

Case Study: Visualization of Local Rainfall from Weather Radar Measurements

Thomas Gerstner¹, Dirk Meetschen², Susanne Crewell², Michael Griebel¹, Clemens Simmer²

Department for Applied Mathematics¹
University of Bonn
Wegelerstr. 6, D-53115 Bonn
{gerstner;griebel}@iam.uni-bonn.de

Meteorological Institute²
University of Bonn
Auf dem Hügel 20, D-53121 Bonn
{dmeetsch;screwel;csimmer}@uni-bonn.de

Abstract

Weather radars can measure the backscatter from rain drops in the atmosphere. A complete radar scan provides three-dimensional precipitation information. For the understanding of the underlying atmospheric processes the interactive visualization of these data sets is necessary. This task is, however, difficult due to the size and structure of the data and due to the necessity of putting the data into a context, e.g. by the display of the surrounding terrain. In this paper, a multiresolution approach for real-time simultaneous visualization of radar measurements together with the corresponding terrain data is illustrated.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling: Curve, Surface, Solid, and Object Representations J.2 [Computer Applications]: Physical Sciences and Engineering: Earth and Atmospheric Sciences

1. Introduction

Remote sensing of meteorological quantities is necessary for the understanding and modelling of atmospheric processes. An important example are weather radar devices which are able to measure the backscatter from atmospheric hydrometeors (e.g. rain drops, hail, or snow flakes). Such a radar antenna is typically mounted on a hill or a high building (see Figure 1) and is busy scanning its surrounding sky. The measured radar reflectivities are protocolled as scalar values in 3D space. Since the scans are repeated regularly this leads to very large time-dependent scalar fields which have to be stored and processed.

The radar reflectivity is proportional to the diameter of the backscattering particle (to the sixth power), but the distribution of the particles is unknown. Thus, the determination of the amount of rain which will ultimately arrive at the ground – which one is of course most interested in – is difficult. The three-dimensional structure contains information which can improve rainfall rate estimates which is very important for hydrological modelling and simulations.

For the understanding and conveyance of the data, interactive volume visualization tools^{5, 7, 16} within geographic infor-

mation systems^{9, 11} are required. However, due to the large size of the data it is difficult to achieve real-time visualization performance. The task is complicated by the necessity that the data has to be put in a context for better comprehension, e.g. by additionally rendering the local terrain^{1, 2, 8}. Thereby, the amount of data which has to be displayed is further increased and the simultaneous display of volume and surface data leads to algorithmic visualization problems.

In this paper we will address these problems by adaptive multiresolution algorithms. Multiresolution methods allow a fast coarse display of the data for overview images or interaction while higher resolution is used when zooming into the data or for detail images. The algorithms are adaptive in the sense that the resolution does not have to be uniform everywhere but may be variable in space. So for example in smooth areas a lower resolution can be sufficient to provide a good impression of the data while in areas with greater variance a higher resolution is used. The adaptive refinement is controlled by suitable error indicators which can be specified by the user.

A multiresolution hierarchy is constructed for both the volume and the surface data. Although these multiresolu-



Figure 1: A scenic view of the weather radar antenna on top of a high building in Bonn, Germany.

tion models are built independently, they are constructed in a matching way. In our case, the 3D model is based on recursive tetrahedral bisection and the 2D model is based on recursive triangle bisection. This way, in principle, visibility problems which result from the merging of the data sets might be solved as well. In our experience so far mutual visibility did not pose a real problem, though.

The employed volume visualization algorithm is based on the simultaneous extraction of several transparent isosurfaces. This choice is motivated by the experience that the display of isosurfaces appears to be more meaningful to people from the atmospheric sciences than direct volume rendering. As an added benefit, no special-purpose hardware (apart from a standard graphics card) has to be used.

The outline of this paper is as follows. In Section 2 we will give some technical details on the weather radar device and the derived data. Section 3 covers the preprocessing of the data. The multiresolution algorithm is explained in Section 4. In Section 5 we will illustrate the performance of the system with some examples. We conclude with a few remarks on further extensions in Section 6.

2. About the Weather Radar

The radar antenna[†] is mounted on top of a 7-storey residential building in Bonn (geographical position $70^{\circ}04' E$, $50^{\circ}43' N$) in a height of 98.5 meters over sea level. It is operated in a continuous mode (save maintenance). The antenna is able to rotate 360 degrees around its vertical axis (azimuth angle) and tilt up to 90 degrees with respect to the horizontal plane (elevation angle). For each combination of azimuth and elevation angle reflectivities along the shot ray are obtained. The measurement radius is 50 km.

