

IBM Research Report

Web-Based Dissemination and Visualization of Mesoscale Weather Models for Business Operations

Lloyd A. Treinish

IBM Research Division

Thomas J. Watson Research Center

P.O. Box 218

Yorktown Heights, NY 10598



Research Division

Almaden - Austin - Beijing - Delhi - Haifa - India - T. J. Watson - Tokyo - Zurich

WEB-BASED DISSEMINATION AND VISUALIZATION OF MESOSCALE WEATHER MODELS FOR BUSINESS OPERATIONS

Lloyd A. Treinish*

IBM Thomas J. Watson Research Center, Yorktown Heights, NY

1. INTRODUCTION

Visualization is critical to the effective analysis, dissemination and assessment of data generated by numerical weather prediction. In that regard, consider two aspects of our previous work. First is the need to develop appropriate mapping of user goals to the design of pictorial content by considering both the underlying data characteristics and the perception of the visualization (Treinish, 2001). The second is the adaptation of these ideas from workstation or PC/game-class three-dimensional graphics systems with sufficient bandwidth for timely access to the model data to remote access via the world-wide-web (Treinish, 2002b). In this situation, the limitation in bandwidth is the primary bottleneck since desktop systems can support interactive visualization of typical model data.

2. TECHNIQUES FOR THE DISSEMINATION OF VISUALIZATIONS

Popular approaches to web-based dissemination of weather forecasting data are problematic as they impose enough compromises in time for access, display fidelity or interactivity to minimize effectiveness. Some utilize a standard workstation-based application after data have been lossy compressed for faster transmission (e.g., Hibbard, 1998). But such data reduction eliminates critical information from weather models operating at cloud-scale resolution. Others utilize web browsers by providing static images or flip-book animations, often presented via Javascript-based players or browser plug-ins (e.g., Wolfinbarger, 2002). Two additional problems are raised by this approach when used alone. The first is the need to heavily sample the images to reduce their size and number to be downloaded. This creates a gross mismatch to the time step used in the simulation, resulting in available images that may miss important results.

ENVIRONMENT	CONTENT	TIME STEP SIZE	INTERACTION	FIDELITY
<i>3d workstation</i>	data	5 - 10 MB	full 3d	full
<i>VRML/MPEG-4</i>	geometry	0.5 - 2 MB	full 3d	reduced
<i>IBR</i>	bitmap	100 - 200 KB	limited 3d	reduced
<i>image player</i>	bitmap	50 - 200 KB	2d	high
<i>video player</i>	bitmap	10 - 50 KB	2d	reduced

Table I. Trade-Offs in Dissemination Methods.

The second is the tendency to provide visualizations of most of the variables derived from the model. This can create a situation where users may have difficulty finding images of relevance. In addition, many of these visualizations may never be used. In some cases, the latter problem has been partially addressed by

* Corresponding author address: Lloyd A. Treinish, IBM T. J. Watson Research Center, P. O. Box 218, Yorktown Heights, NY 10598, lloyd@us.ibm.com, <http://www.research.ibm.com/people/l/treinish>

enabling products to be generated on demand, which has not been considered herein. In this effort, it has been assumed that there are a large number of clients using a relatively limited web server. Hence, the focus is on richer static content that is regularly refreshed.

To summarize some potential methods, consider Table I. Each row corresponds to a different mechanism to provide visualizations on the web. All of the environments shown except the first utilize a reduction in quality. The goal for each is to present an "interactive" visualization of three-dimensional time-varying data.

To put this table in context, consider the problem of disseminating visual browse products (i.e., as described in Treinish, 2001), where the data per time step is relatively modest with sequences of a few hundred time steps. Each approach delivers source material for interaction via different content. The top row is for a traditional interaction using OpenGL hardware acceleration. While high fidelity is inherent in the visualization, a fast network is required.

Typical realizations are composed of several polygonal surfaces, which are generated independently and combined at render time, but may result in a few hundred thousand triangles -- too large to download. To address this problem, each individual component is separately simplified to reduce the bulk by an order of magnitude for storage as VRML (second row). Another approach is available within MPEG-4 utilizing geometric compression to achieve roughly similar reduction. In both cases, it is insufficient for hundreds of time steps.

