Dual-Polarization Radar Principles and System Operations

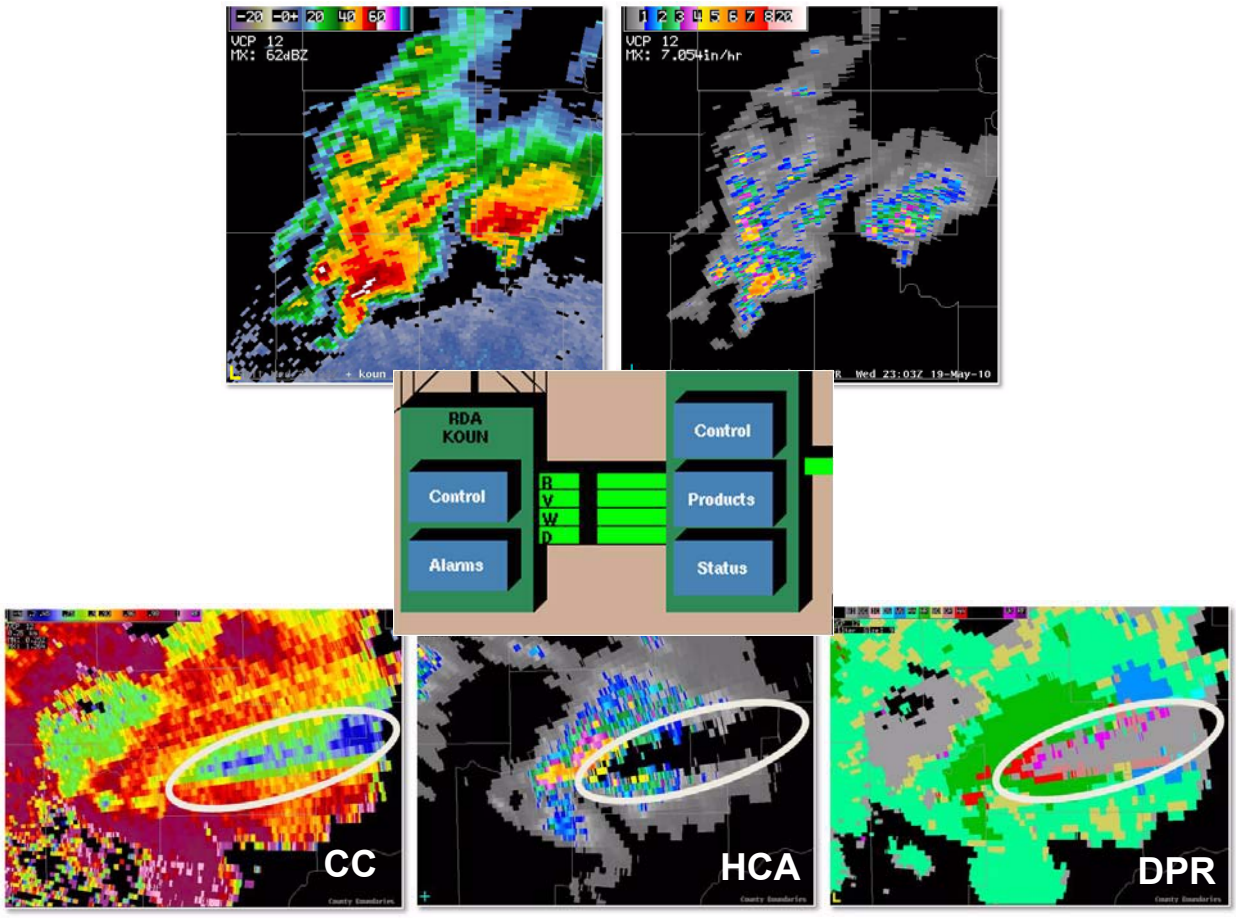

Presented by the
Warning Decision Training Branch

Version: 1109

Table of Contents

Dual-Polarization Radar Principles 5
What Dual-Pol Does Not Change 5
Moments vs. Variables 6
RDA Signal Processing and the Dual-Pol Variables 7
RDA Generation of Differential Reflectivity 8
RDA Generation of Correlation Coefficient 9
Differential Phase (ΦDP) and Specific Differential Phase (KDP) 16

Dual-Polarization Radar Principles 23
Sensitivity Differences 23
Radar Sensitivity 24
Why the Dual-Pol Upgrade Lowers the WSR-88D Sensitivity 25
Calibration 28
Attenuation 31
Depolarization 32
Beam Filling Issues 34
Wet Radomes 38
ZDR and the Trees 39

Life Without CMD and the Dual-Pol Preprocessor 41
Wideband on RPG HCI 41
Clutter Management Without Clutter Mitigation and Decision (CMD) 42
Product Resolution 52
The RPG Dual-Pol Preprocessor 54

Hydrometeor Classification Algorithm (HCA) and Melting Layer Detection Algorithm (MLDA) 61
MLDA Adaptable Parameter 68
Summary 69
Hydrometeor Classification Algorithm (HCA) 69
Summary 78

Quantitative Precipitation Estimation (QPE) Algorithm 79
QPE and PPS in the RPG 79
QPE and PPS Product Data Levels 80
Similar: Storm Total Accumulations 81
Similar: Exclusion Zones 84
Different: One Hour Product 85
Different: QPE Input 86

Different: Pre-product Product 86
Different: Rain Rate Equations 88
QPE and Melting Layer 90
QPE & R(Z,ZDR) 90
QPE and Bright Band Contamination 91
QPE and Hail Contamination 91
QPE and Non-Uniform Beam Filling 92
QPE and ZDR Calibration 93
QPE Strengths 93
QPE Limitations 93
Ongoing QPE Research 94

What Dual-Pol Does Not Change

Though dual-pol is a significant upgrade, it is important to remember that many things do not change. Some of the key hardware components, such as the transmitter, do not change, including:

- Frequency
- Beamwidth
- Range sampling resolution

The structure of the VCPs does not change, including:

- Elevations sampled
- Azimuthal sampling resolution
- Update time
- Number of pulses per radial

All of the pre-dual-pol products will continue to be generated, though there are slight differences which are presented in RDA Lesson 2.

Dual-pol has many benefits, but it will take a lot of time to develop expertise. It is important to bring the dual-pol products into your methodology at your own pace and in a way that keeps you from being overwhelmed.

How have the dual-pol data been integrated into the VCPs? For the Split Cut elevations, the dual-pol data are generated from the Contiguous Surveillance (CS) rotation. This takes advantage of a low PRF with a long R_{max}, meaning that multiple trip echoes are highly unlikely. For the Batch elevations, the dual-pol data are generated from the Doppler pulses, because there are more of them available per radial. For the elevations above Batch, only Contiguous Doppler (CD) mode is used, because range folding is not a concern. Figure 1-1 shows the breakdown for VCP 12.

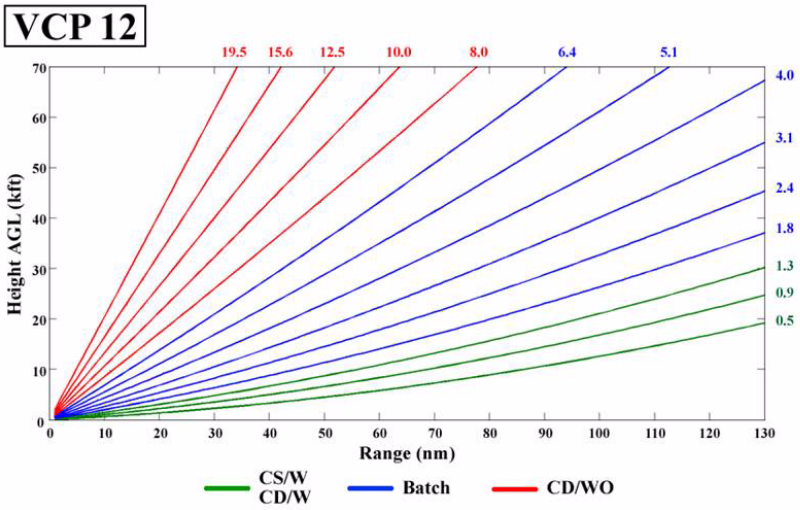

Figure 1-1. Dual-pol data are generated from Contiguous Surveillance (CS) rotation on the Split Cut elevations, Doppler pulses for Batch elevations, and Contiguous Doppler (CD) mode above Batch cuts.

Moments vs. Variables

Dual-pol base data are sometimes referred to as “Variables”, while reflectivity, velocity, and spectrum width are referred to as “Moments”. This is because reflectivity, velocity, and spectrum width are moments of the Doppler Spectrum, which is a distribution of returned signal power as a function of the Doppler Velocity (see Fig. 1-2). You might remember the Doppler Spectrum from previous training modules. We’ll need this tool once again for concepts related to dual-pol.

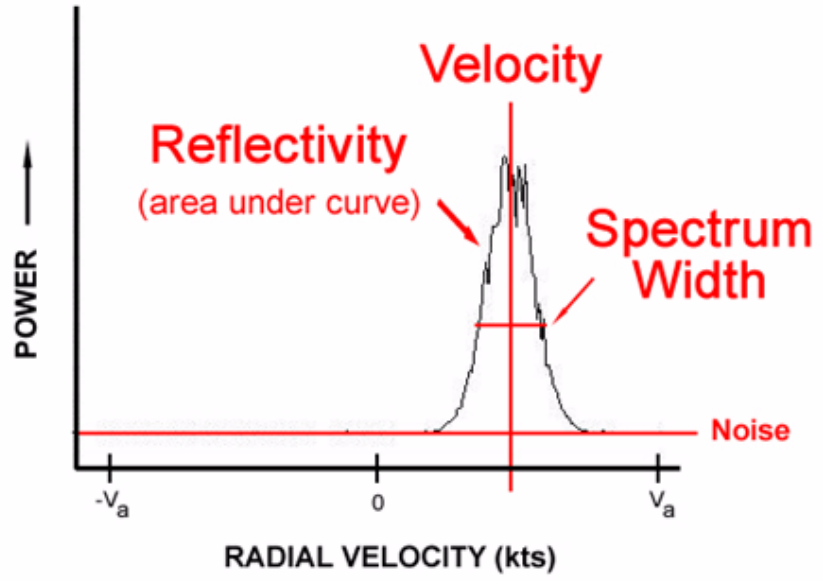

Figure 1-2. Graph showing returned signal power as a function of Doppler Velocity. Reflectivity, velocity and spectrum width are all moments of the Doppler Spectrum.

If you are interested in the specifics, reflectivity is the 0^{th} moment, velocity is the 1^{st} moment, and spectrum width is the 2^{nd} central moment.

RDA Signal Processing and the Dual-Pol Variables

In contrast, we refer to the dual-pol base products as “Variables”. This module presents the basics of dual-pol base data generation at the RDA signal processor, and the products that are then built by the RPG.

The first two, differential reflectivity (ZDR) and correlation coefficient (CC) are generated by the RDA signal processor, then sent to the RPG and generated into a base product.

Differential phase (Φ_{DP}), is also generated by the RDA signal processor. It is base data, but not a product. The Φ_{DP} values are transmitted to the RPG along with CC and ZDR. The RPG generates

the specific differential phase (KDP) product. There is no Φ_{DP} base product.

RDA Generation of Differential Reflectivity

ZDR is defined as the difference between the horizontal and vertical reflectivities, each expressed in units of dB. However, Z_H and Z_V are not calculated directly. The equation in Figure 1-3 is a ***definition*** of ZDR, while ZDR is ***calculated*** from the mean returned power from both the H and V channels.

$$ZDR = Z_H - Z_V$$

Figure 1-3. Equation expressing the definition of differential reflectivity.

Base reflectivity, Z, is still calculated from the returned power in the horizontal channel only.

Using 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 (see Fig. 1-4).

$$ZDR = 10 \log (Z_H \div Z_V)$$
$$= 10 \log (P_H \div P_V) - 10 \log (C_H \div C_V)$$

Figure 1-4. The significance of calibration to both channels.

The ZDR equation can be written as ZDR = 10log (Z_H/Z_V). When the Zs are substituted with returned power, the radar constants and the range terms cancel. Calibration of both channels is very important for an accurate ZDR. The significance of ZDR calibration will be explored in RDA Lesson 2.

RDA Generation of Correlation Coefficient

Correlation coefficient measures the consistency of the H and V returned power and phase for each pulse. This “cross correlation” looks at how the power and phase of one channel compares to the other channel. If the consistency is high, changes with one channel are similar to changes with the other.

CC provides information on the quality of the dual-pol base data estimate and (even better for us) implies information on the nature of the scatterers!

This is in some ways similar to the relationship between spectrum width and velocity. Spectrum width measures the consistency of the phase shifts from one pulse to next, which then relates to the reliability of the associated velocity value.

CC and spectrum width are analogous, but there are some important differences. Base velocity and spectrum width are both calculated from the H channel ***only***. Phase shifts from one pulse to the next are compared. A series of pulse pair phase shifts are averaged for the range bin as a vector sum (see Fig. 1-5).

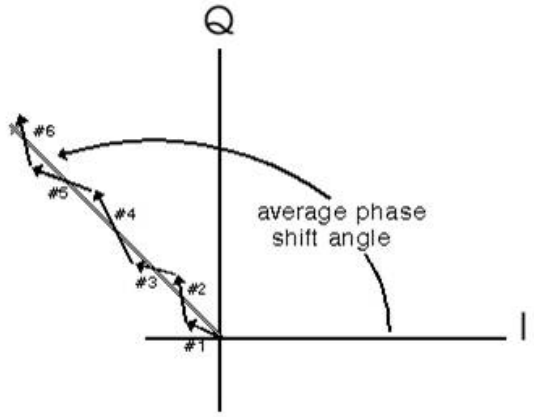

Figure 1-5. Graphical representation of the averaging of pulse pair phase shifts for base velocity and spectrum width.

The greater the variation of these phase shifts, the greater the spectrum width. In the reflectivity image on the left in Figure 1-6, the white circle is an area of weak signal close to an intense supercell. The middle image shows high spectrum width due to both the weak signal and the turbulence. A high spectrum width implies a low consistency of pulse to pulse phase shifts. It is an inverse relationship. Notice that the velocity field on the right is noisy in the same area.

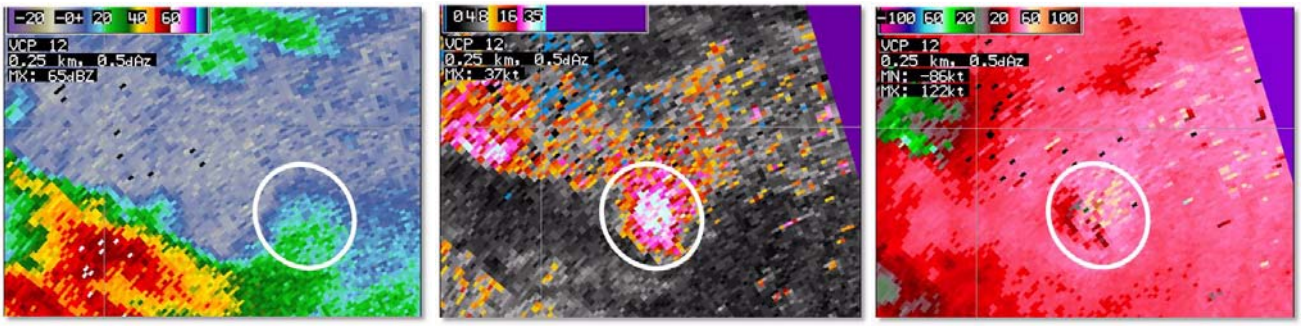

Figure 1-6. Reflectivity (left), spectrum width (middle), and velocity (right). The white circle highlights an area of weak signal in the vicinity of a supercell resulting in high spectrum width and noisy velocity returns.

For each pulse, the returned power and phase from the H and V channels are compared to ***one another***. This is different from the type of correlation that gives us velocity and spectrum width, which is from one pulse to the next.

The consistency that CC measures is based on the angle, Φ_{DP}, between the H and V vectors, which can be determined by vector multiplication (see Fig. 1-7). This “cross correlation” of H and V is checked for each pulse. The vector multiplication of H and V results in the "cross correlation" vector for each pulse, and Φ_{DP} is the angle from the positive x-axis of this vector.

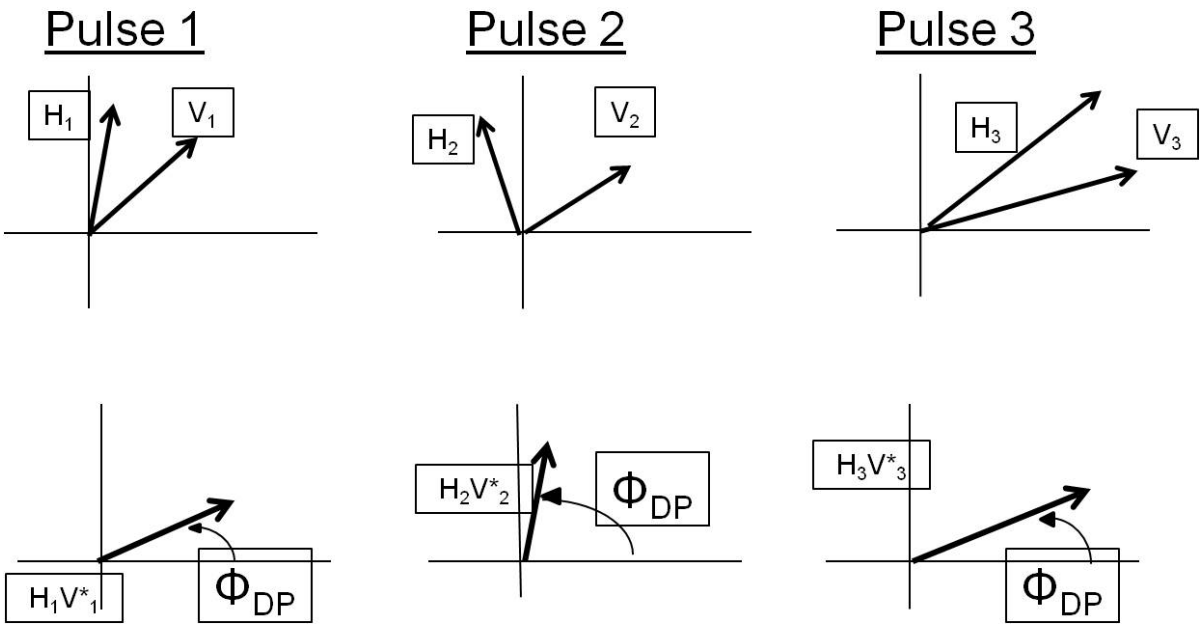

Figure 1-7. Horizontal and vertical vectors from three sample pulses.

Φ_{DP} is important for two of our dual-pol variables.

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

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. Shown in Figure 1-8, this vector sum (red arrow) is what's needed for the remaining two dual-pol variables. 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.

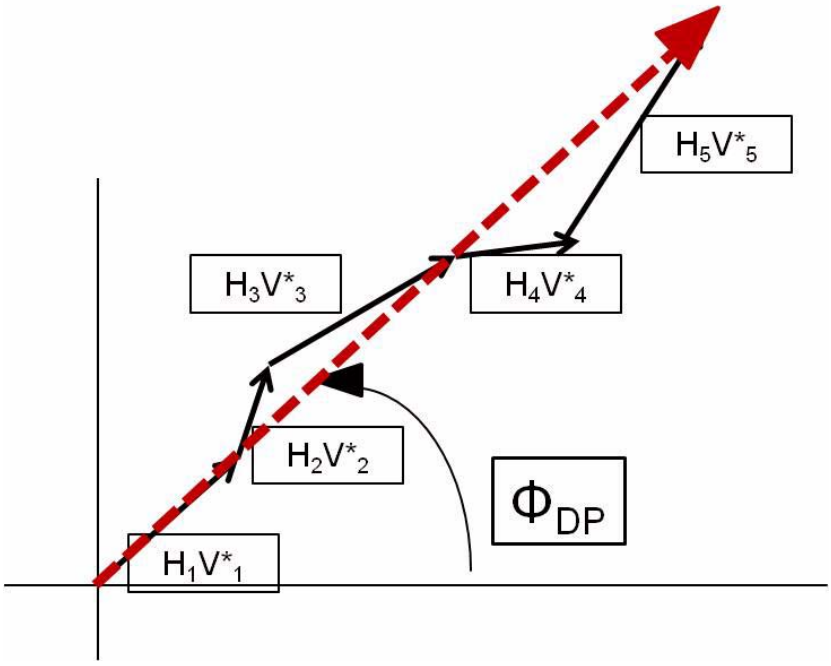

Figure 1-8. Vector sum for the given range bin of five sample pulses, which is known as Φ_{DP}.

CC is calculated by taking the length (or amplitude) of this vector and dividing it by the averaged H and V powers. Though CC is expressed as a number between 0 and 1, you can think of it as a fraction of "perfect" consistency of scatterers. Keep in mind, that CC is never exactly equal to 0 or 1.

If pure rain is being sampled, there is minimal variation between the H and V channels, the cross correlation vectors line up nicely like in the image on the left in Figure 1-9, and CC is close to 1. The more diverse the scatterers, the more variation

with the cross correlation vectors, and CC gets closer to 0.

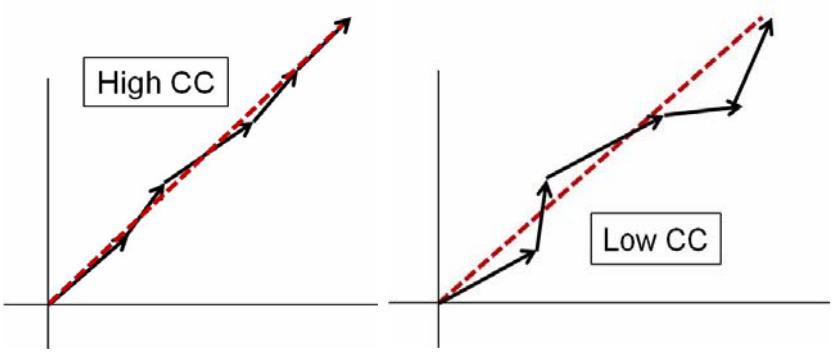

Figure 1-9. Examples of cross correlation vectors for high CC (left) and low CC (right).

A potentially confusing point about the meaning of CC vs. spectrum width is that high CC means high consistency, while a high spectrum width means low consistency. Essentially spectrum width has an inverse relationship with consistency, while CC has a direct one.

Why Does CC Matter?

Low CC (<0.70) implies low consistency between H and V in the estimate and lot of diversity of the scatterers. In fact, CC <0.70 is so diverse the scatterers are likely to be non-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.

In the case of mesoscale melting layer detection (no convective cells), the radar beam is sampling a layer with a mixture of liquid and solid hydrometeors, such as rain and melting snow. This mix of scatterers within the melting layer

creates low consistency of H to V, which lowers the CC value (see Fig. 1-10).

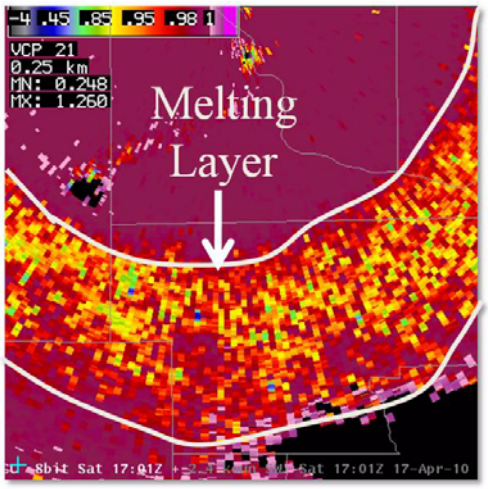

Figure 1-10. CC image showing a clearly defined melting layer.

In areas of weak signal, CC is often noisy in appearance and the magnitudes can vary. In Figure 1-11, near the radar (yellow box) 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 (white circles), 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.

