Performance of a Proposed Hybrid Spectrum Width Estimator for the NEXRAD ORDA

GREGORY MEYMARIS*
National Center for Atmospheric Research, Boulder, Colorado

JOHN K. WILLIAMS
National Center for Atmospheric Research, Boulder, Colorado

JOHN C. HUBBERT
National Center for Atmospheric Research, Boulder, Colorado

1. INTRODUCTION

With the advent of the Open Radar Data Acquisition (ORDA) system on WSR-88D radars and the introduction of significantly more powerful signal processing hardware comes the opportunity to improve the method used for estimating the spectrum width, a measure of the variability of radial wind velocities within a measurement pulse volume. In addition, the implementation of new operational modes for improved data quality, including SZ phase coding, will involve very different signal processing techniques and hence may require novel methods to meet the WSR-88D specifications. While spectrum width has not been used extensively by radar meteorologists in the past, the new NEXRAD Turbulence Detection Algorithm (NTDA), developed under direction and funding from the FAA's Aviation Weather Research Program, will soon be using the WSR-88D spectrum width as a key input for providing in-cloud turbulence estimates (eddy dissipation rate, EDR) for an operational aviation decision support system (Williams et al. 2005). Achieving improved spectrum width estimator performance would directly benefit the accuracy of the NTDA product.

This paper addresses these issues by evaluating performance characteristics of several spectrum width estimators, including the pulse-pair estimator currently used in the WSR-88D. Evaluations are performed using simulated radar time-series data representing a variety of scenarios for different signal-to-noise ratios, overlaid power ratios, and spectrum widths. A hybrid algorithm combining three spectrum width estimators is proposed, and it is shown that this algorithm, while slightly more computationally intensive, is more accurate and robust than any method alone.

2. Methodology

To evaluate and compare different spectrum width estimators we generated random complex time-series data for various true spectrum width, signal-to-noise ratio (SNR) and overlaid power ratio (PR) scenarios. We used an I&Q simulation technique based on the method described in Frehlich and Yadlowsky (1994); Frehlich (2000); Frehlich et al. (2001) except that the autocorrelation function is that of a weather echo as defined in Doviak and Zrnić (1993, p. 125). This is a preferable method for generating complex time-series with a given average autocorrelation function, as opposed to what is described by Zrnić (1975), because it is not necessary to generate as long of a time-series in order to get the correct temporal statistics.

In what follows, the simulator input ("true") spectrum width will be denoted as W, while the estimated spectrum width will be denoted as $\hat{W}$ with a modifying subscript specifying the estimation technique used. Estimation errors were calculated by subtracting the simulator input values from the estimated values (i.e. $\hat{W} - W$). In this paper, we do not show plots of standard errors (i.e., RMS errors); rather, we break out the error analysis into biases and standard deviations, which have quite different

* *Corresponding author address*:
Gregory Meymaris, RAL, NCAR
1850 Table Mesa Dr., Boulder, CO 80305
E-mail: meymaris at ucar.edu

implications for turbulence detection since bias cannot be mitigated by averaging while random unbiased errors can. However, RMS error estimates may be obtained by taking the square root of the sum of the squared biases and standard deviations.

3. Spectrum Width Estimators

In this section, we used the simulator to generate short PRT data with the following characteristics, unless otherwise noted: wavelength $\lambda = 10.5$ cm, the pulse repetition time (PRT) is 987 μs, the number of samples per time-series (M) is 88, signal-to-noise ratios (SNR) of 10 dB, 20 dB, and 30 dB, and varying input spectrum widths. This corresponds to the NEXRAD volume control pattern (VCP) 21, which is commonly used for storms that are not expected to evolve quickly.

a. The $R0/R1$ Pulse Pair Estimator

The standard spectrum width estimator currently used in the WSR-88D radars on short PRT data is the $R0/R1$ estimator (Doviak and Zrnić 1993), so named because it utilizes the ratio of the first two lags of the autocorrelation function:

$$\hat{W}_{s01} = \left(\sqrt{2}/\pi\right) V_a \left|\log\left(P_S/\left|R_1\right|\right)\right|^{1/2} \qquad (1)$$

The "s" in the subscript "s01" indicates that the short PRT data are used. Here V_a is the Nyquist velocity, P_S is the average power of the signal with noise removed, and R_1 is the first lag of the autocorrelation function (i.e. $R_1 = (n-1)^{-1} \sum_{k=1}^{n-1} V^*(k) V(k+1)$ where $V(k)$ are the complex-valued I&Q radar time-series). In the event that $|R_1| < P_S$, in which case the log has a negative argument, the spectrum width is set to 0 as is done on the WSR-88D.

