


**Project Report
ATC-310**

Evaluation of TDWR Range-Velocity Ambiguity Mitigation Techniques

J. Y. N. Cho

4 April 2003

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LExINGTON, MASSACHUSETTS



Prepared for the Federal Aviation Administration,
Washington, D.C. 20591

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

This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof.

1. Report No. ATC-310	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle Evaluation of TDWR Range-Velocity Ambiguity Mitigation Techniques		5. Report Date 4 April 2003	
		6. Performing Organization Code	
7. Author(s) John Y.N. Cho		8. Performing Organization Report No. ATC-310	
9. Performing Organization Name and Address MIT Lincoln Laboratory 244 Wood Street Lexington, MA 02420-9108		10. Work Unit No. (TRAVIS)	
		11. Contract or Grant No.	
12. Sponsoring Agency Name and Address Department of Transportation Federal Aviation Administration 800 Independence Ave., S.W. Washington, DC 20591		13. Type of Report and Period Covered Project Report	
		14. Sponsoring Agency Code	
15. Supplementary Notes This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology, under Air Force Contract F19628-00-C-0002.			
16. Abstract Range and velocity ambiguities pose significant data quality challenges for the Terminal Doppler Weather Radar (TDWR). For typical pulse repetition frequencies (PRFs) of 1-2 kHz, the radar is subject to both range-ambiguous precipitation returns and velocity aliasing. Experience shows that these are a major contributor to failures of the system's wind shear detection algorithms. Here we evaluate the degree of mitigation offered by existing phase diversity methods to these problems. Using optimized processing techniques, we analyze the performance of two particular phase codes that are best suited for application to TDWRs—random and SZ(8/64) [Sachidananda and Zrnic', 1999]—in the protection of weak-trip power, velocity, and spectral width estimates. Results from both simulated and real weather data indicate that the SZ(8/64) code generally outperforms the random code, except for protection of 1 st trip from 5 th trip interference. However, the SZ code estimates require <i>a priori</i> knowledge of out-of-trip spectral widths for censoring. This information cannot be provided adequately by a separate scan with a Pulse Repetition Frequency (PRF) low enough to unambiguously cover the entire range of detectable weather, because then the upper limit of measurable spectral width is only about 2 m/s. For this reason we conclude that SZ phase codes are not appropriate for TDWR use. For velocity ambiguity resolution, the random phase code could be transmitted at two PRFs on alternating dwells. Assuming the velocity changes little between two consecutive dwells, a Chinese remainder type of approach can be used to dealias the velocities. Strong ground clutter at close range, however, disables this scheme for gates at the beginning of the 2 nd trip of the higher PRF. We offer an alternative scheme for range-velocity ambiguity mitigation: Multistaggered Pulse Processing (MSPP). Yielding excellent velocity dealiasing capabilities, the MSPP method should also provide protection from patchy, small-scale out-of-trip weather. To obtain maximum performance in both range and velocity dealiasing, we suggest that information from the initial low-PRF scan be used to decide the best waveform to transmit in the following scan—random phase code with alternating-dwell PRFs or MSPP. Such an adaptive approach presages future developments in weather radar, for example electronically scanned arrays allow selective probing of relevant weather events.			
17. Key Words		18. Distribution Statement This document is available to the public through the National Technical Information Service, Springfield, VA 22161.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 60	22. Price

ABSTRACT

Range and velocity ambiguities pose significant data quality challenges for the Terminal Doppler Weather Radar (TDWR). For typical pulse repetition frequencies (PRFs) of 1-2 kHz, the radar is subject to both range-ambiguous precipitation returns and velocity aliasing. Experience shows that these are a major contributor to failures of the system's wind shear detection algorithms. Here we evaluate the degree of mitigation offered by existing phase diversity methods to these problems. Using optimized processing techniques, we analyze the performance of two particular phase codes that are best suited for application to TDWRs—random and SZ(8/64) [*Sachidananda and Zrnic*, 1999]—in the protection of weak-trip power, velocity, and spectral width estimates. Results from both simulated and real weather data indicate that the SZ(8/64) code generally outperforms the random code, except for protection of 1st trip from 5th trip interference. However, the SZ code estimates require *a priori* knowledge of out-of-trip spectral widths for censoring. This information cannot be provided adequately by a separate scan with a Pulse Repetition Frequency (PRF) low enough to unambiguously cover the entire range of detectable weather, because then the upper limit of measurable spectral width is only about 2 m/s. For this reason we conclude that SZ phase codes are not appropriate for TDWR use.

For velocity ambiguity resolution, the random phase code could be transmitted at two PRFs on alternating dwells. Assuming the velocity changes little between two consecutive dwells, a Chinese remainder type of approach can be used to dealias the velocities. Strong ground clutter at close range, however, disables this scheme for gates at the beginning of the 2nd trip of the higher PRF.

We offer an alternative scheme for range-velocity ambiguity mitigation: Multistaggered Pulse Processing (MSPP). Yielding excellent velocity dealiasing capabilities, the MSPP method should also provide protection from patchy, small-scale out-of-trip weather.

To obtain maximum performance in both range and velocity dealiasing, we suggest that information from the initial low-PRF scan be used to decide the best waveform to transmit in the following scan—random phase code with alternating-dwell PRFs or MSPP. Such an adaptive approach presages future developments in weather radar, for example electronically scanned arrays allow selective probing of relevant weather events.

TABLE OF CONTENTS

	Page
ABSTRACT	iii
List of Illustrations	vii
List of Tables	xi
1. INTRODUCTION	1
2. PHASE CODE PROCESSING OF SIMULATED DATA	5
2.1 Data Simulation	5
2.2 Random Phase Code Processing	5
2.3 SZ Phase Code Processing	6
2.4 Comparison of Random vs. SZ Phase Code Processing Results	17
3. PHASE CODE ISSUES ASSOCIATED WITH PROCESSING ACTUAL DATA	19
3.1 Multiple Trip Overlay	19
3.2 Data Censoring	22
3.3 Ground Clutter Filtering	22
3.4 Phase Errors	22
3.5 Code Truncation	23
4. PHASE CODE PROCESSING OF ACTUAL DATA	25
4.1 Reflectivity Estimate Comparisons	26
4.2 Velocity Estimate Comparisons	30
4.3 Spectral Width Estimate Comparisons	30
4.4 Data Censoring	32
4.5 Reasons for Recovery Failures	34
4.6 Multiple Trip Recovery	37
5. MULTISTAGGERED PULSE PROCESSING (MSPP)	39
6. SUMMARY	43
GLOSSARY	45
References	47

LIST OF ILLUSTRATIONS

Figure No.	Page
Figure 1. Beam height versus elevation angle and range. The dashed lines show multiple trip ranges for PRF = 1930 Hz.	1
Figure 2. Illustration of random-phase code processing using simulated data. PRF = 1930 Hz, M = 64.	8
Figure 3. Illustration of SZ-phase code processing using simulated data.	12
Figure 4. Weak-trip velocity estimation error statistics (standard deviation) for SZ(8/64). 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval.	13
Figure 5. Weak-trip velocity estimation error statistics (standard deviation) for SZ(8/64) and random phase codes. 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval.	14
Figure 6. Weak-trip power estimation error statistics (standard deviation) for SZ(8/64) and random phase codes. 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval.	15
Figure 7. Weak-trip spectral width estimation error statistics (standard deviation) for SZ(8/64) and random phase codes. 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval.	16
Figure 8. Illustration of the spectral replication properties of the SZ(8/64) phase code using simulated data. The input parameters are $p_1 = 15$ dB, $v_1 = 15$ m/s, $w_1 = 1$ m/s, $p_2 = 0$ dB, $v_2 = -15$ m/s, $w_3 = 1$ m/s. There is no added noise power.	20
Figure 9. Weak-trip velocity estimation error statistics (standard deviation) for SZ(8/64) (for various trip differences) and random phase codes. The strong-trip numbers are indicated in the legend. 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval. The input weak-trip spectral width was 4 m/s and the weak-trip SNR was 20 dB.	21
Figure 10. Reflectivity estimates for (a) PRF = 500 Hz, full 300-km range displayed, (b) PRF = 500 Hz, only first 100 km shown, (c) PRF = 1500 Hz, random phase coding, no phase code processing, (d) PRF = 1500 Hz, random phase coding, phase code processing with low-PRF data, (e) PRF = 1500 Hz, random phase coding, phase code processing without low-PRF data, and (f) PRF = 1500 Hz, SZ(8/64) phase code processing.	27

LIST OF ILLUSTRATIONS (Continued)

Figure No.	Page
Figure 11. Velocity estimates for (a) random phase coding, no phase code processing, (b) SZ(8/64) phase code processing, (c) random phase coding, phase code processing with low-PRF data, weak-trip recovery gates only, (d) random phase coding, phase code processing without low-PRF data, weak-trip recovery gates only, and (e) SZ(8/64) phase code processing, weak-trip recovery gates only.	29
Figure 12. Spectral width estimates for (a) the low-PRF scan, (b) random phase code processing, (c) low-PRF data, weak-trip recovery gates only, (d) random phase processing, weak-trip recovery gates only, and (e) SZ(8/64) phase code processing, weak-trip recovery gates only.	31
Figure 13. Random phase code processing censored estimates of (a) reflectivity using low-PRF data, (b) reflectivity without using low-PRF data, (c) velocity using low-PRF data, weak-trip recovery gates only, (d) velocity without using low-PRF data, weak-trip recovery gates only, (e) spectral width using low-PRF data, weak-trip recovery gates only, and (f) spectral width without using low-PRF data, weak-trip recovery gates only.	33
Figure 14. SZ(8/64) phase code processing censored estimates (weak-trip recovery gates only) for (a) velocity and (b) spectral width.	34
Figure 15. Velocity estimation error statistics (standard deviation) for SZ(8/64) versus variation in strong-trip (blue) and weak-trip (red) spectral width. 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval.	36
Figure 16. Spectral width estimates versus input spectral widths. The straight line indicates perfect agreement. 1000 simulations were run per data point and the results averaged. The input velocity values were varied randomly across the Nyquist interval. 15 points were used for the computation, which corresponds to a $21.6^\circ/\text{s}$ antenna rotation rate for a dwell of 1° .	36
Figure 17. (a) Power ratio of strong to weak trip, (b) 1 st trip SNR, (c) Strong-trip spectral width, (d) Weak-trip spectral width.	37
Figure 18. Three-trip SZ(8/64) phase code processing censored estimates for (a) reflectivity and (b) velocity.	38
Figure 19. Illustration of MSPP velocity dealiasing.	41

LIST OF ILLUSTRATIONS (Continued)

Figure No.	Page
Figure 20. Magnitude response (top) and average phase difference (bottom) versus frequency for a multi-PRI clutter filter designed using the Chornoboy algorithm. The waveform parameters are given in the caption to Table 2. In the frequency axis, -0.5 to 0.5 corresponds to the full Nyquist span of the longest PRI. In this case, -1.5 to 1.5 corresponds to -43 to 43 m/s in Doppler velocity space.	42

LIST OF TABLES

Table No.		Page
1	Standard deviations of power estimate errors in dB.	28
2	Velocity dealiasing success rates (%) for $T_0 = 518 \mu\text{s}$, $\delta T = 60 \mu\text{s}$, $m = 8$, and $N = 8$.	41

1. INTRODUCTION

The primary mission of the Terminal Doppler Weather Radar (TDWR) is the continuous surveillance of weather in and around the terminal area. This mission, however, is hampered by range and velocity ambiguities in the radar data. The former results when significant weather exists outside the unambiguous range interval (“out-of-trip” weather) and its signal folds into the first trip, thereby biasing or obscuring the desired close-range signal. The latter problem occurs when the radial Doppler velocity exceeds and wraps around the Nyquist interval. Because of the given operating parameters of the TDWR, such as wavelength and maximum pulse repetition frequency (PRF), it is impossible to fully resolve both ambiguities simultaneously.

This dilemma can be summarized mathematically as follows: $r_a v_a = c\lambda/8$, where r_a is the unambiguous range, v_a is the unambiguous velocity, c is the speed of light, and λ is the radar wavelength. In other words, extending unambiguous range results in decreased unambiguous velocity interval and vice versa. Because the TDWR operates at a 5-cm wavelength, its unambiguous range-velocity product is only half that of a 10-cm radar like the Next Generation Weather Radar (NEXRAD).

From geometric and atmospheric considerations, the required r_a is about 460 km for the lowest elevation scans. This is illustrated in Figure 1. Because of the curvature of the Earth and atmospheric refraction, a 0° elevation beam passes out of the mid-latitude troposphere (where precipitation echoes may be present) at about 460 km in range. However, from the above range-velocity relation, it is clear that such a large r_a will keep v_a at an unacceptably low 4 m/s.




Figure 1. Beam height versus elevation angle and range. The dashed lines show multiple trip ranges for PRF = 1930 Hz.

Even if range folding was not a problem, transmitter limitations impose a maximum PRF of 1930 Hz, which, through the relation, $v_a = \lambda (\text{PRF})/4$, restricts the unambiguous velocity interval to no more than ± 26 m/s. This is still a concern, since wind speeds can exceed these values.

Currently, the procedure used to minimize ambiguity problems with TDWR is to first conduct a near-surface scan at a low PRF (326 Hz), which covers the needed 460-km range. Then the reflectivity data from this scan is used to choose two values from 18 PRFs in the range 1066 to 1672 Hz. The first PRF is chosen to minimize range obscuration that could degrade wind shear detection algorithm performance around the airport, while the second PRF is chosen on the basis of best velocity dealiasing performance when used in conjunction with the first PRF. Two subsequent surface scans use the selected PRFs; see *Crocker* [1988] and *Wieler and Hu* [1993] for details of the procedure. Note that the upper PRF limit of 1672 Hz is set by the Federal Aviation Administration (FAA) requirement to have first-trip coverage out to 48 nmi, while the lower limit involves not only the unambiguous velocity range, but also the signal coherence criterion, i.e., the accuracy of the velocity estimate degrades as the PRF is lowered. Note also that although 18 PRFs are available, because the velocity dealiasing works much better for certain PRF ratios (integral ratios such as 2:3 or 3:4), the range of options is somewhat more restricted.

The present technique helps to cut down range obscuration events, but does not eliminate them. In cases where obscuration is unavoidable, those gates affected are edited out, which can cause wind shear to go undetected. Also the current algorithm requires three consecutive surface scans, which takes up a significant portion of the scan cycle. It would clearly be desirable to improve on the range-velocity ambiguity mitigation scheme and to reduce the amount of time spent on obtaining the low-elevation velocity data.

One approach that has not yet been used on the TDWR 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, notch out 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, which we will discuss later.

The purpose of this report, then, is to present the results of tests conducted on phase code processing techniques. Both random and periodic codes are examined, as well as the variations on the signal processing methods. Simulated data are used to compare the different codes and methods, and data taken

with the FAA's Program Support Facility (PSF) TDWR in Oklahoma City are used for preliminary real-life comparisons.

Finally, an altogether different technique using multiple PRFs for range-velocity ambiguity mitigation is discussed. Simulation results for velocity dealiasing using this method will be presented.

