

A Real-Time, Three-Dimensional, Rapidly Updating, Heterogeneous Radar Merger Technique for Reflectivity, Velocity, and Derived Products

VALLIAPPA LAKSHMANAN AND TRAVIS SMITH

Cooperative Institute of Mesoscale Meteorological Studies, University of Oklahoma, and National Severe Storms Laboratory, Norman, Oklahoma

KURT HONDL

National Severe Storms Laboratory, Norman, Oklahoma

GREGORY J. STUMPF

Cooperative Institute of Mesoscale Meteorological Studies, University of Oklahoma, Norman, Oklahoma, and Meteorological Development Laboratory, National Weather Service, Silver Spring, Maryland

ARTHUR WITT

National Severe Storms Laboratory, Norman, Oklahoma

(Manuscript received 5 April 2005, in final form 18 January 2006)

ABSTRACT

With the advent of real-time streaming data from various radar networks, including most Weather Surveillance Radars-1988 Doppler and several Terminal Doppler Weather Radars, it is now possible to combine data in real time to form 3D multiple-radar grids. Herein, a technique for taking the base radar data (reflectivity and radial velocity) and derived products from multiple radars and combining them in real time into a rapidly updating 3D merged grid is described. An estimate of that radar product combined from all the different radars can be extracted from the 3D grid at any time. This is accomplished through a formulation that accounts for the varying radar beam geometry with range, vertical gaps between radar scans, the lack of time synchronization between radars, storm movement, varying beam resolutions between different types of radars, beam blockage due to terrain, differing radar calibration, and inaccurate time stamps on radar data. Techniques for merging scalar products like reflectivity, and innovative, real-time techniques for combining velocity and velocity-derived products are demonstrated. Precomputation techniques that can be utilized to perform the merger in real time and derived products that can be computed from these three-dimensional merger grids are described.

1. Introduction

The Weather Surveillance Radar-1988 Doppler (WSR-88D) network now covers most of the continental United States, and full-resolution base data from nearly all the WSR-88D radars are compressed using block encoding (Burrows and Wheeler 1994) and transmitted in real time to interested users (Droegemeier et

al. 2002). This makes it possible for clients of this data stream to consider combining the information from multiple radars to alleviate problems arising from radar geometry (cone of silence, beam spreading, beam height, beam blockage, etc.). Greater accuracy in radar measurements can be achieved by oversampling weather signatures using more than one radar.

In addition, data from several Terminal Doppler Weather Radars (TDWRs) are being transmitted in real time. Because the TDWRs provide higher spatial resolution close to urban areas, it is advantageous to incorporate them into the merged grids as well. The radar network resulting from the National Science

Corresponding author address: Valliappa Lakshmanan, CIMMS, OU/NSSL, 120 David L. Boren Blvd., Norman, OK 73072.
E-mail: lakshman@ou.edu

Foundation-sponsored Center for the Collaborative Adaptive Sensing of the Atmosphere (CASA) promises to provide data to fill out the <3 km umbrella of the Next-Generation Weather Radar network (Brotzge et al. 2005). Developing a real-time heterogeneous radar data combination technique is essential to the utility of the CASA network.

Further, with the advent of phased-array radars with no set volume coverage patterns and the possibility of highly adaptive scanning strategies, a real-time, three-dimensional, rapidly updating merger technique to place the scanned data into a earth-relative context is extremely important for downstream applications of the data.

A combination of information from multiple radars has typically been attempted in two dimensions. For example, the National Weather Service creates a 10-km national reflectivity mosaic (Charba and Liang 2005). We, however, are interested in creating a 3D combined grid because such a 3D grid would be suitable for creating severe weather algorithm products and for determining the direction individual CASA radars need to point.

Spatial interpolation techniques used to create 3D multiple-radar grids have been examined by Trapp and Doswell (2000) and Askelson et al. (2000). Zhang et al. (2005) consider spatial interpolation techniques from the point of retaining, as much as possible, the underlying storm structures evident in the single-radar data. Specifically, they show that a technique with the following characteristics suffices to create spatially smooth multiradar 3D mosaics: (a) nearest-neighbor mapping in range and azimuth, (b) linear vertical or bilinear vertical-horizonal interpolation, and (c) weighting individual range gates' reflectivity values with an exponential function of distance. In this paper, we build upon those results by describing the following:

- 1) how to combine the data using "virtual volumes"—a rapidly updating grid such that the merged grid has an (theoretically) infinitesimal temporal resolution;
- 2) how to account for time asynchronicity between radars;
- 3) a technique of precomputations in order to keep the latency down to a fraction of a second;
- 4) extracting subdomain precomputations from larger precomputations in order to create domains of merged radar products on demand;
- 5) a real-time technique of combining not just radar reflectivity data, but also velocity data; and
- 6) the derivation of multiple-radar algorithm products.

a. Motivation

Many radar algorithms are currently written to work with data from a single radar. However, such algorithms for everything from estimating hail sizes to estimating precipitation can perform much better if data from nearby radars and other sensors are considered. Using data from other radars would help mitigate radar geometry problems, achieve a much better vertical resolution, attain much better spatial coverage, and obtain data at faster time steps. Figure 1 demonstrates a case where information on vertical structure was unavailable from the closest radar, but where the use of data from adjacent radars filled in that information. Using data from other sensors and numerical models would help provide information about the near-storm environment and temperature profiles. Considering the numerous advantages of using all of the available data in conjunction, and considering that technology has evolved to the point where such data can be transmitted and effectively used in real time, there is little reason to consider single-radar vertical integrated liquid (VIL) or hail diagnosis.

b. Challenges in merging radar data

One of the challenges with merging data from multiple radars is that radars within the WSR-88D network are not synchronized. First, the clocks of the radars are not in sync. This problem will be fixed in the next major upgrade to the radar signal processors. A challenge faced by algorithms that integrate data from multiple radars is that the radar scanning strategies are not the same. Thus, a radar in Nebraska might be in precipitation mode, scanning 11 tilt angles every 5 min, while the adjacent radar in Kansas might be in clear-air mode, scanning just 5 angles every 10 min. Of course, this is actually beneficial. If the radar volume coverage patterns (VCPs) were synchronized, then we would not be able to sample storms from multiple radars at different heights almost simultaneously (because, in general, the same storm will be at different distances from different radars). Therefore, as opposed to the WSR-88D network radars being synchronized, it would be beneficial if the nonsynchronicity was actually planned, in the form of multiradar adaptive strategies.

As a result of radar beam geometry, each range gate from each radar contributes to the final grid value at multiple points within a 3D dynamic grid. Because some of the data might be 10 min old, while other data might be only a few seconds old, a naive combination of data will result in spatial errors. Therefore, the range gate data from the radars need to be advected based on a time synchronization with the resulting 3D grid. The




FIG. 1. Using data from multiple nearby radars can mitigate the cone of silence and other radar geometry issues. (a) A vertical slice through a full volume of data from the Dyess Air Force Base, TX (KDAY), radar on 6 Feb 2005. Note the cone of silence; this is information unavailable to applications processing only KDAY data. (b) The lowest elevation scan from the KDAY radar. (c) An equivalent vertical slice through the merged data from KDAY; Frederick, OK (KFDR); Lubbock, TX (KLBB); Midland, TX (KMAF); and San Angelo, TX (KSJT). Nearby radars have filled in the cone of silence from KDAY. (d) A horizontal slice at 3 km above mean sea level through the merged data.

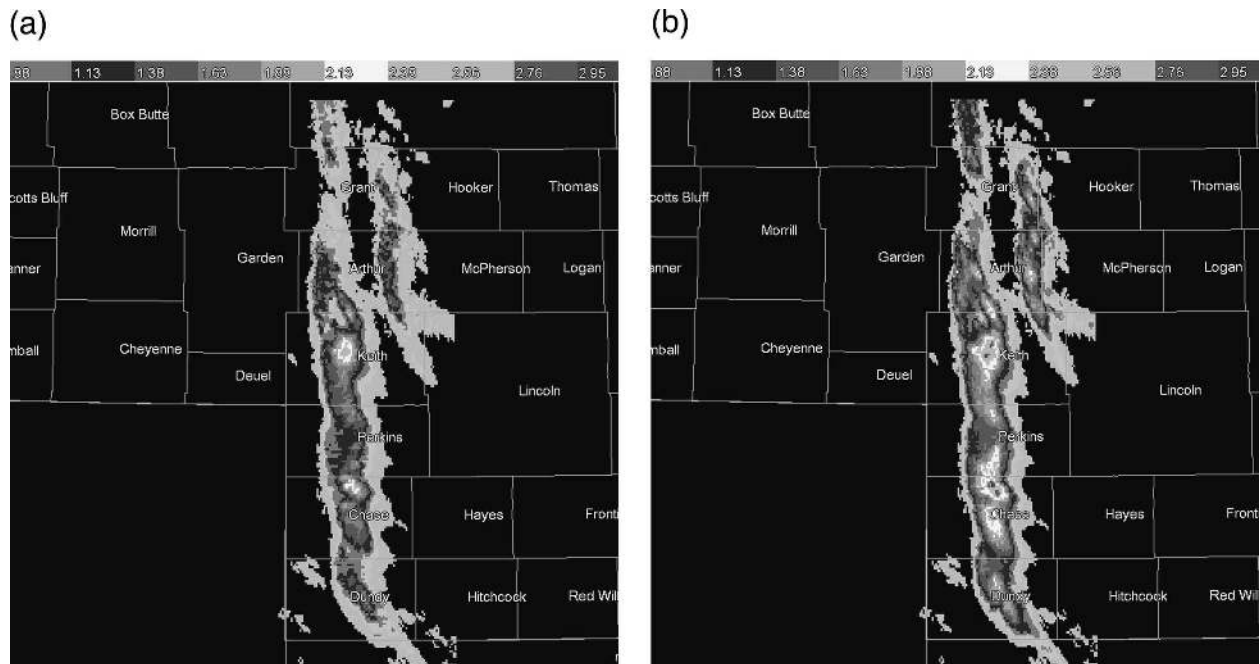


FIG. 2. Maximum expected size of hail (in.; 1 in. = 2.54 cm), calculated from reflectivity at various temperature levels. The temperature levels were obtained from the Rapid Update Cycle analysis grid. The radar reflectivity was estimated from KFDR; Amarillo, TX (KAMA); KLBB; and Dallas/Fort Worth, TX (KFWS) on 3 May 2003 (a) without advection correction, and (b) with advection correction. When advection correction is applied, the cores of the storms sensed by the different radars line up and the resulting merged products are less diffuse.

data need to be moved to the position in which the storm is anticipated to be. This adjustment is different for each point in the 3D grid. At any one point on the 3D grid there could be multiple radar estimates from each of the different radars.

The typical objective analysis technique used to create multiradar mosaics has been to utilize the latest complete volume of data from each of the radars, for example, the method of Charba and Liang (2005). In areas of overlap, the contributions from the separate radars are weighted by the distance of the grid point from the radar because beams from the closer radar suffer less beam spreading. There are two problems with this approach that we address. First, because the elevation scans are repeated once every 5–6 min, the same area of the atmosphere could be sensed as much as 5 min apart by different radars, leading to poor spatial location or smudging of the storm cells if blended without regard for temporal differences. Figure 2 demonstrates the smudging that can happen in multiradar scans, and how accounting for storm movement can significantly reduce such smudging. A second problem is the reliance on completed volume scans; a periodicity of 5 min is sufficient for estimating precipitation but is not enough for severe weather diagnosis and warning. For severe weather diagnosis and warning, spatially ac-

curate grids at a periodicity of at least once every 60 s are preferred by forecasters (Adrianto et al. 2005). In the technique described in this paper, a constantly updating grid is employed to get around the reliance on completed volume scans. Data sensed at earlier times are advected to the time of the output grid so as to resolve the storms better.

