

**Product Development Team
for
NEXRAD Algorithms**

Quarterly Report – 3rd Quarter FY 03

03.6.2 Polarization and Frequency Diversity

Algorithms based on polarimetry will meet the aviation needs for information about the volumetric extent of hail, freezing rain, snow, and icing conditions, as well as non-hydrometeor scatterers. The biggest potential payoff is enhanced data quality. For all practical purposes, polarimetric techniques will eliminate problems associated with sea-clutter, ground clutter, AP, and biological scatterers.

a) Current Efforts

(NSSL):

The most significant activity during Q3 was the implementation of real time display at the WFO of a "non-precipitation removed" reflectivity product. Example before-and-after images are attached. We are estimating that about 80% of the non-precipitating echoes are removed from display. As can be seen, some ground clutter is still getting through, and that is being worked on.

Ron Larkin visited Norman during the week of 21 May. KOUN radar collected clear air bird and insect returns while Ron made field observations of birds and insects over several nights. Using these data, better estimates will be made of the boundaries in the polarimetric signatures separating birds from insects so that contamination of Doppler wind velocity profiles from birds can be mitigated.

Data with biological scatterers were also collected on 1 May, 4 May, 5 May, 8-9 May, and 19 May. Data with significant AP were collected on 1 May and 19 May.

Finally, collected data were on large-hail producing storms on more than a dozen different days during April-June. One storm contained hail more than 13 cm in diameter at the surface, which provides enough data to explore forecasting hail size with a polarimetric radar.

(NCAR):

TASK 03.6.2.10: Winter storm case studies

The collection of video disdrometer observations from winter storms ended. Significant datasets were obtained from a total of 17 events. Two events involved snow pellets and two events had periods of freezing drizzle. The

dataset includes measurements of hydrometeor sizes, aspect ratios, canting angles, terminal velocities, wind speed and direction, temperature, humidity, visibility, liquid water content. Calibration factors for the disdrometer are being applied. The data were collected in an effort to improve the designation of frozen hydrometeors and better quantify winter precipitation with radar. Analysis is underway.

TASK 03.6.2.11: Implement improved hail detection algorithm

The literature on hail detection with polarimetric radar was reviewed and examples produced in preparation of our milestone progress report.

TASK 03.6.2.12: Test NCAR freezing level algorithm

Reviews of two papers regarding freezing level detection were received and the papers revised accordingly. The utility of the differential propagation phase measurement was reviewed. Because the measurement is relatively insensitive to beam blockage and clutter, it was hoped that the measurement would provide useful information when freezing levels were near ground. Testing with the IMPROVE (Oregon) dataset revealed performance varied considerably, presumably due to measurement error and the presence of Mie scatterers. Fig. 1 shows an example from the PRECIP98 (Florida) field experiment. Pronounced melting layer maxima for reflectivity (Z_H), differential reflectivity (Z_{DR}), and linear depolarization ratio (LDR) are clearly evident. In contrast, the signature for differential phase is more ambiguous. Further, the parameter lacks the clarity of LDR for defining the melting layer. At times the parameter shows a well-defined minimum at the height of reflectivity maximum with distinct flanking maxima. The variability in the profiles makes it difficult to model. Consequently, the parameter is not used in the current algorithm.

A paper entitled "Freezing level determinations with polarimetric radar: Retrieval model and applications (Ikeda and Brandes, authors) was prepared for presentation at the 31st International Radar Conference. A PDF file is attached.

b) Planned Efforts

Continue HCA development and testing using KOUN cases from 2003.

c) Problems/Issues

None.

d) Interface with other Organizations

None.

Date: 091798
Time: 221322 UTC

Observed Profiles

Averaged
Elev: 7.5

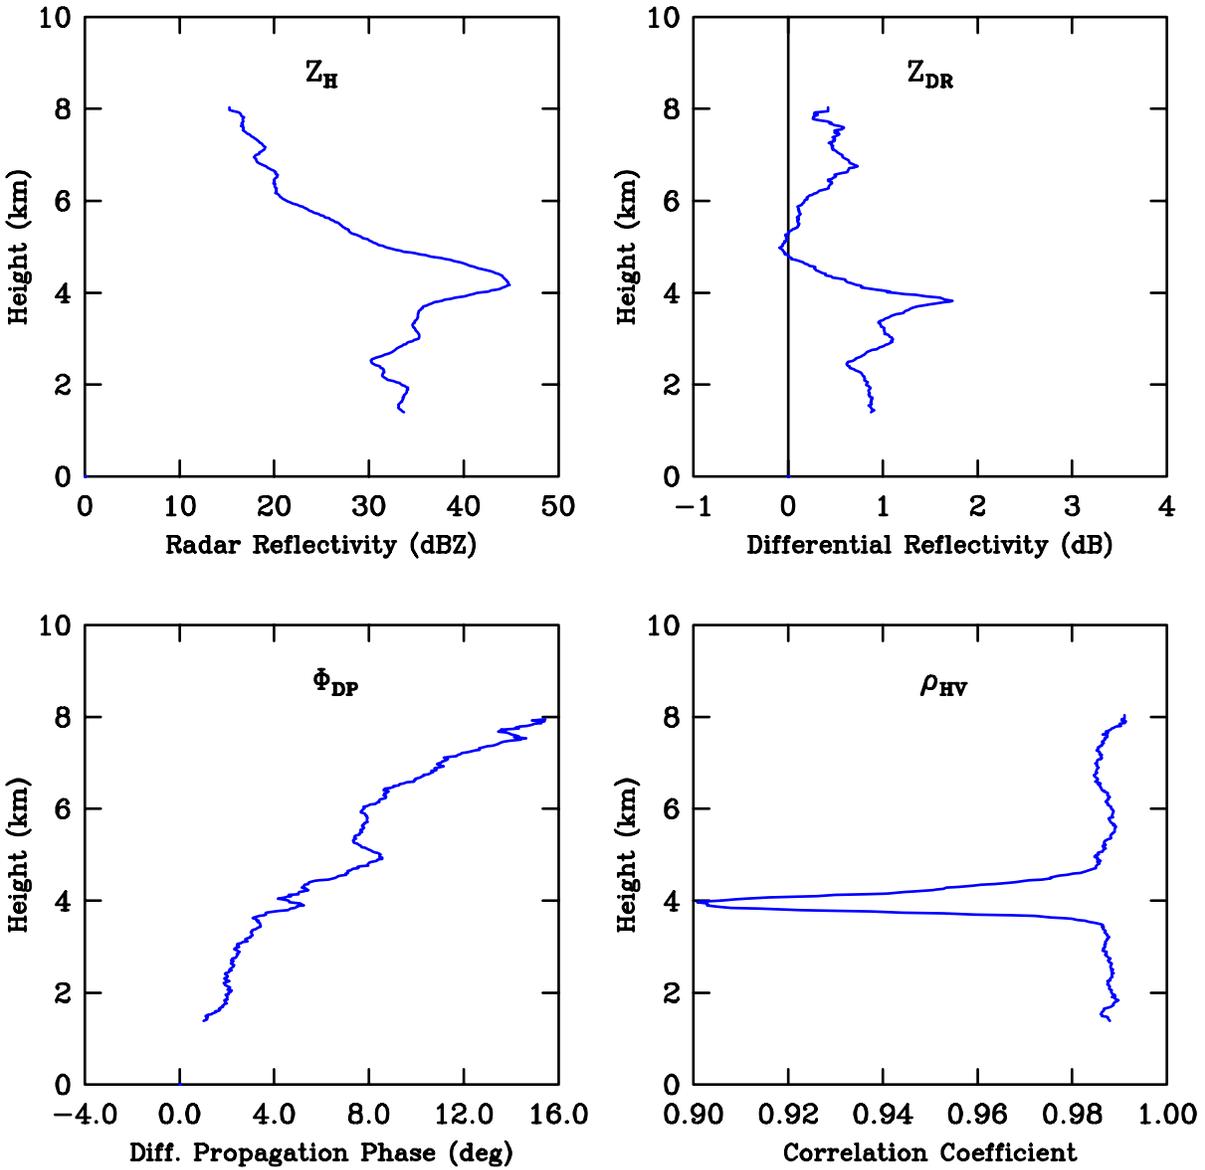


Figure 1. Profiles of radar reflectivity (dBZ), differential reflectivity (Z_{DR}), differential phase (ϕ_{DP}), and correlation (r_{HV}), from the PRECIP98 experiment.

e) Activity Schedule Changes

None.

03.6.3 Circulations

Particularly violent or long-lived storms tend to possess certain notable qualities, including, for example, mesocyclones. The current WSR-88D algorithms have a very high false alarm rate. Controllers find such high false alarm rates unacceptable. To mitigate this problem, new more robust and reliable circulation detection algorithms will be developed. Algorithms that use circulations to diagnose storm severity or estimate storm longevity will be considerably improved by this work.

a) Current Efforts

Prepared and submitted a paper titled "Detection of rotation and boundaries using two-dimensional, local, linear, least squares estimates of velocity derivatives" to the 31st Radar Conference to be held in Seattle, WA. The paper, in PDF, is attached to this report.

b) Planned Efforts

For Q4, NSSL will apply a weighting technique to the calculation of velocity derivatives to reduce discontinuities in the azimuthal shear field with respect to range. NSSL will also apply the LLSD technique to radial velocity models of divergence and convergence, such as might be observed along boundaries or in microbursts.

c) Problems/Issues

None.

d) Interface with other Organizations

None.

e) Activity Schedule Changes

None.

03.6.4 Technical Facilitation

Technical facilitation supports the NEPDT algorithm development. There is currently no standard vehicle outside of NSSL for algorithm development support. The interface being developed at the NSSL, the WDSS-II, provides a way to develop, validate, verify and demonstrate the NEXRAD algorithms developed within this PDT. Additionally, WDSS-II provides a route into the Open Radar Product Generation (ORPG) system. WDSS-II will support and incorporate the MITRE Common Operations Development Environment (CODE). WDSS-II is an important ingredient for the overall success of the NEPDT because, in consonance with CODE, coding and testing standards at the application prototype level are enforced. Transfer of single-radar algorithms to the ROC will be straightforward, as anything within WDSS-II must also conform to CODE standards. Overall, NEPDT efforts will inevitably enhance the algorithms that have been accepted or will be implemented by the ROC as part of the WSR 88D system.

a) Current Efforts

(1) Display support for research tasks:

(a) Several bugs related to visualizing virtual volumes and tables were fixed.

(b) Convenient way to visualize all the tilts of a radar virtual volume, or forecasts of gridded fields was developed.

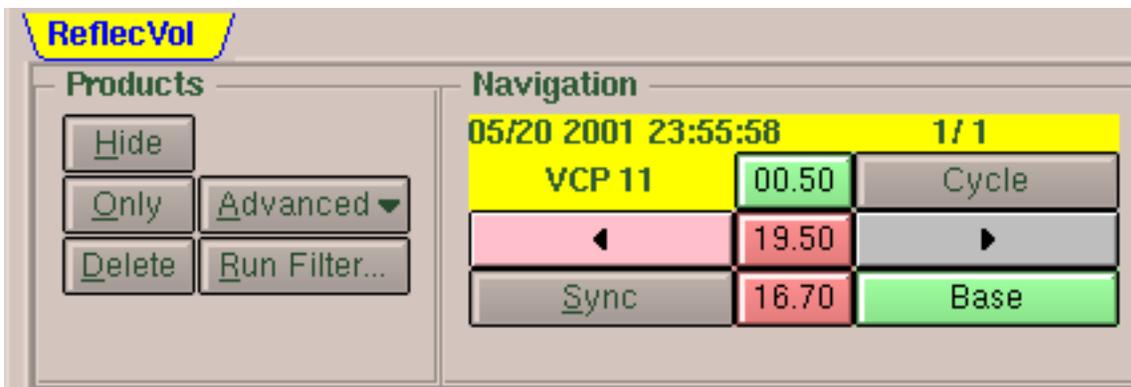


Figure 2. Example of GUI for examining reflectivity products.

(c) Annotation capability was added to the display to allow marking and saving of region information.

(d) Support was added for some of the new digital products in the ORPG.

(2) A quality-control neural network was developed. It incorporates all three radar moments and produces cleaned elevation scans. The neural network product removes AP but leaves storms unaffected.

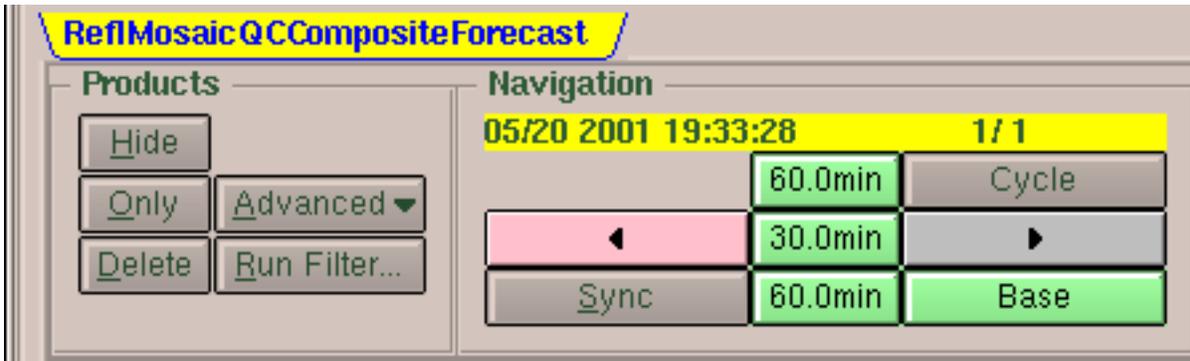


Figure 3. GUI Example for controlling displa of QC composite images.

(3) We implemented the Radar Echo Classifier (REC, from NCAR) and compared the two QC algorithms.

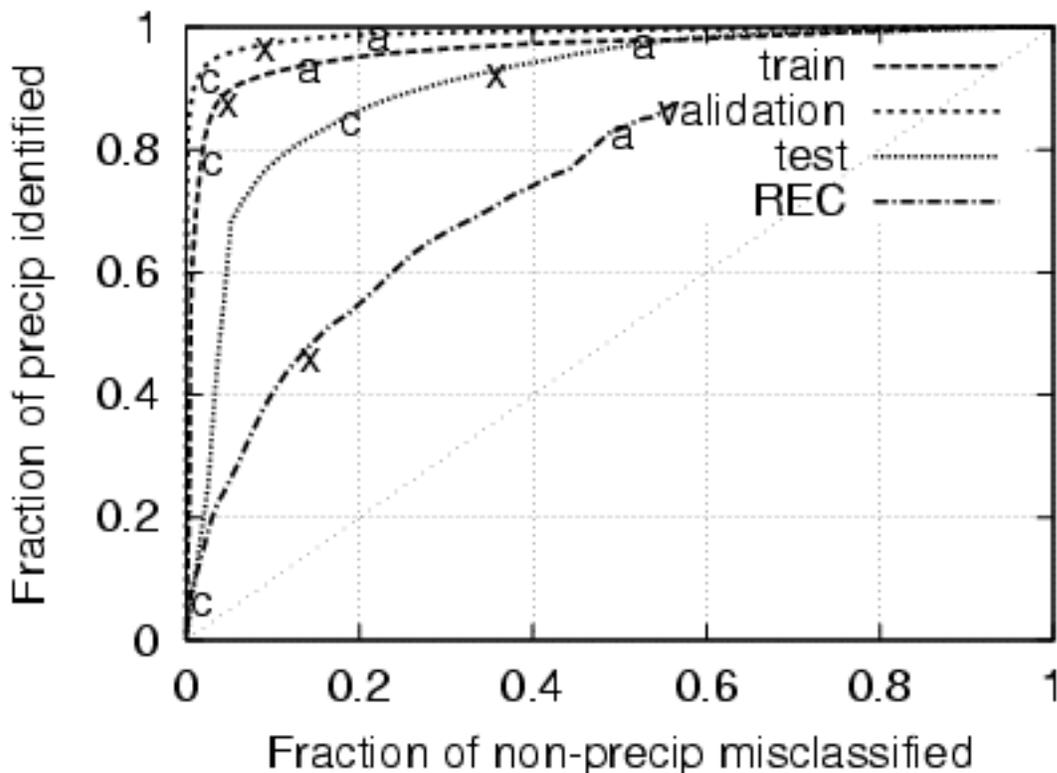


Figure 4. ROC diagram showing performance of various AP quality control techniques.

The closer the curve is to the top-and-left, the better the algorithm. One way to read the graph is this: Suppose we want to retain 80% of all precipitation pixels. With the REC, we have to live with removing only 50% of all AP/clear-air detection pixels. With the NSSL-developed QC neural network, we can remove 90% of AP/clear-air return.

(4) Several algorithms were modified to optionally use the QC'ed reflectivity data instead of the original. This increases the skill of these other algorithms:

(a) Meso-cyclone Detection Algorithm

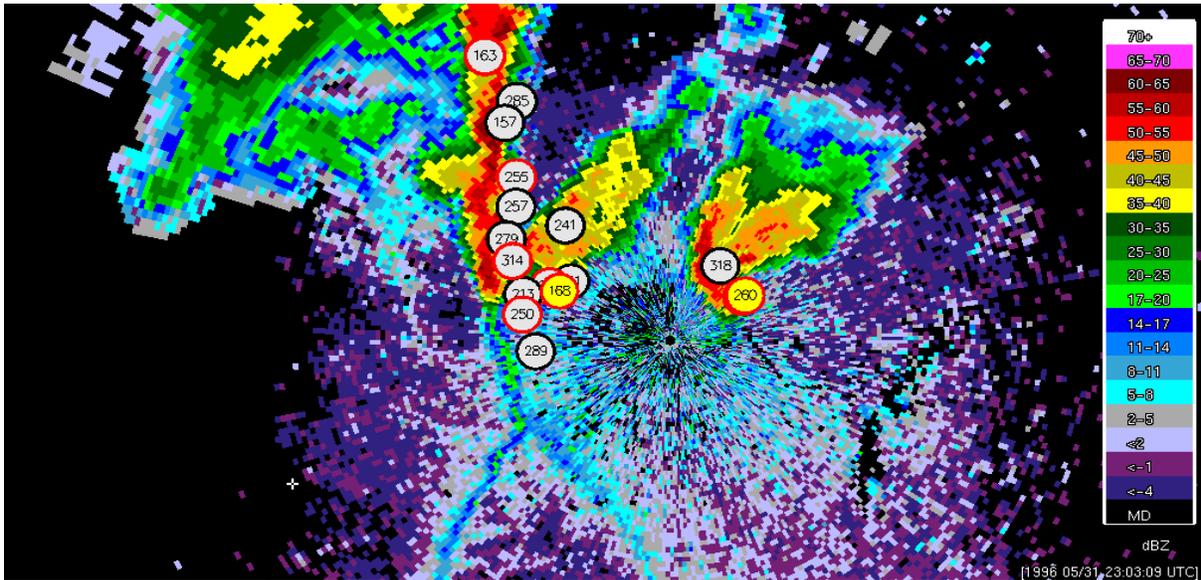


Figure 5. MDA output prior to neural net QC. Note various false detection in clear air.

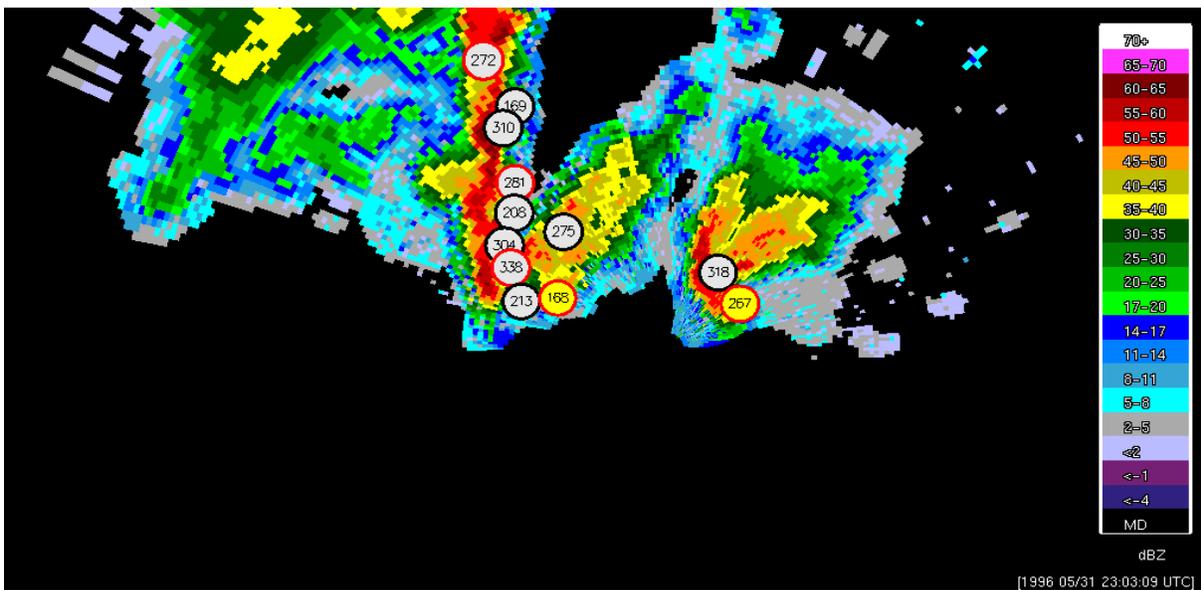


Figure 6. Same as Fig. 5, but for MDA output after neural net QC. All false detections in clear air have been eliminated.

(b) Storm-cell Identification and Tracking Algorithm

Centroid estimates improve when precip is embedded inside AP.

(c) Motion Estimation and Forecast

A 30min forecast and what actually happened in 30minutes are shown. Note that the forecast is a lot better once the AP is removed.

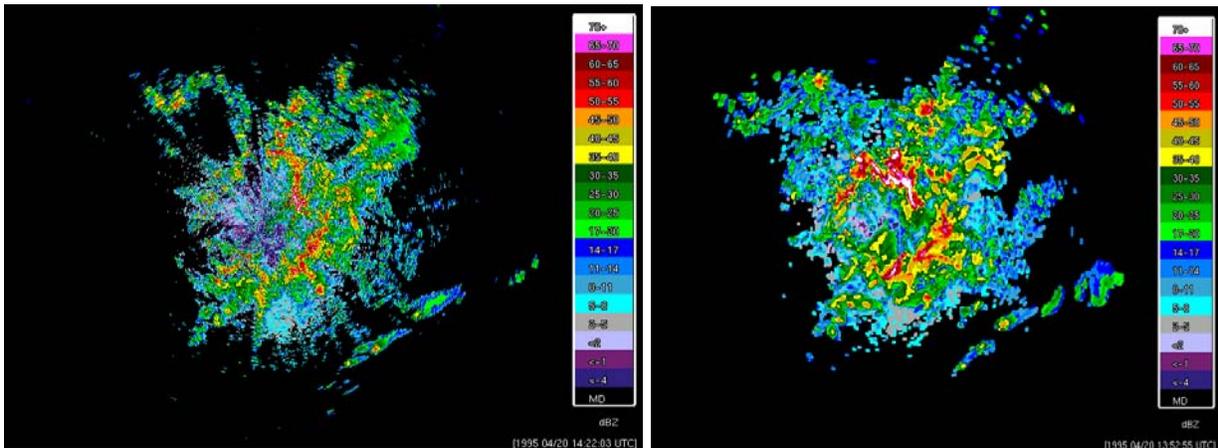


Figure 7. Original radar data (without neural net QC) is on the left, resulting 30 min forecast is on right.

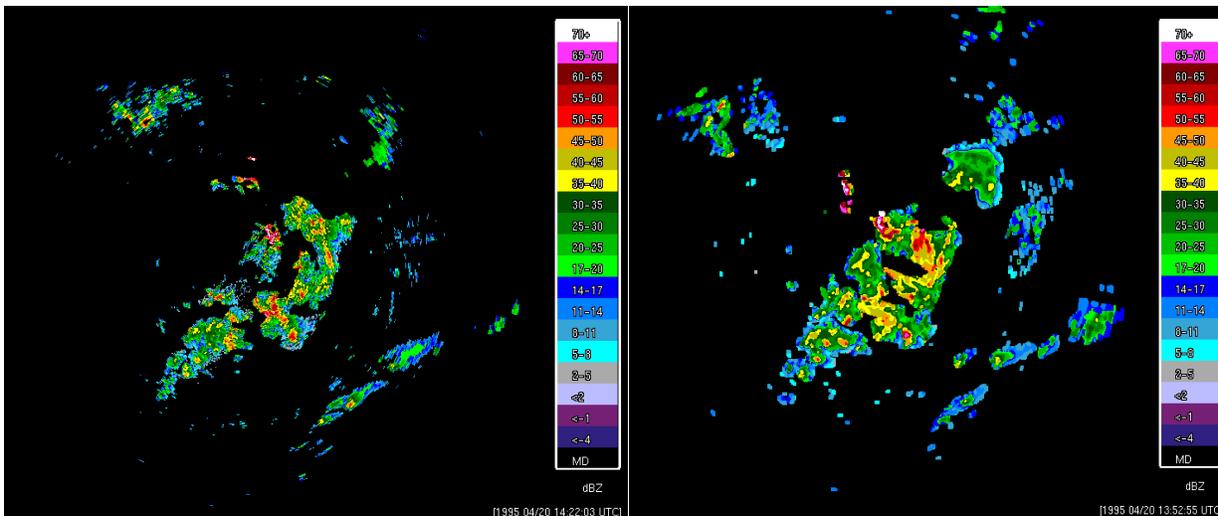


Figure 8. Same as Fig. 7, but for data that has undergone the neural net QC process.