A circular scan for each elevation angle forms an – albeit very flat – cone (see Figure 2). One complete 3D scan takes 5 minutes and is repeated every 30 minutes (in the meantime



Figure 2: A schematic view of the weather radar scan grid.

the radar is not idle but scans only the smallest elevation angle in a higher time resolution). The scanned volume is roughly $55,000 \text{ km}^3$ covering about $7,900 \text{ km}^2$.

The raw data sets have been obtained for a 1 degree resolution of the azimuth angle, for elevation angles of 1.5, 2.5, 3.5, 4.5, 6.0, 8.0, 10.0, and 12.0 degrees, and for 200 equidistant points along each ray. In summary, each time-slice of raw data consists of $360 \times 200 \times 9$ floating point values.

3. Preprocessing of the Raw Data

The raw radar data is prone to various sources of error and thus has to be filtered before further processing. In this section we will shortly illustrate the various necessary steps.

3.1. Clutter Filtering

Clutter are unwanted signals, e.g. resulting from reflections at non-meteorological obstacles such as mountains, buildings, and industrial plants (ground clutter), or at birds and airplanes (moving clutter). Such obstacles alter the signal not only at their location, but also behind them with respect to the position of the antenna.

The basic way to filter the ground clutter is to record the measured signals on a dry day in a so called cluttermap. The cluttermap will be subtracted from the operationally measured signals. Thereby, each measured reflectivity value is compared to its pendant of the cluttermap. Values with precipitation are reduced if they are supposed to be contaminated by clutter. Afterwards, a weak low pass filter is applied to the image to eliminate residual clutter signals due to apparent displacement of ground echoes caused by variations of the refractive index within the atmosphere¹³. An example for the behaviour of the clutter filtering algorithm for a single scan of the second lowest elevation angle is shown in Figure 3.

3.2. Attenuation Correction

When hydrometeors reflect the energy of the pulse sent by the antenna, they also weaken the incoming pulse for the following volumes. This is not considered in the radar equation but can be corrected successively for each measured ray⁶.

[†] type Selenia METEOR-200

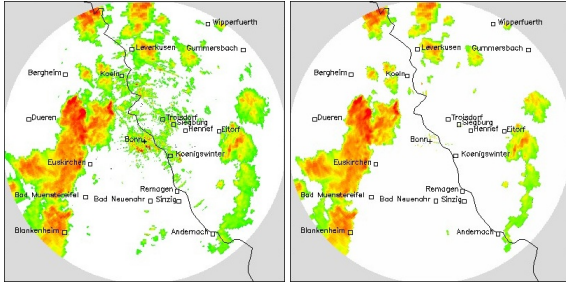


Figure 3: 2D plots of the rainfall data for the second lowest elevation angle: original measurements and after clutter filtering.

But, because attenuation affects the signal just in case of very strong precipitation and this kind of correction is very sensitive to measurements errors, it is employed in special cases only (as in the example of Section 5.1).

3.3. Z–R Relationship

The receiver of the radar system measures the power of the reflected pulse. To relieve the measurements from their dependency on range and the properties of the radar system, the data is usually converted into the so called radar reflectivity (factor) Z . The link between Z and the rain rate R is called Z–R relationship and depends on the unknown drop size distribution within the illuminated volume. Using approximations for the drop size distribution and the fall velocity formulas of the form $Z = aR^b$ can be derived. We have used $a = 300$ and $b = 1.5$ for our examples¹⁰.

3.4. Interpolation

As illustrated in Section 2, the raw data is given in a conical coordinate system. Although the data could be visualized in this coordinate system, this is impractical for two reasons. First, it would be hard to match the three–dimensional mesh with the two–dimensional terrain grid. This would be necessary for the solution of mutual visibility problems of the rainfall and terrain data. Second, the conical mesh is circular along the azimuth angle and thus not homeomorphic to a cubical mesh. This would require a special treatment of the circularity which is difficult.

In order to avoid these problems, we interpolated the data set onto a cubical grid with $400 \times 400 \times 66$ grid points using trilinear interpolation. In comparison to the conical grid, the resolution of the cubical grid is about a factor of two higher in lateral direction and a factor of seven higher in elevation direction in order to minimize the interpolation error. Though the interpolated data set is about 16 times larger than the raw data, it is much easier to handle.




Figure 4: The recursive bisection hierarchy in 2D and in 3D. The upper row shows the respective refinement rules, the lower row the initial grids and the first refinement step.

