P5.2 MOVING CLUTTER SPECTRAL FILTER FOR TERMINAL DOPPLER WEATHER RADAR

John Y. N. Cho
MIT Lincoln Laboratory, Lexington, Massachusetts

1. INTRODUCTION

Detecting low-altitude wind shear in support of aviation safety and efficiency is the primary mission of the Terminal Doppler Weather Radar (TDWR). The wind-shear detection performance depends directly on the quality of the data produced by the TDWR. At times the data quality suffers from the presence of clutter. Although stationary ground clutter signals can be removed by a high-pass filter, moving clutter such as birds and roadway traffic cannot be attenuated using the same technique because their signal power can exist anywhere in the Doppler velocity spectrum. Furthermore, because the TDWR is a single-polarization radar, polarimetry cannot be used to discriminate these types of clutter from atmospheric signals.

The moving clutter problem is exacerbated at Western sites with dry microbursts, because their low signal-to-noise ratios (SNRs) are more easily masked by unwanted moving clutter. For Las Vegas (LAS), Nevada, the offending clutter is traffic on roads that are oriented along the radar line of sight near the airport. The radar is located at a significantly higher altitude than the town, improving the visibility to the roads, and giving LAS the worst road clutter problem of all TDWR sites. The Salt Lake City (SLC), Utah, airport is located near the Great Salt Lake, which is the biggest inland staging area for migrating seabirds in the country. It, therefore, suffers from bird clutter, which not only can obscure wind shear signatures but can also mimic them to trigger false alarms. The TDWR "dry" site issues are discussed in more detail by Cho (2008).

In order to mitigate these problems, we developed a moving clutter spectral filter (MCSF). In this paper we describe the algorithm and present preliminary test results.

2. THE PROBLEM

Figure 1 illustrates one of the difficulties with not filtering out moving clutter signals from TDWR data. A widespread coherent bird flight event like this has two of the characteristics of a microburst—outflow of heightened reflectivity away from a central source and velocity divergence along the radials. Thus, a microburst detector is in danger of issuing a false alert in this case.

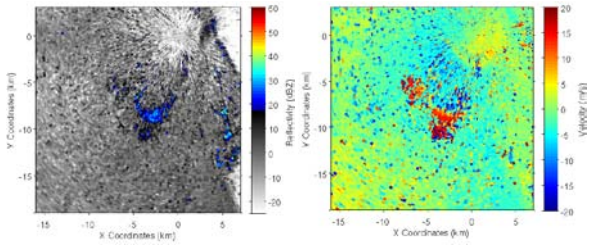

Figure 1. Example of bird clutter observed with the SLC TDWR at 0.5° elevation. Left panel shows reflectivity, right panel shows radial velocity.

Figure 2 shows Doppler velocity spectra vs. range for a microburst event (left) and a case with birds in many different range gates (right). Not only are there bird signatures that are isolated in range-Doppler, there is a feature in range gates ~40-50 that is quasi-continuous and looks similar to a strong wind shear. The challenge is to filter the bird signals in the right-hand case but not the microburst in the left-hand case.

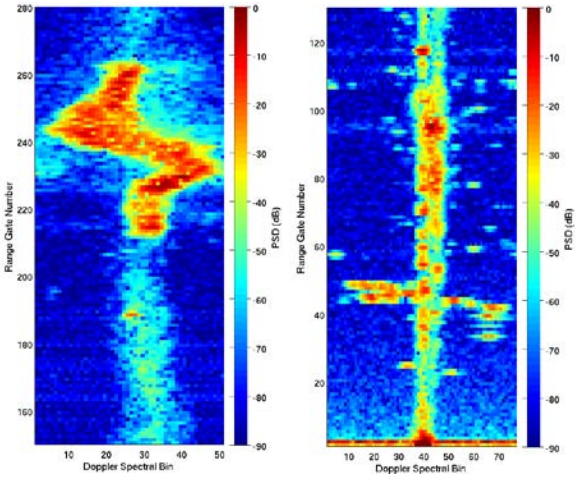

Figure 2. Doppler velocity spectra vs. range for a microburst case observed with the Program Support Facility (PSF) TDWR in Oklahoma City, OK at 0.3° elevation (left), and a widespread bird contamination case observed with the SLC TDWR at 0.5° elevation (right).

