1.5

THE ROLE OF MESO-γ-SCALE NUMERICAL WEATHER PREDICTION AND VISUALIZATION FOR WEATHER-SENSITIVE DECISION MAKING

Lloyd A. Treinish* and Anthony P. Praino
IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA

1. INTRODUCTION

Weather-sensitive business operations are often reactive to short-range, local conditions due to unavailability of appropriate predicted data at this temporal and spatial scale. This situation is commonplace in a number of applications, some of which address either planning for or response to hazards or disasters. These include but are not limited to transportation, agriculture, energy, insurance, entertainment, construction, communications and emergency management. Typically, what 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- to meso-β-scale weather models. This time range is typically beyond what is feasible with modern nowcasting techniques. Hence, near-real-time assessment of observations of current weather conditions may have the appropriate geographic locality, but by its very nature is only directly suitable for reactive response.

Alternatively, meso-γ-scale (cloud-scale) numerical weather models operating at higher resolution in space and time with more detailed physics has shown "promise" for many years as a potential enabler of proactive decision making for both economic and societal value. They may offer greater precision and accuracy within a limited geographic region for problems with short-term weather sensitivity. In principle, such forecasts can be used for competitive advantage or to improve operational efficiency and safety by enhancing both the quality and lead time of such information. In particular, they appear to be well suited toward improving economic and safety factors of concern for transportation applications of interest to state and local highway administrations and airport terminal operators. They are also relevant to other state and local agencies responsible for emergency management due to the effects of severe weather. Among others, such factors relate to 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. Among others, such factors relate to 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, and short-term environmental impact (e.g., Changnon, 2003 and Dutton, 2002).

However, a number of open questions exist (e.g., Mass et al, 2002; Gall and Shapiro, 2000, de Elía and Laprise, 2003). For example, can both business and meteorological value be demonstrated beyond physical realism that such models clearly provide? Such "realism" is based upon the generation of small-scale features not present in background fields used for initial and boundary conditions. Further, can a practical and usable system be implemented at reasonable cost?

*Corresponding author address: Lloyd A. Treinish, IBM T. J. Watson Research Center, 1101 Kitchawan Road, Yorktown Heights, NY 10598, USA, lloydt@us.ibm.com, http://www.research.ibm.com/weather/DT.html.

To begin to address these issues, a prototype system, dubbed "*Deep Thunder*", was first implemented for the New York City metropolitan area. Later, it was extended via customizations for applications to specific weather-sensitive decision making problems and for forecasting in other metropolitan areas in the United States. The operational forecasts produced by this system are in evaluation with respect to their effectiveness for these applications.

2. FORECAST MODEL DESCRIPTION

The model used for this effort is non-hydrostatic with a terrain-following coordinate system and includes interactive, nested grids. It is a highly modified version of the Regional Atmospheric Modeling System or RAMS (Pielke et al, 1992), which is derived from earlier work supporting the 1996 Centennial Olympic Games in Atlanta (Snook et al, 1998). Typically for operational purposes, a 3-way nested configuration is utilized via stereographic projection in a 4:1 spatial and temporal resolution ratio. For the initial configuration focused on New York City, each nest is a 62 x 62 grid at 16, 4 and 1 km resolution, respectively (i.e., 976 x 976 km^2, 244 x 244 km^2 and 61 x 61 km^2). For the other three geographic areas, the three-way nests are 66 x 66 at 32, 8 and 2 km resolution, respectively. All of the operational domains are illustrated in Figure 1. The specific locations of the various configurations were chosen to include the major airports operating in the particular metropolitan area within the highest-resolution (1 or 2 km) nest as well as to have good coverage for a number of other weather-sensitive applications in each geographic region.

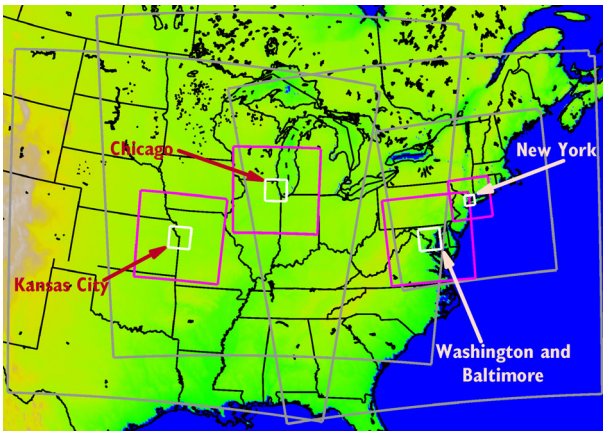

Figure 1. Model Nesting Configurations.

Figure 1 places all of the forecast domains in a geographic context, which shows a map of the eastern two-thirds of the continental United States. On the map are three regions associated with each of the four aforementioned metropolitan areas. They correspond to the triply nested, multiple resolution forecasting domains used to produce each high-resolution weather forecast. The outer nests are in gray, the intermediate nests are in magenta and the inner nests are in white.

The model configuration includes full bulk cloud microphysics (e.g., liquid and ice) for all nests to enable

explicit prediction of precipitation, and hence, does not utilize any cumulus parameterization. The three nests employ 48, 12 and 3 second time steps, respectively for New York and 100, 25, 6.25 second time steps, respectively for the other areas. 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, for each region, typically initiated at 0Z and 12Z or 6Z and 18Z. The 24-hour integration is done for all three nests. Additional runs are scheduled with initialization at other times either on-demand or during interesting weather events.

Currently, the data for both boundary and initial conditions for each model execution are derived from the North American Model (NAM, formerly known as Eta) operated by the United States 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 United States National Weather Service (NWS) NOAAport data transmission system in a number of formats and resolutions for the continental United States in a Lambert-Conformal projection. In addition, the model lateral boundaries are nudged every three hours, using these data. 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 (at 10-minute resolution) and sea surface temperature (one-degree resolution). The latter is updated to use data corresponding to the particular month in which the forecast is made. 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.