An alternative is through image-based rendering (IBR) represented by the third row. At a cost of less three-dimensional flexibility, the expense per time step is reduced by another order of magnitude. Some of that flexibility can be regained through hybrid image and geometry rendering and scaled to larger data per time step. However, the reduction in size is still insufficient.

The fourth and fifth rows represent a more traditional method of delivering sequences as bitmaps. The former is via flip-book animation, where each image frame is independently lossy compressed. Although image quality can be high, the cost per frame is also high. The fifth row implies a compressed video or animation stream that utilizes interframe similarity to reduce overall volume (e.g., MPEG-1). At resolutions sufficiently high to capture the detail in this type of visualization, the cost for a sequence at best is acceptable. Both of these cases come with a significant trade-off, no three-dimensional interactivity.

3. APPROACH

To begin to address these problems, visualizations are presented on a web page as a meta-representation of the model output and serve as an index to simplify finding other visualizations of relevance. To provide consis-

tency with extant interactive products and to leverage their cost of development, the aforementioned applications are adapted to automatically populate a web site with images and interactions for an operational weather forecasting system. Previously, we discussed specific meta-representations such as an interactive 3d image spreadsheet and utilization of specialized MPEG-compressed video sequences to enable products with time sampling consistent with model time steps (i.e., Treinish, 2002b). These efforts have been extended in two ways. First, they have been adapted to address the coupling of specific business processes and mesoscale simulations (i.e., Treinish, 2002a). The other extension is a response to the artifacts introduced by significant data sampling imposed by the bandwidth constraints. The implementation is consistent with the earlier work, but enhanced to incorporate the extensions discussed herein.

To enable effective assessment and appropriate decisions, focused visualizations are designed to integrate business and weather model data, yet still be driven by user goals. Thus, the resultant visualizations may not show forecasts of weather phenomena directly but the derived properties, which are influenced by weather, and are of direct relevance to the decision maker or industry specialist. Such focused visualizations may involve two-dimensional or three-dimensional strategies (Treinish, 2002a).

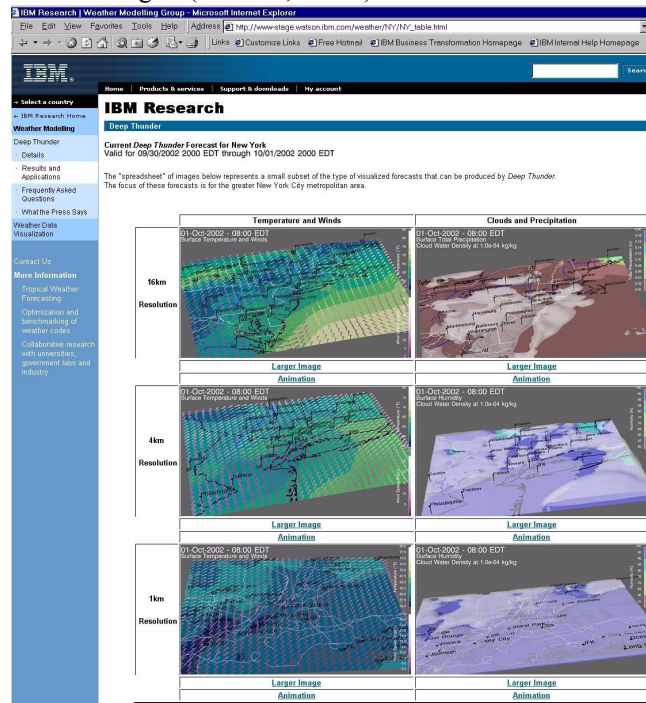


Figure 1. Web-Based Three-Dimensional Image Spreadsheet for Browse Products.

The resolution of the visualization must match that of the scale of the model to build usable products that are perceptually and scientifically coherent. The choice of realization geometry is also affected by the resolution of the data so that perceptual artifacts do not dominate the presentation, especially in animation (Treinish, 2001). Web-based dissemination exacerbates this situation. Therefore, abstraction and compression are introduced to capture sufficient content and reduce

underlying bulk, respectively, while utilizing the methods outlined in Table I.

