The Feasibility of Data Whitening to Improve Performance of Weather Radar

A. C. KOIVUNEN AND A. B. KOSTINSKI

Department of Physics, Michigan Technological University, Houghton, Michigan

(Manuscript received 5 March 1998, in final form 19 August 1998)

ABSTRACT

The problem of efficient processing of correlated weather radar echoes off precipitation is considered. An approach based on signal whitening was recently proposed that has the potential to significantly improve power estimation at a fixed pulse repetition rate/scan rate, or to allow higher scan rates at a given level of accuracy. However, the previous work has been mostly theoretical and subject to the following restrictions: 1) the autocorrelation function (ACF) of the process must be known precisely and 2) infinite signal-to-noise ratio is assumed. Here a computational feasibility study of the whitening algorithm when the ACF is estimated and in the presence of noise is discussed.

In the course of this investigation numerical instability to the ACF behavior at large lags (tails) was encountered. In particular, the commonly made assumption of the Gaussian power spectrum and, therefore, Gaussian ACF yields numerically ill-conditioned covariance matrices. The origin of this difficulty, rooted in the violation of the requirement of positive Fourier transform of the ACF, is discussed. It is found that small departures from the Gaussian form of the covariance matrix result in greatly reduced ill conditioning of the matrices and robustness with respect to noise. The performance of the whitening technique for various meteorologically reasonable scenarios is then examined. The effects of additive noise are also investigated. The approach, which uses time series to *estimate* the ACF from which the whitener is constructed, shows up to an order of magnitude improvement in the mean-squared error of the estimated power for a range of parameter values corresponding to typical meteorological situations.

1. Introduction

Meteorological radars typically sample many times in order to form an estimate of power as an arithmetic mean. As is well known, the statistical error (standard deviation) of such an estimate decreases as the inverse square root of the number of independent samples. However, consecutive radar echoes from precipitation are highly correlated (typical interpulse separation on the order of 1 ms and decorrelation or ''reshuffling'' times on the order of 0.1 s). Consequently, in most meteorological applications of coherent radar, the total number of samples greatly exceeds the number of statistically independent samples, for example, the number of independent samples (10–100 samples s^{-1}) can be 10–100 times smaller than the total number of samples and loss of information can be substantial (see, e.g., Atlas 1964).

Yet, it is usually taken for granted that only the so-called effective number of independent samples (call it N_I) matters and convergence of the average power estimate scales as $\propto 1/\sqrt{N_I}$. This is probably so because one of the following two circumstances is always assumed:

- the standard sample variance estimator is employed (whose variance increases with the increasing decorrelation time), and
- the autocorrelation function (ACF) of the process is not known.

However, even when the ACF is estimated from data (e.g., Gaussian spectrum), it is a common procedure throughout remote sensing to use the estimated autocorrelation function to convert correlated time series to the equivalent number of independent samples (e.g., Ulaby et al. 1986, p. 488; Doviak and Zrnić 1993, p. 127; Nathanson 1990, p. 93; Sauvageot 1992, p. 53). This is done to indicate the types of errors that are likely, based on the $1/\sqrt{N}$ estimator variance dependence. *Why is the ''equivalent number of independent samples'' approach taken for granted in radar meteorology and in other radar applications*?

For example, in the case of incoherent radar, why should one assume that the arithmetic mean estimator of mean power is the optimal one even in cases where the radar echoes are strongly correlated? Can one do more with the ACF information than count one's losses? For instance, could one use a weighted mean where the

Corresponding author address: Prof. A. B. Kostinski, Dept. of Physics, Michigan Technological University, 1400 Townsend Dr., Houghton, MI 49931-1295.
E-mail: kostinsk@phy.mtu.edu

© 1999 American Meteorological Society

weights ''know'' about the estimated ACF? Recall that the ACF is often at least approximately known in radar meteorology [i.e., Gaussian shape of the spectrum is commonly assumed (e.g., Doviak and Zrnić 1993)]. Yet, rather than to find a better estimator, the estimated autocorrelation function is used to compute an equivalent number of independent samples. Hence, a natural question arises: Is there a power estimator that can ''extract'' information from all of the correlated samples if the autocorrelation function of the process is known? Also, given the inevitable Rayleigh fading, what is the best accuracy of determining echo intensity for a given number of (correlated) echoes?

The theoretical results obtained in Monakov (1994) and in Schulz and Kostinski (1997) show that in the coherent radar case (multivariate Gaussian probability density of the complex echo amplitudes) such an estimator does exist. The appropriate convergence bounds are derived and the whitening approach is also given in Schulz and Kostinski (1997). The reason for the existence of such an estimator can be understood in retrospect by realizing that the original amplitude data can be ''whitened'' by a linear transformation whose coefficients know the ACF. Furthermore, this transformation leaves the variance (power) of the amplitude time series unchanged. The resultant independent samples can then be used in a regular manner. In other words, a variance estimator exists whose coefficients are ACF dependent (so that all of the information is taken into account), which converges as $1/\sqrt{N}$, where N is the total number of (correlated) samples.

However, the theoretical method developed in Schulz and Kostinski (1997) relies on the following possibly restrictive assumptions: 1) infinite signal-to-noise ratio (SNR) and 2) ACF is known precisely. The assumption of infinite SNR is not strictly valid, but it is often approximately valid, for example, severe weather signal can exceed 70 dB. In weather applications, the only time the signals are so weak that the SNR is comparable to unity is when the precipitation is either quite weak (e.g., light drizzle) or very distant. But meteorologically most interesting and practically most important cases are those involving severe weather not too far away and then SNR is high. As far as the ACF restriction, in defense of the whitening method, it can be pointed out that similar objections can also be advanced against other methods in use which rely, for example, on the Gaussian shape of the spectrum such as the pulse pair processor. It is widely and effectively used in spite of the fact that the spectrum is often not Gaussian and, in fact, is bimodal in 25% of the cases. See, for instance, Janssen and van der Spek (1985, pp. 208–219).

