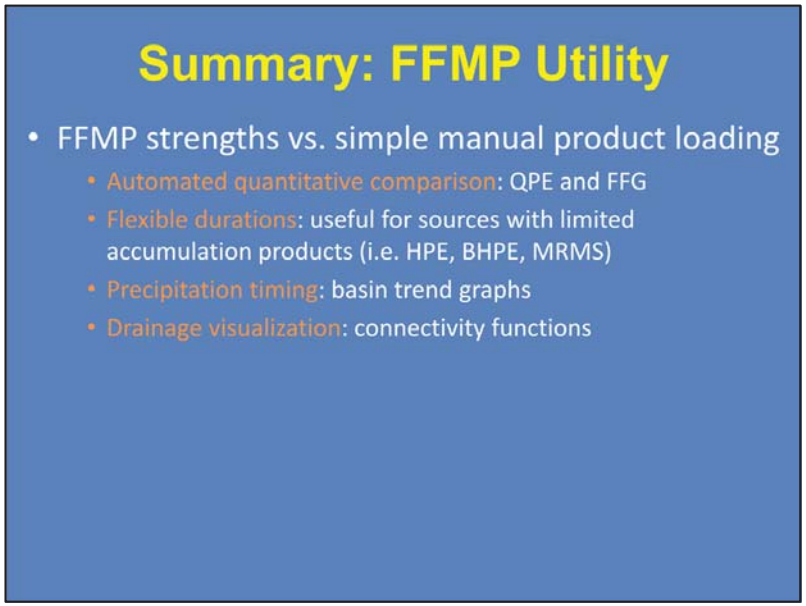


To summarize, the basic approach to flash flood decision making first involves assessing how much rain has fallen, and when. To do this you need to evaluate multiple precip sources and choose the best precip source based off factors like:
-The coverage of the product.
-Is the source Dual-Pol?
-Whether there is a bias correction, and if the biases seem reasonable.
-What is the resolution?
-And how are the rain rate relationships calculated?

Just as important as your initial precip source selection, is to routinely compare all precip sources with surface observations, gauges, and spotter reports. The best precip source can change over the course of an event.

The second part of the process involves comparing your QPE with rules of thumb for local flashy basins, or to the RFC FFG. While you can evaluate some of the raw data with manual loading of precip products in D2D, FFMP provides more robust ways to interrogate multiple times and durations.

Always anticipate significant uncertainty in QPE and FFG. QPE uncertainties exist due to numerous radar sampling limitations. FFG is coarse, only over the CONUS, and has artifacts across some RFC boundaries.



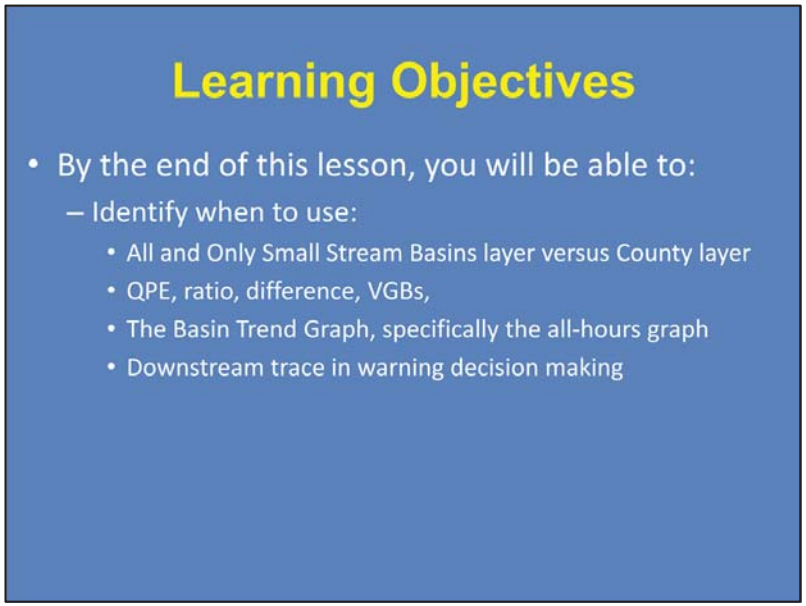
And finally, because FFMP has a number of unique strengths relative to simple manual product loading, Part 2 will focus on how to effectively use it.

The strengths include having automated quantitative comparison of QPE and FFG for any duration. It has flexible durations, which is useful when you want to consider accumulations outside of what is available as a gridded product. Precip timing is also more readily available using the Basin Trend Graphs. And you can better visualize drainage flow using the basin connectivity functions.

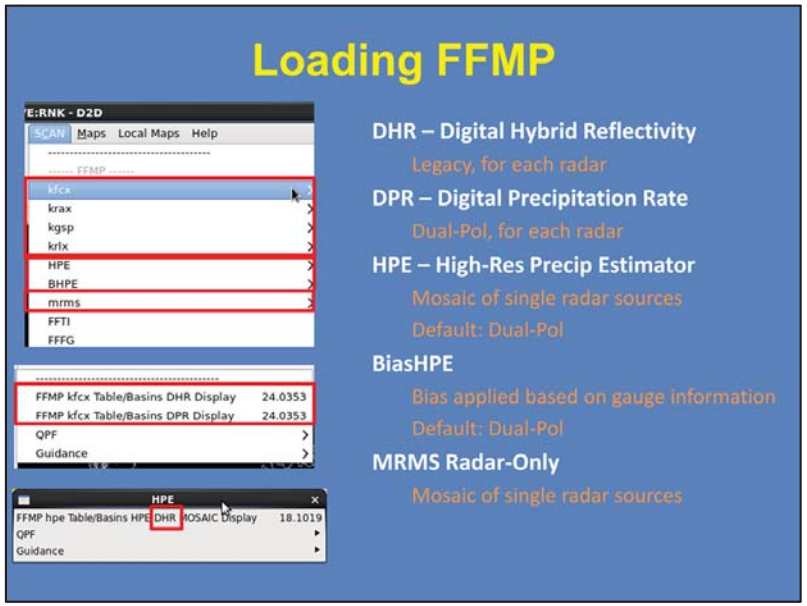
Alright, that's the end of Part 1 of the Flash Flood Decision Making lessons. When you are ready, move onto the next slide to take the quiz and receive credit on the LMS.



Hi, my name is Jill Hardy and welcome to this lesson. This is Part 2 of 2 about flash flood decision-making. In Part 1, we introduced the general approach to flash flood decision making, as well as compared the different precip and guidance sources available. This lesson will now focus on using FFMP to diagnose a flash flood threat.



Here are the learning objectives for this lesson. When you have finished reading them, please continue to the next slide.

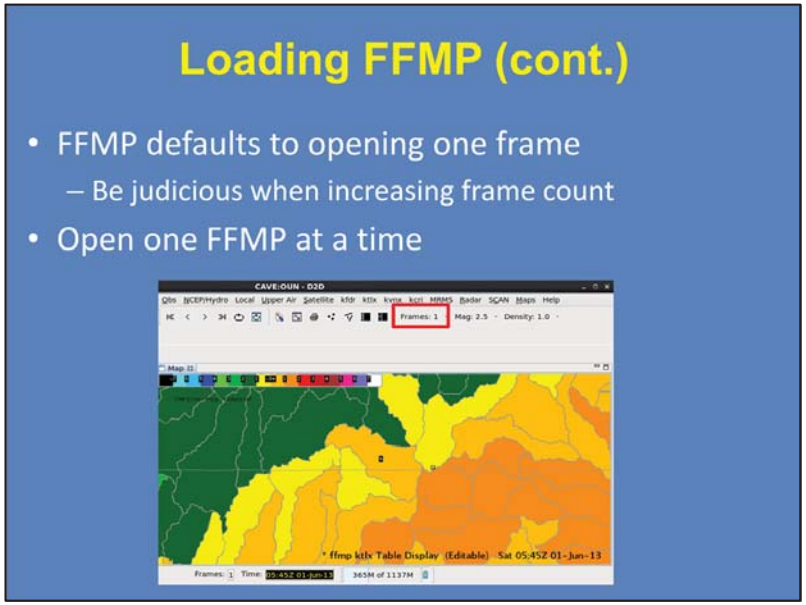


First thing's first: loading FFMP with your desired precip source from the SCAN menu.

The single-radar products are available under the menu referenced by the radar name. As seen here, under each radar submenu, there is the DHR source (which is Legacy), and the DPR source (for Dual-Pol).

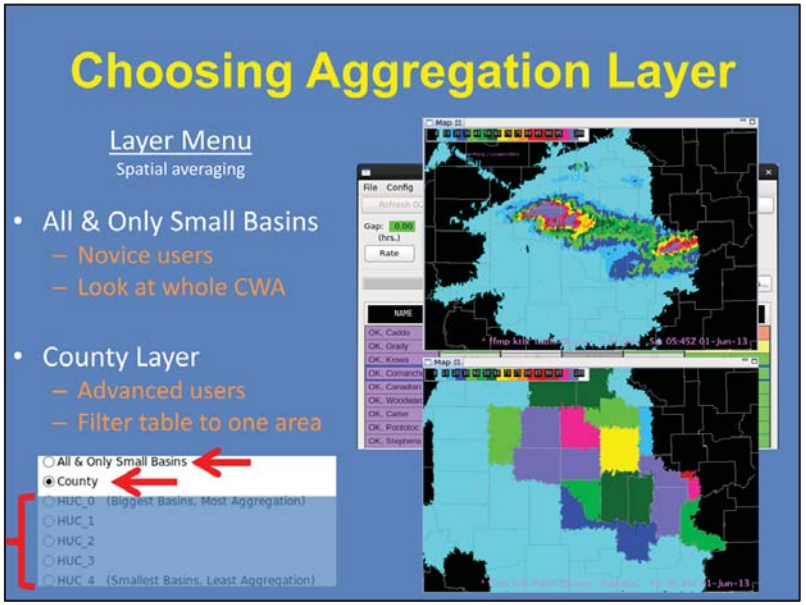
Next, since the HPE and Bias HPE products are mosaics, they are identified as HPE and BHPE on the SCAN FFMP menu. Keep in mind the labeling error for these products. Both use Dual-Pol in their creation, however, their submenu reads "DHR MOSAIC". Just be aware that this is a typo.

Finally, the MRMS Radar Only mosaic is also available from the SCAN FFMP menu.



Once you have decided which precip source to use, you can load FFMP with that source. Note that FFMP defaults to loading just one frame. This is because FFMP uses a lot of system resources. You can increase this manually, just be judicious for performance reasons.

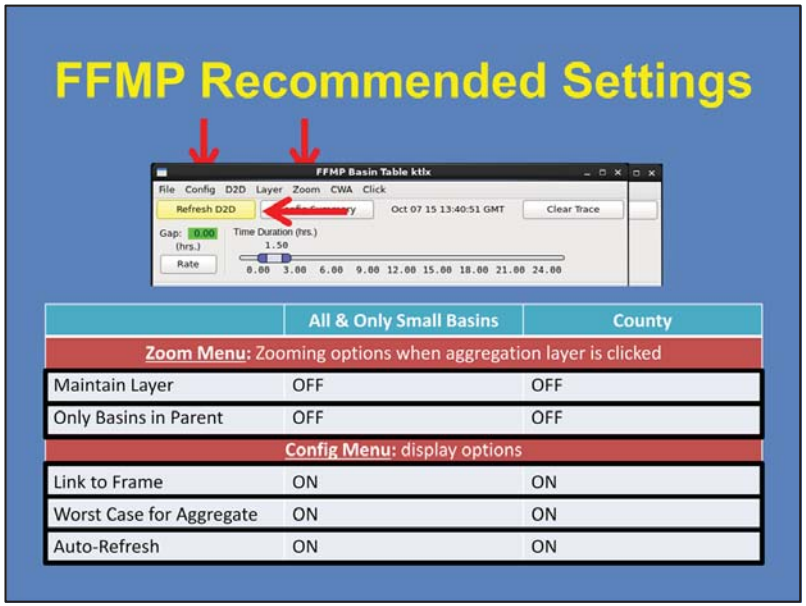
Also, while it is possible to load more than one FFMP at a time (using a different precip source), it can be hard to keep straight which table goes with which D2D display, not to mention adding to potential performance issues. So we don't recommend doing this. Rather, use reports and observations to calibrate whether you need to change your FFMP source.

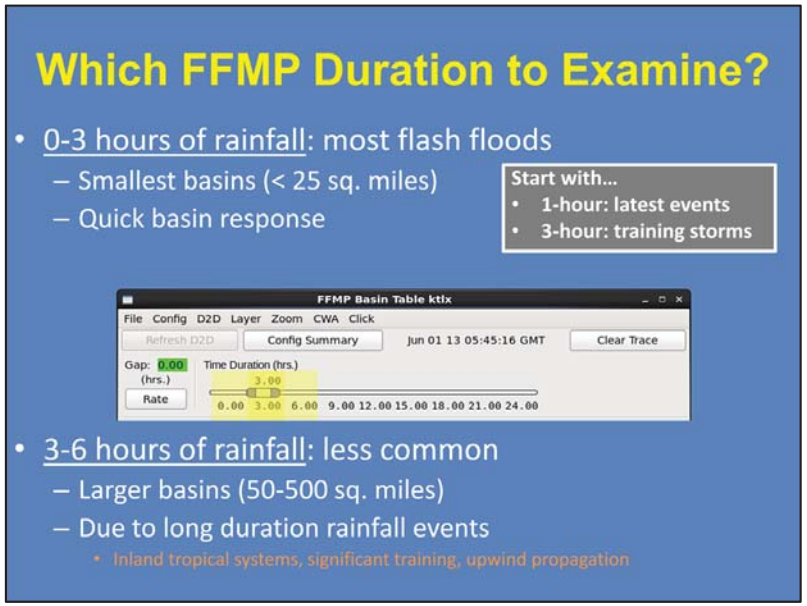


One of the most important steps in using FFMP is choosing the aggregation layer, which defines what spatial scale FFMP averages the QPE and FFG. The two layers we will focus on for flash flooding are the All & Only Small Basins and County.

When you first open FFMP and begin your flash flood interrogation, we recommend starting with the "All and Only Small Basins" layer option. This will give you a simple look over the whole CWA on the most relevant hydrologic scale to see what basins stand out in QPE, ratio, or diff.

As you become more advanced with FFMP, you can switch to using the "County" layer option to organize the basins in the FFMP table by county. This is done to make it easier to find particular basins and virtual gauge basins that allow you to overlay the METAR/mesonet observations in the basin trends. There are a lot more settings to pay attention to when using county layer, though. The HUC layers are collections of small stream basins for larger scales and are not used frequently for flash flood decision making.





There are several FFMP Table menu options that can enhance your D2D display, as well as your FFMP basin table display.

First, the Zoom menu controls how FFMP zooms into smaller basins when an aggregation layer is clicked in the table. In either "All and Only Small Basins" or "County" layer, we recommend these options be turned off. They do not have an effect on your display when using "All and Only Small Basins". But if you use "County" layer, with the "Maintain Layer" option OFF, the D2D will not maintain the county layer and will instead show the individual basins. Setting "Only Basins in Parent" to OFF with the "County" layer allows any neighboring basins outside of the county to be displayed in D2D, so you can see flash flood threats crossing the county line.

Next, the Config menu helps with general display of the data. The "Link to Frame" ensures the D2D and the table are kept in sync when stepping through multiple frames.

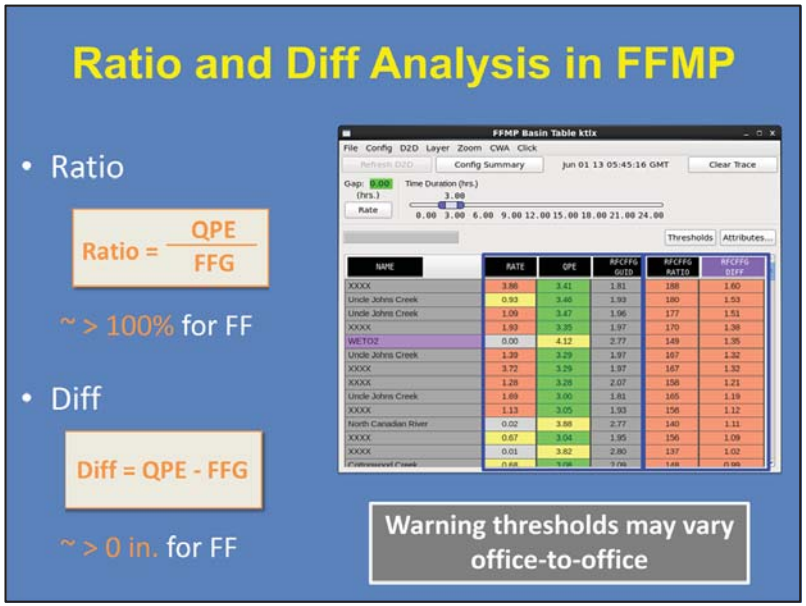
Next is the "Worst Case for Aggregate" option. This option only comes into play when you have chosen a layer larger than "All and Only Small Basins", like "County". When turned on, this option sets the values in the FFMP table to show the "worst case" value for **any** basin within the larger aggregate layer. However, be aware that the "worst case" values may not always be within the same basins.

Finally, there is the Auto-Refresh option. This automatically updates the D2D display with any configuration changes made to the FFMP table. We recommend this be turned ON. However, if you notice performance problems with FFMP, you may consider turning it OFF. When it is off, you need to remember to click on the "Refresh D2D" button to manually update the display after making changes.

Now that we have reviewed the FFMP settings, let's talk about what durations are good to examine.

The majority of flash flood events take place due to less than 3 hours of rainfall, sometimes less than an hour of heavy rainfall. This is because they occur in basins smaller than 25 square miles. These tiny basins have quick responses to the rainfall, and thus inundate rapidly. For RAC, we suggest you focus on the 1-hour duration from the FFMP basin table for the latest events, and the 3-hour duration for training storms. The duration slider bar feature in FFMP makes this easy to do.

However, certain meteorological environments are conducive to flash flooding larger basins, say 50-500 square miles, and thus require a longer duration to get things going. Inland tropical storms, significant cell training, and upwind propagation along a quasi-stationary boundary are examples of long duration heavy rainfall events that may result in flash flooding of large basins. In these types of set-ups, in addition to looking at 1 and 3 hour duration, it would be wise to also check out the 6 hour duration information from the basin table.





Okay, so you have your settings the way you want, and you know the duration you want to examine. So what should you look at? In addition to instantaneous rate, basin-averaged QPE, and basin-averaged FFG, FFMP has two other options for what can be displayed in the table: Ratio and Difference.

By default, these two take into account the Flash Flood Guidance values, and thus, are useful for analyzing exceedance threat. So let's take a look.

For Ratio, it is QPE divided by FFG. So, as Ratio approaches and exceeds 100%, that means QPE is near or exceeding FFG, and thus, the theoretical flash flood threat increases. To calculate the Difference, it is QPE minus FFG. So as Diff approaches zero or becomes positive, similarly, the theoretical flash flood threat increases. But remember, all QPE sources have uncertainty and RFC FFG accuracy varies significantly across the country and over time. So you may find that ratio and difference warning thresholds vary from office-to-office. For instance, at some WFOs, flash flooding may typically start at 0.5 inches over FFG while another may start near flash flood guidance.

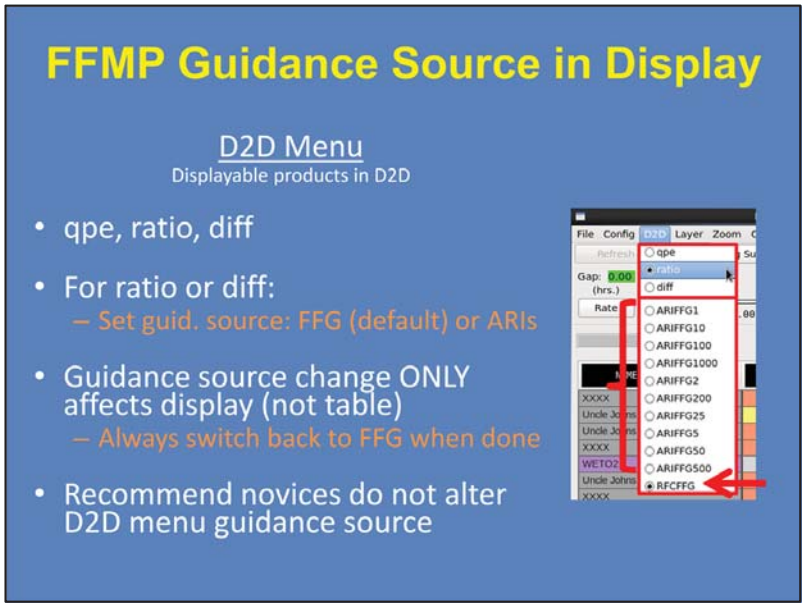
But let's take a minute to review an example of QPE, ratio, and difference values.

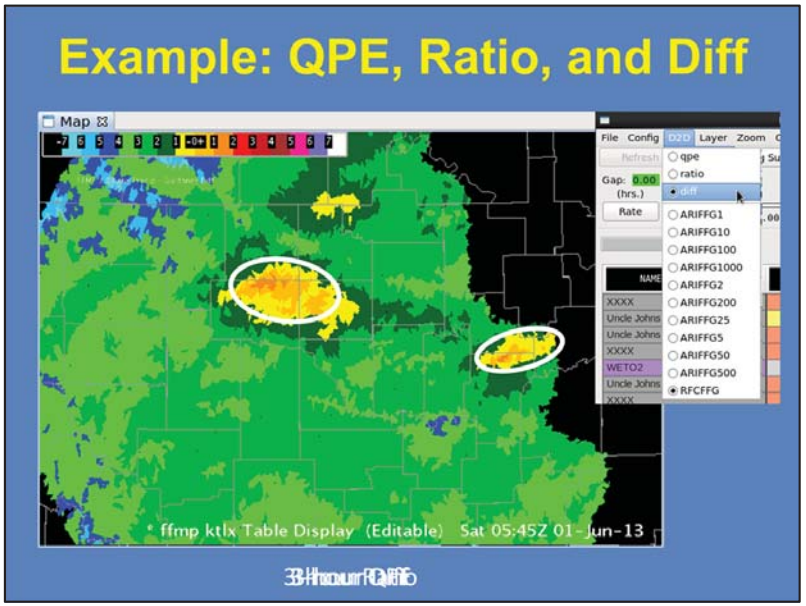
We have a hypothetical basin that has basin-averaged rainfall of 4 inches in one hour, and the flash flood guidance is 2 inches in one hour. Thus, the Diff value would be 2 inches while the ratio would be 200%. Now imagine that for a different rainfall event, the same basin receives 1 inch of rain in an hour, and the FFG is only 0.50 inches.

The Difference is now 0.5 inches, and the ratio is 200%. This Ratio value could lead you to believe a significant flash flood was possible, as in the first case. However, comparing the two Difference values, the 1st event would have much more significant flash flooding given that FFG was exceeded by 2 inches, rather than only 0.50 inches during the second event.

Ratio can be used as a quick awareness tool for basins that are close to or already exceeded flash flood guidance. While, the Diff values give information on the potential magnitude of the flash flooding.

After identifying your areas of precip using QPE, we recommend that you start with viewing the Ratio, in order to pinpoint areas that may be approaching or exceeding FFG. Then, switch over to the Difference to tell how much you are over or under in those areas.





The last menu to discuss is the D2D menu, which determines what is being displayed in D2D. The three options are QPE, Ratio, and Diff, with the default being QPE.

Based on the best practice from the previous slide, it is usually good to start with QPE to get a feel for the high precip areas, and then move to Ratio and Diff to analyze the flash flood threat. When you switch the D2D menu option from QPE to Ratio or Diff, FFMP wants to determine what guidance source is being used for the ratio calculation. The default is RFC FFG, as shown here.

However, as mentioned in Part 1, Average Recurrence Intervals (or ARIs) are a new guidance source option in FFMP. To force the D2D display to show ratio (or diff) calculated against ARI, simply choose one of the various ARI options. Keep in mind this change ONLY affects the display, and not the table values. Therefore, it can become confusing if you set the *display* to use ARIs to calculate ratio and diff, but your *table* uses FFG. So do NOT forget to always switch your D2D display back to FFG once you are done with the ARIs.

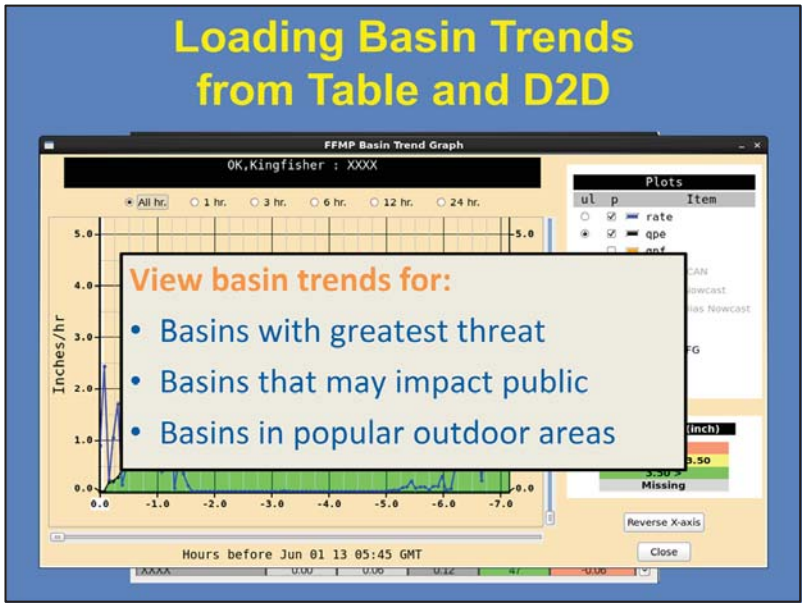
To reiterate from Part 1, there is a lot to learn about ARIs before you start using them in AWIPS. For novices, we recommend you do not use ARIs in FFMP, and therefore do not alter the D2D menu guidance source. Rather, wait until you take WOC Flash Flood to learn more.

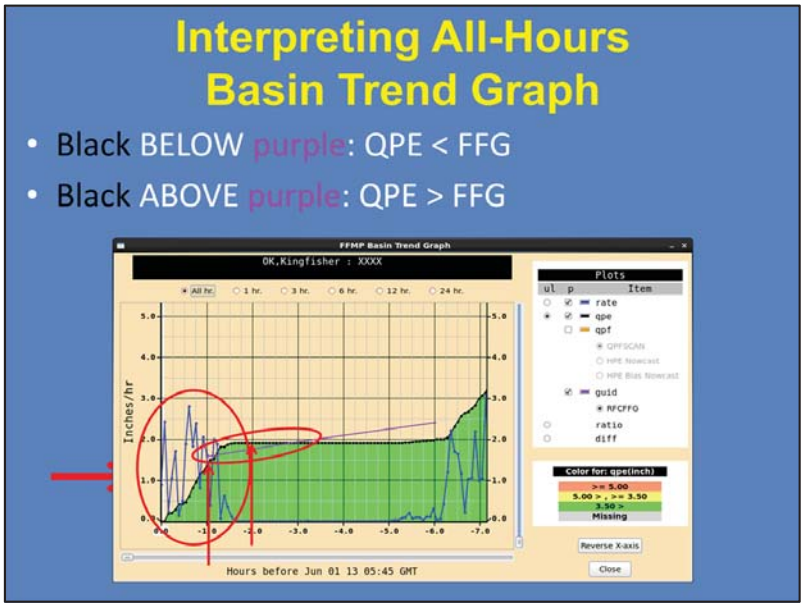
Here is an example of how to interpret the D2D options.

So FFMP defaults to displaying QPE. Simply looking at this output, we see there is a large area of greater than 1 inch in three hours, with isolated areas having upwards of 3 inches of rain in three hours. This information is useful for situational awareness, however it does not tell us anything about the hydrological response.

Therefore, your next move is to look at the Ratio product. Remember, for this, we are interested in areas that are approaching or have exceeded 100%. Begin by seeking out the purple areas, or those that have exceeded FFG. If rain is continuing, then also consider the dark reds and pinks, since those areas are approaching exceeding FFG. With this methodology, we have narrowed our threats to the circled regions.

Finally, use the Difference display to see by how much FFG has been exceeded. Within our areas of interest, light orange is 0-1 inches, and darker orange is 1-2 inches, which is where the more significant flash flooding threat is located.





By this point you have zeroed in on the primary threat areas using Ratio and Diff and by monitoring rain rates. Though basin trend graphs in FFMP take some time to get used to, they contain critical information on the timing and relationship between the QPE and guidance for different durations.

To load a Basin Trend, there are two options: First, you can load it by right-clicking on a small basin name from the basin table. Second, you can set the Click menu option to “Basin Trend”, then go to the D2D pane with FFMP and make the display editable by middle clicking on the text in the legend, and then right-click on any basin in the display to load a basin trend for that basin.

Because there could be tens of thousands of small basins in your localization and it would be time intensive to view large numbers of basin trends, it is best to view the basin trends for basins that: 1) have the greatest current or projected threat; 2) perhaps those basins that might significantly impact the general public (like urban basins); or 3) basins in a National Park that normally contain numerous hikers and campers.

The basin trend graph gives you an easy way to view history of rainfall rates and accumulations as they relate to flash flood guidance. Now let’s look at how to interpret one.

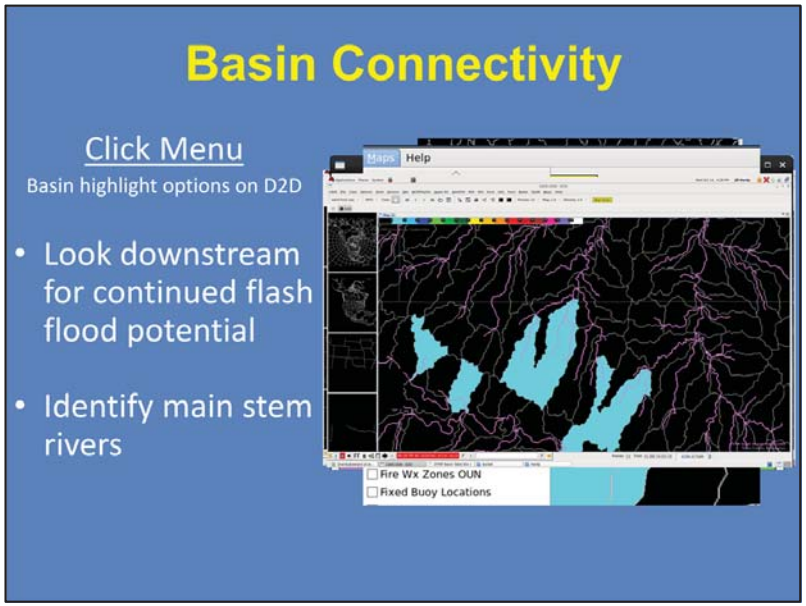
Okay, let’s take a look at this basin, which is only three-hundredths of an inch away from exceeding the 3-hour FFG and is currently experiencing instantaneous rain rates of 0.89 in/hr. As a warning forecaster, I would like to know when within the three hours 1.90 inches of QPE fell, so I right-click on the basin to load a basin trend graph shown here.

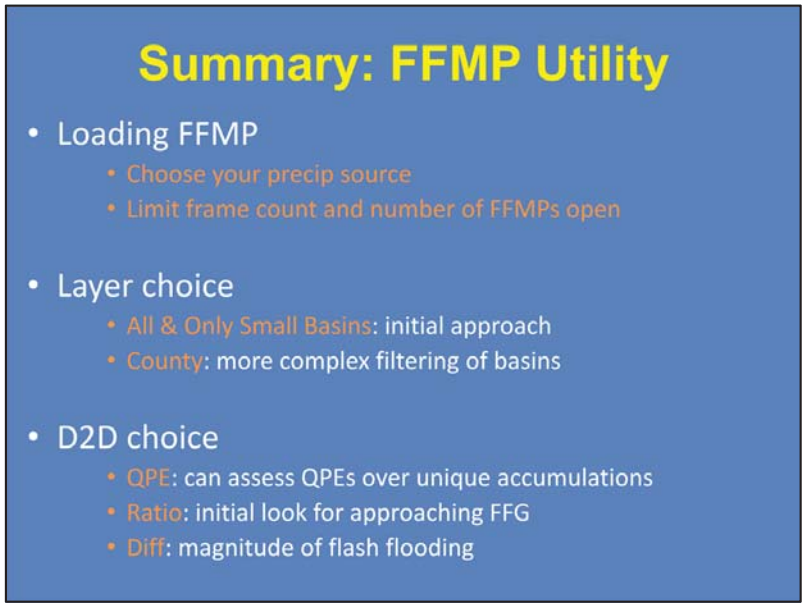
First, you want to look at the blue line, which is the instantaneous rate trend. Each blue dot represents the instantaneous rate for a particular volume scan. From this we see that rates of ~ 2in/hr occurred primarily over the last 1.5 hrs, and there was no precip 3 hrs ago and 4 hrs ago.

Next, the black line is the precip accumulation for different durations. You will notice the accumulations increase every time there is an instantaneous precipitation rate > 0. The instantaneous rate is multiplied by the volume scan time step in order to increase the accumulation. To interpret this line, we see about 1.3” has accumulated over the 1-hourr duration, while 1.9” have accumulated over the most recent 2-hour duration. We see the 1.9” accumulation lasts through the 5-hour duration, because there was no precip falling between 2 and 5 hours ago.

Finally, there is the purple line, which shows FFG for the 1-, 3-, and 6-hour durations. Whenever the black QPE line is BELOW the purple FFG line, QPE is less than FFG for that duration interval. When the black line is ABOVE the purple line, QPE is greater than FFG. Here, QPE is always below FFG, except for durations between 1 and 3 hours where FFG is exceeded by about 0.25” for the 1.5-hour duration. This may be enough to cause flash flooding, particularly since the instantaneous rates are continuing at the current time, and the longer duration FFG values (like 3- and 6-hour) are going to be exceeded more and more as that continues.

Now we’re going to take a few minutes and let you have some practice with basin trend graphs. The following quiz is NOT graded.





Finally, FFMP allows you to see basin connectivity on the D2D display as configured in the "Click" menu. For instance if you select "Downstream" from this menu, then you go to the D2D display and make the FFMP display editable. Once editable, your right-click button will highlight the basins downstream of the basin you selected. If it is hard to see the highlighting, you can change the color of the trace, like I did to green. And whenever you want to get rid of your trace or change the type of trace, simply "Clear Trace" on the FFMP Table.

Once you have pinpointed your current threat area, it is important to look downstream to see where the runoff will go. If upstream accumulation is great enough and the downstream basin is flashy, those downstream basins can have flash flooding even without receiving a drop of rain. Knowing this can help you adjust your warning polygons to account for the future threat.

Additionally, major main stem rivers typically don't *flash* flood, so using the upstream and downstream connectivity to help identify them and can be useful in adjusting your polygon.

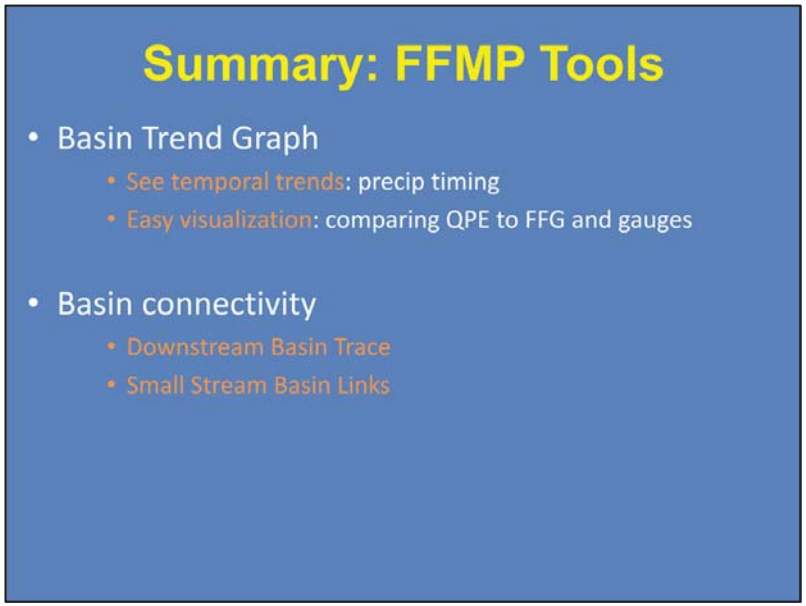
A useful tool for visualizing downstream flow is by overlaying the "FFMP Small Stream Basin Links" from the Map menu in D2D. I made them purple in this graphic. As you can see, all of the basins with any response are in the most upstream headwaters, with no other basins feeding into them. These headwater basins tend to be the smallest, but with the quickest responses to rain, and thus, the highest risk for rapid flash flooding.

Because FFMP has a number of unique strengths, we focused this lesson on using it to its fullest in flash flood warning operations.

First, when loading FFMP, make sure to follow the guidance in Part 1 to consider all of your available precip sources. Also, be diligent with frame count, and only have one FFMP open at a time.

Next, start using the All & Only Small Basins layer to identify areas where QPE is approaching or exceeding FFG. You may change to County layer when you need more complex filtering of basins in the table so you can find particular basins or the virtual gauge basins that can be used to compare METAR/mesonet to precipitation accumulations.

Within FFMP, D2D can be configured to show QPE in order to uniquely be able to assess things like HPE, Bias HPE, and MRMS accumulation durations that aren't usually readily available. The ratio product is one the best ways to identify areas of flash flooding threat so we recommend starting there, and using Diff to help assess the potential magnitude of flash flooding.



FFMP has a lot of useful functionality, as long as you know how to use it.

The Basin Trend Graph allows you to see temporal trends for rainfall rate, accumulation, and Flash Flood Guidance for a selected basin. As well as, provides easy visualization when comparing QPE to FFG, and to gauges when using VGBs.

FFMP also has basin connectivity features to help visualize where the flash flood threat may be evolving, including the Downstream Basin Trace that highlights basins downstream of the selected basin, as well as the Small Stream Basin Links feature for identifying main stem rivers.

This concludes Part 2 of the Flash Flood Decision-Making modules. When you are ready, please move onto the next slide to take the quiz and receive credit on the LMS.

TAB

Radar & Applications Course

Topic:

Storm-Based Warning Fundamentals

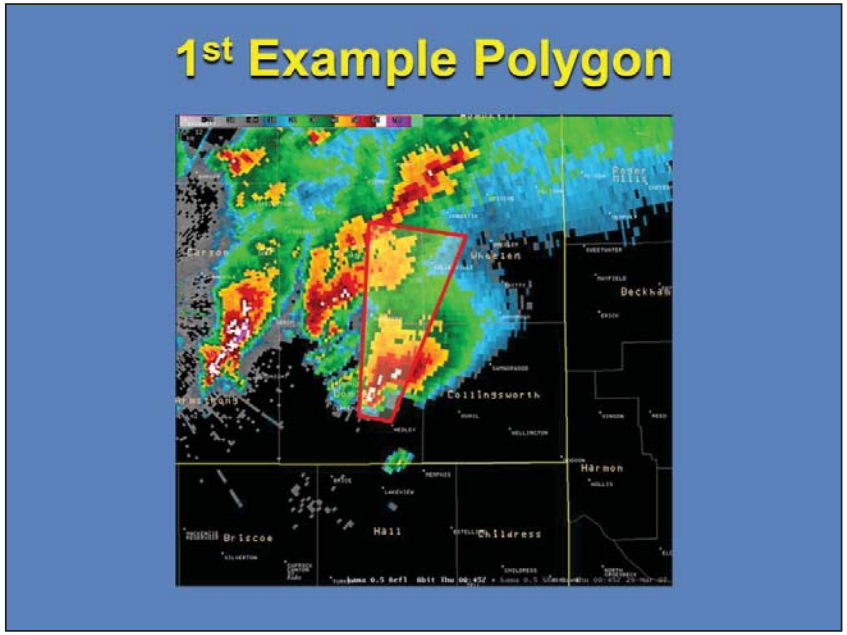
Table of Contents
Topic: Storm-Based Warning Fundamentals

Click to jump to lesson

Lesson 1	**TOR for Isolated Tornado Threat**
Lesson 2	**TOR for QLCS Tornado Threat**
Lesson 3	**TOR/SVR for Training Storms**
Lesson 4	**SVR for Pulse Storms (Low Shear)**
Lesson 5	**SVR for Squall Line Systems**
Lesson 6	**Special Considerations**
Lesson 7	**Two TORs in Close Proximity**
Lesson 8	**Non-Linear Motion**
Lesson 9	**SBWs with Merging Storms**
Lesson 10	**Limiting the Number of Counties in Warnings**
Lesson 11	**Pathcasts in Storm-Based Warnings**
Lesson 12	**Flash Flood Warning Fundamentals**
Lesson 13	**Impact Based Warnings: Overview**
Lesson 14	**Impact Based Warnings: Validation and Application**

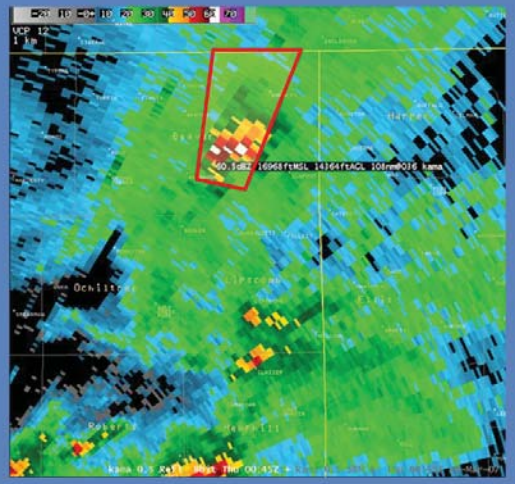


Welcome to the Storm-Based Warning Fundamentals lesson entitled “TOR for Isolated Tornado Threat.” The lesson is narrated by former instructor, Paul Schlatter. Since this module was part of the original storm-based warning course from 2011, there are occasional references to objective “T1” . Please disregard those verbal references .The purpose of this lesson is to show examples of possible ways to issue tornado warnings for individual and potentially tornadic supercells.



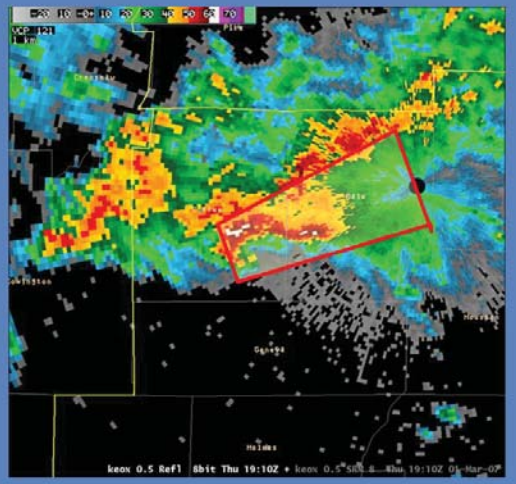
This first example shows several supercells east of the Amarillo, Texas radar. Strongest reflectivity returns from the hail cores are moving from 200 degrees at 35 knots, as are the low-level circulations give or take a few degrees and knots. The furthest east supercell in Donley County has already produced significant tornadoes and shows no signs of weakening or even becoming less tornadic. A storm-based tornado warning for the next 40 minutes for this supercell may look like this, covering the area I feel is most at risk to experience a tornado based on near-storm environment and an analysis of all-tilts SRM. It’s fairly narrow because of the steady-state behavior of the circulation and storm over the last couple of hours. Toggling over to reflectivity, the tornado warning includes just about all of the reflectivity and hail core to the north and west of the inflow notch but it may be a good idea to make the polygon a little wider to include the potential for hail since drawing a tiny severe thunderstorm warning on the fringes of the tornado warning is impractical and would be confusing to our customers for this storm. Thus, it would be a good idea to include expected hail size in the text part of the warning. Additionally, the south end of the polygon is south of the circulation to account for the potential for severe rear flank downdraft winds. It is a good idea to keep a buffer in the polygon to account for the potential of RFD winds.

Polygon Far from Radar

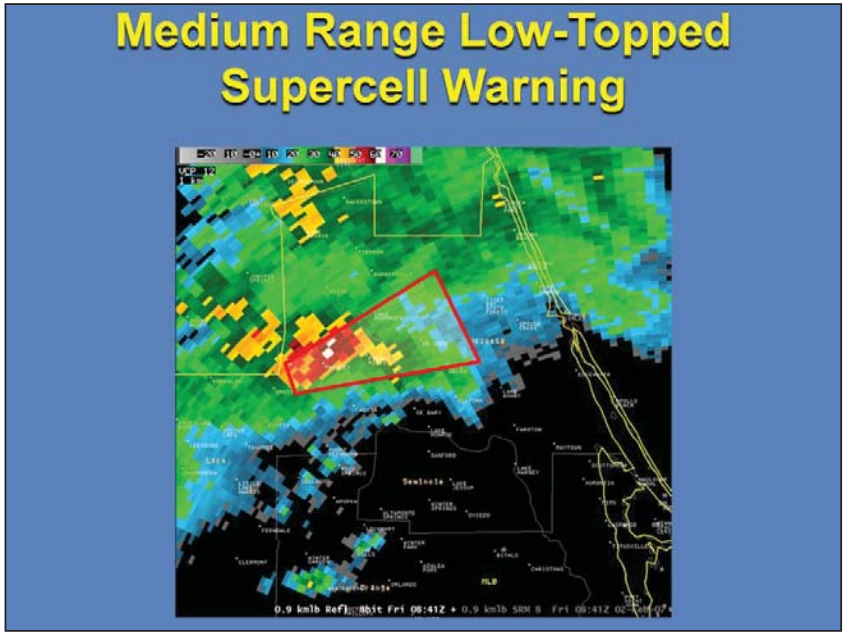


Try this example far from the radar. A tornado warning polygon for this storm 100 nm from KAMA, which is also moving NNE at 35 knots, may look like this for a 30 minute warning. The polygon extends to the CWA border on the north and includes the impressive gate to gate signature. (Z toggle) This storm produced a significant tornado within this polygon and also includes all significant reflectivity associated with the hail core.

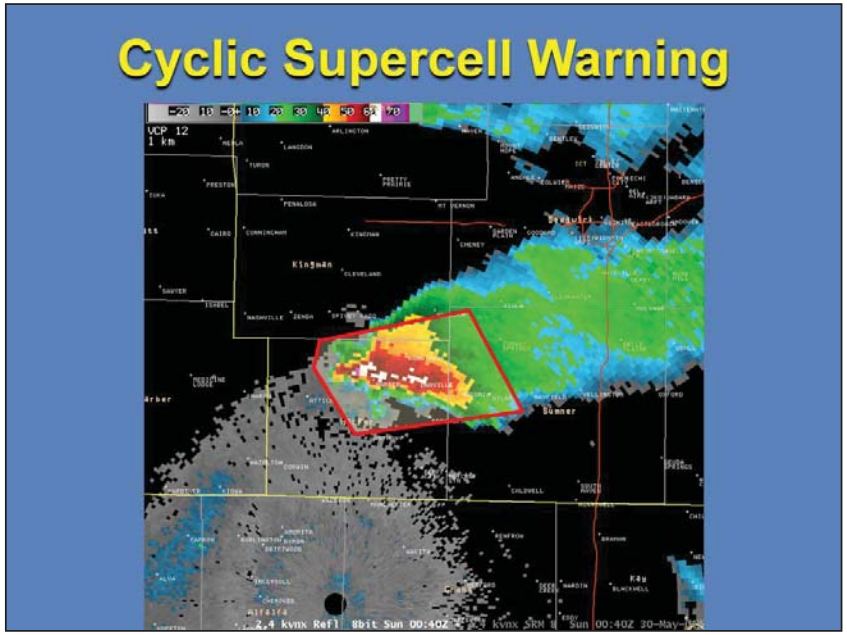
Polygon Close to Radar



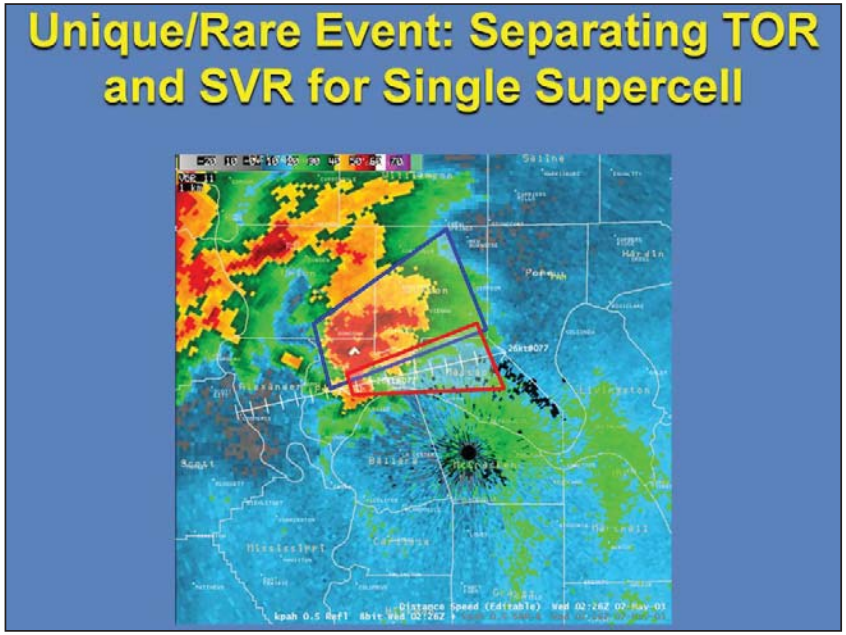
Another example, this time very close to the radar. We are looking at southeast Alabama and the KEOX radar, for the northeast part of the Tallahassee CWA. The low-level circulation is moving from about 230 degrees at 43 knots. Toggling over the reflectivity, we can see that the low-level mesocyclone is moving into an area of 60 dBZ to the east-northeast. There are also 60 dBZ returns just north of the circulation. Thus, a single storm based tornado warning for this event may look like this, for the next 30 minutes. This is the Enterprise, AL tornado so you know how strong and persistent the tornado from this point forward. Let me toggle back to SRM just to give you an idea of the amount of buffer included around and ahead of the circulation, given the uncertainty of the tornado track over the next 30 minutes.



The next case is at a medium range from radar, about 70 nm, this is from KMLB and the storm is over eastern Lake County, Florida. I am showing the 0.9 degree tilt since the 0.5 tilt is contaminated by being right at the end of the first trip. This was a low-CAPE environment but contained off the wind charts shear. This supercell had been tornadic for over a half an hour by this time. A storm-based tornado warning may look something like this. I tracked the mesocyclone for storm motion, and it is moving out of the WSW at 46 knots. This is a 30 minute warning, taking the eastern edge of the polygon into central Volusia County. Toggling over to reflectivity…nearly all strong reflectivity is expected to remain within the tornado polygon.



Let's try a hindsight example to illustrate the need to NOT make warnings too narrow in the new storm based world. I say hindsight because I know the location and movement of the tornadoes from this radar image forward. Let's look at a cyclic supercell in Southern Kansas. Often times, especially with cyclic supercells, older, occluded mesocyclones with or without tornadoes move well to the left of the parent storm track. In this image, the older mesocyclone is moving NNE as indicated by the white arrow, and in fact produced a tornado that persisted for over 15 minutes following this volume scan. The new mesocyclone, clearly dominant based on an all-tilts analysis, is shown here. The new mesocyclone track is indicated with the white arrow, and it went on to produce a strong, long track tornado from here into Sumner County. A narrow, storm-based tornado warning that would merely cover the southern flank of the storm and hook echo would not cut it with this storm, as a significant tornado would be missed, and possibly subsequent tornadoes. Here is the polygon I have drawn. It is pretty wide in the north/south direction to account for the older mesocyclone and the possibility that during this 30-min warning the new meso may also occlude and moves left of the storm track. With cyclic supercells, it's a good idea to make the tornado polygon account for the possibility of tornadoes to the left of the storm track. Toggling to Z, we find that the polygon contains all of the RF gustfront and hail core so those threats should be covered in the text portion of the product. If you are unsure if a particular supercell is cyclic, it is probably better to play it safe and keep a nice buffer to the right and left of the rotational track for your tornado warning polygon.



Here is a final example of a very rare storm, with which you could potentially split up the severe thunderstorm and tornado warnings. Recall from previous examples that given the high uncertainty of tornado location and movement, distance from radar, and/or the location and track of the tornado relative to hail core, issuing separate tornado and severe thunderstorm warnings wouldn't be a viable option. However, there are in fact a few storms where this may be done effectively. Here is one such storm from Paducah, Kentucky, with the track of the TVS indicated with the distance speed tool.

A few things are unique to this and other rarely seen tornadic supercells like it:

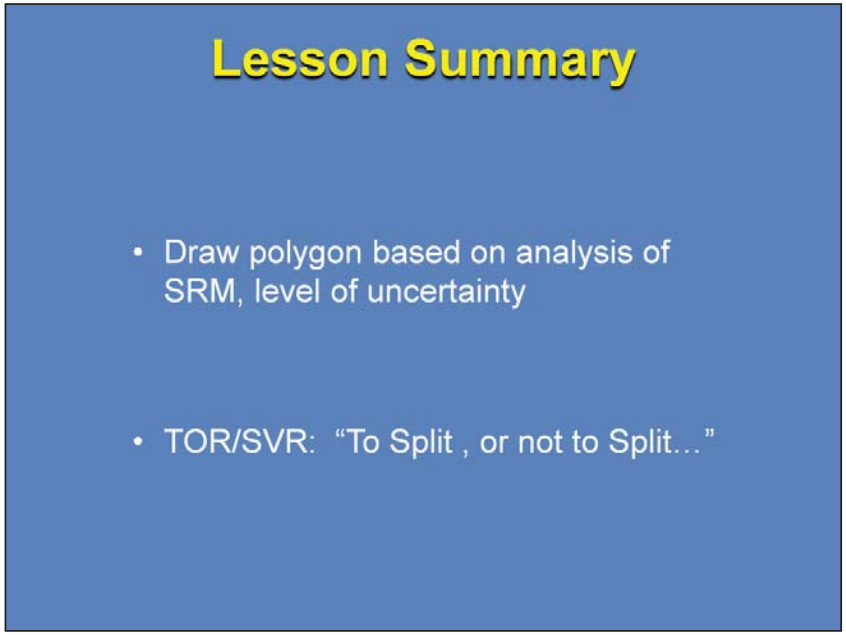
•Close to the radar

•Nearly Steady state, non-cyclic supercell

•Classic reflectivity structure with sufficient separation between hail core and tornado

•Tornado location and track not overlapping with hail core track

With these in mind, I first drew the tornado warning polygon based on all-tilts SRM for 30 minutes. There actually isn't much uncertainty with the location and movement of this tornado because the supercell has been nearly steady state and the TVS movement hasn't varied more than 5-10 degrees over it's already significant lifetime. Toggling over to reflectivity, clearly the hail core is not within the tornado polygon and is expected to track along and to the north of the northern edge of the polygon. Thus, a severe thunderstorm warning for hail up to the size of golfballs could be issued like this, based on storm history and an analysis of all tilts base reflectivity. This storm did produce severe hail and a long track F4 tornado over the next hour, cutting a swatch across Pulaski, Massac and Pope Counties. A bit of overlap is essential between the severe thunderstorm and tornado warnings.



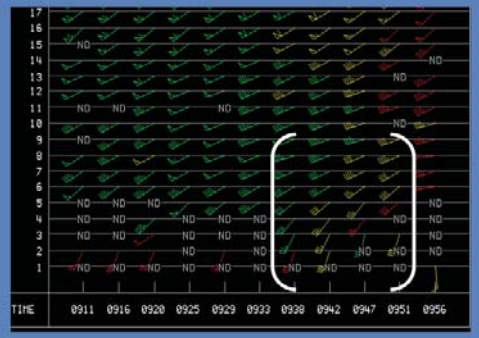


Draw your storm-based tornado warning polygon for a supercell based on a thorough temporal and spatial analysis of all-tilts SRM, incorporating your uncertainty about the location and movement of the potential tornado. Be sure to include wind or hail threats in the text part of the product, as time permits. There are rare but high-end storm events where separating severe thunderstorm and tornado warnings is a viable option. A classic, steady state tornadic supercell close to the radar is a candidate for splitting up the warnings. The vast majority of tornadic supercells however would be well handled by a single storm-based tornado warning with a buffer around the low-level circulation to account for RFD winds to the right of the track, and hail or deviant tornado motions to the left of the circulation track. This concludes the presentation for objective T1, thanks for listening.

Welcome to the Storm-Based Warning Fundamentals lesson on Tornado Warnings for Quasi-Linear Convective System (QLCS) Tornado Threat. This lesson was used as a previous Objective T3 for the original Storm–Based Warning Course.

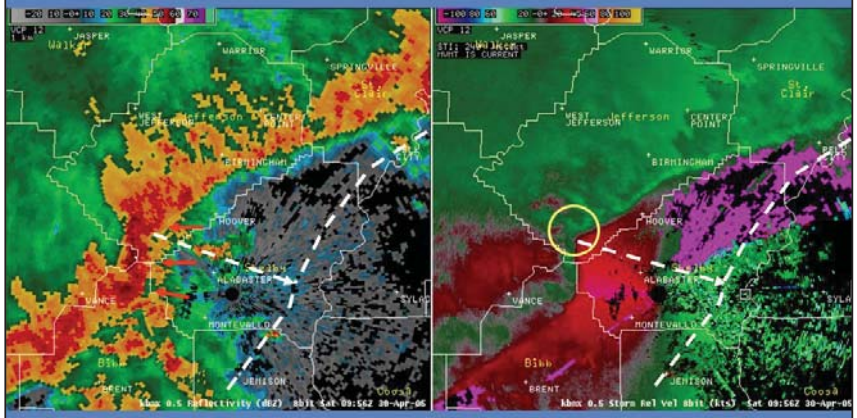
Threat Assessment

- VWP indicates strong directional and speed shear
- Previously warned (SVR) squall line moving east at 30 mph

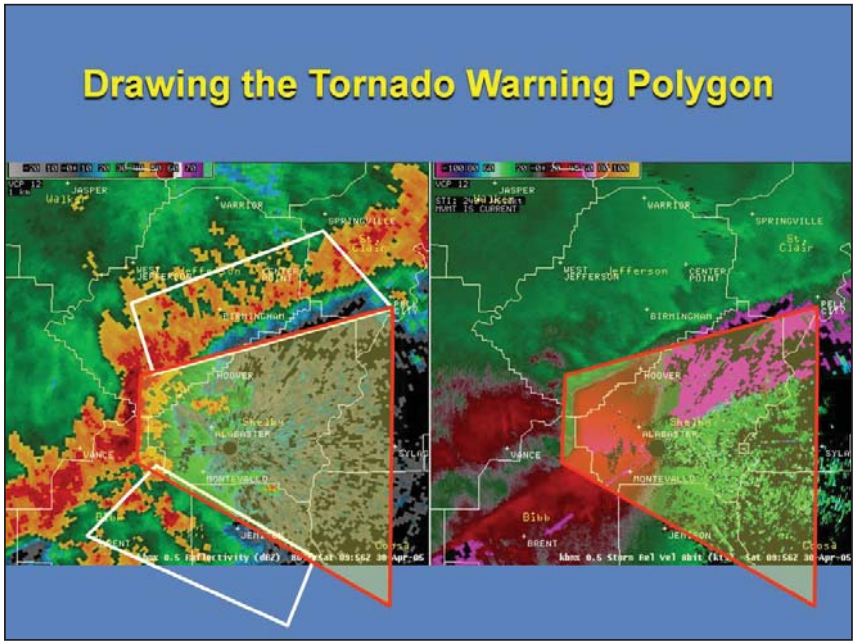


The environment for this event supported tornadic potential due to very strong directional and speed shear as you can see from the BMX VAD wind Profile (VWP). A squall line, which had a history of producing damaging winds and hail was moving eastward at around 30 miles per hour through central AL.

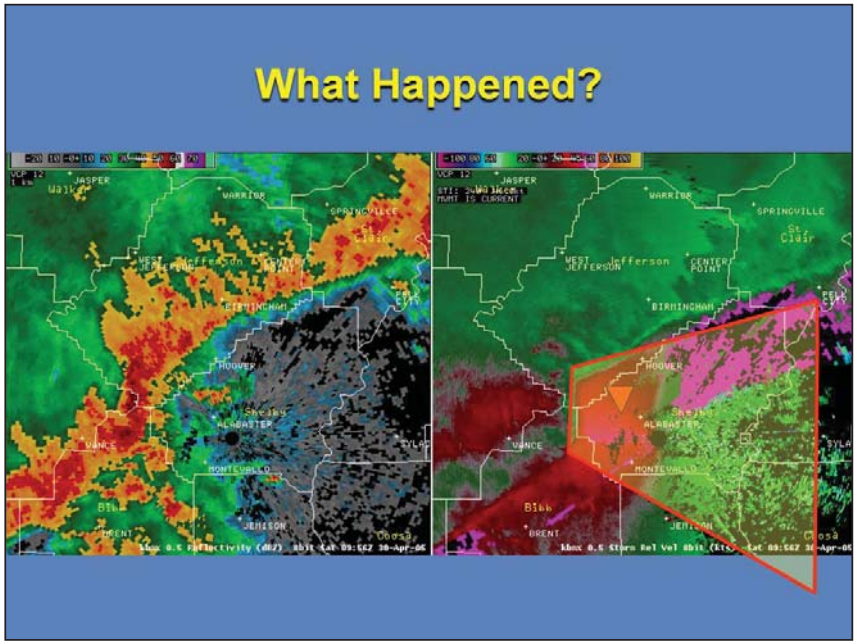
Threat Assessment



Along this squall line, you can see that there is bowing segment starting to develop just to the west of the radar. Also, inflow notches are becoming better defined along and just to the south of a kink in the squall line. The kink is a location for possible rapid, low-level mesocyclone development, which could lead to a squall line tornado. The location mentioned is in southern Jefferson County denoted by the yellow circle. So, how do we lay out the storm threat given the location of the potential tornado threat and additional severe thunderstorm threat (mainly high winds and hail) along the southern portion of the line? The 30 min projected squall line motion is denoted by the dashed white line.



The red polygon is a potential tornado warning for 30 minutes for the squall line tornado threat area. You could also lay out separate severe thunderstorm warnings for adjacent line segments and storms developing along and ahead of the squall line. These severe areas are shown in white (note the slight overlap). For follow-up statements, as the line moves eastward, you would want to trim the back edges of the warnings but be mindful that new development could occur anywhere along the line.



This storm was one of 9 tornadoes that occurred in Birmingham's CWA on April 30, 2005. This particular tornadic storm that affected the Helena, Alabaster, Pelham, and Chelsea areas was rated an F1 with winds estimated around 75 miles an hour. The tornado first touched down just west of County Road 93 near the Cahaba Wildlife Management Area in Helena. The tornado moved generally eastward and crossed County Road 17, County Road 58, US 31, Interstate 65 and County Road 11 before ending near County Road 39. The tornado damage path was approximately 11.6 miles long and 100 yards wide at its widest point. The tornado was on the ground from approximately 504 AM CDT to 518 AM CDT.

Summary

- Recognize heightened tornadic threat in squall lines by monitoring low-level shear
- Use lowest tilt radar indications to help identify and track potential tornadic storm threat areas in squall lines
  - Inflow notches
  - line intersections, kinks
  - Velocity couplets
- Lay out separate severe thunderstorm warnings for adjacent line segments and storms developing along and ahead of the squall line

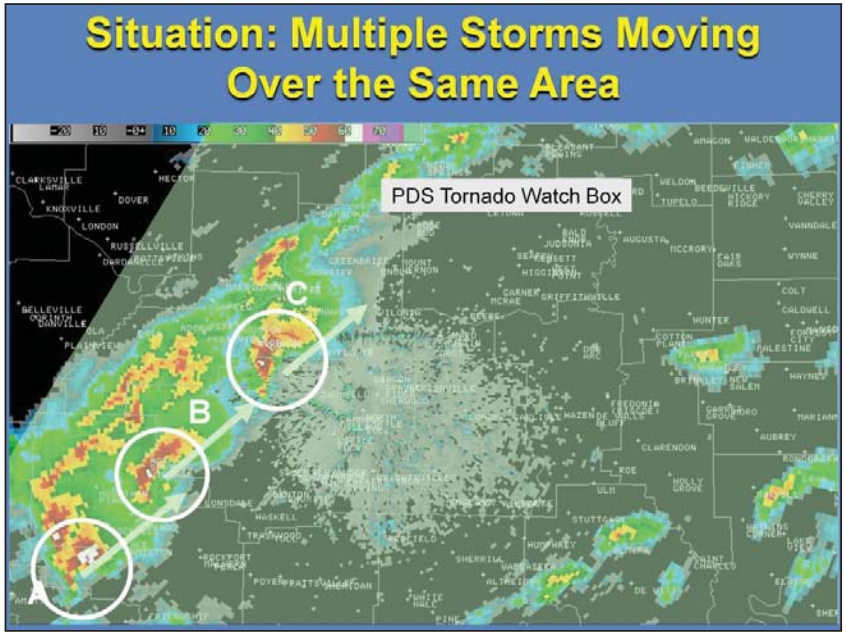
For squall line tornado SBWs, here are a couple of points to remember:

First, recognize tornadic threat in squall lines by evaluating environmental low-level shear, like from VAD wind profiler data. Use lowest tilt radar indications (Base data works best) to help identify and track potential tornadic storm threat areas in squall lines such as: inflow notches, line intersections, velocity couplets at kinks in line.

Lay out separate severe thunderstorm warnings (or tornado warnings) for adjacent line segments and storms developing along and ahead of the squall line.

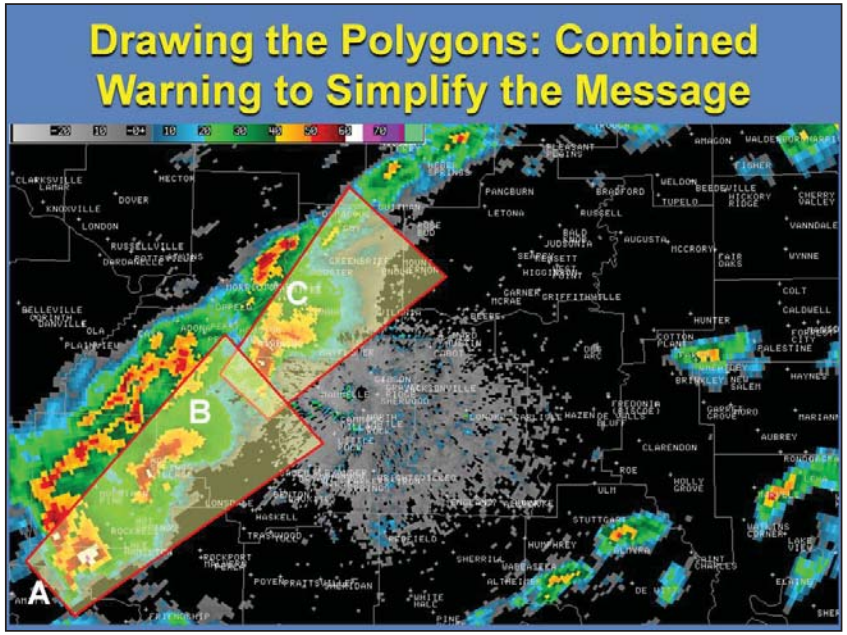


This brief lesson will illustrate some of the issues involved with drawing storm-based warning polygons for storms training over the same region. The lesson will consist of an example of several storms whose individual motions will bring them over the same areas one after another. Additionally, there will be a summary slide with the key points to walk away with from this example.