The performance statistics obtained via simulation for the short PRT (913 μs) $R0/R1$ spectrum width estimator in the case of (essentially) no overlaid echoes is shown in Figure 1 for various input spectrum widths and SNRs. The biases are shown in Figure1(a), and the standard deviation of the errors $\hat{W}_{s01} - W$ is depicted in 1(b). The error standard deviation plot agrees reasonably well with that in Doviak and Zrnić (1993), although there are some differences. These may be caused by different approaches to dealing with the cases where $|R_1| < P_S$, or to different methods used to generate time-series segments for analysis. The biases and standard deviations show that for low SNRs (0 and 4 dB) this estimator is very poor, with large error standard deviations and large and variable bias values. As SNR increases to 10 dB and greater, the bias relative to the input spectrum width improves dramatically for all but rather small or quite large input spectrum widths, and the error standard deviations improve for small and, especially, medium spectrum width values. For large input spectrum widths, the spectrum width estimator eventually saturates, as can be seen from the increasing negative bias for all SNR levels.

b. The $R1/R2$ Pulse Pair Estimator

Another estimator described by Doviak and Zrnić (1993) is the $R1/R2$ estimator, which is based on the ratio of the first and second lags of the autocorrelation function:

$$\hat{W}_{s12} = \left(2/\left(\pi\sqrt{6}\right)\right) V_a \left|\log\left(\left|R_1/R_2\right|\right)\right|^{1/2} \qquad (2)$$

where R_2 is the second lag of the autocorrelation function (i.e. $R_2 = (n-2)^{-1} \sum_{k=1}^{n-2} V^*(k) V(k+2)$. In the event that $|R_2| < |R_1|$, the spectrum width is set to 0.

The performance statistics obtained via simulation for the short PRT (913 μs) $R1/R2$ spectrum width estimator in the case of (essentially) no overlay for various input spectrum widths and SNRs is shown in Figure 2. The biases are shown in Figure2(a), and the error standard deviation in 2(b). The biases and error standard deviations show that for 0 dB SNR this estimator is very poor, but the performance for 4 dB is much improved over the $R0/R1$ estimator. Again, there are biases for small input spectrum widths, but the performance is significantly better than the $R0/R1$ estimator in this regime, particularly for SNRs of 10 dB or higher. In fact, the estimator as a whole performs better than $R0/R1$ until the input spectrum width approaches 8 m/sec. At that point the $R1/R2$ estimator saturates, leading to severe negative biases.

c. The $R1/R3$ Pulse Pair Estimator

The $R1/R3$ estimator, is derived in the same way that the above pulse-pair estimators. It is based on the ratio of the first and third lags of the autocorrelation function:

$$\hat{W}_{s13} = \left(1/\left(2\pi\right)\right) V_a \left|\log\left(\left|R_1/R_3\right|\right)\right|^{1/2} \qquad (3)$$

where R_3 is the second lag of the autocorrelation function (i.e. $R_3 = (n-3)^{-1} \sum_{k=1}^{n-3} V^*(k) V(k+3)$. In the event that $|R_3| < |R_1|$, the spectrum width is set to 0. This estimator behaves much like the

