


**Project Report
ATC-191**

Optimum Time-Varying FIR Filter Designs for the Airport Surveillance Radar Wind Shear Processor

E. S. Chornoboy

28 December 1994

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS



Prepared for the Federal Aviation Administration,
Washington, D.C. 20591

This document is available to the public through
the National Technical Information Service,
Springfield, VA 22161

This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof.

1. Report No. ATC-191	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle Optimum Time-Varying FIR Filter Designs for the Airport Surveillance Radar Wind Shear Processor		5. Report Date 28 December 1994	
		6. Performing Organization Code	
7. Author(s) Edward S. Chornoboy		8. Performing Organization Report No. ATC-191	
9. Performing Organization Name and Address Lincoln Laboratory, MIT P.O. Box 9108 Lexington, MA 02173-9108		10. Work Unit No. (TRAIS)	
		11. Contract or Grant No. F19628-95-C-0002	
12. Sponsoring Agency Name and Address Department of Transportation Federal Aviation Administration Systems Research and Development Service Washington, DC 20591		13. Type of Report and Period Covered Project Report	
		14. Sponsoring Agency Code	
15. Supplementary Notes The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology under Air Force Contract F19628-95-C-0002.			
16. Abstract We have developed new design algorithms for finite impulse response (FIR) filters that compensate for arbitrary input spacing and that allow for arbitrary group delay specification. The potential of these new designs to work with the ASR-9 staggered pulse spacing is examined in the context of the ASR-9 wind-shear processor (WSP). Benefits derived from the new designs include an improved (optimal) stop-band design, an increased yield in pulse samples for moments estimation, and the retention of pulse-stagger phase information, which can be used for velocity dealiasing. These improvements are demonstrated using simulated and test-bed data, the latter acquired during 1991/1992 Orlando operations. Filter utilization, in the context of a pre-existing adaptive selection scheme [1] and the Orlando (FL) clutter environment, is examined using the new filters, and areas for improvement are identified.			
17. Key Words finite impulse response wind-shear processor		18. Distribution Statement This document is available to the public through the National Technical Information Service, Springfield, VA 22161.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 59	22. Price

ABSTRACT

We have developed new design algorithms for finite impulse response (FIR) filters that compensate for arbitrary input spacing and that allow for arbitrary group delay specification. The potential of these new designs to work with the ASR-9 staggered pulse spacing is examined in the context of the ASR-9 wind-shear processor (WSP).

Benefits derived from the new designs include an improved (optimal) stop-band design, an increased yield in pulse samples for moments estimation, and the retention of pulse-stagger phase information, which can be used for velocity dealiasing. These improvements are demonstrated using simulated and test-bed data, the latter acquired during 1991/1992 Orlando operations.

Filter utilization, in the context of a pre-existing adaptive selection scheme [1] and the Orlando (FL) clutter environment, is examined using the new filters, and areas for improvement are identified.

TABLE OF CONTENTS

Section	Page
Abstract	iii
List of Illustrations	vii
1. INTRODUCTION	1
2. GROUND CLUTTER PROCESSING THEORY	3
2.1 Clutter Filter Design	3
2.2 Clutter-Induced Weather-Parameter Estimate Biases and Filter Selection	7
3. OBSERVED CLUTTER FILTER PERFORMANCE	13
3.1 Filter Utilization: Orlando Case Study	13
3.2 Preservation of Pulse Stagger: Velocity Dealiasing Example	25
4. CONCLUSIONS	35
APPENDIX A - TIME-VARYING FILTER RESPONSE PROFILES	37
REFERENCES	51