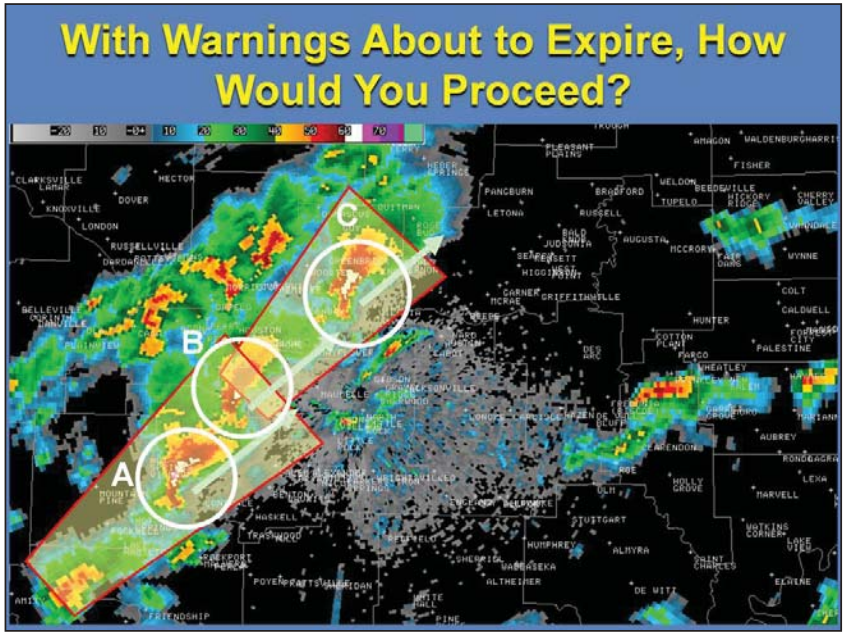


The next few slides will present a situation with potentially tornadic storms moving over the same areas. [A, B, & C labels appear] The discussion will focus on two storms, labeled Storm A and Storm B. Additionally, some mention will be made of a third storm, Storm C. [watch box graphic appears] All three storms are located in a High Risk area that is covered by a PDS Tornado Watch. [watch box graphic fades] …pause…

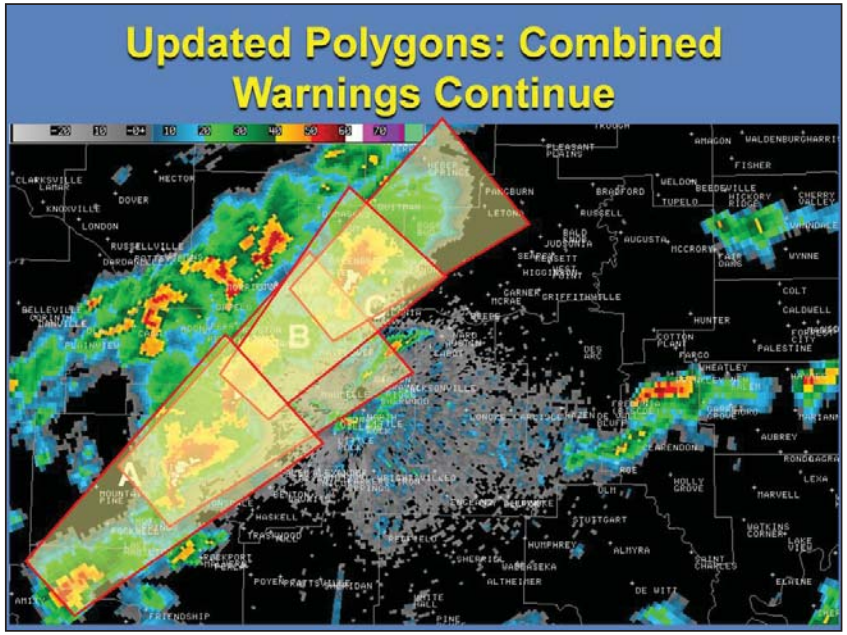
[circles and storm motion vectors for storms appear] In this case, all three storms are moving along in nearly the same direction with Storm A moving slightly faster than Storms B and C (green arrows indicate storm motion vector). All three storms have storm structure on radar indicative of supercells with the potential to produce tornadoes in the near future. While the near-storm environment is more supportive of tornadogenesis further to the north, the warning forecaster thinks that it would be prudent to issue a TOR for all three storms. [circles move forward to show future storm location] Looking forward in time, we can see approximately where the storms will be in the next 30 minutes or so. These future positions just highlight how difficult drawing polygons can be for this case.



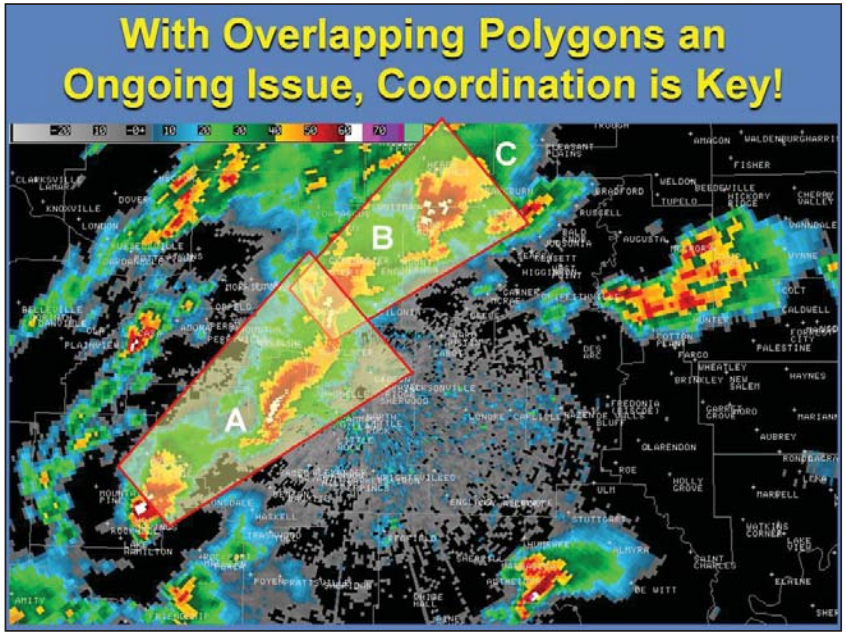
The primary warning issue is how to draw your polygons so that they minimize confusion from overlap. This example could produce several different polygon possibilities. [warning polygon appears] The solution presented here, with Storms A and B encompassed by one warning and Storm C covered by a separate warning, is a good one because it minimizes overlap. Additionally, this product combination would allow for some different wording for the different areas as Storms A and B are impacting more recreational and rural areas while the Storm C is impacting the suburbs of a large metro area.



Moving ahead in time about a half-hour, it's time to reissue our warnings. [warning polygons fade] During that time, the storms have continued to intensify with reports of tornadoes with Storm C. [circles & arrows for storms appear] The storm motion for these three storms continues to be roughly the same as it was earlier. While the decision to continue having all three storms covered by TORs is straight-forward, do you continue to use the same polygon strategy for these three storms? [circles move forward in time] Looking ahead in time, you can see the storms appear to being moving along the same paths still.



…pause… [fade warning polygons to dashed lines] In this example, the subsequent warning polygons that were drawn are similar to the original ones [new warning polygons appear] (with Storms A and B in one warning) as the previous warnings. That will not always be the case, however. If storm motions had changed significantly enough between the two storms, issuing separate polygons may have been prudent. This change in warning issuance would also have been necessary if the threat for either storm changed significantly (say if, in this case, only a SVR was warranted for one of these storms).



And here are Storms A, B, and C 30 minutes later. It appears in this image that these storms will continue to be a significant threat. Additionally, there is another storm further to the southwest along this same path to worry about. This is obviously an extreme example. However, one can see from such an example why intra-office coordination between warning forecasters is so important for these cases.

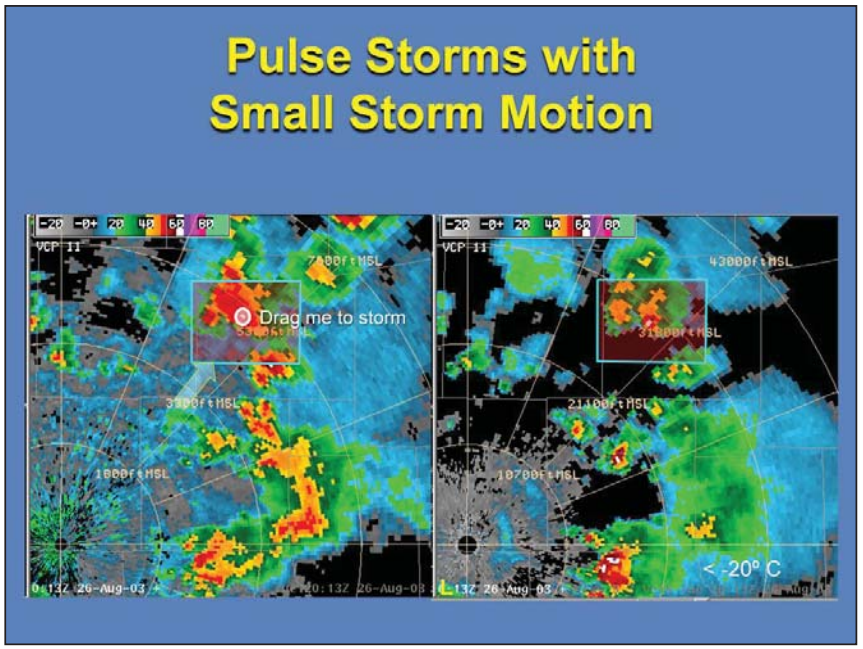
Summary

- Minimize confusion by minimizing overlap
- It's OK to include multiple storms in close proximity with similar threat potential in a single warning
- When combining storms, remember to re-evaluate threat potential with each warning issuance

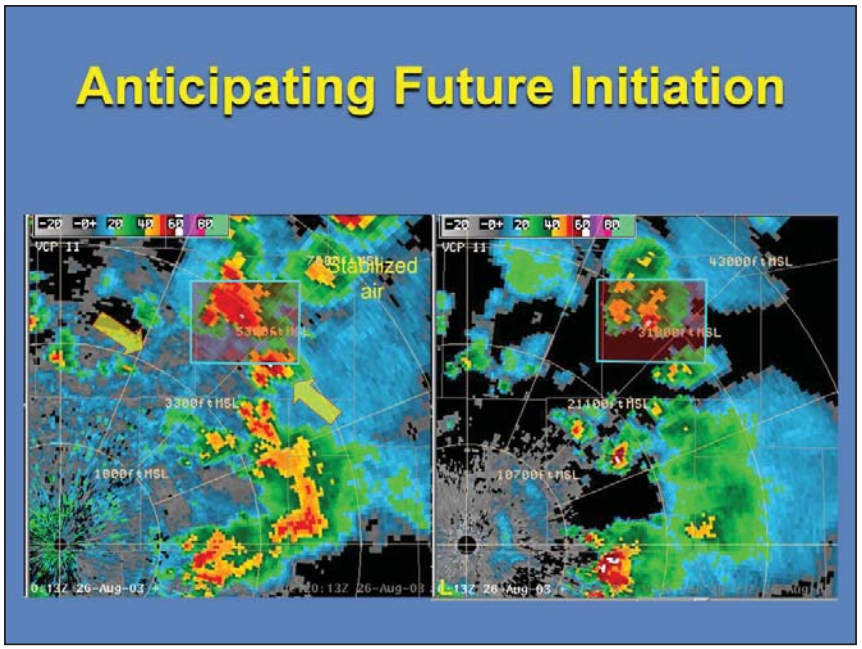
In summary, here are the key points from this example. [first bullet appears] With several storms training over the same area, it will often be confusing to our customers what warning covers what storm. One of the best ways to minimize confusion for them is to minimize the overlap in our warning polygons. [second bullet appears] Even though we are putting a lot of emphasis on these warnings being "storm-based", they are still threat based as well. When multiple storms are in close proximity to each other with similar threats, it's acceptable (if not sometimes preferable) to include them in a single warning polygon. [third bullet appears] When you do combine storms in a single polygon, you need to remember to continually re-evaluate the threat potential of the storms during each warning cycle.



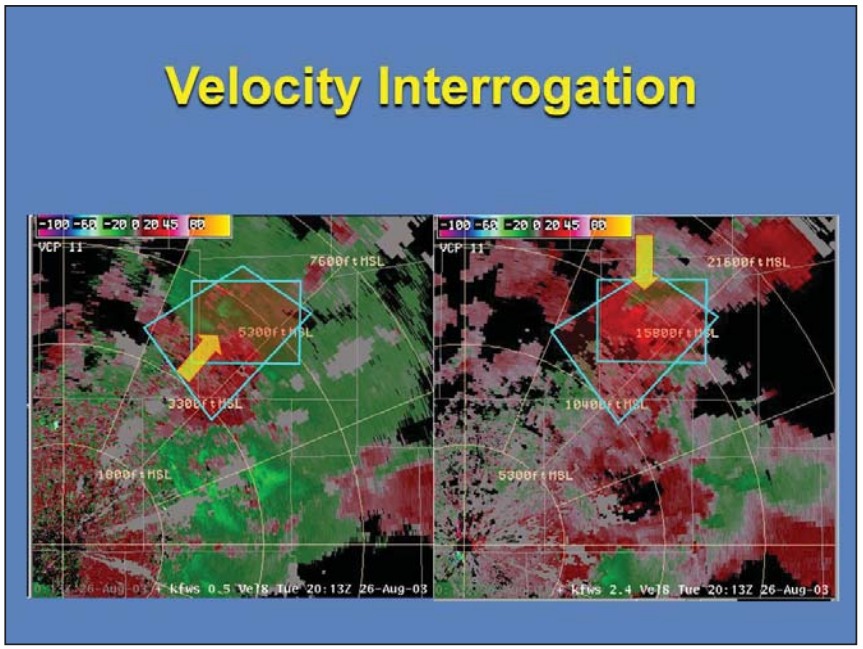
This is a storm-based warning lesson where we discuss what to do for low shear pulse severe storms. I thank Gary Woodall MIC, Phoenix for reviewing this module.



This is a pulse severe thunderstorm environment with little wind shear and weak steering layer flow. The individual cell lifecycle will be dominated by buoyancy process and thus will be relatively short-lived. I have a two panel display, one showing low-level reflectivity (left) and the other showing reflectivity above the -20º C level (right). The storm in question is located northeast of the radar. By the second scan, an updraft pulse develops >60 dBZ reflectivity in the right panel. You may consider issuing a warning at this time assuming that the subsequent downburst expands equally outward in all directions. You draft a warning polygon by taking the 'drag me to storm' icon over the center of the reflectivity core. Adding zero storm motion yields a square box. Your warning's in draft mode pending further consideration. You've got a minute to consider more issues.

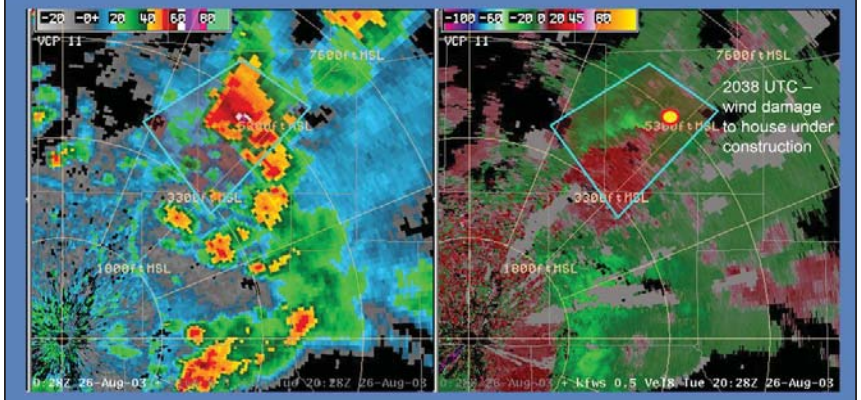


Upon closer inspection, you notice that previous convection is stabilizing the air northeast of your storm. At the same time, there's another cell to the southeast moving to the northwest. There may be a cell merger. Also, you noticed a boundary collision to the southwest of your storm. Boundary collisions would be prime areas for new convective initiation.



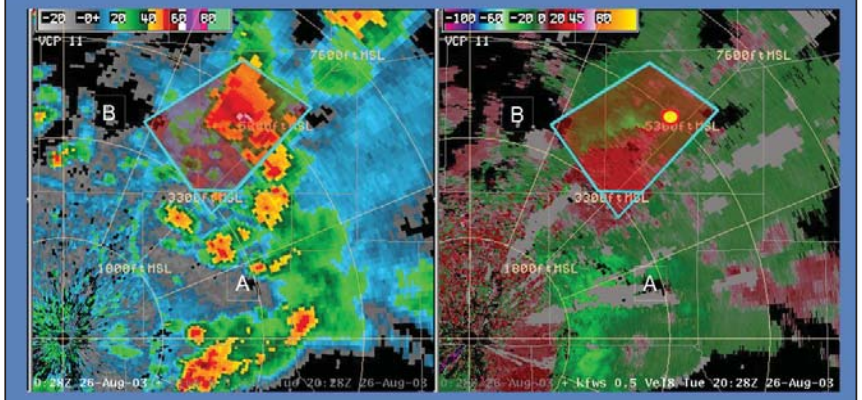
You look at the velocity data and by 2013 UTC you see a MARC signature in the right panel; right above and behind where you see low-level convergence. The MARC starts showing up at the same time as the 60dBZ at 30 kft ARL. There are also indications of a gust front associated with the low-level convergence moving to the southwest. Perhaps it may be necessary to redraw the polygon to reflect the possibility that new cell initiation and subsequent downburst may occur to the southwest and during your warning. You expand the width of the severe thunderstorm polygon on the southwestward end to reflect the greater uncertainty of where that new cell and downburst may occur.

Reshaping the Polygon



The precursor signatures that we saw of a severe microburst in this case included the rapid upward growth of high reflectivities followed by a MARC signature. The downburst did strike ground starting near the 2028 UTC volume scan and spread mainly to the southwest. Damage was reported to a house under construction at 2038 UTC. The northeast part of the warning polygon verified, however the rest of the polygon warning covered the area most likely for subsequent initiation along the colliding boundaries to the southwest of the original storm. A storm can easily undergo initiation and produce another severe downburst within the warning polygon verification time.

Accounting for Impacts of Dissemination



After considering a purely storm-based warning, there are still issues that need consideration with how your dissemination is impacted by the shape of your warning polygon. At times, you may have a purely storm-based warning that may slightly cross a political border. In this case, the southern tip of the warning crossed into another county border (labeled A). In order to prevent NOAA weather radios from tone alerting for every owner of one in county A, would you consider altering the warning polygon to remove the warning from there? If you believe that county A is not sufficiently threatened by the weather, then yes, edit the polygon.

For county B, you may believe the threat remains high enough to have a warning. You could create another polygon for county B but this strategy would diverge away from the spirit of storm-based warnings.

The final shape of this polygon adheres to the storm-based warning concept.

Summary

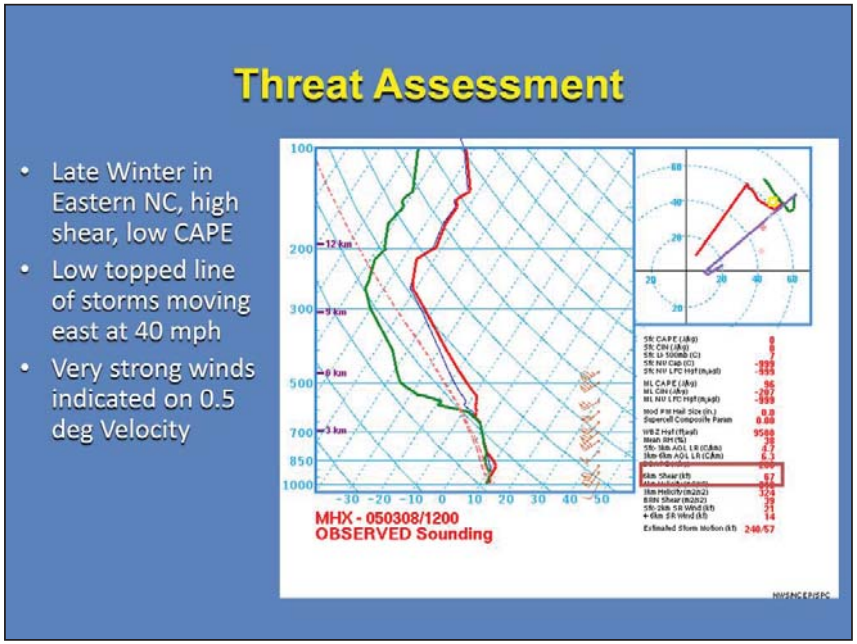
- Center the drag-me-to-storm icon on the reflectivity core of the new storm and then track the individual cell motion to initially shape the polygon.
- The warning polygon should be modified to anticipate new cell development.
- Consider slightly editing your polygon to account for warning dissemination issues that are not storm-based.

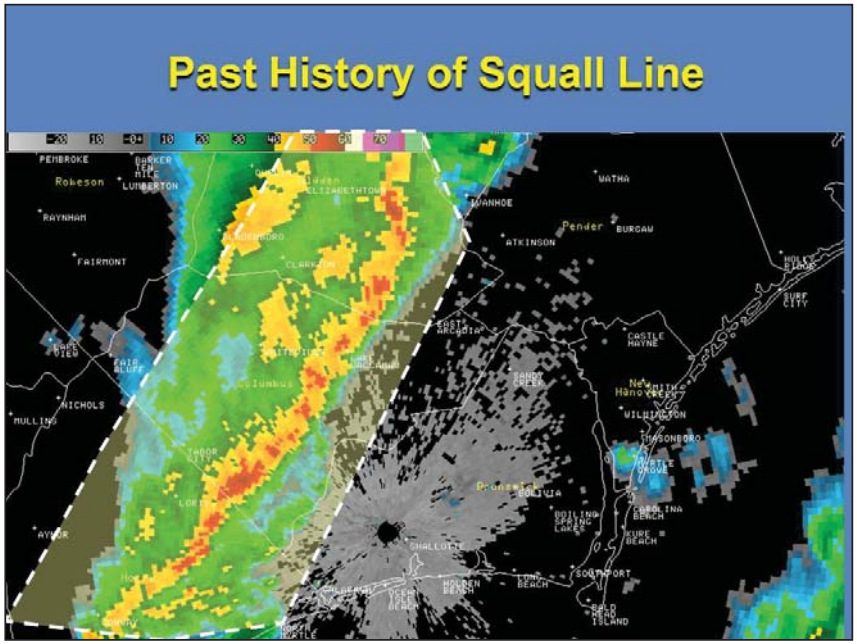
Center the drag-me-to-storm icon on the reflectivity core of the new storm and then track the individual cell motion to initially shape the polygon.

The warning polygon should be modified to anticipate new cell development. Consider slightly editing your polygon to account for warning dissemination issues that are not storm-based.



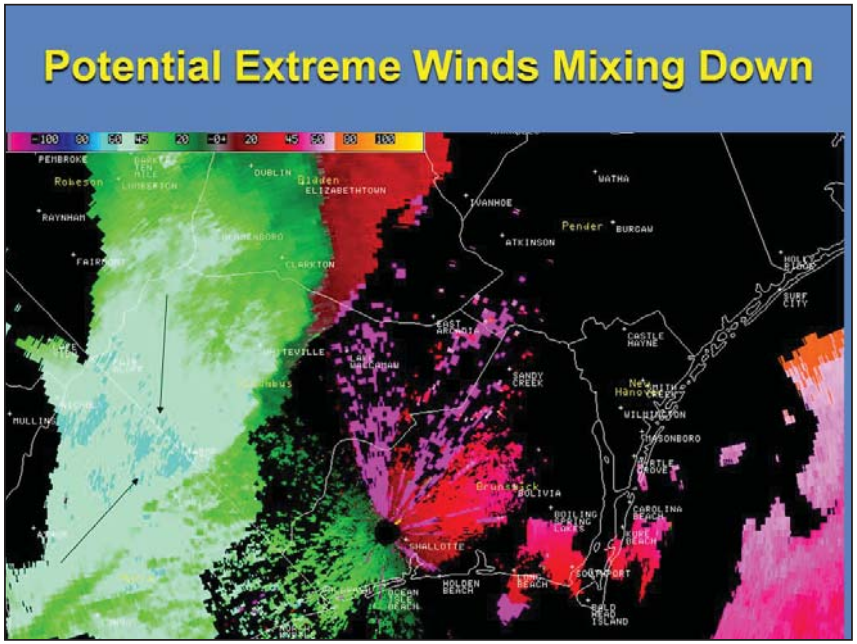
Welcome to the storm-based warning fundamentals lesson on How to Issue Severe thunderstorm Warnings for Squall Line Systems. This is another rather short lesson.



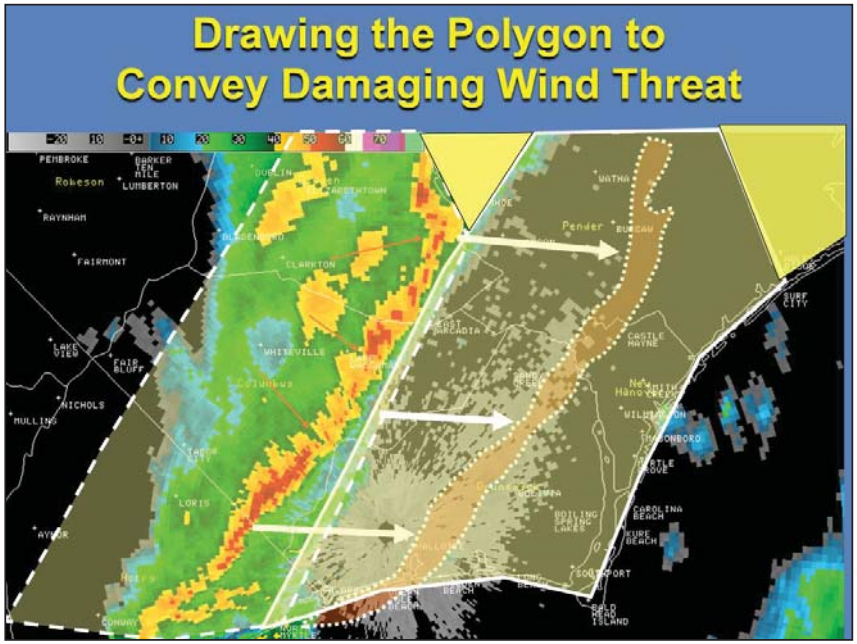


This is a late winter convective warning situation in eastern North Carolina. The environment was characterized by very strong 0-6 km shear (67kts) as seen in the 1200 UTC sounding from MHX. CAPE was very marginal but winds right above the surface were quite strong as indicated on the sounding and the 0.5 degree Velocity product, which you will see on the very next slide. The squall line was moving eastward at around 40 miles per hour.

To get an idea of the movement of the squall line for this storm-based warning example, here (three clicks)is the 0.5 deg Reflectivity imagery for 3 volume scans from the WSR-88D at Wilmington, NC for 1505 UTC to 1514 UTC, March 8, 2005. (click) The current severe thunderstorm warning is shown in dashed white line. Many reports of severe damaging winds had already been reported up to now, but forecasters are concerned for an increase in wind potential as the line moved eastward and approached the coast.



This is the 0.5 deg velocity product at 1505 UTC . (click) Note the large area of > 60 knot winds above the surface moving into western Columbus County.



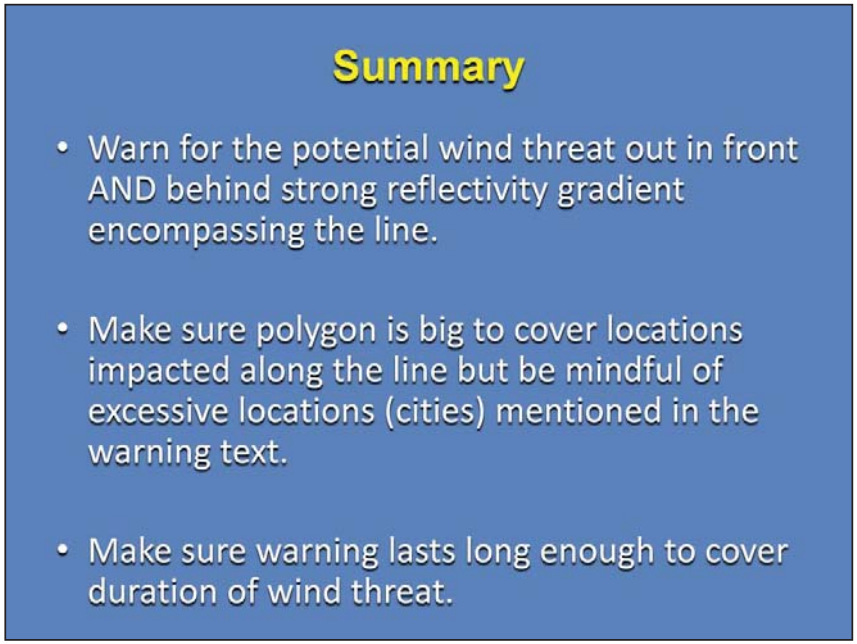
Figuring out where to place the warning polygon for this line of storms is a little tricky since the damaging winds at the surface (which are initially aloft) likely lag behind the leading edge of the strong reflectivity gradient. The gust front is aligned very close to this leading edge. Note the rear-inflow notches behind the leading edge reflectivity gradient indicative of a descending rear-inflow jet. Thus, since this is NOT your first warning (previous warning shown in dashed white), you will need to lay out the severe wind threat area ahead of the line AND a bit behind the leading edge to account for the lag. Since most portions of the line look capable of producing winds to severe limits, you will end up having to issue a pretty large warning to cover the threat. This is preferred in most line case unless you can break up the threat into adjacent polygons. Note: try to eliminate excess overlap with the new warning.

Thirty minute storm (line) motion vectors are shown with the large white arrows with projected leading line echo configuration in dashed brown. Note that based on this projection, the southern portion of the line will clear the coast, so a Special Marine Warning will also likely need to be issued now, if not previously, to enable boaters and other impacted offshore interests sufficient lead time.

In terms of duration of this storm-based warning, a 45 min. warning might be appropriate given the duration of gusty damaging winds in the wake of the leading line.

Use the WarnGen line tool to help you draw the big polygon by aligning the squall line threat area perpendicular to the line orientation, with the northern portion of the line slightly bowed out to account for accelerated line motion. Also, make sure you don't draw the WarnGen polygon into marine areas. Some important considerations: This squall line impacts counties in your adjacent CWAs, so make sure you coordinate with those offices on your storm-based warning. One big consideration when issuing storm-based warning for long squall lines: make

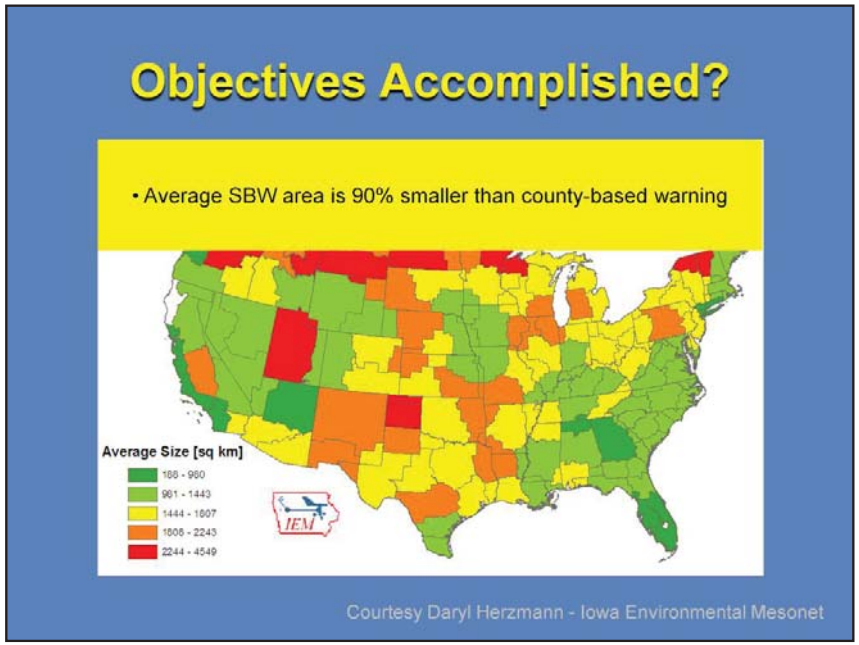
sure you don't have excessive number of locations impacted listed in the Warning Text.



So, in summary, for this type of warning event, warn for the wind threat out ahead of the line AND a bit behind the strong reflectivity gradient. Make your polygon big to cover locations impacted along the line but be mindful of excessive locations (cities) mentioned in the Warning Text. Make sure your warning has a long enough fuse to account for longer duration winds. And in this particular case where the synoptic scale forcing will overcome any effects of a stabilizing effect of the marine layer, you would need to issue a special marine warning even before your final land warning.



Welcome to the Storm-Based Warning Fundamentals lesson entitled, "Special Considerations." This lesson addresses some of bigger challenges in storm-based warnings that have shown up on service assessments from 2009 to the present.

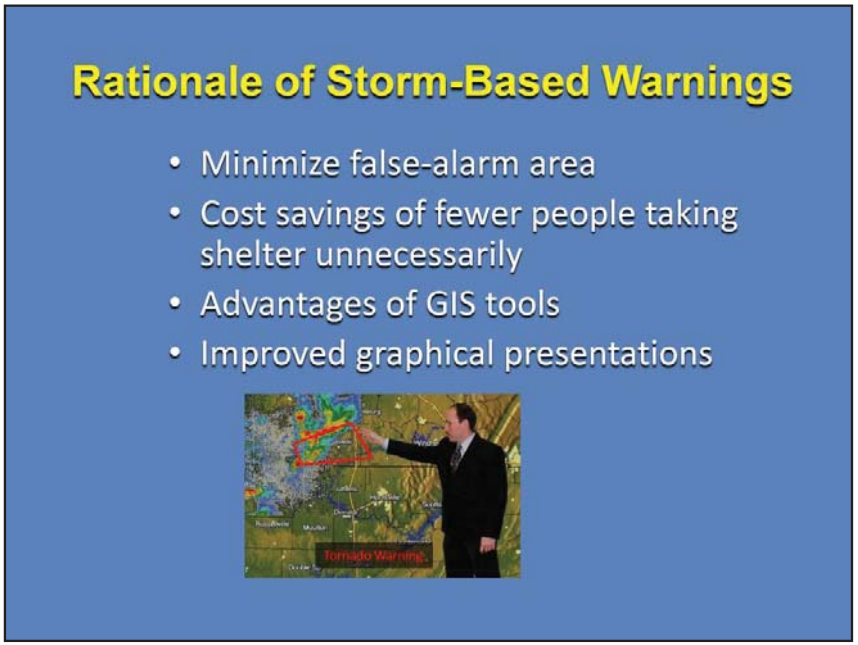


Storm-Based Warnings, or SBWs, became official NWS policy on October 1, 2007. From that point to the end of September 2008, over 31,000 SBWs have been issued. One of the primary purposes of SBWs is the reduction of Falsely Alarmed Area (FAA), which according to national statistics, has been at around 67% when compared to County-Based warnings. Other warning performance metrics such as Probability of Detection, Lead time, and False Alarm Ratio are all exceeding the SBW goals that have been set. (CLICK)

Here is a graphic from the Iowa Environment Mesonet (IEM), which is produced by Iowa State University Department of Agronomy. It shows average storm-based warning size from 1 Oct 2007 to 30 September 2008. Note that many warnings in the Western U.S. are naturally larger due to poor radar coverage and a lack of population and spotters. This graphic provides a context to the following statistics.

Reduction of false alarm areas have been most noticeable in the Western U.S., where the average SBW area is more than 90% smaller than the county-based warning area. Tracking improvement in reduction of the Falsely Alarmed Area (FAA) at the WFO level is more suited to a year-by-year analysis than by comparing WFO CWAs to each other. There are lots of issues when drawing polygons so the size varies from place to place and situation to situation.

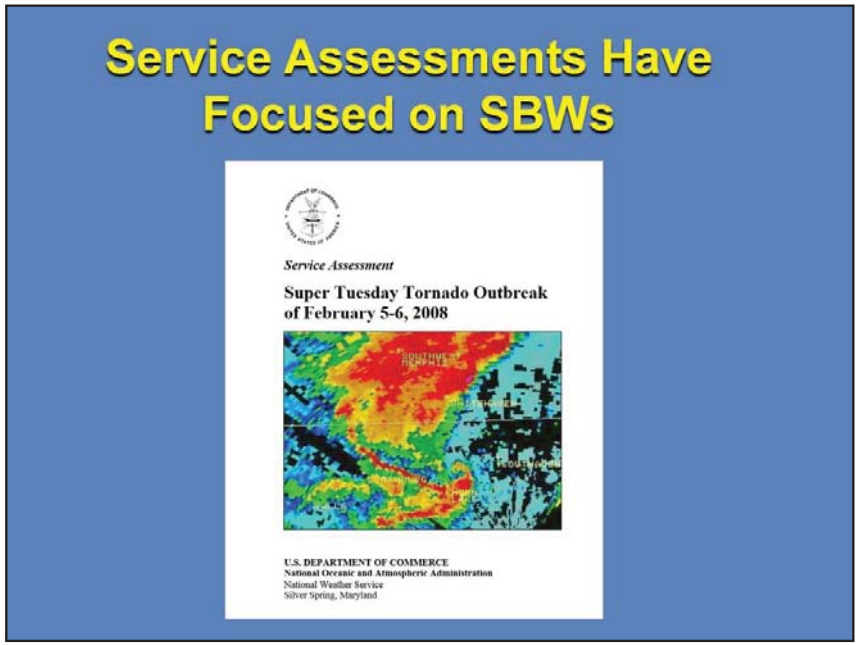
Despite very effective performance measures for SBWs in 2007-08, some issues still exist. This module will address some of these issues.



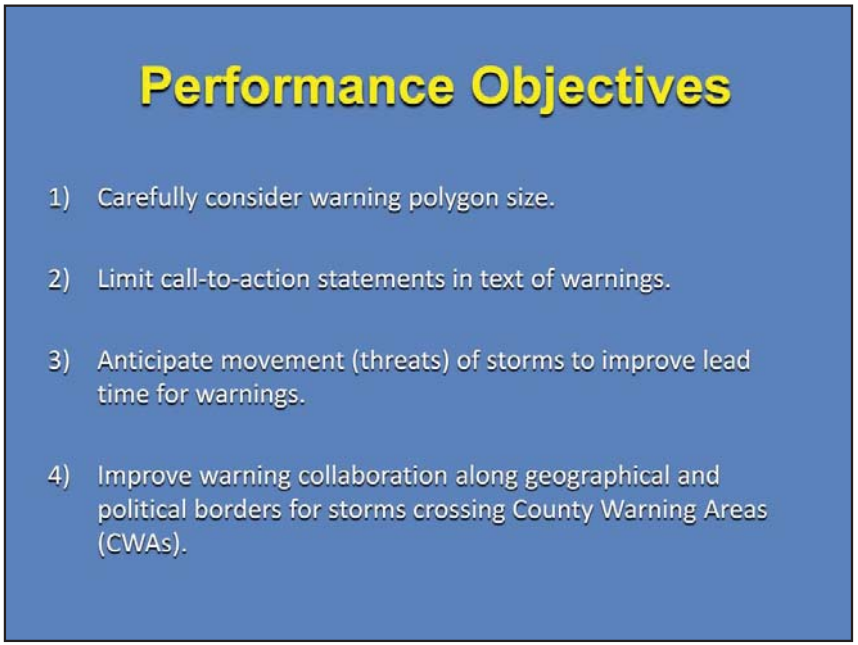
The goal of SBWs remains the same as when it became operational in 2007: to provide geographically concise, timely, and meteorologically accurate, short-fuse warning information.

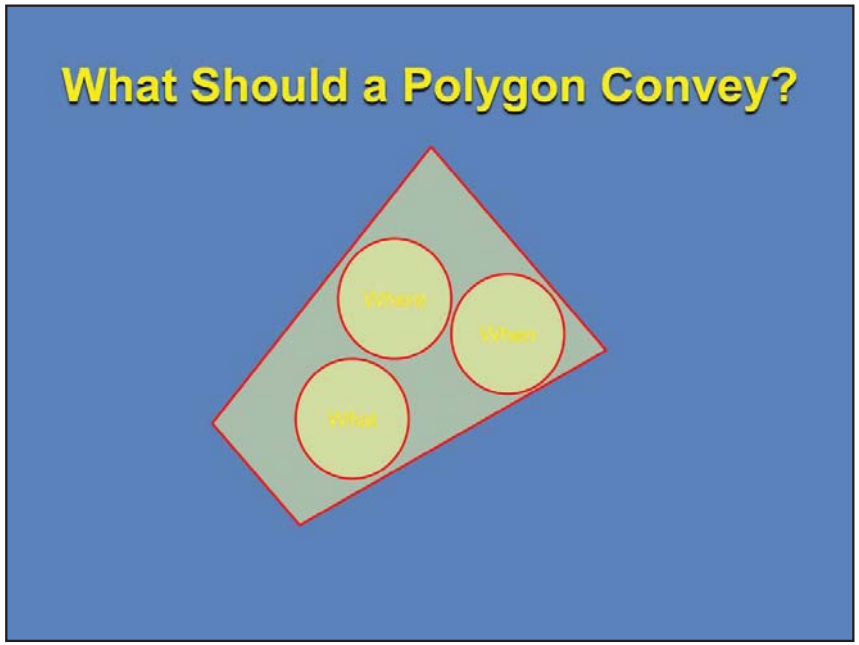
These four bullets encapsulate the primary goals and benefits of the storm-based warning system: Minimize false-alarm area, provide cost savings with fewer people taking shelter unnecessarily , take advantages of GIS tools, and improved graphical presentations of warning information.

While issues and limitations still exist in the storm-based warning system, especially in our dissemination methods, improvements such as impact-based warnings and eventually FACETs will evolve the warning messaging process into the future.



Training basis for storm-based warning has been on the operational assessment of performance from NWS service assessments such as the Super Tuesday Tornado Outbreak and the Service Assessment for the June 6-7, 2008 severe weather outbreak. More recent assessments have focused on communication and messaging issues which are a focus in the warning operations course.





These are the main points we want to emphasize in this storm-based warning fundamentals lesson.

1. Carefully consider warning polygon size. We want to minimize the use of extremely large SBW polygons in most, but not all situations. Some polygons contain so many counties/parishes that the associated warning text runs over the character limitations for some of our partners.

2. Reduce the amount of text in call-to-action statements.

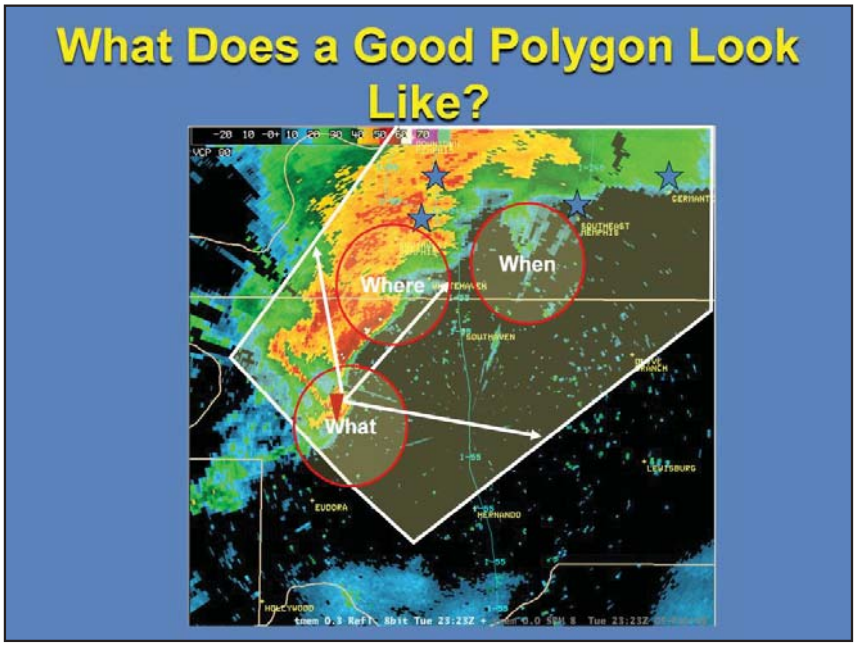
The third objective listed here, anticipate movement of storms, was actually a recommendation from the Super Tuesday 2008 Service Assessment. In the Assessment, it was found that many guidelines presented in the first storm-based warning training were not followed.

For example, many polygons in the Super Tuesday Outbreak were truncated along county boundaries, with some shortened on the downstream end, reducing potential lead time. There was also a tendency to wait for a storm to be near the downstream edge of a polygon before the next warning polygon was issued. Thus, for the long track tornadoes, the lead time from one polygon to the next was often reduced.

Thus, for the case of fast-moving storms, we want to stress the importance of anticipating the need for new warnings well before a given storm moves out of a current polygon, and , we want to encourage forecasters to not remove counties from a polygon unless the forecaster has total confidence that the storm will not impact that area.

Finally, we want to address the need to clarify the coordination process for issuing SBWs at boundaries of Weather Forecast Office (WFO) responsibility, especially where complex boundaries were involved (such as rivers). There were several instances where confusing warning products were issued between adjoining WFOs County Warning Areas (CWAs) as a storm crossed from one CWA to the next.

The storm-based warning polygon is a crucial piece of the warning conveyance mechanism. It helps describe graphically the what, where and when of the warning.



When you create a polygon, it tells the user that the entire area in the defined polygon is under risk of specific threat or threats. Try to treat the polygon as a threat event area, not a single point.

This is vitally important as you want people within the warned area to take precautions, not just along the pathcast. A good example of a well-designed polygon is this image from an event during the Super Tuesday outbreak in 2008.

The polygon uses several good practices of a storm-based warning:

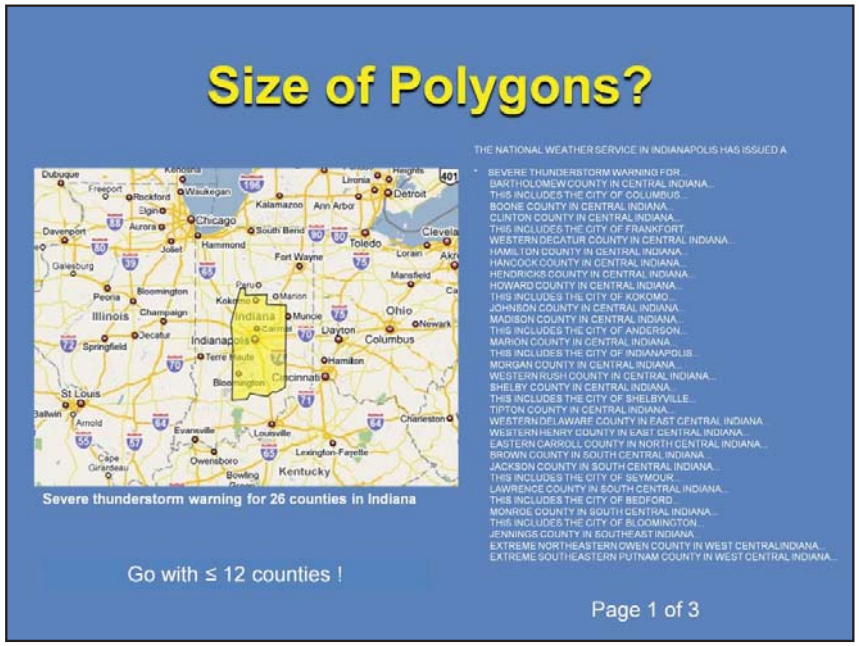
1) Depicts what the threat is – potential tornadic storm

2) Depicts where the threat is south of Walls.

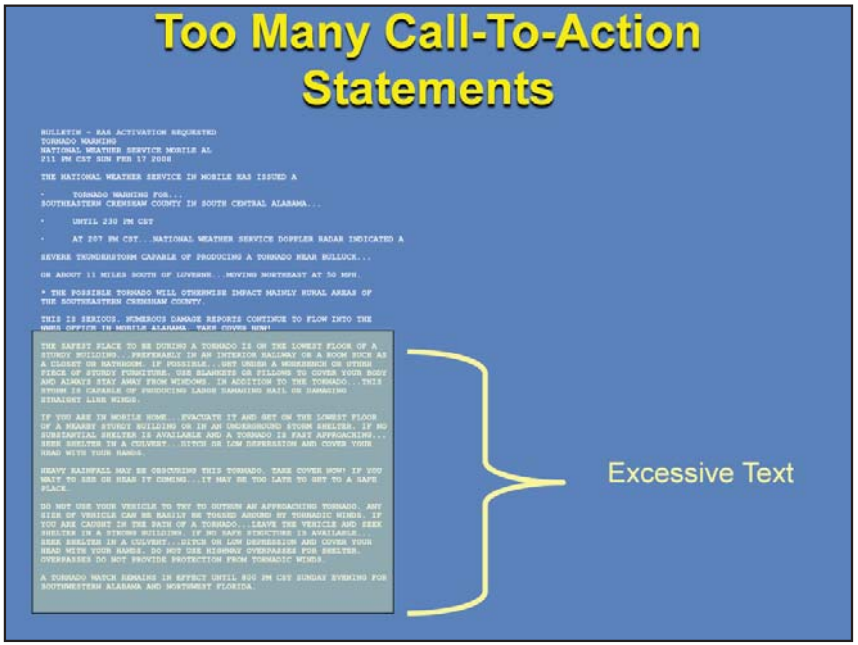
3) Depicts where the threat can occur …allows for uncertainty in forecast movement of the threat and downstream propagation, thus a fairly large (3071 km2) polygon was used.

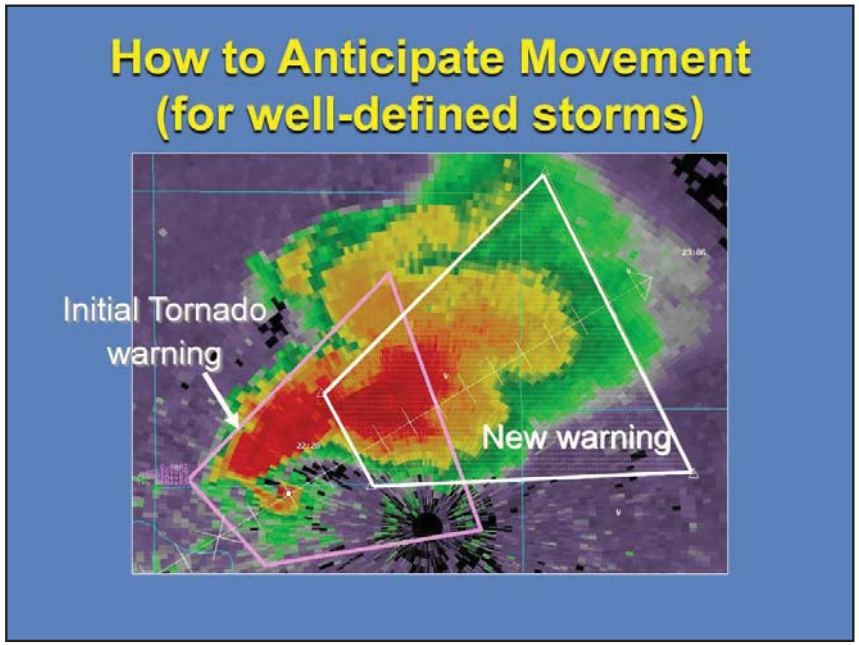
4) Depicts when the threat will occur, thus allows for a long enough lead time for decision makers.

5) All major populations areas expected to be impacted are included in area (such as Memphis and surrounding suburbs to the North and East)



Based on experiences during this first year of SBW use, a small percentage of severe thunderstorm warnings and flash flood warnings contained more than a dozen counties or parishes. For many partners, this causes problems in transmitting text warnings over mobile devices and television text crawls. There are so many locations listed that the basis of the warning is delayed or in some cases truncated altogether. This example from 2007 shows a severe thunderstorm warning in Indiana that was over 20,000 km2. Plans are to include new instruction in NWSI 10-511 and NWSI 10-922 that states tornado, severe thunderstorm and flash flood warnings should be limited to 12 counties/parishes.





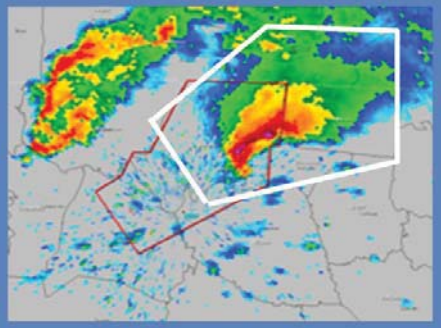
By the same token , having too many call-to-action statements can add to excessive text.

It is a best practice to include only one call-to-action statement per warning unless the situation dictates otherwise. AWIPS "Call-To-Action" tags before and after these statements now allow partners to more easily parse these statements, and if desired, remove them from cell phone text and television text crawls.

In favorable environment and high-confidence situations, it is important to issue downstream warnings well in advance. Don't wait until a storm is close to exiting an existing old warning. Go ahead and issue a new warning downstream. Tornadic supercells are the one obvious situation.

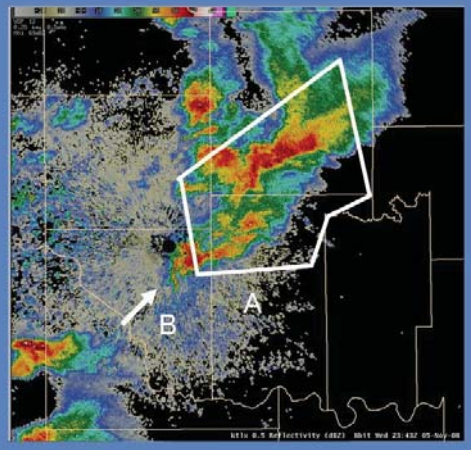
Don't Wait Too Late

- Storm that produced 22 deaths in TN
- Re-issue midway through polygon if storm hasn't changed intensity



Here's another example from the Super Tuesday Tornado Outbreak where the existing warning was allowed to continue a little too long, especially with a potential killer tornadic storm nearing the edge of the warning.

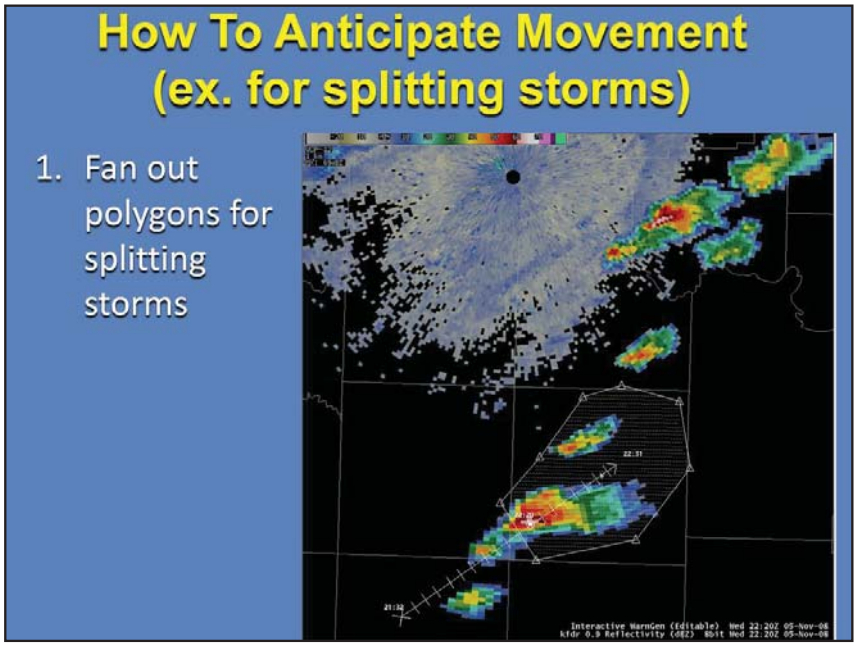
What About Removing Counties?



Based on the Super Tuesday Tornado Outbreak Service Assessment, some offices were too hasty in removing counties from warnings. This example illustrates a situation where despite the fact the principal hail threat is not in county A, it is advisable to include the portion of the county in the polygon (and not remove it) because of the potential, albeit slight, that the storm could still impact the area. Thus, for the reasons stated, and the fact that a weaker storm upstream in County B is moving into County A, how about this configuration for a polygon? (click). 20 minutes, it looks like the right decision as new development has occurred in northern portions of County A.



It's important to anticipate movement in creating storm-based warnings for fast-moving storms such as this one which produced an EF4 tornado in Jackson, TN, on Feb. 5, 2008.



Storm-based warning polygons are most useful when forecasters incorporate the anticipated movement either due to propagation or advection. We have two examples to illustrate the method.

First example is to fan out the polygons when storm morphology is undergoing an evolution such as storms splitting, which can be anticipated by assessing the environment. Before the storm evolves and undergoes splitting, you can capture the threats by fanning the polygon out as depicted here. Keep in mind that , as with most storm splits the right mover will be stronger due internal dynamic of the supercell process.

How To Anticipate Movement (for complex storms)

2. Capture multiple threats in single polygon



The second example shows an example which fans out the polygon when storm morphology, such as multiple threats in close proximity, is undergoing complex evolutions. As these storms start to develop cross the Mississippi River, notice how the two lead storms develop discrete updrafts . Thus, you can still capture the threats by putting them in one polygon. It might be a good idea to maintain a slightly longer warning duration, say 45 min, in this situation since you have propagation effects adding to the total duration time of the threats in these areas.

Issues along CWA and Geographical Boundaries



Two special marine warnings (in blue) for Chesapeake Bay and Potomac River. Note the small spike for the Patuxent River.

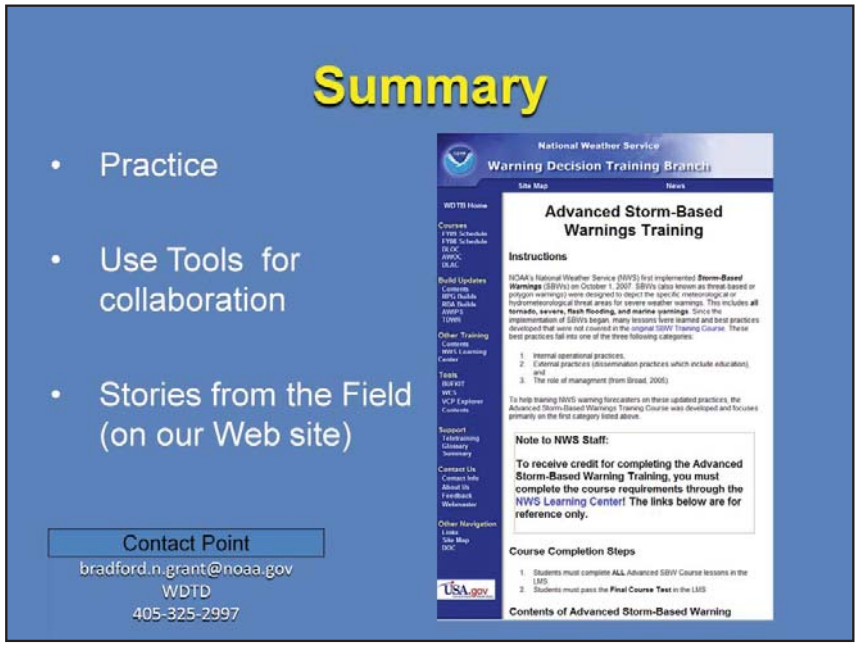
Most boundaries between WFOs county warning areas (CWA) are political boundaries that follow county lines. These are often rivers or other irregularly shaped boundaries. Since all warnings stop at the CWA boundary, this produces odd shaped polygons, and more importantly an inconsistent service for communities along these boundaries such as in this image of a severe thunderstorm polygon in yellow from the Ft. Worth WFO. The storm was moving east northeast along the Red River so the skinny connection was drawn in WarnGen due to the irregular borders along TX/OK.

Another example (click) is the marine warning shown in blue. When there are complex marine or land boundaries, due to the limit of 20 points to a polygon, when you graphically try to describe these complex boundaries with only a few points, it can produce very odd results.

As a frequent and reoccurring problem we recommend that to alleviate this , offices coordinate better to avoid odd shaped polygons.



When a storm is about to cross over into an other CWA, or about to move into your CWA, it is recommended that you call the warning team and collaborate on the characteristics of the threat (especially movement) to coordinate polygon placement to maximize lead time. Don't wait for the storm to reach the CWA border, issue the warning well in advance. Use NWS Chat frequently to coordinate on aspects of storm movement.

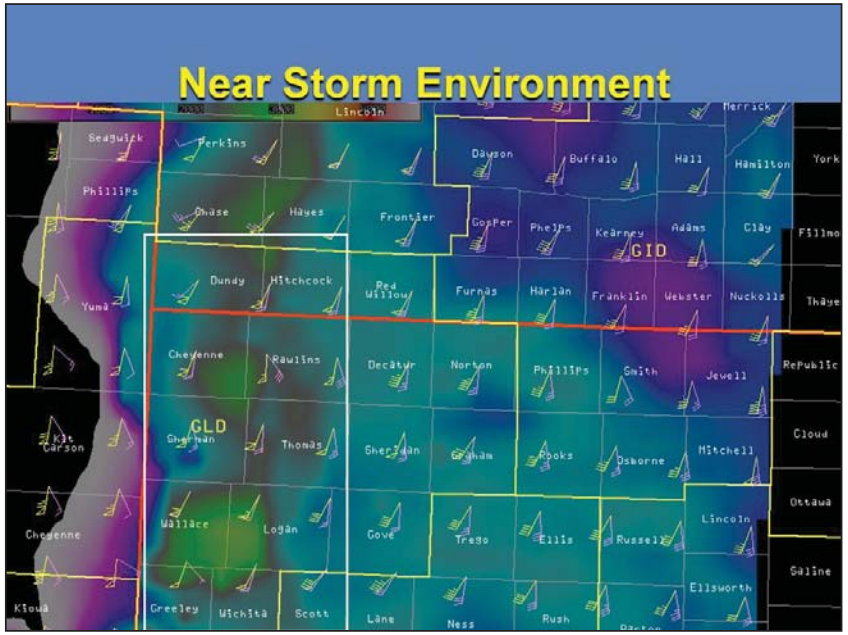


In summary, we have addressed some of the current issues NWS offices still struggle with in issuing storm-based warnings. Most of these will be overcome with more experience, practice, and good office to office collaboration. In addition to this Articulate, please see the stories from the field, which contain real-life warning cases, to help you learn more about issuing effective storm-based warnings. And finally, as with all WDTB training, please contact us at the information listed for rapid response to any questions.

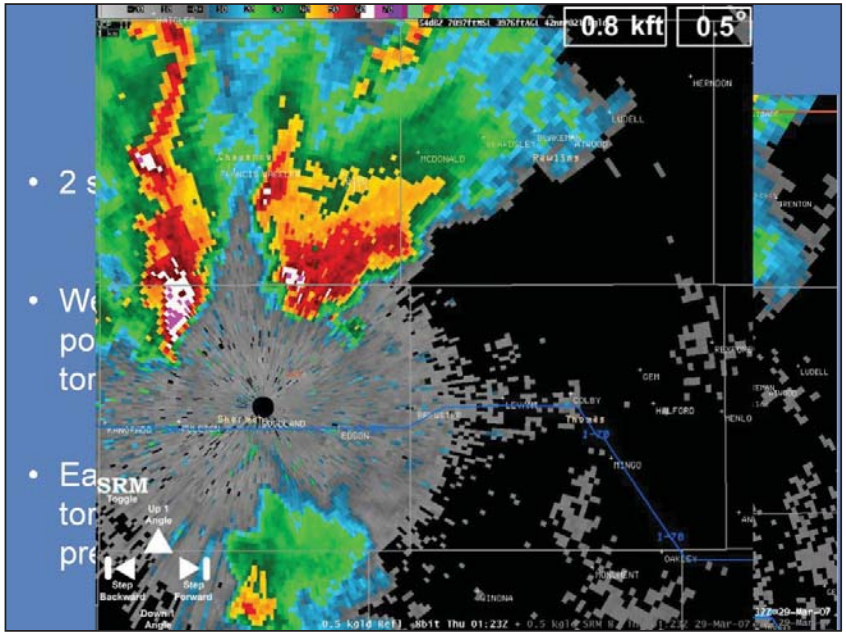
Thank you.

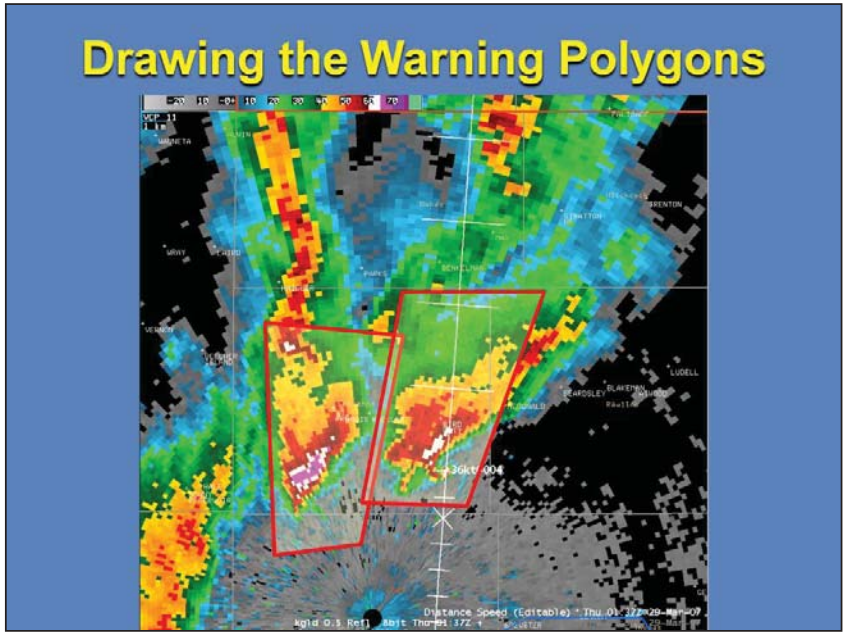


This short lesson in the RAC topic on storm-based warning fundamentals is entitled “Two TORs in Close Proximity.” It is another instructional example showing you how to handle basic warning polygon situations. The material shown here is from Objective T2 of Lesson 2 of the original Storm-Based Warnings course, and it will be narrated by former instructor, Paul Schlatter. We would like to thank Al Pietrycha (SOO) and the rest of the Goodland office for the data and for reviewing this objective.



We are concerned with the Goodland, KS CWA, and in this D2D graphic is outlined in yellow, while state boundaries are in red. The environment was very supportive of supercells, firing along the dryline located along the Colorado/Kansas border. This is LAPS data for 0000 UTC on March 29th, 2007. Mixed layer CAPE is very high just east of the dryline across NW Kansas. Yellow wind barbs are 0-6 km wind shear, while pink wind barbs are 0-1 km wind shear. The western-most counties in Kansas and Southeast Nebraska outlined by this white box have very favorable shear for supercell tornadoes, and plenty of instability. SPC issued a Tornado Watch for this area, and a moderate risk. Storms have already fired along the dryline, and right moving supercells are moving just east of due north at 30-40 kts. Let’s examine the radar situation.





