Multi-PRI Signal Processing for the Terminal Doppler Weather Radar. Part II: Range–Velocity Ambiguity Mitigation

JOHN Y. N. CHO

Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, Massachusetts

(Manuscript received 26 January 2005, in final form 15 April 2005)

ABSTRACT

Multiple pulse-repetition interval (multi-PRI) transmission is part of an adaptive signal transmission and processing algorithm being developed to combat range–velocity (RV) ambiguity for the Terminal Doppler Weather Radar (TDWR). In Part I of this two-part paper, an adaptive clutter filtering procedure that yields low biases in the moments estimates was presented. In this part, algorithms for simultaneously providing range-overlay protection and velocity dealiasing using multi-PRI signal transmission and processing are presented. The effectiveness of the multi-PRI RV ambiguity mitigation scheme is demonstrated using simulated and real weather radar data, with excellent results. Combined with the adaptive clutter filter, this technique will be used within the larger context of an adaptive signal transmission and processing scheme in which phase-code processing will be a complementary alternative.

1. Introduction

The Terminal Doppler Weather Radar (TDWR) radar data acquisition (RDA) subsystem is being replaced as part of a broader Federal Aviation Administration (FAA) program to improve the supportability of the system. An engineering prototype RDA has been developed with a scalable, open-systems hardware platform (Cho et al. 2004). With the dramatically increased computing power and more flexible transmitter control, modern signal processing algorithms can be implemented to improve the quality of the reflectivity, Doppler velocity, and spectral width estimates.

One of the most challenging data quality issues is range–velocity (RV) ambiguity. Compared to S-band radars such as the Weather Surveillance Radar-1988 Doppler (WSR-88D), the ambiguity is more severe for C-band radars such as the TDWR. This is illustrated in Fig. 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, $v_a \geq \pi\sigma_v$ (Doviak and Zrnić 1993), assuming a maximum Doppler velocity spectral width of $\sigma_v = 4$ 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.

With the availability of diversity in PRI and pulse-transmission phase, multi-PRI and phase-code techniques can be applied to this problem. These two approaches have complementary strengths and weaknesses for range-overlay protection (Cho et al. 2003). Multi-PRI 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

Corresponding author address: John Y. N. Cho, MIT Lincoln Laboratory, 244 Wood St., S1-659, Lexington, MA 02420-9185.
E-mail: jync@mit.edu

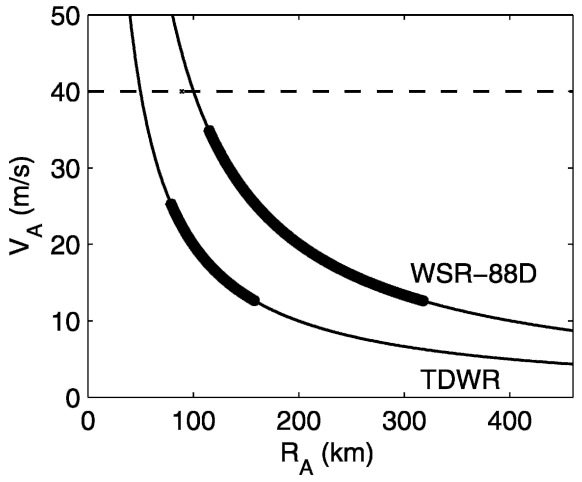

FIG. 1. Unambiguous velocity vs unambiguous range for the WSR-88D and TDWR. The thick lines indicate the operating ranges for velocity estimation mode as bounded on top by the minimum PRI allowed by the transmitter and on bottom by the signal coherency limit. The dashed line at 40 m s^{-1} marks the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

weather does not continuously span a long radial distance. Phase-code processing (e.g., Siggia 1983; Sachidananda and Zrnić 1999) 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. Therefore, we proposed an adaptive solution where, for low-elevation tilts, information from an initial long-PRI scan would be used to select a multi-PRI or phase-code signal transmission and processing on a radial-by-radial basis in the subsequent scan (Cho 2003).

Since the operational application of multi-PRI techniques had been hampered, in large part, by the difficulties involved with clutter filtering, we presented an adaptive clutter-filtering scheme for multiblock-staggered (MBS) PRI sequences that produced acceptable reflectivity, velocity, and spectral-width bias characteristics in Cho and Chornoboy (2005, hereafter Part I). In this paper we continue by presenting techniques for range-overlay protection and velocity dealiasing using MBS PRI signals.

2. Range-overlay protection

The primary mission of the TDWR is to detect weather-induced aviation hazards (especially wind shear phenomena) around airports and approach/departure corridors. Therefore, it is crucial to protect near-range observations from obscuration by range-aliased signals of more distant storm cells. Nationwide, range-overlay contamination is the top cause of missed and false detection of low-altitude wind shear phenomena (microbursts and gust fronts) by TDWR-based algorithms.

In principle, range-overlay protection using multi-PRI transmission is straightforward. In this scheme, more than one PRI is used within a dwell. Pseudorandom phase is encoded upon transmission for a Klystron or solid-state transmitter 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 up to, say, 460 km, 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. In the absence of a long-PRI scan, it is also possible to use statistical methods on the power distribution of the collected pulses within a dwell to eliminate outliers corresponding to contamination. Here we will focus on the former approach, since an initial long-PRI scan will be used for the adaptive clutter filter algorithm. This strategy will be applied to near-surface scans where range overlays and ground clutter are most problematic.

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 at range gate R_i and time index k is computed from the following equations:

$$S_{Oik} = \sum_{j=1}^{L} S(r_j) \quad \text{and} \tag{1}$$

$$r_j = R_i + \frac{c}{2} \sum_{k=1}^{j} T_{-k}, \tag{2}$$

where S is signal power. Here, L is set to the maximum integer that keeps r_j within maximum range of the long-PRI scan. This received pulse is marked clean if $S_{Oik} < LP_n$ or $S_{ik} > \alpha S_{Oik}$, where P_n is noise power and S_{ik} is first-trip signal at range gate R_i and time index k. 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 α. The value of α 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 signal-to-noise ratio (SNR). If we posit that the SOR must be at least 0 dB for minimum moments-estimate accuracy, then $\alpha = 1$. (For reference, the 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 $\alpha = (n/2)^{1/2}$, where n is the number of pulses per PRI. Since estimate error also depends on the normalized Doppler spectral width, σ_v/v_a, there is really no single value for α that is optimal for all moments estimates under all circumstances, but tests using simulated and real data show this expression to yield good results.

a. Clutter filter problem

An additional complication to range-overlay protection is the ground clutter filter (GCF). 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 following is the procedure we use to minimize interference from the GCF for each range gate. First, apply the adaptive multi-PRI GCF (Part I) if the clutter signal is present in the long-PRI scan data and all pulses are clean from range aliasing or the estimated clutter power exceeds the average range-aliased power. Then if the power removed by the GCF is nonnegligible (i.e., greater than the noise power), use the filtered time series for further processing; otherwise, use the unfiltered data. Second, if there is nonnegligible clutter power present, but the multi-PRI GCF was not applied because of range-overlay contamination, then the means are subtracted from the in-phase and quadrature (I&Q) components (clean pulses only) of each PRI set separately. In this way, the zero-Doppler power is removed from each PRI set without cross-contamination. Obviously, the clutter suppression is limited in this case, but it is an improvement over no suppression.

b. PRI set selection

What PRIs should be used for optimal RV ambiguity mitigation? Initially, we examined combinations of simple integral ratios for velocity dealiasing using the Chinese remainder theorem (CRT). However, it turned out that a clustering algorithm applied to the dealiased velocities (Trunk and Brockett 1993) performed as well as the CRT-based method (see section 3). This result then freed us to choose any combination of PRIs.