The rest of the paper is organized as follows. Section 2 describes the technique, starting off with a description of “intelligent agents,” a formulation we use to address the challenges described above. Using this technique to merge radar reflectivity, radial velocity, and derived products is then described in section 3. Section 4 describes extensions to the basic technique to handle beam blockage, quality control of input radar data, time correction, and optimizations for real-time use. Results and future applications of this work are summarized in section 5.

2. Intelligent agents

Our formulation of the problem of combining data from multiple radars, in the form of intelligent agents, provides a way of addressing the merger of both scalar (such as reflectivity) and vector (such as velocity–wind field) data.

Intelligent agents, sometimes called autonomous agents, are computational systems that inhabit some complex dynamic environment, sense and act autonomously in this environment, and by doing so realize a set of goals or tasks for which they are designed (Maes 1995). Smith et al. (1994) concisely characterize intelligent agents as persistent entities that are dedicated to specific purposes—the persistence requirement differentiates agents from formulas [Smith et al. (1994) use the term “subroutine”] because formulas are not long lived. The specific purpose characterization differentiates them from more complex “multifunction” algorithms.

In the context of merging radar data from multiple, possibly heterogeneous radars, each range gate of the radar serves as the impetus for the creation of one or more intelligent agents. Each intelligent agent monitors the movement of the storm at the position that it is currently in, and finds a place in the resulting grid based on time difference. Then, at the next time instant, the range gate migrates to its new position in the grid. The agents remove themselves when they expect to have been superseded. When new storm motion estimates are available, the agent updates itself with the new motion vector. When multiple agents all have an answer for a given point in the 3D grid, they collaborate to come up with a single value following strategies specified by the end user. These strategies may correspond to typical objective-analysis weighting schemes or techniques more suitable to the actual product being merged.

It should be clear that in the absence of time correction via advection of older data, the intelligent agent technique resolves directly to objective-analysis or multiple-Doppler techniques. Thus, the mathematical formulas in this paper will be presented in those more traditional terms.

The use of intelligent agents creates a flexible, scalable system that is not bogged down even by a highly “weather active” domain (e.g., a hurricane where most radar grid points have significant power returns). This technique of using an intelligent agent for each range gate with data from every radar in a given domain was proven to be robust and scalable even during Hurricanes Frances and Ivan in Florida (September 2004). The intelligent agents (about 1.3 million of them at one time) all collaborated flawlessly to create the high-resolution mosaic of data, shown in Fig. 3, of Hurricane Ivan from six different radars in Florida. Since the hurricane is over water and quite far from the coastline, no individual radar could have captured as much of the hurricane as the merged data have.

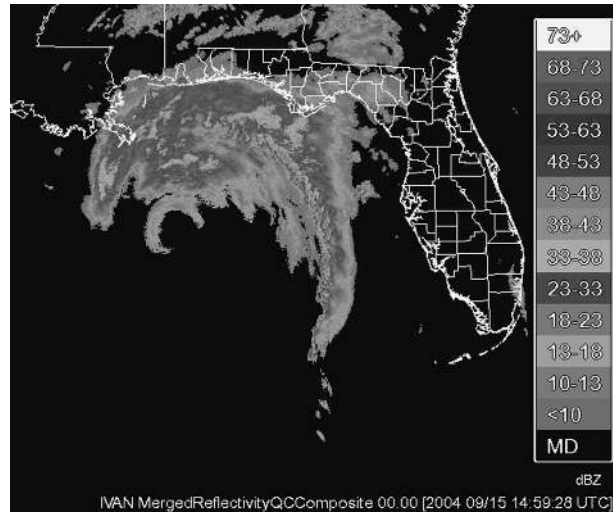


FIG. 3. Image of Hurricane Ivan consisting of combined data from six WSR-88D radars [Lake Charles, LA (KLCH); New Orleans/Baton Rouge (KLIX); Mobile, AL (KMOB); Tallahassee, FL (KTLH); Tampa Bay, FL (KTBW); and Key West, FL (KBYX)]. The images were created from the latest available WSR-88D data every 60 s (at approximately $1 \text{ km} \times 1 \text{ km} \times 1 \text{ km}$ resolution).

a. Agent model

Each range gate of the radar with a valid value serves as the impetus to the creation of one or more intelligent agents. The radar-pointing azimuths, often termed “rays” or “radials,” are considered objects within a radar volume that is constantly updating, and possibly with no semblance of regularity. Not relying on the regularity of radials is necessary to be able to combine data from phased-array radars (McNellis et al. 2005) and adaptive scanning strategies (Brotzge et al. 2005).

When an agent is created, it extracts some information from the underlying radial: (a) its coordinates in the radar-centric spherical coordinate system (range, azimuth, and elevation angle), (b) the radial start time, and (c) the radar the observation came from. All this information is typically available in the radar data regardless of the type of radar or the presence or absence of “volume coverage patterns.” The agent’s coordinates in the earth-centric latitude–longitude–height coordinate system of the resulting 3D grid have to be computed, however. The agents obtain these values following the 4/3 effective earth-radius model of Doviak and Zrnic (1993) (assuming standard beam propagation).

For a grid point in the resulting 3D grid (“voxel”) at latitude α_g , longitude β_g , and height h_g above mean sea level, the range gate that fills it, under standard

atmospheric conditions, is the range gate that is a distance r from the radar [located at (α_r, β_r, h_r) in 3D space] on a radial at an angle a from due north and on a scan tilted e to the earth's surface where a is given by

$$a = \sin^{-1}[\sin(\pi/2 - \alpha_g) \sin(\beta_g - \beta_r) / \sin(s/R)], \quad (1)$$

where R is the radius of the earth and s is the great-circle distance, given by (Beyer 1987)

$$s = R \cos^{-1}[\cos(\pi/2 - \alpha_r) \cos(\pi/2 - \alpha_g) + \sin(\pi/2 - \alpha_r) \sin(\pi/2 - \alpha_g) \cos(\beta_g - \beta_r)]. \quad (2)$$

The elevation angle e is given by (Doviak and Zrnice 1993)

$$e = \tan^{-1} \frac{\cos(s/IR) - \frac{IR}{IR + h_g - h_r}}{\sin(s/IR)}, \quad (3)$$

where I is the index of refraction, which under the same standard atmospheric conditions may be assumed to be 4/3 (Doviak and Zrnice 1993), and the range r is given by

$$r = \sin(s/IR)(IR + h_g - h_r) / \cos(e). \quad (4)$$

Because the \sin^{-1} function has a range of $[-\pi, \pi]$, the azimuth a is mapped to the correct quadrant ($[0, 2\pi]$) by considering the signs of $\alpha_g - \alpha_r$ and $\beta_g - \beta_r$. The voxel at α_g, β_g, h_g can be affected by any range gate that includes (a, r, e) .

An agent's life cycle depends on the radar scan. If the radar does not change its VCP, an agent can expect to be replaced $T_{\text{tot}} + T_i + T_{\text{lat}}$ later, where T_{tot} is the expected length of the volume scan, T_i the time taken to collect the scan that the agent belongs to, and T_{lat} is the latency in arrival of the scan after it has been collected. In practice, T_{lat} is estimated to be on the same order as T_i , because otherwise, it would not be possible to keep up with the data stream. At the end of its life cycle, the agent destroys itself. The time of the latest input radar data is used as the current time. If the radar VCP changes, there will be redundant agents if the periodicity of the VCPs decreases (because the older agents will be around a tad longer). On the other hand, if the periodicity lengthens, there will be a short interval of time with no agents. Because VCPs are set up such that periodicity decreases with the onset of weather and lengthens when weather moves away, there is no problem as long as the technique can reliably deal with redundant agents from the same radar. Redundant agents from the same radar can be dealt with using a time-weighting mechanism. By explicitly building in an allowance for redundant agents, radars with nonstandard scanning strategies such as a phased-array radar can oversample certain regions temporally and sense other regions less often.

Whenever an output 3D grid is desired, all of the

existing agents collaborate to create the 3D grid. Because heterogeneous radar networks are typically not synchronized, the agents will have to account for time differences. Naively combining all the data at a particular grid location in latitude–longitude–height space will lead to older data from one radar being combined with newer data from another radar. This will lead to problems in the initiation and decay phases of storms resulting in severe problems in the case of fast-moving storms. The agents, therefore, move to where they expect to be at the time of the grid. For example, an agent corresponding to a radar scan t seconds ago would move to x^1, y^1 , where

$$\begin{aligned} x^1 &= x + u_{xy} \times t \\ y^1 &= y + v_{xy} \times t, \end{aligned} \quad (5)$$

where x, y are the coordinates obtained from the raw azimuth–range–elevation values, and u_{xy}, v_{xy} are horizontal motion vectors at x, y (scaled to the units of the grid). Because of the difficulty of obtaining a vertical motion estimate, the vertical motion is assumed to be zero. The motion estimate may be obtained either from a numerical model or from a radar-based tracking technique such as in Lakshmanan et al. (2003b). We use the latter in the results reported in this paper.

b. Virtual volume

Whenever a new elevation scan is received from any of the radars contributing to the 3D grid, a set of agents is created. The elevation scan radial data are scale filtered to fit the resolution of the target 3D grid. For example, if the radial data are at 0.25-km resolution but the 3D grid is at a 1-km resolution, a moving average of four gates is used to yield the effective elevation scan at the desired scale. Instead of creating an agent corresponding to each range gate with valid values, we invert the problem and create an agent for each voxel in the 3D latitude–longitude–height grid that each such range gate impacts. The influence of a range gate from the center of the range gate extends to half the beamwidth




FIG. 4. Weighting function for interpolation between the centers of radar beams; α is the fraction of the distance to the next higher or lower beam.

in the azimuthal direction and half the gate width in the range direction. This is a nearest-neighbor analysis in the azimuth and range directions. In the elevation direction, the influence of this observation is given by δ_e , where

$$\delta_e = \exp[\alpha^3 \ln(0.005)], \quad \text{with} \quad \alpha = \frac{e - \theta_i}{|\theta_{i \pm 1} - \theta_i| V b_i}. \quad (6)$$

Here, α is the angular separation of the voxel from the center of the beam of an elevation scan (at elevation e) as a fraction either of the angular distance to the next higher or lower beam or the beamwidth. The V is a maximum operation, b_i the beamwidth of the elevation scan, and θ_i the elevation angle of its center. This function, shown in Fig. 4, is motivated by the fact that it is 1 at the center of the beam, goes to 0.5 at half-beamwidth, and falls to below 0.01 at the beamwidth (beyond which the influence of the range gate can be disregarded). It can be seen that where the elevation weight δ_e is less than 0.5, the voxel is outside the effective beamwidth of the radial. Points within the effective beamwidth are weighted more than they would be in a linear weighting scheme. The direction of the weighting is neither vertical nor horizontal, but tangential to the direction of the beam.

It should be noted that this interpolation is in the spherical coordinate system, in a direction orthogonal to the range–azimuth plane. Zhang et al. (2005) interpolate either in the vertical or in both the vertical and horizontal directions (i.e., in the coordinate system of the resulting 3D grid). Our technique is more efficient computationally, but could fail in the presence of strong vertical gradients such as bright bands, stratiform rain,




FIG. 5. Interpolation in the spherical coordinate system mitigates brightband effects in this stratiform event. The data are from KDYX, KFDR, KLBB, KMAF, and KSJT on 6 Feb 2005.

or convective anvils. Figure 5 demonstrates one case where spherical coordinate interpolation may suffice, but closer examination of a larger number of cases is needed. The optimal method for interpolation might be neither the spherical interpolation nor the vertical–horizontal interpolation, but to interpolate in a direction normal to the gradient direction (in 3D space).

