Real-time airborne radar wind-profiling algorithm for the Integrated Wind and Rain Airborne Profiler (IWRAP)

S J Frasier*and T Chu
Microwave Remote Sensing Laboratory
University of Massachusetts
Amherst, MA 01003

1. Introduction

The Integrated Wind and Rain Airborne Profiler (IWRAP) Esteban-Fernandez et al. (2005) is a dual-frequency dual-polarization radar developed by the Microwave Remote Sensing Laboratory at the University of Massachusetts that is installed routinely on the NOAA WP-3D research aircraft. Primarily designed to study the signature of the ocean surface under wind forcing, it has recently been modified to enable profiling through precipitation through the use of pulse-compression. The radar scans conically below the aircraft with a scan rate of 60 rpm, thereby enabling VAD wind retrievals at a 1 s update rate.

Without compensation, wind retrievals at altitudes below approximately 300 m are subject to contamination by scattering from the sea surface through antenna sidelobes, inducing a negative bias in lower boundary layer winds. We have implemented two algorithms that are amenable to real-time implementation that attempt to suppress this surface echo to enable improved retrieval of lower boundary-layer winds. One algorithm employs an adaptive notch-filter that rejects the surface echo, enabling fast covariance-based estimation of the precipitation echo Doppler shift. This process is performed entirely in the time-domain. Another algorithm employs more traditional spectral estimation with a narrow window about the anticipated range of Doppler frequencies.

2. Measurement Geometry

The typical IWRAP radar installation on the NOAA-P3 is shown in figure 1. Two radars (one C-band and one Ku-band) scan conically below the aircraft at incidence angles usually between $30°$ and $50°$. Each radar is capable of implementing up to four simultaneous beams, however, two simultaneous beams per radar is the normal mode of operation. The radar beamwidths vary depending upon the incidence angle, but are typically in the range of $5-10°$. Both H and V polarizations are available, and are selected based upon mission requirements. Figure 2 shows samples of reflectivity and Doppler velocity obtained during flights.

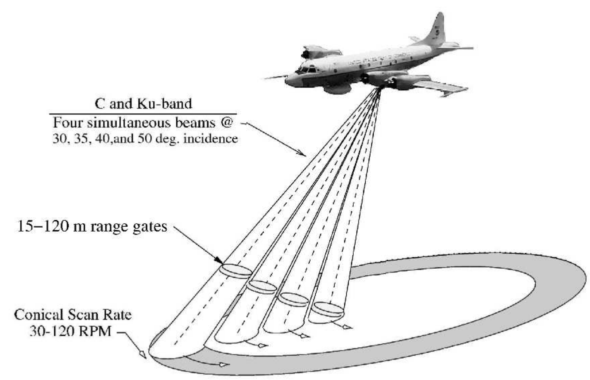

Figure 1: IWRAP radar measurement geometry on the NOAA P-3 research aircraft.

The IWRAP data acquisition system consists of a four-channel digital receiver sampling the intermediate frequencies of the C- and Ku-band radars. The receivers demodulate the echoes to baseband I- and Q-components which are stored to high speed RAID disks and simultaneously broadcast via UDP over a gigabit ethernet network. A real-time processor and display unit monitors all the broadcasts and computes Doppler spectral moments in real-time. The sustained data rate for the four channels collectively is 80 MB/s.

This paper is concerned with improved methods to compute these moments accounting for the influence of surface contamination for range bins near the sea-surface. We describe two algorithms which have been implemented on the real-time processor, and we provide some initial assessment of performance.

*frasier@ecs.umass.edu

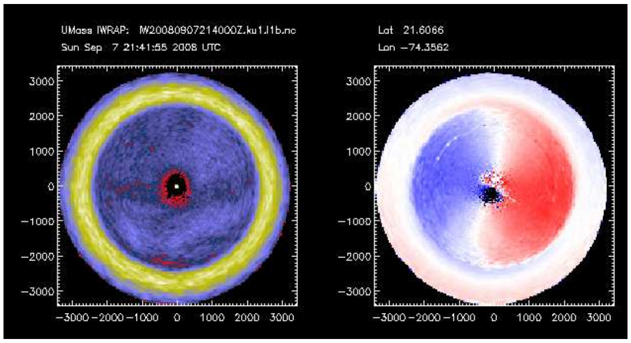

Figure 2: Aircraft-relative PPI displays of Ku-band reflectivity (left) and Doppler velocity (right) observed at $50°$ incidence. Range is horizontal ground range from the aircraft. The yellow ring in reflectivity is the sea-surface echo, with precipitation echo in blue. The aircraft velocity has been removed from the Doppler (by subtracting the surface velocity along each radial). Winds are from the left to right.