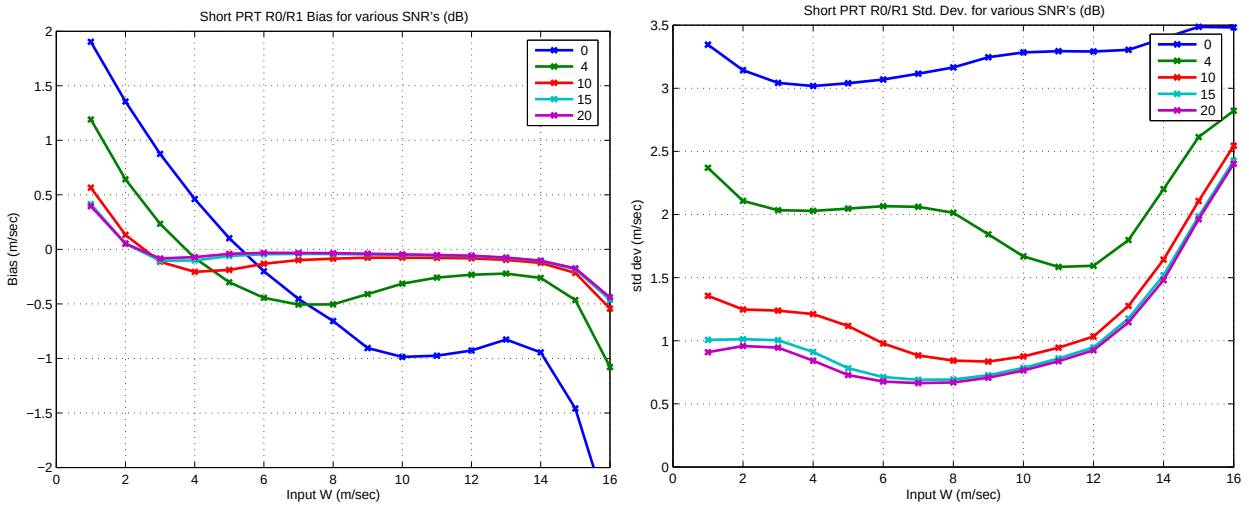

(a) Bias of $\hat{W}_{s01}$

(b) Standard Deviation of $\hat{W}_{s01}$

Figure 1: Bias and error standard deviation plots of the short PRT $R0/R1$ spectrum width estimator for varying input spectrum widths and SNRs (0, 4, 10, 15 and 20 dB shown). The PR in this data is set at 30 dB, low enough such that the weak trip does not significantly impact the statistics.

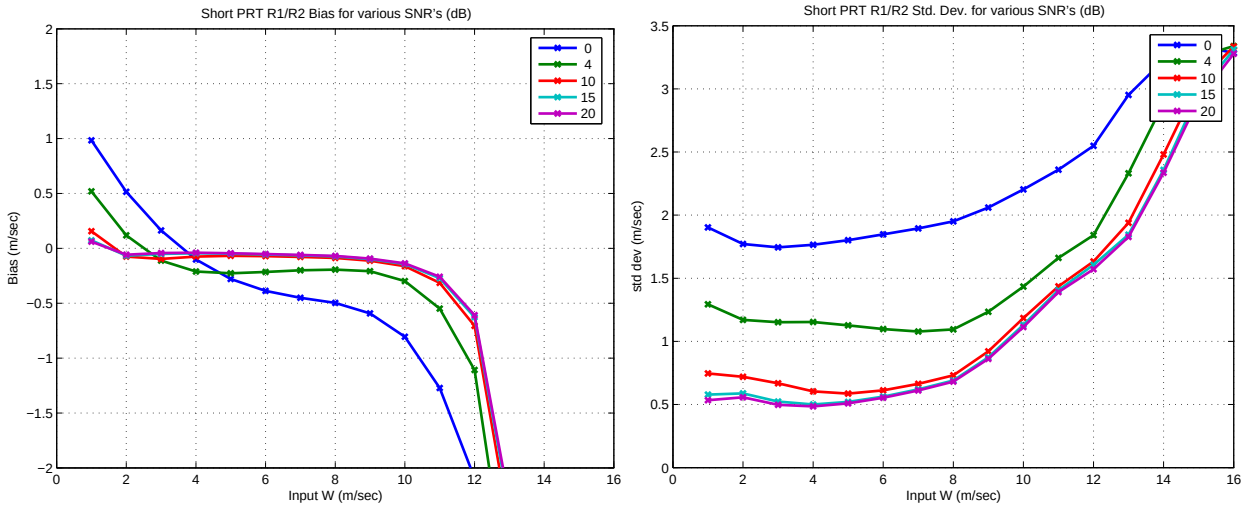

(a) Bias of $\hat{W}_{s12}$

(b) Standard Deviation of $\hat{W}_{s12}$

Figure 2: Bias and error standard deviation plots of the short PRT $R1/R2$ spectrum width estimator for varying input spectrum widths and SNRs (0, 4, 10, 15 and 20 dB shown). The PR in this data is set at 30 dB, low enough that the weak trip does not significantly impact the statistics.

$R1/R2$ estimator except that it performs better at very narrow spectrum widths, but also saturates very quickly.

d. The PPLS2 Estimator

All of the above models assume that the autocorrelation function is a Gaussian, and this one does also. In the above ones the fit of the Gaussian is exactly determined by 2 points (lags). However, the pulse-pair least squares estimator (PPLS2) uses 3 points, lags 0, 1, and 2, and is thus over-determined. The log of the autocorrelation function, if Gaussian, is a concave-down quadratic, and thus a least squares fit can be found efficiently. This estimator behaves somewhere between the $R0/R1$ and the $R1/R2$, as might be expected. It performs better than $R0/R1$ for narrow spectrum widths, although not as good as $R1/R2$, and worse than $R0/R1$ for wide spectrum widths, although better than $R1/R2$.

