

CUSTOMIZATION OF A MESOSCALE NUMERICAL WEATHER PREDICTION SYSTEM FOR TRANSPORTATION APPLICATIONS

Lloyd A. Treinish* and Anthony P. Praino

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

1. INTRODUCTION

A wide variety of operations in the transportation industry are weather-sensitive to local conditions in the short-term (3 to 36 hours). Typically, they are reactive due to unavailability of appropriate predicted data at the required temporal and spatial scale. Hence, optimization that is applied to these processes to enable proactive efforts utilize either historical weather data as a predictor of trends or the results of synoptic-scale weather models. While near-real-time assessment of observations of current weather conditions may have the appropriate geographic locality, by its very nature is only directly suitable for reactive response. Alternatively, mesoscale numerical weather models operating at higher resolution in space and time with more detailed physics may offer greater precision and accuracy within a limited geographic region for problems with short-term weather sensitivity (e.g., Mass et al, 2002; Gall and Shapiro, 2000). Such forecasts can be used for competitive advantage or to improve operational efficiency and safety. In particular, they appear to be well suited toward improving economic and safety factors of concern to state and local highway administrations, and operators of airports. They are also relevant to other state and local agencies responsible for emergency management due to the effects of severe weather.

1.1 Economic and Societal Impacts

Adverse weather conditions have a significant effect on the operation of roads at all scales. According to the United States Department of Transportation, about 7000 people are killed and 800,000 are injured each year in weather-related accidents on US highways. The economic impact of these and other weather-related problems on the roads are estimated to lead to 544M vehicle-hours of delay, and an economic impact of about \$42B annually. A major snowstorm in a large city or state can cost millions of dollars per day due to clean-up efforts, wasted salaries and lost revenue and taxes. Hence, highway administrations are concerned with routine and emergency planning for snow (e.g., removal, crew and equipment deployment, selection of deicing material), road repair, maintenance and construction, repair of downed power lines and trees along roads due to severe winds, evacuation from and other precautions for areas of potential flooding, etc.

According to the Air Transportation Association, air traffic delays caused by weather were about \$4.2B in 2000, of which \$1.3B was estimated to be avoidable. Conceptually, improvements in the quality and lead-time of local weather forecasts could enable air traffic controllers and dispatchers to develop more effective alternative flight paths to reroute aircraft around danger-

*Corresponding author address: Lloyd A. Treinish, IBM T. J. Watson Research Center, 1101 Kitchawan Road, Yorktown Heights, NY 10598, lloyd@us.ibm.com, <http://www.research.ibm.com/people/l/treinish>

ous weather. Airline officials could initiate recovery plans before weather-induced disruptions actually occur, rescheduling passengers and aircraft to avoid congestion in affected areas and to improve safety. Airport operators could more efficiently schedule and staff aircraft deicing and snow removal operations during the winter. (e.g., Changnon, 2003 and Dutton, 2002)

2. PREVIOUS WORK

To begin to address some of these issues, we build upon our earlier work, the implementation of an operational testbed, dubbed "*Deep Thunder*". This prototype provides nested 24-hour forecasts for the New York City metropolitan area to 1 km resolution, which are updated twice daily. The work began with building a capability sufficient for operational use. In particular, the goal is to provide weather forecasts at a level of precision and fast enough to address specific business problems. Hence, the focus has been on high-performance computing, visualization, and automation while designing, evaluating and optimizing an integrated system that includes receiving and processing data, modelling, and post-processing analysis and dissemination. Part of the rationale for this focus is practicality. Given the time-critical nature of weather-sensitive transportation operations, if the weather prediction can not be completed fast enough, then it has no value. Such predictive simulations need to be completed at least an order of magnitude faster than real-time. But rapid computation is insufficient if the results can not be easily and quickly utilized. Thus, a variety of fixed and highly interactive flexible visualizations have also been implemented, including ones focused on support of operational decision making in transportation (Treinish et al, 2003).

The concept behind *Deep Thunder* in this context is clearly to be complementary to what the National Weather Service (NWS) does and to leverage their investment in making data, both observations and models, available. It is therefore also complementary to the deployment of Road Weather Information System (RWIS) stations by state highway administrations to monitor real-time weather conditions along roads. The idea, however, is to have highly focused modelling by geography with a greater level of precision and detail than what is ordinarily available.

3. RELATED EFFORTS

There are a number of independent projects involved in the utilization of mesoscale modelling for transportation applications. Only a few of those efforts will be summarized herein to help characterize them in comparison to the *Deep Thunder* work.

3.1 Maintenance Decision Support System

A major project supported by the US Federal Highway Administration (FHWA) considers weather impacts on road maintenance and operations. Their initial moti-

vation was based upon the lack of integration between weather information, including forecasts and decisions made by winter maintenance managers. This led to a comprehensive and very focused effort called the (winter) Maintenance Decision Support System (MDSS), which integrates weather observations, detailed weather forecasts, observed road conditions, forecasted road conditions, recent maintenance practices, and available maintenance resources (e.g., Pisano et al, 2004).

The early modelling work for this project included a complete data collection, data analysis and numerical weather prediction system for the state of Iowa. Hourly assimilation of surface observations from the National Weather Service and RWIS stations in Iowa was accomplished via the Local Analysis and Prediction System (LAPS). The resultant analyses were then used as initialization to the Penn State/NCAR Mesoscale Model (MM5). The model grid used a double nest configuration with a 30 km grid spacing on the outer grid, and a 10 km inner grid centered on Iowa with coverage over much of the Midwest. Model forecasts were generated four times per day out to 30 hours. Web-based images were created from standard meteorological analysis and forecast fields and interpolated to road locations. The data were also input into a custom road-surface model, which provided predictions of pavement temperature and surface snow/ice build-up which were then available as a support tool for highway chemical treatment decisions (Snook, 2001).

Currently, this capability has evolved into an operational prototype, which was used in the 2002-2003 winter and planned to be used in the upcoming 2003-2004 winter. It utilizes an ensemble of six mesoscale model runs four times daily, focused on a 12-24 hour forecast. The ensemble consisted of three different mesoscale models (MM5, Regional Atmospheric Modeling System -- RAMS, and Weather Research and Forecasting Model -- WRF) using two different models for lateral boundaries (Eta and AVN). The models are configured with identical grid size, resolution, and geometry, centered on Iowa (Pisano et al, 2004)

3.2 University Projects

Given the near-real-time availability of input data, lowered cost of high-performance computing systems and growing quality of modelling codes, the capability of utilizing numerical weather prediction systems for transportation applications has been established at a number of universities. Several of these efforts have included extensive work in forecast verification and development of data assimilation methods as well as refinement of model physics. A few notable projects are summarized briefly herein. This discussion is not meant to be either comprehensive nor complete, but merely to indicate the scale of related activities in the academic community.

Mass et al, 2003 outlines the application of an MM5-based system in the northwestern United States, which includes 36, 12 and 4 km nests, as well as some ensemble runs at the coarser resolutions. The forecast results are being used for a number of applications, including providing road surface temperatures for Washington State Department of Transportation.

Colle et al, 2003 utilize a similar configuration of

MM5 for their capabilities, but focused on the northeastern United States. Their efforts have included detailed analysis of the model results, comparison to synoptic-scale forecasts utilizing Eta, contrasting forecasts utilizing boundary conditions derived from Eta or AVN, and evaluation for use for operational forecasting.

