

**Project Report
ATC-363**

**Signal Processing Algorithms for the
Terminal Doppler Weather Radar: Build 2**

John Y. N. Cho

February 21, 2010

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Lexington, Massachusetts



**Prepared for the Federal Aviation Administration,
Washington, DC 20591**

**This document is available to the public through
the National Technical Information Service,
Springfield, VA 22161**

ABSTRACT

As a new radar data acquisition system (RDA) was developed for the Terminal Doppler Weather Radar (TDWR), enhanced signal processing algorithms taking advantage of its increased capabilities were also developed. The primary goals of protecting the base data estimates from range-aliased signals and providing reliable velocity dealiasing were achieved through multiple pulse repetition interval (PRI) and phase coding methods. An innovative radial-by-radial adaptive selection process was used to take full advantage of the different techniques, the first time such an approach has been implemented for weather radars. Improvement in clutter filtering was also achieved. This report discusses in detail these new RDA signal processing algorithms.

ACKNOWLEDGMENTS

The in-dwell noise estimator in Section 3.2.1, the initial version of the Dijkstra algorithm in Section 3.2.2, and the pseudorandom phase code sequence in Section 3.3.2 were provided by Nathan Parker.

TABLE OF CONTENTS

	Page
ABSTRACT	iii
ACKNOWLEDGMENTS	v
List of Illustrations	ix
List of Tables	xi
1. INTRODUCTION	1
2. OVERVIEW OF SIGNAL TRANSMISSION AND PROCESSING SCHEME	3
3. DESCRIPTION OF ALGORITHMS	9
3.1 Overview	9
3.2 IQM Processes	12
3.3 IQS Processes	14
3.4 Collector Processes	53
4. SUGGESTIONS FOR FUTURE ENHANCEMENTS	63
5. SUMMARY	67
APPENDIX A	69
Glossary	75
References	77

LIST OF ILLUSTRATIONS

Figure No.		Page
1-1	Unambiguous velocity versus unambiguous range for the WSR-88D and TDWR.	2
2-1	Illustration of the MP transmission and processing technique.	3
2-2	Illustration of the DP transmission and processing technique.	4
2-3	Illustration of the adaptive mode selection process.	5
2-4	(a) Reflectivity with a constant-PRI scan, a la the legacy system. (b) Reflectivity with the new adaptive scan. (c) Radial velocity with a constant-PRI scan. (d) Radial velocity with the new adaptive scan. Censoring was not applied to the data.	6
2-5	Illustration of the SP transmission and processing scheme.	7
3-1	Overview of RDA software processes.	9
3-2	Overview of I&Q signal processing in the RDA.	11
3-3	All possible mode sequences for the first five radials diagrammed as linked nodes.	13
3-4	Flow diagram for LP processing.	15
3-5	Illustration of the spectral GCF process (adapted from Siggia and Passarelli (2004)).	22
3-6	Doppler spectrum from a range gate containing a water tower target.	25
3-7	Flow diagram for phase-code processing with LP data, part 1.	29
3-8	Flow diagram for phase-code processing with LP data, part 2.	30
3-9	Flow diagram for phase-code processing without LP data, part 1.	37
3-10	Flow diagram for phase-code processing without LP data, part 2.	38
3-11	Flow diagram for multi-PRI processing, part 1.	43
3-12	Flow diagram for multi-PRI processing, part 2.	44
3-13	Flow diagram for staggered-PRI processing.	51

LIST OF ILLUSTRATIONS
(Continued)

Figure No.		Page
4-1	Censoring of the velocity field produced by the legacy RDA at Salt Lake City using a 5×7 (azimuth \times range) filter.	64
4-2	(a) Long-PRI reflectivity providing “truth.” (b) Multi-PRI range-dealiased reflectivity. (c) Multi-PRI range-dealiased radial velocity. (d) Multi-PRI range-dealiased Doppler spectral width.	66

LIST OF TABLES

Table No.		Page
2-1	Signal Transmission and Processing Mode vs. Elevation Angle	7
3-1	Radial Buffering and Processing Scheme	54
A-1	FIR Filter Specifications and Characteristics for MP Type 4A	70
A-2	FIR Filter Specifications and Characteristics for MP Type 4B	71
A-3	FIR Filter Specifications and Characteristics for MP Type 8	72
A-4	Valid Antenna Rotation Rates	73

1. INTRODUCTION

The original radar data acquisition (RDA) system of the Terminal Doppler Weather Radar (TDWR) contained many custom boards, which made the long-term maintenance of this radar problematic. In response, the Federal Aviation Administration (FAA) commissioned the Massachusetts Institute of Technology Lincoln Laboratory (MIT-LL) to design a replacement RDA that would be supportable for an extended period of time (Elkin et al. 2002). The new RDA design uses mainly commercial off-the-shelf (COTS) components for maintainability and an open, scalable computing architecture capable of supporting new, more complex signal processing algorithms (Cho et al. 2005). Therefore, the required hardware upgrade provided an opportunity for a corresponding enhancement in signal processing that could improve the quality of data produced by the TDWR.

Of the various TDWR base data quality issues, range-velocity (RV) ambiguity was deemed to be the most severe challenge nationwide. Compared to S-band radars such as the Weather Surveillance Radar-1988 Doppler (WSR-88D), the ambiguity is worse for C-band radars such as the TDWR. This is illustrated in Figure 1-1. The two curves indicate unambiguous range $r_a = cT/2$ versus unambiguous velocity $v_a = \lambda/(4T)$ for wavelengths corresponding to the WSR-88D and TDWR as given by the relation $r_a v_a = c\lambda/8$, where c is the speed of light, λ is the radar wavelength, and T is the pulse-repetition interval (PRI). The thick lines superimposed on the curves represent the operational ranges for velocity estimation of the two radars, which are bounded on top by the minimum allowable PRI of the transmitters and on bottom by the signal coherence limit, $v_a \geq \pi W$ (Doviak and Zrnić 1993), assuming a maximum Doppler velocity spectral width of $W = 4 \text{ m s}^{-1}$. The FAA's velocity measurement requirement for the TDWR is 40 m s^{-1} , so clearly this need cannot be met without a velocity dealiasing scheme. Note that the range coverage requirement for velocity estimation is 90 km for the TDWR and 230 km for the WSR-88D. For surface scans the radar beam does not reach above the tropopause until about 460 km in range, so multiple trips of weather signals can alias into the first trip with the TDWR. Contrast this to the WSR-88D case where an operating point can be chosen such that only the second trip could alias into the first trip (albeit at the expense of lower unambiguous velocity). Therefore, a more aggressive approach must be taken to mitigate RV ambiguity for the TDWR.

Ground clutter is another critical data quality challenge for all weather radars, especially a system like the TDWR that has as its primary mission the detection of low-altitude wind shear. Surface scans for detecting microbursts and gust fronts inevitably contain strong ground clutter in many range-azimuth cells, and the signal processing must effectively filter out the clutter contamination from the desired meteorological data.

In this report we describe the first generation of enhanced signal processing algorithms inserted into the upgraded RDA. We dubbed this implementation Build 2, because the first software version (Build 1) was an emulation of the legacy processing algorithm. RV ambiguity mitigation and improved clutter filtering were the focus of Build 2. Further rounds of enhancements in the future are possible, because of

the scalable and open design of the RDA. If significant algorithm upgrades are made, follow-on reports will be issued as necessary.




Figure 1-1. Unambiguous velocity versus unambiguous range for the WSR-88D and TDWR. The thick lines indicate the operating ranges for velocity estimation mode as bounded on top by the minimum PRI allowed by the transmitter and on bottom by the signal coherency limit. The dashed line at 40 m s^{-1} marks the FAA's velocity measurement requirement for the TDWR. Note that this requirement cannot be met by the TDWR without a velocity dealiasing scheme.

2. OVERVIEW OF SIGNAL TRANSMISSION AND PROCESSING SCHEME

To counter the RV ambiguity problem, we exploited the diversity available in PRI and pulse phase with multi-PRI (MP) and phase-code transmission and processing. In MP mode, a multiple number of PRIs are transmitted within one dwell (Figure 2-1). The advantage of MP transmission is that for a given range gate, each set of PRI pulses corresponds to different out-of-trip range gates. Thus, one needs to only use the base data estimates resulting from the PRI sets with no range folding present. Velocity dealiasing can be performed within each radial using the “clean” estimates.

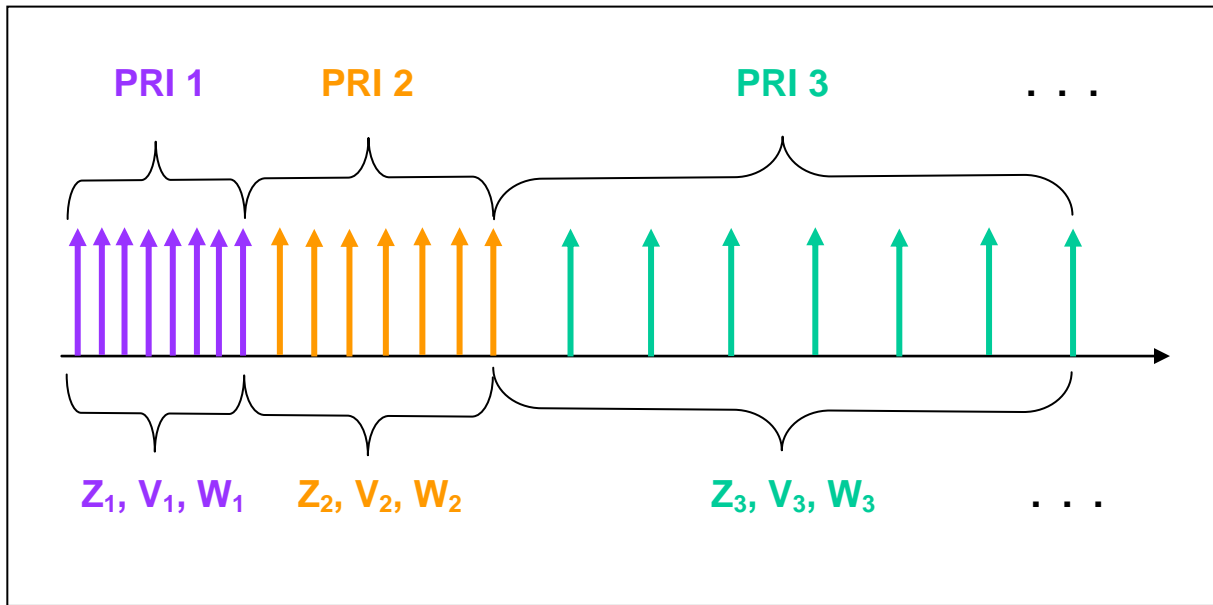


Figure 2-1. Illustration of the MP transmission and processing technique.

In phase-code mode, each pulse is tagged with a characteristic phase so that one can cohere to the unwanted trip signal and filter it out before recovering to the desired trip signal (Siggia 1983) (Figure 2-2). There are different ways of performing this filtering operation, as well as a variety of phase codes that can be used, such as pseudorandom or periodic phase codes (Sachidananda and Zrnić 1999). Although periodic phase codes can yield superior performance relative to random codes, we concluded that they have limited applicability to the TDWR because of three factors (Cho 2003)—the failure or reduced ability to provide first-trip protection against certain trips, the requirement for a particular number of points for spectral processing limiting clutter filter performance, and the need for accurate knowledge of the spectral widths for both the desired and unwanted signals for effective data quality censorship. Unlike

MP processing, phase-code processing does not provide velocity dealiasing. In order to meet the required $\pm 40 \text{ m s}^{-1}$ velocity output range, we decided to switch the PRI between two values on every dwell (radial), and then perform velocity dealiasing across adjacent radials. We refer to this transmission and processing scheme as dual-PRI phase-code (DP) mode. Note that periodic phase coding requiring certain numbers of data points per dwell is also more difficult to combine seamlessly with DP due to the significantly different number of pulses transmitted on neighboring radials.

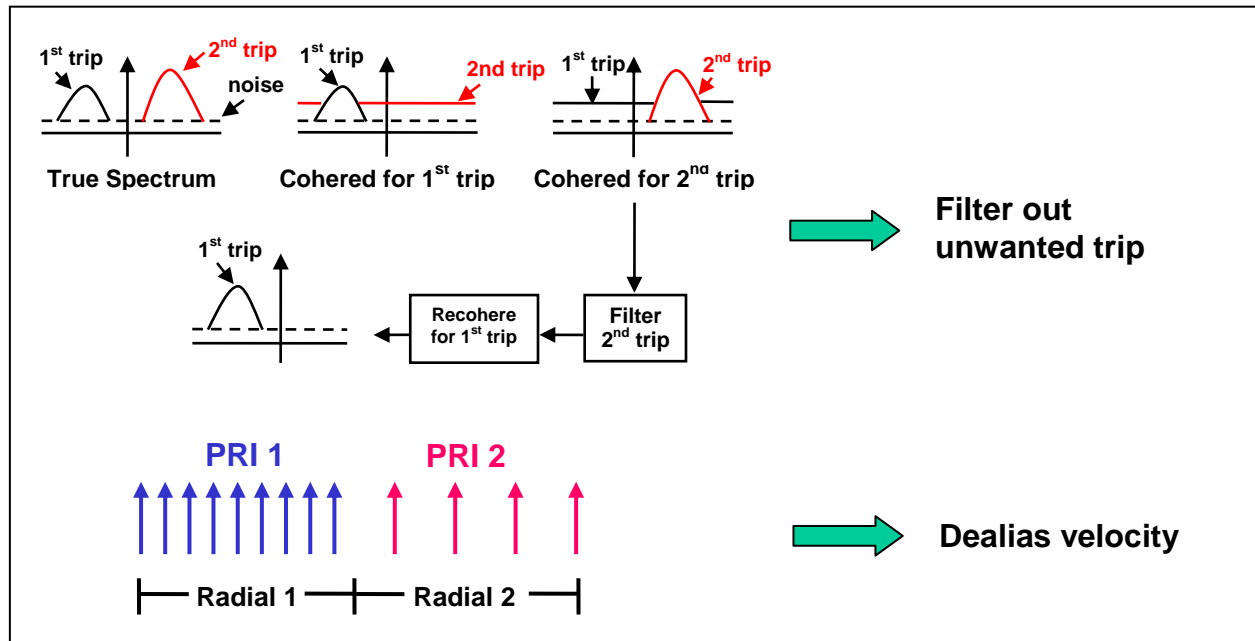


Figure 2-2. Illustration of the DP transmission and processing technique.

These two approaches (MP and DP) have complementary strengths and weaknesses for range-overlay protection (Cho et al. 2003). MP signals can be processed to effectively separate different-trip weather even if the overlaid powers are strong or spectrally wide, as long as the overlaid weather does not continuously span a long radial distance. DP processing works well for trip separation even if the overlaid storm has a long continuous radial range, but breaks down in cases of strong and/or spectrally wide overlays, and also if there are simultaneous overlays from different trips. In order to take maximum advantage of both methods, we implemented an adaptive solution where, for the surface scan, information from an initial long-PRI (LP) scan is used to select MP or DP signal transmission and processing on a radial-by-radial basis in the subsequent scan (Figure 2-3). This is a logical extension of the legacy

processing scheme where the initial LP scan is used to select two constant PRIs for the following two scans to provide correct overlay censoring (Crocker 1988) and velocity dealiasing (Wieler and Hu 1993). With the new scheme, the second dealiasing scan is eliminated and better range-overlay protection is provided. An example of the new RV ambiguity mitigation schemes uncovering a gust front that would have been obscured in a legacy-style constant-PRI scan is shown in Figure 2-4.

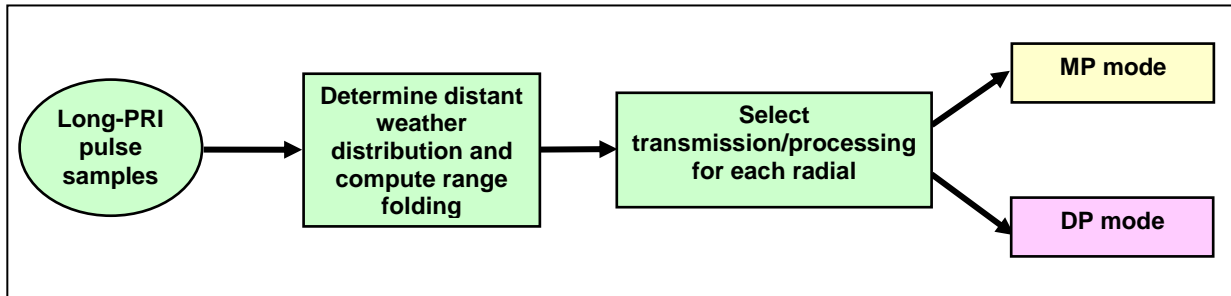


Figure 2-3. Illustration of the adaptive mode selection process.

For high-elevation tilts where range ambiguity ceases to be an issue (because the first trip covers the entire slant-range from which weather returns are possible), we implemented the staggered PRI (SP) signal transmission and processing technique (Figure 2-5). SP processing allows intradwell velocity dealiasing and reduced Doppler estimate variance (due to increased pulse-pair independence) relative to adjacent pulse-pair processing. If ground clutter is present, split time series spectral processing is used to filter it out; otherwise, pulse-pair processing is performed on each PRI and the results are used to generate a dealiased velocity.

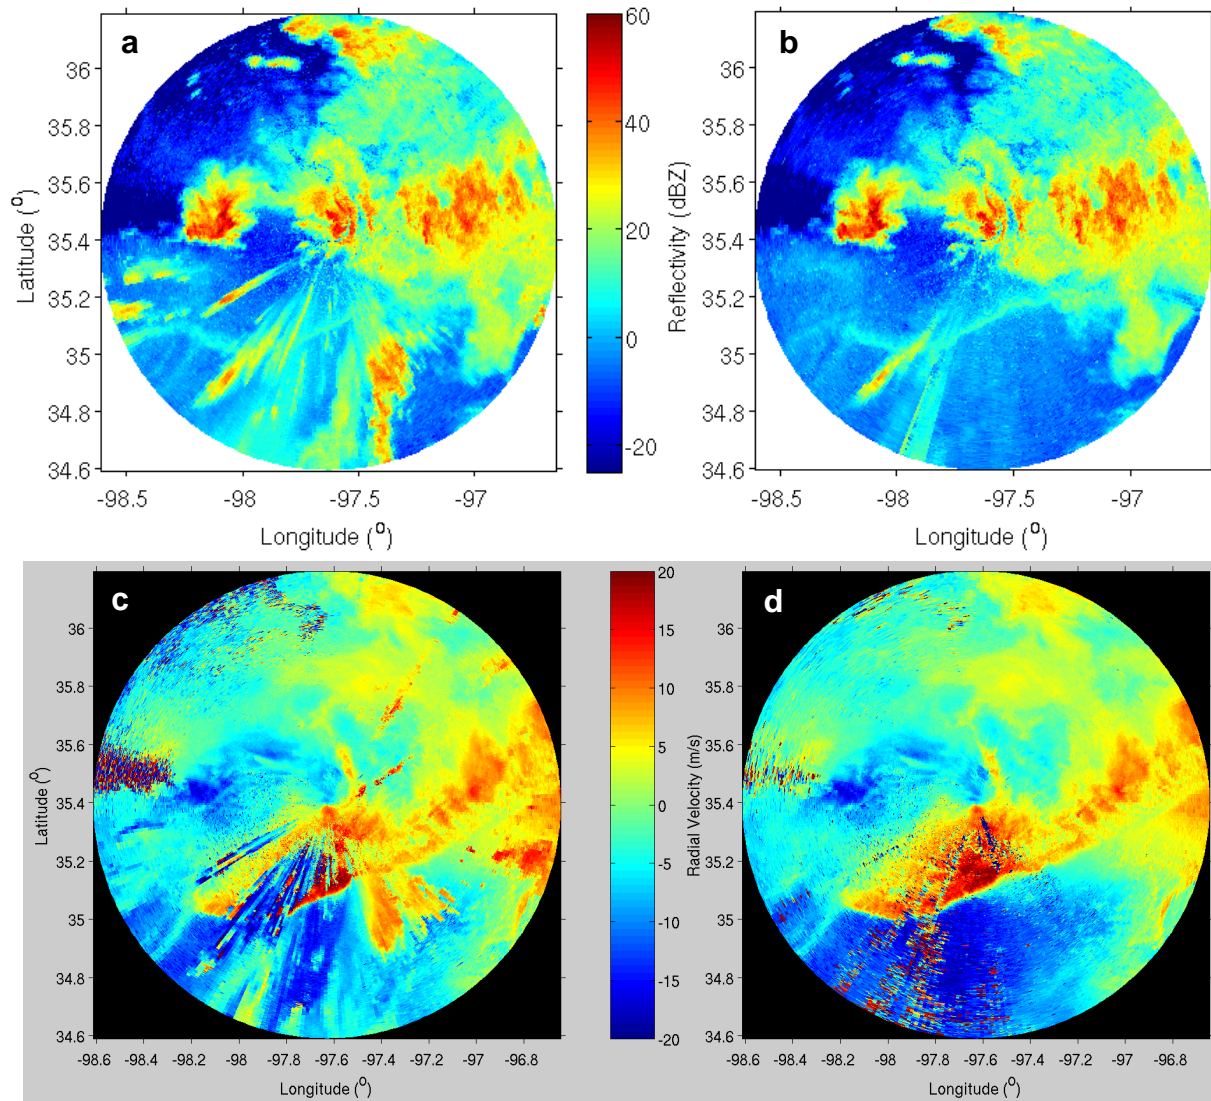


Figure 2-4. (a) Reflectivity with a constant-PRI scan, a la the legacy system. (b) Reflectivity with the new adaptive scan. (c) Radial velocity with a constant-PRI scan. (d) Radial velocity with the new adaptive scan. Censoring was not applied to the data. This 0.3° scan was taken at 03:00 Z, 14 May 2005, with the Program Support Facility (PSF) TDWR in Oklahoma City using the prototype RDA. See Cho et al. (2005) for further details.

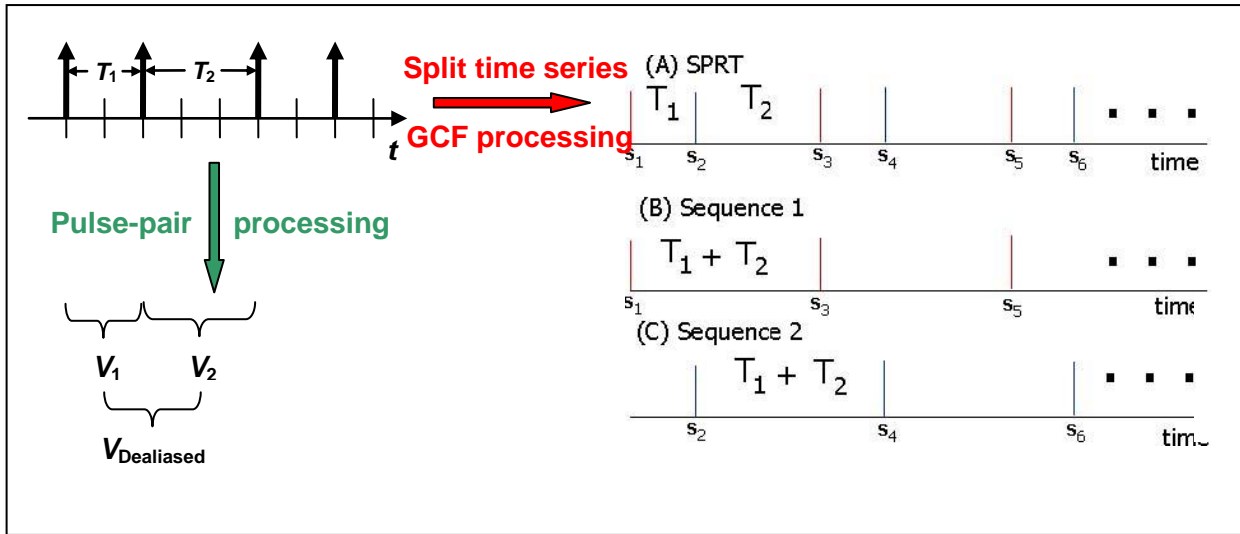


Figure 2-5. Illustration of the SP transmission and processing scheme. Split time series plot is from Figure 1 of Meymaris et al. (2009).