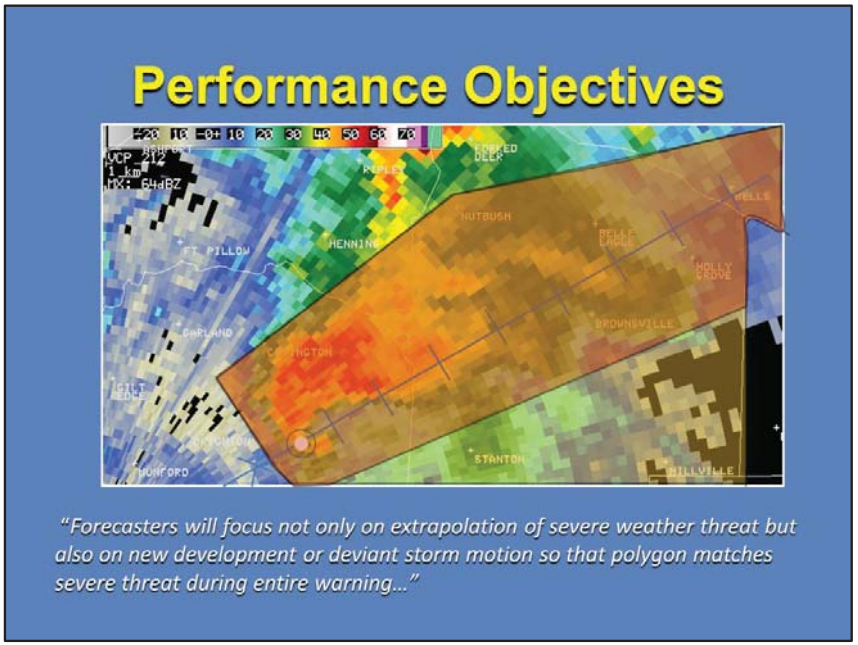
Here is the current situation at 0137 Z, with 2 impressive looking supercells moving in essentially the same direction and speed, across southern Cheyenne County, Kansas. The western storm has only produced severe hail up to this point, but in the last few volume scans it has intensified in terms of reflectivity and structure, and the low level mesocyclone has intensified, such that a tornado warning would be warranted by this time and probably even 10 minutes prior to this time. Distance speed tool is included for the eastern storm, notice that the storm is moving 4 degrees east of due north. Radar signatures and spotter reports have indicated that a tornado was occurring at this time. A new TOR warning replacing the old one will be required for the eastern storm, and a new TOR, essentially an upgrade, will be required for the western storm. To give you an more thorough radar overview of both storms, a "poor person's" all tilts flash graphic should now load in a separate window where you can toggle between Z/SRM, and step up and down in tilts for 4 volume scans. To cut down on analysis time and bandwidth, I've only included the 0.5, 2.4, 4.3, 7.5, and 12.0 degree tilts, which hit both storms at reasonable height increments AGL. As you step through in time and elevation angle, pay particular attention to: The debris associated with a tornado from the eastern supercell associated with a deep well-defined mesocyclone, the cyclic nature of the eastern supercell with 2 mesos evolving over time, the high reflectivities and strengthening mesocyclone in the western supercell, and WER/BWERS on both storms.

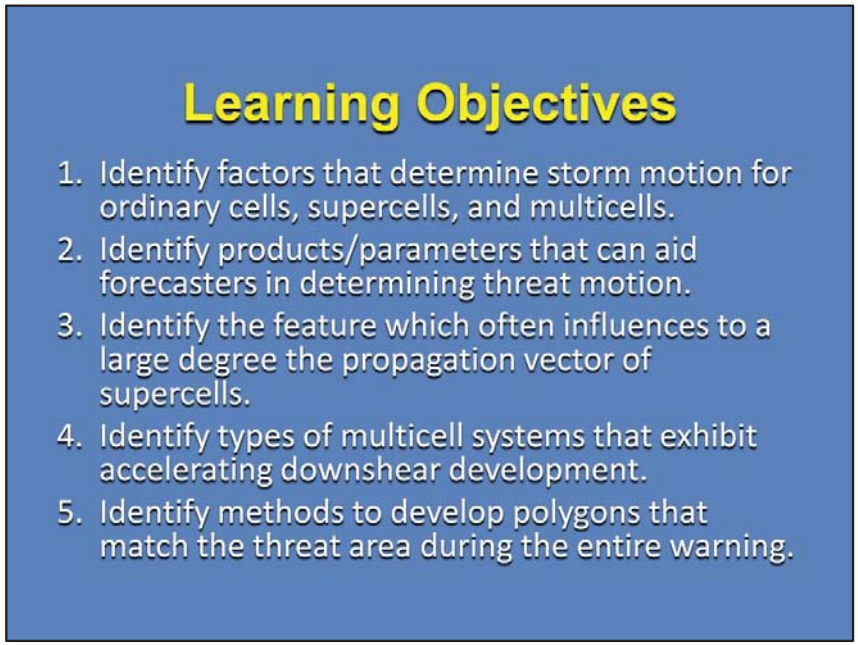
The environment ahead of these storms continues to be favorable for severe weather and tornadoes. With an expected storm motion of 36 knots, and in the interest of not making the polygons too big, it was decided to go with 30 minute tornado warnings for both of the supercells. In this event, it wasn't prudent to draw separate SVR and TOR boxes since the rotational signatures are moving along and just behind the areas expected to have the greatest hail threat. For the western storm , this was the tornado warning and it was issued at 125 UTC (12 minutes prior to this scan). St. Francis would be specifically mentioned in the text product, as would a mention of the threat of hail up to the size of tennis balls. Playing it safe, there is not a gap in between the 2 polygons . For the eastern storm, I couldn't rule out the possibility of the older mesocyclone containing a tornado at this volume scan, though clearly the newer mesocyclone located on the southeast flank of the storm, complete with a debris signature, is the stronger of the two. Thus, the tornado warning is further west to include the older mesocyclone. For this storm, Bird City would be mentioned as in the path of the tornado, and hail up to the size of golfballs would also be possible with this storm. These two warnings overlap across central Cheyenne County, thus a good deal of coordination would be required among warning forecasters and with their customers. The western polygon is more flared out because I am less sure about storm/mesocyclone motion than I am with the right polygon. An clear problem arises with NOAA weather radio with this type of situation: 2 warnings are in the same part of a county valid at

the same time, potentially confusing our customers. This concludes objective T2, thanks for your attention.



Welcome to RAC storm-based warning fundamentals lesson on non-linear motion. This instruction was designed to support the NWS training requirement to provide best practices for developing storm-based warnings. This lesson will deal with effective placement of warnings for storms with non-linear motion.

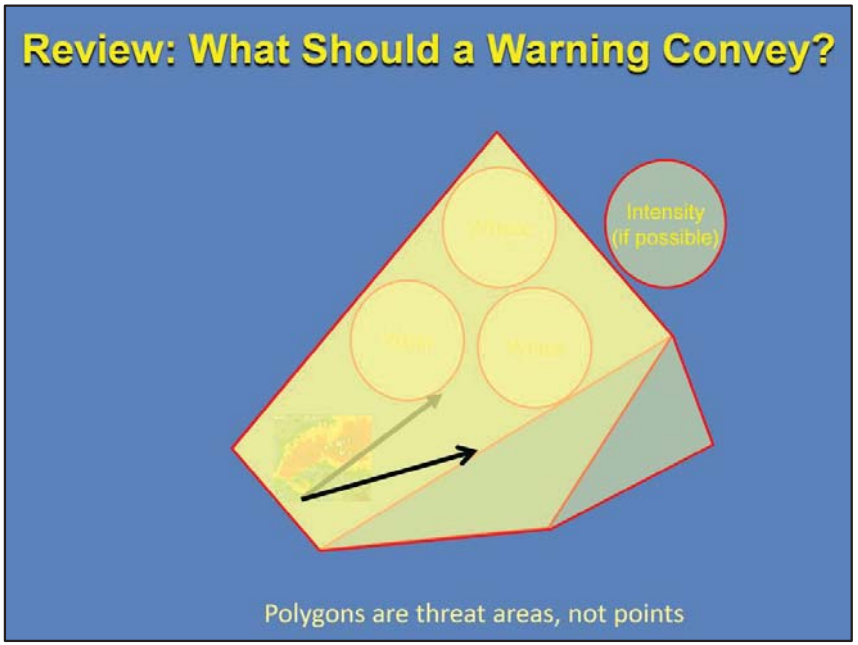


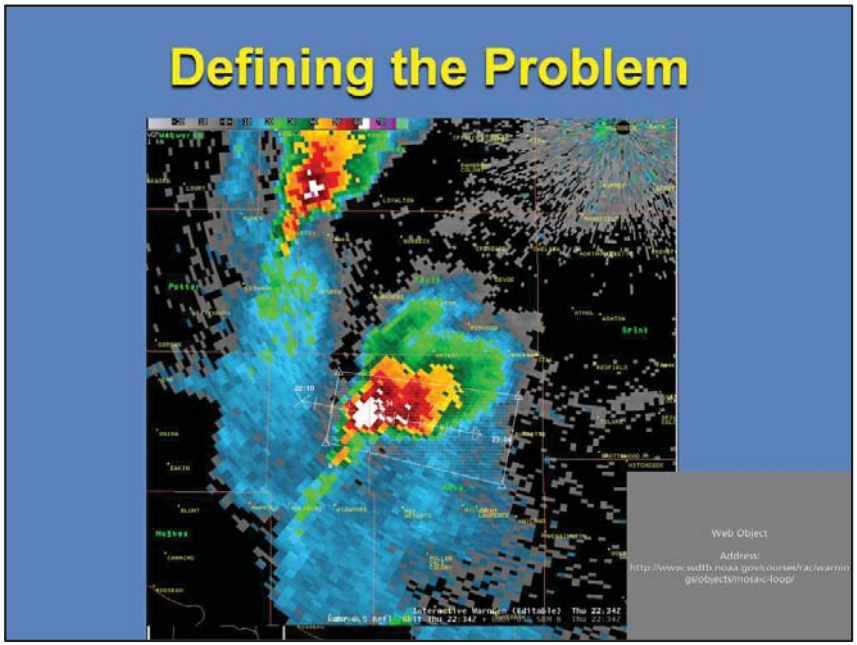


The performance objective for this particular lesson contains another Best Practice for issuing warnings:

Forecasters will focus not only on the extrapolation of the severe weather threat, but also on new development or deviant storm motion so that the polygon properly matches the severe threat during the entire warning time frame. So, to accomplish this, we are going to demonstrate ways to incorporate the use of AWIPS capabilities in conjunction with radar views and diagnostic fields as the polygon is being created. In addition, we are going to talk about ways you can issue polygons ahead of the threat when storms are moving in from an adjacent CWA.

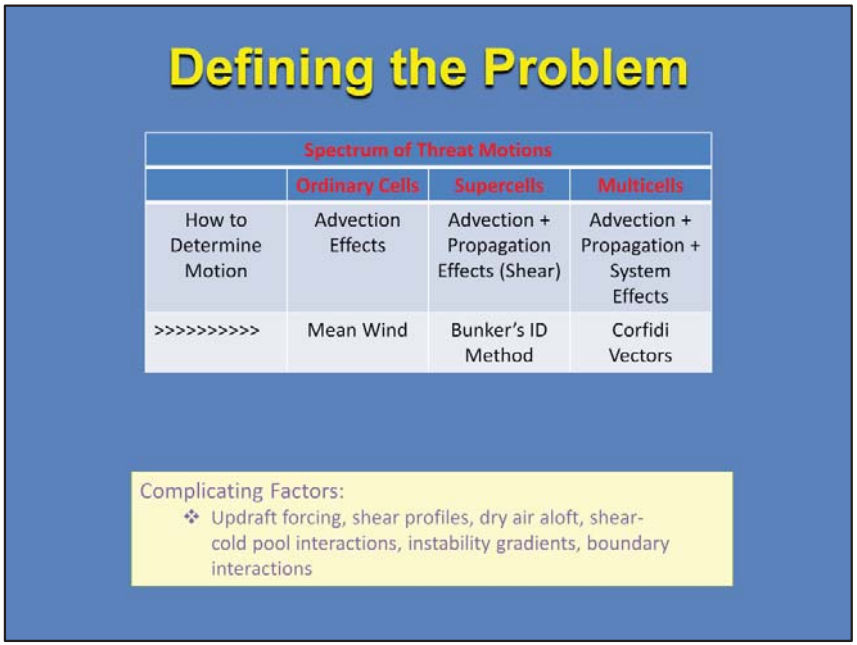
Here are the learning objectives. After you finish reading these, please advance to the next page.

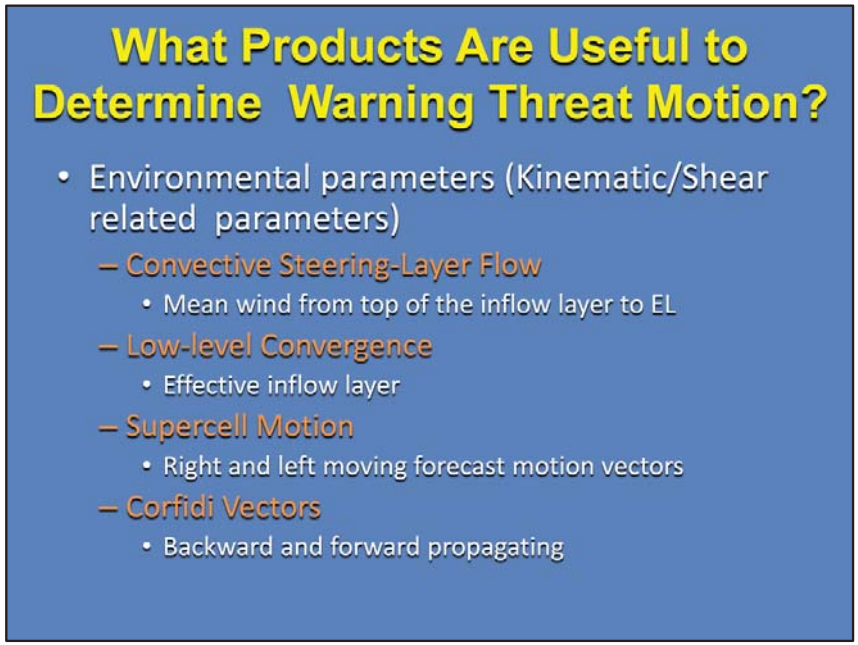




The warning conveyance mechanism should describe the what, where and when of the threat. The text should also help to convey the intensity of the threat (if possible) , and urgencies for taking precautionary measures. The where part is particularly difficult as determination of the forecast track dictates 2/3rds of the warning product. So, given this example supercell storm, if you start with a simple linear extrapolation, you'll typically get a track and associated polygon as shown with four vertices. Due to typical non-linear turning of the supercell, it is wise before you issue the first warning, to add a 5th vertex on the southern side of the polygon. If acceleration is expected , or downstream development is anticipated, an even better approach is to add a 6th vertex on the eastern side of the polygon as shown . Remember, you can't expand the original polygon, but as the threat evolves, you are encouraged to issue frequent updates. From a workload management perspective, if you start off with the best possible polygon configuration, it will help you over the duration of the entire warning. And finally, my advice is to use the storm-based warning polygon as an areal threat forecast, not a point threat update.

So again, the problem is how to determine storm motion so that the threat area in a storm-based polygon depicts the three basic elements of a warning: What, When, and Where. On paper, fairly simplistic, but in reality, quite complex. It is rarely as easy as moving the drag me to storm icon back a few frames and then letting WarnGen compute the polygon and track. Hey, if it were that easy, warning forecasters could get replaced by a computer application. Determining an accurate short term forecast area of the storm threat involves a 4-Dimensional analysis of observed conditions and a forecast of expected conditions. There are many mesoscale and storm scale factors that influence the forecast such that rarely will conditions allow a simple extrapolation of past radar motion to provide an effective warning threat area. This is especially true in cases where storms are exhibiting rapid, downwind propagation such as in this example from the May 8, 2009 multicell severe weather event which moved across several CWAs. This loop illustrates many of the complicating, non-linear factors associated with determining storm motion. It is extremely challenging to determine an accurate threat motion and then incorporate that into warning polygon and track creation. But, we are going to look at a few things that might help.





Determining the motion of threats is complicated and it obviously varies with storm type. In a simplistic sense, the spectrum of threat motions starts with ordinary cells, which are primarily influenced by advection of mass through the mean wind. With increasing amounts of shear, storm movement becomes a function of not only advection but propagation (due to the interaction of vertical wind shear on the updraft). This is where Bunker's ID Method (which is covered in DLOC Topic 7) of forecasting supercell motion is an important aid. Finally, with the most complex storm system, multicells , for determining motion, you also have to take into account system effects, such as cold pool strength and influences of Rear-Inflow Jet and the Coriolis Force, which influences system movement after several hours of evolution. For both backward propagating multicells and forward propagation, Corfidi Vectors are available to estimate system movement. Complication factors to determining motion include updraft forcing, depth of shear profiles, amount of dry air aloft, shear-cold pool interactions, instability gradients, and boundary interactions, which can override many of these factors by themselves. Topic 7 of DLOC contains several lessons on multicell motion, which would be a good review for warning forecasters.

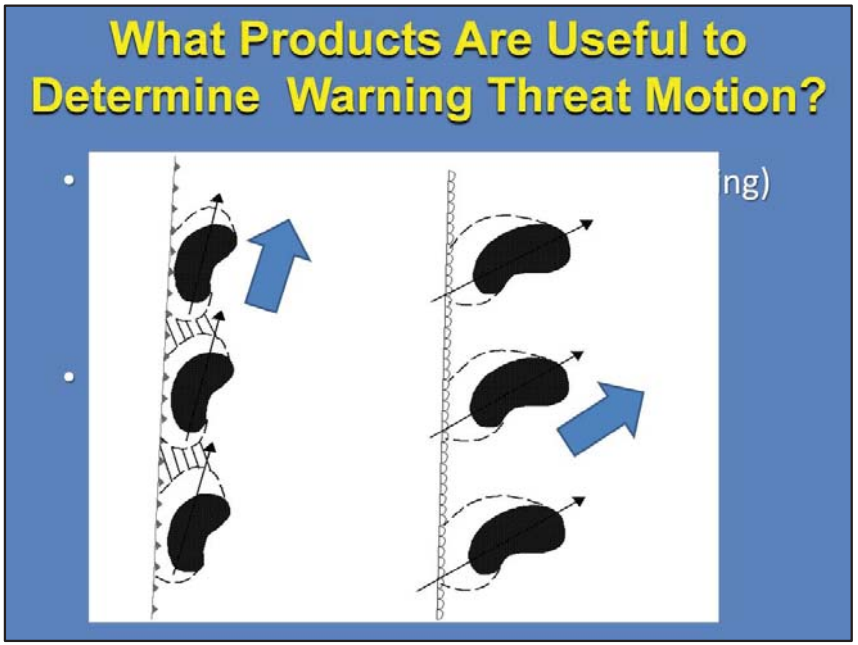
Here are some environmental kinematic parameters that can be useful for determining threat motion over the entire duration of the warning:

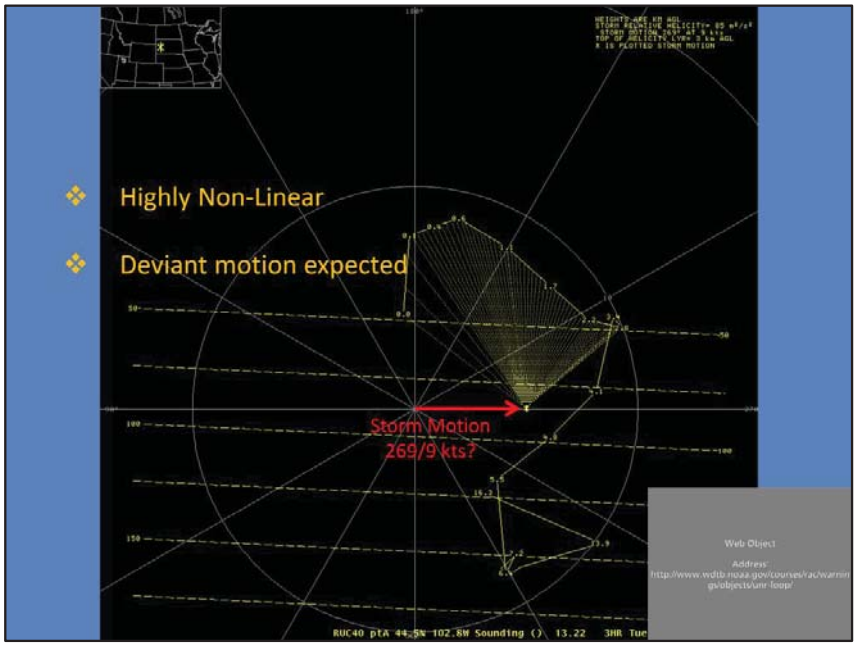
1) Convective Steering-Layer Flow, typically estimated by assessing the Mean Wind from top of the inflow layer to equilibrium level.
2) Low-level convergence, which influences updraft location and resulting storm/system propagation and maintenance, among other things. Typically, for storm motion, you will want to assess low-level convergence throughout the inflow layer of the storm threat area, usually in the lowest 1 to 2 kilometers. For storms whose updraft parcels are not rooted in the boundary layer, such as in situations of elevated convection, for determining motion, you may want to estimate effects from convergence at a higher depth above ground, say up to 3-5 km AGL.

3) Supercell Motion, use Bunker's Storm Motion estimate vectors for right and left moving supercells. Keep in mind storm-scale interactions such as collision of left and right movers, and instability gradients will likely alter the deviant motion estimate.

4) Corfidi Vectors, this is best for multicell motion, use both Backward and Forward Propagating vectors.

Note: there are lots of other complicating factors that influence multicell motion, and



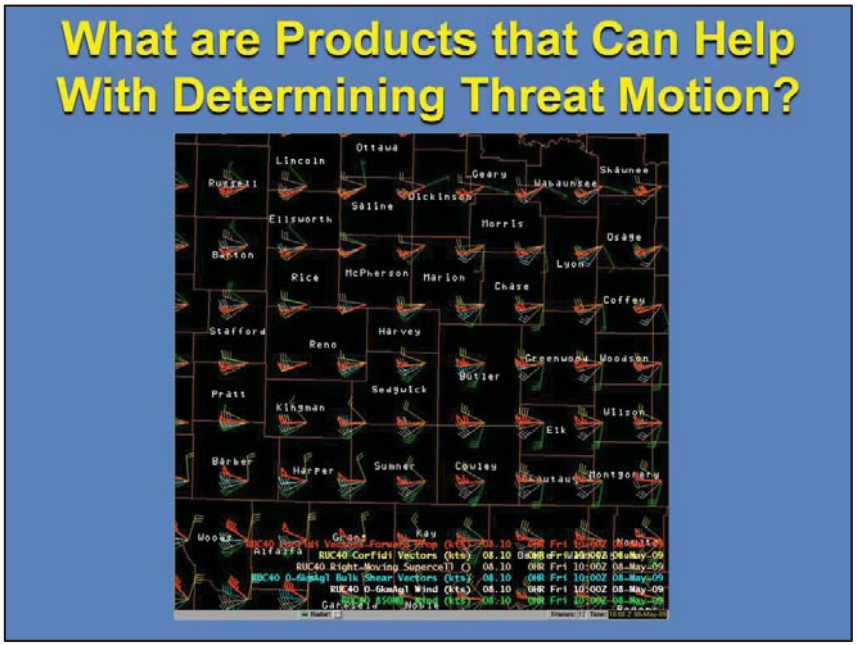


Many environmental parameters modulate updraft forcing and depending on the particular situation can significantly influence a storm or storm system motion, especially after storm initiation. Influences of CAPE and CIN and these other updraft related parameters are pretty well documented and have been covered in DLOC and AWOC. The role of boundary interactions is probably not as well known and can over-modulate all other factors even with multicells.

One of the big influences on resulting storm type is the orientation of the shear vector on the boundary itself. For example, in general, boundary parallel shear usually leads to multicells whereas boundary normal shear often leads to discrete supercells. This configuration is normally valid only in the initial phase of convective development before a storm system becomes mature.

As a start, we'll look at some factors associated with supercell motion. Observations and numerical modeling simulations suggest that supercells frequently do not exhibit purely linear motion and steady-state evolution throughout its entire lifetime. In fact, most supercells exhibit characteristics that are highly non-linear owing to effects of shear on the updraft. This deviant motion is typically estimated by evaluating a proximity storm hodograph like the one shown. For example, in a sheared environment characterized by a cyclonically curved hodograph (as shown here by a series of RUC forecast hodographs), supercells would be expected to deviate to the right of the shear vector. As an example to illustrate the effects of Supercell propagation on motion, note how this supercell propagated southeastward across western South Dakota on 13 July 2009. In determining an accurate polygon threat motion for this event, I overlaid the Right-Moving Supercell vector from the AWIPS volume browser, plotted on the radar loop by the blue vector, derived from the NAM. Even though the ID method provides a physically viable supercell motion estimates for both Right and Left-moving vectors (which are available in AWIPS volume browser) and are superior to the motion plotted on the AWIPS Skew-T hodographs, there are still some uncertainties on what the actual deviant motion is because of other mechanisms, like boundary influences.

In this example, the location of the rear flank downdraft and associated gust front acts as a focal point for the storm updraft to ingest vertical vorticity and propagate along the boundary. The storm itself appears to decelerate and turn to the right toward the end of the loop, so that the actual path of the potential tornado vortex moves to near Cottonwood, whereas a purely linear track follows a track further north. From a service standpoint, examples like these lend themselves to issuing frequent follow-up statements which adjust the storm's track.



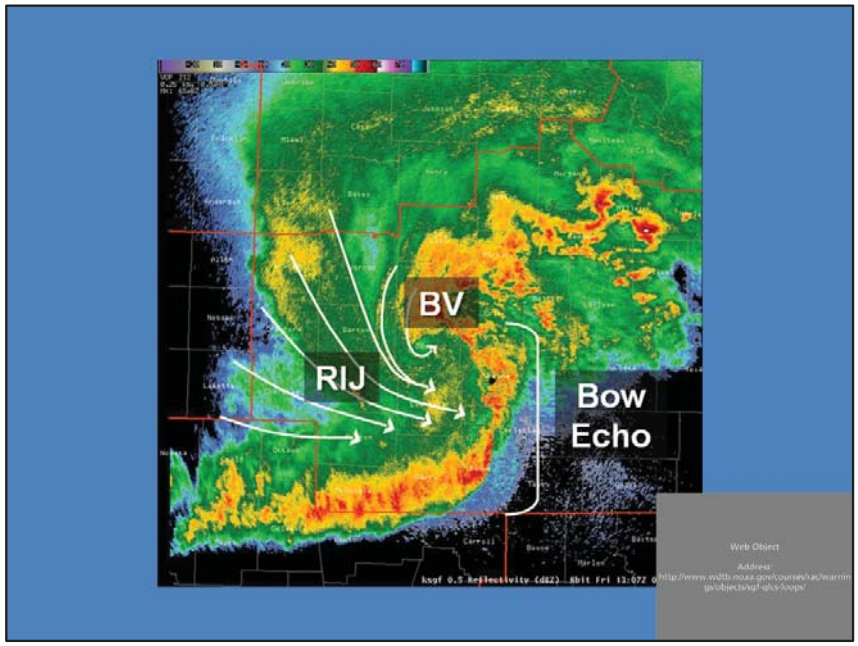


Now, as we transition to multicells and determining threat motion, again, estimating threat motion gets more complicated. You'll typically assess a number of potential factors that provide guidance. You can use the Volume Browser to create Procedures which display storm motion estimates such as Bunker's Right-Moving Supercell and Left-Moving Supercell vectors, Corfidi Vectors for multicell motion, both for backward (just called Corfidi Vectors) and forward propagation. In this example, I am showing plan-view plots of various Motion Vectors as forecast from the RUC. These vectors are useful in helping forecasters determine a reasonable estimate for threat motion for ALL storm types. The case shown (08 May '09) will be used to illustrate how difficult it is to determine storm motion in a complex storm evolution when linear extrapolation falls short.

When the multicell system first evolved and congealed into a linear structure in south central Kansas, it appeared that it would move SEWD into OK. However, owing to a development of the system cold pool, low-level convergence along the leading edge of the system, and a strong Rear-Inflow Jet, the system began to turn east and accelerate into SE Kansas and eventually into SW Missouri. Along the way, you can observe non-linear motion as transverse bands of convection formed by the strong isentropic lift and convergence of large mass of moisture due to the low-level jet. In addition, a bow echo developed just northeast of the radar site and helped to absorb a couple of supercells which at the time were moving ENE.

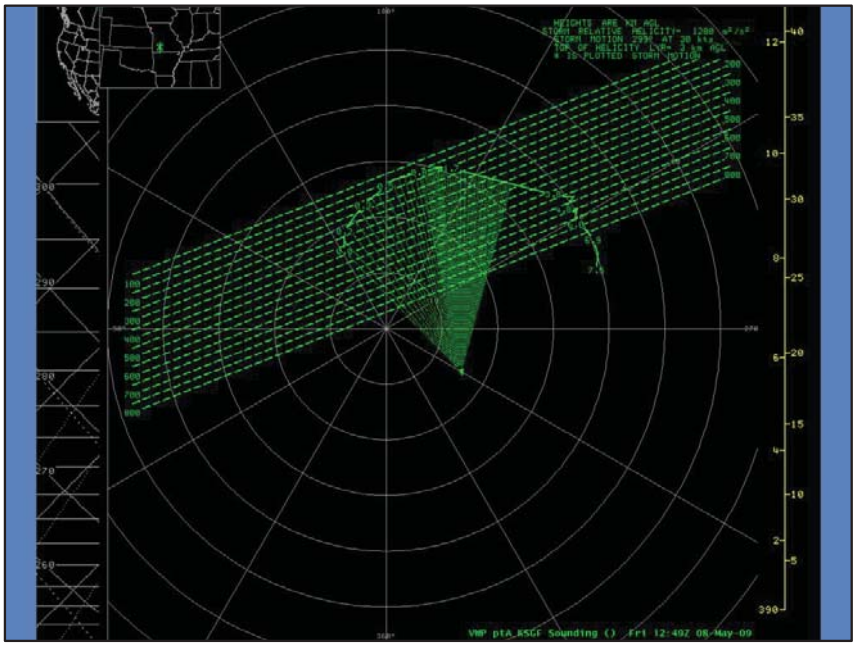
Web Object

Address:
http://www.wdtb.noaa.gov/courses/rac/warnings/objects/kict-motion-aided/

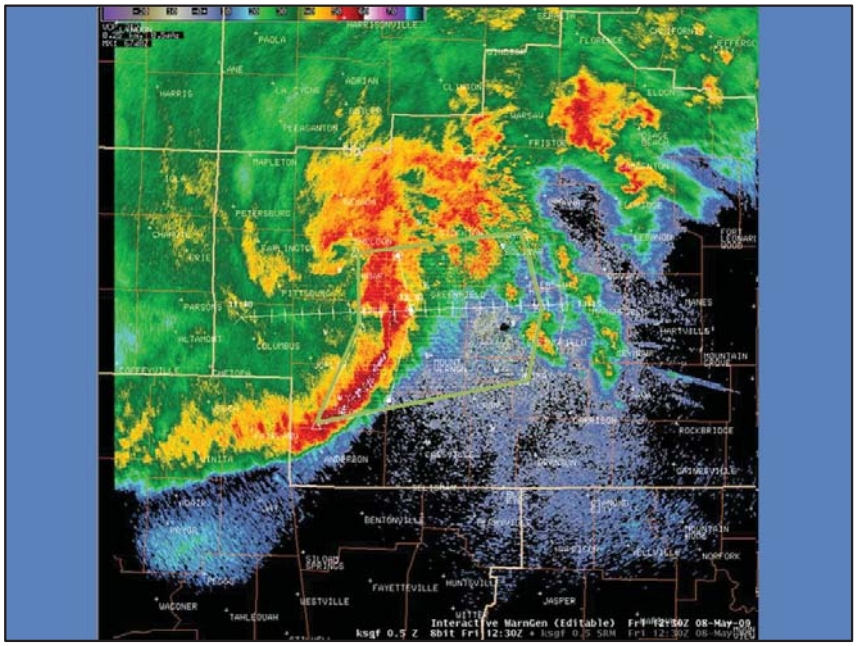


This next loop shows how the storms are moving relative to the 0-6 km mean wind (the yellow vectors). When a multicell system begins to exhibit rapid downshear movement due to increasing low-level moisture flux convergence, and strong mid to upper level shear, it is usually necessary to estimate motion using the Corfidi Forward Propagation Vectors, shown here in red from the LAPS. Again, simple extrapolation will underestimate the complex motion of individual updrafts and resulting cold pool motion of the system.

For this case, over the next several hours, there were many factors influencing the motion of the system and subsequent threats. If you were trying to track the motion linearly like the WarnGen track and box shown, you might mis-represent the highest impact areas. There are many external factors which are influencing the asymmetric shape of the multicell complex and variable threat motion including: for example, the development of at least two accelerating bow echo segments, a tremendously powerful RIJ (Rear-Inflow Jet) as evidenced in Velocity data and several rear-inflow notches in Reflectivity data, and a big, bookend vortex (BV) on the northern end of the system which is causing the entire line to turn to a more easterly direction and take the major wind damage directly toward Springfield in Green County. Please take note of mesovortices along the line which move considerably different than the line itself.

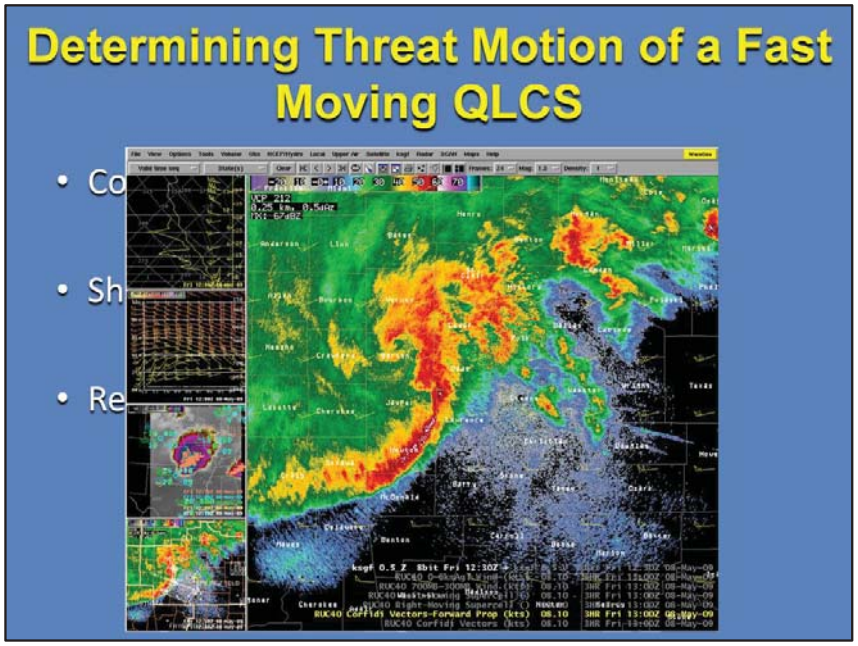


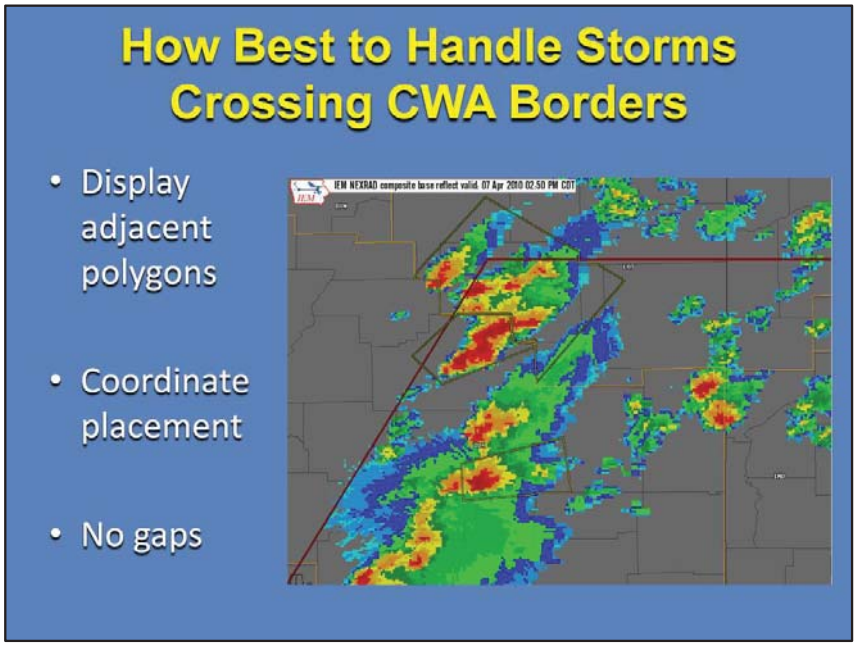
So, as in this particular case, if you take into account some of the environmental data, such as the 12z sounding and accompanying VWP hodograph from the KSGF radar, you can start to recognize some of the factors that are affecting your storm motion, that again, in this case, is causing the entire system to move more easterly and develop potential tornadic mesovortices along the way.



Now, let's incorporate some of the environmental factors into the warning motion of the leading line, specifically, the Corfidi Vectors for Forward Propagation, we can see what a non-linear threat area might look like, at the ending time of the previous loop.

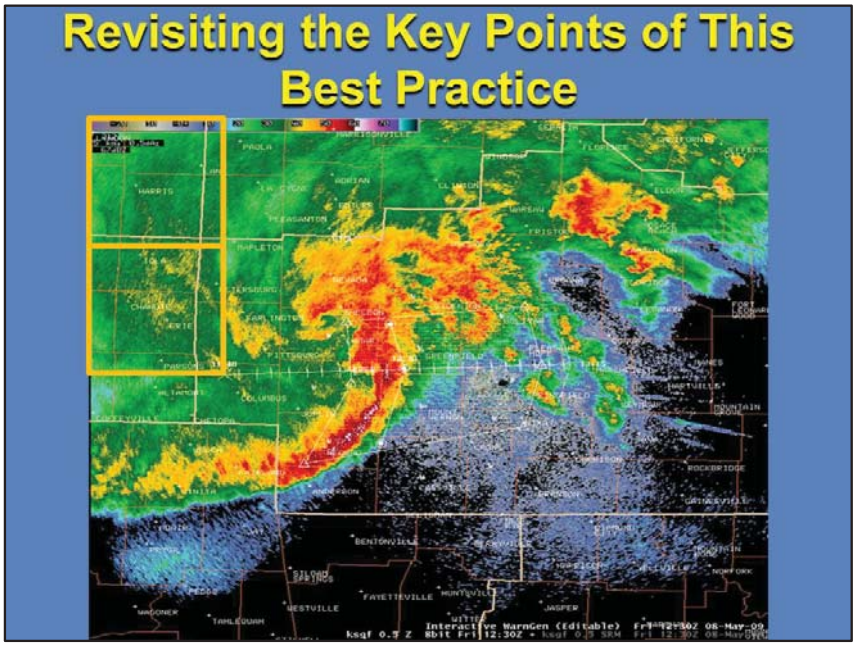
If we use the Forward propagation vectors as a reasonable first estimate, we can draw a slightly different threat area and track.





In summary, when the Cold Pool is large and effecting a mature multicell system, the deep layer shear very strong, and a RIJ is persistent , use the Corfidi Vectors for Forward Propagation to help determine non-linear system threat motions.

One last thing that I want to mention is about situations when storms are crossing CWA borders. As a best practice, it is important to include warning polygons from local and adjacent CWAs on your warning generation display monitors. In this example, note how Chicago (LOT) issued an unusual multi-sided polygon that matched up well and allowed for downstream movement from ILX's severe thunderstorm warning. In many situations, current displays like these in the warning preparation phase will help minimize unintended gaps between warnings and provide a seamless warning service for our users.

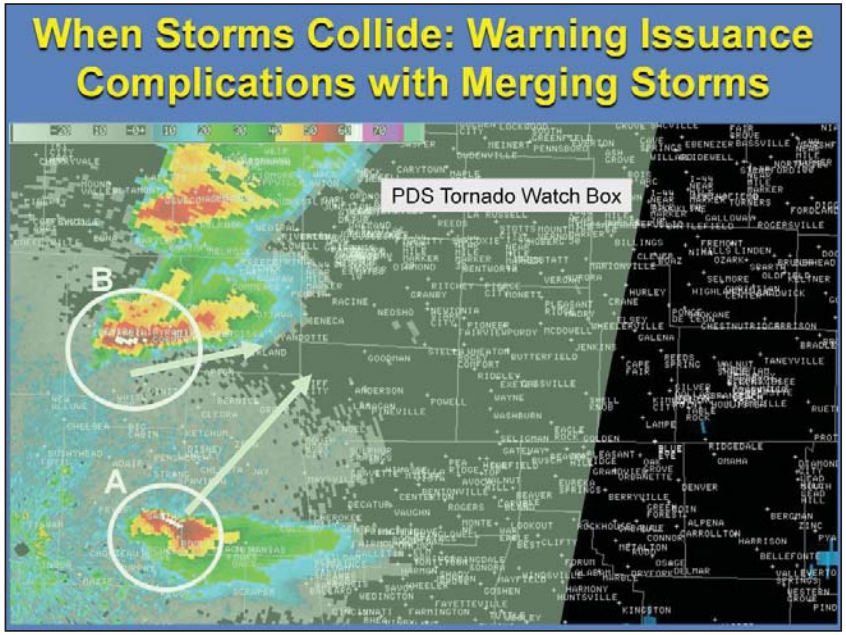


Many factors influence storm motion: updraft forcing, shear, boundary location and orientation, among others.

Know the relationship of these physical factors on storm motion and use those relationships to help guide the effective placement of storm-based polygons so that they incorporate non-linear and/or downstream development.

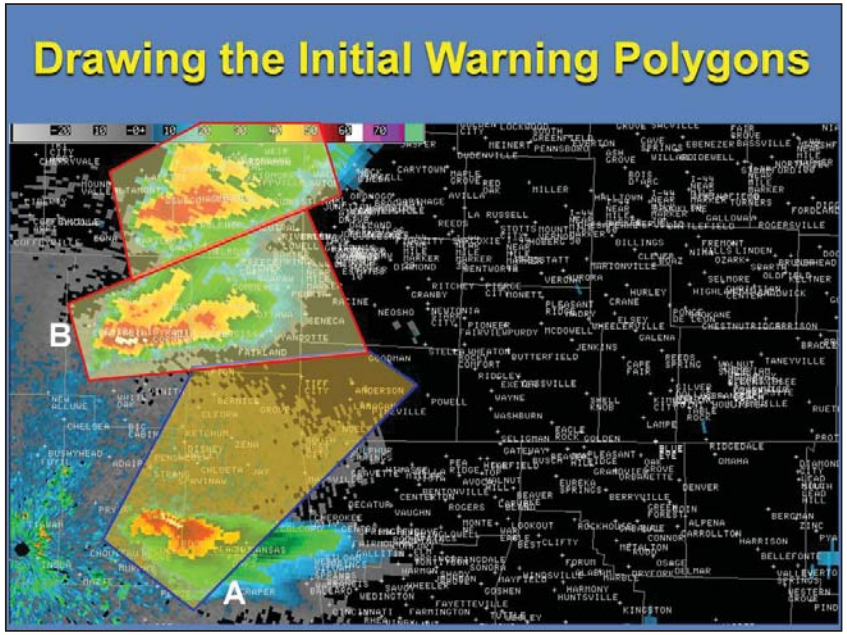


This brief lesson will illustrate some of the issues involved with drawing storm-based warning polygons for merging storms. The lesson will consist of an example of two storms whose individual motions will bring them to roughly the same location at the same time. There will also be a summary slide with the key points to walk away with from this example.



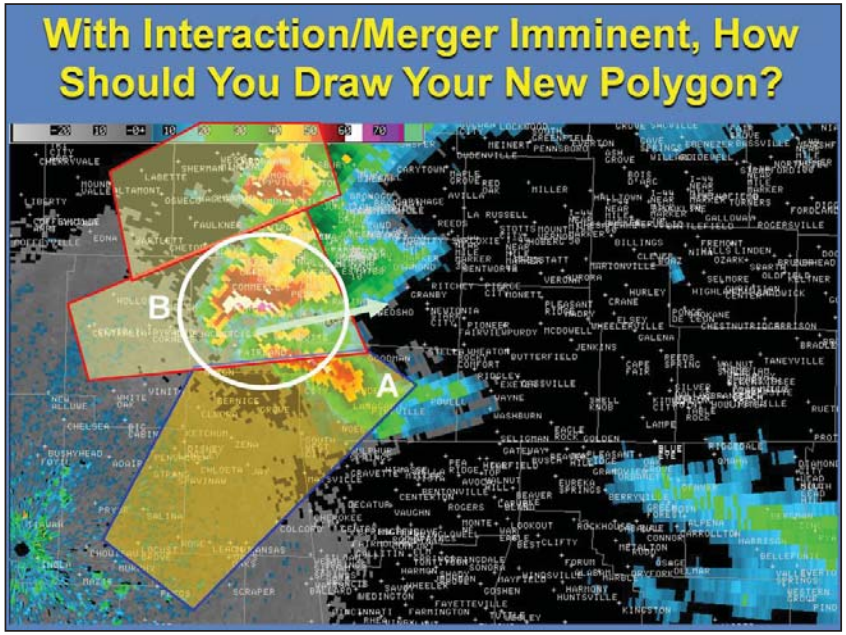
For this example, we will present a case where two storms will likely interact, or merge, near the end of the current warning time period. [A & B labels appear] The discussion will focus on storms A and B. [watch box becomes visible] Both storms are in an area covered by a PDS Tornado Watch with MLCAPE around 3000 J/kg and strong deep (0-6 km shear ~ 40-50 kts) and low-level (0-3 km SRH ~ 300 m2s-2) shear. [watch box disappears] …brief pause...

[storm A circle and arrow appear] Storm A is a strong left-moving supercell that formed during a storm split further to the south (green arrows represent storm motion vectors). [storm B circle and arrow appear] Storm B is an even stronger right-moving supercell. Your office has received several severe hail reports from Storm A. With little observed change in storm structure, the warning forecaster has decided to reissue a SVR for this storm. Storm B has had persistent, deep rotation for several volume scans but no reported tornadoes. Nonetheless, the warning forecaster thinks that reissuing a TOR for this storm is the best course of action. [circles move along arrow to show future storm location] By putting the circles in motion, we can see the approximate area where the storms should be in the next 30 minutes or so.

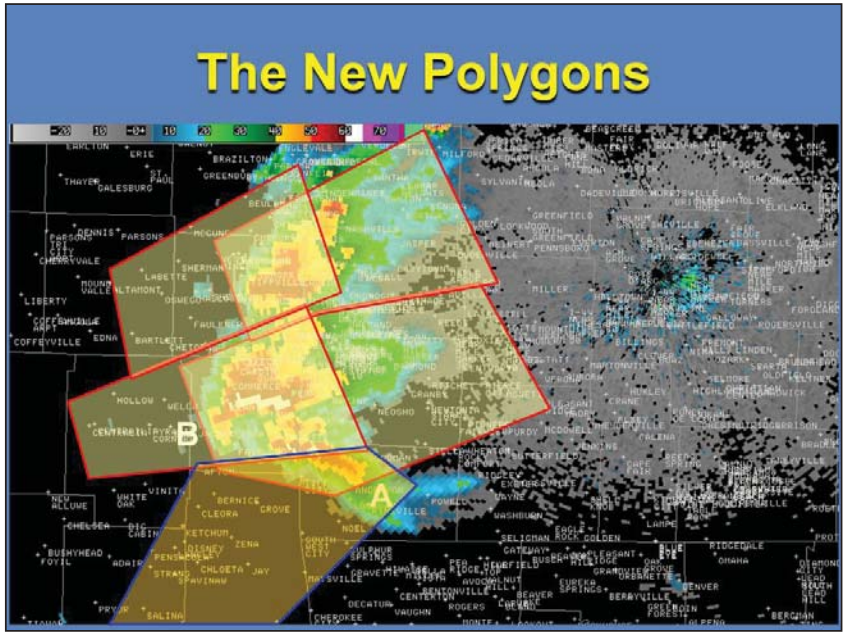


The primary issue here is how to draw the polygons so that all threatened areas are covered with a minimal overlap to limit potential confusion. [warning polygons appear] This case is a little more clear cut than most. The TOR (with Storm B) is the more significant threat, so the SVR (with Storm A) was drawn to cut off at the TOR boundary. To make this solution work best, you would want to either have (a) two separate warning forecasters coordinating their warning polygons or (b) the TOR for Storm B is issued prior to the SVR for Storm A.

Let's move ahead 30 minutes or so and see what to do with the next warning decision.

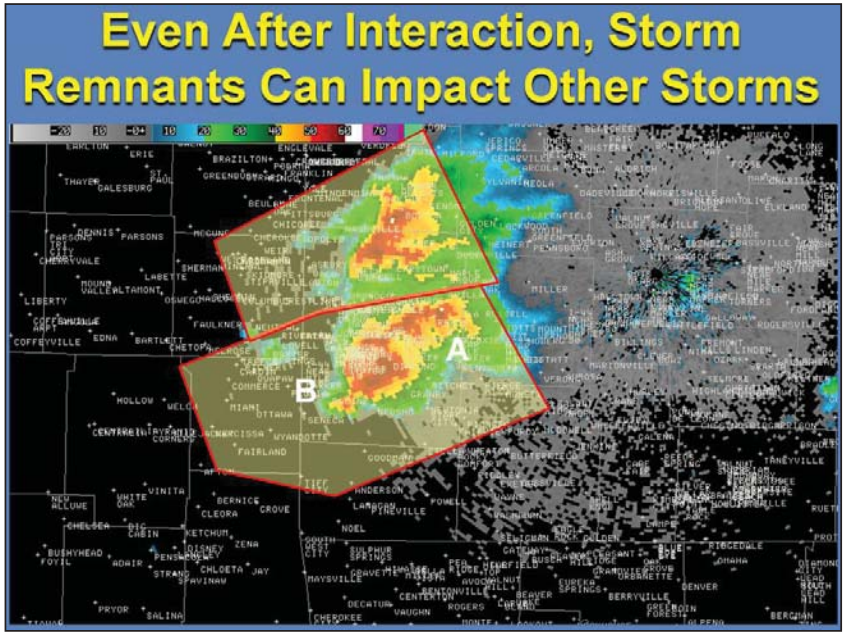


With the warnings due to expire, you can see that the storm interaction/merger is imminent. [warning polygons fade] During the last half-hour or so, Storm A has weakened slightly while Storm B appears to have intensified. Your office staff has had some good, vigorous discussion and are not sure if these storms will actually merge. The consensus opinion is that Storm A will likely continue to weaken or dissipate as the storms interact. [storm B circle and arrow appear] Taking this into account, the warning forecaster has decided to issue a single TOR following the general path of Storm B. [storm B circle moves forward in time] Once again, we will move the circle forward in time to see where our storm should be, approximately, in the next half-hour or so.



To display the next polygons downstream, we need to shift the image a little. [old polygons fade to dashed lines] We will leave an outline of the old warnings up to provide a frame of reference. In addition to moving the image, the storms are a little closer to the next radar downstream, so we have switched views to that radar.

[new polygons appear] Here is the subsequent warning for Storm's A and B. The polygon has been drawn a little larger due to the uncertainty issue of the merger/interaction. At the same time the polygon was drawn such that the warning for the storm to the north could be drawn with as little overlap as possible (but no unintentionally unwarned areas).



Here is an image of Storm's A and B about another 30 minutes later. Storm A has continued to weaken, but appears to remain an independent storm from Storm B. In fact, it appears that Storm A may merge/interact with the storm to the north. Although we could continue, I think you have seen enough to get the point of this example. Let's move on to the summary.

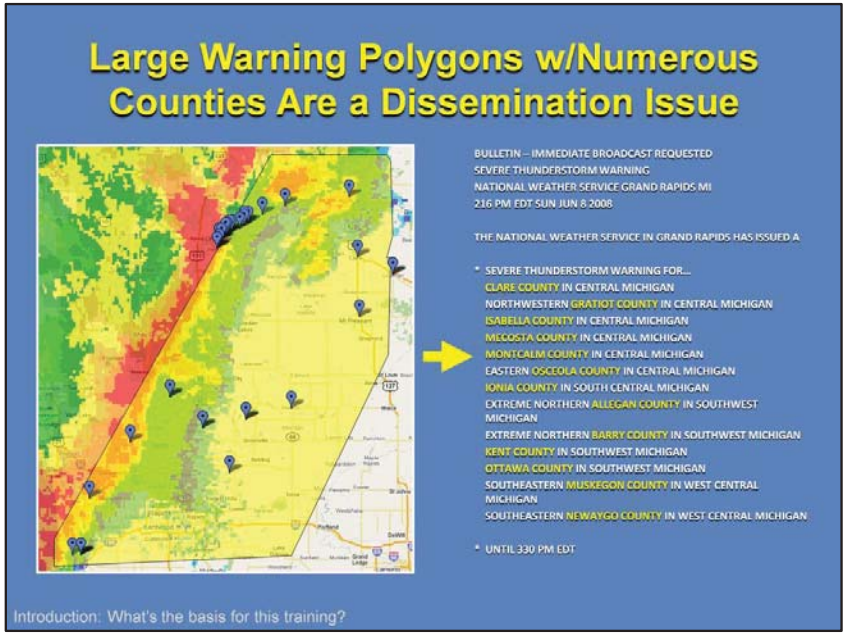
Summary

- Minimize areas of overlap
- In areas of significant overlap, the warning polygon for more significant threat should cover area
- Increase polygon size, as necessary, for uncertainty issues

In summary, here are the key points from this example. [first bullet appears] Just like other situations where you have multiple storms, you want to minimize any areas of overlap between different warning polygons. [second bullet appears] In areas where there would be significant overlap between the polygons due to interaction or merger of the storms, you will need to cut back one or more of the polygons to minimize the overlap. In these cases, the polygon for the more significant threat should cover the overlap area. [third bullet appears] Lastly, remember to take into account uncertainty issues with these situations by increasing the polygon size as needed, especially for warnings that cover post storm merger, or interaction, evolution.

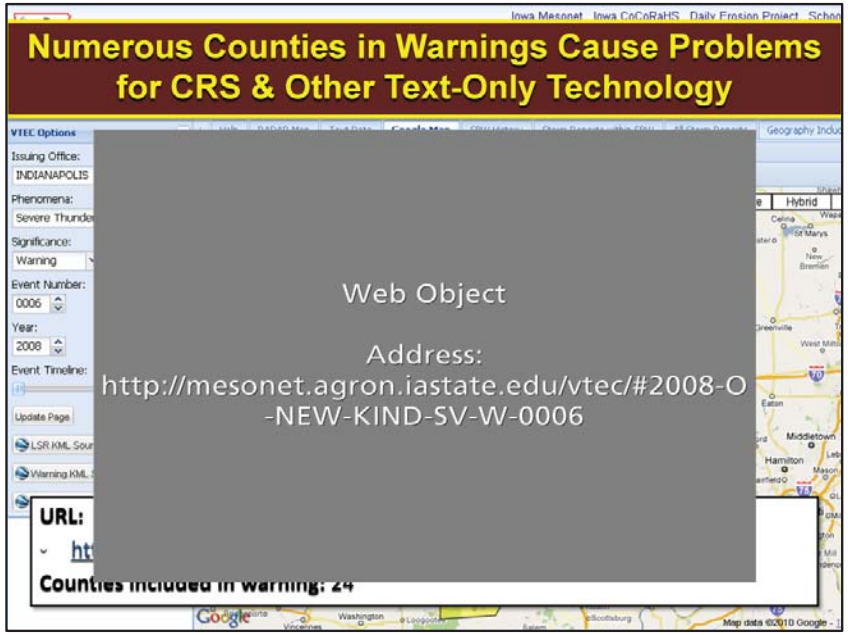


Welcome to the RAC Storm-Based Warning Fundamentals lesson on limiting the number of counties in warnings. This lesson addresses the NWS policy on the size of warning polygons. Lesson addresses when a convective threat area is too large for a single warning polygon. Then there will be a discussion on guidelines for dividing a phenomenon into multiple threat areas for warning polygons. This lesson should take approximately 30 minutes to complete.

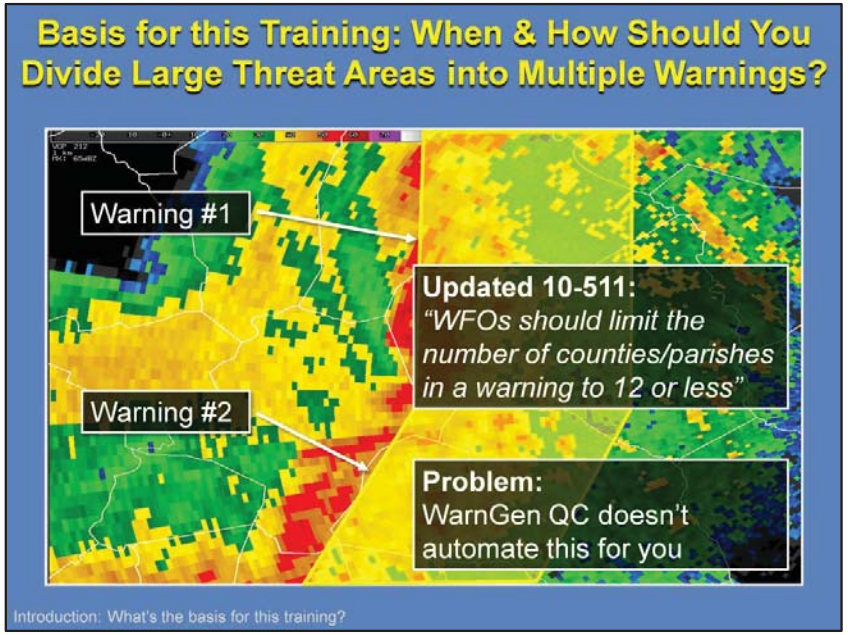


So why are large warning polygons a problem? Large polygons often include lots of counties and result in longer text products. That can lead to dissemination issues. *[show text & arrow]* In the example shown, the warning polygon included parts of 13 counties. The warning text would take any reading technology (e.g., CRS, human readers, and TV text scrolls) a long time to communicate.

From the graphic, you can see the office received reports throughout the warned area. So, the locations included in the warning indicated the potential threat area well. However, multiple warning polygons would have better communicated the threat.

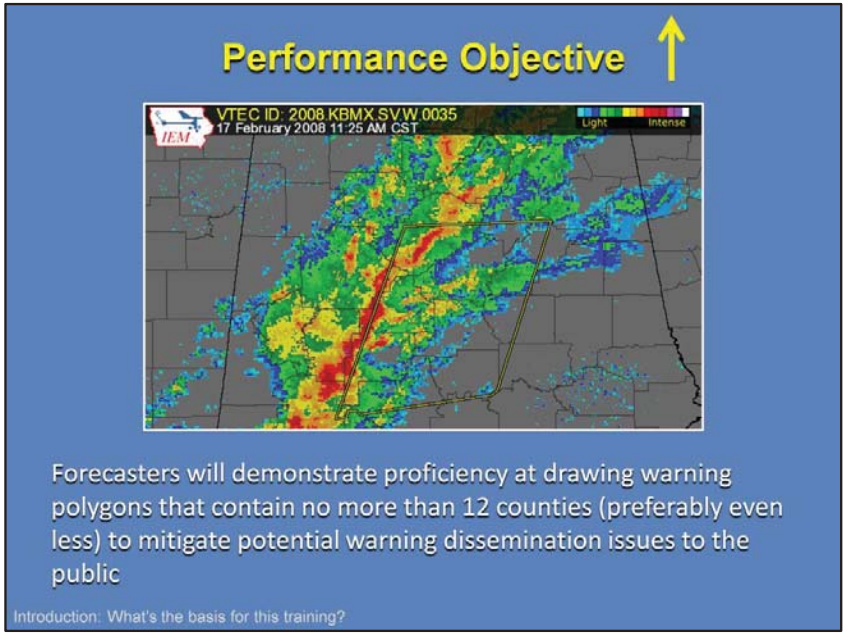


In the external browser window, you'll see another example of a large warning polygon. The display uses the Iowa Environmental Mesonet VTEC web page to better see all of the warning details. If you are not familiar with this web page, I highly recommend it for reviewing warnings as part of any post-mortem process. *[show arrow]* Besides showing the warning polygon, this site also gives you the opportunity to view the warning text, *[move arrow]* overlay storm reports, and review other data associated with the warning.

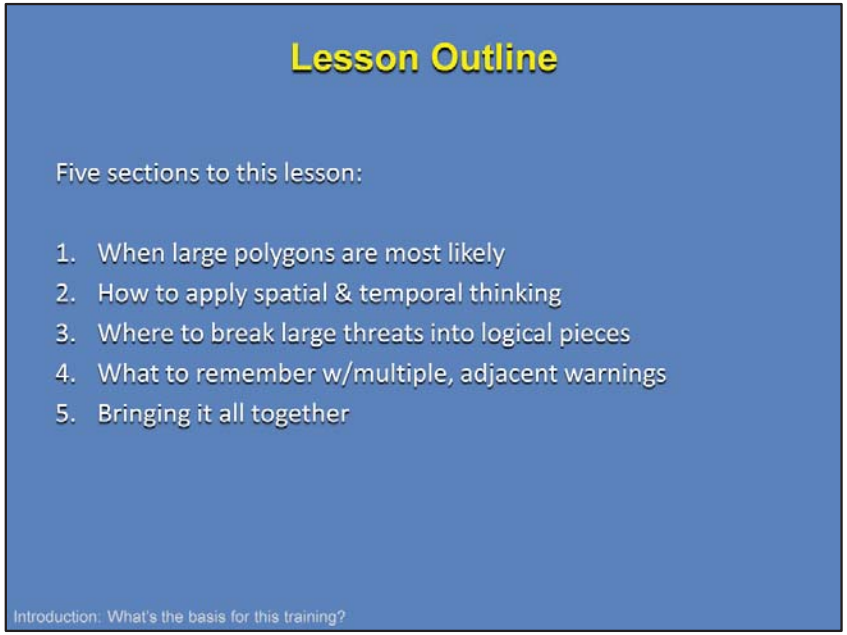


Hopefully, it's clear to you how large warnings can lead to dissemination problems? To prevent future issues, NWS Instruction 10-511 was updated in April 2010 to limit the size of Severe Thunderstorm or Tornado Warning polygons. *[show 1st text box]* Specifically, these warning polygons should contain 12 counties or less. *[show 2nd text box]* Unfortunately, WarnGen doesn't quality control warning text for the number of counties. Forecasters need to be diligent when creating warning products to ensure they meet this guideline.

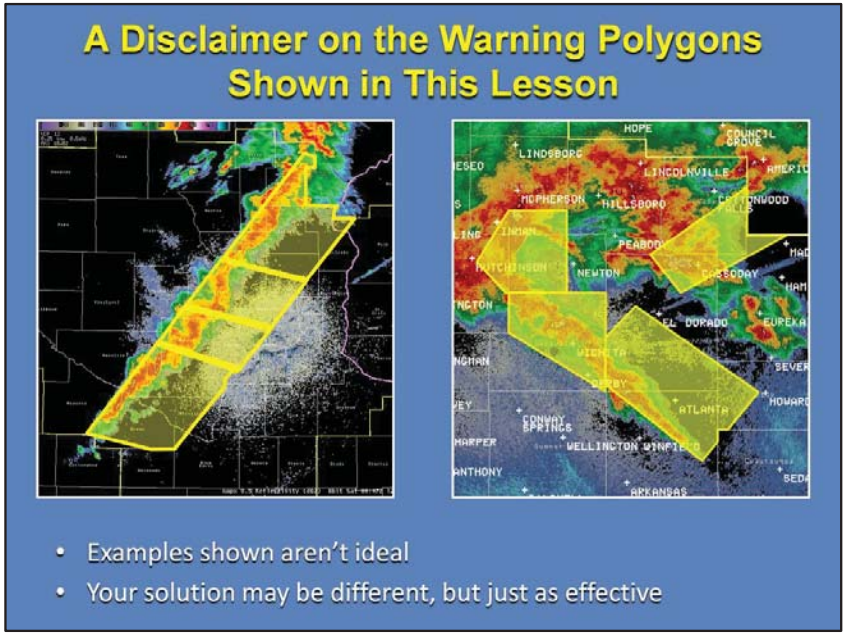
So now that you know when you should break up larger threat areas, the next question is how? That's what the rest of this lesson will discuss.

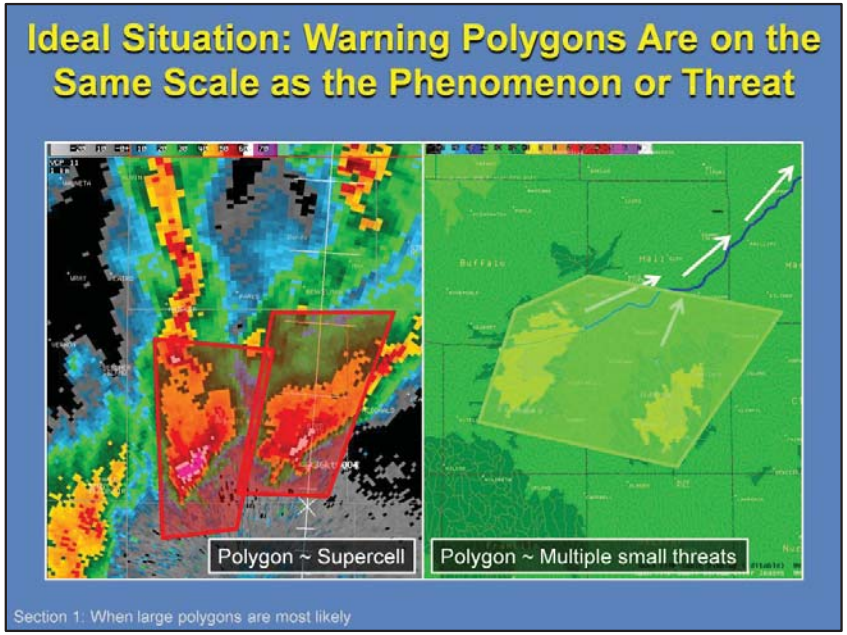


Take a moment to read the performance objective shown on the slide. *[show arrow]* You can also review the learning objectives for this lesson by clicking the tab in the upper-right corner of the window. These objectives, which are accessible throughout the presentation, will be covered by the quiz questions in the LMS upon completion of this course. *[hide arrow]* Once you have reviewed the objective, click the "Next Slide" button to proceed with this lesson.



The remainder of this lesson consists of five sections. *[show 1st bullet]* First, we'll cover when large warning polygons are most likely. *[show 2nd bullet]* Next, we'll discuss how spatial & temporal factors influence the size of a warning polygon. *[show 3rd bullet]* Then, we'll talk about breaking up large threats into logical pieces. *[show 4th bullet]* The fourth section will cover some important fundamentals about drawing adjacent warnings. *[show 5th bullet]* Lastly, I'll you some quiz questions to help bring all of these topics together.