These ideas are integrated into an operational mesoscale numerical weather prediction system dubbed “*Deep Thunder*” (Treinish et al, 2003). Currently, two 24-hour forecasts are produced each day on a 3-way nested configuration of 62 x 62 x 31 at 16, 4 and 1 km resolution focused on New York City. (See <http://www.research.ibm.com/weather/NY>)

4. BROWSE PRODUCTS

First, a set of visualizations are presented as an interactive, three-dimensional spreadsheet as discussed in Treinish, 2002b, where the spreadsheet becomes an abstraction of a model run. Thus, the rows and columns are organized at a high level (e.g., meteorological characteristics vs. model features) to simplify finding relevant visualizations. An example from the *Deep Thunder* web site is shown in Figure 1. These visualizations then become a meta-representation of the model output and serve as an index for more visualizations and interactions.

Each statically-defined cell in the spreadsheet shown in Figure 1 contains a three-dimensional scene generated from one time step of the model output. The left column focuses on atmospheric motion and dynamics by illustrating surface temperature and winds. Each image contains a shaded terrain surface that is colored by contour bands of temperature (°F.). The terrain is overlaid with maps of coastlines, county and state boundaries and rivers. Individual landmarks and cities are shown by name with the predicted temperature for that time step. The map is also overlaid with arrows for forecasted winds indicating speed by color and direction by orientation.

The right column focuses on moisture by showing precipitation at the surface and cloud properties. Each image contains a terrain map in a three-dimensional scene with predicted clouds. The clouds are shown as a translucent white isosurface derived from a threshold of total cloud water density of 10^{-4} kg water/kg air. If the model predicts severe weather, such as convective activity that could lead to the formation of thunderstorms, then a translucent cyan surface of predicted reflectivities may be visible within the clouds. The region within this surface corresponds to where precipitation is forming. The local terrain is a shaded surface that is colored by contour bands of total precipitation (as rain in inches), following the scale to the upper right, where brown implies no rain (dry) and heavier rainfall is in darker shades of blue. If the model predicts no precipitation then a similar visualization of humidity will be shown instead. If areas where precipitation is forecasted are sufficiently cold, then they may be marked to indicate snow.

Each cell also supports limited interaction. Via the mouse, one can appear to navigate inside the presented scene by changing a pseudo-camera view. This effect is provided via a Java applet which presents one of nine distinct images with a fixed angular separation (10° in this case) depending upon the mouse position. The result is context for the product in order to facilitate the selection of alternative visualizations. Currently, only

two choices are available. One is to examine a similar visualization, but with images of higher-resolution and greater fidelity.

The other choice addresses the time-sampling problem. It provides an animation with frames every 10 minutes of forecast time presented as an MPEG-1-compressed video sequence viewable on a web page via a Java-based player or through a plug-in. To preserve the fidelity of the animation and to keep the size to only a few MB, source animations are generated at 720p resolution and then interpolated to one-fourth the number of pixels prior to MPEG encoding. An example is shown in Figure 2.

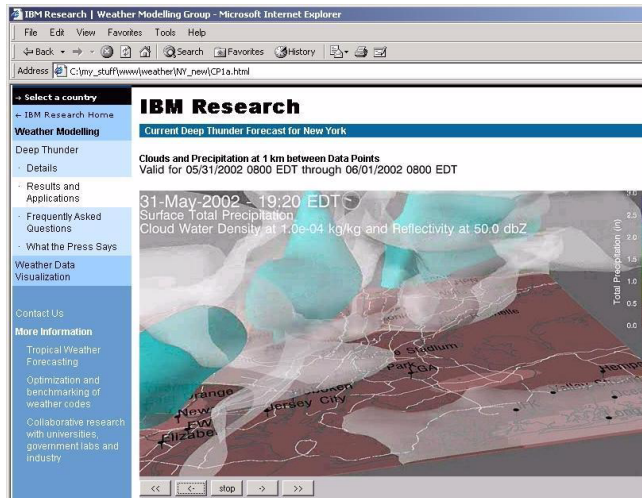


Figure 2. MPEG-1 Animation within a Web Page.

