Back to Top

Table of Contents

Topic: Principles of Radar

Click to jump to lesson

Back to Top

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!

Back to Top

1.3 Roadmap

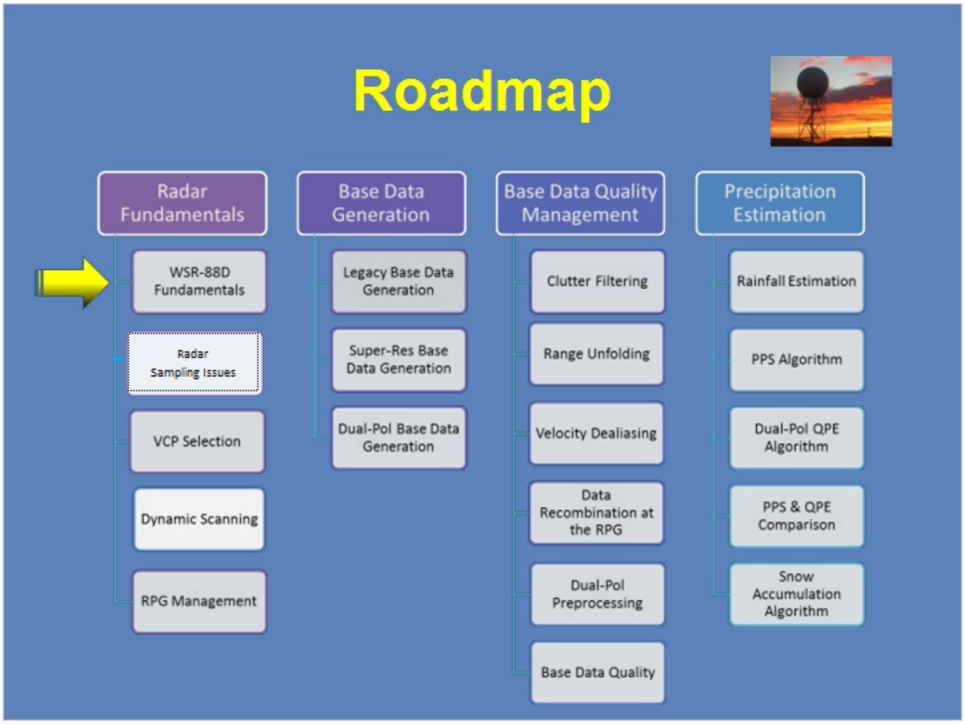

Notes:

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!

Back to Top

1.4 Learning Objectives

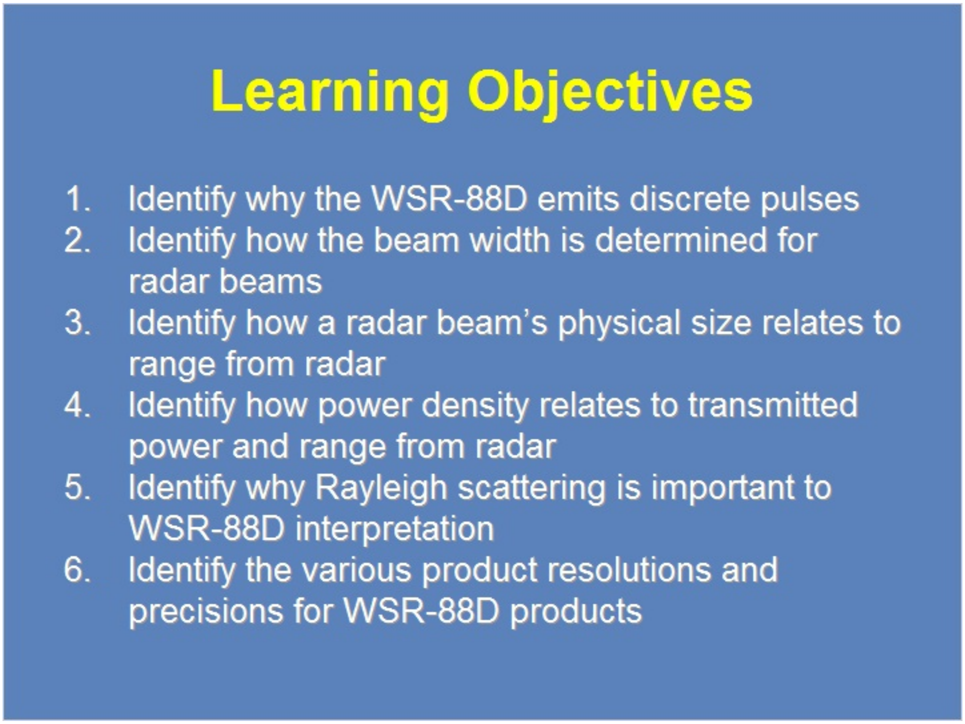

Notes:

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

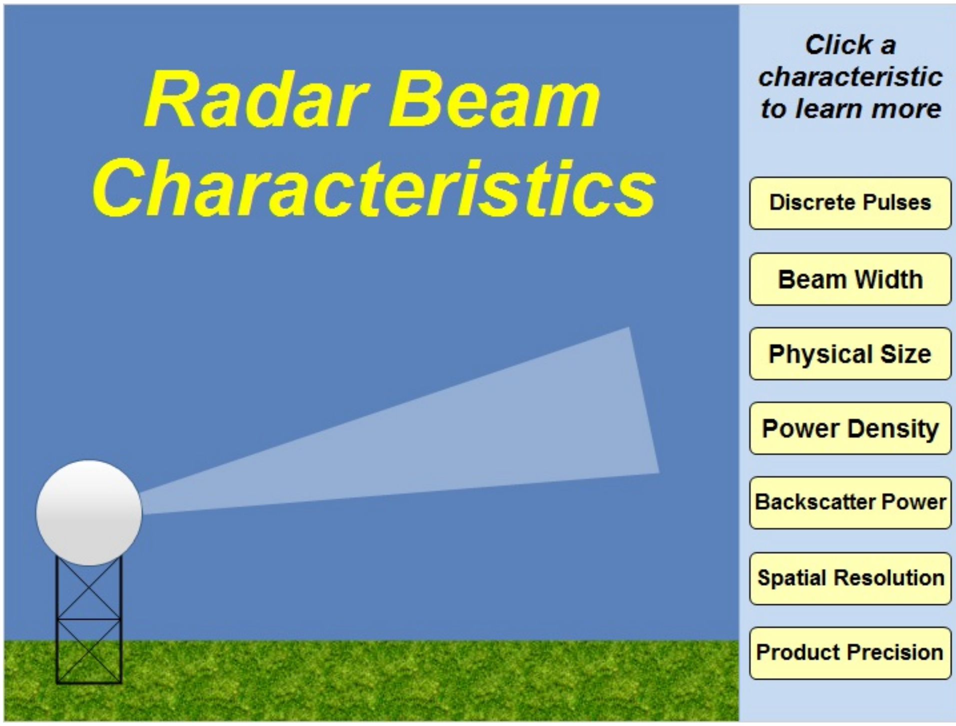

Notes:

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.

Back to Top

2.2 Discrete Pulses

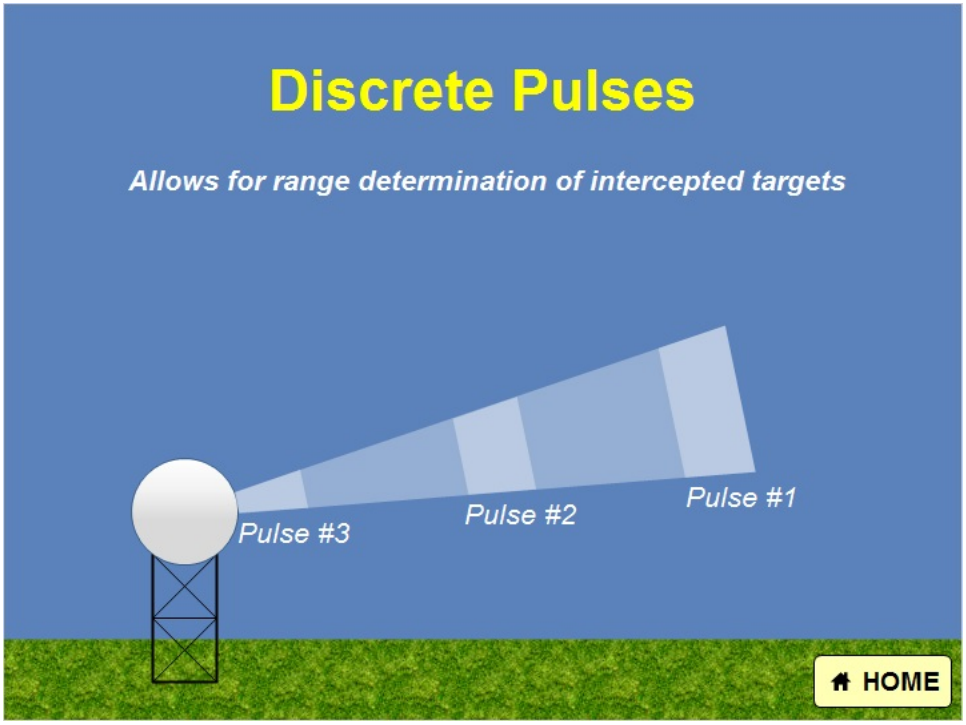

Notes:

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

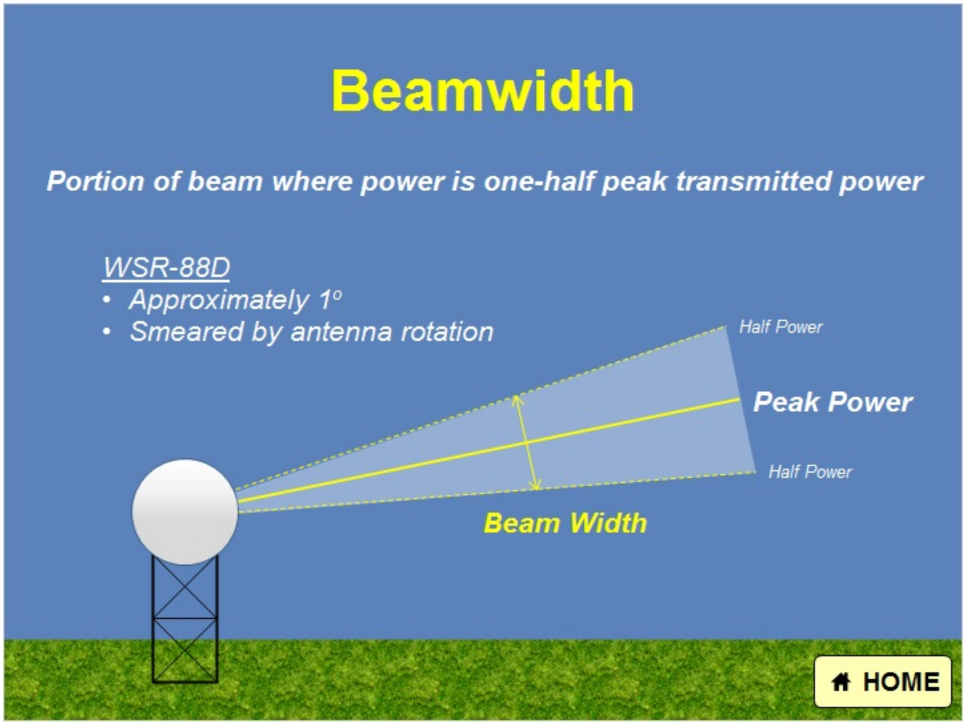

Notes:

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

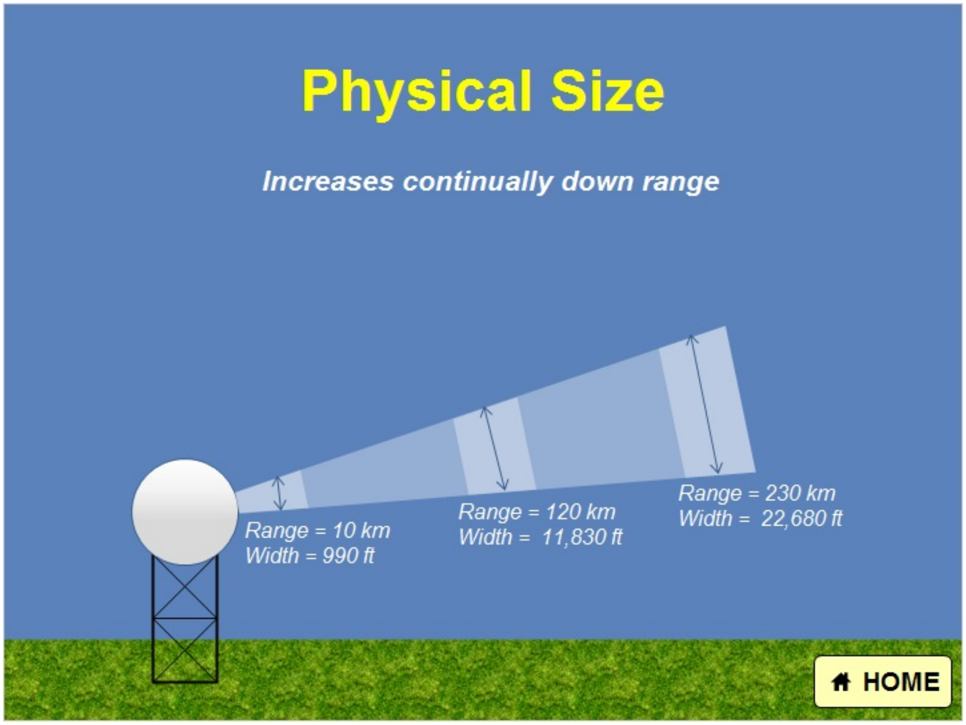

Notes:

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!

Back to Top

2.5 Power Density

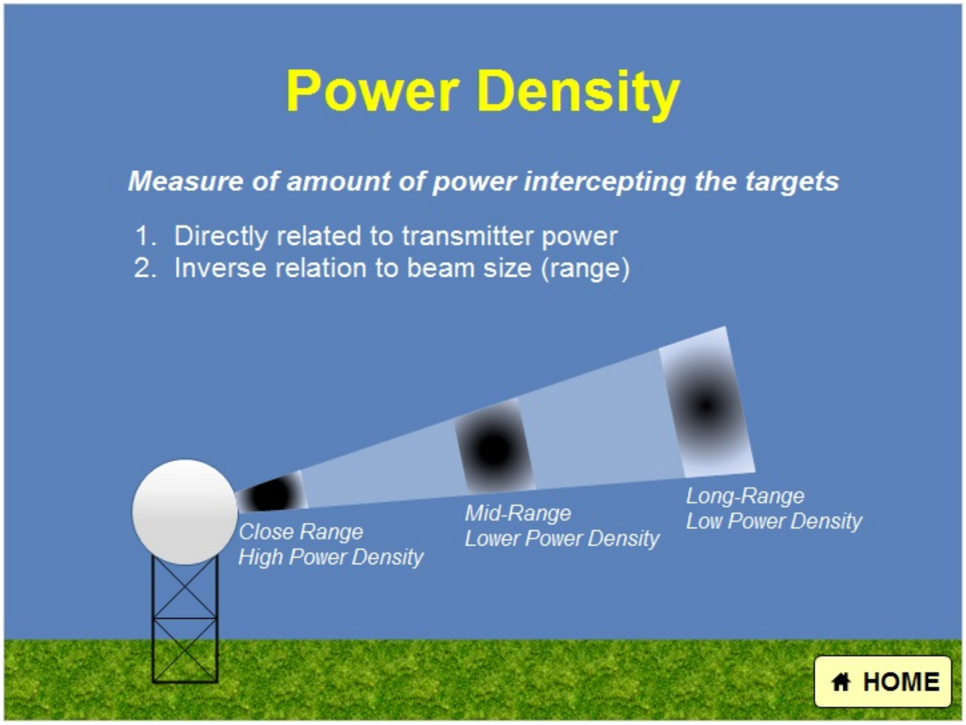

Notes:

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.

Back to Top

2.6 Backscattered Power

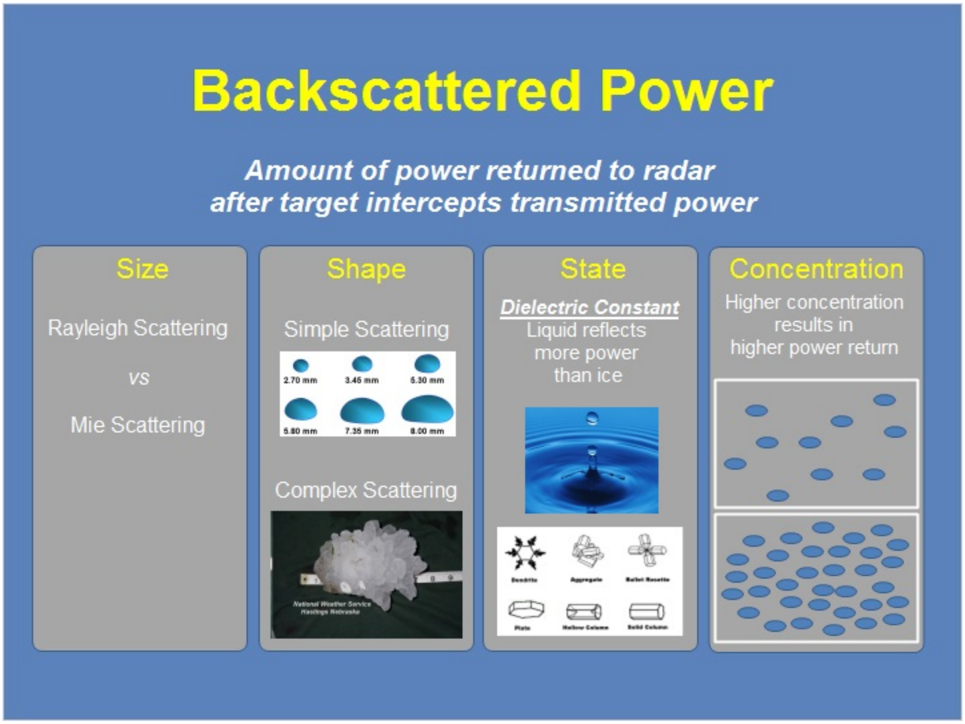

Notes:

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.

Back to Top

2.7 Rayleigh Scattering

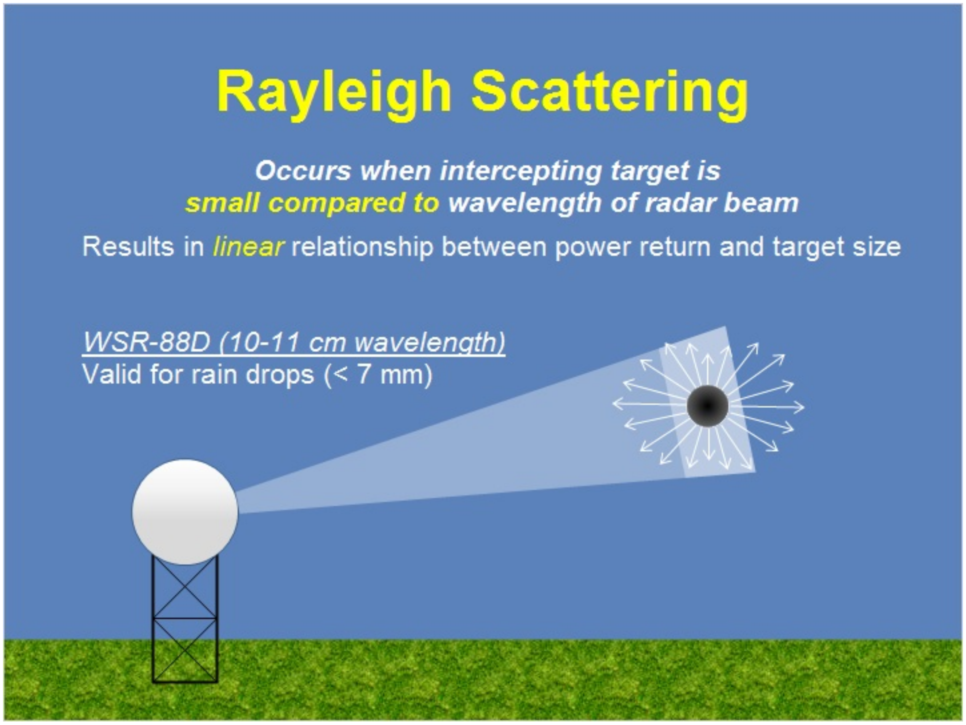

Notes:

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.

Back to Top

2.8 Mie Scattering

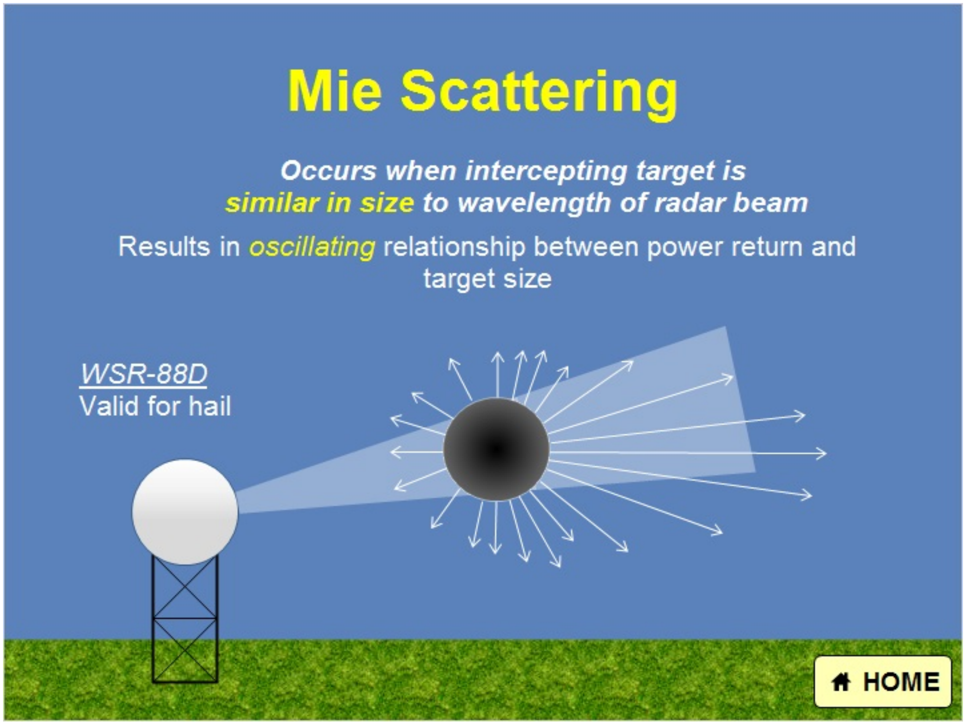

Notes:

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.

Back to Top

2.9 Spatial Resolutions

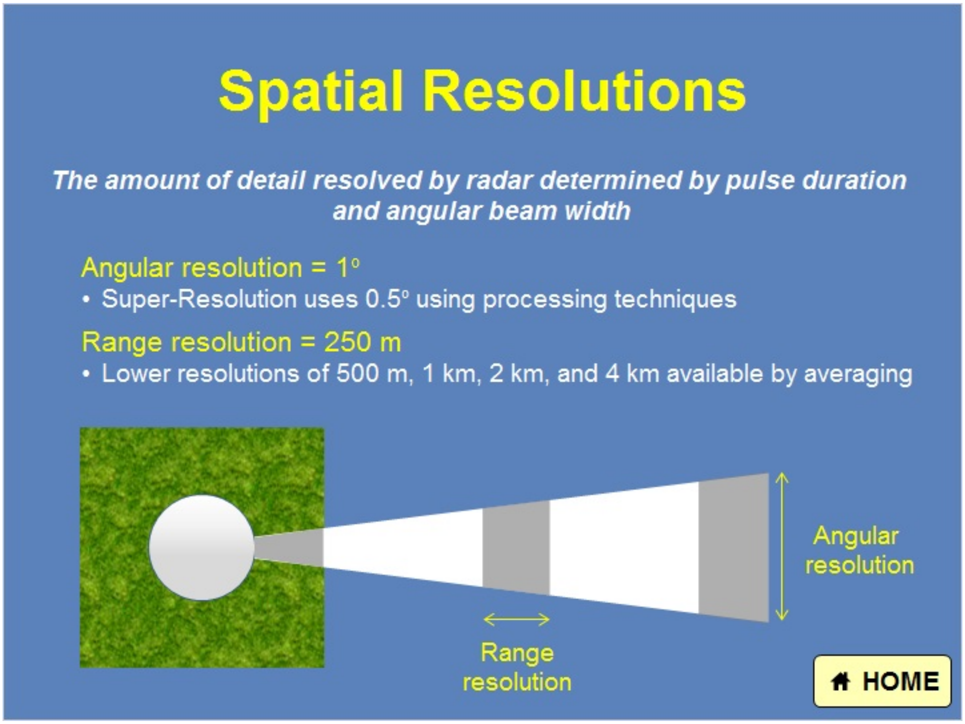

Notes:

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.

Back to Top

2.10 Product Precision

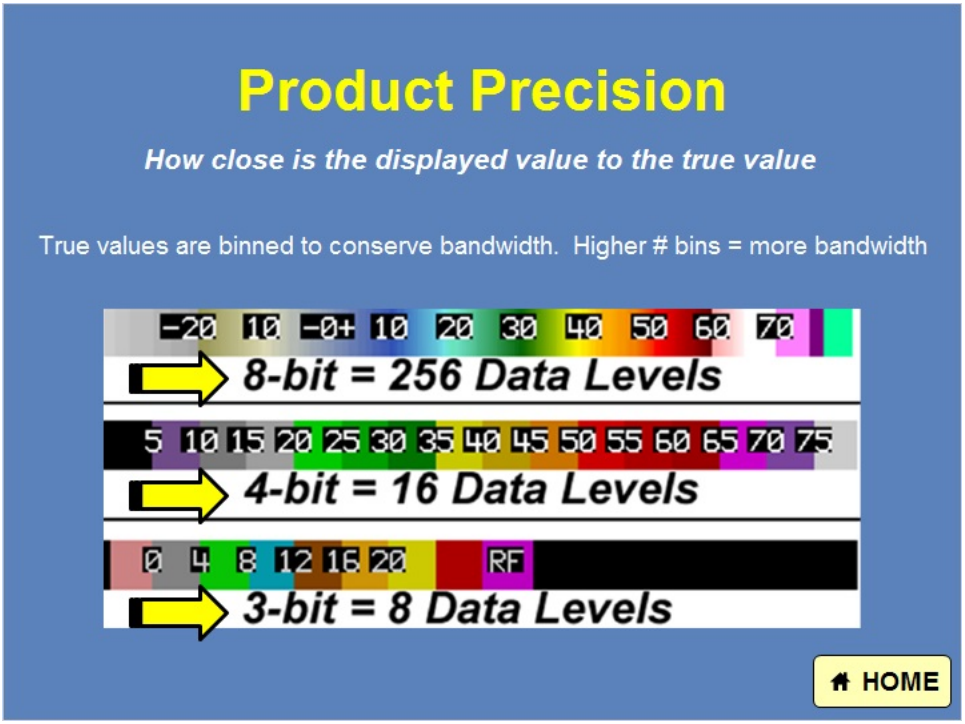

Notes:

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.

4. Completion

4.1 Completion!

Notes:

Thanks for your attention! You are now complete for this lesson.

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!

Back to Top

1.3 Roadmap

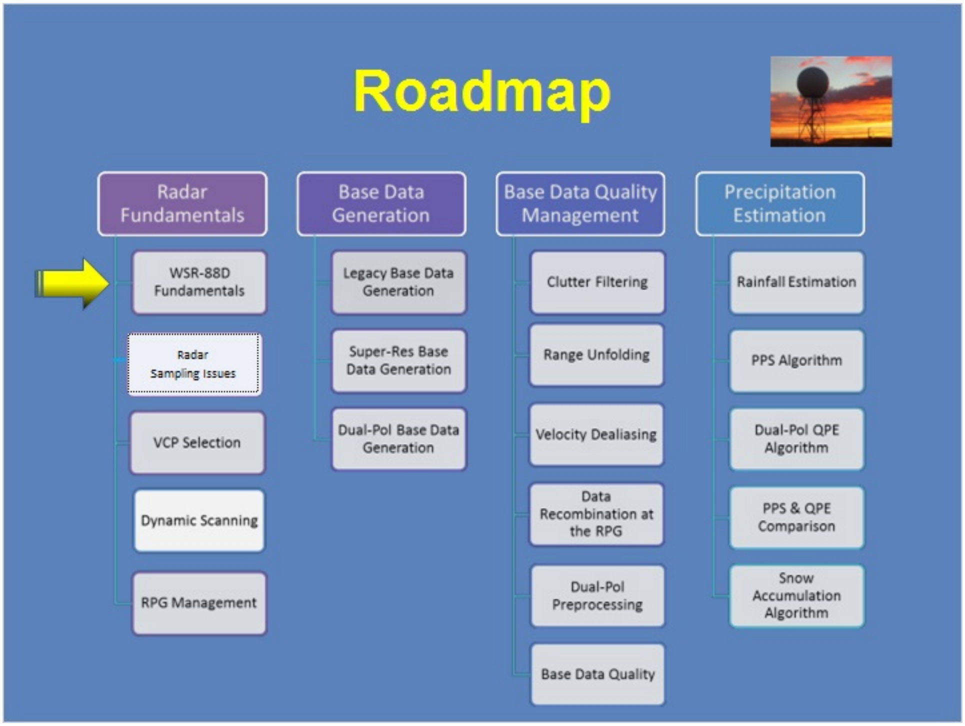

Notes:

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!

Back to Top

1.4 Learning Objectives

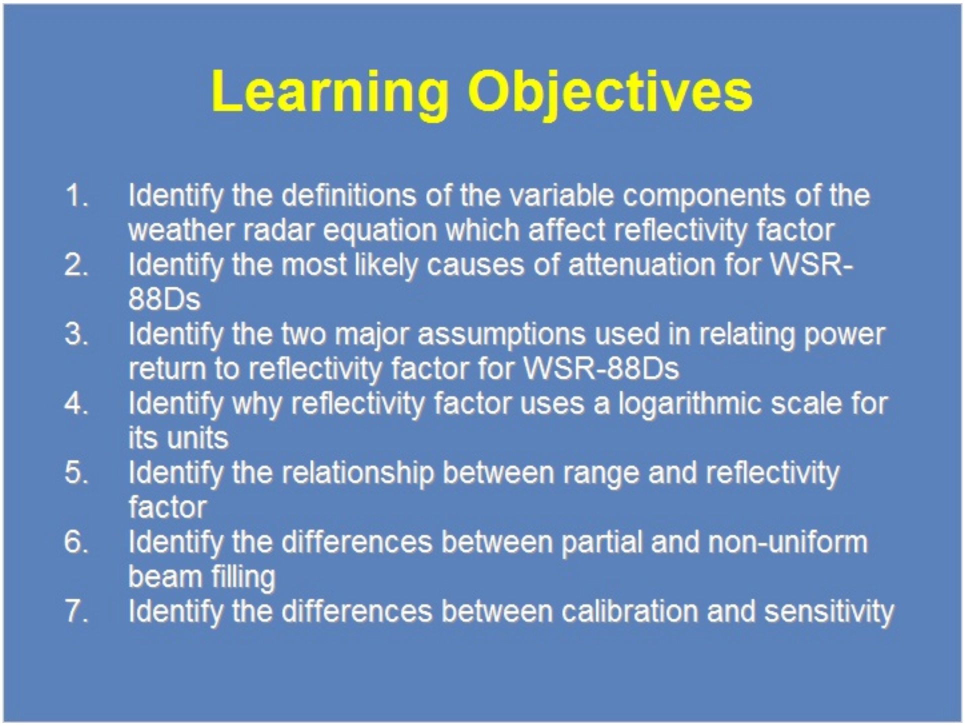

Notes:

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.

Back to Top

2. Main Radar Equation

2.1 Weather Radar Equation

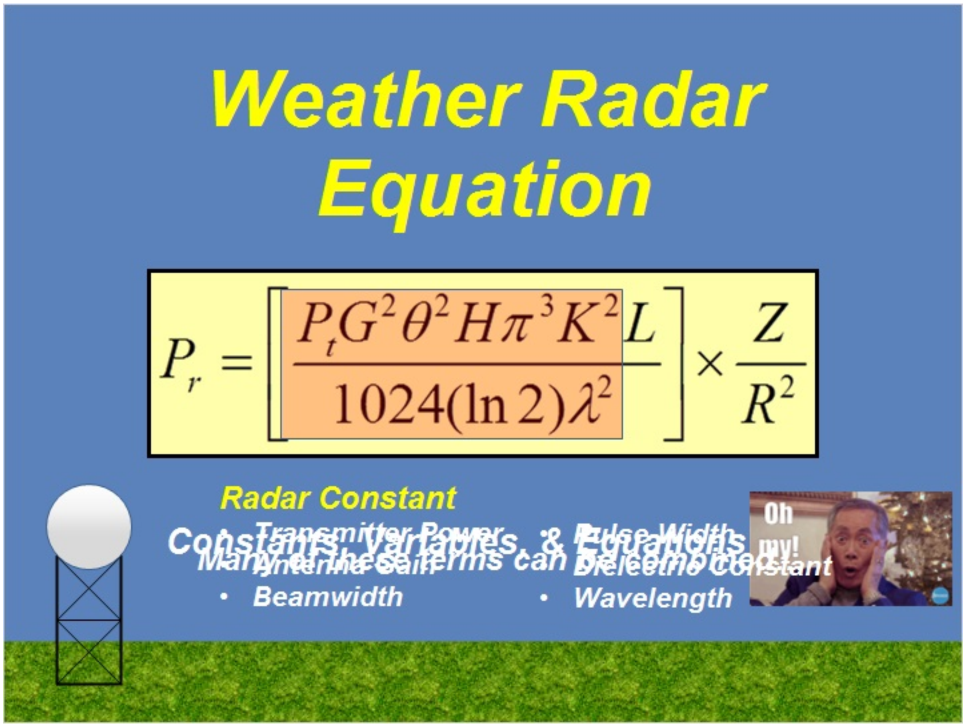

Notes:

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…

Back to Top

2.2 Weather Radar Equation HOME

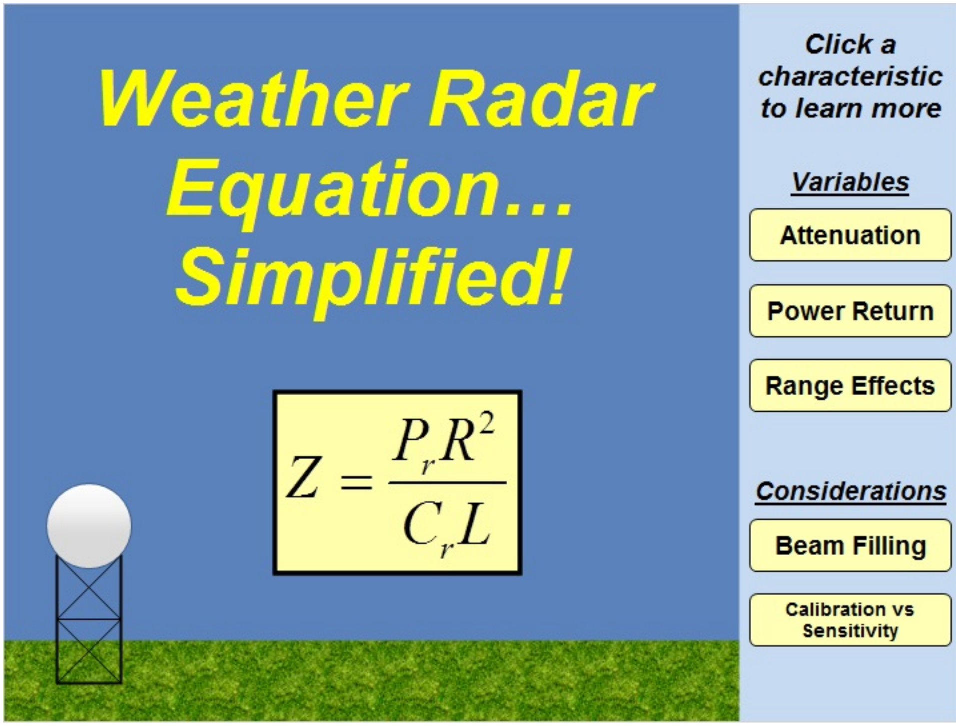

Notes:

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

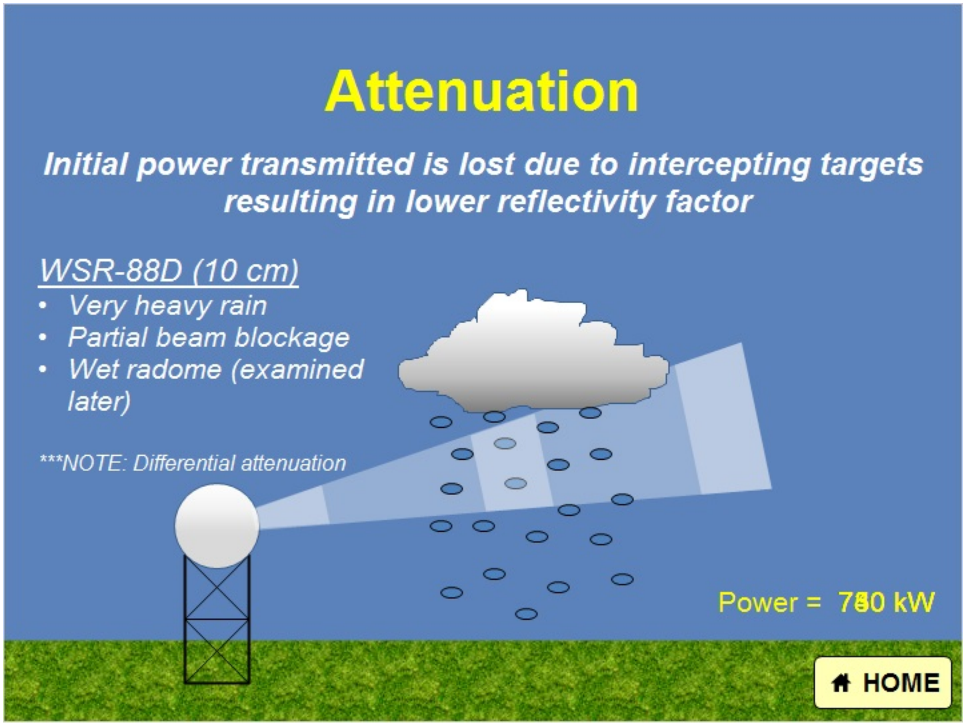

Notes:

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 Reflectivity Factor Units

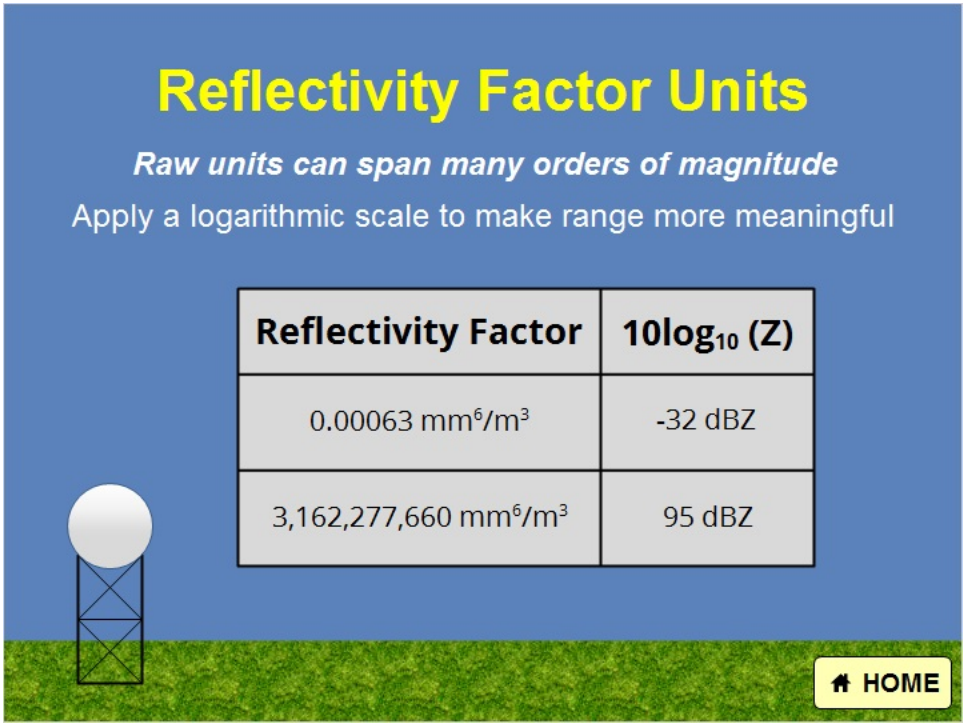

Notes:

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!

2.5 Range Effects

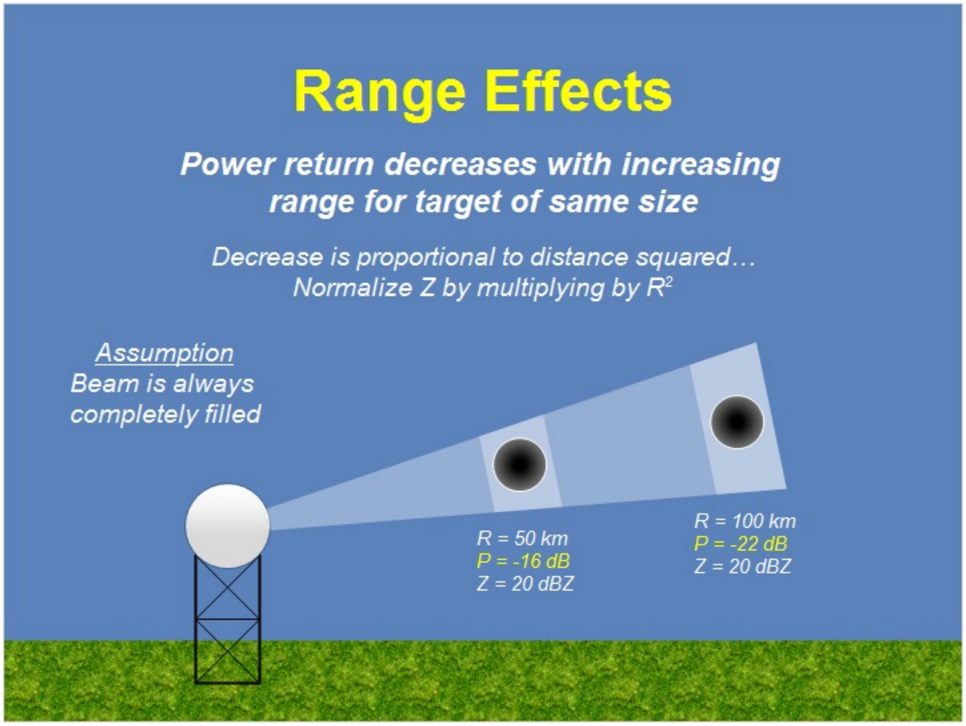

Notes:

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.

Back to Top

2.6 Power Return

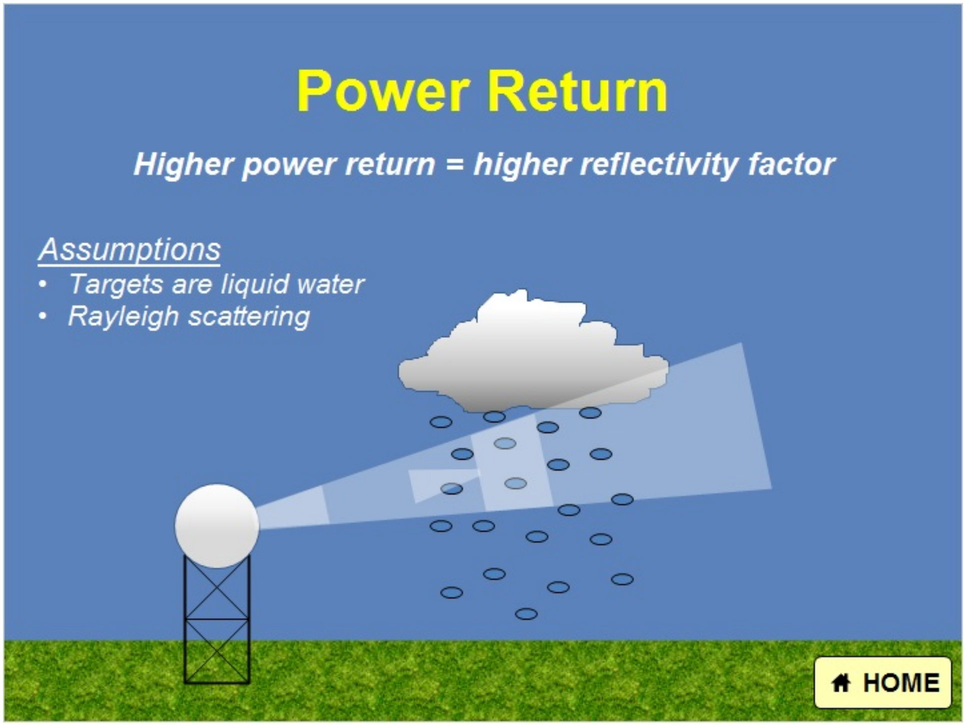

Notes:

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.

Back to Top

2.7 Partial Beam Filling

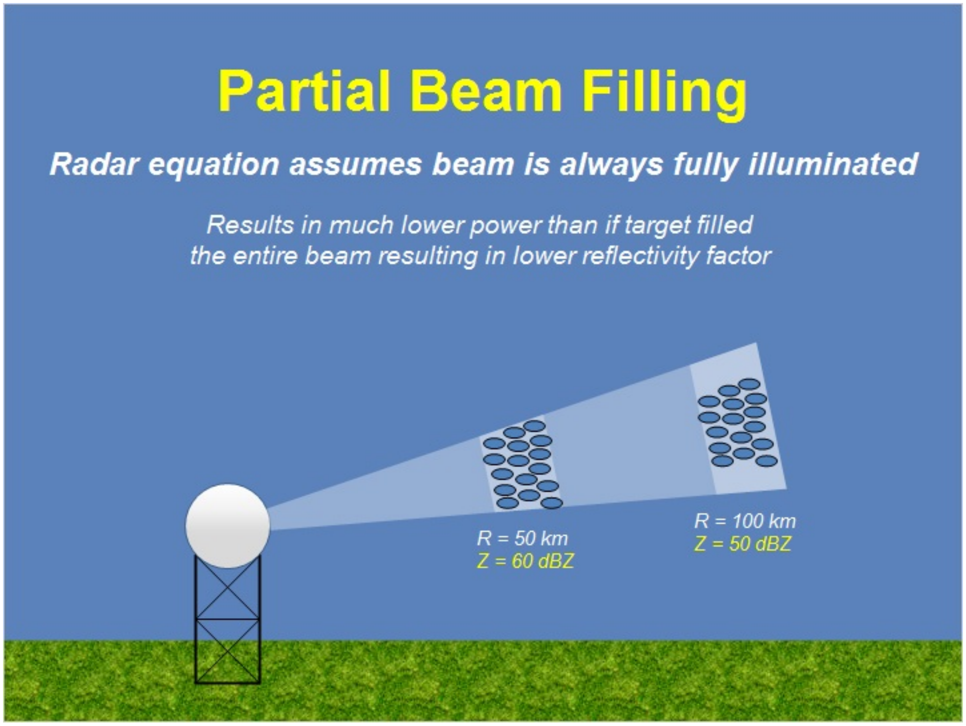

Notes:

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.

Back to Top

2.8 Non-Uniform Beam Filling

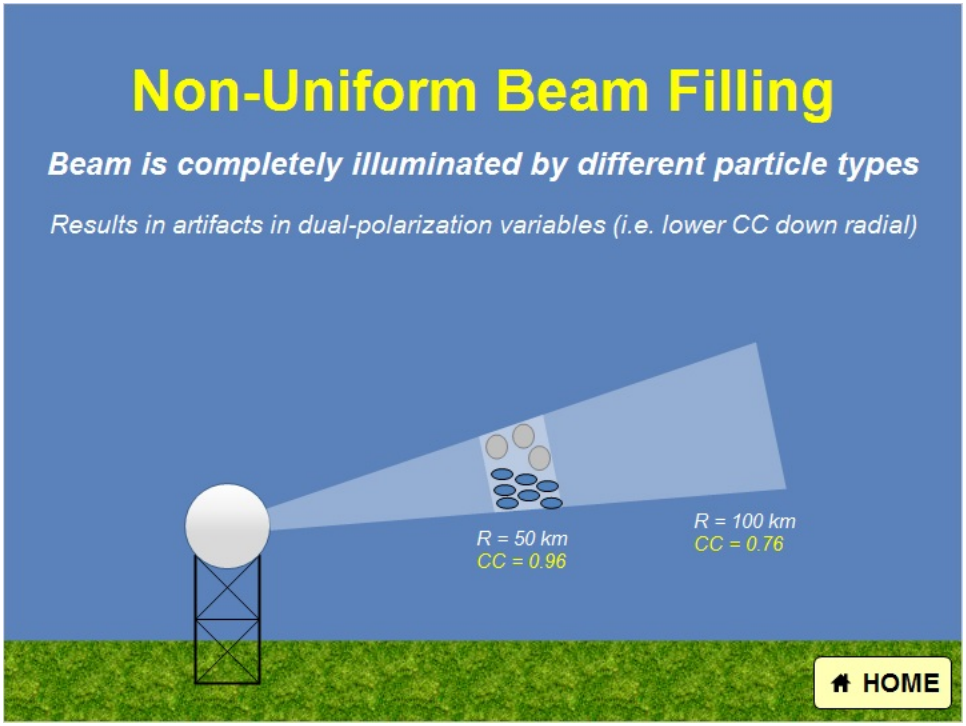

Notes:

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.

Back to Top

2.9 Calibration

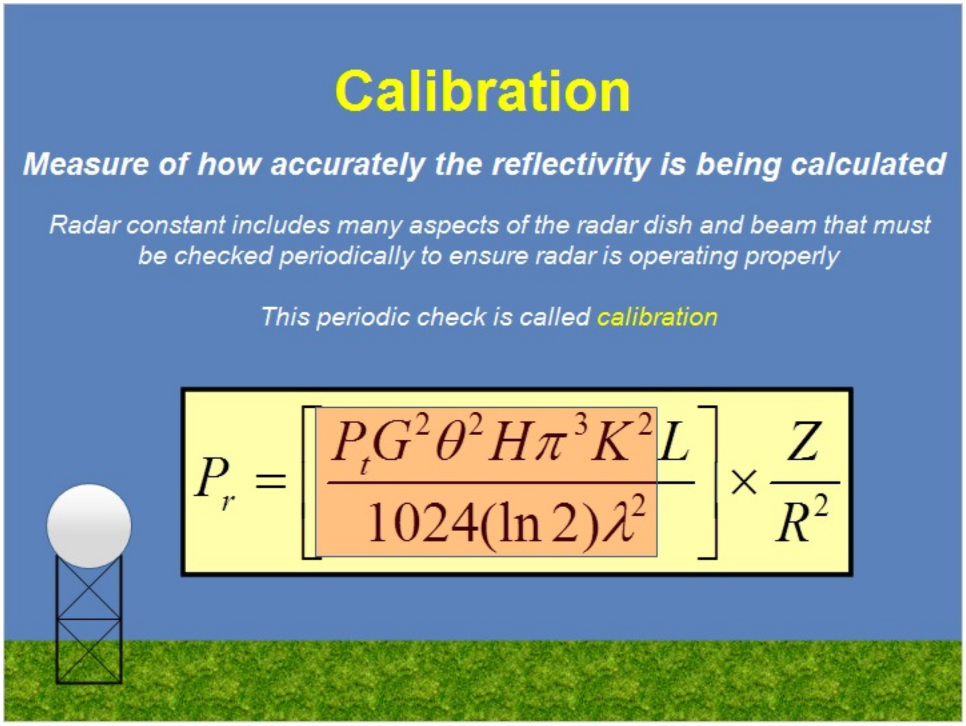

Notes:

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.

Back to Top

2.10 Sensitivity

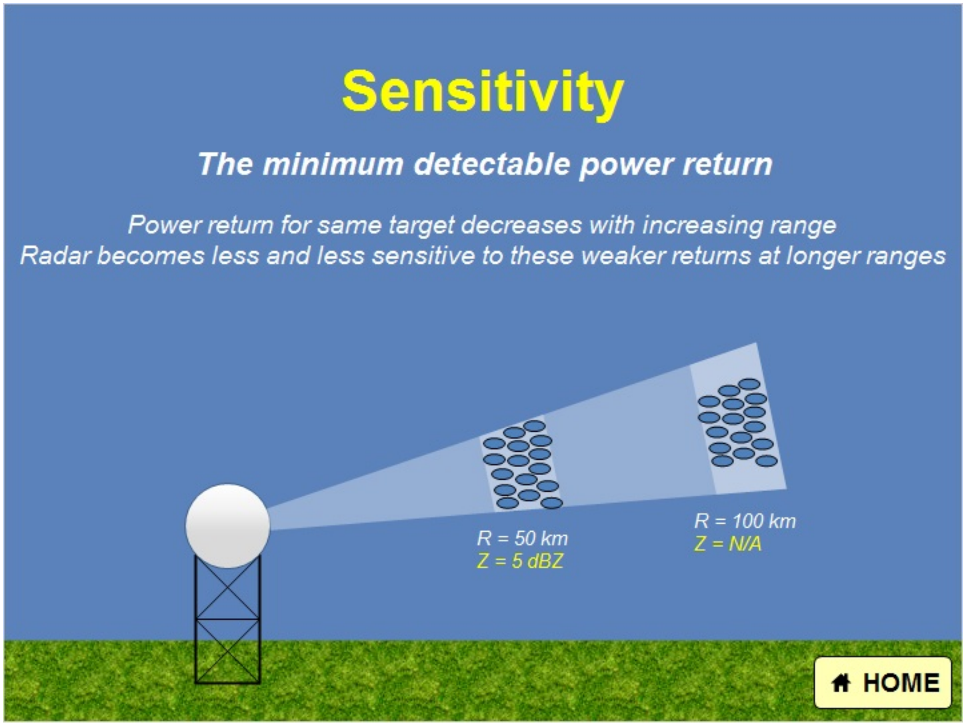

Notes:

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.

Back to Top

4. Completion

4.1 Completion!

Notes:

Thanks for your attention! You are now complete for this lesson.

WSR-88D Fundamentals Part 3: Transmitting & Receiving 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 3: Transmitting and Receiving Characteristics. It is presented by the Warning Decision Training Division (WDTD). Let's get started!

Back to Top

1.3 Roadmap

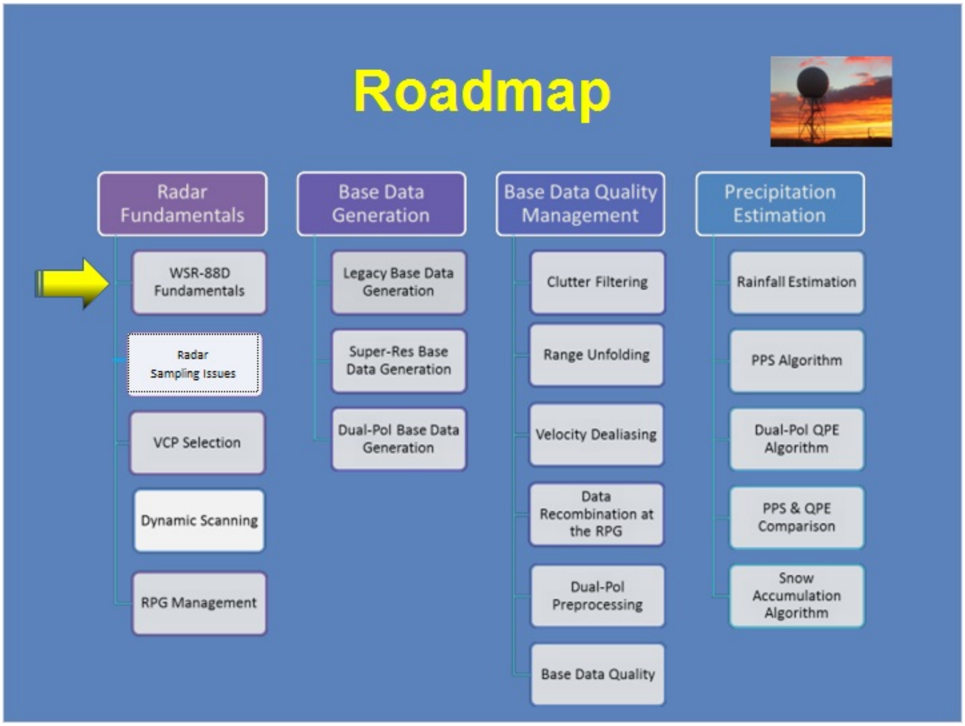

Notes:

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!

Back to Top

1.4 Learning Objectives

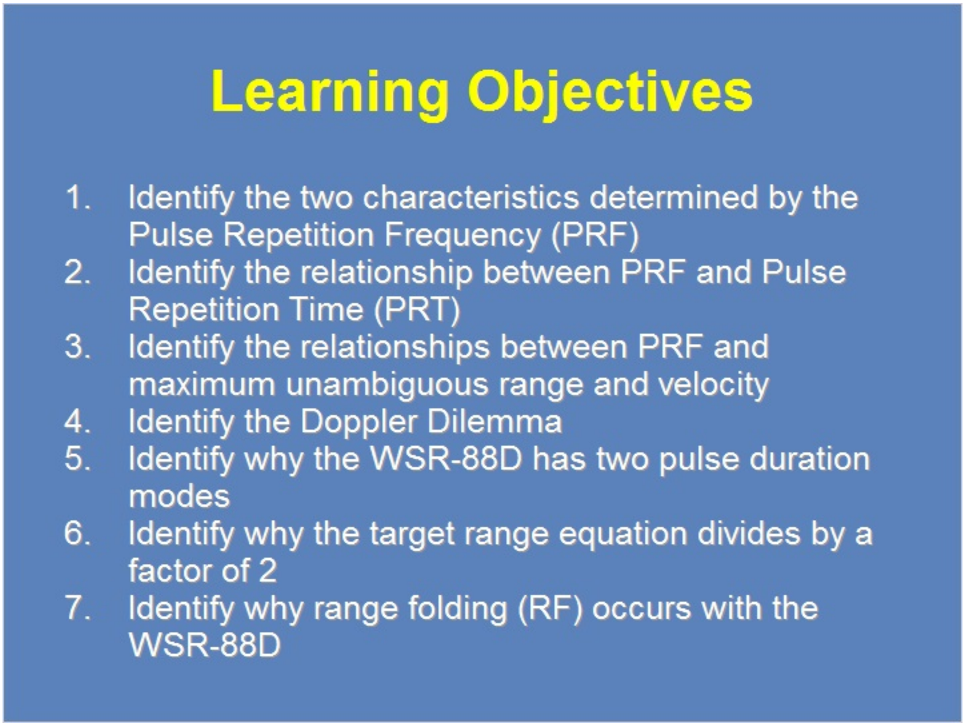

Notes:

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.

Back to Top

2. Transmitting & Receiving Characteristics

2.1 Radar Beam Characteristics HOME

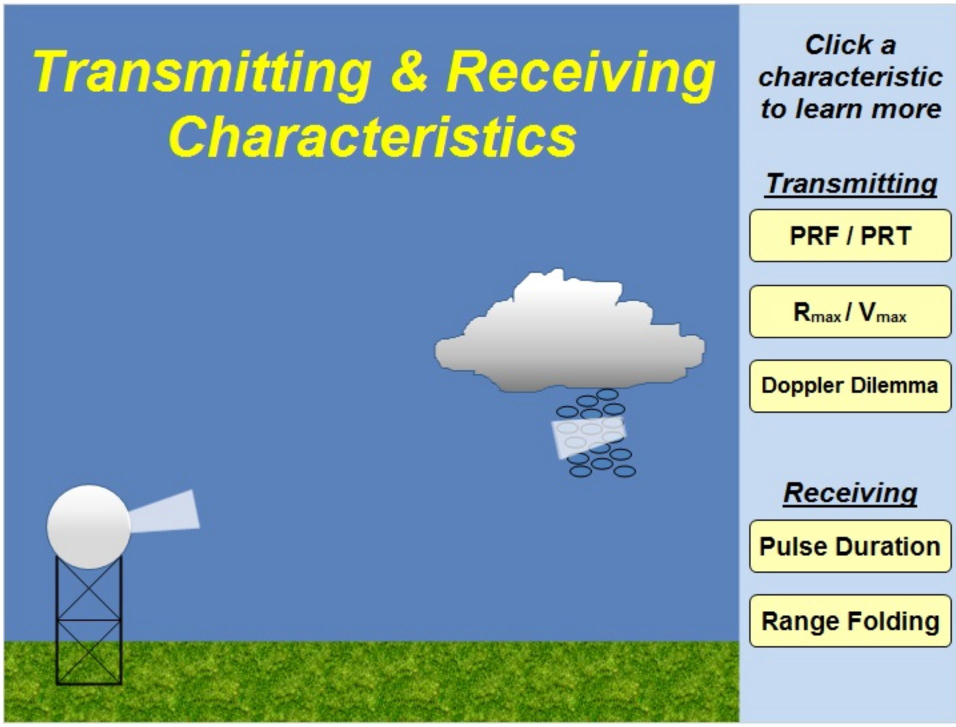

Notes:

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.

Back to Top

2.2 Pulse Repetition Frequency (PRF)

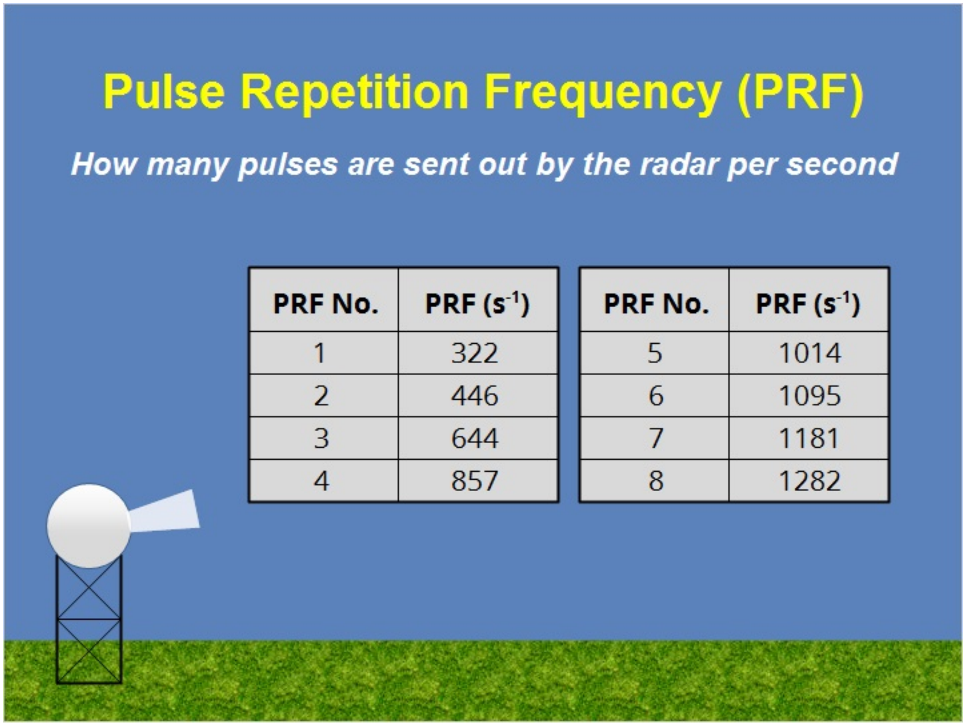

PRF No.	PRF (s^{-1})
1	322
2	446
3	644
4	857

PRF No.	PRF (s^{-1})
5	1014
6	1095
7	1181
8	1282

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!

Back to Top

2.3 Pulse Repetition Time (PRT)

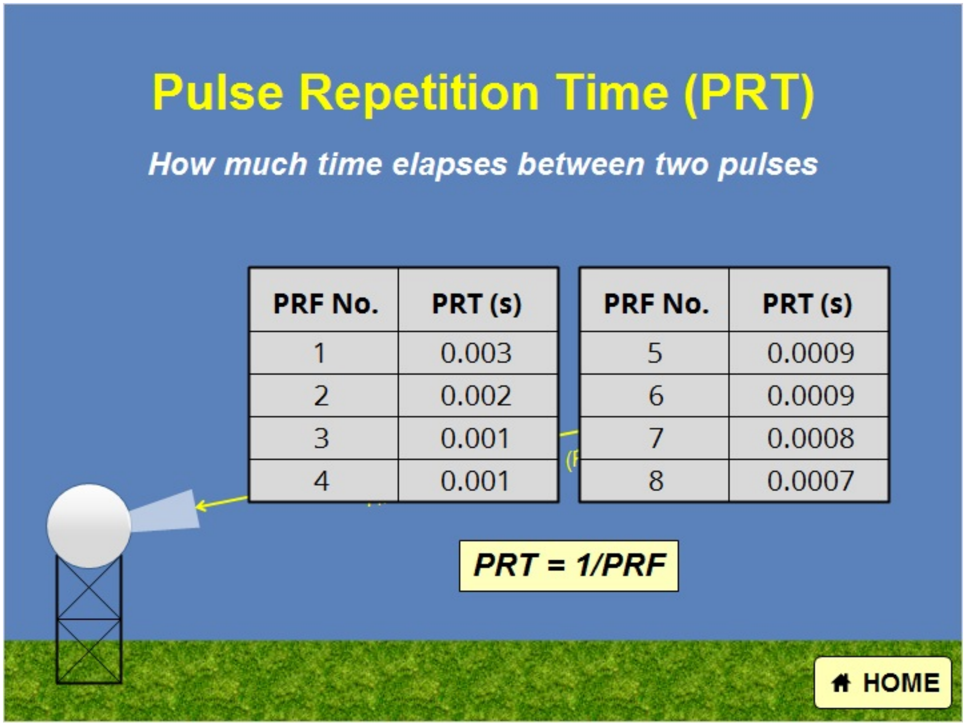

Notes:

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. Notice, there isn't a lot of time in between pulses.

Back to Top

2.4 Maximum Unambiguous Range (Rmax)

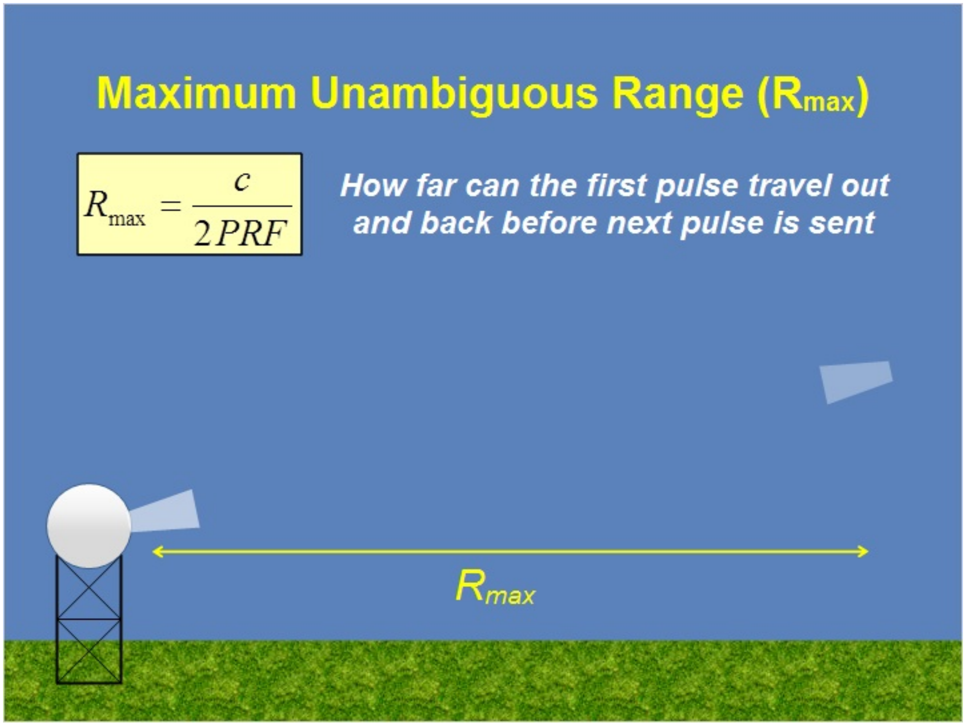

Notes:

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)

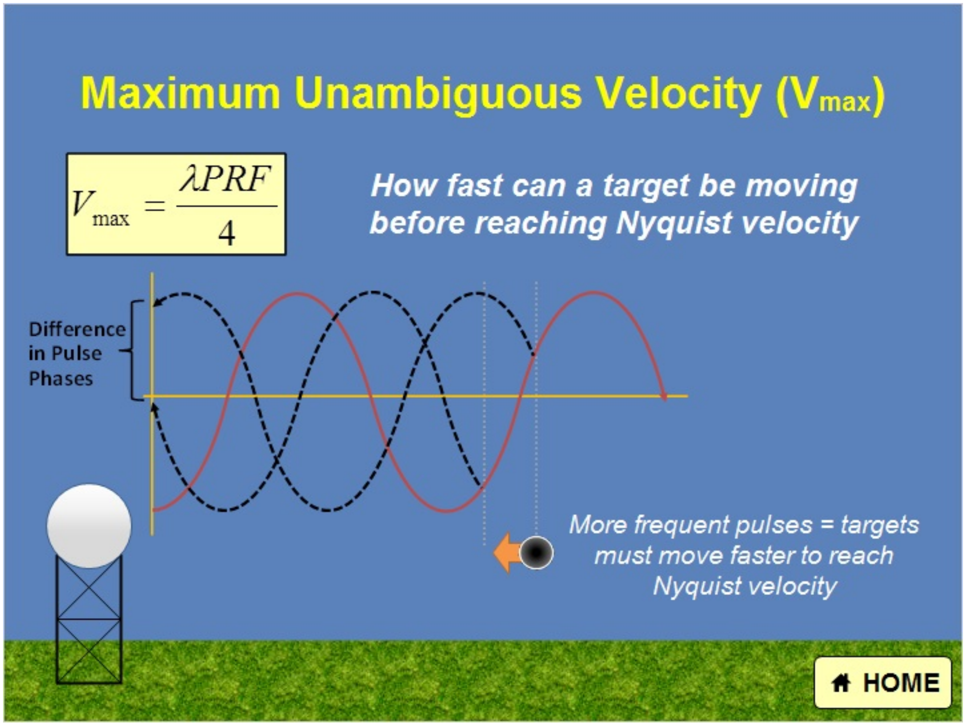

Notes:

When it comes to measuring velocity, the radar measures the phase shift between pulses through a technique called pulse-pair processing. This technique uses the difference in phase of the return signal from one pulse to the next. A target's phase will change from pulse to pulse because the target intercepts the transmitted wave at a different phase position along the wave. However, if the target moves too fast, phase ambiguity can result in an ambiguous determination for velocity. For a given sampling rate, the speed limit where this ambiguity is reached is called the Nyquist Velocity. An easy fix to this is to sample more often, so that scatterers haven't moved far enough between pulses to cause ambiguity. To put it another way, the higher the PRF, the higher the Vmax since targets can move faster before reaching the problematic Nyquist velocity.

Back to Top

2.6 Doppler Dilemma

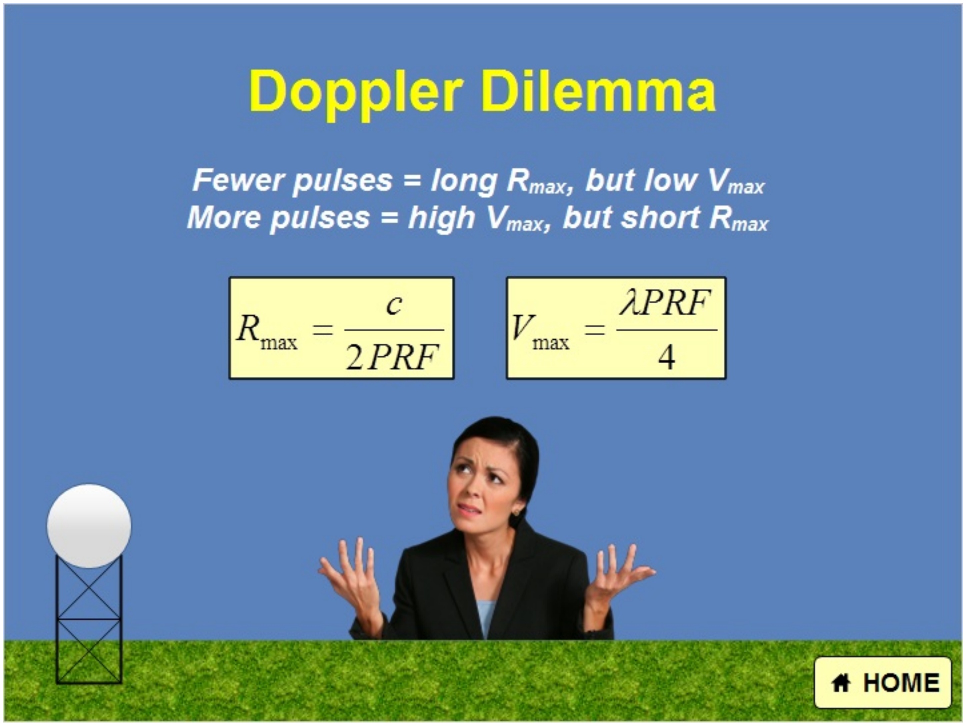

Notes:

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.

Back to Top

2.7 Target Range Equation

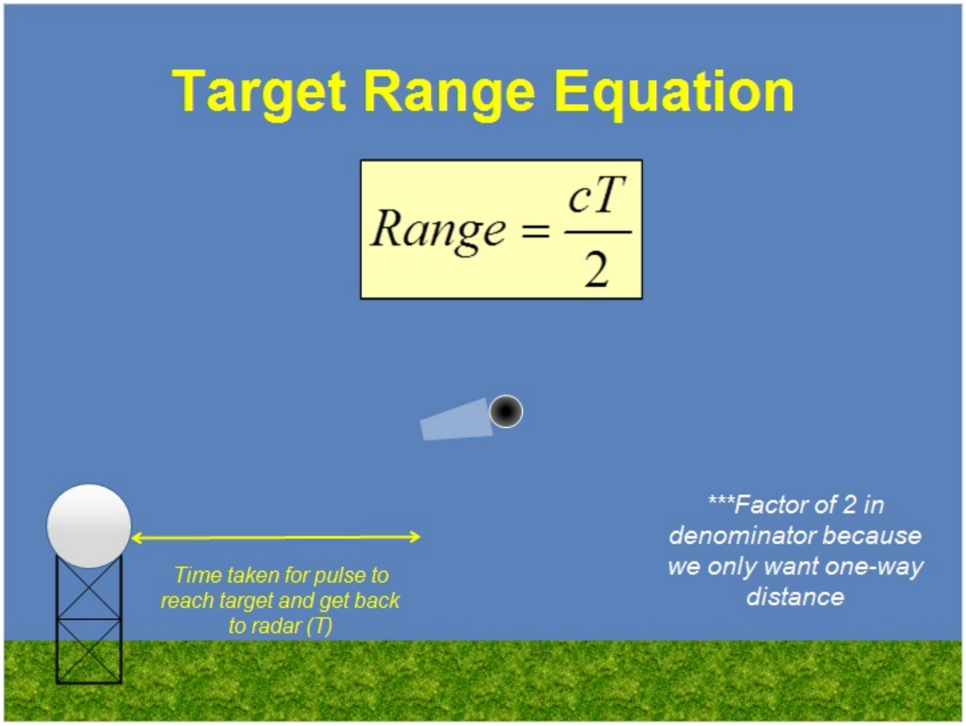

Notes:

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.

Back to Top

2.8 Pulse Duration

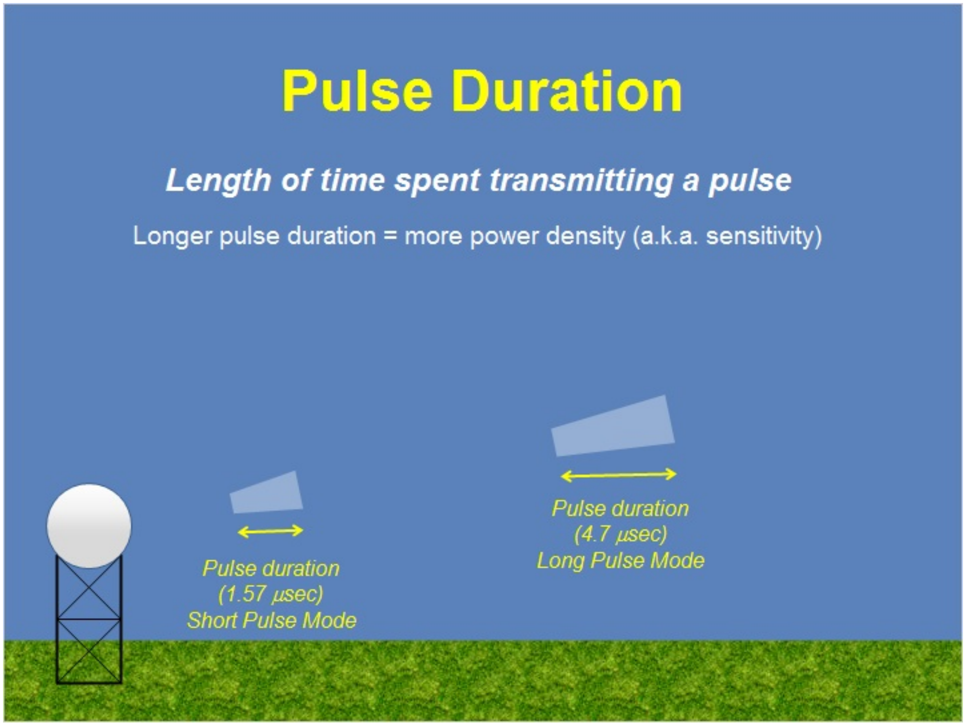

Notes:

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…

Back to Top

2.9 Range Folding

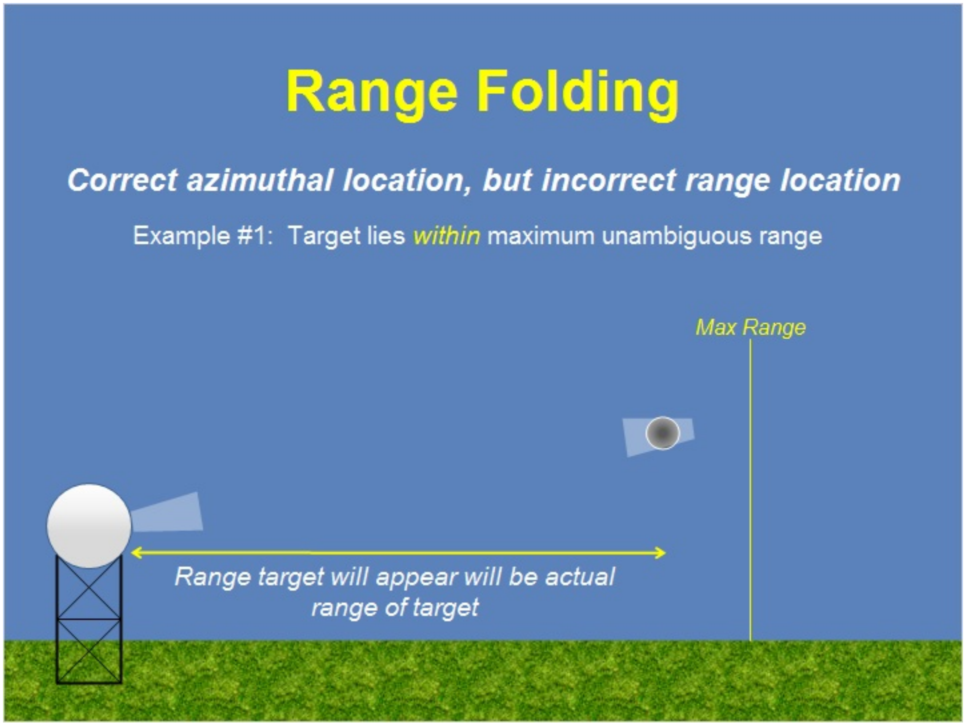

Notes:

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 radar 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.

Back to Top

2.10 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.

2.11 Range Folding

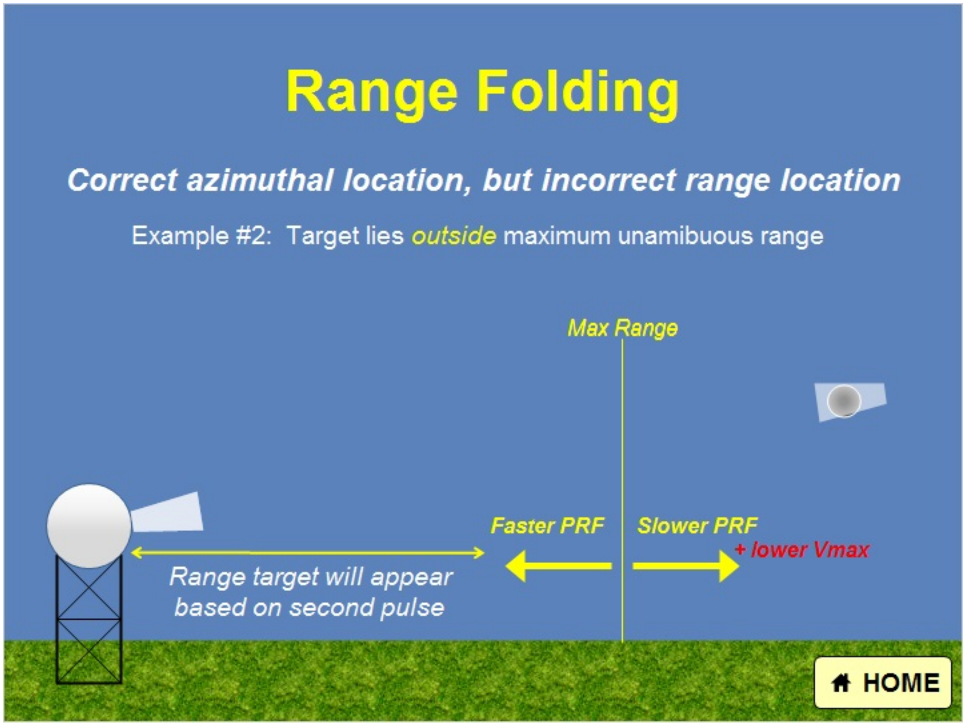

Notes:

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.

But remember from earlier how our selection of the pulse repetition frequency, or PRF, affected max range? A slower PRF results in a farther max range, since it gives echoes more time to return before a new one is sent. So by adjusting your PRF (such as with a different VCP) you can sometimes reduce problems with range folding. But remember that this comes at the cost of a more restricted maximum ambiguous velocity, thanks to the Doppler Dilemma.

Back to Top

4. Completion

4.1 Completion!

Notes:

Thanks for your attention! You are now complete for this lesson.

Back to Top

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!

1.3 Roadmap

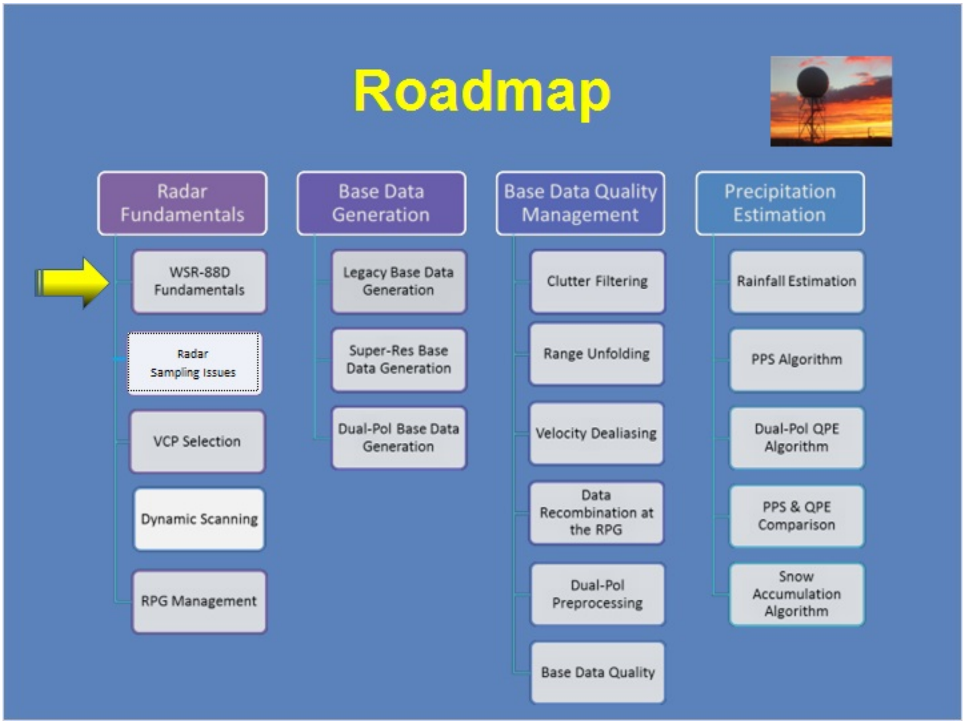

Notes:

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!

1.4 Learning Objectives

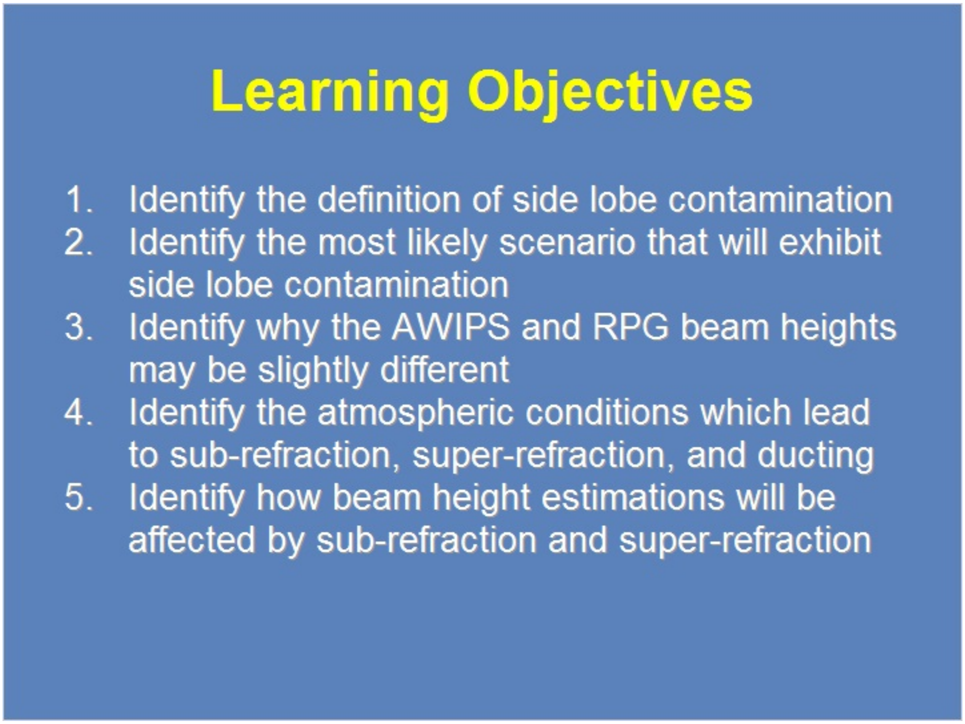

Notes:

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.

Back to Top

2. Transmitting & Receiving Characteristics

2.1 Non-Standard HOME

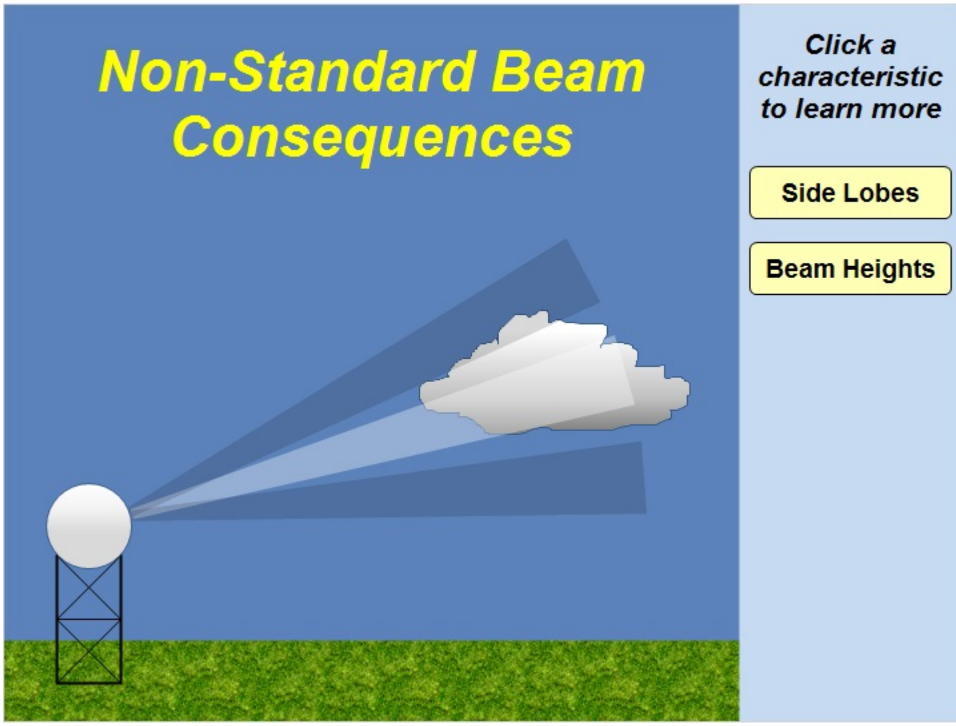

Notes:

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.

Back to Top

2.2 Side Lobe Definition

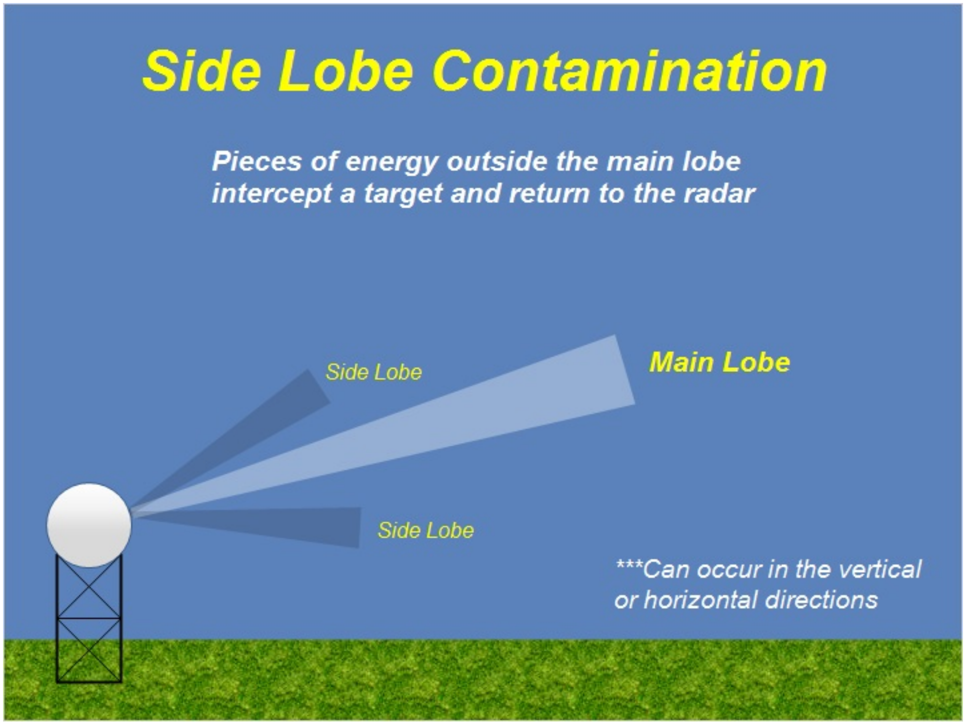

Notes:

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.

Back to Top

2.3 Side Lobe Horizontal

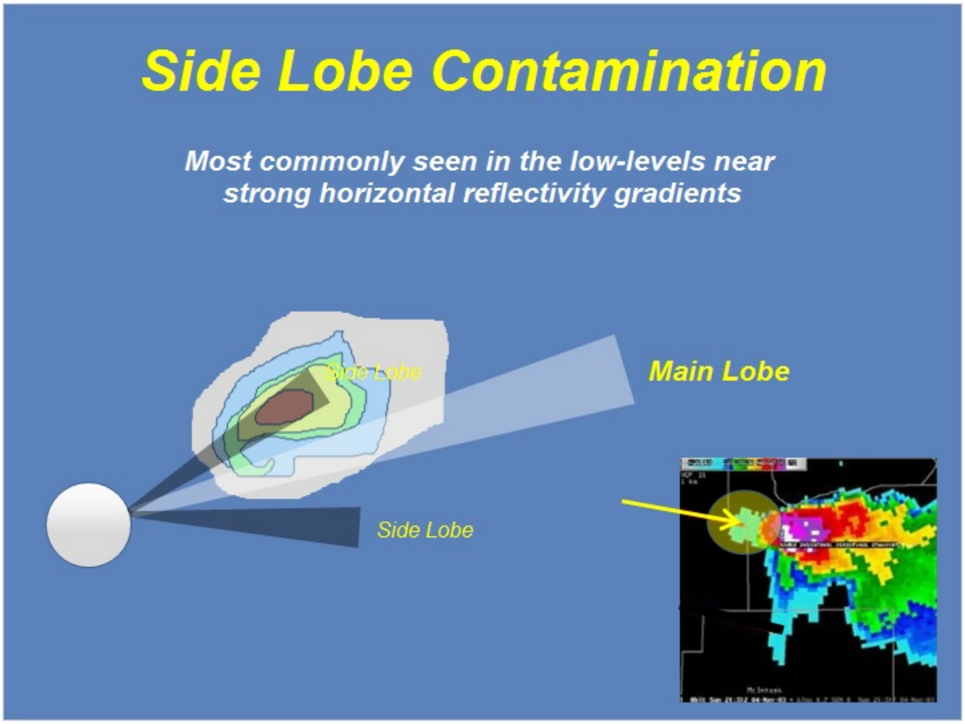

Notes:

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

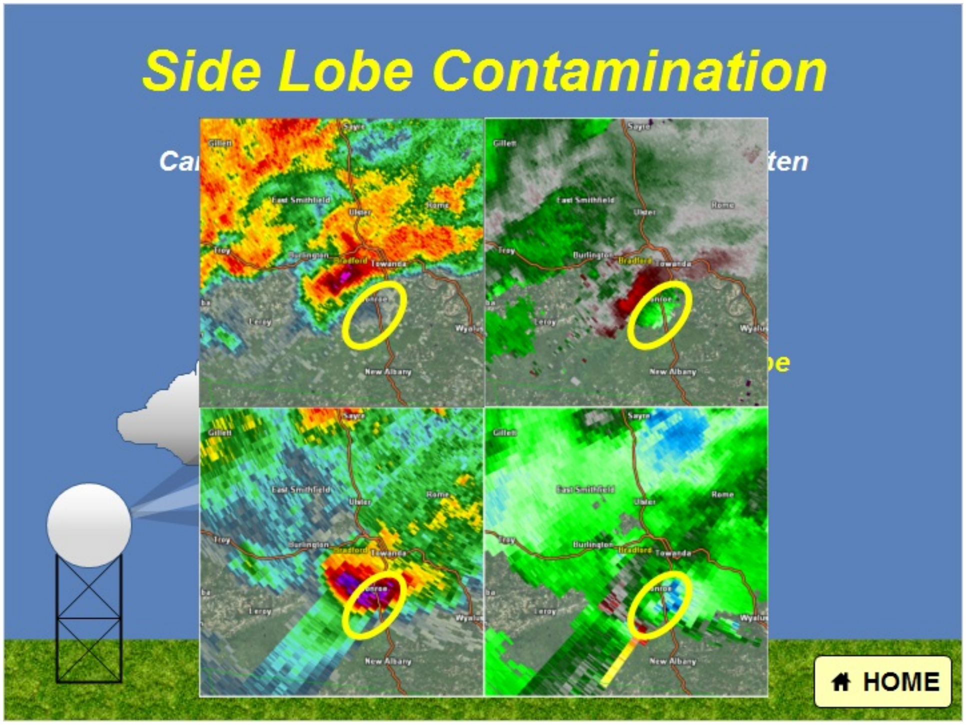

Notes:

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. 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.

Back to Top

2.5 Beam Height Calcs

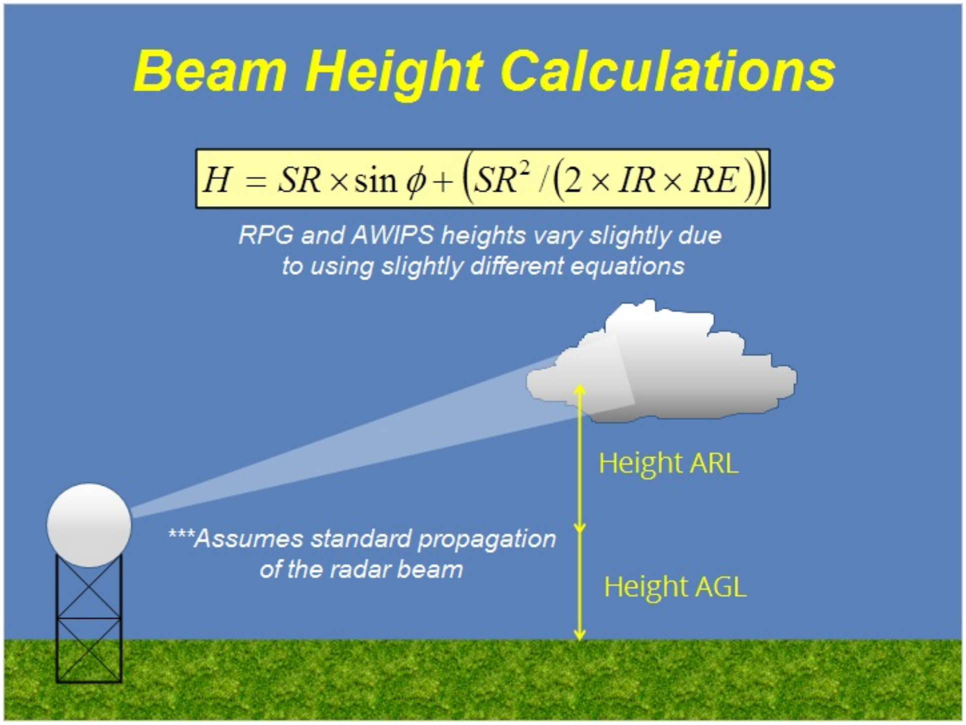

Notes:

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.

Back to Top

2.6 Super-Refraction

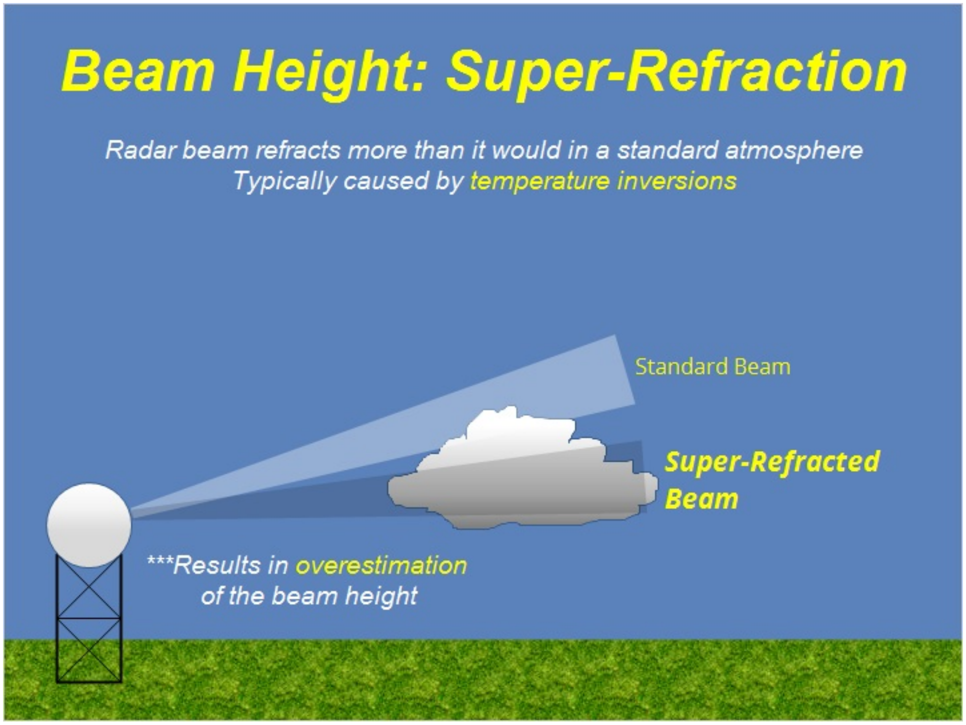

Notes:

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.

2.7 Sub-Refraction

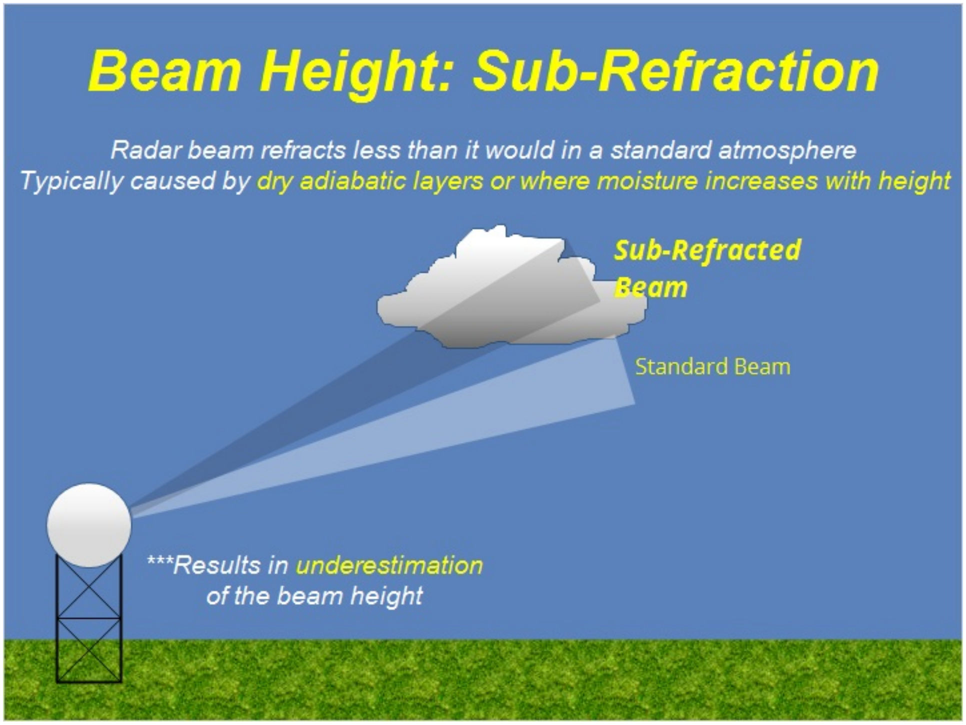

Notes:

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.

2.8 Ducting

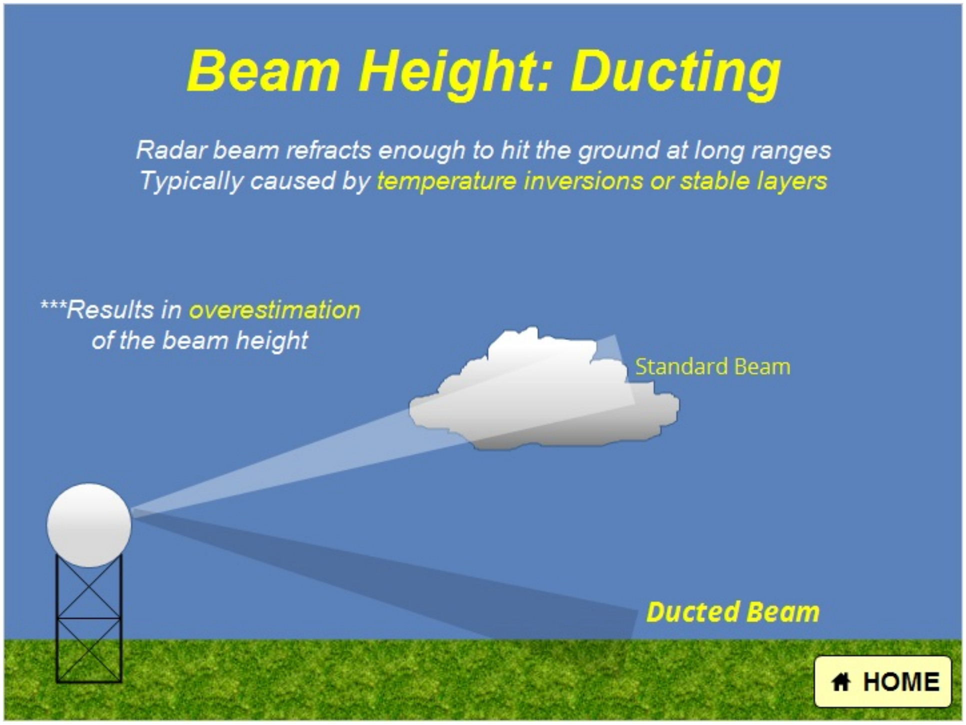

Notes:

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.

4. Completion

4.1 Completion!

Notes:

Thanks for your attention! You are now complete for this lesson.

Back to Top

WSR-88D Fundamentals Part 5: Data Collection

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 5: Data Collection. It is presented by the Warning Decision Training Division (WDTD). Let's get started!

1.3 Roadmap

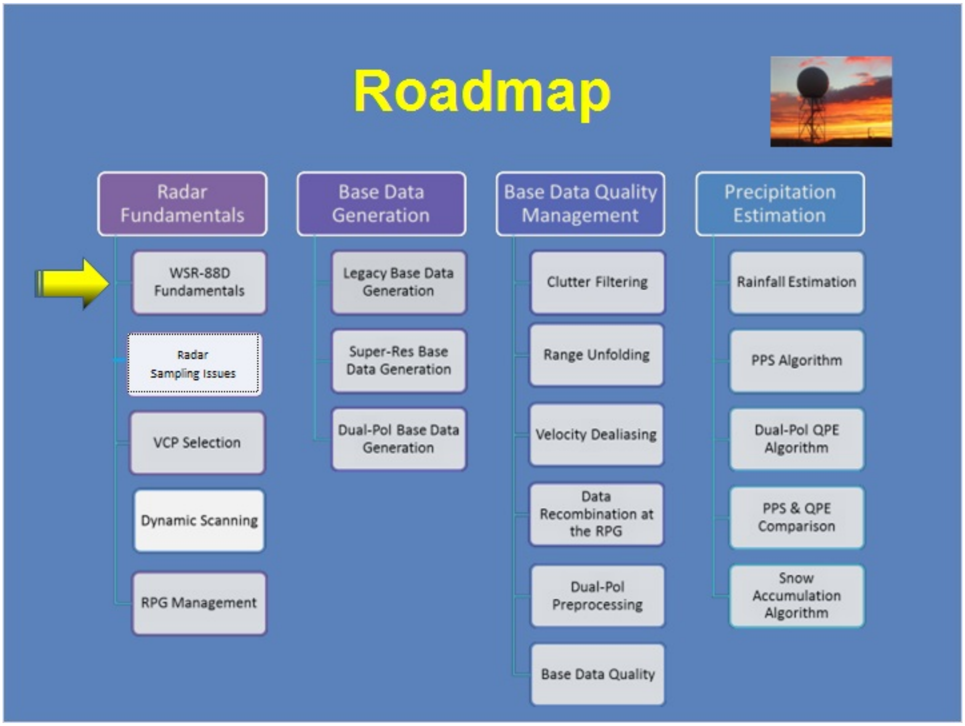

Notes:

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!

1.4 Learning Objectives

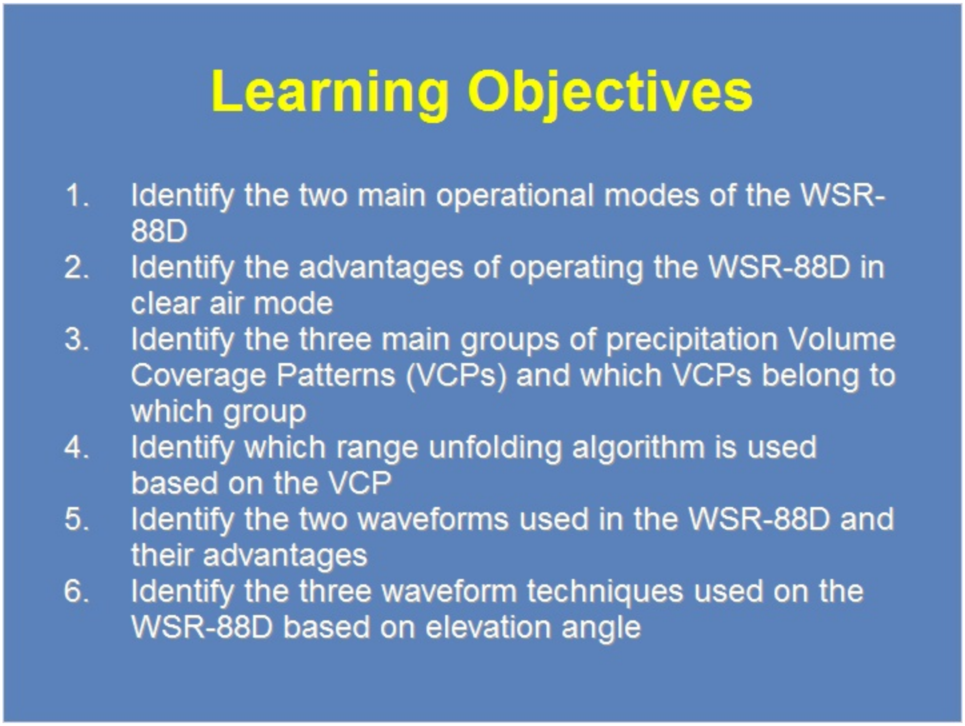

Notes:

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.

Back to Top

2. Transmitting & Receiving Characteristics

2.1 Data Collection HOME

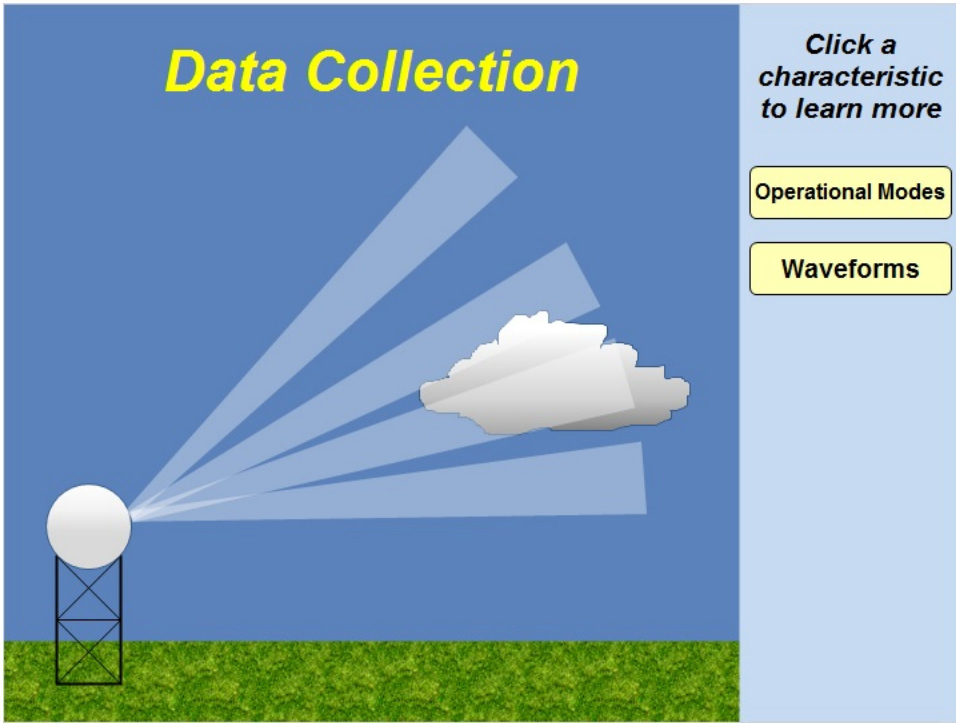

Notes:

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.

Back to Top

2.2 Operational Modes

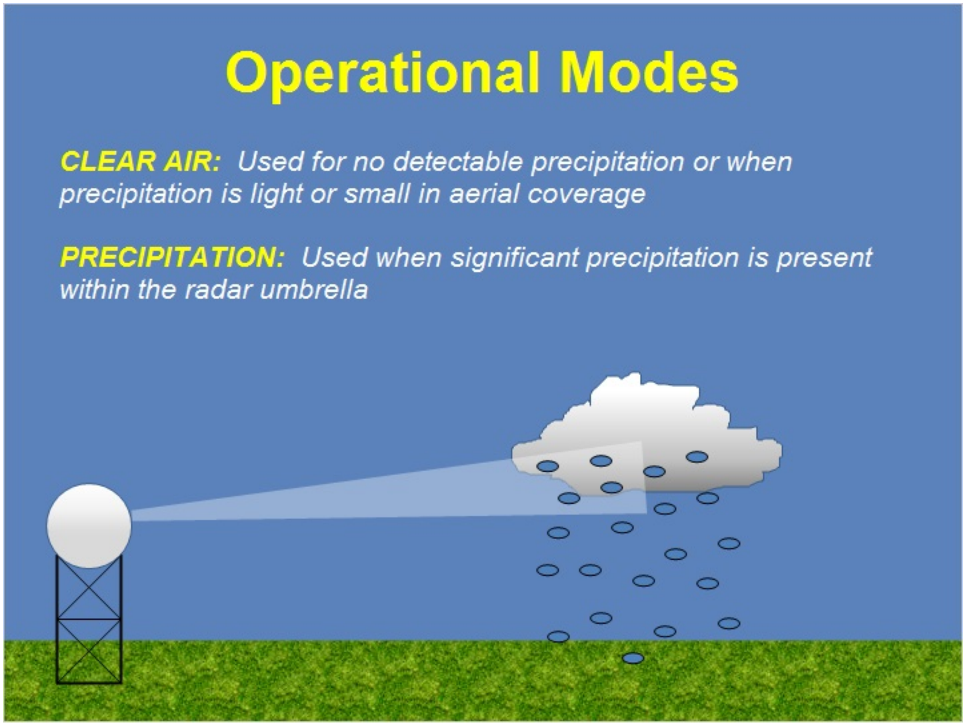

Notes:

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

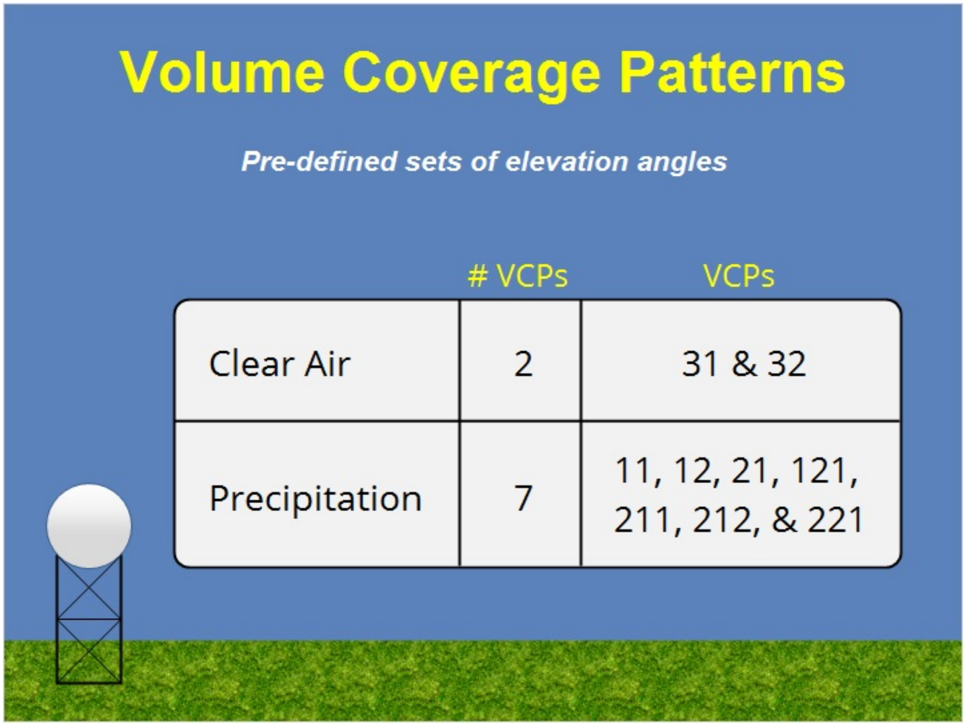

Notes:

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…

Back to Top

2.4 Clear Air VCPs

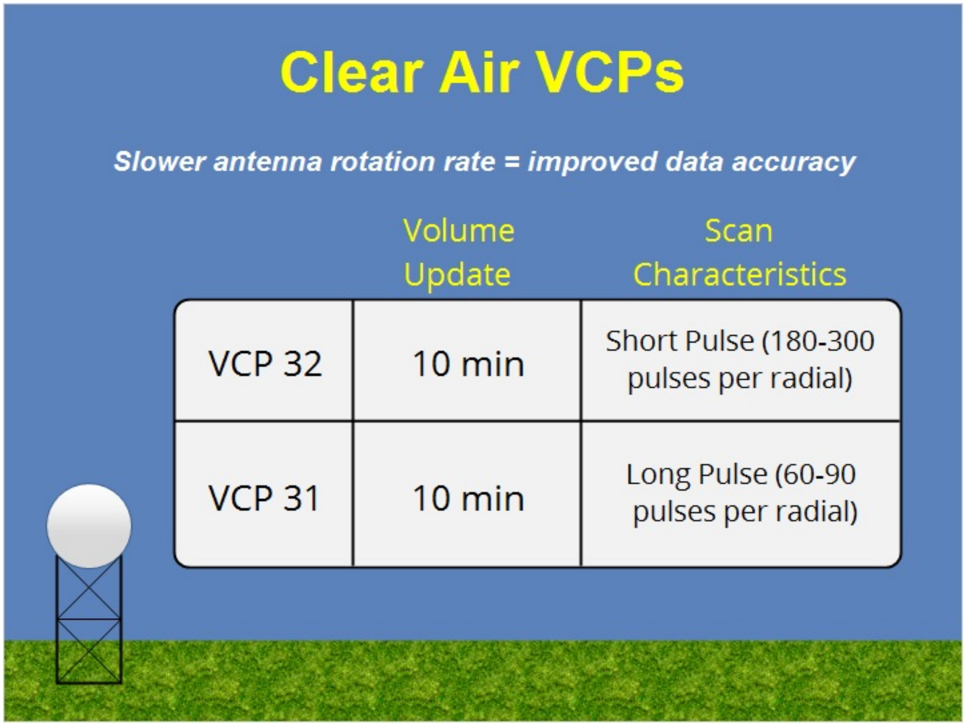

Notes:

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…

Back to Top

2.5 Precip VCPs

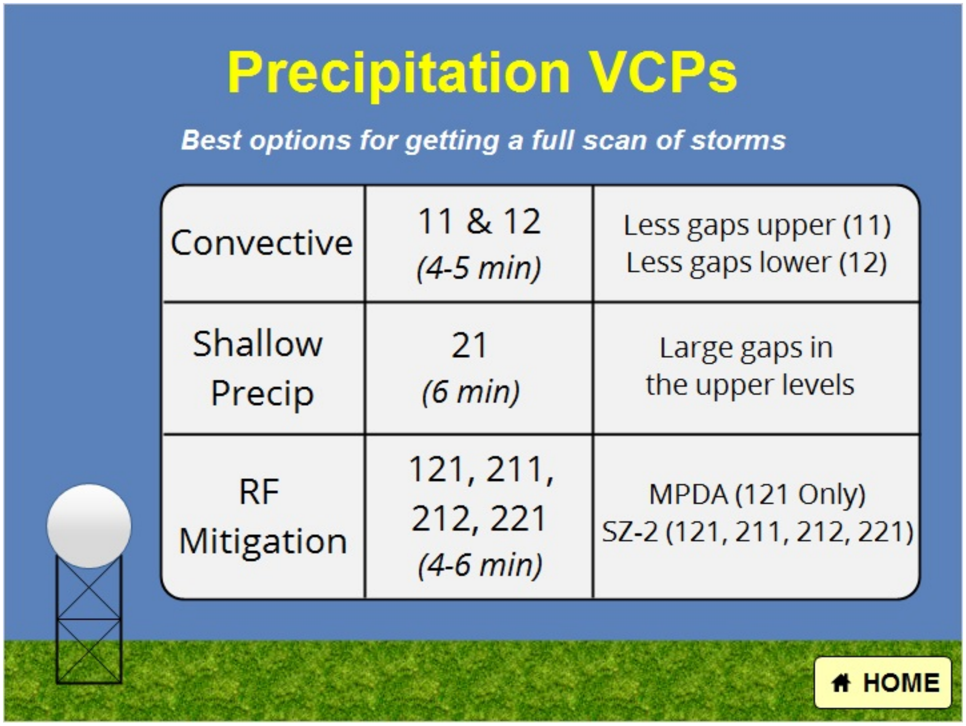

Notes:

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.

Back to Top

2.6 Waveforms

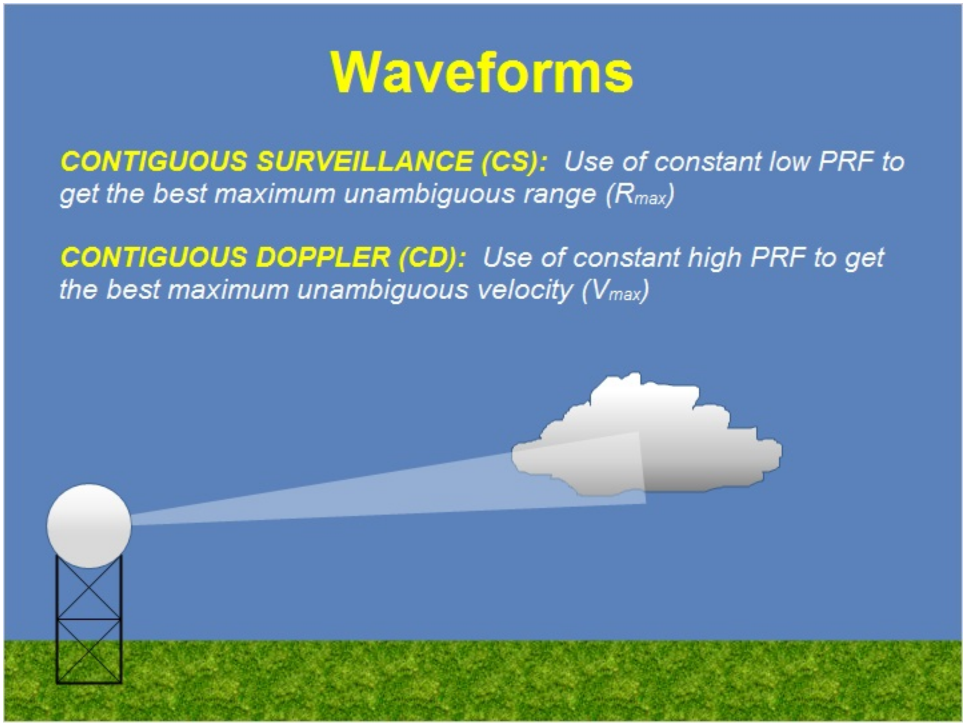

Notes:

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

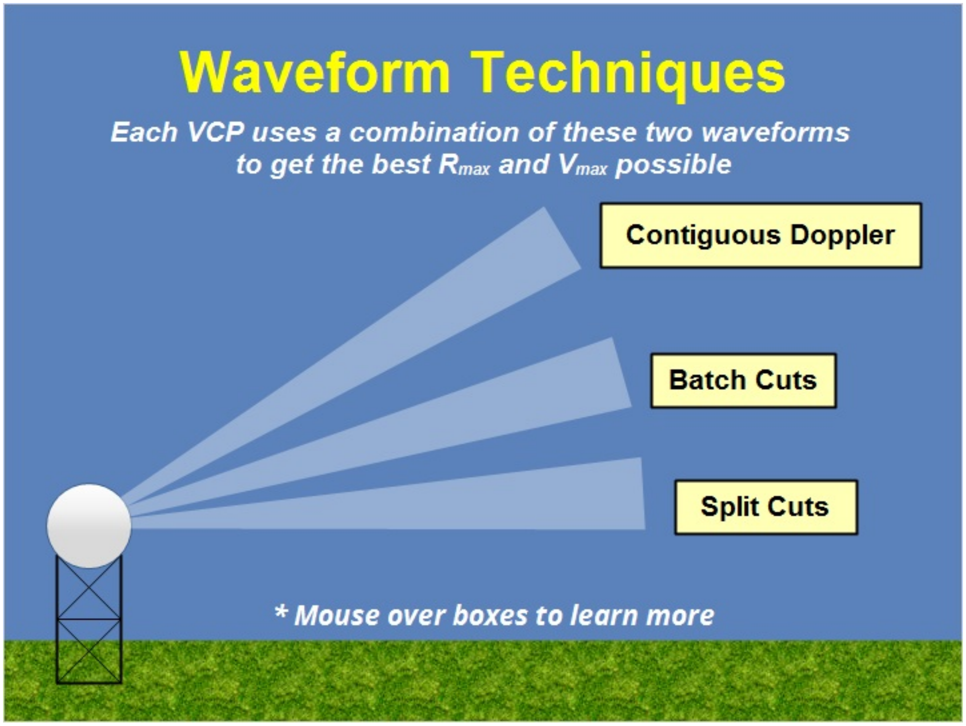

Notes:

There are three main grouping of elevation angles with the WSR-88D. Mouse over each of the descriptions on the right to see which elevations are included and what the scanning strategy is for each of those elevations. 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. 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. Finally, all elevations above 7 degrees perform one scan using the CD mode only. Now it's time for your quiz, so go back home and click the Quiz button to test your knowledge!

4. Completion

4.1 Completion!

Notes:

Thanks for your attention! You are now complete for this lesson.

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 covers common radar sampling issues with the WSR-88D. Let's get started!