The range of usable PRIs for velocity estimation on the TDWR is limited at the upper end by the coherence criterion, $T < \lambda/(4\pi\sigma_v)$, where $\lambda \sim 5.3$ cm for TDWRs. For $\sigma_v = 4$ m s^{-1} [corresponding to median values observed in squall lines; Fang et al. (2004)] this yields $T < 1050$ μs. (The current operational RDA has a maximum PRI for velocity estimation of 938 μs.) For high-elevation scans, the shortest PRI is determined by the transmitter capability (518 μs), but for low-elevation scans the shortest PRI is limited by the FAA operational coverage requirement of 48 n mi in range (594 μs).

For maximum range-overlay protection and velocity dealiasing accuracy, the spread of PRIs should be as big as possible. Furthermore, without the need to form simple integral ratios, and with no preferred locations from which range aliasing takes place, the PRIs should be spaced evenly to avoid gaping "holes" in protection. The question then is how many PRIs to place within that spread.

Suppose there is a strong point target at true-range r_t. For T_1 with unambiguous range $r_{a1} < r_t$, there will be range-overlay contamination at $r_{f1} = r_t - r_{a1}$. If we add T_2 with unambiguous range $r_{a2} < r_t$ to the transmission sequence, then $r_{f2} = r_t - r_{a2}$ will also have range-overlay contamination. However, because r_{f1} and r_{f2} also have "clean" pulses from which the moments can be estimated, we have reduced the number of contaminated range gates from 1 to 0 by adding a PRI. In this way, increasing the number of PRIs can provide increased range-overlay protection. On the other hand, if there was strong ground clutter at r_{f1} and r_{f2}, then the interaction with the GCF will destroy the range-overlay protection capability and data from both range gates will be corrupted, thus increasing the number of "bad" range gates from 1 to 2. Furthermore, by increasing the number of PRIs, the number of pulses per PRI is necessarily decreased for a constant dwell time. This in turn increases the estimate variance coming out of each PRI set, which is detrimental to velocity dealiasing. Therefore, there is no single number of PRIs that is optimal for all situations. In the eventual adaptive processing system, we will have a number of different PRI sets available, from which the algorithm will select an optimum set for each radial based on the range-aliasing and ground-clutter distributions provided by the initial long-PRI scan. It will also preferentially protect areas

that are most important to the TDWR's mission, such as the airport and approach/departure corridors.

3. Velocity dealiasing

a. CRT versus clustering

Velocity dealiasing for Doppler weather radars utilizing PRI diversity is usually limited to two PRIs, and the technique employed is the Chinese reminder theorem (e.g., Ding et al. 1996). For a pair of PRIs with a ratio reducible to relatively prime integers, CRT can be used to prove that the difference of the respective velocity estimates maps uniquely onto an unambiguous velocity interval larger than either of their associated Nyquist intervals. In this way the effective unambiguous velocity can be extended beyond the Nyquist limit. For example, the weather channel on the operational Airport Surveillance Radar-9 (ASR-9) processes a block-staggered dual-PRI sequence (Weber 2002), while the commercial SIGMET RVP8 radar signal processing system employs an alternating-dwell dual-PRI waveform to dealias velocity across radials (SIGMET 2005). Both systems use the difference in velocity estimates from the two PRIs to compute the dealiased velocity based on rules generated by the CRT.

For our case, however, velocity estimates from more than two PRIs may be produced. In the beginning we tried to extend the traditional dual-PRI approach by choosing PRI sets that included simple integral ratios that were amenable to the CRT approach on a pair-by-pair basis. However, it became clear that a more general approach was desirable, since censoring of range-overlay contaminated PRI sets can break up those special pairs. Direct application of the CRT to multiple PRIs with a large common divisor also did not produce good results. We decided, therefore, to use the clustering algorithm (Trunk and Brockett 1993), which does not have a preference for any particular PRI relationship and can be applied in the same manner to any number of PRIs.

Suppose there are velocity estimates from m PRIs. In the clustering algorithm, for each PRI velocity estimate, all possible unfolded velocities are computed up to $\pm v_{\mathrm{MAX}}$, which is set by the user. Then all the velocity values are sorted from smallest to largest and the average squared error is computed in a sliding window of length m that is incremented across the entire list. The median value in the window with the smallest error (the "best cluster") 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. Decreasing v_{MAX} increases the dealiasing success rate as long as most of the velocity distribution lies within $\pm v_{\mathrm{MAX}}$. With the rule-based CRT technique, such a trade-off can only be realized in large, discontinuous jumps and is dependent on the particular PRI ratio used (Torres et al. 2004).

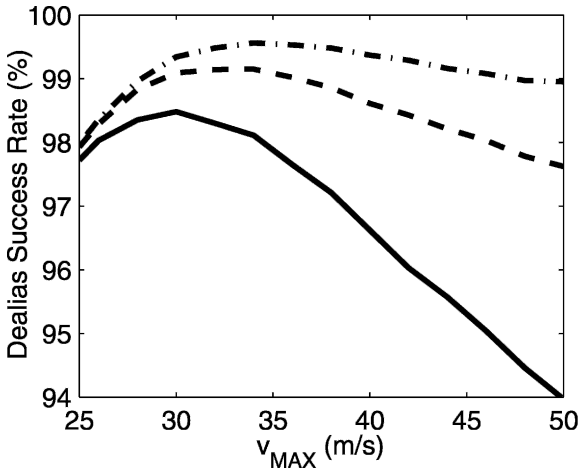

FIG. 2. Velocity dealiasing success rate vs v_{MAX} for dual-PRI using the clustering algorithm. Here, T_1 was set to 600 μs and T_2 was set to 1.3 (solid), 1.4 (dashed), and 1.5 (dashed–dotted) times T_1.

Figure 2 shows the velocity dealiasing success rate versus v_{MAX} for dual-PRI using the clustering algorithm. Here, T_1 was set to 600 μs and T_2 was set to 1.3 (solid), 1.4 (dashed), and 1.5 (dashed–dotted) times T_1. The input velocity had a zero-mean Gaussian distribution with a standard deviation of 10 m s^{-1}. On top of this, an independent Gaussian random error with standard deviation of 2 m s^{-1} was added to each PRI input velocity. The resulting velocities were aliased into the corresponding Nyquist intervals. These velocity pairs were then fed into the clustering algorithm. Dealiasing was deemed to be successful if the absolute difference between the input and dealiased velocities was smaller than the unambiguous velocity v_{a2} corresponding to T_2. Ten thousand Monte Carlo runs were averaged for each data point. Note that at small v_{MAX} the success rate increases with v_{MAX}, because there are enough velocities in the tail of the Gaussian distribution greater than v_{MAX}. At large v_{MAX}, the success rate decreases with v_{MAX}, because there are virtually no velocities in this region and increasing v_{MAX} only serves to raise the chances of false dealiasing. For reference, the percentage of input velocities for PRI T_1 that lies within $\pm v_{a1}$ ($\pm$22 m s^{-1}) is 97%. This means that velocity dealiasing success rates above this figure are an improvement over

what could be achieved using only the shorter PRI for this particular velocity distribution and error.

