© 1994 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

Variable-PRI Processing for Meteorologic Doppler Radars [1]

Edward S. Chornoboy and Mark E. Weber

MIT Lincoln Laboratory
Lexington, Massachusetts 02173

Abstract. In this communication we described how, with nonuniform sampling, the concept of bandlimited extrapolation can be used to obtain unambiguous Doppler velocity estimates in the supra-Nyquist region. The proposed method coherently processes a multi-PRI sample using a generalized form of periodogram analysis. The work is described in the context of meteorologic Doppler processing and includes a discussion of effective suppression for stationary ground clutter when multi-PRI schemes are used.

I. Introduction

A. Meteorologic Doppler Processing

Reflectivity and velocity estimation errors caused by range and Doppler ambiguities are a fundamental problem in weather radar. The use of variable pulse-repetition intervals (PRIs) to support resolution of velocity ambiguities has been discussed in the meteorologic context ([1], for example), but dedicated weather radars typically settle on constant-PRI waveforms to facilitate both ground-clutter suppression and coherent processing across the entire sample. The latter issue is important because these radars typically scan at high rates and multiple elevation angles. Hence, Doppler estimation is limited to coherent sample vectors on the order of only 20-40 pulses, which can be a significant limitation given the lower returned power and distributed nature of meteorologic events. Since weather radars often estimate weather-echo spectral moments using a single-lag autocorrelation method (e.g., "pulse pair"), splitting the available "time on target" among multiple PRIs, a reduction of coherent averaging, represents an additional compromise. Although periodogram based methods are sometimes used to estimate weather spectral moments, there currently is no effort in the meteorologic community directed at spectral estimation using multi-PRI signals.

[1] The work described has been sponsored by the Federal Aviation Administration. The U.S. Government assumes no liability for its contents or use thereof.

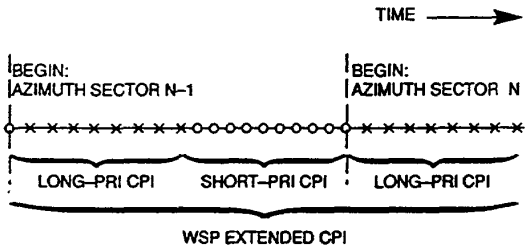

Figure 1: ASR-9 block-staggered waveform.

B. The ASR-9 Wind Shear Processor

Our work is motivated by the following application – estimation of weather spectrum moments using a Wind Shear Processor (WSP) modification to the Federal Aviation Administration's (FAA) Airport Surveillance Radar (ASR-9). Lincoln Laboratory, under FAA sponsorship, has developed a prototype WSP for the testing of candidate low-altitude radial-wind algorithms. These algorithms require both aggressive suppression of stationary ground clutter and spectral moment estimation that accounts for the possibility of velocity ambiguity. The requirements are difficult to meet because the ASR-9's characteristics include a 12.5 RPM scan rate and a block-staggered sampling scheme. Figure 1 illustrates the ASR-9 variable-PRI waveform, which is employed to mitigate "blind" speeds for aircraft targets. In the time it takes to scan one beamwidth, the ASR-9 transmits two coherent processing intervals (CPIs): one eight-pulse long-PRI (~1 msec) block followed by one ten-pulse short-PRI block. The ratio of PRIs is 7:9, and the associated Nyquist velocity for this S-band radar is only 26 m/s corresponding to the long-PRI rate. In order to obtain a sufficient number of samples for clutter suppression and weather-Doppler estimation, the WSP operates coherently across three of the individual pulse blocks, resulting in the "8-10-8" sampling pattern shown in Fig. 1. This "extended" CPI spans 27 pulses and is the longest deterministic waveform available – owing to the use of fill pulses which compensate for wind loading while ensuring antenna registration in azimuth.

0-7803-1438-7/94/$3.00 © 1994 IEEE

Authorized licensed use limited to: MIT Libraries. Downloaded on April 26,2010 at 20:53:54 UTC from IEEE Xplore. Restrictions apply.