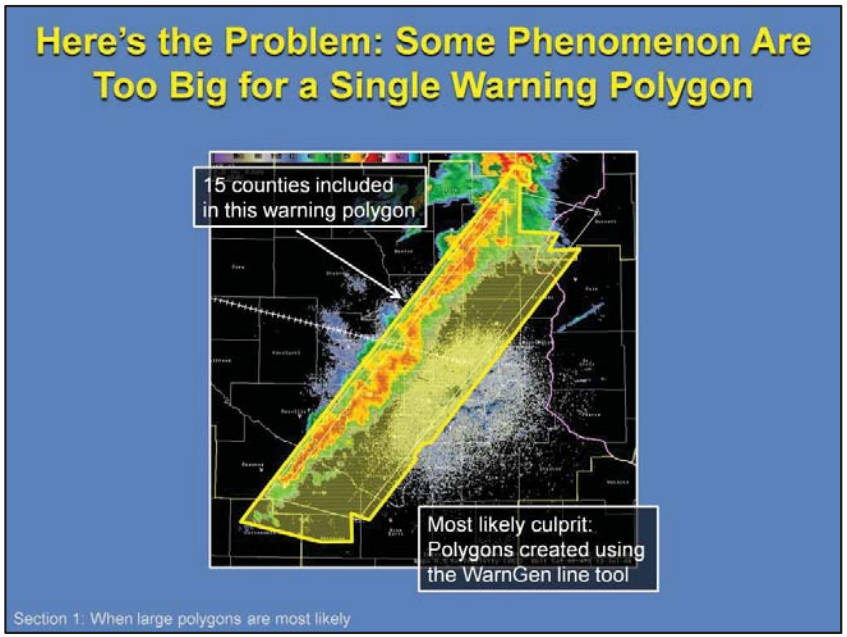


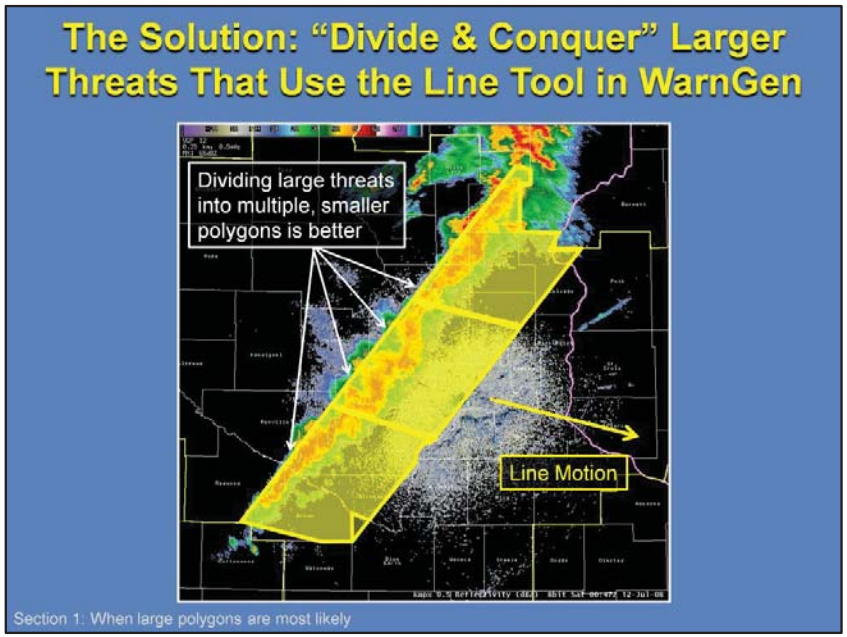
As you go through this lesson, you will see many warning polygon examples. Some of these examples involve actual warning polygons while other polygons were drawn by the instructor to support the goal of this lesson. *[show 1st bullet]* Regardless of the source, the warnings shown here shouldn't be interpreted as the ideal solution. They show just one way to break up a large threat area into multiple polygons. *[show 2nd bullet]* You might draw polygons that are different, but just as effective, as the ones shown here.

Ideally, warning polygons will be roughly the same size as the threat they cover. Actually, they should be a little larger than that to account for various forms of data uncertainty. But, they would be on the same relative scale.

In reality, this process is easier said than done. *[show TOR polygons]* For instance, Tornado Warnings are generally drawn on the same scale as the tornadic supercell. If I ask 10 forecasters to draw a warning polygon for a tornadic supercell, the results will likely vary somewhat. However, the polygons' areas should be fairly similar.

[show FFW polygon] Flash Flood Warnings, on the other hand, are basin-based phenomenon. Due to forecaster and data uncertainties, these warning polygons are often larger in scale than the actual threat area. As in the example here, a single FFW polygon may cover several, separate threat areas that flow into the same drainage basin. When threat areas are small, that's OK.

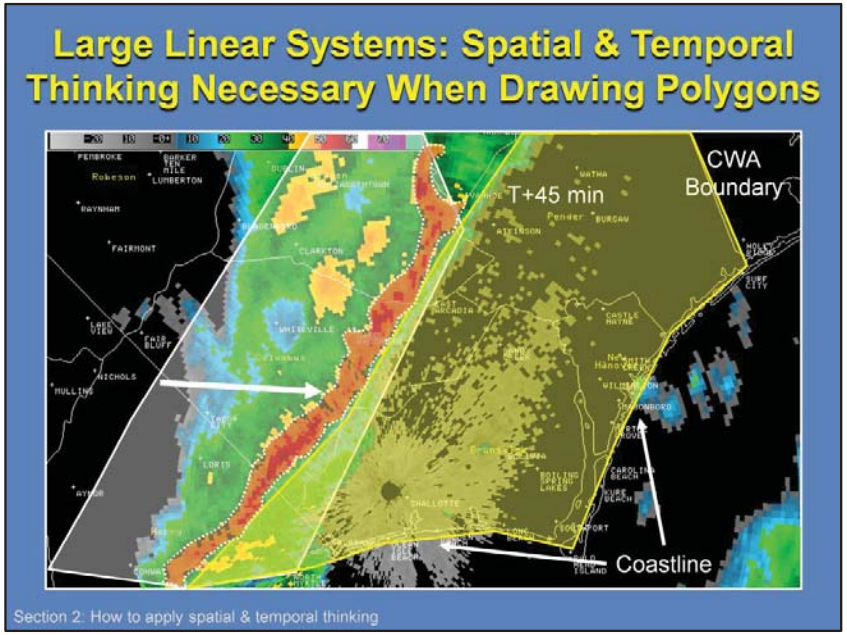


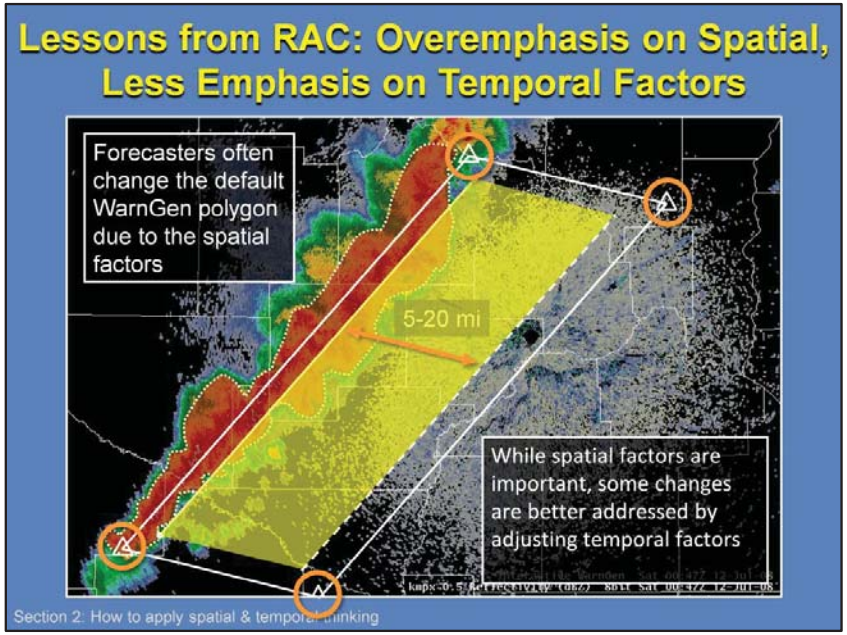


However, when the threat area is big due to a large meteorological phenomena, then warning size is a problem. Take this squall line, for example. *[show SVR polygon]* If you draw the polygon on the same scale as the phenomenon, it might look something like this. Since the squall line extends from one side of the CWA to the other, so does the warning. *[show 1st textbox]* The polygon shown here includes 15 counties.

[show 2nd textbox] From the examples I've shown, you may already realize the following: Having too many counties in a polygon is strongly linked to use of the "Line of Storms" tool in WarnGen. As a result, this lesson will focus primarily on its use in warning generation.

With that problem in mind, let's talk about the solution: Breaking up the phenomenon into multiple, smaller threat areas. What I'm advocating here is the old adage of divide and conquer. *[show smaller polygons]* For this example, the original SVR polygon might look something like this. The same general area is covered by these smaller polygons. *[show textbox]* The benefit of the smaller warning is that the dissemination issues seen with the larger warnings are less likely.





[show textbox] There are four primary factors impacting warning polygon size. Some have primarily spatial importance, while others are more temporal in importance. These factors are:

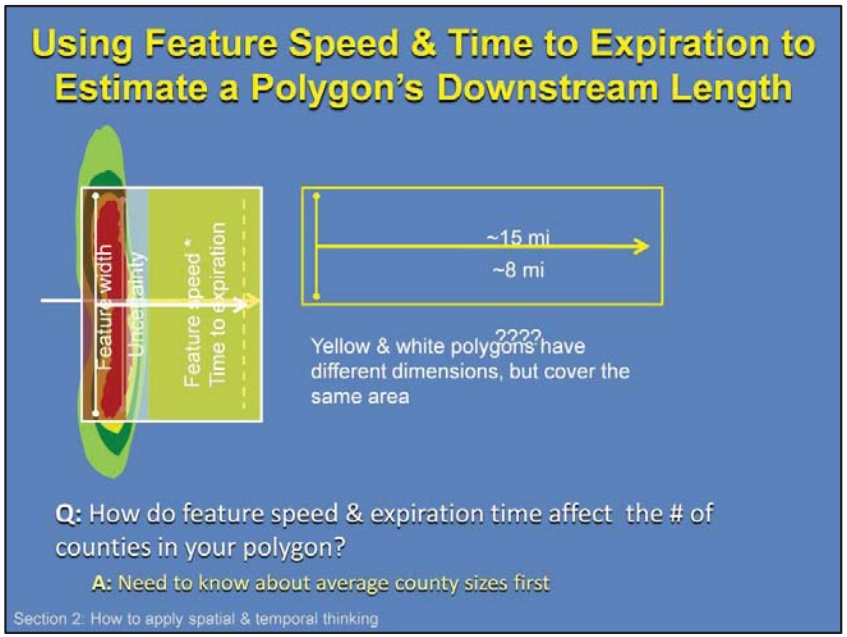
- The spatial area of the threat itself,
- The velocity of that threat,
- The time until expiration of the warning, and
- Any uncertainty the forecaster has about the threat over time.

[Hide text & show old warning] We'll use this example to illustrate. The polygon shown will expire in approximately 10 minutes. We'll reissue the SVR for this storm. *[show arrow & storm outline]* The line is moving in this general direction and, *[move storm outline]* in 45 minutes, the line should be about here. Most of the LSRs have resulted from the area along the leading edge of the storm. Storm motion and evolution has been pretty steady, so uncertainty is fairly low. *[show new warning]* We'll draw our new SVR polygon like so. *[show text & arrows]* In this case, the warning extends to the coastline & boundary with another CWA. So, those boundaries will limit the extent of the warning as well.

So how do these different factors impact warning size? Well, it depends on the situation. *[show textbox]* What I can tell you is an observation I've noticed over the years. Novice forecasters focus their warning efforts on manually adjusting for spatial factors much more than on temporal factors. *[show circles]* They often change the polygon's vertices to customize its shape when *[show arrow]* changing the storm motion vector or the expiration time would be better.

For example, if we use the WarnGen line tool for the outlined convection, then the polygon and lines shown are the default. *[show outline of squall line]* Now this squall line – depending on where you measure - is roughly 5-20 miles across. Using the larger value, that's about 40% of the polygon's downstream length. *[move outline]* Compare that to where WarnGen predicts the line will be when the warning expires based on storm motion. This distance is almost 40 miles.

Now, not all large threats will have those same proportions. *[show textbox & yellow area in polygon]* But often a large portion of the area covered by a warning is due to storm motion and time to expiration. Don't forget this info when customizing your warning polygon!



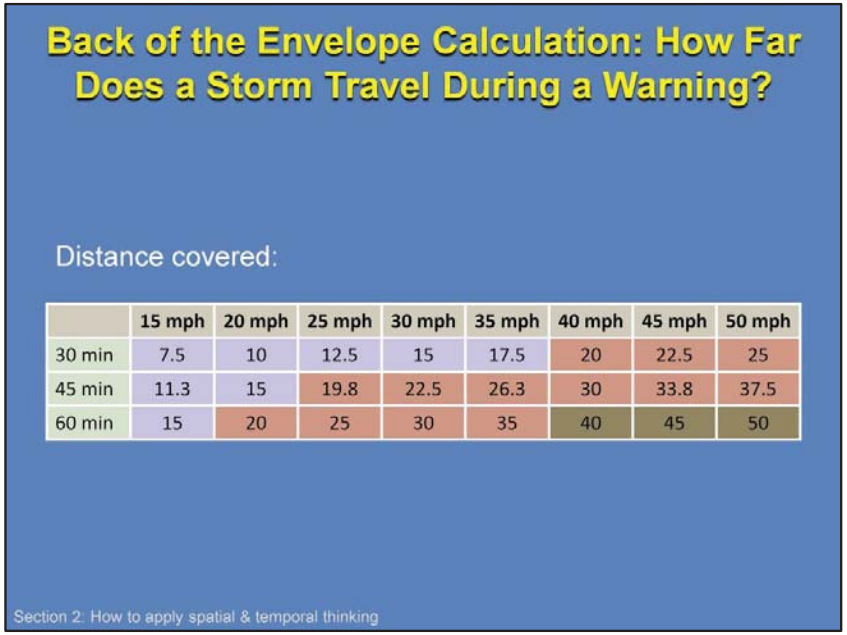


You've probably wondering where I'm going with all of this spatial and temporal factors stuff. *[show polygons & text]* Well, if you're polygon has a maximum size, then the cross-stream and downstream dimensions of your warnings are inversely linked. *[show arrows & lines]* The longer the downstream propagation vector, the shorter your line tool needs to be to meet the policy requirement. So, we've come up with some guidance in this area, but first let me show you how we came up with it.

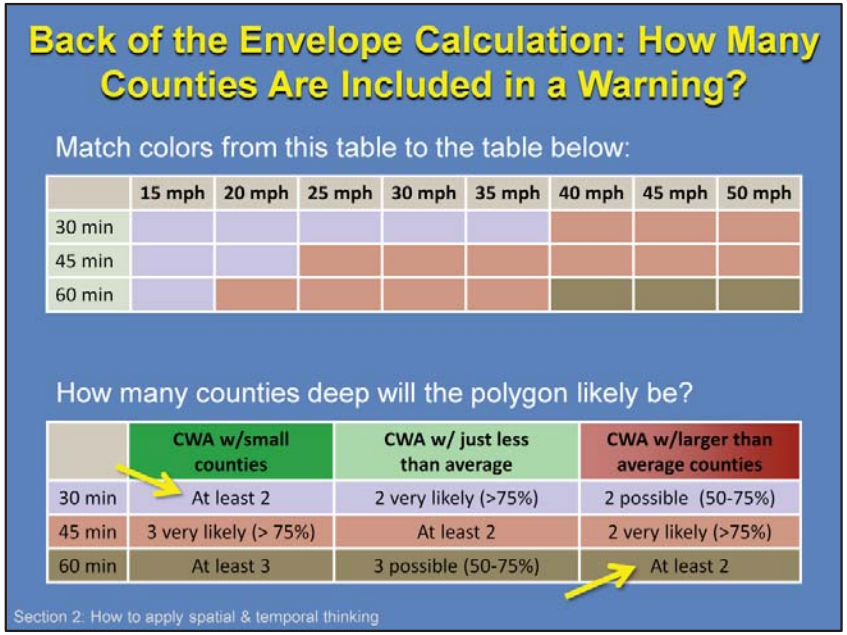
[show polygon] Say this polygon is a warning generated for this storm. *[fade in & drag polygon fills]* Breaking up the downstream distance into separate terms, it illustrates the role that the temporal components play. *[rotate fills & fade in text]* If we assign some "average" values to feature width and uncertainty, then it allows us to solve a problem: *[show Q & A text]* How do the feature speed & expiration time affect the number of counties in your warning polygon? Before we can solve this problem, first we need to know what the average county size is.

Do you know how big the counties are in your CWA? We used ArcGIS to figure that out for each CWA in the Contiguous 48 States and group them into four categories.

Use this interaction to learn more about each of the categories. Each category has an average for the county dimensions in that area. We'll use those dimensions for some back of the envelope calculations on the next slide.

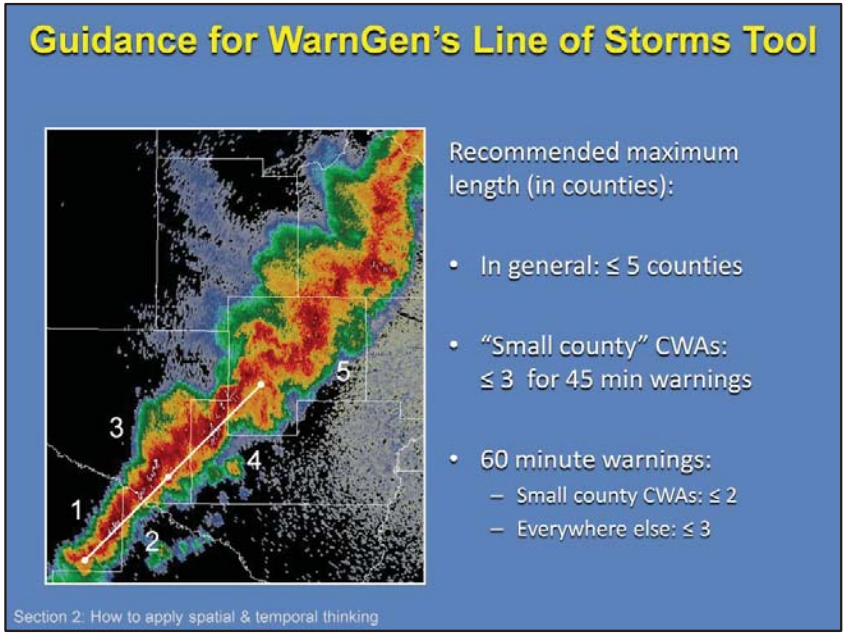


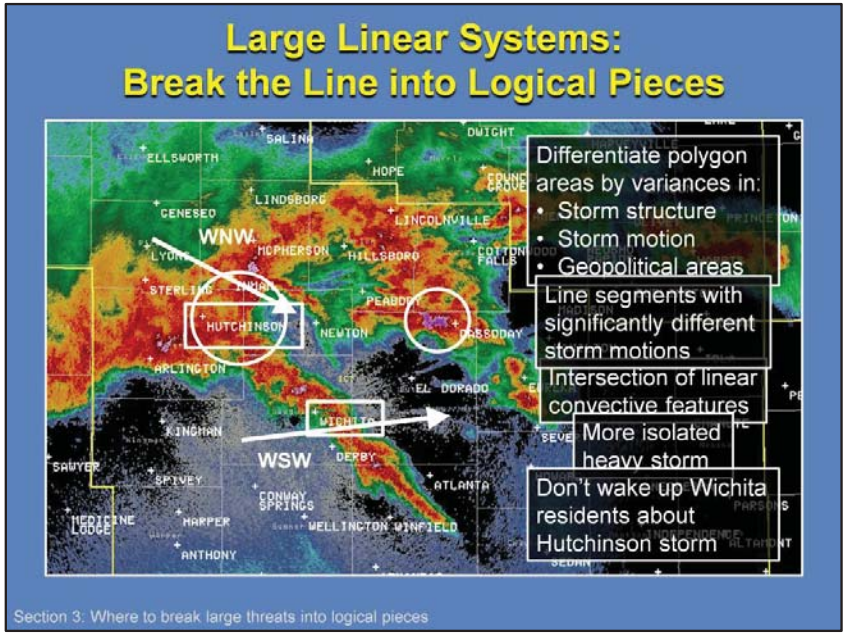
Next, we need to determine how far a storm travels during the warning period. The table on this slide shows these distances for a handful of speeds and expiration times. To simplify things, I grouped the values into three categories and color coded them. We'll discuss the relevance of the color-coded categories on the next slide.



The colors in the table near the top of the screen help indicate the downstream length of your warning polygon in counties. This number varies depending on the average size of counties in your CWA. *[show 1st arrow]* For instance, even at slow speeds & short durations, WFOs with small average county sizes often issue warnings that are at least two counties deep. Just to be clear, that's the county the storm is currently in and the next one downstream. *[show 2nd arrow]* On the other hand, if you have larger than average county sizes in your area, then it usually takes faster storms with longer warnings to include a comparable number of counties.

Based on these values, we'll discuss some basic guidance on the next slide.





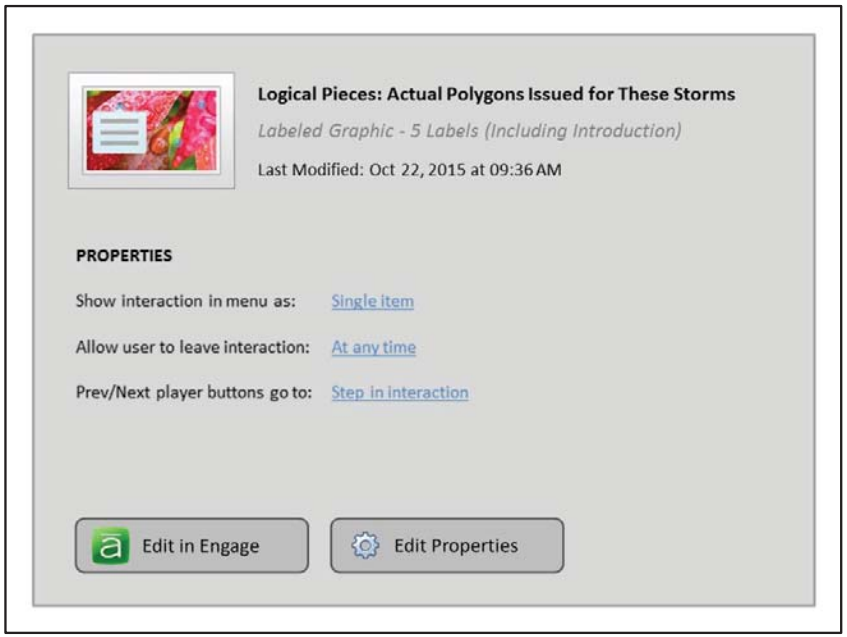
So, now that I've gone through this exercise, what does it all mean? *[show 1st bullet]* When using the Line of Storms tool in WarnGen, add up the counties that the tool touches. To ensure that your warning polygon contains less than the maximum in the directive, you generally want it to touch five counties or less. *[show 2nd bullet]* If you are issuing warnings for an area with "small counties", limit the number of counties the tool touches to three. At least for warnings of 45 minutes. *[show 3rd bullet]* If you are issuing a 60 minute warning, then:

• The tool should only cover two counties in a "small county" CWA, otherwise
• Limit the number of counties the tool touches to three or less.

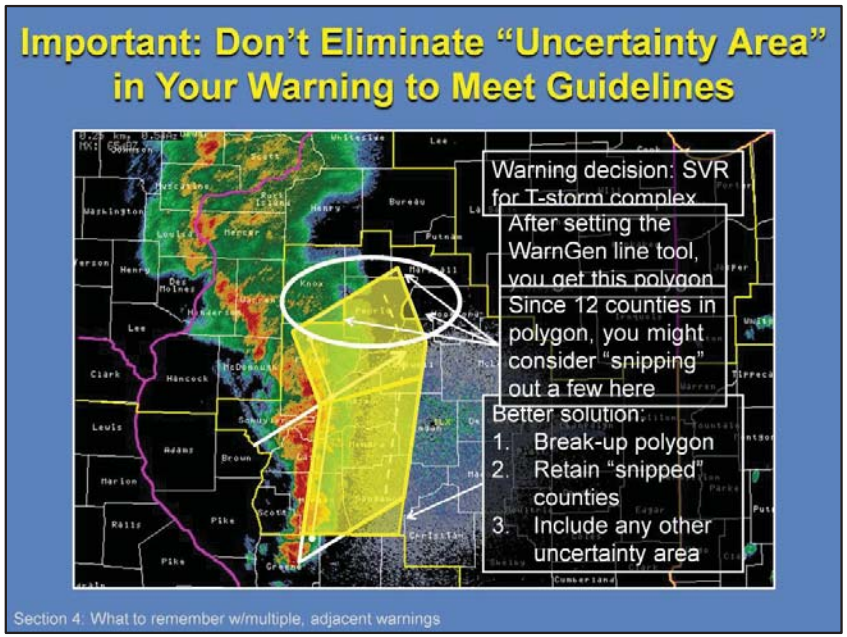
So, if you need to break up a large threat into separate threat areas, how do you do it? *[show 1st text box]* When identifying these threat areas, it's best to look for variations in storm structure, storm motion, and significant geopolitical features. Let's look at an example to show you what I mean.

In the case shown, there are two lines of convection that intersect near Hutchinson. *[show 2nd text box]* The SW-NE line is moving from the WNW while the NW-SE line is moving almost WSW. So, that's one logical break. *[show 3rd text box]* The intersection point between the two features will likely be a focal point of severe weather, so it should have a separate warning. *[show 4th text box]* While not a part of either of these lines, this isolated heavy storm near Cassoday will likely require a warning as well. *[show arrow]* Since the NW-SE line is moving in that general direction, you will likely want to draw your polygons for each threat appropriately.

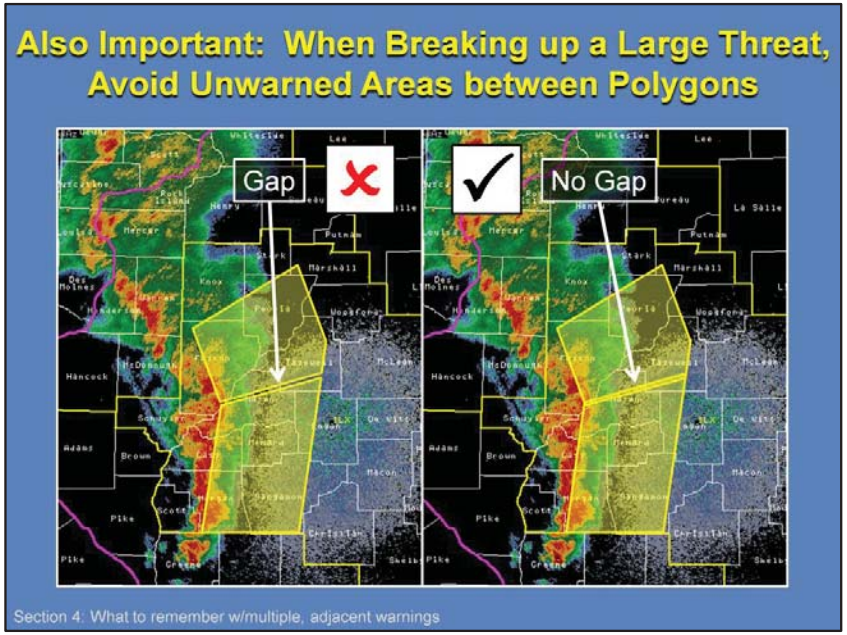
[show 5th text box] Lastly, we look at the area around Hutchinson and Wichita. These cities – about 25 miles apart - are the two largest in south-central KS. If you were to receive reports from one area, but not the other, you risk over stating the threat to the second city if you include them in one polygon.



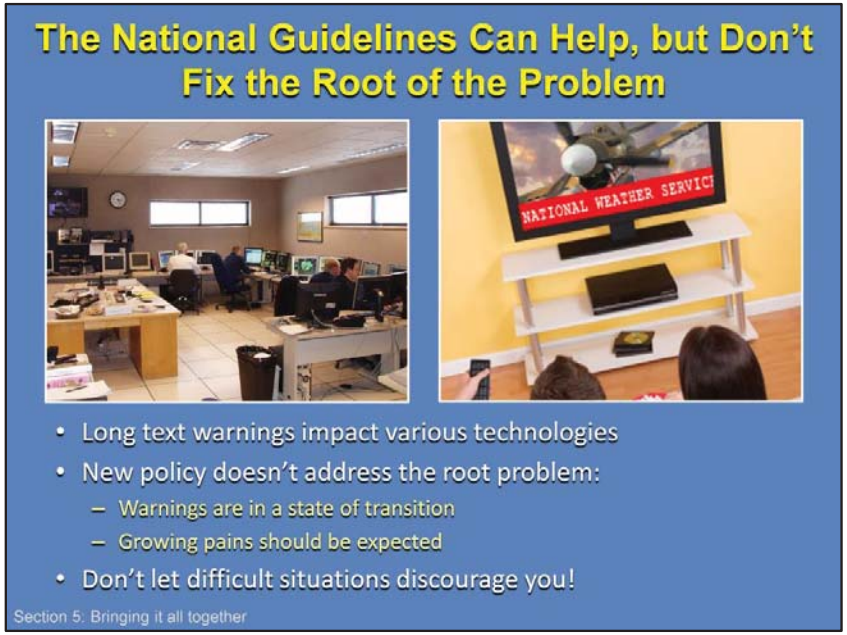
This interactive graphic shows how the Wichita WFO broke up their warnings for a 15 minute period for the example shown on the previous slide. This case was very complex. The polygons illustrate some issues that can come up when breaking up large threats into multiple warnings. Click on the arrows in the upper right-hand corner of the slide to step through details of each warning. You can also view a specific polygon's info by clicking on its red icon. Click on the next slide button to proceed with the lesson.



Let's say you are confronted with the following situation. *[show 1st text box]* You decide to issue a SVR for the thunderstorm complex shown on the slide. *[show 2nd text box]* After setting up the WarnGen Line of Storms tool, you get the polygon shown on the screen. *[show 3rd text box]* This polygon contains parts of 12 counties, so you might be tempted to manually remove some of the counties. *[show 4th text box]* If you find yourself in this situation, consider breaking up the polygon instead. In this case two separate polygons will do. *[show 2nd bullet]* Then, you can easily retain these areas in the warning that provide some room for uncertainty. *[show 3rd bullet]* Doing so covers any additional development that occurs along the southern edge of the line prior to the warning's expiration time.

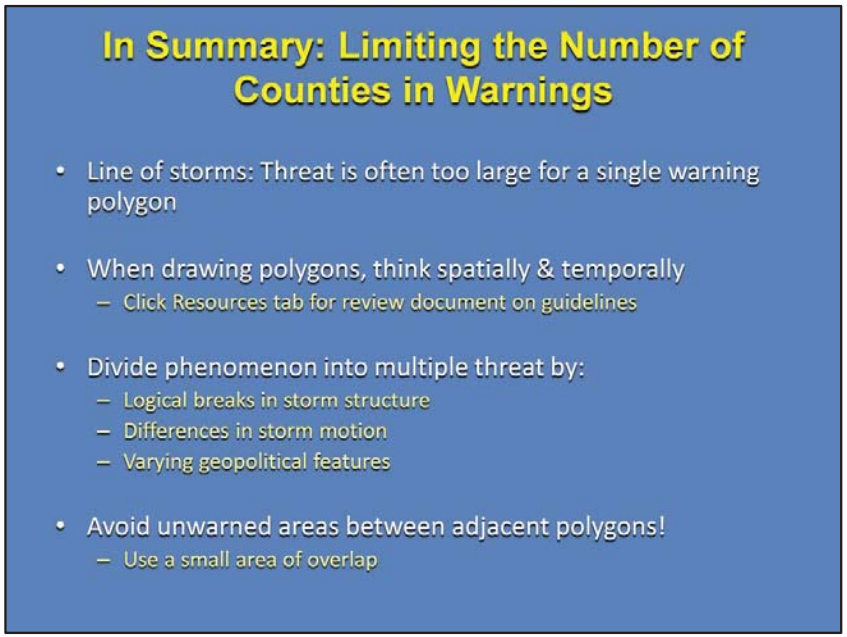


When drawing adjacent warning polygons, it's important to avoid unwarned areas between the polygons. *[show 1st set of polygons]* The best way to avoid this problem is to have a small area of overlap between the two warning polygons. *[show 2nd set of polygons]* This small overlap area will ensure that no area is let out of the warning and minimizes any confusion that may occur in the overlap area.



The 12 county maximum policy was written to help NWS forecasters provide good customer service. *[show 1st bullet]* Large polygons hamper technologies such as Specific Area Message Encoding, text crawls, and cell phones apps where users can customize NWS products to their local area. These same products can cause issues for older technologies like NOAA weather radio.

[show 2nd bullet] Even with this new policy, we haven't fully addressed the root problem. NWS warnings are in a state of transition. We need to support older tools such as NOAA weather radio while moving forward with newer technologies. Growing pains should be expected. Even the best warning solution may still seem awkward. *[show 3rd bullet]* Don't let these situations discourage you from providing the best service possible. Weather happens.



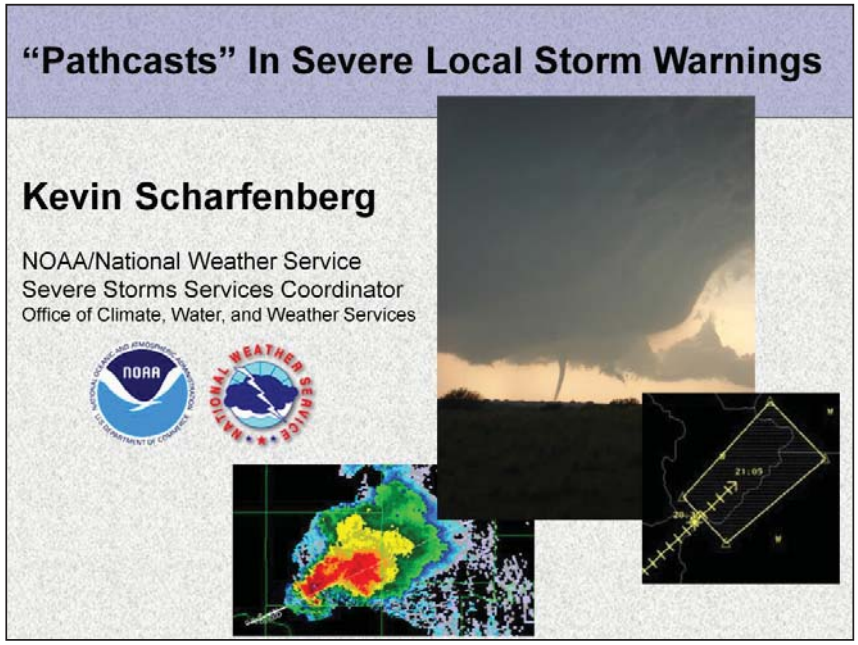
In summary, using the Line of Storms tool in WarnGen can result in warning polygons that exceed NWS national guidelines. *[show 1st bullet]* In these cases, the threat is simply too large for a single warning polygon.

[show 2nd bullet] When using the Line of Storms tool, it's important to think spatially and temporally. Novice forecasters often spend less time on the temporal factors on warning size even though they often result in the largest portion of the polygon's area. *[show sub-bullet]* Several guidelines were provided for using the Line of Storms tool depending on the average county sizes in your local area.

[show 3rd bullet] When dividing large phenomena into multiple threat areas, many times forecasters will intuitively know what to do. *[show sub-bullets]* When intuition fails, look for variations in storm structure, storm motion, and significant geopolitical features to identify logical breaks.

[show 4th bullet] Regardless of warning type, or how the warning is generated, adjacent warning polygons can result in some unwarned areas if forecasters aren't careful. *[show sub-bullet]* Avoid these unwarned areas by having a small area of overlap between the adjacent polygons.

Following these simple steps can help you know when multiple polygons are better than one for a specific threat.



Welcome to this informational seminar about the use of "Pathcasts" in severe local storm warnings. My name is Kevin Scharfenberg, and I'm working with the Office of Climate, Water, and Weather Services as NWS severe storms services coordinator. This presentation should last about 25 minutes.

Outline

Objective:

- Be able to describe the error sources inherent to the creation of "pathcasts" in severe local storm warnings

- Verification/validation studies of "Pathcasts"

- Operational "best practice" recommendations

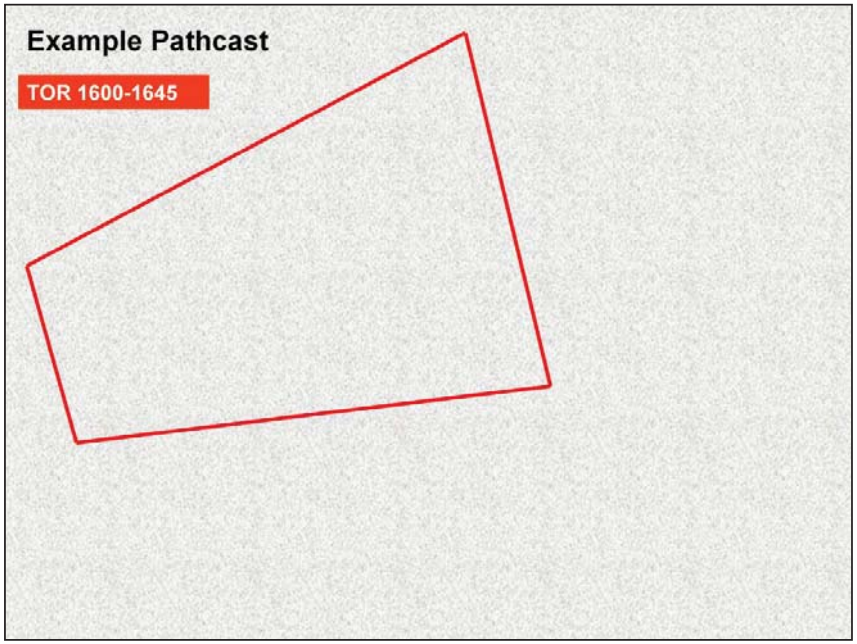
The objective of this course is to be able to describe the error sources a forecaster can expect to face when attempting to create a "pathcast" in a severe thunderstorm warning, tornado warning, or follow-up severe weather statement. First we will take a look at results from recent and ongoing research, and use those results to make some "best practice" recommendations for operations.

Error Sources in Pathcasts

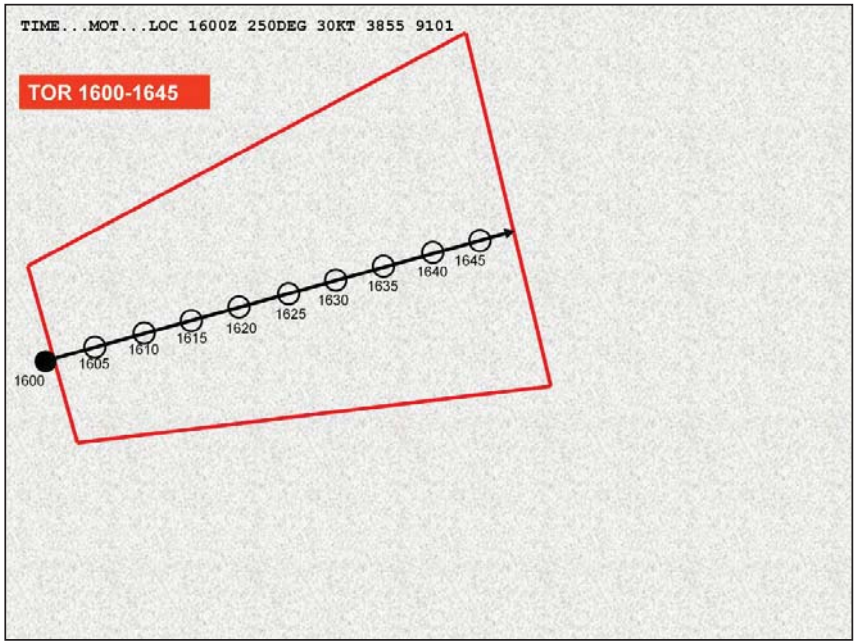
1. **Mapping in AWIPS**
2. **Radar**
3. **Storm**



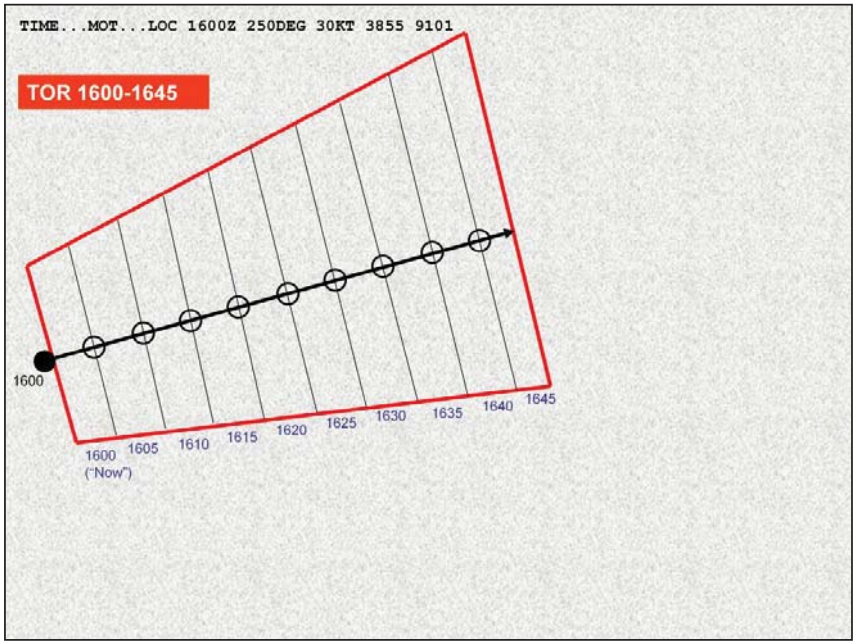
Studies have found there are three main error sources inherent to pathcasts. The first error source is in the way AWIPS handles maps. We have found that many cities, towns, and landmarks in AWIPS are simply in the wrong location. Some forecast offices have spent a great deal of time quality-controlling their location databases, and have had to make dozens of corrections. Even so, many locations are treated as a single point instead of areas, even some relatively large communities are treated as single points. We'll talk more about that in a minute. Finally, there are some implicit uncertainties in describing locations because AWIPS uses an 8-point compass, as well as 1-mile precision and 2-decimal place geocoordinate precision. The next major source of error is due to radar sampling. We know storms have some tilt with height, and in areas of poor low-level sampling that can turn out to be pretty significant. The strongest radar signature can be somewhat displaced from the area of greatest severe weather at the surface. Sometimes the radar data are not mapped absolutely perfectly, particularly far from the radar, and that can lead to small errors. The radar signatures may be mapped correctly, but beam-filling issues or messy and ambiguous signatures can lead to a lot of trouble figuring out where to put the "drag me to storm" dot. For example, a tornado is generally much smaller than the mesocyclone being sampled by radar, and it's not always in the center of the signature. Finally, storm processes can lead to significant errors when creating pathcasts. Warngen treats a pathcast as a linear extrapolation, but we know storms don't often work that way. Instead complex processes are often at work, including curving paths, occlusion of circulations, propagation, and so on. When these processes dominate, the skill of a pathcast can go out the window pretty quickly. Finally, we have to remember that warngen uses a single point to represent the storm, when really the severe weather threat we are trying to track can have a relatively large area associated with it.

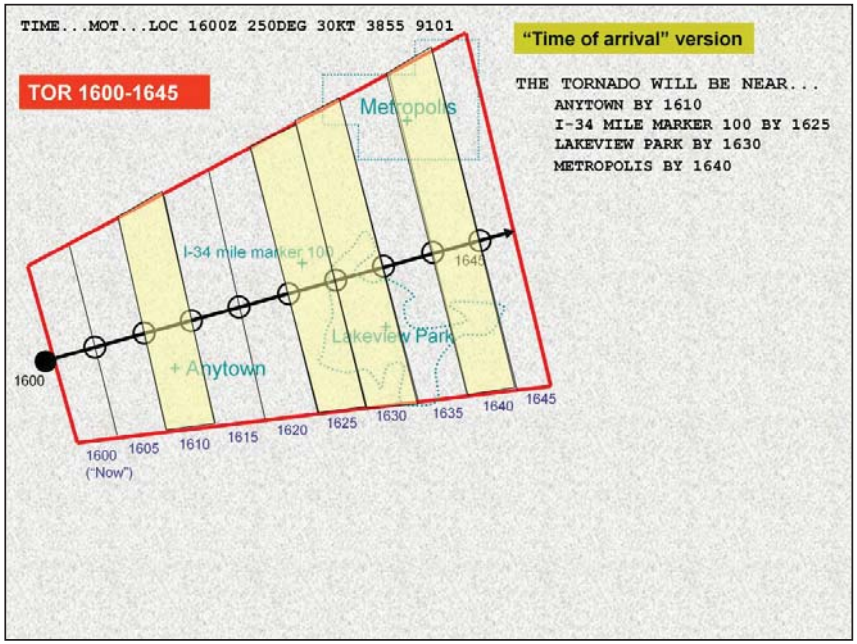


Before we go any further, let's stop and take a look at how a pathcast is created by the warngen software. Suppose we have a tornado warning polygon that looks like this, and let's say it's in effect from 1600 to 1645.



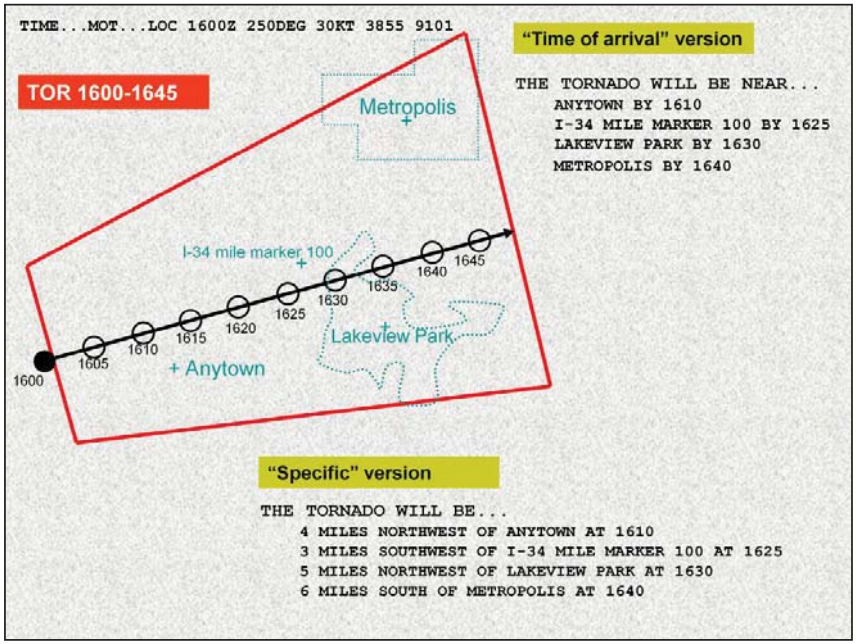
To create the warning, the forecaster drags the dot to a current storm location at 1600, then goes back a few volume scans and drags the dot back to an old location. This creates a linear extrapolated path, which goes out to the end of the warning time, in this case 45 minutes.

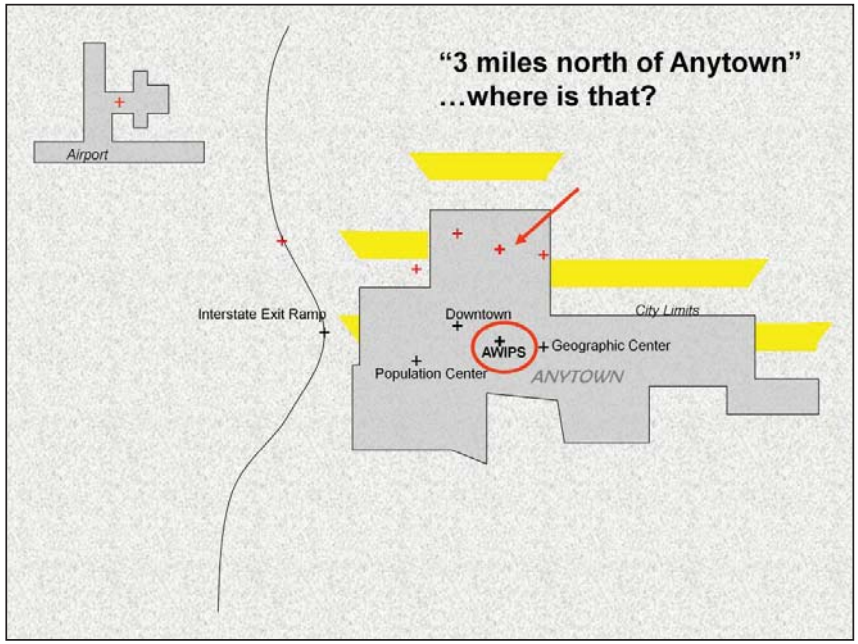




Next, the polygon is divided up into a bunch of skinny polygons normal to the path of the storm, corresponding to five minutes of storm motion. This is how the “time of arrival” is calculated. In the first upstream skinny polygon, the time of arrival is the current time, 1600. In the next skinny polygon downstream, it’s 1605, and so on, to the most downstream polygon which in this case is 1645.

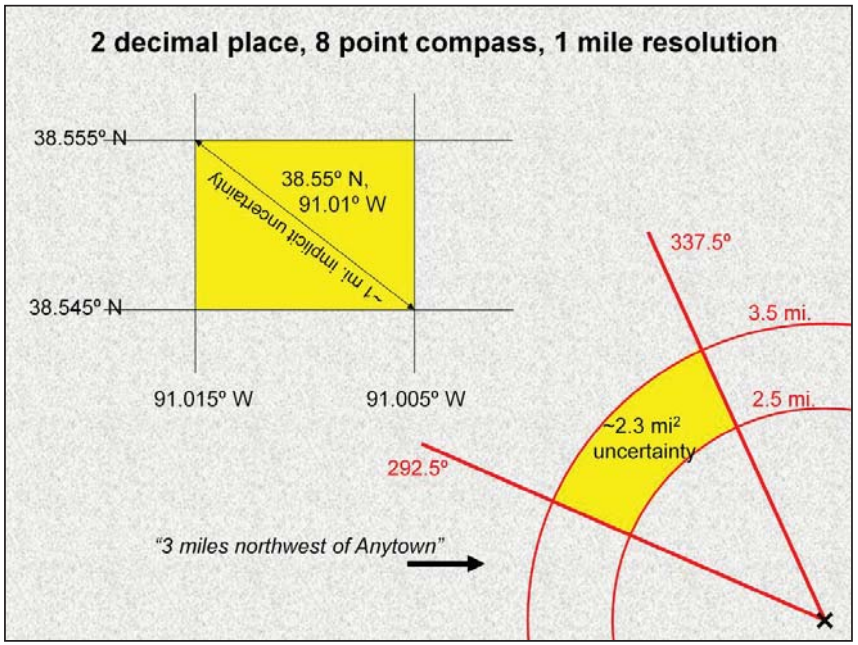
Now let’s see how locations are determined. For locations that are treated as single points, which is the majority of locations, the time of arrival is simply the time of the skinny polygon that location is inside. In this example, the time of arrival for “Anytown” will be listed as 1610. The same goes for any point location, which might include any sort of landmark programmed into the map database, such as an interstate mile marker. Notice that we’re treating metropolis and Lakeview Park as single points that get a single arrival time, even though those locations might have a large surface area. Warngen can be made to show a list of locations and times of arrival, and for the purposes of this presentation we will call this the “time of arrival” version of the pathcast.

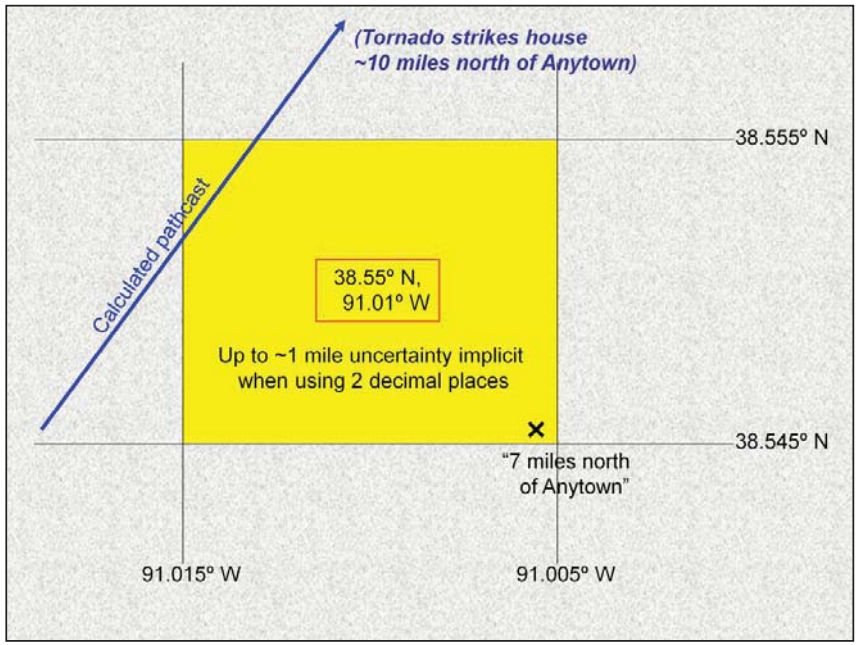




Depending on how warngen is configured, a "specific" version of the pathcast can be used instead. This version gives specific locations of the forecast storm centroid along with the times of arrival. In this example, warngen might output that the storm or tornado will be 4 miles northwest of Anytown at 1610, 3 miles southwest of the interstate mile marker at 1625, and so on. Next, we will discuss reasons why this version of the pathcast should not be used operationally.

When reading a "specific" pathcast, one question that might come to mind right away is: what exactly does "3 miles north of Anytown" mean, anyway? Now those of us with an AWIPS workstation can see where the point for "anytown" is located and measure out the location 3 miles to the north, but our end-users don't have access to the AWIPS location. They may think it refers to the population center, the geographic center, or the downtown. Some end-users might even think it refers to the location 3 miles north of the town's highway exit ramp, or 3 miles north of the airport because that's where the weather station is often located. Many others can perceive it to mean 3 miles north of the city limits, which could be any of the yellow area on this example map.



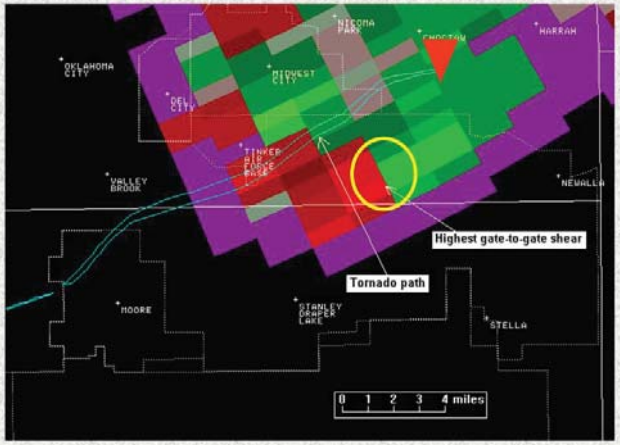


Another source of mapping error is due to the uncertainty associated with AWIPS precision in describing geocoordinates. Since AWIPS only uses 2 decimal places for latitude and longitude in warngen, we know we are not actually describing a point on earth's surface, but a two-dimensional area with four sides. In most of North America, that automatically implies about 1 mile of location uncertainty in the diagonal direction. Another source of mapping uncertainty is due to the 8-point compass and 1 mile precision used in warngen calculations. Because of this, the description "3 miles northwest of anytown" actually describes a possible area on earth's surface about 2.3 square miles in size. Obviously if we were to say "7 miles northwest of anytown" that describes a possible surface area that is even larger.

Here is an example of how these mapping errors can cause problems, even with a perfectly accurate pathcast. Let's suppose warngen calculates a tornado pathcast along the blue line and the forecaster chooses to include a "specific" pathcast in the warning. Warngen takes the projected tornado path to the coordinates 38.55 north, 91.01 west, and let's suppose that AWIPS interprets that location to be 7 miles north of Anytown. In reality, the pathcast clipped the northwest corner of that latitude-longitude box, and the actual AWIPS location of 7 miles north of Anytown is more toward the southeast corner of the same latitude-longitude box, about a mile away. Even if the tornado ended up moving right down the blue line, it could end up striking a house located more like ten miles north of town. You can see how the wording "7 miles north of anytown" can cause a lot of confusion in this case, even for a perfect pathcast, because of the way we handle maps.

Example radar location uncertainty – F4 tornado

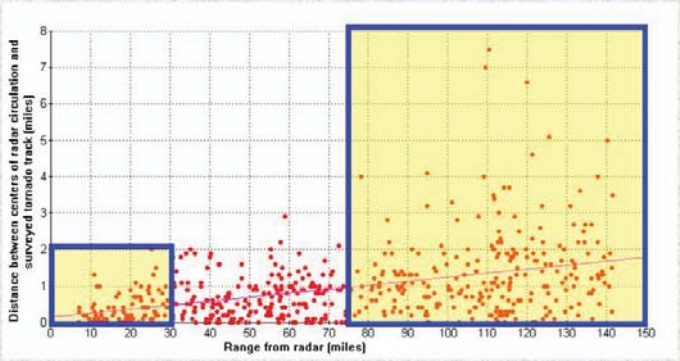
RDA 104 miles from storm -- From Speheger & Smith 2006



Next we'll move on to the second major error source associated with pathcasts, and that involves radar sampling. This image from a paper by Doug Speheger and Rick Smith is a good example of a case where the severe weather at the surface was displaced from the best radar signature. This storm was producing an F4 tornado at the time of the image, at a range of approximately 104 miles from the RDA. The tornado damage path is within the blue contours, and the tornado at this time was located near the tip of the error marked "tornado path". From the radar's perspective, however, the highest gate-to-gate shear was displaced several miles to the southeast of the tornado. In this case the mesocyclone was fairly large, there was storm tilt with height, and the tornado was occluded back in the northwest part of the circulation. You can see how creating a warning with a "specific" pathcast at this moment might create warning text that says the tornado will pass well south of the town of Harrah, when in reality the ongoing violent tornado is aimed right for Harrah.

Radar sampling limitations

Distance between radar circulation center and assoc. tornado damage



Speheger and Smith (2006) repeated this process for a lot of tornadoes, measuring the distance between the location of the tornado and the location of the strongest gate-to-gate shear. They found that even with tornadoes happening within 30 miles of the radar, the best shear signature is often displaced one to two miles away from the actual surface tornado damage location. For events farther from the radar, beyond 75 miles, the error can range from 3 miles to 6 miles or even more.

Example linear extrapolation error

"Tornado will be 3 miles west of Tuttle at 2345Z"



Error Sources in Pathcasts

Mapping in AWIPS - observed to be 3+ miles in some cases
- City/town/landmark location error
- Treating most locations as single points instead of areas
- 8-point compass, 1 mile, 2 decimal place lat/lon uncertainty

Radar - observed to be 3+ miles in some cases
- Storm tilt / poor low-level sampling
- Radar data mapping errors
- Radar beam-filling issues (large mesocyclone, small tornado)
- Messy/ambiguous signatures

Storm - observed to be 3+ miles in some cases
- Cyclic occlusion processes, curving paths, acceleration (deceleration)
- Discrete propagation, boundary collisions, etc.
- The storm is not a single "point"

→Preliminary study results: errors of 5-10 miles/minutes are frequently observed in Pathcasts!

Our third and final major source of pathcast error is due to non-linear storm processes. In this example, the storm is about 50 miles from the nearest radar and showing very strong and clear signatures. Suppose we're issuing a tornado warning for a storm at 2255 UTC located near the southwest part of this map. In Warngen we would drag the dot to the 2255 UTC location of strongest shear, then go back a couple of volume scans and drag the dot back to the 2245 UTC location. This would create a linear path represented by the solid red line, and extrapolated locations shown by open red circles at 5 minute intervals downstream. Note in this case that the signature was offset by about one mile to the northwest of the damage path, which is shown in purple. Let's suppose we ran our tornado warning out to 2345 UTC, and we used the "specific" version of the pathcast. This would probably cause warngen to create text that said "the tornado will be 3 miles west of Tuttle at 2345". In reality, just after we created the warning, the storm took a subtle turn to the right, at about an angle of 10 to 15 degrees. The resulting F5 tornado at 2345 UTC was actually located well south-southeast of Tuttle, about 8 miles from the original extrapolated path. These sorts of errors were observed many times in the verification of pathcasts issued by the NWS.

Looking at our detailed list of error sources in pathcasts, we have seen in ongoing verification and validation studies that each error source can be independently responsible for errors of 3 or more miles. In a worst case scenario, they may add up to 9 or 10 miles. As a matter of fact, we've found that errors of 5 miles or so are quite frequent in "specific" pathcasts, and 10 miles are not unheard of. For a storm moving at 60 miles per hour, that can correspond to errors of 5 to 10 minutes when forecasting the time of arrival.

Problems with Interpreting Pathcasts

The "skinny black line" issue:

- *Which is the warning: the pathcast or the polygon?*
- Invites users to focus on an exact storm path forecast?
- Minimizes the threat elsewhere in the warning polygon?

Communication barriers:

- Where is "6 miles south of Metropolis?"
- "Lakeview Park" is a big place!
- Am I safe until the time in the Pathcast? Safe afterwards?
- I'm south of Anytown and NWS says the tornado will be north of Anytown, so I must be safe

Can imply forecast precision we do not have!

Real-world examples from 2008

TOR from 510-615 pm: THE TORNADO WILL BE NEAR…

TOR 510pm: 8 MILES WEST OF TOWN A AROUND 535 PM CDT...
TOWN B AROUND 545 PM CDT...

SVS 515pm: TOWN A AROUND 535 PM CDT…
TOWN B AROUND 545 PM CDT…

SVS 519pm: TOWN A AROUND 545 PM CDT…
TOWN B AROUND 555 PM CDT…

SVS 534pm: TOWN A AROUND 555 PM CDT…
TOWN B AROUND 605 PM CDT…

SVS 550pm: TOWN A AROUND 610 PM CDT…
TOWN B AROUND 615 PM CDT…

(Result: Tornado passed 2-3 miles west of Town A at ~557pm)

In addition to the quantitative errors that can be associated with pathcasts, there is some concern about the end-user interpretations of pathcasts. First, in hurricane forecasts, there is something we call the "skinny black line" issue. The famous hurricane track forecast maps show both a thin black line that represents the forecaster's best track forecast of the eye of the storm, and also a "cone of uncertainty" representing the area the hurricane might track. We want all the people in the cone of uncertainty and warning to be taking action, regardless of the exact path forecast of the hurricane's eye. The same issue comes into play when pathcasts are issued in severe local storm warnings. There is a similar concern that including the path forecast invites users to focus on the skinny black line of the pathcast when really we want everyone in the warning polygon to be taking action. There is also a concern that communications barriers exist when trying to convey pathcast information in text and in audio broadcasts. For example, saying a tornado will be 3 miles north of Anytown might convey that people in the town itself and south of town are safe, even if they are in the polygon. Also, we don't have good information about how people use exact timing information included with some warnings. The important thing to remember is that by including exact times and locations we run the risk of implying forecast precision we do not have.

The following slides show a few examples of pathcasts issued operationally in 2008. In this example a tornado warning was issued at 5:10 pm effective until 6:15 pm. Note that's 65 minutes is considered too long for most tornado warnings, we recommend limiting the valid time of tornado warnings to 30 to 45 minutes. Now this supercell thunderstorm was about 60 to 70 miles from the radar with relatively strong signatures. The initial tornado warning state the tornado would pass 8 miles west of town A, but the follow-up severe weather statement 5 minutes later just said "Town A". Then the next SVS a few minutes later changed the times of arrival by 10 minutes. The next SVS 15 minutes later changed the times again, and finally the last SVS issued at 5:50 pm extended the times of arrival even more, so that over 40 minutes of warnings, four different times of arrival were stated for each town and they changed by a total of 30-35 minutes. It turns out a tornado passed a few miles west of town A at about 5:57 pm, so you can see that the last SVS issued was actually less skillful than the one issued at 5:34 pm. This example shows that even time of arrival information can be very unstable when going through the pathcast creation process several times in quick succession.

Real-world examples from 2008

SVS 1237pm: TOWN A BY 1255 PM CDT...
TOWN B AND 6 MILES NORTH OF TOWN C BY 100 PM CDT...
TOWN D BY 105 PM CDT...

SVS 1242pm: 8 MILES SOUTH OF TOWN A BY 1245 PM CDT...
TOWN C AND 8 MILES SOUTH OF TOWN B BY 1255 PM CDT...
8 MILES SOUTH OF TOWN D BY 100 PM CDT...

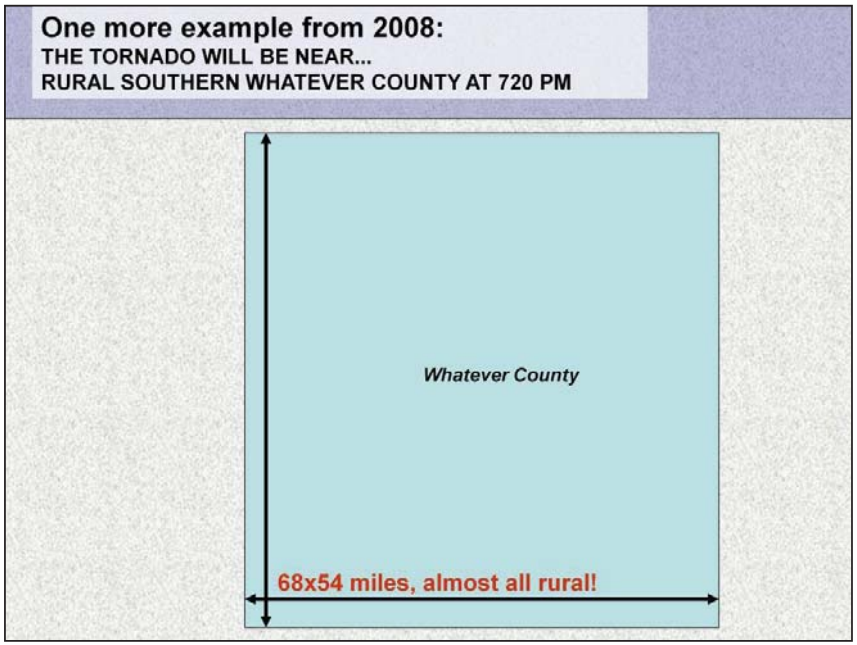
Here's another example from 2008, of a couple of severe weather statements issued 5 minutes apart as follow-up to a tornado warning. In this case, in just one volume scan all 4 locations and all 3 times changed radically. This again shows how unstable the results can be when repeating the pathcast process.