3. Combining radar data

The methodology of performing objective analysis on radar reflectivity data is relatively clear. In section 2, we introduced a formulation of interpolating radar data onto a uniform grid that can address problems that arise in being able to perform the required computations in real time. In this section, we describe the application of the formulation described above to combining radar reflectivity, radial velocity, and derived products.

a. Combining scalar data

All of the intelligent agents that impact a particular voxel need to collaborate to determine the value at that voxel. In the case of scalar data, if each intelligent agent is aware of its influence weight, this reduces to determining a weighted average. The agent can determine its weight simply as an exponentially declining function of its distance from the radar. However, there could be multiple agents from the same radar that affect a voxel. In the case of a radar like TDWR, there could be multiple scans at the same elevation angle within a volume

scan. In the case of voxels that do not lie within a beam-width, there may be two agents corresponding to the adjacent elevation scans that straddle this voxel. Also, decaying or slowly moving storms often pile up agents into the same voxel because of the resolution of the grid points. Using all of these observations together, regardless of the reason that there are multiple agents from the same radar, should lead to a statistically more valid estimate.

There are two broad strategies for resolving the problem of having multiple observations (agents) from the same radar. One strategy is to devise a best estimate from each radar and then combine these estimates into an estimate for all the radars. The other is to combine all the agents regardless of the radar from which they come. In the case of velocity data, we use the former technique, while in the case of scalar data, the source of the data is ignored. However, to mitigate the problem of multiple, repeated elevation scans, the data are weighted both by time and distance. The influence weight of an observation is given by

$$\delta = \delta_e \exp[-(t^2 r^2 / \beta)], \quad (7)$$

where δ_e is the elevation angle weight [see Eq. (6)]; t is the time difference between the time the agent was created (when the radar observation was made) and the time of the grid; r is the range of the range gate from which this observation was extracted; and β is a constant of $17.36 \text{ s}^2 \text{ km}^2$, a number that was chosen through experimentation. Among a range of factors that we tried, this value seemed to provide the smoothest transitions while retaining much of the resolution of the original data.

The weighted sum of all the observations that impact a voxel is then assigned to that voxel in the 3D grid. This grid is used for subsequent severe weather product computations.

b. Combining velocity data

Unlike the methodology of combining scalar data, the methodology of combining velocity data in real time is not clear. Traditional multi-Doppler wind field retrieval is computationally intractable, because of its reliance on numerical convergence. In this paper, we present the following three potential solutions to this problem: (a) computing an “inverse” velocity azimuth display (VAD); (b) performing a multi-Doppler analysis, with certain approximations in order to keep it tractable; and (c) forgoing the wind field retrieval altogether, but merging shear, a scalar field derived from the velocity data. All three of the above techniques are applied after dealiasing the velocity data. For both the

WSR-88D and TDWR data, we apply the operational WSR-88D dealiasing algorithm. The intelligent agents used for combining the velocity data are created in the same manner as in the case of combining scalar data, but their collaborative technique is not an objective analysis one.

The multi-Doppler technique is based on the over-determined dual-Doppler technique of Kessinger et al. (1987). The terminal velocity was estimated from the equivalent radar reflectivity factor (Foote and duToit 1969). We initially attempted a full 3D version of the multi-Doppler technique, because severe storms do contain regions of strong vertical motion, and it would be advantageous to estimate the full 3D wind field. Unfortunately, test results for the full 3D technique were unsatisfactory. The vertical velocities in that technique, computed via the mass continuity equation, turned out to be numerically unstable and propagated errors into the horizontal wind fields. Instead of abandoning the process entirely, we decided to try a simplified version of the technique, which assumes that $w = 0$. Despite the fact that this assumption of $w = 0$ will not be valid for some regions of severe storms, initial test results for this 2D version of the technique were promising (Witt et al. 2005). Test results for the 2D version of the technique on two severe weather cases showed very good agreement between the calculated horizontal wind field and the corresponding radial velocity data. The 2D wind field also closely matched what conceptual models of the airflow in severe storms would suggest.

An example of merging reflectivity and velocity data from two heterogeneous radars—a WSR-88D at Oklahoma City/Twin Lakes, Oklahoma (KTLX), and the Oklahoma City, Oklahoma (OKC), TDWR—is shown in Fig. 6. A single horizontal level, at 1.5 km above mean sea level, is shown for the merged products. It should be noted that the reflectivity images from the two radars, KTLX and OKC, are different because the KTLX scan is 5 min older than the OKC one and the storm has moved in the time interval. The merged radar grid does have the storm in the right location at the reference time. Note also that the wind field retrieval has correctly identified the rotation signature.

Because the vertical motion computed from an over-determined dual-Doppler technique is not useful, we sought to examine if a more direct way of computing 2D horizontal wind fields from the radar data was feasible. The VAD technique may be used to estimate the U and V wind components from the radial velocity observations using a discrete Fourier transform (Lhermitte and Atlas 1961; Browning and Wexler 1968; Rabin and Zrnic 1980). The VAD technique uses a

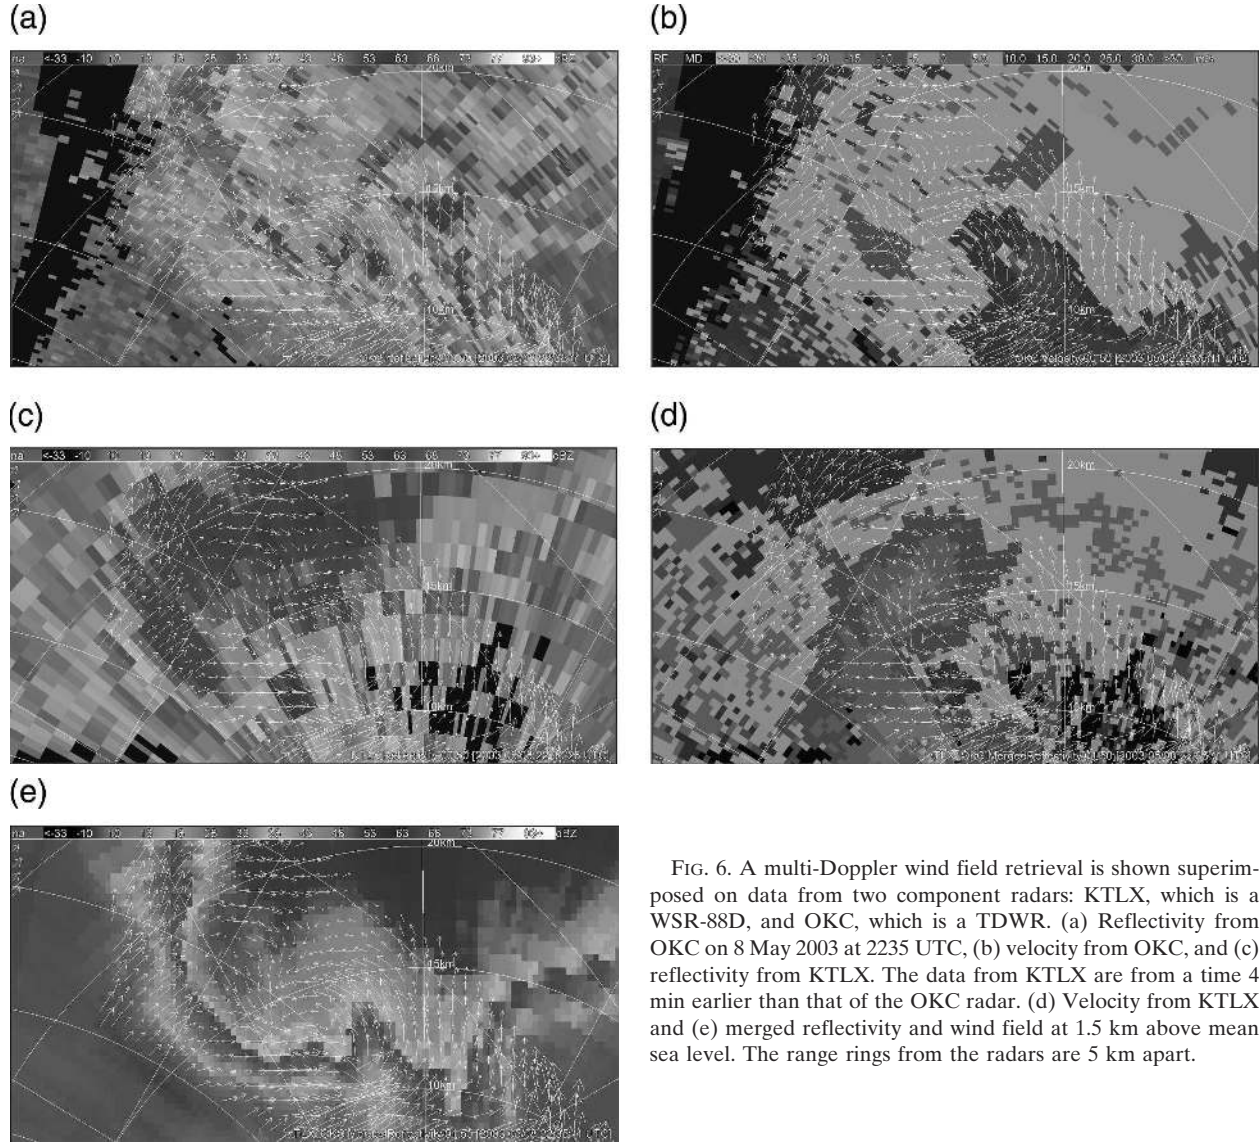


FIG. 6. A multi-Doppler wind field retrieval is shown superimposed on data from two component radars: KTLX, which is a WSR-88D, and OKC, which is a TDWR. (a) Reflectivity from OKC on 8 May 2003 at 2235 UTC, (b) velocity from OKC, and (c) reflectivity from KTLX. The data from KTLX are from a time 4 min earlier than that of the OKC radar. (d) Velocity from KTLX and (e) merged reflectivity and wind field at 1.5 km above mean sea level. The range rings from the radars are 5 km apart.

least squares approximation to calculate the first harmonic from the radial velocity observations at different azimuths as observed from the radar location. The inverse VAD technique similarly uses a least squares solution to determine the 2D wind components at a point in space from the radial velocity observations from different azimuths (i.e., as observed from different radars).

The radial velocity observed from the i th radar v_i is dependent on the wind components u_o and v_o and the observing angles ϕ_i of the radars,

$$\mathbf{V}_{n \times 1} = \mathbf{P}_{n \times 2} [u_o v_o]^T, \quad (8)$$

where

$$\mathbf{V}_{n \times 1} = [v_1, v_2, v_3, \dots, v_n]^T \quad (9)$$

and $\mathbf{P}_{n \times 2}$ is given by

$$\begin{bmatrix} \sin \phi_1 & \cos \phi_1 \\ \sin \phi_2 & \cos \phi_2 \\ \dots & \dots \\ \sin \phi_n & \cos \phi_n \end{bmatrix}. \quad (10)$$

If one has the radial velocity observations and the azimuth of the observations, a least squares approximation of the wind components may be estimated using a least squares formulation as

$$[u_o v_o]^T = (\mathbf{P}_{n \times 2}^T \mathbf{P}_{n \times 2})^{-1} (\mathbf{P}_{n \times 2}^T \mathbf{V}_{n \times 1}). \quad (11)$$

The inverse VAD technique is viable as long as there are some radial velocity observations from nearly orthogonal angles.

In Fig. 7, we demonstrate the technique on a simulation of three CASA radars. The radial velocity data were created from a network of simulated radars observing the numerical simulation (Biggerstaff and May 2005) of a tornadic storm. The inverse-VAD wind field technique was used to retrieve the 2D wind field from the Doppler velocity corresponding to the three simulated radars. The circulation signature is clearly identifiable in the wind vector plot and correlates with the location of the shear signature in the radial velocity data.

While it is useful to be able to perform multi-Doppler velocity analysis or inverse-VAD analysis in real time to retrieve wind fields, the applicability of such a processing is limited to radars that are somewhat proximate to each other, and situated such that the radars' viewing angles are nearly orthogonal to each other.