At intermediate tilts, where range-folding is possible but LP surveillance scan is not conducted, the DP scheme is used. The choice of available modes vs. elevation angle is summarized in Table 2-1.

TABLE 2-1
Signal Transmission and Processing Mode vs. Elevation Angle

Elevation	Mode	RV Ambiguity
Surface	LP, DP/MP*	RV
Surface < EL < 11.9°	DP	RV
≥ 11.9°	SP	V

*Adaptive selection.

Velocity is dealiased using the unfolded-velocity matching (UVM) algorithm (Trunk and Brockett 1993). The UVM technique performs better than the commonly used Chinese remainder theorem approach for number of PRIs greater than two, and provides more flexibility in the choice of PRIs and the maximum dealiased velocity interval (Cho 2005). Finally, because velocity dealiasing inevitably generates some incorrectly dealiased data, we developed a two-dimensional (2D) false-dealias correction (FDC) filter (Cho 2005).

The ground clutter filter (GCF) used depends on the transmission and processing mode. For the LP and DP modes, we applied a procedure similar to the Gaussian model adaptive processing (GMAP) algorithm (Siggia and Passarelli 2004). For the MP mode, we developed an adaptive finite impulse response (FIR) filter selection algorithm (Cho and Chornoboy 2005). For the SP mode, we split the time series into two evenly spaced sequences (Figure 2-5) before applying the same spectral GCF used in the LP and DP modes (Meymaris et al. 2009).

Note that all modes utilize pseudorandom phase coding on transmission. Even when the returned signals are not phase-code processed, if they are cohered to the first trip then the other trip signals will be rendered incoherent (i.e., white noise in the spectral domain), which removes the velocity estimation bias associated with a range-overlaid signal (Laird 1981). The ability to cohere to the measured phase of the transmitted signal (taken from the burst pulse sample) is a new feature available in the upgraded RDA.

All of these modes and algorithms will be explained in more detail in the rest of the report. In the next section we begin the description at the top level then subsequently drill down to lower levels.

3. DESCRIPTION OF ALGORITHMS

3.1 OVERVIEW

The digital signal processing (DSP) algorithms described in this report reside entirely in the RDA, specifically in the in-phase and quadrature (I&Q) master (IQM) and slaves (IQS), and the collector (shaded domain in Figure 3-1). The purpose of these algorithms is to generate base data from I&Q data. The DSP algorithms in the Vaisala Sigmet RVP9 that convert intermediate-frequency (IF) signals to base band are not discussed here. We refer the reader to Vaisala documentation for a description of those algorithms (Vaisala 2009).

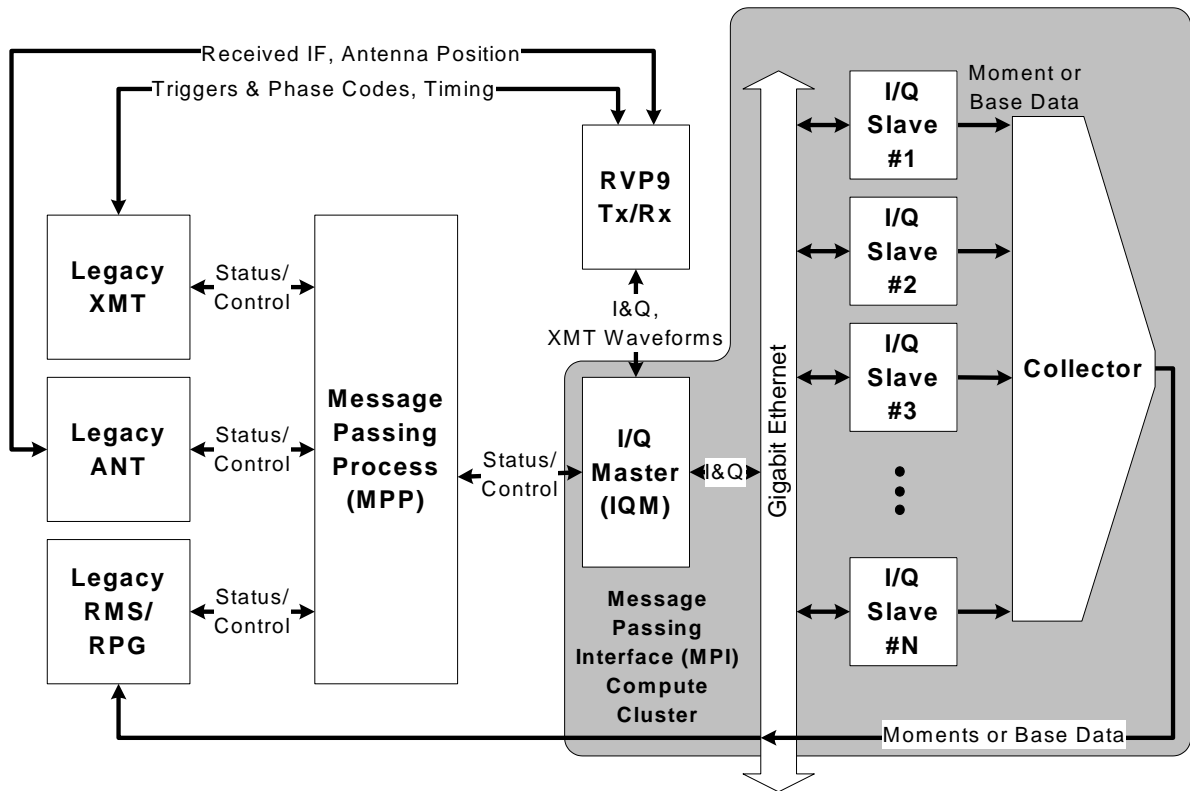


Figure 3-1. Overview of RDA software processes.

From the RVP9 the IQM receives I&Q data as well as the burst-pulse samples, which yield the transmitted pulse phases. The IQM distributes the data to the IQS for parallel gate-by-gate processing.

Processed data from all gates are assembled in the collector for 1D (range) and 2D (range-azimuth) filtering. The collector output consists of the base data and data quality flags. This is different from the legacy RDA, which only produced moments data; the radar product generator (RPG) then converted them to base data and generated the flags. These changes will be discussed more explicitly in Section 3.4.6.

Figure 3-2 provides a high-level schematic of the various I&Q DSP tasks. I&Q data for one radial (dwell) goes in and base data for one radial comes out. There is, however, a latency associated with a 4-radial buffering process in the collector. This buffering is necessary for the 2D data quality filter. Because these processes require adjacent radials on both sides of the output dwell, the number of input radials is actually 362 per 360° scan for all non-LP modes. The first and last radials are used only for providing the required information in the 2D processing and are not output to the RPG.

Dashed boxes denote processes that occur only once per full-circle scan. Dashed connectors indicate data transfers that are buffered in memory over 360 radials. Noise power is estimated from data collected at the beginning of each tilt with the transmitter turned off. On surface tilts, the processed data from the LP scan are fed back as auxiliary input to the DP and MP processing in the subsequent adaptive scan(s). The LP data are also used to determine the transmission and processing modes in the subsequent adaptive scan(s) on a radial-by-radial basis by the adaptive mode scoring and selection tasks.

The clutter residue map (CREM) editing requires maps to be produced off-line in the RPG. The algorithm for CREM generation has not changed and it still resides in the RPG, so it is not discussed in this report.

Velocity dealiasing for the MP and SP modes are done in the IQS processes, while velocity dealiasing for the DP mode is performed in the collector. A “despoking” filter is also applied in DP mode as part of the 2D data quality filter, because the PRI switching on every radial can lead to the data quality level also oscillating with every radial. For all modes, an SNR cutoff is applied at 0 dB to be consistent with legacy base data output and minimize disruptions for downstream users.

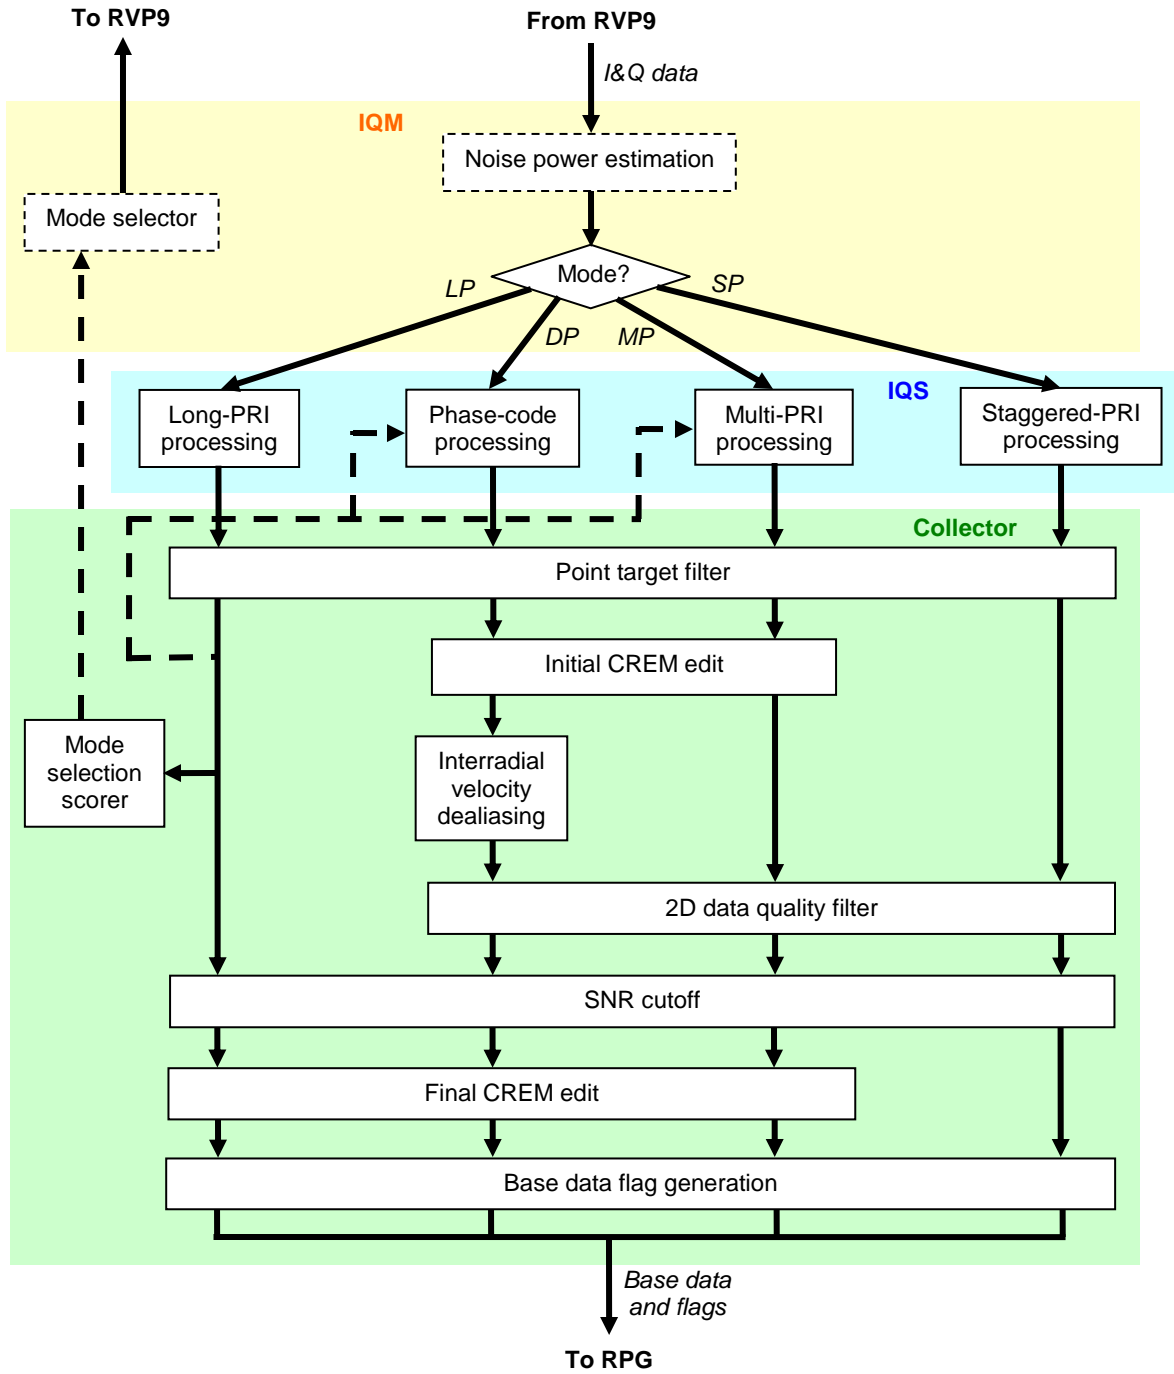


Figure 3-2. Overview of I&Q signal processing in the RDA.

3.2 IQM PROCESSES

The I&Q data stream is distributed in range-gate chunks to the IQSs for parallel processing. The IQM handles this distribution as well as other sundry low-level tasks. It also generates noise power estimates and optimized mode sequences for the adaptive scans. These two tasks are discussed in the following sections, since they fall under the aegis of RDA signal processing.

3.2.1 Noise Power Estimation

At the beginning of every elevation scan, there is a dwell (which can be more than 1° of azimuth) with the transmitter turned off. During this dwell, samples are collected for noise estimation. The noise power is computed as $P_N = \text{median}(|s_{kl}|^2) / (\ln 2)$, where s is the complex I&Q signal, k is the range gate number, and l is the pulse time index. The median is taken over all range and time indices in the dwell, which helps to filter out sporadic interference. The $\ln 2$ factor converts the median to a mean for an exponential distribution function. The I&Q noise power distribution is expected to be exponential, because the initial Gaussian distribution in voltage is transformed to a Rayleigh distribution by the intermediate frequency (IF) narrowband filter, then the absolute-value-squared operation results in an exponential distribution.

The noise power, in general, is a combination of the radar system noise and the external noise from the ground, atmosphere, space, and any other sources with energy in the C band. Thus, P_N varies with elevation angle, with the value tending to decrease with increasing angle (mostly near the surface). This is the reason for estimating noise at each elevation angle. Interference that is persistent during the noise dwell can also temporarily elevate the noise estimate. To filter out such effects, P_N values are stored at each elevation for 3 consecutive scans, and the median value is output for current use. At start-up, the default noise power values are loaded into the two previous elements per elevation.

3.2.2 Mode Selector

This process is only used after the LP surface scan. It takes the radial scores computed by the mode selection scorer (Section 3.4.7) and generates the optimal radial-by-radial mode sequence to be transmitted and processed in the subsequent adaptive surface scan(s). The scores indicate the expected quality of the velocity estimates averaged appropriately over each radial for a given mode. However, we cannot simply choose the mode with the best score for each radial, because the DP mode requires at least two consecutive radials for interradial velocity dealiasing. We have, therefore, devised an algorithm based on a shortest-path method (Dijkstra 1959) to find the optimum mode sequence over 360 radials.

The difference between choosing the DP mode vs. MP mode is that whenever a switch is made from the latter to the former, two consecutive DP radials must be used. In all other cases—DP to DP, MP to MP, and DP to MP—only one radial at a time needs to be set. This situation is illustrated in Figure 3-3 for the first five radials. The quantities shown in the figure are “distances” (or more generally, costs) taken to be the negative of the radial scores associated with the move from one node to the next. The goal is to find the shortest (least costly) path between the start point and end point at the 360th radial.




Figure 3-3. All possible mode sequences for the first five radials diagrammed as linked nodes.

To convert the radial mode scores passed from the scorer to costs, we do the following. For MP to MP costs, $L_M(n) = -\max[Y_{MP1}(n), Y_{MP2}(n), Y_{MP3}(n)]$, where n is the radial index. Keep track of which MP mode was used for each radial, so the right ones can be assigned at the end. For DP to DP costs, $L_D(n) = -Y_{DP}^{PN}(n)$. For DP to MP costs, $L_{DM}(n) = -Y_{MP}(n) + Y_{DP}^{PN}(n - 1) - Y_{DP}^P(n - 1)$. The last two terms in this expression is a needed correction for the cumulative cost to this point, because switching from DP to MP alters the last DP radial score from Y_{DP}^{PN} to Y_{DP}^P (see Section 3.4.7). For MP to DP costs, $L_{MD}(n) = -Y_{DP}^N(n - 1) + Y_{DP}^{PN}(n)$.

The Dijkstra algorithm increments from the first to the last radial and keeps track of the route of the smallest cumulative cost to each node for the DP and MP sides. At each point, there are at most two possible previous paths—from the same mode or from the other mode. When the process reaches the final radial, the node with the smaller cost is chosen and the route taken is traced back to the beginning. If any MP nodes were chosen, then the correct MP types are assigned to those radials. Of course, the whole process could have been applied directly to the scores, rather than to their negatives, and the objective inverted to maximize the cost, but we implemented the algorithm for minimization, because that is how the problem is usually couched and solved.

3.3 IQS PROCESSES

This is where the heart of the I&Q DSP resides. Because of the heavy computational burden, the IQM divides the load across the available IQS processors according to the number of range gates. The processing is dependent on the data type (Figure 3-2). Each range gate computation is independent of others. Processing that require results from other azimuths and/or other range gates are done in the collector.

3.3.1 Long-PRI Processing

This is the mode used in the first surface tilt of every monitor or hazard volume scan. The purpose is to obtain unambiguous reflectivity (Z) data to 460 km. The SNR, clutter power (P_C), and spectral width (W) data are used to select on a radial-by-radial basis the transmission modes in the following surface scan, as well as in the processing of those modes.

The flowchart for LP processing is shown in Figure 3-4. This process is repeated for each range gate's worth of data. The number of pulses is dependent on the PRI and the antenna rotation rate, but it can vary slightly from radial to radial. I&Q data, transmitted pulse phase angle, PRI, antenna rotation rate, radar parameters, and noise power are passed in from the IQM. SNR, Z , W , and P_C are output to the collector. No internal flags are generated in this mode.

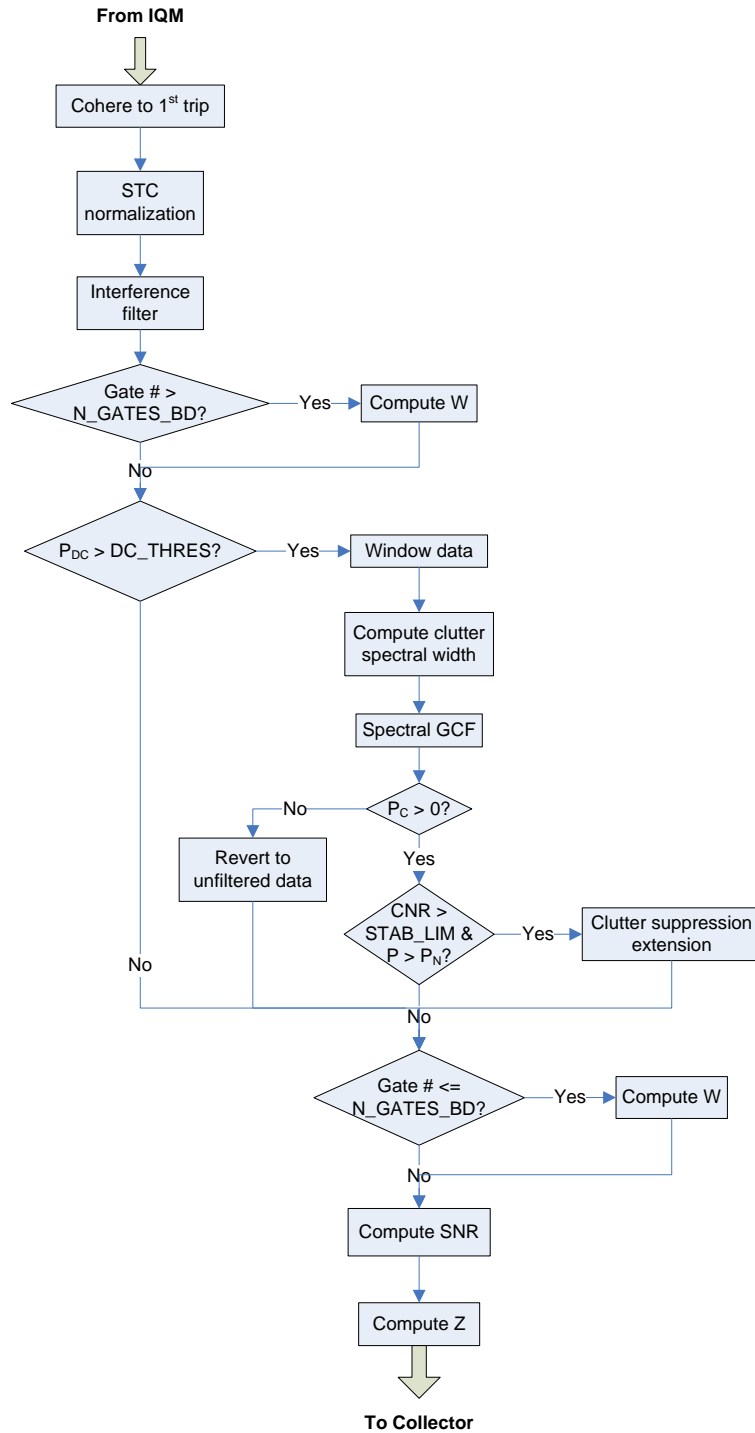


Figure 3-4. Flow diagram for LP processing.

3.3.1.1 1st trip coherence

This function coheres the received pulse data to the phase of the most recently transmitted pulse (the first trip). Other trip returns will then be rendered incoherent, because all transmission modes utilize pseudorandom phase coding. The I&Q data cohered to the first trip are given by

$$s_{kl} = s_{kl} e^{-j\phi_l} , \quad (3-2)$$

where ϕ is the measured transmitted pulse phase angle.

3.3.1.2 STC normalization

In order to reduce receiver saturation at close range, the TDWR employs a sensitivity time control (STC) device. The attenuation curve used by the STC is proportional to r^{-2} in power, or r^{-1} in amplitude. Only the first 60 range gates are affected by the STC. To reverse the STC attenuation the I&Q data are normalized as

$$s_{kl} = s_{kl} \alpha_k , \quad (3-3)$$

where

$$\alpha_k = \begin{cases} 20 & k \leq 3 \\ 60/k & 3 < k \leq 60 \\ 1 & k > 60 \end{cases} . \quad (3-4)$$

These theoretical values for α_k can be slightly adjusted in real time based on actual measurements.