4. Multiresolution Visualization

The here used multiresolution visualization algorithms (both in 2D and 3D) have been presented in previous publications^{3,4}. At this point, we will just repeat the main steps in order to illustrate the special adaptations required for the application under consideration.

4.1. Multiresolution Hierarchy

Let us consider a hierarchy of triangular, respectively tetrahedral grids generated by recursive bisection. To this end, the midpoint of a predetermined edge is chosen as a refinement vertex and the triangle, respectively tetrahedron, is split in two by connecting the refinement vertex with its opposing vertices. Starting with initial meshes consisting of two triangles, respectively six tetrahedra, and recursively splitting all triangles or tetrahedra, nested hierarchical meshes are constructed (Figure 4).

An adaptive mesh can be extracted from such a hierarchy by selection and computation of a suited error indicator. Then, all triangles, respectively tetrahedra, whose associated error indicator is below a user–prescribed error threshold are visited. If this error indicator is saturated then the resulting adaptive meshes cannot contain so–called hanging nodes which would lead to holes during the visualization^{4,17}.

4.2. Terrain Visualization

For a two–dimensional scalar field, such as terrain, it is sufficient to draw and shade the triangles on the finest local resolution. An often applied extension is view–dependent refinement¹² which is usually used for low flyovers. Here, the error threshold is not uniformly distributed but increases with the distance to the viewer. For this application, view–dependent refinement was not necessary since, due to the nature of the data, a top or bird’s view is mostly used (compare Figure 6). For visualization purposes the triangles are color shaded according to elevation value using a simple geographical colormap ranging from green (low elevation) to white (high elevation).

4.3. Rainfield Visualization

Three-dimensional scalar fields require special volume visualization techniques such as volume rendering (e.g. using tetrahedral splats¹⁵) or isosurface extraction (e.g. using marching tetrahedra¹⁴). Isosurfaces can be efficiently extracted hierarchically if additionally bounds for the minimum and maximum data value inside each tetrahedron are available. The extraction algorithm then traverses the hierarchical mesh searching only through those regions where the isosurface is actually located (see Figure 7) and extracts the local isosurface using a simple lookup-table on the finest local resolution¹⁷. Multiple transparent isosurfaces can be rendered using alpha blending if the tetrahedra and isosurface components are sorted in visibility order³.

In our examples, we visualize the rainfall by 20 transparent isosurfaces with low opacity and color ranging from white (low rainfall) to blue (high rainfall). The isovalues are equidistant and the maximum isovalue is determined by the maximum rainfall in the data set (single event), respectively the maximum rainfall during the time interval (time series). The isosurface triangles are Gouraud shaded based on the gradients of the 3D data set.

4.4. Integration

In this application, the user focus is clearly on the rainfield. The terrain is mainly displayed for orientation and spatial analysis. Nevertheless, the terrain has to cover a far greater area than the extent of the radar on the ground in order to be able to also provide realistic side views of the rainfield. In our experience, the terrain should extend at least twice the rainfield size in all directions. This way, however, the two meshes cannot be constructed in a matching way in the sense that the z -projection of the 3D mesh will fall exactly onto the 2D mesh.

In principle, intervisibility between the two data sets is a problem due to the transparency of the isosurfaces which requires alpha blending. For example, if part of the rainfield is located in between two mountains, then the more distant mountain has to be drawn first, then the rainfield, and the closer mountain last (see Figure 5). This case can occur since the viewing direction does not have to be identical to the direction of the radar ray. However, it did not pose a real problem here since, due to the nature of the radar, the measured rainfield is usually located high above the ground (compare Figure 6d). The problem per se is interesting though and would require an interleaved traversal of the two multiresolution hierarchies which is not a trivial problem and subject to current research.

Due to these two problems we treat the rainfield and the terrain meshes independently from each other as two separate multiresolution hierarchies. In order to achieve sufficiently correct results the terrain is rendered first and the rainfield field second.




Figure 5: *The mutual occlusion problem: part of the rainfield is located in between two mountains. Rendering the opaque terrain first and then the transparent rainfield does not lead to correct results. In principle, an interleaved traversal of the terrain and rainfield meshes is required.*

However, the coupling of the two error thresholds has to be addressed. Although the two thresholds could be chosen independently, the user should not have to adjust both of them. We solved the problem here by providing the user with a global error ruler and a second ruler which controls the relation between the two errors. This way, the user will adjust the accuracy of the terrain in comparison to the rainfield once, and use a global error ruler to control fidelity. Latter threshold could be determined automatically based on the performance, but we did not feel it necessary here.