If we were to consider the WSR-88D network alone, such situations are rare and made more so by the fact that unlike reflectivity data where the surveillance scans go out to 460 km, velocity data go out only to about 230 km. Thus, there are few places where such analysis may be performed. In this paper, we suggested the use of TDWR radars in addition to the WSR-88D network to increase areas of overlap and showed that the results were promising, at least for horizontal wind vectors. However, the combination of data from radars of different wavelengths requires further study. We demonstrated the merging of data from an S-band radar and a C-band radar in this paper, but in that particular instance there was no noticeable attenuation in the C-band data. Besides, even in the presence of attenuation in reflectivity, velocity data might still be usable. As CASA radars (where there will be considerable overlap between the radars in the network) are deployed more widely, the real-time multi-Doppler analysis methods described here will become more practical.

Because of the limitations of merging velocity data to retrieve wind fields, many researchers (e.g., Liou 2002) have examined the use of single-Doppler velocity retrieval methods. It is possible that a merger of single-radar-retrieved wind fields may prove beneficial. It is also likely that applying single-Doppler velocity retrieval methods to data from a spatially distributed set of radars might yield robust estimates of wind fields. We have limited ourselves, in this paper, to describing multiradar retrievals of wind fields that can be performed in real time using well-understood techniques.

Traditional methods of calculating vorticity and di-

vergence from Doppler radial velocity data can yield unreliable results. We use a two-dimensional, local, linear least squares (LLSD) method to minimize the large variances in rotational and divergent shear calculations (Smith and Elmore 2004). Besides creating greater confidence in the value of the 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, the azimuthal shear that results from the LLSD is a scalar quantity that can be combined using the same technique as is used for radar reflectivity as shown in Fig. 8.

c. *Derived products*

In addition to merging reflectivity and velocity data from individual radars into a multiradar grid, it is possible to apply the same method of merging reflectivity data to any scalar field derived from the radar moment data. When doing so, it is essential to justify the reason for combining derived products instead of deriving the product from the combination of moment data since the latter is more efficient computationally.

Shear is a scalar quantity that needs to be computed in a radar-centric coordinate system, because it is computed on Doppler velocity data taking the direction of the radial beam into account. Therefore, it is necessary to compute the shear on data from individual radars and then combine them. Having computed the shear, we can accumulate the shear values at a certain range gate over a time interval (typically 2–6 h), and then merge the maximum value over that time period from individual radars. The strategy of blending such a maximum value from multiple radars is to take the value whose magnitude (disregarding the sign) is maximum. Such a "rotation track field" (shown in Fig. 9) is useful for conducting poststorm damage surveys.

Once a 3D merged grid of radar reflectivity is obtained, it is possible to run many severe weather algorithms on this grid. For example, a multiradar vertical composite product is shown in Fig. 3. The VIL was introduced in Greene and Clark (1972) using data from a single radar. A multiradar VIL product is shown in Fig. 10. VIL estimated from storm cells identified from multiple radars is a more robust estimate than VIL estimated using just one radar (Stumpf et al. 2004). Figure 11 demonstrates that the multiradar VIL is longer lived and more robust.

Three-dimensional grids of reflectivity are created at constant altitudes above mean sea level. By integrating numerical model data, it is possible to obtain an estimate of isotherm heights. Thus, it is possible to com-

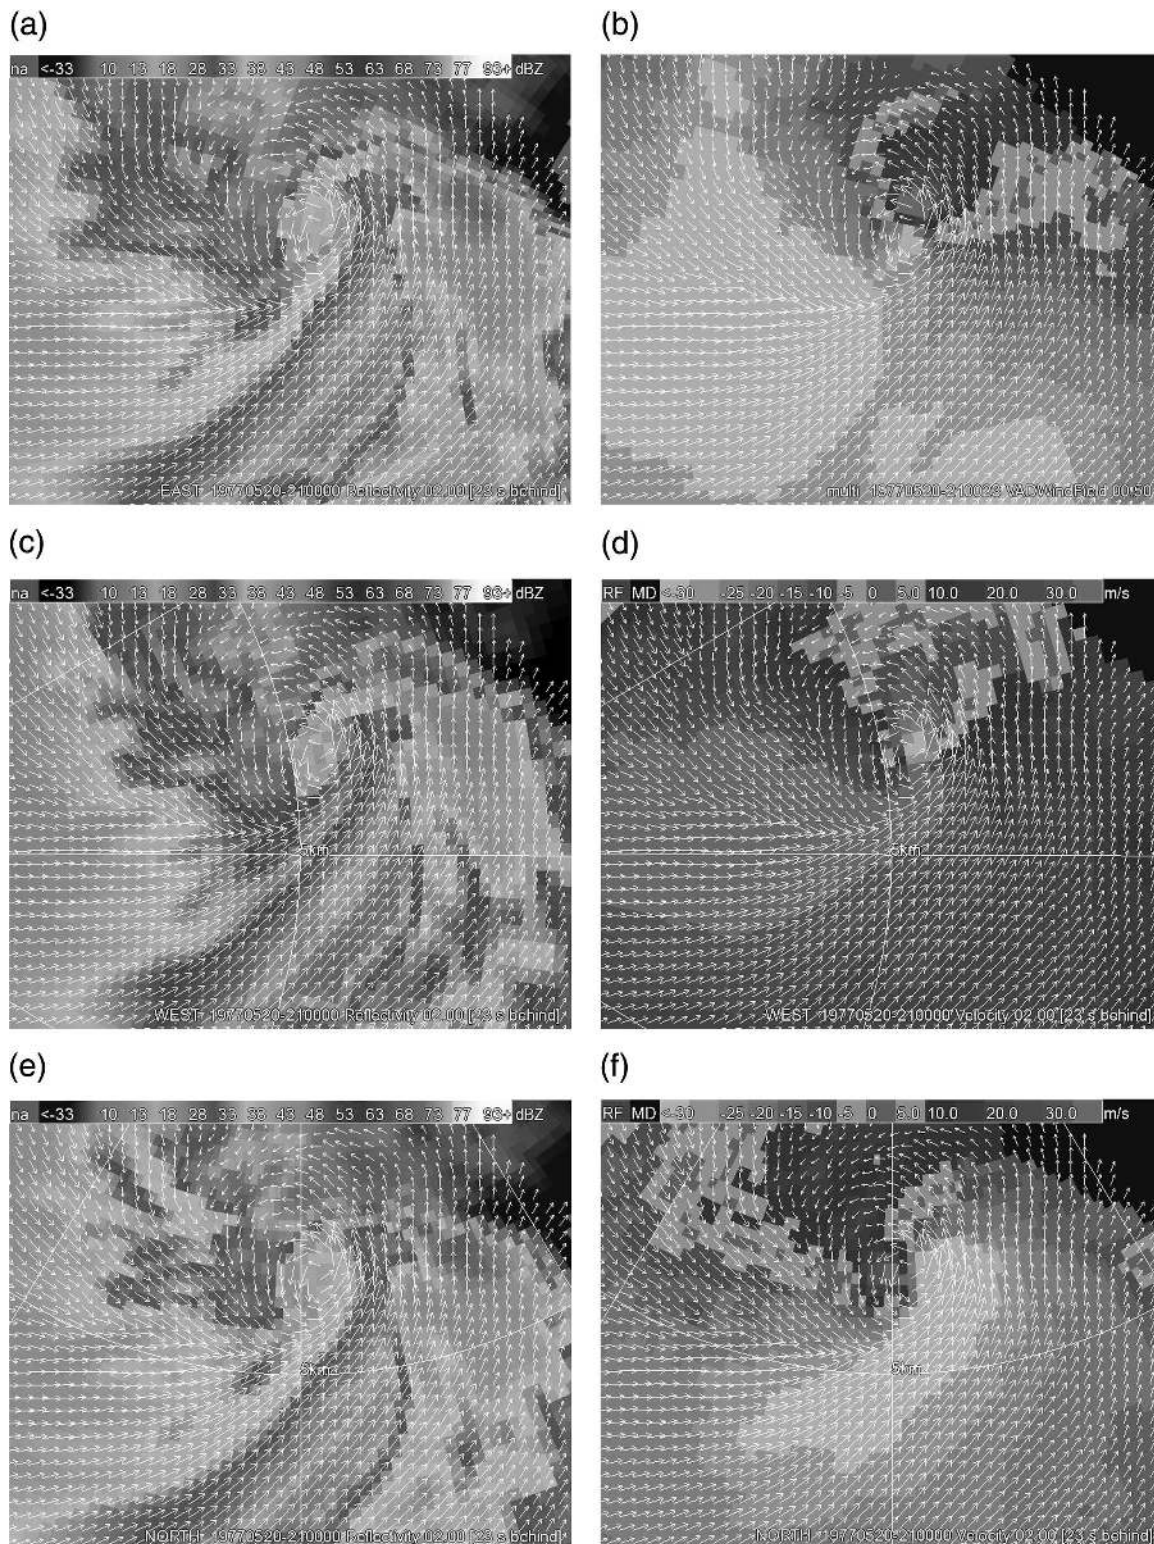


FIG. 7. An inverse-VAD wind field retrieval is shown superimposed on data from the component radars. (left) Reflectivity and (right) velocity are from simulated CASA radars (top) east, (middle) west, and (bottom) north of the area of interest. The range rings are at 5-km intervals. Note that the wind field retrieved using the inverse-VAD technique has captured the thunderstorm's rotation.

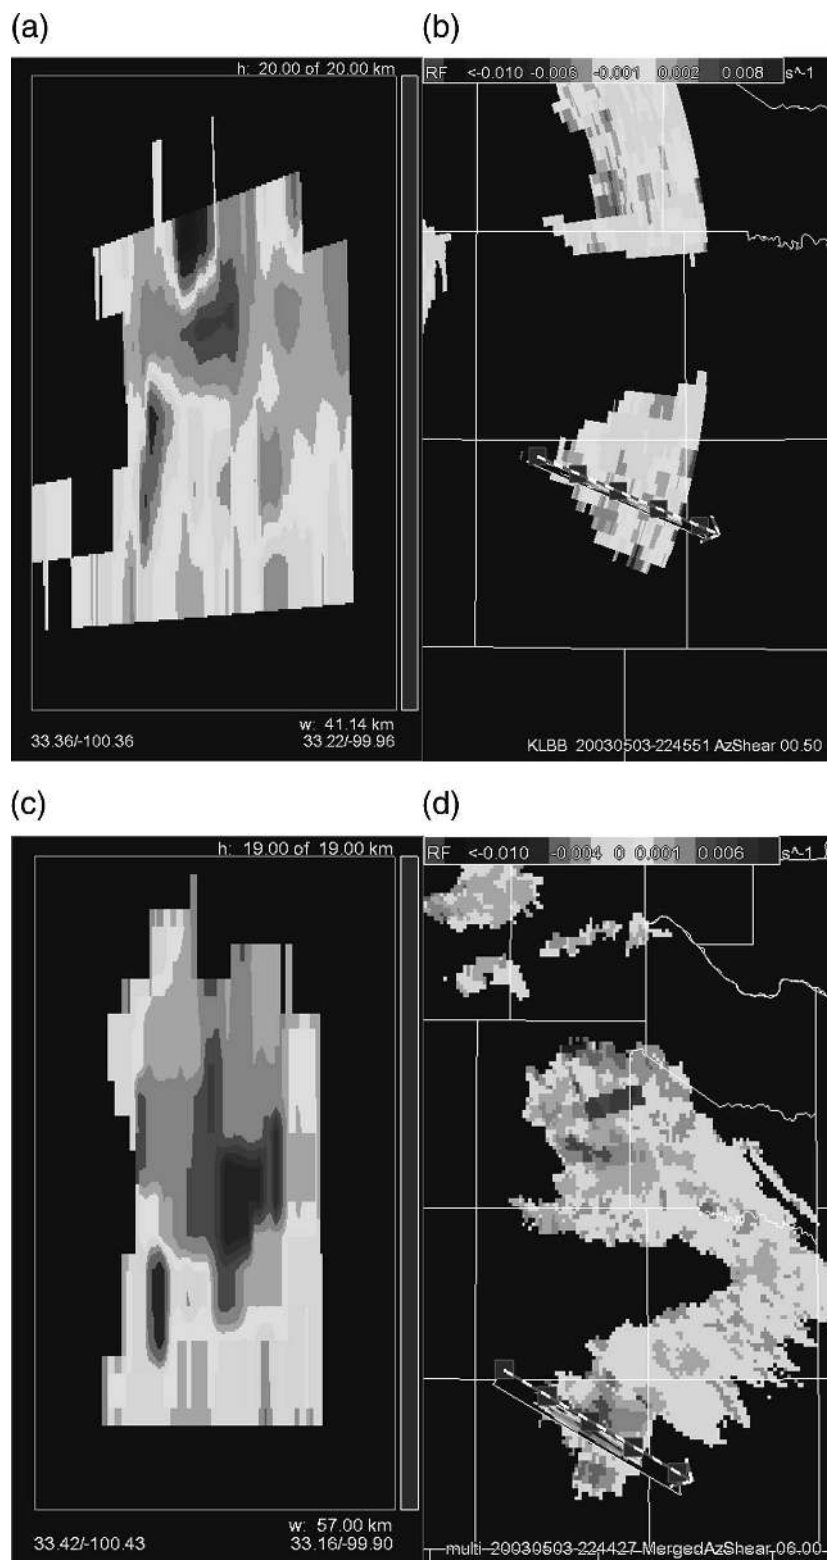


FIG. 8. (a) Vertical slice through the azimuthal shear computed from a volume of data from the KLBB radar on 3 May 2003. (b) Azimuthal shear computed from a single elevation scan. (c) Vertical slice through the multiradar merged azimuthal shear from KFDR, KAMA, KLBB, and KFWS. (d) A 6-km horizontal slice through the multiradar data.