Because the STC is located on the antenna side of the low-noise amplifier (LNA), the noise power must also be normalized with range as

$$P_{Nk} = \begin{cases} P_N^0 + P_N^{\text{high}} \left(\frac{2}{k} - 1 \right) & k \leq 60 \\ P_N^0 & k > 60 \end{cases} , \quad (3-5)$$

Where P_N^0 is noise power for the current tilt computed in the IQM and P_N^{high} is the noise power from the highest elevation tilt that was computed and stored previously. The assumption is that P_N^{high} represents the receiver noise power only, whereas P_N^0 includes receiver noise plus all noise that came from sources in front of the STC. This is not exactly correct, but it is a good approximation. Equation 3-

5 then corrects for the amount of receiver noise amplification introduced by the I&Q normalization of Equation 3-4. In computing the SNR, both P_N^0 and P_N are used.

3.3.1.3 Interference filter

The purpose of this filter is to suppress spikes in the I&Q signal due to intermittent radio frequency (RF) interference. The algorithm was adapted from the RVP9's built-in interference filter #3 (Vaisala 2009). We decided to implement it within the RDA processors, because the interference filter can, in turn, interfere with range-overlay protection. Therefore, we wanted to be able to turn it on or off for a given range gate.

For a given gate number k , if $|20\log(|s_{l-1}|/|s_{l-2}|)| < C_1$ and $10\log[2|s_l|^2/(|s_{l-1}|^2 + |s_{l-2}|^2)] > C_2$, then replace s_l with $(s_{l-1} + s_{l+1})/2$, where $C_1 = 10$ dB and $C_2 = 12$ dB. Note that if $l < 3$ or $l = M$, the number of pulses in the dwell, then these criteria cannot be applied because the indices go out of bounds. However, since the spike detector works whether time is run forwards or backwards, we can take care of the boundary cases by simply running the indices backwards from the boundary. So for $l = 1$, if $|20\log(|s_2|/|s_3|)| < C_1$ and $10\log[2|s_1|^2/(|s_2|^2 + |s_3|^2)] > C_2$, then replace s_1 with s_2 . For $l = 2$, if $|20\log(|s_3|/|s_4|)| < C_1$ and $10\log[2|s_2|^2/(|s_3|^2 + |s_4|^2)] > C_2$, then replace s_2 with $(s_1 + s_3)/2$. For $l = M$, if $|20\log(|s_{M-1}|/|s_{M-2}|)| < C_1$ and $10\log[2|s_M|^2/(|s_{M-1}|^2 + |s_{M-2}|^2)] > C_2$, then replace s_M with s_{M-1} .

With these thresholds, Vaisala estimates a detection probability of 94.6% and a false alarm probability of 0.85% for 16-dB interference. See the RVP9 user's manual Section 5.1.5 for further information (Vaisala 2009).

3.3.1.4 Spectral width computation decision

Spectral width in LP mode is not output to the RPG; it is only used internally in the RDA by the mode selection algorithm. Specifically, it is used in estimating the velocity estimate variance for first-trip gates and for determining the limit of range-fold protection provided by phase-code processing. For these purposes, the LP spectral width needs to be calculated without GCF processing for gates beyond the number of base data gates output to the RPG for the subsequent surface Doppler scan (N_GATES_BD), and with any necessary GCF processing for gate numbers $\leq N_GATES_BD$. The spectral width is computed using the P_s/R_1 method.

3.3.1.5 Spectral width computation (P_s/R_1)

For a given range gate, the spectral width is computed from the expression (Doviak and Zrnić 1993)

$$W = \frac{\lambda}{2\sqrt{2\pi T}} \left| \ln \left(\frac{P_S}{|R_1|} \right) \right|^{\frac{1}{2}} \operatorname{sgn} \left[\ln \left(\frac{P_S}{|R_1|} \right) \right], \quad (3-6)$$

where

$$P_S = P - P_N = \frac{1}{M} \sum_{l=1}^M |s_l|^2 - P_N \quad (3-7)$$

is the signal power,

$$R_1 = \frac{\sum_{l=1}^{M-1} s_l s_{l+1}^*}{\sum_{l=1}^{M-1} a_l a_{l+1}} \quad (3-8)$$

is the first lag of the autocorrelation function, a_l is the amplitude of any windowing function applied to the I&Q data, and asterisk denotes the complex conjugate. If no windowing was applied, the denominator in Equation 3-8 reduces to M . To avoid logarithms of zero, P_S is cut off at a very small positive number. Also, since negative spectral width values are not physically meaningful, we set W to a minimum of zero.

Note that, due to the rather narrow Nyquist velocity range, wide spectral widths cannot be properly estimated in the LP mode. For a PRI of 3066 μs and 15 data points per dwell, the spectral width estimate saturates at about 2.5 m s^{-1} with this estimator (Cho 2003). The limitation is even more severe for other estimators.

3.3.1.6 DC power threshold decision

In general, the application of GCF degrades the output data quality if the amount of clutter power present is negligible. Thus, to avoid unnecessary clutter filtering, we test the incoming I&Q data for power at zero Doppler (DC). The DC power is computed as

$$P_{\text{DC}} = \left| \frac{1}{M} \sum_{l=1}^M s_l \right|^2. \quad (3-9)$$

The criterion for attempting a GCF is $P_{\text{DC}} > \beta P_N / (MT)$, where β is a constant. The rationale for this expression is as follows. We wish to filter the clutter if it can be distinguished from noise. This suggests the criterion should be proportional to P_N / M if the clutter power is entirely contained in the DC spectral bin. However, ground clutter, in general, is not a steady spike at DC, but a fluctuating spectrum with an

exponential distribution (Billingsley 2002). So we divide by T to account for the spillage of clutter power into non-DC bins as the Doppler resolution is increased. This is by no means a rigorously derived expression, but simulations show that the ability of this criterion to detect ground clutter scales well with the different transmission modes used in the TDWR. We set $\beta = 0.005$ for our application.

Note, also, that this is not the only criterion used in choosing whether or not to use the clutter-filtered data for further processing. Therefore, it is acceptable to not set the GCF criterion at this stage too strictly.

3.3.1.7 Data windowing

Performing a digital Fourier transform (DFT) on time-series data requires careful windowing. (We say DFT rather than Fast Fourier transform (FFT), because we do not restrict the number of points in the transform to any particular set such as powers of 2.) Otherwise, sidelobes can contaminate processing in the Doppler spectral domain. The trade-off is the more aggressive the window, the more sidelobe suppression but also more loss of information. The degree of sidelobe suppression needed depends on the relative strengths of the desired signal vs. unwanted signal. An iterative algorithm a la the RVP9 internal code could be applied to converge to an optimal window selection, but under the current RDA processing hardware it could not be guaranteed that all data would be processed in real time under worst case conditions using such an algorithm. Therefore, for now, we have decided to implement a simpler window selection algorithm.

We consider two windows, the Hamming and the Blackman. For an M_{DFT} -length time-series, the Hamming is defined as

$$a_l = 0.54 - 0.46 \cos \left(2\pi \frac{l-1}{M_{\text{DFT}} - 1} \right), \quad (3-10)$$

while the Blackman we define as

$$a_l = 0.42 - 0.5 \cos \left(\frac{2\pi l}{M_{\text{DFT}} + 1} \right) + 0.08 \cos \left(\frac{2\pi l}{M_{\text{DFT}} + 1} \right). \quad (3-11)$$

This is the common Blackman function with the end-point zeros eliminated.

The windows are also normalized so that the total power after windowing is the same as before windowing. The window is multiplied by the normalization factor, which is computed as

$$\sqrt{\frac{M_{\text{DFT}}}{\sum_{l=1}^{M_{\text{DFT}}} a_l^2}} \quad (3-12)$$

Of the two, the Blackman is the more aggressive window. It allows up to about 60 dB of clutter suppression. The Hamming has higher spectral sidelobes, thus the clutter suppression capability is reduced, but more information is retained, so the base data estimate variances are decreased (assuming equal clutter suppression was achieved). To select between the two we use the SNR, P_S/P_N , of the raw data. If $\text{SNR} > 200$, we choose the Blackman, otherwise we use the Hamming. This is a fairly conservative threshold that errs on the side of more sidelobe suppression. The I&Q data are then multiplied by the chosen window, $a_l s_l$, before spectral processing.

3.3.1.8 Clutter spectral width computation

The spectral GCF needs as input the clutter spectral width. For a radar with a rotating antenna, it is usually assumed that the clutter spectrum is Gaussian in form and its width given by (Doviak and Zrnić 1993)

$$\sigma_R = 0.1325 \frac{\lambda \gamma}{\theta_B} \quad \text{m s}^{-1}, \quad (3-13)$$

where γ is the antenna rotation rate in deg s^{-1} and θ_B is the beamwidth in deg . Ground clutter in general, however, is not perfectly stationary. Winds can make vegetation flutter and power lines swing, effectively widening the clutter spectrum by generating extended tails (Billingsley 2002). This type of spectral widening varies with the weather condition and its magnitude cannot be known a priori. To account for this effect, the TDWR specifications call for an “internal motion” spectral width component of 0.1 m s^{-1} to be added to the rotational clutter width (Raytheon 1992).

Windowing the time series data before Fourier transforming them introduces a further widening bias in the Doppler spectral domain. To compute this bias we use the expression (Doviak and Zrnić 1993)

$$\sigma_W^2 = \frac{\sum_{m=1}^{M_{\text{DFT}}} v_m^2 A_m^2}{\sum_{m=1}^{M_{\text{DFT}}} A_m^2}, \quad (3-14)$$

where $v_m = 2v_a[m - 1 - \text{floor}(M_{\text{DFT}}/2)]/M_{\text{DFT}}$, “floor” denotes rounding down to the nearest integer, and A is the DFT coefficients of the window function. Thus, the final clutter spectral width is given by the sum of the rotational, window bias, and internal motion components:

$$\sigma_c = \sqrt{\sigma_R^2 + \sigma_W^2 + (0.1)^2} \quad \text{m s}^{-1}. \quad (3-15)$$

3.3.1.9 Spectral GCF

The GCF that we use for spectral processing is a modified version of the GMAP filter (Siggia and Passarelli 2004) that is used internally by the RVP9. Based on an assumed clutter spectral width and the power present in the spectrum near zero Doppler, the spectral GCF computes the theoretical Gaussian form of the clutter spectrum and removes the points for which this function is greater than the noise level. A Gaussian function is then generated using the computed spectral moments from the remaining points under the assumption that the clutter has been removed and only weather signals remain. The gap around zero Doppler is filled in using the spectral points of the Gaussian. The moments are recomputed and the gap refilled until there is reasonable convergence. (Clearly, it is assumed that the weather spectrum can be adequately represented by a single Gaussian.) This process is illustrated in Figure 3-5. The aim of this spectral GCF is to reduce the clutter filter bias by filling in the stop band with spectral points that are modeled to follow the remaining weather spectrum.

Real ground clutter spectra, however, are not necessarily Gaussian. We, therefore, added a feature to search the spectrum outward, starting from either the zero-Doppler bin (weak type) or the points where the presumed Gaussian falls to the noise level (strong type), for upward inflection points. The purpose is to extend the clutter window, if necessary, to follow a non-Gaussian tail. The strong type is specified for all function calls, except in the phase-code processing mode when there is no clutter power in the corresponding LP range gate data. For both types, the search for inflection points is limited to N_HUNT bins outside the spectral bin where the computed clutter spectrum falls to the noise level. N_HUNT is currently set to 1.




Figure 3-5. Illustration of the spectral GCF process (adapted from Siggia and Passarelli (2004)).

The following steps are performed in the spectral GCF:

1. Compute mean power, P_{unfilt} , in the windowed I&Q data, $a_l s_l$.
2. DFT $a_l s_l$ to get S_m . Spectral bins are arranged to have the DC point in the middle.
3. Get $|S_{DC}|$, the maximum of $|S_m|$ within ± 1 spectral bin of DC. This is different from the original GMAP, where the average of those three points is taken (Figure 3-5, middle plot).
4. Calculate clutter half width in spectral points

$$L_H = \text{floor} \left\{ \frac{M_{\text{DFT}}}{2v_a} \sqrt{2\sigma_c^2 \ln \left[\max \left(1, \frac{|S_{DC}^2|}{M_{\text{DFT}} P_N} \right) \right]} \right\} . \quad (3-16)$$

5. If in wide clutter mode (used on higher elevation tilts lacking CREMs to reduce sidelobe clutter), $L_H = \max\{L_H, \text{floor}[M_{\text{DFT}}\sigma_{SC}/(4v_a)]\}$. (See Equation 3-23 for the definition of σ_{SC} .)
6. Search outward from the middle until an upward inflection is found in $|S_m|$. For weak type, begin both left and right searches at the DC bin index given by $m_{DC} = \text{floor}(1 + M_{\text{DFT}}/2)$. For strong type, begin searches at $m_{DC} \pm L_H$. In either case do not exceed $m_{DC} \pm (L_H + N_{\text{HUNT}})$ or go beyond the end points of the spectrum. The points to be replaced by the GCF, the clutter gap, are defined as the points interior to the end points found in this search.
7. Compute normalized power spectral components, $Q_m = |S_m|^2/M_{\text{DFT}}^2$, and replace the clutter gap points by the spectral noise level, P_N/M_{DFT} .

8. This is the beginning of the iterative loop. Compute signal power $P_S = \max(\sum Q_m - P_N, \text{EPS})$ and first autocorrelation lag $R_1 = \sum Q_m \exp\{j2\pi[(m - 1 - \text{floor}(M_{\text{DFT}}/2)) \bmod M_{\text{DFT}}]/M_{\text{DFT}}\}$, where EPS is the smallest nonzero number supported by the computer in the data type being used for computation.
9. If $P_S = \text{EPS}$, go to step 18.
10. Compute the mean frequency estimator, μ_F , in units of spectral bins. If M_{DFT} is odd, $\mu_F = \text{round}[M_{\text{DFT}} \angle R_1 / (2\pi)]$. If M_{DFT} is even, $\mu_F = \text{floor}[M_{\text{DFT}} \angle R_1 / (2\pi)]$, where $\angle R_1$ is computed in the range $\pm\pi$. The result is quantized so the peak falls into the center of the bin.
11. Compute the signal spectral width in units of spectral bins: $\sigma_S = \max\{\text{EPS}, M_{\text{DFT}}[\max(0, |\ln(P_S/|R_1|)\text{sgn}(\ln(P_S/|R_1|))/2)|]^{1/2}/\pi\}$.
12. Calculate the coefficients of the Gaussian model fit to the signal spectrum: $Q_{Gm} = \exp\{-[\text{ceil}(m - 1 - M_{\text{DFT}}/2) - \mu_F]^2 / (2\sigma_S^2)\} / [(2\pi)^{1/2} \sigma_S]$.
13. Normalize to the signal power: $Q_{Gm} = P_S Q_{Gm} / \sum Q_{Gm}$.
14. Replace the clutter gap spectral coefficients with the Gaussian coefficients + noise ($Q_{Gm} + P_N/M_{\text{DFT}}$).
15. If not first time through iterative loop, check for exit condition: If $|\angle R_1 - \angle R_1^{\text{old}}| < \text{R1_PROGRESS}$ and $P_S < (\text{SIG_PROGRESS})P_S^{\text{old}}$ and $P_S^{\text{old}} < (\text{SIG_PROGRESS})P_S$, then go to step 18. Currently the constants are set to be $\text{R1_PROGRESS} = 0.005$ and $\text{SIG_PROGRESS} = 1.04$.
16. $P_S^{\text{old}} = P_S$; $\angle R_1^{\text{old}} = \angle R_1$.
17. End of loop; go to step 8 if number of iterations is less than MAX_ITER (currently set to 12).
18. Replace the clutter gap coefficients in the complex spectrum with the magnitude of the Gaussian model coefficients, but keep the phases from the original input: $M_{\text{DFT}}(Q_{Gm} + P_N/M_{\text{DFT}})^{1/2} S_m / |S_m|$.
19. Inverse DFT to get the clutter filtered time series, s_m^{GCF} .
20. Compute clutter power removed: $P_C = \max[0, P_{\text{unfilt}} - (\sum |s_m^{\text{GCF}}|^2) / M_{\text{DFT}}]$.

3.3.1.10 GCF decision

At this point if $P_C > 0$, then the clutter-filtered time series is used for further processing and $P = (\sum |s_m^{GCF}|^2) / M_{\text{DFT}}$. Otherwise, the original unfiltered and unwindowed data are used, and $P = P_{\text{unfilt}}$.

3.3.1.11 Clutter suppression extension

The amount of ground clutter that can be filtered from the time series is limited by the radar system stability. As the clutter power increases beyond the stability limit, the excess power goes into raising the spectral noise floor, rather than into a coherent signal around DC. The legacy TDWR specifications state stability figures of 63.7 dB for the transmitter, 64.7 dB for the receiver-exciter (REX), and 71.0 dB 12-bit quantization noise for the analog-to-digital converter (ADC), to yield an effective system stability of 60.7 dB (Raytheon 1992). In the upgraded RDA the REX and the ADC have been replaced to yield improved stability on the receiving end (exact figures are unknown). Although the transmitter remains the same, by cohering the I&Q data to the measured transmitted pulse phase, much of the transmitter phase instability can be removed. With these updates we have measured as high as 66 dB clutter suppression on a point scan of a close-range water tower target using a Blackman window and spectral GCF (Figure 3-6).

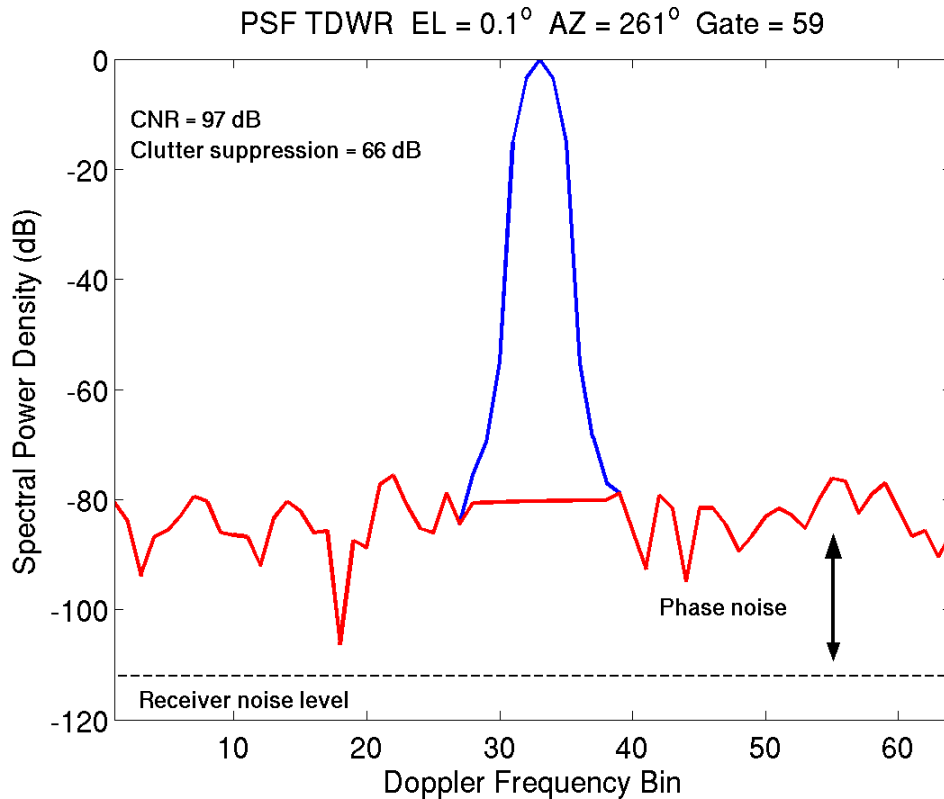


Figure 3-6. Doppler spectrum from a range gate containing a water tower target.

We cannot filter any more of the clutter signal from the time series data than dictated by this basic limitation. However, it is possible to estimate the amount of excess power that gets injected into the spectral noise floor (see Figure 3-6) and subtract this from the signal power estimate, which will improve the SNR and reflectivity estimates, and sometimes the spectral width estimate, but not the velocity estimate. We can then subtract this excess power from the original power estimate following these steps:

1. Proceed if $P_C/P_N > \text{STAB_LIM}$ and $P > P_N$. For phase-code processing, the condition that the 1st trip be the strongest trip is added to this entry criterion, because out-of-trip overlay can also appear as increased spectral noise and that will be the dominant “noise” if the 1st trip is not the strongest trip. The entry criterion helps keep the following computationally intense code from being executed unnecessarily. STAB_LIM is currently set to 3.16×10^5 (55 dB).
2. Estimate noise level, p_{SN} (see Section 3.3.1.12), of the spectrum $|\text{DFT}(s_m^{GCF})/M_{\text{DFT}}|^2$.

3. Assign $P' = P$ and $P'_C = P_C$.
4. If $p_{SN}M_{DFT} > P_N$, then assign $P = \max(P_N, P' - p_{SN}M_{DFT} + P_N)$ and $P_C = P'_C + P' - P$.

3.3.1.12 Spectral noise level estimation

Although the noise power estimation is conducted up front in the IQM (Section 3.2.1), there are other spectral “noise” contributions (system instability residue and out-of-trip signal overlay) that are not accounted for by that computation. The Hildebrand and Sekhon (1974) method can be used to directly estimate the total white noise level (p_{SN} , noise power per spectral bin). In this technique, the spectral bins are sorted by power and the bins are eliminated one by one from the strongest on down until a statistical test indicates that the remaining spectrum has the characteristics of noise.

1. Sort the spectral power coefficients, Q_m , into ascending order to get Q_m^{sort} .
2. Let $M_C = M_{DFT}$ and $b_0 = 0$.
3. While $M_C > 2$ and $b_0 < 1$, do the following:
 4. $M_C = M_C - 1$.
 5. $b_1 = \frac{1}{M_C} \sum_{m=1}^{M_C} Q_m^{\text{sort}} - b_2$, where $b_2 = \left(\frac{1}{M_C} \sum_{m=1}^{M_C} Q_m^{\text{sort}} \right)^2$.
 6. If $b_1 \leq 0$, then go to step 8.
 7. $b_0 = b_2/b_1$; end of while loop.
8. $p_{SN} = \frac{1}{M_C} \sum_{m=1}^{M_C} Q_m^{\text{sort}}$.

In addition to the noise level, some calls to this function require the sorted indices and the index to the weakest-power bin that is deemed to have a coherent signal component.

3.3.1.13 SNR computation

The SNR is computed as $(P - P_N)/P_N^0$. This is unchanged from the legacy system (Raytheon 1992). As discussed in Section 3.3.1.2, P_N and P_N^0 are only different for range gates affected by the STC. For base ata, the SNR is converted to decibel units and cut off at 0 dB (Section 3.4.5).

3.3.1.14 Reflectivity computation

This computation is also unchanged from the legacy system. See Section 8.2.1 in the Raytheon (1992) document for details.

3.3.2 Phase-code Processing

As explained in Section 1, range ambiguity is one of the TDWR data quality problems we wished to alleviate with new approaches. One such approach is to exploit phase diversity to discriminate between signals returned from different pulses. Simply put, each transmitted pulse is tagged with a particular phase value, and on reception the signal is cohered to the phase matched to one pulse back, two pulses back, etc., depending on the trip of interest. In the spectral domain, the cohered signal is reconstituted while the uncohered signals appear as noise. This procedure alone does not completely prevent range obscuration, because the uncohered signals can be so strong that the corresponding “noise” swamps the desired cohered spectrum. However, by taking advantage of the expectation that a weather spectrum is compact (i.e., narrow with respect to the Nyquist interval), one can cohere first to the undesired trip, remove the resulting out-of-trip weather spectrum, then cohere to the desired trip. This is not a perfect solution, since some of the desired trip signal is inevitably lost during the notching process.