LIST OF ILLUSTRATIONS

Figure No.		Page
1.	Three levels of suppression used in the current WSP adaptive scheme.	5
2.	New filter designs to fit the WSP adaptive scheme.	6
3.	Response characterizations for standard FIR designs with interpolation: 0 dB signal-to-clutter power ratio.	9
4.	Response characterization for standard FIR designs with interpolation: -30 dB signal-to-clutter power ratio.	10
5.	Response characterizations for optimum staggered FIR designs: 0 dB signal-to-clutter power ratio.	11
6.	Response characterizations for optimum staggered FIR designs: -30 dB signal-to-clutter power ratio.	12
7.	1991 Orlando clutter environment: computed dBZ reflectivity map.	13
8.	1991 Orlando clutter environment: computed CNR histograms	15
9.	Filter usage by range: August 32, 1991.	17
10.	Single-gate Doppler spectrum example I.	18
11.	Scatter plot of input vs. output estimated CNR for filter “failures.”	19
12.	Single-gate Doppler spectrum example II.	20
13.	Ideal filter usage map: Orlando (FL).	21
14.	Clutter residue after “ideal” filtering.	23
15.	Adaptive filter usage given low-power weather returns.	25
16.	Clutter residue given low-power weather returns and after adaptive filter selection.	27
17.	Pulse-pair velocity estimates with 17-coefficient FIR filters.	29
18.	Pulse-pair velocity estimates with new time-varying filters.	31

LIST OF ILLUSTRATIONS
(Continued)

Figure No.	Page
19. Velocity dealiasing combined with clutter filtering.	33
A-1. MSE Design #1 (20 dB): frequency magnitude and phase error response.	38
A-2 MSE Design #2 (40 dB): frequency magnitude and phase error response.	42
A-3 MSE Design #3 (60 dB): frequency magnitude and phase error response.	46

1. INTRODUCTION

In Project Report ATC-143 [1], Weber provided an analysis of ground clutter rejection requirements for low-altitude wind measurements using an airport surveillance radar (ASR). A number of contributing factors make this a difficult and challenging problem. For example, the high scan rate of an ASR results in a scan-modulated clutter spectrum comparable in width to valid weather. The low number of pulses per azimuth radial and the use of a block-staggered signal design severely limit the stop- and transition-band widths obtainable with conventional (FIR) filter designs. Furthermore, as the effects of filter transients on Doppler parameter estimation are not well characterized, the use of FIR prefilters typically results in the decision to discard an amount of data equivalent to the length of the filter; in the case of an ASR, this would amount to over half of all I/Q data. The use of a single prefilter matched to the strongest clutter source is clearly overly aggressive. A bank of prefilters “matched” to cover a range of clutter sources is better but still aggressive because the crude filter notches will often remove a substantial amount of power corresponding to weather echo. Weber argued that in situations where the weather return power dominates any clutter component (10 dB), the prudent strategy is to not filter the clutter because its biasing effect is negligible. Hence, an adaptive filtering method was proposed whereby *a priori* clutter information (residue maps) would be used to balance the removal of clutter against the measured weather content. Weber demonstrated the design of a three-filter (plus all-pass) system taking into consideration the limiting parameters of the testbed radar in use at that time. In addition, he documented the time-varying characteristics of ground clutter as seen by an ASR and developed an appropriate threshold strategy that takes this into consideration.

Two developments, in the time since [1], motivate the present report:

1. in [2], new design algorithms for optimal¹ time-varying filters were developed, and
2. initial success of the ASR-WSP in wet environments prompts the effort to extend its capabilities to dry environments and, in general, improve performance with weak-signal events, such as leading-edge gust fronts.

The new filter design algorithms of [2] are significant in that they take into consideration the pulse stagger spacing *during* design and permit arbitrary specification of the filter group delay. The potential benefits resulting from this include the design of (time-varying FIR) filters characterized by

- greater control over stopband width,
- sharper transition regions, and
- retention of all I/Q data for moments estimation.

¹Chebyshev and Mean-Squared Error.

In [1], the clear emphasis was on the balancing of clutter suppression in the face of significant weather-echo returns. The new clutter filters address improved low-signal strength estimation by minimizing the data loss when prefiltering is required. As an addendum to [1], we would like to emphasize those aspects of Weber's adaptive strategy that are important for low-signal strength estimation.

This report evaluates the new filters in two ways: one, it repeats the suppression analysis first considered in [1]; and two, it examines the "optimal" matching of filters to the clutter environment when weather-echo returns do not dominate. The threshold test described in [1] did not explicitly emphasize the need for limiting the "maximum" filter used at any range/azimuth location (based on prior observation); we do so here. One new benefit derived from these filters was not listed above. That is the ability to use dual-PRF velocity dealiasing methods with filtered weather-channel data. We intend this as the subject for a future report but include one example here because it is the new approach to prefiltering that enables the capability.