2. PHASE CODE PROCESSING OF SIMULATED DATA

2.1 DATA SIMULATION

Weather spectra for a single range gate were produced following *Zrnić* [1975]. These were then sampled with the specified waveform. To simulate weather signals coming from different trips, we modulated each time series by the appropriately shifted phase code sequence and then summed them. For example, if the first-trip weather spectrum was contained in the complex time series s_{1k} and the second-trip was in s_{2k} , then the resulting I/Q signal would be $s_{1k}\exp(j\phi_k) + s_{2k}\exp(j\phi_{k-1})$, where ϕ_k is the phase code sequence. The parameters that could be set for each simulated signal were power, velocity, spectral width (all for each trip), and noise. Unless otherwise noted, the PRF was fixed at 1930 Hz, 64 points were generated per run, and the transmitter frequency was assumed to be 5.6 GHz. For the test runs discussed below, we generated only two-trip data, with the stronger weather signal in the second trip.

2.2 RANDOM PHASE CODE PROCESSING

The procedure we adopted to decode and process the random-phase-coded data is similar to the one outlined by *Sachidananda* [1997]. The main difference is that we decided to use an adaptive rather than a fixed-width notch filter to eliminate the stronger-trip signal for velocity estimation. Overall this approach yielded better performance. Below is an outline of the algorithm. $M = 64$ is the length of the simulated I/Q series and note again that the stronger weather signal is in the second trip for our test runs. If the stronger trip is put into the first trip, simply reassign the “w” subscript to trip 2 quantities and the “s” subscript to trip1 quantities.

1. Cohere to first and second trips to produce series s_w (weak) and s_s (strong).
2. Autocovariance process to get p_s , v_s , w_s .
3. Compute strong-trip spectrum with Hann windowing.
4. Notch out $3M/4$ spectral bins centered on v_s , calculate power in remaining bins, and multiply by 4 to get p_w .
5. Calculate power ratio $R = 10 \log(p_s/p_w)$.
6. If $R < 25$ dB, correct strong-trip power estimate: $p_s = p_s - p_w$.
7. Get noise floor of unnotched spectrum, smooth the spectrum, start from v_s , and find where spectrum drops to noise floor. Notch out this apparently coherent signal part of the spectrum, but do not exceed $M/2$ bins.
8. Transform back to time domain and cohere to the weak trip.
9. Compute weak-trip velocity, v_w .

10. Transform to spectral domain.
11. Get noise floor of spectrum, smooth the spectrum, start from ν_w , and find where spectrum drops to noise floor. Keep this apparently coherent signal part of the spectrum and zero out the noise bins.
12. Compute weak-trip spectral width, w_w .

Commentary on some of the steps is in order. For real-life data, ground clutter may be present, which needs to be filtered. This can be done after the time series is cohered to the first trip. Then the filtered stream is cohered to the second trip for further processing. In step 3, a window with low spectral sidelobes is required, because otherwise the notch filter will not be effective in eliminating the strong-trip signal. In the notch filter process, there is a trade off between the need to get rid of as much of the strong-trip signal as possible (widest possible notch) versus the need to retain as much of the weak-trip information as possible (narrowest possible notch). For weak-trip power estimation, the emphasis is placed on the former requirement, because in principle the power can be estimated from just a few spectral bins since the uncohered spectrum is white. Hence, we use the $\frac{3}{4}M$ notch. For weak-trip velocity and spectral width estimation, the latter requirement becomes more important because the phase relation between all the spectral bins is needed to fully recover the cohered spectrum. Hence, we use the adaptive notch that cannot exceed $M/2$. The notching of the uncohered spectrum inevitably introduces “noise” in the cohered spectrum, which we attempt to eliminate in step 11 before estimating the weak-trip spectral width. The weak-trip velocity is not biased by this noise, since it is white.

Noise floor estimates are made with the method given by *Hildebrand and Sekhon* [1974]. A 5-point boxcar filter is used for the spectral smoothing operations. (This is done to minimize spurious detection of coherent spectral edges due to noise.)

Figure 2 gives an example of how the spectra look at various stages of data processing. Figures 2a and b show the input spectra for trips 1 and 2 before they are combined. Figure 2c shows the combined signal made coherent to the second trip. The dashed line indicates the computed noise floor. In Figure 2d the strong-trip signal has been notched out. Note that in reality the uncohered spectrum is not very white. Figure 2e shows the spectrum of the weak trip after the strong trip was notched out. Again, the dashed line indicates the calculated noise floor. Finally, Figure 2f shows the weak-trip spectrum after the “noise” was zeroed. The estimates of power, velocity, and spectral width for both trips are also given in the final plot.

2.3 SZ PHASE CODE PROCESSING

As alluded to above, one of the weaknesses in random phase code processing is that in notching out the strong-trip signal one inevitably loses information about the weak-trip signal. In an effort to minimize this loss of information, *Sachidananda and Zrnić* [1999] applied a periodic phase code to the problem. This so-called SZ code has the useful property that the uncohered spectrum is repeated periodically across the Nyquist interval instead of being scrambled throughout as white noise. If the uncohered spectrum is narrow enough not to overlap each other in the periodic replication and there is no system noise, then retaining just two consecutive spectral replicas is enough to reconstruct the original spectrum perfectly. This means that, in principle, one can notch out the cohered spectrum without losing any information

about the uncohered spectrum. In reality this is never quite accomplished because 1) there is always receiver noise, 2) the uncohered spectra may be wide enough to overlap each other in their periodic replication, and 3) the sidelobes from the cohered trip may extend beyond the notch thus introducing more “noise.”

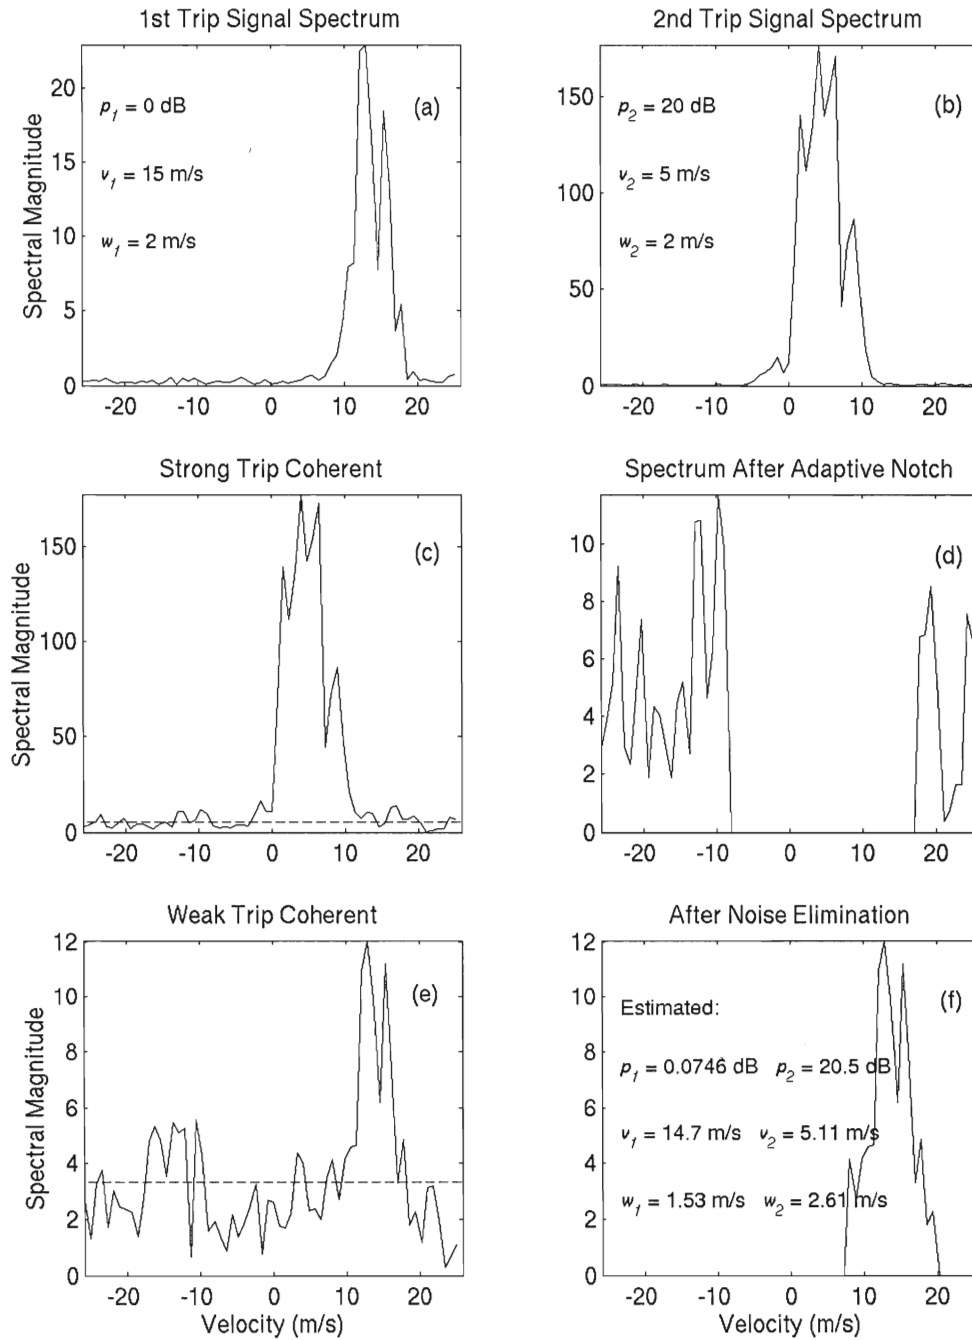


Figure 2. Illustration of random-phase code processing using simulated data. PRF = 1930 Hz, M = 64.

The SZ code is actually a family of codes, $SZ(n/M)$, defined by $\phi_k = \exp[-j\sum(n\pi m^2/M)]$, where the summation is taken from $m = 0$ to k , and n and M are integers to be selected appropriately. For our application $M = 64$ is equal to the number of I/Q data points to be processed. Then if n is chosen to be a factor of 64, the spectral replication number is M/n . The trade-off between various n values is that if it is set higher then the spectral “bandwidth” will be wider so there is less chance for spectral overlap, but two consecutive spectral replicas will take up more bins so the maximum notch filter width must be set narrower, thus allowing for more leakage of the coherent-trip signal. For realistic weather spectral widths, $n = 8$ and 16 are the only viable candidates. As we shall see later, however, for TDWR applications, $n = 8$ is the only choice. Therefore, we selected the $SZ(8/64)$ code for testing.

As for the data processing procedure, there is more than one approach. But first we will present the algorithm that we chose, and then we will discuss the possible variations. Here is the list of steps used to process the data.

1. Cohere to first and second trips to produce series s_w (weak) and s_s (strong).
2. Autocovariance process to get p_s, v_s, w_s .
3. Compute strong-trip spectrum with Hann windowing.
4. Notch out $M-2n$ spectral bins centered on v_s , calculate power in remaining bins, and multiply by $M/(2n)$ to get p_w .
5. Get noise floor of unnotched spectrum, smooth the spectrum, start from v_s , and find where spectrum drops to noise floor. Notch out this apparently coherent signal part of the spectrum, but do not exceed $|1-2n/M|$ bins.
6. Calculate power ratio $R = 10 \log(p_s/p_w)$.
7. If $R < 25$ dB, correct strong-trip power estimate: $p_s = p_s - p_w$.
8. Transform back to time domain and cohere to the weak trip.
9. Compute weak-trip velocity, v_w .
10. Compute magnitude spectrum and multiply by deconvolution matrix.
11. Get noise floor of spectrum, smooth the spectrum, start from v_w , and find where spectrum drops to noise floor. Keep this apparently coherent signal part of the spectral and zero out the noise bins.
12. Compute weak-trip spectral width, w_w .

Most of the comments made on the individual steps of the random-phase-code processing apply here also. The maximum allowable notch width is $|1-2n/M|$, which leaves at least two consecutive uncohered spectral replicas. The boxcar smoothing was increased to 7 points, because the “noise” for the SZ-coded spectra is actually quasiperiodic with a period of 8 bins. The deconvolution matrix is defined by $|\tilde{A}^* \text{diag}(F) A|^{-1}$, where A is the phase modulation code spectral convolution matrix, $\text{diag}(F)$ is a diagonal

matrix with elements of the notch filter spectral coefficients along the diagonal, \sim represents the reversal operation, and $*$ denotes complex conjugation.

Figure 3 gives an example of how the spectra look at various stages of data processing. Figures 3a and b show the input spectra for trips 1 and 2 before they are combined. Figure 3c shows the combined signal made coherent to the second trip. The dashed line indicates the computed “noise floor.” In Figure 3d the strong-trip signal has been notched out. Note that in reality the uncohered spectrum is not perfectly periodic. Figure 3e shows the spectrum of the weak trip after the strong trip was notched out. Note that the spectrum is widened by periodic duplication, which is an effect of the notching. Finally, Figure 3f shows the weak-trip spectrum after the notching artifact has been minimized through a deconvolution and noise subtraction procedure.

There are two particular areas in the processing stream for which there are significant variants. First, the maximum-width notch filter could be used in velocity and spectral width recovery of the weak-trip signal as advocated by *Sachidananda and Zrnić* [1999]. We agree that this would be optimal if the two remaining spectral replicas of the uncohered signal contained all the needed information for reconstruction. However, this is not the case in general. Our simulation study shows that an adaptive notch filter yields superior results for velocity estimation (Figure 4) for narrow strong-trip spectra. This is consistent with the speculations of *Frush et al.* [2002]. The only drawback is increased computation time.

The second area is the mitigation of the notch artifact. Instead of trying to deconvolve the consequences of the notching operation, one can refill the notch in an optimal way so as to avoid that particular problem altogether. Theoretically, the uncohered spectrum is replicated following a known sequence of phase shifts (which is dependent on the phase code pattern), so given at least two spectral replicas one can reconstruct the other replicas with the correct phase differences [*Frush*, 1999]. This technique is called spectral substitution. The same caveats discussed before, however, apply: the spectral “replicas” are never exactly alike due to noise, spectral overlap, and leakage of the strong-trip signal. Therefore, the reconstruction cannot be exact. Note that spectral substitution is mainly a competitor to deconvolution in the estimation of spectral width, since the notch filter does not bias the velocity estimate.

Different approaches can be taken to implement spectral substitution. *Siggia and Passarelli* [2002] use normalized cross-correlation as a criterion to determine which two spectral replicas to use as the “key” for filling out the rest of the spectrum by octants. *Sachidananda* [2001] takes the spectral replicant pair furthest from the strong-trip velocity peak as the key and refills the gap on a bin-by-bin basis. Our simulation tests showed that the latter approach produced better estimates, so we chose it for comparison with the deconvolution results.

The simulation results did not produce a clear winner. Spectral width estimation performance comparisons for deconvolution and spectral substitution depended on the set of input signal parameters. Although we selected the deconvolution method in the analysis of actual data, the spectral substitution technique could just as well have been used. In either case, however, we found a tendency toward overestimation due to the spectral widening effects during reconstruction. Therefore, we introduced the “noise” deletion scheme (step 11), which significantly reduced this bias.

For velocity estimation, spectral substitution did slightly worse than direct computation from the notched and recohered data. Therefore, if the spectral substitution method is adopted, it is better to estimate the velocity from the notched spectral data rather than from the reconstructed spectrum.