The “noise” in the spectrum generated by an uncohered signal is white if the phase code sequence is random. If particular periodic phase sequences are used, however, the “noise” is periodic replicas of the uncohered signal spectrum. The latter has an advantage in that less of the signal information is lost during the notching process. There are, however, disadvantages to periodic phase codes. An earlier report (Cho 2003) compared the pros and cons of the pseudorandom vs. a particular periodic (SZ) phase code, and concluded that the pseudorandom code was more suitable for the TDWR. Since the TDWR uses a klystron transmitter, the transmit phases must be specified.

The actual cyclical phase angle sequence specified on transmission are $2\pi[79, 217, 194, 184, 20, 87, 49, 59, 164, 155, 219, 123, 97, 99, 0, 24, 114, 157, 23, 174, 55, 242, 141, 173, 104, 78, 9, 91, 70, 236, 80, 226, 78, 165, 167, 244, 40, 187, 138, 32, 118, 7, 78, 92, 177, 112, 81, 244, 2, 101, 75, 17, 204, 0, 104, 229, 144, 240, 9, 174, 34, 194, 175, 152, 135, 137, 46, 176, 163, 166, 199, 190, 65, 127, 221, 207, 73, 54, 153, 223, 133, 161, 178, 168, 215, 124, 52, 141, 248, 26, 35, 14, 252, 111, 209, 167, 91, 204, 31, 39, 45, 21, 15, 119, 170, 40, 138, 5, 212, 114, 162, 179, 184, 10, 167, 81, 72, 222, 115, 162, 238, 76, 3, 122, 80, 69, 193, 172]/255$. This sequence was found using a program that searched for spectra that were flat as possible for the code itself and for increment differences (modulation codes) corresponding up to the 6th trip. There are 128 values in the sequence, which is longer than any dwell that is expected to be used under normal operation. The values are quantized at 8 bits as specified for the RVP9. For signal processing, the measured burst-pulse phases are used for cohering the I&Q data, instead of the specified phases.

As discussed in Section 2, only the surface tilt employs a long-PRI surveillance scan. Thus, there can be phase-code processing with LP data or without, depending on the elevation angle. We discuss these two cases separately in the following sections.

3.3.2.1 Phase-code processing with LP data

This mode can be selected during the adaptive surface tilt on a radial-by-radial basis. If there is no range aliasing detected at any range gate, then this mode will be chosen over the MP mode, because it tends to yield estimates with lower variances. The flowchart for phase-code processing with LP data is shown in Figures 3-7 and 3-8. This process is repeated for each range gate's worth of data. The number of pulses is dependent on the PRI and the antenna rotation rate, but it can vary slightly from radial to radial. I&Q data, transmitted pulse phase angles, PRI, antenna rotation rate, radar parameters, and noise power are passed in from the IQM. LP SNR and clutter power are passed back from the collector. SNR, Z , V_{raw} , W , SQI, and internal flags are output to the collector. Velocity dealiasing takes place in the collector.

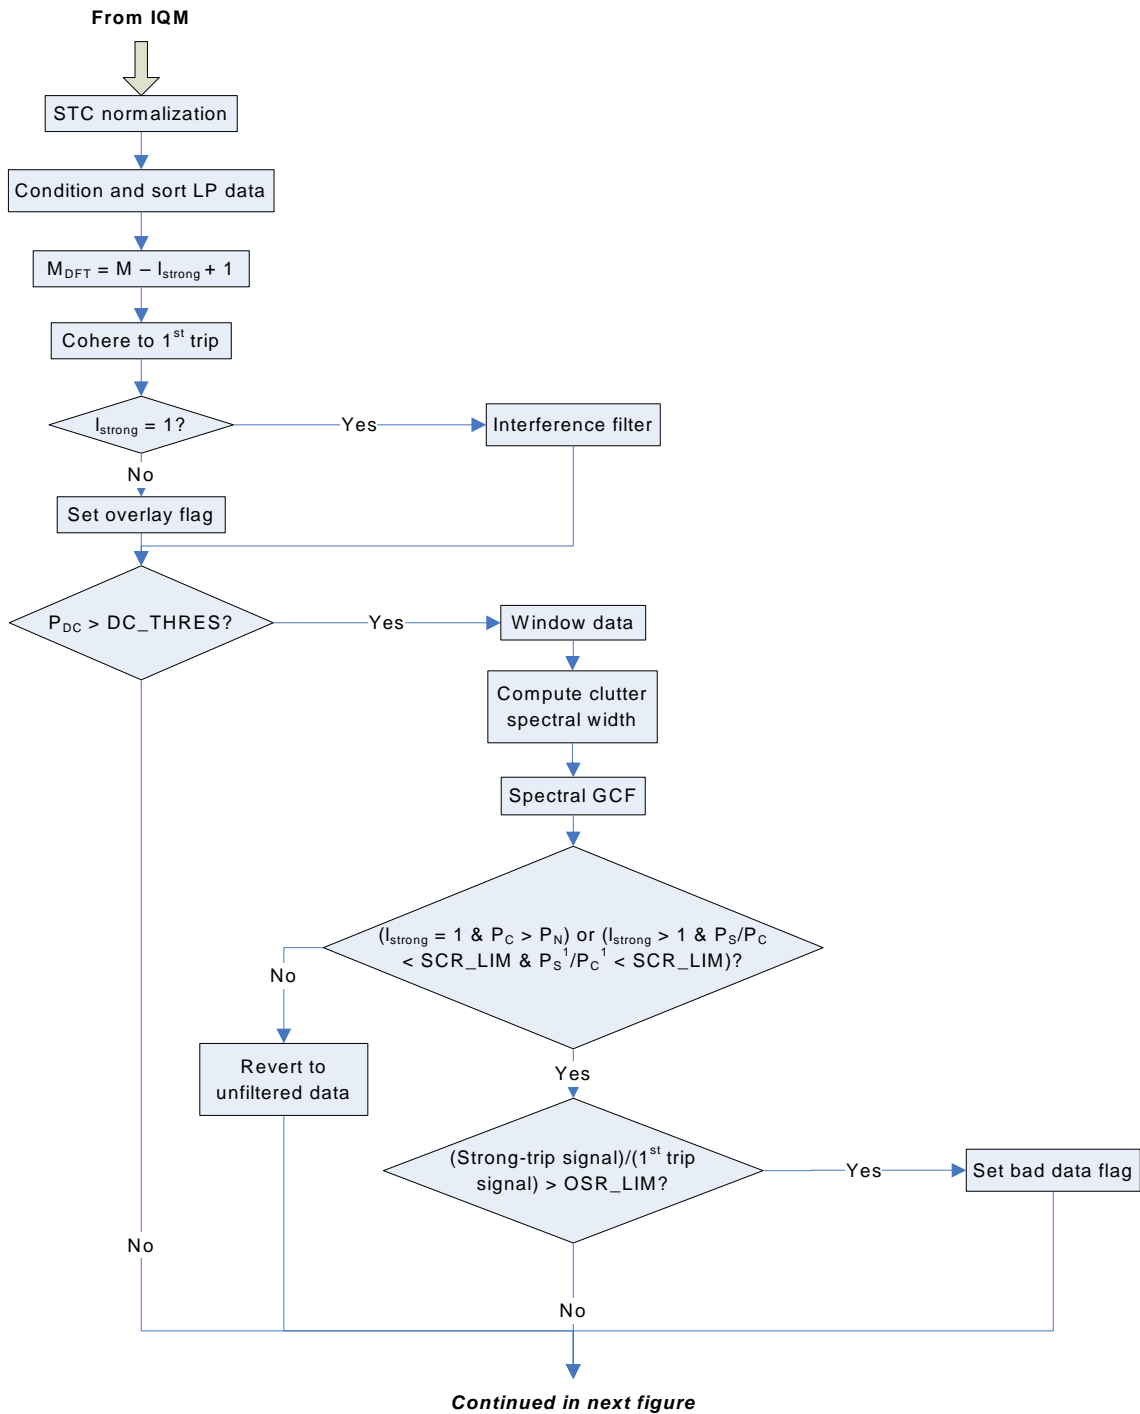


Figure 3-7. Flow diagram for phase-code processing with LP data, part 1.

Continued from previous figure

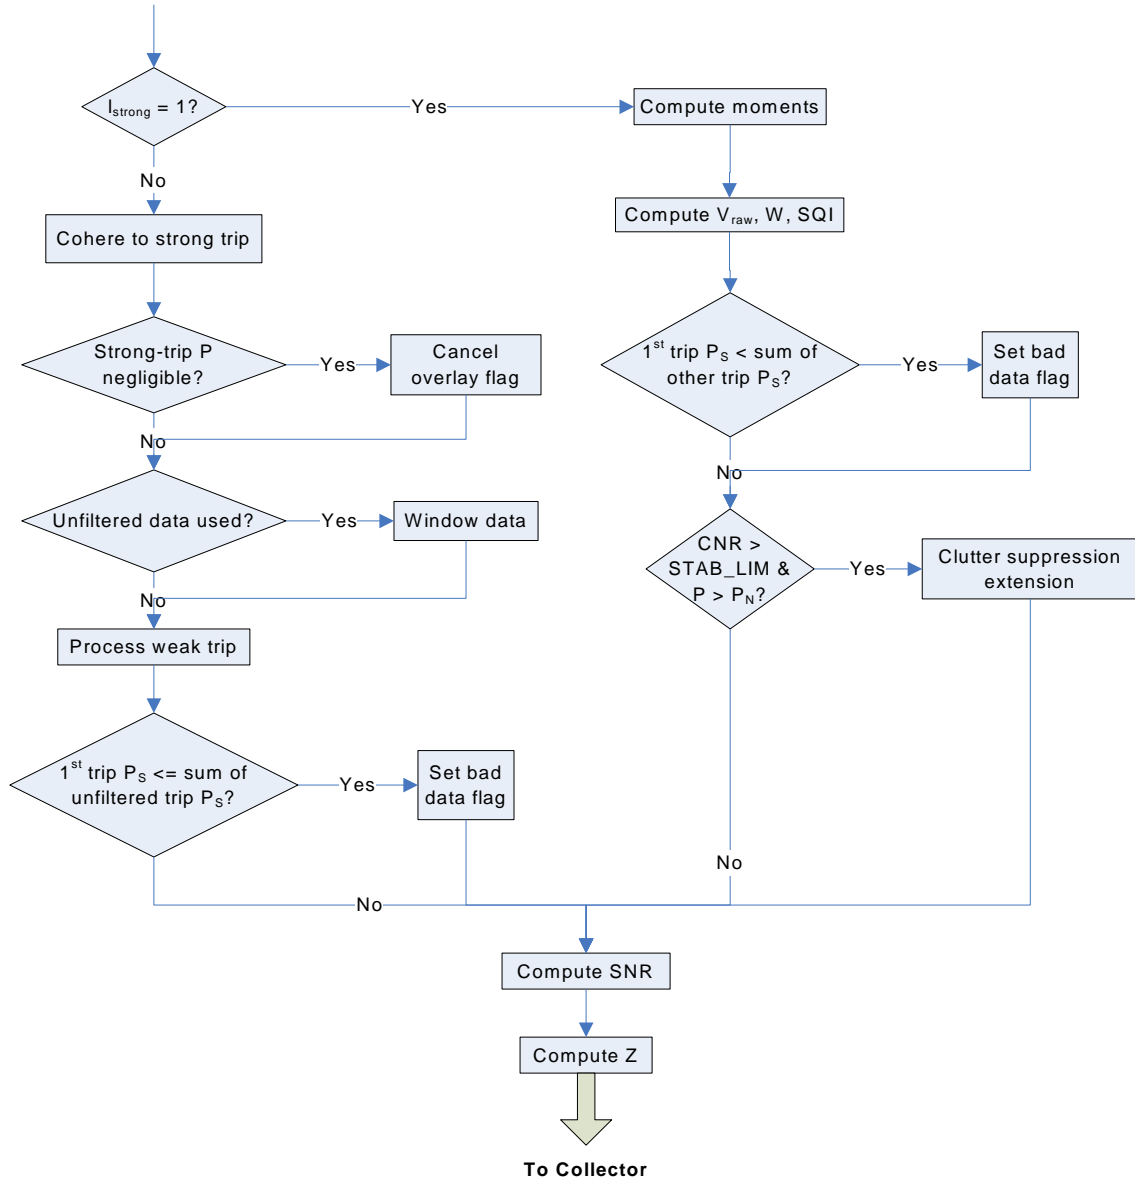


Figure 3-8. Flow diagram for phase-code processing with LP data, part 2.

3.3.2.1.1 STC normalization

This process is identical to Section 3.3.1.2.

3.3.2.1.2 LP data conditioning and sorting

The LP SNR and clutter power data for the current azimuth position are passed back from the collector following the LP scan. These go out to 460 km in range. The LP SNR is converted to signal power as $P_S = P_N 10^{(SNR/10)}$, because the SNR is stored in dB units. Since the TDWR range gates are sampled every 1 μ s, the number of gates per trip is given by $10^6 T$. Thus, at range gate k , the corresponding LP data gates for trip $i = 1, 2, 3 \dots$ are $k, k + 10^6 T, k + (2 \times 10^6 T) \dots$ etc. The number of trips available from the LP data is given by $n_{\text{trip}} = \text{floor}[(N_{\text{LP}} - k)/10^6 T] + 1$, where N_{LP} is the total number of LP range gates.

We now have the signal and clutter powers from trip $i = 1$ to n_{trip} , denoted as P_S^i and P_C^i . For $i > 1$, we add the clutter power back to the signal power, because any range-overlaid signal will include the clutter power. We then sort the n_{trip} signal powers in order of strength. However, if the signal powers are all very small (comparable to the noise power), then we want to designate the first trip as being the strongest. So if the strongest trip is not 1 and the strongest P_S^i is less than P_N , then assign the first trip to be the strongest and bump the rest down by one rank. Finally we designate the strongest trip as the “strong trip” and the second strongest trip as the “weak trip.”

3.3.2.1.3 Set dwell length

If the strong trip is 1, we set $M_{\text{DFT}} = M$; otherwise, we set $M_{\text{DFT}} = M - l_{\text{strong}} + 1$, where l_{strong} is the trip number of the strong trip. If weak-trip processing is not needed, then we want to utilize all the data points in this dwell. However, if weak-trip processing is needed, then for maximum coherence of the strong trip (in order to filter it out) we need to eliminate the pulses at the beginning of the dwell that do not contain signal from the strong trip. This boundary condition exists because we change the PRI on every dwell.

3.3.2.1.4 1st trip coherence

Cohere to 1st trip as in Section 3.3.1.1 for $l = l_{\text{strong}}$ to M to get M_{DFT} points.

3.3.2.1.5 Interference filter and overlay flag decision

Range-overlaid signal that is not cohered because of the pseudorandom pulse phase coding is decorrelated in the time domain. The interference filter can confuse some of the spiky points of this signal as interference and interpolate over them. This, in turn, can interfere with the phase-code processing for the weak trip. Therefore, we only apply the interference filter (Section 3.3.1.3) if the $l_{\text{strong}} = 1$ (i.e., no weak-trip processing will be performed).

If $l_{\text{strong}} > 1$, we set the overlay flag, which is an internal flag that is used in the collector during 2D processing.

3.3.2.1.6 DC power threshold decision

This process is identical to Section 3.3.1.6.

3.3.2.1.7 Data windowing

This process is identical to Section 3.3.1.7.

3.3.2.1.8 Clutter spectral width computation

This process is identical to Section 3.3.1.8.

3.3.2.1.9 Spectral GCF

This process is identical to Section 3.3.1.9. If the LP clutter power in the first trip, P_C^1 , is greater than the noise power, P_N , then the strong type is specified; otherwise, the weak type is called. This is another safeguard against taking out unnecessary clutter.

3.3.2.1.10 GCF decision

Phase-code processing is a filtering process in the spectral domain. It needs the signal from the trip to be filtered out to be coherent. Clutter filtering interferes with this process, because it changes the values of some spectral points, which leads to loss of information that degrades the coherence of the other-trip signal spread across the spectrum. The GCF is applied to the data coerced to the 1st trip. If $l_{\text{strong}} = 1$, then we do not need phase-code processing to filter out other-trip signals, so there is no need to worry about GCF/phase-code interference. The only criterion for using the clutter-filtered data or not is whether there was any clutter removed. If $l_{\text{strong}} > 1$, however, we have to make a decision whether it is better to use the clutter-filtered data or the unfiltered data for weak-trip processing.

The criterion that we use is: If ($l_{\text{strong}} = 1$ and $P_C > P_N$) or ($l_{\text{strong}} > 1$ and $P_S/P_C < \text{SCR_LIM}$ and $P_S^1/P_C^1 < \text{SCR_LIM}$) then use the clutter-filtered data. SCR_LIM is currently set to 10. Both the present and LP signal-to-clutter ratios are used because the short-PRI clutter can be contaminated by range aliasing whereas the LP clutter can be contaminated by velocity aliasing.

Note that there is no need to attempt to apply the GCF to out-of-trip signals. Any clutter present in those signals will be treated like weather signals and be filtered out by the phase-code processing.

3.3.2.1.11 Bad data flag decision #1

If both the GCF and phase-code processing for the weak-trip are to be applied, then the resulting base data estimates have the potential to be of poor quality. Here we set a bad data flag if the ratio of the

strong-trip signal to the 1st trip signal is greater than OSR_LIM (overlay-to-signal ratio limit). Currently OSR_LIM is set to 1.

3.3.2.1.12 Strong- or weak-trip processing decision

Here the process bifurcates for a while depending on whether the strong trip is the 1st trip or not. Sections 3.3.2.1.13 through 3.3.2.1.16 discuss the $l_{\text{strong}} = 1$ case, while Sections 3.3.2.1.17 through 3.3.2.1.21 detail the $l_{\text{strong}} > 1$ case.

3.3.2.1.13 Moments computation

The expressions for computing the zeroth moment, P , and the first moment, R_1 , are given in Equations 3-7 and 3-8, with $M = M_{\text{DFT}}$. The second moment is computed as

$$R_2 = \frac{\sum_{l=1}^{M_{\text{DFT}}-2} s_l^* s_{l+2}}{\sum_{l=1}^{M_{\text{DFT}}-2} a_l a_{l+2}} \quad . \quad (3-17)$$

3.3.2.1.14 Doppler computation

The standard pulse-pair estimator (Doviak and Zrnić 1993) is used to compute the raw radial velocity,

$$V_{\text{raw}} = \frac{\lambda \angle R_1}{4\pi T} \quad . \quad (3-18)$$

The spectral width is computed using the R_1/R_2 estimator (Doviak and Zrnić 1993),

$$W = \frac{\lambda}{2\sqrt{6}\pi T} \left| \ln \left| \frac{R_1}{R_2} \right| \right|^{\frac{1}{2}} \text{sgn} \left(\ln \left| \frac{R_1}{R_2} \right| \right) \quad . \quad (3-19)$$

We choose this estimator, because it is unbiased against white noise, which is helpful in the presence of range-overlaid signal. In principle, the P_S/R_1 method is a better estimator in the absence of range overlay, but switching between two estimators creates transition artifacts, so we stick with one estimator.

We also calculate, as a measure of the Doppler data quality, a quantity dubbed the signal quality index (SQI) by Vaisala,

$$SQI = \frac{|R_1|}{P} . \quad (3-20)$$

3.3.2.1.15 Bad data flag decision #2

Even though the 1st trip is the strongest here, there is a possibility that the 1st trip signal power is less than the sum of the signals from the other trips. In this case, the bad data flag is set.

3.3.2.1.16 Clutter suppression extension

This process is identical to Section 3.3.1.11.

3.3.2.1.17 Strong-trip coherence

We are now in the $l_{\text{strong}} > 1$ branch in Figure 3-8. At this point the I&Q data are still cohered to the 1st trip. To cohere this time series to the strong trip, we do

$$s_l = s_l e^{j\phi_l - \phi_{l_{\text{strong}}+1}} . \quad (3-21)$$

3.3.2.1.18 Overlay flag cancellation

If the mean signal power in the time series at this point is less than or equal to the noise power, then cancel the overlay flag.

3.3.2.1.19 Data windowing for weak-trip processing

If the data stream here is unfiltered, then it has not been windowed. In that case, apply windowing as per Section 3.3.1.7.

3.3.2.1.20 Weak-trip processing

This process coheres and filters out the unwanted strong-trip signal in the spectral domain, and computes the weak-trip (1st trip) power, velocity, spectral width, and SQI. Conceptually it is similar to Vaisala's RVP9 phase-code processing, but details differ. The basic idea is to determine which spectral bins have the characteristics of a coherent signal and remove them, then transform back to the time domain and recohere to the 1st trip. The following steps describe the algorithm.

1. DFT the windowed time series data, $a_l s_l$, to get S_m .
2. Estimate the noise level, p_{SN} (see Section 3.3.1.12), of the spectrum $|S_m/M_{\text{DFT}}|^2$. The sorted indices (arranged in increasing order of power per spectral bin) are also obtained.

3. Assign $P = p_{SN}M_{DFT}$.
4. Set a spectral filter mask for bins that exceed p_{SN} using the sorted indices. Start down the sorted list from strongest, but do not go beyond $2/3$ of M_{DFT} .
5. Set $S_m = 0$ at indices where the spectral filter mask has “true” values.
6. If there are any spectral points that are singular, i.e., a zero flanked by nonzero bins, or a nonzero point flanked by zero bins, set these ambiguous points to amplitude equal to p_{SN} . Retain their original phase values.
7. Inverse DFT to the time domain.
8. Cohere to the weak (1st) trip. This is accomplished by the inverse operation of Equation 3-21 (i.e., reverse the sign of the exponential multiplier).
9. Compute V_{raw} , W , and SQI (see Section 3.3.2.1.14).

3.3.2.1.21 Bad data flag decision #3

Phase-code processing can only effectively filter out one trip. If there is significant signal power in trips other than the strong trip that was filtered out in Section 3.3.2.1.20 (i.e., $P_S \leq$ sum of unfiltered trip signal powers), then the bad data flag is set.

3.3.2.1.22 SNR computation

This process is identical to Section 3.3.1.13.

3.3.2.1.23 Reflectivity computation

This process is identical to Section 3.3.1.14.

3.3.2.2 Phase-code processing without LP data

In this mode, the phase-code processing does not have a priori range-overlay information from the LP data. Therefore, the signal coherence (using SQI) is computed for each trip and is used in lieu of the LP signal power to select the “strong” and “weak” trips. Another difference from phase-code processing with LP data is that a CREM may not be available. (CREMs should always be available for the surface scan, but may or may not be present for any higher elevation tilts.) Without a CREM for editing, antenna sidelobe clutter may contaminate the base data estimates. Sidelobe clutter, unlike normal ground clutter, can have substantial non-zero Doppler shift, which can leave a strong residue after the standard clutter filtering. The apparent radial velocity results from the feed horn of the antenna being offset from the axis of rotation, which introduces a radial motion between targets in the azimuthal sidelobes and the feed horn

(Rinehart 1991). Thus, we need to widen the clutter filter width based on certain criteria such as elevation angle, range to target, and antenna rotation rate.

The flowchart for phase-code processing without LP data is shown in Figures 3-9 and 3-10. This process is repeated for each range gate's worth of data. The number of pulses is dependent on the PRI and the antenna rotation rate, but it can vary slightly from radial to radial. I&Q data, transmitted pulse phase angles, PRI, antenna rotation rate, antenna elevation angle, radar parameters, and noise power are passed in from the IQM. SNR, Z , V_{raw} , W , SQI, and internal flags are output to the collector. Velocity dealiasing takes place in the collector.