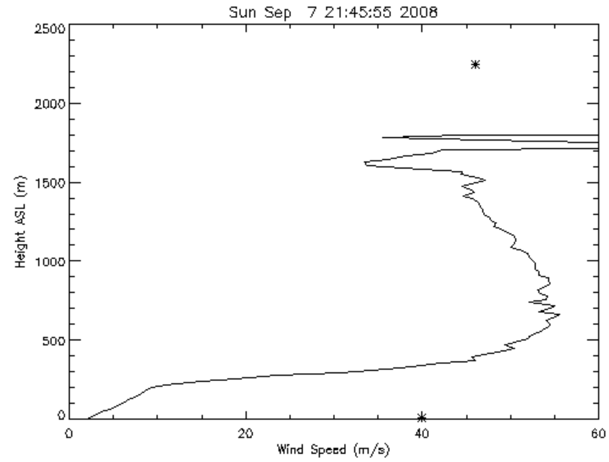

Figure 3: Wind speed profile obtained on 7 Sept 2008 (Hurricane Ike). Asterisks indicate aircraft flight-level winds and surface winds from microwave radiometer. Winds are corrected for aircraft motion by subtracting the surface echo's apparent velocity. Surface contamination is evident in the wind profile from the surface up to approximately 300 m.

3. VAD measurement

Wind profiles are most directly obtained via the Velocity Azimuth Display technique after correction for aircraft attitude and velocity Lee et al. (1994). Both algorithms described herein rely on an initial estimate of the surface Doppler echo obtained via the pulse-pair technique. The single lag covariance is accumulated over 125 pulses, corresponding to an integration time of 8.33 ms for a 15 kHz pulse rate (typical). During this interval, the antenna has rotated aproximately 3 degrees.

Pulse-pair techniques are sufficient to obtain accurate Doppler measurements for most altitudes. It is only for ranges between that of the nadir echo (shortest distance to the surface) and beyond that we find it necessary to more sophisticated methods. In these cases, the Doppler spectrum is often no longer unimodal, but is more often bimodal. Thus it is necessary to identify the correct feature in the Doppler spectrum corresponding to the weather echo and reject that of the surface clutter. Because the measurements are made from a moving platform with a rotating antenna, both of these sources have widely varying Doppler shifts. While it is often easy to see, visually, the correct peak in the Doppler spectrum to choose, it is much more difficult to devise an automated algorithm.

4. Adaptive Notch Filter Algorithm

The notch filter algorithm attempts to identify the surface clutter, build a notch filter around its center frequency, and filter it out thereby permitting a conventional pulse-pair analysis on the remaining signal, which is assumed unimodal. The objective is to implement a filter of low order (requiring minimal computation) that is sufficient to reject the surface clutter, and that is sufficiently selective to avoid rejecting the desired signal. Because of the low order of the filter, it is faster to implement the convolution in the time domain.

The surface echo is identified as the strongest echo beyond the delay corresponding to the aircraft altitude, which is known separately. The Doppler shift of this echo is determined from the phase of the covariance in the usual way.

This techinque is based on a simple second-order infinite-impulse-response (IIR) filter designed by placing poles and zeros about the desired notch frequency as illustrated in Fig 4. The filter is given by the difference equation

$$y_n = K_0(x_n - Ax_{n-1} + Bx_{n-2}) + Cy_{n-1} - Dy_{n-2} \quad (1)$$

where

$$\begin{aligned} A &= z_1 + z_2 \quad (2) \\ B &= z_1 z_2 \\ C &= p_1 + p_2 \\ D &= p_1 p_2 \end{aligned}$$

and where $z_1, z_2, p_1,$ and p_2 denote the values of the zeros and poles. The transfer function of the filter is given

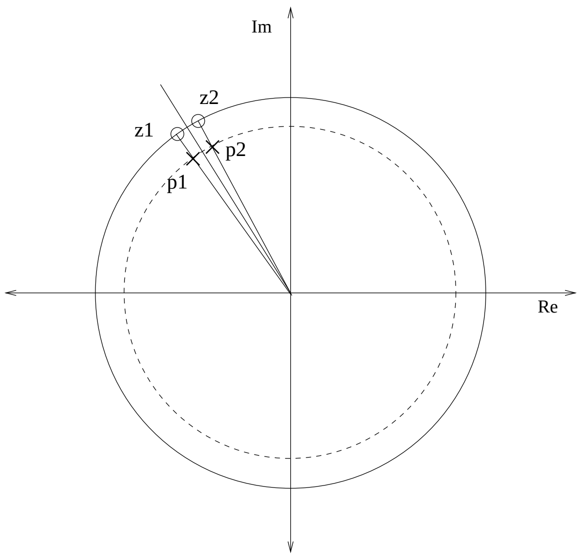

Figure 4: Z-plane pole-zero plot of the 2nd-order filter used to implement a notch filter at the given radian frequency. Zeros are placed on the unit circle (solid) about the desired frequency; poles are placed at the same angles but at a smaller radius (dashed), as required for stability.