3. ARCHITECTURE AND IMPLEMENTATION

This effort began with building a capability sufficient for operational use for the New York City metropolitan area. 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 business decisions, 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. They range from techniques to enable more effective analysis to strategies focused on the applications of the forecasts. The focus is on nested 24-hour forecasts, which are typically updated twice daily. In 2004, the system was extended to provide forecasts for the Chicago, Baltimore/Washington and Kansas City metropolitan areas at 2 km resolution as outlined in Figure 1. In 2005, extensions were made to enable experimental forecasts for the San Diego metropolitan area to 1 km resolution and the Miami-Fort Lauderdale area to 1.5 km resolution. The former uses a nest configuration similar to the one used for the New York forecasts. The Florida configuration uses three 74 x 74 nests at 24, 6 and 1.5 km resolution, respectively. All of the processing, modelling and visualization are completed in 30 to 60 minutes on relatively modest hardware to enable sufficiently timely dissemination of forecast products at reasonable cost.

With such goals, the system has also evolved from its initial implementation. Hence, the discussion herein outlines the current approach, whose components are shown schematically in Figure 2, and are described below from left to right.

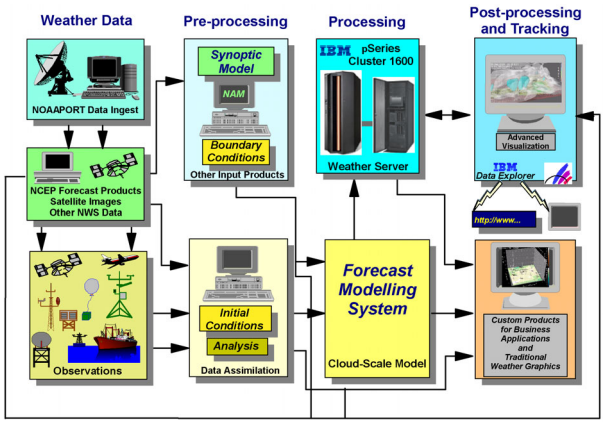

Figure 2. *Deep Thunder* Architecture.

3.1 Data

The NOAAport system provides a number of different data sources as disseminated by the NWS. These include *in situ* and remotely sensed observations used currently for forecast verification as well as the aforementioned NAM data for model boundary and initial conditions. For the *Deep Thunder* system, a four-channel facility manufactured by Planetary Data, Incorporated, is utilized, which was initially installed at the IBM Thomas J. Watson Research Center in 2000. It has gone through a succession of upgrades to accommodate newer computer systems and the migration of the satellite broadcasts to DVB-S technology.

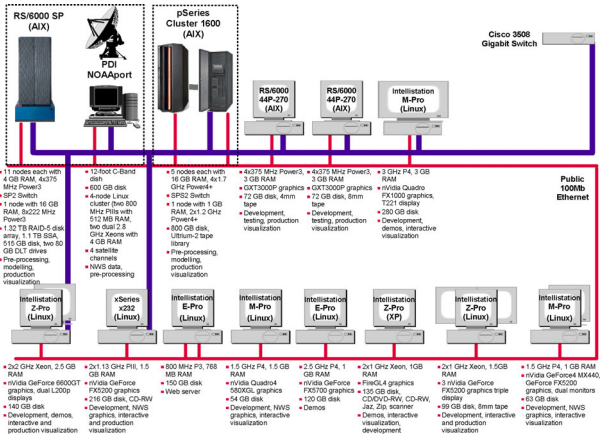

Figure 3. *Deep Thunder* Hardware Environment.

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.

3.2 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. Specialized processing and analysis has been implemented to assure quality control and appropriate utilization of these data in the model pre-processing. The details concerning this approach, support of NAM data in GRIB-1 and GRIB-2 formats and related issues are discussed in Treinish et al, 2005. However, the data and procedural flow of these processes is outlined in Figure 4. Most of them run serially on, although compiler-optimized for IBM Power or Intel Xeon processors. Other aspects related to forecast verification and product visualization are discussed in subsequent sections.

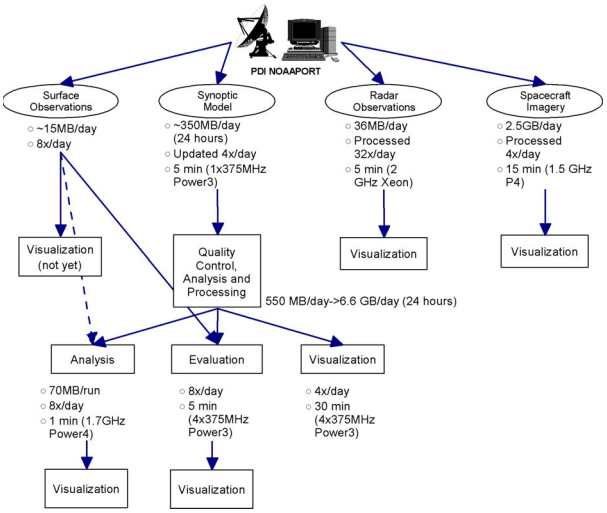

Figure 4. Pre-Processing Procedural and Data Flow.