Carpenter et al 1999 discusses the use of the Advanced Regional Prediction System (ARPS) to support airport terminal operations. The nested forecasts at 27, 9 and 3 km resolution focus on specific large airports in the midwestern United States.

4. FORECAST MODEL DESCRIPTION

The model used for the *Deep Thunder* project is non-hydrostatic with a terrain-following coordinate system and includes interactive, nested grids. It is a highly modified version of RAMS (Pielke et al, 1992), which is derived from earlier work supporting the 1996 Centennial Olympic Games in Atlanta (Snook et al, 1998). It includes full cloud microphysics (e.g., liquid and ice) to enable explicit prediction of precipitation, and hence, does not utilize any cumulus parameterization. Operationally, a 3-way nested configuration is utilized via stereographic projection. Each nest is a 62 x 62 grid at 16, 4 and 1 km resolution, respectively (i.e., 976 x 976 km², 244 x 244 km² and 61 x 61 km²), focused on New York City, which is illustrated in Figure 1. The specific location was chosen to include the major airports operating in the New York City metropolitan area within the 1 km nest. The three nests employ 48, 12 and 3 second time steps, respectively. The time steps were chosen to ensure computational stability and to also accommodate strong vertical motions that can occur during modelling of severe convection. Each nest employs the same vertical grid using 31 stretched levels with the lowest level at 48 m above the ground, a minimum vertical grid spacing of 100 m, a stretch factor of 1.12 and a maximum grid spacing of 1000 m. At the present time, two 24-hour forecasts are produced daily, typically initiated at 0Z and 12Z. Additional runs are scheduled with initialization at 6Z and/or 18Z either on-demand or during interesting weather events.

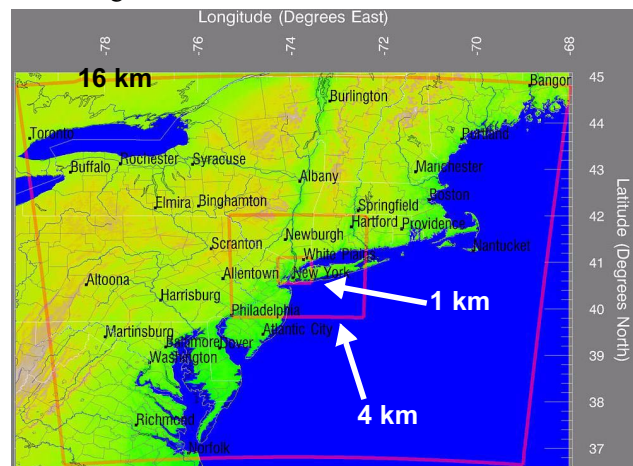


Figure 1. Model Nesting Configuration.

Currently, the data for both boundary and initial conditions for each model execution are derived from the Eta synoptic-scale model operated by the National

Centers for Environmental Prediction (NCEP), which covers all of North America and surrounding oceans at 12 km resolution and 60 vertical levels. These data are made available via the National Weather Service NOAAport data transmission system after sampling to 40 km resolution on the AWIPS 212 grid and interpolated to 27 isobaric levels for the continental United States in a Lambert-Conformal projection. In addition, the model lateral boundaries are nudged every three hours, using the Eta-212 grids. Static surface coverage data sets provided by the United States Geological Survey at 30-second resolution are used to characterize topography and vegetation coverage. Similar but lower-resolution data are used to define land use and coverage and sea surface temperature. The static and dynamic data are processed via an isentropic analysis package to generate three-dimensional data on the model nested grids for direct utilization by the modelling code.

5. ARCHITECTURE AND IMPLEMENTATION

The components of *Deep Thunder* are shown schematically in Figure 2, and are described below from left to right.

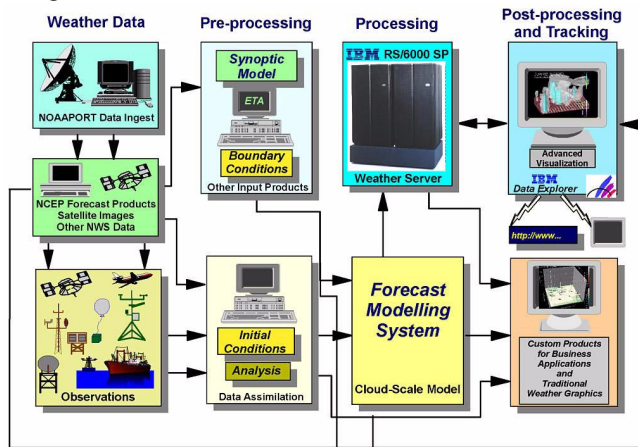


Figure 2. *Deep Thunder* Architecture.

5.1 Data

The NOAAport system provides a number of different data sources, including *in situ* and remotely sensed observations used currently for forecast verification as well as the aforementioned Eta data for model boundary and initial conditions. A three-channel facility manufactured by Planetary Data, Incorporated is utilized, which was installed at the IBM Thomas J. Watson Research Center in Yorktown Heights, NY in 2000 and upgraded in 2003. The NOAAport and other hardware that supports this project is shown in Figure 3. This NOAAport receiver system, based upon Linux, has a very flexible design, enabling the type of customization and integration necessary to satisfy the project goals. The various files transmitted via NOAAport are converted into conventional files in Unix filesystems in their native format, accessible via NFS mounting on other hardware systems via a private gigabit ethernet.

5.2 Pre-Processing

The pre-processing consists of two parts. The first is essentially a parsing of the data received via

NOAAport into usable formats to be used by the second part -- analysis and visualization, which is discussed in Treinish et al, 2003. The data and procedural flow of these processes is outlined in Figure 4. Other aspects related to forecast verification and product visualization are discussed in subsequent sections.

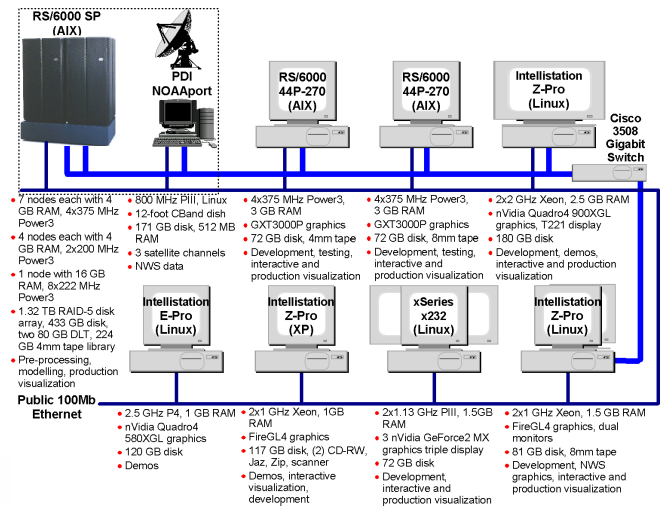


Figure 3. *Deep Thunder* Hardware Environment.