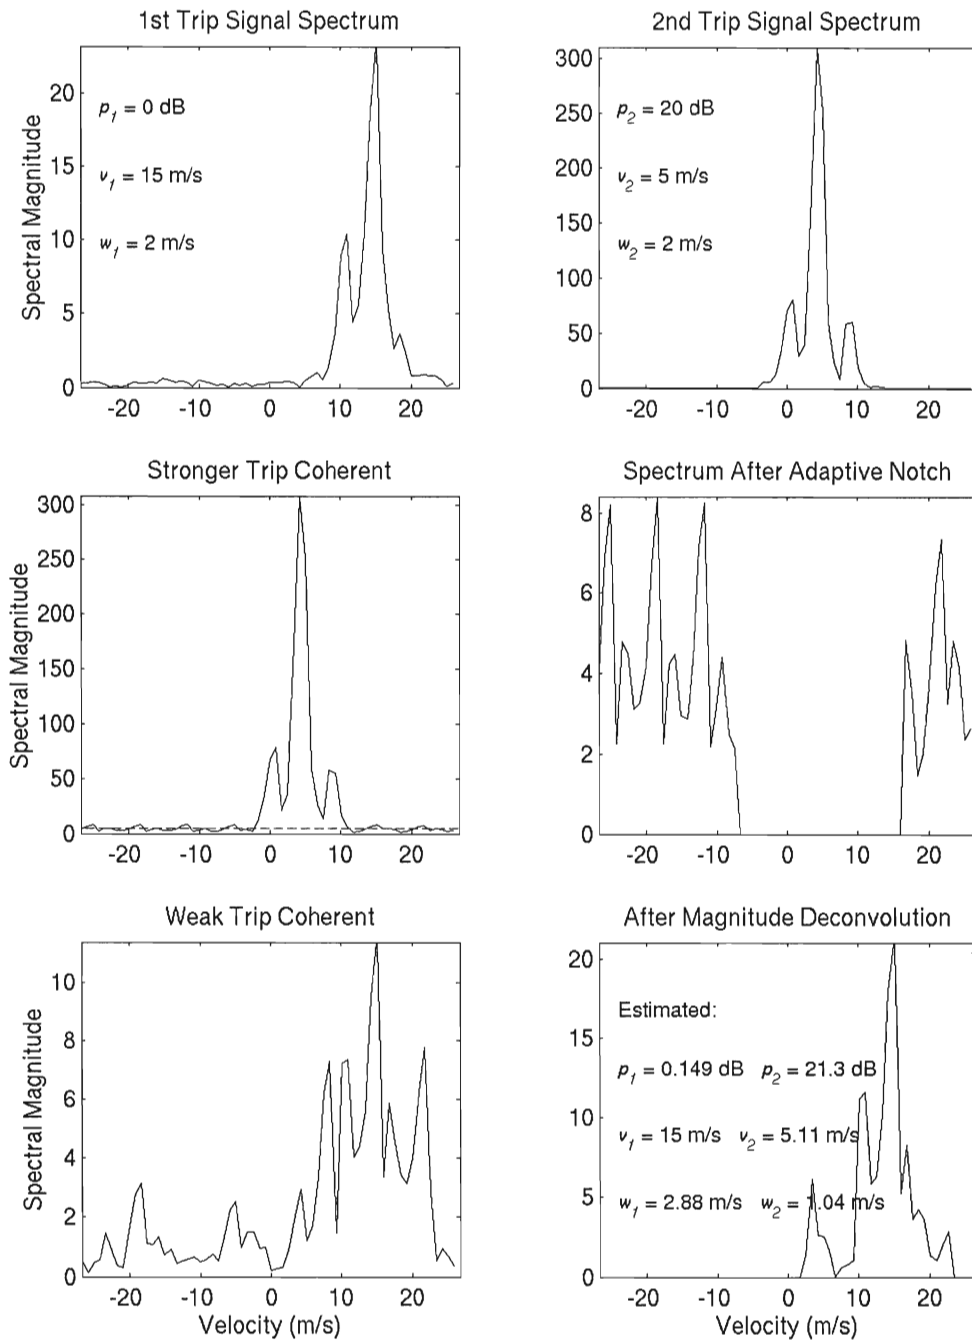


Figure 3. Illustration of SZ-phase code processing using simulated data.

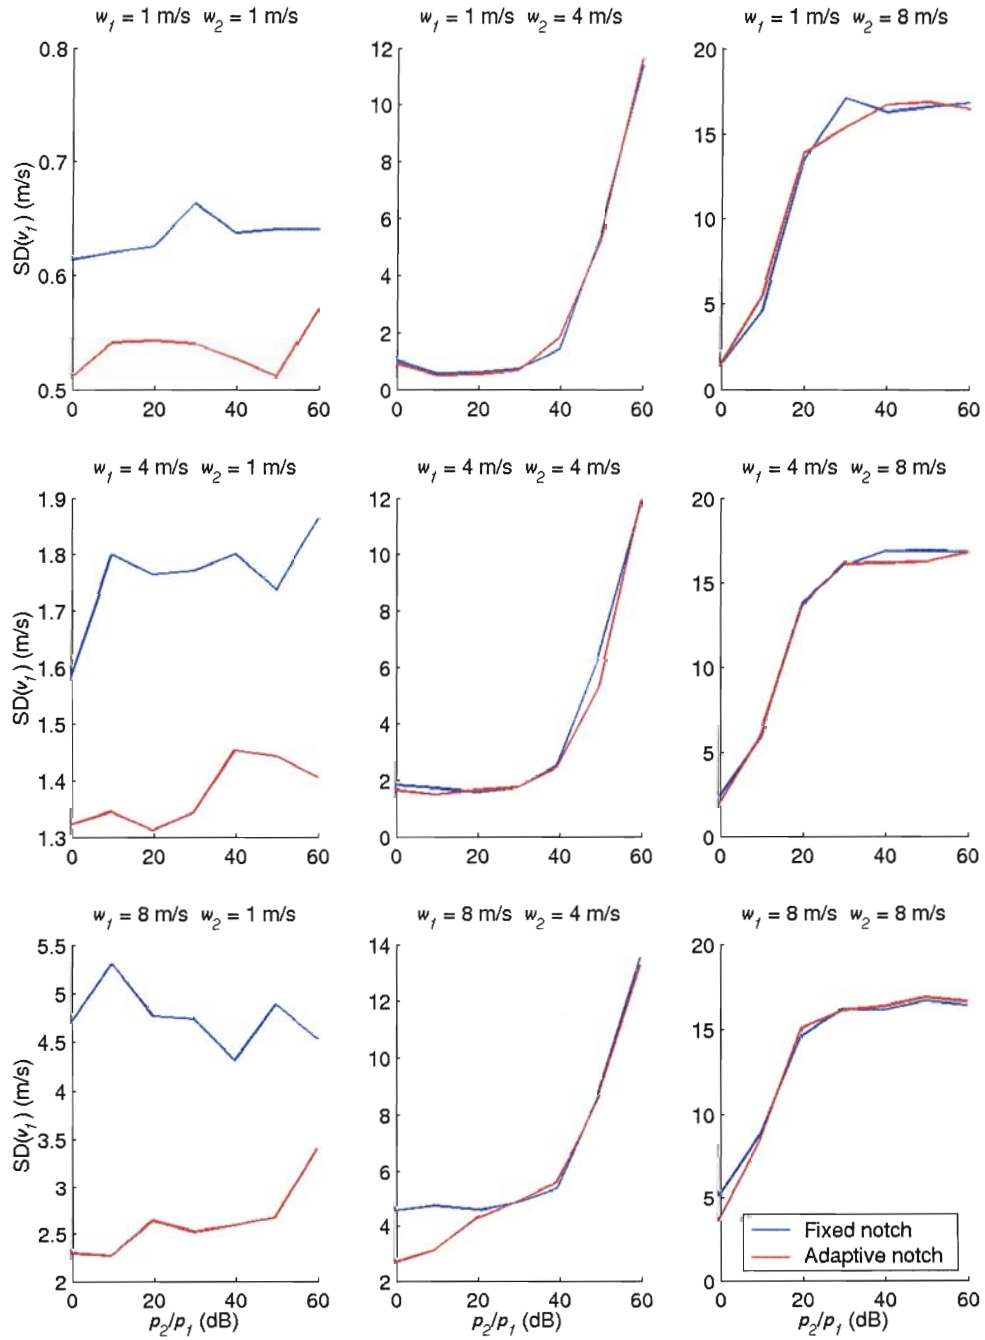


Figure 4. Weak-trip velocity estimation error statistics (standard deviation) for SZ(8/64). 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval.

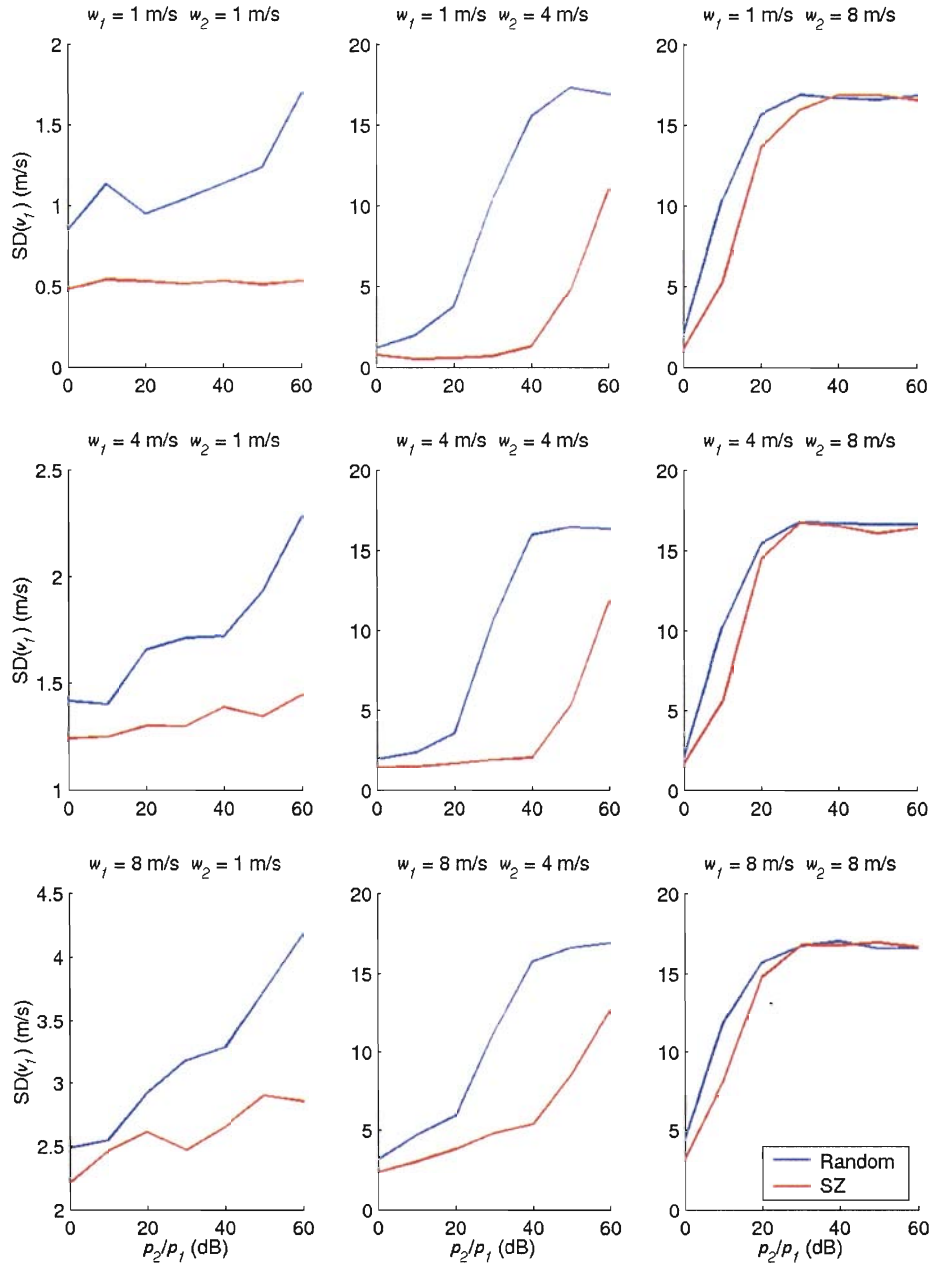


Figure 5. Weak-trip velocity estimation error statistics (standard deviation) for SZ(8/64) and random phase codes. 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval.

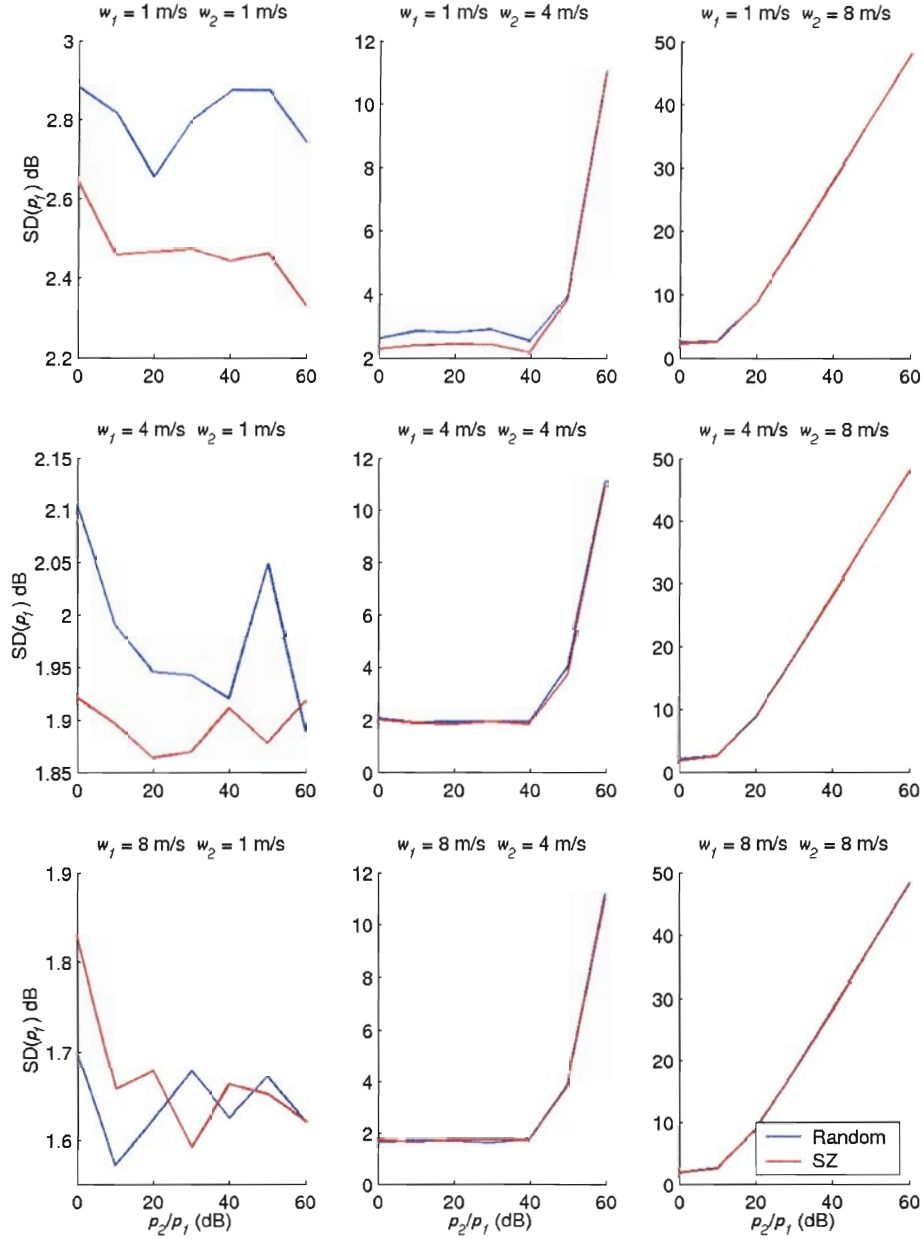


Figure 6. Weak-trip power estimation error statistics (standard deviation) for SZ(8/64) and random phase codes. 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval.