5. Examples

We will now illustrate the performance and results of the visualization tool with some examples. The first example shows a single rainfall event while the second example covers a time series. In all examples the rainfall data is displayed on top of a 50 m digital elevation model[‡] of roughly three times the extent of the measurement area. The height of the terrain data as well as the rainfield data is exaggerated by a factor of 5 in the images. Of course, the height exaggeration can be controlled by the user but this value proved to be the most useful.

5.1. Single Event

We will start with a highly localized and heavy rainfall event which occurred during May 3rd, 2001. This event is important since it caused severe damage in the region. The spatial extent of the event on the ground is roughly 100 km².

Figure 6 (a) depicts the location of the measurement area of the radar in northwestern Germany. In Figure 6 (b) this area (circular with 100 km diameter) as viewed from directly above is shown. In the center is the city of Bonn, the river Rhine runs from southeast (Koblenz) to northwest (Cologne). Note that this top view is already more illustrative to the meteorologist than single horizontal slices (as Figure 3). The 3D structure of the rainfall is really revealed through interactive rotation and movement through the data

[‡] courtesy of SFB350, University of Bonn

set. As an example, Figure 6 (c) shows a 70 degree angle view of the area. A closeup of the rainfall event as seen from a flat 87 degree angle facing northwest is shown in Figure 6 (d). Here, the typical formation of rain tubes can be seen. On the upper border aliasing artefacts which result from the employed interpolation algorithm (see Section 3.4) are visible.

In Figure 7 we show the multiresolution performance of the visualization system. Here we show the corresponding triangular and tetrahedral grids next to the terrain and rainfield. From left to right the global error threshold (see Section 4.4) was decreased by a factor of ten resulting in an increase of the number of triangles also by a factor of ten (in this example, the number of isosurface and terrain triangles were roughly the same). As can be seen, the isosurface extraction algorithm refines only in the vicinity of the isosurface. The terrain is also much coarser in smooth, less mountainous areas such as in the upper part of the images which substantially increases the overall performance. For the coarser images in addition the number of displayed isosurfaces were reduced to 5, respectively 10. This number was also coupled to the global error threshold.

The whole system is interactive for large error thresholds (as in Figure 7 left and middle) while higher resolution images (such as Figure 7 right) require a few seconds rendering time on an Intel Pentium III with a standard nVidia graphics card running Linux. For the rightmost image several million triangles were rendered.

5.2. Time Series

As a second example will serve a heavy convective rainfall event which has occurred on July 20th, 2001. In Figure 8 we show three images at 11:52, 12:22, and 12:52. The wind direction is from the southwest and thus the rain front moves roughly from left to right in the images. Well visible here is the deficiency that the radar does not cover the area directly above the antenna (which is located in the middle of the images). This problem could be solved by increasing the maximum elevation angle, or using several radar devices with overlapping measurement areas.

For larger time series the rainfall data should be stored in compressed form (not only on disk but also in memory). Since large parts of the rainfield are typically empty (zero rainfall) this can be achieved by the construction of adaptive tetrahedral meshes which capture the rain data (such as in Figure 7 lower right). Then, only the data values in the adaptive mesh together with the corresponding tree structure (which requires only a few bits per vertex) have to be stored. However, this was not necessary for the small time series in the example.

6. Concluding Remarks

We have illustrated a meteorological application of volume visualization for rainfall measurements from weather radar.

Here multiresolution algorithms were necessary in order to be able to handle the large amounts of data in an interactive visualization environment.

We are currently working on further improvements of the visualization system such as the inclusion of symbols and rainfall animation. Certainly a great improvement would be a multiresolution hierarchy on the native (conical) coordinate system of the radar measurements, because this way interpolation errors and aliasing artefacts can be avoided. This, however, will require special visualization algorithms for cyclic meshes, and visibility problems will be still harder to solve. Also, 3D data compression schemes which allow the visualization based on compressed rainfall data will be an interesting research direction.

Acknowledgement

The authors would like to thank SFB 350 "Interaction between and modeling of continental geosystems" of the German research foundation (DFG) for previous financial support.