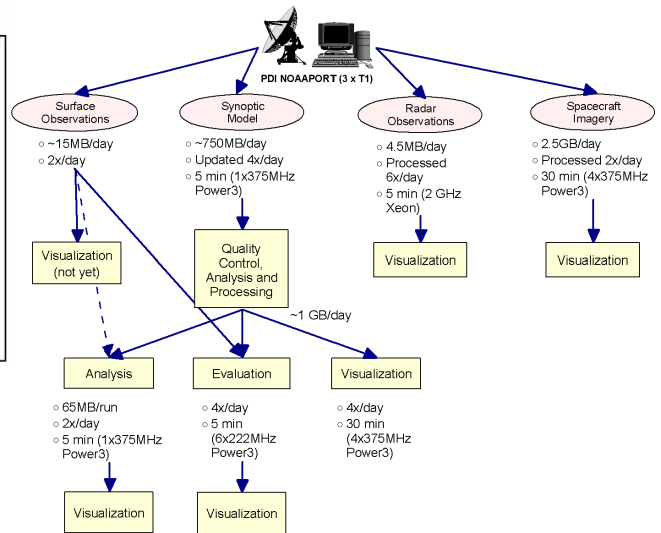


Figure 4. Pre-Processing Procedural and Data Flow.

5.3 Processing

To enable timely execution of the forecast models, the simulation is parallelized on a high-performance computing system. For this effort, an IBM RS/6000 Scalable Power Parallel (SP) is employed. This is IBM's previous generation of supercomputer systems, which is in common use at many operational centers for numerical weather prediction or has been upgraded with newer IBM systems built with a similar architecture. The SP is a distributed memory MIMD computer, consisting of two to 512 RS/6000 processor nodes, that communicate via a high-speed, multi-stage interconnect (the SP Switch). Each node has an SMP configuration of two to 16 Power3 processors. In the current implementation for the *Deep Thunder* effort, there are seven nodes of four 375 MHz Power3 processors, four nodes

of two 200 MHz Power3 processors and one node of eight 222 MHz Power3 processors as shown in Figure 3. The modelling software is parallelized using the Scalable Modelling System/Nearest-Neighbor Tool described by Edwards et al (1997) for single model domains. It has been extended to support multiple nests for the current operational efforts. The modelling domain for all nests is spatially decomposed for each processor to be utilized, which is mapped to an MPI task. Within each node, there are four MPI tasks, which communicate via shared memory. The SP switch fabric enables communications between nodes. None of these tasks do I/O. Instead an additional processor is utilized to collect results from the MPI tasks and perform disk output asynchronously. This enables an efficient utilization of the SP platform for the modelling code. For current operations, seven nodes of four 375 MHz processors each are used for computing and a single 222 MHz cpu of another node is used for I/O. A typical model run with the aforementioned configuration requires about 1.8 hours to 2.2 hours to complete. This variation is due to the relative dominance of radiative vs. microphysics calculations, respectively for a particular forecast run.

5.4 Post-Processing

Post-processing essentially operates on the raw model output to provide useful products. There are several aspects of post-processing, the most important of which is visualization. Since large volumes of data are produced, which are used for a number of applications, the use of traditional graphical representations of data for forecasters can be burdensome. Alternative methods are developed from a perspective of understanding how the weather forecasts are to be used in order to create task-specific designs. This approach for *Deep Thunder* and details of its implementation are discussed in Treinish, 2001.

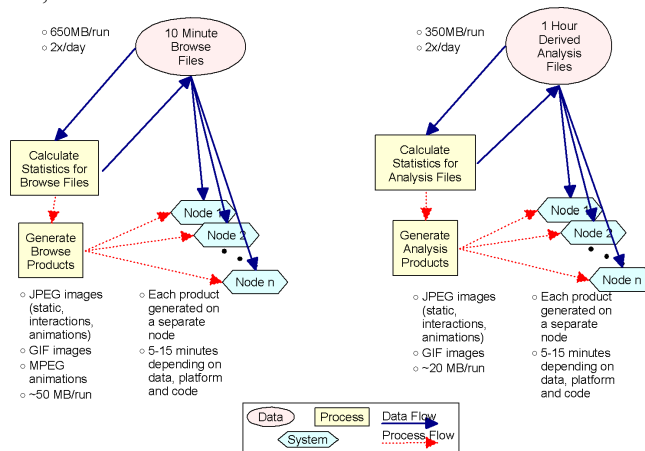


Figure 5. Visualization Post-Processing Procedural and Data Flow.

To enable timely availability of the visualizations, the parallel computing system used for the model execution is also utilized for post-processing. This approach is outlined schematically in Figure 5. Two types of output data are generated by the model. The first, is a comprehensive set of variables at hourly resolution (analysis) files for each nest, which are further pro-

cessed to generate derived products and interpolated to isobaric levels from the model terrain-following coordinates.

The second output is a subset of variables relevant to the applications of the model output produced every 10 minutes of forecast time. The finer temporal spacing is required to better match the model time step in all nests as well as to capture salient features being simulated at the higher resolutions. A subset is chosen to minimize the impact of I/O on the processing throughput. These browse files are also generated to enable visualization of model results during execution for quality control and simulation tracking.

Two classes of visualizations are provided as part of the *Deep Thunder* system. The first is a suite of highly interactive applications utilizing the workstation hardware shown in Figure 3, including ultra-high-resolution and multi-panel displays (Treinish, 2001).

The second is a set of web-based visualizations, which are generated automatically after each model execution via a set of hierarchical scripts (Treinish, 2002 and Treinish, 2003). That work is also illustrated schematically in Figure 5. This simple parallelism, including intranode parallelism, enables the independent generation of various products for placement on a web server to be completed in five to fifteen minutes.

An approach similar to that used for visualization is employed for forecast verification, which yields a set of evaluation tables as well as visualizations for each model run as well as the aggregation of all model runs during the previous week. The later are presented via web pages in a manner following that of the model visualizations. The details of this approach and examples are discussed in Praino et al, 2003, Treinish et al, 2003, and Praino and Treinish, 2004.

5.5 Integration

All of the components are operated by a master script, implemented in the Perl scripting language. Model executions are set up via a simple spreadsheet identifying basic run characteristics such as start time, length, location, resolution, etc. A Unix crontab is used to initiate the script. In addition to bookkeeping and quality control and logging, it polls input data availability whose arrival via the NOAAport is variable, does all the necessary pre-processing steps, initiates the parallel modelling job and then launches the parallel visualization post-processing.

6. CUSTOMIZATION FOR TRANSPORTATION

The aforementioned infrastructure provides a foundation for customization for transportation applications. To date, this has been done primarily through visualization. A few of the types of focused products that have been implemented and are operational, are discussed below.

6.1 Highway Operations

The first examples are of very focused visualizations by both geography and application for use by highway supervisors. A user has a limited choice of specific two-dimensional map products on a web site. A portion

of a sample page is shown in Figure 6, which illustrates forecasted temperatures as fixed colored contour bands

