

Revisiting the Optimum Receiver Filter Bandwidth for Range-Oversampling Processing

SEBASTIÁN M. TORRES AND CHRISTOPHER D. CURTIS

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

(Manuscript received 12 April 2019, in final form 10 October 2019)

ABSTRACT

For weather radars, range-oversampling processing was proposed as an effective way either to reduce the variance of radar-variable estimates without increasing scan times or to reduce scan times without increasing the variance of estimates. Range oversampling entails acquiring the received signals at a rate L times as fast as the reciprocal of the pulse width (the conventional rate), where L is referred to as the range-oversampling factor. To accommodate the L -times-as-fast sampling, the original formulation of range-oversampling processing required a receiver filter with a bandwidth L times as wide as that of the matched filter (the conventional receiver filter). In this case, the noise at the output of the receiver filter can still be assumed to be white, resulting in a simplified formulation of the technique but also, and more important, in a more difficult practical implementation since the receiver filter in operational weather radars is typically matched to the transmitted pulse. In this work, we revisit the role of the receiver filter in the performance of range-oversampling processing and show that using a receiver matched filter not only facilitates the implementation of range-oversampling processing but also results in the lowest variance of radar-variable estimates.

1. Introduction

Weather radars typically sample the received signals at a rate given by the inverse of the pulse width. Range-oversampling processing operates on time series samples that are acquired at an L -times-as-fast rate, where L is the range-oversampling factor. A generalized model for this type of processing involves two stages: transformation and estimation (Torres and Curtis 2012). In the transformation stage, an L -by- M matrix of time series samples \mathbf{V} is transformed as $\mathbf{X} = \mathbf{T}\mathbf{V}$, where \mathbf{T} is a complex-valued L -by- L transformation matrix. Here, \mathbf{V} contains a dwell worth of time series data [i.e., the in-phase and quadrature (IQ) signals], where L is the number of consecutive samples in range time, and M is the number of samples (or transmitted pulses) in the dwell time. For the estimation stage, sets of L correlations (one set for each required lag) are estimated from \mathbf{X} . The L correlation estimates from each set are then incoherently averaged and used to compute radar-variable estimates with reduced variance. For example, a whitening transformation \mathbf{T} produces uncorrelated samples and leads to maximum variance reduction when the signal-to-noise ratio (SNR) is high. A whitening transformation can be easily obtained from the

range correlation of the oversampled signals (Torres and Zrnić 2003), which depends on the transmitter pulse, the distribution of scatterers in the resolution volume, and the receiver filter. Under the assumption of uniform distribution of scatterers in the resolution volume (i.e., no reflectivity gradients), the range correlation is only a function of the modified pulse, which is defined as the convolution of the transmitter pulse and the receiver-filter impulse response. Unfortunately, in addition to decorrelating the signal, the whitening transformation increases the noise power; at low SNR, this can offset the benefits of averaging independent estimates. The noise-power gain of the transformation is referred to as the noise enhancement factor (NEF).

The original assumption used in the development of range-oversampling processing was that the receiver bandwidth would be increased by the range-oversampling factor (i.e., L times). This introduces negligible range correlation in the noise, which is typically white at the input of the receiver filter (Torres and Zrnić 2003). The white-noise assumption was used to derive theoretical expressions for the variance of radar-variable estimates and to quantify the performance of range-oversampling processing in terms of variance reduction. Also, theoretical variance expressions based on the white-noise assumption were used to develop

Corresponding author: Sebastián Torres, sebastian.torres@noaa.gov

DOI: 10.1175/JTECH-D-19-0057.1

© 2020 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the AMS Copyright Policy (www.ametsoc.org/PUBSReuseLicenses).

closed-form solutions for adaptive pseudowhitening (Curtis and Torres 2014). The adaptive-pseudowhitening technique uses estimates of the input-signal characteristics to select an optimal pseudowhitening transformation that produces radar-variable estimates with minimum variance by limiting the amount of noise enhancement at the cost of less decorrelation. For example, adaptive pseudowhitening uses a whitening transformation at high SNRs when the noise enhancement effect is negligible and gradually transitions into a matched-filter-like transformation at low SNRs to reduce the noise power when it is more critical.

In general, the adaptive-pseudowhitening transformation \mathbf{T} is different for every radar variable. Therefore, a direct application of these different transformations results in multiple versions of the time series data, and other signal-processing techniques that operate on the time series data, such as ground-clutter filtering, must be applied to each version. Rather than operating on several versions of time series data, an efficient implementation was developed to reduce the computational complexity (Curtis and Torres 2011). In the efficient implementation, there is a single transformation stage for all the radar variables that consists of a decorrelating transformation, and the estimation stage utilizes a weighted average specific to each radar variable. Thus, the transformation stage is given by $\mathbf{X} = \mathbf{U}^* \mathbf{T} \mathbf{V}$, where the decorrelating transformation \mathbf{U}^{*T} is a unitary matrix that comes from the eigendecomposition of the range correlation of oversampled signals \mathbf{C}_V . In mathematical terms, $\mathbf{C}_V = \mathbf{U}^* \mathbf{A} \mathbf{U}^T$, where the asterisk and T superscripts denote complex conjugation and matrix transpose, respectively. In practice, \mathbf{C}_V can be measured in real time from the data using the procedure described by Curtis and Torres (2013). The vector of radar-variable-specific adaptive weights \mathbf{d}_θ depends on estimated signal characteristics and is used to average the range-oversampled correlation estimates needed for the computation of each radar variable. That is, $R_X(k) = \sum_{l=0}^{L-1} d_\theta(l) R_X^{(l)}(k)$, where θ is a generic radar variable, k is the correlation lag, and $R_X(k)$ is any of the correlation estimates needed to compute θ . Therefore, $R_X(k)$ is calculated from a weighted average of L range-oversampled correlations $R_X^{(l)}(k)$ [cf. Eq. (10) in Curtis and Torres (2011)].