(d) Multi-radar merged Reflectivity with quality-control of reflectivity and internal advection individual radar views: ktlx.gif, kinx.gif, ksrx.gif. Note the differing resolutions, as well as the extensive clear-air return.

quality-controlled, merged product: multi.gif

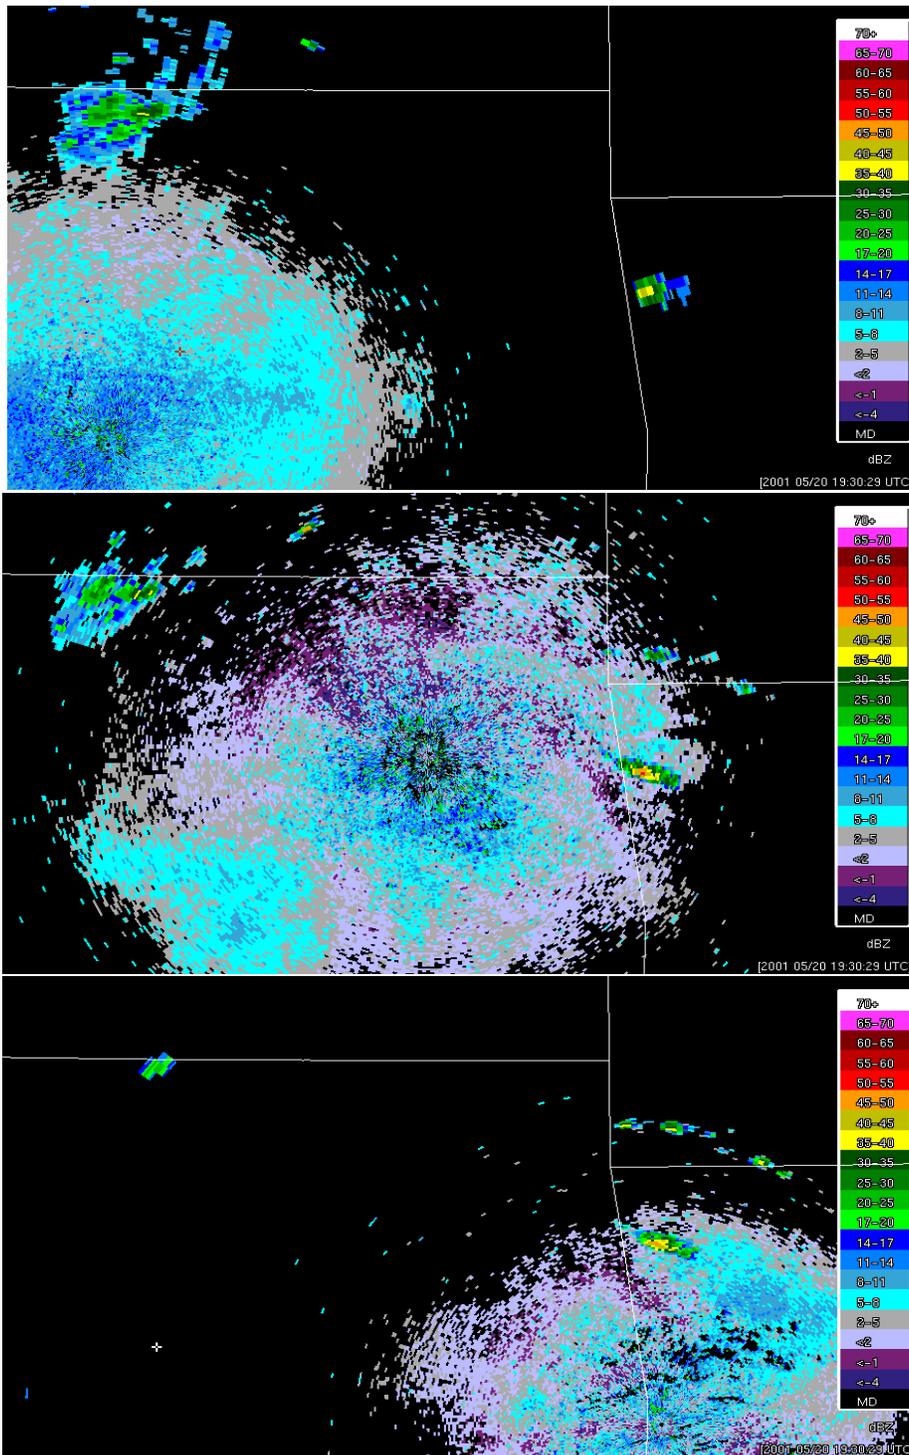


Figure 9. Individual views of non-QC radar data from three different radars: top is KTLS, center is KINX, and bottom is KSRT.

The storm cells are at high resolution, in the right spot, taking the time differences, and viewing angles of each of the radars into account.

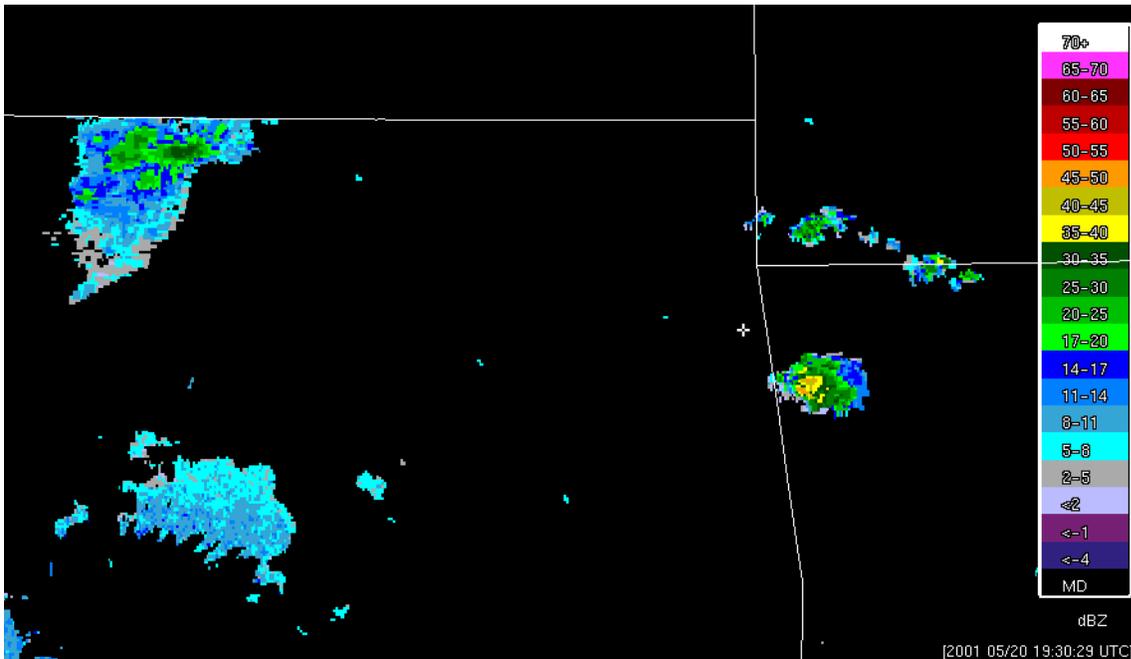


Figure 10. Resultant merger of data from three different radars after neural net QC procedure.

(4) The multi-radar merger process was modified to take terrain and beam-blockage into account.

The following papers have been submitted and are attaced to the end of this report:

Lakshmanan, V. 2003 Motion estimator based on hierarchical clusters. 19th IIPS Conference, Amer. Meteo. Soc., Long Beach, CA.

Lakshmanan, V. 2003 Real-time merging of multi-source data. 19th IIPS Conference, Amer. Meteo. Soc., Long Beach, CA.

Lakshmanan, V., K. Hondl, G. Stumpf, and T. Smith: 2003 Quality control of WSR-88D data. 31st Radar Conference, Amer. Meteo. Soc., Seattle.

Lakshmanan, V., R. Rabin, and V. DeBrunner, 2003: Multiscale storm identification and forecast. J. Atmospheric Research, In Press.

b) Planned Efforts

Continue development and evaluation of neural net QC procedures.

c) Problems/Issues

None.

d) Interface with other Organizations

None.

e) Activity Schedule Changes

None.

03.6.12 Product Implementation

Product implementation is a the process by which implementation paths are explored and defined within the aviation community systems that are best for NEPDT products. This process includes collaboration with other PDTs to help define the nature of WSR-88D they need. Technical facilitation also includes the low-level process of defining technical details (formats, data set file structures, etc.) of the products developed NEPDT.

a) Current Efforts

Discussions held with NCEP and FSL concerning implementation of 3D gridded radar mosaic.

b) Planned Efforts

Develop work plan for implementing 3D gridded radar mosaic to implement in FY 04 (an addition to the original seven year plan).

c) Problems/Issues

None.

d) Interface with other Organizations

NCAR, MIT/LL, FSL.

e) Activity Schedule Changes

None

03.6.14 Multi-radar Composites

The area for which any arbitrary ARTCC has responsibility likely encompasses the coverage area of several WSR-88D installations. Neither the ROC nor the NWS has plans to treat the various WSR-88D installations as a single network, so there are no existing algorithms that use data from more than one radar. This is a serious limitation, because treating each radar separately leads to ambiguities when the radar data overlap. Currently, the users must independently mitigate these ambiguities, which requires significant knowledge about meteorological radar data and the nature of the algorithms that are run on these data. Aviation users generally do not possess this knowledge, so for the WSR-88Ds to be treated as a network, algorithms and techniques aimed specifically at multiple radar composites must be developed.

a) Current Efforts

03.6.14.1-2 Continued test and refinement of real-time 3-D mosaic for the FAA CIWS region

The activities for this quarter include continued monitoring of the real-time 3-D mosaic for the CIWS domain. The system has been up for more than 97% of the time since the installation of the 3D mosaic (August 30, 2002). Real-time 3D mosaic data as well as single radar radial velocity data are kept online for 5-running days (extended from 2 days, see NAPDT 2003 1st quarterly report). Other FAA PDTs (e.g., CWPDT and WWRPDT) have been accessing these data via ftp.

The 3D mosaic algorithm has been further refined based on real-time case studies. It is found that when storms are very close to radar, their intensities in the 3D analyses could be weaker than they are in the single radar field. Even though the current 3D mosaic grid resolution (i.e., 1km x 1km x 500m) is much finer than radar polar grid in majority regions of radar umbrellas, it is coarser than the polar grid resolution at closer ranges. This is due to the fact that radar data has extremely non-uniform spatial resolutions. Figures 10 and 11 show a horizontal and a vertical cross section where Cartesian grid (indicated by white lines) is overlaid on polar grid (color shaded radar bins). At the range of 10 km, there are ~6 radar bins in one grid cell (Fig.10) in horizontal and ~3 bins in vertical (Fig. 11). If a local maximum reflectivity is observed in these regions, the intensity could be reduced in the analysis field if a weighted mean scheme is used. Figure 12 shows composite reflectivities from KVNx radar and from the 3D mosaic grid that covers the same region. The high reflectivities (>60 dbZ) as indicated by arrows in Fig.12a were not shown in the mosaic field (Fig. 12b). In addition, the echo area of 55 dbZ and higher (red-color shaded, Fig. 12a) is larger in the single radar field than in the mosaic grid. This could pose problems for severe storm applications since magnitude of storm intensities is an important parameter. To assure the conservation of storm intensities in analyses, a separate 2D analysis and mosaic scheme is used to derive composite reflectivity on the Cartesian grid. The scheme uses a "taking maximum" instead

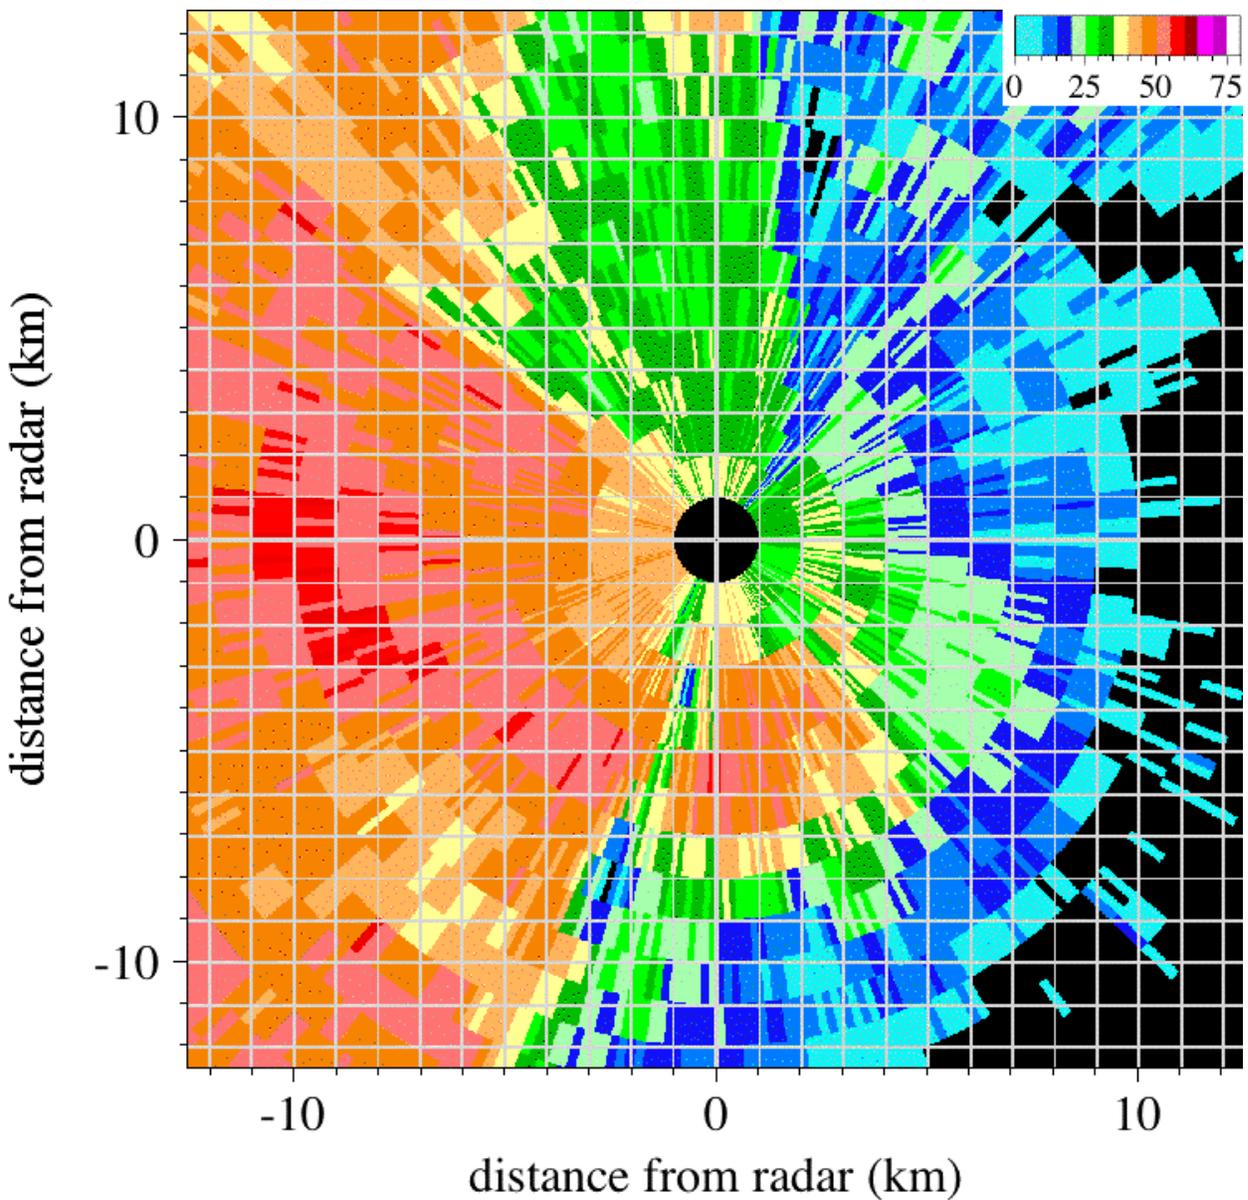


Figure 11. A reflectivity PPI image shows relative size of radar bins ($1^\circ \times 1\text{km}$, color shaded) and Cartesian grid cells ($1\text{km} \times 1\text{km}$, white lines). At ranges of 10 and 15km, one grid cell can cover ~6 and ~3-4 bins, respectively. Damping of storm intensities can occur in these regions when a weighted mean analysis scheme is used.

of "weighted mean" approach when deriving the composite reflectivity field. The new scheme successfully retained intensities of storms (Fig. 13).

03.6.14.3-5 Provide the 3D mosaic grid data for case studies requested by fellow PDTs.

In addition to making the real-time 3D mosaic data available to other FAA PDTs, the NAPDT has performed reanalysis on archived radar data for the CWPDT.

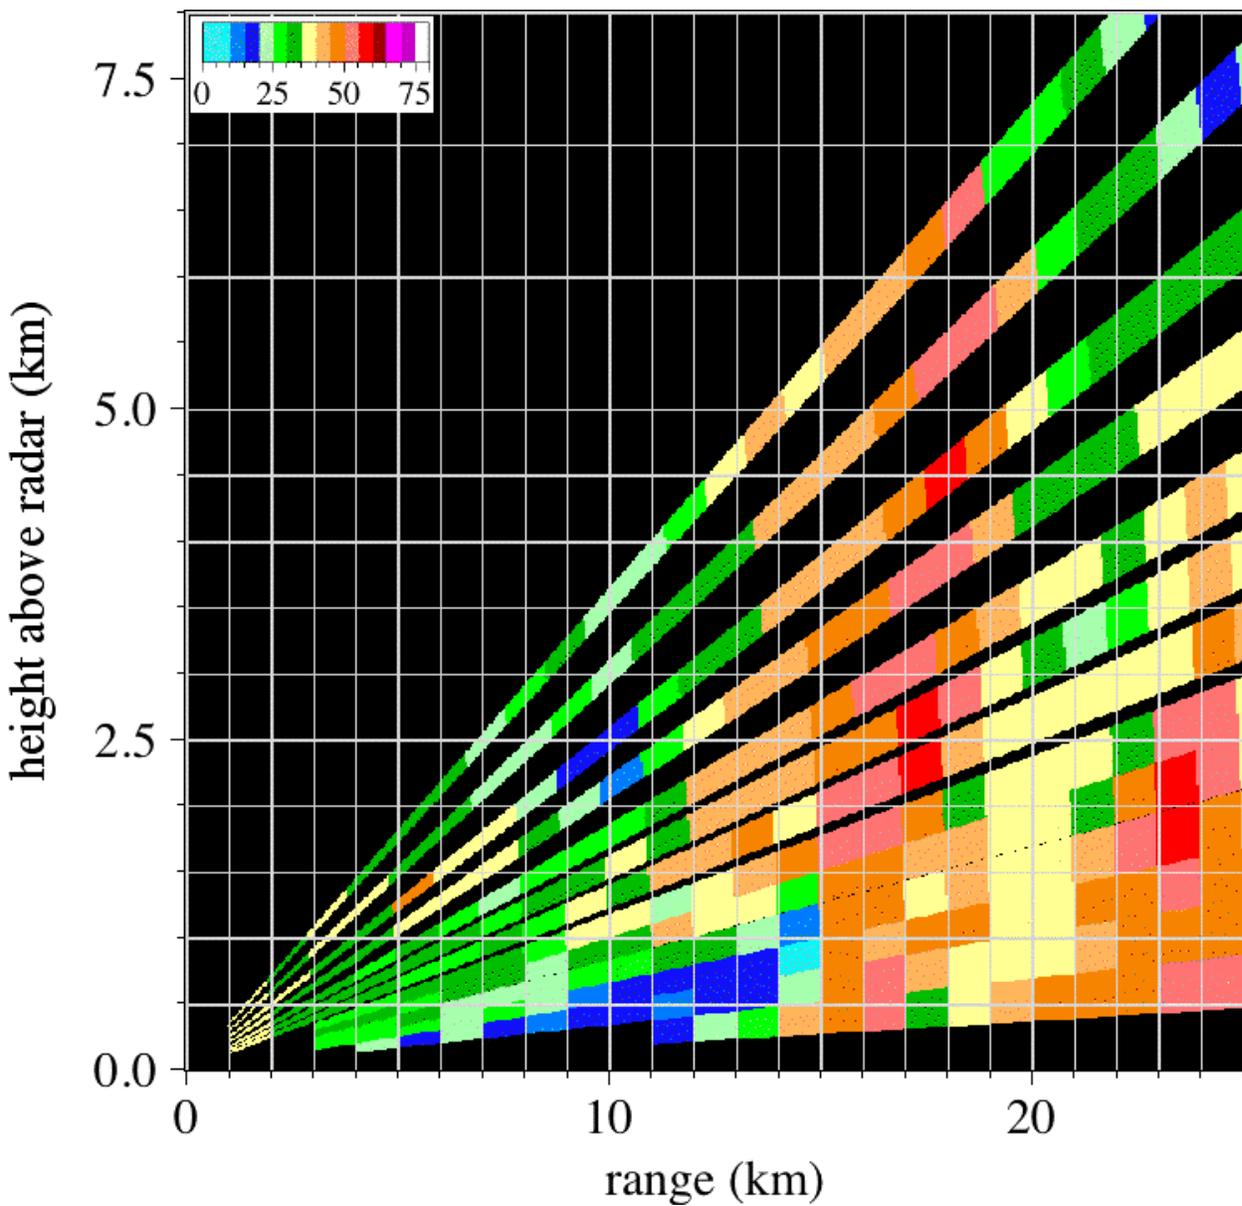


Figure 12. A reflectivity RHI image shows relative size of radar bins (color shaded) and Cartesian grid cells (white lines). Note that the radar bins size is $1^\circ \times 1\text{km}$ and there are gaps between the higher tilts (VCP 11). The Cartesian grid has a horizontal resolution of 1km and vertical resolutions of 500m below 5km height and 1km above 5km height. At ranges of 10 and 15km, one grid cell can cover ~ 3 and ~ 2 bins, respectively. Damping of storm intensities will occur in these regions when a weighted mean analysis scheme is used.

By the end of Q3, 3D mosaic reanalysis for all 8 cases have been completed (Table 1, data sizes are calculated after compression using gzip). The reanalysis data sets have been provided to the CWPDT via ftp.

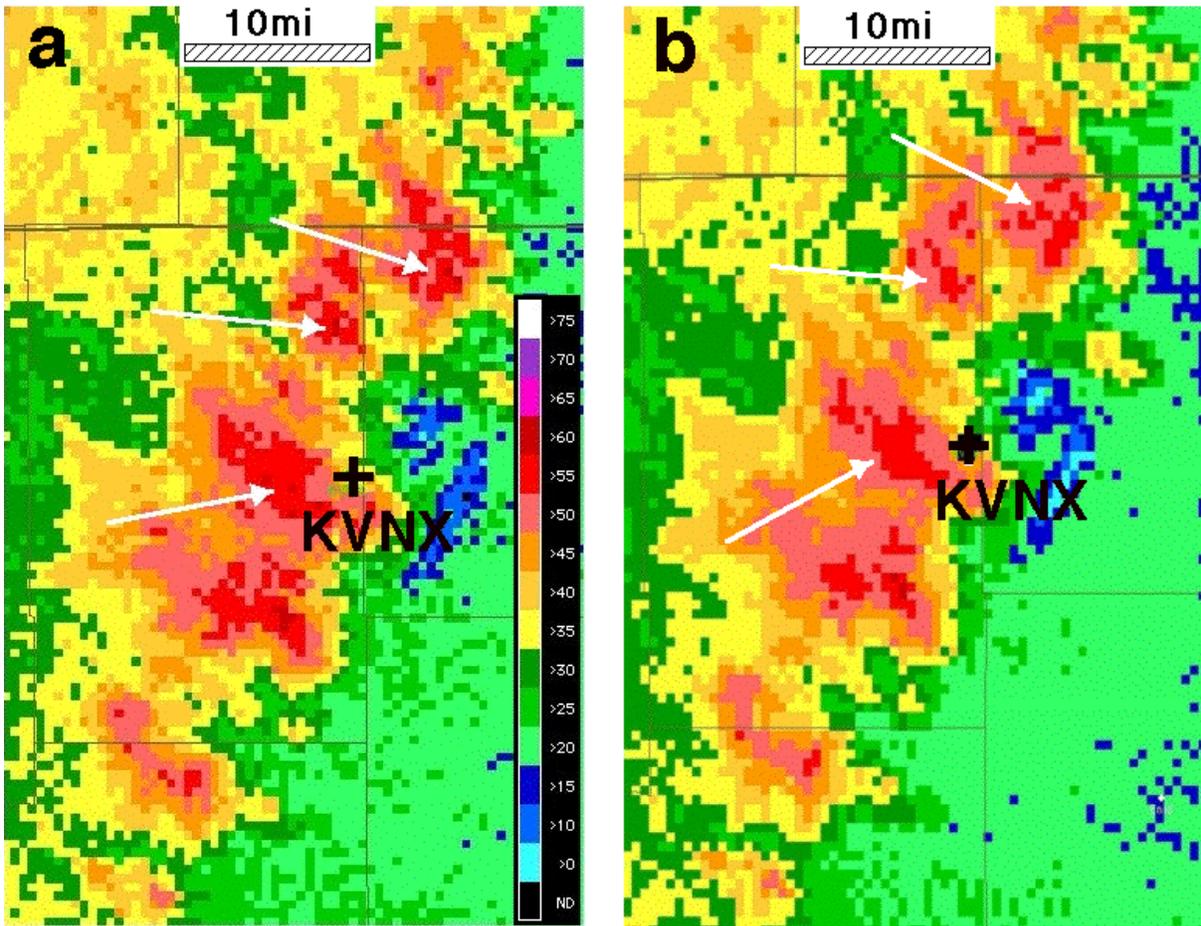


Figure 13. Composite reflectivity fields from KVNx radar (0358Z, May 27, 2002, panel a) and from 3D mosaic grid (panel b) using distance weighted mean analysis. Since the storms are very close to the radar (within 25km of range), the analysis scheme damped the storm intensities (see white arrows).

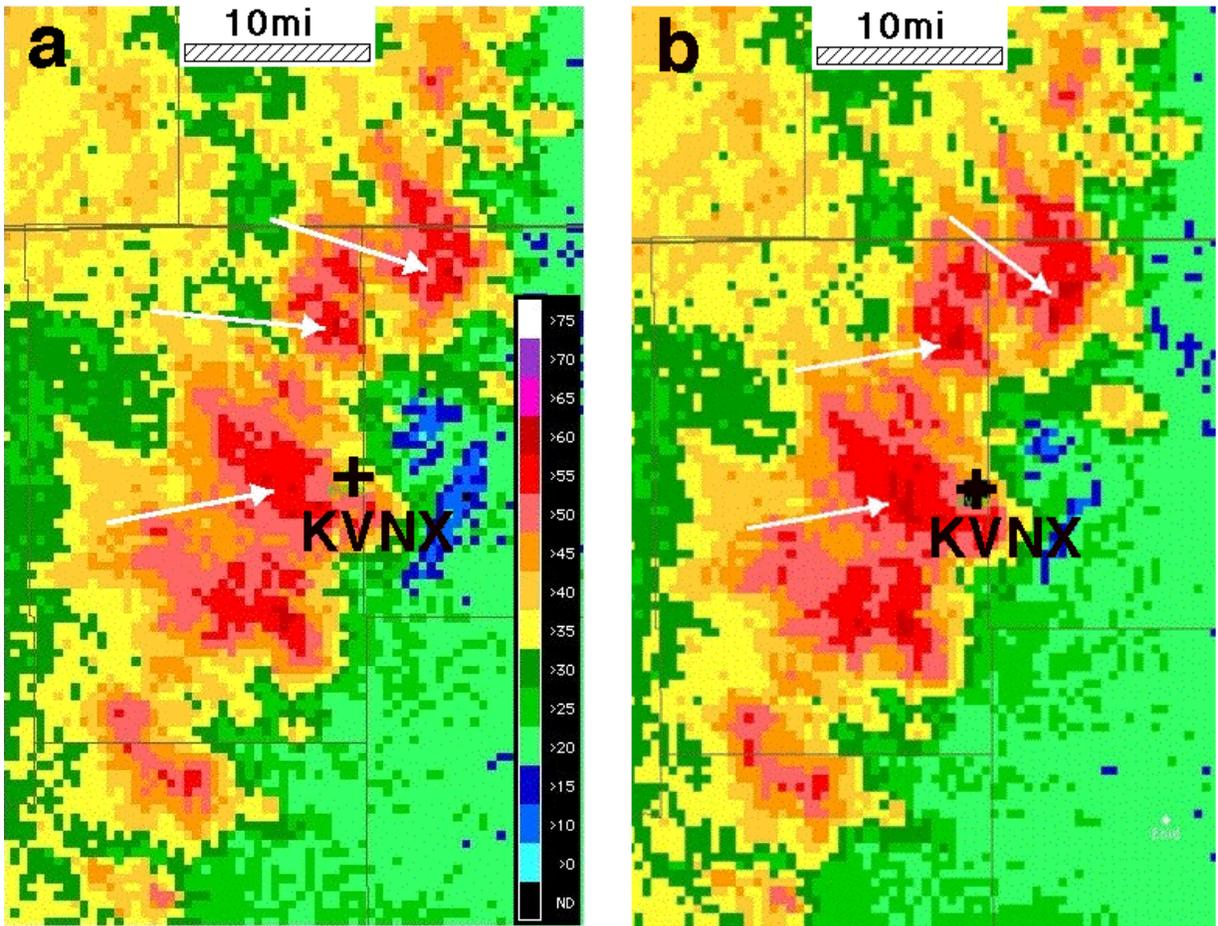


Figure 14. Composite reflectivity fields from KVNx radar (0358Z, May 27, 2002, panel a) and from a 2D mosaic grid (panel b) using the taking maximum approach. The new 2D mosaic scheme successfully retained storm intensities and coverage.