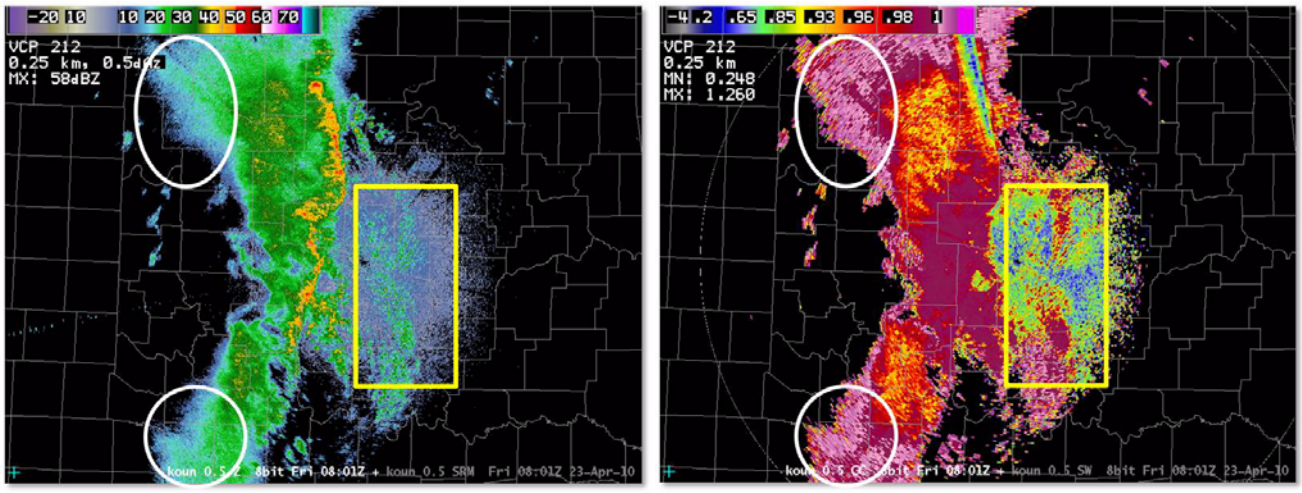

Figure 1-11. Reflectivity (left) and CC (right) images showing the effects of weak signal on CC (white circles).

Though the Z values near the radar and at long range are similar, recall that dual-pol base data are generated from returned power. Thus the weak signal areas are at longer ranges and the associated CC values are less reliable.

CC and ZDR together provide valuable information (see Fig. 1-12). In areas where CC is high, you know that the scatterers are uniform, and the base data estimate error is low. You would expect ZDR to be relatively smooth in appearance, such as the white circled area of pure rain below the melting level. On the other hand, within the melting level (yellow circled area), scatterers are more diverse, CC is lower and ZDR is more noisy in appearance.

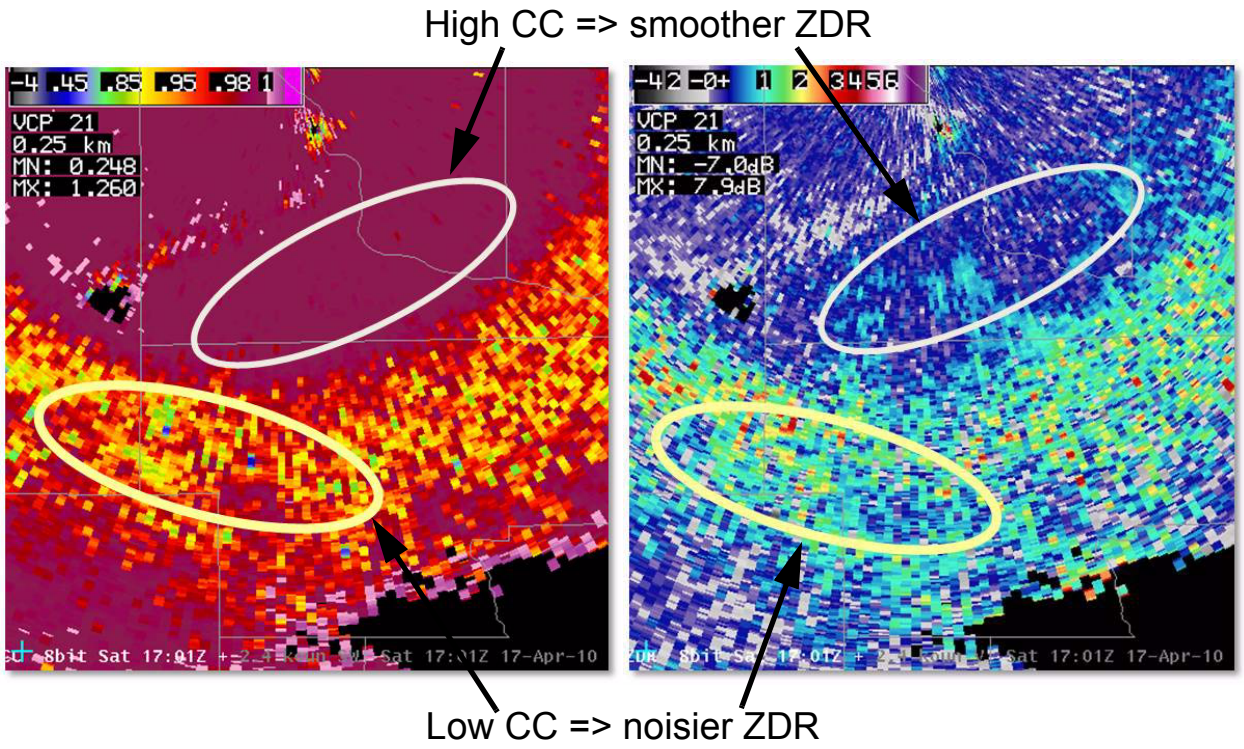

Figure 1-12. CC (left) and ZDR (right) images showing the effects of high values (white oval) and low values (yellow oval) of CC on ZDR.

Differential Phase (Φ_{DP}) and Specific Differential Phase (KDP)

The third 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. 1-13). Since Φ_{DP} is not available on AWIPS, we think of KDP as a dual-pol base product.

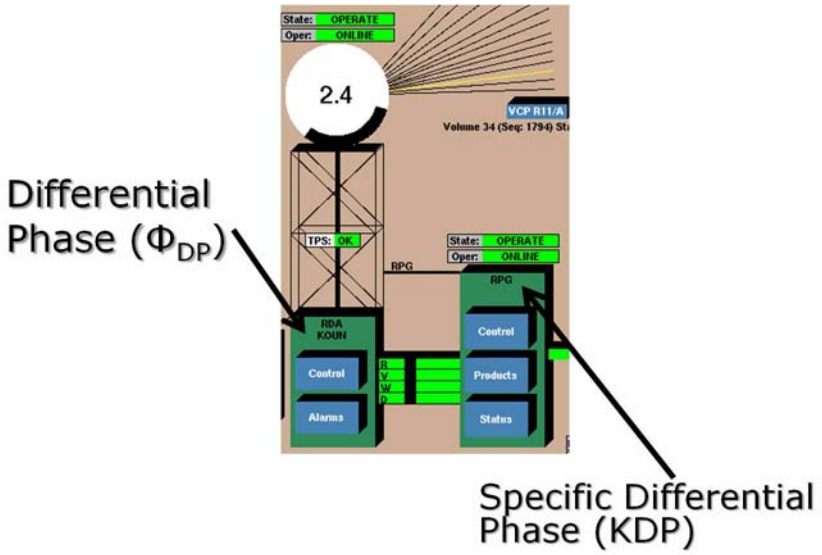

Figure 1-13. Differential phase (Φ_{DP}) is generated at the RDA and sent to the RPG which then generates specific differential phase (KDP).

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}. Recall that for a single pulse, Φ_{DP} is the angle of the cross correlation vector (see Figure 1-14 on page 1-17).

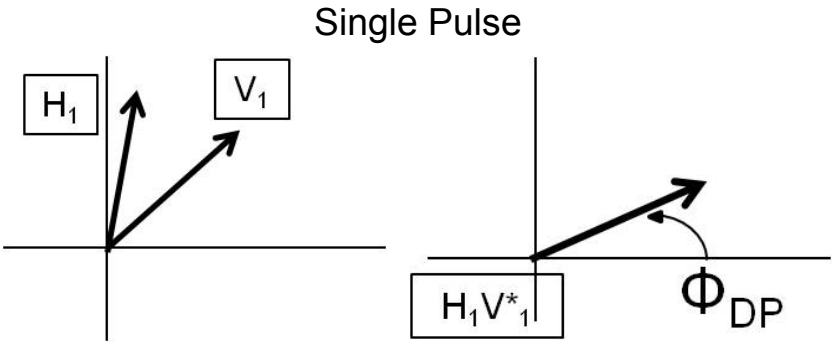

Figure 1-14. Φ_{DP} for a single pulse is the angle of the cross correlation vector.

For a series of pulses, Φ_{DP} is the angle of the vector sum of the cross correlation vectors (see Fig. 1-15).

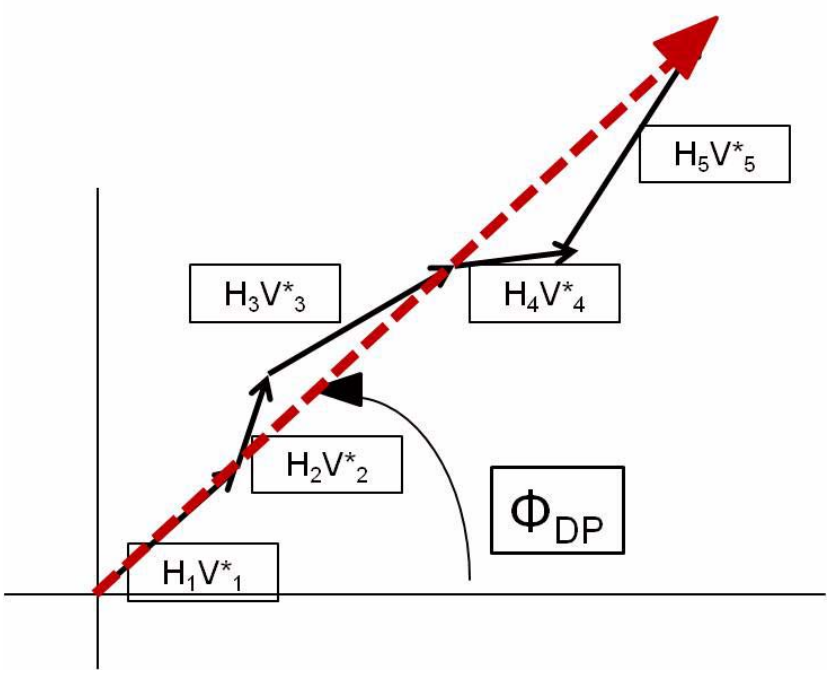

Figure 1-15. Vector sum for the given range bin of five sample pulses, which is known as Φ_{DP} (same as Figure 1-8).

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 H and V, we can compare how the “H delay” differs from the “V

delay". This provides valuable information on the nature of the scatterers that are being sampled.

Liquid water provides "resistance" to the outgoing pulse. In Figure 1-16, the pulse is passing through raindrops, which have a larger horizontal extent than vertical. There is more resistance in the H direction compared to V, creating a longer delay in the H phase compared to the V phase. The returned phase value for H will be greater than for V and Φ_{DP} >0.

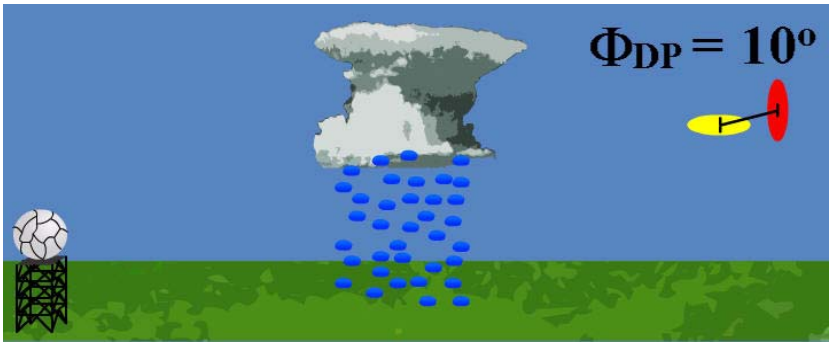

Figure 1-16. Differences in delays from horizontal and vertical channels can provide information about the medium through which the pulse has passed.

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

Below are two examples that would result in a slower H propagation compared to V, and thus a higher Φ_{DP}:

1. If the beam is passing through large raindrops, there is more propagation delay in the H direction than in the V direction (see Fig. 1-17, left).
2. Assume the same size and shape raindrops in each of two sample volumes. However, there is a greater concentration of them in Figure 1-17 on the left. This greater concentration, which means more liquid water content, also

creates a greater propagation delay in the H direction than in the V direction.

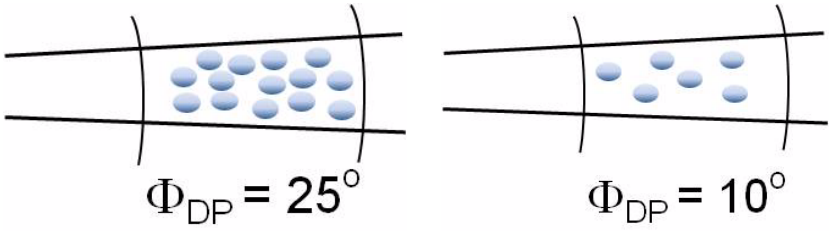

Figure 1-17. Greater concentration of raindrops leads to greater propogation delay in the H direction than in the V direction.

This relationship of Φ_{DP} to liquid water content is what makes this dual-pol variable so important!

How Φ_{DP} Changes Along the Radial

Since Φ_{DP} is dependent on propagation, the values accumulate down radial. There is no way to "reset" the phase shift as the pulse travels outbound, encountering one or more areas of precipitation.

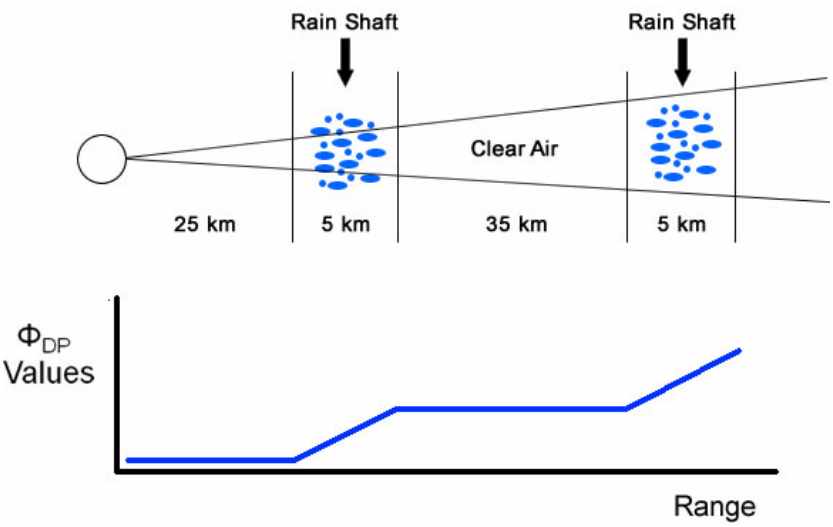

Figure 1-18. Because Φ_{DP} cannot be reset as the pulse travels outbound, the value is increased with each area of liquid scatterers it encounters.

In Figure 1-18, 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. Through this process, the Φ_{DP} value does not “reset” to 0. This makes interpretation of Φ_{DP} as a base product very challenging.

Why KDP? Specific differential phase or 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 the Φ_{DP} changes over very short ranges, which gives us more useful information. You can think of it as a “local” variable. Thus the units for KDP are degrees per km.

The equation in Figure 1-19 does not represent the actual calculation of KDP. It is used to represent the concept of subtracting Φ_{DP} over a range interval. The actual calculation involves a least squares fit of multiple differences along the radial, centered at the range bin.

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

Figure 1-19. Equation representing the concept of subtracting Φ_{DP} over a range interval.

The span of range bins used for the KDP calculation is dependent on the Z value:

- Z ≥ 40 dBZ
  - KDP is based on an integration of 9 bins (4 bins back and 4 bins forward along the radial)
  - There is less smoothing required, and fewer bins are used

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

The KDP calculation fits a line to the varying Φ_{DP} values along the 9 or 25 bin chunk of the radial.

Figure 1-20 (left) depicts Φ_{DP} values with range for a low diversity of hydrometeors (high CC), such as pure rain or snow. The graphic on the right depicts Φ_{DP} values with range for a high diversity of hydrometeors (low CC), such as mixed rain and snow. In both cases, KDP is the slope of these lines.

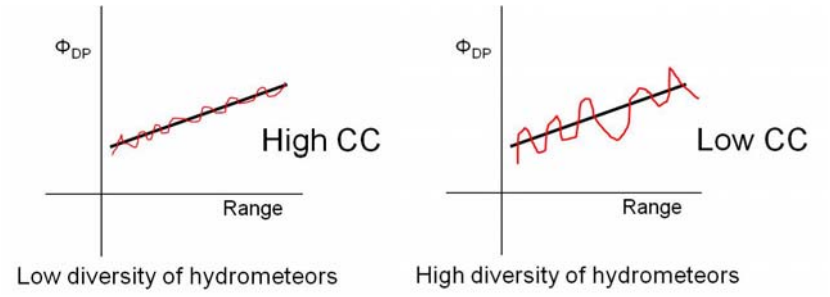

Figure 1-20. KDP is the slope of a line representing Φ_{DP} as a function of range.

KDP is only calculated for bins with CC > 0.90. The idea is to limit KDP to range bins which are likely to contain precipitation.

KDP and CC From their respective computation methods, KDP and CC are related to one another (see Fig. 1-21). When CC is high, close to 0.99, the hydrometeors are uniform and both CC and KDP are very smooth in appearance (white circles). When CC approaches 0.90, the hydrometeors are mixed in size and shape, and CC and KDP are noisier in appearance, such as in the melting layer (white boxes). Also notice that within the melting layer there are some KDP range gates with no data. This is because the CC values are < 0.90, and KDP is not generated for bins where CC < 0.90.

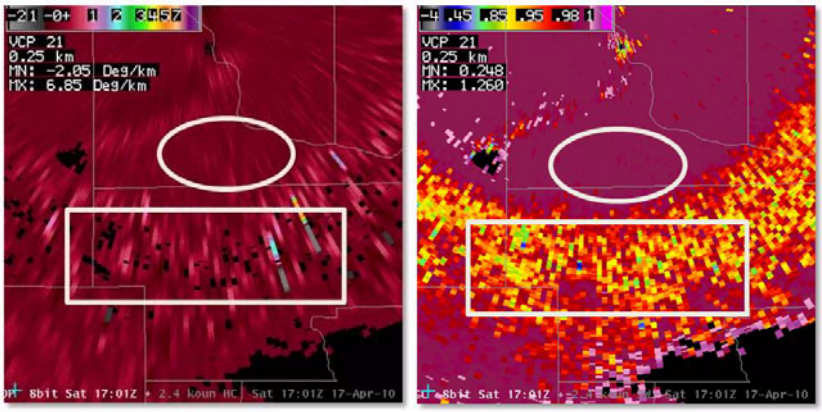

Figure 1-21. KDP appears smoother with high CC (white circles). For CC values close to 0.90, KDP appears noisy. KDP is not generated for bins where CC < 0.90.

Sensitivity Differences

As a result of their design, VCPs 31 and 32 have had significant sensitivity differences since the inception of the WSR-88D. Visible differences in areal coverage between these VCPs may have been noticed without awareness that sensitivity is the underlying cause. VCP 31 uses long pulse, which means power is transmitted for a longer time. This increases the "power on target" and thus increases the returned power. VCP 31 is about 4.8 dB more sensitive than VCP 32 (and the other short pulse VCPs). This means that for any given range, there may be targets that are detectable by VCP 31, but ***not*** by VCP 32.

Figure 2-1 shows an example of these two VCPs during a freezing drizzle event. There is more coverage of this very light precipitation with VCP 31 because of its better sensitivity. The significance of sensitivity is also dependent on the type of weather. The weather events that are most impacted by the decrease in sensitivity after the dual-pol upgrade are freezing drizzle and light snow.

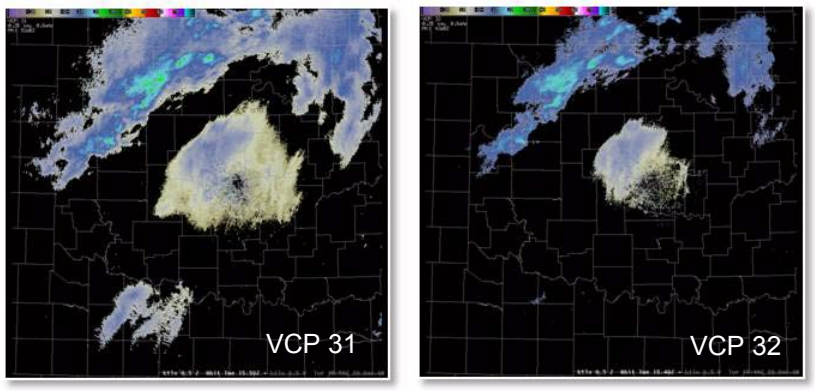

Figure 2-1. Imagery from a freezing drizzle event showing the differences between long pulse VCP31 (left) and short pulse VCP 32 (right).

Radar Sensitivity

Sensitivity defines the minimum signal that a radar can detect at a given range. Sensitivity is expressed as a dBZ value at a given range for a returned signal that is above the system noise by a threshold amount. We'll call this the power threshold.

Transmitted power and system noise are major contributors to sensitivity. The more power transmitted, the more power received, and the more likely we are to have returns above the power threshold, like the image on the left in Figure 2-2. The greater the system noise, the greater the power threshold and the less likely the returned power is sufficiently high. On the right, the system noise is much higher than any radar you would ever want to use!

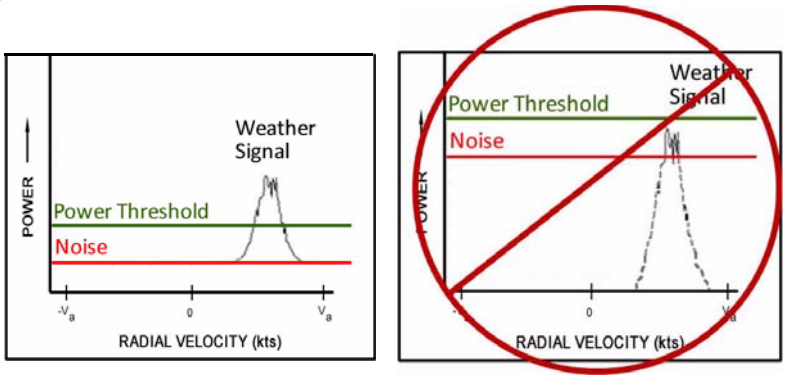

Figure 2-2. Only returned power above the power threshold (left) can be used to generate valid base data. If the system noise and power threshold are too high (right), the error is too high and the base data are not sent to the RPG.

The WSR-88D can only "see" targets with returned power that exceeds the power threshold.

The combination of transmitted power and system noise determines whether or not the returned signal of a given target is sufficiently strong for valid weather data. If the returned signal is below the power threshold, it would have too much error to be acceptable, and the base data are not

transmitted to the RPG. For a given target, sensitivity determines whether or not data are assigned to the relevant gate on a radar product.

For the same system noise, the higher the transmitted power, the higher the returned power, and the better the sensitivity. The two graphics in Figure 2-3 show a strong signal for single vs. dual-pol, but the dual-pol returned power is lower.

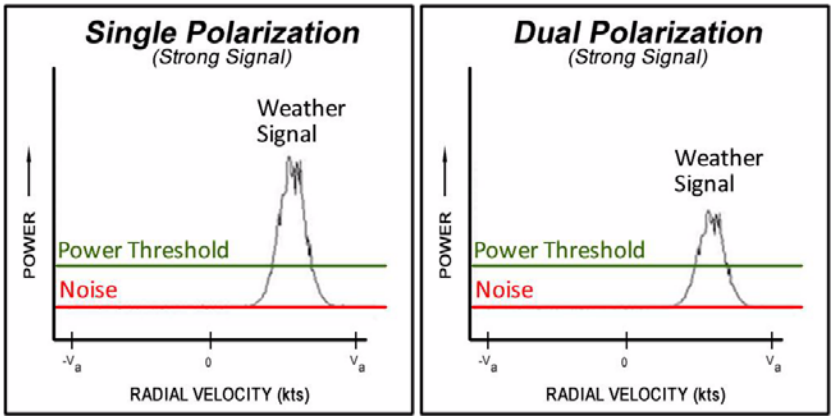

Figure 2-3. For a strong signal, the returned power is higher for single polarization (left) than for dual polarization (right).

Why the Dual-Pol Upgrade Lowers the WSR-88D Sensitivity

The reduction in sensitivity starts with a basic requirement of dual-pol: the WSR-88D must be able to transmit and receive waves that are polarized both horizontally and vertically oriented as shown in Figure 2-4.