2. GROUND CLUTTER PROCESSING THEORY

2.1 Clutter Filter Design

The WSP adaptive method for obtaining “sufficient, but not excessive” clutter attenuation was previously discussed in [1]. For this method, three filters (plus one all-pass option) are designed to provide successive levels of attenuation centered at zero Doppler. Both depth of attenuation and filter width are adjusted to provide approximate (inverse) “matches” to expected scan-modulated clutter spectra at three levels of returned power. Pre-computed “clear-weather” maps of post-filtering clutter residue are used to indicate the expected effectiveness of each filter at each range/azimuth location. Although, on the one hand, the objective is to match the filter notch to the clutter spectral component, less attenuation can be used if the “weather signal” component is sufficiently large to render clutter-induced bias negligible². A comparison between filter-output power and (expected) clutter-residue power (for each filter/gate pair) is the basis for the adaptive “minimum filtering” strategy (see [1]). Central to the analysis and discussion in [1], therefore, is the trade-off between clutter suppression and weather-parameter estimation, as characterized by computing (theoretical) velocity and reflectivity measurement biases, indexed by filter type and weather/clutter content.

The filters examined in [1] are shown in Fig. 1, as characterized by their “average” frequency response measures. These filters are “equally spaced” at levels of suppression between 0 and 45 dB (15.5, 31.2, and 45.3 dB) with 3 dB bandwidths of 2.8, 4.7, and 6.1 m/s. These filters were derived from standard (equispaced) Chebyshev design. At the time of their design, it was determined that the instability residue of the ASR-8 transmitter then in use would make useless any effort at an attenuation greater than 45 dB.

The filters examined in this report were obtained by optimizing an MSE design criterion, as discussed in [2]. All designs are for a data block showing an 8-10-8 (low-high-low, 7:9:7 PRF ratio) pulse stagger, as discussed in [1] and [2]. Time-averaged frequency response profiles for these filters are shown in Fig. 2. Detailed (individual pulse) frequency-response characterizations are also provided in the accompanying appendix, Figs. A-1—A-3. The designs were formulated around the “ideal” filter responses indicated by dotted lines in the figures; if desired, more complicated response profiles – such as an attempt to exactly match the falloff of the scanned antenna pattern – could be supported by the new design algorithms. Anticipating that improvement in the instability residue would be achieved, we designed new filters at equally spaced intervals between 0 and 60 dB suppression (20, 40, and 60 dB). Filter stop-band widths were adjusted in accordance with stop-band suppression to provide 3 dB bandwidths of 2.0, 3.2, and 4.5 m/s. In Fig. 2, both 3 and 10 dB notch widths are recorded, as well as the notch width requested at design (again, the ideal filter

²The specifics of this consideration of course depend on the method used to estimate the spectral moment.

response is indicated by a dotted line). Transition regions sharper than those shown are possible at the expense of greater pass-band ripple. As discussed in [2], arbitrary control of the filter phase response (toward the linear-phase ideal) is also possible at the expense of magnitude response. For the present designs, we opted for modest phase-response control (phase errors less than 5% Nyquist, see response profiles in Appendix) to increase the yield of usable output pulses. With the current WSP method, processing by any of the three filters yields only 11 output samples (from a 27-pulse input block). The number of output pulses that can be retained with the new method is variable depending on the depth of suppression: 27 for “all pass” and the 20 dB filter, and 23 otherwise. These yields were selected based on examination of the filter response characterizations shown in Figs. A-1-A-3. The notch widths and output block size of the new filters all represent substantial improvements relative to the previous processing design.




Figure 1. Three levels of suppression used in the current WSP adaptive scheme (time-averaged responses taken from [1], Fig. II-4).




Figure 2. New filter designs to fit the WSP adaptive scheme (time-averaged responses obtained from appendix responses; 26 m/s corresponds to the Nyquist velocity in this case).

