Distance Learning Operations Course

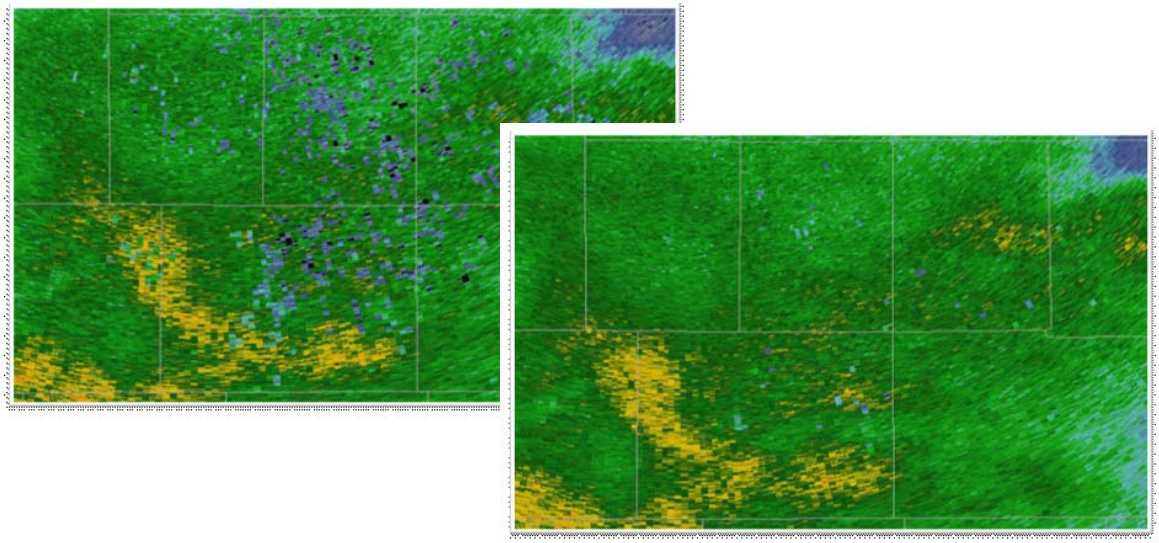

Topic 3: Principles of Meteorological Doppler Radar

Presented by the

Warning Decision Training Branch

Version: 1209

Table of Contents

Lesson 1: WSR-88D Fundamentals (self guided web module)..... 3 - 8
Objectives.....3 - 8

Lesson 2: Base Data Generation ..... 3 - 12
Objectives.....3 - 12
The Doppler Effect.....3 - 12
Sound Waves and Doppler Shift ..... 3 - 13
WSR-88D and Doppler Shift..... 3 - 14
Signal Phasor ..... 3 - 15
Phasor for a Single Pulse ..... 3 - 16
Phasors for Two Pulses..... 3 - 16
Maximum Unambiguous Velocity (Vmax).....3 - 18
Phase Shift and Radial Speed Relationship ..... 3 - 18
First Guess and Alias Velocities ..... 3 - 20
Super Resolution Signal Processing.....3 - 22
Base Reflectivity (Z) Generation .....3 - 26
Base Velocity (V) Generation.....3 - 27
Base Spectrum Width (SW) Generation.....3 - 28
Z Generated from PrH .....3 - 30
ZDR Generated from PrH and PrV.....3 - 31

Lesson 3: More Base Data Generation ..... 3 - 34
Objectives.....3 - 34
Differential Reflectivity (ZDR).....3 - 34
RDA Generation of ZDR..... 3 - 34
Dual-Pol on Split Cuts..... 3 - 35
Dual-Pol on Batch and Above ..... 3 - 36
Correlation Coefficient.....3 - 37
RDA Generation of CC ..... 3 - 37
Correlation Coefficient, ΦDP, and Cross Correlation Vectors..... 3 - 38
Correlation Coefficient is Based on ΦDP ..... 3 - 39
Why Does CC Matter? ..... 3 - 40
Correlation Coefficient and a Weak Returned Signal ..... 3 - 41
Correlation Coefficient and Spectrum Width are Similar but Different..... 3 - 42
Specific Differential Phase (KDP).....3 - 44
RDA and RPG Roles for KDP ..... 3 - 44
RDA Generation of ΦDP ..... 3 - 44
ΦDP Affected by Liquid Water..... 3 - 45
How ΦDP Changes Along the Radial..... 3 - 46
Why Specific Differential Phase (KDP)?..... 3 - 47
Specific Differential Phase (KDP) and Reflectivity..... 3 - 48

Specific Differential Phase (KDP) and Liquid Water ..... 3 - 49
Summary of the Dual-Pol Base Products ..... 3 - 49
Recombination at RPG ..... 3 - 50
Reflectivity Recombination ..... 3 - 51
Velocity and Spectrum Width Recombination ..... 3 - 52

Lesson 4: Filtering Clutter and Range Unfolding Velocity ..... 3 - 56
A Note on Terminology ..... 3 - 56
Objectives ..... 3 - 56
Ground Clutter Contamination ..... 3 - 56
Clutter Filtering Algorithms ..... 3 - 58
Clutter Mitigation Decision (CMD) ..... 3 - 58
The Clutter Filter Control (CFC) and Elevation Segments ..... 3 - 62
The Gaussian Model Adaptive Processing (GMAP) Algorithm ..... 3 - 66
CMD and GMAP Summary ..... 3 - 69
Range Unfolding of Velocity and Spectrum Width ..... 3 - 70
Range Folding Rarely on Reflectivity ..... 3 - 71
Legacy Range Unfolding Algorithm ..... 3 - 72
Non-overlaid Echoes Case ..... 3 - 72
Echoes are Overlaid ..... 3 - 76
Distribution of RF (Purple) ..... 3 - 79
Legacy Range Unfolding Algorithm Strengths ..... 3 - 80
Legacy Range Unfolding Algorithm Limitations ..... 3 - 80
SZ-2 Range Unfolding Algorithm ..... 3 - 81
Legacy vs. SZ-2 ..... 3 - 82
SZ-2 Technique ..... 3 - 83

Lesson 5: Dealiasing Velocity and Pre-processing Dual-Pol ..... 3 - 88
Objectives ..... 3 - 88
Improperly Dealiased Velocities ..... 3 - 88
Product Characteristics ..... 3 - 89
Velocity Dealiasing Algorithm ..... 3 - 90
Environmental Winds Table ..... 3 - 91
Environmental Data Editor Window ..... 3 - 92
Velocity Dealiasing Algorithm Strengths ..... 3 - 95
Velocity Dealiasing Algorithm Limitations ..... 3 - 95
Velocity Dealiasing Algorithm Considerations ..... 3 - 96
Multiple PRF Dealiasing Algorithm (MPDA) ..... 3 - 96
VCP 121 ..... 3 - 97
Strengths of the Multiple PRF Dealiasing Algorithm, VCP 121 ..... 3 - 98
Limitations of the Multiple PRF Dealiasing Algorithm, VCP 121 ..... 3 - 98
Dual-Pol Preprocessor ..... 3 - 98
Preprocessing for Differential Reflectivity (ZDR) ..... 3 - 99

Preprocessing for Correlation Coefficient (CC) ..... 3 - 100
Preprocessing and Differential Phase (ΦDP) ..... 3 - 100
Dual-Pol Preprocessor Artifact ..... 3 - 101

Lesson 6: Operational Impacts of Base Data Quality ..... 3 - 106
Objectives ..... 3 - 106
CMD False Detections ..... 3 - 106
Range Folding Mitigation VCPs ..... 3 - 108
VCPs and Manual PRF Edits ..... 3 - 110
Auto PRF ..... 3 - 115
Auto PRF, MPDA, and SZ-2 ..... 3 - 116
Attenuation of Z ..... 3 - 116
Differential Attenuation of ZDR ..... 3 - 117
What's in the Beam? ..... 3 - 118
Uniform Beam Filling and Correlation Coefficient (CC) ..... 3 - 119
Non-uniform Beam Filling ..... 3 - 120
Differential Reflectivity (ZDR) and Depolarization ..... 3 - 126
WSR-88D Data Quality ..... 3 - 128
1. VCPs 12, 212, 121 ..... 3 - 128
2. CMD ..... 3 - 128
3. GMAP ..... 3 - 129
4. Super Resolution Processing ..... 3 - 129
Putting It Together: Stratiform Rain, VCP 12 ..... 3 - 130

Lesson 7: Precipitation Estimation ..... 3 - 132
Objectives ..... 3 - 132
Heavy Rainfall Detection vs. Rainfall Estimation ..... 3 - 132
Issues of Using Radar to Estimate Rainfall ..... 3 - 133
RPG Algorithms to Estimate Rainfall ..... 3 - 139
Rain Rate Equation Input: Z ..... 3 - 140
RainRate Equation Input: Z and ZDR ..... 3 - 143
RainRate Equation Input: KDP ..... 3 - 144
PPS Tour ..... 3 - 144
PPS Strengths ..... 3 - 152
PPS Limitations ..... 3 - 152

Lesson 8: More Precipitation Estimation ..... 3 - 154
Objectives ..... 3 - 154
QPE Tour ..... 3 - 154
QPE Inputs ..... 3 - 154
QPE Tour: Hydroclass ..... 3 - 155
QPE Tour: Melting Layer ..... 3 - 156
Importance of HHC ..... 3 - 157
How QPE Assigns Rain Rate ..... 3 - 157

Extra for the Experts: QPE Rain Rate Equations ..... 3 - 159
Extra for the Experts: How QPE Builds DPR ..... 3 - 160
QPE Expectations ..... 3 - 165
QPE and Bright Band Contamination ..... 3 - 165
QPE Above the Melting Layer ..... 3 - 166
QPE and Hail Contamination ..... 3 - 166
QPE and Non-Uniform Beam Filling ..... 3 - 167
QPE and Differential Reflectivity (ZDR) Calibration ..... 3 - 168
QPE Strengths ..... 3 - 168
QPE Limitations ..... 3 - 169
Similarities and Differences between PPS and QPE ..... 3 - 170
Similar: Storm Total Accumulations ..... 3 - 170
Similar: Storm Total Accumulations Manual Reset ..... 3 - 171
Similar: Exclusion Zones ..... 3 - 171
Difference: Input ..... 3 - 172
Difference: Rain Rate Equations ..... 3 - 173
Difference: One Hour Product ..... 3 - 174
Difference: “Pre-product Product” ..... 3 - 174
Similar and Different: PPS, QPE and Valid Bin for Rainfall Estimation ..... 3 - 175
Snow Accumulation Algorithm (SAA) ..... 3 - 176
Snow Accumulation Algorithm (SAA) Design ..... 3 - 176
Z-S Relationships ..... 3 - 177
Snow Accumulation Algorithm (SAA) and Snow Ratio ..... 3 - 178
Snow Accumulation Algorithm (SAA) Adaptable Parameters ..... 3 - 178
Snow Accumulation Algorithm (SAA) Products ..... 3 - 178
Begin and End of Snowfall Accumulations ..... 3 - 179
Snow Accumulation Algorithm (SAA) Strengths ..... 3 - 179
Snow Accumulation Algorithm (SAA) Limitations ..... 3 - 180

Distance Learning Operations Course

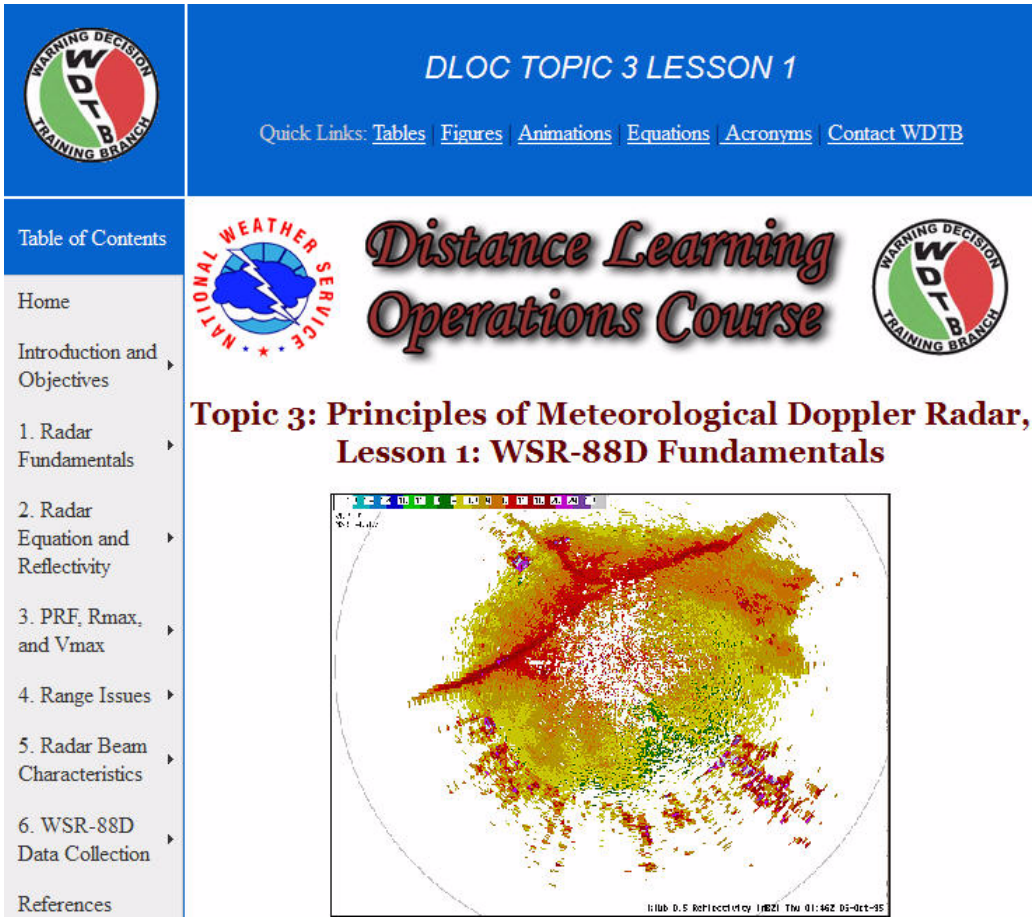

Topic 3, Lesson 1: WSR-88D Fundamentals

Presented by the
Warning Decision Training Branch

Version: 1210

Lesson 1: WSR-88D Fundamentals (self guided web module)

This lesson will present information that is fundamental to weather radars in general, such as beam propagation and the Probert-Jones radar equation. It will also present information that is very specific to the WSR-88D, such as the characteristics of a particular Volume Coverage Pattern (VCP).

Lesson 1 is unlike any other part of Topic 3, because it is a self guided web module. The remaining web modules, Lessons 2 through 8, are instructor guided, a sequence of annotated slides with accompanying audio for each slide. Lesson 1 has no audio, and it is designed as much for future reference as it is for initial learning. Lessons 7 and 8 of Topic 3 focus on the challenges of and techniques for estimating precipitation accumulation with the WSR-88D.

The amount of time needed to complete Lesson 1 will vary depending on your experience with the WSR-88D. The ***average*** completion time is estimated to be 2 hours.

Objectives

1. Identify how the variables below impact the size of the radar beam, the strength of the radar pulse, and the resulting values of Reflectivity:
   a. Range from the radar
   b. Transmitted power
   c. Drop-size distribution
2. Identify the resolution (i.e.range and azimuth) and precision (e.g. 3-bit, 4-bit, or 8-bit) values for WSR-88D base products.

3. Identify the difference between radar sensitivity and calibration.
4. Identify the definition and radar product impacts of attenuation, differential attenuation, and non-uniform beam filling.
5. Identify the definition of the Doppler dilemma, the definition of range folding, and the relationship pulse repetition frequency (PRF) has on calculations of target range and velocity.
6. Identify the two situations where sidelobe contamination is likely to occur.
7. Identify meteorological conditions conducive to superrefraction and subrefraction of the radar beam, and the resultant operational impacts.
8. Identify the azimuthal width that defines super resolution and legacy base data and the VCP waveform that is necessary to generate super resolution base data.

Distance Learning Operations Course

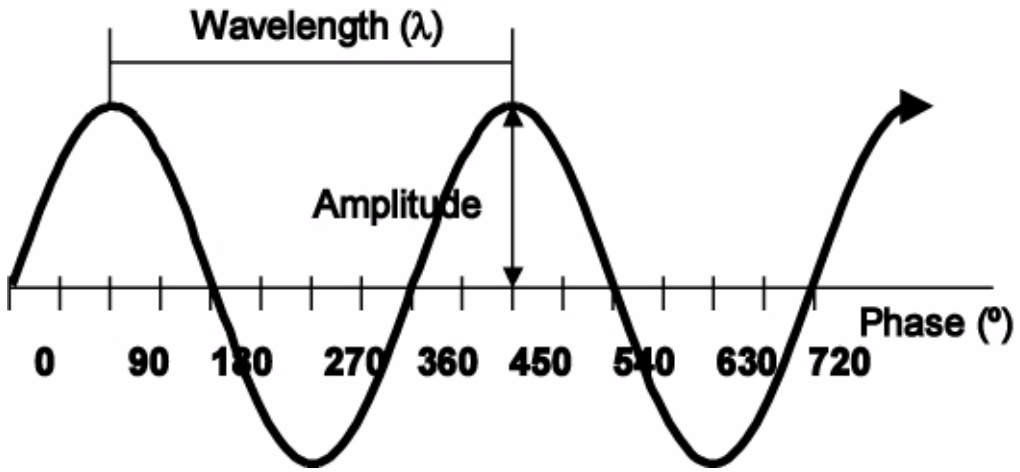

Topic 3, Lesson 2: Base Data Generation

Presented by the
Warning Decision Training Branch

Version: 1210

Lesson 2: Base Data Generation

Base data are the input for the generation of ***all*** radar products. It is therefore critical that the base data be of the highest possible quality. There are data quality issues inherent in meteorological Doppler weather radars. Proper interpretation of the WSR-88D products requires an understanding of data quality. Lessons 2 - 6 of Topic 3 provide the necessary foundation for this understanding.

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 the operational impacts of the signal processing techniques used to produce super resolution base data.
4. Identify how the returned signal is used to generate
   a. Reflectivity (Z)
   b. Radial Velocity (V)
   c. Spectrum Width (SW)
5. Identify the returned pulse characteristics used to calculate ZDR vs. Z.

The Doppler Effect

The Doppler Effect is defined as *the change in frequency with which energy reaches a receiver when the receiver and the energy source are in motion relative to each other*. Determining the Doppler Effect or shift is straightforward when the energy transmission source is stationary and the target being sampled is moving (or stationary). Any frequency shifts would be solely the result of

the target moving toward or away from the energy transmission source.

The relationship between the speed of transmitted electromagnetic (E-M) energy and the frequency and wavelength of that energy is

$$\mathbf{c = f\lambda} \qquad (1)$$

where c is the speed of light (assumed to be constant), f is the frequency and λ is the wavelength of the energy. Since c is constant, f and λ have an inverse relationship. For example, if λ is increased, f must decrease, and vice versa.

The equation used to represent the Doppler motion of a target sampled by a weather radar is

$$\mathbf{V_r = \frac{-(f_{dop}\lambda)}{2}} \quad \text{or} \quad \mathbf{f_{dop} = \frac{2V_r}{\lambda}} \qquad (2)$$

where V_r is the target's radial velocity, f_{dop} is the Doppler frequency shift due to target motion either toward or away from the radar, and λ is the wavelength of the transmitted energy. For a stationary target, there will be no change in wavelength or frequency. The factor of two is there because the signal is transmitted and interacts with the target, then is reflected back. The minus sign is there for target direction. By convention, inbound velocities are negative and outbound velocities are positive. For example, an inbound target produces a positive Doppler shift, making the velocity negative. An outbound target produces a negative Doppler shift, making the velocity positive.

Sound Waves and Doppler Shift

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 detectable by the human ear, which is not as precise an instrument as the WSR-88D.

WSR-88D and Doppler Shift

Now for the same target motion, but the listening device is the WSR-88D and the signal is no longer sound waves. The frequency of transmission, 2850 MHz, is very high compared to sound. The frequency shift, ±487Hz, 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.

If the frequency shift isn't used to determine target motion, then what is? Figure 3-1 shows 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.

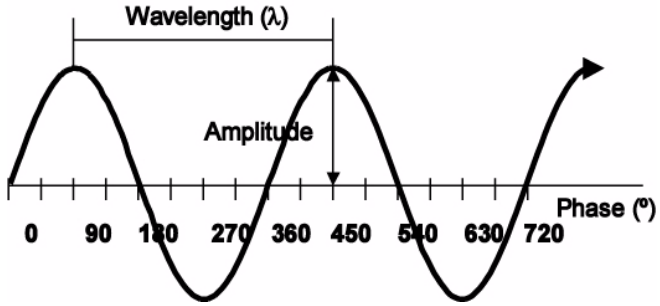

Figure 3-1. Radar wave characteristics.

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.

Signal Phasor

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 (see Fig. 3-2). The phase is the angle 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.

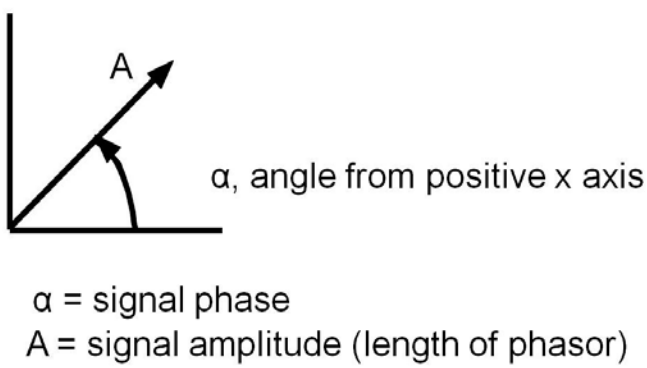

Figure 3-2. A phasor represents the necessary information from a returned pulse including signal phase and amplitude.

Phasor for a Single Pulse

In Figure 3-3, 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°.

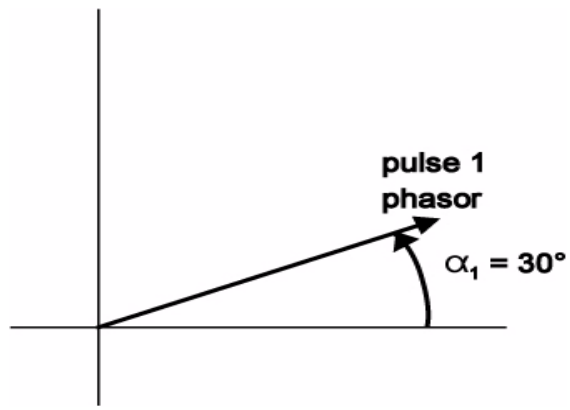

Figure 3-3. Pulse 1 phasor with amplitude and phase.

Phasors for Two Pulses

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.

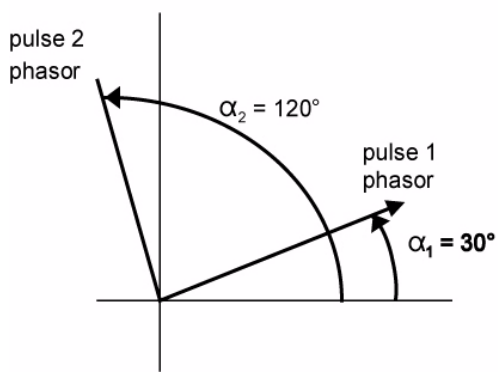

Figure 3-4. Two pulse phasors.

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.

Determining Target Direction

The target direction, inbound vs. outbound, is determined by the phasor rotation from pulse 1 to pulse 2 (Figure 3-5). If the rotation is counterclockwise, the direction is inbound. If clockwise, the direction is outbound. 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.

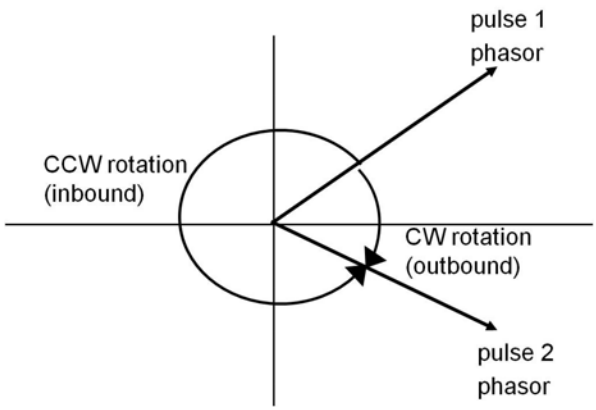

Figure 3-5. If the rotation is counterclockwise, the direction is inbound. If clockwise, the direction is outbound.

Pulse to Pulse Phase Change

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 (see Fig. 3-6). The angle less than 180° is always used.

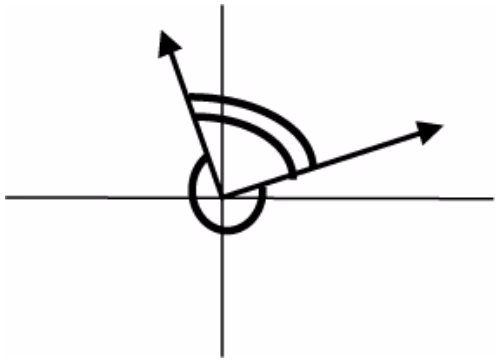

Figure 3-6. Two possible angles between the two phasors.

Why less than 180°?

A phase shift of exactly 180° introduces ambiguity (see Fig. 3-7). 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 is greater than or equal to 180°, there is ambiguity in determining the velocity.

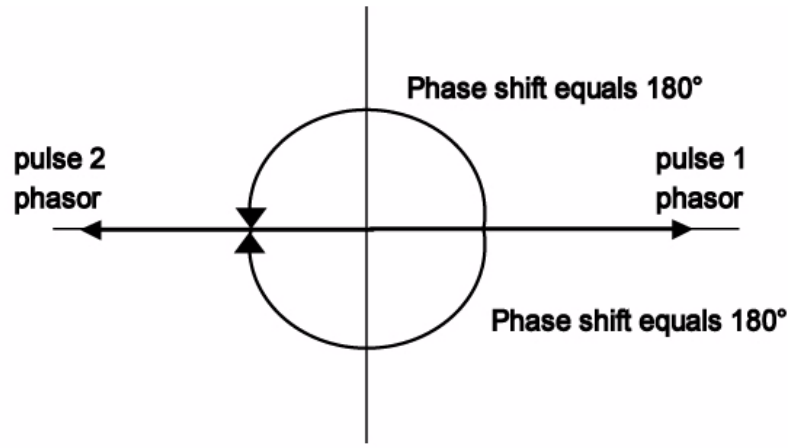

Figure 3-7. A phase shift of greater than or equal to 180° introduces ambiguity.

Maximum Unambiguous Velocity (V_{max})

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 PRF. With the WSR-88D, V_{max} values range from about 16 to about 64 kts.

Phase Shift and Radial Speed Relationship

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}.

A pulse pair phase shift and radial speed are related by the ratio

$$\frac{\mathbf{P.S.}}{\mathbf{180°}} = \frac{|V_r|}{|V_{max}|} \qquad (3)$$

where **P.S.** is the amount of pulse pair phase shift, $|\mathbf{V_r}|$ is the target radial speed, and $|\mathbf{V_{max}}|$ is the maximum unambiguous speed (magnitude of V_{max}). For any given V_{max}, target speed is directly related to the amount of pulse pair phase shift that occurs.

Example 1: V_{max} = 60 kts

Here's an example where V_{max} = 60 kts. In this case, the phase shift between pulses is 90°, which is one half of the maximum of 180°. The radial speed is then going to be one half of 60 kts (the maximum), or 30 kts.

$$\frac{90°}{180°} = \frac{|V_r|}{60}$$

$$60(1/2) = |V_r|$$

$$|V_r| = 30kts$$

Example 2: V_{max} = 50 kts

In this example V_{max} = 50 kts and the phase shift between pulses is again 90°. The radial speed is going to be one half of 50 kts (the maximum), or 25 kts.

$$\frac{90°}{180°} = \frac{|V_r|}{50}$$

$$50(1/2) = |V_r|$$

$$|V_r| = 25kts$$

Example 3: V_{max} = 54 kts

This time V_{max} = 54 kts and the phase shift between pulses is 30°, which is one sixth of the maximum shift of 180°. The radial speed is going to be one sixth of 54 kts, or 9 kts.

$$\frac{30°}{180°} = \frac{|V_r|}{54}$$

$$54(1/6) = |V_r|$$

$$|V_r| = 9kts$$

First Guess and Alias Velocities

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} also 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.

First Guess Correct

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 (see Fig. 3-8).

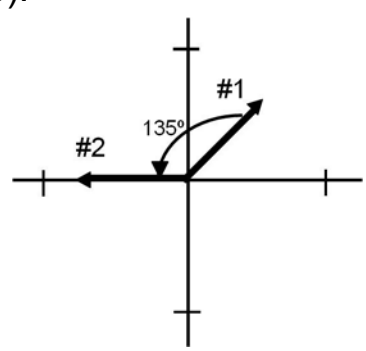

Figure 3-8. A pulse pair with an angle less than 180°.

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.

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

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

First Guess Incorrect

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 the example shown in Figure 3-9, 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 possible velocity or alias. How these aliases are used to find the true radial velocity is discussed later in Topic 3.

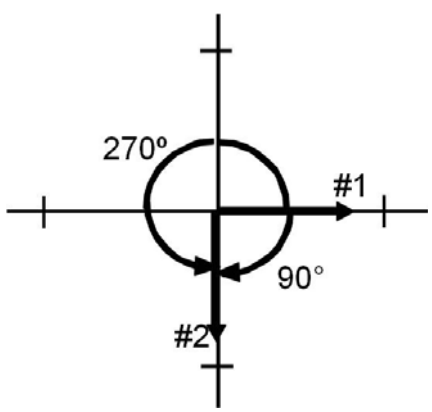

Figure 3-9. Pulse pair with a phase shift of 90°.

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.

Super Resolution Signal Processing

Super resolution (SR) is defined as having 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° 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" presented in Topic 3 Lesson 1 of ~1.0° is based on the antenna being stationary. Antenna movement produces what is know as the effective beamwidth.

Effective Beamwidth

The physical beamwidth for the WSR-88D for a single pulse is 1.0°. However, the antenna is rotating as pulses are transmitted and data are assigned to a 1.0° radial. 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 Figure 3-10, each dot represents a single pulse and the location of the beam centerline. There are 11 pulses that fall within the 0° to 1° radial in this simplified example. As the antenna rotates, the pulse on the inside edge of the radial still has a physical beamwidth of 1°. That means that the pulse is sampling a volume that is both

inside and outside the radial. Only when the beam is centered on this radial at 0.5° is the associated pulse contained within the radial. This process of capturing volumes outside a radial as base data are collected results in what is known as the effective beamwidth. For the WSR-88D, the effective beamwidth is about 1.4°.

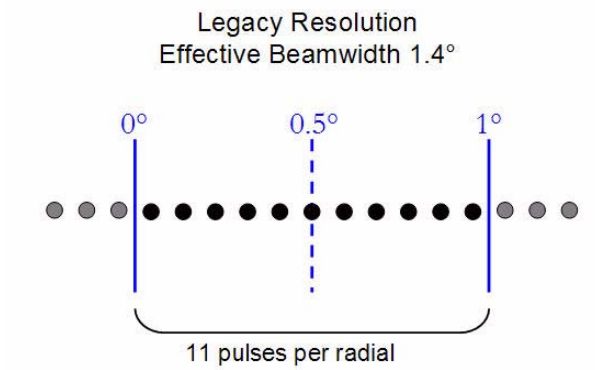

Figure 3-10. Effective beamwidth of legacy resolution.

Overlapping Radials

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. (see Fig. 3-11). 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 effective beamwidth needs to be narrowed.

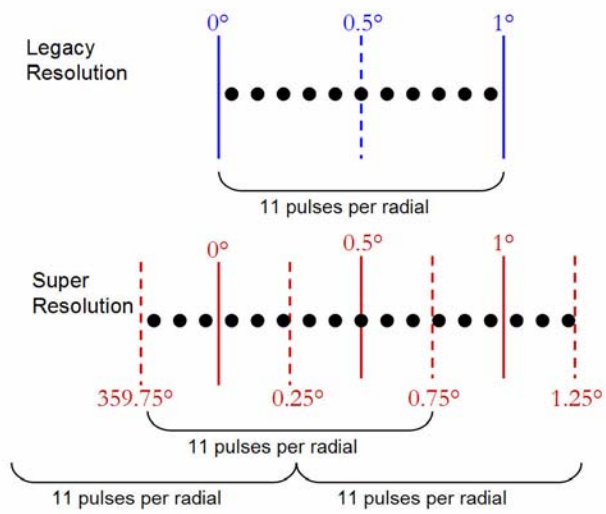

Figure 3-11. Overlapping radials for super resolution.

Data Windowing

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 (see Fig. 3-12).

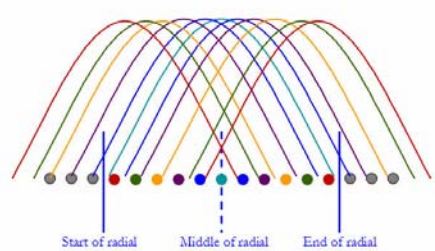

Figure 3-12. Rectangular window, where all the pulses have equal weight.

What is needed for super resolution is a window that narrows the effective beamwidth (see Fig. 3-13). 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.

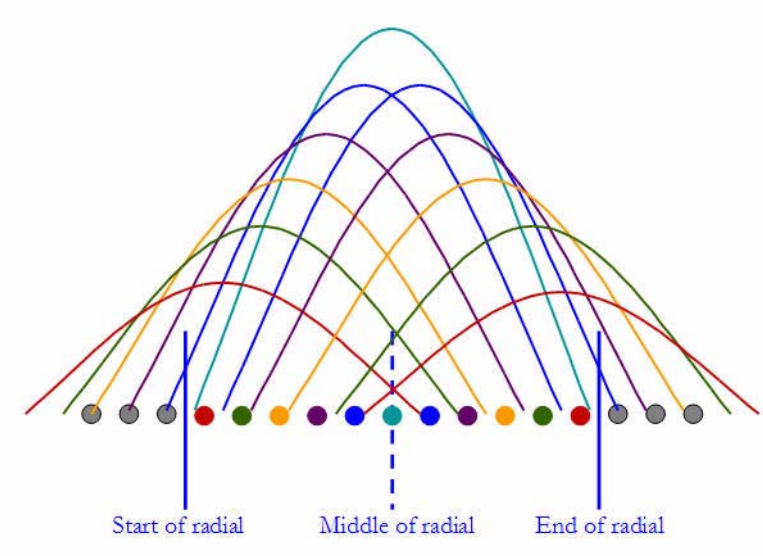

Figure 3-13. Windowing technique which narrows the effective beamwidth.

Super Resolution Base Data Quality

The 0.5° azimuthal resolution is obtained through signal processing techniques, specifically overlapping radials and data windowing. 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 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 cannot ingest super resolution base data due to this noisiness and error.

Base Reflectivity (Z) Generation

For each range and azimuth, there are multiple pulses used to generate base reflectivity. The returned power for these pulses is averaged for each 0.25 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 is 0.5° azimuth by 0.25 km (super resolution). For the Batch or higher elevation Z products, the best resolution to 1.0° azimuth by 0.25 km (legacy resolution).

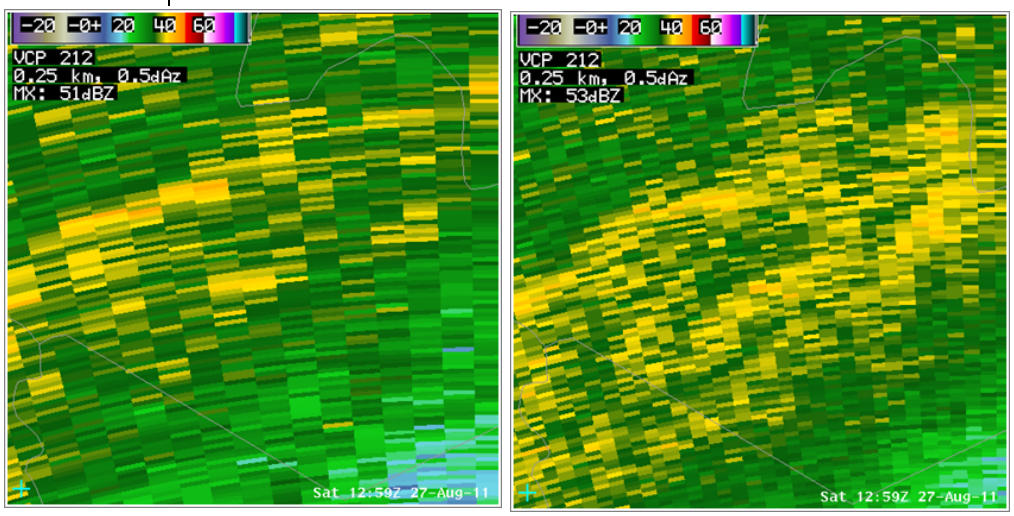

Figure 3-14. Legacy reflectivity (left) and super resolution reflectivity (right).

Recall that PRF & R_{max} have an inverse relationship with one another. Though use of a low PRF provides a long R_{max} and reduces multiple trip echoes, high PRFs are needed for velocity data collection. To mitigate this problem, Split Cut mode is used for the lowest 2 or 3 elevations for all the VCPs except VCP 121 (see Fig. 3-15). 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. Though range folding is com-

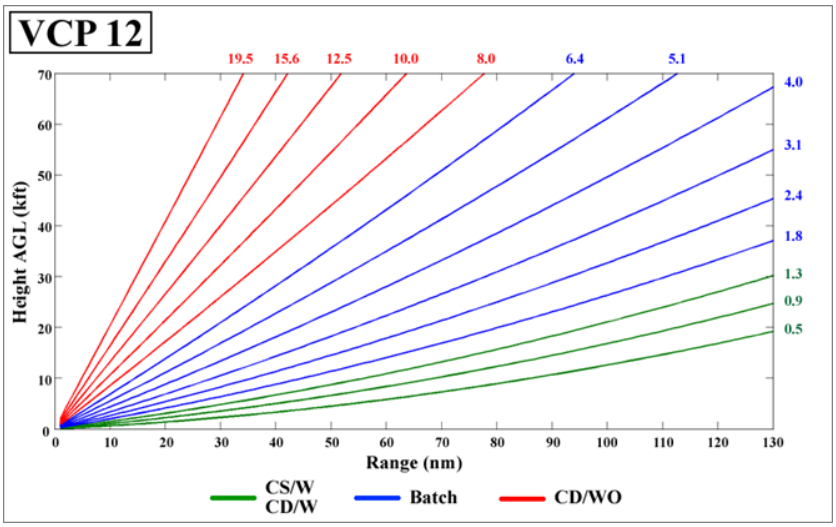

Figure 3-15. Split cut elevations for VCP 12.

mon, base velocity and spectrum width are generated from the CD rotation.