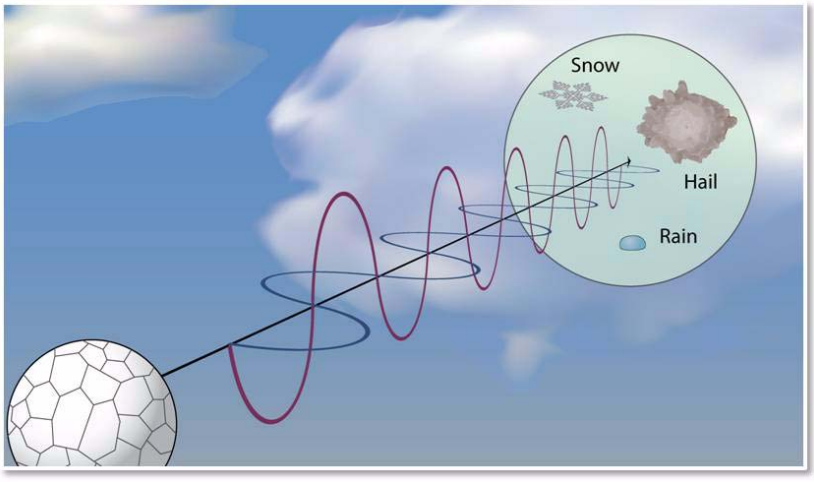

Figure 2-4. The very definition of dual-pol, it's horizontal and vertical waves, cause a reduction in sensitivity.

It is important to remember that returned power is not the same as dBZ, and this will be addressed as part of the discussion on the calibration of Z.

Splitting Transmitted Power Into H and V Channels

So how to transmit a Horizontal (H) and Vertical (V) signal without buying an extra transmitter? The solution is to split the transmitted pulse into the H and V channels, then send them to the antenna. Upon return, the H and V channel pulses are analyzed separately (see Fig. 2-5).

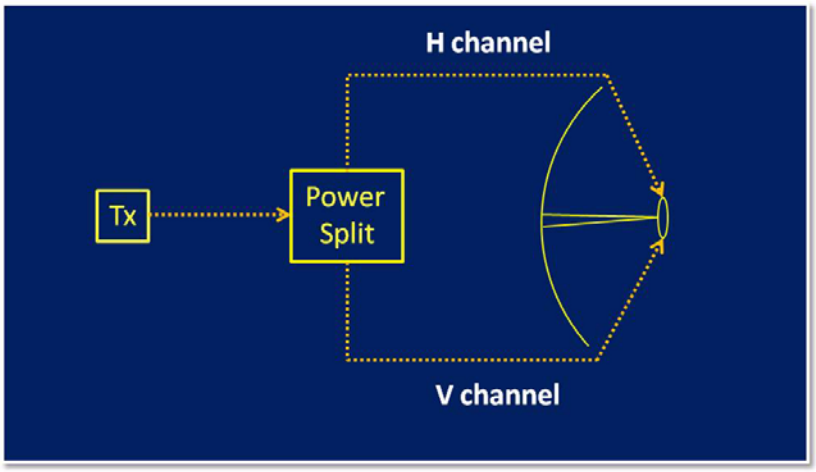

Figure 2-5. Instead of adding another transmitter, the transmitted pulse is split into H and V channels which are then analyzed separately.

Splitting the transmitted power, with one half going into each channel results in a 3 dB loss in transmitted power per channel. The new hardware that is part of the dual-pol upgrade contributes a bit more loss. The engineering group at the Radar Operations Center has performed a very careful analysis of the expected sensitivity loss due to the installation of dual-pol. For any given WSR-88D site, the sensitivity loss due to the dual-pol upgrade is expected to be 3.5 to 4 dB.

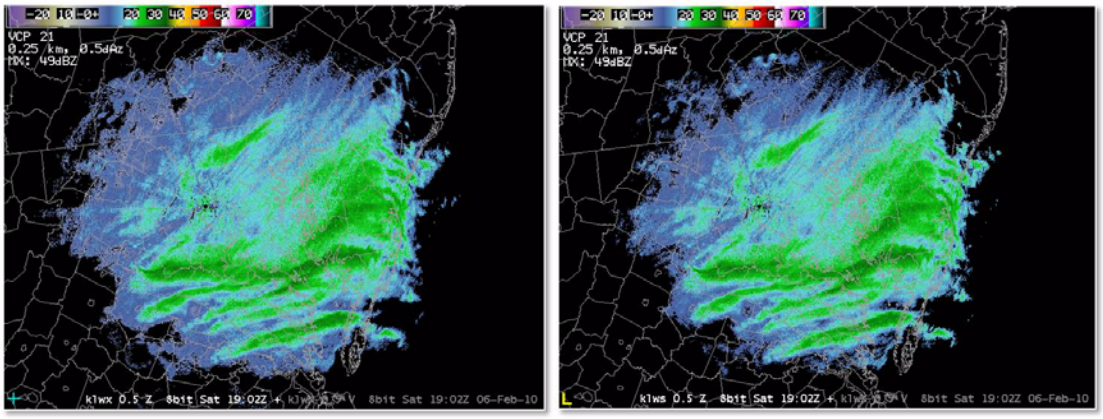

Figure 2-6. Images from a snow event showing full sensitivity (left) and the same image adjusted for a 4dB sensitivity loss (right).

Impacts of the Loss of Sensitivity

So what does a 4 dB loss in sensitivity look like? First of all, it is a lot less than the difference between VCP 31 and VCP 32. Figure 2-6 shows an example of 4 dB loss with a snow event. The image on the right has the same data, but it has been adjusted to simulate a 4 dB sensitivity loss.

A 4 dB sensitivity loss ***does not*** mean that all the dBZ values decrease by 4 dB. What it ***does*** mean is that there are a few less gates of data in the weak signal areas. The impact of sensitivity loss is limited to very weak signal areas.

There has been a series of evaluations of the operational impact of this sensitivity loss for several years now. The final evaluation was conducted by a group of 20 NWS forecasters who had sufficient training to review a variety of weather cases with dual-pol data included. In each of these evaluations, the groups concluded that the benefits of dual-pol outweigh the impacts of the loss of sensitivity.

Figure 2-7 presents results from the evaluation of the group of 20 forecasters, most from the NWS. The yellow bars represent their perception of the effectiveness of the WSR-88D without dual-pol for each of the weather hazards. The green bars represent their perception of the effectiveness of the WSR-88D with dual-pol.

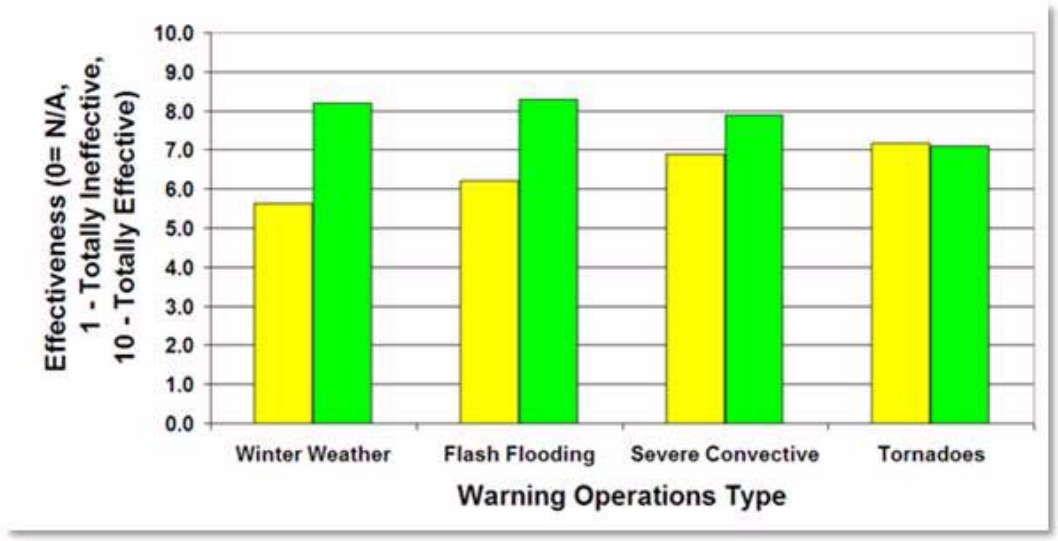

Figure 2-7. Forecasters' average response for evaluating effectiveness of WSR-88D data without dual-pol (yellow bars) and WSR-88D data with dual-pol (green bars).

Calibration

In the most general sense, calibration means assigning a "correct" value. Single-pol radars are calibrated to assign the "correct" dBZ value to any given range gate. Correct is in quotes because there are always trade offs that prevent perfection, but any error is within acceptable limits.

With dual-pol, reflectivity (Z) and differential reflectivity (ZDR) require calibration. Though the calibration of Z will remain transparent, ZDR calibration does have an impact.

In the snow storm example in Figure 2-6, the loss of data on the "fringes" is due to the lowered sensitivity. For the bins that have data assigned, notice that the dBZ values are the same. The

magnitude of the dBZ value is dependent on the radar's Z calibration.

Reflectivity

It's always been tempting to draw conclusions about Z calibration by comparing your radar data to an adjacent radar sampling the "same" location. Differences such as atmospheric propagation, beam blockage, and frequencies must be accounted for. It is not expected that the dual-pol upgrade will cause any additional differences when comparing a radar to its neighbor.

When large targets are sampled, the frequency difference alone can cause differences in dBZ values due to non-Rayleigh scattering. So when comparing the dBZ values in a storm core with hail that is equidistant from two radars, the frequency difference may account for the difference in the dBZ values. This has nothing to do with dual-pol.

Sensitivity vs. Calibration

Sensitivity is a characteristic of a radar system, whereas calibration is a process. Sensitivity ultimately determines whether data are assigned to a given range gate on a product. With the upgrade to dual-pol, there will be a little less return from the weak signal areas. Another way to look at it is that sensitivity determines the "footprint" of all the radar base data. Z calibration determines the magnitude of the value that is assigned. Calibration determines the magnitude of the Z or ZDR value that is assigned.

Differential Reflectivity

Calibration of ZDR is a much more challenging process than Z, because both the H and V channels must be calibrated separately. The goal of ZDR calibration is to stay within 0.1 dB of the true ZDR value, which is stringent. Dual-pol algorithms, particularly Quantitative Precipitation Estimation (QPE), are very sensitive to ZDR

calibration errors. An imperfect ZDR calibration may be acceptable for ZDR base product interpretation by humans, but has a much greater impact on algorithm performance.

Calibrating Z and ZDR

Calibration of *Z* and ZDR will continue to be a combination of on-line and off-line procedures. There are two on-line checks: a short procedure at the end of each volume scan, and a longer procedure every 8 hours. The 8-hour Performance Check is not new, but thus far has been short enough to be nearly transparent to you. With Dual-Pol, the 8 hour check has more tasks and takes just over two minutes to complete. The 8-hour interval is set at each radar independently. It is not synched throughout the network.

There is no 8-hour check notification on the RPG HCI main page. The 19.5° trace on the radome simply pauses. If you notice that your radar data seem delayed, first check the RPG HCI to see if there is a 19.5° donut in the radome as in Figure 2-8. That pause is likely associated with the 8-hour check.

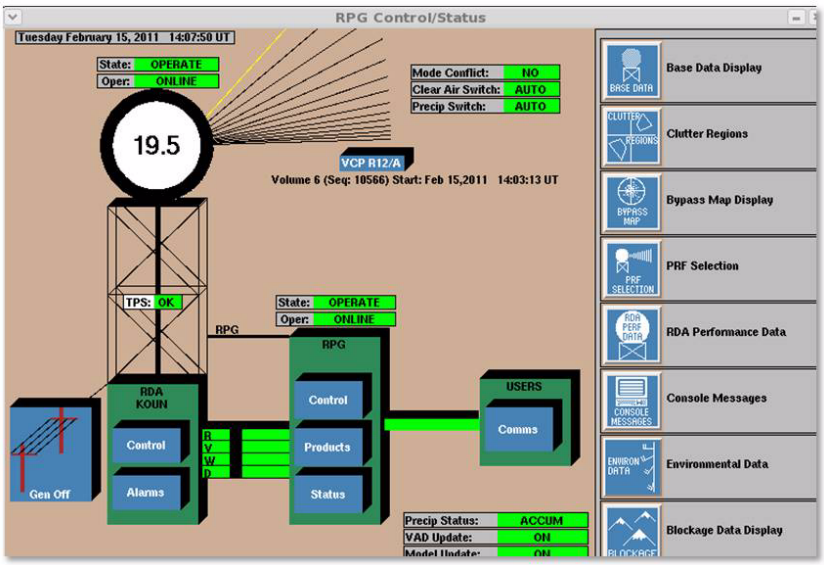

Figure 2-8. If RPG HCI appears to be delayed at 19.5°, it's likely due to the 8 hour check.

Attenuation

Reflectivity

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

Figure 2-9 shows base reflectivity, Z, generated from the two test radars, KCRI (single-pol, left) and KOUN (dual-pol, right). There is a squall line parallel to several radials and there is attenuation down radial in both of the Z products. Once the signal is attenuated, the loss cannot be recovered and propagates down radial.

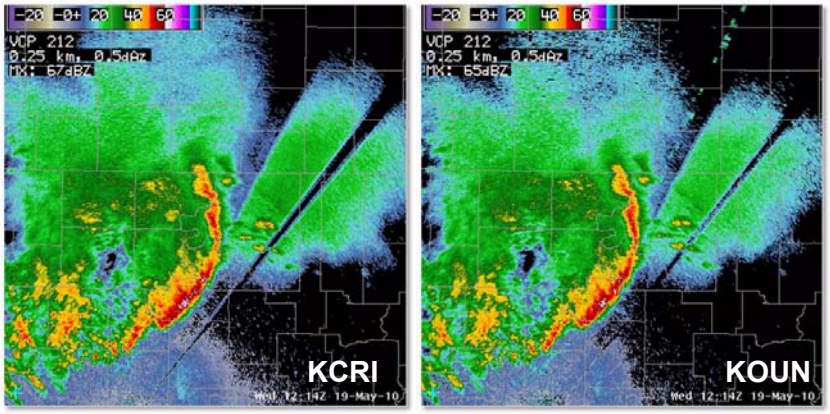

Figure 2-9. Images showing attenuation of Z with both KCRI (left) and KOUN (right) radars.

Differential Reflectivity

With this same squall line case, there are very low ZDR values down radial (Figure 2-9, right) that visually correlate with the Z attenuation. With ZDR, it is possible to have “differential attenuation”. ZDR is computed from the returned power of the H and V channels. Heavy rain with large drops results in more attenuation in the H direction compared to the V. This causes underestimation of ZDR down radial, just like the attenuation of Z on the left image. Once the signal is attenuated, the loss in

ZDR cannot be recovered and propagates down radial, as seen in Figure 2-10.

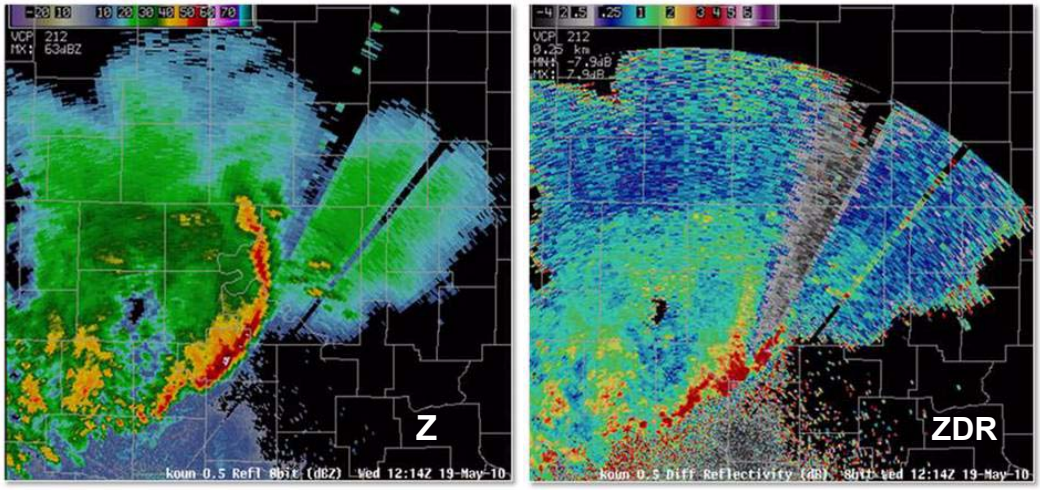

Figure 2-10. Images showing the attenuation of ZDR (right) relative to Z (left).

Depolarization

Depolarization is a phenomenon that has always occurred with radar, but will be apparent on the 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!

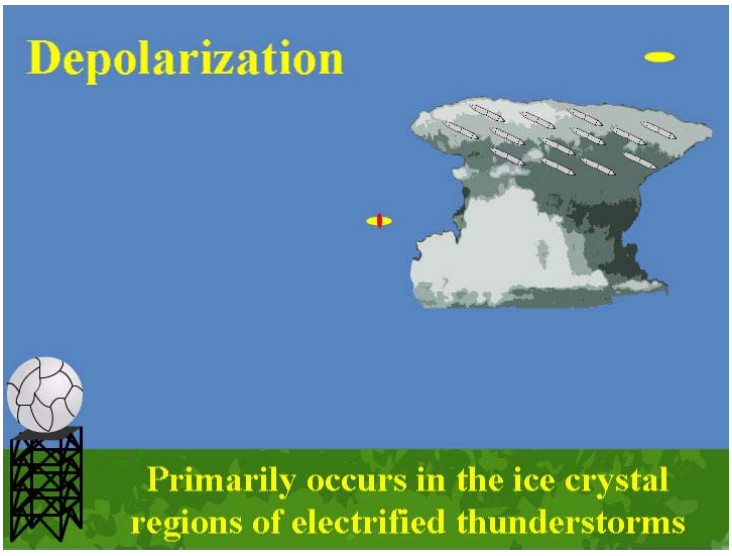

Figure 2-11. Image showing a horizontal pulse being reflected at least partially in the vertical direction after encountering ice crystals.

Figure 2-11 shows a horizontal pulse that has been reflected back at least partially in the vertical direction. The original horizontal 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. 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. Though it may rarely occur in hail, depolarization is far more likely to happen in the upper regions of thunderstorms when the electrification causes canting of the ice crystals. Since the electrification varies with time, so does the impact of depolarization.

Fortunately, regions that are down radial from thunderstorm tops are usually of low operational significance (see Fig. 2-12). Be aware that this is a known ZDR data artifact, and is not a cause for concern.

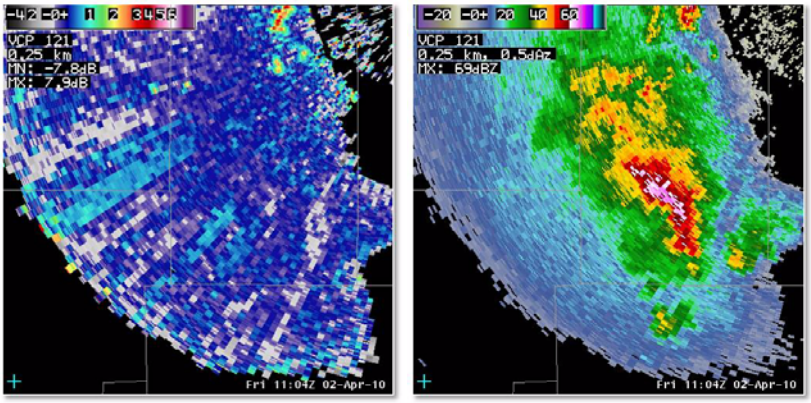

Figure 2-12. Depolarization is seen in ZDR (left) in the area down radial of thunderstorm tops as seen in Z (right).

Beam Filling Issues

Beam filling has always had implications with weather radar data quality, especially as range increases. It turns out that with dual-pol, new data artifacts can occur due to specific types of beam filling patterns.

In the graphics in Figure 2-13, the black circle represents the radar beam from the prospective of the RDA looking outbound. The image on the left represents partial beam filling, which is familiar, resulting in underestimated Z values. In the image on the right, the beam is filled, but by a mix of precipitation sizes and types. The mix may be varying raindrop or hail sizes or varying precipitation types such as rain/snow or rain/hail. The nature of this mix is relevant for dual-pol.

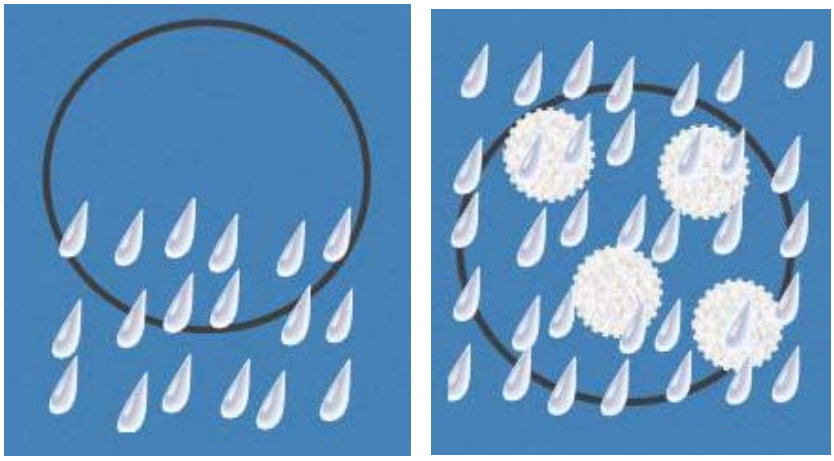

Figure 2-13. Partial beam filling (left) and mixed hydrometeors (right can cause data quality issues for dual-pol.

Non-Uniform Beam Filling (NBF)

It turns out that dual-pol products are negatively impacted by what is called Non-Uniform Beam Filling (NBF). With a uniform mixture, there is a relatively even distribution of the drop sizes and/or types.

In Figure 2-14, there is a supercell close to the radar and the associated CC product is on the right. Note that the CC values are lower within the core areas of the storm, which is expected when the radar samples a uniformly distributed mixture of rain and hail across the radar beam cross section.

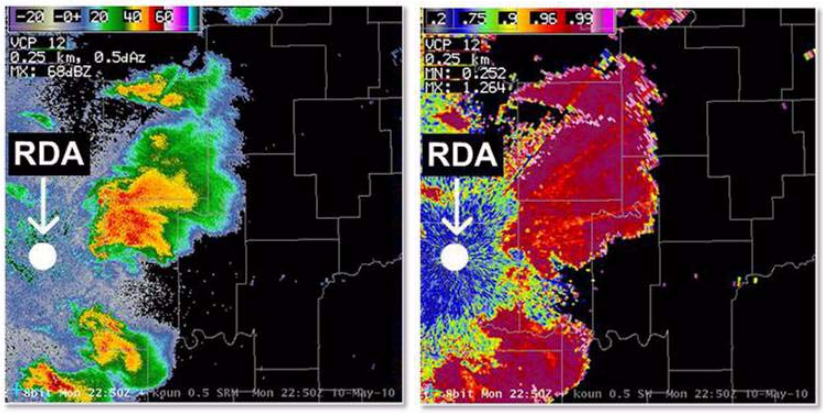

Figure 2-14. Z image (left) of a supercell at close range and CC image (right) of corresponding low CC values within the core of the storm.

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, in the graphic in Figure 2-15, 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. Recall that Φ_{DP} contributes to both CC and KDP.

Figure 2-15. A non-uniform mixture can produce a gradient of precipitation types within the beam (black circle).

Figure 2-16 represents the variation of Φ_{DP} from the top to the bottom of the beam, but we do not have the vertical resolution to measure it. The gradient of precipitation types and the associated gradient of Φ_{DP} is the bottom line for low CC values locally and down radial.

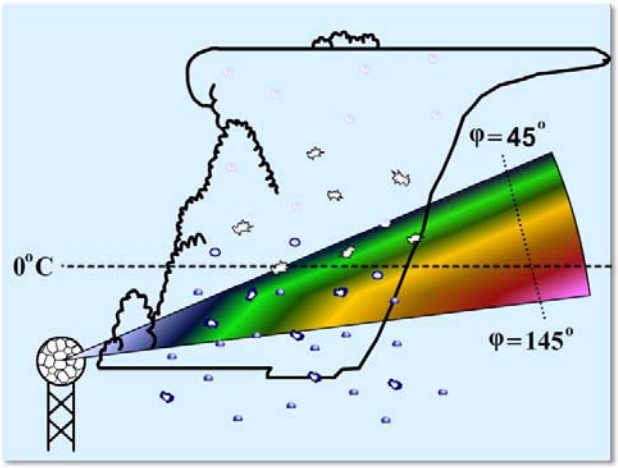

Figure 2-16. Colors represent a variation in Φ_{DP} over the height of the beam.

In Figure 2-17, the supercell has moved away from the radar. There are radial swaths of low CC that originate from the storm core. This is an example of non-uniform beam filling and its impact on the CC product. This has impacts on other dual-pol products, with examples coming up.