C. Remainder Theorem Dealiasing

When multi-PRI is used to extend the range of unambiguous Doppler, computations based on the so-called Chinese Remainder Theorem are usually employed. In such cases, measurements corresponding to each PRI are pooled and processed to obtain single-lag velocity estimates. In the case of the ASR-9 this produces two estimates: $\hat{V}_L$ and $\hat{V}_H$. Fig. 2 illustrates the type of logic that is relied upon. The top panel shows the folding behavior of low and high velocity estimates. The middle panel illustrates a form of correction logic based on the difference between the estimates. The mapping in the middle panel — between V_{True} and the ideal states of $\Delta V = V_H - V_L$ — is unique out to 9 times V_L (7 times V_H). In principle, by testing the difference velocity, one can infer the proper correction for one or both estimated velocities. The performance of such a scheme, in the context of the ASR-9, is illustrated in the bottom panel which plots the percentage of correct (POC) decisions (i.e., dealiasing corrections) made in a Monte Carlo experiment consisting of 5,000 realizations. This plot is for the case where the signal has a Gaussian-shaped power spectrum, the spectral spread is 2 m/s, measurement noise is white, and the signal-to-noise (SNR) ratio is 10 dB. Clearly, for this example, the performance of this method is quite erratic beyond the long-PRI Nyquist bound.

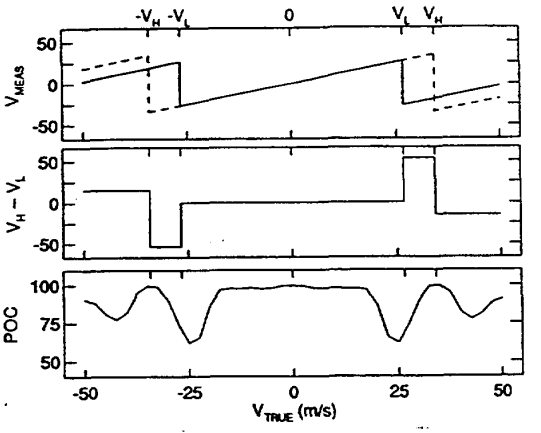

Figure 2: Remainder Theorem dealiasing and the ASR-9.

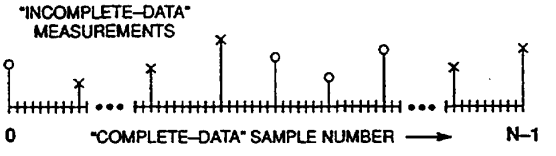

Figure 3: ASR-9 data as an incomplete sampling: symbols "×" and "o" represent locations of long- and short-PRI samples.

II. Bandlimited Extrapolation and Dealiasing

A. Bandlimited Extrapolation

A signal y is bandlimited (BL) if its Fourier transform $Y(\omega)$ is zero for $|\omega| > \omega_c$. Band-limited extrapolation refers to a class of problems wherein one attempts to make use of this specific prior information to extend a signal beyond the original measurement interval. The discrete case is treated nicely in [2]; the essential components are as follows. A finite, or "incomplete", measurement vector z results from application of a sampling matrix $\mathcal{S}$ to a "complete" data vector sample y:

$$z = \mathcal{S}y \,. \tag{1}$$

Because y is bandlimited, there exists an appropriate ideal low-pass filter matrix L such that

$$y = Ly \,. \tag{2}$$

The operator L is a projection operator for the known band of frequencies; L is both idempotent and self adjoint. Estimation of y from measured z is underdetermined, but with constraint L and an appeal to the minimum-norm least-squares principle, one may select as solution

$$\hat{y} = L\mathcal{S}^T(\mathcal{S}L\mathcal{S}^T)^{-1}z \tag{3}$$

because this has minimum norm among all feasible reconstructions. Interpolation to missing measurement values is an equally valid interpretation of the above, and this is the focal point for our variable-PRI application.

B. BL Interpolation with Nonuniform Sampling

If, as in the case of the ASR-9, the non-uniform sampling of multi-PRI occurs at integer multiples of some fundamental sampling interval, then the measured radar returns can be viewed as an incomplete measurement from a much finer uniform sampling grid. This situation is illustrated in Fig. 3 for the case of the ASR-9. Here, the uniform sampling grid is at a rate 9 times higher than the long-PRI sampling rate and would have a corresponding Nyquist velocity of 234 m/s.

Authorized licensed use limited to: MIT Libraries. Downloaded on April 26,2010 at 20:53:54 UTC from IEEE Xplore. Restrictions apply.