The data collected from these two rotations are used together to “range unfold” velocity and spectrum width. The range unfolding techniques are presented in later Topic 3 lessons.

Base Velocity (V) Generation

Velocity also has a 0.25 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. This is because the goal is to assess the mean flow within

a volume of the atmosphere, and the larger scatterers are more likely to move with the mean flow.

Calculating this is surprisingly easy. Each phasor in Figure 3-16 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.

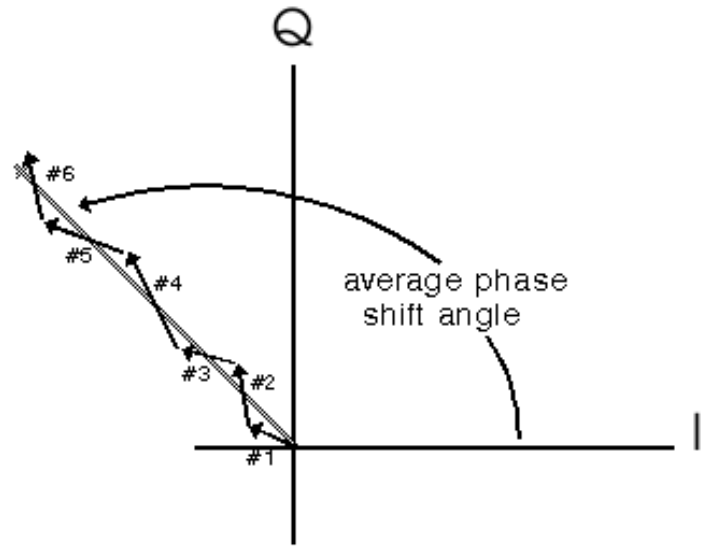

Figure 3-16. A vector sum of the phasors results in a power-weighted average pulse pair phase shift.

Base Spectrum Width (SW) Generation

Spectrum width also has a 0.25 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 pro-

portional 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.

The technique used to calculate Spectrum Width is called Autocorrelation. The series of phase shifts from one pulse to the next are what's being correlated. 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, and will be used for a variety of applications in Topic 3.

Doppler Power Spectrum

The Doppler Power Spectrum (examples in Fig. 3-17 and Fig. 3-18 on page 3-30) 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 as "spectral coefficients." A bell curve is fit to these coefficients and the average returned power (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 the spectrum width.

Low vs. High 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 (Fig. 3-17). There is minimal variation in pulse pair phase shifts and the Doppler Power Spectrum

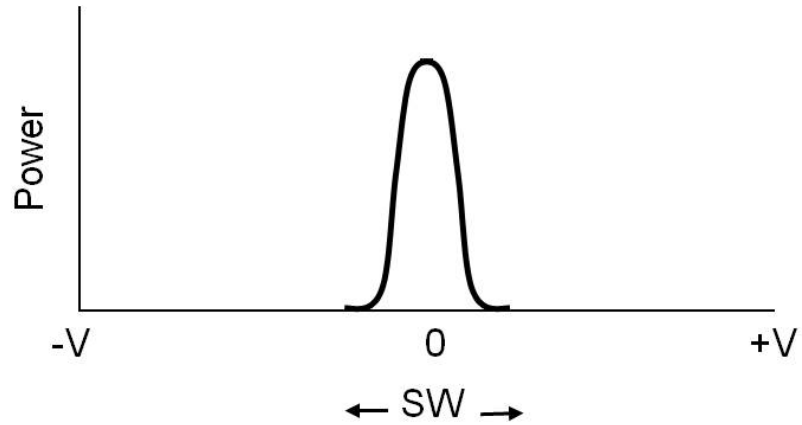

Figure 3-17. Strong power and near zero velocity leads to a narrow spike in the Doppler Power Spectrum and low spectrum width.

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 (Fig. 3-18). 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.

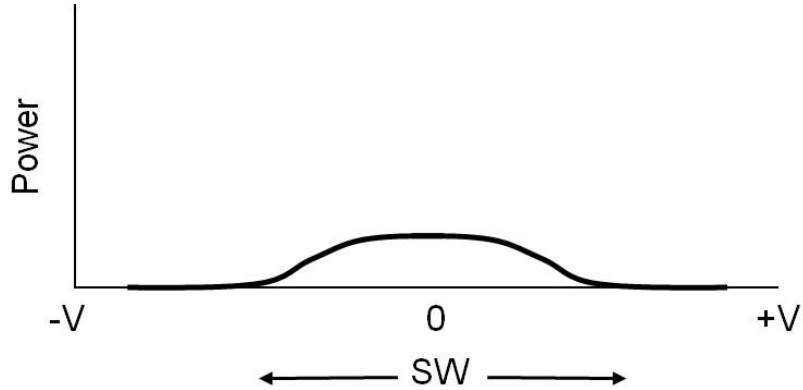

Figure 3-18. Low power across a wide variety of velocities results in a wide and flat Doppler Power Spectrum and high spectrum width.

Z Generated from P_{rH}

Base reflectivity is calculated from the average returned power, that is then converted to reflectivity from the Probert-Jones radar equation [see

equation (4)]. 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 base data are built from the horizontal channel information only.

$$P_r = \left[\frac{P_t G^2 \theta^2 H \pi^3 K^2 L}{1024(\ln 2)\lambda^2}\right] \times \frac{Z}{R^2} \qquad (4)$$

ZDR Generated from P_{rH} and P_{rV}

Differential Reflectivity, ZDR, is defined as the difference between the horizontal and vertical reflectivities, each expressed in units of dB. This equation is presented as a definition of ZDR, not how it is actually calculated [see equation (5)]. The larger the raindrop, the greater its horizontal extent compared to the vertical, thus the ZDR value is highly positive. On the other hand, drizzle drops are nearly spherical, and ZDR is near zero.

$$ZDR = Z_H - Z_V \qquad (5)$$

There is no returned pulse phase information used to calculate ZDR, only the returned power from both the horizontal and vertical channels.

Distance Learning Operations Course

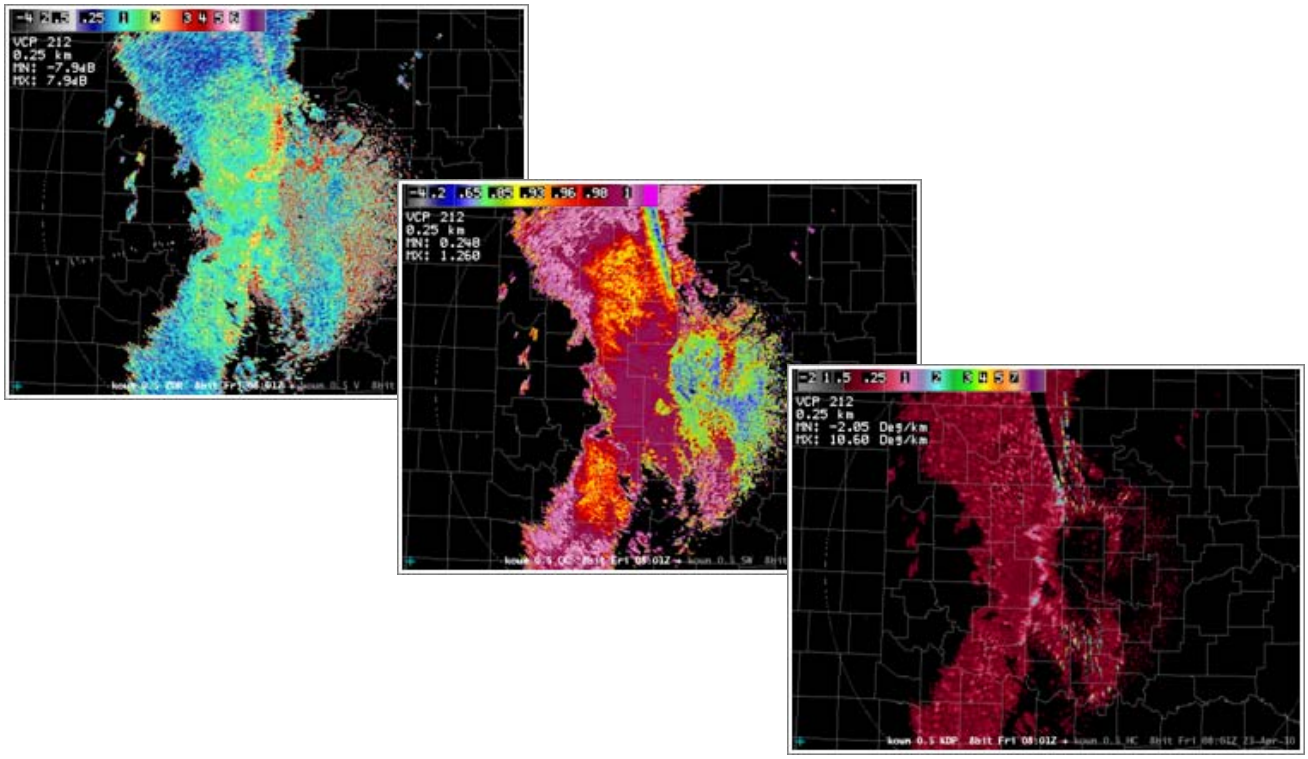

Topic 3, Lesson 3: More Base Data Generation

Presented by the
Warning Decision Training Branch

Version: 1210

Lesson 3: More Base Data Generation

Generating base data of the highest quality is an obvious priority. Once generated, the base data are sent, via the wideband, to the RPG. At the RPG, the base data are input to *all* the algorithms, building *all* the Base and Derived products. This lesson completes the necessary foundation for studying WSR-88D data quality issues.

A Note on Terminology

Though often used interchangeably, range gate and range bin have different meanings. Here, range gate will be used when describing the appearance on WSR-88D products and range bin will be used when describing algorithm functions.

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.
4. Identify the purpose of recombination at the RPG.

Differential Reflectivity (ZDR)

RDA Generation of ZDR

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 (ZH/ZV). When the Zs are substituted with returned power, the radar constants and the range, the range terms cancel [equation (6)].

$$ZDR = 10\log_{10}\left(\frac{Z_H}{Z_v}\right) = 10\log_{10}\left(\frac{P_h}{P_v}\right) + 10\log_{10}\left(\frac{C_h}{C_v}\right) \quad (6)$$

Notice how the logarithmic separation of the returned power and the radar constant terms underscores 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 in DLOC Topic 6.

Dual-Pol on Split Cuts

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 dual-pol 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 left in Figure 3-19 is a ZDR image built from the 0.5° azimuthal resolution Level II data. On the right is that same data that has been recom-

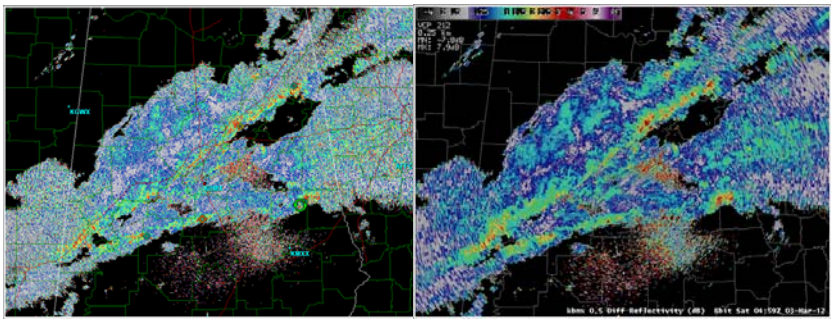

Figure 3-19. Differential reflectivity image build from the 0.5° azimuthal resolution Level II data (left) and the same data recombined to 1.0° azimuth and smoothed.

bined to 1.0° azimuth and smoothed. This is how ZDR appears through the AWIPS display.

All of the dual-pol base products are recombined at the RPG to 1.0° azimuth and smoothed. There will be more about this “preprocessing” later in Topic 3.

Dual-Pol on Batch and Above

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 CD/WO pulses because that is all that is used. Since multiple trip echoes are so unlikely at these higher elevations, the CD/WO means Contiguous Doppler with no range unfolding.

Correlation Coefficient

RDA Generation of CC

Correlation coefficient (CC) measures the consistency of the H and V returned power and phase for each pulse (see Fig. 3-20). 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 (SW) 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.

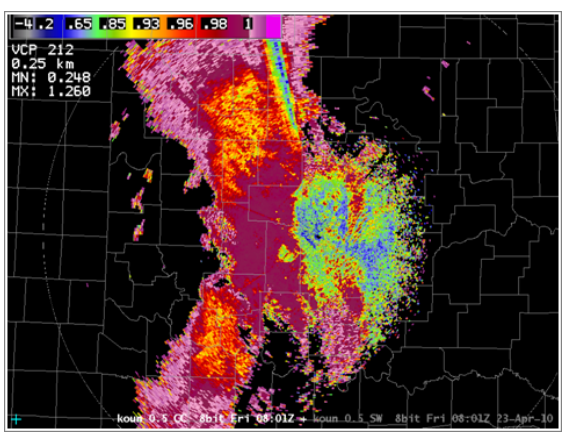

Figure 3-20. Correlation coefficient (CC) measure the consistency of the H and V returned power and phase for each pulse. In some ways, it is similar to spectrum width.

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 known 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 of H and V is checked for each pulse.

Correlation Coefficient, Φ_{DP}, and Cross Correlation Vectors

Cross correlation vectors are at the heart of dual-pol base data. Figure 3-21 shows 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.

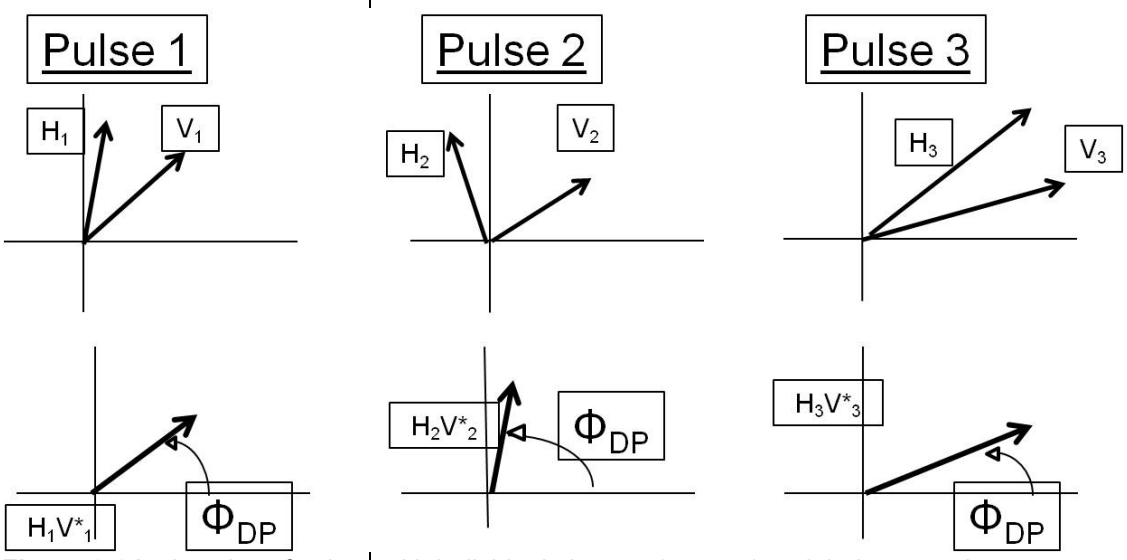

Figure 3-21. A series of pulses with individual phasors (top row) and their respective cross correlation vector (bottom row).

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 technically 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.

Correlation Coefficient is Based on Φ_{DP}

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 (Figure 3-22, 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.

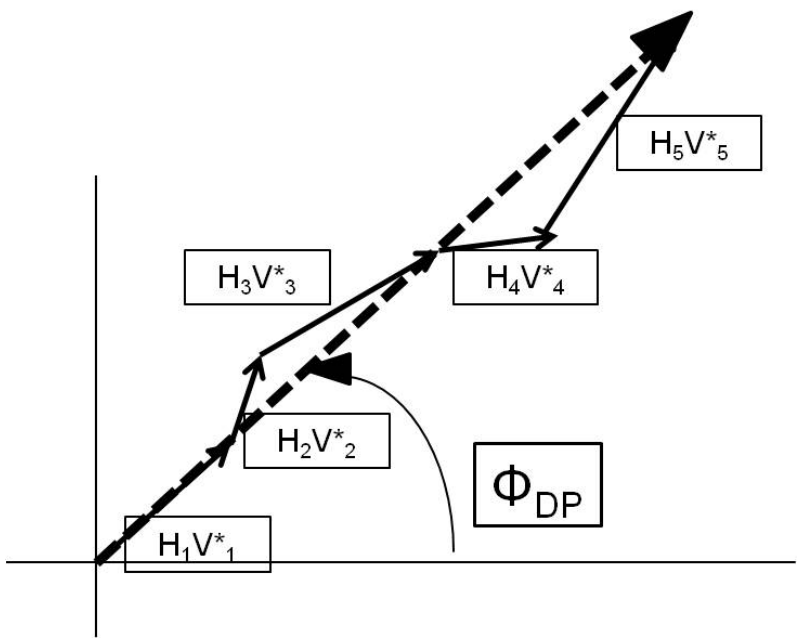

Figure 3-22. Differential Phase, Φ_{DP}, is the angle of the sum of the cross correlation vectors for a series of pulses.

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. The vector-summed Φ_{DP} is base data generated at the RDA and sent to the RPG. Specific Differential Phase, or KDP, is derived from it.

CC is calculated by taking the length (or amplitude) of this vector sum 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 (see Fig. 3-23).

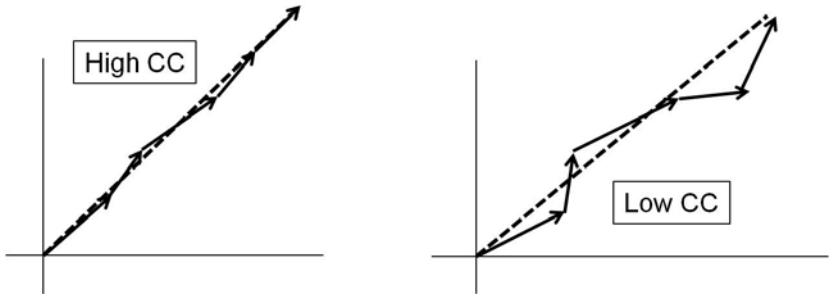

Figure 3-23. The less variation in the individual cross correlation vectors, the higher the CC values (left). The more variation, the higher the CC (right).

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 (left, Figure 3-23). The more diverse the scatterers, the more variation with the cross correlation vectors, and CC gets closer to 0. However, CC is never exactly equal to 0 or 1.

Why Does CC Matter?

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

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

Correlation Coefficient and a Weak Returned Signal

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, Figure 3-24) 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.

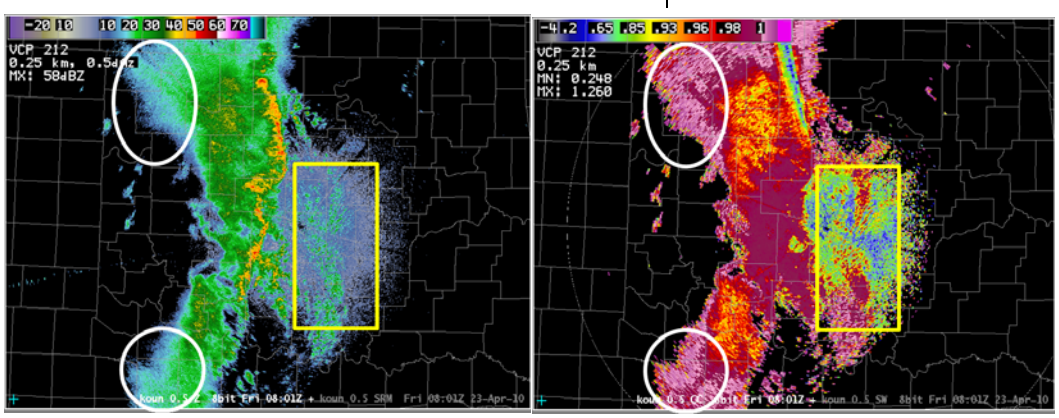

Figure 3-24. In areas of weak signal, correlation coefficient (CC, right) will be noisy making it easier to find areas of weak signal in reflectivity (Z, left).

Though the Z values near the radar and at long range are similar in Figure 3-24, the CC values tell 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 Similar but Different

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 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.

In Figure 3-25, the circled area is weak signal close to an intense supercell. There is high spectrum width due to both the weak signal and the turbulence. Notice that the velocity field in this same area is noisy.

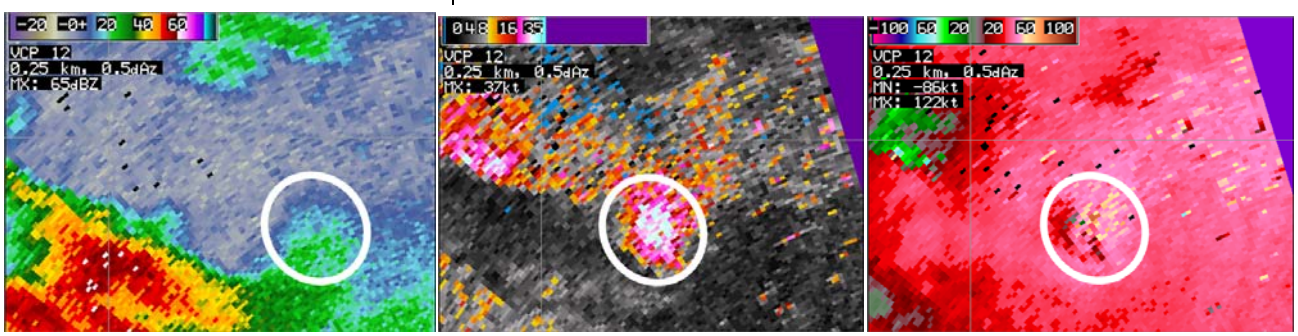

Figure 3-25. An area of weak signal close to an intense thunderstorm (white circle) has high spectrum width (SW, middle) and slightly noisy velocity (right).

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

Unlike spectrum width, correlation coefficient is calculated from both the horizontal and vertical channels, and tells us something about the nature of the dual-pol data quality and 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 a high consistency.

In Figure 3-26, 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.

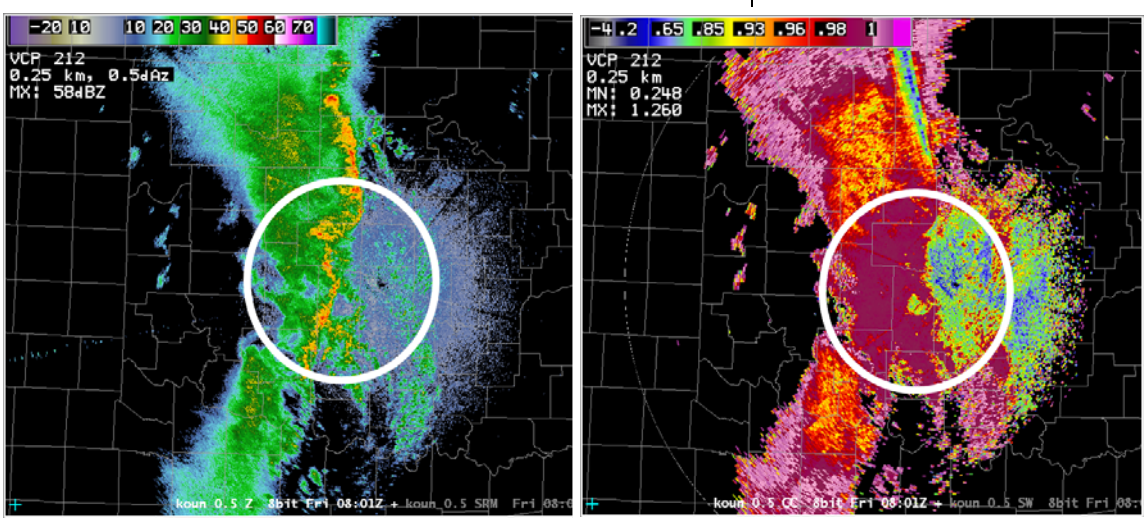

Figure 3-26. The white circle highlights an area of both precipitation and clutter near the radar. In the reflectivity (Z, left) data, it's not clear where the precipitation is versus clutter. In the correlation coefficient (CC, right) data, high values of CC indicate precipitation and low values represent clutter.

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

Specific Differential Phase (KDP)

RDA and RPG Roles for KDP

The last 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 Differential Phase, or Φ_{DP}, which is generated by the RDA and sent to the RPG (see Fig. 3-27). Since Φ_{DP} is not available on AWIPS, we think of KDP as a dual-pol base product.

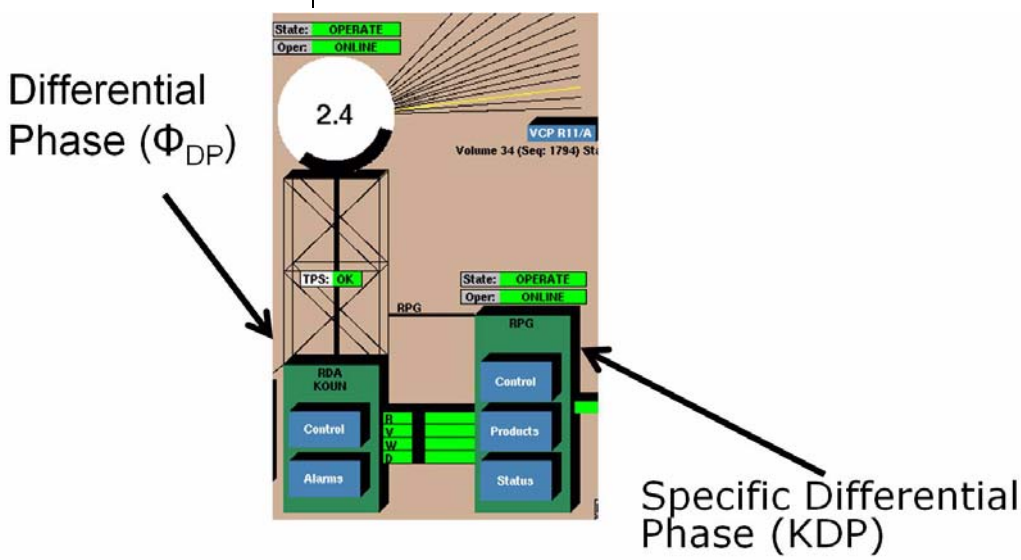

Figure 3-27. Differential phase, or Φ_{DP}, is generated by the RDA and then sent to the RPG where KDP data is built from the Φ_{DP}.

RDA Generation of Φ_{DP}

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.

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

As the pulse propagates through different mediums (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 Figure 3-28, 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.

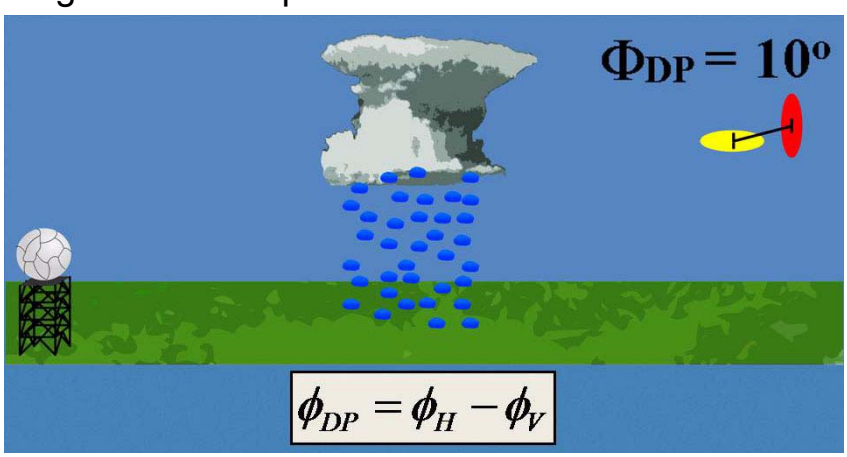

Figure 3-28. As the pulse propagates through different mediums, there is a delay that is apparent in the phase of the returned pulse.

Φ_{DP} Affected by Liquid Water

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 two 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, there is more propagation delay in the horizontal direction than in the vertical. That was the example on Figure 3-29.

2. Assume the same size and shape raindrops in each of two volumes as in Figure 3-29. 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.

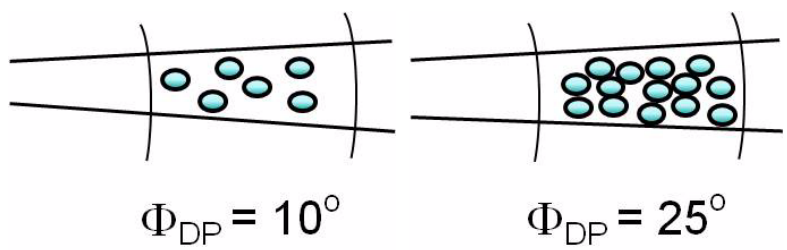

Figure 3-29. Given the same size and shape raindrops, a higher concentration will lead to more propagation delay in the horizontal and thus a higher differential phase.

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

How Φ_{DP} Changes Along the Radial

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 in Figure 3-30, 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.

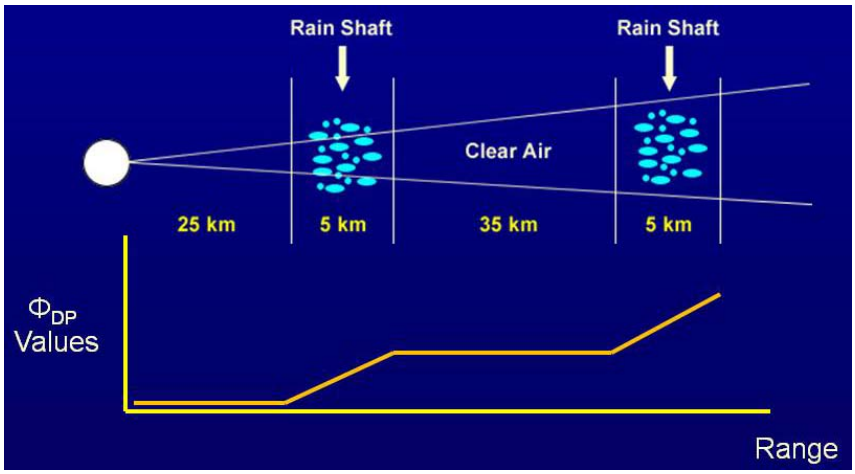

Figure 3-30. The value of Φ_{DP} propagates down radial, accumulating with range.

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

Why Specific Differential Phase (KDP)?

Differential phase, Φ_{DP}, is sent as base data from the RDA, but is not generated as a base product. Instead, we have specific differential phase or KDP. KDP is defined as the range derivative of Φ_{DP}, and is the dual-pol base product seen in AWIPS. 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.

$$KDP = \frac{\phi_{DP}(r_2) - \phi_{DP}(r_1)}{2(r_2 - r_1)} \qquad (7)$$

Equation (7) 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.

Specific Differential Phase (KDP) and Reflectivity

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.

Specific Differential Phase (KDP) and Liquid Water

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. In Figure 3-31, 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 in AWIPS. The higher KDP values correspond to the areas of highest Φ_{DP} gradients along the radial.

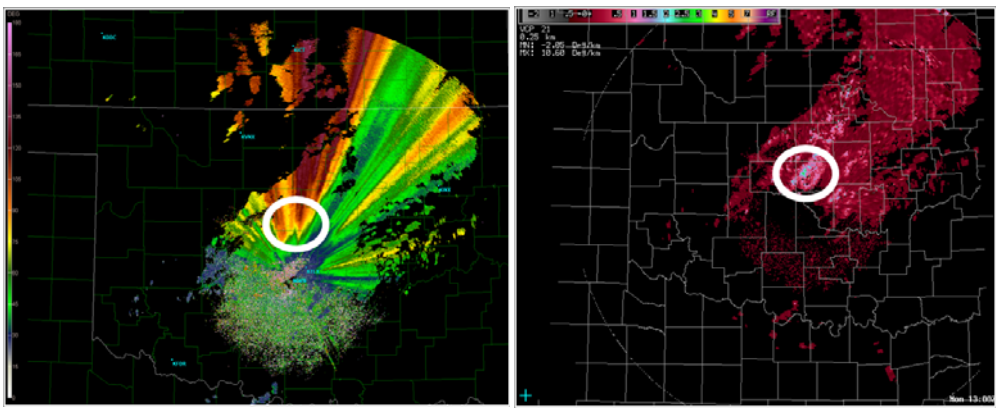

Figure 3-31. High values of specific differential phase (KDP, right) correspond to the areas of highest differential phase (Φ_{DP}, left) gradients along the radial.

Summary of the Dual-Pol Base Products

Since dual-pol data are still new to the WSR-88D, here's a quick summary of each of these new base products.

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.

Recombination at RPG

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

For the Split Cut elevations, the dual-pol base data are processed at the RDA with an azimuthal resolution of 0.5°, then sent via the wideband to the RPG. This includes differential reflectivity (ZDR), correlation coefficient (CC), and differential phase (Φ_{DP}), all of which are too visually noisy for direct product generation.

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

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

The dual-pol base data are also “preprocessed”, which includes smoothing and converting differential phase, Φ_{DP}, into specific differential phase,

(KDP). The Dual-Pol RPG Preprocessor algorithm will be discussed later in Topic 3.

In Figure 3-32, 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 in AWIPS. The data have been recombined to 1.0°, as well as preprocessed.

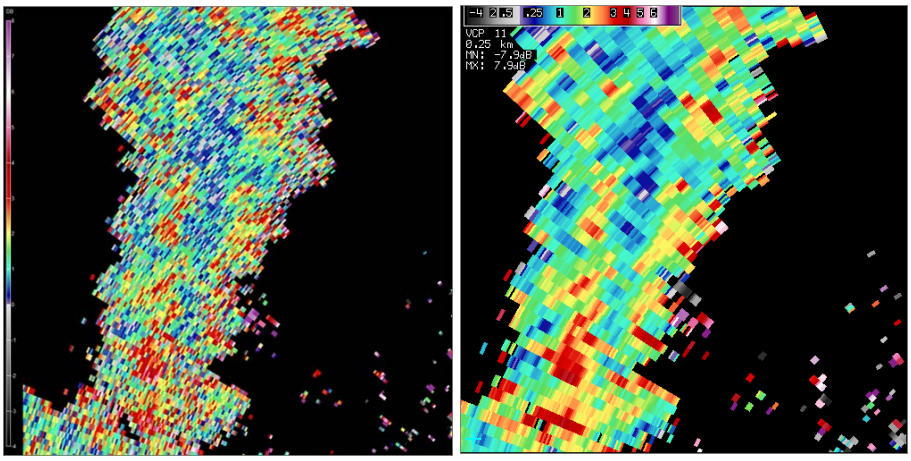

Figure 3-32. On the left, ZDR base data as seen in a Level II viewer (GR Analyst) and on the right, ZDR base data that have ben recombined to 1.0°, as well as preprocessed.

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 derived products that were designed for legacy resolution base data such as the tornado detection algorithm, precipitation processing, vertically integrated liquid, and many others. One exception is the Mesocyclone Detection Algorithm (MDA). You will learn much more about the derived products later in DLOC, Topic 5.

Reflectivity Recombination

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

For reflectivity, the recombination process is a linear average of the eight super resolution bins into the corresponding one legacy resolution bin (see Fig. 3-33). For bins that are assigned No Data, the associated power is estimated, converted to Z and included in the average.

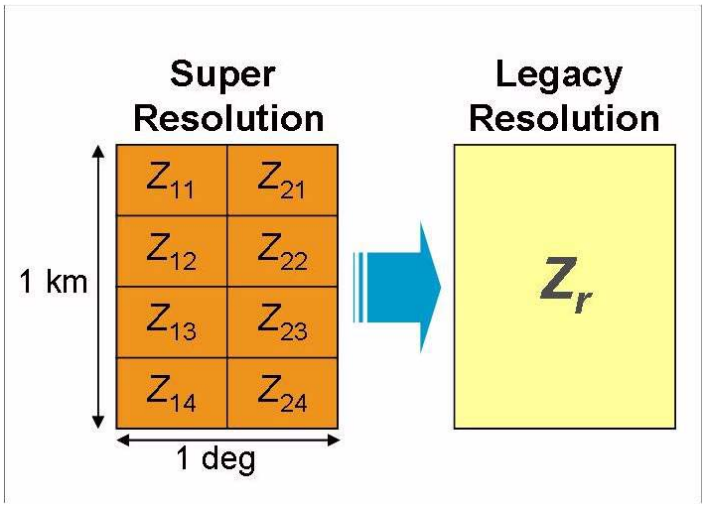

Figure 3-33. Reflectivity data is recombined by a linear average of eight super resolution bins into the corresponding one legacy resolution bin.

Velocity and Spectrum Width Recombination

This step is to support RPG algorithms that require 1.0° azimuth by 0.25 km range resolution for velocity and spectrum width input.

For velocity and spectrum width, the recombination process is a power weighted average of the two super resolution bins into the corresponding one legacy resolution bin (see Fig. 3-34). There is an additional step in the spectrum width processing to account for the variance of the two corresponding velocity values. If both bins are assigned No Data or RF, then the legacy bin is also assigned No Data or RF. If one of the two super resolution bins is assigned No Data or RF, then the remaining valid bin is assigned as the legacy resolution value.