Table 1. Data statistics for the CWPDT cases.

Events	# of hours	# of radars	Data processed (GB*)		
			Level-2	Intermediate (QC&display)	3D mosaic (NetCDF)
2002-06-25	14	42	8.5	7.6	0.85
2002-06-27	13	43	9.0	9.9	1.03
2002-07-08	19	43	6.8	6.1	0.65
2002-07-22	20	42	8.1	7.3	0.63
2002-07-23	13	42	5.0	4.6	0.46
2002-08-19	15	42	7.4	5.8	0.83
2002-08-22	13	42	6.3	6.0	0.57
2003-05-01	15	39	6.4	7.0	0.55
Total:			57.5	54.3	5.57

03.6.14.6-7 A 4-D dynamic data analysis and visualization system.

The formulation for the 4-D dynamic grid (4DDG) has been finalized. The method is defined by the following equation:

$$f_g = \frac{\sum_{n=1}^N w_n \cdot f_n^0}{\sum_{n=1}^N w_n}$$

Where f_g represents the analysis value on a given Cartesian grid, f_n represents an observation value at a radar bin, N is the total number of radar bins in the volume of influence of the grid cell, and w_n is a 4-D weighting function given to the n th observation (i.e., f_n) in the volume of influence. The weighting function is defined by:

$$w = w_{remap} \cdot w_{mosaic} \cdot w_t \cdot w_{calib}$$

Here w_{remap} represents a spatial weighting function for remapping single radar reflectivity field from polar coordinates to Cartesian coordinates, w_{mosaic} and w_t

represent spatial and temporal weighting functions for mosaicking multiple radar reflectivities, respectively, and w_{calib} is an adjustment factor for calibration differences among different radars.

The spatial weighting function, w_{remap} , for remapping single radar reflectivity field onto a Cartesian grid is the same as the 3D mosaic algorithm, that is:

$$w_{remap} = w_r \bullet w_\phi \bullet w_\theta \bullet w_{gapfil}$$

where w_r , w_ϕ , and w_θ are weighting functions in range, azimuth and elevation directions, respectively (ref: Zhang et al., 2001: Three-dimensional gridding and mosaic of reflectivities from multiple WSR-88D radars. Preprints, The 30th International Radar Conference, July 19-24, Munich, Germany, 719-721). w_{gapfil} is a weighting function for the gap-filling scheme (see the NEPDT 2001 2nd and 2002 2nd Quarterly reports).

The pseudo code for the 4DDG has been developed based on the formulation described above. Various weighting functions will be examined and optimal ones will be used for the final 4DDG system prototype.

b) Planned Efforts

The refined 3D mosaic algorithms will be implemented in the real-time CIWS domain. The development of a prototype 4-D dynamic grid will continue.

c) Problems/Issues

Due to the case study work for the CWPDT, the NAPDT has requested some changes to the original FY03 tasks in late March, 2003. This quarterly report is based on the revised task document.

d) Interface with other Organizations

Worked with WWRPDT and CWPDT on the ftp access of the 3D mosaic data for case studies.

e) Activity Schedule Changes

None.

03.6.15 WARP Activities

The WARP is integral to AT controller displays. Warp is significant in that it shifts the burden of displaying weather radar returns to an instrument specifically designed as a weather radar: the WSR-88D. However, due to the nature of its mission and hardware, the WSR-88D cannot take the same approaches to data quality control as do the long-range L-band radars currently used by ATC. New approaches to data quality control need to be developed so the users have confidence in the weather data products displayed to them.

a) Current Efforts

Work continues on providing the National Air Traffic Controllers Association (NATCA) with guidance on the accuracy of ORPG AP (anomalous propagation)-mitigated composite reflectivity products. While the AP removal algorithm can be effective in removing non-precipitation returns, reflectivity returns associated with storm cells can be adversely affected. AP-mitigated composite reflectivity products have been examined in a variety of precipitation events to examine the potential for a reduction in the areal coverage and magnitude of reflectivity returns as compared to composite reflectivity products without AP removal.

Figures 15-23 show examples of WSR-88D Level II reflectivity, ORPG product 36 (0-70,000 composite reflectivity without AP mitigation), and product 96 (0-70,000 ft composite reflectivity with AP mitigation) for three different convective events. Figure 15 shows one of the more striking examples of how the AP-mitigation algorithm can reduce reflectivity returns. In this convective event, numerous small (~20 km in diameter), individual convective cells exist east and northeast of the Corpus Christi, Texas WSR-88D (KCRP) on June 10, 2003, with many of the cells containing regions of >50 dBZ echoes (Figure 16). However, the AP-mitigation algorithm reduces the maximum reflectivity returns by 5-15 dBZ for most storms. In this example, the reduction is very clear because there are numerous cells containing localized (~4 km) regions of >50 dBZ regions shown in red while no cells contain regions of >50 dBZ in the AP-mitigated product (Fig. 17).

In Figure 18, a handful of cells are evident northeast of the Pueblo, Colorado (KPUX) WSR-88D on May 31, 1996. The cells are slightly larger (~30-40 km in diameter) than those shown in Figs. 15-17 and also have correspondingly larger regions of >50 dBZ returns (Figure 19). Once again, the AP-mitigated algorithm reduces the intensity of the cells (Figure 20), although in this case, some >50 dBZ pixels remain. Figure 21 shows reflectivity returns at approximately 0000 UTC, May 4, 1999, associated with a central Oklahoma tornadic outbreak. The mature convective cells in this case are much larger with diameters of 60 km or greater (Fig. 22). In addition, the regions of maximum intensity are also much larger (~20 to 30 km). ORPG product 96 (Figure 23) does not reduce the intensity of the reflectivity returns to the degree it does in the first two cases, however, areal coverage of >50 dBZ returns is still reduced.

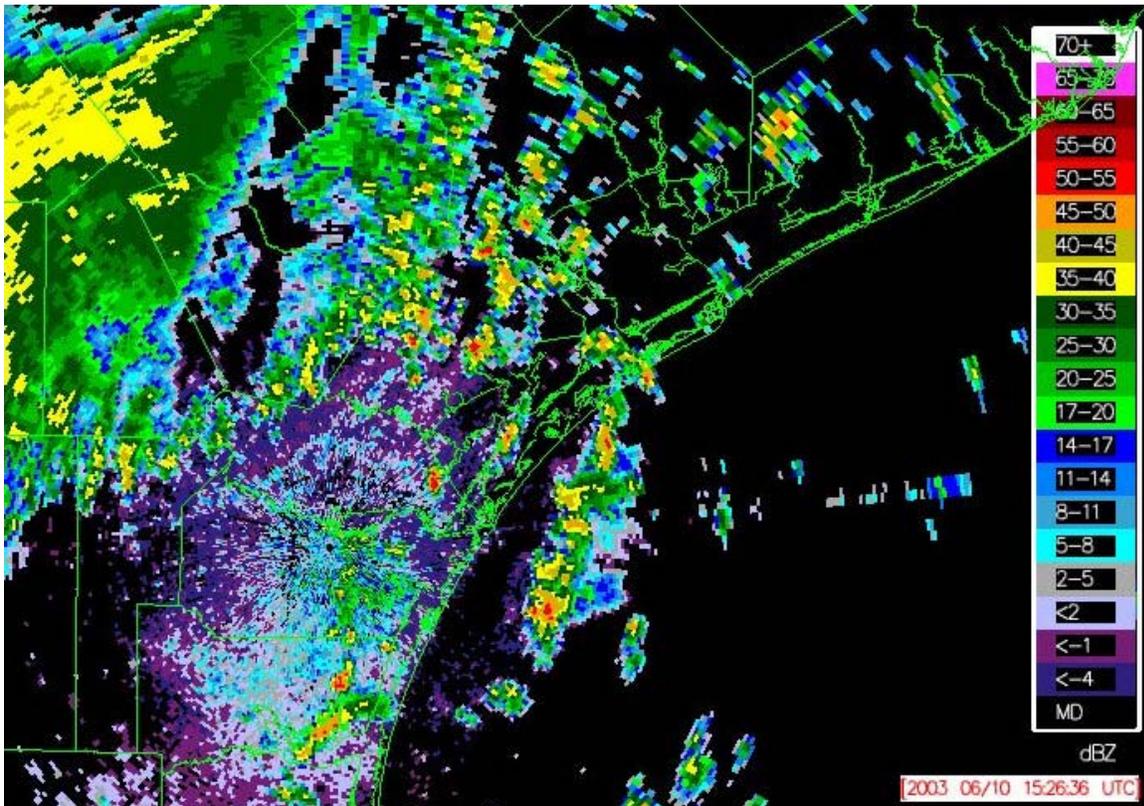


Figure 15. Corpus Christ, TX (KCRP) WSR-88D Level II reflectivity image from 1527 UTC, June 10, 2003.

Subjective examination of AP-mitigated products suggests some smoothing of reflectivity data is occurring causing a reduction in magnitude and areal coverage of maximum reflectivity returns.

b) Planned Efforts

Investigate the ORPG AP-mitigated code to determine if the smoothing is inherent to the AP algorithm or if it can be removed, providing a simple solution to this potentially dangerous problem.

c) Problems/Issues

None.

d) Interface with other Organizations

None.

e) Activity Schedule Changes

None.

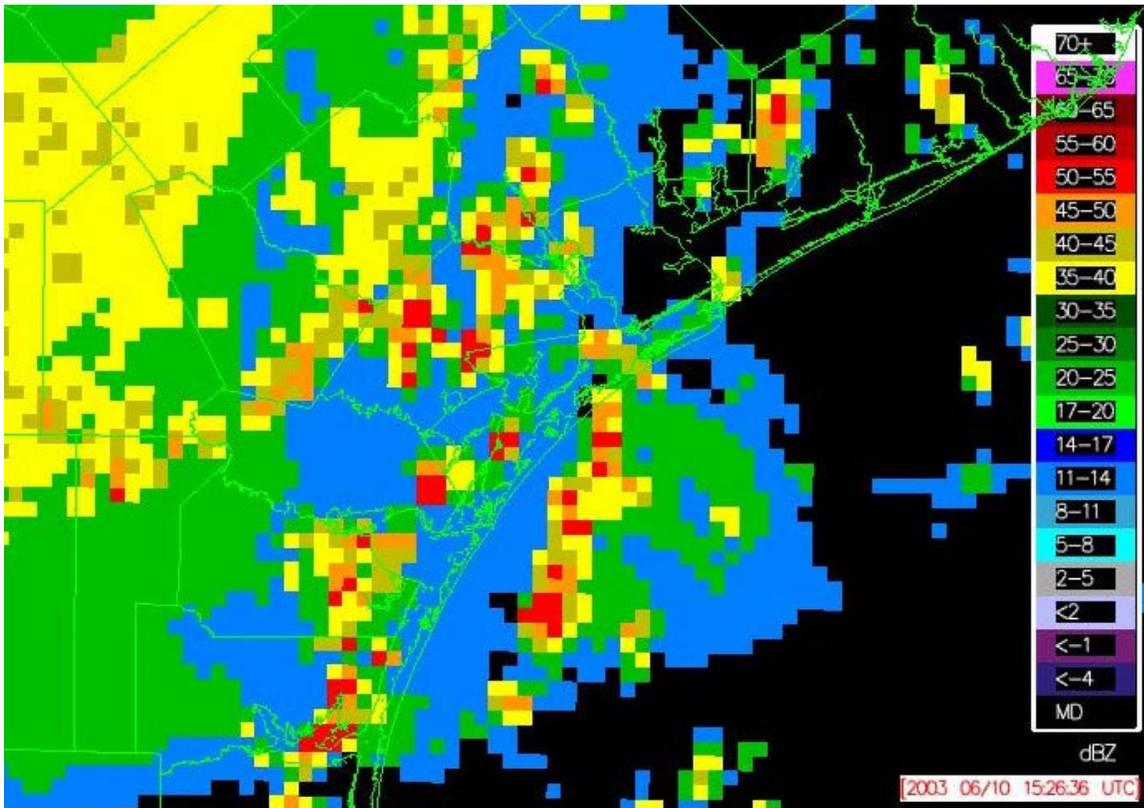


Figure 16. Same as Figure 15, except for ORPG product 36 (0-70,000 composite reflectivity without AP mitigation).

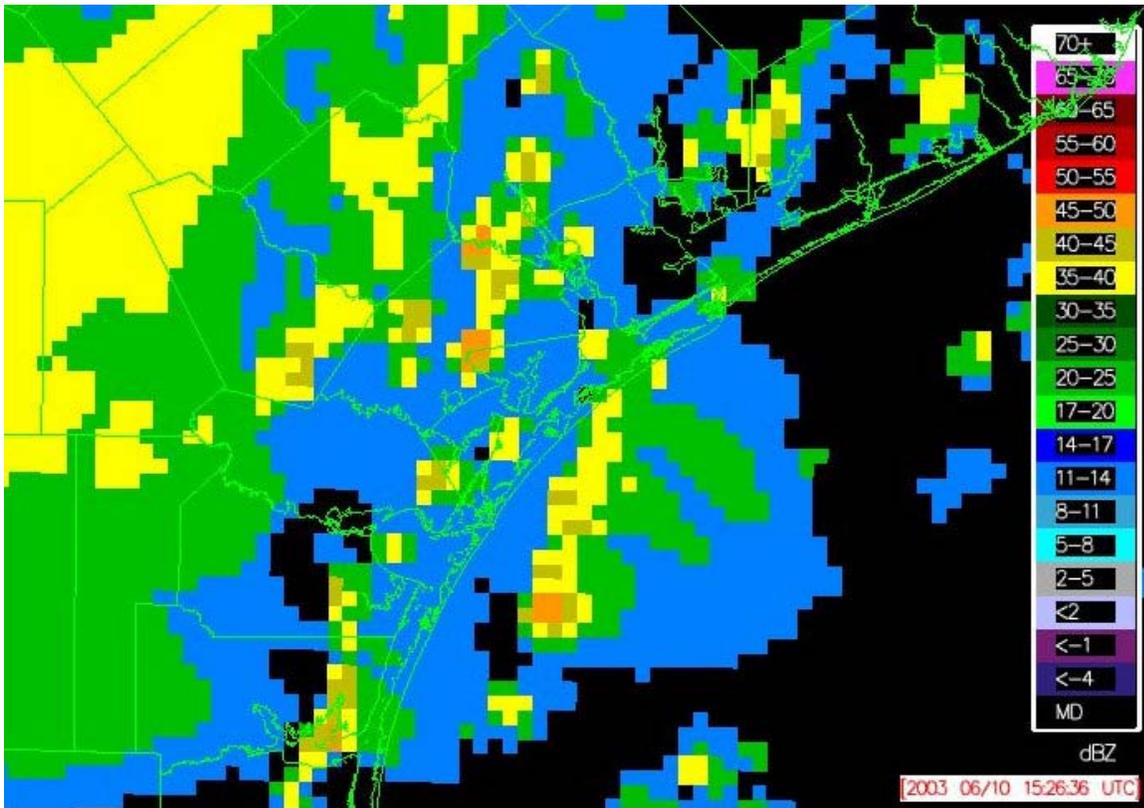


Figure 17. Same as Figure 15, except for ORPG product 96 (0-70,000 composite reflectivity with AP mitigation).

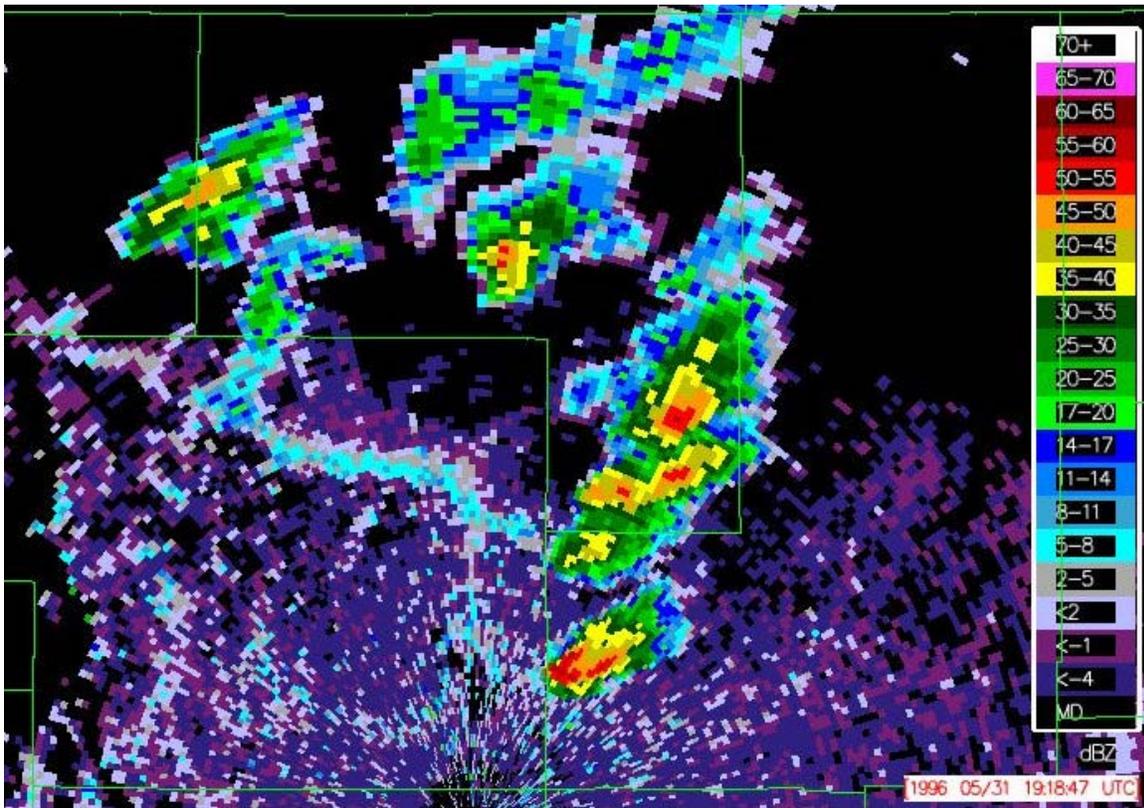


Figure 18. Pueblo, CO (KPUX) WSR-88D Level II reflectivity image from 1919 UTC, May 31, 1996.

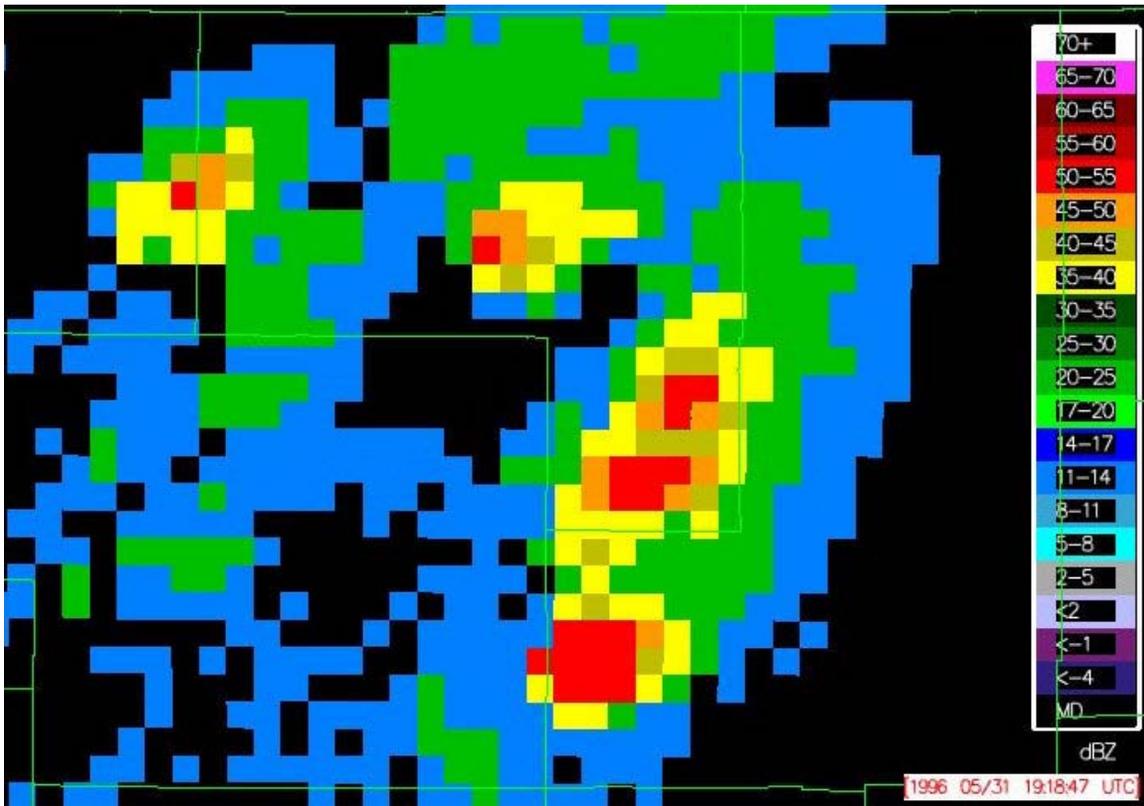


Figure 19. Same as Figure 18, except for ORPG product 36 (0-70,000 composite reflectivity without AP mitigation).

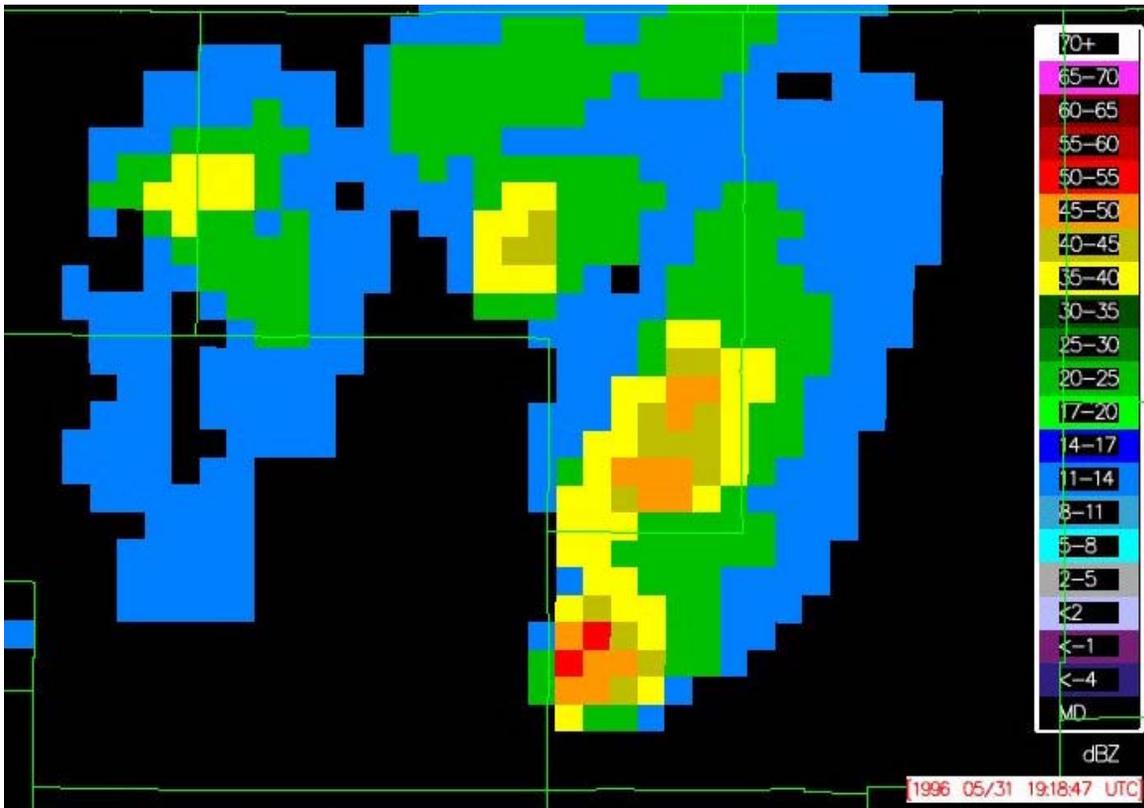


Figure 20. Same as Figure 18, except for ORPG product 96 (0-70,000 composite reflectivity with AP mitigation).

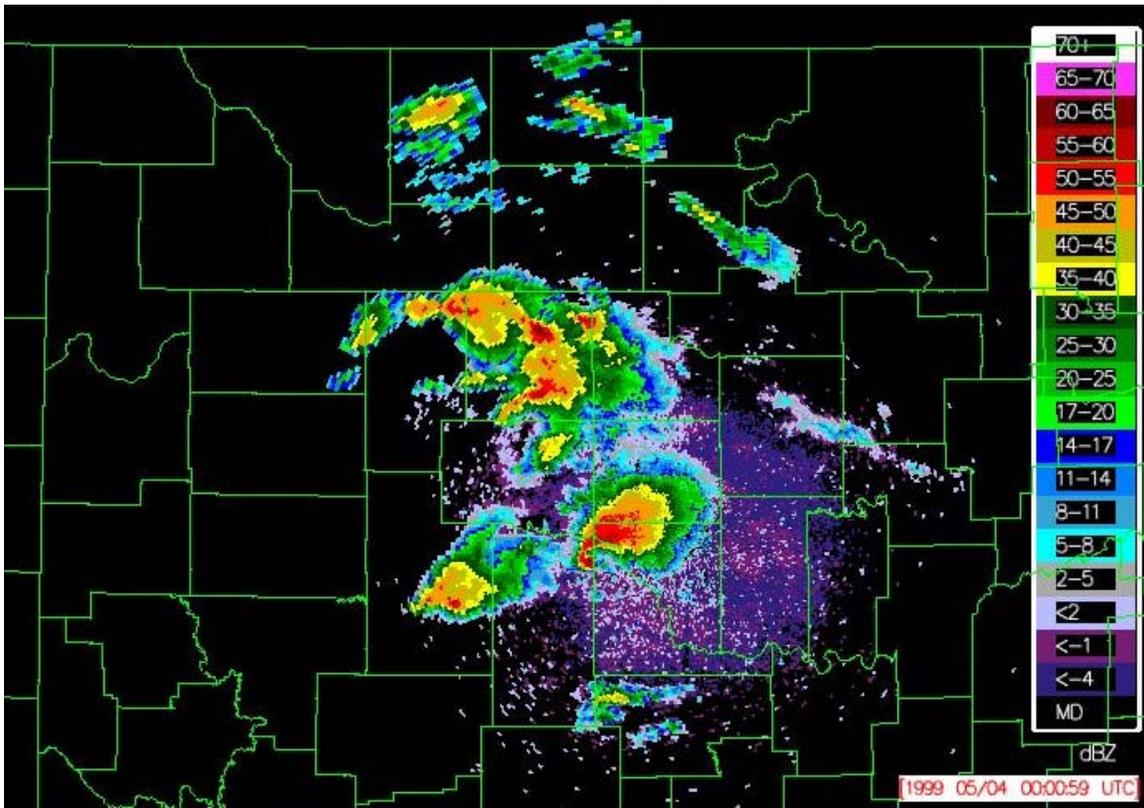


Figure 21. Twin Lakes, OK (KTLX) WSR-88D Level II reflectivity image from 0001 UTC, May 4, 1999.