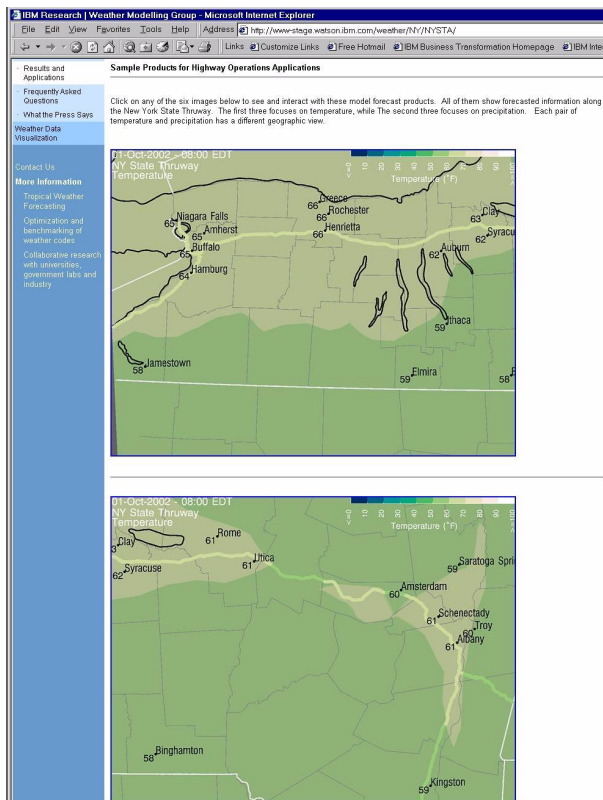


Figure 3. Custom Weather Map Selection Page.

5. EXTENSIONS OF BROWSE PRODUCTS

The techniques discussed above have been used to incorporate several additional products into the operational web site. The first provides very focused visualizations by both geography and application. A user has a limited choice of specific two-dimensional map products. A sample is shown in Figure 3, which illustrate forecasted temperatures as fixed colored contour bands for various parts of New York State. This type of map also serve as an index to further visualizations, for example an animation in Figure 4. Unlike Figure 2, the animation is via a set of individual JPEG images presented with a Javascript player. This more traditional approach is utilized given that there is one hour of forecast time between each frame, and greater temporal resolution is available in the browse product animations to aid in feature identification. The figure illustrates temperature predictions for Westchester County, NY.

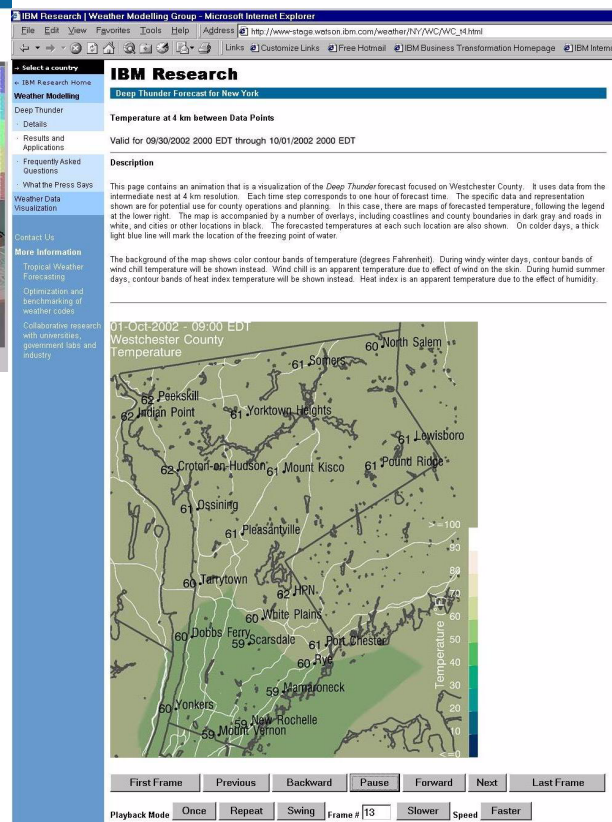


Figure 4. Animated Temperature Map.