2.2 Clutter-Induced Weather-Parameter Estimate Biases and Filter Selection

Following the convention set in [1], we computed ideal power spectra corresponding to “typical” weather and clutter signals. Both spectral components were Gaussian shaped: the clutter component was given a spectral width of 0.75 m/s (assuming a Nyquist interval of 26 m/s) and the weather component was given a spectral width of 2.0 m/s. Two combinations of weather and clutter were considered: a 0 dB signal-to-clutter power ratio and a -30 dB ratio. Ideal composite spectra were multiplied by the magnitude “frequency response” curves of the time-varying filters. Because the frequency response curves are different for each of the filter output pulses, this process was done for each response profile, and individual errors were calculated and averaged over all the output pulses for a given filter implementation. In contrast to [1], however, calculation of a velocity estimate was not based on a pulse-pair computation. Instead, spectral means were computed directly from the synthesized (and filter-masked) power spectra. For the type of analysis considered in [1] and here, this discrepancy should not matter.

For comparison, we begin with a presentation of filter responses based on FIR filters similar to those presented in [1]. Figures 3 and 4 “repeat” the characterizations shown in Figs. II-5(a) and (b) of [1]. However, note the modification that Figs. 3 and 4 cover an extended velocity interval of $\pm 2\pi$, corresponding to ± 56 m/s. For $|\omega| > \pi$, these plots “ignore” the issue of velocity folding, since the issue here is clutter-residue power vs. weather-return power. The filters of Figs. 3 and 4 were the test-bed filters in use as of this report date. Some differences in magnitude responses were observed in comparison to Figure II-4 of [1], which reflect design modifications since the reporting date of [1]. For example, the filters examined in [1] had 3 dB pass-band edges of 2.8, 4.7, and 6.1 m/s while the test-bed filters shown here had computed values of 2.8, 5.5, and 7.3 m/s (assuming a 26 m/s Nyquist velocity). Since these filter designs only provide 11 usable output pulses, the responses shown are the average of computations for those 11 output samples only. As stated above, velocities were obtained by computing the mean of the composite spectral density, and “Average Velocity Bias” represents the average of the computed velocity bias corresponding to the response profiles for each of the 11 sampled output pulses. Reflectivity “bias” represents the ratio of power in the weather component to the power in the filtered (signal+clutter) output. The plotted average is a geometric average for the 11 output pulses. The last plotted quantity is based on Equation 5 of [1],

$$\begin{aligned} R &= \frac{\text{WeatherPower} + \text{ClutterResidue}}{\text{ClutterResidue}} \\ &= \frac{\text{WeatherPower}}{\text{ClutterResidue}} + 1, \end{aligned} \tag{1}$$

and plots the (arithmetic) average of the measure R , which is then expressed in dB. Although numeric differences are evident (the studied filters are not exactly the same), Figs. 3 and 4 clearly verify the trends and conclusions of [1]. In Fig. 3, the plot of “Average Velocity Bias” in particular illustrates why it is desirable to keep the applied suppression as low as possible: proper filter matching minimizes velocity and reflectivity biases when weather spectra overlap the filter’s stop band. Fig. 4 implies that the test threshold for R lies between 10 and 20 dB, although probably closer to

10 dB. The testbed currently operates with an R threshold of 10.4 dB. The most attenuating filter is required for operation with a -30 dB signal-to-clutter power ratio.

A similar characterization for the new filters is presented in Figs. 5 and 6. For this analysis, averages over 23 output samples were used for all filters. In Fig. 6, the widest filter is still required to suppress the 30 dB clutter component, although the 40 dB filter is just borderline from being sufficient. The new filters confine the region of significant reflectivity or velocity bias to smaller intervals about zero; this does not require a corresponding compromise in any other aspect of the performance characterization.




Figure 3. Response characterizations for standard FIR designs with interpolation: 0 dB signal-to-clutter power ratio.




Figure 4. Response characterizations for standard FIR designs with interpolation: -30 dB signal-to-clutter power ratio.




Figure 5. Response characterizations for optimum staggered FIR designs: 0 dB signal-to-clutter power ratio.