FIG. 9. A rotation track field created by merging the maximum observed shear over time from single radars. The overlaid thin lines indicate the paths observed in a postevent damage survey. Note that the computed path of high shear corresponds nicely with the damage observed. Data from 3 May 1999 in the Oklahoma City area are shown.

pute the reflectivity value from multiple radars and interpolate it to points not on a constant-altitude plane, but on a constant-temperature level. This information, updated every 60 s in real time, is valuable for forecasting hail and lightning (Stumpf et al. 2005). The inputs for a lightning forecasting application that makes use of these data are shown in Fig. 12.

The technique to map reflectivity levels to constant temperature altitudes is used to transform the technique of the hail detection algorithm (Witt et al. 1998) from a cell-based technique to a gridded field. A quantity known as the severe hail index (SHI) vertically integrates reflectivity data with height in a fashion similar to VIL. However, the integration is weighted based on the altitudes of several temperature levels. In a cell-based technique, this is done using the maximum dBZ values for the 2D cell feature detected at each elevation scan. For a grid-based technique, the dBZ values at each vertical level in the 3D grid are used and compared with the constant temperature altitudes. From SHI we compute the probability of severe hail and the maximum expected hail size values, which are also plotted on a grid (see Fig. 2). Having hail size estimates on a geospatial grid allows warning forecasters to understand precisely where the largest hail is falling. These grids can also be compared across a time interval to map the swathes of the largest hail or estimate the hail

damage by combining the hail size and duration of the hail fall. These hail swathes can later be used to enhance warning verification. They can also be used to provide 2D aerial locations of hail damage to first responders in emergency management and in the insurance industry.

The 3D reflectivity grids can also be used to identify and track severe storm cells, and to trend their attributes. This is presently done using two techniques. The first is a method similar to that developed for the storm cell identification and tracking algorithm that uses a simple clustering method to extract cells of a given area and vertical extent (Johnson et al. 1998). Another technique is to compute the motion estimate directly from derived products on the 3D grid following Lakshmanan et al. (2003b). For example, either the multiradar VIL or the multiradar reflectivity vertical composite may be used as the input frames to the motion estimation technique. A *K*-means clustering technique is used to identify components in these fields.

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. Typically, frames 10 or 15 min apart are chosen. Given the motion estimates for each of the regions in the image, the motion estimate at each pixel is determined through interpolation. This motion




FIG. 10. Multiradar VIL product at high spatial (approximately $1 \text{ km} \times 1 \text{ km}$) and temporal (60-s update) resolution from the latest reflectivity data from four radars: KFDR, KAMA, KLBB, and KFWS on 3 May 2003.

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). This motion estimation technique is used




FIG. 11. The VDL computed on cells detected from data blended from individual radars is more robust than VIL computed on storms cells identified from a single radar. As the storm approaches the radar, the single-radar VIL drops because higher-elevation data are unavailable. The multiradar product does not have that problem. Image courtesy of Stumpf et al. (2004).

as a source of u_{xy} , v_{xy} , the anticipated movement of the intelligent agent currently at x, y in the 3D grid [see Eq. (5)].

4. Extensions to method

While the basic technique of creating intelligent agents and combining them yields reasonable results, we extended the basic technique by correcting for beam blockage from individual radars, improving the quality of the input data, correcting for time differences between the radars, and devising optimizations to enable the technique to be used in real time.

a. Corrections for beam blockage

A range gate from a radar elevation scan is assumed to not impact a voxel if it falls within a beam-blockage umbrella because of terrain. Currently, for reasons that will become evident in section 4c, a standard atmospheric model with an effective $4/3$ earth radius (Doviak and Zrnic 1993) is assumed to determine which radar beams will be blocked by terrain features. A terrain digital elevation map (DEM) at the scale of the desired grid (approximately $1 \text{ km} \times 1 \text{ km}$) was used for the results presented here, but the technique would apply even if higher-resolution terrain maps were used.

The assumptions for a beam being blocked very closely follow the techniques of O'Bannon (1997) and Fulton et al. (1998). Interested readers should consult those papers for further information. For each point in the DEM, the azimuth, range, and elevation angle of a thin radar beam are computed following the standard beam propagation assumptions and tolerances as given by O'Bannon (1997). Any thin beam above this elevation angle passes above this terrain point unblocked. Other beams are assumed to be blocked by this elevation point. Every radial from the radar is then considered a numerical integrand of all the thin beams that fall within its range of azimuths and elevations. The influence of the individual thin beams follows the power density function of Doviak and Zrnic (1993). Thus, a fraction of the beam that is blocked is known at each range gate. If this fraction is greater than 0.5, the entire range gate is assumed to be blocked; the data from such range gates are not used to create new intelligent agents. However, because of advection, it is possible that a voxel that would be beam blocked might get a value due to the movement of an agent created from a nonblocked range gate. Beam-blocked voxels can also get filled in by data from other nearby radars, leading to a more uniform spatial coverage than what is possible using just one radar. Figure 13 demonstrates the

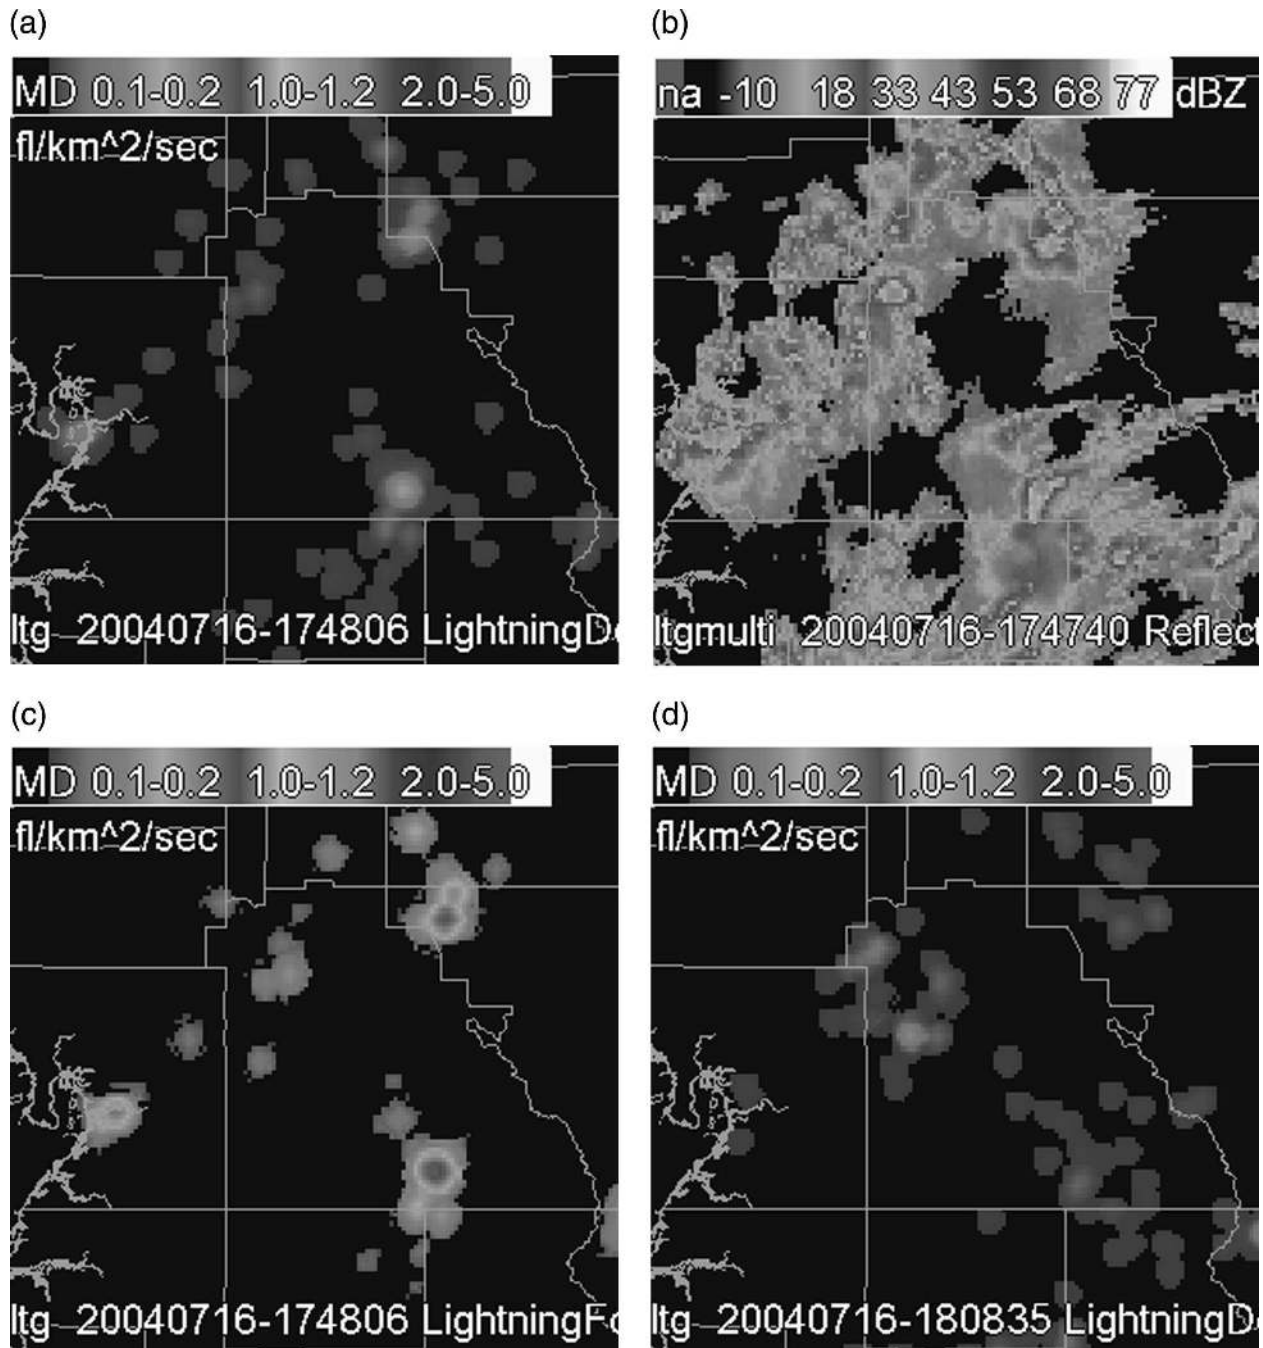


FIG. 12. (a) Current lightning density; (b) reflectivity at a height of -10°C temperature, used as input to the lightning prediction; (c) the 30-min lightning forecast; and (d) the actual lightning 30 min later.

filling of data from adjacent radars to yield better spatial coverage when some parts of a radar domain are beam blocked.