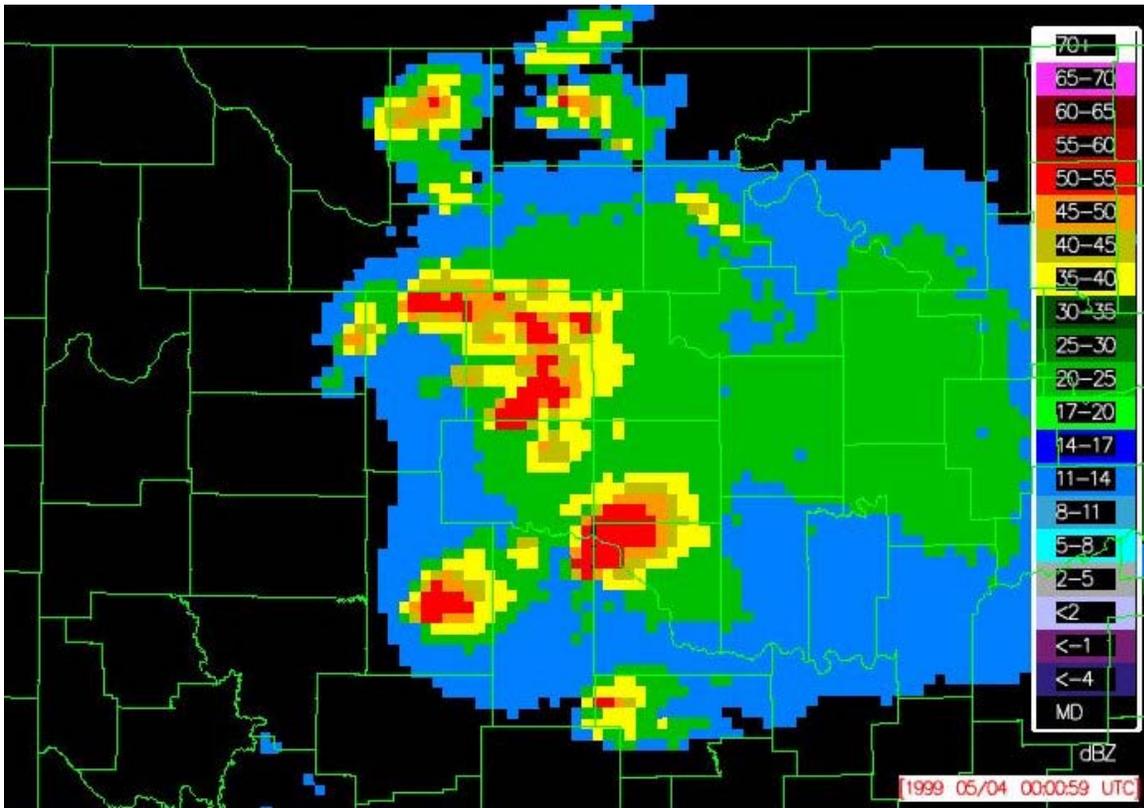


Figure 22. Same as Figure 21, except for ORPG product 36 (0-70,000 composite reflectivity without AP mitigation).

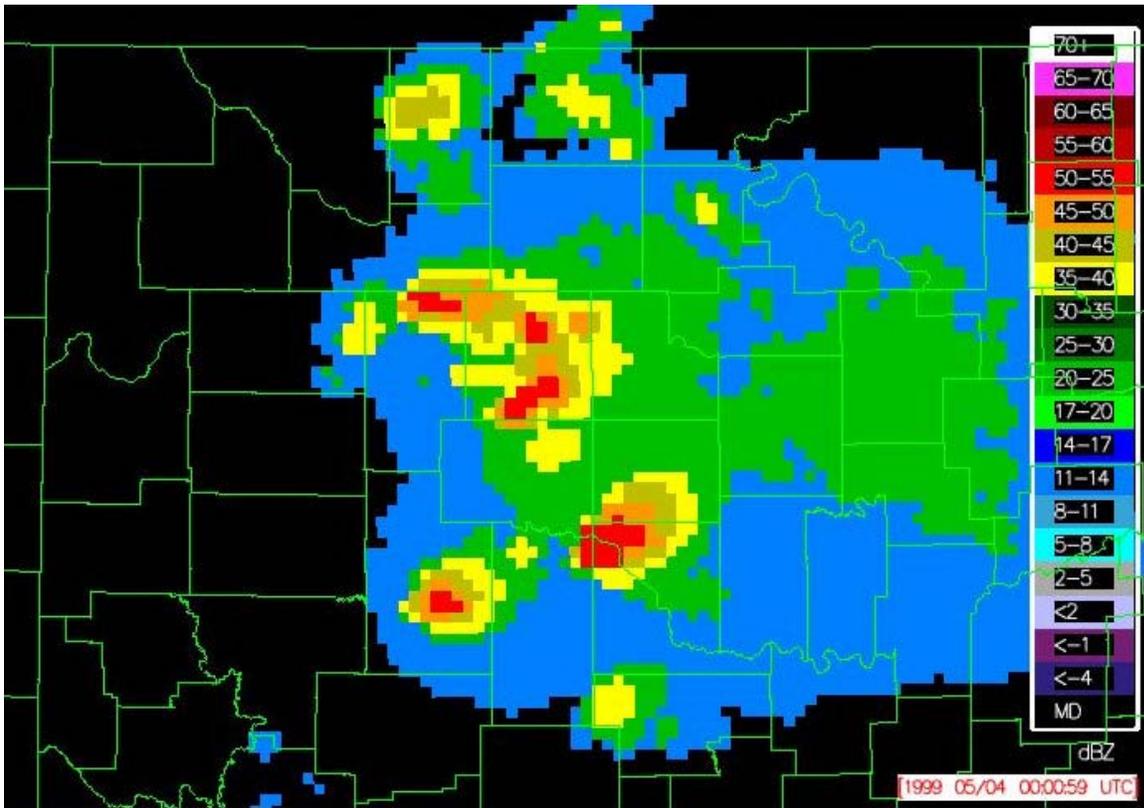


Figure 23. Same as Figure 21, except for ORPG product 96 (0-70,000 composite reflectivity with AP mitigation).

Freezing Level Determinations with Polarimetric Radar: Retrieval Model and Application

Kyoko Ikeda* and Edward A. Brandes
National Center for Atmospheric Research, Boulder, Colorado

1. Introduction

Backscattered signals from dual-polarization radars provide information regarding hydrometeor size, shape, orientation, and thermodynamic phase in precipitating storms. Because polarimetric measurements are particularly sensitive to the presence of large, wetted particles that characterize melting layers, response signatures in the measurements can be used to designate freezing levels. Previous studies to determine freezing levels have focused on reflectivity measurements (White et al. 2002; Mittermaier and Illingworth 2003). In this study we present a freezing level detection algorithm that exploits melting layer signatures in vertical profiles of radar reflectivity (Z_H), linear depolarization ratio (LDR), and co-polar correlation coefficient (ρ_{HV}). Response signals from LDR and ρ_{HV} during hydrometeor phase changes are more pronounced than that for reflectivity, often allowing determination of the freezing level when signatures in Z_H are absent. Moreover, consensus estimates from the three parameters should reduce errors in the estimates compared with those derived solely from Z_H .

The algorithm has been applied to radar data collected from a number of field campaigns. Performance is demonstrated with a multiple freezing level event observed in the Oregon Cascades during the Improvement of Microphysical PaRametrization through Observational Experiments (IMPROVE) II field campaign. Although melting and freezing occur at the wet-bulb 0°C level, the current algorithm is designed to retrieve dry-bulb zeroes. For saturated environments, as expected with upslope conditions in the Cascades, both temperatures occur at the same height. For unsaturated conditions, wet-bulb zeroes occur at a lower level. Examination of high resolution aircraft data from several events occurring during other field programs revealed that wet-bulb zero depressions were 0-105 m. The accuracy in freezing level heights deduced with the proposed algorithm is believed to be 100-200 m.

2. Overview of the freezing level detection algorithm

Typical vertical radar profiles through the melting layer disclose radar reflectivity (Z_H), linear depolarization ratio (LDR), and differential reflectivity (Z_{DR}) maxima and a co-polar correlation coefficient (ρ_{HV}) minimum just below the 0°C level (Fig. 1). The onset of melting changes the density and dielectric factor of

frozen hydrometeors (Fabry and Zawadzki 1995) causing them to behave as raindrops of equivalent size and increasing their reflectivity (Battan 1973; Chapter 10). Further increase in Z_H results as wetted hydrometeors become sticky and aggregation occurs. Eventually, melting hydrometeors collapse into raindrops and an increase in terminal velocity removes them from the sample volume. The net effect is a rapid decrease in reflectivity.

Melting snowflakes and aggregates wobble as they fall creating a distribution of canting angles. Particle canting causes a small portion of the transmitted energy to "leak" into the orthogonal direction. The leakage is enhanced for large mixed-phased particles and creates a distinct LDR maximum in the melting layer. Correlations between horizontally and vertically polarized signals are typically close to unity for pristine ice crystals and raindrops. However, changes in hydrometeor shapes and the presence of mixed-phased precipitation cause ρ_{HV} to decrease to <0.93 in the melting layer. Snow aggregates do not have preferred orientations; therefore they have small Z_{DR} (<0.5 dB). Raindrops are flattened and tend to orient themselves with their major axes near horizontal, causing Z_{DR} to be 0.3–4 dB. Often Z_{DR} has a maximum value in the melting layer indicating large aspect ratios associated with partly-melted hydrometeors.

The principal idea behind the freezing level detection algorithm lies in identifying the heights at which the melting layer signature extremes discussed above exist and using statistical relationships between the signatures and the 0°C level. Fig. 2 shows the depression of melting layer signature extremes for LDR, ρ_{HV} , and Z_{DR} from the height of the reflectivity bright band maximum. The data are from over 300 radar profiles obtained from constant antenna elevation scans collected on 5 September 1998 during the PRECIP98 field campaign. Vertical distributions of each parameter were calculated by averaging measurements over 10-degree sectors at elevation angles between 4 and 12 degrees. Examination reveals that LDR and ρ_{HV} extremes typically occur 200 m below the Z_H maximum while the maximum Z_{DR} occurs 200-300 m below the Z_H maximum. The depression distributions are narrow for LDR and ρ_{HV} , and their extremes generally occur at the same heights. The depression distributions for Z_{DR} , on the other hand, are broad and often 500 m or more below that for reflectivity. Also, there is no melting layer signature for Z_{DR} at vertical incidence or in many convective situations. Consequently, only Z_H , LDR, and ρ_{HV} are used for designating the 0°C level in the proposed algorithm.

* *Corresponding author address:* Kyoko Ikeda,
 National Center for Atmospheric Research, P.O. Box
 3000, Boulder, CO 80307
 E-mail: kyoko@ncar.ucar.edu

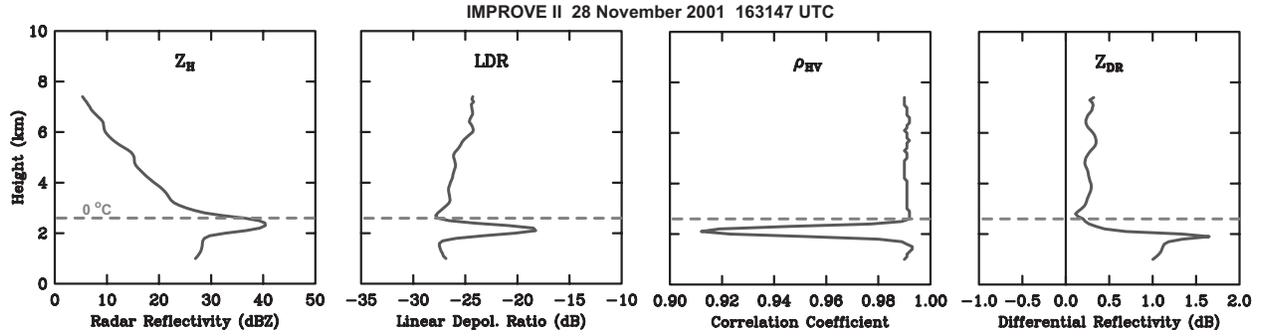


Figure 1: Profiles of radar reflectivity (dBZ), linear depolarization (dB), correlation coefficient, and differential reflectivity (dB) from 28 November 2001 163147 UTC. The 0°C level at 2.53 km is an average of two soundings.

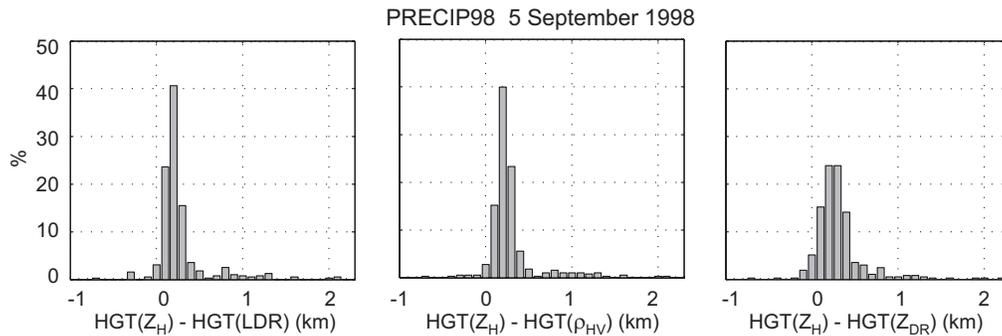


Figure 2: Frequency distributions showing the offsets of the LDR, ρ_{HV} , and Z_{DR} extremes [HGT(LDR), HGT(ρ_{HV}), and HGT(Z_{DR})] from the Z_H maxima [HGT(Z_H)].

A composite of LDR and ρ_{HV} depressions from PRECIP98 (Florida), STEPS (eastern Colorado), CASES (Kansas), TRMM-LBM (Brazil), MAP (Italy), and IMPROVE (northwestern U.S.) field projects suggests that offsets of the LDR and ρ_{HV} extremes from the Z_H maximum are typically 200 m for warm season datasets, whereas an offset of 100 m is likely for cold season datasets (MAP and IMPROVE). However, examination of individual profiles reveals that larger offsets associate with more intense precipitation.

Fig. 3 shows model profiles developed from the observed relations of melting layer signatures discussed above. Freezing level designations begin by identifying the melting layer extremes in the observed vertical profiles. Extremes in the observed and modeled profiles are aligned, and the correlation coefficient between the observed and modeled profile is calculated to determine the degree of fit. When the observed distribution roughly matches the melting layer signature in the model, the correlation is high. Melting layer signatures are accepted by the algorithm when correlation coefficients are greater than 0.7. This threshold value was chosen because one half of the variability in the observations is explained by the model at this value. Freezing level determination is then made by knowing the statistical offset between the typical 0°C level and the heights of the Z_H , LDR, and ρ_{HV} extremes. The offset was 300 m for Z_H and 500 m for LDR and ρ_{HV} for the PRECIP98 dataset. The offsets of

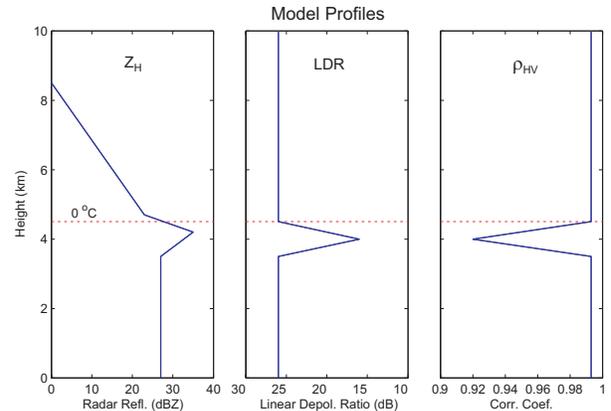


Figure 3: Sample model vertical profiles of radar reflectivity, linear depolarization ratio, and correlation coefficient.

200 m for Z_H and 300 m for LDR and ρ_{HV} were used for the IMPROVE dataset. Then the estimated freezing level height for the i th parameter (with 1 for Z_H , 2 for LDR, and 3 for ρ_{HV}) is

$$h_i = ht_i + offset_i \quad (1)$$

for the height of the melting layer signature extreme ht_i . Resulting estimates from the three parameters typically vary. A consensus (weighted) estimate is then computed from

$$h_{fzlv} = \frac{\sum_{i=1}^n h_i r_i^2}{\sum_{i=1}^n r_i^2} \quad (2)$$

where n is the number of parameters with a correlation coefficient (r) above the threshold value, and h_i and r_i are the freezing level height and correlation coefficient for the i th parameter, respectively. When only one profile meets the threshold criteria, the consensus freezing level height is not calculated. The measure of scatter among the estimates is expressed with the standard deviation σ ,

$$\sigma = \sqrt{\frac{1}{n} \sum_{i=1}^n (h_i - h_{fzlv})^2} \quad (3)$$

This value serves as confidence factor for the retrieval.

3. Application of the freezing level algorithm

In this section algorithm performance is demonstrated with a dataset collected over the Oregon Cascades during the IMPROVE II field campaign. On 28 November 2001 a large synoptic low pressure system approached the Cascades region from the northeastern Pacific Ocean. By 1200 UTC, the low pressure center was located to the west of Vancouver Island, British Columbia and the associated warm and cold fronts were advancing across the northwestern coastal U.S.

Radar data indicated that warm frontal precipitation over the IMPROVE domain was most intense between 0930 and 1000 UTC. Soon afterward the radar displayed multiple freezing levels for 20–40 minutes. Fig. 4 is a vertical cross-section from 1009 UTC for an azimuth of 250° . In the figure the lower melting layer is located ~ 0.7 km above ground level (AGL). The elevated melting layer is at ~ 2.3 km AGL. The figure also shows a sloped band of high reflectivity between 1.5 km AGL and the surface at a range of 11–17 km. [The slope is different from that expected for a rain streak because the horizontal wind was generally from the southwest (toward the right in the figure). Consequently, the sloping band is also believed to be connected with melting hydrometeors.]

Fig. 5 is a time series of Z_H and LDR profiles between 0906 and 1121 UTC. Consensus freezing levels deduced from the algorithm are superimposed. ρ_{HV} profiles are not shown because the identified melting layers were essentially the same as those from LDR measurements. The profiles were constructed from range-height-indicator scans (RHI) by averaging measurements 3 km in the horizontal and 0.2 km in the vertical at a distance of 8 km for the same azimuth as in Fig. 4. Multiple freezing levels were detectable for all quadrants in which RHI scans were performed, but measurements from the azimuth of 250° were least contaminated by ground clutter from nearby mountains. Fig. 5 suggests that the lower freezing level rose from 1.05 km at 0906 UTC to 1.3 km MSL at 1009 UTC. Afterward, it descended slightly, dissipating by 1029 UTC.

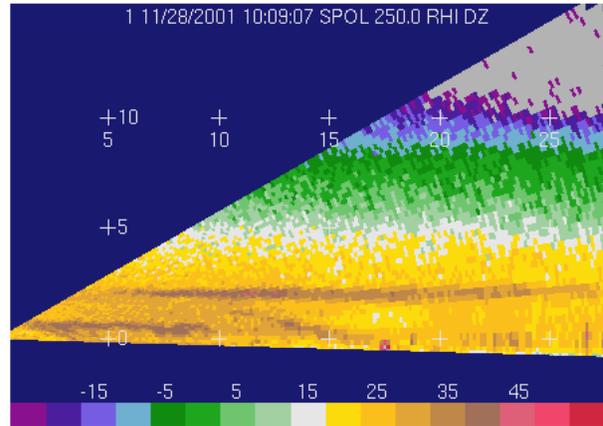
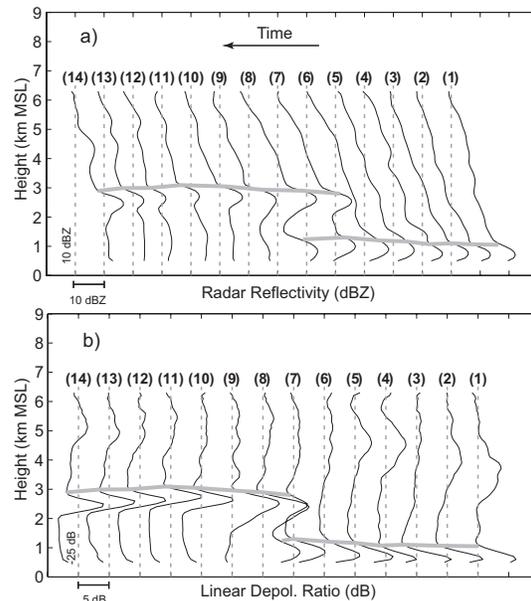


Figure 4: Vertical cross-section of radar reflectivity at an azimuth of 250° . The ordinate is height above ground level. The radar site is at an elevation of 0.457 km.



(1) 0906 UTC (4) 0937 UTC (7) 1009 UTC (10) 1040 UTC (13) 1110 UTC
(2) 0917 UTC (5) 0948 UTC (8) 1019 UTC (11) 1050 UTC (14) 1121 UTC
(3) 0927 UTC (6) 0958 UTC (9) 1029 UTC (12) 1100 UTC

Figure 5: Vertical profiles of a) radar reflectivity (dBZ) and b) linear depolarization ratio (dB) from 0906 to 1121 UTC at an azimuth of 250° and a range of 8 km. Gray bold lines are consensus freezing level heights from Z_H , LDR and ρ_{HV} . Dashed vertical lines represent a) 10 dBZ and b) -25 dB with respect to each profiles. Tick mark increments are equivalent to a) 10 dBZ and b) 5 dB.

An elevated freezing level first appeared in the Z_H and LDR profiles at 1009 UTC and a height of 2.73 km MSL. The Z_H signature for the elevated freezing level was the strongest at this time and weakened as precipitation decreased. The LDR signature intensified to -17.3 dB by 1029 UTC. According to the consensus estimates, the elevated freezing level rose initially to 3.09 km at 1050 UTC and descended to 2.89 km MSL by 1121 UTC.

Synoptic analyses indicate that evolution of the two freezing levels responded to motion of the low pressure center and its associated warm front. While the surface warm front was well upstream of the radar, a single freezing level was low to the ground. As the warm air advanced, the elevated freezing level appeared. The lower freezing level eventually dissipated as the surface warm front passed the profile location, and only the elevated freezing level remained. A frontal temperature inversion near the 0°C level was observed with a sounding released at 1100 UTC from Salem, Oregon (60 km northwest of the radar site). The sounding found the upper and lower freezing levels to be 2.70 and 1.03 km MSL, respectively, in agreement with the estimated freezing level heights.

Over the next six hours the elevated freezing level fell an additional 0.3 km. With cold frontal passage between 1700-1800 UTC, the freezing level rapidly descended to 1.7 km MSL by 2200 UTC (Fig. 6). Data points in Fig. 6 are estimates of freezing level heights over a gridded domain (35x35 km) in the eastern sector of the radar. The 0°C levels from observations are also plotted. Generally, freezing level heights from Z_H have larger variability than those from other parameters. Scatter in the designations are more pronounced after 1730 UTC for all parameters. This is due to broadening freezing level height distributions over the domain from weakening precipitation and descending freezing level caused by passage of the surface cold front.

4. Summary and concluding remarks

A freezing level detection algorithm that utilizes polarimetric measurements has been described. The algorithm looks for melting layer signatures in the radar parameters and applies a statistical offset between the 0°C level and the signature heights. Determinations are typically good when radar signals are strong. Moreover, sensitivity of LDR and ρ_{HV} to mixed-phased hydrometeors often make designations possible when signatures are absent in Z_H measurements. The error with the algorithm is on the order of 100 to 200 m. Algorithm enhancement is possible by taking into account the precipitation intensity dependencies of the 0°C-depressions. The algorithm readily makes retrievals for distances to ~50 km. At greater distances radar beam broadening make precise determinations difficult.

The algorithm was applied to a dataset collected over the northwestern U.S. on 28 November 2001. Evolution of the two observed freezing levels responded to the movement of the synoptic low pressure center and its associated fronts. Initially, a single freezing level existed near ground. An advancing warm front warmed the lower atmosphere and created the elevated freezing level. Passage of the surface front and continued warming eventually overwhelmed the lower melting layer. Later, cold front passage caused a rapid lowering of the elevated freezing layer.

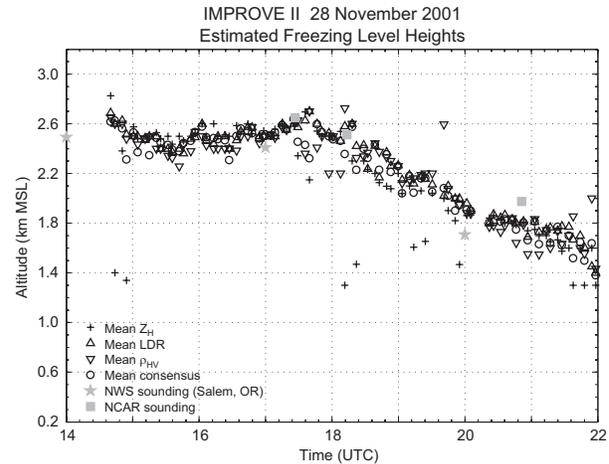


Figure 6: Time series of freezing level estimates for 1400-2200 UTC on 28 November 2001. The estimates from Z_H , LDR, and ρ_{HV} were obtained over a gridded domain in the eastern sector of the radar. "Mean" of each parameter is the weighted mean over the domain, computed using the correlation coefficient between the model and observed vertical profile. "Mean consensus" values were obtained similarly, except the reciprocal of σ from Equation 3 was used for weighting. The NCAR sounding location on this day was ~52 km southwest of the radar site, at an elevation of 0.24 km.

Acknowledgement: This work was supported by the Federal Aviation Weather Research Program and is in response to requirements of the FAA. The views expressed are those of the authors and do not necessarily represent the official policy of the U.S. government.

References

- Battan, L. J., 1973: *Radar Observation of the Atmosphere*, University of Chicago Press, 324pp.
- Fabry, F., and I. Zawadzki, 1995: Long-term radar observations of the melting layer of precipitation and their interpretation. *J. Atmos. Sci.*, **52**, 838-851.
- Mittermaier, M. P., and A. J. Illingworth, 2003: Comparison of model-derived and radar-observed freezing-level heights: Implications for vertical reflectivity profile-correction schemes. *Quart. J. Roy. Meteor. Soc.*, **129**, 83-95.
- White, A. B., D. J. Gattas, E. T. Strem, F. M. Ralph, and P. J. Neiman, 2002: An automated brightband height detection algorithm for use with Doppler radar spectral moments. *J. Atmos. Oceanic Technol.*, **19**, 687-697.

P2C.9

DETECTION OF ROTATION AND BOUNDARIES USING TWO-DIMENSIONAL, LOCAL, LINEAR LEAST SQUARES ESTIMATES OF VELOCITY DERIVATIVES

Travis M. Smith^{1,2,*}, Kimberly L. Elmore^{1,2}, and Gregory J. Stumpf^{1,2}, and V. Lakshmanan^{1,2}

¹Cooperative Institute for Mesoscale Meteorology Studies, University Of Oklahoma, Norman, Oklahoma.

²NOAA/National Severe Storms Laboratory, Norman, Oklahoma.