Back to Top

1.3 Roadmap

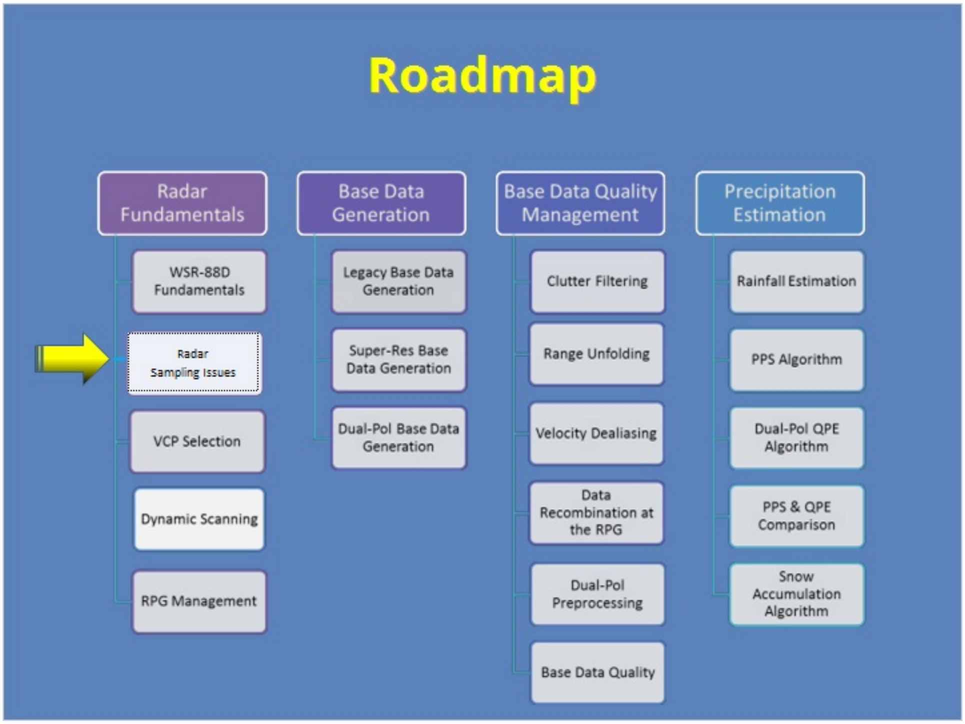

Notes:

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 topic. Let’s keep going!

Back to Top

1.4 Learning Objectives

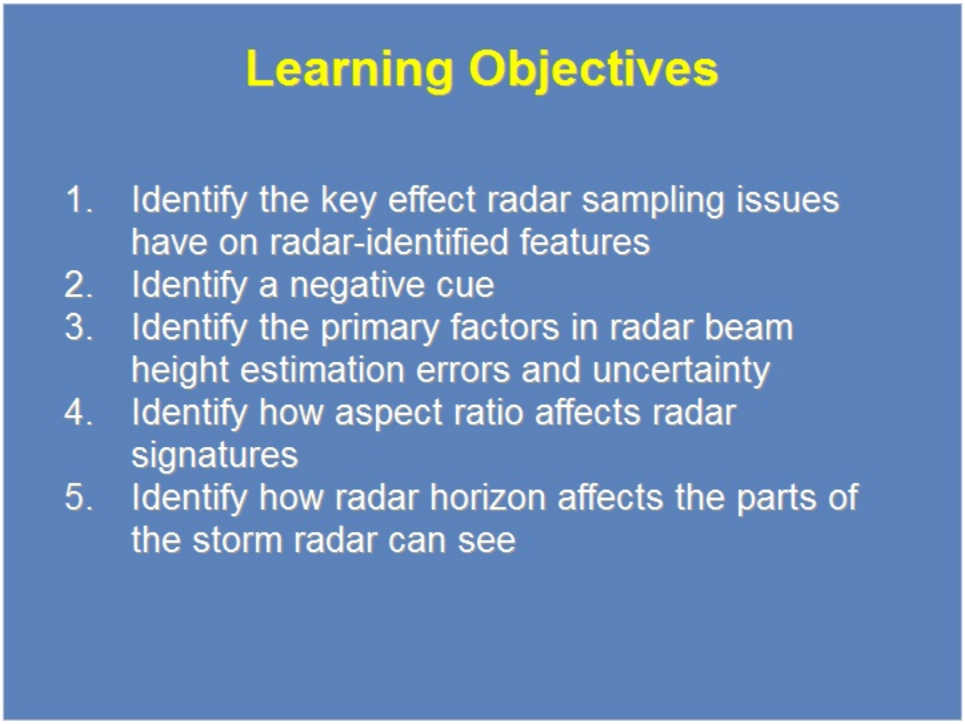

Notes:

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

Back to Top

1.5 Learning Objectives

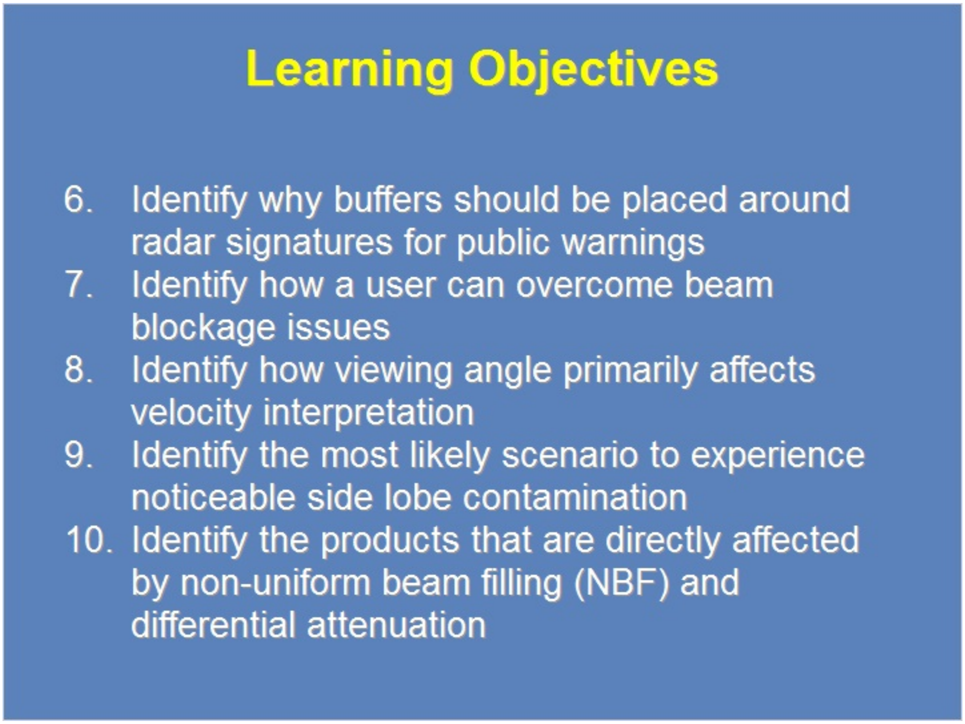

Notes:

Here are the last five learning objectives. Feel free to go back and forth between this slide and the previous one as needed.

1.6 Motivation for This Lesson

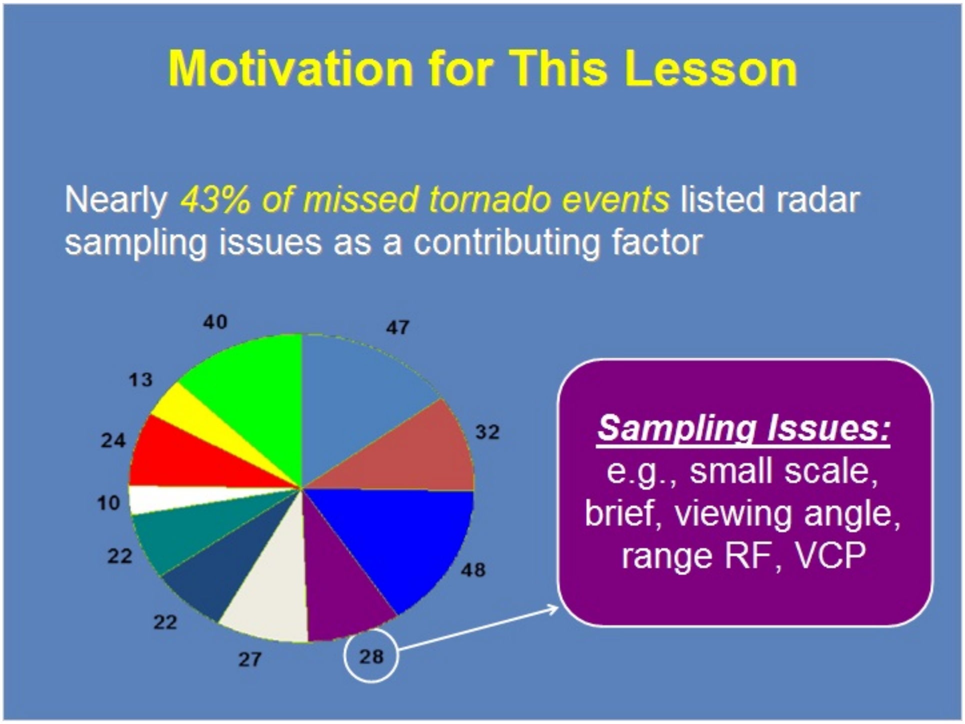

Notes:

A root cause analysis study revealed that in 65 missed tornado events, 28 of those events (or nearly 43%) listed radar sampling issues as a contributing factor. As a result of this research, understanding common radar sampling issues helps forecasters mitigate these issues and avoid future missed events. Let's take a look at these common issues.

2. Transmitting & Receiving Characteristics

2.1 Sampling Issues HOME

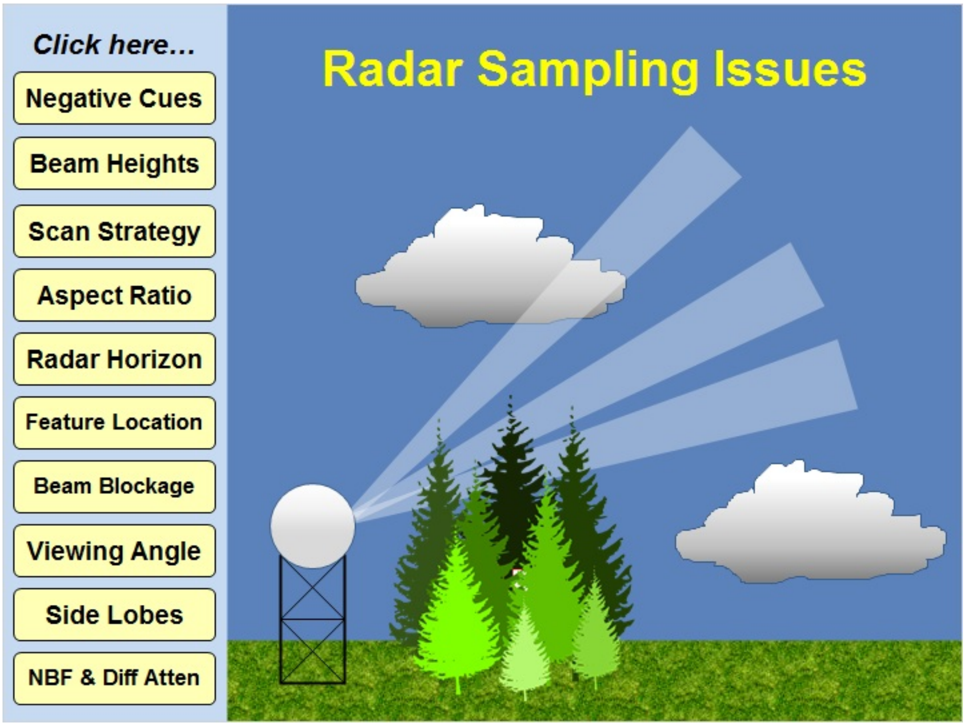

Notes:

Radar sampling issues come in all shapes and sizes. Some can make interpretation almost impossible, while others are just annoyances. However, when we understand how these situations can occur, then we can take steps to mitigate their impacts (if possible). This lesson introduces 10 common issues that are shown on the right. Click on the button for each issue to learn more. When you are done with each section, you will be directed back to this page. After completing all 10 sections, a button shall appear to take you to the quiz.

Back to Top

2.2 Negative Cues

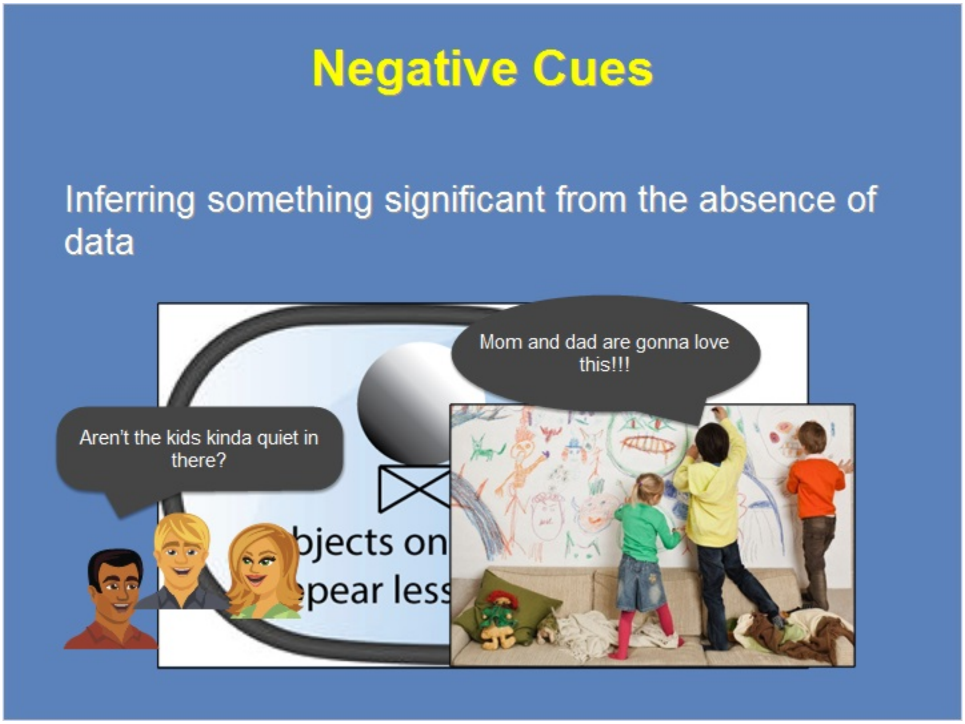

Notes:

Before we tackle actual sampling limitations, we need to discuss the concept of negative cues. Many of the sampling limitations we'll discuss in this lesson impact forecasters by masking radar signatures and making them appear LESS IMPRESSIVE than they are in actuality. In some situations, the sampling issue masks the feature altogether. In these situations, we must infer that something significant is present by the absence of data. This task requires the observation of negative cues.

To better understand concept of negative cues, let's look at an everyday situation. Let's say a group of families gather at a house. The parents talk in one room while the kids play in another room. After some time, the parents notice the kids' room is quiet. I mean too quiet. The parents notice there is NO noise coming from the kids' room which gets their attention. The absence of noise when noise is expected provides a negative cue to the parents that something significant could be occurring in the kids' room.

Back to Top

2.3 Negative Cues in Meteorology

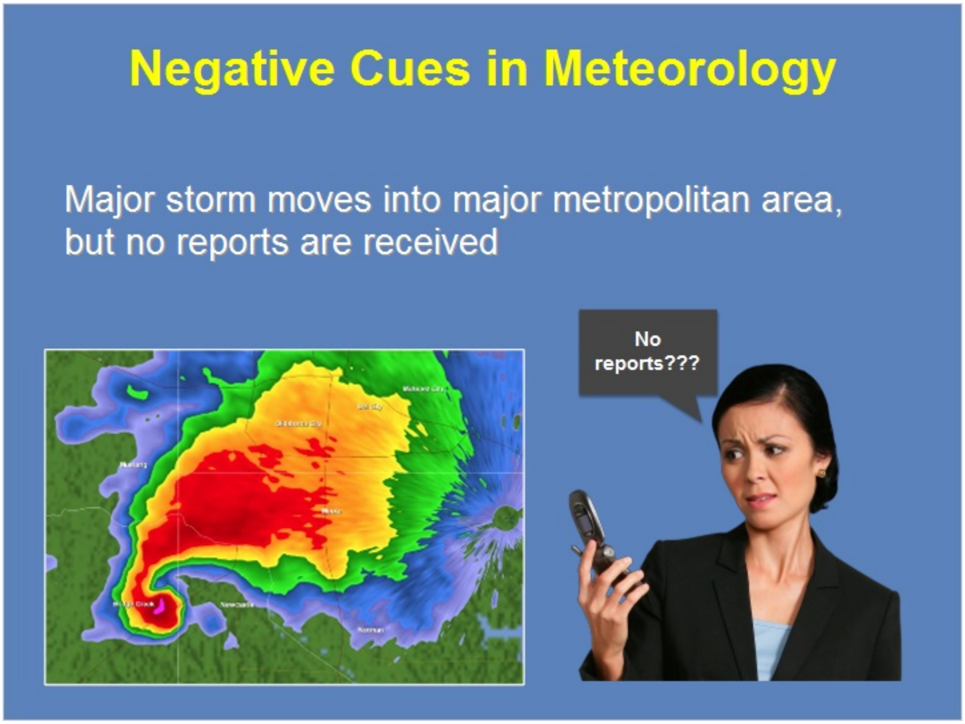

Notes:

Let's take a look at some examples of negative cues in the warning world. Our first example shows a classic supercell with a hook echo moving into a major metropolitan area. With such a large population being affected by this storm, you would certainly expect reports coming into the office, right? What if your office received no reports at all? Would that seem suspicious? The lack of reports from what appears to be a very severe storm should signal that something isn't right. The lack of calls may be due to the storm's damage, either because people aren't thinking of calling the office. Maybe you've lost your phone lines 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 could be happening.

Back to Top

2.4 Negative Cues in Meteorology

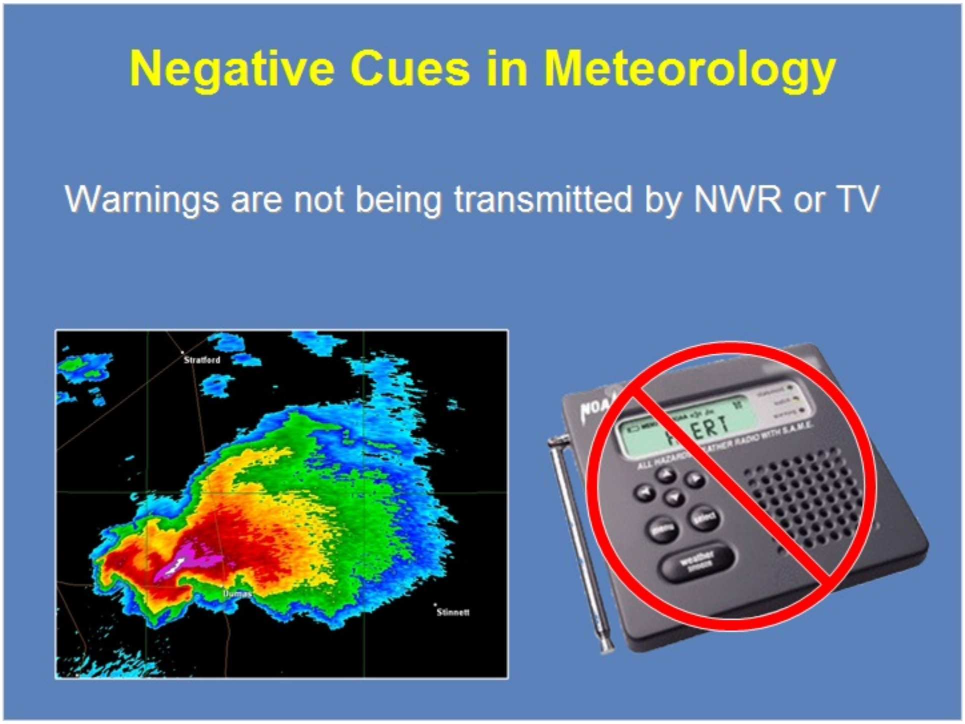

Notes:

Here's another example from real life. Your office sees this storm shown on the left and decides to issue a tornado warning (which is a good decision). However, you keep getting calls and tweets from the public and media partners asking, "Are you going to issue a Tornado Warning?" When you tell them that you have issued one, they proceed to tell you that they are getting no indication of a warning…even from weather radio. Oh, and by the way, they are also getting reports of a tornado.

In this situation, your partner is observing the negative cue with this storm that should have a Tornado Warning. What does this situation tell you? Maybe you issued a warning for the wrong storm. Maybe you forgot to click that last button to submit your product in WarnGen. Maybe the NWS network is down and warnings aren't being disseminated to the public. You can learn a lot from the negative cues that others observe.

Back to Top

2.5 Negative Cues in Meteorology

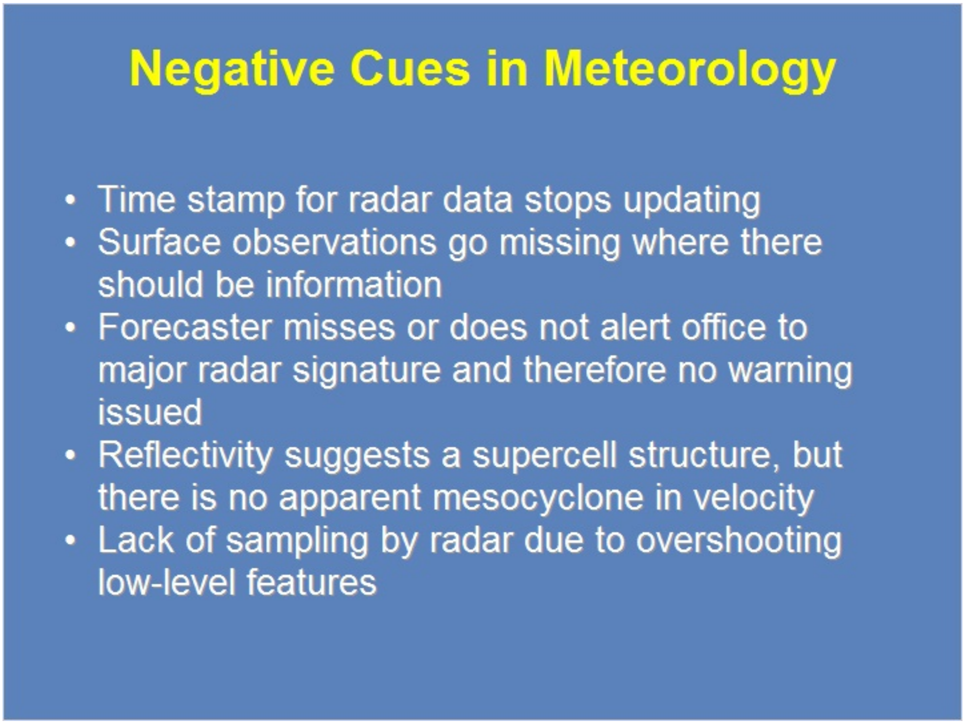

Notes:

Numerous examples of negative cues exist in warning operations. However, how do these apply to radar sampling issues? Well, this slide lists some negative cues that are related to radar interpretation. One of the main points of this lesson is radar sampling issues often lead to negative cues in a warning environment. Therefore, have a good grasp on radar sampling limitations and you'll be a better warning forecaster because you will be better equipped to observe and respond to negative cues.

Back to Top

2.6 Beam Heights & Uncertainty

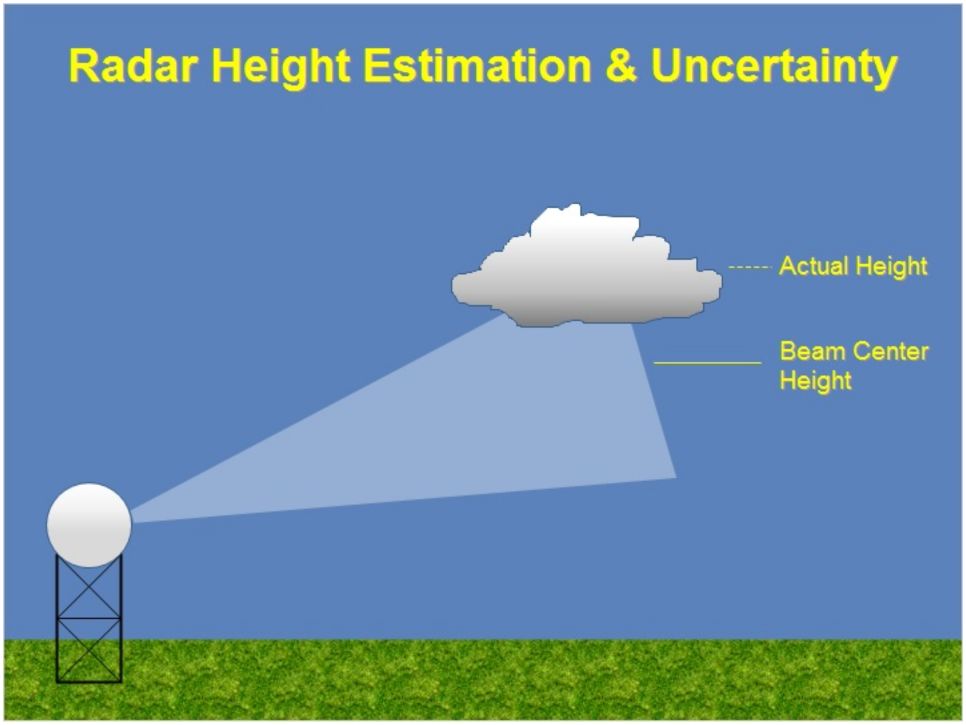

Notes:

Recall the radar beam spreads out as the pulse travels away from the radar, making the sampling area wider and wider. For locations as close as 60-70 miles from the radar, the beam's width extends 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 even though it is much higher. Just keep this fact in mind when interpreting radar echo heights in AWIPS.

Back to Top

2.7 Radar Height Estimates

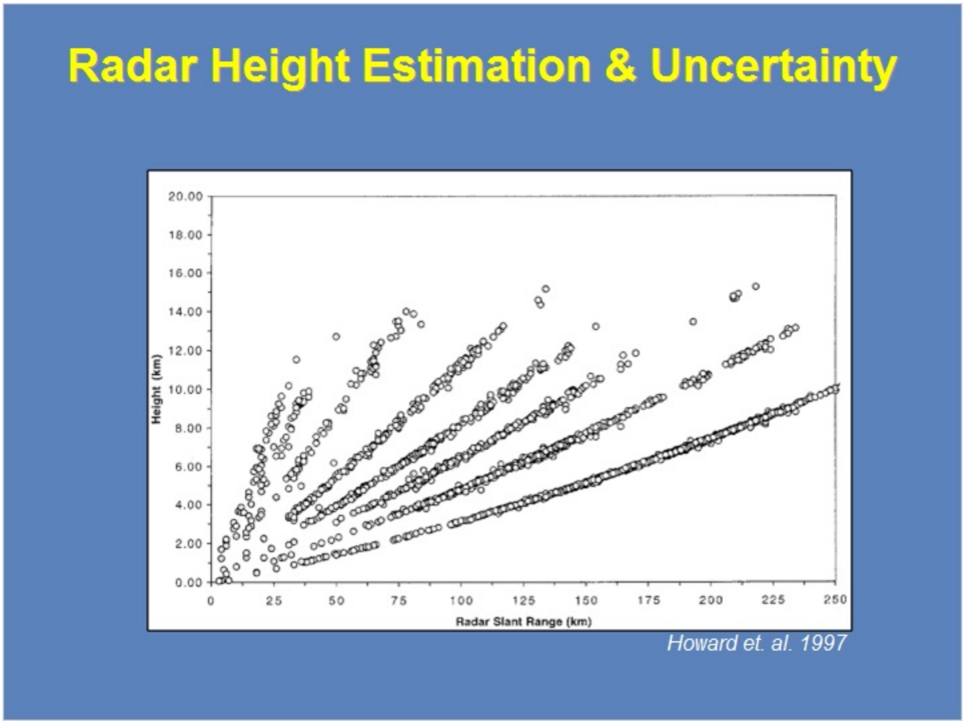

Notes:

A study done by Howard et. al. (1997) actually plotted the echo top heights as a function of range from the radar. Notice how all those radar echo top heights fall nicely along straight lines. Do you think this plot shows reality? Absolutely not! This graphic illustrates how using the beam center height to compute echo top heights is not the most accurate. In reality, storms exist in a continuous space, but the radar can only measure echo heights at discrete levels. So, what uncertainty exists in these height estimates? We'll look at that next.

2.8 Radar Height Uncertainties

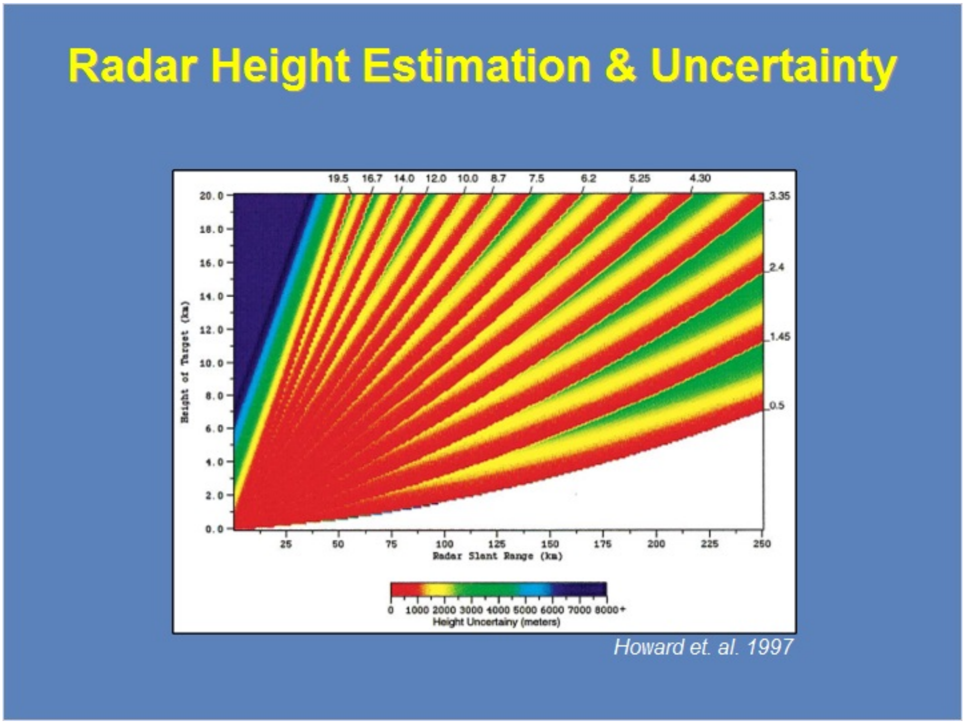

Notes:

This first graphic shows the uncertainty (in meters) of the target height estimate for each elevation angle in Volume Coverage Pattern (or VCP) 21. This color scheme is a little unusual in that red shades are good, and green and blue tones are bad! Notice how accurate the estimates are within 50 km range and below 4 km height. Almost all estimates are within 1000 m. As you got higher in height, or farther in range, the uncertainty increases to as much as 3 to 7 km due to the gaps in the scanning strategy.

Now, look at a similar graphic for VCP 11. The uncertainty increase happens at further ranges (past 75 km) and higher heights (above 8 km). So, choosing a scanning strategy with more elevation scans reduces the likelihood of uncertainty in your echo top height estimates.

Back to Top

2.9 Scanning Strategies

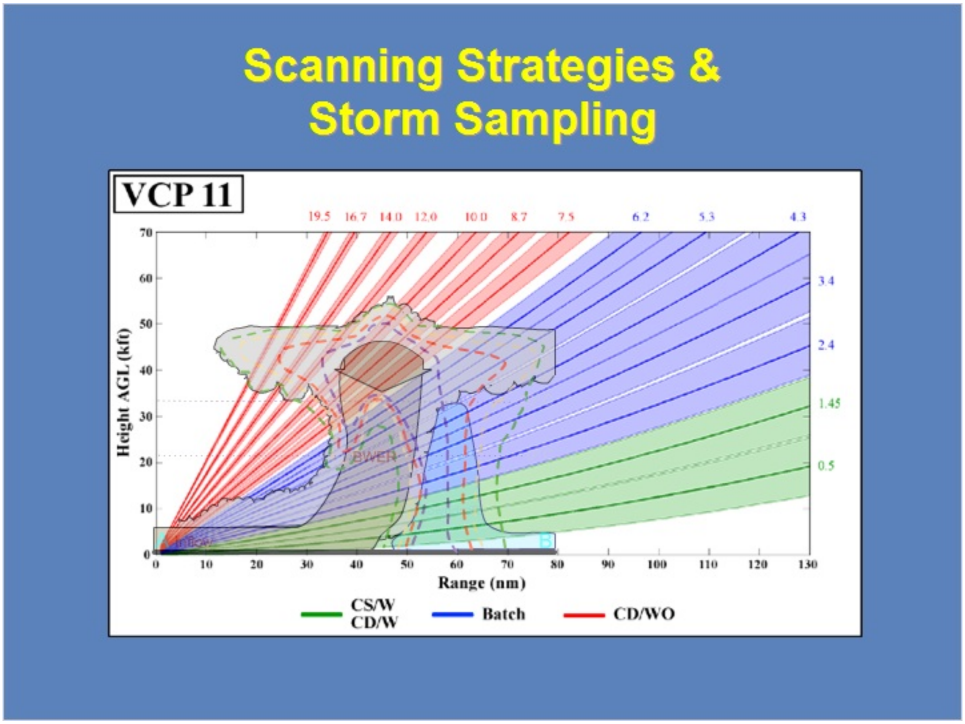

Notes:

To see how scanning strategies impact sampling, let's start with an example. Suppose we have a supercell that is observable by a radar. If we were to choose Volume Coverage Pattern (or VCP) 21 to sample this storm, what would it look like? Notice how the low level sampling is decent. Not many gaps are present. However, significant gaps can be seen aloft. We will likely miss some important features with this scanning strategy. If we switch to VCP 11, the gaps aloft disappear. If you are interested in features near the top of this storm, like an elevated hail core or storm top divergence, clearly VCP 11 gives you a better picture of what is going on up there. So, your scanning strategy choices can lead to better or worse sampling. Be mindful of the scanning strategy you choose.

Back to Top

2.10 VIL Differences between Strategies

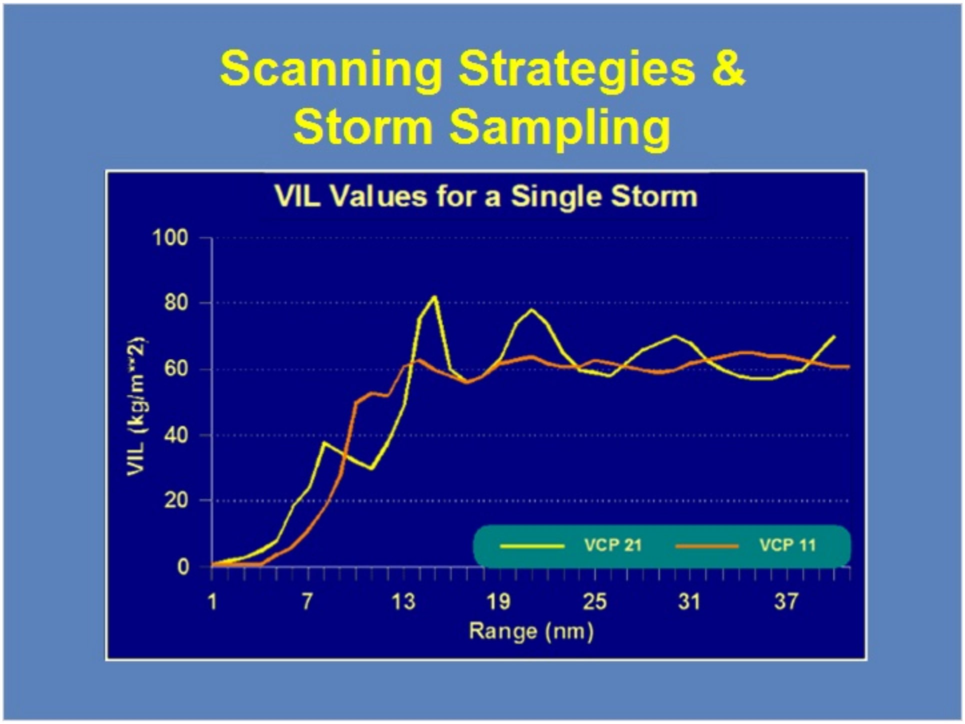

Notes:

The scanning strategy you choose also impact the radar algorithms. Here's an example using Vertically Integrated Liquid, or VIL. The chart shows VIL values as a function of range for two different scanning strategies: VCPs 21 and 11. Notice how the VIL values are rarely the same. In fact, the difference in values between the two VCPs can be quite large at certain ranges. Now, we know there's no difference in the storm's intensity just because of our VCP choice. The differences are strictly a function of scanning strategy.

2.11 Aspect Ratio

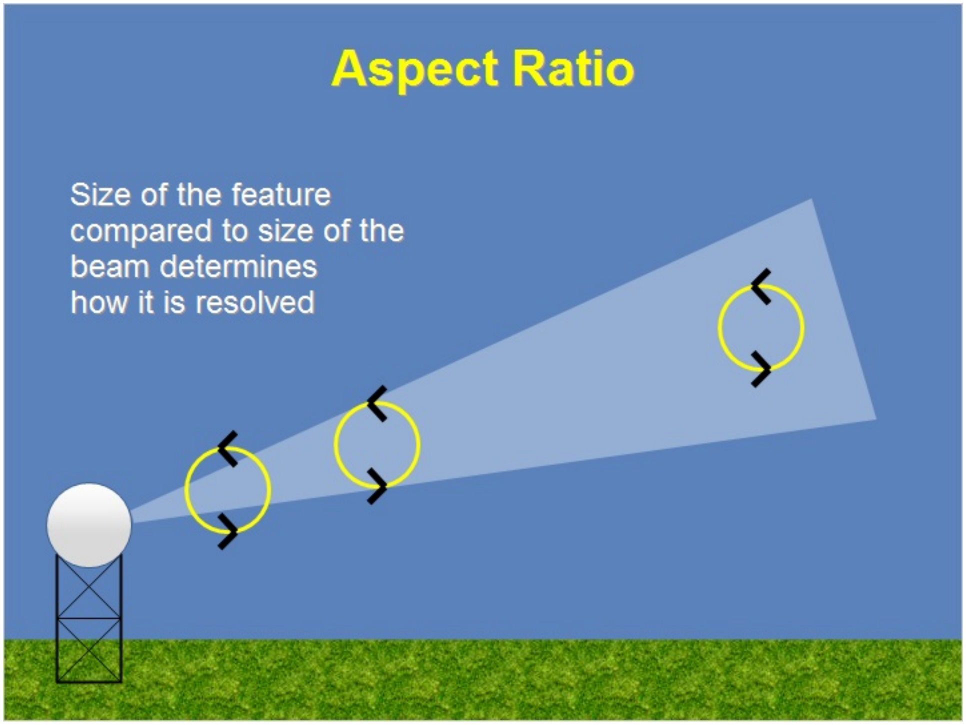

Notes:

A target's size compared to the size of the radar beam is called the aspect ratio. Apsect ratio matters because the radar beam spreads out with increasing range. So, a target of constant size will look different in the radar base data at different ranges. Take, for example, the idealized circulation shown on this slide. Close to the radar, the radar needs multiple beams to sample the circulation, providing a more detailed view of the feature. At medium ranges, the circulation may be roughly the size of the beam. So, you still get some detail, but not as much as you would at close ranges. At far ranges, only a portion of the beam is sampling the feature. In these cases, you'll either get very little detail or the feature may not be resolvable at all. In other words, the greater the targets range, the less impressive a feature will appear because less detail is apparent.

2.12 Aspect Ratio EXAMPLE

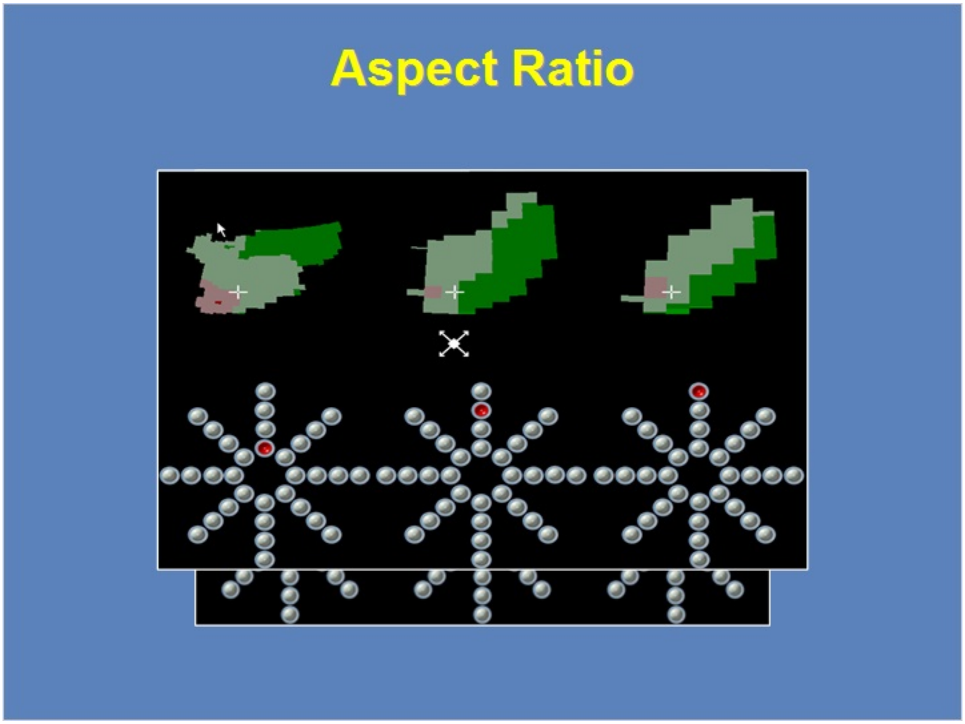

Notes:

To further illustrate the concept of aspect ratio, we show an example of the same storm that is viewed to the south of a radar at three different ranges. I should note this example is from an old application that predates both the dual-pol and super-res upgrades to WSR-88D, so we will focus on the reflectivity and velocity displays. The displays show what the same storm looks like with range from the radar increasing from left to right among the examples. Notice the details you can see on the left at the closest ranges, especially the reflectivity gradient on the southeast side of the storm. At further ranges, the reflectivity appears less detailed, almost blocky. Remember that the beam will be observing the storm at a higher altitude at these farther ranges. Still, you can see how the storm might appear less intense when looking at the farther ranges as compared to the closer distances.

Switching over to velocity shows a similar pattern. On the closest look, you can see more details of the circulation. At the further ranges, the circulation appears weaker.

Back to Top

2.14 Radar Horizon

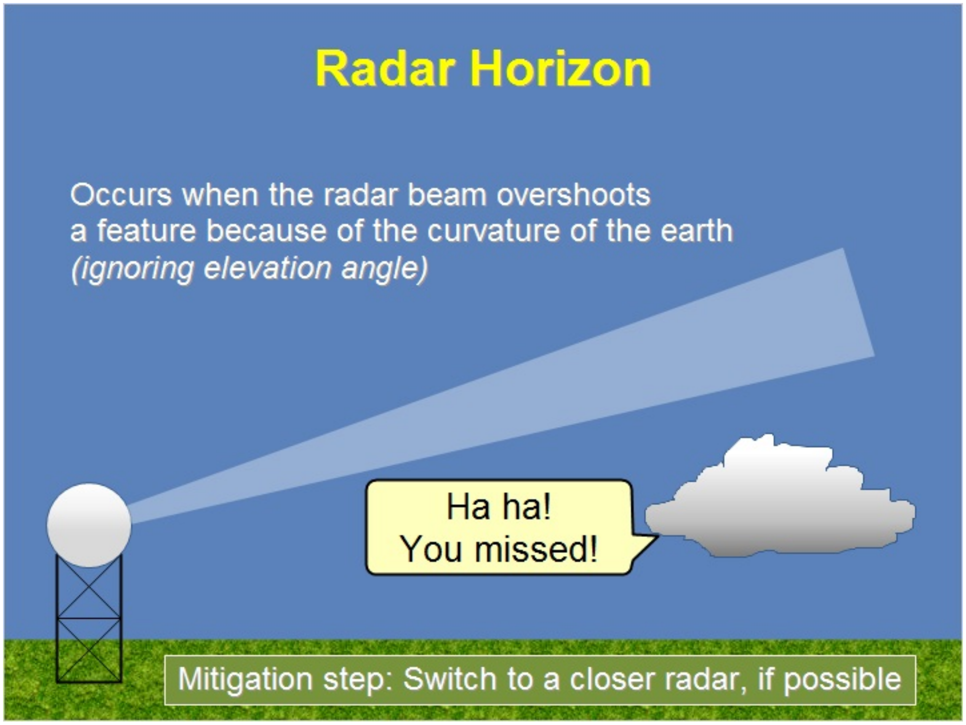

Notes:

Recall, as the beam propagates away from the radar, the earth's curvature causes the radar pulse to sample areas of increasing height as its range from the radar increases. In practical terms, this relationship of increasing beam height with increasing range means low level features will not be seen by the radar at far ranges. To mitigate this problem, forecasters can switch to a closer radar when one is available. When you can't switch to a closer radar, forecasters are forced to draw conclusions about unseen storm features based on the data you do have. Let's look at an example on the next slide.

Back to Top

2.13 Radar Horizon EXAMPLE

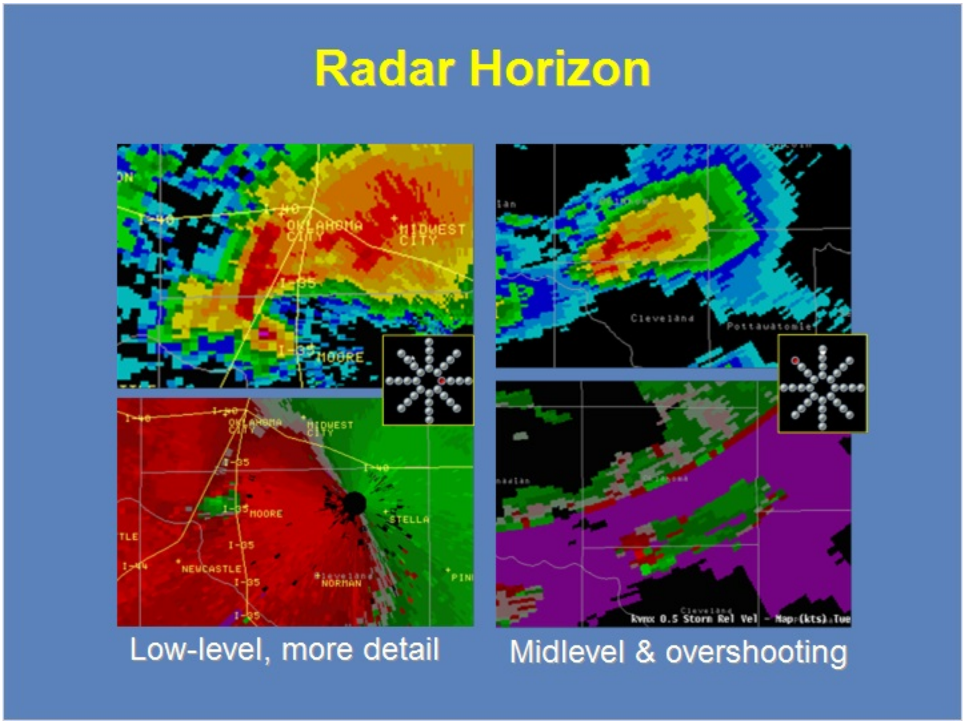

Notes:

This example shows a supercell thunderstorm producing a tornado moving through a major metro area. The reflectivity and velocity images on the left show the storm from a radar located to the east at very close range. The comparable images on the right display data from a radar located much further away to the northwest of the storm. Notice how the images on the left show more details in both products. The images on the right show the storm's structure at midlevels, where the radar beam is overshooting the low-level hook echo in reflectivity, among other features. Therefore, if low-level features are your primary concern, then make sure you keep this limitation in mind and use the closest radar to observe these features.

Back to Top

2.15 Feature Location

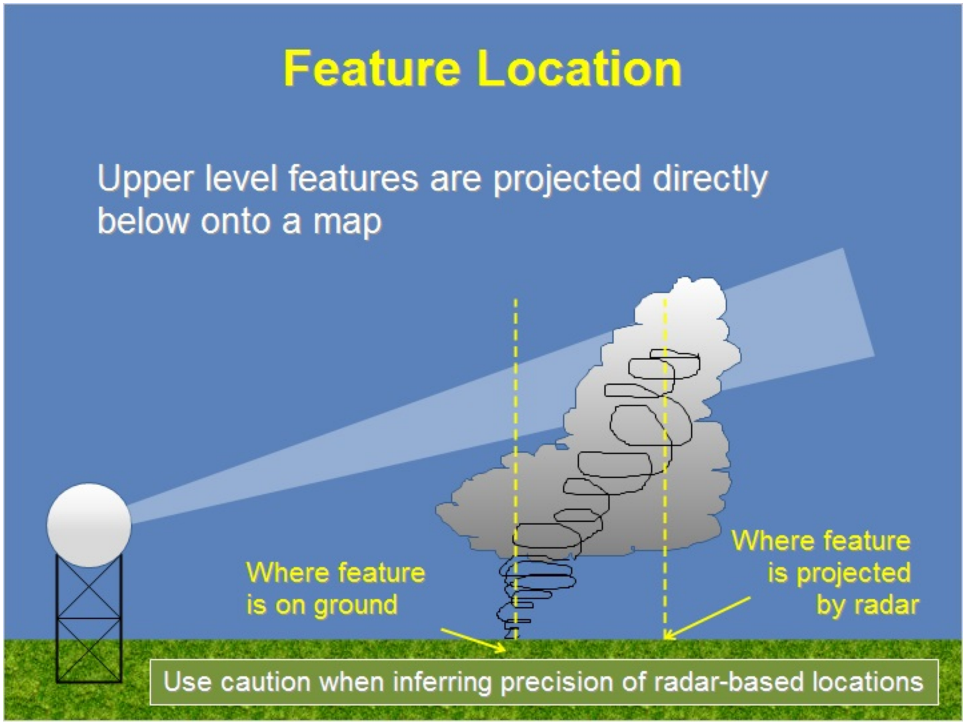

Notes:

As mentioned in another area of this lesson, the radar beam samples higher elevations of the atmosphere as its range from the radar increases. However, radar imagery projects these data onto the ground surface directly below that point. Many meteorological targets tilt vertically. In other words, features observed in the radar beam are not vertically stacked. For example, tornadic circulations will likely not be located directly underneath the midlevel mesocyclone. You can see from the illustration on the slide that the surface projection based on radar is here, but the actual surface feature is located here. Another example would be hail cores, which travel horizontally as they fall and not be located directly underneath the hail core aloft. So, what's the lesson here? Be careful about how much precision you infer from a circulation's location based on what you see on radar.

Back to Top

2.16 Feature Location EXAMPLE

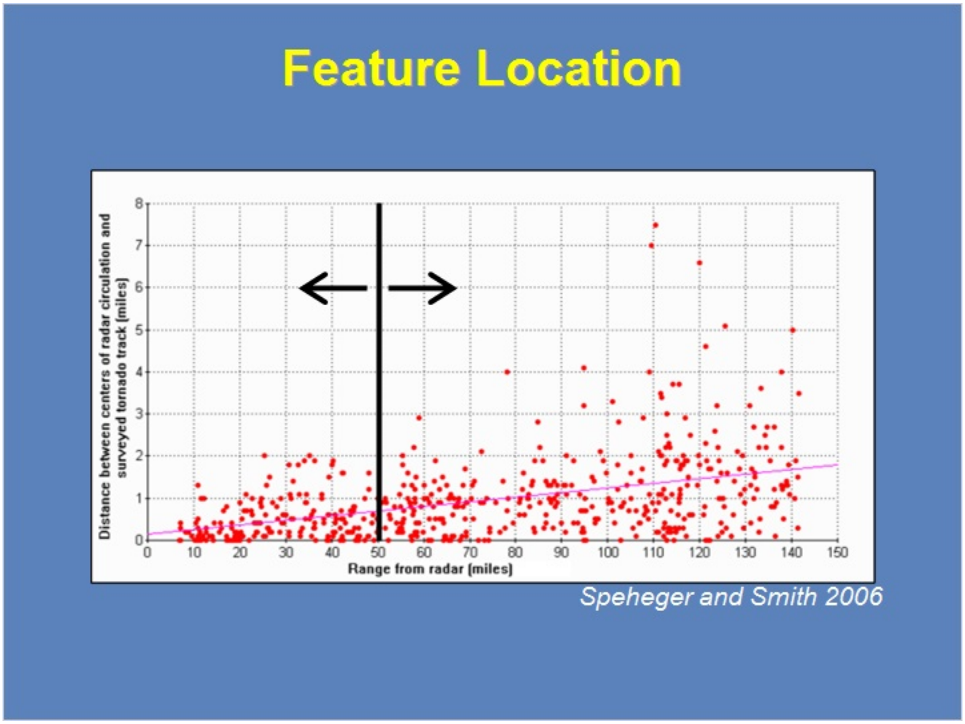

Notes:

The scatter plot shown here comes from a study by Speheger and Smith (2006). The diagram identifies the distance between where a radar identified circulation exists relative to the actual circulation location at the surface. Within 50 miles of the radar, errors remain small, averaging less than 3/4 mile and usually less than 2 miles. Once you get past a range of 50 miles, the distribution becomes more noticable. In this study, several circulations had errors as large as 4-8 miles! Therefore, be careful with how precisely you portray feature locations on radar compared with where they occur on the ground.

Back to Top

2.17 Feature Location EXAMPLE 2

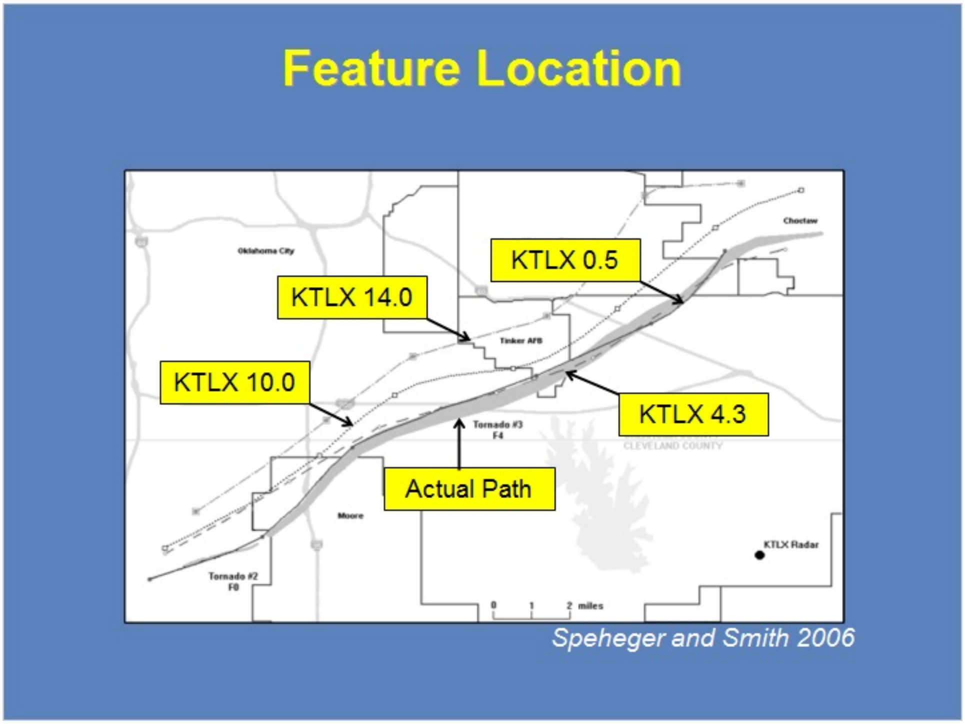

Notes:

Let's look at a specific example illustrating this issue, also from Speheger and Smith (2006). Notice how, as the radar elevation angle increases, the distance between the radar circulation location and the actual tornado path increases. So, again, be careful assuming that a radar-determined location precisely identifies the surface circulation location, especially on the higher tilts.

Back to Top

2.18 Beam Blockage

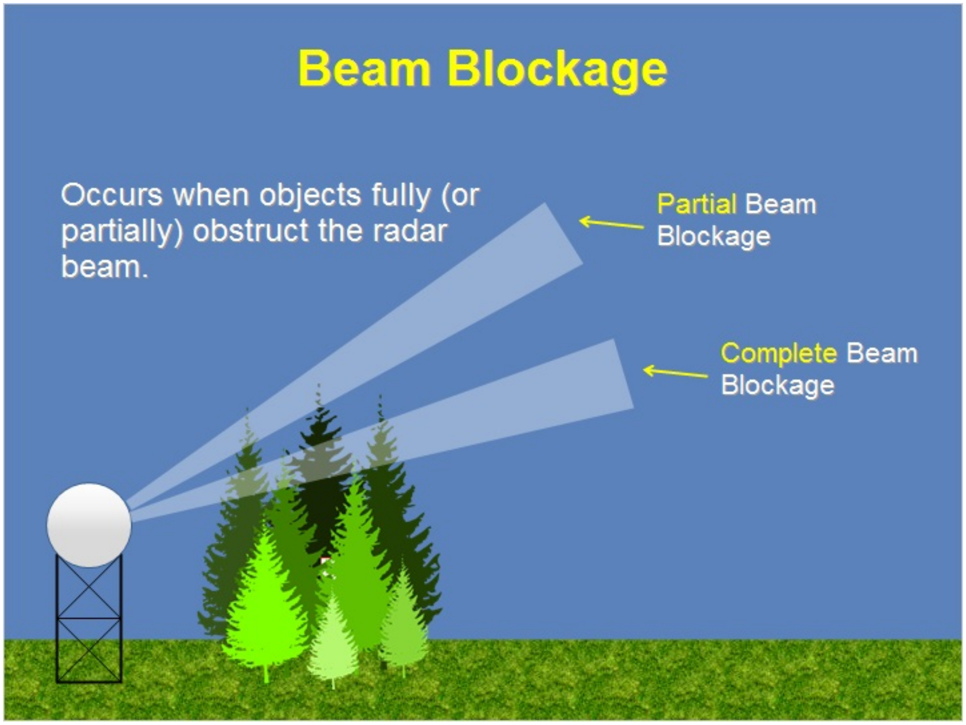

Notes:

When tall objects reside near the radar, these objects can obstruct the beam and cause a power loss significant enough where either no targets down radial from the object will be visible at all, or the targets will appear much weaker than if the obstruction was not present. Objects that can obstruct the beam include nearby trees, buildings, or even a mountain range. In the example shown, the trees block the entire beam where as the higher beam experiences only partial beam blockage.

Back to Top

2.20 Beam Blockage EXAMPLES

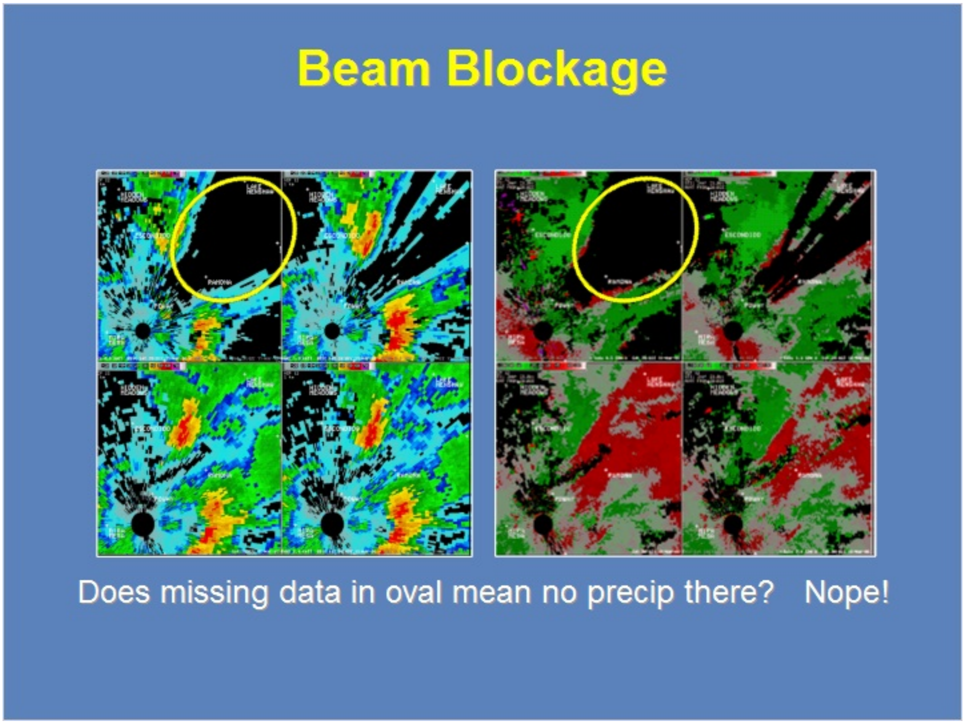

Notes:

Let’s look at an actual example of beam blockage near a radar. The lowest tilt (top left panel of both images) shows a data gap to the northeast of the radar. Does this gap mean no precipitation is present here? Nope. An obstacles blocks the radar from seeing down radial targets in this area. In fact, the blockage looks complete. So, what might we have missed due to this beam blockage? In this case, there was a tornado with the storm located on the left edge of the beam blockage. Some cyclonic shear is visible aloft. However, we don’t know if the radar would have shown the circulation better near the surface because the data are unavailable.

Back to Top

2.21 Beam Blockage EXAMPLES

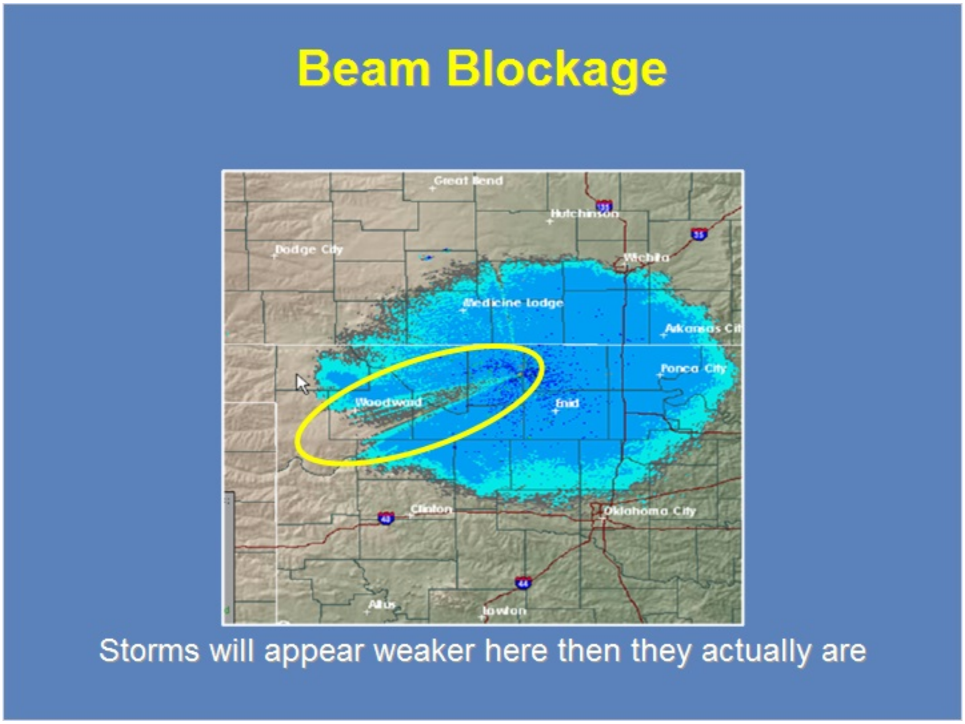

Notes:

This example shows a situation where there is partial beam blockage to the west-southwest of the radar. In instances like this one, remember that any storms down radial of the blockage will appear weaker than if no blockage was present. Therefore, you will need to make some mental adjustments to your expectations.

Back to Top

2.19 Side Lobe Contamination

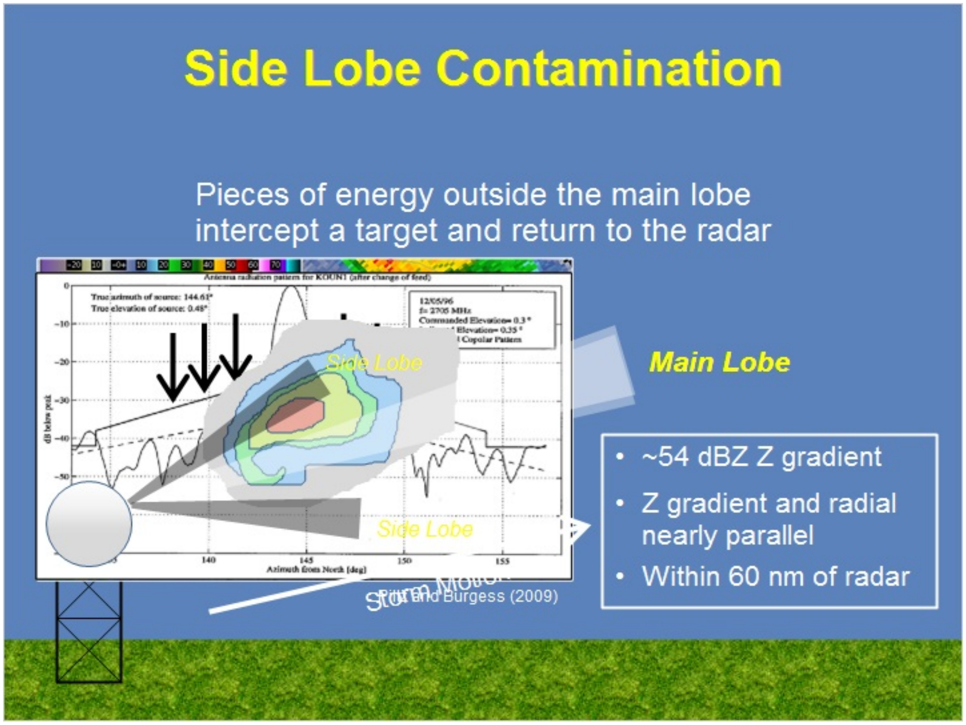

Notes:

Recall from the WSR-88D Fundamentals lessons that some of the transmitted energy propagates outside the main lobe in areas called side lobes. During some situations, targets reflect this energy back to the radar and corrupt the returns collected in the main lobe (Piltz and Burgess, 2009). While these situations are rare, forecasters should know when side lobe contamination tends to occur. Side lobes can occur when a strong storm has a significant azimuthal reflectivity gradient…usually where low-level inflow enters the storm's updraft. The gradient should be at least 54 dBZ over 3 degrees of azimuth of the main beam for the primary side lobe to contribute. Secondary and tertiary side lobes exist between 5 and 8 degrees, but reflectivity gradients would need to be over 70 dBZ across those azimuths for contamination to occur. Likewise the storm core in question needs to be relatively close to the RDA, say 60 nm or less, for side lobe contamination to be possible.

Lastly, when the storm motion vector also aligns nearly parallel to the reflectivity gradient and the radial, side lobe contamination may occur over several volume scans. It's during these situations that side lobes cause the most problems. Let's look at some examples to see why that is the case.

Back to Top

2.27 Side Lobes EXAMPLE

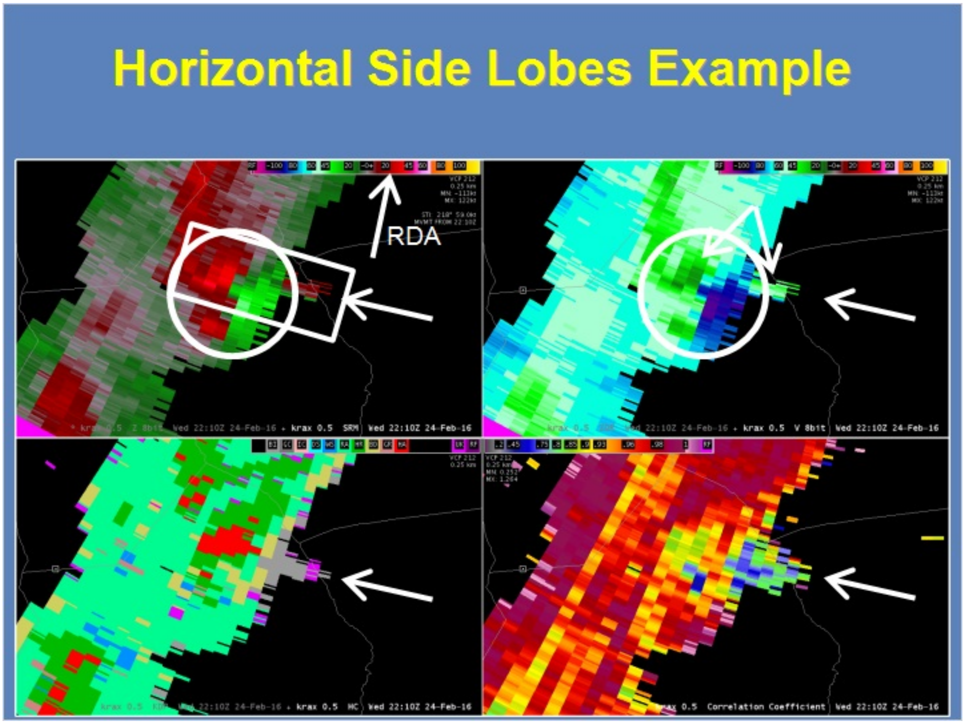

Notes:

Here's a good example of primarily horizontal side lobe contamination. The radar (located to the NNE) samples the storm at a range of ~50 nm. A fairly tight gradient of reflectivity exists on the east side of the storm. Directly to the east of, and at the same range as, the storm core, a weak reflectivity feature protrudes out to the east of the storm. In this area, side lobe contamination has occurred. Notice the ZDR and CC values in this area aren't indicative of precipitation. If we switch over to the other products, you see the HC algorithm has tagged the data as biological. Notice how the radial velocity values in this area better match those from the storm core then those in the adjacent radials.

The storm motion (218 degrees at 59 kts) takes this storm on a path roughly parallel to the radial. This orientation allowed the side lobe contamination to occur for several volume scans. Fortunately, forecasters can quickly identify the corrupted data and see that it doesn't impact their interpretation of the key feature in the velocity data: the mesocyclone located directly to the west of the bad data. However, data corruption can be more difficult to spot and result in poor warning decisions if forecasters are not careful. Let's look at another example to see how that can happen.

Back to Top

2.22 Side Lobes EXAMPLE 2

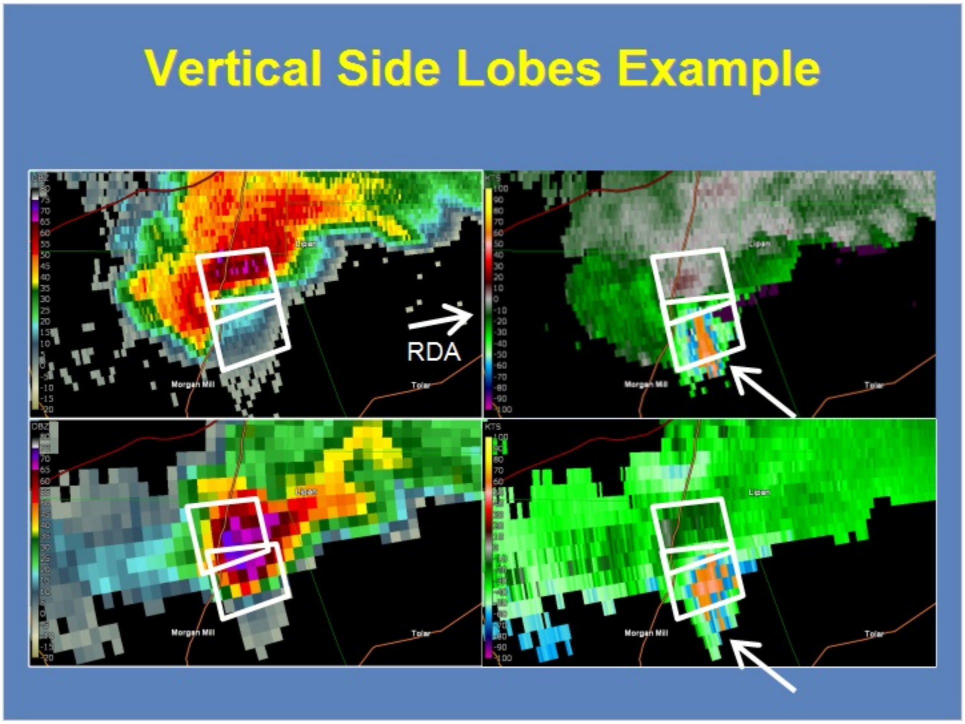

Notes:

This next example results from side lobe contamination in mostly the vertical direction. The radar (located to the ENE) samples the storm at a range of 40 nm. As in the previous example, the reflectivity gradient on the inflow side of the storm lines up nearly parallel to the radar beam at 0.5 degrees. Unlike the previous example, the horizontal reflectivity gradient doesn't meet the criteria as its only around 40 dBZ over 3 degrees. Yet, the velocity and storm-relative motion look suspicious. If we look aloft, at 3 .4 degrees, we see stronger Reflectivity values than we did in the storm core at 0.5 degrees. Notice, also, how the velocity values in these areas match the core aloft better than the core at the surface. The Spectrum Width is very high, also, in the same area where the velocity values are anomalously high.

2.23 Final Note on Side Lobes

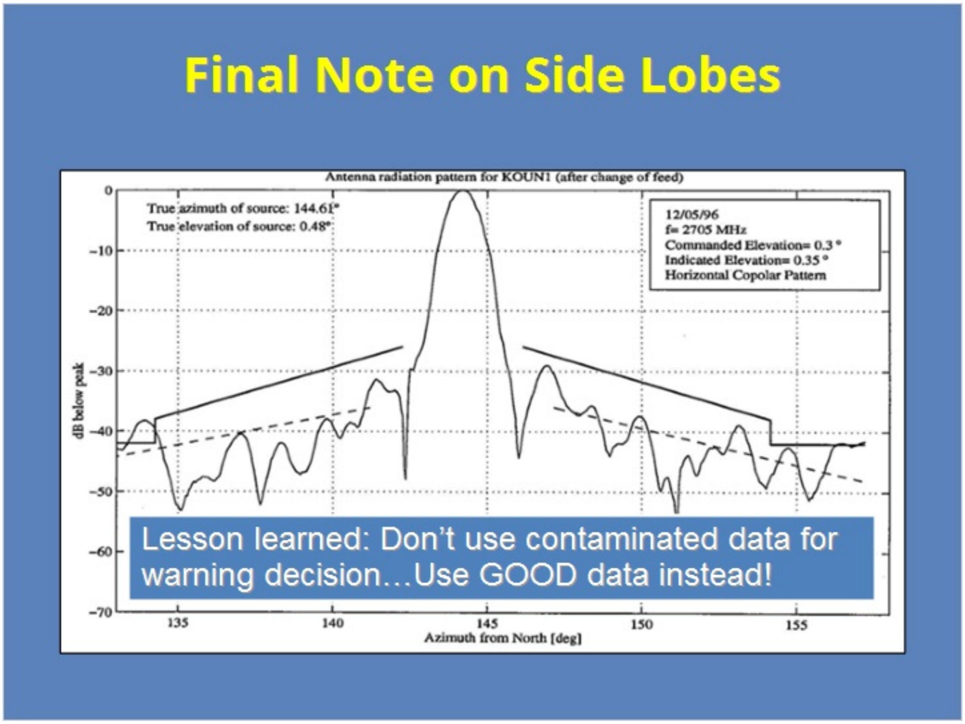

Notes:

Let me make one final note on side lobe contamination. Side lobe contamination has not been thoroughly researched, and supposition and inference both play a role when investigating errors related to this topic in live radar data (Burgess, 2017). We assume all WSR-88D antenna patterns look like the one I show here. With over 150 radars in the fleet (and with most of them having been in service for over 20 years), some variances likely exist between different radars. Some sites may have side lobes that run a little hot (i.e., where contamination happens at slightly weaker gradients). Other sites may have a misaligned feedhorn that could result in the azimuth angle of the primary side lobe being a little more or less than 3 degrees. Your office knows your radar better than just about anyone else, so use that local knowledge to help you know when data are contaminated.

So what's the key takeaway when it comes to side lobe contamination? Well, you need to know when data are contaminated and then don't use that data as part of your warning decision. Now, I'm not saying that storms with side lobe contamination should never have a warning. Instead, use the other GOOD data available to you to make that decision. More on this subject will be discussed in later topics of this course.

2.24 Viewing Angle

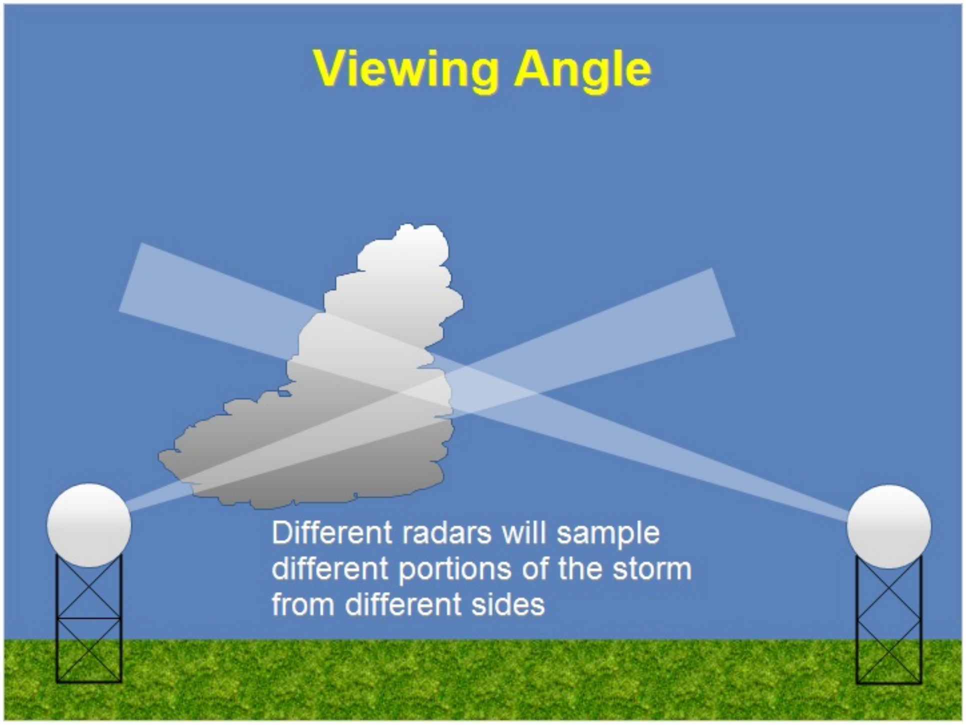

Notes:

In several places in the US, especially in the Central & Eastern US, storms are often sampled by multiple radars. Storms lack symmetrical structure both in the horizontal and in the vertical, so getting multiple views of a storm improves storm sampling and generally helps you see what you want to see. Forecasters need to know when a better viewing angle for a storm is available. Let's look at a few examples.

Back to Top

2.25 Viewing Angle EXAMPLE

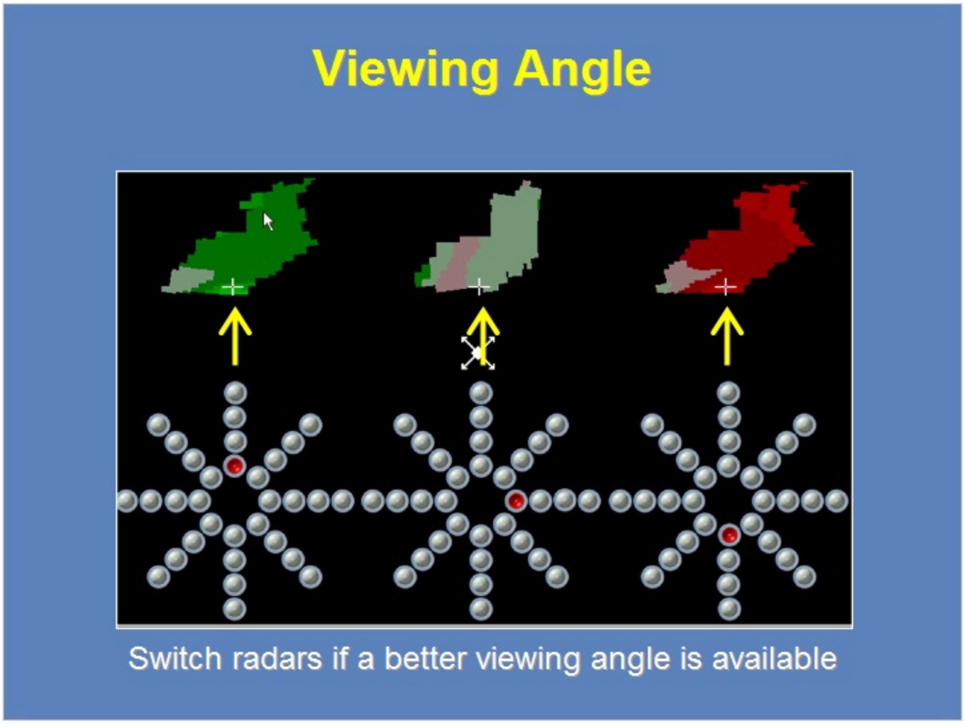

Notes:

Recall that your radial velocity measurements depend on your viewing angle. When the radar's beam aligns with the wind direction, wind speed measurement accuracy maximizes. When the radial looks perpendicular to the wind direction, the radar can't measure the wind speed at all. Usually, the radar observes an individual storm at some angle in between.

This image shows the same radial velocity data when viewed from 3 radars: one located north of the echo on the left, one located east of the echo in the center, and one located south of the echo on the right. Notice how the radial velocity changes depending on the radar's viewing angle, even though the actual wind field remains the same. Therefore, know your radar's location relative to your area of interest, and switch radars when another site provides a better look…especially when looking at velocity signatures.

Back to Top

2.26 Viewing Angle EXAMPLE

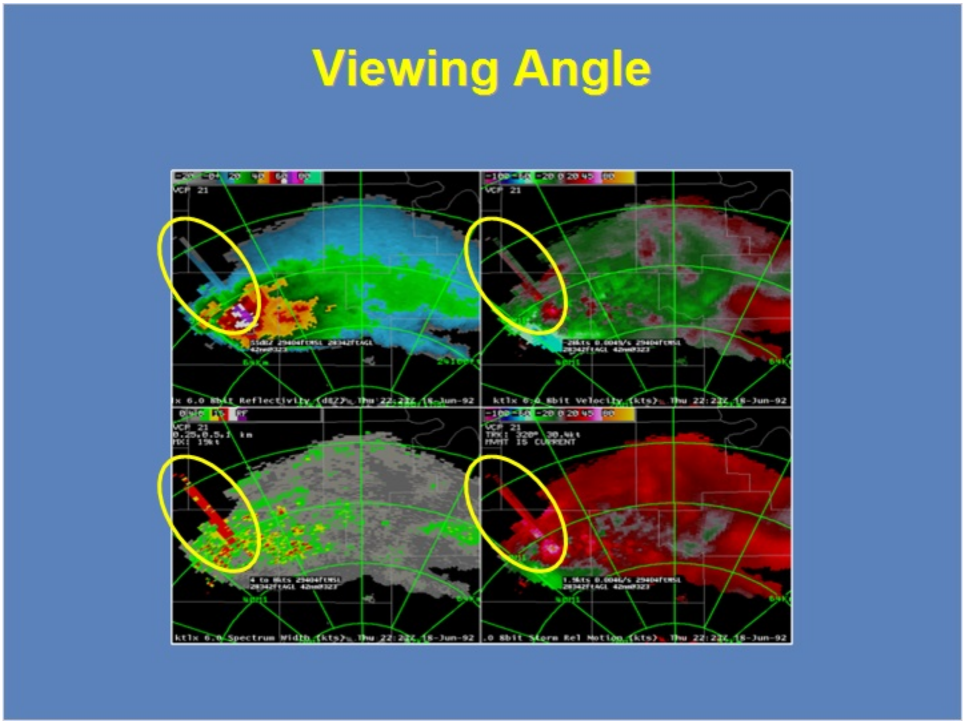

Notes:

Viewing angle impacts Radial Velocity the most frequently of all the base data, but other products, such as Reflectivity, are not immune. Three-body scatter spikes (TBSSs) often appear down radial of significant hail cores. The example on the slide show an example of a TBSS to the northwest of the storm's core. If another storm was located in this area, this signature might not be visible from this radar. To see the TBSS, you would likely need to switch to a radar with a different viewing angle. Likewise, data from other products in a TBSS region can be corrupted. If the radar beam passes through a hail core and then through a mesocyclone, then the velocity data for the mesocyclone could be negatively impacted by the TBSS. The solution to this problem: Pick a different viewing angle.

Back to Top

2.28 Non-Uniform Beam Filling & Differential Attenuation

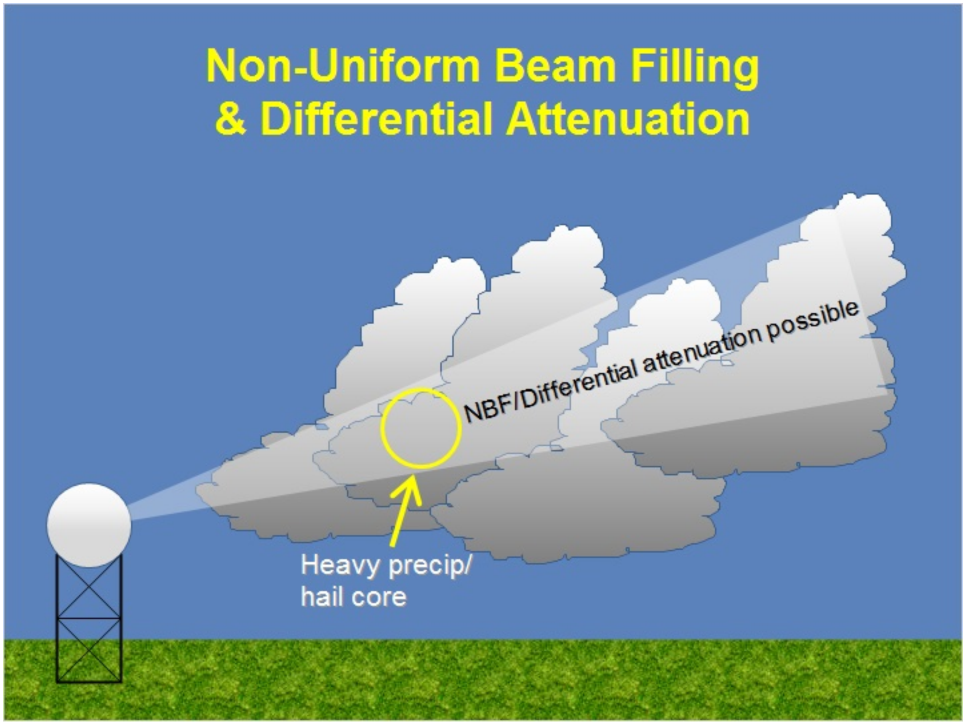

Notes:

The last two sampling issues on our list impact the dual-pol variables of Correlation Coefficient (CC) and Differential Reflectivity (ZDR), albeit in different ways. The first artifact, non-uniform beam filling, impacts CC while the second artifact, differential attenuation, impacts ZDR. As you'll see on the next few slides, these two sampling issues often occur in tandem, but they are caused by completely different processes. The most common situation occurs when strong hail cores exist, especially when the cores are aligned down a radial from the radar. Let's look at both of these issues in more detail.

Back to Top

2.29 Non-uniform Beam Filling

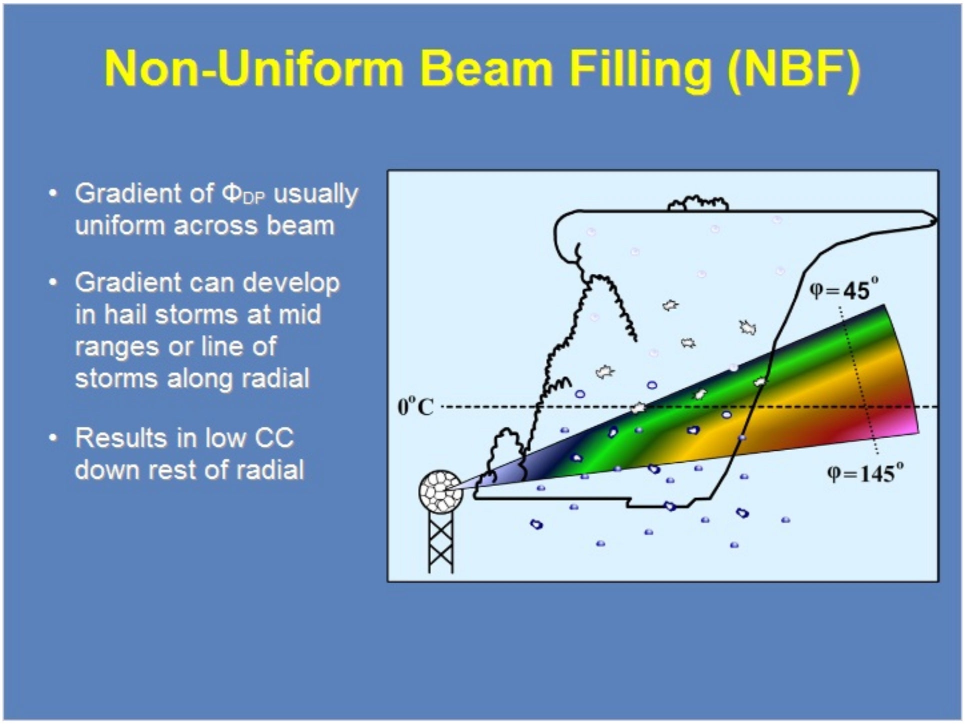

Notes:

In most circumstances, differential phase remains uniform across the radar beam. One instance where this condition isn't true occurs when hail storms are observed at medium ranges from the radar. Under these conditions, the radar beam can sample significant hail melt across the beam. When this happens, differential phase experiences little phase shift at the top of the beam, but a large phase shift toward the bottom of the beam where more liquid hydrometeors exist. This gradient in phase shift across the beam causes Correlation Coefficient to be reduced from that point and the rest of the gates located down radial. This condition, called non-uniform beam filling (or NBF), occurs down radial from the hail core. Now let's look a similar situation where NBF can occur.

2.30 NBF Example

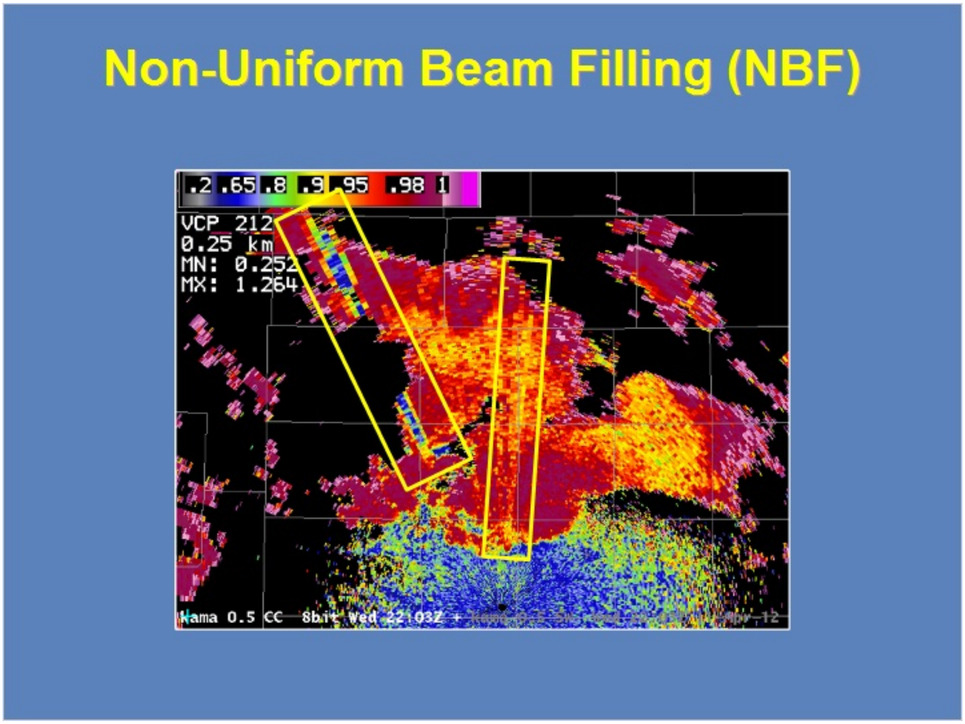

Notes:

This image shows storms that are causing non-uniform beam filling issues. Can you determine which storms are causing problems? Does it help when we switch over to CC? Hopefully, you said yes. The storm to the northwest caused significant NBF issues, as is visible in the reduced CC values down radial from the storm's core. Just north of the radar, another storm appears to have some NBF issues, too. The CC values drop less prominently, but they are still noticeable.

Back to Top

2.31 Differential Attenuation

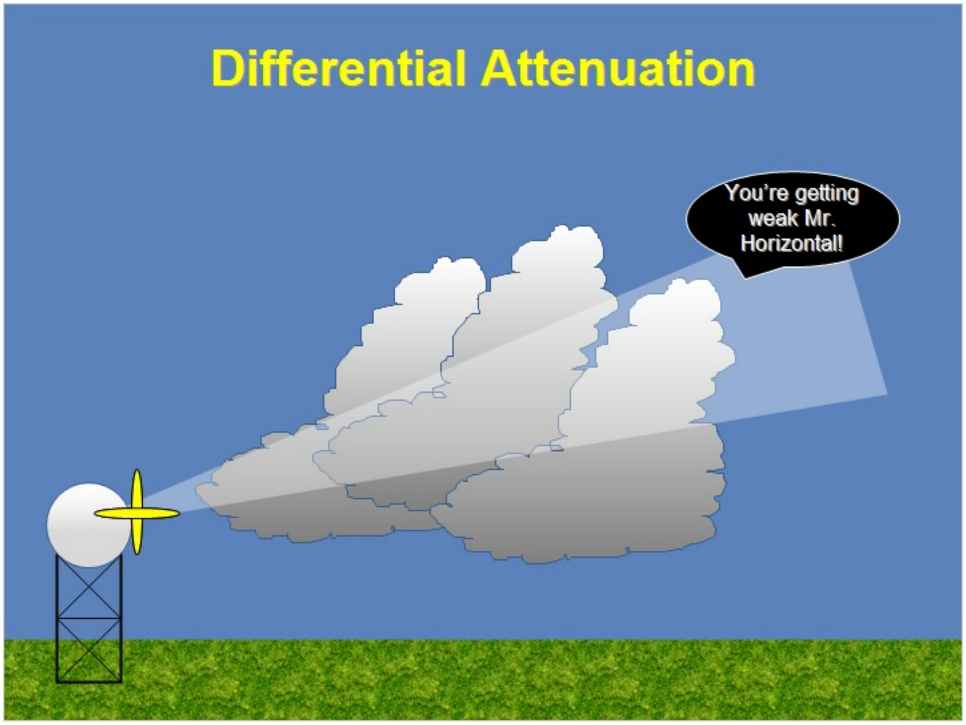

Notes:

Now let's move on to the second, related sampling issue that impacts ZDR. The radar transmits its pulse with two polarizations. As the pulse propagates through the atmosphere, each polarization usually attenuates (or slows down and weakens) at the same rate because most hydrometeors don't attenuate S-band radiation significantly. In some cases, however, the horizontal pulse will attenuate significantly, but the vertical pulse will not. These situations occur when the pulse travels through storms with very heavy rain and hail cores, or when several storms are aligned along a radial. In these cases, the precipitation impacts the vertical pulse far less than the horizontal pulse. As a result, the differential reflectivity skews toward lower values down radial of differential attenuation regions than they would otherwise be. In fact, the values are often negative. Let's look at an example.

Back to Top

2.32 DA Example

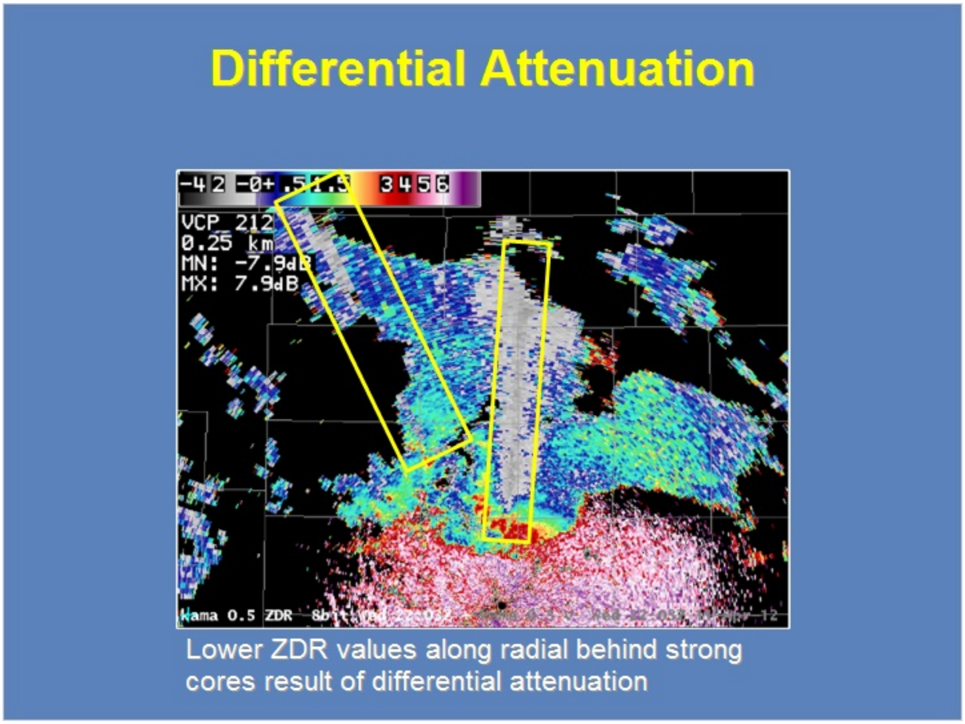

Notes:

This example should look familiar. It's the same case we just showed for the non-uniform beam filling example. Remember how we said that NBF and differential attenuation often occur in the same areas. Well, let's look downstream of the two storms we noted before. We don't see a noticeable change in Reflectivity even though there is probably some attenuation happening in that channel. Switching over to ZDR, we can see the attenuation more clearly down radial of both of these heavy rain and hail cores. These areas of negative ZDR oriented along the radial result from differential attenuation.

Back to Top

2.33 NBF/DA Differences

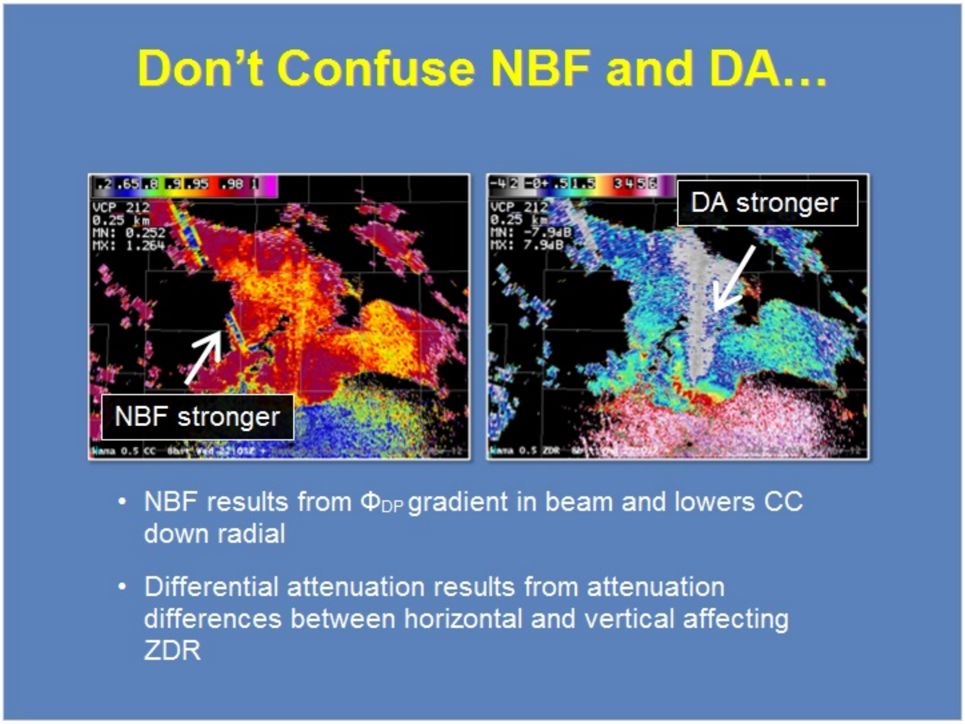

Notes:

As I showed previously, NBF and differential attenuation often occur in similar situations. The causes include the radar beam passing through a heavy rain and hail core or strong storms aligned along a radial. However, the physical reasons they occur are very different. NBF results from a gradient in Differential Phase, or PhiDP, within the radar beam that reduces correlation coefficient values for all the down radial range bins. Differential attenuation results from the horizontal channel being attenuated more in these situations than the vertical channel, causing ZDR to be lower than expected. It may not seem like a big deal that these are different processes, but it helps to know the differences when you only observe one phenomenon, but not the other.

Back to Top

4. Completion

4.1 Completion!

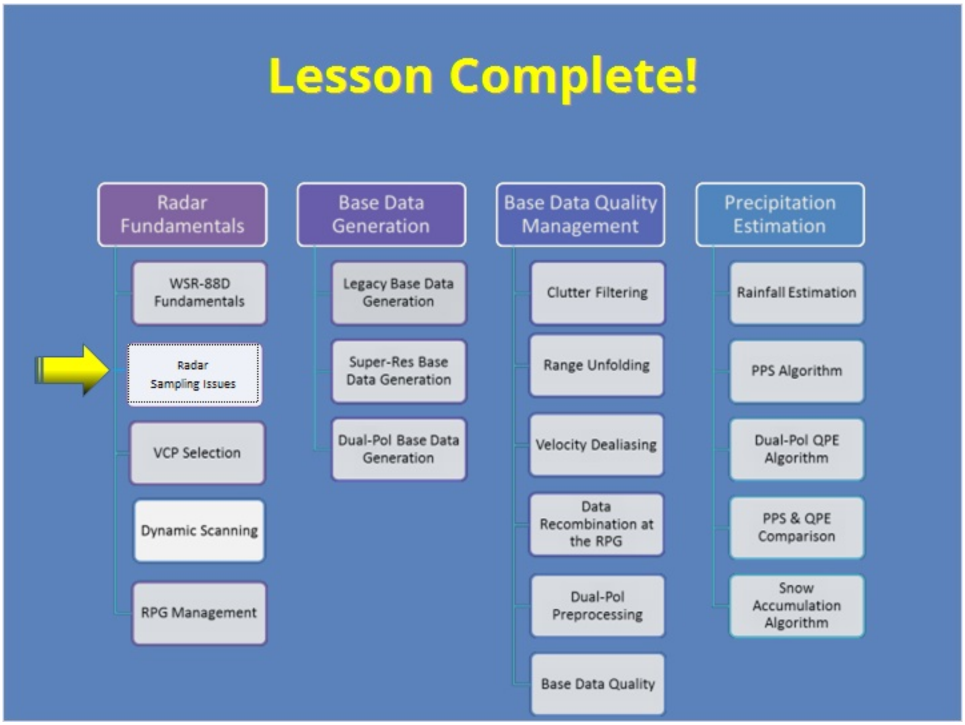

Notes:

Thanks for your attention! You are now complete for this lesson.

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.

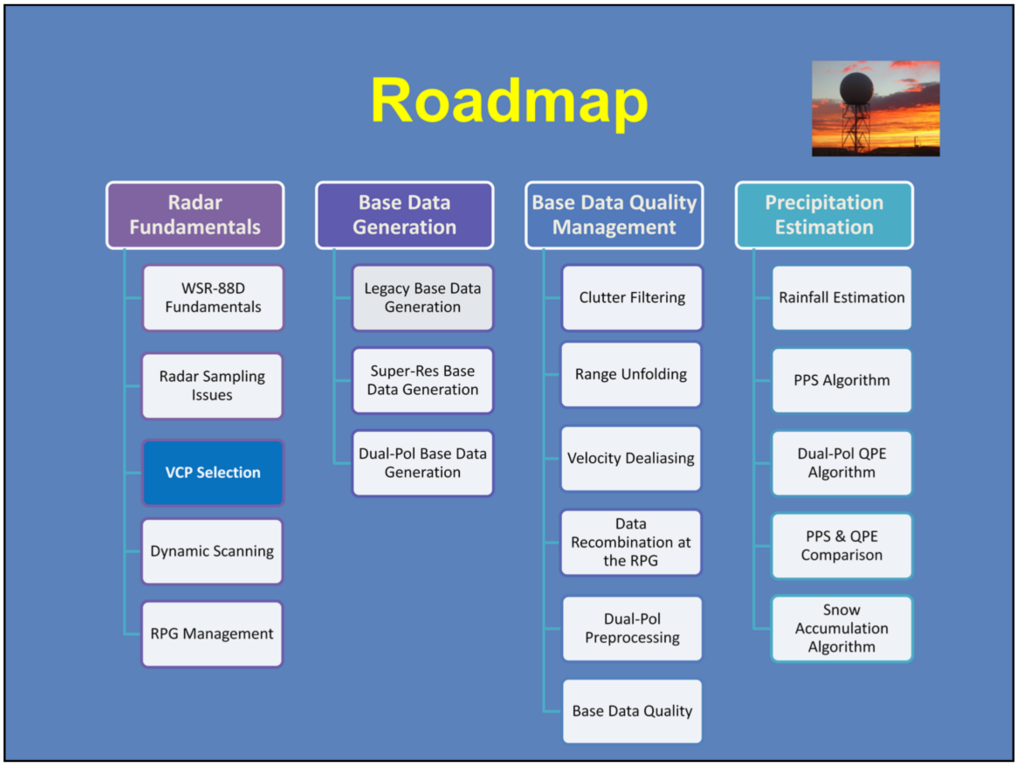

Here is the “roadmap” with your current location.

Back to Top

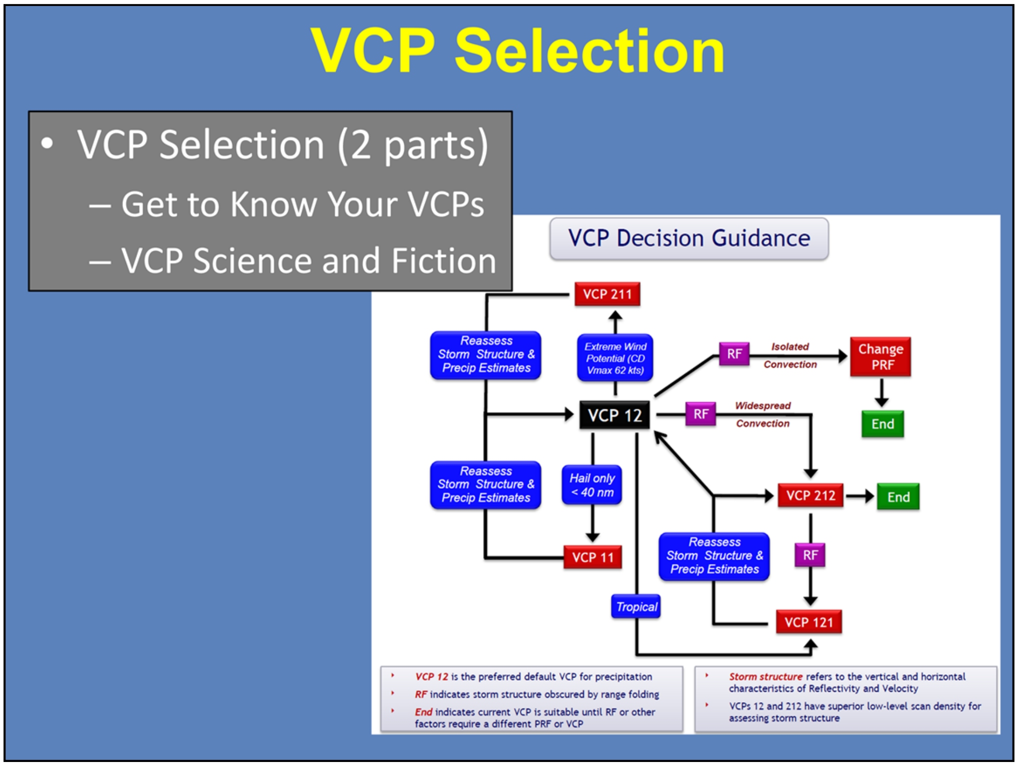

In this lesson on VCP Selection, we will cover the two main parts we want to address in the VCP decision-making process. First, you have to know your VCPs, what their strengths and limitations are – and how they compare to one another. Then, we'll spend some time debunking some VCP myths or fiction in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

Back to Top

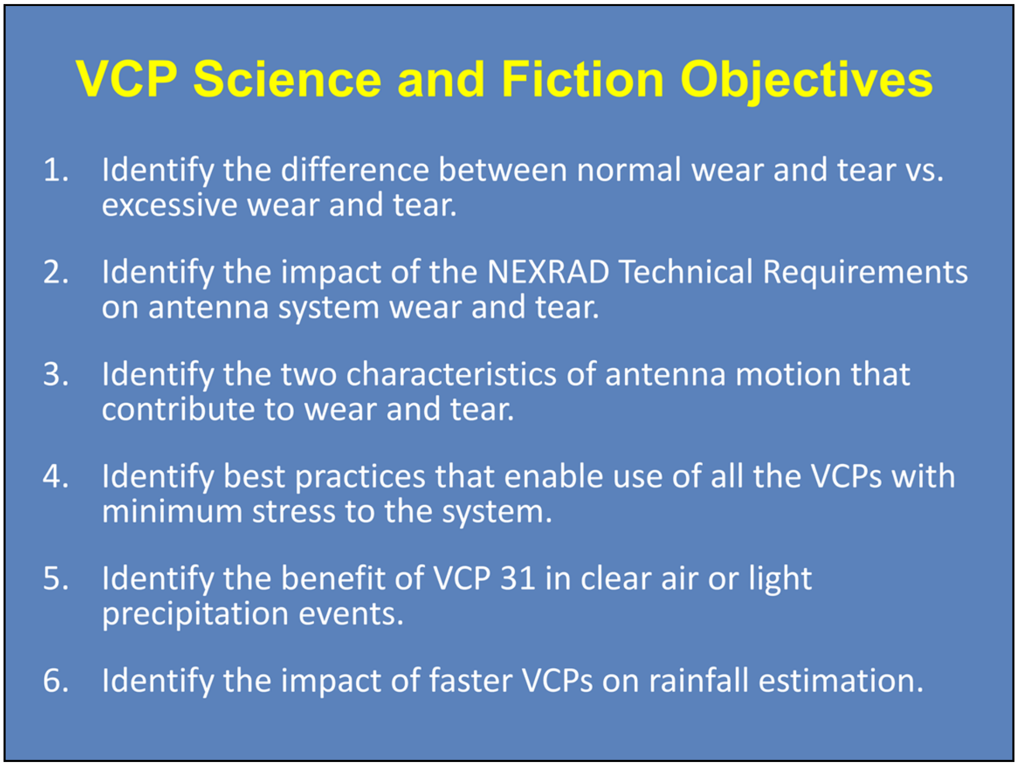

Here are the learning objectives for VCP Science and Fiction. Notice that the first four focus on VCP choice and how it relates to wear and tear on the WSR-88D hardware, which has been a significant concern with the introduction of the faster VCPs. The remaining objectives address the use of VCP 31 and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

Back to Top

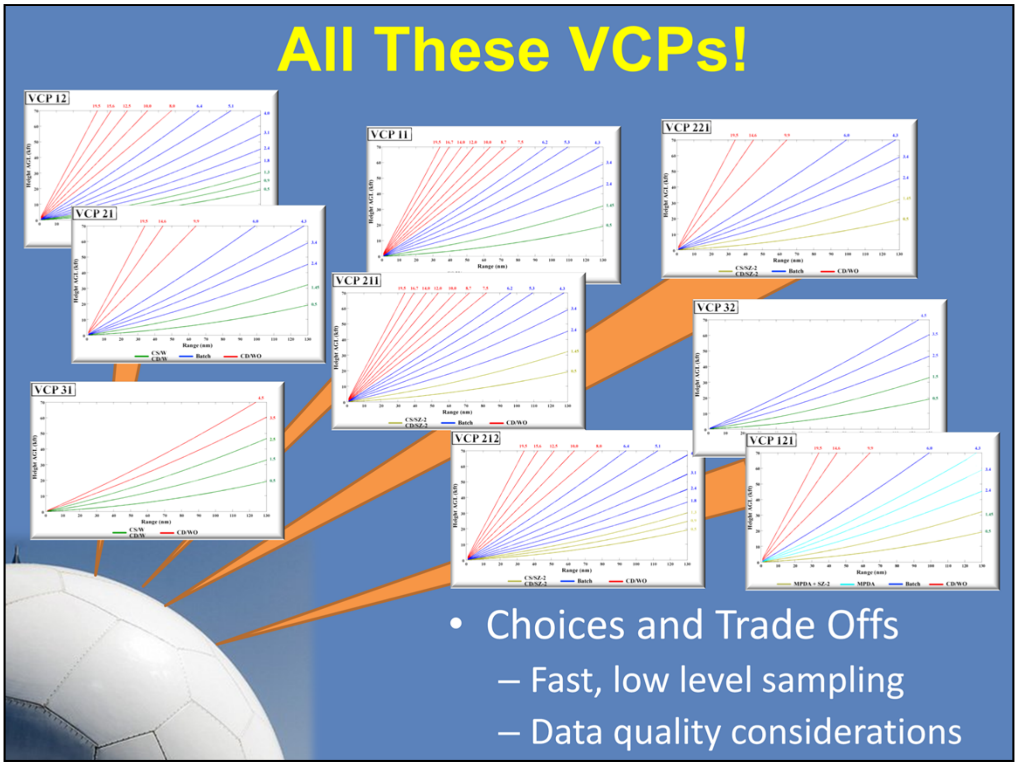

Currently, there are nine different VCPs to choose from, which has evolved over the past couple of decades with from changing needs and derivative capabilities of older VCPs. Your choices involve trade offs, such as the benefits of fast, low level sampling vs. the impact on data quality of fewer pulses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

Back to Top

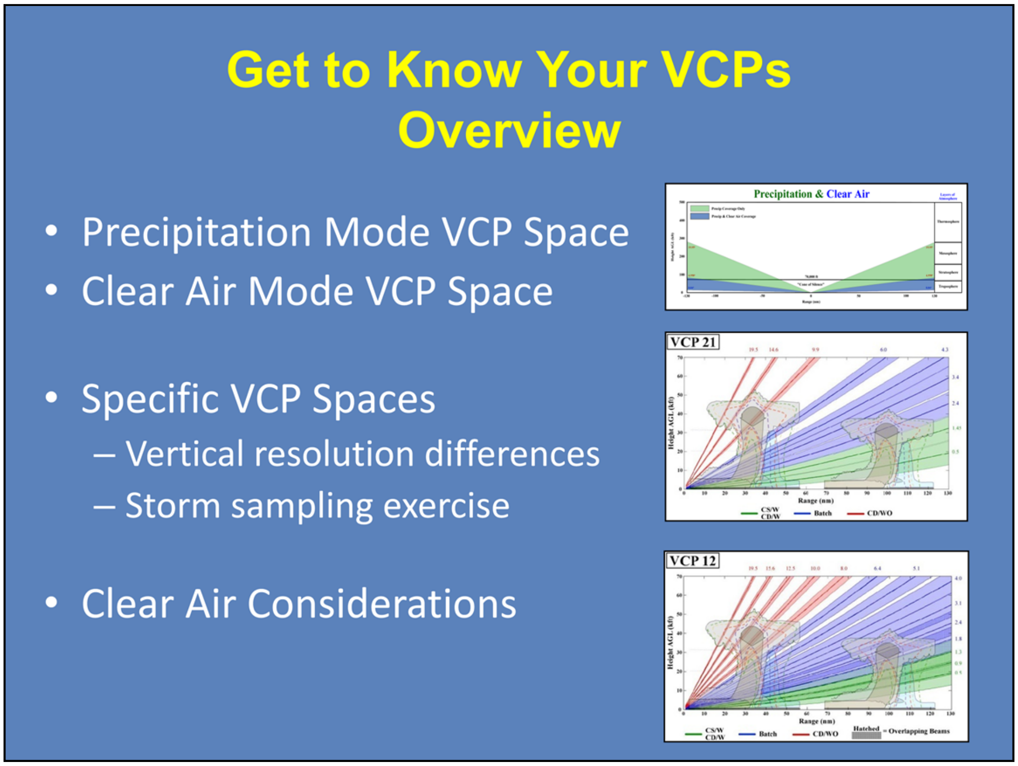

To set the foundation for getting to know your VCPs, we'll start with the atmospheric space that Precipitation Mode and Clear Air mode sample. Then the space sampled for each of the VCPs is presented, along with animations of the angles used during a volume scan. These animations will help you to visualize 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.

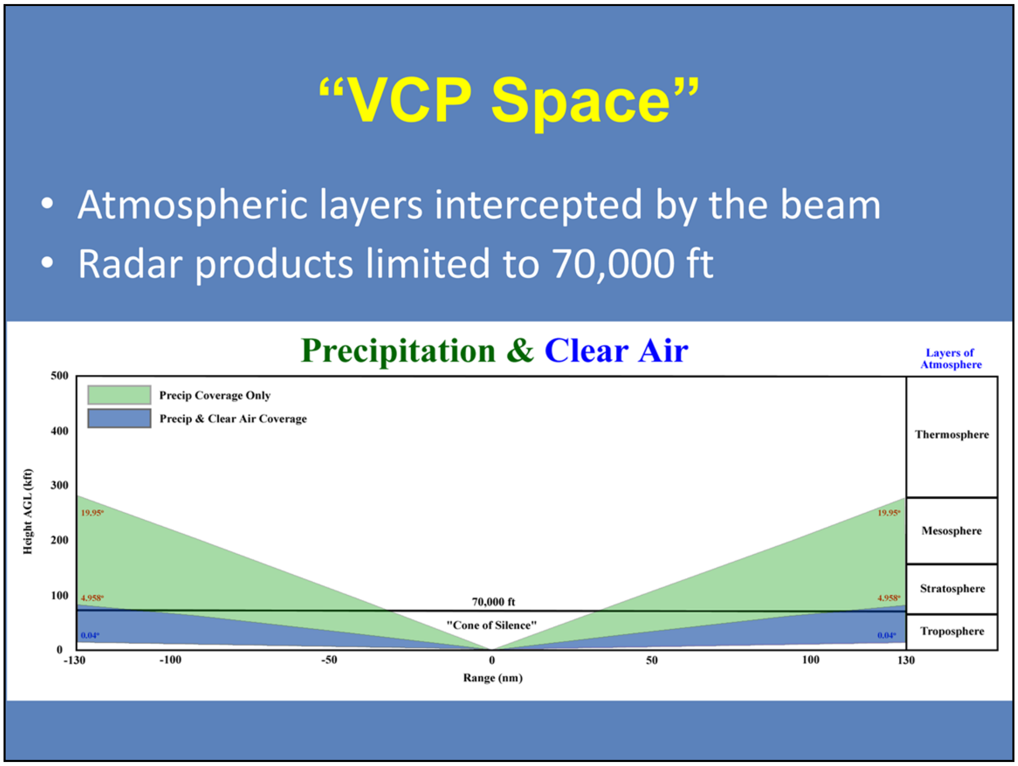

Let’s look at a quick cross-section of the quote VCP space, or the amount of atmosphere that’s actually being sample by a doppler radar, both in clear air and precip modes – and showing just how vast the cone of silence is above a particular radar. This image is mostly a reminder of how thin the Tropospheric 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!

Back to Top

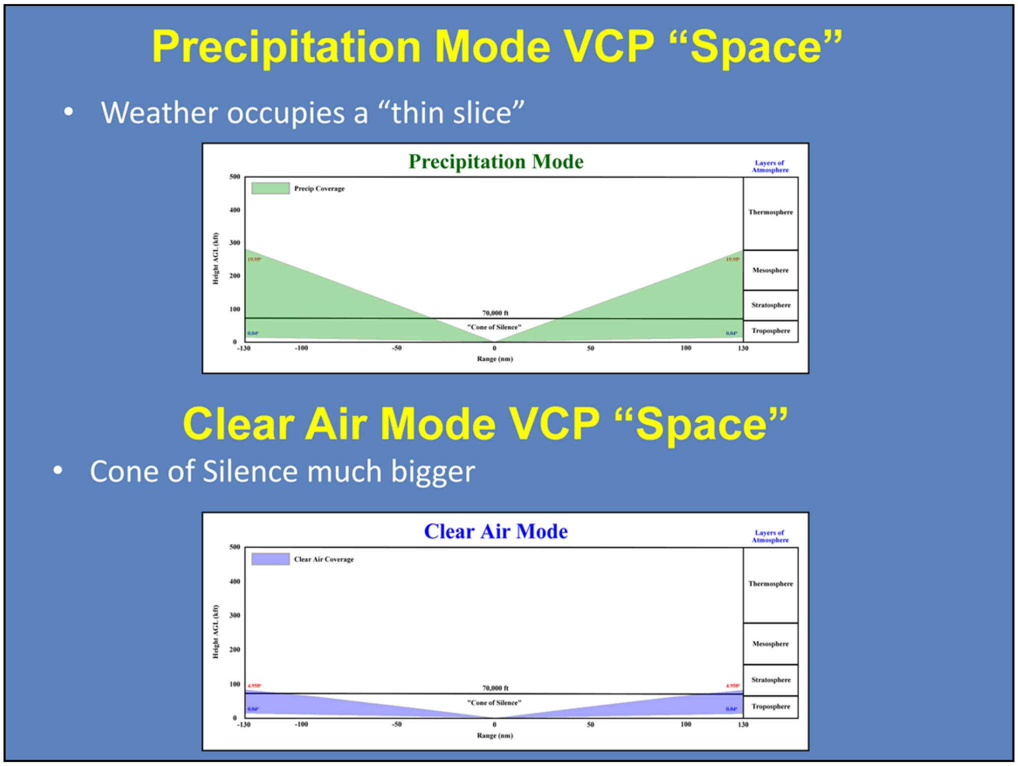

Here's a look at the both modes separately. The cone of silence is always something to remember and hopefully you have sufficient adjacent radars to mitigate some of the data loss. Compare the size of the cone of silence for Clear Air mode. By design, it covers much less atmosphere because nothing is 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.

Back to Top

Specific VCP Spaces

- Launch animation from each slide
- Vertical sampling of "typical" supercell (Precipitation Mode VCPs) or boundary (Clear Air Mode VCPs)

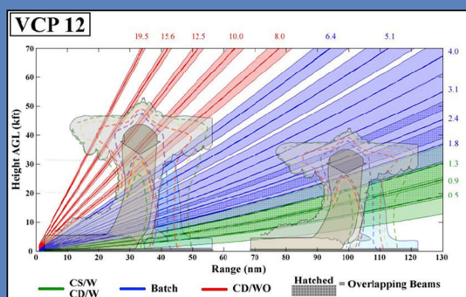

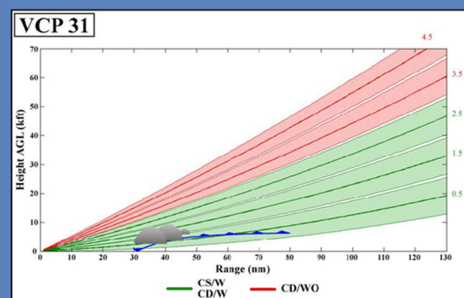

Now for a series of animations of the spaces that each VCP occupies. The following slides will each present a particular VCP, and you can launch the animation from that slide. For the Precipitation mode VCPs, you'll see a couple of "typical" (not too big, not too small) supercells, for you to see the 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.

Back to Top

VCP 31 Space

- Long pulse with highest sensitivity and uniform low level sampling
- Doppler PRFs not editable

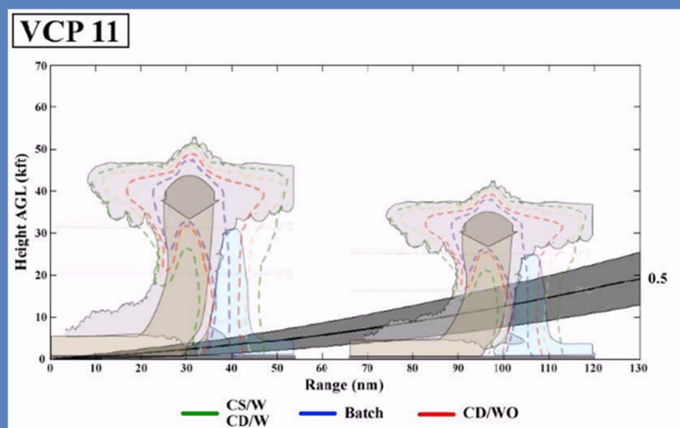

We'll start off with the Clear Air Mode VCPs. VCP 31 is the only VCP that uses long pulse, offering the highest sensitivity with uniform low level sampling. The Doppler PRFs are held constant and are not editable.

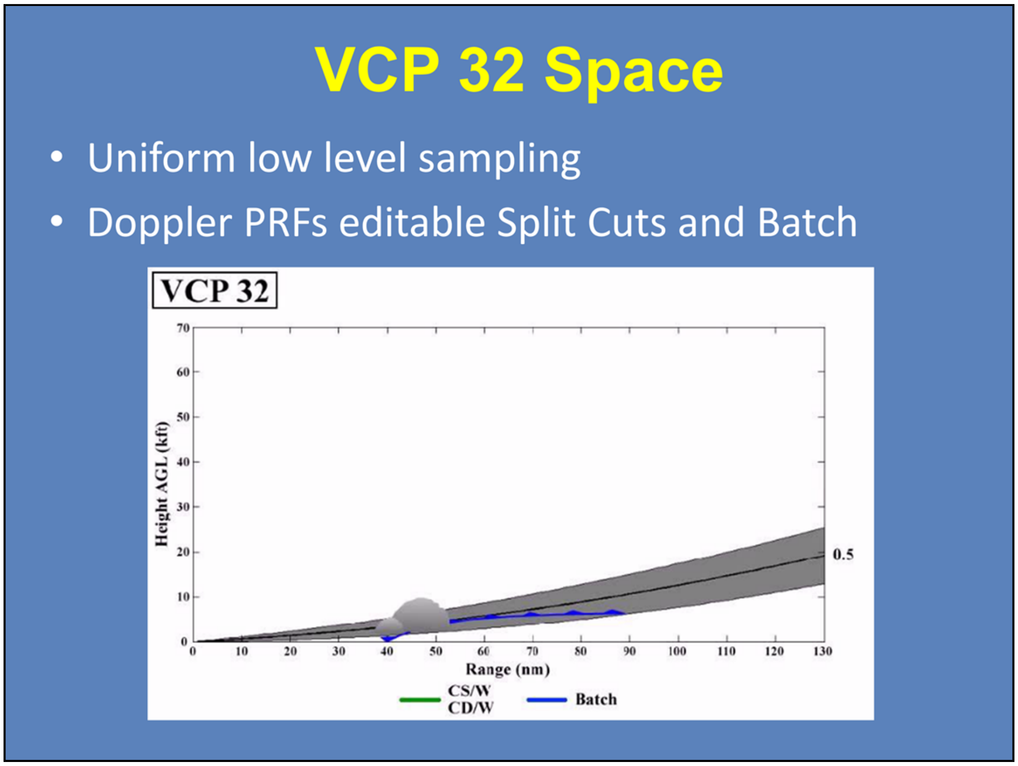

VCP 32 is the short pulse Clear Air mode VCP, with uniform low level sampling. The Doppler PRF is editable for the Split Cut and the Batch elevations.