Real-world examples from 2008

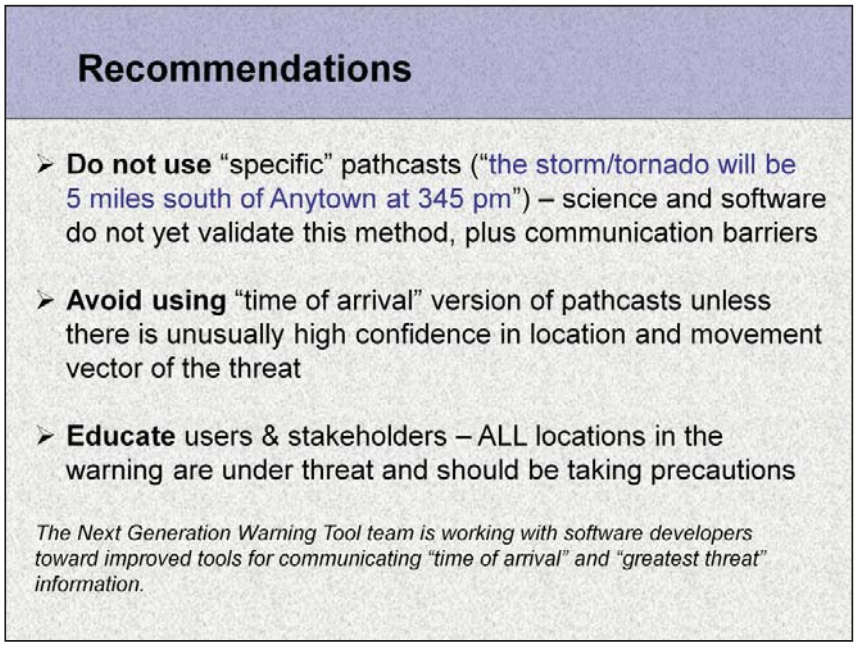
SVS 553pm:
Radar-indicated tornado located "NEAR TOWN A"
will be "8 MILES NORTHWEST OF TOWN B BY 615 PM."

Tornado was actually located ~6 miles southeast of Town A at the time and passed ~2 miles south of Town B at ~618 pm.

Here is another example severe weather statement in 2008 associated with a violent tornado event. The initial location in the basis statement of the warning said the tornado was near Town A, and would be 8 miles northwest of town B in 22 minutes. The forecaster didn't know it at the time but the tornado was ongoing at the time this statement was issued, but was actually 6 miles southeast of Town A. It passed just south of town B 25 minutes later. As it turns out the forecaster was using the "specific" version of the pathcast despite the fact the radar signatures were quite messy and the storm was far from the radar. This shows how errors can be large even at the initial time step, and can grow very quickly during the forecast period.



One final example from 2008. In this case the warngen-created wording in the pathcast was “the tornado will be near rural southern whatever county at 7:20 pm” but that makes little sense when you consider the county is 68 miles tall by 54 miles across, and almost entirely rural. In this case the software was asked to create pathcast wording but the polygon did not intercept any landmarks, so some rather ridiculous wording ended up being created and transmitted.



Based on this research, here are a few “best practice” recommendations for operations. First, do not use the specific version of pathcasts, such as stating in warnings the tornado will end up 5 miles south of Anytown at 345 pm. The science does not validate this approach is yet possible, and the software does not yet take into account the necessary uncertainty. In addition, there is great concern that end-users are not interpreting this sort of information properly. Second, it’s generally a good idea to avoid using “time of arrival” pathcast functionality, except in unusual situations where you are very confident about the location and movement vector of the threat, and nonlinear processes seem to be minor. But again, generally it’s best just to leave it out. Finally, and most importantly, we need to reinforce to users and stakeholders that ALL locations in the warning are considered to be threatened and everyone inside the polygon should be taking immediate precautionary action, regardless of any times and specific locations listed. It should be noted that the next generation warning tool team is working with software developers and social scientists to improve the software in AWIPS II so that we can perhaps have some better tools to convey time of arrival and greatest threat area information.



This concludes the informational seminar on pathcasts in severe local storm warnings. Thanks for taking a look. If you have any comments or questions please be in touch, again my name is Kevin Scharfenberg and my e-mail address is there on the screen.



Hello, my name is Jill Hardy and welcome to this lesson on Flash Flood Warnings. Already in this course we have discussed the precipitation estimation products available, fundamentals in flash flood meteorology and hydrology, and how to utilize FFMP and the precip sources.

But before you arrive at the workshop, it is vital that you understand some flash flood warning fundamentals. This lesson will focus on the polygonology and warning text fundamentals, in order to give you the last piece of the puzzle for great flash flood decision making.

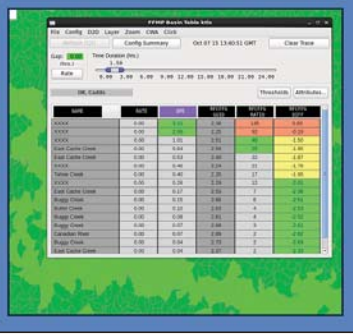
Learning Objectives

- By the end of this lesson, you will be able to:
  - Identify when to issue a Flash Flood Warning (FFW) versus Flood Advisory or Areal Flood Warning
  - Identify appropriate polygon sizes and FFW durations
  - Identify when to use the automated basin list in WarnGen
  - Identify appropriate basis and call-to-action details to include in a FFW
  - Identify how to follow-up a FFW
  - Identify the criteria for a Flash Flood Emergency

Here are the learning objectives for this lesson. When you have finished reading them, please continue to the next slide.

Defining the Threat Area Using FFMP

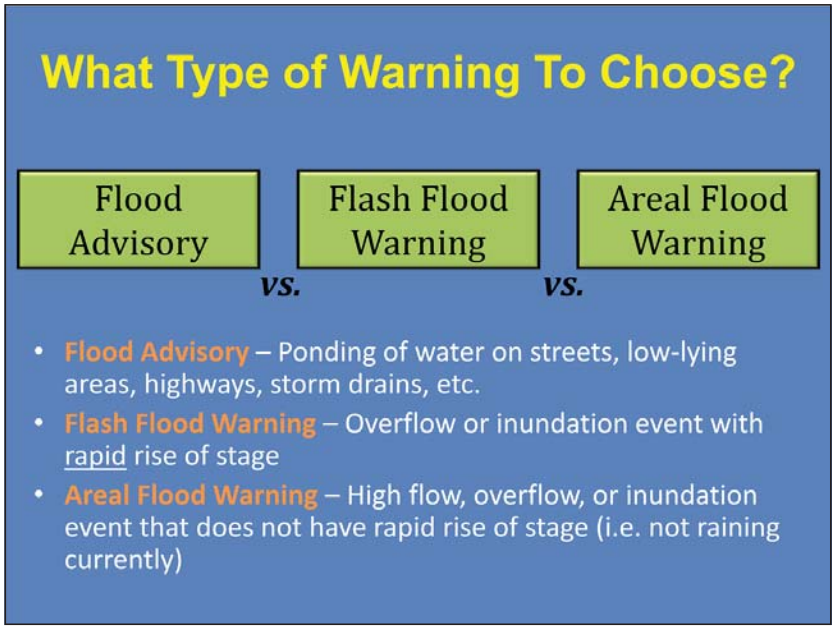
- How does QPE and FFG vary across your CWA?
- Where has FFG been exceeded and by how much?
- Where/How is the threat evolving?
- Think ahead!



Before we jump into the details of creating polygons and text warnings, there are several questions you should always be asking yourself when defining the threat area.

First, how does your selected QPE source and flash flood guidance vary across the CWA? Is there an area of very high QPEs moving into an area of low FFG? Or, where has flash flood guidance been exceeded and by how much? Do you have an ongoing event? Next, where and how is the threat evolving? Again, is the event moving into an area of low FFG? Or, are training storms a concern?

The main take away is to always be thinking ahead! FFMP focuses heavily on what is happening now, based mainly on QPE. But as a forecaster, you must mentally extrapolate storm movement and threat evolution to generate proper lead time, particularly for rapid runoff in urban areas.



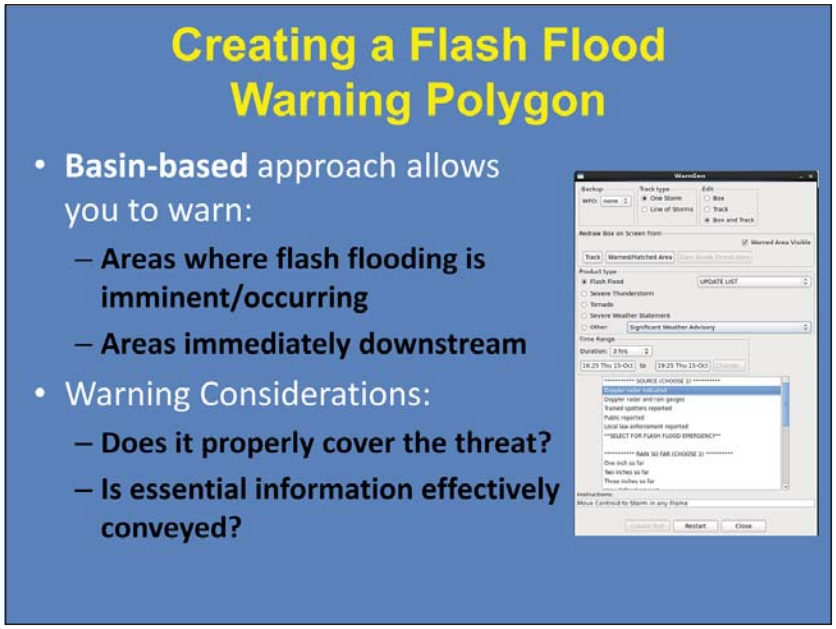
So let's start moving through the things to think about while on the hydro desk. Let's say you've diagnosed that the potential for flooding is likely…So what do you issue?

There are a few routes you can take: a Flood Advisory, Flash Flood Warning, or Areal Flood Warning. Make sure to talk to your office to see if they have a protocol in place.

Generally, a Flood Advisory covers any sort of ponding that is not life-threatening. A FFW should be used when there is a RAPID rise of water, within 6 hours. Whereas, an Areal Flood Warning would be used if there is high flow, but it is not a rapid rise.

So remember, for a FFW, you're looking for a rapid rise of swift-moving water. If this criterion is not met, one of the other two is probably a better choice.

A best practice: Let's say you are expecting a widespread, long duration rainfall event with marginal rain rates. Putting out a 6 to 24-hour Flood Advisory or Areal Flood Warning may be an effective product. But let's say that during this long duration event there are small areas of localized heavy rain rates that could lead to life-threatening flash flooding. Here, you would want to embed FFWs already under the Flood Advisory or Areal Flood Warning. This set-up properly explains the different threats. Once the flash flooding warnings have expired, Flood Advisories or Areal Flood Warnings can be continued for basins that are still seeing non-life-threatening, general flooding.

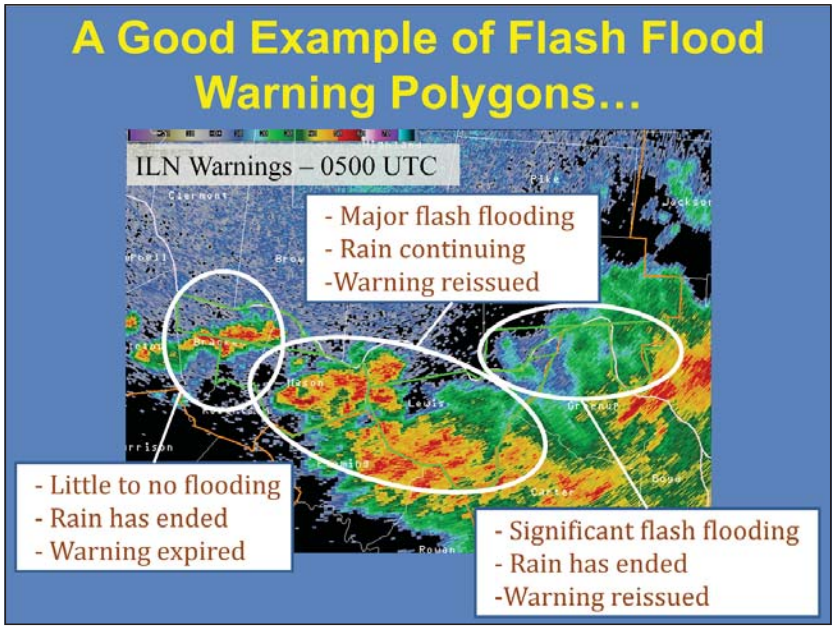


Okay, so you've decided that a FFW is necessary. So let's look into some of the warning polygon fundamentals.

First, your warnings should be basin-based, rather than storm-based. Remember, flash flooding is a two-headed beast…You must consider the meteorological **and** hydrological factors. An area with heavy rainfall will not produce flash flooding unless the hydrological criteria are also met. Basin-based warnings allow you to warn the areas where flash flooding is imminent or already occurring, as well as areas immediately downstream.

Additionally, you should consider if the polygon properly covers the threat area…Not only right now, but in a few hours too. Where is it moving? What is the hydrology like there?

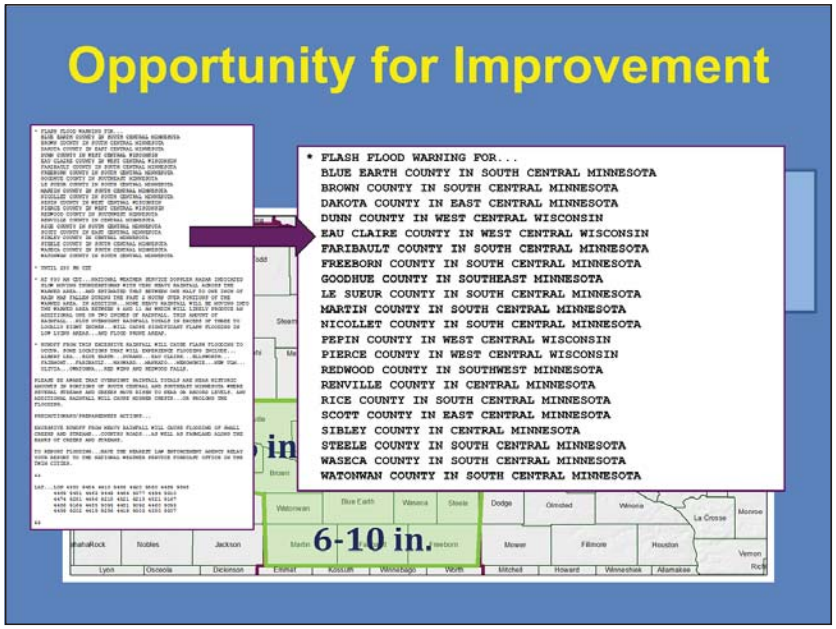
And, is essential information effectively conveyed? You don't want the threat to be overshadowed by wordy warnings listing obscure basin names.



So let's look at a good example of flash flood warning polygons from the Wilmington, OH office. We start at 0500 UTC, and we see 4 active warnings. Let's go ahead and circle the different threat areas, combining the two warnings in the middle.

Now, let's move forward an hour to 0600 UTC. For the far western threat area, there was little to no flooding, and the rain has ended. So this warning was allowed to expire. Moving east, there was major flash flooding occurring at this time, and the rain is continuing. Therefore, the office reissued the warning, and combined the threat area. Finally, the far eastern area had significant flash flooding, even though the rain had ended. Therefore, the warning was reissued.

In this case, the warnings were properly itemized, so to explain the evolving threats. And notice how the 4 original warnings were made to expire at the same time, which helped in the reissuing process at the later time.

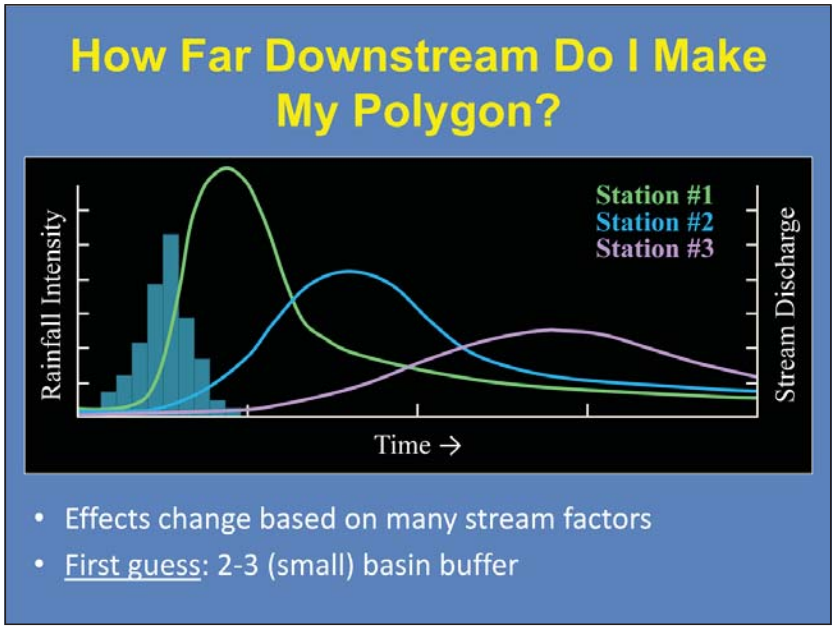


What about this example? Does the warning properly cover the threat area? And is the warning text effective?

In this case, there are 21 counties included in one warning. Would you expect the threat to be the same in all locations? Likely not.

In fact, there was a spread of 1-10 inches across the warning. Yet, all areas are receiving the same message. This is not ideal.

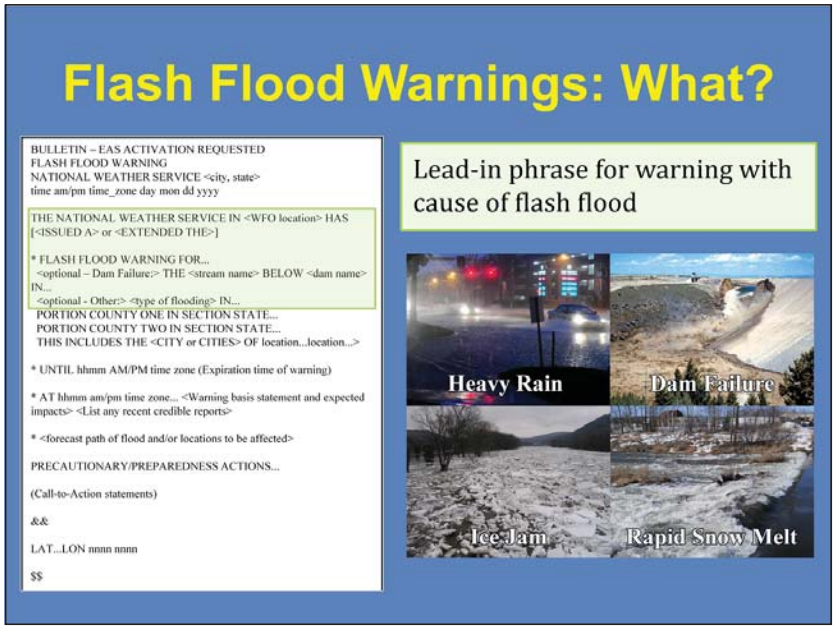
And in the warning text, all 21 counties are listed, which would be difficult for the existing systems to disseminate, like NOAA Weather Radio. Or even for the public to have to decipher themselves.



So you figured out your immediate threat area, but how far downstream is necessary to cover the evolving threat?

Here we have a plot of rainfall intensity over time. Let's say Station #1 is very close to where the rainfall fell. It is going to respond quickly to the rainfall, with a sharp jump in stage. Let's say Station #2 is a little further downstream, so its response is later than the first, with a more gradual rise in stage that isn't as high as Station 1. Finally, Station #3 is the farthest downstream, so it only sees a slight rise in stage.

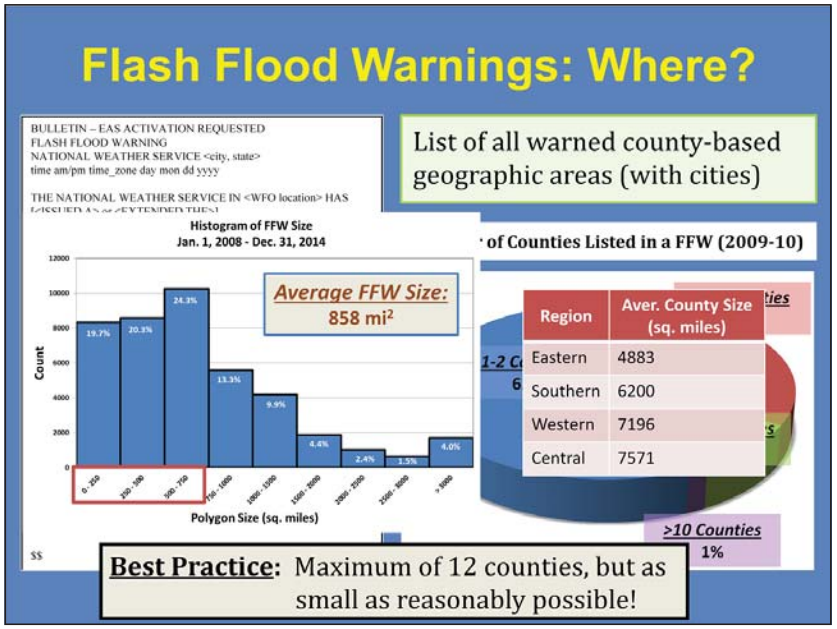
This progression can be expected in most cases, however effects can change based on many stream factors. A good place to start in the absence of any local hydrological knowledge is to expand your FFW 2-3 basins downstream to account for runoff, not 2-3 counties! This is in addition to the expanding threat due to training storms and the short-term movement of precipitation areas.



Okay, so you've drawn the perfect polygon, now let's start looking at the warning text best practices.

The first question is: What is the cause of the flash flooding? This will be the "lead-in" phrase for your warning.

In this course, you'll just focus on heavy rain. However, be aware that there are several other causes that you should learn more about from your office.

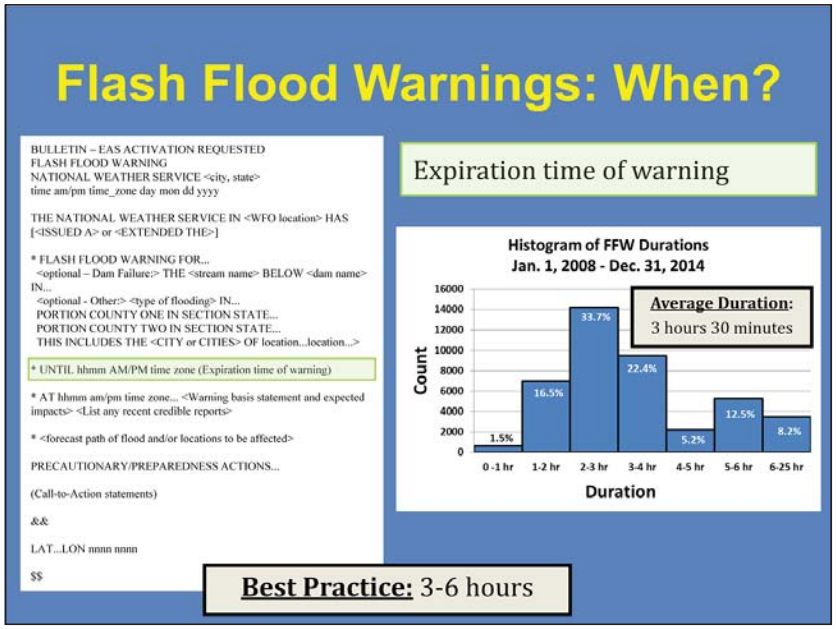


Then you have to add in the “where”. Where is this threat imminent or already occurring?

When it comes to how many counties you should include, consider this pie chart. It is the number of counties listed in a FFW, using data from all FFWs for 2009 and 2010. This chart shows that the vast majority of FFWs include less than 6 counties.

Since county size varies a lot across the country, consider this histogram, which uses data for all FFWs for 2008 through 2014. It shows the distribution of warnings based on size, in square miles. Here you can see that the majority of FFWs are less than 750 square miles, with an average warning size of around 850 square miles. To put this in perspective, here is the average county size for each region in the CONUS. So the good news is that our warnings, on average, are quite a bit smaller than the average county. However, there were several warnings between 18 and 24,000 square miles. So we still need to minimize warnings of these sizes, due to the reasons described earlier: what's the likelihood of everyone in these warnings having similar threats?

To give a number, best practice is no more than 12 counties included in your warning, again, understanding this will vary with county size. But ideally, your goal as a forecaster issuing FFWs is to make the warning as small as reasonably possible to cover the threat area and how it will evolve with time.



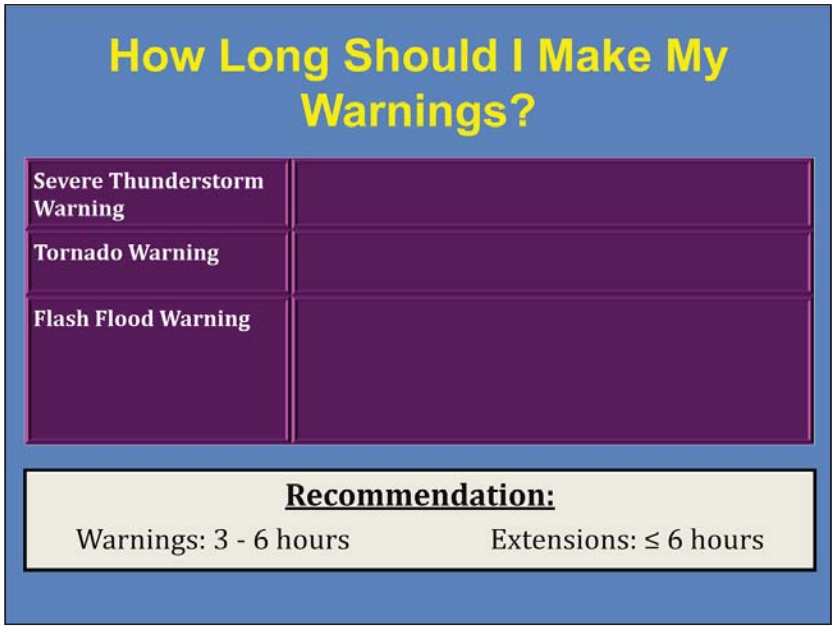
Next you have to add the “when”. When should the warning expire?

This is another selection that will vary significantly based on the cause, location, if storms are training, etc.

Using data for all FFWs from 2008 through 2014, we see the distribution of warning durations. The average for this 7-year dataset was 3 hours and 30 minutes.

In the absence of unusual circumstances, a best practice is for FFWs to be between 3 and 6 hours. For routine FFWs, 3 hours allows for one hour for the event to begin and the rain to fall, one hour for runoff and the stage to crest, and one hour for the flood to recede.

But let’s see what the directive says…



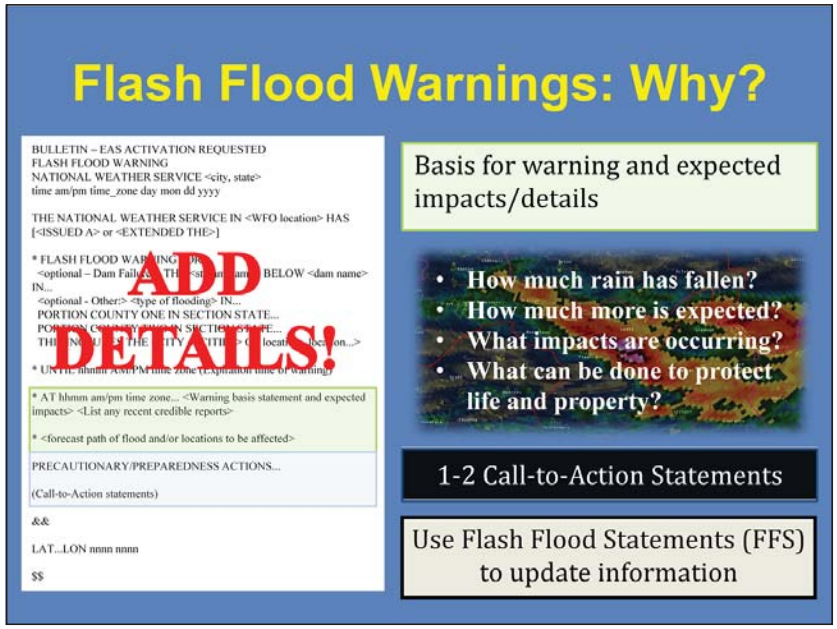
Pop quiz! Let's quickly review the different directives.

For a severe thunderstorm warning, what do you think the directive says? The answer is 30-60 minutes. How about a tornado warning? Know that one? It's 15-45 minutes.

Alright, now a FFW. Got a guess? Here is the answer: A flash flood warning will be valid from the time of issuance until the time when flooding (requiring immediate actions to protect life and property) is expected to end.

Hmmmm, so what does that mean? What would you say is a minimum? Maximum? Well, I can't give you an exact rule to live by, but here is what we recommend: For heavy rain threats, make your initial warning a minimum of 3 hours, as explained on the previous slide. The recommended maximum is 6 hours, if you expect repeated cores of heavy rain to move through the area.

In rare cases where there is long-term excessive rainfall where life-threatening flash flooding continues beyond 6 hours, extensions of no more than 6 more hours can be issued, if needed. Any longer than that, and you're getting into the realm of areal flood warnings.



Finally, we need to explain the "why".

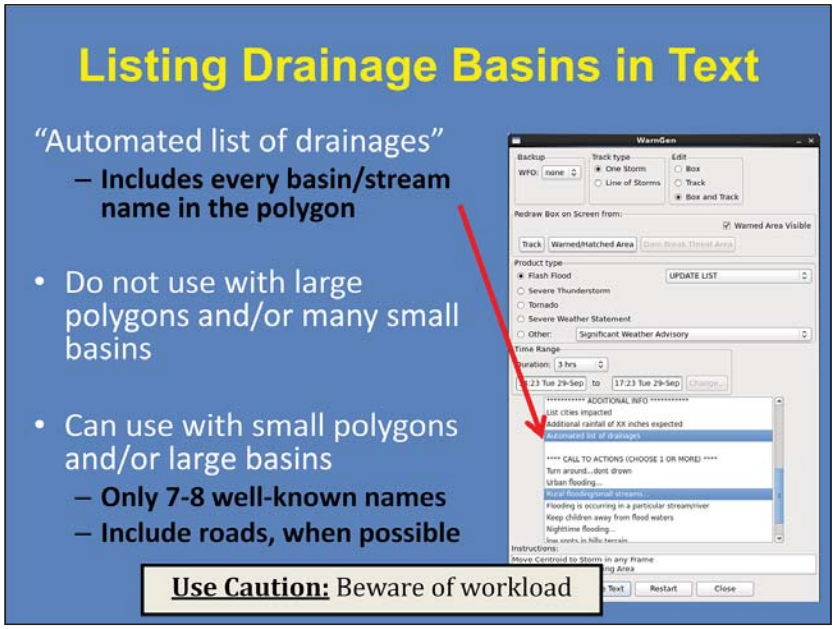
At this point, you may be thinking "Man, there's a lot to remember to forecast and warn on flash flooding!" But you know one good thing we've got going for us? A little bit of extra time!

Be sure to put in the details in the basis and call-to-action statements!

Answer these questions: How much rain has already fallen? How much more is expected? What impacts are occurring or can be expected? Include any relevant reports. And how can you protect yourself and your property?

Talk to your office about how many call-to-actions are good, but generally 1-2 per warning is best practice. You can include more if it's a significant event, long-lasting, in a metropolitan area, etc. Just do whatever properly disseminates the threat information.

Finally, don't forget to use Flash Flood Statements to update an ongoing event. Flash flood warnings are relatively long, recommended at 3 to 6 hours. Flash Flood Statements can be very useful in disseminating new information as the threat evolves.



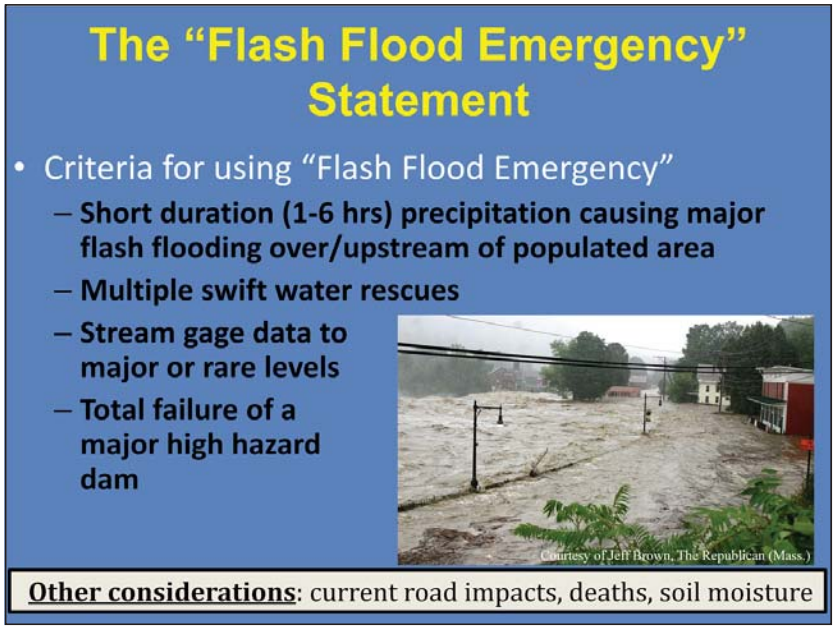
An important side note about one particular call-to-action.

Basin names can be automatically inserted into FFWs by clicking the "Automated list of drainages" option under the Calls to Action, shown here. What this does is include every single basin/stream name that falls within the warning polygon. There are positives and negatives to this.

For one, if a polygon is too big or in an area with small basins, this option can lead to a list of hundreds of basins. Since there is no geographic organization to this list, many of the basin and stream names may be unfamiliar to the general public. This can lead to un-needed and unwanted text in your warning. Reducing this list would require massive amounts of text editing, and simply not feasible during warning operations.

However, if a polygon is small or basins are rather large, this option may prove useful. Only keep drainage names of the creeks under the biggest threat and those well-known to the average customer, and/or those creeks and rivers that are well identified by signage for travelers to the area, if that information is indeed known. Try to reduce the number of names down to about 7-8. Also, include known road crossings that may be affected by flash flooding since the general public and media would recognize those even better than most creek names.

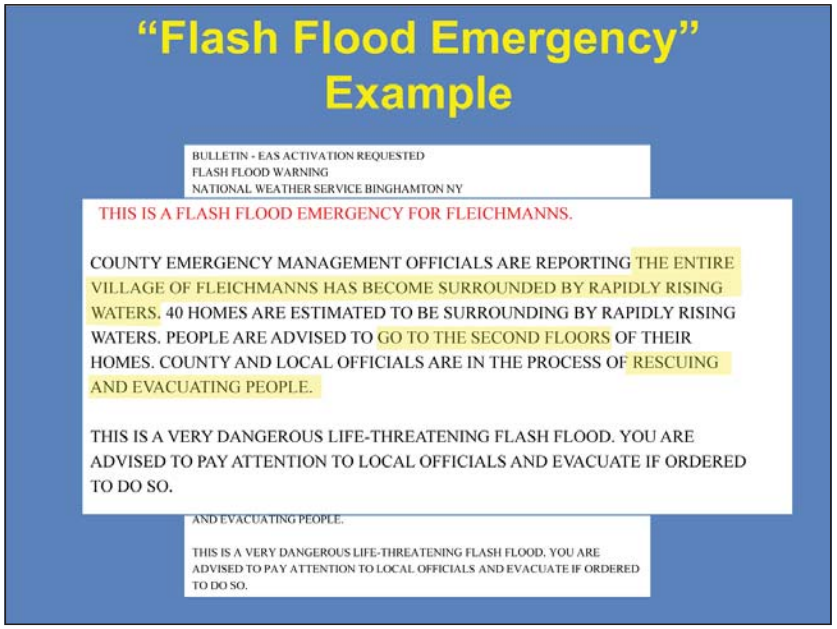
Simply use caution when considering this option. Determine if your hydrologic knowledge and workload can handle the edits needed to make the information pertinent to the public.



The last question is "When should I use a Flash Flood Emergency" statement? Well, to put it in perspective…Out of about 6000 FFWs per year, only about 5-10 are emergencies. Here are the criteria for a FF Emergency.

Basically, there must be an imminent or ongoing **elevated** threat to human life and/or **catastrophic** damage to property. Other considerations include road impacts, reported deaths, and soil moisture.

Remember: This is not a forecast! It should be issued only after you have reports. However, keep in mind that while you want to wait until you know the event is worthy of an Emergency tag, you want to declare an Emergency while it is still early enough in the event to be useful in terms of life and property impacts.

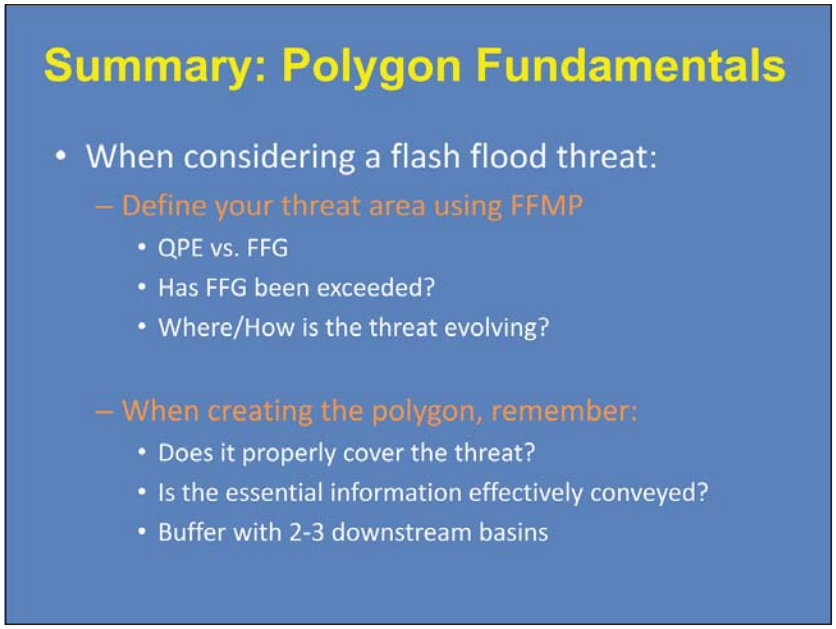


Here is an example of a Flash Flood Emergency statement.

Notice the language used: "…the entire village…has become surrounded by rapidly rising waters", "go to the second floors", "rescuing and evacuating people".

In these circumstances, if you are going to use the Emergency statement, you want people to change their behavior by using strong wording.

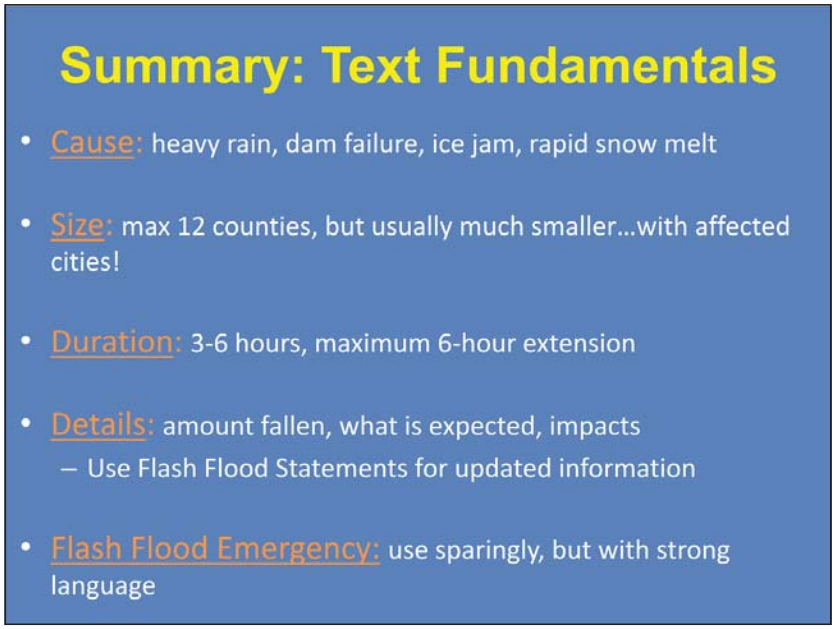
More information about Flash Flood Emergencies is covered in the WOC Flash Flood Course.



In summary, there are several things to think about at the start of your flash flood threat assessment.

Before even drawing your polygon, define the threat area using FFMP. Consider how QPE and FFG compare. See if FFG has been exceeded and by how much. Then, think about where the threat is evolving. Does the current storm motion, rain rates, and FFG values help point out the next area to expect impacts?

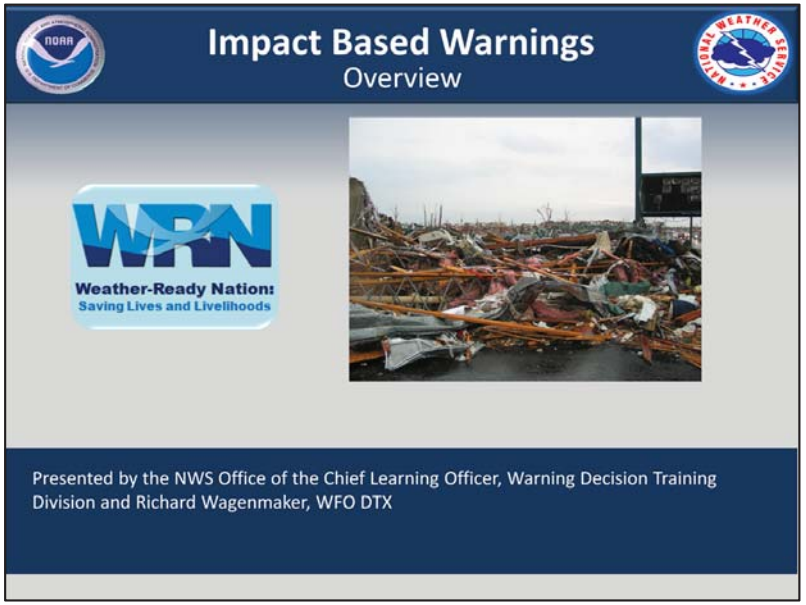
Once you start creating your polygon, ask yourself two important questions: Does the polygon properly cover the threat? And is the essential information effectively conveyed? Too large of polygons can lead to problems with impacts not being relayed to the correct audience, as well as cause alert systems to have issues reading the lists of county and basin names. Oppositely, too small of polygons can accidentally leave threatened areas out. Remember to account for routing effects by including a 2-3 basin downstream buffer to your polygon, in addition to the expanding threat.



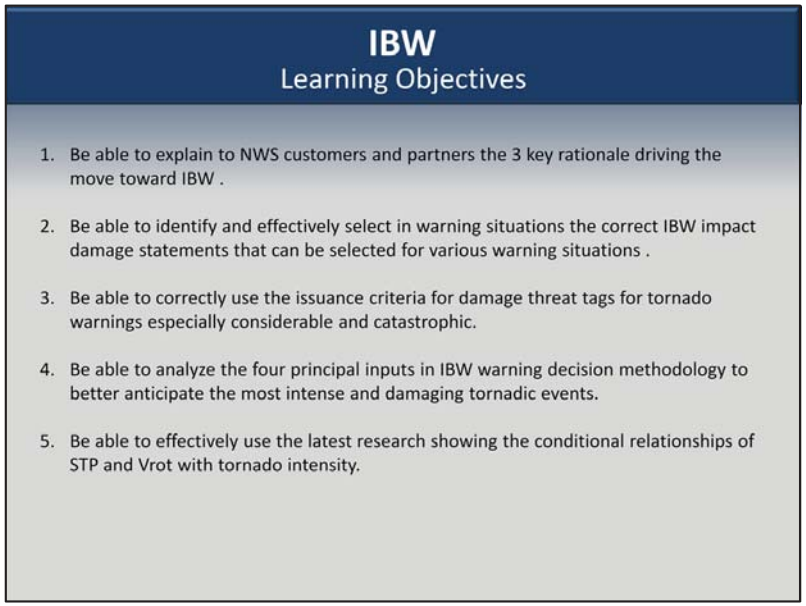
Once you are happy with your polygon, it's time to start thinking about the warning text.

1. Include the cause of the flash flooding, which is generally one of these four things. Talk to your office about how they handle each of these options.
2. As for size, you don't want any more than 12 counties listed, as it can make alerting difficult. But of course, try to make the warnings as small as reasonably possible, so that you are properly covering the threat. And, make sure to include any affected cities, since this is important, easily understood information.
3. FFWs, generally, should not be any shorter than 3 hours or any longer than 6 hours, and should not be extended more than 6 hours at a time. The event needs enough time to evolve, but also shouldn't be so long that the impacts are more areal flood-related.
4. Finally, and perhaps most importantly, make sure to take your time and include the necessary details about the event. Include how much rainfall has already fallen and what else is expected over the warning duration. And include current and expected impacts. Follow the warning with at least one Flash Flood Statement, including any reports and relevant rainfall and hydrological updates.
5. And hopefully you won't have to worry about it, but use Flash Flood Emergency wording sparingly. If you do meet the criteria and decide to use it, strong language should be used to notify the public of the elevated risk and to take action to protect themselves and their property.

This lesson was meant to provide useful recommendations when you are creating flash flood warning polygons and text. But don't forget to work with your local office to learn more about their protocol, as well.



Hello. This is Brad Grant of the Warning Decision Training Division of the Office of the Chief Learning Officer. This is a 2-part course on Impact Based Warnings, or IBW. The first part of the course is intended to be an overview of IBW to help NWS warning forecasters become more familiar with the updated warning practices brought about by the implementation of IBW. You will be hearing from Dick Wagenmaker, MIC of NWS DTX who has been one of the leaders for the IBW demonstration project, first started in Central Region.



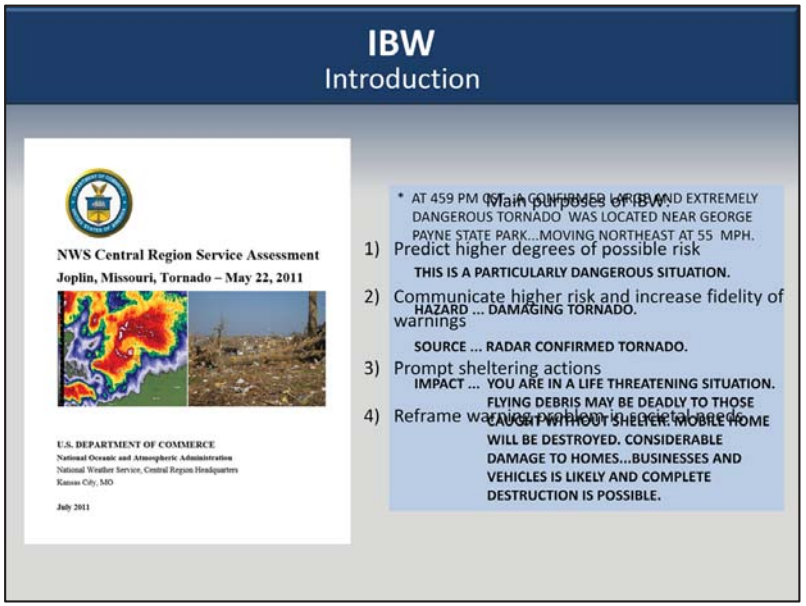
These are the learning objectives for the IBW Course: Be able to explain to NWS customers and partners the 3 key rationale driving the move toward IBW.

Be able to identify and effectively select in warning situations the correct IBW impact damage statements that can be selected for various warning situations.

Be able to correctly use the issuance criteria for damage threat tags for tornado warnings especially considerable and catastrophic.

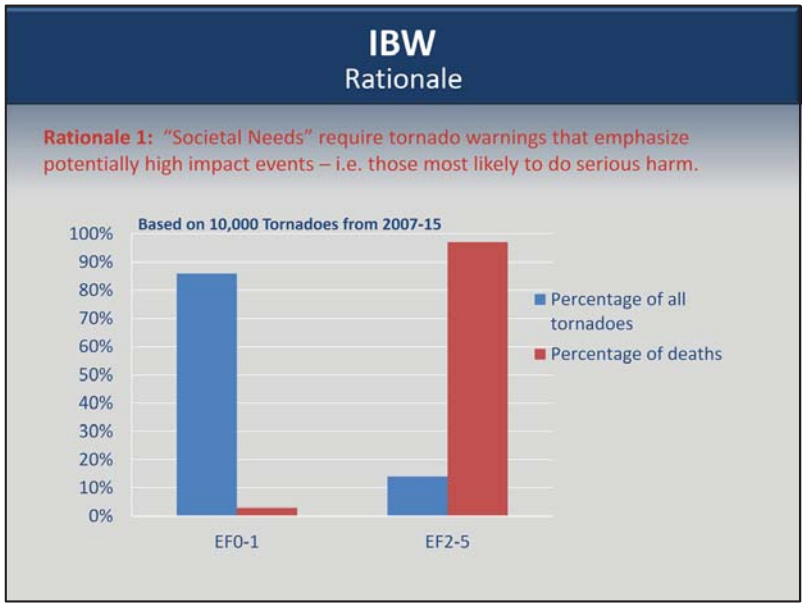
Be able to analyze the four principal inputs in IBW warning decision methodology to better anticipate the most intense and damaging tornadic events.

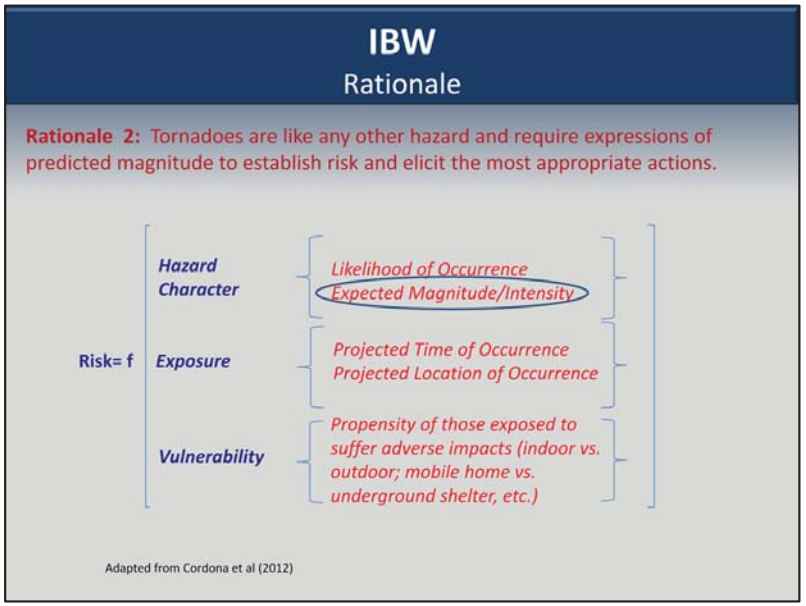
Be able to effectively use the latest research showing the conditional relationships of STP and Vrot with tornado intensity.



Hi. This is Dick Wagenmaker of NWS WFO DTX. Welcome to this introduction and overview of the Impact Based Warning training course. As most of you know, the content and construct of our severe weather warnings have changed little in 50 years. Following the tornado disasters in the spring of 2011, the Joplin, MO Tornado Service Assessment used its findings to propose exploring an evolution of the existing NWS warning system to facilitate improved public response and decision making in the most life-threatening weather events. IBW is intended to be a simple incremental first step in the evolution of these warnings to provide a better service. This is not an evolutionary leap. Since the initial proposals, noteworthy progress in both research and operational aspects of NWS tornado warnings have led to spirited debate in order to improve the utility of our warnings.

Essentially, IBW is a “risk-based” approach designed to 1) predict higher degrees of risk when possible (like any other type of warning); 2) communicate higher risk and increase the fidelity of warnings by telling people what we know; 3) prompt sheltering actions by adding emphasis for the most life-threatening weather events; and 4) reframe the warning problem in terms of societal needs.





What is meant by reframing the warning problem in terms of societal needs is rationale number 1: That the public as a whole require tornado warnings that highlight events that have the most potential to do serious harm. It is important to note here that, yes, all tornadoes are dangerous - as are all severe thunderstorms - and the IBW project does not imply otherwise. However, the numbers supporting risk-based warning concepts are compelling. Nationally, over the most recent 8 year period, which included over 10,000 tornadoes, just 14% were rated EF2-5… but these result in 97% of the fatalities. Contrast that with EF0-1 tornadoes which constitute 86% of all tornadoes, but only 3% of the fatalities. Only 2 fatalities resulted from a very large number of EF0's, suggesting a mortality rate for weak tornadoes is roughly equivalent to the mortality rate from severe thunderstorm winds. Clearly, while all tornadoes are dangerous to degrees, all tornadoes are not the same. And based on mortality, there is a clear societal need for tornado warnings that emphasize potential high impact events. Over the past two decades much of the research in this field has focused largely on the tornadogenesis problem and distinguishing tornadic storms from non-tornadic storms. IBW simply tries to reframe the warning problem into better distinguishing strong tornadoes from weak tornadoes; thereby correcting what many view as a flaw in the legacy tornado warning paradigm , and one that leaves the public exposed to the dangers of high end tornadoes.

Following on these ideas, the second rationale for conducting the IBW project is to help provide clarity on risk assessment for users of our warnings. At the beginning of IBW, project social scientists surveyed Emergency Management personnel in Kansas and Missouri concerning risk communication. Their response was that knowing the potential intensity or magnitude of the tornado was an important factor in helping them determine a course of action. This makes obvious sense. What if we issued a flood warning for the Red River in Fargo, but refused to say how high above flood stage the river would get? Or, if we predicted hurricane landfall without offering max wind speeds or storm surge; or if we forecasted fog at a major airport without providing an expected visibility. We could on and on with examples. (Click) On this slide the simple risk paradigm adapted from Cordona et al. (2012) provides a summary of the basic information people need to assess their personal risk from any weather hazard. In this model, risk is expressed as a function of hazard character, exposure to the hazard, and vulnerability to the hazard. From the point of view of those issuing warnings, this includes identifying the likelihood of the hazard, the magnitude of the hazard, the expected time of occurrence, and the expected location of the hazard. The missing piece of the risk paradigm in tornado warnings is information on hazard magnitude. Rationale #2 is that tornadoes are like any other hazard and require expressions of magnitude to establish a level of risk and elicit the most appropriate actions. This should especially resonate given the huge differences in mortality between strong tornadoes and weak tornadoes.

IBW
Rationale

Rationale 3: Clear and credible risk communication is necessary for people to take immediate protective action.

Key findings from 2013 National Institute Standards and Technology (NIST) Report on the Joplin tornado and recent NWS Service Assessments:

- High-intensity cues (risk signals) prompt people to take action; outside of IBW, there are few mechanisms to elevate threats within NWS tornado warnings.
- Most seek confirmation from additional sources before seeking shelter. Thus, consistency of message is important. Conflicting or incomplete information delays sheltering actions.
- Existing dissemination systems not fully compatible with storm-based warning polygons can cause confusion over threat location when there are multiple polygons.
- Perceptions of false alarms may negatively impact warning credibility.

IBW
Overview

What are Impact Statements?

- Provide end-users high intensity cues for particularly dangerous situations
- Describe what may happen **IF** hazard impacts people or structures

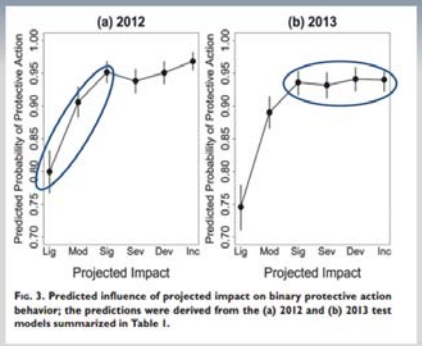


FIG. 3. Predicted influence of projected impact on binary protective action behavior; the predictions were derived from the (a) 2012 and (b) 2013 test models summarized in Table 1.

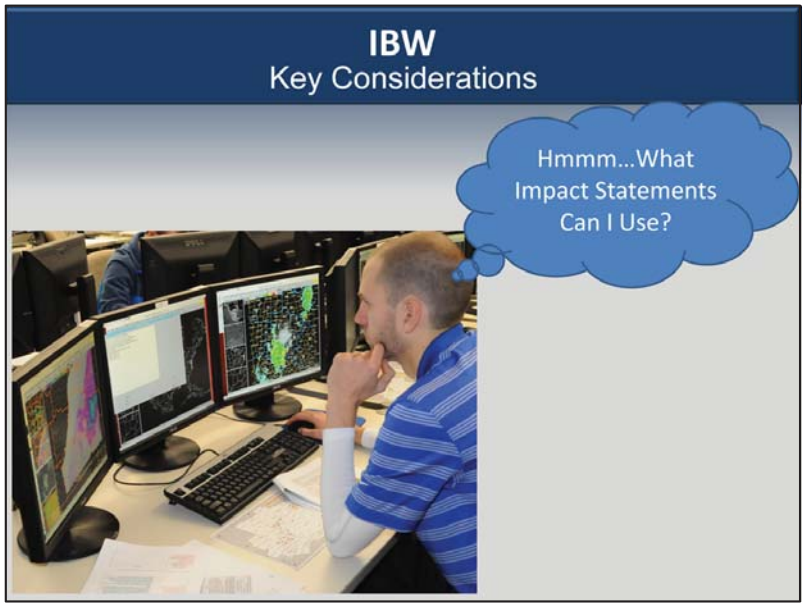
Ripberger et al (2014)

Finally, clear and credible risk communication is necessary for people to take timely protective action. Social scientists tell us that improving communication of risk is the prime public warning challenge for events like tornadoes. A key is converting people's natural perception of safety (which is called optimism bias) to a perception of risk… and thus speed-up risk assessment and sheltering actions. To do this people need clarity on impacts, as in does this affect me ? And how severe is it going to be? The intent is not to scare people (as many incorrectly suggest IBW intends), but to create fidelity in the warning message. What is meant by this is simply that we should inform people of what we know, and by doing so help people make quick and proper sheltering decisions.

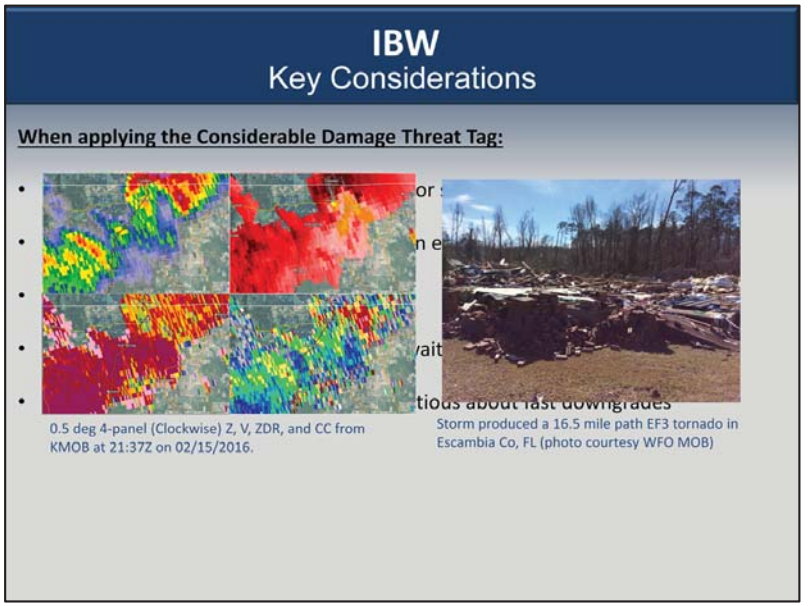
Going back almost a half century, the Lubbock Tornado Service Assessment from 1970 was the first to recommend a different siren tone for tornadoes as a way of elevating the threat. More recently, a 2½ year study by NIST of the Joplin tornado echoed NWS Service Assessment findings that showed high intensity cues are what prompted people to take sheltering actions…and that people will seek confirmation from additional sources before sheltering. These studies also stressed the importance of consistent messaging across the weather enterprise. Inconsistent or incomplete messaging can result in delayed or incomplete sheltering.

Existing dissemination systems are also not fully compatible with storm-based warning polygons and these can cause confusion over threat location when there are multiple polygons, especially overlapping ones. (click) Lastly, a recent study from Ripberger et al. showed that the credibility of the warning is important and that perceptions of false alarms and missed events can play a role in public response. We'll talk a little later in this presentation about IBW warnings and confidence markers.

As previously mentioned, a goal of the IBW demonstration is to provide high intensity cues to emergency managers when we know it for particularly dangerous situations. To do this tags have been expanded from severe thunderstorm warnings to tornado warnings so that national users can code software to read details about the warning without having to do extensive word searching of the warning text. Impact statements commensurate with the damage threat indicators are sometimes referred to as consequence-based messages and are intended to serve as the high-intensity cues as referred to in the NIST report. A 2014 study by Ripberger et al. presented empirical evidence that consequence-based language can have an important and desirable effect on people's propensity to take sheltering actions - up to a point. Their findings specifically showed that vulnerable residents told to expect a high-consequence event were more likely to take sheltering actions than residents who were told to expect a low-consequence event. This is shown in the figure on the right showing the probability of protective action rises sharply with higher-consequence messaging but thereafter levels off with more dire consequence messaging. To correct for that, starting in 2014, extreme wording in the impact statements in the tornado emergency was scaled back to match that for the considerable tag. The severe and tornado impact statements are meant to be conditional (that is, what may occur should the expected tornado strike infrastructure, trees, etc.). They fall directly out of your choice of tag and are commensurate with damage threats associated with the EF spectrum. These were formulated through a Regional Labor Council effort with input from social scientists and meteorologists. During the decision-making process you should be less concerned with the impact statements (which again are designed as cues for end-product users)… and focus on the meteorology associated with distinguishing strong tornadoes vs weak tornadoes vs. no tornado.

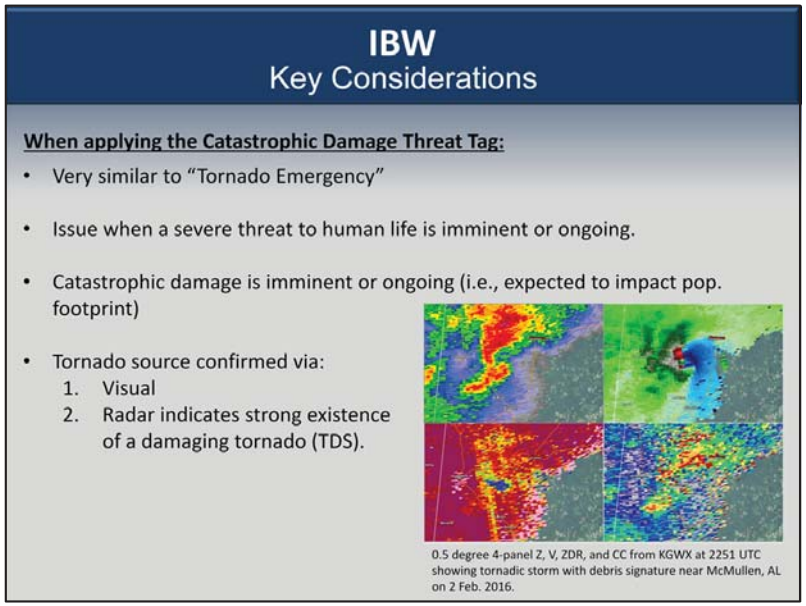


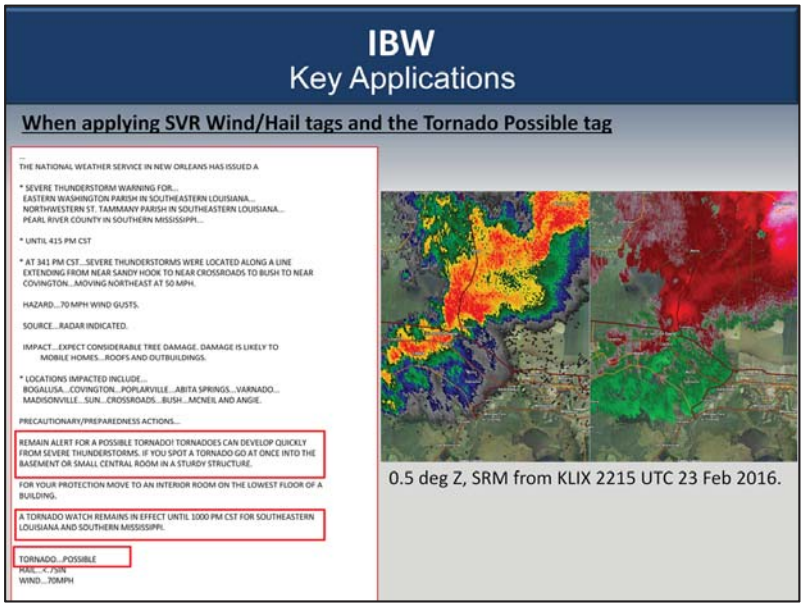
Next we are going to provide a series Engage interactions that describe the new IBW SVR and TOR tags and impact statements for the various types of tornado, wind, and hail hazards.



When applying the considerable damage tag, remember, the intent of IBW is to warn for high-impact events rather than try to predict actual storm impacts. So, the primary IBW tool to alert for high-impact tornadoes is the "Considerable" damage threat tag. As warning forecasters, your target range for the "Considerable" tag are EF2-5 tornadoes. This is also where enhanced, conditional, impact statements kick in to provide needed high-intensity cues for end-users and partners. By the way, for those wondering, the phrase "considerable" for this damage threat indicator was selected by emergency managers as a more descriptive proxy for "significant" as defined by SPC). The considerable tag should be selected only rarely and for those tornado warnings where the storm information (from near storm environment and radar signatures, or even spotter reports) suggest the possibility of a strong tornado. Don't try to pinpoint EF scale. Its perfectly acceptable if an EF1 occurs on a Considerable tag or an EF2 occurs on a Base tier warning. Radar signatures are the primary method for distinguishing between significant tornadoes and small tornadoes. But you do not need to wait for a report of a tornado.

You can upgrade a tornado warning using the SVS option, but be cautious about downgrading too soon. Here's an example of a tornadic storm that occurred with an outbreak severe weather in the southeastern U.S. on President's Day 2016. This storm moved across western portions of the Florida Panhandle and Southern Alabama with a swath of damage.



