Velocity Estimation Improvements for the ASR-9 Weather Systems Processor

James M. Kurdzo[1,2,3], John Y. N. Cho[3], and Robert S. Frankel[3]

[1] *Advanced Radar Research Center, University of Oklahoma, Norman, Oklahoma*
[2] *School of Meteorology, University of Oklahoma, Norman, Oklahoma*
[3] *Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, Massachusetts*

ABSTRACT

The Airport Surveillance Radar (ASR-9) is a rapid-scanning terminal aircraft detection system deployed at airports around the United States. To provide cost-effective wind shear detection capability at medium-density airports, the Weather Systems Processor (WSP) was developed and added on to the ASR-9 at 35 sites. The WSP on the ASR-9 is capable of utilizing dual fan-beam estimates of reflectivity and velocity in order to detect low-level features such as gust fronts, wind shear, and microbursts, which would normally be best detectable by a low-scanning pencil-beam radar. An upgrade to the ASR-9 WSP, which is currently ongoing, allows for additional computational complexity in the digital signal processing algorithms compared to previous iterations of the system. This paper will explore ideas for improving velocity estimates, with a focus on dealiasing. A discussion of the unique challenges afforded by the ASR-9's block-stagger pulse repetition time is presented, along with thoughts on possible applications to other systems.

1. Introduction

In the past 20 years, 35 Airport Surveillance Radars (ASR-9s) have been equipped with a Weather Systems Processor (WSP) capable of utilizing opposite-sense circular polarization returns for weather observations (Newell 2000). The WSP makes use of the ASR-9's dual fan-beam antenna in order to generate estimates near ground level, similar to what would be observed with a more traditional pencil-beam weather radar. These ground level estimates are made via beam-weighting techniques, as well as the inference of vertical motion observed over time (specifically for detection of microburst phenomena, which are critical to near-airport operations; Weber (2002)).

The original implementation of the WSP made use of relatively low availability of computational power, resulting in the necessity for simple, efficient calculations. Due to these restrictions, certain features were left out of the final WSP implementation. Ongoing upgrades to the WSP framework will allow for more computational complexity in real-time calculations, meaning the possibility for added capabilities can be explored.

Weber (2002) outlined some of the possible additions to the WSP if upgrades were completed in the future. Of principal concern within these ideas was improvement of velocity estimates. While the microburst detection algorithm (Newell and Cullen 1993) and gust front detection algorithm (Delanoy and Troxel 1993) have shown strong performance since the implementation of the original WSP framework, these algorithms will always be limited by the quality of data fed into them. Since microburst and gust front detection were the critical reasons for the WSP implementation, and their algorithms rely heavily on velocity estimates, this was a prudent area to explore for future improvement.

This paper focuses on experimental methods for improving velocity estimates within the ASR-9 WSP architecture. Specifically, the issue of velocity aliasing in strong microburst and tornado cases near airports was found to flag velocity data as erroneous in the calculation of dual-beam (near-ground) velocity. The belief within the microburst detection algorithm that the data were in error precluded them from being used in the detection algorithm, which has led to missed detections.

Corresponding author address: James M. Kurdzo, Advanced Radar Research Center, University of Oklahoma, 120 David L. Boren Blvd., Suite 4600, Norman, OK 73072.
E-mail: kurdzo@ou.edu

This work was sponsored by the Federal Aviation Administration under Air Force Contract No. FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

Section 2 describes an adapted clustering method which utilizes multiple levels of filtering and thresholding in order to reliably dealias high velocity events within the WSP, despite the challenging pulse structure utilized by the ASR-9 platform. Section 3 shows results of this method with a microburst case and an EF-5 tornado case using simulated IQ data generated from ASR-9 base data. Finally, Section 4 provides discussion of these results, while outlining future work and ongoing work regarding beam-weighting advancements within the WSP.