Similar to the beam blockage correction, it should be possible to dynamically correct for inaccurate clocks on individual radars and for heavy attenuation from individual radars. Our current implementation does neither

because of our initial focus on the WSR-88D network. The inaccurate clocks on the WSR-88D network are to be fixed with automatic time-correction software in the radar sites. WSR-88Ds, being S-band radars, do not attenuate as severely as X-band radars. For C-band radars such as TDWR and X-band radars such as those that will be used in the CASA network, an attenuation

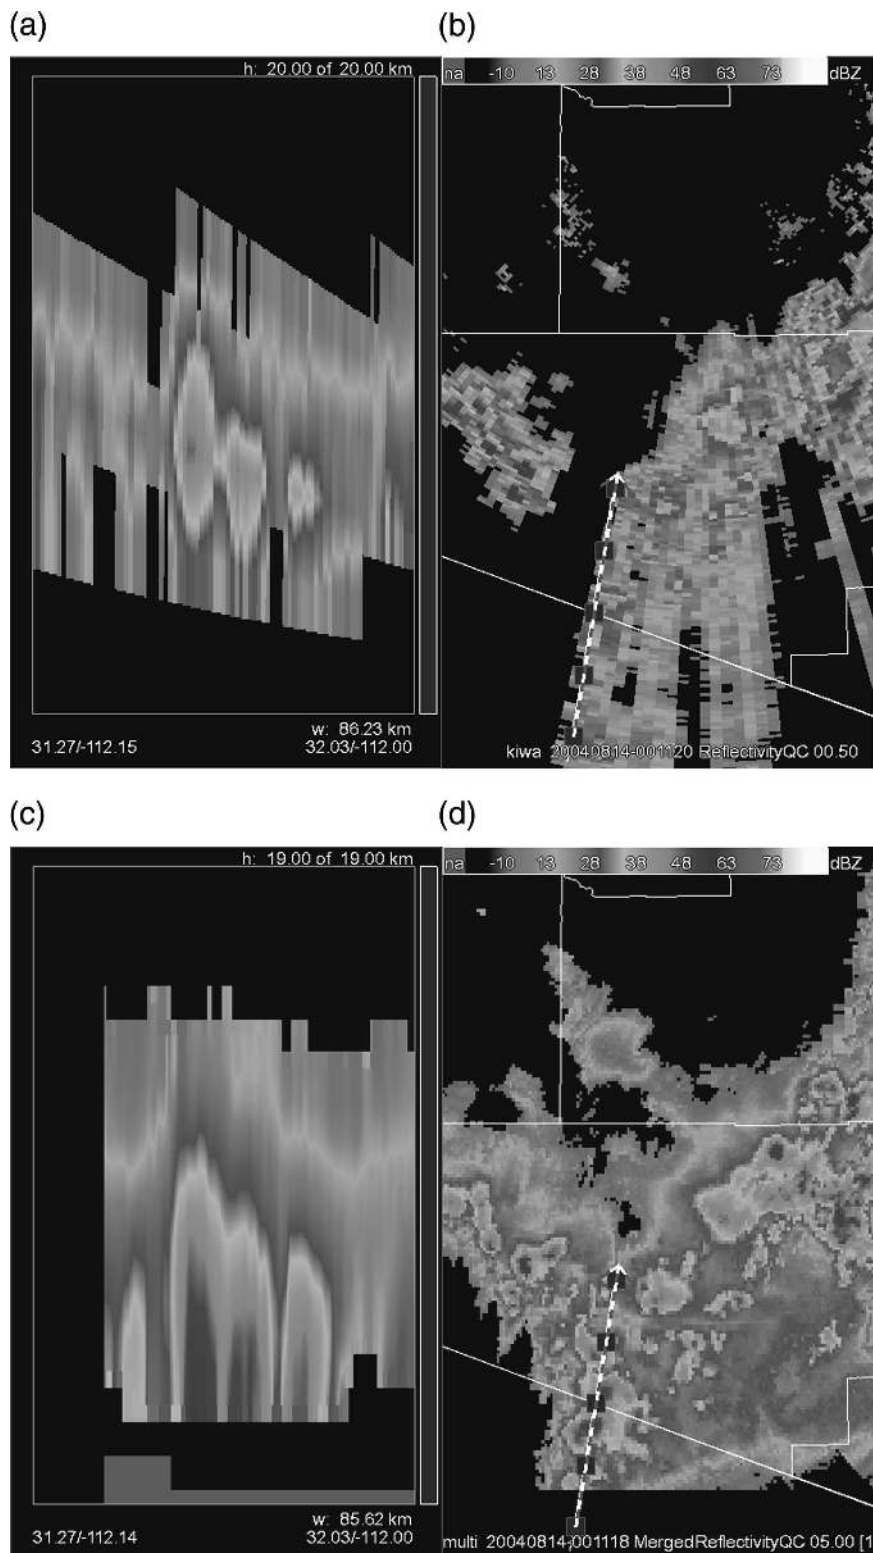


FIG. 13. (a) Vertical slice through data from the Phoenix, AZ (KIWA), radar on 14 Aug 2004. (b) Elevation scan from the KIWA radar, note the extensive beam blockage. (c) Vertical slice through multiradar data from KIWA; Tucson, AZ (KEMX); Yuma, AZ (KYUX); Flagstaff, AZ (KFSX); and Albuquerque, NM (KABX) covering the same domain. (d) A 5-km horizontal slice through multiradar data. Note that the entire domain is filled in.

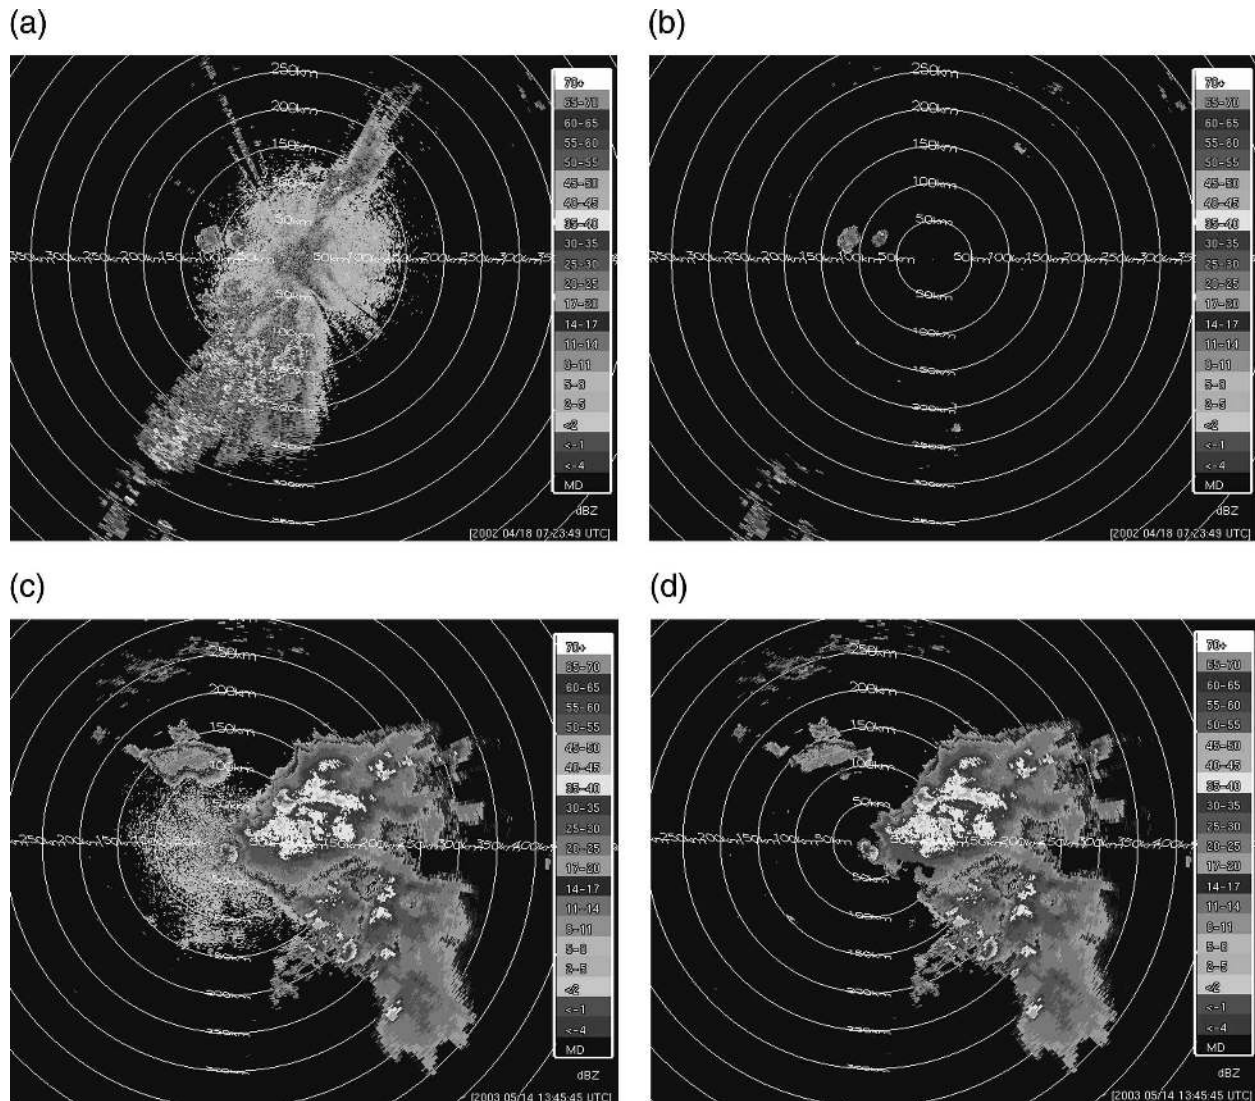


FIG. 14. Independent test cases for the Quality Control Neural Network (QCNN). (a) A data case from KAMA with significant AP. The range rings are 50 km apart. (b) Edited using the neural network; note that all of the AP has been removed, but the storm cells northwest of the radar are retained. (c) A typical spring precipitation case. (d) Edited using the neural network; note that even the storm cell almost directly overhead the radar has been retained, but biological scatter has been removed.

correction will have to be put in place to avoid an underreporting bias in the multiple-radar grids.

b. Quality control of input data

Weather radar data are subject to many contaminants, mainly due to nonprecipitating targets (such as insects and wind-borne particles) and anomalous propagation (AP) or ground clutter. If the radar data are directly placed into the 3D merged grid, these artifacts lower the quality of the gridded data. Hence, the radar velocity data need to be dealiased and the radar reflectivity data need to be quality controlled. We used

the standard WSR-88D dealiasing algorithm to dealias the velocity data.

In radar reflectivity data, several local texture features and image processing steps can be used to discriminate some types of contaminants (Kessinger et al. 2003). However, none of these features by themselves can discriminate between precipitating and nonprecipitating areas. A neural network is used to combine these features into a discrimination function (Lakshmanan et al. 2003a). Figure 14 demonstrates that the neural network is able to identify, and remove, echoes due to anomalous propagation while retaining echoes due to precipitation.