Although adaptive pseudowhitening produces data with improved quality for a wide range of signal characteristics, the original implementation relies on explicit expressions of the variances of radar-variable estimates. To remove this limitation, Curtis and Torres (2017) developed a procedure that extends adaptive pseudowhitening to be compatible with nontraditional radar-variable estimators, which do not have explicit variance expressions.

That is, instead of computing the adaptive weights for the estimation stage using analytical solutions based on explicit variance expressions, the so-called lookup-table (LUT) adaptive-pseudowhitening technique obtains the adaptive weights from tables created through simulations. As a side benefit, LUT adaptive pseudowhitening removes the need to assume the traditional model of a signal in additive *white* noise. Because the lookup tables are derived from simulations, the noise can be modeled as having any range correlation. Therefore, with the introduction of LUT adaptive pseudowhitening, it is possible to forgo the white-noise assumption and consider the use of receiver filters with narrower bandwidths, even if they introduce range correlation in the noise.

The receiver filter has an important role in determining the performance of range-oversampling processing. It contributes to the range correlation of the signal, and it determines the range correlation of the noise. Furthermore, its bandwidth determines the noise power in the baseband time series data. Also, narrower receiver bandwidths introduce additional correlation in the signal. Because the signal is more correlated, it is more “difficult” to decorrelate, and this results in a larger NEF. However, because the receiver bandwidth is narrower, it lets less noise through. In this work, we examine this trade-off to determine the optimum receiver bandwidth when using LUT adaptive pseudowhitening.

The rest of paper is organized as follows. Section 2 describes the necessary modifications to the LUT adaptive-pseudowhitening technique to accommodate correlated noise. In section 3, realistic simulations that include the effects of the receiver filter are used to identify the optimum receiver bandwidth for LUT adaptive pseudowhitening. These results are verified in section 4 using data collected with the research KOUN radar in Norman, Oklahoma. Section 5 concludes with a summary of findings and recommendations for the implementation of range oversampling on operational weather radars.

2. Modifications to LUT adaptive pseudowhitening

The current implementation of LUT adaptive pseudowhitening assumes that the noise prior to the transformation stage is white. However, in this work, we consider the use of receiver filters with narrower bandwidths, which may result in noise that is correlated in the range-time dimension. Hence, it is important to first identify any modifications to the technique needed to accommodate this departure from initial assumptions.

As mentioned before, LUT adaptive pseudowhitening relies on lookup tables to determine the adaptive weights for the estimation stage of range-oversampling

processing. A one-time process to generate lookup tables using simulations was described by [Curtis and Torres \(2017\)](#). In essence, the original steps involve 1) simulating a large number of realizations of range-oversampled weather signals with prescribed range correlation and additive white noise, and 2) finding optimal adaptive weights that minimize the mean-square error of radar-variable estimates for each set of signal characteristics. Thus, to accommodate the use of receiver filters with narrower bandwidths, the simulation step must be modified to use correlated noise instead of white noise. A relatively simple way to simulate noise with the proper range correlation is to pass white noise through the desired receiver filter. This can be accomplished by convolving a simulated white noise sequence with the receiver-filter impulse response. To complete the simulation step, the resulting correlated noise must be added to the weather signal.

Further, the LUT adaptive-pseudowhitening processing steps include the computation of radar variables using both a digital matched filter (DMF) and appropriate pseudowhitening transformations. Some of these radar variables require a noise-power estimate (e.g., reflectivity), which is obtained from the measured input noise power scaled by the transformation NEF. In general, the NEF is defined as the ratio of the noise powers before and after range-oversampling processing. A closed-form expression for the NEF when using adaptive pseudowhitening can be derived assuming only noise is present at the input. Let \mathbf{V}_N be an L -by- M matrix of time series noise samples after the receiver filter but before adaptive-pseudowhitening processing; these samples could potentially be correlated. Assuming the noise samples are zero mean, the L -by- L range correlation matrix is defined as $\mathbf{R}_{V_N} = E[\mathbf{V}_N^* \mathbf{V}_N^T]$, where $E[\cdot]$ denotes expected value, and the corresponding normalized range correlation matrix is $\mathbf{C}_{V_N} = \mathbf{R}_{V_N}/N_V$, where N_V is the noise power before processing. After the transformation stage, $\mathbf{X}_N = \mathbf{U}^{*T} \mathbf{V}_N$, the range correlation of the transformed data is given by

$$\begin{aligned} \mathbf{R}_{\tilde{\mathbf{X}}_N} &= E[\tilde{\mathbf{X}}_N^* \tilde{\mathbf{X}}_N^T] = E[\mathbf{U}^T \mathbf{V}_N^* \mathbf{V}_N^T \mathbf{U}^*] = \mathbf{U}^T \mathbf{R}_{V_N} \mathbf{U}^* \\ &= N_V (\mathbf{U}^T \mathbf{C}_{V_N} \mathbf{U}^*). \end{aligned}$$