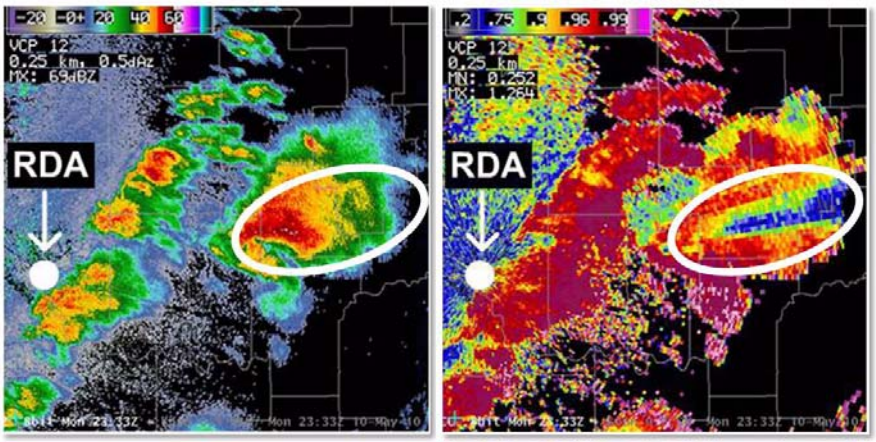

Figure 2-17. Z (left) and CC (right) images from the same supercell in Figure 2-14 after it has moved further away from the radar. Notice the wedge of low CC values east of the storm.

Once the storm is at a longer range, the non-uniform beam filling results in low CC values over a large wedge. We know from the associated Z product that these low CC values do not make sense.

First recall that Φ_{DP} values propagate down radial. When the hydrometeors are uniformly distributed, all is well. Φ_{DP} increases down radial as the beam passes through areas of pure rain. When sampling a convective storm at longer range or a squall line along a radial, there is an increasing chance of capturing a gradient of precipitation types within the beam. At the top can be hail and/or graupel, while the bottom of the beam is sampling liquid drops (see Fig. 2-18).

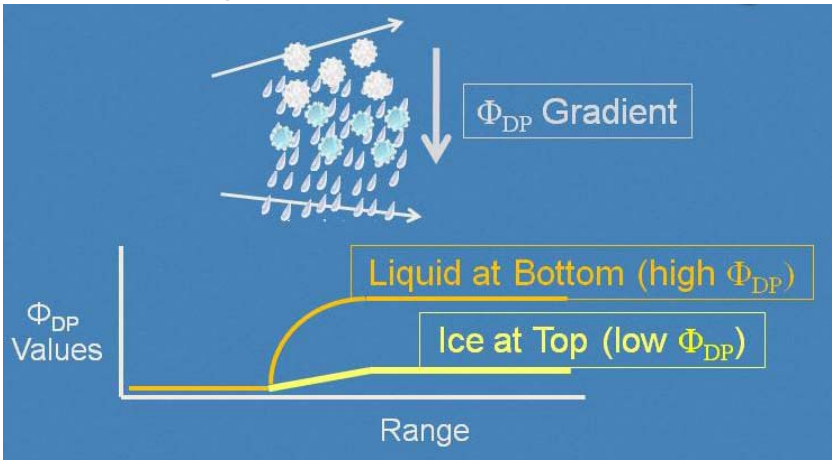

Figure 2-18. Higher Φ_{DP} values continue down radial when the beam is sampling non-uniformly distributed hydrometeors.

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.

The artifact of a swath of low CC values due to non-uniform beam filling 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. Figure 1-19 shows how low CC values affect the Hydroclass value that gets assigned, which then affects whether or nor rainfall is accumulated.

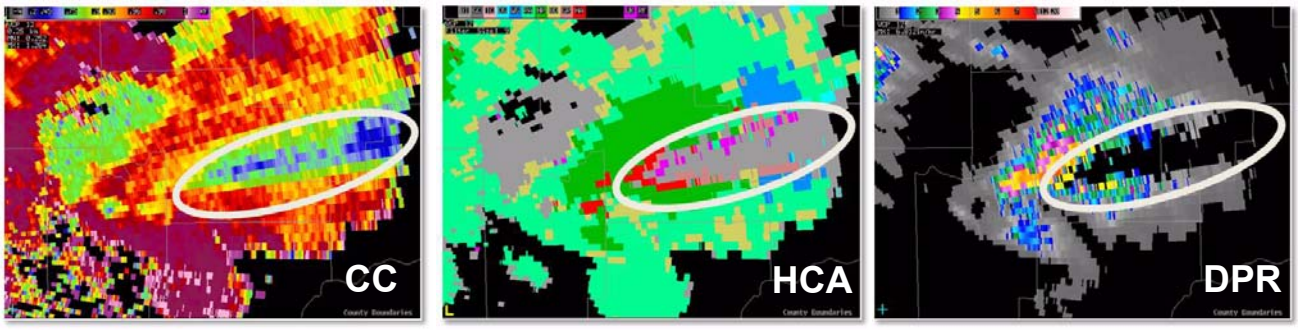

Figure 2-19. Low CC values (left) due to non-uniform beam filling can also impact the derived products like the Hydroclassification product (middle) and the Digital Precipitation Rate (right).

Wet Radomes

Wet radomes have always caused periodic data problems with weather radars and WSR-88D dual-pol data are no exception. The WSR-88D radomes are designed to be “hydrophobic” so the surface is designed to repel water, much like a well-waxed

car exterior. This promotes “rivulets” that drain off the dome. This vertically oriented water has been observed to cause ZDR changes for several volume scans. Since the rivulets are vertically oriented, there is more signal attenuation in the vertical, usually increasing the ZDR values.

ZDR and the Trees

Another data artifact that has been recently discovered is the behavior of ZDR where there is partial beam blockage due to deciduous trees. Figure 2-20 shows a Z product on the left with weak showers to the west of the radar. The associated ZDR product on the right has a swath of higher ZDR values along radials subject to partial beam blockage by deciduous trees. Since the vertical portion of the trees is more prevalent, there is some attenuation in the vertical compared to the horizontal. This has the effect of increasing the ZDR value. As of this writing, it is expected that the ZDR values in this wedge will be lower once the trees leaf out. It is not known if relatively higher ZDR values will persist in this wedge.

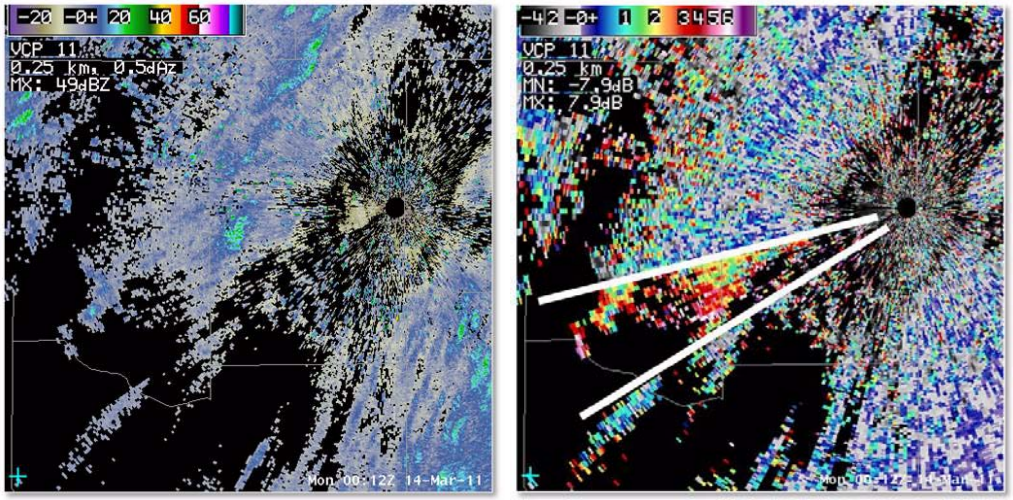

Figure 2-20. Z (left) and ZDR (right) images showing the effects on ZDR of partial beam blockage due to trees.

RPG Lesson 1: Life Without CMD and the Dual-Pol Preprocessor

Wideband on RPG HCI

The addition of the dual-pol data stream changes the appearance and function of the wideband link on the RPG HCI main page. There is a new channel, D, for dual-pol. The wideband link shows different behavior depending on which elevation angles are being sampled.

For the split cut elevations, the dual-pol data are generated using the Contiguous Surveillance (CS) rotation. Though there are fewer pulses per radial in CS, dual-pol base data are generated from this rotation because there is very little chance of range folding. For the CS rotation, the wideband link shows the R and D channels green and animated, while the V and W channels are white. For the Contiguous Doppler (CD) rotation, the R, V, and W channels are green and animated, while the D channel is white. For the Batch elevations and higher, all 4 channels are green and animated. Figure 1-1 shows the HCI main page display for CS, CD and Batch elevations and higher.

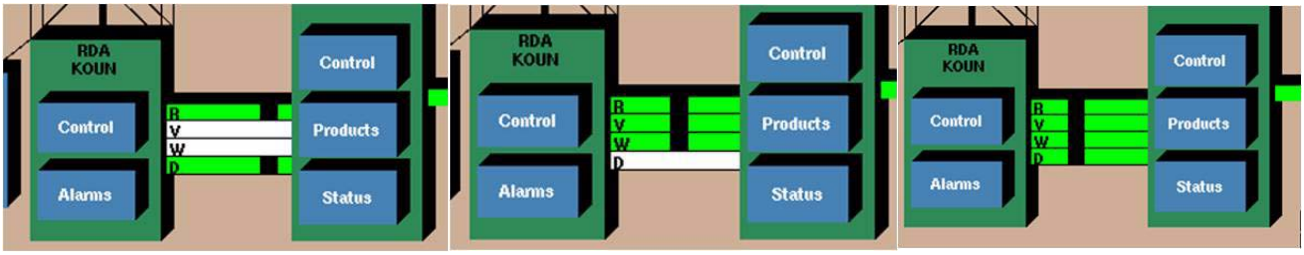

Contiguous Surveillance (CS) Contiguous Doppler (CD) Batch Elevations and Higher

Figure 1-1. HCI Main Page wideband link display for CS (left), CD (middle) and Batch elevations and higher (right).

Clutter Management Without Clutter Mitigation and Decision (CMD)

So why does the Clutter Mitigation and Decision (CMD) algorithm go away with dual-pol? The government had to give the dual-pol contractor a version of RDA software to serve as a “baseline”. That RDA software was Build 10.0, which did not yet have CMD.

A re-implementation of CMD is expected with the first post-dual-pol upgrade, RDA/RPG Build 13.0. However, getting CMD back is not “plug and play.” Its going to take quite a bit of software engineering and the work is ongoing. Until then, management of clutter filtering reverts back to the “pre-CMD” days.

Clutter 101

If a target is moving, it is not going to be identified as clutter. In order for a range bin to be identified as containing clutter, it must have near zero velocity and spectrum width. There are many targets out there that are not weather, but they move, such as wind turbine blades, traffic on roads, birds, bats, and insects (see Fig. 1-2).

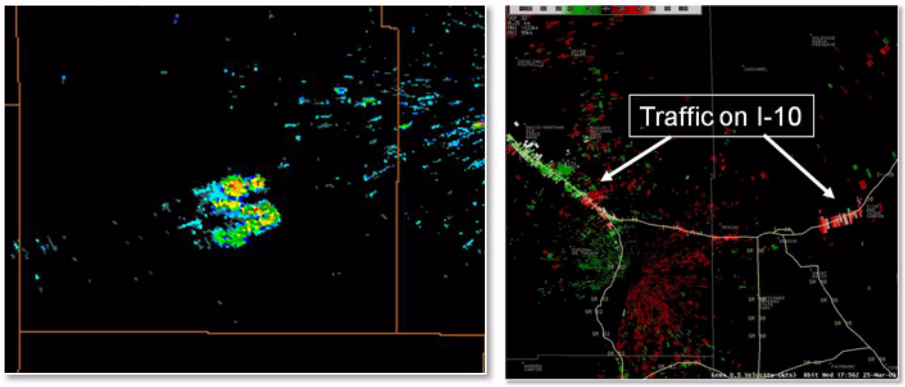

Wind Farm Clutter Traffic Clutter

Figure 1-2. Reflectivity and velocity images of wind farm clutter (left) and traffic clutter (right), respectively.

Even with clutter suppression applied everywhere, these type targets will ***not*** be removed by the clutter filter.

CMD's Role in Clutter Filtering

CMD identifies the range bins that contain clutter, whether it be normal clutter, such as a mountain range, or AP clutter. CMD does ***not*** perform the actual power removal.

The Gaussian Model Adaptive Processing, or GMAP, removes the power, but only for the bins identified by CMD (see Fig. 1-3, right). The loss of CMD has no impact on the removal of the clutter signal, just where it gets applied (see Fig. 1-3, left).

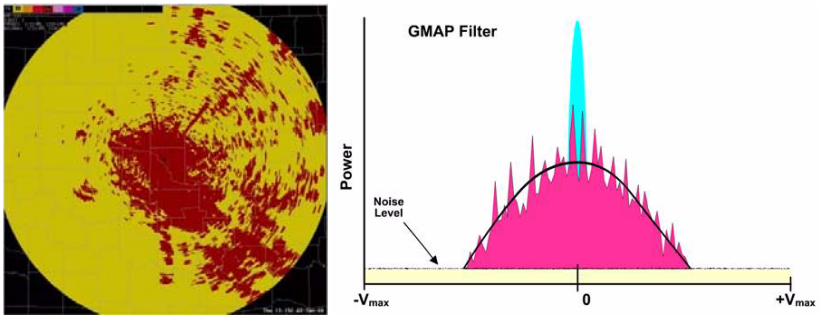

Figure 1-3. CMD (left) identifies clutter, GMAP (right) filters it.

Gaussian Model Adaptive Processing (GMAP)

The purpose of the Bypass Map vs. All Bins definitions is to identify where clutter filtering will be applied. GMAP performs the actual filtering by identifying a narrow width near zero velocity in an attempt to isolate the portion of the returned signal associated with clutter (see Fig. 1-4). Of course ground targets that move, such as cars and wind turbine blades, are not going to fall into this width and will not be filtered.

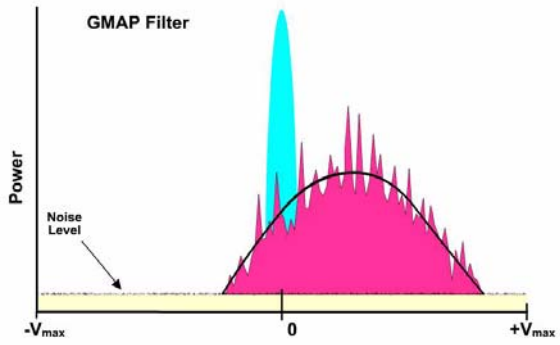

Figure 1-4. GMAP identifies a narrow width near zero velocity in an attempt to isolate ground clutter.

After signal removal, GMAP has the capability to rebuild the weather signal if sufficient pulses remain. On the graphic on the left hand side of Figure 1-5, there is no precipitation present. GMAP is applying filtering to an area of higher elevation to the southwest of the radar, which shows up in the No Data pixels. Some time later, a squall line passes through, with light to moderate rain behind it. As the precipitation passes over the higher elevation, GMAP rebuilds the weather signal, assigning valid base data instead of No Data.

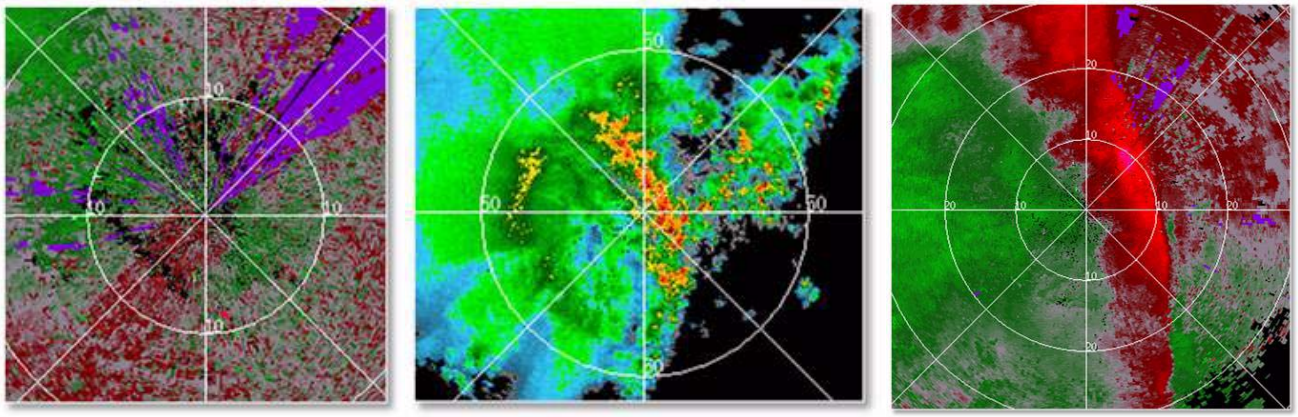

Higher Elevation SW of Radar: No Precipitation

Squall Line Passes Through

GMAP Rebuilds: Precipitation instead of No Data

Figure 1-5. Before a squall line passes, GMAP is filtering for higher elevation southwest of the radar (left). Once the squall line moves through the area (middle), GMAP rebuilds to show valid base data where there was No Data before.

CMD Status On the RPG HCI main page, there are two possible status categories for CMD: PENDING or DISABLED.

PENDING shows up when the software is initially installed or after VCP 211, 212, or 221 has been downloaded. When you download any of these VCPs, the RPG still tries to turn on CMD. Clicking the PENDING button tells the RPG to stop waiting for CMD. The RPG will then show DISABLED as the status. (See Fig. 1-6 on page 3-45.)

Keeping the CMD status as DISABLED may be helpful for remembering that manual clutter identification is necessary.

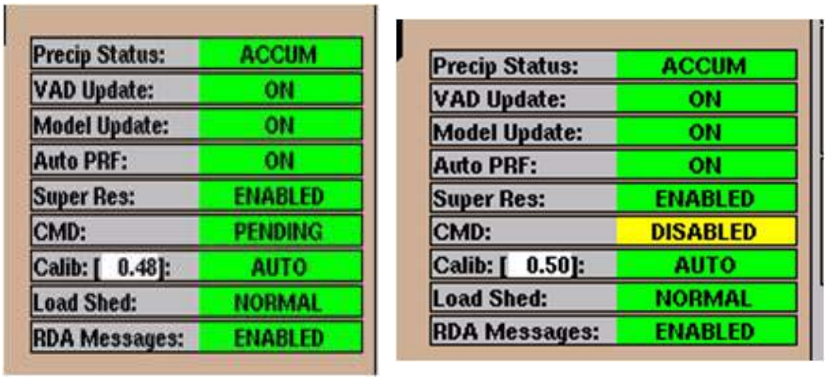

Figure 1-6. When the software is initially installed, or after certain VCPs are downloaded, the CMD status will be "PENDING" (left). Click on the word "PENDING" to disable CMD (right).

Clutter Files and Elevation Segments

Without CMD, identifying the location of clutter targets must be done manually. A clutter regions file, shown in Figure 1-7, tells GMAP where to apply filtering, for all azimuths and all elevations. Used appropriately, clutter files apply Bypass Map filtering where there is normal clutter, and All Bins where there is AP Clutter.

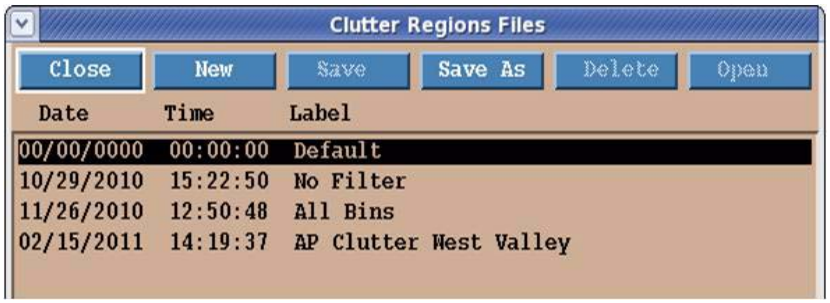

Figure 1-7. Clutter regions files display.

Each clutter file defines clutter suppression for all the elevation angles. This is done through five “clutter elevation segments” (see Fig. 1-8), each one having different clutter filtering. It isn’t necessary to apply suppression to the same bins at 6° that are applied at 0.5°. Elevation segments allow for better vertical resolution of the application of clutter filtering.

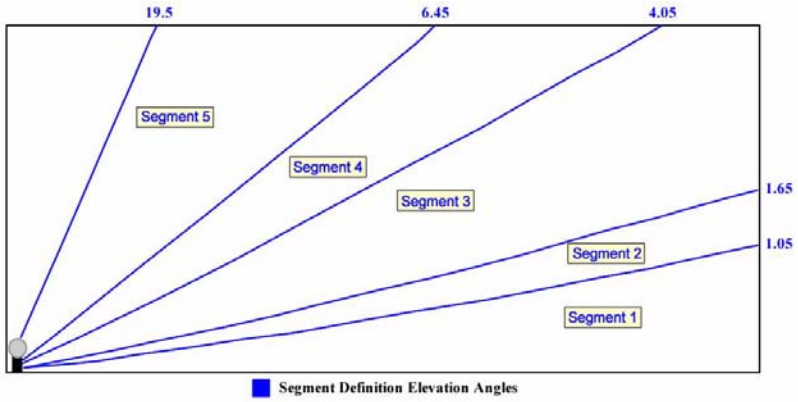

Figure 1-8. Clutter suppression is applied through five “clutter elevation segments”.

Bypass Maps are generated off-line. During this process, there is one Bypass Map built for each elevation segment. The Bypass Map for Segment 3 will likely identify a lot less clutter than the Bypass Map for Segment 1. The default angles used to generate the maps are shown in Figure 1-9 (red lines).

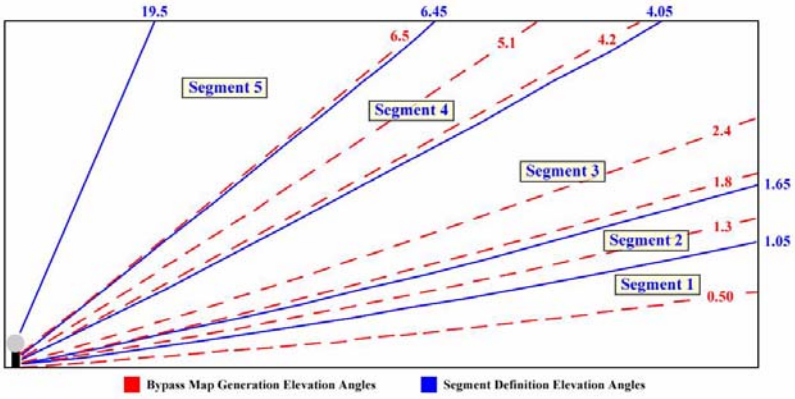

Figure 1-9. Default angles use to generate Bypass Maps (red lines).

The Clutter Regions Editor

Clutter regions files are accessed by first clicking on the Clutter Regions application button on the RPG HCI main page. This brings up the Clutter Regions window (see Fig. 1-10).