Figure 6. Response characterizations for optimum staggered FIR designs: -30 dB signal-to-clutter power ratio.

3. OBSERVED CLUTTER FILTER PERFORMANCE

3.1 Filter Utilization: Orlando Case Study

In examining filter usage, one can ask (at least) two questions. One, how well do the designed filters “match” the expected clutter environment; and two, how well does the adaptive scheme utilize the filters in an operational setting? This section examines the new filters by looking at aspects of both questions. As Weber ([1]) included an analysis on temporal clutter variation, we will not consider that here.

The first data examined are from a clear-day operation at Lincoln Laboratory’s Orlando (FL) testbed site (August 31, 1991). Figure 7 shows the computed dBZ low-beam reflectivity for this day, essentially representing the clutter environment for Orlando. In addition to strong returns from near the radar (image center), other areas of interest include reflections from east-west oriented power-lines and downtown Orlando to the NNW. The first priority is to alleviate strong clutter




Figure 7. 1991 Orlando clutter environment: computed dBZ reflectivity map.

where it significantly interferes with even the strongest of weather echoes, and the tendency is to focus on these regions alone. However the need to improve gust-front visibility, which is generally characterized by low-power returns, suggests that areas of moderate clutter be viewed with equal importance. It is perhaps for the cases of low-power returns combined with moderate clutter that the adaptive scheme can be of the most help. For this reason, the present analysis includes the totality of the scan data shown in Fig. 7.

3.1.1 Optimal effectiveness of a three-filter bank

Given the absence of significant weather returns for this data set (this is our model for low-signal strength processing), we rely on an estimated clutter-to-noise ratio (CNR) as gauge of filter performance. Figure 8 contains a histogram of computed CNRs for the data scan of Fig. 7 (256 radials $\times$ 240 range gates). In Fig. 8, vertical dashed lines have been added to mark the 0 dB CNR




Figure 8. 1991 Orlando clutter environment: computed CNR histograms.

threshold, the 60 dB threshold, and divisions that partition the histogram bin counts between 0 and 60 dB into three equal (in number) regions. A three-level adaptive filter bank directed at equalized filter *usage* would contain filters matched for the 0-12 dB, 12-30 dB, and 30-60 dB CNR ranges, as indicated by the (low-beam data) partition. The filters of Fig. 2 are only partially consistent with this, indicating that a slight redesign of stop-band depth may be more “optimal”. Fully one-third of the gates with some clutter return need never be “served” by anything more than the shallowest of filters.

How well do the filters of Fig. 2 match and suppress the clutter of Figs. 7 and 8? Using the criterion of a *post-filtering* CNR (PCNR) less than 0 dB as a measure for “success”, histogram distributions showing the number and location (by range) of gates that were serviced by each filter (0-3, with Filter #0 corresponding to “all-pass”) were constructed. These data are presented in Fig. 9, including the distribution of gates that could not pass the 0 dB PCNR criterion with any filter (i.e., “Failed at Filter #3”)³. The figure shows that (in the absence of weather) the all-pass filter is selected primarily outside the 100 gate range (almost exclusively outside the 150 gate range). The usage of Filters #1 and #2 is relatively balanced in number and illustrates a similar correlation in range. Gates that passed at Filter #3 are decreased in expected number (relative to any prediction based on Fig. 8); hence, there is a significant portion of gates that fail the criterion altogether; much greater than that predicted on the basis of Fig. 8, where only gates with CNR greater than 60 dB would fail. Although the 0 dB criterion is strict, and many of the “failures” represent statistical fluctuations above the mean noise threshold, there are many gates that do not fall into this category.

In the gate range 100-256, failure of Filter #3 is mostly due to the presence of contaminating Doppler returns outside the filter stop band. Figure 10 illustrates one such example where the filter was successful but failed the test. The figure shows an estimated single-gate (staggered) Doppler spectrum with and without the implementation of a clutter filter.