Let us now compare the velocity dealiasing performances of the clustering algorithm and the rule-based CRT technique (Torres et al. 2004). Figure 3 shows the velocity dealiasing success rates for dual-PRI using the clustering (solid) and CRT (dashed) algorithms versus the standard deviation of the input zero-mean Gaussian velocity distribution: $T_1 = 600$ μs and $T_2 = 900$ μs. As with Fig. 2, a Gaussian random error with standard deviation of 2 m s^{-1} was added for each PRI input velocity and aliased into the corresponding Nyquist intervals. These velocity pairs were then fed into the respective algorithms, and the outputs were the means of the dealiased pairs. Ten thousand Monte Carlo runs were averaged for each data point. Dealiasing success was measured in the same way as for Fig. 2. For the CRT method, a PRI ratio of 2:3 yields a maximum extended unambiguous velocity of $3v_{a2} = 44.5$ m s^{-1}, so v_{MAX} was also set to this value. The small difference in the performance is due to the difference in the output velocity range of the two techniques. With the CRT approach, v_1 has a maximum dealiased range of $\pm 3v_{a1}$ and v_2 has a maximum dealiased range of $\pm 3v_{a2}$, even though the dealiasing rules are only unambiguous up to $\pm 3v_{a2}$, which is narrower than $\pm 3v_{a1}$. With the clustering algorithm, $\pm v_{\mathrm{MAX}}$ is the dealiasing range for both v_1 and v_2. Thus, the clustering algorithm performs slightly better if the input velocity distribution rarely exceeds $\pm 3v_{a2}$, and the CRT algorithm works slightly better as the velocity distribution tail extends beyond $\pm 3v_{a2}$.

For number of PRIs greater than 2, we tried various simple integral ratio combinations to use with the CRT technique, but none was able to outperform the clustering algorithm. Note that Trunk and Brockett (1993) also showed the clustering algorithm to perform better than a CRT approach for range dealiasing using three PRIs. Therefore, also considering the flexibility of PRI selection and maximum dealiasing range that the clustering algorithm allows, we conclude that the clustering algorithm should generally be chosen over the CRT method, unless computational speed is the number one priority.

b. Estimation performance dependence on PRI loss

Let us now examine the effect that PRI-set censorship has on velocity dealiasing performance. For example, consider a four-PRI sequence with $T = 600$, 700, 800, and 900 μs. Velocity dealiasing can be performed on the resulting velocity estimates as long as at least two PRI sets are clean from range-overlay contamination. However, we would expect the dealiasing performance to be different for different combinations. This expectation is supported by simulation results. Figure 4 shows dealiased velocity estimation error versus different available PRI combinations. The input velocity generation procedure was the same as that used for Figs. 2 and 3, with a zero-mean Gaussian distribution of standard deviation 15 m s^{-1} and an additive Gaussian error with standard deviation of 2 m s^{-1}. Here, v_{MAX} was set to 40 m s^{-1} and 10 000 Monte Carlo runs were made for each data point. The four-digit binary number over each point indicates which PRI velocity estimates were used; for example, 1001 means that velocity estimates

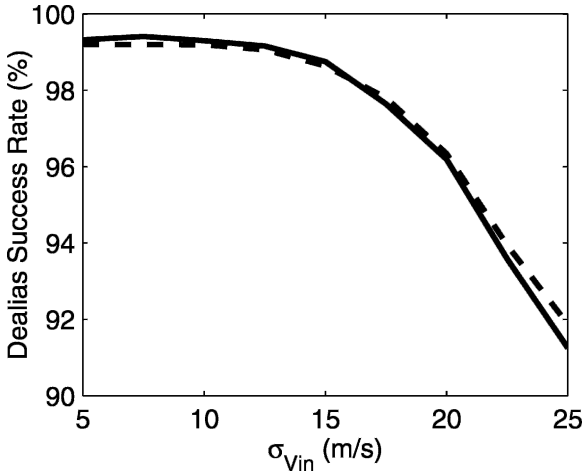

FIG. 3. The velocity dealiasing success rates for dual-PRI using the clustering (solid) and CRT (dashed) algorithms vs the standard deviation of the input zero-mean Gaussian velocity distribution: $T_1 = 600$ μs and $T_2 = 900$ μs.

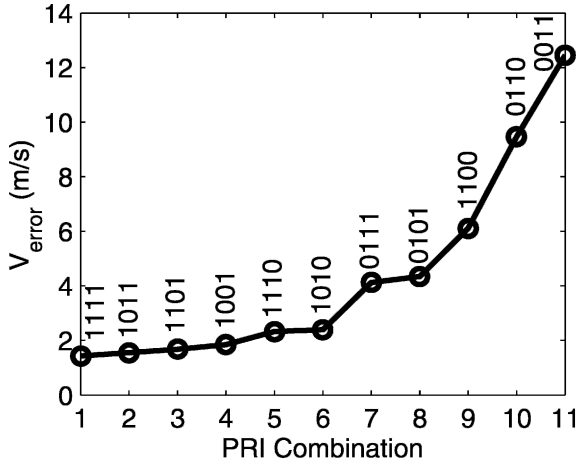

FIG. 4. Dealised velocity estimation error vs different available PRI combinations. The available PRIs were 600, 700, 800, and 900 μs. The four-digit binary number over each point indicates which PRI velocity estimates were used; e.g., 1001 means that velocity estimates from $T = 600$ and 900 μs were used to compute the dealiased velocity.

from $T = 600$ and 900 μs were used to compute the dealiased velocity. The results were sorted for the plot in order of increasing velocity estimation error.

As seen in Fig. 4, the primary PRI-combination factor in determining velocity dealiasing performance is the maximum difference in available unambiguous velocities. The unambiguous velocities corresponding to the four PRIs used here are $v_a = 22.3$, 19.1, 16.7, and 14.8 m s^{-1}. So the maximum difference for a pair is $v_{a1} - v_{a4}$ followed by $v_{a1} - v_{a3}$, $v_{a2} - v_{a4}$, $v_{a1} - v_{a2}$, $v_{a2} - v_{a3}$, and $v_{a3} - v_{a4}$. A secondary factor is the total number of PRIs available.

Not included in this analysis is the dependence of the velocity error on the PRI itself, because that is, in turn, dependent on the characteristic of the weather spectrum [e.g., Zrnić (1977), for pulse-pair velocity estimation]. In most cases smaller PRIs would yield more accurate velocity estimates, so the ordering in Fig. 4 should not be affected by this factor.

c. False dealias correction

Because of the inherent variance in real weather radar data, false velocity dealiasing is unavoidable for some fraction of cases no matter what technique is used. The human eye can usually detect such errors, because of the available contextual information in space and/or time. Similarly, automated algorithms can also detect and correct such errors based on continuity. We developed the following algorithm for a two-dimensional (2D) range–azimuth field of dealiased velocity data.

First, for every range cell, three values of velocity are stored: the undealiased (raw) velocity, the dealiased velocity, and the second-choice dealiased velocity. The second-choice dealiased velocity corresponds to the dealiased velocity cluster with the second-smallest error given by the clustering algorithm. Of these three velocities, none or two or three could be the same, depending on the value of v_{MAX} relative to the v_as. Second, for each cell a weighted median (Arce 1998) of the dealiased velocity is computed over a 2D range–azimuth grid (e.g., 3×3) centered on that cell. The weights are provided by the magnitude of the correlation coefficients of the time series at lag one [denoted as signal quality index (SQI) in the SIGMET manuals], a quantity that varies between zero and one. Because of the discontinuous nature of dealiased velocity data, the median is a more accurate measure of the background value than the mean. The weighting by the SQI diminishes the contribution to the median by less reliable data points. Third, of the three possible velocity values, the one closest to the weighted median is chosen. In this way, falsely dealiased velocity values are restored to the correct value, as long as the bad values are not clustered too densely compared to the grid dimensions used. Results from this algorithm will be presented in the following section.