Given the potential benefits of the decorrelation technique, it seems worthwhile to proceed with the quantitative evaluation of the sensitivity of the whitening method to deviations from the two assumptions above. What happens when the ACF must be estimated or is known imprecisely and SNR is not infinite? Is the ACF estimation so poor or the noise sensitivity so severe as to render any practical application useless? An attempt to answer these questions is the main objective of this work.

2. Signal model

As is well known, constant ''reshuffling'' of hydrometeors and randomly constructive and destructive interference cause the radar echo to fluctuate. Doppler weather radar signals are usually modeled as a zero-mean complex Gaussian random process z, with discrete samples $z_m = s_m \exp(\omega_D i T_s m) + n_m$, where $s_m \exp(\omega_D i T_s m)$ is the precipitation echo signal returned to the receiver, n_m is uncorrelated noise, ω_D is the mean Doppler shift frequency, T_s is the sampling interval, and m is the sampling index. If we, for example, assume a Gaussian power spectrum for the returned signal, we may write

$$S(\nu) = \frac{S}{\sqrt{2\pi}\sigma_\nu} \exp\left[\frac{-(\nu - \overline{\nu})^2}{2\sigma_\nu^2}\right] + \frac{2\kappa T_S}{\lambda}, \quad (1)$$

where $S = \langle |s_m|^2 \rangle$ is the average signal power, σ_ν is the velocity spectral width, $\overline{\nu}$ is the mean radial velocity of the scattering particles, λ is the wavelength, and κ is the noise power. The corresponding autocorrelation is written,

$$\rho(\tau) = S \exp[-8(\pi\sigma_\nu \tau/\lambda)^2] \exp[-4\pi i \overline{\nu}\tau/\lambda] + \kappa\delta(\tau), \quad (2)$$

or, discretely,

$$\rho(mT_S) = S \exp[-8(\pi\sigma_\nu m T_S/\lambda)^2] \exp[-4\pi i \overline{\nu} m T_S/\lambda] + \kappa\delta(mT_S), \quad (3)$$

where m is the sampling index.

The covariance matrix, K_ρ, with elements $K_{j,m} = \langle z_j z_m^* \rangle$, where $K_{j,m} = K_{j,-m}^*$, is Hermitian. The signal spectral width, σ_ω, in the frequency domain is related to the spectral width in the velocity domain by the relation

$$\sigma_\nu = \sigma_\omega \lambda/2. \quad (4)$$

The variance of our estimate of the average scattered power, $\sigma_{\overline{P}}^2$, depends on the number of *independent* samples in the average according to

$$\sigma_{\overline{P}}^2 = \frac{\sigma_P^2}{N_I}, \quad (5)$$

where N_I is the number of *independent* samples in the average.

Weather radars typically sample at about 1 kHz, while the time to independence is in the range of 0.1–0.01 s (on the order of time it takes for hydrometeors to change relative positions by about a radar wavelength). Hence, consecutive samples are typically strongly correlated.

Increasing the independence of the samples in our estimate of the average power would reduce the variance of the estimate, thereby increasing the amount of information extracted from a given set of data. The information needed to decrease the correlation between the samples used in the estimate of the average scattered power is contained in the scattered power ACF. Can one use the estimated ACF information to decorrelate dependent samples through *whitening*? Theoretically, at least, the answer is yes (Schulz and Kostinski 1997) because the Gaussian quadrature amplitude probability density function (PDF) can remain unchanged with linear filtering, that is, whitening. Below, we show that the method is practically feasible as well.

3. Whitening

For the reader's convenience we briefly review some definitions of the random process theory needed later. Consider the autocorrelation function of a random process z [sequence of radar echoes in the complex amplitude format (e.g., see Papoulis 1984)]:

$$\rho(t_1, t_2) = E[z(t_1)z^*(t_2)] \tag{6}$$

$$\rho(t_1, t_2) = \rho(t_1 - t_2) = \rho(\tau) = E[z(t)z^*(t + \tau)]. \tag{7}$$

The average power is then

$$\overline{P} = E[z^2(t)] = \rho(\tau = 0) \tag{8}$$

and the variance of the process can be written

$$\begin{aligned}\sigma_z^2 &= E\{[z(t) - \overline{z}]^2\}\\ &= E[z^2(t)] - 2E[z(t)\overline{z}] + E(\overline{z^2}),\end{aligned} \tag{9}$$

where $\overline{z}$ is the mean. For a zero-mean process such as in-phase and quadrature components of radar echoes

$$\overline{P} = E[z^2(t)] = \sigma_z^2, \tag{10}$$

so the average return power equals the variance of the amplitude process.

Next we present a brief summary of the relevant development from Schulz and Kostinski (1997). Consider a time series of correlated echoes $\{z_n\}_{n=0}^{N-1}$. Given the autocorrelation function $\rho(\tau)$, can we "whiten" this time series without affecting its variance (power)? To that end, construct the covariance matrix $\mathbf{K}\rho$ as follows:

$$\mathbf{K}_\rho = \begin{bmatrix} \rho(0) & \rho(1) & \cdots & \rho(N-1) \\ \rho(-1) & \rho(0) & \cdots & \rho(N-2) \\ \vdots & \vdots & \ddots & \vdots \\ \rho(1-N) & \rho(2-N) & \cdots & \rho(0) \end{bmatrix}, \tag{11}$$

which is Hermitian and Toeplitz (elements along each of the diagonals have the same value). This matrix can always be factored as

$$\mathbf{K}_\rho = \mathbf{U}_\rho \boldsymbol{\Lambda}_\rho \mathbf{U}_\rho^{\mathrm{T}}, \tag{12}$$

where $\mathbf{U}_\rho$ is orthonormal, $\mathbf{U}_\rho\mathbf{U}_\rho^{\mathrm{T}} = \mathbf{I}$, $\mathbf{U}_\rho^{\mathrm{T}}$ is the transpose of $\mathbf{U}_\rho$, and

$$\boldsymbol{\Lambda}_\rho = \begin{bmatrix} \lambda_0 & 0 & \cdots & 0 \\ 0 & \lambda_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_{N-1} \end{bmatrix} \tag{13}$$

is a diagonal matrix with the eigenvalues of $\mathbf{K}_\rho$ on the diagonal. If we form the matrix $\mathbf{H} = \boldsymbol{\Lambda}_\rho^{-1/2}\mathbf{U}_\rho$, where $\boldsymbol{\Lambda}_\rho^{-1/2}$ is a diagonal matrix with the elements $\lambda_\rho^{-1/2}$ on the diagonal, then the transformed data, $Y_k = \Sigma_{l=0}^{N-1} H_{k,l}X_l$, $k = 0, \ldots, N - 1$, are independent, identically distributed random variables.