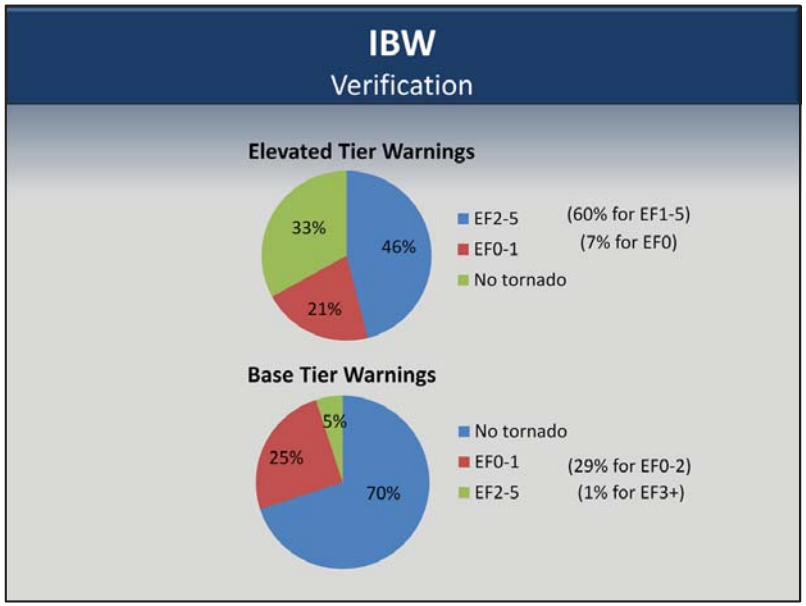


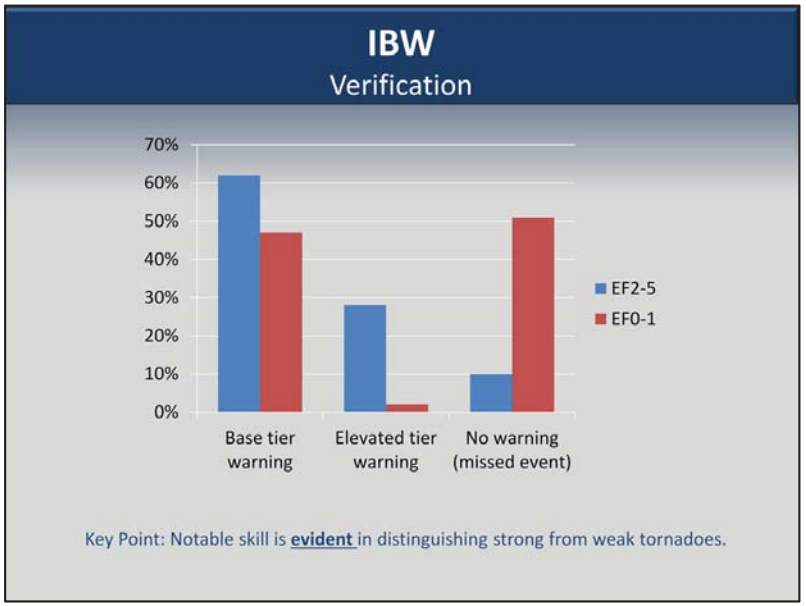
When applying the catastrophic damage threat tag, here are the important considerations: First, recall the term "tornado emergency" has been in the forecaster toolbox for 15 years and so in the IBW framework, we've adopted the "catastrophic" damage threat tag to be very similar to the use of Tornado Emergency criteria. In other words, the catastrophic damage tag is appropriate for the warning situation if all of the following criteria are met:

a. A severe threat to human life is imminent or ongoing,
b. Catastrophic damage is imminent or ongoing , AND you expect the tornado to impact a population footprint.
c. Reliable sources confirm the tornado , either by a visual or via radar imagery, which strongly suggests the existence of the damaging tornado (e.g. debris ball signatures).
Since the interpretation of the first part of the catastrophic criteria is somewhat subjective, especially the determination of the size of population footprint impacted, it is requested that you work within your CWA partners and Regional Severe Weather Program Focal points to determine the best use of these rare situations.

Severe thunderstorm warnings and SVSs continue to have wind and hail tags. The value next to the tag is an indication of how strong the winds may be. Although most warnings will have a 60 or 70 mph value in the tag, the tag allows forecasters to express much stronger winds. This is useful for describing impacts from high winds associated with a derecho, for example. Hail tags are optional in tornado warnings - but wind tags are not used in tornado warnings to avoid confusion with regard to tornadic winds. Impact statements are commensurate with the expected hazard are also included in severe thunderstorm warnings.

The "tornado possible" tag is used in severe thunderstorm warnings for situations where a severe thunderstorm has some potential for producing a brief, small tornado, but forecaster confidence is not high enough to issue a Tornado Warning. This tag has also been in the forecaster toolbox for years according to NWS directives, and is typically used in QLCS severe thunderstorm events, or in severe thunderstorm warnings within tornado watches.





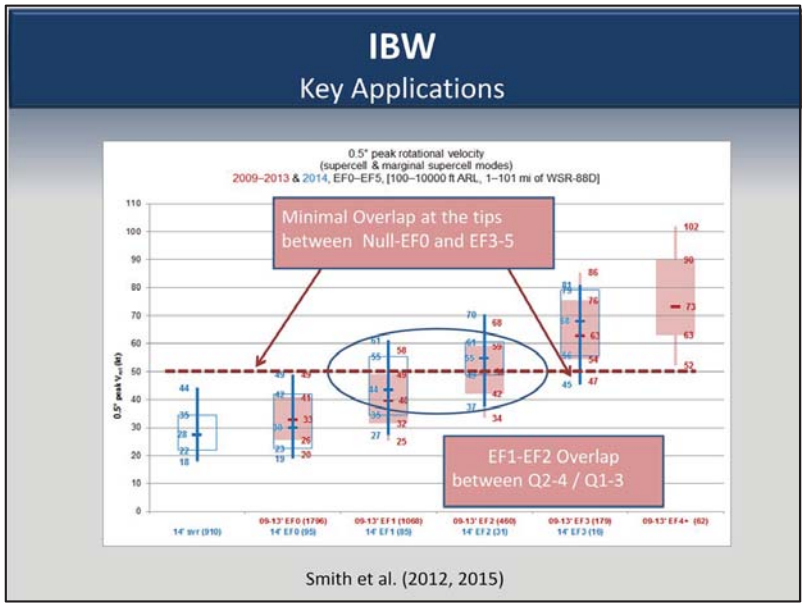
The IBW demonstration project included a verification project from 2012 to 2014 to verify both warning perspective metrics such as success ratios and false alarm ratios as well as event perspective metrics. The comparisons summarized on the next two slides here show value was added using warnings with damage threat tags (so-called "elevated tier" tornado warnings) over legacy warnings, or the so-called "base tier" warnings. In particular, the project showed that since the NWS warns for nearly all EF3-5's and can use elevated tier tags for half of them, the FAR is less than half than that for the base tier warnings. For example, the most likely outcome when an Elevated Tier Tornado Warning was issued for EF2-5 tornado occurrence was 46% (60% for EF1-5 occurrence), 21% for EF0-1 (7% for EF0 occurrence), and 33% for no tornado (i.e., a False Alarm). Contrast those statistics showing the most likely outcome when a Base Tier Warning is issued: 70% no tornado outcome, 25% of an EF0-1, and only 5% of an EF2 or greater.

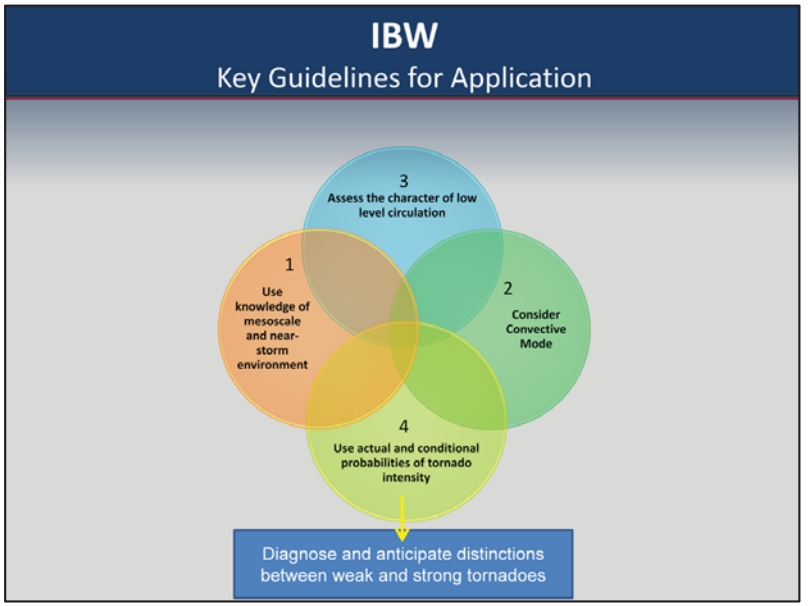
What does this data all mean? Well, for most base tier warnings issued, the most likely outcome is a false alarm, by a 70% majority. This is compared to the 33% false alarm ratio for elevated tier warnings. In short, you can see that false alarms in NWS warnings are largely the result of trying to warn for weak tornadoes.

From an event perspective, looking at tornado detection ratios and missed event ratios, when a EF2-5 tornado occurs (the blue series in the histogram) it is associated with a base tornado warning – 62% of time (51% when EF3+ event occurs), an elevated tier warning – 28% of the time (47% when EF3+ event occurs), and no warning (Missed Event), 10% of the time (2% when EF3+ event occurs). On the other hand, when a EF0-1 tornado occurs (red bar graphs) , it is associated with no warnings 51% of the time (54% when EF0 event occurs), 47% of the time with a base tornado warning (45% when EF0 event occurs), and only 2% with an elevated tier warning (1% when EF0 event occurs).

What does this all mean? Approximately 90% of EF2-5 tornadoes are warned with a warning of any tier. This is also true for 98% of EF3-5 tornadoes. Roughly half of EF0-1 tornadoes are warned and most of those are warned with base tier warnings. The verification also found that forecasters underutilize enhanced tags and/or issue them late. While the vast majority of strong tornadoes are warned, they are mostly covered with base tier warnings and only 28% by elevated tier warnings. This increases to 47% of EF3-5 tornadoes. While 13% of all tornadoes are of the strong variety, elevated tags are included in just 5% of all warnings.

Finally, skill in specifically predicting EF2-5 and EF0-1s independently about the same as distinguishing a tornado from no tornado. However, when using near misses (that is, partial credit for EF1s in elevated tier warnings and for EF2s in base tier warnings) notable skill is evident in broadly distinguishing strong from weak tornadoes.





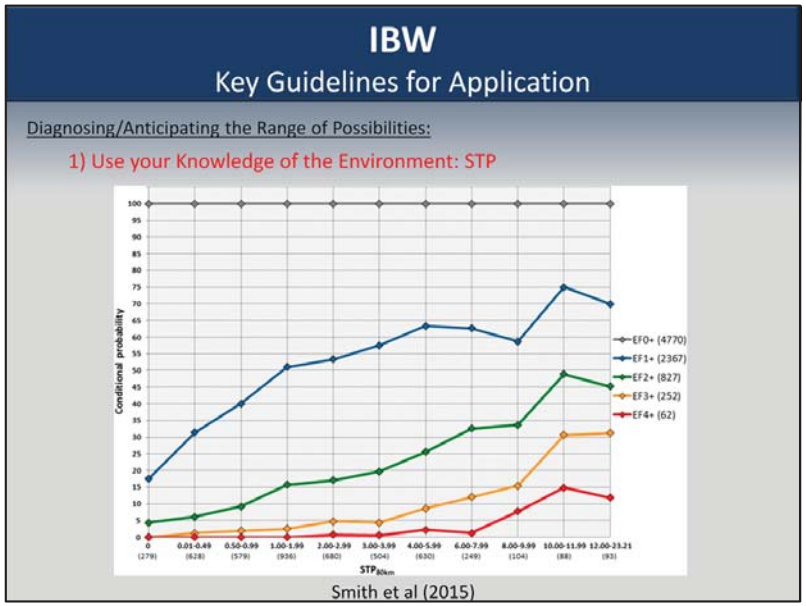
This SPC relational climatology by Smith et al from 2012, 2015 shows a significant relationship between increasing maximum 0.5 degree rotational velocity and increasing maximum EF scale of occurring tornadoes. The box and whiskers chart shown is for max low level rotational velocity signatures below 10000 feet and for supercell and marginal supercell convective modes. The light red is data from 2009-2013, while the blue is data from 2014 and includes rotational velocities from non-tornadic severe supercells.

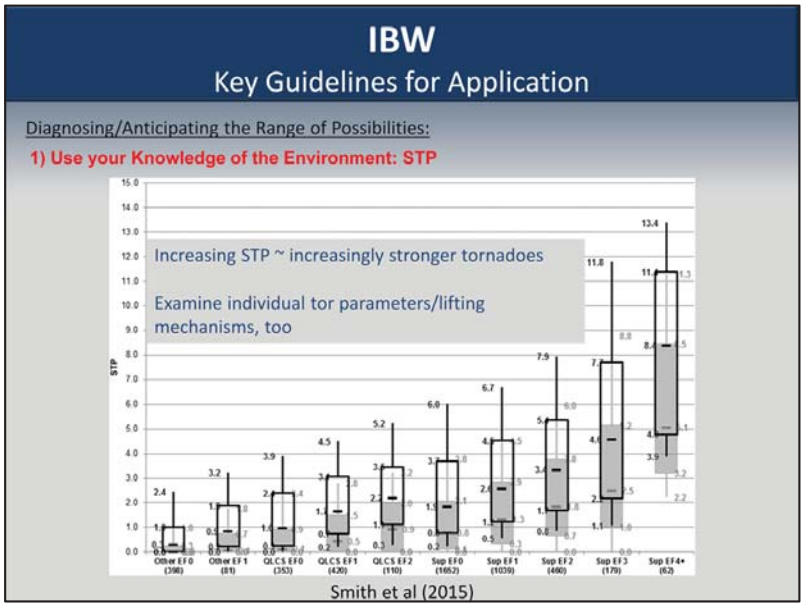
The box and whiskers have standard configurations with the box bounded by the bottom of the 2nd quartile and top of the 3rd quartile - and the tips of the whiskers represent the 10th and 90th percentiles. The dash in the middle of each box is the median or top of the 2nd quartile. It is important to note that for each event the study recorded rotational velocities immediately prior to tornado touchdown until just prior to tornado dissipation.

Again, we are not trying to pinpoint tornado intensity by EF scale – just "ring the bell" a little louder for more significant tornado events. This chart shows why. You'll note there is plenty of overlap in max low level rotational velocity associated with high end EF1 and low end EF2 – but you can also see there is little overlap between null events/EF0's and EF3-5's. Again, it won't be unusual for an EF1 to occur on Elevated Tier Warnings, nor unusual for EF2's to occur on Base Tier Warnings. We should avoid as much as possible having EF3's or greater occur on Base Warnings or No Warning, and avoid EF0's and Null events occur on Elevated Tier warnings. The graphic here certainly hints at some capacity for distinguishing between weak and strong tornadoes in that respect… and also hints at the viability of probabilistic approaches to the IBW warning process.

There is no perfect answer to the question of when to use a "considerable" tag, but that is rarely the case for anything in operational meteorology. Keep in mind that the value of your role in the warning decision process comes into play by staying situationally aware and considering a wide variety of factors to stay one step ahead of the tornado threat.

At this point we'll lay out some basic guidelines for distinguishing weak tornadoes from strong tornadoes. Essentially, there are 4 steps in the process which is diagnosing and anticipating the range of possibilities. First, use your situational awareness to assess the mesoscale and near-storm environments. Second, use your knowledge and understanding of convective modes and storm evolution as they relate to the environment. Third, use your understanding of the 4-dimensional character of radar-depicted mesocyclone circulations, especially the strength of low level rotational velocity, and 4thly, use your understanding of actual and conditional probabilities of tornado intensity (null vs. weak vs. strong) as related to low-level rotational velocity.





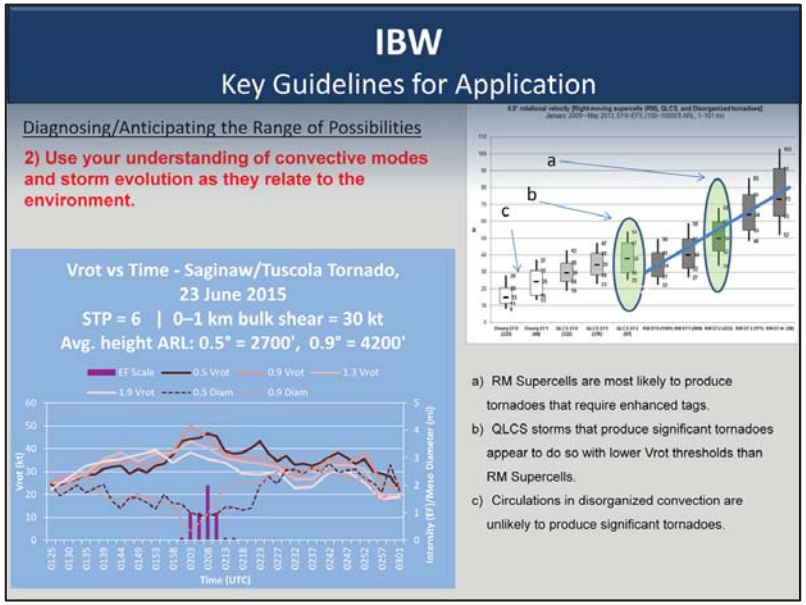
We'll first look at environmental parameters that serve as a baseline for determining the likelihood that strong tornadoes may occur in a given situation. The Significant Tornado Parameter (STP) 80 km is a multiple component parameter meant to highlight co-existence of low-level CAPE and shear which are crucial ingredients for right moving supercells. Based on past studies, STP exhibits greater skill in discriminating between nontornadic and significantly tornadic supercell environments compared to any of its individual components or any other parameters among the 38-variable database at the SPC

There is a relational climatology between increasing **conditional tornado intensity** and increasing values of STP 80km as measured on hourly SPC mesoscale objective analysis. Results from examining environmental and radar attributes, Smith et al (2015) found that increasing conditional probability for greater EF-scale damage, both STP and 0.5° peak Vrot increase, especially with supercells. This figure shows conditional probability of meeting or exceeding a given EF-scale rating (for the 5 series shown in the legend) for STP 80km for all convective mode tornado events from 2009–13.

The conditional probabilities may give you an idea ahead of time on how aggressive you can be on potential considerable tornado tags during the warning process.

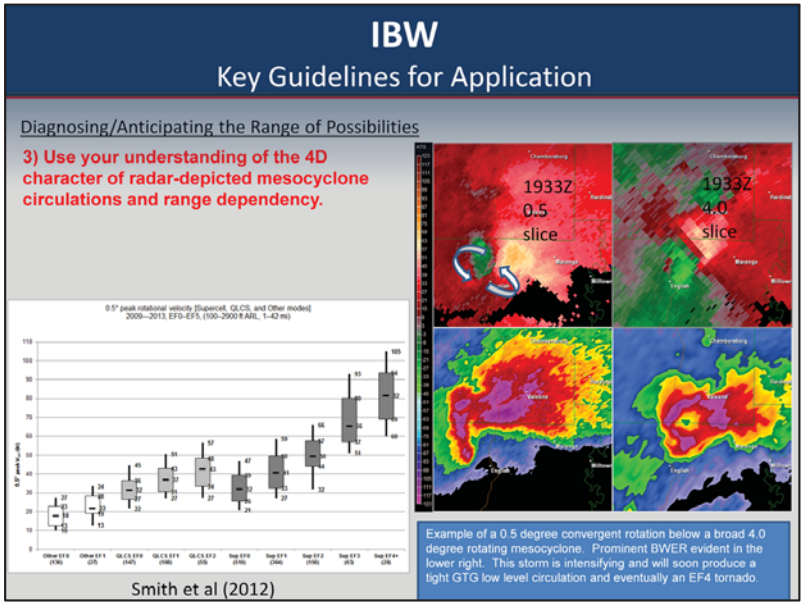
Next to further illustrate the relationship, here is a box and whiskers chart showing effective-layer STP (a dimensionless number) for all supercell, QLCS EF0–EF2 tornadoes, and other modes EF0 and EF1 tornadoes by EF-scale damage rating classes. The 40-km grid data are shaded gray, labels on right. The black overlays (with labels on left) denote STP 80 km grid point values, at the analysis time immediately preceding the event time. Using either the nearest STP grid point value or the neighborhood maximum value (which is the STP 80km) STP increases as tornado damage classifications increase. Supercell events tended to exhibit higher STP values than QLCS and other modes for the same EF-Scale damage rating. The STP for supercell, QLCS, and other modes tended to increase monotonically with increasing damage class ratings (aside from the 10th percentile). Substantial overlap exists in the distributions between adjacent EF-scale ratings, though the higher values of STP 80km (i.e., > 6) are more common for a greater proportion of supercell events at higher EF-scale rating classes (i.e., EF3+).

It must be stressed that composite parameters such as the STP 80 km should not be examined alone, but rather in concert with the individual components in the STP that identify important supercell tornado ingredients. Despite STP utility as a composite tornado predictor, there is no replacement for a thorough diagnosis of the mesoscale environment. You have to be aware of rapid storm interactions due to low-level boundary interactions which produce rapid tilting and/or stretching of local vorticity maxima. Also, closely monitor 0-1 km effective bulk shear, subsequent updraft helicity, and LCL heights < 1500 m AGL which are all important ingredients in the tornado development process.



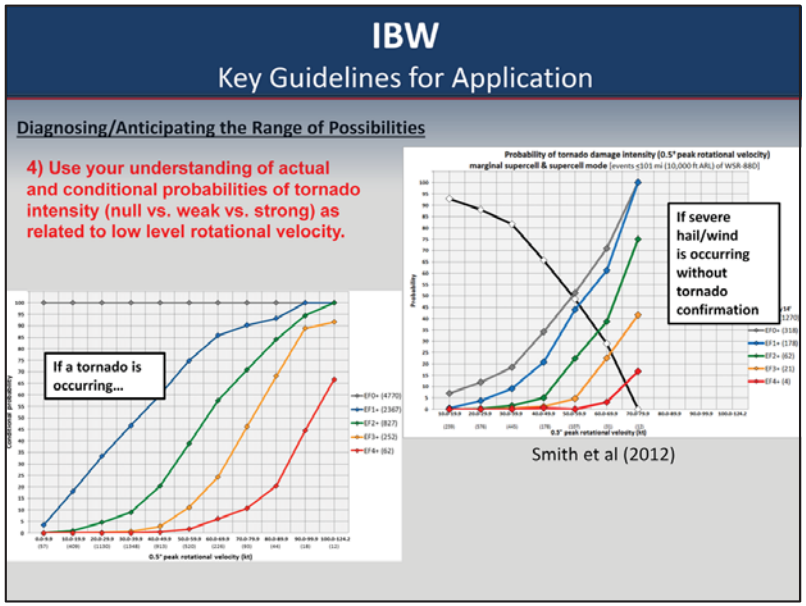
Here we'll discuss convective modes and rotational velocity evolution, or Vrot evolution. The graphic on the left is a marginal supercell tornado example from a study by Frey and Thompson in 2015. It shows the time evolution of Vr at various elevation slices through a storm. It also shows the time evolution of mesocyclone diameter through the storm cycle. On this graphic you can see a broad peak in 0.5 degree Vr around 45 knots starting near the time of tornadogenesis and a 0.9 degree peak near 50 knots right before a brief EF2 occurrence. Also note the max Vr starts at higher elevation slices 1.3 and 1.9 degrees but shifts to the lower elevation slices just prior to tornadogenesis. Additionally note how the meso diameter tightens as we approach tornadogenesis. This is a nice presentation of how you can view the full time evolution of a circulation to help make a determination of tornado development and potential max intensity.

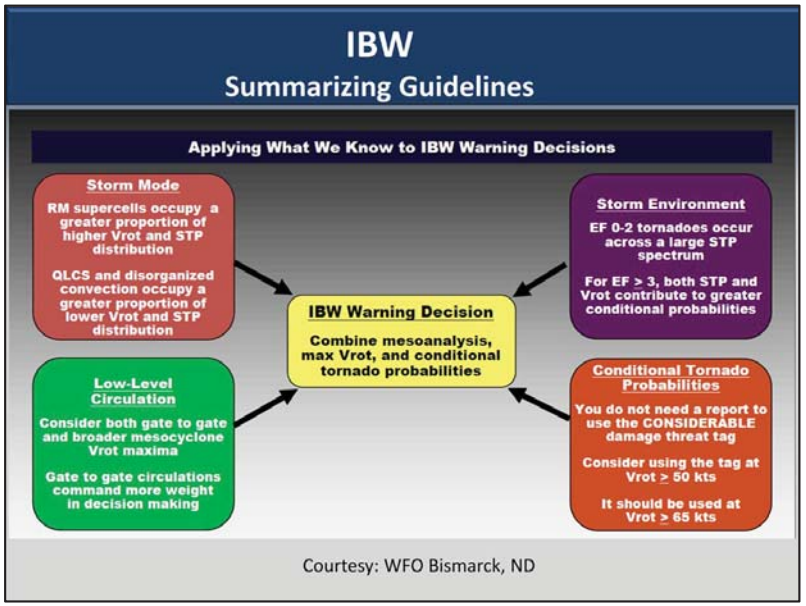
On the right side of the slide is the Smith et al. study breaking down max Vrot distributions vs. Max EF-Scale for each convective mode. Obviously, supercell modes have the strongest relationship and most strong tornadoes occur with these modes. On the opposite end of the spectrum, it is highly unlikely to have a strong tornado occur with weak disorganized convection. Last, it is interesting to note that strong tornadoes can occur with QLCS modes – but QLCS storms that do produce significant tornadoes appear to do so with lower Vrot values than RM Supercells. This possibly due to enhanced forward motion vector contributions on right flanks of low level circulations.



Thirdly, you must use your understanding of the 4-dimensional character of the radar-depicted mesocyclone and range dependency. On the right hand side of the screen is an example of how you should anticipate how convergent low level circulations will behave given the near-storm environment. The Smith et al study uses both broad Vrot maxima and Gate-to-Gate Vrot maxima, depending on which is strongest for a given case. It goes without saying that a Gate-to-Gate Vrot maxima should operationally command more weight and a lower Vrot probability threshold for EF2+ events. This example shows a 0.5 degree convergent rotation below a broad 4.0 degree rotating mesocyclone. Note the prominent BWER signature evident in the lower right. This storm is intensifying and will soon produce a tight GTG low level circulation and eventually an EF4 tornado.

On the left is the range dependency of the dataset. The relationship between Vr and EF scale is not as pronounced at longer ranges from the radar, and very pronounced at shorter ranges. Probabilities of tornado intensity increase as range from the radar decreases.





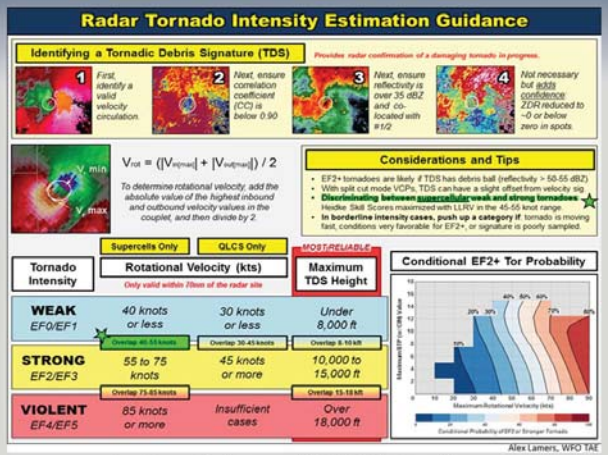
Lastly, here are two figures representing the raw probability space from the box and whiskers graphic that was previously shown. First, since there are several caveats that apply to the dataset, these probabilities should be considered ESTIMATES of tornado damage intensity as related to max 0.5 degree rotational velocity. Second, these data are only for supercells and marginal supercells. Data from most QLCS and non-mesocyclone tornadoes are not factored into the probability calculations. There is range dependency not accounted for in these graphs, which we'll explore more in the next slides. There is no filtering of the data to account for likely underestimation of EF scale that might be associated with tornadoes in areas sparsely populated with damage indicators, such as in the high plains. Time evolution of rotational velocity is not accounted for. Instead, the climatology relates max EF-Scale to max Rotational Velocity only. Circulation diameter is loosely accounted for, but specific distinctions between broader signatures and gate-to-gate signatures are not fully accounted for. And last, the null sample only includes events corresponding to significant severe events for large hail and winds – and this data was collected only from 2014.

Despite these caveats, the probability estimates contained here can be very useful if applied generally and in conjunction with other tools. In the upper right are the actual probabilities of EF-Scale supercell tornado intensity while on the lower left are conditional probabilities of tornado intensity (provided a tornado is occurring). As mentioned you can use each of these to help diagnose the need to issue a "considerable" tag while also considering range dependency, the character of the low level circulation, the favorability of the ambient environment, and time evolution of the supercell. These continue to be refined as more data is collected.

An interesting observation between the box and whiskers chart previously shown and the probability chart here on the right is that distinguishing between null severe events and EF0 tornado events using rotational velocity is nearly impossible. The probability distribution for null events vs EF0's is almost completely driven by the comparatively high volume of null events. In fact, once the probability of a tornado reaches 50%, the most likely intensity outcome quickly approaches EF2 or greater.

Here is a one-pager summarizing the various applications to IBW warning decisions, looking at environment, mode, circulation strength, and conditional probabilities. This summary come from WFO BIS.

IBW
Summarizing Guidelines



Courtesy: Alex Lamers, WFO Tallahassee, FL

Here's another one-page summary from SR that's provides the most recent methodology of evaluating Tornado intensity using radar only signatures including TDS height. All of the considerations and tips are based on peer-reviewed papers. Notice the categories of Weak, Strong, and Violent Tornadoes somewhat coincide with the selection of no damage tag, considerable, and catastrophic.

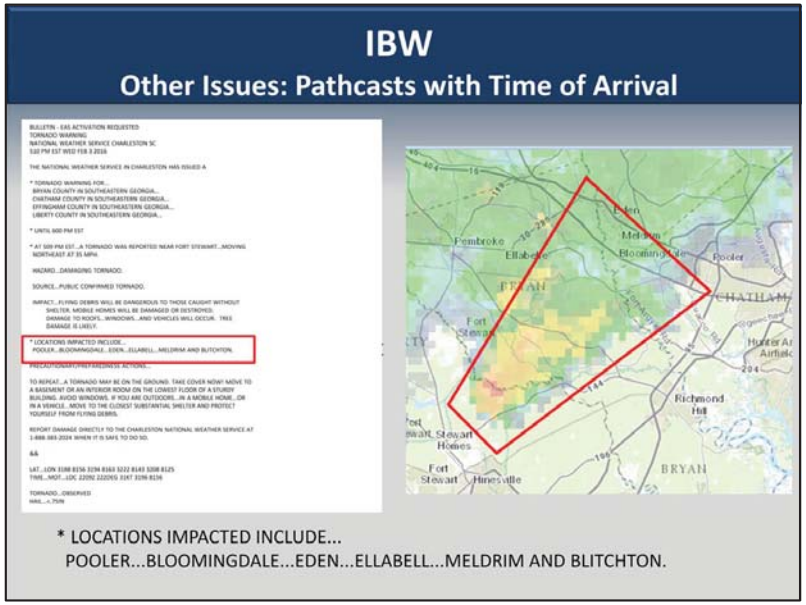
IBW
Other Issues: Pathcasts

- Recall technical limitations:
  - Radar resolution and range
  - Placing the "Drag Me To Storm" dot correctly
  - Large/Irregularly shaped cities
- Meteorological Limitations
  - Non-Linear Storm Motion
  - Multiple Threats from a Single Storm



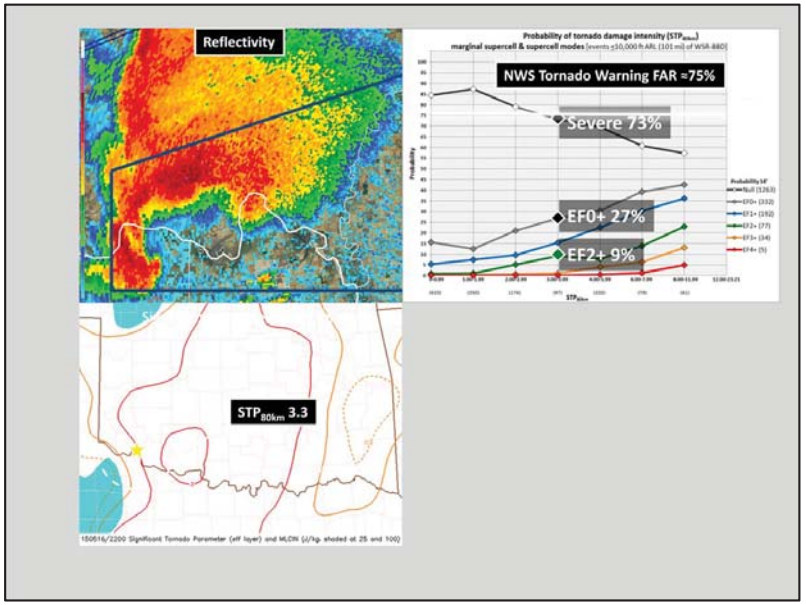
Town "A" is in the pathcast; Town "B" is not.

Since pathcasts are turned on with the default setting in WarnGen IBW templates, you will need to be very careful with using the time of arrival option version. Recall the technological limitations due to radar resolution and range, issues when you don't place the "Drag Me To Storm" dot correctly, and of course, all the problems caused in the shapefiles due to large and/or irregularly shaped cities. In addition, there are known meteorological limitations with using specific time of arrival locations in warnings such as non-linear storm motion, and conveying multiple threats from a single storm. Please review some of the specific practices regarding the misuse of pathcasts at wdtd.noaa.gov.



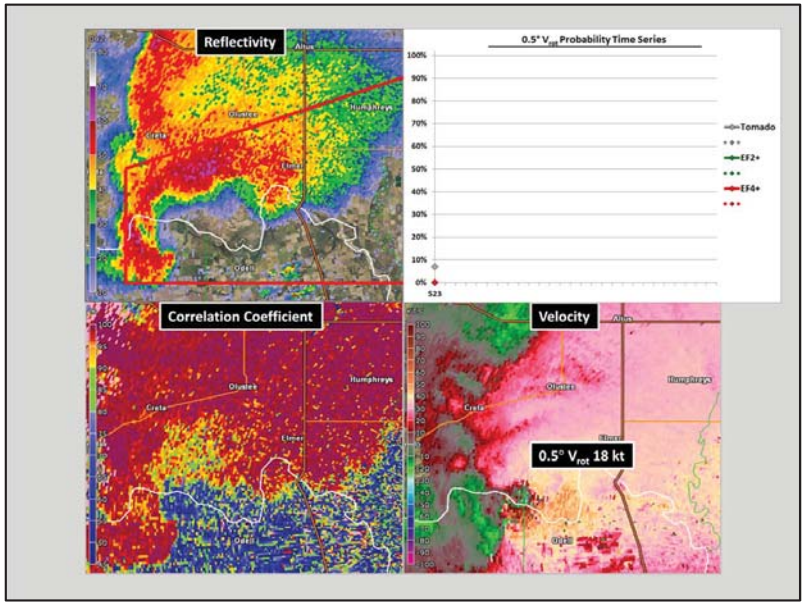
Some best practices of using pathcasts, as reiterated from previous storm-based warning training from WDTD are:

- Always keep your storm track vector accurate (this keeps your threat motion attributes in the best possible location)
- Issue frequent SVS updates (keep trimming back unnecessary areas)
- Configure your WarnGen customization files for the best local configuration of cities and towns to be included along a track.
- And finally, when in doubt about time of arrival, leave it out.

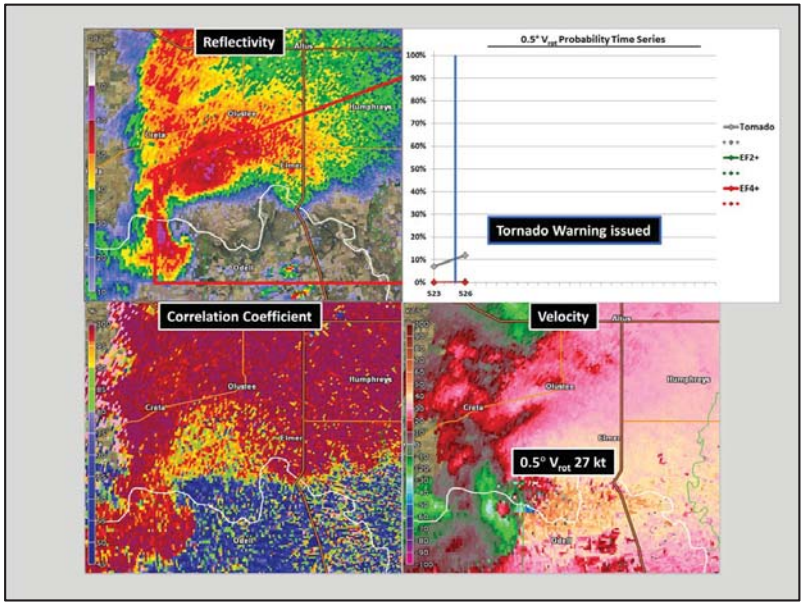


Finally, we're going to demonstrate a way in which the probability information can be applied. Only a few snapshots in the evolution of a tornadic storm are going to be shown for training purposes. Obviously more information would be available in a real-time setting and we recognize that the warning meteorologist is considering more than 0.5° reflectivity/velocity/CC data. Focus on the application of the probabilities and the resulting decision-making process. Keep in mind this is a basic example to illustrate how to apply the underlying concepts.

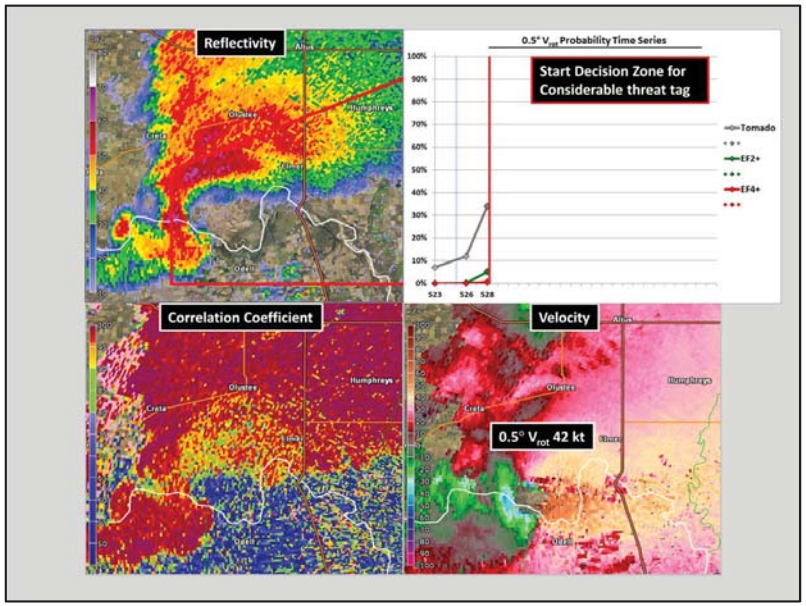
In this case, the STP is rather low and the probability of an EF2 or greater occurring in the region is quite low. There is an isolated supercell in southwest Oklahoma that will be the focus of the IBW decision. Distance to the radar is less than 40 miles, so 0.5 Vr relationships to EF-Scale potential are pretty reliable. Note there is already a SVR issued for this storm based on a detection of a severe updraft.



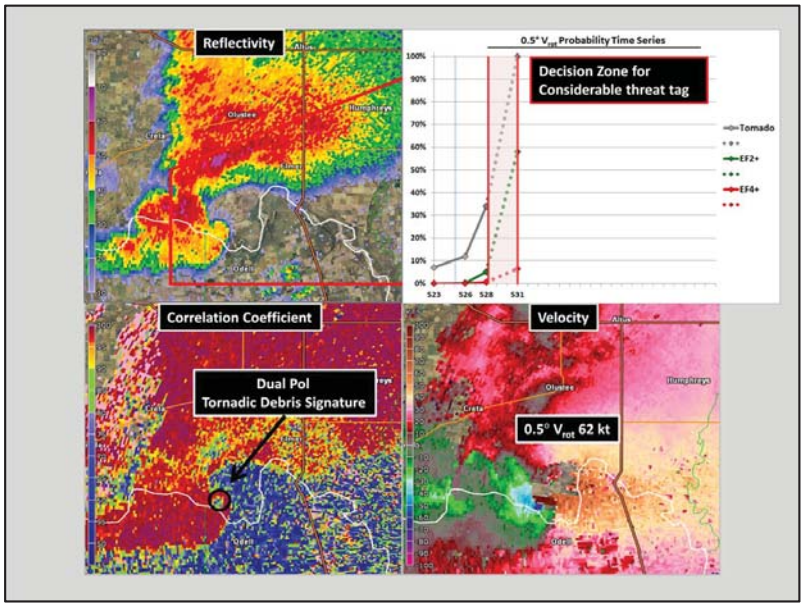
At 523 PM probabilities are very low and low level Vrot is pretty paltry. Nonetheless, strengthening was anticipated and a tornado warning was issued at 525.



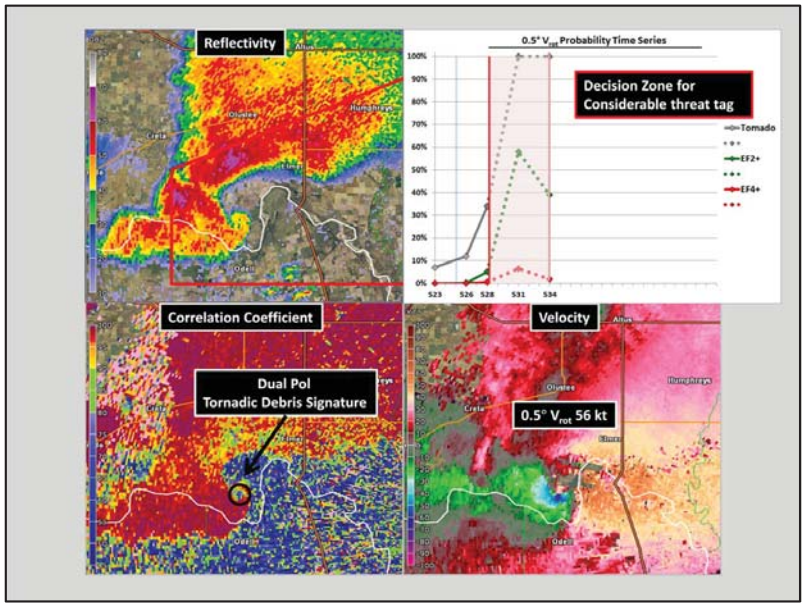
At 525 PM the tornado warning is issued and by 526 pm Vrot was 27 knots and probabilities have risen slightly but are still pretty low.



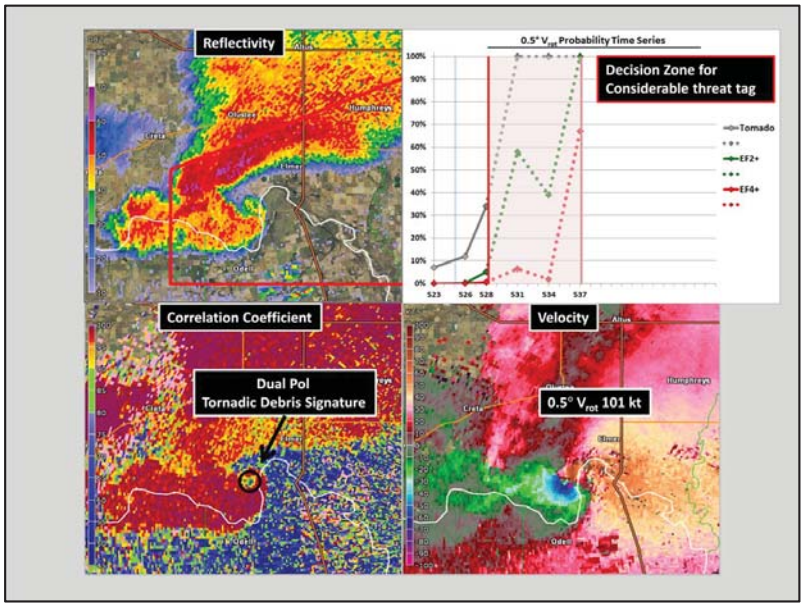
At this point, full attention should be paid toward the potential need for a considerable tag. Vrot is 42 knots and probability of a tornado is 35% and probability of a stronger tornado is 5%. Things are trending upward though and you should be anticipating very closely.



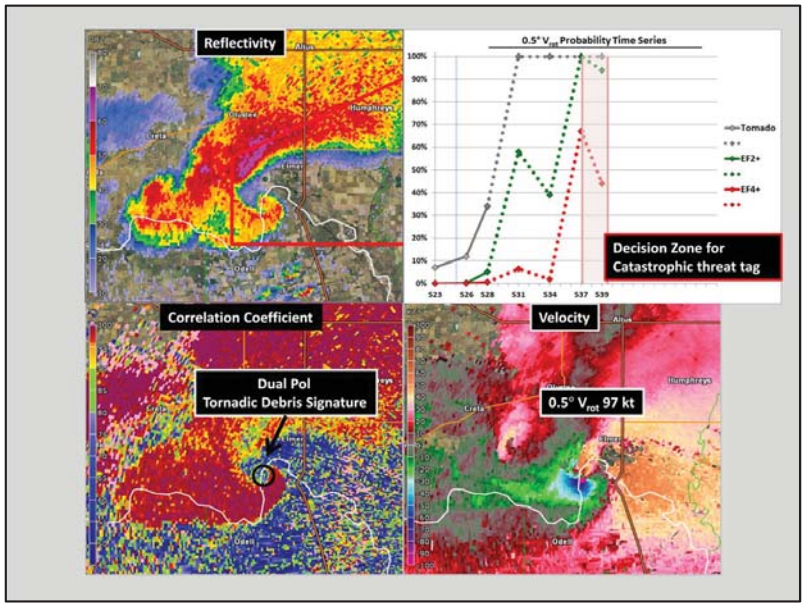
At the 531 PM CDT volume scan, the first sign of a CC minima is evident, at which point probability source should switch to the conditional probability graph. Vrot has risen substantially and quickly to 62 knots and the probability of a strong tornado has climbed to 60%. At this point the most likely outcome is an EF2 or stronger tornado. If you haven't already done so, a considerable tag upgrade is a likely choice.



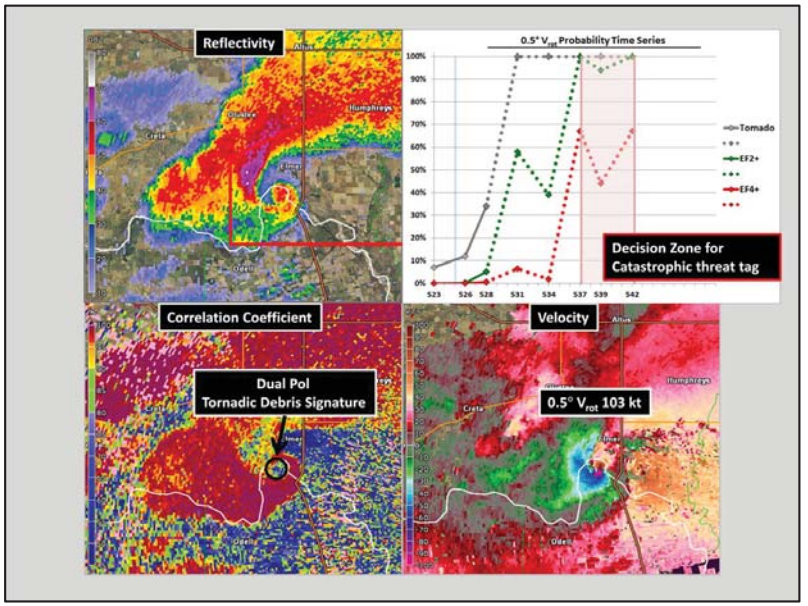
At the 534 PM CDT volume scan, TDS signature continues with Vrot at 56 kt. If you haven't already done so, once again, the considerable tag upgrade is a likely choice.



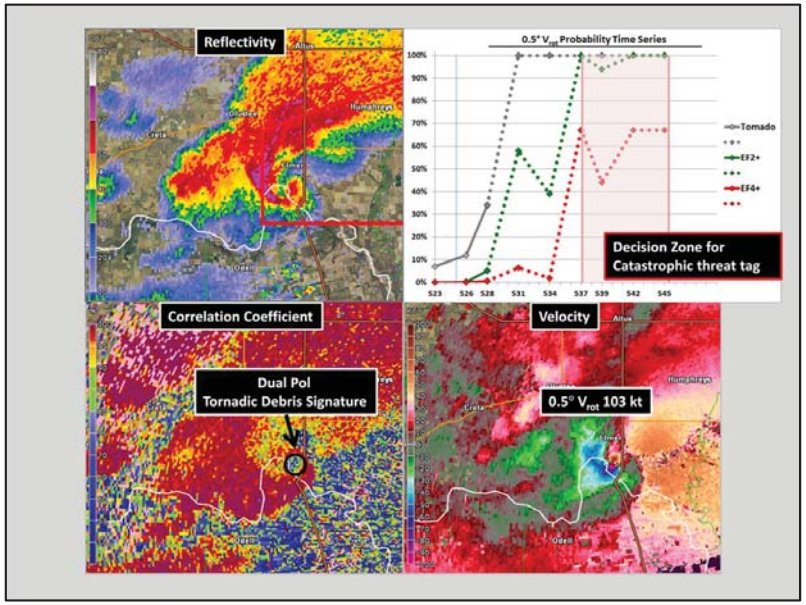
At the 537 PM CDT volume scan, TDS signature continues with Vrot now up to 101 kts. If you haven't already done so, once again, the considerable tag upgrade is a likely choice.



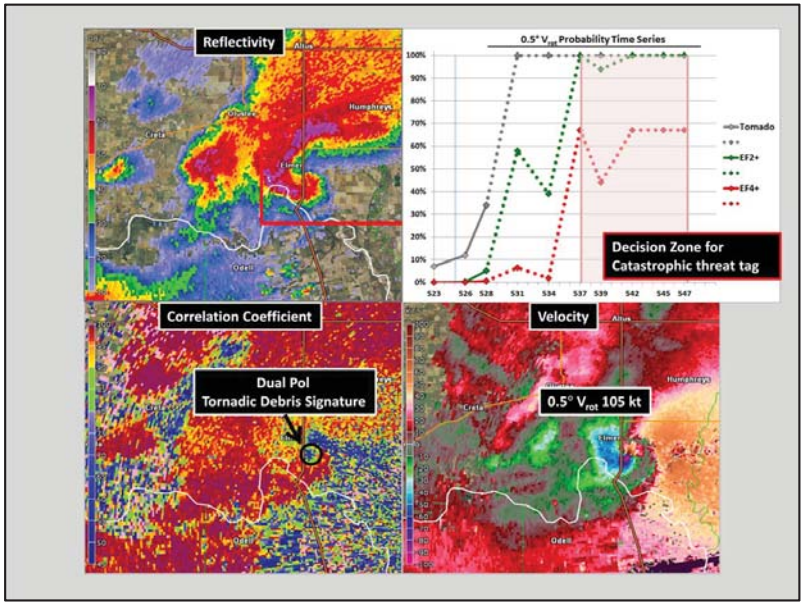
At the 539 PM CDT volume scan, TDS signature continues with Vrot now at 97 kts. As the tornado is approaching Elmer, you should start considering upgrading to a catastrophic threat tag.



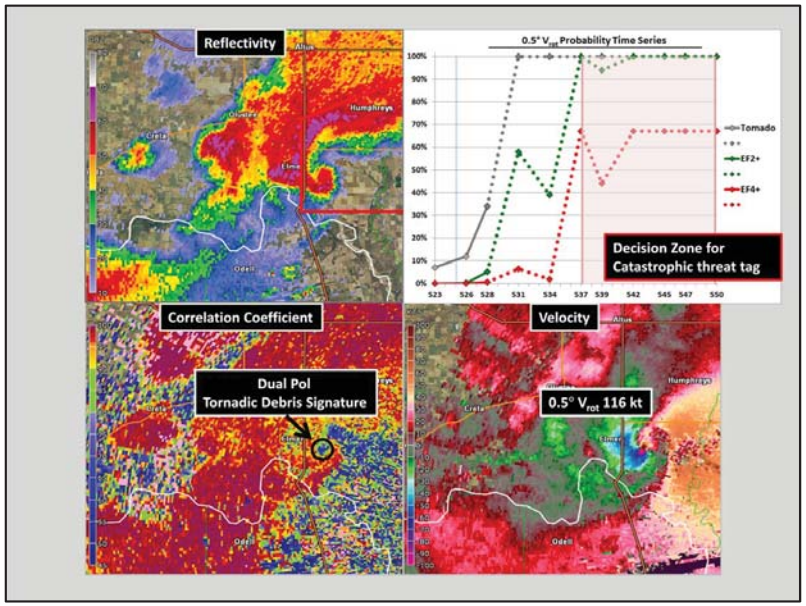
At the 542 PM CDT volume scan, TDS signature continues with Vrot now at 103 kts. As the tornado is approaching Elmer, you should start considering upgrading to a catastrophic threat tag



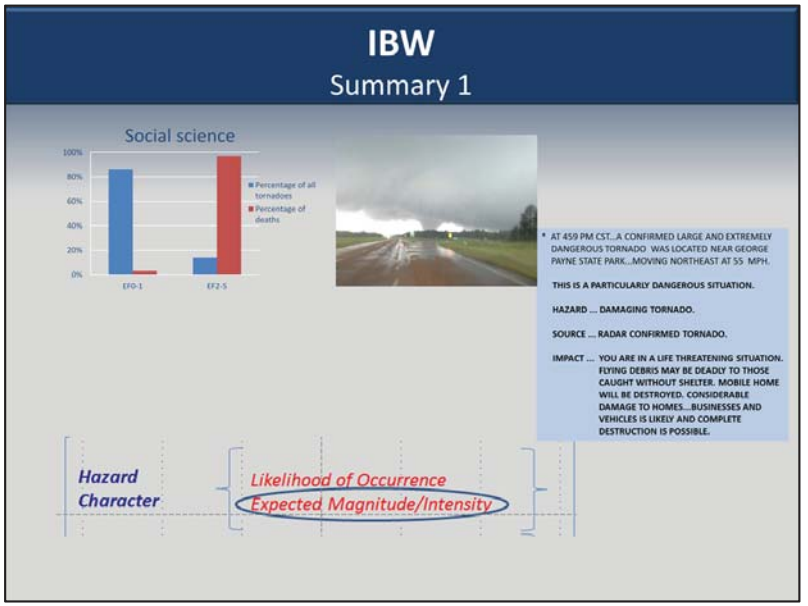
At the 545 PM CDT volume scan, TDS signature continues with Vrot still at 103 kts. As the tornado is entering Elmer, you should start considering upgrading to a catastrophic threat tag.



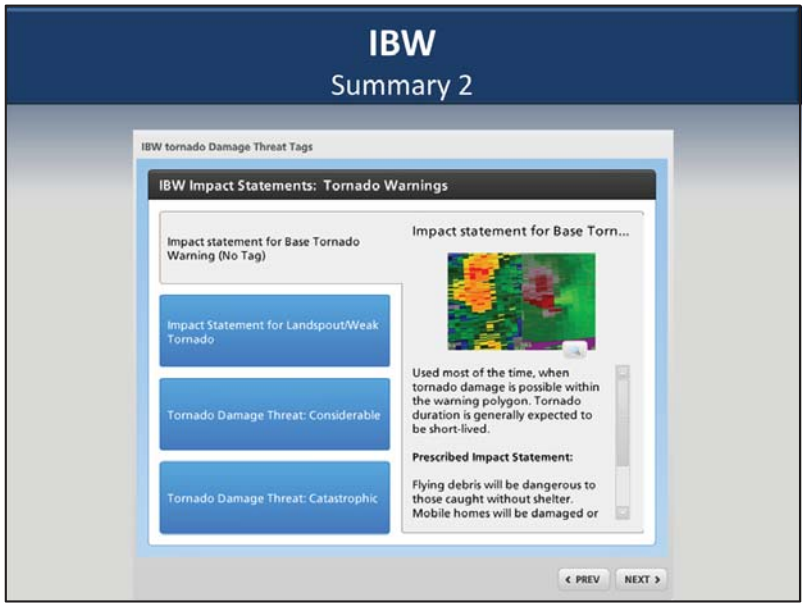
At the 547 PM CDT volume scan, TDS signature continues with Vrot up to 105 kts. As the tornado is entering Elmer, you are in the decision zone for considering a catastrophic threat tag.



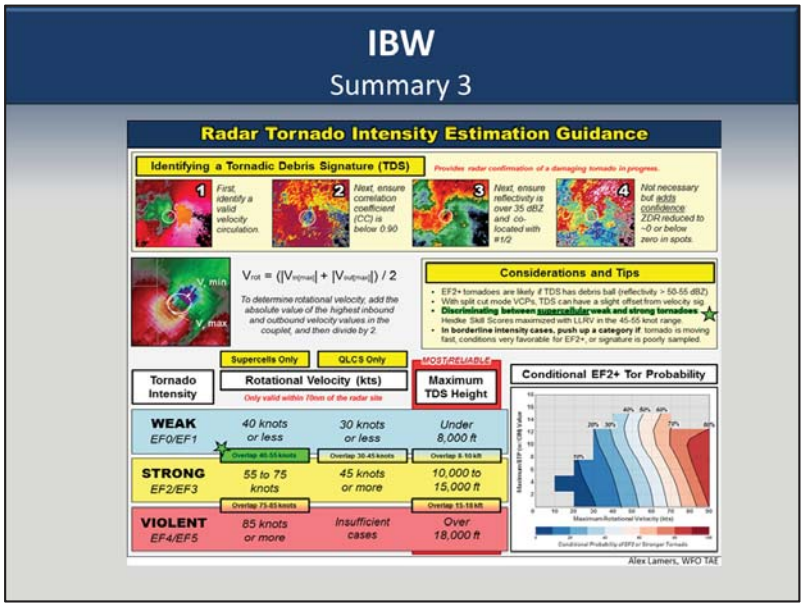
The 116 kt 0.5° V_{rot} is the 3rd highest of the super-resolution radar era (mid 2008–present). The Elmer, OK tornado was 1 mile wide and the sparse density of damage indicators yielded EF3 damage. The Norman Forecast Office in their damage assessment mentioned the tornado was likely violent (EF4+) based on video and radar presentation.



In summary, IBW provides an updated framework for providing high intensity cues for warnings that is based on social science and risk communication principals which require tornado warnings to highlight events that have the most potential to do serious harm. As part of that risk communication process, the character of a weather hazard like a tornado must have expressions of magnitude to establish a level of risk to elicit the most appropriate actions. High-intensity cues are the risk signals that prompt people to take action. These cues form the basis of the impact statements that seek to elevate threats within NWS warnings.



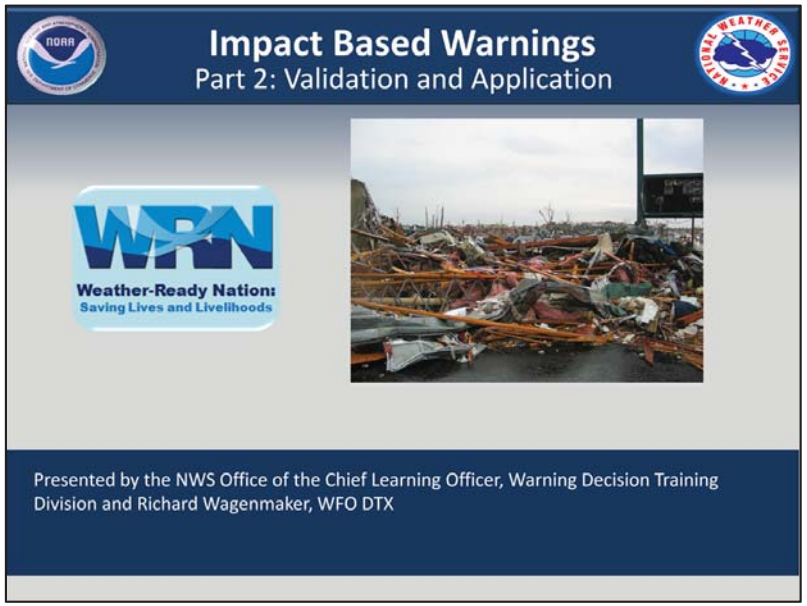
To highlight specific risks and impacts, IBW provides some options based on expected severe hazard risks and impacts. These are the impact statements and associated tags. Make the selections for the appropriate impact statement in your warnings to highlight the potential risk of damage expected.



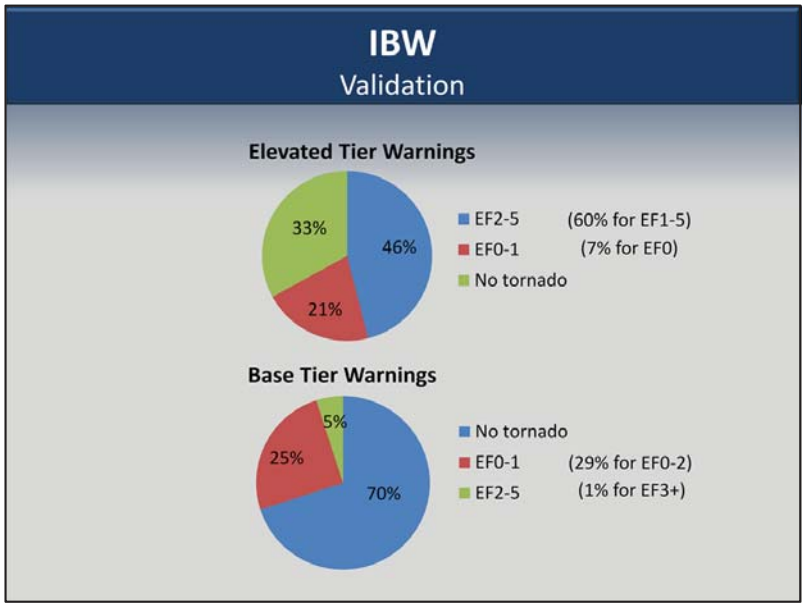
Finally, IBW is not just some half-baked policy update. The warning paradigm shift brought about by IBW is based on sound observational research which show a degree of skill in discriminating between tornado intensity. The use of STP along with other near storm environment considerations, rotational velocity, and tornado debris signatures in a probabilistic sense should be used to develop your warning methodology to better anticipate and warn for situations where the most intense and destructive tornado events can occur.

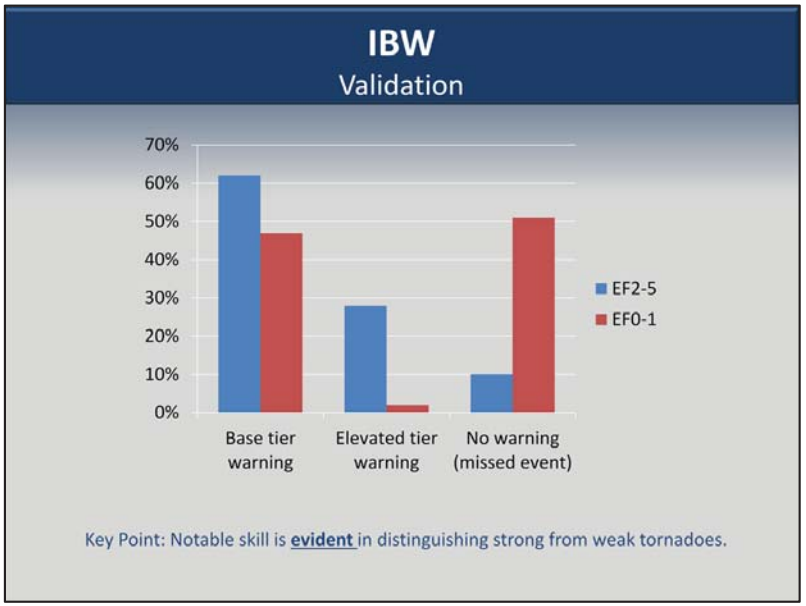


A complete list of references used in this presentation are available at http://www.wdtd.noaa.gov/courses/ibw/references.php.



Welcome to part II of the NWS course on Overview and Introduction to Impact Based Warnings. In the first part, we talked about the rationale and motivation for moving to IBW and showed you the specific impact damage statements and tags options. In this second part of the course, Dick Wagenmaker will look at the validation of IBW including some verification studies showing the relationship of occurrence of tornadoes via EF scale and tiered warnings, and the scientific research behind key environmental and radar guidelines for successfully diagnosing and anticipating the range of possibilities for tornado intensity. Finally, there will a short example illustrating how you might apply the enhanced tags in a warning situation.



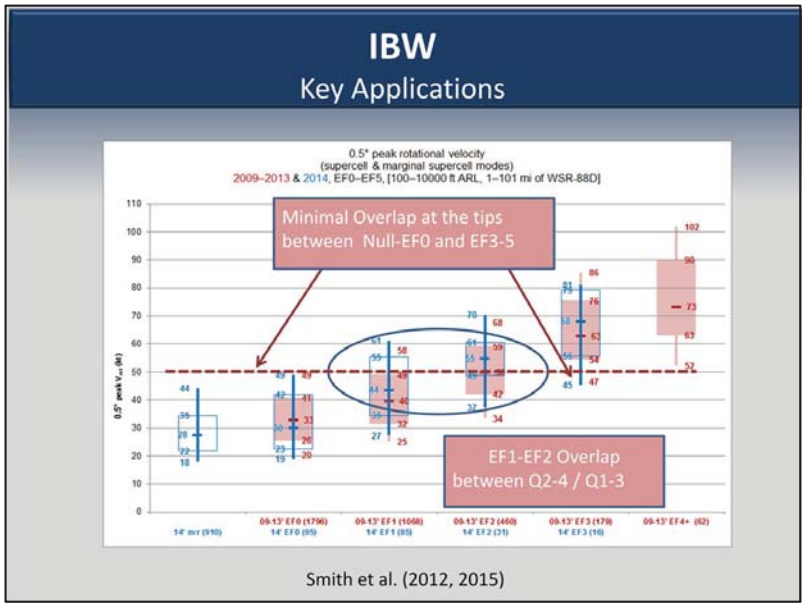


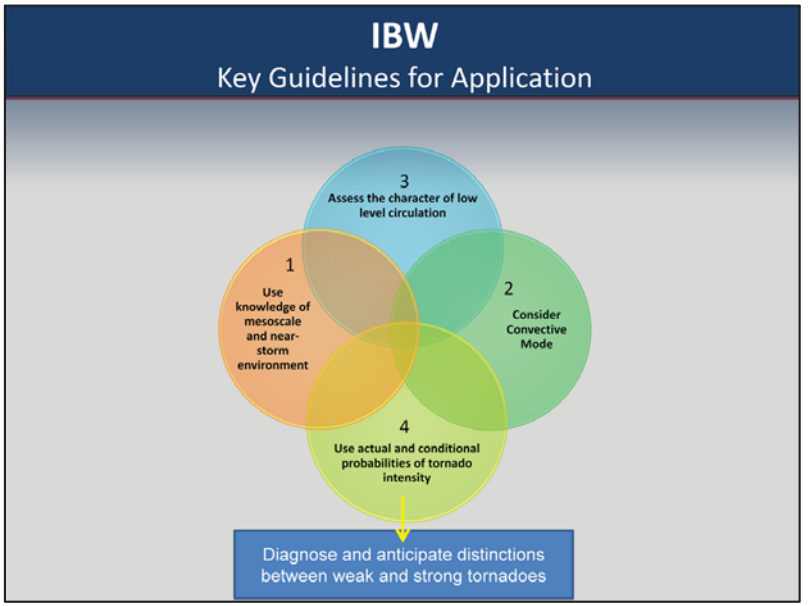
The IBW demonstration project included a verification period from 2012 to 2014 to verify both warning perspective metrics such as success ratios and false alarm ratios as well as event perspective metrics. The comparisons summarized on the next two slides here show value was added using warnings with damage threat tags (so-called "elevated tier" tornado warnings) over legacy warnings, or the so-called "base tier" warnings. In particular, the project showed that since the NWS warns for nearly all EF3-5's and can use elevated tier tags for half of them, the FAR is less than half than that for the base tier warnings. For example, the most likely outcome when an Elevated Tier Tornado Warning was issued for EF2-5 tornado occurrence was 46% (60% for EF1-5 occurrence), 21% for EF0-1 (7% for EF0 occurrence), and 33% for no tornado (i.e., a False Alarm). Contrast those statistics showing the most likely outcome when a Base Tier Warning is issued: 70% no tornado outcome, 25% of an EF0-1, and only 5% of an EF2 or greater. What does this data all mean? Well, for most base tier warnings issued, the most likely outcome is a false alarm, by a 70% majority. This is compared to the 33% false alarm ratio for elevated tier warnings. In short, you can see that false alarms in NWS warnings are largely the result of trying to warn for weak tornadoes.