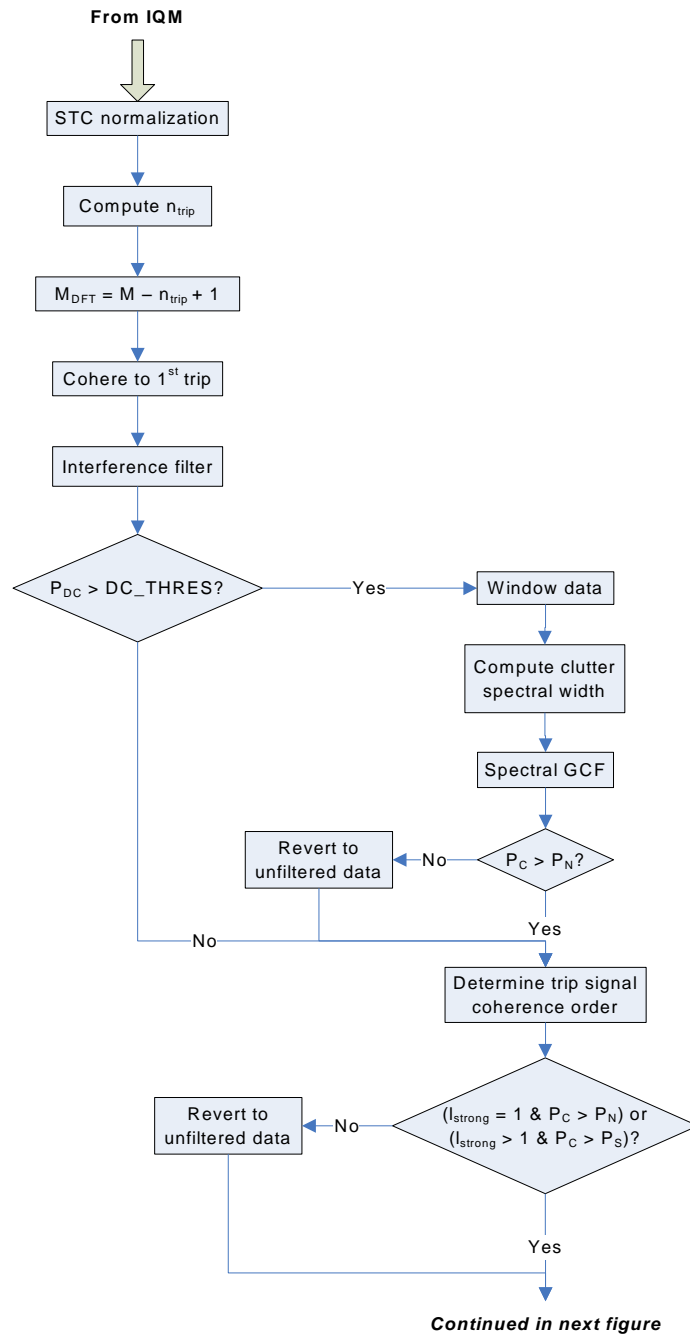


Figure 3-9. Flow diagram for phase-code processing without LP data, part 1.

Continued from previous figure

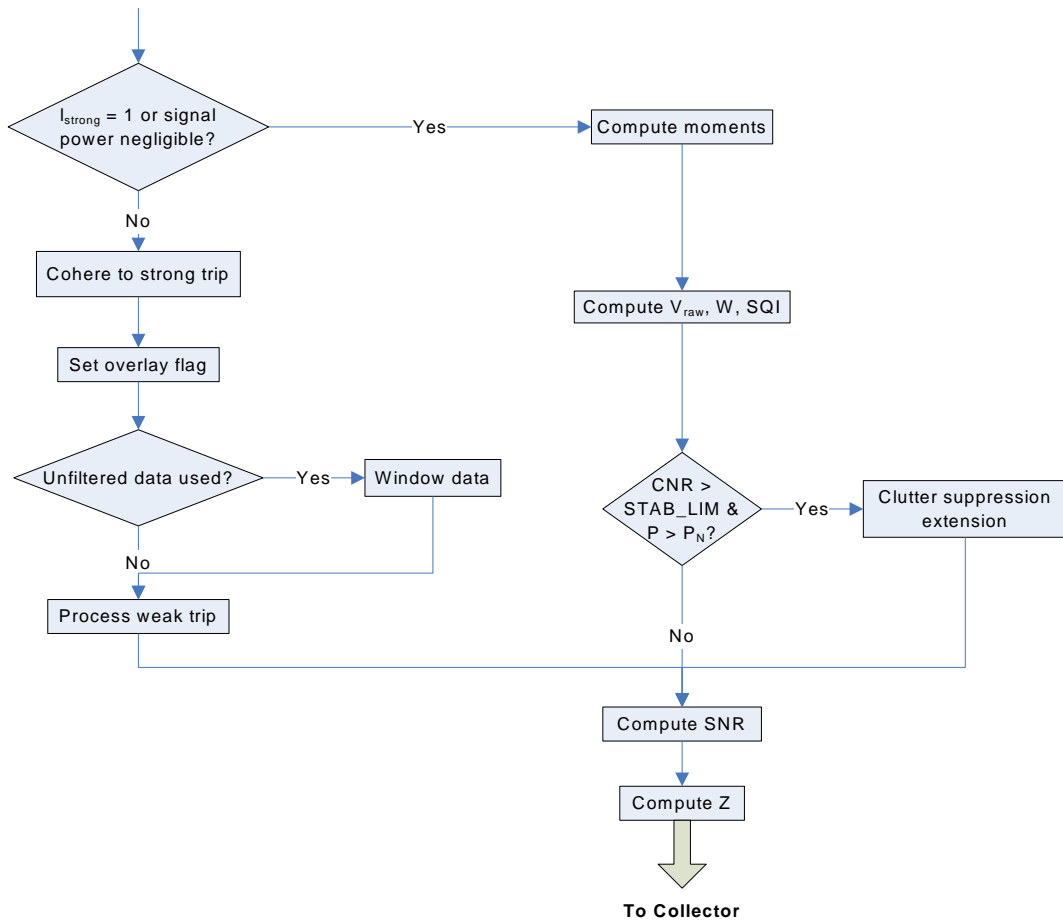


Figure 3-10. Flow diagram for phase-code processing without LP data, part 2.

3.3.2.2.1 STC normalization

This process is identical to Section 3.3.1.2.

3.3.2.2.2 Computation of n_{trip}

To proceed with the phase-code processing, we need to compute the number of trips back from which weather signal can be returned. This varies with the antenna beam elevation angle and PRI. We use the expression, $n_{\text{trip}} = 1 + \max\{0, \text{floor}[(r_{\text{ceil}} - (k - 1)\Delta r)/(10^6 T \Delta r)]\}$, where $\Delta r = 150$ m is the range sampling interval, and the range to the weather ceiling, in meters, is given by $r_{\text{ceil}} = [(\tan^2 \theta_{\text{el}} + 4CH)^{1/2} - \tan \theta_{\text{el}}]/(2C \cos \theta_{\text{el}})$, where θ_{el} is the elevation angle in degrees, $C = 5.887 \times 10^{-8}$, and H is the maximum

expected height of weather, which we choose to be 18,288 m (60,000 ft). This expression incorporates the standard 4/3-Earth-radius approximation for atmospheric refraction effects.

3.3.2.2.3 Set dwell length

Without a priori knowledge of the trip signal strength ordering, we need to assume the worst-case scenario in which the farthest trip back would be phase-code processed. Therefore, we set $M_{\text{DFT}} = M - n_{\text{trip}} + 1$.

3.3.2.2.4 1st trip coherence

Cohere to 1st trip as in Section 3.3.1.1 for $l = n_{\text{trip}}$ to M to get M_{DFT} points.

3.3.2.2.5 Interference filter

This process is identical to Section 3.3.1.3. Without a priori knowledge of the presence or absence of range-overlaid signals, we take interference filtering to be a priority over withholding it for possible undesired removal of out-of-trip information.

3.3.2.2.6 DC power threshold decision

This process is identical to Section 3.3.1.6.

3.3.2.2.7 Data windowing

This process is identical to Section 3.3.1.7.

3.3.2.2.8 Clutter spectral width computation

As explained in Section 3.3.2.2, for scan elevations with no CREM, we attempt to filter out sidelobe clutter. The maximum Doppler shift that antenna sidelobe clutter can have is given by (Rinehart 1991)

$$V_{sc} = \frac{2\pi d \gamma \cos \theta_{el}}{360^\circ} \text{ m s}^{-1}, \quad (3-22)$$

where $d = 5.1$ m is the distance from the axis of antenna rotation to the feed horn (Michelson et al. 1990). Taking into account positive and negative excursions, we can set the widened clutter spectral width to be (see Equation 3-15)

$$\sigma_{sc} = \sqrt{\sigma_c^2 + (V_{sc})^2} \text{ m s}^{-1}. \quad (3-23)$$

For range gates beyond SIDELOBE_CLUTTER_GATE_LIM, this clutter spectral width widening is not applied (σ_{SC} is set to zero). Currently SIDELOBE_CLUTTER_GATE_LIM is set to 200.

3.3.2.2.9 Spectral GCF

This process is identical to Section 3.3.1.9. The strong type GCF is always specified.

3.3.2.2.10 GCF decision #1

At this point if $P_C > P_N$, then the clutter-filtered time series is used for further processing. Otherwise, the original unfiltered and unwindowed data are passed on to subsequent processing.

3.3.2.2.11 Trip signal coherence ordering

Since we do not have LP data for determining the order of the various trip signal strengths, we do the next best thing: Compute the SQI values (Equation 3-20) after cohering to each trip, up to n_{trip} . The presumption is that the higher the SQI, the more coherent signal is present in the corresponding trip. Then we designate the trip with the highest SQI to be the “strong” trip and the trip with the second highest SQI to be the “weak” trip for further processing.

3.3.2.2.12 GCF decision #2

Now that we have the relative trip signal strength ordering, we decide again whether or not to use the GCF. If ($l_{\text{strong}} = 1$ and $P_C > P_N$) or ($l_{\text{strong}} > 1$ and $P_C > P_S$) then use the clutter-filtered data; otherwise, revert to the unfiltered data.

3.3.2.2.13 Strong- or weak-trip processing decision

Here the process bifurcates for a while depending on whether the strong trip is the 1st trip or not. Sections 3.3.2.2.14 through 3.3.2.2.16 discuss the $l_{\text{strong}} = 1$ case, while Sections 3.3.2.2.17 through 3.3.2.2.20 detail the $l_{\text{strong}} > 1$ case. Also, if the signal power is negligible ($P_S \leq P_N$), the $l_{\text{strong}} = 1$ case is chosen.

3.3.2.2.14 Moments computation

This process is identical to Section 3.3.2.1.13.

3.3.2.2.15 Doppler computation

This process is identical to Section 3.3.2.1.14.

3.3.2.2.16 Clutter suppression extension

This process is identical to Section 3.3.1.11.

3.3.2.2.17 Strong-trip coherence

This process is identical to Section 3.3.2.1.17.

3.3.2.2.18 Set overlay flag

If we are here it means that there is at least some overlay signal power, so set the overlay flag.

3.3.2.2.19 Data windowing for weak-trip processing

This process is identical to Section 3.3.2.1.19.

3.3.2.2.20 Weak-trip processing

This process is identical to Section 3.3.2.1.20.

3.3.2.2.21 SNR computation

This process is identical to Section 3.3.1.13.

3.3.2.2.22 Reflectivity computation

This process is identical to Section 3.3.1.14.

3.3.3 Multi-PRI Processing

Another approach to RV ambiguity mitigation is through MP transmission and processing. In this scheme, more than one PRI is used within a dwell. Pseudorandom phase is encoded on transmission to render out-of-trip return signals incoherent upon coherence to the first trip. Since the range-aliasing interval changes with the PRI, signals transmitted at varying PRIs will return with potential overlays from different range intervals. If the backscattered signal power distribution in range is known from an initial long-PRI scan, then it is a simple matter to flag all pulses at any given range gate that are contaminated by out-of-trip signals. Only the remaining “clean” pulses are processed for the moments data.

Currently there are three multi-PRI sets used by the mode selection algorithm. All are in a block-staggered pattern, i.e., M_p pulses transmitted consecutively for N_p PRIs. There are two sets with $N_p = 4$: $T_p = 600, 670, 740, 810 \mu\text{s}$, and $T_p = 698, 798, 898, 998 \mu\text{s}$. There is one set with $N_p = 8$: $T_p = 600, 648, 696, 744, 792, 840, 888, 936 \mu\text{s}$. The values of M_p are pre-selected using the expression

$$M_p = \text{floor} \left(\frac{\frac{1}{\gamma} - \varepsilon_t}{\sum_{p=1}^{N_p} T_p} \right), \quad (3-24)$$

where ε_t is a small value (currently set to 1 ms) that provides a margin to prevent fluctuations in the antenna rotation rate from cutting off the transmission of the full PRI sequence within the 1° dwell. In cases where the fluctuation in γ generates extra pulses at the end of the dwell, the signal processing ignores these points. Therefore, the number of I&Q data points processed per dwell for the MP mode is $M = M_p N_p$.

This mode is only used in the adaptive surface tilt. The flowchart for MP processing is shown in Figures 3-11 and 3-12. This process is repeated for each range gate's worth of data. The number of pulses is dependent on the PRIs and the antenna rotation rate, but it can vary slightly from radial to radial. I&Q data, transmitted pulse phase angles, PRIs, antenna rotation rate, radar parameters, and noise power are passed in from the IQM. LP SNR and clutter power are passed back from the collector. SNR, Z, V, V_{raw} , V_{alt} , W, SQI, and internal flags are output to the collector.

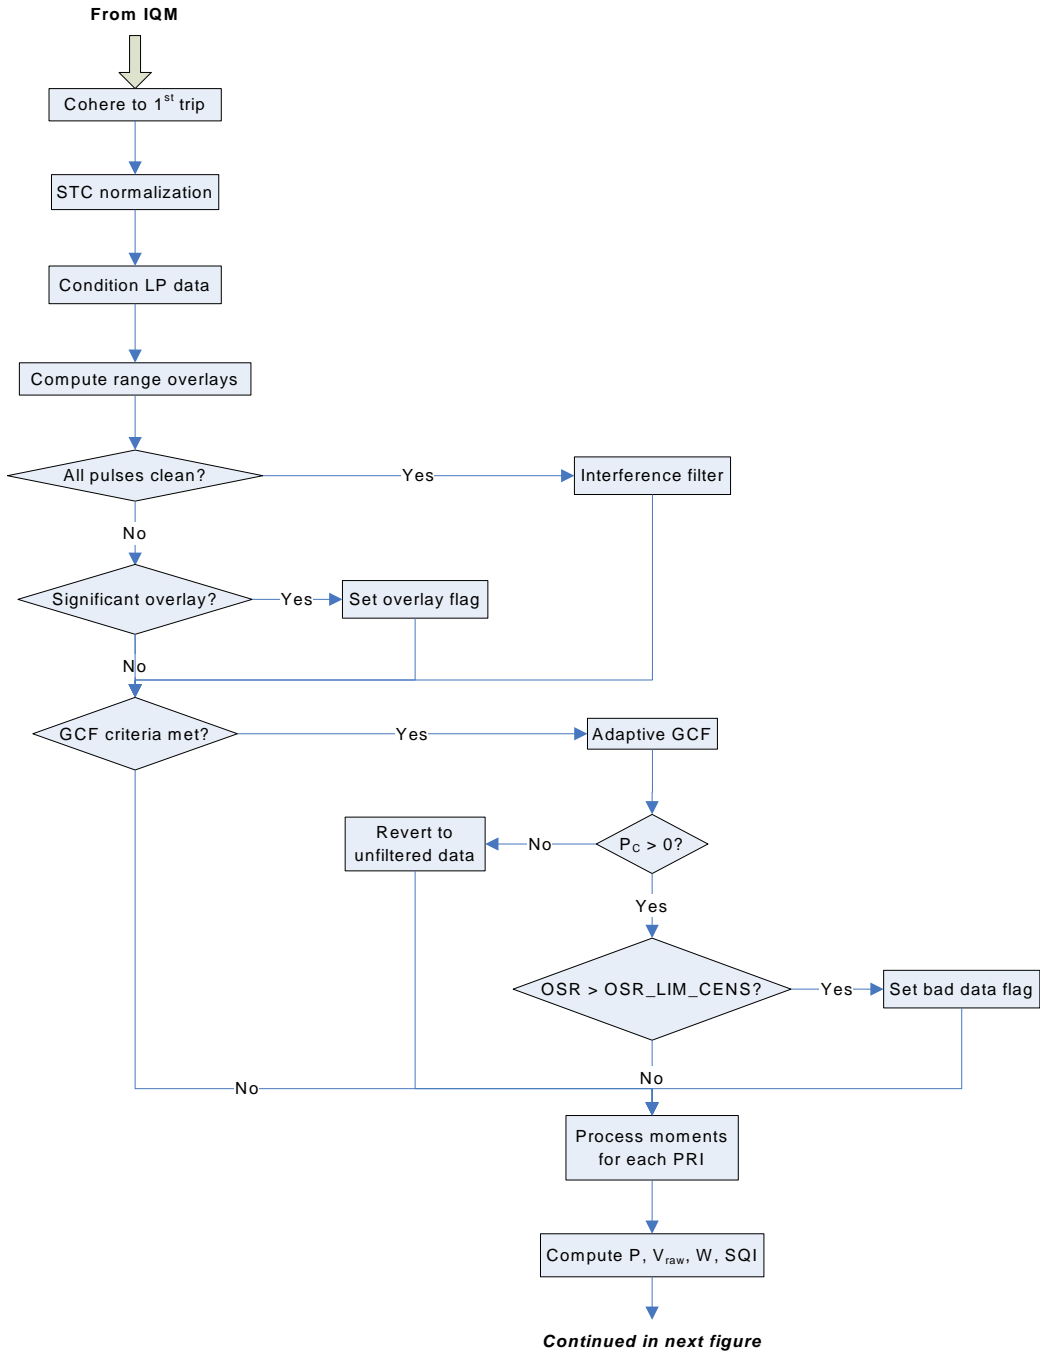


Figure 3-11. Flow diagram for multi-PRI processing, part 1.

Continued from previous figure

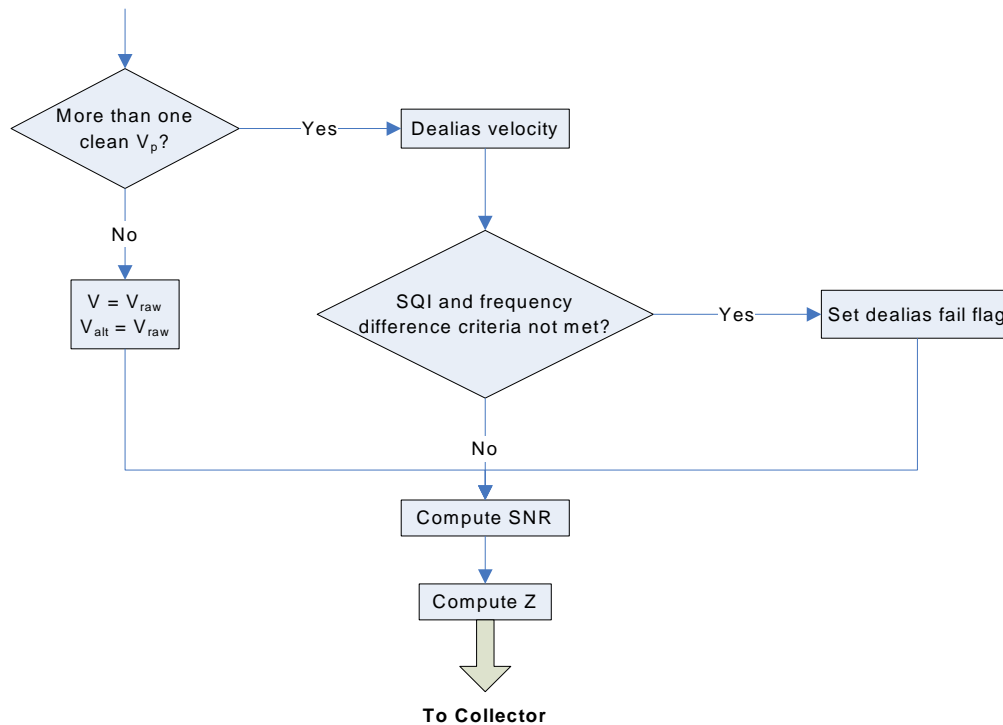


Figure 3-12. Flow diagram for multi-PRI processing, part 2.

3.3.3.1 1st trip coherence

This process is identical to Section 3.3.1.1.

3.3.3.2 STC normalization

This process is identical to Section 3.3.1.2.

3.3.3.3 LP data conditioning

The LP SNR and clutter power data for the current azimuth position are passed back from the collector following the LP scan. These go out to 460 km in range. The LP SNR is converted to signal power as $P_S = P_N 10^{(SNR/10)}$, since the SNR is stored in dB units. For range gates $k = 1$ to N_GATES_BD we add the clutter power back to the signal power, because any range-overlaid signal will include the clutter power.

3.3.3.4 Range overlay computation

Suppose we have a PRI sequence T_0, T_1, T_2 , etc. Then the sum of all the out-of-trip signals aliasing into the first-trip (overlay signal power) at range gate k and pulse index l is computed from the following equations,

$$P_{kl}^O = \sum_{i=1}^{n_{kl}^{\text{trip}}-1} P_S(r_k + r_i') \quad (3-25)$$

$$r_i' = \frac{c}{2} \sum_{l=1}^i T_{-l} \quad (3-26)$$

where r_k is the range to gate k and $n_{kl}^{\text{trip}} = \text{floor}[(N_{\text{LP}} - k)/10^6 T_l] + 1$. This received pulse is marked clean if $P_{kl}^O < (n_{kl}^{\text{trip}} - 1)P_n$ or $P_{Sk} > \text{OVERLAY_LIM_FAC} * P_{kl}^O$, where P_{Sl} is first-trip signal at pulse index l . The first condition corresponds to negligible range aliasing while the second condition allows some overlay power as long as the ratio of first-trip to out-of-trip powers exceeds OVERLAY_LIM_FAC. The value of OVERLAY_LIM_FAC is determined as follows.

Note that we apply a pseudorandom phase code on transmission, which means that, in theory, out-of-trip signal will appear as white noise when cohered to the first trip phase. Thus, the signal-to-overlay ratio (SOR) can be thought of as a type of SNR. If we posit that the SOR must be at least 0 dB for minimum moments-estimate accuracy, then OVERLAY_LIM_FAC = 1. (For reference, the legacy operational TDWR acceptance threshold for velocity is SNR = 0.5 dB). However, the moments estimate will be computed from an average of all the clean pulses corresponding to each PRI. The averaging will effectively decrease the estimate error by the inverse square root of the number of points averaged. If we conservatively assume that only half the pulses in a PRI set will be clean, then we arrive at OVERLAY_LIM_FAC = $(M_p/2)^{1/2}$. This is the expression currently used. Since estimate error also depends on the normalized Doppler spectral width, σ_v/v_a , there is really no single value for OVERLAY_LIM_FAC that is optimal for all moments estimates under all circumstances, but tests using simulated and real data show this expression to yield good results.