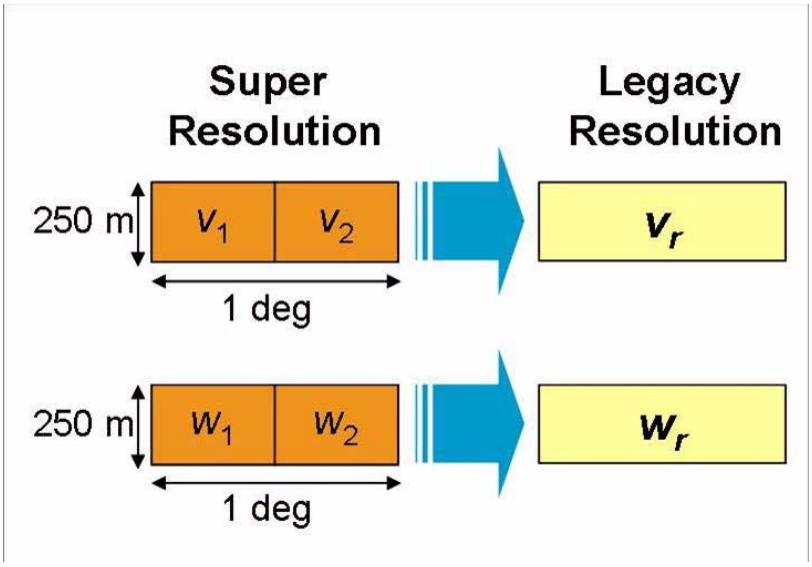

Figure 3-34. Velocity and spectrum width are recombined using a power weighted average of the two super resolution bins corresponding to one legacy resolution bin.

Distance Learning Operations Course

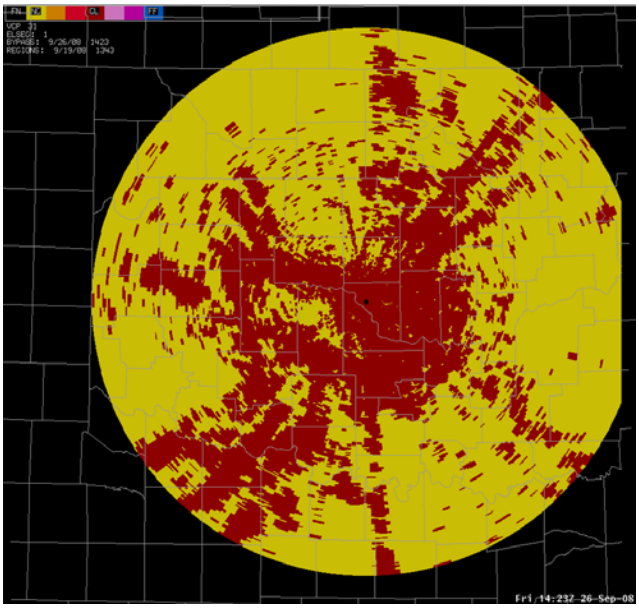

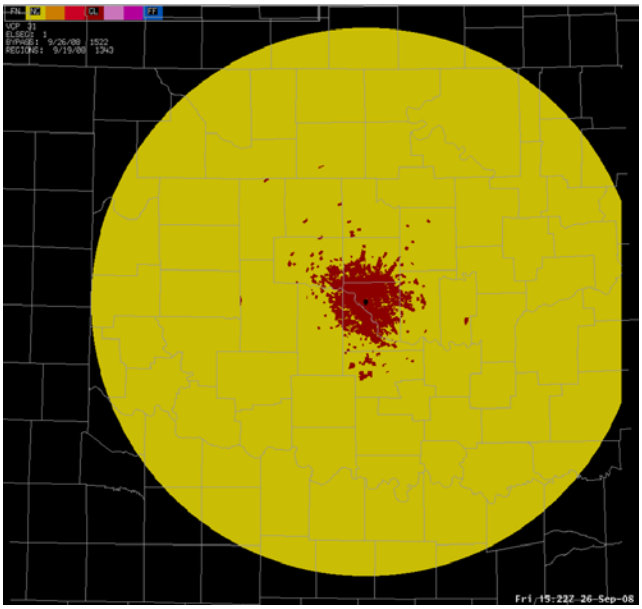

Topic 3, Lesson 4: Filtering Clutter and Range Unfolding Velocity

Presented by the
Warning Decision Training Branch

Version: 1210

Lesson 4: Filtering Clutter and Range Unfolding Velocity

Clutter Filtering is one of the many important functions of the signal processor. It is applied to the returned signal ***before*** the base data are generated. Range Unfolding of velocity and spectrum width is a fundamental algorithm applied to the base data before any of the radar products are built.

A Note on Terminology

Though often used interchangeably, range gate and range bin have different meanings. Here, range gate will be used when describing the appearance on WSR-88D products, and range bin will be used when describing algorithm functions.

Objectives

1. Identify the purpose, strengths and limitations of the following clutter filtering algorithms:

d. Clutter Mitigation Decision (CMD)

e. Gaussian Model Adaptive Processing (GMAP)

2. Identify examples of moving ground-based targets that cannot be identified by CMD.

3. Identify the purpose, strengths and limitations of the following range unfolding algorithms:

a. Legacy Range Unfolding

b. SZ-2 Range Unfolding

Ground Clutter Contamination

In general, ground clutter on WSR-88D radar products is return from stationary or nearly stationary ground targets that has not been filtered. Clutter filtering is applied at the signal processor just before the base data are built. Since all the base and derived products are built from the base data, unfiltered clutter will negatively impact all the products. There are two types of clutter contamination. The first is normal ground clutter, meaning fea-

tures that are present all the time, such as terrain, buildings, etc. The second type is transient and dependent on beam propagation, known as anomalous propagation clutter.

For the reflectivity (left) and velocity (right) images in Figure 3-35, 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.

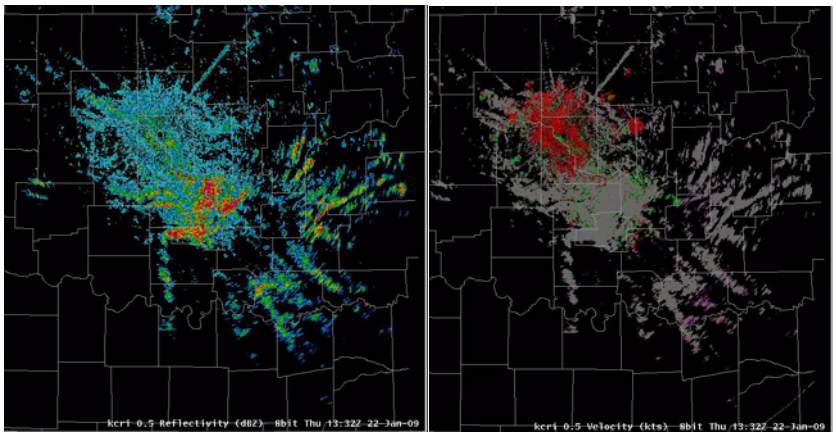

Figure 3-35. Reflectivity (Z, left) and velocity (V, right) data with clutter filtering only applied very close to the radar.

If something is moving, it is not clutter. The WSR-88D clutter algorithms are designed to detect and remove 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 (see Fig. 3-36).

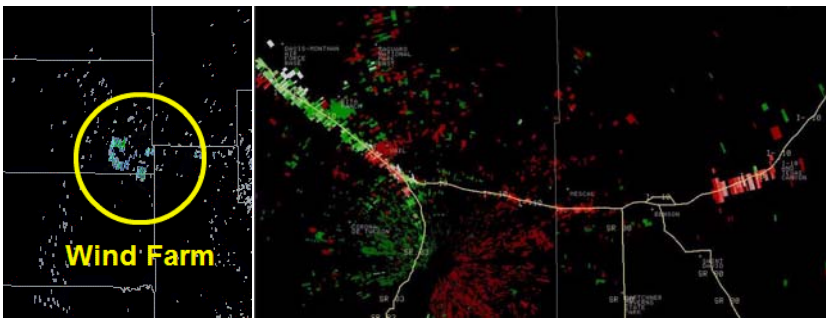

Figure 3-36. Contamination of radar data from wind farms (left) and traffic (right).

Though research is ongoing, there is not yet any algorithm that can discriminate weather from wind farms or road traffic, both of which have non-zero velocity most of the time.

Clutter Filtering Algorithms

There are two different algorithms in the clutter filtering process. The data are first processed by the Clutter Mitigation Decision (CMD) algorithm (Figure 3-37, left), 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 (Figure 3-37, right) then applies the signal reduction of the clutter. Both CMD and GMAP are run at the RDA signal processor.

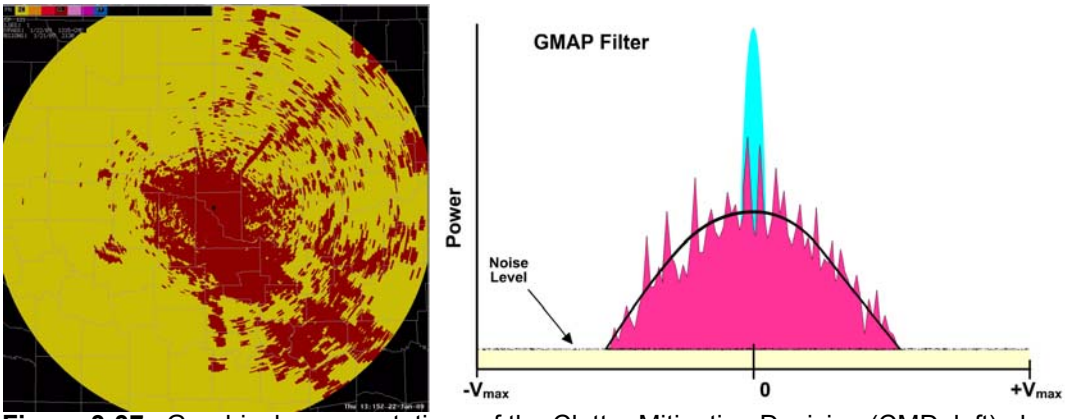

Figure 3-37. Graphical representations of the Clutter Mitigation Decision (CMD, left) algorithm results and the Gaussian Model adaptive Processing (GMAP) algorithm processing

Clutter Mitigation Decision (CMD)

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

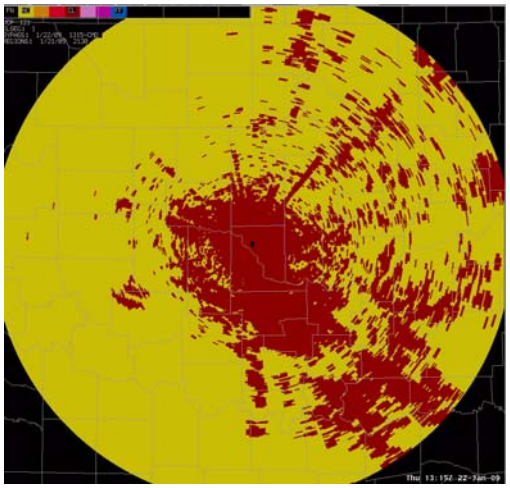

Figure 3-38. A visualization of the CMD generated Bypass Map.

Figure 3-38 is a visualization of a CMD generated Bypass Map. It is called the Clutter Filter Control (CFC) product. Each of the red bins has been identified by CMD as containing clutter. The yellow bins do not contain clutter, based on CMD's analysis.

CMD Inputs

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

The inputs based on Reflectivity, Z, include:

1. Z texture, which is smooth (a low value) for a meteorological target and rough (a high value) for a clutter target
2. Z spin, which looks at the number of times that the Z gradient changes sign along a portion of the radial. The higher the number of sign changes, the higher the likelihood of clutter.

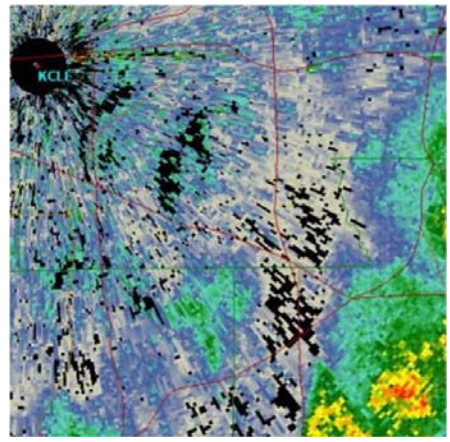

Figure 3-39. Areas of weak signal near the radar lead to noisy data.

In the reflectivity image in Figure 3-39, the weak signal areas close to the radar have a visually rougher texture with more frequent changes in sign of the Z gradient.

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

CMD Inputs - Dual-Pol Base Data

CMD has 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 Figure 3-40, you can likely tell by looking at the smoothness of the ZDR data where the clutter (yellow circle) vs. the weather (white oval) is located. That variation (or lack of it) is captured in the standard deviation data. The higher the stan-

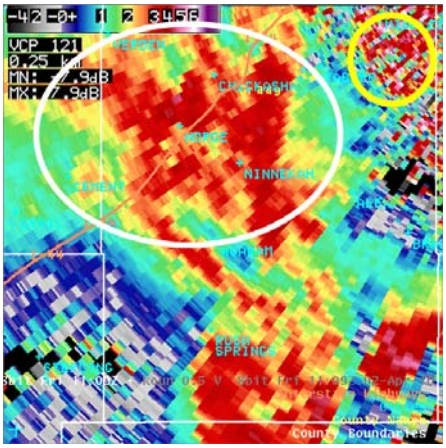

Figure 3-40. Smooth ZDR data (white oval) is usually weather, while noisier ZDR (yellow circle) is more than likely clutter.

dard deviation of ZDR, the more likely it is that the bin contains clutter.

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

Though Φ_{DP} is not a base product displayed in AWIPS, it has significant value. The higher the standard deviation of Φ_{DP}, the higher the likelihood that the range bin contains clutter. In Figure 3-41, compare the noisiness of the Φ_{DP} data surrounding the radar (yellow oval) to the areas of precipitation to the north through the east (white oval).

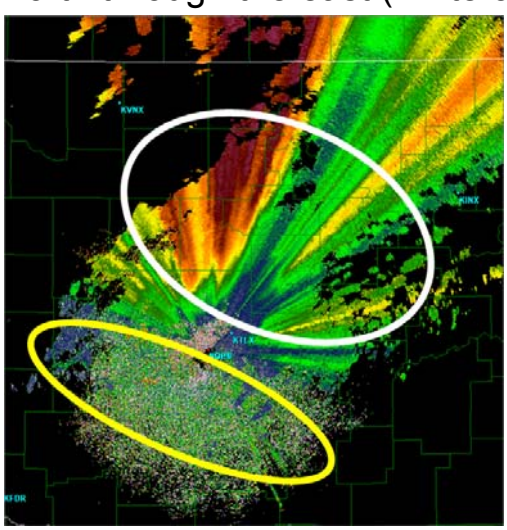

Figure 3-41. Areas of noisy Φ_{DP} (yellow oval) are likely clutter, while smoother Φ_{DP} (white oval) is probably precipitation.

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.

The Clutter Filter Control (CFC) and Elevation Segments

The Clutter Filter Control (CFC) product is useful to understand clutter filtering in general, and CMD performance in particular. The CFC product is generated each volume scan, with one product built for each clutter elevation segment.

There are five elevation segments, bounded by the angles in blue in Figure 3-42. These clutter elevation segments exist to improve the vertical resolution of clutter filtering. For example, the filtering applied to bins for the 4.5° elevation data shouldn't also be applied to the same bins for the 0.5° elevation. It isn't important to remember the specific angles that define these elevation segments, just that there are five of them. For each of these, there is only one CFC product built from the lowest elevation in that segment.

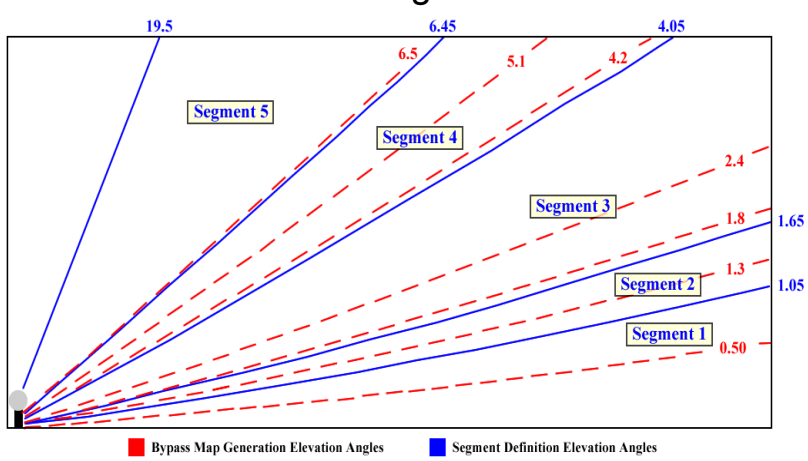

Figure 3-42. A CFC product is build from the lowest elevation in each of the five elevation segments (blue lines).

For each elevation segment, the CFC displays the type of clutter filtering that has been invoked for the lowest elevation in that segment. The area with

a red/yellow combination are where the Bypass Map is in control (see Fig. 3-43). That means there is bin by bin identification of clutter, with red for clutter (to be filtered) and yellow for no clutter (to be left alone). The blue area defines All Bins filtering, which means every bin is going to filtered.

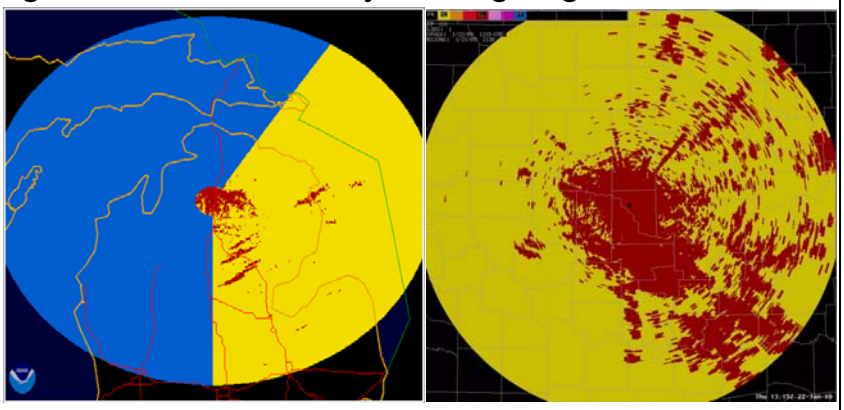

Figure 3-43. On a CFC display, red areas indicate clutter to be filtered, yellow areas indicate no clutter and blue areas indicate All Bins filtering.

The CFC on the left in Figure 3-43 shows a clutter filtering scheme that has All Bins (blue) in control over a large area from south through northeast. Outside of the All Bins area, the Bypass Map is on control, with red bins identified as containing clutter and yellow bins identified as not containing clutter.

The CFC on the right in Figure 3-43 shows a clutter filtering scheme that has the Bypass Map in control everywhere.

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

CMD Implementation

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 Contig-

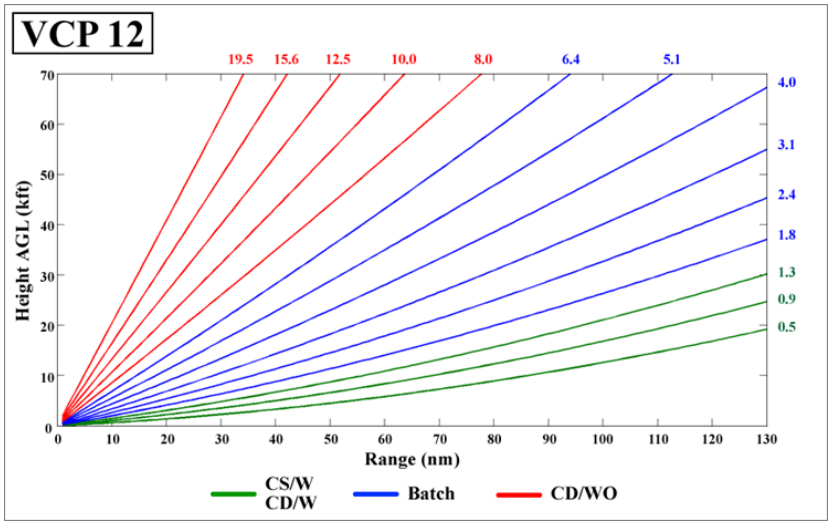

Figure 3-44. CMD builds a different map for each Contiguous Surveillance (CS) and Contiguous Doppler (CD) map.

uous Doppler (CD). CMD builds a different map for each of these two rotations (green curves in Figure 3-44). For the remaining elevations in the volume scan, CMD builds a new map each time.

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

CMD 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 Figure 3-45) to create a local version of this file, “SUPPRESSION ONLY W/ BYPASS MAP”. Whether it is the Default file, or a locally defined version, the key point is to

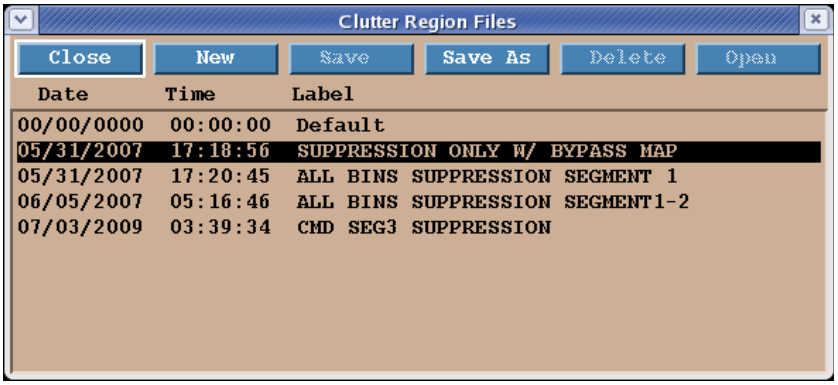

Figure 3-45. For CMD to build maps for every elevation, the Bypass Map must be in control for the entire display for all elevations.

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 effect on the amount of power removed by GMAP when the filtering is performed.

CMD Summary

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. The Clutter Filter Control product is available to help visualize where CMD has identified clutter. However, CFC products are limited to one per elevation segment, only showing a subset of the maps actually built by CMD.

The number of pulses per radial impacts CMD's performance, especially with respect to discriminating clutter from weather with little movement, like stratiform rain. The faster VCPs, 12, 212, and 121, have the lowest number of pulses per radial. The impact of faster VCPs on data quality will be explored later in Topic 3.

The Gaussian Model Adaptive Processing (GMAP) Algorithm

Now that clutter has been identified by CMD, it is time to perform the 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.

Weather and clutter signals have different characteristics. A clutter signal (blue green spike in Figure 3-46) has high power, is centered at zero velocity and has a narrow spectrum width. A weather signal (broader pink area) will have varying power, velocity and spectrum width.

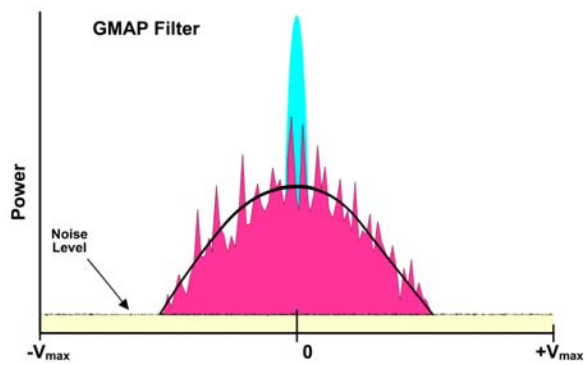

Figure 3-46. Clutter signals have high power and are centered around zero (blue area), whereas weather signals have varying power, velocity and spectrum width (pink area).

This difference between clutter signals and weather signals can be used to filter out the clutter signal with minimal damage to the weather signal. Another aspect that is part of GMAP's design is that both clutter and weather signals can be well represented by Gaussian curves. This is a simplified graphic. For example, weather signals are not usually centered at zero velocity.

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

Once the width for signal removal is determined, filtering is applied to all of the signal within that width. In this case, both the clutter and weather signal within this interval will be removed. However, GMAP has the ability to rebuild the weather signal that was lost (see Fig. 3-47). 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.

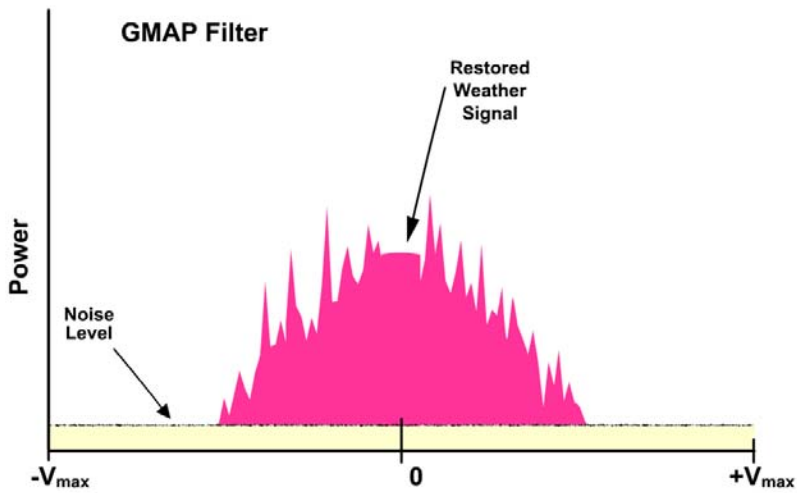

Figure 3-47. With sufficient remaining data, GMAP can rebuild the weather signal after removing power from the signal near zero.

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 in Topic 3.

GMAP Radar Example: No Weather Present

As an example of GMAP performance, we start with no weather with either of these images in Figure 3-48. 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.

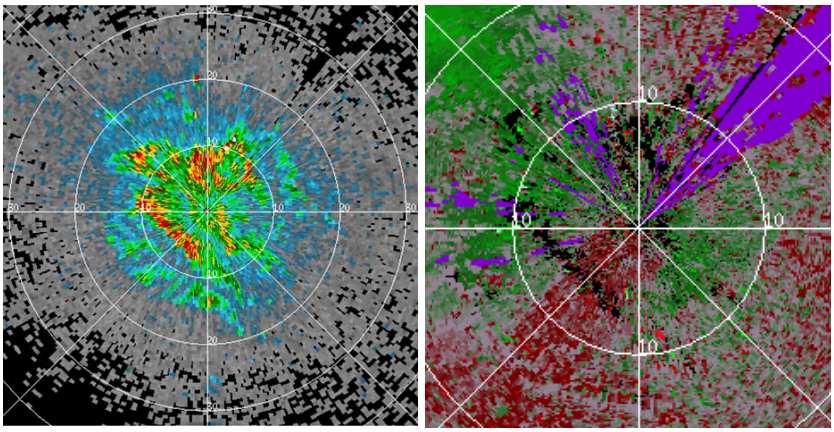

Figure 3-48. Reflectivity (Z, left) without clutter filtering and Velocity (V, right) without clutter filtering

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

At a later time, a squall line has passed through the area. In Figure 3-49, the reflectivity on the left and the velocity product on the right show the squall line. Compared to the reflectivity, the velocity product has been zoomed in.

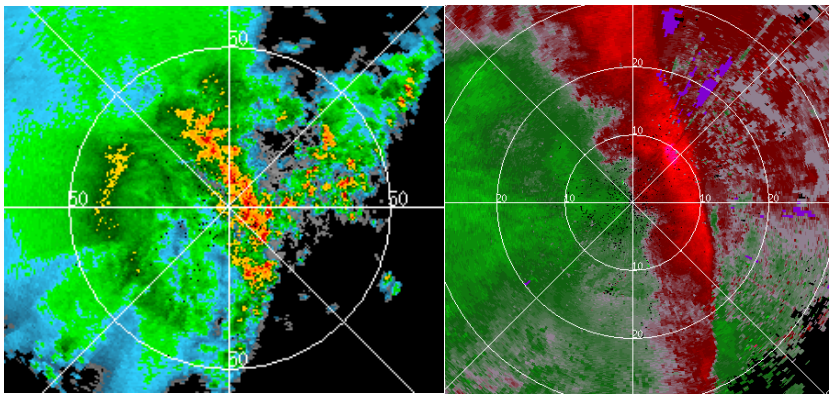

Figure 3-49. Reflectivity (Z, left) and velocity (V, right) data from the same area shown in Figure 3-48, this time with a squall line passing through.

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

CMD and GMAP Summary

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

The performance of both CMD and GMAP are impacted by the number of pulses per radial. The faster VCPs, 12, 212, and 121, have the lowest

number of pulses per radial. For these VCPs, it is more difficult for CMD to discriminate clutter from weather with little movement, like stratiform rain. For these VCPs, it is more difficult for GMAP to rebuild the weather signal after the clutter has been removed.

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

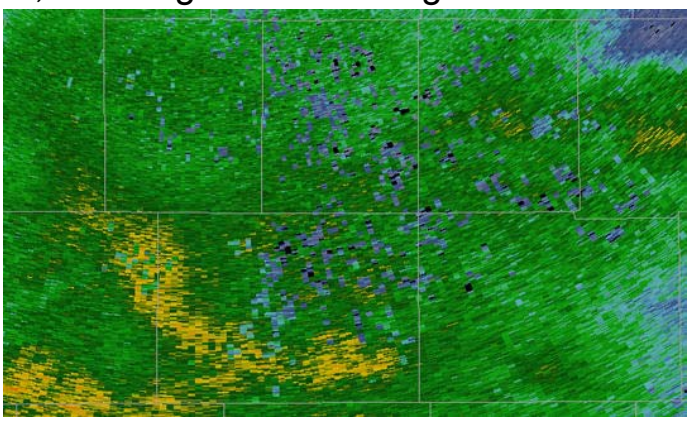

Figure 3-50. Invoking VCP 12 in stratiform rain makes it harder for CMD and GMAP and results in numerous gates of “lost” data.

Range Unfolding of Velocity and Spectrum Width

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

In Figure 3-51, notice the discontinuity in the velocity data at a fixed range. You will often see this on velocity and spectrum width as long as there is sufficient areal coverage of echo. This range is the R_{max} for the Contiguous Doppler/high PRF mode that is used for velocity data collection. It's showing the end of the first trip and the beginning of the second trip. The question is "how did any velocity data end up being assigned beyond the first trip?"

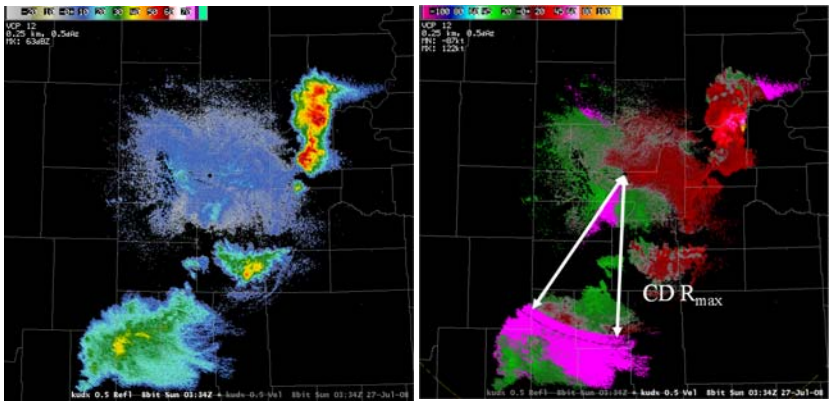

Figure 3-51. Reflectivity (Z, left) and corresponding velocity (V, right) images showing

Range Folding Rarely on Reflectivity

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

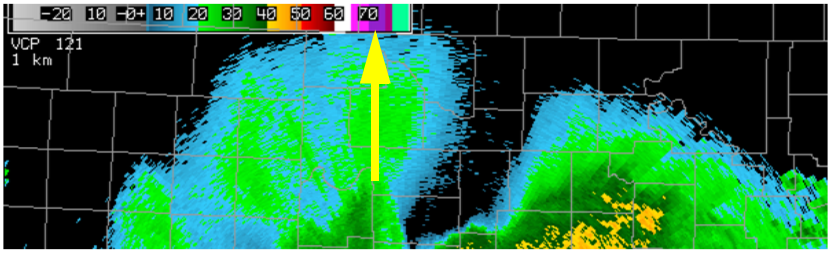

Figure 3-52. In reflectivity data, the color purple represents high dBZ values, not range folding as in velocity data.

The color purple on reflectivity products indicates very high dBZ values, not RF (see Fig. 3-52).

Legacy Range Unfolding Algorithm

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

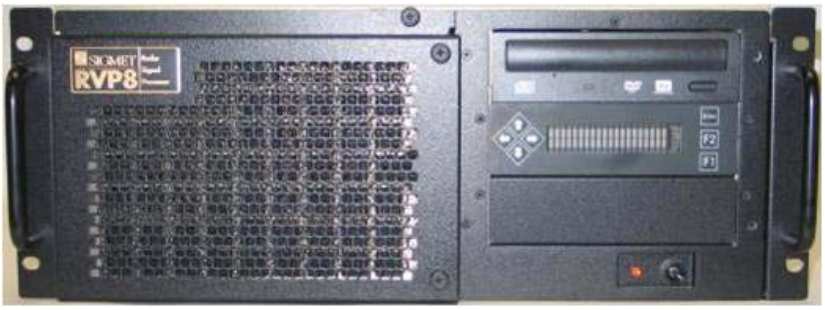

Figure 3-53. A photo of the RDA signal processor, where the Legacy Range Unfolding algorithm is run.

Non-overlaid Echoes Case

In the non-overlaid echoes case, echoes are positioned along a radial such that when the CD mode is employed and R_{max} is short, there are no echoes occupying the same apparent range in the first trip.

After the low PRF (CS) data have been collected, each target's range and returned power is known. However, the velocity estimates would be based on a V_{max} of 16 kts (first guess velocities from -16 kts to +16 kts) and are ***not*** used for velocity product generation (see Fig. 3-54).

Step 1: True Range and Possible Ranges

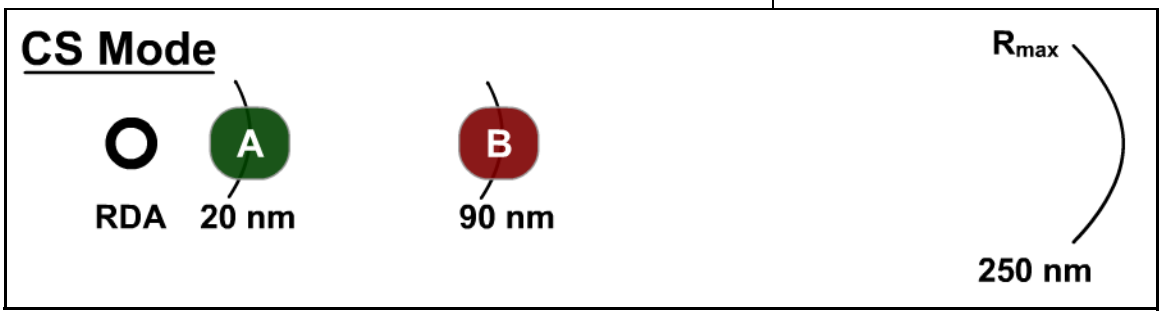

Figure 3-54. Looking along a radial, the range and power of each target is detected while in CS mode.

For this radial, the algorithm computes the ***apparent*** range that each target will have when the high PRF (CD) mode is employed. In this case, the CD PRF will result in an R_{max} of 60 nm. The apparent range for both targets will be within the first trip. See Figure 3-55.

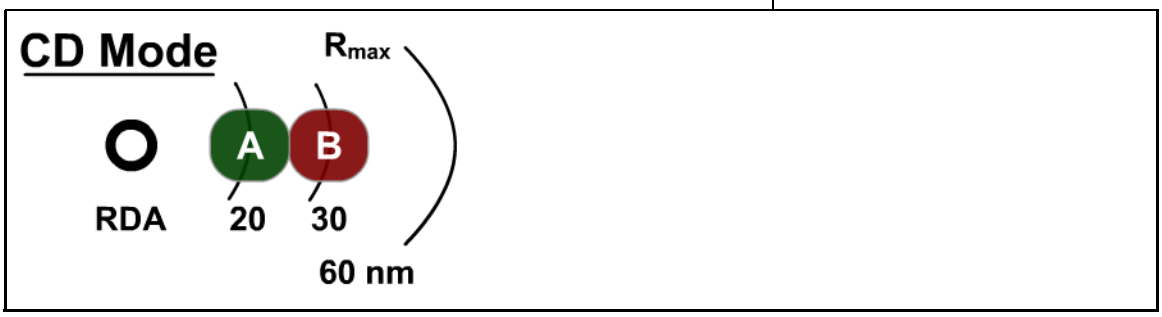

Figure 3-55. The computed ***apparent*** ranges in CD mode of each target. The ranges are calculated prior to collecting the CD data.

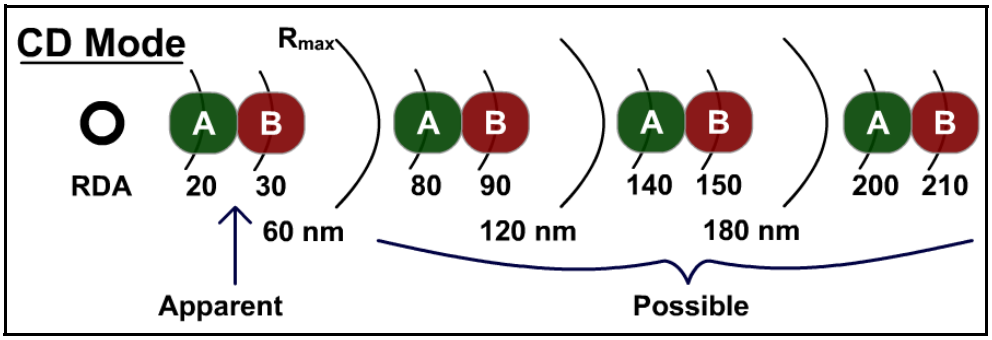

Figure 3-56. The ***apparent and possible*** ranges in CD mode of each target. The ranges are determined prior to collecting the CD data.