4. Results using simulated weather radar data

To test the effectiveness of the multi-PRI signal processing in mitigating RV ambiguity under controlled conditions, we simulated I&Q radar data for 360° scans. Within a range–azimuth space of 460 km $\times$ 360°, at 0.15-km and 1° resolution, reflectivity, velocity, and Doppler spectral width were specified for each cell. This specification was carried out through the definition of a background (constant for all space) plus any number of compact "patches." The patches were meant to mimic, in a crude way, storm cells and ground clutter regions. These reflectivity patches were specified as 2D Gaussians, so that their location, size, and shape were determined by the mean and standard deviation in the two dimensions. Cutoff boundaries were also defined so that computation for each patch would not have to be carried out over the entire domain. The velocity and spectral width of a patch were constant, except for a special type designed to look like a microburst, in which the velocity field was perturbed as a symmetric radial divergence with the perturbation magnitude decaying as a cosine from 0° to 90° with distance. Then for each range–azimuth cell, the resulting signal strengths and radial velocities (from the background and any patches) were computed, corresponding Doppler velocity spectra were generated and sampled with the PRI sequence using the standard technique (Zrnić 1975) modified for nonuniform time sampling, and the resulting time domain series summed together if needed. For short PRIs, out-of-trip signals were added to the first-trip signal using the appropriate phase multiplier associated with the pulse phase code. (For our study, we used a pseudorandom code.) Finally, white noise was added to simulate receiver noise.

Figure 5 shows the input reflectivity (top panel) and radial velocity (bottom panel) fields for an example simulated scan. The white ring indicates the unambiguous range (90 km) associated with the FAA requirement for range coverage. The plot limit is twice this range (180 km); outside of this limit there were no patches. So this example is limited to range aliasing from the second trip. The patch in the center mimics ground clutter with zero Doppler velocity and a spectrum width of 0.25 m s^{-1}, which corresponds to clutter spectral width generated by an antenna rotation rate of 19° s^{-1} (the current operational monitor scan rate on the TDWR). The patch to the north in the first trip is a

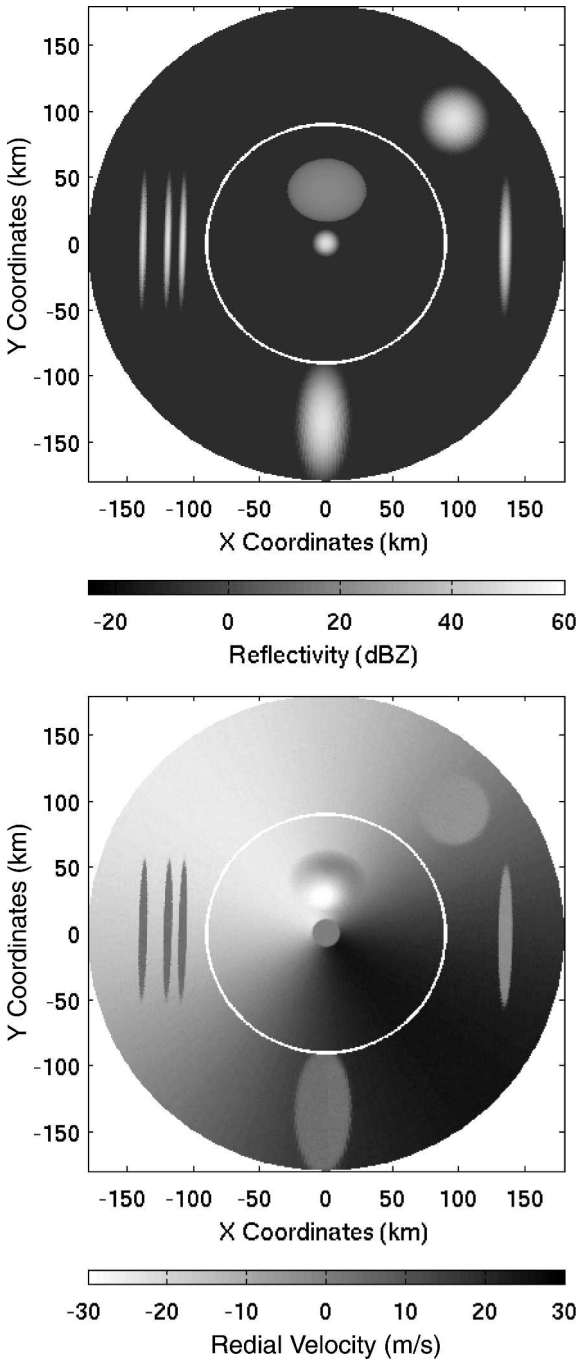

FIG. 5. The (top) input reflectivity and (bottom) radial velocity fields for the simulated weather radar data. The white ring has a radius corresponding to 90 km, which will essentially be the "first trip" range of the short-PRI scans.

"microburst" with a divergent velocity perturbation. All patches except for ground clutter have spectral widths of 4 m s^{-1}. The eastern and western second-trip patches have relatively short radial dimensions, so the multi-PRI processing should be able to filter them out from the first-trip estimates. The northeastern and southern second-trip patches, however, are a challenge because of their long radial extent.

The I&Q data for the long-PRI scan were produced with a sampling period of 3066 μs and 16 pulses per 1° dwell. This scan was then processed with an adaptive spectral-domain GCF called Gaussian model adaptive processing (GMAP; Siggia and Passarelli 2004) to yield clutter power estimates and a standard pulse-pair algorithm for signal power and spectral width estimates. These estimates provided input to the subsequent multi-PRI signal processing. However, let us first examine the results from a conventional single-PRI scan.

Figure 6 shows the results of processing simulated I&Q data generated using the moments field of Fig. 5 as input and sampled at a constant period of 598 μs and 88 pulses per 1° dwell. The corresponding unambiguous velocity was 22 m s^{-1}. Again, an adaptive GCF and standard pulse-pair processing (no phase-code processing) were used. In the reflectivity plot in Fig. 6 (top panel) one can see that the second-trip overlays contaminate the first-trip signal, which destroys the velocity estimates in those areas (bottom panel). The velocity is aliased in the northwest and southeast sectors as well as in the southern part of the microburst. The central ground clutter patch is filtered perfectly.

An example of multi-PRI processing as outlined in this paper using four PRIs in MBS format (647, 734, 821, and 908 μs at 16 pulses each per 1° dwell) is shown in Fig. 7. The plot range was set to match the Fig. 6 plots. Adaptive clutter filtering with a six-level GCF (10, 20, 30, 40, 50, and 60 dB) was applied as outlined by Part I. For both reflectivity (Fig. 6, top panel) and velocity (Fig. 6, bottom panel), the eastern and western overlays are filtered out nicely. The northeastern overlay is mostly filtered out in the reflectivity, but the velocity field indicates contamination over a wider range. This is because loss of PRI sets leads to degraded velocity estimates as demonstrated in section 3b. The southern overlay is not eliminated, because the radial patch extent is too long. This is the type of overlay that is better filtered out with phase-code processing. The clutter is removed quite well.