The diagonal elements of this matrix are the L noise powers or lag-0 range correlations $R_{\tilde{\mathbf{X}}_N}^{(l)}(0)$ described in the previous section. Hence, the noise power after the weighted average in the estimation stage is

$$N_X = R_{X_N}(0) = \sum_{l=0}^{L-1} d_\theta(l) R_{\tilde{\mathbf{X}}_N}^{(l)}(0) = \sum_{l=0}^{L-1} d_\theta(l) (\mathbf{R}_{\tilde{\mathbf{X}}_N})_{ll},$$

where $d_\theta(l)$ are the corresponding radar-variable-specific adaptive weights obtained from the lookup tables, and

$(\mathbf{A})_{ll}$ denotes the l th element in the diagonal of matrix \mathbf{A} . We can now compute the NEF for a variable θ as follows:

$$\begin{aligned} \text{NEF}(\theta) &= \frac{N_X}{N_V} = \frac{1}{N_V} \sum_{l=0}^{L-1} d_\theta(l) (\mathbf{R}_{\tilde{\mathbf{X}}_N})_{ll} \\ &= \sum_{l=0}^{L-1} d_\theta(l) (\mathbf{U}^T \mathbf{C}_{V_N} \mathbf{U}^*)_{ll}. \end{aligned} \quad (1)$$

If the noise is white, \mathbf{C}_{V_N} and $\mathbf{U}^T \mathbf{U}^*$ are both identity matrices (\mathbf{U} is unitary), and Eq. (1) reduces to

$$\text{NEF}(\theta) = \sum_{l=0}^{L-1} d_\theta(l), \quad (2)$$

which agrees with Eq. (13) in [Curtis and Torres \(2011\)](#).

In summary, two modifications are needed to use LUT adaptive pseudowhitening with range-correlated noise: 1) properly simulating the range-correlated noise when producing the lookup tables and 2) using the correct noise power when estimating the radar variables in the processing stage.

3. Simulation analysis

Having a modified version of the LUT adaptive-pseudowhitening technique that can accommodate range-correlated noise, we now use realistic simulations to quantify the impacts of using receiver filters with different bandwidths: from the conventional matched filter, to a bandwidth that is L times as wide. We begin by describing the signal-simulation approach used both for the one-time generation of lookup tables and for the actual performance evaluation.

a. Signal simulation

There are two main considerations for the simulation of baseband range-oversampled IQ signals corresponding to different receiver filters. The first one is that noise samples at the output of the receiver filter may be correlated in range. That is, the noise at the input of the receiver filter is white but becomes correlated after being filtered. The second consideration is that receiver filters in typical modern radar systems operate on signals that are sampled at the intermediate frequency (IF). Thus, a realistic simulation that includes the effects of practical receiver filters should use signals sampled at the IF. In summary, the signal simulation procedure involves three steps: 1) the simulation of conventional weather signals in additive white noise, 2) the convolution of these signals with the impulse response of the desired receiver filter, and 3) the decimation of the resulting signal by an appropriate factor to produce IQ

samples at the baseband (range oversampled) rate. This process is described next in more detail.

Using methods analogous to those developed by [Curtis \(2018\)](#), we start by simulating R sets of $L \times D + N_t$ independent realizations of M -sample, dual-polarization, weather signals with Gaussian Doppler spectrum and predefined signal characteristics (i.e., SNR, Doppler velocity, spectrum width, differential reflectivity, differential phase, and correlation coefficient). Here, R is the number of independent measurements used to compute radar-variable estimate statistics, which needs to be large enough to provide robustness to the empirical statistics obtained from the Monte Carlo simulation; L is the desired range-oversampling factor; and D is the required decimation factor to go from IF to baseband sampling. Last, $N_t = N_p + F - 2$ is the number of extra samples needed to accommodate the convolution with the N_p -sample transmitter pulse and the F -tap finite-impulse-response (FIR) receiver filter. Note that the assumed range-time sampling rate is given by the IF, and the sample-time sampling rate is given by the pulse repetition time (PRT).