Bats can also cause similar problems for radar wind-shear detection when they leave their roosts en masse at dusk.

Road and rail traffic are also sources of moving clutter. In this case, the range-azimuth cells affected are known a priori, so the current method of dealing with the problem is to periodically generate a clutter residue map (CREM) and then censor the base data where the

---

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 U.S. Government.

Corresponding author address: John Y. N. Cho, MIT Lincoln Laboratory, 244 Wood St., S1-539, Lexington, MA 02420-9185; e-mail: jync@ll.mit.edu.

reflectivity does not exceed a certain amount over the CREM reflectivity. The LAS road clutter problem is shown in Figure 3. Note that some of the clearest road echoes line up with the radar line-of-sight radials. This phenomenon is also observed at other sites, and we call it the building canyon effect. In an urban environment roads are lined with buildings, so traffic is often not visible to the radar unless the beam shoots down along the road itself.

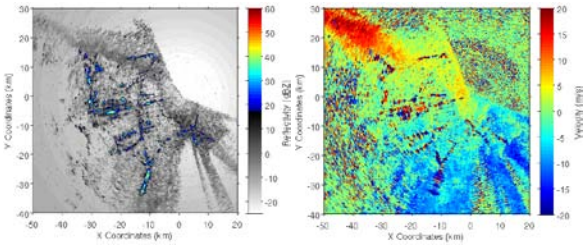

Figure 3. LAS road clutter at 0.8° elevation. Left panel shows reflectivity, right panel shows radial velocity. The large wedge of blank reflectivity to the northeast of the radar is due to terrain blockage.

Airplanes, of course, are also moving clutter. Because they are isolated targets, a point target filter deals with them effectively. Although the current operational algorithm merely censors these points, we plan to interpolate the base data across these points in the next major radar data acquisition (RDA) system software revision. The filtering algorithm presented in this paper should be an even better solution in the future.

Finally, there are other moving targets that appear in TDWR data such as sea clutter and the spinning blades of a wind power turbine. The latter is expected to be an ever-growing source of clutter that is still in search of a solution. The work presented in this paper does not specifically address these phenomena, as the focus has been on the data quality issues at SLC and LAS.

3. MOVING CLUTTER SPECTRAL FILTER

Stationary clutter filters are typically applied on one range-azimuth dwell data at a time. Recognizing that more contextual information is necessary to identify moving clutter, techniques were developed that utilize multiple range gates of data. The point target filters employed on both the TDWR and the Weather Surveillance Radar-1988 Doppler (WSR-88D, more commonly known as NEXRAD) look for targets with reflectivity much higher than the neighboring range gates. This type of filter can remove aircraft and isolated avian signals, but the weather returns in the same cells are also lost.

A more modern approach filters data in the two-dimensional (2D), Doppler velocity spectrum vs. range, domain (e.g., Sasaoka 2003; Meymaris 2007). Because moving clutter spectral signals tend to be spectrally compact and discontinuous in range, these techniques would look for a range-continuous signal (weather) and discard other spectral components (clutter). (Of course, stationary ground clutter can also be continuous in range, so it would have to be filtered out first.) Two drawbacks to this approach are that it is computationally intensive relative to traditional weather radar signal processing, and that the output base data are smoothed in range.

With the development of a new RDA system (Cho et al. 2005), which is currently running operationally at LAS and SLC, there is considerable spare computational capacity for the implementation of future algorithms. The scalable architecture easily allows for additional (and/or upgraded) central processing units (CPUs) to be installed to meet further computational demand.

With regard to the second drawback of the existing 2D spectral processing approach for moving clutter filtering, spatial smoothing of the data runs counter to the TDWR's mission of detecting wind shear. This is the reason why we decided to develop our own moving clutter spectral filter for the TDWR.

First, we will give a brief outline of the MCSF algorithm, and then delve into more details.

