

Christopher D. Curtis and Sebastián M. Torres

Cooperative Institute for Mesoscale Meteorological Studies, The University of Oklahoma, and
NOAA/OAR National Severe Storms Laboratory, Norman, Oklahoma

1. INTRODUCTION

Range oversampling techniques are a practical way to decrease weather radar update times without increasing estimation errors. One of these techniques, adaptive pseudowhitening, was implemented on the National Weather Radar Testbed (NWRT) during the 2010 spring season. Adaptive pseudowhitening utilizes different pseudowhitening matrices based on the estimated signal-to-noise ratio (SNR) and spectrum width at each range gate. This approach leads to better performance than a fixed transformation over widely varying conditions. On the NWRT, scan update times were reduced by roughly a factor of two without a noticeable increase in errors. At higher SNR values, there was even some improvement in data quality.

This novel real-time implementation is computationally efficient and simplifies the utilization of different transformations for all of the basic spectral moments. With the forthcoming dual polarization upgrade to the Next Generation Radar (NEXRAD) system, this novel implementation could also be extended to dual polarization variables. An overview of this real-time implementation will be presented along with some practical suggestions for SNR censoring.

2. BACKGROUND

In general, range oversampling describes the process of sampling in range at a rate greater than the inherent range resolution determined by the length of the transmit pulse. When oversampling by a factor of L , the L samples in range will be correlated because the inherent range resolution is not changed. All of the range oversampling techniques covered in this paper are based on applying a linear transformation to the correlated time series samples in order to reduce the range-time correlation. The autocorrelations computed

from the transformed samples are then averaged and processed to produce moments at the original non-oversampled resolution.

At each range gate, the linear transformation, \mathbf{W} , is applied to \mathbf{V} which is an $L \times M$ matrix of time series data. L is the oversampling factor and M is the number of pulses in the radial. The result is the transformed matrix of time series data, \mathbf{X} .

$$\mathbf{X} = \mathbf{WV} \quad (1)$$

Since \mathbf{V} is made up of both a signal and a noise component, $\mathbf{V} = \mathbf{V}_S + \mathbf{V}_N$ and $\mathbf{X} = \mathbf{WV}_S + \mathbf{WV}_N$. The linear transformation of the noise component can increase the noise resulting in noise enhancement. More details can be found in Torres and Zrnić (2003).

3. ADAPTIVE PSEUDOWHITENING

The easiest way to visualize the effects of noise enhancement is to use a time series simulation (Zrnić 1975) showing the change in the standard deviation of a spectral moment versus the signal-to-noise ratio (SNR). In this paper, the simulated time series are utilized to estimate reflectivity performance using the parameters for a surveillance cut on the NWRT. The parameters were chosen to match the weather data plotted in Section 5: number of pulses, $M = 12$, pulse repetition time or PRT, $T_S = 3$ ms, and spectrum width, $\sigma_v = 4$ m s⁻¹. The operating frequency of the NWRT is 3.2 GHz.

Fig. 1 compares traditional matched-filter based (MFB) processing against three range oversampling techniques. Whitening transformation based (WTB) processing (the blue curve) shows significantly better performance at high SNR, but performs worse than MFB processing at low SNR. One way to mitigate the poor performance for whitening at low SNR values is to use fixed pseudowhitening transformation based (PTB) processing that trades off some performance at high SNR to gain performance at low SNR. This approach is illustrated by the magenta curve. Although a fixed pseudowhitening

Corresponding author address: Christopher D. Curtis,
120 David L. Boren Blvd, NWC Room 4417, Norman,
OK, 73072; e-mail: Chris.Curtis@noaa.gov

transformation improves performance at low SNR, it still does not perform as well as MFB or WTB processing in some regions. The optimal approach would be to use the best possible transformation at each SNR so that a whitening-type transformation is used at high SNR, a matched filter-type transformation is used at low SNR, and an appropriate pseudowhitening transformation is used in between. Adaptive pseudowhitening is a practical implementation of this optimal approach.