1. Introduction

Traditional methods of calculating rotational shear from Doppler radial velocity data can give results that vary widely from the true value of shear for a vortex. Some factors that must be considered include noisy data and the offset of sample volumes from the center of rotation. This work illustrates preliminary results from a two-dimensional, local, linear least squares (LLSD) method to minimize the large variances in rotational and divergent shear calculations.

Besides creating greater confidence in the value of intensity of meteorological features that are sampled, the LLSD method for calculating shear values has several other advantages. The LLSD removes many of the radar dependencies involved in the detection of rotation and radial divergence (or radial convergence) signatures. Thus, these derivatives of the radial velocity field may be viewed in three-dimensional space or used as input to multi-sensor meteorological applications that are not single-radar based. Additionally, fields of these radial estimates of rotation and divergence have specific signatures when boundaries or circulations are sampled. This manuscript describes how the derivatives are calculated as well as how the rotational LLSD compares with the less-robust (but operationally used) "peak-to-peak" estimates of azimuthal shear. The accompanying poster presentation will describe the divergent LLSD and examples of data mining techniques that use LLSD for boundary and rotation detection.

2. Local, Linear, Least Squares Derivatives

Elmore et al. (1994) describe a method for estimating divergent shear from single Doppler radar data for use in calculating headwind loss estimates for aircraft that encounter microbursts. The rotation portion of the derivative was also

derived by Elmore et al., but not utilized for microburst detection. The LLSD technique was implemented in NSSL's Damaging Downburst Prediction and Detection Algorithm (Smith et al. 2003) for detecting low-level outflows and mid-level convergence and rotation in storm cells. Mitchell and Elmore (1998) first explored the uses of the LLSD for identifying regions of high shear in mesocyclones and tornadic vortex signatures.

Elmore et al. (1994) show that the estimates of radial divergence (u_r) and rotational shear (u_s) can be calculated on a local neighborhood surrounding each range gate, where:

$$u_r = \frac{\sum i V_{ij} w_{ij}}{\Delta r \sum i^2 w_{ij}}, \text{ and}$$

$$u_s = \frac{\sum s_{ij} V_{ij} w_{ij}}{\sum (\Delta s_{ij})^2 w_{ij}}$$

Here, V_{ij} is the radial velocity, Δr is the pulse volume width, s_{ij} is the azimuthal distance from the center of the kernel to the point (i,j), and w_{ij} is a uniform weight function. Because u_r and u_s are derived from only the radial component of the wind, they are approximations of one half the horizontal divergence and vertical vorticity ("half vorticity", hereafter), respectively, assuming a symmetric wind field.

In order to make LLSD calculations on a field of radial velocities, the data are passed through a 3x3 median filter to reduce noise. Then, because the LLSD calculations require a complete kernel of data in order to produce a result, missing radial velocities are filled in with the median of the four adjacent range gates. Finally, the fields of u_r and u_s are calculated.

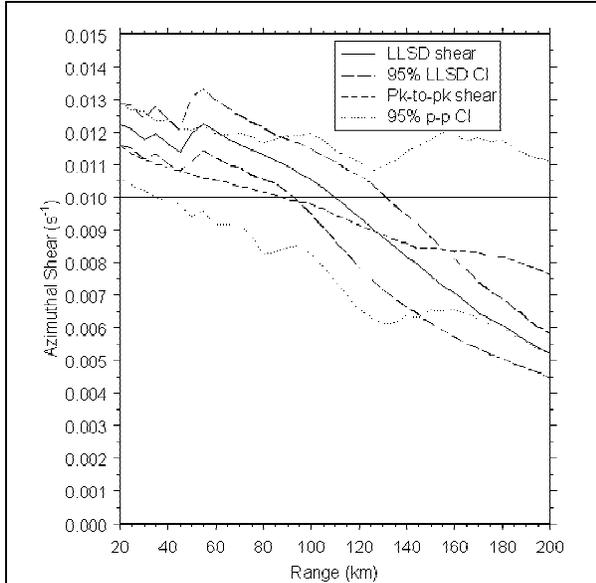


Figure 1: Mean and 95% confidence limits for LLSD and peak-to-peak estimates of azimuthal shear at the vortex center for a 5 km diameter vortex with half vorticity of 0.01 s^{-1} .

The size of the kernel that is used in the calculation is variable, and is described below.

3. LLSD of Rotational Shear

LLSD of rotational, or azimuthal, shears are calculated for simulated circulation signatures of different sizes and at different ranges from a hypothetical radar in order to compare with traditional methods of estimating the strength of circulations. We use a Rankine combined vortex model to generate simulated circulation signatures in the Doppler radial velocity field (Wood and Brown 1997). We superimpose 2 ms^{-1} uniform noise on the radial velocity field to test the robustness of the LLSD calculations.

We compare the LLSD values to the more traditional “peak-to-peak” azimuthal shear calculation, given by

$$u_{as} = \frac{V_{max} - V_{min}}{d}$$

where V_{max} and V_{min} are the maximum outbound and minimum inbound radial velocities (on opposite sides of a circulation), respectively, and d is the distance between those two peaks. For the LLSD calculations, we choose three different kernel sizes that are each 3 range gates deep and approximately 2500 m, 5000 m, or 8000 m

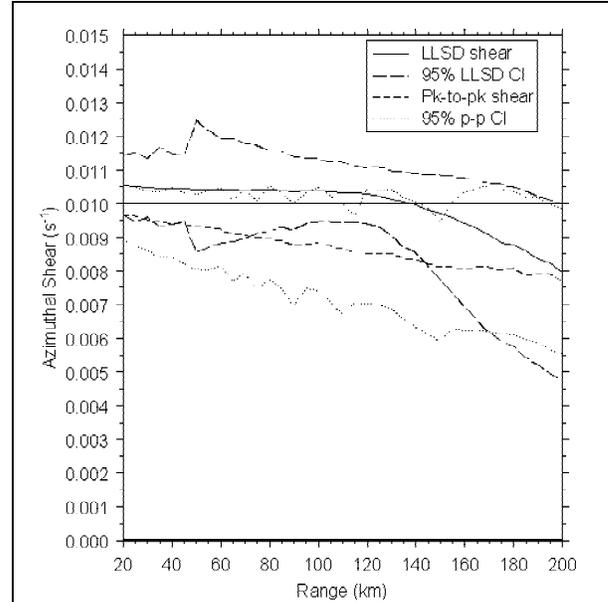


Figure 2: Mean and 95% confidence limits for LLSD and peak-to-peak estimates of azimuthal shear at the vortex center for a 8 km diameter vortex with half vorticity of 0.01 s^{-1} .

wide. Thus the number of radials used in the calculation varies with range from the radar, although a minimum of three radials of data are required for a complete calculation. Kernels that use a fixed number of radials at all ranges usually only provide good shear estimates in a small percentage of those ranges.

To test the variability of the three LLSD kernels, we generate synthetic radial velocity signatures of vortices at ranges every 5 km from 20 km to 200 km. Because radar data suffer from many imperfections, including noise and sampling issues that can affect azimuthal shear values (Wood and Brown 2000), 1000 vortices of the same size and strength are generated at each range, each with different noise patterns and azimuthal offsets to the center of the simulated vortex. This allows for calculation of mean azimuthal shear values and 95% confidence intervals for the three LLSD kernels and u_{as} . We test these calculations on simulated vortices of three different diameters: 2.5 km, 5 km, and 8 km.

Figure 1 shows the 2500m LLSD kernel and peak-to-peak azimuthal shear estimates for a 5 km diameter vortex with half vorticity of 0.01 s^{-1} . In this case, the mean LLSD value is within about 20% of the true value out to about 140 km, with a much smaller variance than that of the peak-to-peak azimuthal shear calculations.

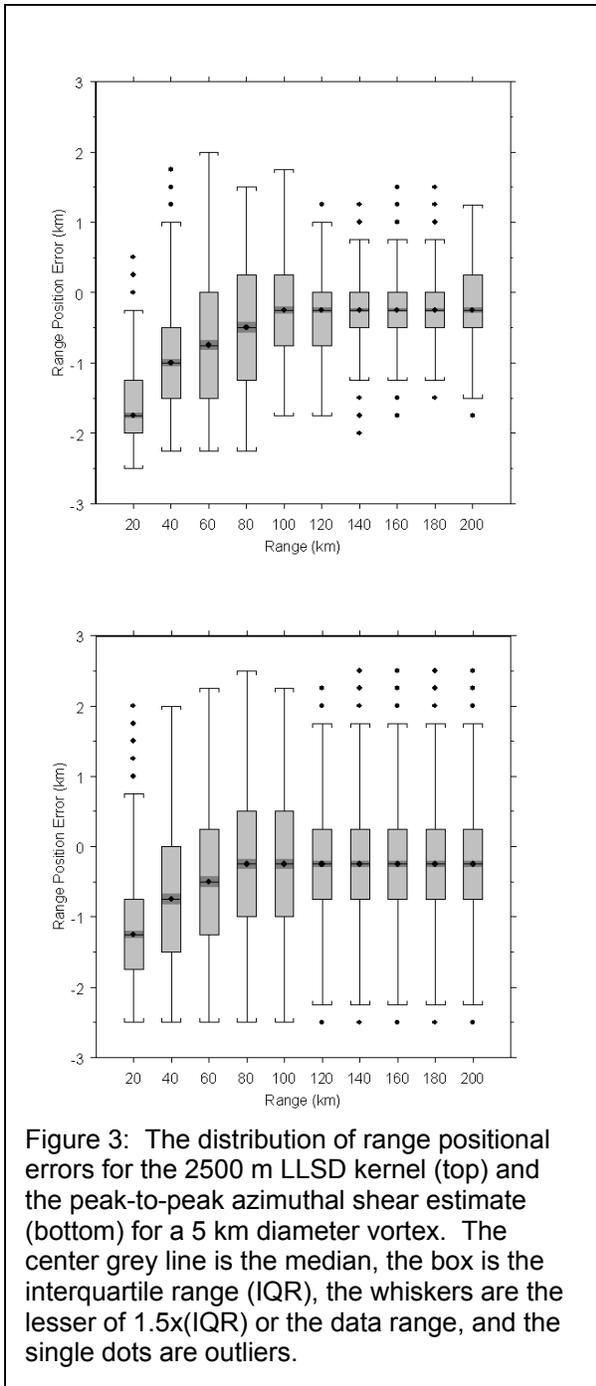


Figure 3: The distribution of range positional errors for the 2500 m LLSD kernel (top) and the peak-to-peak azimuthal shear estimate (bottom) for a 5 km diameter vortex. The center grey line is the median, the box is the interquartile range (IQR), the whiskers are the lesser of 1.5x(IQR) or the data range, and the single dots are outliers.

These values drop with range because of the geometry of the radar beam – circulations are not well sampled at long ranges. For a larger-scale 8 km diameter vortex (Fig. 2) sampled with the 2500 m kernel, the mean LLSD values are within 5% of the true value out to about 150 km. For brevity, results from the 5000 m and 8000 m kernel are not shown. However, these larger

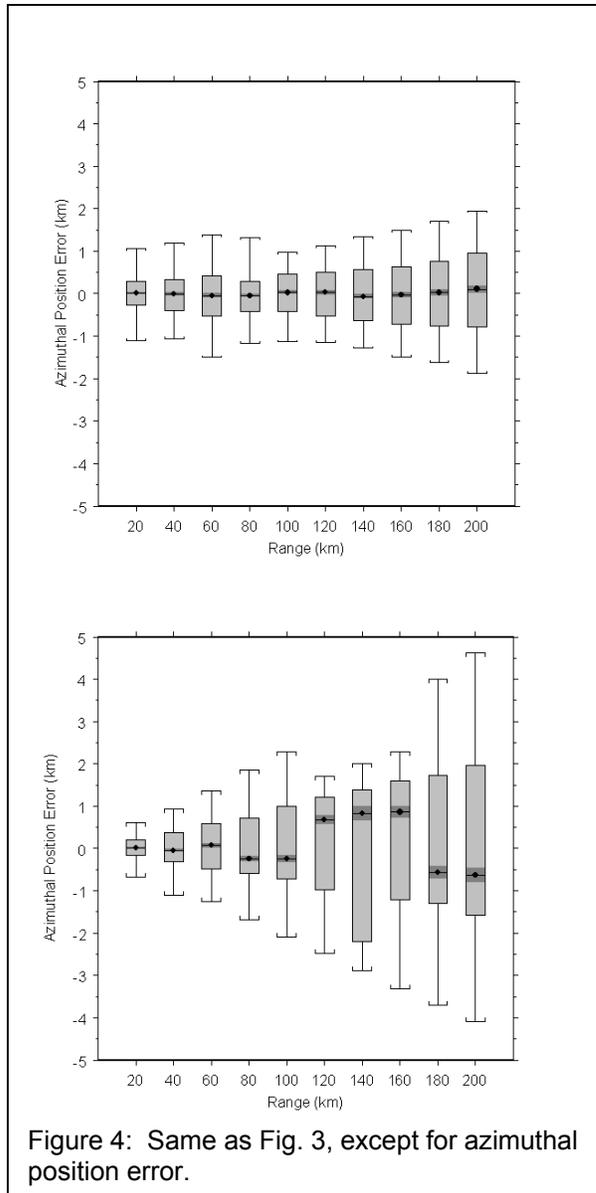
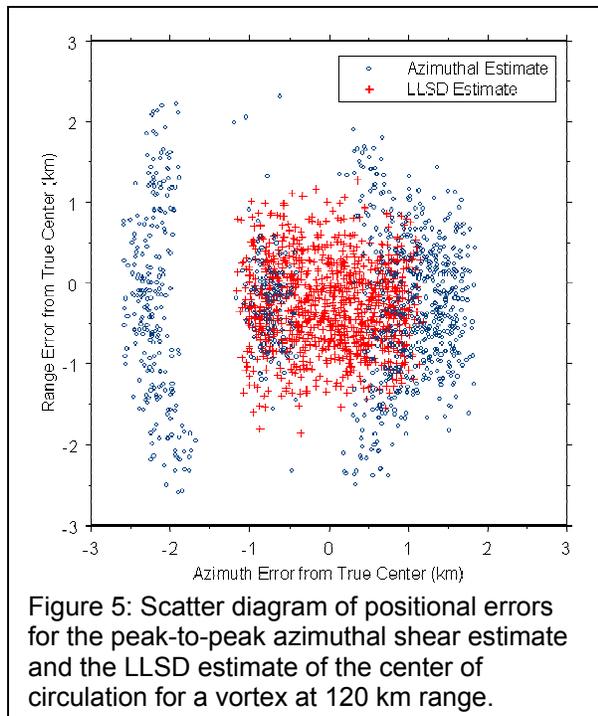


Figure 4: Same as Fig. 3, except for azimuthal position error.

kernels tended to underestimate the strength of smaller vortices compared to the 2500 m kernel.

Because we use synthetic radar data, the true location of the center of the circulation is known. Range and azimuthal position errors were calculated for both the LLSD and peak-to-peak methods. For azimuthal shear, the center of circulation was considered to be halfway between velocity absolute maxima on each side of the circulation. The NSSL Mesocyclone Detection Algorithm (Stumpf et al 1998) uses this method to determine the center of a circulation. For the LLSD rotation, the center of circulation was considered to be at the LLSD rotation maximum.



The errors in range (Fig. 3) for both methods are quite similar, although the variance is smaller for the LLSD estimate. However, the azimuthal distance errors (Fig. 4) for the peak-to-peak method are significantly larger than the LLSD. Additionally, the distribution of the peak-to-peak location estimates is not Gaussian. This is illustrated in Fig. 5. While the LLSD position estimates are clustered around the center of the diagram, there are three distinct groupings for the peak-to-peak data. Because the peak-to-peak method only uses two data points in its calculations, it is highly susceptible to errors caused by the radial offset from the center of the circulation and noise.

4. Conclusion

The local, linear, least squares approach to calculating radial velocity derivatives is a vast improvement over the frequently-used but simplistic and grossly inaccurate method of calculating shear from two data points. The LLSD provides relatively smooth fields that may be used in other applications to identify features such as boundaries and vortices, as well as to accurately assess their strength and position.

5. Acknowledgements

Funding for this research was provided under NOAA-OU Cooperative Agreement

NA17RJ1227, National Science Foundation Grants 9982299 and 0205628, and from the Federal Aviation Administration.

6. References

- Elmore, K.M., E.D. Albo, R.K. Goodrich, and D.J. Peters, 1994: NASA/NCAR airborne and ground-based wind shear studies. Final Report, contract no. NCC1-155, 343 pp.
- Mitchell, E.D., and K.E. Elmore, 1998: A technique for identifying regions of high shear associated with mesocyclones and tornadic vortex signatures. Preprints, 14th Intl. Conf. on Interactive Information and Processing Systems for Meteor., Oceanography, and Hydrology, Phoenix, AZ, Amer. Met. Soc., 312-315.
- Smith, T.M., K.E. Elmore, and S.A. Myers-Dulin, 2003: A Damaging Downburst Prediction and Detection Algorithm for the WSR-88D. *Wea. Forecasting* (in press).
- Stumpf, G.J., A. Witt, E.D. Mitchell, P.L. Spencer, J.T. Johnson, M.D. Eilts, K.W. Thomas, and D.W. Burgess, 1998: The National Severe Storms Laboratory Mesocyclone Detection Algorithm for the WSR-88D. *Wea. Forecasting*, **13**, 304-326.
- Wood, V. T., and R. A. Brown, 1997: Effects of radar sampling on single-Doppler velocity signatures of mesocyclones and tornadoes. *Wea. Forecasting*, **12**, 928-938.
- Wood, V. T., and R. A. Brown, 2000: Oscillations in mesocyclone signatures with range owing to azimuthal radar sampling. *J. Atmos. Oceanic Tech.*, **17**, 90-95.

7.1: Quality Control of WSR-88D Data

V Lakshmanan¹, Kurt Hondl², Gregory Stumpf¹, Travis Smith^{1*}

Abstract

We describe a real-time algorithm for removing non-precipitating echoes from WSR-88D reflectivity data using vertical profiles of reflectivity, Doppler velocity and Spectrum Width. Various attributes, described in this paper, were computed from the radar moments. A few volume scans of radar data were chosen and used to train a feed-forward neural network trained in a supervised manner using resilient-propagation and adaptive weight decay. The trained network was tested on other, independent data cases. The performance of the algorithm was compared to existing methods of performing automated quality control on radar reflectivity data using these independent test cases. We find that the neural-net outperforms the other techniques handily. We also explore the limits on how general the neural network is and make suggestions for quality control in general situations.

1. Introduction

From the point of view of automated applications operating on weather data, echoes in radar reflectivity may be contaminated. These applications require that echoes in the radar reflectivity moment correspond, broadly, to “weather”. By removing ground clutter contamination, estimates of rainfall from the radar data using the National Weather Service (NWS) Weather Surveillance Radar-Doppler 1998 (WSR-88D) can be improved (Fulton et al. 1998; Kessinger et al. 2003). A large number of false positives for the Mesocyclone Detection Algorithm (Stumpf et al. 1995) are caused in regions of clear-air return (McGrath et al. 2002). Poor segmentation and forecasting are achieved by a hierarchical motion estimation technique in regions of ground clutter (Lakshmanan et al. 2003; Lakshmanan 2001). Hence, a completely automated algorithm that can remove regions of ground clutter, anomalous propagation and clear-air returns from the radar reflectivity field would be very useful in improving the performance of other automated weather algorithms.

The problem of examining the radar moments for automated removal of non-precipitating echoes has been

the focus of some research by Steiner and Smith (2002) and Kessinger et al. (2003). The Radar Echo Classifier (REC) described in (Kessinger et al. 2003) has been implemented into the operational Open Radar Product Generator (ORPG). It incorporates the ideas introduced in (Steiner and Smith 2002) and hence, provides a good baseline for comparison. In this paper, the development of a neural network (NN) to do the same task is described. We compare the neural network’s performance on independent cases with the Radar Echo Classifier (REC).

2. The Neural Networks

The final set of features used in the network for which results are reported were: for the lowest scan of velocity, spectrum width and the second lowest scan of reflectivity: local mean, local variance, difference between the data value and the mean, for the lowest scan of reflectivity: local mean, local variance, difference between the data value and the local mean, REC Texture (Kessinger et al. 2003), homogeneity, SPIN (Steiner

¹V Lakshmanan, Gregory Stumpf and Travis Smith are with the Cooperative Institute of Mesoscale Meteorological Studies (CIMMS), University of Oklahoma. ²Kurt Hondl is with the National Severe Storms Laboratory, Norman, OK

and Smith 2002), number of inflections at a 2dBZ threshold, SIGN (Kessinger et al. 2003), echo size. Features related to the vertical profile of reflectivity were the maximum value, weighted average, difference between data values at the two lowest scans, echo top height at a 5dBZ thresholds.

To decorrelate the data value from the mean and median, the difference between the data value and the local mean was used. The weighted average of the reflectivity values is computed over all the elevations where the weight of each data point is given by the height of that pixel above the radar. This takes into account the entire vertical profile instead of just the first two elevations. The homogeneity of the reflectivity field is defined as:

$$hom_{xy} = \frac{\sum_{i \in N_{xy}} \frac{1}{1 + (\frac{I_{xy} - I_i}{I_{xy}})^2}}{card(N_{xy}) - 1} \quad (1)$$

where N_{xy} is the set of valid pixels (I_i) in the neighborhood, N_{xy} , of the pixel at (x, y) in the image, I_{xy} is the pixel value and $card(N_{xy})$ is the number of such neighbors. Echo-size is defined as the fraction of neighbors whose values are within 10dBZ of this pixel's reflectivity value. An inflection point is defined similar to the SPIN (Steiner and Smith 2002) except that the inflection is defined not in a polar neighborhood, but along the entire radial until that point.

We used two separate neural networks – one in regions where Doppler Velocity and Spectrum Width are available, and another where they aren't. All the neural network inputs were scaled such that each feature in the training data exhibited a zero mean and a unit variance when the mean and variance are computed across all patterns.

a. Network Architecture

We used a resilient backpropagation neural network (RPROP) described in Riedmiller and Braun (1993). The RPROP network was trained using supervised batch learning in a multi-layer perceptron (MLP) network. There was one hidden layer. Every input unit was connected to every hidden unit, and every hidden unit to the output unit. In addition, there was a short-circuit connection from the input units directly to the output unit, to capture any linear relationships. Every hidden node

had a “tanh” activation function, chosen because of its signed range. The output unit had a sigmoidal activation function so that the outputs of the networks could be interpreted as posterior probabilities (Bishop 1995). Each non-input node had, associated with it, a bias value which was also part of the training.

The error function that was minimized was a weighted sum of the cross-entropy (which Bishop (1995) suggests is the best measure of error in binary classification problems) and the squared sum of all the weights in the network. The weight decay term improves generalization (Krogh and Hertz 1992). The relative weight, λ , of the two measures is computed within a Bayesian framework (MacKay 1992; Bishop 1995).

The with-velocity network had 22 inputs, 5 hidden nodes and one output while the reflectivity-only network had 16 inputs, 4 hidden nodes and one output.

A validation set ensures a network's generalization, typically through the use of early stopping methods (Bishop 1995). In the neural network literature, a validation set is usually utilized to select the architecture of the neural network. We did not use a separate validation set, mainly because we did not have enough training data in order to do so. Because we lacked a validation set, we did not consider any alternate network topologies. A different network topology may outperform our neural network. Weight decay, rewarding smaller weights, is an alternative way to ensure generalization (Krogh and Hertz 1992; Bishop 1995).

We did use an independent testing set, as described in Section 3.. An extended form of this study, could involve training with the current training set, validating with the current testing set, and then testing on a newly gathered and truthed set of cases. Such a study could make use of early stopping as well.

Just eight volumes of WSR-88D data were selected to encompass different scenarios – strong convection, stratiform rain, ice-coating, low-topped cells, etc. A human interpreter examined these volume scans and drew polygons using the WDSS-II display (Hondl 2002) to select “bad” echo regions. An automated procedure used these human-generated polygons to classify every pixel into the two categories (precipitating and non-precipitating).

Radar Echo Classifier							
Number	nulls	false-alarms	miss	hit	POD	FAR	CSI
1	48554	2573	579	512	0.47	0.83	0.14
2	1187	161	0	0	nan	1	0
3	13440	23	24648	17927	0.42	0	0.42
4	46124	20474	1033	1126	0.52	0.95	0.05
5	10420	14	20798	13828	0.40	0	0.40
6	29731	629	4965	7562	0.60	0.08	0.57
total	149456	23874	52023	40955	0.44	0.37	0.35

Neural Network							
Number	nulls	false-alarms	miss	hit	POD	FAR	CSI
1	50346	781	271	820	0.75	0.49	0.44
2	419	929	0	0	nan	1	0
3	13363	100	5489	37086	0.87	0	0.87
4	25517	41081	7	2152	1	0.95	0.05
5	10259	175	3828	30798	0.89	0	0.88
6	14697	15663	76	12451	0.99	0.56	0.44
total	114601	58729	9671	83307	0.90	0.41	0.55

Table 1: Skill scores when classifying using the Radar Echo Classifier and when using the Neural Network.

3. Results and Conclusions

For testing, a diverse set of volume scans of weather data were chosen and bad echoes marked on these volume scans by a human observer. The volume scans are listed below:

1. KAMA Apr 18, 2002 07:19:10 to 07:23:49 – Significant AP.
2. KFSX Jan 10, 2003 16:12:09 to 16:20:20 – Terrain-induced ground clutter.
3. KTLX May 14, 2003 13:41:08 to 13:45:45 – Strong convection with sharp gradients.
4. KTLX May 19, 2003 08:57:52 to 09:06:02 – AP and spatially smooth clear-air return.
5. KINX May 20, 2003 10:25:52 to 10:30:31 – Strong convection close to the radar.
6. KTLX May 20, 2003 16:39:14 to 16:44:30 – Clear-air return at several elevations.

The volume scans were processed using the trained neural network and using the Radar Echo Classifier (Kessinger et al. 2003). Comparisons were made on a pixel-by-pixel basis of all pixels for which at least one of the elevations had a reflectivity value greater than zero

The neural network identifies regions of precipitation with high skill. It is able to identify bad echoes (AP/GC) when they are similar to cases that it has seen before, but is not able to deduce unfamiliar situations (terrain-

dBZ. Nulls refer to pixels which are not weather-related, and are correctly classified. Hits refer to pixels which are weather-related and are correctly classified. Misses refer to weather-related echoes that are missed, while false-alarms refer to non-weather echoes that are incorrectly classified.

The confusion matrices for each of the volume scans are shown in Table 1 for the two algorithms being compared. The first row of Figure shows a case of significant AP/GC while the third row shows a significant precipitation event. Looking at these images, it is possible to put the quantitative measures (cases 1 and 3) in context. We see that a lot of good data is misclassified by the Radar Echo Classifier. At the same time, the neural network makes its mistakes on lower reflectivity values, but gets higher reflectivity values (whether AP/GC or good data) correct more often. This is a consequence of the cost factors used in the network error equation.