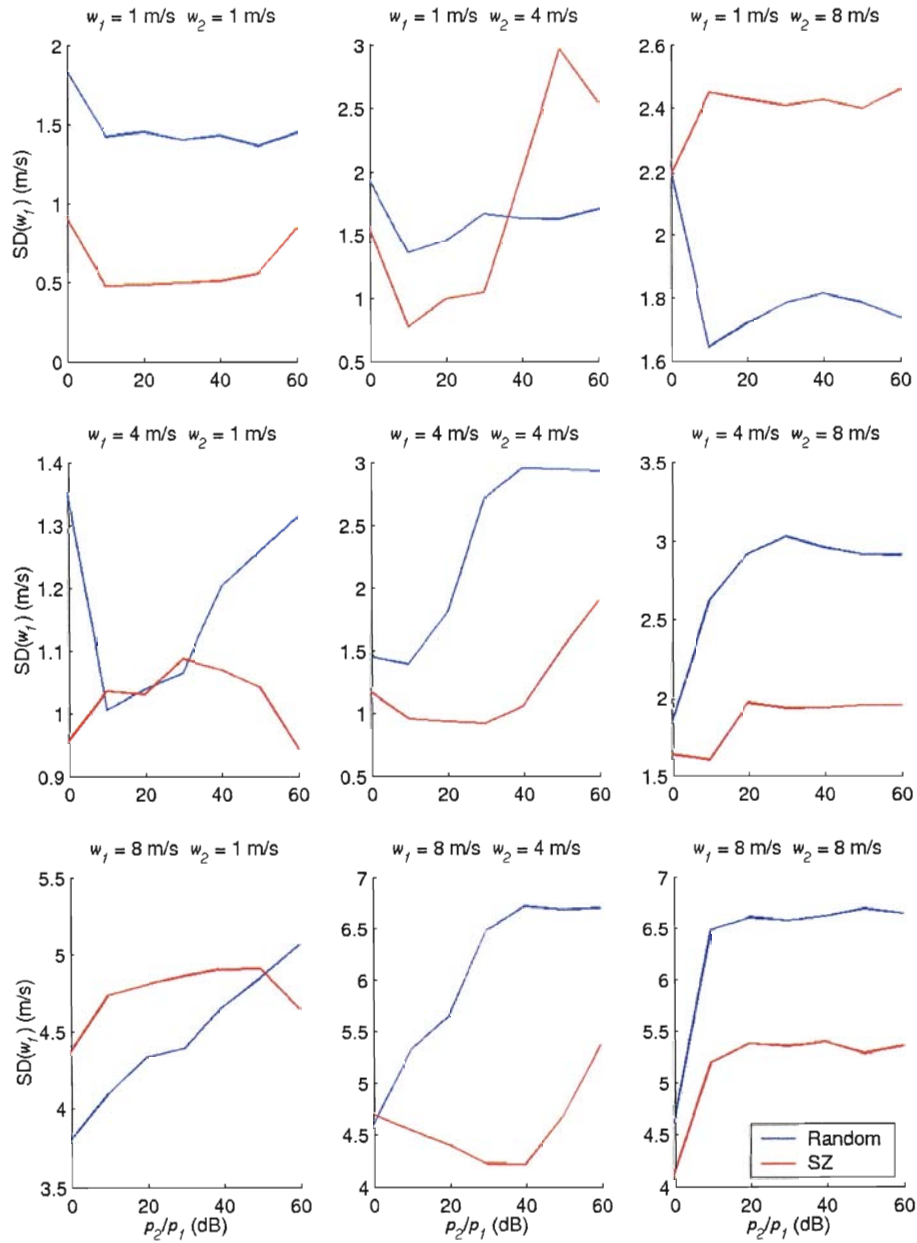


Figure 7. Weak-trip spectral width estimation error statistics (standard deviation) for SZ(8/64) and random phase codes. 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval.

2.4 COMPARISON OF RANDOM VS. SZ PHASE CODE PROCESSING RESULTS

The velocity estimation error statistics for random and SZ phase code processing are compared in Figure 5. The SZ phase code performs better overall, with significant differences for the midranges of spectral width and power ratio. Note that for either code, the velocity estimate errors become unacceptably large if the strong-trip spectral width is as wide as 8 m/s.

The power estimation procedure is essentially identical for random and SZ phase codes. Interestingly, however, SZ-phase code has a slight edge for very narrow input spectra (Figure 6). We speculate that this might be caused by a departure of the uncohered random-phase-code spectrum from the theoretical white noise spectrum. Since the weak-trip power estimate is made from only $\frac{1}{4}$ of the spectral bins assuming a flat spectrum, a deviation from a flat spectrum would introduce an error in the power estimate. This error may be secondary to the error caused by leakage of the strong-trip signal into the unnotched spectral bins when the strong-trip spectrum is wider, but it may become noticeable when the input spectra are narrow.

The SZ code slightly outperformed the random code in estimating weak-trip spectral width (Figure 7). Note, however, that overall performance for either method was not especially good.

Note that the parameter estimation errors are strongly dependent on the strong-trip spectral width but not the weak-trip spectral width. We also varied the noise level and found that, as expected, estimate errors increased with noise. However, the effects were not very significant until the SNR for the weak trip dropped below about 10 dB.

All the strong-trip parameter estimation errors were essentially the same for both code types.

3. PHASE CODE ISSUES ASSOCIATED WITH PROCESSING ACTUAL DATA

There are problems associated with the processing of real data that we did not address in the simulation tests. Before we go on to examine example data from the field, let us discuss these issues.

3.1 MULTIPLE TRIP OVERLAY

Although we only tested the candidate phase code processing for two-trip signal overlays, TDWR surface scans can contain weather signals up to 6 trips for PRF = 1930 Hz (Figure 1). In principle one could extend the filtering method used for two-trip range overlays to an arbitrary number of trip overlays for either SZ or random phase coded spectra. In practice, however, this does not work well. What happens is this. Begin with a spectrum that is cohered to trip A. Filter out the signal from this trip, which, inevitably, filters out some of the signal from all other trips as well. Now you cohere for trip B, but in addition to losing some of the trip B signal, the notching process has spread out its spectrum, making it more difficult to filter *it* cleanly.

We ran some simulations for the 3-trip overlay case ($p_1 < p_2 < p_3$) to recover v_1 , but the results were not encouraging. Even for $p_1 = 0$ dB, $p_2 = 5$ dB, and $p_3 = 10$ dB, v_1 recovery was not satisfactory. If, however, p_1 was bigger than the sum of all other powers or second biggest (but still larger than the sum of the powers of all the weaker trips), recovery became possible.

What this means is that we need to know the relative signal powers from all the trips that have to be considered. That would enable us to say whether the 1st trip parameters are recoverable. Can this information be obtained from the phase-coded data stream itself? The answer is no for SZ coded data. There is no way to obtain this information, because the uncohered-trip spectral data look “coherent,” thus confusing any attempt to separate out the relative powers from different trips. Consequently, an SZ-phase-code surface scan must be accompanied by a low-PRF scan that gives unambiguous power estimates up to 460 km.

With the random-phase code, the uncohered trip spectrum looks like noise, so it is possible to estimate the relative trip powers using a measure such as $|R_1|/R_0$, where R_0 and R_1 are the zeroth and first autocovariance function lags. (This quantity is known as the SQI or signal quality index). How well this technique works in practice will be discussed in Section 4.

SZ(8/64) Phase Code

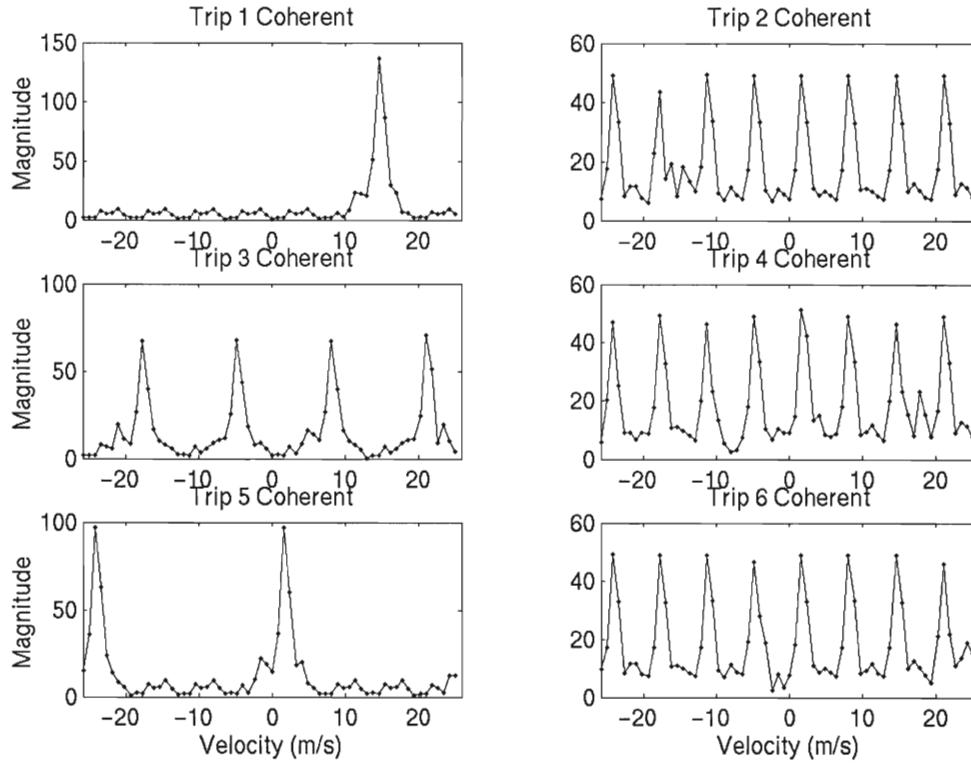


Figure 8. Illustration of the spectral replication properties of the SZ(8/64) phase code using simulated data. The input parameters are $p_1 = 15$ dB, $v_1 = 15$ m/s, $w_1 = 1$ m/s, $p_2 = 0$ dB, $v_2 = -15$ m/s, $w_3 = 1$ m/s. There is no added noise power.

There is another drawback to using the SZ code for an arbitrary k^{th} trip overlay filtering. Although the SZ(8/64) code will not be coherent again until the 9th trip, it will have trouble recovering the 1st trip signal from the filtered 5th trip (or, more generally, for any trip difference of 4). This occurs because the spectral replication property that makes the SZ coding superior to random phase coding changes with the difference in trip numbers. For reconstruction of the 1st trip signal, there must be at least two spectral replicas remaining after the out-of-trip signal has been notched out of the spectrum. For example, if the 2nd trip is cohered, there will be 8 spectral replicas of the 1st trip, so even after notching out the 2nd trip there should be enough information left over to reconstruct the 1st trip. However, when the 5th trip is cohered, there are only 2 spectral replicas of the first trip. Therefore, if anything is notched out, the ability to reconstruct the 1st trip signal deteriorates immediately. The spectral replication properties of the SZ(8/64) code are illustrated in Figure 8. The SZ(16/64) code is even worse, as it produces only 2 spectral replicas for the 3rd trip before becoming coherent with the 1st trip on the 5th trip. This is why we are only considering the use of the SZ(8/64) code.

The parameter estimation performance of the SZ code is correspondingly dependent on the trip differences. Considering only trip differences up to 6, for trip differences of 3 and 5, the performance is the same as that already presented for a trip difference of 1 (Figures 5 to 7). The estimation performance for a trip difference of 2 corresponds to that of the SZ(16/64) code. For a trip difference of 4, the performance corresponds to that of the SZ(32/64) code. Rather than presenting the statistics for each case in detail, we show how the velocity estimation performance changes for a given set of parameters (Figure 9).

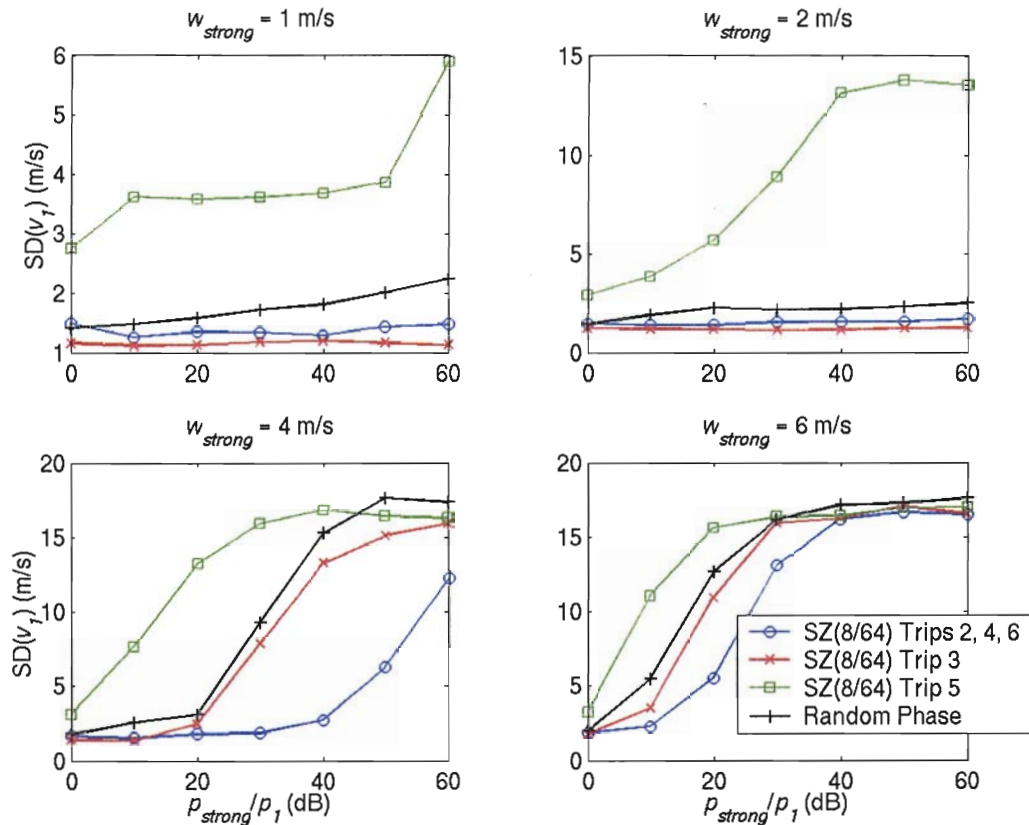


Figure 9. Weak-trip velocity estimation error statistics (standard deviation) for SZ(8/64) (for various trip differences) and random phase codes. The strong-trip numbers are indicated in the legend. 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval. The input weak-trip spectral width was 4 m/s and the weak-trip SNR was 20 dB.