To set up the processing environment, we form a Doppler velocity power spectrum at each range gate with the usual window-and-DFT approach. These spectra are stored in a 2D matrix (Figure 2, right). The following steps are then followed.

1. Remove stationary ground clutter with the spectral filter of choice (Figure 5).
2. Remove positive power anomalies along range at each spectral bin in a manner similar to a point target filter (Figure 6). Clutter signals that are not continuous in range are removed in this way.
3. After suitable data massaging, determine the number and location of spectral modes (statistically significant peaks) at each range gate. A measure called the normalized circular excess mass (NCEM) is used to winnow out "insignificant" peaks in the spectra.
4. Select one mode at each range gate such that the path connecting the modes has the shortest possible overall distance (Figure 7). This process favors retaining globally continuous (in range) signals over only locally continuous features. The assumption is that atmospheric signals will be continuous over longer ranges than moving clutter features.
5. For modes other than the ones selected by the shortest connecting path, reduce the number of "extra" modes using a stricter NCEM threshold. This step is needed to reduce the incidence of these "extra" modes being part of the atmospheric spectrum.
6. Any leftover "extra" modes are deemed to be moving clutter and are removed by filling in with the computed spectral "noise" floor (Figure 8).

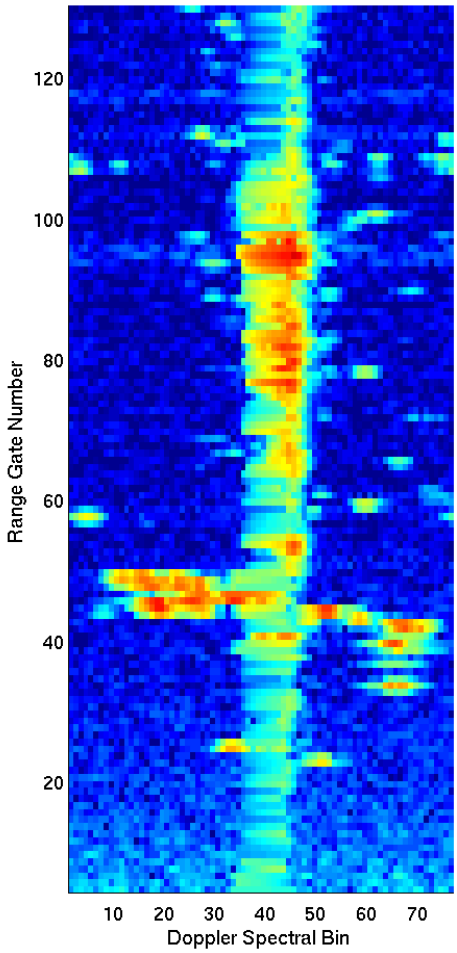

Figure 5. Stationary ground clutter is removed from the spectra shown in Figure 2, right.

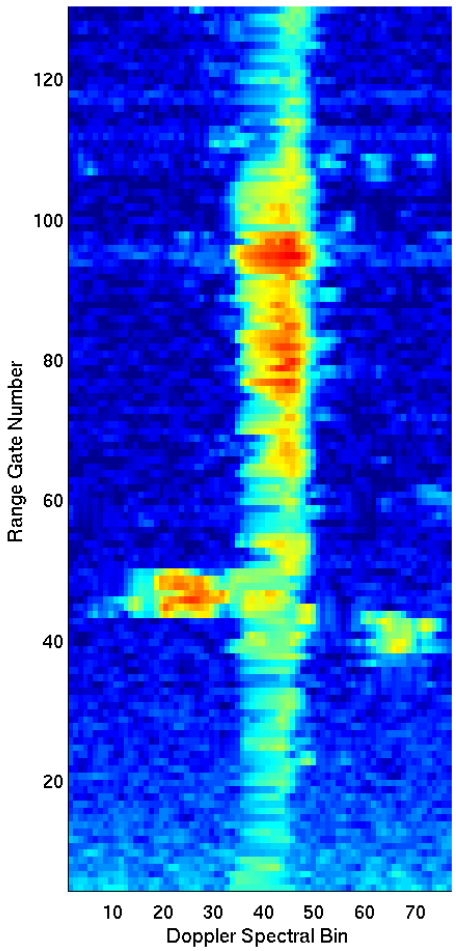

Figure 6. Isolated spectral peaks are removed from the spectra shown in Figure 5.