Figure 1. Standard deviation of power estimates for matched-filter based (MFB), whitening transform based (WTB), fixed pseudowhitening transform based (PTB), and adaptive pseudowhitening transform based (APTB) processing.

Adaptive pseudowhitening uses estimates of the SNR (and of the spectrum width) to pick an optimal transformation at each range gate. The transformation is chosen based on the variance formulas in Torres et al. (2004). These formulas give expressions for the variance of all three moments in terms of the range covariance matrix, the SNR, and the normalized spectrum width. If uniform reflectivity is assumed, the range covariance matrix is known a priori. Then, MFB estimates of the SNR and normalized spectrum width can be used to produce moment-specific transformation matrices. The results from this adaptive pseudowhitening transform based (APTB) approach are shown with the green curve in Fig. 1. APTB processing performs well at all SNRs even though it uses estimates to pick the appropriate transformation. Adaptive pseudowhitening was implemented on the NWRT in the spring of 2010, and some of the details of the implementation are described in the next section.

4. REAL-TIME IMPLEMENTATION

The main difference between adaptive pseudowhitening and the single-transformation processing proposed in previous works is that a meteorological-variable-specific transformation is needed at every range resolution cell which leads to a more complicated real-time implementation. A brute-force approach to adaptive pseudowhitening is presented in this section, but issues with computational complexity and clutter filtering make it less suitable for a real-time implementation. A novel, more efficient approach is then described to address the shortcomings of the brute-force method. The computational complexity of both approaches are compared, and some additional implementation issues are addressed.

The brute force approach follows the steps briefly described in the description of adaptive pseudowhitening. Assume that at each range gate we start with an $L \times M$ matrix, \mathbf{V} , of time series (IQ) data. The final output will be the moments, in this case, reflectivity, velocity, and spectrum width.

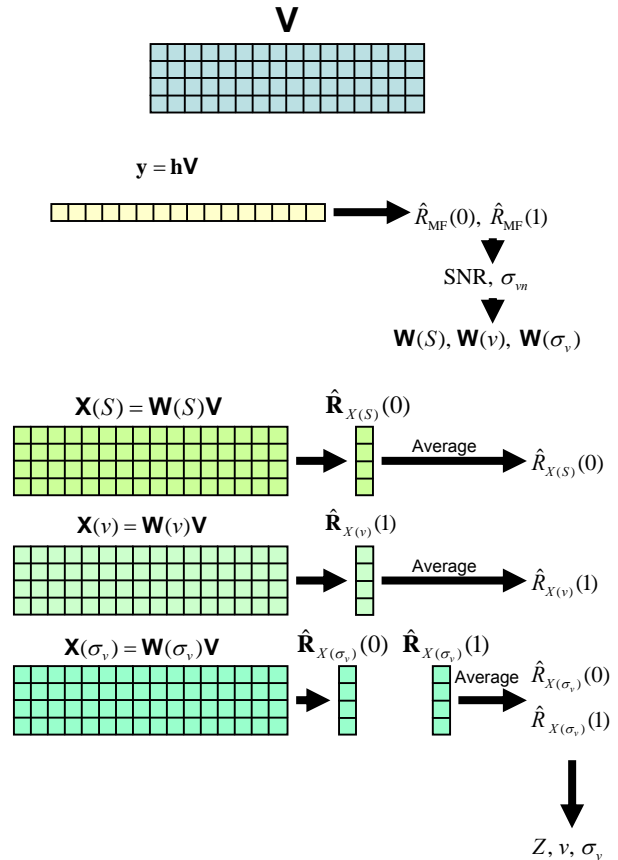


Figure 2. Graphical representation of the brute-force approach to adaptive pseudowhitening.