The processing described next is performed in the range-time dimension and is independently repeated along the sample-time and polarization-channel dimensions. First, the weather signals are convolved with the N_p -sample transmitter pulse (also sampled at the IF rate) to produce R sets of $L \times D + F - 1$ signals with the required range correlation at the input of the receiver filter. Next, R sets of $L \times D + F - 1$ independent realizations of M -sample white noise for each polarization channel (horizontal and vertical) are simulated and added to the weather signals. These signal-plus-noise samples are convolved with the F -sample receiver-filter impulse response; this completes the *shaping* of the weather signal's range correlation and imposes the corresponding range correlation to the noise. Last, the resulting R sets of $L \times D$ receiver-filter outputs are decimated by a factor of D to obtain R sets of L samples at the desired range-oversampling rate of $cL/(2\Delta r)$, where c is the speed of light and Δr is the desired range spacing of the radar variables (i.e., after signal processing). In practice, D must be chosen so that $f_{IF} = cLD/(2\Delta r)$, which makes sense because the radar IF is typically selected as an integer multiple of the reciprocal of the required IQ range-time sample spacing $[2\Delta r/(cL)]$. Ultimately, this process produces R independent sets of L -by- M dual-polarization signal-plus-noise samples with appropriate range, sample, and cross-channel correlations.

For our analysis, we selected the IF receiver filter that matches the one used on the Weather Surveillance Radar-1988 Doppler (WSR-88D); it is an FIR filter designed using the well-known window method (e.g.,

[Oppenheim and Schaffer 1989](#)). This method starts with the impulse response of an ideal low-pass filter:

$$h(n) = \frac{\sin(\pi Bn)}{\pi n}, \quad n \in \mathbb{Z}, \quad (3)$$

where B is the normalized bandwidth of the low-pass filter in hertz. Because this filter is noncausal and infinitely long, it is truncated and tapered with a window. Thus, the desired FIR filter taps are obtained as

$$h_{\text{FIR}}(n) = w[n - (F - 1)/2]h[n - (F - 1)/2], \\ n = 0, 1, \dots, F - 1, \quad (4)$$

where F is the length of the FIR filter (assumed to be odd to simplify this description), w is the F -sample window (the signal processor on the WSR-88D uses the Hamming window), and the indexing on the right-hand side takes care of the necessary shift to make the filter causal. Note that, because of the application of a window, the resulting filter bandwidth is generally larger than B . Because we are looking for filters with specific bandwidths, we automated the filter design by means of a minimization process where the objective function is the square of the difference between the actual and desired filter bandwidths.

b. Lookup-table generation

In general, the lookup tables for adaptive pseudowhitening are robust to changes in the modified pulse ([Curtis and Torres 2017](#)). However, for different receiver filters, both the modified pulse and the range correlation of the noise are different. As will be verified later, this requires unique sets of lookup tables. Thus, for our analysis, we generated lookup tables for each receiver filter using the procedure introduced by [Curtis and Torres \(2017\)](#) with the modification in [section 2](#). That is, the simulation step was carried out using the procedure described in [section 3a](#) above. Despite the fact that we had to generate several sets of lookup tables (producing each set of lookup tables takes about 1 h in a conventional desktop computer running the MATLAB software), this procedure has to be done only once: the generated lookup tables are saved, and they are simply recalled during LUT adaptive-pseudowhitening processing.

c. Performance evaluation

Here, we use the simulation framework described in [section 3a](#) to determine the optimum receiver bandwidth for LUT adaptive pseudowhitening. In this context, optimality is assessed in terms of the variance of radar-variable estimates obtained after processing.




FIG. 1. Frequency response of the five simulated receiver filters. The filter bandwidth increases with the $B_6\tau$ product. Assuming white noise at the input, the noise power at the output of these filters also increases with the $B_6\tau$ product.

That is, the optimum receiver bandwidth is the one that leads to the smallest variance of estimates when using LUT adaptive-pseudowhitening processing. To do this, we simulated signals as described in section 3a, generated lookup tables as described in section 3b, and produced radar-variable estimates using the LUT adaptive pseudowhitening processing with the modifications described in the latter part of section 2.

The receiver filter was designed with $F = 200$ taps, and we used the nominal $1.54\text{-}\mu\text{s}$ WSR-88D “short” transmitter pulse with 200-ns rise and fall times. We repeated the simulation for five different receiver filters with $B_6\tau$ products of 1, 2, 3, 4, and 5, where B_6 is the 6-dB receiver-filter bandwidth and τ is the 6-dB transmitter pulse width. Note that $B_6\tau = 1$ corresponds to an approximately matched receiver filter, which is typically used on weather radars (Zrnić and Doviak 1978); $B_6\tau = 5$ corresponds to a receiver filter with a bandwidth that is 5 times as wide as the matched filter (i.e., it is wider by the range-oversampling factor) as was originally proposed to accommodate range-oversampling processing. Figure 1 shows the frequency responses of the five receiver filters in the simulation; Fig. 2 depicts the corresponding modified pulses, and Fig. 3 shows the corresponding normalized range correlations for range-oversampled weather signals. As expected, when the receiver filter bandwidth gets smaller, the modified pulse has smoother transitions, and the range correlation at higher lags is larger.

For the simulation, we started with signals sampled at the WSR-88D IF of $f_{\text{IF}} = 95\,915\,167$ Hz and ended with radar variables spaced by $\Delta r = 250$ m after processing.