3.3.3.5 Interference filter

In the MP mode, the signal level can abruptly change going from one pulse to the next. This happens because the pattern of PRIs preceding each pulse is different; thus, the amount of range-overlaid signal can be different from pulse to pulse. The interference filter can confuse the resulting spiky data as interference and interpolate over them. This, in turn, can interfere with the subsequent signal processing. Therefore, we only apply the interference filter (Section 3.3.1.3) if all pulses in the dwell are deemed to be clean of overlays (as determined in Section 3.3.3.4).

3.3.3.6 Overlay flag decision

If the mean overlay signal,

$$P_O = \sum_{l=1}^M P_l^O \quad , \quad (3-27)$$

is greater than P_N , then set the overlay flag.

3.3.3.7 GCF decision #1

As with phase-code processing, clutter filtering interferes with range-overlay protection. Applying a GCF coherently across all PRI pulse sets convolves information from different pulses and destroys the independence of range aliasing between PRI sets. In other words, even if only one PRI set is contaminated by an overlaid signal, application of the GCF will mix some of this unwanted signal into the time series of all the other PRI sets. Therefore, the GCF should only be applied when absolutely necessary.

The GCF is tried if the following conditions are met:

1. DC power is present (see Section 3.3.1.6; replace T by the mean of all T s).
2. ($P_C^{-1} > P_N$ and all pulses are clean) or $DC_CLUT_FAC * N_p^2 * P_O / P_C^{-1} < OSR_LIM$.

DC_CLUT_FAC is currently set to 0.125.

3.3.3.8 Adaptive GCF

The operational application of MP techniques has been hampered, in large part, by the difficulty of clutter filtering. Because nonuniform sampling aliases power from nonzero Doppler frequencies to the ground-clutter band around zero, the clutter filtering also removes power from the aliased frequencies and distorts the phase response at those frequencies. The phase distortion in turn leads to degradation of velocity estimates. To solve this problem, a finite impulse response (FIR) clutter filter design yielding an excellent balance of magnitude response and phase linearity was introduced by Chornoboy (1993) for block-staggered PRI pulse trains.

A drawback of clutter filtering in the time domain is the difficulty of determining the optimal clutter suppression level for a given range-azimuth cell in real time. Clutter residue maps are sometimes used to select a filter from a range of suppression levels. However, ground clutter is not static over time. Vegetation changes with the seasons, construction and demolition alter buildings, roads, bridges, and towers, weather transfers water and ice to and from structures, and winds sway signs and rustle leaves. AP can also introduce ground clutter to range gates that are usually clutter-free. Therefore, ideally, the

clutter suppression level should be determined without resorting to a preset clutter residue map. Such a technique was developed by Cho and Chornoboy (2005) and has been applied to MP processing for the TDWR.

The idea is to use clutter power estimated from the initial LP scan (Section 3.3.1.9) to select the filter with the appropriate suppression level. The suppression levels available are 20, 35, 50, and 65 dB. (For details on the filter coefficient sets, see Appendix A.) The clutter filter level is selected according to the estimated clutter-to-noise ratio (CNR). For $\text{CNR} < 0$ dB, no clutter filter is applied. For $0 \text{ dB} \leq \text{CNR} < 20$ dB, the 20-dB filter is applied. For $20 \text{ dB} \leq \text{CNR} < 35$ dB, the 35-dB filter is applied. For $35 \text{ dB} \leq \text{CNR} < 50$ dB, the 50-dB filter is applied. For $\text{CNR} \geq 50$ dB, the 65-dB filter is applied.

3.3.3.9 GCF decision #2

If $P_C > 0$, then the clutter-filtered data set is used for further processing; otherwise, the unfiltered data are passed on.

3.3.3.10 Bad data flag decision #1

If the clutter-filtered data stream was selected for processing, then protection against overlaid signal will be limited. So set the bad data flag if $P_O/P_S^1 > \text{OSR_LIM_CENS}$. OSR_LIM_CENS is currently set to 10.

3.3.3.11 Moments processing for each PRI

This process calculates the first three moments for each PRI pulse set. They are computed using all pulses and using only clean pulses. The DC component is subtracted first from the latter set if needed.

1. Set the DC-remove mode to on if $P_C = 0$ and $P_C^1 > 0$ and DC power is present (see Section 3.3.3.7, step 1).
2. Set bad flag if $P_C^1/P_S^1 > \text{DC_CLUT_FAC_CENS} * N_p^2$, i.e., DC clutter suppression will not be enough to yield good moment estimates. Currently, DC_CLUT_FAC_CENS is set to 1.25.
3. Compute the moments for each PRI pulse set using all pulses (P_p , R_{1p} , and R_{2p}) and using only clean pulses (P_p^C , R_{1p}^C , and R_{2p}^C) (Equations 3-7, 3-8, and 3-17). Subtract the DC components from the “clean” moments if the DC-remove mode is on and there are at least MIN_PTS_DC clean pulses in the PRI set. MIN_PTS_DC is currently set to 2. Also compute the powers using the same pulses used for the first-moment computations, i.e., including the last pulse of the previous PRI set. These values (P_p^R and P_p^{CR}) are used for the SQI calculations.
4. Compute the number of clean pulse-pairs per PRI pulse set, C_p , in computing R_{1p}^C . If the number of PRI sets with $C_p \geq M_p/2$ is less than 2, then set the velocity dealias fail flag.

5. If the DC-remove mode is on, compute clutter power removed.
6. Compute SQI for each PRI set: $|R_{1p}|/P_p^R$.

3.3.3.12 P , V_{raw} , W , and SQI computation

These quantities are averaged over the PRI sets.

1. If at least one P_p^C exists, take the median to get P ; otherwise, $P = \min(P_p)$ and set bad data flag.
2. If at least one R_{1p}^C exists, take the median of V_p (given by Equation 3-18 for each PRI) to get V_{raw} ; otherwise, use R_{1p} with the highest SQI to compute V_{raw} .
3. If at least one pair of R_{1p}^C and R_{2p}^C exists, take the median of W_p (given by Equation 3-19 for each PRI); otherwise, use R_{1p} and R_{2p} with the highest SQI to compute W .
4. If at least one R_{1p}^C exists, $\text{SQI} = |\text{mean}(R_{1p}^C)|/\text{mean}(P_p^{CR})$; otherwise, take the maximum value of the SQIs computed in Section 3.3.3.11, step 5.

3.3.3.13 Velocity dealias decision

If there are two or more sets of clean PRI velocity sets, then perform velocity dealiasing; otherwise, set $V = V_{\text{raw}}$ and $V_{\text{alt}} = V_{\text{raw}}$. See next section for meaning of V_{alt} .

3.3.3.14 Velocity dealiasing

Our velocity dealiasing algorithm is based on the unfolded velocity matching technique (Trunk and Brockett 1993). In this method, for each PRI velocity estimate, all possible unfolded velocities are computed up to $\pm V_{\text{max}}$. Then all the velocity values are sorted from smallest to largest and the average squared error is computed in a sliding window of length N_p that is incremented across the entire list. The median value in the window with the smallest error (the “best match”) is the dealiased velocity. One of the advantages of this algorithm is that V_{max} can be set to any value. In other words, the trade off between the maximum speed that can be dealiased and dealiasing error can be adjusted in a smooth, continuous fashion (see Cho 2005). Decreasing V_{max} increases the dealiasing success rate as long as most of the velocity distribution lies within $\pm V_{\text{max}}$. Currently, V_{max} is set to 48 m s^{-1} for MP processing.

We made some changes to this basic algorithm for improved performance. The entire procedure is as follows.

1. For each PRI velocity estimate, V_p , add and subtract v_{ap} , the unambiguous velocity for that PRI, to get the unfolded velocity set in the interval $\pm V_{\text{max}}$.
2. Sort all the values in the unfolded velocity sets from smallest to largest to get V_q^{US} .

3. Compute the weighted median absolute deviation (WMAD) over a window of length N_p incremented across all the sorted values. The WMAD for index q in the sorted list is given by

$$\frac{\sum_{i=q}^{q+N_p-1} |V_i^{US} - V_q^{WM}| w_i}{\sum_{i=q}^{q+N_p-1} w_i}, \quad (3-28)$$

where the weights, w_i , are given by the SQI values associated with the corresponding V_i^{US} , and V_q^{WM} is the weighted median of V_i^{US} for $i = q$ to $q + N_p - 1$, again weighted by w_i . For an explanation of the weighted median, see Arce (1998). The dealiased velocity, V , is assigned the V_q^{WM} corresponding to the smallest WMAD. If the smallest WMAD is greater than the `VD_DEV_LIM`, then the dealias-fail flag is set. Currently, `VD_DEV_LIM` is set to 2.5 m s^{-1} . The V_q^{WM} corresponding to the next smallest WMAD, if it exists, is assigned to the alternative dealiased velocity, V_{alt} . This quantity is used by the 2D dealias correction filter in the collector.

3.3.3.15 Velocity dealias fail flag decision

Even if there were two or more “clean” V_p for velocity dealiasing, if their SQIs were too low or the minimum frequency separation between the PRIs were too small for reliable dealiasing, then we set the dealias fail flag. The minimum frequency difference, `FREQ_DIFF_LIM`, is currently set to 250 Hz. The SQI criterion is that if the number of clean PRI sets with SQIs \geq `SQI_OK_LIM` is less than or equal to half the number of clean PRI sets, then set the dealias fail flag. `SQI_OK_LIM` is currently set to 0.4.

3.3.3.16 SNR computation

This process is identical to Section 3.3.1.13.

3.3.3.17 Reflectivity computation

This process is identical to Section 3.3.1.14.

3.3.4 Staggered-PRI Processing

At high elevation scan angles, range aliasing ceases to be a concern, because the extent of weather echoes is contained entirely within the first trip. Ground clutter is also much less of a problem, although sidelobe clutter can be very strong in some locations. Velocity aliasing is, perhaps, even more of a concern compared to low-elevation scans, because wind speed tends to increase with height. Taking all these factors into account, we decided that the best transmission strategy would be the staggered-PRI mode.

In SP mode, pulses are transmitted with alternating PRIs. The main advantage over a block-staggered pattern is that the increased separation between pulse pairs increases their statistical independence, which leads to lower estimate variance. The pattern is also invariant with respect to the antenna rotation rate, whereas with the block-staggered pattern the number of pulses per PRI must be changed to fit the dwell.

Standard pulse-pair processing is applied separately to the two PRI sets and the resulting pair of velocity estimates is dealiased with the UVM method. If clutter is present, the time series is split into two evenly spaced sequences (Meymaris et al. 2009) and the same spectral GCF as used in the LP and DP modes is applied. The resulting filtered data are recombined in the time domain before pulse-pair processing.

At the highest elevation angles, with the required range coverage decreasing below 90 km, the PRIs can be reduced correspondingly. Transmitter limitations keep the shorter PRI to be no less than 518 μ s, however.

The flowchart for SP processing is shown in Figure 3-13. This process is repeated for each range gate's worth of data. The number of pulses is dependent on the PRIs and the antenna rotation rate, but it can vary slightly from radial to radial. I&Q data, transmitted pulse phase angles, PRIs, antenna rotation rate, radar parameters, and noise power are passed in from the IQM. SNR, Z , V , V_{raw} , V_{alt} , W , SQI, and internal flags are output to the collector.

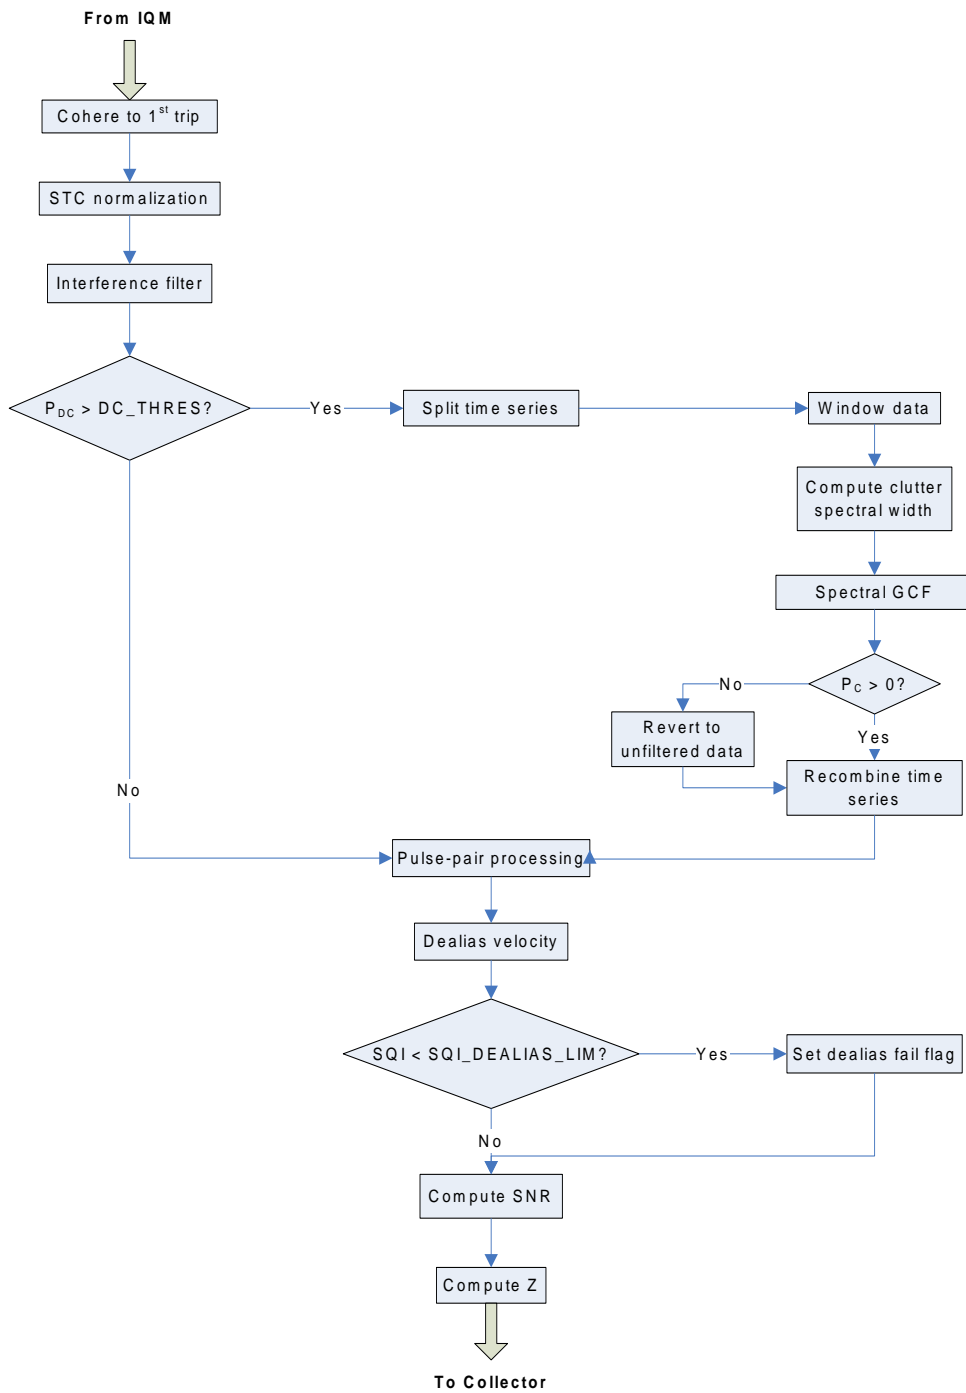


Figure 3-13. Flow diagram for staggered-PRI processing.

3.3.4.1 1st trip coherence

This process is identical to Section 3.3.1.1.

3.3.4.2 STC normalization

This process is identical to Section 3.3.1.2.

3.3.4.3 Interference filter

This process is identical to Section 3.3.1.3.

3.3.4.4 DC power threshold decision

The criterion used for decision is the one given in Section 3.3.1.6 plus the clutter detection result from Section 3.2.3. If the answer from the former criterion is yes (try GCF) and $switchGCF = 1$, go to the next process; otherwise, skip to Section 3.3.4.11.

3.3.4.5 Split time series

The time series is split into two sequences—one with odd time indices, the other with even time indices (Figure 2-5). Both subseries now have evenly spaced samples with an effective PRI that is the sum of the short and long PRIs. From this point on to 3.3.4.9, the processes described are applied to both subseries.

3.3.4.6 Data windowing

This process is identical to Section 3.3.1.7.

3.3.4.7 Clutter spectral width computation

This process is identical to Section 3.3.2.2.8.

3.3.4.8 Spectral GCF

This process is identical to Section 3.3.1.9. The strong type GCF is always specified.

3.3.4.9 GCF decision

If $P_C > 0$, revert to the original no-GCF time series.

3.3.4.10 Recombine time series

Recombine the two evenly spaced time series into the original staggered sequential order (see Section 3.3.4.5), but retain a copy of the split time series for the next process.

3.3.4.11 Pulse-pair processing

To compute the power, use Equation 3-7. However, for velocity we wish to make the calculations separately for each PRI. So calculate R_1 for each PRI (where the summation indices in Equation 3-8 is odd for one PRI and even for the other PRI) and use Equation 3-18 to compute the velocity (V_p) for each PRI; the mean is assigned to V_{raw} . For the SQIs, divide the absolute value of the R_1 s by the R_0 s calculated for each PRI (using the same points used for the R_1 summation); the overall SQI is computed from the absolute value of the mean of the R_1 s divided by the power. For spectral width, compute R_1 for each split time series, average them, and then calculate W using Equation 3-6. Using the longer time interval somewhat degrades the dynamic range of W , but improves consistency because a single PRI value is used.

3.3.4.12 Velocity dealiasing

As in Section 3.3.3.14, the UVM technique is used to dealias the velocity. However, in this case, we use it to select the correct unfolding, but assign the unfolded velocity from the V_p with the higher SQI to V (instead of averaging the two unfolded velocities). The maximum unfolded velocity, V_{max} , is set to three times the smaller of the two Nyquist velocities.

3.3.4.13 Velocity dealias fail flag decision

If the spectral processing route was taken, set the velocity dealias fail flag if the overall SQI was less than `SQI_DEALIAS_LIM`. If the pulse-pair processing route was taken, set the velocity dealias fail flag if either of the individual SQIs was less than `SQI_DEALIAS_LIM`. `SQI_DEALIAS_LIM` is currently set to 0.4.

3.3.4.14 SNR computation

This process is identical to Section 3.3.1.13.

3.3.4.15 Reflectivity computation

This process is identical to Section 3.3.1.14.

3.4 COLLECTOR PROCESSES

Whereas the data were processed on a gate-by-gate basis in the IQS, processes that require contextual information in the range or azimuth dimensions are conducted in the collector. As such, a

number of dwells (radials) worth of data is buffered in the collector. Currently, the buffer is set for 4 dwells. If more complex 2D processing is implemented in the future, the buffer size could expand as well.

362 radials are processed in a given scan. Data from the first and last radials are only used internally in the collector for processing—they are not output to the RPG. As can be seen in Figure 3-2, the DP mode requires the most processing in the collector; this is because the PRI is changed on every radial and velocity dealiasing must be performed between radials.

The radial buffering and processing scheme is illustrated in Table 3-1. First, for every new radial that comes into the buffer, point target filtering and CREM flagging (PC) are done. Next, for DP radials, interrational velocity dealiasing (VD) is conducted. This process is delayed by a radial, because at least two consecutive radials are required. (Three consecutive radials are better, because the velocity dealiasing algorithm can then choose the neighbor with the better SQI; this is possible except on the first and last radials of the scan.) Then the 2D data quality filter (2DF) is applied using three consecutive radials and the resulting base data for the middle radial is output (O). When the last (362nd) radial is reached, the 2DF and O are repeated for the penultimate radial to yield a total of 360 radials for the scan.

TABLE 3-1
Radial Buffering and Processing Scheme

Radial	Radial Buffer Index			
	<i>i</i> - 3	<i>i</i> - 2	<i>i</i> - 1	<i>i</i>
1				PC
2			PC VD	PC
3		PC VD	PC VD	PC
4	PC VD	PC VD 2DF O	PC VD	PC
5	PC VD 2DF	PC VD 2DF O	PC VD	PC
⋮	⋮	⋮	⋮	⋮
362	PC VD 2DF	PC VD 2DF O	PC VD 2DF O	PC VD

In the legacy system, the LP base data were output at 150-m resolution for the first 900 range gates and at 300-m resolution beyond that point. We will maintain this scheme by averaging the LP results every two gates beyond the 900th gate for output to the RPG. For internal processing, however, the 150-m resolution will be maintained at all ranges.

3.4.1 Point Target Filter

The point target filter algorithm is unchanged from the legacy system. The procedure is given in Section 8.5.1 of Raytheon (1992). The only difference is that in the legacy algorithm, the nearest-neighbor interpolation technique was used. We also use nearest-neighbor for the velocity fields (in case of aliasing) and for the internal flags, but we use linear interpolation for all the other quantities.

3.4.2 CREM Editing

Currently, CREMs are generated and stored for the lowest few elevation scans (the exact number being site dependent), which means that data produced by SP processing will not be CREM edited. There is a limited amount of memory allocated for CREMs. As in the legacy system, CREMs only go out to 70 km (467 range gates). It is possible that, in the future, high-elevation CREM editing will be implemented at sites where sidelobe clutter is severe, which would require changes in the CREM memory allocation. Also, at this time, CREM editing is usually not performed on the LP scan. However, the capability to do so is present.

Editing based on CREMs is straightforward. The CREM edit flag is set if the difference between the computed reflectivity and the clutter residue reflectivity is less than CREM_BREAKTHROUGH_Z, which is currently set to 8 dBZ, a value typically used in the legacy system.

For the modes that use the 2D data quality filter, CREM edit flags are generated first so that the information can be used in the filter. However, the CREM edit flags are regenerated after the SNR cutoff (in the base data flagging stage) because some reflectivity values are changed due to the cutoff.

3.4.3 Interradial Velocity Dealiasing

This process is only applicable to the data produced by the DP mode, which is used on at least two consecutive radials. (Isolated DP radials are not allowed by the mode selection algorithm.) The (unfolded) velocity estimate from the current radial will be assigned to the dealiased velocity, and a velocity estimate from a neighboring radial will be used to select the best unfolding based on the UVM technique. Which neighbor to use is based on the SQIs. The procedure is as follows.

Let the current range-azimuth cell be denoted by range gate k and radial index n . (Note that for the purposes of this process, the SQI values used are penalized by subtracting 1 from them if CREM edit

flags are set.) First, check if the SQI at $(k, n - 1)$ or $(k, n + 1)$ is greater than or equal to SQI_DEALIAS_LIM. If yes, then choose the one with the higher SQI. If no, check if the SQI of $(k - 1, n - 1)$, $(k + 1, n - 1)$, $(k - 1, n + 1)$, or $(k + 1, n + 1)$ is greater than or equal to SQI_DEALIAS_LIM. If yes, select the one with the highest SQI. If no, that is, none of the examined positions had high enough SQI, then choose the one with the higher SQI from $(k, n - 1)$ or $(k, n + 1)$. In the case of only two consecutive DP radials, the same procedure is followed except that either the $n - 1$ or $n + 1$ option is eliminated from consideration. If $k = 1$, the $k - 1$ choices are eliminated, while for $k = K$ (the last available gate) the $k + 1$ choices are eliminated.

