Proceedings of ERAD (2004): 67–73 © Copernicus GmbH 2004

ERAD 2004

Gaussian model adaptive processing (GMAP) for improved ground clutter cancellation and moment calculation

A. D. Siggia and R. E. Passarelli, Jr.

SIGMET Inc., 2 Park Drive, Unit 1, Westford, MA 01886, USA

Abstract. The new generation of signal processors, such as the SIGMET, Inc. RVP8, provides greater processing power and programming flexibility than could be achieved on previous systems. This provides the opportunity for greatly improved algorithms as opposed to simple pulse-pair processing with a fixed-width IIR clutter filter. The new processors allow the implementation of algorithms that adapt to the weather and clutter that is present, often repeating the algorithm several times with various parameter settings until an optimal result is achieved. GMAP is a frequency domain approach that uses a Gaussian clutter model to remove ground clutter over a variable number of spectral components that is dependent on the assumed clutter width, signal power, Nyquist interval and number of samples. A Gaussian weather model is then used to iteratively interpolate over the components that have been removed, if any, thus restoring any overlapped weather spectrum with minimal bias caused by the clutter filter. GMAP uses a DFT rather than an FFT approach to achieve the highest possible spectrum resolution. The algorithm is first performed with a Hamming window and then, based on the outcome, the Hamming results are kept or the algorithm is repeated with either the rectangular or Blackman window. This allows the least aggressive spectrum window to be used, depending on the strength of the ground clutter, to minimize the negative impact of more aggressive windows on the variance of the moment estimates. The technique is fully implemented and operational on the SIGMET RVP8 and has been evaluated in comprehensive tests by the US WSR88D ORDA and found to meet all NEXRAD requirements.

1 Introduction

With the advent of low-cost, flexible, high-speed processors, it is now possible to perform algorithms that would have been impractical to implement even ten years ago. In the past, most weather radar processors have been built using the approach of a fixed notch-width IIR clutter filter followed by time-domain autocorrelation processing (so-called pulse-pair processing). These techniques require minimal storage and very few computational MAC's (multiply accumulate steps) per pulse per range bin. The algorithms are well suited for real time implementation and indeed, since there was no ability to buffer large number of I and Q samples, there was little choice. The use of this approach is widespread, e.g. US NEXRAD, US TDWR, FMI Network Finland, DWD Network, Germany. SIGMET's RVP5 (1985), RVP6 (1992) and RVP7 (1997) processors used the IIR/Pulse-Pair approach. The major drawbacks of this approach are:

– The impulse response of the IIR filter is, as the name implies, infinite. This means that perturbations that are encountered, such as a very large point clutter target or change in the PRF will effect the filter output for many pulses sometimes effecting the output for several beamwidths. The use of clearing pulses or filter initialization can mitigate the effect of this at the expense of effectively reducing the number of pulses.

– The filter width that is necessary to remove clutter bias depends on the strength of the clutter. If the clutter is very strong, then a wider filter is required since the clutter power will exceed the noise power for a greater fraction of the Nyquist interval. In other words, the fixed notch-width is guaranteed to be either not aggressive enough for strong clutter and overly aggressive in removing weather echoes even when there is no clutter. When no clutter is present, the filter will bias the intensity and velocity estimates when the weather target is in the stop band of the filter (overlapped).

– Operators must manually select a filter that is sufficiently wide to remove the clutter without being too wide that the filter attenuates weather. Clutter filter maps have been used in some systems to try and address

Correspondence to: R. E. Passarelli, Jr.
(rich@sigmet.com)

this, but the fact is that scan rates and local target properties (width) may not be the same from day-to-day. AP is a good example of this.