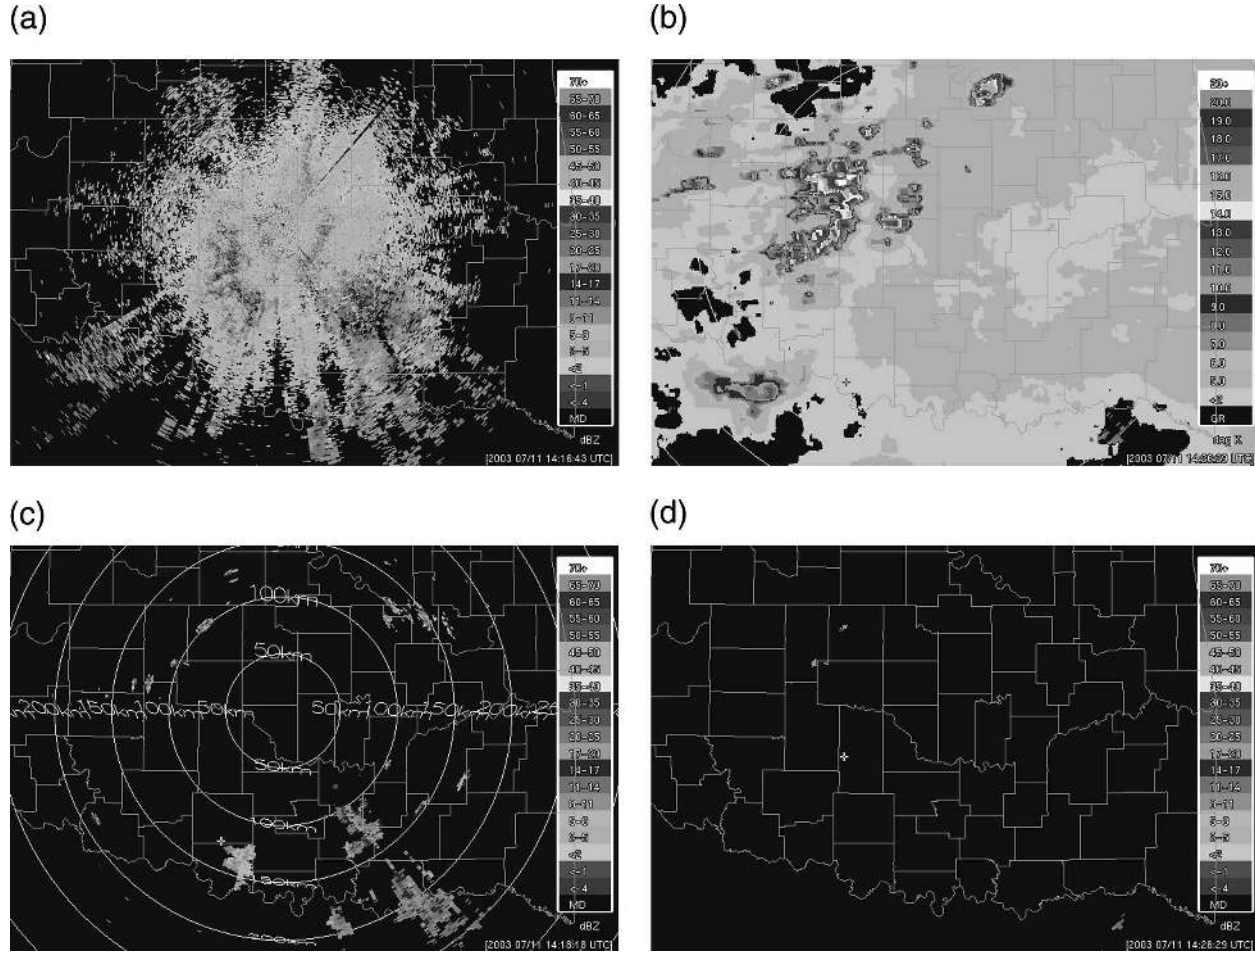


FIG. 15. (a) KTLX reflectivity composite showing the effects of biological contamination. (b) The cloud cover field derived from satellite data and surface observations. (c) The effect of the radar-only quality control neural network. (d) The effect of using both the radar-only neural network and the cloud cover field. Note that the small cells to the northwest of the radar are unaffected, but the biological targets to the south of the radar are removed.

No current technique using only single radar-data (ignoring polarimetric data) can discriminate between shallow precipitation and spatially smooth clear-air returns (Lakshmanan and Valente 2004). The radar-only techniques also have problems removing some biological targets, chaff, and terrain-induced ground clutter far away from the radar. In addition to the radar-only quality control above, an additional stage of multisensor quality control is applied. This uses satellite data and surface temperature data to remove clear-air echoes. Figure 15 demonstrates that biological returns may be removed by using such cloud cover information. For more details, the reader is directed to Lakshmanan and Valente (2004).

The radar reflectivity elevation scans are quality controlled, either using the radar-only quality control technique described in Lakshmanan et al. (2003a) or using the multisensor technique described in Lakshmanan

and Valente (2004). It is these quality-controlled reflectivity data that are presented to the agent framework for incorporation into the constantly updating grid.

c. Precomputation

The coordinate system transformation to go from each voxel α_g, β_g, h_g to the spherical coordinate system (a, r, e) can be precomputed as long as we limit the input radars to be nonmobile units (so that the radar position does not change). If the radar follows one of a set number of VCPs, then the elevation scan can be determined from e . If the elevation scans are indexed to always start at a specific azimuth and each beam constrained to an exact beamwidth (WSR-88D scans are not), then a, r can also be mapped to specific locations in the radar array. In fact, the presence of a VCP is not required to precompute the effect of data from an elevation angle; all that is required is that the radar op-

erate only at a limited number of elevation angles, perhaps in a range of $[0^\circ, 20^\circ]$ in increments of 0.1° . Then, the voxels impacted by data from an elevation scan can be computed beforehand and stored as part of the computer's permanent memory (i.e., on the hard drive) for immediate recall whenever that elevation angle from that radar is received. Intelligent agents can then be created, with coordinates assigned to them in both coordinate systems at the time of creation. If the VCPs are known or if the radar operates only at set elevations, it is possible to precompute the elevation weight δ_e . However, because of the dependence on t , a time-delay variable, it is necessary to compute δ at the time of grid creation. Similarly, the presence of t in the advection equations necessitates that the movement of the agents to their new positions be dynamic and not precomputed.

These precomputations have to be performed for every possible elevation angle from every radar that will be used as an input to the merged 3D grid. For a grid of approximately $800 \text{ km} \times 800 \text{ km}$ and about 10 radars, a typical regional domain, the precomputation can take up to 8 h on a workstation with a 1.8-GHz Intel Xeon processor, 2 GB of random access memory, and a 512-MB cache (referred to from hereon as simply Intel Xeon). Thus, it is necessary to decide upon a regional domain well in advance of the storm event, or to have enough hardware available to process a large radar domain reflecting the threat 6–8 h in advance.

d. Extraction of subdomains

Naturally, determining the domain of interest 6–8 h in advance is not a trivial task. Is it possible to reuse precomputations? Because the coordinate system of the output 3D grid is a rectilinear system (the α_g, β_g, h_g are additive), we can precompute the transformation of data from any radar in the country at every elevation angle possible onto a domain the size of the entire continental United States (CONUS). The CONUS domain can be created at the desired resolution. We currently use 0.01° in latitude and longitude and 1 km in height, approximately $1 \text{ km} \times 1 \text{ km} \times 1 \text{ km}$ at midlatitudes. Then, the coordinates of the range gate in a subdomain of the CONUS domain can be computed from its coordinates in the CONUS domain using the offsets of the corners. If the northwest corner of the CONUS domain is α_c, β_c, h_c and that of the desired subdomain is α_s, β_s, h_s , then the additive correction to the α_g, β_g, h_g entries in the CONUS cache to yield correct subdomain entries is

$$\begin{aligned}\delta_\alpha &= (\alpha_g - \alpha_c)/\text{res}_\alpha \\ \delta_\beta &= (\beta_c - \beta_g)/\text{res}_\beta\end{aligned}\quad (12)$$

where $\text{res}_\alpha, \text{res}_\beta$ are the resolution of the 3D grid in the latitudinal and longitudinal directions. The only condition is that the above operations have to yield integers, because they will be indexes into the CONUS domain arrays. Thus, the limitation is that subdomains have to be defined at $0.01^\circ/1\text{-km}$ increments. Computation of the CONUS domain takes about 120 h on a dual-processor workstation and requires 77 GB of disk space. However, this needs to be done only once and can be farmed out to a bank of such machines to reduce the computation time. With this precomputed CONUS domain, we gain the ability to quickly switch domains. See Levit et al. (2004) for how this capability is used to get real-time access to merged 3D radar data from anywhere in the country using a single workstation.

e. Time correction

Motion estimates obtained from the technique described in Lakshmanan et al. (2003b) are used to correct for time differences between the sensing of the same storms by different radars using Eq. (5). This dramatically improves the value of derived products computed from the 3D grid because, by correcting the locations of the storms, the areas sensed by different radars line up in space and time. Figure 2 demonstrates the significant difference between an uncorrected image and a time-corrected one.

The characteristics of the motion estimation technique impact the results in the merged grid. If the motion of a few storms is underestimated, then those storms will appear to jump whenever newer data are obtained, as the intelligent agents correct their positions. Because the motion estimation technique of Lakshmanan et al. (2003b) is biased toward estimating the movement of larger storms correctly, smaller cells and cells with erratic movements will tend to not provide smooth transitions. However, the transition in such cases will typically still be less than the transition that would have resulted if no advection correction had been applied.

Weighting individual observations by time [in addition to distance; see Eq. (7)] can have undesirable side effects if the radars are not calibrated identically. If one radar is “hotter” than all the others, then in those time frames where data from that radar are the latest available, the merged image will appear hotter. This problem affects radar data from the WSR-88D network because of large calibration differences between individual radars (Gourley et al. 2003) and is most evident when viewing time sequences of merged radar data.

In the absence of time weighting, however, the merged data will be a biased estimate in areas where the storms are exhibiting fast variations in intensity be-

TABLE 1. Timing statistics collected in real time for reading radar data and creating 3D merged grids of reflectivity.

Aspect	CPU time (s)	Clock time (s)	Scaling
Read radar scan	0.012 ± 0.001	0.254 ± 0.023	One elevation scan
Update grid with scan	0.075 ± 0.007	1.28 ± 0.135	One elevation scan
Create 3D grid	0.594 ± 0.127	9.179 ± 2.056	8-m voxels
Write 3D grid	1.257 ± 0.268	20.68 ± 4.62	8-m voxels

cause older data are still retained. The older data are advected to their correct positions, but no compensation is made for potential changes in intensity because the automated extraction of the growth/decay of the storm was found to possess very little skill (Lakshmanan et al. 2003b). Therefore, calibration differences still remain evident when looking at features that migrate from the domain of one radar to the domain of another, especially when these features are accumulated over time.

One solution to the problem of combining time weighting with calibration differences between radars is to retain time weighting but to apply calibration corrections to the data from individual radars. We intend to implement such a calibration correction in future research.

f. Timing

The time taken to read individual elevation scans, update the 3D grid by creating intelligent agents, and write the current state of the 3D grid periodically depends on a variety of factors. We carried out a test where the individual elevation scans and the output grid are both compressed (as will be the case in a networked environment, to conserve bandwidth), so that all of these timing data reflect the time taken to uncompress while reading and compress when writing. We carried out this test for a $650 \times 700 \times 18$ regional domain with convective activity in most parts of the region and using data from three WSR-88Ds. Different tasks scale to either the domain size or the number of radars. These are marked along with the timing information in Table 1. The test was carried out on the aforementioned Intel Xeon workstation.

Table 1 can be used to determine the computing power needed for different domain sizes and numbers of radars. On average, we require 0.5 seconds per elevation scan and 2.75 seconds per million voxels. It should be noted that merging radar data using the technique described in this paper is an “embarrassingly par-

allel” problem; that is, if multiple machines are put to work on the problem, each machine can concentrate on a subdomain and the output grids can be cheaply stitched together again. The subdomains do have to partially overlap to take into account advection effects.