2. Method

Among the most challenging aspects of velocity dealiasing with the ASR-9 is the pulse structure. This pulse structure exists due to the legacy method used for aircraft detection, which can not be changed for weather observations. The primary mission for the ASR-9 is near-airport aircraft detection, meaning any weather or WSP-related calculations must be made with "as-is" data. Instead of a traditional staggered pulse repetition time (PRT), 3 PRT blocks are used. The first block involves the transmission of 8 pulses with a single PRT (PRT 1). The second block transmits 10 pulses with a different single PRT (PRT 2). Finally, the third block transmits 8 pulses with the first PRT (PRT 1). This results in 26 total pulses per azimuth, with two different PRTs, and at blocks of differing length.

Clearly, this results in a pulse pair pattern not suitable for staggered-PRT dealiasing. However, since the pulse structure is predictable, other methods may be suitable. For this paper, an adaptation of the clustering algorithm is made for use on the ASR-9 WSP. The clustering algorithm is most commonly known in weather radar for its planned use within the TDWR platform (Cho 2005). Dealiasing via clustering is performed by grouping PRT blocks (usually two), extrapolating the returned phase shifts to all possible aliased velocities, and running a moving window over the possibilities in order to find the area with the least error.

Clustering, like any estimation technique, works significantly better with dependable observations. In general, stochastic theory shows that more samples result in better estimates. Unfortunately, with PRT blocks as small as 8 pulses, the ASR-9 does not provide large enough blocks to use traditional clustering for dealiasing. The lack of reliable estimates results in a high dealiasing error rate, defeating the overall purpose of the technique.

In order to overcome these issues, an adaptation of the clustering method was required for use within the WSP. First, in order to avoid widespread error, application of clustering in only localized areas where aliasing is likely can make for significantly better performance. This stage of the adapted clustering method is called the "filtering" phase. Second, in order to correct errors that do happen in areas determined by filtering to be candidates for aliasing, phase two implements a local variance-based speckle correction. This stage is called the "cor-

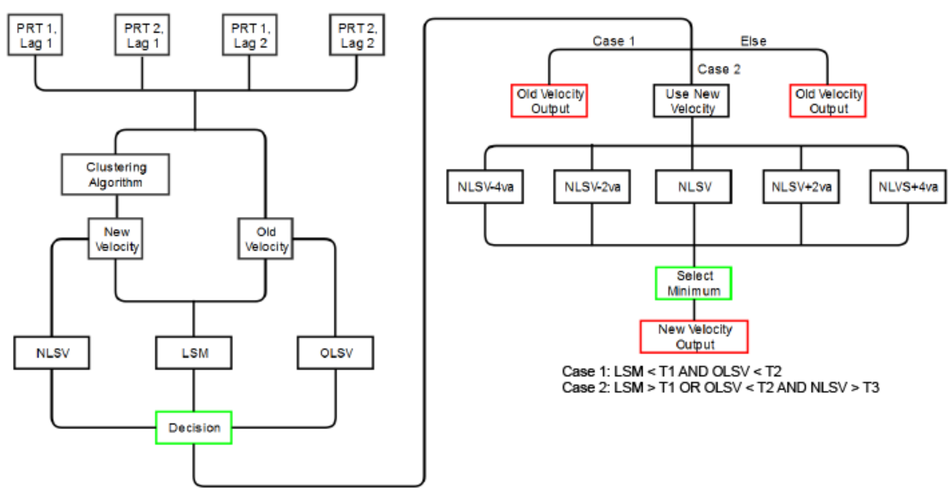

FIG. 1. Adapted clustering algorithm flowchart. T1, T2, and T3 are pre-determined threshold levels.

rection" phase, and is only implemented in areas determined to meet a threshold for possible aliasing based on the filtering phase. It is important to note that WSP processing can only take place along radials, meaning that each step described from this point forward applies along-radial only (not across azimuths). An overview flowchart of the adapted algorithm is provided in Fig. 1.

