Adaptive Range Oversampling Processing for Nontraditional Radar-Variable Estimators

CHRISTOPHER D. CURTIS AND SEBASTIÁN M. TORRES

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

(Manuscript received 19 February 2016, in final form 31 March 2017)

ABSTRACT

Adaptive range oversampling processing can be used either to reduce the variance of radar-variable estimates without increasing scan times or to reduce scan times without increasing the variance of estimates. For example, an implementation of adaptive pseudowhitening on the National Weather Radar Testbed Phased-Array Radar (NWRT PAR) led to a twofold reduction in scan times. Conversely, a proposed implementation of adaptive pseudowhitening the U.S. Next Generation Weather Radar (NEXRAD) network would reduce the variance of dual-polarization estimates while keeping current scan times. However, the original version of adaptive pseudowhitening is not compatible with radar-variable estimators for which an explicit variance expression is not readily available. One such nontraditional estimator is the hybrid spectrum-width estimator, which is currently used in the NEXRAD network. In this paper, an extension of adaptive pseudowhitening is proposed that utilizes lookup tables (rather than analytical solutions based on explicit variance expressions) to obtain range oversampling transformations. After describing this lookup table (LUT) adaptive pseudowhitening technique, its performance is compared to that of the original version of adaptive pseudowhitening using traditional radar-variable estimators. LUT adaptive pseudowhitening is then applied to the hybrid spectrum-width estimator, and simulation results are confirmed with a qualitative analysis of radar data.

1. Introduction

On weather radars, the variance of radar-variable estimates and the scan (or update) time are typically coupled—that is, improving one usually comes at the expense of degrading the other. For example, one of the simplest ways to reduce the variance of estimates is by increasing the number of samples that are used in the estimation process; this leads to longer dwell times (for the same pulse repetition time) and overall longer scan times (for the same volumetric coverage). Range oversampling processing was proposed by Torres and Zrnić (2003a,b) as a way to reduce the variance of radar-variable estimates without increasing scan times or, conversely, to reduce scan times without degrading the variance of estimates. Range oversampling consists of sampling the received signals at rates faster than the inverse of the pulse width so that more samples (in range time) are available for processing with the same number of transmitted pulses per dwell (or the same dwell times). Because the transmitted pulse width remains the same as with conventional sampling, so does the size of the radar resolution volumes. Thus, range oversampling creates overlapped radar resolution volumes, which lead to a significant correlation of signals along range time. If auto- and cross-correlation estimates were incoherently averaged in range, this would not result in maximum variance reduction. However, under the assumption of a uniform distribution of hydrometeors, the range correlation of oversampled signals can be measured or computed a priori (Curtis and Torres 2013), and range-oversampled signals can be decorrelated (or whitened) using a linear transformation so that incoherent averaging results in maximum variance reduction for all radar-variable estimates.

Although a whitening transformation on range-oversampled signals achieves complete decorrelation and hence maximum variance reduction, it also leads to an enhancement of the noise that limits its application at low-to-moderate signal-to-noise ratios (SNR). To overcome this limitation, pseudowhitening was introduced (Torres et al. 2004) as a natural way to trade variance reduction for less noise enhancement. In practice, the best compromise is obtained when the pseudowhitening transformation is tailored to the signal characteristics; this is referred to as adaptive range

Corresponding author: Christopher Curtis, chris.curtis@noaa.gov

DOI: 10.1175/JTECH-D-16-0051.1

oversampling processing or simply adaptive pseudowhitening. Minimization of explicit variance expressions for the spectral moments and polarimetric variables leads to closed-form solutions for pseudowhitening transformations that depend on signal characteristics such as the SNR, the normalized spectrum width, the differential reflectivity, and the cross-correlation coefficient (Curtis and Torres 2011, 2014). Currently, adaptive pseudowhitening relies on these closed-form solutions and on estimates of signal characteristics to obtain the proper pseudowhitening transformation for each radar variable at each radar gate.

In recent years, nontraditional radar-variable estimators have gained popularity, since they effectively address some of the deficiencies of their traditional counterparts and exhibit overall improved performance. Examples of nontraditional estimators are the hybrid spectrum-width estimator (Meymaris et al. 2009), the multilag polarimetric variable estimators (Lei et al. 2012), and the improved correlation-coefficient estimator (Ivić 2014). Unfortunately, explicit variance expressions for these estimators either do not exist or are difficult to derive. Thus, the original adaptive pseudowhitening algorithm cannot be applied to these nontraditional estimators because without closed-form solutions, there is no mechanism by which to obtain suitable pseudowhitening transformations.

In this paper, we extend adaptive pseudowhitening so that transformations can be obtained without the need for explicit variance expressions. The proposed extension makes adaptive pseudowhitening compatible with any radar-variable estimator and, more importantly, enables its implementation on systems that use nontraditional estimators for which explicit variance expressions are not available. For example, the hybrid spectrum-width estimator has recently replaced the traditional spectrum-width estimator (the one based on the ratio of autocorrelation estimates at lags 0 and 1) on the Weather Surveillance Radar-1988 Doppler (WSR-88D) signal processor, and other nontraditional estimators are currently being considered for future upgrades. Hence, an operational implementation of adaptive pseudowhitening on the WSR-88D signal processor is contingent upon its compatibility with nontraditional estimators (at the time of writing this paper, the hybrid spectrum-width estimator is the only nontraditional estimator used operationally on the WSR-88D).

The paper is organized as follows. Section 2 describes the proposed extension of adaptive pseudowhitening. In section 3, we validate the proposed extension by comparing its performance to that of the original algorithm for traditional estimators and show that it is robust to hardware drifts or radar-to-radar variations. In section 4, we use simulations and data collected with the KOUN radar (Norman, Oklahoma) to illustrate the performance of the proposed extension. We show that the proposed extension of adaptive pseudowhitening can be used to improve spectrum-width estimates that are obtained with the hybrid spectrum-width estimator, which does not have an explicit variance expression. In section 5, we end the paper with a summary and final remarks.

2. Lookup table adaptive pseudowhitening

In this section, we give a short introduction to adaptive pseudowhitening and then extend it by replacing the closed-form variance-minimization solutions used in the original algorithm with lookup tables. After presenting a suitable approach, the procedure for producing the lookup tables is described, the lookup table adaptive pseudowhitening algorithm is presented, and some practical issues are considered.

The basic idea behind whitening or pseudowhitening is to multiply the range-oversampled time series data for a particular range gate by a linear transformation to produce transformed time series data that can be used to obtain estimates with lower variance. This is captured with the following expression:

$$\mathbf{X} = \mathbf{W}\mathbf{V}, \tag{1}$$

where $\mathbf{V}$ is an L-by-M matrix of range-oversampled complex time series data, $\mathbf{W}$ is a complex-valued L-by-L linear transformation, and $\mathbf{X}$ is an L-by-M matrix of transformed time series data. The constant L is the range oversampling factor, and M is the number of samples in the dwell. For each correlation needed to estimate a particular radar variable [e.g., only the lag-0 autocorrelation is needed for the reflectivity estimator, but both lag-0 (R0) and lag-1 (R1) autocorrelations are needed for the traditional R0–R1 spectrum-width estimator], L correlation estimates are calculated, corresponding to the rows of the transformed data matrix $\mathbf{X}$. These L correlation estimates are then averaged to form one estimate of each correlation; these estimates are finally used to compute the single radar-variable estimate for the range gate.

The goal of adaptive pseudowhitening is to find the best possible $\mathbf{W}$ that minimizes the variance of the radar-variable estimator. In the case of original adaptive pseudowhitening, we start with an explicit expression for the variance of a particular radar variable θ. Generically, this is given as

$$\mathrm{Var}(\hat{\theta}) = D\{A\mathrm{tr}[(\mathbf{W}^*\mathbf{C}_V\mathbf{W}^\mathrm{T})^2] + B\mathrm{tr}[(\mathbf{W}^*\mathbf{C}_V\mathbf{W}^\mathrm{T})(\mathbf{W}^*\mathbf{W}^\mathrm{T})] + C\mathrm{tr}[(\mathbf{W}^*\mathbf{W}^\mathrm{T})^2]\}, \tag{2}$$

where tr(·) is the trace operator, superscript T denotes the matrix transpose, the superscript asterisk (*) denotes complex conjugation, and $\mathbf{C}_V$ is the normalized range correlation matrix of the time series data before the linear transformation (Curtis and Torres 2014).[1] Assuming a uniform distribution of hydrometeors in the radar volume, the normalized range correlation matrix can be precomputed, since it depends only on the modified pulse (recall that the modified pulse is determined by the transmitted pulse envelope and the receiver impulse response). The A–D constants are radar-variable specific; for example, the constants used for the signal power estimator are $A = (1/2)\sigma_{vn}\pi^{1/2}$, $B = 2/\mathrm{SNR}_0$, $C = 1/\mathrm{SNR}_0^2$, and $D = S^2/ML^2$, where σ_{vn} is the spectrum width normalized by the Nyquist cointerval, SNR_0 is the signal-to-noise ratio at the output of the digital receiver (linear units), and S is the signal power (linear units). The D constant is useful to accurately calculate the variance but is not needed for the minimization, so the constants of interest are A–C. For the signal power estimator, these constants depend on two values: σ_{vn} and SNR_0. As shown in Curtis and Torres (2014), the variances of the signal power estimator S, the radial-velocity estimator v, and the spectrum-width estimator σ_v all depend on those two variables. The variances of the dual-polarization estimators depend on those same two variables and also on the differential reflectivity Z_{DR} and the correlation coefficient ρ_{HV}.

At each range gate, adaptive pseudowhitening finds a nearly optimal linear transformation that minimizes (2). A truly optimal linear transformation cannot be obtained because theoretical variance expressions are derived using approximations and because they depend on values that are unknown (σ_{vn}, SNR_0, Z_{DR}, and ρ_{HV}). In practice, these unknown values are estimated as part of the process of finding the linear transformation. Experience has shown that adaptive pseudowhitening performs almost identically whether estimates or true values are used (Curtis and Torres 2011, 2014). Based on (2) and on estimates of σ_{vn} and SNR_0 (and Z_{DR} and ρ_{HV} for the dual-polarization variables), a linear transformation is calculated that attempts to minimize the variance to provide a nearly optimal estimate of the radar variable. However, if it is difficult or impossible to find an explicit variance expression for a particular radar variable, then an alternate approach must be developed. In this case, the unknown values still need to be estimated but a replacement for the variance expression is required. We propose using a lookup table indexed by estimated values of σ_{vn} and SNR_0 (and Z_{DR} and ρ_{HV} for the dual-polarization variables) to provide the appropriate pseudowhitening transformation. This approach results in the proposed extension, herein referred to as lookup table (LUT) adaptive pseudowhitening.