Clearly, if there is significant weather in trip 5, it is better to use random phase coding rather than SZ coding. To determine if there are such out-of-trip signals, once again we see the need to have an accompanying low-PRF surface scan.

3.2 DATA CENSORING

Parameter estimates are not very useful unless one knows how valid they are. If bad estimates cannot be flagged, the data user cannot have any confidence in the product. A reflectivity estimate may be flagged if the ground clutter in that gate is too strong, while velocity and spectral width estimates may be censored according to a criterion such as SQI. For random-phase-code processed velocity and spectral width estimates, this type of “self measure” appears to work well. For SZ-phase-code processed data, however, it does not. The end spectrum can “look” good (“coherent”), but the answer might be completely off. The way to get around this problem is, again, to have an accompanying low-PRF scan. Then one can use the power and spectral width data (for all relevant trips) from this range-unambiguous scan to censor the velocity and spectral width estimates using the standard deviation data computed from simulations. We implemented such a look-up table approach in censoring SZ-code parameter estimates in Section 4.

3.3 GROUND CLUTTER FILTERING

Recognition of the presence of ground clutter is more difficult for SZ-coded data. For example, let's say that the second-trip signal is much stronger than the first-trip signal. If the time series of SZ(8/64) coded data is cohered for the first trip, then the second-trip spectrum will be prominently replicated 8 times across the Nyquist range. If one of these second-trip spectral replicas happens to fall on zero Doppler, then it could obscure or be mistaken for the ground clutter. This problem does not occur for random phase coding, because the second-trip spectrum would appear as noise when cohered for the first trip.

The filtering of ground clutter in general degrades the ability of the phase code to cohere to other trips. This is because the clutter filter inevitably deletes (or distorts) the information needed to completely reconstruct the signal from the uncohered trips. In this case the SZ code has an advantage in that the spectral gap generated by the clutter filter can be optimally refilled for a chosen trip using the spectral replication property. This is analogous to the spectral substitution method discussed earlier. We did not implement this approach at this time. Details are given in *Sachidananda* [2002].

3.4 PHASE ERRORS

In real life, the phase shifts introduced to the transmitted pulses may not be exactly as requested by the controlling program. For example, the phase shifter currently installed on the FAA's PSF TDWR in Oklahoma City yields an error of about 1.8° . Much of this is due to a bias of -1.6° . After removing the bias, the random standard error is still about 0.85° . This is, of course, not a problem for random phase codes, since the exact values (and their relative sequential differences) do not matter. For systematic phase coding, however, there is a potential that a perturbation to the prescribed phase sequence could affect important spectral replication properties, even if the measured phases are used for decoding.

Fortunately, the results of simulations for SZ(8/64) that we ran showed that the velocity recovery was not significantly affected until the random error increased to the order of 10° if the measured phases were used in decoding and processing. Therefore, phase errors should not be a concern in selecting a phase code. In fact, the upgraded TDWR system will have a digital IF transmit waveform synthesizer that should produce negligible phase errors.

3.5 CODE TRUNCATION

Periodic phase coding works best when the full code length is processed. For example, the SZ(8/64) code performs optimally when the dwell length is 64 points (actually, $64+k-1$ for full k -trip recovery). Performance begins to degrade when the full length is not used, but this fall-off is actually not so steep (Fig. 4.6 in *Sachidananda* [1998]). Thus, it does not appear to be a significant problem in synchronizing the dwell with azimuth sectors. One can also overlap the edges of the dwells to make sure the full code length is processed. Random phase code processing is not affected by segment length beyond the usual increase in variance associated with fewer samples.

4. PHASE CODE PROCESSING OF ACTUAL DATA

FAA's PSF TDWR in Oklahoma City, equipped with a phase shifter and I/Q data recorder, provided field data using both random and SZ(8/64) phase coded waveforms. In particular, we analyzed a case study from May 23, 2002, around 1700 UTC. The scans were run at an elevation of 0.3° and the antenna rotation rate was $22^\circ/\text{s}$. The key characteristics of the TDWR are as follows: frequency = 5.6 GHz, peak power = 250 kW, receiver noise = 53 dB, pulse width = 1 μs , antenna beamwidth = 0.5° , and range gate spacing = 125 m. Consecutive scans were run using the random and SZ(8/64) phase coded waveforms at high PRF and an uncoded waveform at low PRF. At the time, the low PRF was set to 500 Hz and the high PRF to 1500 Hz, so the former had an unambiguous range of only 300 km covering 3 trips of the latter. For later data collection, the low PRF scans have been reset to cover 460 km unambiguously. The combination of the PRFs and the rotation rate yielded 1° dwell lengths of about 23 and 69 points for the low and high PRFs. For the low PRF and random phase modes, we used arbitrary-length DFTs, and for the SZ code mode we used 64-point FFTs.

In processing the phase-coded data, the low-PRF results were first looked up to determine the signal strength in each trip. The ground clutter was filtered if needed, then the strongest two trips were processed according to the enumerated random and SZ phase code algorithm steps outlined in Sections 2.2 and 2.3. For the weakest trip, the parameter estimates are not expected to be good, but we computed them anyway, raising a flag to mark this condition. The weakest-trip velocity and spectral width were simply computed from the notched and recohered (to this trip) data. The weakest-trip power was set equal to the residual power filtered out as "noise" from the reconstructed weak-trip spectrum or the receiver noise level, whichever was greater.

As discussed earlier, clutter filtering adversely affects phase code processing, so the following scheme was implemented for random-phase-code processing to minimize filtering without use of a clutter map, which was not available at this range resolution.

1. Cohere to first trip.
2. Hann window and DFT.
3. Calculate total power (p_t).
4. Ground clutter filter.
5. Calculate remaining power (p_r).
6. If $p_t - p_r > \text{noise power}$, inverse DFT and use filtered data for rest of processing. Else, revert to unfiltered, unwindowed data.

In this way, the clutter filter is applied only if the estimated clutter power is significant, i.e., above noise level. As discussed earlier, for SZ-phase-code processing a concentration of power around zero Doppler does not necessarily indicate the presence of clutter, so the clutter power computed from the low-PRF scan was used to determine whether or not to apply the clutter filter.

The ground clutter filter works in the following way. First, a "core" clutter spectral width is (pre)calculated using the radar parameters, $\sigma_c = (\ln 2)^{1/2} \alpha \lambda |\cos \theta_e| / (2\pi \theta_{BW})$, where α is the antenna rotation rate, λ is the radar wavelength, θ_e is the beam elevation angle, and θ_{BW} is the one-way half-power beamwidth [Doviak and Zrnić, 1993]. We added 0.1 m/s to this value to account for wind effects. This number is then increased using an analytical function that calculates the finite-window spreading effect over the DFT to arrive at the spectral filter width, $\sigma_f = [2\sigma_c^2 \ln(p_p/p_n)]^{1/2}$, where p_p is the peak power at zero Doppler and p_n is the noise power [Sachidananda, 2002]. Then the spectrum is interpolated across the filter width (instead of nulling) in order to minimize bias in near-zero-Doppler weather spectra. Tests with real data showed this technique to work quite well. Very narrow weather spectra centered on zero Doppler can, however, fool the algorithm.

As an aside, we note that the apparent clutter widening effect of the sequence length depicted in Figures 3.1 and 3.2 of Sachidananda [2002] is, in fact, an artifact. In the averaging of the spectral width, the absolute value operation was used, which biased the means high for small spectral widths. This occurs because the autocovariance estimator for spectral width can produce negative values when the width is narrow (near zero). For statistical averaging, these negative values should not be flipped to positive values. The upshot is that their "realistic" clutter widths are overestimates.

4.1 REFLECTIVITY ESTIMATE COMPARISONS

For power estimates we can use the data from the low-PRF scan as the reference in evaluating the performance of the phase code trip separation. The reference is obviously not perfect due to the time difference in the scans. Velocity comparisons cannot be made, because of the severe aliasing and loss of correlation in the low-PRF scans.

In addition to comparing random and SZ phase code processing, we can also examine the performance of the random phase code processing with and without the use of the low-PRF data. In the former case, the low-PRF SNR is used to determine the order of the trip recovery sequence: process the strongest trip, filter out the strongest trip, and then re-cohere for the second strongest trip. In the latter case, the SQI from the high-PRF spectra cohered for each trip determines the relative ordering of the powers. The comparison of the two methods is "fair" because the same data input is used for both. The results should indicate how important it is to have the low-PRF scan data for random phase processing of multiple trips.




Figure 10. Reflectivity estimates for (a) PRF = 500 Hz, full 300-km range displayed, (b) PRF = 500 Hz, only first 100 km shown, (c) PRF = 1500 Hz, random phase coding, no phase code processing, (d) PRF = 1500 Hz, random phase coding, phase code processing with low-PRF data, (e) PRF = 1500 Hz, random phase coding, phase code processing without low-PRF data, and (f) PRF = 1500 Hz, SZ(8/64) phase code processing.

The low-PRF scan clearly showed significant weather signal in all three trips (Figure 10a). The concentric circles indicate 100-km (one trip) in range. Figure 10b shows a zoom on the first trip reflectivity only. Since protection of first-trip data is of paramount importance, we will concentrate on this range. Figure 10c shows high-PRF reflectivity without any phase-code processing. Note the range-aliased power in the western sector, especially along some WSW and NWN radials. Figure 10d shows random-phase-code processed data that utilized relative trip power information from the low-PRF scan. Figure 10e display the result of processing that did not incorporate the low-PRF scan data. Finally, Figure 10f shows the reflectivity derived from SZ-phase-code processing. The SZ-code scan actually ran first, then the low-PRF scan, then the random-phase-code scan.

Visual inspection of these results confirms the effectiveness of phase-code processing in removing range-aliased signal from the first trip. Note that the SZ code appears to remove a bit more of the westward streaks than the random code. For random-phase reflectivity estimation, the availability of the low-PRF data for trip-strength ordering does not seem to be a significant factor. Note, however, that phase-code processing failed to remove all the range-folded power. Referring to Figure 10a, the regions with significant power in both the 2nd and 3rd trips in the WSW direction leaked signal into the first trip. This is as expected from theory and simulation—the first trip parameters cannot be recovered when there are two or more other trips with greater powers. These areas need to be recognized and censored.

Note that since the random-phase-code algorithm that does not rely on the low-PRF power can be set to recover any number of trips, we ran it for both 3-trip and 5-trip recovery (500 km is beyond the maximum possible extent of weather signals). The results were virtually identical for the first trip, which implies that there were probably no significant weather signals beyond 300 km during these scans.

We can also provide a more quantitative measure of the power estimation performance. By comparing the estimates on a gate-by-gate basis with the low-PRF scan, we can compute the standard deviations. Since we are mainly interested in the algorithm performance for weak-trip recovery, we restricted the computations to only those gates that were the second-strongest trip. This left about 15% of the total number (792 gates x 360 radials = 285,120) of range/azimuth cells. As with the figures, only 1st trip gates were used.

TABLE 1
Standard deviations of power estimate errors in dB.

Censoring	Random w/Low PRF	Random w/out Low PRF	SZ(8/64)
None	5.8	5.6	5.3
P-RANK	5.5	5.5	5.3
CCOR	5.6	5.5	4.8
P-RANK/CCOR	5.2	5.2	4.8

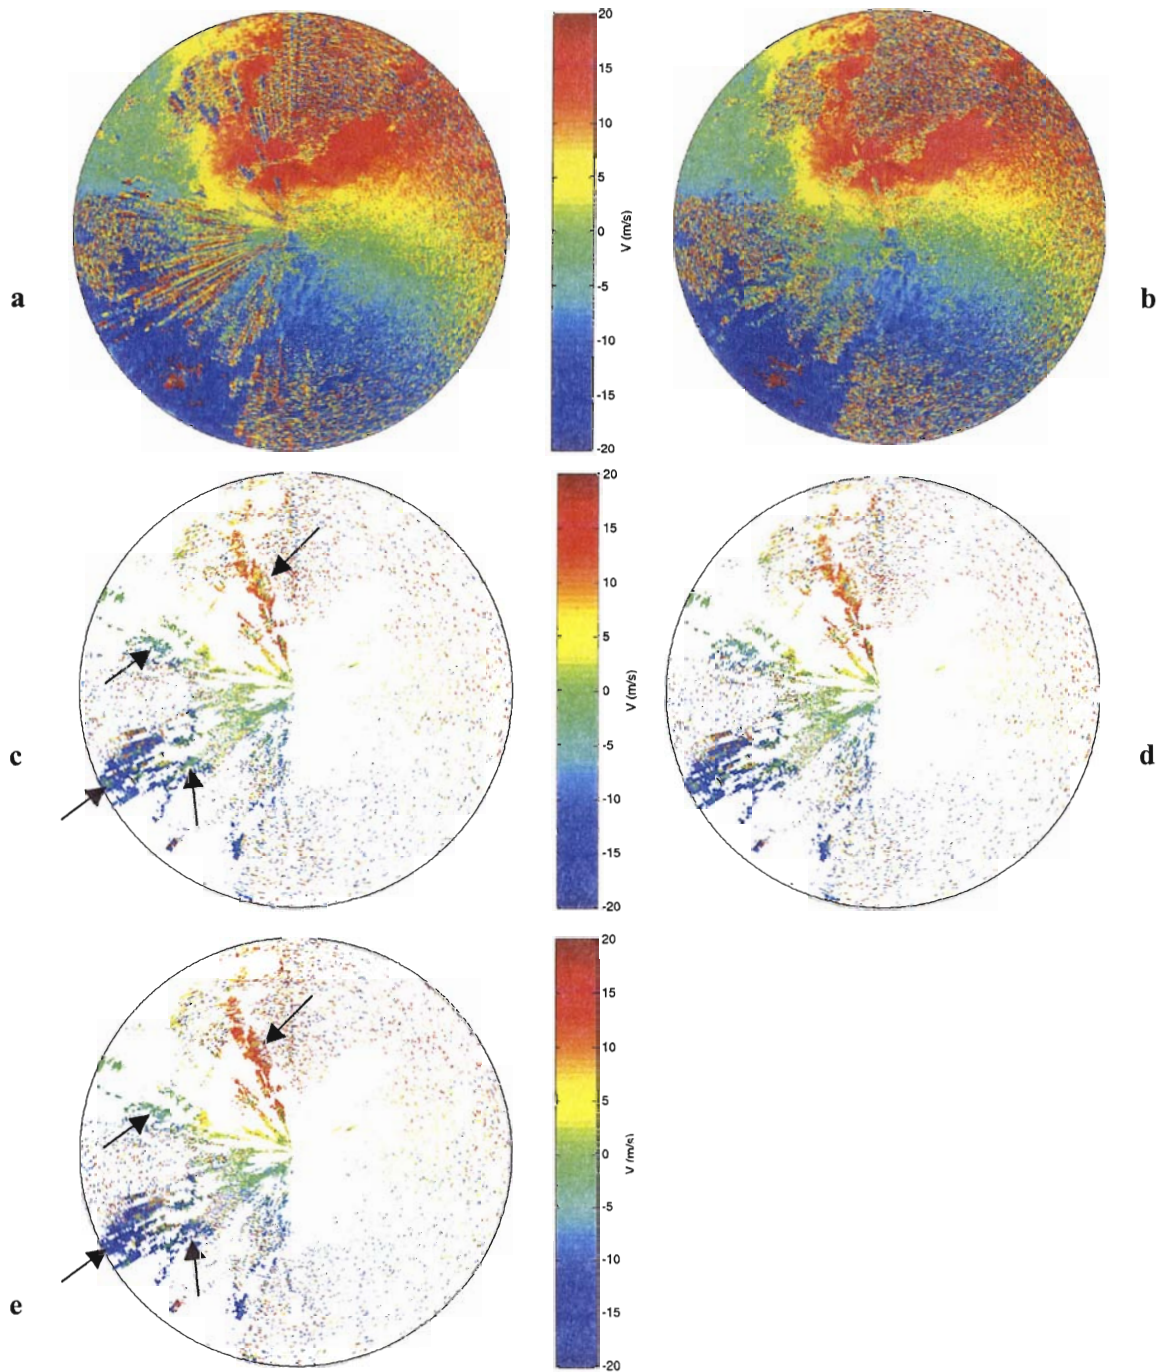


Figure 11. Velocity estimates for (a) random phase coding, no phase code processing, (b) SZ(8/64) phase code processing, (c) random phase coding, phase code processing with low-PRF data, weak-trip recovery gates only, (d) random phase coding, phase code processing without low-PRF data, weak-trip recovery gates only, and (e) SZ(8/64) phase code processing, weak-trip recovery gates only.