Using the information in Table 1, we can estimate the computational requirements necessary to create a $5000 \times 3000 \times 20$ domain every 60 s in real time using information from 130 WSR-88Ds. Let us estimate that we will receive a new elevation scan from each of the radars every 30 s (in areas having significant weather, this will be around 20 s while in areas of little weather, the interval approaches 120 s). Thus, we will have 260 elevation scans per minute to read and update, which will take about 130 wall-clock seconds on our single workstation. Because our $5000 \times 3000 \times 20$ domain contains 300 million voxels, creating and writing out the grid will take 1125 more seconds. Our single workstation, if outfitted with enough computer memory, will be able to create and write 3D merged grids for this domain in 1250 s. To maintain our update interval of 60 s, we would require about 20 such workstations. The hardware requirements for several scenarios are provided in Table 2. Levit et al. (2004) used the regional configuration in their study.

5. Summary

In this paper, we have described a technique based on an intelligent agent formulation for taking the base radar data (reflectivity and radial velocity), and derived products, from multiple radars and combining them in real time into a rapidly updating 3D merged grid. We demonstrated that the intelligent formulation accounts for the varying radar beam geometry with range, vertical gaps between radar scans, the lack of time synchronization between radars, storm movement, varying beam resolutions between different types of radars, beam blockage due to terrain, differing radar calibration, and inaccurate time stamps on radar data.

The techniques described in this paper of merging moment data from individual radars have been tested in real time and on archived data cases in diverse storm regimes. They have been tested on different types of radars as well. For example, Fig. 3 shows the technique operating in a hurricane event. Figure 13 shows the technique operating on late-summer monsoon events in an area with terrain. Figures 2, 9, and 10 illustrate the use of this technique in convective situations, while Fig. 5 illustrates a stratiform one. One of the radars in Fig. 6 is a TDWR while the others are WSR-88Ds, while Fig. 7 shows simulated CASA radars.

With the continuing improvements in computer pro-

TABLE 2. Estimated hardware requirements for implementing this technique. The machine referred to here is a 1.8-GHz Intel Xeon processor with 2 GB of RAM.

	Update interval	Domain size	No. of radars	Domain resolution	No. of machines
CONUS	1 min	5000 km \times 3000 km \times 20 km	130	1 km \times 1 km \times 1 km	20
Regional	1 min	800 km \times 800 km \times 20 km	10	1 km \times 1 km \times 1 km	1
CONUS	5 min	5000 km \times 3000 km \times 20 km	130	1 km \times 1 km \times 1 km	4

processors, operating systems, and input–output performance, preliminary tests indicate that a 2-min CONUS 3D merger can be implemented using a single 3-GHz 64-bit processor with 8 GB of RAM. Thus, we plan to start utilizing this technique to merge scalar fields, both reflectivity and derived shear products, from all the WSR-88D radars in the continental United States. We now have the ability to ship the merged products in real time to Advanced Weather Interactive Processing System (AWIPS) and National Centers for Environmental Prediction AWIPS (N-AWIPS) workstations at weather forecast offices and national centers. Because these workstations will be unable to handle the data at its highest resolution and spatial extent, we may have to subsect the data before shipping it to operational centers. The merged products created using the technique described in this paper are now available in real time (online at <http://wdssw.nssl.noaa.gov/>).

Also note that the software implementing this technique of combining data from multiple radars is freely available online (at <http://www.wdssii.org/>) for research and academic use.

Acknowledgments. Funding for this research was provided under NOAA–OU Cooperative Agreement NA17RJ1227, National Science Foundation Grants ITR 0205628 and 0313747, and the National Severe Storms Laboratory. Part of the work was performed in collaboration with the CASA NSF Engineering Research Center (EEC-0313747). We thank Kiel Ortega for creating some of our data cases, Jian Zhang for her useful pointers at several stages of this research, Rodger Brown for guidance on the multi-Doppler velocity combination approach and the anonymous reviewers whose inputs improved the clarity of our paper. The statements, findings, conclusions, and recommendations are those of the authors and do not necessarily reflect the views of the National Severe Storms Laboratory, the National Science Foundation, or the U.S. Department of Commerce.

REFERENCES

Adrianto, I., T. M. Smith, K. A. Scharfenberg, and T. Trafalis, 2005: Evaluation of various algorithms and display concepts

for weather forecasting. Preprints, *21st Int. Conf. on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, San Diego, CA, Amer. Meteor. Soc., CD-ROM, 5.7.

Askelson, M., J. Aubagnac, and J. Straka, 2000: An adaptation of the Barnes filter applied to the objective analysis of radar data. *Mon. Wea. Rev.*, **128**, 3050–3082.

Beyer, W., Ed., 1987: *CRC Standard Math Tables*. 18th ed. CRC Press, 487 pp.

Biggerstaff, M. I., and R. M. May, 2005: A meteorological radar emulator for education and research. Preprints, *21st Int. Conf. on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, San Diego, CA, Amer. Meteor. Soc., CD-ROM, 5.6.

Brotzge, J., D. Westbrook, K. Brewster, K. Hondl, and M. Zink, 2005: The meteorological command and control structure of a dynamic, collaborative, automated radar network. Preprints, *21st Int. Conf. on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, San Diego, CA, Amer. Meteor. Soc., CD-ROM, 19.15.

Brown, R., and P. Hwang, 1997: *Introduction to Random Signals and Applied Kalman Filtering*. John Wiley and Sons, 480 pp.

Browning, K., and A. Wexler, 1968: The determination of kinematic properties of a wind field using Doppler radar. *J. Appl. Meteor.*, **7**, 105–113.

Burrows, M., and D. Wheeler, 1994: A block-sorting lossless data compression algorithm. Digital Equipment Corporation Tech. Rep. 124, Palo Alto, CA, 18 pp. [Available online at gatekeeper.dec.com/pub/DEC/SRC/research-reports/abstracts/src-rr-124.html.]

Charba, J., and F. Liang, 2005: Quality control of gridded national radar reflectivity data. Preprints, *21st Conf. on Weather Analysis and Forecasting/17th Conf. on Numerical Weather Prediction*, Washington, DC, Amer. Meteor. Soc., CD-ROM, 6A.5.

Doviak, R., and D. Zrnic, 1993: *Doppler Radar and Weather Observations*. 2d ed. Academic Press, 458 pp.

Droegemeier, K., and Coauthors, 2002: Project CRAFT: A test bed for demonstrating the real time acquisition and archival of WSR-88D level II data. Preprints, *18th Int. Conf. on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, Orlando, FL, Amer. Meteor. Soc., 136–139.

Foote, G., and P. S. duToit, 1969: Terminal velocity of raindrops aloft. *J. Appl. Meteor.*, **8**, 249–253.

Fulton, R., D. Breidenback, D. Miller, and T. O'Bannon, 1998: The WSR-88D rainfall algorithm. *Wea. Forecasting*, **13**, 377–395.

Gourley, J. J., B. Kaney, and R. A. Maddox, 2003: Evaluating the calibrations of radars: A software approach. Preprints, *31st Int. Conf. on Radar Meteorology*, Hyannis, MA, Amer. Meteor. Soc., CD-ROM, P3C.1.

- Greene, D. R., and R. A. Clark, 1972: Vertically integrated liquid water—A new analysis tool. *Mon. Wea. Rev.*, **100**, 548–552.
- 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. *Wea. Forecasting*, **13**, 263–276.
- Kalman, R., 1960: A new approach to linear filtering and prediction problems. *Trans. ASME J. Basic Eng.*, **82D**, 35–45.
- Kessinger, C. J., P. Ray, and C. E. Hane, 1987: The Oklahoma squall line of 19 May 1977. Part I: A multiple Doppler analysis of convective and stratiform structure. *J. Atmos. Sci.*, **44**, 2840–2864.
- , S. Ellis, and J. Van Andel, 2003: The radar echo classifier: A fuzzy logic algorithm for the WSR-88D. Preprints, *Third Conf. on Artificial Applications to the Environmental Sciences*, Long Beach, CA, Amer. Meteor. Soc., CD-ROM, P1.6.
- Lakshmanan, V., and M. Valente, 2004: Quality control of radar reflectivity data using satellite data and surface observations. Preprints, *20th Int. Conf. on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, Seattle, WA, Amer. Meteor. Soc., CD-ROM, 12.2.
- , K. Hondl, G. Stumpf, and T. Smith, 2003a: Quality control of weather radar data using texture features and a neural network. *Fifth Int. Conf. on Advances in Pattern Recognition*, Calcutta, India, IEEE–Indian Statistical Institute, 54–59.
- , R. Rabin, and V. DeBrunner, 2003b: Multiscale storm identification and forecast. *Atmos. Res.*, **67/68**, 367–380.
- Levit, J. J., V. Lakshmanan, K. L. Manross, and R. Schneider, 2004: Integration of the Warning Decision Support System-Integrated Information (WDSS-II) into the NOAA Storm Prediction Center. Preprints, *22d Conf. on Severe Local Storms*, Hyannis, MA, Amer. Meteor. Soc., CD-ROM, 8B4.
- Lhermitte, R., and D. Atlas, 1961: Precipitation motion by pulse Doppler. Preprints, *Ninth Conf. on Radar Meteorology*, Boston, MA, Amer. Meteor. Soc., 218–223.
- Liou, Y.-C., 2002: An explanation of the wind speed underestimation obtained from a least squares type single-Doppler radar velocity retrieval method. *J. Appl. Meteor.*, **41**, 811–823.
- Maes, P., 1995: Artificial life meets entertainment: Life-like autonomous agents. *Commun. ACM*, **38**, 108–114.
- McNellis, T., S. Katz, M. Campbell, and J. Melody, 2005: Recent advances in phased array radar for meteorological applications. Preprints, *21st Int. Conf. on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, San Diego, CA, Amer. Meteor. Soc., CD-ROM, 19.6.
- O'Bannon, T., 1997: Using a terrain-based hybrid scan to improve WSR-88D precipitation estimates. Preprints, *28th Conf. on Radar Meteorology*, Austin, TX, Amer. Meteor. Soc., 506–507.
- Rabin, R., and D. Zrnic, 1980: Subsynoptic-scale vertical wind revealed by dual Doppler-radar and VAD analysis. *J. Atmos. Sci.*, **37**, 644–654.
- Smith, D. C., A. Cypher, and J. Spohrer, 1994: KidSim: Programming agents without a programming language. *Commun. ACM*, **37**, 55–67.
- Smith, T. M., and K. L. Elmore, 2004: The use of radial velocity derivatives to diagnose rotation and divergence. Preprints, *11th Conf. on Aviation, Range, and Aerospace*, Hyannis, MA, Amer. Meteor. Soc., CD-ROM, P5.6.
- Stumpf, G. J., T. M. Smith, and J. Hocker, 2004: New hail diagnostic parameters derived by integrating multiple radars and multiple sensors. Preprints, *22d Conf. on Severe Local Storms*, Hyannis, MA, Amer. Meteor. Soc., CD-ROM, P7.8.
- , S. Smith, and K. Kelleher, 2005: Collaborative activities of the NWS MDL and NSSL to improve and develop new multiple-sensor severe weather warning guidance applications. Preprints, *21st Int. Conf. on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, San Diego, CA, Amer. Meteor. Soc., CD-ROM, P2.13.
- Trapp, R. J., and C. A. Doswell, 2000: Radar data objective analysis. *J. Atmos. Oceanic Technol.*, **17**, 105–120.
- Witt, A., M. Eilts, G. Stumpf, J. Johnson, E. Mitchell, and K. Thomas, 1998: An enhanced hail detection algorithm for the WSR-88D. *Wea. Forecasting*, **13**, 286–303.
- , R. A. Brown, and V. Lakshmanan, 2005: Real-time calculation of horizontal winds using multiple Doppler radars: A new WDSS-II module. Preprints, *32d Conf. on Radar Meteorology*, Albuquerque, NM, Amer. Meteor. Soc., CD-ROM, P8R7.
- Zhang, J., K. Howard, and J. J. Gourley, 2005: Constructing three-dimensional multiple-radar reflectivity mosaics: Examples of convective storms and stratiform rain echoes. *J. Atmos. Oceanic Technol.*, **22**, 30–42.