7. Carry on with the rest of the usual signal processing for base data generation.

Let us now elaborate on each step. (Note: The term "circular" used in this paper refers to the fact that the Doppler velocity spectrum aliases, or wraps around, at the end points.) There is nothing special about the first step of stationary clutter filtering. One of various different techniques can be used. However, note that the Gaussian model adaptive processing (GMAP) filter (Siggia and Passarelli 2004) assumes the absence of spectral signal other than a stationary ground clutter peak and a Gaussian weather component. Therefore, the presence of moving clutter signals distorts the weather moment estimation and can spuriously widen the "weather" spectrum within the filled ground clutter gap. We, thus, recommend against the use of Gaussian fitting and clutter-gap filling in the suspected presence of moving clutter.

The isolated spectral peak filter in step 2 is similar to the TDWR point target filter, except that it runs along each spectral bin vs. range, and the number of neighbors considered in the spectral dimension expands with range away from the range gate of interest. This expansion prevents filtering of cases where the spectral-range signature has a shape like "\" "/" ">" or "<". So only features shaped like "—" or "•" are filtered. To be more specific, at range gate i and spectral index j, compare the spectral power $P(i,j)$ to $PP(i-m,j)$ and $PP(i+m,j)$, and mark $P(i,j)$ for removal if it is stronger than both PPs by a threshold $T(m)$. The gate difference m goes from 1 to 3, and the threshold values used currently are $T(1)$ = 8 dB, $T(2)$ = 10 dB, and $T(3)$ = 17 dB. Removal is accomplished by linear interpolation over range. PP is the peak value of spectral power taken over bins $j - n$ to $j + n$ (circularly wrapped around to the other side of the spectrum if needed). As stated earlier, this spectral neighborhood expands with range such that n = 2 for m = 1, n = 4 for m = 2, and n = 6 for m = 3.

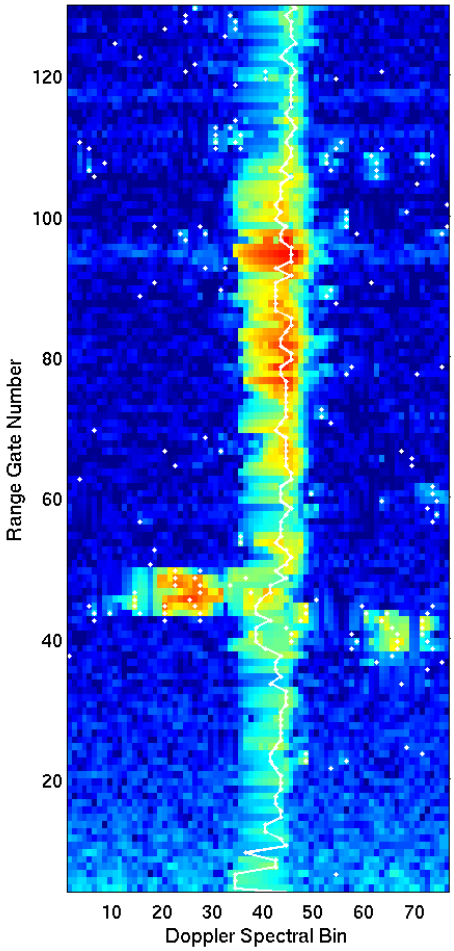

Figure 7. Spectral modes (white dots) are identified at each range gate, and the shortest path (white line) connecting them in range is determined.