3.3 Processing

To enable timely execution of the forecast models, which is required for operations, the simulation is parallelized on an high-performance computing system. For this effort, an IBM RS/6000 Scalable Power Parallel (SP) and an IBM pSeries Cluster 1600 are employed. Both are from earlier generations of IBM supercomputer systems, which are in common use at many operational centers for numerical weather prediction or have been upgraded with newer IBM systems built with a similar architecture. These and the current generation are distributed memory MIMD computers, typically consisting of two to 512 processor nodes, that communicate via a proprietary, high-speed, multi-stage, low-latency interconnect. Depending on the flavor of the system, each node has an SMP configuration of two to 64 processors. The older Power3-based systems only supported up to 16-way nodes while current Power5 systems support up to 64-way nodes. In the current implementation for the *Deep Thunder* effort, the SP has eleven nodes of four 375 MHz Power3 processors and one node of eight 222 MHz Power3 processors. The Cluster 1600 has five nodes of four 1.7 GHz Power4 processors and one node with two 1.2 GHz Power4 processors. The latter has a faster interconnect compared to the older SP. Both systems are shown in Figure 3.

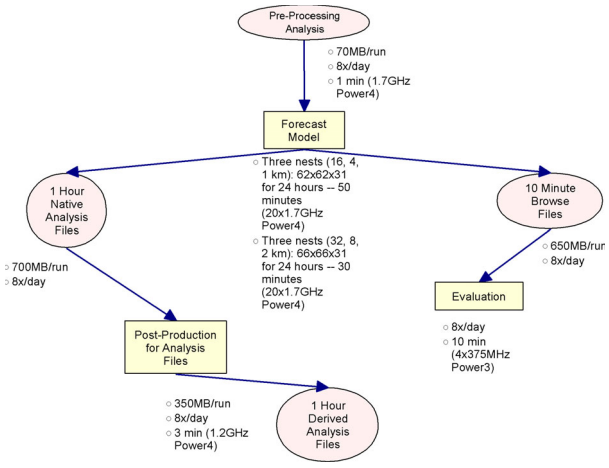

Figure 5. Processing Procedural and Data Flow.

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 interconnect 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 platform for the modelling code. For current operations, forty-two 375 MHz processors (eleven nodes) are used for computing and a single 222 MHz cpu of the remaining node is used for I/O on the SP. Similarly, twenty 1.7 GHz processors (five nodes) are used for computing and a single 1.2 GHz cpu for I/O are used on the Cluster 1600. On average, the latter system has about 1.7 times the throughput of the older, much larger SP system. A typical model run with the 32, 8 and 2 km configuration on the Power4 Cluster 1600 system requires about 30 to 40 minutes to complete a 24-hour forecast. This variation is due to the relative dominance of radiative vs. microphysics calculations, respectively for a particular run. The data and procedural flow of these processes is outlined in Figure 5.

3.4 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, as suggested earlier. 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. In many cases, a "natural" coordinate system is used to provide a context for three-dimensional analysis, viewing and interaction. These visualizations provide representations of the state of the atmosphere, registered with relevant terrain and political boundary maps. This approach for *Deep Thunder* and details of its implementation are discussed in Treinish, 2001.

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 out-

lined schematically in Figure 6. 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 processed to generate derived products and interpolated to isobaric levels from the model terrain-following coordinates.

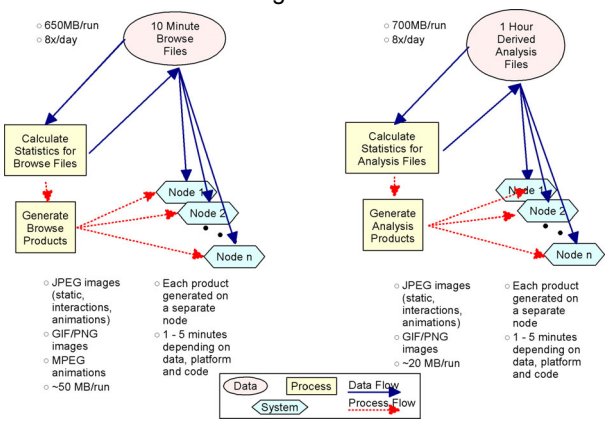

Figure 6. Visualization Post-Processing Procedural and Data Flow.

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). That work is also illustrated schematically in Figure 6. The processes to create individual products (i.e., image or animation files) are split up among the available nodes to run simultaneously. This simple parallelism, including intranode parallelism, enables the independent generation of various products for placement on a web server to be completed in a few minutes.

An approach similar to that used for visualization is employed for forecast verification. After each model run, the results of all three nests combined in a multi-resolution structure (Treinish, 2000) are bilinearly interpolated to the locations of the NWS observing (e.g., metar) stations, whose data are available through the NOAAport receiver. The same approach is applied to the locations of surface observation data acquired from other sources, including the small mesonet operated as part of the *Deep Thunder* project. An analogous process is applied to each NAM grid as part of the automated pre-processing. After the observations corresponding to each model run become available, a verification process is initiated in which these spatially interpolated results are statistically analyzed and compared to parsed and quality-control-checked surface observations. This yields a set of evaluation tables and statistical summaries as well as visualizations for each model run. In addition, similar results from an aggregation of all model runs during the previous week are also calculated. The visualizations are provided on web pages in a manner similar to those generated from the model output. The details of this approach and examples are discussed in Praino et al, 2003.

3.5 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.