Although the variance expression is minimized by finding a linear transformation $\mathbf{W}$, our efficient real-time implementation of adaptive pseudowhitening does not explicitly compute $\mathbf{W}$. Following a similar approach to the one in Curtis and Torres (2014), the lookup table formulation of adaptive pseudowhitening divides the linear transformation into two parts:

$$\mathbf{W} = \mathbf{D}^{1/2}\mathbf{U}^{*\mathrm{T}}. \tag{3}$$

Rather than applying $\mathbf{W}$ directly, the efficient implementation first decorrelates the time series data using $\mathbf{U}^{*\mathrm{T}}$ and then employs a weighted average $\mathbf{d}$ to combine L correlation estimates. Matrix $\mathbf{U}$ comes from the eigendecomposition of the normalized range correlation matrix $\mathbf{C}_V = \mathbf{U}^*\mathbf{\Lambda}\mathbf{U}^{\mathrm{T}}$ and only changes when the range correlation matrix changes. The weighted average is derived directly from the real-valued diagonal matrix $\mathbf{D}^{1/2}$; this is the adaptive part of the transformation that is computed separately for each radar variable and at each range gate. The main advantage of the efficient implementation is that the range-oversampled data can be clutter filtered before the second part of the transformation is applied; that is, the clutter filter is applied only once. Otherwise, the clutter filter would need to be applied separately to transformed data corresponding to each radar-variable-specific transformation.

Mathematically, the first step of the efficient implementation for each range gate is to premultiply the L-by-M time series matrix $\mathbf{V}$ by $\mathbf{U}^{*\mathrm{T}}$. The partially transformed (and decorrelated) matrix $\tilde{\mathbf{X}} = \mathbf{U}^{*\mathrm{T}}\mathbf{V}$ is used to compute sets of L correlations (one set for each correlation is needed to compute a radar variable). Figure 1 graphically depicts the steps of the efficient implementation of adaptive pseudowhitening processing for the computation of the spectral moments; steps 1 and 2 show the partial transformation being applied to $\mathbf{V}$ and the computation of the sets of L correlations. The third step is the application of a fixed matched-filter weight vector $\mathbf{d}_{MF}$ to obtain the correlations needed to compute the initial estimates of σ_{vn} and SNR_0 (and Z_{DR} and ρ_{HV} for the dual-polarization variables). These initial estimates are used in the next step (step 4 in Fig. 1) to calculate the radar-variable-specific weight vectors that capture the adaptive part of adaptive pseudowhitening and are the focus of both the original and LUT versions.

[1] This formulation assumes that the polarimetric channels are matched so that the normalized range correlation matrices for the horizontal and vertical channels are the same. The reader is referred to Torres (2009) for the case with mismatched channels.

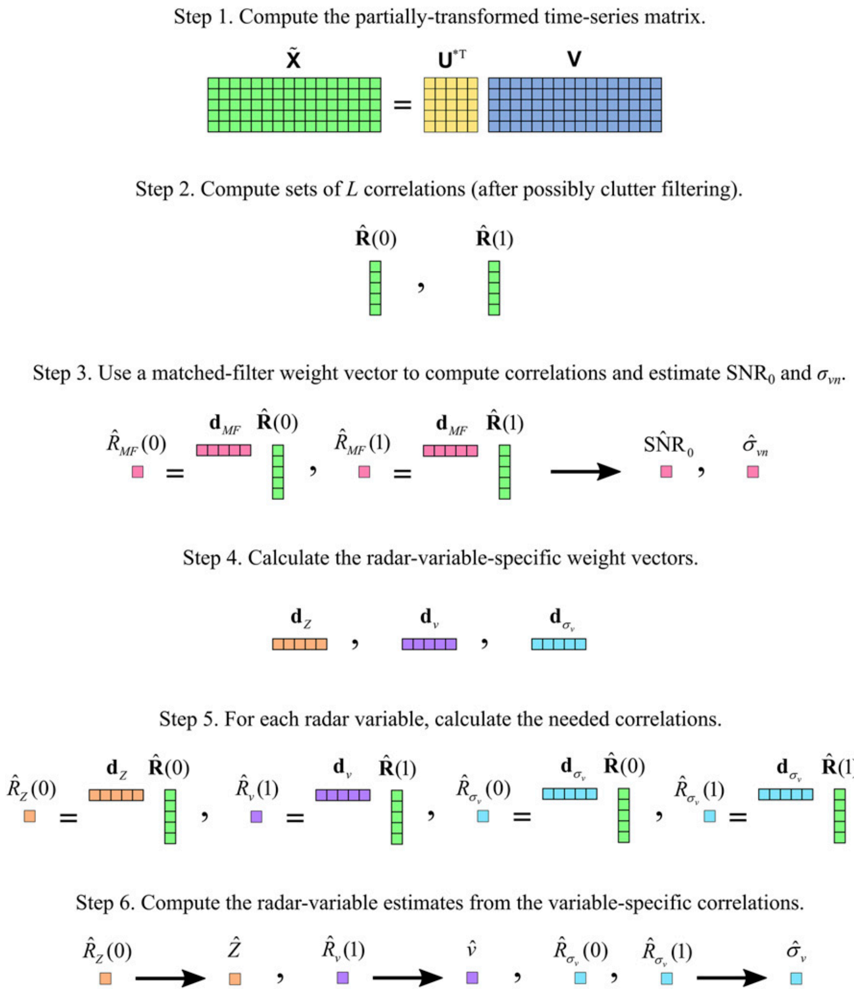

FIG. 1. Graphical depiction of the efficient implementation of adaptive pseudowhitening processing. Only the steps to produce the single-polarization spectral moments are included for simplicity.

The weight vector for each radar variable $\mathbf{d}_\theta$ is closely related to the second part of $\mathbf{W}$, $\mathbf{D}^{1/2}$. It combines the sets of L correlation estimates to obtain the correlations needed to calculate the final radar-variable estimates (step 5 in Fig. 1). The elements of $\mathbf{d}_\theta = [d_0, d_1, \ldots, d_{L-1}]$ are properly scaled values of the elements along the diagonal matrix $\mathbf{D}$. We use the squared matrix $\mathbf{D}$ instead of $\mathbf{D}^{1/2}$ because the correlations are computed from products of time series samples; the details of this derivation can be found in Curtis and Torres (2011).

For the original version of adaptive pseudowhitening, the d_l values are given by

$$d_l = g\frac{\lambda_l}{A\lambda_l^2 + B\lambda_l + C}, \qquad (4)$$

where λ_l are the eigenvalues of $\mathbf{C}_V$ and A–C are from the radar-variable-specific variance expression and depend on σ_{vn} and SNR_0 (and also Z_{DR} and ρ_{HV} for the dual-polarization radar variables), and g is a power-preserving factor. When using the original version of adaptive pseudowhitening, the initial estimates of σ_{vn}, SNR_0, Z_{DR}, and ρ_{HV} are computed and substituted into the formulas for A–C from the variance expression [(2)].

For the LUT version, we need to find a way to compute the d_l values without using an explicit variance expression. One way of doing this would be to calculate three different lookup tables for A–C that are based on the estimates of σ_{vn}, SNR_0, Z_{DR}, and ρ_{HV}. Instead of using a variance expression, the lookup tables can be constructed using an optimization procedure based on Monte Carlo simulations of radar signals with varying characteristics. After trying three variations of this procedure (an L-parameter version to find $\mathbf{d}_\theta$ directly, a three-parameter version to find A–C, and a single-parameter version), we concluded that the single-parameter formula for the d_l values is sufficient, since all three versions performed nearly identically. This single-parameter formula is related to the sharpening filter described in Torres et al. (2004) and is given by

$$d_l = g\frac{\lambda_l}{[p\lambda_l + (1-p)]^2}. \qquad (5)$$

This formula also includes a power-preserving factor and depends on the eigenvalues of the normalized range correlation matrix, but there is only one parameter, p $(0 \le p \le 1)$, instead of the three parameters in (4). This greatly simplifies the optimization procedure, and we will show in section 3 that the one-parameter version of the LUT adaptive pseudowhitening algorithm performs comparably to the original adaptive pseudowhitening algorithm.

The dimensionality of the lookup tables is determined by the number of unknown values that are found in the formulas for A–C. For the spectral moments, lookup tables are two-dimensional and are based on σ_{vn} and SNR_0; for the dual-polarization variables, they are four-dimensional and are based on σ_{vn}, SNR_0, Z_{DR}, and ρ_{HV}. The number of samples M is not an independent variable in the lookup tables because the variance expressions that were originally used do not have a dependence on M for A, B, or C. Although the approximations used to derive the formulas for A–C are less valid for very small values of M, the formulas currently work well for a wide range of signal parameters. If there were a strong dependence on M for a future radar-variable estimator, then this procedure could be expanded to include M as an independent variable in the lookup tables. For this work, we analyzed the lookup tables corresponding to the traditional radar-variable estimators for different values of M and observed an overall weak dependence that vanishes as M increases. That is, the reflectivity LUT converges for $M \ge 4$, the ones for the polarimetric variables for $M \ge 16$, and those for radial velocity and spectrum width for $M \ge 32$. Thus, for all radar variables, the smallest value of M beyond which all LUTs look virtually identical is about the same or smaller than the typical values used operationally. Thus, to simplify the process and still achieve accurate performance for all radar variables, we produced the LUTs using a fixed value of $M = 32$.

These are the basic steps to produce a lookup table for θ:

1) Simulate a large number of time series realizations assuming a Gaussian weather signal model (50 000 were used in this case) and a given set of signal characteristics (i.e., σ_{vn} and SNR_0 for the spectral moments plus Z_{DR} and ρ_{HV} for the dual-polarization variables) using a representative value of M and the desired L (in this work $M = 32$ and $L = 5$). This includes imposing the appropriate range correlation on the data by convolving independent time series realizations with the radar's modified pulse, which can be measured a priori (in this work we used the modified pulse measured on the KOUN radar). White noise is added after applying the range correlation to the data because the noise is assumed to be uncorrelated with the weather signal.
2) Apply $\mathbf{U}^{*T}$ to each realization of $\mathbf{V}$ (L-by-M samples) and calculate the correlations (sets of L) needed for the particular radar-variable estimator.
3) Find the p value that minimizes the mean squared error of radar-variable estimates; this can be done by brute-force search or by using a standard single-parameter optimization algorithm (in this work we used the well-known golden-section search). Regardless of the method used to minimize the mean squared error of estimates, for each value of p being tested, find the corresponding $\mathbf{d}_\theta$ and apply it to each set of L correlations. Estimate the radar variable for each realization and compute the mean squared error of estimates using the results for all realizations.
4) Store the optimum value of p in a lookup table that is indexed by σ_{vn} and SNR_0 (and Z_{DR} and ρ_{HV} for the dual-polarization variables).
5) Repeat steps 1–4 for each set of signal characteristics: SNR_0 and σ_{vn} for the spectral moments and Z_{DR} and ρ_{HV} for the dual-polarization variables.