We can take y to be the finite uniformly sampled vector of length N_y, whereas the incomplete sample z has length $N_z < N_y$. For the ASR-9 example, $N_y = 215$ and $N_z = 27$. Filter L can be decomposed using a discrete Fourier transform (DFT) D so that [2]

$$L = D^\dagger \Lambda D \,, \tag{4}$$

where Λ is a diagonal matrix containing the filter weights as its entries; for an ideal filter, these equal either zero or one. Solution (3) now has the form

$$\hat{y} = D^\dagger \Lambda D \mathcal{S}^T (\mathcal{S} D^\dagger \Lambda D \mathcal{S}^T)^{-1} z \,; \tag{5}$$

or, since we are potentially interested in frequency-domain computation, we can work with the DFT of $\hat{y}$:

$$\hat{Y} = D\hat{y} = \Lambda D \mathcal{S}^T (\mathcal{S} D^\dagger \Lambda D \mathcal{S}^T)^{-1} z \,. \tag{6}$$

The $N_y \times N_z$ matrix $D\mathcal{S}^T$ is the submatrix obtained from the DFT matrix by taking the N_z columns corresponding to the actual data sample times. Equation 6 also indicates how one obtains a periodogram spectral estimate from nonuniformly sampled data.

C. BL Interpolation and Dealiasing

Suppose that y is bandlimited such that the total support of its frequency content is confined to a region of length $\sigma \ll \omega_c$ but that *a priori* we do not know its location. The problem of identifying a most likely location was considered in [3], where an iterative scheme was proposed in the context of BL extrapolation with unknown sinusoidal components. The method begins by taking $\Lambda = \Lambda_0$ such that all $\omega < \omega_c$ are likely, and one estimates $\hat{Y}$ according to (6). The resulting estimate $\hat{Y}$ is then thresholded in magnitude to eliminate low-energy components and this result is used to collapse the support of Λ_0 to produce Λ_1, where $[\Lambda_1]_{kk} = 1$ if $|\hat{Y}_k| > \epsilon$ and is zero otherwise. In this sense, the method is similar to other forms of signal-subspace estimators.

Good empirical results were reported in [3]. How does the idea hold up with non-sinusoidal signals? Figure 4 is an ASR-9 simulation that illustrates the method. The bottom-most panel plots the source power spectral density for a simulated ASR-9 signal centered at the supra-Nyquist velocity of 35 m/s with spectral width 2 m/s. A realization simulating a 10 dB SNR was analyzed as

[2]Here, we use "†" to indicate conjugate transpose.

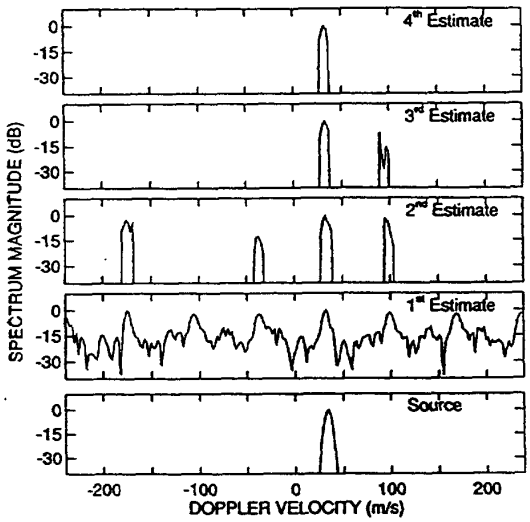

Figure 4: Bandlimited dealiasing of an ASR-9 signal.

described above, and the results are summarized in the upper four panels of Fig. 4. From the bottom up, the first minimum norm estimate considers all frequencies in the range ± 234 m/s equally likely. Multiple peaks due to aliasing are clearly evident; however, nonuniform sampling introduces significant distortion such that one region can be distinguished as more likely than the others. The top three panels illustrate that successive contraction of the support of Λ, based solely on power thresholding, correctly identifies the target interval and provides an unambiguous estimate for this example. Surprisingly, the method holds up quite well over the full range of Doppler velocities shown in the figure.

III. Weather Radar Dealiasing

A. Testing for Spectral Support

For Fig. 4, ω_c was taken to be the Nyquist velocity corresponding to the uniform sampling rate — a value much higher than needed for most meteorologic applications. A more practical version might only consider Doppler velocities in the range ± 50 m/s, as initially shown for the Remainder Theorem dealiasing in Fig. 2. When this is done it becomes apparent that there is information, other than peak power, which can be used to distinguish candidate support for meteorologic targets. Figure 5 shows an analysis for a simulation as in Fig. 4 but with the reduced ω_c of

Authorized licensed use limited to: MIT Libraries. Downloaded on April 26,2010 at 20:53:54 UTC from IEEE Xplore. Restrictions apply.