The algorithm also computes *all possible ranges* within subsequent trips for each target. Therefore, when the Doppler data are collected, all possible target ranges will have been previously identified (see Fig. 3-56).

Step 2: Power Comparison

The algorithm determines if there will be any echoes "folded" into the same range bin when the CD waveform is employed. If so, the power returns of the two or more echoes are compared. In this case, there are no echoes folded into the same bin, hence no echo overlay.

Step 3: Unfolding

The high PRF (CD) data are collected. The velocity data will be accurate but, due to the short R_{max}, any echoes beyond R_{max} will be ***folded*** into the first trip. The unfolding step involves checking the CS data (power and range) bin by bin against the CD data (velocity and spectrum width values). The apparent range of a velocity value in the CD data is verified by checking that same range in the CS data (see Fig. 3-57).

Target A is assigned a range of 20 nm, since there was something at that range in the CS data. Since the CS data has no target at 30 nm, the possible ranges of target B are next compared to the CS data (Figure 3-58).

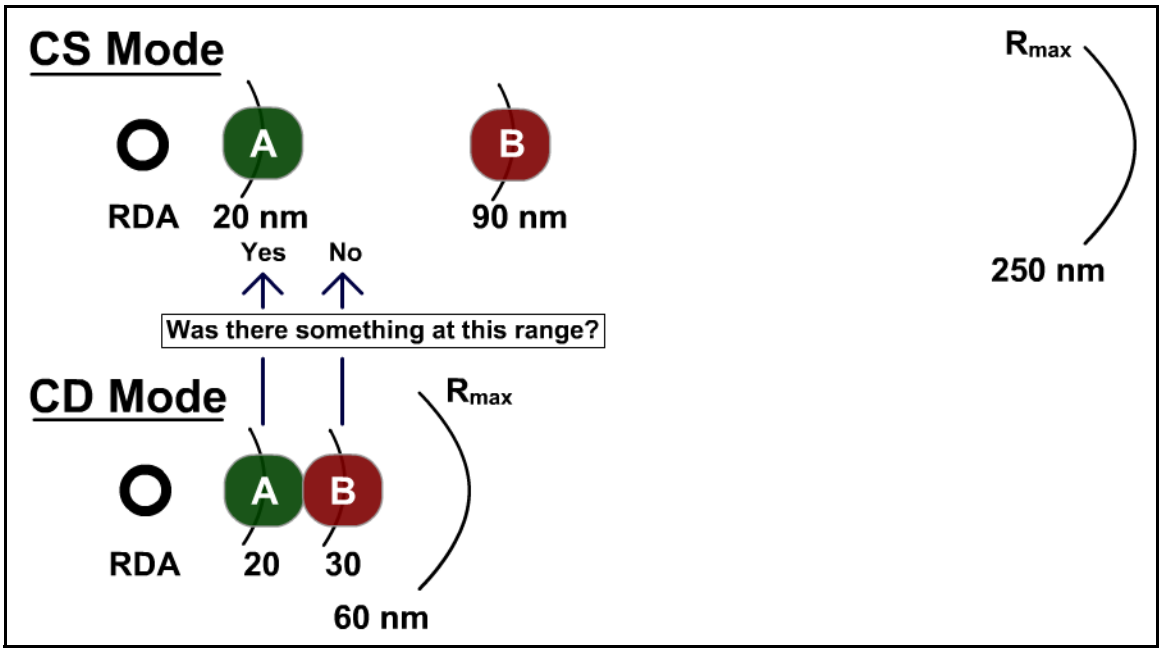

Figure 3-57. The apparent range of a velocity value in the CD data is compared to that range in the CS data.

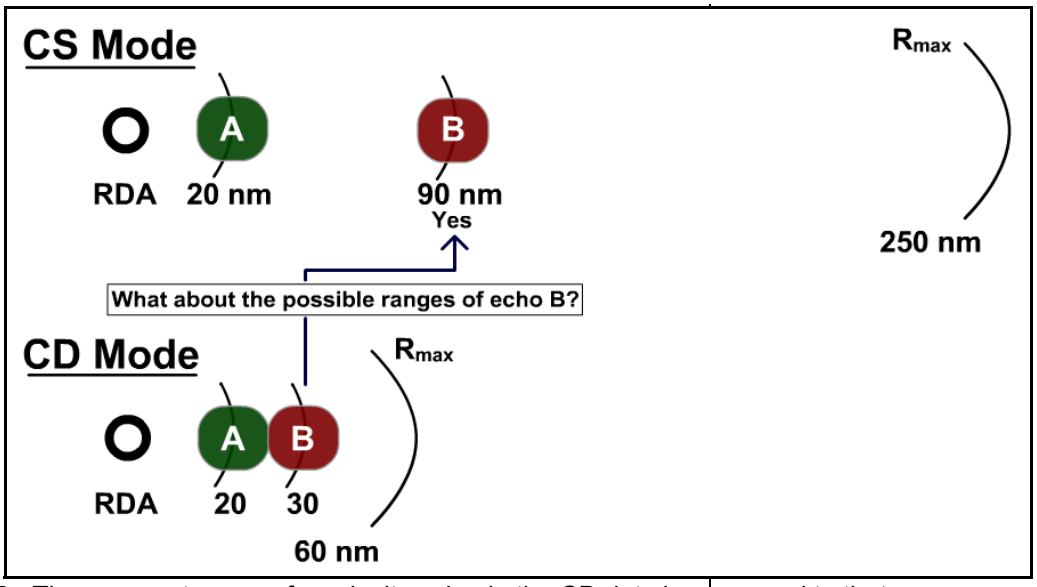

Figure 3-58. The apparent range of a velocity value in the CD data is compared to that same range in the CS data.

Echoes are Overlaid

In this case, the target positions along the radial are such that in CD mode there ***will*** be echoes folded into the same range bin.

Step 1: True Range and Possible Ranges

As in the previous case, once the CS data have been collected, power and range information for each target along a radial is known (see Fig. 3-59).

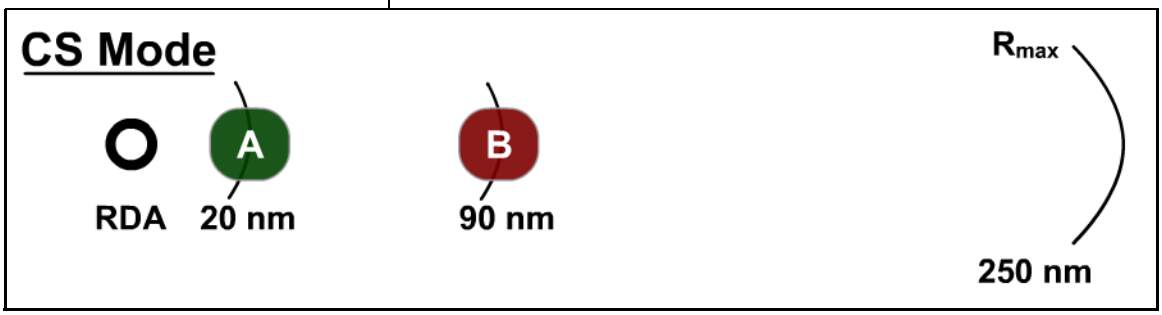

Figure 3-59. Looking along a radial, note the range of each target while is CS mode.

In this case, the apparent ranges of targets A and B are computed based on a CD R_{max} of 70 nm. For A and B with a CD R_{max} = 70 nm, the true ranges of 20 nm and 90 nm would ***both*** have apparent ranges of 20 nm (see Fig. 3-60).

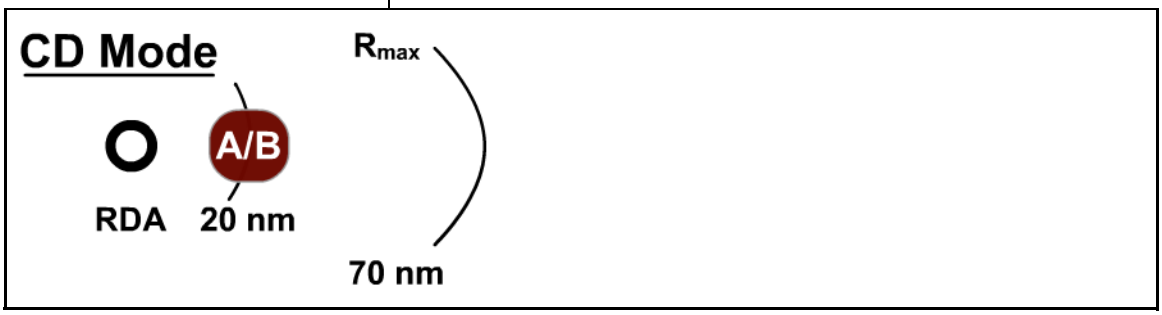

Figure 3-60. The computed apparent ranges in CD mode, calculated prior to collecting the CD data. Note that echoes A and B will be overlaid.

The algorithm also computes other possible ranges within subsequent trips for each target. Therefore, when the CD data are collected, all possible target ranges will have been previously identified (see Fig. 3-61).

Note that when the velocity data are collected, the velocity and spectrum width estimates that are at

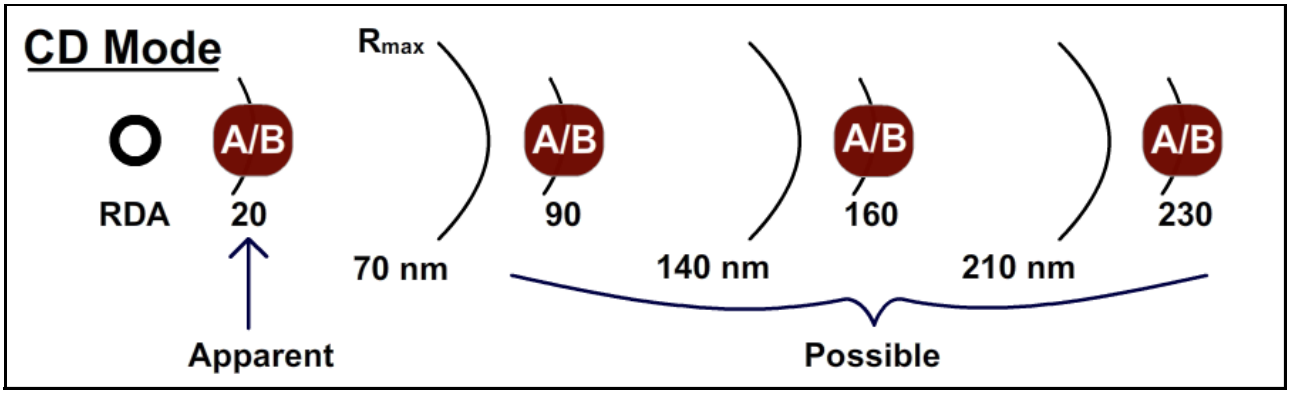

Figure 3-61. The apparent and possible ranges in CD mode of each target. The ranges are determined prior to collecting the CD data. Note the repeated echo overlay.

an apparent range of 20 nm will be composed of the returns from ***both*** targets A and B. Pulses from both of these targets will reach the antenna at the same time. Since velocity estimates are power weighted, velocity data associated with an overlaid echo will be representative of one of the two echoes ***only if*** one of them has returned significantly greater power than the other.

In Figure 3-62, targets A and B are represented as independent signals, but that is for our understanding only. The signal processor analyzes all of the available pulses from that range bin, not knowing that they come from two different targets.

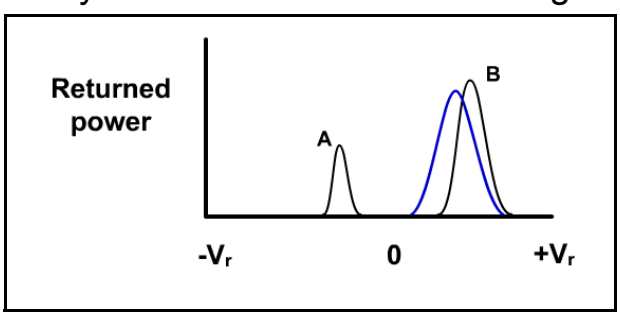

Figure 3-62. Since the pulses from both targets A and B return to the RDA at the same time, they are analyzed as a single target (blue curve).

The blue curve is the base data analysis of that range bin, and the associated velocity estimate is more representative of B than A. Target B has returned significantly higher power than target A.

Since velocity estimates are power weighted, assigning a “blended” velocity estimate to a target that has returned significantly greater power than the other will result in a reasonably accurate estimate. However, “greater power” needs to be quantified.

Step 2: Power Comparisons

From the CS data, the returned power of the two echoes is known. The power ratio (higher to lower) in dB is computed.

$$dB = 10\log\left(\frac{P_{high}}{P_{low}}\right)$$

- If the power ratio exceeds TOVER (an adaptable parameter with a default value of 5 dB), then velocity data will be assigned to the echo returning the higher power, while the other will be assigned RF.
- If the power ratio does not exceed TOVER, ***both*** overlaid echoes will be assigned RF.

Step 3: Unfolding

The CD data are collected. As with the non-overlaid echo example, the apparent and possible ranges of a velocity value in the CD data are compared to the true range in the CS data.

For this example, assume the power ratio of target B over A exceeds TOVER. Since the range of both echoes is known, the blended velocity estimate will be assigned to B at 90 nm, with RF assigned to A at 20 nm (Figure 3-63). Note that the velocity estimate assigned to the range of B is slightly biased by the returns from A.

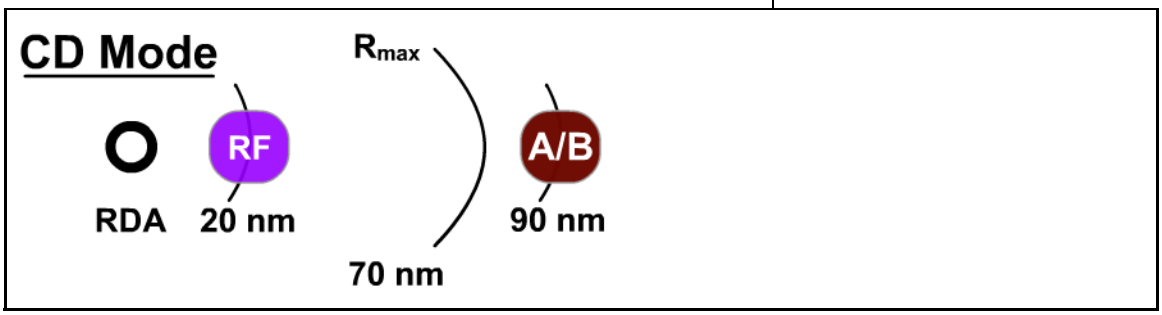

Figure 3-63. For the range bins where echo overlay has occurred, the assignment of velocity data or purple (RF) is determined by the power ratio and TOVER.

Distribution of RF (Purple)

Now for another look at the distribution of RF on the velocity product, given the location of the echoes on the reflectivity. In Figure 3-64, the strongest storms to the northeast are within the first trip, so there is no issue with the availability of the associated velocity data. There is an area of weaker echo to the south and southwest of the radar, some of which falls within the first trip of the velocity data, some of which falls within the second trip.

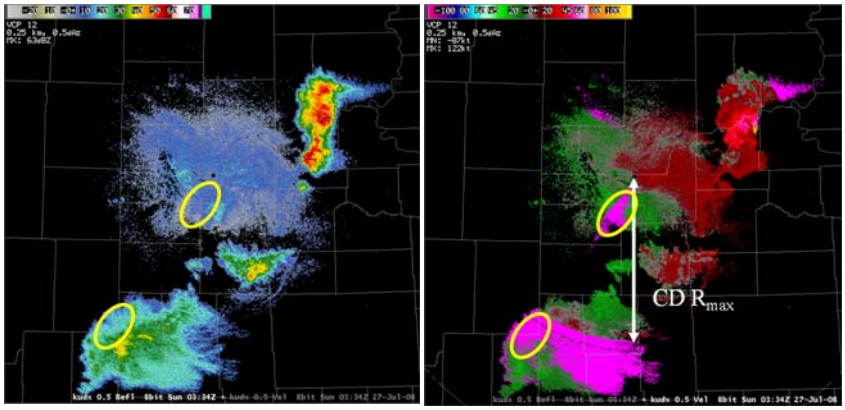

Figure 3-64. Comparing the location of RF data in the velocity (V, right) data compared to the reflectivity (Z, left) data.

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

Legacy Range Unfolding Algorithm Strengths

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

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

Legacy Range Unfolding Algorithm Limitations

The Legacy Range Unfolding algorithm is going to have the greatest limitation on its performance when there is extensive echo coverage aligned along radials, through both the first and second trips in the CD data. This situation maximizes the

echo overlay along the radial. Since the Legacy Range Unfolding algorithm can assign velocity and spectrum width to, at best, one of the overlaid echoes, the result is a lot of RF data. In Figure 3-65, the echo coverage is extensive for both the first and second trips to the north and the south of the radar. There is extensive RF coverage in the second trip because the TOVER requirement was met by the echoes in the first trip.

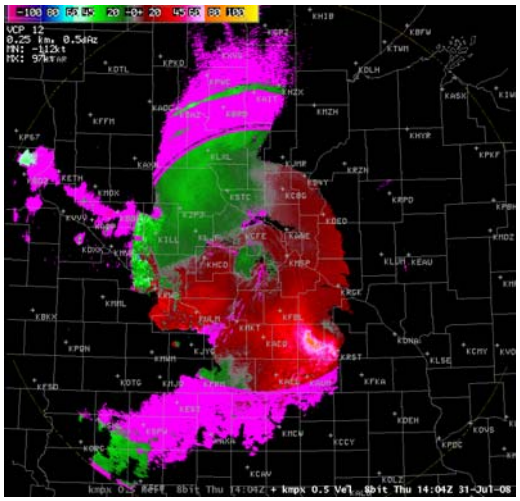

Figure 3-65. Velocity image with extensive RF in the second trip because the TOVER requirement was met by the echoes in the first trip.

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

SZ-2 Range Unfolding Algorithm

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

SZ-2 is available only for the Split Cut elevations of the three SZ-2 VCPs, 211, 221, and 212. The “2” as the first digit in the VCP name indicates that it is an SZ-2 VCP. The second two digits in the name

tells you the structure of the VCP with respect to the angles used and the update rate. For example, VCP 212 is just like VCP 12 except that SZ-2 is used to range unfold velocity and spectrum width on the Split Cuts.

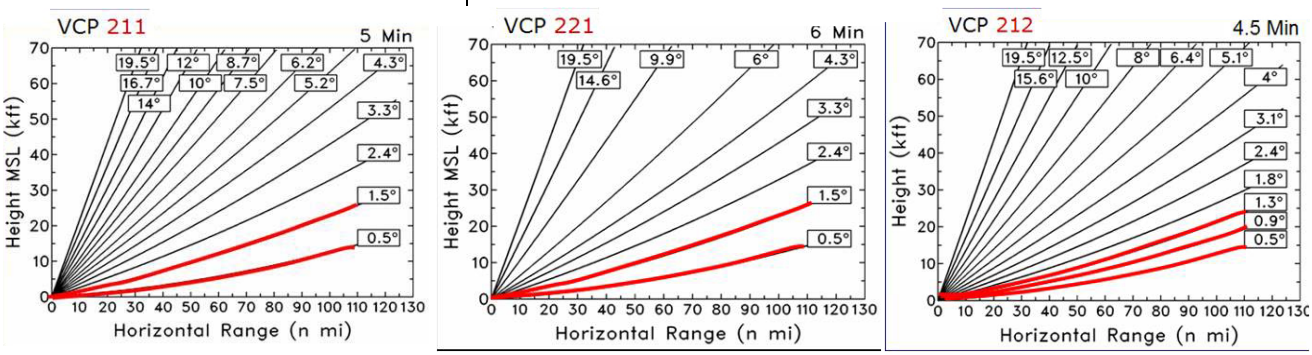

Figure 3-66. SZ-2 Range Unfolding is only available for the Split Cut elevations (shown in red) for VCPs 211, 221, and 212.

Legacy vs. SZ-2

The SZ-2 Range Unfolding algorithm is particularly effective with widespread echo coverage. In Figure 3-67, the Legacy Range Unfolding algorithm (VCP 12) is used for the data on the left, with SZ-2 (VCP 212) on the right. There is a dramatic improvement in the availability of velocity and spectrum width when VCP 212 is used.

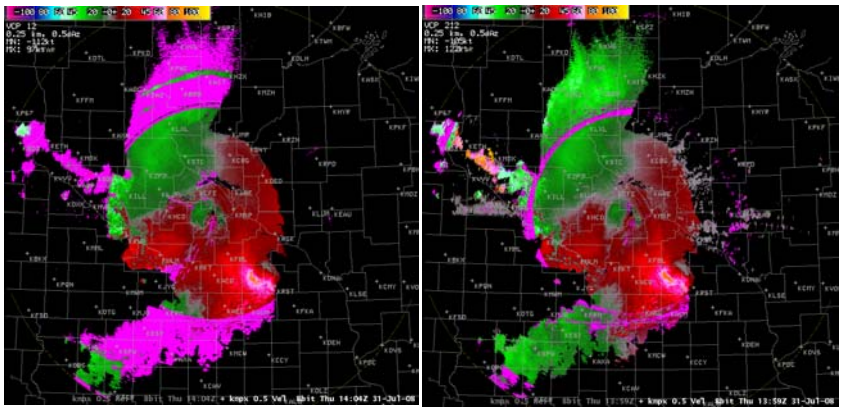

Figure 3-67. Two velocity images: one using the Legacy Range Unfolding algorithm (left) and the other using the SZ-2 algorithm (right).

This example also demonstrates the key difference between the two range unfolding techniques currently available when there is echo overlay, i.e.

returns from multiple trips in the CD data. When there is echo overlay, the Legacy Range Unfolding algorithm can recover, at best, one of the echoes. On the other hand, SZ-2 can usually recover both overlaid echoes.

Figure 3-68 shows another example of the benefits of the SZ-2 VCPs. This case was taken from two different, but nearby radars, so there is some echo displacement. There was a supercell in the second trip for both radars, with Legacy Range Unfolding on the left and SZ-2 Range Unfolding on the right. Notice that SZ-2 can reveal the circulation that is otherwise obscured by RF.

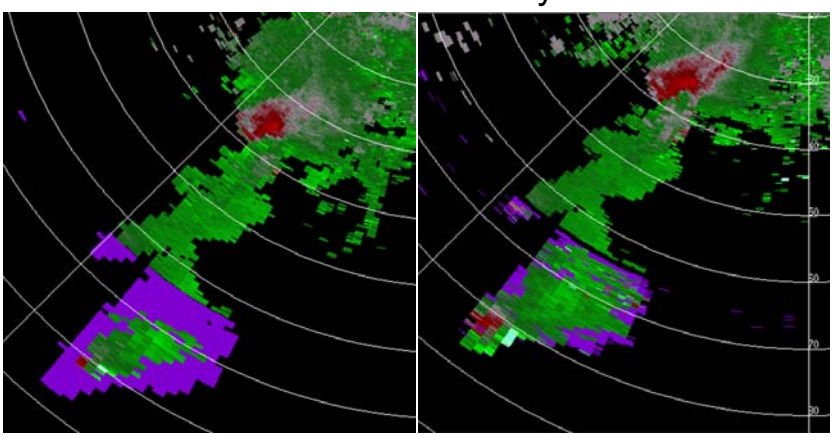

Figure 3-68. The circulation within a nearby supercell is visible in the velocity data using SZ-2 Range Unfolding (right), but not in the data using Legacy Range Unfolding (left).

SZ-2 Technique

When in CD mode, the initial phase of each transmitted pulse is changed in a specific sequence. This allows for much better comparison of pulses from multiple trips. SZ-2 requires exactly 64 Doppler pulses per radial. Though there is clear benefit with the overall increase in availability of velocity data, there are implications of the requirement for exactly 64 pulses per radial. The CD R_{max} will not change from volume scan to volume scan, either automatically or manually. You cannot edit the

Doppler PRF for the Split Cut elevations of the SZ-2 VCPs.

SZ-2 Strengths

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

SZ-2 Limitations

With CMD active, the need for All Bins clutter suppression is nearly zero. SZ-2 is another reason to avoid All Bins suppression, as it can degrade the velocity data when SZ-2 is applied. The RPG software has some steps in place to avoid having All Bins and SZ-2 active at the same time. The following points assume that your RDA configuration is such that CMD is available.

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

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

The other limitation for the use of SZ-2 VCPs is that the R_{max} for the Split Cut elevations does not move from volume scan to volume scan, and the Doppler PRFs cannot be edited. If there is a significant storm near the first trip/second trip RF “ring”, it may be necessary to switch to another VCP. This

could be a different SZ-2 VCP or a non-SZ-2 VCP that allows for editing of the Doppler PRF. You'll see more about these Doppler PRF manual edits later in Topic 3.

Distance Learning Operations Course

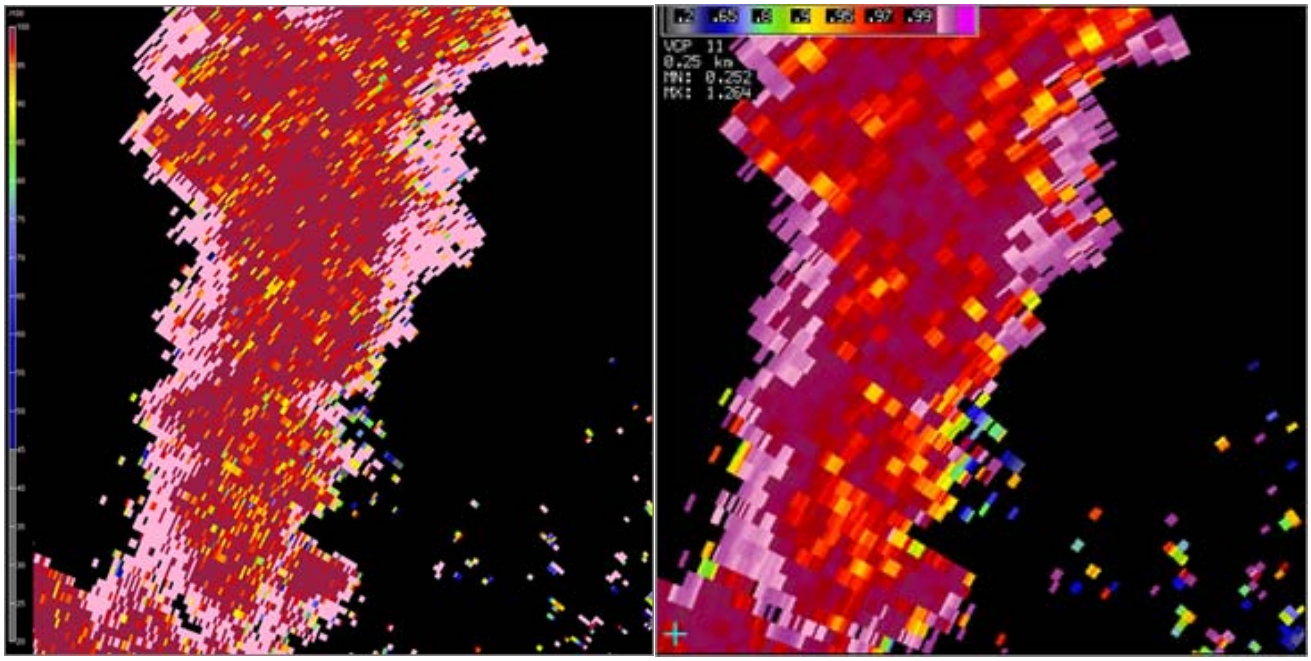

Topic 3, Lesson 5: WSR-88D Dealiasing Velocity and Preprocessing Dual-Pol

Presented by the
Warning Decision Training Branch

Version: 1210

Lesson 5: Dealiasing Velocity and Pre-processing Dual-Pol

After clutter filtering and range unfolding, there is more work to be done to improve the quality of the base data. This lesson first looks at the process of dealiasing the first guess velocity estimates. The dual-pol base data that are sent from the RDA require preprocessing before products can be generated. This lesson describes the need for and the primary tasks of the Dual-Pol Pre-processing algorithm at the RPG.

A Note on Terminology

Though often used interchangeably, range gate and range bin have different meanings. Here, range gate will be used when describing the appearance on WSR-88D products, and range bin will be used when describing algorithm functions.

Objectives

1. Identify the purpose, strengths and limitations of the following techniques to dealias velocity data.
   a. Velocity Dealiasing
   b. VCP 121: Multiple PRF Dealiasing Algorithm
2. Identify the primary tasks of the Dual-Pol Pre-processor at the RPG.
3. Identify data artifacts due to preprocessing and which Dual-Pol products are affected.

Improperly Dealiased Velocities

Lesson 4 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 Lesson 2, Base Data Generation, had a first guess velocity of +30

kts, while the correct velocity was -90 kts (see Fig. 3-69).

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°.

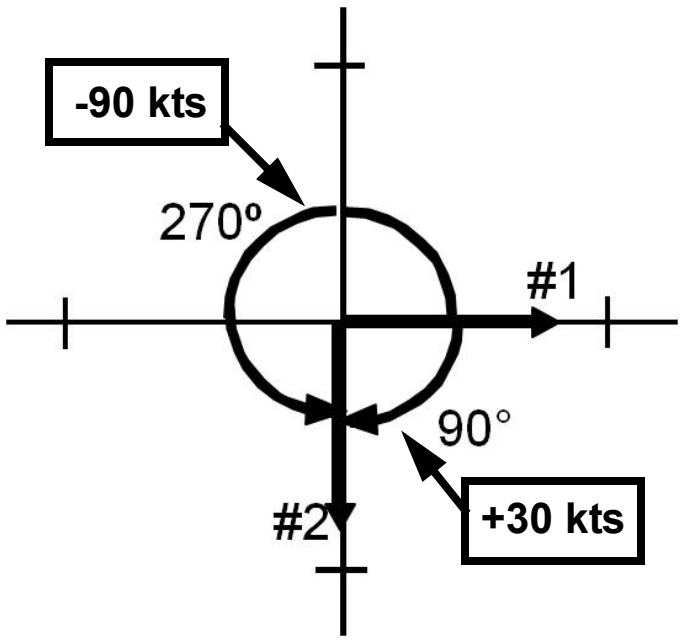

Figure 3-69. Two possible phase shifts between two phasors.

Product Characteristics

Before looking at the algorithm that "de-aliases" velocities, consider 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 surrounding data often occur. Further out, blocks or wedges of values opposite in direction from the surrounding data are generally larger. In some cases where shears appear along an azimuth, it is difficult to determine if the shears are meteorological or the result of improperly dealiased velocities.

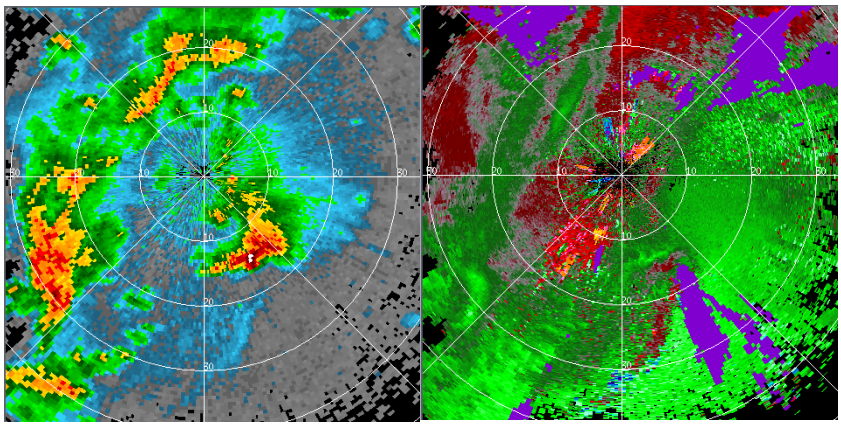

Figure 3-70. Velocity (V, right) image showing improperly dealiased velocities. The corresponding reflectivity (Z) image is shown on the left.

Figure 3-70 shows 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.

Figure 3-71 is an example of improperly dealiased velocities resulting in at least 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.

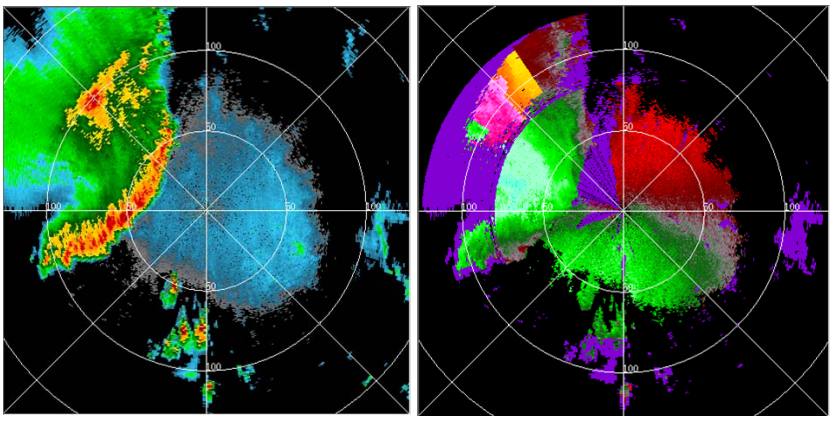

Figure 3-71. Example of dealiasing failures in velocity data (V, right) with the corresponding reflectivity image (Z, left).

Velocity Dealiasing Algorithm

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).

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 the 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 so dealiasing failures can be mitigated.

Environmental Winds Table

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, is a realistic representation of the state of the winds aloft.

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

Environmental Data Editor Window

When initially accessing Environmental Data, this window provides multiple types of data (see Fig. 3-72). The local winds aloft are 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), formerly known as the Rapid Update Cycle (RUC) model. The VAD is an RPG algorithm that uses WSR-88D data to generate winds at a series of heights. Though the RUC has been recently replaced by the RAP, “Model Update” on this RPG window is referring to the RAP.

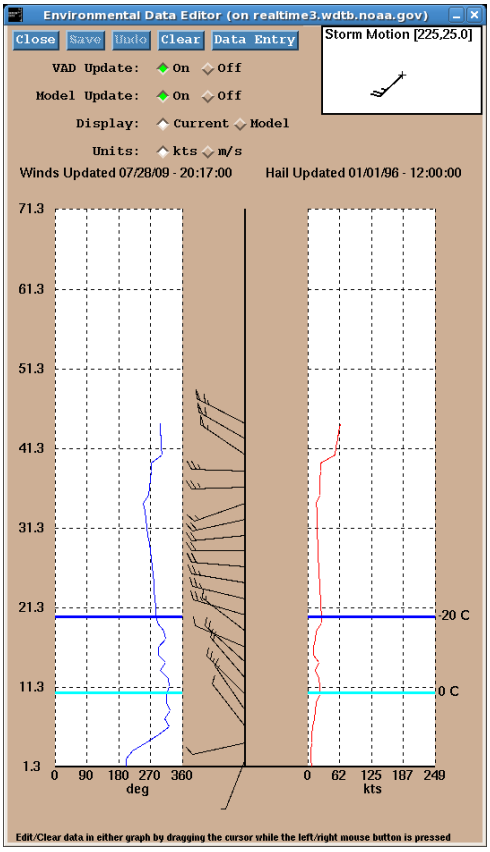

Figure 3-72. The Environmental Data Editor window includes information about local winds aloft, default storm motion, and hail temperature heights.

The Environmental Data Editor window also has the default storm motion and the hail temperature heights. The use of both of these data sources will be discussed in DLOC Topic 5.

Managing the Environmental Winds Table (EWT)

Figure 3-73 offers 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.

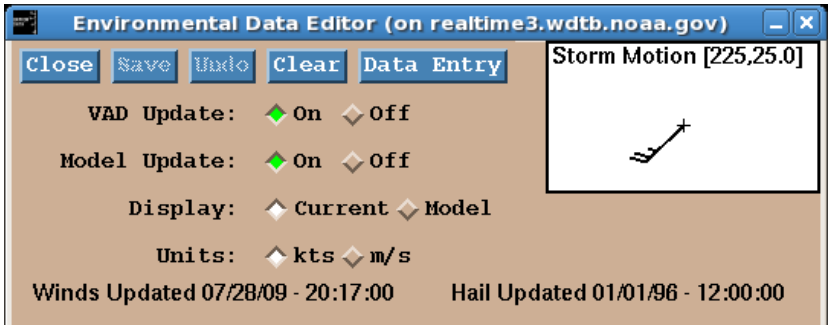

Figure 3-73. A closer look at the top of the Environmental Data Editor Window and the buttons that control the inputs to the Environmental Winds Table.

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 unreliable, Model

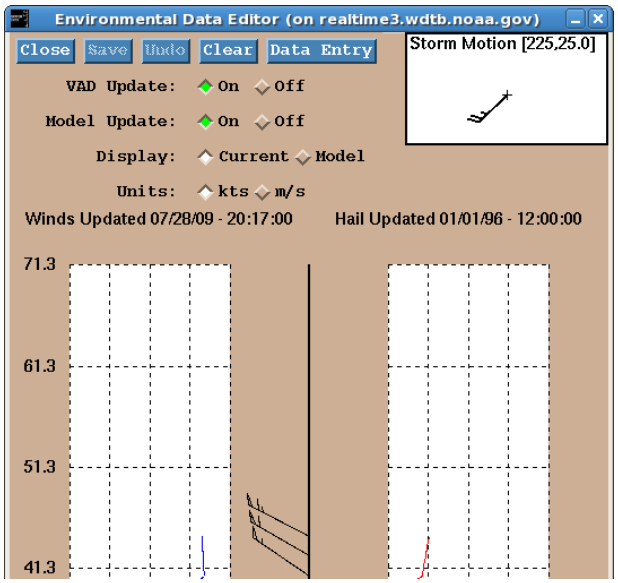

Figure 3-74. The Display button allows the user to view the current state of the EWT and model data to verify reliability of the data.

Update can be set to "Off."

Velocity Dealiasing Algorithm Strengths

The Velocity Dealiasing algorithm is designed to provide the best possible velocity data, primarily 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.

Velocity Dealiasing Algorithm Limitations

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. Dealiasing failures can also contaminate the results of the RPG algorithms, such as false circu-

lations 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.