One lookup table for each radar-variable estimator is stored for later use by LUT adaptive pseudowhitening; a lookup table needs to be recomputed only if a particular radar-variable estimator changes.

Although the steps for implementing LUT adaptive pseudowhitening are similar to the steps for the original version, they are included here to show the particular details of the LUT version and also some slight changes from previous implementations (Curtis and Torres 2011, 2014). The steps directly follow the ones depicted in Fig. 1 but also include the dual-polarization radar variables.

1) Compute the partially transformed matrices of time series data, $\tilde{\mathbf{X}}_{H,V} = \mathbf{U}^{*T}\mathbf{V}_{H,V}$, using the $\mathbf{U}$ computed from the eigendecomposition of $\mathbf{C}_V = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$. Ground clutter filtering can be applied to these two partially transformed data matrices, where subscript H,V indicates a quantity that corresponds either to the horizontal or vertical polarization channel.
2) Compute sets of L range-oversampled auto- and cross correlations from the $\tilde{\mathbf{X}}_{H,V}$ needed to compute SNR_0, σ_{vn}, Z_{DR}, and ρ_{HV}, and any additional correlations needed for other radar-variable estimators.

3) Compute the matched-filtered weight vector $\mathbf{d}_0$ (a particular instance of $\mathbf{d}_{\mathrm{MF}}$) from (5) using $p = 0$ and apply it to each set of L correlations. Using the resulting correlations and the adjusted noise powers, obtain initial estimates of SNR_0, σ_{vn}, Z_{DR}, and ρ_{HV}. Horizontal- and vertical-channel noise powers are adjusted by the corresponding noise enhancement factor, which is given by the sum of the elements of $\mathbf{d}_0$.
4) Using the initial estimates of SNR_0, σ_{vn}, Z_{DR}, and ρ_{HV}, retrieve radar-variable-specific values of p from the lookup tables and calculate the corresponding weight vectors: $\mathbf{d}_Z$, $\mathbf{d}_v$, $\mathbf{d}_{\sigma_v}$, $\mathbf{d}_{Z_{\mathrm{DR}}}$, $\mathbf{d}_{\Phi_{\mathrm{DP}}}$, and $\mathbf{d}_{\rho_{\mathrm{HV}}}$.
5) For each θ, apply the appropriate $\mathbf{d}_\theta$ to each set of L correlations needed to compute the final radar-variable estimate.
6) Compute final radar-variable estimates from the variable-specific correlations obtained in the previous step. Use the appropriate variable-specific-adjusted noise powers when necessary, where $N_{\theta(H,V)} = \mathrm{NEF}_\theta \times N_{H,V}$ and $\mathrm{NEF}_\theta = \sum_{l=0}^{L-1} d_\theta(l)$.

It should be noted that step 3 differs from the corresponding step in Curtis and Torres (2014). We now use pseudowhitening estimators with $p = 0$ to obtain the initial estimates of σ_{vn}, SNR_0, Z_{DR}, and ρ_{HV} to access the LUTs, since they outperform their digital-matched-filter counterparts in the useful range of SNR values. This also leads to a simplification of the algorithm because step 3 better matches the processing in steps 4–6; the only difference is that p is 0 in step 3, whereas p is retrieved from the lookup tables in step 4. However, there is a trade-off: pseudowhitening estimators with $p = 0$ exhibit slightly degraded range resolution compared to their digital-matched-filter counterparts, which can be quantified using the range weighting function formulation from Torres and Curtis (2012). This is typically not a problem, since it applies only to the initial estimates used to access the LUTs and not to the final radar-variable estimates.

Although extending adaptive pseudowhitening to use lookup tables is straightforward, there are some practical issues that need to be addressed. The first is that the lookup tables are computed using a particular modified pulse. If the lookup tables are sensitive to changes in the range correlation, they would require periodic recalculation, which would be a major obstacle to implementing LUT adaptive pseudowhitening. This could also be an issue for a network of radars, since there is variation in the modified pulse across the network. Fortunately, we will show in the next section that the lookup tables are not especially sensitive to changes in the modified pulse. Thus, a single set of lookup tables (one per radar variable) computed using a representative pulse should work well over time and across a network of radars.

The second practical issue is the number and range of values of σ_{vn} and SNR_0 (and Z_{DR} and ρ_{HV} for the dual-polarization variables) needed to produce the lookup tables. Based on a sensitivity analysis with varying resolutions for each independent variable, we identified practical sets of values. For SNR_0, we used [−5, −1, 3, 7, 11, 15, 19, 23, 27, 31, 35] (dB), and we chose a nonuniform grid for the normalized spectrum width σ_{vn}: [0.01, 0.02, 0.03, 0.04, 0.05, 0.07, 0.09, 0.13, 0.17, 0.25]. When accessing the lookup tables, linear interpolation is used to extract the appropriate value of p; however, if a value outside of the lookup table is needed, the extrapolated p is just the nearest neighbor. For dual-polarization variables, the values for Z_{DR} (dB) are [−6, −5, −4, −3, −2, −1, 0, 1, 2, 3, 4, 5, 6], and a nonuniform grid is used for ρ_{HV} [0.2, 0.4, 0.6, 0.7, 0.8, 0.9, 0.94, 0.98, 0.99, 1]. These sets of values worked well for the estimators we tested, but they could be easily adjusted if necessary.

In the next section, LUT adaptive pseudowhitening is validated using (traditional) estimators that have explicit variance expressions. The effects of changes in the normalized range correlation matrix are also explored.

3. Validation of LUT adaptive pseudowhitening

The most straightforward way to validate LUT adaptive pseudowhitening is to compare its performance to that of the original version of adaptive pseudowhitening. Since this requires an estimator with an explicit variance expression, we chose to compare the performance of these two techniques for both traditional spectral-moment estimators of reflectivity, radial velocity, and spectrum width and traditional dual-polarization estimators of differential reflectivity, differential phase, and correlation coefficient. These are all estimators that have an explicit variance expression, and the values for A–C that are needed to compute $\mathbf{d}_\theta$ for the original version of adaptive pseudowhitening from Eq. (4) can be found in Curtis and Torres (2014). For LUT adaptive pseudowhitening, lookup tables were produced using the procedure outlined in section 2. From the lookup tables, a value of p is retrieved and used to compute the LUT adaptive pseudowhitening $\mathbf{d}_\theta$ from (5).

To compare the performance of the original and LUT adaptive pseudowhitening, 50 000 realizations of time series data were simulated and processed with both techniques using a range oversampling factor of $L = 5$ (the recommended value for the WSR-88D) and the modified pulse of the KOUN radar (Fig. 4). The parameters were set to closely match the lowest elevation cut of volume coverage pattern (VCP) 12 on the NEXRAD network. For the computation of reflectivity and the three dual-polarization radar variables, the maximum unambiguous velocity is $\upsilon_a = 8.33\,\mathrm{m\,s^{-1}}$, and there are $M = 15$ samples per dwell. For radial velocity and spectrum width, $\upsilon_a = 25\,\mathrm{m\,s^{-1}}$ and $M = 40$. As mentioned previously, the LUTs were produced using a fixed value of $M = 32$. The weather signal characteristics were chosen to be representative of a typical rain case: $\sigma_\upsilon = 2\,\mathrm{m\,s^{-1}}$, $Z_{\mathrm{DR}} = 0.25\,\mathrm{dB}$, and $\rho_{\mathrm{HV}} = 0.99$. Since they have no bearing on the statistical performance, radial velocity and differential phase were arbitrarily chosen to be zero; the SNR was systematically varied from 0 to 30 dB in steps of 0.5 dB. The resulting radar-variable estimates corresponding to each technique were then used to compute biases and standard deviations.

The results for the spectral moments are shown in Fig. 2; curves for whitening-transformation-based (WTB) and digital-matched-filter-based (MFB) estimates are also included for comparison. WTB estimates achieve the maximum variance reduction at high SNRs and are obtained with a fixed (nonadaptive) transformation with parameter $p = 1$. In contrast, MFB estimates use a digital matched filter that maximizes the SNR, also maximizing performance at low SNRs. Details about both can be found in Curtis and Torres (2011). In all cases, LUT adaptive pseudowhitening transformation based (LUT APTB) estimates perform comparably to original APTB estimates. LUT adaptive pseudowhitening actually performs slightly better for the spectrum-width estimator; this is most likely due to the approximations inherent in the derivation of the explicit variance expressions. Because the lookup tables are computed from simulated data, the computations avoid any approximations and depend only on the simulation assumptions. For reflectivity, whereas the biases are comparable for LUT and original adaptive pseudowhitening, both techniques exhibit a small negative bias at low SNR. These SNR-dependent biases can be attributed to the larger variability of p values at low SNRs, which also results in larger variability of power-preserving factors. The p values are more variable at low SNRs because they depend on estimates of the signal power, and these become less precise as the SNR decreases. In addition to this variability, the distribution of the power-preserving factors at a given SNR is not symmetric about the mean. The combination of larger variability and asymmetric distribution creates a small bias in power estimates. At high SNRs, the p values are concentrated in a narrow interval as are the associated power-preserving factors so the lack of symmetry is irrelevant, and the performance approaches the unbiased behavior observed when using a fixed value of p (e.g., whitening with $p = 1$ in the same figure).

Similar to Fig. 2, Fig. 3 shows the results for the polarimetric variables, where it is evident that LUT APTB estimates perform comparably to original APTB estimates. Although we show results for only one set of signal characteristics, several cases with varying characteristics were compared, and LUT adaptive pseudowhitening performed comparably in all of them. For all six of the traditional radar-variable estimators, the lookup tables provide equivalent performance to the explicit variance expressions. This suggests that the LUT adaptive pseudowhitening will work well for estimators without explicit variance expressions.

The LUTs used in the previous simulations were obtained for one particular modified pulse (or range correlation). We want to examine the effect of using LUTs computed from one modified pulse on data with a mismatched modified pulse. To be clear, this is not the same as producing pseudowhitening transformations from a mismatched range correlation matrix. We have previously looked at the importance of accurately measuring the range correlation for adaptive pseudowhitening, and using a mismatched range correlation matrix may lead to biases in reflectivity estimates and overall degraded performance because of a less-than-optimum variance reduction (Torres and Curtis 2013). For this validation of LUT adaptive pseudowhitening, we assume that the range correlation is measured accurately using a method like the one found in Curtis and Torres (2013).

For both original and LUT adaptive pseudowhitening, the measured range correlation matrix is the same and leads to the same unitary transformation and eigenvalues. The difference in the efficient implementation occurs when computing the adaptive $\mathbf{d}_\theta$. As discussed in section 2, the weight vector for original adaptive pseudowhitening is computed using three parameters A–C that are based on the initial estimates of σ_{vn} and SNR_0 (and Z_{DR} and ρ_{HV} for the dual-polarization variables). The A–C values are completely independent of the range correlation (or modified pulse). For LUT adaptive pseudowhitening, a single parameter p is used that relies on the same estimates. Since p serves a similar function to A–C, it is probably not especially dependent on the range correlation, but we would still like to confirm this single-parameter conjecture by examining the effects of a mismatched modified pulse.