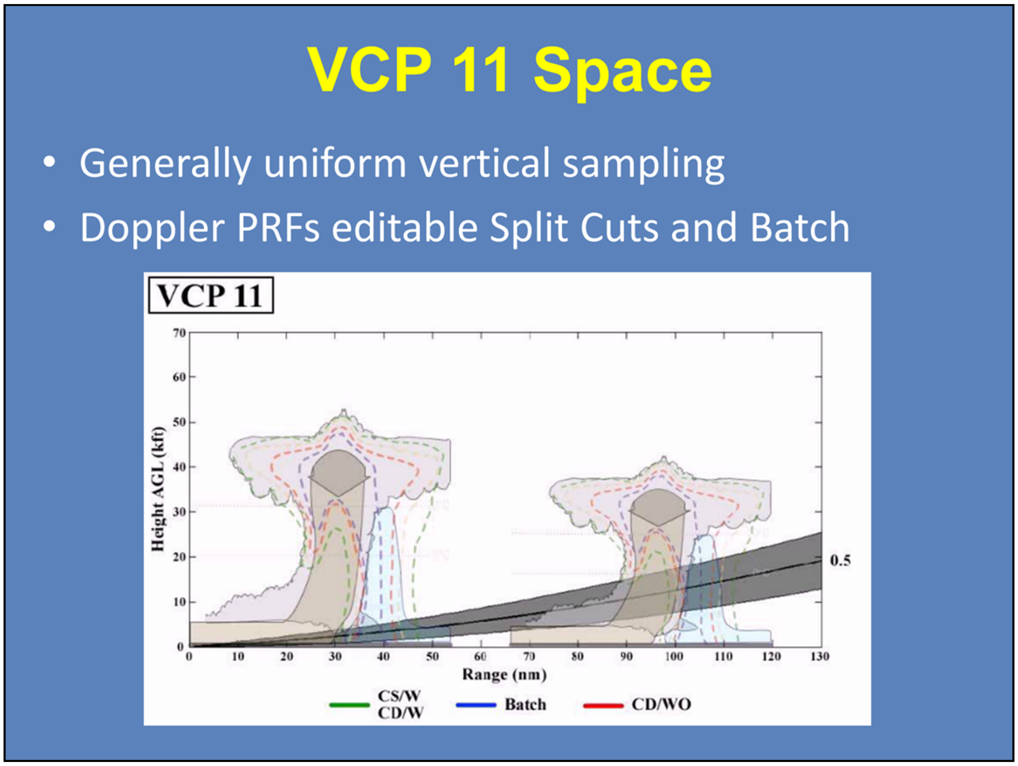

Now to the Precip Mode VCPs. We'll start off with VCP 11, which was actually the original convection VCP with relatively uniform vertical sampling. The Doppler PRF is editable for the Split Cut and the Batch elevations.

VCP 21 Space

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

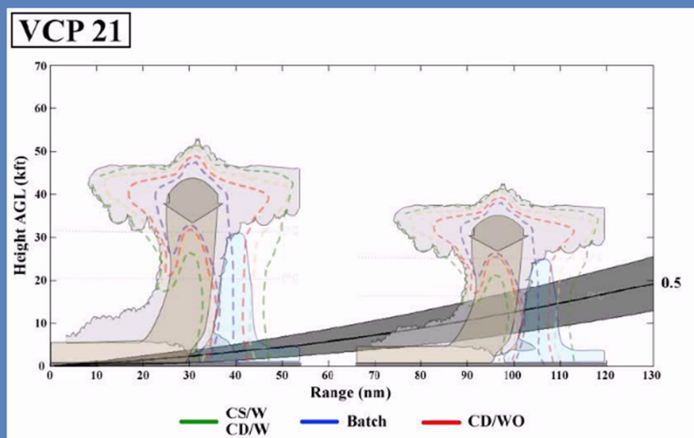

The widely used and widely defaulted, VCP 21 was the original stratiform precipitation VCP with uniform low level sampling, but significant gaps aloft. The Doppler PRF is editable for the Split Cut and the Batch elevations.

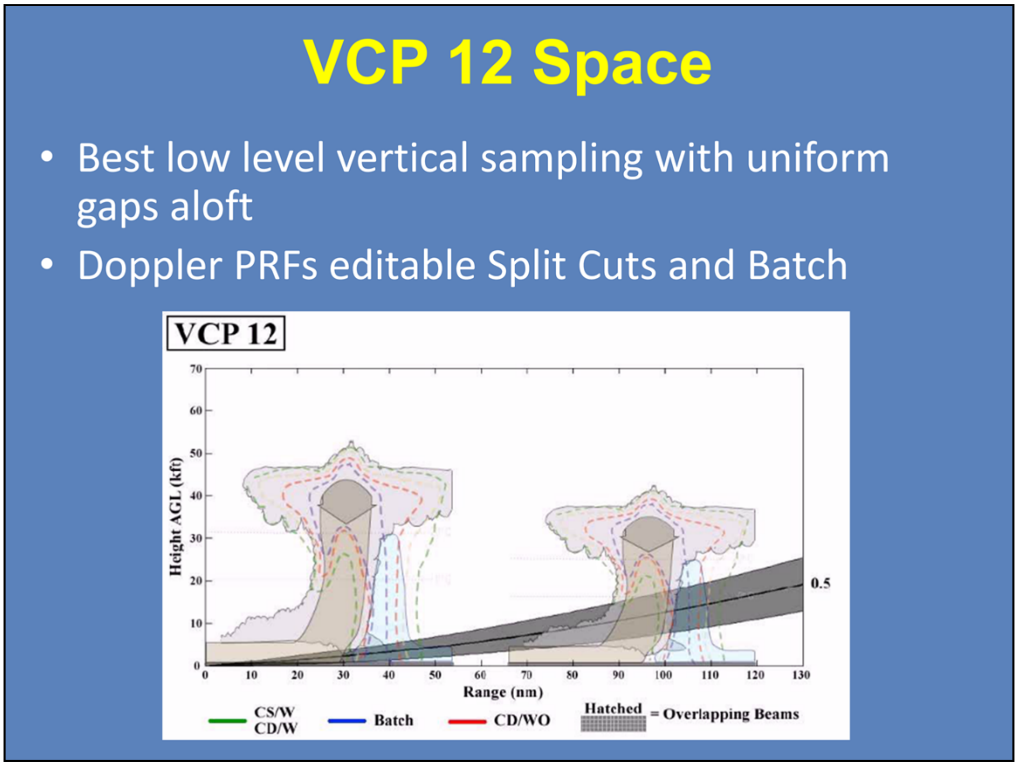

VCP 12 has the best low level vertical sampling with uniform gaps aloft. The Doppler PRF is editable for the Split Cut and the Batch elevations.

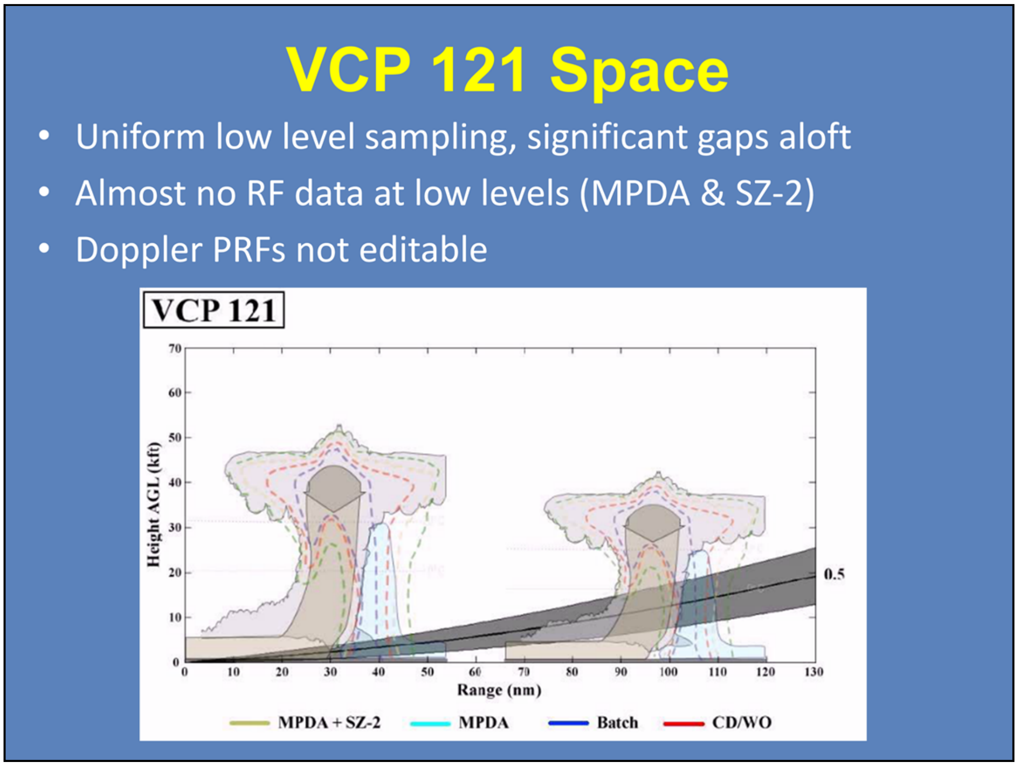

VCP 121 has uniform low level sampling with significant gaps aloft (just like VCP 21!). It's strength lies in the processing of the lower elevations, combining the SZ-2 and MPDA techniques. The result is almost no RF data for the lowest two elevation angles. The Doppler PRFs are held constant and are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

Back to Top

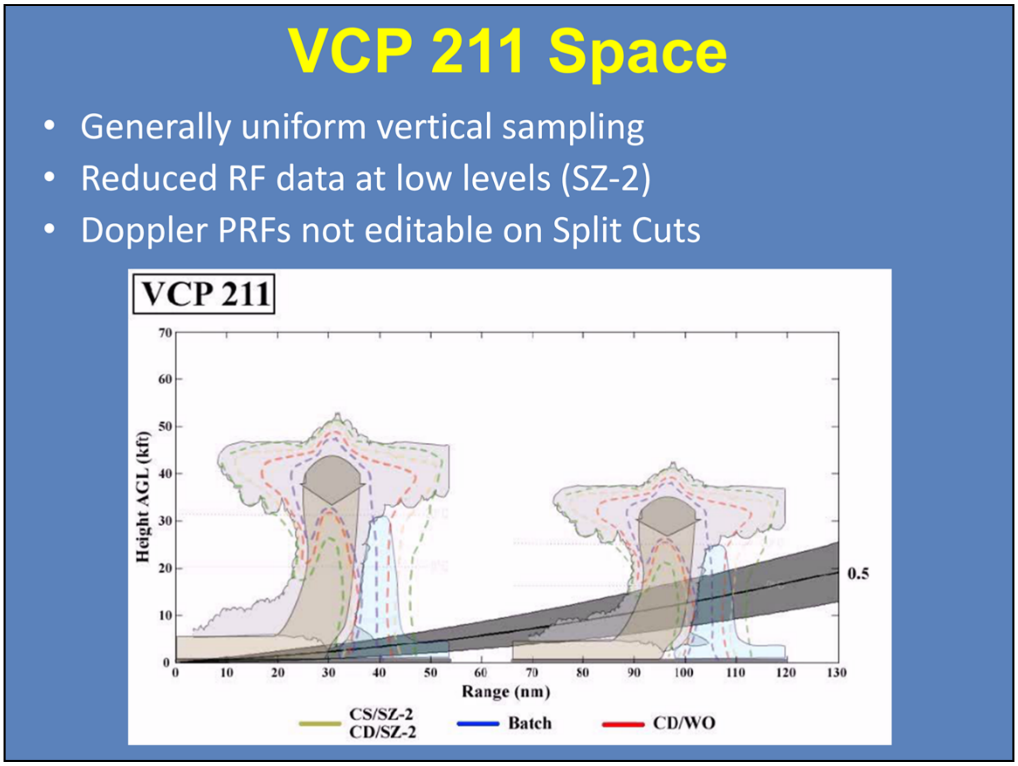

VCP 211 has relatively uniform vertical sampling, with reduced RF data due to SZ-2 processing on the lowest 2 elevations. The Doppler PRFs are not editable on the Split Cuts, but can be edited on the Batch elevations.

Back to Top

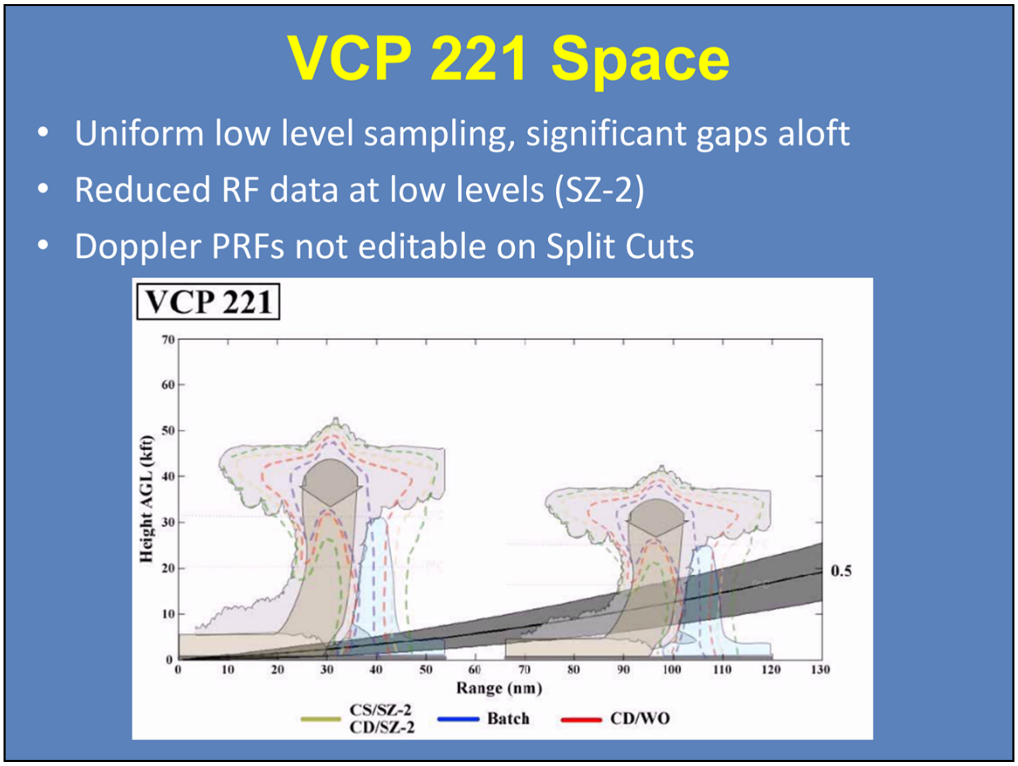

Like VCP 21, VCP 221 has uniform low level sampling with significant gaps aloft. Due to SZ-2 algorithm processing on the lowest 2 elevations, RF data is reduced. The Doppler PRFs are not editable on the Split Cuts, but can be edited on the Batch elevations.

Back to Top

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

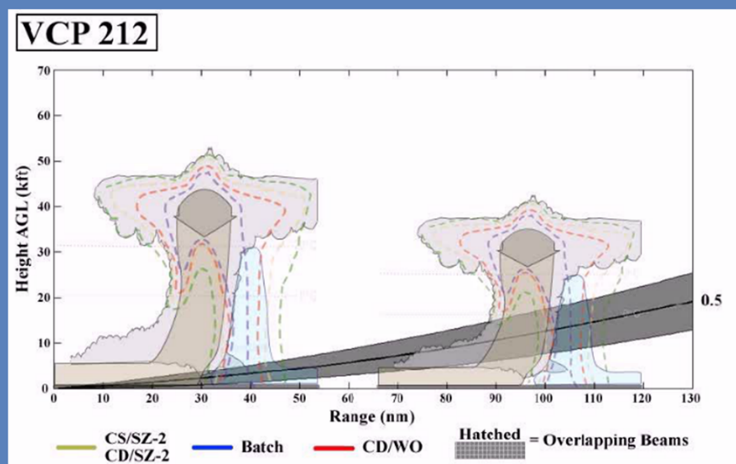

VCP 212 has the best low level vertical sampling with uniform gaps aloft. Due to SZ-2 processing on the lowest 3 elevations, RF data is reduced. The Doppler PRFs are not editable on the Split Cuts, but can be edited on the Batch elevations.

Back to Top

VCP Usage:
Radar Operations Center survey

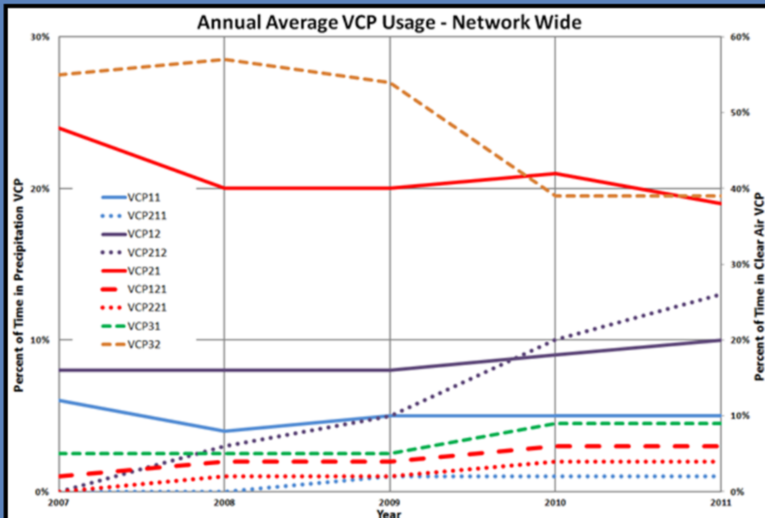

Figure 1: Annual VCP Usage (2007-2011).

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.

Back to Top

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.

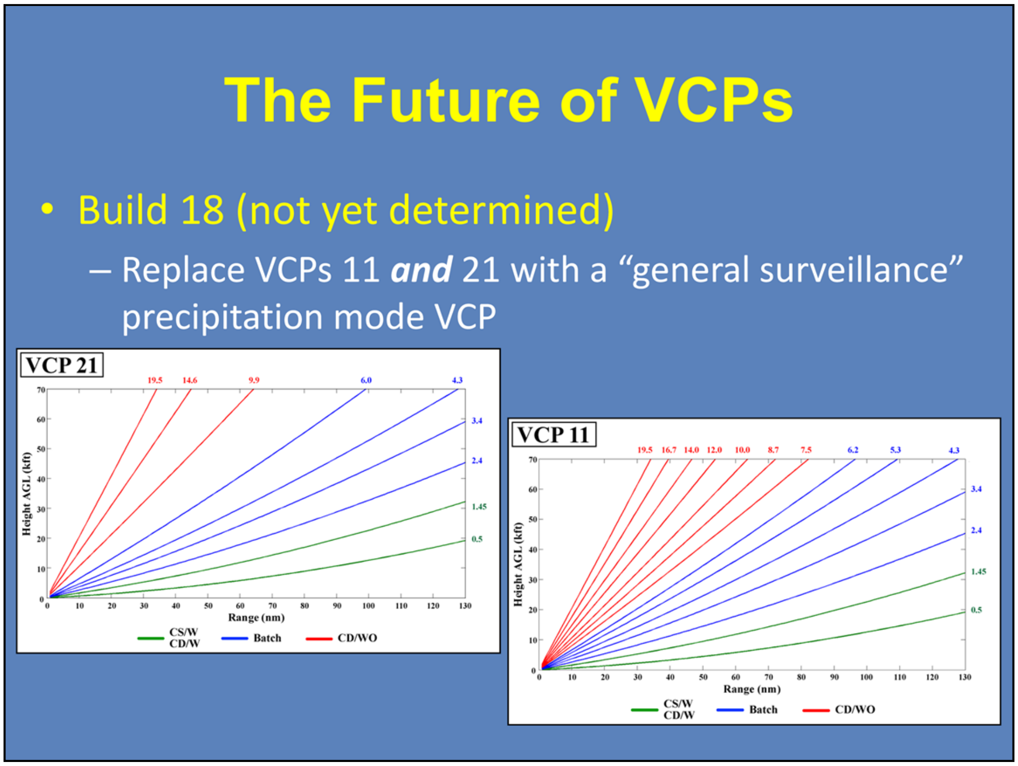

As of this recording, it is already in the works to combine VCPs 11 and 21, creating a new general surveillance VCP. This combination takes the uniform angle spacing of VCP 11 and the slow antenna speeds of VCP 21 to generate a VCP with better vertical sampling and data quality for non-severe precipitation.

Back to Top

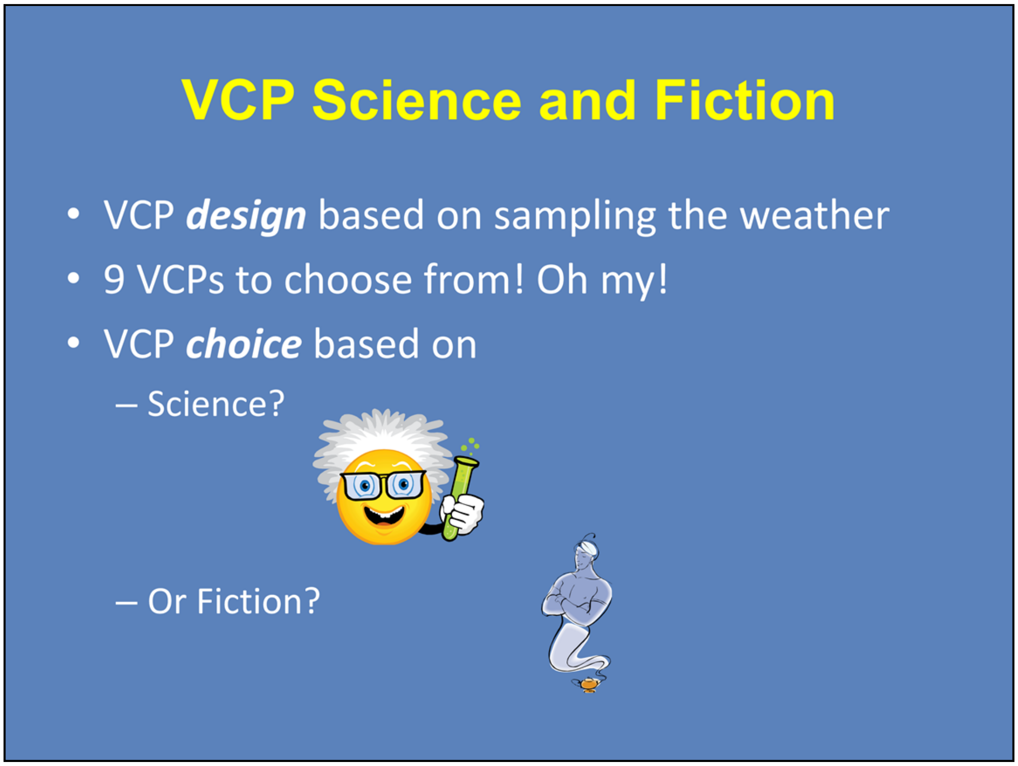

Given all of these VCP choices, we'll now examine some of the factors affecting VCP selection in real time. Each VCP was designed to sample the atmosphere in a different way. What we want to address here is whether VCP choices are based on science or fiction.

Back to Top

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.

Back to Top

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.

Back to Top

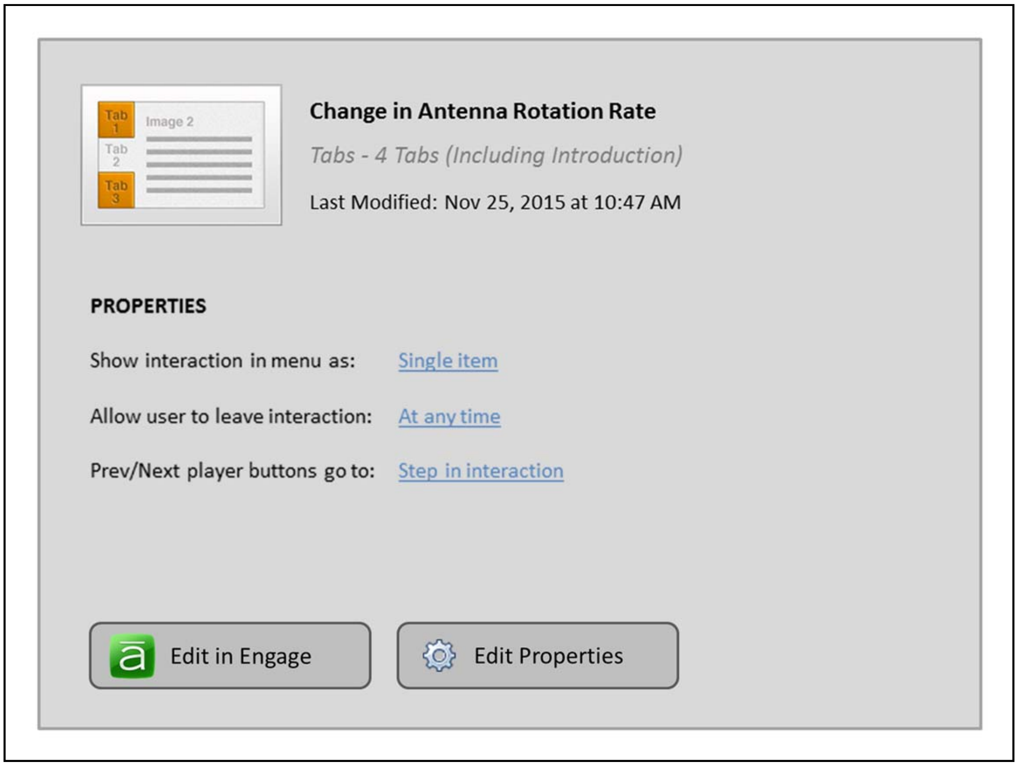

This short interaction will help you to explore the number of times that the antenna rotation rate changes significantly among the three fastest VCPs. You may be surprised to discover which one has greatest number of changes in antenna rotation.

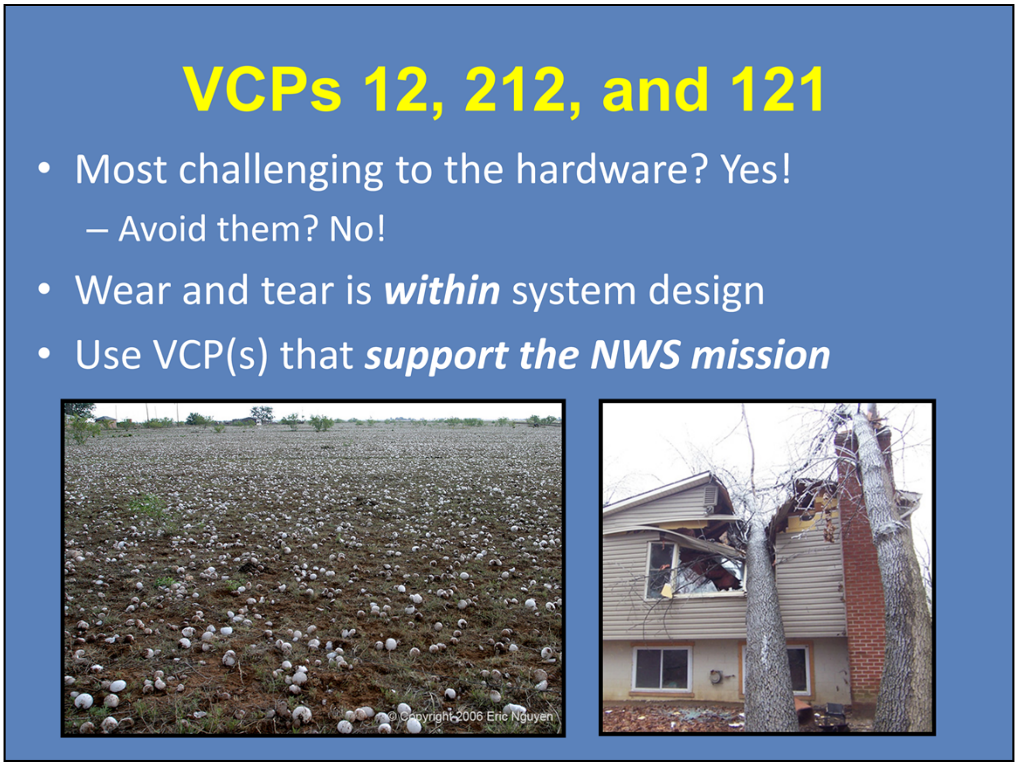

It is true that VCPs 12, 212 and 121 are the most stressful to the hardware. That does not mean that they should be avoided. Any associated wear and tear is within system design, and this will be explored in subsequent slides. Given thoughtful overall VCP usage, at appropriate times, VCPs 12, 212 and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

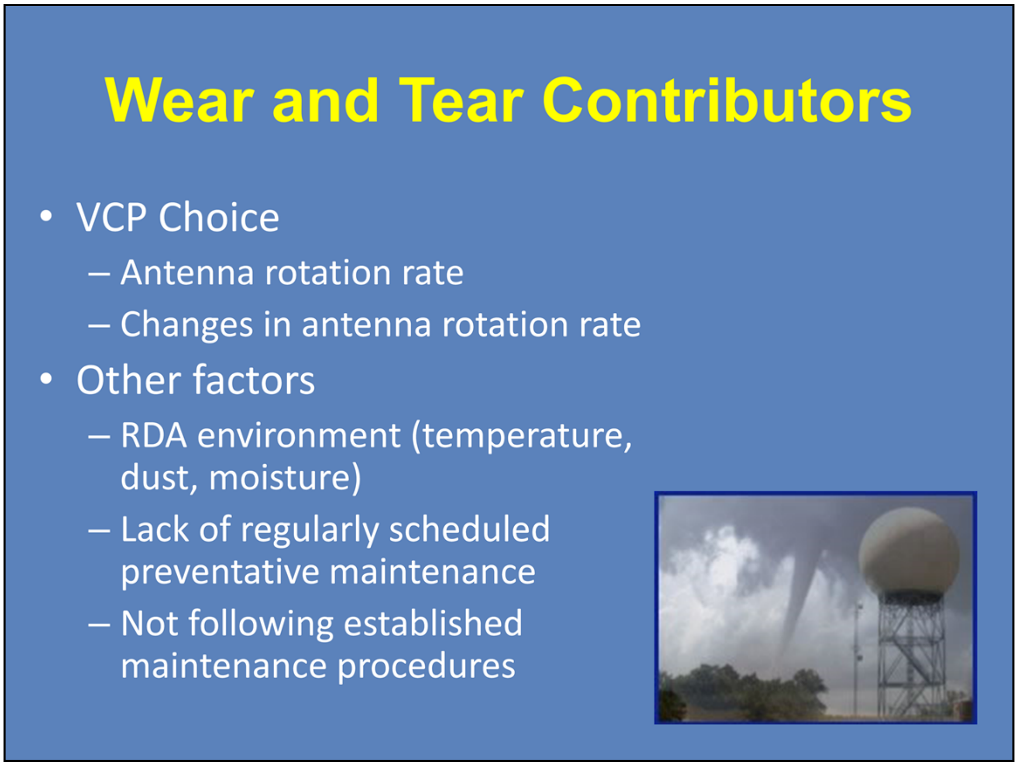

VCP choice does contribute to wear and tear, specifically fast antenna rotation rates and frequent changes in antenna rotation rates. There are other, equally important factors. Some RDAs are located in pretty tough environments with respect to extremes in temperature, dust, or moisture. The greater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

Back to Top

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.

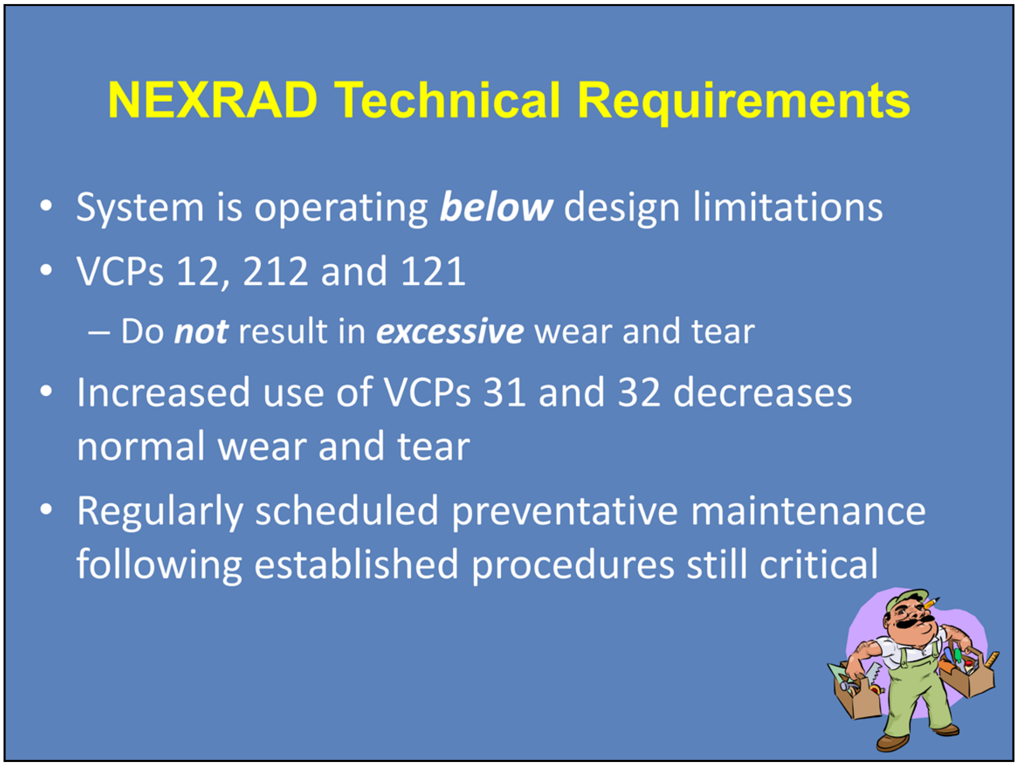

Given the NEXRAD Technical Requirements, even the fastest VCPs are operating well within the system design. VCPs 12, 212 and 121 do not result in excessive wear and tear. Also, the Mode Selection Function has resulted in an increased use of VCPs 31 and 32, which are the slowest, decreasing normal wear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

Back to Top

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!

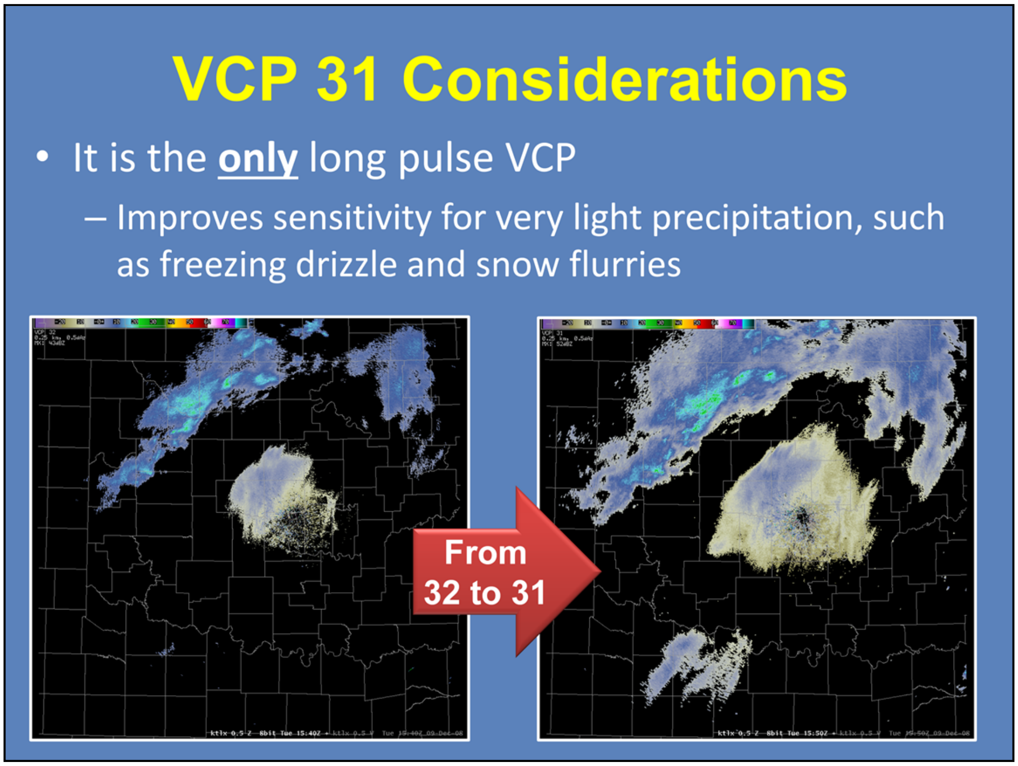

What is the operational advantage of VCP 31? Long pulse provides greater sensitivity, which means seeing returns from very light precipitation that you would not see with short pulse. In this freezing drizzle event, VCP 31 detects this very light, but high impact type of precipitation much more effectively 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.

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.

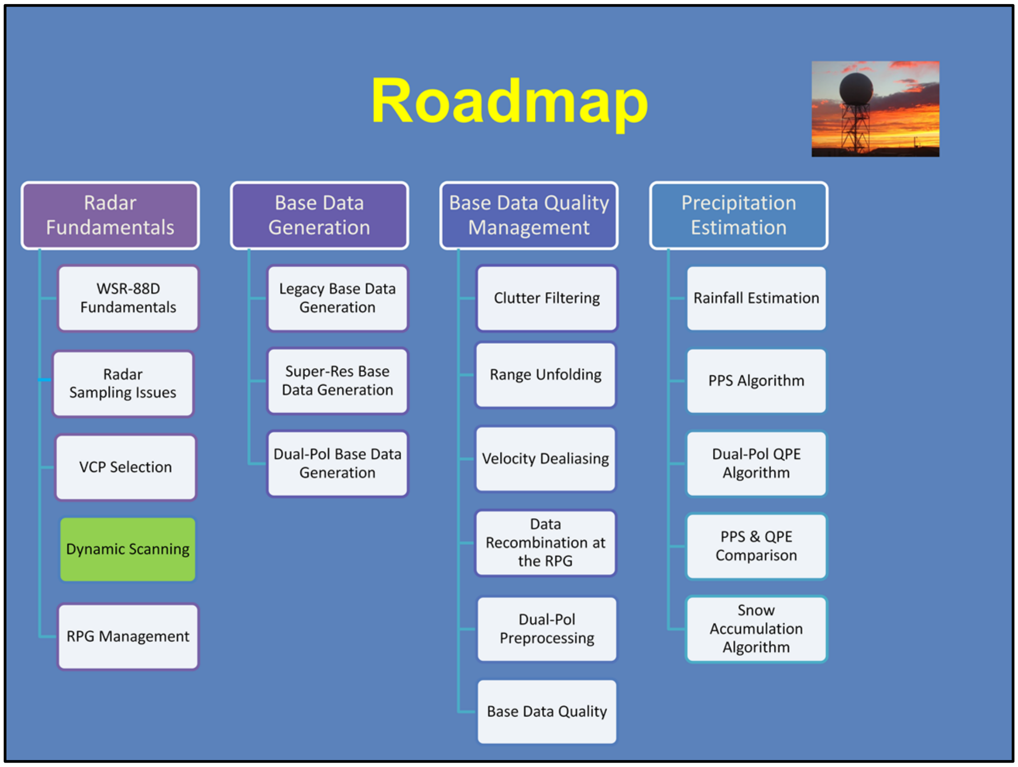

Here is the “roadmap” with your current location.

Back to Top

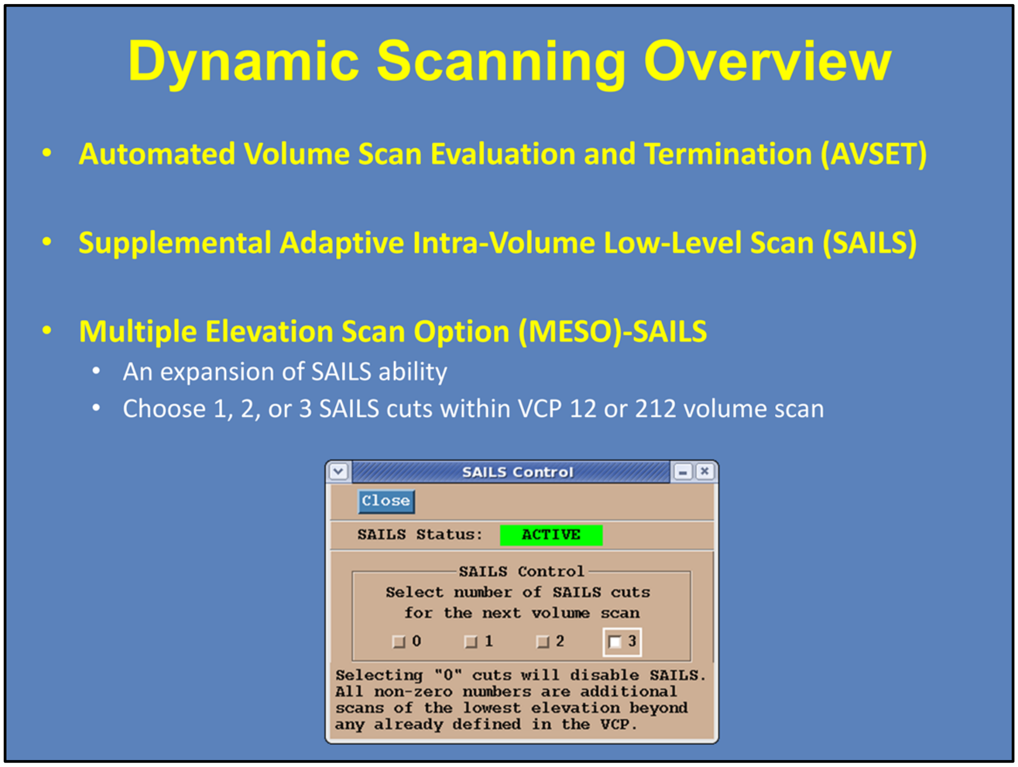

Before we define the lesson objectives, we will spell out the components of this Dynamic Scanning ability of the WSR-88Ds. The first application is called the Automated Volume Scan Evaluation and Termination, or AVSET. This allows any of the Precip mode VCPs to truncate the upper level tilts of a particular 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 0.5⁰ 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 0.5⁰ cuts in one volume. With three SAILS cuts, you get almost one-minute updates on the 0.5⁰ products.

Back to Top

Dynamic Scanning

Learning Objectives

1. Identify the purpose of AVSET, and which VCPs utilize it.
2. Identify how AVSET's design mitigates potentially missed detection of developing elevated convection.
3. Identify the purpose of MESO-SAILS, and which VCPs utilize it.
4. Identify why MESO-SAILS will not result in excessive hardware wear and tear.

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

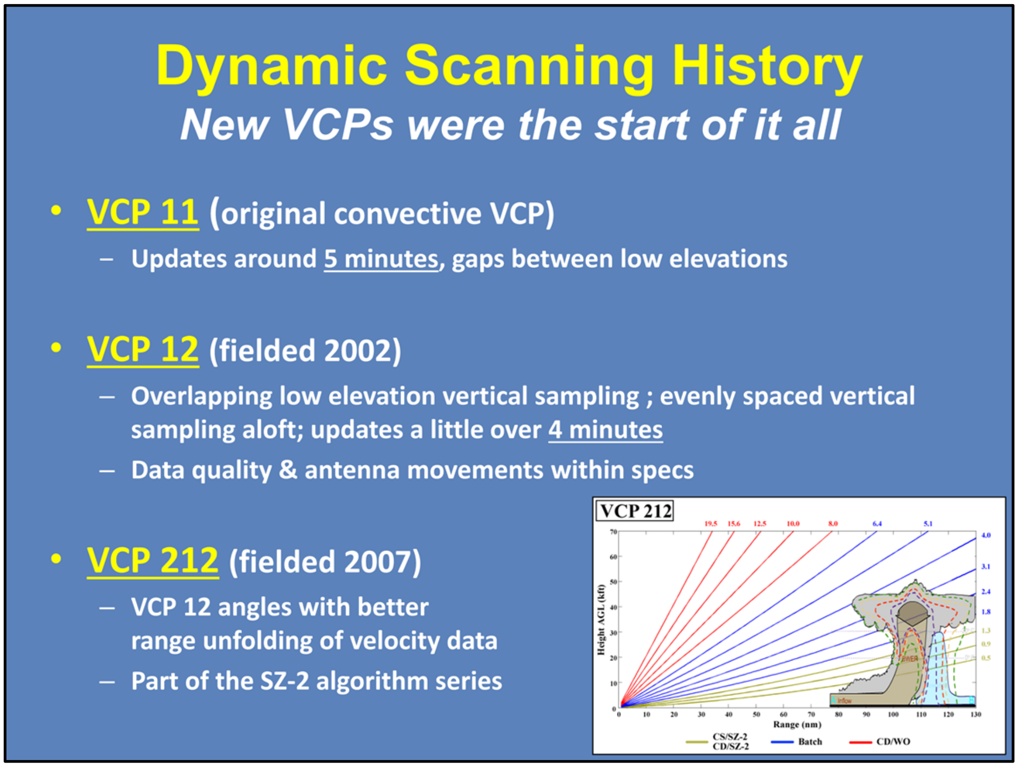

Let's start off with a short background of how we even got to this era of Dynamic Scanning. Some of the VCPs currently in use did not exist from the beginning. They were developed over time and the Radar Operations Center folks looked for more effective ways to use the radars. VCP 11 was the original 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.

Back to Top

Dynamic Scanning History

The Need for Faster Updates at Low Levels

- **AVSET** (fielded 2011)
  - Precip Mode VCPs terminate early based on depth of weather
  - Low level products update more often

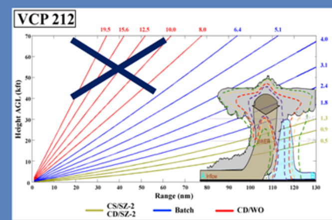

- **SAILS** (fielded 2014)
  - Antenna rescans 0.5° during middle of VCP
  - Low level products update twice as fast

- **MESO-SAILS** (fielded 2016)
  - Selectable number of SAILS scans, up to 3
  - With AVSET, low level products update rate almost 1 minute

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.

Back to Top

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)

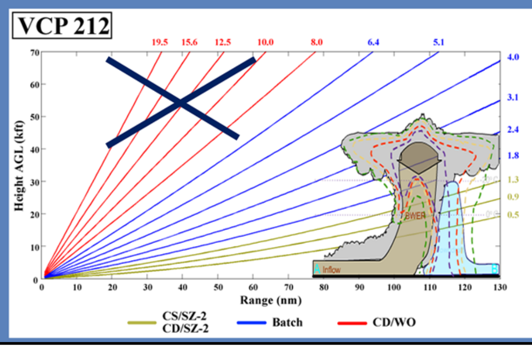

As we learned in the history of these applications, there has long been a desire for faster VCP updates. AVSET can meet this need when the weather return is limited to the lower elevation angles, though it is dependent on the depth of any detected precipitation and its range from the radar. This image 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.

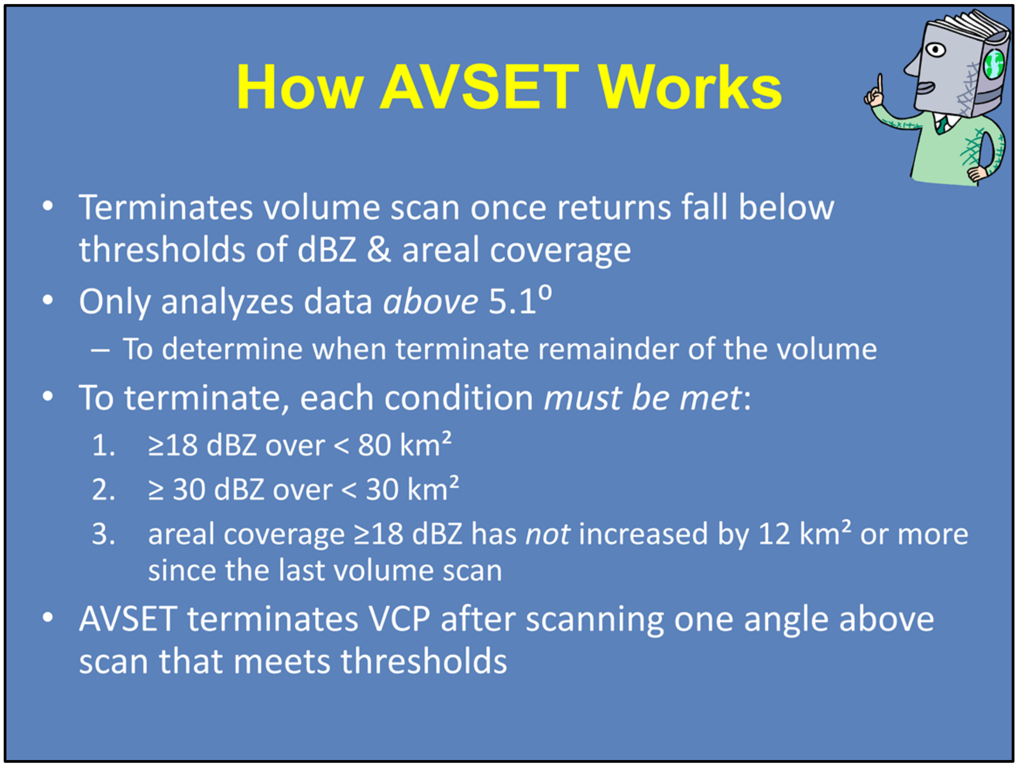

How does this process work? Once the radar gets above 5.1^{o}, it starts to calculate whether the sampling of precipitation meets certain dBZ and areal coverage thresholds. AVSET checks the ***total*** areal coverage of returns above both 18 and 30 dBZ. To terminate the volume scan, there are three conditions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

The coverage at or above 18 dBZ must be less than 80 km^2.
The coverage at or above 30 dBZ must be less than 30 km^2.
The coverage at or above 18 dBZ has not increased by 12 km^2 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.

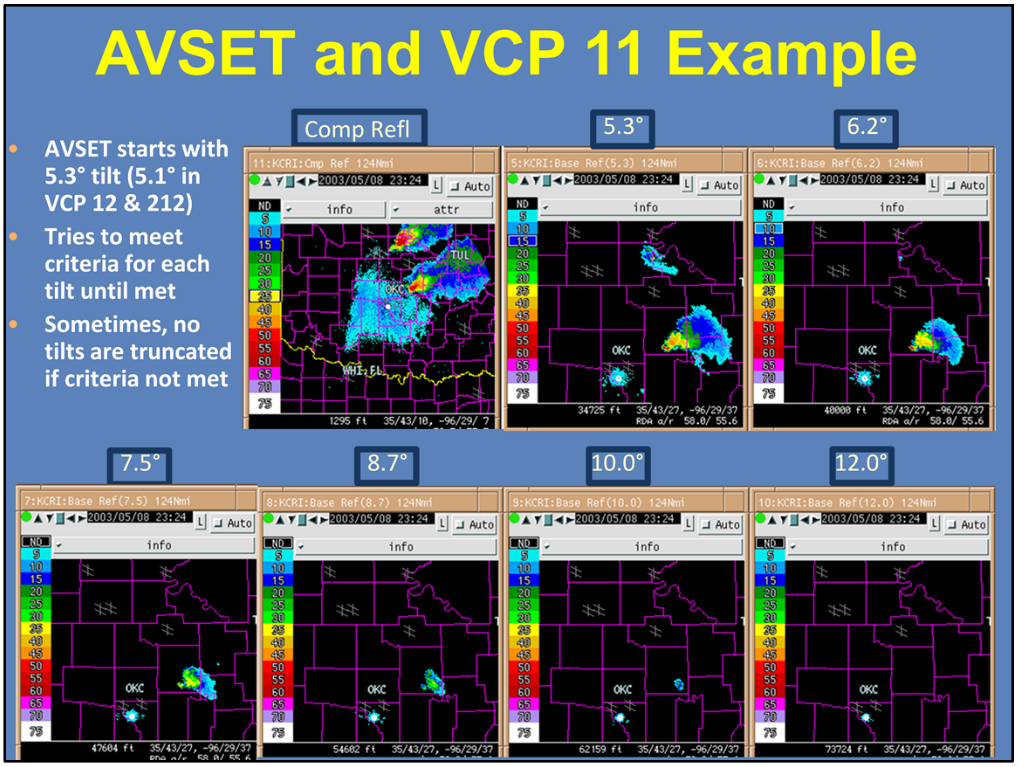

Here is a quick example of how AVSET works with a couple of storms on the display. You can see in the Composite Reflectivity, the first image on the top-left, there are a couple of storms to the north and east of the radar. Starting with the 5.3° tilt, which is 5.1° degrees in VCPs 12 and 212, AVSET 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.

Back to Top

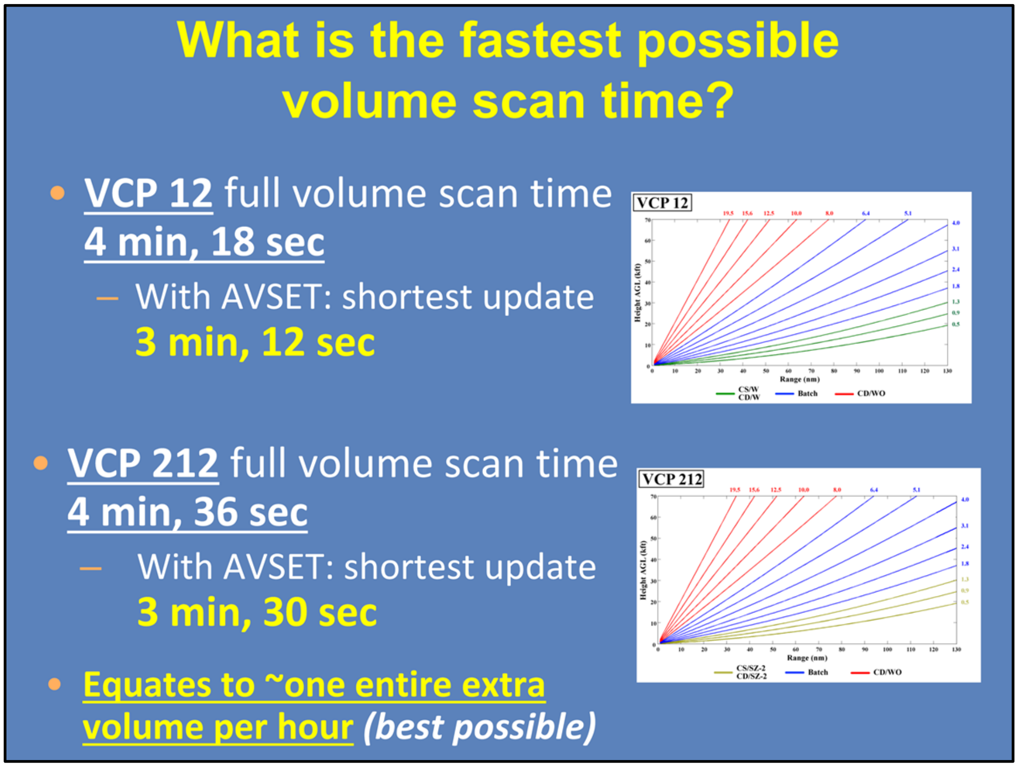

So, after all this talk of speeding up VCPs, what ***IS*** the fastest possible volume scan time with AVSET enabled? Here, we use the two fastest VCPs, 12 and 212. With AVSET disabled and performing a full volume scan, the times are 4 and 4.5 minutes, respectively. But with AVSET enabled and able to 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.

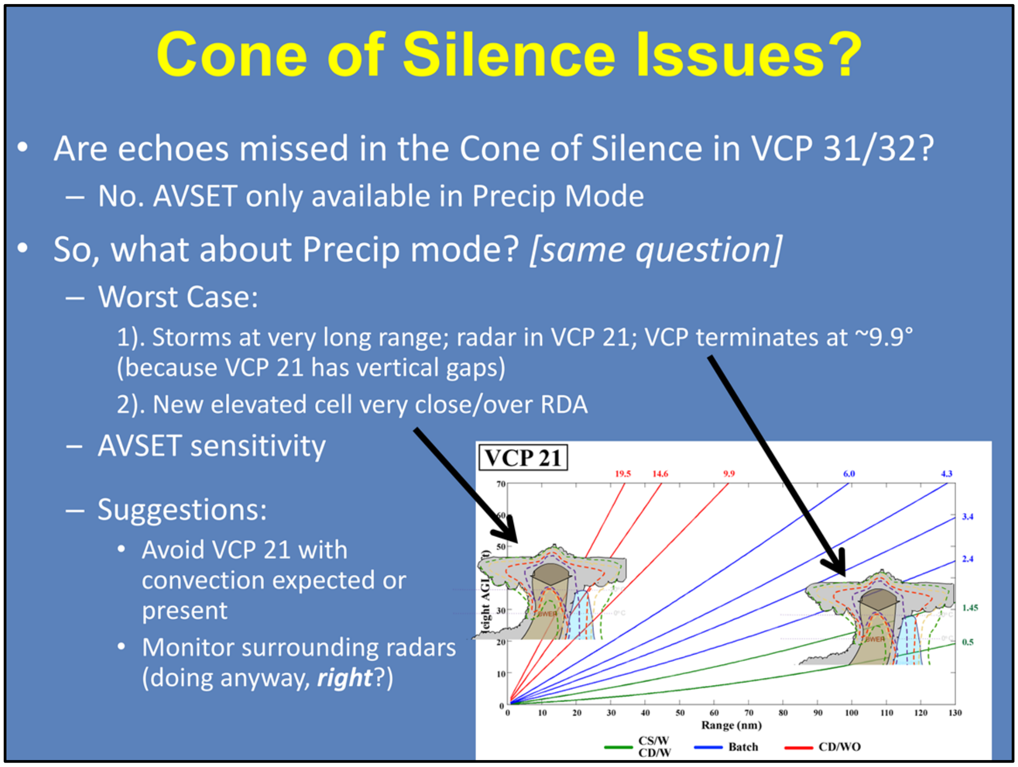

One more consideration for AVSET is the Cone of Silence. Since AVSET starts at 5.3^{o}, it has no affect on VCPs 31 and 32. There is no change in detection of elevated convection in Clear Air Mode.

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.

Back to Top

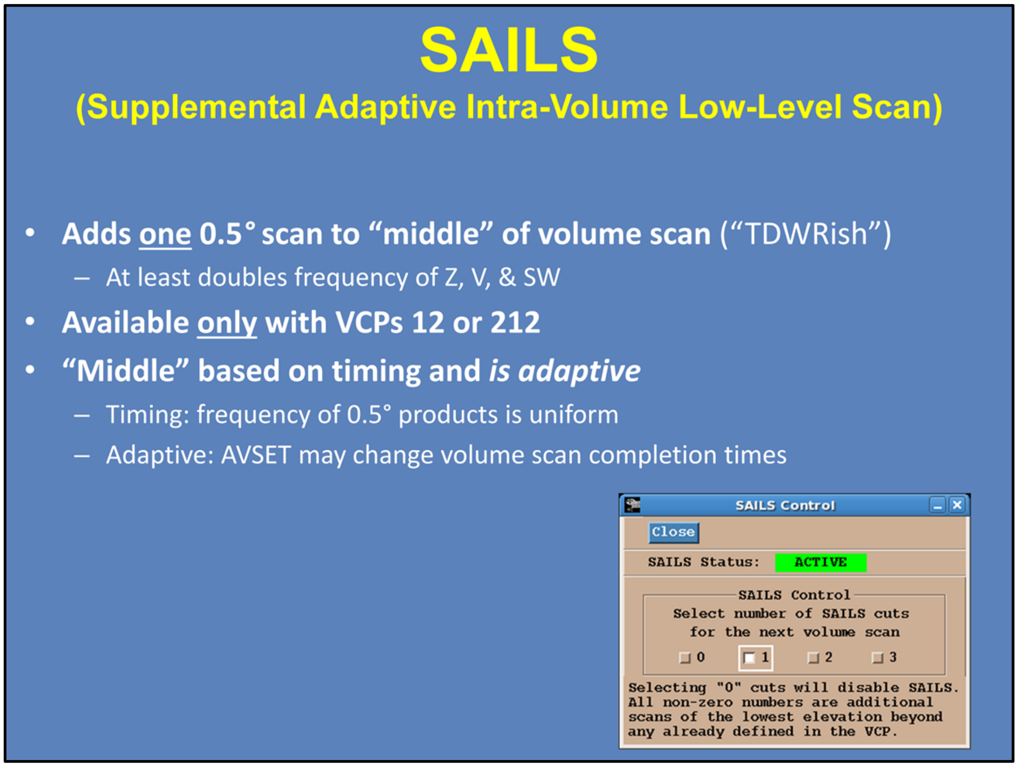

Next in our utility belt of Dynamic Scanning is SAILS, which stands for – Supplemental Adaptive Intra-Volume Low-Level Scan. The lowest degree tilt is by far the most widely used tilt, especially when scanning for low-level features that contribute to tornadogenesis. TDWRs are well known for their one-minute 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.

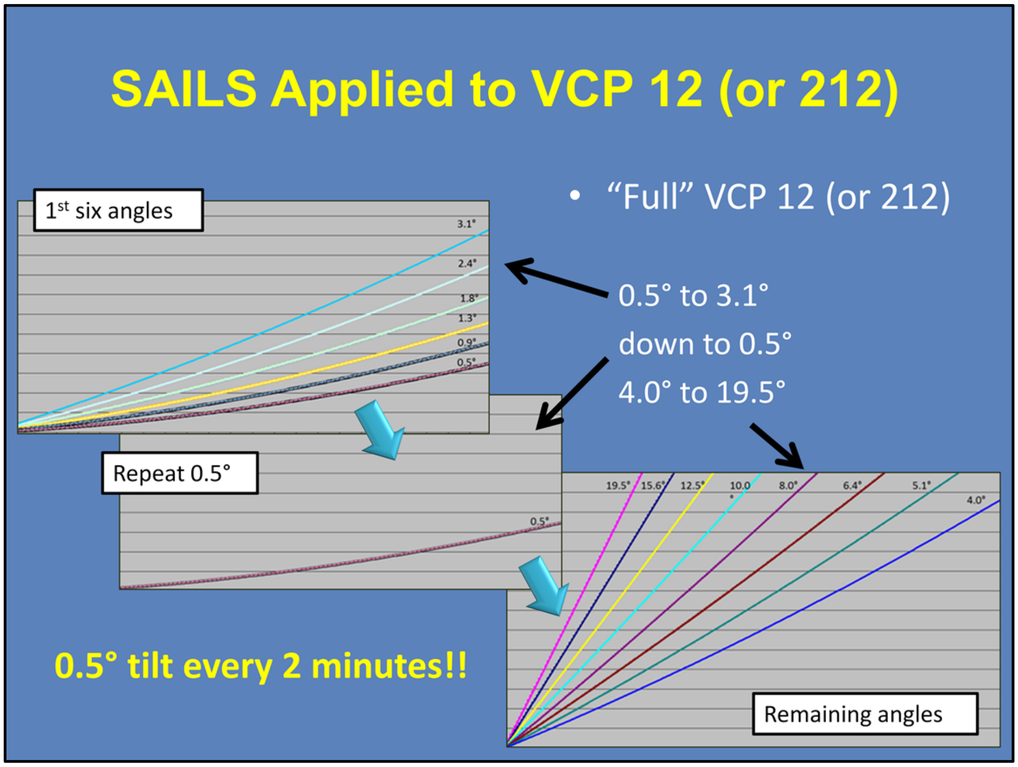

Here's an example of SAILS applied to VCP 12 – or 212, with all the angles sampled, so this would be either with AVSET disabled or there are enough storms that no upper tilts are truncated. The first six angles are samples as normal, elevations 0.5° through 3.1°. The antenna then drops down to sample 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.

SAILS and AVSET

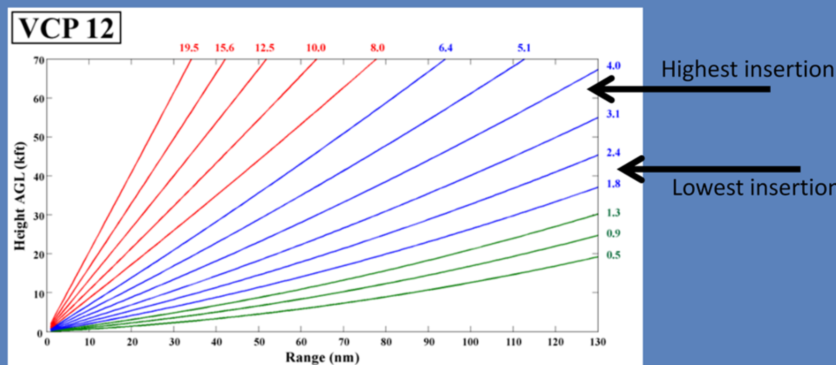

- Insertion dependent on AVSET's termination angle
  - Lowest is between 1.8° – 2.4°
  - Highest is between 3.1° – 4.0°
- Want SAILS cut products to update "timing middle" of VCP

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.

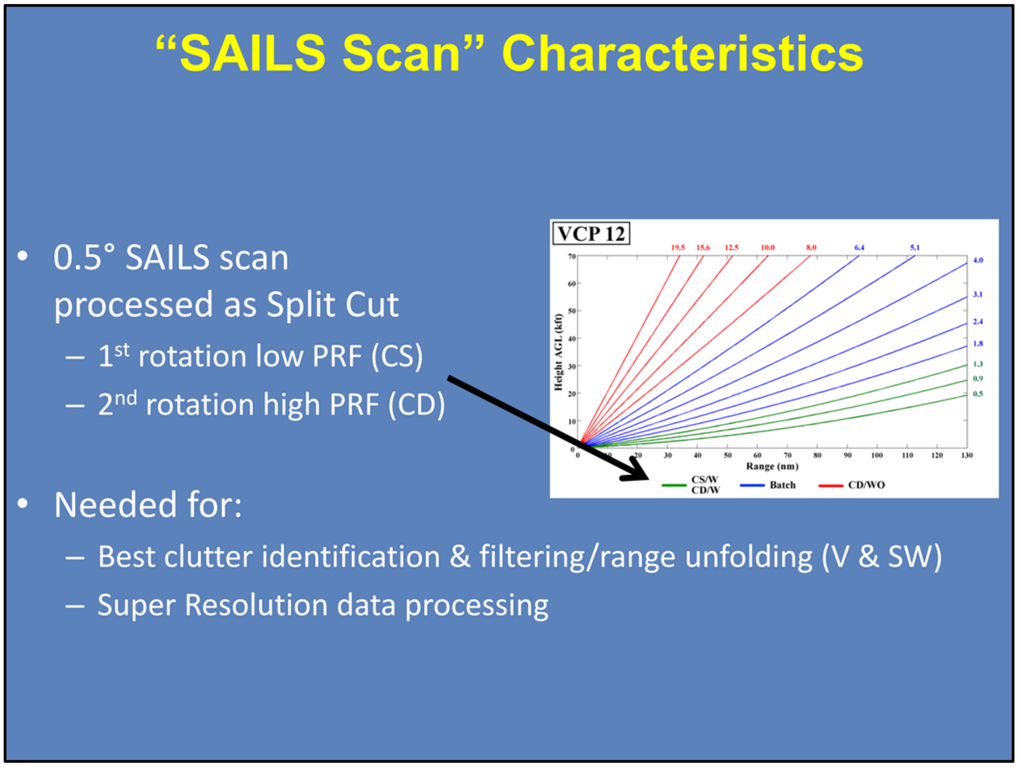

This extra 0.5° scan is known as the SAILS cut. It is processed as a Split Cut, which means the first rotation is at low PRF (Contiguous Surveillance), followed by a second rotation at high PRF (Contiguous Doppler). The Split Cut sampling is needed for the best data quality, better clutter identification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

Back to Top

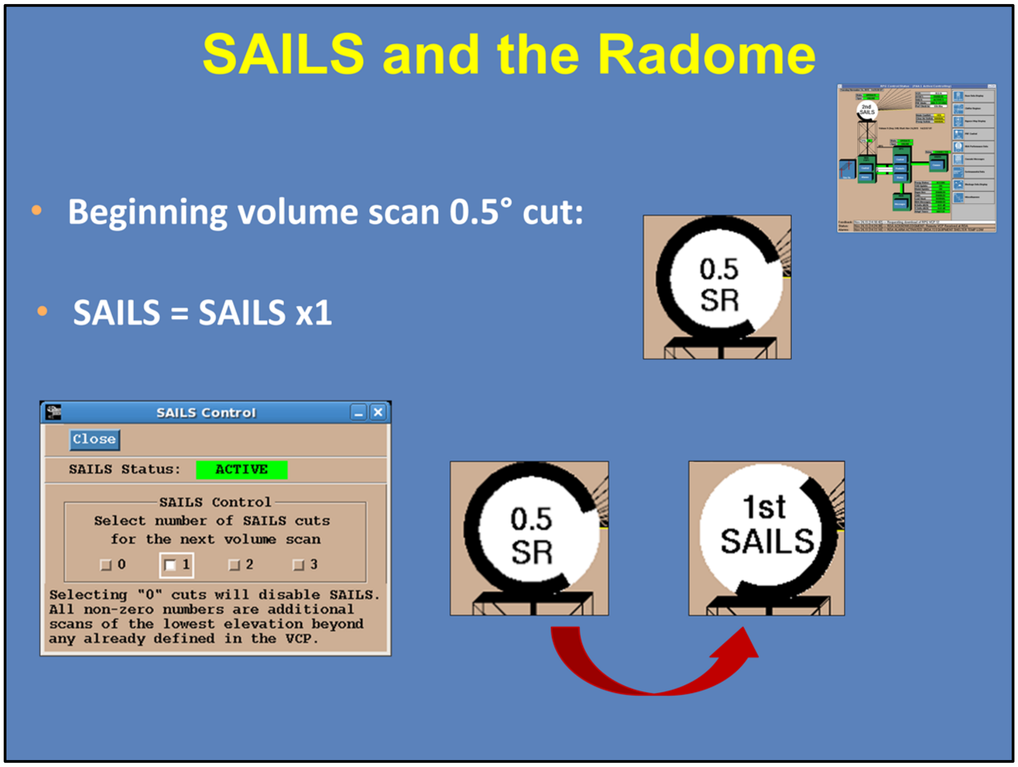

With both SAILS and AVSET active, there is a need to know when the SAILS cut is executed. To support this, there are two features on the radome on the RPG Control/Status Window. At the beginning of the volume scan, the radome displays 0.5° and SR. In this example, SAILS x1 has been selected, which is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

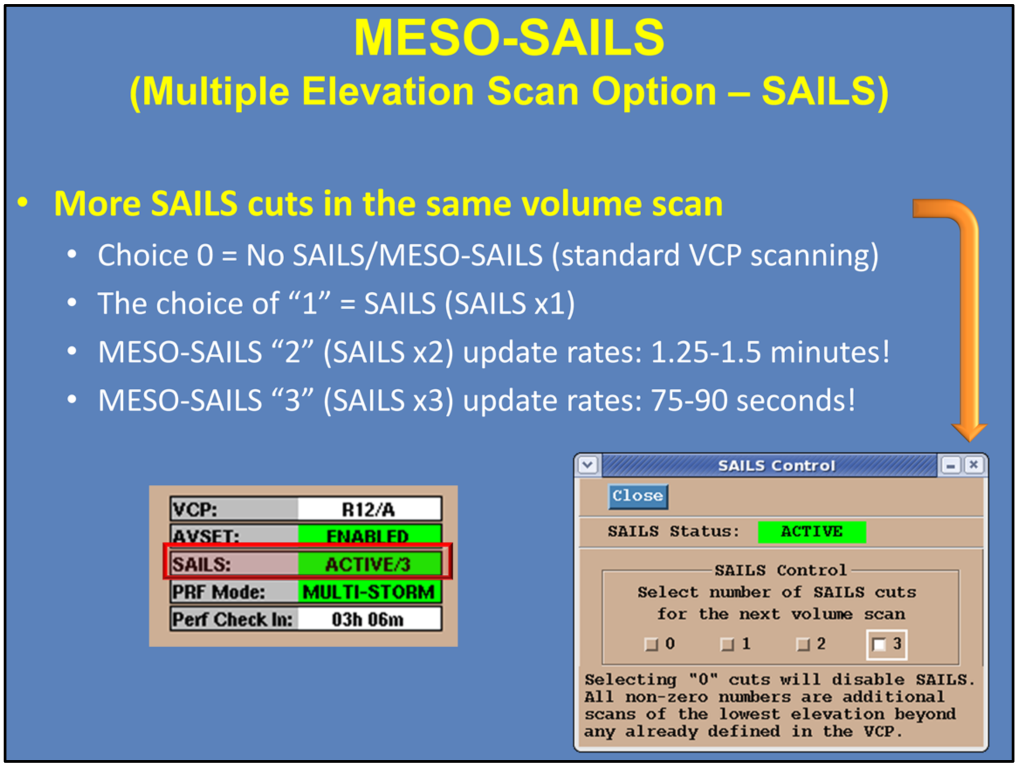

While SAILS adds one additional 0.5° cut in the "middle" of VCP 12 or 212, MESO-SAILS takes this two and even three steps further, giving you the option of having an additional 2 or 3 0.5° cuts. From the image in the lower-right, you can see the options for 0, 1, 2, or 3. Zero, means that SAILS and MESO-SAILS 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!

SAILS x2 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
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

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.

Back to Top

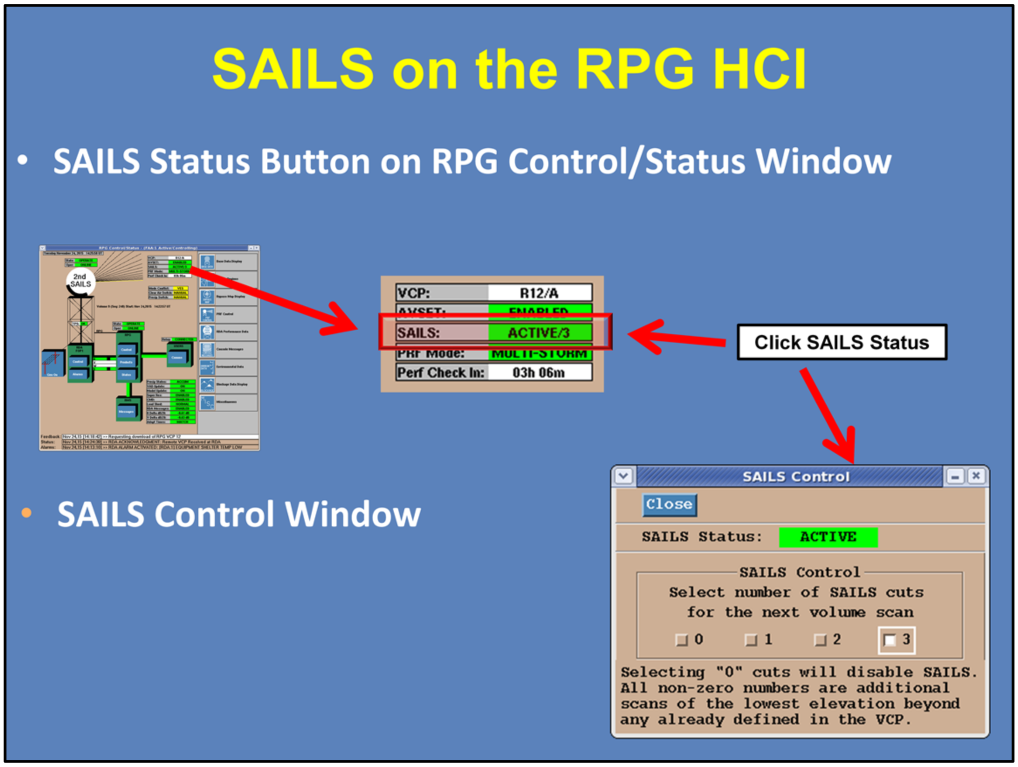

How do you invoke the number of SAILS cuts desired? Starting on the RPG Control/Status Window, the SAILS Status button tells you the current setting. In this example, it is SAILS x3, and since we're in VCP 12, SAILS is ACTIVE.

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.

Back to Top

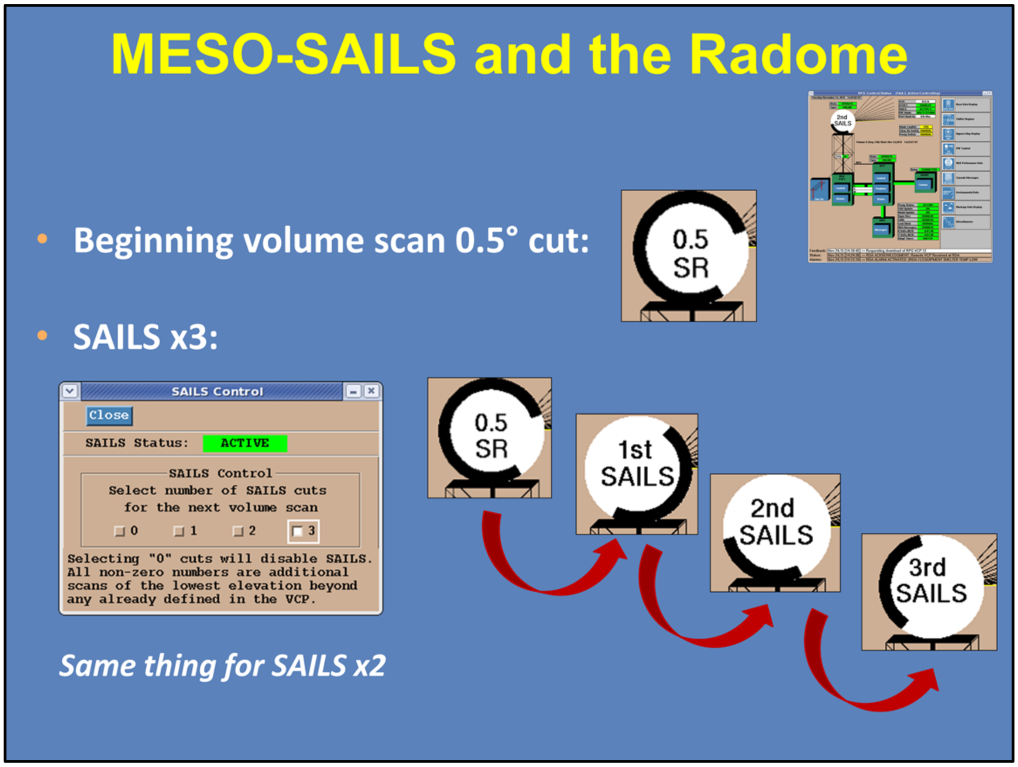

As with the SAILS x1 example, the radome tells you when the SAILS cuts are scanned, identified as 1st, 2nd, or 3rd, depending on the number of cuts selected. Here we have an example of each of the SAILS cuts on the radome when SAILS x3 has been selected.

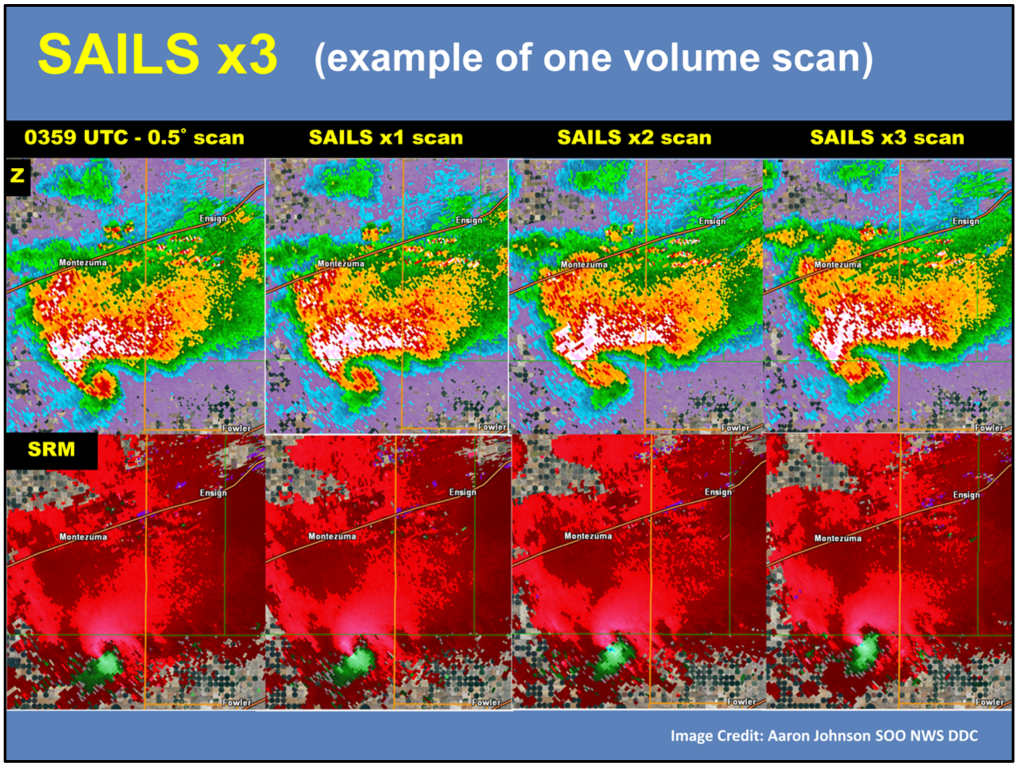

It is fairly intuitive what one additional 0.5° set of products would look like, updating at twice the "normal" rate. But what does almost one-minute 0.5° data look like and what could it mean in terms of watching the evolution of features?

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.

Back to Top

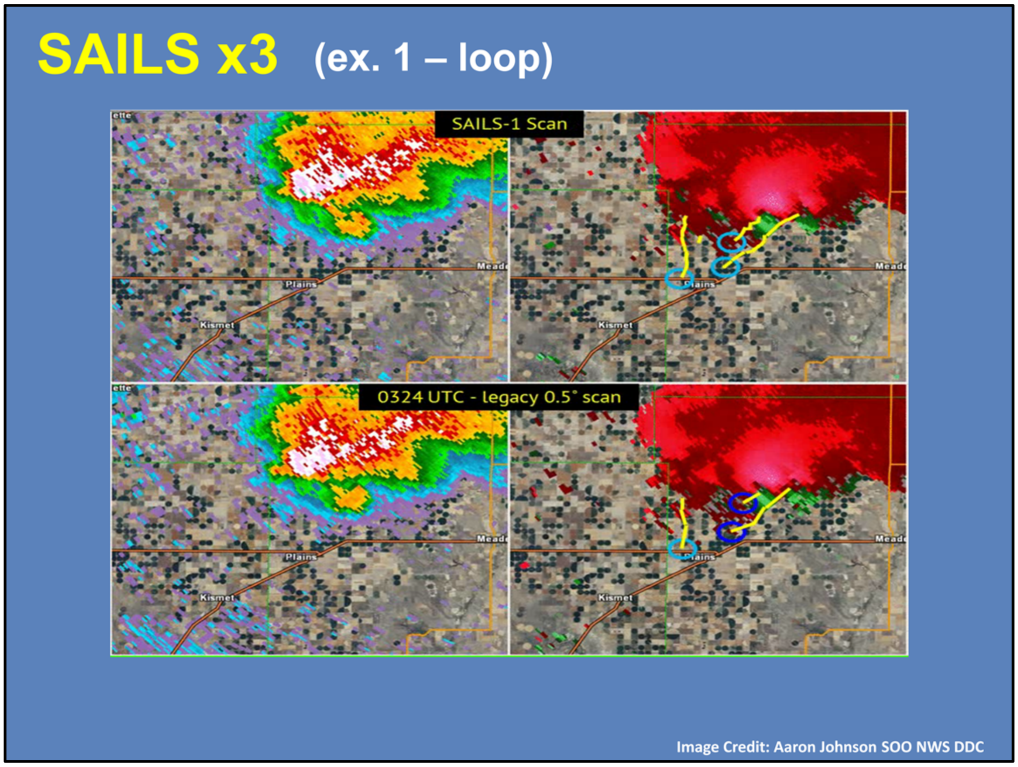

This is the first of 3 examples that are all short gif animations, that will keep replaying from start to finish. SAILS x3 is enabled, with Base Reflectivity on the left and SRM on the right. The top half of these loops are the SAILS x3 cuts and the bottom half show just the first 0.5° cut from each 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.

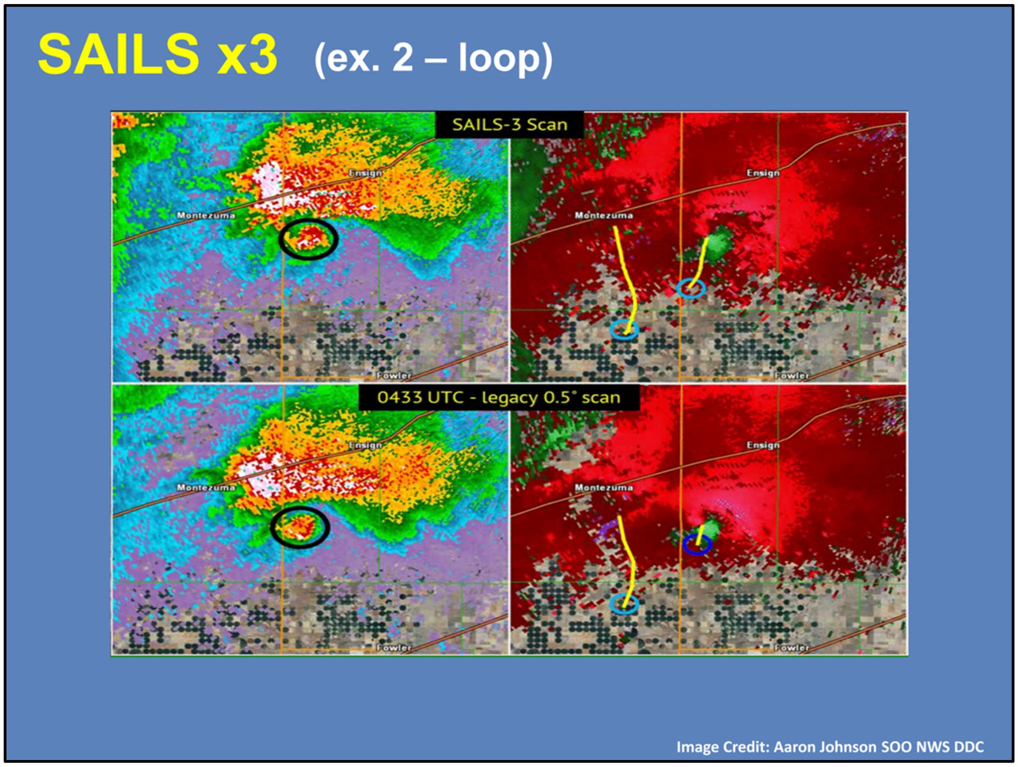

Now for the second example, again showing the tracks of multiple possible mesocyclones. Not only are we able to better track a more exact path of a potential tornado or at least the best area of rotational velocities, but we now have a bit more lead time on the precursors of these types of signals. As 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.

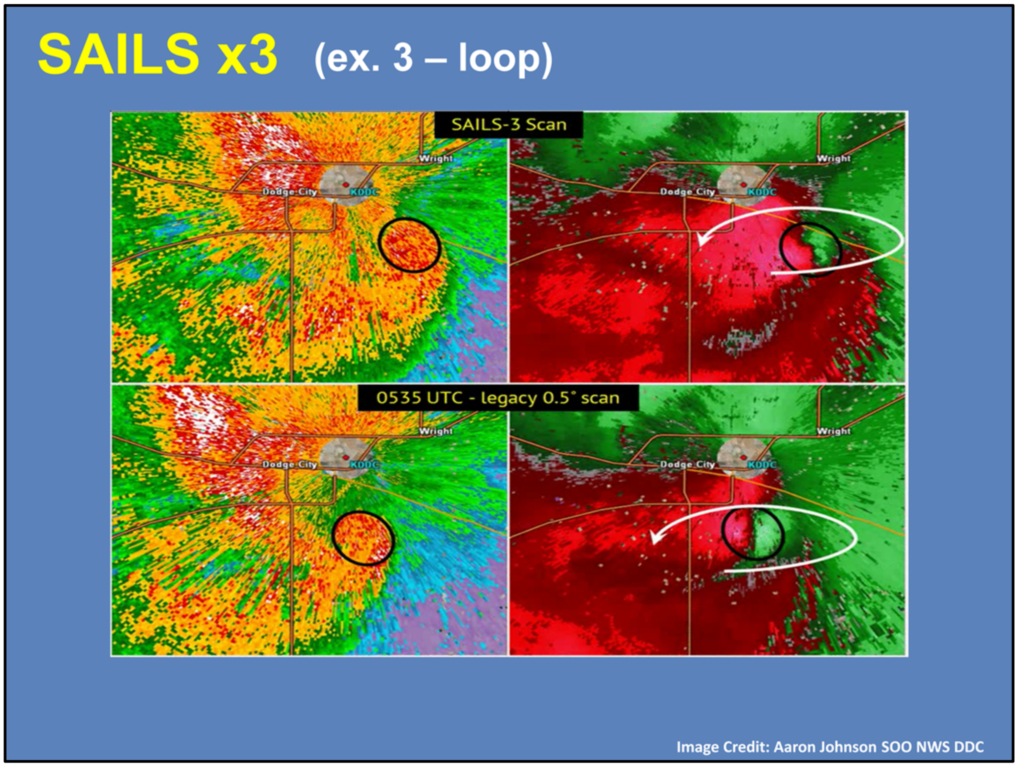

The last example is of an RFD surge that swept just south of Dodge City, KS – just a few miles away from the RDA. There are some brief and embedded couplets with this surge. Notice the difference in the evolution of this feature with SAILS x3 enabled and compare it to the legacy scan timing on the bottom. With a feature like this so close to the radar, you are probably <u>not</u> 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.

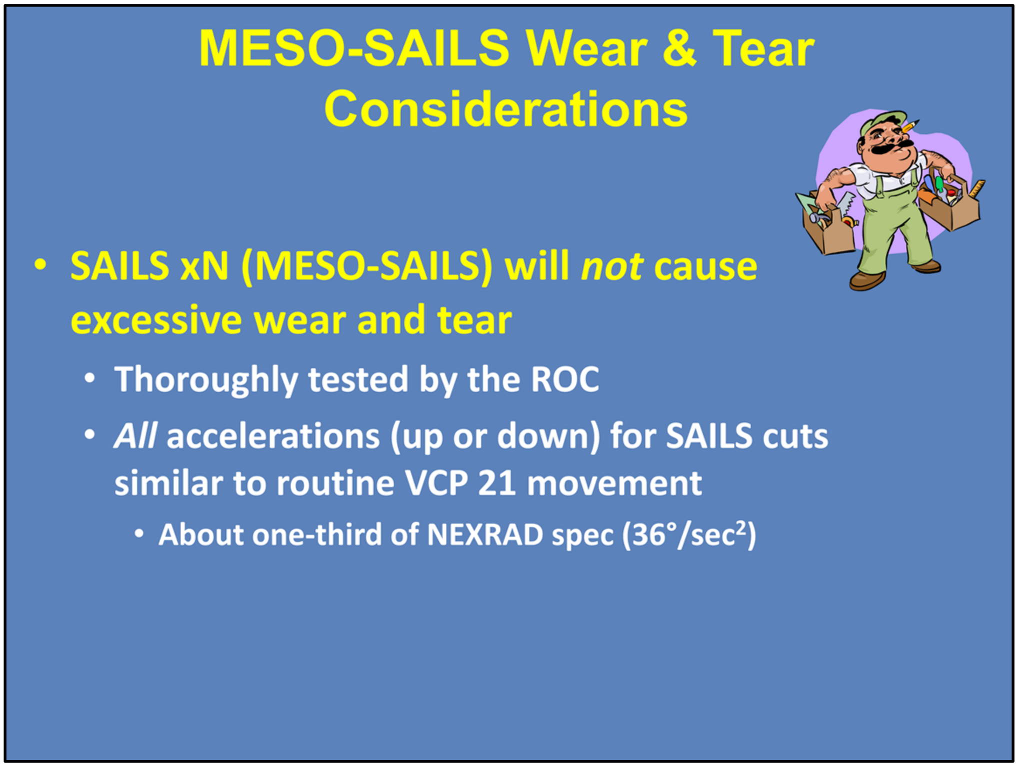

When the faster VCPs, 12 and 212, were fielded, there was initial concern about more wear and tear on the moving hardware of the WSR-88Ds. When SAILS and MESO-SAILS were being tested before deployment, it was understood that no feature was going to be deployed that resulted in excessive wear and tear. 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.

Welcome to this lesson on the Radar Product Generator Human Control Interface Controls, which is part of the Radar & Applications Course topic on Principles of Radar. Let’s get started!

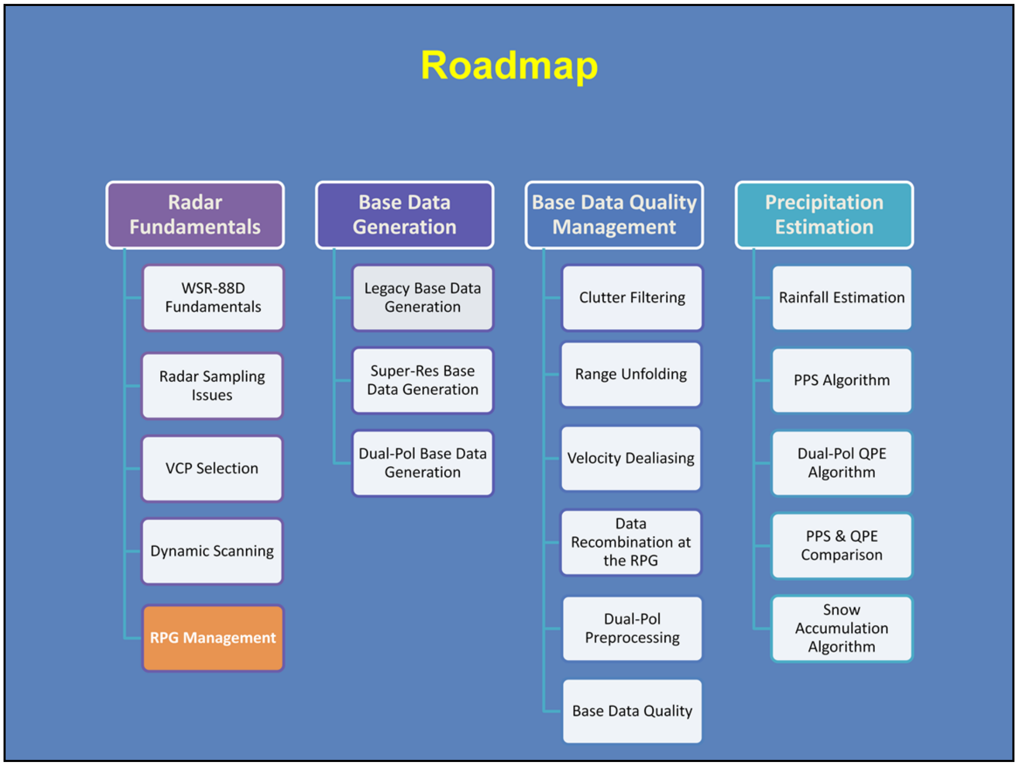

Here is a “roadmap” for the Principles of Radar topic, with your current location highlighted.

Learning Objectives

RDA Controls:

- Identify the 4 RDA control options
- Identify the 3 RDA control permissions groups
- Identify the 3 ways RDA status messages can be filtered

RPG Controls:

- Identify the 4 commands to control the RPG
- Identify the 2 operational modes for the RPG
- Identify the 2 repositories for data archiving
- Identify the 3 RPG alarm groups seen when checking the RPG status
- Identify the purpose of the wideband & narrowband lines
- Identify the location for console message dissemination

Here are the learning objectives for this lesson. I listed the objectives depending on whether they apply to controlling the Radar Data Acquisition Unit (or RDA) or the Radar Product Generator (or RPG). Please take a few moments to look over these objectives and advance to the next slide when ready.

RDA Control & Monitoring

- HCI allows user to view, modify, & monitor radar
- Users can control RDA, RPG, & comms
- Start w/RDA Control & Monitoring

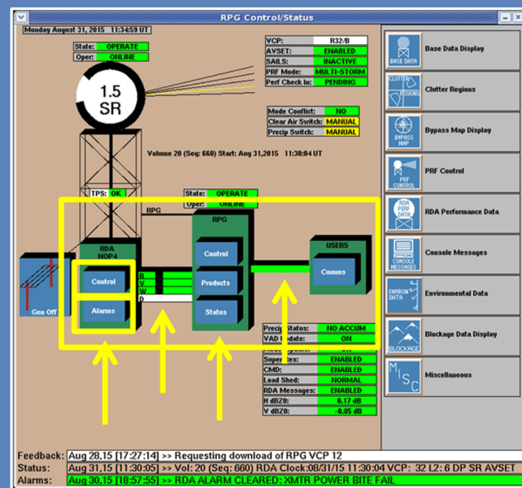

The RPG Human Control Interface (or HCI) shown here serves as a control panel of sorts for the WSR-88D. From this display, forecasters can view, modify, and monitor radar settings and ensure the system is functioning properly. The interface provides access to controls for the RDA (where base data are originally collected), the RPG (where the products you view are made), as well as the connections between the two and users. If we focus on the HCI a little further, you can see blue buttons located in green boxes labeled RDA and RPG. This lesson will discuss some of the common control features available from these buttons. Let's start with the control and monitoring capabilities the RPG HCI has for the RDA.

Back to Top

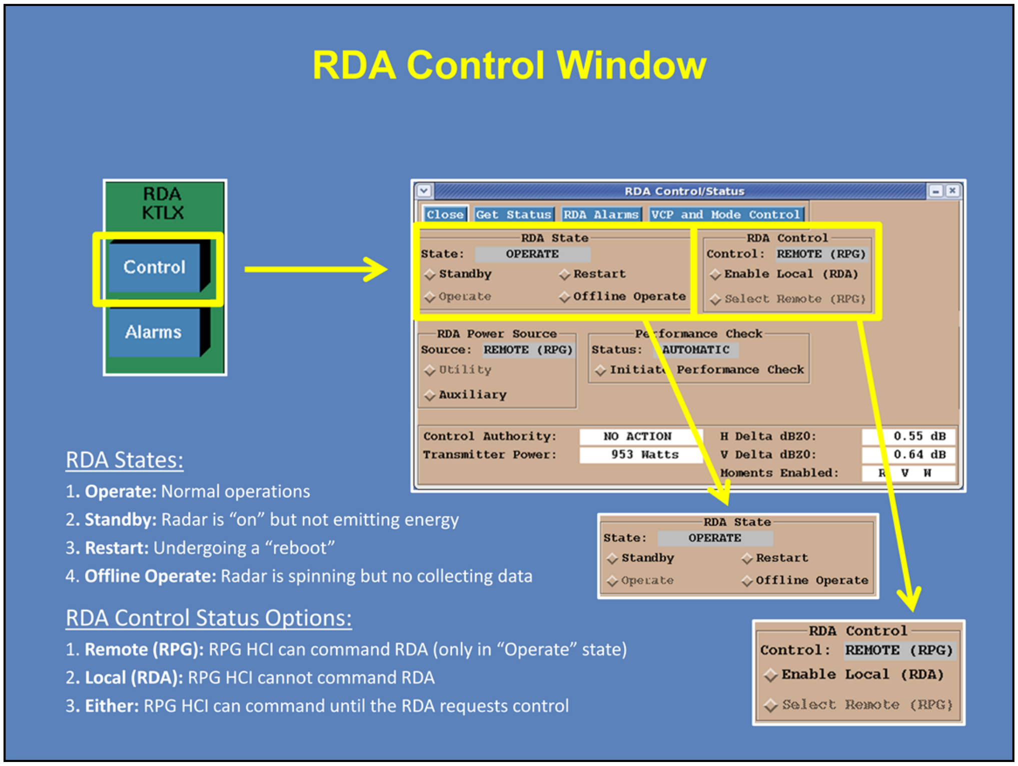

Clicking on the RDA Control button in the RPG HCI opens the RDA Control/Status Window. This GUI controls the basic functions of the RDA. The RDA State describes how the RDA is currently functioning. When running routinely, the state will be "Operate". If you need to change the RDA State (and remote RPG control is enabled) you can click on the check boxes next to each status. Besides "Operate", there are three other RDA states: Standby, Restart, and Offline Operate. Likewise, there are three RDA Control states: Enable Local (RDA), Select Remote (RPG), and Either. These states are described in a little more detail on the lower-left portion of the slide. You can look through these descriptions and advance to the next slide when ready.

Back to Top

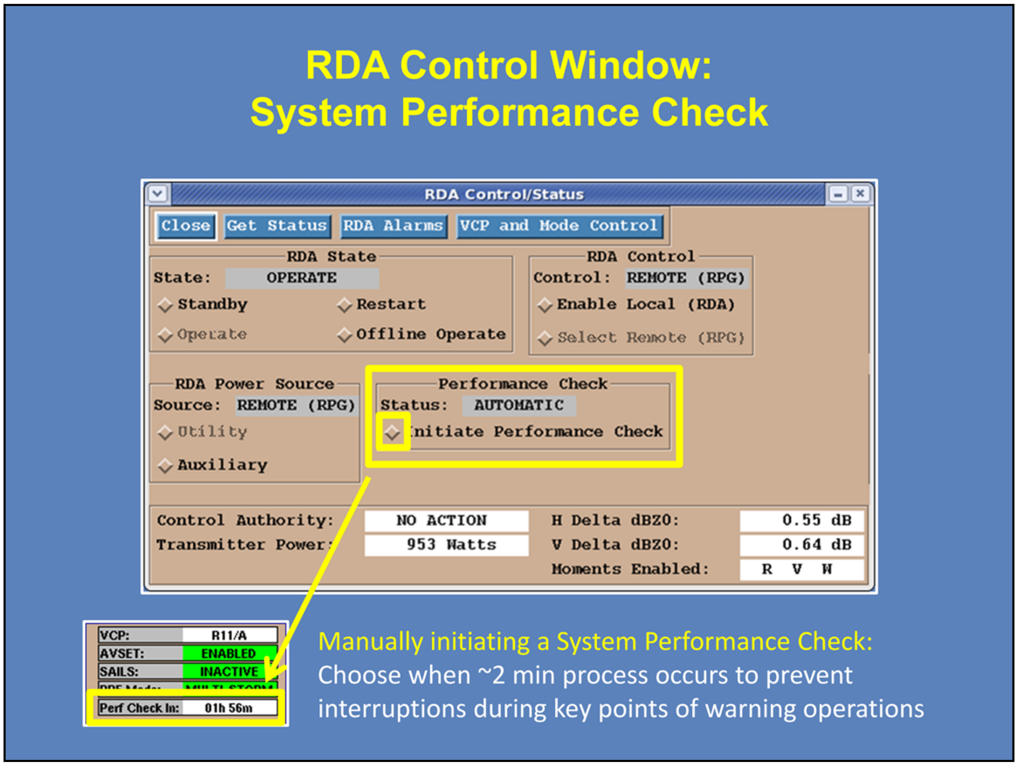

The RDA Control/Status GUI has another feature that forecasters may need to use more frequently than changing the RDA state or control settings. The radar runs a performance check every 8 hours to ensure the system operates correctly. During this check, the radar collects no data. The next scheduled performance check could occur when significant weather is expected, which is not desired. When this happens, forecasters can initiate a performance check during a quieter period to ensure data are not lost at a more critical time.

The next scheduled performance check will occur based on the timer listed on the main RPG HCI. To initiate an earlier system check, just click the box next to “Initiate Performance Check”. Doing so will push back the subsequent system check 8 hours. Just make sure you click on the correct box!

Back to Top

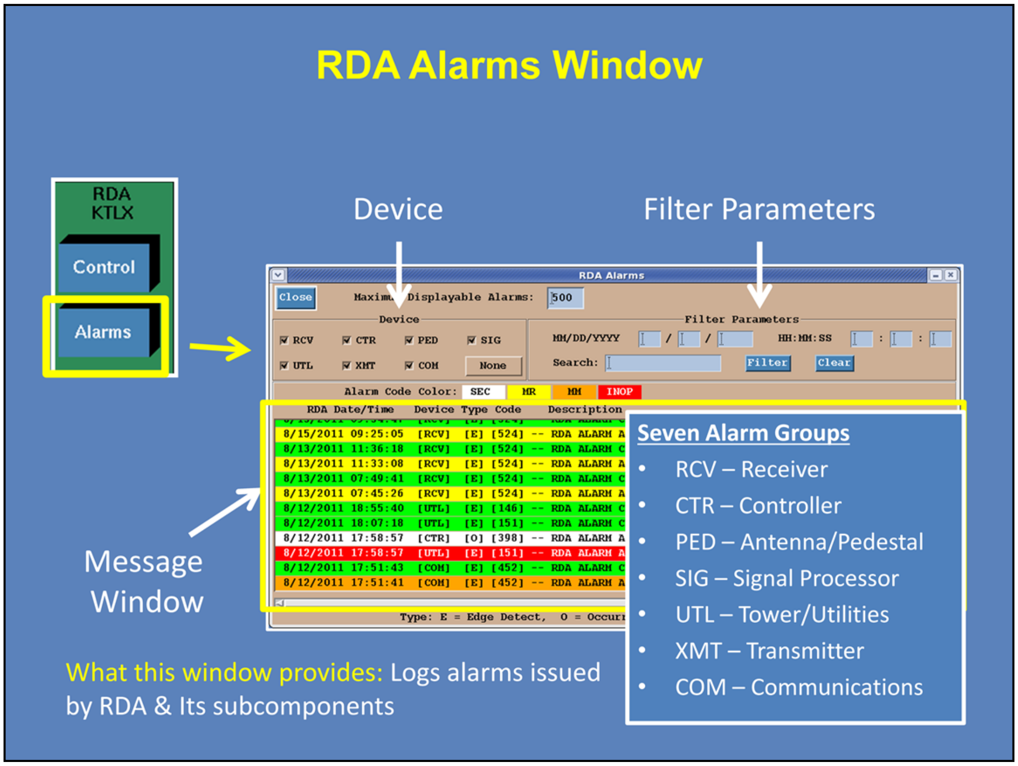

Just as you can control features of the RDA from the RPG HCI, you can also monitor RDA performance by clicking on the “Alarms” button. From the RDA Alarms window, you can view and filter any alarms issued by the RDA. The alarms will appear in the Message Window in the main part of the GUI, which includes a scroll bar in case you need to look for older alarms that are not visible. You can sort the alarms by device using the seven checkboxes in the Device panel, by time/date using the Filter Parameters fields, or by text content using the Search field. NOTE: The abbreviations in the Device panel correspond to alarms in these 7 groups shown on the screen. Remember that toggling a device “off” will make the alarm disappear, but not make the underlying problem go away! So, you generally want to keep all of these groups turned on.

The yellow and green colored alarm messages can generally be ignored as they are usually just status messages, not actual alarms. Orange and red labeled messages are another matter. If these alarms don’t clear themselves in a couple of messages, you should alert the Electronics (or El) Tech on duty or shift supervisor so that they can troubleshoot the issue further.

Back to Top

RPG Control & Status Checking

- 3 buttons for the RPG:
  - RPG Control
  - RPG Status
  - Products (not covered in this lesson)

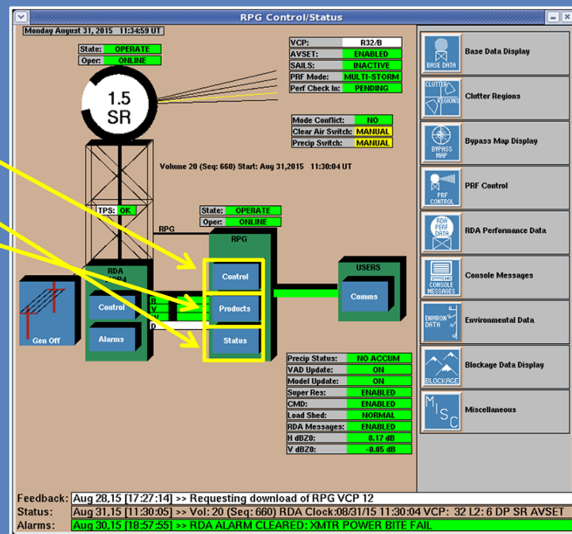

So, we've covered the main features of the RDA buttons on the main HCI display. Now let's talk about the RPG buttons. To control the RPG, users can access features through the "Control" button. Likewise, monitoring features can be accessed through the Status button. The third button, labeled Products, allows you to control the products generated by the RPG. For the sake of time, we will only discuss the Control and Status features in this lesson. Let's take a look at them!

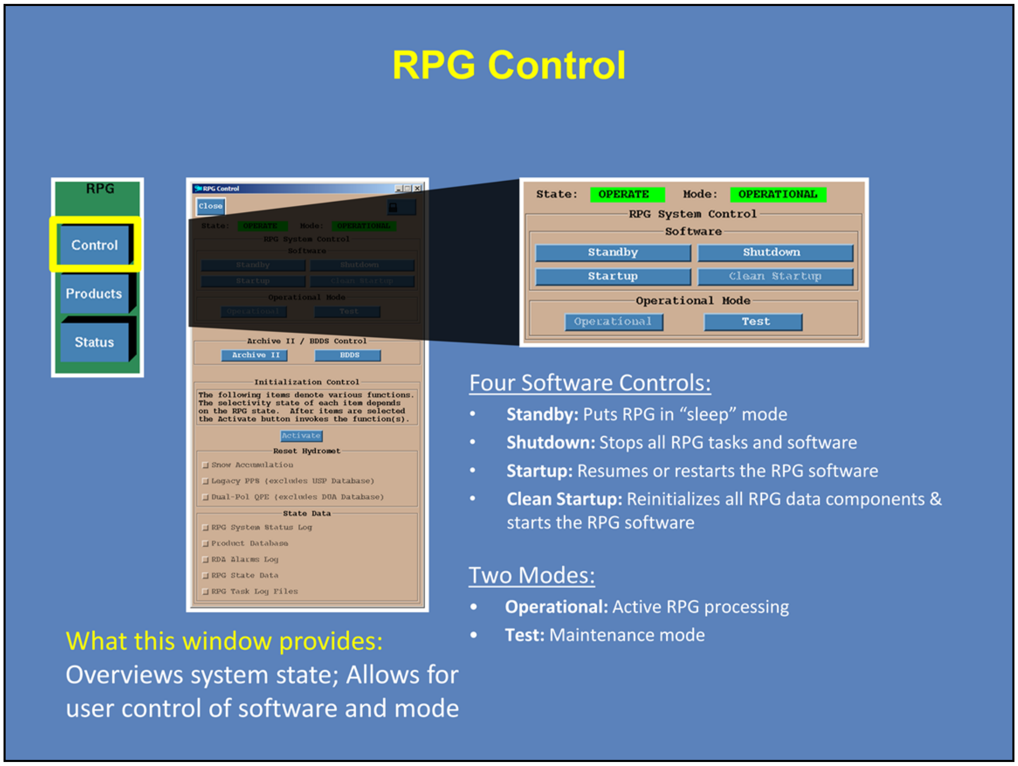

Selecting the Control button on the main HCI display will bring you to the RPG Control GUI. As with the comparable RDA GUI, you can see the current state of the RPG. Users can also initiate four software control commands: Standby, Shutdown, Startup and Clean Startup. You can also initiate two operational modes (not to be confused with the operational modes related to VCP selection): Operational or Test.

A basic explanation of what these controls and modes mean are provided on the slide. Take a moment to look them over, then advance to the next slide when ready.

Back to Top

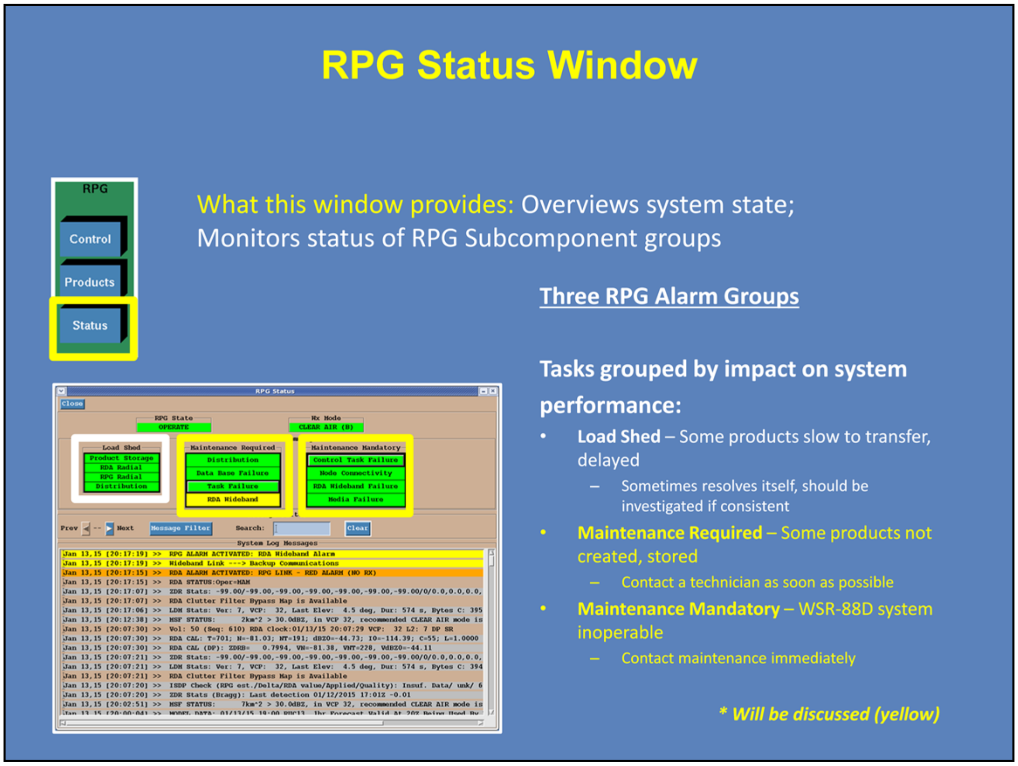

As with the RDA, the Status button for the RPG will open the RPG Status GUI and allow users to monitor status messages and alarms. The messages are grouped into three categories based on their impact on system importance. These categories are: Load Shed, Maintenance Required, and Maintenance Mandatory. Load shedding alarms are usually less critical than the others. Many times, they resolve themselves without intervention from the WFO staff. If consistent load shedding errors appear, forecasters should investigate further.

For the sake of time, that is all we will say about load shedding in this lesson. The other alarms will be discussed in further detail on the next slide.

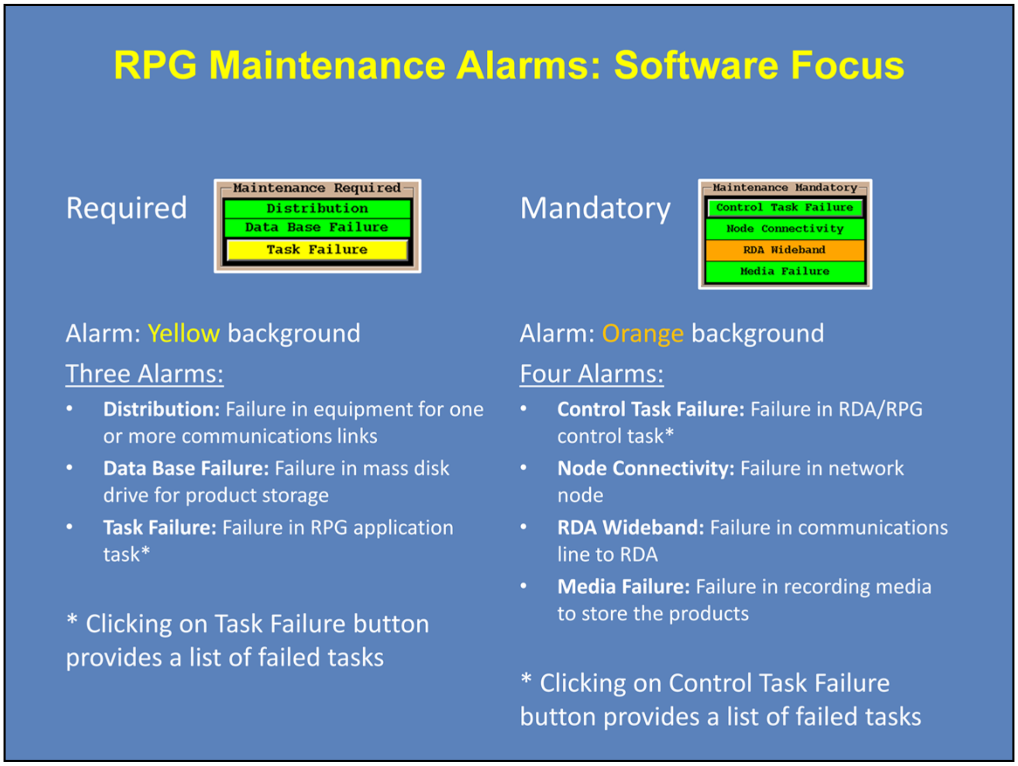

The Maintenance Required and Maintenance Mandatory alarms require users to take action, unless they clear themselves automatically. Unlike the RDA alarms (which focus more on hardware), the RPG alarms usually relate to software issues. Common issues include a backlog of products, improper generation of products and connectivity issues. The Required alarms on the left appear in yellow and are a little less serious, but still important. The Mandatory alarms on the right appear in orange and are more time sensitive because the issue prevents the RPG from doing its job. I've listed the categories for both types of alarms. Take a moment to look through the list and advance to the next slide when ready.

Back to Top

Overview of RPG "Comms" (Communication)

Communication lines necessary for transfer of WSR-88D Products

Topics Covered:

1. RDA -> RPG (Wideband Line)
2. RPG -> Users (Narrowband Line)
3. RPG -> Archive (RPG Control/Archive)
4. Operator -> Users (Console Messages)

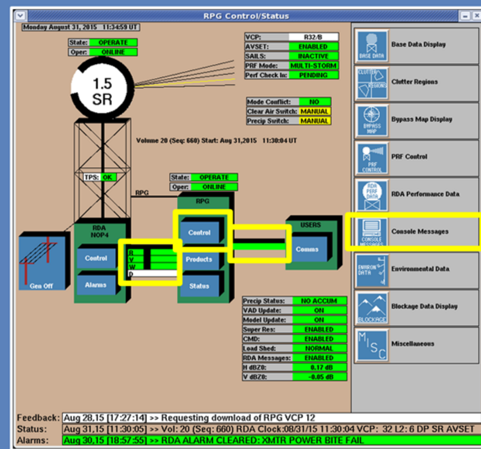

The last area we will discuss involves communications (or comms) to and from the RPG. Communication lines are important conduits for data through the WSR-88D system. While the WSR-88D implements numerous communication procedures, we will focus on four:

1) Wideband comms between the RDA and the RPG,

2) Narrowband comms between the RPG and users,

3) Archiving of WSR-88D data, and

4) Free Text Messages from the radar operator to other users.

Back to Top

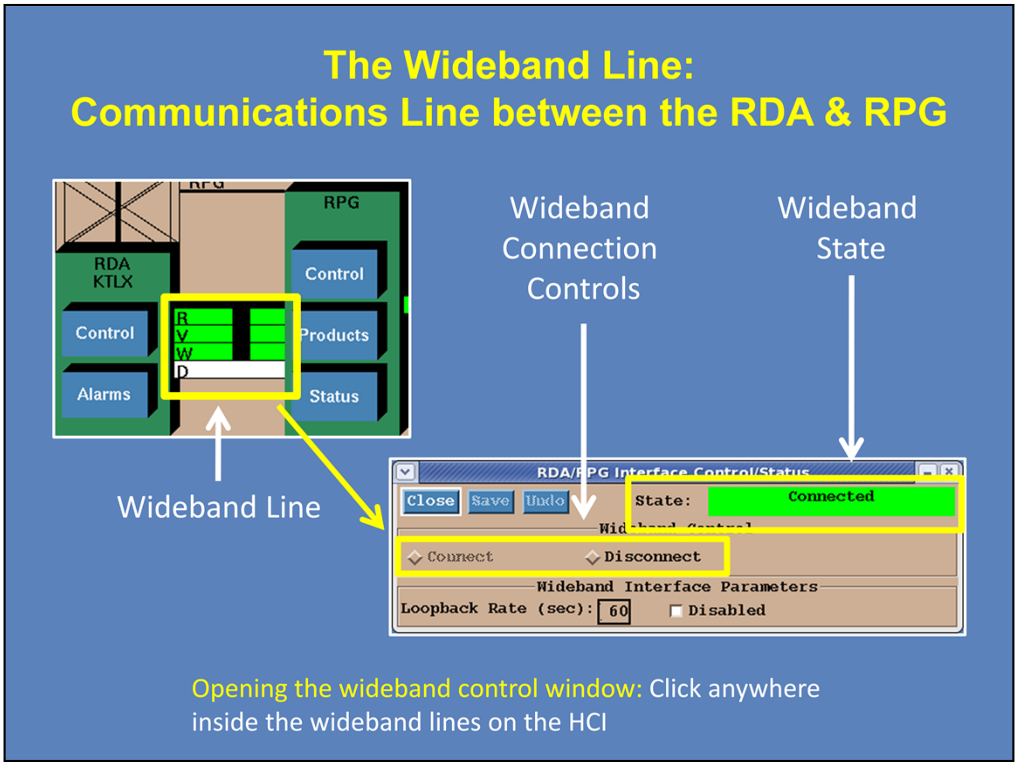

The wideband comms line connects the RDA and the RPG. The RPG HCI depicts this connection as four lines between the RDA and RPG control boxes labeled R (for Reflectivity), V (for Velocity), W (for Spectrum Width), and D (for Dual-Pol). The RPG needs all four of these streams to generate the base and derived products for users each volume scan.

Clicking on the icon in the RPG HCI will open the RDA/RPG Interface Control/Status GUI. From here, you can check the state of the comms line and change its status. NOTE: in general, connecting and disconnecting the wideband comms should be handled by your El Tech. However, forecasters should be familiar with how to access this functionality in case they need to reactivate the connection for some reason.

Back to Top

The Narrowband Lines: Communication Lines between the RPG & External Users

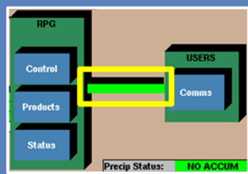

Opening the narrowband control window: Click on either the connection line or the Comms button in the USERS box on the HCI

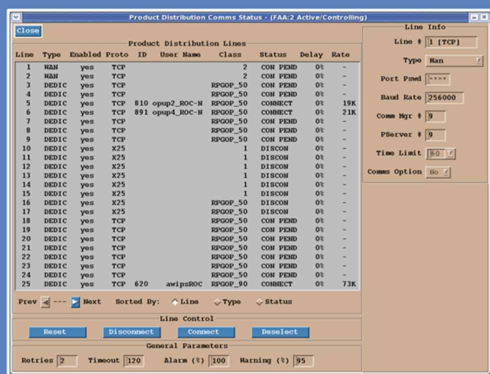

Single Connection per Line

- Connection Types:
  - **DEDIC:** Dedicated
  - **DIALIN:** Dial-Up
  - **WAN:** Wide Area Network
- Connection Status:
  - **CONNECT:** Connected
  - **DISCON:** Disconnected
  - **CON PEND:** Connection Pending

The product distribution (or "narrowband") comms connect the RPG with its various users. The "narrowband" moniker dates back to when these connections were significantly narrower in bandwidth than the wideband comms lines. The RPG HCI depicts the connection as a single line between the RPG and User control boxes. Clicking on this icon opens the Product Distribution Comms Status window. This window shows all of the connection lines available to users and their current status. Your Radar Program leader should have a list of what's called "Dedicated Users" who have a pre-defined line on this display. Other users will access the RPG via dial-up or wide area network (or WAN) connections.

This displays can help forecasters troubleshoot connection issues for other users. If another user has a connection issue, the status should say "DISCON" for disconnected or "COM PEND" for connection pending. Not every line with these states are a problem, though. Just keep that in mind if you routinely check this window as part of a radar operations checklist.

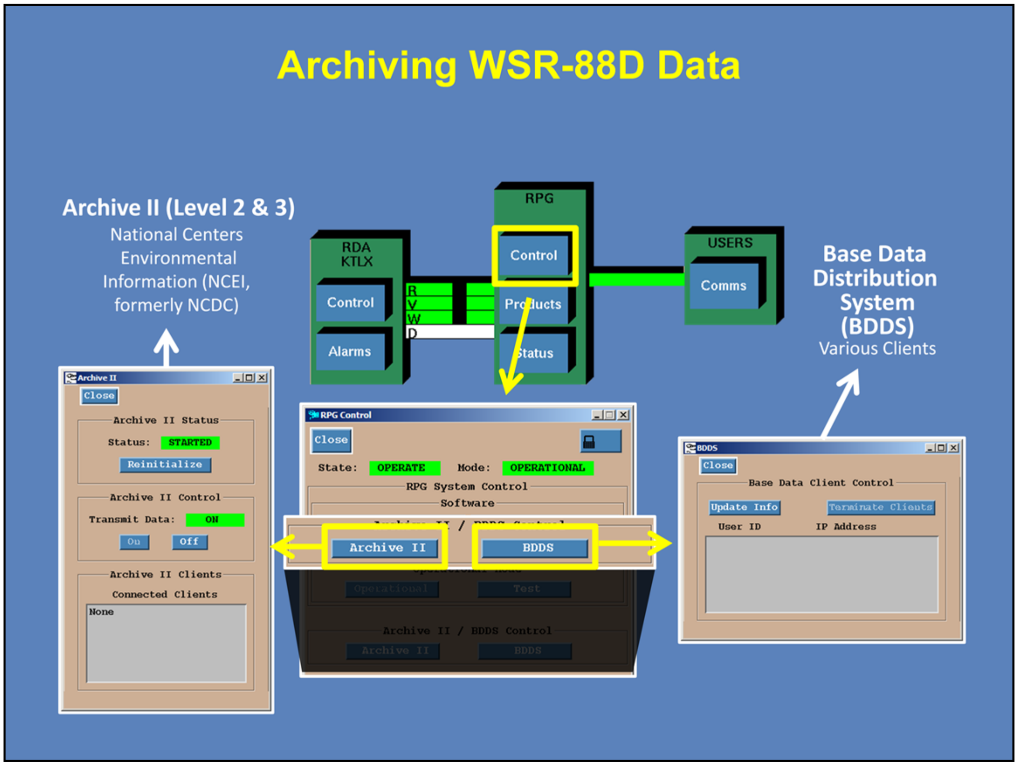

Two archiving methods exist for the WSR-88D, and both are important. The first method involves the National Centers for Environmental Information (NCEI), formally known as the National Climatic Data Center. They archive both the Level II and Level III data and make them available for public distribution. The second method involves the distribution of Level II data only to specific clients through the Base Data Distribution System (or BDDS). The status of these archives can be viewed and modified by clicking on the "Archive II" and "BDDS" buttons, respectively. Each archive has its own configuration window where you can monitor and manage the status of each process. Forecasters shouldn't need to change these settings unless a significant issue has occurred.