a. Filtering Phase

In the first phase of the algorithm, PRT blocks 1 and 3 are combined in order to provide better initial velocity estimates, while PRT block 2 is left alone. The combination of block 1 and 3 results in 16 total pulses, while block 2 supplies 10 total pulses. Each set (set "A" with 16 pulses, and set "B" with 10 pulses) is used to calculate a velocity estimate at lag 1 and lag 2. This results in four velocity estimates, two at each PRT. It is important to note that while lag 2 cuts the aliasing velocity in half (offering no extra clustering points in theory), slightly improved performance was observed in our data by using both lags. While the theory behind this is still under development and beyond the scope of this paper, we believe that this has to do with the ability to use a wider clustering window (Trunk and Brockett 1993), possibly resulting in fewer errors.

After the four velocity estimates are calculated, they are fed into a standard clustering algorithm in order to determine a dealiased velocity. This single number per gate is referred to as the "new" velocity. The four velocity estimates are averaged together in a traditional calculation of the "old" velocity estimate, which is representative of what would be displayed in the current WSP iteration. In practice, this would mean that only one of these calculations would result in added computational complexity, since the "old" estimate already exists within the WSP framework.

Once the new velocity and old velocity estimates are calculated along an entire radial, three filters are implemented along the radial. These filters are not particularly useful individually, but when combined, they paint a critical picture for the likelihood of aliasing. First, a local spatial variance filter is applied to the old velocity estimates. This filter finds the edges of folding along a radial, and also detects thin lines (possible gust fronts which we do not want to disrupt). Second, a local spatial variance filter is applied to the new velocity estimates. This filter returns high values in heavily folded areas, telling the algorithm that aliasing is ongoing along the radial. Finally, a local spatial mean filter is applied to the difference between the old and new estimates. This "differential spatial mean" filter returns high values in areas of strong folding, but does not return high values for edges or thin lines.

Through extensive experimentation (and in the future, the possibility of optimization), a series of thresholds were found for common aliasing features within each filter. When these thresholds are met in all three filters, the gate is flagged as contaminated, and dealiasing via the clustering algorithm is applied (meaning the "new" velocity is kept for the gate). If the threshold criteria are not met, the gate is assigned the original "old" velocity. It is important to note that these thresholds can be chosen to be conservative or aggressive. For the purposes of this paper, relatively conservative values were chosen, meaning that only areas of obvious folding were touched by the algorithm in the final results.

b. Correction Phase

Despite the clustering algorithm only being applied in areas of strong folding via the thresholds set in the filtering phase, errors are still likely in folded areas due to the relatively low number of samples used for each estimate. Therefore, the correction phase is necessary to clean up errors which occurred in clustering. Correction is only applied where folding was detected (i.e., where clustering was utilized).

Correction is achieved via a local spatial variance filter and a sliding window. The sliding window tests for variance differences when the center gate is changed to different possible aliasing velocity shifts. The center gate is changed to the possible aliased velocities for each PRT, in both directions, and for multiple folds of intensity. If any change results in a smaller variance for the entire sliding window, the center gate is changed to the velocity which resulted in the smallest variance. It is important to choose a small enough sliding window for speckle correction so as to not have variance bias from areas too far from the folding. A small window (seven-point filter) was found to be most effective for this type of correction.

For the dual-beam velocity estimates, the same process is used in the filtering and correction phases, except the local spatial variance filter within the correction phase uses the dual-beam velocity at the center point, and dealiased low-beam velocity at the surrounding points within the window. This is because dual-beam velocity is designed to determine low-level velocities, and the already-corrected low-beam velocities provide a sufficient comparison of variance for proper error correction.