The velocity estimate variance in Fig. 7 is noticeably larger, especially with increasing range, compared to Fig. 6. (The SNR resulting from the -10 dBZ background reflectivity at 90 km is -2 dB.) This is the price to be paid for decreasing the number of pulses per PRI and for constantly attempting to dealias the velocity to an interval wider than the intrinsic Nyquist intervals. The false dealiasing correction procedure was developed in section 3c to combat this degradation. The result of applying this correction technique is shown in

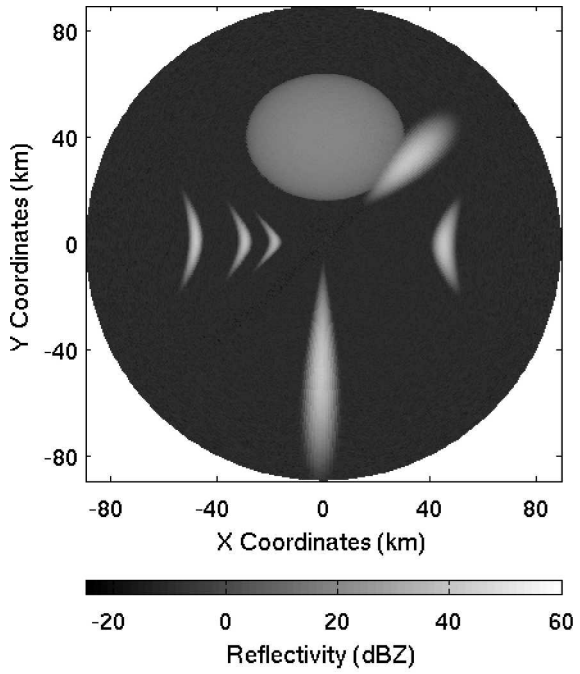

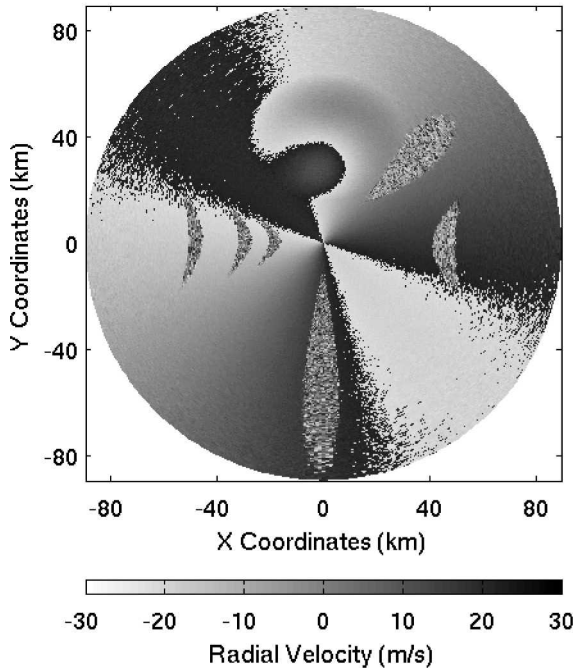

FIG. 6. (top) Reflectivity and (bottom) radial velocity estimates produced by processing simulated I&Q data corresponding to Fig. 5 input sampled at a constant PRI of 598 μs and 88 pulses per 1° dwell. Note that data are shown for only the inner (90 km) range. The corresponding unambiguous velocity was 22 m s^{-1}. An adaptive GCF and standard pulse-pair processing (no phase-code processing) were used.

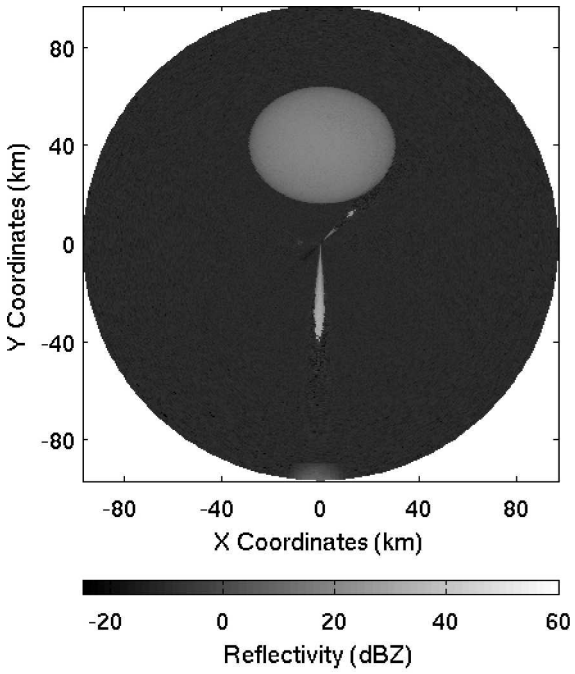

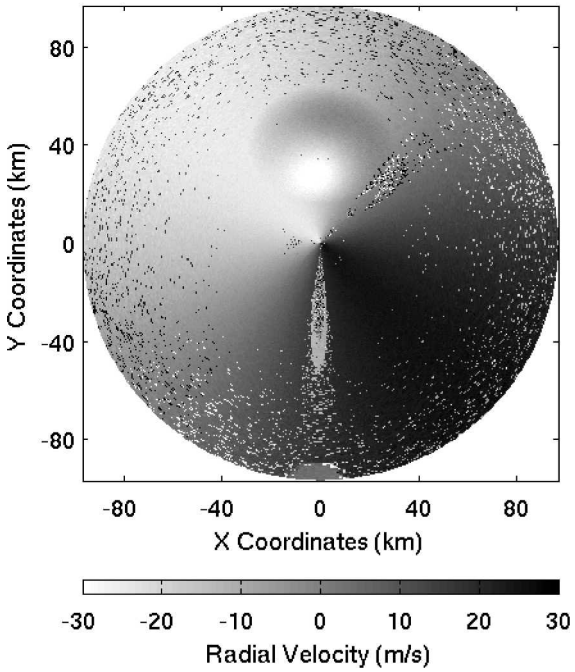

FIG. 7. (top) Reflectivity and (bottom) radial velocity estimates produced by processing simulated I&Q data corresponding to Fig. 5 input sampled with a multi-PRI sequence (647, 734, 821, and 908 μs at 16 pulses each per 1° dwell). Adaptive clutter filtering with a six-level GCF (10, 20, 30, 40, 50, and 60 dB) was used. The false velocity dealiasing correction algorithm was not applied.

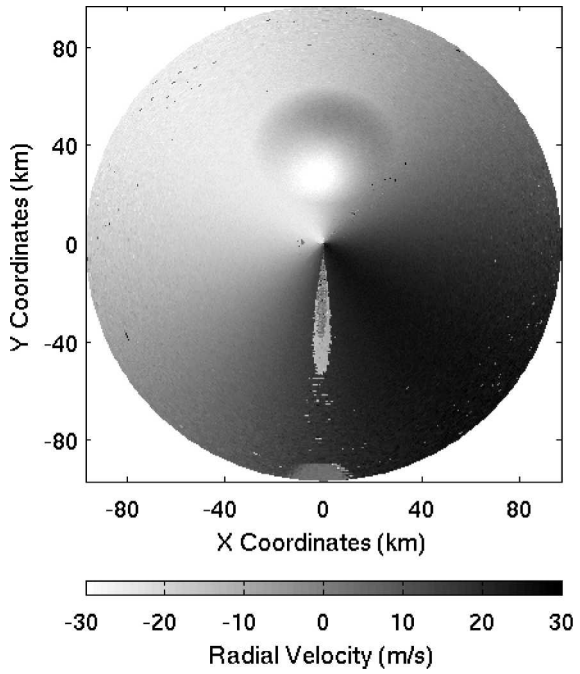

FIG. 8. Radial velocity estimate data from Fig. 7 with the false velocity dealiasing correction algorithm applied.