FIG. 2. Modified pulse corresponding to each of the five simulated receiver filters when using the “narrow” $1.54\text{-}\mu\text{s}$ transmitted WSR-88D pulse. The modified pulse becomes smoother as the $B_6\tau$ product decreases.

Also, we utilized a range-oversampling factor of $L = 5$, which is the recommended value for a future operational implementation of range-oversampling processing. Thus, the decimation factor was obtained as $D = 2\Delta r f_{\text{IF}} / (cL) = 32$, which leads to baseband IQ samples with 50-m range spacing (i.e., 5-times oversampling in range). Signals were simulated ($R = 7000$) with the dwell parameters of the lowest elevation angle in NEXRAD’s Volume Coverage Pattern 12. That is, for the estimation




FIG. 3. Normalized range correlation of the weather signal corresponding to each of the five simulated receiver filters. The range correlation at higher lags increases as the $B_6\tau$ product decreases. For a given transformation, the noise enhancement factor also increases as the $B_6\tau$ product decreases.

of reflectivity and the polarimetric variables, we used $M = 15$ samples per dwell and a PRT of 3.1 ms (leading to a Nyquist velocity $v_a = 8.6 \text{ m s}^{-1}$). For the estimation of velocity and spectrum width, we used $M = 40$ and a PRT of 1 ms ($v_a = 27.1 \text{ m s}^{-1}$). For the weather signal, the Doppler velocity was fixed at 0 m s^{-1} , the spectrum width was fixed at 2 m s^{-1} , the differential reflectivity was fixed at 0 dB, the differential phase was fixed at 0° , and the correlation coefficient was fixed at 0.99. Note that the arbitrary choice of these dwell parameters and signal characteristics does not preclude us from generalizing our results. Whereas changes to these parameters may result in changes to the variance of radar-variable estimates, in this work, we are just interested in the relative performance of different receiver filters with respect to range-oversampling processing, both of which operate solely in the range-time dimension. Thus, neither dwell parameters (M and the PRT) nor sample-time signal characteristics (Doppler velocity, spectrum width, and the polarimetric characteristics) influence our results. The only relevant parameters for our analysis are the SNR, the range correlation of the baseband IQ weather and noise samples, and the lookup tables used in range-oversampling processing.

Figure 4 shows the variance improvement factor as a function of the SNR for estimates of reflectivity, velocity, spectrum width, differential reflectivity, differential phase, and correlation coefficient obtained with LUT adaptive-pseudowhitening processing for each of the five receiver filters with respect to estimates obtained with a receiver filter with 5-times-as-wide bandwidth. The SNR prior to the receiver filter ranges from -20 to 24 dB in steps of 2 dB , and we assumed that the noise powers in the horizontal and vertical channels are equal. For the abscissa, we chose the SNR after the receiver filter with $B_6\tau = 1$ since it is the one conventionally used in operations for radar-data thresholding (censoring). It can be seen from this figure that the best performance (i.e., the largest variance improvement factor or, equivalently, the smallest variance) for all radar variables is consistently achieved when the receiver filter is matched to the width of the transmitter pulse (i.e., $B_6\tau = 1$ and the solid blue line). In addition, the dashed blue curve shows the performance for $B_6\tau = 1$ with lookup tables obtained using white noise instead of correlated noise. This corresponds to using the original lookup-table generation procedure that does not take into account the range correlation of the noise in the simulations. The observed suboptimal performance (relative to the solid blue curve) comes from using lookup tables generated with white noise and applying them to data with correlated noise, which illustrates the need for the modifications to the lookup-table generation procedure

as described in section 2. Although not shown here, the optimality of the matched filter was also observed for different dwell parameters (M and v_a) and signal characteristics. We verified this by spot-checking several cases using relatively extreme parameters; in all of them, using a matched filter resulted in the largest variance improvement factor.

The optimality of the matched filter was indeed a surprising finding given that the original assumption was that the receiver filter bandwidth had to increase by the range-oversampling factor to accommodate the faster receiver sampling rates required for range oversampling. It turns out that the reduction in noise power provided by the matched filter more than makes up for the increase in noise enhancement due to the more correlated signal. On the basis of these results, upgrading a radar system to incorporate range oversampling should require no modifications to the receiver filter, which is already matched to the transmitted pulse in most systems.