4. A Hybrid Approach

The three estimators ($\hat{W}_{s01}$, $\hat{W}_{s12}$, and $\hat{W}_{s13}$) each performs well in certain regimes. $\hat{W}_{s01}$ performs well in higher SNRs and for larger spectrum widths, whereas $\hat{W}_{s12}$ performs well for slightly lower SNRs and medium-valued spectrum widths. The estimator $\hat{W}_{s13}$ performs the best for very narrow spectrum widths. These complementary regimes of relatively good performance suggest that a hybrid approach where the appropriate estimator is used depending on the true (but unknown) spectrum width, might achieve good overall performance. Because the true spectrum width is unknown, a guess is made by calculating different estimators ($R0/R1$, $R1/R3$, and PPLS2) and then using a heuristic algorithm. Once the decision (guess) is made whether the true spectrum width is narrow, medium, or wide, then the appropriate estimator ($R1/R3$, $R1/R2$, and $R0/R1$, respectively) is used to calculate the final spectrum width estimate.

a. Algorithm

- The spectrum width estimators $\hat{W}_{s01}$, $\hat{W}_{s13}$, and $\hat{W}_{sPPLS2}$ are calculated.
- Based on n, the number of samples in the time-series, a table lookup of the wide *normalized* spectrum width threshold, w_{tw} is performed. By *normalized* we mean that the spectrum width threshold must be multiplied by V_a in order to be directly compared to the spectrum width estimators.
- If $\frac{1}{2}\left(\hat{W}_{s01} + \hat{W}_{sPPLS2}\right) > V_a w_{tw}$ then the spectrum width is guessed to be large. In which case, $\hat{W}_{s01}$ is the final output.
- Otherwise, another table lookup is performed (again based on n) to find the narrow normalized spectrum with threshold, w_{tn}.
- If $\hat{W}_{s13} < V_a w_{tn}$ then the spectrum width is guessed to be small. In which case, $\hat{W}_{s13}$ is the final output. For smaller values of n ($n \leq 58$), w_{tn} is set to -1, in which case this comparison is always false, and the algorithm proceeds to the next step. This is done because for smaller values of n, the capabilities of any tested estimator (including $\hat{W}_{s13}$) for discriminating between narrow and medium spectrum widths is poor. Since it is better to guess that a narrow spectrum width is medium-sized than vice verse, the algorithm errs on the side of guessing that the spectrum width is medium-sized.
- Otherwise, the spectrum width is guessed to be medium-sized. $\hat{W}_{s12}$ is calculated and returned as the final output.

In figure 3, the thresholds as a function of n are shown. These thresholds were obtained in a automated way by running simulation data through a classification decision tree. The costs associated with misclassifications were set to reflect that guessing that a spectrum width is too big is, in general, better than guessing that a wide spectrum width is narrow. This is true both from an estimator comparison standpoint as well as from the fact that wide spectrum widths are associated with hazards and so occasional over-warning is generally better than under-warning.

5. Results

Statistical comparisons of the $R0/R1$ estimator and the proposed hybrid estimator are shown in figures 4 - 6 (for $n = 38, 88,$ and 278, resp., with PRT of 780 μs). The left plots (a) show the results from the $R0/R1$ estimator, and the right (b) shows those from the hybrid estimator. As can be seen, there are substantial improvements, in both bias and standard deviation, for narrower spectrum widths. It would not

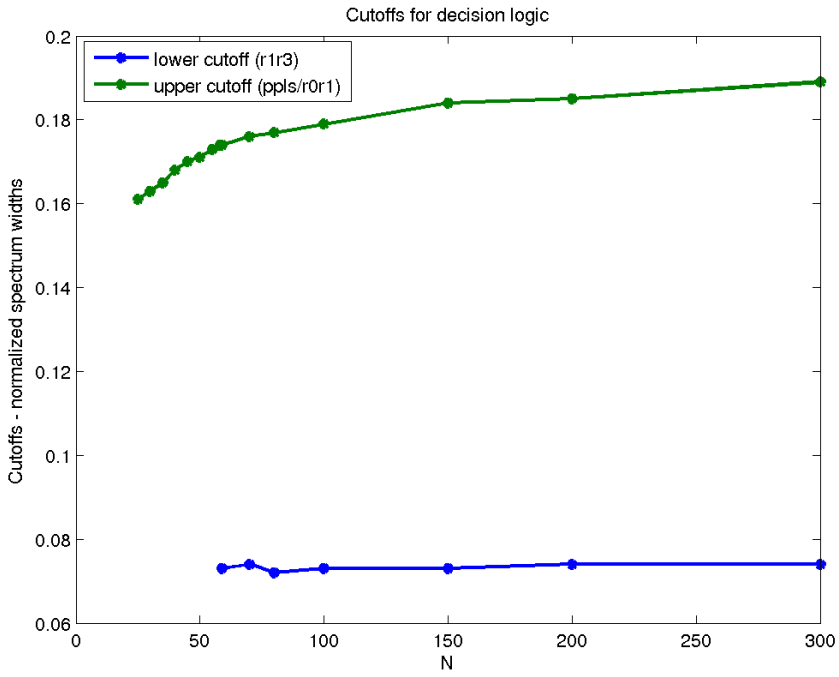

Figure 3: Plot of the thresholds (cutoffs) for narrow, medium, and wide spectrum widths as a function of n, the number of pulses in the time-series. The cutoffs are normalized and so must be multiplied by the Nyquist velocity to be compared to the estimators. The upper cutoff is compared first to the average of the PPLS2 and $R0/R1$ estimators. If the spectrum width is deemed *not* wide, then the $R1/R3$ estimator is compared to the lower cutoff to determine if the spectrum width is narrow.