To test the effects of hardware drift or hardware variations across a network on the performance of LUT adaptive pseudowhitening (with fixed LUTs), we used a similar approach to the one in Torres and Curtis (2013). Unlike the work in that paper, different modified pulses

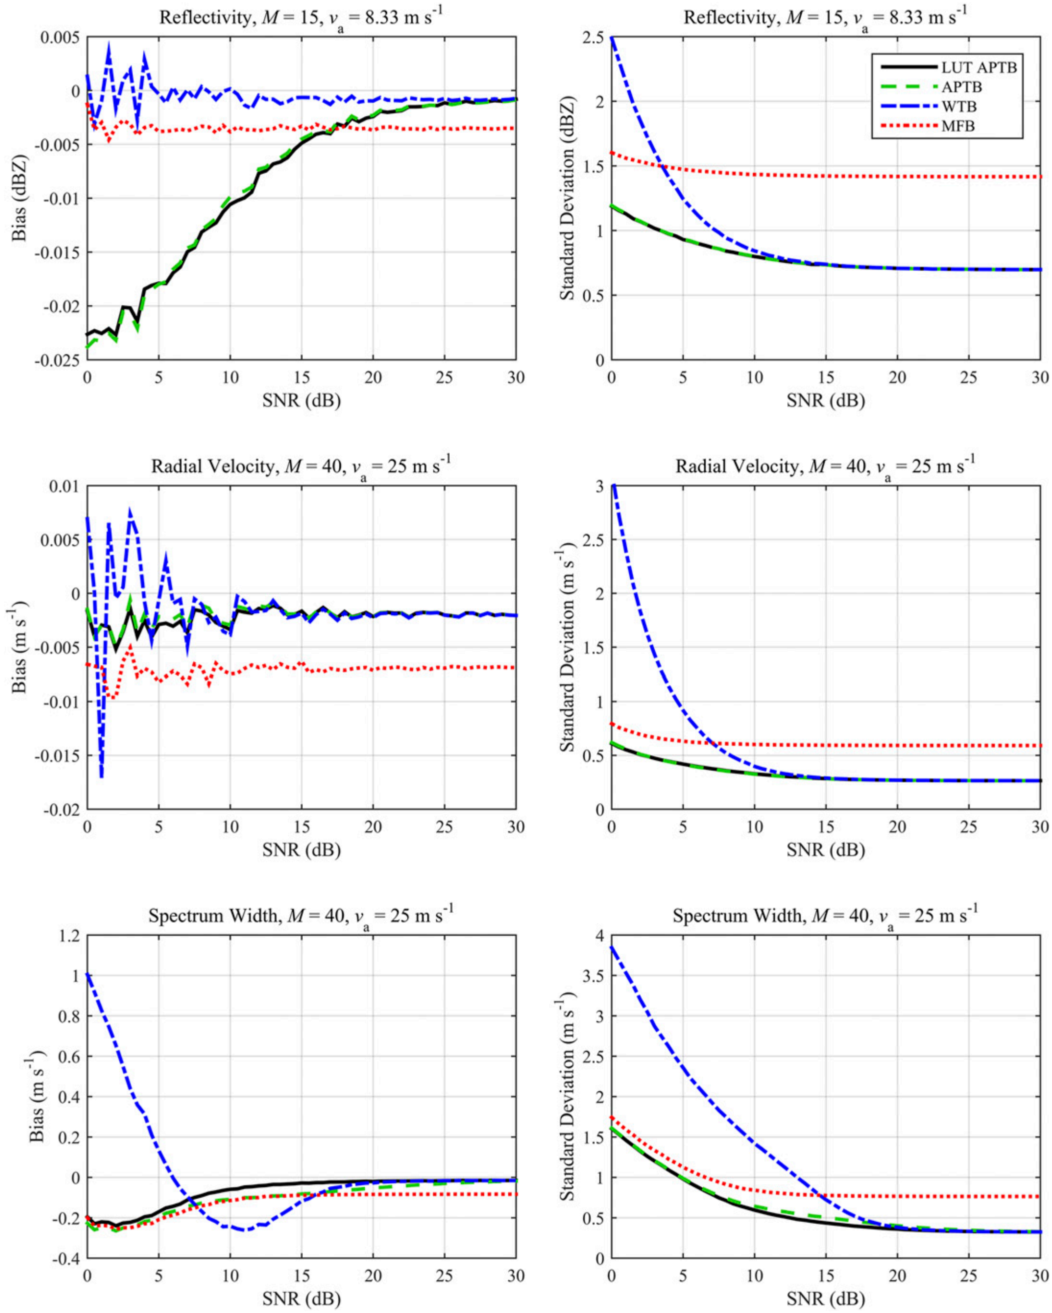

FIG. 2. (left) Bias and (right) standard deviation of LUT APTB (black), original APTB (green dashed), MFB (red dotted), and WTB (blue dashed–dotted line) estimates as a function of SNR for the spectral moments. Simulations use $L = 5$, $\sigma_v = 2$ m s^{-1}, (top) $M = 15$ with a Nyquist velocity $v_a = 8.33$ m s^{-1} ($\sigma_{vn} = 0.12$) for reflectivity, and $M = 40$ with $v_a = 25$ m s^{-1} ($\sigma_{vn} = 0.04$) for (middle) radial velocity and (bottom) spectrum width.

are utilized here to quantify the performance degradation when using mismatched LUTs (but not mismatched transformations); that is, we will explore the robustness of the lookup tables to mismatched range correlations. Figure 4 shows the magnitude and phase of three modified pulses. The original (black) pulse is the one measured on the KOUN radar and also the one used to produce the lookup tables in the previous simulations. The wide (green) pulse and the narrow (red) pulse have different shapes and different phases from the original pulse, leading to different range correlations (also depicted in Fig. 4). Based on our experience, we contend that both the wide and narrow pulses are farther from the original pulse than the typical variation of pulses in

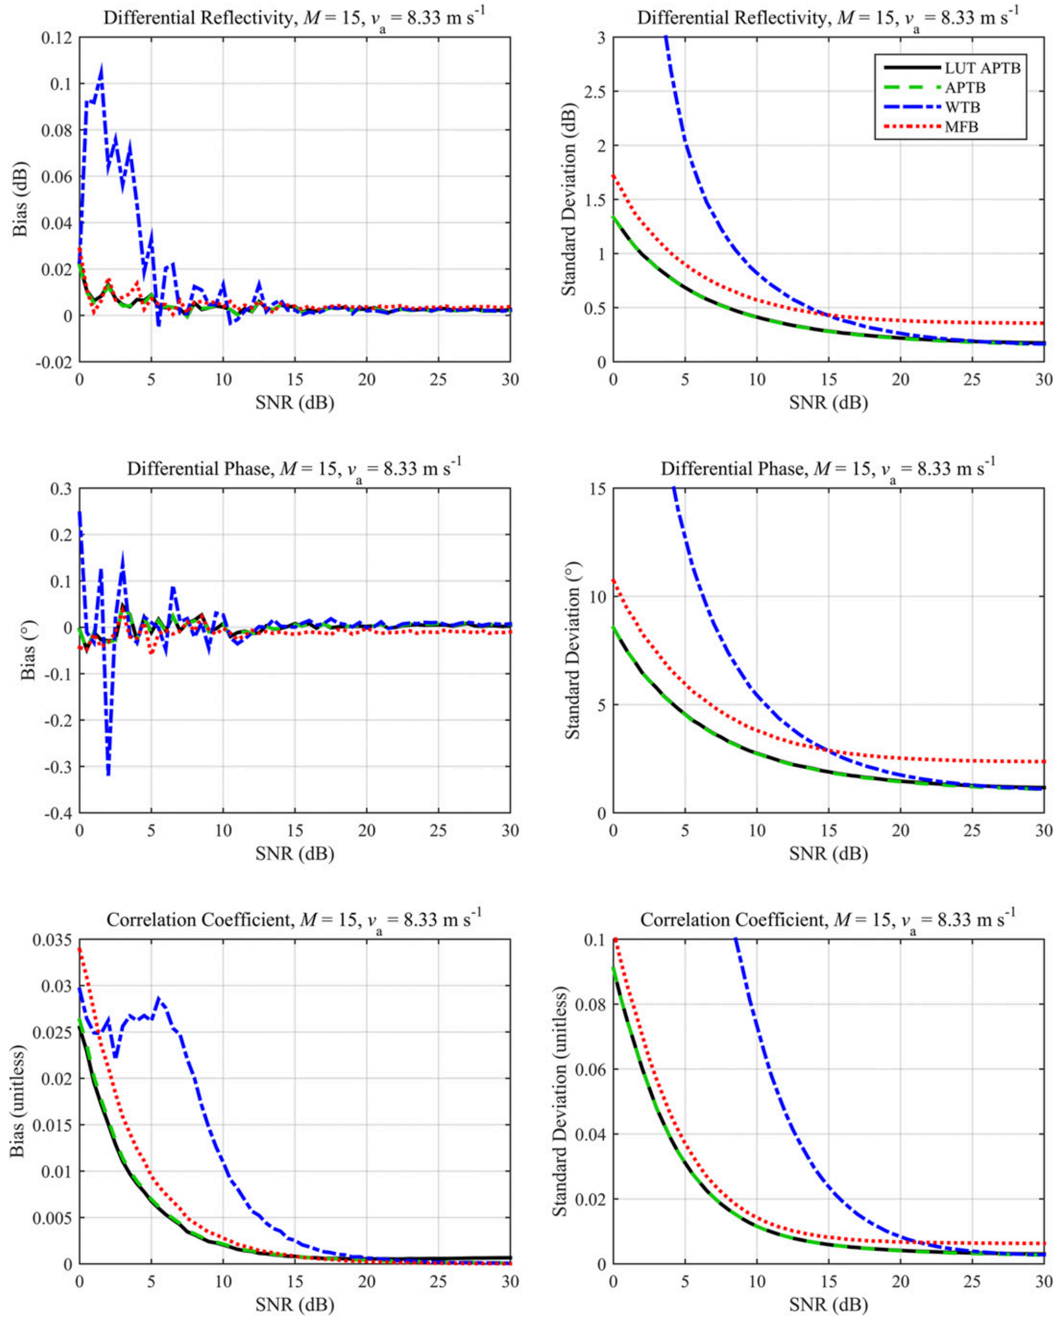

FIG. 3. As in Fig. 2, but for (top to bottom) the polarimetric variables: differential reflectivity and phase, and correlation coefficient. Simulations use $L = 5$, $\sigma_v = 2\,\mathrm{m\,s^{-1}}$, $Z_{\mathrm{DR}} = 0.25\,\mathrm{dB}$, $\rho_{\mathrm{HV}} = 0.99$, and $M = 15$ with $v_a = 8.33\,\mathrm{m\,s^{-1}}$.