The scattered power can thus be estimated using the transformed data, where the number of *independent* samples has been increased to the total number of samples, N. The resulting variance in the scattered power estimator will be reduced, improving the estimate of average scattered power. However, infinite SNR and precisely known ACF have been assumed. Our objective is to explore, via computer calculations, the sensitivity of the whitening technique to the assumptions of perfectly known ACF and infinite SNR. To that end, in the remainder of this paper we construct examples of time series with specifically prescribed correlation functions and SNRs, whiten them with transforms derived from ACFs estimated from sample time series, and compare the power in the time series before and after whitening. Next, let us discuss the numerical stability of the whitening transform calculation versus the assumed form of the ACF.

4. The Gaussian anomaly

To motivate two particular ACF parameterizations used in the remainder of this paper, we next describe what we termed the "Gaussian anomaly." In the course of our whitening calculations we noticed that there is a surprising sensitivity to the assumed form of the autocorrelation function. In particular, whitening worked well in the case of the exponential ACF (which is in agreement with Monakov's results) but less well in the case of the Gaussian one. After numerical experimentation, we found that the Gaussian case is numerically ill conditioned in the sense that the condition number (ratio of maximal to minimal eigenvalue) is extremely large for the Gaussian covariance matrix. The root of the difficulty surprised us and it is related to the Wiener–Khintchin theorem. *The numerical instability turns out to be caused by the violation of the requirement of a positive Fourier transform of the ACF.*

Let us illustrate it on the following example. Consider an ACF of the form $\exp[-\tau^\alpha]$. Then, it is a valid ACF if and only if $0 < \alpha \le 2$ (e.g., Yaglom 1987). Hence, the Gaussian ACF is right at the boundary ($\alpha = 2$) and therefore results in ill-conditioned covariance matrices in the presence of even the smallest computer noise. We present calculations to prove this in section 5 on the condition number. The important thing to note now is that the above restriction on α dictates our choice of the ACF functional form in sections to follow. In other words, consider a continuous autocorrelation function

$$\rho(\tau) = \exp[-8(\pi\sigma_\nu\tau/\lambda)^\alpha] = \exp[-\beta\tau^\alpha], \quad (14)$$

where this form is chosen to explore the Gaussian anomaly as α approaches 2. The discrete version of this ACF, sampled at equally spaced time intervals, is given by

$$\rho(mT_s) = \exp[-8(\pi\sigma_\nu mT_s/\lambda)^\alpha] = \exp[-\beta m^\alpha] = r^{m^\alpha}, \quad (15)$$

where $\beta = 8(\pi\sigma_\nu T_s/\lambda)^\alpha$ and $r = \exp[-\beta]$. For example, if T_s is 1 ms, λ is 10 cm, σ_ν is 3.6 m s^{-1}, and the autocorrelation is Gaussian so that $\alpha = 2$, then $\beta \cong 0.11$ and $r \cong 0.9$. Hence, under these meteorological conditions the correlation coefficient between consecutive samples is about 0.9.

5. Condition number

To implement the whitening procedure described earlier, one needs to perform an inversion of the covariance matrix. The latter is ill conditioned when the ratio of maximal-to-minimal eigenvalue (the condition number) is very large. As discussed before, our computations have shown that the condition number is much larger in the case of the Gaussian ACF than in other cases we tried. In this section we present calculations to explain this Gaussian anomaly.

We explore the effects of departure from the Gaussian form by parameterizing the ACF, as indicated in Eq. (14) or using the discrete version (sampled at evenly spaced intervals) of Eq. (15), where $r \in (0, 1)$ is the correlation coefficient between two consecutive samples and α determines how Gaussian the ACF is (with 2 being pure Gaussian). Valid autocorrelation sequences must posses a nonnegative Fourier transform (physically meaningful power spectrum). The eigenvalues of the corresponding covariance matrices, λ_k, are nonnegative, which turns out to force $0 < \alpha \leq 2$ for ACFs of the family $\rho(\tau) = r^{\tau^\alpha}$ (Yaglom 1987).

Note that as $r \to 1$, the maximum eigenvalue approaches N, the length of the autocorrelation sequence, and all other eigenvalues approach zero. Physically, at $r = 1$, all measurements are perfectly correlated and no new information is provided by additional measurements. Then the covariance matrix is singular. Thus, both r and α greatly affect the numerical stability of the whitening transform. The rate at which the covariance matrix approaches singularity as r approaches unity depends strongly on α, with the Gaussian case $\alpha = 2$ forcing the approach to singularity most rapidly.

We now proceed to illustrate the Gaussian anomaly quantitatively. Let us quote from (Priestly 1981, p. 261) ''eigenvalues of the covariance matrix . . . are proportional to the values of the (formal) 'spectral density function' at the frequencies $\{2\pi k/N\}$,'' where the spectral density values, ν_k, are computed as follows:

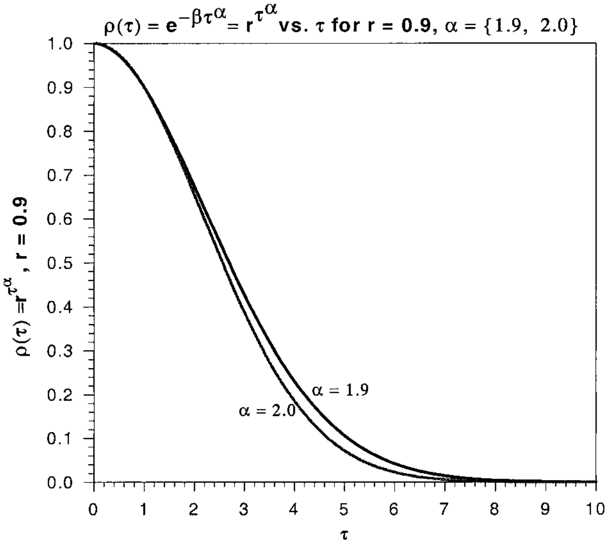

FIG. 1. Gaussian and nearly Gaussian functions. Small changes in the parameter α result in small changes in the shape of the curve $\rho(\tau) = (0.9)^{\tau^\alpha}$.

$$\nu_k = \frac{1}{\sqrt{N}} \sum_{r=1}^{N} \rho(r) e^{-i2\pi kr/N} \qquad k = 1, 2, \ldots, N. \quad (16)$$

Since the values of power spectral density cannot be negative, the eigenvalues of the covariance matrix must also be nonnegative.