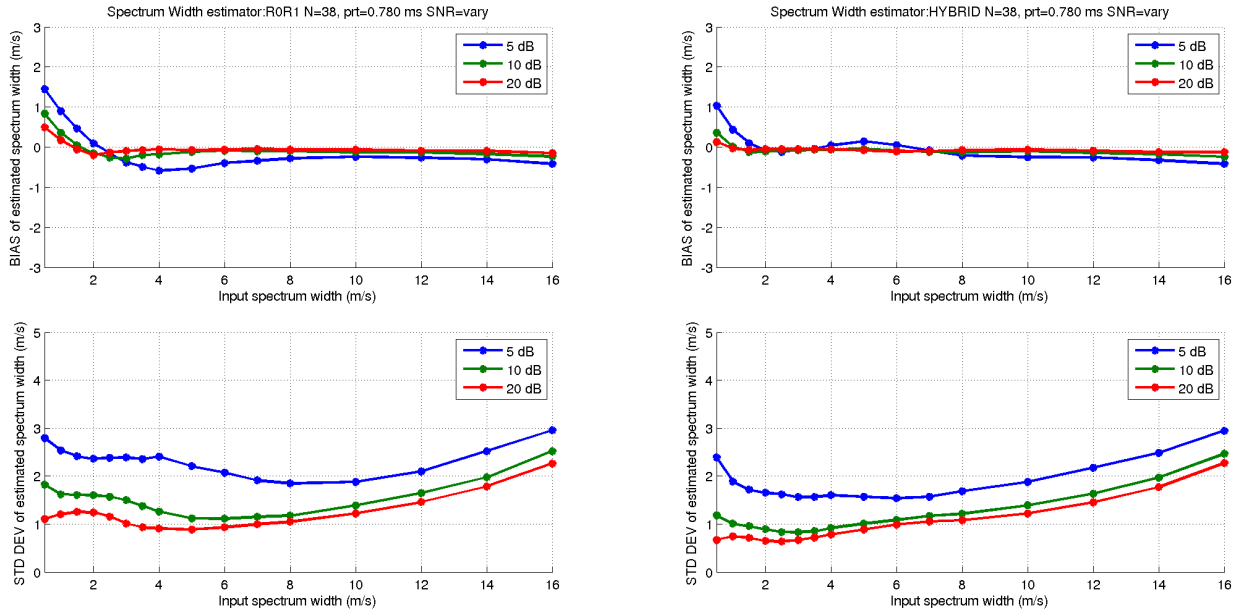

(a) Bias and Std of $\hat{W}_{s01}$

(b) Bias and Std of $\hat{W}_{hyb}$

Figure 4: Bias and error standard deviation plots of the $R0/R1$ (a), and hybrid (b) spectrum width estimators for varying input spectrum widths and SNRs (5, 10, and 20 dB) as a function of true input spectrum width. The number of samples per time-series is 38 and the PRT is 780 μs. The biases are shown in the top panels and the standard deviations are shown in the bottom panels.

be expected to see improvements for larger spectrum widths because the $R0/R1$ is *used* in the hybrid estimator for wide spectrum widths because it performs the best of all the estimators tested in that regime.

Another way of comparing the performance of the two estimators is via 2-D histograms of true input spectrum width versus the output from the estimators. This is similar to a scatter plot comparison of the data. These are shown in figures 7 - 9, with the left plots (a) again showing the results from the $R0/R1$ estimator, and the right (b) showing those from the hybrid estimator. The simulation parameters are the same as the above results, but focused solely on a SNR of 20 dB. As can be seen, the hybrid again performs better than $R0/R1$ for the narrower spectrum widths. It is possible to see a few more outliers, especially for $n = 38$ (figure 7) around an input width of 10 m/s. This is caused by wide spectrum widths being wrongly diagnosed as medium-sized. However, the number of outliers is quite small.

6. Conclusions

The simulation results presented in this paper have shown that the $R0/R1$ spectrum width estimator currently used on WSR-88Ds does not perform as well as the $R1/R2$ estimator or $R1/R3$ estimator in certain regimes. A hybrid approach that combines these methods using weights appropriate to each regime shows great promise in producing improved overall performance. While knowledge of the true spectrum width would allow determining the ideal estimator, an alternative that uses spectrum width estimates to try to decide the general magnitude of the true spectrum width was proposed as a practical alternative. The hybrid estimator presented in this paper was shown to outperform all three spectrum width estimators in most cases, and at the least did no worse than the $R0/R1$ estimator currently used by the WSR-88D. For larger spectrum widths, it outperforms the $R1/R2$ and $R1/R3$ estimators. Computationally, the hybrid algorithm is fairly modest, requiring fewer operations than the FFT needed by a spectral technique.