Terrain-induced ground-clutter (Jan. 10, 2003 from KFSX, shown in the second row of Figure) was not part of the training regimen of the neural network, and does pose problems. In mountain regions, terrain heights, or the height of the echo above terrain, could be part of the inputs to the network, instead of, as currently, simply the height above the radar. The network would also have to use texture statistics from the second tilt of the radar, and use vertical differences for the lowest three tilts.

induced GC, vertically continuous clear-air return, chaff, etc.) Even with these limitations, however, the neural network greatly outperforms existing automated techniques. At locations where one or more of these bad-echo forma-

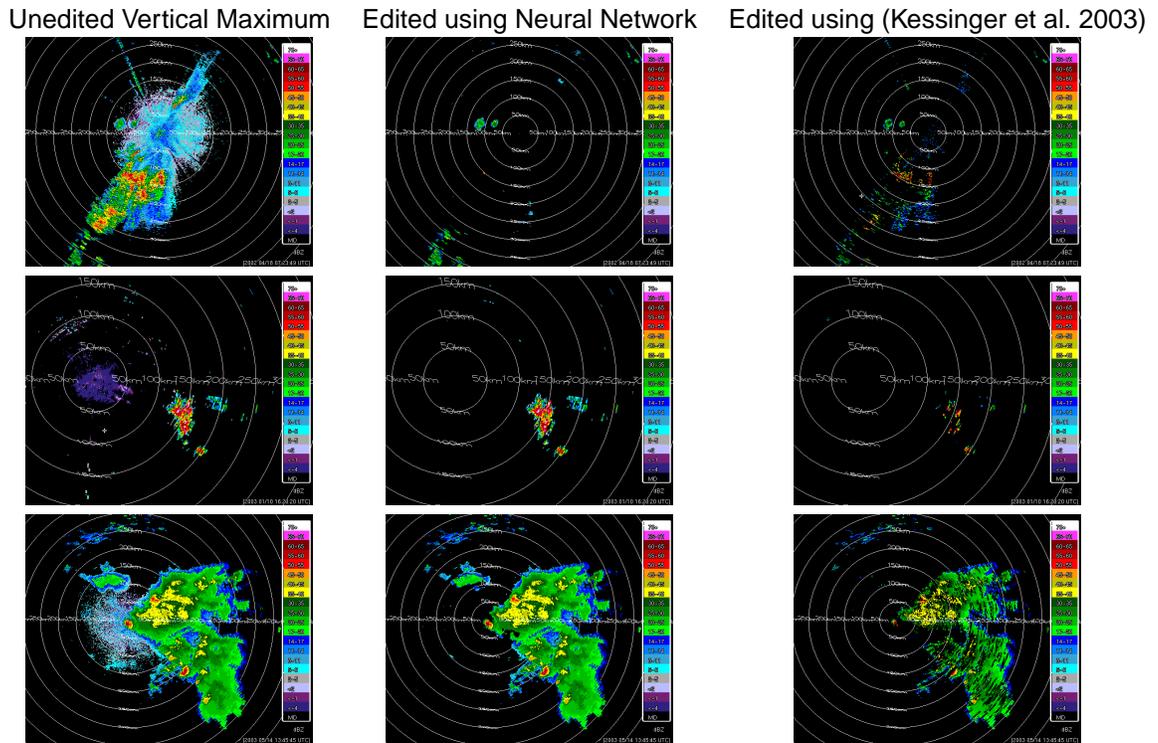


Figure 1: Selected testing cases: KAMA 4/18/2002, KFSX 1/10/2003 and KTLX 5/14/2003. Performance on a data case with significant AP/GC is shown in the first row. The second row illustrates that the Neural Network performs poorly on cases from the Mountain West, where it was not trained, while the third row shows a typical spring precipitation event from the Great Plains.

tions is frequent, the neural network should be trained with those. The network, if trained judiciously, shows high skill in “remembering” and removing bad-echo patterns that it has been trained on.

Acknowledgement Funding for this research was provided under NOAA-OU Cooperative Agreement NA17RJ227, FAA Phased Array Research MOU, and the National Science Foundation Grants 9982299 and 0205628.

References

- Bishop, C., 1995: *Neural Networks for Pattern Recognition*. Oxford.
- Fulton, R., D. Breidenback, D. Miller, and T. O'Bannon, 1998: The WSR-88D rainfall algorithm. *Weather and Forecasting*, **13**, 377–395.
- Hondl, K.: 2002, Current and planned activities for the warning decision support system-integrated information (WDSS-II). *21st Conference on Severe Local Storms*, Amer. Meteor. Soc., San Antonio, TX.
- Kessinger, C., S. Ellis, and J. Van Andel: 2003, The radar echo classifier: A fuzzy logic algorithm for the WSR-88D. *19th IIPS Conference*, Amer. Meteor. Soc., Long Beach, CA.
- Krogh, A. and J. Hertz: 1992, A simple weight decay can improve generalization. *Advances In Neural Information Processing Systems*, S. H. Moody, J. and R. Lippman, eds., Morgan Kaufmann, volume 4, 950–957.
- Lakshmanan, V., 2001: *A Hierarchical, Multiscale Texture Segmentation Algorithm for Real-World Scenes*. Ph.D. thesis, U. Oklahoma, Norman, OK.
- Lakshmanan, V.: 2003, Real-time quality control of reflectivity using satellite infrared channel and surface observations. *Royal Met. Soc. Conf. 2003*, Norwich.
- Lakshmanan, V., R. Rabin, and V. DeBrunner, 2003: Multiscale storm identification and forecast. *J. Atmospheric Research*, ??
- MacKay, D. J. C.: 1992, A practical Bayesian framework for backprop networks. *Advances in Neural Information Processing Systems 4*, J. E. Moody, S. J. Hanson, and R. P. Lippmann, eds., 839–846.
- McGrath, K., T. Jones, and J. Snow: 2002, Increasing the usefulness of a mesocyclone climatology. *21st Conference on Severe Local Storms*, Amer. Meteor. Soc., San Antonio, TX.
- Moller, M., 1993: A scaled conjugate gradient algorithm for fast supervised learning. *Neural Networks*, **6**, 525–533.
- Riedmiller, M. and H. Braun: 1993, A direct adaptive method for faster backpropagation learning: The RPROP algorithm. *Proc. IEEE Conf. on Neural Networks*.
- Steiner, M. and J. Smith, 2002: Use of three-dimensional reflectivity structure for automated detection and removal of non-precipitating echoes in radar data. *J. Atmos. Ocea. Tech.*, 673–686.
- Stumpf, G., C. Marzban, and E. Rasmussen: 1995, The new NSSL mesocyclone detection algorithm: a paradigm shift in the understanding of storm-scale circulation detection. *27th Conference on Radar Meteorology*.
- Verikas, A., K. Malmqvist, A. Lipnickas, M. Bacauskiene, and A. Gelzinis, 1999: Soft combination of neural classifiers. *Pattern Recognition Letters*, **20**, 429–444.

15.14: Motion Estimator Based On Hierarchical Clusters

V Lakshmanan*

Cooperative Institute of Mesoscale Meteorological Studies
University of Oklahoma

Abstract

In this paper, we describe the use of statistically derived hierarchical clusters of weather data to derive movement estimates from pairs of frames in a time sequence. We show that the use of hierarchical clusters enables small cells to be tracked over short periods of time while using the movement of the larger scale features they are embedded in for longer periods.

The motion estimator has been applied both to reflectivity data obtained from the National Weather Service Radar (WSR-88D) and to cloud-top infrared temperatures obtained from the GOES-11 satellite. We demonstrate the results on both these sensors.

1. Short-term forecast methods

The operational way of identifying storms from radar images involves the use of multiple thresholds and counting runs of values above a threshold along a radial. The centroids are then used as proxy for the storms (Johnson et al. 1998) and tracked either on the basis of proximity to expected position or through a linear programming approach (Dixon 1994). Change in position is extrapolated.

A second technique is to use rectangular sub-grids and find the maximum correlation within a search radius (Rinehart and Garvey 1978; Tuttle and Gall 1999). A modification of this technique is to pre-filter the data so as to track only the larger scales (Wolfson et al. 1999; Lakshmanan 2000). It is also possible to use sub-grids ranging in size from that of the entire image to small

16km x 16km grids and to compute motion estimates at each of these scales. Smoothness criteria can be used to constrain these estimates at different scales.

Identifying, matching and extrapolating storm core locations is suitable for small scale storms. The large scale features and cross-correlation technique is suitable for longer forecasts, but with loss of detailed motion estimates. An assumption here is that the storms are of the scale of the sub-grid, not larger. The multiscale estimation is suitable also for large scale forecasts, but with less precise detailed motion estimates.

When used for advection, all the correlation techniques rely on reverse projection, so there needs to be wind speed at the spot where the storm is moving to. The image template methods also assume that all pixels within a grid are moving together.

We use a hybrid approach where motion estimates are made for groups of storms (rather than for sub-grids of the image), but at various scales. The motion estimate for a storm cell is the movement that minimizes the mean-absolute-error between the current frame and corresponding pixels in the previous frame, except that the template is not a sub-grid of the image, but is instead the actual shape of the storm cell.

Instead of simply matching storm cells across frames, motion estimates are made by finding the best match for the storm-template. Thus, the major steps in the technique are:

1. Find storms at different scales.
2. Estimate motion at the various scales.
3. Forecast for different periods using motion at different scales.

*Corresponding author address: lakshman@ou.edu, also affiliated with NOAA/OAR/National Severe Storms Laboratory.

2. Identifying storms

A K-Means clustering technique from Lakshmanan (2001); Lakshmanan et al. (2002) is used to identify components in vector fields. The technique provides nested partitions, i.e. the identified storms structures are strictly hierarchical. The technique works by clustering image values (reflectivity/infrared temperature, etc.) in the neighborhood of a pixel on two opposing criteria:

- Belong to same cluster as your neighbors.
- Belong to cluster whose mean is closest to your value.

Hierarchical segmentation is incorporated into the K-Means clustering technique by steadily relaxing inter-cluster distances.

K of this K-Means clustering is not the number of regions in the final segmented output. It is the number of central vectors about which we do the clustering. The number of regions is determined by the spatial location. As the number K increases, the clusters cover a smaller range in the texture space. In case the number of regions is not known a priori, a very high value of K may be chosen. The most detailed segmentation may have too many regions, but coarser levels might yield the desired result. This is one advantage of using a hierarchical technique.

We iteratively move pixels minimizing

$$E(k) = \lambda d_m(k) + (1 - \lambda) d_c(k) \quad 0 \leq \lambda \leq 1 \quad (1)$$

where the distance in the measurement space is:

$$d_m(k) = \| \mu_k^n - T_{xy} \| \quad (2)$$

and the discontinuity measure is::

$$d_c(k) = \sum_{ij \in N_{xy}} (1 - \delta(S_{ij}^n - k)) \quad (3)$$

A region growing algorithm is employed to build a set of connected regions, where each region consists of 8-connected pixels that belong to the same K-Means cluster. If a connected region is too small, then its cluster mean (the mean of the texture vectors at each pixel in the region) is compared to the cluster means of the adjoining regions and the small region is merged with the

closest mean. The result of the K-Means segmentation, region growing and region merge steps is the most detailed segmentation of the image.

The inter-cluster distances of all adjacent clusters (or regions) in the image are computed. A threshold is set such that half the pairs fall below this threshold. If a pair of clusters differs by less than this threshold, the clusters are merged and cluster means updated. This process is continued until no two adjacent regions are closer in cluster space than the threshold. When this process is complete, we have the next coarser scale of the segmentation. This process is repeated until no changes happen.

3. Motion Estimation

Once the storms have been identified from the images, these storms are used as a template and the movement that minimizes the absolute-error between two frames is computed. For radar images, we used consecutive (5 min) volume scans. For satellite imagery, we used frames 400 seconds apart.

Motion estimation is done by moving a template of the identified cluster at the appropriate scale around in the previous image. A matrix of mean absolute error at the different positions is obtained as shown in Figure 1

The field is minimized by weighting each pixel by how much it differs from the absolute minimum and finding the centroid.

For each storm template, we also get a growth/decay estimate. This is based on how much the average value inside the template changes based on the template at the best match.

4. Short-term Forecast

The forecast of the fields is done based on the motion estimates, growth and decay heuristic and the current data. Forecasts can be made on fields other than the tracked field. For example, motion estimates can be derived from VIL and applied to radar reflectivity and probability fields of lightning and hail.

The forecast is done in three steps:

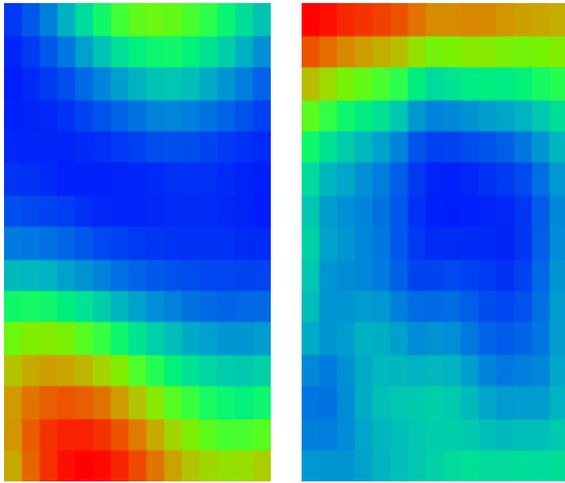


Figure 1: Matrix of mean absolute error by position. Larger errors are “hotter”. Two different locations are shown.

1. Forward: project data forward in time to a spatial location given by the motion estimate at their current location and the elapsed time.
2. Define a background (global) motion estimate given by the mean storm motion.
3. Reverse: obtain data at a spatial point in the future based on the current wind direction at that spot and current spatial distribution of data.

The skill of this technique is quantitatively measured by comparing, for example, the 30 minute forecast against the actual field closest to 30 minutes ahead. Results over a 750 minute period on reflectivity from the Fort Worth radar on April 1, 1995 are shown in Figures 2 and 3.

The CSI seems to indicate the technique performs a lot better than persistence. The MAE, especially in the longer-range forecasts, doesn't show much difference between the two. The reason is that the MAE takes into account actual reflectivity values. We are good at predicting storm location, but not so good at growth/decay.

A forecast based on satellite infrared temperature is shown in Figure 4.

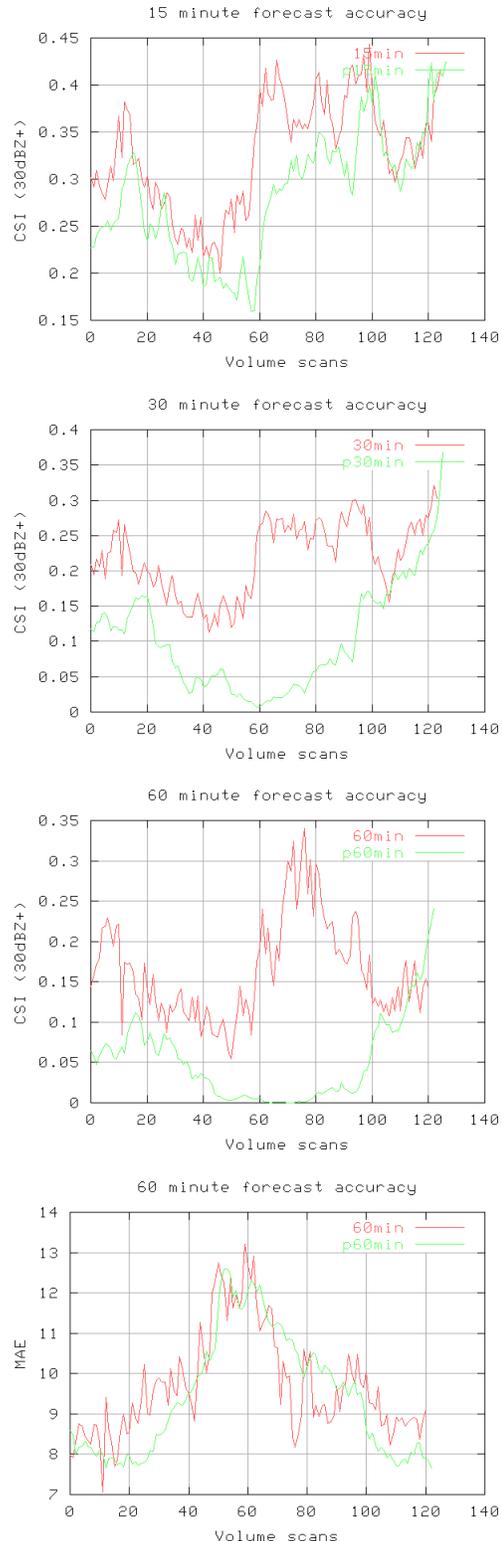


Figure 2: Skill (red) at forecasting a radar reflectivity field compared to a persistence forecast (green). (a) Values 30dBZ and above for 15 minutes (b) Values 30dBZ and above for 30 minutes (c) Values 30dBZ and above for 60 minutes (d) Mean absolute error in 60 minute forecast

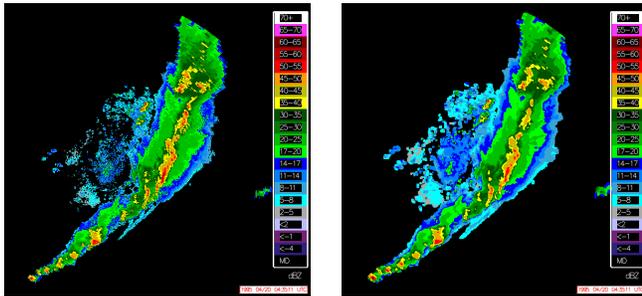


Figure 3: The original (left) and a 15 minute forecast on KFWS reflectivity data from April 1995.

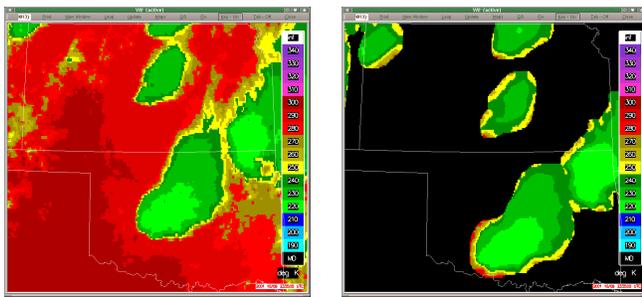


Figure 4: The original (left) and a 30 minute forecast of infrared temperature from Oct. 1999.

Further work is needed in these areas:

1. High bias – associated with splatting during forward projection.
2. Poor forecast of actual data values (high MAE), i.e. poor growth/decay estimate.
3. A better choice of scale for making forecasts.
4. Assimilation of mesoscale model wind speeds.
5. Use of Doppler radar velocity estimates.
6. Images look unrealistic beyond 60 minutes.

5. Acknowledgements

Funding for this research was provided under NOAA-OU Cooperative Agreement NA17RJ1227, FAA Phased Ar-

ray Research MOU, and the National Science Foundation Grants 9982299 and 0205628.

References

- Dixon, M.: 1994, *Automated Storm Identification, Tracking and Forecasting – A Radar-Based Method*. Ph.D. thesis, University of Colorado and National Center for Atmospheric Research.
- Johnson, J., P. Mackeen, A. Witt, E. Mitchell, G. Stumpf, M. Eilts, and K. Thomas: 1998, The storm cell identification and tracking algorithm: An enhanced WSR-88D algorithm. *Weather and Forecasting*, **13**, 263–276.
- Lakshmanan, V.: 2000, Speeding up a large scale filter. *Journal of Oceanographic and Atmospheric Technology*, **17**, 468–473.
- 2001, *A Heirarchical, Multiscale Texture Segmentation Algorithm for Real-World Scenes*. Ph.D. thesis, U. Oklahoma, Norman, OK.
- Lakshmanan, V., V. DeBrunner, and R. Rabin: 2002, Nested partitions using texture segmentation. *Southwest Symposium on Image Analysis and Interpretation*, IEEE, Santa Fe, New Mexico.
- Rinehart, R. and E. Garvey: 1978, Three-dimensional storm motion detection by conventional weather radar. *Nature*, **273**, 287–289.
- Tuttle, J. and R. Gall: 1999, A single-radar technique for estimating the winds in tropical cyclones. *Bull. Amer. Met. Soc.*, **80**.
- Wolfson, M., B. Forman, R. Hallowell, and M. Moore: 1999, The growth and decay storm tracker. *8th Conference on Aviation*, Amer. Meteor. Soc., Dallas, TX, 58–62.

Multiscale Storm Identification and Forecast

V Lakshmanan¹, R. Rabin², V. DeBrunner^{3*}

Abstract

We describe a recently developed hierarchical K-Means clustering method for weather images. that can be employed to identify storms at different scales. We describe an error-minimization technique to identify movement between successive frames of a sequence and show that we can use the K-Means clusters as the minimization kernel. A Kalman filter is used to provide smooth estimates of velocity at a pixel through time. Using this technique in combination with the K-Means clusters, we can identify storm motion at different scales and choose different scales to forecast based on the time scale of interest.

The motion estimator has been applied both to reflectivity data obtained from the National Weather Service Radar (WSR-88D) and to cloud-top infrared temperatures obtained from GOES satellites. We demonstrate results on both these sensors.

1. Introduction

The segmentation of weather imagery is a fundamental problem to automated weather analysis, as has been pointed out in Peak and Tag (1994); Lakshmanan et al. (2000); Johnson et al. (1998). There are numerous pattern recognition algorithms that have been developed on weather images, such as for rainfall estimates (Lai et al. 2000), and cloud classification (Lee et al. 1990), but segmentation techniques for weather images have not been addressed. This is true even of work that attempts to factor out weather-related effects in satellite imagery of land (Markus and Cavalieri 2000; Narasimhan and Nayar 2000).

In the meteorological community, the importance of multiscale segmentation has been often noted (Johnson et al. 1998; Wolfson et al. 1999; Lakshmanan et al. 2000). In (Peak and Tag 1994), the authors detail the difficulties that traditional segmentation algorithms have with satellite weather images because of the textural nature of clouds. As a result, a complex technique consisting of a sequence of fixed thresholds, followed by a neural network that decides how and when to prune or merge the resulting regions

*¹ Corresponding author address: lakshman@ou.edu. V Lakshmanan is at the Cooperative Institute of Mesoscale Meteorological Studies, U. Oklahoma (OU) and is also affiliated with NOAA/OAR/National Severe Storms Laboratory (NSSL). ²Robert Rabin is with NSSL and the Cooperative Institute of Meteorological Satellite Studies (CIMSS) at the University of Wisconsin, Madison. ³Victor DeBrunner is with the Department of Electrical and Computer Engineering at OU.

is proposed (Peak and Tag 1994). We show here that using a hierarchical technique in combination with a texture segmentation algorithm makes segmentation of satellite weather images possible such that even small cloud features can be identified.

The textural nature of weather imagery makes robust segmentation for storm tracking purposes very difficult. For storm tracking to be useful, the identification and tracking algorithm should be completely automated. The identification algorithm should not require training, i.e. the algorithm should not expect to see examples of all the “objects” it must identify. Storm “cells” (small scale features) should be capable of being identified. Because the notion of scale is natural in the storm tracking context, we would like to add the requirement that storms at various scales be identified, with their hierarchical structure intact. A multiscale tracking algorithm would be a significant improvement over current tracking schemes which concentrate either on small scales (e.g: (Johnson et al. 1998)) or on large scales (e.g: (Wolfson et al. 1999)).

In the United States, the operational way to identify storms from radar images involves the use of multiple thresholds and counting runs of values above a threshold along a radial. The centroids are then used as a proxy for the storms (Johnson et al. 1998) and tracked either on the basis of proximity to expected position or through a linear programming approach (Dixon 1994). The change in position is extrapolated.

Identification and tracking algorithms for satellite weather imagery have been implemented for mesoscale convective systems (Morel et al. 1997) – where the features of interest (storm anvils colder than 240K) are on the scale of about $10,000 \text{ km}^2$. Our goal, however, is to identify storm scale features, features on the scale of about 10 km^2 .

Another technique (French et al. 1992) is to use neural networks to model input reflectivity fields as a set of nodes and to forecast reflectivity locations in the future based on the evolution of the nodes required to model successive frames. The problems with this technique are that it requires training a neural neural network in real-time, and that a trained neural network can not be used to forecast fields which have not been tracked.

A third technique is to use rectangular sub-grids and to find the maximum correlation within a search radius (Rinehart and Garvey 1978; Tuttle and Gall 1999). A modification of this technique is to pre-filter the data so as to track only the larger scales (Wolfson et al. 1999; Lakshmanan 2000). It is also possible to use sub-grids ranging in size from that of the entire image to small (say, $16\text{km} \times 16\text{km}$) grids, and to compute motion estimates at each of these scales. Smoothness criteria can be used to constrain these estimates at different scales.

Identifying, matching and extrapolating storm core locations is suitable for small scale storms. The large scale features and cross-correlation technique is suitable for longer forecasts, but with loss of detailed motion estimates. An assumption here is that the storms are of the scale of the sub-grid, not larger. The multiscale estimation is suitable also for large scale forecasts, but with less precise detailed motion estimates.

When used for advection, all the correlation techniques rely on reverse projection, so there needs to be wind speed at the spot where the storm is moving to. Methods rely on correlation estimates of rectangular templates also assume that all pixels within that rectangular template are moving together.

We use a hybrid approach where motion estimates are made for groups of storms (rather than for sub-grids of the image), but at various scales. The motion estimate

for a storm cell is the movement that minimizes the mean-absolute-error between the current frame and corresponding pixels in the previous frame, except that the template is not a rectangular sub-grid of the image, but is instead the actual shape of the storm cell.

Instead of simply matching storm cells across frames, motion estimates are made by finding the best match for the storm-template. Thus, the major steps in the technique are:

1. Find storms at different scales.
2. Estimate motion at the various scales.
3. Forecast for different periods using motion at different scales.

Because the motion estimates are made for storms, it is possible to interpolate between storm boundaries to obtain motion estimates at every part of the domain.

2. Hierarchical texture segmentation to identify storms

A K-Means clustering technique from Lakshmanan (2001); Lakshmanan et al. (2002) is used to identify components in vector fields. The technique provides nested partitions, i.e. the identified storms structures are strictly hierarchical. The technique works by clustering image values (reflectivity/infrared temperature, etc.) in the neighborhood of a pixel on two opposing criteria:

- Belong to same cluster as your neighbors.
- Belong to cluster whose mean is closest to your value.

Hierarchical segmentation is incorporated into the K-Means clustering technique by steadily relaxing inter-cluster distances.

The technique works by iteratively moving pixels between clusters minimizing

$$E(k) = \lambda d_m(k) + (1 - \lambda) d_c(k) \quad 0 \leq \lambda \leq 1 \quad (1)$$

where the distance in the measurement space is:

$$d_m(k) = \| \mu_k^n - T_{xy} \| \quad (2)$$

and the discontinuity measure is::

$$d_c(k) = \sum_{ij \in N_{xy}} (1 - \delta(S_{ij}^n - k)) \quad (3)$$

A region growing algorithm is employed to build a set of connected regions, where each region consists of 8-connected pixels that belong to the same K-Means cluster. If a connected region is too small, then its cluster mean (the mean of the texture vectors at each pixel in the region) is compared to the cluster means of the adjoining regions and the small region is merged with the closest mean. The result of the K-Means

segmentation, region growing and region merge steps is the most detailed segmentation of the image.

The inter-cluster distances of all adjacent clusters (or regions) in the image are computed. A threshold is set such that half the pairs fall below this threshold. If a pair of clusters differs by less than this threshold, the clusters are merged and cluster means updated. This process is continued until no two adjacent regions are closer in cluster space than the threshold. When this process is complete, we have the next coarser scale of the segmentation. This process is repeated until no changes happen.

a. Weather Radar Images

Texture segmentation using Markov Random Field (MRF) models has been utilized to segment synthetic aperture radar (SAR) images, mainly because SAR images are characterized by a lot of speckle (Dong et al. 2001; Schroder et al. 1998; Smits and Dellepine 1999), a problem which is resolved through the use of neighborhood statistics. Another reason for using texture segmentation on SAR imagery is that the same MRF model used for segmentation can also be used for classifying the identified segments (Dong et al. 2001).