Figure 1 illustrates the parameterization and displays the two ACFs $\rho_{1,2}(\tau) = r^{\tau^{\alpha_{1,2}}}$, with $r = 0.9$, $\alpha_1 = 1.9$, and $\alpha_2 = 2.0$. The slight difference between the curves lies well within the variation of meteorological ACF sequences (Janssen and van der Spek 1985), and the corresponding condition numbers for the two covariance matrices derived from 32-point sequences of these forms, are about 1.1×10^3 for $\alpha_1 = 1.9$ and about 1.4×10^9 for $\alpha_2 = 2.0$. This is striking, as small differences in α produce very large differences in condition number and consequently greatly affect the computability of the whitening transform. Figure 2 displays the effect of α on condition number for 4-, 8-, 16-, 32-, and 64-point ACF sequences where $r = 0.98$. [An example four-point ACF sequence would be $(0.98^{0^{1.9}}, 0.98^{1^{1.9}}, 0.98^{2^{1.9}}, 0.98^{3^{1.9}})$]. Note that the condition number is well behaved up to $\alpha \approx 1.95$ for 32- and 64-point sequences (condition number $\approx 2.7 \times 10^4$ and 2.8×10^4, respectively), indicating relatively easy computation of the corresponding whitening transforms.

6. Potential improvement over the equivalent number of independent samples

As already mentioned in the introduction, the degree of correlation between echoes depends on the distribution of raindrop velocities in the resolution volume, which in turn depend on turbulence, wind shear, beam broadening, and fall velocity (e.g., see Nathanson 1968).

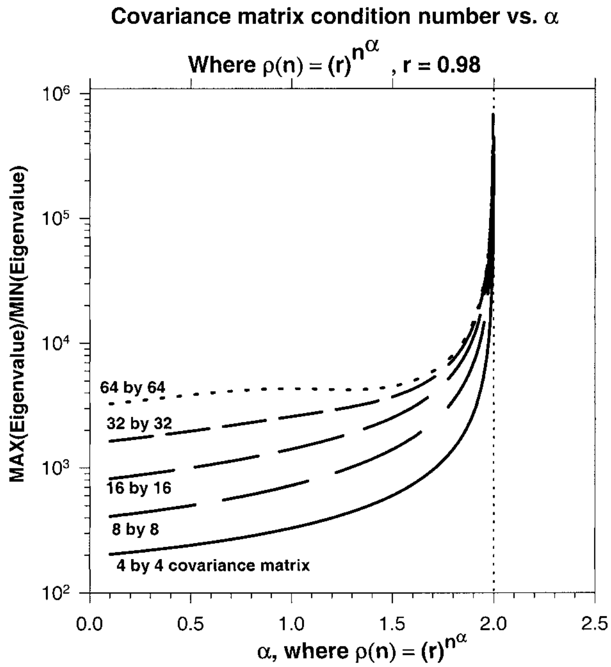

FIG. 2. Condition number vs α for covariance matrices derived from autocorrelation functions of the form $\rho(n) = \exp[-\beta n^{\alpha}] = (r)^{n^{\alpha}}$ and $r = 0.98$.

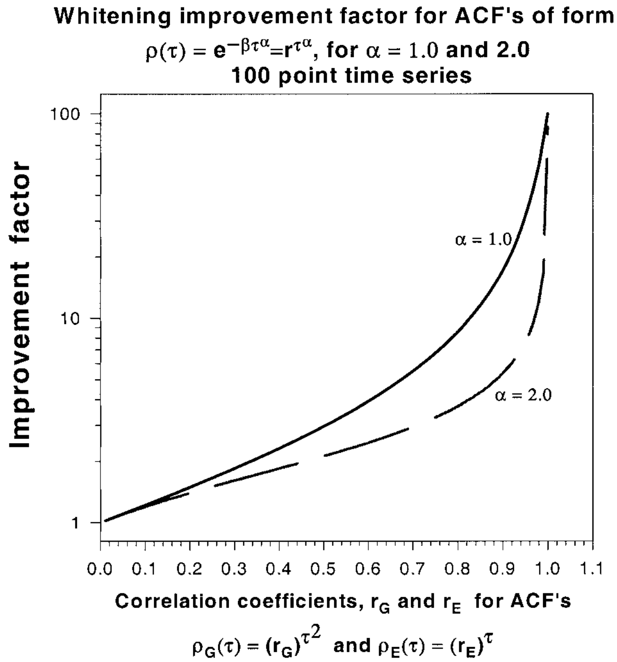

FIG. 3. Factor by which perfect, noise-free whitening increases the number of independent samples in correlated 100-point time series for time series with exponential and Gaussian autocorrelations.