4. Real-data analysis

In this section, we empirically validate the results from section 3 using real data. As mentioned before, the results in section 3 only depend on the SNR, the range correlation of the weather and noise samples, and the lookup tables used for processing. Of these, only the SNR depends on the weather scenario. Thus, a single case with a wide range of SNR values is sufficient to validate the simulation results. At approximately 0103 UTC 1 December 2018, we collected range-oversampled IQ data from severe storms in central Oklahoma using the S-band, dual-polarization, research KOUN radar in Norman. Two sets of 50-m oversampled data ($L = 5$) were collected with both a matched filter ($B_6\tau = 1$) and a filter with 5-times-as-wide bandwidth ($B_6\tau = 5$). Thus, after range-oversampling processing, the range gate spacing was 250 m. The antenna was stationary and pointed at an azimuth of 193° to capture a wide range of SNRs while avoiding range-overlaid echoes. Under the assumption that storms do not significantly change in the few seconds that it takes to collect the data, this mode of operation provides a practical means to obtain multiple dwells from the same storm configuration. These dwells can be thought of as approximate independent “realizations” of IQ signals that can be processed to extract meaningful statistical information about radar-variable estimates. A PRT of 1 ms was employed, resulting in a Nyquist velocity $v_a = 27.7 \text{ m s}^{-1}$. We collected and processed 200 dwells of data for each receiver-filter configuration using 16 samples per dwell. That is, for each receiver configuration,