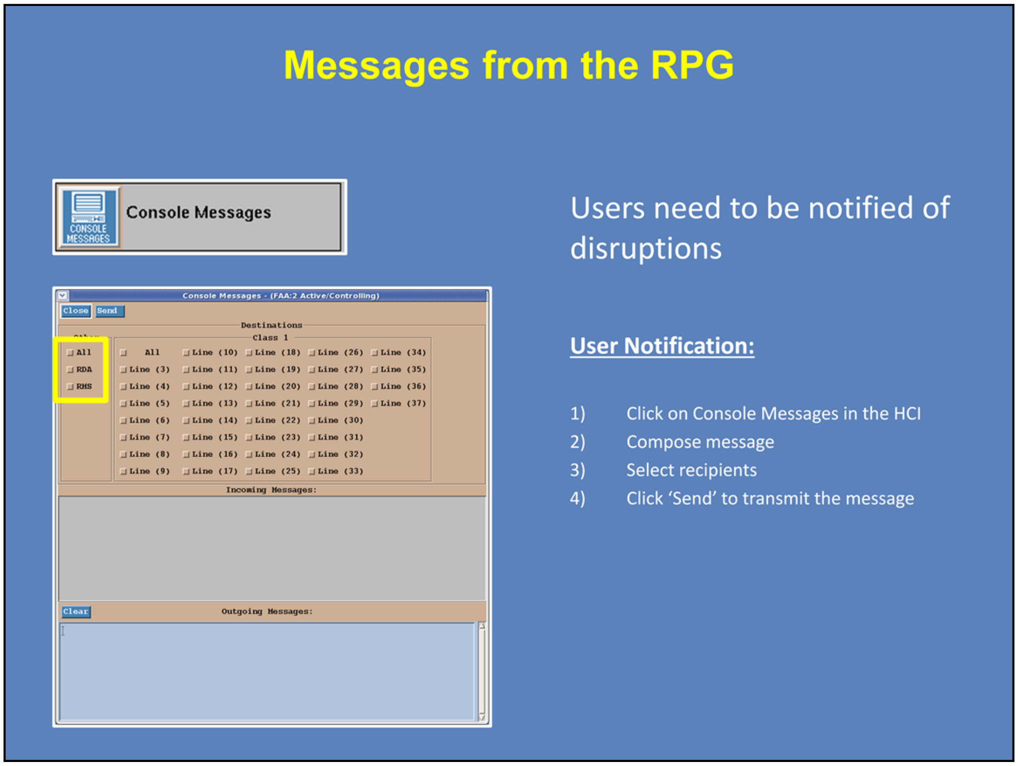

When there is change in a radar operability status, it is a best practice to notify all users of the change. A quick way to notify users involves sending a Free Text Message product, or FTM. Use of FTMs is a best practice when the radar has failed or will be taken down for routine maintenance. These messages can be generated from the RPG HCI by clicking the “Console Messages” button on the right-hand side of the window, causing the Console Messages window to appear. Make sure to click on each checkbox in the “Other” section to send the FTM to all external users.

Once you have finished the message, click “Send”. You won't see anything happen on your end. However, you can view the FTM in an AWIPS text editor window. In fact, we recommend you do look at the message there just to ensure it got sent out.

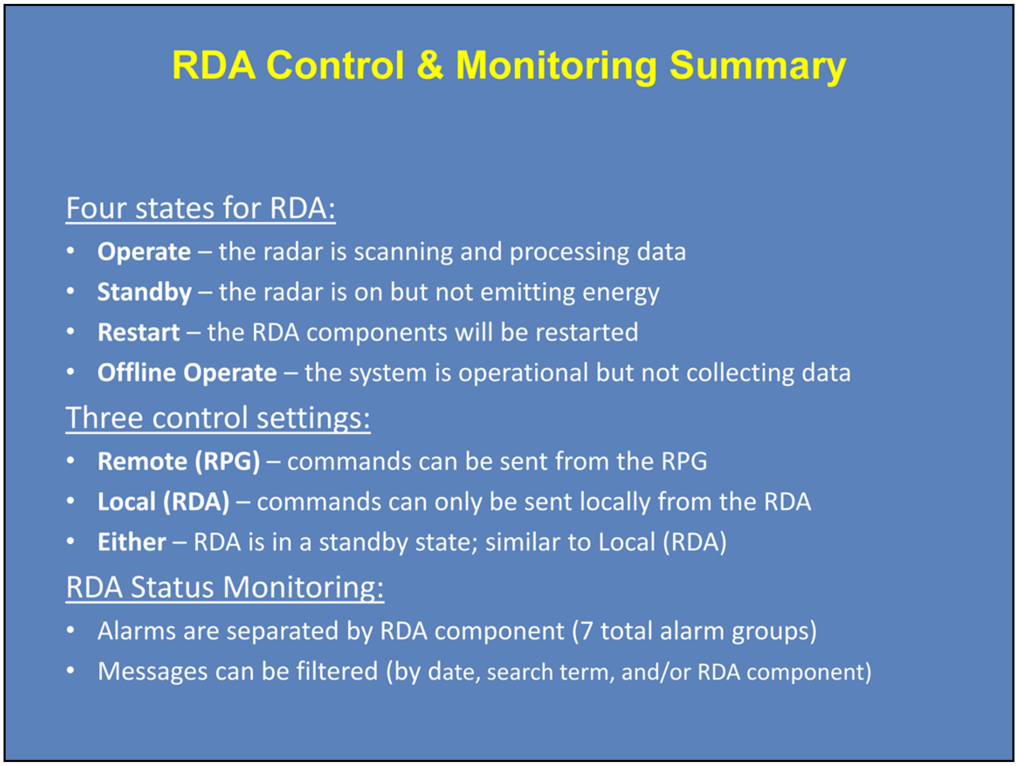

This lesson covered three different areas of the RPG HCI, so I will summarize each section separately. We started off discussing the RDA control and monitoring functionality. The RDA has four states: Operate, Standby, Restart, and Offline Operate. If the radar is not in "Operate", then some form of maintenance or troubleshooting is on-going. The user's ability to change the RDA state at the RPG HCI will depend on the three RDA control settings: Remote RPG, Local RDA, or Either. Users can monitor RDA alarms from the status window. You have the ability to filter alarms by RDA device (of which there are 7 groups), or you can use other filtering tools.

Back to Top

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

Users can initiate four software system controls from the RPG HCI: Standby, shutdown, startup, and clean startup. The RPG also has two operational modes (not to confused with the operational modes associated with VCPs): Operational and test. Just as with the RDA, users can monitor alarms from the RPG by using the status window. RPG Alarms are grouped into three categories: Load shed, maintenance required, and maintenance mandatory. Alarms in the latter two categories require immediate attention. Load shed alarms should be investigated if they continue repeatedly for several minutes.

Back to Top

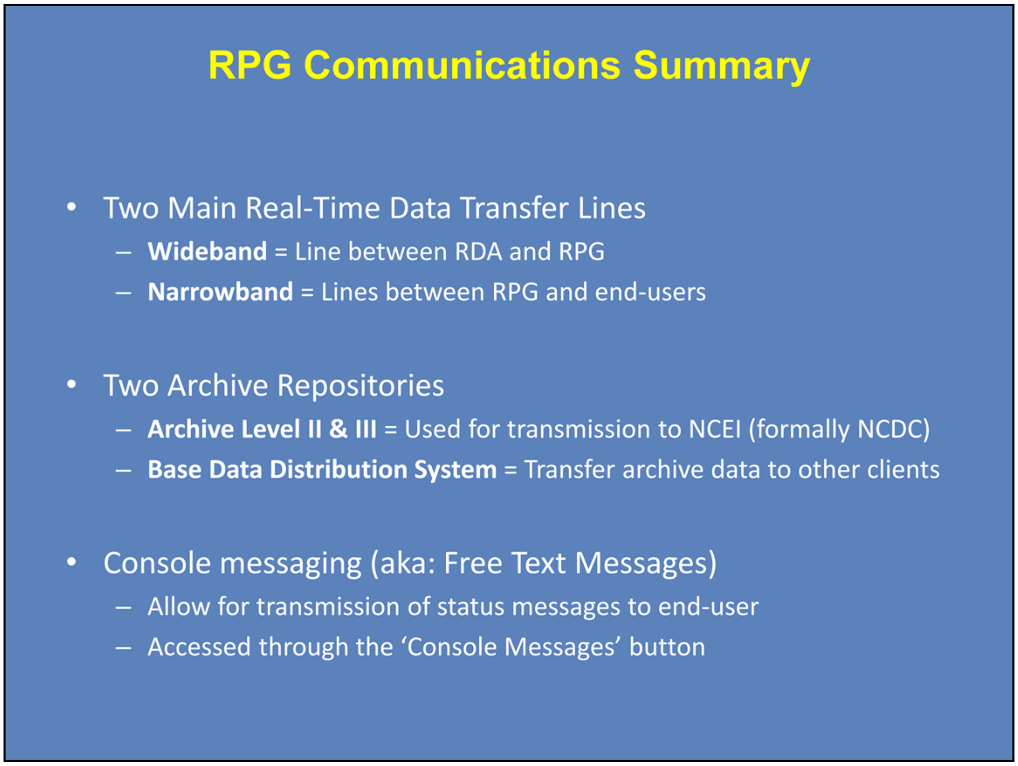

The last topic covered was the communications between the different components of the WSR-88D system. The two main comms lines are the wideband connection between the RDA and RPG and the narrowband connection that distributes products from the RPG to users. The WSR-88D provides two ways for data to be archived: Archive Level II and Level III at NCEI and the Base Data Distribution System that transfers the base data to specific clients. Lastly, when changes to radar operations are needed (or happen without warning), console messages (called Free Text Messages) are quick way to communicate those status changes to other users.

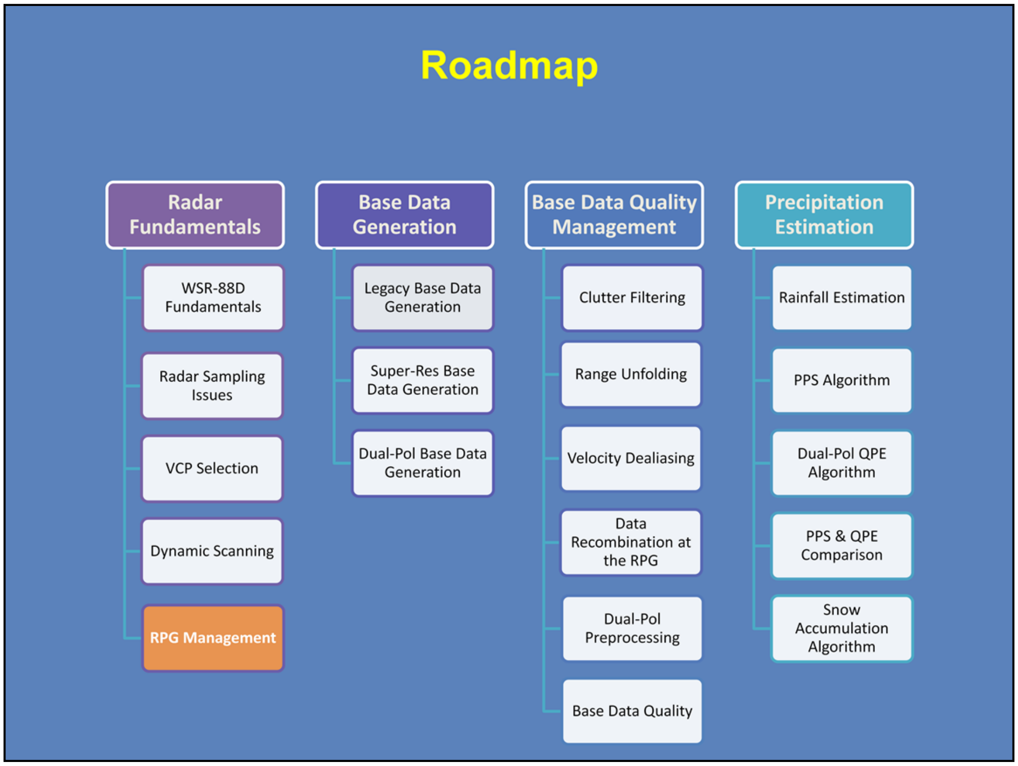

Thanks for your time. Good luck with the remaining lessons in the Principles of Radar topic of RAC.

Welcome to this lesson on Radar Product Generator (or RPG) Human Control Interface (or HCI) Functions, which is a part of the Principles of Radar topic in RAC. This lesson will discuss some of the core functionality that the RPG HCI provides radar operators. Let's begin!

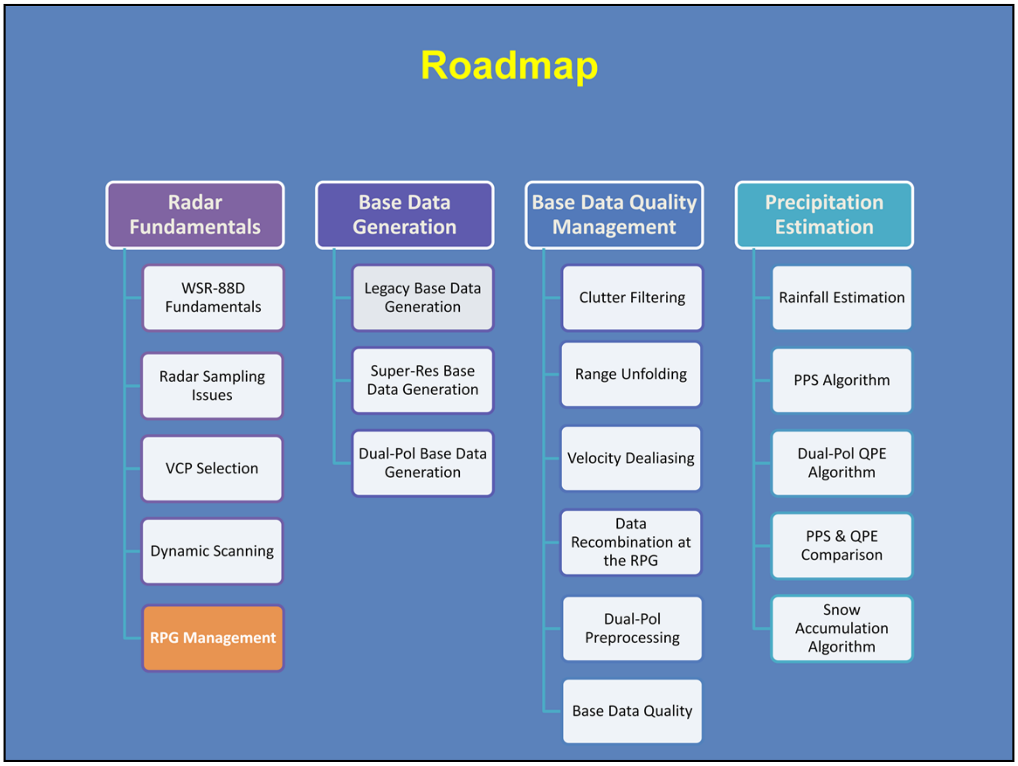

Here is the “roadmap” for the course with your current location highlighted.

Back to Top

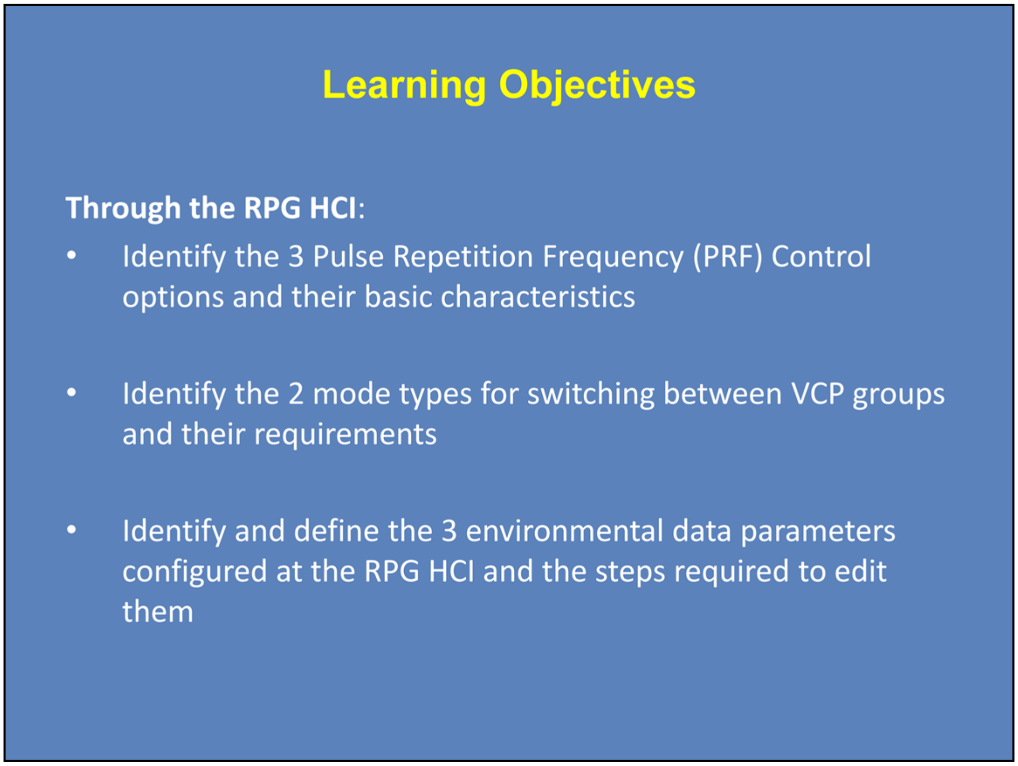

There are three learning objectives with this lesson. Please take a moment to read through these and move to the next slide when you're ready.

Back to Top

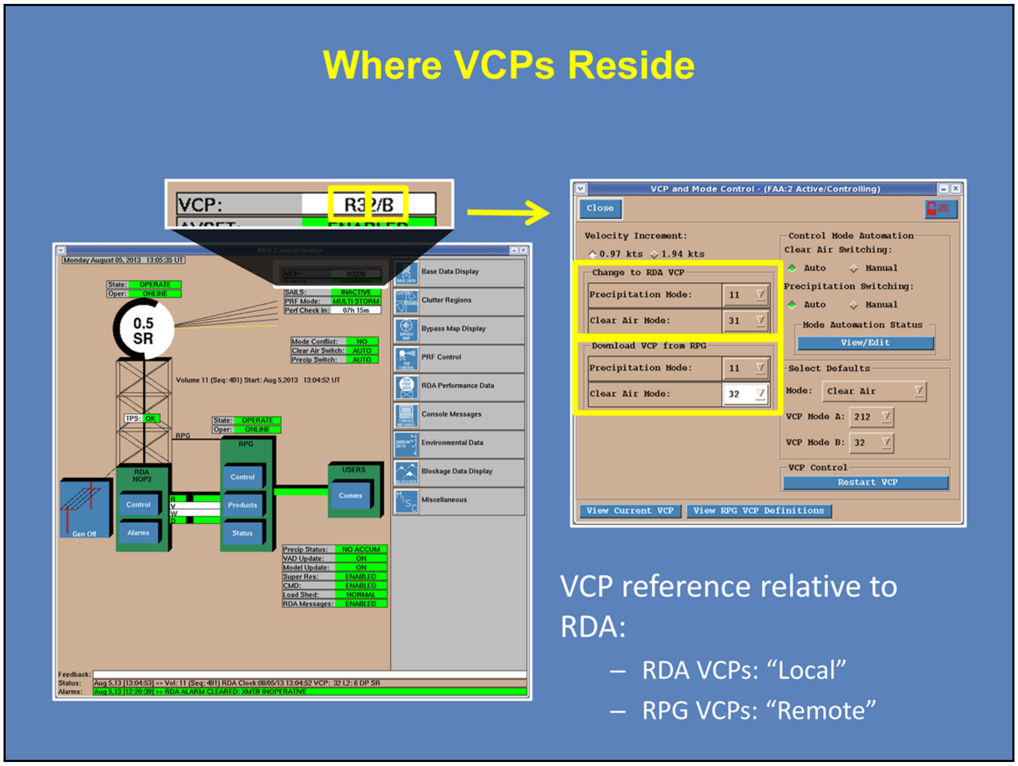

The WSR-88D uses volume control patterns (or VCPs) to collect data. Information about the current VCP in use can be found in the main RPG HCI window. To make changes to the current VCP (or VCP default settings), click on the VCP info near the top of the HCI to reveal the VCP and Mode Control window.

When talking about VCPs, you need to realize that VCPs are stored at both the Radar Data Acquisition Unit (or RDA) and at the RPG. The original VCPs (11, 21, 31, and 32) have been stored at the RDA since the original WSR-88D deployment. The RPG stores all of those VCPs as well as the newer ones such as 12 and 212. Because there are two copies of some VCPs, the GUI refers to them by their location relative to the RDA since that's where they are ultimately implemented. So, local VCPs refer to copies stored at the RDA, while remote VCPs are stored at the RPG. If you look back at the main HCI window, notice the extra characters listed with the VCP name. In this example, the “R” that precedes VCP 32 means it's the remote copy of VCP 32. The “B” that follows indicates that VCP 32 is a clear air mode VCP.

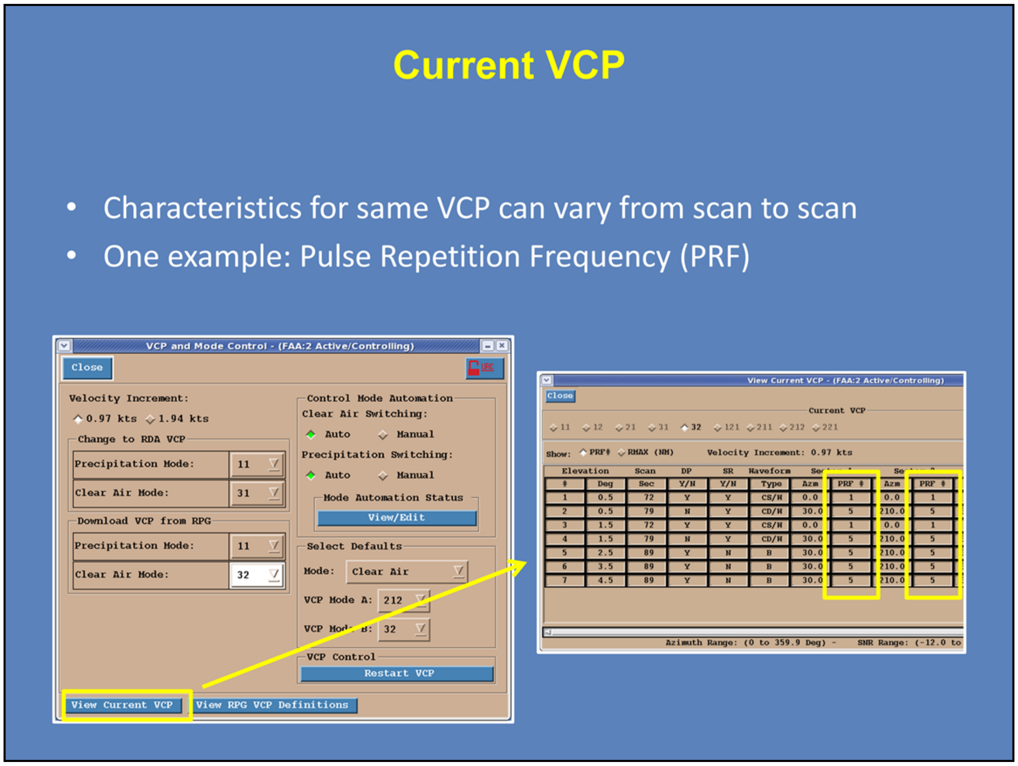

Each volume scan has a “Current VCP”. VCP 32 was the current VCP in the previous example. To see the details of the current VCP, just click on the “View Current VCP” button in the lower left-hand corner of the VCP and Mode Control window.

Even though the same VCP is used for multiple scans, some characteristics of the Current VCP may be different than previous scans. One such characteristic that can change from scan to scan is the Doppler pulse repetition frequency (or PRF). Let’s look at this characteristic in more detail to better understand why it changes.

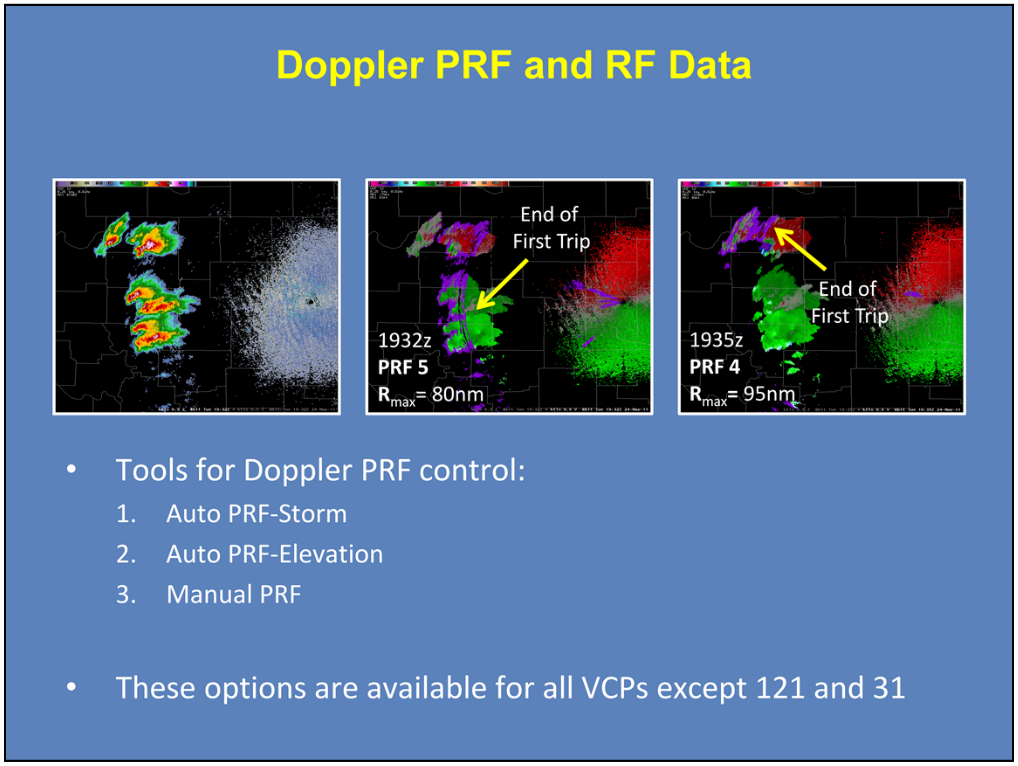

The WSR-88D uses multiple PRFs to give users the ability to minimize range folded velocity data. The system routinely displays velocity data beyond the first trip as a result of the five different Doppler PRFs that are used. Each PRF has a different resulting maximum unambiguous range (or R_{max}). The end of the first trip can range from 65 to 95 nm, depending on the PRF used. Users have three tools at their disposal to control the PRF used: Auto PRF-Storm, Auto PRF-Elevation, and Manual PRF. It should be noted that two VCPs (121 and 31) don't allow the user to change the PRF used for reasons related to features of those scanning strategies.

Back to Top

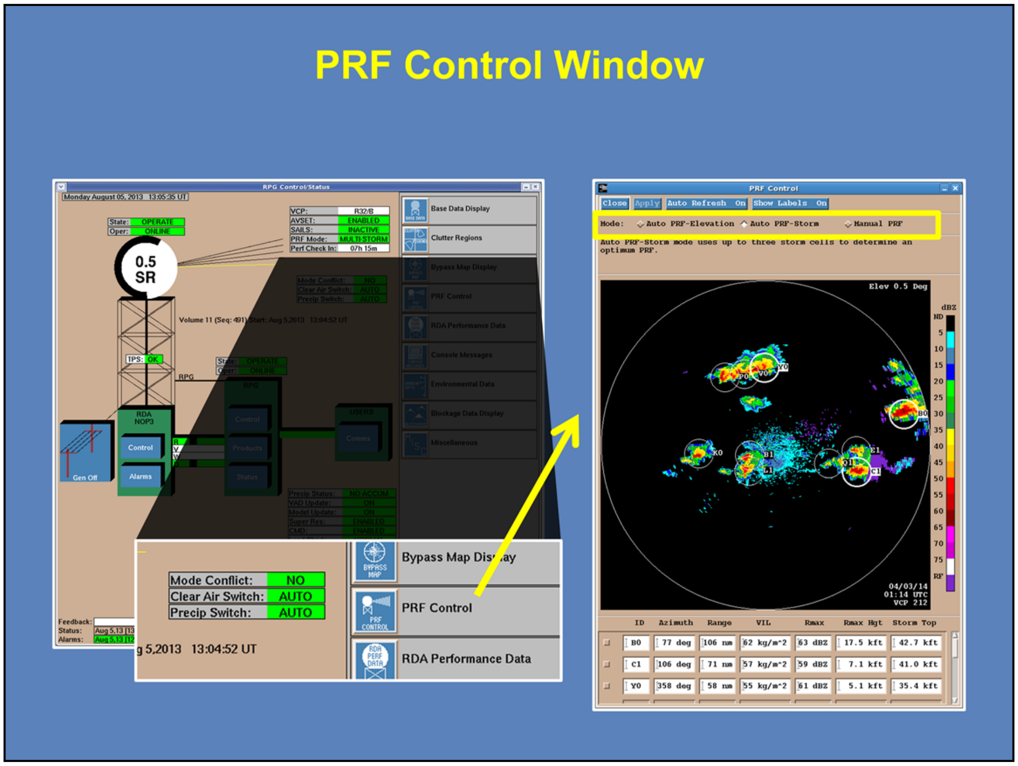

To access the PRF control techniques, select the “PRF Control” button from the main RPG HCI. Each control method is listed at the top of the window. The window shows a graphic with the 0.5° Reflectivity and the RF from the Velocity data only, with a table listing storm IDs and traits at the bottom of the window. More on these features in a moment.

Back to Top

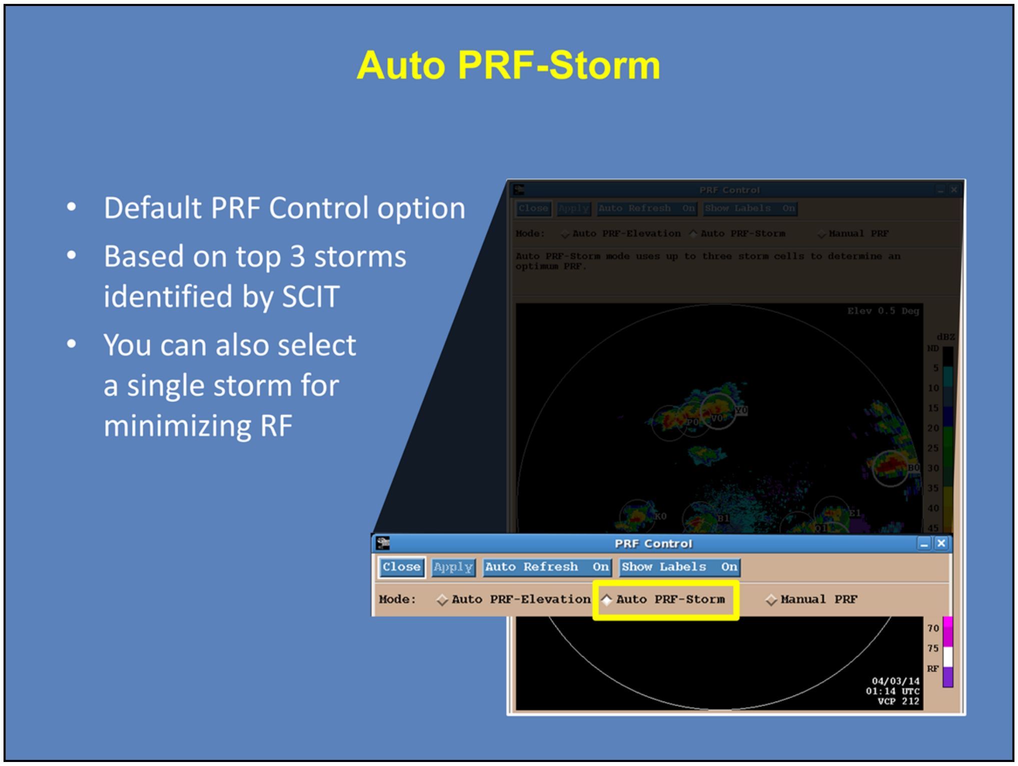

The RPG uses a default mode of “Auto PRF-Storm”. This mode chooses the Doppler PRF that results in the least amount of RF data for the top 3 storms in the table as identified by the Storm Cell Identification and Track (or SCIT) algorithm. Users can also adjust the settings to minimize RF data for a single storm with Auto PRF-Storm.

Back to Top

Auto PRF-Storm Process

- Default is top 3 storms
  - Ranking by Storm-Based VIL ($\geq 20\ kg/m^2$)
- Computes RF area in each 20 km circle around storms
  - Compare each Doppler PRF
  - Smallest area of RF for all circles is chosen
- Process repeated every volume scan

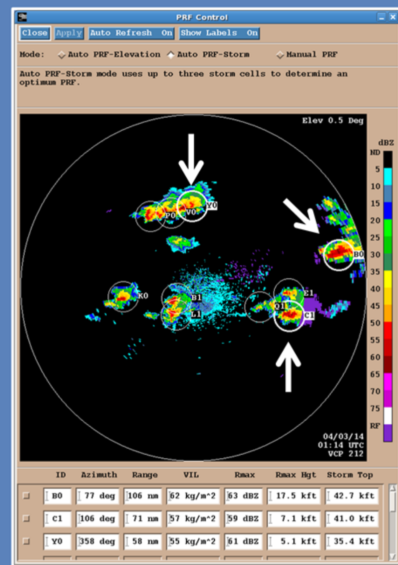

The three top storms in the table are the ones that SCIT thinks are the most intense based on their Vertically Integrated Liquid (or VIL) values. These storms can be identified in the Reflectivity graphic by the white circles and are labeled with their ID from the table. The RDA compares the range folded area inside the 20 km radius circle around each storm centroid and picks the one that results in the smallest area. Each volume scan, this process repeats.

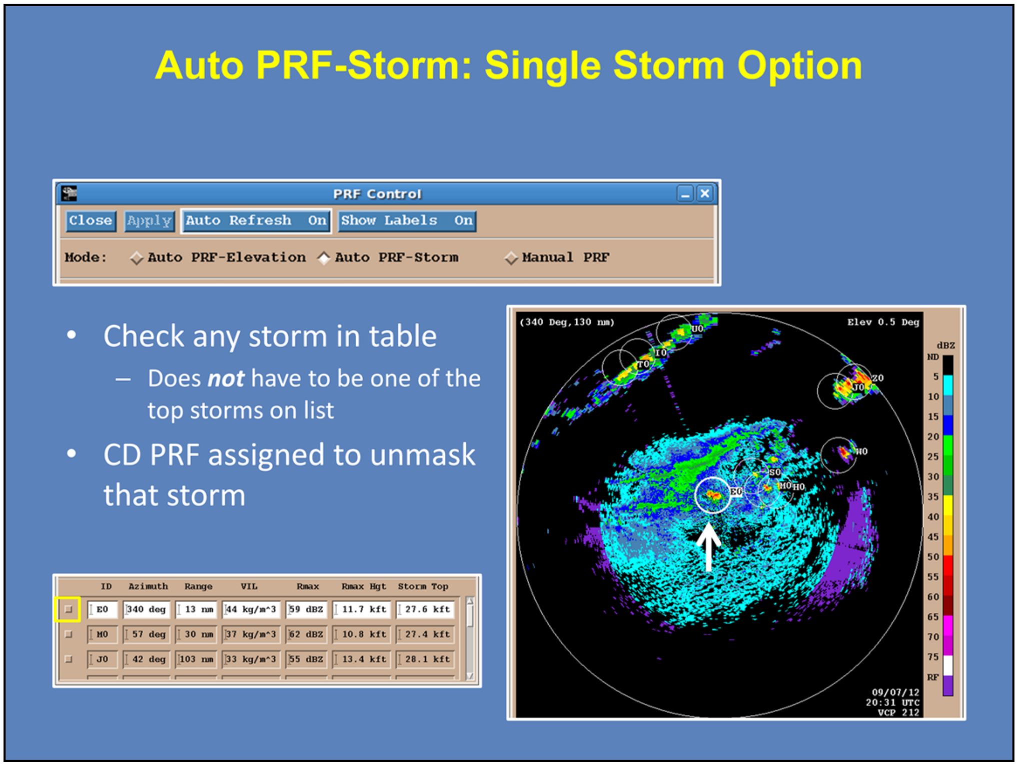

As mentioned previously, a single storm can be the focus of Auto PRF-Storm's range folding mitigation technique. Users just need to select the check box next to the storm in question, which doesn't have to be one of the three strongest. The result can be seen in the Reflectivity display window, where only one storm now has the white circle around it.

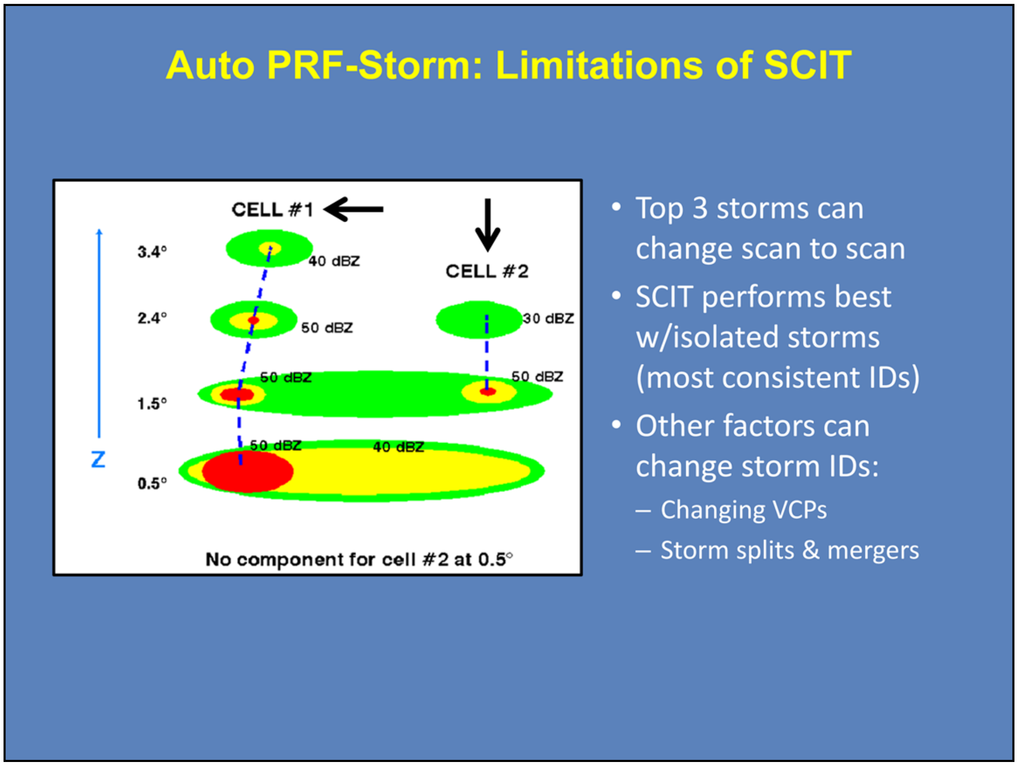

Since Auto PRF-Storm uses output from the SCIT algorithm to rank storms, so the limitations of SCIT should be kept in mind. The storms with the highest VIL values may change from scan to scan because of algorithm performance. The diagram on the slide shows two storms in close proximity. SCIT has a hard time tracking cells like Cell #2 because of its proximity to Cell #1. The SCIT performs best with discrete, isolated storms. These storms result in the most consistent storm IDs over time. When storms aren't isolated, then the best PRF may change frequently. Likewise, other factors (such as changing VCPs and storm splits and mergers) can result in storm IDs changing, also.

Back to Top

Auto PRF-Elevation

- Based on 0.5° Z and V
- Chooses Doppler PRF with least RF over ***entire*** display
- PRF chosen regardless of any particular storm's intensity

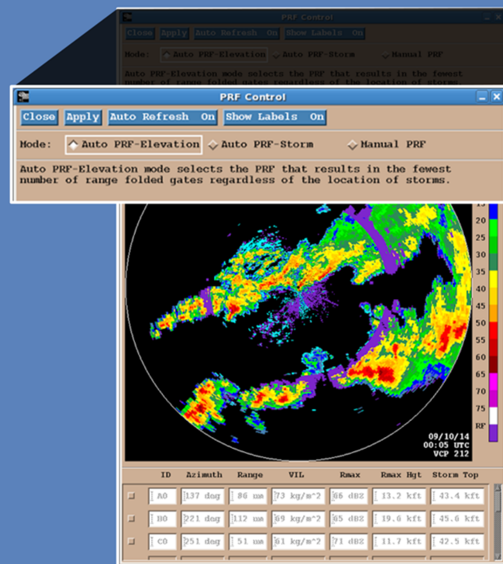

A second PRF Control mode is "Auto PRF-Elevation". This mode examines the total areal coverage in Reflectivity and Velocity data at 0.5°. This mode uses the PRF with the lowest areal coverage of range folding over the entire display. Auto PRF-Elevation pays no attention to actual storms or their intensities, so keep that in mind when using this mode.

Back to Top

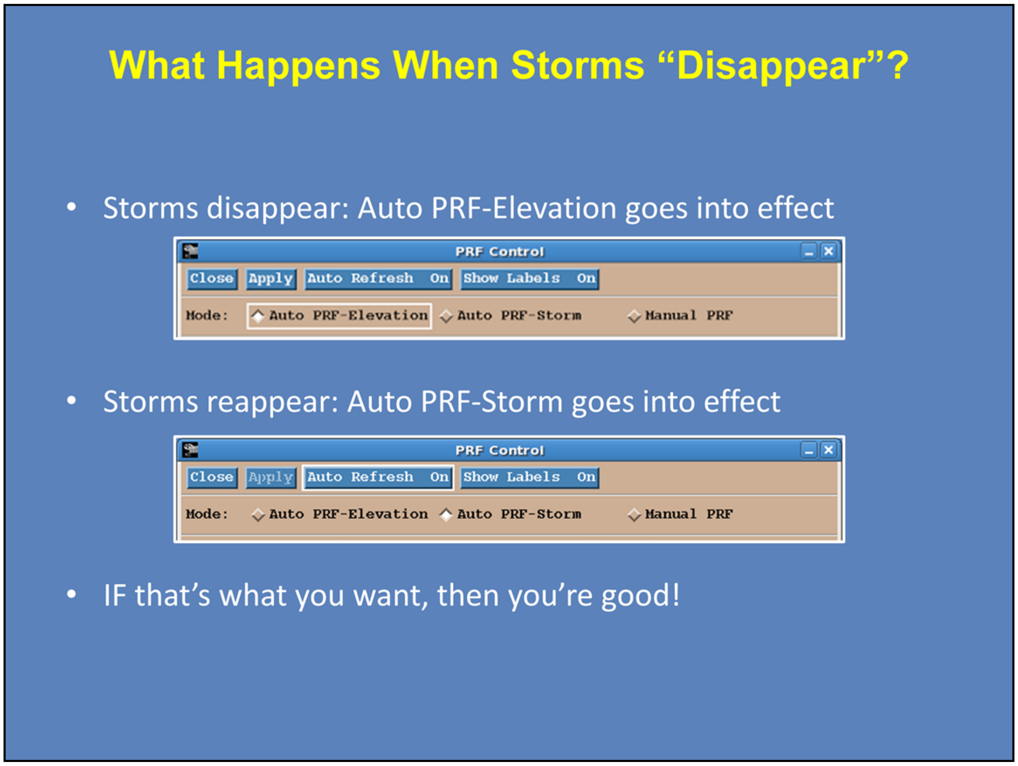

You may have asked yourself during the Auto PRF-Storm discussion: What happens when storms disappear? Good question. When storms are no longer visible to SCIT, the RPG switches to Auto PRF-Elevation automatically to control the PRF chosen. Once storms redevelop, Auto PRF-Storm takes over. If you wish to use those modes, you don't need to do anything!

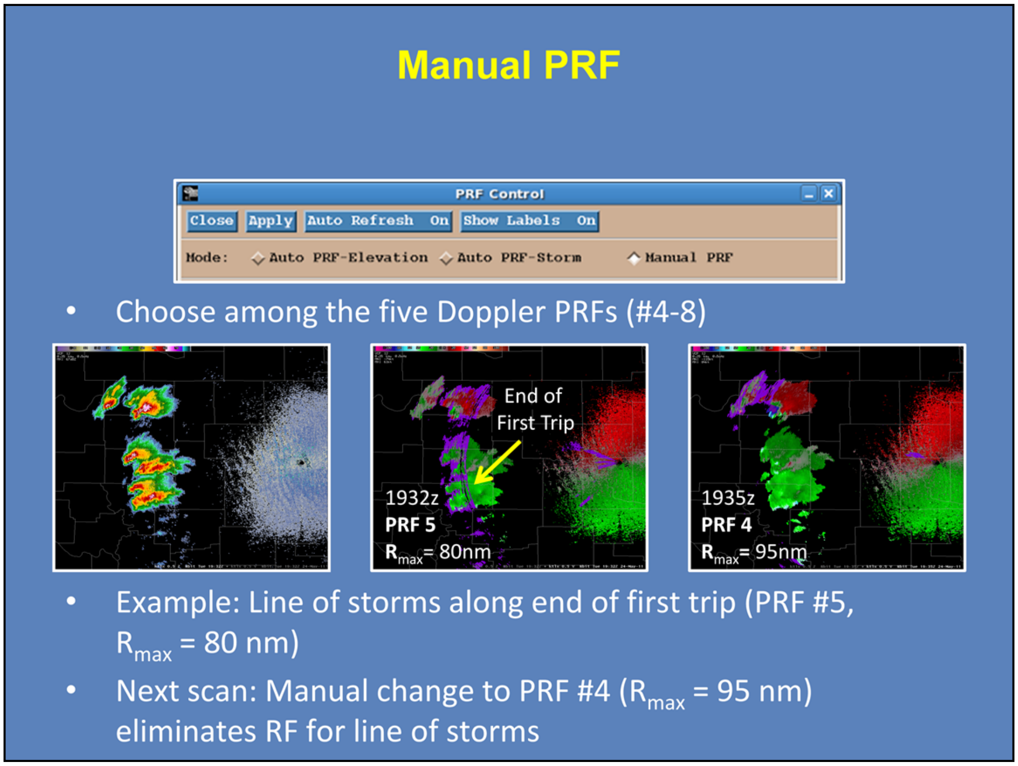

The last option for PRF control mode is Manual PRF. This mode allows the user to specify the exact Doppler PRF they want. In this example, a line of storms approaches the first trip for PRF #5 (which is in use). Switching to PRF #4, which has a larger maximum unambiguous range, you can see range folding no longer obscures the velocity data. This technique of manually switching PRFs can be very helpful when there are multiple strong storms that require your attention, but the storm(s) most requiring good velocity information are obscured.

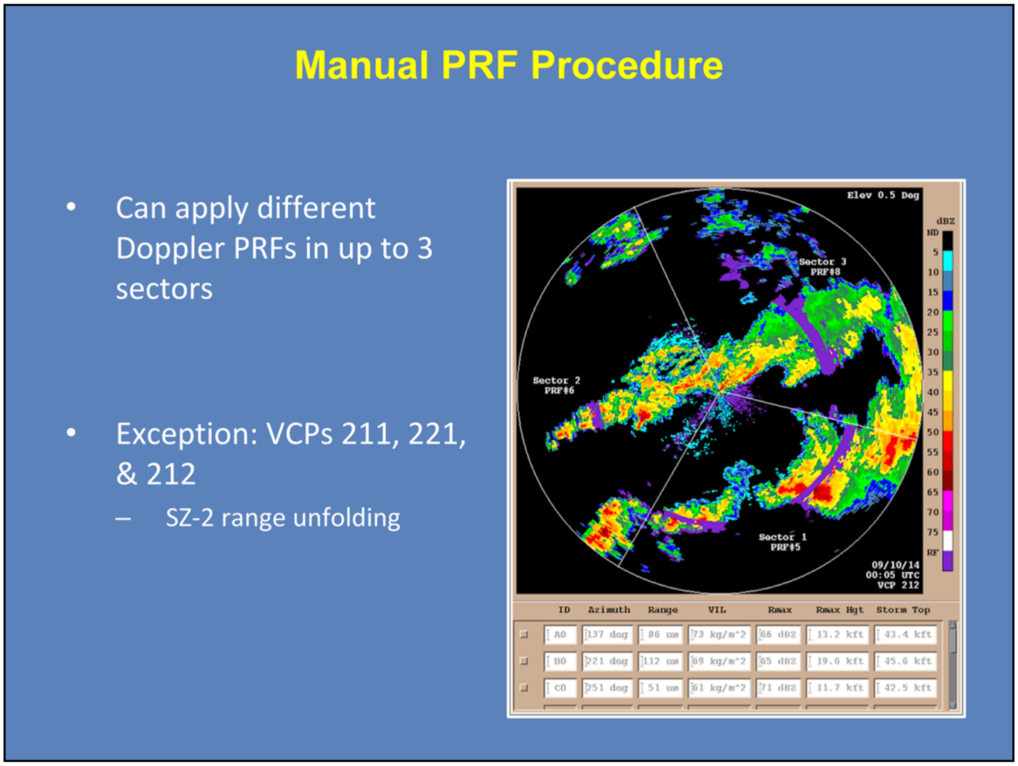

Most VCPs allow forecasters to specify up to three different Doppler PRFs in configurable sectors. You can perform this configuration within the display using a variety of mouse clicks. Your goal should be to have the least concealment of the velocity data by range folding for the strongest storms. NOTE: the SZ-2 VCPs (211, 212, and 221) don't allow users to sectorize Doppler PRFs as they use the SZ-2 range unfolding technique to minimize range folding.

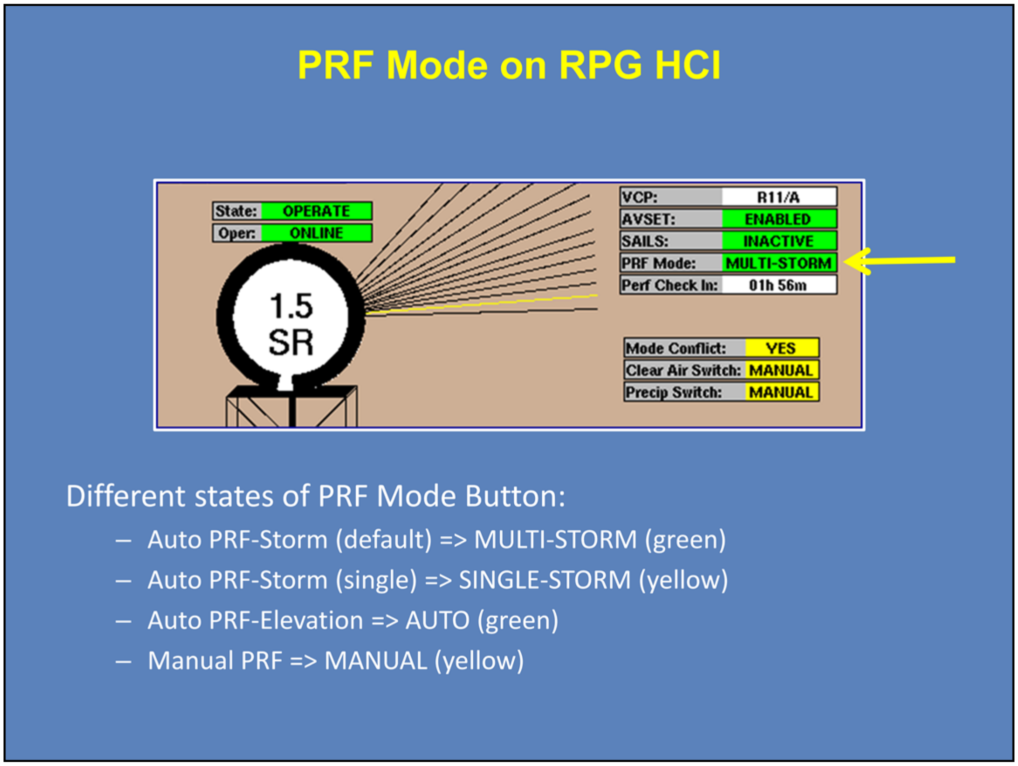

Four possible states may appear on the PRF Mode button of the RPG HCI depending on which PRF control mode you have selected:

1) For the default Auto PRF-Storm, the button reads “MULTI-STORM: on a green background like the example shown,

2) For the single-storm Auto PRF-Storm, the button reads “SINGLE-STORM” on a yellow background,

3) For Auto PRF-Elevation, the button reads “AUTO” with a green background, and

4) For Manual PRF, the button reads “MANUAL” with a yellow background.

Back to Top

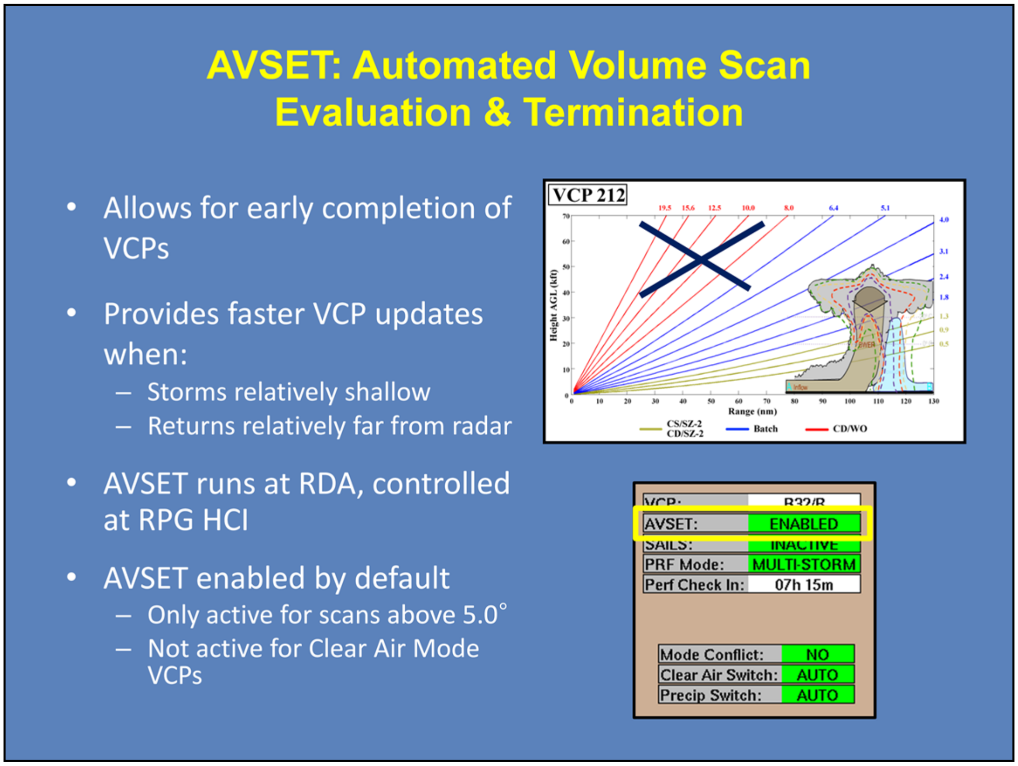

Now let’s move on from PRFs to a tool called AVSET. AVSET stands for the Automated Volume Scan Evaluation and Termination. This tool can be activated on precipitation mode VCPs to allow for the early completion of volume scans when weather return is limited to the lower elevation angles. AVSET meets a need for faster VCP update times, but only when weather isn’t particularly deep and the returns range is relatively far from the radar.

This graphic shows an example of how AVSET works. Once the radar starts scanning above 5 degrees, it encounters an elevation angle with no returns. When AVSET is enabled, the current volume scan would end early and the next volume scan would begin. AVSET runs at the RDA, but its status is controlled at the RPG HCI. AVSET is enabled by default. Since AVSET begins checking the areal coverage of returns above , you can enable it for clear air mode VCPs, but it will not terminate any volume scans early.

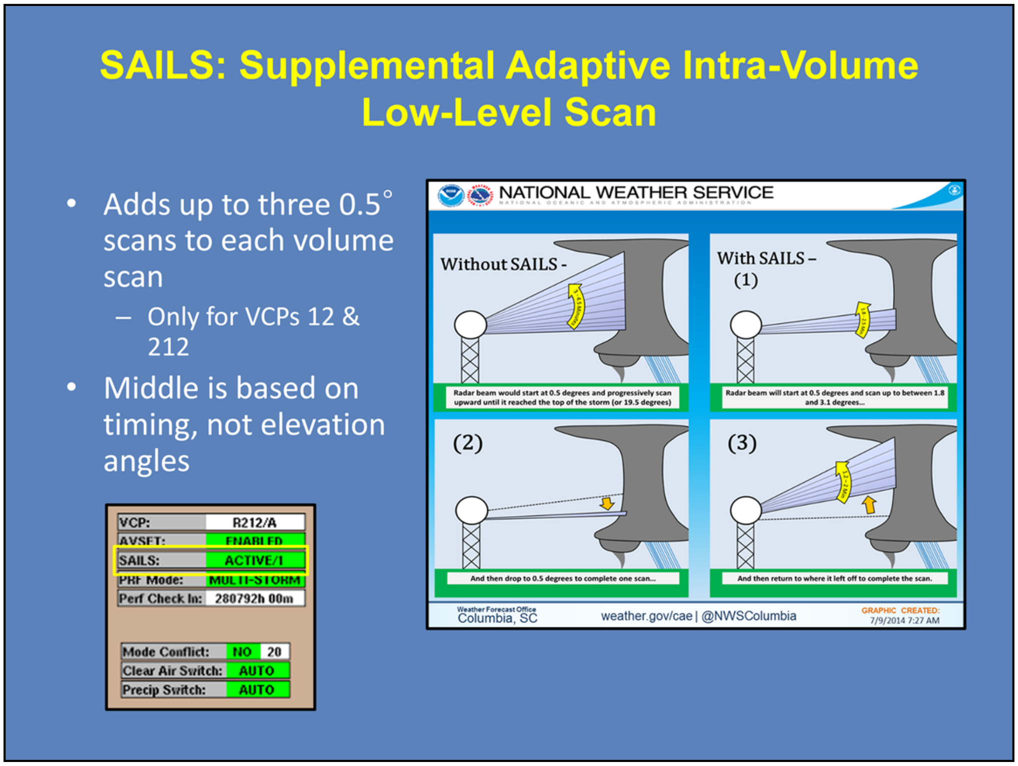

The Supplemental Adaptive Intra-Volume Low-Level Scan (SAILS) is another significant tool that makes VCPs dynamic. On VCPs 12 and 212, SAILS adds up to three additional 0.5° scans in the “middle” of the volume scan when active. You can tell how many SAILS scans you will get when its enabled by the number that appears after “ACTIVE/” in the RPG HCI. The “middle” of the volume scan is defined not by elevation angle, but by uniformly spacing the low-level scans in time. With AVSET enabled, the SAILS scan times change further because the length of the volume scan is dynamic.

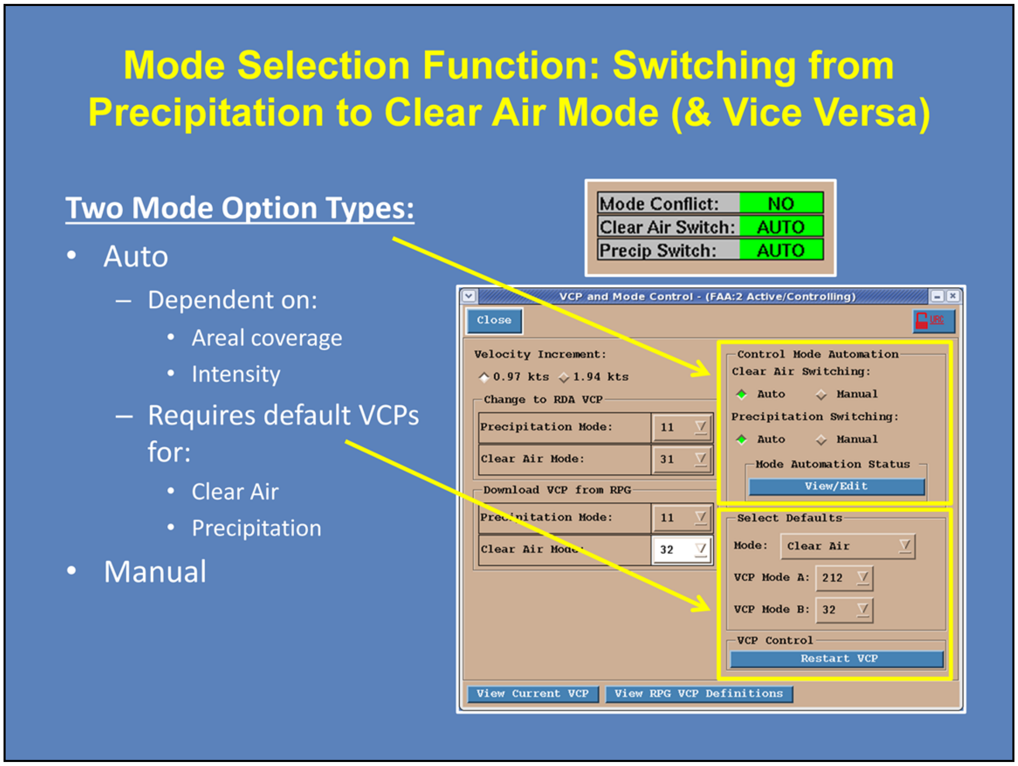

Another important function is the Mode Selection Function, which controls the transitions between Clear Air and Precipitation modes. To access the controls for the Mode Selection Function, we need to go back to the VCP and Mode Control window. Both modes have two control options: Auto and Manual. The automatic method decides when to switch based on the areal coverage of returns. When in automatic mode, you also need to specify a default VCP for both Clear Air and Precipitation modes. Usually, changes to the default VCPs are done seasonally, as needed, and requires URC password permission. The manual method requires more oversight and communication from shift to shift to ensure switches between the two operational modes occurs as needed.

Back to Top

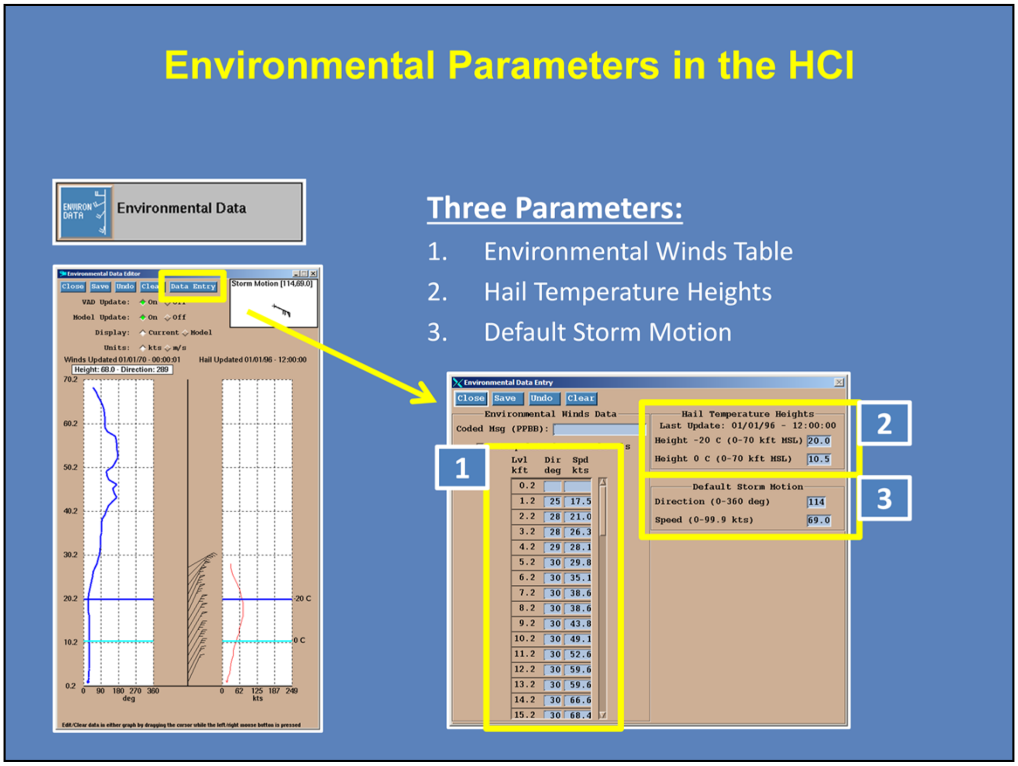

The final RPG functionality we will discuss involves the environmental data used by the radar. Several algorithms used by the radar require environmental data. To monitor environmental data used by the RPG, select the “Environmental Data” button on the right-hand side of the RPG HCI. To monitor and configure the environmental data, you need to select the “Data Entry” button. Three different environmental parameters can be accessed from the Environmental Data Entry GUI here: Environmental Winds Table, Hail Temperature Heights, and Default Storm Motion. Let’s take a look at each in more detail.

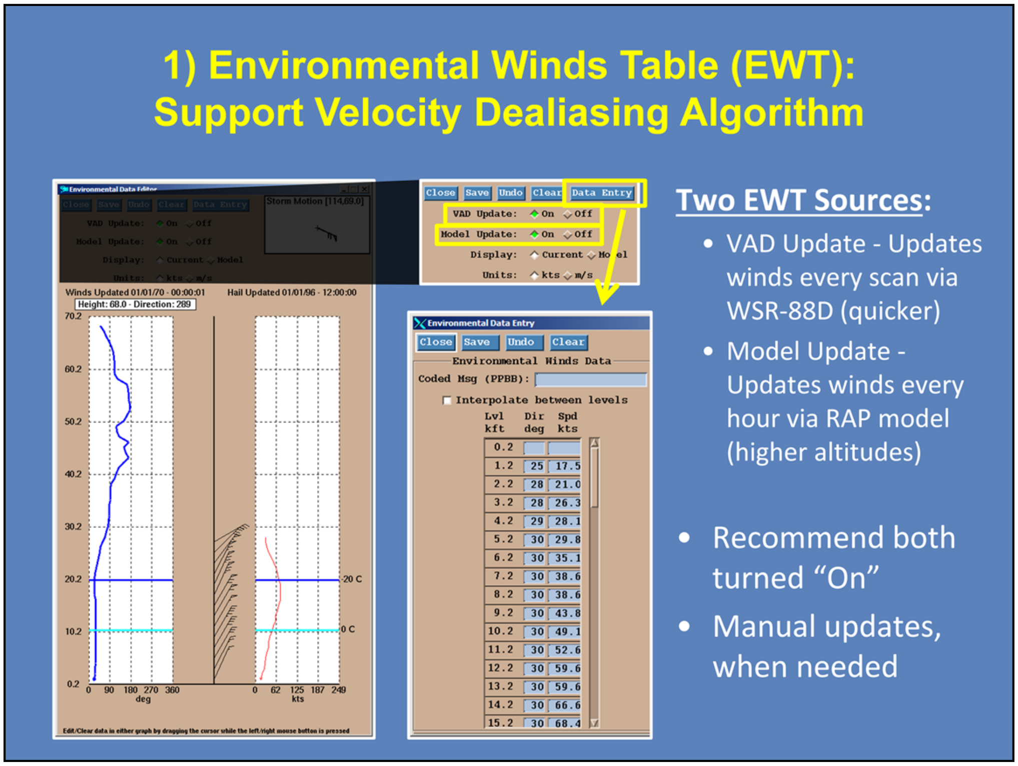

The Environmental Winds Table contains wind speeds and direction at multiple levels that supports the velocity dealiasing algorithm. The Environmental Data Editor window displays these data graphically. Control buttons at the top of that GUI allow for automatic updates from the VAD algorithm and model data. Toggling on both of these options is recommended for the best quality environmental wind data. The VAD update occurs every volume scan, but the model update from the RAP is hourly.

When you need to manually modify the environmental wind data, select the "Data Entry" button to open the Environmental Data Entry GUI. You can also click and drag on the Environmental Data Editor window. From the Environmental Data Entry GUI, wind speed and directions are separated by 1000 ft intervals. Once you have edited the values, users just save and close the window.

Back to Top

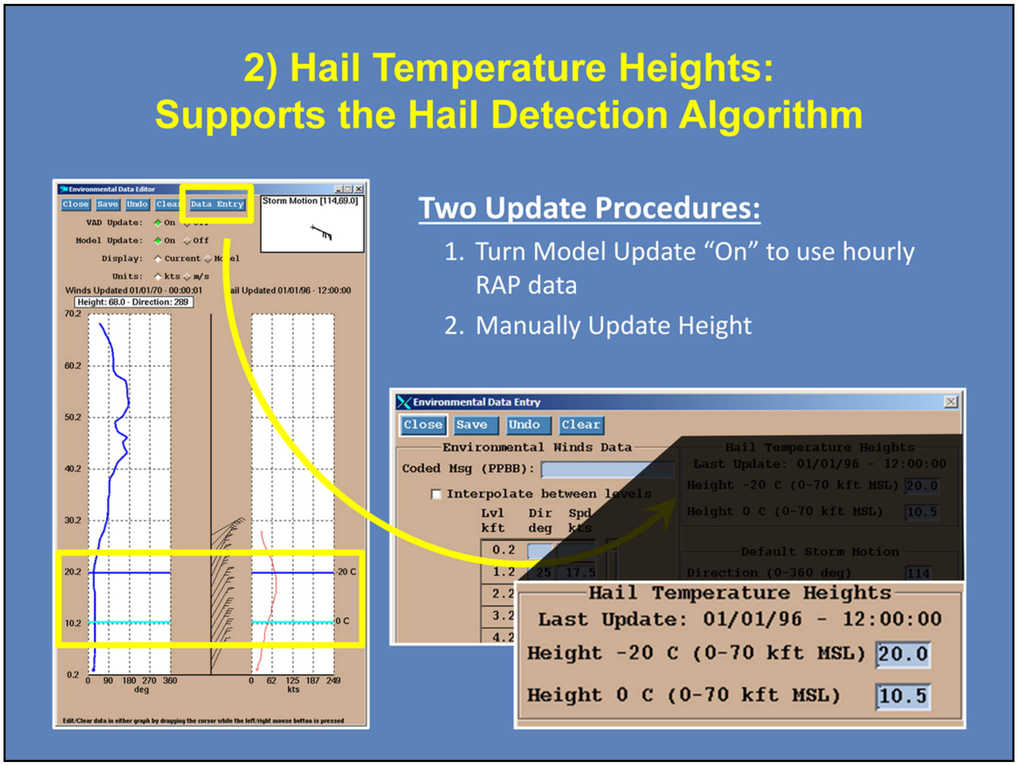

The Hail Temperature Heights has two parameters that support the Hail Detection Algorithm: Freezing level height and the height of the -20° C isotherm. The RPG populates these parameters in two ways:

1) If Model Update is "On", the RPG updates the values from RAP hourly temperature analyses.

2) Otherwise, the radar operator can manually enter the values based on nearby sounding data.

To manually edit these values, click on the "Data Entry" button. The parameters can be edited using the text fields in the Hail Temperature Heights panel of the Environmental Data Entry GUI. You can also click and drag on the lines displayed in the Environmental Data Editor window, but this method is much less precise.

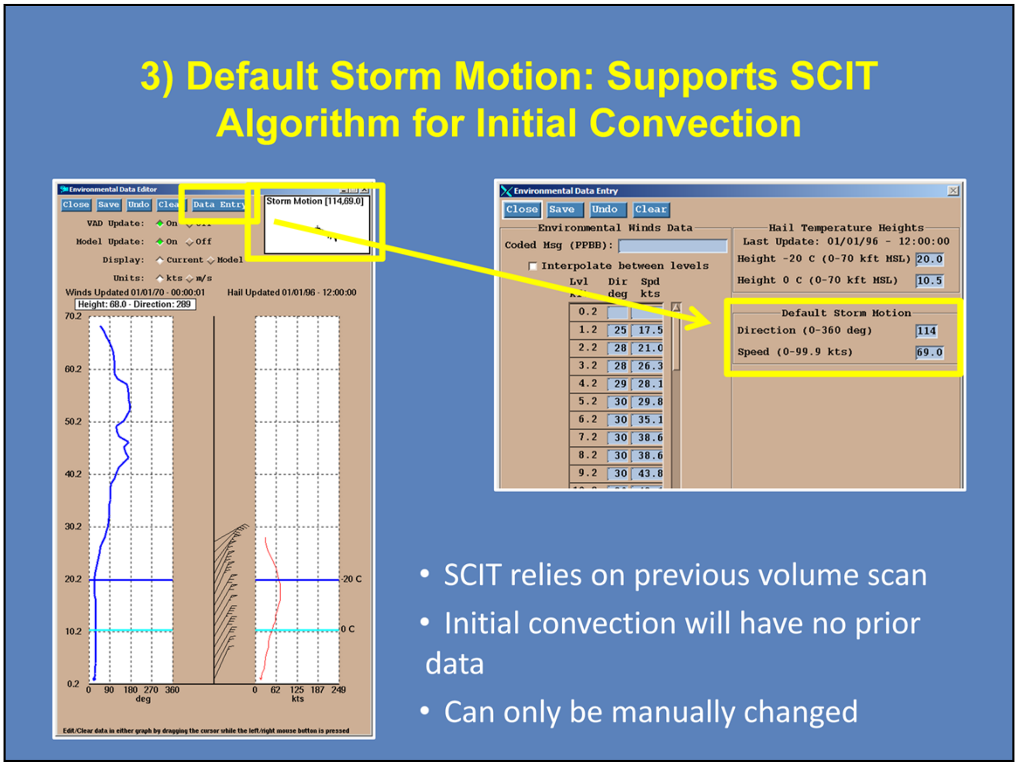

The Default Storm Motion parameters support the Storm Cell Identification and Tracking (or SCIT) algorithm. The wind speed and direction defined here helps the algorithm estimate storm motion. Normally, the SCIT algorithm can figure a cell's storm motion from data collected on previous volume scans. So, these parameters help the algorithm primarily with newly identified cells.

The RPG has no automated way to update the values, so you need to routinely monitor the values and update them, as needed. To make a change, you will once again select the "Data Entry" button and modify the text fields in the Default Storm Motion panel. Don't forget to click the "Save" button when finished.

Summary

- The RPG HCI provides three options to control Doppler PRF:
  1. Auto PRF-Storm
  2. Auto PRF-Elevation
  3. Manual PRF
- Two methods to switch between operational modes:
  1. Auto
  2. Manual
- Three environmental parameter sets to monitor/edit at RPG:
  1. Environmental Winds Table (velocity dealiasing)
  2. Hail Temperature Heights (hail detection)
  3. Default Storm Motion (new storms in SCIT)

This lesson highlighted several functions radar operators can perform at the RPG HCI. Radar operators have three options control the Doppler PRF: Auto PRF-Storm, Auto PRF-Elevation, and Manual PRF. Each option has applicability depending the users' needs. You also have two methods for switching between the operational modes of the radar: Auto and Manual. Remember when using "Auto" that you need to define default VCPs for both Clear Air and Precip modes. Lastly, the RPG has three important environmental parameter sets to monitor (and sometimes edit): the Environmental Winds Table, the Hail Temperature Heights, and Default Storm Motion. These parameters ensure that several RPG-based algorithms run properly and directly impact data quality.

Back to Top

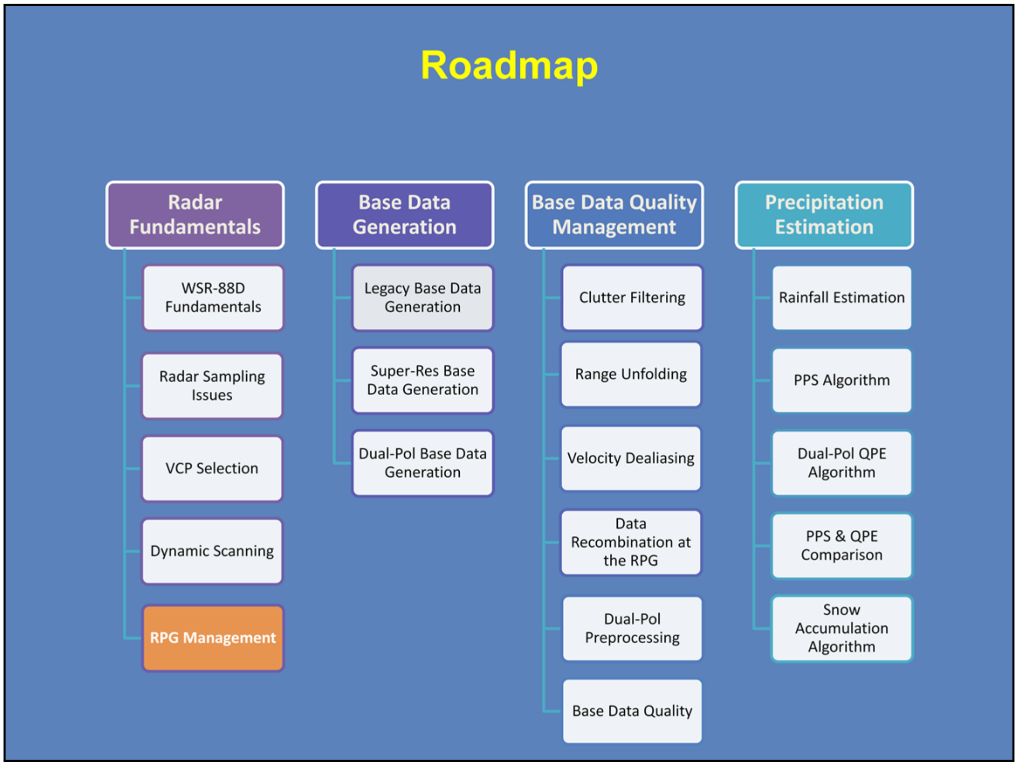

Thank you for your time. Good luck with the remaining lessons in the Principles of Radar topic of RAC!

Welcome to Legacy Base Data Generation.

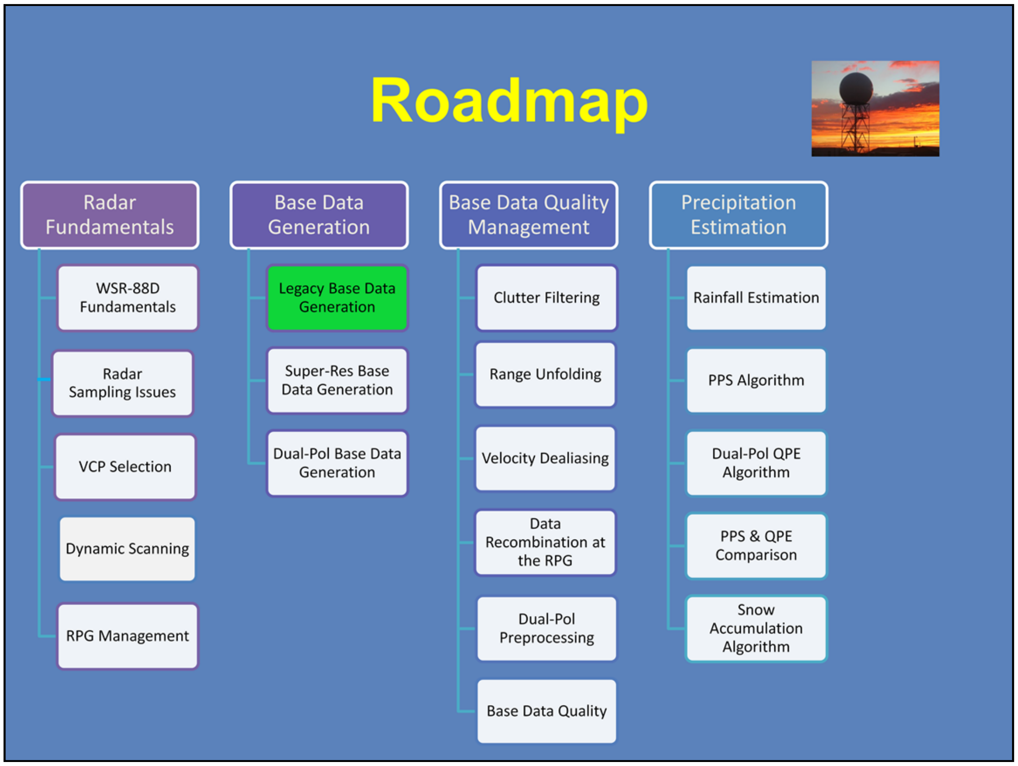

Here is the “roadmap” with your current location.

Back to Top

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.

Back to Top

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.

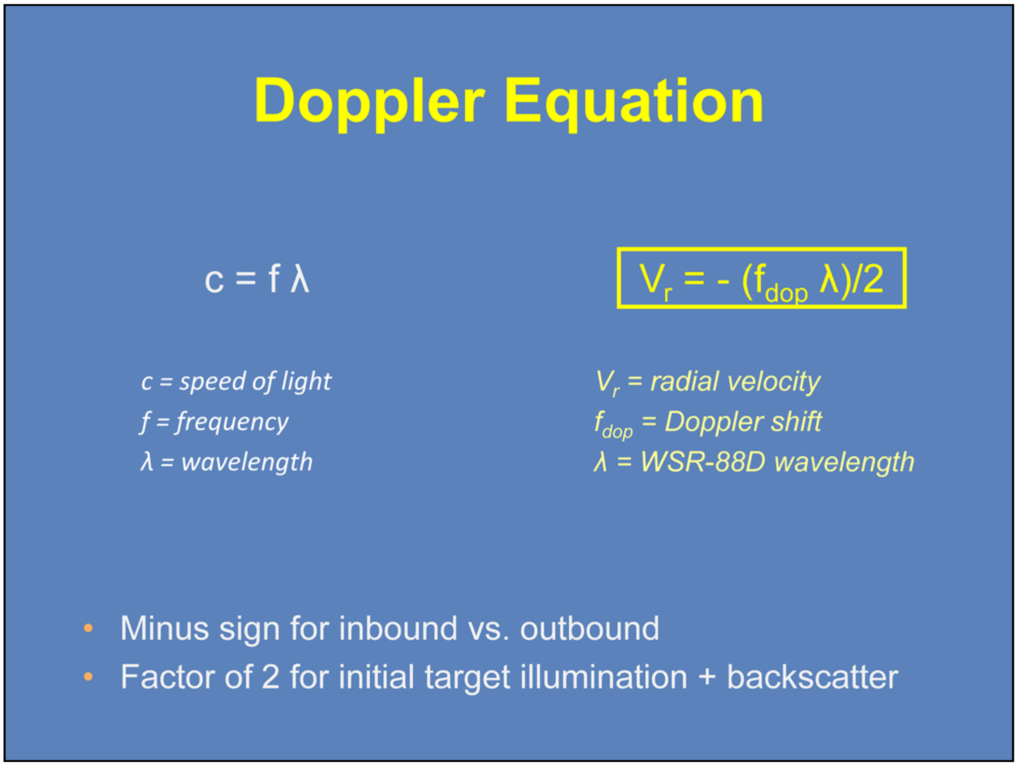

On the left is the basic relationship of frequency and wavelength to the speed of light. Based on that equation, on the right, is the relationship of the Doppler shift to radial velocity for Doppler weather radar (derivation not required!).

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.

Back to Top

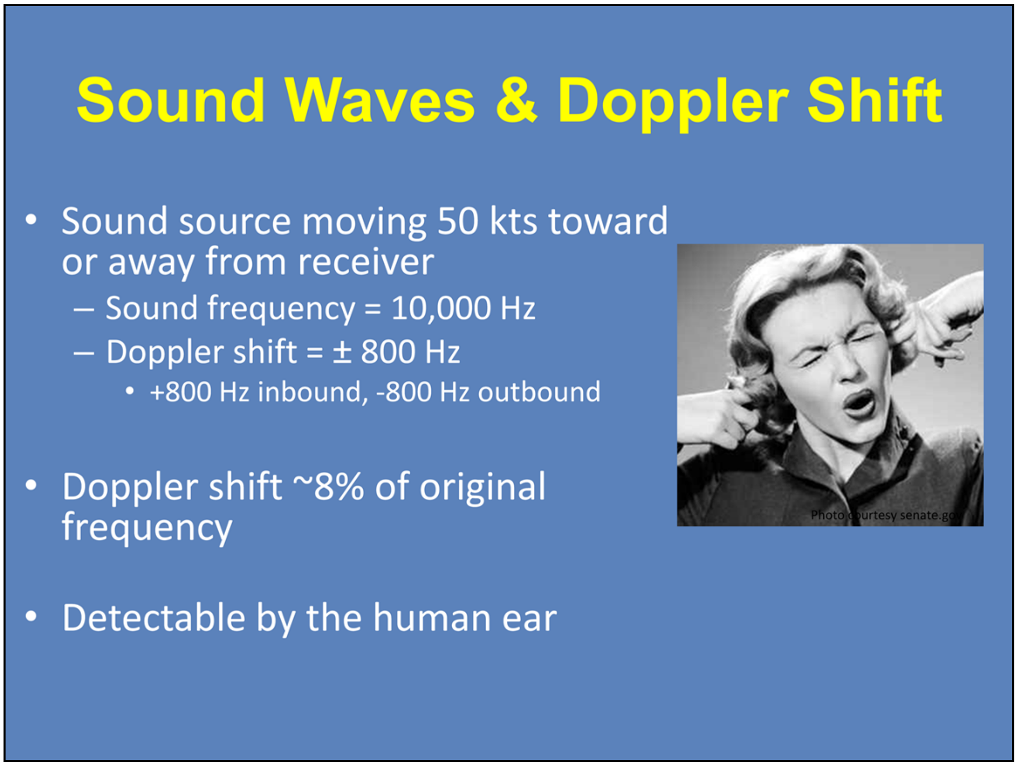

A common way to demonstrate the Doppler effect is with the change in pitch of the sound of a train or ambulance as it first moves toward you, then moves away. In this example with a speed of 50 kts, the frequency shift is 800 Hz, +800 when the sound source is moving toward you and -800 Hz when the sound 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!

Back to Top

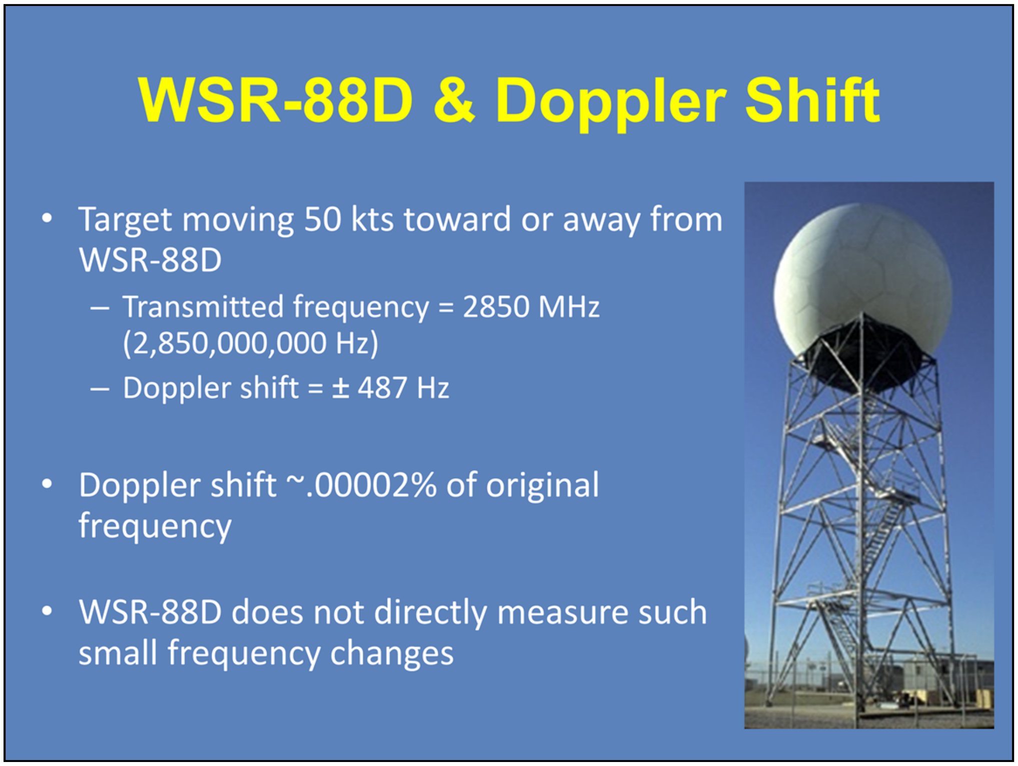

Now for the same target motion of 50 kts, but the listening device is the WSR-88D. We are no longer dealing with sound waves. Note that the frequency of transmission is very high compared to the frequency of a sound wave. The frequency shift is tiny compared to the original frequency…too small to be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

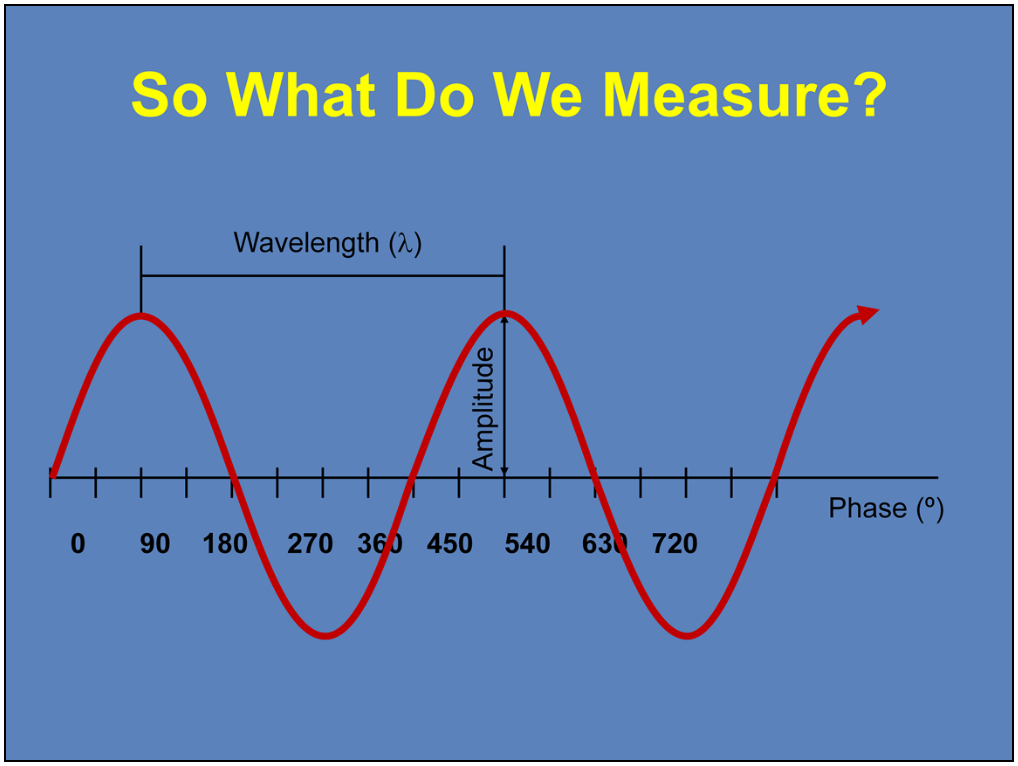

Since we don't measure the frequency shift, what else is there? Here is a reminder of the various characteristics of wave energy. The wavelength is the distance for one complete cycle, which is about 10 cm with the WSR-88D. The amplitude is the signal strength, which is directly related to reflectivity.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

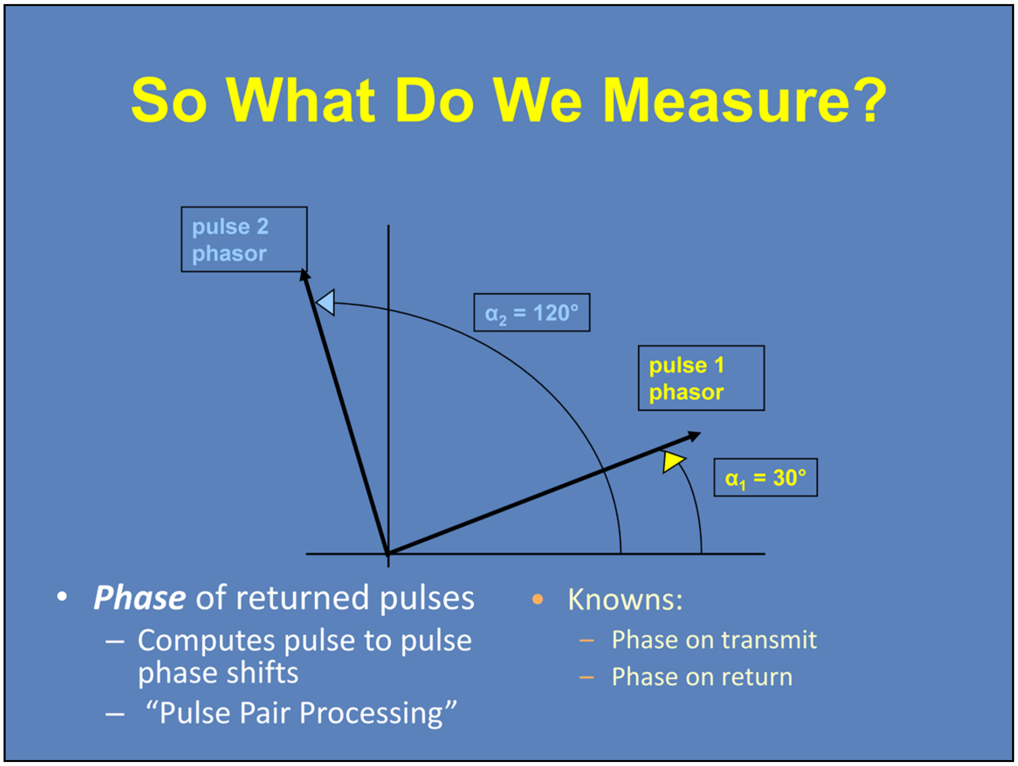

It turns out that the WSR-88D measures the phase of each returned pulse and is able to compare the phase values from one pulse to the next. The phase shift from one pulse to the next is directly related to the radial velocity. This technique is called Pulse Pair Processing.

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.

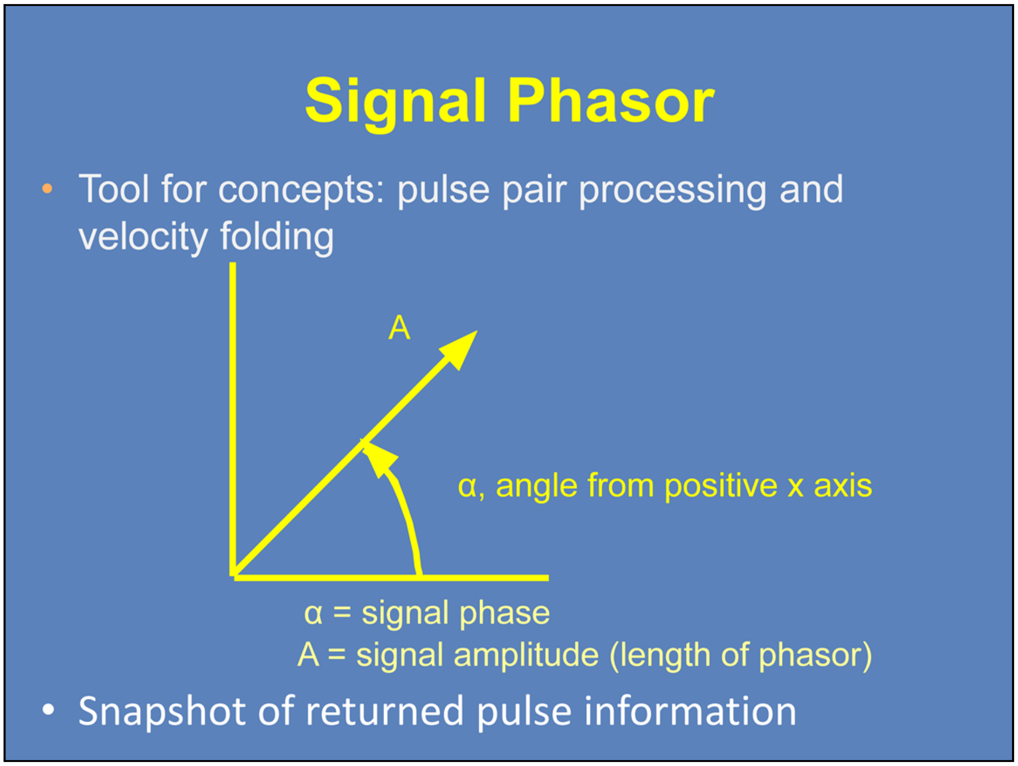

One way to represent the concept of a pulse pair phase shift is to use phasors. A phasor is a tool for temporary use only to support your understanding of one of the fundamental ambiguities with Doppler weather radar: velocity folding or aliasing. A phasor represents the necessary information from each 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.

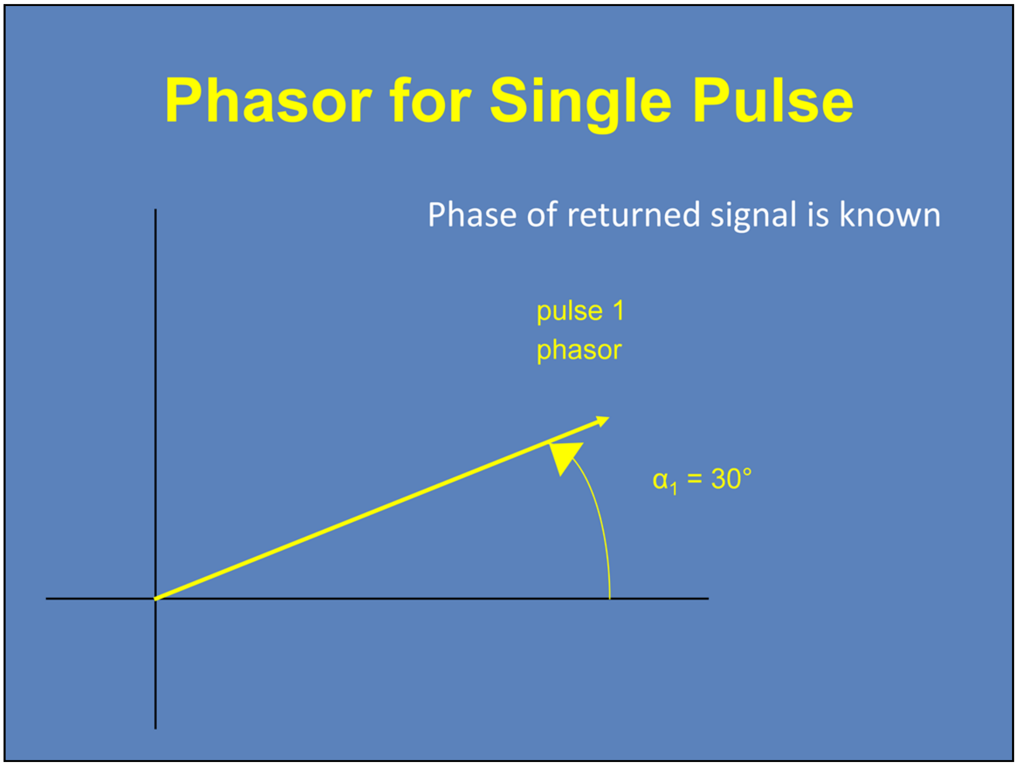

In this example, pulse 1 has been transmitted, has interacted with a target, and the returned signal has been processed. The phase value for pulse 1 is 30°.

Back to Top

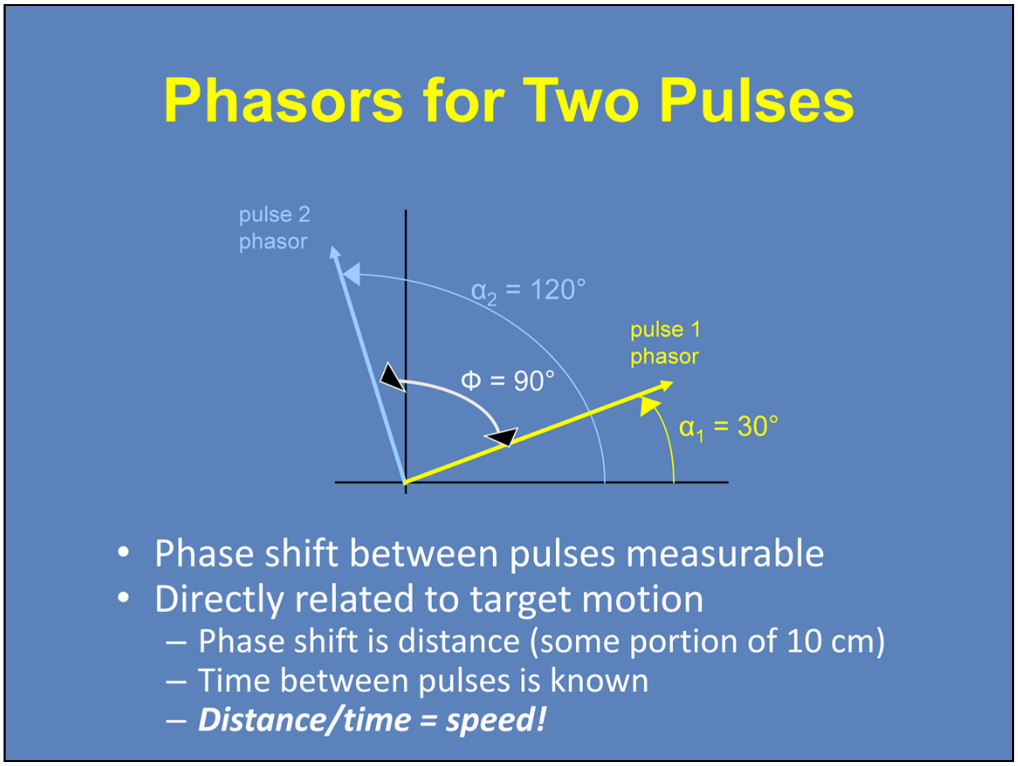

Using the pulse 1 phase of 30°, assume that the target is in motion and the phase value for pulse 2 is 120°. The angle between the two phasors (90°) is called the pulse pair phase shift.

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.

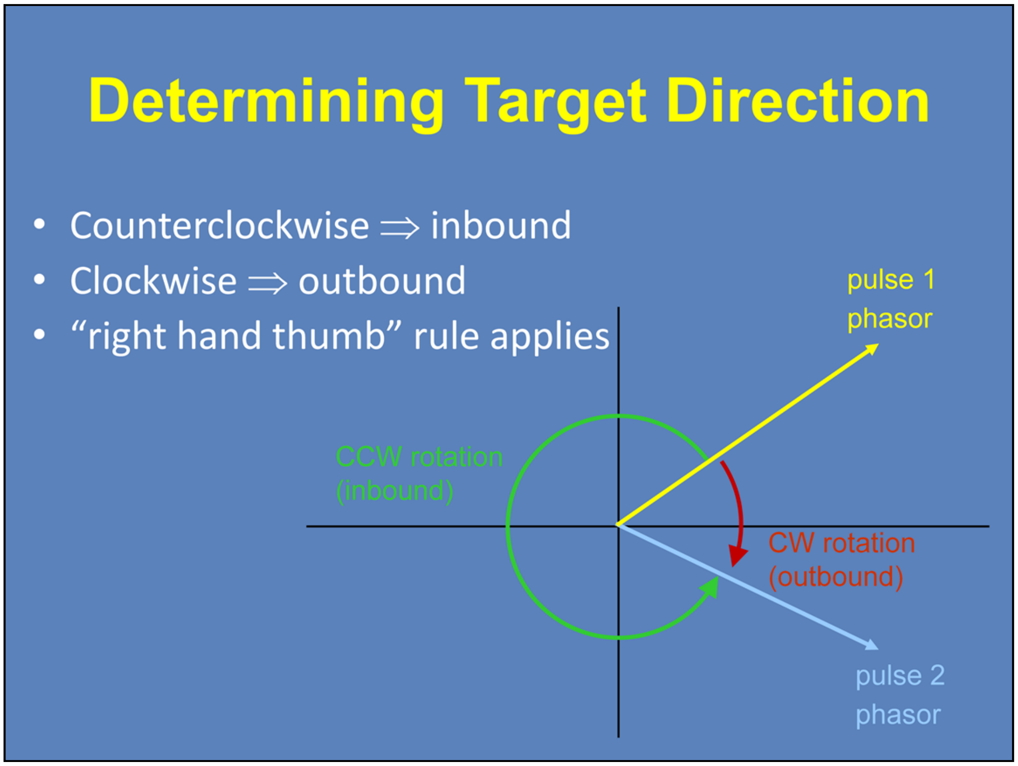

The target direction, inbound vs. outbound, is determined by the phasor rotation from pulse 1 to pulse 2. If clockwise, the direction is outbound. If the rotation is counterclockwise, the direction is inbound. You can also use the right hand thumb rule for the cross product of two vectors. Using the 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.

Back to Top

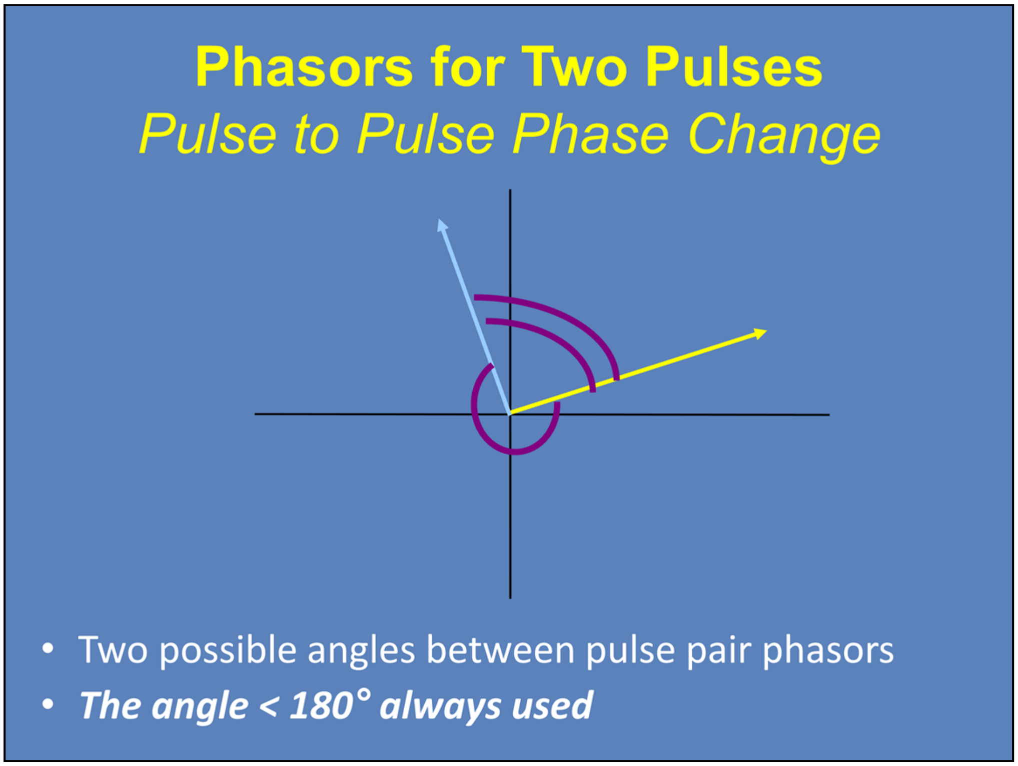

There are two phasors representing the information for two different pulses. Since these phasors are snapshots, some assumptions have to be made about what happened in between the two pulses. There are two possible angles between the phasors that represent these pulses. Which one is used? It is always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

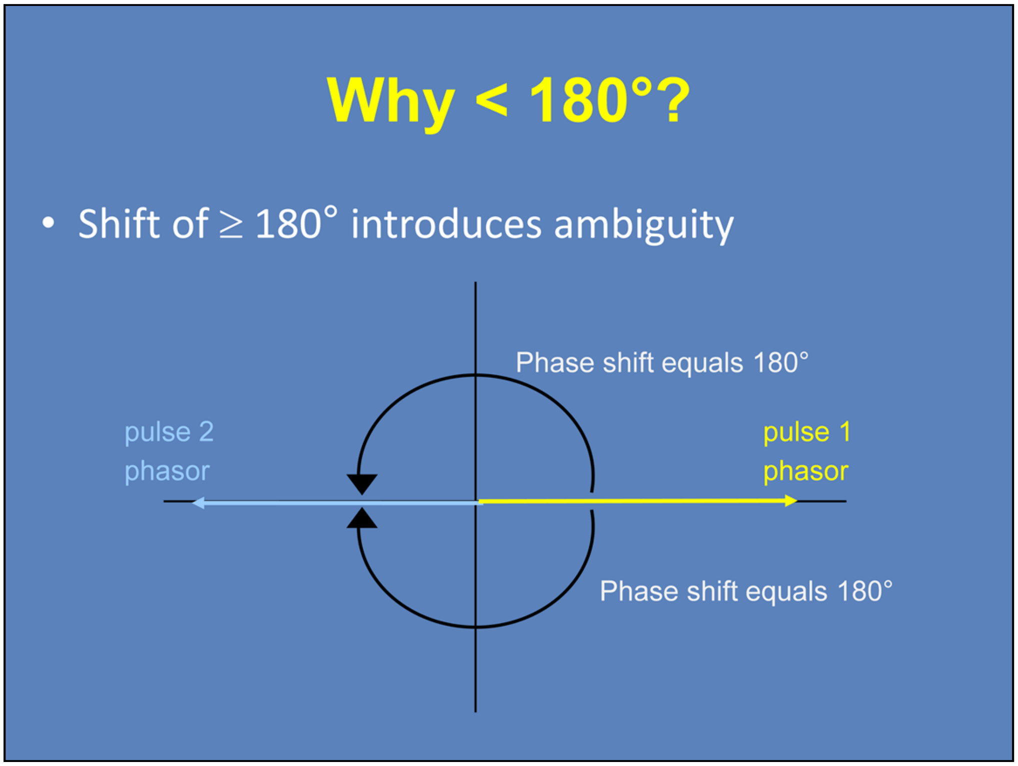

Why use the angle < 180°? A phase shift of exactly 180° introduces ambiguity…it is unknown which direction the phasor rotated to get from pulse 1 to pulse 2. If the target moves so much between pulses that the true phase shift ≥ 180°, there is ambiguity in determining the velocity. Stay tuned for how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

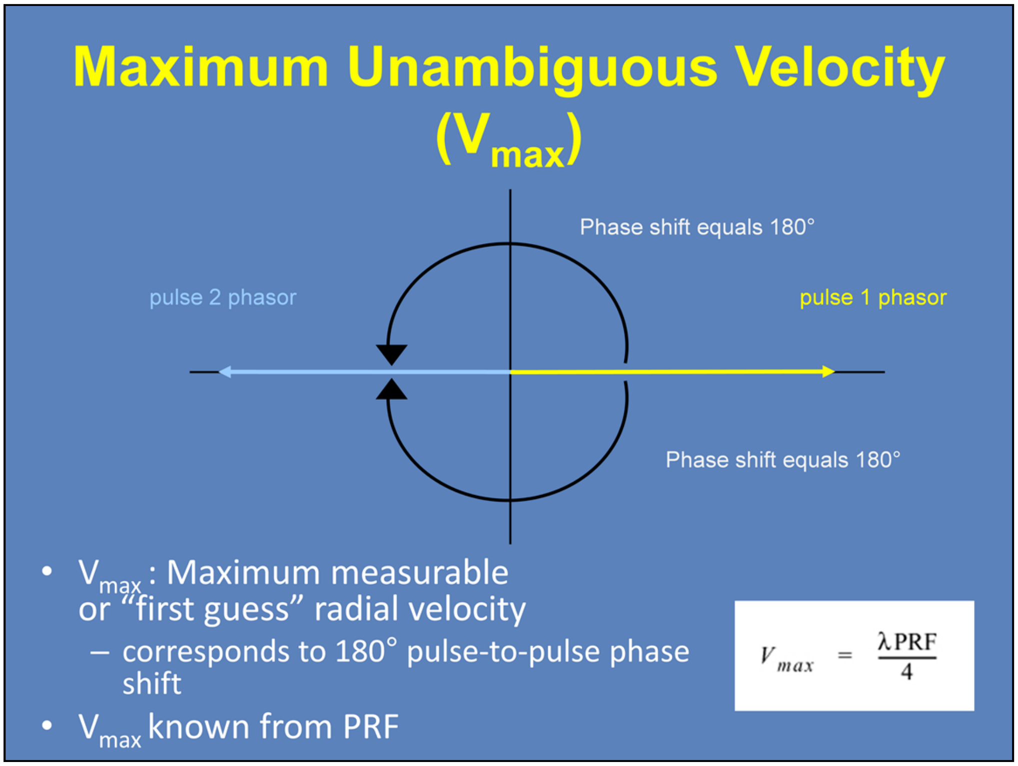

The maximum velocity that can be measured is called the maximum unambiguous velocity. It corresponds to a pulse pair phase shift of 180° (actually 179.99999…°), and is dependent on the pulse repetition frequency (PRF). With the WSR-88D, V_{max} values range from about 16 to about 64 kts.

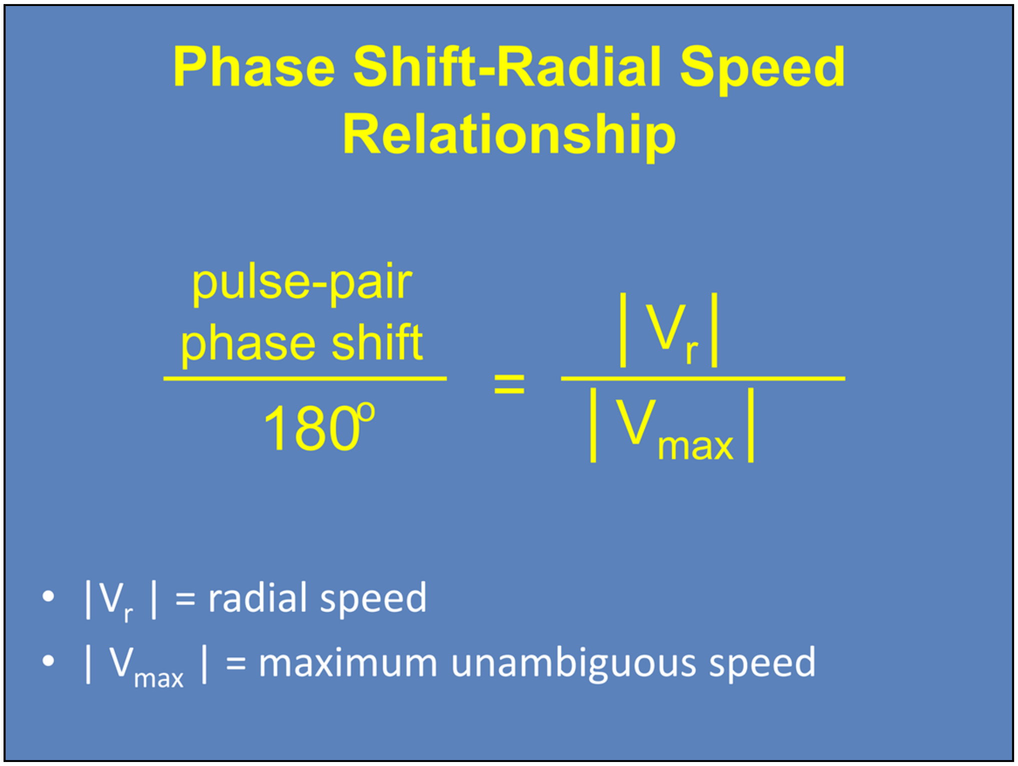

Once the pulse-pair phase shift and the V_{max} are known, computing the first guess radial speed is straightforward. That’s because the pulse-pair phase shift is some portion of the maximum shift of 180°, and the radial speed is that same portion of the maximum speed, or absolute value of V_{max}.

Back to Top

Phase Shift - Radial Speed

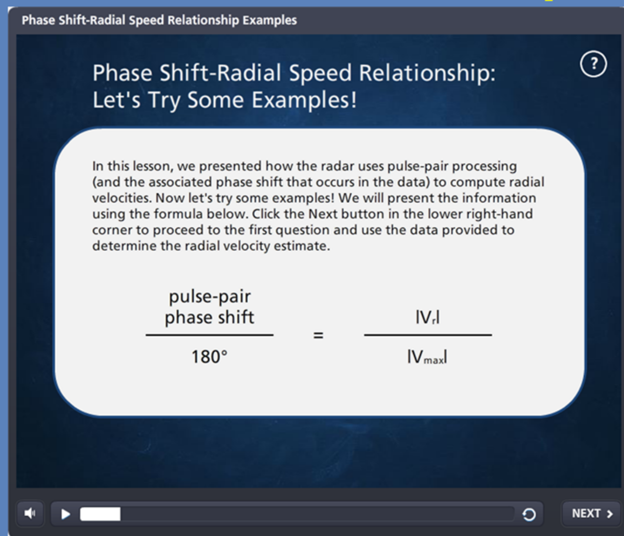

If no pop-up window appears that looks like the above, open a browser and go to:
http://www.wdtd.noaa.gov/courses/rac/principles/interactions/phaseshift-radialspeed/

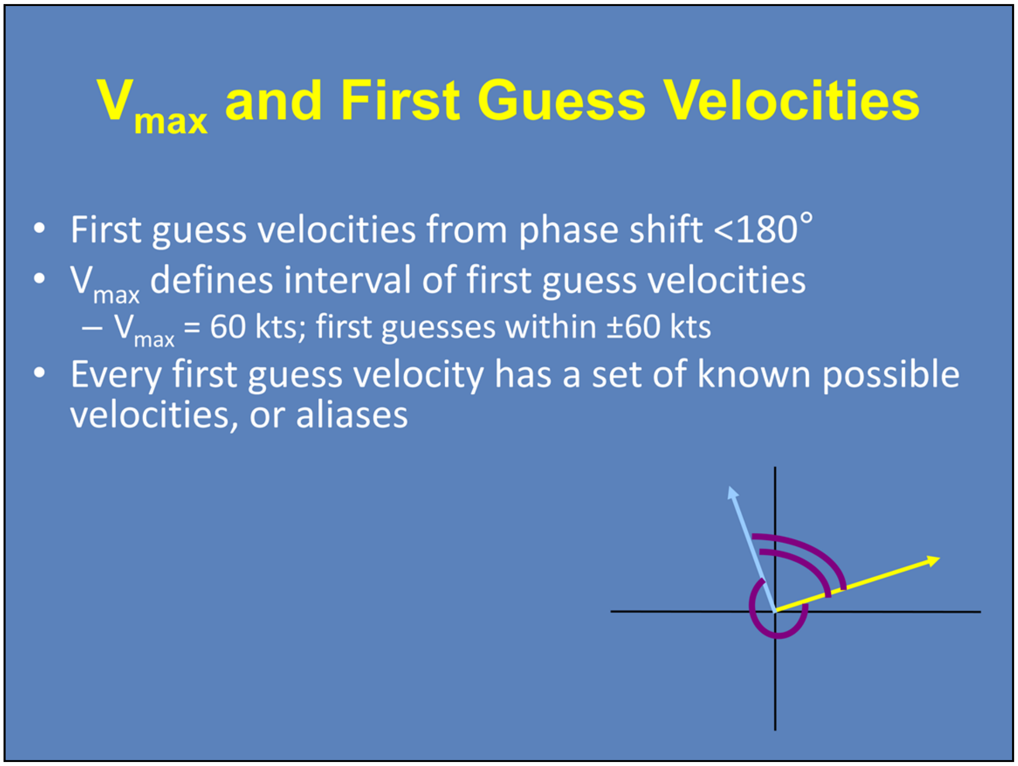

The previous examples were all based on the pulse pair phase shift that is <180°. A velocity that is based on this assumption is called the "first guess velocity". Since V_{max} is associated with 180°, V_{max} then defines an interval of first guess velocities. For example, when V_{max} = 60 kts, the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

Back to Top

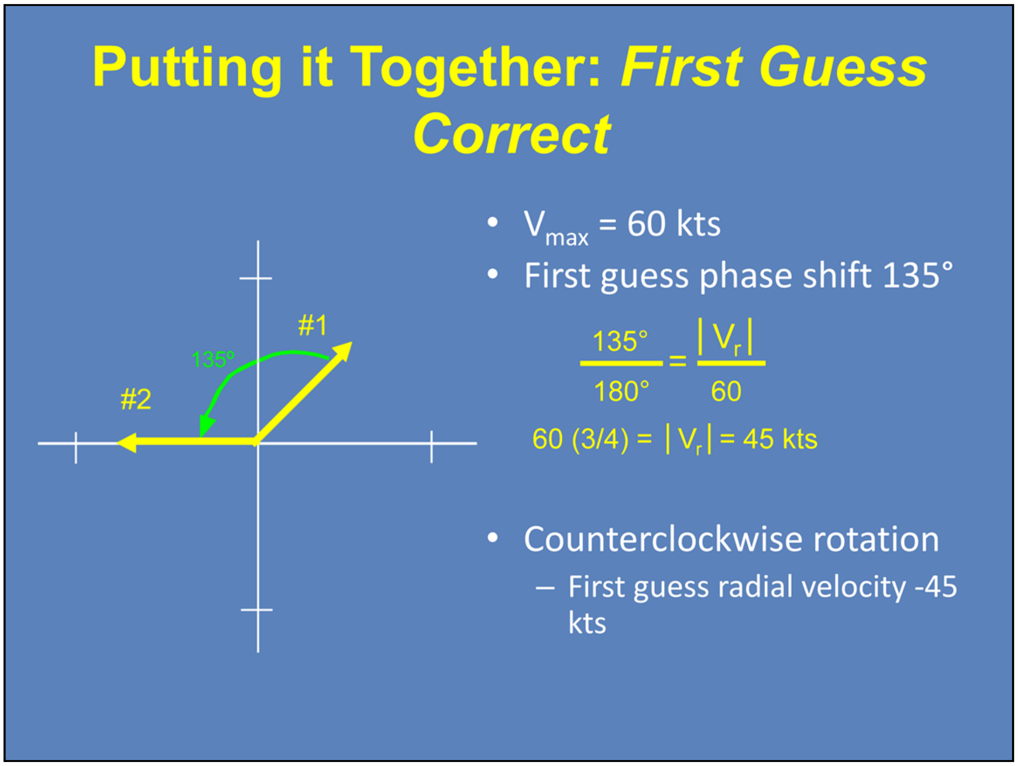

This example combines the concepts of pulse pair shift plus V_{max}, which gives us the first guess speed, along with phasor rotation which gives us target direction. In this case, we'll assume that the true phase angle between pulses is the one that is <180°, and that it represents the true target motion.

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

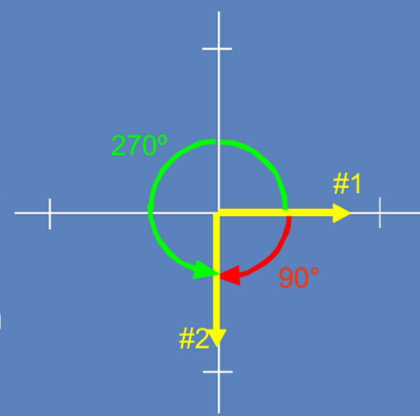

So what happens when the first guess velocity is not correct, i.e. the actual phase shift is >180°? The good news is that for every first guess velocity, there are other possible velocities which are known.

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.

Back to Top

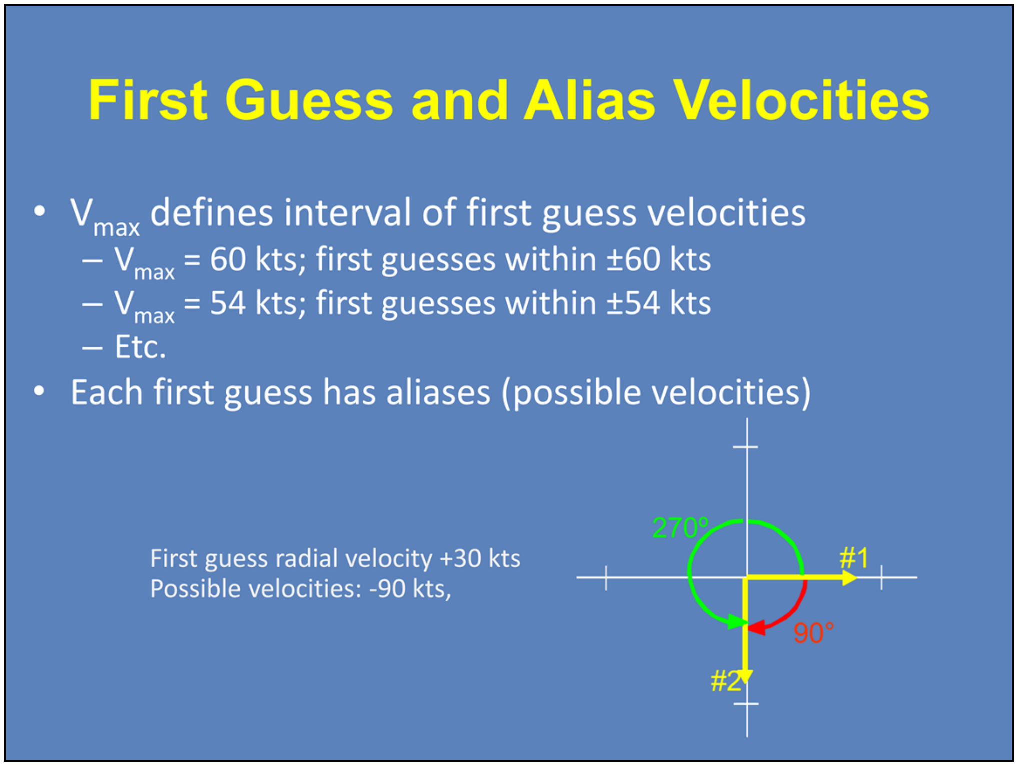

First guess velocities are based on the phase shift <180°, and V_{max} is the maximum unambiguous velocity, associated with a phase shift of 180. Each V_{max} thus defines an interval of first guess velocities. For example, for V_{max} = 60 kts, first guess velocities range from -60 kts to +60 kts, for 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).

Back to Top

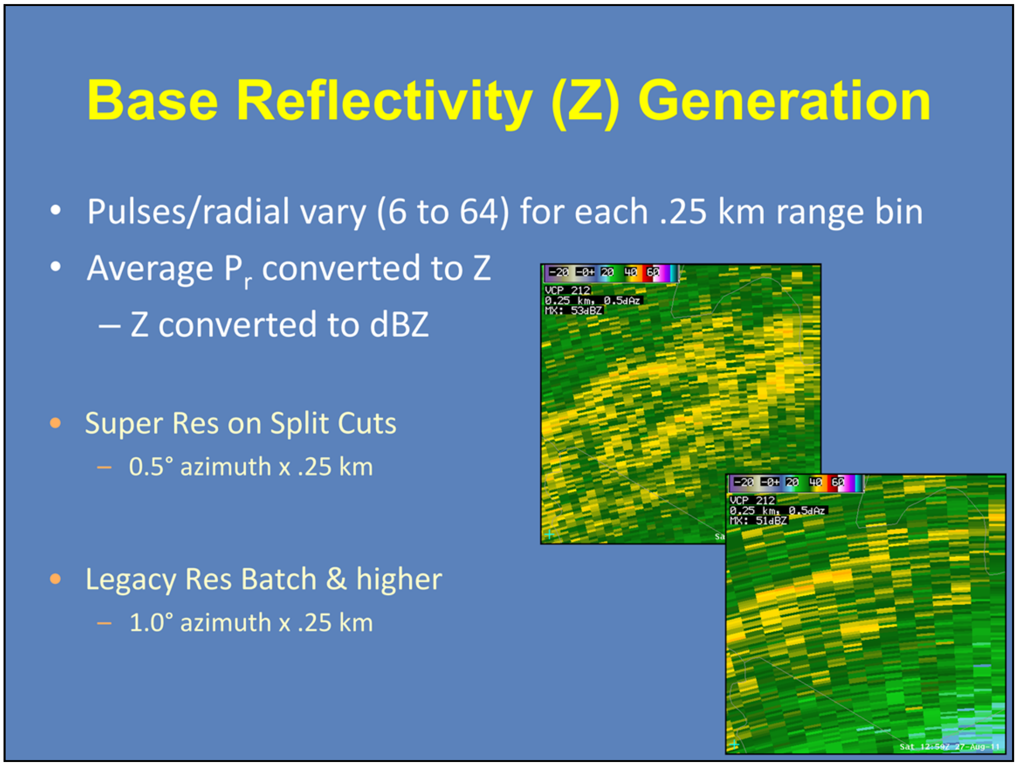

For each range and azimuth, there are multiple pulses used to generate base reflectivity. The returned power for these pulses is averaged for each .25 km range bin, then converted to Z using the Probert-Jones radar equation. The Z value is next converted to dBZ for product generation. For the Split Cuts, 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.