time and across a well-maintained radar network. Hence, if the lookup tables computed for the original pulse produce satisfactory results for both the narrow and wide pulses, then a single set of lookup tables based on a nominal, representative pulse should accommodate expected range correlation mismatches from hardware drift and across a radar network. Note that the range correlation is measured independently for each of the three pulses, and a range correlation matrix that is matched to each pulse is used to compute the unitary transformation and eigenvalues for both LUT and original adaptive pseudowhitening. The only part that is being validated is the robustness of the LUTs computed using the original pulse. The results will be deemed

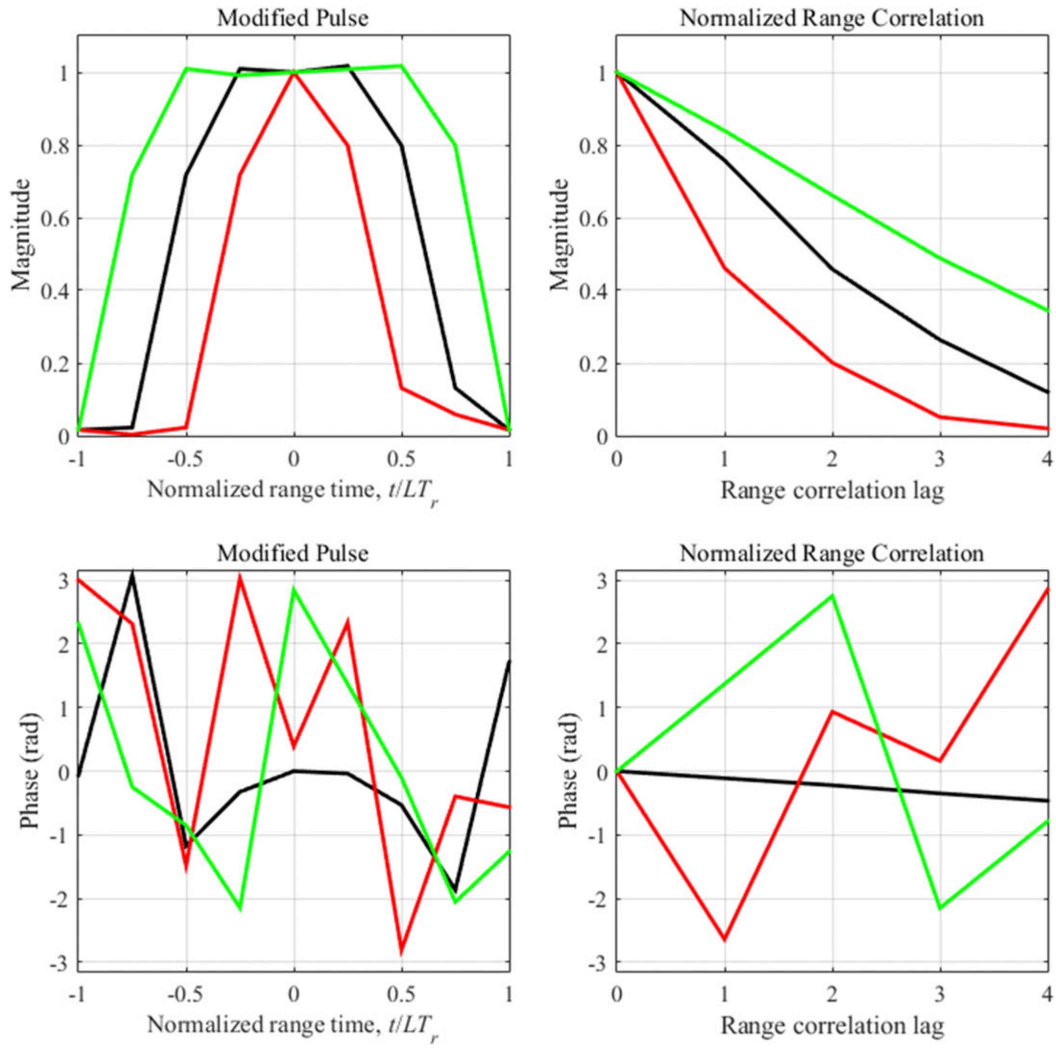

FIG. 4. (top) Magnitude and (bottom) phase of (left) three modified pulses: actual (black), wide (green), and narrow (red), and (right) their corresponding normalized range correlations.

satisfactory if, for all three pulses considered here, the performance of LUT adaptive pseudowhitening is comparable to the performance of original adaptive pseudowhitening.

Figures 5 and 6 show the results of this experiment for the traditional spectral-moment estimators and the traditional polarimetric-variable estimators, respectively. The bias and standard deviation for the six radar variables are displayed for the original pulse and for wide and narrow versions of it. For each case, the performance of both LUT adaptive pseudowhitening (solid) and original adaptive pseudowhitening (dotted) is shown. The simulation parameters are the same as those used for Figs. 2 and 3. The pseudowhitening transformations are calculated based on the (true) range correlation corresponding to each pulse (original, narrow, and wide). In this context, the differences in performance among the three pulses are not the focus of the comparison and can be explained by the differences in range correlations (due to the different modified pulses). The key comparison is between LUT adaptive pseudowhitening and original adaptive pseudowhitening for each pulse. A visual comparison of the solid and dotted curves of the same color (i.e., the same modified pulse) confirms that LUT adaptive pseudowhitening performs comparably to the original version for each of the three modified pulses. This indicates that the LUTs are not especially sensitive to significantly mismatched range correlation matrices, should not require periodic recalculation, and should perform well across a well-maintained radar network. However, the range correlation matrix should still be accurately measured in real time to prevent reflectivity biases and overall degraded variance reduction performance.

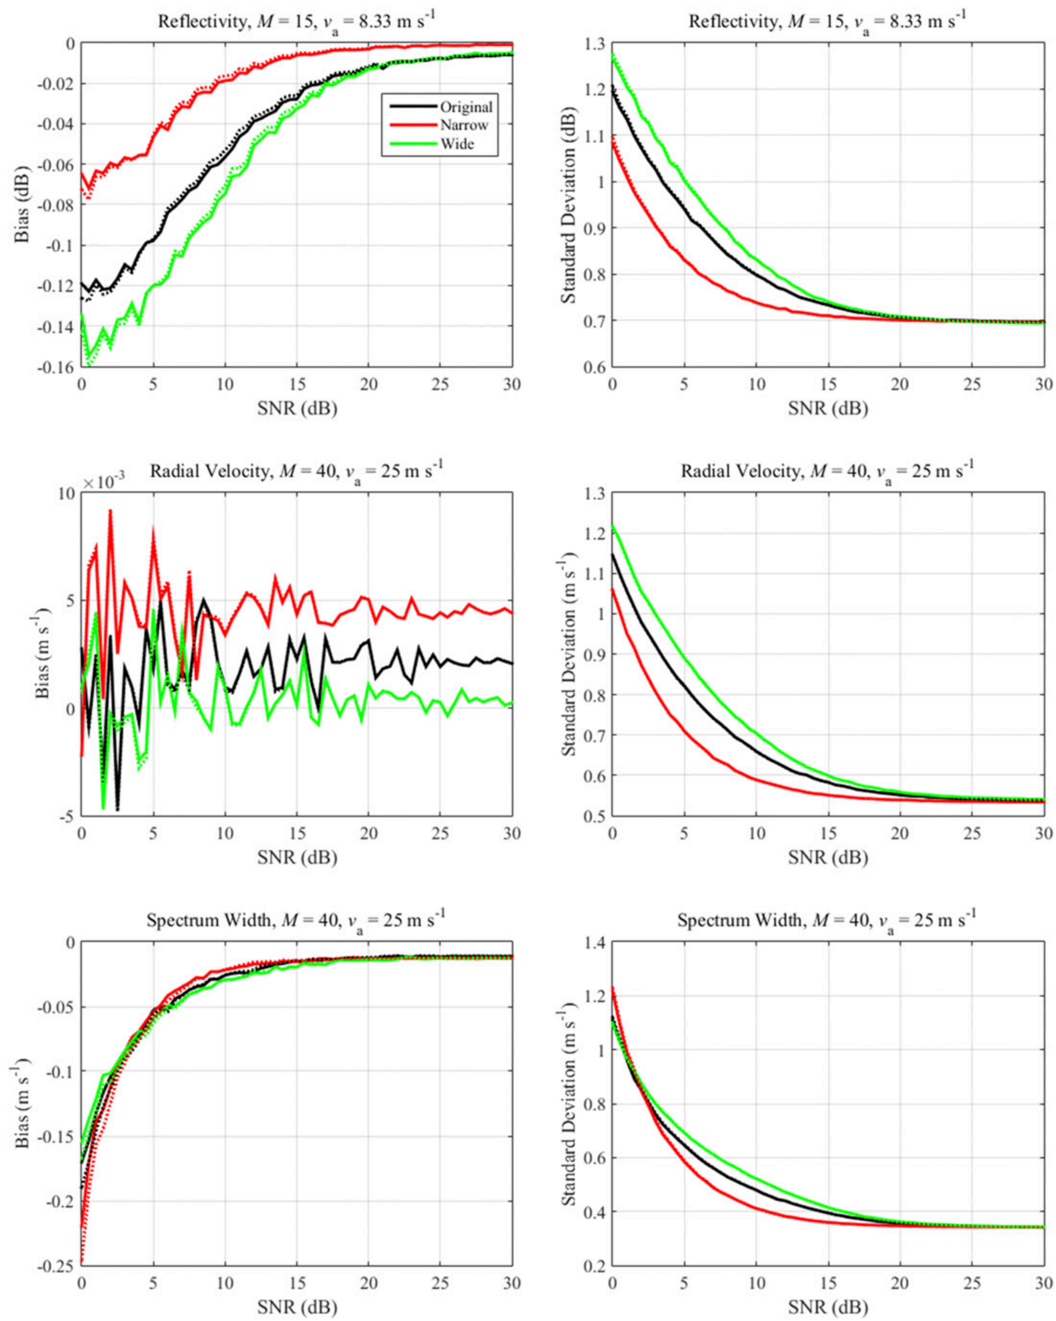

FIG. 5. (left) Bias and (right) standard deviation vs SNR for the spectral-moment estimators using the actual (black), the wide (green), and the narrow pulse (red) in Fig. 4 for the same simulation parameters used in Fig. 2. Shown is the performance of LUT APTB estimates (solid lines) and original APTB estimates (dotted lines).