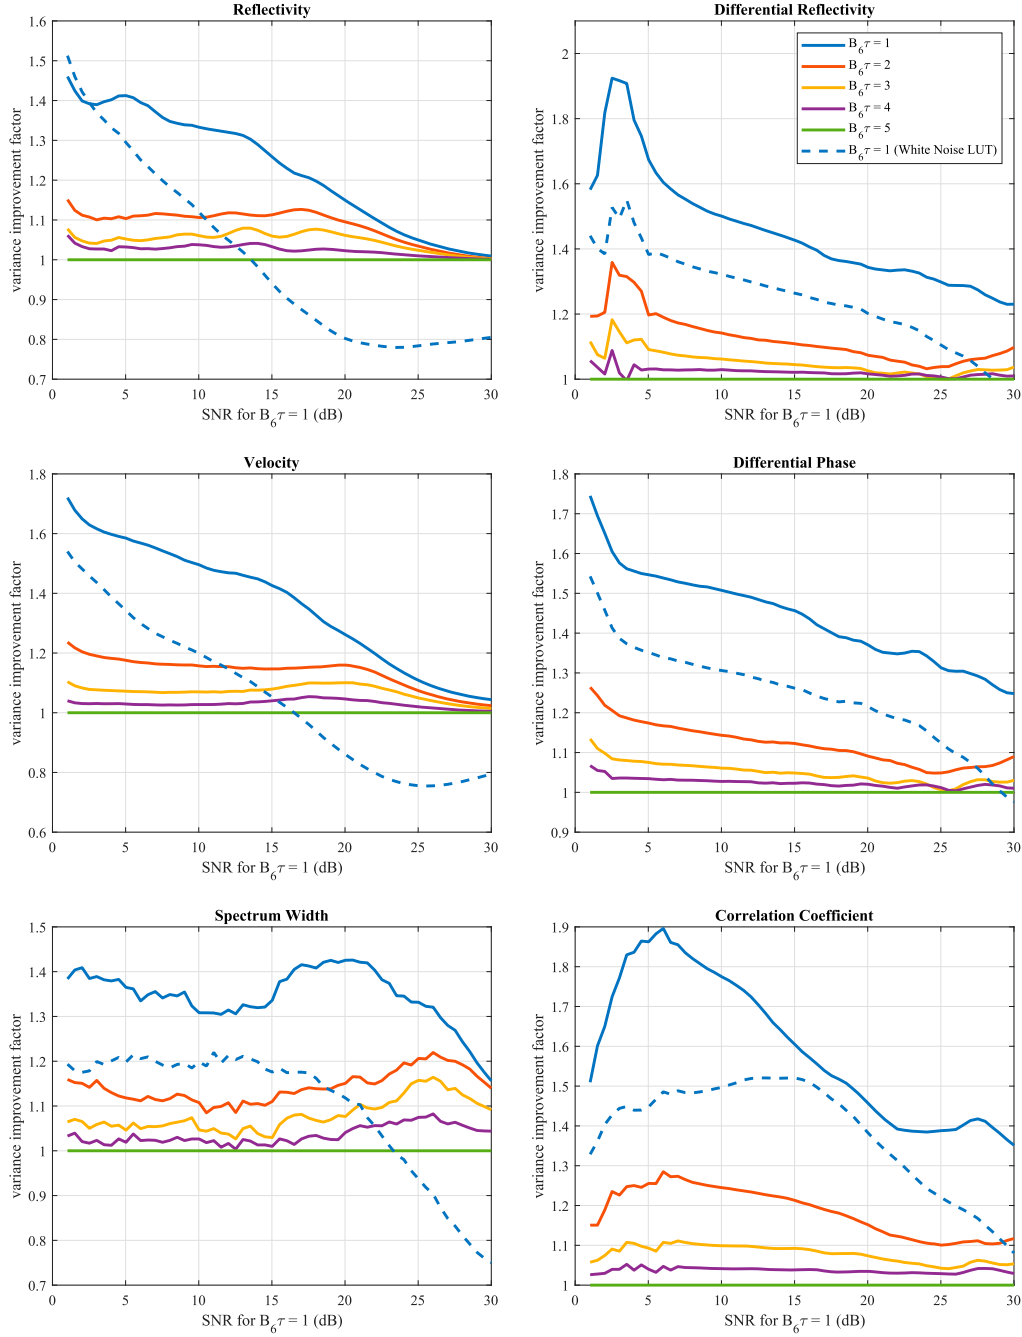


FIG. 4. Variance improvement factor corresponding to estimates of (top left), reflectivity, (middle left), Doppler velocity, (bottom left) spectrum width, (top right) differential reflectivity, (middle right) differential phase, and (bottom right) correlation coefficient using LUT adaptive-pseudowhening processing with different receiver filters as a function of the SNR at the output of the receiver matched filter. The variance improvement factor is computed relative to the performance of the receiver filter with a 5-times-as-wide bandwidth. The solid curves correspond to using LUTs generated for the appropriate receiver filter with the modifications introduced in section 2 (i.e., properly simulating range-correlated noise at the output of the receiver filter). The dashed curve corresponds to a receiver matched filter ($B_6\tau = 1$) but using LUTs generated under the (incorrect) assumption of white noise at the output of the receiver filter.

the collection time was 3.2 s; however, there were approximately 45 s between collections that allowed for the reconfiguration of the receiver filter. Last, we measured the range correlation (Curtis and Torres 2013) from each set of data to ensure that the decorrelation transformation (\mathbf{U}^{*T}) was accurately computed from the appropriate \mathbf{C}_V for both cases. The same lookup tables that were computed for the different filter bandwidths in section 3 were used to implement LUT adaptive pseudowhitening ($B_6\tau = 1$ and $B_6\tau = 5$).

Figure 5 shows the mean SNR (Fig. 5, top panel), the reflectivity (Fig. 5, middle panel), and the standard deviation of reflectivity estimates (Fig. 5, bottom panel) corresponding to the data from both receiver filters. The results are plotted for 100 range gates located between 57.5 and 82.25 km to capture a wide range of SNR values. To account for advection between collections, the data from the second collection (corresponding to the larger receiver bandwidth) were shifted by four range gates (1 km). The SNR plots show that the average signal power was similar for both cases with some minor differences related to echo fluctuations and storm evolution in the time between collections. To compare estimation errors, the standard deviation of reflectivity estimates was selected to ensure that the power differences would not affect the results. From 57.5 km (the start of the data) to about 70 km, the narrower, matched filter bandwidth results in better performance. Beyond about 70 km, the performance is similar for both receiver filters. This is consistent with the simulation results in Fig. 4, where the $B_6\tau = 1$ standard deviation is noticeably smaller for SNRs below 20 dB but is comparable to the $B_6\tau = 5$ standard deviation above 20 dB. In addition to agreeing with the simulation results, this corroborates that range-oversampling processing can be implemented without having to increase the receiver bandwidth. If LUT adaptive pseudowhitening is used, the performance is better with the receiver matched filter, especially at lower SNRs.

5. Conclusions

We investigated the role of the receiver filter in the performance of range-oversampling processing and determined that a receiver matched filter leads to the best performance in terms of the variance of radar-variable estimates and, more importantly, leads to a simpler practical implementation. However, whereas previous formulations of range-oversampling processing relied on the assumption that the noise was white prior to the transformation stage, the use of a receiver matched filter on range-oversampled data leads to correlated noise. Fortunately, the recently developed

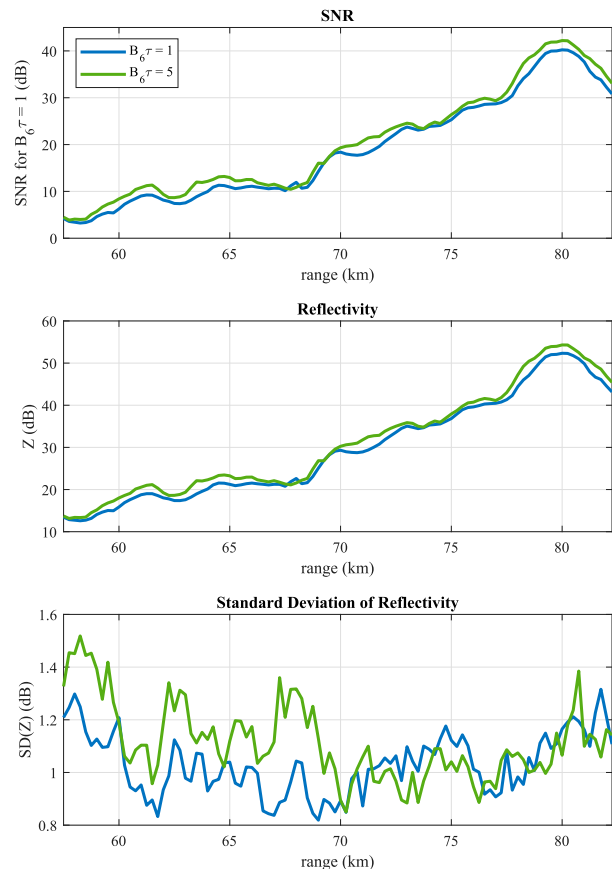


FIG. 5. Data collected with the KOUN radar on 1 Dec 2018 using a stationary antenna with a receiver matched filter ($B_6\tau = 1$) and a filter with 5-times-as-wide bandwidth ($B_6\tau = 5$). For each case, 200 dwells with $M = 16$ samples each were processed using LUT adaptive pseudowhitening. Shown are (top) the ratio of the mean signal power to the matched-filter noise power (same as the abscissa in Fig. 4), (middle) the corresponding reflectivity estimates, and (bottom) the standard deviation of reflectivity estimates, all as a function of range in kilometers. The range-gate spacing is 250 m.