The variance of the power estimate is inversely proportional to the number of *independent* samples in the estimate; that is, $\sigma_{\hat{P}}^2 = \sigma_P^2/N_I$. Current practice in radar meteorology is to use estimated ACFs or power spectra to determine the effective number of independent samples [see Eq. (17)] and to keep sampling until this effective number is large enough to drive the variance of the estimate, $\sigma_{\hat{P}}^2$, down to an acceptable level. For a stationary random process, the relation between the total number of samples N, the effective number of independent samples N_I, and the ACF $\rho(n)$ is expressed (e.g., see Doviak and Zrnić 1993):

$$\frac{\sigma_{\hat{P}}^2}{\sigma_P^2} = N_I^{-1} = \sum_{n=-(N-1)}^{N-1} \frac{N - |n|}{N^2}\rho(n). \tag{17}$$

Meteorological time series are commonly assumed to have Gaussian ACF (Gaussian power spectrum). The assumption is often made for convenience—the Gaussian form is analytically tractable and a good approximation for the near-center portion of the power spectrum (e.g., see Srivastava et al. 1979). However, Janssen and van der Spek (1985) find that this is not always a good approximation even in the center and especially near the tails. To investigate the effect of whitening applied to non-Gaussian autocorrelations, we will parameterize the autocorrelation (because of the Gaussian anomaly) as $\rho(\tau) = \exp[-\beta\tau^{\alpha}] = r^{\tau^{\alpha}}$.

Figure 3 displays the ratio N/N_I versus correlation coefficient $r_{E,G}$ for ACFs of the form $\rho_E(n) = r_E^n$ and $\rho_G(n) = r_G^{n^2}$, where N is 100. If the ACF were known exactly, the improvement in estimation variance, due to whitening, would follow these curves. That is, if a sample 100-point time series possessed an exponential ACF with a correlation coefficient of $r_E = 0.8$, there would be about 11 effective independent samples in the dataset. Perfect whitening would raise the number of independent samples to 100. From this figure we see the potential efficiency improvement in the use of data, particularly for highly correlated samples—rapid sampling and/or light drizzle conditions. This increased efficiency could be used to improve radar echo power estimate accuracy at a given sample rate or to allow higher sample rates, and therefore higher scan rates, at currently acceptable levels of accuracy. This may be particularly beneficial in spaceborne radar applications (e.g., see Meneghini and Kozu 1990, pp. 46–47).

7. Generation of time series with specified correlation

Let us now return to the exploration of the effects of the use of ACFs estimated from data, and the addition of noise on data whitening. To test this, we must first construct time series with a prescribed autocorrelation sequence. Following a procedure described in Johnson (1994), we write

$$X_n = \sum_{m=0}^{N-1} T_{n,m}u_m, \tag{18}$$

where the u_m are independent, normally distributed, zero-mean, unit variance random draws. The constructed random draws, X_n, have the desired correlations if the matrix $\mathbf{T}$ is a root of the covariance matrix $\mathbf{C}$,

$$\mathbf{T}\mathbf{T}^{\mathrm{T}} = \mathbf{C}, \tag{19}$$

where the covariance matrix $\mathbf{C}$ describes the desired time series correlation.

8. Whitening with an estimated ACF

In a realistic application of the whitening technique, the user must obtain an ACF to be used to calculate the whitener. The candidate ACF will likely be estimated in the presence of noise and probably will not be exactly Gaussian or any other analytic form. In this section we investigate a way to apply the whitening technique to noisy measured data. To simulate the actual use of the whitening transform, we first simulate measured time series using the method of section 7, which allows us to control the variance and autocorrelation of the time series. We then add zero-mean uncorrelated Gaussian random noise to the time series—scaling the variance of this noise to control the time series SNR.

We generated up to 50 realizations of the noisy time series, computed the sample autocorrelations of each time series using the standard routine of Marple (1987, p. 168), and averaged them. This average sample ACF, $\overline{\rho}(n)$, was then used to construct a whitener that was applied to subsequent time series. We did not use the average ACF to *directly* compute the whitening transform, but fitted $\overline{\rho}(n)$ to an equation of the form $\rho(n) = r_w^{n^{\alpha_w}}$ and then constructed the whitener from this analytic form. We very simply solved for the parameters α_w and r_w using lags 1 and 2 of $\overline{\rho}(n)$. That is, we set $r_w = \overline{\rho}(1)$ and $\alpha_w = \ln[\overline{\rho}(2)]/\ln(r_w^2)$.To test the whitener, we then generated up to 10 000 realizations of the noisy time series and computed the mean-squared error (mse) in the variance estimate both before and after whitening and compared the two errors. The mean-squared error is defined as $-\text{mse} = M^{-1}\sum_{j=1}^{M}(P_j - P_0)^2$, where P_0 is the average power in the noise-free, unwhitened time series (unity) and P_j is the power in the jth time series realization. The mse was chosen as a measure of the error in the estimate because it detects both systematic bias error and the fluctuation about the mean due to finite sampling.

This computational procedure is divided into two main parts:

1) Prepare the whitener.
   (a) Generate unit variance time series (we generated 50 realizations).
   (b) Shape the time series with prescribed ACF.
   (c) Add the desired amount of noise to the time series.
   (d) Compute the sample ACF for each realization.
   (e) Average the ACFs.
   (f) Fit a whitener ACF of the form $\rho_w(n) = r_w^{n^{\alpha_w}}$ to the average ACF.
   (g) Build the whitening transform with $\rho_w(n) = r_w^{n^{\alpha_w}}$.
2) Test the whitener.
   (a) Generate up to 10 000 realizations of noisy time series with the same statistics used to construct the whitener.
   (b) Whiten the sample time series.
   (c) Compute the mse's of variance estimates for the whitened and unwhitened time series, one realization at a time.
   (d) Compare the mse's to evaluate the whitening technique.

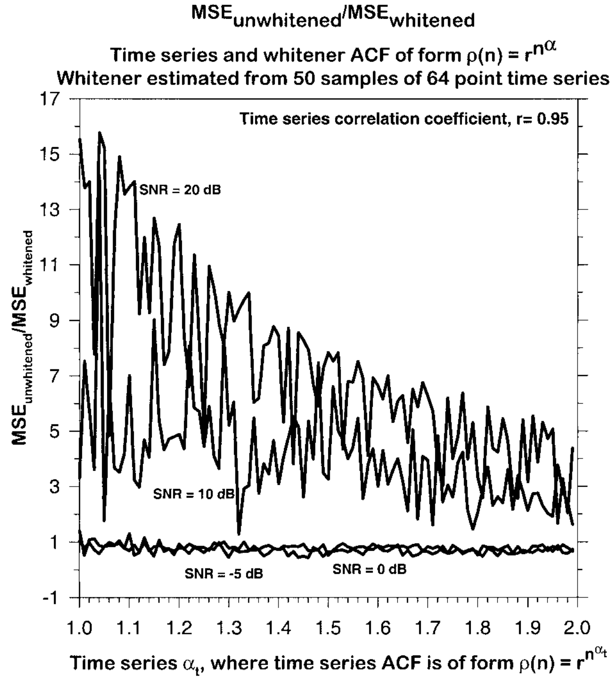

FIG. 4. Ratio of unwhitened to whitened mse. The whitener ACF was estimated from 50 samples of the noisy time series, assuming a form $\rho(n) = r^{n^{\alpha}}$.