From an event perspective, looking at tornado detection ratios and missed event ratios, when a EF2-5 tornado occurs (the blue series in the histogram) it is associated with a base tornado warning – 62% of time (51% when EF3+ event occurs), an elevated tier warning – 28% of the time (47% when EF3+ event occurs), and no warning (Missed Event), 10% of the time (2% when EF3+ event occurs). On the other hand, when a EF0-1 tornado occurs (red bar graphs) , it is associated with no warnings 51% of the time (54% when EF0 event occurs), 47% of the time with a base tornado warning (45% when EF0 event occurs), and only 2% with an elevated tier warning (1% when EF0 event occurs).

What does this all mean? Approximately 90% of EF2-5 tornadoes are warned with a warning of any tier. This is also true for 98% of EF3-5 tornadoes. Roughly half of EF0-1 tornadoes are warned and most of those are warned with base tier warnings. The verification also found that forecasters underutilize enhanced tags and/or issue them late. While the vast majority of strong tornadoes are warned, they are mostly covered with base tier warnings and only 28% by elevated tier warnings. This increases to 47% of EF3-5 tornadoes. While 13% of all tornadoes are of the strong variety, elevated tags are included in just 5% of all warnings.

Finally, skill in specifically predicting EF2-5 and EF0-1s independently about the same as distinguishing a tornado from no tornado. However, when using near misses (that is, partial credit for EF1s in elevated tier warnings and for EF2s in base tier warnings) notable skill is evident in broadly distinguishing strong from weak tornadoes.





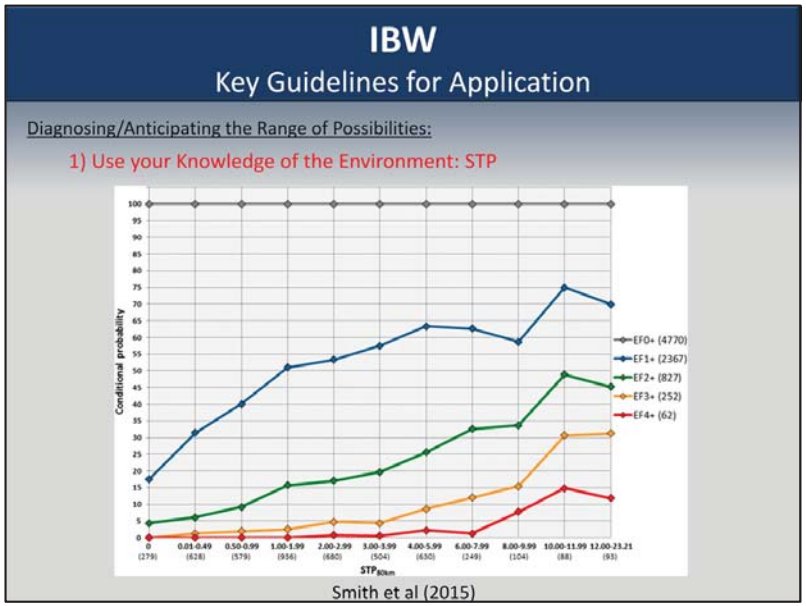
This SPC relational climatology by Smith et al from 2012, 2015 shows a significant relationship between increasing maximum 0.5 degree rotational velocity and increasing maximum EF scale of occurring tornadoes. The box and whiskers chart shown is for max low level rotational velocity signatures below 10000 feet and for supercell and marginal supercell convective modes. The light red is data from 2009-2013, while the blue is data from 2014 and includes rotational velocities from non-tornadic severe supercells.

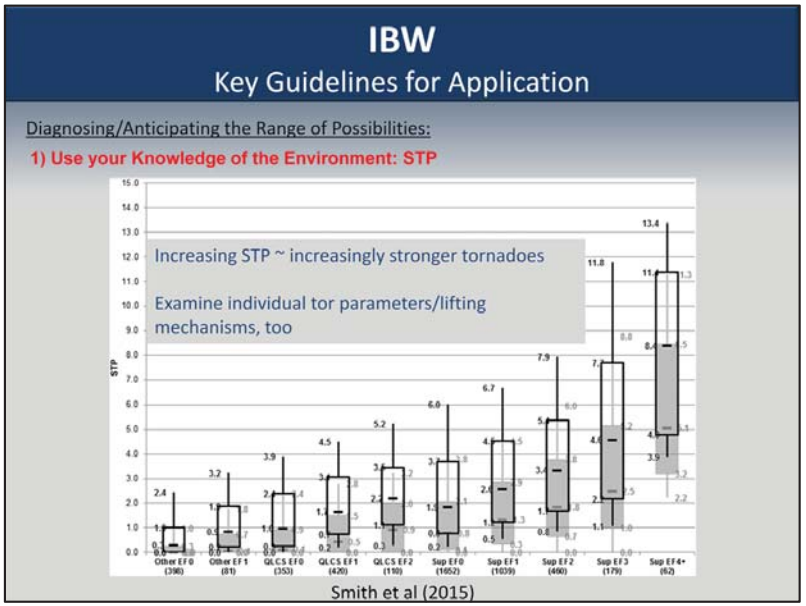
The box and whiskers have standard configurations with the box bounded by the bottom of the 2nd quartile and top of the 3rd quartile - and the tips of the whiskers represent the 10th and 90th percentiles. The dash in the middle of each box is the median or top of the 2nd quartile. It is important to note that for each event the study recorded rotational velocities immediately prior to tornado touchdown until just prior to tornado dissipation.

Again, we are not trying to pinpoint tornado intensity by EF scale – just "ring the bell" a little louder for more significant tornado events. This chart shows why. You'll note there is plenty of overlap in max low level rotational velocity associated with high end EF1 and low end EF2 – but you can also see there is little overlap between null events/EF0's and EF3-5's. Again, it won't be unusual for an EF1 to occur on Elevated Tier Warnings, nor unusual for EF2's to occur on Base Tier Warnings. We should avoid as much as possible having EF3's or greater occur on Base Warnings or No Warning, and avoid EF0's and Null events occur on Elevated Tier warnings. The graphic here certainly hints at some capacity for distinguishing between weak and strong tornadoes in that respect… and also hints at the viability of probabilistic approaches to the IBW warning process.

There is no perfect answer to the question of when to use a "considerable" tag, but that is rarely the case for anything in operational meteorology. Keep in mind that the value of your role in the warning decision process comes into play by staying situationally aware and considering a wide variety of factors to stay one step ahead of the tornado threat.

At this point we'll lay out some basic guidelines for distinguishing weak tornadoes from strong tornadoes. Essentially, there are 4 steps in the process which is diagnosing and anticipating the range of possibilities. First, use your situational awareness to assess the mesoscale and near-storm environments. Second, use your knowledge and understanding of convective modes and storm evolution as they relate to the environment. Third, use your understanding of the 4-dimensional character of radar-depicted mesocyclone circulations, especially the strength of low level rotational velocity, and 4thly, use your understanding of actual and conditional probabilities of tornado intensity (null vs. weak vs. strong) as related to low-level rotational velocity.





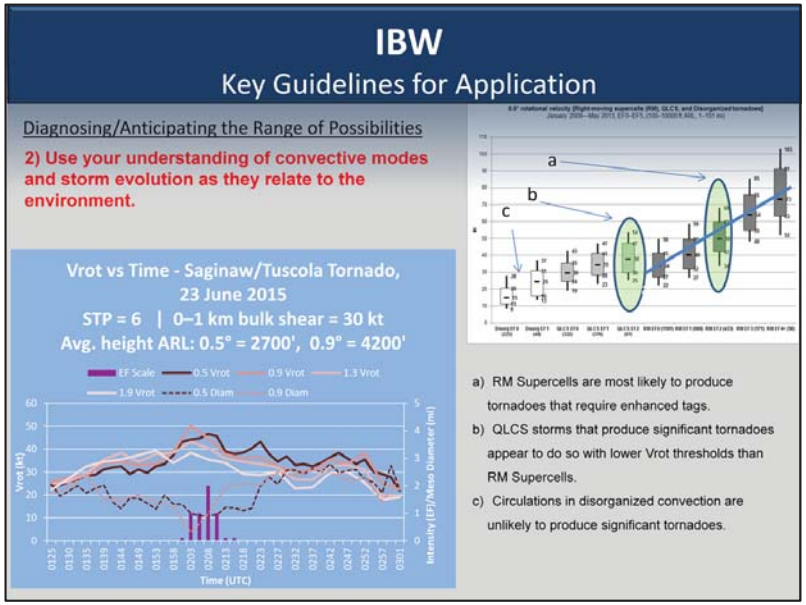
We'll first look at environmental parameters that serve as a baseline for determining the likelihood that strong tornadoes may occur in a given situation. The Significant Tornado Parameter (STP) 80 km is a multiple component parameter meant to highlight co-existence of low-level CAPE and shear which are crucial ingredients for right moving supercells. Based on past studies, STP exhibits greater skill in discriminating between nontornadic and significantly tornadic supercell environments compared to any of its individual components or any other parameters among the 38-variable database at the SPC.

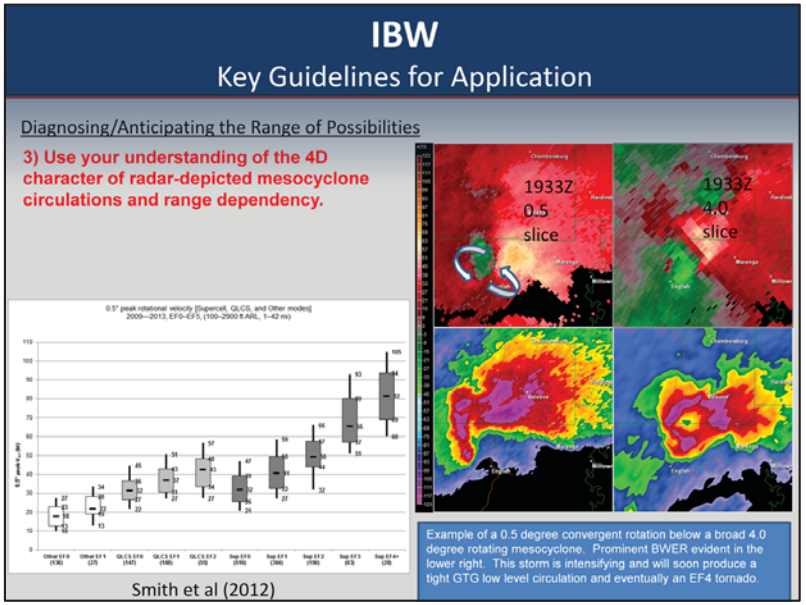
There is a relational climatology between increasing **conditional tornado intensity** and increasing values of STP 80km as measured on hourly SPC mesoscale objective analysis. Results from examining environmental and radar attributes, Smith et al (2015) found that increasing conditional probability for greater EF-scale damage, both STP and 0.5° peak Vrot increase, especially with supercells. This figure shows conditional probability of meeting or exceeding a given EF-scale rating (for the 5 series shown in the legend) for STP 80km for all convective mode tornado events from 2009–13.

The conditional probabilities may give you an idea ahead of time on how aggressive you can be on potential considerable tornado tags during the warning process.

Next to further illustrate the relationship, here is a box and whiskers chart showing effective-layer STP (a dimensionless number) for all supercell, QLCS EF0–EF2 tornadoes, and other modes EF0 and EF1 tornadoes by EF-scale damage rating classes. The 40-km grid data are shaded gray, labels on right. The black overlays (with labels on left) denote STP 80 km grid point values, at the analysis time immediately preceding the event time. Using either the nearest STP grid point value or the neighborhood maximum value (which is the STP 80km) STP increases as tornado damage classifications increase. Supercell events tended to exhibit higher STP values than QLCS and other modes for the same EF-Scale damage rating. The STP for supercell, QLCS, and other modes tended to increase monotonically with increasing damage class ratings (aside from the 10th percentile). Substantial overlap exists in the distributions between adjacent EF-scale ratings, though the higher values of STP 80km (i.e., > 6) are more common for a greater proportion of supercell events at higher EF-scale rating classes (i.e., EF3+).

It must be stressed that composite parameters such as the STP 80 km should not be examined alone, but rather in concert with the individual components in the STP that identify important supercell tornado ingredients. Despite STP utility as a composite tornado predictor, there is no replacement for a thorough diagnosis of the mesoscale environment. You have to be aware of rapid storm interactions due to low-level boundary interactions which produce rapid tilting and/or stretching of local vorticity maxima. Also, closely monitor 0-1 km effective bulk shear, subsequent updraft helicity, and LCL heights < 1500 m AGL which are all important ingredients in the tornado development process.



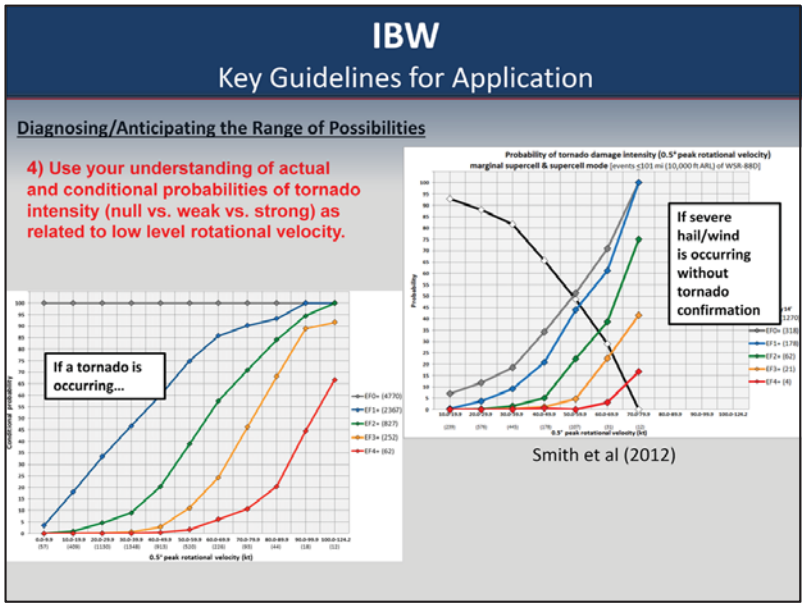


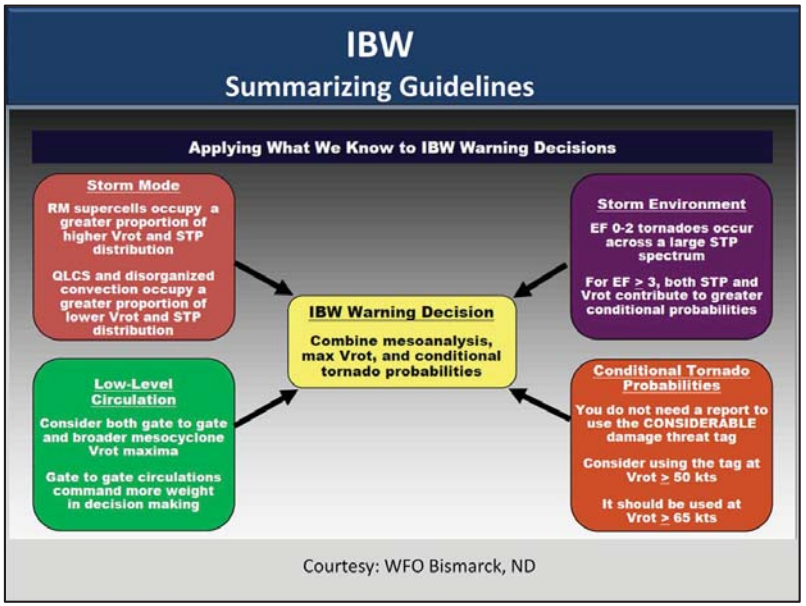
Here we'll discuss convective modes and rotational velocity evolution, or Vrot evolution. The graphic on the left is a marginal supercell tornado example from a study by Frey and Thompson in 2015. It shows the time evolution of Vr at various elevation slices through a storm. It also shows the time evolution of mesocyclone diameter through the storm cycle. On this graphic you can see a broad peak in 0.5 degree Vr around 45 knots starting near the time of tornadogenesis and a 0.9 degree peak near 50 knots right before a brief EF2 occurrence. Also note the max Vr starts at higher elevation slices 1.3 and 1.9 degrees but shifts to the lower elevation slices just prior to tornadogenesis. Additionally note how the meso diameter tightens as we approach tornadogenesis. This is a nice presentation of how you can view the full time evolution of a circulation to help make a determination of tornado development and potential max intensity.

On the right side of the slide is the Smith et al. study breaking down max Vrot distributions vs. Max EF-Scale for each convective mode. Obviously, supercell modes have the strongest relationship and most strong tornadoes occur with these modes. On the opposite end of the spectrum, it is highly unlikely to have a strong tornado occur with weak disorganized convection. Last, it is interesting to note that strong tornadoes can occur with QLCS modes – but QLCS storms that do produce significant tornadoes appear to do so with lower Vrot values than RM Supercells. This possibly due to enhanced forward motion vector contributions on right flanks of low level circulations.

Thirdly, you must use your understanding of the 4-dimensional character of the radar-depicted mesocyclone and range dependency. On the right hand side of the screen is an example of how you should anticipate how convergent low level circulations will behave given the near-storm environment. The Smith et al study uses both broad Vrot maxima and Gate-to-Gate Vrot maxima, depending on which is strongest for a given case. It goes without saying that a Gate-to-Gate Vrot maxima should operationally command more weight and a lower Vrot probability threshold for EF2+ events. This example shows a 0.5 degree convergent rotation below a broad 4.0 degree rotating mesocyclone. Note the prominent BWER signature evident in the lower right. This storm is intensifying and will soon produce a tight GTG low level circulation and eventually an EF4 tornado.

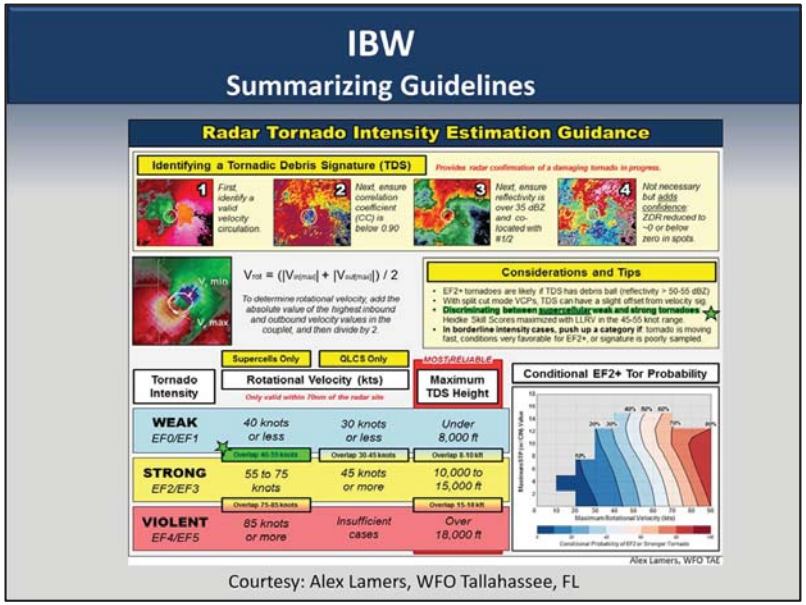
On the left is the range dependency of the dataset. The relationship between Vr and EF scale is not as pronounced at longer ranges from the radar, and very pronounced at shorter ranges. Probabilities of tornado intensity increase as range from the radar decreases.





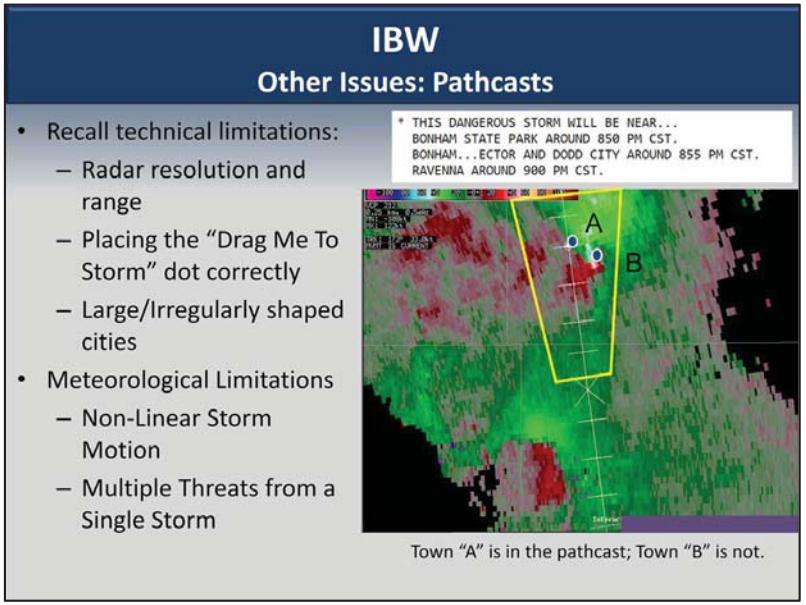
Lastly, here are two figures representing the raw probability space from the box and whiskers graphic that was previously shown. First, since there are several caveats that apply to the dataset, these probabilities should be considered ESTIMATES of tornado damage intensity as related to max 0.5 degree rotational velocity. Second, these data are only for supercells and marginal supercells. Data from most QLCS and non-mesocyclone tornadoes are not factored into the probability calculations. There is range dependency not accounted for in these graphs, which we'll explore more in the next slides. There is no filtering of the data to account for likely underestimation of EF scale that might be associated with tornadoes in areas sparsely populated with damage indicators, such as in the high plains. Time evolution of rotational velocity is not accounted for. Instead, the climatology relates max EF-Scale to max Rotational Velocity only. Circulation diameter is loosely accounted for, but specific distinctions between broader signatures and gate-to-gate signatures are not fully accounted for. And last, the null sample only includes events corresponding to significant severe events for large hail and winds – and this data was collected only from 2014. Despite these caveats, the probability estimates contained here can be very useful if applied generally and in conjunction with other tools. In the upper right are the actual probabilities of EF-Scale supercell tornado intensity while on the lower left are conditional probabilities of tornado intensity (provided a tornado is occurring). As mentioned you can use each of these to help diagnose the need to issue a "considerable" tag while also considering range dependency, the character of the low level circulation, the favorability of the ambient environment, and time evolution of the supercell. These continue to be refined as more data is collected. An interesting observation between the box and whiskers chart previously shown and the probability chart here on the right is that distinguishing between null severe events and EF0 tornado events using rotational velocity is nearly impossible. The probability distribution for null events vs EF0's is almost completely driven by the comparatively high volume of null events. In fact, once the probability of a tornado reaches 50%, the most likely intensity outcome quickly approaches EF2 or greater.

Here is a one-pager summarizing the various applications to IBW warning decisions, looking at environment, mode, circulation strength, and conditional probabilities. This summary comes from WFO BIS.



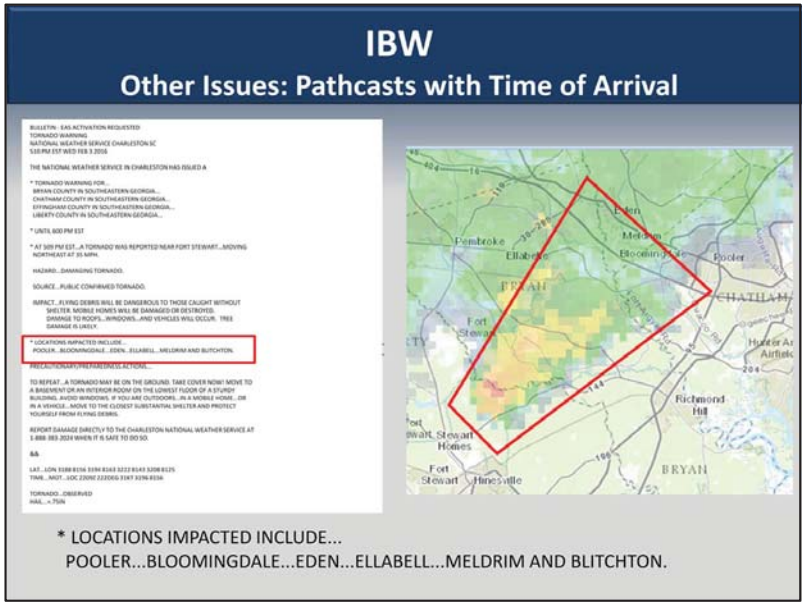
Courtesy: Alex Lamers, WFO Tallahassee, FL

Here's another one-page summary from SR that's provides the most recent methodology of evaluating Tornado intensity using radar only signatures including TDS height. All of the considerations and tips are based on peer-reviewed papers. Notice the categories of Weak, Strong, and Violent Tornadoes somewhat coincide with the selection of no damage tag, considerable, and catastrophic.



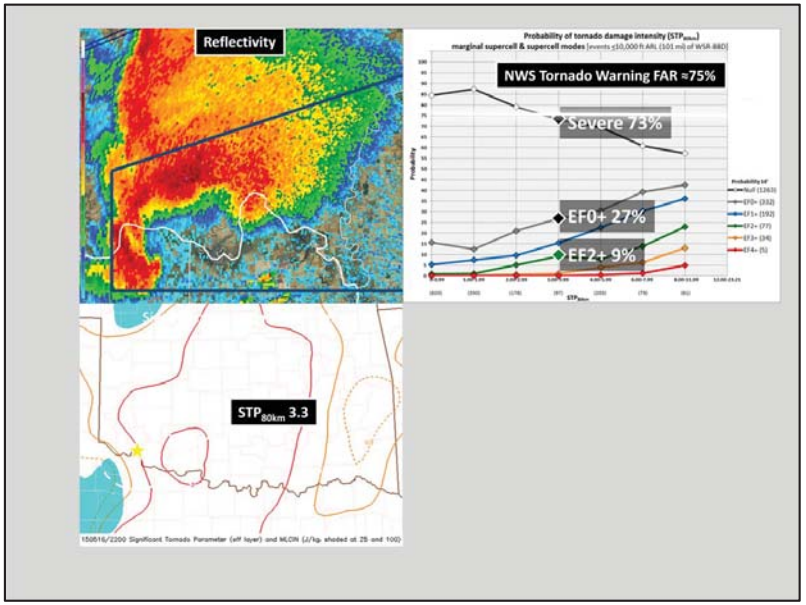
Town "A" is in the pathcast; Town "B" is not.

Since pathcasts are turned on with the default setting in WarnGen IBW templates, you will need to be very careful with using the time of arrival option version. Recall the technological limitations due to radar resolution and range, issues when you don't place the "Drag Me To Storm" dot correctly, and of course, all the problems caused in the shapefiles due to large and/or irregularly shaped cities. In addition, there are known meteorological limitations with using specific time of arrival locations in warnings such as non-linear storm motion, and conveying multiple threats from a single storm. Please review some of the specific best practices regarding the use of pathcasts at wdtd.noaa.gov.



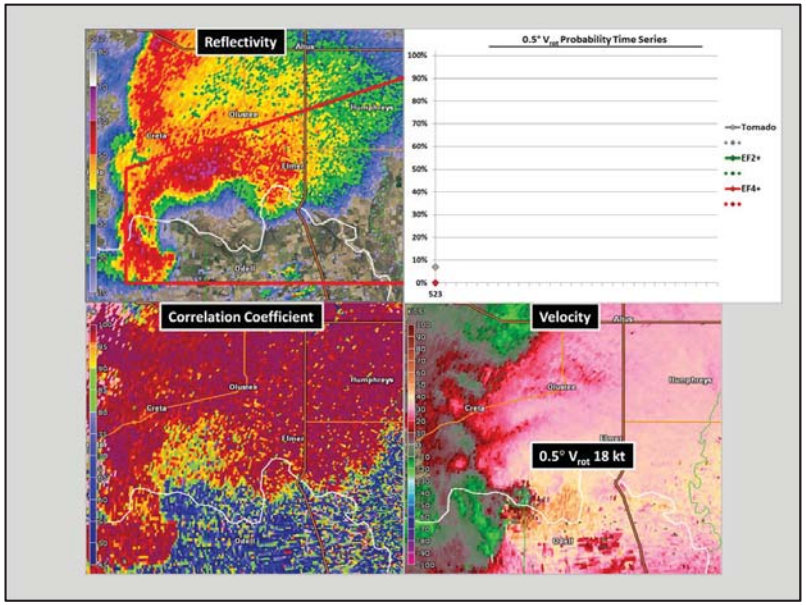
Some best practices of using pathcasts, as reiterated from previous storm-based warning training from WDTD are:

- Always keep your storm track vector accurate (this keeps your threat motion attributes in the best possible location)
- Issue frequent SVS updates (keep trimming back unnecessary areas)
- Configure your WarnGen customization files for the best local configuration of cities and towns to be included along a track.
- And finally, when in doubt about time of arrival, leave it out.

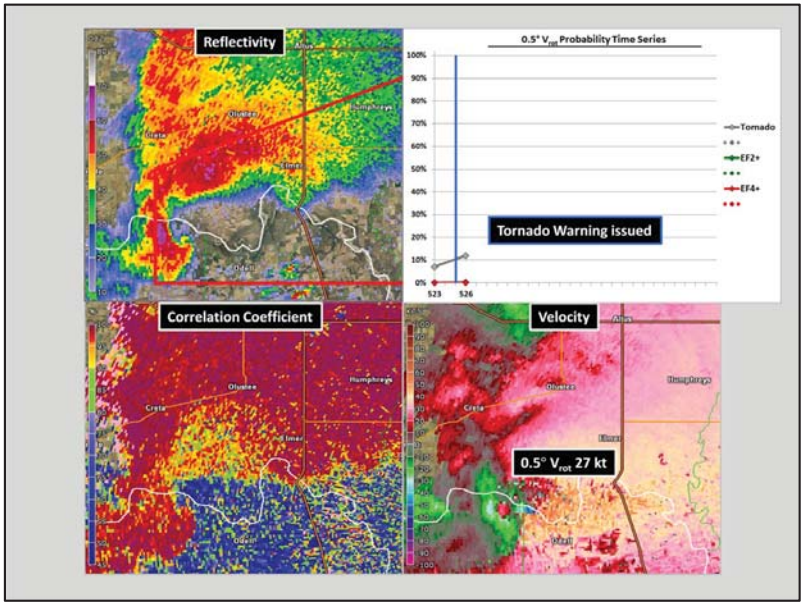


Finally, we're going to demonstrate a way in which the probability information can be applied. Only a few snapshots in the evolution of a tornadic storm are going to be shown for training purposes. Obviously more information would be available in a real-time setting and we recognize that the warning meteorologist is considering more than 0.5° reflectivity/velocity/CC data. Focus on the application of the probabilities and the resulting decision-making process. Keep in mind this is a basic example to illustrate how to apply the underlying concepts.

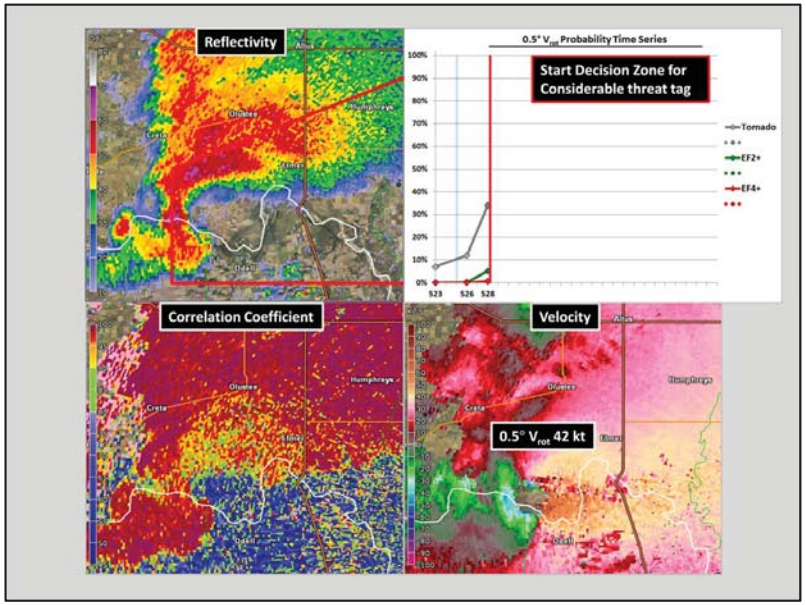
In this case, the STP is rather low and the probability of an EF2 or greater occurring in the region is quite low. There is an isolated supercell in southwest Oklahoma that will be the focus of the IBW decision. Distance to the radar is less than 40 miles, so 0.5 Vr relationships to EF-Scale potential are pretty reliable. Note there is already a SVR issued for this storm based on a detection of a severe updraft.



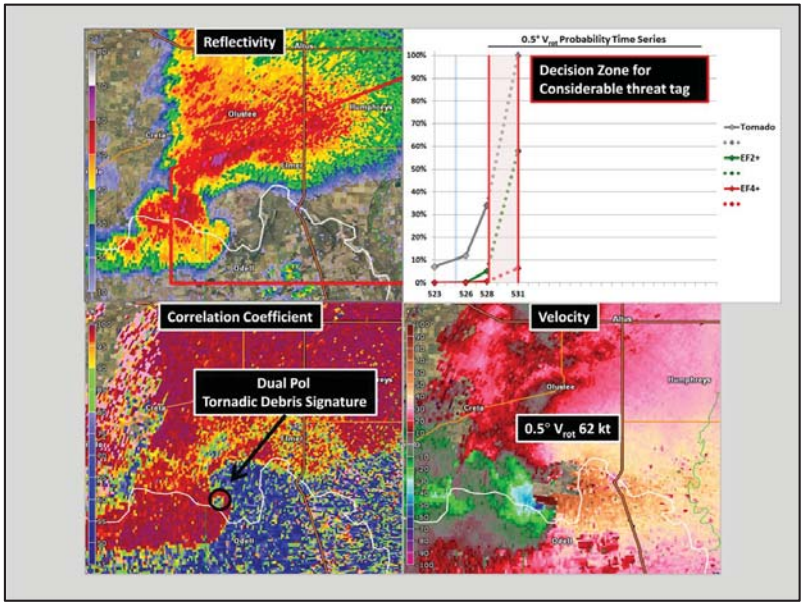
At 523 PM probabilities are very low and low level Vrot is pretty paltry. Nonetheless, strengthening was anticipated and a tornado warning was issued at 525.



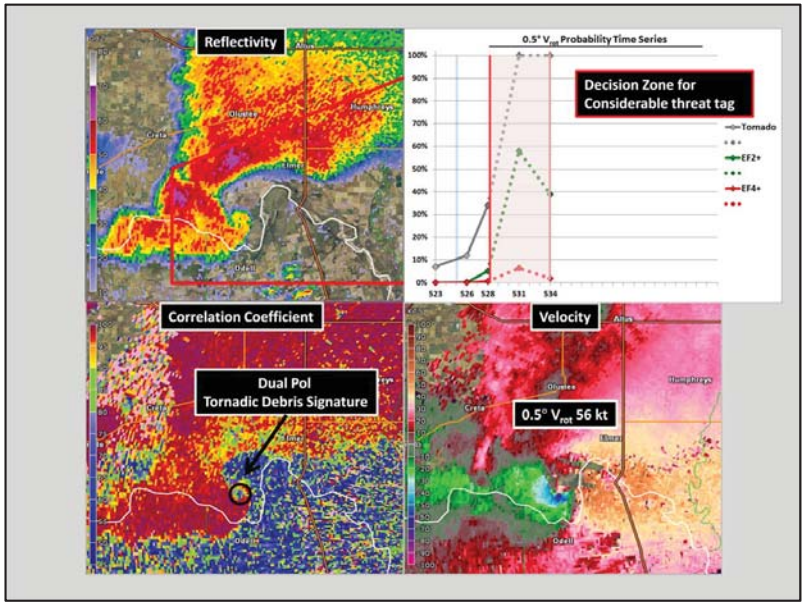
At 525 PM the tornado warning is issued and by 526 pm Vrot was 27 knots and probabilities have risen slightly but are still pretty low.



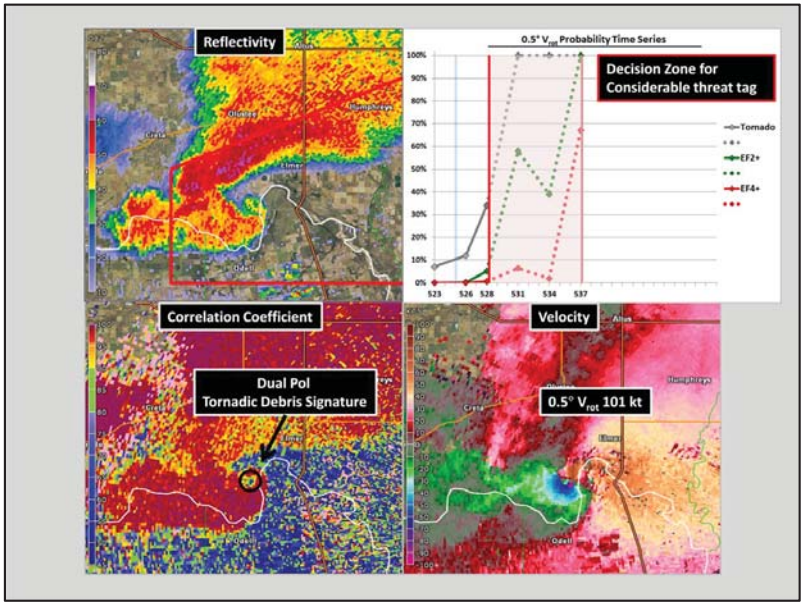
At this point, full attention should be paid toward the potential need for a considerable tag. Vrot is 42 knots and probability of a tornado is 35% and probability of a stronger tornado is 5%. Things are trending upward though and you should be anticipating very closely.



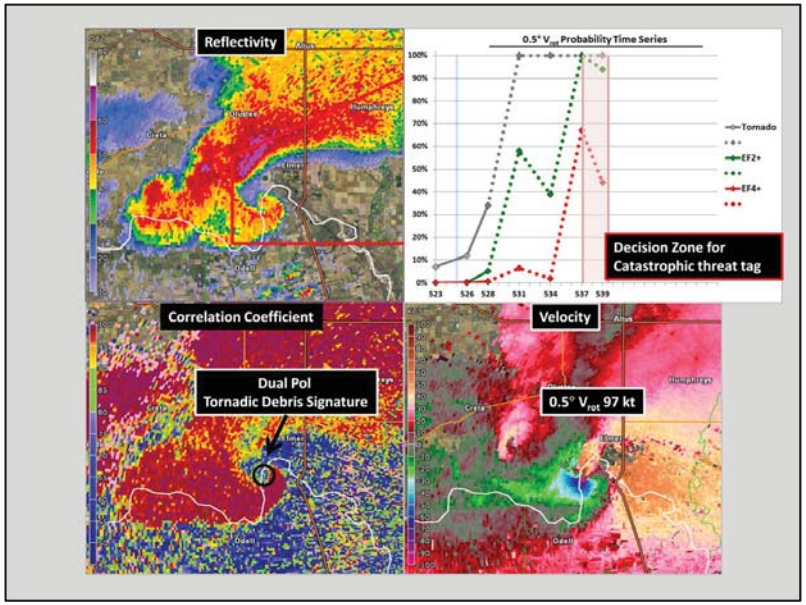
At the 531 PM CDT volume scan, the first sign of a CC minima is evident, at which point probability source should switch to the conditional probability graph. Vrot has risen substantially and quickly to 62 knots and the probability of a strong tornado has climbed to 60%. At this point the most likely outcome is an EF2 or stronger tornado. If you haven't already done so, a considerable tag upgrade is a likely choice.



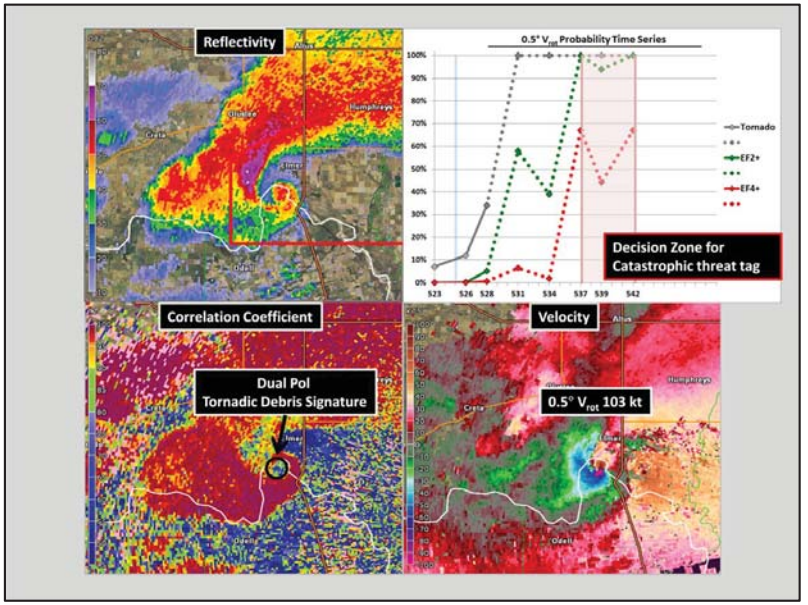
At the 534 PM CDT volume scan, TDS signature continues with Vrot at 56 kt. If you haven't already done so, once again, the considerable tag upgrade is a likely choice.



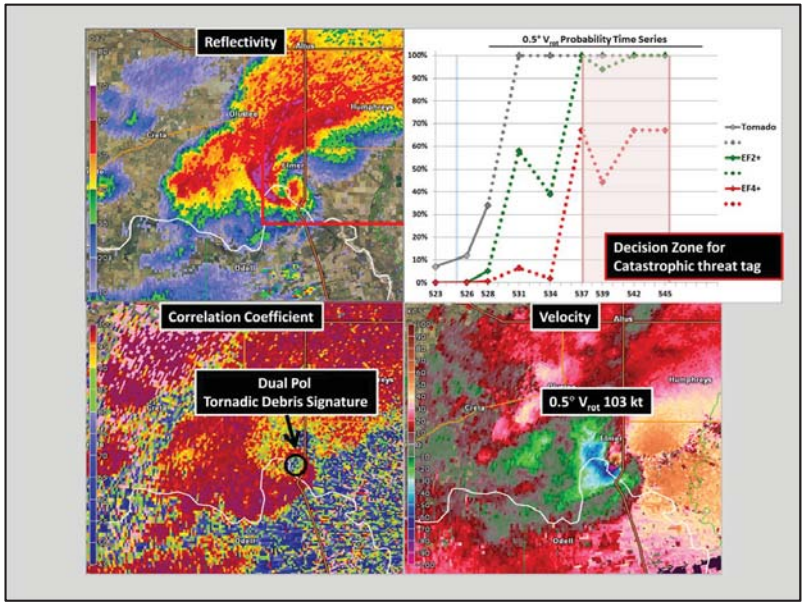
At the 537 PM CDT volume scan, TDS signature continues with Vrot now up to 101 kts. If you haven't already done so, once again, the considerable tag upgrade is a likely choice.



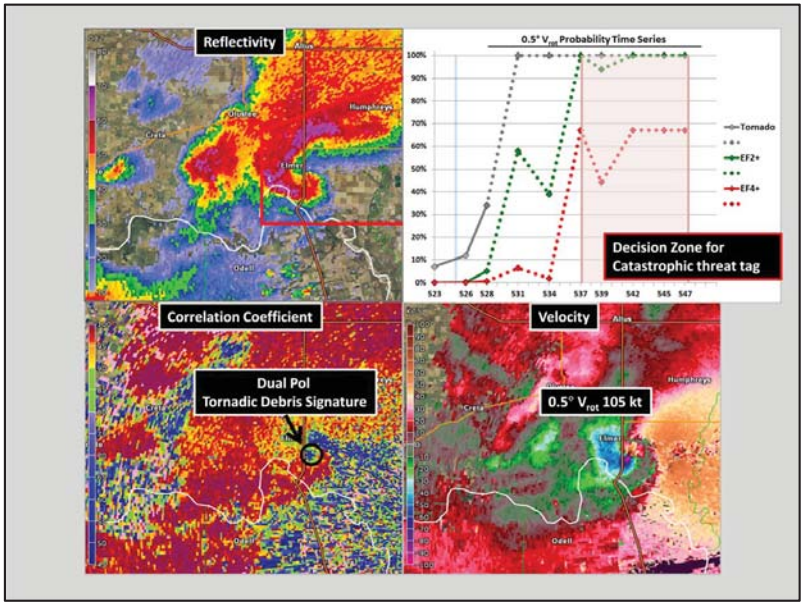
At the 539 PM CDT volume scan, TDS signature continues with Vrot now at 97 kts. As the tornado is approaching Elmer, you should start considering upgrading to a catastrophic threat tag.



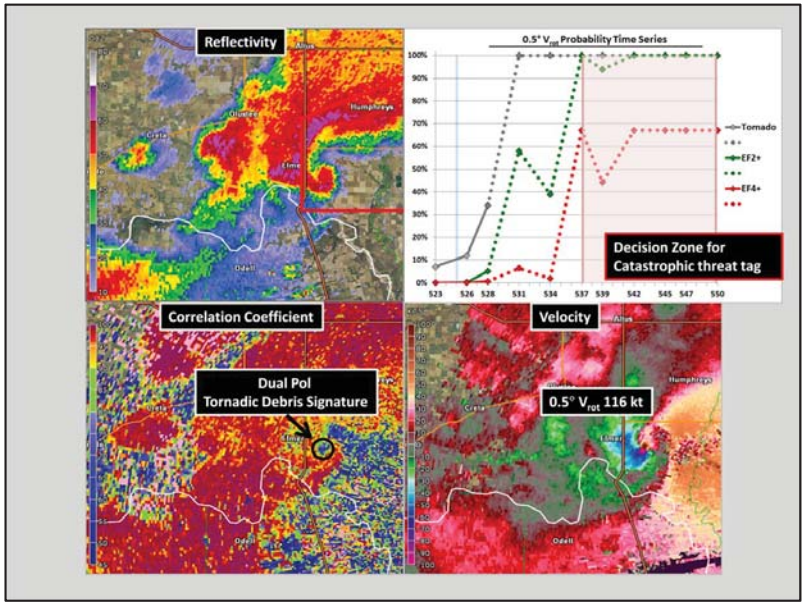
At the 542 PM CDT volume scan, TDS signature continues with Vrot now at 103 kts. As the tornado is approaching Elmer, you should start considering upgrading to a catastrophic threat tag



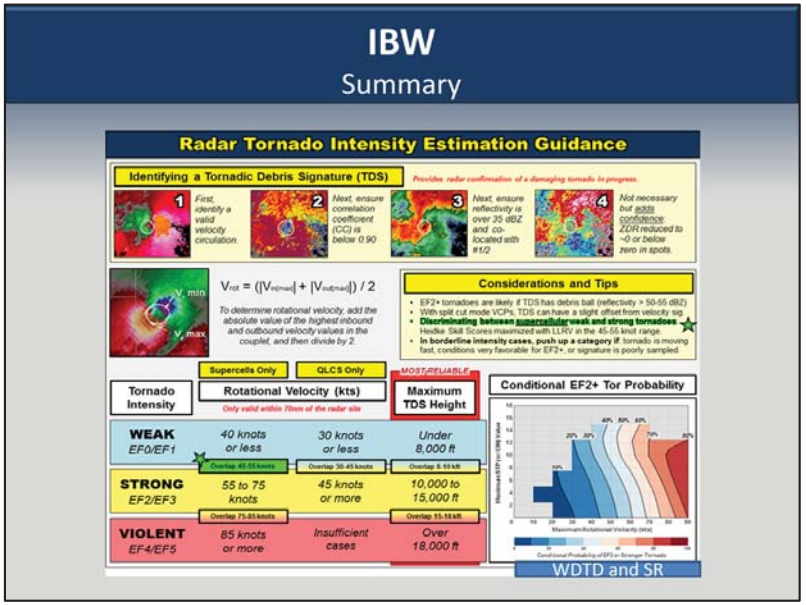
At the 545 PM CDT volume scan, TDS signature continues with Vrot still at 103 kts. As the tornado is entering Elmer, you should start considering upgrading to a catastrophic threat tag.



At the 547 PM CDT volume scan, TDS signature continues with Vrot up to 105 kts. As the tornado is entering Elmer, you are in the decision zone for considering a catastrophic threat tag.



The 116 kt 0.5° V_{rot} is the 3rd highest of the super-resolution radar era (mid 2008–present). The Elmer, OK tornado was 1 mile wide and the sparse density of damage indicators yielded EF3 damage. The Norman Forecast Office in their damage assessment mentioned the tornado was likely violent (EF4+) based on video and radar presentation.



Finally, IBW is not just some half-baked policy update. The warning paradigm shift brought about by IBW is based on sound observational research which show a degree of skill in discriminating between tornado intensity. The use of STP along with other near-storm environment considerations, rotational velocity, and tornado debris signatures in a probabilistic sense should be used to develop your warning methodology to better anticipate and warn for situations where the most intense and destructive tornado events can occur.



A complete list of references used in this presentation are available at http://www.wdtd.noaa.gov/courses/ibw/references.php.

TAB

Radar & Applications Course

Glossary

Term	Defintion
Above Radar Level (ARL)	The altitude measured with repsect to the radar.
Advanced Weather Interactive Processing System (AWIPS)	A processing, display, and telecommunications system that is the cornerstone of the United States National Weather Service's (NWS) operations.
Anafront	A front at which the warm air is ascending the frontal surface up to high altitudes. With anafronts, precipitation may occur to the rear of the front and is sometimes associated with cyclogenesis. *Compare* katafront. American Meteorological Society, cited 2013: “term.” Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/anafront]
Anvil Cloud	The flat, spreading top of a cumulonimbus cloud, often shaped like an anvil block tool. Thunderstorm anvils may spread hundreds of miles downwind from the thunderstorm itself, and sometimes may spread upwind.
Atmospheric Boundary Layer	The bottom layer of the troposphere that is in contact with the surface of the earth. It is often turbulent and is capped by a statically stable layer of air or temperature inversion. The ABL depth (i.e., the inversion height) is variable in time and space, ranging from tens of meters in strongly statically stable situations, to several kilometers in convective conditions over deserts. During fair weather over land, the ABL has a marked diurnal cycle. During daytime, a mixed layer of vigorous turbulence grows in depth, capped by a statically stable entrainment zone of intermittent turbulence. Near sunset, turbulence decays, leaving a residual layer in place of the mixed layer. During nighttime, the bottom of the residual layer is transformed into a statically stable boundary layer by contact with the radiatively cooled surface. Cumulus and stratocumulus clouds can form within the top portion of a humid ABL, while fog can form at the bottom of a stable boundary layer. The bottom 10% of the ABL is called the planetary boundary layer. American Meteorological Society, cited 2013: Atmospheric Boundary Layer. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Atmospheric_boundary_layer]
Backing Winds	Winds which shift in a counterclockwise direction with time at a given location (e.g. from southerly to southeasterly), or change direction in a counterclockwise sense with height (e.g. westerly at the surface but becoming more southerly aloft). The opposite of veering winds. National Weather Service Glossary, cited 2013.

Baroclinic Zone	A region in which a temperature gradient exists on a constant pressure surface. Baroclinic zones are favored regions for the development of extratropical cyclones. Also, wind shear is characteristic of a baroclinic zone.
Beaver('s) Tail	[Slang] A particular type of inflow band with a relatively broad, flat appearance suggestive of a beaver's tail. It is attached to a supercell's general updraft and is oriented roughly parallel to the pseudo-warm front (forward-flank downdraft gust front). As with any inflow band, cloud elements move toward the updraft. Its size and shape change as the strength of the inflow changes. National Weather Service Glossary, cited 2013.
Book-end Vortices	Mesoscale vortices observed at the ends of a line segment of convective cells, usually cyclonic on the northern end of the system and anticyclonic on the southern end, for an environment of westerly vertical wind shear (in the Northern Hemisphere). The vortices are generally strongest between 2 and 4 km above ground level, but may extend from near the surface to about 8 km above ground level. They have been observed at scales between 10 and 200 km, and often have lifetimes of several hours. In extreme cases, the larger cyclonic vortices may become balanced with the Coriolis force and last for several days. *See also* bow echo and comma echo. American Meteorological Society, cited 2013:Book-end Vortices. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/http://glossary.ametsoc.org/wiki/Book-end_vortices]
Boundary	A transition zone of temperature and/or moisture and/or wind in the atmosphere. Fronts, drylines, and outflow boundaries are examples.
Bounded Weak Echo Region (BWER)	(*Also called* vault). A radar signature within a convective storm characterized by a conically-shaped, nearly vertical channel of low reflectivity echo, surrounded on the sides and capped by significantly higher reflectivity echo. The BWER is associated with the core of an intense convective updraft that carries newly formed hydrometeors to high levels before they can grow into precipitation sized particles. It's capped by large concentrations of supercooled liquid water and rapidly growing (typically ≥2-inch diameter) hail. It's typically found imbedded in the slopping echo overhang and aloft above the apex of the low-level inflow notch. A persistent BWER indicates the parent convective storm is likely a supercell, but the BWER itself is not associated with updraft roation. The BWER is typically found 3-10 km (10-33 kft.) AGL and is a few kilometers (1-4 nm) in horizontal diameter. However, on rare occasions, BWERs have been observed up to 5-6 nm wide and

	extending to storm summit. Because of its size, BWERs are rarely detected beyond about 80 mb due to radar beamwidth resolution limitations.
Bow Echo	A radar relectivity echo which is linear, but bent outward into the shape of an archer's bow. The strongest straight-line winds often occur near the "crest" or apex of the bow. Key structural features include an intense rear-inflow jet impinging on the core of the bow, with book-end vortices on both sides of the rear-inflow jet, behind the ends of the bowed convective segment. Bow echoes have been observed with scales between 20 and 200 km, and often have lifetimes between 3 and 6 h. At early stages in their evolution, both cyclonic and anticyclonic book-end vortices tend to be of similar strength, but later in the evolution, the northern cyclonic vortex often dominates (in the Northern Hemisphere), giving the convective system a comma-shaped appearance (see comma echo). Tornadoes sometimes occur, especially with the cyclonic vortex.
Buoyancy	The tendency of a body to float or to rise when submerged in a fluid; the power of a fluid to exert an upward force on a body placed in it. National Weather Service Glossary, cited 2013.
Burn Scar	The mark left on a landscape by fire.
C Band Radar	A radar which operates in the 4-8 GHz frequency and 3.75-7.5 cm wavelength ranges. Because of these characteristics, the antenna size does not need to be as large as an S Band radar to achieve a smaller beam width. This makes C band radars more affordable. The signal is more easily attenuated, so this type of radar is best used for short range weather observation. The Terminal Doppler Weather Radar (TDWR) is a C band radar. *See also* S band radar, X band radar.
Cap	(*Also called* capping layer, capping inversion, or lid.). A region of negative buoyancy below an existing level of free convection (LFC) where energy must be supplied to the parcel to maintain its ascent. This tends to inhibit the development of convection until some physical mechanism can lift a parcel to its LFC. The intensity of the cap is measured by its convective inhibition. The term capping inversion is sometimes used, but an inversion is not necessary for the conditions producing convective inhibition to exist. American Meteorological Society, cited 2013: Cap. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Cap]

Cell	(*Also called* convective cell.) In radar usage, a deep, moist, convectively-induced local maximum in precipitation density that undergoes a life cycle of growth and decay. The rising portion of the reflectivity maximum is indicative of updraft, and the later descending portion is indicative of a precipitation downdraft. Cells in ordinary convective storms last from 20 to 30 min, but often form longer-lasting multicell convective storms. Cells in supercell storms are more steady and last considerably longer. *See also* thunderstorm cell.
Cloud Seeding	The addition of agents (aerosol, small ice particles) that will alter the phase and size distribution of cloud particles, with the intent of influencing precipitation. American Meteorological Society, cited 2013: Cloud Seeding. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Cloud_seeding]
Cold Pool	A region of relatively cold air surrounded by warmer air. It is represented on a weather map analysis as a relative minimum in temperature surrounded by closed isotherms. Cold pools aloft represent regions of relatively low stability, while surface-based cold pools are regions of relatively stable air.
Cold Rain Process	In cloud physics, precipitation generation which is dominated by deposition and the Bergeron Process. *Compar*e warm rain process.
Comma Echo	A radar signature characterized by a convective storm which has evolved into a comma-like shape. It often appears during latter stages in the life cycle of a bow echo due to coriolis forcing. *See also* book-end vorticies.
Convection	Generally, transport of heat and moisture by the movement of a fluid. In meteorology, the term is used specifically to describe vertical transport of heat and moisture in the atmosphere, especially by updrafts and downdrafts in an unstable atmosphere. The terms "convection" and "thunderstorms" often are used interchangeably, although thunderstorms are only one form of convection. Cumulonimbus, towering cumulus clouds, and ACCAS clouds all are visible forms of convection. However, convection is not always made visible by clouds. Convection which occurs without cloud formation is called dry convection, while the visible convection processes referred to above are forms of moist convection. The term 'deep, moist convection' is more accurate to describe precipitating convection most often accompanying thunderstorms.
County Warning Area (CWA)	The group of counties for which a National Weather Service (NWS) Weather Forecast Office (WFO) is responsible for issuing warnings.

Critical Success Index (CSI)	(*Also called* threat score (TS)). A verification measure of categorical forecast performance equal to the total number of correct event forecasts (hits) divided by the total number of storm forecasts plus the number of misses (hits + false alarms + misses). The CSI is not affected by the number of non-event forecasts that verify (correct rejections). However, the CSI is a biased score that is dependent upon the frequency of the event. National Weather Service, Space Weather Prediction Center: Forecast Vertification Glossary, cited 2013.
Crosswise Vorticity	The component of the vorticity vector that is perpendicular to the flow velocity vector. Compare streamwise vorticity. American Meteorological Society, cited 2013: Crosswise Vorticity. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Crosswise_vorticity]
Debris Ball	A dual-polarization tornado debris signature (DPTDS) which contains an isolated reflectivity maximum greater than 51-55 dBZ, co-located with both a tornado vortex signature (TVS) or tornado signature (TS), or other, and low (<0.72) correlation coefficients. As with the broader class of DPTDSs, debris ball identification becomes increasingly difficult with range and detection is generally limited to 80 nm. *See also* dual-polarization debris signature (DPTDS), tornado debris signature (TDS).
Debris Cloud	A rotating "cloud" of dust or debris, near or on the ground, often appearing beneath a condensation funnel and surrounding the base of a tornado. This term is similar to dust whirl, although the latter typically refers to a circulation which contains dust but not necessarily any debris. A dust plume, on the other hand, does not rotate. Note that a debris cloud appearing beneath a thunderstorm will confirm the presence of a tornado, even in the absence of a condensation funnel. National Weather Service Glossary, cited 2013.
Deep Convergence Zone (DCZ)	A narrow interface of high shear and turbulence found along the leading edge of both the rear and forward flank gust fronts of a high precipitation (HP) supercell that extends vertically though midlevels of the storm (~10 km/ 33 kft). The DCZ behaves like a “fluid wall” separating the major storm drafts; dry, potentially cold mid-level inflow feeding downdraft on one side and very warm, moist, low-level inflow feeding the updraft on the other side. Air stream mixing is effectively confined to this narrow zone which shields the supercell updraft from destructive mixing effects and allows the undiluted updraft to approach parcel theory values supportive of significant and often giant hail.
Deep Moist Convection (DMC)	A subset of convection in which air parcels rise above their level of free convection (LFC) to release convective instability through a substantial fraction of the depth of the troposphere in the form of a cumulonimbus cloud.

Deep-Layer Vertical Shear	The vertical wind shear through the lowest half of a deep, moist convective storm. It is best determined using the effective bulk vertical shear, but often the 0-6 km bulk shear is used for simplicity.
Derecho	A widespread convectively induced straight-line windstorm. Specifically, the term is defined as any family of downburst clusters produced by an extratropical mesoscale convective system. Two sub-categories are commonly identified: Serial and progressive. A serial derecho consists of an extensive squall line which is oriented such that the angle between the mean wind direction and the squall line axis is small. A series of LEWPs and bow echoes move along the line. The downburst activity is associated with the LEWPs and bows. It is associated with a linear type mesoscale convective system that moves along and in advance of a cold front or dry line. These boundaries are often associated with a strong, migratory surface low pressure system and strong short wave trough at 500 mb (strong dynamic forcing). Lifted Indices are typically -6 or lower and the advection of dry air in the mid-troposphere (3-7 km above ground) by relatively strong winds leads to high convective instability and increased downdraft potential. The bow echoes move along the line in the direction of the mean flow, often southwest to northeast. These storms move at speeds exceeding 35 knots. Squall line movement is often less than 30 knots. A progressive derecho is characterized by a short curved squall line oriented nearly perpendicular to the mean wind direction with a bulge in the general direction of the mean flow. Downburst activity occurs along the bulging portion of the line. This type of derecho typically occurs in the warm season (May through August) and is most frequent in a zone extending from eastern South Dakota to the upper Ohio Valley. The severe wind storms typically begin during the afternoon and continue into the evening hours. Several hours typically pass between initial convection and the first wind damage report. *See also* serial derecho, progressive derecho. American Meteorological Society, cited 2013: Derecho. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Derecho] National Weather Service Glossary, cited 2013.
Differential Reflectivity (ZDR)	The ratio of radar reflectivity measured by means of two signals that differ in one attribute, for example, polarization or wavelength. As applied to polarimetric radar observations, the differential reflectivity is the ratio of the reflectivity observed with transmitted and received signals of horizontal polarization to that observed with signals of vertical polarization. It is commonly represented by the symbol ZDR. American Meteorological Society, cited 2013: Differential Reflectivity (ZDR). Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Zdr]

Downburst	An area of strong, often damaging winds produced by one or more convective downdrafts over an area from less than 1 to 400 km in horizontal dimensions. *See also* macroburst, microburst. American Meteorological Society, cited 2013: Downburst. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Downburst] Buoyancy induced downbursts typically have spatial dimensions in the meso-B to meso-gamma scales. *See also* downdraft.
Downdraft	Small-scale downward moving air current, most often forced by negative buoyancy processes, in a cumulonimbus cloud. *Compare* updraft.
Downshear	In the same direction as the shear vector within a specified layer. *Compare* upshear.
Downwind	In the same direction as the wind flow, or toward the direction in which the wind is moving. *Compare* upwind. National Weather Service Glossary, cited 2013.
Dual-Polarization Radar	A radar capable of transmitting and receiving two orthogonal polarizations. The transmitted polarization must be switchable at a rate that is fast compared with the timescale of changes in the scattering properties of the target and the propagation medium. American Meteorological Society, cited 2013: Dual-Polarization Radar. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Dual-polarization_radar]
Dust Devil	A well-developed dust whirl; a small but vigorous whirlwind, usually of short duration, rendered visible by dust, sand, and debris picked up from the ground. Dust devils occasionally are strong enough to cause minor damage (up to F1 on the Fujita scale). Diameters range from about 3 m to greater than 30 m; their average height is about 200 m, but a few have been observed as high as 1 km or more. They have been observed to rotate anticyclonically as well as cyclonically. Although the vertical velocity is predominantly upward, the flow along the axis of large dust devils may be downward. Large dust devils may also contain secondary vortices. Dust devils are best developed on a hot, calm afternoon with clear skies, in a dry region when intense surface

	heating causes a very steep lapse rate of temperature in the lowest 100 m of the atmosphere.: American Meteorological Society, cited 2013: Dust Devil. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Dust_devil]
Dust Plume	A non-rotating "cloud" of dust raised by straight-line winds. Often seen in a microburst or behind a gust front. National Weather Service Glossary, cited 2013.
Dust Whirl	(Also called dancing dervish, dancing devil, devil, satan, shaitan; and, over desert areas, desert devil, sand auger, sand devil.) A rapidly rotating column of air (whirlwind) over a dry and dusty or shady area, carrying dust, leaves, and other light material picked up from the ground. When well developed it is known as a dust devil. Dust whirls typically form as the result of strong convection during sunny, hot, calm summer afternoons. This type is generally several meters in diameter at the base, narrowing for a short distance upward and then expanding again, like two cones apex to apex. Their height varies; normally it is only 30–100 m, but in hot desert country it may be as high as 1 km. Rotation may be either clockwise or counterclockwise. Dust whirls move erratically, from one patch of heated air to another, and generally slowly. In desert country it is not unusual for three or more desert devils to be visible at the same time. Another type of vigorous dust whirl occurs under the bases of cumulonimbus or cumulus clouds, almost always on or near a wind-shift line. These vortices often inflict little or no damage and are short- lived, but occasionally represent the first visible sign of a developing tornado. Another form of dust whirl, often seen at street corners, is merely an eddy caused by the meeting of winds blowing along two intersecting streets. Such whirls are small and very short-lived. American Meteorological Society, cited 2013: Dust Whirl. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Dust-whirl]
Echo	In radar, a general term for the appearance, on a radar display, of the radio signal scattered or reflected from a target. The characteristics of a radar echo are determined by 1) the waveform, frequency, and power of the incident wave; 2) the range and velocity of the target with respect to the radar; and 3) the size, shape, and composition of the target. American Meteorological Society, cited 2013: Echo. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Echo]

Echo Overhang	(*Also called* overhang). In the radar reflectivity echo associated with a convective storm, that portion of the echo that is located above the weak-echo region on the low-level inflow side of the storm. The overhang consists of precipitation particles diverging from the storm's summit that descend as they are carried downwind. If the storm echo develops a bounded weak-echo region, it is found within the echo overhang. *See also* bounded weak echo region (BWER). American Meteorological Society, cited 2013: Overhang. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Overhang]
Echo Top	The height above ground of the center of the radar beam using the tilt, or scan, that contains the highest elevation where reflectivities greater than 18 dBZ can be detected. National Weather Service Glossary, cited 2013.
Elevated Convection	Convection whose ascending air parcels are rooted within an atmospheric layer above the boundary layer. Elevated convection often occurs when air near the ground is relatively cool and stable (e.g., during periods of isentropic lift), but an unstable layer of air is present aloft. In these cases, stability indices based on near-surface measurements (such as the lifted index) will underestimate the amount of instability present. Severe weather (particularly hail) is possible from elevated convection, but is less likely than it is from surface-based convection.
Enhanced Fujita (EF) Scale	(Formerly called Fujita (F) Scale). A scale of wind damage intensity in which wind speeds are inferred from an analysis of wind damage. It uses three-second gusts estimated at the point of damage based on a judgment of eight (8) levels of damage to the twenty eight (28) damage indicators (DIs). All tornadoes, and most other severe local windstorms, are assigned a single number from this scale according to the most intense damage caused by the storm. EF0 (weak): 65- 85 mph, Gale EF1 (weak): 85-110 mph, Weak EF2 (strong): 111-135 mph, Strong EF3 (strong): 136-165 mph, Severe EF4 (violent): 166-200 mph, Devastating EF5 (violent): >200 mph, Incredible
Enhanced "V"	A signature on infrared satellite imagery that depicts a warm, wedge-shaped region stretching from the upshear edge of a thunderstorm anvil, downshear along its long axis. This so-called warm wake is surrounded by long, narrow regions of colder pixels along either side, forming an apparent V shape. The warmer pixels may form from 1) Stratospheric cirrus "blow-off" from

	overshooting tops, meaning that they are higher than the mean anvil height, even though they are warmer. 2) Mixing of saturated updraft air with warmer air at higher altitudes resulting in a narrow, cloudy, warmer interface.The signature typically forms on convective storms possessing extremely strong updrafts. Enhanced V should not be confused with V notch, which is a radar signature.
Enhanced Wording	1. An option used by the Storm Prediction Center (SPC) in tornado and severe thunderstorm watches when the potential for strong/violent tornadoes, or unusually widespread damaging straight-line winds, is high. The text that accompanies a watch of this type will include the line "THIS IS A PARTICULARLY DANGEROUS SITUATION." 2. Strong wording used in National Weather Service (NWS) Weather Forecast Office (WFO) products (such as in warnings and statments) to convey the urgency of the action needed from impending weather threat. An example is "This is a Tornado Emergency." Enhanced wording should only be used in extremely rare, life-threatening situations. National Weather Service Glossary, cited 2013.
Entrainment	In meteorology, the mixing of environmental air into a preexisting organized air current so that the environmental air becomes part of the current; the opposite of detrainment. Entrainment of air into clouds, especially cumulus, is said to be inhomogeneous when the timescale for mixing of environmental air is very much greater than the timescale for drop evaporation. Under these conditions, which are often found when environmental air is first entrained into cumulus, regions of cloud and entrained air are intertwined, with evaporation occurring only on the edges of the interface between the cloudy and entrained environmental air. See also entrainment zone. American Meteorological Society, cited 2013: Entrainment. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Entrainment]
Entrainment Zone	A shallow region at the top of a mixed layer where fluid is entrained into the growing atmospheric boundary layer from the overlying fluid by the collapse of rising convective plumes or bubbles. *See also* entrainment. National Weather Service Glossary, cited 2013.
Equilbrium Level (EL)	On a sounding, the level above the level of free convection (LFC) at which the temperature of a rising air parcel again equals the temperature of the environment. Thus, according to parcel theory, the height of the EL is the height at which thunderstorm updrafts no longer accelerate upward. The EL is often considered the layer at which thunderstorm anvil tops are found.