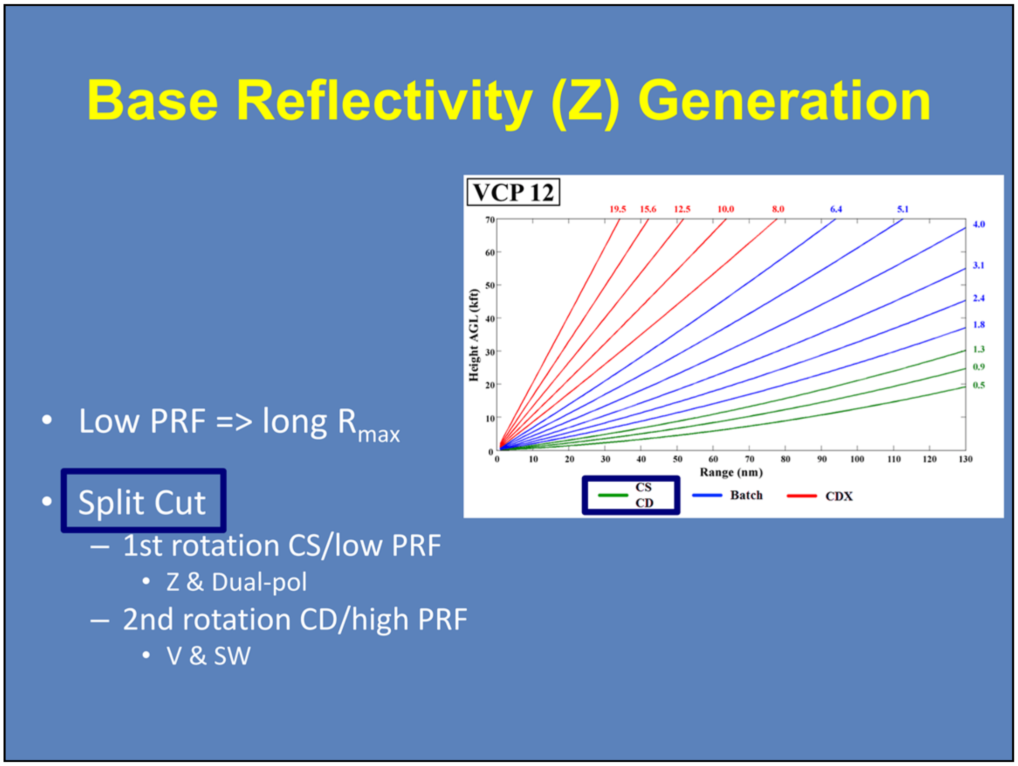

A low PRF provides a long R_{max} and is used for Reflectivity data. Split Cut mode is used for the lowest 2 or 3 elevations for all the VCPs except VCP 121. Split Cut first uses one rotation in Contiguous Surveillance (CS), which is a low PRF mode. Base reflectivity and the dual-pol data are generated 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.

Back to Top

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.

Back to Top

Base Velocity (V) Generation

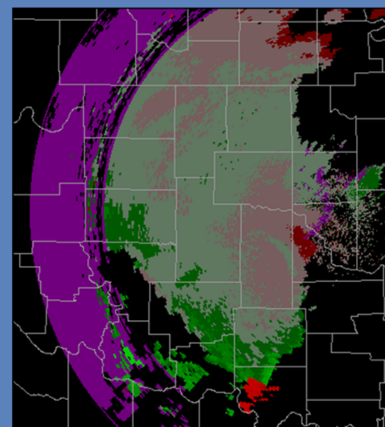

- Super Res on Split Cuts
  - 0.5° azimuth x .25 km
- Legacy Res Batch & higher
  - 1.0° azimuth x .25 km
- Maximum display range 162 nm
- Pulse pair processing
  - Phase changes between successive returned pulses averaged

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.

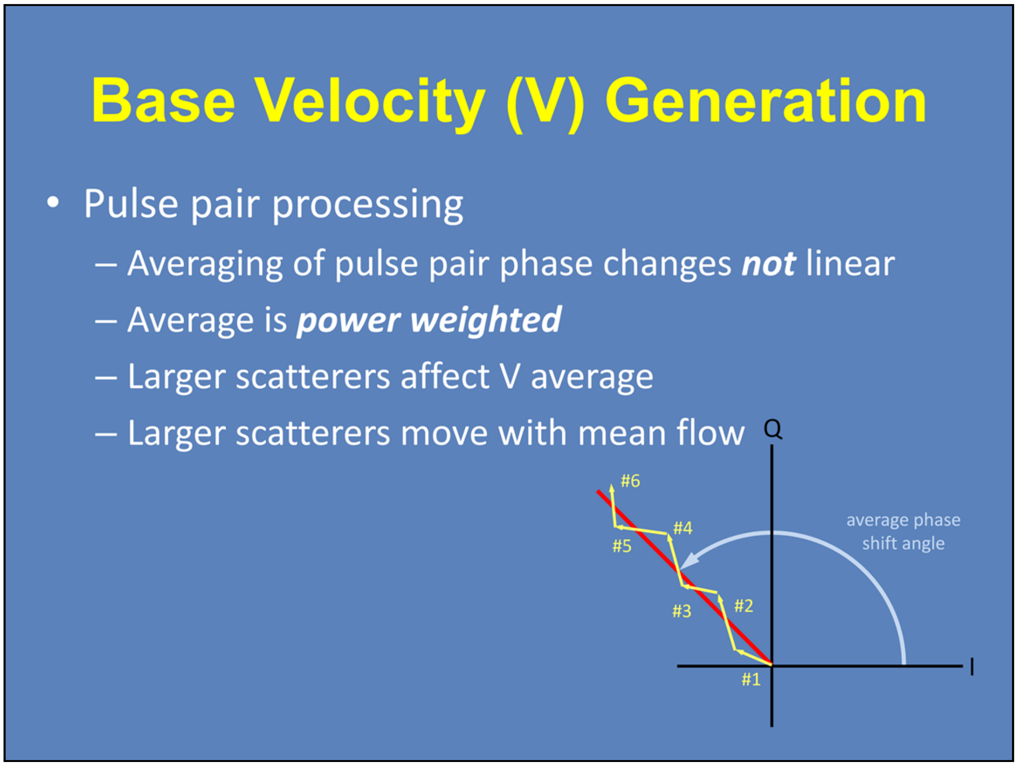

Instead of a linear average of pulse pair phase shifts, the average is weighted toward those pulses that return higher power. This means that the larger scatterers in the volume will have a greater influence on the velocity estimate, and the larger scatterers are more likely to move with the mean flow.

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

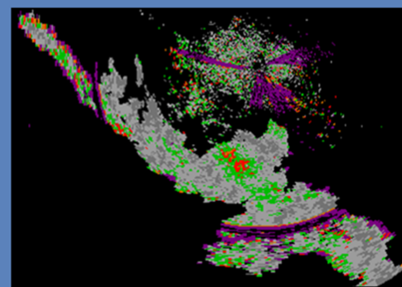

- Super Res on Split Cuts
  - 0.5° azimuth x .25 km
- Legacy Res Batch & higher
  - 1.0° azimuth x .25 km
- Maximum display range 162 nm
- Measure of velocity dispersion
  - Proportional to variation in wind speed/direction
- SWs typically high with
  - Boundaries, thunderstorms, high shear

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.

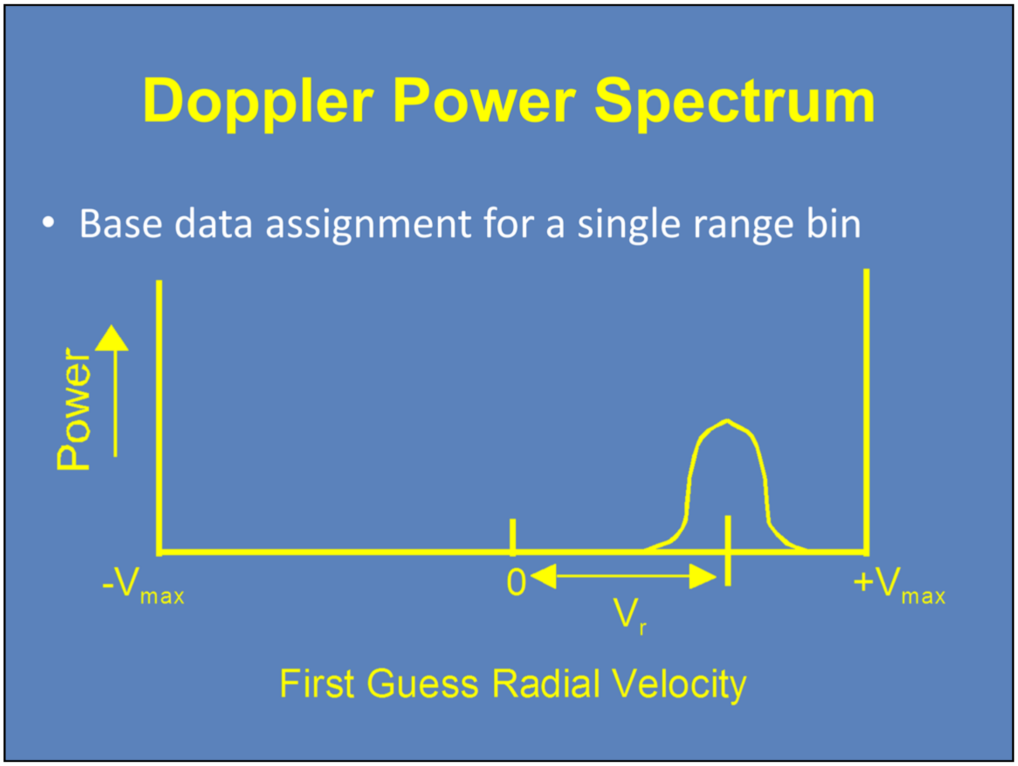

The Doppler Power spectrum represents the base data analysis process for a single range bin. The power and velocity information from a series of pulses are converted to points known at “spectral coefficients”. A bell curve is fit to these coefficients. The average returned power (thus reflectivity) is 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.

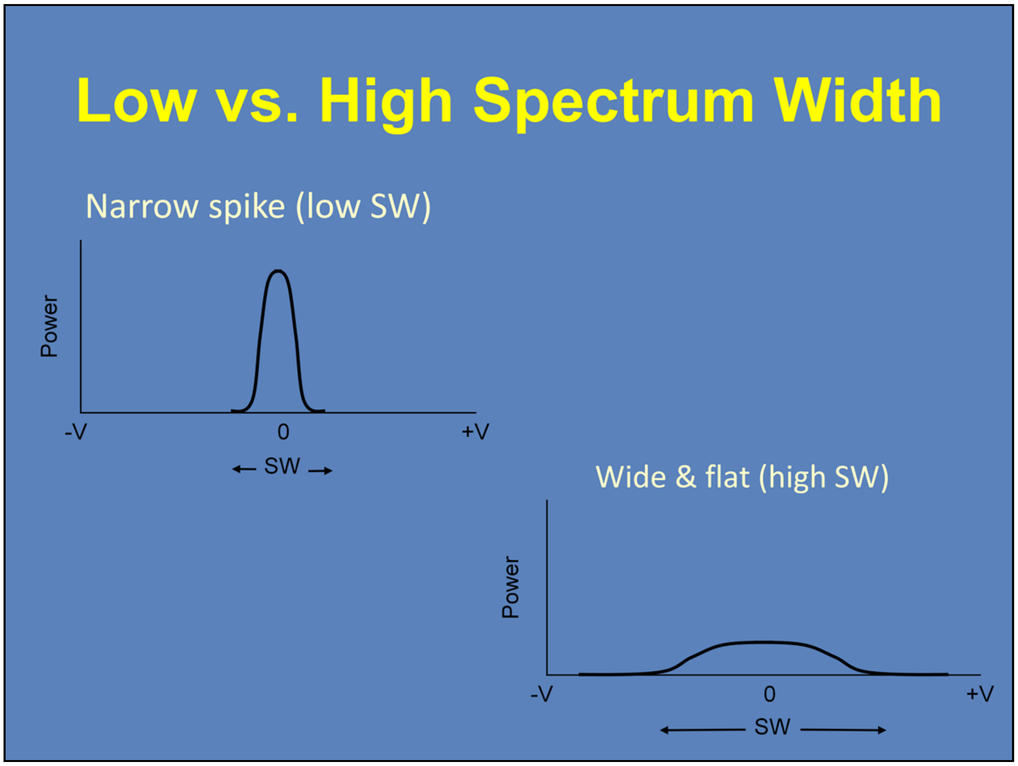

The magnitude of spectrum width will vary depending on the shape of the power spectrum. Returned pulses from ground clutter will likely have strong power and near zero velocity (upper left image). There is minimal variation in pulse pair phase shifts and the Doppler Power Spectrum curve is narrow and 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.

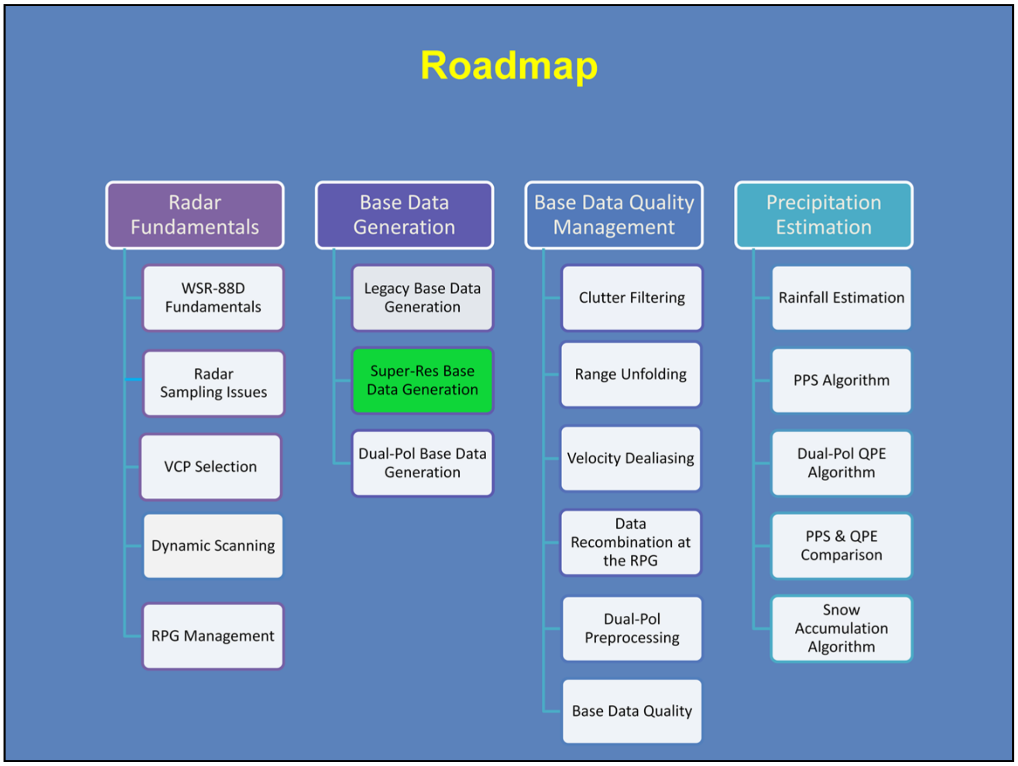

Here is the “roadmap” with your current location.

Back to Top

Super Res Base Data Generation: Objective

1. Identify the operational impacts of the signal processing techniques used to produce super resolution base data.

There is one objective in Super Resolution Base Data Generation.

Back to Top

Super Resolution Signal Processing

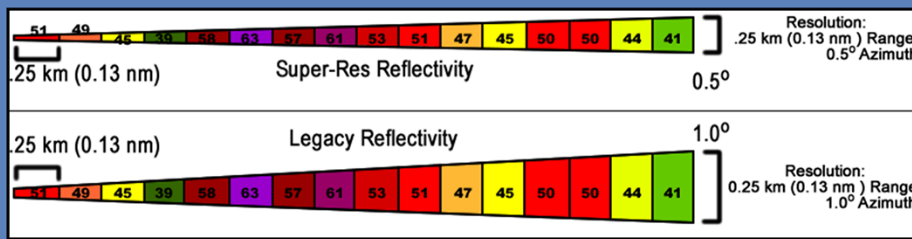

- Super Resolution: 0.5° azimuth for legacy base data on Split Cuts
- How to narrow azimuthal resolution to 0.5°?
  - Overlapping radials
  - Data windowing
- "Effective" beamwidth
  - Physical beamwidth (single pulse) plus the antenna is moving

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.

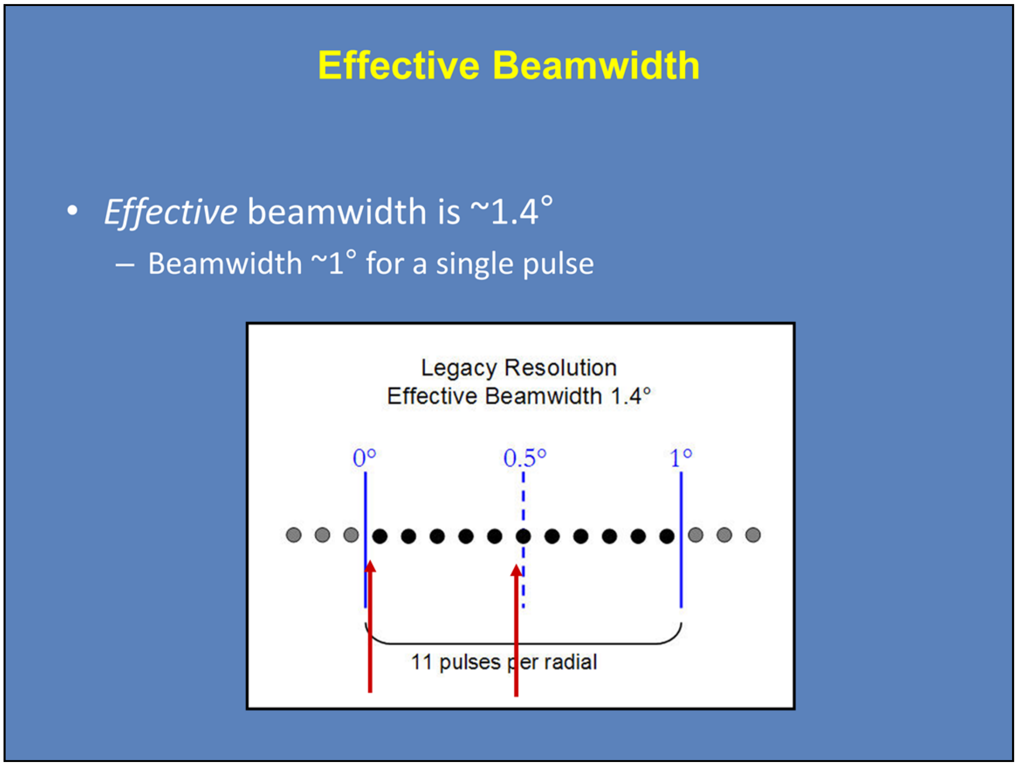

In order for a pulse to be used for the base data estimate for a radial, the beam centerline must be somewhere within that radial. In this example, each dot represents the location of the beam centerline for a single pulse. There are 11 pulses that fall within this simplified 1° radial. As the antenna 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°.

Back to Top

Overlapping Radials

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

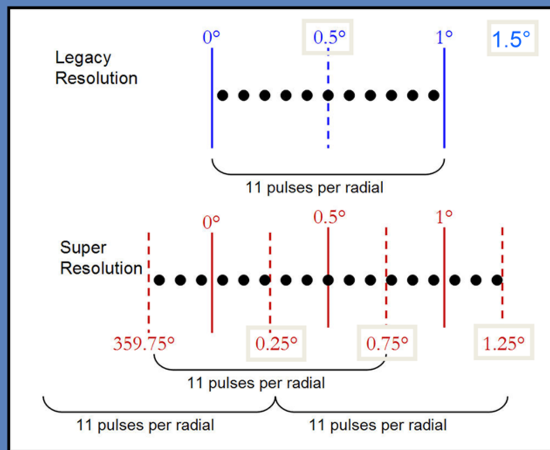

The process of overlapping radials is simply changing the definition of the center of each radial. The radial centers for legacy resolution are 0.5°, 1.5°, 2.5°, etc. The radial centers for super resolution are 0.25°, 0.75°, 1.25°, 1.75°, etc. However, since the number of pulses per radial cannot decrease 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.

Back to Top

Data Windowing

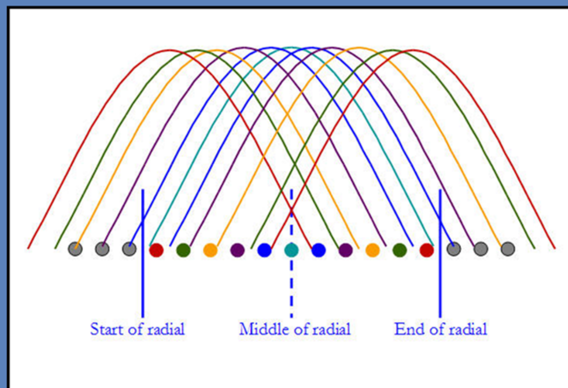

- Data Windowing
  - Apply weighting function to pulses that comprise base data for a radial
- Rectangular window: all pulses have equal weight

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.

Data Windowing

- Window with pulses close to center more weight than pulses away
- Narrows effective beamwidth
- Increases variance of estimate; data noisier

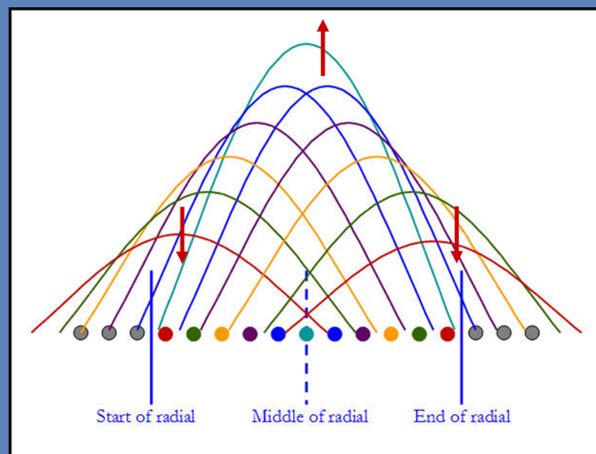

What is needed for super resolution is a window that narrows the effective beamwidth. This is accomplished by giving pulses near the center of the radial more weight, progressing to less weight applied to pulses away from the center. This technique meets the need of narrowing the effective beamwidth, 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.

Back to Top

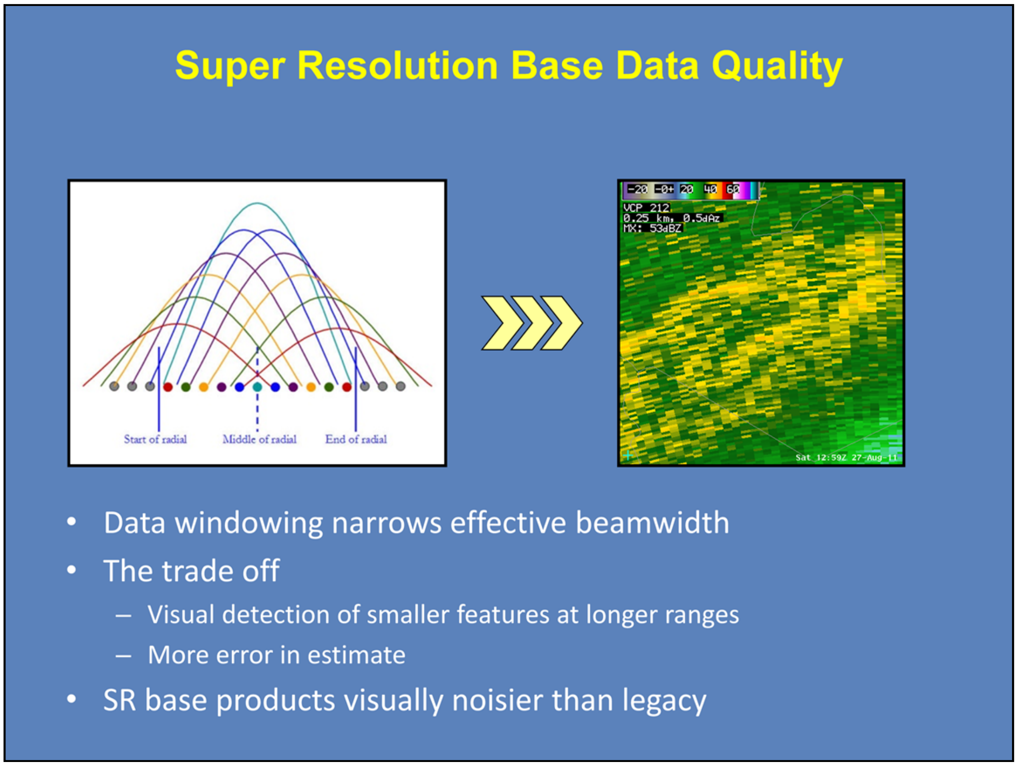

The 0.5° azimuthal resolution is obtained through signal processing techniques, specifically data windowing. This technique narrows the effective beamwidth, but there is a trade off. Smaller scale features are visually detectable at longer ranges with super resolution, but there is also more error in 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.

Back to Top

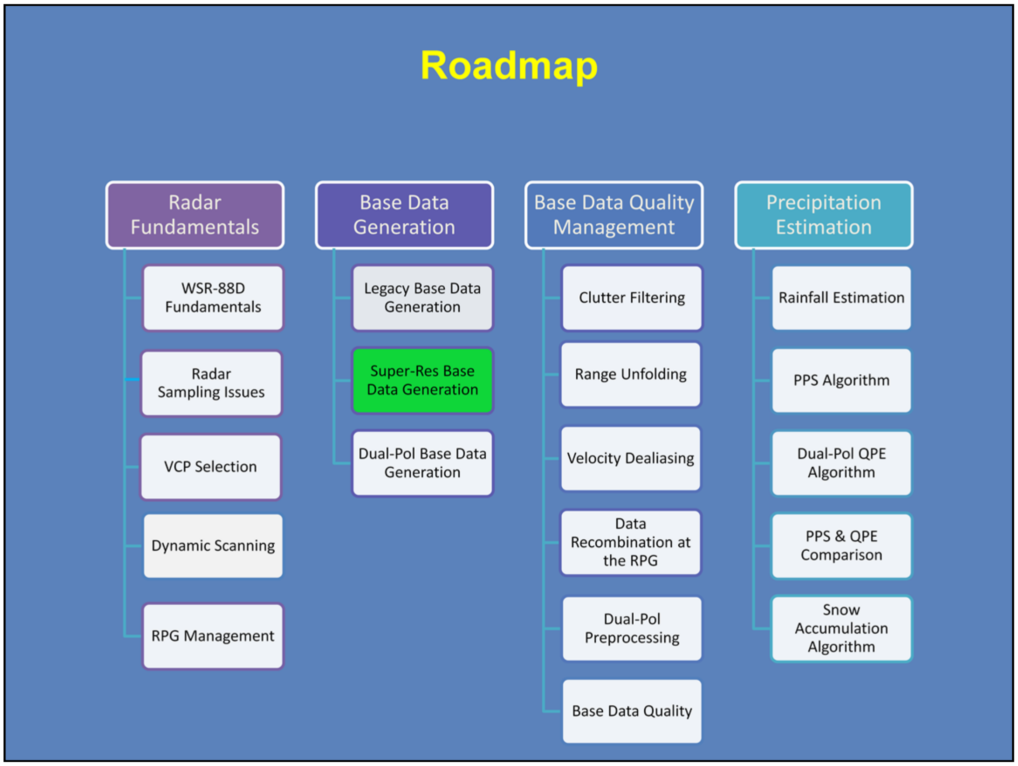

This concludes this lesson and here is the “roadmap” with your current location.

Welcome to Dual Polarization Base Data Generation

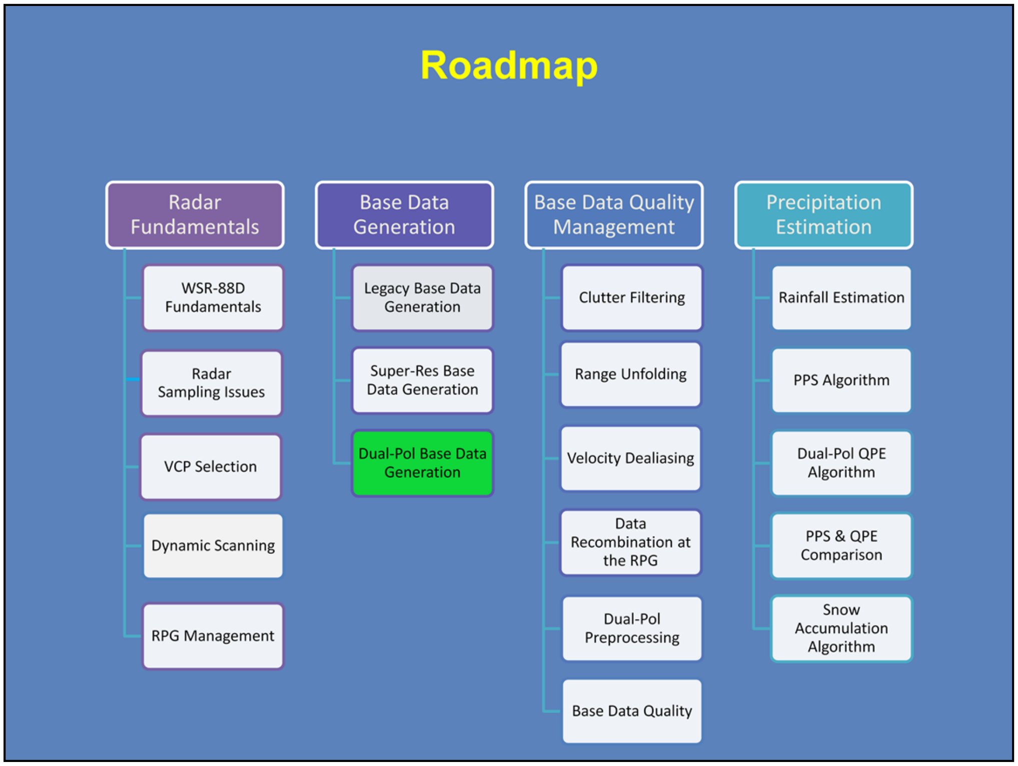

Here is the “roadmap” with your current location.

Learning 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.

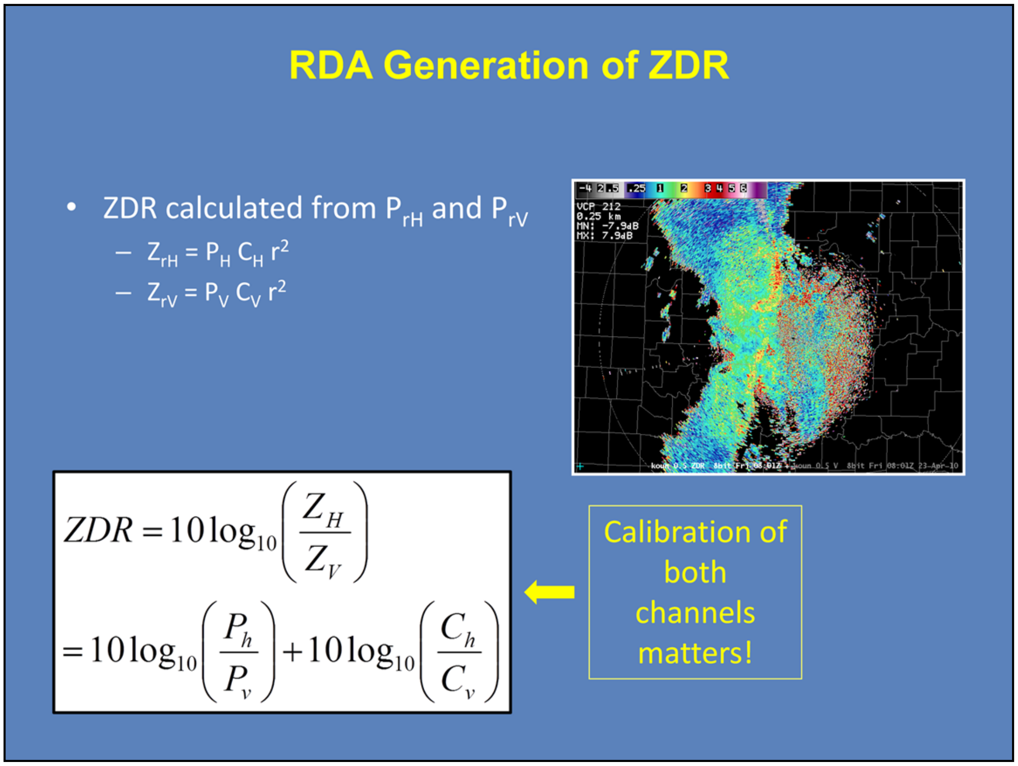

From the Probert-Jones radar equation, the Z value for each channel is equal to the returned power, times the range squared times a constant that is based on the radar's calibration.

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.

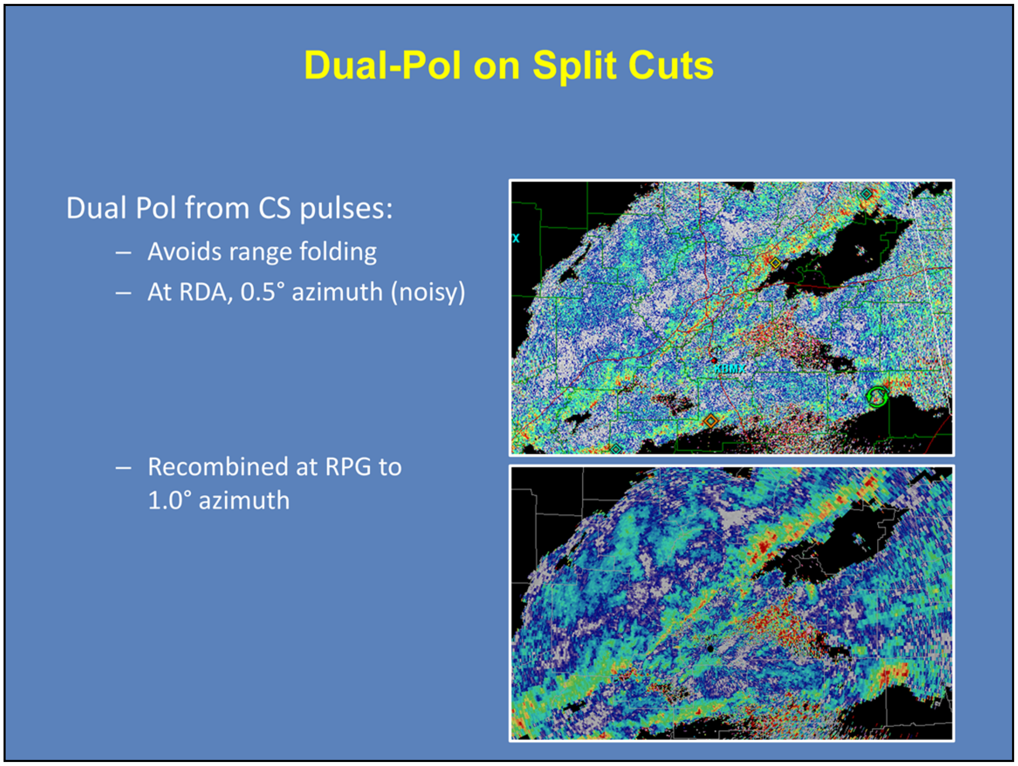

For the Split Cut elevations, ZDR and the other Dual Pol data are built from the low PRF, Contiguous Surveillance pulses on the first rotation. This is to avoid multiple tripping or range folding, and thus the need to “range unfold” the ZDR data. As with the legacy base moments (Z, V, & SW), the azimuthal 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

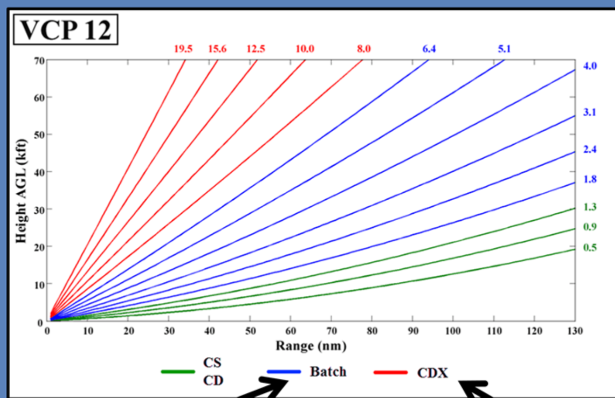

- Batch: CD pulses for dual pol
  - More CD pulses than CS
  - 1.0° azimuth
- Above Batch: dual-pol built from CDX pulses
  - That's all there is!

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.

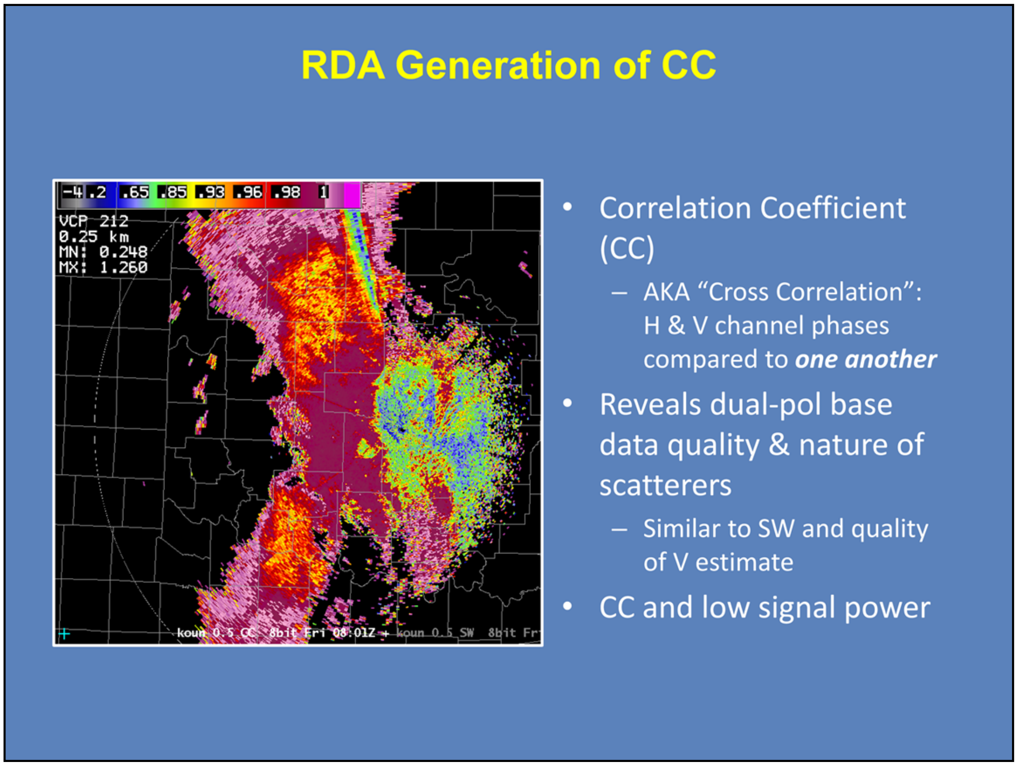

Correlation Coefficient (CC) measures the consistency of the H and V returned power and phase with one another for each pulse. This "cross correlation" looks at how the returned power and phase of one channel compares to the other channel. If the consistency is high (for example, stratiform light rain 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.

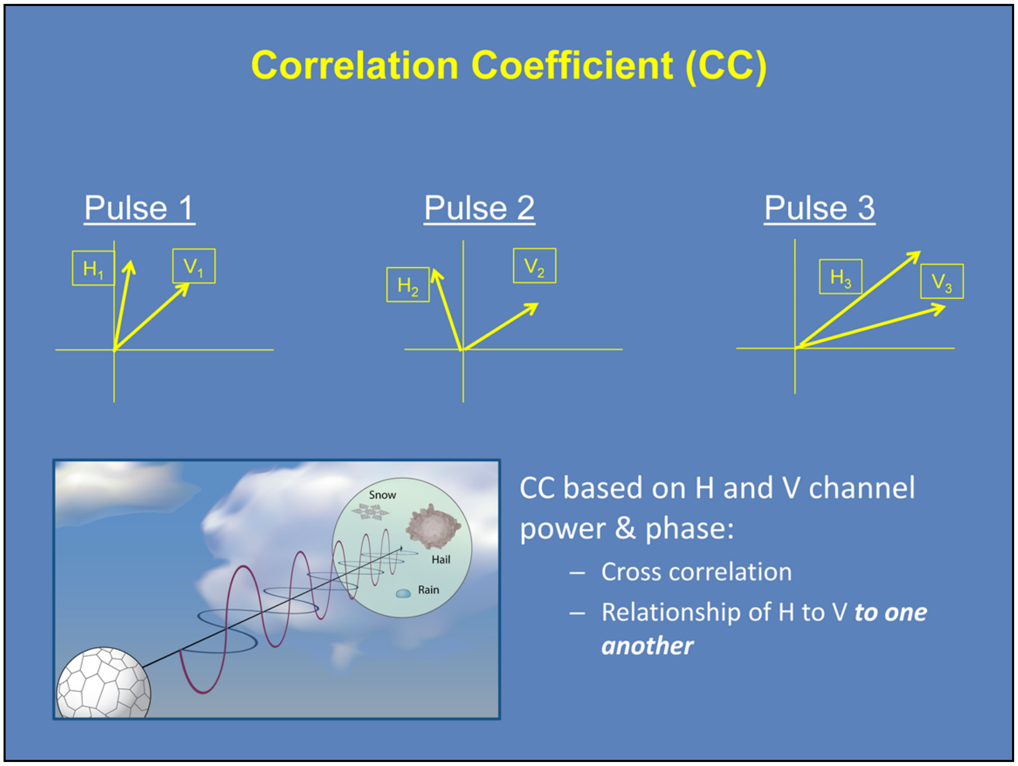

For each pulse, the returned power and phase from the H and V channels is known and can then be compared to one another. This type of comparison is know as a cross correlation. Of interest here is the magnitude of and the angle between the H and V vectors, which can be determined by vector multiplication.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

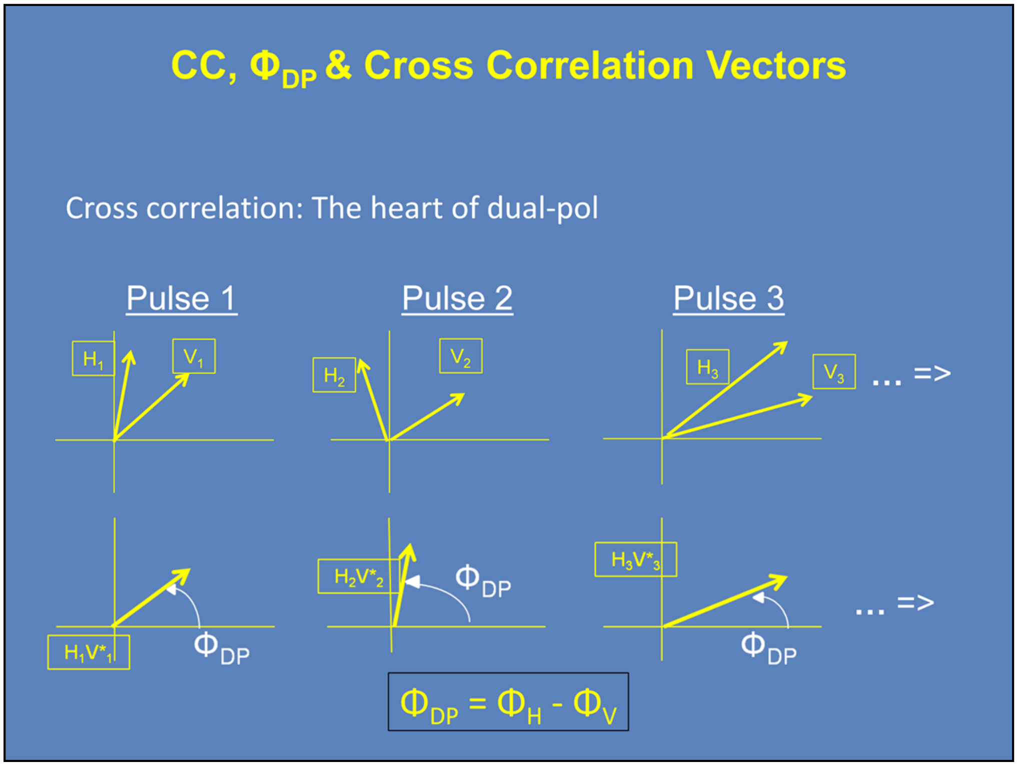

Cross correlation vectors are at the heart of Dual Pol base data. Here are a series of pulses, with the individual phasors for the H and V channels of each pulse. The cross correlation vector for each pulse captures how the horizontal and vertical information relate to one anther.

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.

Back to Top

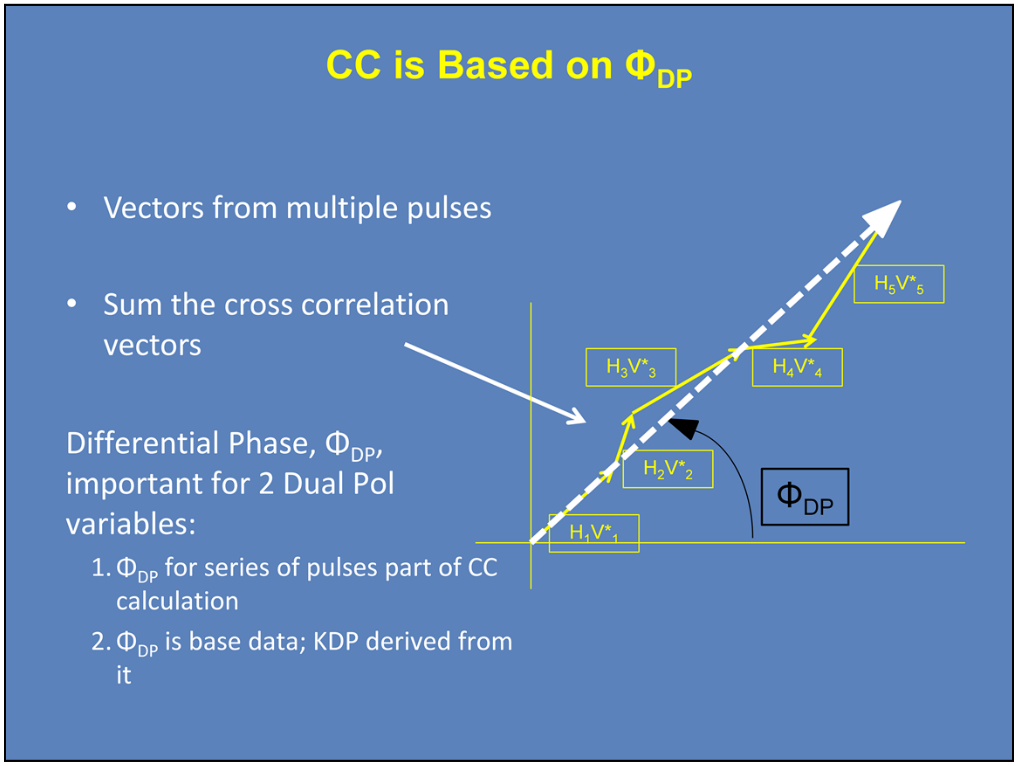

Since we don't assign any type of base data with just one pulse, the cross correlation vectors for a series of pulses are summed. This vector sum (white dashed arrow) is what's needed for the remaining two Dual Pol variables. Differential Phase, Φ_{DP}, is the angle of this vector sum, and it is part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

Back to Top

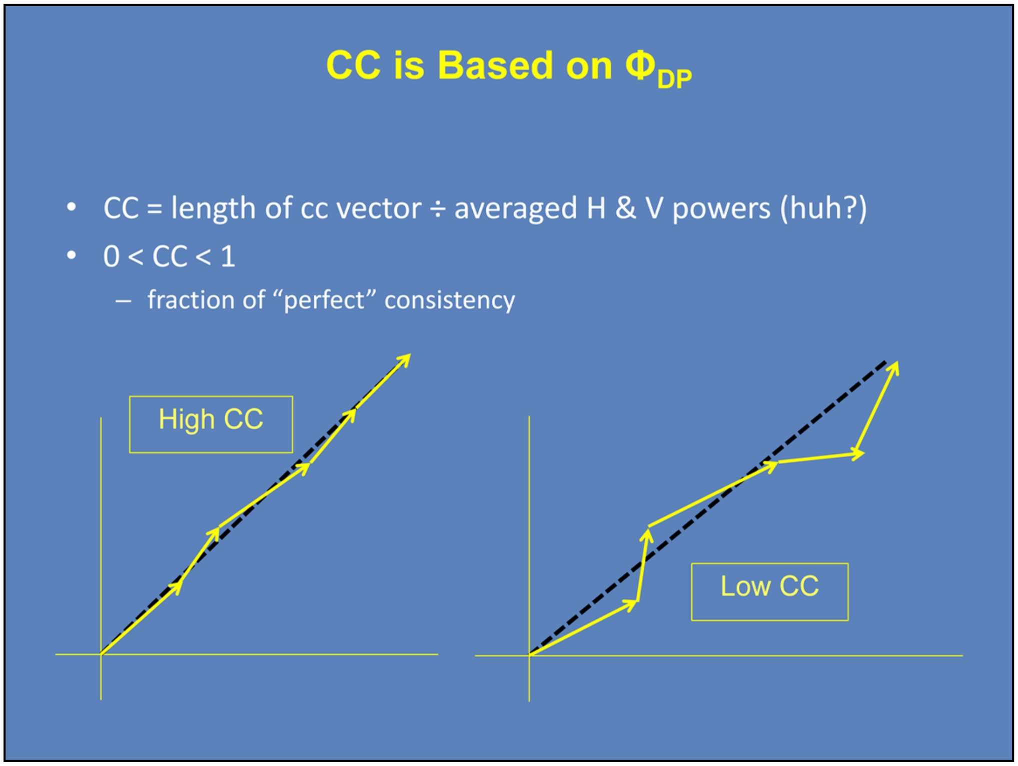

For a little more detail, CC is calculated by taking the length (or amplitude) of the vector sum of the cross correlation vectors and dividing it by the averaged H and V powers. This calculation captures the variation of the individual cross correlation vectors that contributed to the sum.

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.

Back to Top

CC-PhiDP Relationship

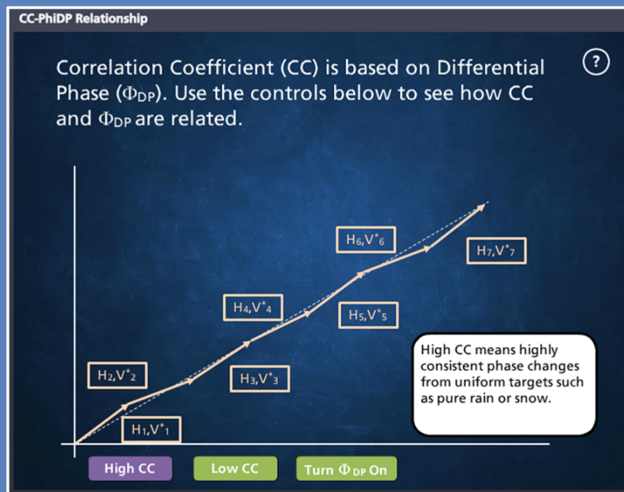

If no pop-up window appears that looks like the above, open a browser and go to: http://www.wdtd.noaa.gov/courses/rac/principles/interactions/cc-phidp/

Back to Top

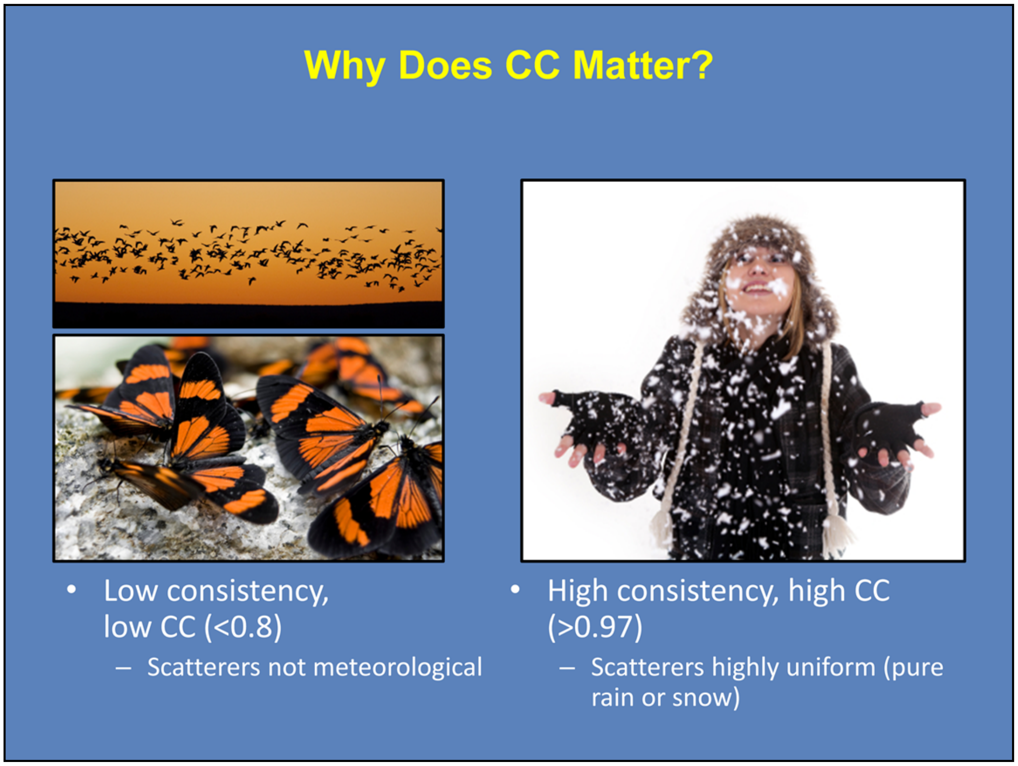

What does CC tell us? Low CC (<0.80) implies low consistency between H and V in the estimate and lot of diversity of the scatterers. In fact, CC <0.80 is so diverse the scatterers are unlikely to be meteorological, such as birds or insects. This distinction between biological and meteorological targets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

CC & Weak Returned Signal

CC and weak returned signal:
- Noisy with CCs > 1??

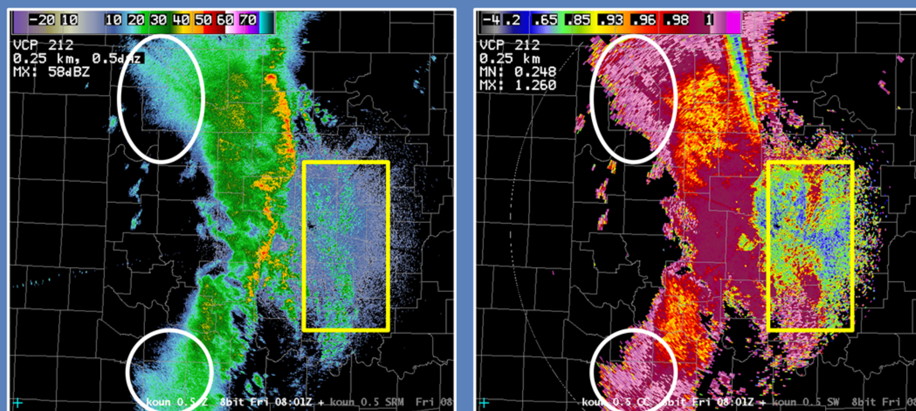

Correlation Coefficient is an important indicator of Dual Pol data quality. In general, the Dual Pol data will be noisier, and less reliable in weak signal areas than the legacy data. The good news is that CC can help to identify where the Dual Pol data are least reliable.

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.

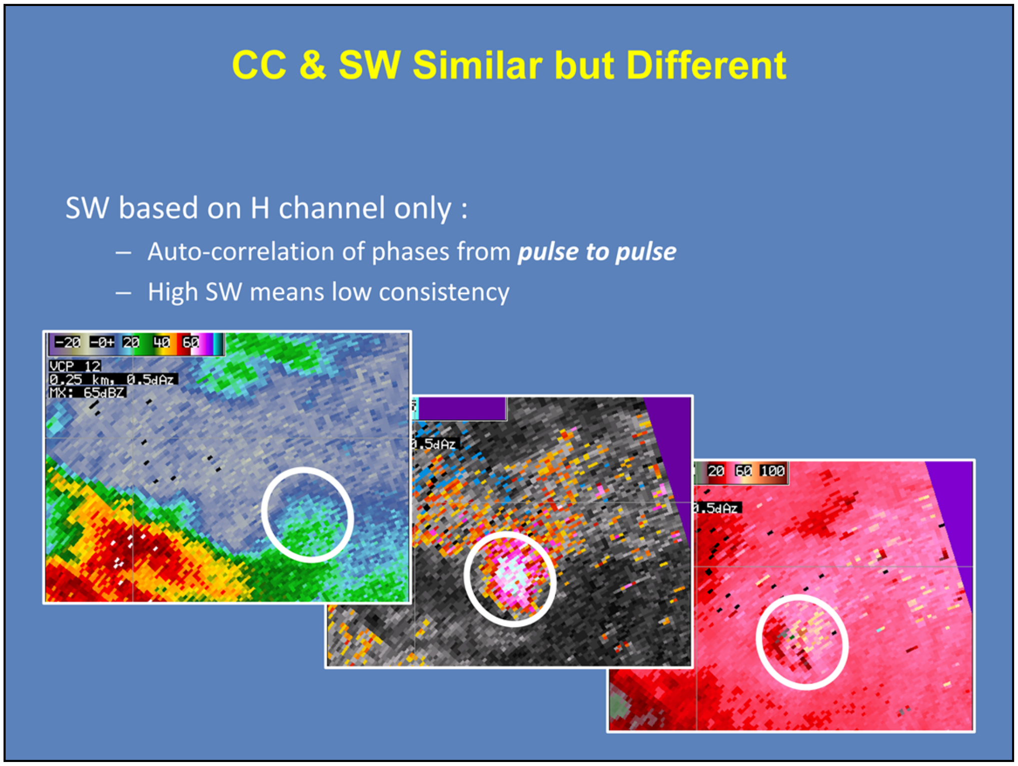

Correlation coefficient and spectrum width are analogous, but there are some important differences. Spectrum width is calculated from the horizontal channel only, and tells us something about the nature of the data from which the base velocity was calculated. Spectrum width is derived from auto-correlation,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

Back to Top

CC & SW Similar but Different

CC based on H & V channels:

- Cross correlation of H & V phases to one another
- High CC means high consistency

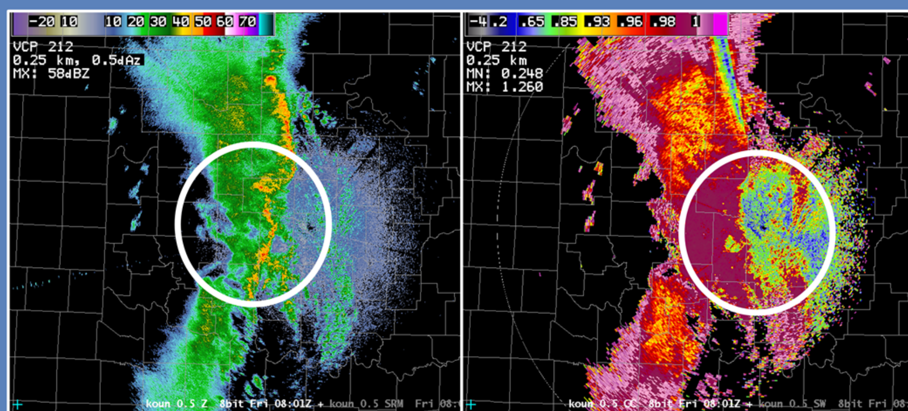

Unlike spectrum width, correlation coefficient is calculated from both the horizontal and vertical channels, and tells us something about the nature of the scatterers that were sampled. CC is derived from cross-correlation, which compares phases from the horizontal and vertical channel pulses to one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

Back to Top

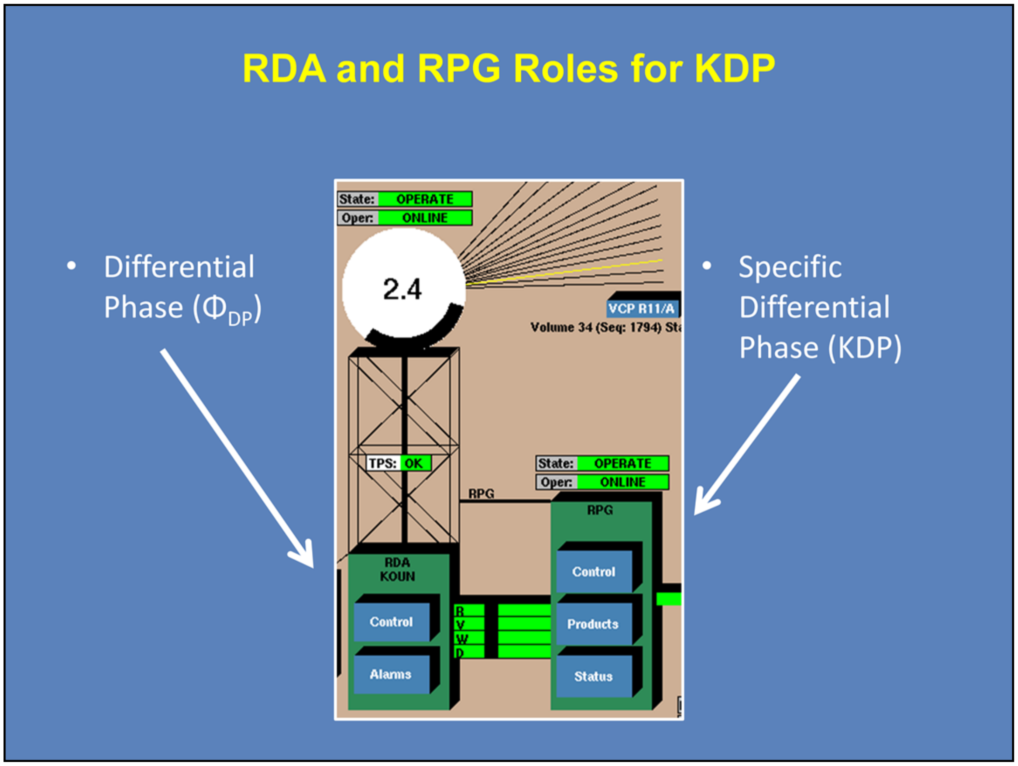

Another Dual Pol base product is Specific Differential Phase, or KDP, though it is technically a derived product generated at the RPG.

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

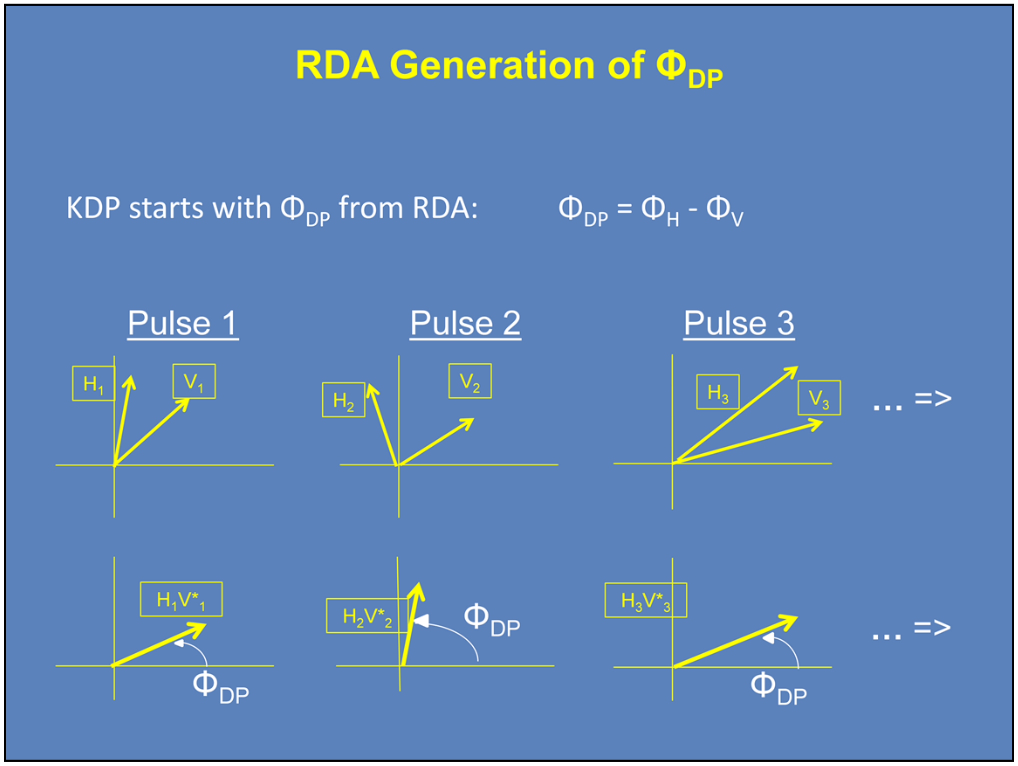

The benefit of KDP is that it tells you something about the type of medium (light rain? heavy rain?) that the beam has passed through.

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).

Back to Top

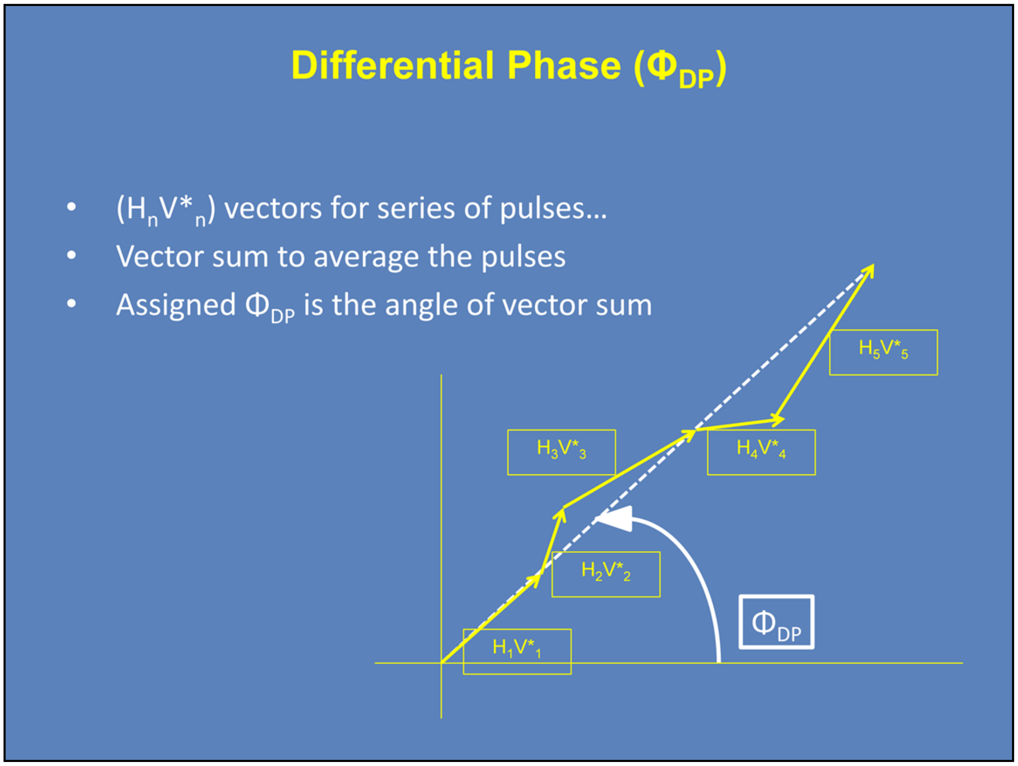

The differential phase, or Φ_{DP} value assigned as base data comes from the vector sum of the single pulse cross correlation vectors. The angle of this vector sum is the assigned Φ_{DP} for that range bin.

Back to Top

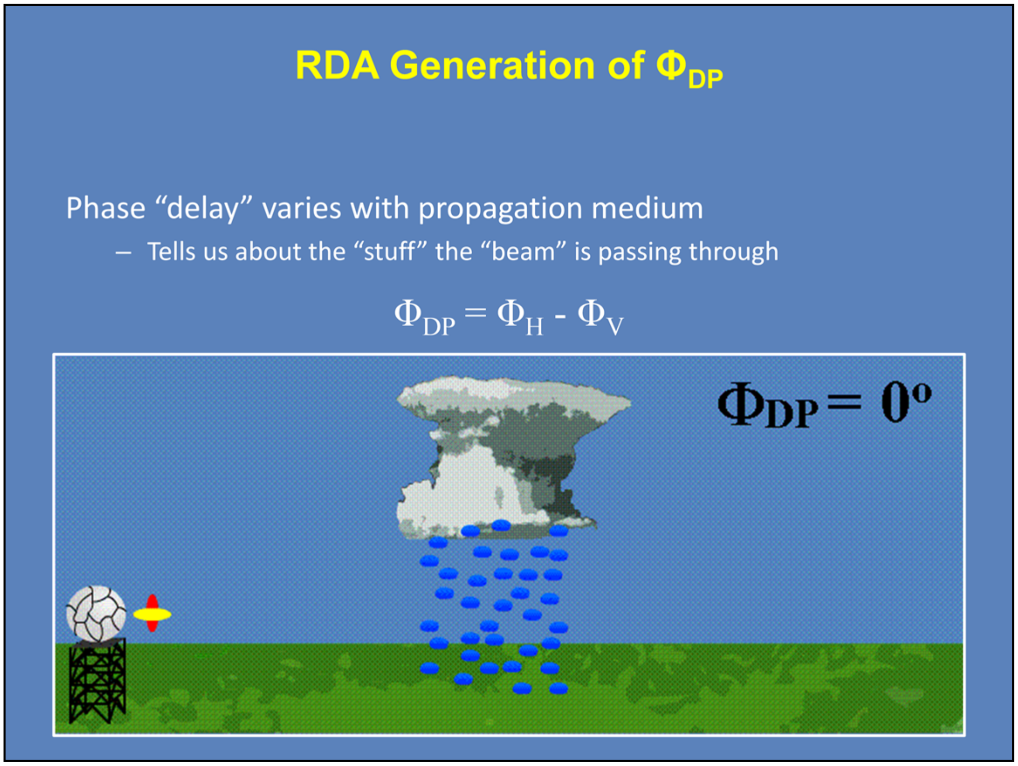

As the pulse propagates through different media (light rain, heavy rain, etc.), there is a delay that is apparent in the phase of the returned pulse. Since we have both horizontal and vertical phase, we can compare how the "H delay" differs from the "V delay". This gives us valuable information on the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

Back to Top

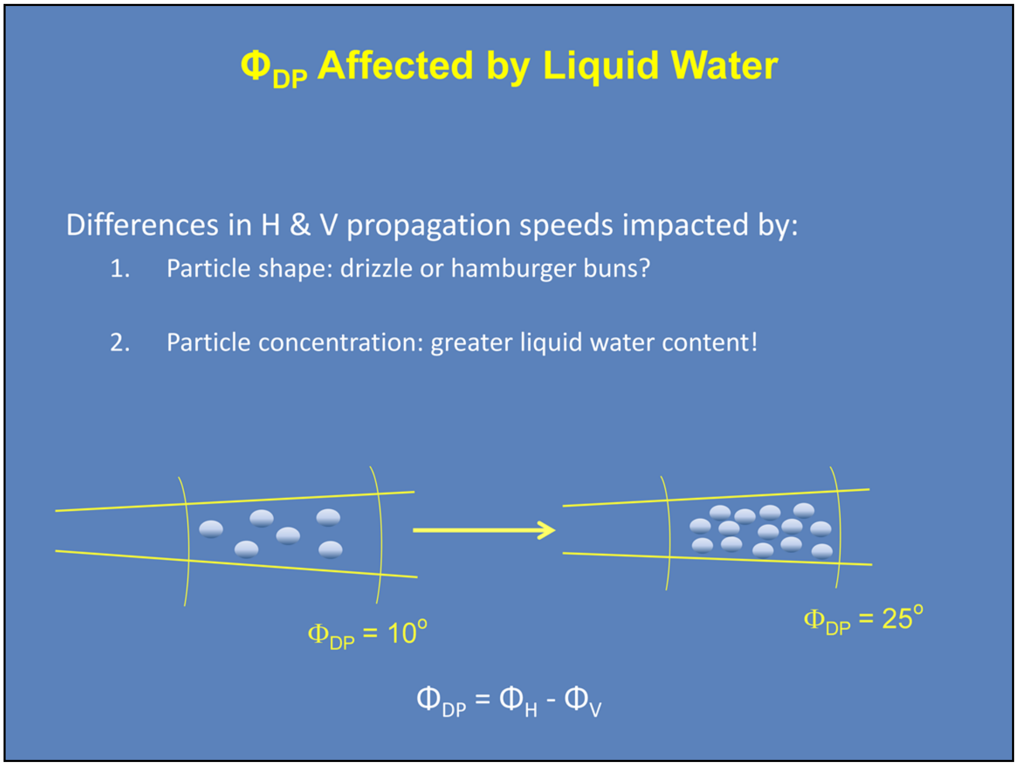

For any given atmospheric volume, the value of Φ_{DP} is affected by differences in propagation speeds of the horizontal and vertical waves. Propagation is slowed by both particle shape and/or by particle concentration.

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.

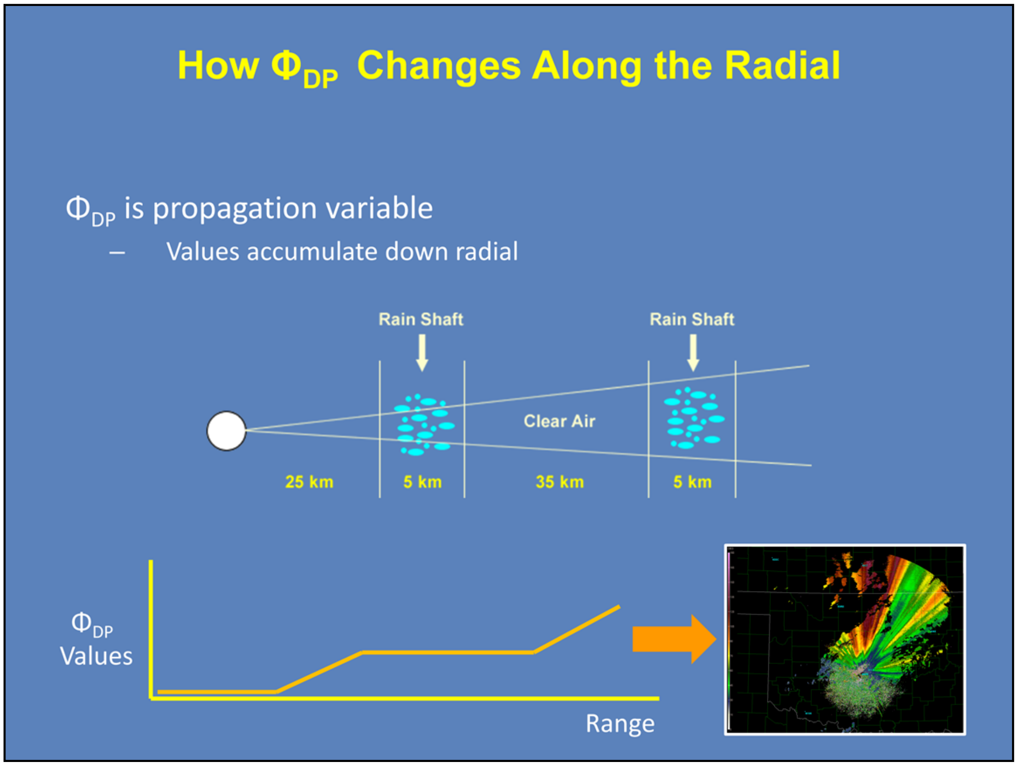

The value of Φ_{DP} propagates down radial, accumulating with range. There is no way to "reset" Φ_{DP} as the pulse travels outbound, encountering one or more areas of precipitation.

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.

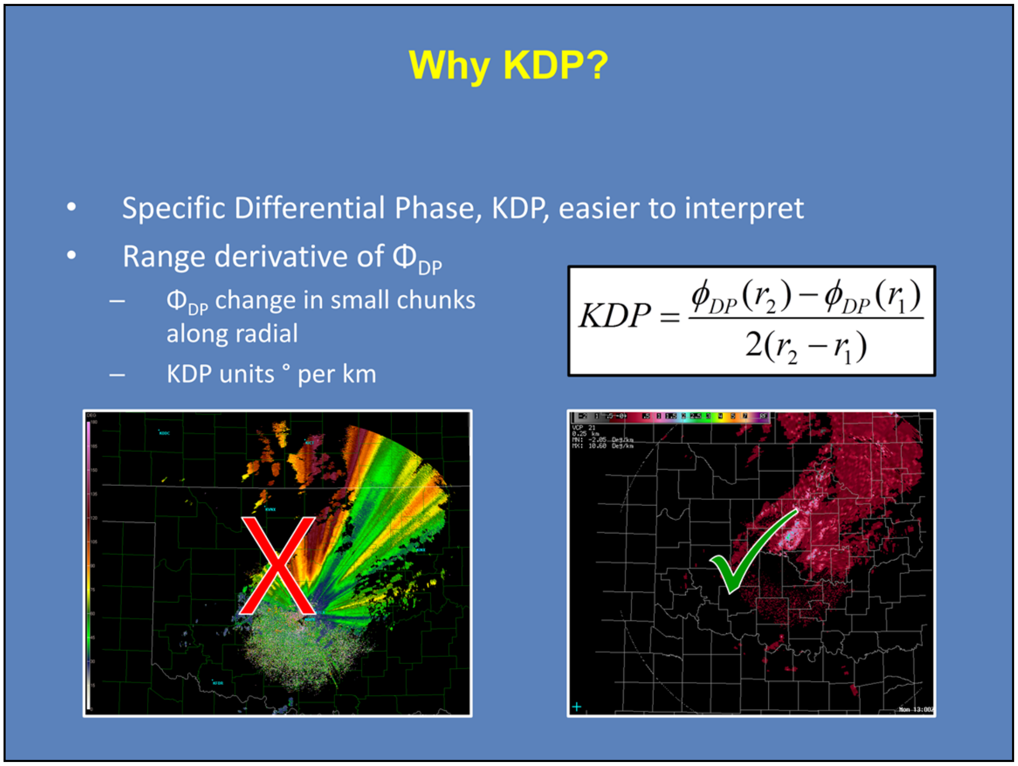

Differential phase, Φ_{DP}, is sent as base data from the RDA. Though both KDP and Φ_{DP} are available in AWIPS, KDP is easier to interpret. KDP is defined as the range derivative of Φ_{DP}. KDP is a way to capture how Φ_{DP} changes over very short ranges, which gives us more useful information. Thus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

Back to Top

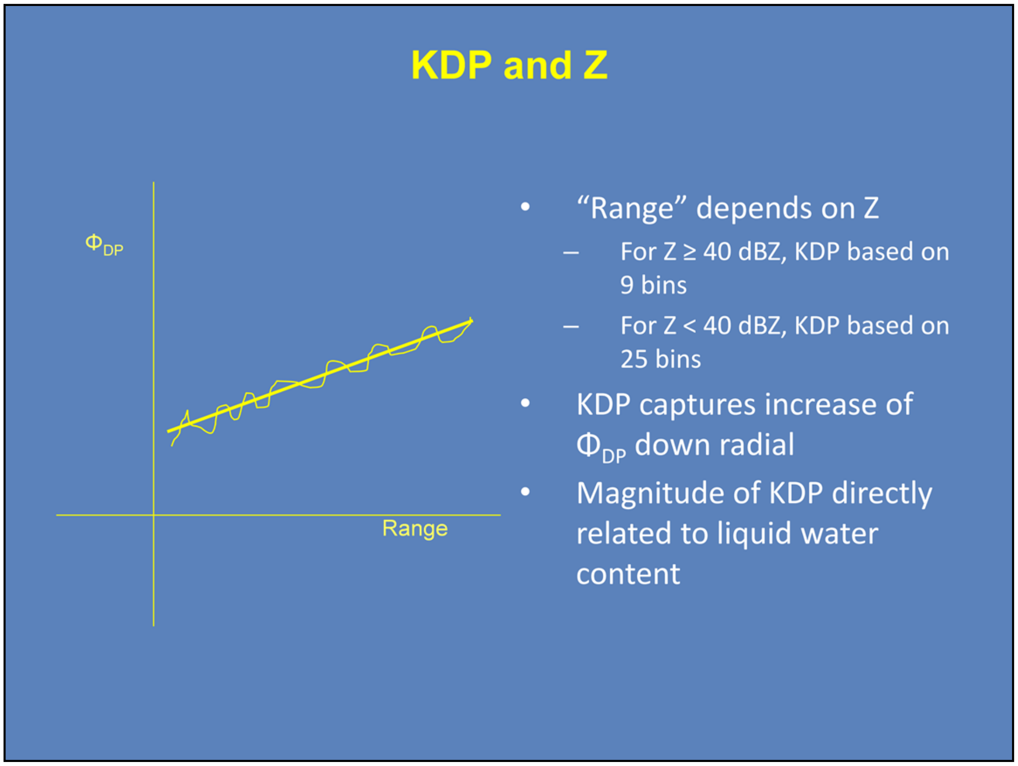

The span of range bins used for the KDP calculation is dependent on the Z value. Higher reflectivity generally corresponds to better data quality, and less noisy Φ_{DP} values. Thus for higher Z, KDP is calculated along a shorter interval of the radial, as compared to lower Z values.

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 and Liquid Water

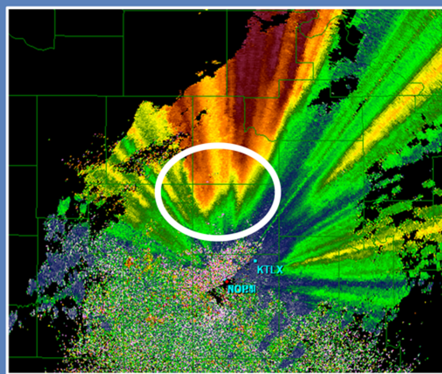

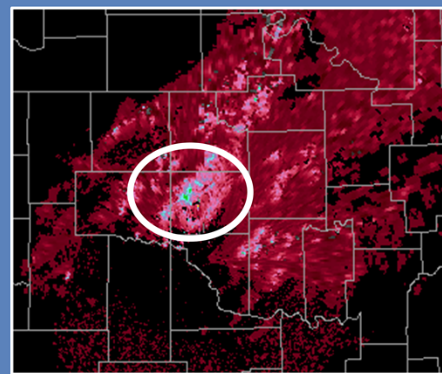

- High KDP means high local increase in Φ_{DP}
- KDP depicts relative liquid water content sampled by beam

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.

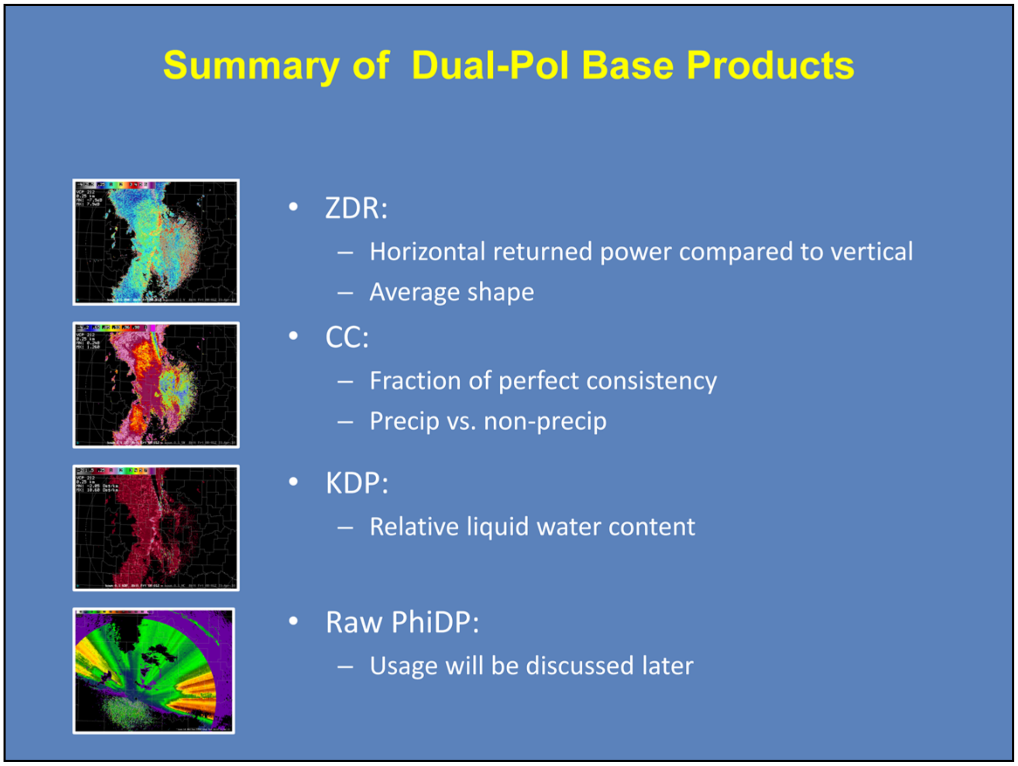

Differential Reflectivity, ZDR, is computed from the returned power in the horizontal and vertical channels. It does not use returned phase information. ZDR tells us about the average shape of the scatterers.

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.

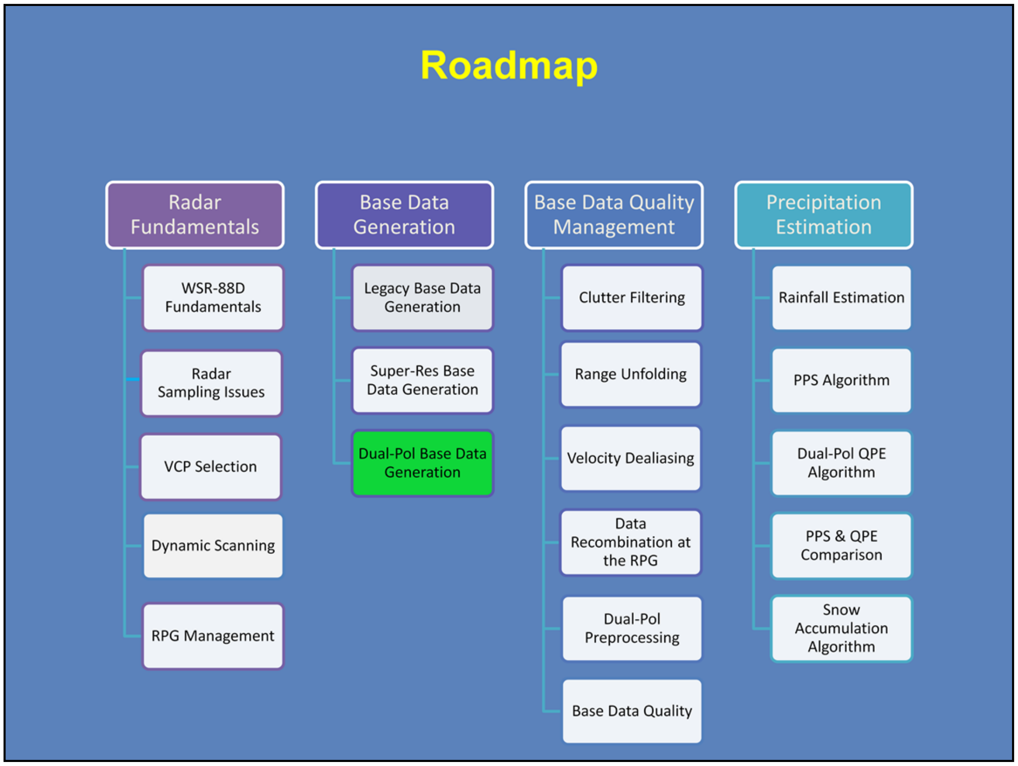

This concludes the lesson and here is the “roadmap” with your current location.

Back to Top

Welcome to this lesson on Clutter Filtering.

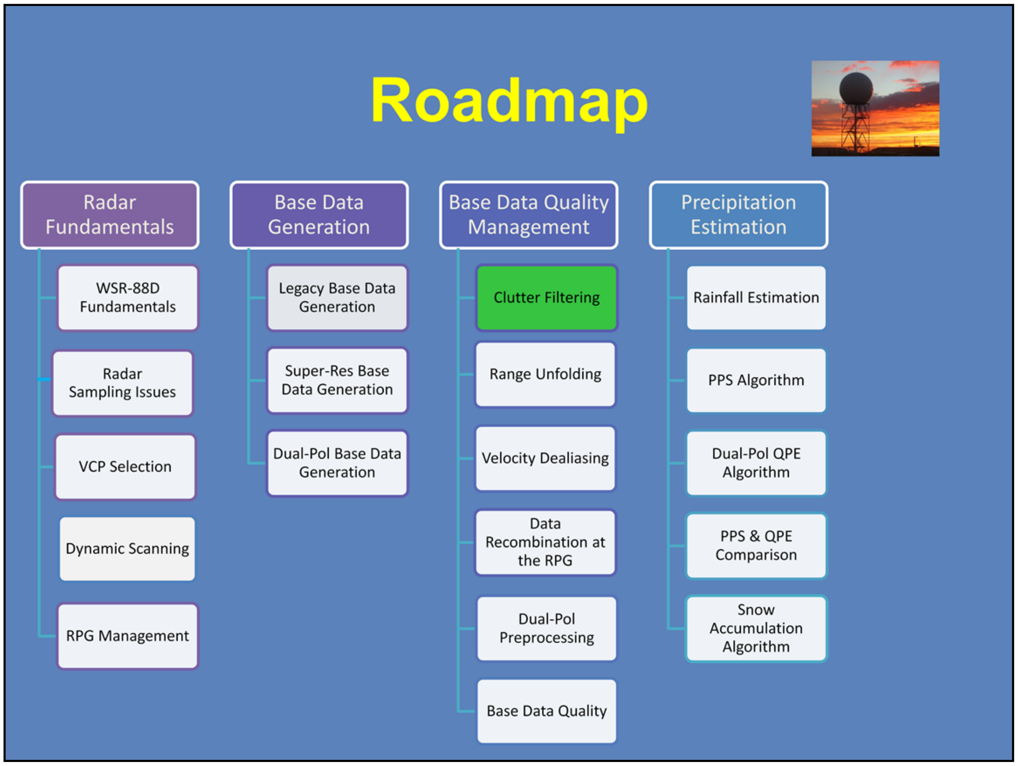

Here is the “roadmap” with your current location.

Back to Top

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.

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

Back to Top

Ground Clutter Contamination

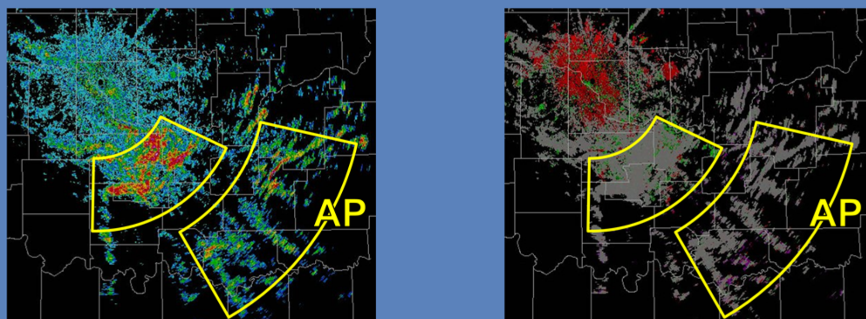

- Returns from non-moving ground targets
  - Affects ***all*** radar products
- Two types of clutter contamination
  - Normal ground clutter
  - Anomalous propagation clutter

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. So, unfiltered clutter is going to negatively impact all the radar 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 a transient type 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.

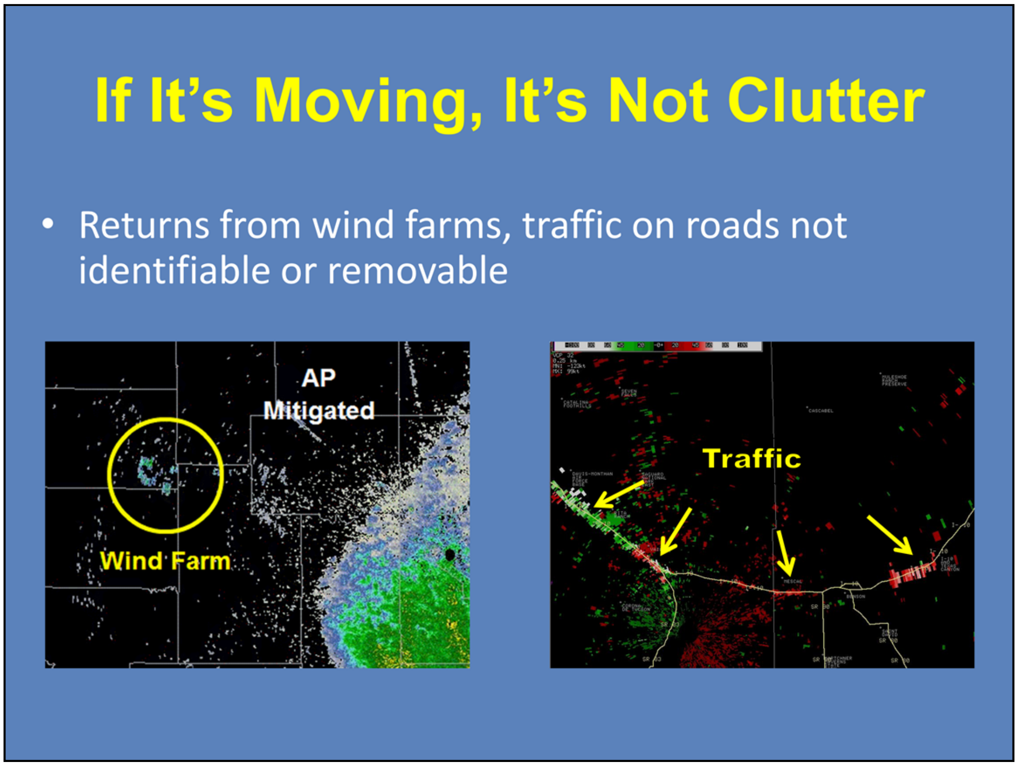

It's important to remember that if something is moving it is not considered clutter. The WSR-88D clutter algorithms are designed to detect ground clutter, which means near zero velocity and spectrum width. There will still be contamination from moving ground targets such as wind farms and traffic on highways.

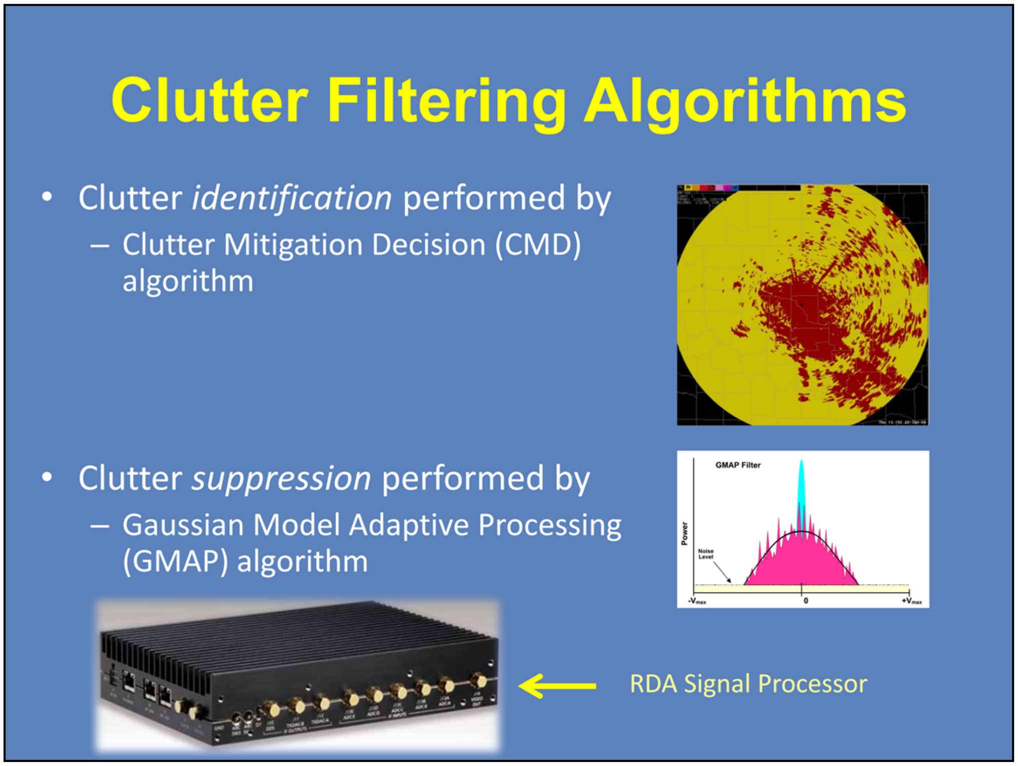

There are two different algorithms in the clutter filtering process. The data are first processed by the Clutter Mitigation Decision (CMD) algorithm, which does the job of identifying clutter on a bin by bin basis. For each bin identified by CMD, the Gaussian Model Adaptive Processing (GMAP) algorithm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

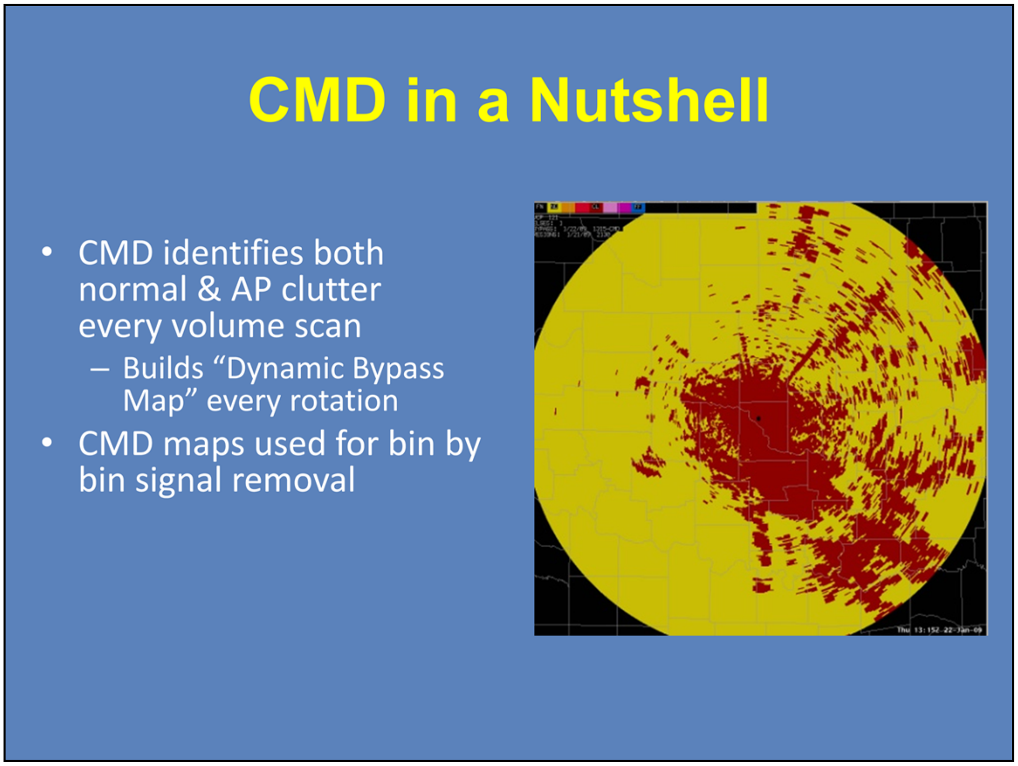

Clutter Mitigation Decision (CMD) offers an automated approach to the management of clutter filtering. CMD can identify both normal and AP clutter every volume scan, which eliminates the need for manually defining and downloading regions files to address AP clutter. CMD builds what's known as a dynamic 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.

Back to Top

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 (i.e. low variation) = clutter
  - Low CPA (i.e. highly variable) = weather

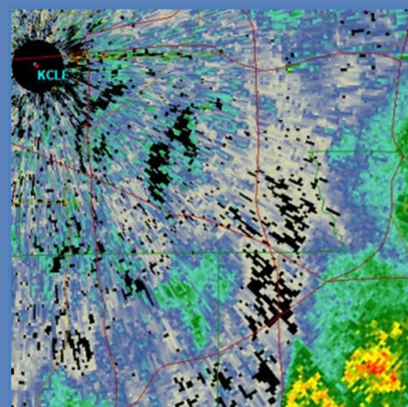

CMD is a fuzzy logic algorithm with multiple inputs. Here are a couple of inputs based on reflectivity.

1. The first one is (Z) texture, and if you think about it, weather tends to have a much smoother texture than clutter.

2. Another one of the inputs is called Z spin, and this is about how the reflectivity gradient changes sign as you move along the radial.

3. Another input is called the Clutter Phase Alignment (CPA), and this captures the variance of pulse to pulse phase changes. If you have a high CPA, that means there’s a good alignment of the phasors (a low variation in those returned phase values) and a higher likelihood of clutter.

Back to Top

CMD Inputs

Standard deviation of ZDR

– Low STDZDR = weather; High STDZDR = clutter

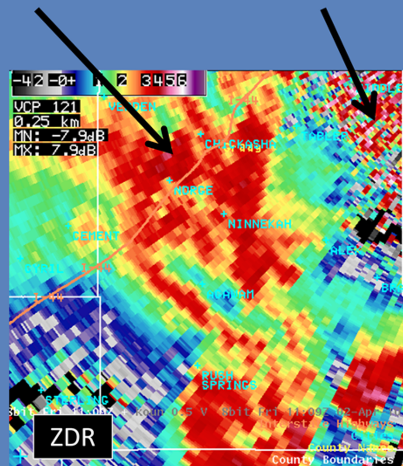

There are two inputs based on Dual-Pol base data. They are standard deviations of ZDR and Φ_{DP}.

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 is most likely (the upper right-hand corner) 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.

Back to Top

CMD Inputs

- Standard Deviation of Φ_{DP}
  - Low STDPHI = weather
  - High STDPHI = clutter

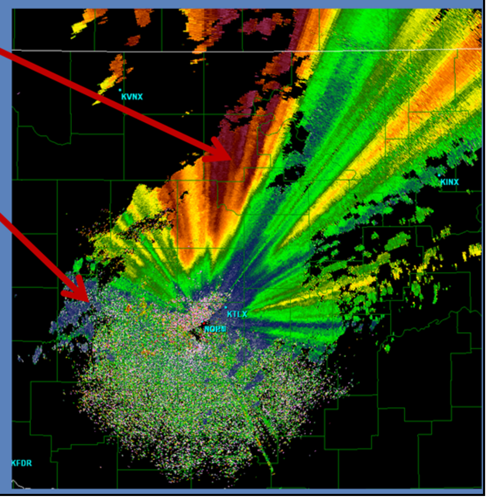

The second CMD input based on Dual-Pol base data is the standard deviation of Φ_{DP}.

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.

Back to Top

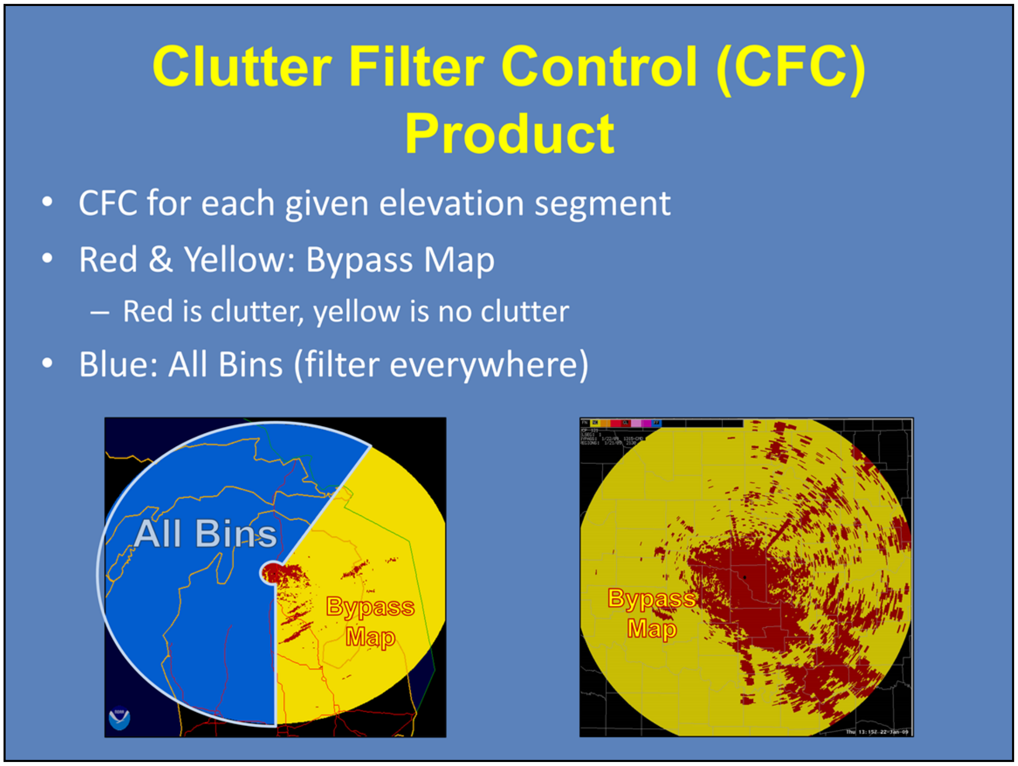

For each elevation segment, the CFC product displays the type of clutter filtering that has been invoked for the lowest elevation in that segment. The area with a red/yellow combination are where the Bypass Map is in control. That means there is a bin by bin identification of clutter, with red for clutter (to be filtered) and yellow for no clutter (to be left alone). Where you see blue is where All Bins filtering has been manually implemented, and that means you're applying clutter filtering to every single bin within that area that's blue.

These varying clutter filter schemes are implemented manually by downloading clutter region files form the RPG to the RDA.

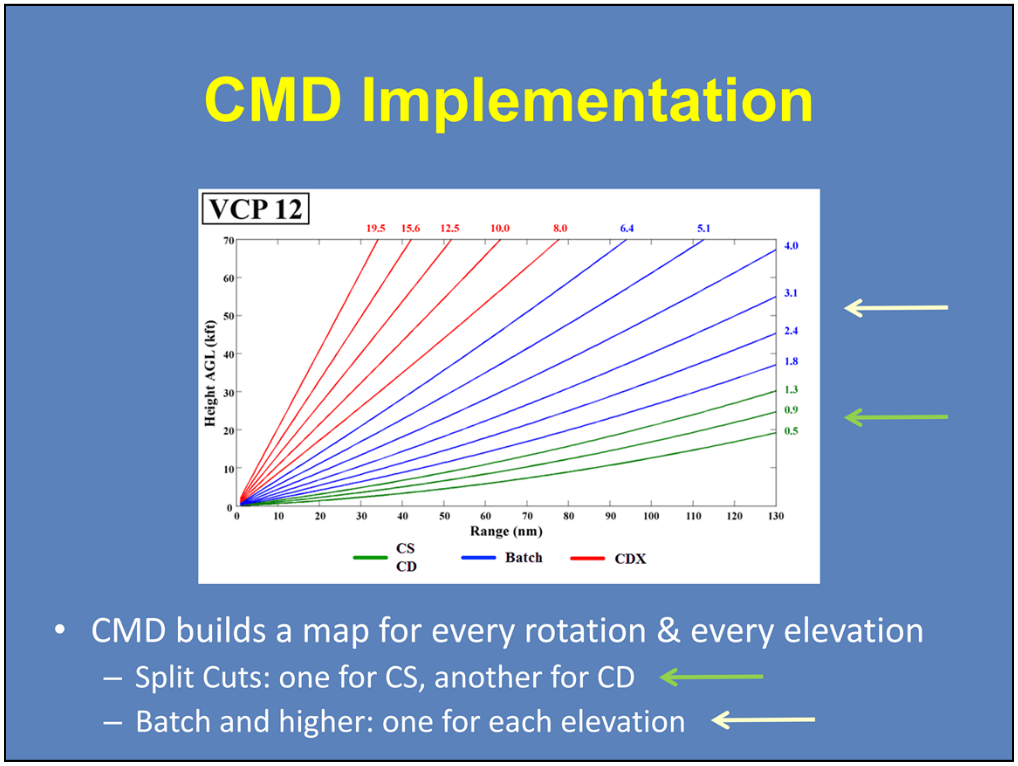

CMD is active for every rotation and every elevation, building a bypass map each time. For the Split Cuts, there are two rotations at the same elevation, Contiguous Surveillance (CS), then Contiguous Doppler (CD). CMD builds a different map for each of these two rotations. For the remaining elevations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

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

Back to Top

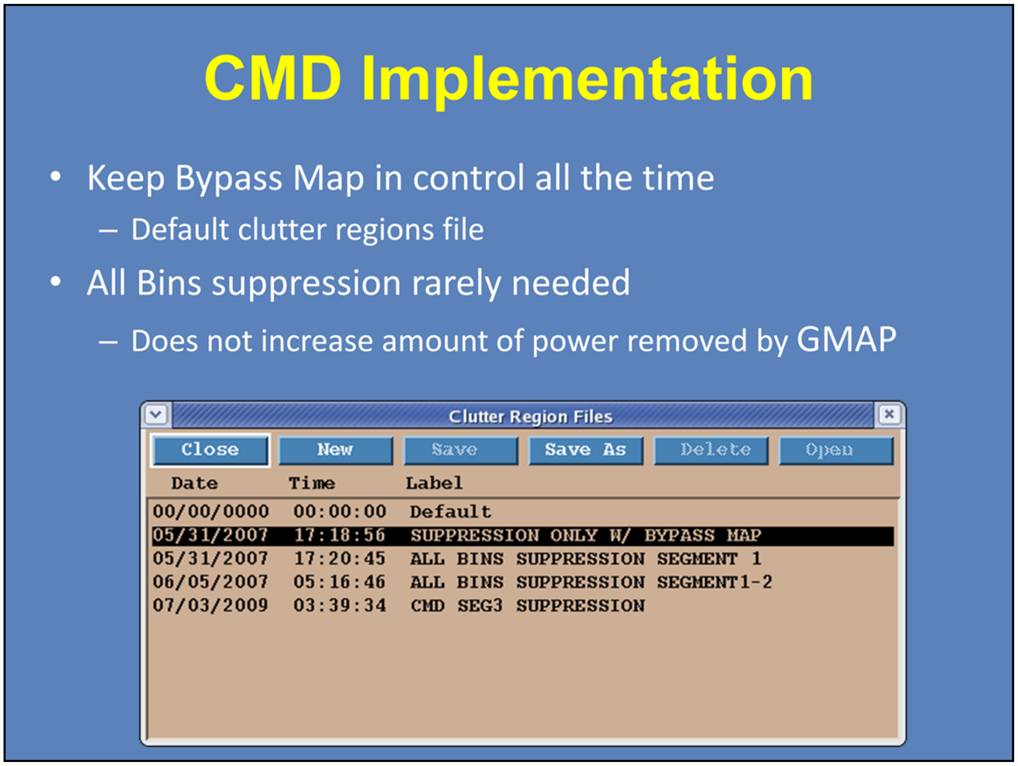

CMD offers “hands off” clutter suppression, though some data quality monitoring may be needed from time to time. In order for CMD to build maps for every elevation, the Bypass Map must be in control for the entire display for all elevations. The “Default” clutter regions file is designed to do just that. It may be necessary (as in this example) to create a local version of this file (e.g. “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

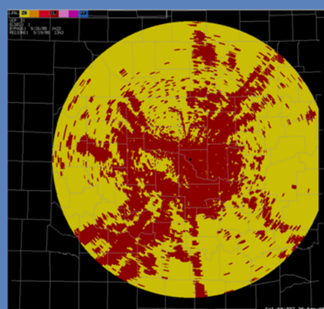

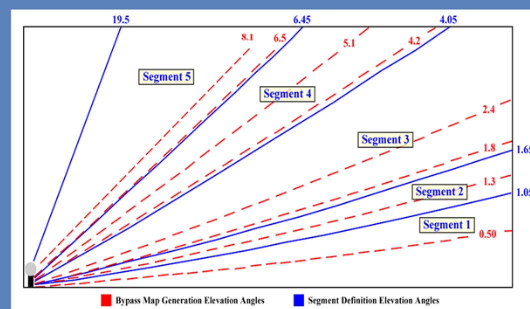

In summary, CMD's job is to identify the bins that contain clutter, both normal and AP clutter. CMD performs this identification and builds a clutter map for every rotation and every elevation of every VCP (that's a lot of maps). The Clutter Filter Control product is available to help visualize where 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 regard 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.

Back to Top

Clutter Filtering Algorithms

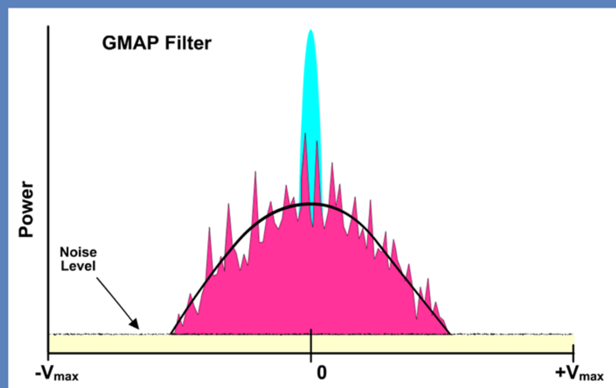

- CMD has ***identified*** clutter
- For the bins identified by CMD, filtering performed by Gaussian Model Adaptive Processing (GMAP) algorithm

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

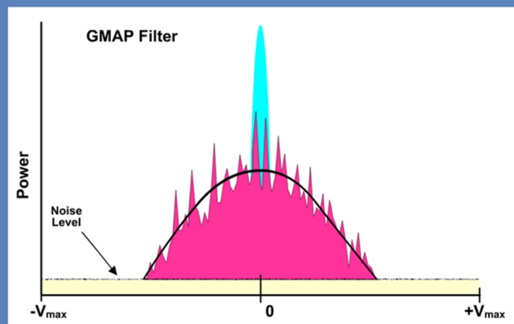

- Weather and clutter signals look different
  - Can be approximated by Gaussian curve
- Remove power from narrow spike near zero velocity
- Once power removed, GMAP attempts to rebuild lost weather signal

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. We're doing that in this super simple graphic. And one thing to keep in mind here is that 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.

Back to Top

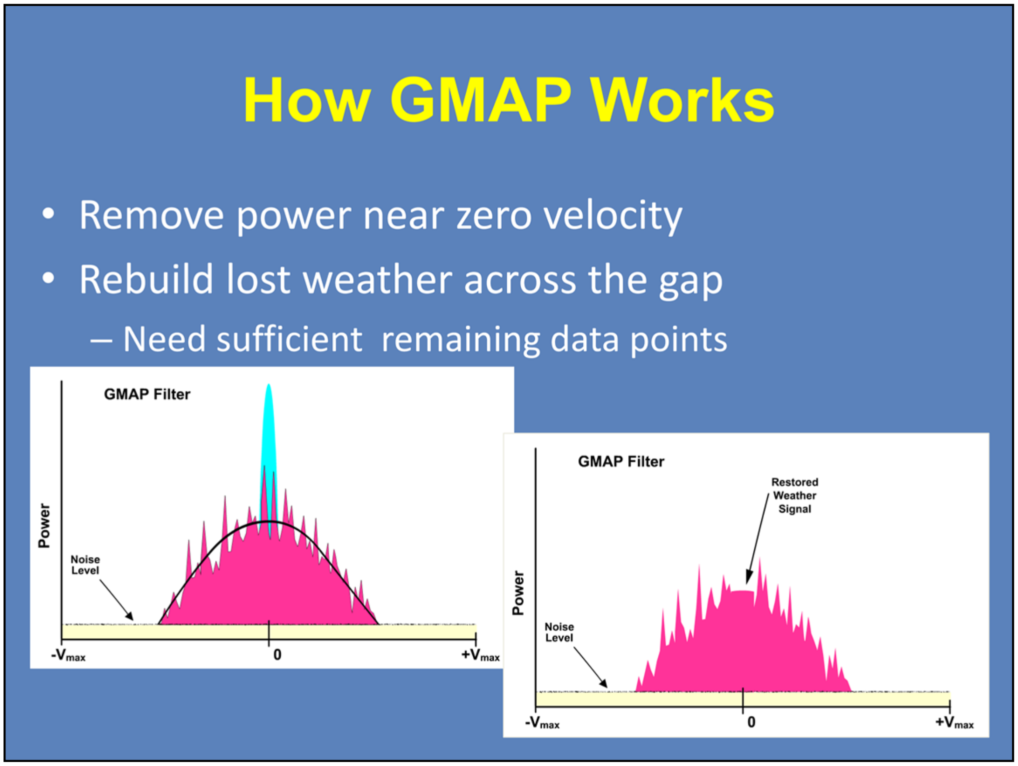

Once the width for signal removal is determined, filtering is applied to all of the signal within that width. In this case, both the clutter and weather signal within this interval will be removed. However, GMAP has the ability to rebuild the weather signal that was lost. This is dependent on the availability 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.

Back to Top

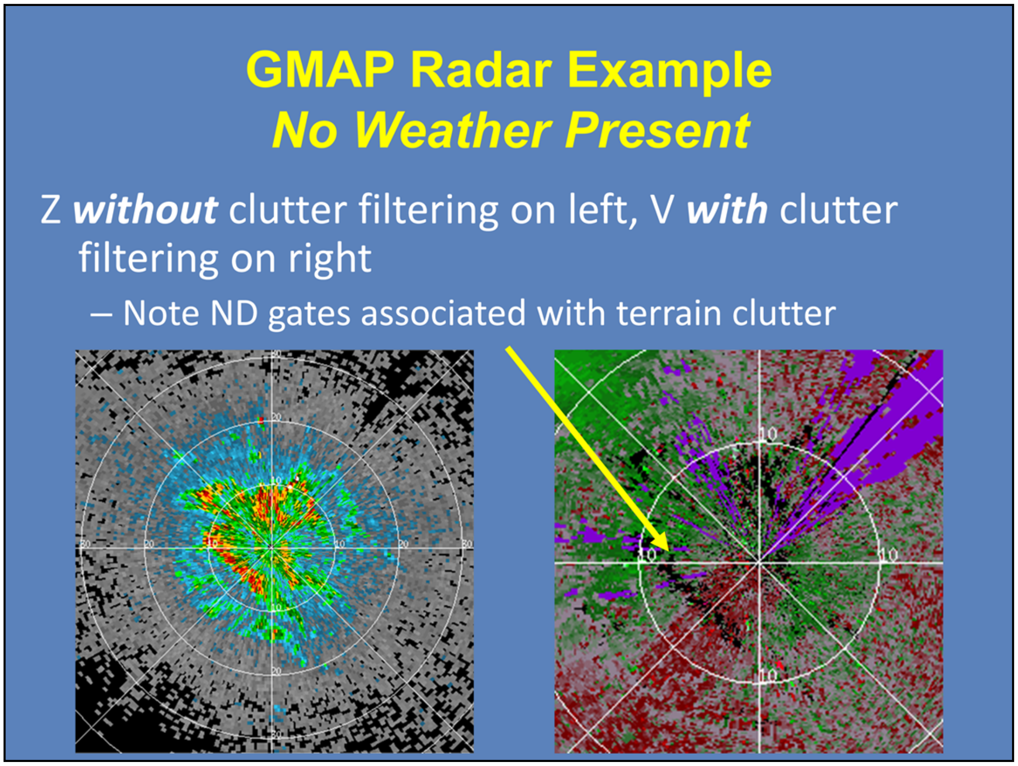

As an example of GMAP performance, we start with no weather with either of these images. On the reflectivity on the left, clutter filtering has been turned off to identify the local terrain clutter. Of particular importance is the ridge line to the southwest. On the velocity on the right, clutter filtering 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.

GMAP Radar Example
Squall Line Passes Through

Z & V as squall line passes through:

- Clutter has been filled in with data

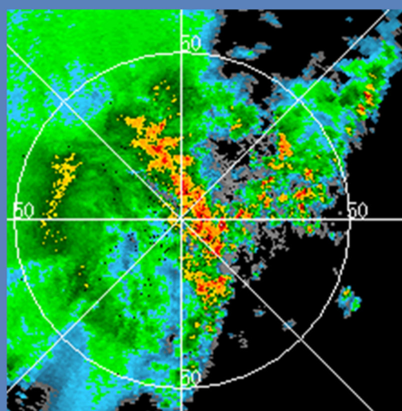

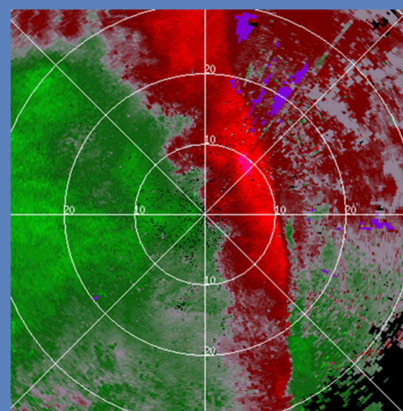

At a later time, a squall line has passed through the area. The reflectivity on the left and the velocity product on the right show the squall line, but the velocity product has been zoomed in compared to the reflectivity.

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

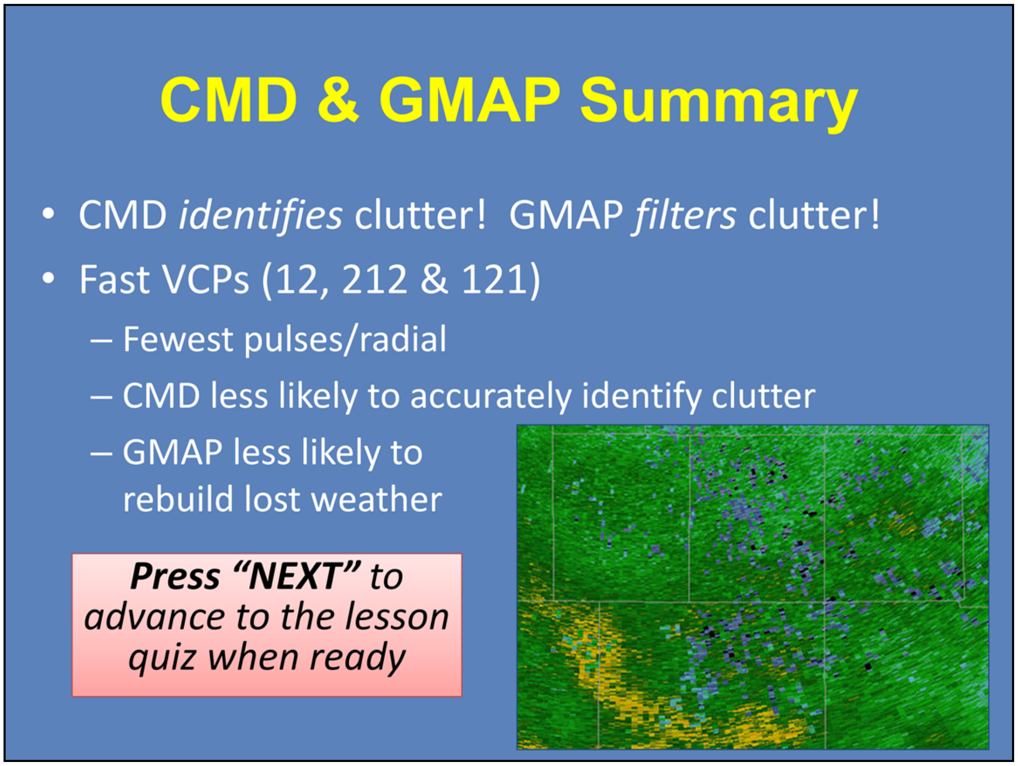

In summary, Clutter Mitigation Decision (CMD) and Gaussian Model Adaptive Processing (GMAP), work together to identify and suppress clutter. CMD identifies the clutter on a bin by bin basis. And GMAP performs the suppression by applying filtering only to the bins identified by CMD and removing power from the signal near zero velocity. GMAP can rebuild 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 fewest pulses per radial. For these VCPs, it is more difficult for CMD to discriminate clutter from weather with little movement, like stratiform rain. 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 used, even though VCP 12 is designed for intense convective precipitation. The low number of pulses per radial from VCP 12 in stratiform rain makes it harder for CMD and GMAP, resulting in numerous gates of "lost" data.

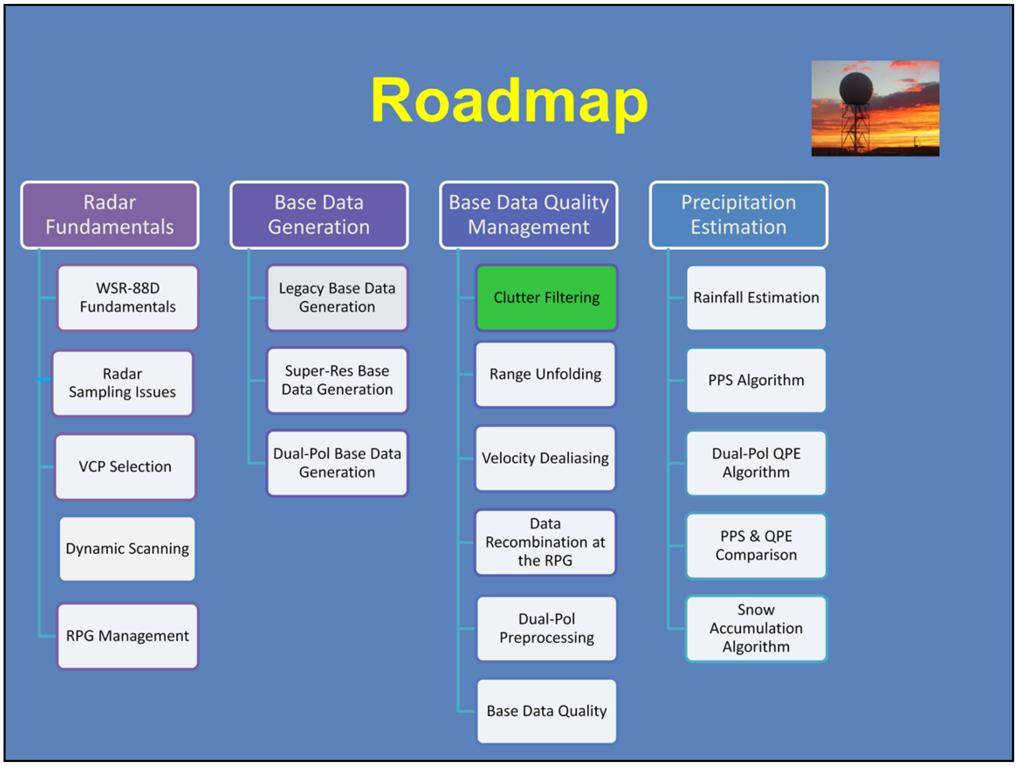

Here is the “roadmap” with your current location.

Welcome to this lesson on the Principles of Meteorological Doppler Radar. This lesson will discuss range unfolding of precipitation echoes.