3. Results

While multiple cases and simulations were used for threshold determinations, only two cases are shown in this paper for brevity. Each case was tested by simulating IQ data from available low-beam and high-beam base data (power and radial velocity). A noise-added Doppler spectrum was applied to base data to simulate time-series data, which could be generated based on the different PRT blocks used with the ASR-9.

a. Case 1: 28 June 2000 Microburst, Austin, TX

Microbursts are a significant hazard to commercial and private aircraft (Wilson et al. 1984). The need to detect and warn for impending microbursts and high wind activity precipitated the development of the TDWR platform, as well as the WSP addition to the ASR-9 at medium-density airports. On 28 June 2000, an exceptionally strong microburst event occurred just to the west of the Austin-Bergstrom International Airport in Austin, Texas. Due to the strength of low-level winds as the microburst made contact with the ground, radial velocities estimated by the ASR-9 in Austin were aliased. Due to the aliasing, an erroneous data flag was triggered within the microburst detection algorithm, and the microburst was not detected. Given the speed of ground-level winds near the airport, this missed detection represents a critical area for improvement in WSP performance.

Fig. 2a shows the low-beam power observed at the time of ground impact. The microburst is located on the east side of the highest reflectivity

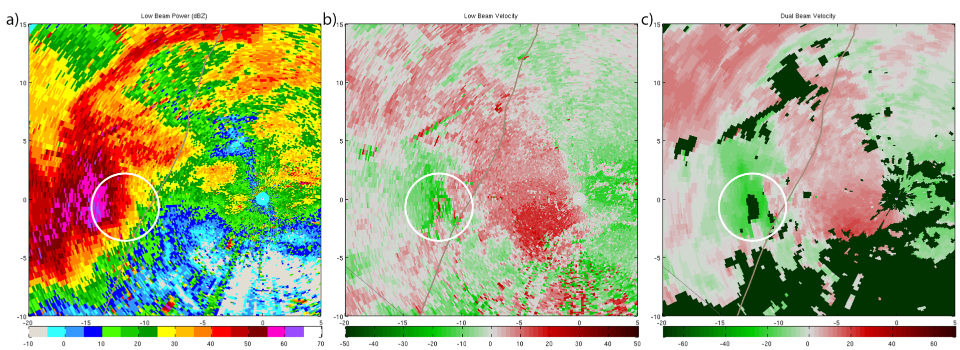

FIG. 2. 28 June 2000 Austin, TX microburst ASR-9 data. (a) Low-beam power (dBZ), (b) low-beam velocity ($\mathrm{m\,s^{-1}}$), (c) dual-beam velocity ($\mathrm{m\,s^{-1}}$)).

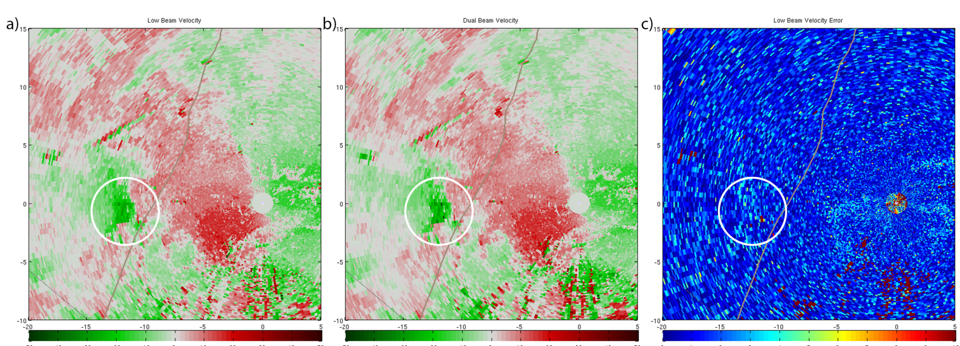

FIG. 3. Dealiased data from Fig. 2. (a) Low-beam velocity, (b) dual-beam velocity, (c) low-beam velocity absolute error (all $\mathrm{m\,s^{-1}}$)).

core, pointing towards the airport. Fig. 2b shows the original low-beam velocity estimation, indicating velocity folding. Fig. 2c shows the original dual-beam velocity estimation, which shows an area of censored velocity data in the immediate area of the microburst. These censored data were delivered to the microburst detection algorithm, resulting in the missed detection.