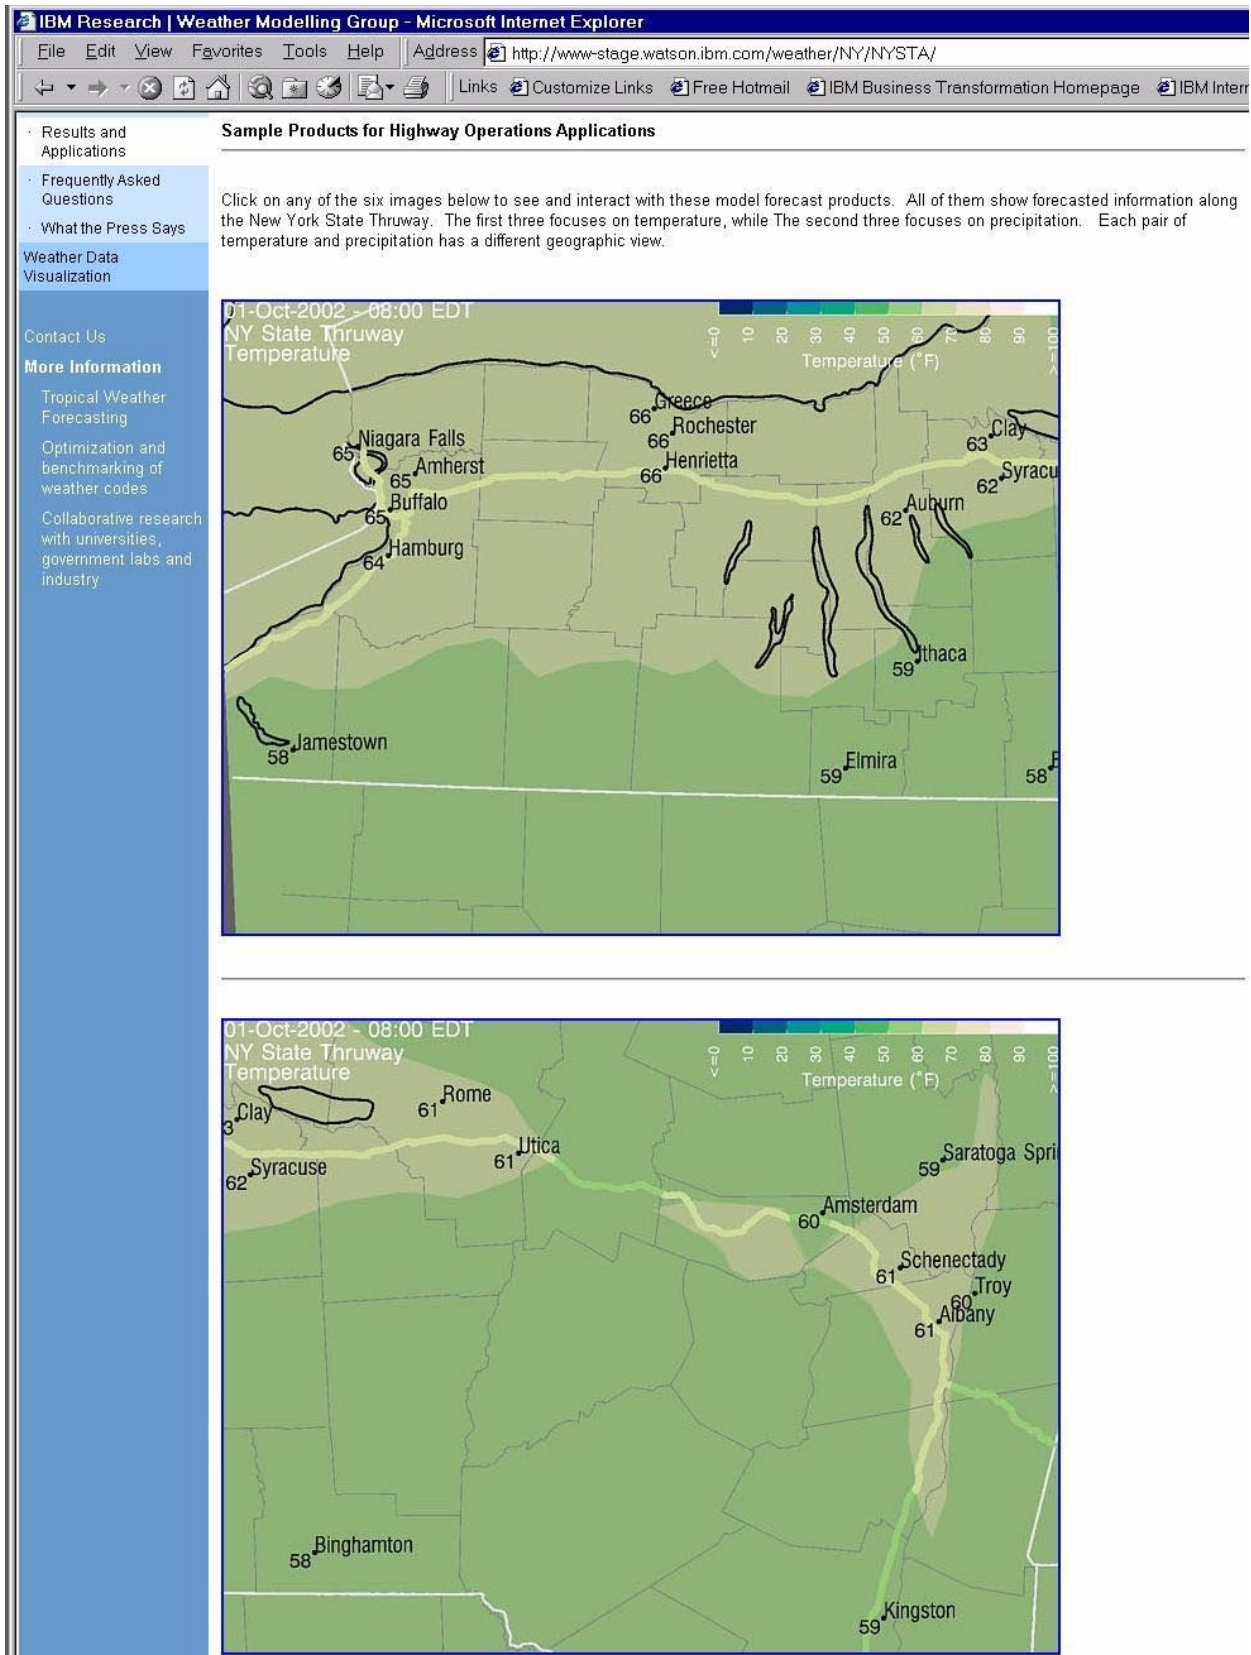


Figure 6. Portion of a Custom Weather Map Selection Web Page for Highway Operation