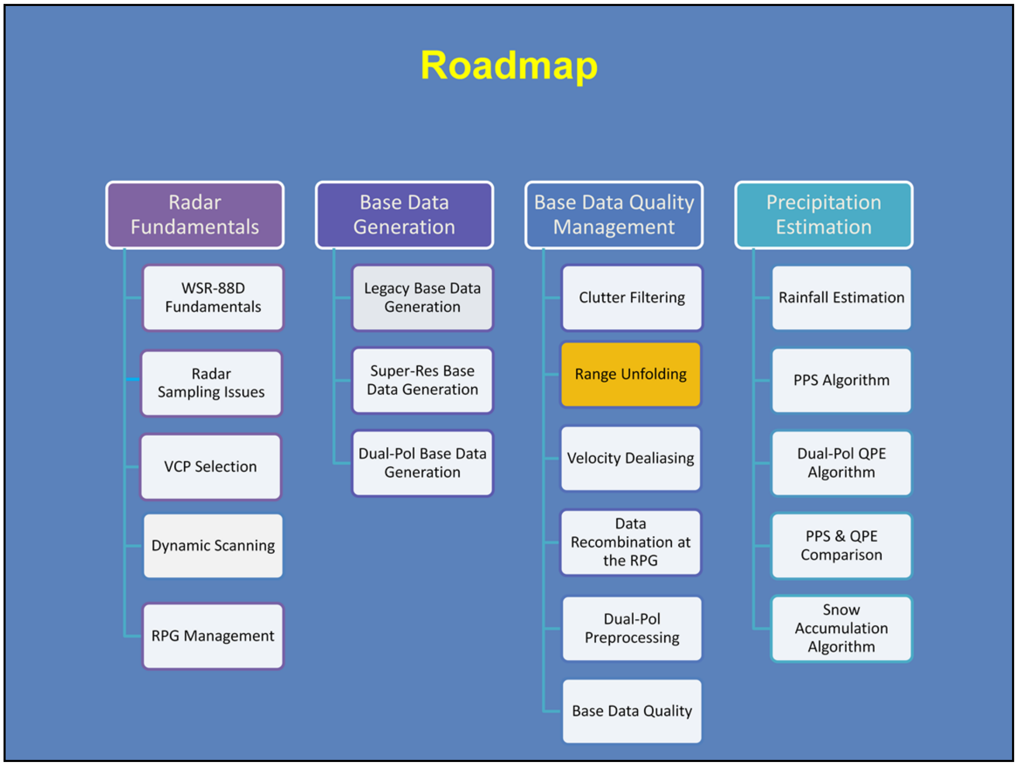

Here is the “roadmap” for this topic with your current lesson highlighted.

Back to Top

Objective

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 this lesson. Please read it over and advance to the next slide when you are ready.

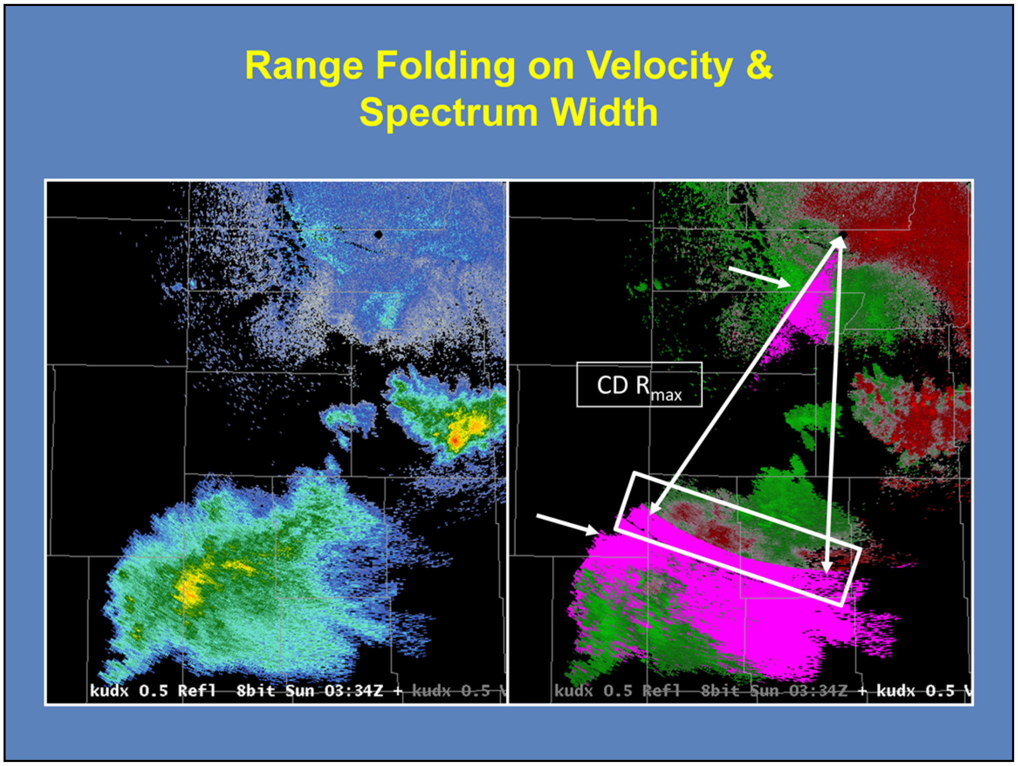

Recall from our discussions of the Doppler dilemma that high PRFs are needed to allow for high Nyquist velocities, but also result in short maximum unambiguous ranges (or R_{max}). When precipitation echoes are detected past the first "trip", it is necessary to unfold the data to its appropriate range.

Through the next several slides, we will explain how the data are unfolded and the purple "RF" data areas are assigned on the Radial Velocity and Spectrum Width products. An important feature to notice in these products is the discontinuity in the velocity data at a fixed range as seen in the example shown here. You will often see this discontinuity in velocity and spectrum width data as long as there is sufficient areal coverage of echo over multiple trips. This discontinuity occurs at the range for R_{max} on the Contiguous Doppler (or high PRF) cut used for velocity data collection. The two important questions to ask are: "how did any velocity data end up being assigned beyond the first trip?" and "why does the radar assign some of those regions with valid velocity data and others with the RF 'purple haze'?"

Back to Top

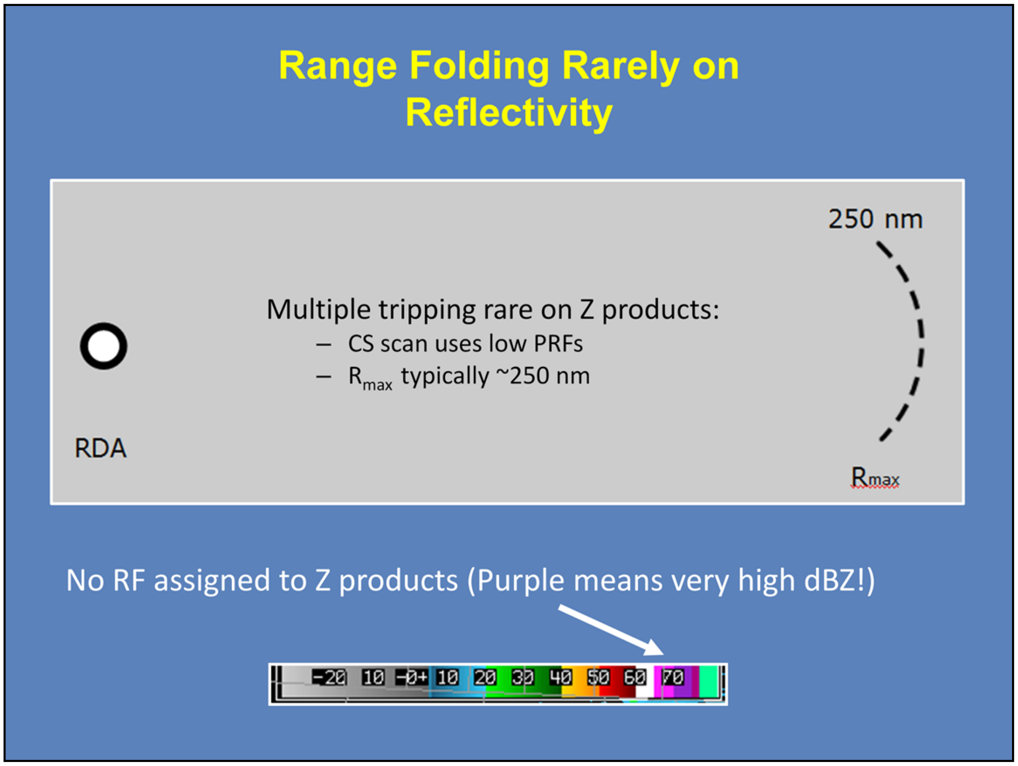

Remember that the Continuous Surveillance scans (where reflectivity data are collected) use low PRFs that result in long maximum unambiguous ranges. So, these cuts produce reflectivity data with a R_{max} of ~250 nm and very little data that gets collected outside of the first trip. As a result, the color purple on Reflectivity products indicates very high dBZ values, not range folding.

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 several VCPs:
  - VCPs 11, 12, 21, 31, & 32

We will discuss two range unfolding techniques in this lesson, the first of which is the technique released initially with the WSR-88D. We will refer to this range unfolding algorithm as the Legacy Range Unfolding Algorithm.

The goal of the Legacy Range Unfolding Algorithm is to assign velocity and spectrum width data to its appropriate range. That range can be beyond the maximum unambiguous range for the Continuous Doppler scans. So, the precipitation echoes might be in the first trip, the second trip, or (in rare cases) even the third trip. This algorithm runs at the RDA signal processor, which is literally a black box, for VCPs 11, 12, 21, 31, and 32.

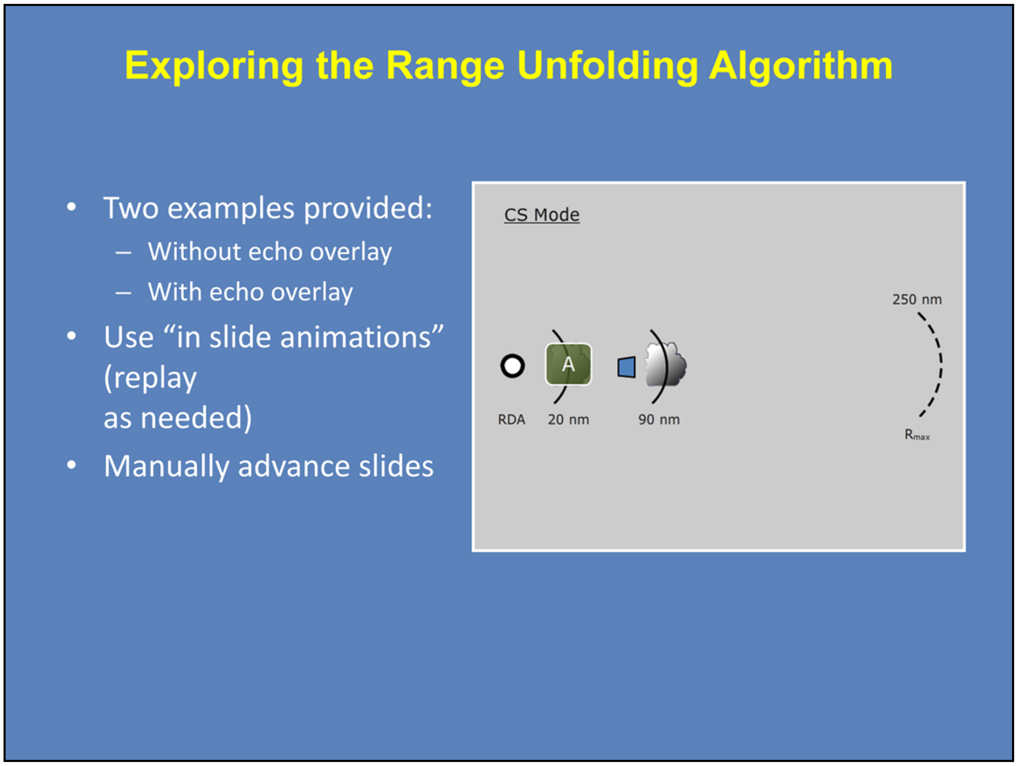

The next several slides will demonstrate two examples of how the Range Unfolding algorithm works: One with no echo overlay and one with overlay. We will walk through each example using numerous steps. The steps will be shown using in slide animations that you can replay as many times as you like, if needed. For each of these slides, you will need to manually advance the slide to move to the next step. We also provide some troubleshooting steps on the next slide in case the animations do not work initially for you.

Back to Top

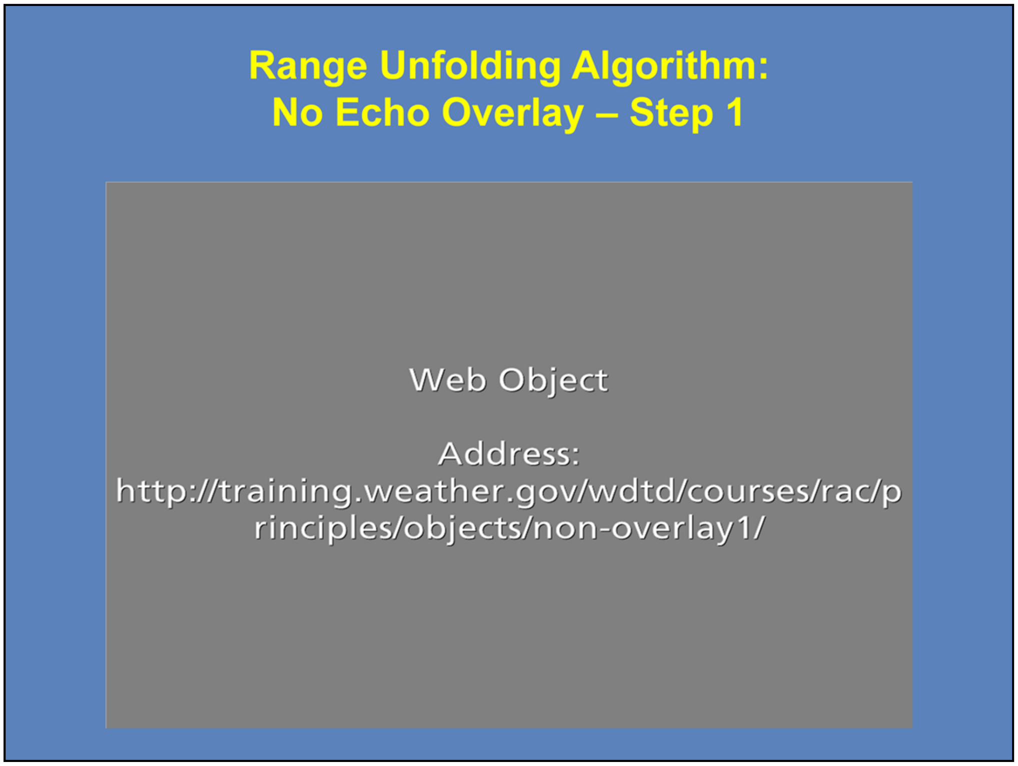

Step 1: We look down a single radial in Continuous Surveillance mode. That means we have a nice long unambiguous range. We have two targets along this radial. The first one is at 20 nm, and the second one is at 90 nm. Since both of these targets are within the first trip, we know their true range and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

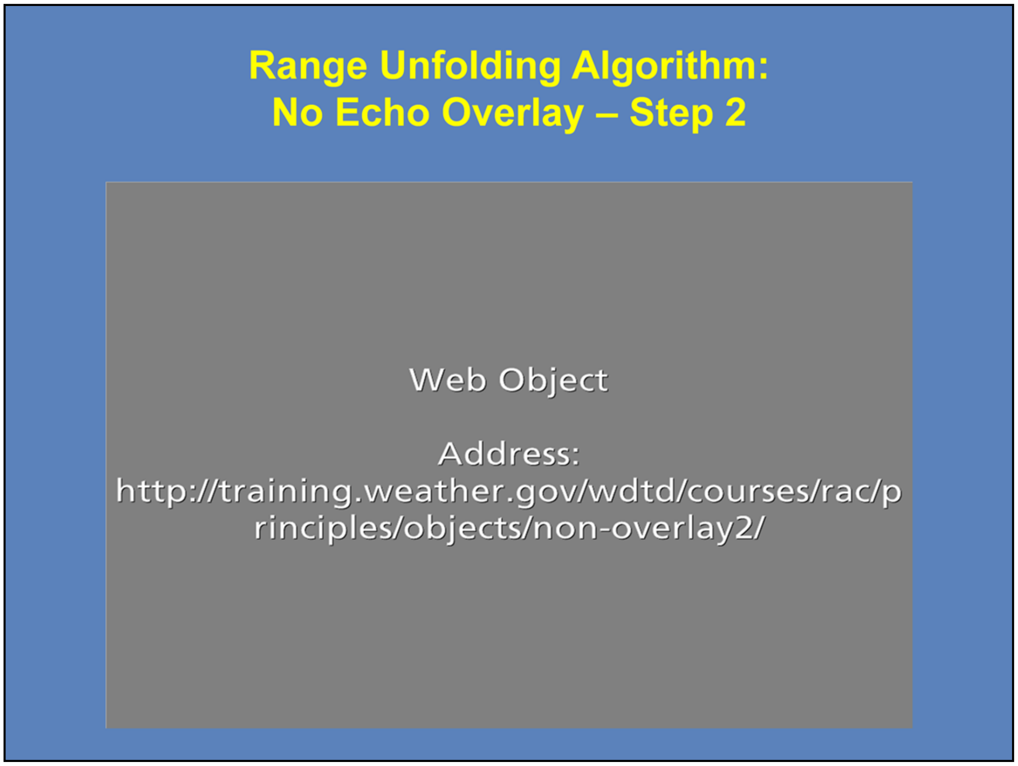

Step 2. Prior to the switch to Continuous Doppler mode on the next rotation, the algorithm computes what the apparent ranges of these two targets will be. In this case, Our R_{max} of 60 nm in CD mode results in target A staying at 20 nm while target B will be folded into an apparent range of 30 nm.

Range Unfolding Algorithm: No Echo Overlay – Step 3

Web Object

Address: http://training.weather.gov/wdtd/courses/rac/principles/objects/non-overlay3/

Step 3: We know target B has an apparent range of 30 nm. At this step, the algorithm computes all the possible ranges for each target that shows up in the CD mode first trip, and these possible range go out into the second and third trips.

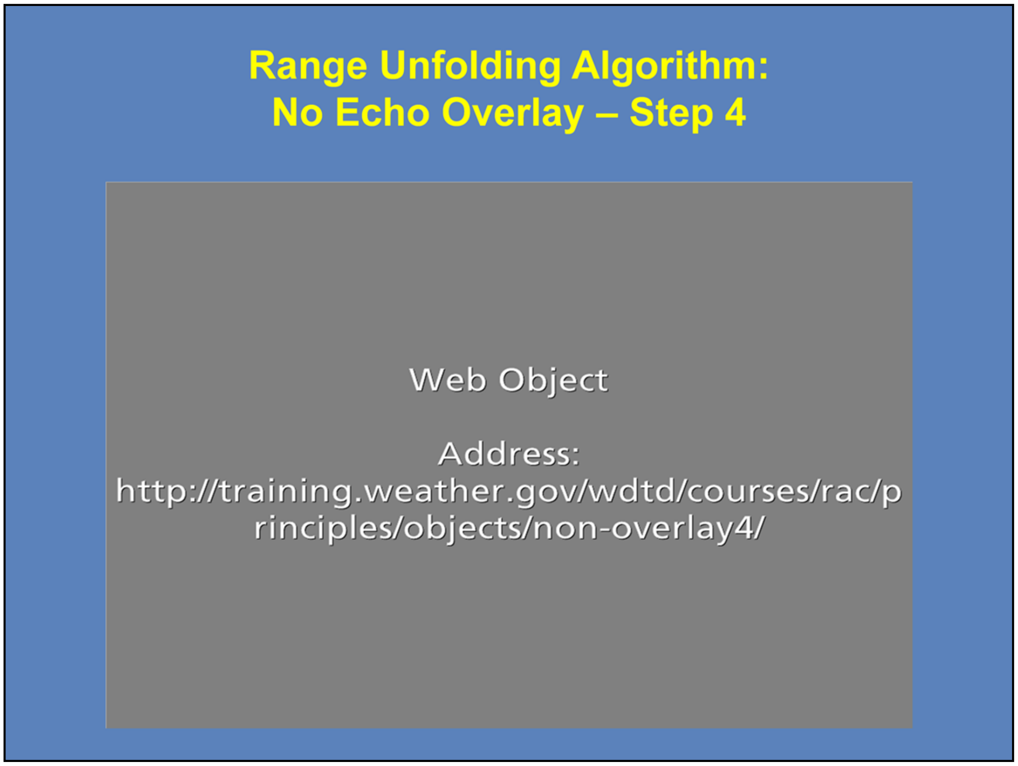

Step 4. With those calculations out of the way, the radar can collect data in Continuous Doppler mode. The energy from pulse 1 for target B (which was at 90 nm) propagates back to the RDA, but it does not arrive before the energy from pulse 2 is transmitted. That is how we get an apparent range for target B at 30 nm in CD mode.

Back to Top

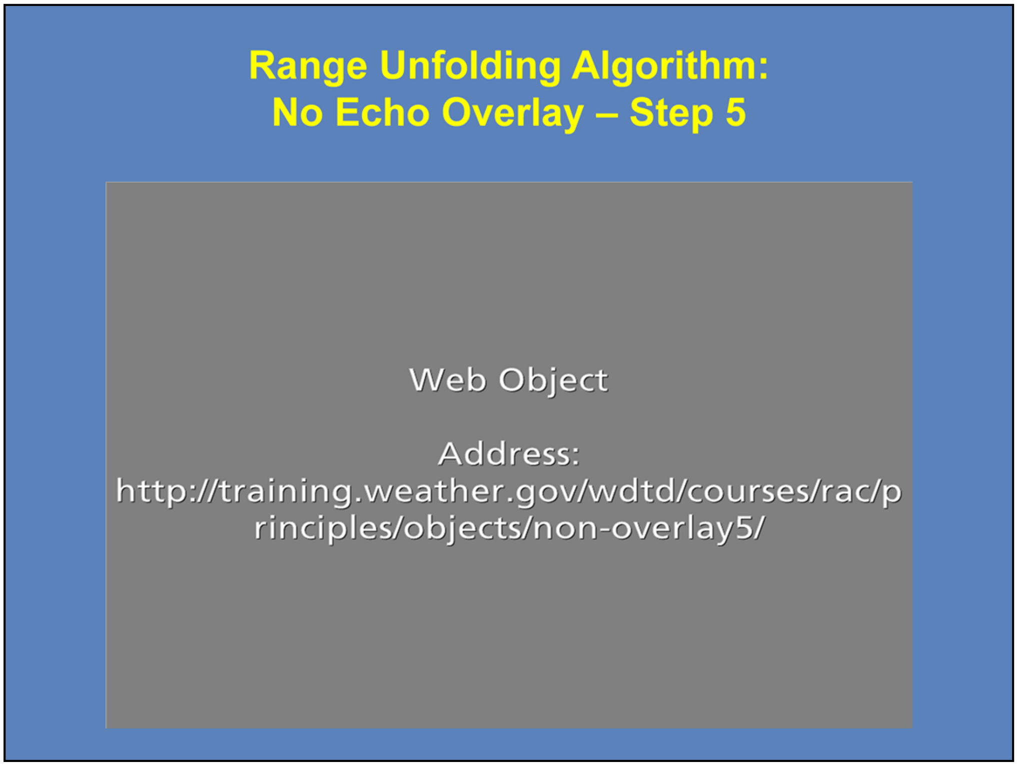

Step 5: Now that the RDA can begin the actual down radial comparison of the CS and CD data. The algorithm checks to see wherever there is a target in the CD data if there was something at that same range in CS data. So, when we get to target A, the algorithm checks if there was something at 20 nm in the CS data? And the answer is Yes, so we know the velocity for target A belongs at the range of 20 nm.

Back to Top

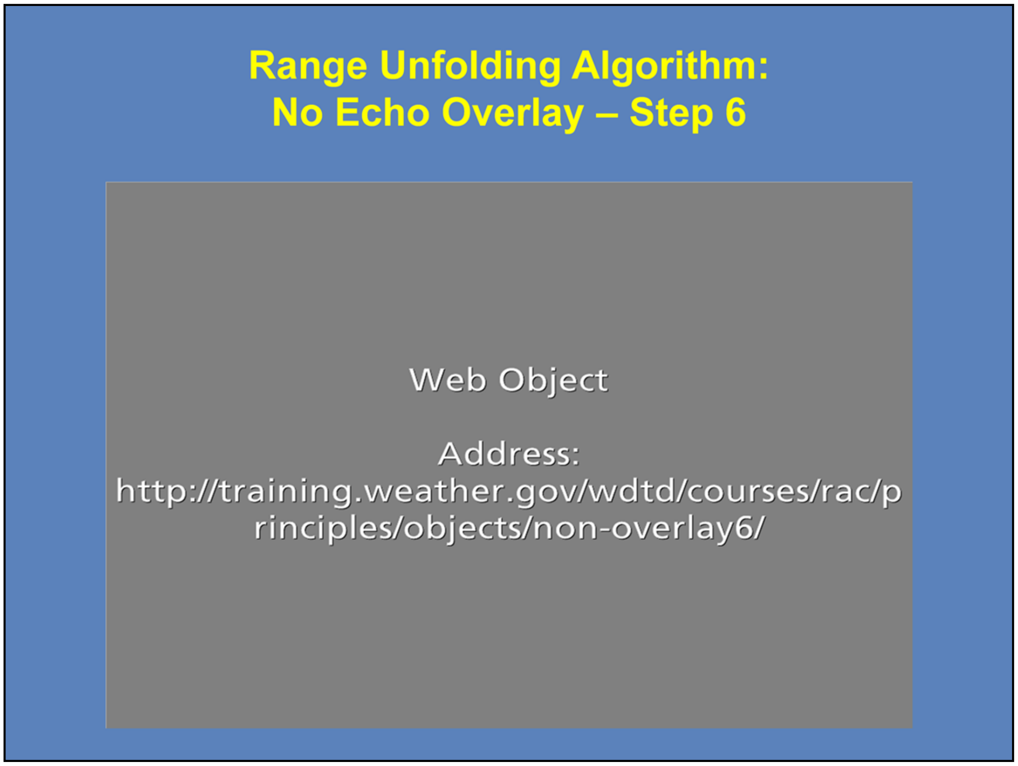

Step 6: When the algorithm gets to target B at an apparent range of 30 nm, it compares that location to the CS data. In this case, no there wasn’t a target at 30 nm in the CS data.

Back to Top

Range Unfolding Algorithm: No Echo Overlay – Step 7

Web Object

Address: http://training.weather.gov/wdtd/courses/rac/principles/objects/non-overlay7/

Step 7: Since no targets were detected at 30 nm in the CS data, the algorithm checks the next possible range which is 90 nm. And, yes, a target was detected 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.

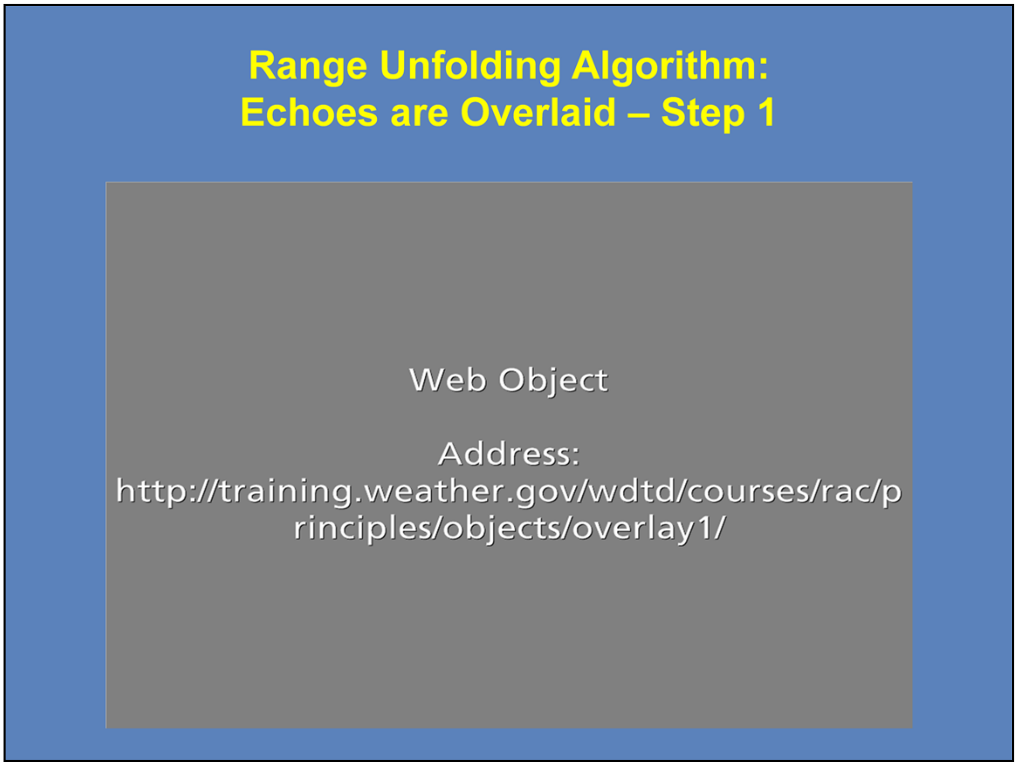

Now let’s look at our second example where echoes are overlaid on top of each other.

Step 1: We start off with the same initial conditions. In Continuous Surveillance mode, the RDA detects targets at 20 and 90 nm. The maximum unambiguous range is 250 nm, so we know the returned power and range for each of these targets.

Back to Top

Range Unfolding Algorithm: Echoes are Overlaid – Step 2

Web Object

Address:
http://training.weather.gov/wdtd/courses/rac/principles/objects/overlay2/

Step 2: This time, when the radar switches to Continuous Doppler mode, the maximum unambiguous range will be 70 nm. That means both of the targets will have an apparent range of 20 nm. It also means that pulses are returning back to the RDA at the same time, so this is an overlay situation.

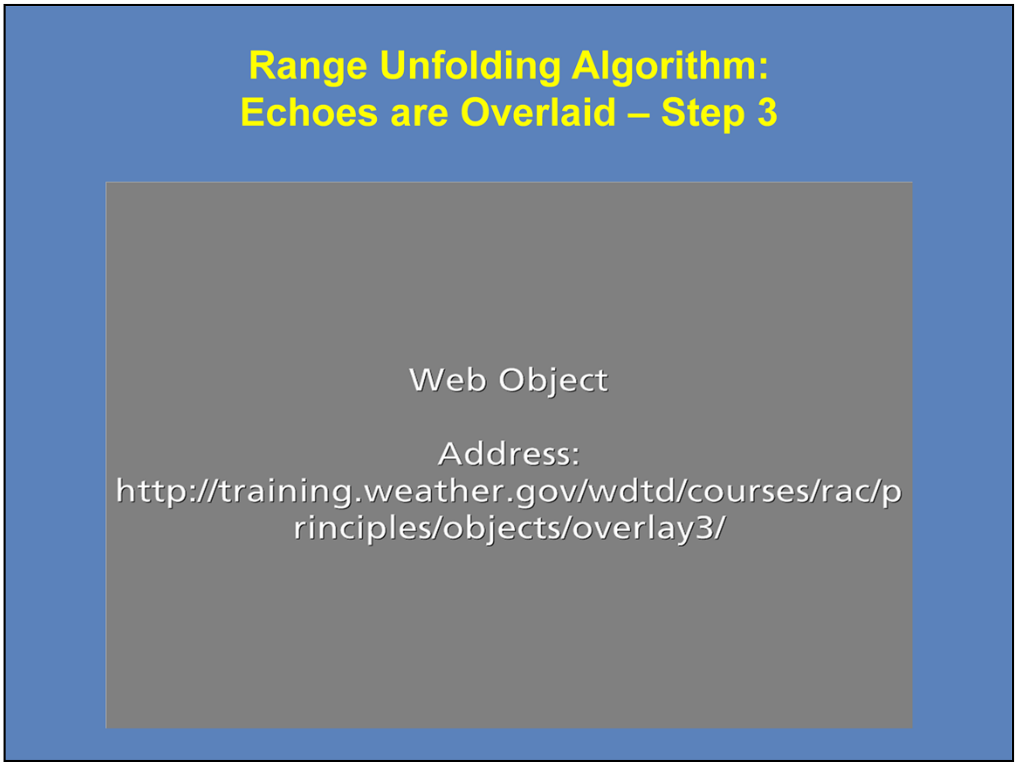

Step 3: Just as before, the algorithm accounts for the apparent ranges of the targets across the first, second and third trips.

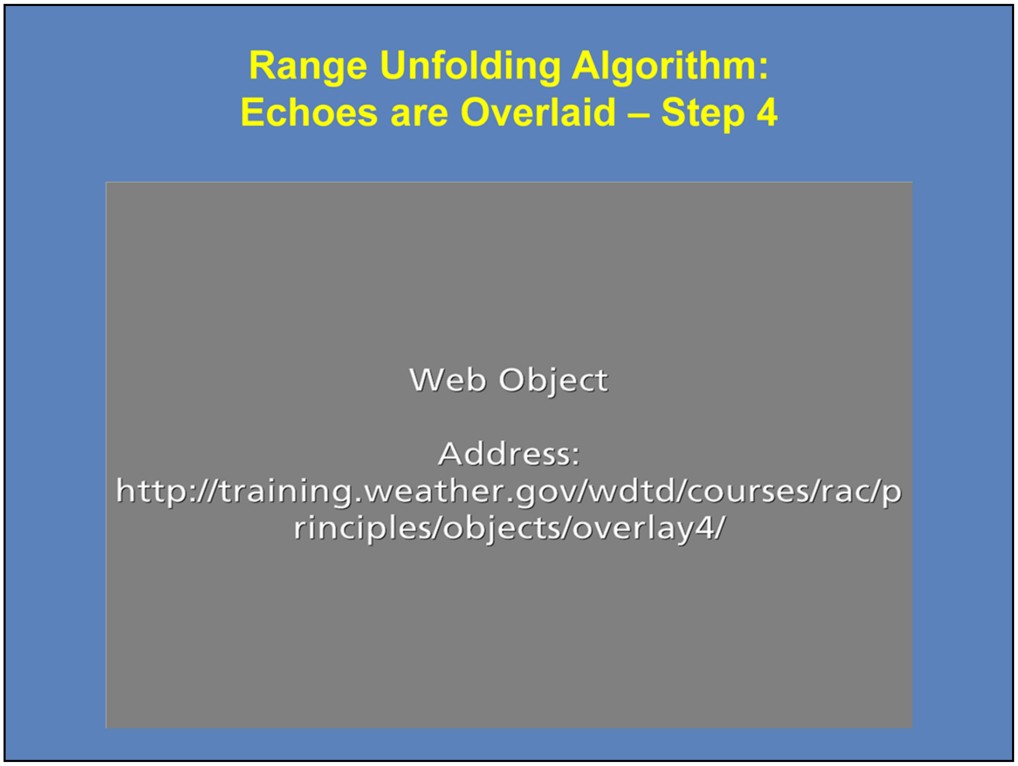

Step 4: Here's where the algorithm takes advantage of the fact that it knows the returned power for each of these two targets and their original range. The algorithm compares those returned powers to each other, and it computes a threshold (called TOVER) to determine whether the velocity and spectrum width data are assigned to one of these two targets. So, the higher powered target is compared to the lower power target, and if difference reaches this threshold, the algorithm can assign the velocity and spectrum width values to one of these overlaid targets.

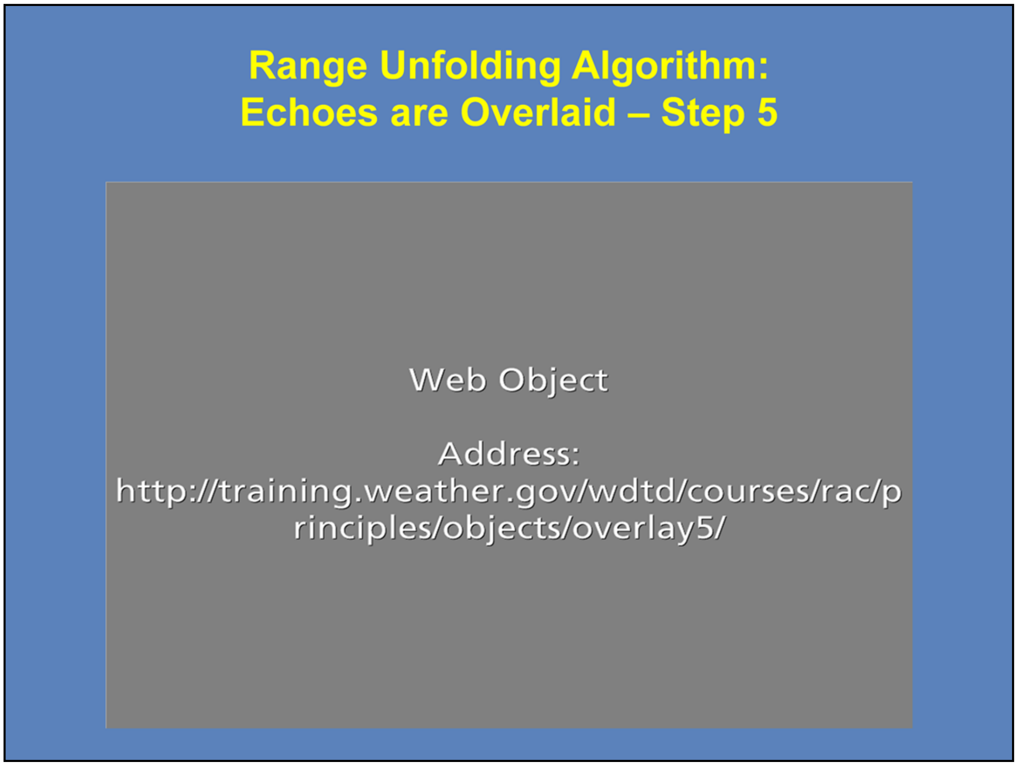

Step 5: Now the actual Doppler data collection begins, and notice that information from pulses 1 and 2 come back to the radar at the same time. So the velocity and spectrum width values have an apparent range of 20 nm, that is composed of data from returned pulses for both targets A and B.

Range Unfolding Algorithm: Echoes are Overlaid – Step 6

Web Object

Address: http://training.weather.gov/wdtd/courses/rac/principles/objects/overlay6/

Step 6: Now the algorithm begins the comparison of the Continuous Doppler to the Continuous Surveillance data. The algorithm answers the question whether a target was at a specific range, but also determines if that target that had the higher returned power. In this particular example, Target B had the lower power, so the answer is “No” at a range of 20 nm.

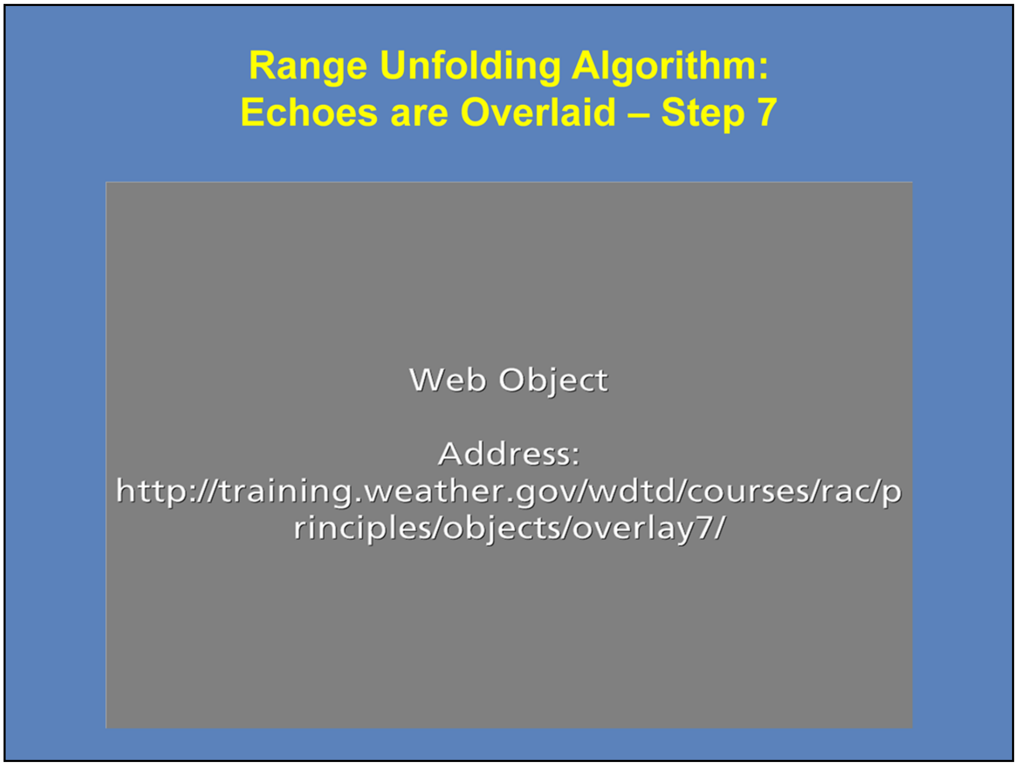

Step 7: Since the answer was no the first time, the algorithm looks at other possible ranges for this target. Was there something in the CS data at 90 nm? Yes. And was this the higher power target? Yes.

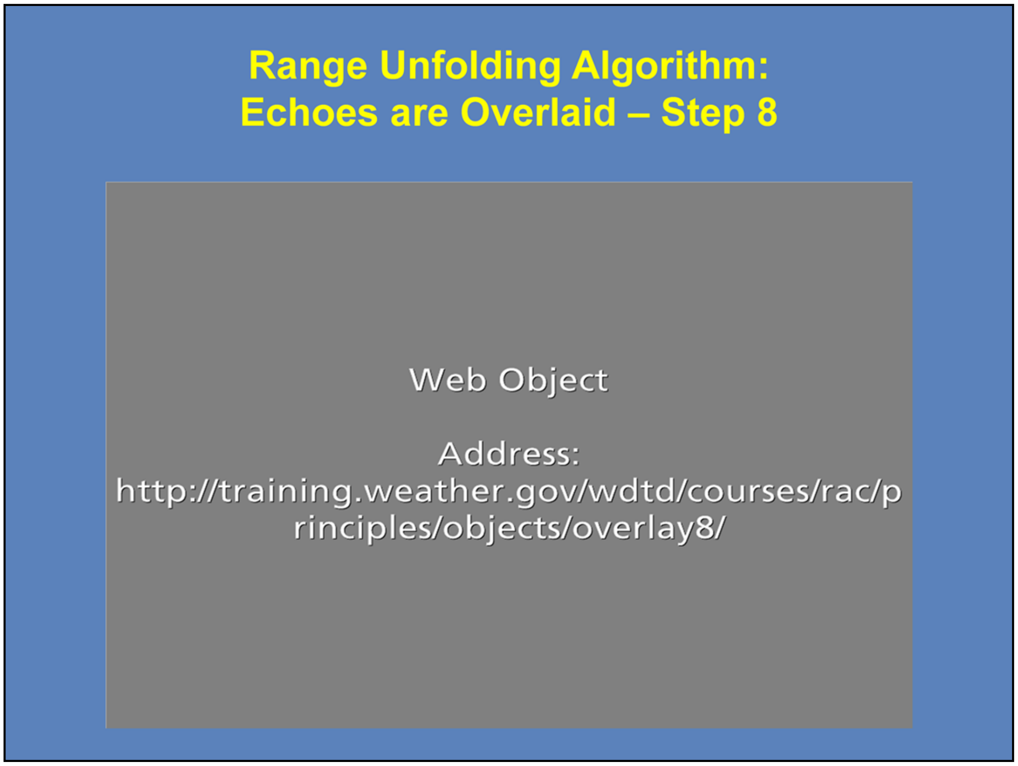

Step 8: Now that we know this velocity and spectrum width data at a range of 20 nm is more representative of the target that was at 90 nm. So, we assign the observed data to the bin at 90 nm, and range folding, or purple, to the echoes at 20 nm.

Back to Top

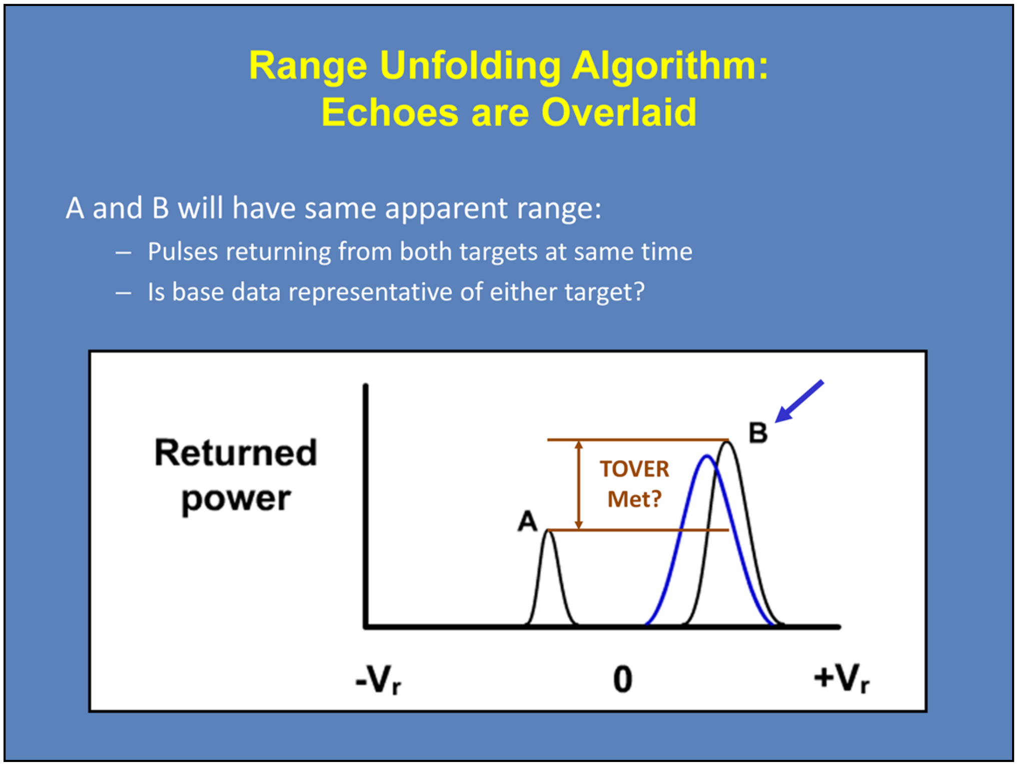

This graphic represents the base data estimation process for a range bin when there is an echo overlay case such as in our previous example. Pulses from two different targets are received and processed at the same time, apparently from 20 nm. If the pulses from these two different targets could be analyzed separately, the result would be the black bell curves for targets A and B. However, the returned pulses 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 since it is returning significantly more power than target A. Recall that velocity estimates are power weighted, so this “hybrid” velocity estimate will more closely resemble the more “powerful” target B.

Lastly, the algorithm needs to determine if the power returned by B is sufficiently greater than A for the “hybrid” velocity to be assigned to target B.

Back to Top

Range Unfolding Algorithm: Echoes are Overlaid

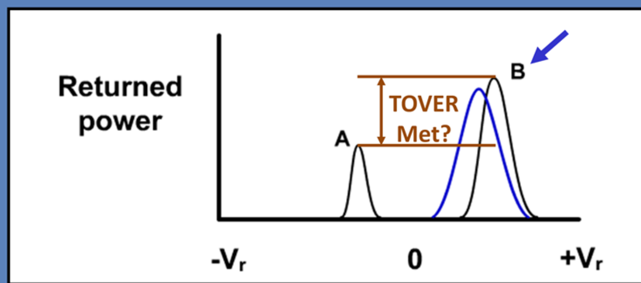

- If power ratio exceeds TOVER (5 dB)
  - V & SW assigned to echo with higher power
  - Other echo assigned RF
- If power ratio does not exceed TOVER
  - Both echoes assigned RF

The Range Unfolding Algorithm uses a parameter known as “threshold over”, or TOVER, to determine if the overlaid echoes have a sufficiently high difference in returned power to assign the hybrid velocity to one of them.

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

If the power ratio does not exceed TOVER, the algorithm assigns RF to both of the overlaid echoes.

Back to Top

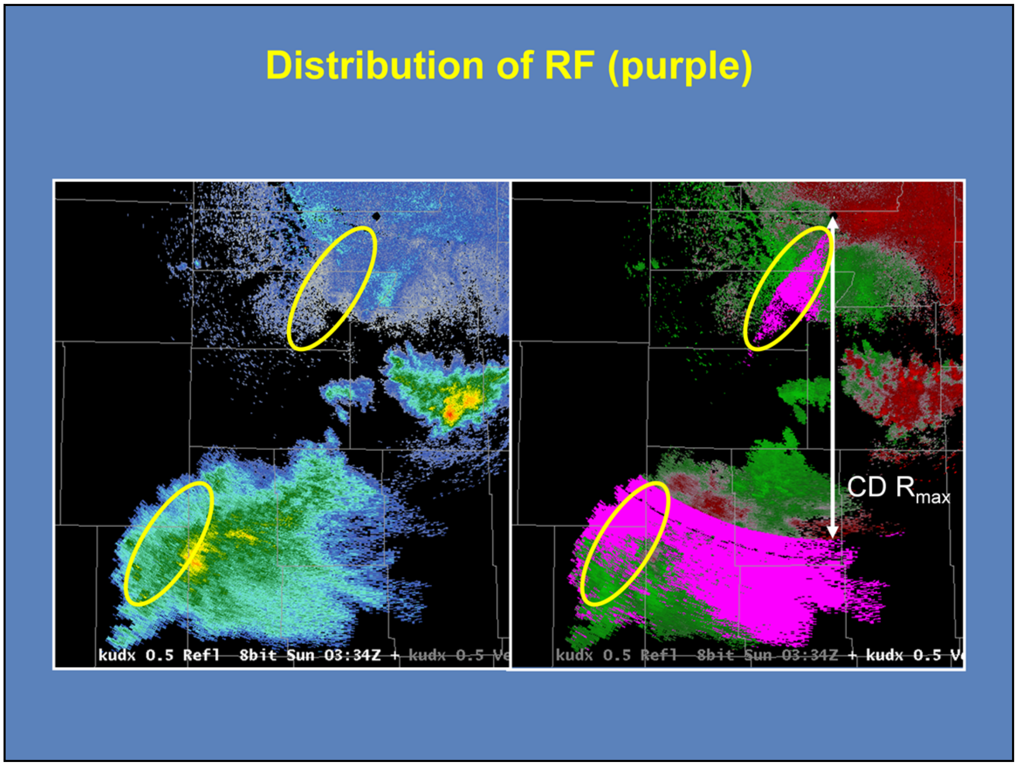

Now for another look at the distribution of RF in Base Velocity, given the location of the echoes in Base Reflectivity. The strongest storms south of the radar are within the first trip, so there is no issue with the availability of the associated velocity data. There is an area of weaker echo to the south and southwest of the radar, some of which falls within both the first and second trip of the Continuous Doppler scan data.

The yellow ovals highlight a group of gates in the same relative positions in the first and second trips, with purple assigned to both of the bins in most cases. In general, 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.

Legacy Range Unfolding Algorithm: Strengths

1. Places V & SW at proper range
   - May be beyond CD R_{max}
   - Echoes overlaid and TOVER exceeded:
     - One assigned V and SW
     - Other assigned RF
2. Mitigating “Doppler Dilemma”
   - Low PRF for target range and intensity
   - High PRF for velocity and spectrum width

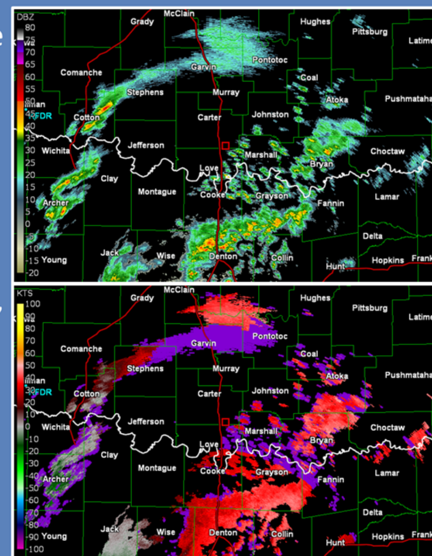

Now for a summation of Legacy Range Unfolding Algorithms strengths. In most instances, this algorithm achieves its objective of assigning velocity and spectrum width data to the appropriate range, which may be beyond the R_{max} of the Continuous Doppler (CD) scan. When echoes are overlaid, the Legacy Range Unfolding Algorithm can, at best, assign velocity and spectrum width to one of the overlaid echoes. The algorithm decides whether to assign the values based on the returned power of the separate echoes that contribute to the overlay. If TOVER is exceeded, then one of the echoes has returned sufficiently more power that the others and can be assigned the velocity and spectrum width values. The other echoes are assigned RF.

The Legacy Range Unfolding Algorithm mitigates the Doppler Dilemma in this manner. 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 more accurate velocity and spectrum width data. Comparing both data sources allows for “unfolding” the velocity and spectrum width to its appropriate range along the radial.

Legacy Range Unfolding Algorithm: Limitations

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

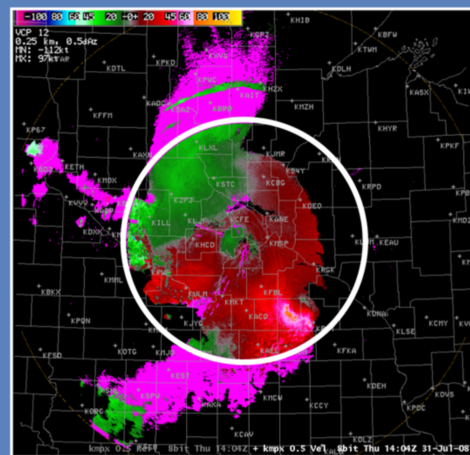

Extensive echo coverage aligned along radials limits the Legacy Range Unfolding Algorithm ability to unambiguous unfold velocity CD data. These situations maximize echo overlay along a radial through the first, second, and third trips. The Legacy Range Unfolding algorithm can only assign overlaid data to one of the echoes, so the result of these situations is extensive range folded data. If the returned power from the overlaid echoes is too similar, than both echoes will be assigned RF values.

This example shows a case where the TOVER requirement was met for most of the echoes in the first trip.

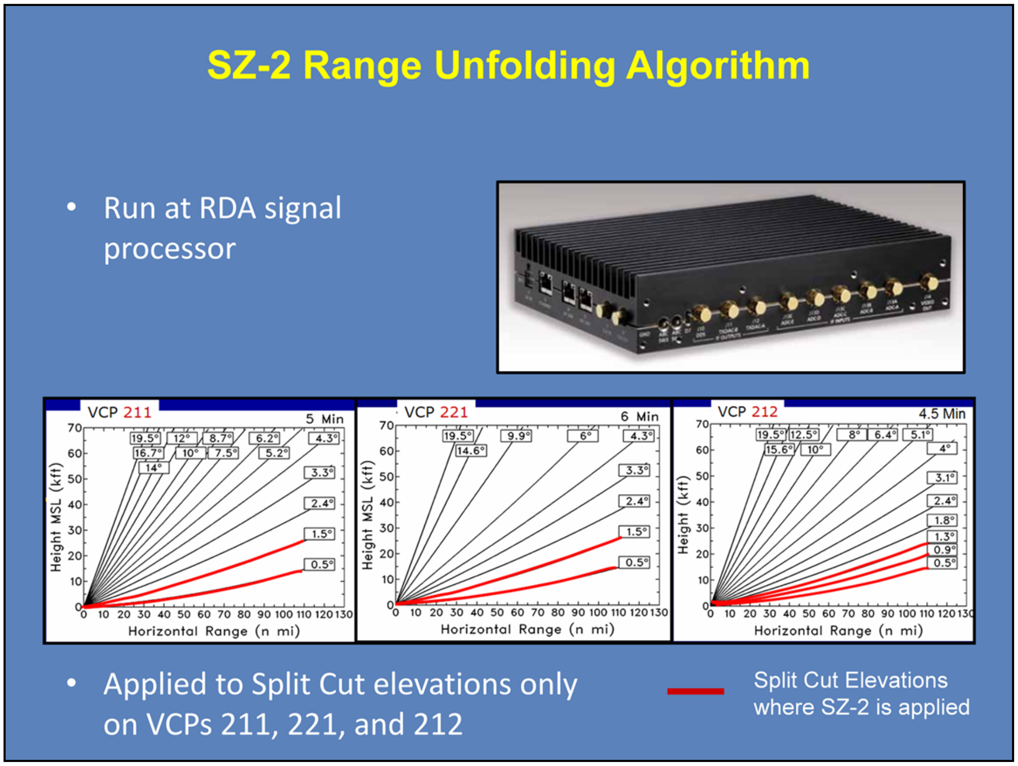

The second range unfolding technique for the velocity and spectrum width data is known as SZ-2 Range Unfolding and was fielded in 2007. This technique is named for the two research scientists who developed it: Mangalore Sachidananda and Dusan Zrnic. This algorithm runs at the RDA signal processor, the literal black box, just like the Legacy Range Unfolding Algorithm.

The SZ-2 algorithm is available only for the Split Cut elevations of the three SZ-2 VCPs: 211, 221, and 212. In fact, the first “2” in these VCP names indicates it uses the SZ-2 algorithm. The remaining numbers indicates the other characteristics of the VCP. For example, VCP 212 is simply the VCP 12 that uses the SZ-2 Range Unfolding Algorithm velocity and spectrum width on the Split Cuts.

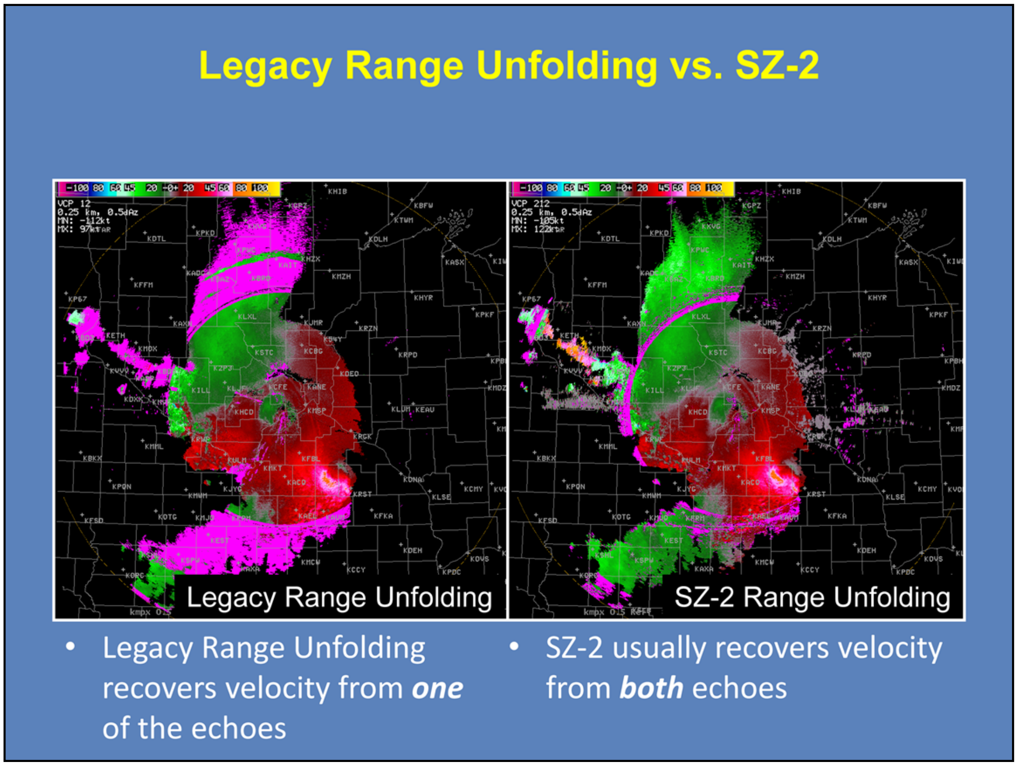

The SZ-2 Range Unfolding algorithm is particularly effective with widespread echo coverage. In this case, the Legacy Range Unfolding algorithm (VCP 12) is used for the data on the left, with SZ-2 (VCP 212) on the right. Use of VCP 212 dramatically improves the availability of velocity and spectrum width. Notice from the data that Continuous Doppler data are available from multiple trips over large areas, resulting in extensive echo overlay. Where there is echo overlay, the Legacy Range Unfolding Algorithm can recover, at best, one of the echoes. On the other hand, the SZ-2 technique usually recovers both overlaid echoes.

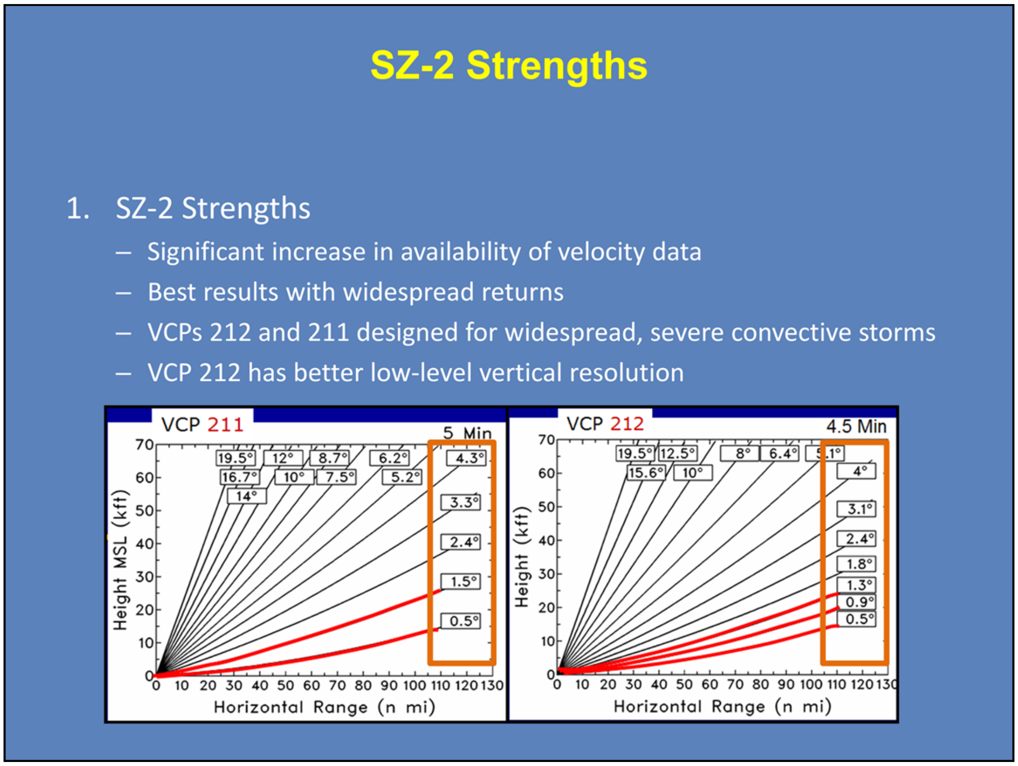

The SZ-2 Range Unfolding Algorithm's strength results from its increased availability of velocity data for multiple trips. The best results come with events with widespread echo coverage. VCPs 212 and 211 are designed for widespread rapidly evolving severe convective storms. However, VCP 212 has superior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

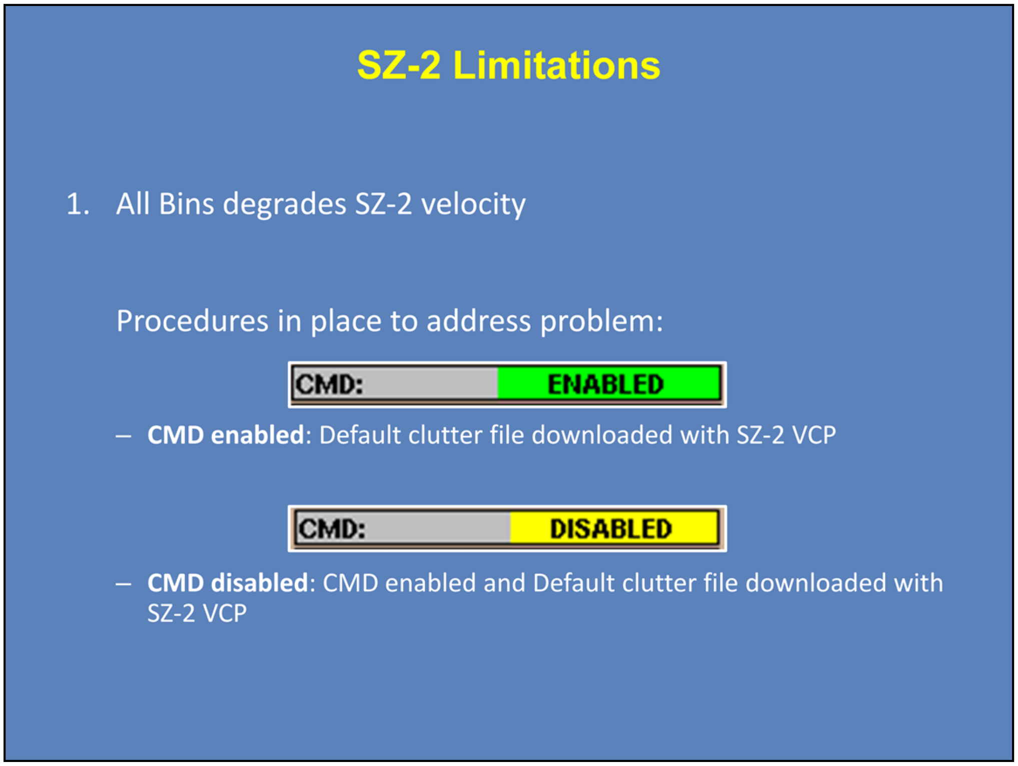

With the Clutter Mitigation Detection algorithm active, the radar almost never needs All Bins clutter suppression. SZ-2 usage adds another reason to avoid All Bins suppression, as All Bins suppression can degrade the velocity data when SZ-2 is applied. In fact, the RPG software has built in procedures to avoid having All Bins and SZ-2 active at the same time.

If CMD is enabled and an SZ-2 VCP gets downloaded to the RDA, the Default clutter regions file will be downloaded along with it. This procedure ensures that the Bypass Map is in control everywhere, overwriting any All Bins filtering that may be in use at that time. If CMD has been disabled and an SZ-2 VCP gets downloaded to the RDA, CMD will be enabled automatically, and the Default clutter regions file will again be downloaded along with the SZ-2 VCP.

Back to Top

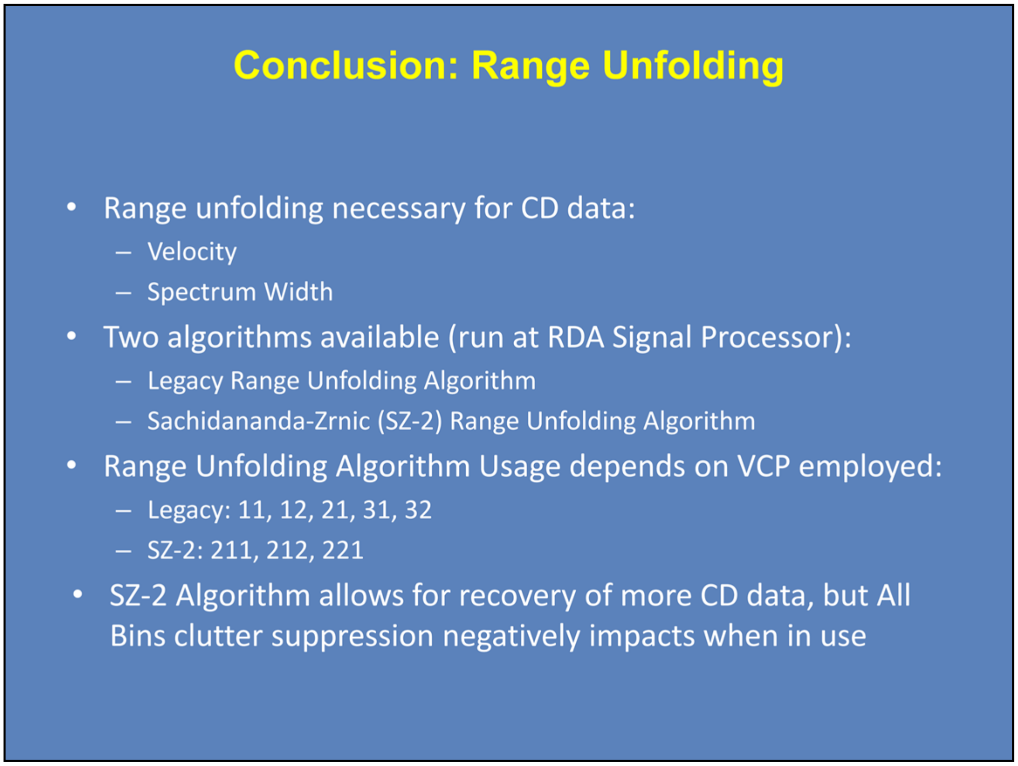

To wrap things up, the higher PRFs used during the Continuous Doppler scan makes range unfolding of velocity and spectrum width data necessary because of the short maximum unambiguous ranges. The WSR-88D uses two different range unfolding algorithms at the signal processor to handle this process. The Legacy Range Unfolding Algorithm addresses the issue for VCPs 11, 12, 21, 31, and 32. The Sachidananda-Zrnic (or SZ-2) algorithm unfolds the CD scan data for VCPs 211, 212, and 221. The SZ-2 Algorithm generally outperforms the Legacy algorithm because it allows for the recovery of more CD velocity and spectrum width data during overlay situations. However, the downside to the SZ-2 Algorithm is All Bins clutter suppression negatively impacts data quality when SZ-2 is in use.

Back to Top

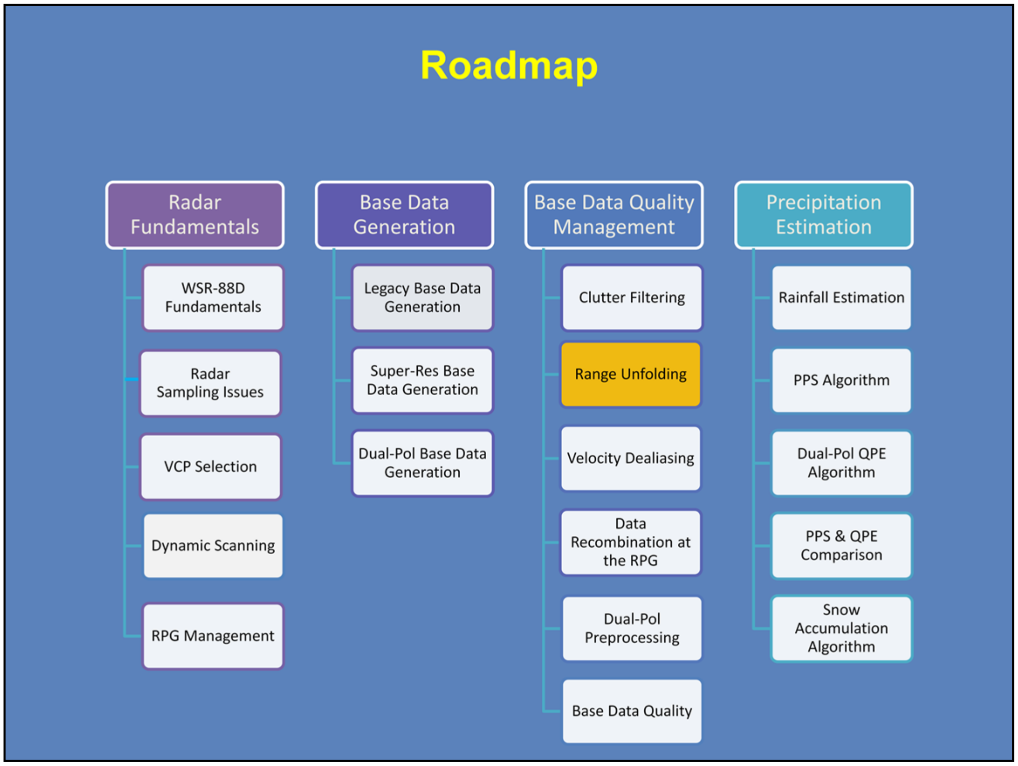

You have now completed this lesson on the range unfolding of velocity data. You are ready to proceed to the next lesson.

Back to Top

Welcome to this lesson on Velocity Dealiasing.

Back to Top

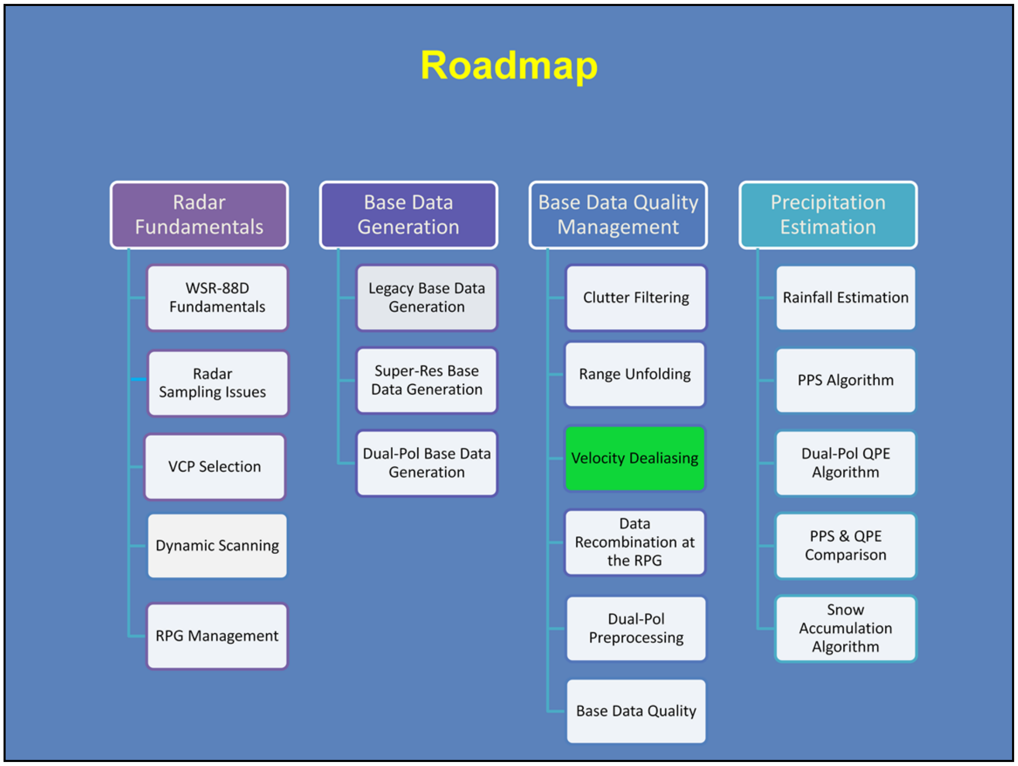

Here is the “roadmap” with your current location.

Back to Top

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.

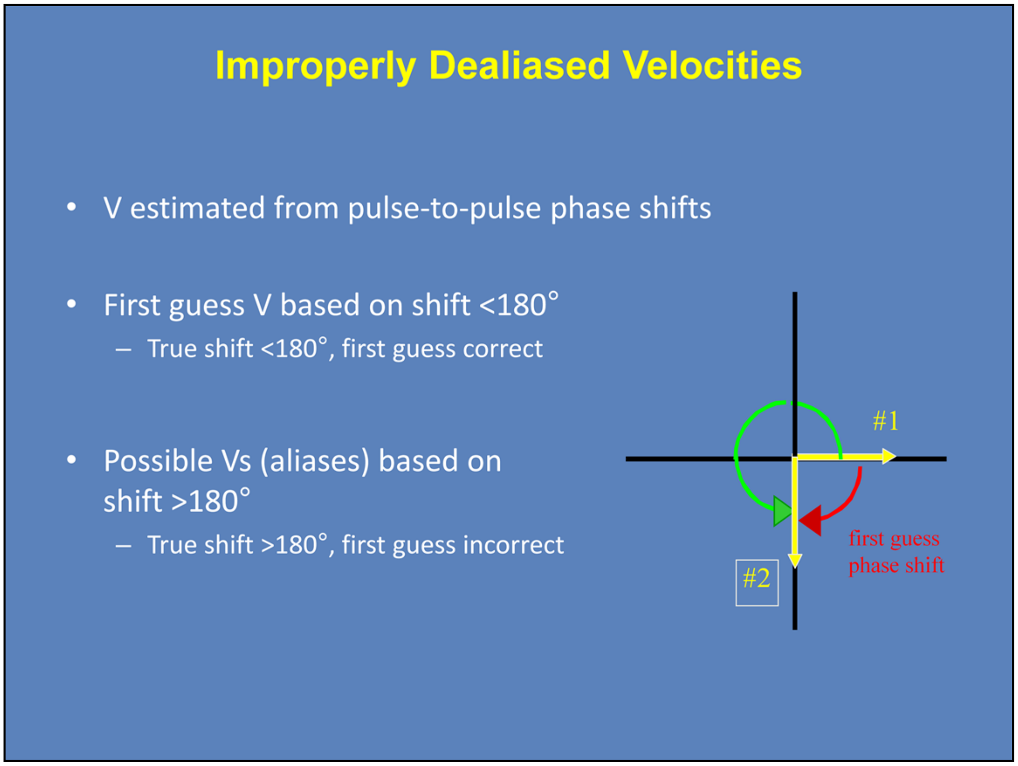

The previous lesson presented the process for assigning velocity and spectrum width to its appropriate range. The radial velocity value itself is the first guess velocity, which may not be correct. Perhaps one of the aliases of that first guess is the correct velocity. The example used in Base Data Generation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

Back to Top

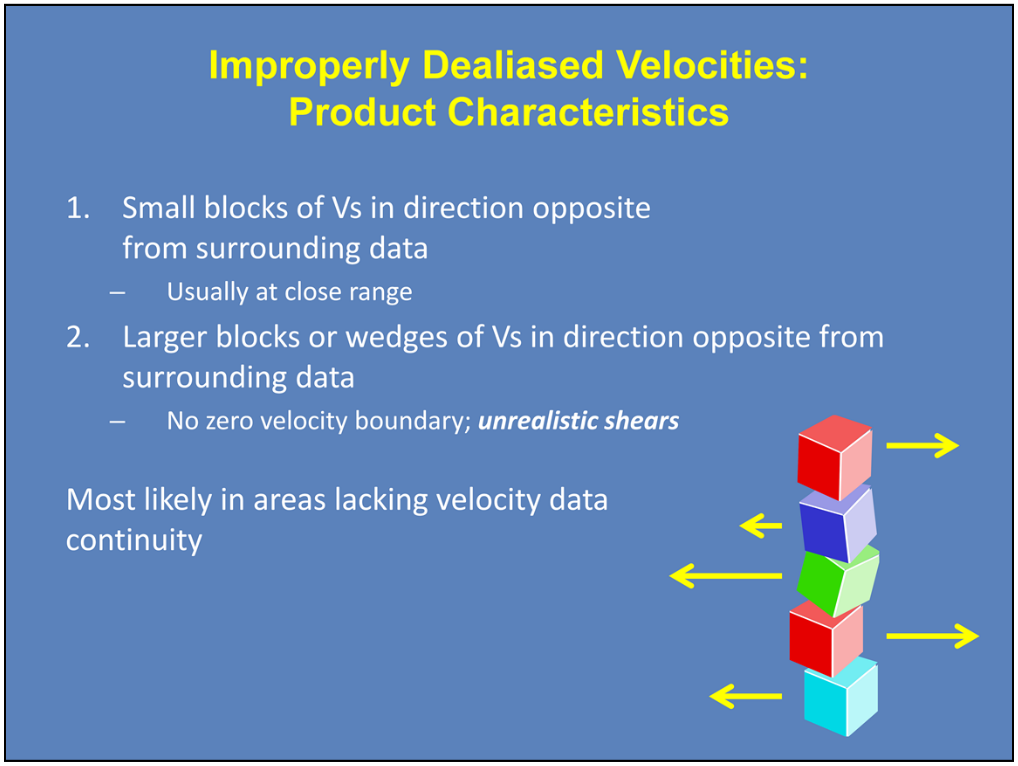

Before looking at the algorithm that “de-aliases” velocities, we first look at the impact of improperly dealiased velocities on the radar products. There are two types of improperly dealiased velocities. At close range, especially in residual ground clutter, small blocks of velocity values opposite in 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.

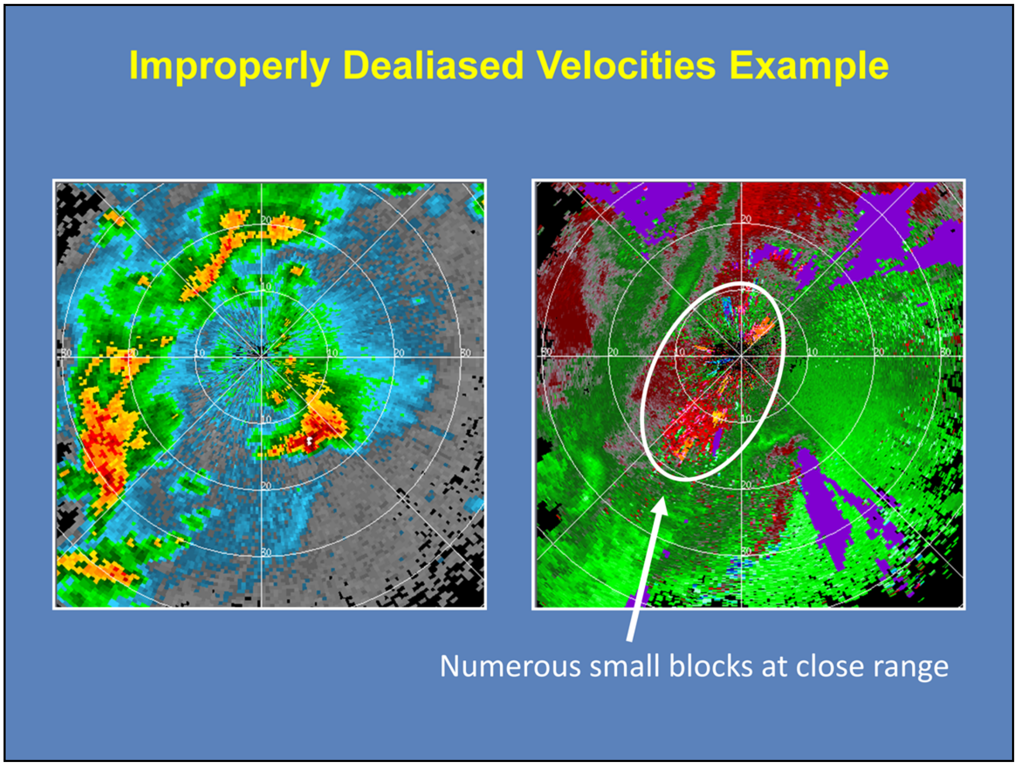

Here is an example of the type of improperly dealiased velocities frequently seen near the radar in the residual clutter region. They are usually not operationally significant.

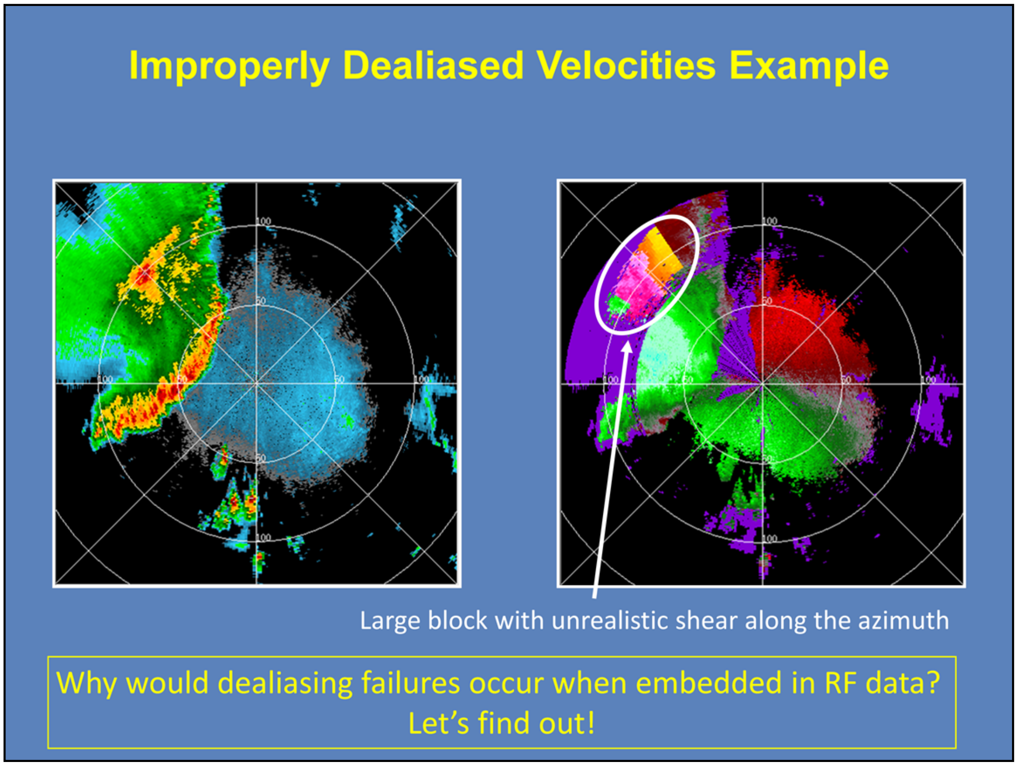

This is an example of improperly dealiased velocities resulting in two azimuthal shears that are clearly not meteorological. Another term for this data artifact is dealiasing failures. The fact that these failures occur while embedded in an area where RF has been assigned is related to how the dealiasing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

Back to Top

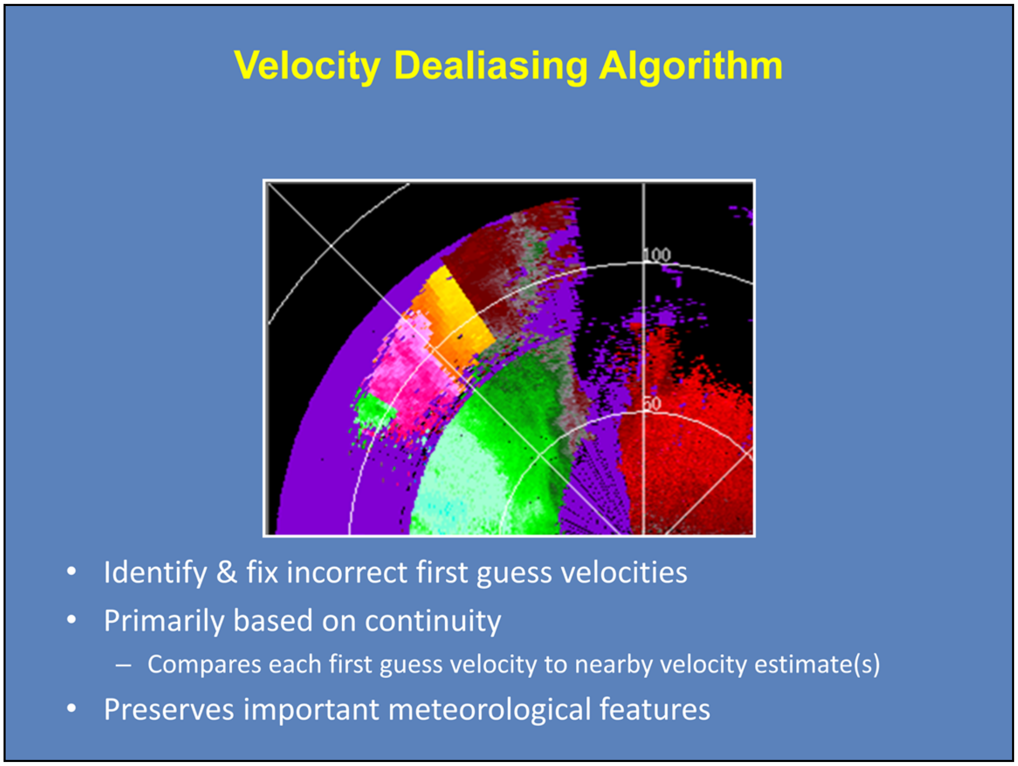

The goal of the Velocity Dealiasing algorithm is to assign the correct radial velocity to each range bin. It first identifies any first guess velocities that are “suspect”, then attempts to assign one of the aliases if that is more appropriate. The decision making for what is “appropriate” comes from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

Back to Top

Velocity Dealiasing Algorithm: "Check Your Neighbor"

- 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.

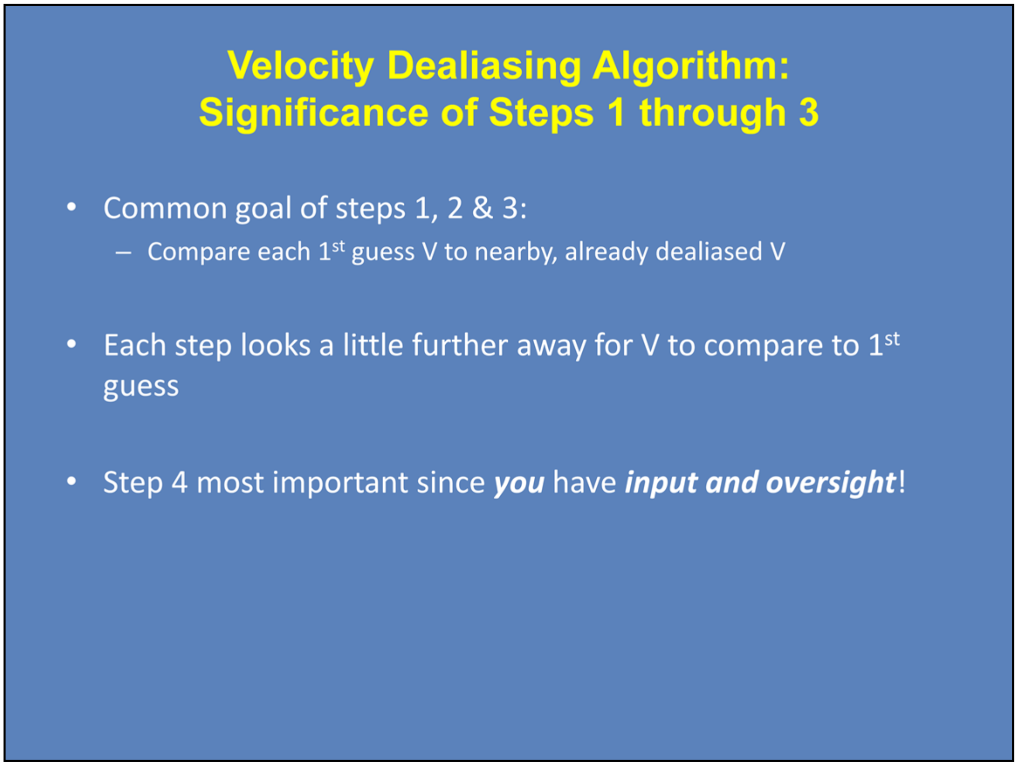

The first three steps of the Velocity Dealiasing algorithm have the common goal of comparing each first guess velocity to a “velocity neighbor”. That “velocity neighbor” varies with each step, getting a little further away from the first guess velocity. With respect to everyday operations, the most important 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.

Back to Top

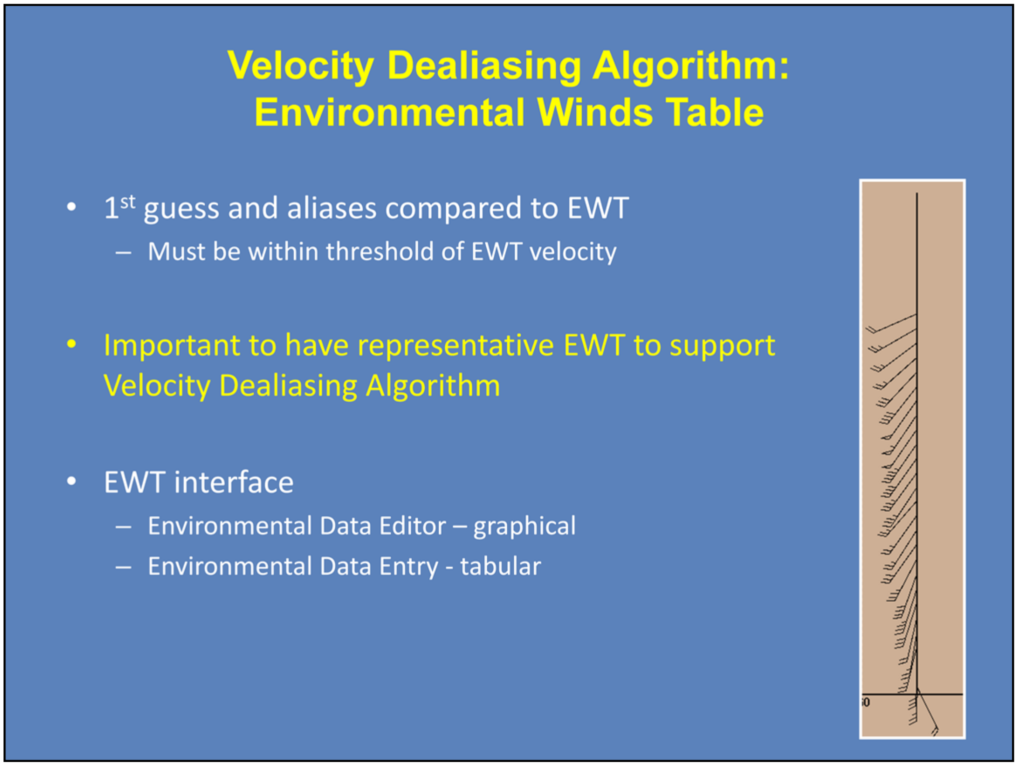

When the previous steps do not resolve a possible incorrect first guess velocity, the Velocity Dealiasing algorithm accesses the Environmental Winds Table to find a representative velocity for the given range and height. It is important that the Environmental Winds Table, which is stored and updated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

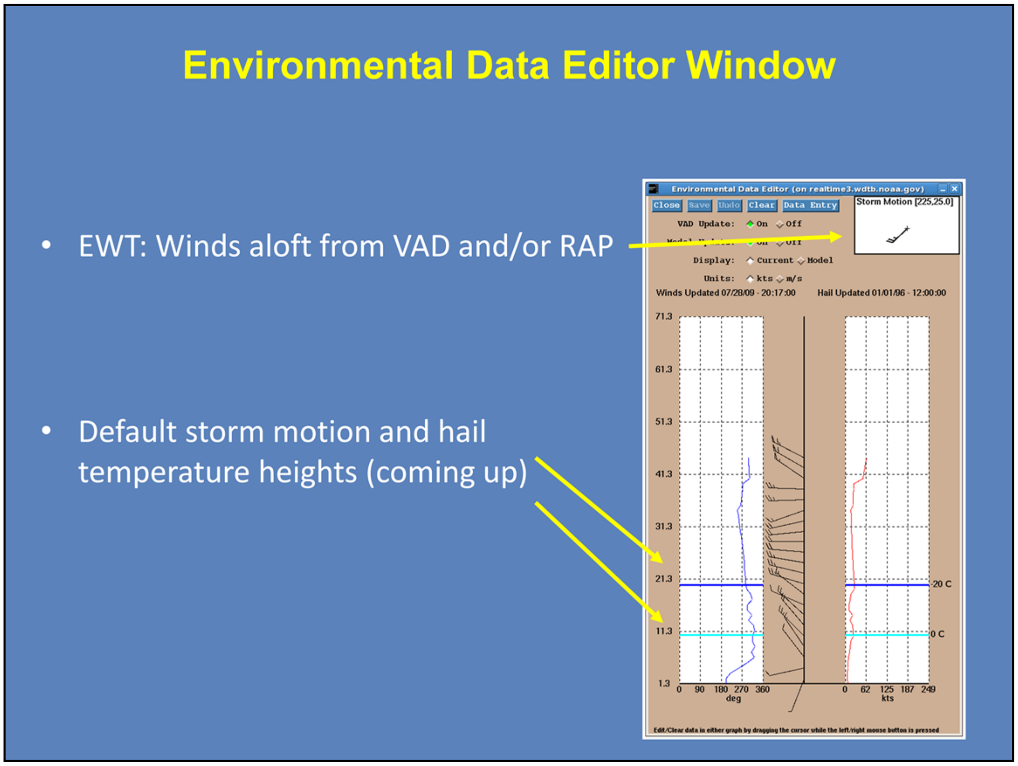

When initially accessing Environmental Data, this window provides multiple types of data. The local winds aloft is presented with the wind barbs in the center, as well as the wind direction and speed represented on the graphs. The input for the wind information is usually a combination of input from 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.

Back to Top

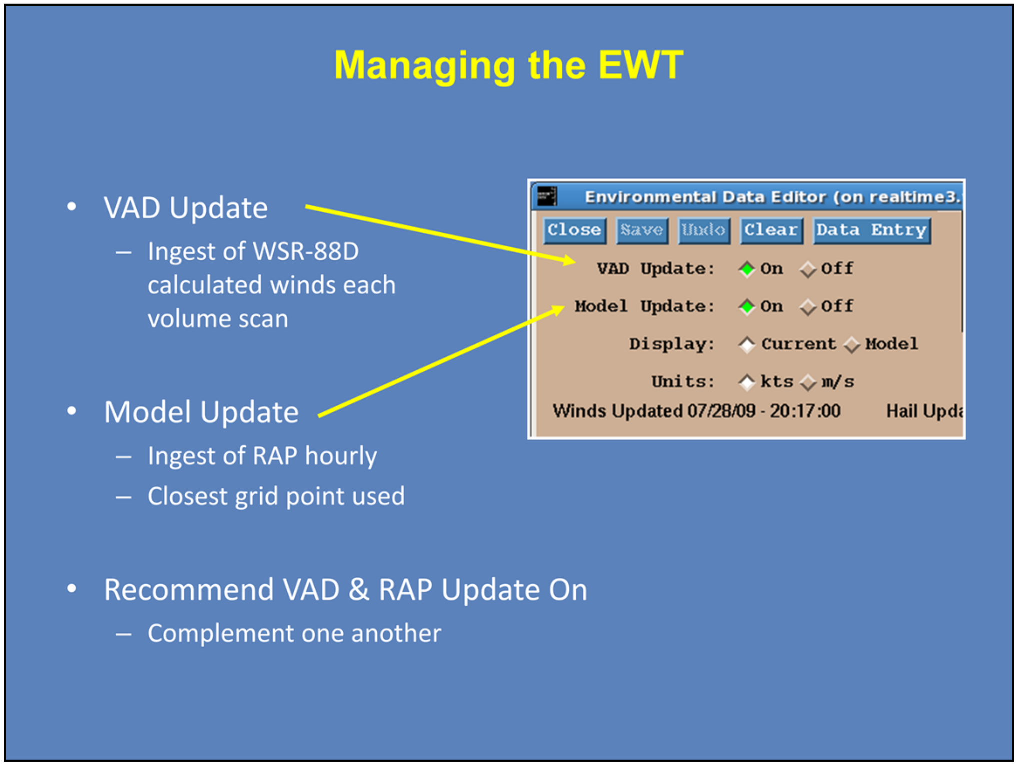

Here's a closer look at the buttons that control inputs to the Environmental Winds Table. When VAD Update is set to On, the EWT is updated every volume scan by VAD-generated winds aloft. When Model Update is set to On, the RAP data for the closest grid point are used hourly to update the EWT.

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.

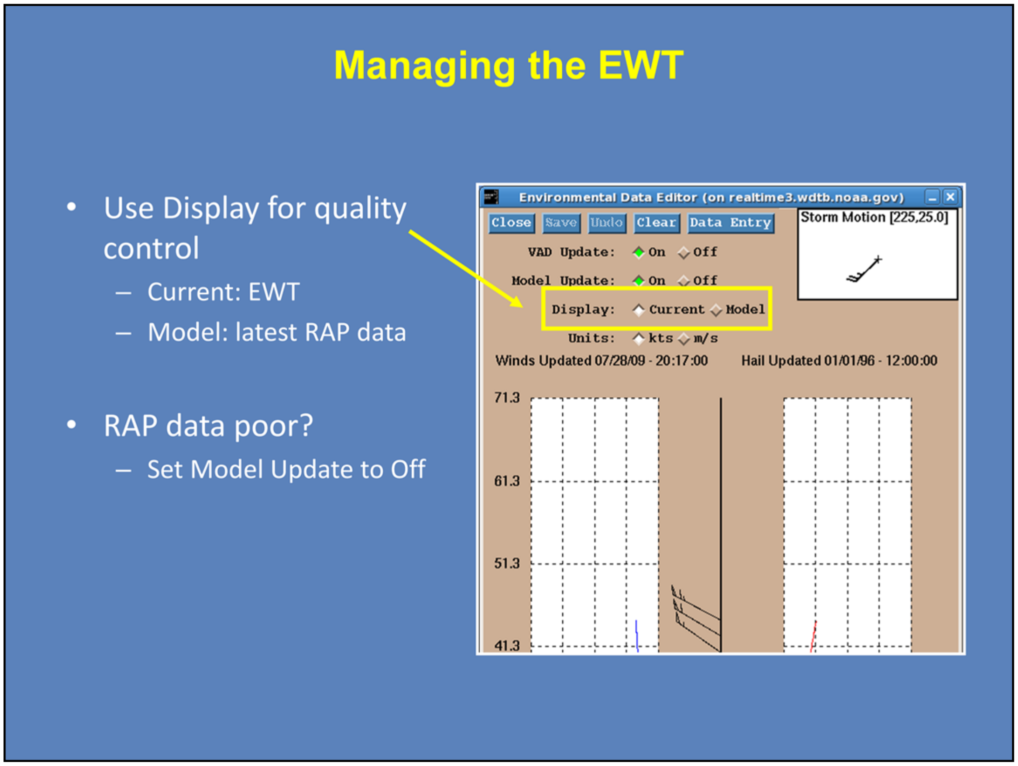

Managing the EWT is a task of monitoring it for relevance, and the Display button can help. When Current is selected, the window is displaying the current state of the EWT. Based on knowledge of winds aloft from other sources, you can verify if the EWT is representative. If there is a need to check the 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.

Back to Top

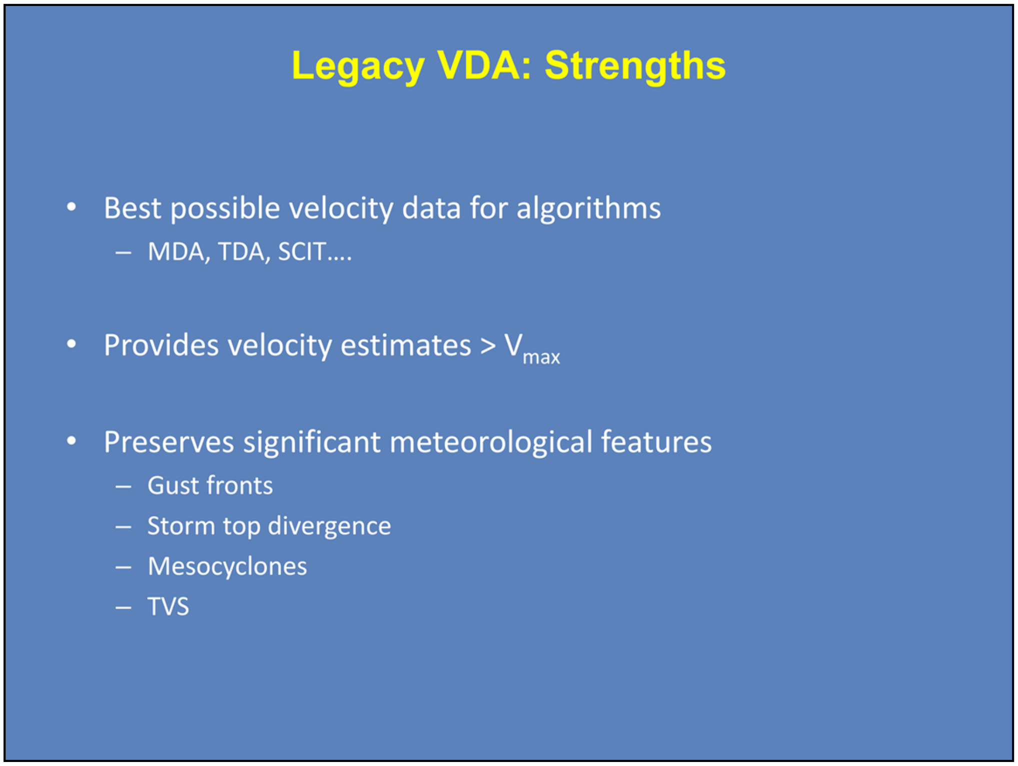

The Velocity Dealiasing algorithm is designed to provide the best possible velocity data, primary to support the RPG algorithms that look for significant features such as circulations. The Velocity Dealiasing algorithm attempts to assign the true radial velocity, even when it exceeds V_{max}. There are 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.

Back to Top

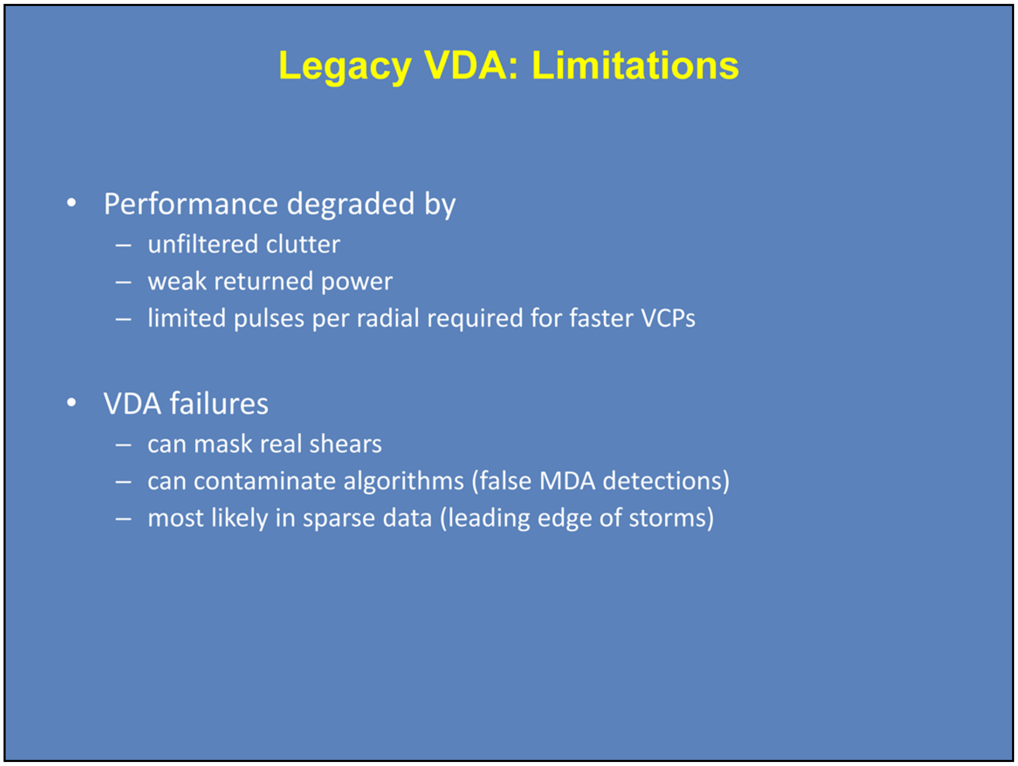

The performance of the Velocity Dealiasing algorithm is degraded by data quality problems that reduce the reliability of the first guess velocity. This includes unfiltered clutter, weak returned power, and in some cases, the low number of pulses per radial with the faster VCPs. When dealiasing failures 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.

Legacy VDA: Considerations

- Suspect improperly dealiased velocities?
  - Change elevation angle
  - Examine previous or later volume scan
  - Anticipate flow based on synoptic conditions
- VCP 31: low PRF (#2) for CD mode
  - Switch to VCP 32

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.

Back to Top

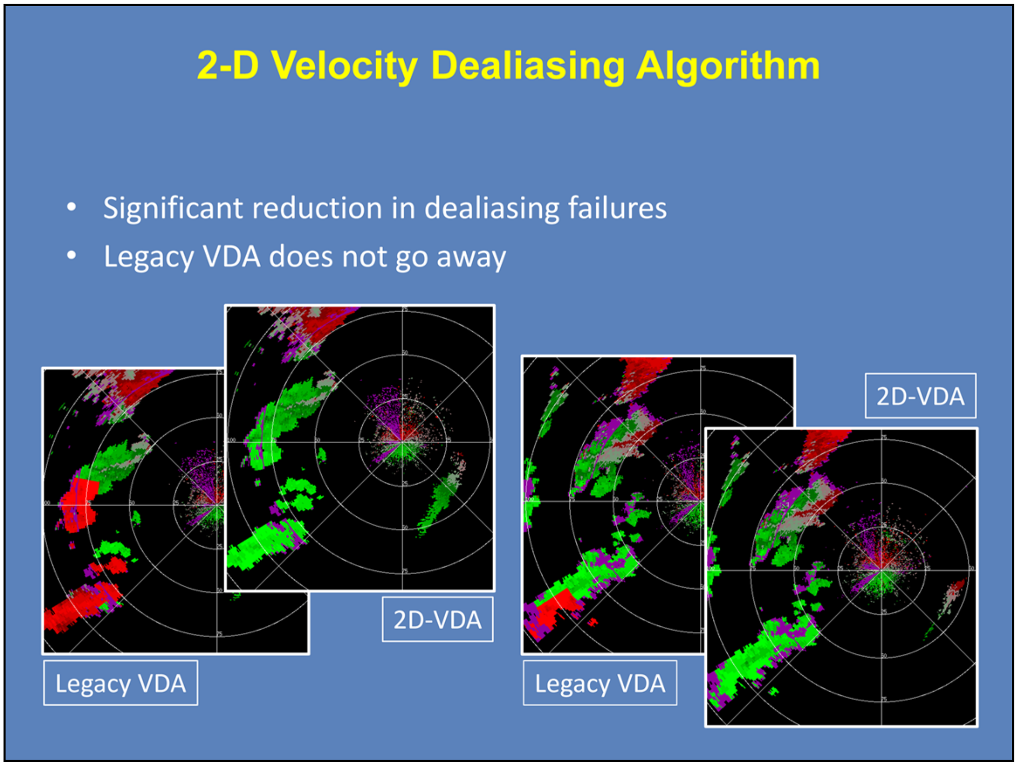

The Two Dimensional Velocity Dealiasing Algorithm (2D-VDA) was deployed in 2013, and offers a significant reduction in dealiasing failures. These examples come from an event with storms that were moving so fast that the legacy Velocity Dealiasing Algorithm had numerous dealiasing failures, while the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

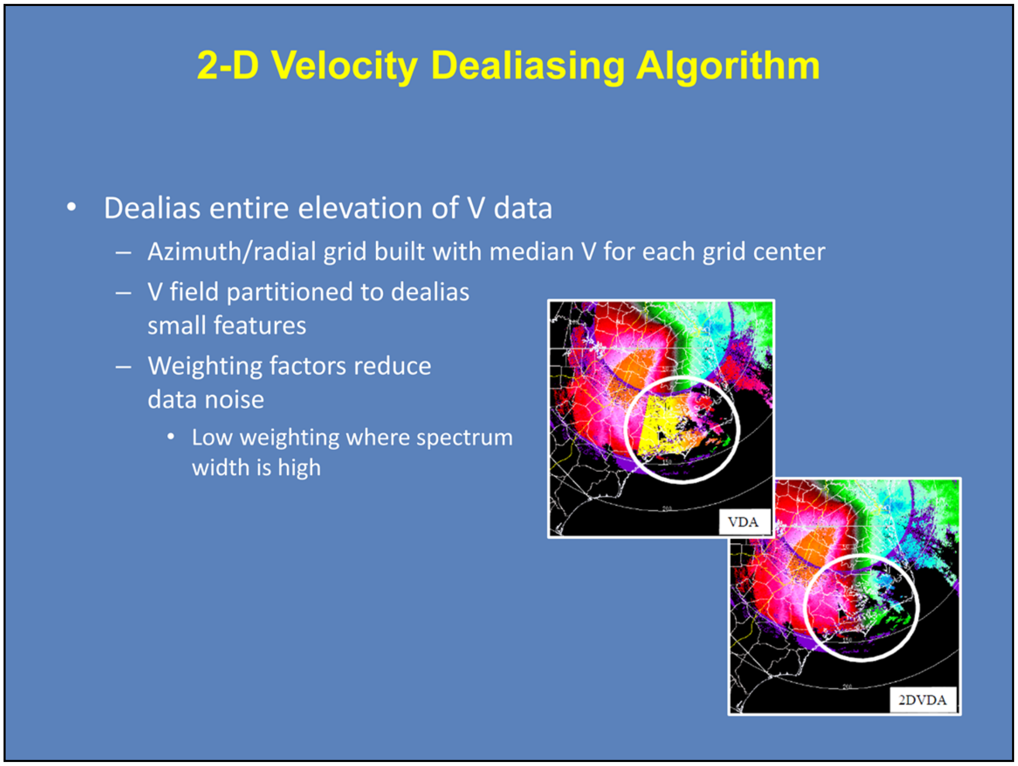

The 2D-VDA is much more robust than the Legacy VDA, and this slide provides a very brief overview. The overall approach is to use a least squares method to minimize errors in the velocity. For each elevation, 2D-VDA first builds a 2 dimensional grid (azimuth and radial) of the velocity data, with a median 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.

Back to Top

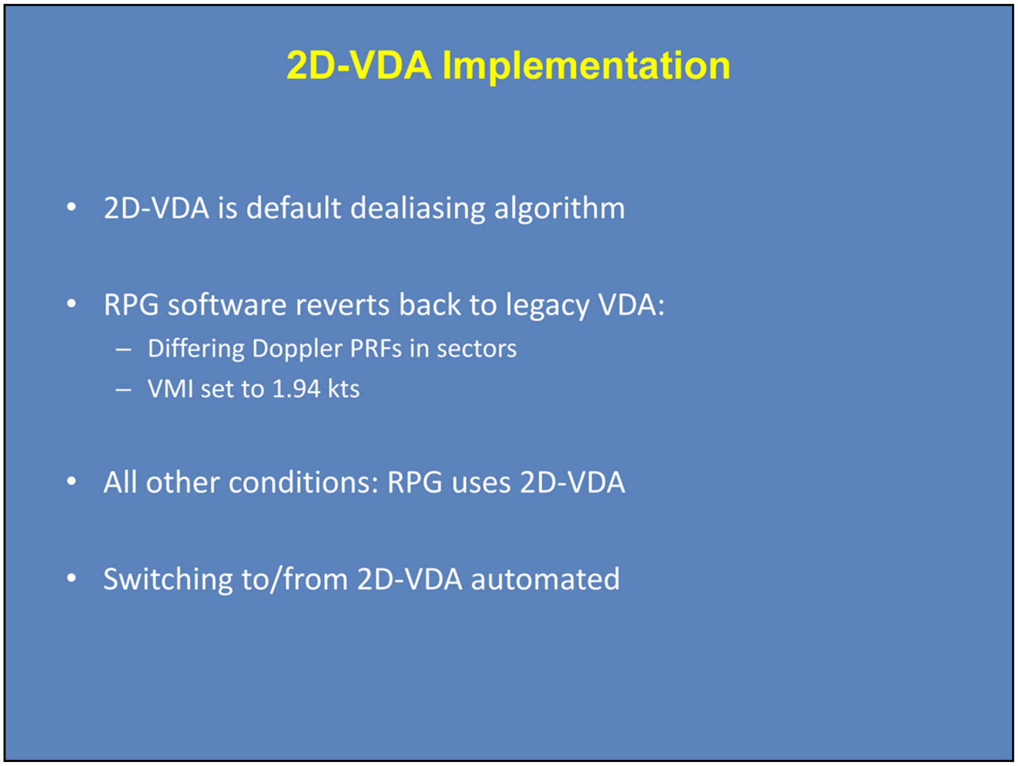

It is important to be aware of how the 2D-VDA has been implemented. Unless it is turned off, the 2D-VDA is the default velocity dealiasing algorithm. There are two conditions where the RPG software will automatically revert back to the Legacy VDA. The first is when a VCP with different Doppler PRFs in 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.

Back to Top

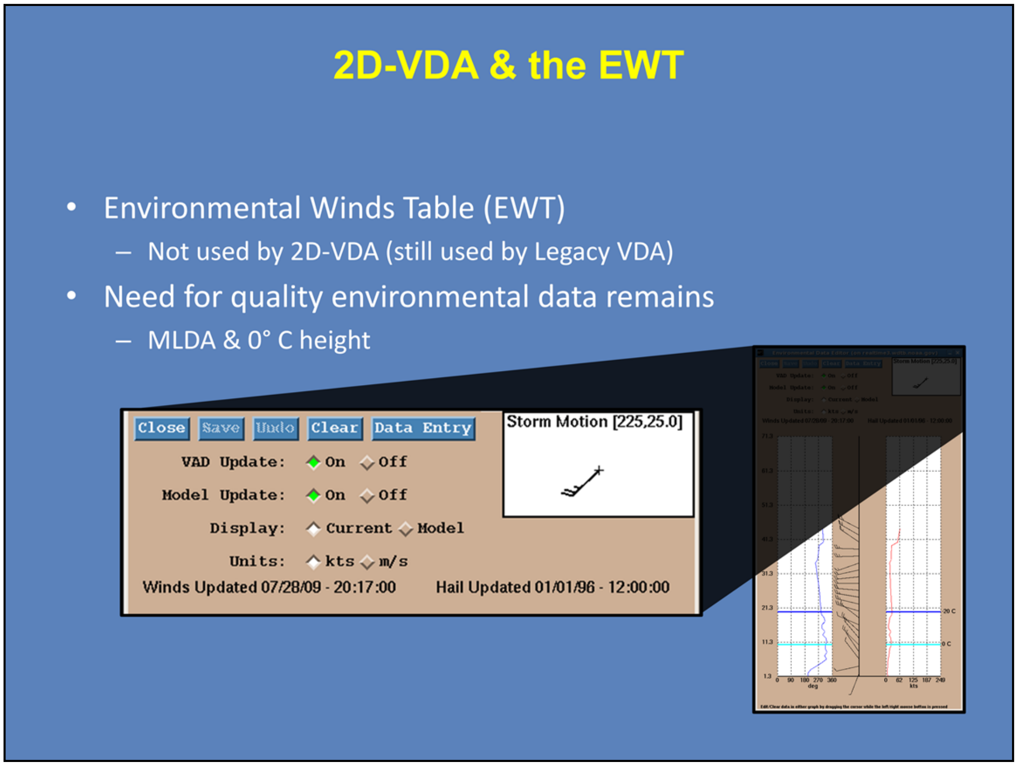

The Environmental Winds Table (EWT) supports the performance of the Legacy VDA, and will continue to do so whenever the Legacy VDA is running. The 2D-VDA does not rely on the EWT.

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.

Back to Top

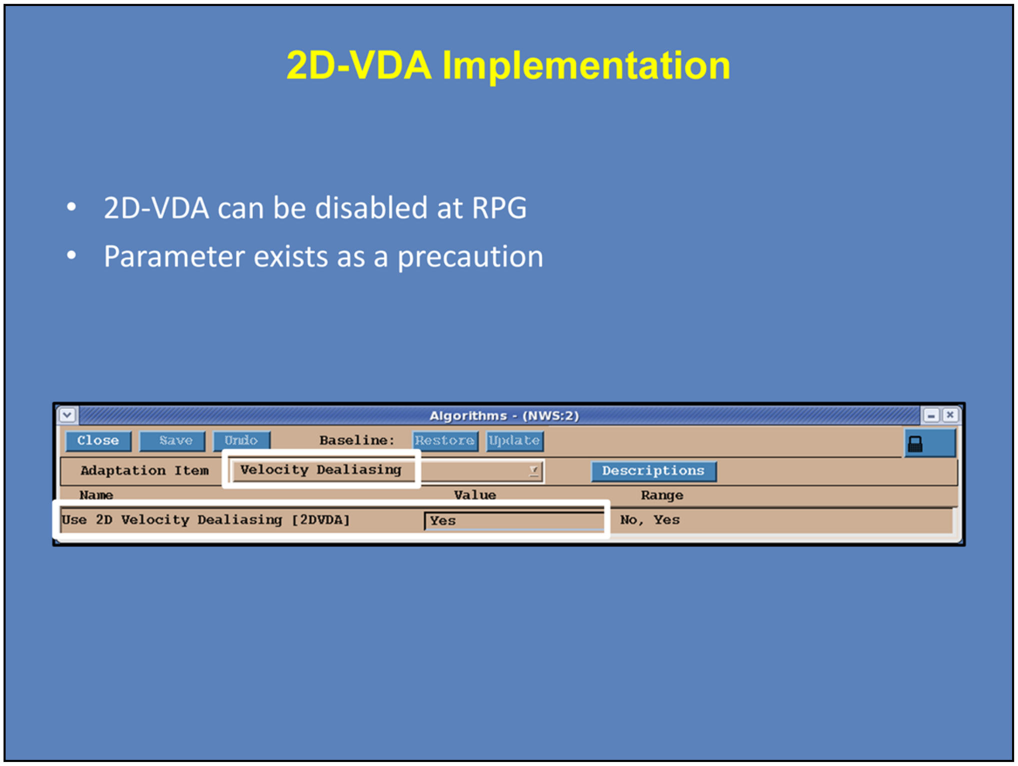

There is a new entry at the Algorithms window at the RPG, called "Velocity Dealiasing". The parameter, "Use 2D Velocity Dealiasing", controls whether the 2D-VDA is used, and the default setting is Yes. Setting this parameter to No is not expected to be needed, but is available as a precaution.

Back to Top

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

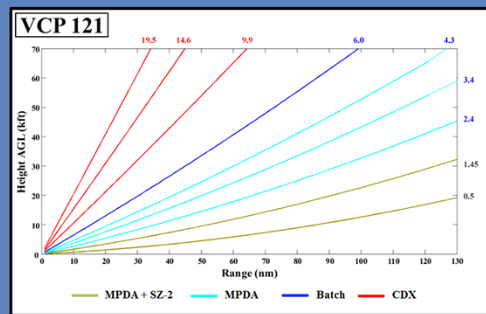

The Multiple PRF Dealiasing Algorithm (MPDA) is a special application uniquely used for VCP 121. It is designed to mitigate both range folding and improperly dealiased velocities, with range folding mitigation the most apparent. VCP 121 samples the same elevations as VCP 21, and thus has the same limitations. 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

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							51	55	59
3.35	21.599							35	38	40
3.35	27.400							40	43	46
3.35	21.401							51	55	59
4.3	16.304							52	56	61
4.3	29.498							37	40	43
6.0	20.204							43	47	51
9.9	29.498							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

3 CD antenna rotations:
- SZ-2
- Legacy Range Unfolding
- extra CD rotations

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, but ***fastest antenna rotations***
  - 20 rotations in 5 mins 45 secs

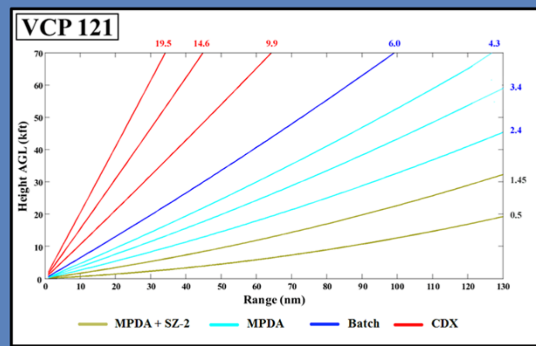

- VCP 121 **NOT** appropriate for fast moving or rapidly evolving storms
  - Use VCP 12 or 212
  - Better vertical sampling & faster updates

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.

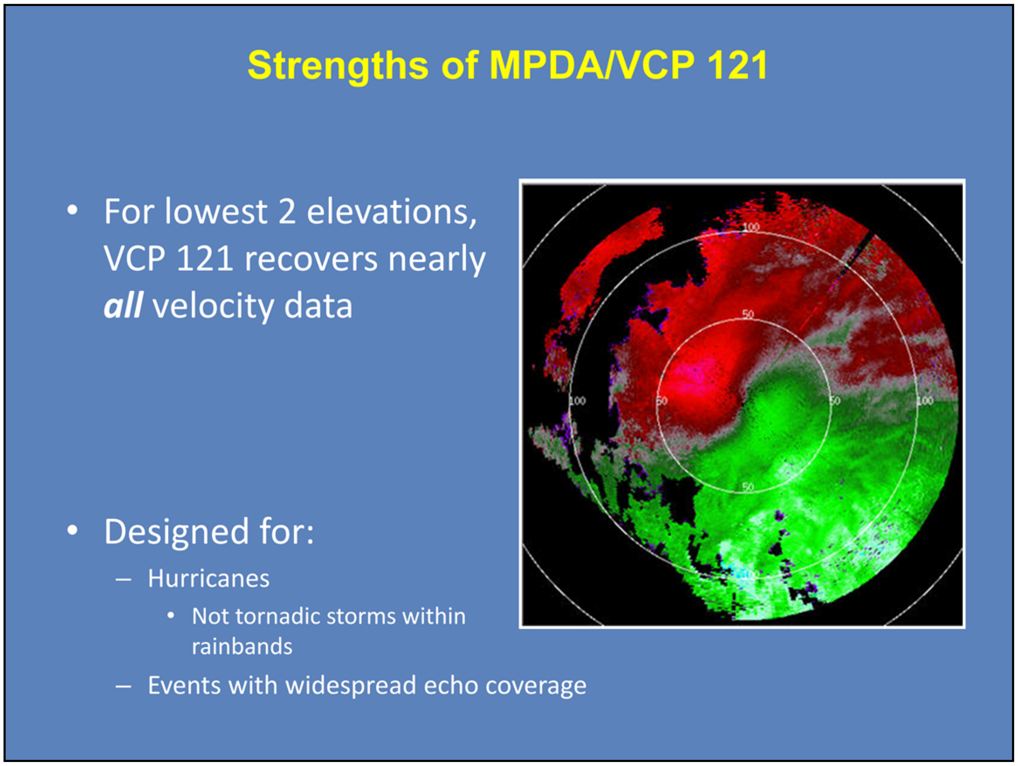

MPDA, aka VCP 121 is capable of recovering nearly all velocity data for the lowest two elevations. It is designed for sampling hurricanes while still offshore when large scale velocity structure is the priority. Once operations shift to looking for potentially tornadic storms within the rainbands, VCP 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.