The first step is to compute estimates of the SNR and normalized spectrum width. This is done by producing digital matched-filtered data, \mathbf{y} , from \mathbf{V} . The SNR and normalized spectrum width are calculated from the lag-0 and lag-1 autocovariances, and three moment-specific pseudowhiting matrices are formed: $\mathbf{W}(S)$, $\mathbf{W}(v)$, and $\mathbf{W}(\sigma_v)$. The three whitening matrices are then applied to \mathbf{V} to produce three moment-specific, pseudowhitened time series matrices: $\mathbf{X}(S)$, $\mathbf{X}(v)$, and $\mathbf{X}(\sigma_v)$. Next, L autocovariances are computed for each autocovariance needed to calculate the moments. These correspond to the L rows of the time-series matrices. The L autocovariances are averaged to produce the final moment-specific autocovariances, and the moments are then calculated from these autocovariances.

For this implementation, three different time-series matrices are used to compute the moment-specific autocovariances although this does not seem especially efficient. The bigger issue is with clutter filtering, and there are basically two options. The first is to use a clutter filter that results in a filtered version of \mathbf{V} that can then be processed as before. This would come before the digital matched filter. Although this is the most efficient place to apply the clutter filter since it will result in clutter filtered versions of \mathbf{y} , $\mathbf{X}(S)$, $\mathbf{X}(v)$, and $\mathbf{X}(\sigma_v)$, it also depends on a specific type of filter that returns filtered time-series data. This restricts the type of clutter filter that can be used and necessitates a more complicated conversion back to the time domain if a spectral domain filter is utilized. The second option is to apply the clutter filter to \mathbf{y} , $\mathbf{X}(S)$, $\mathbf{X}(v)$, and $\mathbf{X}(\sigma_v)$ separately. This removes the restrictions on the clutter filter but greatly increases the computational complexity (by a factor of $3L+1$ compared to traditional processing). The first option has restrictions on the type of clutter filtering used but only increases the computational complexity by a factor of L compared to traditional processing. These issues led to a search for a more efficient implementation.

The basic idea behind the efficient adaptive pseudowhiting implementation is to split the pseudowhiting matrix into two parts. In this way, the part that is common to all of the transformations, $\mathbf{W}(S)$, $\mathbf{W}(v)$, and $\mathbf{W}(\sigma_v)$, can be applied first before clutter filtering and then the moment-specific part can be applied later. This can be seen by looking at the definition of the transformation matrix described in Torres et al. (2004):

$$\mathbf{W} = \gamma \mathbf{\Sigma}^{\dagger} \mathbf{P}^{*T} \quad (1)$$

where γ is the transformation gain, $\mathbf{\Sigma}^{\dagger}$ is a diagonal real-valued matrix, and \mathbf{P} is a unitary matrix from the eigenvalue decomposition of the normalized range-correlation matrix, \mathbf{C}_v . The mathematical derivations for this and the rest of the efficient implementation can be found in Curtis and Torres (2010). The \mathbf{P}^{*T} matrix is common to all three of the $\mathbf{W}(S)$, $\mathbf{W}(v)$, and $\mathbf{W}(\sigma_v)$ matrices, but the $\gamma \mathbf{\Sigma}^{\dagger}$ part is moment-specific. The efficient implementation applies both of these parts separately.

Fig. 3 describes the steps of the efficient implementation of adaptive pseudowhiting.




Figure 3. Graphical representation of the efficient approach to adaptive pseudowhiting.

The efficient implementation starts by forming the partially-transformed times series matrix

$$\tilde{\mathbf{X}} = \mathbf{P}^{*T} \mathbf{V}. \quad (2)$$

The clutter filter can be applied to the partially-transformed matrix with the only restriction on the filter being that L lag-0 and lag-1 autocovariances can be computed from the rows of $\tilde{\mathbf{X}}$. The first row of $\tilde{\mathbf{X}}$ corresponds to the matched-filtered data so the SNR and normalized spectrum width can be computed from the properly-scaled, first value of the autocovariances. This avoids the extra step of computing the matched-filtered vector, \mathbf{y} . From the estimates, moment-specific weight vectors, $\mathbf{d}(S)$, $\mathbf{d}(v)$, and $\mathbf{d}(\sigma_v)$, can be computed that correspond to the second part of the transformation. The moment-specific autocovariances can be calculated from a weighted average of the autocovariances computed from $\tilde{\mathbf{X}}$. Finally, the moments are calculated from the moment-specific autocovariances as in the brute-force approach.