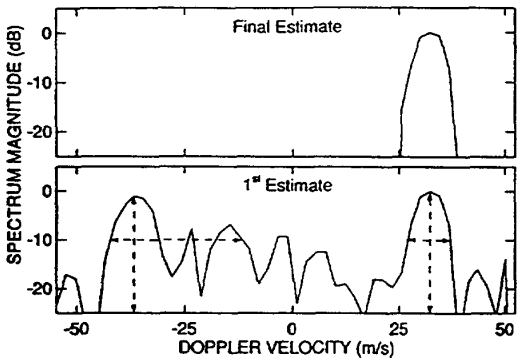

Figure 5: Peak power and width statistics can serve to identify aliased components: aliased energy exhibits decreased peaks and local whitening.

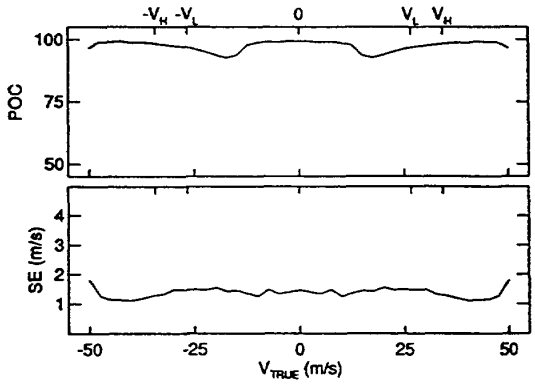

Figure 6: Probability of correct decision and standard error results for a spectrum width of 2 m/s and SNR of 10 dB.

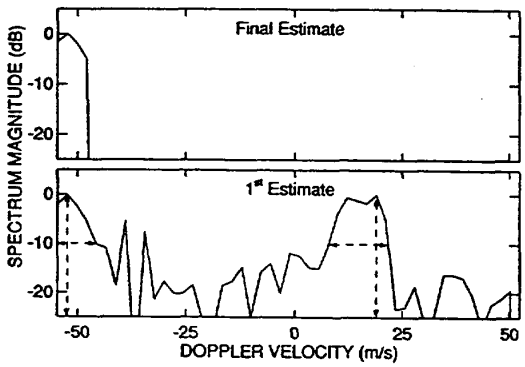

Figure 7: For the ASR-9, a true V near 17 m/s has aliased energy near -50 m/s which, when partially cut-off, can lead to erroneous decisions.

50 m/s. As before, the true source location is the supra-Nyquist 35 m/s and the main competition is from aliased energy appearing around -35 m/s. However, aliased energy generally appears as a "whitened" version of the true source due to a mismatch with the phase encoding of the known nonuniform sampling. This suggests that a statistic based on a ratio of local peak power to spectral spread can be used to collapse the support of A. We therefore investigated the following modification to the method proposed in [3]. In the bottom panel of Fig. 5 horizontal arrows indicate, first, that a threshold test at -10 dB was used to identify significant energy levels. Because we assume our source spectrum to be distributed, we extend the support at the -10 dB threshold (i.e., fill gaps) to ensure that support intervals have at least a minimum length; in the present case, the minimum support length was taken to be nine contiguous DFT coefficients. The horizontal arrows in Fig. 5 therefore indicate the energy extent over two likely support regions after one iteration. Local spectral moments over individual support segments were computed and the ratio of peak local power (vertical arrows in Fig. 5) to local spread was used to select the most likely support interval. In this way, the method was used to arrive at a decision after one iteration, and a second iteration could be used to provided a final spectral estimate.

B. Performance Results

The feasibility of the above decision/estimation procedure was explored using a Monte Carlo analysis. Spectrum widths in the range 1-4 m/s and SNR ratios between 0-20 dB were examined. All performance estimates were computed from 5,000 experiment realizations. The results for a spectrum width of 2 m/s and SNR of 10 dB are shown in Fig. 6. The lower panel plots the estimated standard error (SE) for the spectral mean (i.e., velocity estimate) conditioned on the event that the correct Nyquist interval had been selected[3]. The upper-panel is the corresponding POC performance curve, which compares very favorably with the equivalent POC curve for Residue Theorem dealiasing, shown in Fig. 2. Generally the POC values in Fig. 6 are at 98% or above. However, the POC curve does dip to 93% at a velocity of ~17.5 m/s. The reason for this is shown in Fig. 7 and points out a shortcoming in the described decision/estimation procedure. Signal energy near 17.5 m/s is aliased near -50 m/s, which was the

[3]Here, a correct Nyquist interval decision means that the estimated spectral mean was within ±26 m/s of the true Doppler velocity.