It should be noted that the dual-beam velocity product is qualitatively different than the raw velocity estimates from the low and high beams. This is partially due to the use of a beam weighting algorithm which estimates the low-level winds that a pencil beam scanning at a low elevation angle would observe. The beam weighting described in Weber (2002) utilizes spatial averaging, which accounts for a slightly smoother look to the data. In addition, the dual-beam calculation censors data which are believed to be erroneous. This can be due to low signal-to-noise ratio, ground clutter, aliasing, or other reasons. Censored velocities are displayed in dark green in the raw dual-beam velocity shown in Fig. 2c. No censored velocities are shown in re-calculated dual-beam velocities (Figs. 3b and 5c).

After applying the method described previously, the resulting dealiased low-beam velocity, dealiased dual-beam velocity, and low-beam velocity absolute error are shown in Figs. 3a, 3b, and 3c, respectively. High-beam velocity is not shown due to the lack of aliasing (which is a common signature during microburst impact). The low-beam velocities show a successfully dealiased microburst signature, and the dual-beam velocities are no longer flagged as erroneous. In fact, the dual-beam signature is generally

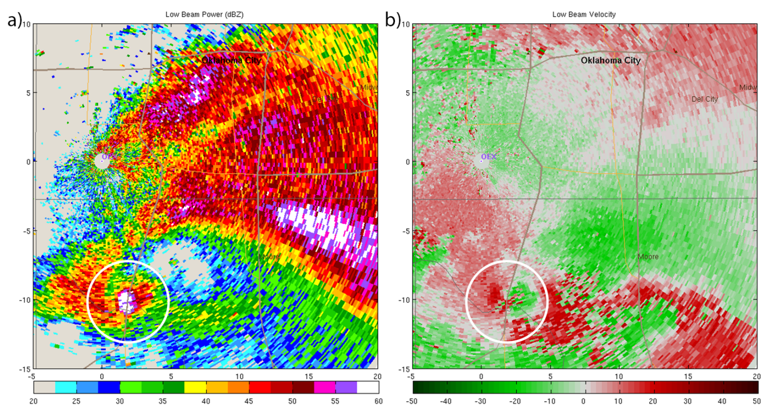

FIG. 4. 20 May 2013 Oklahoma City, OK EF-5 tornado ASR-9 data. (a) Low-beam power (dBZ), (b) low-beam velocity ($\mathrm{m\,s^{-1}}$).

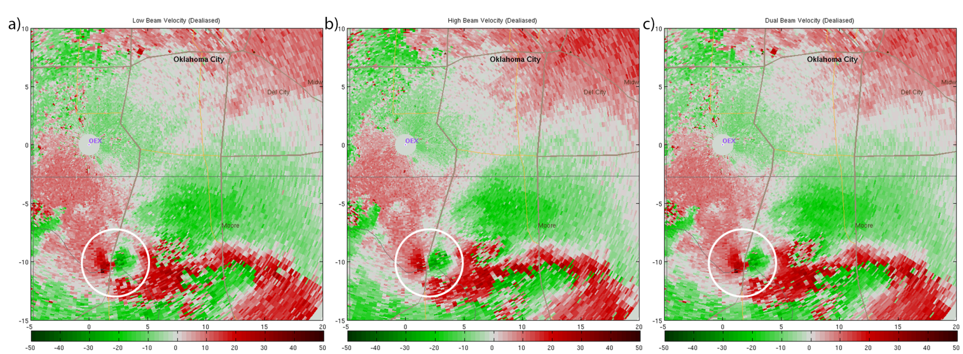

FIG. 5. Dealiased data from Fig. 4. (a) Low-beam velocity, (b) high-beam velocity, (c) dual-beam velocity (all $\mathrm{m\,s^{-1}}$).

stronger, which is a promising signature which we expect to see at the lowest levels during microburst impact. The error within the low-beam estimates is calculated using manually-dealiased data, and in general, does not display a significantly increased concentration of error near the microburst signature.