Velocity Dealiasing Algorithm Considerations

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 usually not 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.

Multiple PRF Dealiasing Algorithm (MPDA)

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 (see Fig. 3-75) 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 three CD rotations. Each of these CD rotations uses a different Doppler PRF.

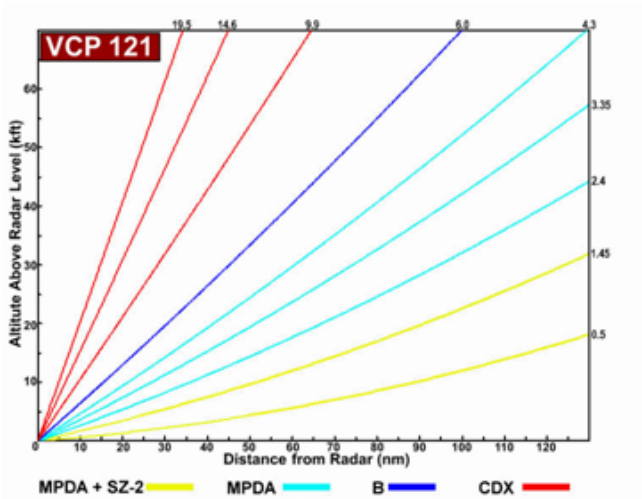

Figure 3-75. the Multiple PRF Dealiasing Algorithm (MPDA) is used for VCP 121 only to mitigate range folding and improperly dealiased velocities.

VCP 121

Figure 3-76 is 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.

VOLUME COVERAGE PATTERN 121
SCAN STRATEGY MPDA/SZ2 SHORT PULSE

Scan				Surveillance		Doppler PRF No.				
Elevation (deg)	AZ Rate (deg/sec)	Period (sec)	WF Type	PRF No.	No Pulses	4 No. Pulses	5 No. Pulses	6 No. Pulses	7 No. Pulses	8 No. Pulses
0.5	18.677	19.28	SZCS	1	17	-	-	-	-	-
0.5	19.754	18.22	SZCD	8	-	43	51	55	59	64
0.5	27.400	13.14	CD	6	-	31	37	40	43	46
0.5	21.401	16.82	CD	4	-	40	47	51	55	59
1.45	19.842	18.14	SZCS	1	16	-	-	-	-	-
1.45	19.754	18.22	SZCD	8	-	43	51	55	59	64
1.45	27.400	13.14	CD	6	-	31	37	40	43	46
1.45	21.401	16.82	CD	4	-	40	47	51	55	59
2.4	19.205	18.75	B	1,8	6	27	32	34	37	40
2.4	27.400	13.14	CD	6	-	31	37	40	43	46
2.4	21.401	16.82	CD	4	-	40	47	51	55	59
3.35	21.599	16.67	B	2,8	6	28	33	35	38	40
3.35	27.400	13.14	CD	6	-	31	37	40	43	46
3.35	21.401	16.82	CD	4	-	40	47	51	55	59
4.3	16.304	22.08	B	2,4	6	40	48	52	56	61
4.3	29.498	12.20	CD	7		29	34	37	40	43
6.0	20.204	17.82	B	3,5	6	34	40	43	47	51
9.9	29.498	12.20	CD	7	-	28	34	37	40	43
14.6	29.795	12.08	CD	8	-	28	33	36	39	43
19.5	29.795	12.08	CD	8	-	28	33	36	39	43

Figure 3-76. A table of VCP 121.

Since VCP 121 has additional rotations for several elevations, it has the greatest number of total rota-

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

Strengths of the Multiple PRF Dealiasing Algorithm, VCP 121

MPDA, or 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 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 the Multiple PRF Dealiasing Algorithm, VCP 121

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.

Dual-Pol Preprocessor

The RPG Dual-Pol Preprocessor prepares the Dual-Pol base data for base product generation and for input into the RPG Dual-Pol algorithms: the Hydrometeor Classification Algorithm (HCA), the Melting Layer Detection Algorithm (MLDA), and the Quantitative Precipitation Estimation Algorithm (QPE).

The “raw” Dual-Pol base data has a 0.5° azimuthal resolution. The data are noisy in appearance, gen-

erally too noisy for human interpretation; certainly too noisy for algorithm input. For any given range bin, the 0.5° Dual-Pol base data from the RDA are first recombined to 1.0°.

The Preprocessor first smoothes reflectivity (Z), differential reflectivity (ZDR), correlation coefficient (CC) & differential phase (Φ_{DP}) data along the radial. The smoothed Z data are only used for input to the Dual-Pol RPG algorithms. There is no change to the Z values used to generate all the legacy reflectivity-based products.

The remaining tasks for the Preprocessor are computing the Specific Differential Phase (KDP) values and something called Attenuation Correction.

Preprocessing for Differential Reflectivity (ZDR)

On the left in Figure 3-77 is raw Differential Reflectivity, ZDR, from the RDA, at 0.5° azimuthal resolution and 0.25 km range resolution. This image is from GR Analyst, showing the raw Level II data. 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.

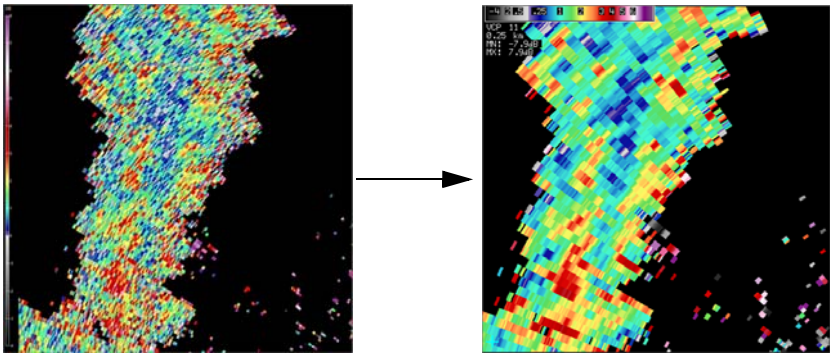

Figure 3-77. Differential reflectivity (ZDR) images from before (left) and after (right) Preprocessor smoothing.

Preprocessing for Correlation Coefficient (CC)

Figure 3-78 shows a similar comparison for correlation coefficient, CC. Raw CC is on the left and the recombined and smoothed CC that becomes a product in AWIPS is on the right. As with ZDR, the Preprocessor smoothing technique applies a linear average to a segment (of varying length) of data along the radial. This average value is then assigned to the original range bin, which is at the center of the segment.

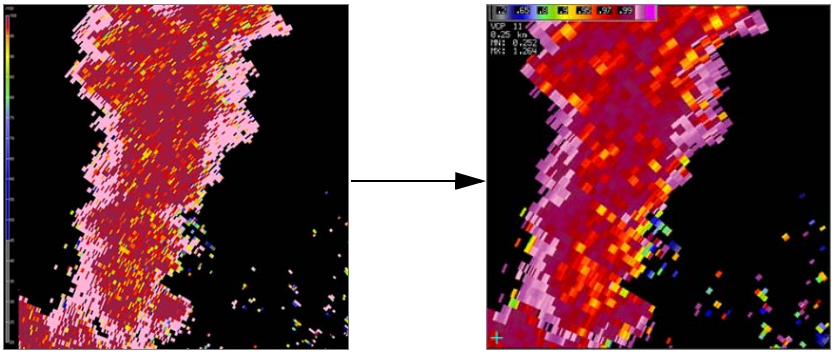

Figure 3-78. Raw correlation coefficient (CC, left) data is smoothed by the Preprocessor to generate the product seen in AWIPS (right).

Preprocessing and Differential Phase (Φ_{DP})

As with ZDR and CC, the Differential Phase, Φ_{DP}, base data are first recombined, then smoothed. On the left in Figure 3-79 is an example of Φ_{DP} base data, not yet recombined or smoothed. Φ_{DP} is not displayed in AWIPS

Once the Φ_{DP} data have been smoothed, the Preprocessor then calculates specific differential

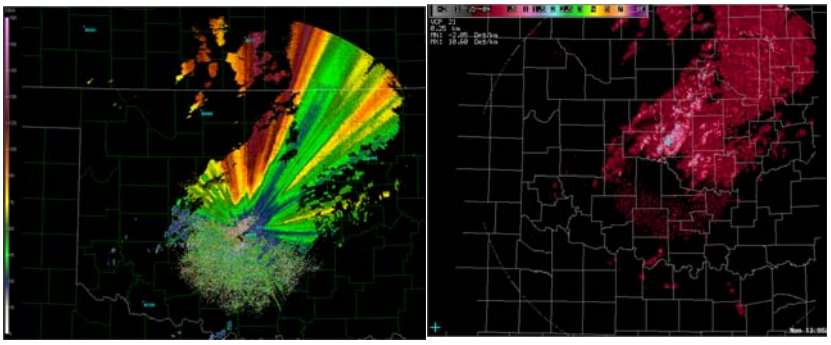

Figure 3-79. Once the Φ_{DP} data (left) have been smoothed, the Preprocessor calculates specific differential phase (KDP) values used to generate the KDP product in AWIPS.

phase, or KDP. The KDP values are then available for generation of the KDP product (Figure 3-79, right) and for input to the Dual-Pol algorithms.

These two images are a good example of why Φ_{DP} is not generated into a base product. It is much more difficult to interpret than KDP.

Dual-Pol Preprocessor Artifact

The next topic is an artifact of the Preprocessor algorithm that shows up mainly on ZDR. Notice the radially-oriented stripes that do not make meteorological sense in Figure 3-80.

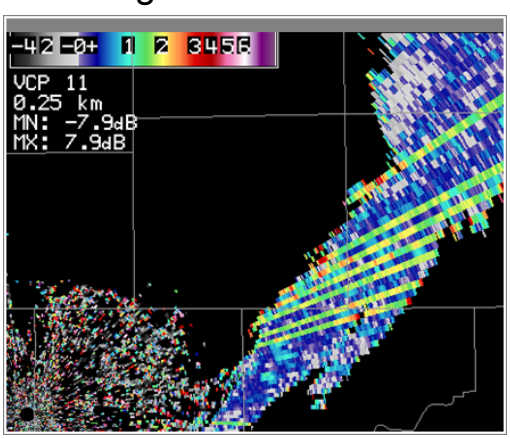

Figure 3-80. The radially-oriented stripes in the differential reflectivity (ZDR) image don't make much sense meteorologically.

First, some background information. Recall that Φ_{DP} is a phase value from 0° to 360° which accumulates down radial (see Fig. 3-81). Φ_{DP} has the

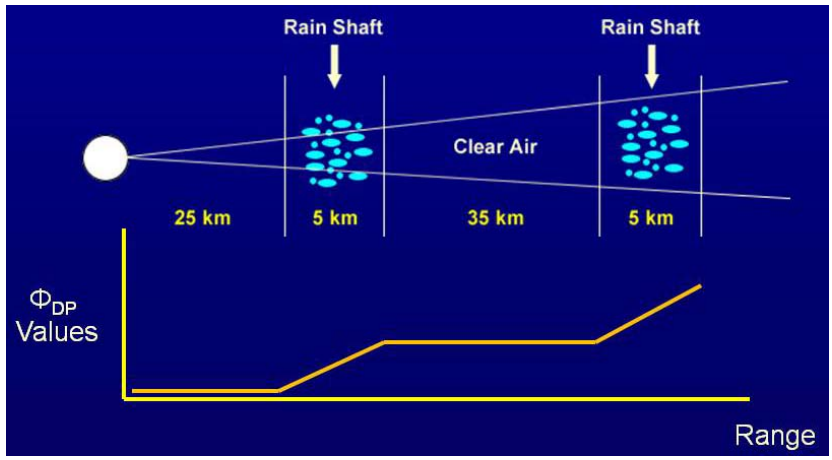

Figure 3-81. Recall that Φ_{DP} accumulates down radial.

greatest increase with range where the beam encounters large amounts of liquid water.

The “ZDR stripes” data artifact is related to how Φ_{DP} is processed along each radial, specifically in between segments of “weather”. Since the Preprocessor is an algorithm, its ability to judge what is or is not weather is limited. “Weather”, as defined by the Preprocessor, resides in bins where CC>0.85. In Figure 3-81, there are two segments along the radial that would be tagged as rain, because of their associated high CC values.

The graphic of Φ_{DP} increasing with range shows how Φ_{DP} ideally behaves in “clear air”, but there are exceptions. Clear air returns can sometimes result in CC>0.85 and thus be Preprocessed as weather.

There can be bins of data in clear air where CC>0.85 is calculated, and from the standpoint of the Preprocessor that means weather. The risk of CC>0.85 clear air is greatest in residual clutter near the radar.

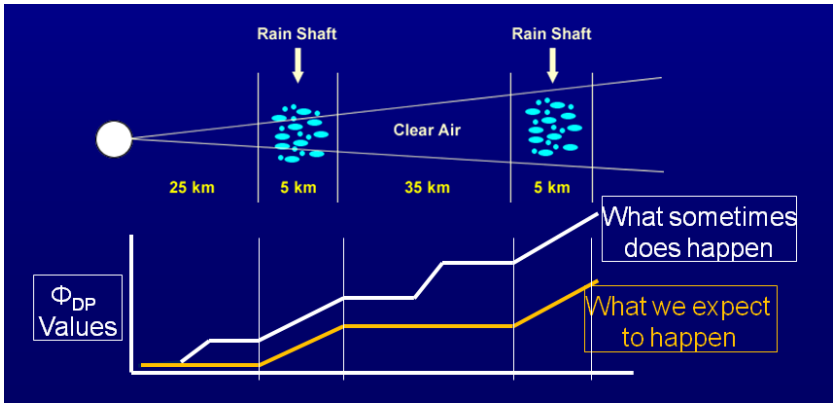

Figure 3-82. Even bins with clear air can have correlation coefficient values less than 0.85, increasing Φ_{DP} values along the way.

The overall result of CC>0.85 in clear air is that Φ_{DP} is then artificially increased down radial (see Fig. 3-82).

The artifact (stripes on ZDR, Figure 3-80) are related to the final task of the Preprocessor, which is Attenuation Correction for Z and ZDR. This correction when applied to Z is only for the Z data used by the Dual-Pol algorithms.

When the radar beam encounters high water content, two things happen: the returned power is attenuated and Φ_{DP} increases. It turns out that the relationship between Φ_{DP} increasing and the amount of attenuation is strong enough that Φ_{DP} can be used as a basis for attenuation correction, and the Preprocessor does just that.

The problem of Φ_{DP} artificially increasing in clear air can result in an overcorrection for attenuation that shows up as radial stripes on ZDR. When overcorrected, the ZDR values are unrealistically high down radial.

These Preprocessor stripes are radially oriented and are transient in time and space. With ZDR, the spikes originate where the beam encounters pre-

cipitation, then persist with range (see Fig. 3-83, bottom right). This is due to Φ_{DP} values that are too high with respect to the ZDR values within the precipitation. Based on this unrealistically high Φ_{DP}, the associated amount of attenuation correction is too high. When applied to ZDR, the attenuation correction falsely increases the ZDR values along these radials.

With KDP, the spikes originate near the RDA, because of noisy Φ_{DP} values (see Fig. 3-83, bottom left). The spikes in KDP usually smooth out if the beam intercepts precipitation. The best smoothing occurs with rain only, since Φ_{DP} behaves best in rain.

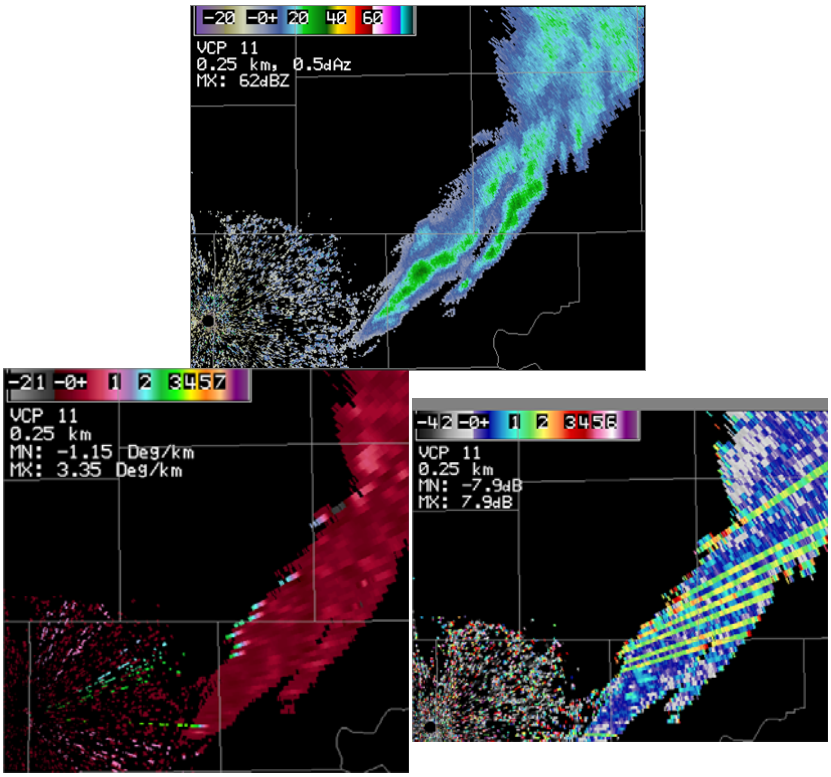

Figure 3-83. Reflectivity (Z, top), specific differential phase (KDP, bottom left), and differential reflectivity (ZDR) images showing the Preprocessor Data Artifact.

Distance Learning Operations Course

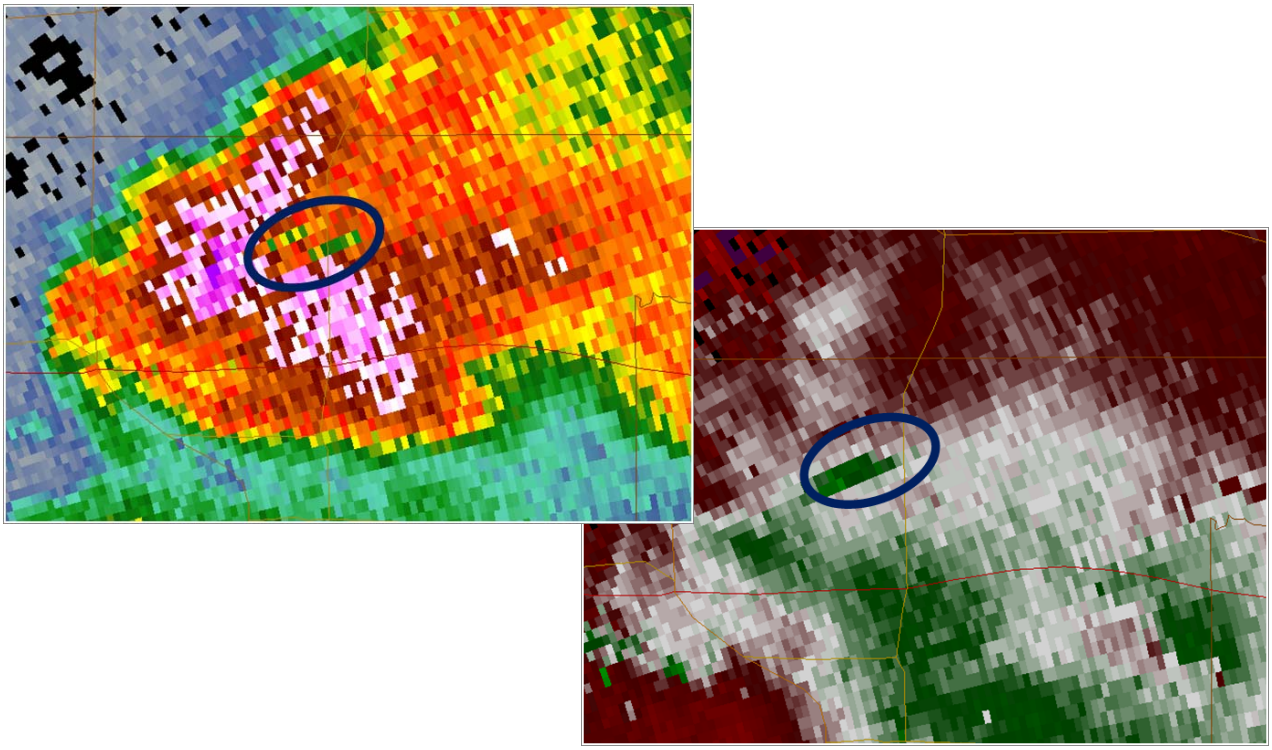

Topic 3, Lesson 6: Operational Impacts of Base Data Quality

Presented by the
Warning Decision Training Branch

Version: 1210

Lesson 6: Operational Impacts of Base Data Quality

This lesson culminates the previous lessons in Topic 3, and the overall message is rather simple: “if your base data ain’t any good, nothin’ else is gonna be.” This lesson brings together all the previous Topic 3 concepts, exploring how you can optimize your base data quality, as well as recognizing the trade-offs between optimal base data versus operational needs for fast updates

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 for minimal RF.
3. Identify the VCPs that allow for manual Doppler PRF adjustment.
4. Identify the purpose of Auto PRF and how it is applied within the VCPs.
5. Identify the impact of differential attenuation, non-uniform beam filling, and depolarization on the Dual-Pol products.
6. Identify the “trade-offs” involved with producing high quality base data vs. meeting operational constraints.

CMD False Detections

CMD is a wonderful automated tool for managing clutter filtering, but it is not perfect. It 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

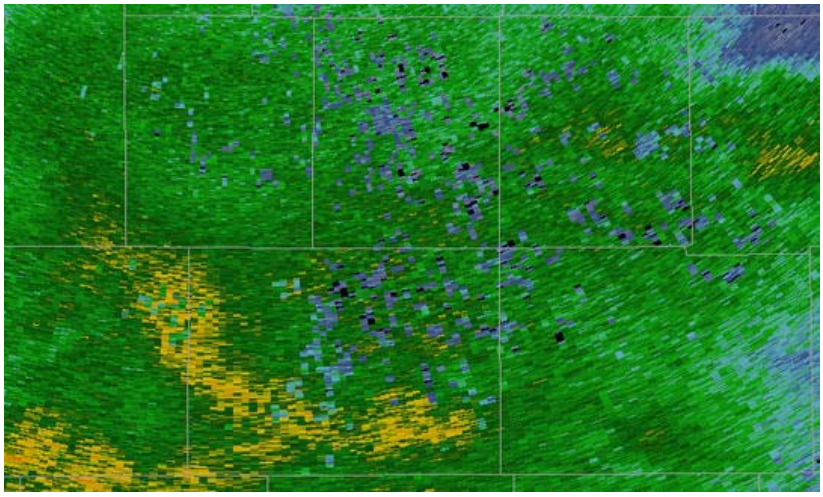

Figure 3-84. Example of noisy reflectivity data from VCP 12.

result in noisy data with sporadic gates of signal removed that are not clutter (see Fig. 3-84).

CMD false detections can even occur in or near convective storms (see Fig. 3-85). 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, versus the best performance of the clutter filtering algorithms: CMD for identification and GMAP for removal.

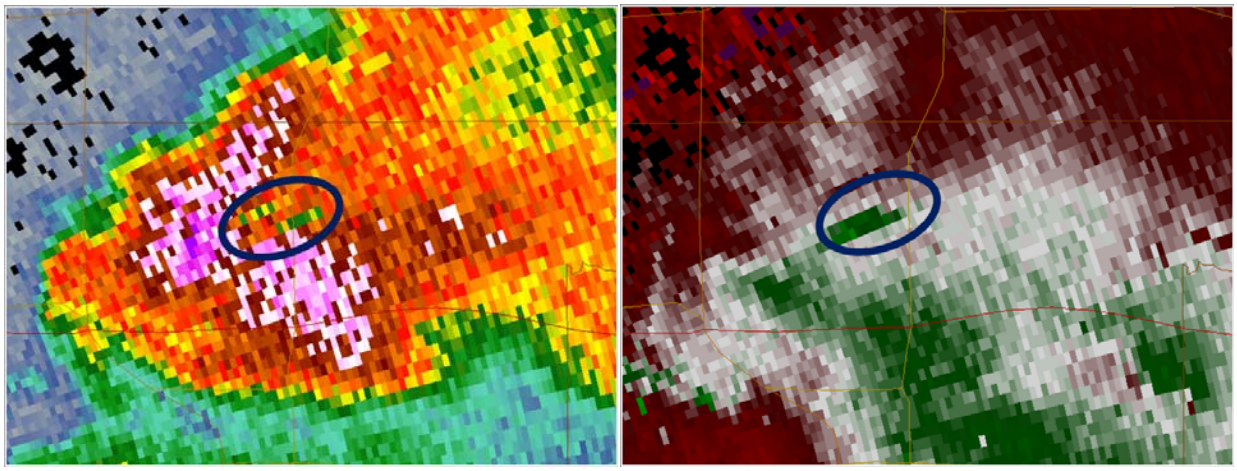

Figure 3-85. Reflectivity (Z, left) and velocity (V, right) data with CMD false detections (blue circle) during a convective event.

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.

Range Folding Mitigation VCPs

There is a class of VCPs that is 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 (see Fig. 3-86). The advantage of SZ-2 is for echo overlay conditions. When echoes are overlaid, the Legacy Range Unfolding algorithm can often provide velocity data for one of the original echoes. Given the same echo overlay situation, SZ-2 Range Unfolding can often provide velocity data for both the echoes.

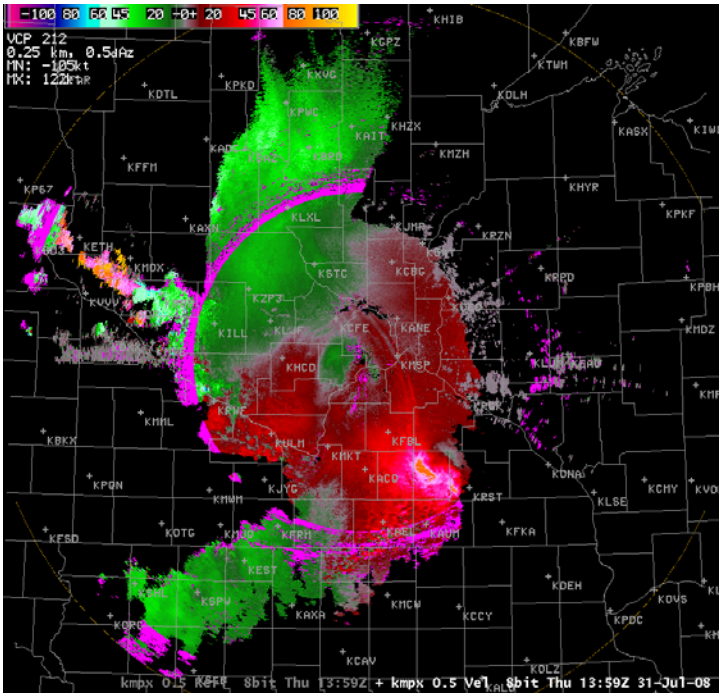

Figure 3-86. An example of SZ-2 Range Unfolding in VCP 212 velocity data.

VCP 212 is the most frequently used VCP of this group. It is usually best for widespread severe convection, with an update rate of about 4.5 minutes. The elevation angles provide good vertical resolution, especially for the lower elevations (see Fig. 3-87). VCP 212 has fast antenna rotation rates, which can degrade data quality, especially when used for events other than severe convection. Also there will be a small band of RF at the end of the first trip that will not move from volume scan to volume scan, with some chance that velocity data could still be masked by RF for a particular storm.

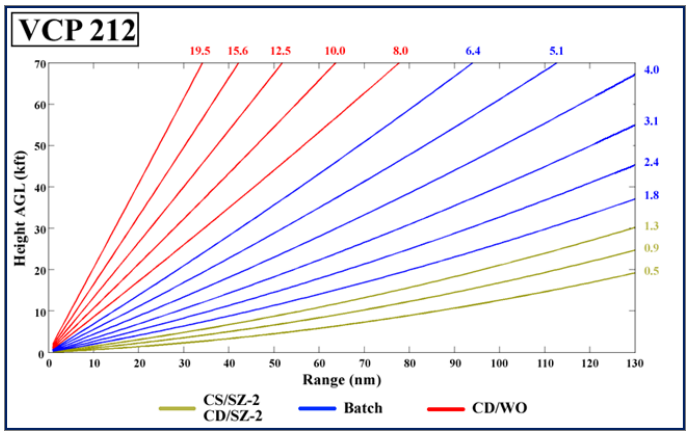

Figure 3-87. Graphical representation of VCP 212.

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 (see Fig. 3-88).

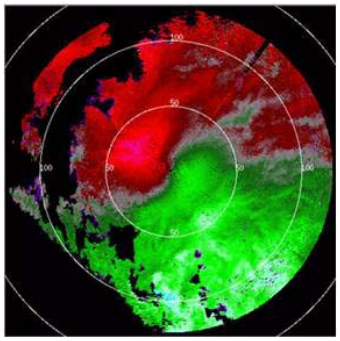

Figure 3-88. In VCP 121, nearly all the data in the lowest two elevations are recovered.

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 (see Fig. 3-89). The update rate of almost six 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.

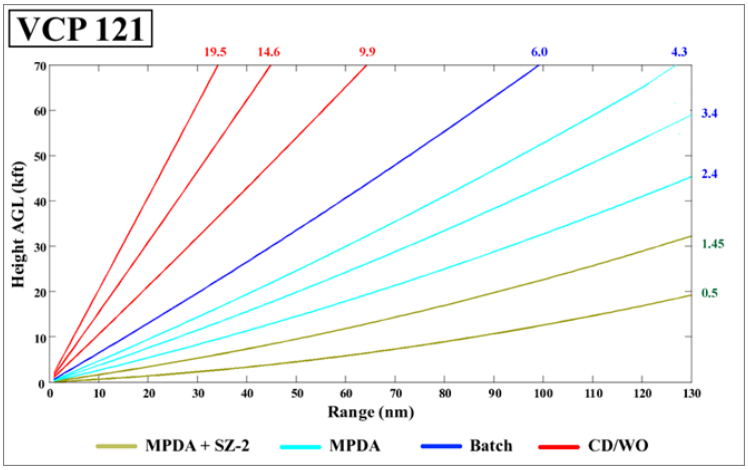

Figure 3-89. Graphical representation of VCP 121.

VCPs and Manual PRF Edits

There are a group of VCPs that allow you to manually select one of the Doppler PRFs, which are PRFs 4, 5, 6, 7, & 8. The process for manual selection will be presented shortly. For now, the point is that manual PRF selection is limited to this

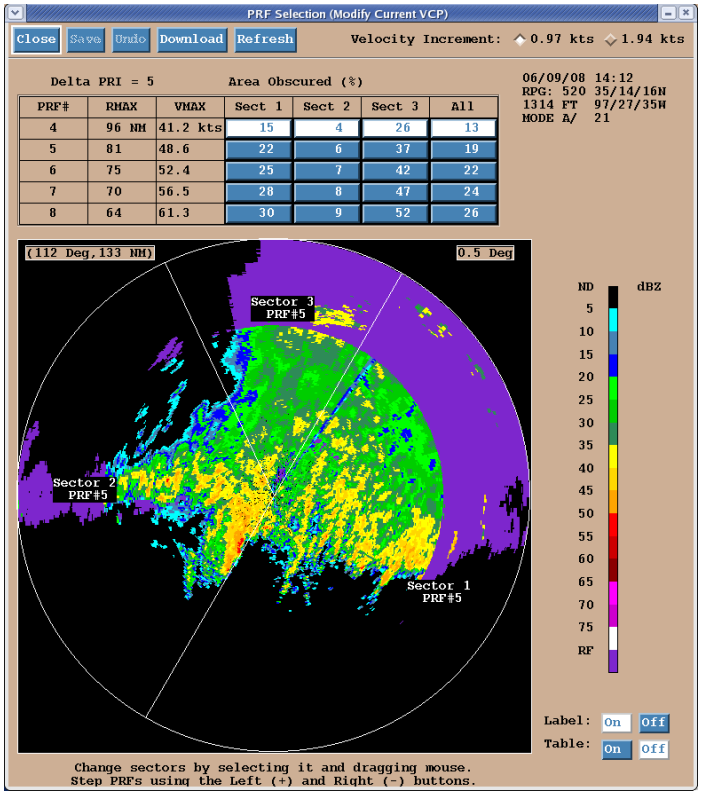

Figure 3-90. PRF Selection window.

group of VCPs: 11, 12, 21, and 32. Figure 3-90 is the window used for PRF editing, which will be explored in the next few slides.

Among this group, VCP 12 (see Fig. 3-91) is often the best choice for isolated to scattered severe convection. Using VCP 12 during a severe weather event and manually editing the Doppler PRFs is the most precise way to prevent a particular storm from being obscured by RF data.

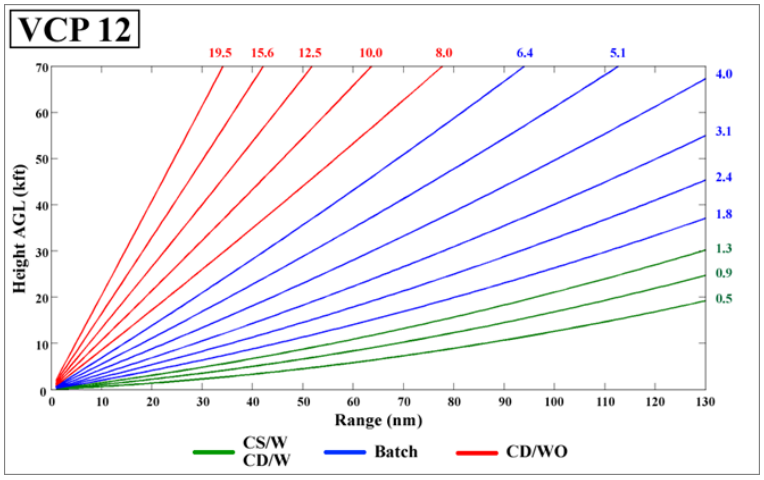

Figure 3-91. Graphical representation of VCP 12.

Though the need to adjust the Doppler PRF is likely rare, this is a good place to mention that VCP 21 (see Fig. 3-92) is often the best choice for stratiform rain. Since it is low-topped, the issues with gaps between the higher elevation angles are minimal. VCP 21 has relatively slow antenna rotations, which means more pulses per radial. This in turn means better performance from the data quality algorithms such as CMD, GMAP, and Super Resolution processing.

There are multiple ways to perform a Doppler PRF selection in the PRF Selection window. For now,

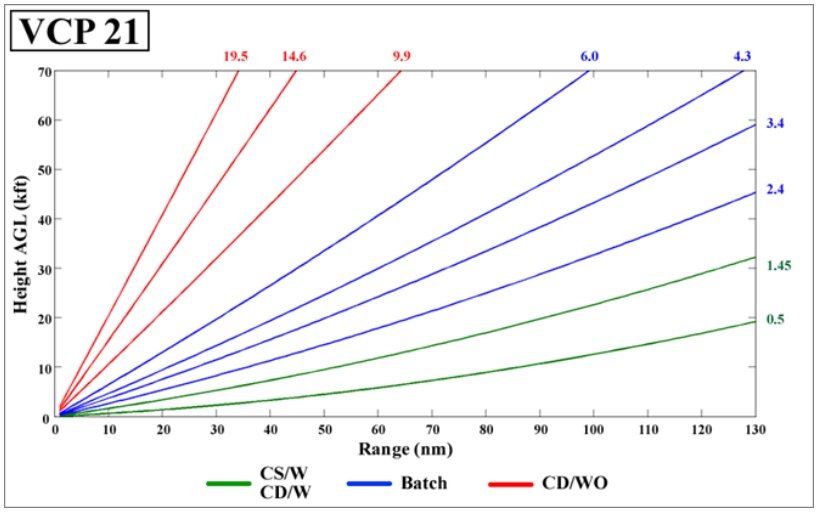

Figure 3-92. Graphical representation of VCP 21.

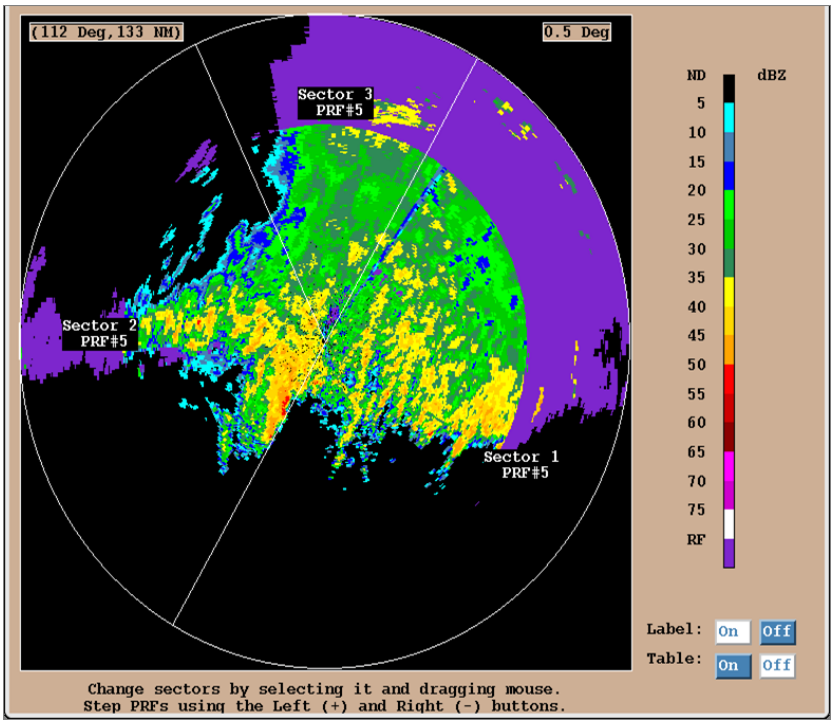

Figure 3-93. The PRF Selection window displays the latest 0.5° reflectivity data and the RG data from the 0.5° velocity data.

consider the graphical area of the PRF Selection window (see Fig. 3-93). It displays the latest 0.5° reflectivity data along with the distribution of the RF data from the 0.5° velocity. This viewer allows the user to see the coverage of RF for each of the Doppler PRFs with respect to the reflectivity. This allows you to see which Doppler PRF protects the most important storms from being obscured by RF.