The efficient approach addresses both of the issues mentioned earlier. The first is that only one matrix of time-series data needs to be computed from \mathbf{V} instead of three. The second is that there are no longer restrictions on the type of clutter filter that can be used as long as the lag-0 and lag-1 autocovariances are computed. This reduces the computational complexity of the efficient implementation by a factor of two compared to the brute-force approach (without taking clutter filtering into account). In the end, the efficient implementation increases the computational complexity by a factor of L over traditional processing with or without clutter filtering.

The last practical issue that needs to be addressed is data censoring (or thresholding). Data censoring is a data quality issue whereby only data from significant weather returns are utilized (e.g., displayed and sent to algorithms) and data from noise-like returns are not. A common way to censor weather radar data is to use an SNR threshold. Data corresponding to signals above a particular SNR threshold are treated as significant, and data below the threshold are treated as non-significant or noise. For adaptive pseudowhitening, this is complicated by the fact that the noise powers are different for each range resolution cell and the corresponding SNR threshold should be adjusted because of the lower variance of APTB estimates. For our implementation, we chose a simple approach and used the SNR computed from the MFB data and a fixed threshold. This censors the data in the same

way as the matched-filter approach, but the estimates that are designated as significant will often have better quality than the ones computed with the MFB data. By taking advantage of the improved estimates from the APTB processing, it may be possible to preserve more data that are considered significant, but this will require additional research.

5. NWRT WEATHER DATA




Figure 4. Comparison of digital matched filter (top) and adaptive pseudowhitening (bottom) using NWRT data from 10:54 UTC, April 2, 2010.

Fig. 4 shows part of two 90° sectors collected using the NWRT and processed using a digital matched filter (top) and adaptive pseudowhitening (bottom). The images should look similar, but the key difference is that the total collection time at each beam position was 99.2 ms for the digital matched filter and only 56.8 ms when using adaptive pseudowhitening (this includes the time for collecting data for all of the moments). Overall, a scanning strategy was developed for adaptive pseudowhitening that takes about 54% of the time of a traditional scan. Thus, similar data quality can be produced in roughly half the time on the NWRT when using adaptive pseudowhitening and the corresponding scanning strategy.

6. CONCLUSIONS

Efficient adaptive pseudowhitening is less computationally complex than a brute-force approach and simplifies the application of a clutter filter. The real-time implementation was used in the spring of 2010 to produce data on the NWRT with similar data quality in roughly half the time. This approach will also extend nicely to dual polarization variables and could be used to increase data quality on the NEXRAD network after the dual polarization upgrade is completed.

7. REFERENCES

- Curtis, C. D., and S. M. Torres, 2010: Adaptive Range Oversampling to Achieve Faster Scanning on the National Weather Radar Testbed Phased Array Radar. *J. Atmos. Oceanic Technol.*, submitted for publication.
- Torres, S. M., C. D. Curtis, J. R. Cruz, 2004: Pseudowhitening of weather radar signals to improve spectral moment and polarimetric variable estimates at low signal-to-noise ratios. *IEEE Transactions on Geoscience and Remote Sensing*, 42, 941-949.
- Torres, S. M., and D. S. Zrnić, 2003: Whitening in range to improve weather radar spectral moment estimates. Part I: Formulation and simulation. *J. Atmos. Oceanic Technol.*, **20**, 1433–1448.
- Zrnić, D. S., 1975: Simulation of weatherlike Doppler spectra and signals. *J. Appl. Meteor.*, **14**, 619–620.

ACKNOWLEDGEMENTS

This conference paper was prepared by Christopher D. Curtis with funding provided by NOAA/Office of Oceanic and Atmospheric Research under NOAA-University of Oklahoma Cooperative Agreement #NA17RJ1227, U.S. Department of Commerce. The statements, findings, conclusions, and recommendations are those of the authors and do not necessarily reflect the views of NOAA or the U.S. Department of Commerce.