Back to Top

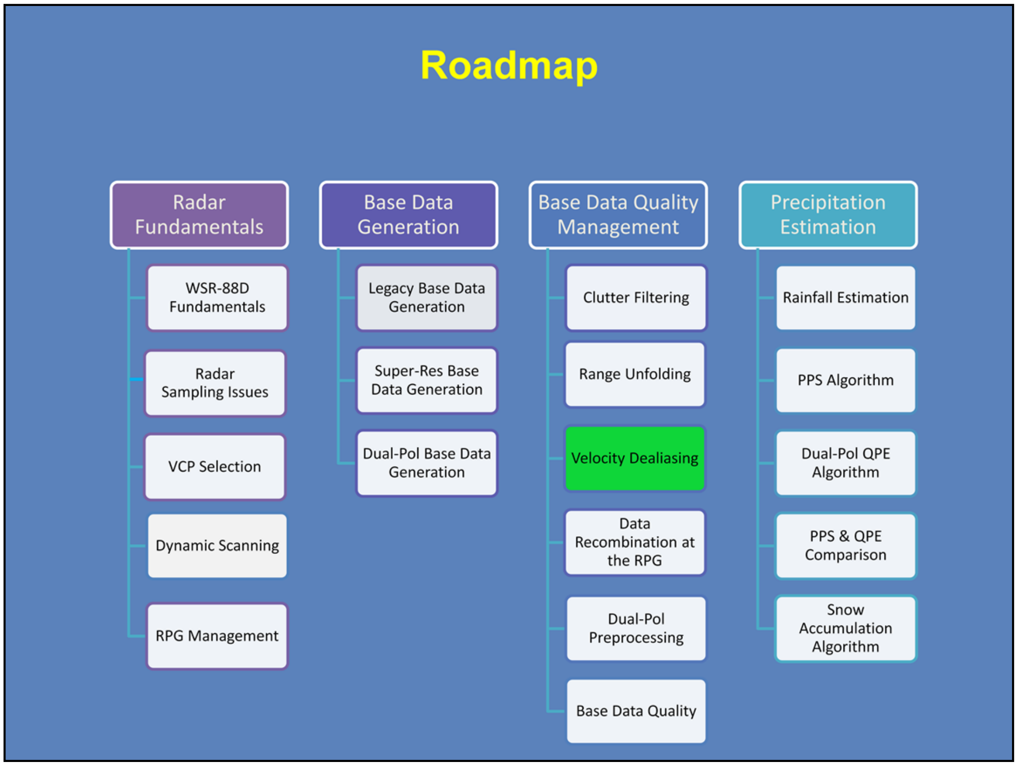

Here is the “roadmap” with your current location.

Welcome to the lesson on Data Recombination at the RPG.

Course Completion Info

Tabs - 4 Tabs (Including Introduction)

Last Modified: Jul 25, 2017 at 02:42 PM

PROPERTIES

Show interaction in menu as: Single item

Allow user to leave interaction: At any time

Prev/Next player buttons go to: Step in interaction

Edit in Engage

Edit Properties

Back to Top

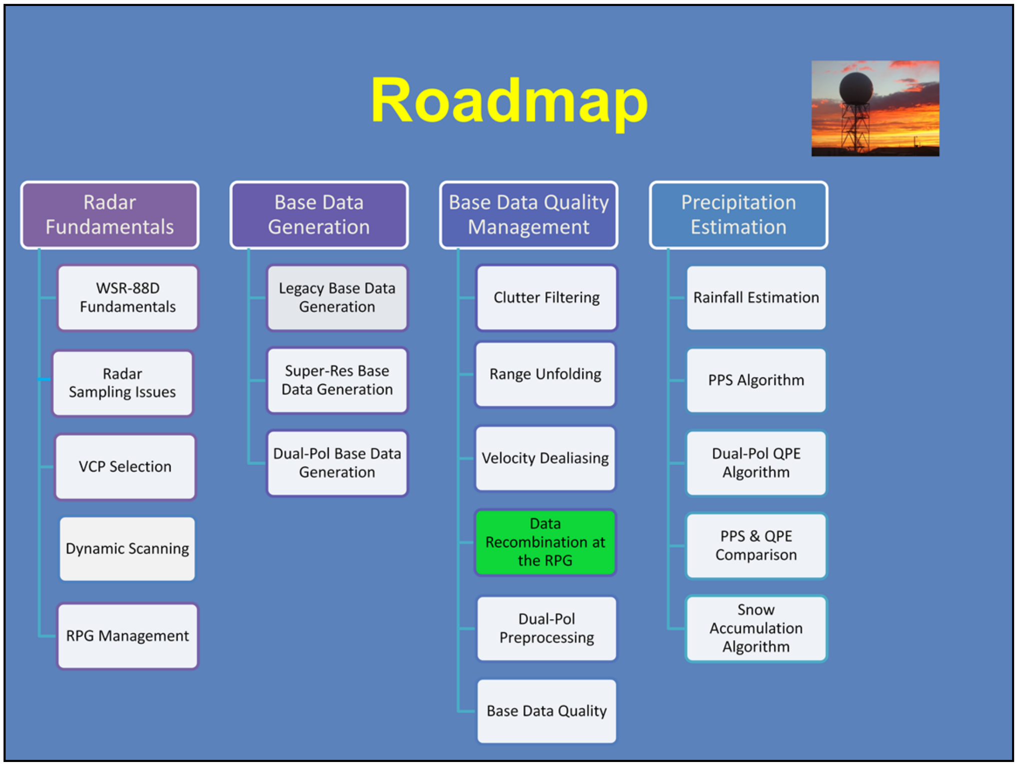

Here is the “roadmap” with your current location.

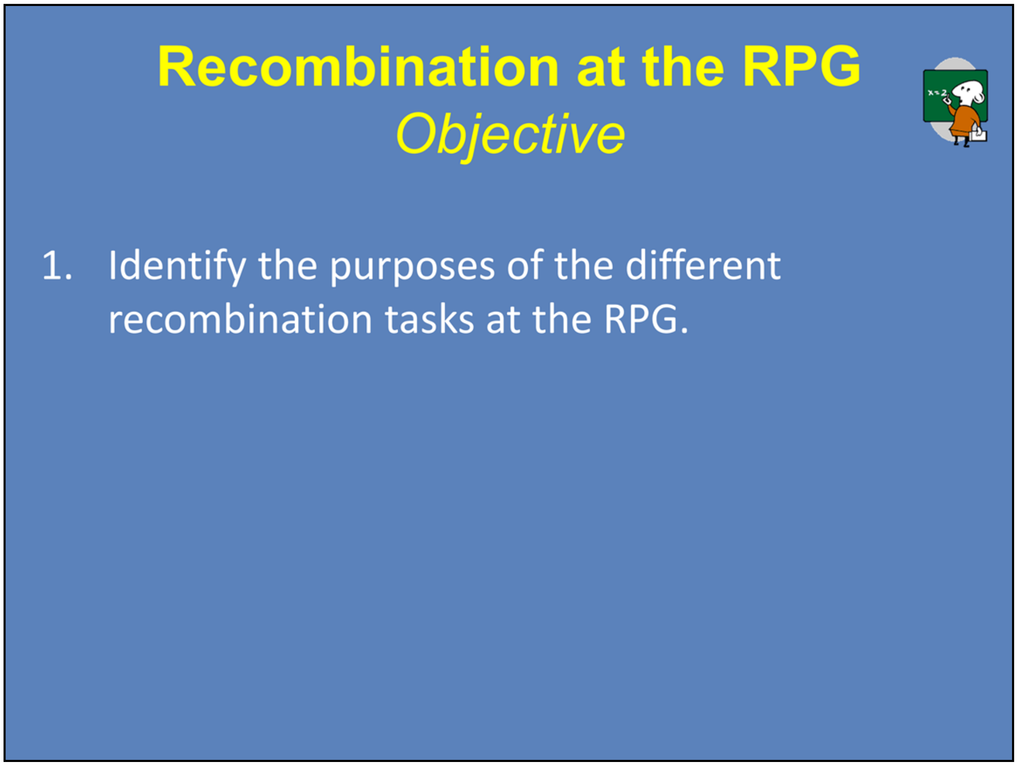

There is one objective in Recombination at the RPG.

Back to Top

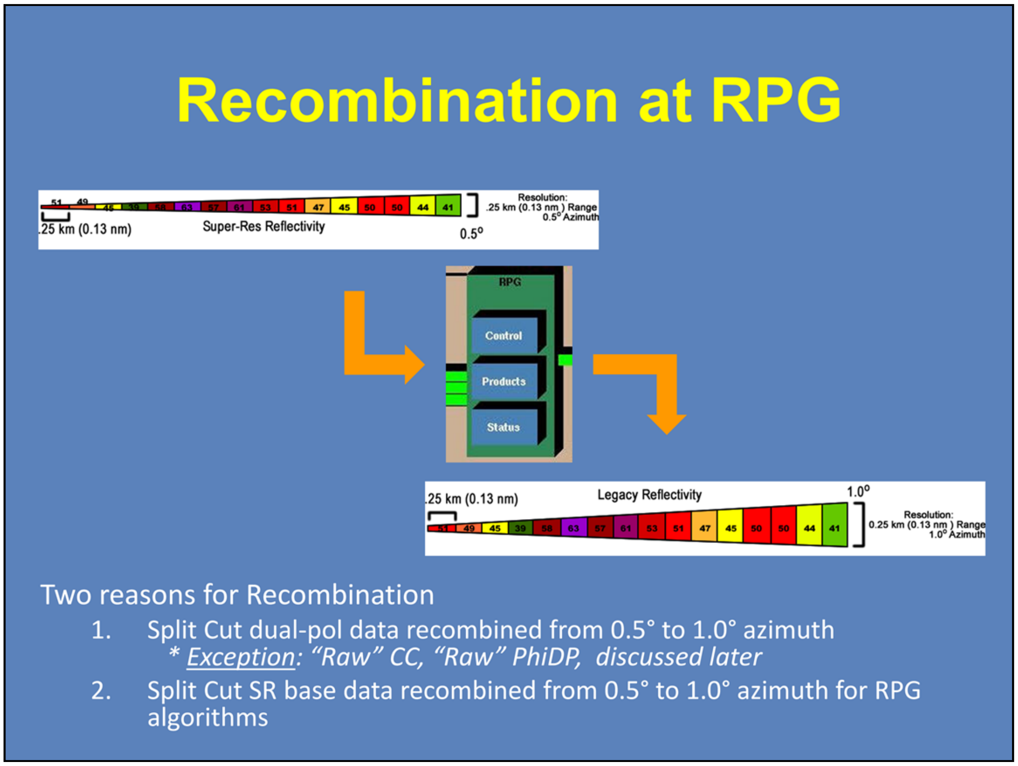

Now for recombination at the RPG, which is a process that is needed for two reasons, to generate products themselves or for input into algorithms.

1. For the Split Cut elevations, the dual-pol base data are processed at the RDA with an azimuthal resolution of 0.5°, and that is simply too noisy. The dual pol base data are recombined to a 1.0° azimuth before the products are built.

2. Also for the Split Cut elevations, the super res base data (reflectivity, velocity and spectrum width) are recombined from 0.5° to 1.0° to support some of the RPG algorithms which cannot ingest a resolution that is that high.

Back to Top

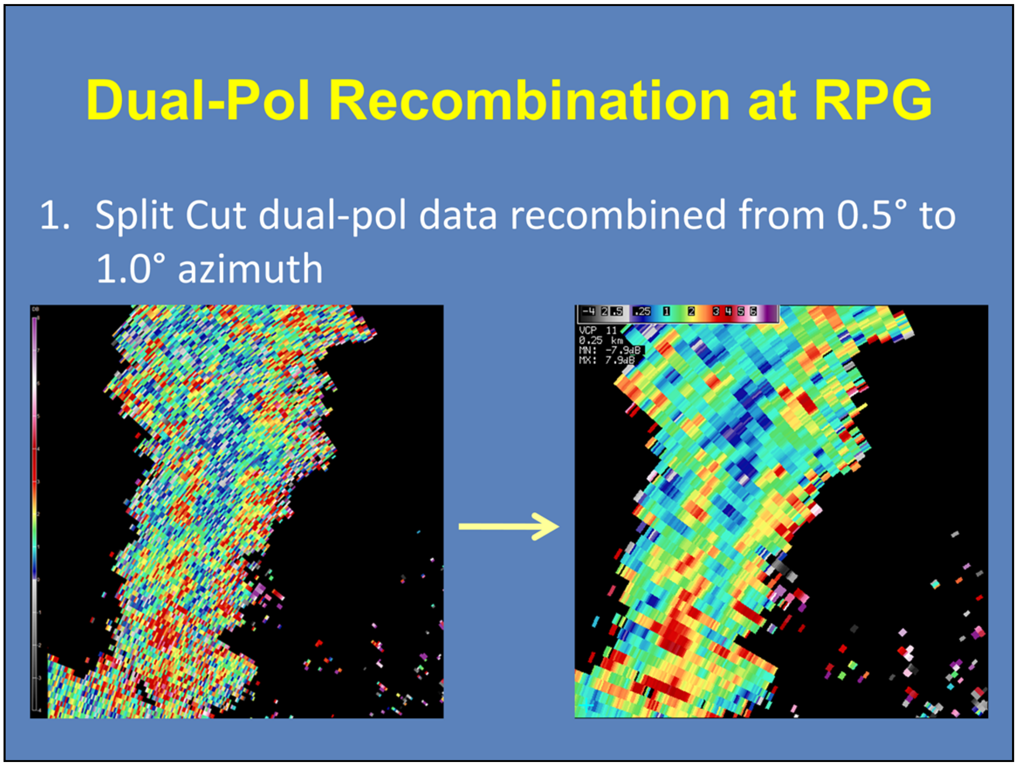

The dual-pol base data arriving from the RDA are ZDR, CC, and Φ_{DP}. For the Split Cuts, all of these data have an azimuthal resolution of 0.5°, and are too visually noisy for direct product generation. These data are recombined to an azimuth of 1.0°.

The dual-pol base data are also “preprocessed”, and this involves smoothing as well as 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.

Back to Top

Recombination at RPG

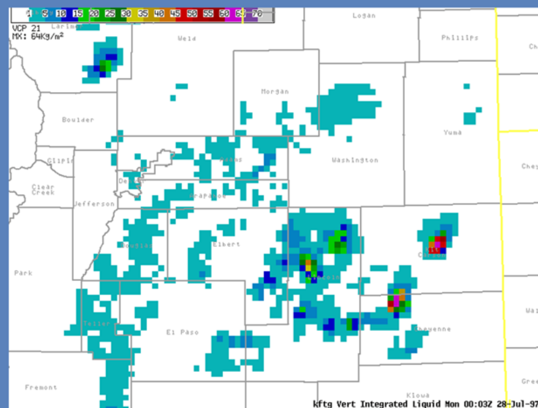

2. Many RPG algorithms (VIL, PPS, TDA,…) cannot ingest SR base data
   – MDA is an exception

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, the legacy 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

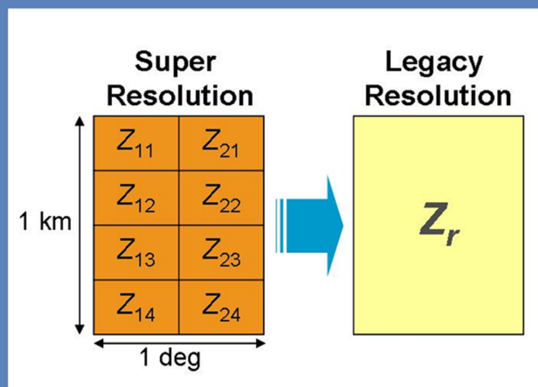

What you see here is to support the RPG algorithms that require 1.0° azimuth by 1 km range resolution for reflectivity input.

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.

Back to Top

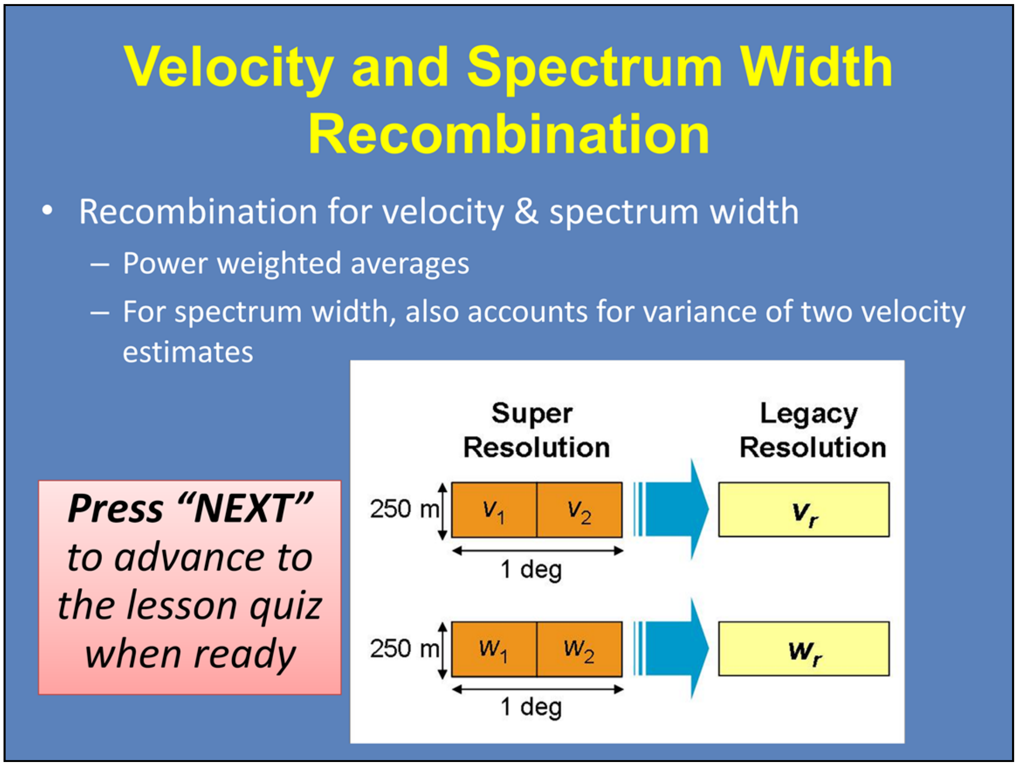

This step is for RPG algorithms that require 1.0° azimuth by .25 km range resolution for velocity and spectrum width input.

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 range-folded, then the legacy bin is also assigned No Data or range-folded (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.

Back to Top

Data Recombination - Final Quiz

Quiz - 2 questions

Last Modified: Oct 09, 2015 at 11:57 AM

PROPERTIES

On passing, 'Finish' button:	Goes to Next Slide
On failing, 'Finish' button:	Goes to Next Slide
Allow user to leave quiz:	After user has completed quiz
User may view slides after quiz:	At any time
Show in menu as:	Single item

Edit in Quizmaker | Edit Properties

Back to Top

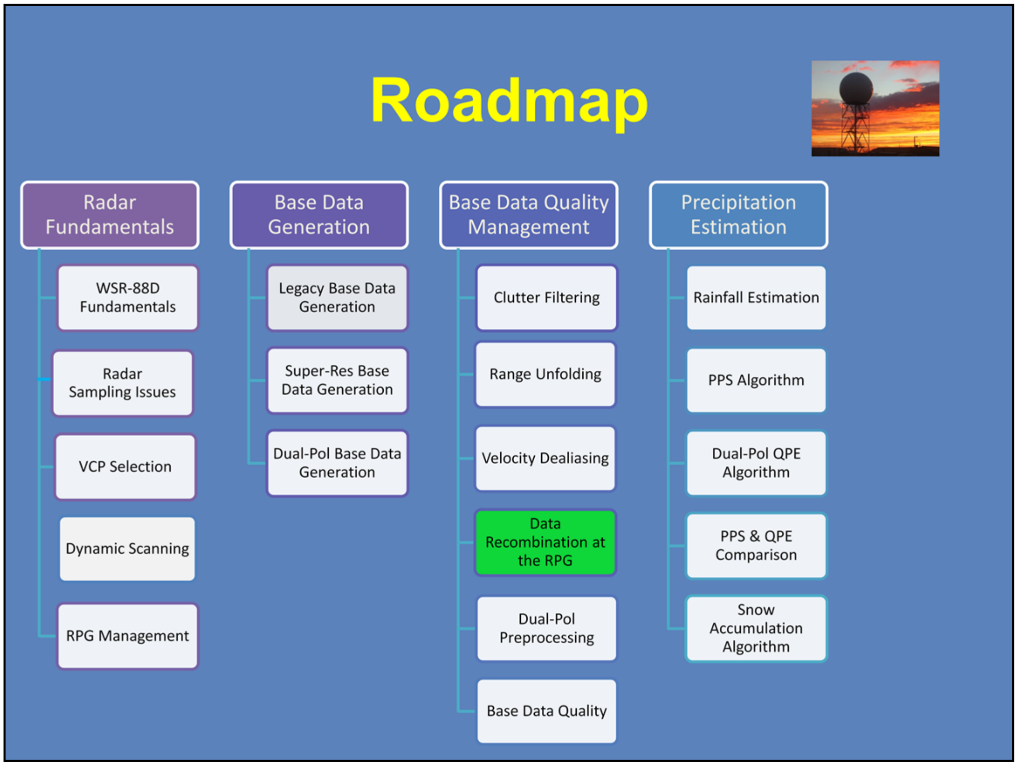

This concludes the lesson, and here is the “roadmap” with your current location.

Back to Top

Welcome to Dual-Pol Base Data Preprocessing at the RPG

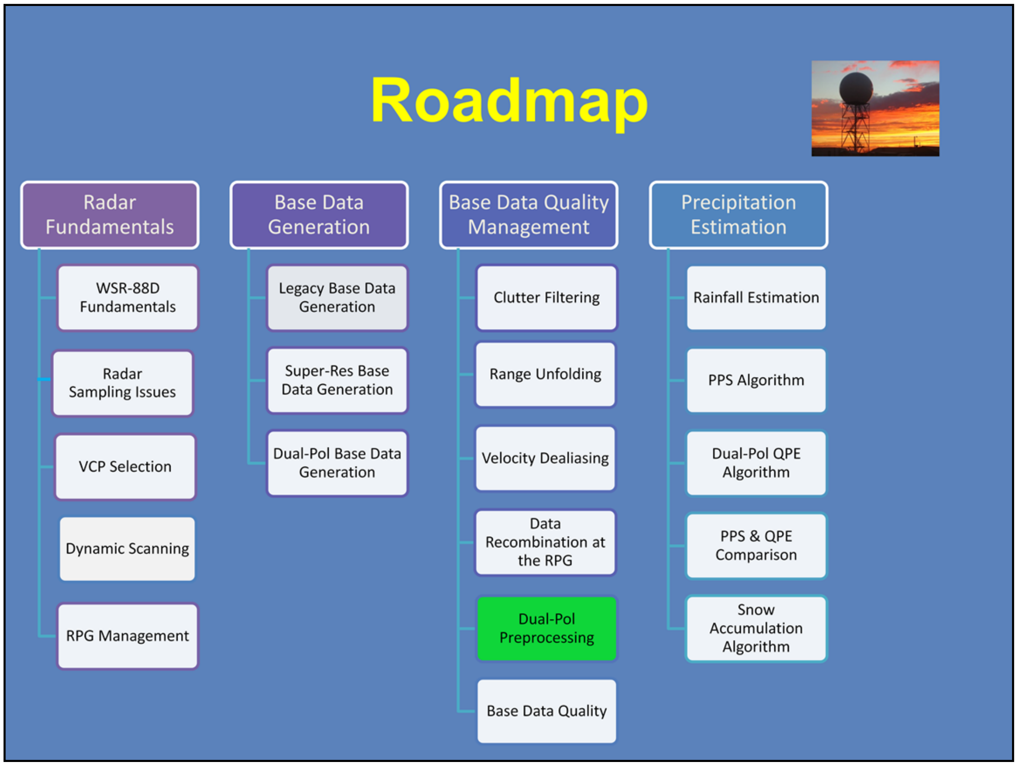

Here is the “roadmap” with your current location.

Back to Top

Dual Pol Preprocessing at the RPG

Objectives

1. Identify the primary tasks of the Dual-Pol Preprocessor at the RPG.

There is one objective for Dual Pol Preprocessing.

Back to Top

Dual-Pol Preprocessor at the RPG

- Goal: prepare ZDR, CC & Φ_{DP} for
  - Dual-Pol base product generation
  - HCA, MLDA, and QPE input
- Tasks
  - Smooth Z*, ZDR, CC & Φ_{DP}
  - Compute KDP

*Smoothed Z used ***only*** for input to Dual-Pol RPG algorithms!

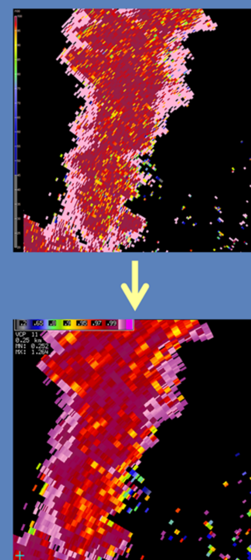

The Dual-Pol Preprocessor is an RPG algorithm. It's purpose is to prepare the Dual-Pol base data for two things: base product generation for the Dual-Pol products, as well as input into the RPG Dual-Pol algorithms, i.e. Hydrometeor Classification Algorithm (HCA), the Melting Layer Detection Algorithm (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 generally too noisy for human interpretation and certainly for 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 the legacy base products that you are familiar with.

The remaining task for the Preprocessor is to compute the Specific Differential Phase (KDP) values, before we generate the KDP product.

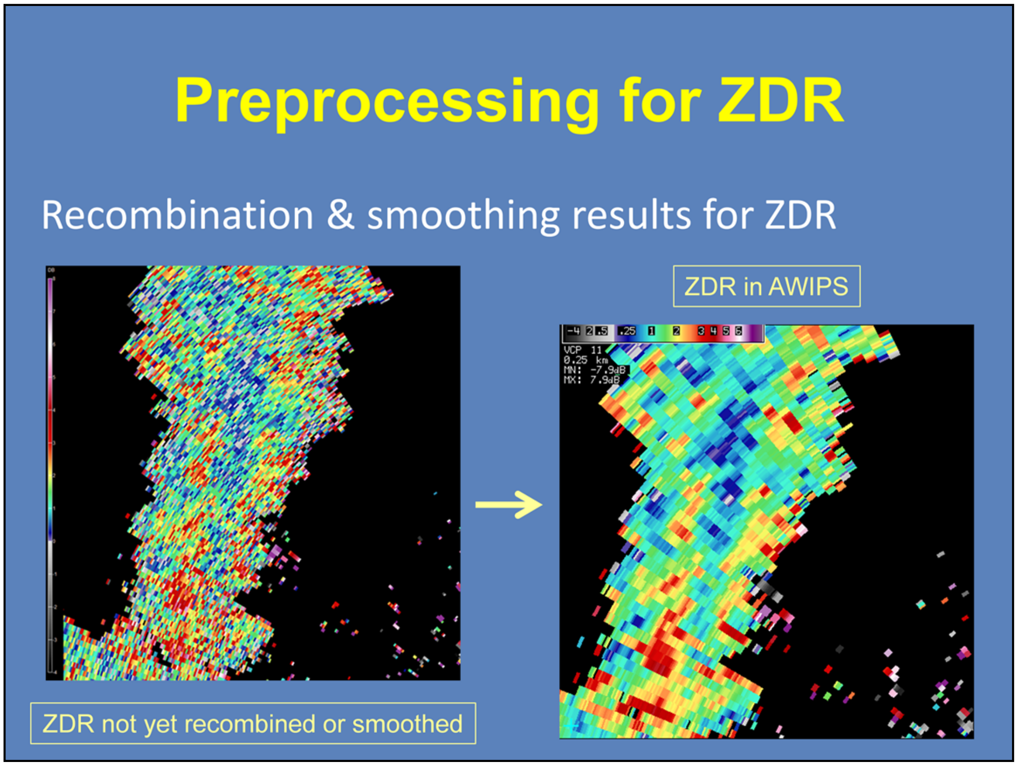

The example on the left is raw Differential Reflectivity, ZDR, from the RDA, at 0.5° azimuthal resolution and .25 km range resolution. It has not yet been recombined or smoothed. It is pretty noisy for even human interpretation. The image on the right is the same data displayed in AWIPS after recombination 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.

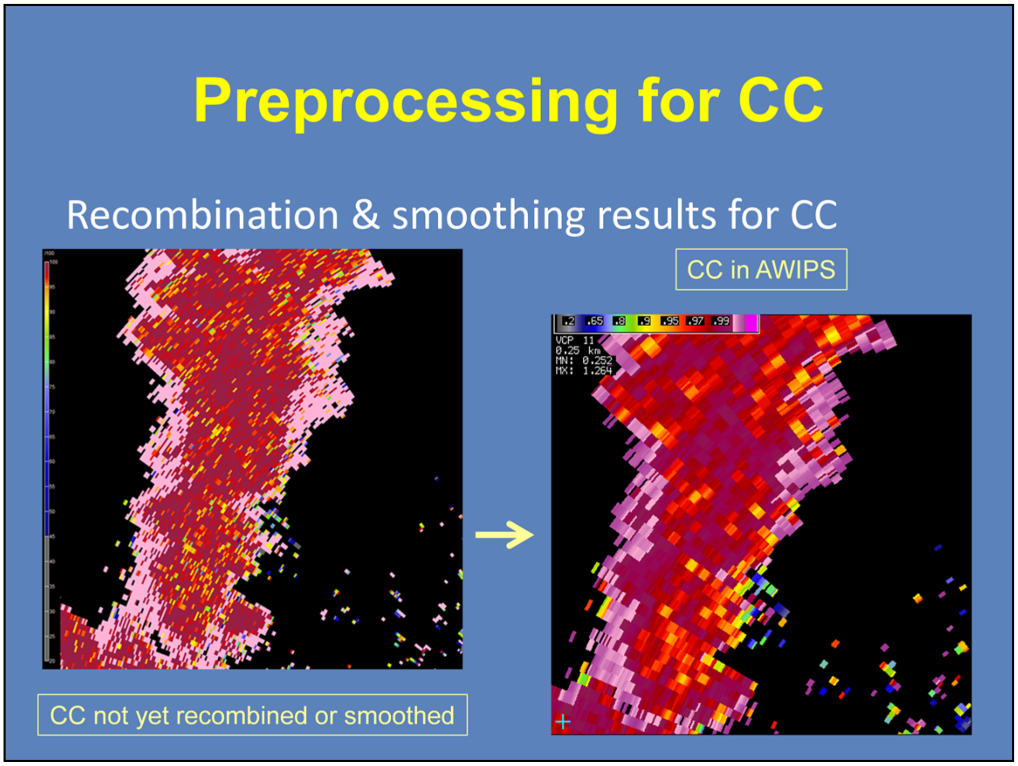

Here's a similar comparison for Correlation Coefficient, CC. We have raw CC on the left and the recombined and smoothed CC on the right. As with ZDR, the same smoothing technique is applied. It's a linear average to a segment (of varying length) along the radial. That average value is then assigned to the original range bin at the center of the segment.

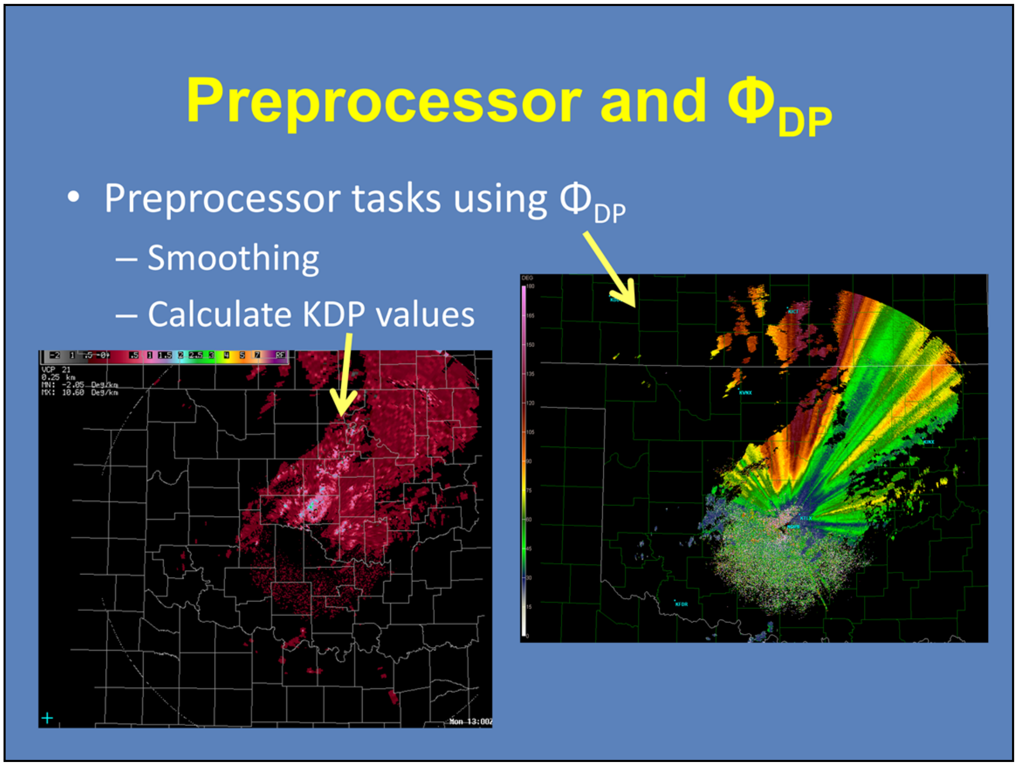

As with ZDR and CC, the Differential Phase, Φ_{DP}, base data are first recombined, then smoothed. On the right is an example of Φ_{DP} base data, not yet recombined or smoothed. This image is from GR Analyst, showing the raw Level II data.

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.

Back to Top

PhiDP: The Good, the Bad, and the Ugly

PhiDP: The Good, the Bad, & the Ugly

Differential Phase (Φ_{DP}): The Good

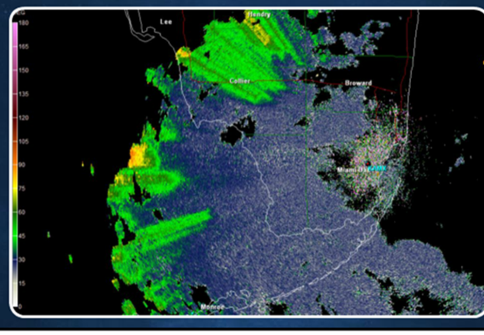

Differential Phase (Φ_{DP}) is a dual-pol, base data product that is not available for viewing in AWIPS. As a result, Φ_{DP} data interpretation can be both difficult and undervalued. Even if you never look at a Φ_{DP} product, it's important to understand what Φ_{DP} is and how it impacts other base data products. Use the buttons below to navigate through the information provided on "the good, bad, and ugly" of Φ_{DP}.

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

(Click to be linked to sub-lesson on PhiDP)

Back to Top

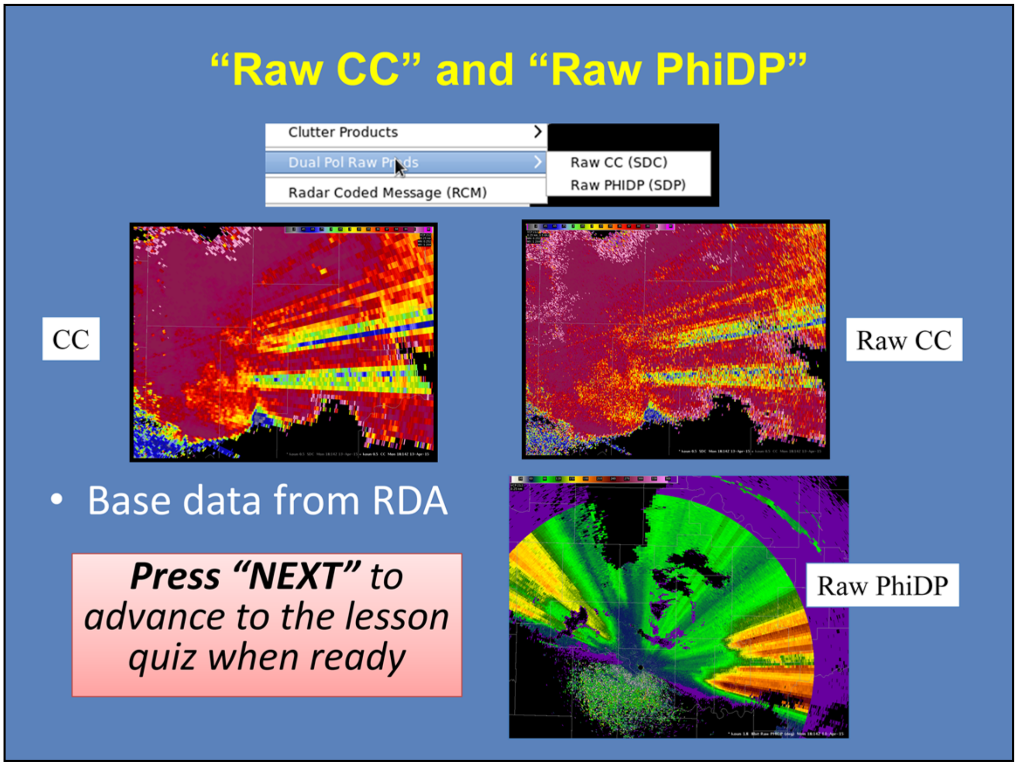

There are two dual pol products in AWIPS that *not* Preprocessed. They are both titled "Raw" to indicate that you are seeing *only* the base data sent from the RDA to the RPG.

The Raw CC has higher azimuthal resolution (0.5°), however it is not a substitute for the regular CC product. 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). That's why this product was made available.

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

Back to Top

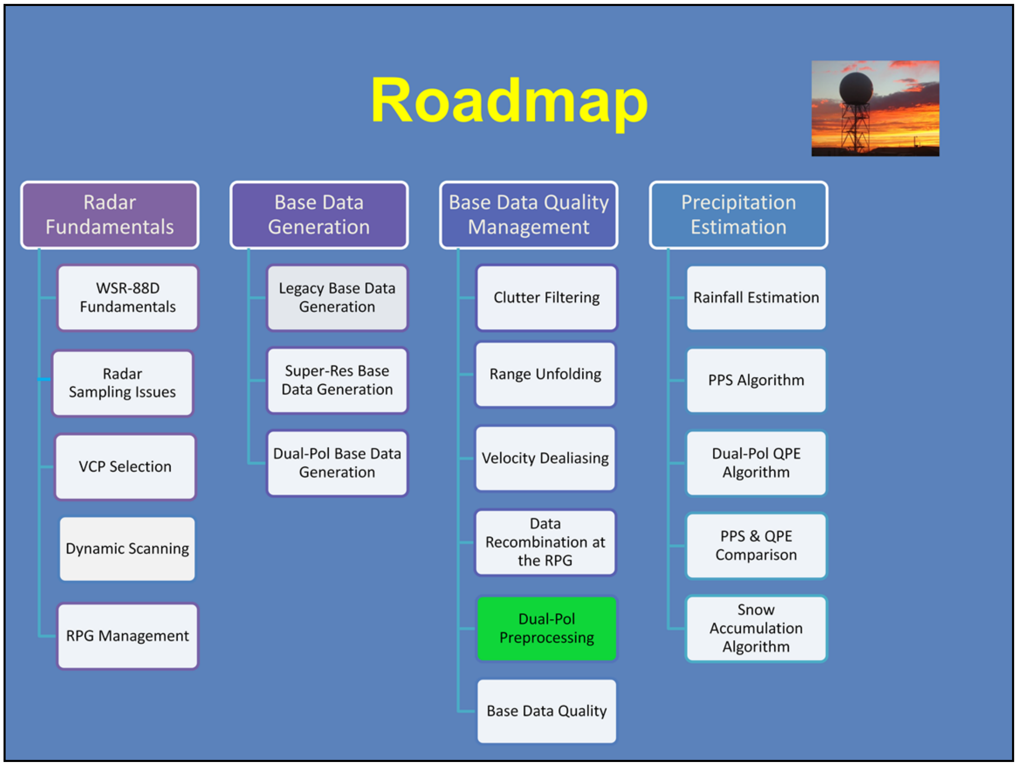

This concludes the lesson, and here is the "roadmap" with your current location.

Welcome to WSR-88D Base Data Quality

Back to Top

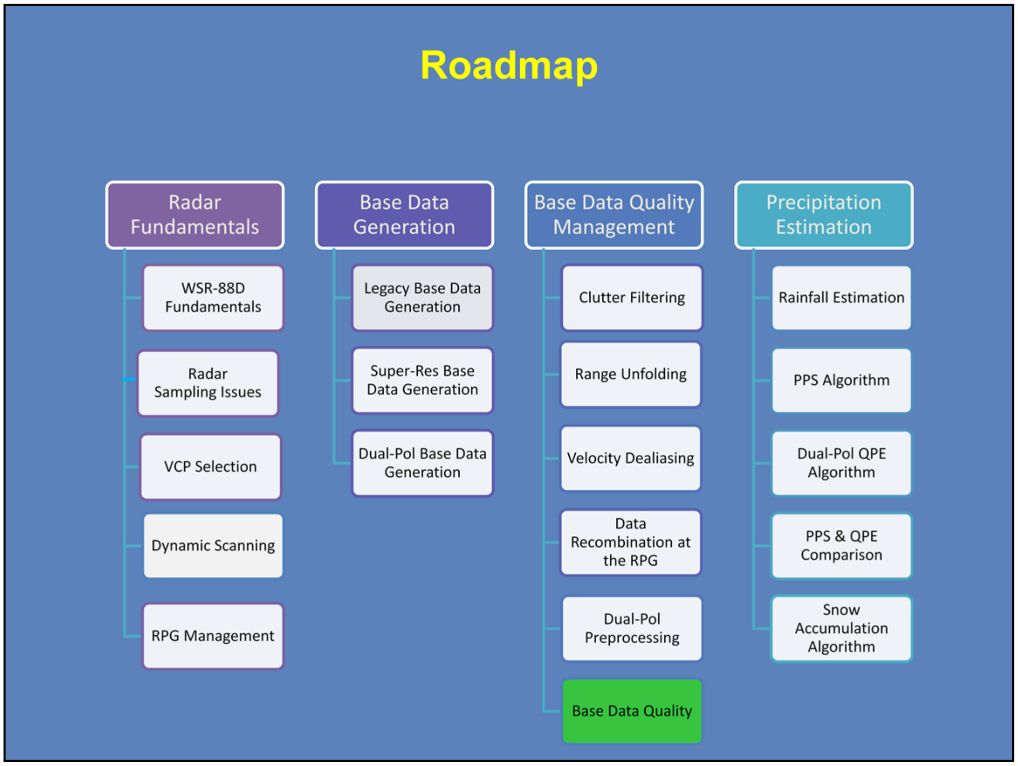

Here is the “roadmap” with your current location.

Back to Top

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.

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.

Back to Top

CMD False Detections

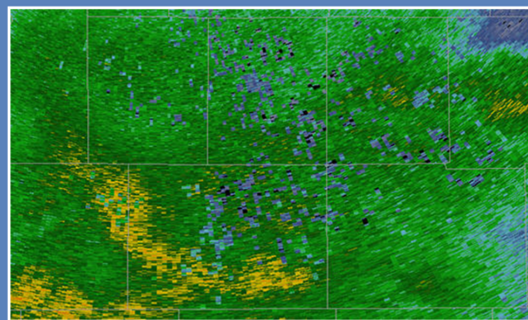

- CMD uses fuzzy logic and multiple inputs to identify bins with clutter
- False detections: seemingly random data loss
  - Stratiform rain
  - Faster VCPs with fewer pulses per radial

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.

Back to Top

CMD False Detections

Also with convection:

– Trade off: fast product updates vs. best clutter identification and removal

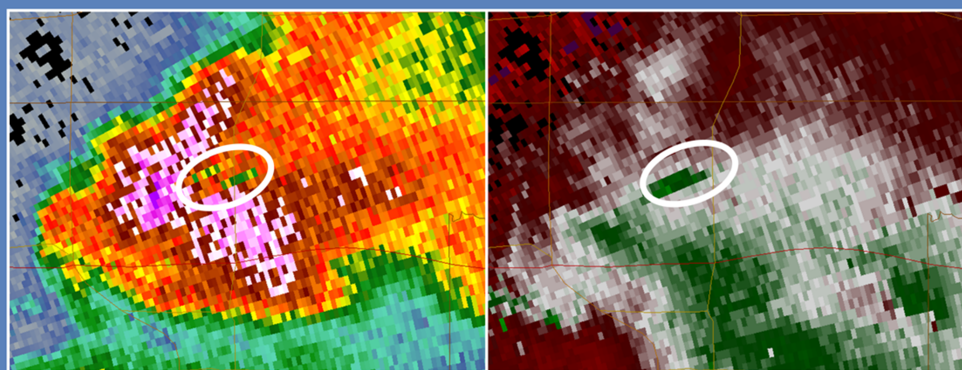

CMD false detections can even occur in or near convective storms. The result is seemingly random blocks of data loss that do not persist in space or time. The trade off at work here is the need for fast product updates with convective events, which means fewer pulses per radial, vs. the best performance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

Back to Top

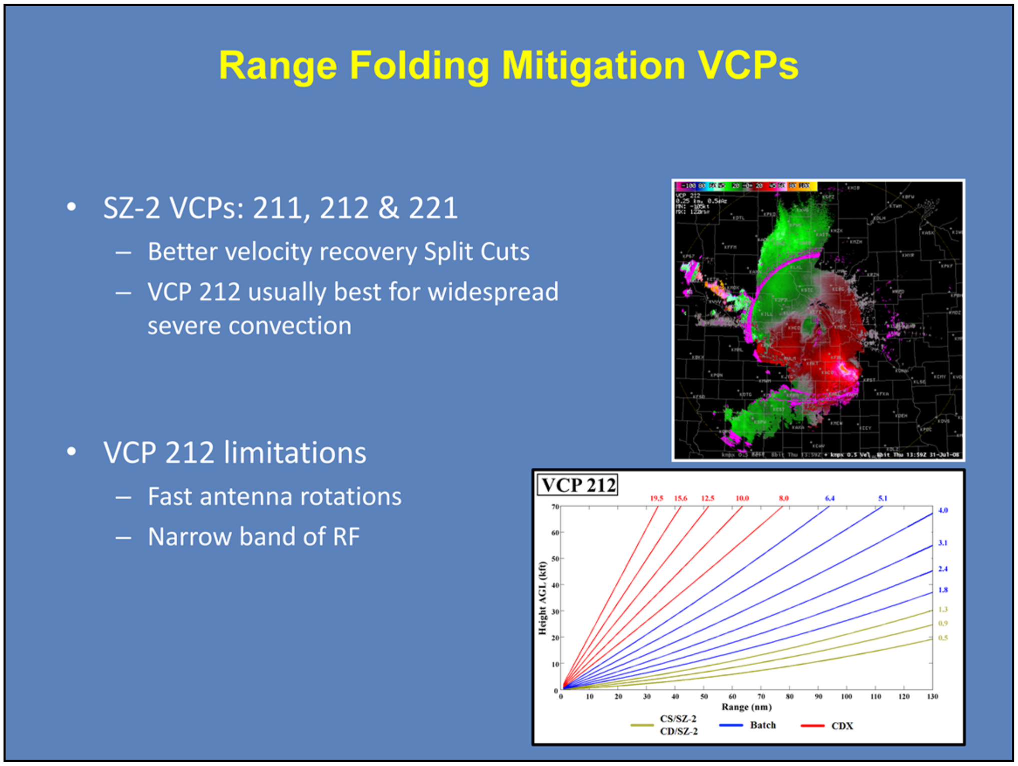

There is a class of VCPs that are designed to mitigate or minimize range folding in the velocity data. The first three in this group, VCPs 211, 212, and 221, share the fact that SZ-2 Range Unfolding is applied on the Split Cut elevations. The advantage of SZ-2 is much greater availability of velocity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

Back to Top

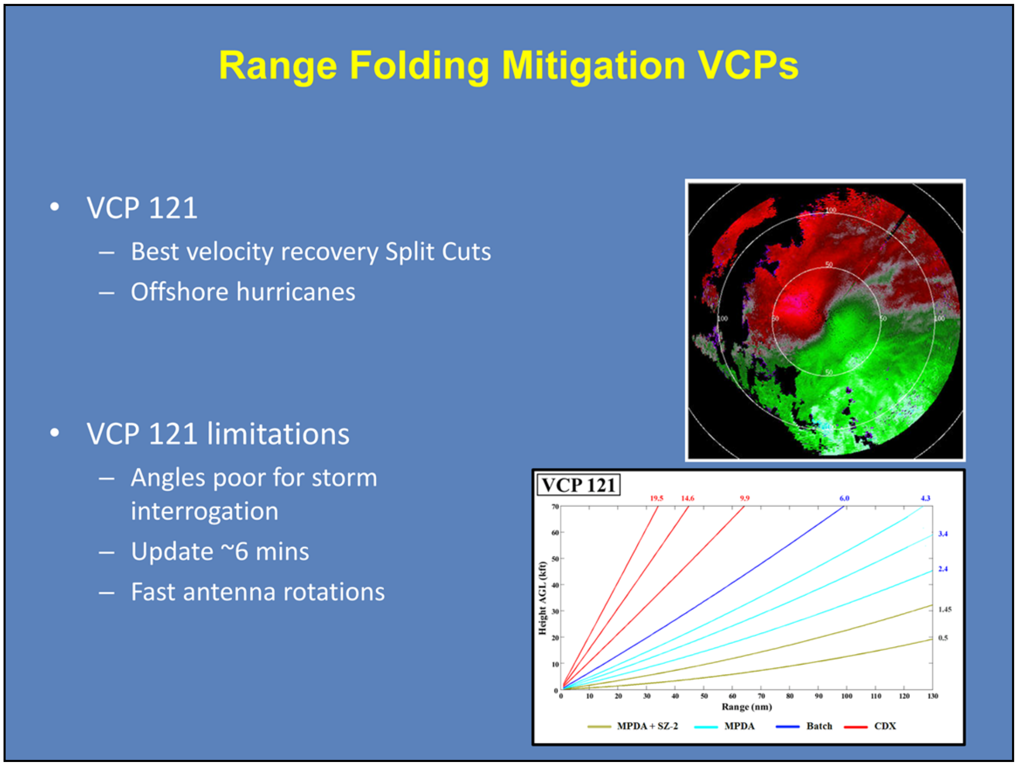

The other VCP in the group of range folding mitigation VCPs is 121. The primary benefit of VCP 121 is that for the lowest two elevations, nearly all the velocity data are recovered. VCP 121 is usually the best choice for offshore hurricanes or widespread non-severe convection.

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

Even with 10 cm, Z attenuation happens!

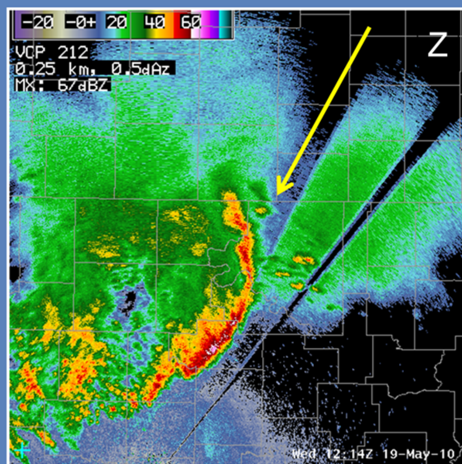

KCRI: Single-Pol

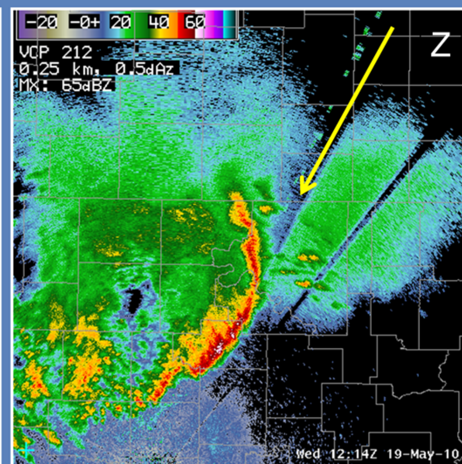

KOUN: Dual-Pol

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.

Differential Attenuation of ZDR

Differential attenuation happens, too (in ZDR)!

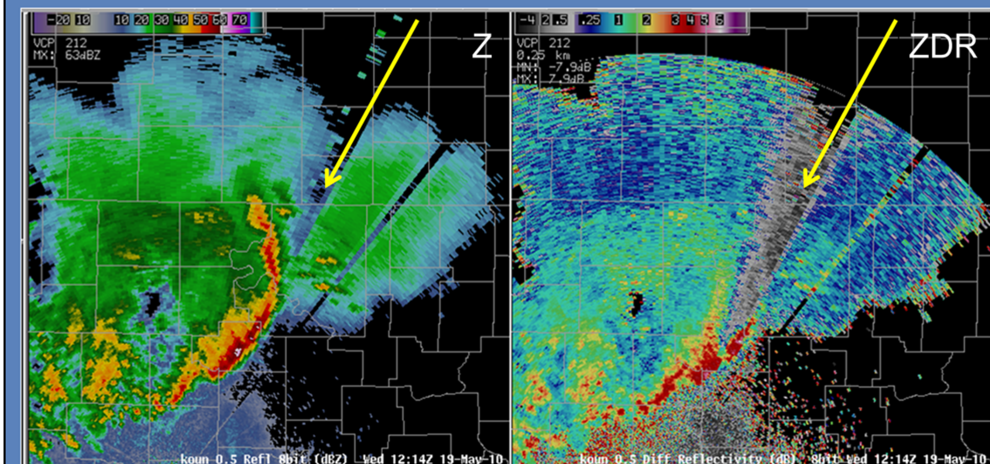

With this same squall line case, there are very low ZDR values down radial (right image) that visually correlate with the Z attenuation. With ZDR, it is possible to have “differential attenuation”.

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.

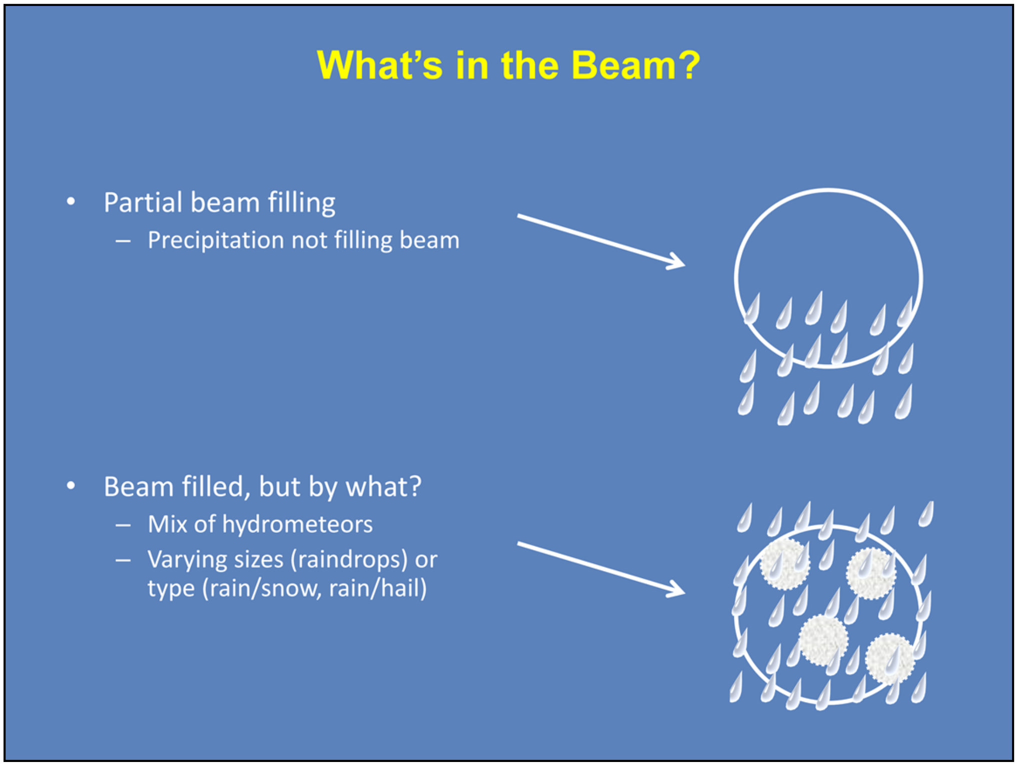

Non-uniform beam filling often occurs, and has always had implications with weather radar data quality, especially as range increases. However, the impacts have not been as apparent as it can sometimes be with Dual-Pol data. Even with Dual-Pol data, a specific type of non-uniform beam filling is required 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.

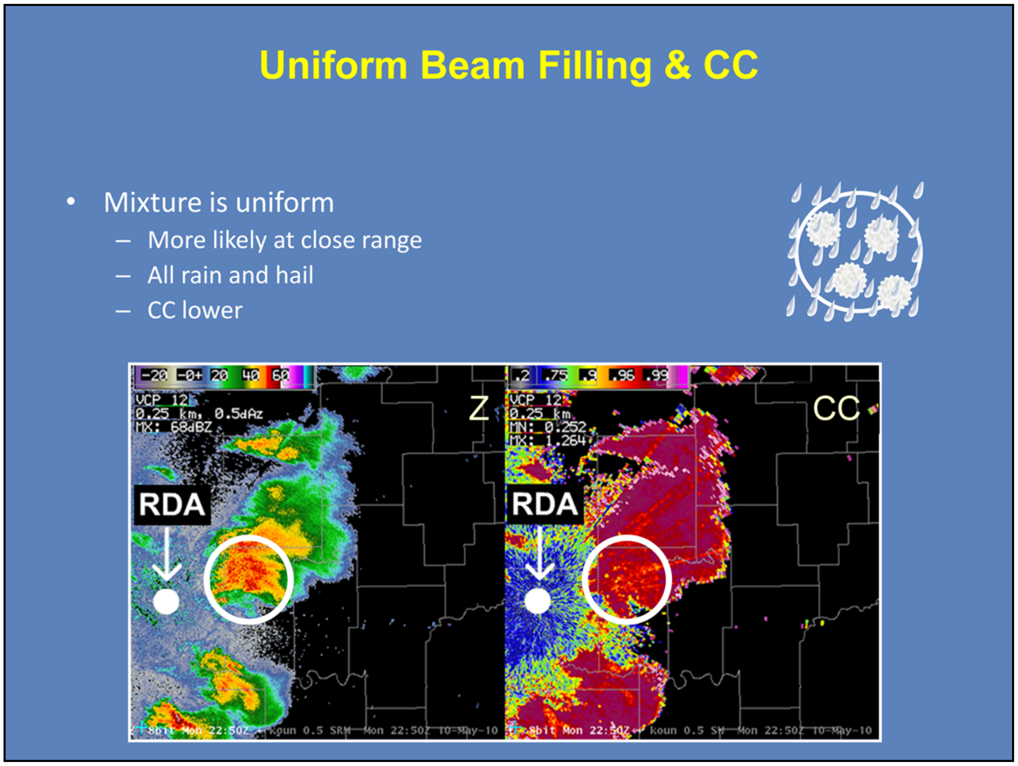

It turns out that Dual-Pol products are negatively impacted by what is called Non-Uniform Beam Filling (NBF), and there are examples coming up. Though non-uniform beam filling in the literal sense occurs frequently, we also use NBF to describe a specific type of signature on Dual-Pol products that results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

Partial-Uniform Beam Filling

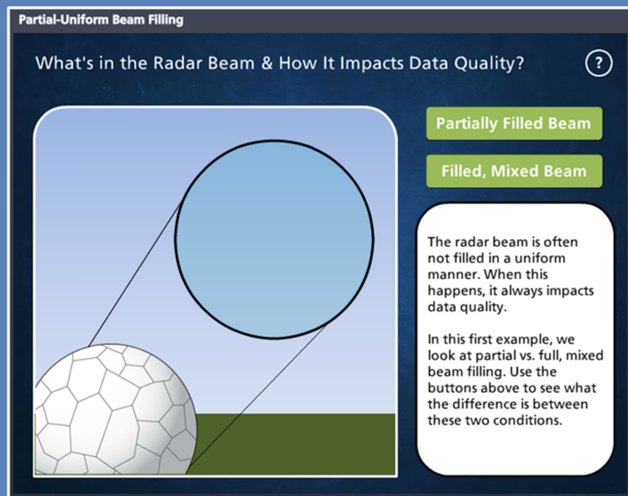

If no pop-up window appears that looks like the above, open a browser and go to:
http://training.weather.gov/wdtd/courses/rac/principles/interactions/partialbf/

Back to Top

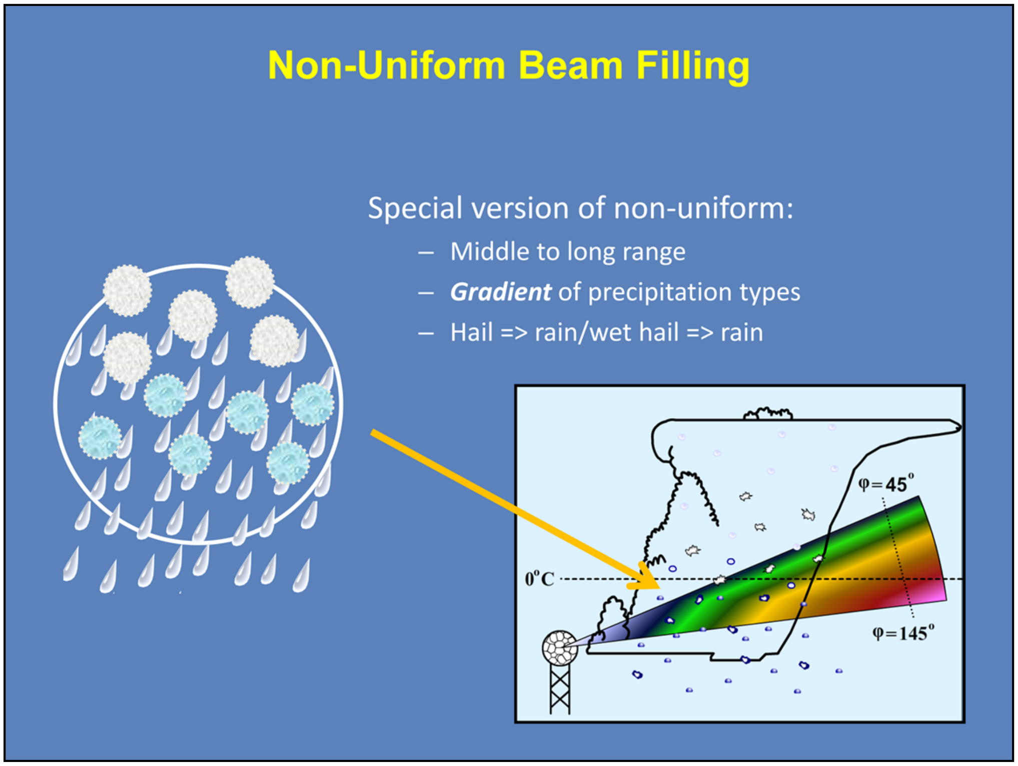

A non-uniform mixture can produce a gradient of precipitation types within the beam. This is more likely to occur at middle to long range. For example, the top of the beam may be sampling mostly hail, the middle sampling rain and wet hail, and the bottom of the beam sampling rain only. This gradient 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.

Back to Top

Non-Uniform Beam Filling & CC

Supercell at longer range:

- CC low at storm core and down radial due to NBF

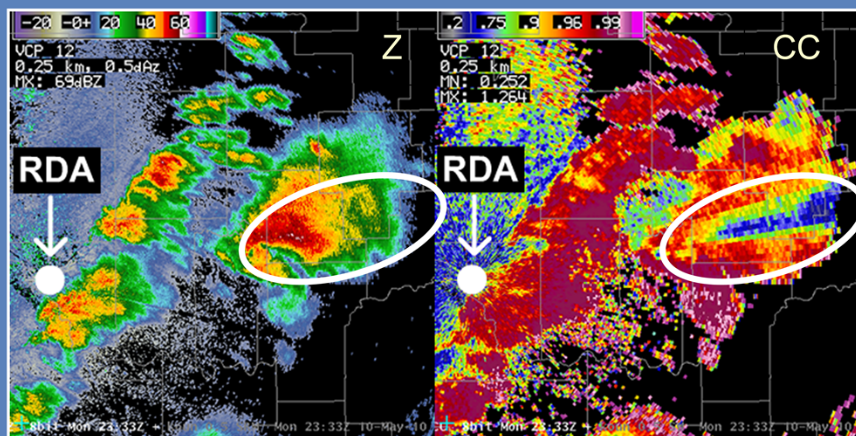

In the radar example, the supercell has moved to the east and is at a longer range, with the beam sampling a larger volume of the storm. There are radial swaths of low CC that originate from the storm core. This is an example of non-uniform beam filling and its impact on the CC product. This has impacts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

Back to Top

Non-Uniform Beam Filling and CC

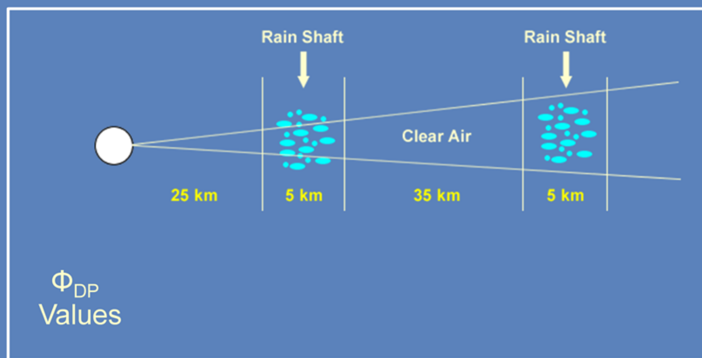

Φ_{DP} propagates down radial:

- Increase proportional to liquid water
- Hydrometeors uniformly distributed

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.

Back to Top

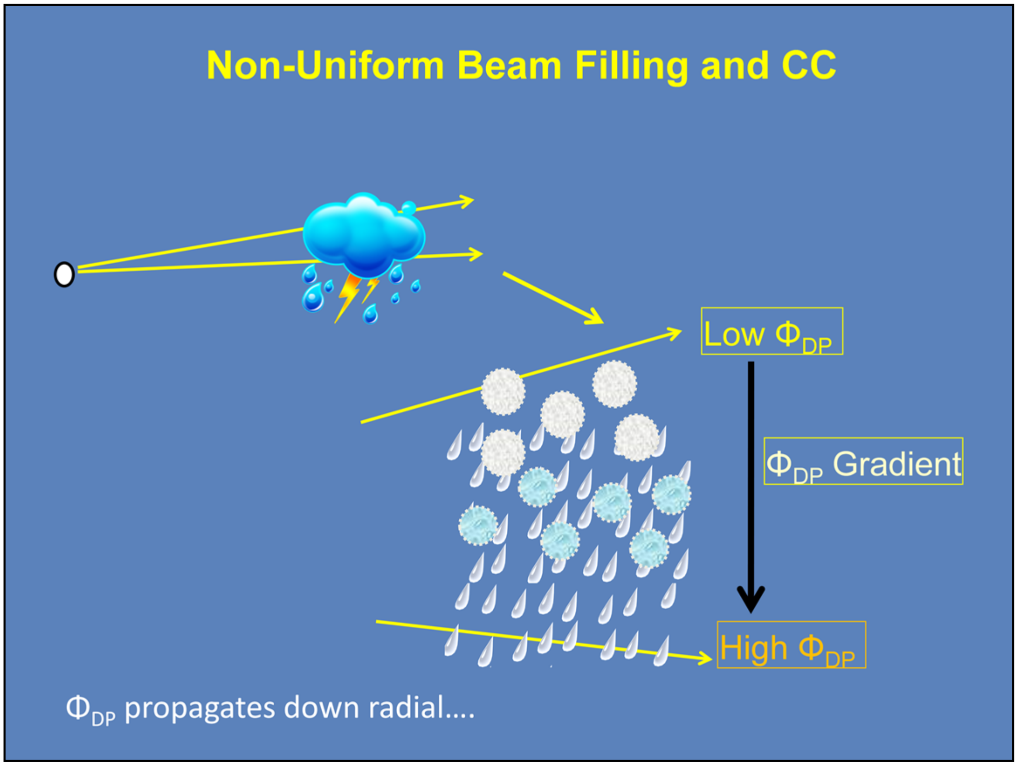

When sampling a convective storm at longer range or a squall line along a radial, there is an increasing chance of capturing a gradient of precipitation types within the beam. At the top can be hail and/or graupel, while the bottom of the beam is sampling liquid drops.

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.

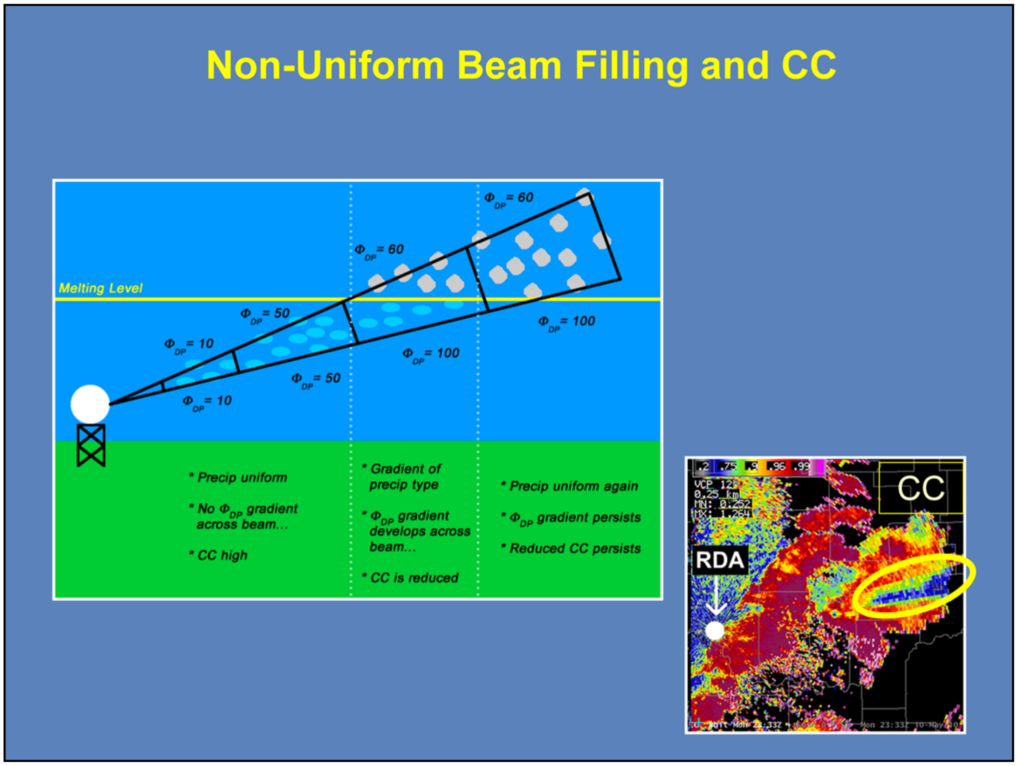

Here's a super simple example of what happens to the CC down radial with only four range bins. Note the Φ_{DP} values at the top of the beam and at the bottom of the beam for each of these bins.

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.

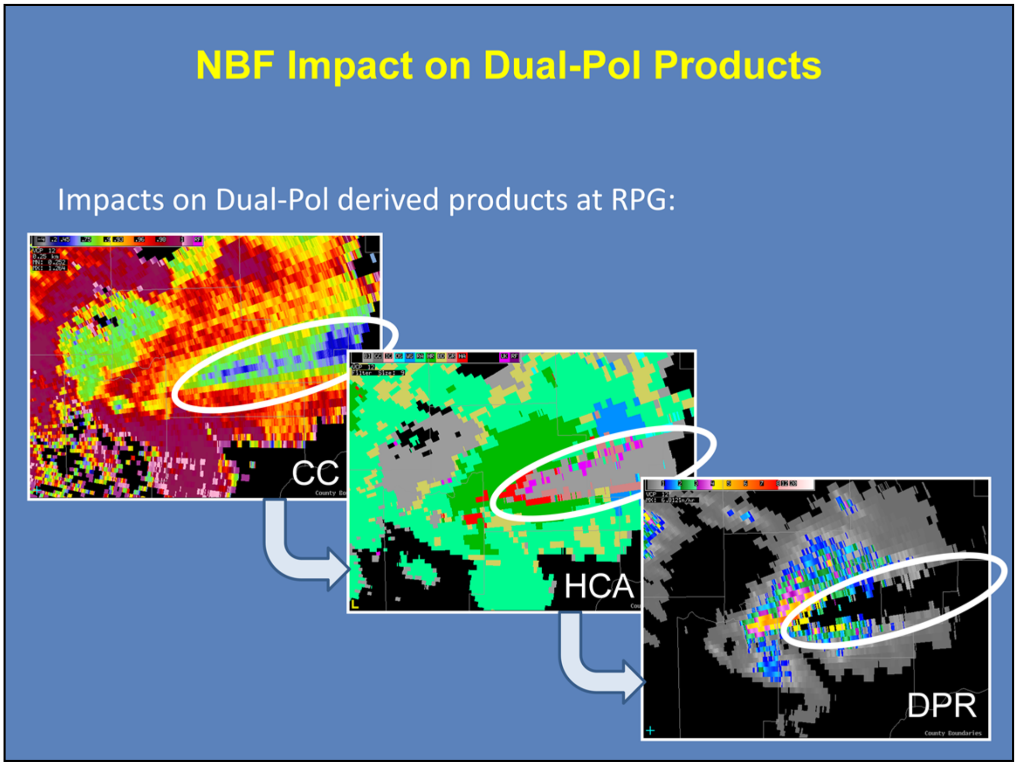

The artifact of a swath of low CC values due to NBF can be either easy to spot or subtle. By comparing it to other radar data and understanding the environment, you can ask yourself if the CC values make sense.

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.

Back to Top

Non-Uniform Beam Filling

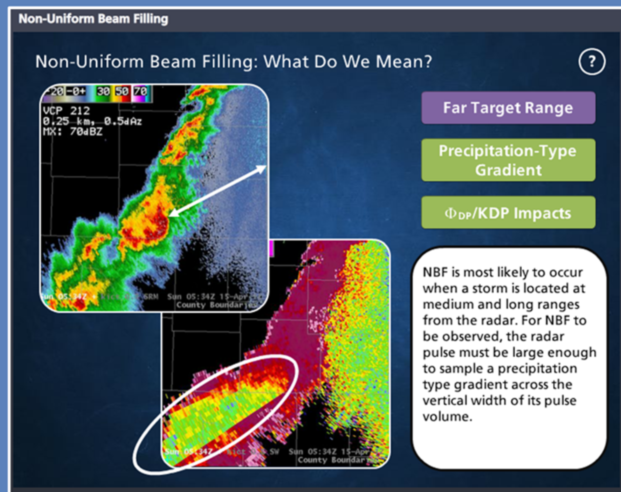

If no pop-up window appears that looks like the above, open a browser and go to:
http://training.weather.gov/wdtd/courses/rac/principles/interactions/nbf/

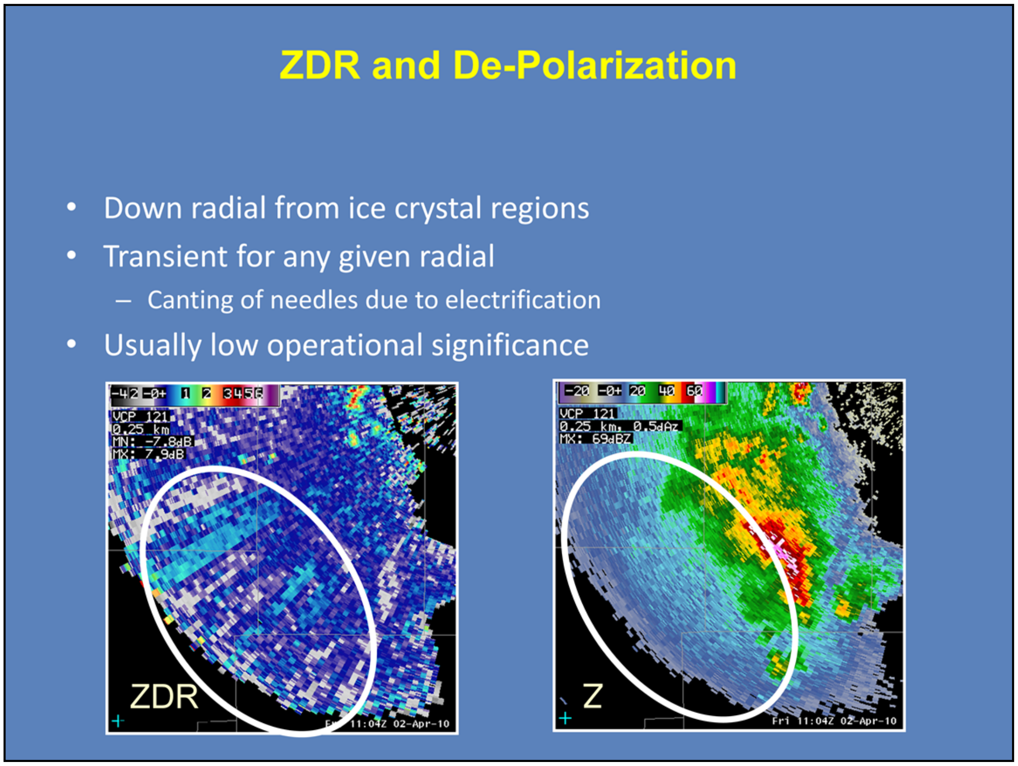

Depolarization will sometimes be apparent on the ZDR product. Depolarization means that the reflected energy from a particle switches polarization, from horizontal to vertical, vertical to horizontal, are maybe both.

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.

Back to Top

WSR-88D Data Quality

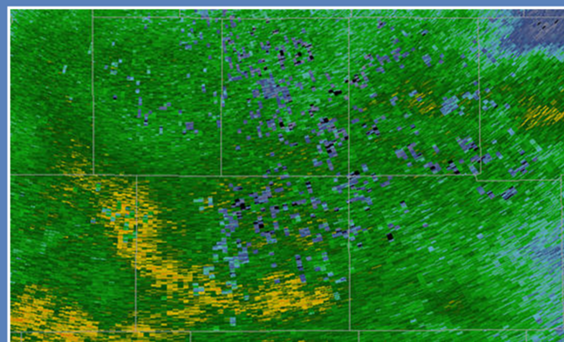

- Trade off:
  - high quality base data vs.
  - low level sampling and fast product updates
- Benefits of VCPs 12, 212, & 121 are obvious
  - Use ‘em when you need ‘em!

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.

Back to Top

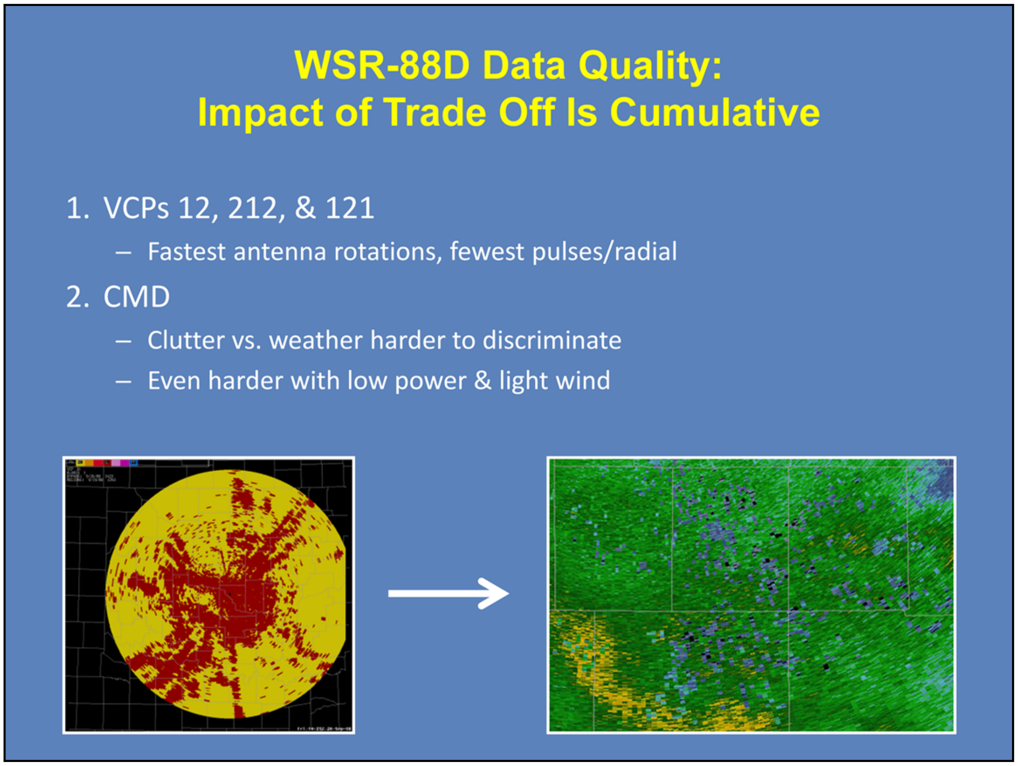

The impact of this trade off is cumulative, and I focus on VCPs 12, 212, and 121, since they have the fastest antenna rotation rates and thus the fewest number of pulses per radial.

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.

WSR-88D Data Quality: Impact of Trade Off Is Cumulative

3. GMAP
   - Filtering applied only to bins identified by CMD
   - Less effective at rebuilding lost weather signal with fewer pulses
   - More bins with significant signal loss

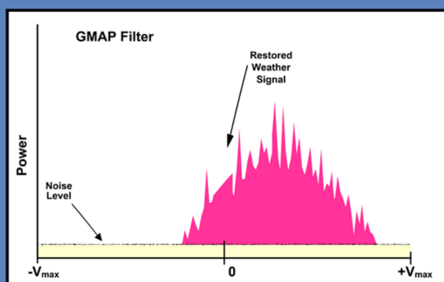

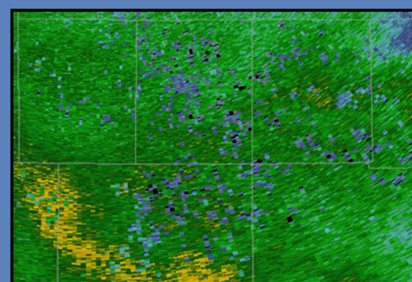

Once CMD identify which bins need to have clutter suppression applied, GMAP does the actual signal removal. It first isolates the clutter signal near zero velocity, then removes power just from the notch, or interval, around zero velocity. One of the strengths of GMAP is its ability to rebuild a lost 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.

WSR-88D Data Quality:
Impact of Trade Off Is Cumulative

4. Super Resolution
   - Better spatial resolution, but…
   - Windowing used for Super Res increases error/noisiness in base data
   - Noisiness can approach human tolerance

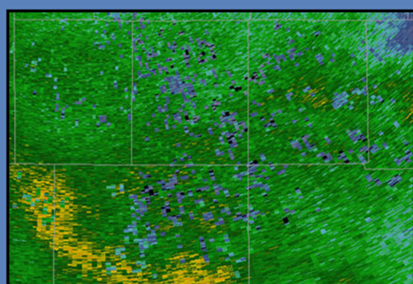

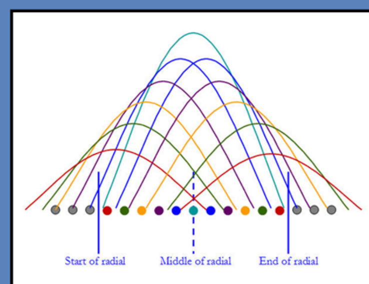

Super resolution processing is another trade off. The benefit of better spatial resolution is obvious, but super resolution processing includes a windowing technique that introduces some error in the base data estimate. The cost is an increase in noisiness in the base data. The fewer the pulses per radial,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

Back to Top

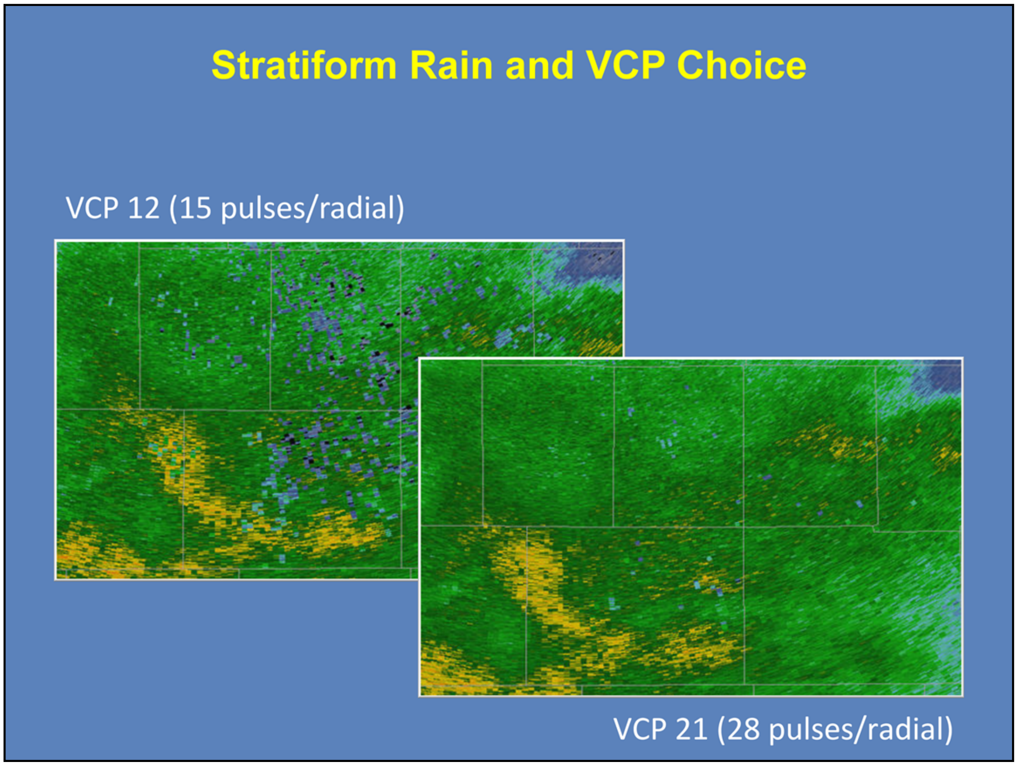

Now for the grand finale on the impact of this trade off between the need for high quality base data, along with fast product updates and low level sampling. This is a stratiform rain event, initially with VCP 12, which is the image on the left hand side. The staff noticed numerous gates of data loss 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.

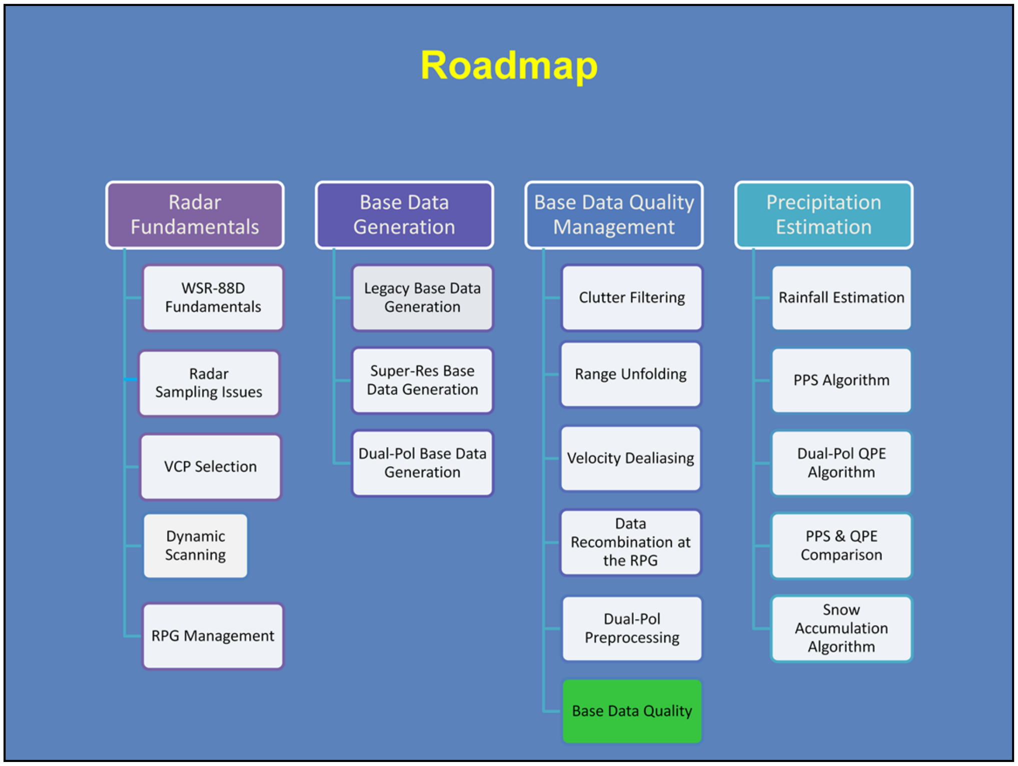

Here is the “roadmap” with your current location.

Welcome to Radar Rainfall Estimation Errors.

Back to Top

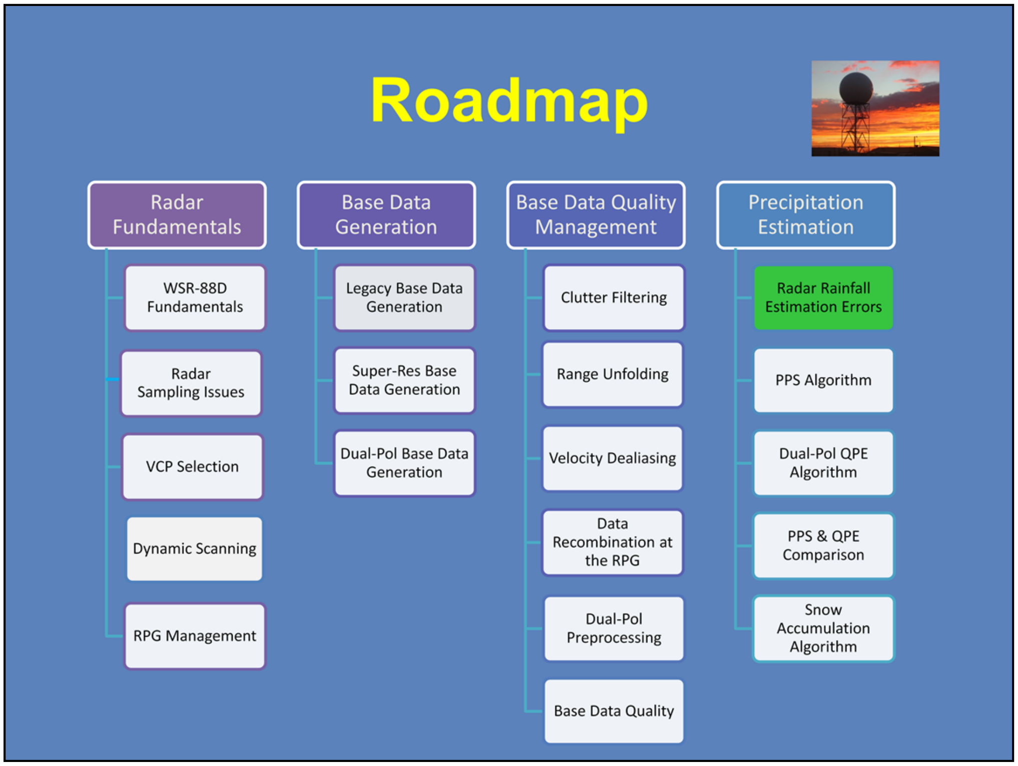

Here is the “roadmap” of the “Principles of Radar” topic in RAC. This lesson on Rainfall Estimation is the first of five lessons on precipitation estimation.

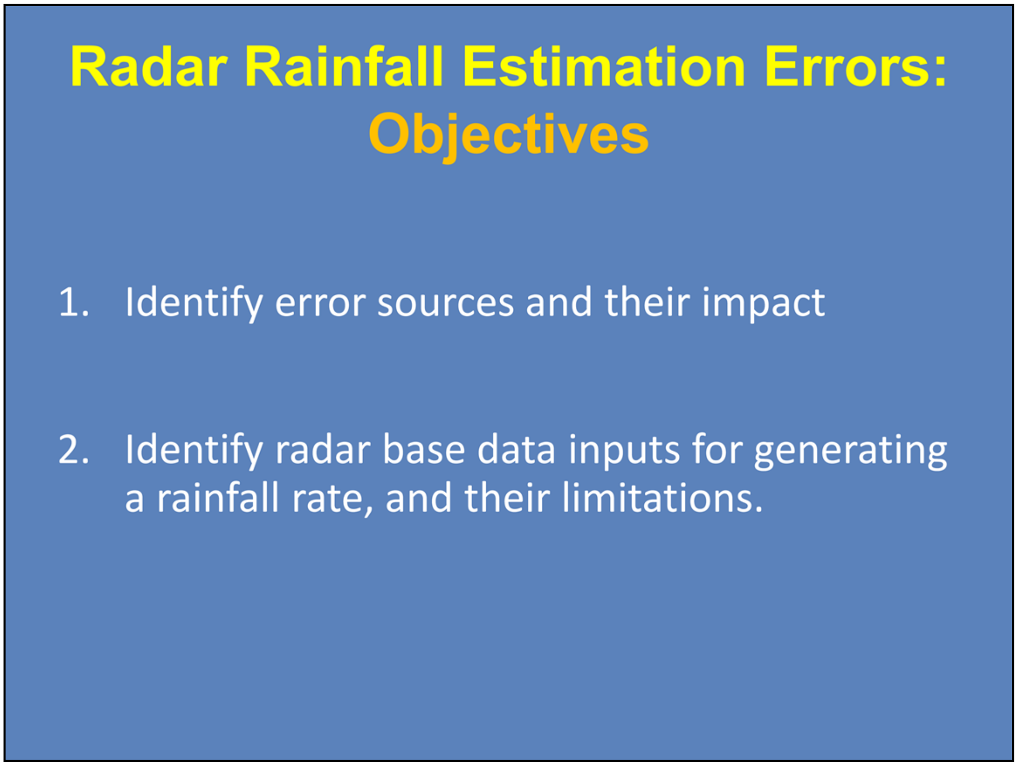

These are the two learning objectives for this lesson.

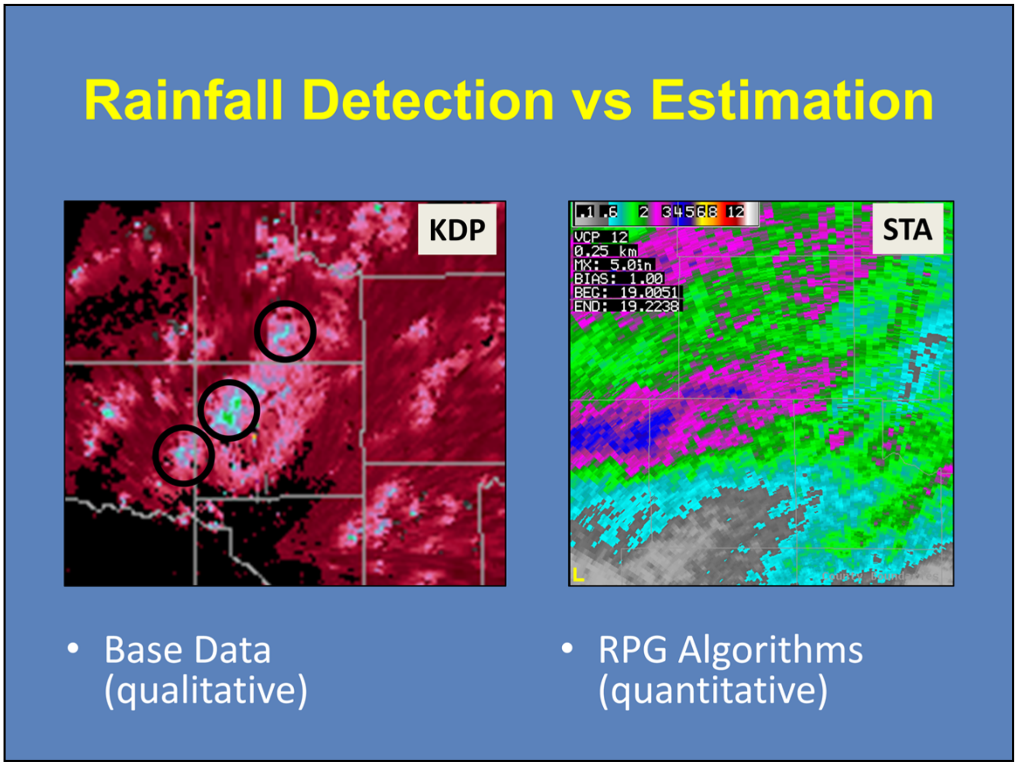

Let's start by making a distinction between using the base products to identify areas of heavy rainfall vs. using the output from an RPG algorithm to estimate rainfall amounts.

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.

Back to Top

Issues When 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.

Back to Top

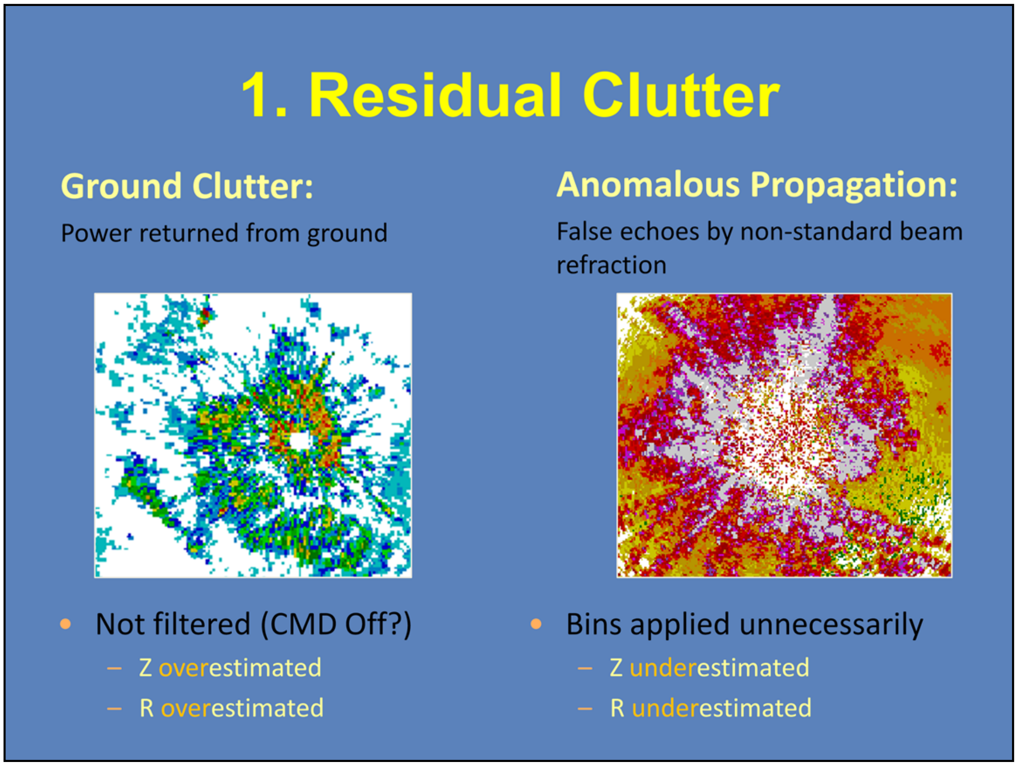

Residual Clutter describes unfiltered clutter, and can occur when clutter filtering has been turned off or when clutter remains after filtering is applied. Ground clutter indicates returns from certain ground based targets are always present, such as buildings and terrain. Technically, Anomalous Propagation (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.

Back to Top

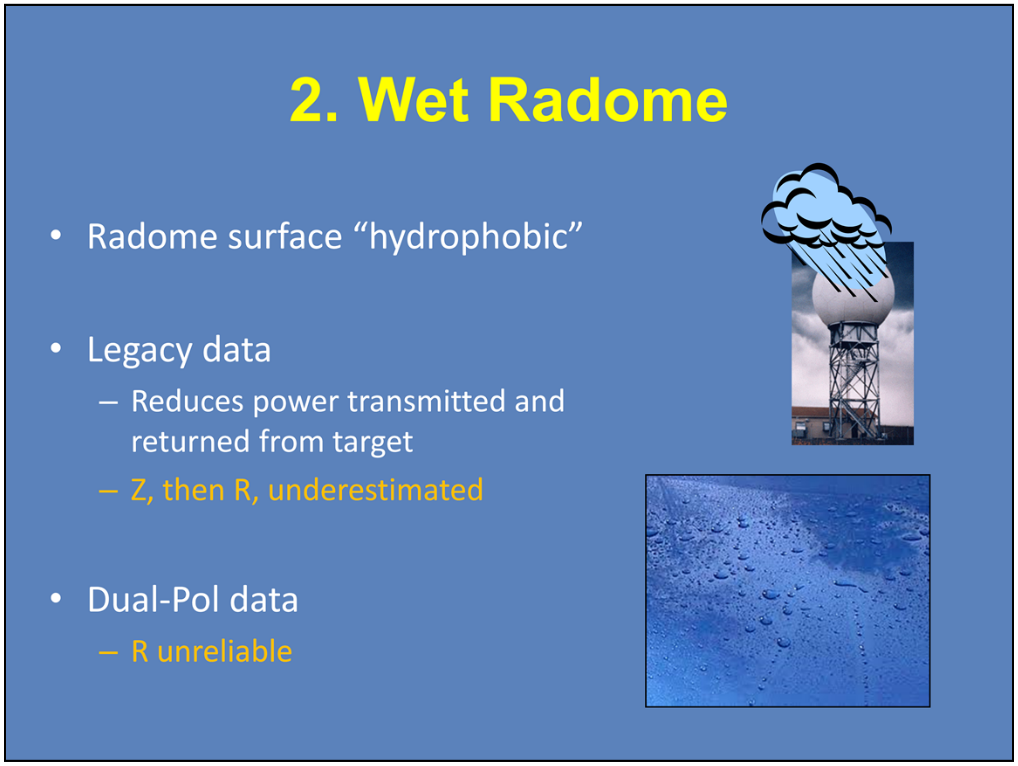

A wet radome can cause multiple issues with base data. The radome surfaces are designed to be "hydrophobic", repelling water like wax on a car, in order to prevent water coating the surface. It still happens, of course, but the condition is usually transient.

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.

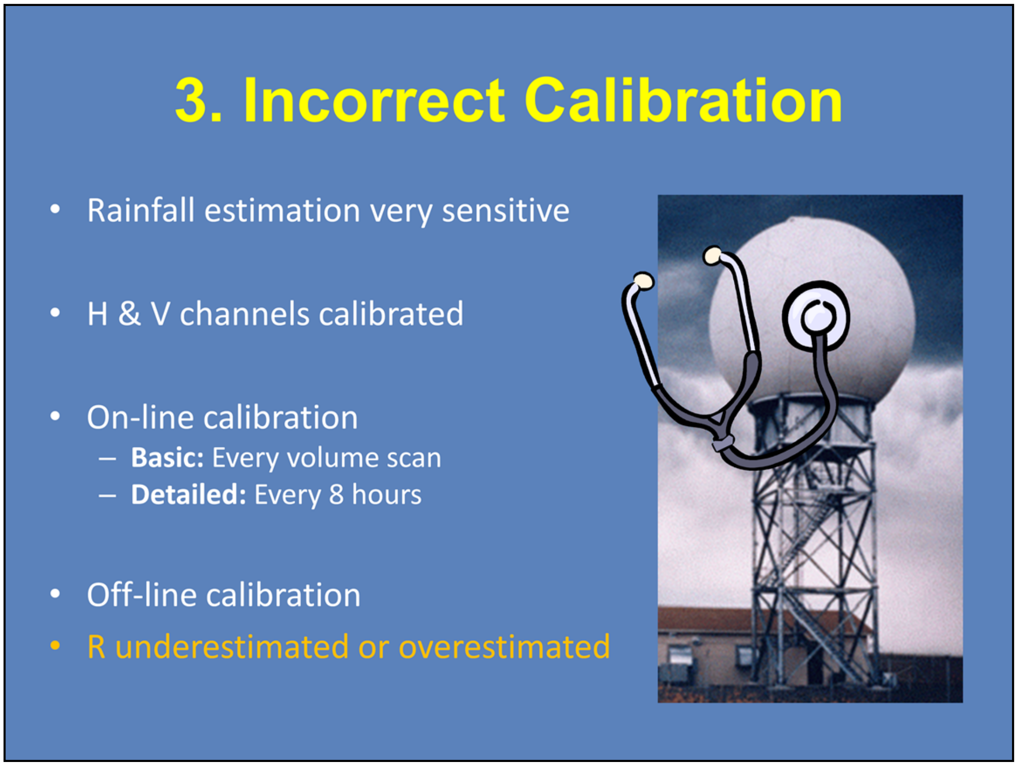

Rainfall estimation is particularly sensitive to calibration errors. A valid Z value is dependent on a well calibrated horizontal channel, while a valid ZDR value depends on each of the horizontal and vertical channels being well calibrated.

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.

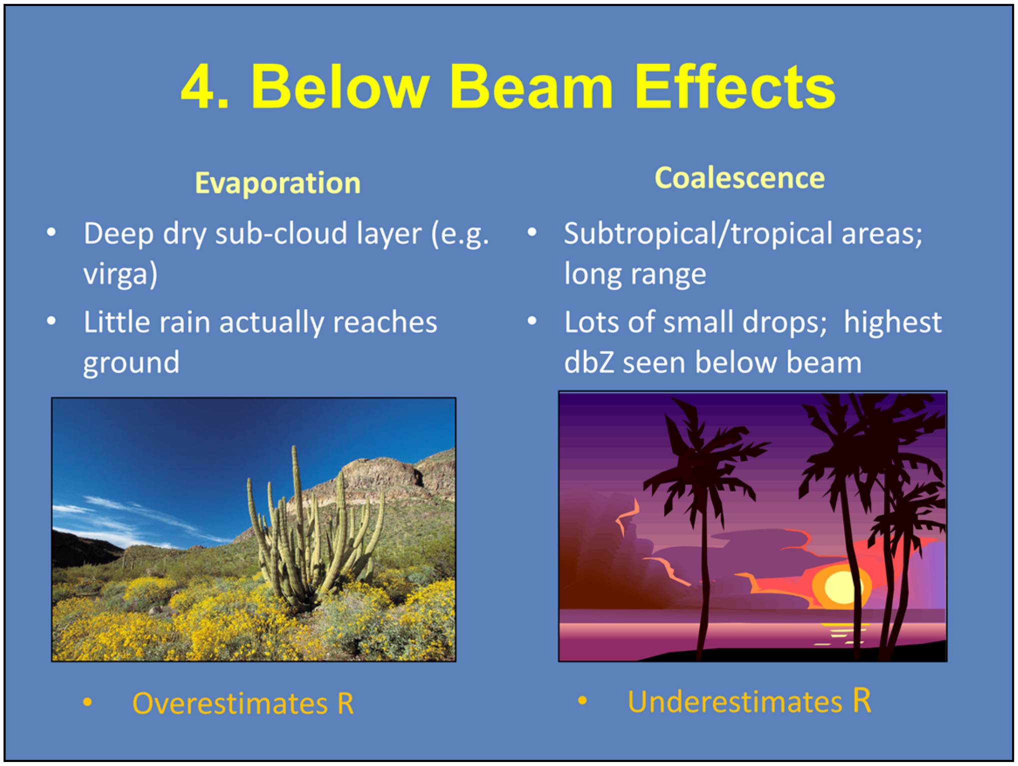

Below beam effects are a fundamental challenge when using a radar to estimate rainfall. We must remember that rainfall estimates are based on the hydrometeors that were sampled within the radar beam. Depending on range, it may be a very long way down to reach the ground! Consideration of the layer below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

Back to Top

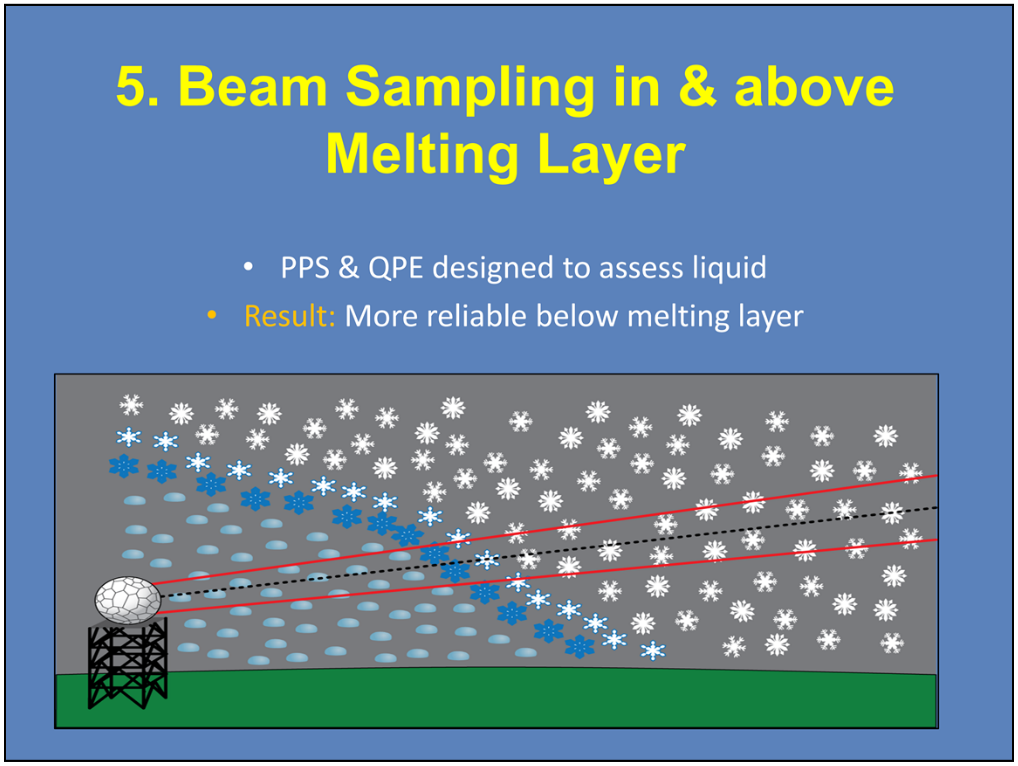

It is important to remember that both the RPG rainfall algorithms are designed to estimate liquid rainfall. Algorithm performance is much more reliable at locations where the radar beam is intercepting liquid hydrometeors below the melting layer. Within the melting layer and above, hydrometeors are frozen (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

$$Z = 300R^{1.4}$$

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

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

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

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

- 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.

Back to Top

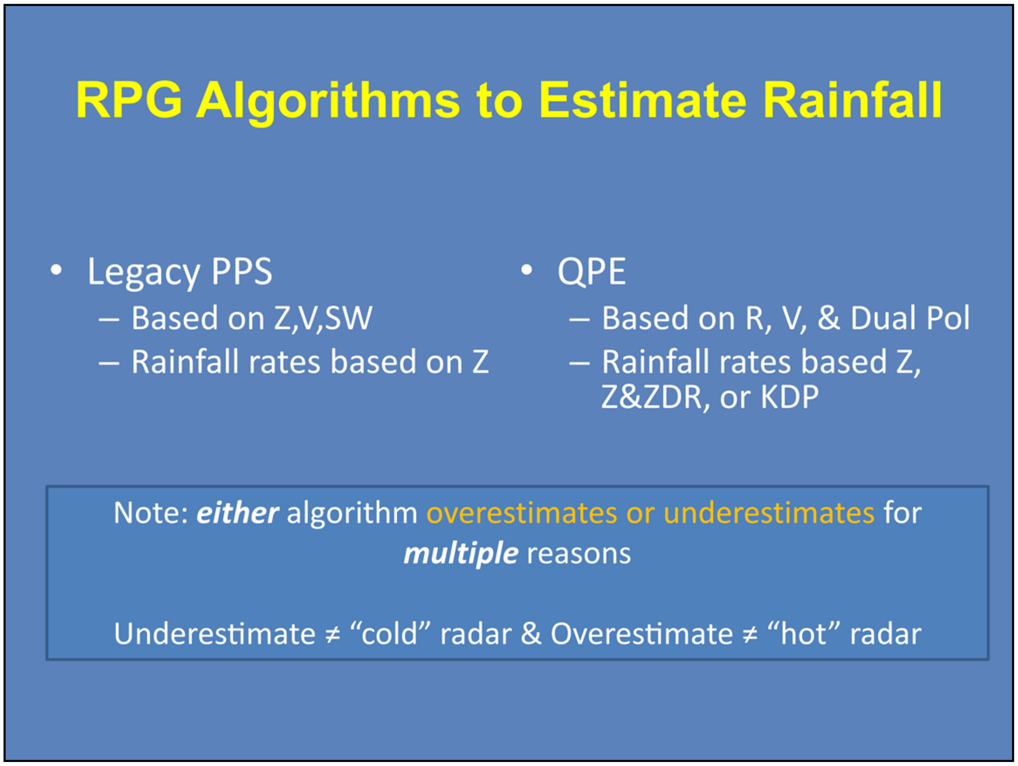

There are two RPG algorithms designed to estimate liquid rainfall at the surface. The first is called the Legacy Precipitation Processing SubSystem (PPS), and it has been in place since the original deployment of the WSR-88D. The PPS has seen substantial design changes over the years, as with many of 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.

Back to Top

Z is Estimated

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

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

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

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

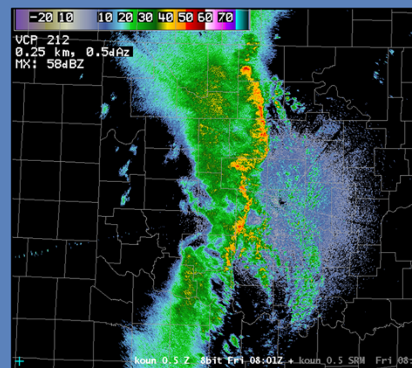

Estimating rainfall from radar is a very complicated business. It's amazing it works as well as it does! Remember that there are potential errors in the Z values even before a rain rate is calculated. Here are the limitations on the validity of the Z value itself.

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.

Back to Top

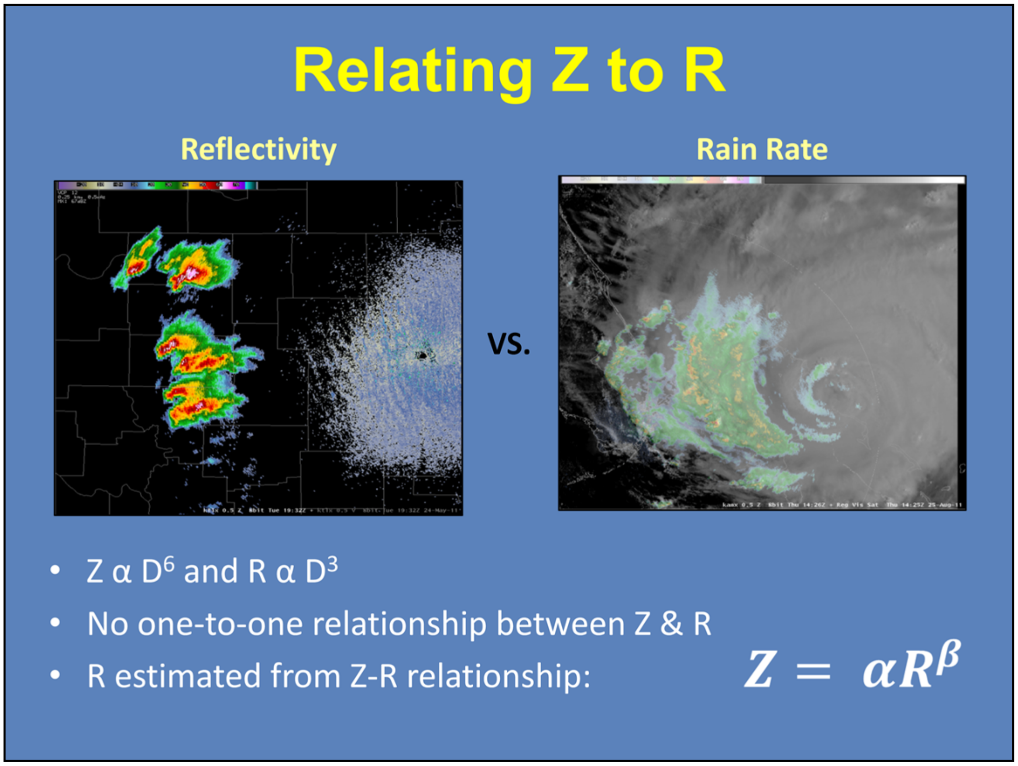

Both Z and R are dependent on the dropsize distribution, which is unknown, but with different dependencies. Z is proportional to the drop diameter to the sixth power, while the rainfall rate is proportional to the drop diameter to the third power. There is no one-to-one relationship between Z and R. 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.

Back to Top

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.

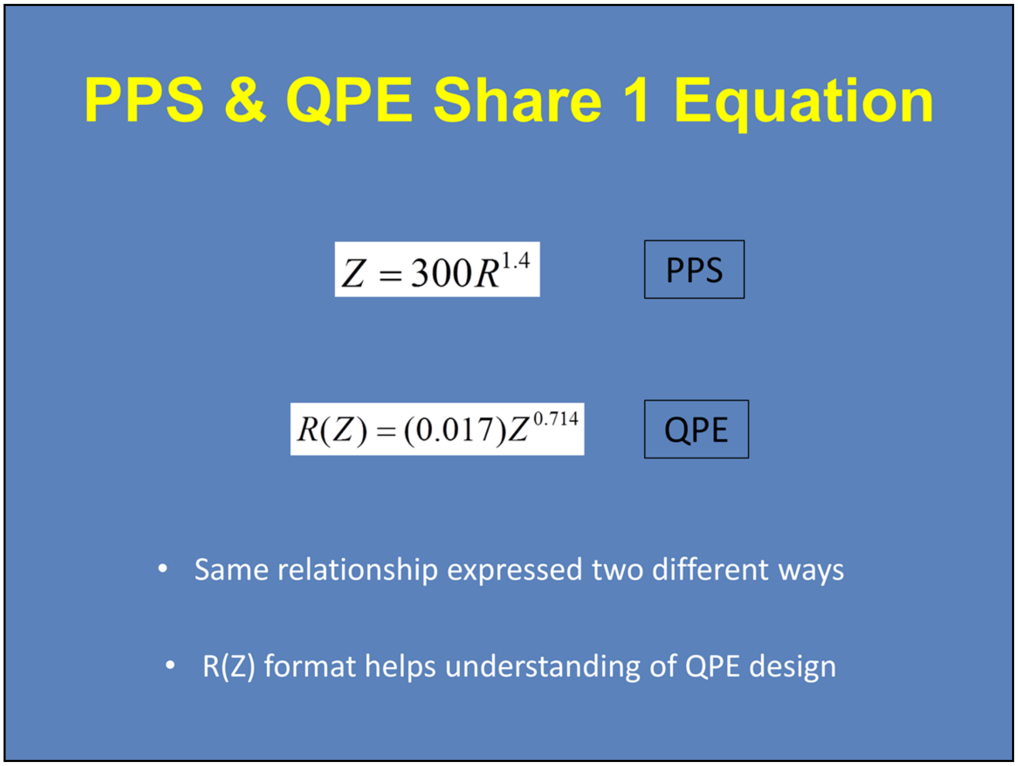

The legacy PPS and the Dual-Pol Quantitative Precipitation Estimation (QPE) algorithm share one equation for converting Z to rainfall rate. The first equation shows the more familiar format, solved for Z. The second equation is the same as the first, just rewritten in a format that is solved for R. We 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.

Back to Top

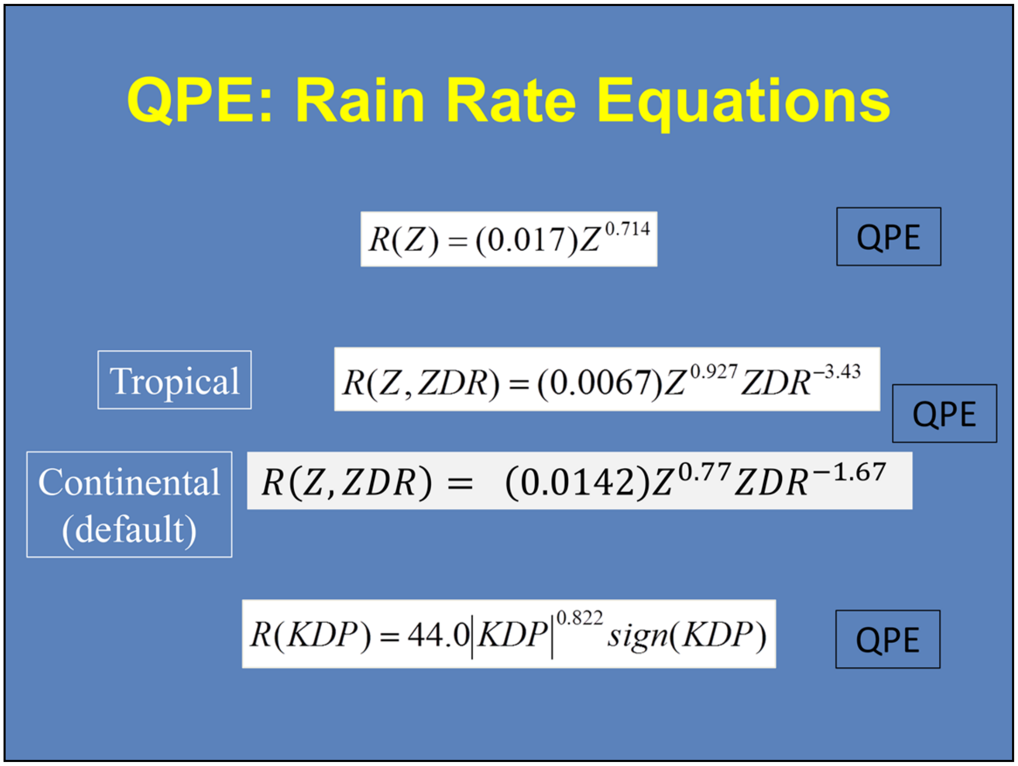

There are four equations for computing rain rate used by QPE. How QPE determines which of these equations to use will be explored in a later lesson.

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.

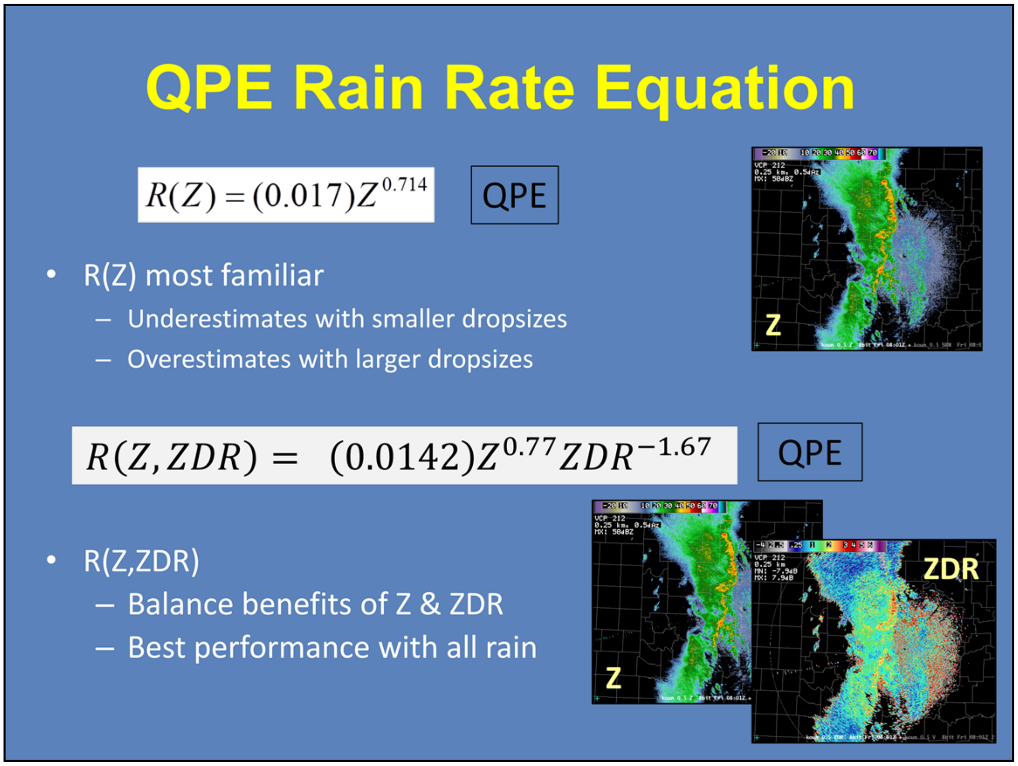

The limitations of $Z = 300R^{1.4}$, now presented in its R(Z) format, are probably familiar. It often underestimates rainfall in a warm rain dominant event, while overestimates the water at the surface from mixed precipitation types, such as wet snow and hail.

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.

Back to Top

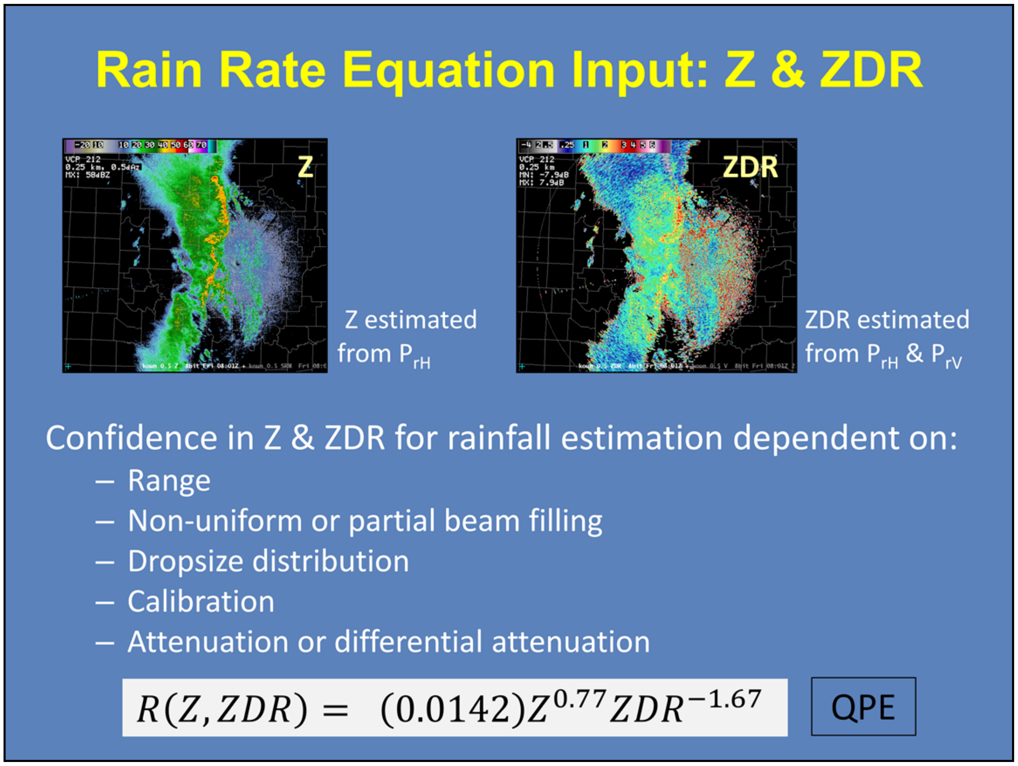

Using Z and ZDR as inputs for computing rain rate have similar dependencies that affect their reliability for rainfall estimation. As a reminder, Z is estimated from the returned power in the horizontal channel, while ZDR is estimated from the returned power in both the horizontal and vertical channels.

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.