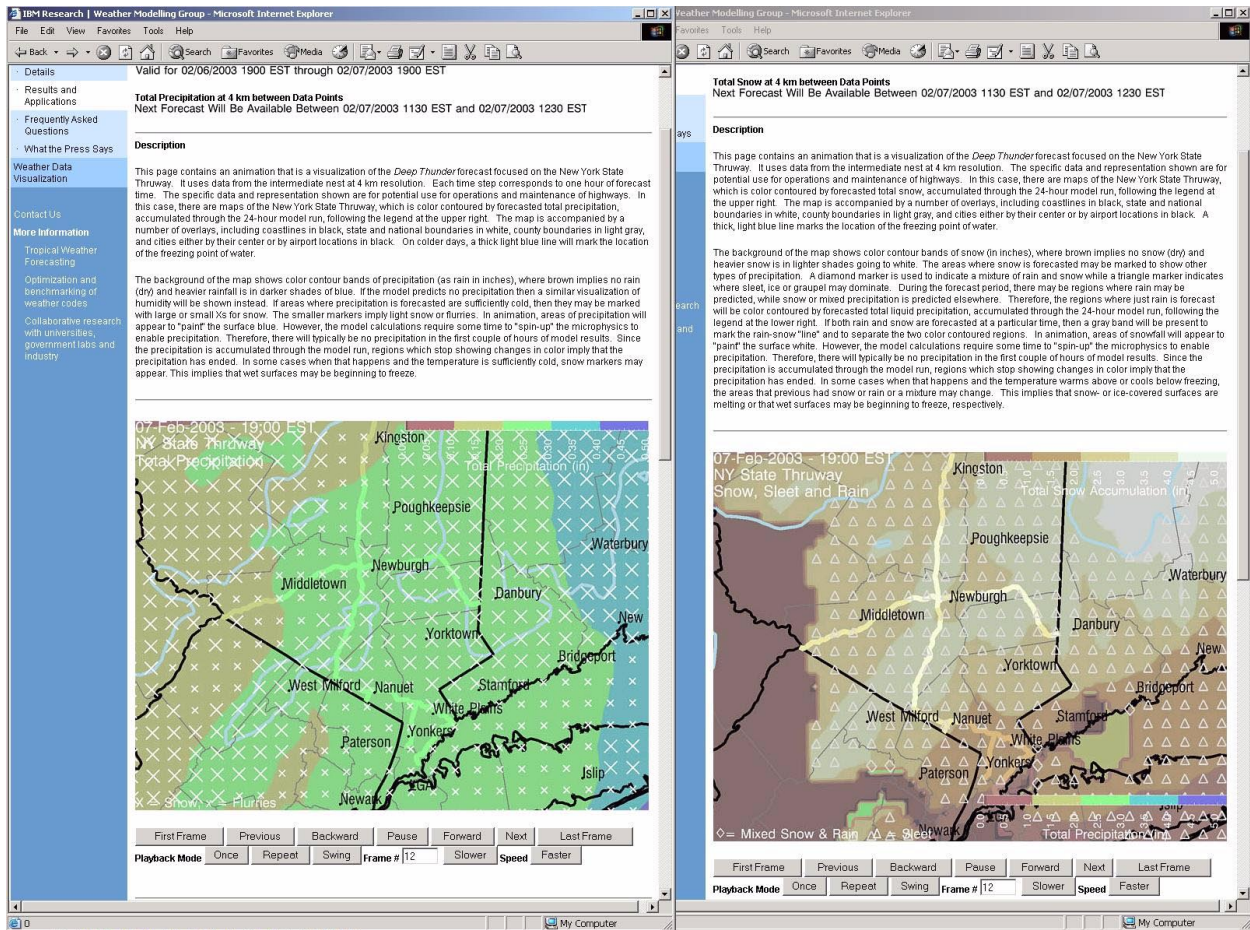


Figure 7. Precipitation Animation Web Pages.

for various parts of New York State. Similar maps are available for precipitation forecasts. This type of map also serves as an index to further animations of the particular variable and geographic area. Each time step in each animation corresponds to one hour of forecast time. There is a map of the New York State Thruway, which is color contoured by forecasted temperature, following the legend at the upper right. The map is accompanied by a number of overlays, including coastlines in black, state and national boundaries in white, county boundaries in light gray, and cities either by the center or airport locations in black. The forecasted temperatures at each city locations are also shown. The city locations and values will be in white when lower temperatures are forecast to avoid blending with the background colors. On colder days, a thick light blue line will mark the location of the freezing point of water. The background of the map shows color contour bands of temperature (degrees Fahrenheit). During windy winter days, contour bands of wind chill temperature will be shown instead. During humid summer days, contour bands of heat index temperature will be shown instead.

The second set of animations are of precipitation such as the left half of Figure 7. They are constructed in a similar fashion as the temperature maps. The map of the New York State Thruway is color contoured by forecasted total precipitation, accumulated through the 24

hour model run, following the legend at the upper right. The background of the map shows color contour bands of precipitation (as rain in inches), 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 with large or small Xs for snow. The smaller markers imply light snow or flurries. In an animation, areas of precipitation will appear to "paint" the surface blue. However, the model calculations require some time to "spin-up" the microphysics to enable precipitation. Therefore, there will typically be no precipitation in the first couple of hours of model results. Since the precipitation is accumulated through the model run, regions which stop showing changes in color imply that the precipitation has ended. In some cases when that happens and the temperature is sufficiently cold, snow markers may appear. This implies that wet surfaces may be beginning to freeze.

If the current model forecast indicates that snow, sleet or snow mixed with rain is predicted then an additional similar animation will be generated, as shown in the right half of Figure 7. The map of the New York State Thruway is color contoured by forecasted total snow, accumulated through the 24-hour model run, following the legend at the upper right. The background of the map shows color contour bands of snow (in inches),