Authorized licensed use limited to: MIT Libraries. Downloaded on April 26,2010 at 20:53:54 UTC from IEEE Xplore. Restrictions apply.

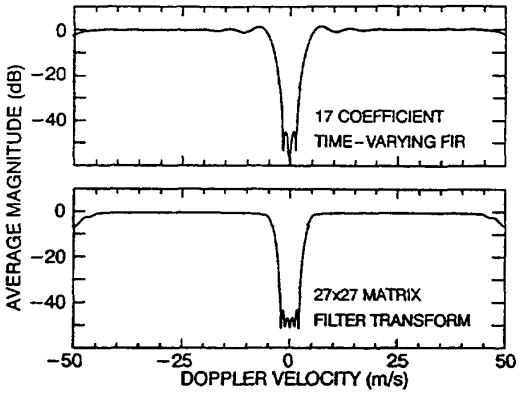

Figure 8: High-pass clutter filters for use with multi-PRI signals.

(arbitrarily) chosen cutoff for spectral estimation. The artificial cutting of the aliased energy has an obvious negative effect on our statistic for selecting spectral support by spectral spread. While Fig. 6 shows the method to be promising, Fig. 7 shows that there are still problems that must be dealt with. Overall, within the range of spectral widths and SNR values examined, POCs of better than 90% appear possible.

IV. Suppression of Stationary Ground Clutter

A. Options

As implied in the introduction, the suppression of stationary clutter is a major concern for weather radars that measure horizontal wind fields near ground level. There are, at least, three options that can be used in the context of the above spectral processing; they are as follows.

Bandlimited Estimation: In view of the discussion of Section II, it is clear that if the clutter spectrum is centered at zero and confined to frequencies $|\omega| < \omega_{clutter}$, one can appropriately modify the filter weighting function Λ so that estimation of signals only within the range $\omega_{clutter} < |\omega| < \omega_c$ is enabled. This was considered and discussed in [4], for example, in the context of BL-extrapolation spectral estimation. This could be sufficient if zero-Doppler energy does not alias into any other velocity region of interest; for the ± 50 m/s examples considered here, it doesn't.

High-Pass Filtering: To be free of the effects of aliased ground-clutter returns, it may be necessary to preprocess the data with a high-pass filter. Because the samples are no longer uniformly spaced, this implies the need for time-varying filters in order to preserve the "phase encoding" of that sampling. Recent work by one of us [5] examined design techniques for time-varying finite impulse response (FIR) filters that achieve Chebyshev or mean-squared (MS) optimality while preserving the sampling information of the multi-PRI waveform. Ignoring for the moment that the ASR-9 uses fill pulses to maintain antenna alignment, it would require a periodically time-varying structure with period 18 to filter the ASR-9 WSP data (see Fig. 1). The average magnitude response[4] for such a 17-coefficient MS-design time-varying FIR filter is shown in Fig. 8 (upper panel).

Matrix Transform: In a slight deviation from traditional filter forms, one could also consider the design of a matrix transform H, applied to data z, as a means of filtering. Clearly FIR and IIR structures can be implemented as matrix transforms, but in [5] we examined the complex domain design of the matrix transform directly in an attempt to benefit from any added degrees of freedom. The average magnitude response for one of these transforms is shown in the lower panel of Fig. 8. The potential of this method is that filter H can be concatenated with the operators in Equations 5 and 6:

$$\hat{y} = D^{\dagger}\Lambda D\mathcal{S}^{T}(\mathcal{S}D^{\dagger}\Lambda D\mathcal{S}^{T})^{-1}Hz \qquad (7)$$

and

$$\hat{Y} = \Lambda D\mathcal{S}^{T}(\mathcal{S}D^{\dagger}\Lambda D\mathcal{S}^{T})^{-1}Hz\,. \qquad (8)$$

This invites an alternative idea to the iterative method described in Section III. To the extent that a set of predefined masks Λ_k ($k = 0, \ldots, K-1$) can be specified, the filtering operation can be embedded in the computations of the frequency-domain mapping. For example, with the above ASR-9 example, one could predefine operators based on the overlapping support regions [-50, 0], [-25, 25], and [25, 50] m/s. The dealiasing problem then becomes a problem of making the *a posteriori* decision of best model fit. Figure 5 indicates that such a strategy might also succeed, at least in the context of our ASR-9 application.