Texture segmentation has not been applied to weather radar data before. In weather radar data, especially in cases where there is significant precipitation, the problem of speckle does not arise except in the immediate vicinity of the radar. Hence, traditional texture segmentation provides no significant advantage. In fact, as shown in Figure 1e, even a scalar segmentation approach works quite well. What neither the scalar segmentation approaches, for example (Johnson et al. 1998), nor standard texture segmentation approaches (Blum and Rosenblat 1972; Hofmann et al. 1996; Ma and Manjunath 1997) can provide is a nested partition of identified segments. The watershed segmentation approach of Najman and Schmitt (1996) can provide a nested partition, but does not segment weather data well (See Figure 1f). As shown in Lakshmanan et al. (2002); Lakshmanan (2001), multiscale segmentation can be achieved by agglomerative K-Means clustering of texture vectors and slow relaxation of the allowed inter-cluster distance.

We wish to segment the reflectivity moment of radar elevation scans obtained from a Doppler Weather Service Radar (WSR-88D). The data have been mapped from polar coordinates into a Cartesian grid tangential to the earth's surface at the radar location where each pixel is a square area of one kilometer on each side. The pixel values, in dBZ, range from about $-7dBZ$ to about $64dBZ$, with the reflectivity values for some pixels missing. Missing values and all reflectivity values less than $0dBZ$ were thresholded to be $0dBZ$ before the segmentation process.

The radar elevations scans in this study were collected every 5-6 minutes. The weather surveillance radars used by the National Weather Service scan through thunderstorms starting at a low elevation angle, 0.5° for Volume Coverage Pattern (VCP) 21, and after completing a full 360° azimuthal sweep, progressively increase the elevation angle until an upper limit is reached (19.5° in VCP 21). See Figure 2 (Crum and Alberty 1993; Smith 1995). The data were remapped to a Cartesian plane and were then segmented using the K-Means clustering technique.

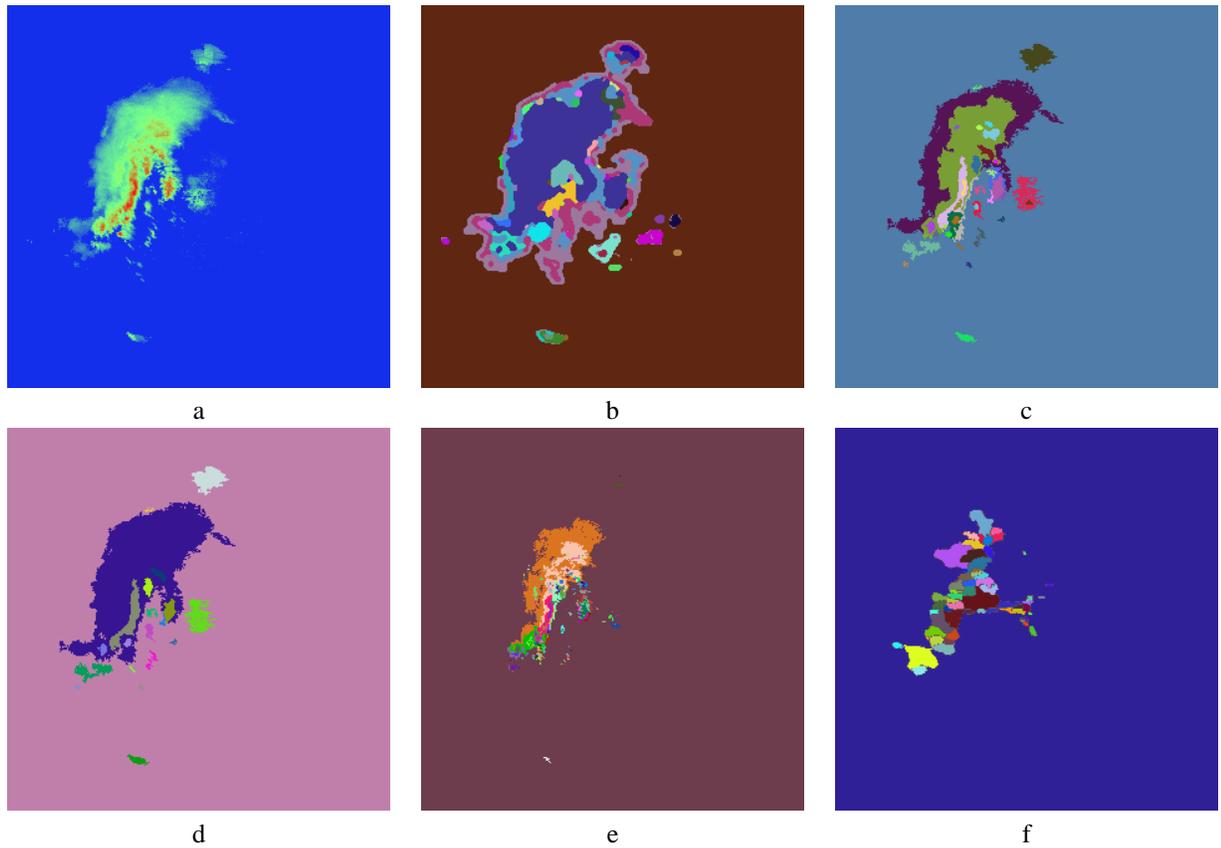


Figure 1: Segmenting a radar reflectivity image. (a) A radar reflectivity image, from Fort Worth May 5, 1995. (b) The result of segmenting the radar reflectivity image using the Markov Random Field (MRF) approach of Blum and Rosenblat (1972). (c) The result of segmenting the image using the method of this paper, tweaked to process the reflectivity range of interest. The most detailed scale is shown. (d) The next higher scale of segmentation using the method of this paper. (e) Simply separating the image into contiguous bands of $10dBZ$. (f) Using the watershed approach of Najman and Schmitt (1996).

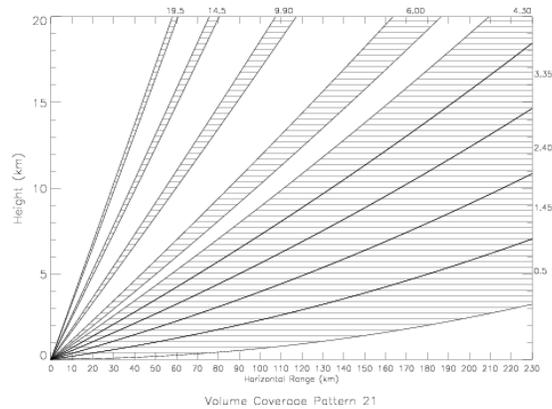


Figure 2: Volume Coverage Pattern (VCP) 21 of the WSR-88D, a weather surveillance radar used by the National Weather Service. The volume coverage is shown. The beamwidth is 0.95 degrees and there are 9 elevation scans in this VCP. Figure from Smith (1995).

b. Satellite Infrared Images

We demonstrate results of segmenting the infrared window channel (11μ) of GOES satellite imager data. The images are 200×300 with each pixel representing a $4\text{km} \times 4\text{km}$. The images are projected onto a plane tangential to the surface of the earth. The satellite data were collected over the continental United States using GOES-11 on March 29, 1998. The pixel values were also mapped from radiance values to equivalent black body temperature in degrees Kelvin before the segmentation. Images of the sequence are available at eight minute intervals.

The sequence of satellite images captures a day of significant thunderstorm activity. Several thunderstorms grow and decay during the day. The temperatures and sizes of the cloud tops in the images show relate in a bulk sense to the magnitude and extent of clusters of storm updrafts within the anvil clouds. Segmentation of this sequence should be able to consistently identify the thunderstorm cloud tops in the images. Ideally, when cloud tops appear to split or merge, the corresponding segmented regions should do the same. A very important requirement is that small changes in the storm structure should be reflected as small changes in the segmented region corresponding to the storm.

Studies (Browning 1979; Bellon and Zawadzki 1994) have shown that a single storm cell grows and decays in under an hour. Therefore, a storm cell can be expected to stay for no more than seven frames of the satellite sequence. However, a line of thunderstorms within which these cells crop up can be expected (Browning 1979) to persist for up to six hours. Also, the cloud top (anvil) may persist after the cell on radar decays. Therefore, the segmentation should lend itself to segmenting regions corresponding to larger scale features while identifying small scale features that are contained within the large scale feature but have shorter life-times.

A single infrared image was segmented using various segmentation methods in the literature. The results are shown in Figure 3c and d. The results of segmentation using the other approaches (Figure 3b,e and f. are poor in terms of the scale of the resulting regions. This is not surprising because the infrared satellite weather imagery has several characteristics that make it hard to segment: very low dynamic range (from about 225K to 240K) for the regions of interest, poor resolution as compared to the scale of the phenomena of interest, and high pixel value variance, even in the absence of edges. It is instructive to compare the poor performance of these algorithms on the satellite image (see Figure 3) with the performance of the same algorithms on radar reflectivity images in Section a.

The poor spatial resolution of the satellite image affects our algorithm also, in the scale of features that we can detect. Although we can detect features as small as 10 pixels in the image, this translates to about 40 km^2 , a mid-size storm cell (although significantly more detailed than what could be obtained using earlier approaches). The pruning threshold of 10 pixels was set in the algorithm so that any statistics collected are somewhat reliable. One possible way to relax this threshold is by creating a pseudo-high resolution form of the original image, thus getting less square kilometers in the 10-pixel threshold. Unfortunately, on satellite weather images, even a pseudo-high resolution technique (Yao 1999) introduces unacceptable smoothing (Lakshmanan 2001), resulting in worse performance. A second possibility, one that we have not yet looked into because of the prohibitive cost for a continuously running system, is to obtain weather satellite data that has higher spatial resolution. A third possibility is to use the multi-channel nature of satellite weather information to form the pixel representation (instead of using a texture vector based on neighborhood statistics).

Instead of using only texture measurements from only the infrared channel, we used texture measurements (mean and variance) computed on four channels corresponding to 3.9, 6.7, 11 and 12 microns (near infrared, water vapor, window and “dirty window” respectively (Menzel and Purdom 1994)). Since every pixel of the segmented output actually corresponds to four relatively independent measurements (rather than just one), the minimum pruning size in the algorithm can be reduced from about 10 pixels to about 3. The result of using multi-channel information and a lower size threshold is shown in Figure 5 where it is compared to the segmented result if only the 11 micron image had been used.

Notice that the result of segmenting using all four channels (Figure 5f) has smaller regions than the result that uses only the infrared window channel. It is not clear, however, how significant these smaller features are in the context of thunderstorms.

3. Motion Estimation

Once the storms have been identified from the images, these storms are used as a template and the movement that minimizes the absolute-error between two frames is computed. For radar images, we used consecutive (5-6 min) volume scans.

Motion estimation is done by moving a template of the identified region at the appropriate scale around in the previous image. A matrix of mean absolute error at the different positions is obtained as shown in Figure 6

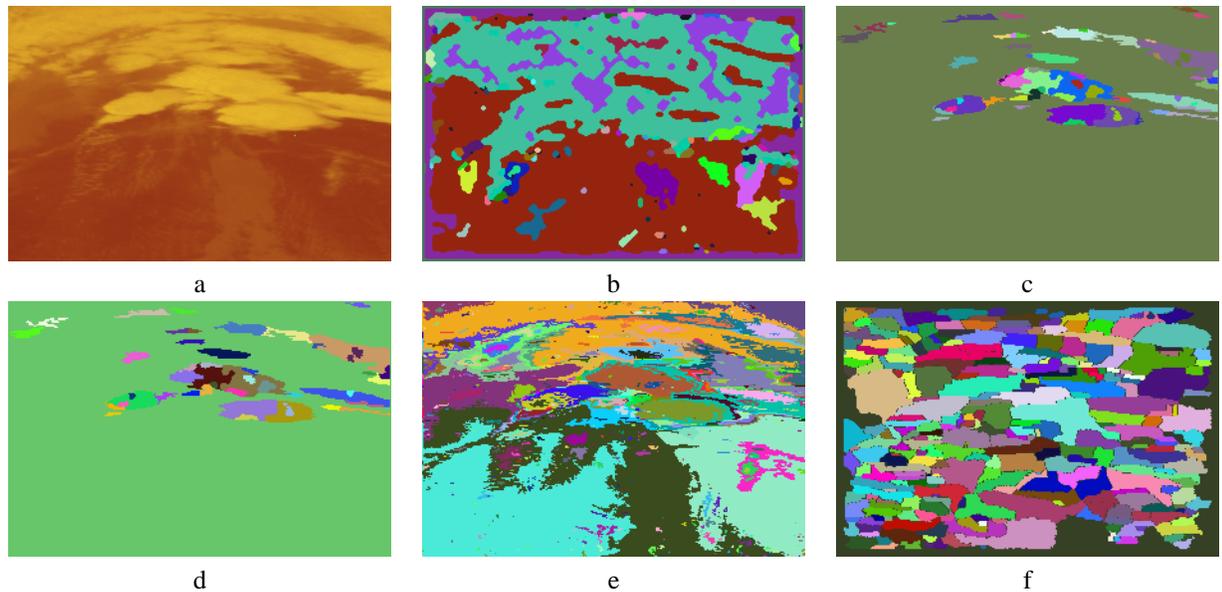


Figure 3: Segmenting an infrared satellite weather image. (a) The infrared image being segmented. Notice the various storms at the top of the image. The darker areas in the bottom correspond to ground. (b) The result of segmenting the image using the Markov Random Field (MRF) approach of Blum and Rosenblat (1972). There is no detail – it is effectively a binary segmentation. (c) The result of segmenting the image using the method of this paper (the most detailed scale). Notice the fine detail within the clouds. (d) The next higher scale of segmentation using the method of this paper. The strong storm cells being significantly colder are retained – the large cloud masses are merged. (e) Simply separating the image into contiguous bands of 1 Kelvin . There is a lot of detail, but no organization. This is what you get using hierarchical thresholds. (f) Using the watershed segmentation approach of Najman and Schmitt (1996). Because of the textural nature of the data, the watershed algorithm has very poor performance.

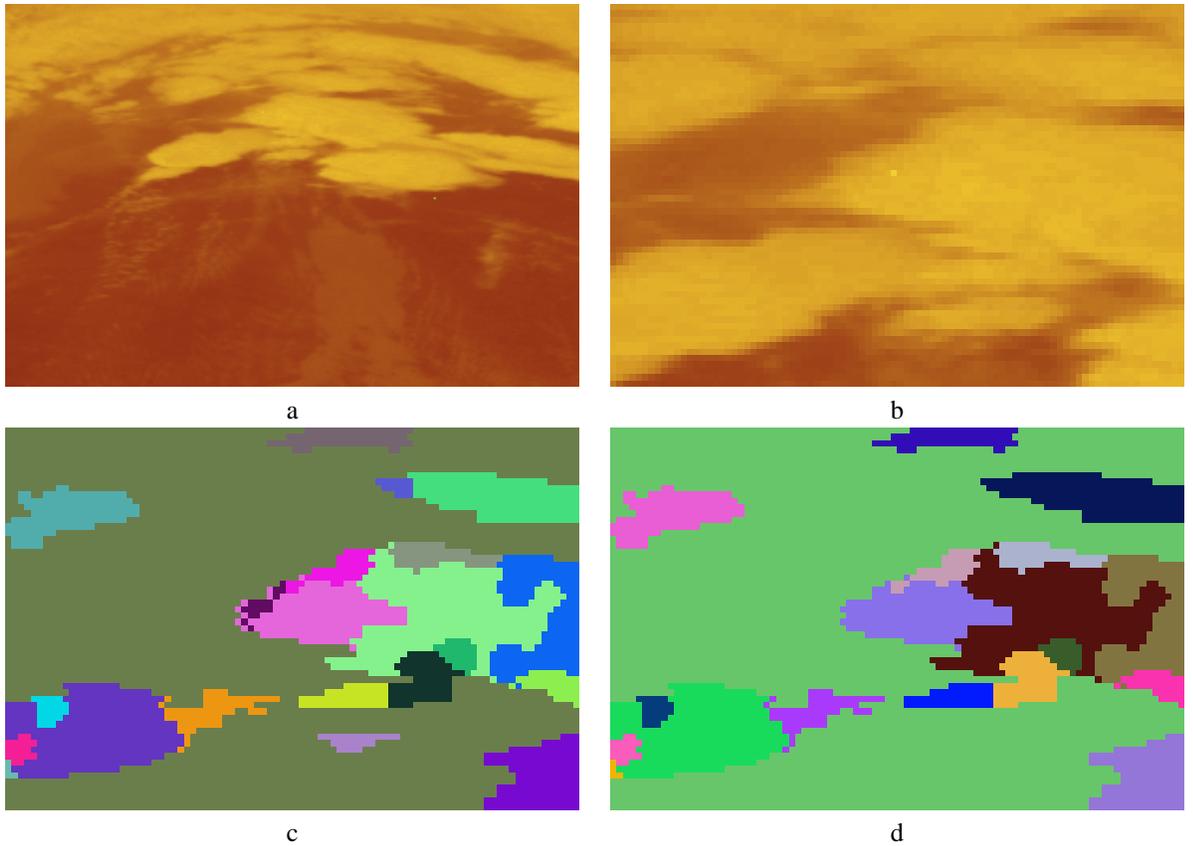


Figure 4: A close-in look at the results shown in Figure 3.. (a) The infrared image being segmented (same as Figure 3a). (b) A close-in look at the input satellite infrared image of (a). (c) A close-in look at the result of segmenting the image using the method of this paper (the most detailed scale). Notice the fine detail within the clouds. (d) A close-in look at the next higher scale of segmentation using the method of this paper.

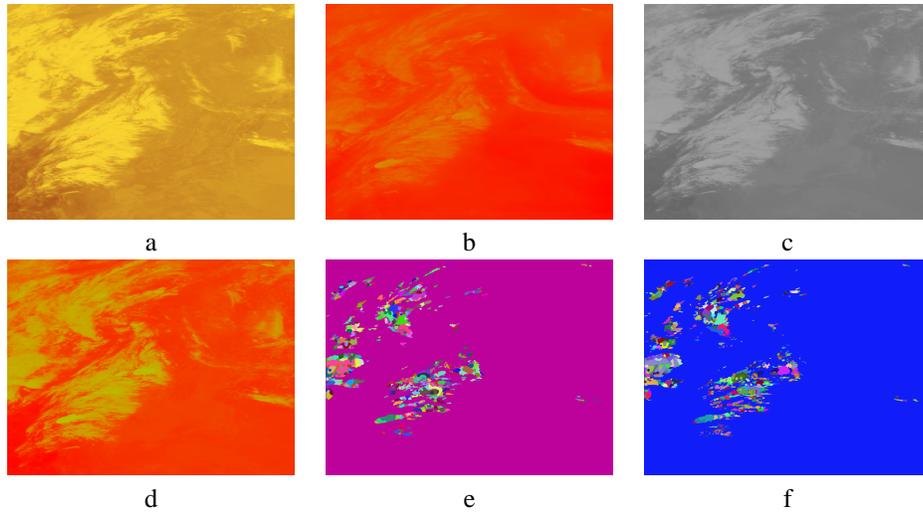


Figure 5: Using multi-channel satellite data for segmentation. (a) 3.9 micron infrared (b) 6.7 micron water vapor (c) 11 micron window (d) 12 micron dirty window channels of data. (e) Most detailed segmentation using only the 11 micron image (f) Most detailed segmentation using all four channels. The segmentation is more detailed than the segmentation that was achieved in (e), but whether these extra details are useful is yet to be determined.

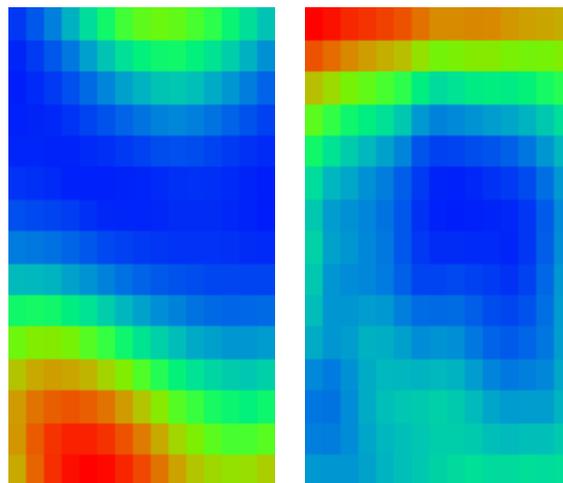


Figure 6: Matrix of mean absolute error by position. Larger errors are “warmer”. Two different locations are shown.

Instead of simply finding the absolute minimum, a smoother minimum of the absolute error field is sought. The field of absolute errors is minimized by weighting each value by how much it differs from the absolute minimum and finding the centroid.

For each storm template, we also get a growth/decay estimate. This is based on how much the average value inside the template changes based on the template at the best match.

Given the motion estimates for each of the regions in the image, the motion estimate at each pixel is determined through interpolation. At the pixel xy , the motion estimate u_{xy} is given by

$$u_{xy} = \frac{\sum_i u_i w_{ixy}}{\sum_i w_{ixy}} \quad (4)$$

where u_i is the motion estimate for the i^{th} region and the weight of this estimate at the point xy is given by:

$$w_{ixy} = \frac{N_i}{\|xy - c_i\|^2} \quad (5)$$

N_i is the number of pixels in the i^{th} region, c_i its centroid and $\|$ denotes the Euclidean distance between the two points.

This motion estimate is for the pair of frames that were used in the comparison. We do temporal smoothing of these estimates by running a Kalman filter Kalman (1960) at each pixel of the motion estimate. The Kalman estimator is built around a constant acceleration model with the standard Kalman update equations Brown and Hwang (1997).

a. Short-term Forecast

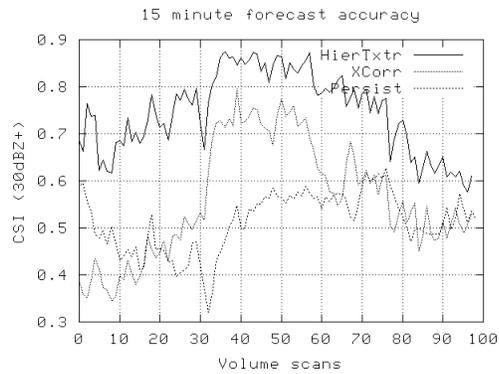
The forecast of the fields is done based on the motion estimates, growth and decay heuristic and the current data. Forecasts can be made on fields other than the tracked field. For example, motion estimates can be derived from VIL and applied to radar reflectivity and probability fields of lightning and hail.

The forecast is done by first project data forward in time to a spatial location given by the motion estimate at their current location and the elapsed time. Locations not filled by this forward projection are filled by interpolating using an inverse square-distance metric of nearby filled locations.

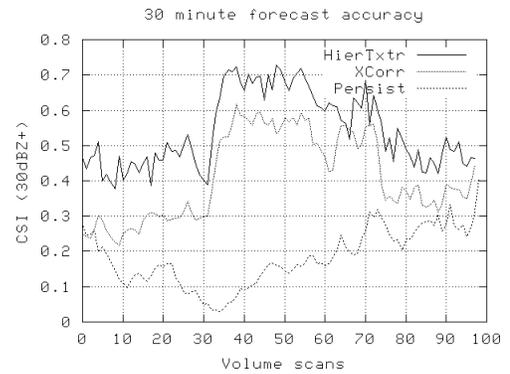
The skill of this technique is quantitatively measured by comparing, for example, the 30 minute forecast against the actual field closest to 30 minutes ahead. For the mean absolute error results, the actual values are used. For the critical success index (CSI) results, the best match with a 5x5 window is used. Comparisons are made with a plain persistence, and with motion estimates derived by minimizing the correlation of a 5x5 template between the frames.

Results over a 60minute period on reflectivity from the Fort Worth radar on April 201995 are shown in Figures 7 and 8.

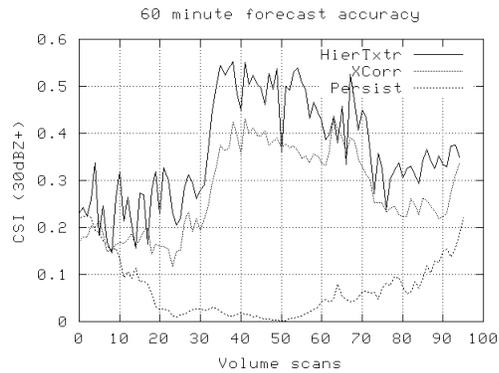
The CSI and MAE measure different aspects of the forecast accuracy. The MAE takes into account actual reflectivity values and is, therefore, a measure of how good the growth-and-decay aspect is. The CSI is a measure of predicting storm location. We are good at predicting storm location, but not so good at growth/decay.



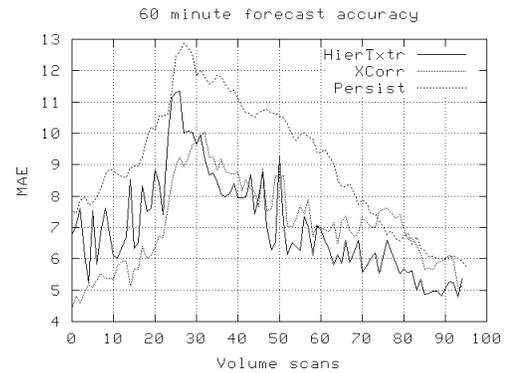
a



b



c



d

Figure 7: Skill at forecasting a radar reflectivity field compared to a persistence forecast and to a local correlation approach. (a) Values 30dBZ and above for 15 minutes (b) Values 30dBZ and above for 30 minutes (c) Values 30dBZ and above for 60 minutes (d) Mean absolute error in 60 minute forecast

A forecast based on satellite infrared temperature is shown in Figure 9. The data are taken from GOES-12 imagery on Oct. 9, 2001. The data provided were 100 seconds apart; we used every 4th frame of the sequence to compute motion estimates. Work to compute skill scores on satellite data is underway.

A version of this paper with color illustrations is available online at

http://www.cimms.ou.edu/~lakshman/Papers/kmeans_motion.pdf

4. Conclusions

It is possible to use a K-Means clustering to provide hierarchical identification of storms. The clusters can then be used to estimate the movement of the storm cores. A forecast that projects the movement of the storm cores linearly possesses some skill.

5. Acknowledgements

Funding for this research was provided under NOAA-OU Cooperative Agreement NA17RJ1227, FAA Phased Array Research MOU, and the National Science Foundation Grants 9982299 and 0205628.

References

- Bellon, A. and I. Zawadzki: 1994, Forecasting of hourly accumulations of precipitation by optimal extrapolation of radar maps. *Journal of Hydrology*, **157**, 211–233.
- Blum, J. and J. Rosenblat: 1972, *Probability and Statistics*. W.B. Saunders Company, 549 pp.
- Brown, R. and P. Hwang: 1997, *Introduction to Random Signals and Applied Kalman Filtering*. John Wiley and Sons.
- Browning, K.: 1979, The FRONTIERS plan: A strategy for using radar and satellite imagery for very-short-range precipitation forecasting. *The Meteorological Magazine*, **108**, 161–184.
- Crum, T. and R. Alberty: 1993, The WSR-88D and the WSR-88D operational support facility. *Bulletin of the American Meteorological Society*, **74**, 1669–1687.
- Dixon, M.: 1994, *Automated Storm Identification, Tracking and Forecasting – A Radar-Based Method*. Ph.D. thesis, University of Colorado and National Center for Atmospheric Research.
- Dong, Y., A. Milne, and B. Forster: 2001, Segmentation and classification of vegetated areas using polarimetric sar image data. *IEEE Trans. Geosci. Remote Sensing*, **39**, 321–329.