Virtually every gate within the 0-50 gate range fails the 0 dB PCNR criterion. For a large number of these, instability residue appears to be the cause. Figure 11 is a scatter plot of estimated input vs. output CNR for those gates that failed at the 60 dB filter. There is a prominent concentration of points near the “ $x = y$ ” identity mapping, corresponding to those gates that had Doppler returns outside the filter stopband. The most striking feature of the scatter plot in Fig. 11 is the break and linear rise that begins at an estimated input CNR of (approximately) 45 dB. This indicates that there is an instability on the order of -45 dB, rather than the assumed -60 dB. Figure 12 shows an estimated spectrum from a gate close to the radar. The estimated spectrum is obtained from a single-gate (27-pulse) sample (23-pulse sample in the case of post filtering), and the effect of the instability residue is to introduce an apparent (elevated) noise floor. Nearly all

³Note, in Fig. 9, for each range value, there is a maximum count of 256 azimuth locations, corresponding to the ASR 1.4 degree beamwidth.




Figure 9. Filter usage by range: August 31, 1991.




Figure 10. Single-gate Doppler spectrum example I. Top: estimated single-gate spectrum without clutter filtering; bottom: single-gate spectrum after Filter #3. (Input data were Blackman weighted).




Figure 11. Scatter plot of input vs. output estimated CNR for filter "failures".




Figure 12. Single-gate Doppler spectrum example II. Top: estimated single-gate spectrum without clutter filtering; bottom: single-gate spectrum after Filter #3. Instability residue effectively introduces an elevated “noise” floor.

gates within the 0-50 gate interval exhibit such an elevated noise floor. No further investigation of this phenomena was carried out for this report.

If there cannot be further improvement in the instability residue, then the implication is that the 60 dB filter can be reduced to put it in line with the expected residue threshold. There is a clear correlation between range, ideal filter choice, and residue “breakdown” that can be incorporated in the operational selection of clutter filter. As an additional note, frequency domain velocity estimators that need to do noise “subtraction” before moment calculation will perform better if provided a map indicating expected residue noise levels.

Figure 13 shows by location the clutter filter that best matched each clutter return of Fig. 7. This map shows the ideal filter to use given a very small weather-return signal. The theme of “filter




Figure 13. Ideal filter usage map: Orlando (FL).

less given strong weather returns” implies that such a utilization map should tend towards lower numbered filters as storms enter the radar’s view. Figure 14 shows the clutter residue that remains

after filtering as directed by the ideal map of Fig. 13. Figures 15 and 16 show the filter choice and




Figure 14. Clutter residue after "ideal" filtering.

residue when the current adaptive selection method is used; the data being the same clear-weather case. This example illustrates that in the situation of low-power weather returns, such as with gust fronts, the current trend is to overfilter. This suggests that a modification be made to the adaptive scheme to limit filter selection based on the *a priori* clutter information, as manifest in Fig. 13.

3.1.2 Improved sensitivity from filter yield

The new filters provide almost twice the number of data samples for moment estimation. Figures 17 and 18 contain pulse-pair velocity estimates for a case having a weak sea-breeze front east of and moving toward the radar. This case occurred on July 11, 1992. Figure 17 was obtained after filtering using the standard 17-coefficient FIR filters. Figure 18 was obtained after using the new time-varying filters. An increased sensitivity to low power winds is clearly indicated.




Figure 15. Adaptive filter usage given low-power weather returns.

3.2 Preservation of Pulse Stagger: Velocity Dealiasing Example

Another advantage of the new design is that the high-low-high pulse stagger relationship is maintained in the output. Hence, velocity dealiasing schemes that use dual velocity estimates and Nyquist residue checking (such as “Chinese Remainder Theorem”) can be applied to data that have been filtered. Figure 19 shows the test of one such dealiasing method. The upper-left panel shows a scan which has not been filtered for clutter. Instances of folded velocities exist near azimuth 200° between 9 and 13 nm. The upper-right panel corrects the folded (positive red) velocities using separate high- and low-PRF velocity estimates and a correction based on residue difference. The lower-left panel shows the same data, filtered using clutter Filter #1. The lower-right panel has been filtered for clutter and has had the velocity dealiasing procedure applied. Clearly, folded velocities have been detected and corrected in the filtered data.




Figure 16. Clutter residue given low-power weather returns and after adaptive filter selection.