Table 1 lists the results of the power estimate error calculations. They confirm the visual impressions—that low-PRF data are not important for random-phase-code processing of power estimates, and that SZ(8/64) code outperforms the random code. In the table, "None" indicates no data censoring, "P-RANK" indicates censoring of recovered trips for cases where the power of the recovered trip was less than the power from the other two trips, and "CCOR" indicates censoring of gates associated with clutter correction factor < -25 dB, where $CCOR = 10 \log [p/(p + p_c)]$, p_c is clutter power, and "P-RANK/CCOR" indicates use of both P-RANK and CCOR for censoring. As expected, the data censoring lowered the estimation error.

4.2 VELOCITY ESTIMATE COMPARISONS

Velocity is a more important quantity than power in our application (since power is available from the low-PRF scan!), but harder to compare because the low-PRF scan does not provide a reference. We will have to rely on visuals for qualitative comparisons. Unfortunately, direct gate-to-gate comparisons are difficult because of the time difference between the two scans of about 90 s. Since the advection velocity was nearly 20 m/s that means a spatial separation of almost 2 km for the two scans.

Figure 11a shows the first-trip velocity estimates for random-phase-coded data with no phase-code processing. Note the streaks of multicolored noise that correspond to the out-of-trip power folding observed in Figure 10. Figure 11b shows that SZ-phase-code processing can remove some of these artifacts.

To make comparisons on only the weak-trip recovery cases, we plot only those gates for which the trip power was the second strongest. Figure 11c displays the velocity estimates from random-phase-code processing using the low-PRF power data. Figure 11d shows the velocity estimates from random-phase-code processing without the use of the low-PRF data. These two plots are indistinguishable by eye. Finally, Figure 11e shows the velocity estimates from the SZ(8/64) code processing. The results are very similar to those of the random code, but there are subtle differences. The SZ-code plot appears to have more consistent coloration in some areas, which is more realistic than isolated pixel patches with different wind speed. Some of these locations are pointed out with arrows in Figures 11c and e.

4.3 SPECTRAL WIDTH ESTIMATE COMPARISONS

At this time, spectral width is not used as input to TDWR operational algorithms that estimate wind shear and other aviation hazards near the terminal. There is potential utility, however, since high-quality estimates of spectral width, if processed carefully, could yield estimates of small-scale turbulent intensity.

Figure 12a shows the spectral width estimates from the low-PRF scan. However, because actual spectral widths are a significant fraction of the low-PRF Nyquist interval, accurate estimation of the widest spectra becomes impossible (see, e.g., Figure 6.6 by *Doviak and Zrnić* [1993]). One can see this by comparing Figures 12a and b in the areas indicated by arrows. (Figure 12b shows SZ-phase-code processed high-PRF scan estimates of spectral width.) Where the high-PRF data yield widths of 7-8 m/s, the low-PRF data only indicate widths of the order 4-5 m/s. Therefore, the low-PRF spectral-width estimates cannot be used as reference for quantitative comparisons between random and SZ code spectral width estimates. There are also repercussions for the censoring of SZ-code processed results as will be discussed later.

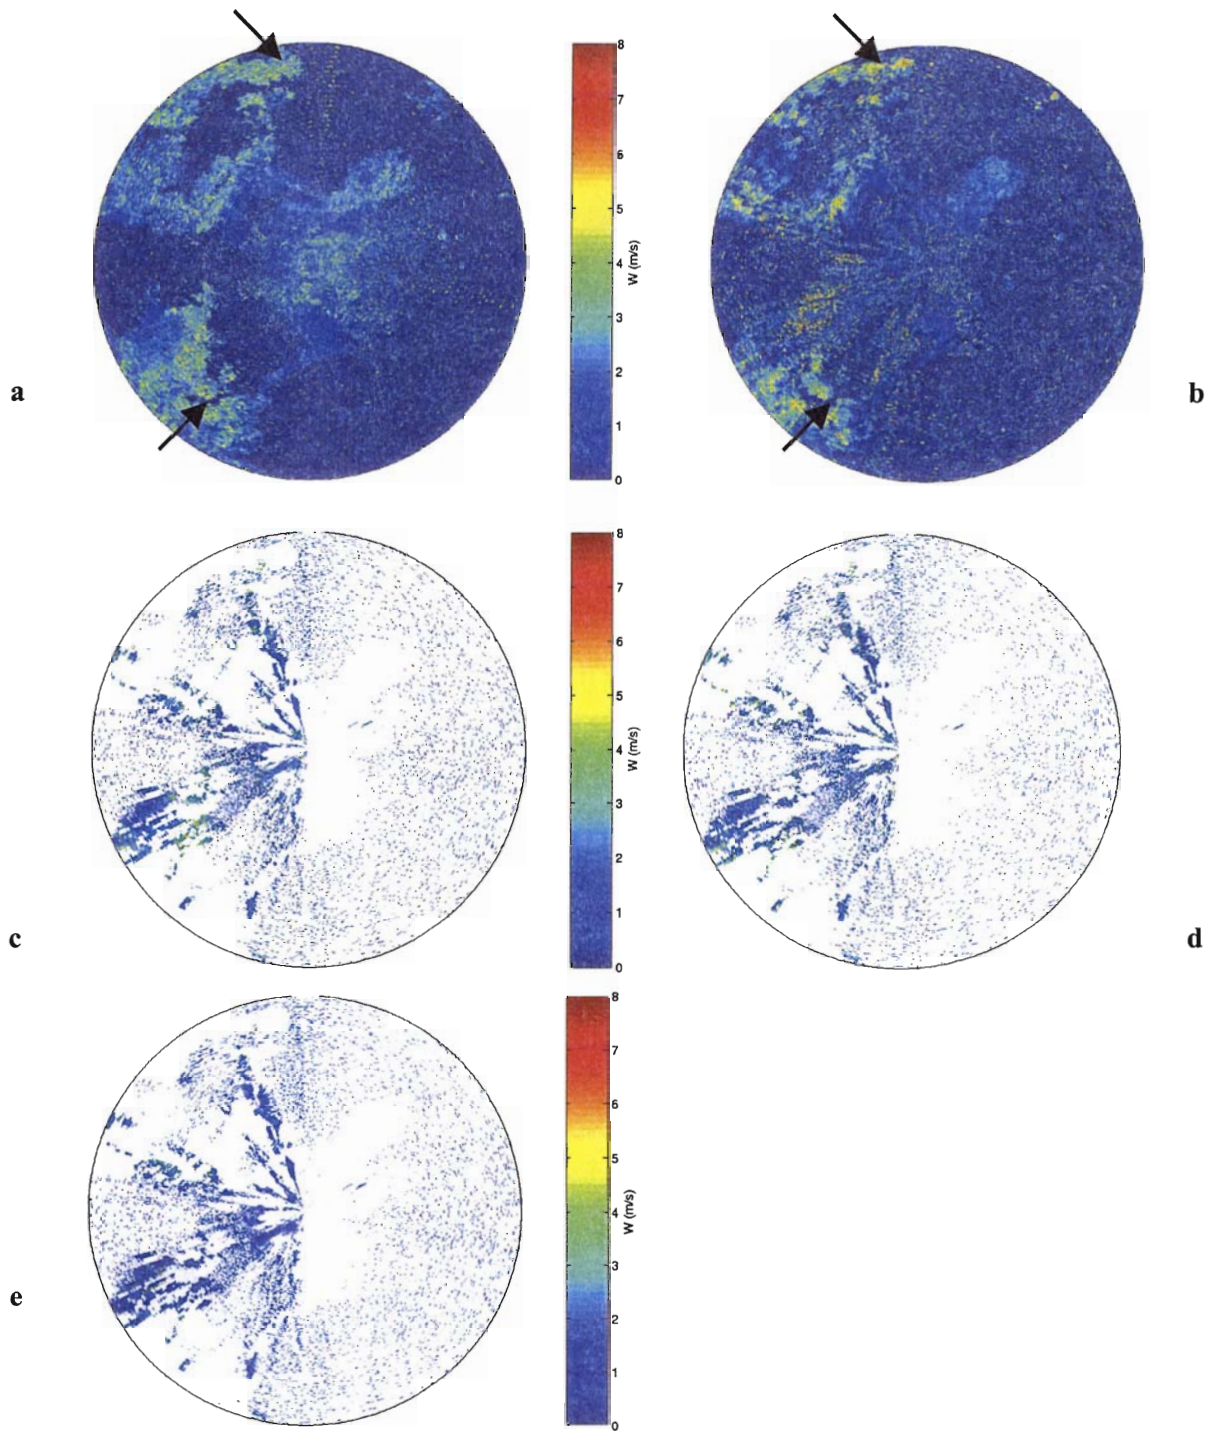


Figure 12. Spectral width estimates for (a) the low-PRF scan, (b) random phase code processing, (c) low-PRF data, weak-trip recovery gates only, (d) random phase processing, weak-trip recovery gates only, and (e) SZ(8/64) phase code processing, weak-trip recovery gates only.

For comparison, again we will only display the weak-trip recovery cells. The low-PRF estimates are shown in Figure 12c. As discussed above, the low-PRF estimates are not necessarily correct, but we can think of them as lower bounds. The random (processed using the low-PRF power data) and SZ phase code processed results are shown in Figures 12d and e. In the middle of the WSW streak, both the random and SZ code results underestimate the spectral widths. These are likely due to the penultimate “noise” elimination step, which was adopted to counteract the upward bias in the original spectral width algorithms. It appears that the random-phase-code results may have slightly less underestimation problems than the SZ-phase-code results. In any case, not much effort was put into optimizing the spectral width estimates due to the lack of end users at this time. There is scope for improvement in the future if interest in this parameter increases.

4.4 DATA CENSORING

For operational use of radar-derived meteorological parameters, the ability to censor bad estimates is critical. Improvements in estimation quality are useful only if failure of the process can be recognized.

We have seen in earlier sections that 1st trip processing of random-phase-coded data appears to work just as well without the use of the low-PRF power data. Will the ability to censor bad estimates be affected by the availability of a low-PRF scan? Figure 13 shows censored parameter estimates for random-phase-code data processed with (left column) and without (right column) the low-PRF power data.

The criteria used for eliminating “bad” estimates from random-phase-coded data were as follows. For all parameters: (1) CCOR < -25 dB, (2) the trip-rank of p was 3rd or below. For spectral width only: $p < 10$ dB. For velocity and spectral width: (1) SQI < 0.4, (2) $p < \text{sum of powers from all weaker trips}$. This last criterion was not available without the low-PRF data.

Comparison of Figures 13a and b (and for reference, 10b) shows that the unavailability of the low-PRF data does make a difference in censoring random-phase-coded reflectivity estimates. As indicated by the arrow in Figure 13b, there is power where there should not be. There is not much difference, though, in the velocity and spectral width censorship (Figures 13c to f). This is probably because SQI is the most crucial criterion for these parameters.

For SZ-code estimates, the censorship criteria list above were used except for SQI, which is not a quality discriminator for SZ-code-processed spectra. Additional censoring was accomplished by calculating the following five quantities for each gate: the power ratio of strong trip to weak trip, the SNR of the weak trip, the spectral width of the strong trip, the spectral width of the weak trip, and the trip difference between weak and strong (e.g., if the strong trip was 3 and the weak trip was 1, then the difference was 2). These parameters, computed from the low-PRF data, were used to look up an interpolated value in a precomputed multidimensional table of velocity and spectral width estimate standard deviations. The table was compiled from simulations of 1000 runs for each combination of input parameters. The velocity estimate was rejected if the estimated standard deviation was greater than 2 m/s, while the spectral width estimate was thrown out if the estimated standard deviation was greater than 1 m/s.