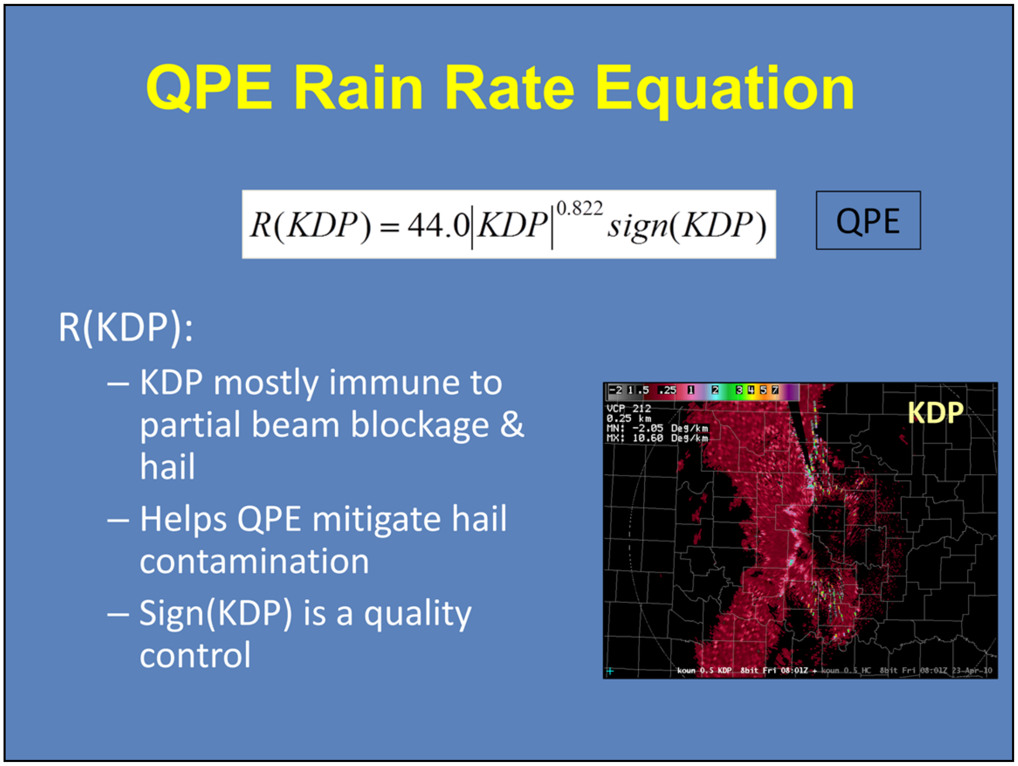

R(KDP) has the advantage that KDP is mostly immune to both partial beam blockage and hail contamination. The magnitude of KDP is related to the liquid water content in the volume, but not to ice. For example, if a volume contains mostly hail, KDP values can be quite low. On the other hand, a large quantity 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.

Back to Top

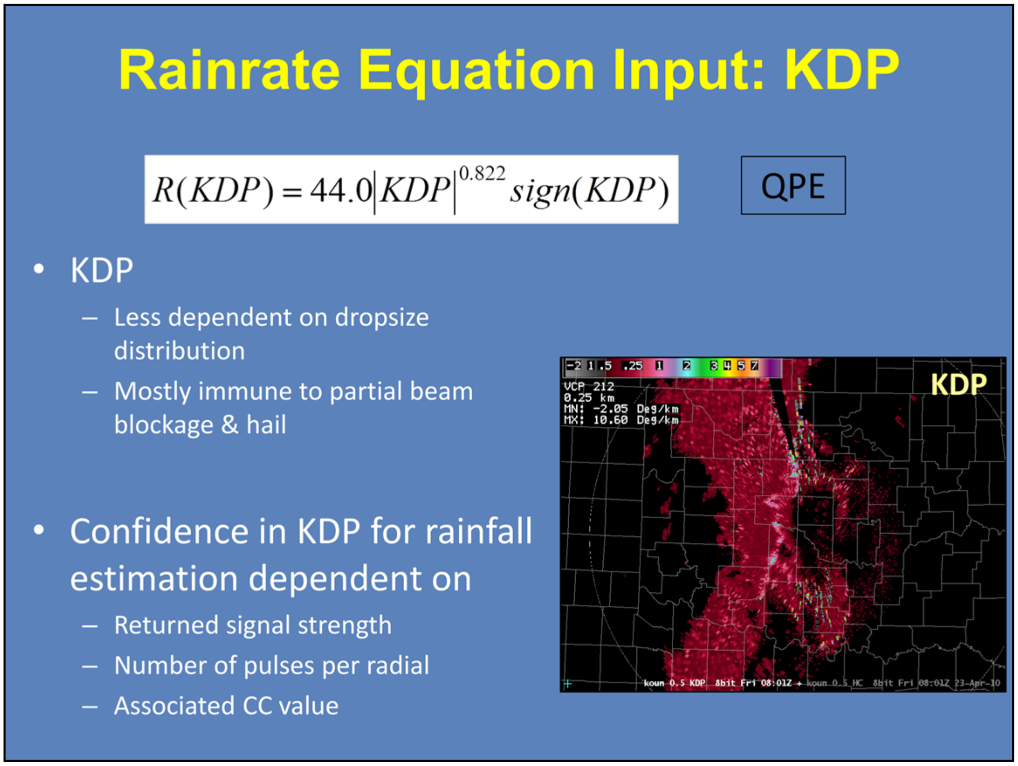

KDP as input for computing rainrate is less dependent on dropsize distribution than Z and ZDR. KDP can be useful where hail is present or there is partial beam blockage.

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.

Back to Top

Welcome to Legacy Precipitation Processing Subsystem (PPS) Algorithm.

Back to Top

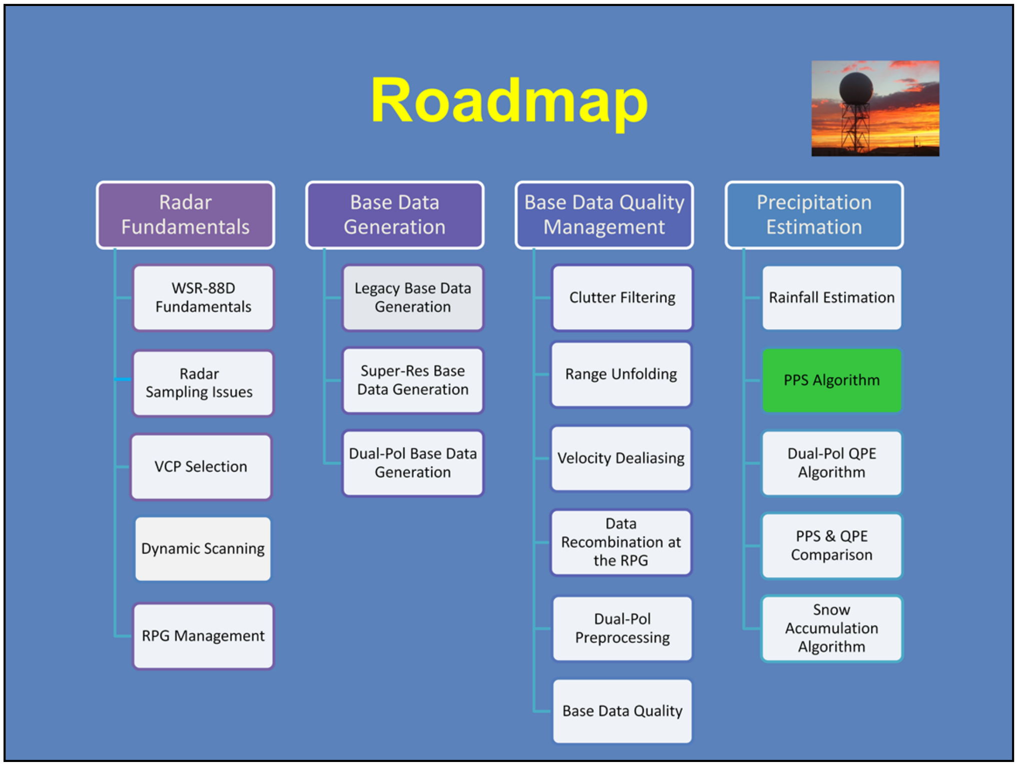

Here is the “roadmap” with your current location.

Learning Objectives

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

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

Back to Top

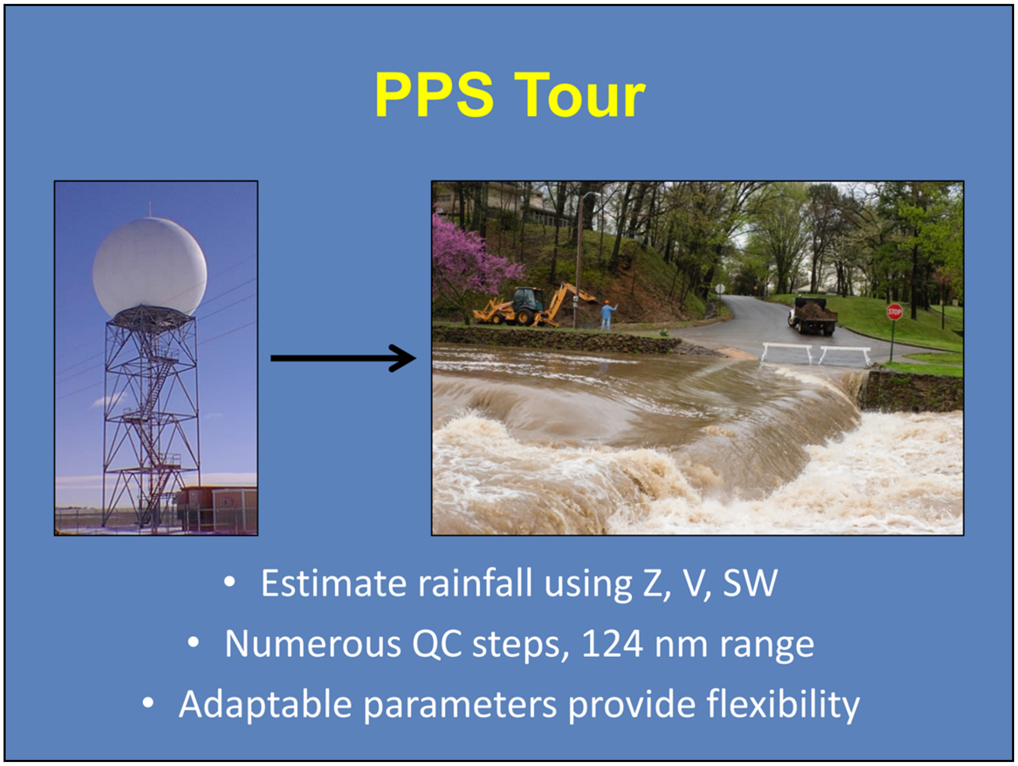

Now for a “tour” of the Legacy (aka been around awhile) Precipitation Processing Subsystem (PPS). The PPS is a series of algorithms that use base reflectivity, velocity and spectrum width data as input, then estimate rainfall and generate rainfall products. This process contains several quality control 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.

Back to Top

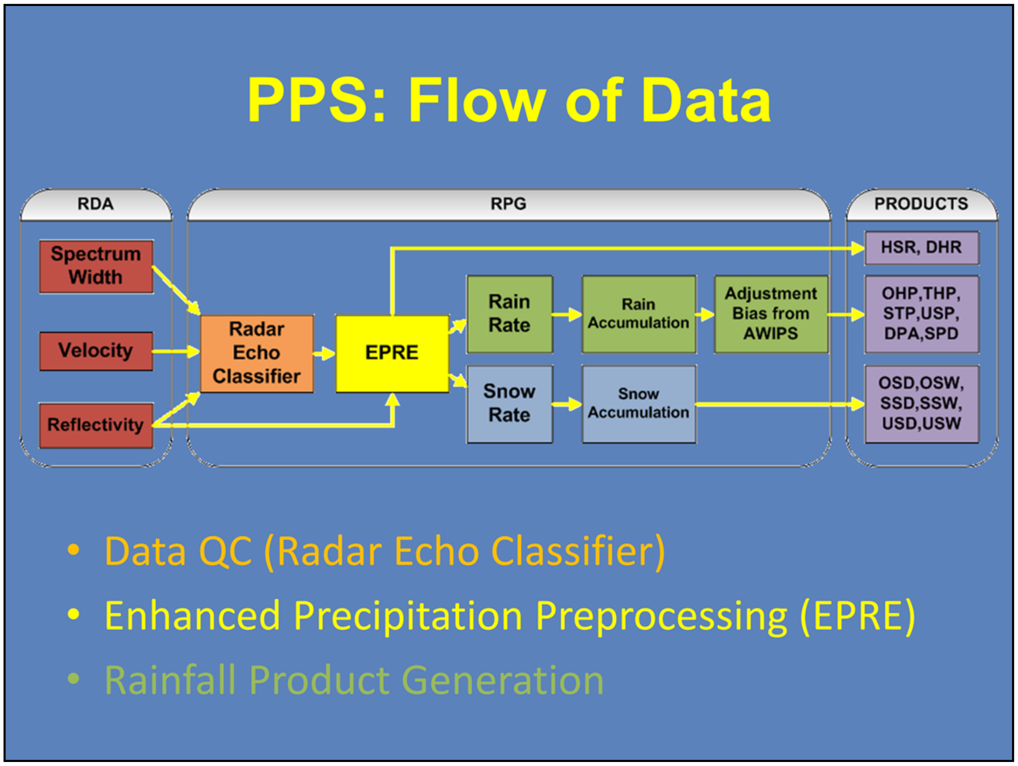

This graphic provides a high-level look at the flow of information through the Precipitation Processing Subsystem (PPS) and will serve as an outline for the PPS components. There are multiple quality control steps that will be highlighted, including a look at how the base data are “prepared” by the Enhanced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

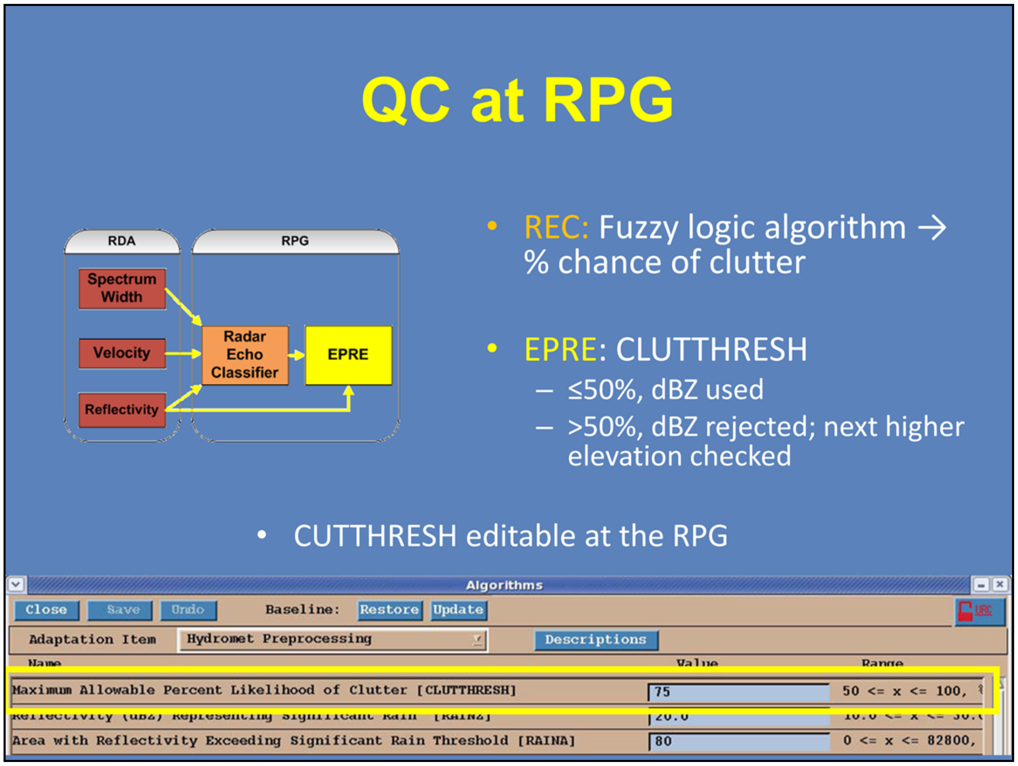

The Radar Echo Classifier is a fuzzy logic algorithm that uses reflectivity, velocity, and spectrum width to assign a likelihood that a particular bin contains clutter.

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.

Back to Top

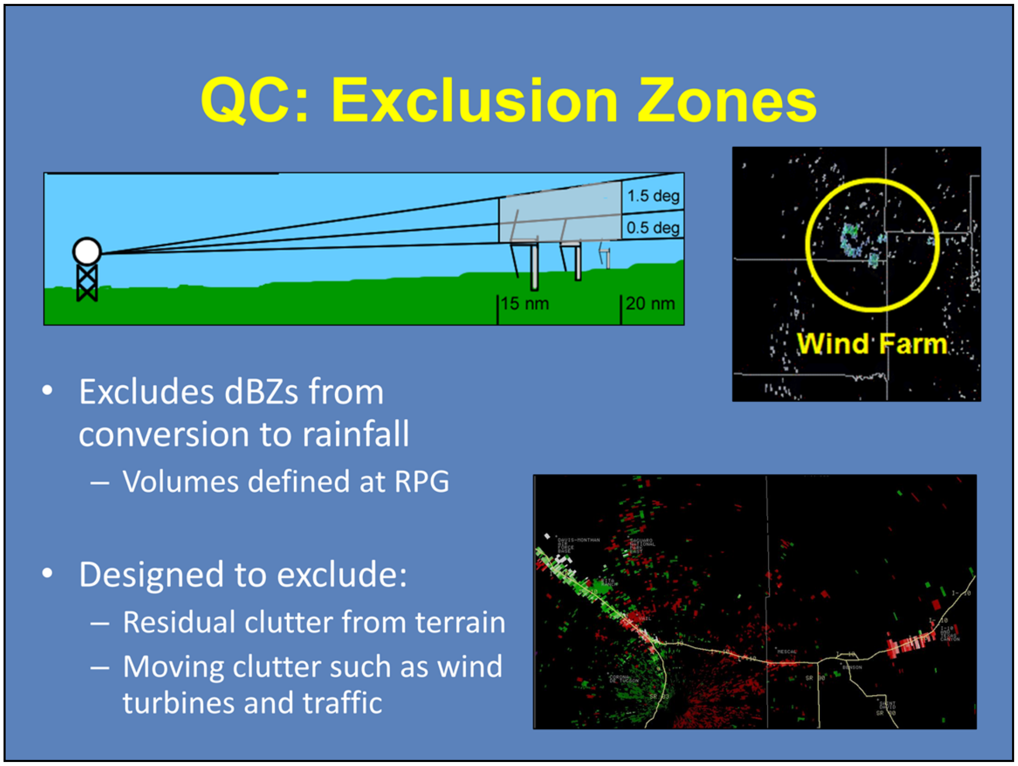

Another quality control option that is part of EPRE is the application of Exclusion Zones, which are defined locally. Exclusion zones are used to prevent reflectivity from specific areas and elevations from being converted to rainfall. An exclusion zone is actually a volume, defined from azimuth to azimuth,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

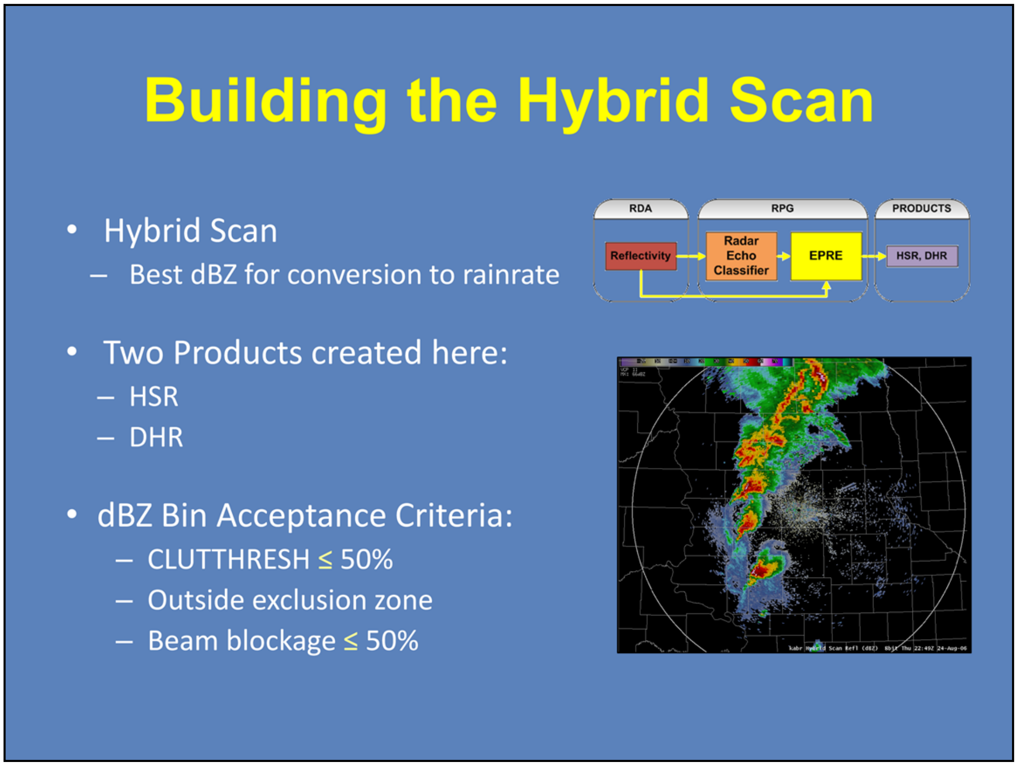

The "grand finale" of the EPRE algorithm is the building of the hybrid scan. The idea of a hybrid scan is to find an optimal dBZ value at each range bin only for the purpose of converting to rainrate. There are two Hybrid Scan products that represent the reflectivity field that was used by the PPS for 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%

Back to Top

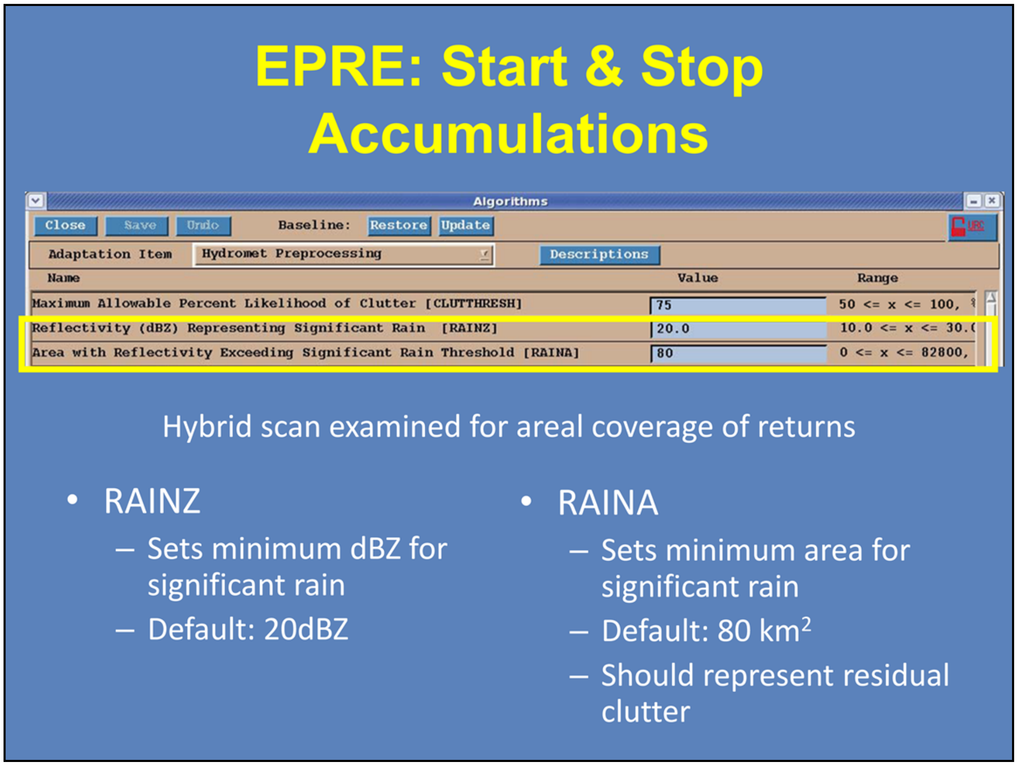

Based on the hybrid scan, EPRE determines when accumulations begin and when they end. The idea of whether or not it is "raining" is based on the areal coverage of returns above a certain dBZ value. There are two EPRE adaptable parameters that govern the start and stop of rainfall, RAINZ and RAINA. RAINZ 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.

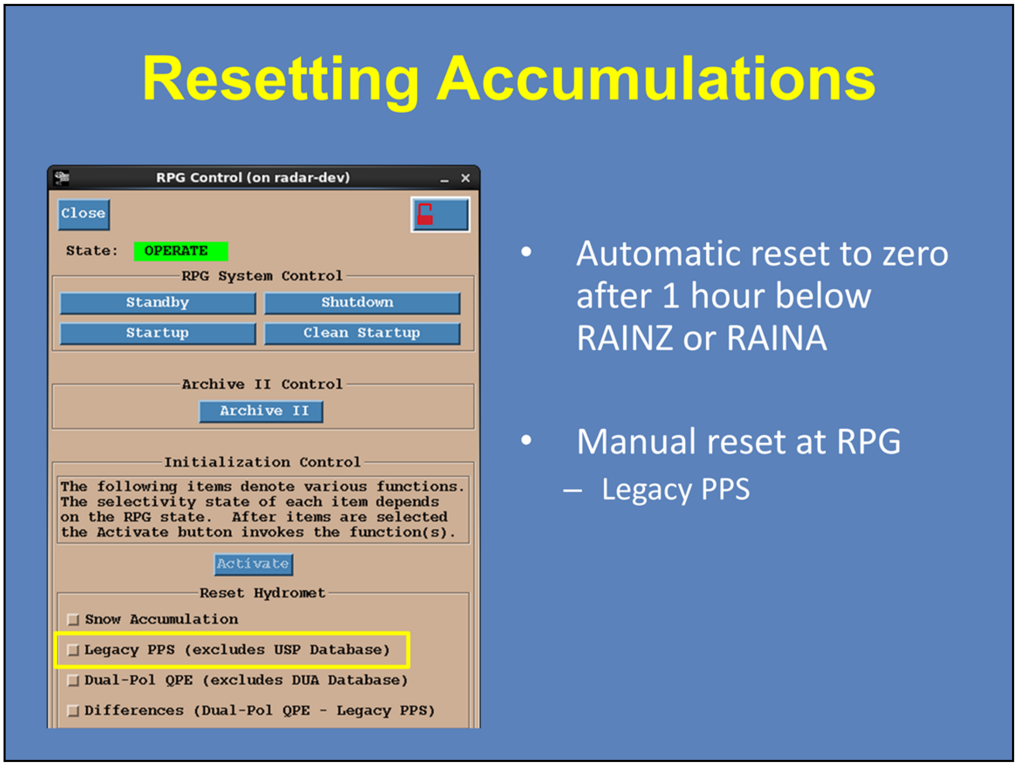

There are two approaches to resetting rainfall accumulations to zero. The first is the automatic reset when the conditions fall below RAINZ and RAINA for one hour. The storm total is a type of accumulation that continues as long as RAINZ and RAINA are exceeded. For some locations, this can be too long!

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”.

Back to Top

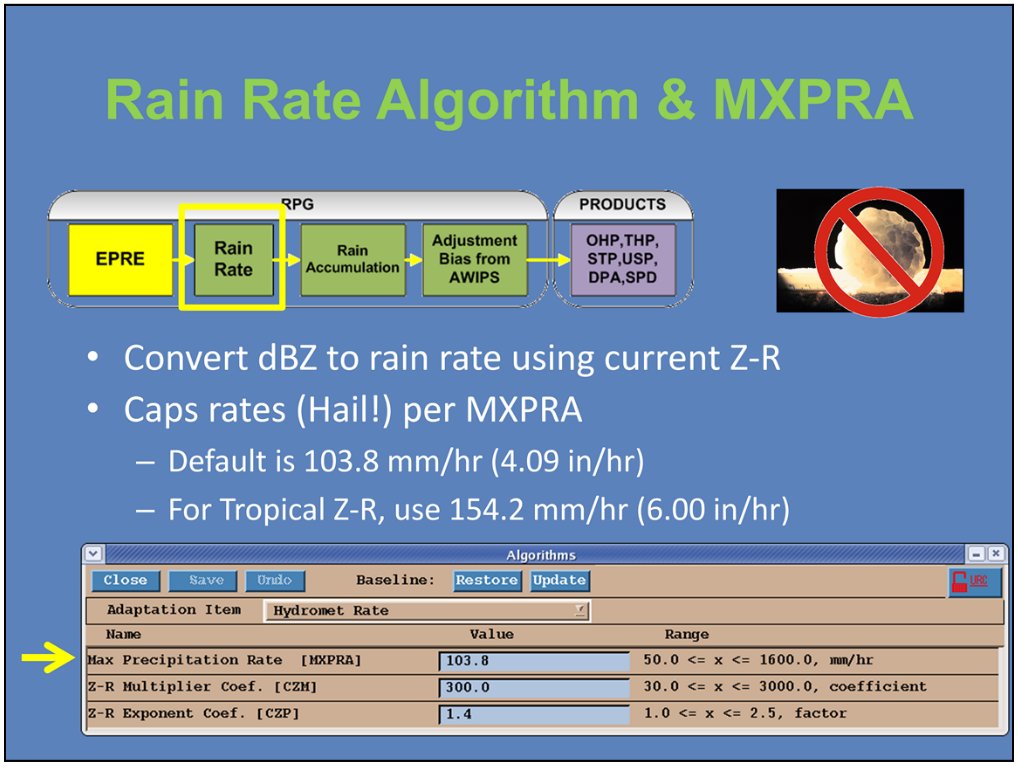

The Rain Rate algorithm converts the dBZ values from the hybrid scan to rain rate using the current Z-R relationship applied at the RPG. The Rate Algorithm also applies a parameter called the Max Precipitation Rate (MXPRA). MXPRA works as a cap to prevent hail contamination. The default setting for MXPRA 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.

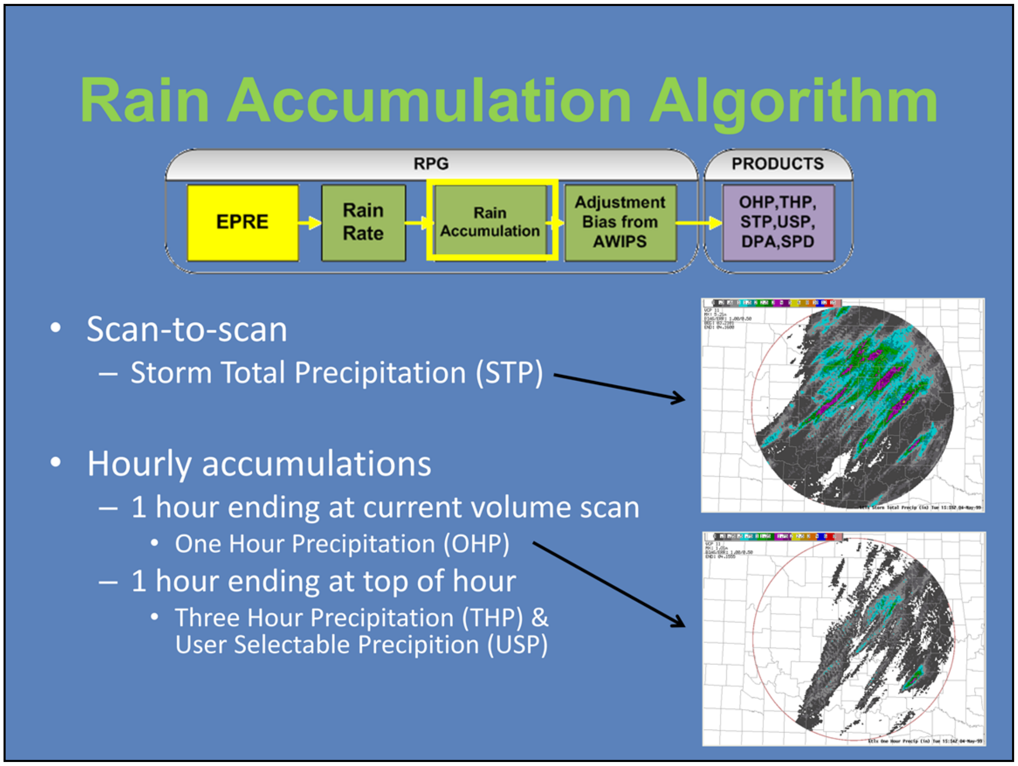

The Accumulation Algorithm, uses the calculated rain rates and differing durations to accumulate rainfall. There are two different types of accumulations.

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.

Back to Top

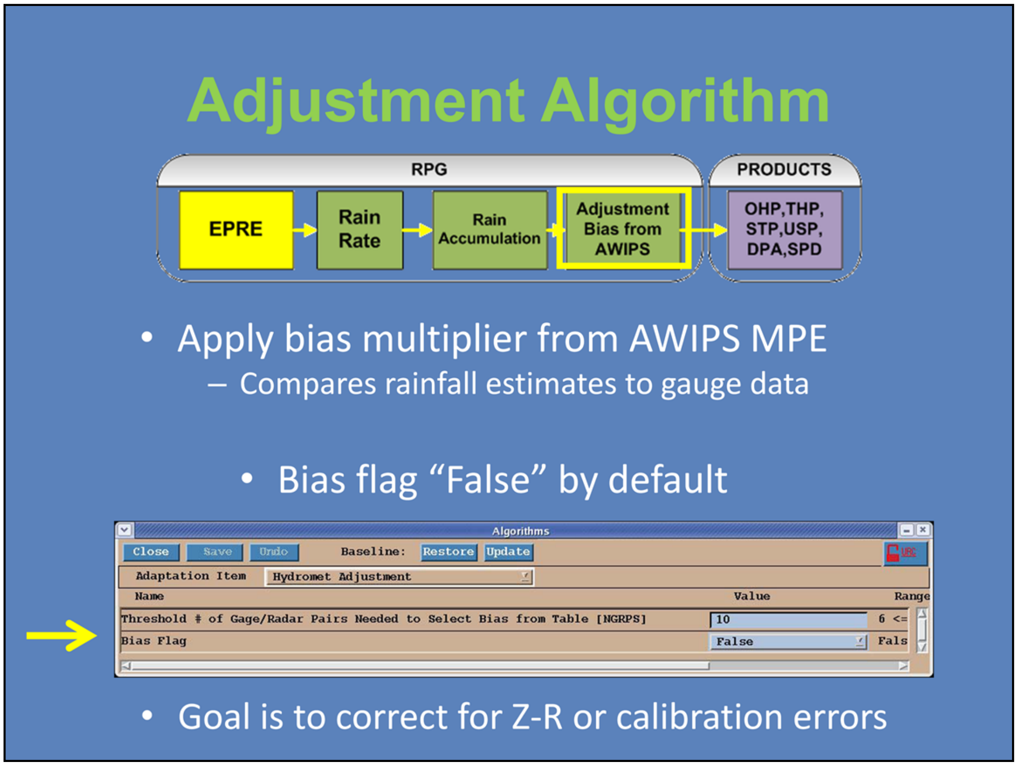

The adjustment algorithm is the last of the PPS algorithms, providing the option of applying a bias multiplier to the rainfall accumulations. The AWIPS Multi-Sensor Precipitation Estimator (or MPE) compares radar rainfall estimates to gauge data and sends a bias table to the RPG once an hour. Applying 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.

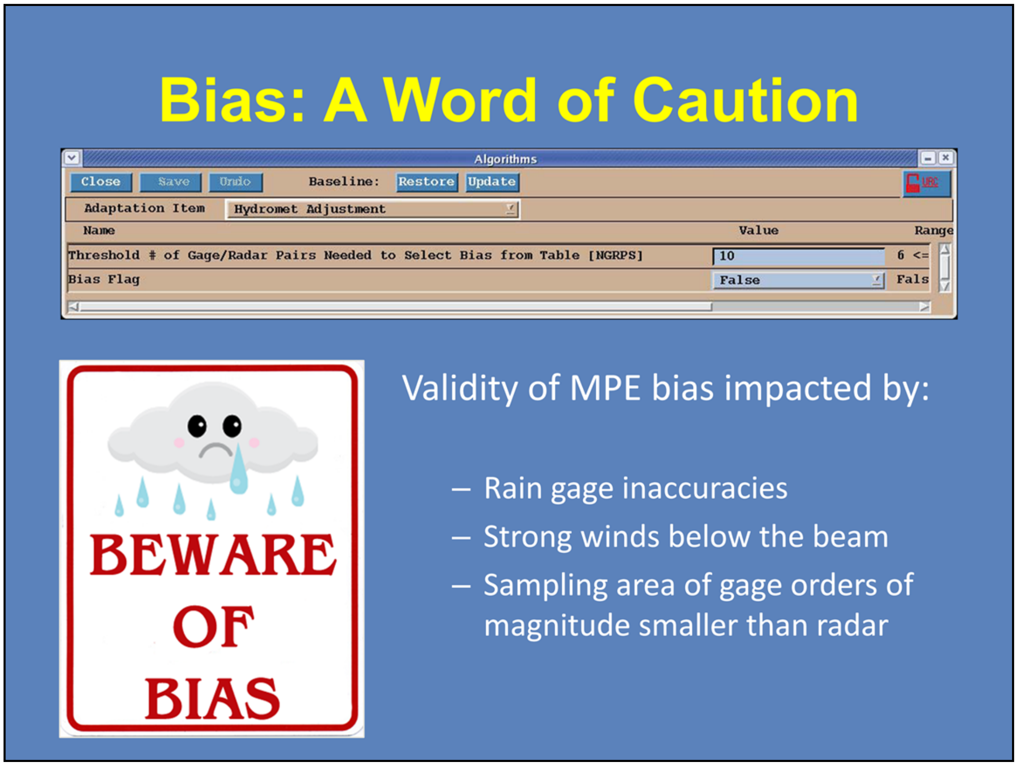

The MPE bias output isn't perfect and there are some scenarios to consider that could affect the validity of this bias. Rain gage values can be inaccurate for a variety of reasons. There may be strong winds below the beam. The most important consideration is that the rain gage sampling area is orders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

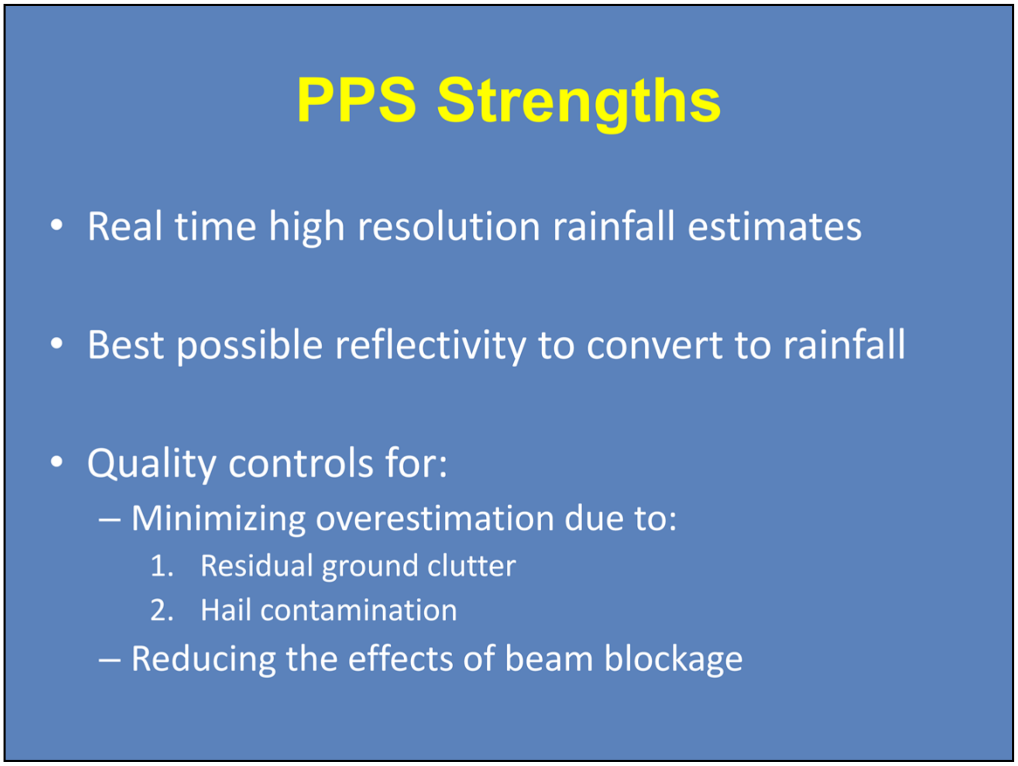

Here are the strengths of the PPS algorithm. This is the only source of real time high resolution rainfall estimates. It is important to remember that the qualitative spatial information of the rainfall pattern can be valuable.

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.

PPS Limitations

- Cannot account for
  - Below beam effects
  - Non-uniform dropsize distributions
    - Uniform Z-R relationship applied over domain
- Estimates may still be contaminated by
  - Bright band
  - Hail

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.

Back to Top

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

Back to Top

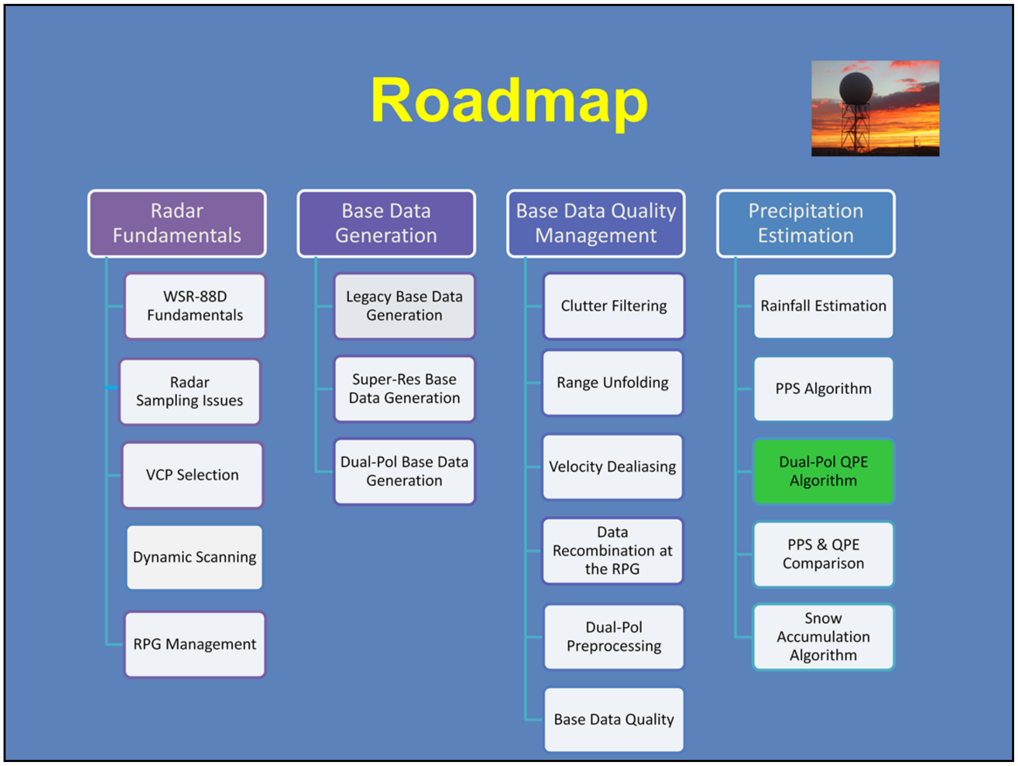

Here is the “roadmap” with your current location.

Learning Objectives

- 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.

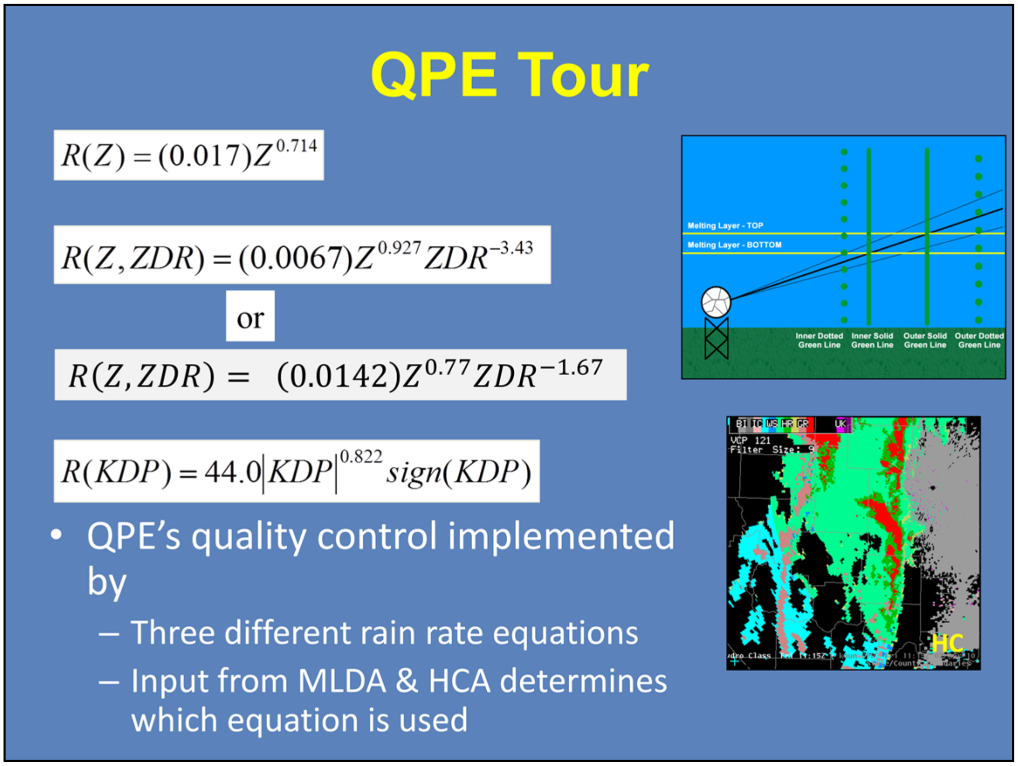

The QPE Tour will be somewhat different from the PPS Tour. The Legacy Precipitation Processing Subsystem (PPS) was presented as a sequence of quality control steps. The QPE Tour will be less sequential, but focused on key design elements of the QPE. This QPE Tour also provides more information on how 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.

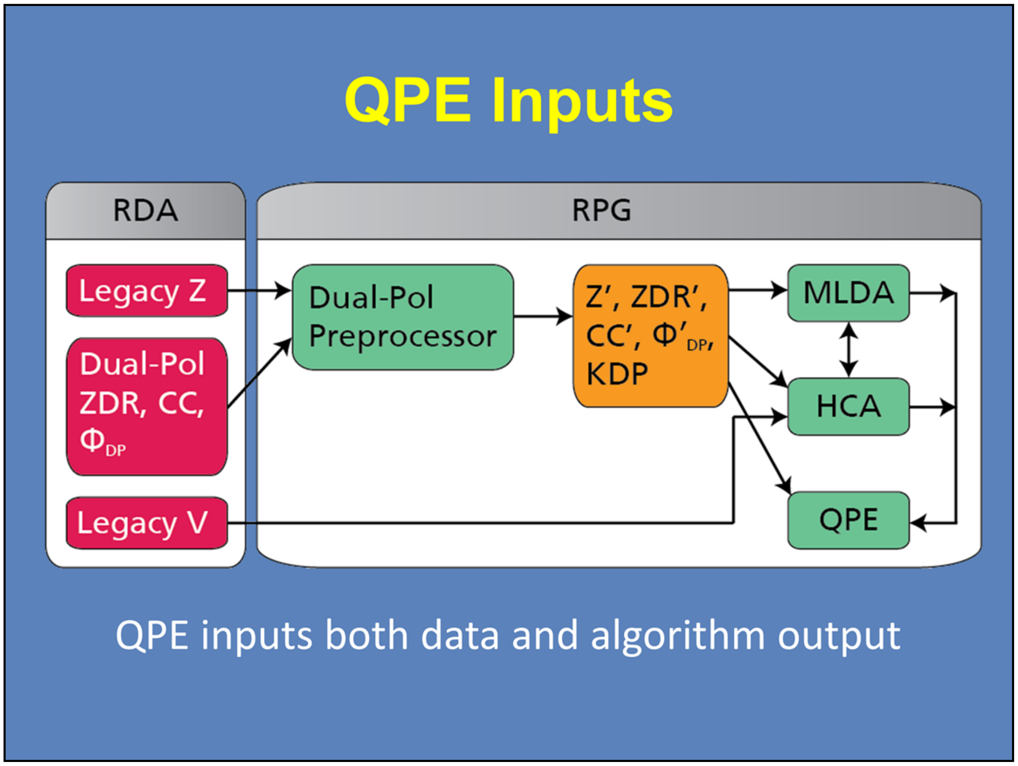

The Dual-Pol Quantitative Precipitation Estimation (QPE) algorithm is the most complex RPG algorithm yet implemented. QPE relies on multiple base data inputs, as well as the output from two new RPG algorithms. These algorithms are the Melting Layer Detection Algorithm (MLDA) and the Hydrometeor Classification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

Back to Top

QPE Tour: Hydroclass (HCA)

- QPE uses HCA for:
  1. Clutter & bio identification
  2. Estimated type of hydrometeor
- QPE's QC implemented by:
  – HHC built by QPE, based on HCA input

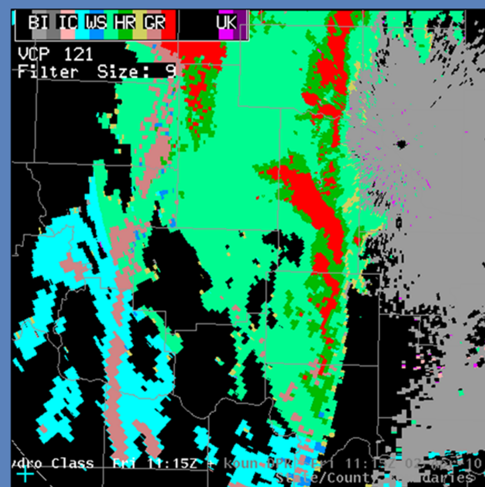

The QPE relies on the HCA in two ways. Bins that have clutter or biological targets identified by HCA are not converted to rainfall by the QPE. For bins that HCA identifies as having some type of precipitable hydrometeor, the QPE uses one of the different rain rate equations based on the type of hydrometeor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

Back to Top

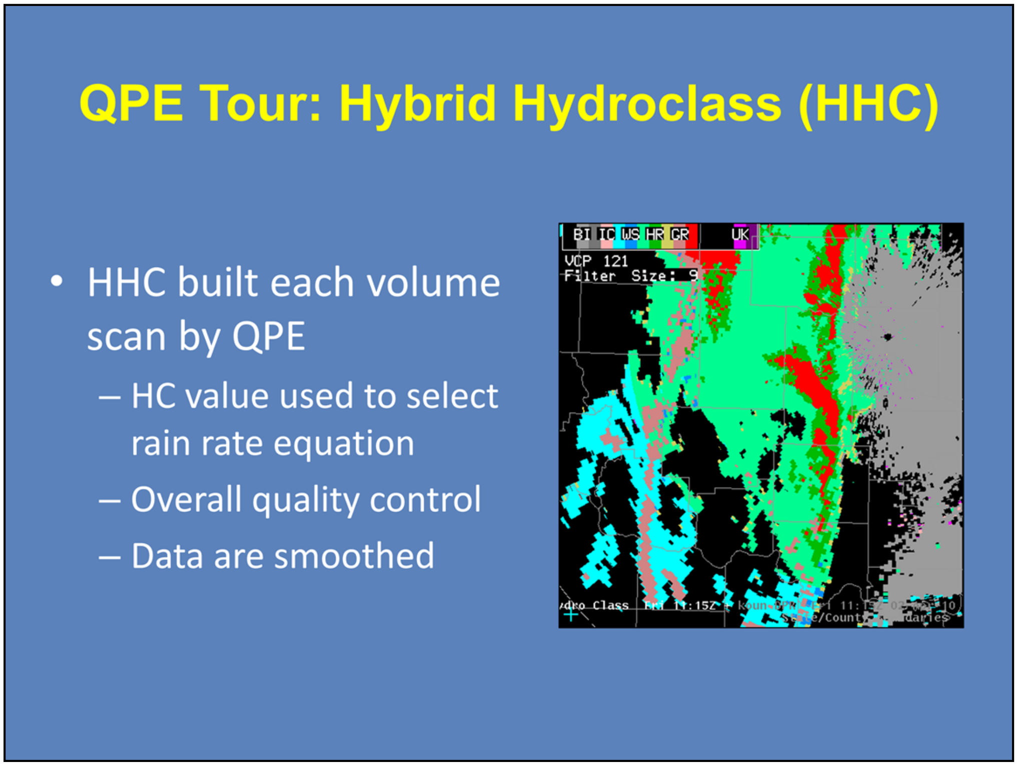

The Hybrid Hydroclass (HHC) product is built by the QPE and represents the hydrometeor classification values that were used to determine the rain rate equation applied, on a bin by bin basis, for each volume scan.

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.

QPE Tour: Melting Layer (MLDA)

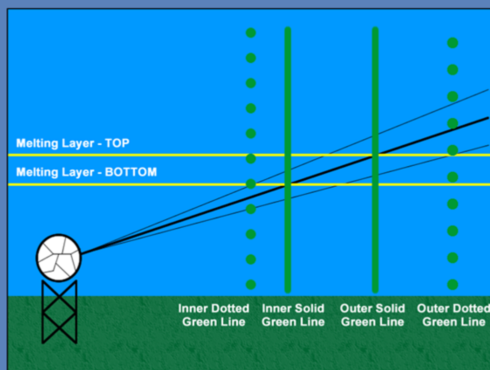

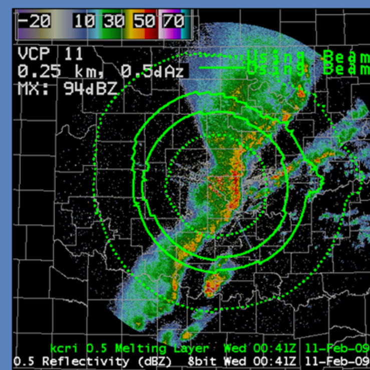

QPE's quality control implemented by Melting Layer Heights from MLDA

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.

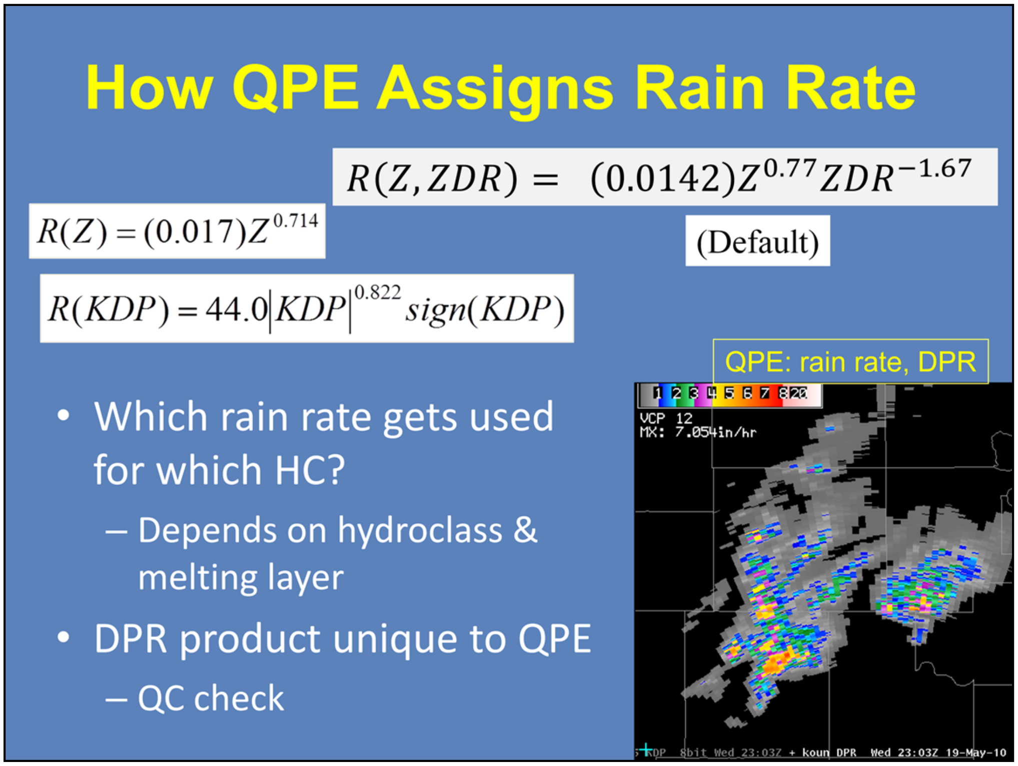

There are a number of steps involved in determining which rain rate equation is used, given the Hydroclass value and position of the range bin with respect to the melting layer. In some cases, a multiplier is used with the Z-R equation, such as 0.8* R(Z). Variations of the R(Z) are used in or above the 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.

Back to Top

How QPE Assigns Rain Rate

Classifications & Conditions	Equation
No Echo (ND) or Biological (BI)	0
Light/Moderate Rain (RA) or Big Drops (BD)	R(Z, ZDR)
Heavy Rain (HR) *and* Z ≤ 45 dBZ	R(Z, ZDR)
Heavy Rain (HR) *and* Z > 45 dBZ	R(KDP)
Rain/Hail (HA)	R(KDP)
Rain/Hail (HA) *and* echo is *at or below* top of ML	R(KDP)
Rain/Hail (HA) *and* echo is *above* top of ML	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* top of ML	R(Z)
Dry Snow (DS) *and* echo is *above* top of ML	2.8*R(Z)
Ice Crystals (IC)	2.8*R(Z)

Editable

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.

Back to Top

QPE Expectations

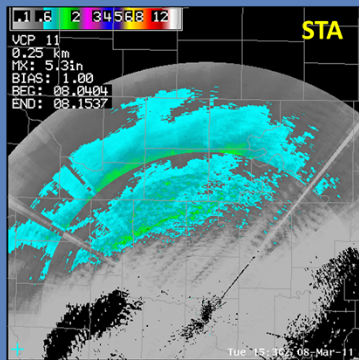

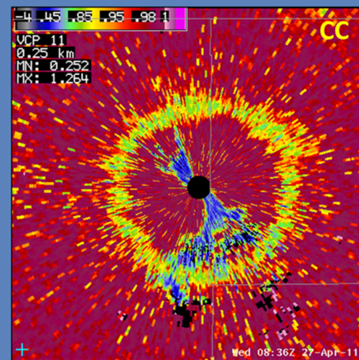

- Bright band contamination
- Above the melting layer
- Hail contamination
- Non-uniform beam filling

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

Back to Top

QPE & Bright Band Contamination

- Does HHC have Wet Snow (WS) within this band?
  - QPE uses 0.6R(Z)

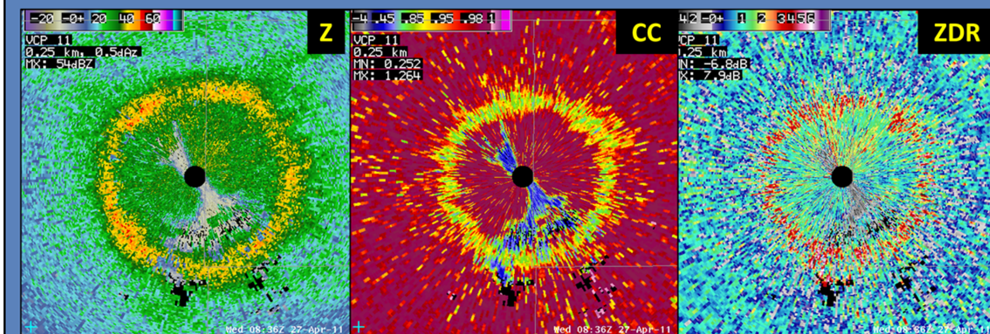

Bright band contamination has long been a challenge for using a radar to estimate rainfall. The QPE approach is to adjust the rain rate equation for water coated frozen hydrometeors such as wet snow that are typically located within a mesoscale melting layer (stratiform rain event). A multiplier of 0.6 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.

Back to Top

QPE Above the Melting Layer

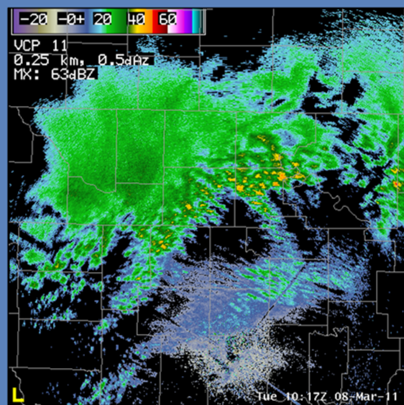

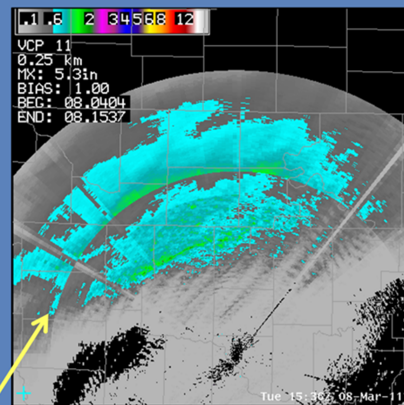

- R(Z) and 2.8R(Z) used near melting layer (defined by MLDA)
  - Discontinuity on QPE products
- 2.8R(Z) editable given local research

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.

Back to Top

QPE & Hail Contamination

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

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

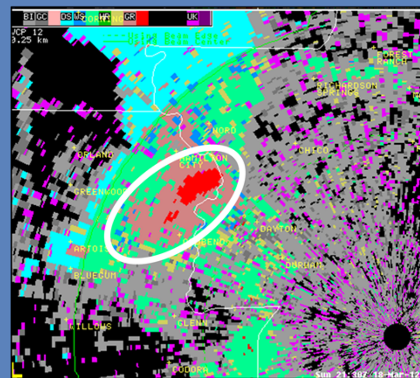

- Does HHC have Graupel (GR)?
  - Use 0.8R(Z)
- Does HHC have Hail Possibly Mixed with Rain (HA)?
  - Use R(KDP) if positive or 0.8R(Z)

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.

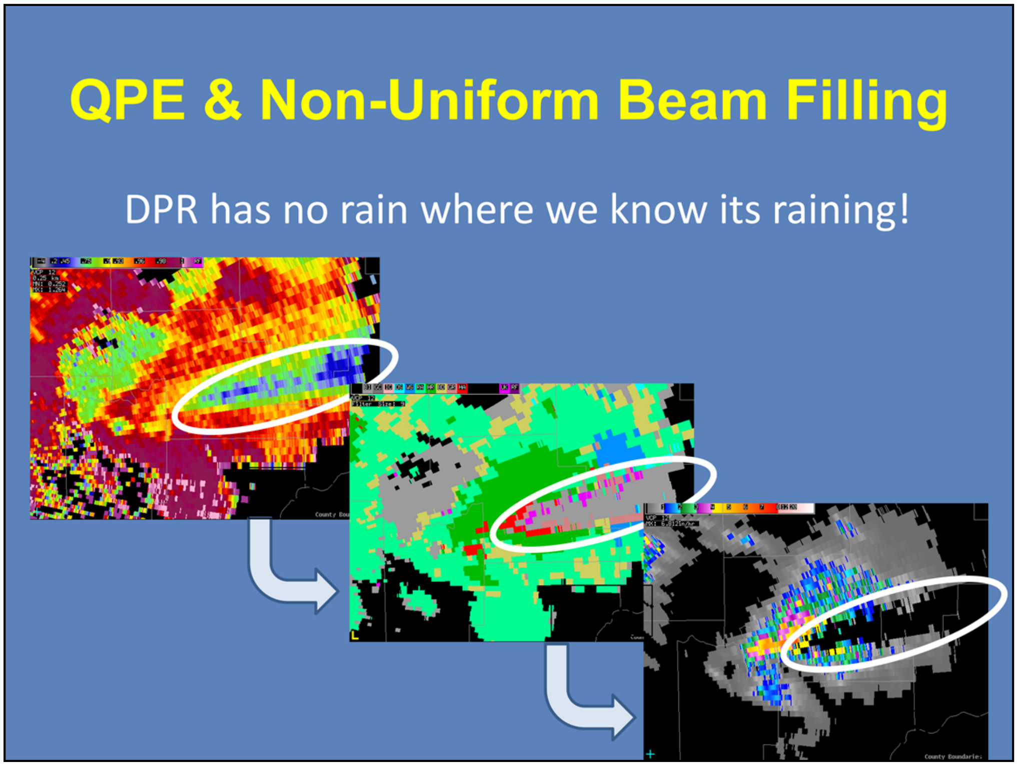

The artifact of a swath of low CC values due to non-uniform beam filling (NBF) can be either easy to spot or subtle. By comparing it to other radar data and understanding the environment, you can ask yourself if the CC values make sense.

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.

Back to Top

QPE & ZDR Calibration

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

- QPE designed with assumption of ZDR within ±0.1 Db
- Calibration sufficient for human interpretation of ZDR base product

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.

Back to Top

QPE "Ground Rules"

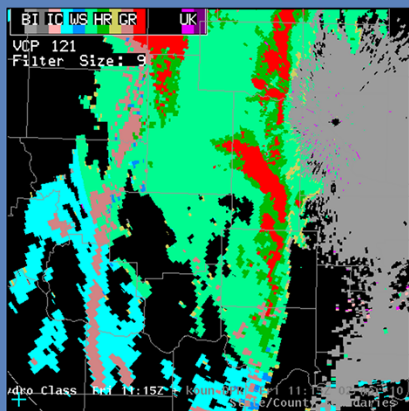

- Always check if ML and HHC makes sense
- Recognize situations when QPE performance typically impacted
  - Better performance below the melting layer

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.

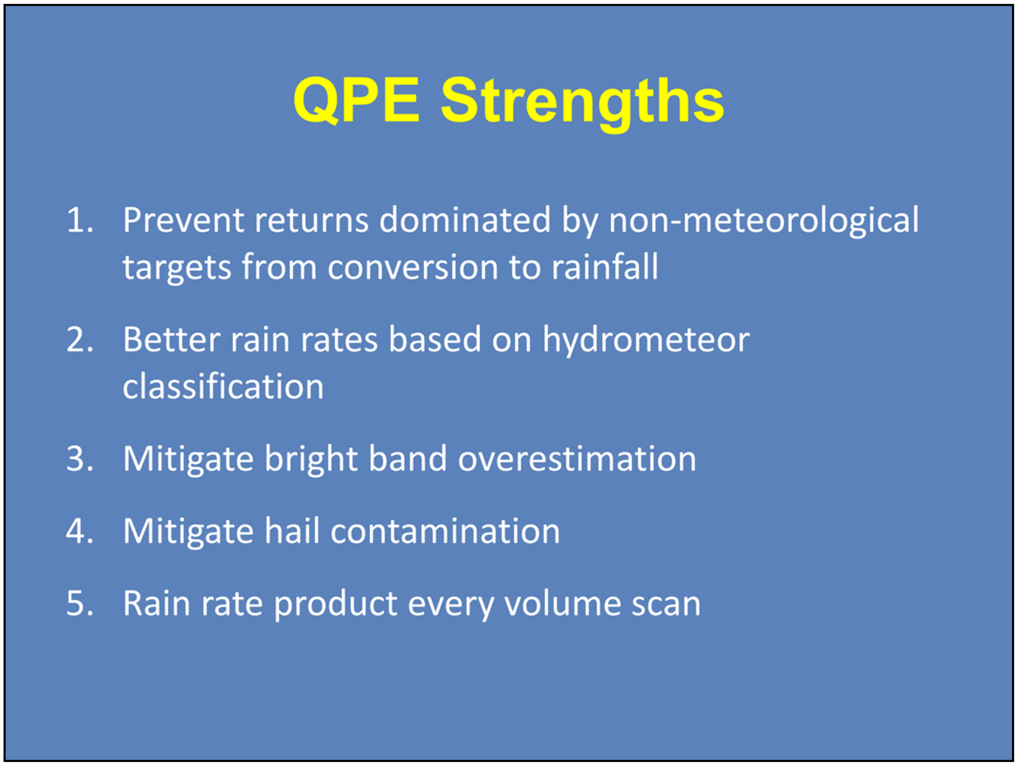

The strengths of QPE are mostly based on using the benefits of Dual-Pol to mitigate long standing challenges with using any radar to estimate rainfall.

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.

Back to Top

QPE Limitations

1. QPE "tuning" continues
2. Non-uniform beam filling unique impact on Dual-Pol base data
3. Sensitivity to ZDR calibration
4. Invalid hydroclass increases error
5. Discontinuity at top of melting layer
6. Reliance on MLDA and HCA
   – Must monitor these to ensure they "make sense"

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

Back to Top

Lesson Objectives

1. Identify design similarities between the PPS and QPE

2. Identify design differences between the PPS and QPE

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

Similarities: *Storm Total Start & Stop Accumulations*

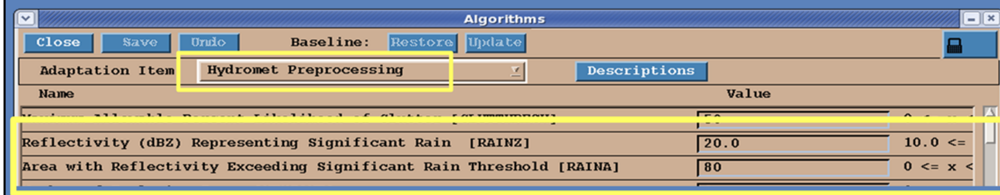

- Storm total start & stop based on coverage & intensity thresholds
- PPS uses
  - RAINA: areal coverage (default 80 km^2)
  - RAINZ: intensity (default 20 dBZ)

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”.

Back to Top

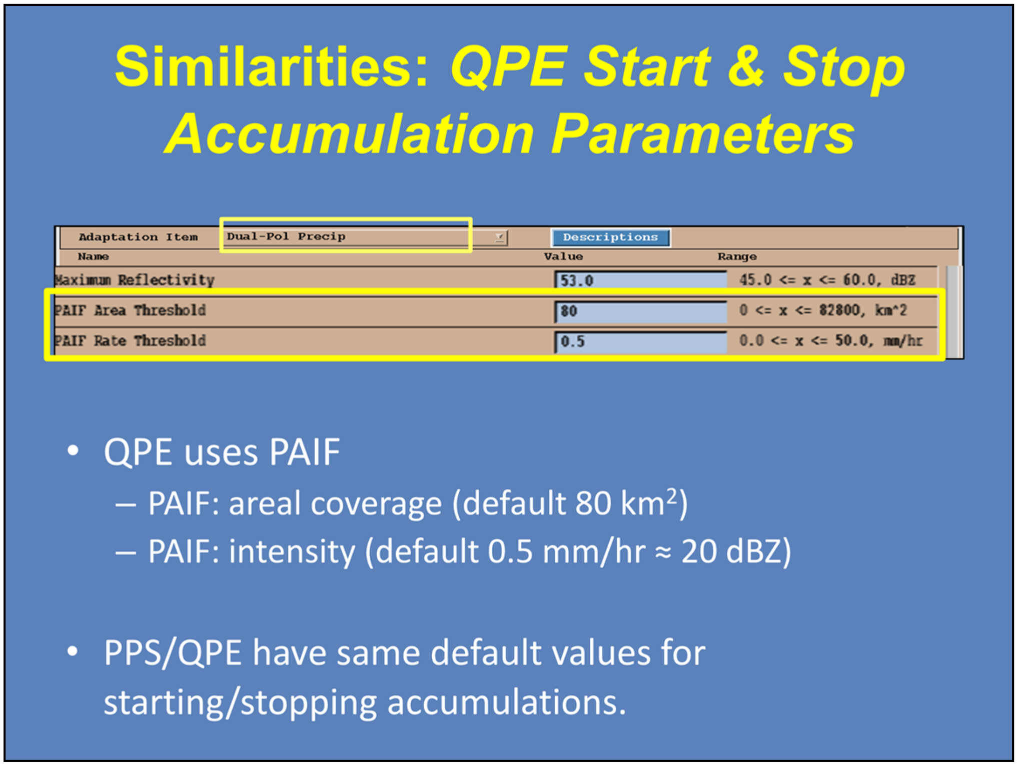

QPE uses the same concept for start and stop of storm total accumulations, it just uses different threshold names. The QPE thresholds start with the Precipitation Accumulation Initiation Function (PAIF). The PAIF Area Threshold is analogous to RAINA, while the PAIF Rate Threshold is analogous to RAINZ. 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.

Back to Top

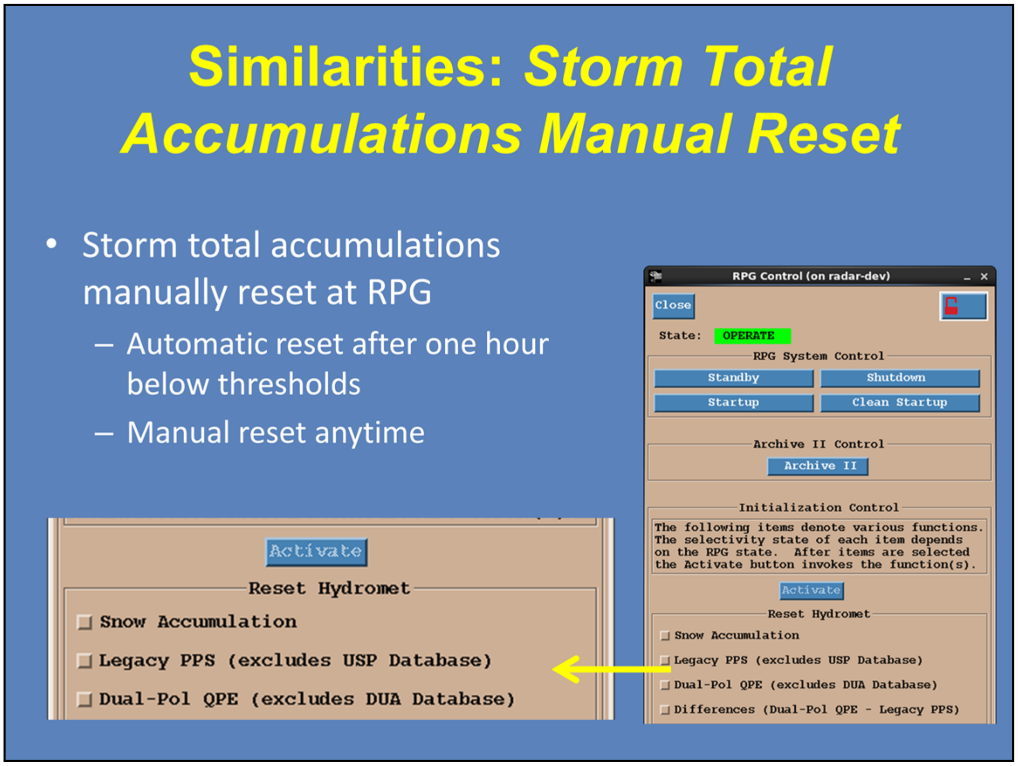

Both the PPS and the QPE have an automatic reset of the storm total accumulations after one hour of radar returns below their respective thresholds. Both the PPS and the QPE allow for a manual reset, which can be done from the RPG Control Window, selecting the “Legacy PPS” and the “Dual-Pol QPE”, respectively.

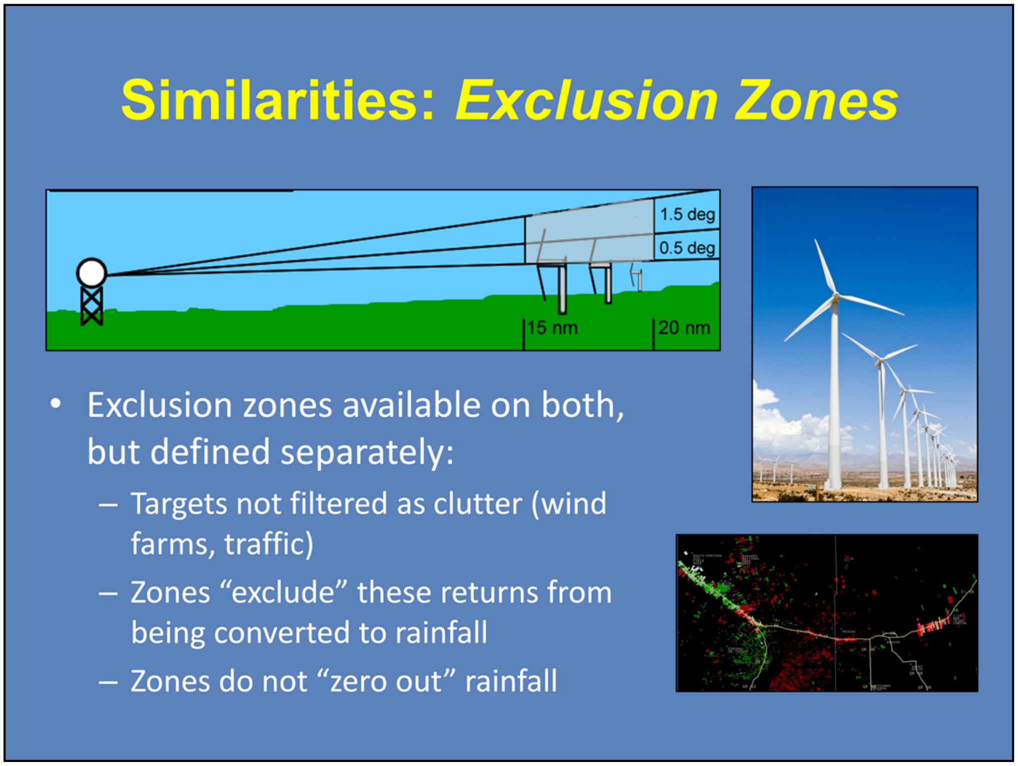

The next similarity between the PPS and the QPE is the use of exclusion zones. There are moving ground targets that are not filtered as clutter, such as rotating wind turbine blades and traffic on roads. Returns from targets like this usually exceed the intensity thresholds (RAINZ and PAIF Rate). Exclusion 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.

Back to Top

Similarities: *Exclusion Zones and HHC*

- Higher elevations used above exclusion zones are sometimes visible

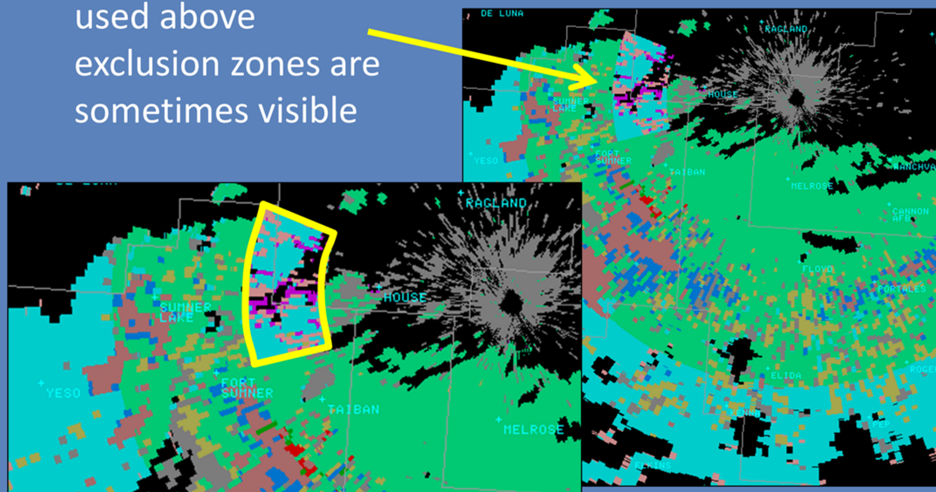

Here is an example where the use of exclusion zones for QPE becomes evident in the HHC product. Higher elevations are used by QPE for the azimuths and ranges within the exclusion zone, and that can sometimes be reflected in the hydroclass types on the HHC. In this case, you can see the transition from 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.

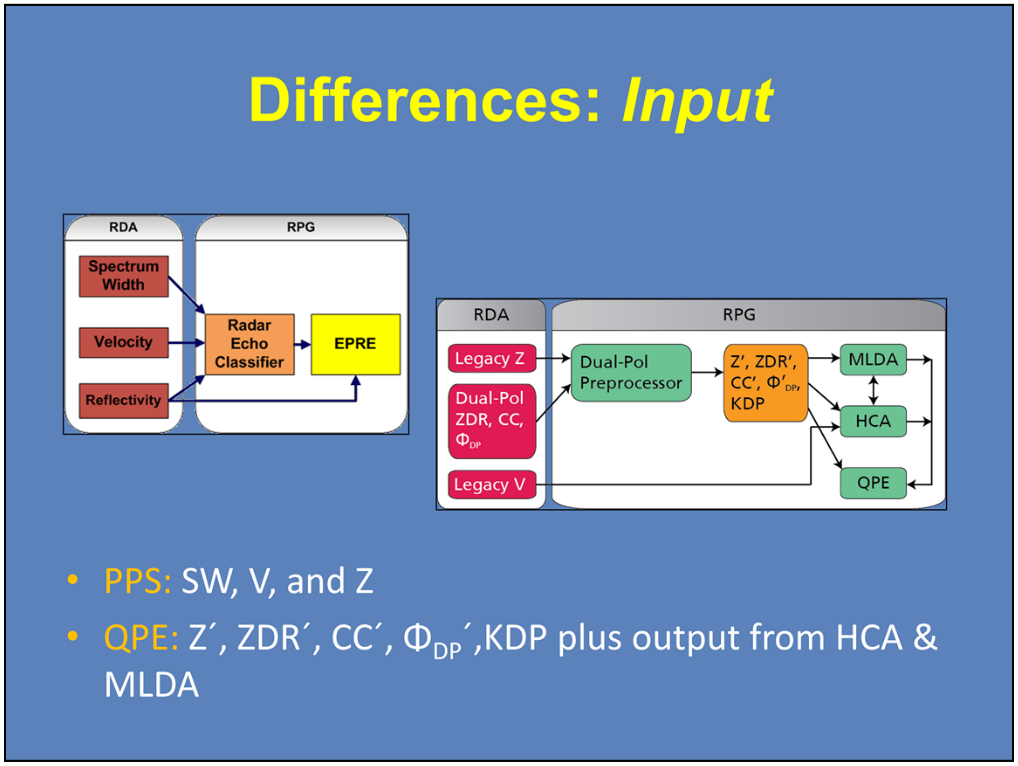

The Legacy PPS relies on the legacy base data: spectrum width, velocity and reflectivity. Though reflectivity is the only input for conversion to rainfall, spectrum width and velocity are used by the Radar Echo Classifier to identify bins that likely contain clutter.

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.

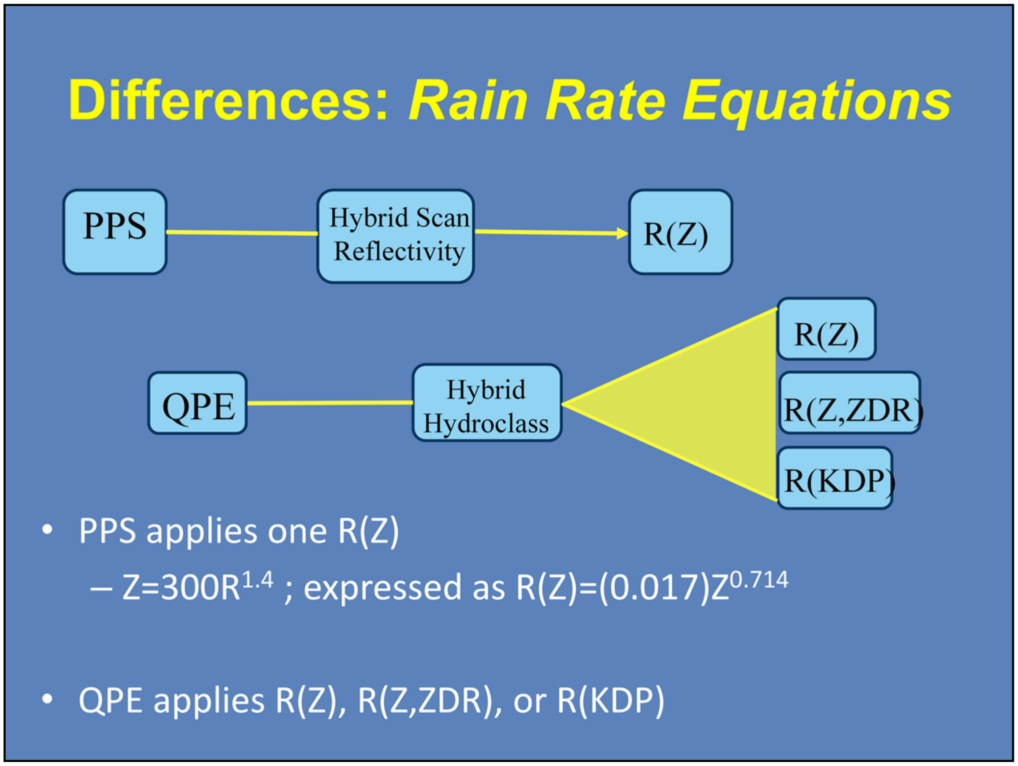

A key difference between the PPS and the QPE is the approach for converting base data to rainfall rate.

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.

Differences: *"Pre-Product Product"*

PPS: hybrid scan Z, DHR

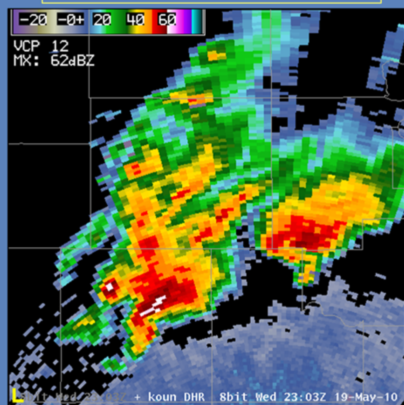

QPE: rain rate, DPR

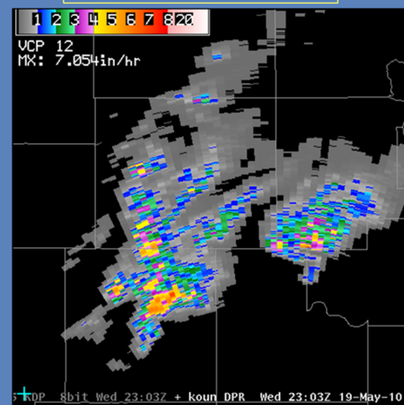

- PPS builds DHR, then converts to rate
- QPE builds DPR

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).

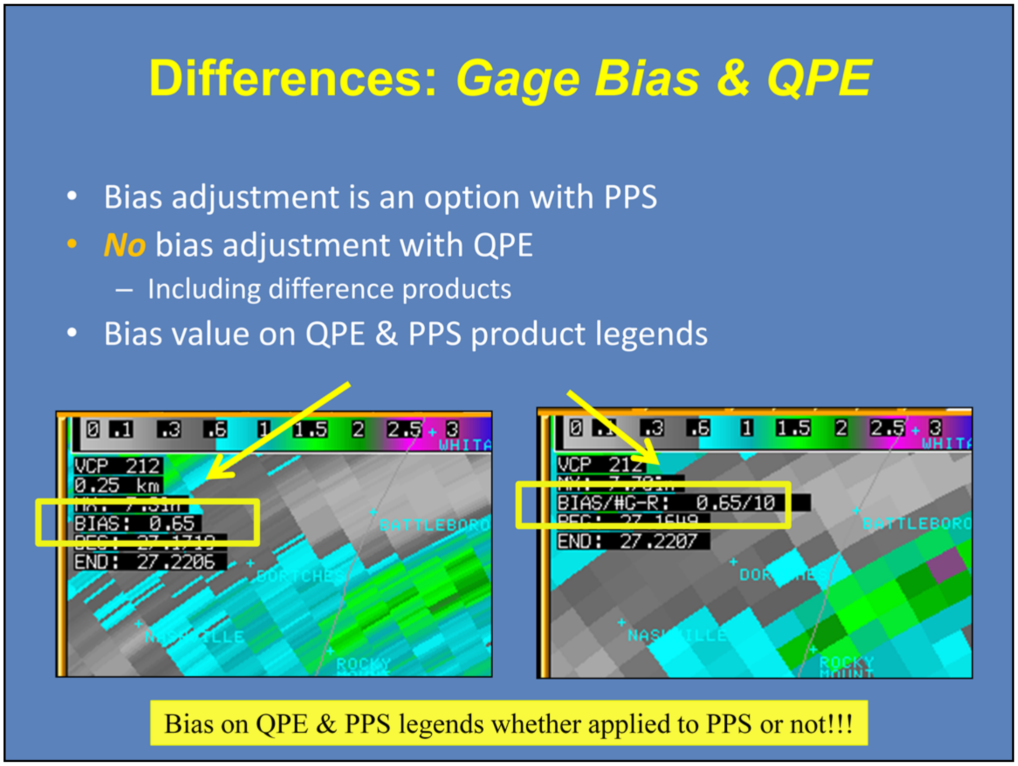

There is also a difference between the PPS and the QPE with respect to the application of a rain-gage bias. Bias adjustment is an option with the PPS, but there is no such option to apply a rain-gage bias to the QPE accumulations. All of the QPE-generated products are un-biased, including the QPE Difference 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.

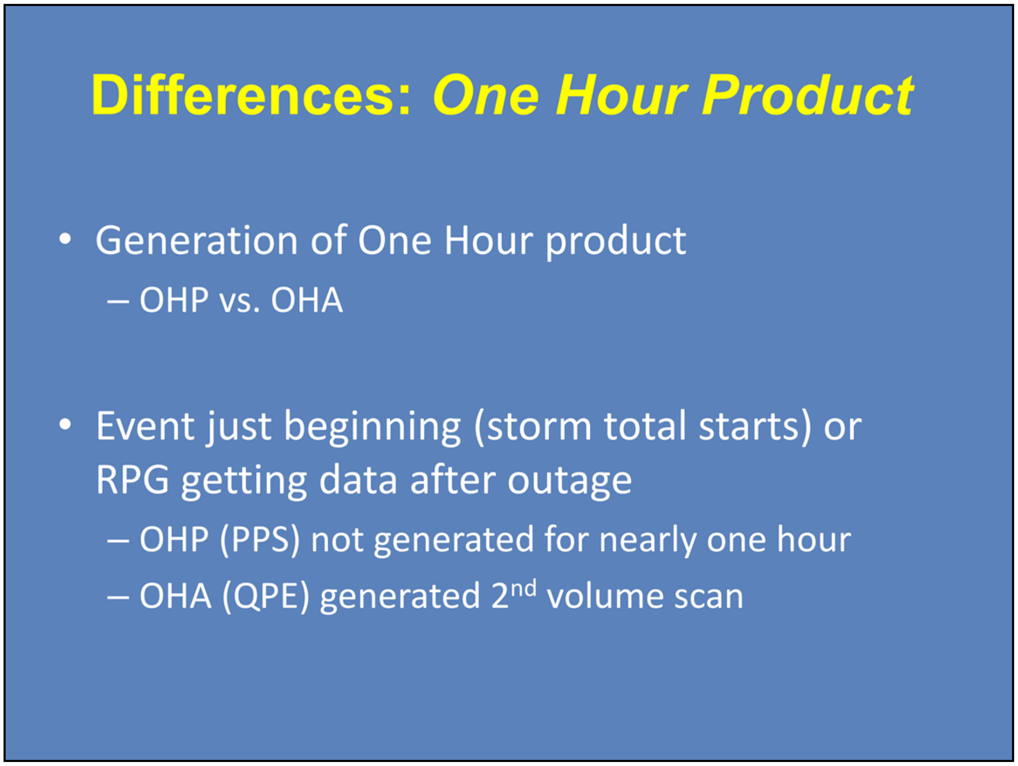

There is a significant difference between the PPS and QPE on the generation of their respective one hour products, the OHP and the OHA. This difference occurs for the beginning of an rainfall event or the return of base data to the RPG after an outage. The beginning of an event means that the storm total 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 2^{nd} full volume scan.

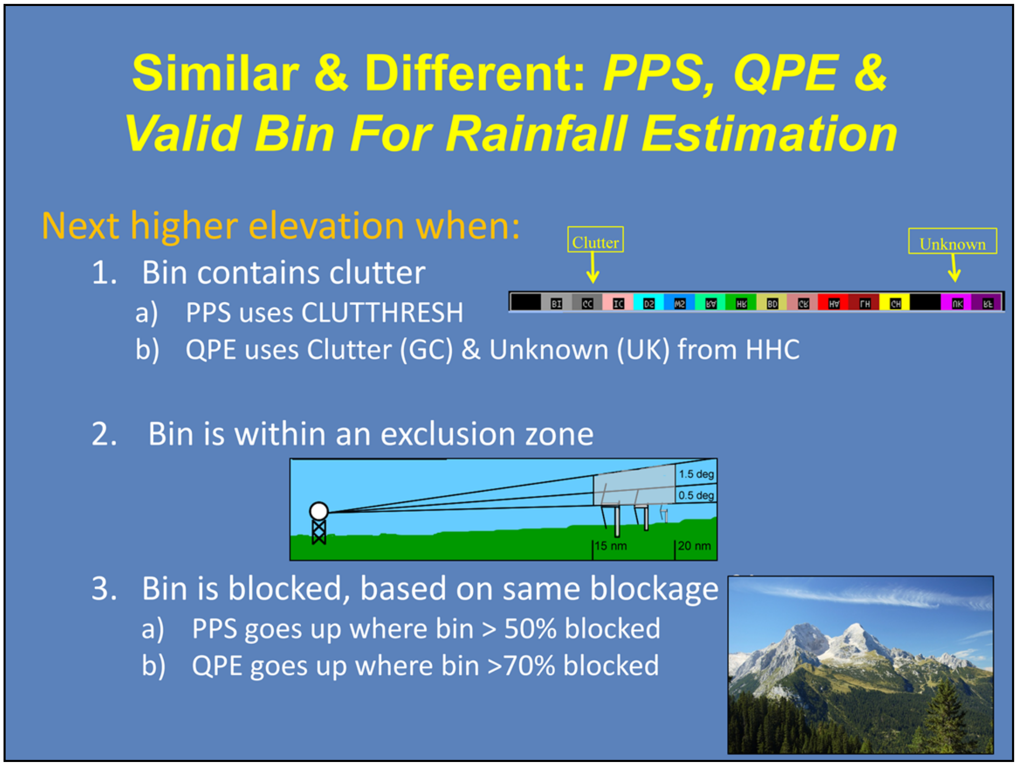

In a general way, the approach for selecting a bin to be converted to rainfall is similar for the PPS and the QPE. There is some difference in the implementation. There are three requirements that can cause the PPS or the QPE to use a higher elevation angle.

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.

Back to Top

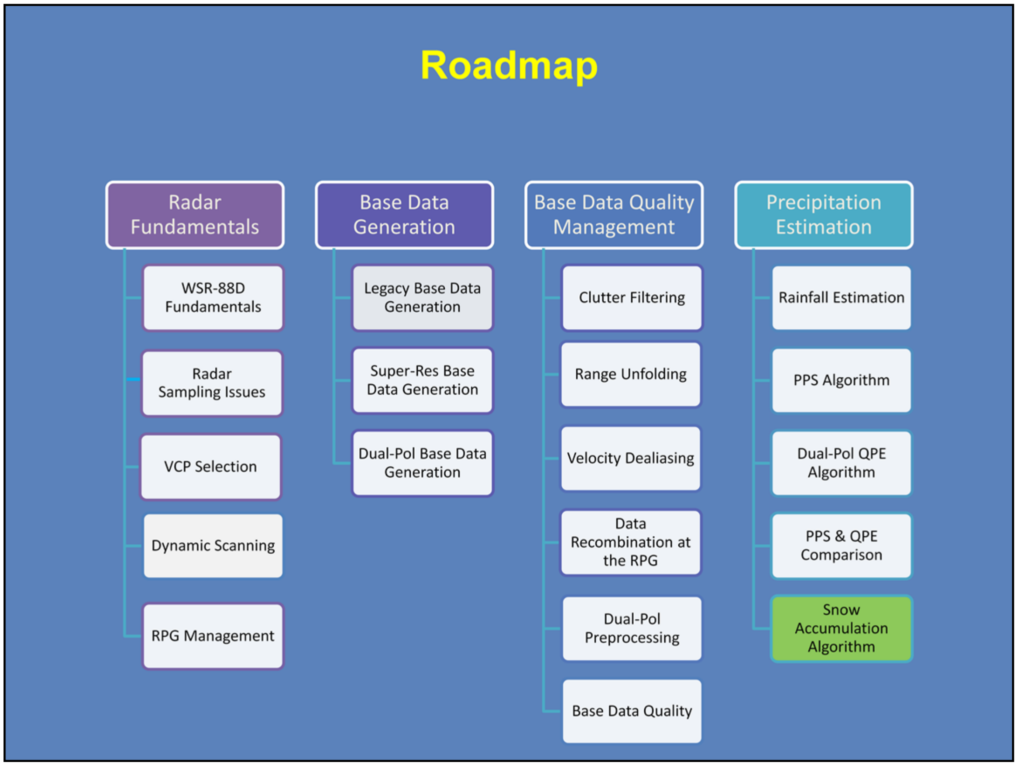

Here is the Roadmap with your current location.

Learning Objectives

1. Identify strengths & limitations of the Snow Accumulation Algorithm

There is one objective for this lesson.

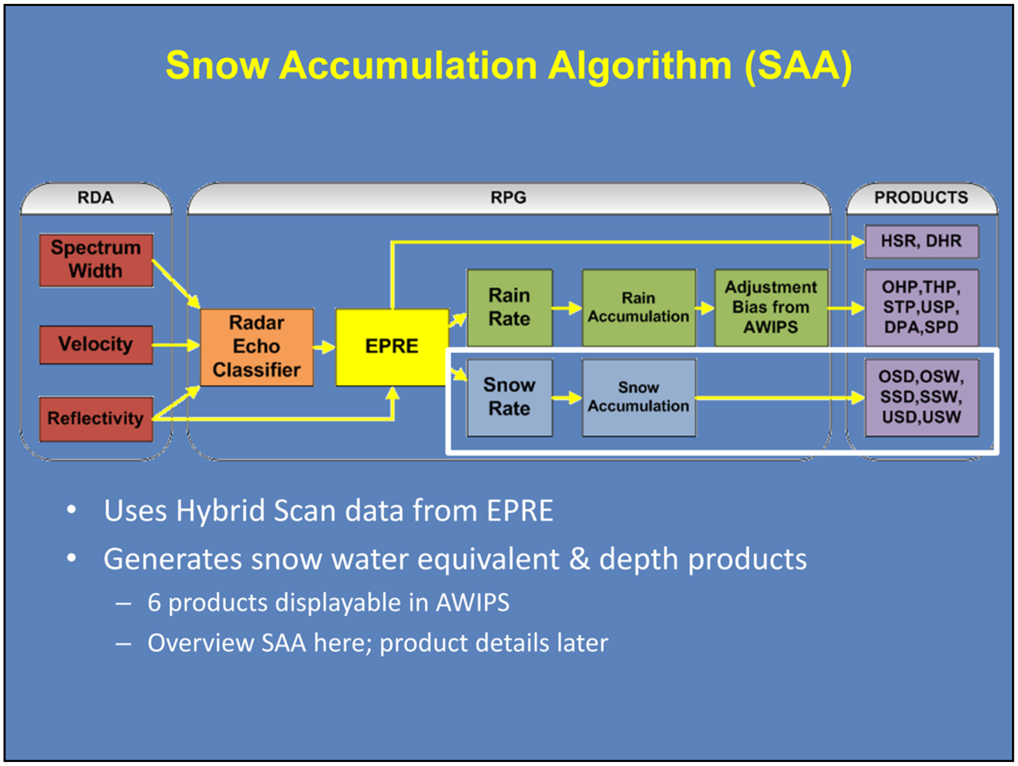

We now focus on the generation of snow water equivalent and snow depth products from the Snow Accumulation Algorithm (or SAA). The SAA uses hybrid scan data from the Enhance Preprocessing Algorithm (EPRE)

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.

Back to Top

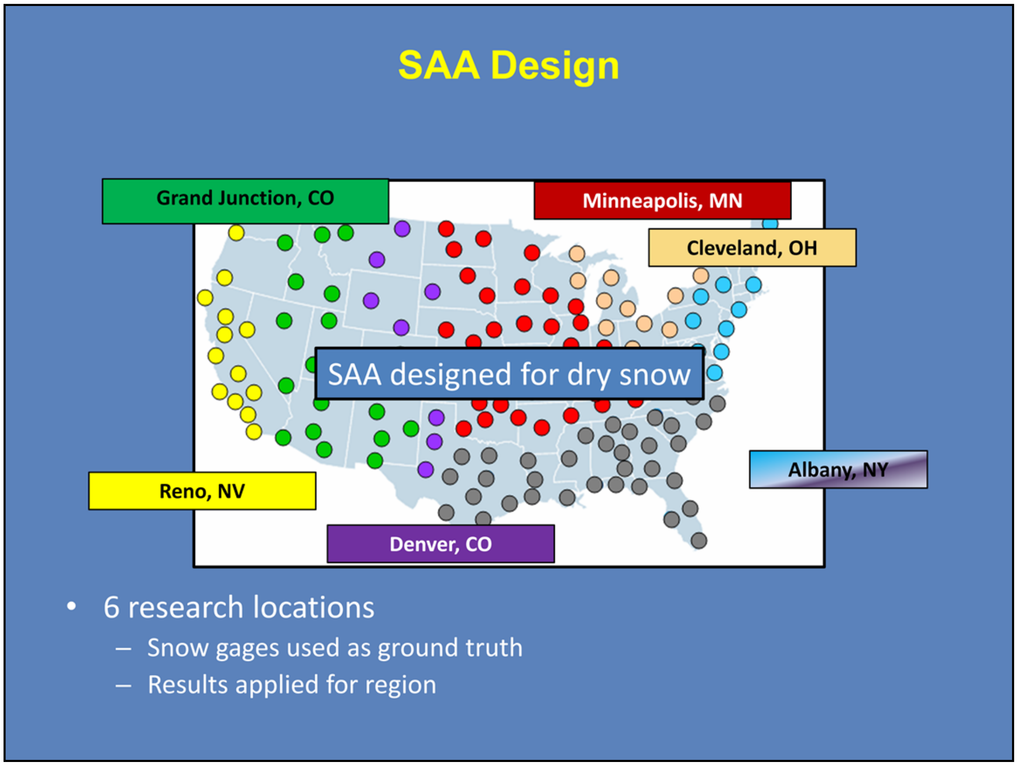

The SAA was developed from data collected at 6 different research locations. At each of these locations, a network of high quality snow gages was used as ground truth against the radar snowfall estimates. The output of this research provided default adaptable parameters that are used at each of the regions 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

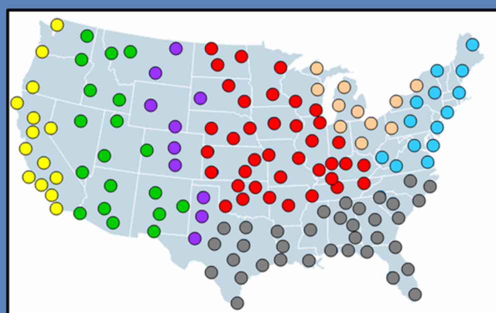

Research Location	Relationship
Albany, NY	$Z = 120\ S^2$
Cleveland, OH	$Z = 180\ S^2$
Minneapolis, MN	$Z = 180\ S^2$
Denver, CO	$Z = 130\ S^2$
Grand Junction, CO	$Z = 40\ S^2$
Reno, NV	$Z = 222\ S^2$

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

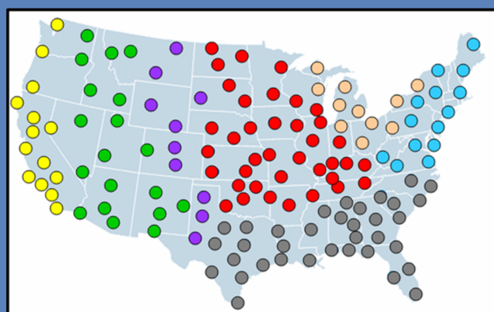

Research Location	Snow Ratio
Albany, NY	11.8
Cleveland, OH	16.7
Minneapolis, MN	11.8
Denver, CO	13.3
Grand Junction, CO	14.3
Reno, NV	8.0

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.

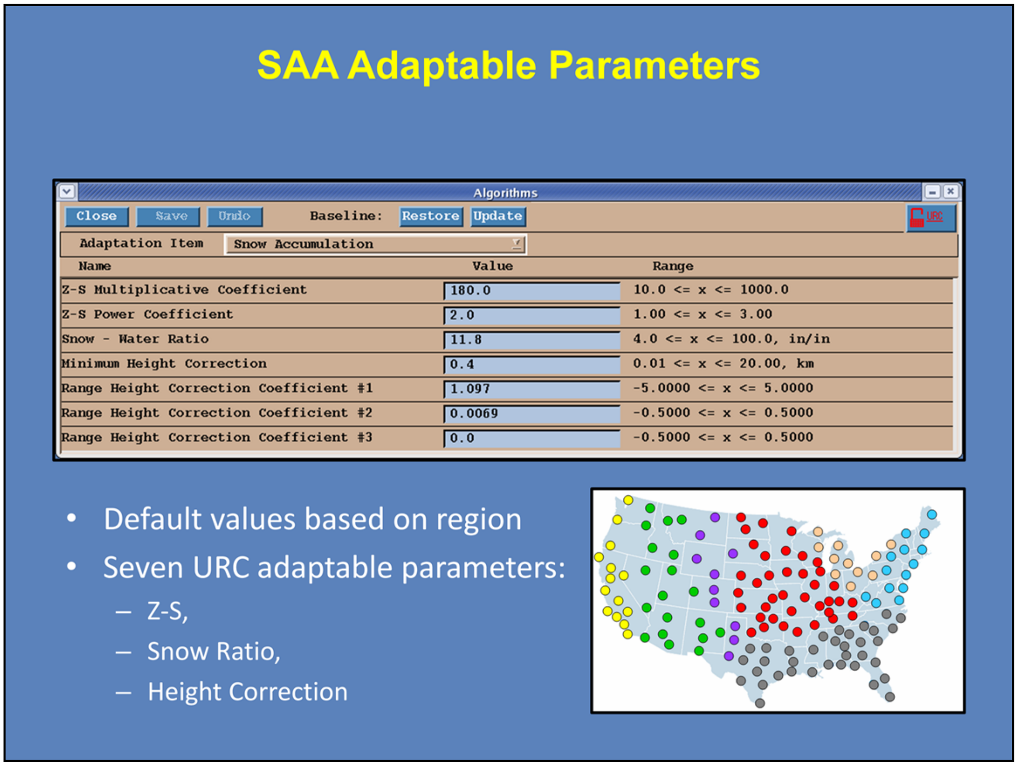

All the sites within a given region have the same default values for the SAA adaptable parameters. There are seven SAA adaptable parameters that are editable under URC guidelines in the RPG. Here you can edit the coefficients in the Z-S relationship, and modify the snow-water ratio.

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

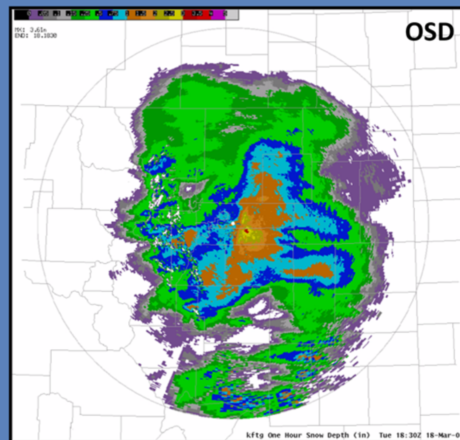

There are six snow products generated, all with 16 data levels, a resolution of 1km by 1 degree and a range of 124 nm. There are three durations: one hour, storm total and user selectable. For each of these durations, there is a snow water equivalent and a snow depth product. Examples of the products will be shown with more information in a later lesson.

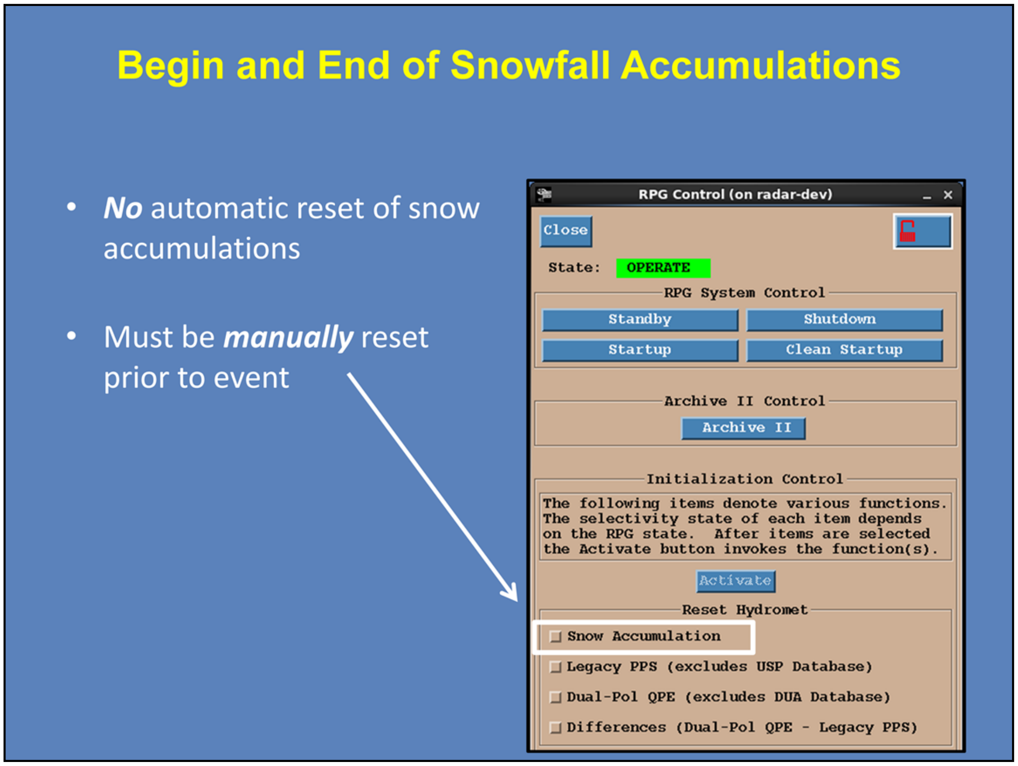

The SAA is designed to be event driven, and there is no automatic reset of the accumulations. This means the snow accumulations must be reset to zero at the beginning of an event. Resetting is done at the RPG Control window, just as with a reset of the PPS or QPE storm total accumulations.

Back to Top

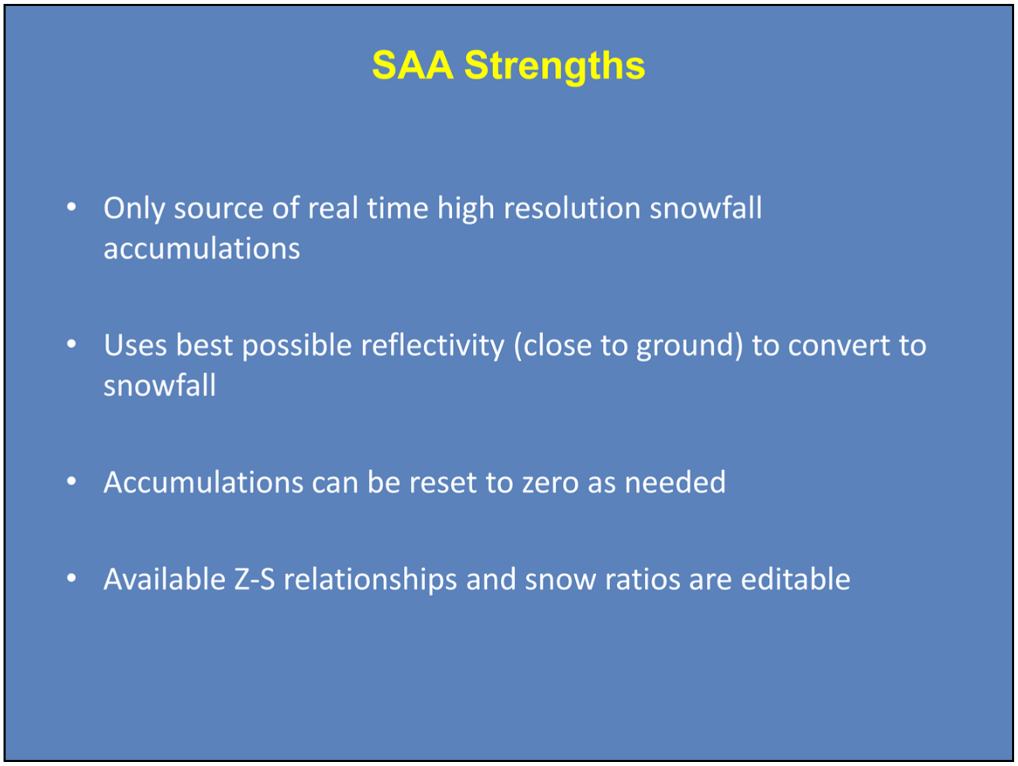

The first strength of the SAA is that it is the only source for real time high resolution snowfall accumulations.

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.

Back to Top

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.

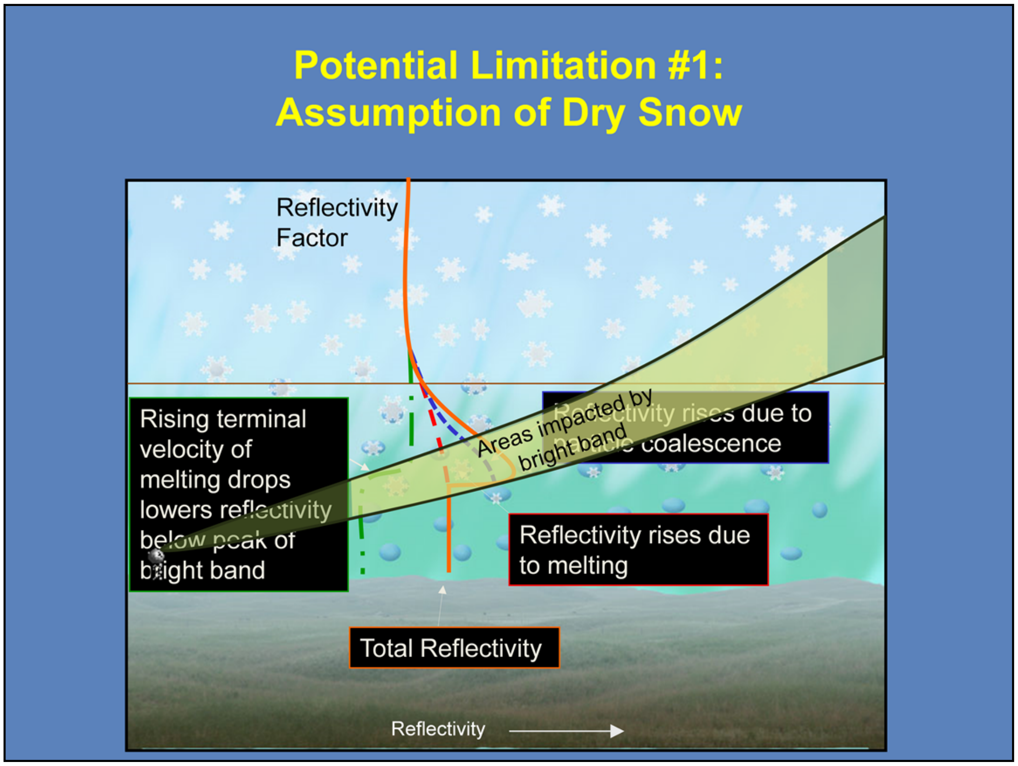

The Z-S equation in the Snow Accumulation Algorithm is based on the assumption that the precipitation at the surface is dry snow. We know this assumption isn't true in the bright band, and here is a closer look at what's going on.

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.

Back to Top

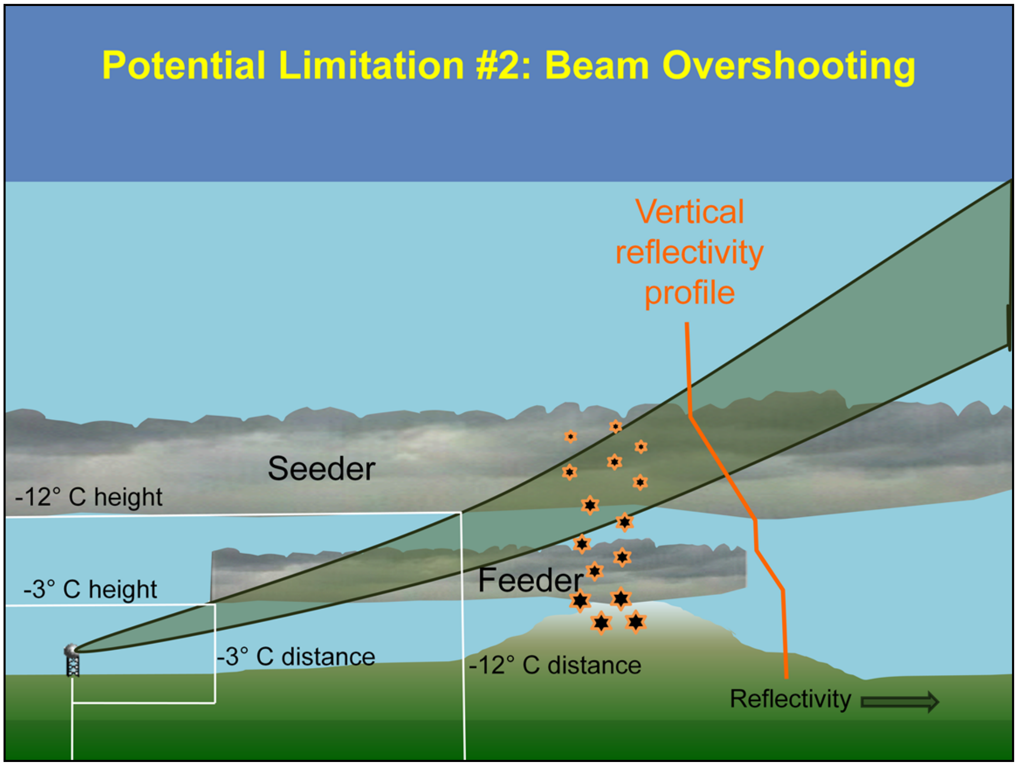

Assume we have a cloud pictured here actively generating precipitation so the intensity increases from top to bottom. Reflectivity begins to degrade once the top of the radar beam is above this precipitation production layer, and the signal is lost once the bottom of the beam is above the cloud.

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.

Back to Top

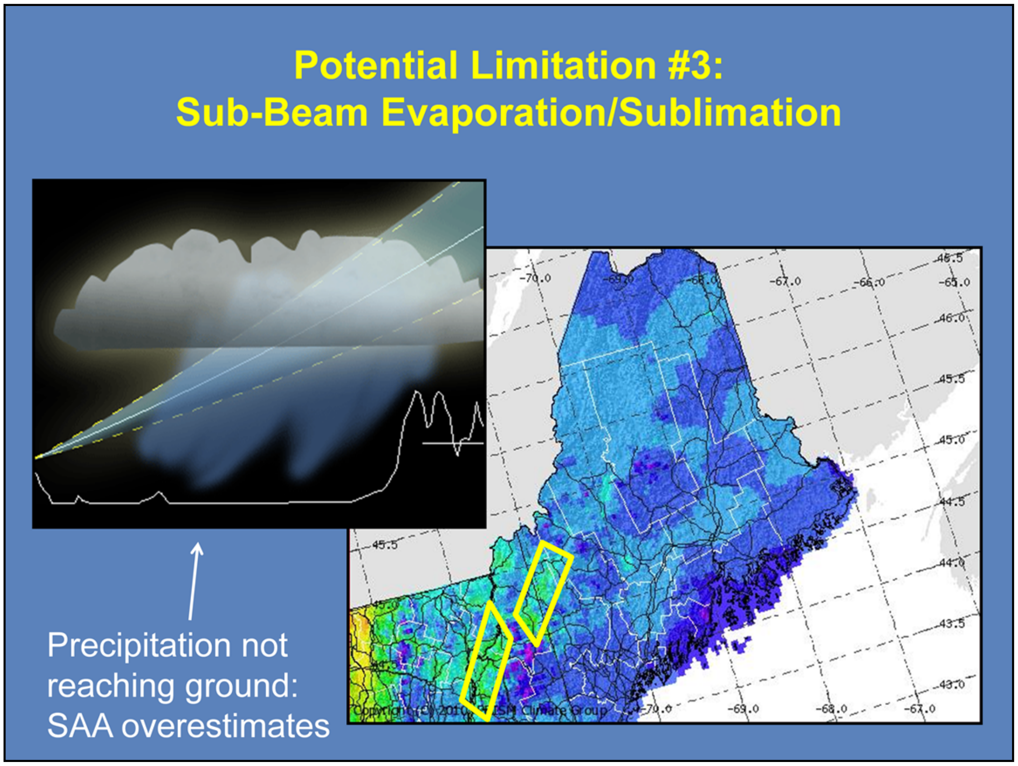

Another data quality issue to consider is sub-beam evaporation (or sublimation). When sub-beam evaporation/sublimation occurs, expect the Snow Accumulation Algorithm to overestimate the liquid equivalent precipitation at the surface. The problem is most common when the near-surface airmass is dry (such 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.

Back to Top

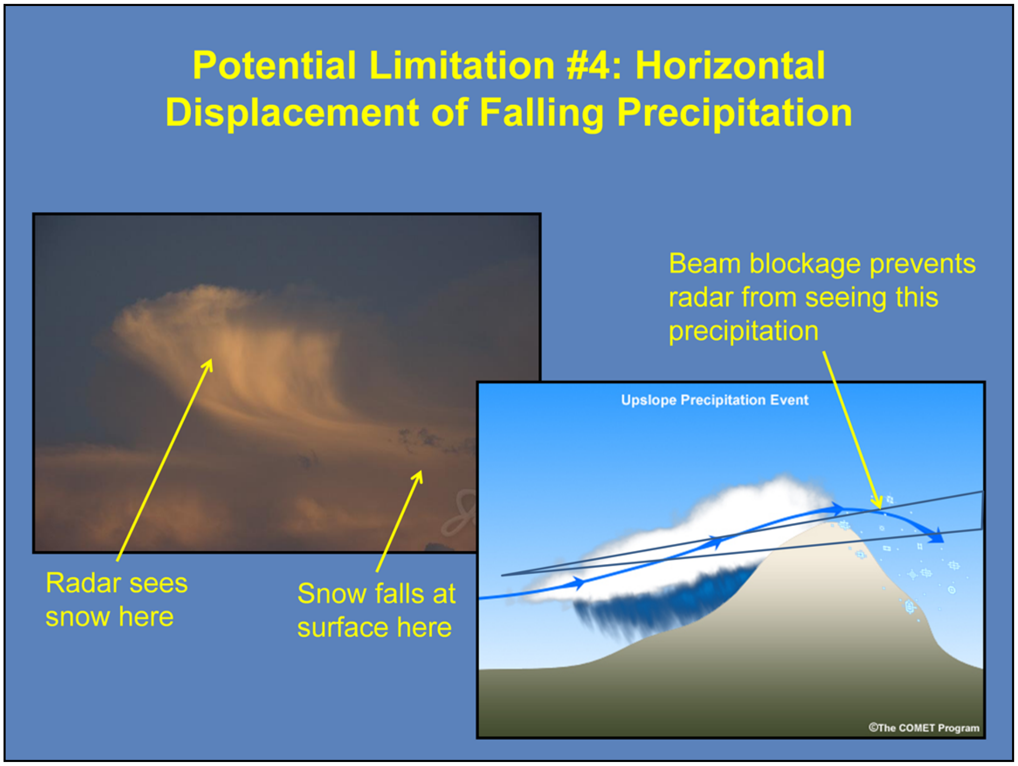

Under strong, sub-beam horizontal wind conditions (or areas of strong sub-beam vertical wind shear), precipitation may drift horizontally a long distance before reaching the ground. In some cases, the radar may not observe the precipitation that reaches the ground, such as when there is lee side spillover of orographic precipitation.

Back to Top

Potential Limitation #5: Unusual Precipitation Particle Shapes

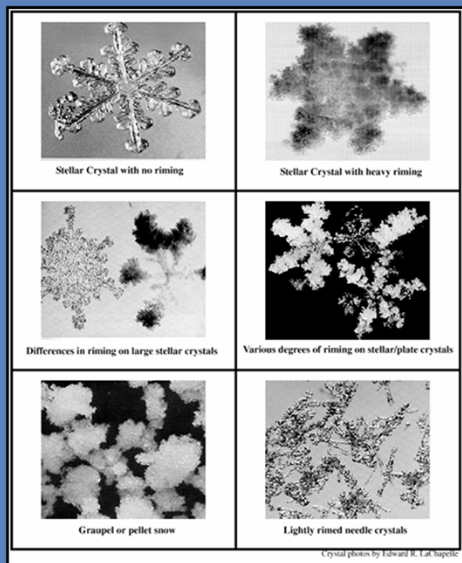

- Changes in shape, size can impact Z, LE differently
- Events can contain mix of stratiform, convective elements
- Variations can be difficult to detect, let alone predict
  - SAA doesn't account for these variations

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.

Back to Top

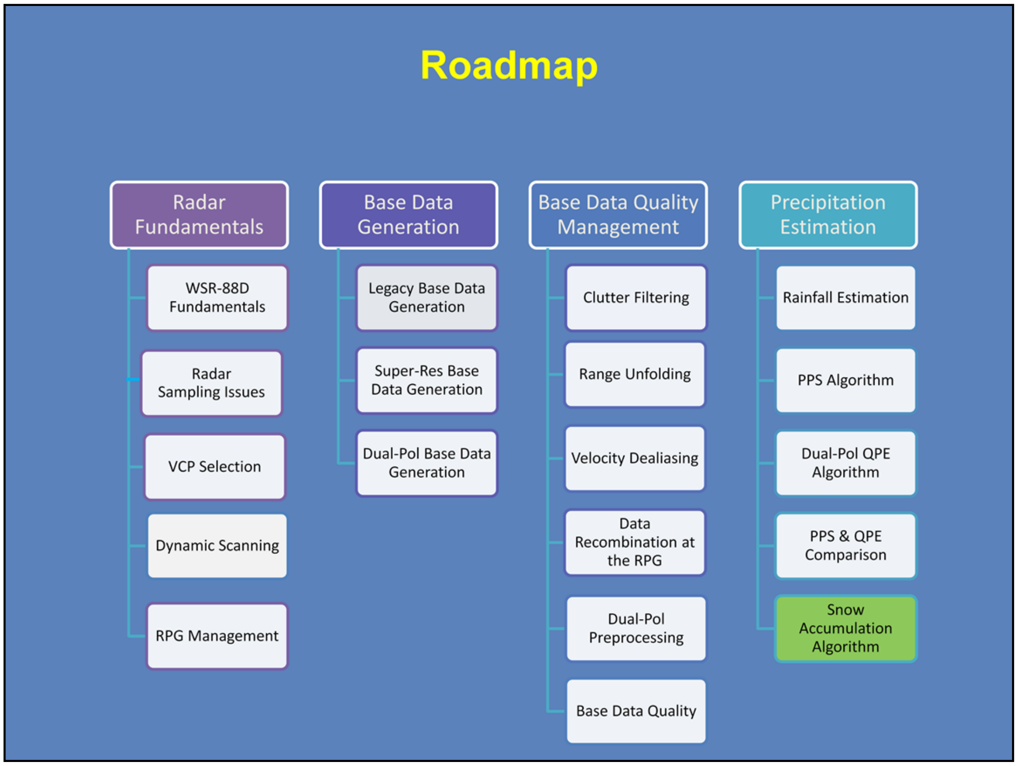

This concludes the lesson and here is the Roadmap with your current location.