The boundaries form three sectors, and can be adjusted by clicking on these lines and dragging the mouse. Within each sector a different Doppler PRF can be chosen.

PRF Selection Example

Consider this example of a sectorized Doppler PRF edit during operations. Fig. 3-94 on page 3-114 shows a squall line moving through the Shreveport CWA. The white arc on reflectivity and velocity corresponds to the end of the first trip in the velocity. The squall line is very close to the end

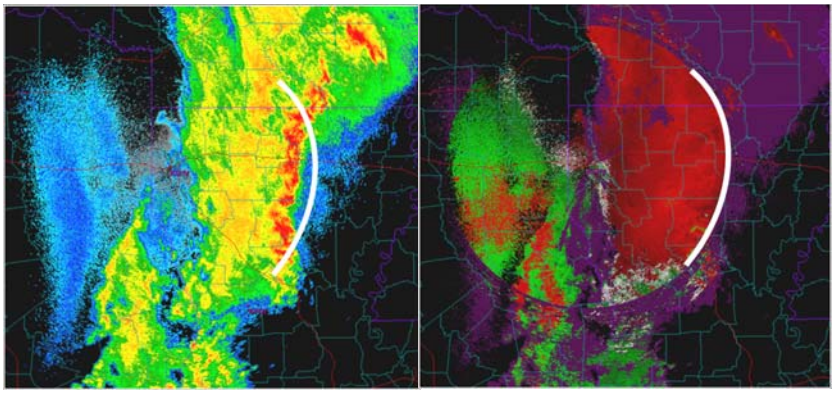

Figure 3-94. Reflectivity (Z, left) and velocity (V, right) data as a squall line moves through. The white arc represents the end of the first trip in the velocity.

of the first trip and there is the risk of the loss of velocity data in the next few volume scans. At this time, the CD R_{max} is at the same range in all directions, so no sectorized PRF editing has happened yet.

At the next volume scan, the squall line has moved further east and there is a “notch out” of the CD R_{max} in the velocity data to the east of the radar. The Doppler PRF has been edited such that the CD R_{max} is longer for a sector from northeast through southeast. This makes velocity data related to the squall line available for a little longer.

The PRF Selection interaction is part of DLOC Topic 6. It allows you to practice the steps of performing a manual Doppler PRF change. For now,

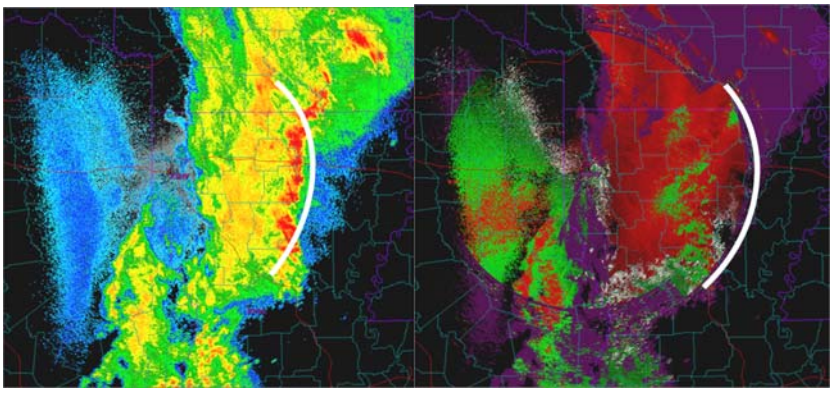

Figure 3-95. Reflectivity (Z, left) and velocity (V, right) images one volume scan after Figure 3-94.

the key point is that VCPs 11, 12, 21 and 32 allow for manual Doppler PRF adjustments. Editing a Doppler PRF manually gives you the most precise control over which storms have available velocity information.

Auto PRF

The goal of Auto PRF is to choose a Doppler PRF that minimizes range folding over the entire display. This is very different from adjusting the Doppler PRF manually for a particular storm. Though you aren't likely to be aware of it, Auto PRF is operating most of the time. With the exception of VCPs 31 and 121, Auto PRF is a function that runs for some portion of all the VCPs. For VCPs 11, 12, 21, and 32, Auto PRF is usually active for the Split Cut and Batch elevations. For VCPs 211, 212, and 221, Auto PRF is usually active for the Batch elevations.

It turns out that the same interface that allows you to manually edit the Doppler PRF also provides a table that shows the output of Auto PRF (see Fig. 3-96, right). For the entire display (the All column), the percent area obscured by RF on the velocity data is listed as a function of each of the Doppler PRFs. This is a numeric version of the qualitative display in the graphical area. Among PRFs 5, 6, 7, and 8, Auto PRF chooses the PRF with the lowest percent area obscuration, and downloads that information to the RDA for the next volume scan.

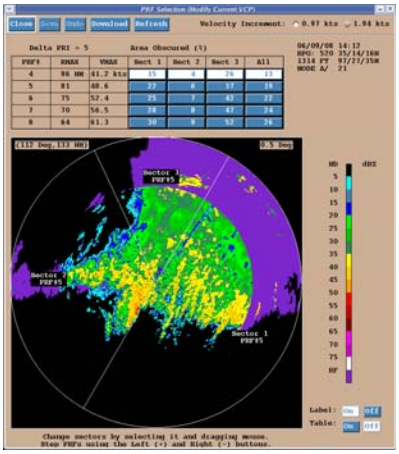

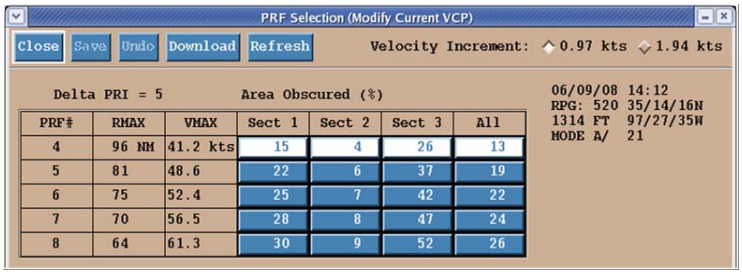

Figure 3-96. The same interface where the Doppler PRF is edited (right) also provides a table that shows the output of Auto PRF (right).

Auto PRF, MPDA, and SZ-2

Recall that for the SZ-2 and MPDA VCPs, 212, 211, 221, and 121, respectively, the Doppler PRFs are held constant for many of the elevation angles. Auto PRF has no effect on this group of VCPs. With Auto PRF set to “On” while any of these VCPs is active, it will continue to calculate the percent area of RF on the velocity data. However, Auto PRF will not download the updated Doppler PRF information to the RDA. It will not overwrite the fixed PRFs used by VCPs 212, 211, 221, and 121.

Attenuation of Z

Attenuation of Z sometimes occurs, and will continue to do so with Dual-Pol. It is fortunate that the WSR-88D is a 10 cm radar, which attenuates much less than 5 cm radars. Of course, attenuation still happens and here is a look at how the Dual-Pol variables are impacted.

Figure 3-97 shows Base Reflectivity, Z, generated from two test radars, KOUN (Dual-Pol) on the right and KCRI (Single-Pol) on the left, which are very close to one another. In fact, they are across the street from one another. There is a squall line parallel to several radials and you can see the attenuation down radial in both of the Z products (white

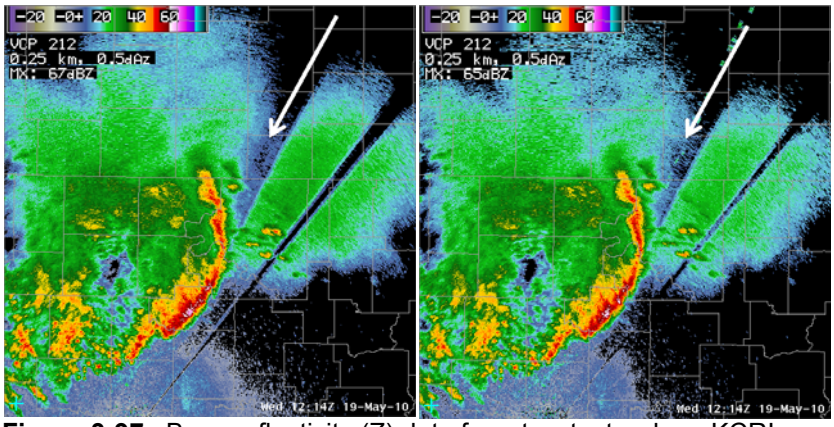

Figure 3-97. Base reflectivity (Z) data from two test radars: KCRI (Single-Pol, left) and KOUN (Dual-Pol, right).

arrows). Once the signal is attenuated, that loss cannot be recovered and propagates down radial.

Differential Attenuation of ZDR

With this same squall line case (Fig. 3-98 on page 3-118), 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 result 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 than expected. In areas where we know large to medium sized raindrops exist, large positive ZDR values are expected. With differential attenuation, ZDR values are generally negative, and extend down radial from the storm cores along the radial. Once the signal is attenu-

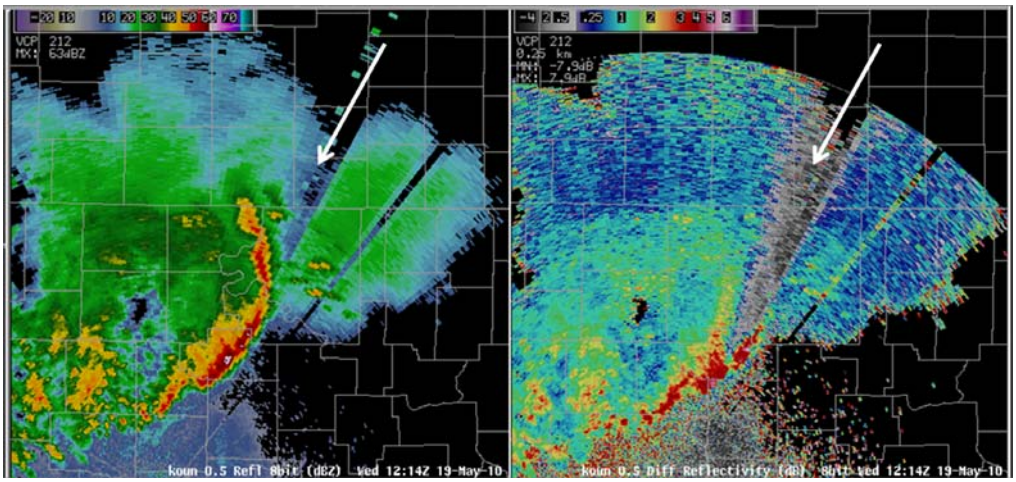

Figure 3-98. Reflectivity (Z, left) image with attenuation of Z and differential reflectivity (ZDR, right) image showing differential attenuation of ZDR.

ated, the loss in ZDR cannot be recovered and propagates down radial.

What's in the Beam?

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 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 Figure 3-99 and Figure 3-100, the circle represents the radar beam as if you were standing at the RDA looking outbound along a radial. Figure 3-99 represents partial beam filling, which is familiar, resulting in underestimated Z values.

Figure 3-99. Partial beam filling as seen from the RDA looking outbound.

Figure 3-100. The beam is filled, but by a mix of precipitation sizes and types.

In Figure 3-100, 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 base data quality.

Uniform Beam Filling and Correlation Coefficient (CC)

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 correlation coefficient (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 uni-

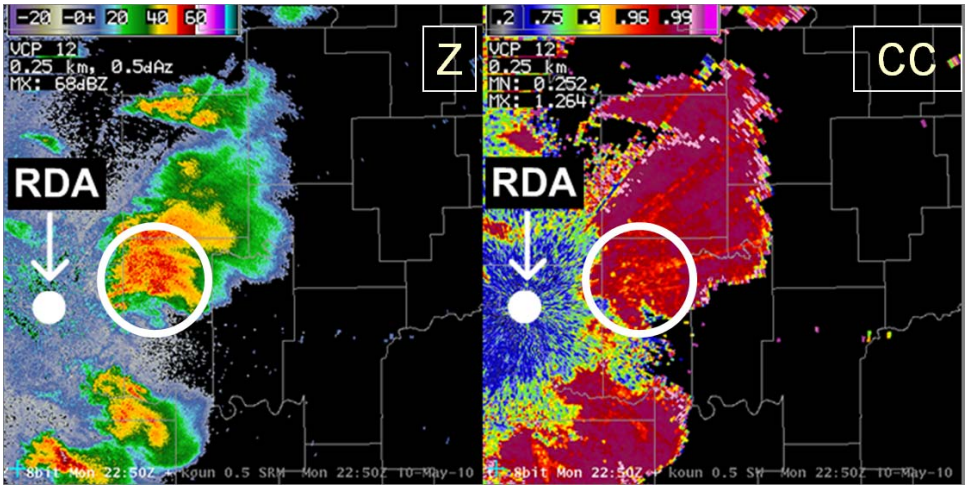

Figure 3-101. Reflectivity (Z, left) data and correlation coefficient (CC, right) data showing a supercell close to the radar. Lower CC values within the supercell are likely due to a mix of rain and hail.

formly distributed across the radar beam cross section.

Non-uniform Beam Filling

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, Figure 3-102 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

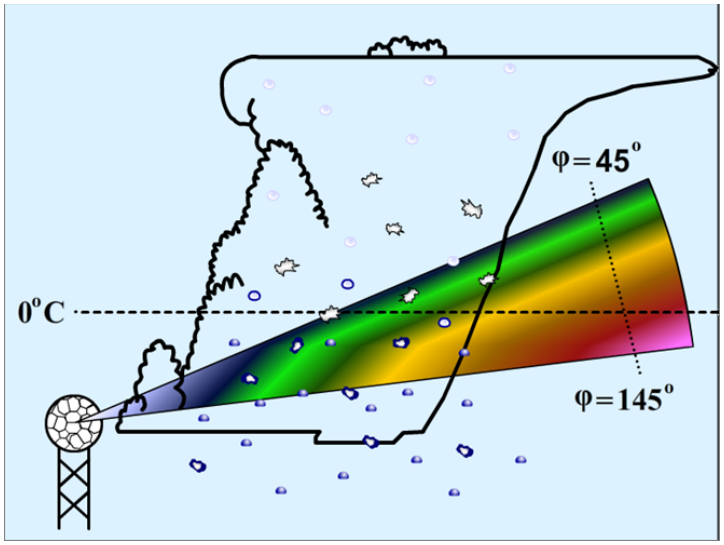

Figure 3-102. For storms far from the radar, the beam will sample a gradient of precipitation, resulting in a gradient in Φ_{DP} values

of Φ_{DP} is the bottom line for low CC values locally and down radial.

Non-uniform Beam Filling and Correlation Coefficient (CC)

From the example in Figure 3-101, the supercell has moved to the east and is at a longer range, with the beam sampling a larger volume of the storm (see Fig. 3-103). There are radial swaths of low CC that originate from the storm core (white ovals). This is an example of non-uniform beam filling and its impact on the CC product. It is important to be aware of the potential for NBF on the CC products, because it has consequences for other Dual-Pol products.

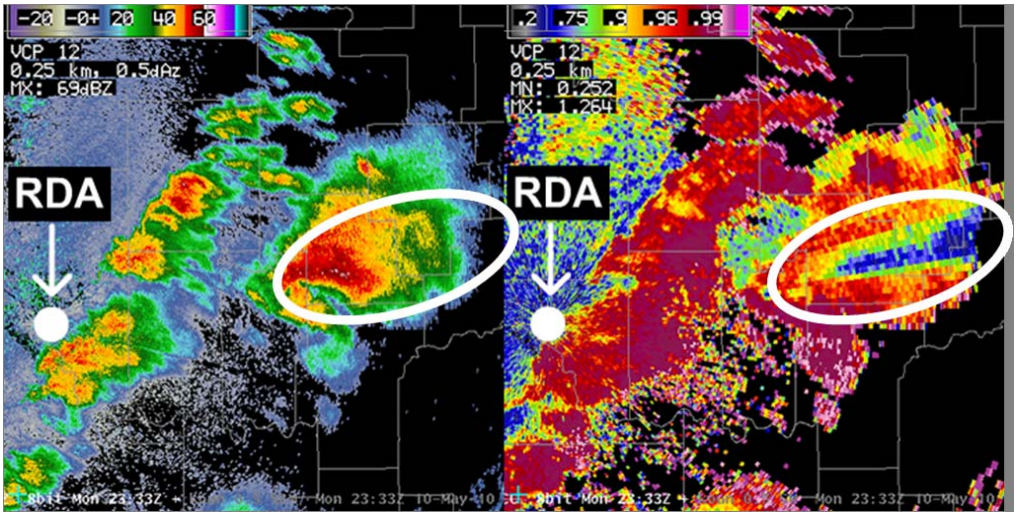

Figure 3-103. Reflectivity (Z, left) and correlation coefficient (CC, right) data from the same supercell shown in Figure 3-101 after the storm has moved east.

Extra for Experts: How NBF Lowers CC Downradial

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 (see Fig. 3-104).

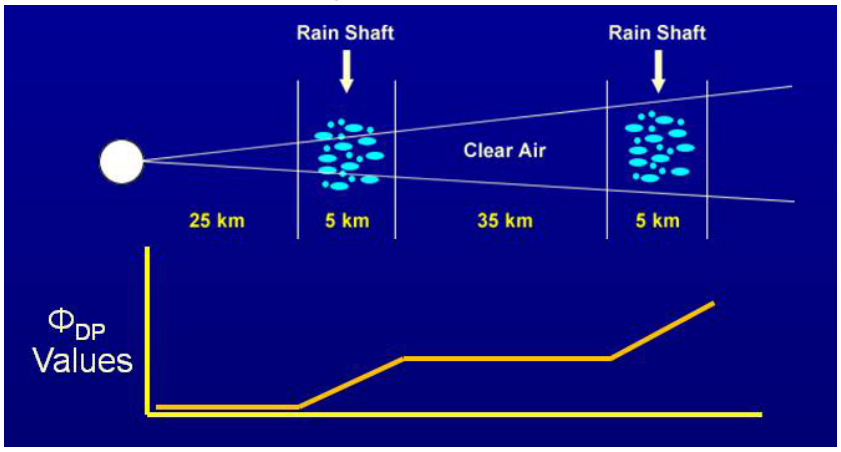

Figure 3-104. Recall that Φ_{DP} values are cumulative down radial.

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

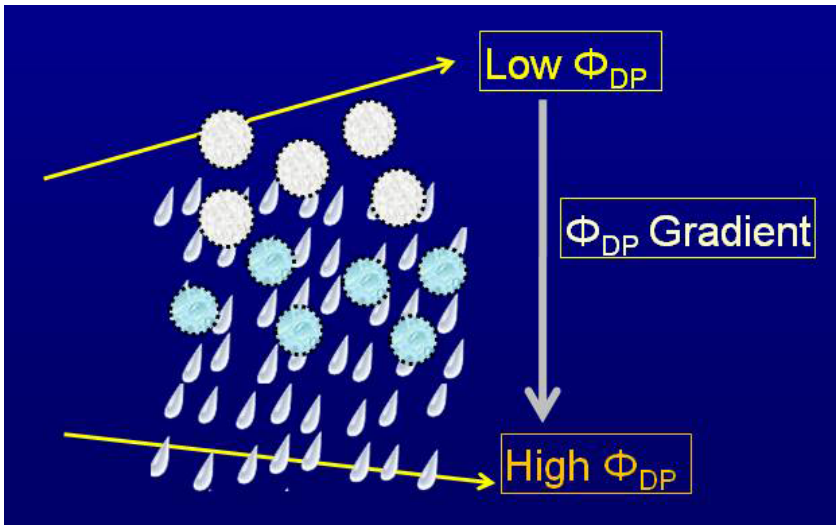

Figure 3-105. Sampling convective storms at longer ranges or squall lines along a radial can lead to capturing a gradient of precipitation types within the beam.

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.

If we could double the vertical resolution, the Φ_{DP} values for the top of the beam could be measured separately from the bottom. The Φ_{DP} value at the top of the beam only slightly increases since mostly ice is being sampled. The Φ_{DP} value at the bottom of the beam increases significantly since mostly water is being sampled. This dramatic Φ_{DP} gradient within the beam continues down radial. If it is significant enough, the CC value is also lowered down radial, because Φ_{DP} does not reset.

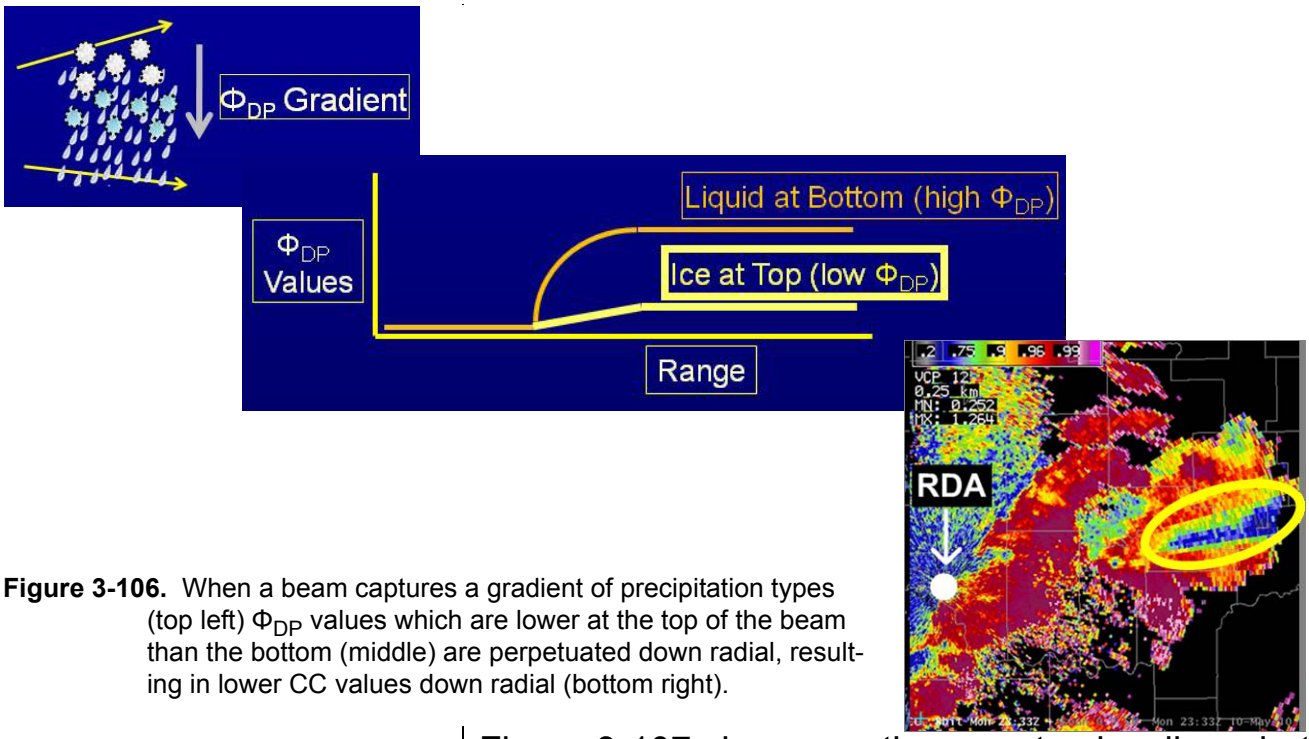

Figure 3-106. When a beam captures a gradient of precipitation types (top left) Φ_{DP} values which are lower at the top of the beam than the bottom (middle) are perpetuated down radial, resulting in lower CC values down radial (bottom right).

Figure 3-107 shows another way to visualize what happens to the CC down radial from this type of non-uniform beam filling. This is a simplified example with only four range bins. We are showing 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

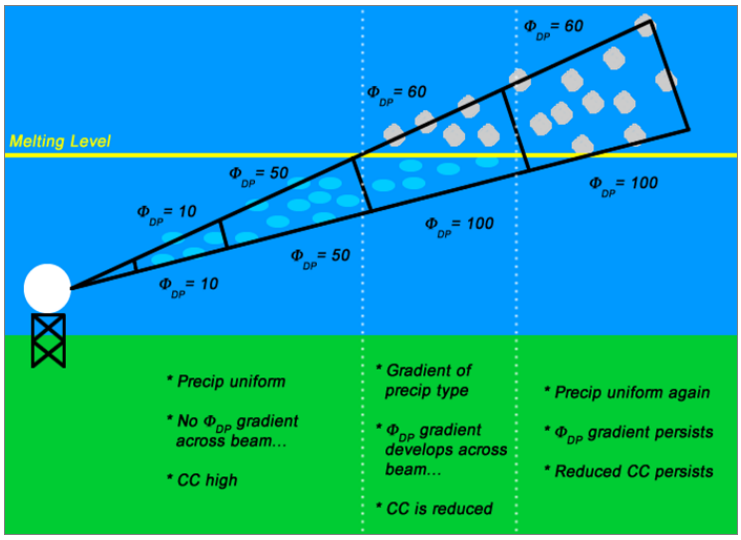

Figure 3-107. Simplified example of how CC is affected by non-uniform beam filling.

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.

NBF Impact on Dual-Pol Products

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 (left, Fig. 3-108) affects the Hydroclass value (middle, Fig. 3-108) that gets assigned, which then affects whether or nor rainfall is accumulated (DPR, right, Fig. 3-108).

You will see more about NBF later in DLOC.

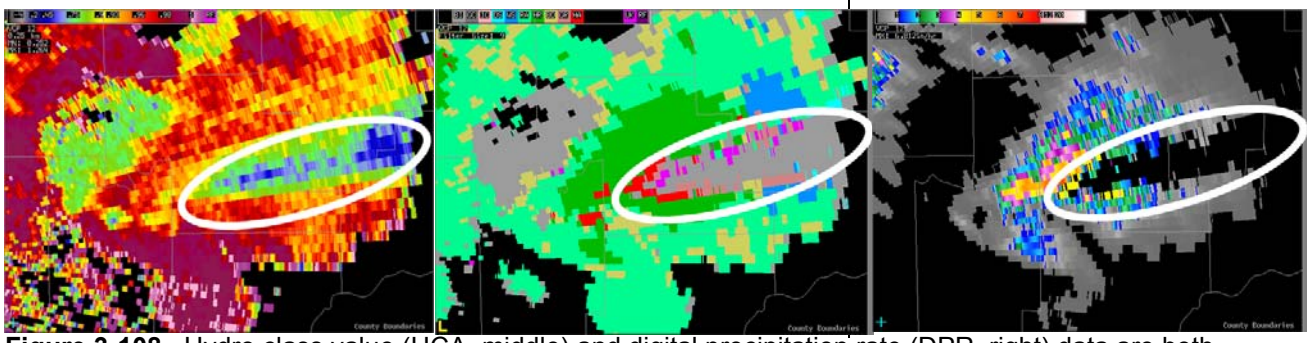

Figure 3-108. Hydro class value (HCA, middle) and digital precipitation rate (DPR, right) data are both affected by non-uniform beam filling in correlation coefficient (CC, left) data.

Differential Reflectivity (ZDR) and Depolarization

Depolarization is a phenomenon that has always occurred with radar, but will be apparent on the differential reflectivity (ZDR) product. Depolarization means that the reflected energy from a particle switches polarization, from horizontal to vertical, vertical to horizontal, or even more fun, both at the same time!

In Figure 3-109, for simplicity, we are looking at a transmitted pulse that is polarized only in the horizontal direction (yellow oval). The beam encounters needle-shaped ice crystals which are canted at an angle such that the energy reflected back to the radar is at least partially in the vertical direction (yellow and red ovals). A Dual-Pol radar is going to process the vertical along with the horizontal, thus affecting the ZDR value.

Depolarization only affects the ZDR product. It appears as radial spikes that can be either high or low ZDR values which are transient with time (see white oval in Fig. 3-110). Though it may rarely

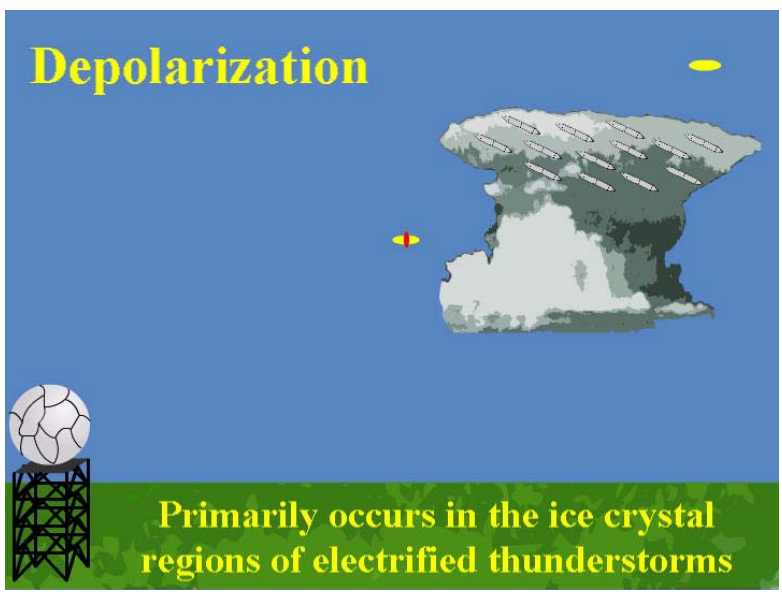

Figure 3-109. When the beam encounters needle-shaped ice crystals, which are canted at an angle, the energy reflected back is partially in the vertical. This affects the ZDR product.

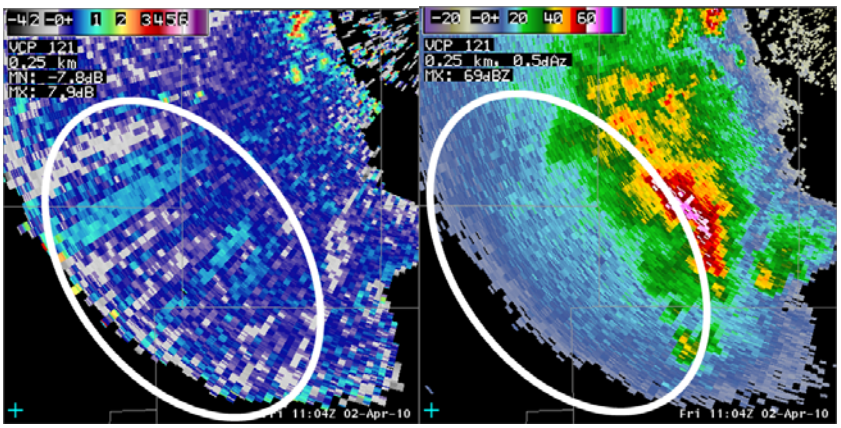

Figure 3-110. Differential reflectivity (ZDR, left) data and reflectivity (Z, right) data showing the effects of depolarization on ZDR (white ovals).

occur in hail, depolarization is far more likely to happen in the upper regions of thunderstorms when the electrification causes canting of the ice crystals. Since the electrification varies with time, so does the impact of depolarization.

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

WSR-88D Data Quality

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 versus 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.

1. VCPs 12, 212, 121

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

2. CMD

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 Figure 3-111, 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

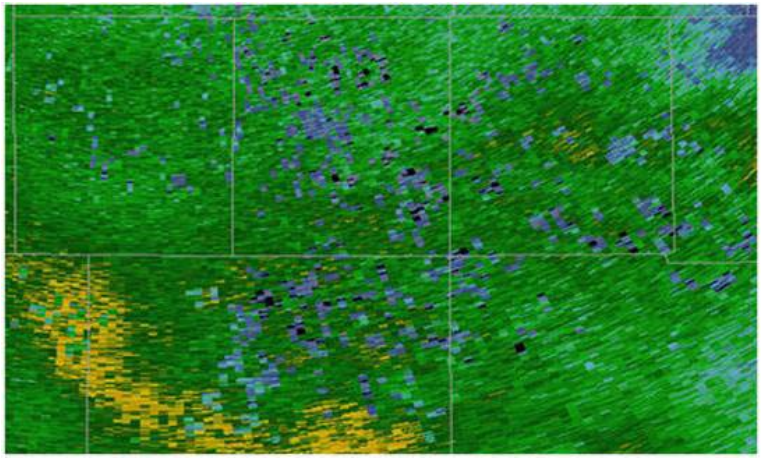

Figure 3-111. Reflectivity (Z) data with VCP12 in effect. The random distribution of gates with reduced power or no data at all is, in part, due to CMD.

image. CMD is one of the contributors, by falsely identifying bins that contain clutter.

3. GMAP

Once CMD identifies 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.

4. Super Resolution Processing

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. For many of the RPG algorithms, the error is too high for input, and the data are recombined before the algorithms “do their thing”. For human interpretation of the base products, the noisiness can sometimes approach even human tolerance, even given the variation in tolerance from one human to the next.

Putting It Together: Stratiform Rain, VCP 12

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 has 28 pulses per radial for the lowest elevation. VCP 12 has 15 pulses per radial, almost half the number of VCP 21. WIth VCP 21, base data quality improved significantly.

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 some 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.

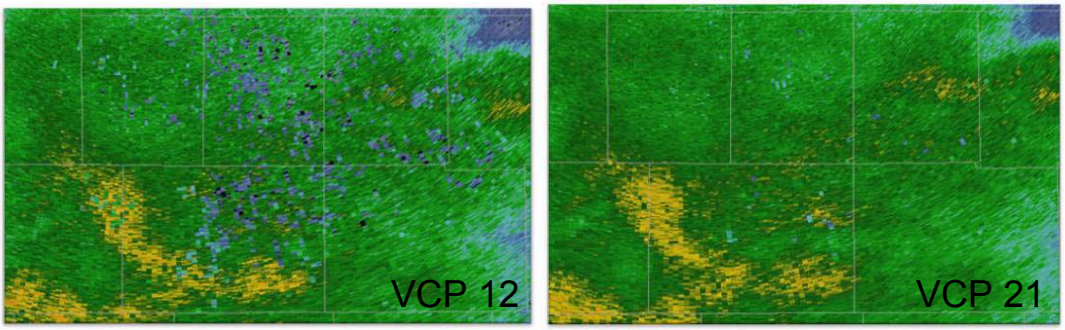

Figure 3-112. Reflectivity (Z) from a stratiform rain event using VCP 12 (left) and VCP 21 (right).

Distance Learning Operations Course

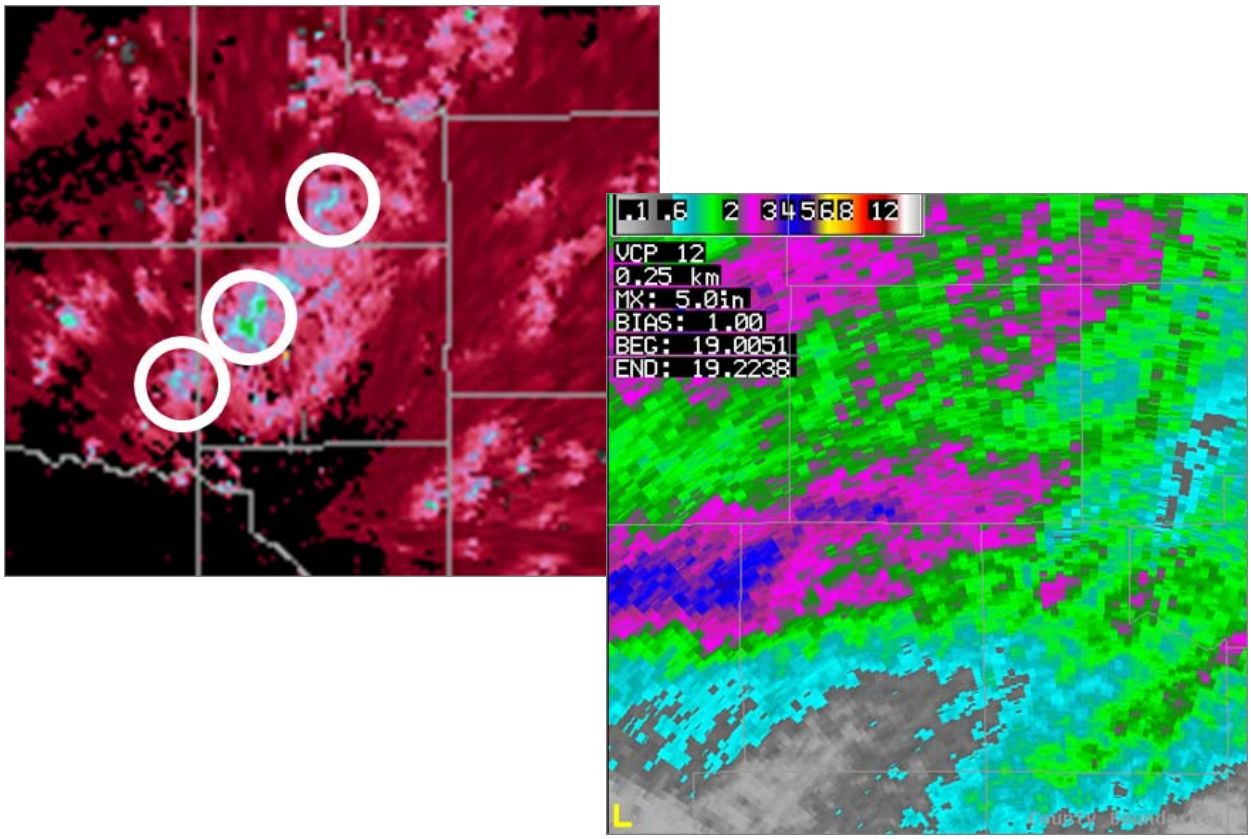

Topic 3, Lesson 7: Precipitation Estimation

Presented by the
Warning Decision Training Branch

Version: 1210

Lesson 7: Precipitation Estimation

Lessons 7 and 8 look at the overall challenges of using a radar to estimate the amount of rain or snow that reaches the surface. These lessons also give you a "tour" of the RPG algorithms that estimate rainfall (two of these) and snowfall (one of these).