Future work includes improving the performance for small spectrum widths, where the less than opti-

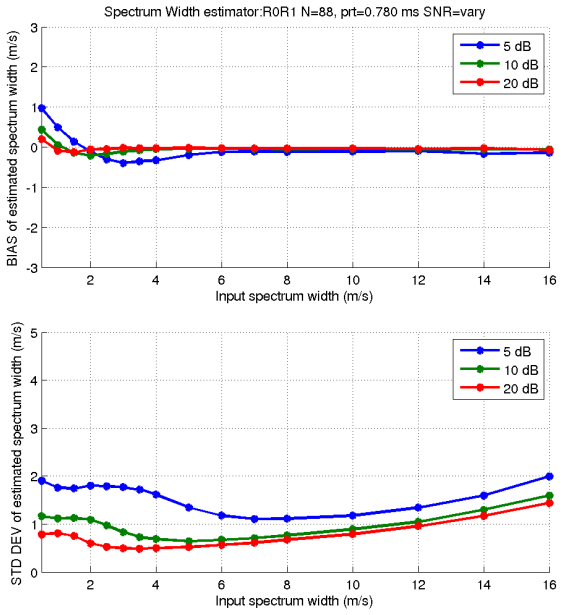

(a) Bias and Std of $\hat{W}_{s01}$

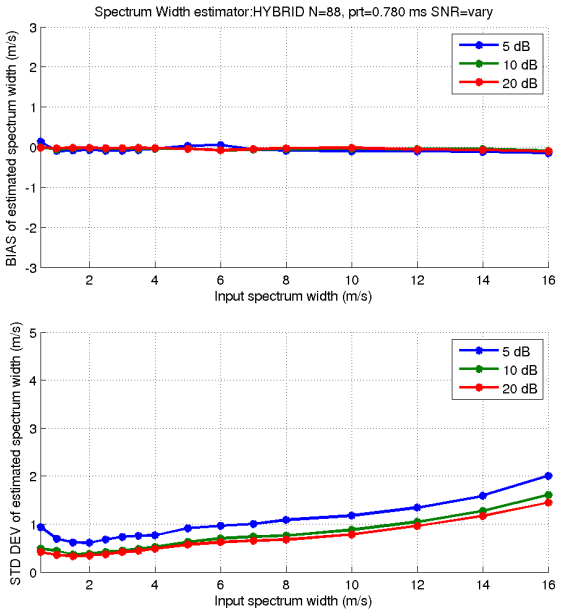

(b) Bias and Std of $\hat{W}_{hyb}$

Figure 5: Bias and error standard deviation plots of the $R0/R1$ (a), and hybrid (b) spectrum width estimators for varying input spectrum widths and SNRs (5, 10, and 20 dB) as a function of true input spectrum width. The number of samples per time-series is 88 and the PRT is 780 μs. The biases are shown in the top panels and the standard deviations are shown in the bottom panels.

mal quality seems to be due to the wrong decision about the general size of the true spectrum width. In addition, other spectrum width estimators such as spectral or maximum likelihood methods could easily be integrated into the general framework developed here, and this hybrid approach can be applied to other VCPs including those that involve phase-coded signals.

7. Acknowledgements

This research was supported in part by the ROC (Radar Operations Center) of Norman OK. Any opinions, findings and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the ROC.

References

Doviak, R. J. and D. S. Zrnić: 1993, *Doppler Radar and Weather Observations*, Academic Press, San Diego, California. 2nd edition.

Frehlich, R., 2000: Simulation of coherent doppler lidar performance for space-based platforms. *Journal of Applied Meteorology*, **39**, 245–262.

Frehlich, R., L. B. Cornman, and R. Sharman, 2001: Simulation of three-dimensional turbulent velocity fields. *Journal of Applied Meteorology*, **40**, 246–258.

Frehlich, R. and M. J. Yadlowsky, 1994: Performance of mean-frequency estimators for doppler radar and lidar. *Journal of Atmospheric and Oceanic Technology*, **11**, 1217–1230, corrigenda, **12**, 445–446.

Williams, J. K., L. Cornman, J. Yee, S. G. Carson, and A. Cotter: 2005, Real-time remote detection of convectively-induced turbulence. AMS 32nd Radar Meteorology Conference, Albuquerque, NM.

Zrnić, D. S., 1975: Simulation of weatherlike doppler spectra and signals. *Journal of Applied Meteorology*, **14**, 619–620.