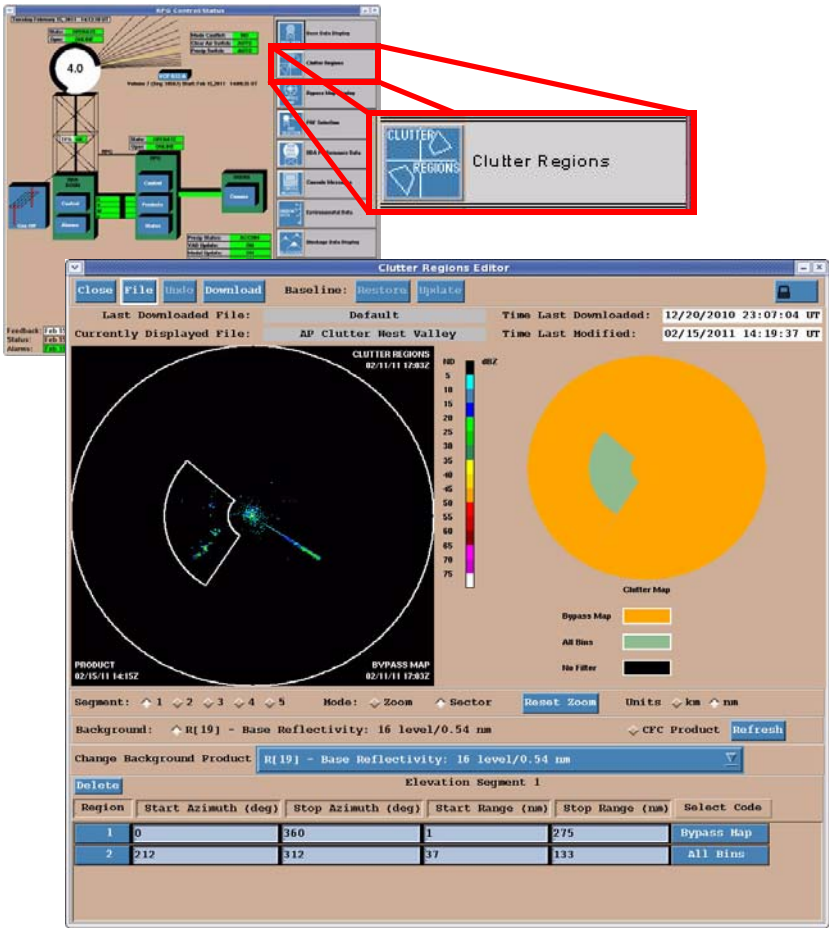

Figure 1-10. Clicking on the “Clutter Regions” button on the HCI Main Page (top image, red box) will open the Clutter Regions window (bottom image)

Each clutter file defines clutter suppression for all five elevation segments. Each segment can have one or more regions, which are volumes defined from azimuth to azimuth and range to range, encompassing the elevation angles that fall within

that elevation segment (see Fig. 1-11). Each clutter file can have up to 25 regions defined.

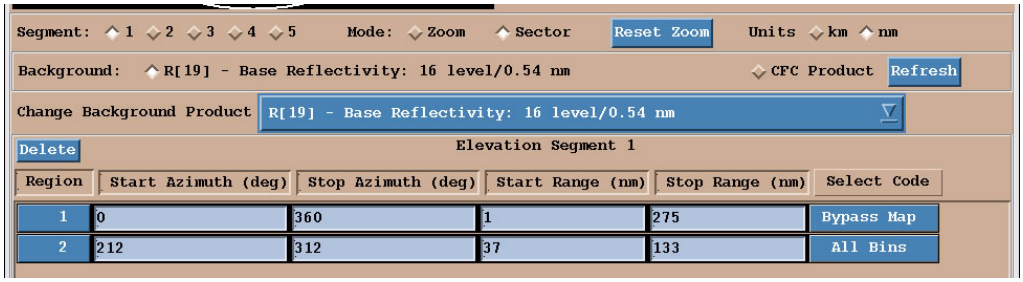

Figure 1-11. Clutter regions are defined in the bottom section of the Clutter Regions file display.

For the example Clutter Regions file in Figure 1-12, assume that segment 1 has the Bypass Map in control everywhere (Region 1), with one area of All Bins to the northeast (Region 2), and another to the southeast (Region 3). The clutter file in Figure 1-12 has 3 regions in segment 1, 2 regions in segment 2, and 1 region in segments 3 through 5. There are 8 regions total defined for this file.

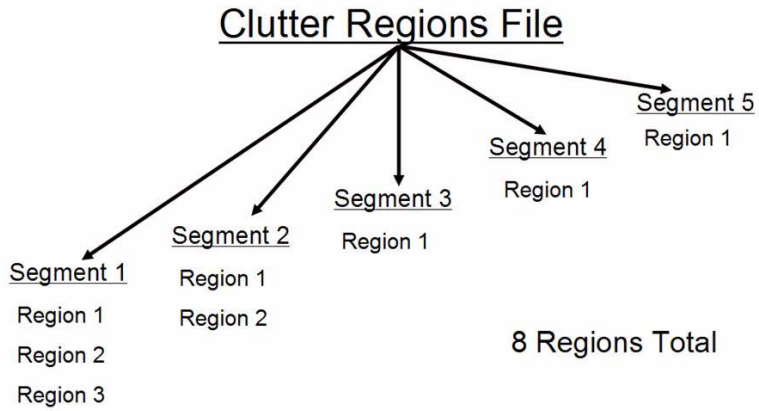

Figure 1-12. Sample clutter regions file with eight total regions. A clutter regions file can have up to twenty-five regions.

The two lines of text at the top of the Clutter Regions Editor window, shown in Fig. 1-13 on page 3-49, help you keep track of which file is displayed on the window, and which file was last downloaded to the RDA. The first line indicates the name of the clutter file that was last downloaded to the RDA and at what time it was downloaded. The

next line tells you which clutter file is currently displayed on the window.

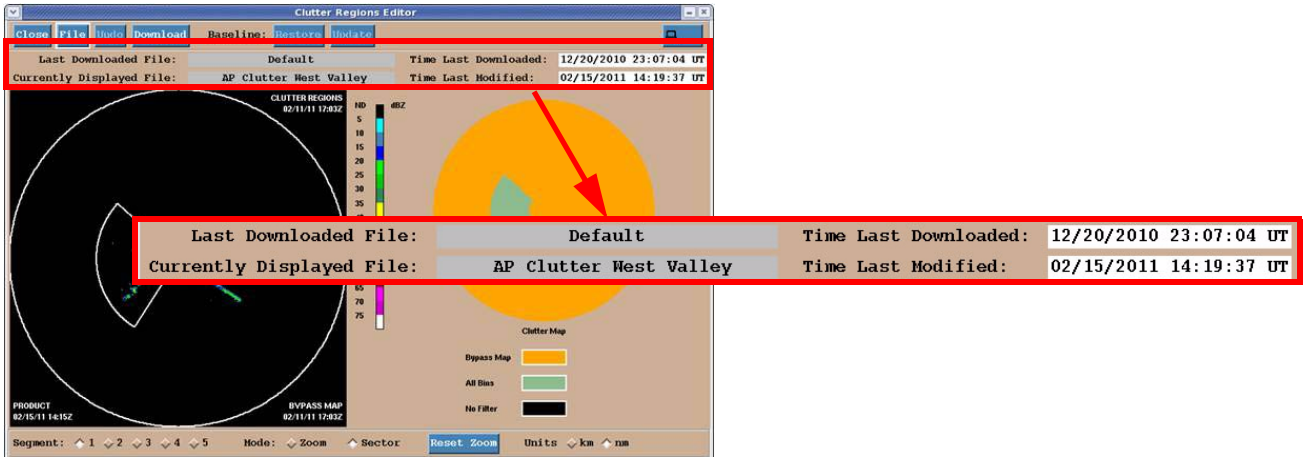

Figure 1-13. The top of the display indicates which file is displayed and the last file downloaded to the RDA.

Within each clutter file displayed, the configuration for each of the five segments is available. Figure 1-14 shows Segment 1, the lowest segment. Any region that is defined needs a Select Code, either Bypass Map or All Bins.

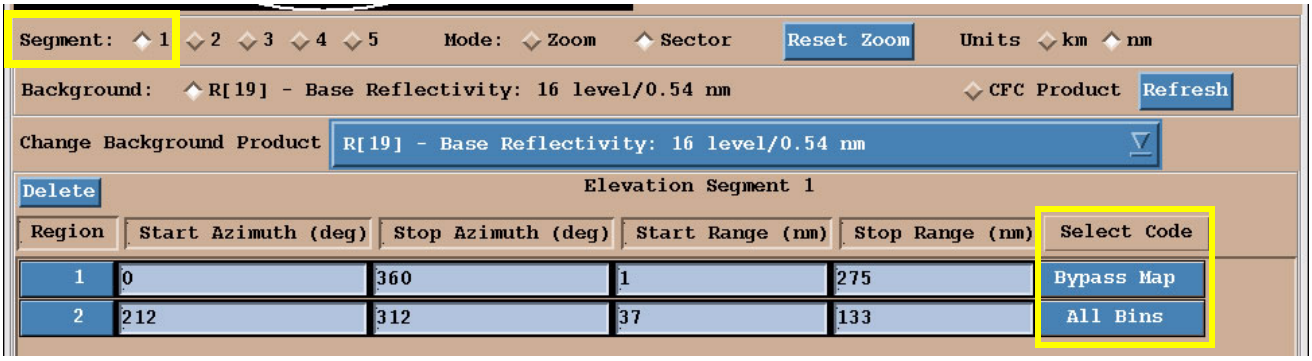

Figure 1-14. Each region needs a Select Code, either Bypass Map or All Bins (yellow box on right). A region is defined by start and stop azimuths and start and stop ranges.

In line 1, which should be the same for every clutter file, there is a region from 0 to 360 degrees and from 1 to 275 nm, which is everywhere, with the Bypass Map in control.

After line 1, which has the Bypass Map in control everywhere, additional regions can be defined with All Bins if needed. If there are areas with

predictable AP Clutter, files can be defined in advance and downloaded as needed. For example, diurnal AP frequently occurs in valleys and over bodies of water, so the associated AP clutter is in a predictable location. Otherwise, regions that apply All Bins must be defined as needed before downloading to the RDA.

Bypass Maps are generated offline, stored, and do not update until new maps are generated. Bypass Maps should be generated under propagation conditions that are “normal” for a given site. This means no AP or precipitation going on. It is recommended that Bypass Maps be checked periodically for relevance, and updated at least seasonally.

The list of available clutter regions files is accessed from the File button at the top of the Clutter Regions window (see Fig. 1-15). The strategy for identifying clutter without CMD involves downloading existing files and creating new ones as needed. The Default file is a pre-defined file that cannot be deleted. It is part of the baseline and designed to apply Bypass Map filtering for all azimuths and elevations. It is intended to be used for all non-AP situations.

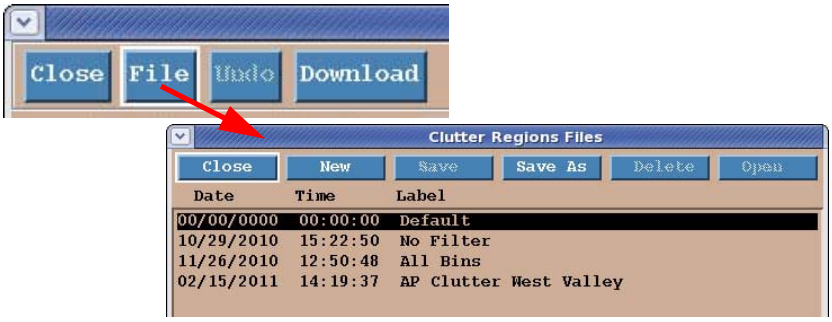

Figure 1-15. Click the File button at the top of the Clutter Regions Editor window to access the list of available clutter regions files.

Pre-defined clutter regions files that address predictable locations for AP clutter, such as the "AP Clutter West Valley" file, are recommended.

To download a clutter file to the RDA, it must first be displayed at the Clutter Regions window. Click on the name of the file, select Open, and the file will be displayed. Once it is displayed on the Clutter Regions Editor window, it can then be downloaded to the RDA by clicking the download button at the top of the Clutter Regions Editor window (see Fig. 1-16).

Figure 1-16. Once the correct Clutter Regions file has been selected, it can be downloaded to the RDA by clicking the Download button (yellow box).

In this example, the file "AP Clutter West Valley" has been displayed and is ready for download.

Clutter Management Conclusion

Bypass Maps will need to be generated often enough to keep them representative. Having predefined clutter files to address predictable areas of AP clutter is recommended. In general, clutter management without CMD requires pro-active downloading of clutter files. The Default file is recommended when there is no AP present, and new files can be created and downloaded as needed to address AP clutter.

Product Resolution

RPG Build 12.2 brings a change in the range resolution to some products. Super Resolution means an azimuthal resolution of 0.5° and it is still limited to Z, V and SW on the Split Cut elevations. The display range for the Split Cut V and SW products is out to 300 km (162 nm) or 70 kft, whichever comes first. The remainder of the Split Cut products, including all the new dual-pol products, have an azimuthal resolution of 1.0° and a range resolution of 0.25 km. Figure 1-17 shows the difference in resolution between the Z, V and SW Split Cut elevations (top row) and differential reflectivity and correlation coefficient (bottom row).

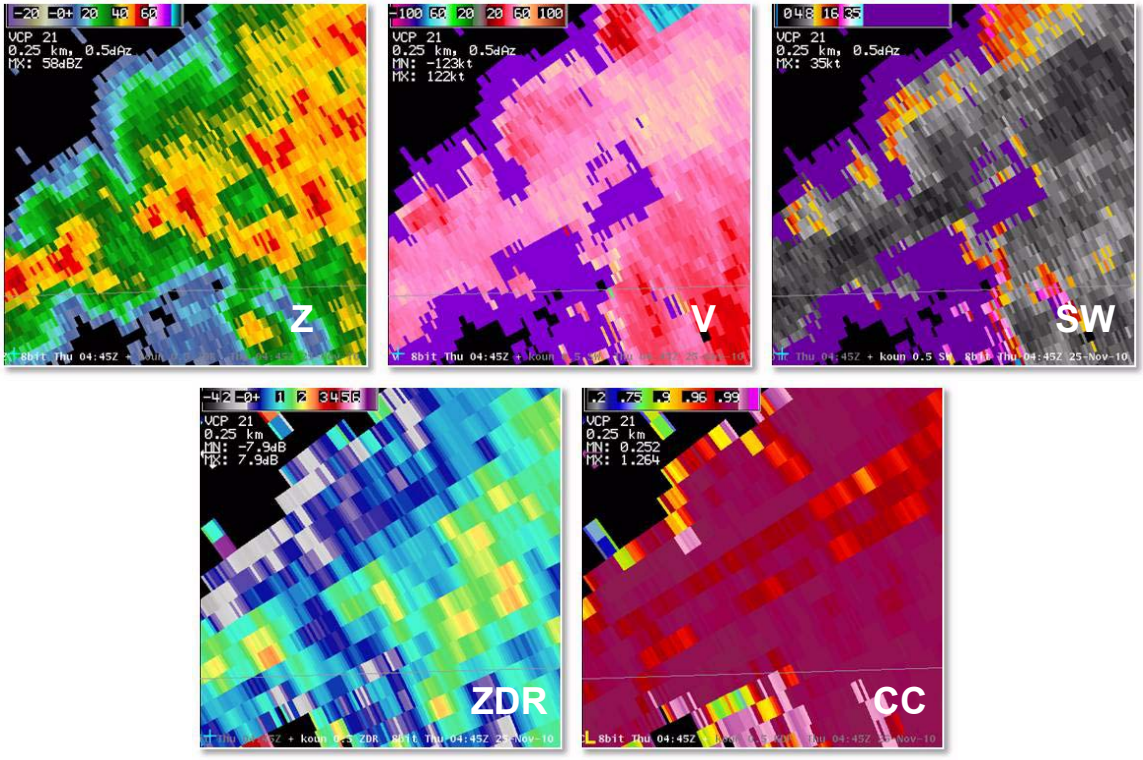

Figure 1-17. For the Split Cut elevations, Z (upper left), V (middle), and SW (upper right) have 0.5 azimuthal resolution, while ZDR (lower left) and CC (lower right) have 1.0° azimuthal resolution.

Split Cuts

The Split Cut elevation dual-pol base data that arrives from the RDA has an azimuthal resolution of 0.5°. The left image in Figure 1-18 is the “raw” CC data from the RDA. It is too noisy for both human interpretation and algorithm ingest. At the RPG, the dual-pol base data are first recombined to 1° azimuth. The recombination is performed on each channel separately: recombine H power, then recombine V power, then the cross correlation is recombined. The ZDR, CC, and Φ_{DP} base data are all recombined to 1° azimuth. The smoothing is done by the RPG Dual-Pol Preprocessor.

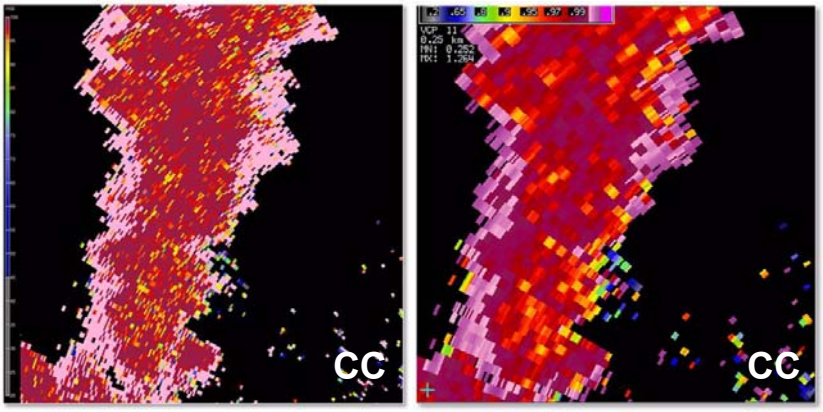

Figure 1-18. Dual-pol base data arrives from the RDA with 0.5^{o} azimuthal resolution (left) and is recombined to 1^{o} azimuthal resolution (right).

Batch Elevations and Above

For the base and derived products of the Batch elevations and above, the azimuthal resolution is 1.0° (see Fig. 1-19). The range resolution is 0.25 km, which is an increase in resolution for many products for these elevations. Also, the V and SW products have a display range out to 300 km (162 nm) or 70 kft, whichever comes first.

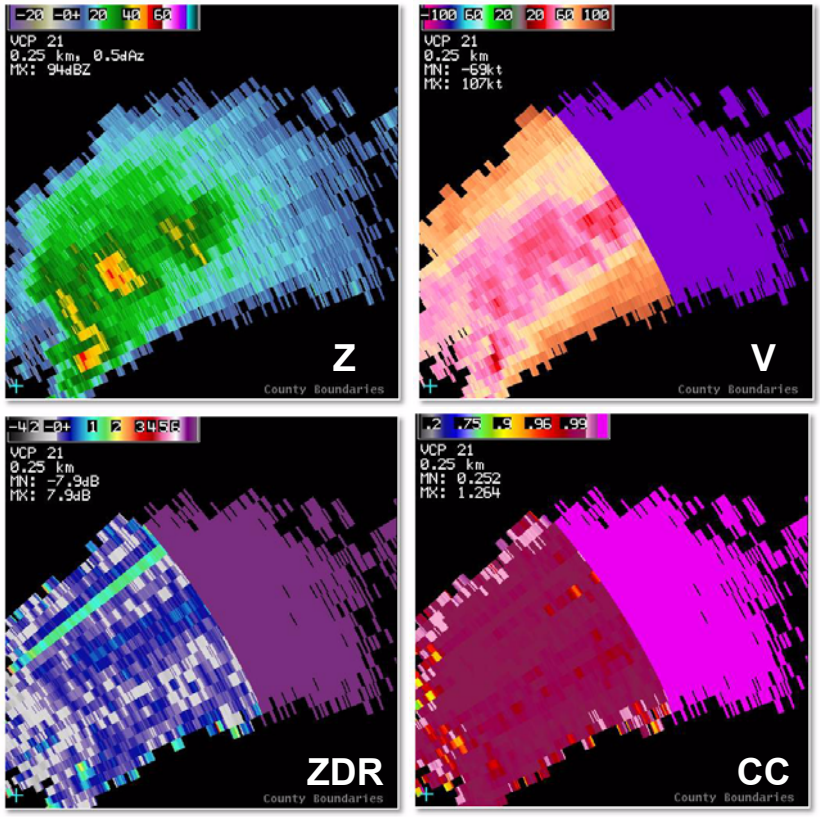

Figure 1-19. Base and derived products of the Batch elevations and above have an azimuthal resolution of 1.0° and a range resolution of 0.25km.

The RPG Dual-Pol Preprocessor

The RPG Dual-Pol Preprocessor prepares the dual-pol base data for base product generation and for input into the 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 are noisy in appearance. For any given range bin, 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.

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

Smoothing is also applied to the Z base data, but only for input to the dual-pol RPG algorithms. There is no change to the Z values used to generate all the base Z products.

Preprocessor and ZDR/CC

In Figure 1-20, the image on the left is raw ZDR from the RDA, which is not yet recombined or smoothed. The image on the right is the same data displayed in AWIPS after recombination and Preprocessor smoothing.

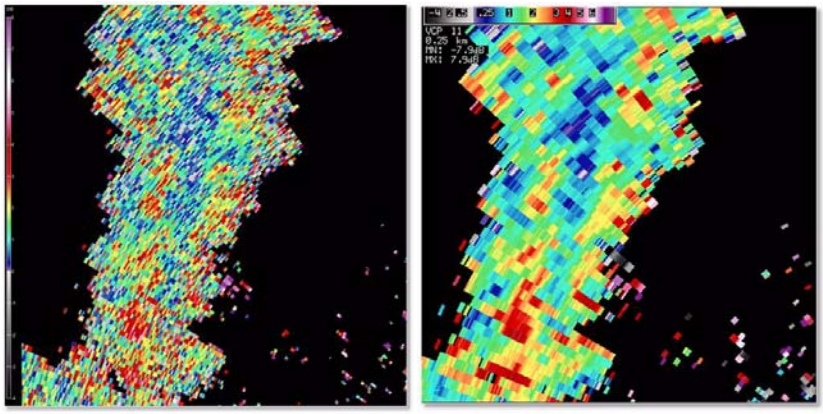

Figure 1-20. Raw ZDR (left) from the RDA is recombined and smoothed before it is displayed in AWIPS (right).

Figure 1-21 shows before and after preprocessing of correlation coefficient.

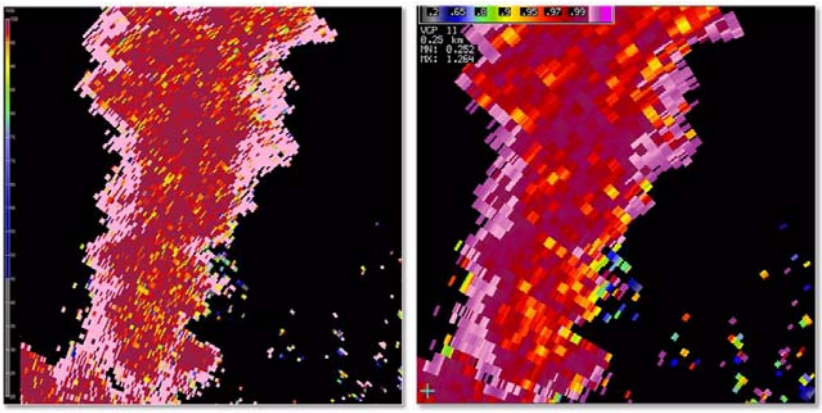

Figure 1-21. Raw CC (left) is recombined and smoothed (right) by the RPG Dual-Pol Preprocessor.

Preprocessor and Φ_{DP}

Since Φ_{DP} is at the heart of dual-pol base data, it gets special treatment from the Preprocessor!

On the right in Figure 1-22 is an example of Φ_{DP} values, which are not generated into a product at AWIPS. This image is from GR Analyst, showing the raw Level II data. Once the Φ_{DP} data have been smoothed, the Preprocessor calculates the KDP values. The KDP values are then available for generation of the KDP product (see Fig. 1-22, left) and for input to the dual-pol algorithms.

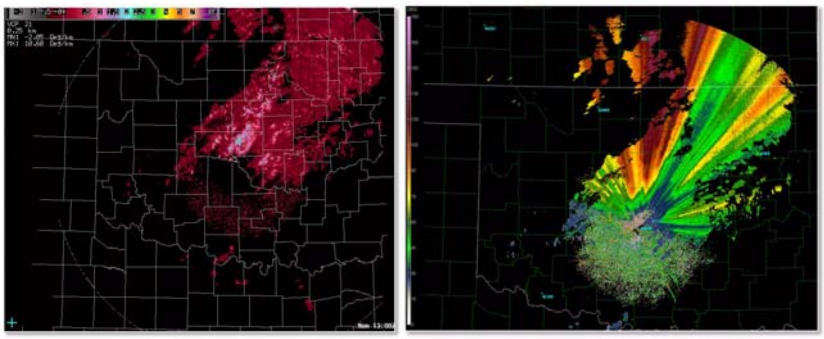

Figure 1-22. Φ_{DP} values as seen in GR Analyst (right) and KDP values calculated by the Preprocessor (left).

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.

As you may recall from the RDA training, Φ_{DP} is a phase value which accumulates down radial (see Fig. 1-23). Φ_{DP} has the greatest increase with range where the beam encounters large amounts of liquid water.