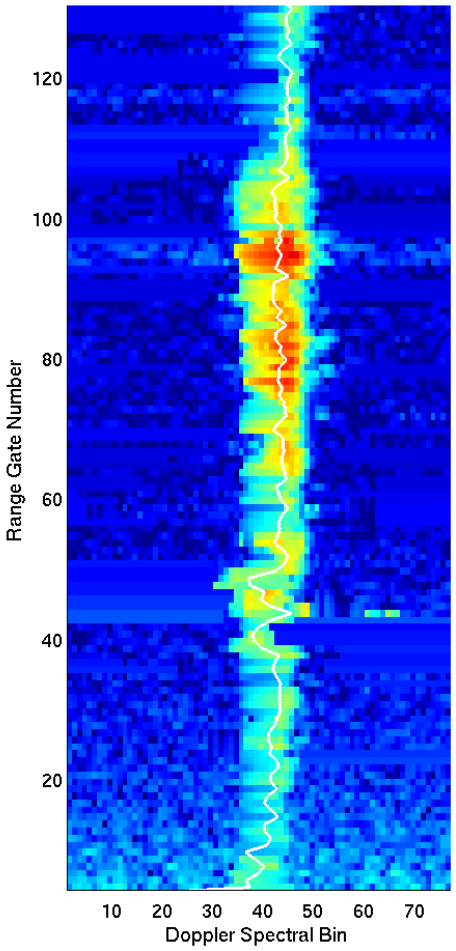

Figure 8. Undesirable spectral modes are eliminated. The first moment (white line) is shown for reference.

In step 3, in order to avoid identification of spurious spectral peaks as modes, we first smooth the spectra over the spectral bins. For this purpose we use a Gaussian kernel of width $0.3N_{DFT}/(2v_a)$, where N_{DFT} is the number of DFT bins and v_a is the Nyquist velocity. We also perform a three-point triangular moving average over range; this minimal smoothing is the only averaging that we do in the range dimension. Then we estimate the spectral noise floor (power per bin) (e.g., Hildebrand and Sekhon 1974).

Now, for a spectrum at a range gate, we locate all the local maxima. Then for each maximum, we compute the NCEM. The excess mass is essentially the area under the peak bounded by the curve itself and a horizontal-line lower limit defined by the highest value taken from the following set—either of the neighboring minima or the spectral noise floor. (If a local maximum is below the noise floor, it is deleted from the maxima list.) A nice visual illustration of the excess mass is given in Figure 2 of a paper by Fisher and Marron (2001). NCEM is the excess mass divided by the per bin spectral noise power, with the distribution curve wrapping around at the ends.

Local maxima with NCEM below a threshold (currently 5 dB) are eliminated. Elimination of a maximum expands the size of a neighboring mode, so the NCEMs are recalculated for the remaining maxima. The process of NCEM computation and maxima elimination is repeated until no more reduction takes place. The remaining modes are deemed to be significant. Finally, the pointers to the modes are adjusted to be at the center of mass of the modes, not at the peaks.

In step 4, one mode at each range gate is selected such that the path connecting the modes has the shortest possible overall distance from the first to the last gate. If there is no mode in a range gate, that gate is skipped. "Distance" is defined to be $\alpha(\Delta i)^2 + (\Delta j)^2$, where Δi is the range gate difference, Δj is the circular difference in spectral bin numbers between modes, and α is a constant determining the relationship between range distance and spectral bin distance. Currently α is set to one; however, the algorithm is not very sensitive to α.

The shortest overall path is determined using Dijkstra's method (Dijkstra 1959). We want the endpoints to have only one mode, so if an endpoint has more than one

mode, we compute the circular median (Fisher 1993) of the mode locations over ten range gates from the endpoint and select the mode closest to the median.

In step 5, we attempt to consolidate the spectral modes that are not part of the shortest path. The reason for doing this is that the initial mode detection process allows for fairly minor modes to be kept, and thus the weather spectrum may be tagged with more than one mode. We want it to be tagged with only one mode so that the final mode elimination step will not erode the weather spectrum. To do this, we repeat the mode identification process of step 3, but with the NCEM threshold raised to 20 dB and the mode that is part of the shortest path protected from elimination.

Finally, in step 6, the modes that are not part of the shortest overall connection path are eliminated by filling them in with the spectral noise floor power.