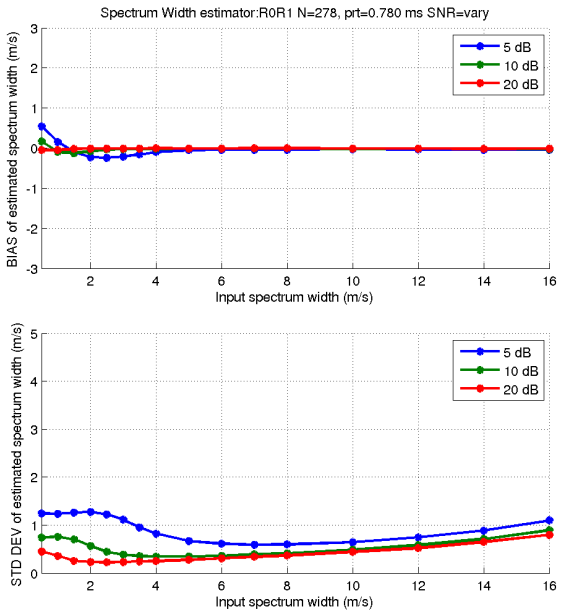

(a) Bias and Std of $\hat{W}_{s01}$

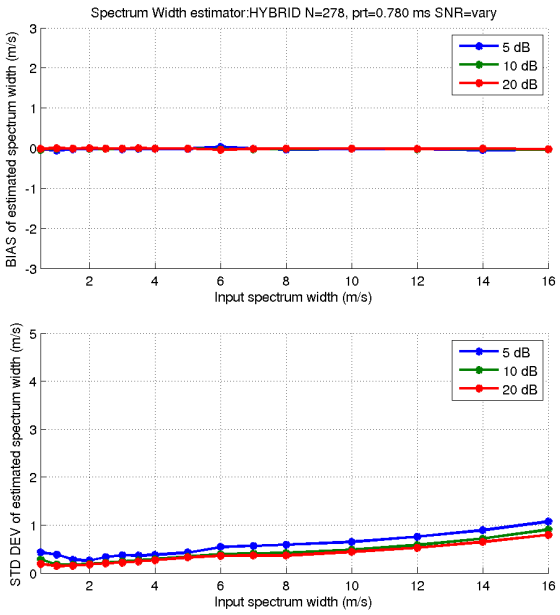

(b) Bias and Std of $\hat{W}_{hyb}$

Figure 6: Bias and error standard deviation plots of the $R0/R1$ (a), and hybrid (b) spectrum width estimators for varying input spectrum widths and SNRs (5, 10, and 20 dB) as a function of true input spectrum width. The number of samples per time-series is 278 and the PRT is 780 μs. The biases are shown in the top panels and the standard deviations are shown in the bottom panels.

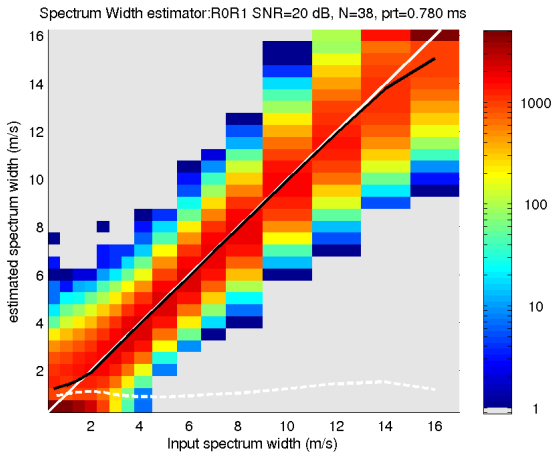

(a) $\hat{W}_{s01}$

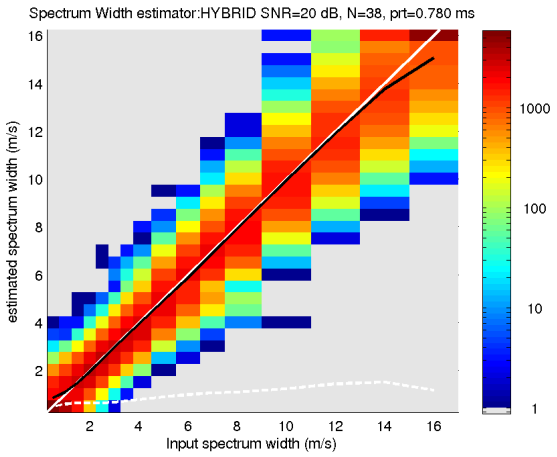

(b) $\hat{W}_{hyb}$

Figure 7: 2-D Histograms of true input spectrum width versus $R0/R1$ (a) and hybrid (b) estimators. The color corresponds to the frequency counts within the bins. Note that the color scale is logarithmic. The white line is the 1-to-1 line, the black line show follows the mean for each "column", and the dashed white line corresponds the standard deviation for each "column". The number of samples per time-series is 38 and the PRT is 780 μs. The biases are shown in the top panels and the standard deviations are shown in the bottom panels.