We generated samples of 64-point time series with ACFs of the form $\rho_T(n) = 0.95^{n^{\alpha_t}}$, with α_t varying from 1 (exponential) to 2 (Gaussian) and added noise so that we had SNRs of −5, 0, 10, and 20 dB. We used 50 realizations of the time series to ''train'' the whitener—corresponding to about 3.2 s of data at a 1-kHz pulse repetition rate. We then generated 10 000 samples of 64-point time series at each value of α_t and computed the mse's in the variance estimates for the unwhitened and whitened time series. The ratio of these errors are plotted versus α_t in Fig. 4.

Values of the ratio greater than 1.0 indicate improvement in the power estimate due to whitening. In fact, the y-axis values indicate the factor by which the mse is improved by whitening. A value of $\text{mse}_{\text{unwhitened}}/\text{mse}_{\text{whitened}} = 2.0$ means that the error in the unwhitened power estimate is twice that in the whitened power estimate. It is apparent from the figure that except for

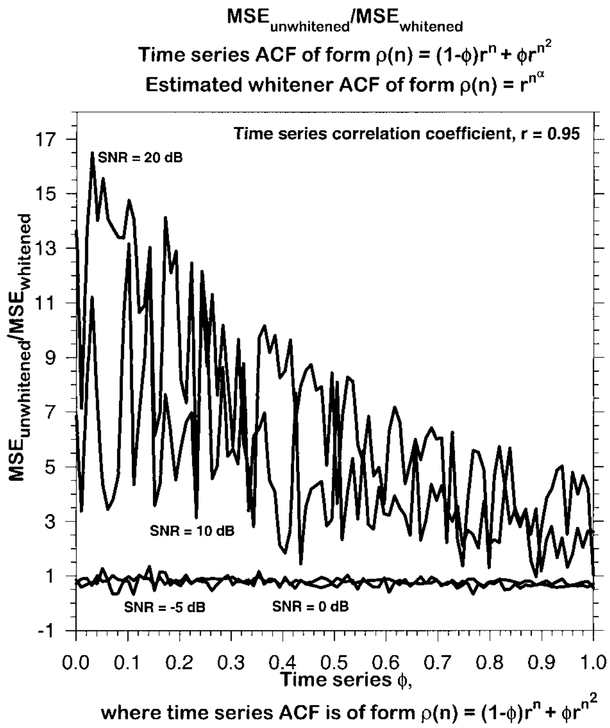

FIG. 5. Ratio of unwhitened to whitened mse's. The whitener ACF was estimated from 50 samples of the noisy time series, assuming a form $\rho(n) = r^{n^\alpha}$.

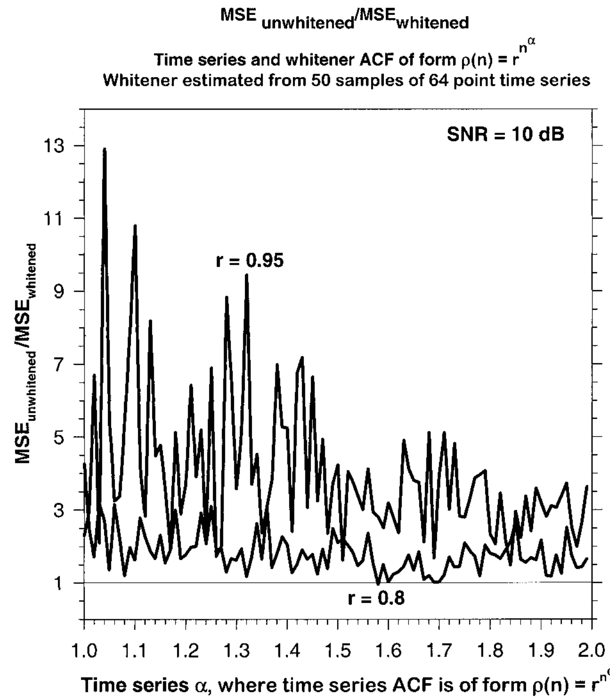

FIG. 6. Ratio of unwhitened to whitened power estimate mse vs α, where both time series and whitener ACF are of the form $\rho(n) = r^{n^\alpha}$. The whitener is estimated from 50 samples of 64-point time series.

SNR $\leq$ 0 dB, whitening improves the power estimate, in the mean-squared sense, over the full range of α_t. The great variability in $\text{mse}_{\text{unwhitened}}/\text{mse}_{\text{whitened}}$ as α_t varies is due to the relatively small number of time series realizations, corresponding to $\cong$3.2 s of data at 1 kHz pulse repetition frequency, used to train the whitener. However, this training can be continuously and adaptively updated with an ongoing data stream.

Furthermore, despite the variability, it seems clear from Fig. 4 that whitening has potential to substantially improve power estimates. As the time series ACF gets increasingly Gaussian ($\alpha_t \rightarrow 2$), the benefit due to whitening decreases but is still quite good, with an improvement factor of approximately 2 even at $\alpha_t = 1.98$ and SNR $=$ 10 dB. We carried out the same calculations with $r = 0.8$ and found improvement for all $\alpha_t \in (1.0, 2.0)$, for SNR $>$ 0 dB.