Figure 17. Pulse-pair velocity estimates with 17-coefficient FIR filters. Sea-breeze front is perceptible due east of center at 14 km.




Figure 18. Pulse-pair velocity estimates with new time-varying filters.




Figure 19. Velocity dealiasing combined with clutter filtering. Top-left panel: low-PRF velocity estimates, unfiltered data (white streak is due to clutter bias). Top-right panel: corrected velocity estimates (low- and high-PRF data), unfiltered data. Bottom-left panel: low-PRF velocity estimates, filtered data. Bottom-right panel: corrected velocity estimates, filtered data.

4. CONCLUSIONS

This report supplements a previous analysis [1] of ground clutter rejection requirements for low-altitude wind measurements with an airport surveillance radar. The present results validate that the new time-varying filter designs easily meet the WSP filtering criteria as set forth in [1]. The most significant advantage of these new designs is that they provide a greater number of pulse samples for parameter estimation, which should improve estimation accuracy. In addition, since these filters maintain the pulse stagger relationship on output, we can now implement dual-PRF velocity dealiasing schemes with clutter filtering. Previously, this was not possible.

The new filter designs are optimal and could easily provide up to 80 dB suppression. The instability residue of the transmitter remains a limiting factor, currently degrading the desired performance with input signals of power 45 dB or greater. For returns of high power, the instability residue has the effect of introducing an elevated “noise floor” in spectral representations. Virtually every range gate within the interval 0-50 exhibits a spectral floor dominated by the instability residue. It may be necessary to account for this when using spectral-domain velocity estimators that attempt to compensate for the spectral noise floor. Barring reduction of the instability residue, a redesign of the filters may be in order to better distribute the filtering “load”.

Finally, it should be remarked that the new clutter filters are time varying and implemented as matrix multiplies. This does represent an increase in processing as compared to the current 17-coefficient FIR method. However, this processing overhead can be combined with frequency transformation operators that might also lead to more effective (frequency domain) velocity estimation methods. An increased sensitivity in estimation was indicated by one low-signal strength example. This and the incorporation of optimal velocity estimation algorithms, which can be implemented in the frequency domain, may improve WSP detection of gust fronts. Further research is needed to study this possibility.

APPENDIX A

TIME-VARYING FILTER RESPONSE PROFILES

The output response of each filter is measured at each output time in response to complex sinusoidal frequency input. Since the ASR-9 WSP 8-10-8 block is symmetric about pulse number 13 (output pulses numbered 0 – 26), the response profiles are also “symmetric” about output pulse 13 (phase response, being an odd function, is “flipped” for the corresponding pulses). Hence, only responses for pulses 2-13 are shown. Response profiles for output pulses 0 and 1 are not shown as they were not used for the average profiles shown earlier (except in the case of filter #1). In all cases, pulse #13 achieves exact linear phase because it is the only output pulse corresponding to a symmetrical placement of input pulses. Single-pulse phase error is generally kept within five percent of Nyquist. Because the phase errors do not “sum coherently”, the actual effect on velocity error is generally less. Although near linear phase can be achieved at the cost of a degraded magnitude response, additional phase control was not needed for the test data examined.




Figure A-1. MSE Design #1 (20 dB): frequency magnitude and phase error response.




Figure A-1 (Continued).




Figure A-1 (Continued).




Figure A-1 (Continued).




Figure A-2. MSE Design #2 (40 dB): frequency magnitude and phase error response.




Figure A-2 (Continued).




Figure A-2 (Continued).




Figure A-2 (Continued).




Figure A-3. MSE Design #3 (60 dB): frequency magnitude and phase error response.




Figure A-3 (Continued).




Figure A-3 (Continued).




Figure A-3 (Continued).

REFERENCES

1. M.E. Weber, "Ground clutter processing for wind measurements with airport surveillance radars," MIT Lincoln Laboratory, Lexington, Massachusetts, Project Report ATC-143 (DOT/FAA/PM-87/21), 1987.
2. E.S. Chornoboy, "FIR design methods for non-uniformly sampled signals," MIT Lincoln Laboratory, Lexington, Massachusetts, Technical Report 976, 1994. **In Revision.**