Now that a dealiasing “partner” has been chosen, the velocity value from that position is passed along with the velocity from (k, n) , along with their respective PRIs, to the dealiasing process outlined in Section 3.3.3.14. In this case, however, we use the UVM technique to select the correct unfolding, but assign the unfolded velocity from (k, n) or the partner radial instead of averaging the two unfolded velocities. The maximum unfolded velocity, V_{\max} , is set to three times the smaller of the two Nyquist velocities. There are two factors to consider when choosing which unfolded velocity to use. For best azimuthal accuracy and resolution the current (k, n) unfolded velocity should be chosen, because otherwise the dealiased velocity value would effectively be assigned from a neighboring radial. However, if there is range-aliased contamination present in alternating radials (as often happens due to the alternating PRI), then the velocity field can display an artificial spoking pattern caused by the alternating data quality. In this case, the overall data quality is improved if the unfolded velocity from the radial with the higher SQI is used. Therefore, in order to achieve the best of both worlds, we choose the unfolded velocity from (k, n) unless its SQI is below a certain threshold ($2 * \text{SQI_DEALIAS_LIM}$), is also less than the SQI corresponding to the partner radial, and the range overlay flag is set. The effect is that we get maximum radial accuracy and resolution at high SQIs and eliminate spoking when there is range-aliased contamination present.

Finally, we set the velocity dealias fail flag if the SQI from the current or partner cell had a value less than SQI_DEALIAS_LIM.

3.4.4 2D Data Quality Filter

This process is applied to all modes except LP, which does not generate a velocity field and is not affected by range aliasing. It is a 2D filter that operates in an azimuth-range window of dimension $3 \times \text{N_GATES_2D_FILTER}$. At this time, $\text{N_GATES_2D_FILTER} = 5$. Since one radial on each side is required, the first and last radials of the 362 transmitted and processed in a scan are used in the filter but do not produce base data. Currently, the filter serves four purposes.

1. *False dealias correction.* The FDC process is necessary to mitigate the incorrect dealiasing that is inevitably produced by any velocity dealiasing method. It is a type of weighted consensus filter, with the SQIs and data quality flags as the weights.

2. *Despoking.* In the DP mode, the PRI changes on every radial. This means that, for a given range, the out-of-trip overlay power will generally be different on alternating radials, which can lead to a “spoking” effect (alternating “good” and “bad” radials). Even though the phase-code processing can effectively filter out overlaid power in many cases, there are times when the protection breaks down. The despoking filter attempts to interpolate over the bad radial using information from neighboring good radials.
3. *Spectral width smoothing.* As the highest order Doppler moment computed, the spectral width is the most sensitive to differences in the modes used. This can lead to artificial transition features on radial boundaries during an adaptive scan. The 2D filter helps smooth out such non-physical features.
4. *Fill in holes.* If a cell flagged as bad is surrounded by good cells, then we can fill in the hole by intelligent interpolation.

In addition, there is a 3 x 3 filter for the SP mode that removes isolated bad range-azimuth cells. This despeckling operation helps to clean up sidelobe clutter cells that did not get filtered completely in low SNR regions.

In the future more features may be added to the 2D filter. For example, texture and “spin” information could be used to further censor unwanted clutter (Ellis et al. 2005). This could also be extended to include processing on LP data.

Due to boundary conditions, the filter is applied only to gates $k = \text{floor}(\text{N_GATES_2D_FILTER}) + 1$ to $K - \text{floor}(\text{N_GATES_2D_FILTER}/2)$. The data in the “current” radial are stored before filtering, so that the original values are used in the filter for the next radial, not the filtered values. All median operations are taken over the $3 \times \text{N_GATES_2D_FILTER}$ grid. The filtering procedure (applied to each gate) is outlined below.

1. Define a special case: If the radials are DP processed without LP data, $P_S/P_C < \text{SCR_LIM}$, $P_S(k, n - 1)/P_S(k, n) > \text{DP_NOLP_PRATIO_LIM}$, and $P_S(k, n + 1)/P_S(k, n) > \text{DP_NOLP_PRATIO_LIM}$, then *AliasClutter* is set to true. *DP_NOLP_PRATIO_LIM* is currently set to 2. This condition marks the possibility that a narrow weather spectrum had velocity-aliased to DC and was filtered out on this radial (but not the neighboring ones, because of the different PRI).
2. Assign weights, g_{pq} , to points in the $3 \times \text{N_GATES_2D_FILTER}$ grid. If all three radials are DP and *AliasClutter* is false, $g_{pq} = 1$ for $p = k$, $g_{pq} = 0.5$ for $p = k - 1$ and $k + 1$; otherwise, $g_{pq} = 1$. In the all-DP case, these weights are the key to despoking by having the sum of the weights in the current radial equal the sum of the weights in the neighboring radials. This means that the information from the radial(s) with the overall higher SQI will win out.

3. Compute the weighted median velocity, V_{WM} , using g_{pq} as weights.
4. Compute weighted SQIs, $SQI_{pq}^W = g_{pq}SQI_{pq}$.
5. Replace V by V_{WM} and SQI by the weighted median of SQI (using g_{pq}) if *AliasClutter* is true or $SQI < SQI_DEALIAS_LIM$ or any of the following flags are set: bad data, dealias fail, CREM edit. If the replacement was made and if at least half the points in the grid had none of the offending flags and $SQI \geq SQI_DEALIAS_LIM$, then clear the dealias-fail flag.
6. If V was replaced in step 5, skip to step 7. Otherwise, pick the velocity that is closest to V_{WM} : V , V_{alt} , or V_{raw} . If the difference between the selected velocity and V_{WM} is less than or equal to the minimum v_a (corresponding to the maximum PRI used within the grid), then assign that velocity to V ; otherwise, let $V = V_{WM}$ and SQI equal the weighted median of SQI (using g_{pq}). This step is essentially the FDC filter.
7. If *AliasClutter* is true or the bad data flag is set or the CREM edit flag is set or $SQI < SQI_LIM$ or (all radials are DP and $SQI < SQI_Z_LIM$), then replace Z and SNR by their weighted median values (using SQI_{pq}^W). If the replacement was made and if at least half the points in the grid had none of the offending flags and $SQI \geq SQI_LIM$, then clear the bad-data and CREM-edit flags. Currently, SQI_LIM is set to 0.3 and SQI_Z_LIM is set to 0.6. This step and step 5 perform the “fill in the bad data hole” function.
8. If all three radials are DP, replace W with the weighted median of W (using SQI_{pq}^W); otherwise, replace W with the median of W .

If this is SP mode, then go through another loop over gates (2 to $K - 1$). For each gate, if there are no internal data quality flags set and if all of the 8 neighboring cells in the 3 x 3 grid have at least one internal flag set, then let $SNR = 0$ dB for this gate.

3.4.5 SNR Cutoff

This process limits the SNR values output to the RPG to 0 dB or greater. Any smaller values are changed to 0 dB and the corresponding reflectivity values are recomputed to match $SNR = 0$ dB. The only reason why this cutoff is implemented is because the legacy system used it and downstream users (and FAA technicians who work on the system) expect it. Although the cutoff generates an artificial “ramp” function with range in the reflectivity field when the signal is below the threshold, and there may be some usable Doppler information at $SNR < 0$ dB, at this time we wished to minimize disruptions caused by changes in the base data output.

3.4.6 Base Data Flag Generation

In the legacy system, the base data and base data flags were generated in the RPG. Now both are generated in the RDA. Because of the changes in the way the data are processed, there are also changes in the way flags are generated. In fact, there are flags that are no longer used. In the legacy system there were 6 edit flags: valid data, valid velocity, point target, dealias fail, range dealias fail, and azimuth dealias fail. In the new system only 3 edit flags are used: valid data, valid velocity, and dealias fail. Instead of censoring by flags, point targets are simply interpolated over in the base data (Section 3.4.1). In the legacy system, the interpolated data were available in the RPG but were not passed on to downstream users. The range dealias fail and azimuth dealias fail flags are also no longer used, because of the changes in the way velocity dealiasing is performed.

The base data flag generation procedure depends on the mode. For LP mode data, only the valid data flag is relevant. This flag is set if the CREM edit flag was not set. For SP mode data, the valid data flag is always set, because there are no range overlays or CREM editing. The valid velocity flag is set if $SQI \geq SQI_LIM$ and $SNR > 0$ dB. The dealias fail flag is set if the internal dealias fail flag is set. For DP and MP mode data, the CREM edit flags are regenerated, because the SNR cutoff has changed reflectivity values. The valid data flag is set if the bad data flag is not set and the CREM edit flag is not set and (the overlay flag was not set or $SQI \geq SQI_LIM$). The valid velocity flag is set if the valid data flag is set and $SQI \geq SQI_LIM$. (For MP mode data, the valid velocity flag also requires the internal dealias fail flag to be clear. This criterion was added, because it was observed that the spectral width field in MP mode was often of poor quality when the dealias fail flag was set.) The dealias fail flag is set if the internal dealias fail flag is set.

3.4.7 Mode Selection Scorer

This process is only relevant to the surface scan, where the data collected in LP mode are used to select on a radial-by-radial basis from the DP or one of the three MP (see Section 3.3.3) modes. The selection process is divided into two parts. This first part (described in this section) assigns a score to each of the four possible mode selections for each gate, then a weighted overall score for the radial. The second part, which selects the optimal modes for the whole scan based on the radial scores, is processed in the IQM and is described in Section 3.2.2.

The score per gate, $y(k)$, spans a range of -1 to 1. Its purpose is to quantify the expected quality of the velocity estimate at gate k for different modes. The primary scores are -1 (bad) or 1 (good), depending on whether the range-azimuth cell is free from overlay and/or clutter problems. For MP mode types, expected velocity dealiasing failure leads to a score of 0, which is analogous to the DP case of averaging a bad and a good score over two radials. There is a secondary score adjustment using the expected velocity estimate error, which spreads the score across non-integer values. This adjustment establishes the “default” hierarchy of mode preference when there is no overlay in any of the range gates. The following steps are used to compute y_k .

1. The LP SNR is converted to signal power as $P_S = P_N 10^{(SNR/10)}$, since the SNR is stored in dB units. For $k > N_GATES_BD$, we add the clutter power back to the signal power, because any range-overlaid signal will include the clutter power.
2. The LP spectral width, W_k , is smoothed by a boxcar mean over three gates to yield W_{Sk} , because it is a noisy quantity.
3. In steps 3 through 5 we compute the score for the DP mode. Look up the signal and clutter powers from trip $n = 1$ to n_{trip} for this gate (see Section 3.3.2.1.2). Assign $y_{DP}(k) = 1$ if the 1st trip is the “strong” trip and go to step 11.
4. If the 1st trip is not the strongest trip and there is significant signal in more than one other trip, assign $y_{DP}(k) = -1$ and go to step 11.
5. If (the smoothed spectral width is greater than SPECWIDTH_LIM and the ratio of the maximum out-of-trip signal to the 1st trip signal exceeds PRATIO_LIM) or the signal-to-clutter ratio is less than SCR_LIM, then assign $y_{DP}(k) = -1$; otherwise, set $y_{DP}(k) = 1$. Go to step 11. Currently SPECWIDTH_LIM is set to 2 m s^{-1} and PRATIO_LIM is set to 30 dB for the shorter PRI and 20 dB for the longer PRI in the DP PRI pair.
6. In steps 6 through 10 we compute the scores for the MP modes. Each step outlined is repeated for the three MP modes. First, compute the range overlays. This process is the same as given in Section 3.3.3.4.
7. If all the pulses are clean or the 1st trip signal is greater than or equal to the mean overlay signal, then assign $y_{MPi}(k) = 1$ and go to step 11.
8. Determine which PRI sets are clean for velocity calculation. There must be at least two clean pulses per set and at least one of the pulse pairs must be consecutive pulses.
9. If (the clutter is strong enough to trigger use of the GCF and the overlay-to-signal ratio is greater than OSR_LIM) or a DC subtraction is not good enough to produce a valid velocity estimate or there are no clean PRI sets, then set $y_{MPi}(k) = -1$ and go to step 11.
10. If the maximum frequency difference in the clean PRI sets is less than FREQ_DIFF_LIM, then assign $y_{MPi}(k) = 0$; otherwise, set $y_{MPi}(k) = 1$.
11. If $SNR_k > 1/SD_NSR_LIM$, then adjust the score by subtracting $\sigma_{vk}/SD_VELOCITY_LIM$, but do not let $y(k)$ go below -1. SD_NSR_LIM is currently set to 1 and SD_VELOCITY_LIM is currently set to 10 m s^{-1} . σ_{vk} is the theoretical velocity estimate error given by (Doviak and Zrnić 1993)

$$\sigma_{vk} = \frac{\lambda e^{8\pi W_{sk} \tau / \lambda}}{4\pi\tau \sqrt{2} \sqrt{1 - e^{-8\pi W_{sk} \tau / \lambda}}} \sqrt{\frac{\lambda \left(1 - e^{-8\pi W_{sk} \tau / \lambda}\right)}{4\sqrt{\pi W_{sk} \tau}} + \frac{1}{SNR_k^2} + \frac{2}{SNR_k} \left(1 - e^{-32\pi W_{sk} \tau / \lambda}\right)}, \quad (3-29)$$

where $\tau = T$ and $\mu = M$ for DP mode, and τ is assigned the mean of the T s and $\mu = M_{pi}$ in MP mode.

Because of the need for interrational velocity dealiasing, the DP mode always requires at least two consecutive radials. Therefore, the DP scores should be averaged in some way over two radials, the minimum selection unit, to account for dealiasing success or failure. There are three possibilities: 1) only the previous radial is DP, 2) only the next radial is DP, and 3) both previous and next radials are DP. For case 1, we average the current radial scores with the previous radial scores, $y_{DP}^P(k, l) = [y_{DP}(k, l) + y_{DP}(k, l - 1)]/2$. For case 2, we average the current radial scores with the next radial scores, $y_{DP}^N(k, l) = [y_{DP}(k, l) + y_{DP}(k, l + 1)]/2$. Case 3 is more complicated, because the interrational dealiasing algorithm chooses the better option (previous or next radial) on a gate-by-gate basis. Thus, we set $y_{DP}^{PN}(k, l) = \{y_{DP}(k, l) + \max[y_{DP}(k, l - 1), y_{DP}(k, l + 1)]\}/2$.

We now have the gate scores assigned relative to the expected velocity estimate quality. However, not all gates have equal importance to the TDWR's mission. Therefore, the gate scores should be weighted appropriately before being averaged over the radial. The area noted for attention (ARENA) is the most important region for surveillance, followed by the microburst and gust-front product generation regions. Also, since the area of a range-azimuth cell increases linearly with distance, there should be a linear range weighting. The following steps are taken to form the gate score weights.

1. Assign unity value to ARENA cells in array A_1 ; all other cells are zero. Assign 1 and 0.5 to microburst and gust-front product cells, respectively, in array A_2 ; all other cells are zero. Assign values of $(k - 1)/(N_GATES_BD - 1)$ in array A_3 . All three arrays are of dimension $N_GATES_BD \times 360$.
2. Compute weights, $h(k, l) = [ARENA_FAC * A_1(k, l) + MBGF_FAC * A_2(k, l) + RANGE_FAC * A_3(k, l)] / (ARENA_FAC + MBGF_FAC + RANGE_FAC)$. Currently, ARENA_FAC is set to 8, MBGF_FAC is set to 4, and RANGE_FAC is set to 1.
3. Set $h(k, l) = 0$ at $k = 1$ and 2. (The data at these gates are usually invalid.)
4. Normalize $h(k, l)$ so that their sum over each radial is unity.

Now we multiply $h(k, l)$ with each type of $y(k, l)$ and sum over range to get the weighted radial scores: Y_{MP1} , Y_{MP2} , Y_{MP3} , Y_{DP}^P , Y_{DP}^N , and Y_{DP}^{PN} . These 360-element arrays are passed to the mode selection processor (Section 3.2.2) in the IQM.

4. SUGGESTIONS FOR FUTURE ENHANCEMENTS

In this first round of signal processing enhancement, we focused on RV ambiguity mitigation. We also worked on improving clutter filtering and censoring, but there are still ways in which this front can be advanced. Here are some ideas.

1. *Adaptive windowing.* Data windowing presents a compromise between increasing sidelobe suppression and increasing estimate errors. Although we have in place a simple window selection process (Section 3.3.1.7), a more optimal window can be chosen if iterations are made based on the clutter suppression achieved. Such a technique is used in the RVP9 (Vaisala 2009) and can be fairly easily implemented in the RDA given more computing capacity.
2. *Dual-stream processing.* Without any spatial context information, there is no way to distinguish between a narrow weather spectrum centered at DC and a ground clutter spectrum. That is why the zero-Doppler line during stratiform rain presents anomalously low reflectivity when clutter filtering is turned on. To circumvent this problem, it is possible to delay the GCF decision until the 2D “image processing” stage by generating two streams of data—unfiltered and clutter filtered. The correct selection can be made based on spatial texture and other contextual data. When there is range overlay and ground clutter, the decision whether or not to use the GCF can also be made more intelligently after the fact. The price will be a doubling in computational burden between the GCF and 2D processing stages. A simpler version (a front-end 1D algorithm) of such a GCF decision scheme is being implemented for the Weather Surveillance Radar-1988 Doppler (WSR-88D) (Hubbert et al. 2009).
3. *Higher elevation CREMs.* Although we have a GCF widening feature to help eliminate antenna sidelobe clutter, it may not be enough for sites with severe sidelobe clutter or it may remove too much weather information in areas with no sidelobe clutter. A straightforward alternative is to provide CREMs for all elevations. The cost will be more memory allocation for CREMs and an increase in labor for the people who run the CREM generation program.
4. *More 2D filtering.* Currently we use a 3×5 (azimuth \times range) filter. It may be of benefit to increase the filter size so that more contextual information can be incorporated for processing. For example, wind-shear false alarms are a noted problem at Salt Lake City due to suspected bird flocks causing chaotic velocity patterns (there is a bird sanctuary just north of the airport). We have shown that a 5×7 filter using texture and “spin” fields (Ellis et al. 2005) has the potential to effectively censor such cases (Figure 4-1). Clutter residue might also be recognized using other criteria, which will help censor breakthroughs during times of high winds that widen the spectra of not-so-stationary reflectors such as vegetation, as well as intense sidelobe clutter. The price to be paid here is increased radial buffering (memory) and computation time.

5. *3D filtering.* Contextual information in the vertical dimension could also be used for perhaps even more effective clutter censoring. The short latency requirement prohibits the withholding of volume scan data in the RDA without outputting, but we can keep all the latest elevation scans stored for use by the filter. This has been an active area of research with several papers published in the last few years (Charalampidis et al. 2002; Steiner and Smith 2002; Lee et al. 2005; Berenguer et al. 2006). Again, memory allocation and computation time will increase.
6. *Contextual spectral processing.* For transmission and processing modes where spectral analysis is possible, one can also use spatial contextual information to filter out unwanted signals in the spectral domain. For example, many range gates worth of spectra can be collectively processed (Cho 2009). The potential exists to produce a larger percentage of “clean” base data relative to the case where clutter-contaminated cells are recognized only after the base data have already been generated. Moving clutter such as vehicular traffic and birds may be effectively filtered out in this way. This type of processing would require a change in our parallel processing architecture to allow “cross talk” between range gates during spectral processing. Once again, more computational time is required.

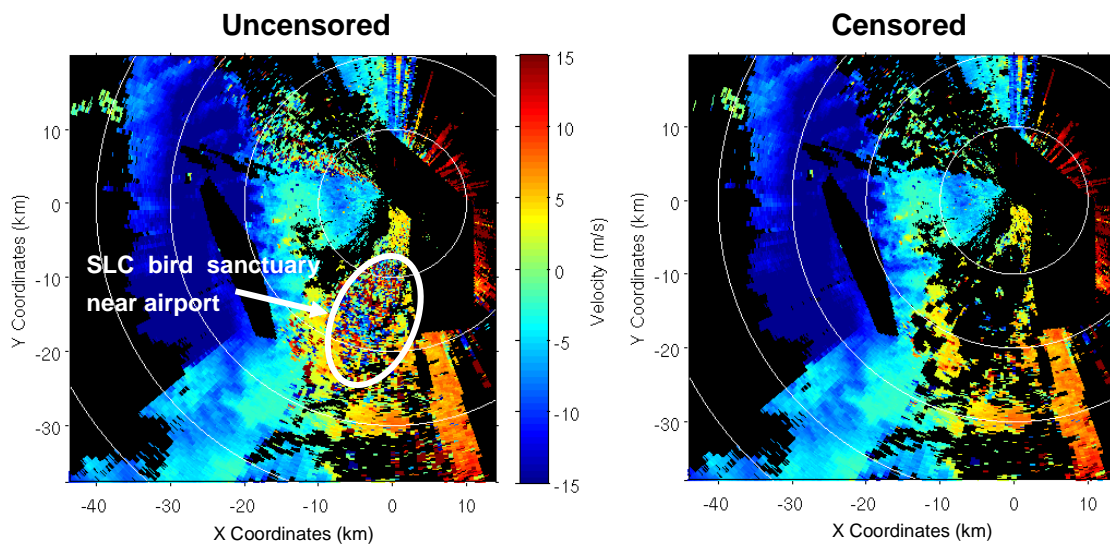


Figure 4-1. Censoring of the velocity field produced by the legacy RDA at Salt Lake City using a 5×7 (azimuth \times range) filter. The scan is from 01:35 Z, 26 February 1999, at 0.3° elevation.

Here are some other suggestions besides improvements in clutter removal. By oversampling and pseudowhitening in range, one can reduce the base data estimate errors (Torres et al. 2004). The capability to oversample already exists in the RVP9. However, since the computational load scales

roughly as the oversampling ratio, the implementation of this technique definitely requires an upgrade to dramatically faster and/or a larger number of processors.

Dwell lengths could be adaptively changed. The antenna beamwidth of the TDWR is 0.55° and yet the base data azimuthal resolution is 1° . Beyond a range of about 8.6 km, the azimuthal resolution becomes coarser than the range resolution. Shorter dwells could be used in cases where the weather signal is high compared to clutter and range overlaid signals to provide a sharper picture of the event. In principle, this is a desirable feature. However, the main difficulty with this scheme is that all downstream programs must be revised to accommodate variable resolution data, which may not be an acceptable task for all involved.

Although the “T” in TDWR stands for “terminal,” it is a powerful radar that can sense weather out to the full LP surveillance range of 460 km. With the new range dealiasing transmission and processing schemes in place, it is possible to generate base data beyond the FAA-required range of 90 km. In fact, it is a fairly straightforward algorithm extension that has been demonstrated off-line (Cho 2006). See Figure 4-2 for an example. It is essentially “free” data that are currently being thrown away. Besides the increase in computational load, the extra data will obviously widen the output bandwidth, which may require an upgrade in the communication hardware.

Weather radars rely on the presence of radiowave scattering entities such as hydrometeors and insects to sense the dynamic evolution of the atmosphere. Under clear-air, low-reflectivity conditions, when no such “visible” tracers are present, air mass boundaries such as the outflow edge of a dry microburst may go undetected. A radar data processing technique has now been developed to estimate the near-ground atmospheric refractivity field using ground targets (Fabry et al. 1997). Refractivity is dependent on the moist thermodynamic variables of the atmosphere and, thus, can be used to detect air mass changes and boundaries. As most of the experiments using this technique has been at S band, it is worth investigating to see whether the TDWR is capable of generating good quality refractivity fields.

For the WSR-88D, dual polarization is the next major upgrade. Benefits of polarization diversity are many, including clutter discrimination, identification of hydrometeor types, and improved precipitation estimates (Zrnić and Ryzhkov 1999). The same benefits would be obtained for the TDWR if it is upgraded to a dual-polarization system.