References

1. P. Chen. Climate and Weather Simulations and Data Visualization using a Supercomputer, Workstations, and Microcomputers. In *Visual Data Exploration and Analysis III (Proc. SPIE '95)*, volume 2656, pages 254–264. SPIE, 1996.
2. J. Doick and A. Holt. The Full Volume Simultaneous Visualization of Weather Radar Echoes with Digital Terrain Map Data. In R. Earnshaw, J. Vince, and H. Jones, editors, *Visualization and Modeling*, pages 73–86. Academic Press, 1997.
3. T. Gerstner. Fast Multiresolution Extraction of Multiple Transparent Isosurfaces. In D. Ebert, J. Favre, and R. Peikert, editors, *Data Visualization '01*, pages 35–44. Springer, 2001.
4. T. Gerstner, M. Rumpf, and U. Weikard. Error Indicators for Multilevel Visualization and Computing on Nested Grids. *Computers & Graphics*, 24(3):363–373, 2000.
5. H. Haase, M. Bock, E. Hergenröther, C. Knöpfle, H.-J. Koppert, F. Schröder, A. Trembilski, and J. Weidenhausen. Meteorology Meets Computer Graphics – A Look at a Wide Range of Weather Visualizations for Diverse Audiences. *Computers & Graphics*, 24(3):391–397, 2000.
6. S. Hacker. Probleme mit der Dämpfungskorrektur von Radardaten. Master's thesis, Meteorological Institute, University of Bonn, 1996.
7. W. Hibbard, J. Anderson, I. Foster, B. Paul, R. Jacob, C. Schafer, and M. Tyree. Exploring coupled

- Atmosphere–Ocean Models using Vis5D. *International Journal of Supercomputing Applications and High Performance Computing*, 10(2):211–222, 1996.
8. C. James, S. Brodzik, H. Edmon, R. Houze Jr, and S. Yuter. Radar Data Processing and Visualization over Complex Terrain. *Weather and Forecasting*, 15:327–338, 2000.
 9. T. Jiang, W. Ribarsky, T. Wasilewski, N. Faust, B. Hanigan, and M. Parry. Acquisition and Display of Real–Time Atmospheric Data on Terrain. In D. Ebert, J. Favre, and R. Peikert, editors, *Data Visualization '01*, pages 15–24. Springer, 2001.
 10. J. Joss and A. Waldvogel. A Method to Improve the Accuracy of Radar-Measured Amounts of Precipitation. In *Proc. 14th International Conference on Radar Meteorology*, pages 237–238. American Meteorological Society, 1970.
 11. D. Koller, P. Lindstrom, W. Ribarsky, L. Hodges, N. Faust, and G. Turner. Virtual GIS: A Real–Time 3D Geographic Information System. In *Proceedings IEEE Visualization '95*, pages 94–100. IEEE Computer Society Press, 1995.
 12. P. Lindstrom, D. Koller, W. Ribarsky, L. Hodges, N. Faust, and G. Turner. Real–Time, Continuous Level of Detail Rendering of Height Fields. *Computer Graphics (SIGGRAPH '96 Proceedings)*, pages 109–118, 1996.
 13. D. Meetschen. Erkennung, Nutzung und Entfernung von Clutter zur Verbesserung der Niederschlagsmessung mit dem Bonner Radar. Master's thesis, Meteorological Institute, University of Bonn, 1999.
 14. B. Payne and A. Toga. Surface Mapping Brain Function on 3D Models. *IEEE Computer Graphics and Applications*, 10(5):33–41, 1990.
 15. P. Shirley and A. Tuchman. A Polygonal Approximation to Direct Scalar Volume Rendering. *ACM Computer Graphics (Proc. of Volume Visualization '90)*, 24(5):63–70, 1990.
 16. M. Toussaint, M. Malkomes, M. Hagen, H. Höller, and P. Meischner. A Real–Time Visualization, Analysis and Management Toolkit for Multi–Parameter, Multi–Static Weather Radar Data. In *Proc. 30th International Conference on Radar Meteorology*, pages 56–57. American Meteorological Society, 2001.
 17. Y. Zhou, B. Chen, and A. Kaufman. Multiresolution Tetrahedral Framework for Visualizing Volume Data. In *Proc. IEEE Visualization '97*, pages 135–142. IEEE Computer Society Press, 1997.




Figure 6: Four images visualizing a heavy local rainfall event during May 3rd, 2001 in northwestern Germany. Displayed is the terrain and 20 transparent isosurfaces with low opacity and color ranging from white to blue. All heights are exaggerated by a factor of 5.




Figure 7: The multiresolution approach for the terrain as well as for the isosurfaces allows scalable performance for the interaction with the data. The images show terrain and isosurface renderings as well as the corresponding 2D and 3D grids for varying error thresholds. The number of triangles and tetrahedra increase by a factor of 10 in between the images.




Figure 8: Three snapshots of a heavy convective rainfall event on July 20th, 2001 (moving from left to right). The time between the images is 30 minutes. Note that directly above the radar antenna (in the middle of the images) no rainfall information is available.