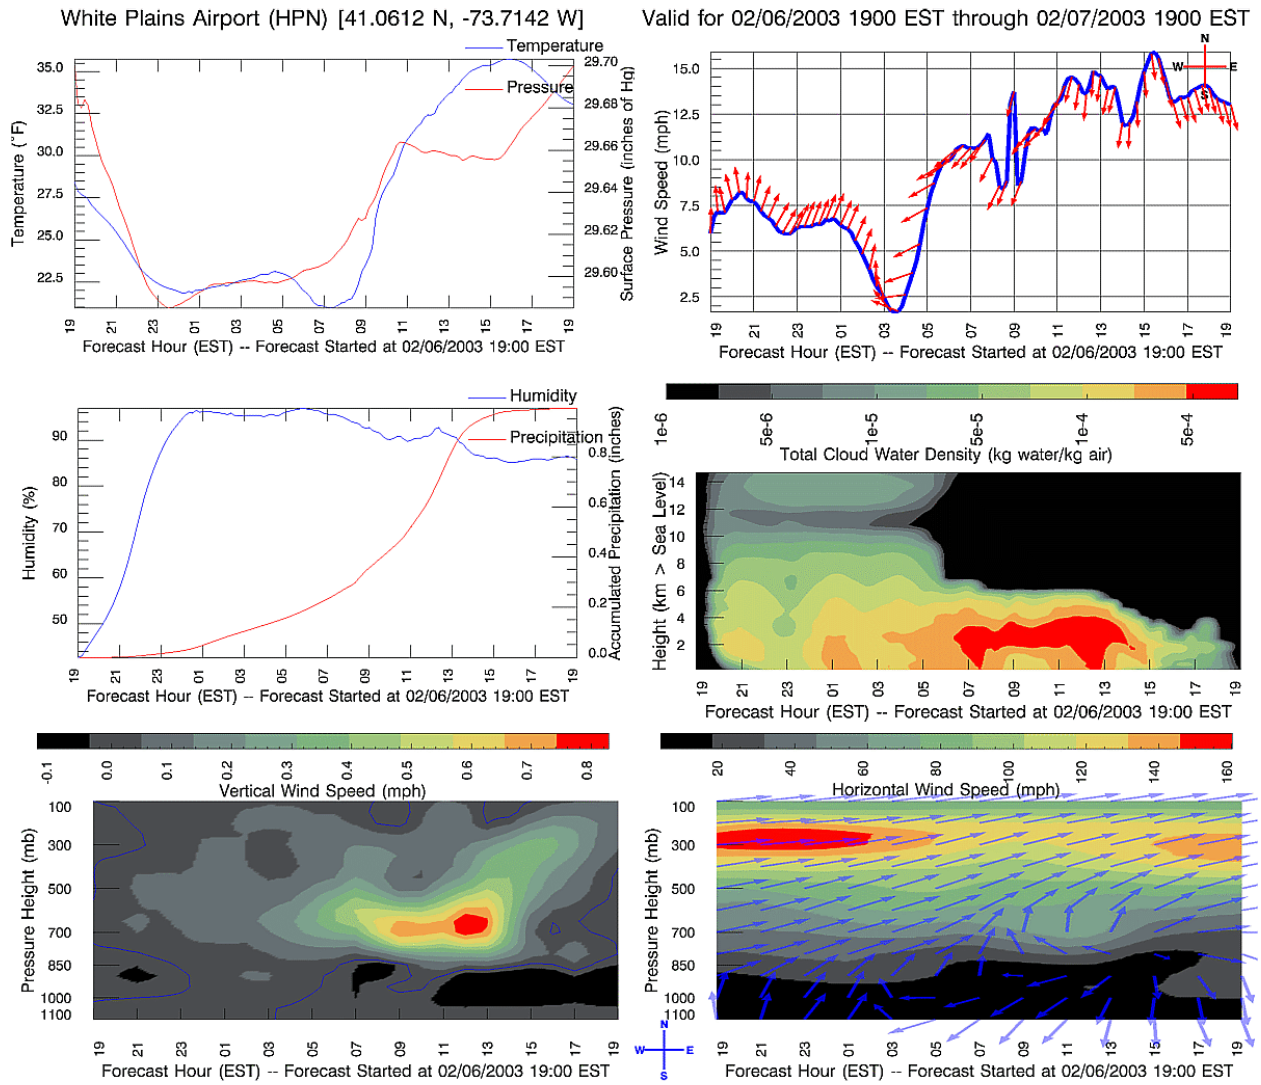


Figure 8. Sample Site-Specific One- and Two-Dimensional Visualizations for an Airport Terminal.

where brown implies no snow (dry) and heavier snow is in lighter shades going to white. The areas where snow is forecasted may be marked to show other types of precipitation. A diamond marker is used to indicate a mixture of rain and snow while a triangle marker indicates where sleet, ice or graupel may dominate. During the forecast period, there may be regions where rain may be predicted, while snow or mixed precipitation is predicted elsewhere. Therefore, the regions where just rain is forecast will be color contoured by forecasted total liquid precipitation, accumulated through the 24-hour model run, following the legend at the lower right. If both rain and snow are forecasted at a particular time, then a gray band will be present to mark the rain-snow "line" and to separate the two color contoured regions. In animation, areas of snowfall will appear to "paint" the surface white. Since the precipitation is accumulated through the model run, regions which stop showing changes in color imply that the precipitation has ended. In some cases when that happens and the temperature warms above or cools below freezing, the areas that previous had snow or rain or a mixture may change. This

implies that snow- or ice-covered surfaces are melting or that wet surfaces may be beginning to freeze, respectively.

6.2 Airport Operations

Other types of visualization are designed for use in supporting airport operations. The first one is a specialized meteogram shown in Figure 8. It consists of three panels showing surface data and three panels to illustrate upper air data. In all cases, the variables are shown as a function of time interpolated to a specific location (White Plains Airport within the 1 km nest). The upper and middle 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). Since the precipitation is accumulated through the model run, the slope of the curve will be indicative of the predicted rate of precipitation. Therefore, when the slope is zero, it is not raining (or snowing). The top right plot illustrates forecasted surface 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 (compass) direction to which the wind is going. The middle right plot is a colored contour map of forecasted 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 on the top of the panel. The bottom two panels show upper air winds using some of the same techniques. The lower left shows contours of vertical




Figure 9. Cloud, Wind and Visibility Visualizations for Airport Operations.

winds as a function of time and pressure following the legend above it. In addition, the zero velocity contour is shown in blue. At the lower right, is a contour map of horizontal wind speed also as a function of time and pressure. It is overlaid with arrows (blue) to illustrate the predicted compass wind direction. This example illustrates the forecast of snow in the White Plains terminal area from the same model run shown in Figure 7.

Figure 9 has two panels each of which 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. This surface is not a cloud boundary. 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 increments of 2 km, corresponding to the height in meters of the forecasted cloud base as shown in the color legend to the lower right. Areas in gray imply no cloud data. The cloud base contours are overlaid with maps of coastlines and state boundaries in black and rivers in blue. The volume is marked at the locations of major airports with set of colored poles. The top panel is from the 16 km nest, marking DCA, PHL (3), LGA and BOS (4). The bottom panel is from the 4 km nest, marking PHL (1), EWR (2) and HPN (3). The poles are color contoured by the derived visibility using the color legend to the lower left. At each of 21 pressure levels, the horizontal wind is shown via arrows. The arrows are colored by horizontal wind speed following the legend to the upper left. The arrow length also corresponds to speed. This approach presents information relevant to flight planning as opposed to direct meteorological analysis. These images are from animation sequences of forecasts showing Hurricane Isabel moving into the northeastern United States. These animations are also presented on the web as shown in Figure 10 from the 4 km nest.

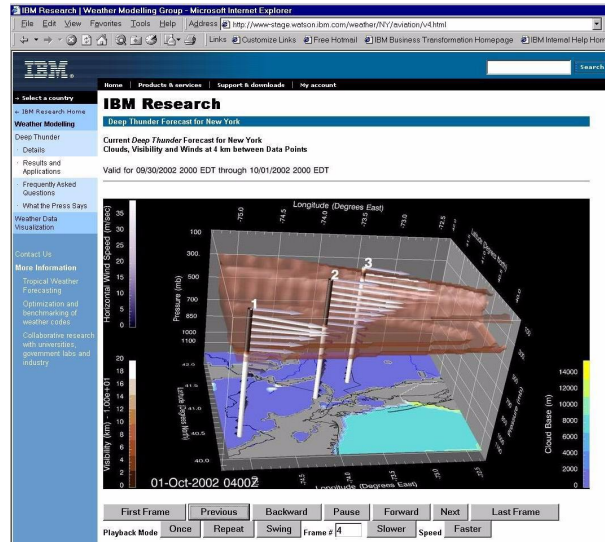


Figure 10. Animation of Cloud, Wind and Visibility.

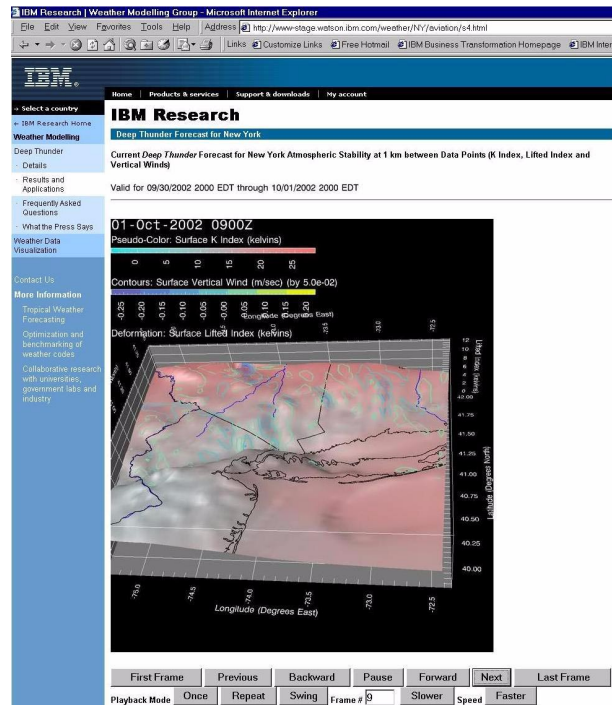


Figure 11. Animation of Atmospheric Stability.

Another type of visualization is shown in Figure 11 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 (4 km) is shown. 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. As with the previous example, the visualization is available as an hourly animation.