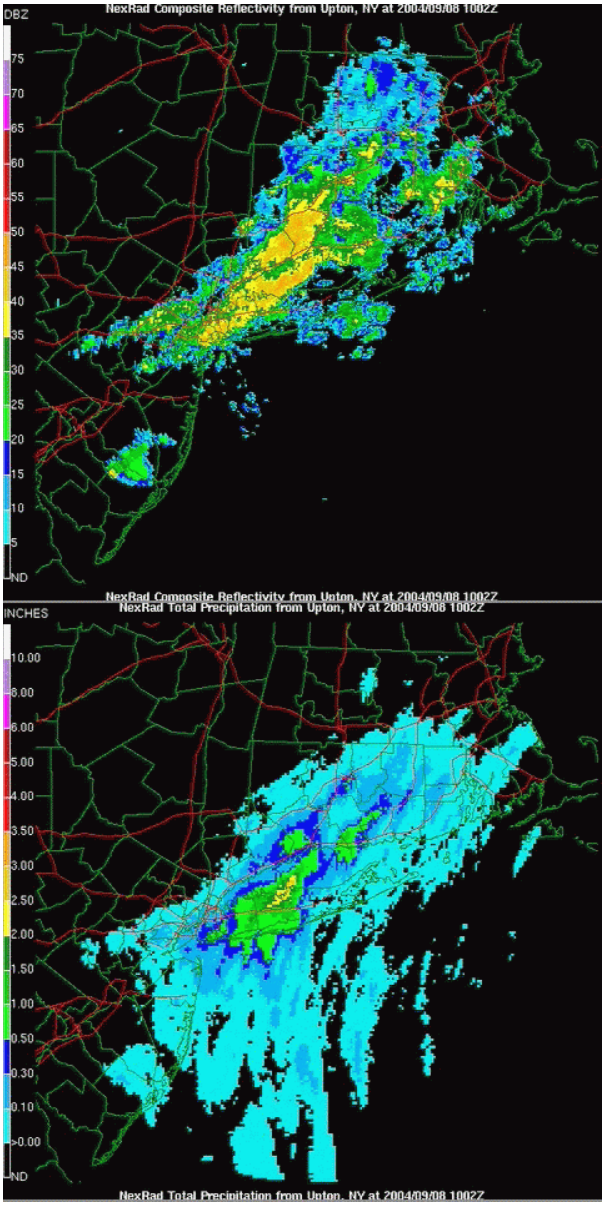

Figure 7. Radar Reflectivity (Top) and Precipitation (Bottom) at 1002Z, 8 September 2004 (Hurricane Frances in the New York City Metropolitan Area).

4. EXAMPLE RESULTS

To illustrate some of the range of capabilities that have been implemented, a few visualization products that *Deep Thunder* can generate automatically are shown herein. These are shown in the context of operational forecasts of two interesting weather events that had significant economic and societal impact, primarily on transportation systems.

4.1 Extratropical Event

Early in the morning of 8 September 2004, the remnants of Hurricane Frances moved into the New York City metropolitan area. The heaviest rainfall occurred in an area stretching from northeastern New Jersey through central Westchester County, NY. Amounts ranged from 2.5 to 15 cm, which caused extensive flash flooding across the region. This led to widespread disruption of transportation systems (e.g., road closures, flooded subways, airport delays). Figure 7 is a snapshot of local radar observations from the nearby NWS Office during the event (0602 local time). Composite reflectivity is shown at the top and estimated accumulated rainfall for the event at that time. One of the rain bands that deposited significant precipitation is clearly seen.

Figures 8 through 12 show different aspects of the *Deep Thunder* model results for a 24-hour forecast for this event. It was initialized using 0Z data from that day and was available at about midnight local time (4Z). Figure 8 illustrates predicted accumulated hourly liquid precipitation as two-dimensional maps for the 4 km and 1 km nests combined in a multi-resolution fashion (Treinish, 2000). Each panel shows the hourly change in rainfall from 9Z to 12Z on September 8, which corresponds to the period of heavy rainfall (predicted and observed). A set of colored contour bands following the legend to