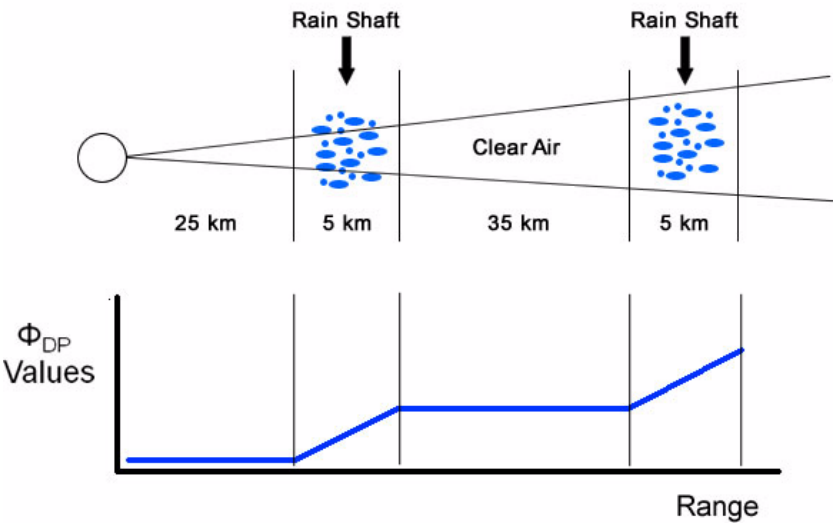

Figure 1-23. This graphic shows how Φ_{DP} values accumulate down radial.

You will likely see a data artifact that is related to how Φ_{DP} is processed along each radial, specifically in between segments of “weather.”

“Weather” is identified by CC>0.85. In the example in Figure 1-23, there are two segments along the radial that are tagged as rain, because of their associated high CC values. Figure 1-23 shows how Φ_{DP} typically behaves in “clear air”, but there are exceptions.

Sometimes CC>0.85 is assigned to non-meteorological returns, such as residual clutter near the radar. This can then artificially increase the Φ_{DP} values down radial (see Fig. 1-24). There is an artifact that you will see on the ZDR and KDP products that is related to the behavior of Φ_{DP} down radial and another job of the Preprocessor: Attenuation Correction.

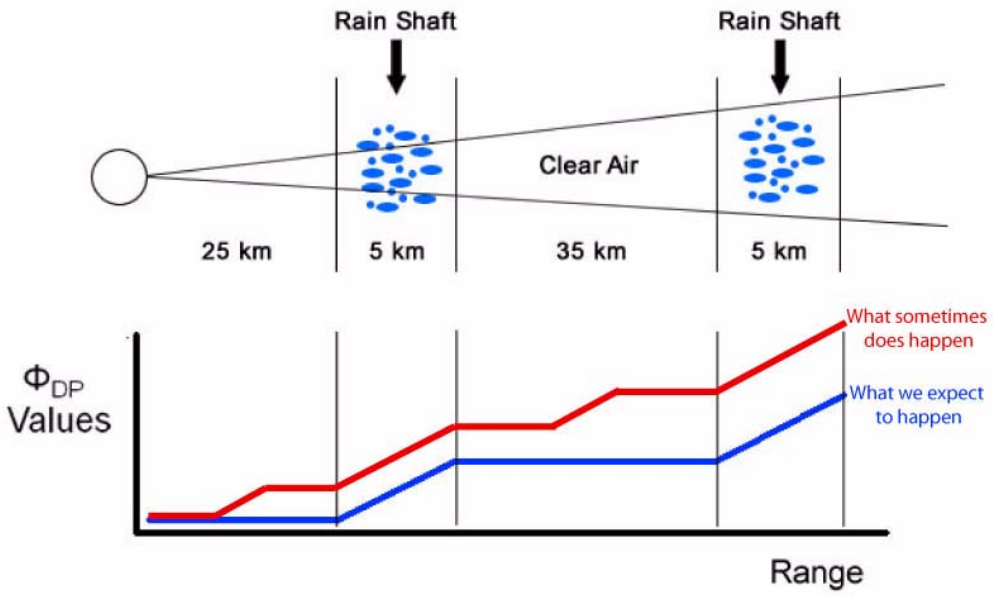

Figure 1-24. Non-meteorological clutter can artificially increase Φ_{DP} values down radial (red line).

Attenuation Correction

The final task of the Preprocessor is Attenuation Correction for Z and ZDR. As before, the Z value that has this correction applied is only used for the dual-pol algorithms.

When the radar beam encounters high water content, the returned power can be attenuated, and Φ_{DP} increases. The relationship between Φ_{DP} and the amount of attenuation is strong enough that the increase in Φ_{DP} can be used as an attenuation correction. Empirical relationships have been developed between Φ_{DP} and the

attenuation of Z and between Φ_{DP} and the attenuation of ZDR. These relationships are currently based on research in Oklahoma and will likely be changed as knowledge from other regions is gained.

ZDR and KDP Spikes

If you see spikes on ZDR and KDP, they are known artifacts of the Preprocessor. There is research underway to adjust the Preprocessor parameters to minimize this artifact in the future.

The result is radially oriented spikes that are transient in time and space as shown in Figure 1-25. With ZDR, the spikes originate where the beam encounters precipitation and persist with range. This is due to Φ_{DP} values that are too high with respect to the ZDR values within the precipitation. There is too much attenuation correction applied to ZDR, which then falsely increases the ZDR values along these radials.

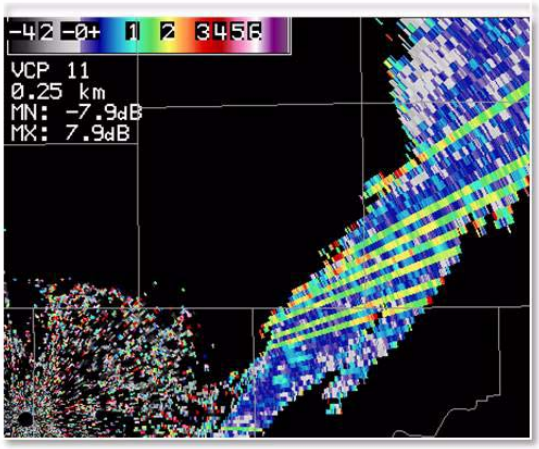

Figure 1-25. ZDR shows spikes where Φ_{DP} values are too high with respect to ZDR values within precipitation.

With KDP, the spikes originate near the RDA, because of noisy Φ_{DP} values (see Fig. 1-26). 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.

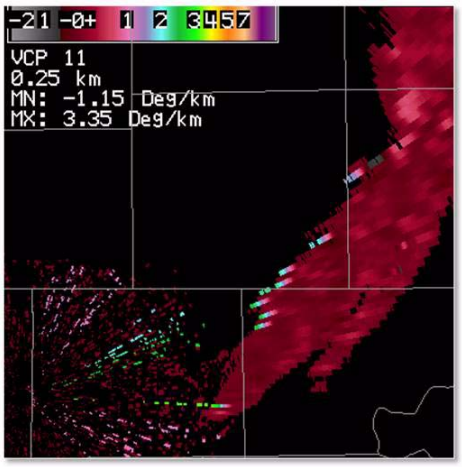

Figure 1-26. Spikes in KDP values can occur with noisy Φ_{DP} values, but usually smooth out when the beam intercepts precipitation.

The ZDR and KDP spikes are more likely to occur with residual clutter near the RDA. The need to manage clutter manually (without CMD) will make this more challenging in the short term. These spikes are most likely to occur on the lower elevations. If there is rain at close range, spikes are rare, since Φ_{DP} is well behaved in rain.

RPG Lesson 2: Hydrometeor Classification Algorithm (HCA) and Melting Layer Detection Algorithm (MLDA)

Until now, the only way to identify a melting layer via radar was to see a bright band in reflectivity, but it was never a guarantee that the bright band would be visible. With dual-pol radars, the melting layer stands out much more often in the products of CC and ZDR because they have very unique signatures in the melting layer. The melting layer manifests itself as a ring (or partial ring) of lower CC, typically around 0.9 to 0.95, and higher ZDR, somewhere between 0.5 and 2.5 dB. In the reflectivity image in Figure 2-1 there is not a readily visible bright band, but very prominent rings of lower CC and high ZDR.

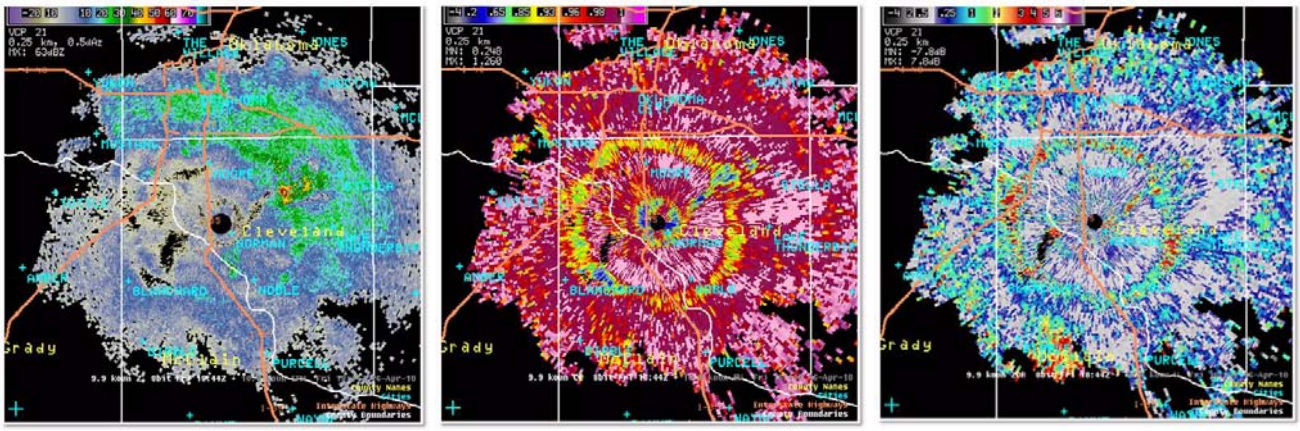

Figure 2-1. An example of a bright bands that isn't well-defined in Z (left), but is easily seen in CC (middle) and ZDR (right).

The Melting Layer Detection Algorithm (MLDA) uses these unique signatures of CC and ZDR within the melting layer to identify the heights of the top and bottom of the melting layer for each radial. These heights are updated every volume scan and are used to produce a Melting Layer (ML) graphic product. The ML product is an overlay and is available every volume scan for every elevation angle (see Fig. 2-2).

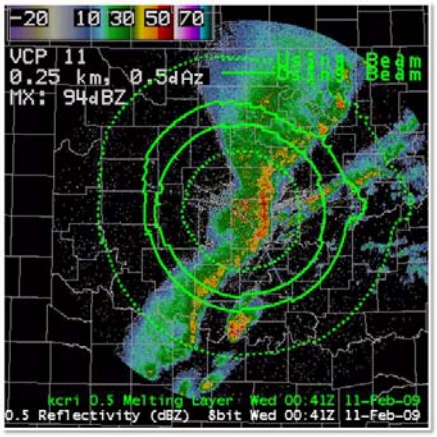

Figure 2-2. Image of the Melting Layer (ML) product.

Figure 2-3 shows the process of calculating the MLDA. The idea of the MLDA is to identify the rings of CC and ZDR in an automated way providing the height of the top and bottom of the melting layer that is updated every volume scan and can be used by other algorithms such as the Hydrometeor Classification Algorithm (HCA).

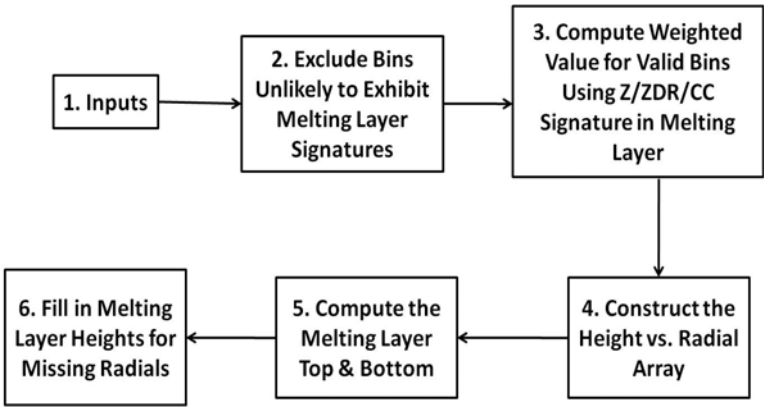

Figure 2-3. Flow chart providing a high-level overview of how the MLDA works.

1. Inputs for the MLDA

Table 4-1 lists the inputs for the MLDA. The algorithm ingests radial-based data from the elevations of 4° through 10°. Below the 4° elevation angle, the signature of the melting layer is smeared due to beam broadening, thus accurate detections are not possible. Above the 10° elevation angle, there is little data. In addition to the radial-based data, the MLDA uses MLDA data from previous volume scans. The number of volume scans used depends on the VCP mode (either precipitation or clear-air mode). Finally, a default 0 Celsius height is ingested which either comes from a user-defined height set at the RPG or the height from the latest RUC model run.

Elevations 4.0° through 10°	• Reflectivity (Z) • Differential Reflectivity (ZDR) • Correlation Coefficient (CC) • Signal-to-Noise Ratio (SNR) • Hydrometeor Classification (HC)
MLDA	• Previous 2 Vol Scans (Precip Mode) • Previous 5 Vol Scan (Clear-Air Mode)
Default 0 Celsius Height	• User-defined RPG Adaptation Data • Latest RUC Model

Table 4-1: Inputs used to calculate the ML product.

2. Exclude Bins from Radial for Computation of Melting Layer Top/Bottom

The first step in the MLDA is to eliminate any bins in a radial that will not be used for identifying the melting layer top and bottom heights for that radial. Bins identified as ground clutter or biologicals by the Hydrometeor Classification Algorithm (HCA) are excluded. Also, bins where the HCA is uncertain about a classification type, so either identified as ND or UK, are excluded. Finally, bins where the signal is too low (SNR < 5 dB) and any bins where the slant range height is above the maximum climatological height for a melting layer, which is set to 6 km (or 19 thousand feet), are excluded.

The graphic in Figure 2-4 shows an example radial where bins meeting the exclusion criteria are marked with a red X, and bins that will be used to identify the melting layer top and bottom are marked with a green check mark.

Characteristic	Value
HydroClass (HC)	GC, BI, ND and UK
Signal-to-Noise Ratio (SNR)	< 5 dB
Slant Range Height (SRH)	> 6 km (19kft)

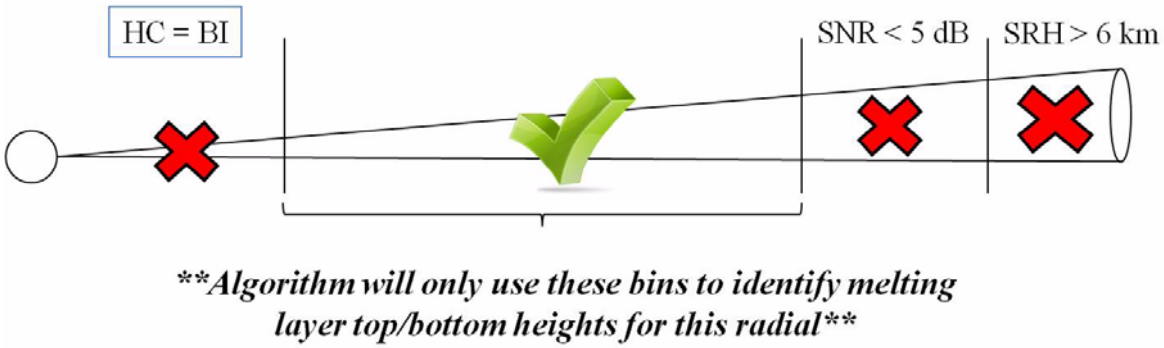

Figure 2-4. The algorithm filters bins according to their hydroclass, signal-to-noise ratio, and slant range height.

3. Compute Weighted Value for Each Remaining Bin

For each remaining bin from the previous step, the MLDA will compute a weighted value based on the likeliness of that bin being within the melting layer. The MLDA uses Z, ZDR, CC and the elevation angle to compute this weighted value, and the higher the weighted value, the more likely it is that the bin exhibits a melting layer signature.

As mentioned before, the motivation behind the MLDA is the observation that the ZDR increases and CC decreases within the melting layer producing rings of increased ZDR and decreased CC coincident with the melting layer. The weighted value will be computed using this logic. Since these two rings are not always in the exact same

location, the algorithm will look at the given bin and a few bins down radial from it and look for Z between 30 and 47 dBZ, ZDR between 0.8 and 2.2 dB and CC between 0.9 and 0.97. If these conditions pass, a non-zero value is assigned to the given bin that is weighted according to its elevation angle. The higher the elevation angle, the higher the weight. This is because higher elevation data are less susceptible to beam broadening effects, giving a more accurate depiction of where the melting layer is truly located.

4. Sort Weighted Values According to Their Height and Radial

Once all the valid bins have been assigned a weighted value, these values are sorted according to their height and radial into a height vs. radial array with a vertical resolution of 0.1 km and radial resolution of 1 degree. A portion of an example array is shown in Figure 2-5.

Radial (deg)

Height (km)	275°	276°	277°	278°
4.0 km	400	398	408	410
3.9 km	387	375	390	392
3.8 km	325	329	340	332
3.7 km	299	289	300	293
3.6 km	250	258	248	255
3.5 km	200	198	199	201
3.4 km	178	168	170	173
3.3 km	100	95	90	93
3.2 km	25	20	15	23
3.1 km	0	5	1	3
3.0 km	0	0	0	0

Figure 2-5. Portion of an example height vs. radial array

The array is a combination of weighted values from the current volume scan and the previous 2 volume scans for when the current volume scan is a precipitation VCP, and the previous 5 volume scans when the current volume scan is a clear-air

VCP. Keep in mind that the array is reset, or wiped clean, if the time between the current volume scan and the previous volume scan is greater than 30 min.

5. Compute ML Top and Bottom from Weighted Values

Next, the MLDA computes the melting layer top and bottom heights for each radial from the weighted values if a threshold for the weighted values is met. As an example, for a given radial, the algorithm first sums all the weighted values in a window surrounding the given radial. This window includes the radials 10 degrees either side of the given radial, and 1 km either side of the height of the previous melting layer top. The sum is done from the bottom up. If this sum is above a defined threshold (i.e. 1500), which indicates that there are enough melting layer detections to confidently identify a melting layer from the dual-pol data, the melting layer top and bottom are computed from the melting layer detections. The bottom is defined as where the sum is 20% of the total sum, and the top is defined as where the sum is 80% of the total sum.

6. Fill in Missing Radials

If the total sum from the previous step is below the defined threshold (i.e. 1500), the melting layer top and bottom heights are not computed from the melting layer detections. The melting layer top and bottom will instead be generated by one of two other methods. First, it will see if other radials in the current scan had enough melting layer detections to determine a melting layer top and bottom, and use the average top and bottom heights of those radials.

If no radials have a melting layer top and bottom height determined from melting layer detections, then it will use the default 0 Celsius height defined either in the Environmental Data Entry window, or by the latest RUC model. The top will be defined as the 0 Celsius height, and the bottom will be 500 meters below that. It is recommended to have the model update “On” as shown in Figure 2-6, as this will limit how often you will need to update the user-defined 0 Celsius height seen in the Environmental Data Entry window (see Fig. 2-7).

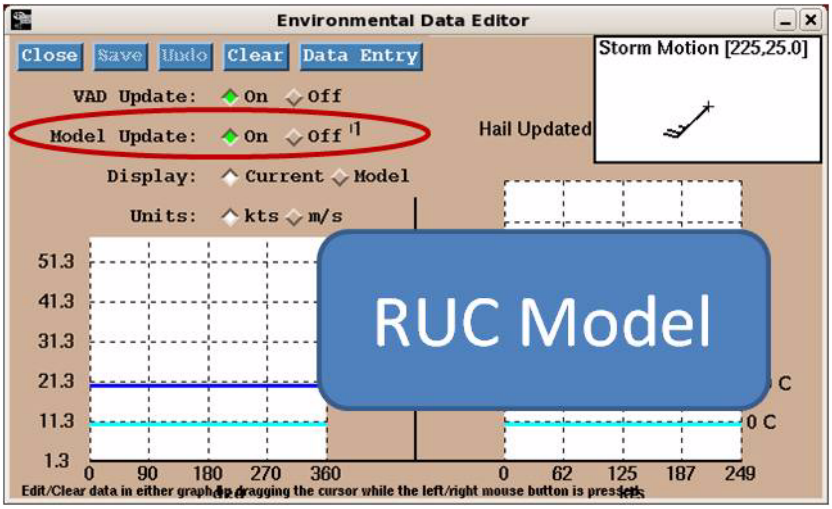

Figure 2-6. The RUC Model is used if the MLDA cannot determine a melting layer from the available dual-pol data.

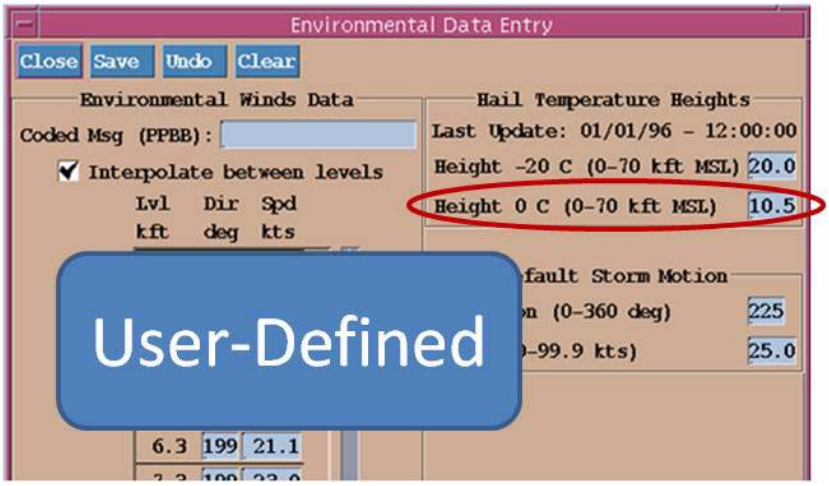

Figure 2-7. 0 Celsius height can be manually defined in the Environmental Data Entry window.

MLDA Adaptable Parameter

There is one editable adaptable parameter with the MLDA. Under the Algorithms window in the RPG, choose MLDA as the “Adaptation Item” (see Fig. 2-8).

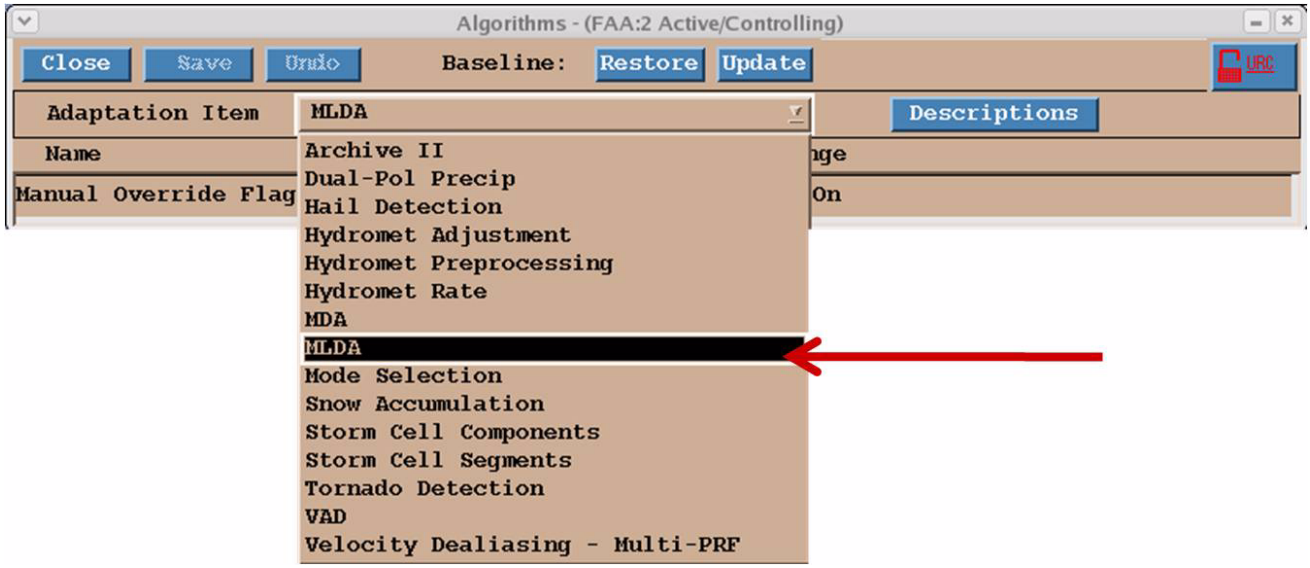

Figure 2-8. The adaptable parameter associated with the MLDA can be found in the Algorithms window.