7. DISCUSSION

Even though the overall system and implementation is still evolving, the type of products that *Deep Thunder* can generate has provided a valuable platform to investigate a number of practical transportation applications. To aid in that evaluation, the products have been made available to several local agencies to assist in their operational decision making with various weather-sensitive problems. Aspects of the visualization and web-based access to these products are discussed in Treinish, 2003. A preliminary analysis of the operational forecasts for significant snow events during the 2002-2003 winter has also been done to support the evaluation (Praino and Treinish 2004).

The feedback from these users coupled with more rigorous verification (Praino et al, 2003) has raised a number of comments and issues. In general, there has been a very favorable view of the ability of the overall system to provide useful and timely forecasts of severe weather including convective events and high winds

with greater precision. The user-driven design of visualization products has enabled effective utilization of the model output. However, improved throughput is required to enable more timely access to the forecast products, which need to cover broader areas at higher resolution. These suggest the direction for continued work and improvement of the utility of the system.

8. CONCLUSIONS AND FUTURE WORK

This is an on-going effort. The results to date illustrate a practical and useful implementation with automatically generated user-application-oriented forecast visualizations on the world-wide-web. But they also point to several next steps besides refining the quality of the model results, improving the degree of automation, and developing new methods of visualization and dissemination.

To enable more timely availability of forecast products, the number of model runs each day will expand, including coverage of other geographic areas at high resolution. Overall throughput is limited by the capacity of the current hardware. Although plans for a modest increase will be applied, the model configuration may need to be adjusted to somewhat lower resolution with broader coverage focused on the specific geographic areas of concern expressed by the current set of users at local government agencies in order to meet delivery times for decision makers (e.g., highway supervisors). To support this effort, the operational environment will be enhanced to include an initial model run at relatively lower resolution (e.g., 16 to 20 km) covering the eastern United States as well as most of the Great Lakes and a significant portion of the western Atlantic Ocean to accommodate coastal effects. This model run, which would be completed in the order of 10 minutes of wall-clock time would utilize the lateral boundary conditions from Eta as in the current operations. The results would be used to generate hourly lateral boundaries for a set of nested forecasts, targeted for specific regions at different resolutions (e.g., 4 and 1 km, 8 and 2 km) that could be run simultaneously or in sequence, depending on available computing resources.

One of the applications for winter road maintenance and operations is to enable some predictive capability for lake-effect snow. The understanding of the morphology of such snow storms remains a significant challenge (Laird and Kristovich, 2004). However, the utilization of the class of mesoscale modelling used for *Deep Thunder* may provide useful information on the potential to support highway operations. In that regard, some experimental hindcasts have been done to gauge that potential. An example is shown in Figure 12.

In this case, a model run was established, which simulated an operational scenario for a forecast on December 28, 2001. The forecast was initialized at 0Z on that date, to cover the second Lake Erie phase of the *Bald Eagle* lake-effect snow event that lasted from December 24, 2001 through January 1, 2002. This phase was characterized by an early morning snow (0100 - 0700 EST) in Buffalo, NY (15.2 inches), followed by a mid-day snowfall of 7 inches (1100 - 1300 EST).

The current model configuration was only modified

to operate at 18, 6 and 2 km resolution with three 68 x 68 x 31 grids, centered on Buffalo, with time steps of 60, 20 and 6.67 seconds. A qualitative, yet comprehensive, three-dimensional visualization of the results from the 2 km nest at two different times is shown in Figure 12. Both panels show a terrain map, colored by contour bands of a forecast of total snow, where lighter shades indicate heavier accumulations. The map is marked with the location of major cities or airports as well as river, coastline, and national, state and county boundaries. In addition, there are colored arrows indicating predicted winds, with the lighter color being faster winds and the arrow direction corresponding to the direction to which the wind is flowing. Above the terrain is a forecast of clouds, represented by a three-dimensional translucent white surface of total cloud water density (water and ice) at a threshold of 10^{-4} kg water/kg air. Within the cloud surface is translucent cyan surface of forecast reflectivities at a threshold of 30 dbZ.

The top panel of Figure 12 shows the forecast after the snow band has moved north. The bottom panel shows the forecast for nine hours later, when the snow band has reversed direction and moved back to the south. A light blue line on Lake Erie outlines a region where the air temperature has gone above freezing providing additional moisture for the snow band. Although the total snow amounts were low compared to measurements reported by snow spotters, the model run resolved the lake-effect snow band and the reversal in direction of the flow over Lake Erie with a positive bias of about an hour in time. Given these results and the fact that the forecast could have been provided operationally at about 0400 UTC, and thus, a long lead-time for highway supervisors is very encouraging for its potential for winter road maintenance.

As these customized capabilities are made available to assist in weather-sensitive transportation operations, efforts will also be addressed to determine and apply appropriate metrics for measuring business value. These will serve to provide an evaluation of *Deep Thunder* that is complementary to the traditional meteorological verification.

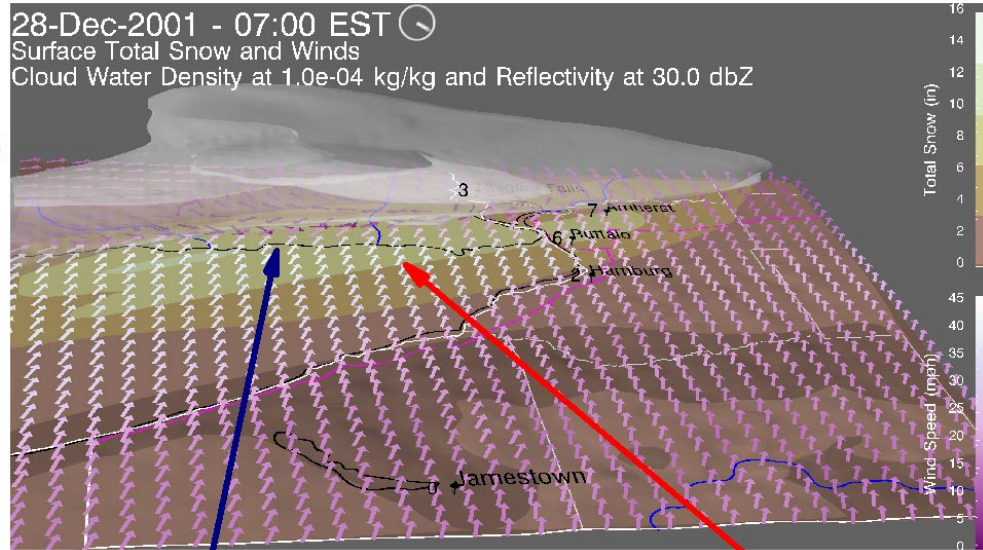
9. ACKNOWLEDGEMENTS

The authors wish to acknowledge the contributions to this work by three students. Annette Samuelsen, a Ph.D. candidate in oceanography at Florida State University, implemented much of the software that is used for the verification portion of the *Deep Thunder* system during her summer 2002 internship at the IBM Thomas J. Watson Research Center.

Craig Tashman, an undergraduate physics student at Rensselaer Polytechnic Institute, has an on-going internship at IBM Research. He began by working on techniques of graphics compression for the dissemination of model results, which are not discussed herein. More recently, he has expanded his work to include the implementation of the Java-based processing of NOAAport-received Eta GriB files.