In Fig. 4 we assumed the ''correct'' form for the whitener ACF, that is, both time series and whitener ACFs of the form $\rho(n) = r^{n^\alpha}$. We also consider the results when we fit an ''incorrect'' form of the whitener ACF to the sample ACF. In Fig. 5 the time series were generated with ACF $\rho_t(n) = \Phi 0.8^{n^2} + (1 - \Phi)0.8^n$, a mixture of exponential and Gaussian ACFs, but the whitener was still constructed with an ACF of form $\rho_w(n) = r_w^{n^{\alpha_w}}$. Here again, we fit the whitener ACF form to the average ACF computed with 50 realizations of the time series. We see that despite using an incorrect parametric form of the ACF to build the whitener, we still reduce the mse substantially by whitening the time series before computing the variance.

We illustrate the dependence of whitening on the magnitude of the correlation coefficient in Fig. 6. Here we plot $\text{mse}_{\text{unwhitened}}/\text{mse}_{\text{whitened}}$ versus α for time series generated with ACF of the form $\rho_T(n) = r^{n^\alpha}$, with $r =$ (0.85, 0.90, and 0.95) and whitened with a transform derived from an estimated ACF of the same form, as described earlier in this section. We set the signal to noise ratio at SNR $=$ 10 dB. Again, there is great variability in the ratio as α is varied, but the ratio exceeds unity throughout the range of α, indicating power estimate improvement. The improvement is greatest for highly correlated time series. This is to be expected because highly correlated time series have few independent measurements and consequently high estimator variance. Whitening then has greater potential to reduce the correlation between successive measurements and increase the number of independent samples, thereby reducing the variance of the power estimator. We also see that the benefits of whitening are greater for time series with more nearly exponential ACFs. Note, however, that even at $\alpha = 1.98$, $\text{mse}_{\text{unwhitened}}/\text{mse}_{\text{whitened}}$ is still greater than unity.

9. Concluding remarks

In this paper we explored the practical feasibility of the ''ideal case'' theoretical whitening approach proposed in Schulz and Kostinski (1997) to improve meteorological radar echo power estimates. Three issues not adequately addressed in that paper were 1) the numerical instability of commonly used covariance matrices, 2) the lack of precisely known ACF and the necessity of estimating the correlation function to compute the whitening transform, and 3) the effects of noise.

All three points have been addressed in this paper. In section 5, the numerical instability and ill conditioning were related to the nonnegativity requirement of the ACF and the Gaussian anomaly (the fact that the Gaussian ACF is right at the edge of what is allowed by the nonnegativity condition). Furthermore, a simple parameterization was suggested to avoid this problem.

Concerning the other two points, the results of section 8 strongly suggest that the whitening transform is a robust technique (with respect to noise), which is capable of significant improvement of power estimates. We showed that estimates of the time series ACF, in the presence of noise and even assuming an incorrect functional form for the time series ACF, can be used to design a whitener that results in improved power estimates for SNRs greater than or equal to about 5 dB. No claim is made here that the proposed approach can be applied to all meteorological situations and all kinds of autocorrelation functions. However, we have demonstrated here that it can be applied to a variety of situations: varying correlation coefficient magnitude, ACF functional form, and SNR. The standard two-lag ACF estimator we used here is not optimal, but it is simple and practical. Perhaps the potential benefit may justify the computational overhead involved in implementing a more sophisticated ACF estimator. Recall, however, that only a few lower lags of the Doppler spectrum are required for most estimators widely used in radar meteorology such as 2 for the pulse-pair estimator (see, e.g., Passarelli and Siggia 1983, p. 1783).

Our simulation results suggest that SNR values of about 5 dB or greater may be necessary for noticeable improvement via whitening. This is not very restrictive as severe weather echoes are often in excess of 50 dB. Recall that drops reshuffle quickly in severe weather cases, implying that small interpulse correlation coefficients and nearly singular values of, say, $r = 0.999$, are rather unlikely, for example, see the discussion in Kostinski (1994, p. 1179). Several initial sets of time series could be taken and used to compute a whitening transform. Subsequent sets of time series could then be whitened with this transform. An adaptive procedure such as periodic reestimation of the time series ACF could be carried out for updating the whitening transform as the decorrelation time changes. Assuming a PRT of 1 kHz and the use of fifty 64-point (correlated) time series to train the whitener, 3–4 s of the ongoing data stream could be used for whitener "training," and a fraction of a second in computing the whitening transform (which, for 64 samples, involves the inversion of a 64×64 Hermitian–Toeplitz matrix). The cost would also depend on the rate at which the signal statistics (ACF) change, that is, how locally stationary the signal is (see the appendix for a corresponding sensitivity study). This additional computational burden appears rather small relative to the potential benefit.

To conclude, the question is whether the simulations carried out in this paper accurately reflect real-world conditions. We feel that given the range of ACF forms and correlation coefficients examined, the simplicity of the technique, and the potential to produce up to an order of magnitude improvement in accuracy for weather radars or other remote sensing instruments, an experimental study is justified.

Acknowledgments. This work was supported by National Science Foundation Grant ATM95-12685. ABK acknowledges sabbatical leave support from the USRA–NASA/Goddard Space Flight Center Laboratory for Atmospheres Fellowship program. We would like to thank I. Pinelis and T. Schulz and the three anonymous reviewers for many helpful comments.

APPENDIX

Sensitivity to Incorrect Whitener ACF

To provide an additional test of the method via comparison with Monakov's results (Monakov 1994) and to provide the worst-case scenario in which highly variable weather renders any whitening estimation impossible in real time, we provide the following sensitivity study in this appendix. Here we let the time series and whitener correlation coefficients (between two consecutive samples) differ widely in magnitude and examine the conditions where whitening might still prove useful.

We then come back to the question: How accurately does the r (correlation between consecutive samples) of the ACF have to be known in order for whitening to produce promising results? To that end, we generate noise-free time series with prescribed ACF and correlation parameter r and then whiten the series with a *different* r and compute the resulting error in estimated power. We report the results of computer simulations for the case where both the time series and whitener ACFs are of the same form. Specifically, we construct time series with exponential autocorrelation, that is, the time series autocorrelation is of the form $\rho_T(n) = r_T^n$. The average power in the unwhitened time series is unity in all cases. Whiteners are then constructed using autocorrelations $\rho_w(n) = r_w^n$, with correlation coefficients $r_w = 0.8$, 0.85, 0.9, 0.95, and 0.98. We expect whitening to be most effective when $r_w = r_T$.