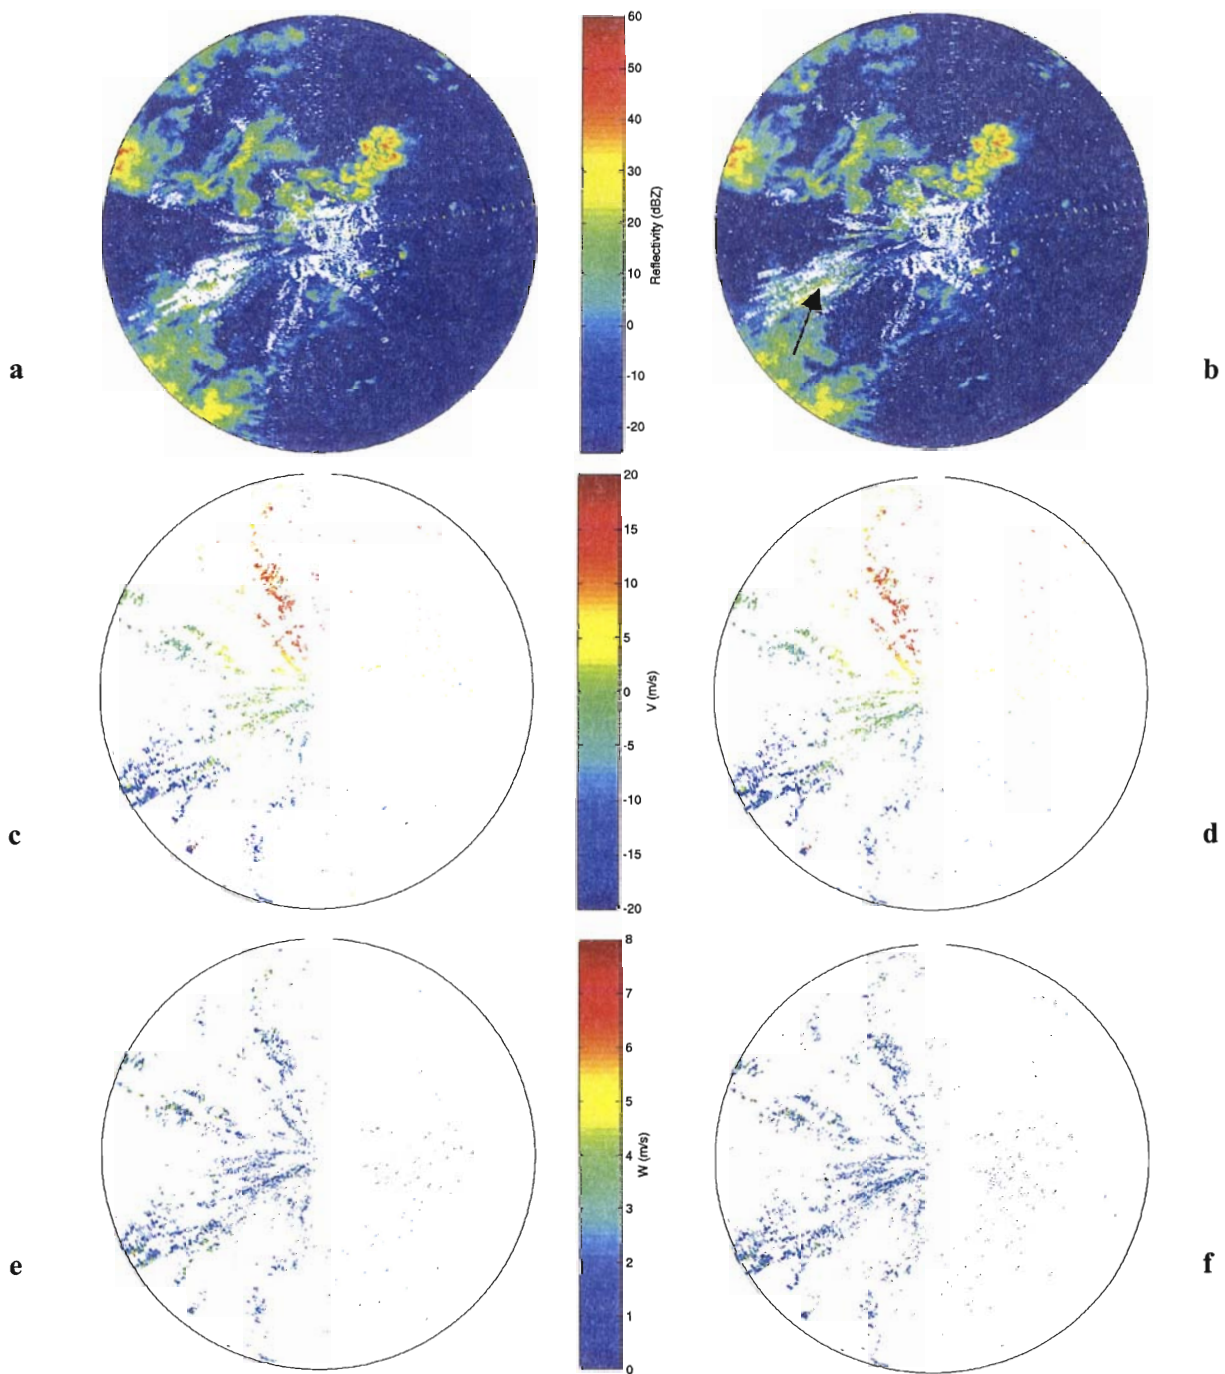


Figure 13. Random phase code processing censored estimates of (a) reflectivity using low-PRF data, (b) reflectivity without using low-PRF data, (c) velocity using low-PRF data, weak-trip recovery gates only, (d) velocity without using low-PRF data, weak-trip recovery gates only, (e) spectral width using low-PRF data, weak-trip recovery gates only, and (f) spectral width without using low-PRF data, weak-trip recovery gates only.

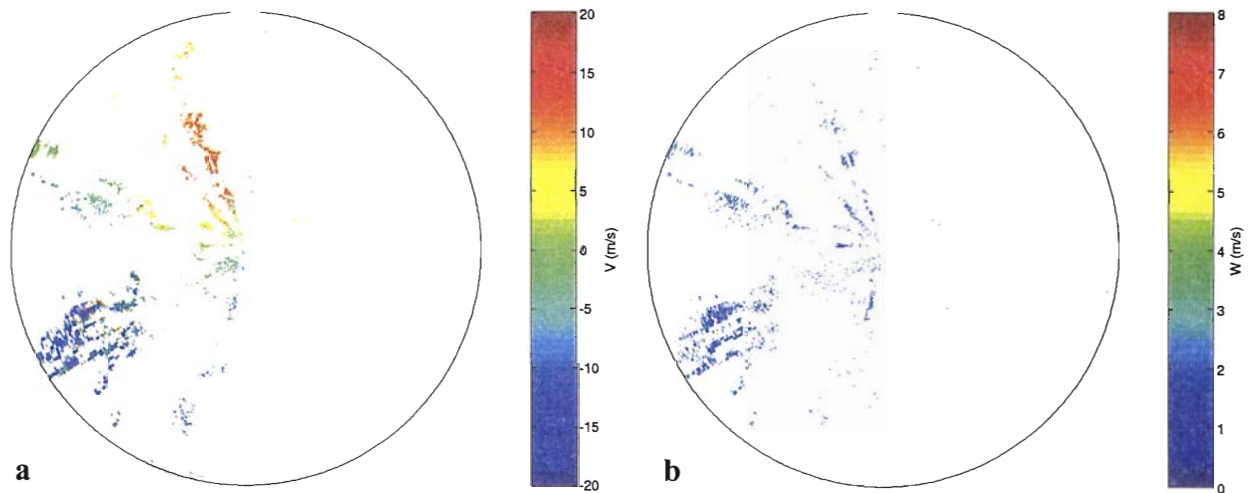


Figure 14. SZ(8/64) phase code processing censored estimates (weak-trip recovery gates only) for (a) velocity and (b) spectral width.

The results are displayed in Figure 14. Very roughly speaking, the two censoring methods for random and SZ code processed estimates (SQI vs. simulation table) produced similar results, even though they are completely different techniques. Looking at the details, however, there are some differences. Of course, the degree of censorship depends on the threshold values that are picked, but qualitatively there seems to be some real differences. For velocity censoring, the SZ results appear to leave more suspicious color “speckling,” suggesting a failure of the censoring mechanism. One possible cause of such a failure is the reliance on the low-PRF data for spectral widths. As discussed before, wide spectra cannot be accurately estimated from the low-PRF scan (e.g., Figure 12a and b). An underestimate in spectral widths would lead to an underestimate in standard deviation of velocity estimate, and, therefore, failure to censor bad data. This is not a trivial problem given the dependence of velocity estimate error on the spectral widths, especially on strong-trip width (Figure 15). The problem would be even more severe for a low PRF of 326 Hz (Figure 16), which is what we would want to use to cover up to 460 km in range. Note that spectral widths can only be estimated up to about 2 m/s, which is not enough to properly censor the SZ-code-derived estimates.

The spectral width censorship also differs somewhat, but it is difficult to say whether one technique yields better results than the other.

4.5 REASONS FOR RECOVERY FAILURES

What were the reasons behind the instances of bad parameter estimates seen above? Let's look at the conditions crucial for successful trip separation computed from the low-PRF data. Figure 17a shows the power ratios (strong trip/weak trip) and Figure 17b displays the 1st trip SNR in the weak-trip recovery gates. There were some quite high power ratios (up to about 50 dB) at close range in the west-southwest streak region, coupled with very low 1st trip SNRs. Such conditions make 1st trip recovery very difficult. This probably accounts for the suspicious patch of green (near-zero) velocity seen at close range in these gates in the velocity estimates (Figures 11c to e). The outer ranges of the west-southwest streak, on the other hand, have low power ratios and higher 1st-trip SNR, but have spectral widening in either the strong or weak trip (Figures 17c and d). The northwest streaks also have spectral widening in the strong trip. Note also that the low-PRF spectral width data can be an underestimate due to the limited Nyquist

velocity range. Taking these factors into account, it seems reasonable to tentatively conclude that the algorithms are working more-or-less as expected, given the power and spectral width conditions.




Figure 15. Velocity estimation error statistics (standard deviation) for SZ(8/64) versus variation in strong-trip (blue) and weak-trip (red) spectral width. 1000 simulations were run per data point. The input velocity values were varied randomly across the Nyquist interval.




Figure 16. Spectral width estimates versus input spectral widths. The straight line indicates perfect agreement. 1000 simulations were run per data point and the results averaged. The input velocity values were varied randomly across the Nyquist interval. 15 samples were used for the computation, which corresponds to a 21.6%/s antenna rotation rate for a dwell of 1°.

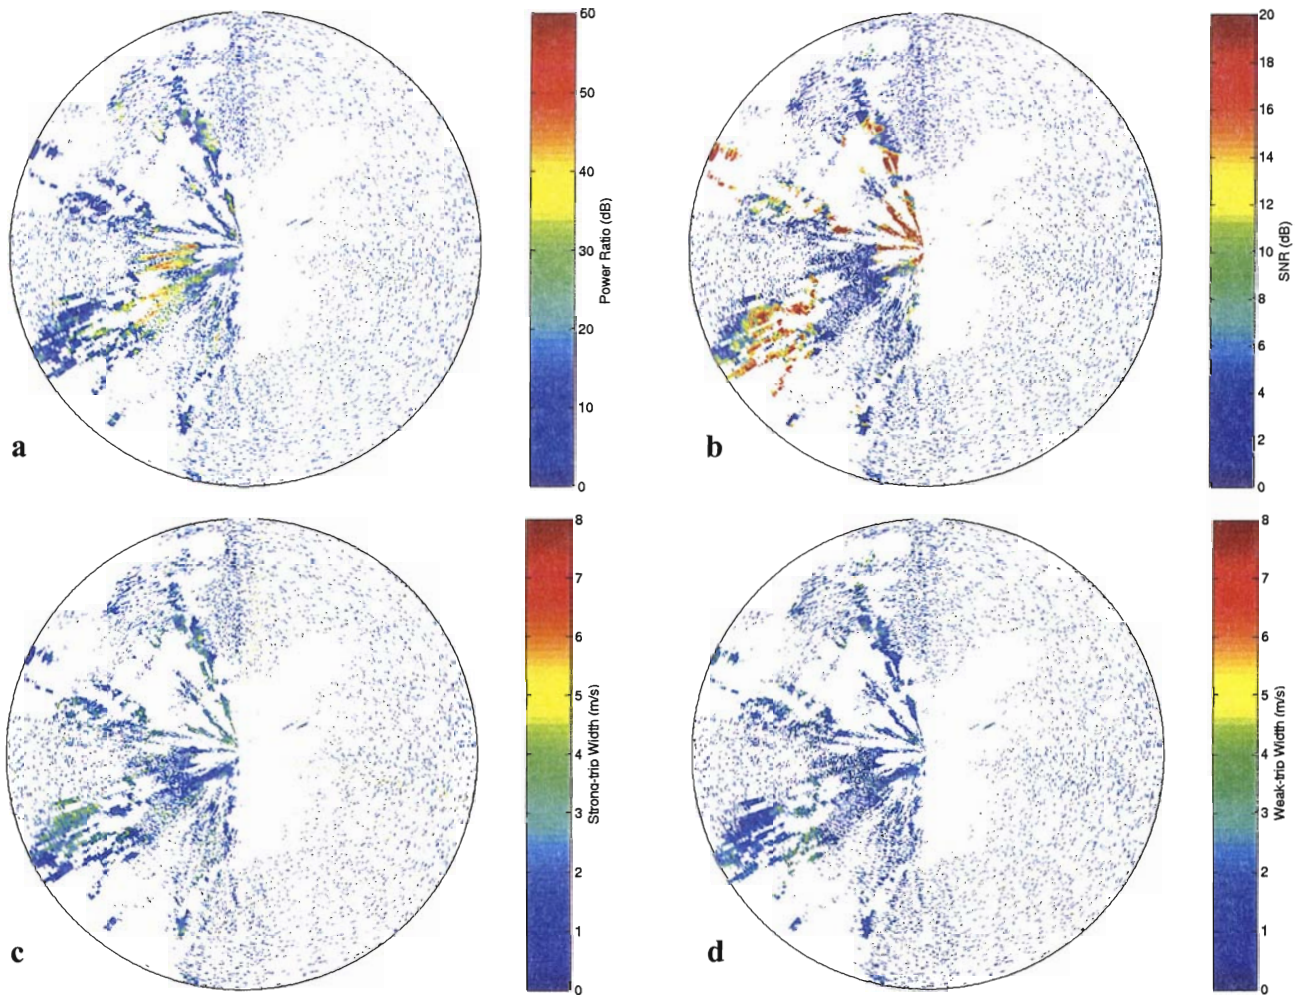


Figure 17. (a) Power ratio of strong to weak trip, (b) 1st trip SNR, (c) Strong-trip spectral width, (d) Weak-trip spectral width.

4.6 MULTIPLE TRIP RECOVERY

Although we have been mainly concerned with first-trip protection and recovery, future application of TDWR data may include expanded spatial coverage. The potential for such expansion can be seen in Figures 18a and b. SZ(8/64) phase code transmitted at a PRF of 1500 Hz was processed for three trips to provide an effective range of 300 km. The reflectivity compares favorably with the low-PRF scan (Figure 10a). Note the limitations such as the data gaps that show up as rings at the beginning of the second and third trips. These “shadow” zones are produced by the extremely strong ground clutter at close range in the first trip, which prevents the clean recovery of corresponding gates in the outer trips. One can also see probable velocity aliasing in the southwest sector of the third trip.

From a computational load viewpoint, the recovery of the outer trip parameters does not cost much extra. The bottleneck is more likely to be the data transfer bandwidth after the fact.

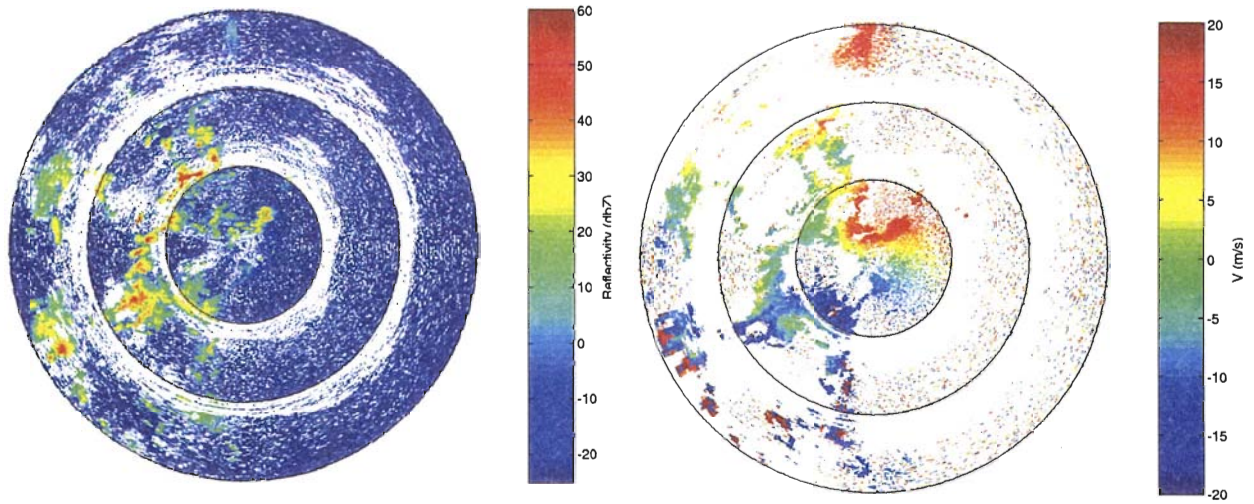


Figure 18. Three-trip SZ(8/64) phase code processing censored estimates for (a) reflectivity and (b) velocity.