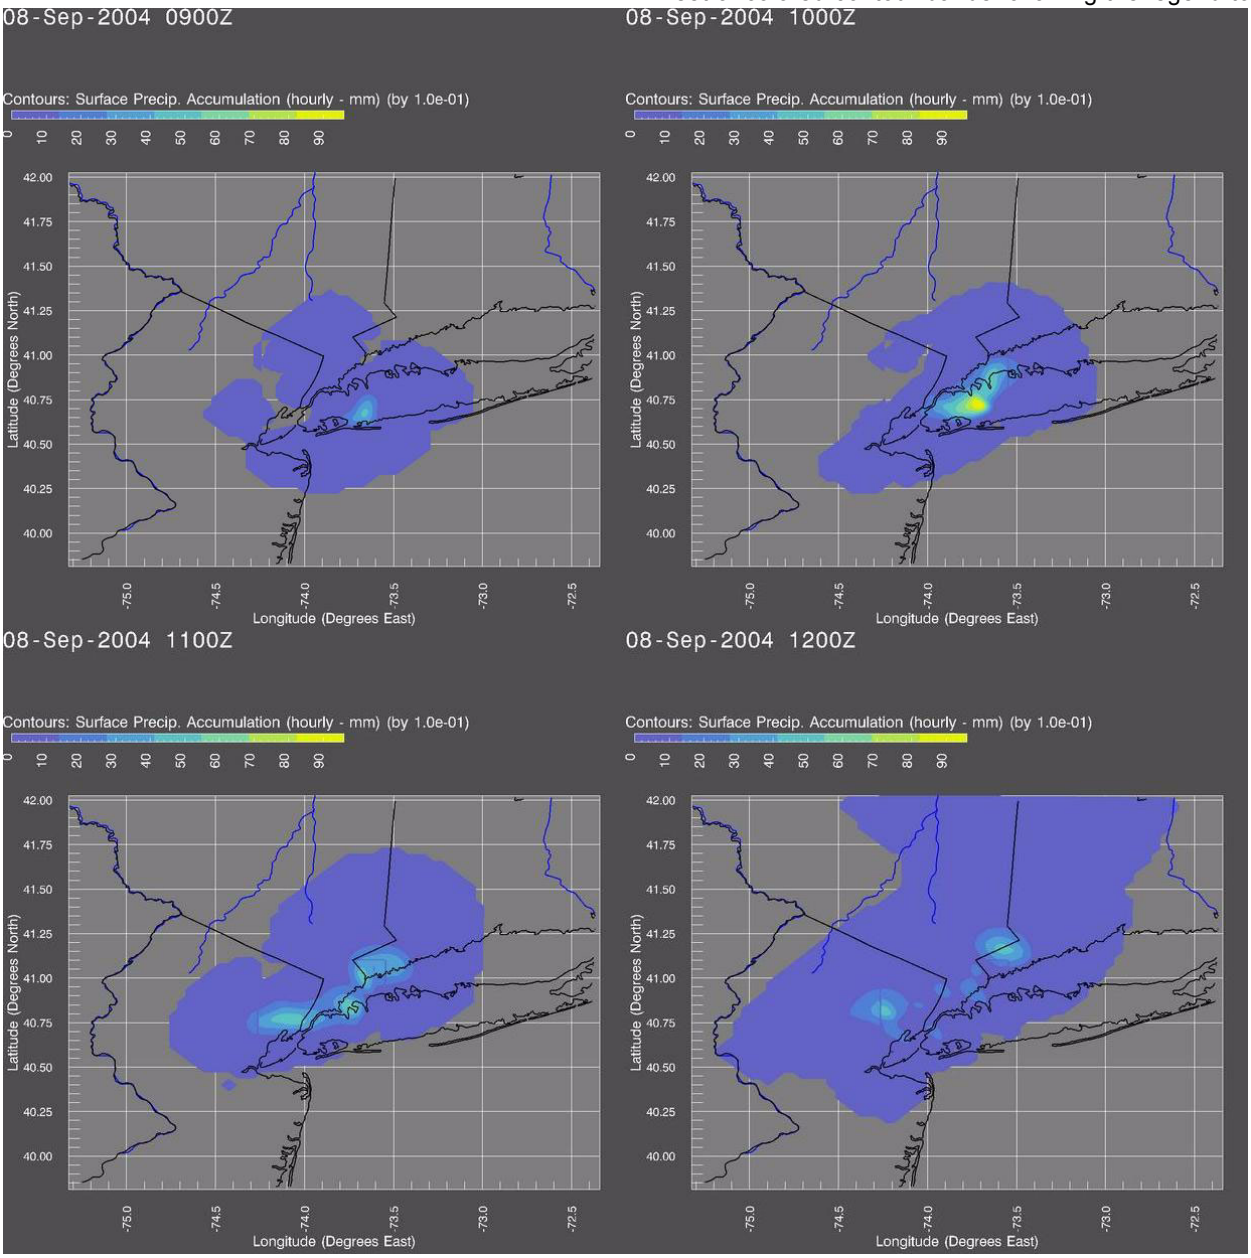

Figure 8. Two-Dimensional Visualization of Predicted Hourly Precipitation from the 4 km and 1 km Nests (Initialized at 0Z) for the Remnants of Hurricane Frances in the New York City Metropolitan Area on 8 September 2004.

the upper left are overlaid with the location of state boundaries and coastlines (black) and rivers (blue). With the exception of Figure 9, animations of these visualizations are available through interactive applications or web browsers. The former also permit flexible viewing and data selection to enable customized presentations (Treinish, 2001).

Figure 9 represents a class of meteogram that is oriented toward interpretation by the non-meterologist. 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 (southern Manhattan in the center of New York City's financial district 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). In addition, the model calculations require some time to "spin-up" the microphysics to enable precipitation. Hence, there will typically be no precipitation in the first hour or two of model results. 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 (compass) direction to which the wind is going. The middle right plot is a colored log-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 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. The forecast for this location shows two significant rain bands, one starting at