Udam Dewaraja, an undergraduate computer engineering student at the University of Washington, implemented new capabilities for web-based and other

- Air temperature below freezing
- Snow band already moved north



Winds shift direction

Lake Erie

- Air temperature rises above freezing over Lake Erie
- Snow band moves south from Ontario and Niagara

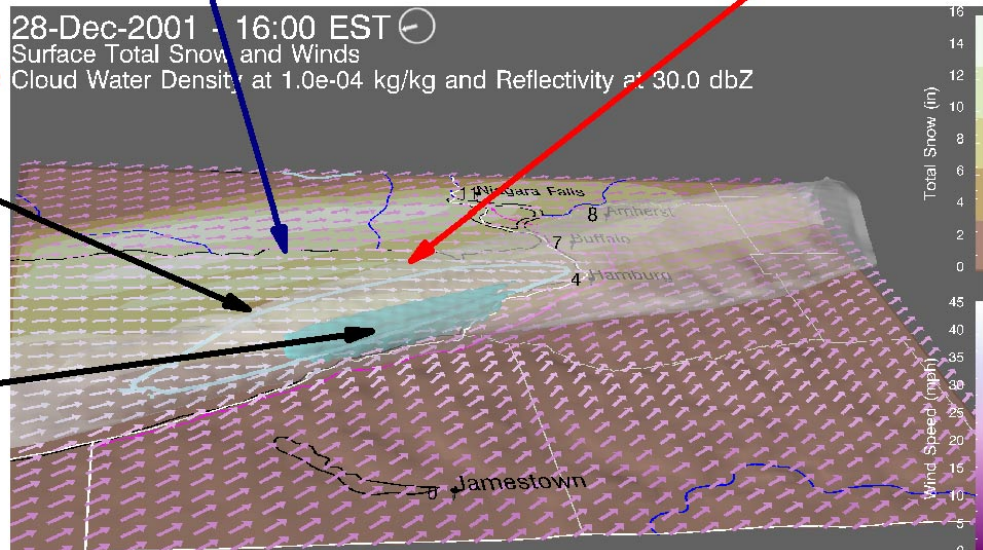


Figure 12. Experimental Forecast for Lake Effect Snow.

methods of forecast dissemination during his summer 2003 internship at the IBM Thomas J. Watson Research Center.

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

10. REFERENCES

Changnon, S. D. *Measures of Economic Impacts of Weather Extremes*. **Bulletin of the American Meteorological Society**, 84, no. 9, pp. 1231-1235, September 2003.

Carpenter, R.L. Jr., K. K. Droegemeier, G. M. Bassett, S. S. Weygandt, D. E. Jahn, S. Stevenson, W. Qualley and R. Strasser. *Storm-scale numerical weather prediction for commercial and military aviation, Part I: Results from operational tests in 1998*. **Proceedings of the Eighth Conference on Aviation, Range and Aerospace Meteorology**, January, 1999, Dallas, TX, pp. 209-211.

Changnon, S. D. *Measures of Economic Impacts of Weather Extremes*. **Bulletin of the American Meteorological Society**, 84, no. 9, pp. 1231-1235, September 2003.

Colle, B. A., J. B. Olson and J. S. Tongue. *Multiseason Verification of the MM5. Part I: Comparison with the Eta Model over the Central and Eastern United States and Impact of MM5 Resolution*. **Weather and Forecasting**, 18, no. 6, pp. 431-457, June 2003.

Colle, B. A., J. B. Olson and J. S. Tongue. *Multiseason Verification of the MM5. Part II: Evaluation of High-Resolution Precipitation Forecasts over the Northeastern United States*. **Weather and Fore-**

casting, 18, no. 6, pp. 431-457, June 2003.

- casting, **18**, no. 6, pp. 458-480, June 2003.
- Dutton, J. A. *Opportunities and Priorities in a New Era for Weather and Climate Services*. **Bulletin of the American Meteorological Society**, **83**, no. 9, pp. 1303-1311, September 2002.
- Edwards, J., J. S. Snook and Z. Christidis. *Forecasting for the 1996 Summer Olympic Games with the SMS-RAMS Parallel Model*. **Proceedings of the Thirteenth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, February 1997, Long Beach, CA, pp. 19-21.
- Gall, R. and M. Shapiro. *The Influence of Carl-Gustaf Rossby on Mesoscale Weather Prediction and an Outlook for the Future*. **Bulletin of the American Meteorological Society**, **81**, no. 7, pp. 1507-1523, July 2000.
- Laird, N. F. and D. Kristovich. *Utilizing Idealized Mesoscale Model Simulations to Aid in the Prediction of Lake-Effect Snowstorms*. To be published in **Proceedings of the 20th Conference on Weather Analysis and Forecasting/16th Conference on Numerical Weather Prediction**, January 2004, Seattle, WA.
- Mass, C. F., D. Owens, K. Westrick and B. A. Colle. *Does Increasing Horizontal Resolution Produce More Skillful Forecasts*. **Bulletin of the American Meteorological Society**, **83**, no. 3, pp. 407-430, March 2002.
- Mass, C. F., M. Albright, D. Owens, R. Steed, M. MacIver, E. Gritter, T. Eckel, B. Lamb, J. Vaughan, K. Westrick, P. Storck, B. Colman, C. Hill, N. Maykut, M. Gilroy, S. A. Ferguson, J. Yetter, J. M. Sierchio, C. Bowman, R. Stender, R. Wilson and W. Brown. *Regional Environmental Prediction Over the Pacific Northwest*. **Bulletin of the American Meteorological Society**, **84**, no. 10, pp. 1353-1366, October 2003.
- Pielke, R. A., W. R. Cotton, R. L. Walko, C. J. Tremback, W. A. Lyons, L. D. Grasso, M. E. Nicholls, M.-D. Moran, D. A. Wesley, T. J. Lee and J. H. Copeland. *A Comprehensive Meteorological Modeling System - RAMS*. **Meteorology and Atmospheric Physics**, **49**, 1992, pp. 69-91.
- Pisano, P. A., A. D. Stern and W. P. Mahoney. *The Winter Maintenance Decision Support System (MDSS): Demonstrations and Future Plans*. To be published in **Proceedings of the 20th International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, January 2004, Seattle, WA.
- Praino, A. P., L. A. Treinish, Z. D. Christidis and A. Samuelsen. *Case Studies of an Operational Mesoscale Numerical Weather Prediction System in the Northeast United States*. **Proceedings of the Nineteenth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, February 2003, Long Beach, CA.
- Praino, A. P. and L. A. Treinish. *Winter Forecast Performance of an Operational Mesoscale Numerical Modelling System in the Northeast U.S. -- Winter 2002-2003*. To be published in **Proceedings of the 20th Conference on Weather Analysis and Forecasting/16th Conference on Numerical Weather Prediction**, January 2004, Seattle, WA.
- Snook, J. S., P. A. Stamus, J. Edwards, Z. Christidis and J. A. McGinley. *Local-Domain Mesoscale Analysis and Forecast Model Support for the 1996 Centennial Olympic Games*. **Weather and Forecasting**, **13**, no. 1, pp. 138-150, January 1998.
- Snook, J. S. *Foretell: An operational forecasting system designed for the surface transportation community*. **Proceedings of the 14th Conference on Numerical Weather Prediction**, July 2001, Ft. Lauderdale, FL.
- 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., A. P. Praino and Z. D. Christidis. *Implementation of Mesoscale Numerical Weather Prediction for Weather-Sensitive Business Operations*. **Proceedings of the Nineteenth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, February 2003, Long Beach, CA.
- Treinish, L. *Web-Based Dissemination and Visualization of Mesoscale Weather Models for 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.