Figure A1 illustrates the effect of whitening with varying r_w on the error in the power estimates. The horizontal axis is the correlation coefficient of the time series being whitened, r_T. The vertical axis is the mean-squared error (mse) of the estimated power for the whitened and unwhitened times series.

Time series of 64 points were used with up to 50 000

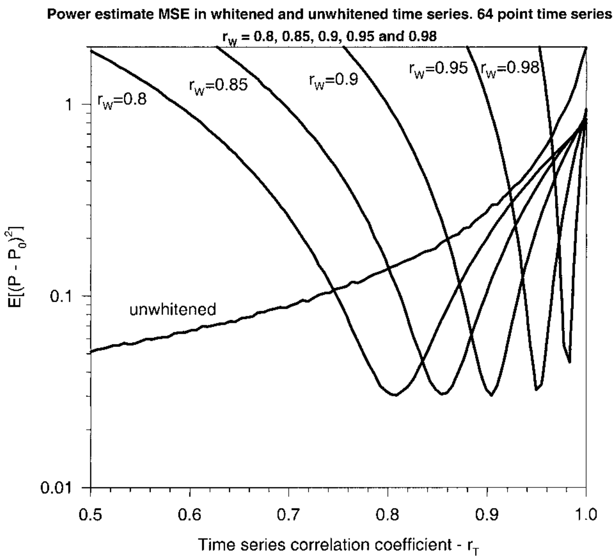

FIG. A1. Curves of power estimate mse in 64-point time series with and without whitening. Whitener and time series ACF both of form $\rho(n) = r^n$.

separate time series realizations in the mse sum. The "ripple" in the curves is due to the fact that finite numbers of time series realizations were used to compute the mean-squared error at each value of r_T. Increasing this number would reduce the ripple. Note that in the case of *unwhitened* time series, as r_T increases, the correlation between points of the time series increases and the independence of successive points in the time series decreases. This increases the mse in the power estimate and is apparent in the monotonic increase of the unwhitened mse with r_T. This is expected. The high error in power estimation associated with high correlation is important in meteorological applications where correlations can often be greater than 0.95. The five other curves of Fig. A1 plot the power estimate mse for whitened time series. The mse's for the whitened time series are minimum at $r_T \cong r_w$. The best whitener, in the mse sense, is expected to be the one derived from the time series ACF where $r_T \cong r_w$. Note in Fig. A1 that although the best power estimate occurs when the whitener correlation coefficient is approximately the same as the time series autocorrelation coefficient, the whitened power estimate is always better than the unwhitened power estimate when the whitener correlation coefficient is either less than, equal to, or just *slightly* greater than the time series correlation coefficient. Although we might expect the greatest improvement to occur when $r_T = r_W$, we see that whitened mse is minimum when r_W is *slightly less than* r_T. Monakov (1994) also reported similar results. When we estimate the average power in the time series without explicit whitening (i.e., with the standard sample mean), we are implicitly "whitening" because the sample mean is optimal for the case of independent (white) samples. This amounts to *always underestimating the correlation coefficient.* We pay a penalty if we overestimate the time series correlation coefficient, however. For example, if $r_T = 0.8$ and $r_W = 0.9$, the unwhitened power mse is about 0.14, but the whitened power mse is about 0.93.

REFERENCES

Atlas, D., 1964: *Advances in Radar Meteorology.* Vol. 10, *Advances in Geophysics,* Academic Press, 488 pp.

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

Janssen, L., and G. van der Spek, 1985: The shape of Doppler spectra from precipitation. *IEEE Trans. Aerosp. Electron. Syst.,* **21,** 208–219.

Johnson, G. E., 1994: Constructions of particular random processes. *Proc. IEEE,* **82,** 270–285.

Kostinski, A. B., 1994: Fluctuations of differential phase and radar measurements of precipitation. *J. Appl. Meteor.,* **33,** 1176–1181.

Marple, S. L., 1987: *Digital Spectral Analysis with Applications.* Prentice-Hall, 492 pp.

Meneghini, R., and T. Kozu, 1990: *Spaceborne Weather Radar.* Artech House, 199 pp.

Monakov, A., 1994: Estimation of the covariance matrix for dependent signal samples: Polarization diversity systems. *IEEE Trans. Aerosp. Electron. Syst.,* **30,** 484–492.

Nathanson, F. E., 1968: Radar precipitation echoes. *IEEE Trans. Aerosp. Electron. Syst.,* **4,** 505–514.

——, 1990: *Radar Design Principles.* 2d ed. McGraw-Hill, 720 pp.

Papoulis, A., 1984: *Probability, Random Variables, and Stochastic Processes.* McGraw-Hill, 576 pp.

Passarelli, R. E., Jr., and A. D. Siggia, 1983: The autocorrelation function and Doppler spectral moments. Geometric and asymptotic interpretations. *J. Climate Appl. Meteor.,* **22,** 1776–1787.

Priestley, M. B., 1981: *Spectral Analysis and Time Series.* Academic Press, 890 pp.

Sauvageot, H., 1992: *Radar Meteorology.* Artech House, 366 pp.

Schulz, T., and A. B. Kostinski, 1997: Variance bounds on the estimation of reflectivity and polarization parameters in radar meteorology. *IEEE Trans. Geosci. Remote Sens.,* **35,** 248–255.

Srivastava, R. C., A. R. Jameson, and P. H. Hildebrand, 1979: Time-domain computation of mean and variance of Doppler spectra. *J. Appl. Meteor.,* **18,** 189–194.

Ulaby, F. T., R. K. Moore, and A. K. Fung, 1986: *Microwave Remote Sensing: Active and Passive.* Vol. 2. Artech House, 608 pp.

Yaglom, A. M., 1987: *Correlation Theory of Stationary and Related Random Functions I.* Springer-Verlag, 526 pp.