Some other systems have employed FFT processing such the INM Network in Spain, SMHI/FV network in Sweden and EC network in Canada. SIGMET's RVP6 and RVP7 processors provided FFT processing. The advantage of an FFT approach is that the ground clutter filtering can be made adaptive by searching in the frequency domain to determine the boundary between the system noise level and the ground clutter. The FFT is inherently an FIR block processing approach that does not have the transient behaviour problems of the IIR filter. It is possible to interpolate over spectrum components that are removed to minimize the effects of filter bias. However, the FFT approach has two distinct disadvantages:

– The spectrum resolution is limited by the number of points in the FFT which is constrained to be a power of 2 (e.g. 16, 32, 64, 128). Operational systems typically use 32 or 64-point FFT's. If the number of points is low, then clutter will be spread over a larger fraction of the Nyquist domain thus obscuring weather targets.

– To provide the best performance, a time-domain window is applied to the IQ values prior to performing the FFT. When the clutter is very weak, even a rectangular window (uniform weighting) is adequate for removing clutter. However, when the clutter is strong, then a more aggressive window such as the Hamming or the very aggressive Blackman must be used to have any hope of isolating weather signals from strong clutter. The drawback of these windows is that they effectively reduce the number of samples that are processed since the points at the beginning and end of a time series are weighted less than those in the middle. This translates into estimates with a higher variance.

Also, if a fixed notch-width rather than an adaptive width approach is used, then the same types of problems inherent in the IIR filter will exist.

The GMAP approach, used in SIGMET's new RVP8 digital receiver and signal processor (O'Hora and Passarelli, 2002) addresses all of these issues. IQ time series values are calculated on a PCI digital receiver card and passed over the PCI bus into memory where they are effectively processed "off-line" rather than in real-time, with the constraint that the average throughput not exceed the data acquisition time.

With the CPU and memory speed of the RVP8 it is now possible to perform adaptive algorithms that operate iteratively to automatically determine the optimal processing approach which does minimal damage to overlapped weather while providing aggressive cancellation when strong clutter requires it. An important design feature for unattended operational networks is that GMAP does not require any local operator intervention to select the best clutter filter.

2 GMAP Assumptions

The approach makes several assumptions about clutter, weather and noise, i.e.

– The spectrum width of the weather signal is greater than that of the clutter. This is a fundamental assumption required of all Doppler clutter filters.

– The Doppler spectrum consists of ground clutter, a single weather target and noise. Bi-modal weather targets, aircraft or birds mixed with weather would violate this assumption.

– The width of the clutter is approximately known. This is determined primarily by the scan speed and to a lesser extent by the climatology of the local clutter targets. The assumed width is used to determine how many interior clutter points are removed.

– The shape of the clutter is approximately Gaussian. This shape is used to calculate how many interior clutter points are removed.

– The shape of the weather is approximately Gaussian. This shape is used to reconstruct filtered points in overlapped weather.

3 Algorithm Description

The steps used to implement the GMAP approach are shown schematically in Fig. 1 and summarized below.

– **Step 1: Window and DFT**
First a Hamming window weighting function is applied to the IQ values and a discrete Fourier transform (DFT) is then performed. This provides better spectrum resolution than a fast Fourier Transform (FFT) which requires that the number of IQ values be a power of 2. Note that if the requested number of samples is exactly a power of 2, then an FFT is used.

As mentioned in Section 1, when there is no or very little clutter, use of a rectangular weighting function leads to the lowest-variance estimates of intensity, mean velocity and spectrum width. When there is a very large amount of clutter, then the aggressive Blackman window is required to reduce the "spill-over" of power from the clutter target into the sidelobes of the impulse response function. The Hamming window is used as the first guess. After the first pass GMAP analysis is complete, a decision is made to either accept the Hamming results, or recalculate for either rectangular or Blackman depending on the clutter-to-signal ratio (CSR) computed from the Hamming analysis. The recalculated results are then checked to determine whether to use these or the original Hamming result (see Fig. 1 for details).