In the next section, LUT adaptive pseudowhitening is applied to the hybrid spectrum-width estimator, which does not currently have an explicit variance expression. The performance cannot be quantified in the same way as the conventional estimators but meaningful comparisons can still be made.

4. Application of LUT adaptive pseudowhitening to the hybrid spectrum-width estimator

Compared to the traditional spectrum-width (TSW) estimator that is based on the ratio of the autocorrelation at lags 0 and 1 (Doviak and Zrnić 1993), the hybrid

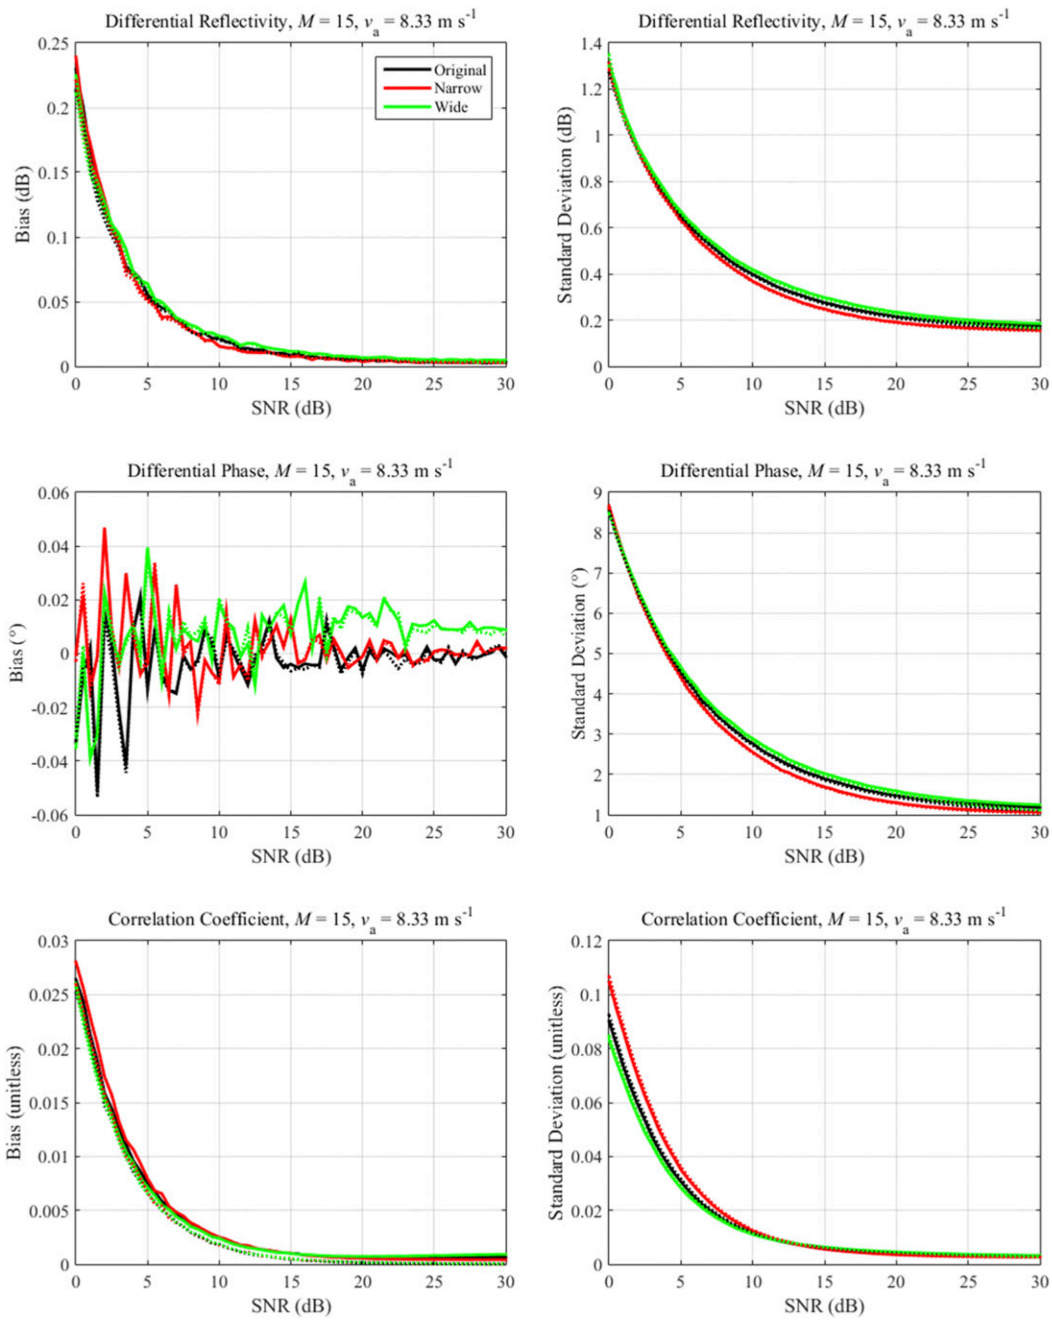

FIG. 6. As in Fig. 5, but for the polarimetric variables and the same simulation parameters used in Fig. 3.

spectrum-width (HSW) estimator introduced by Meymaris et al. (2009) is less biased and has a lower standard deviation for narrow spectrum widths. For this reason, it has recently been implemented on the WSR-88D signal processor. The HSW estimator is a nontraditional estimator that uses one of three possible spectrum-width estimators depending on whether the expected spectrum width is in a small, medium, or large category. The three possible estimators are based on the ratio of the autocorrelation at lags 0 and 1, 1 and 2, and 1 and 3, respectively. Whereas each of these possible spectrum-width estimators has an explicit variance expression (e.g., Zrnić 1979), the variance of the resulting hybrid estimator depends on how the individual estimators are selected, and a closed-form expression is not easily obtained. Thus, the HSW estimator is a perfect candidate to illustrate the performance of LUT adaptive pseudowhitening.

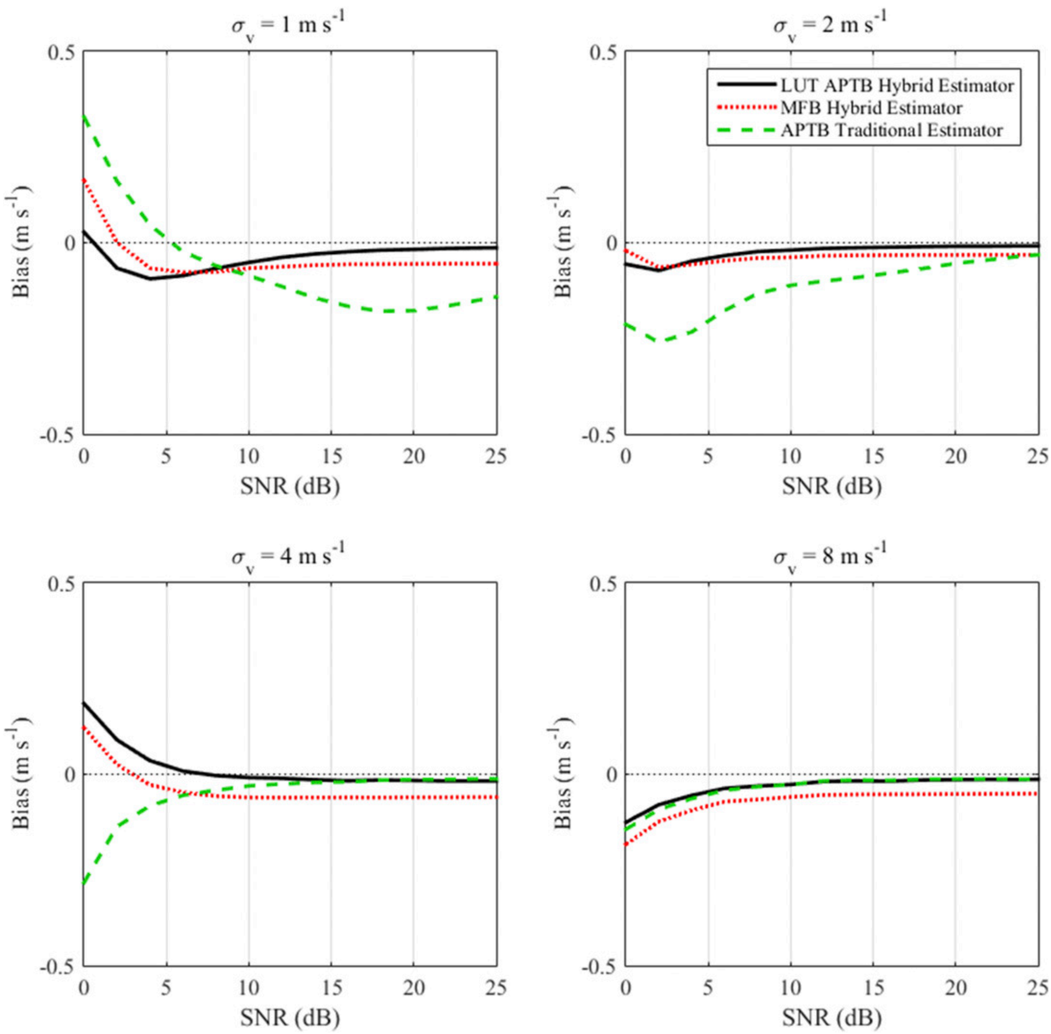

FIG. 7. Bias of spectrum-width estimates as a function of the SNR for the LUT APTB hybrid estimator (black), the MFB hybrid estimator (dotted red), and the APTB traditional estimator (dashed green) with true spectrum widths of (top left) 1, (top right) 2, (bottom left) 4, and (bottom right) 8 m s^{-1}. Statistics for each SNR and spectrum-width value are computed using 50 000 realizations with $L = 5$, the modified pulse of the KOUN radar, $M = 40$, and $\upsilon_a = 25$ m s^{-1}.