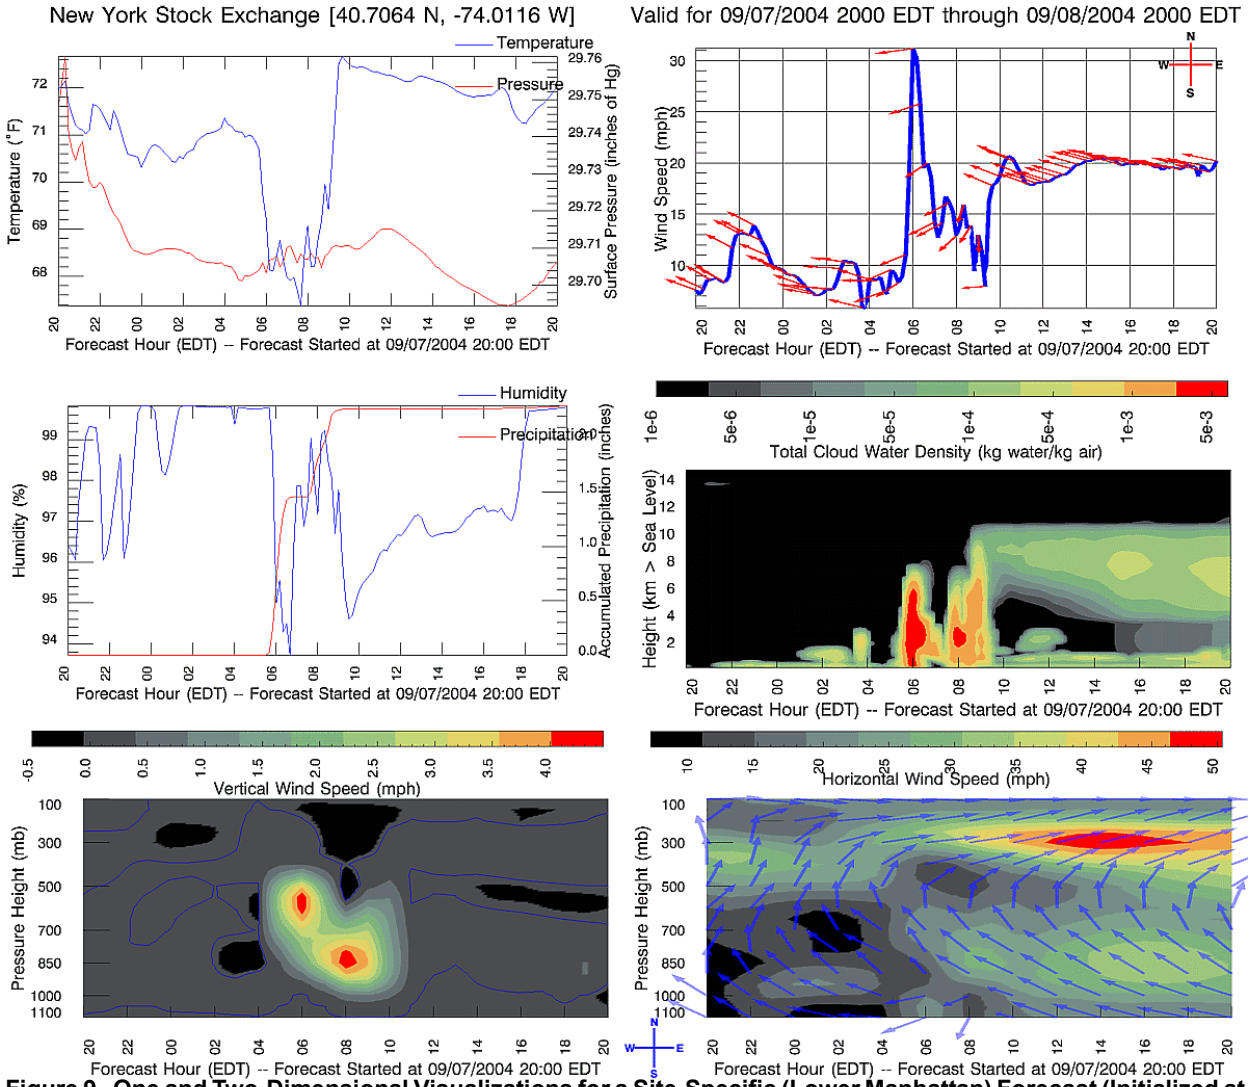

Figure 9. One and Two-Dimensional Visualizations for a Site-Specific (Lower Manhattan) Forecast (Initialized at 0Z) in the 1 km Nest for the Remnants of Hurricane Frances on 8 September 2004.

about 0545 and the other at about 0730 local time, which are consistent with the limited available observations for that area and time.

Figure 10 shows the total accumulated rainfall forecasted for the 24-hour model forecast period starting at 0Z (2000 local time, September 7 to 2000, September 8). Both panels show a terrain map, colored by contour bands of precipitation, where darker shades of blue indicate heavier accumulations. The top presents the 4 km nest while the bottom corresponds to the 1 km nest. In both cases, the model orography is used for the terrain. The maps are marked with the location of major cities or airports as well as river, coastline and county boundaries within the 4 km nest. For the bottom map, major roads are also shown.

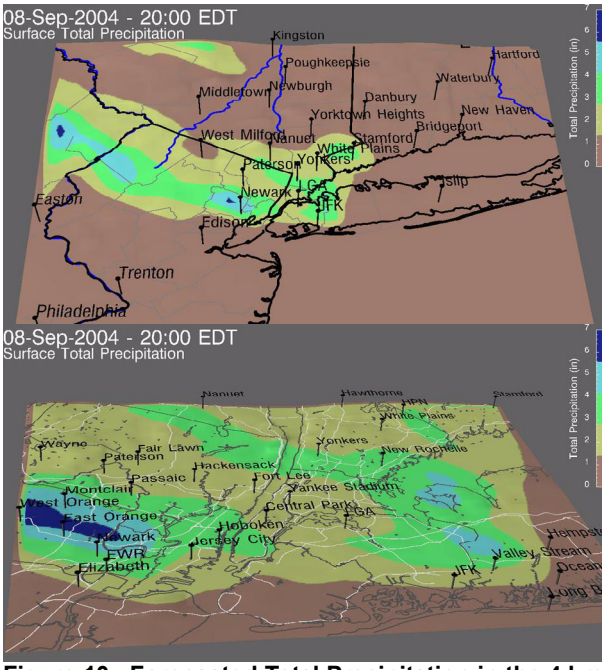

Figure 10. Forecasted Total Precipitation in the 4 km (Top) and 1 km (Bottom) Nests for 8 September 2004.

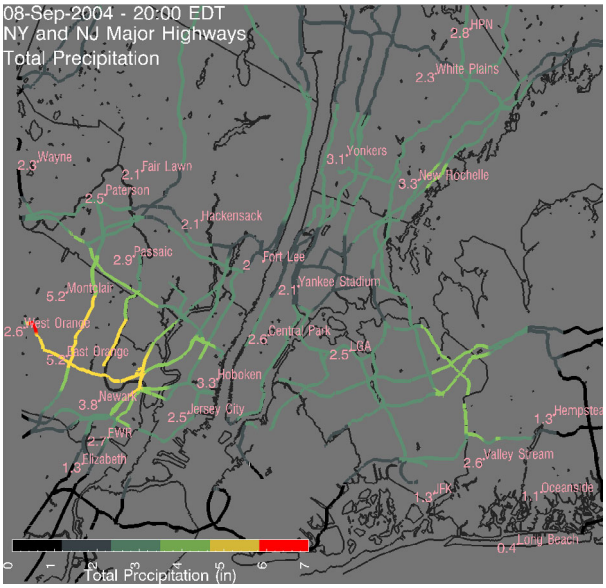

Figure 11. Forecasted Total Precipitation at 1 km Resolution on Major Roads for 8 September 2004.