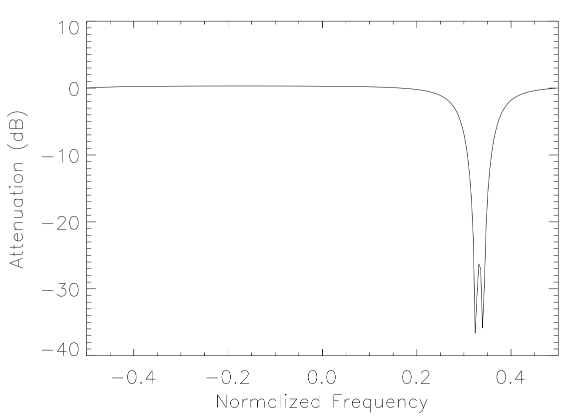

Figure 5: Amplitude response of the IIR filter depicted in Fig 1

by its Z-transform

$$H(z) = K_0 \frac{(z - z1)(z - z2)}{(z - p1)(z - p2)} \quad (3)$$

where K_0 is a gain factor. Although the IIR filter has a non-linear phase response (unlike FIR filters which can be perfectly linear), this is not a significant deterrent to using it. Ultimately, we compute Doppler shifts from phase differences, and not the phases themselves. The advantage of the IIR filter is its simplicity and its sharp cutoff considering the low filter order which minimizes computation.

The algorithm can be summarized by the following steps:

1. Accumulate pulse-pair covariances for all range bins.

2. Identify surface echo (strongest echo).

3. Compute filter zero/pole locations in accordance with surface echo covariance phase.

4. Convolve time-series with filter for each affected range bin.

5. Re-accumulate pulse-pair covariance on filtered time-series.

This simple algorithm suffers from a few drawbacks. The primary one is that the filter rejection and width needs to be reasonably matched to the intensity and spectral width of the surface echo. In the absence of any significant weather echo the algorithm tends to produce a velocity of opposite velocity to the surface echo. This is a consequence of the “hole” placed in the spectrum at the filter frequency.

5. Sliding Spectral Analysis Window Algorithm

In this method, there is no attempt to null or reject the surface echo, but rather to track the atmospheric echo. The algorithm works from near ranges to far ones keying on the Doppler centroid at moderate heights that are free of surface clutter. Once the Doppler centroid is identified, an integration window approximately 1/10 of the full Nyquist interval is used to calculate the centroid for the next range bin, and the interrogation window for the subsequent bin is based upon the current and prior bin’s centroid (Fig. 6).

The algorithm is summarized by the following steps:

1. Accumulate pulse-pair covariances for all range bins.

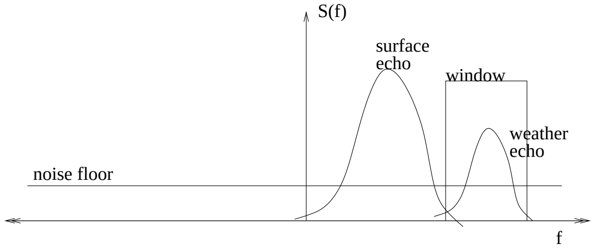

Figure 6: Doppler spectrum at a range bin containing both surface echo and weather echo. At ranges corresponding to nadir return, surface echo is centered at zero frequency. Otherwise, both weather and surface echoes are at varying frequencies depending upon radar azimuth and aircraft and wind speeds. A narrow interrogation window is determined

2. Choose initial range.
3. Calculate Doppler centroid (weather-only echo).
4. Form interrogation window (FFT bins about centroid)
5. Compute FFT of time-series from next range bin.
6. Compute integrated power and Doppler centroid within interrogation window only.
7. Update interrogation window based on new Doppler centroid.
8. Repeat 5-7 until surface range bin is reached.

The adaptation of the window with height allows for tracking some directional shear. In practice, a running average of the centroids of the previous few range bins provides better control of the window location. This method appears quite effective in following the the weather echo (when present) down to the surface, except when the weather echo and surface echoes overlap in frequency in which case the window may lock onto the surface echo. This happens twice per revolution of the antenna.

As in the previous method, in the absence of any weather signal, this algorithm still tends to produce a smooth looking velocity profile, which is a consequence of the initial choice of the interrogation window. Thus, signal-to-noise ratio or another metric must be used to insure the validity of the velocity profile. Although this technique requires computation of the full Doppler spectrum and is therefore more computationally intensive than the prior method, it is still tractable for real-time operation.

Evaluation of these two approaches is ongoing. Of the two methods described, the latter method appears to be preferable and less prone to pathological behavior.

6. Acknowledgments

This work is supported by a NASA Earth and Space Systems Fellowship (T. Chu), by the NASA Ocean Vector Winds Science Team, and in collaboration with NOAA-NESDIS (P. Chang).

References

Esteban-Fernandez, D., E. A. Kerr, A. Castells, S. J. Frasier, J. R. Carswell, P. S. Chang, P. Black, and F. D. Marks, 2005: IWRAP: The imaging wind and rain airborne profiler for remote sensing of the ocean and the atmospheric boundary layer within tropical cyclones. *IEEE Trans. Geosci. Remote Sensing*, **43**, doi:10.1109/TGRS.2005.851640.

Lee, W.-C., P. Dodge, F. D. Marks, and P. H. Hildebrand, 1994: Mapping of airborne doppler radar data. *J. Atmos. Oceanic. Tech.*, **11**, 572–578.