Another choice is shown in Figure 5, which contains forecasted variables as a function of time interpolated to a specific location (Yankee Stadium). It consists of four panels showing surface data and supplemented with two panels to illustrate upper air winds. The plots on the left each show two variables while the rest each show one. The top left plot presents temperature (blue) and pressure (red). The middle left panel shows humidity (blue) and total precipitation (red). The top right plot illustrates forecasted winds -- speed (blue) and direction (red). The wind direction is shown via the arrows that are attached to the wind speed plot. The arrows indicate the predicted direction to which the wind is going. The bottom right plot is a colored contour map of forecasted

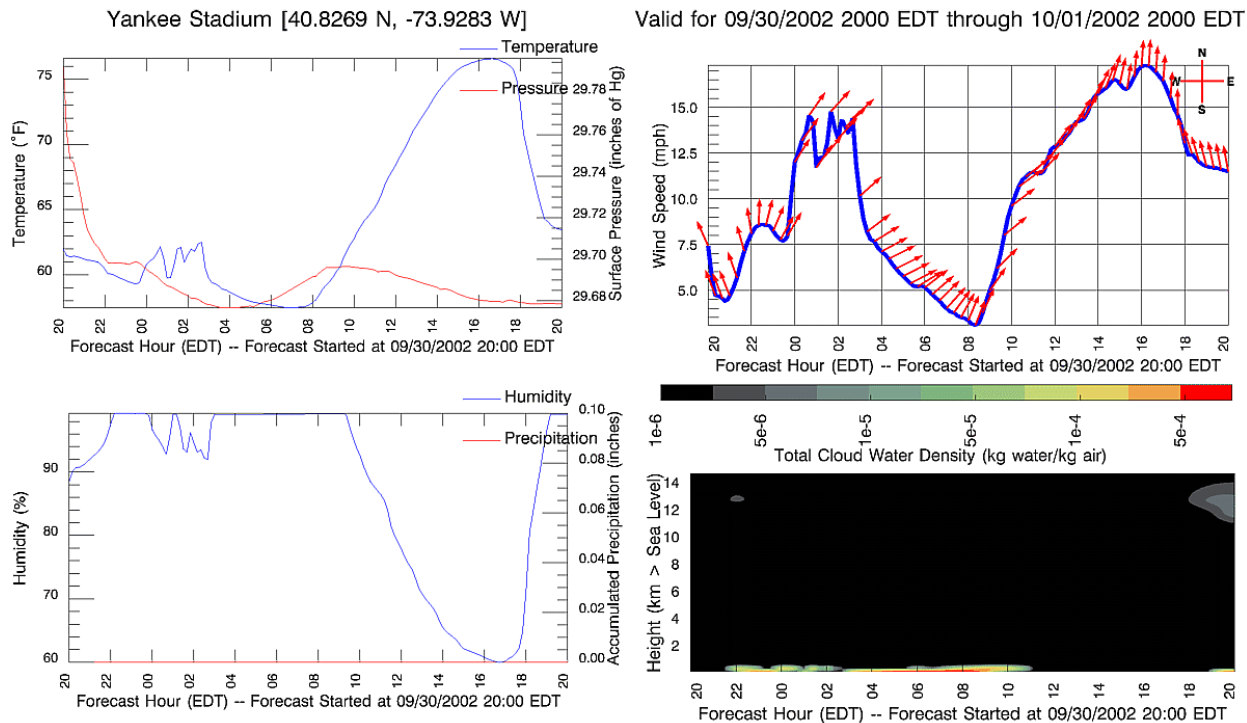


Figure 5. Location-Specific Forecast Products.

total (water and ice) cloud water density as a function of elevation and time. This "cross-sectional" slice can provide information related to storms, fog, visibility, etc. predicted at this location. Portions of the plot in black imply time or elevations where there are little or no clouds. Areas in yellow, orange and red imply when and where the relatively densest clouds are forecasted, following the color legend above the panel.