Fig. 8. We can see that most of the bad velocity points are restored, even in the patch caused by the northeastern overlay. If desired, a 3×3 speckle filter could be applied to remove the remaining few bad points.

5. Results using real weather radar data

Collection of I&Q data was made possible with the installation of the RDA prototype at the FAA's Program Support Facility (PSF) TDWR in Oklahoma City, Oklahoma. In this prototype, dual-processor (Intel Xeon) servers running Linux perform both the signal processing and system control functions. (The server has since been upgraded to a quad-processor unit.) The system control computer houses a SIGMET RVP8, which provides the digital receiver, digital waveform shaping, and timing functions in three PCI cards each with several field programmable gate array (FPGA) chips. A combination of interrupt-driven software and FPGA code allows the system to change PRI and phase coding on a radial-by-radial basis, a key feature for the planned adaptive signal transmission and processing algorithms. For these tests, the I&Q data were merely recorded and were processed later offline.

The TDWR transmits a peak power of 250 kW. The antenna beamwidth is 0.55°, and the pulse length is 1.1 μs. The PSF TDWR operates at a frequency of 5.62 GHz. Although the operational system samples range at 150-m resolution, the first version of the RDA prototype used in this study sampled at 125-m resolution. (This has been updated to 150 m in the current version.) For further details on the TDWR, see Michelson et al. (1990).

a. Range-overlay protection

The first set of scans was taken on 17 March 2003 starting at 2040 UTC, while convective storm cells were active in the vicinity. The scan elevation angle was 0.3° with an antenna rotation rate of 21.6° s^{-1}. Figure 9 shows the reflectivity fields computed from the long-PRI (3.06 ms) scan for the full 460-km radius (top panel) and zoomed in to a radius of 77 km (bottom panel). Very strong scattering targets exist at many different ranges and at azimuths, thus making this a challenging case for range-overlay protection. In some azimuths there are multiple trips aliasing into the first trip. The resulting near-range reflectivity suffers much contamination as can be seen in the single-PRI (667 μs) scan (Fig. 10, top panel). The MBS (518, 578, 638, 698, 758, 818, 878, 938 μs $\times$ 16 pulses each) scan reflectivity (Fig. 10, bottom panel), however, looks very similar to the "truth" provided by the long-PRI scan (Fig. 9, bottom panel). There is still some unfiltered overlay power in the south-southeast and west-northwest sectors, but otherwise the differences are small. These problem areas are associated with out-of-trip patches that have relatively long continuous radial extent (Fig. 9, top panel). As expected these areas are protected better in a constant-PRI phase-code processed scan (not shown), which supports our plan for an adaptive signal transmission and processing scheme.

Here we insert a note about the clutter filtering. As with the simulated data processing, we used the GMAP GCF for the constant-PRI scans. For the multi-PRI scans, however, the full, multilevel adaptive GCF as presented in Part I was not applied. In the initial prototype version, the PRI transmission sequence was not synched to the beginning of each 1° azimuthal dwell boundary. Consequently, for each dwell, different filter coefficient matrices had to be defined for each possible sequence permutation, in this case 64. Therefore, we only generated one set of coefficient matrices at a 60-dB suppression. We used the estimated clutter power from the long-PRI scan to turn this GCF on or off for a given cell. Furthermore, the initial SIGMET receiver had limited dynamic range compared to the legacy operational receiver, which led to saturation on very strong clutter targets. The current version has a dynamic range slightly better than the legacy system. The present RDA prototype is also capable of initializing the PRI sequence at every 1° azimuthal dwell bound-