This will display one URC parameter (see Fig. 2-9), which essentially allows the MLDA to compute the heights from the dual-polarization data, which is what was just described, or to skip that whole process and use only the input from the RPG Environmental Data instead. It is expected that the need to set this parameter to “No” will most likely be seasonal.

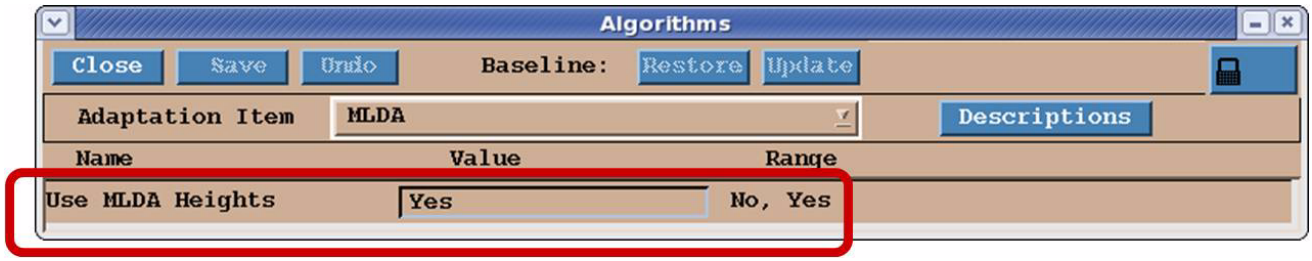

Figure 2-9. The adaptable parameter tells the MLDA whether or not to compute the heights from the dual-pol data.

The MLDA was developed primarily using warm season events from central Oklahoma. However, some cool season events have been looked at and

it appears that when the melting level is near the ground or there are two melting layers, the algorithm has some issues and it is preferable to bypass the algorithm and use a uniform melting layer representative of the actual environment.

Summary

In summary, we learned that the MLDA is based on the unique signatures of CC and ZDR in the melting layer. It searches for bins exhibiting these unique signatures from the elevations of 4° through 10°, and assigns a melting layer top and bottom height for each radial if there are enough melting layer detections. If there are not enough melting layer detections, it uses either the average top and bottom heights from the other radials in the volume scan, or the default 0 Celsius height defined in the RPG. There is one editable adaptable parameter with this algorithm and it is used to determine whether the MLDA results come from the dual-pol data or from the RPG defined 0 Celsius height.

Hydrometeor Classification Algorithm (HCA)

The dual-pol base products reveal additional, and useful, information about the characteristics of the scatterers inside the radar resolution volume beyond what the current legacy system already tells us. The Hydrometeor Classification Algorithm (HCA) works by using this additional information to determine the dominant scatterer type within that resolution volume. The product output by the HCA is called the hydrometeor classification (HC) and is produced for every elevation angle.

Before diving into the details of the HCA, here are the 12 classes currently defined in the HCA (see Fig. 2-10): biological scatterers (BI), ground clutter / anomalous propagation (GC), ice crystals (IC), dry snow (DS), wet snow (WS), light/moderate rain (RA), heavy rain (HR), big drops (BD), graupel (GR) and hail possibly mixed with rain (HA). The last two classifications are unknown (UK) and no data (ND) and will be discussed later in this lesson. RF stands for range folding and is actually not part of the HC classification.

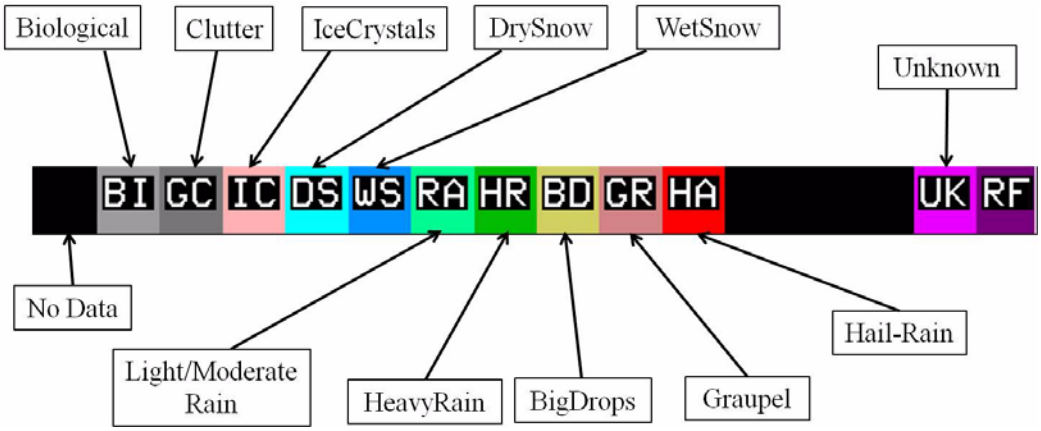

Figure 2-10. HC product color legend from AWIPS display.

Figure 2-11 on page 4-71 shows a very high-level overview of how the HCA works. The HCA assigns a hydrometeor type for each radar bin based on the dual-pol inputs. The dual-pol inputs are combined to compute a likelihood value for each hydrometeor type based on pre-defined thresholds for the various hydrometeor types, and the hydrometeor type with the highest likelihood value will be assigned to that bin.

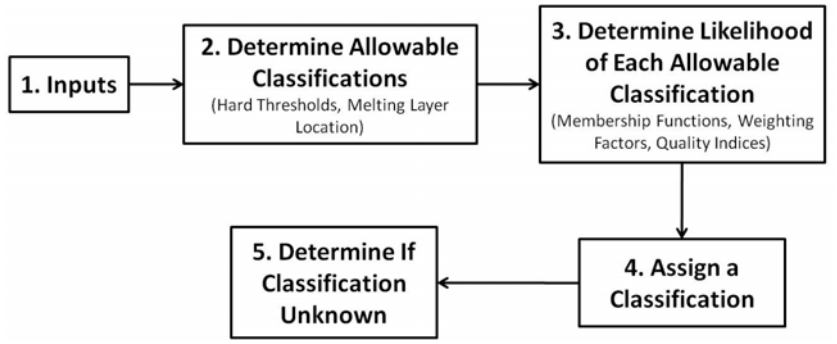

Figure 2-11. High-level overview of how the Hydrometeor Classification Algorithm (HCA) works.

1. Inputs for the HCA

The inputs for the HCA are listed in Table 4-2 on page 4-71. For the base moments, reflectivity (Z) and velocity (V) are used. For the dual-pol variables, differential reflectivity (ZDR), correlation coefficient (CC), differential phase (Φ_{DP}), and specific differential phase (KDP) are used. From Z and Φ_{DP}, two texture parameters are computed: SD(Z) and SD(Φ_{DP}). These variables have been shown to be excellent discriminators between non-meteorological and meteorological echoes. Finally, two other inputs that are used are signal-to-noise ratio (SNR) in the horizontal channel, and the melting layer top and bottom heights as determined by the melting layer detection algorithm (MLDA).

Base Moments	• Reflectivity (Z) • Velocity (V)
Dual-Pol Variables	• Differential Reflectivity (ZDR) • Correlation Coefficient (CC) • Differential Phase (Φ_{DP}) • Specific Differential Phase (KDP)
Texture Parameters	• Reflectivity (SD(Z)) • Differential Phase (SD(Φ_{DP}))
Others	• Signal-to-Noise Ratio (SNR) • Melting Layer Top & Bottom (MLDA)

Table 4-2: Inputs used to calculate the HC product.

2. Determine Allowable Classifications

Hard Thresholds

For any given range bin, the first step in the algorithm is to determine which classifications are allowed based solely on the input data. There are two ways the algorithm does this. The first is by applying hard thresholds. Table 4-3 lists the current hard thresholds, and which classification is deemed invalid based on the thresholds. For example, the classification for hail possibly mixed with rain is invalid for a bin if the reflectivity is below 30 dBZ. There is a "special case" example, however, based on SNR. If SNR is below 5 dB for a bin, then all classifications are deemed invalid, and the bin is automatically set to ND and the algorithm moves on to the next bin.

Discarded Class	Threshold (s)
All Classifications	1) SNR < 5 dB → ***HC = ND***
BI	1) CC > 0.97;
GC	1) V > 1 ms^{-1};
IC	1) Z > 40 dBZ;
DS	1) ZDR > 2 dB;
WS	1) Z < 20 dBZ *OR* ZDR < 0 dB;
RA	1) Z > 50 dBZ; 2) CC < 0.94 *AND* Φ_{DP} < 100 degrees;
HR	1) Z < 30 dBZ *OR* ZDR < 1 dB;
BD	1) Z < 10 dBZ *OR* ZDR < 0.5 dB;
GR	1) Z < 10 dBZ *OR* Z > 60 dBZ *OR* ZDR > 2 dB;
HA	1) Z < 30 dBZ;

Table 4-3: Hard threshold for each hydrometeor classification.

Melting Layer Location

The other method used to determine allowable classifications is determining where the bin of interest is located relative to the melting layer. Figure 2-12 illustrates this method. First, the HCA defines four heights based on the top and bottom of the melting layer which are labeled on the left side of the graphic. They are the heights where:

- The top of the beam enters the bottom of the melting layer
- The center of the beam enters the bottom of the melting layer
- The center of the beam exits the top of the melting layer
- The bottom of the beam exits the top of the melting layer.

Figure 2-12. Graphic illustrating the method of using the melting layer to determine hydrometeor classification.

Where the radar bin falls relative to these heights determines which classifications are allowed.

These regions are defined along with the allowable classifications in the graphic. For example, if a bin's height falls in between the height where the beam center has entered the melting layer bottom but has not crossed the melting layer top, this is defined as being "mostly within" the melting layer and the allowable classifications are GC, BI, DS, WS, GR, BD, and HA.

3. Determine Likelihood Values for the Allowable Classifications

After the allowable classifications have been identified, the algorithm then iterates through each allowable classification and calculates a likelihood value for each of those classifications. That likelihood value is a value ranging from 0 to 1 and is computed combining the following 3 factors:

- Membership functions
- Weighting factors
- Quality indices

Membership Functions

Membership functions can be thought of as lookup tables. The purpose of these lookup tables is to determine how likely a classification is based on an individual input. For example, hail typically has low ZDR due to its tumbling nature, therefore, if ZDR is near zero for a bin, then the probability of hail is high. Obviously, the other inputs need to be looked at in context, also. This will be done a couple steps later in the algorithm. At this point, the algorithm is just gathering information based solely on the individual inputs. NOTE: the lookup tables are based on research primarily conducted in central Oklahoma. Therefore, they may not be optimal for all locations when the entire network is upgraded.

Weighting Factors

Weighting factors are used to determine which inputs are best at discriminating the different hydrometeor classifications. Table 4-4 lists the weighting factors used in the algorithm. Looking at ground clutter (GC) and the input CC, the weighting factor is 1.0, but for Z, the weighting factor is 0.2. This indicates that CC is much better at discriminating GC than is Z. Again, like the lookup tables, the weighting factors have been optimized primarily using central Oklahoma cases. Therefore, these values may not be optimal for all locations. These values, and the lookup table values, are not yet editable adaptable parameters. It is not yet known which future RPG software build will provide these parameters in an editable format.

	Radar Variables					
Hydro Class Type	**Z**	**ZDR**	**CC**	**LKDP**	**SD(Z)**	**SD(Φ_{DP})**
GC	0.2	0.4	1.0	0.0	0.6	0.8
BI	0.4	0.6	1.0	0.0	0.8	0.8
DS	1.0	0.8	0.6	0.0	0.2	0.2
WS	0.6	0.8	1.0	0.0	0.2	0.2
IC	1.0	0.6	0.4	0.5	0.2	0.2
GR	0.8	1.0	0.4	0.0	0.2	0.2
BD	0.8	1.0	0.6	0.0	0.2	0.2
RA	1.0	0.8	0.6	0.0	0.2	0.2
HR	1.0	0.8	0.6	1.0	0.2	0.2
HA	1.0	0.8	0.6	1.0	0.2	0.2

Table 4-4: Weighting factors used in the algorithm.

Quality Indices

Through extensive research it has been shown that the accuracy of the dual-pol products must be quite good in order to reliably determine the most likely classification for a radar bin. Thus, the quality indices for Z, ZDR, CC, KDP and SD(Z)

and SD(Φ_{DP}) are used. These are values between 0 and 1 that indicate the quality of each of those inputs with 0 meaning the quality is lacking to 1 meaning the quality is excellent. The quality indices for each of the above inputs are dependent upon some or all of the following characteristics:

- Attenuation
- Non-uniform beam filling
- Partial beam blockage
- Magnitude of the CC
- Receiver noise, or SNR

Compute Overall Likelihood Values

Next the algorithm puts the inputs in context with one another. It combines the lookup table information with the weighting factors and quality indices for all inputs and determines an overall likelihood value for each allowable hydrometeor type. For example, suppose the allowable classifications are ground clutter, biological scatterers, heavy rain, and hail possibly mixed with rain. An overall likelihood for each of these classifications would be computed, and an example is shown in Table 4-5.

GC	BI	HR	HA
0.9861	0.7851	0.7761	0.9885

Table 4-5: Example of overall likelihood for ground clutter, biological scatterers, heavy rain, and hail.

4. Assigning a Classification

To determine which classification gets assigned by the HCA, the algorithm iterates through each classification and finds the classification with the maximum likelihood value. For the example above, iterating through these values, the algorithm would determine that HA has the highest likelihood value and would assign HA as the classification for the HC product for that range bin.

5. Unknown Classification

There are two situations where the HCA will assign a classification of unknown, and this is the final step in the HCA. The first situation is in the case of low data quality. If the maximum likelihood value is below 0.4, this indicates that the data used to compute the likelihood values was probably very poor, thus the classification will be set to unknown. Typical areas for this result are near the radar in the absence of any real strong clutter or biological returns as seen in Figure 2-13 (left) with all the pink bins in the HC product.

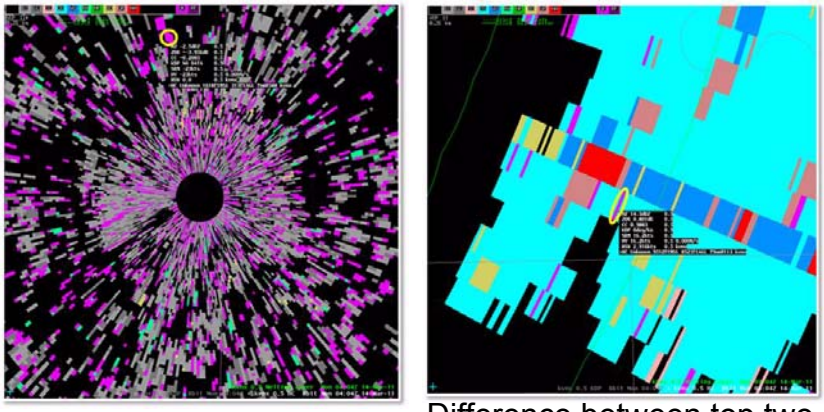

Max likelihood of < 0.4

Difference between top two classifications is < 0.0001

Figure 2-13. Example of bin with a classification of “unknown” as a result of low data quality (left) and a bin with a classification of “unknown” as a result of an overlap in characteristics (right).

The other situation is in the case of overlapping characteristics. If the top two classifications have very similar overall likelihood values and if this difference in likelihood values is less than 0.001, then the classification will be set to unknown because the algorithm cannot confidently decide which classification is dominant in that range bin. This will most likely be less common than the low data quality example, and is illustrated on the right in Figure 2-13. Note the sparse bins of pink mixed in with the other bins identified as a dry snow, or graupel, etc.

Summary

In summary, dual-polarization data reveal additional information about the characteristics of precipitation not previously shown by conventional radar. The HCA works by using this additional information provided by the base data to determine a likelihood value for each pre-defined hydrometeor type. The hydrometeor type with the highest likelihood value is tagged as the hydrometeor type for that range bin. If the confidence in the decision is deemed too low, then the range bin gets tagged as unknown. To make sure some obvious misclassifications do not sneak through, some sanity checks are put in place. For the foreseeable future, there are no editable adaptable parameters associated with the HCA.

RPG Lesson 3: Quantitative Precipitation Estimation (QPE) Algorithm

The Quantitative Precipitation Estimation (QPE) Algorithm applies the benefits of dual-pol to the very complicated task of using a radar to estimate rainfall. For example, dual-pol can better identify ground clutter and biological targets which should not be accumulated as rainfall. Range bins dominated by these non-precipitation returns are prevented from being converted to rainfall.

It will take time for QPE to reach it's full potential. In the meantime, the legacy Precipitation Processing Subsystem (PPS) and all the associated products will continue to be available for the foreseeable future.

QPE is run at the RPG, using both legacy and dual-pol base data. There are many similarities between the QPE and PPS, along with significant differences. QPE is also dependent on the results from the MLDA and the HCA, which were presented in RPG Lesson 2.

QPE and PPS in the RPG

Figure 3-1 shows an overview of the data flow for both the PPS and the QPE. The PPS is based on Z, V and SW base data, which are then processed

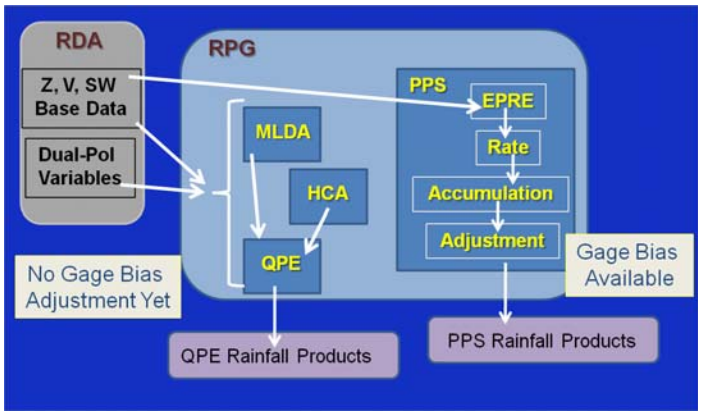

Figure 3-1. Overview of the data flow for both PPS and QPE.

through a series of functions that result in the PPS suite of rainfall products. PPS also includes a gage bias adjustment applied through the Adjustment algorithm.

All the base data (including the dual-pol variables) are available for MLDA, HCA and QPE, though the actual inputs vary. Specific inputs to QPE are: Z, ZDR, CC and KDP, as well as the results from the MLDA and the HCA. There is no gage bias adjustment yet implemented with QPE.

QPE and PPS Product Data Levels

Both the QPE and the PPS have one-hour rainfall products (see Fig. 3-2). One Hour Accumulation (OHA) is the dual-pol version of the legacy One Hour Precipitation (OHP) product. The data levels for the OHP have been editable at the RPG for many years. Now the data levels selected for the OHP also apply to the OHA product.

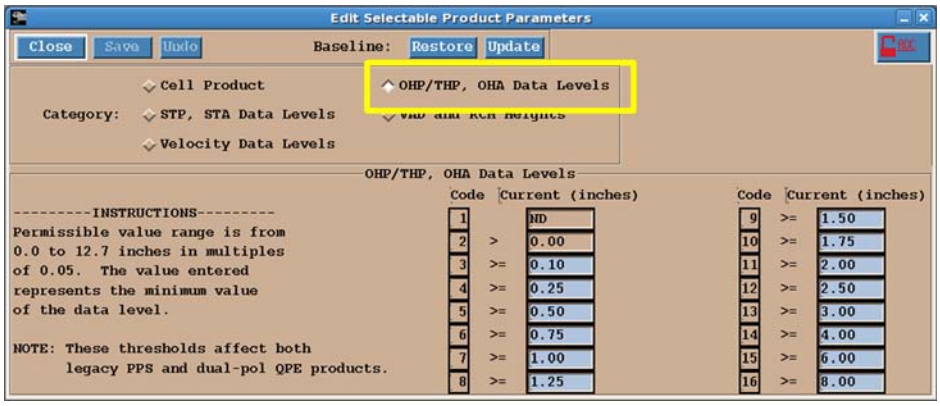

Figure 3-2. Legacy and dual-pol One Hour rainfall product (yellow box) data levels are editable at the RPG.

Similarly, both the QPE and the PPS have storm total rainfall products (see Figure 3-3 on page 5-81). Storm Total Accumulation (STA) is the dual-pol version of the legacy Storm Total Precipitation (STP) product. The data levels for the STP have been editable at the RPG for many years. Now the data levels selected for the STP also apply to the STA product.

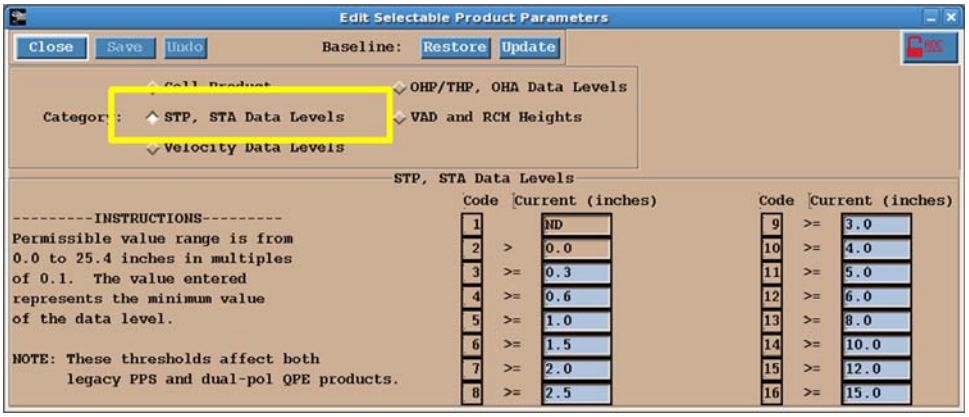

Figure 3-3. Legacy and dual-pol storm total rainfall product (yellow box) data levels are editable at the RPG.

Similar: Storm Total Accumulations

The first similarity between the PPS and QPE is how the storm total accumulations start and stop. For both algorithms, if the returns exceed thresholds of areal coverage and intensity, accumulations begin. Once the returns fall below the thresholds for one hour, accumulations stop.

PPS Parameters

The thresholds for the PPS shown in Figure 3-4 should look familiar. RAINA is the areal coverage, while RAINZ is the intensity (in dBZ) threshold. RAINA and RAINZ are accessed from the Algorithms window at the RPG under “Hydromet Preprocessing”.

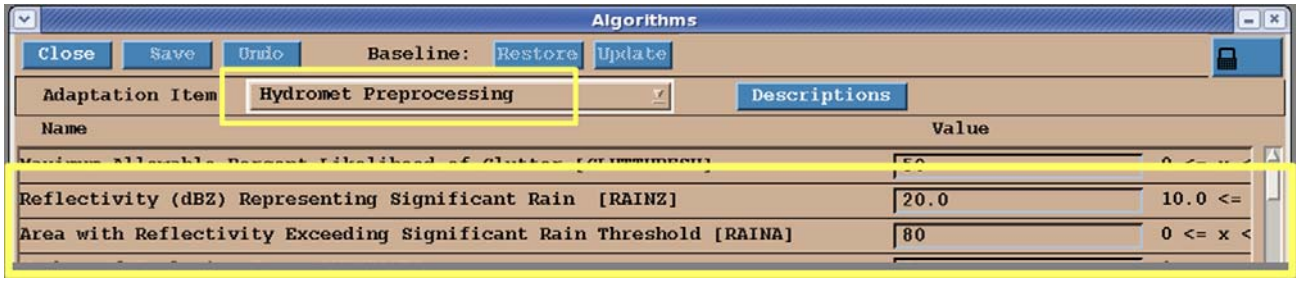

Figure 3-4. RAINA and RAINZ (bottom yellow box) indicate the area coverage and intensity thresholds, which can be accessed from the Algorithms window under “Hydromet Preprocessing” (top yellow box).

With the PPS, the total areal coverage of returns that exceed RAINZ are compared against RAINA every volume scan. If RAINA is exceeded,

accumulations either begin or continue. Storm total accumulations reset to zero automatically when returns have been below RAINA and RAINZ for one hour.

QPE Parameters

QPE uses the same concept for start and stop of storm total accumulations, but different threshold names. The QPE thresholds start with the Precipitation Accumulation Initiation Function (PAIF). The PAIF Area Threshold is analogous to RAINA, while the PAIF Rate Threshold is analogous to RAINZ. Though the units differ, the default values are the same as for the PPS. The PAIF area threshold is 80 km^2, while the PAIF rate threshold is 0.5 mm/hr, which is equivalent to 20 dBZ for $Z = 300R^{1.4}$. The PAIF thresholds are accessed from the Algorithms window at the RPG under “Dual-Pol Precip” as shown in Figure 3-5.