Another application of the techniques used in the web-based spreadsheet is for qualitative verification of model forecasts using remotely-sensed observations. Figure 6 illustrates an animation of 1-km-per-pixel-resolution GOES-8 visible imagery for the forecasting domain currently used for *Deep Thunder*. The 4 km and 1 km nests are overlaid as well as coastline, river and political boundary maps. The raw imagery are processed to enhance underlying detail prior to reprojection, pseudo-color mapping, rendering and MPEG-1 encoding. A similar technique is also applied to local radar observations of composite reflectivity and total (storm) precipitation.

6. ANALYSIS PRODUCTS

The ideas used for browse products have also been applied in a limited fashion to the generation of images oriented toward analysis. The first case is a single column of interactive cells with an aviation product, corresponding to the three modelling nests, from top to bottom, respectively. A user can select a single nest to display an animation using the technique employed in Figure 1. An example is shown in Figure 7.

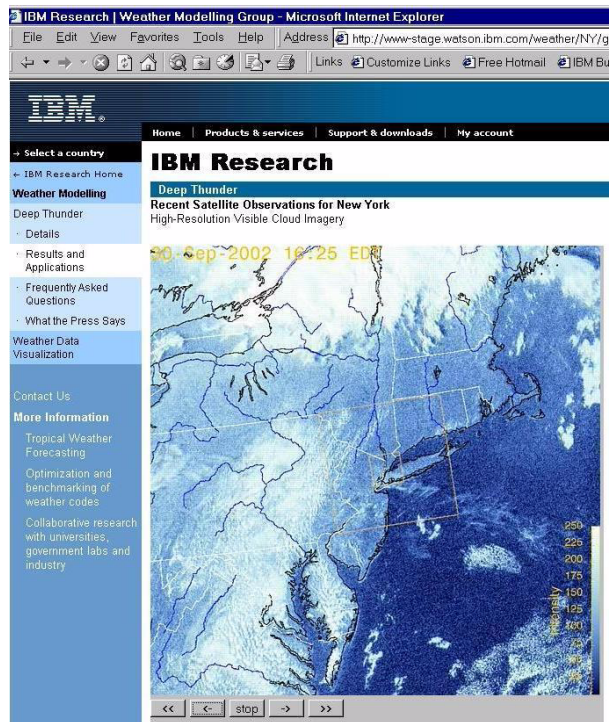


Figure 6. MPEG-1 GOES-8 Web-Based Animation.

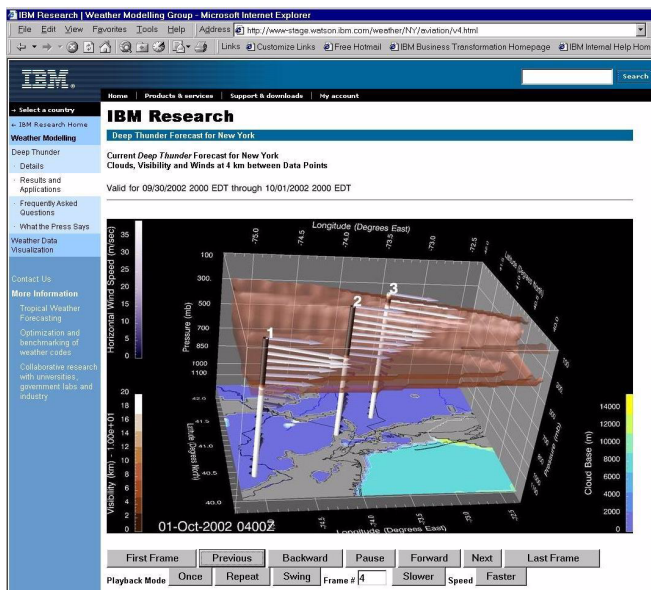


Figure 7. Animation of Cloud/Visibility Products.



Figure 8. Animation of Atmospheric Stability.