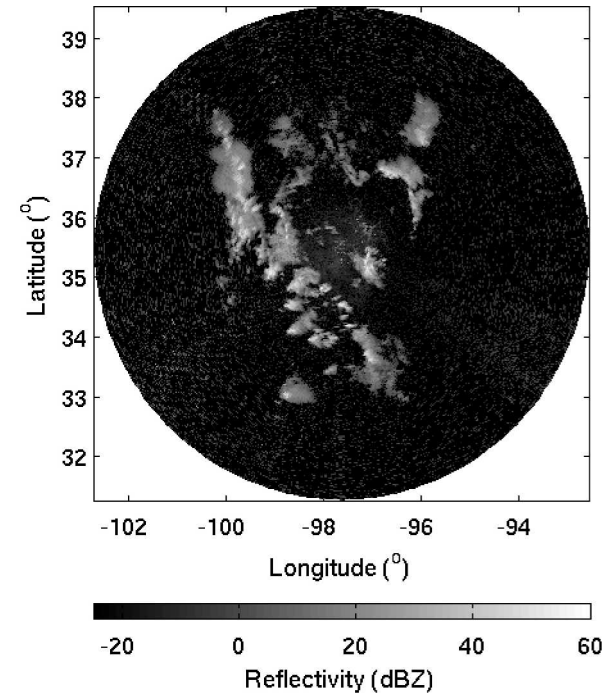

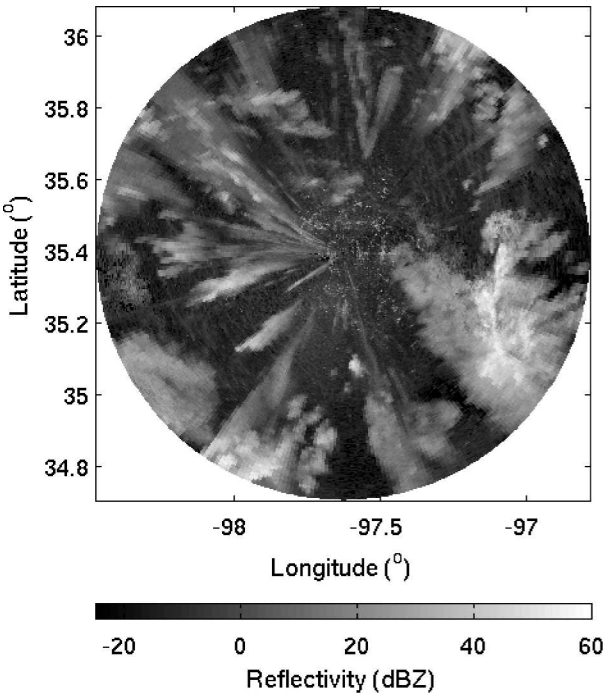

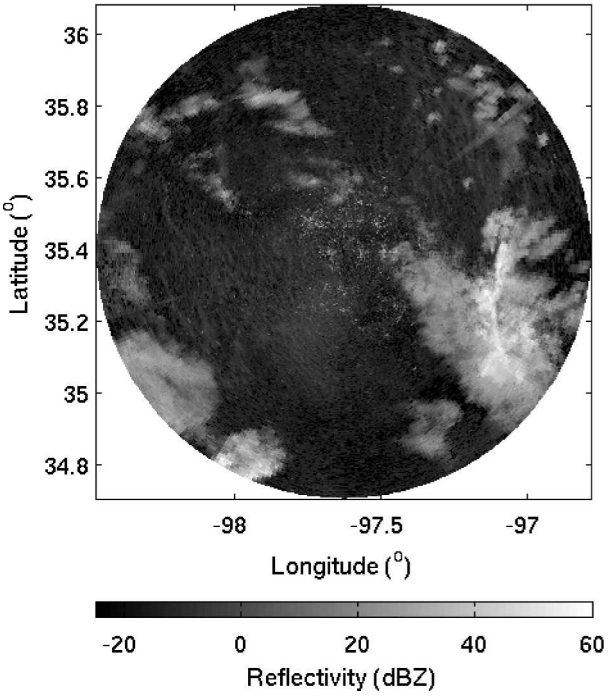

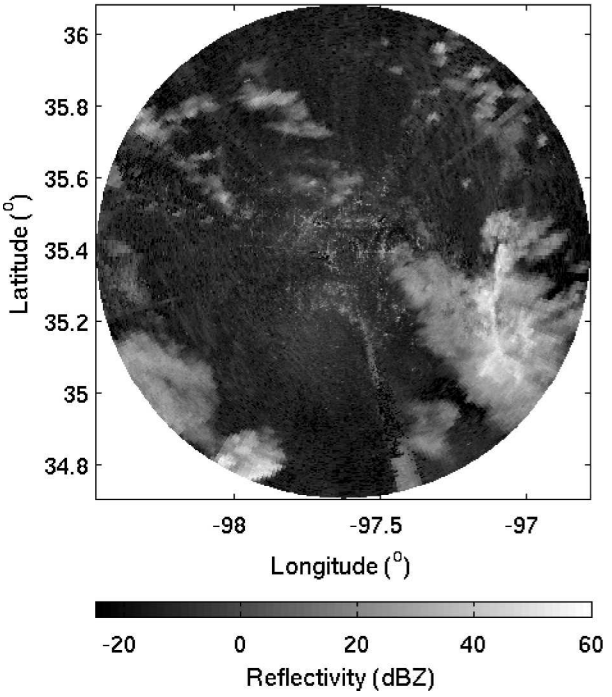

FIG. 9. Reflectivity estimates from a long-PRI (3.06 ms) scan for (top) the full 460-km radius and (bottom) zoomed in to a radius of 77 km. The dataset was collected with the PSF TDWR in Oklahoma City on 17 Mar 2003 starting at 2040 UTC using our initial RDA prototype. The scan elevation angle was 0.3° with an antenna rotation rate of 21.6° s^{-1}. An adaptive GCF and standard pulse-pair processing were used.

FIG. 10. Reflectivity estimates using (top) a single PRI of 667 μs and (bottom) a multi-PRI sequence of 518, 578, 638, 698, 758, 818, 878, and 938 μs at 16 pulses each. The other parameters are as given for Fig. 9. For the single-PRI scan, adaptive GCF and standard pulse-pair processing (no phase-code processing) were used.

ary, allowing for easy implementation of a multilevel adaptive GCF for multi-PRI signals. Therefore, the overall clutter filtering performance with the present prototype is better than what is shown here. We used the older data, because they contained the best examples of range aliasing.

The crucial parameter that must be protected from range-overlay contamination is not reflectivity, which is essentially available from the long-PRI scan, but velocity. Figure 11 shows the comparison for velocity between the single-PRI scan (top panel) and the multi-PRI scan (bottom panel). Although there is no truth available for comparison, it is clear that the multi-PRI scan has protected many of the areas that were corrupted by range aliasing in the single-PRI case.

The single-PRI scan shown here is representative of estimates produced by the current operational TDWR. Even though the TDWR utilizes an adaptive PRI selection algorithm at the lowest-elevation scan in an attempt to minimize range-overlay obscuration in certain key areas, such as the airport and approach/departure corridors, it is still a single PRI for the entire scan (Crocker 1988). Therefore, the leverage to protect all azimuth and range cells is very limited.

b. Velocity dealiasing

Since the velocities in the previous example were not strong enough to test the velocity dealiasing algorithm, we present another case. The following scans were taken on 3 April 2003 starting at 1820 UTC at an elevation angle of 2.6°. No significant range aliasing was present. The velocity plots are shown in Fig. 12 for the single-PRI (598 μs) scan (top panel) and the multi-PRI (same sequence as before) scan (bottom panel). The unambiguous velocity of the single-PRI scan was 22 m s^{-1}, which was clearly exceeded over significant areas. The multi-PRI dealiasing algorithm does an excellent job of restoring the velocity field. The problem regions (for both scans) to the northwest and to the southwest at close range are areas of very low SNR.

The current operational TDWR also has a velocity dealiasing procedure, but it uses two consecutive scans at the same elevation with two different PRIs (Wieler and Hu 1993). Being able to perform the dealiasing within a single scan will save time and allow a faster volume scan or reduced estimate variance via a slower scan rate. Also, of course, the multi-PRI processing simultaneously provides vastly improved range-overlay protection.

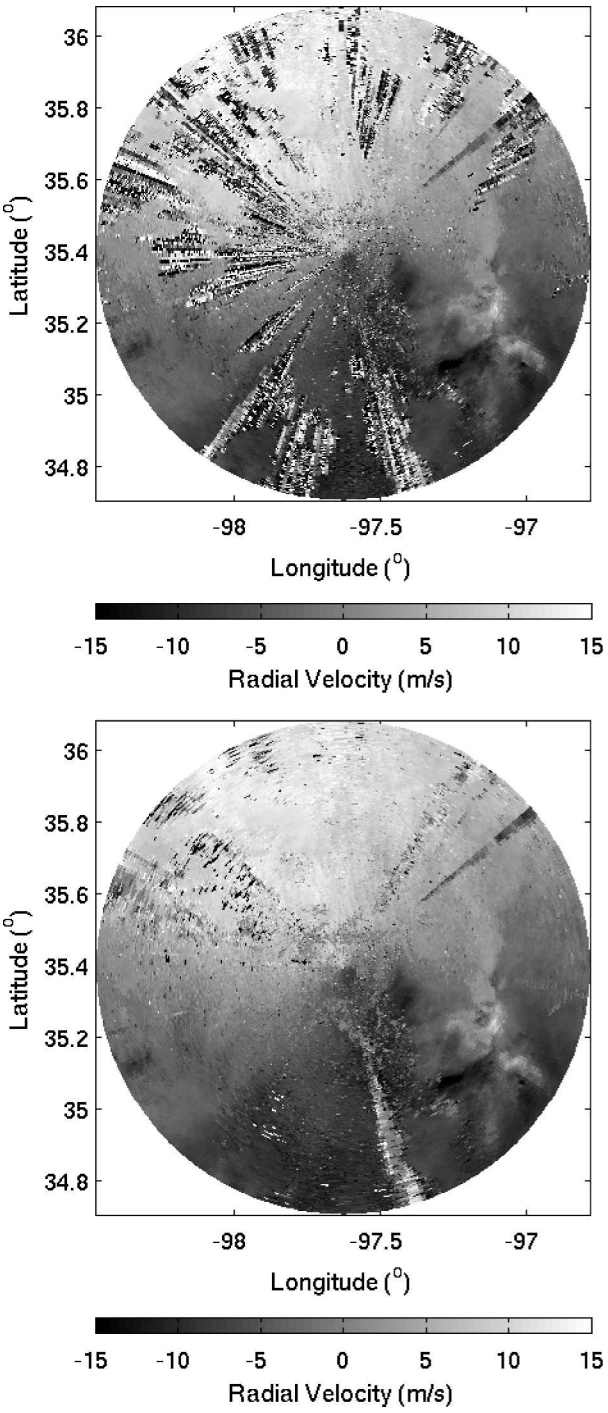

FIG. 11. Velocity estimates for the (top) single- and (bottom) multi-PRI scans. Parameters are the same as for Fig. 10.