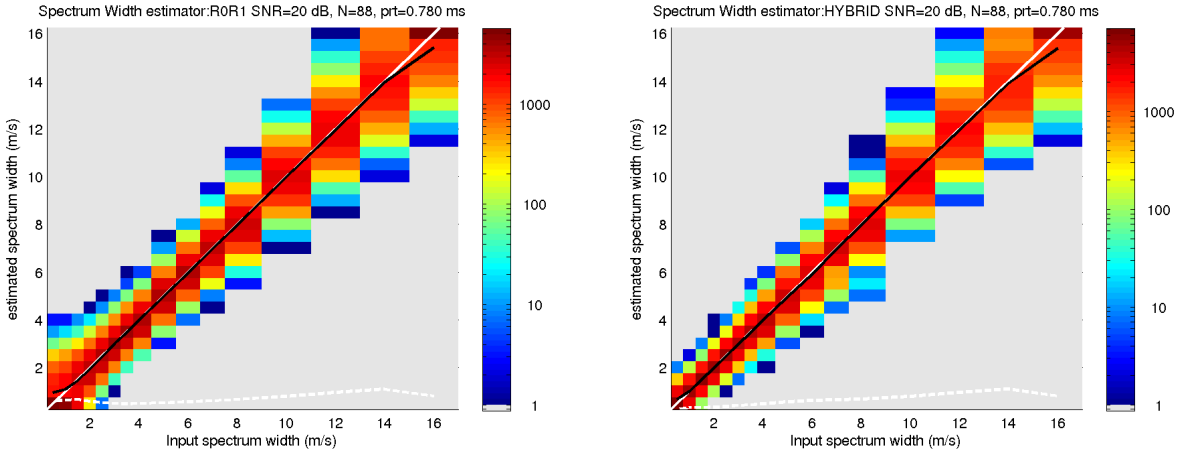

(a) $\hat{W}_{s01}$ (b) $\hat{W}_{hyb}$

Figure 8: 2-D Histograms of true input spectrum width versus $R0/R1$ (a) and hybrid (b) estimators. The color corresponds to the frequency counts within the bins. Note that the color scale is logarithmic. The white line is the 1-to-1 line, the black line show follows the mean for each "column", and the dashed white line corresponds the standard deviation for each "column". The number of samples per time-series is 88 and the PRT is 780 μs. The biases are shown in the top panels and the standard deviations are shown in the bottom panels.

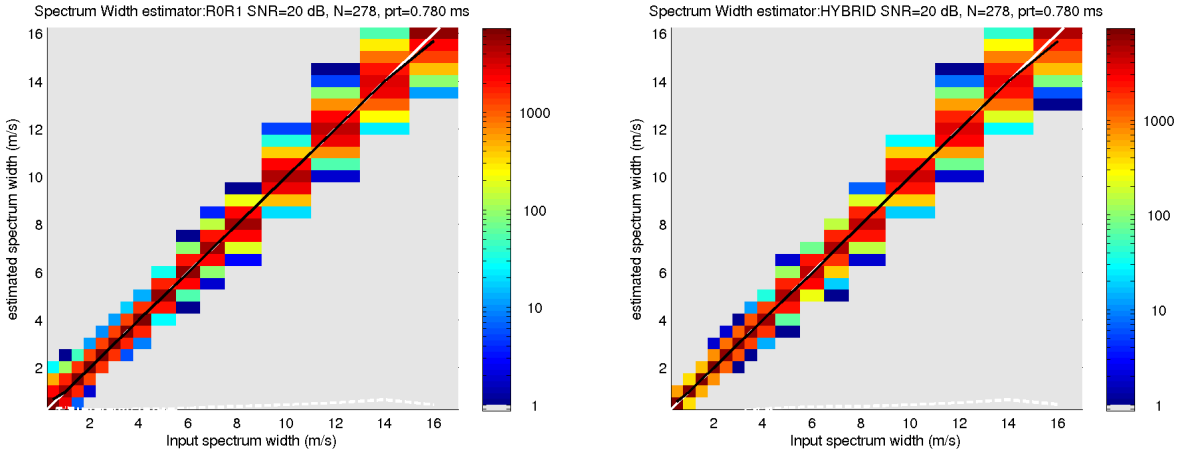

(a) $\hat{W}_{s01}$ (b) $\hat{W}_{hyb}$

Figure 9: 2-D Histograms of true input spectrum width versus $R0/R1$ (a) and hybrid (b) estimators. The color corresponds to the frequency counts within the bins. Note that the color scale is logarithmic. The white line is the 1-to-1 line, the black line show follows the mean for each "column", and the dashed white line corresponds the standard deviation for each "column". The number of samples per time-series is 278 and the PRT is 780 μs. The biases are shown in the top panels and the standard deviations are shown in the bottom panels.