The image contains a brown, translucent, three-dimensional surface shown in vertical pressure coordinates, which corresponds to a boundary where the derived visibility is 10 km. This visibility is based upon extinction properties of cloud water, ice and precipitation, which is determined from the modelled upper air. Thus, the volume inside the surface represents relatively clear air, that is, visibility over 10 km. If no surface is visible then there are no clouds predicted at that time step, and thus, the visibility is high. At the bottom of the scene is a set of colored contours, typically in incre-

ments of 2 km, corresponding to the height of the forecasted cloud base. Areas in gray imply no cloud data. The cloud base contours are overlaid with maps of coastline and state boundaries and rivers. The volume is marked at the locations of major airports with set of poles color contoured by the derived visibility. At each of 21 pressure levels, the horizontal wind is shown via arrows, colored and sized by speed.

The second case also utilizes a single column with a product oriented toward atmospheric stability, especially for indicating the potential for severe weather. As with the previous example, an animation of one nest is shown, in Figure 8. A colored surface is presented, where the color corresponds to K Index, following the top legend. The surface is deformed linearly by Lifted Index and overlaid with a set of contour lines of forecasted vertical wind speed using the second color legend. Significant updrafts (green to yellow contours) in blue "valleys" on the surface would imply regions of convective activity. The surface is also overlaid with maps of coastline and state boundaries and rivers.

The techniques illustrated in Figure 7 have also been used to present results from the Eta synoptic-scale model on the same web site. An example is shown in Figure 9. The visualization shows a three-dimensional scene in vertical pressure coordinates in a Lambert Conformal projection covering North America. Several variables are shown. Near the top there is a translucent orange surface, which is an isosurface of horizontal wind speed at 30 meters/second. Its shape and movement is indicative of the jet stream. Closer to the bottom are white translucent surfaces representing "clouds" as an isosurface of relative humidity at 90%. At the bottom is a colored surface, whose height corresponds to the forecasted surface pressure. The coloring are contour bands of surface temperature using the legend at the lower left. The surface is overlaid with a map of political boundaries and coastlines in black.

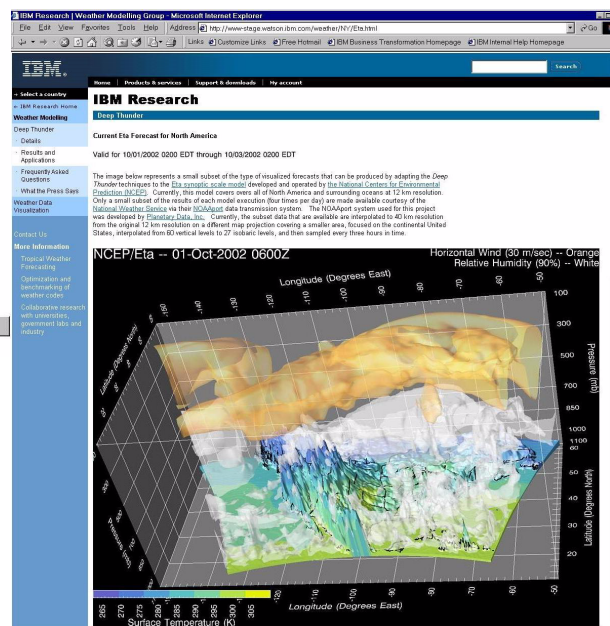


Figure 9. Web-Based Analysis Visualization Techniques Applied to Eta-212 Grids.

7. EXTREME COMPRESSION

Another interaction currently under development is applied to the data-sampling problem. It employs a specialized “*extreme*” level of compression based upon task-specific abstractions of underlying components of a visualization scene composed of several simpler geometric forms that are used multiple times such as in the browse products shown in Figures 1 and 2. They each can be represented by higher-order focused descriptions of the geometry. Hence, renderings of geometry are distributed instead of geometry itself. A typical Mpixel image of this type can be adequately described with only a few KB of data enabling animation to be transmitted and decoded cheaply at a client for interactive viewing.

As an example, consider a reconstructed image of a test data set including clouds, contours and arrows as shown in Figure 10. It only contains a broad variety of shapes and sizes of objects, as well as showing the translucency of those objects, as opposed to have been derived from actual weather model results. This reconstruction in its original size of 1018 x 974 pixels is stored in a file that is only 6 KB in size. Thus, animation sequences of a few hundred frames become very inexpensive to store and thus, transmit (i.e., one order of magnitude smaller than MPEG and two orders of magnitude for JPEG sequences). Since the description is geometric, it provides the potential for limited direct manipulation like the aforementioned spreadsheet cells but unlike a pure animation sequence.