4. RESULTS

We now present some initial testing results of the MCSF algorithm. Figure 9 is a case from SLC where wide-spread bird flock flights contaminated the data. The upper two plots show the results without MCSF, and the lower two plots show the corresponding results with MCSF. The filtering works very well in this case, and we have many other examples of bird clutter that are cleaned up by MCSF in a similar fashion, albeit with small amounts of unfiltered moving clutter in some instances.

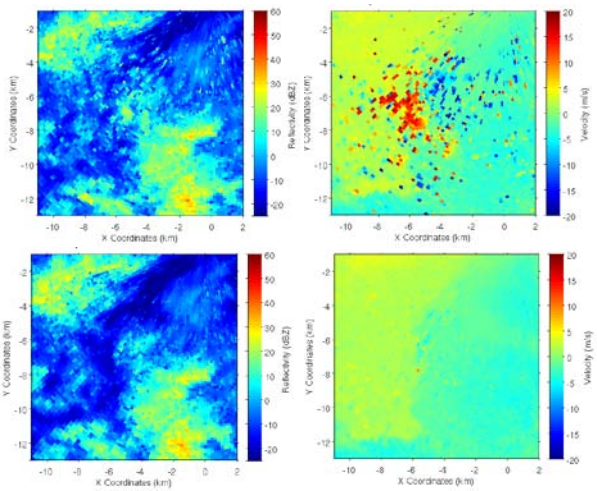

Figure 9. Reflectivity (left) and radial velocity (right) plots for data processed without MCSF (top) and with MCSF (bottom). This was a bird clutter case observed with the SLC TDWR at 0.5° elevation.

Of course, the ability to filter out moving clutter is worthless to the TDWR if wind shear events are also removed or altered. Figure 10 shows a microburst without (top) and with (bottom) MCSF. There is very little change between the two processed results, and MCSF preserves the strong velocity gradients. MCSF also does well on a few other wind shear cases that we have collected so far, but we are in the process of capturing many more cases for testing.

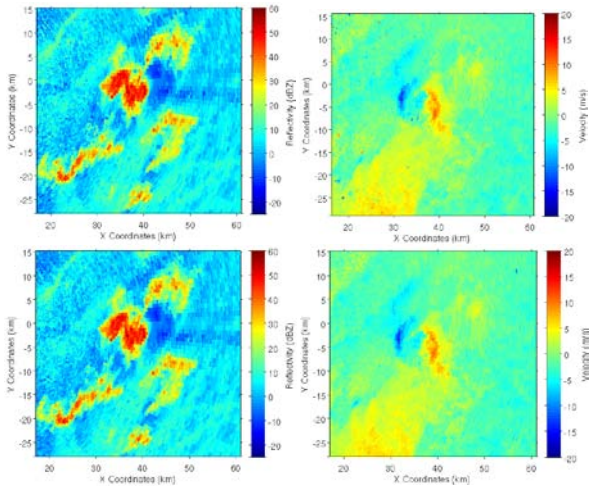

Figure 10. Reflectivity (left) and radial velocity (right) plots for data processed without MCSF (top) and with MCSF (bottom). This was a microburst observed with the PSF TDWR at 0.3° elevation.

The results for road clutter were only partially successful, with MCSF filtering out traffic returns from roads that were not parallel to the radar line of sight. The roads that lined up with the radar radials, however, were largely not filtered. MCSF fails in these cases, because the moving clutter extends continuously in range for long stretches at a time. Thus, MCSF is not a complete solution for road clutter, but it should reduce the size of CREM areas where roads currently force data censoring.

5. SUMMARY DISCUSSION

With the introduction of a new open and scalable RDA to operational TDWRs, enhanced signal processing algorithms can be successively implemented to improve data quality. The first RDA algorithm upgrade focuses on mitigating range-velocity ambiguity (Cho et al. 2005). Dubbed Build 2 (emulation of the legacy processing algorithms is Build 1) it is currently undergoing acceptance testing by the FAA. The MCSF algorithm discussed in this paper may be part of a subsequent software upgrade (Build 3).