Figure 1: GMAP Algorithm Steps

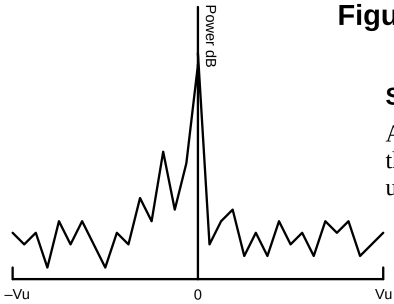

Step 1: Window and DFT

Apply window and DFT the input time series to obtain the Doppler power spectrum. A Hamming window is used for the first trial.

Step 2 (Optional): Dynamic noise power

If the noise level is not known, or if GMAP is recalculated using the Blackman window for CSR>40 dB, then this step is performed. Re-organize the spectrum components in ascending order of intensity. The theoretical relationship for noise is the curved line. The sum of the power in the range 5% to 40% is calculated. This is used to determine the noise level by comparing with the sum value corresponding to the theoretical curve. Next, the power is summed beyond the 40% point for both the actual and theoretical rank spectra. The point where the actual power sum exceeds the theoretical value by 2 dB determines the boundary between the noise region and the signal/clutter region.

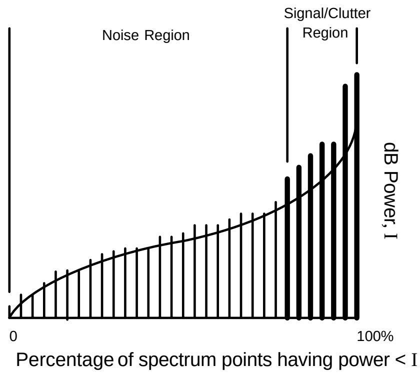

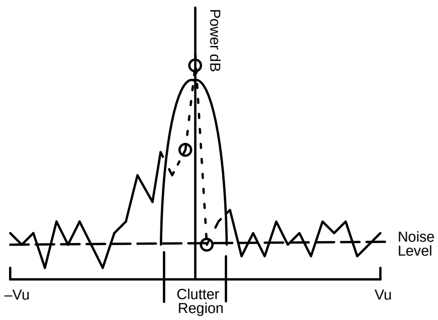

Step 3: Remove clutter points

Use the total power of the three central spectrum points (indicated by the three open circles) to fit a Gaussian having the selected nominal spectrum width in m/s (a function of the number of spectrum samples, PRF and wavelength). The points within the intersection of the Gaussian clutter and the noise level (the "Clutter Region") are discarded (indicated by the dashed lines).

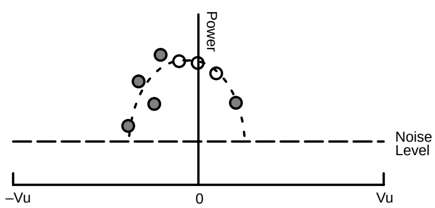

Step 4: Replace clutter points

Dynamic Noise Case: Using the components which have been determined to be neither clutter nor noise (indicated by the filled circles), fit a Gaussian and fill-in the clutter points that were removed in the previous step (indicated by the open circles). Then re-fit the Gaussian with the replacement values inserted. Repeat the iteration until the computed power does not change by more than 0.2dB AND the velocity does not change by more than 0.5% of the Nyquist velocity.

Fixed Noise Case: Similar except the spectrum points that are larger than the noise level are used.

Step 5: Recompute GMAP with optimal window

Determine if the optimal window was used based on the clutter-to-signal ratio (CSR)

IF CSR > 40 dB	repeat GMAP using a Blackman window and dynamic noise calculation.
IF CSR > 20 dB	repeat GMAP using a Blackman window. Then if CSR>25dB use Blackman results.
IF CSR < 2.5 dB	repeat GMAP using a rectangular window. Then if CSR < 1 dB use rectangular results.
ELSE	accept the Hamming window result.

Fig. 1. GMAP Algorithm Steps.