Figure 10. Reconstructed test image based upon “*extreme*” compression.

8. DISCUSSION

These capabilities are used to disseminate results of model-based 24-hour weather predictions produced twice daily by the *Deep Thunder* system via a web site. They include several thousand images per day conceptually similar to the examples shown herein. They are rapidly generated and organized in the hierarchical fashion discussed earlier.

The users accessing the *Deep Thunder* web site range from non-specialists with casual interest for precision forecasts to individuals focused on decision support and planning in transportation, emergency management and energy applications, some of whom are meteorologists. Their diverse needs and successful utilization via a small web server has shown the viability of this approach.

However, feedback from these users has indicated the need for further development. In particular, the spreadsheet established interest in the techniques, but it was insufficient for analysis tasks due to the lack of query or transformation capabilities. This suggests three types of enhancement. The first is to incorporate additional, even very simple, products that could be selected from each cell. The next would be more rows and columns to compensate for the lack of further choices. Finally, the need to have greater interactivity beyond the geometric quick-look technique, which could be provided within the cell or as a selection for each cell. Development of additional visualization content continues to follow the principles of task-specific design to improve effectiveness while minimizing development cost (Treinish, 2001).

9. CONCLUSIONS AND FUTURE WORK

This approach to web-based visualization has shown promise in an operational environment. The spreadsheet/abstraction notion can be further expanded to serve as an index for more traditional web-based visualizations as well as visualizations that can be developed on technologies becoming available to support entertainment and e-business applications (e.g., MPEG-4). In addition, the spreadsheet itself can be enhanced to incorporate other techniques. Additional visualization methods focused on the weather-sensitive applications will be implemented. In addition, the techniques will be used to implement representations of forecast verification statistics. Finally, applications of the *extreme* compression will be incorporated into the operational environment.

10. ACKNOWLEDGEMENTS

The author wishes to acknowledge the contributions to this work by Craig Tashman. He is an undergraduate physics student at Pace University, who has an on-going internship at IBM Thomas J. Watson Research Center. The bulk of his time to date has been spent developing new techniques of graphics compression for the dissemination of model visualizations.

This work is supported by the Mathematical Sciences Department at the IBM Thomas J. Watson Research Center.

11. REFERENCES

- Hibbard, W. *Vis5D Version 5.2*. <ftp://www.ssec.wisc.edu/pub/vis5d-5.2/README>.
- Praino, A. P., L. A. Treinish and Z. D. Christidis. *Evaluation of an Operational Mesoscale Numerical Weather Prediction System in the Northeast U.S.* To be published in **Proceedings of the Nineteenth International Conference on Interactive Infor-**

mation and Processing Systems for Meteorology, Oceanography and Hydrology, February 2003, Long Beach, CA.

Treinish, L. *Coupling of Mesoscale Weather Models to Business Operations Utilizing Visual Data Fusion*. **Proceedings of the Third Symposium on Environmental Applications**, January 2002, Orlando, FL, pp. 94-101.

Treinish, L. *How Can We Build More Effective Weather Visualizations?* **Proceedings of the Eighth ECMWF Workshop on Meteorological Operational Systems**, November 2001, Reading, England, pp. 90-99.

Treinish, L. *Interactive, Web-Based Three-Dimensional Visualizations of Operational Mesoscale Weather Models*. **Proceedings of the Eighteenth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, January 2002, Orlando, FL, pp. J159-161.

Treinish, L., A. P. Praino and Z. D. Christidis. *Implementation of Mesoscale Numerical Weather Prediction for Weather-Sensitive Business Operations*. To be published in **Proceedings of the Nineteenth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, February 2003, Long Beach, CA.

Wolfenbarger, J.M., J.R. Greenfield, T.B. Stanley, and R.A. Young. 2002. *WeatherScope: Interactive Software for Visualizing Web-Based Meteorological Data Sets*. **Proceedings of the Eighteenth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**. January 2002, Orlando, FL, pp. J169-J170.