LUT adaptive-pseudowhitening technique can accommodate correlated noise with two minor modifications. The first modification is in the simulations used for the one-time generation of lookup tables. Instead of white noise, range-correlated noise obtained by applying the receiver matched filter must be added to the simulated weather signals. The second modification is in the real-time processing steps, where the computation of the noise power after the transformation was revised to accommodate the range correlation of the noise.

Simulations involving receiver filters with varying bandwidths revealed that the use of a conventional matched filter leads to the lowest variance of radar-variable estimates when compared to filters with wider bandwidths. The improvement occurs at low-to-medium SNRs because a receiver filter with a narrower bandwidth

results in less noise power at its output. This seems to dominate the noise enhancement effect from the additional range correlation of the weather signal. Simulation results were validated using data collected with the research KOUN radar. The analysis of data collected with both a receiver matched filter and a receiver filter with L -times-as-wide bandwidth confirmed that the standard deviation of reflectivity estimates was lower for the data collected with the receiver matched filter when the SNR was low. As expected, at higher SNR values (above 20 dB), the performance of both receiver filters was very similar.

As a side benefit, relative to using wider receiver bandwidths, the use of a receiver matched filter results in increased data coverage as a result of an improvement in the SNR. That is, with LUT adaptive pseudowhitening, a DMF is used to produce estimates of the SNR that are used for thresholding purposes (i.e., to censor data corresponding to nonsignificant returns). Because the receiver matched filter provides an additional small improvement in post-DMF SNR relative to the widest receiver bandwidth (~ 0.6 dB for the pulse and the filters in this study), it leads to fewer data being thresholded (censored).

In conclusion, range-oversampling processing performs best when using a receiver matched filter. Because most operational weather radars already employ a receiver filter that is matched to the transmitter pulse, implementing range-oversampling processing does not require modifications to the receiver filter. This should facilitate implementations of range-oversampling processing on operational weather radars like the WSR-88D.

Acknowledgments. The authors thank Feng Nai, David Schwartzman, and three anonymous reviewers for providing comments to improve the paper. We also

thank Charles Kuster and Danny Wasielewski for helping with data collection using the KOUN radar. Funding was provided by the NOAA/Office of Oceanic and Atmospheric Research under NOAA–University of Oklahoma Cooperative Agreement NA11OAR4320072 from the U.S. Department of Commerce.

REFERENCES

- Curtis, C., 2018: Weather radar time series simulation: Improving accuracy and performance. *J. Atmos. Oceanic Technol.*, **35**, 2169–2187, <https://doi.org/10.1175/JTECH-D-17-0215.1>.
- , and S. Torres, 2011: Adaptive range oversampling to achieve faster scanning on the National Weather Radar Testbed phased array radar. *J. Atmos. Oceanic Technol.*, **28**, 1581–1597, <https://doi.org/10.1175/JTECH-D-10-05042.1>.
- , and —, 2013: Real-time measurement of the range correlation for range oversampling processing. *J. Atmos. Oceanic Technol.*, **30**, 2885–2895, <https://doi.org/10.1175/JTECH-D-13-00090.1>.
- , and —, 2014: Adaptive range oversampling to improve estimates of polarimetric variables on weather radars. *J. Atmos. Oceanic Technol.*, **31**, 1853–1866, <https://doi.org/10.1175/JTECH-D-13-00216.1>.
- , and —, 2017: Adaptive range oversampling processing for nontraditional radar-variable estimators. *J. Atmos. Oceanic Technol.*, **34**, 1607–1623, <https://doi.org/10.1175/JTECH-D-16-0051.1>.
- Oppenheim, A., and W. Schaffer, 1989: *Discrete-Time Signal Processing*. Prentice Hall, 879 pp.
- Torres, S., and D. Zrnić, 2003: Whitening in range to improve weather radar spectral moment estimates. Part I: Formulation and simulation. *J. Atmos. Oceanic Technol.*, **20**, 1433–1448, [https://doi.org/10.1175/1520-0426\(2003\)020<1433:WIRTIW>2.0.CO;2](https://doi.org/10.1175/1520-0426(2003)020<1433:WIRTIW>2.0.CO;2).
- , and C. Curtis, 2012: The impact of signal processing on the range-weighting function for weather radars. *J. Atmos. Oceanic Technol.*, **29**, 796–806, <https://doi.org/10.1175/JTECH-D-11-00135.1>.
- Zrnić, D., and R. Doviak, 1978: Matched filter criteria and range weighting for weather radar. *IEEE Trans. Aerosp. Electron. Syst.*, **AES-14**, 925–930, <https://doi.org/10.1109/TAES.1978.308559>.