- **Step 2: Determine the noise power**
  In general, the spectrum noise power is known from periodic noise power measurements. Since the receiver is linear and requires no STC or AGC, the noise power is well-behaved at all ranges. The only time that the spectrum noise power will differ from the measured noise power is for very strong clutter targets. In this case, the clutter contributes power to all frequencies, essentially increasing the spectrum noise level. This occurs for two reasons: 1) In the presence of very strong clutter, even a small amount of phase noise causes the spectrum noise level to increase, and 2) There is significant power that occurs in the window side-lobes. For a Hamming window, the window side lobes are down by 40 dB from the peak at zero velocity. Thus 50 dB clutter targets will have spectrum noise that is dominated by the window sidelobes in the Hamming case. The more aggressive Blackman window has approximately 55 dB window sidelobes at the expense of having a wider impulse response and larger negative effect on the variance of the estimates.

  When the noise power is not known, it is optionally computed using a dynamic approach similar to that of Hildebrand and Sekhon (1974). The Doppler spectrum components are first sorted in order of their power. As shown in Fig. 1, the sorting places the weakest component on the left and the strongest component on the right. The vertical axis is the power of the component. The horizontal axis is the percentage of components that have power less than the y-axis power value. Plotted on a dB scale, Poisson distributed noise has a distinct shape, as shown by the curved line in Fig. 1. This shape shows a strong singularity at the left associated with taking the log of numbers near zero, and a strong maximum at the right where there is always a finite probability that a few components will have extremely large values.

  There are generally two regions: a noise region on the left (weaker power) and a signal/clutter region on the right (stronger power). The noise level and the transition between these two regions is determined by first summing the power in the range 5% to 40%. This sum is used to determine the noise level by comparing with the sum value corresponding to the theoretical curve. Next, the power is summed beyond the 40% point for both the actual and theoretical rank spectra. The point where the actual power sum exceeds the theoretical value by 2 dB determines the boundary between the noise region and the signal/clutter region.

  Finally there are two outputs from this step: a spectrum noise level and a list of components that are either signal or clutter.

- **Step 3: Remove the clutter points**
  The inputs for this step are the Doppler power spectrum, the assumed clutter width in m/s and the noise level, either known from noise measurement or optionally calculated from the previous step. First the power in the three central spectrum components is summed (DC +1 component) and compared to the power that would be in the three central components of a normalized Gaussian spectrum having the specified clutter width and discretized in the identical manner. This serves as a basis for normalizing the power in the Gaussian to the observed power. The Gaussian is extended down to the noise level and all spectral components that fall within the Gaussian curve are removed. The power in the components that are removed is the "clutter power".

  A subtle point is the use of the three central points to do the power normalization of the actual vs the idealized spectrum of clutter. This is more robust than using a single point since for some realizations of clutter targets viewed with a scanning antenna, the DC component is not necessarily the maximum. Averaging over the three central components is a more robust way to characterize the clutter power.

  The very substantial algorithmic work that has been done thus far is to eliminate the proper number of central points. The operator only has to specify a nominal clutter width in m/s. This means that the operator does not need to consider the PRF, wavelength or number of spectrum points- GMAP accounts for these automatically.

  A key point is that in the event that the sum of the three central components is less than the corresponding noise power, then it is assumed that there is no clutter and all of the moments are then calculated using a rectangular window. If the power in the three central components is only slightly larger than the noise level, then the computed width for clutter removal will be so narrow that only the central (DC) point shall be removed. This is very important since, if there is no clutter then we want to do nothing or at worst only remove the central component.

  Because of this behaviour, there is no need to do a clutter bypass map, i.e. turn-off the clutter filter at specific ranges, azimuths and elevation for which the map declares that there is no clutter. Because of the day-to-day variations in the clutter and the presence of AP, the clutter map will often be incorrect. Since GMAP determines the no-filter case automatically and then processes accordingly, a clutter map is not required.

– **Step 4: Replace clutter points**
The assumption of a Gaussian weather spectrum now comes into play to replace the points that have been removed by the clutter filter. There are two cases depending on how the noise level is determined under Step 2, i.e. the dynamic noise case and the fixed noise level case.