Unlike the results in the previous section, we cannot assess the optimality of LUT adaptive pseudowhitening, since a variance expression is not available. Instead, we compare the performance of LUT APTB HSW estimates to that of two other practical alternatives: APTB TSW estimates and MFB HSW estimates. Figures 7 and 8 show the bias and standard deviation of the three estimators as a function of the SNR for true spectrum-width values of 1, 2, 4, and 8 m s^{-1}. Statistics for each SNR and spectrum-width value are computed using Monte Carlo simulations with 50 000 realizations using $L = 5$ and the modified pulse of the KOUN radar. Other signal characteristics are the same as in the previous simulations; that is, $M = 40$ samples per dwell and $\upsilon_a = 25$ m s^{-1}. The LUT for the HSW estimator is computed using the procedure described in section 2 using the modified pulse of the KOUN radar. As expected, the bias of the HSW estimator at narrow spectrum widths (1 and 2 m s^{-1}) is lower than that of the TSW estimator. Still, the bias of LUT APTB HSW estimates is the smallest, except for very small SNRs (less than ~5 dB) with narrow-to-medium widths (1, 2, and 4 m s^{-1}). The standard deviation plots illustrate the distinct advantage of adaptive pseudowhitening to improve data quality. In all cases, the standard deviation of LUT APTB HSW estimates is the smallest. A final metric of quality for any spectrum-width estimator is the likelihood of producing invalid values, which is

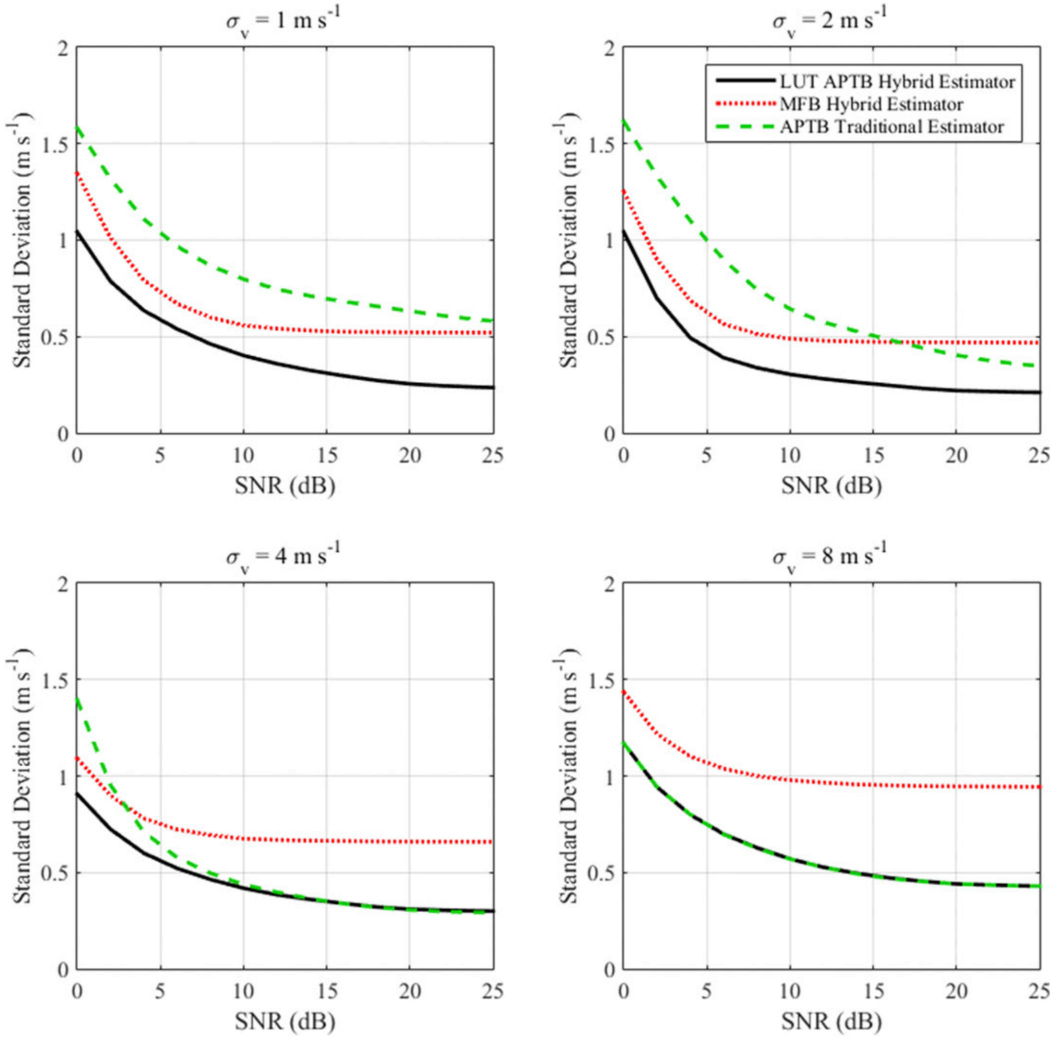

FIG. 8. As in Fig. 7, but for the standard deviation of spectrum-width estimates.

referred to as a failure condition. Invalid spectrum-width values occur when, due to errors of estimates, the autocorrelation ratio is less than one (e.g., for the traditional spectrum-width estimator, this occurs when the lag-1 autocorrelation estimate is larger than the lag-0 estimate, which violates the Gaussian autocorrelation assumption). Figure 9 shows the failure rate for the same parameters as in the previous two figures, where it is clear that LUT APTB HSW estimates are the least likely to produce an invalid result, especially for narrow spectrum widths (1 and 2 m s^{-1}). Next, we validate these results using data collected with the KOUN radar.

On 12 August 2004, the polarimetric S-band KOUN radar sampled a severe storm event southwest of Norman. Figure 10 shows (zoomed in) plan position indicator (PPI) displays at ~2337 UTC of reflectivity (top left), original APTB TSW estimates (top right), conventional MFB HSW estimates (bottom left), and LUT APTB HSW estimates (bottom right). Data shown in this figure correspond to the lowest elevation scan at an elevation of 0.5°. At this elevation, 17 samples were collected at each range resolution volume using a pulse repetition time (PRT) of ~3.1 ms for reflectivity, and 52 samples were collected using a PRT of ~1 ms for radial velocity and spectrum width, which matches the operational parameters of VCP 11 on the NEXRAD network. All spectrum-width fields were obtained using the same time series data and the same ancillary signal processing functions, such as ground clutter filtering and data thresholding. It is important to note that adaptive pseudowhitening was based on range correlation measurements from the data using the technique described by Curtis and Torres (2013).

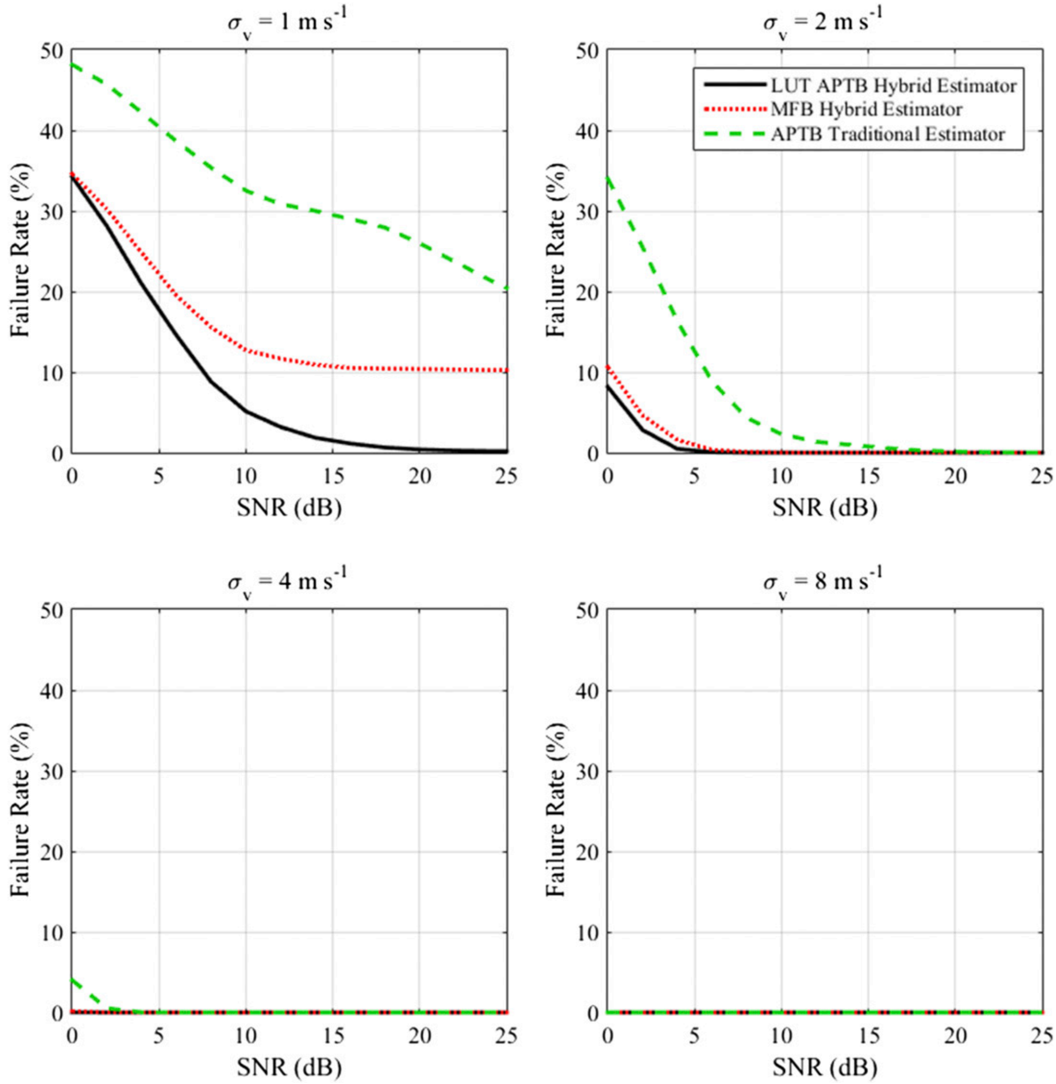

FIG. 9. As in Fig. 7, but for the percentage of estimator failures (i.e., negative spectrum widths).

The three spectrum-width panels in Fig. 10 confirm the results from simulated data. Compared to conventional (no range oversampling) processing (bottom left panel), both APTB fields (right panels) exhibit a smoother texture due to the smaller variance of estimates. However, the field of LUT APTB HSW estimates shows regions of low spectrum widths (deep black colors) more consistently and has many fewer negative estimates (blue). To quantify the improvement afforded by LUT APTB, spatial variances were estimated for the fields of MFB and LUT APTB estimates using a 5×5 mask, and the ratio was taken between the MFB spatial variances and the corresponding LUT APTB ones. The resulting field (shown in Fig. 11) is loosely termed variance reduction factor (VRF); however, it conveys not only the true VRF but also the spatial variability inherent to the fields. Still, VRF values significantly larger than one represent true improvement of LUT APTB over MFB processing. Values close to 5 (the theoretical maximum improvement) are evident in Fig. 11, especially in regions of high SNR and more uniform spectrum widths. In summary, among the three estimators, LUT adaptive pseudowhitening with the HSW estimator exhibits the best quality (in terms of bias and standard deviation) with the least amount of invalid estimates.