There is still considerable testing and development that needs to be conducted before MCSF is deemed acceptable for operational use. For example, MCSF has only been tested on moving clutter with Build 1 data, which uses the legacy transmission scheme of constant pulse repetition interval (PRI) and pulse phase. Starting with Build 2, the transmission scheme will become more complex, with a mix of various multiple PRI and pseudo-random phase-code processing techniques. On surface scans, which are used for wind-shear detection and where MCSF will be most needed, the default mode will be phase-code processing with the PRI changing every dwell. Therefore, the interaction of MCSF with the phase-code processing, which also takes place in the spectral domain, needs to be investigated and the results optimized. There are other new processing features in Build 2 that need to be tested together with

MCSF. However, the collection of Build 2 data at the problematic sites (SLC and LAS) awaits the acceptance and operational fielding of Build 2. So for now, we are only able to investigate MCSF-Build 2 interaction using data from the non-operational radars in Oklahoma City, where the moving clutter problem is not nearly as severe as in SLC and LAS.

Due to the simultaneous processing of spectra across many range gates, the real-time implementation of this algorithm may require changes to the RDA software architecture, which is currently designed to process chunks of range gates in parallel. The computational load will also increase significantly, so it must be determined whether the present CPUs will be able to keep up or whether a hardware upgrade would be needed.

MCSF is certainly not foolproof. If there is no underlying weather signal, it will not be able to filter out the moving clutter. If the 2D spectral-range signature of the moving clutter extends continuously over a long range, MCSF may select to preserve it instead of the weather signal. As we have seen, roads that run parallel to the radar line of sight will not be filtered well. However, the testing so far shows that it can dramatically improve the base data quality during periods when the current signal processing algorithm fails to do anything with moving clutter. As long as it is proven to be robust against the degradation of wind shear information, it would be a valuable addition to the TDWR arsenal in fighting moving clutter.

6. ACKNOWLEDGMENT

I would like to thank Amrita Masurkar for implementing and evaluating the Sasaoka and Meymaris algorithms. I would also like to thank Mike Donovan for identifying test cases, Rich Chafe (SLC) and Pat Hinkle (LAS) for collecting the data, and Jeff Simpson and Robert Schaefer for assistance with data handling.

7. REFERENCES

Cho, J. Y. N., G. R. Elkin, and N. G. Parker, 2005: Enhanced radar data acquisition system and signal processing algorithms for the Terminal Doppler Weather Radar. Preprints, *32nd Conf. on Radar Meteorology*, Albuquerque, NM, Amer. Meteor. Soc., P4R.8, http://ams.confex.com/ams/pdfpapers/96018.pdf.

Cho, J. Y. N., 2008: TDWR dry site problem assessment and RDA software Build 3 recommendation. *Project Memorandum 43PM-Wx-0107*, MIT Lincoln Laboratory, Lexington, MA, 32 pp.

Dijkstra, E. W., 1959: A note on two problems in connexion with graphs. *Numerische Mathematik*, **1**, 269-271.

Fisher, N. I., 1993: *Statistical Analysis of Circular Data*. Cambridge Univ. Press, New York, 277 pp.

Fisher, N. I., and J. S. Marron, 2001: Mode testing via the excess mass estimate. *Biometrika*, **88**, 499-517.

Hildebrand, P. H., and R. S. Sekhon, 1974: Objective determination of the noise level in Doppler spectra. *J. Appl. Meteor.*, **13**, 808-811.

Meymaris, G., 2007: The use of spectral processing to improve radar spectral moment. Preprints, *23rd Conf. on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, San Antonio, TX, Amer. Meteor. Soc, 8A.4, http://ams.confex.com/ams/pdfpapers/119629.pdf.

Sasaoka, M., 2003: Improvement of wind speed estimation with boundary layer radar using a grouping algorithm. *Tenki*, **50**, 161-174.

Siggia, A., and R. E. Passarelli, Jr., 2004: Gaussian model adaptive processing (GMAP) for improved ground clutter cancelation and moment estimation. Preprints, *Third European Conf. on Radar in Meteorology and Hydrology*, Visby, Sweden, Copernicus Gesellschaft, 67-73.