Dynamic noise level case: From Step 2, we know which spectrum components are noise. From Step 3 we know which spectrum components are clutter. Presumably, everything that is left is weather signal. An inverse DFT using only these components is performed to obtain the autocorrelation at lags 0, 1. This is very computationally efficient since there are typically few remaining points and only the first two lags need be calculated. The pulse pair mean velocity and spectrum width are calculated using the Gaussian model (e.g. see Doviak and Zrnic, 1993). Note that since the noise has already been removed, there is no need to do a noise correction. The Gaussian model is then applied using the calculated moments to determine a substitution value for each of the spectrum components that were removed in Step 3. In the case of overlapped weather as shown in the Fig. 1 example, the replacement power is typically too small. For this reason, the algorithm recomputes R0 and R1 using both the observed and the replacement points and computes new replacement points. This procedure is done iteratively until the power difference between two successive iterations is less than 0.2 dB and the velocity difference is less than 0.5% of the Nyquist interval.

In summary of this step, the Gaussian weather model is used to repair the filter bias, i.e. the damage that is caused by removing the clutter points. An IIR filtering approach makes no attempt to repair filter bias, rather the filter simply ”digs a hole” into overlapped weather.

– **Step 5: Check for appropriate window and recalculate the moments if necessary.**
The clutter power is known from the spectrum components that were removed in Step 3. Since the weather spectrum moments and the noise are also known from Step 4, the CSR can be calculated. The value of the CSR, is used to decide whether the Hamming window is the most appropriate. The scenarios are described in Fig. 1. The end result is that very weak clutter is processed using a rectangular window, moderate clutter a Hamming window, while severe clutter requires a Blackman window. Note that if no clutter were removed in Step 3, then the spectrum is processed with a rectangular window.

The benefit of adaptive windowing is that the least aggressive window is used for the calculation of the spectrum moments, resulting in the minimum variance of the moment estimates.

4 Examples of Implementation

GMAP has undergone extensive evaluation for use in the US WSR88D ORDA network upgrade (Ice et al., 2004). They conclude that GMAP meets the ORDA requirements.

Their study was based on a built-in simulator that is provided as part of the RVP8 system. The simulator allows users to construct Doppler spectra, process them and evaluate the results (Sirmans and Bumgarner, 1975). This is an essential tool for evaluating the system performance.

Figure 2 shows an example of the simulations for the very difficult case when the weather has zero velocity, i.e. it is perfectly overlapped with clutter. The upper left graph shows the weather signal with −40 dB power without any clutter and without any GMAP filtering. The graph at the upper right shows the same spectrum with 0 dB of clutter power added for a clutter width of 0.012 (0.3 m/s at S band, 1000 Hz PRF). This is a CSR of 40 dB. The panel at the lower left shows the weather signal after GMAP filtering.

In each of the moment plots, there are several values that are displayed. The left-most number shows the value at the range cursor which is positioned as indicated by the vertical line. To the right, the “m” value is the mean and the “s” value the standard deviation as averaged over all range bins (1000 in this example). For velocity these are in normalized units expressed as a fraction of the Nyquist interval. For reflectivity the values are in dB.

Some key points are:

– The mean velocity is correctly recovered as expected (the “m” value in the plot), but the standard deviation is higher (0.06 vs 0.04 in normalized units).

– The “Cor dBZ” shows 40.2 dB of “C.Rej”. This is the difference between the “Tot dBZ” and the “Cor dBZ” values. The expected value is 40 dB in this case. This indicates that GMAP has recovered the weather signal in spite of the aggressive clutter filtering that is required.

– The standard deviation of the ”Tot dBZ” is greater in the weather plus clutter (4.35 normalized units) as compared to the weather-only case. This is caused by the fluctuations in the clutter power in the Gaussian clutter model.