Objectives

1. Identify 3 of 7 potential error sources associated with using a radar to estimate rainfall.
2. Identify the available radar inputs for generating a rainfall rate, and their limitations.
3. Identify strengths & limitations of the Precipitation Processing Subsystem.

Heavy Rainfall Detection vs. Rainfall Estimation

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

Assume that we are looking at radar data that is below the melting layer, i.e. hydrometeors are liquid. Among the available base products, Specific Differential Phase (KDP) is the best indicator of relative liquid water content (see Fig. 3-113, left). That means that observing KDP over time can identify areas that are most vulnerable to significant rainfall. In addition to identifying areas of potentially significant rainfall, areal coverage patterns over time can be revealing, such as growth and movement, linear vs. circular, etc.

On the other hand, using the output from the RPG algorithms (Fig. 3-113, right) 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.

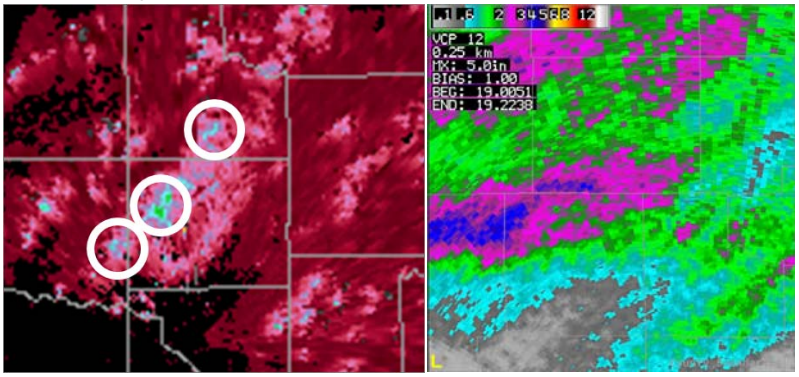

Figure 3-113. Specific differential phase (KDP, left) and legacy PPS Storm Total Precipitation (STP, right)

Issues of Using Radar to Estimate Rainfall

There are many issues that make using a radar to estimate rainfall amounts extraordinarily difficult. This lesson focuses on seven of these issues, which are most relevant to the WSR-88D.

1. Residual Clutter

The term Residual Clutter is used here to describe unfiltered clutter, whether the clutter filtering has been turned off, or there is some clutter remaining after filtering is applied.

Ground clutter (Fig. 3-114, left) is the term for returns from ground based targets that are always present, such as buildings and terrain. Anomalous Propagation (AP) (Fig. 3-114, right) clutter technically refers to the beam superrefracting, but AP is also used to describe the clutter contamination that can result when the beam is striking ground targets at varying ranges due to superrefraction. AP clutter is transient in space and time.

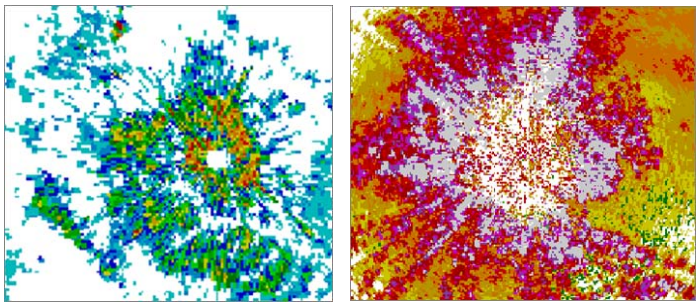

Figure 3-114. Reflectivity data showing ground clutter (left) and anomalous propagation (AP, right).

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

It is also possible to apply clutter suppression where it is not needed, through the use of All Bins suppression. This can result in an underestimate of Z, thus an underestimate of R. The data quality impact of applying All Bins suppression unnecessarily to the Dual-Pol base data is not yet fully understood. However, it is reasonable to assume that unnecessary clutter suppression would negatively impact the quality of the dual-pol rainfall estimates.

2. Wet Radome

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. The goal is to prevent water coating on the surface. It still happens, but is usually a transient condition.

A water coating on the radome reduces the amount of transmitted power. For the horizontal

channel from which the legacy base data are calculated, this results in a reduction in returned power, an underestimate of Z, and an underestimate of R. 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.

3. Incorrect Calibration

Rainfall estimation is particularly sensitive to calibration errors. The horizontal and vertical channels must each be calibrated separately. A valid Z value is dependent on a well-calibrated horizontal channel, while a valid ZDR value is dependent on both the horizontal and vertical channels being well-calibrated.

There are two different types of calibration, on-line and off-line. An 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 eight hours. This 8-hour “Performance Check” also includes multiple tests to assess the “health” of the radar. You will learn a bit more about the 8-hour check in DLOC Topic 6. 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.

4. Below Beam Effect Errors

Below beam effects are a fundamental challenge when using a radar to estimate rainfall. It is important to remember that the 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 with the presence of a deep, dry sub-cloud layer. The desert southwest often has convection with significant evaporation of rain below the beam. This layer will evaporate some or all of the rain leaving the cloud, causing a smaller amount rain to reach the ground. As a result, we have an overestimate of rainfall.

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 high number of smaller drops collide for raindrop growth. Except for short range precipitation, the largest drops are often too low to be sampled by the radar beam, which results in an underestimate of rainfall.

5. Beam is In or Above the Melting Layer

It is important to remember that both of the RPG rainfall algorithms are designed to estimate liquid rainfall. Performance of these algorithms is much more reliable for the ranges where the radar beam is intercepting liquid hydrometeors, below the melting layer. Within the melting layer and above, hydrometeors are frozen or freezing (see Fig. 3-115), and converting to liquid rainfall on the

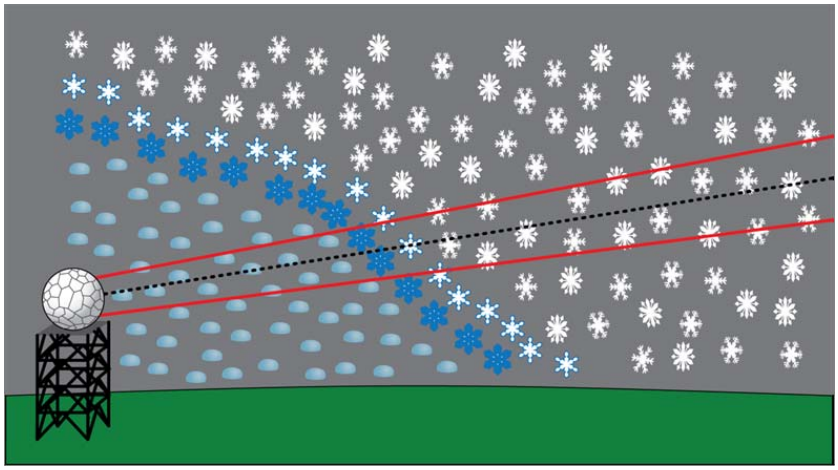

Figure 3-115. Because the RPG rainfall algorithms are designed to estimate liquid precipitation, they perform best below the melting layer, where the beam is sampling purely liquid.

ground really requires a snow conversion algorithm, which will also be discussed in Lesson 8.

6. Partial or Non-Uniform Beam Filling

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

Figure 3-116 shows 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. The reflectivity and rainfall rates are both underestimated. Since the beam volume

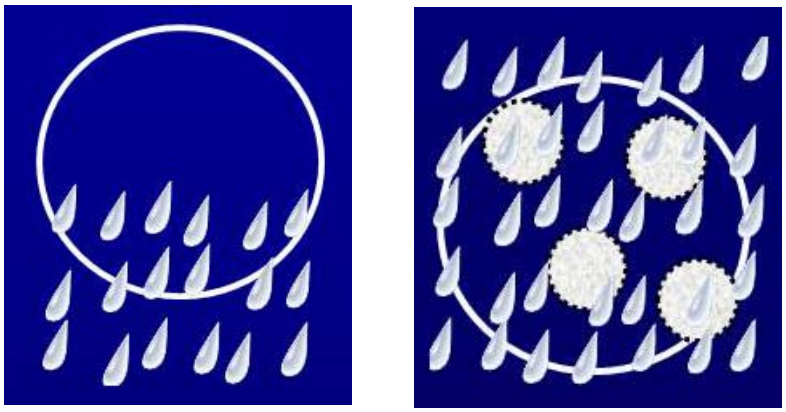

Figure 3-116. If the beam is only partially filled with precipitation (left) or sampling more than one precipitation type (right), precipitation estimates can be affected.

is greater than the actual precipitation volume, the areal coverage is then overestimated.

The literal meaning of non-uniform beam filling is something that most likely happens frequently. It turns out that a particular type of non-uniform beam filling has a significant impact on the Dual-Pol products. Where there is a gradient of precipitation type across the beam, CC can be underestimated down the radial. This means that areas that have active precipitation may not be converted to rainfall. There was an example of this in Topic 3 Lesson 6.

7. Coefficients and Exponents May Vary

The equations below are used to convert radar base data into rainfall rate by the RPG algorithms. These equations will be presented again later in this lesson, as well as in Lesson 8. The point for now is that the coefficients and exponents for these rainrate equations are empirical. The applicability of these equations varies with differing dropsize distributions.

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

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

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

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

RPG Algorithms to Estimate Rainfall

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. Its inputs are reflectivity, velocity and spectrum width. The PPS computes rainfall rates based on Z, from a choice of Z-R relationships. The “tour” of the PPS will be presented in this lesson.

The second rainfall algorithm is the Quantitative Precipitation Estimation (QPE) algorithm. The QPE has been recently fielded as part of the Dual-Pol upgrade. Along with reflectivity and velocity, QPE uses the dual-pol base data for input. The “tour” of the QPE will be presented in Lesson 8, “More Precipitation Estimation”.

Z is Estimated

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

is an estimate that comes from the Probert-Jones radar equation (equation (1)).

$$P_r = \left[\frac{P_t(G^t\theta^2 H\pi^3 K^2 L)}{1024(\ln 2)\lambda^2}\right] \times \frac{Z}{R^2} \qquad (1)$$

Rain Rate Equation Input: Z

Estimating rainfall from radar is a very complicated business. It's important to remember that there are potential errors in the Z value even before it gets converted to rain rate.

As range increases, so does the size of the volume that is sampled by the beam. 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 filled 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.

Relating Z to R

Both Z and the rainfall rate, R, are dependent on the dropsize distribution, which is unknown. Also, the dependence differs. 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 through a Z-R relationship, expressed as a power law equation. In equation (2), alpha and beta are empirical constants that change depending on the meteorological event being analyzed.

$$Z = \alpha R^{\beta} \quad (2)$$

PPS: Z-R Relationships

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. The table in Figure 3-117 summarizes the available WSR-88D Z-R relationships, along with their optimal environments.

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

Figure 3-117. This table lists the available WSR-88D Z-R relationships and 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.

PPS and QPE Share One Equation

The legacy PPS and the Dual-Pol Quantitative Precipitation Estimation (QPE) algorithm share one equation for converting Z to rainfall rate. Equation (3) and equation (4) are the same relationship expressed in two different ways.

$$Z = 300R^{1.4} \quad (3)$$

$$R(Z) = (0.017)Z^{0.714} \quad (4)$$

Equation (3) is the more familiar format, solved for Z. This equation can be rewritten in a format that is solved for R. The R(Z) format in Equation (4) is introduced here to help your understanding of the QPE design. 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 $Z=300R^{1.4}$.

QPE: Rain Rate Equations

Remember, the objectives for this module do not involve memorizing equations! Here are the 3 equations for computing rain rate used by QPE, only to allow for comparison of the PPS vs. QPE approaches for converting to rain rates. How QPE determines which of these equations to use will be explored in Lesson 8.

$$R(Z) = (0.017)Z^{0.714} \quad (5)$$

$$R(Z, ZDR) = (0.0067)Z^{0.927}ZDR^{-3.43} \quad (6)$$

$$R(KDP) = 44.0|KDP|^{0.822}sign(KDP) \quad (7)$$

Equation (5) is the $Z=300R^{1.4}$ relationship solved for R. Equation (6), R(Z,ZDR), combines Z and ZDR as input, and they are in linear units (mm^6/m^3). You may notice that the R(KDP) equation includes the possibility of a negative rain rate. The QPE logic will reject R(KDP) when the rate is negative.

The coefficients and exponents for these equations were developed based on research conducted in Oklahoma. Additional research will be needed to tune these equations for other areas.

The limitations of equation (5), $Z = 300R^{1.4}$ now presented in its R(Z) format, are probably familiar. It often underestimates rainfall in a warm rain dominant event and overestimates the water at the surface from mixed precipitation types, such as wet snow and hail.

Equation (6), 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. For example, in a warm rain process environment, R(Z) is often underestimated. The ZDR value is likely low, and the negative exponent will slightly increase the rain rate. On the other hand, if the dropsize distribution is dominated by large drops, R(Z) alone can often be overestimated, while the high ZDR would slightly decrease the rain rate.

RainRate Equation Input: Z and ZDR

Z and ZDR as inputs for computing rainrate 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, making it harder for the beam to be uniformly filled with the same hydrometeors. Variations in dropsize distribution, especially the presence of hail, affect the Z and/or ZDR values. Both Z and ZDR are sensitive to calibration errors. Even at 10 cm, attenuation or differential attenua-

tion happens, making the Z and/or ZDR estimates less reliable.

QPE Rain Rate Equation

Equation (7), R(KDP), has the advantage that KDP is mostly immune to partial beam blockage and to hail contamination. The magnitude of KDP is related to the liquid water content in the volume, but not to ice. For example, if mostly hail is present in a volume, KDP values can be quite low. On the other hand, if there is a large quantity of small melting hail, KDP can be very large. Since KDP is mostly immune to hail, it can be used to estimate rainfall where hail is present. This equation includes the term “sign(KDP)”. Negative KDP values are possible, and this term is included as a quality control. If R(KDP) < 0, it is not used to estimate rainfall.

RainRate Equation Input: KDP

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 and/or fast VCPs are used, KDP is less reliable. The use of KDP for rainfall estimation is also dependent on the associated CC value. The details of this will be presented in Lesson 8, but for now, if CC is low, R(KDP) is not used for rainfall estimation.

PPS Tour

Now for a “tour” of the Legacy Precipitation Processing Subsystem (PPS). The PPS is a series of algorithms that use base reflectivity, velocity and spectrum width data as input, then estimates 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.

PPS: Flow of Data

Figure 3-118 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 blue boxes that have the word "snow" on them are not related to the PPS. They will be discussed in Lesson 8, which presents the Snow Accumulation Algorithm.

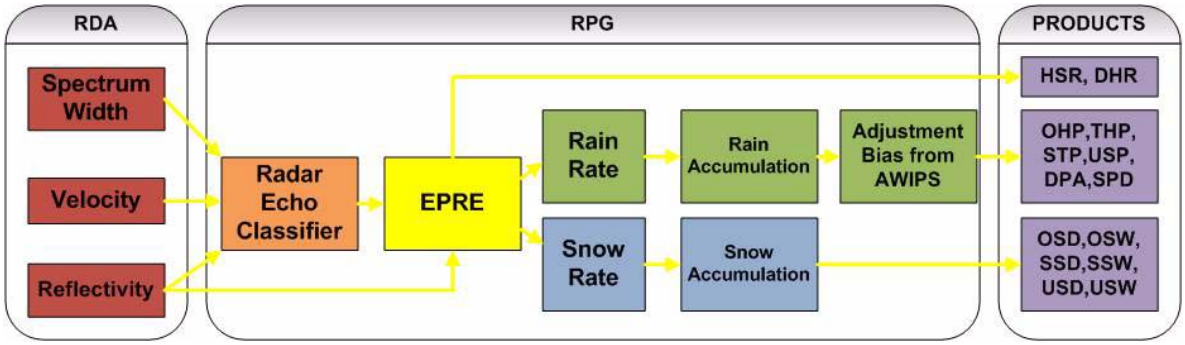

Figure 3-118. Graphical representation of the Precipitation Processing Subsystem (PPS).

Quality Control: Radar Echo Classifier (REC)

The Radar Echo Classifier is a fuzzy logic algorithm that uses reflectivity, velocity, and spectrum

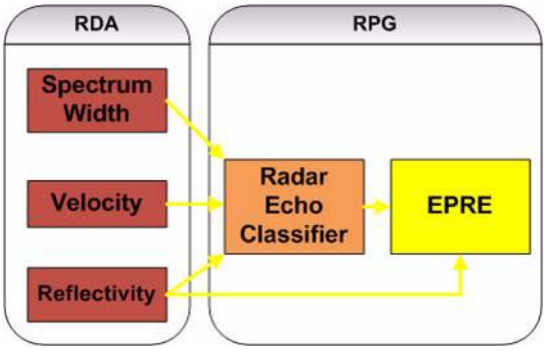

Figure 3-119. The Radar Echo Classifier (REC) uses reflectivity, velocity and spectrum width to assign a likelihood to a bin containing clutter

width to assign a likelihood that a particular bin contains clutter (see Fig. 3-119).

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 percentage 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. Figure 3-120 shows an example where CLUTTRESH was increased to 75% (yellow box) 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.

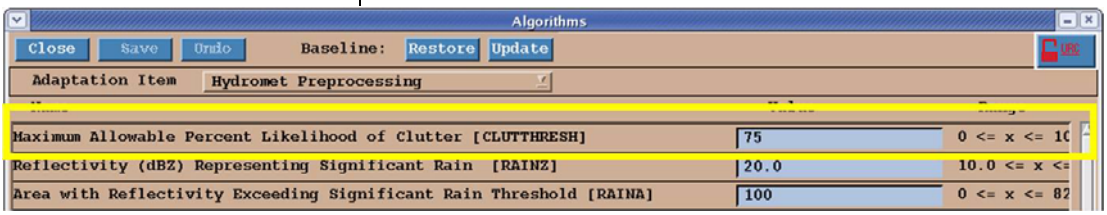

Figure 3-120. CLUTTHRESH (yellow box) is the threshold that determines whether a bin is converted to rainfall or rejected.

Quality Control: Exclusion Zones

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.

Building the Hybrid Scan

The “grand finale” of the EPRE algorithm is the building of the hybrid scan (see Fig. 3-121). 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 Topic 5.

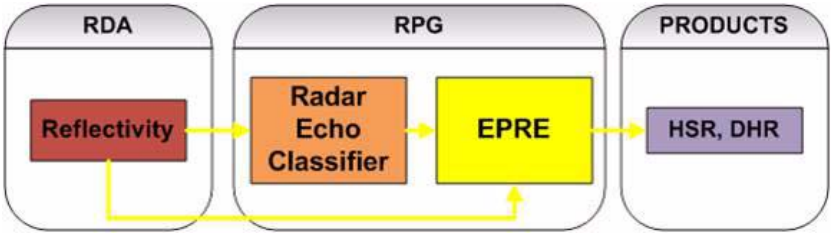

Figure 3-121. The EPRE builds the hybrid scan which finds an optimal dBZ value at every range bin solely for converting to rainrate.

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%

Start and Stop Accumulations

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.

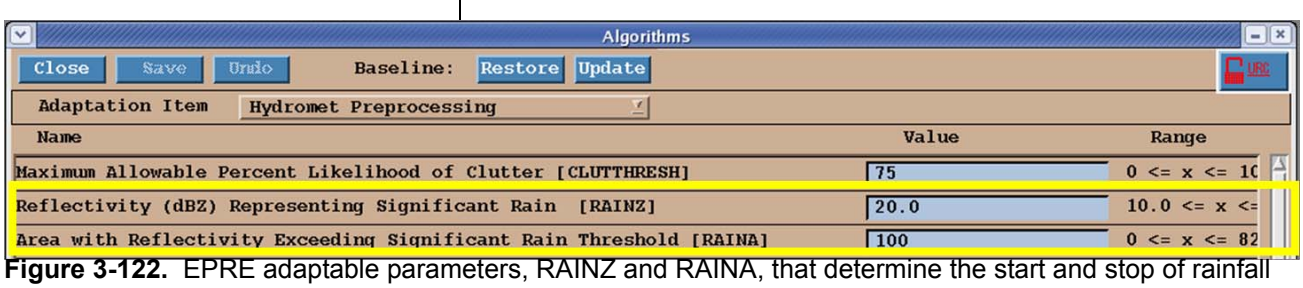

Figure 3-122. EPRE adaptable parameters, RAINZ and RAINA, that determine the start and stop of rainfall accumulations.

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

Resetting Accumulations

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 as the

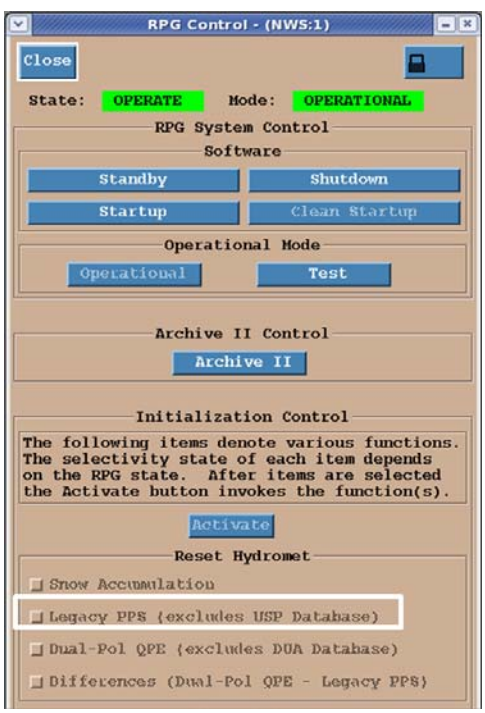

Figure 3-123. Storm total accumulation can be manually reset at the RPG Control Window.

snowfall algorithm, but for now, we focus on the "Legacy PPS" (Fig. 3-123, white box).

Rain Rate Algorithm and the Max Precipitation Rate (MXPRA)

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 Rain 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.

Rain Accumulation Algorithm

The Accumulation Algorithm uses the 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, such as the Storm Total Precipitation (STP) product (Fig. 3-125, left).

The second type of accumulation is hourly. There is a one hour accumulation ending at the current volume scan time. This is the accumulation dis-

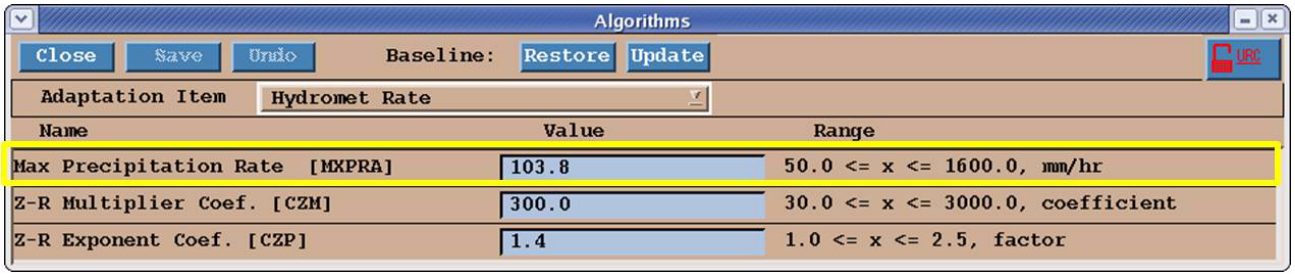

Figure 3-124. The Rain Rate Algorithm applies the Max Precipitation Rate, which acts as a cap to prevent hail contamination.

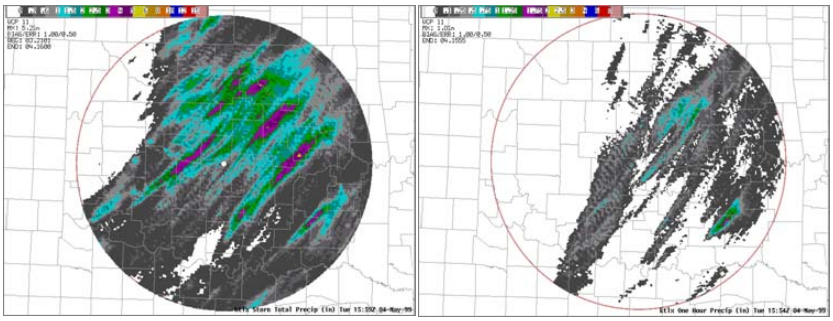

Figure 3-125. Storm Total Precipitation (STP, left) and One Hour Precipitation (OHP, right) products.

played on the one hour products, such as the One Hour Precipitation (OHP) product (Fig. 3-125, right). There is also a one hour accumulation that ends at the top of each hour. These are used to build the Three Hour Precipitation (THP) and User Selectable Precipitation (USP) products.

Adjustment Algorithm

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 (Fig. 3-126, yellow box). 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.

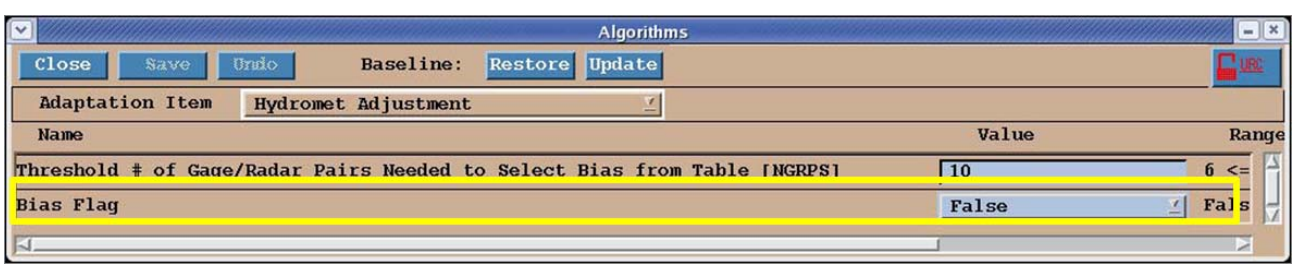

Figure 3-126. Bias is applied by setting the Bias Flag (yellow box) to “True.”

Bias: A Word of Caution

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 than the radar.

PPS Strengths

There are many strengths of the PPS algorithm. It 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

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 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.

Distance Learning Operations Course

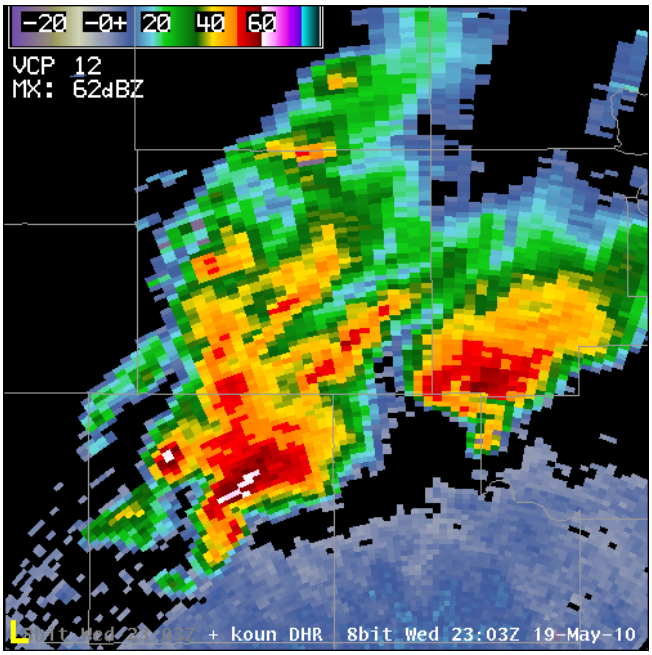

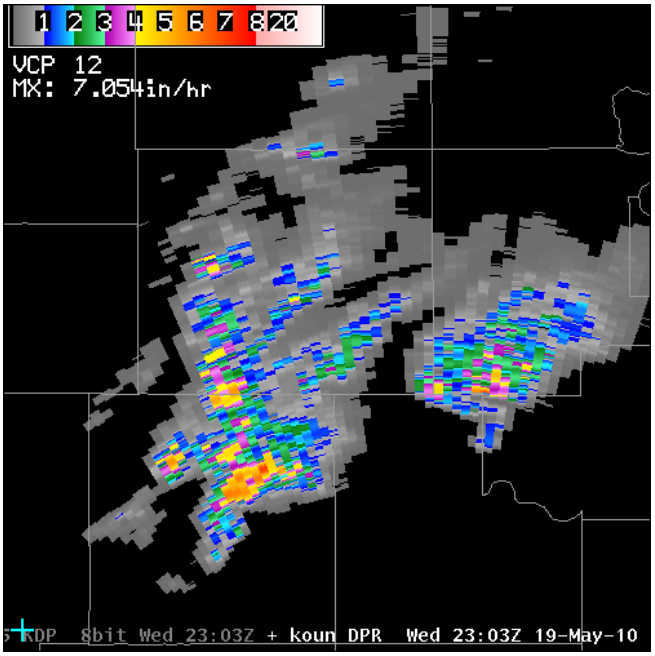

Topic 3, Lesson 8: More Precipitation Estimation

Presented by the
Warning Decision Training Branch

Version: 1210

Lesson 8: More Precipitation Estimation

This is the remainder of the tour of the RPG algorithms that estimate rainfall (two algorithms) and snowfall (one algorithm).

Objectives

1. Identify strengths & limitations of the Quantitative Precipitation Estimation (QPE) algorithm.
2. Identify design similarities between the PPS and QPE.
3. Identify design differences between the PPS and QPE.
4. Identify strengths & limitations of the Snow Accumulation Algorithm.

QPE Tour

The QPE Tour will be somewhat different from the PPS Tour in Lesson 7. The Legacy Precipitation Processing Subsystem (PPS) was presented as a sequence of quality control steps. The QPE Tour is less sequential, but focused on key design elements of the QPE. The three equations that QPE uses to generate rain rates were introduced in Lesson 7. This QPE Tour provides more information on how QPE relies on the output from other Dual-Pol RPG algorithms, which both determine which rain rate equation is used for any given range bin.

QPE Inputs

The 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 (see Fig. 3-127). 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 Preproces-

sor. Reflectivity (Z), Differential Reflectivity (ZDR), Correlation Coefficient (CC), and Differential Phase (Φ_{DP}) are smoothed and, in some cases, corrected for attenuation (see Fig. 3-127). The "prime" notation is 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.

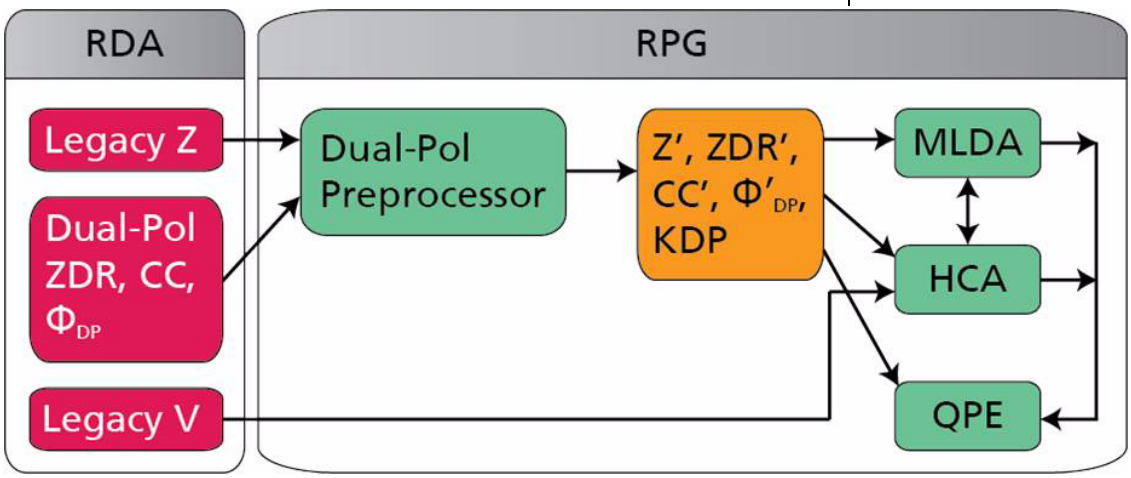

Figure 3-127. Flow chart showing how the QPE product is generated.

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.

QPE Tour: Hydroclass

The QPE relies on the Hydrometeor Classification Algorithm (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.

The Hybrid Hydrometeor Classification (HHC) product is output from the QPE (see Fig. 3-128). It shows you, on a bin by bin basis, which HCA val-

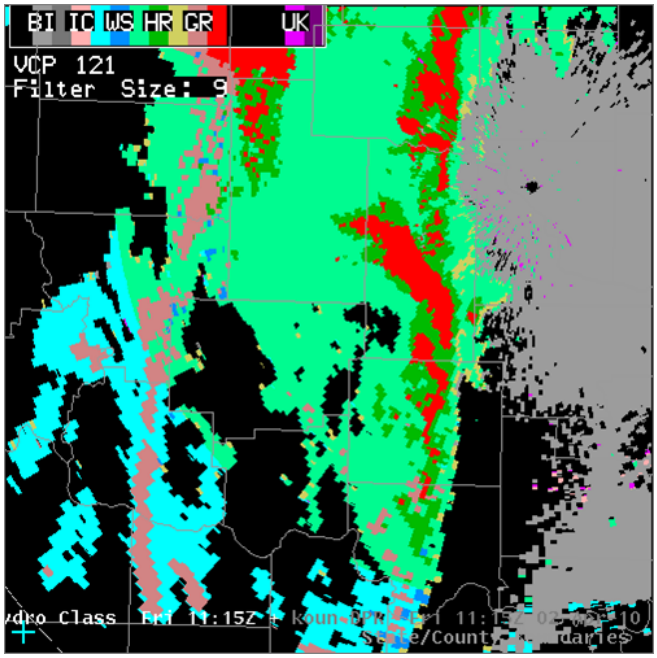

Figure 3-128. Example of the Hybrid Hydrometeor Classification (HHC) product. Notice the biological scatterers near the radar (grey area).

ues 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.

QPE Tour: Melting Layer

The QPE also relies on output from the Melting Layer Detection Algorithm (MLDA). For each elevation angle, the MLDA provides four different heights related to the melting layer (Fig. 3-129, left), which are displayed on AWIPS as an overlay product (Fig. 3-129, right). For the purpose of QPE quality control, the top and bottom of the melting layer are used along with the Hydrometeor Classification values. For example, for the Hydroclass value of Dry Snow (DS), one rain rate equation is used where DS is above the top of melting layer,

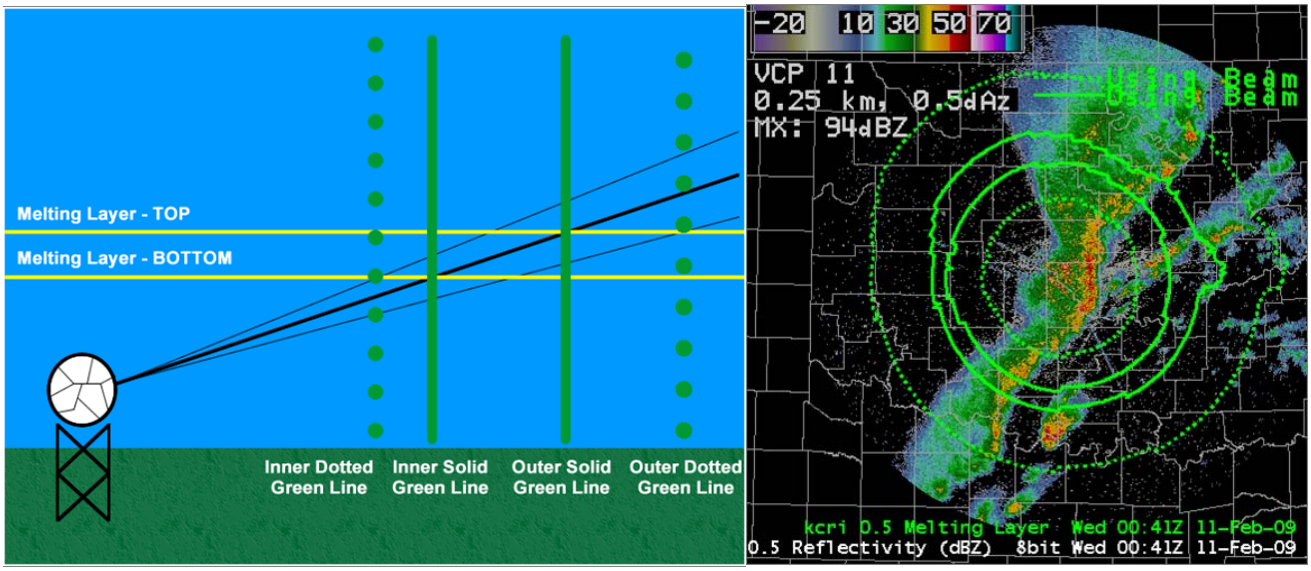

Figure 3-129. A graphical representation of the melting layer detection algorithm (MLDA) product (left) and an image of the MLDA product in AWIPS.

while another equation is used where DS is in or below the melting layer.

Importance of HHC

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 (Fig. 3-128) 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 nine-bin filter. For each bin, the most common hydroclass value for the surrounding nine bins is assigned. This has the effect of reducing speckling on the product.

How QPE Assigns Rain Rate

Based on the HHC, QPE calculates rain rates dependent on the hydroclass value assigned to each bin. One of the most promising aspects of Dual-Pol is the ability to better identify returns to the radar that are from non-meteorological targets. Given that HCA has accurately identified the domi-

Figure 3-130. Color table for the HHC product highlighting Biological, Clutter and Unknown data types which the QPE uses to prevent these returns from being converted to rainfall.

nant radar return for a bin as biological (BI), ground clutter (GC), or Unknown (UK) (see Fig. 3-130), QPE can prevent these returns from being converted to rainfall.

If the bin is assigned biological (BI), the rain rate is set to 0.0 and there is no additional search. It is assumed that precipitation would not be present for any elevation above biological targets.

If the bin is assigned clutter (GC) or unknown (UK), QPE checks the same azimuth and range one elevation higher. The idea is to check for precipitable echo above the clutter.

All of the remaining possible Hydroclass values are some type of precipitation (see Fig. 3-131). QPE uses these Hydroclass values and the height of the bin with respect to the melting layer to determine a rain rate.

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 addition to the three rain rate equations, in some cases, the R(Z) is used, but with a multiplier, such as 0.8* R(Z).