5. MULTISTAGGERED PULSE PROCESSING (MSPP)

In the previous sections, we discussed the results of phase-code processing techniques for range ambiguity mitigation. For TDWRs, however, velocity ambiguity is also a problem. Running the transmitter at the maximum PRF of 1930 Hz yields a Nyquist velocity of 25.6 m/s, which can certainly be exceeded by realistic winds.

By transmitting pulses at variable time intervals, one can unfold aliased velocities. For example, with two PRFs one can use the Chinese remainder theorem to extend the unambiguous velocity range beyond the Nyquist interval of either of the individual PRFs. A straightforward application of this concept to phase-code processing is to transmit at two PRFs alternating between dwells (radials). The velocity estimate can be dealiased, if needed, using the neighboring velocity estimated from the other PRF. *SIGMET* [2001] has implemented this type of scheme for their commercial system.

This method is not without problems. It assumes that the velocity does not change appreciably between radials, i.e, a strong wind shear might result in invalid velocity unfolding. This may be problematic, since the top priority objective of the TDWR is to detect strong wind shears. Also, the region where the higher PRF radial is in its 2nd trip, but the lower PRF radial is still in its 1st trip presents a problem. Because of the often extremely strong ground clutter near the radar, much of the beginning of the 2nd trip is wiped out (e.g., Figure 18). This means that the high-PRF radials are not usually available in these gates, and, thus, velocity dealiasing is not possible.

Here we propose an altogether different approach to simultaneous range and velocity ambiguity mitigation, which we shall call multistaggered pulse processing (MSPP).

Recall that in the current TDWR data acquisition scheme, after the initial low-PRF scan a pair of PRFs for the next two scans is chosen to minimize range folding into the area around the terminal and to optimize velocity dealiasing performance. An alternative path to accomplish the same task (perhaps to do it better) is to simply transmit multiple PRFs per dwell within one scan. In this way one can eliminate the third scan at the same elevation and still dealias the velocity.

Let us take a concrete example. Define a waveform such that the PRI (pulse repetition interval) is T_0 for m pulses, then $T_0 + \delta T$ for m pulses, then $T_0 + 2\delta T$ for m pulses, and so on until the PRI reaches $T_0 + (N-1)\delta T$. Then the total number of pulses is mN and the total (dwell) time is $(mN-1)T_0 + m(N-1)N\delta T/2$. If $T_0 = 518 \mu\text{s}$, $\delta T = 60 \mu\text{s}$, $m = 8$, and $N = 8$, the dwell time would be 46 ms and a matching antenna rotation rate for 1°/dwell would be 21.7°/s, which is close to the monitor scan rate presently used.

There are two advantages to employing a multi-PRI waveform like this one. First, one obtains N estimates of power, velocity, and spectral width that are independent with respect to range folding, since each PRI corresponds to a different unambiguous range. So, unless the out-of-trip target covers a range wider than a small fraction of the PRI spread, one can statistically weed out outliers. For example, one can use a criterion such as "if $p_i - \langle p \rangle > 2\sigma_p$, then throw out p_i ." (Actually, $\log p$ probably yields a more normal distribution than p for statistical purposes.) Then one can take the median of the remaining p values (the median being more robust against outliers than the mean); likewise for v and w . Alternatively, the low-PRF data can be used to choose the PRI set with the least amount of out-of-trip interference.

The second advantage is that velocity dealiasing can be accomplished within each dwell. Since the Nyquist interval changes with PRI, the velocity computed for each PRI is different if it is aliased. The change in aliased velocity is given by $v(T_0 + k\delta T) - v(T_0) = \lambda q [(1/T_0 - 1/(T_0 + k\delta T)]/2$, where q is the

number of alias wrapping. Least-squares fitting of the ν estimates versus PRI to this equation then yields the alias number. Obviously, if there is no aliasing there is no change in ν and $q = 0$. With no noise and zero spectral width, the fit would be perfect and no error would result. Of course, with real-world data this does not hold. Choosing higher values of δT will enable the correspondingly steeper slopes to become more detectable through the noise. Each ν value is then dealiased according to q , and the final result is given by the median of all values. This procedure is outlined graphically in Figure 19.


Fitting becomes more complicated for cases where the alias number is not constant within the range of PRIs used. In that case the fitting has to be performed separately on the points with the same q value.

The two advantageous characteristics listed above hold true in general for multi-PRI waveforms. Therefore, δT need not be multiplied by integer increments and m and N can be chosen to be any number that fit in well with the required dwell time and scan rate. In fact, the initial low-PRF scan could be used to determine an optimal set of PRIs to use for mitigation of range folding. The fitting procedure will work better, however, if the PRI differences are spaced evenly.


The velocity dealiasing performance can be tested through simulations. Table 2 shows the dealiasing success rate for the example waveform described above. 1000 runs were performed for each entry in the table, with the input velocity varying randomly across 3 times the smallest Nyquist span, in this case ± 43 m/s.

To avoid false dealiasing, one could use a quality of fit criterion or pass along a warning flag. Optimal parameter values would have to be determined by testing on real data and constrained by scan rate, data resolution/quality requirements, etc.


The range-overlay protection function can only be tested on actual weather data. This is a task we are currently working on by collecting TDWR I/Q data using multi-PRI waveforms. The pulses are randomly phase coded also to make the out-of-trip signals incoherent within the constant-PRI blocks.



Constant $v \Rightarrow$ no aliasing.



Fit theoretical function to get alias number.
Dealias.



Different alias states.
Fit to majority group.

Figure 19. Illustration of MSPP velocity dealiasing.

TABLE 2

Velocity dealiasing success rates (%) for $T_0 = 518 \mu\text{s}$, $\delta T = 60 \mu\text{s}$, $m = 8$, and $N = 8$.

SNR (dB)	Spectral Width (m/s)				
	1	2	4	6	8
5	95	95	93	86	71
10	99	99	98	94	82
20	100	100	99	96	86
30	100	100	100	95	83
∞	100	100	99	95	85

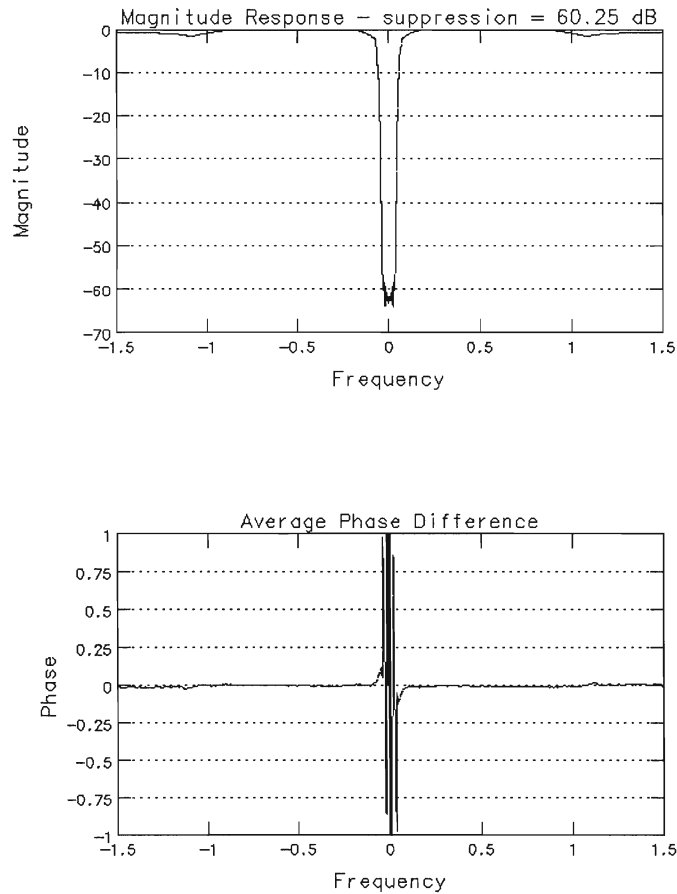


Figure 20. Magnitude response (top) and average phase difference (bottom) versus frequency for a multi-PRI clutter filter designed using the Chornoboy algorithm. The waveform parameters are given in the caption to Table 2. In the frequency axis, -0.5 to 0.5 corresponds to the full Nyquist span of the longest PRI. In this case, -1.5 to 1.5 corresponds to -43 to 43 m/s in Doppler velocity space.

One issue with multi-PRI waveforms is ground clutter filtering. The trick is to avoid having the notch filter around zero Doppler affect the velocity estimates at the corresponding aliased intervals. The perfect filter would have a flat magnitude response for the unfolded velocity spectral interval of interest and a linear phase response for the same interval except at the notch around zero Doppler. Chornoboy [1993] proposed a mean-squared-error design scheme that produces multi-PRI filters with excellent magnitude/phase characteristics, especially for block-staggered pulse trains. Figure 20 shows the magnitude response and phase error for a filter designed through his algorithm for the multi-PRI waveform that we tested above. The stop-band half-width is 1.2 m/s. The clutter suppression of 60 dB was calculated using the “core” clutter width formula given in Section 4 with an antenna rotation rate of $21.7^\circ/\text{s}$ and a beamwidth of 0.5° . This level of suppression for the stop-band width is comparable to that achieved by the currently implemented IIR filters. Outside of the stop-band region, the magnitude response is quite flat and the average phase errors are quite small.

6. SUMMARY

The main purpose behind this investigation was to determine the suitability of existing phase diversity methods for 1st trip protection from range-folded signals in TDWRs. In theory, periodic phase codes can be devised, such as the SZ codes, which outperform random-phase codes in range unfolding. For TDWR applications, we determined that the SZ(8/64) is the best choice of the family of SZ codes, and tested its performance against the random phase code. Except for protection against the 5th trip, the SZ(8/64) code did outperform the random code in weak-trip velocity and power estimation using simulated data. Analysis of real-world data confirmed some of this difference. However, the SZ-phase-code processing method had several disadvantages relative to the random-phase-code technique. The most critical problem was in censoring the resulting estimates. Evaluation of the validity of SZ-code-derived parameter estimates requires knowledge of the powers and spectral widths of the individual trip signals. Although these quantities can be made available from a separate low-PRF scan, the estimate of spectral width has an upper limit determined by its PRF. Since we need a PRF of 326 Hz to unambiguously cover the range extent where signal folding can take place, it means that the low-PRF scan can only provide spectral width estimates up to about 2 m/s. This narrow range is not enough to provide proper censoring capability for the high-PRF SZ-coded scan. It is possible to obtain better estimates of strong-trip spectral width from the SZ-coded data stream itself for high power ratios, but the weak-trip spectral width can only be obtained from a low-PRF scan. Both strong-trip and weak-trip spectral widths are needed for proper data censorship. For this reason, we cannot recommend the use of SZ codes for TDWR applications.

We showed that the random-phase code technique could be used on its own (i.e., with no accompanying low-PRF scan) without loss of performance. Censorship performance for power estimates, however, suffered without the low-PRF data. Coupled with the failure of phase-code processing to remove all out-of-trip overlays, we recommend that the low-PRF scan be kept, even if the random-phase-code method is implemented.

Even with phase-code pulses transmitted at the maximum PRF of 1930 Hz, the Nyquist velocity range is only ± 25.6 m/s. If greater velocities are to be measured, we need a way to dealias them. A straightforward adaptation of the phase-coding technique is to alternate two PRFs on successive radials (interdwell block stagger). In this way the phase-code processing is unaltered and the velocity estimates can be dealiased using the results from the neighboring radial. Preliminary testing of this scheme shows that it can work well, except at the beginning of the 2nd trip for the higher PRF, where severe clutter from the close-range 1st trip gates ruins the recovery of other-trip gates. In this “ring” zone, velocity dealiasing becomes impossible because only estimates from the lower PRF radials are available.

In addition to testing existing phase-coding schemes, we also proposed an intradwell multistaggered pulse processing (MSPP) technique. MSPP can provide excellent velocity dealiasing performance along with mitigation of range folding. The latter aspect is limited by the range of available PRIs, which, for the TDWR is from 518 μ s up to about 940 μ s (the upper limit is mainly determined by the coherency limit for accurate velocity estimation). This yields a difference in range-folding distance of about 70 km. Therefore, out-of-trip weather covering more than this range would certainly disable MSPP’s 1st trip protection mechanism. In this case, the random-phase-code approach has a better chance to eliminate the range ambiguity. For sparse, patchy out-of-trip weather, MSPP may be preferred.

An optimal solution may be to use information from the low-PRF scan to determine whether random-phase-code processing or MSPP should be used in the next scan. In most cases, there will be little out-of-

trip interference and MSPP will provide superior velocity dealiasing capability. If there is widespread out-of-trip weather, then phase-code processing method can be used instead. Further testing of both methods on real data are needed to compare their performance.

GLOSSARY

FAA	Federal Aviation Administration
MSPP	Multistaggered Pulse Processing
NEXRAD	Next Generation Weather Radar
PRF	Pulse Repetition Frequency
PSF	Program Support Facility
SQI	Signal Quality Index
TDWR	Terminal Doppler Weather Radar

REFERENCES

- Chornoboy, E. S., 1993: Clutter filter design for multiple-PRT signals, Preprints, 26th Int. Conf. Radar Meteorol., Am. Meteorol. Soc., 235-237.
- Crocker, S.C., 1988: TDWR PRF selection criteria, LL ATC-147, DOT/FAA/PM-87-25, MIT Lincoln Laboratory, Lexington, Mass., 57 pp.
- Doviak, R. J., and D. S. Zrnić, 1993: *Doppler Radar and Weather Observations*, Academic Press, San Diego, 562 pp.
- Frush, C., 1999: *NEXRAD Range-Velocity: Exploring Selected Mitigation Techniques, First Year Report*, NCAR, 34 pp.
- Frush, C., R. J. Doviak, M. Sachidananda, and D. S. Zrnić, 2002: Application of the SZ phase code to mitigate range-velocity ambiguities in weather radars, *J. Atmos. Oceanic Technol.*, 19, 413-430.
- Hildebrand, P. H., and R. S. Sekhon, 1974: Objective determination of the noise level in Doppler spectra, *J. Appl. Meteorol.*, 13, 808-811.
- Sachidananda, M., 1997: Signal Design and Processing Techniques for WSR-88D Ambiguity Resolution, Part 1, NOAA/NSSL, 100 pp.
- Sachidananda, M., 1998: Signal Design and Processing Techniques for WSR-88D Ambiguity Resolution, Part 2, NOAA/NSSL, 105 pp.
- Sachidananda, M., 2001: Signal Design and Processing Techniques for WSR-88D Ambiguity Resolution, Part 5, NOAA/NSSL, 75 pp.
- Sachidananda, M., 2002: Signal Design and Processing Techniques for WSR-88D Ambiguity Resolution, Part 6, NOAA/NSSL, 56 pp.
- 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, 1351-1362.
- Siggia, A., and R. Passarelli, 2002: SZ(8/64) and other phase codes, SIGMET Internal Notes, 6pp.
- SIGMET, 2001: RVP6 Doppler Signal Processor User's Manual, SIGMET Corp., Westford, Mass.
- Wieler, J.G., and S.-C. Hu, 1993: Elimination of Doppler ambiguities in weather radar data, 1993 IEEE National Radar Conference, 163-166.
- Zrnić, D. S., 1975: Simulation of weatherlike Doppler spectra and signals, *J. Appl. Meteorol.*, 14, 619-620.