Following RV ambiguity, the data quality challenge of most concern is moving clutter. And, because of the low SNR conditions often accompanying wind shear events, as well as the high relief in the surrounding topography, moving clutter presents the worst problems at Las Vegas (LAS) and Salt Lake City (SLC). Highly visible road traffic is the culprit at LAS and birds are the main source of vexation at SLC (Cho 2008). With the enhanced RDA now operating in demonstration mode at these two sites, we are collecting base and I&Q data for an in-depth study and subsequent development of mitigation algorithms. Some of the algorithmic enhancements discussed above may form part of the solution (Cho 2009), but there may be other approaches that will be formulated as well.

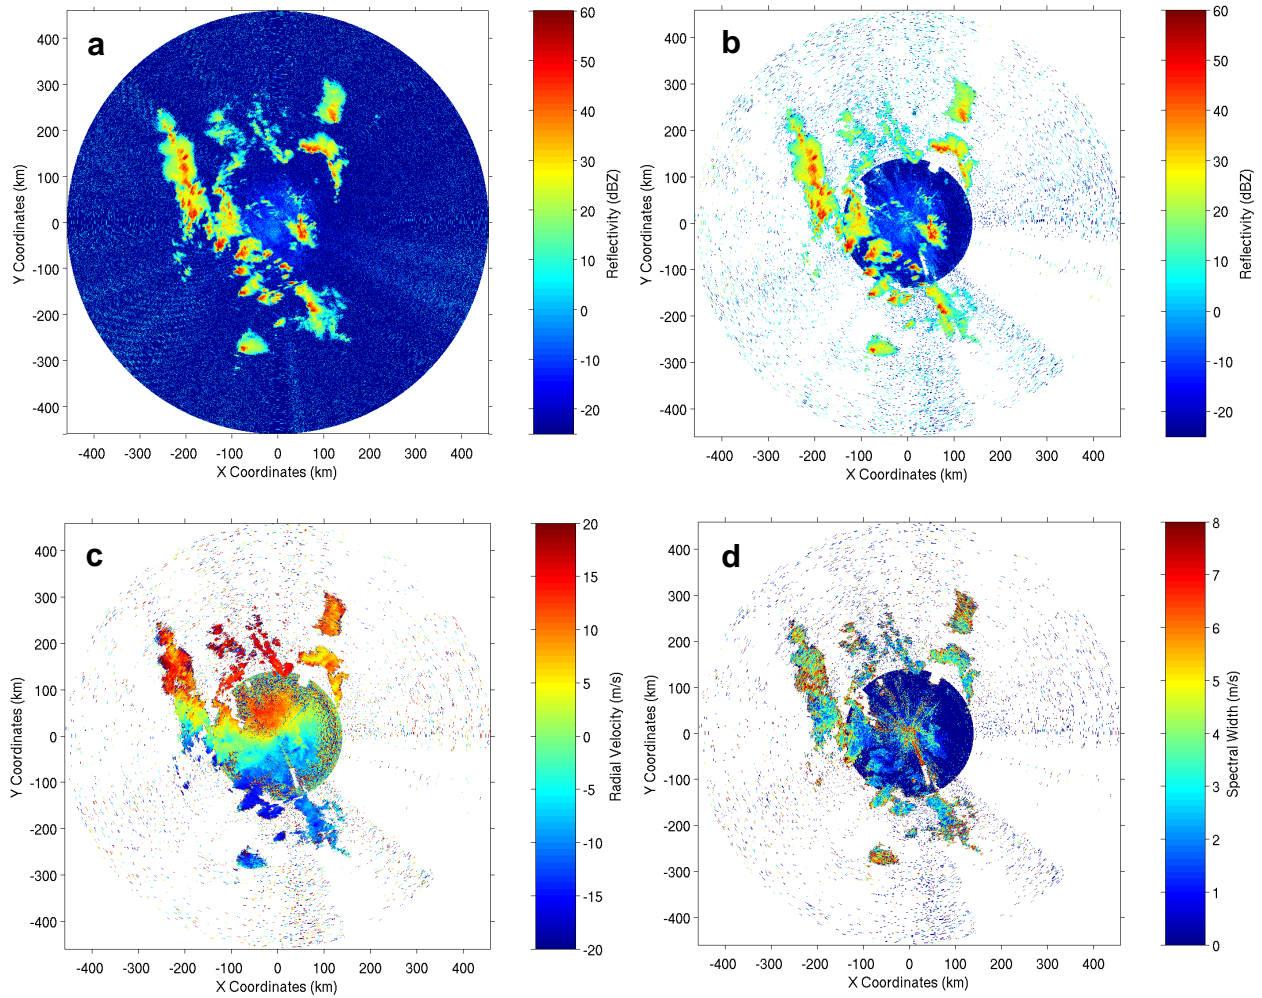


Figure 4-2. (a) Long-PRI reflectivity providing “truth.” (b) Multi-PRI range-dealiased reflectivity. (c) Multi-PRI range-dealiased radial velocity. (d) Multi-PRI range-dealiased Doppler spectral width. These 0.3° scans were taken at 20:40 Z, 17 March 2003, with the PSF TDWR in Oklahoma City using the prototype RDA. Taken from Figures 8 to 11 of Cho (2006). Note that censoring was not applied to these data fields. For further details, see Cho (2006).

5. SUMMARY

Taking advantage of the increased computing power and flexibility in radar control available in the upgraded RDA, we have developed and demonstrated cutting edge signal transmission and processing techniques for the TDWR. Pulse phase and repetition interval diversities were utilized extensively along with single-scan velocity dealiasing and 2D spatial context filtering. Spectral domain processing for range dealiasing and ground clutter filtering were included. For the surface scan, most crucial to the TDWR's mission of detecting low-altitude wind shear, a radial-by-radial adaptive signal transmission and processing scheme was adopted, the first time such an approach has been used in weather radars.

The resulting base data quality is a significant improvement over the data quality provided by the legacy system, especially in the areas of RV ambiguity and AP elimination. The maximum clutter suppression capability has been raised due to the new capability to cohere the received data to the measured transmitter pulse phase. Because of the single-scan velocity dealiasing method, the "dealiasing" tilt has been eliminated from the monitor and hazard volume scans. The time savings will likely be used to slow the antenna rotation rate for the monitor mode (reducing base data estimate errors) and on an additional precipitation tilt for the hazard mode.

This report detailed the enhanced signal processing algorithms implemented in the RDA. Suggestions for future enhancements were also discussed. The open and scalable architecture of the new RDA (along with, no doubt, the continued leverage provided by Moore's Law) makes possible the incorporation of additional signal processing tasks for years to come.

APPENDIX A

MULTI-PRI CLUTTER FILTER SPECIFICATIONS

For a given antenna rotation rate, there are 12 FIR filters used in clutter filtering the MP mode data (3 PRI sequences \times 4 suppression levels). Since the surface scan operates at two rotation rates (monitor and hazard), there are then 24 FIR filters used. However, since the FAA may decide at any time to alter the antenna rotation rates for their operational volume scans, we wish to specify filters for a range of rotation rates. We have done so, but rather than listing the filter coefficients, which would take up an unreasonable number of pages, we summarize in tables the parameters needed to generate the filters using Cho and Chornoboy's (2005) program (available from the first author as Matlab code), and their characteristics (Tables A-1 through A-3).

Given the PRI sequence, the parameters required as input to the filter generation program are the antenna rotation rate, the spectral range to be covered, phase weight, magnitude weight, and stop-band edge. For all the filters, the spectral range is set to be ± 3 times the maximum-PRI v_a , and the phase weight is set to unity. Therefore, only the antenna rotation rate, magnitude weight, and stop-band edge are in the tables. (The specified stop-band edge is normalized to the maximum-PRI v_a .) The actual computed stop and pass band edges (defined by 3-dB roll-off points) are also listed in the tables. In all cases the actual suppression, as computed for a Gaussian clutter spectrum with width $[\sigma_R^2 + (0.1)^2]^{1/2}$ (see Section 3.3.1.8), is within 0.2 dB of the nominal values given in the tables. The three MP PRI sets currently used are 4A (600, 670, 740, and 810 μ s), 4B (698, 798, 898, and 998 μ s), and 8 (600, 648, 696, 744, 792, 840, 888, and 936 μ s).

TABLE A-1

FIR Filter Specifications and Characteristics for MP Type 4A

M_p	γ (°/s)	Suppression (dB)	Magnitude Weight	Normalized Stop Band Edge	Stop Band Edge (m s ⁻¹)	Pass Band Edge (m s ⁻¹)
13	26	65	950,000	0.073	0.96	2.60
		50	80,000	0.066	0.96	2.21
		35	1,600	0.06	0.87	1.73
		20	47	0.056	0.77	1.44
14	24	65	1,100,000	0.067	1.16	2.41
		50	37,000	0.061	0.96	2.02
		35	1,100	0.055	0.87	1.64
		20	50	0.051	0.67	1.25
15	22	65	1,400,000	0.061	1.06	2.21
		50	46,000	0.057	0.96	1.93
		35	1,300	0.052	0.77	1.54
		20	45	0.05	0.77	1.25
16	21	65	900,000	0.061	1.06	2.12
		50	33,000	0.055	0.96	1.83
		35	1,100	0.05	0.77	1.44
		20	50	0.048	0.67	1.16
17	20	65	2,000,000	0.057	0.96	2.02
		50	25,000	0.055	0.96	1.73
		35	920	0.049	0.77	1.44
		20	46	0.047	0.67	1.16
18	19	65	1,100,000	0.055	0.96	1.93
		50	44,000	0.054	0.58	1.64
		35	1,400	0.048	0.67	1.35
		20	43	0.046	0.67	1.06
19	18	65	4,300,000	0.054	0.87	1.83
		50	29,000	0.049	0.58	1.54
		35	1,100	0.046	0.67	1.25
		20	40	0.045	0.67	1.06
20	17	65	1,700,000	0.049	0.87	1.64
		50	69,000	0.046	0.77	1.44
		35	2,000	0.041	0.58	1.16
		20	46	0.04	0.58	0.96
21	16	65	860,000	0.049	0.87	1.64
		50	39,000	0.044	0.77	1.44
		35	1,400	0.04	0.58	1.16
		20	44	0.038	0.58	0.96

TABLE A-2

FIR Filter Specifications and Characteristics for MP Type 4B

M_p	γ (°/s)	Suppression (dB)	Magnitude Weight	Normalized Stop Band Edge	Stop Band Edge (m s ⁻¹)	Pass Band Edge (m s ⁻¹)
11	26	65	2,000,000	0.088	1.09	2.58
		50	43,000	0.083	1.02	2.19
		35	1,000	0.078	0.94	1.80
		20	41	0.073	0.86	1.49
12	23	65	1,100,000	0.08	1.09	2.34
		50	32,000	0.075	0.94	1.95
		35	900	0.07	0.86	1.64
		20	45	0.065	0.70	1.25
13	21	65	1,100,000	0.073	0.78	2.11
		50	33,000	0.068	0.86	1.80
		35	950	0.063	0.78	1.49
		20	47	0.058	0.63	1.17
14	20	65	2,200,000	0.069	0.94	2.03
		50	59,000	0.063	0.78	1.72
		35	1,300	0.059	0.70	1.41
		20	54	0.054	0.55	1.09
15	19	65	1,200,000	0.067	0.94	1.88
		50	38,000	0.061	0.78	1.64
		35	1,100	0.059	0.70	1.33
		20	49	0.053	0.63	1.02
16	18	65	3,300,000	0.066	0.86	1.80
		50	25,000	0.061	0.78	1.56
		35	900	0.056	0.70	1.33
		20	44	0.051	0.63	1.02
17	17	65	1,500,000	0.063	0.86	1.72
		50	42,000	0.058	0.78	1.49
		35	1,300	0.053	0.63	1.17
		20	52	0.048	0.55	0.94
18	16	65	5,000,000	0.06	0.78	1.64
		50	26,000	0.056	0.78	1.41
		35	980	0.051	0.63	1.17
		20	46	0.047	0.55	0.86

TABLE A-3

FIR Filter Specifications and Characteristics for MP Type 8

M_p	γ (°/s)	Suppression (dB)	Magnitude Weight	Normalized Stop Band Edge	Stop Band Edge (m s ⁻¹)	Pass Band Edge (m s ⁻¹)
6	26	65	2,200,000	0.082	1.17	2.58
		50	48,000	0.073	1.00	2.25
		35	2,000	0.065	0.83	1.75
		20	55	0.059	0.67	1.33
7	21	65	1,400,000	0.067	1.00	2.17
		50	40,000	0.061	0.83	1.83
		35	1100	0.056	0.75	1.50
		20	46	0.05	0.58	1.17
8	19	65	1,600,000	0.062	0.92	2.00
		50	54,000	0.056	0.83	1.67
		35	1300	0.054	0.67	1.33
		20	49	0.045	0.50	1.00
9	17	65	2,100,000	0.058	0.83	1.75
		50	24,000	0.055	0.83	1.50
		35	840	0.049	0.67	1.25
		20	39	0.044	0.58	0.92
10	16	65	1,300,000	0.055	0.83	1.67
		50	17,000	0.055	0.58	1.42
		35	1,400	0.047	0.58	1.17
		20	36	0.044	0.58	0.92

Although the clutter suppression levels in Tables A-1 to A-3 correspond to particular antenna rotation rates as shown, each M_p set is applicable to a range of rotation rates. These ranges are calculated so that the full MP sequence is transmitted during a 1° dwell with at least a 1-ms margin to account for fluctuations in the rotation speed. For convenience, these ranges are listed in Table A-4.

TABLE A-4**Valid Antenna Rotation Rates**

Type	M_p	γ_{\min} (°/s)	γ_{\max} (°/s)
4A	13	24.70356	26.55337
	14	23.09469	24.70355
	15	21.68257	23.09468
	16	20.43318	21.68256
	17	19.31994	20.43317
	18	18.32173	19.31993
	19	17.42160	18.32172
	20	16.60578	17.42159
	21	15.86294	16.60577
4B	11	23.97852	26.10148
	12	22.17492	23.97851
	13	20.62366	22.17491
	14	19.27525	20.62365
	15	18.09234	19.27524
	16	17.04623	18.09233
	17	16.11448	17.04622
	18	15.27931	16.11447
8	6	22.72314	26.41031
	7	19.93938	22.72313
	8	17.76325	19.93937
	9	16.01537	17.76324
	10	14.58066	16.01536

GLOSSARY

ADC	analog-to-digital converter
CNR	clutter-to-noise ratio
COTS	commercial off-the-shelf
CREM	clutter residue map
DFT	digital Fourier transform
DSP	digital signal processing
FAA	Federal Aviation Administration
FDC	false-dealias correction
FFT	Fast Fourier Transform
FIR	finite impulse response
GCF	ground clutter filter
GMAP	Gaussian model adaptive processing
I&Q	in-phase and quadrature
IQM	in-phase and quadrature master
IQS	in-phase and quadrature slave
LAS	Las Vegas
MIT-LL	Massachusetts Institute of Technology Lincoln Laboratory
MP	multi-PRI
PRI	pulse repetition interval
RDA	radar data acquisition
RPG	radar product generator
SLC	Salt Lake City
SOR	signal-to-overlay ratio
TDWR	Terminal Doppler Weather Radar

UVM	unfolded-velocity matching
WMAD	weighted median absolute deviation
WSR-88D	Weather Surveillance Radar - 1988 Doppler

REFERENCES

- Arce, G. R., 1998: A general weighted median filter structure admitting negative weights. *IEEE Trans. Signal Processing*, **46**, 3195-3205.
- Berenguer, M., D. Sempere-Torres, C. Corral, and R. Sánchez-Diezma, 2006: A fuzzy logic technique for identifying nonprecipitating echoes in radar scans. *J. Atmos. Oceanic Technol.*, **23**, 1157-1180.
- Billingsley, J. B., 2002: *Low-Angle Radar Land Clutter: Measurements and Empirical Models*. William Andrews, 722 pp.
- Charalampidis, D., T. Kasparis, and W. L. Jones, 2002: Removal of nonprecipitation echoes in weather radar using multifractals and intensity. *IEEE Trans. Geosci. Remote Sens.*, **40**, 1121-1131.
- Cho, J. Y. N., 2003: Evaluation of TDWR range-ambiguity mitigation techniques. Project Rep. ATC-310, MIT Lincoln Laboratory, Lexington, MA, 47 pp.
- Cho, J. Y. N., 2005: Multi-PRI signal processing for the Terminal Doppler Weather Radar. Part II: Range-velocity ambiguity mitigation. *J. Atmos. Oceanic Technol.*, **22**, 1507-1519.
- Cho, J. Y. N., 2006: Extended-range signal recovery using multi-PRI transmission for Doppler weather radars. Project Rep. ATC-322, MIT Lincoln Laboratory, Lexington, MA, 30 pp.
- Cho, J. Y. N., 2008: TDWR dry site problem assessment and RDA Software Build 3 recommendation. Project Memorandum 43PM-Wx-0107, MIT Lincoln Laboratory, Lexington, MA, 23 pp.
- Cho, J. Y. N., 2009: Moving clutter spectral filter for Terminal Doppler Weather Radar. Preprints, *34th Conf. on Radar Meteorology*, Williamsburg, VA, Amer. Meteor. Soc., P5.2, <http://ams.confex.com/ams/pdfpapers/155381.pdf>.
- Cho, J. Y. N., and E. S. Chornoboy, 2005: Multi-PRI signal processing for the Terminal Doppler Weather Radar. Part I: Clutter filtering. *J. Atmos. Oceanic Technol.*, **22**, 575-582.
- Cho, J. Y. N., G. R. Elkin, and N. G. Parker, 2003: Range-velocity ambiguity mitigation schemes for the enhanced Terminal Doppler Weather Radar. Preprints, *31st Conf. on Radar Meteorology*, Seattle, WA, Amer. Meteor. Soc., 463-466, <http://ams.confex.com/ams/pdfpapers/63750.pdf>.
- Cho, J. Y. N., G. R. Elkin, and N. G. Parker, 2005: Enhanced radar data acquisition system and signal processing algorithms for the Terminal Doppler Weather Radar. Preprints, *32nd Conf. on Radar Meteorology*, Albuquerque, NM, Amer. Meteor. Soc., P4R8, <http://ams.confex.com/ams/pdfpapers/96018.pdf>.
- Chornoboy, E. S., 1993: Clutter filter design for multiple-PRT signals. Preprints, *26th Conf. on Radar Meteorology*, Norman, OK, Amer. Meteor. Soc., 235-237.
- Crocker, S. C., 1988: TDWR PRF selection criteria. Project Rep. ATC-147, DOT/FAA/PM-87-25, MIT Lincoln Laboratory, Lexington, MA, 57 pp.

- Dijkstra, E. W., 1959: A note on two problems in connexion with graphs. *Numerische Mathematik*, **1**, 269-271.
- Doviak, R. J., and D. S. Zrnić, 1993: *Doppler Radar and Weather Observations*. Academic Press, San Diego, CA, 562 pp.
- Elkin, G. R., O. J. Newell, and M. E. Weber, 2002: Enhancements to Terminal Doppler Weather Radar to improve aviation weather services. Preprints, *10th Conf. on Aviation, Range, and Aerospace Meteorology*, Portland, OR, Amer. Meteor. Soc., 28-31.
- Ellis, S. M., M. Dixon, G. Meymaris, S. Torres, and J. Hubbert, 2005: Radar range and velocity ambiguity mitigation: Censoring methods for the SZ-1 and SZ-2 phase coding algorithms. Preprints, *21st Conf. on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, San Diego, CA, Amer. Meteor. Soc., 19.3, <http://ams.confex.com/ams/pdfpapers/86439.pdf>.
- Fabry, F., C. Frush, I. Zawadzki, and A. Kilambi, 1997: On the extraction of near-surface index of refraction using radar phase measurements from ground targets. *J. Atmos. Oceanic Technol.*, **14**, 978-987.
- Hildebrand, P. H., and R. S. Sekhon, 1974: Objective determination of the noise level in Doppler spectra. *J. Appl. Meteor.*, **13**, 808-811.
- Hubbert, J. C., M. Dixon, and S. Ellis, 2009: Weather radar ground clutter, Part II: Real time identification and filtering. *J. Atmos. Oceanic Technol.*, **26**, 1181-1197.
- Laird, B. G., 1981: On ambiguity resolution by random phase processing. Preprints, *20th Conf. on Radar Meteorology*, Boston, MA, Amer. Meteor. Soc., 327-331.
- Lee, G., Y.-H. Cho, K.-E. Kim, and I. Zawadzki, 2005: Identification and removal of non-precipitation echoes using the characteristics of radar echoes. Preprints, *32nd Conf. on Radar Meteorology*, Albuquerque, NM, Amer. Meteor. Soc., 4R3, <http://ams.confex.com/ams/pdfpapers/97079.pdf>.
- Meymaris, G., J. C. Hubbert, and G. Gray, 2009: A simplified approach to staggered PRT clutter filtering. Preprints, *25th Conf. on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, Phoenix, AZ, Amer. Meteor. Soc., 10B.4, <http://ams.confex.com/ams/pdfpapers/147955.pdf>.
- Michelson, M., W. W. Shrader, and J. G. Wieler, 1990: Terminal Doppler Weather Radar. *Microwave J.*, **33**, 139-148.
- Raytheon, 1992: Terminal Doppler Weather Radar (TDWR) Build 5 system/segment design document (S/SDD). CDRL Seq. No. B002-Bld5-1, Raytheon Company, Sudbury, MA, 286 pp.
- Rinehart, R. E., 1991: Spurious velocities in Doppler radar data caused by a moving antenna feedhorn. *J. Atmos. Oceanic Technol.*, **8**, 733-745.

- Sachidananda, M., and D. S. Zrnić, 1999: Systematic phase codes for resolving range overlaid signals in a Doppler weather radar. *J. Atmos. Oceanic Technol.*, **16**, 1351-1363.
- Siggia, A., 1983: Processing phase coded radar signals with adaptive digital filters. Preprints, *21st Int. Conf. on Radar Meteorology*, Edmonton, AB, Canada, Amer. Meteor. Soc., 167-172.
- Siggia, A., and R. E. Passarelli, Jr., 2004: Gaussian model adaptive processing (GMAP) for improved ground clutter cancelation and moment estimation. Preprints, *Third European Conf. on Radar in Meteorology and Hydrology*, Visby, Sweden, Copernicus Gesellschaft, 67-73.
- Steiner, M., and J. A. Smith, 2002: Use of three-dimensional reflectivity structure for automated detection and removal of nonprecipitating echoes in radar data. *J. Atmos. Oceanic Technol.*, **19**, 673-686.
- Torres, S. M., C. D. Curtis, and J. R. Cruz, 2004: Pseudowhiting of weather radar signals to improve spectral moment and polarimetric variable estimates at low signal-to-noise ratios. *IEEE Trans. Geosci. Remote Sens.*, **42**, 941-949.
- Trunk, G., and S. Brockett, 1993: Range and velocity ambiguity reduction. Preprints, *1993 IEEE National Radar Conf.*, Lynnfield, MA, IEEE, 146-149.
- Vaisala, 2009: RVP9 Digital IF Receiver/Doppler Signal Processor User's Manual. Vaisala Inc., Westford, MA, 451 pp.
- Wieler, J. G., and S.-C. Hu, 1993: Elimination of Doppler ambiguities in weather radar data. Preprints, *1993 IEEE National Radar Conf.*, Lynnfield, MA, IEEE, 163-166.
- Zrnić, D. S., and A. V. Ryzhkov, 1999: Polarimetry for weather surveillance radars. *Bull. Amer. Meteor. Soc.*, **80**, 389-406.