Figures 11 and 12 are examples of the type of highly specialized visualizations that *Deep Thunder* can produce. Given the flash flooding that occurred along the streets and highways of the region, agencies responsible for road maintenance and operations as well as traffic management can benefit from having the model forecasts visualized in relevant terms. Following the design and implementation principles published earlier (Treinish, 2001 and Treinish and Praino, 2004), Figure 11 presents the accumulated rainfall for the 24-hour forecast from the bottom panel of FIgure 10, interpolated to the local of major roads. In particular, the road locations are extracted from a commercial Geographic Information System and registered in the same coordinate system as the native model output. The interpolated data are reprojected cartographically to minimize linear distortion, when rendered and viewed. The results are then color-contoured following the legend to the lower left. In addition, the forecasted rainfall is shown at specific locations (i.e., towns, airports and landmarks) on the map, along with an overlay of coastlines and political boundaries.

The impact of precipitation may not be the only concern for weather-sensitive decisions from a severe event. Alternatively, consider Figure 12, which only presents forecasted wind information within the 1 km nest. It shows local terrain and water in the model coordinate system. Overlaid on the surface, there are colored arrows indicating predicted winds, with the lighter color being faster winds according to the legend at the lower right. The arrow direction corresponds to the direction to which the wind is flowing. In addition, local political boundaries are overlaid on the map.

The volume is marked with poles at two positions, one of which (at the right) is the same location as the site-specific forecast shown in Figure 9. Using the upper air model wind data interpolated to 21 isobaric levels, the horizontal wind is shown via arrows spaced at 50mb increments vertically along each pole, as a virtual wind profiler. Each of these positions are then used as seed points for calculating wind trajectories based upon the model output. These trajectories are visualized as a set of steady-state stream lines shown as colored ribbons. The arrows and ribbons are colored by horizontal wind speed following the same scale as the surface winds. In addition, the arrow length corresponds to speed.

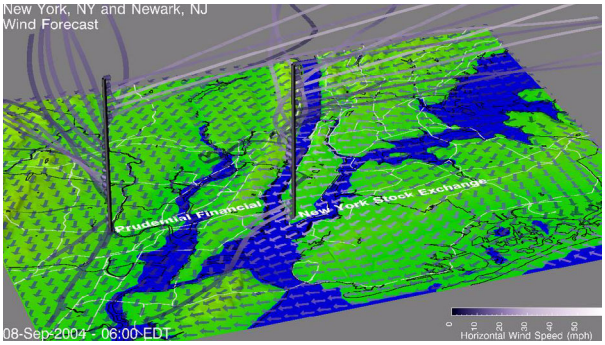

Figure 12. Forecasted Surface and Upper Winds at 1 km Resolution During Extreme Precipitation Event.

4.2 Severe Convective Event

A fast-moving line of late-afternoon thunderstorms occurred along Interstate 95 north of Baltimore, Maryland between 1600 and 1630 local time (2000 to 2030Z) on October 16, 2004. Observers in the area reported heavy rain, zero visibility and “pea-size hail”. The description of the hail suggested that it may have been graupel, especially given the strong cold front with a

sharp temperature gradient that afternoon. In contrast, formation of graupel would have been rather unusual for this region.

The combination of the "hail", rain, sun glare and steam rising off the pavement led to at least 17 multi-car collisions, involving more than 90 vehicles (cars, trucks and one bus). Figure 13 shows some the effects of this event, which sent about 50 people to hospitals. There was widespread disruption of traffic along this heavily travelled highway, which was closed in both directions for several hours. It was the largest multi-vehicle crash in Maryland history.

Figure 13. Vehicle Collisions on Interstate 95, North of Baltimore, Maryland on 16 October 2004.

Figure 14 is a snapshot of local radar composite reflectivity from the nearby NWS Office during the event. A number of convective cells can clearly be seen. The particular time (1609 local time) was when the maximum reflectivity was observed (about 60 dbZ). In addition, the radar information suggested "hail" as large as 1.25 cm.

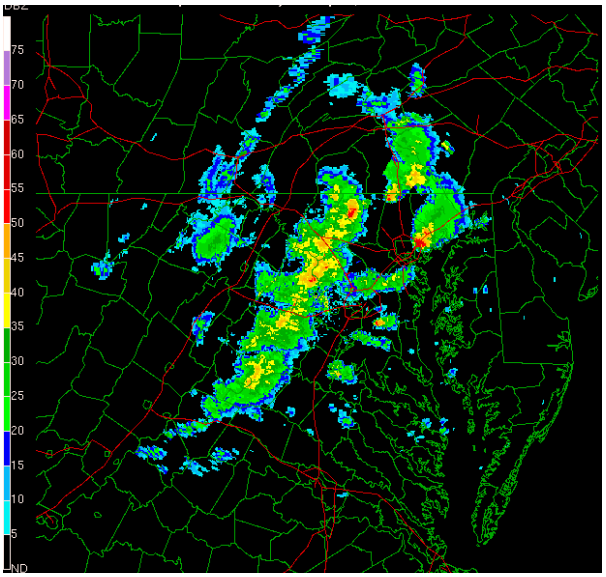

Figure 14. Radar Reflectivity at 2009Z on 16 October 2004 for the Event that Lead to Traffic Disruption.