– The standard deviation of the Cor dBZ after GMAP filtering, while not as low as for the weather-only case are lower than the weather plus clutter case. In other words, the GMAP processing removes some of the high variance in the dBZ estimates that is caused by clutter, but is not quite as good as doing nothing.

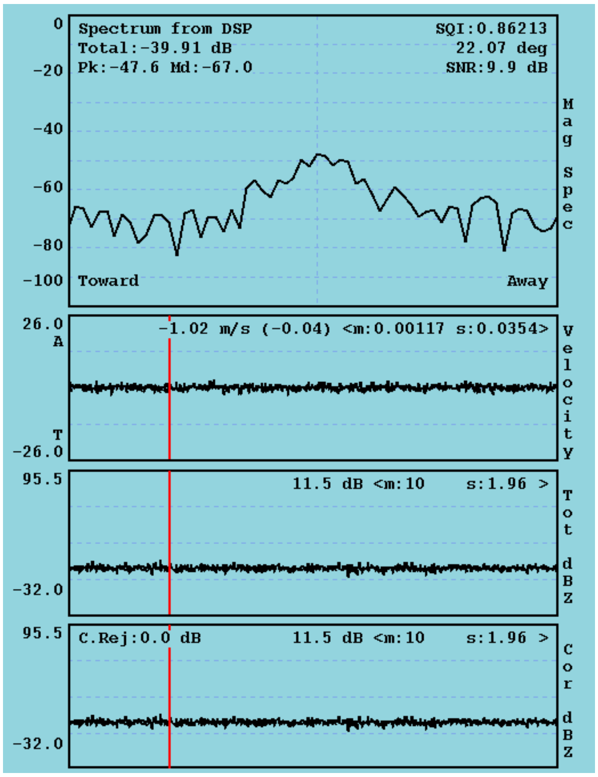

Weather only

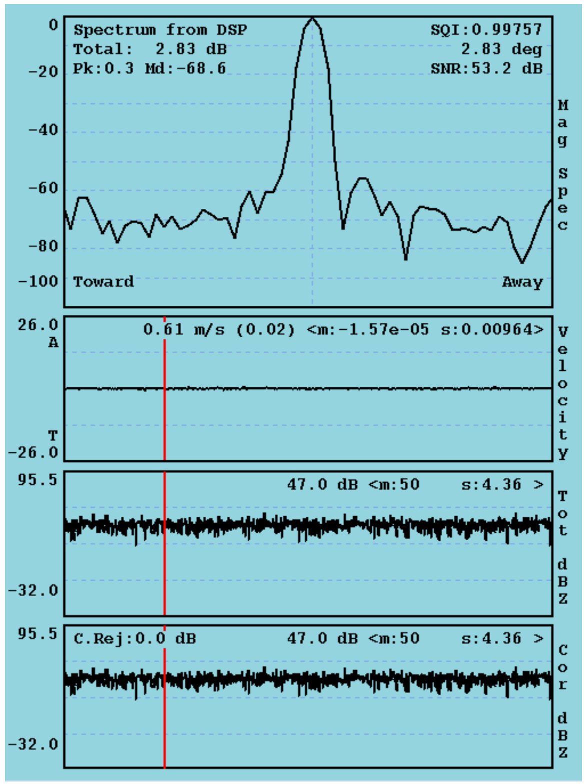

Weather plus clutter

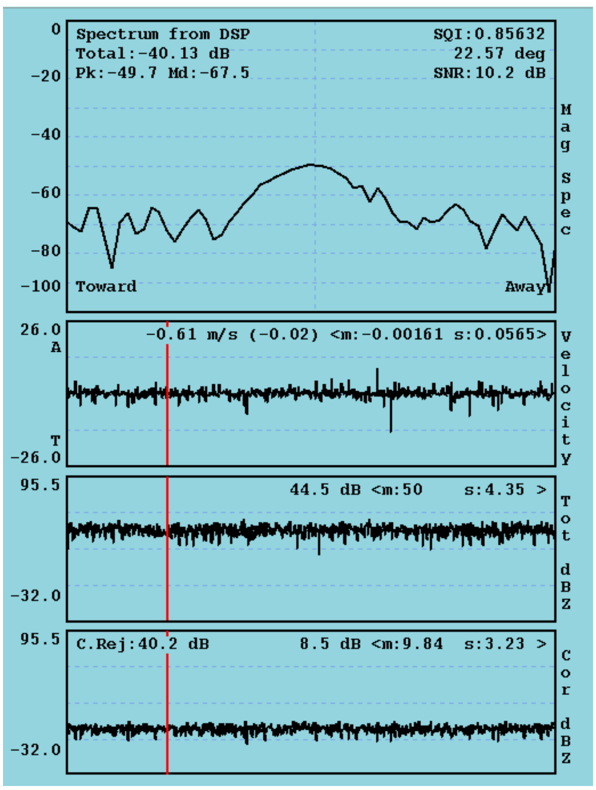

Weather signal after GMAP Filtering

Simulation Characteristics

	Clutter	Weather	Units
Power	0	–40	dB
Vel	0	0	Any
Width	0.012	0.1	Normalized
PRF	1000 Hz	**Window**	Blackman
Mode	FFT	**Samples**	64

"Mag Spec": Doppler Spectrum in dB Units spanning the Nyquist interval.

"Velocity": Mean velocity of the spectrum in over Nyquist interval. Mean "m" and standard deviation values "s" are for the normalized interval ±1.

"Tot dBZ": Power in dB of weather and clutter. Mean "m" and standard deviation values "s" are in dB.

"Cor dBZ": Power in dB after GMAP filtering. Mean "m" and standard deviation values "s" are in dB.

"<m: ...s: ...>" mean and standard deviation over all ranges, in this case 1000 range bins.

Fig. 2. GMAP simulation example.

5 Conclusions

GMAP provides substantial advantages over legacy processing techniques such as pulse-pair processing with fixed IIR or FIR filters. The computation required to perform the GMAP is substantially greater than is required of these other techniques because of the use of DFT's, the optional rank noise calculation and the fact that many times the analysis needs to be recalculated to select the optimal window. However, new processors such as the RVP8 have more than adequate speed to perform GMAP.