5. Conclusions

We developed a procedure that extends adaptive pseudowhitening resulting in compatibility with nontraditional radar-variable estimators that do not have explicit variance expressions. Whereas the original version of adaptive pseudowhitening requires explicit variance

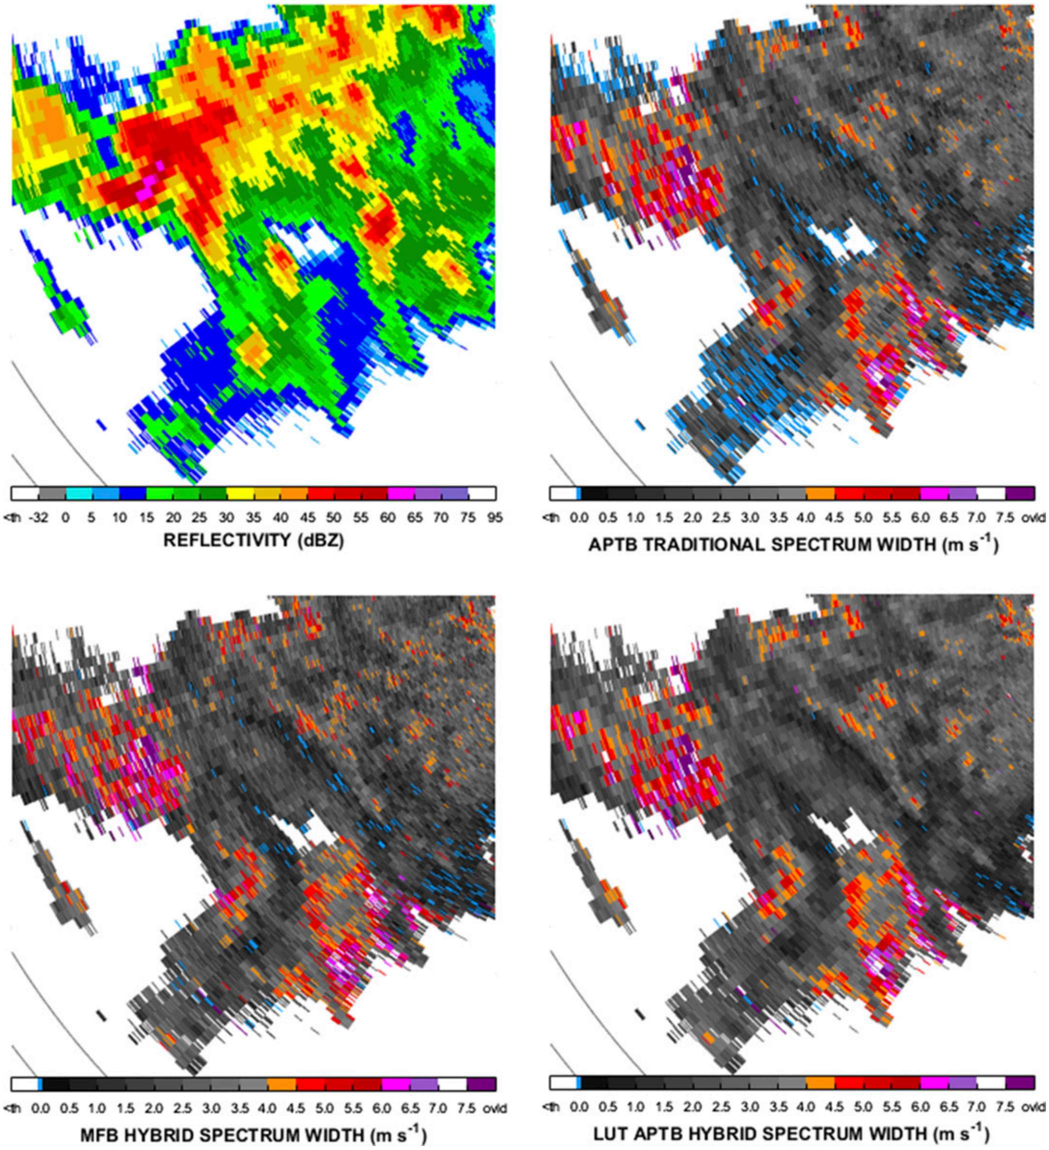

FIG. 10. (top left) Zoomed-in PPI displays of reflectivity, (top right) original APTB TSW estimates, (bottom left) conventional (no range oversampling) MFB HSW estimates, and (bottom right) LUT APTB hybrid spectrum-width estimates. Data were collected with the S-band KOUN radar at ~2337 UTC 12 Aug 2004 at an elevation of 0.5°. Radar used $L = 5$, $M = 17$ with a long PRT of ~3.1 ms ($v_a \approx 8.9\,\mathrm{m\,s^{-1}}$) for reflectivity estimates, and $M = 52$ with a PRT of ~1 ms ($v_a \approx 27.7\,\mathrm{m\,s^{-1}}$) for spectrum-width estimates. The color scale was modified to highlight invalid (negative) spectrum-width estimates (blue).

expressions for each estimator, the proposed lookup table (LUT) approach circumvents this requirement. That is, unlike original adaptive pseudowhitening, which computes nearly optimum pseudowhitening transformations using analytical solutions based on explicit variance expressions, the proposed technique obtains transformations using lookup tables. Although the generation of lookup tables is a computationally complex process, the tables can be precomputed for each radar-variable estimator. Tables need to be recomputed only if the radar-variable estimator changes. For the three spectral moments and the three polarimetric variables, a total of ~45 000 lookup table values would need to be stored (~180 KB with single precision), which are negligible storage requirements for modern signal processors.

The performance of LUT adaptive pseudowhitening using the traditional radar-variable estimators (the three spectral moments and the three polarimetric variables) was shown to be comparable to that of the original technique. Even though both original and LUT adaptive

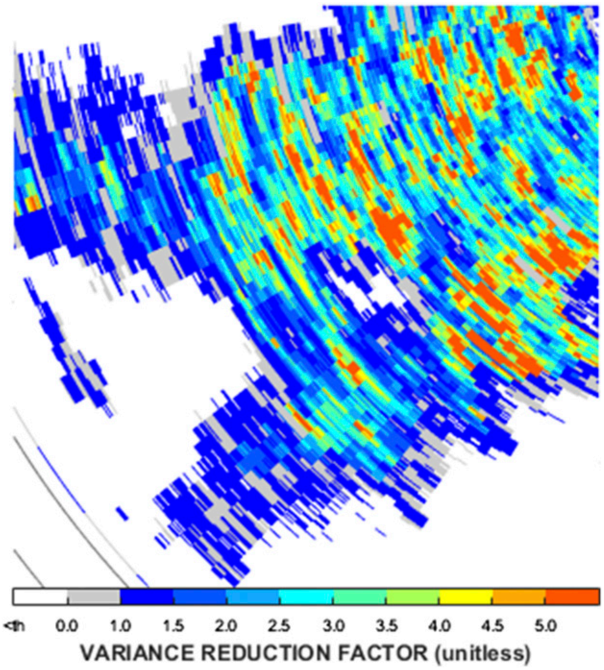

FIG. 11. As in Fig. 10, but for the VRF of LUT APTB HSW estimates with respect to MFB estimates.

pseudowhitening depend on the normalized range correlation matrix, it would not be feasible to recalculate lookup tables in real time to accommodate changes in the range correlation due to hardware drifts. Fortunately, the performance of LUT adaptive pseudowhitening is robust to significant departures of the modified pulse from the one originally used to generate the lookup tables.

The proposed technique was applied to the hybrid spectrum-width estimator, which does not have an explicit variance expression. Compared to two other practical alternatives (the traditional spectrum-width estimator with original adaptive pseudowhitening and the hybrid spectrum-width estimator with conventional matched filter), the hybrid spectrum-width estimator with LUT adaptive pseudowhitening exhibited the smallest bias, the lowest standard deviation, and the fewest number of failures. This was confirmed qualitatively on data collected with the KOUN radar.

In conclusion, the proposed LUT adaptive pseudowhitening technique enables the use of nontraditional radar-variable estimators for which explicit variance expressions are not readily available. This is essential for an operational implementation on the WSR-88D, which already includes the nontraditional hybrid spectrum-width estimator.

Acknowledgments. The authors thank Igor Ivić, David Schvartzman, and the anonymous reviewers for providing comments, which helped to improve the manuscript. Funding was provided by NOAA/Office of Oceanic and Atmospheric Research under NOAA–University of Oklahoma Cooperative Agreement NA11OAR4320072, U.S. Department of Commerce.

REFERENCES

Curtis, C., 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, doi: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, doi: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, doi:10.1175/JTECH-D-13-00216.1.

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

Ivić, I., 2014: Improved correlation coefficient estimator. Preprints, *Eighth European Conf. on Radar Meteorology and Hydrology (ERAD)*, Garmisch-Partenkirchen, Germany, DWD and DLR, 240. [Available online at http://www.pa.op.dlr.de/erad2014/programme/ShortAbstracts/270_short.pdf.]

Lei, L., G. Zhang, R. Doviak, R. Palmer, B. Cheong, M. Xue, Q. Cao, and Y. Li, 2012: Multilag correlation estimators for polarimetric radar measurements in the presence of noise. *J. Atmos. Oceanic Technol.*, **29**, 772–795, doi:10.1175/JTECH-D-11-00010.1.

Meymaris, G., J. Williams, and J. Hubbert, 2009: Performance of a proposed hybrid spectrum width estimator for the NEXRAD ORDA. *25th Int. Conf. on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, Phoenix, AZ, Amer. Meteor. Soc., 11B.1. [Available online at https://ams.confex.com/ams/89annual/techprogram/paper_145958.htm.]

Torres, S., 2009: Processing of oversampled signals in range on polarimetric weather radars with mismatched channels. *J. Atmos. Oceanic Technol.*, **26**, 1289–1301, doi:10.1175/2009JTECHA1209.1.

—— and D. Zrnić, 2003a: Whitening in range to improve weather radar spectral moment estimates. Part I: Formulation and simulation. *J. Atmos. Oceanic Technol.*, **20**, 1433–1448, doi:10.1175/1520-0426(2003)020<1433:WIRTIW>2.0.CO;2.

——, and ——, 2003b: Whitening of signals in range to improve estimates of polarimetric variables. *J. Atmos. Oceanic Technol.*, **20**, 1776–1789, doi:10.1175/1520-0426(2003)020<1776:WOSIRT>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, doi:10.1175/JTECH-D-11-00135.1.

——, and ——, 2013: The importance of accurately measuring the range correlation for range oversampling processing. *J. Atmos. Oceanic Technol.*, **30**, 261–273, doi:10.1175/JTECH-D-12-00085.1.

——, ——, and J. R. Cruz, 2004: Pseudowhitening of weather radar signals to improve spectral moment and polarimetric variable estimates at low signal-to-noise ratios. *IEEE Trans. Geosci. Remote Sens.*, **42**, 941–949, doi:10.1109/TGRS.2004.825579.

Zrnić, D., 1979: Spectrum width estimates for weather echoes. *IEEE Trans. Aerosp. Electron. Syst.*, **AES-15**, 613–619, doi:10.1109/TAES.1979.308849.