P5.1: 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.

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

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, refered 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 an 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 optimizations including the Flyweight pattern (Gamma et al. 1994) and reusing the results of previous computations are available. The merger process was tested on a Linux

Pentium-III desktop; it easily keeps up with a real-time feed from four 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.

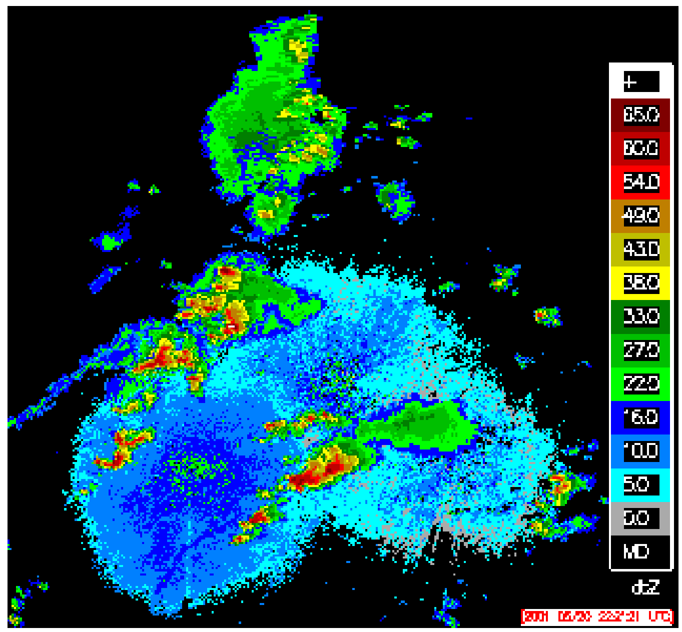

Figure 1: The lowest tilts from the three radars used to demonstrate the results shown here. Data were collected on May 20, 2001 from Oklahoma City, Hot Springs and Tulsa.

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

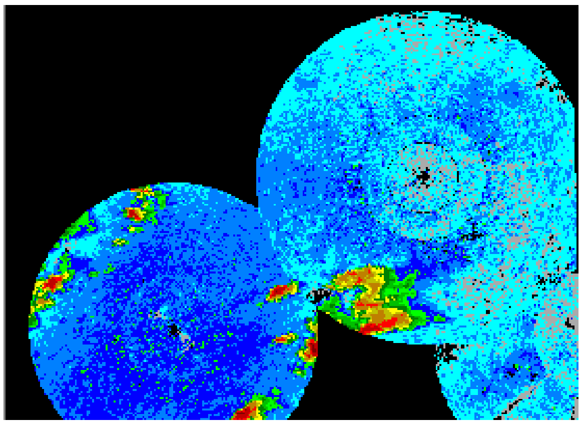

a

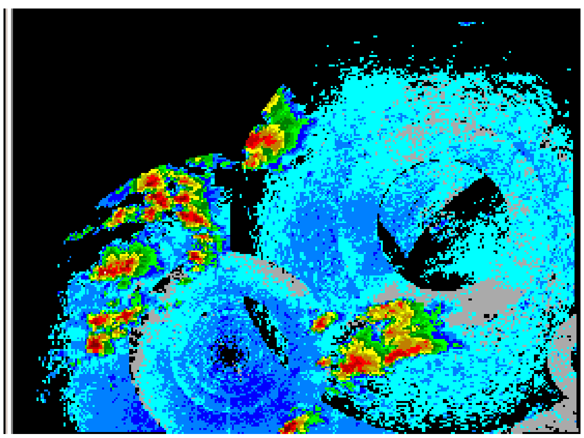

b

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.

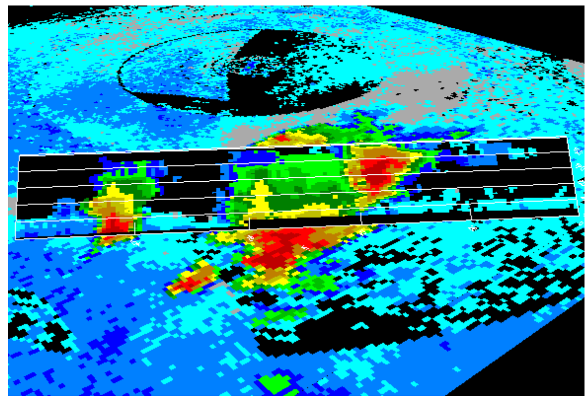

a

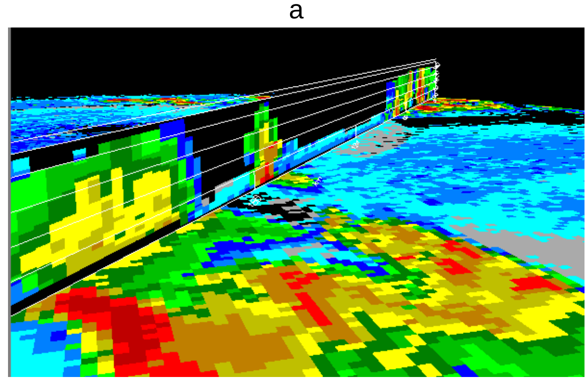

b

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

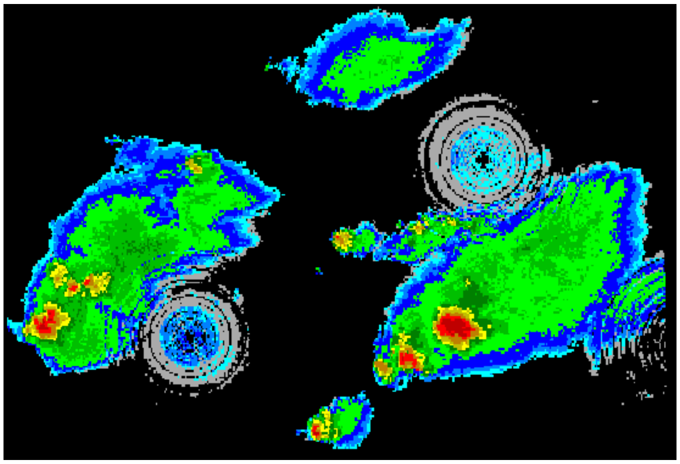

Figure 4: The maximum in the vertical direction at every grid point is an optional product that may be obtained from the 3D grid. Note that there is a new composite product formed every time the 3D grid is updated. This is an example of a multisensor time-update (or virtual volume) algorithm.

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.

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.