- French, M., W. Krajewski, and R. Cuykendall: 1992, Rainfall forecasting in space and time using a neural network. *J. Hydrology*, **137**, 1–31.
- Hofmann, T., J. Puzicha, and J. Buhmann: 1996, A deterministic annealing framework for unsupervised texture segmentation. Technical Report IAI-TR-96-2, Institut für Informatik III, U. Bonn, <http://www-dbv.cs.uni-bonn.de/image/example4.html>.
- Johnson, J., P. Mackeen, A. Witt, E. Mitchell, G. Stumpf, M. Eilts, and K. Thomas: 1998, The storm cell identification and tracking algorithm: An enhanced WSR-88D algorithm. *Weather and Forecasting*, **13**, 263–276.
- Kalman, R.: 1960, A new approach to linear filtering and prediction problems. *Trans. ASME – J. Basic Engr.*, 35–45.
- Lai, E., P. Li, C. Chan, M. Chu, and W. Wong: 2000, Pattern recognition of radar echoes for short-range rainfall forecast. *15th Intl. Conf. on Pattern Recog.*, IEEE, volume 4, 299–302.
- Lakshmanan, V.: 2000, Speeding up a large scale filter. *Journal of Oceanographic and Atmospheric Technology*, **17**, 468–473.
- 2001, *A Heirarchical, Multiscale Texture Segmentation Algorithm for Real-World Scenes*. Ph.D. thesis, U. Oklahoma, Norman, OK.
- Lakshmanan, V., V. DeBrunner, and R. Rabin: 2002, Nested partitions using texture segmentation. *Southwest Symposium on Image Analysis and Interpretation*, IEEE, Santa Fe, New Mexico.
- Lakshmanan, V., R. Rabin, and V. DeBrunner: 2000, Identifying and tracking storms in satellite images. *Second Artificial Intelligence Conference*, American Meteorological Society, Long Beach, CA, 90–95.
- Lee, J., R. Weger, S. Sengupta, and R. Welch: 1990, A neural network approach to cloud classification. *IEEE Trans. on Geoscience and Remote Sensing*, **28**, 846–855.
- Ma, W. and B. Manjunath: 1997, Edge flow: a framework of boundary detection and image segmentation. *Proc. IEEE International Conference on Computer Vision and Pattern Recognition*, San Juan, Puerto Rico, 744–749.
- Markus, T. and D. Cavalieri: 2000, An enhancement of the NASA team sea ice algorithm. *IEEE Trans. on Geoscience and Remote Sensing*, **38**, 1387–1398.
- Menzel, P. and J. Purdom: 1994, The first of a new generation of geostationary operational environmental satellites. *Bull. Amer. Meteo. Soc.*, **75**, 757–782.
- Morel, C., F. Orain, and S. Senesi: 1997, Automated detection and characterization of MCS using the meteosat infrared channel. *Proc. Meteo. Satellite Data Users Conf.*, Brussels, 213–220.
- Najman, L. and M. Schmitt: 1996, Geodesic saliency of watershed contours and hierarchical segmentation. *IEEE Trans. Patt. Anal. and Mach. Intell.*, **18**, 1163–1173.

- Narasimhan, S. and S. Nayar: 2000, Chromatic framework for vision in bad weather. *IEEE Conf. on Comp. Vision and Patt. Recog.*, IEEE, volume 1, 598–605.
- Peak, J. and P. Tag: 1994, Segmentation of satellite weather imagery using hierarchical thresholding and neural networks. *Journal of Applied Meteorology*, **33**, 605–616.
- Rinehart, R. and E. Garvey: 1978, Three-dimensional storm motion detection by conventional weather radar. *Nature*, **273**, 287–289.
- Schroder, M., H. Rehrauer, K. Seidel, and M. Dutcu: 1998, Spatial information retrieval from remote-sensing images – part ii: Gibbs-markov random fields. *IEEE Trans. Geosci. Remote Sensing*, **36**, 1446–1455.
- Smith, T.: 1995, Visualization of WSR-88D data in 3d using application visualization software. *14th Conf. on Weather Forecasting*.
- Smits, P. and S. Dellepine: 1999, Discontinuity-adaptive markov random field model for the segmetnation of intensity sar images. *IEEE Trans. Geosci. Remote Sensing*, **37**, 627–631.
- Tuttle, J. and R. Gall: 1999, A single-radar technique for estimating the winds in tropical cyclones. *Bull. Amer. Met. Soc.*, **80**.
- Wolfson, M., B. Forman, R. Hallowell, and M. Moore: 1999, The growth and decay storm tracker. *8th Conference on Aviation*, Amer. Meteor. Soc., Dallas, TX, 58–62.
- Yao, M.: 1999, *Model-Based Methods for Image Interpolation and Enhancement That Retain Edge Information*. Ph.D. thesis, U. Oklahoma.

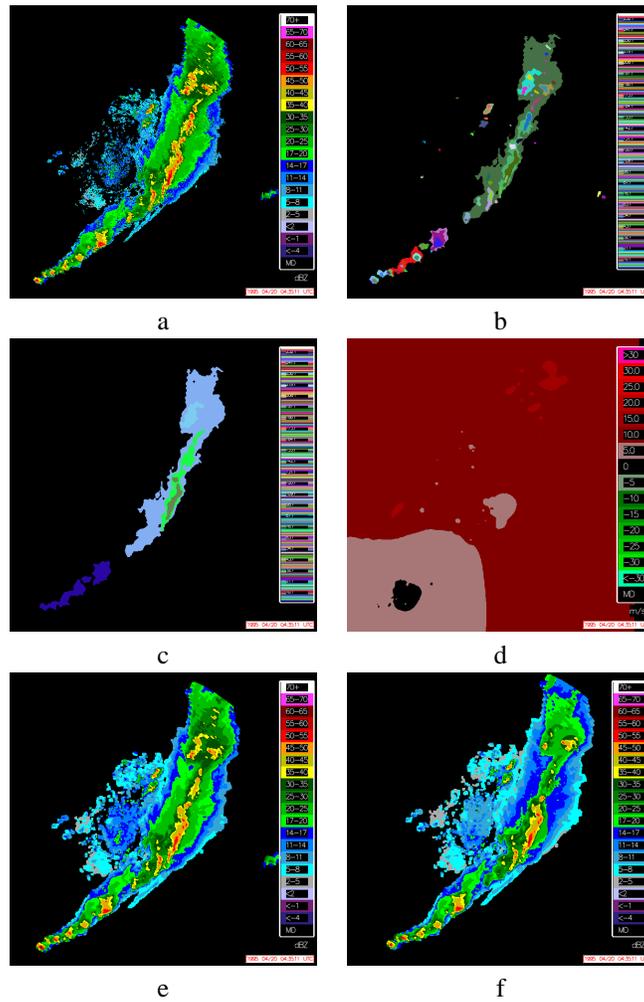


Figure 8: (a) Reflectivity data from KFWS, April 1995. (b) Most detailed scale of segmentation, used in forecasting under 30 minutes. (c) Coarse segmentation, used in forecasting more than 90 minutes. (d) Motion estimate (red is eastward motion) (e) 15min forecast (f) 60min forecast

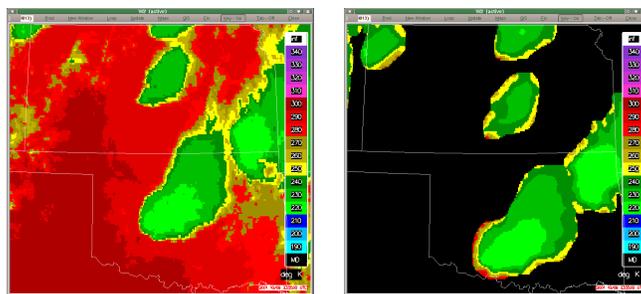


Figure 9: The original (left) and a 30 minute forecast of infrared temperature from Oct. 9, 2001.

14.9: Real-time Merging of Multisource Data

V Lakshmanan*

National Severe Storms Laboratory & University of Oklahoma

Abstract

We describe an extension of the virtual volume concept to multiple sensors. Data from multiple sensors are combined in real-time and mapped into a constantly updating three-dimensional grid. The data are combined in a time-centric manner, with data replacing older data, regardless of the sensor that the older data came from. We discuss scaling problems with this method and how they can be resolved.

We demonstrate this method of merging real-time data on base data (such as radar reflectivity) as well as on derived data (such as linear least square derivatives).

1. Method

a. Time-ordering

Lynn and Lakshmanan (2002) described the concept of virtual volumes of radar data, where the volume of radar data is defined by the latest elevation scans at all times. Such a definition is valid for a single-radar product, as described in (Lynn and Lakshmanan 2002). Is it possible to extend the virtual volume idea to multiple sensors?

Another way to think of the virtual volume is as a time-ordered list of elevation scans. The traditional volume scan is an elevation-ordered list, with the elevations arranged from the the lowest tilt to the current radar scan. If we define a radar volume as a time-ordered list that contains the entire angular space of elevation scans, then the virtual volume results.

With this time-based definition, it is possible to define a multisensor merged grid where each of the parts is updated with the most current sensor input.

b. Mergers and mosaics

Radar reflectivity data, with volumes in the traditional sense have been merged successfully by Zhang et al. (2000). The merging scheme, referred to as “mosaicking”, consists of obtaining volumetric radar data periodically and using Cressman interpolation to create a three-dimensional data set. The data are quality controlled using factors such as the radar range.

In this paper, we propose a different approach to merging data from multiple sensors. Instead of using data that is essentially a snapshot, the entire volume is updated with elevation scan inputs from each radar.

c. Technique

The technique is to connect to multiple data streams and with the arrival of an elevation scan to update the output volume with data from that elevation scan. The WDSS-II Application Programming Interface (Lakshmanan 2002) supports the concept of a listener (or Observer (Gamma et al. 1994)) attached to multiple streams and reacting to the input.

To handle the problem of not receiving data from one or more of the sensors, every grid cell updated is marked with the sensor that it was updated by. When data from a sensor expire, the grid cells that were updated with that data are reset. Naturally, this also handles non-update of grid cells due to factors such as changes in the volume coverage pattern.

Updating a large three dimensional volume with data that arrive every 30 seconds from each radar (in a scenario where we are merging data from upto 4 radars) could be computationally intensive, but several optimizations outside the scope of this paper are available. The merger process was tested on a Linux Pentium-III desktop; it easily keeps up with a real-time feed from four

*Corresponding author address: lakshman@nssl.noaa.gov

weather service radar while using less than 100MB of memory.

d. Outputs

We will demonstrate the technique using data from three weather service radars (KTLX, KSRX and KINX in Oklahoma City, Hot Springs and Tulsa respectively) on May 20, 2001 around 21:45 UTC. Although it is weather service radar that was used to generate the outputs shown here, we have also successfully used a combination of weather service radar and Terminal Doppler Weather Radar (TDWR). In Figure 1, the lowest tilt from each of the three input radars is shown.

The three-dimensional grid was output level by level into separate NetCDF (Jenter and Signell 1992) files and visualized with WDSS-II. In Figure 2a and 2b, the merged data at 2000m and 3000m respectively are shown. Note that the pattern of elevation update is clearly visible in these images. If we use Barnes filtering and confidence-weighting in combination with the purely time-based update used here, we envision higher quality of the resulting data.

The default behaviour is to output the grid with every update. This scales well to two radars, where on the average, a new output is obtained every 15 seconds, but the increased resolution is confusing beyond that. Therefore, the merger process provides an option to write out a new grid only when the time since the last update is greater than, say, 30 seconds.

In addition to the layers of the grid, the merger process puts out a volume product to enable easy navigation up and down the volume and to permit examination by flying through the volume and drawing cross-sections. (See Figure 3).

Finally, the output can optionally include derived fields such as the vertical maximum (composite – See Figure 4) and the vertical average.

e. Blending

The example outputs in the above section were generated by simply updating each grid cell with the latest data from any radar that falls in the spatial volume of that grid cell. Spatial volumes within the atmosphere may be sensed by different tilts of the same radar (considering beam width of the radar beams) and by different radars.

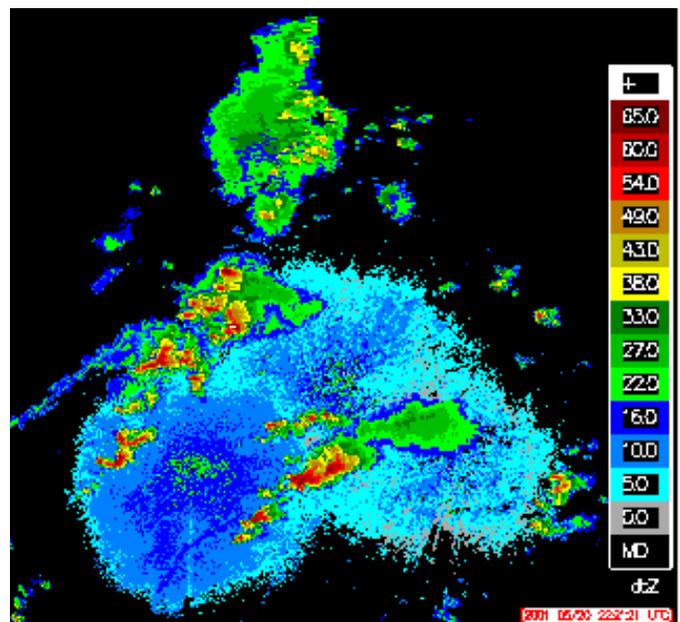
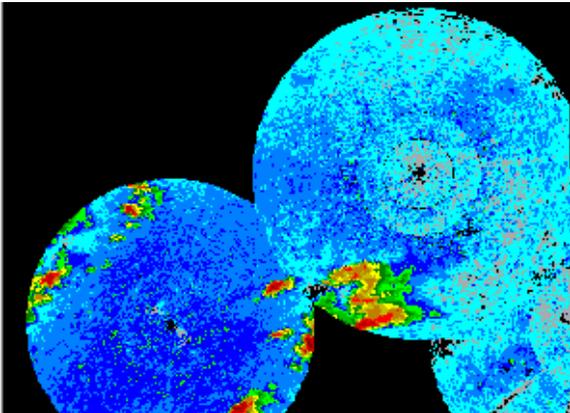
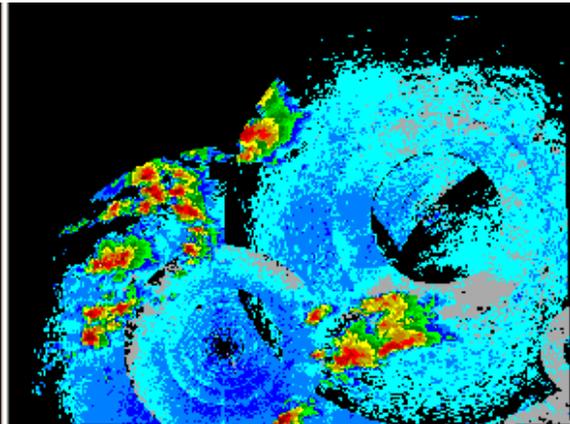


Figure 1: The lowest tilts from the three radars used to demonstrate the results shown here. Data were collected on May 20, 2001 from Oklahoma City, Hot Springs and Tulsa.

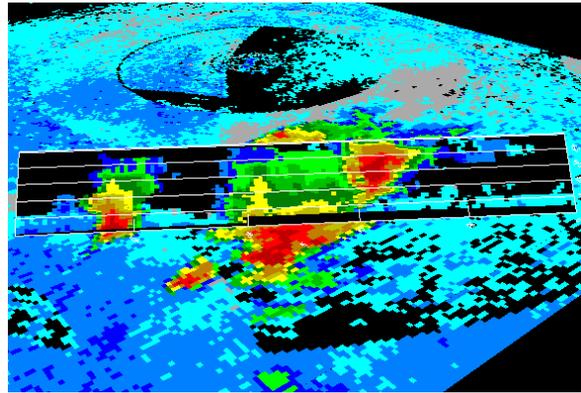


a

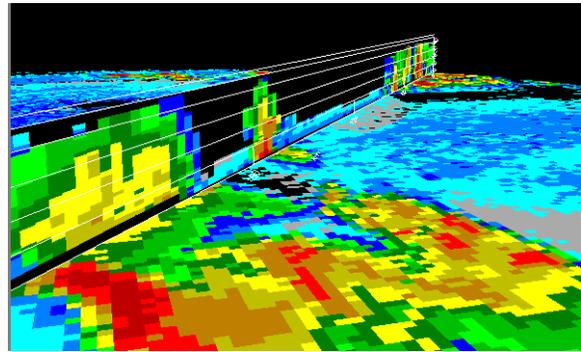


b

Figure 2: The merged 3D grid at 2000m (a, top) and 3000m (b). Notice that the pattern of elevation-by-elevation update is clearly evident. With interpolation and filtering, the quality of the resulting grid can be improved.



a



b

Figure 3: Ways of visualizing the resulting 3D grid: by flying through and by drawing vertical cross-sections.

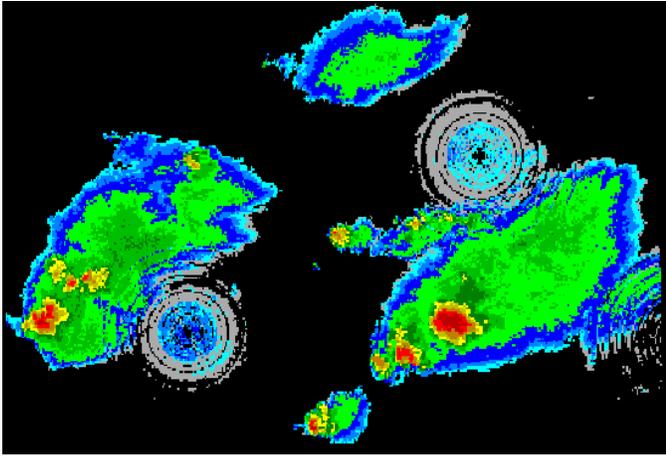


Figure 4: The maximum in the vertical direction at every grid point is an optional product that may be obtained from the 3D grid. Note that there is a new composite product formed every time the 3D grid is updated. This is an example of a multisensor time-update (or virtual volume) algorithm.

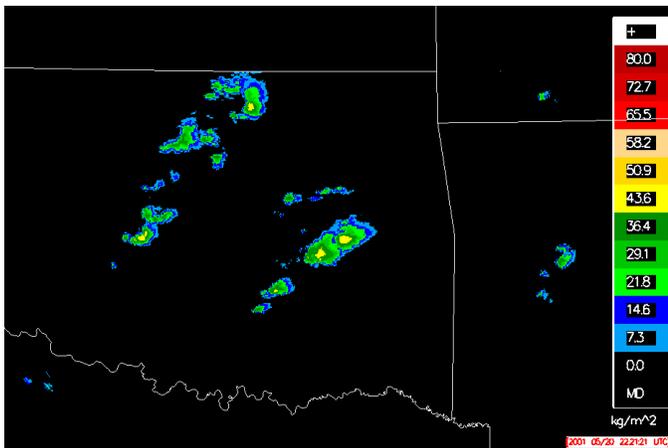


Figure 5: If the data are radar reflectivity, as in these examples, a vertical integrated liquid (VIL) calculation can be computed on the 3D grid. Note that there is a new VIL product formed every time the 3D grid is updated. This is an example of a multisensor time-update (or virtual volume) algorithm.

Thus, when combining data from multiple sensors, it is possible to “blend” the data, rather than simply taking the latest elevation scan. Several strategies have been employed:

1. Taking the maximum reflectivity from any of the radars. This assumes that the radars are uniformly calibrated, but future extensions can correct for calibration errors. An example of this is shown in Figure 6.
2. Blending the data by assigning a weight inversely proportional to the distance from the radar. An example of this is shown in Figure 8.
3. Blending the data by assigning a weight inversely proportional to both the distance from the radar and its age, such that newer data are assigned higher weights. This is shown in Figure 7.
4. Blending the data by spreading the beamwidth, so that the radar data from successive elevation scans are linearly interpolated into the space between them. Although this does not show much difference in Figure 9, it is an effective strategy for single radar volumes.

2. Conclusion

We described a new way of merging data from multiple sensors, by constantly updating a three-dimensional grid of data with data from the sensors. This allows a more current view than traditional merging methods. Future plans include doing quality control of the data and implementing non-uniform and nested grids.

Although demonstrated here on radar reflectivity, the algorithm has been used to merge derived fields such as the local linear least square derivative field (Smith 2002) as well.

Acknowledgement The work detailed in this paper was supported with funds from the National Science Foundation, the Federal Aviation Authority and the National Severe Storms Laboratory (NSSL). The author would like to thank Kurt Hondl, Jason Lynn and Travis Smith for contributions to the ideas underlying this algorithm.

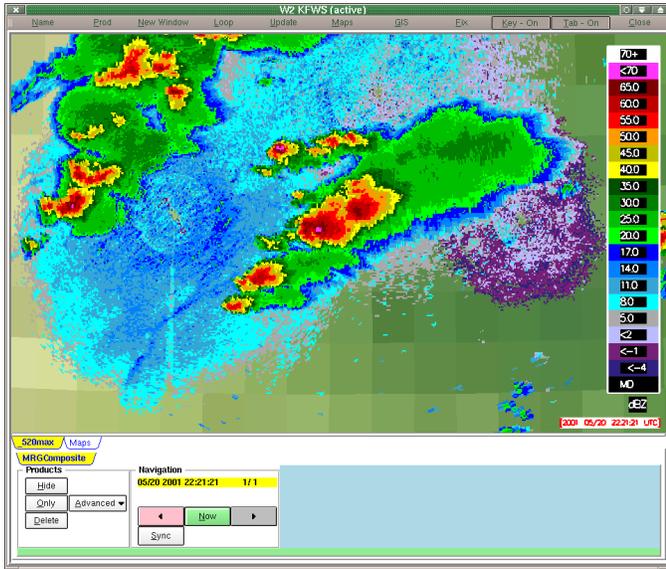


Figure 6: Blending data from two radar by choosing the maximum value from either radar. This strategy is useful for getting around beam-blockage, radar horizon and cone-of-silence issues.

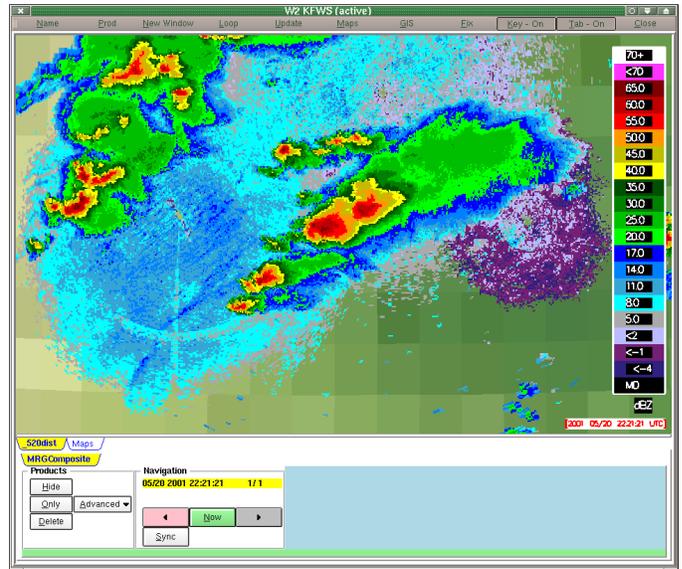


Figure 8: Blending data from two radar by assigning them weights inversely proportional to the distance from the radar.

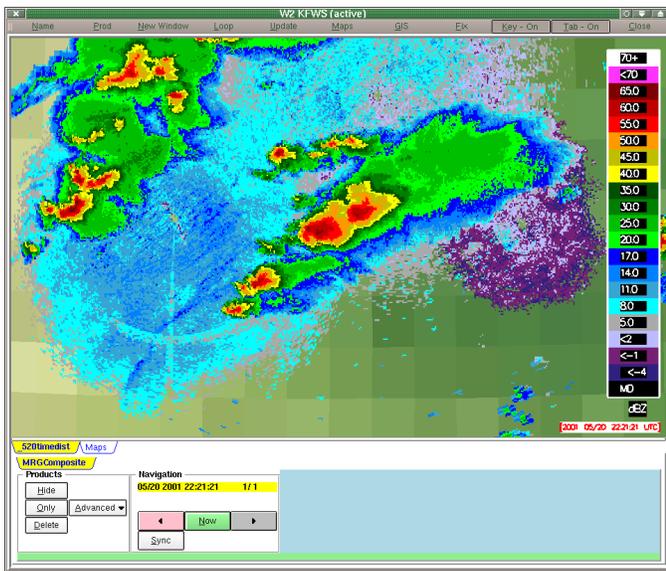


Figure 7: Blending data from two radar by assigning them weights inversely proportional to the distance from the radar and to the age of the data.

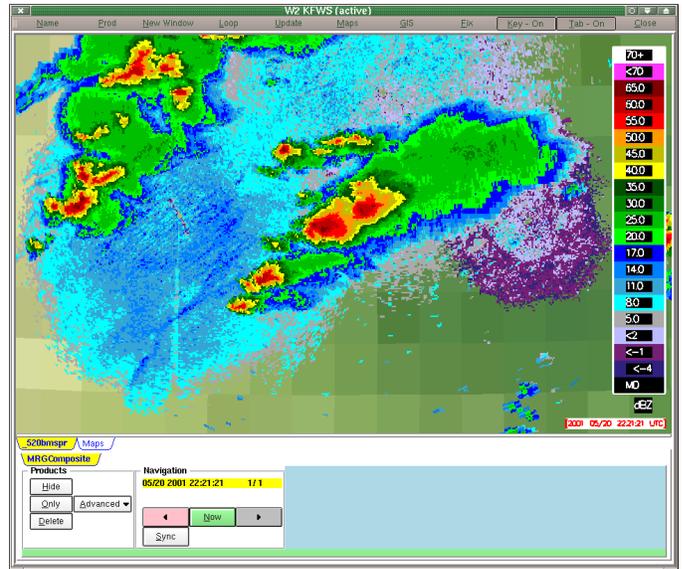


Figure 9: Blending data from different elevation scans of the same radar by linearly interpolating between them. This strategy is useful for single radar volumes.

References

- Gamma, E., R. Helm, R. Johnson, and J. Vlissides: 1994, *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 416 pp.
- Jenter, H. L. and R. P. Signell: 1992, NetCDF: A freely-available software-solution to data-access problems for numerical modelers. *Proceedings of the American Society of Civil Engineers Conference on Estuarine and Coastal Modeling*, Tampa, Florida.
- Lakshmanan, V.: 2002, An extensible, multi-source meteorological algorithm development interface. *21st Conference on Severe Local Storms*, Amer. Meteo. Soc., San Antonio, TX.
- Lynn, R. and V. Lakshmanan: 2002, Virtual radar volumes: Creation, algorithm access and visualization. *21st Conference on Severe Local Storms*, Amer. Meteo. Soc., San Antonio, TX.
- Smith, T.: 2002, A two-dimensional, local, linear, least-squares method of derivative estimates from doppler radial velocity. *21st Conference on Severe Local Storms*, Amer. Meteo. Soc., San Antonio, TX.
- Zhang, J., J. Gourley, K. Howard, and B. Maddox: 2000, Three-dimensional multiple radar reflectivity mosaic. *The Second Southwest Weather Symposium*, Tucson, AZ.