6. Conclusions

Here we have presented procedures for simultaneously providing range-overlay protection and velocity dealiasing using multi-PRI signal transmission and processing for weather radars. For velocity dealiasing we showed that the clustering algorithm is more flexible, more general, and performs just as well as the Chinese remainder theorem approach that is typically

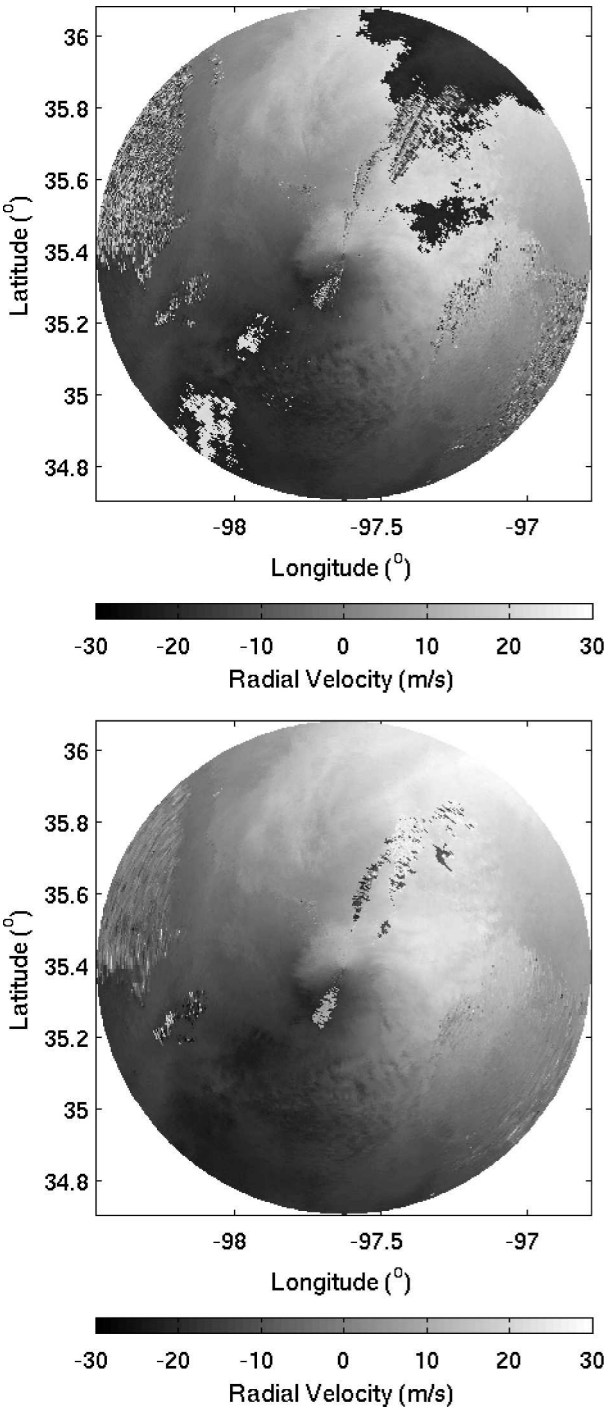

FIG. 12. Velocity estimates for (top) a single PRI of 598 μs and (bottom) a multi-PRI sequence as given in the caption to Fig. 10. The unambiguous velocity of the single-PRI scan was 22 m s^{-1}. For the single-PRI scan, adaptive GCF and standard pulse-pair processing (no phase-code processing) were used. The dataset was collected with the PSF TDWR in Oklahoma City on 3 Apr 2003 starting at 1820 UTC using our initial RDA prototype. The scan elevation angle was 2.6° with an antenna rotation rate of 21.6° s^{-1}.

used with weather radars. This RV ambiguity mitigation scheme was demonstrated to be effective with simulated and real weather radar data. The clutter filtering capability is limited in gates where range-aliased signals need to be removed, but this is also true for phase-code processing techniques. Real-time implementation of these algorithms in the enhanced TDWR RDA is under way, and extensive testing will be conducted under diverse weather conditions. Combined with the adaptive MBS clutter filter presented in Part I, this technique will be used within the larger context of an adaptive signal transmission and processing scheme in which the optimal mode will be selected on a dwell-by-dwell basis. In this scheme, a number of different MBS sequences will be available for selection, as well as constant-PRI phase-code processing that changes PRIs every dwell (for interdwell velocity dealiasing). The latter mode can provide range-overlay protection under conditions in which the multi-PRI technique fails, that is, when the overlay has a long, continuous radial extent. The mode selection algorithm and corresponding processing results will be presented in a future paper.

Acknowledgments. This work was sponsored by the Federal Aviation Administration under Air Force Contract F19628–00-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the author and are not necessarily endorsed by the U.S. government. I would like to acknowledge the efforts of the TDWR RDA upgrade team (Gabe Elkin, Nathan Parker, and Walter Heath) and management (Mark Weber) in the development of the engineering prototype used in this study. The able assistance provided by the personnel at the FAA PSF in Oklahoma City (Robert Schaefer, Rex Crissup, Roger Smith, Rob Hyam, Peter Smith, and Paul Biron) is also appreciated. Finally, thanks to Jeff Simpson, Andrew Hickman, and Stacey Uy for their contribution to the development of the weather radar data simulator under the MIT Lincoln Laboratory–Worcester Polytechnic Institute Major Qualifying Project program.

REFERENCES

Arce, G. R., 1998: A general weighted median filter structure admitting negative weights. *IEEE Trans. Signal Processing,* **46,** 3195–3205.

Cho, J. Y. N., 2003: Evaluation of TDWR range–ambiguity mitigation techniques. Project Rep. ATC-310, MIT Lincoln Laboratory, Lexington, MA, 47 pp.

——, 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.

——, 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.

——, N. G. Parker, and G. R. Elkin, 2004: Improved range–velocity ambiguity mitigation for the Terminal Doppler Weather Radar. Preprints, *11th Conf. on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA, Amer. Meteor. Soc., CD-ROM, 5.4.

Crocker, S. C., 1988: TDWR PRF selection criteria. Project Rep. ATC-147, DOT/FAA/PM-87-25, MIT Lincoln Laboratory, Lexington, MA, 57 pp.

Ding, C., D. Pei, and A. Salomaa, 1996: *Chinese Remainder Theorem: Applications in Computing, Coding, and Cryptography*. World Scientific, 213 pp.

Doviak, R. J., and D. S. Zrnić, 1993: *Doppler Radar and Weather Observations*. Academic Press, 562 pp.

Fang, M., R. J. Doviak, and V. Melnikov, 2004: Spectrum width measured by WSR-88D: Error sources and statistics of various weather phenomena. *J. Atmos. Oceanic Technol.,* **21,** 888–904.

Michelson, M., W. W. Shrader, and J. G. Wieler, 1990: Terminal Doppler Weather Radar. *Microwave J.,* **33,** 139–148.

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.

——, 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.

SIGMET, 2005: RVP8 digital IF receiver/Doppler signal processor user's manual. SIGMET, 451 pp.

Torres, S. M., Y. F. Dubel, and D. S. Zrnić, 2004: Design, implementation, and demonstration of a staggered PRT algorithm for the WSR-88D. *J. Atmos. Oceanic Technol.,* **21,** 1389–1399.

Trunk, G., and S. Brockett, 1993: Range and velocity ambiguity reduction. Preprints, *1993 IEEE National Radar Conf.*, Lynnfield, MA, IEEE, 146–149.

Weber, M. E., 2002: ASR-9 Weather Systems Processor (WSP) signal processing algorithms. Project Rep. ATC-255, MIT Lincoln Laboratory, Lexington, MA, 53 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., 1975: Simulation of weatherlike Doppler spectra and signals. *J. Appl. Meteor.,* **14,** 619–620.

——, 1977: Spectral moment estimates from correlated pulse pairs. *IEEE Trans. Aerosp. Electron. Syst.,* **13,** 344–354.