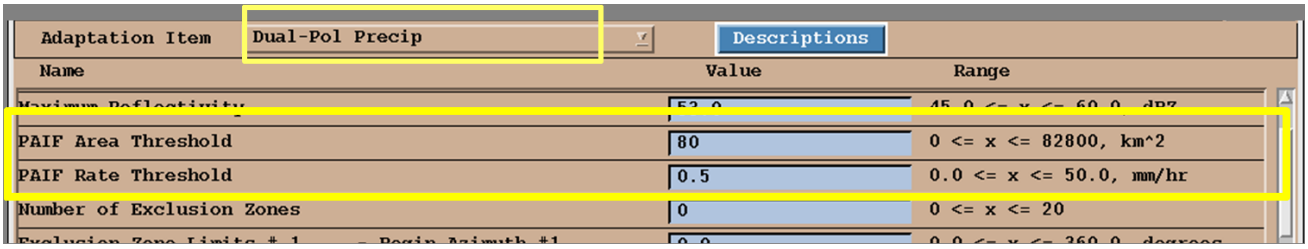

Figure 3-5. The PAIF Area and Rate Thresholds (lower yellow box) are analogous to RAINA and RAINZ, however rain rate values are expressed in mm/hr, not dBZ. These thresholds can be accessed through in the Algorithms window at the RPG under “Dual-Pol Precip.”

PPS & QPE Storm Total Accumulations

The initial recommendation is to set the parameters for both the PPS and the QPE to the same values. RAINA is a threshold value that varies from site to site, and the PAIF Area Threshold should be set to that same local value. Likewise, set the PAIF Rate Threshold to 0.5 mm/hr, which matches the default RAINZ value of 20 dBZ. When using a different RAINZ, dBZ will need to be converted to mm/hr for the PAIF Rate Threshold.

With the thresholds set to the same values, the actual start and stop times of the PPS vs. the QPE are still expected to differ. This is because the QPE can better exclude clutter, biologicals, and hail from being converted to rainfall. However, the extent of any difference in start and stop times is not yet known.

Both the QPE and the PPS have an automatic reset of the storm total accumulations after one hour of returns below their respective thresholds. It is also possible to reset both the PPS and QPE accumulations at the RPG.

From the RPG Control window, provide the URC password. This allows for a reset of the “Legacy PPS” and the “Dual-Pol QPE”. Either or both of the boxes can be selected, followed by clicking the Activate button just above as shown in Figure 3-6.

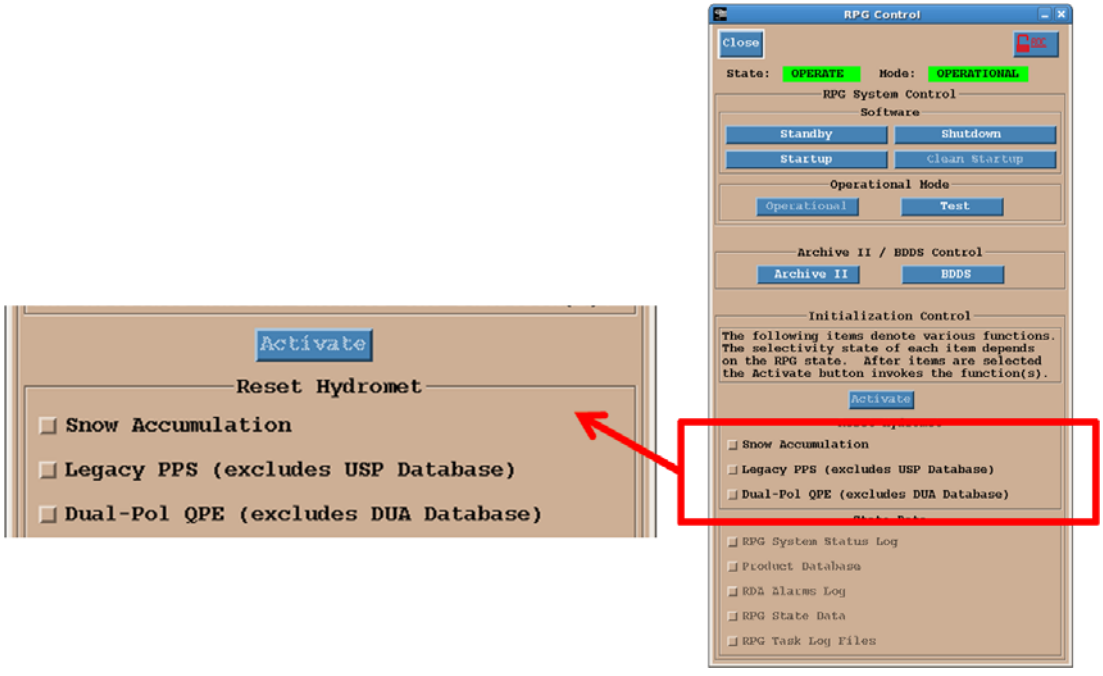

Figure 3-6. Storm total accumulations can be reset manually at the RPG for both PPS and QPE.

Similar: Exclusion Zones

Another similarity between the QPE and the PPS is the use of exclusion zones. There are moving ground targets that are not filtered as clutter. Examples are rotating wind turbine blades and traffic on roads. Returns from these targets usually exceed both the RAINZ and PAIF Rate Thresholds. Exclusion zones can be applied to both the PPS and QPE, preventing this type of radar return from being converted to rainfall.

Exclusion zones are set at the Algorithms window for Hydromet Preprocessing (PPS) and Dual-Pol Precip (QPE). Figure 3-7 on page 5-84 shows the list for QPE, but the list of parameters for defining exclusion zones is the same for each algorithm. The recommendation for the initial use of both the PPS and the QPE is to set Exclusion Zone parameters to the same values. It is expected that QPE will better identify clutter, hail and biological targets, but how that may translate to these threshold settings is not yet known.

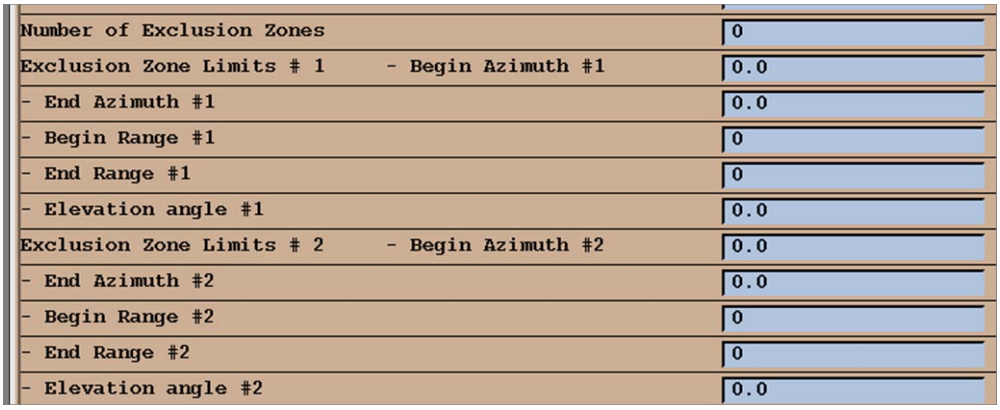

Figure 3-7. Sample list of exclusion zones for QPE.

Exclusion zones are an important tool, defining a volume from azimuth to azimuth, range to range, and up to a maximum elevation angle. The total number of exclusion zones must also be defined.

There is a misconception that exclusion zones "zero out" rainfall estimates within the zone. If a range bin falls within an exclusion zone, QPE and PPS just use the lowest elevation that is above the exclusion zone to estimate rainfall.

Figure 3-8 is an example of an event from the Dodge City CWA. The reflectivity data (center image) are contaminated by a wind farm (circled) that is to the southwest of the radar. Level II data from the event has been processed by an RPG twice. The STP product on the left was generated with an exclusion zone over the wind farm, eliminating false precipitation. The STP product on the right was generated without the exclusion zone. Note the false precipitation that results without the exclusion zone, especially the radially oriented structure.

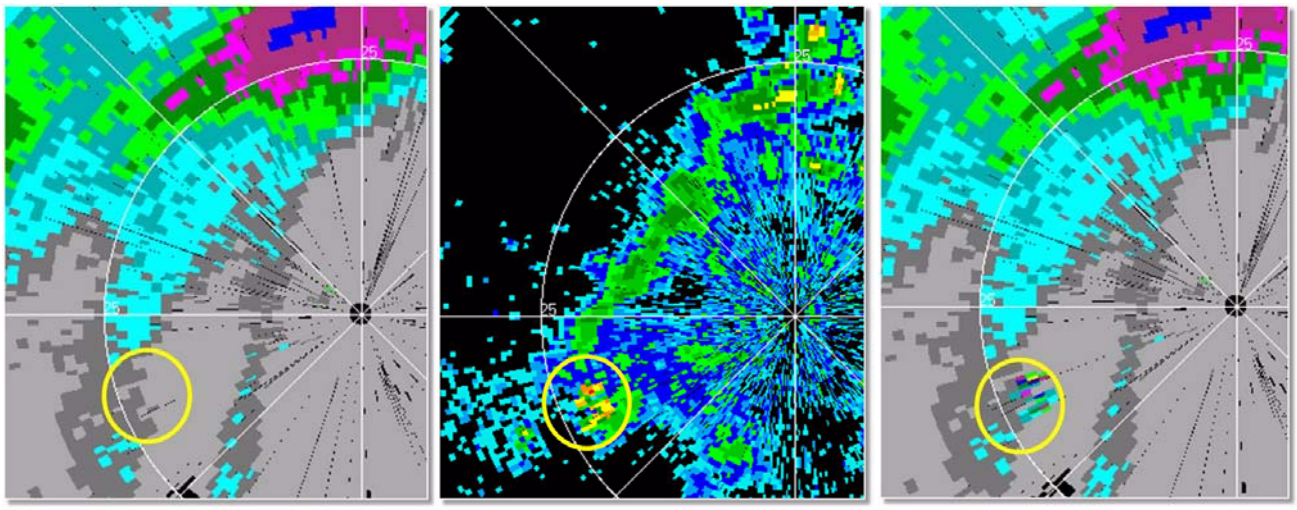

STP with exclusion zone

STP without exclusion zone

Figure 3-8. Example of reflectivity (center) contaminated by a wind farm (yellow circles). Images of STP with an exclusion zone (left) and without an exclusion zone (right) are also shown.

Different: One Hour Product

There is a significant difference between the PPS and QPE on the generation of their respective one-hour products, the OHP and the OHA.

The behavior is the same for the beginning of an event or the return of 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 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.

The difference, in either case, is that the PPS will not generate an OHP for nearly one hour, while the OHA will be available beginning with the second full volume scan.

Different: QPE Input

In addition to the dual-pol base data, QPE uses the output from the Hydrometeor Classification Algorithm (HCA) and the Melting Layer Detection Algorithm (MLDA) to select rain rates 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 computation for the particular range bin.

Different: 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. In terms of a “pre-product” product, the PPS produces the Digital Hybrid Scan Reflectivity (DHR), which is the dBZ value for each range bin before it is converted to rain rate. The QPE produces the Digital Precipitation Rate (DPR), which is the rain rate for each range bin that is used for the product accumulations. Examples of both the DHR and the DPR are presented in Figure 3-9. The DPR product is also available for input to FFMP.

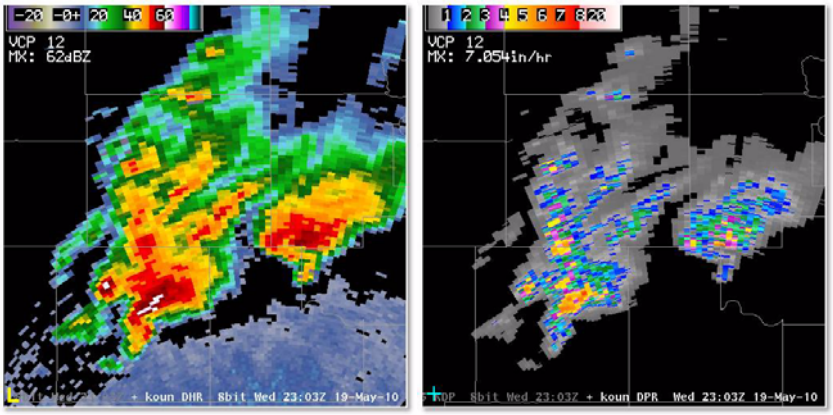

Figure 3-9. In terms of “pre-product” products, the PPS produces the Digital Hybrid Scan Reflectivity (DHR, left) and the QPE produces the Digital Precipitation Rate (DPR, right).

How QPE Builds DPR

Each WSR-88D has a local terrain file that is used to determine the percentage of beam blockage. The PPS uses this file, and for any azimuth and range that is determined to have >50% blockage, the next higher elevation angle is checked.

QPE has the benefit of Specific Differential Phase (KDP), which is immune to partial beam blockage. Thus QPE can use a lower elevation that has greater blockage. For any azimuth and range that is determined to have >70% blockage, QPE then goes to the next higher elevation angle.

Lesson 2 discussed how the HCA assigns a classification to each range bin. 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 an 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.

QPE uses a higher elevation to compute a rain rate where:

- The HCA identifies clutter or unknown, or
- There is an exclusion zone in place, or
- The beam is > 70% blocked.

All of the remaining possible classification values are some type of precipitation. QPE uses these classification values and the height of the bin with respect to the melting layer to determine a rain rate.

Different: Rain Rate Equations

The three different QPE rain rate calculation methods are each based on different inputs, and the notation tells you the input. For example, R(Z) means the conversion is based on Z alone. The familiar $Z=300R^{1.4}$ can be represented as $R(Z)=(0.017)Z^{0.714}$. For the QPE algorithm, the notations R(Z), R(Z, ZDR) and R(KDP) are used to underscore the input data for computing the rain rate. In some cases, R(Z) is used with a multiplier.

The objectives for this module do not involve remembering equations, but Figure 3-10 shows the three equations for computing rain rate used by QPE. The equations are included only to show a comparison of PPS vs. QPE approaches for converting to rain rates.

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

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

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

Figure 3-10. Equations used to calculate rain rate used by QPE.

The choice of equation is dependent on the classification value and the height with respect to the melting layer. For the R(Z,ZDR) equation, the Z and ZDR terms are in linear units (mm^6/m^3); they have not been converted to dB. The R(KDP) equation includes the possibility of a negative rain rate, however 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 $Z = 300R^{1.4}$, now presented in its R(Z) format, are probably familiar. It often underestimates rainfall in a warm rain dominant event, while it overestimates the water at the surface from mixed precipitation types, such as wet snow and hail.

R(Z,ZDR) is expected to perform better in convection than stratiform rain. For a given Z, the associated ZDR accounts for large drops (high ZDR) in low concentration (low rain rate) vs. hail contamination (high Z, but low ZDR).

R(KDP) has the advantage that KDP is mostly immune to partial beam blockage. In convective environments there is the risk of non-uniform beam filling, which can result in lowered CC values down radial from a storm core or a squall line. QPE does not use KDP values associated with CC<0.90, which can happen with non-uniform beam filling.

There are a number of steps involved in determining which rain rate is used, given the

classification value and position of the range bin with respect to the melting layer.

QPE and Melting Layer

Recall that the HCA uses the melting layer to help determine what classifications are possible. Because the QPE uses HCA and applies a modified R(Z) relationship above the melting layer (2.8 * R(Z)) and just R(Z) below the melting layer, this causes a hard transition to appear near the melting layer in the QPE products. For example, Figure 3-11 shows the reflectivity and STA product from a rather long stratiform precipitation event. Notice the sharp transition in the STA product depicted by the red arrow.

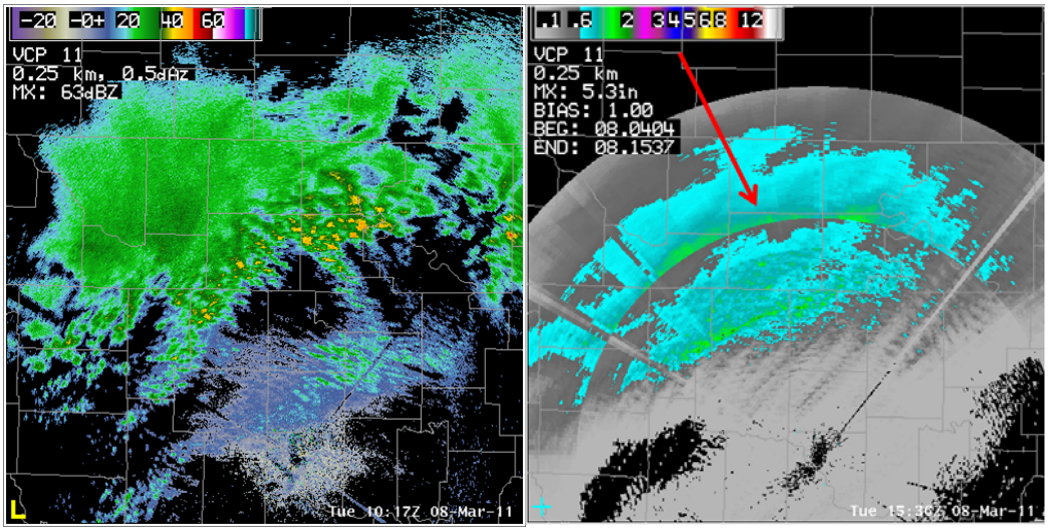

Figure 3-11. The use of R(Z) below the melting layer and 2.8R(Z) above can cause a discontinuity in the QPE product (red arrow).

QPE & R(Z,ZDR)

For Heavy Rain (HR) and light to moderate rain (RA), R(Z,ZDR) is used. Ideally, R(Z,ZDR) allows Z and ZDR to balance one another. For example, an environment with large drops is often overestimated by Z alone. The ZDR term would be large because of the drop size, decreasing the rain rate.

The R(Z,ZDR) relation has thus far shown the best performance with continental rain events, with

poorer performance for tropical events. In the tropical environment, a large number of small drops can make for very high efficiency rainfall, while the associated Z value is often underestimated. Since the ZDR is going to be low for small drops, it is less likely to compensate by increasing the rain rate.

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 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.

The bright band will still be apparent on Z or CC (see Fig. 3-12), but the overestimated Z will be mitigated before conversion to rain rate.

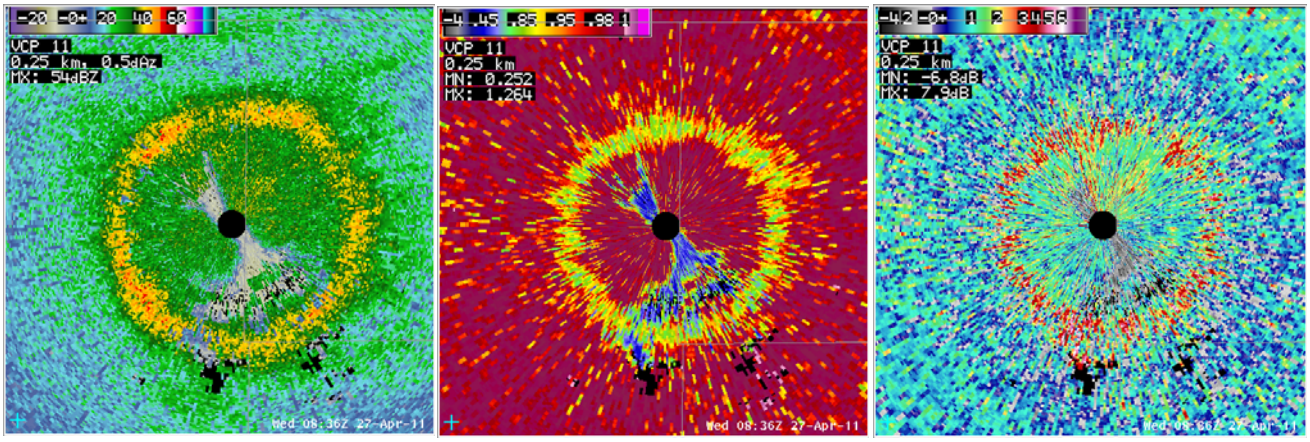

Figure 3-12. Z, CC and ZDR images of a brightband which can contaminate rainfall estimates. The QPE aims to adjust for this.

QPE and Hail Contamination

With the legacy PPS, hail contamination is mitigated by imposing a cap on the maximum rain rate used for accumulations. The limitation to this method is that there are many possible combinations of rain/hail mixtures, with many possible associated Z values. With QPE, the first

step is that range bins that are likely to contain hail or graupel are first identified 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). KDP is the dual-pol variable that indicates the magnitude of liquid water content. 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.

QPE and Non-Uniform Beam Filling

The artifact of a swath of low CC values due to non-uniform beam filling can be either easy to spot or subtle. Comparing it to other radar data and understanding the environment may help evaluate whether the CC values make sense.

It is important to be mindful of this artifact because CC affects the classification value that gets assigned, which then affects whether or nor rainfall is accumulated. In Figure 3-13, biological targets are identified by the HCA, and QPE assigns no rain rate to bins with the Biological classification value.

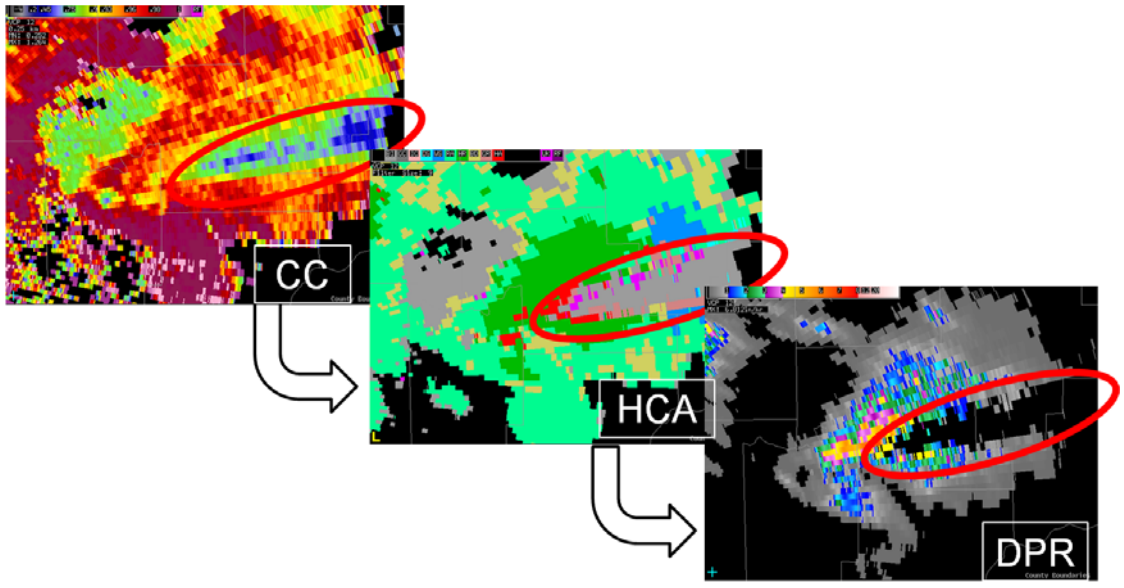

Figure 3-13. Non-uniform beam filling effects can trickle down through other products, eventually affecting the QPE.

There is no impact to the PPS that is as apparent as this swath of no rain. However, non-uniform beam filling is often accompanied by beam attenuation, which can lower the PPS estimates.

QPE and ZDR Calibration

The QPE logic and parameter design was based on an assumption of ZDR values that are calibrated to within ±0.1 dB. At the time of this writing, it is not known if that level of accuracy has been achieved. This will be an area of intense investigation for the foreseeable future.

QPE Strengths

The strengths of QPE listed below 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 well supported by the dual-pol identification of ground clutter and biological returns
- Three different rain rate equations that are applied bin-by-bin based on the classification value
- Mitigation of bright band contamination
- Mitigation of hail contamination
- Rain rate product is generated every volume scan

QPE Limitations

As with most new algorithms, there are a number of limitations to QPE (listed below) that will likely improve over time.

- Many of the QPE 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
- QPE performance is highly dependent on ZDR calibration
- If an assigned classification value is invalid, that will increase the error of the rainfall estimate
- In long term stratiform events, the QPE products are likely to show a discontinuity at the top of the melting layer

Ongoing QPE Research

QPE is a very promising algorithm, though additional research is needed to improve performance. Though ZDR accuracy is sufficient for human interpretation, improving ZDR calibration is, and will continue to be, a high priority for QPE.

QPE parameter adjustment is needed for three reasons.

1. Optimize performance for multiple regions outside of Oklahoma
2. Identify and compensate for non-uniform beam filling
3. Improve performance of the R(Z,ZDR) relationship for tropical events

Research is likely to be needed to mitigate or smooth the discontinuity on QPE products at the top of the melting layer.