B. Effect on Dealiasing

The experiment of Section III-B was repeated using the matrix filter of Fig. 8. For the exam-

[4] The magnitude response to sampled complex sinusoids was averaged over filter number.

Authorized licensed use limited to: MIT Libraries. Downloaded on April 26,2010 at 20:53:54 UTC from IEEE Xplore. Restrictions apply.

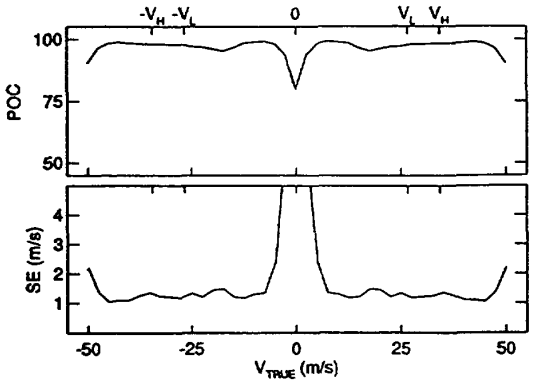

Figure 9: Probability of correct decision and standard error results for filtered ASR-9 data. Signal spectrum width is 2 m/s, SNR is 10 dB, and the lower-panel filter of Fig. 8 was used.

ple shown, the simulation used a signal spectrum width of 2 m/s, an SNR of 10 dB, a clutter to signal ratio (CSR) of 10 dB, and a clutter (also Gaussian shaped) spectral width of 0.75 m/s. The corresponding SE and POC estimates are shown in Fig. 9. Overall, the POC results are slightly better than those of Fig. 6 because the filter has a 5-6 dB dip at ± 50 m/s, which alleviates the correction errors characterized by Fig. 7. This dip in the filter pass-band response is a consequence of the ASR-9 block-staggered spacing. Applications that would allow greater freedom in construction of the variable-PRI sampling scheme have the potential to improve upon this limitation.

V. Summary

The ASR-9 WSP development has led us to attempt some new and novel processing methods with respect to ground-clutter suppression, Doppler velocity estimation, and dealiasing. We here propose a processing scheme which can be viewed as a generalized form of periodogram spectral analysis. Although the method was described as an adaptive form, it is also possible to cast these ideas into a form where multiple precomputed transforms provide a fit for preselected frequency models. The "encoding" of a multi-PRI waveform coupled with an *a posteriori* test, applied to the resulting signal spectrum estimates, could be used to select the frequency interval (model) that best describes the data. These ideas were demonstrated by Monte Carlo analysis, and the resulting performance was shown to be an improvement over a more conventional multi-PRI dealiasing technique. Hence, this processing approach – when used with suitable transmitted waveforms – can provide reliable, unambiguous Doppler spectrum estimates which can be used directly for weather echo moment calculations. Because the method processes an entire data sample coherently, there is no sacrifice of data integration to remove Doppler ambiguity.

To be an effective method, the multi-PRI technique must also be compatible with methods that suppress radar returns from stationary ground clutter. We have implemented algorithms that perform complex domain design and that allow arbitrary specification of the filter group delay. We have successfully designed block-oriented "matrix" filters, demonstrated for the ASR-9 27-pulse sample; and more traditional time-varying FIR filters, such as might be used for Moving Target Indicator (MTI) processing. We have successfully demonstrated the use of the block-oriented filtering scheme with time-domain Doppler processing and the multi-PRI dealiasing method described here.

References

[1] Z. Banjanin and D.S. Zrnić, "Clutter rejection for Doppler weather radars which use staggered pulses," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 29(4), pp. 610–620, 1991.

[2] A.K. Jain and S. Ranganath, "Extrapolation algorithms for discrete signals with application in spectral estimation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-29(4), pp. 830–845, 1981.

[3] A. Papoulis and C. Chamzas, "Detection of hidden periodicities by adaptive extrapolation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-27(5), pp. 492–500, 1979.

[4] C.L. Byrne and R.M. Fitzgerald, "Spectral estimators that extend the maximum entropy and maximum likelihood methods," *SIAM J. Appl. Math.*, vol. 44(2), pp. 425–442, 1984.

[5] E.S. Chornoboy, "FIR design methods for non-uniformly sampled signals," MIT Lincoln Laboratory, Lexington, Massachusetts, Technical Report 976, 1993.

Authorized licensed use limited to: MIT Libraries. Downloaded on April 26,2010 at 20:53:54 UTC from IEEE Xplore. Restrictions apply.