QPE has one product that has no PPS counterpart. It is the Digital Precipitation Rate (DPR) prod-

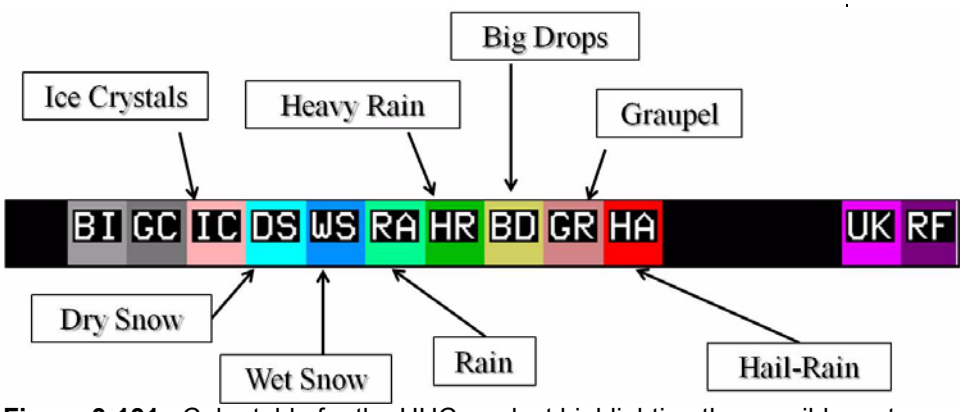

Figure 3-131. Color table for the HHC product highlighting the possible meteorological returns.

uct (see Fig. 3-132). 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.

Extra for the Experts: QPE Rain Rate Equations

Below are the three equations for computing rain rate used by QPE, and the choice of equation is dependent on the Hydroclass value from the HHC,

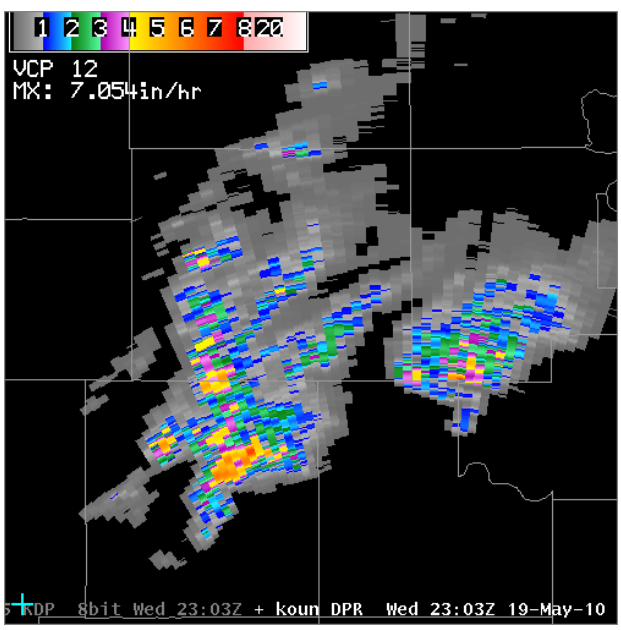

Figure 3-132. The Digital Precipitation Rate (DPR) product presents the rainfall rates that were used to generate the suite of QPE rainfall accumulations products.

the height with respect to the melting layer, and the beam blockage percentage.

$$R(Z) = (0.017)Z^{0.714} \tag{8}$$

$$R(Z, ZDR) = (0.067)Z^{0.927}ZDR^{-3.43} \tag{9}$$

$$R(KDP) = 44.0|KDP|^{0.822}sign(KDP) \tag{10}$$

The coefficients and exponents for these equations were developed based on research conducted in Oklahoma, which is dominated by warm season events. Additional research will be needed to fine tune these equations for other areas.

Extra for the Experts: How QPE Builds DPR

Among the precipitable HHC values, we start with those that are frozen (see Fig. 3-133). All of these rain rates are intended to account for water equivalent at the surface, though bear in mind these multipliers were based on research in central Oklahoma, dominated by warm season events.

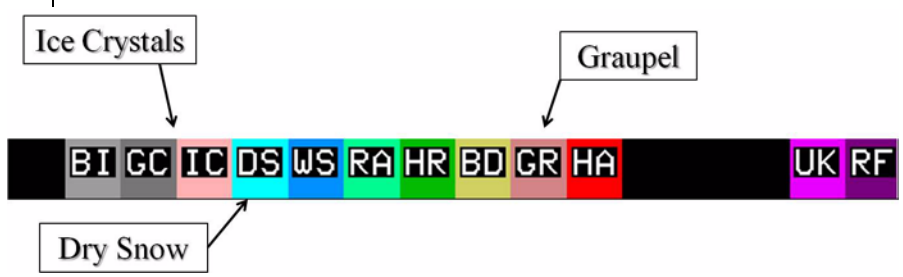

Figure 3-133. Ice Crystals, Dry Snow and Graupel make up the frozen precipitation within the HHC values.

For bins with a hydroclass of Ice Crystals (IC), 2.8R(Z) is used. For Dry Snow (DS) that is above the top of the melting layer, 2.8R(Z) is also used. For Dry Snow (DS) that is in or below the top of the melting layer, R(Z) is used. Finally, bins with a hydroclass of Graupel (GR) use 0.8R(Z).

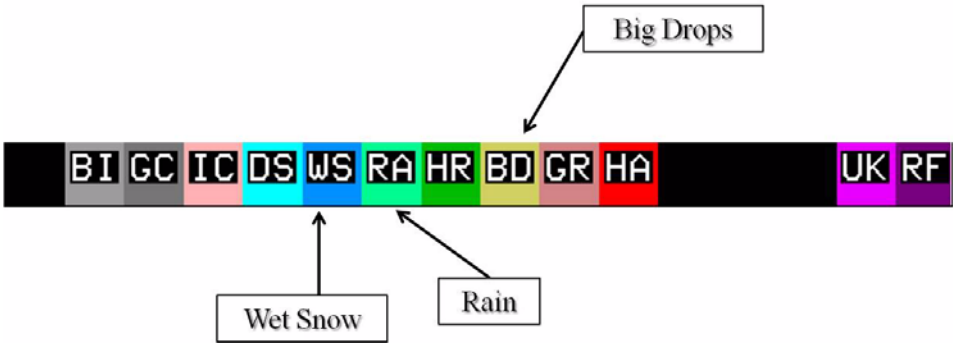

Figure 3-134. Partially frozen or liquid hydrometeors shown in the HHC product include Wet Snow, Rain and Big Drops.

Now for the HHC values of partially frozen or liquid hydrometeors (see Fig. 3-134). For Wet Snow (WS), 0.6 R(Z) is used to calculate the rain rate.

Light to moderate rain (RA), and Big Drops (BD) both use R(Z,ZDR) with no multiplier. For any given Z, the higher the associated ZDR, the lower the rain rate. For example, the highest rain rates come from large numbers of smaller drops, which means a lower ZDR (typically 1.5 to 2.0 dB). As ZDR values get larger, there are more large drops, and the rainfall rate drops. With the R(Z,ZDR) equation, rainfall rate has a direct relationship with Z and an inverse relationship with ZDR. They help to balance one another.

Presented so far have been the rain rate equations for Ice Crystal, Dry Snow, Graupel, Wet Snow, Light to Moderate Rain, and Big Drops. Each of these uses the R(Z) or R(Z,ZDR) equations, in some cases with a multiplier.

Ice Crystals (IC), Dry Snow (DS), Graupel (GR), Rain (RA), Big Drops (BD) and Beam Blockage

In all these cases, there is an adjustment to the Z input if there is beam blockage between 5% and 70%. This “power added back” approach is similar to what the PPS does for partially blocked radials, though the numbers differ. The table in Figure 3-135 has example values from 5% to 70%. If the

Percent Blockage	dBZ Added	Percent Blocked	dBZ Added
5 %	+0.34 dBZ	40 %	+1.97 dBZ
10 %	+0.45 dBZ	45 %	+2.45 dBZ
15 %	+0.58 dBZ	50 %	+3.01 dBZ
20 %	+0.75 dBZ	55 %	+3.65 dBZ
25 %	+0.97 dBZ	60 %	+4.38 dBZ
30 %	+1.24 dBZ	65 %	+5.18 dBZ
35 %	+1.57 dBZ	70 %	+6.05 dBZ

Figure 3-135. Table showing the relationship between percentage of beam blockage and the amount by which the Z input is adjusted.

blockage is greater than 70%, then the algorithm goes to the next higher elevation bin with valid data.

Heavy Rain (HR) as the Hydro Class

If the HHC is Heavy Rain (HR) (see Fig. 3-136), the choice of rain rate equation is dependent on the percentage of beam blockage. If the blockage is greater than 5%, but less than 20%, R(Z,ZDR) is used with power added back to the Z value. If the blockage is between 20% and 70%, R(KDP) is used if possible, since KDP is largely immune to partial beam blockage.

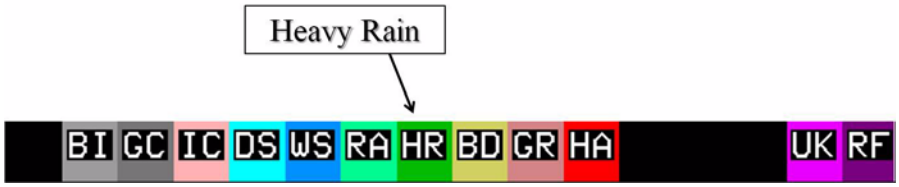

Figure 3-136. When it comes to Heavy Rain, the rain rate equation used depends on the percentage of beam blockage.

Now for a range bin assigned HA, which is best defined as hail possibly mixed with rain (see Fig. 3-137).

Hail Possibly Mixed with Rain (HA) as the Hydroclass

Figure 3-137. Several factors are considered when selecting a rain rate equation when the HHC value is Hail-Rain (HA).

If the bin is above the melting layer and there is <5% blockage, then 0.8R(Z) is used. If the bin is partially blocked (5% to 70%) irrespective of the melting layer height, R(KDP) is used if possible. These are the conditions where R(Z) would be used instead of R(KDP):

1. KDP = No Data
2. CC<0.9 and the Preprocessor algorithm has not identified NBF for that bin
3. R(Z) is less than 10 mm/hr and the Preprocessor algorithm has not identified NBF for that bin
4. R(KDP) is negative

As an example, if HA is tagged below the melting layer, but that bin also has CC < 0.9, R(Z) will be used instead of R(KDP).

When non-uniform beam filling occurs, the most important impact is that the CC values are very low down radial as shown in Figure 3-138. Among

Non-Uniform Beam Filling

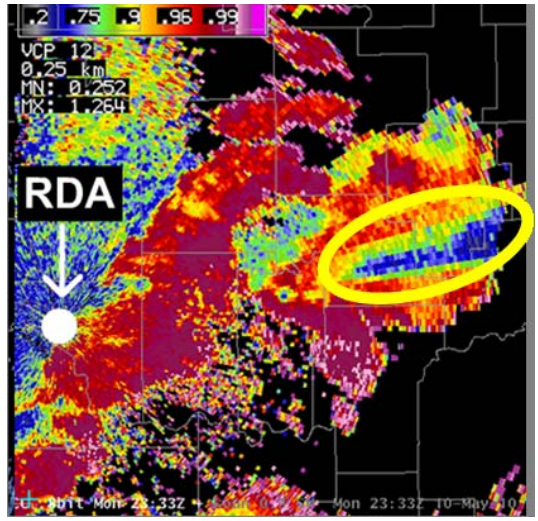

Figure 3-138. Correlation coefficient (CC) data showing the effects of non-uniform beam filling (yellow circle)

many other impacts, this results in bins with No Data (ND) assigned by QPE.

The RPG Preprocessor algorithm has logic that can detect bins that are subject to NBF. For bins that are identified with NBF by the Preprocessor and have ND assigned by QPE, there is a check to see if any of the following has a non-zero rain rate for that bin:

1. R(Z,ZDR)
2. R(Z)
3. R(KDP)

If none of the relationships have a valid rain rate assigned for the bin, it remains ND.

Power is added back to Z if partial beam blockage is present for R(Z,ZDR) and R(Z).

Truncating Rain Rates

Finally, for the generation of the DPR product, there is a universal cap on rain rates, irrespective of the equation used to generate the rate. Any rain rates that exceed 7.87 in/hr, are truncated, such

that the maximum rain rate for the DPR is 7.87 in/hr.

QPE Expectations

The following is a summary of how QPE is designed to mitigate some of the most difficult problems with using a radar to estimate rainfall.

QPE and Bright Band Contamination

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 (WS) 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, CC, and ZDR (see Fig. 3-139), allowing for a base data quality control check for these algorithms. Once the radar is attempting to estimate rainfall from bins that are above the top of the melting layer, the approach is quite different.

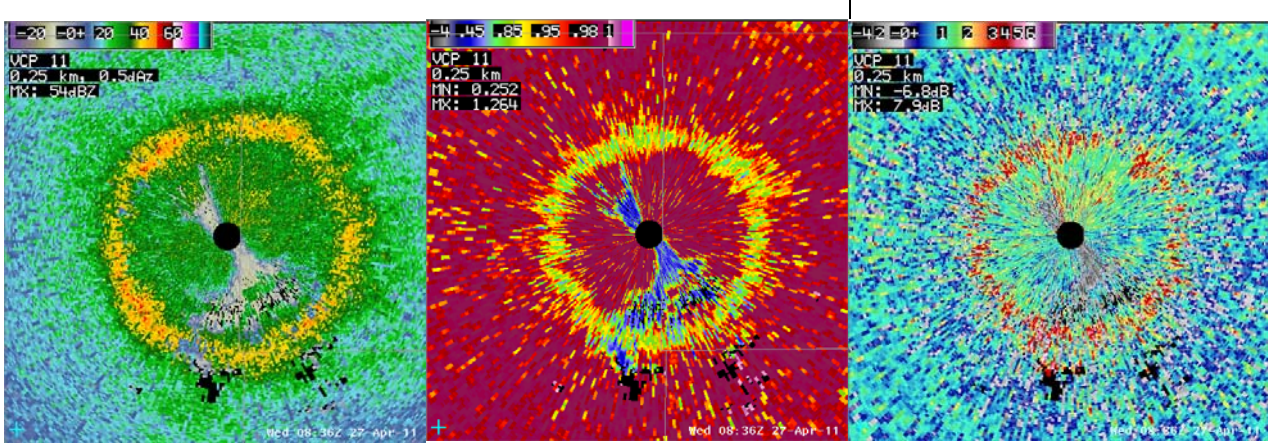

Figure 3-139. Reflectivity (Z, left), correlation coefficient (CC, middle) and differential reflectivity (ZDR, right) data images with a very prominent bright band.

QPE Above the Melting Layer

When the HHC assigns Ice Crystals (IC) or Dry Snow (DS), and the DS is above the melting layer, 2.8R(Z) is used to generate a rainfall rate. This multiplier was developed during testing in central Oklahoma, with mostly warm season events. This is not intended to convert snow to liquid equivalent (i.e. this is not a snow algorithm)!

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 at the top of the MLDA-defined melting layer for long term stratiform events (see Fig. 3-140, right). Here's an example Storm Total Accumulation (STA) product from an lengthy stratiform rain event.

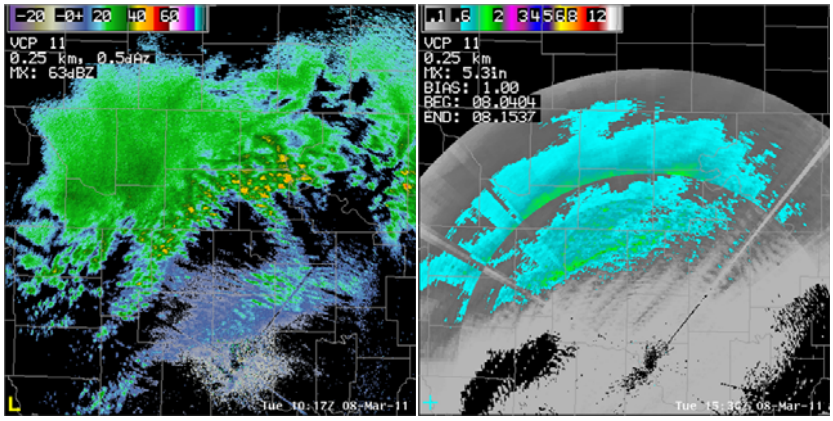

Figure 3-140. The Storm Accumulation Product (STA, right) and the associated reflectivity data (Z, left). Notice the sharp transition in the STA data at the "top" of the melting layer.

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

QPE and Hail Contamination

QPE addresses hail contamination for range bins that are likely to contain hail or graupel, as identi-

fied by the HCA. The approach is to use 0.8R(Z), which lowers the rain rate by the 0.8 multiplier, or to use R(KDP). Specific differential phase (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.

QPE and Non-Uniform Beam Filling

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 Algo-

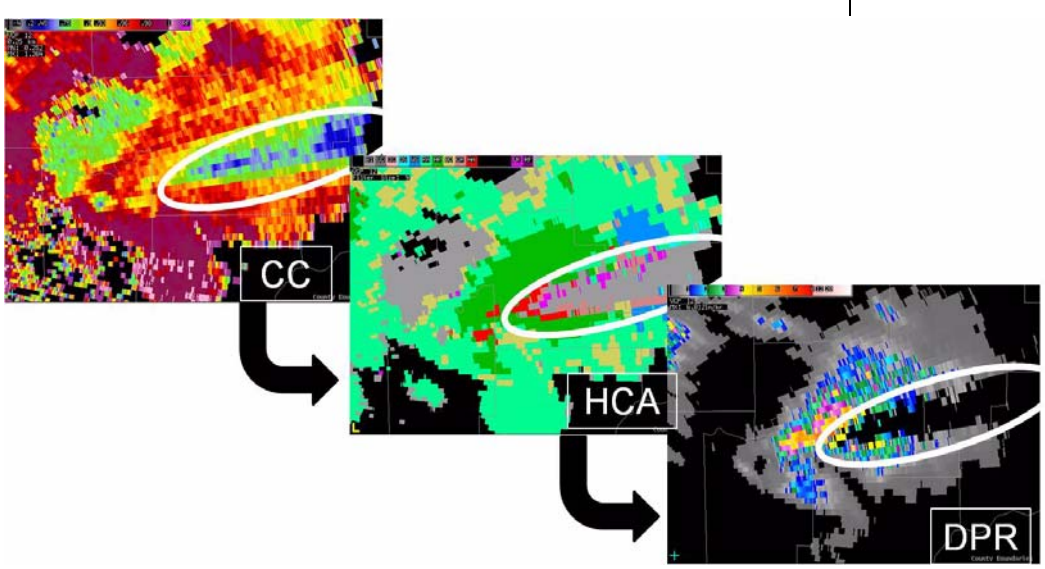

Figure 3-141. Correlation coefficient (CC, left) values affected by non-uniform beam filling will have an effect on the Hydroclass Algorithm (HCA, middle) which then affects the Digital Precipitation Rate (DPR, right).

rithm, and QPE assigns no rain rate to bins with the Biological Hydroclass value (see Fig. 3-141).

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

QPE and Differential Reflectivity (ZDR) Calibration

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. Work is underway to develop techniques to refine calibration of the horizontal and vertical channels.

QPE Strengths

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 the identification of ground clutter and biological returns by the HCA.

For bins that are identified with precipitable returns, QPE has three different rain rate equa-

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

QPE Limitations

As with most new algorithms, there are a number of limitations to QPE that will likely improve over time. The QPE design and associated parameters are based on research in Oklahoma, and regional "tuning" is expected as Dual-Pol is fielded throughout the U.S. 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, due to the difference in rain rate equations at this transition zone.

Similarities and Differences between PPS and QPE

Similar: Storm Total Accumulations

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.

First presented in Lesson 7, Figure 3-142 shows the thresholds for the PPS, along with their default values. RAINA is the areal coverage, while RAINZ is the intensity, or dBZ, threshold (yellow box). RAINA and RAINZ are accessed from the Algorithms window at the RPG under “Hydromet Preprocessing” (white box).

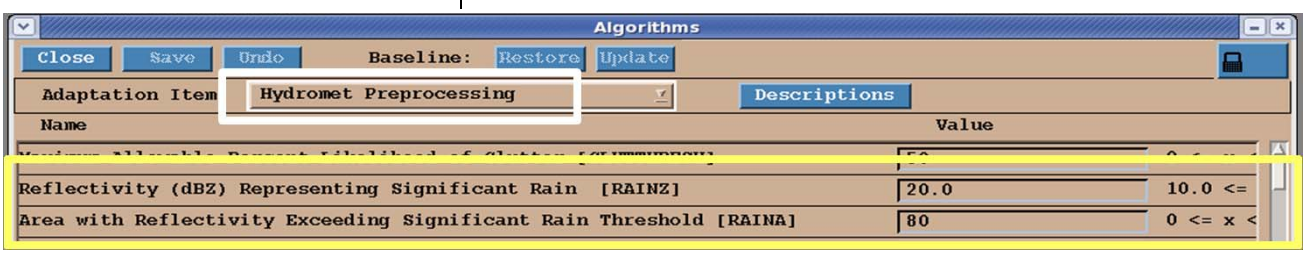

Figure 3-142. The thresholds for PPS can be accessed from the Algorithms window at the RPG under “Hydromet Preprocessing.”

QPE Start and Stop Accumulations Parameters

QPE uses the same concept for start and stop of storm total accumulations, just different threshold names. The QPE thresholds (see Fig. 3-143) 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 (Fig. 3-143, yellow box). 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.

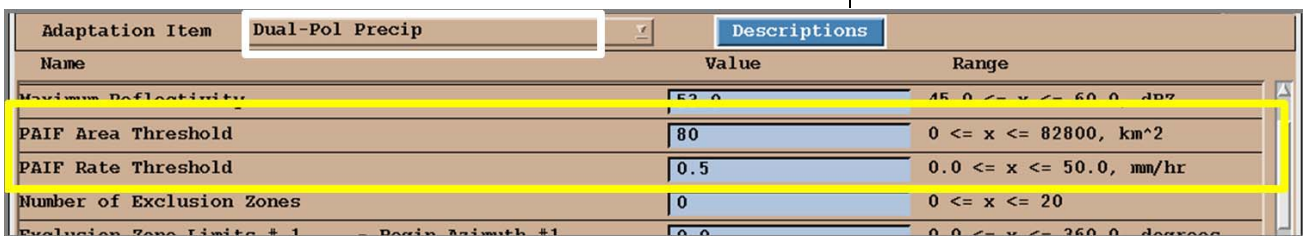

Figure 3-143. To start and stop storm total accumulations, QPE thresholds start with the Precipitation Accumulation Initiation Function (PAIF) (yellow box).

The PAIF thresholds are accessed from the Algorithms window at the RPG under “Dual-Pol Precip” (Fig. 3-143, white box).

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

Similar: Storm Total Accumulations Manual Reset

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 also allow for a manual reset, which can be done from the RPG Control window, selecting the “Legacy PPS” and the “Dual-Pol QPE”, respectively (see Fig. 3-144).

Similar: Exclusion Zones

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.

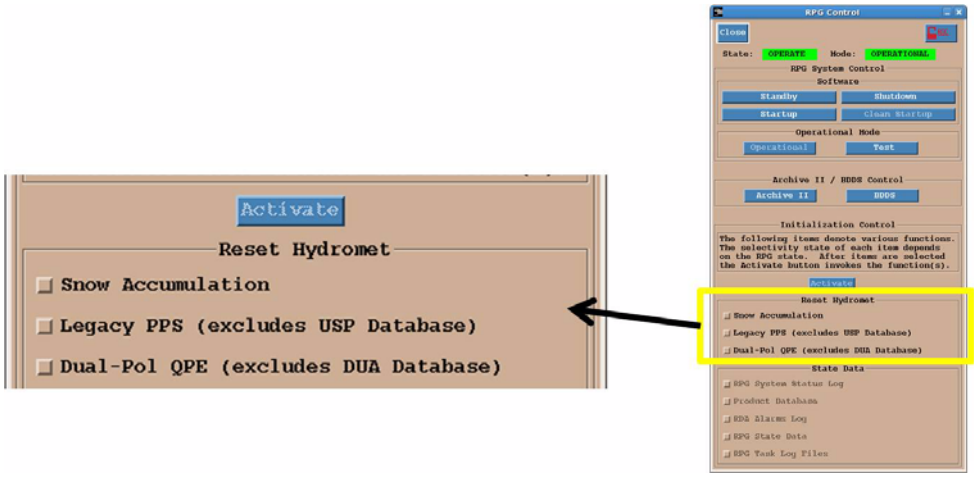

Figure 3-144. Manual resets for both PPS and QPE can be done from the RPG Control window.

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.

Difference: Input

The Legacy PPS relies on the legacy base data: spectrum width, velocity and reflectivity (see Fig. 3-145). 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.

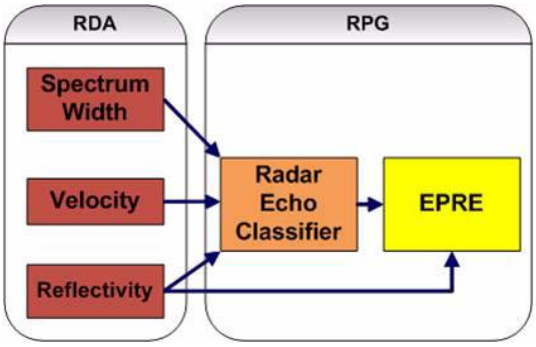

Figure 3-145. The Legacy PPS uses reflectivity, velocity and spectrum width as input.

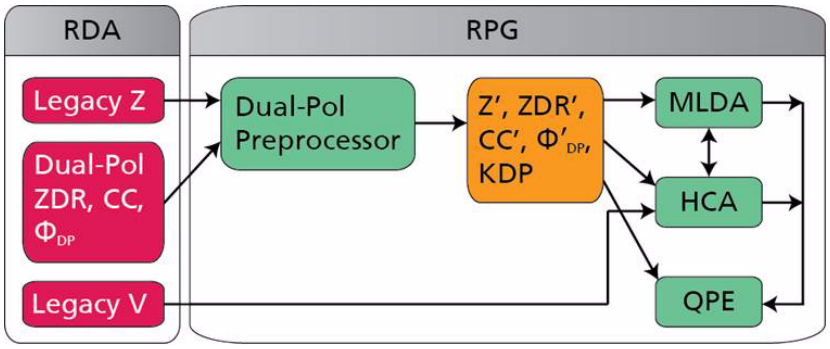

Figure 3-146. QPE uses input from reflectivity, Dual-Pol base data, as well as the Hydroclass Algorithm (HCA) and the Melting Layer Detection Algorithm (MLDA)

QPE relies on the reflectivity and Dual-Pol base data, but only after it has passed through the RPG Dual-Pol Preprocessor algorithm (see Fig. 3-146). 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.

Difference: Rain Rate Equations

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 (see Fig. 3-147).

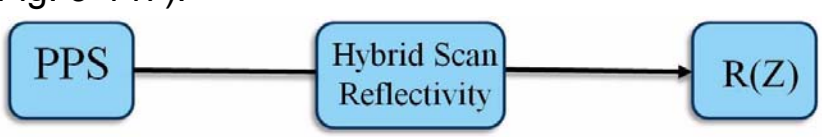

Figure 3-147. The PPS relies solely on the Z-R relationship.

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.

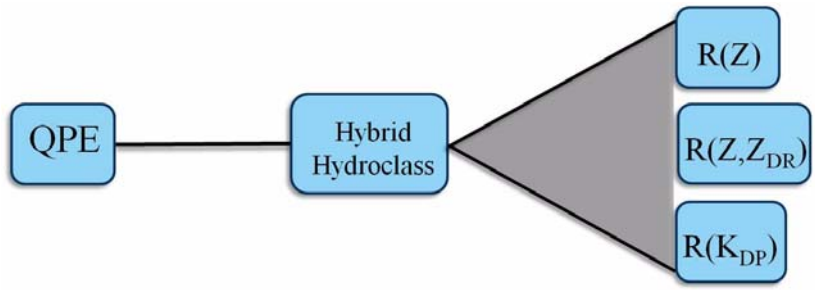

Figure 3-148. Depending on the type of hydroclass indicated by the HHC product, QPE can be calculated by one of three different rain rate calculation methods which are each based on different inputs.

The three different QPE rain rate calculation methods are each based on different inputs, and the notation tells you the input (see Fig. 3-148). 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.

Difference: One Hour Product

There is a significant difference between the PPS and QPE on the initial generation of their respective one hour products, the One Hour Precipitation (OHP) and the One Hour Accumulation (OHA). This difference occurs for the beginning of a 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 QPE OHA will be available beginning with the 2nd full volume scan.

Difference: "Pre-product Product"

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) (see Fig. 3-149, left), which is the dBZ value for each range bin before it is converted to rainfall rate. The QPE gives you the rainfall rate directly, via the Digital Precipitation Rate (DPR) (see Fig. 3-149, right), which is the instantaneous rate for each range bin that is used for the product accumulations.

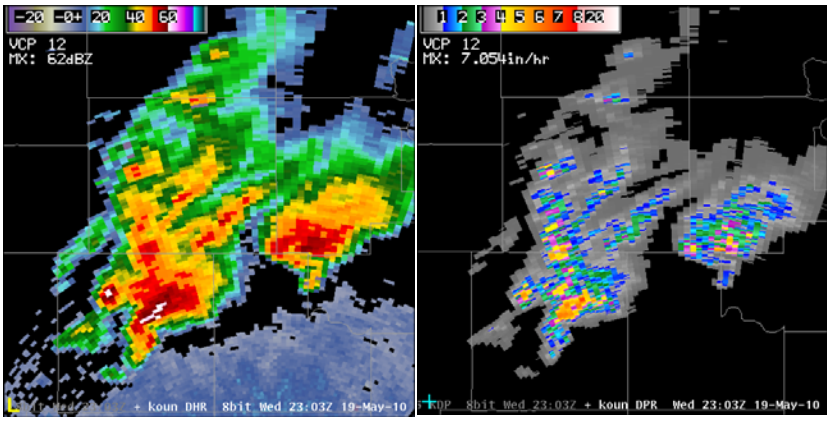

Figure 3-149. In terms of a “pre-product” product, PPS produces the Digital Hybrid Scan Reflectivity (DHR, left) and QPE produces the rainfall directly with the Digital Precipitation Rate (DPR, right).

Similar and Different: PPS, QPE and Valid Bin for Rainfall Estimation

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 (UK) 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%.

Snow Accumulation Algorithm (SAA)

We now focus on the generation of snow water equivalent and snow depth products from the Snow Accumulation Algorithm (SAA) (Fig. 3-150, red box). Like the PPS, the SAA uses EPRE hybrid scan data as input. After data processing, this algorithm produces six total products in AWIPS that represent values for snow water equivalent and depth over various time ranges. More information on the products will be presented in Topic 5.

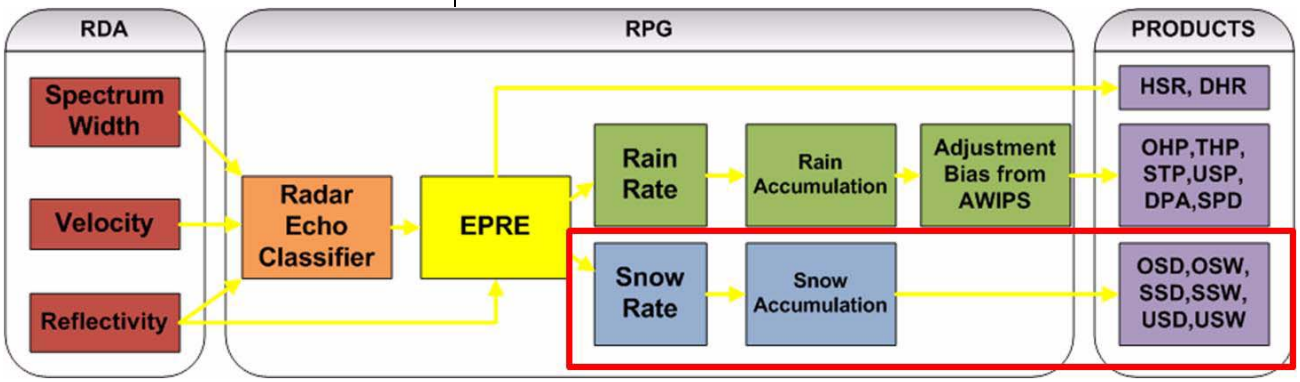

Figure 3-150. The Snow Accumulation Algorithm generates snow water equivalent and snow depth products (red box).

Snow Accumulation Algorithm (SAA) Design

The SAA was developed from data collected at six different research locations. At each of these locations, a network of high-quality snow gauges was used as ground truth against the radar snowfall estimates. The output of this research provided

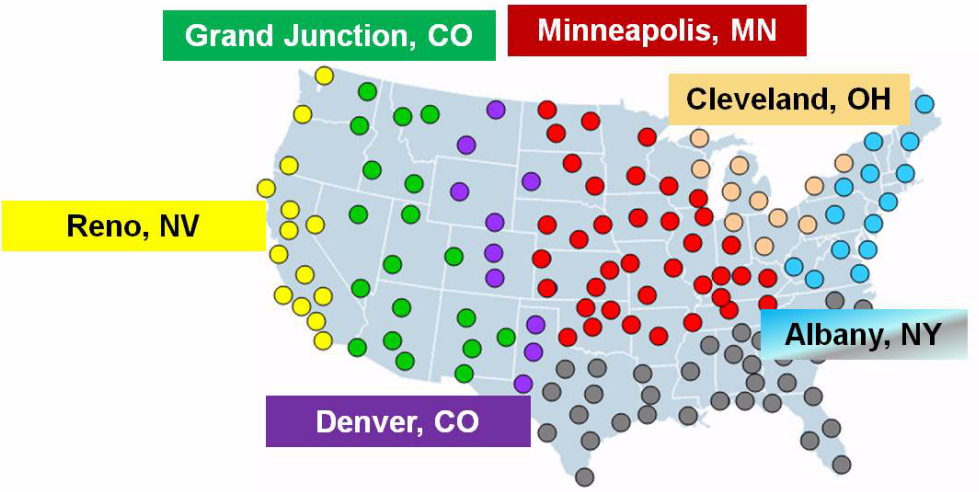

Figure 3-151. The SAA adaptable parameters were developed from data collected at six different research locations.

default adaptable parameters that are used at each of the regions on the map in Figure 3-151.

One of the most important assumptions with the SAA is that it was designed for dry snow events.

Z-S Relationships

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 EPRE Hybrid Scan data as input, the returned power is plugged into a Z-S relationship defined at one of the six research locations and applied to your region. Figure 3-152 lists the default Z-S relationships for each research location.

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

Figure 3-152. Default Z-S relationships for each research location.

Snow Accumulation Algorithm (SAA) and Snow Ratio

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 in Figure 3-153. Even within a given region, it is expected that the appropriate ratio will vary from event to event.

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

Figure 3-153. Default values for converting snow water equivalent to snow depth for each of the research locations.

Snow Accumulation Algorithm (SAA) Adaptable Parameters

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 at the RPG (see Fig. 3-154). Here you can edit the coefficients in the Z-S relationship, and modify the snow-water ratio.

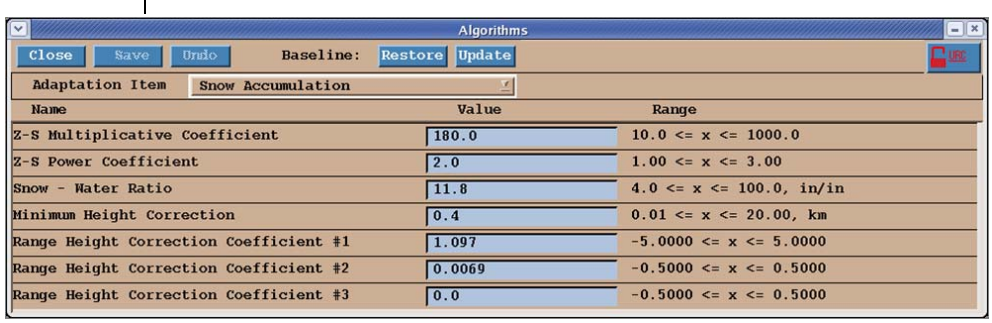

Figure 3-154. There are seven SAA adaptable parameters that are editable under the URC guidelines in the RPG shown here.

Snow Accumulation Algorithm (SAA) Products

There are six snow products generated, all with 16 data levels, a resolution of 1 km by 1° 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 presented in Topic 5.

Begin and End of Snowfall Accumulations

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 (see Fig. 3-155).

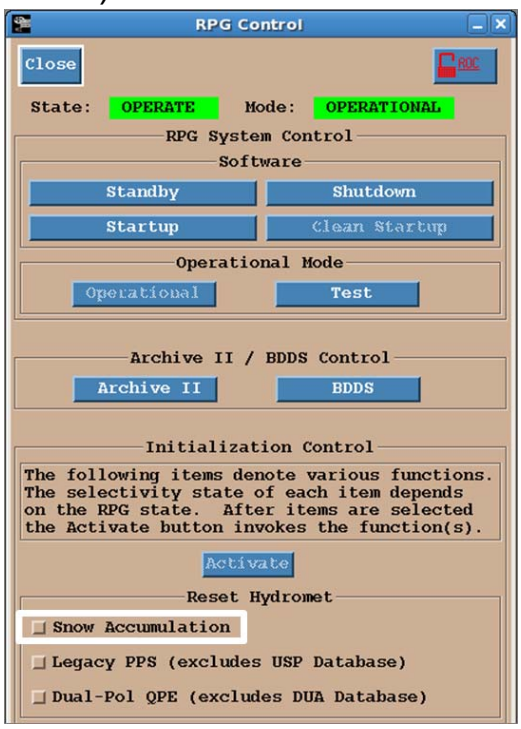

Figure 3-155. Snow accumulations (white box) must be reset at the RPG Control window before each event.

Snow Accumulation Algorithm (SAA) Strengths

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.

Snow Accumulation Algorithm (SAA) Limitations

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 relationships were developed at specific locations and applied across an entire region. As a result, the snow ratio may not be representative for your CWA and will require adjusting.