False Alarm Ratio (FAR)	A verification measure of categorical forecast performance equal to the number of false alarms (e.g., a NWS warning which does not verify) divided by the total number of event forecasts. National Weather Service, Space Weather Prediction Center: Forecast Vertification Glossary, cited 2013.
Family of Tornadoes	(*Also called* tornado family). A sequence of long-lived tornadoes produced by a "cyclic" supercell thunderstorm. Tornadoes touch down at quasi-regular intervals (typically 45 min). Usually a new tornado develops in a new mesocyclone just after an old tornado has decayed in an old, occluded neighboring mesocyclone. Sometimes, two successive tornadoes may overlap in time for awhile. The two mesocyclones may rotate partially around each other. If the damage tracks of the tornadoes appear to form a wavy broken line, the family is classified as a series mode. In the more common parallel-mode family, the damage tracks are parallel arcs with each new tornado forming on the right side of its predecessor. The parallel mode is subcategorized into left turn and right turn, according to the direction in which the paths curve.: American Meteorological Society, cited 2013: Family of Tornadoes. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Family_of_tornadoes]
Flanking Line	An organized lifting zone of cumulus and towering cumulus clouds, connected to and extending outward from the mature updraft tower of a supercell or strong multicell convective storm. The flanking line often has a stair-step appearance, with the tallest clouds adjacent to the mature updraft tower. American Meteorological Society, cited 2013: Flanking Line. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Flanking_line]
Flash Flood	A rapid and extreme flow of high water into a normally dry area, or a rapid water level rise in a stream or creek above a predetermined flood level, beginning within six hours of the causative event (e.g., intense rainfall, dam failure, ice jam). However, the actual time threshold may vary in different parts of the country. Ongoing flooding can intensify to flash flooding in cases where intense rainfall results in a rapid surge of rising flood waters. National Weather Service Glossary, cited 2013.
Flash Flood Guidance (FFG)	Forecast guidance produced by the National Weather Service (NWS) River Forecast Centers (RFCs), often model output, specific to the potential for flash flooding (e.g., how much rainfall over a given area will be required to produce flash flooding). National Weather Service Glossary, cited 2013.

Flash Flood Statement (FFS)	In hydrologic terms, a statement by a National Weather Service (NWS) Weather Forecast Office (WFO) which provides follow-up information on flash flood watches and warnings. National Weather Service Glossary, cited 2013.
Flash Flood Warning	A National Weather Service (NWS) Weather Forecast Office (WFO) product issued to inform the public, emergency management, and other cooperating agencies that flash flooding is in progress, imminent, or highly likely. National Weather Service Glossary, cited 2013.
Flash Flood Watch	A National Weather Service (NWS) Weather Forecast Office (WFO) product issued to indicate current or developing hydrologic conditions that are favorable for flash flooding in and close to the watch area, but the occurrence is neither certain or imminent. National Weather Service Glossary, cited 2013.
Forcing	Energy exerted upon the air which causes a change in its vertical motion.
Forward Flank Downdraft (FFD)	The main region of downdraft in the forward, or leading, part of a supercell, where most of the heavy precipitation is. *See also* rear flank downdraft (RFD). National Weather Service Glossary, cited 2013.
Frontogenesis	1. The initial formation of a front or frontal zone. 2. In general, an increase in the horizontal gradient of an airmass property, principally density, and the development of the accompanying features of the wind field that typify a front.. *Compare* frontolysis. American Meteorological Society, cited 2013: Frontogenesis. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Frontogenesis]
Frontolysis	1. The dissipation of a front or frontal zone. 2. In general, a decrease in the horizontal gradient of an air mass property, principally density, and the dissipation of the accompanying features of the wind field. *Compare* frontogenesis. American Meteorological Society, cited 2013: Frontolysis. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Frontolysis]

Funnel Cloud	A condensation cloud, typically funnel-shaped and extending outward from a cumuliform cloud, associated with a rotating column of air (a vortex) that may or may not be in contact with the surface. If the rotation is violent and in contact with the surface, the vortex is a tornado. Funnel clouds can occur through a variety of processes in association with convection. For example, small funnel clouds are infrequently seen extending from small, dissipating cumulus clouds in environments with significant vertical wind shear in the cloud-bearing layer. American Meteorological Society, cited 2013: Funnel Cloud. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Funnel_cloud]
Graupel	Heavily rimed snow particles, often called snow pellets; often indistinguishable from very small soft hail except for the size convention that hail must have a diameter greater than 5 mm. Sometimes distinguished by shape into conical, hexagonal, and lump (irregular) graupel. American Meteorological Society, cited 2013: Graupel. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Graupel]
Gust Front	The leading edge of a mesoscale pressure dome separating the outflow air in a convective storm from the environmental air. This boundary, which is marked by upward motion along it and downward motion behind it, is followed by a surge of gusty winds on or near the ground. A gust front is often associated with a pressure jump, wind shift, temperature drop, and sometimes with heavy precipitation. Gust fronts are often marked by arcus clouds. American Meteorological Society, cited 2013: Gust Front. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Gust_front]
Gustnado	A short-lived, shallow, generally weak, vertically oriented vortex found along a gust front. Gustnadoes are usually visualized by a rotating dust or debris cloud. Gustnadoes are not associated with storm-scale rotation (i.e. mesocyclones); they are more likely to be associated visually with a shelf cloud than with a wall cloud.
Hagen-Poiseuille Law	(*Also called* the Hagen-Poiseuille Equation, Poiseuille Law, or Poiseuille equation). A statement in physics was states that the volume of water through the pore length is proportional to the pore space.
Hail	Precipitation in the form of balls or irregular lumps of ice more than 5mm in diameter, always produced by convective clouds, nearly always cumulonimbus.

Hail Growth Zone	The region of supercooled water within a deep, moist, convective storm where hailstone growth is maximized, approximately the -10°C to -30°C layer. The growth rate is maximized near -13°C and rapidly diminishes at temperatures approaching -30°C as supercooled water droplets become rare at these colder tempertures.
Heidke Skill Score (HSS)	A skill corrected verification measure of categorical forecast performance. The HSS is equal to the total number of correct forecasts minus the correct random forecasts (hits + correct rejections - correct random forecasts) divided by the total number of forecasts minus the correct forecasts due to chance (hits + false alarms + misses + correct rejections - correct random forecasts). This skill score falls within a (-1, +1) range. No incorrect forecasts give a score of +1, no correct forecasts give a score of -1, and either no events forecast or no events observed give a score of 0. National Weather Service, Space Weather Prediction Center: Forecast Vertification Glossary, cited 2013.
Helicity	A property of a moving fluid which represents the potential for helical flow (i.e. flow which follows the pattern of a corkscrew) to evolve. Helicity is proportional to the strength of the flow, the amount of vertical wind shear, and the amount of turning in the flow (i.e. vorticity). Atmospheric helicity is computed from the vertical wind profile in the lower part of the atmosphere (usually from the surface up to 3 km), and is measured relative to storm motion. Higher values of helicity (generally, around 150 m2/s2 or more) favor the development of mid-level rotation (i.e. mesocyclones). National Weather Service Glossary, cited 2013.
Hodograph	A plot representing the vertical distribution of horizontal winds, using polar coordinates. A hodograph is obtained by plotting the end points of the wind vectors at various altitudes, and connecting these points in order of increasing height. Interpretation of a hodograph can help in forecasting the subsequent evolution of thunderstorms (e.g., squall line vs. supercells, splitting vs. non-splitting storms, tornadic vs. nontornadic storms, etc.). National Weather Service, WFO Norman: A Comprehensive Glossary of Weather Terms for Storm Spotters, cited 2013.
Hook Echo	A radar signature characterized by a curve-shaped band of reflectivity echo caused when precipitation is drawn into the spiral of a mesocyclone. The hook is a fairly shallow feature, typically 3-4 km in height, extending downward as a precipitation streamer from the echo overhang aloft. When high resolution radar imagery is available, the hook is seen to spiral inward forming a sharply defined figure “6.” A tornado, if present, is within the figure “6” or at the tip of the hook echo itself.

Hydrometeor	Any product of condensation or deposition of atmospheric water vapor, whether formed in the free atmosphere or at the earth's surface; also, any water particle blown by the wind from the earth's surface. American Meteorological Society, cited 2013: Hydrometeor. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Hydrometeor]
Index	In meteorology forecasting, a number derived from a formula, used to characterize a set of data. Examples include Lifted Index (LI) and Energy Helicity Index (EHI).
Infiltration	In hydrologic terms, movement of water through the soil surface into the soil. National Weather Service Glossary, cited 2013.
Infiltration Capacity	In hydrologic terms, the maximum rate at which water can enter the soil at a particular point under a given set of conditions. National Weather Service Glossary, cited 2013.
Infiltration Index	In hydrologic terms, an average rate of infiltration, in inches per hour, equal to the average rate of rainfall such as that the volume of rainfall at greater rates equals the total direct runoff. National Weather Service Glossary, cited 2013.
Infiltration Rate	The ability of water to move into the soil from the surface.
Inflow Bands	(*Also called* feeder bands.) Low clouds, arranged in parallel rows to the low-level winds and moving into or toward a thunderstorm. They may indicate the strength of the inflow of moist air into the storm, and, hence, its potential severity. Spotters should be especially wary of inflow bands that are curved in a manner suggesting cyclonic rotation; this pattern may indicate the presence of a mesocyclone. National Weather Service Glossary, cited 2013.
Inflow Jet	Local jets of air near the surface flowing inward toward the base of a tornado. National Weather Service Glossary, cited 2013.
Inflow Notch	A radar signature characterized by an enhanced, low-level, concave reflectivity gradient open to the inflow side of a convective storm. This signature indicates the presence of a very strong updraft with associated enhanced low-level inflow. *See also* rear inflow notch.
Intercell Seeding	A method of feeding hydrometeors to the main updraft(s) of a deep, moist, convective storm by flanking lines or recycling. It improves precipitation efficiency because cells share hydrometeors while the environmental humidity increases. Multicell storms whose updrafts recycle hydrometeors are more efficient than more discrete cells.

Internal Dynamics (ID) Method	(*Also called* Bunker's Method, Bunker's Motion). A method for estimating supercell motion which uses the mechanisms by which updraft and shear interact. It can be used to calculate storm motion for both the cyclonically and anticyclonically rotating supercells resulting from a storm split. Supercell motion is approximately 7.5 m/s right and left of the deep-layer shear vector along a line that passes through the point describing the mean convective steering layer flow.
Katafront	A front (usually a cold front) at which the warm air descends the frontal surface (except, presumably, in the lowest layers). *Compare* anafront. American Meteorological Society, cited 2013: Katafront. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Katafront]
KDP Column	A dual-polarimetric radar signature of specific differential phase (KDP) values > 2-3°/km caused by high concentrations of liquid water in a convective cell. *See also* ZDR column, KPD column.
Kinematics	The branch of dynamics that describes the properties of pure motion without regard to force, momentum, or energy. Translation, advection, vorticity, and deformation are examples of kinematic variables. American Meteorological Society, cited 2013: Kinematics. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Kinematics]
Lag	1) The measure of the time between the center of mass of precipitation to the center of mass of runoff (on the hydrograph); basin lag is a function of not only basin characteristics, but also of storm intensity and movement. Some hydrologic texts define lag from the center of mass of rainfall to the hydrograph peak. 2) The time it takes a flood wave to move downstream. National Weather Service Glossary, cited 2013.
Landspout	A non-mesocyclonic tornado occuring with a parent cloud in its growth stage and with its vorticity originating in the boundary layer. Landspouts typically are observed beneath Cbs or towering cumulus clouds (often as no more than a dust whirl), and essentially are the land-based equivalents of waterspouts.
Left Mover	A thunderstorm (often a supercell) which moves to the left relative to the steering winds, and to other nearby thunderstorms. *Compare* right mover. National Weather Service Glossary, cited 2013.

Level of Free Convection (LFC)	The level at which a parcel of air lifted dry-adiabatically until saturated and saturation-adiabatically thereafter would first become warmer than its surroundings in a conditionally unstable atmosphere. On a thermodynamic diagram the level of free convection is given by the point of intersection of the process curve, representing the process followed by the ascending parcel, and the sounding curve, representing the lapse rate of temperature in the environment. From the level of free convection to the point where the ascending parcel again becomes colder than its surroundings the atmosphere is characterized by latent instability. Throughout this region the parcel will gain kinetic energy as it rises. American Meteorological Society, cited 2013: Level of Free Convection. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Level_of_free_convection]
Lifting Condensation Level (LCL)	The level at which a parcel of moist air lifted dry-adiabatically would become saturated. On a thermodynamic diagram it is located at the point of intersection of the dry adiabat through the point representing the parcel's original pressure and temperature with the saturation mixing ratio line having the same value of the mixing ratio as the parcel. The pressure and temperature at the lifting condensation level are usually called the condensation pressure and condensation temperature, respectively, and the corresponding point on a thermodynamic diagram is called either the characteristic point, adiabatic saturation point, or adiabatic condensation point. American Meteorological Society, cited 2013: Lifting Condensation Level. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Lifting_condensation_level]
Line Echo Wave Pattern (LEWP)	A weather radar formation in which an area of thunderstorms forms a mesoscale low-pressure area with a rotating "head" (cyclonically-rotating bookend vortex) and, typically, a bow echo to its south (or equatorward). LEWPs are often associated with a multiple-bow serial derecho and often produce tornadoes, some of which can be strong, particularly those associated with the rotating "head." *See also* bow echo, comma echo.
Liquid Water Loading	(*Also called* precipitation loading). The amount of liquid water present within an air parcel as cloud droplets, rain, or ice, usually expressed in percent or fraction by weight (e.g., as a liquid water mixing ratio rL) or volume. The higher the liquid water loading, the greater the average density and colder the virtual temperature of the parcel.: American Meteorological Society, cited 2013: Liquid Water Loading. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Liquid_water_loading]

Loaded Gun (Sounding)	[Slang] A sounding characterized by extreme instability but containing a cap, such that explosive thunderstorm development can be expected if the cap can be weakened or the air below it heated sufficiently to overcome it. National Weather Service Glossary, cited 2013.
Low CC column	A dual-pol radar signature used to identify a deep, moist, convective updraft. It appears as an upward extension of the low (< 0.8) Correlation Coefficient (CC) clear air boundary layer echoes into the updraft. The echo is associated with scattering from flying insects, light vegetative debris, and other non-meteorological scatterers. *See also* ZDR column, KPD column.
Low CC Inflow	A dual-pol radar signature of low (< 0.8) Correlation Coefficient (CC) clear air boundary layer echoes that help demarcate the inflow into a deep, moist, convective updraft. The echo is associated with scattering from flying insects, light vegetative debris, and other non-meterological scatterers.
Low CC Ring	A dual-polarimetric radar signature of low (< 0.8) Correlation Coefficient (CC) echoes located at the periphery of a persistent updraft in a sheared environment, just above the freezing level, caused by mixed-phase precipitation which forms as a result of melting between the environment and the updraft. *See also* ZDR ring.
Low-Echo Centroid (LEC) Signature	A radar signature in which the greatest refelectivity of a mature cell is below the freezing (0°C) level.
Low-Level	The lower portion of the tropopshere. No distinct limit is set, but the term can generally be applied to the levels 850 mb and below. *Compare* midlevel. *See also* upper air.
Low-Level Jet (LLJ)	Releatively strong winds concentrated within a narrow band found in the lowest 2–3 km of the troposphere. At night, sometimes called a nocturnal jet. In the United States, it often refers to a southerly wind maximum in the boundary layer, common over the Plains states at night during the warm season (spring and summer).
Low-Level Rotational Velocity (LLDV)	The velocity difference (DV) measured in the lowest radar elevation slice calculated by adding the absolute magnitudes of the lowest and highest radial velocities found within a storm-scale vortex signature (typically a mesocyclone, tornado vortex signature, or tornado signature).
Macroburst	A downburst that covers an area greater than 4 km (2.5 nm) along a side and peak winds lasting between 5 and 20 minutes. Extreme macrobursts may cause wind damage up to EF-3 intensity. *See also* microburst.

Macroscale	Meteorological expression referring to synoptic events occurring on a scale of thousands of kilometers, such as warm and cold fronts. *Compare* mesoscale, microscale. American Meteorological Society, cited 2013: Macroscale. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Macroscale]
Manning Equation	An empirical equation that applies to uniform flow in open channels and is a function of the channel velocity, flow area, and channel slope. Corvallis Forestry Research Community: Hydraulic Reference (http://www.fsl.orst.edu/geowater/FX3/help/FX3_Help.html#8_Hydraulic_Reference/Manning_s_Equation.htm), cited 2013.
Maximum Delta V (MDV)	The maximum velocity difference (DV) for all Doppler radar elevation slices containing a TS/TVS.
Melting Layer	The altitude interval throughout which ice-phase precipitation melts as it descends. The top of the melting layer is the melting level. The melting layer may be several hundred meters deep, reflecting the time it takes for all the hydrometeors to undergo the transition from solid to liquid phase. The temperature of the melting layer is typically 0°C or slightly warmer. *See* bright band. American Meteorological Society, cited 2013: Melting Layer. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Melting_layer]
Meltwater	The water released by the melting of snow or ice, including hail.
Mesocyclogenesis	Any development or strengthening of a mesocyclone.
Mesocyclone	A storm-scale (around 2–10 km in diameter) vertical vorticity maximum partially or fully embedded within an updraft of deep, moist convection (DMC). The vorticity associated with a mesocyclone is often on the order of 10-2 s-1 or greater. A mesocyclone is the defining characteristic of a supercell. Tornadoes sometimes form in mesocyclones. Persistent mesocyclones that have significant vertical extent are detected by Doppler radar as mesocyclone signatures. A mesocyclone can often be implied visually by the existence of updraft striations and/or curved inflow bands.

Mesocyclone Signature	The Doppler velocity pattern of a mesocyclone. In a storm-relative reference frame, the idealized signature is symmetric about the radar viewing direction with marked azimuthal shear across the core region between peak Doppler velocity values of opposite sign. Typical signatures consist of Doppler velocity differences of 25– 75 m s-1 across core diameters of 2–8 km, with resulting azimuthal shear values of 5 × 10-3 s-1 to 2 × 10-2 s-1.: American Meteorological Society, cited 2013:Mesocyclone Signature. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Mesocyclone_signature]
Mesocyclonic Tornado	A tornado that is associated with a mesocyclone. Also called a "supercell tornado." *Compare* non-mesocyclonic tornado.
Mesohigh	A mesoscale area of high atmospheric pressure that typically forms beneath a multicell thunderstorm. It is usually associated with a mesoscale convective system (MCS) or its remnants. *Compare* mesolow.
Mesolow	A mesoscale low pressure center. Mesolow should not be confused with mesocyclone, which is a storm-scale phenomenon.. *Compare* mesohigh. National Weather Service Glossary, cited 2013.
Mesoscale	Pertaining to atmospheric phenomena having horizontal scales ranging from a few to several hundred kilometers, including thunderstorms, squall lines, fronts, precipitation bands in tropical and extratropical cyclones, and topographically generated weather systems such as mountain waves and sea and land breezes. From a dynamical perspective, this term pertains to processes with timescales ranging from the inverse of the Brunt–Väisälä frequency to a pendulum day, encompassing deep moist convection and the full spectrum of inertio-gravity waves but stopping short of synoptic-scale phenomena, which have Rossby numbers less than 1. *Compare* macroscale, microscale. American Meteorological Society, cited 2013: Mesoscale. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Mesoscale]

Mesoscale Beta Elements (MBE) Technique	A procedure used to forecast the short-term (3-6 hour) motion of Mesoscale Convective System (MCS) centroids which builds on the long-established observation that MCS motion is a function of both the advection of existing cells by the mean wind and the propagation of new cells relative to existing storms. MCS centroid forecast motion is calculated to be the vector sum of the mean cloud-layer wind (typically assumed to be the 850-300 mb layer) and the mesoscale beta element (MBE) propagation component. The MBE propagation vector is dictated by the location of the maximum cold-pool gust front convergence in the presence of conditional instability. For an upwind (downwind) propagating MCS, the magnitude and direction is assumed to be equal and opposite (identical) to that of the low-level jet (typically assumed to be the 850 mb wind). Corfidi, Stephen F., 2003: Cold Pools and MCS Propagation: Forecasting the Motion of Downwind-Developing MCSs. Wea. Forecasting, 18, 997–1017.
Mesoscale Convective Complex (MCC)	A subset of mesoscale convective systems (MCS) that exhibits a large, circular (as observed by satellite), long-lived, cold cloud shield with the following physical characteristics: Size: A - Cloud shield with continuously low infrared (IR) temperature ≤ -32°C must have an area ≥ 105 km2; and B - Interior cold cloud region with temperature ≤ -52°C must have an area ≥ 0.5 X 105 km2. Initiate: Size definitions A and B are first satisfied Duration: Size definitions A and B must be met for a period ≥ 6 h. Maximum extent: Contiguous cold cloud shield (IR temperature ≤ -33°C) reaches maximum size. Shape: Eccentricity (minor axis/major axis) ≥ 0.7 at time of maximum extent. Terminate: Size definitions A and B no longer satisfied. MCCs typically form during the afternoon and evening in the form of several isolated thunderstorms, during which time the potential for severe weather is greatest. During peak intensity (usually at night), the primary threat shifts toward heavy rain and flooding.
Mesoscale Convective System (MCS)	A complex of thunderstorms which becomes organized on a scale larger than the individual thunderstorms, produces a contiguous precipitation area on the order of 100 km or more in horizontal scale in at least one direction, and normally persists for several hours or more. An MCS exhibits deep, moist convective overturning contiguous with or embedded within a mesoscale vertical circulation that is at least partially driven by the convective overturning. MCSs may be round or linear in shape, and include systems such as tropical cyclones, squall lines, and MCCs (among others). MCS often is used to describe a cluster of thunderstorms that does not satisfy the criteria of a mesoscale convective complex (MCC).

Mesoscale Convective Vortex (MCV)	A quasi-steady, mesoscale, cyclonic circulation that forms in the mid-troposphere within the stratiform region of a mesoscale convective system (MCS) often persisting after its parent MCS has dissipated. With a core of only 50-100 km wide and 1500-4500 m deep, and MCV is often overlooked in standard weather analyses. And MCV can persist for more than 12 hours upon achieving a balance between pressure gradient and Coriolis forces, sometimes becoming the seed of the next thunderstorm outbreak or, upon reaching tropical waters, serving as the nucleus for a tropical cyclone.
Microburst	A downburst that covers an area less than 4 km (2.5 nm) along a side with peak winds that last 2–5 minutes. Differential velocity across the divergence center is greater than 10 m s-1. The strong wind shears associated with a microburst can result in aircraft accidents. Microburst are commonly sub-classified as wet, dry, or hybrid. A wet microburst is accompanied by heavy precipitation at the surface; an outward distortion of the precipitation near the surface called a rain foot is a common visual sign. A dry microburst produces little or no precipitation which reaches the surface; a dust plume or a ring of blowing dust beneath a local area of virga is a common visual sign. A hybrid microburst contains characteristics of both wet and dry microbursts. *See also* macroburst.
Microphysics	The branch of physics concerned with small objects and systems, such as atoms, molecules, nuclei, and elementary particles. Specifically, in the meteorological field, physics associated with cloud and precipitation processes.
Microscale	Atmospheric motions with Lagrangian Rossby numbers greater than 200 or spatial scales 2 km or less. *Compare* macroscale, mesoscale. American Meteorological Society, cited 2013: Microscale. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Microscale]
Mid-Altitude Radial Convergence (MARC) Velocity Signature	A Doppler radar-velocity signature that serves as a precursor to the initial onset of damaging straight-line winds in a deep, moist, convective cell; Quasi-Linear Convective System (QLCS); or bowing convective system. It may also represent the upper portions of a gust front associated with a deep, convergence zone (DCZ).
Midlevel	The middle portion of the tropopshere. No distinct limits are set, but the term can generally be applied to levels between 700 mb to 500 mb. *Compare* low-level. *See* also upper air.

Mie Scattering	Any scattering produced by spherical particles whose diameters are greater than 1/10 the wavelength of the scattered radiation. This type of scattering causes the clouds to appear white in the sky. Often, hail exhibits in this type of scattering. National Weather Service Glossary, cited 2013.
Misocyclone	A horizontal vortex with a width of between 40 m and 4 km. It is often used to refer to 1) a vortex within a convective storm with a horizontal scale of less than 4 km, and 2) a near-surface vortex with a horizontal scale of less than 4 km along a convergence line. *Compare* mesocyclone. American Meteorological Society, cited 2013: Misocyclone. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Misocyclone]
Mixed Layer (ML)	(Also called convective mixed layer, convective boundary layer (CBL), or mixing layer). A type of atmospheric boundary layer characterized by vigorous turbulence tending to stir and uniformly mix, primarily in the vertical, quantities such as potential temperature and momentum or wind speed. Moisture is often not so well mixed, showing a slight decrease with height. The vigorous turbulence can be caused by either strong winds or wind shears that generate mechanical turbulence (called forced convection), or by buoyant turbulence (called free convection) associated with large thermals. The buoyantly generated mixed layers are usually statically unstable, caused by heating at the bottom boundary such as the earth's surface or radiative cooling at the tops of cloud or fog layers within the mixed layer. American Meteorological Society, cited 2013: Mixed Layer. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Mixed_layer]
Multicell Convective Storm	(*Also called* multicell, multicell thunderstorm). A cluster of ordinary cells and/or supercells at various stages of their life cycle in close enough proximity to at least share a common precipitation area and cold pool (gust front). New cells are generated primarily by either low-level convergence along a preexisting boundary or by lifting at the leading edge of the system-scale cold pool that was produced by the previous cells. The cells move roughly with the mean wind. However, the storm motion usually deviates significantly from the mean wind due to discrete propagation (new cell development) along the gust front. The multicellular nature of the storm is usually apparent on radar with multiple reflectivity cores and maximum tops. Lifetime may be several hours.

Multiple Vortex Tornado	A tornado in which two or more vortices are present at the same time, often rotating about a common center or about each other. Multiple-vortex tornadoes can be especially damaging.
National Hurrican Center (NHC)	One of three branches of the Tropical Prediction Center (TPC). This center maintains a continuous watch on tropical cyclones over the Atlantic, Caribbean, Gulf of Mexico, and the Eastern Pacific. The Center prepares and distributes hurricane watches and warnings for the general public, and also prepares and distributes marine and military advisories for other users. During the "off-season" NHC provides training for U.S. emergency managers and representatives from many other countries that are affected by tropical cyclones. NHC also conducts applied research to evaluate and improve hurricane forecasting techniques, and is involved in public awareness programs. National Weather Service Glossary, cited 2013.
National Severe Storms Laboratory (NSSL)	A federal research facility under the National Oceanic and Atmospheric Administration's (NOAA's) Office of Oceanic and Atmospheric Research (OAR) at the National Weather Center (NWC) in Norman, Oklahoma which serves the nation by working to improve the leadtime and accuracy of severe weather warnings and forecasts in order to save lives and reduce property damage. NSSL's basic and applied research focuses on understanding severe weather processes, developing weather observation technology, and improving forecast tools, with emphasis on: Weather radar, tornadoes, flash floods, lightning, damaging winds, hail, and winter weather.
NEXRAD	NEXt Generation RADar. A National Weather Service (NWS) network of about 159 S-band WSR-88D Doppler radars operating nationwide.
Non-Mesocyclonic Tornado	(*Also called* non-supercell tornado). A tornado which is not associated with a mesocyclone. *Compare* mesocyclonic tornado.*See also* landspout, waterspout.
Ordinary Cell	The most basic component of a convective storm, consisting of a single main updraft that is usually quickly replaced by a downdraft once precipitation begins. Ordinary cells are especially observed in environments with weak vertical wind shear, and typically have lifetimes of 30-50 minutes. Ordinary cells are the primary component of multicell storms. *See also* convective cell, multicell convective storm, thunderstorm. American Meteorological Society, cited 2013: Ordinary Cell. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Ordinary_Cell]
Outflow Boundary	A mesoscale surface boundary separating thunderstorm-cooled air (outflow) from the surrounding air. Outflow boundaries may be short-lived, or persist for longer than a day after the thunderstorms that generated them dissipate, and may travel hundreds of miles from their area of origin. New thunderstorms often

	develop along outflow boundaries, especially near the point of intersection with another boundary (cold front, dry line, another outflow boundary, etc.).
Overshooting Top	(*Also called* anvil dome, penetrating top.) A domelike protrusion above a cumulonimbus anvil, representing the intrusion of an updraft through its equilibrium level (level of neutral buoyancy). It is usually a transient feature because the rising parcel's momentum acquired during its buoyant ascent carries it past the point where it is in equilibrium (i.e., anvil level); the air within it rapidly becomes negatively buoyant and descends. Tall and persistent overshooting tops are frequently observed with strong or severe thunderstorms in which there is a nearly continuous stream of buoyant updrafts.: American Meteorological Society, cited 2013:Overshooting Top. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Overshooting_top]
Parameter	In general, any quantity of a problem that is not an independent variable. More specifically, the term is often used to distinguish, from dependent variables, quantities that may be more or less arbitrarily assigned values for purposes of the problem at hand. Examples include the Vorticity Generation Parameter (VGP), Significant Hail Parameter (SHIP), and Significant Tornado Parameter (STP). American Meteorological Society, cited 2013: Parameter. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Parameter]
Parcel	An imaginary volume of air to which may be assigned various thermodynamic and kinematic quantities. The size of a parcel is arbitrary but is generally much smaller than the characteristic scale of variability of its environment. American Meteorological Society, cited 2013: Parcel. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Parcel]
Partial Beam Filling	A limitation of the rainfall estimation techniques used by weather radar. At far ranges from the radar, a storm may occupy only a portion of the radar beam (which may be several miles across). However, the radiation received by the radar antenna consists of the average reflectivity across the entire beam, so the reflectivity and associated rainfall rates are underestimated. National Weather Service Glossary, cited 2013.

Pendant Echo	Radar signature generally similar to a hook echo, except that the hook shape is not as well defined, typically due to radar range and/or beamwidth resolution limitations. National Weather Service Glossary, cited 2013.
Percolation	The gravity flow of water within soil. American Meteorological Society, cited 2013: Percolation. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Percolation]
Pfafstetter Coding System	A methodology for assigning watershed (basin) identifications based on the topology of the land surface. It is a hierarchical system where watersheds are delineated from junctions on a river network based on levels. The base level (Level 1) corresponds to the continental scale while higher levels (Levels 2, 3, 4, etc.) represent finer details of the watershed. Watershed Topology: The Pfafstetter System by J. Furnans and F. Olivera (http://proceedings.esri.com/library/userconf/proc01/professional/abstracts/a1008.html), cited 2013.
Precipitation Efficiency	The percentage of the total volume of moisture transported upward to the volume of precipitation received at ground-level over the lifetime of a precipitating system.
Probability of Detection (POD)	A verification measure of categorical forecast performance equal to the total number of correct event forecasts (hits) divided by the total number of events observed. Simply stated, it is the percent of events that are forecast. National Weather Service, Space Weather Prediction Center: Forecast Vertification Glossary, cited 2013.
Propagation	1. The component of convective storm motion that does not lie along the passive steering layer flow. It is the result of new cell development and old cell dissipation. 2. The transmission of electromagnetic energy as waves through or along a medium.
Proximity Sounding	Atmospheric properties (i.e. temperature, dew point, and wind) plotted on a thermodynamic diagram (usually a SKEW-T diagram) that represent the atmospheric environment associated with a particular meteorolocial event such as a tornadic supercell.
Pulse Storm	An ordinary cell with a stronger than typical updraft phase, during and immediately after which the storm produces a short episode of severe weather.
Quasi-Linear Convective System (QLCS)	A linear mesoscale convective system (MCS) that has taken the form of a squall line or bow echo.

Radar	(Coined word for radio detection and ranging.) An electronic instrument used for the detection and ranging of distant objects of such composition that they scatter or reflect radio energy. A radar consists of a transmitter, receiver, antenna, display, and associated equipment for control and signal processing. The most common radars are monostatic radars, which use the same antenna for both transmission and reception. These radars depend on backscattering to produce a detectable echo from a target. Bistatic radars have the transmitter and its antenna at one location and the receiver and its antenna at a remote location. These radars depend upon forward scattering to produce a detectable signal. Radio energy emitted by the transmitter and focused by the antenna of a monostatic radar propagates outward through the atmosphere in a narrow beam. Objects lying in the path of the beam reflect, scatter, and absorb the energy. A small portion of the reflected and scattered energy, called the target signal, travels back along the same path through the atmosphere and is intercepted by the receiving antenna. The time delay between the transmitted signal and the target signal is used to determine the distance or slant range of the target from the radar. The direction in which the focused beam is pointing at the instant the target signal is received (i.e., the azimuth and elevation angles of the antenna) determine the direction and height of the target. This information is presented visually as echoes on different types of radar displays. Because hydrometeors scatter radio energy, weather radars, operating in certain radar frequency bands, can detect the presence of precipitation and other weather phenomena at distances up to several hundred kilometers from the radar, depending upon meteorological conditions and the type of radar. MST radars and wind profilers, which operate at longer wavelengths than weather radars, are able to detect echoes from optically clear air that are caused by spatial fluctuations of refractivity. Additional information provided by a radar about a target may include the radial velocity or rate of change of range, as measured by a Doppler radar, or the depolarizing characteristics of the target, as measured by a polarimetric radar. American Meteorological Society, cited 2013: Radar. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Radar]
Radar Beam	The focused electromagnetic emissions from a radar antenna. The beam is defined by the main lobe of the antenna pattern. American Meteorological Society, cited 2013: Radar beam. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Radar_beam]

Radar Frequency Band	A frequency band of microwave radiation within which radars operate. *See also* C band, S band, X band. American Meteorological Society, cited 2013: Radar Frequency Band. Glossary of Meteorology. [Available online at http://http://glossary.ametsoc.org/wiki/Radar_frequency_band]
Radar Mosaic	A product that combines information from multiple radars to give a regional or national view of reflectivity, precipitation, echo top, etc.
Rain Foot	Slang for a horizontal bulging near the surface in a precipitation shaft, forming a foot-shaped prominence. It is a visual indication of a wet microburst. National Weather Service Glossary, cited 2013.
Rain-free Base	A dark, horizontal cloud base with no visible precipitation beneath it. It typically marks the location of the thunderstorm updraft. Tornadoes may develop from wall clouds attached to the rain-free base, or from the rain-free base itself. Note that the rain-free base may not actually be rain free; hail or large rain drops may be falling. For this reason, updraft base is more accurate. National Weather Service Glossary, cited 2013.
Range Aliasing	(*Also called* range folding.) A radar sampling problem that arises when echoes located beyond the maximum unambiguous range (rmax) are received as if they were within this range of the radar. This occurs when the radar receives a signal return from a pulse other than the most recent pulse. In this case, the radar sends out a pulse (a short burst of energy). This pulse will continue to go in a straight line until it strikes a target. When it strikes the target, a portion of the pulse will be back scattered towards the radar. If the target it strikes is well beyond the normal range of the radar, it will take longer for the back scattered energy to arrive back at the radar. As a result, the radar will most likely have sent out another pulse in the same direction before the back scattered energy arrives back at the radar. Therefore, when the radar receives the back scattered energy, it will assume that it came from an object much closer to the radar and it will improperly locate the echo. A multiple-trip return appears at the difference of the true range and a multiple of the unambiguous range, i.e., R_displayed = R_true - n * R_max, where n = 0,1,2,... American Meteorological Society, cited 2013: Range aliasing. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Range_aliasing] National Weather Service Glossary, cited 2013.
Range Gate	The discrete point in range along a single radial of radar data at which the received signal is sampled. Range gates are typically spaced at 100-1000 meter intervals. A "radial" of radar data is composed of successive range gates, out to the maximum unambiguous range. National Weather Service Glossary, cited 2013.

Range Resolution	The least radial separation between two targets in the same direction from a radar that allows them to be distinguished. This separation equals one-half the transmitted pulse length. Targets closer together than this distance are not resolved and appear as a single target on the display. American Meteorological Society, cited 2013: Range Resolution. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Range_resolution]
Range Unfolding	Process of removing range ambiguity in apparent range of a multitrip target on the radar. National Weather Service Glossary, cited 2013.
Rayleigh Scattering	Changes in directions of electromagnetic energy by particles whose diameters are 1/16 wavelength or less. This type of scattering is responsible for the sky being blue. National Weather Service Glossary, cited 2013.
Rear Flank Downdraft (RFD)	A region of dry air subsiding on the back side of, and wrapping around, a mesocyclone. It often is visible as a clear slot wrapping around the wall cloud. Scattered large precipitation particles (rain and hail) at the interface between the clear slot and wall cloud may show up on radar as a hook or pendant; thus the presence of a hook or pendant may indicate the presence of an RFD. *See also* forward flank downdraft (FFD). National Weather Service Glossary, cited 2013.
Rear Flank Downdraft Gust Front (RFDGF)	A gust front associated with a rear flank downdraft (RFD).
Rear Flank Downdraft Internal Surge (RFDIS)	A secondary rear flank downdraft (RFD) embedded within a pre-existing RFD. RFDIS events appear to influence tornado development, intensity, and demise by altering the thermodynamic and kinematic character of the RFD region bounding the pretornadic and tornadic circulations.
Rear Flank Downdraft Internal Surge Boundary (RFDISB)	A gust front associated with a rear flank downdraft internal surge (RFDIS).
Rear Inflow Jet (RIJ)	A mesoscale region of strong winds that originate in the trailing stratiform rainfall region of a squall line near the top of the cold pool and are directed toward the leading edge.
Rear-Inflow Notch	(*Also called* weak echo channel). A radar signature of a rear inflow jet (RIJ) characterized by an indentation or channel in the reflectivity pattern that originates in the trailing stratiform rainfall region of a squall line near the top of the cold pool and is directed toward the leading edge. *See also* inflow notch.

Richards Equation	A tool used to define the threshold for rainfall runoff, and ultimately flash flood guidance, which states that the infiltration rate of a soil is proportional to the ease with which water can move through its pore space.
Right Mover	A thunderstorm (usually a supercell) that moves appreciably to the right relative to the main steering winds and to other nearby thunderstorms. Right movers typically are associated with a high potential for severe weather. *Compare* left mover. National Weather Service Glossary, cited 2013.
River Forecast Center (RFC)	A National Weather Service (NWS) office which provides hydrologic guidance to weather forecast offices (WFO) and is the first echelon office for the preparation of river and flood forecasts and warnings.
Roll Cloud	A low-level, horizontal, tube-shaped arcus cloud associated with a gust front of a convective storm (or occasionally a cold front). Roll clouds are relatively rare; they are completely detached from the convective storm's cloud base, thus differentiating them from the more familiar shelf clouds. Roll clouds appear to be rolling about a horizontal axis because of the shearing effects and horizontal vorticity provided by the differing air masses. American Meteorological Society, cited 2013: Roll Cloud. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Roll_cloud]
Rope	(*Also called* "Rope Funnel") A narrow, often contorted condensation funnel usually associated with the decaying stage of a tornado. *See also* rope stage. National Weather Service Glossary, cited 2013.
Rope Stage	The dissipating stage of a tornado, characterized by thinning and shrinking of the condensation funnel into a rope (or rope funnel). Damage still is possible during this stage. National Weather Service Glossary, cited 2013.
Rotational Velocity (Vr)	The single site Doppler radar intensity of an atmospheric circulation (i.e., mesocyclone, TVS, or TS) as quantified by the sum of the absolute value of the minimum radial velocity (Vmin) and the absolute value of the maximum radial velocity (Vmax) divided by two. *Compare* velocity difference (DV). *See also* rotational velocity (Vr) shear.
Rotational Velocity (Vr) Shear	The single site Doppler radar intensity of an atmospheric circulation (i.e., mesocyclone, TVS, or TS) as quantified by the sum of the absolute value of the minimum radial velocity (Vmin) and the absolute value of the maximum radial velocity (Vmax) divided by the distance

	between velocity peaks. Values are on the order of 10-2 s-1 for mesocyclones.
S-Band Radar	A radar which operates in the 2-4 GHz wavelength and 7.5-15 cm wavelength ranges. Because of these characteristics, S band radars are not easily attenuated. This makes them useful for both near and far range weather observation. The National Weather Service (NWS) WSR-88D is an S band rada which operates on a wavelength of just over 10 cm. The drawback to this band of radar is that it requires a large antenna dish and a large motor to power it. *See also* C band radar, X band radar.
Severe Local Storm	A convective storm that usually covers a relatively small geographic area, or moves in a narrow path, and is sufficiently intense to threaten life and/or property. Examples include severe thunderstorms with large hail, damaging wind, or tornadoes. Although cloud-to-ground lightning is not a criteria for severe local storms, it is acknowledged to be highly dangerous and a leading cause of deaths, injuries, and damage from thunderstorms. A thunderstorm need not be severe to generate frequent cloud-to-ground lightning. Additionally, excessive localized convective rains are not classified as severe storms but often are the product of severe local storms. Such rainfall may result in related phenomena (flash floods) that threaten life and property. National Weather Service Glossary, cited 2013.
Severe Thunderstorm	A thunderstorm that produces a tornado, winds of at least 58 mph (50knots), and/or hail at least 1" in diameter. Structural wind damage mayimply the occurrence of a severe thunderstorm. A thunderstorm wind equalto or greater than 40 mph (35 knots) and/or hail of at least 1" is defined as approaching severe. National Weather Service Glossary, cited 2013.
Severe Thunderstorm Warning	A National Weather Service (NWS) Weather Forecast Office (WFO) product issued when either a severe thunderstorm is indicated by radar or a spotter reports a thunderstorm producing hail one inch or larger in diameter and/or winds equal or exceed 58 miles an hour. They are usually issued for a duration of 30 minutes to one hour. Lightning frequency is not a criteria for issuing a severe thunderstorm warning. If the severe thunderstorm will affect the nearshore or coastal waters, it will be issued as the combined product--Severe Thunderstorm Warning and Special Marine Warning.

Severe Thunderstorm Watch	A National Weather Service (NWS) product issued when conditions are favorable for the development of severe thunderstorms in and close to the watch area. The size of the watch can vary depending on the weather situation. They are normally issued well in advance of the actual occurrence of severe weather with a typical duration of 4 to 8 hours. A Severe Thunderstorm Watch is issued by the Storm Prediction Center in Norman, Oklahoma. Prior to the issuance, SPC will usually contact the affected local NWS Weather Service Forecast Office (WFO) and they will discuss what their current thinking is on the weather situation. Afterwards, SPC will issue a preliminary Severe Thunderstorm Watch and then the affected WFO will then adjust the watch (adding or eliminating counties/parishes) and then issue it to the public by way of a Watch Redefining Statement. During the watch, the WFO will keep the public informed on what is happening in the watch area and also let the public know when the watch has expired or been cancelled. National Weather Service Glossary, cited 2013.
Severe Weather Statement	A National Weather Service (NWS) Weather Forecast Office (WFO) product which provides follow up information on severe weather conditions (severe thunderstorm or tornadoes) which have occurred or are currently occurring. National Weather Service Glossary, cited 2013.
Shear	Variation in wind speed (speed shear) and/or direction (directional shear) over a short distance within the atmosphere. Shear usually refers to vertical wind shear, i.e., the change in wind with height, but the term also is used in Doppler radar to describe changes in radial velocity over short horizontal distances. National Weather Service Glossary, cited 2013.
Shelf Cloud	A low-level, horizontal wedge-shaped arcus cloud, associated with a convective storm's gust front (or occasionally with a cold front, even in the absence of thunderstorms). Unlike the roll cloud, the shelf cloud is attached to the base of the parent cloud above it (usually a thunderstorm). Rising cloud motion often can be seen in the leading (outer) part of the shelf cloud, while the underside often appears turbulent, boiling, and wind-torn. National Weather Service Glossary, cited 2013.
Squall Line	A line of active thunderstorms, either continuous or with breaks, including contiguous precipitation areas resulting from the existence of the thunderstorms. The squall line is a type of mesoscale convective system (MCS) distinguished from other types by a larger length-to-width ratio. American Meteorological Society, cited 2013: Squall line. Glossary of

	Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Squall_line]
Steering Winds	(*Also called* steering currents, steering flow). A prevailing synoptic scale flow which governs the movement of smaller features embedded within it. National Weather Service Glossary, cited 2013.
Storm Motion	The speed and direction at which a thunderstorm travels. National Weather Service Glossary, cited 2013.
Storm Prediction Center (SPC)	A National Weather Service (NWS) National Center for Environmental Prediction (NCEP) in Norman, Oklahoma, which is responsible for providing short-term forecast guidance for severe convection, excessive rainfall (flash flooding), fire weather, and severe winter weather over the contiguous United States.
Storm Scale	Referring to weather systems with sizes on the order of individual thunderstorms. *See also* synoptic scale, mesoscale. National Weather Service Glossary, cited 2013.
Storm-Relative	Measured relative to a moving thunderstorm, usually referring to winds, wind shear, or helicity. National Weather Service Glossary, cited 2013.
Streamwise Vorticity	The component of vorticity that is parallel to the ambient velocity vector. *Compare* crosswise vorticity. *See also* helicity. American Meteorological Society, cited 2013: Streamwise Vorticity. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Streamwise_vorticity]
Striations	Grooves or channels in cloud formations, arranged parallel to the flow of air and therefore depicting the airflow relative to the parent cloud. Striations often reveal the presence of rotation, as in the barber pole or "corkscrew" effect often observed with the rotating updraft of a Low Precipitation (LP) supercell. National Weather Service Glossary, cited 2013.
Suction Vortices	Smaller-scale secondary vortices within a tornado core that orbit around a central axis. The transition of a one-celled vortex into secondary vortices in laboratory and numerical simulations occurs at high swirl ratios. The vortices produce cycloidal swaths within tornado damage tracks and are often used to explain the gradation of wind damage caused by a tornado. Structures in the path of a suction vortex are damaged while others are spared.

	American Meteorological Society, cited 2013: Suction Vortices. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Suction_vortices]
Supercell	An often dangerous convective storm that consists primarily of a single, quasi-steady rotating updraft (i.e., a mesocyclone), which persists for a period of time much longer than it takes an air parcel to rise from the base of the updraft to its summit (often much longer than 10–20 min). Most rotating updrafts (in the Northern Hemisphere) are characterized by cyclonic vorticity (see mesocyclone). The supercell typically has a very organized internal structure that enables it to propagate continuously. It may exist for several hours and usually forms in an environment with strong vertical wind shear. Supercells propagate in a direction and with a speed other than indicated by the mean wind in the environment. Such storms sometimes evolve through a splitting process, which produces a cyclonic, right-moving (with respect to the mean wind), and anticyclonic, left-moving, pair of supercells. Severe weather often accompanies supercells, which are capable of producing high winds, large hail, and strong, long-lived tornadoes. American Meteorological Society, cited 2013: Supercell. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Supercell] Supercells comprise a spectrum, but are often sub-categorized based on the extent to which their mesocyclone is wrapped in precipitation as revealed by their radar and/or visual appearance. A low precipitation (LP) supercell (also called a dryline storm) is dominated by updraft with little precipitation reaching the ground. It is visualized by an exposed updraft and a translucent to nearly transparent precipitation core. Low-level mesocyclones and tornadoes are rare owing to the lack of a well defined rear flank downdraft (RFD). Most of the precipitation is carried well downstream of the updraft by strong (>30 m/s or 58 kt) anvil-layer winds. A high precipitation (HP) supercell (also called an HP storm) is a highly efficient precipitation producer that develops and maintains precipitation-filled rear flank downdrafts (RFDs) that often envelop the mesocyclones. This makes visual identification of any embedded tornadoes difficult and very dangerous. HP supercells often produce large damaging hail, extreme and prolonged downbursts, and flash flooding. A classic supercell, which falls in between these two extremes, exhibits moderate precipitation production. While there may be some precipitation with a classic supercell's rear flank downdraft (RFD) (and hook echo), its radar reflectivities will be lower than its foward flank

	downdraft (FFD) precipitation core.
Supercooled Water	(*Also called* supercooled liquid water). Liquid water at a temperature less than the freezing point. Important in the formation of graupel and hail.
Surface-based Convection	Convection occurring within a surface-based layer, i.e., a layer in which the lowest portion is based at or very near the earth's surface. *Compare* elevated convection. National Weather Service Glossary, cited 2013.
Synoptic Scale	(*Also called* large scale). Size scale which refers to weather systems ranging in size from several hundred kilometers to several thousand kilometers, the scale of migratory high and low pressure systems (frontal cyclones) of the lower troposphere. *Compare* mesoscale, storm-scale. American Meteorological Society, cited 2013: Synoptic scale. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Synoptic_scale]
Tail Cloud	A cloud band, often laminar and tube-shaped, but sometimes ragged and turbulent, the extends from the forward-flank downdraft precipitation cascade region of a supercell toward the wall cloud. Cloud motion in the tail cloud is away from the precipitation and toward the wall cloud, with rapid upward motion often observed near the junction of the tail and wall clouds. Compare with beaver tail, which is a form of inflow band that normally attaches to the storm's main updraft (not to the wall cloud) and has a base at about the same level as the updraft base (not the wall cloud).
Tail-End Charile	[Slang] A thunderstorm at the southernmost end of a squall line or other line or band of thunderstorms. Since low-level southerly inflow of warm, moist air into this storm is relatively unimpeded, such a storm often has a higher probability of strengthening to severe levels than the other storms in the line. National Weather Service Glossary, cited 2013.
Thermodynamics	In general, the relationships between heat and other properties (such as temperature, pressure, density, etc.). In forecast discussions, thermodynamics usually refers to the distribution of temperature and moisture (both vertical and horizontal) as related to the diagnosis of atmospheric instability. National Weather Service Glossary, cited 2013.

Three-Body Scatter Spike (TBSS)	(*Also called* hail spike.) A radar artifact caused by radar microwave scattering associated with large hydrometeors, typically severe hail. The TBSS is strictly an artifact of the electromagnetic radar beam being subject to “Mie scattering” instead of the usual “Rayleigh scattering” process. A TBSS forms as incident energy from the radar is reflected off the hail, down to the ground, then back up to the hail and back to the radar. Because of the delay in reception of the pulses, the radar circuitry displays the TBSS as downrange from the hail core. The TBSS is characterized by a 10-30 km (5-16 nm) long, low reflectivity (< 25 dBZ), echo “spike” aligned radially downrange from a high reflectivity (usually > 63 dBZ) core. The TBSS signature also produces low radial velocities (V), high spectrum widths (SW), extremely low correlation coefficients (CC), and extremely positive Differential Reflectivity (ZDR) transitioning into lower positive or even negative values farther down-radial. The presence of a TBSS with reflectivities greater than 5 dBZ on a S-band (10 cm) radar (such as the WSR-88D) suggests that the thunderstorm possesses severe hail.
Thunderstorm	In general, a local storm, invariably produced by a cumulonimbus cloud and always accompanied by lightning and thunder, usually with strong gusts of wind, heavy rain, and sometimes with hail. A thunderstorm is a consequence of atmospheric instability and constitutes, loosely, an overturning of air layers in order to achieve a more stable density stratification. A strong convective updraft is a distinguishing feature of this storm in its early phases. A strong downdraft in a column of precipitation marks its dissipating stages. Thunderstorms often build to altitudes of 40 000–50 000 ft in midlatitudes and to even greater heights in the Tropics; only the great stability of the lower stratosphere limits their upward growth. American Meteorological Society, cited 2013: Thunderstorm. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Thunderstorm]
Thunderstorm Cell	The convective cell of a cumulonimbus cloud having lightning and thunder. American Meteorological Society, cited 2013: Thunderstorm Cell. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Thunderstorm_cell]
Time of Concentration	The time it takes for runoff to travel from the hydraulically most distant points of the basin to the basin outlet.

Tornadic Vortex Signature (TVS)	A Doppler velocity identified tornado-scale vortex whose diameter, measured by the distance between the maximum and minimum radial velocity, is smaller than the effective beam width. With azimuthal over-sampling, as is the case with WSR-88D super-resolution data, the velocity peaks can be separated by up to three radials. As the signature occurs when the radar beam is wider than the vortex, the measured Doppler velocities are weaker than the rotational velocities within the vortex and the apparent core diameter is larger than that of the vortex. The signature, which may extend throughout a considerable vertical depth, is ideally characterized by extreme Doppler velocity values of opposite sign separated in azimuth by the equivalent of one beamwidth. However, since most radars display and record Doppler velocity values at discrete azimuthal intervals, the extreme Doppler velocity values are usually at azimuthally adjacent positions that are roughly one beamwidth apart. *See also* tornado signature (TS)
Tornado	A violently rotating column of air, in contact with the earth's surface, either pendant from a cumuliform cloud or underneath a cumuliform cloud, and often (but not always) visible as a funnel cloud. When tornadoes do occur without any visible funnel cloud, debris at the surface is usually the indication of the existence of an intense circulation in contact with the ground. *See also* landspout and waterspout.
Tornado Debris Signature (TDS)	A unique radar feature which indicates objects lofted by a violently rotating colum of air in contact with the ground. It typically possesses a localized correlation coefficient (CC) minimum with values less than 0.9 (typically less than 0.8), centered near a storm-scale vortex (i.e., mesocyclone, TVS, or TDS), and reflectivities at least 40 dBZ. Differential Reflectivity (ZDR) values are typically near zero in valid tornado debris, but this signature is not nearly as pronounced as CC. Neverthless, ZDR can bu used as a confirmatory check. *See also* debris ball.
Tornado Emergency	Enhanced wording that National Weather Service (NWS) Weather Forecast Offices (WFOs) can insert into tornado warnings and severe weather statements that describes an exceedingly rare situation when a severe threat to human life and catastrophic damage from a tornado is imminent or ongoing. Use of “TORNADO EMERGENCY” terminology is appropriate for the tornadic situation if all of the following criteria are met: a. Severe threat to human life is imminent or ongoing b. Catastrophic damage is imminent or ongoing c. Reliable sources confirm tornado (either 1 or 2) (1) Visual (2) Radar imagery strongly suggests the existence of a damaging tornado (e.g., debris ball signature).

Tornado Signature (TS)	A Doppler velocity identified tornado-scale vortex which is larger than or equal to the effective beam width. *See also* tornado vortex signature.
Tornado Warning	A National Weather Service (NWS) Weather Forecast Office (WFO) product issued when when a tornado is indicated by weather radar or sighted by spotters. They are usually issued for a duration of 30-45 minutes. If the tornado will affect the nearshore or coastal waters, it will be issued as the combined product--Tornado Warning and Special Marine Warning. *See also* tornado emergency.
Tornado Watch	A National Weather Service (NWS) product issued when conditions are favorable for the development of tornadoes in and close to the watch area. Their size can vary depending on the weather situation.They are normally issued well in advance of the actual occurrence of severe weather with a typical duration of 4 to 8 hours. A Tornado Watch is issued by the Storm Prediction Center (SPC) in Norman, Oklahoma. Prior to the issuance of a Tornado Watch, SPC will usually contact the affected local NWS Weather Forecast Office (WFO) and they will discuss what their current thinking is on the weather situation. Afterwards, SPC will issue a preliminary Tornado Watch and then the affected WFO will then adjust the watch (adding or eliminating counties/parishes) and then issue it to the public. After adjusting the watch, the WFO will let the public know which counties are included by way of a Watch Redefining Statement. During the watch, the NWFO will keep the public informed on what is happening in the watch area and also let the public know when the watch has expired or been cancelled. National Weather Service Glossary, cited 2013.
Tornadogenesis	The process by which a tornado forms. *See also* mesocyclonic tornado, non-mesocyclonic tornado.
Torus	A doughnut-shaped surface generated by a circle rotated about an axis in its plane that does not intersect the circle. For a hailstone, a torus is the meltwater which forms a band around the equator of the hailstone due to drag as it falls through the atmosphere.
Training	(*Also called* train echo). Repeated areas of rain, typically associated with thunderstorms, that move over the same region in a relatively short period of time and are capable of producing excessive rainfall totals. Train(ing) echoes can frequently be a source of flash flooding. National Weather Service Glossary, cited 2013.

Transverse Bands	Bands of clouds oriented perpendicular to the flow in which they are embedded. They often are seen best on satellite photographs. When observed at high levels (i.e., in cirrus formations), they may indicate severe or extreme turbulence. Transverse bands observed at low levels (called transverse rolls or T rolls) often indicate the presence of a temperature inversion (or cap) as well as directional shear in the low- to mid-level winds. These conditions often favor the development of strong to severe thunderstorms. National Weather Service Glossary, cited 2013.
Transverse Rolls	Elongated low-level clouds, arranged in parallel bands and aligned parallel to the low-level winds but perpendicular to the mid-level flow. Transverse rolls are one type of transverse band, and often indicate an environment favorable for the subsequent development of supercells. Since they are aligned parallel to the low-level inflow, they may point toward the region most likely for later storm development. National Weather Service Glossary, cited 2013.
Triple Point	A junction point of three distinct airmasses denoted by the intersection point between two boundaries (dry line, outflow boundary, cold front, etc.), often a focus for thunderstorm development. Triple point also may refer to a point on the gust front of a supercell, where the warm moist inflow, the rain-cooled outflow from the forward flank downdraft, and the rear flank downdraft all intersect; this point is a favored location for tornado development (or redevelopment).
Universal Time (UT)	(*Also called* Coordinated Universal Time or Universal Time Coordinated (UTC), Z time, Zulu Time). By international agreement, the local time at the prime meridian, which passes through Greenwich, England. Prior to 1972, this time was called Greenwich Mean Time (GMT) but is now referred to as Coordinated Universal Time or Universal Time Coordinated (UTC). It is a coordinated time scale, maintained by the Bureau International des Poids et Mesures (BIPM). National Weather Service Glossary, cited 2013.
Updraft	A small-scale current of rising air. If the air is sufficiently moist, then the moisture condenses to become a cumulus cloud or an individual tower of a towering cumulus or Cumulonimbus. *Compare* downdraft. National Weather Service Glossary, cited 2013.
Upper Air	(*Also called* upper level). In synoptic meteorology and in weather observing, that portion of the atmosphere that is above the lower troposphere. No distinct lower limit is set but the term can be generally applied to the levels above 850 mb. *See also* low-level, midlevel.:

	American Meteorological Society, cited 2013: Upper Air. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Upper Air]
Upshear	In the opposite direction as shear vector within a specified layer. *Compare* downshear.
Upwind	In the opposite direction as the wind flow, or opposite the direction in which the wind is moving. *Compare* downwind.
V Notch	A radar reflectivity signature seen as a V-shaped notch in the downwind part of a thunderstorm echo. The V-notch often is seen on supercells, and is thought to be a sign of diverging flow around the main storm updraft (and hence a very strong updraft). This term should not be confused with inflow notch or with enhanced V, although the latter is believed to form by a similar process. National Weather Service Glossary, cited 2013.
Veering Winds	Winds which shift in a clockwise direction with time at a given location (e.g., from southerly to westerly), or which change direction in a clockwise sense with height (e.g., southeasterly at the surface turning to southwesterly aloft). *Compare* backing winds. National Weather Service Glossary, cited 2013.
Velocity Difference (DV)	(*Also called* delta V). The single site Doppler radar intensity of a convergent velocity signature, divergent velocity signature, or atmospheric circulation as quantified by the product of the absolute value of the minimum radial velocity and the absolute value of the maximum radial velocity. *See also* low-level rotational velocity (LLDV).
Vertical Wind Shear	The change in the wind's direction and speed with height. It plays a critical role in the evolution of deep, moist, convection. *See also* wind shear.
Vortex	In its most general use, any flow possessing vorticity. More often the term refers to a flow with closed streamlines or to the idealized case in which all vorticity is concentrated in a vortex filament. American Meteorological Society, cited 2013: Vortex. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Vortex]

Vorticity	A measure of the local rotation in a fluid flow. In weather analysis and forecasting, it usually refers to the vertical component of rotation (i.e., rotation about a vertical axis) and is used most often in reference to synoptic scale or mesoscale weather systems. By convention, positive values indicate cyclonic rotation. *See also* crosswise vorticity, streamwise vorticity.
Wall Cloud	A localized, persistent, often abrupt lowering from a cumulonimbus cloud base into a low-hanging accessory cloud, normally a kilometer or more in diameter. A wall cloud marks the lower portion of a very strong updraft, usually associated with a supercell or severe multicell storm. It typically develops near the precipitation region of the cumulonimbus. Wall clouds that exhibit significant rotation and vertical motions often precede tornado formation by a few minutes to an hour. American Meteorological Society, cited 2013: Wall Cloud. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Wall_cloud]
Warm Cloud	Cloud that is only in the liquid phase; levels are not present with temperature below 0°C (32°F); no ice is present. Any precipitation will originate from droplet coalescence.
Warm Rain Process	In cloud physics, the process producing precipitation through collision between liquid particles (cloud droplets, drizzle drops, and raindrops). The major role of the warm rain process in thunderstorms is to transfer condensed water, in the form of cloud droplets, to precipitable water, in the form of drizzle droplets and raindrops, by the collision–coalescence process.
Waterspout	1. In general, any tornado over a body of water. 2. In its most common form, a non-mesocyclonic tornado over water. *See also* landspout.
Weak Echo Region (WER)	A radar signature characterized by a region of weak reflectivity on the low-altitude inflow side of a thunderstorm topped by stronger reflectivity in the form of a sloping echo overhang directly above. The WER is produced by strong updraft and associated strong storm-summit divergence that carries large amounts of precipitation particles in all directions. This creates a high reflectivity echo-canopy (slopping echo overhang) over the low-level inflow of a strong or intense convective storm.
Weather Forecast Office (WFO)	A National Weather Service (NWS) office responsible for issuing advisories, warnings, statements, and short term forecasts for its county warning area (CWA). National Weather Service Glossary, cited 2013.

Wedge Tornado	(*Also called* wedge.). [Slang] A large tornado with a condensation funnel that is at least as wide (horizontally) at the ground as it is tall (vertically) from the ground to cloud base. The term "wedge" often is used somewhat loosely to describe any large tornado. However, not every large tornado is a wedge. A true wedge tornado, with a funnel at least as wide at the ground as it is tall, is very rare. Wedges often appear with violent tornadoes (EF4 or EF5 on the Fujita Scale), but many documented wedges have been rated lower. And some violent tornadoes may not appear as wedges. Whether or not a tornado achieves "wedge" status depends on several factors other than intensity - in particular, the height of the environmental cloud base and the availability of moisture below cloud base. Therefore, spotters should not estimate wind speeds or EF-scale ratings based on visual appearance alone. However, it generally is safe to assume that most (if not all) wedges have the potential to produce strong (EF2/EF3) or violent (EF4/EF5) damage. National Weather Service Glossary, cited 2013.
Whirlwind	General term for a small-scale, rotating column of air. More specific terms are dust whirl, dust devil, waterspout, and tornado.: American Meteorological Society, cited 2013: Wind Shear. Glossary of Meteorology. [Available online at http://glossary.ametsoc.or/wiki/Wind_shear]
Wind Shear	The rate at which wind velocity changes from point to point in a given direction (as, vertically). The shear can be speed shear (where speed changes between the two points, but not direction), direction shear (where direction changes between the two points, but not speed) or a combination of the two. *See also* vertical wind shear. National Weather Service Glossary, cited 2013.
WSR-88D	Weather Surveillance Radar - 1988 Doppler; NEXRAD unit. National Weather Service Glossary, cited 2013.
X Band Radar	A radar which operates in the 8-12 GHz and 2.5-3.75 cm wavelength ranges. Because of the smaller wavelength, the X band radar is more sensitive and can detect smaller particles. These radars are used for studies on cloud development because they can detect the tiny water particles and also used to detect light precipitation such as snow. X band radars also attenuate very easily, so they are used for only very short range weather observation. Also, due to the small size of the radar, it can therefore be portable. Most major airplanes are equipped with an X band radar to pick up turbulence and other weather phenomenon. This band is also shared with some police speed radars and some space radars.

	See also C band radar, S band radar.
ZDR Arc	A region of high differential reflectivity (ZDR) precipitation echoes that lie along the sharp low-level reflectivity gradient facing the storm-relative inflow. Some of these hydrometeors are from the sloping echo overhang and others are from the edge of the precipitation cascade region. Resarch has theorized that the ZDR arc originates as the precipitation falling from aloft, is sorted by the vertical wind shear present in the environment, and enhanced along the forward flank outflow.
ZDR Column	A dual-polarimetric radar signature of differential reflectivity (ZDR) values above 1-2 dB caused by large liquid drops associated with the updraft of a convective cell+B134
ZDR Ring	A dual-polarimetric radar signature of differential reflectivity (ZDR) values above 1-2 dB horizontally surrounding a bounded weak echo region (BWER) caused by large liquid drops which have been transported around the exterior of a supercell thunderstorm updraft. The ZDR ring is often small and ephemeral making radar detection difficult due to sampling limitations. *See also* low CC ring.
Zulu (Z) Time	For practical purposes, the same as Coordinated Universal Time (UTC). The notation formerly used to identify time Greenwich MeanTime. The word "Zulu" is notation in the phonetic alphabet corresponding to the letter "Z" assigned to the time zone on the Greenwich Prime Meridian. National Weather Service Glossary, cited 2013.