One of the strong features of GMAP is that it does nothing or very little when there is no clutter. This eliminates the need for clutter bypass maps that vary from day-to-day and must be maintained.

Acknowledgement. For their comprehensive GMAP evaluation study the authors thank Richard L. Ice, Rick D. Rhoton, Darcy S. Saxion, and Nita K. Patel of RS Information Systems, Inc. Norman, Oklahoma; Dale Sirmans and David A. Warde of Systems Technology Associates, Norman, Oklahoma; Darrel L. Rachel Management and Engineering Technology International, Norman, Oklahoma; Ron G. Fehlen, Lt, USAF WSR-88D Radar Operations Center, Norman, Oklahoma. At SIGMET the authors thank Fritz O'Hora and Joseph Holmes.

References

Hildebrand, P. H. and Sekhon, R. S.: Objective Determination of Noise level in Doppler spectra. J. Atmos. Sci. 13, 808–811, 1974.

Ice, R. L., Rhoton, R. D., Saxion, D. S., Patel, N. K., Sirmans, D., Warde, D. A., Rachel, D. L., and Fehlen, R. G.: Radar Operations Center (ROC) Evaluation of the WSR-88D Open Radar Data Acquisition (ORDA) System Signal Processing, 20th International Conference on Interactive Information and Processing Systems for Meteorology Oceanography, and Hydrology, 11–15 January 2004, Seattle, WA, AMS, 2004.

Sirmans, D. and Bumgarner, B.: Numerical Comparison of Five Mean Frequency Estimators, Journal of Applied Meteorology, 14, 991–1003, 1975.

Doviak, R. J. and Zrnic, D. S.: Doppler Radar and Weather Observations, p. 136, Academic Press, Inc., New York, 1993.