b. Case 2: 20 May 2013 EF-5 Tornado, Moore, Oklahoma

One of the most challenging tests for a dealiasing algorithm is the estimation of strong tornadic winds. On 20 May 2013, the Oklahoma City ASR-9 observed an EF-5 tornado in Moore, Oklahoma with winds over 90 $\mathrm{m\,s^{-1}}$ near the surface. At the observation elevations, winds as high as 60 $\mathrm{m\,s^{-1}}$ were present, which resulted in aliased velocity estimates within both the low and high beams. Fig. 4a shows the low-beam power estimate, while Fig. 4b shows the original low-beam velocity estimate.

After application of the adapted clustering algorithm, de-aliased velocities at both low and high beams, as well as the dual-beam estimate are shown in Figs. 5a, 5b, and 5c, respectively. While the results are not perfect, a significant recovery of radial velocities is apparent. Additionally, the lack of erroneous flagging within the dual-beam algorithm is critical to many applications the WSP is used for.

4. Conclusions and Future Work

As the ASR-9 WSP is being upgraded to include more computational capability, it is important that we explore areas where additional computing power can be best utilized. It has been shown that a viable area to explore regarding enhancements to the WSP architecture is Doppler velocity estimation. This paper has detailed an adaptation of the clustering algorithm for velocity dealiasing which is capable of correcting the inherent errors associated with a small number of samples per pulse block in the ASR-9.

Ongoing work with these data includes experimentation with adaptive dual-beam weight estimation. More accurate determination of dual-beam, low-level winds will only increase the success rate of critical algorithms such as the microburst detection algorithm and gust front detection algorithm.

Additionally, we plan to implement a retrofitted version of this dealiasing method on the University of Oklahoma Advanced Radar Research Center's PX-1000 X-band polarimetric transportable weather radar (Cheong et al. 2013) in the near future. With promise for improved tornadic velocity dealiasing with advanced PRT structures, there may be additional uses for this method, especially in the expedition of manual dealiasing of tornadic data for research purposes.

Acknowledgments.

The authors would like to thank Will Pughe, John Morgan, Mark Weber, Paul Morin, Aaron Tuttle, and Ilana Steinhorn for their collaborations, ideas, and technical/data support.

REFERENCES

Cheong, B. L., R. Kelley, R. D. Palmer, Y. Zhang, M. Yeary, and T.-Y. Yu, 2013: PX-1000: A solid-state polarimetric X-band radar and time-frequency multiplexed waveform for blind range mitigation. *IEEE Trans. Instrum. Meas.*, **62 (11)**, 3604–3072.

Cho, J. Y. N., 2005: Multi-PRI signal processing for the terminal doppler weather radar. Part II: Range–velocity ambiguity mitigation. *J. Atmos. Oceanic Technol.*, **22 (10)**, 1507–1519.

Delanoy, R. L. and S. W. Troxel, 1993: Machine intelligent gust front algorithm. Tech. Rep. ATC-196, MIT Lincoln Laboratory.

Newell, O., 2000: ASR-9 weather systems processor software overview. Tech. Rep. ATC-264, MIT Lincoln Laboratory.

Newell, O. and J. A. Cullen, 1993: ASR-9 microburst detection algorithm. Tech. Rep. ATC-197, MIT Lincoln Laboratory.

Trunk, G. and S. Brockett, 1993: Range and velocity ambiguity reduction. *1993 IEEE National Radar Conf.*, Lynnfield, MA, IEEE, 146–149.

Weber, M. E., 2002: ASR-9 weather systems processor (WSP) signal processing algorithms. Tech. Rep. ATC-255, MIT Lincoln Laboratory.

Wilson, J. W., R. D. Roberts, C. Kessinger, and J. McCarthy, 1984: Microburst wind structure and evaluation of Doppler radar for airport wind shear detection. *J. Climate Appl. Meteor.*, **23 (6)**, 898–915.