Figure 15 shows one aspect of the *Deep Thunder* forecast results for this event. The model was initialized using 06Z data from that day and was available at about 0600 local time (10Z). Figure 15 is an example of a qualitative, yet comprehensive, three-dimensional visualization. It is part of an animation sequence that was generated automatically in production for presentation in a web browser. Like Figure 10, it shows a terrain map, colored by a forecast of total precipitation, where darker shades of blue indicate heavier accumulations. The map is marked with the location of major cities or airports as well as river, coastline and county boundaries within the 2 km nest for this model domain focused on the greater Washington, DC and Baltimore, MD metropolitan areas. 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 a translucent cyan surface of forecast reflectivities at a threshold of 30 dbZ. This combination is indicative of the line of thunderstorms associated with strong convection.

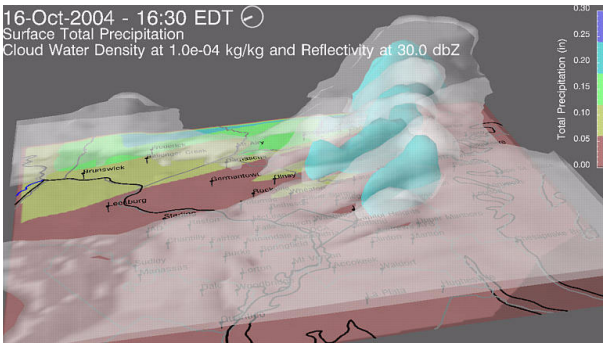

Figure 15. Three-Dimensional Visualization of Severe Thunderstorm Forecast (Initialized at 6Z) at 2 km Resolution for Traffic Disrupting Event on 16 October 2004.

5. DISCUSSION

Although 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 applications related to decision support for risk and environmental impact assessment, especially in response to severe weather events. To aid in that evaluation, a number of forecast products have been made available to several local agencies to assist in their operational decision making with various weather-sensitive problems in transportation and emergency response (e.g., Roohr et al, 2006). In addition, preliminary analysis of *Deep Thunder* forecasts for weather events disrupting local airport operations has been encouraging (Treinish and Praino, 2005). For a number of weather events that affected operational deployment or scheduling of resources during these evaluations, approximate societal or economic value has been determined. This is in conjunction with more generic assessment of the forecasts for winter storms, convective events and extratropical systems (Praino and Treinish, 2004; Praino and Treinish, 2005; Praino and Treinish, 2006, respectively). In particular, consider the two events outlined above.

5.1 Extratropical Event

Other forecasts for the morning of September 8, 2004 indicated "showers and a slight chance of thunderstorms, rain may be heavy at times". At 0748 local time, flash flood watches and warnings were issued throughout the New York City metropolitan area. In contrast, the *Deep Thunder* forecast showed sufficient rainfall to lead to flash floods in several parts of the area. The model results were available about six hours before the heavy rain began and about eight hours before the flood watches and warnings were issued. When the precipitation forecasts, particularly as shown in Figures 8, 10 and 11, are compared to limited, available observations and spotter reports, the timing is good. However, there

appears to be some error in the spatial distribution of the regions of heaviest rainfall (e.g., a bias to the west of about 10 km in northern New Jersey).

5.2 Severe Convective Event

Throughout the day, various forecasts for October 16, 2004 in northeastern Maryland indicated "mostly cloudy with a chance of showers and isolated thunderstorms". Although, the line of thunderstorms started to develop about two hours before the event, as seen from local radar observations, there was no significant change to the local forecasts for the region. Comparison of the *Deep Thunder* results with the radar images shows the model to be in good agreement in timing, intensity and spatial distribution with the exception of the southern portion of the squall line. Operationally, this forecast provided approximately a ten-hour lead-time for the event with initialization data from 14 hours before.

5.3 Overall Utility

Taken at a regional qualitative scale, the results for the aforementioned events as well as others studied in the four distinct geographies are very encouraging with the model showing significant skill predicting the structure, distribution and intensity of convective storms. This has been especially true for severe or unusual events compared to other available forecasts. The model predictions were often available with considerable lead time when compared with other forecast data and with the actual occurrence of the event.

Biases in the model forecasts in terms of timing and/or location appeared to be primarily due to phase errors propagated from inaccurate initial conditions. This class of errors has been discussed in the literature, leading to the suggestion that moving to ensemble solutions is the preferred approach as opposed to higher resolution models (e.g., Zhong et al, 2005 and Roebber et al, 2004).

However, the results of the *Deep Thunder* work to date does suggest that a higher-resolution deterministic forecast can help address gaps in available local weather information for decision-support applications. Applying the computational resources to enable explicit microphysics for all nests, which are integrated for the full forecast period appears to provide realistic information that would be lacking in lower-resolution ensemble forecasts utilizing simpler physics.

On the other hand, the current modelling code is relatively limited compared to newer implementations that can, for example, consider convective precipitation independent of microphysics, more effectively assimilate observations for improved initial conditions or support more accurate representations of the planetary boundary layer.

6. CONCLUSIONS AND FUTURE WORK

The feedback from users coupled with more rigorous verification 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, winter storms, fog 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.

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. These include continuing to refine the quality of the model results, improving the degree of automation, developing new methods of visualization and dissemination, and evaluate the relevance to additional applications.

For example, the aforementioned experimental forecasts for the Miami-Fort Lauderdale area were implemented to evaluate suitability for predictions of local impact of tropical events. An example is illustrated in Figure 16. It shows the total forecasted precipitation for Hurricane Wilma on October 24, 2005 similar to that used for Figure 10. In this case, the intermediate nest at 6 km resolution, covering southern Florida is shown. The model was initialized using 0Z data from that day and was available at about midnight local time (04Z). Although the total amount of predicted rain has a positive bias, particularly along the coast of the Gulf of Mexico, near Fort Myers, the spatial distribution is in reasonable agreement with available estimates of the actual precipitation as shown in Figure 17.

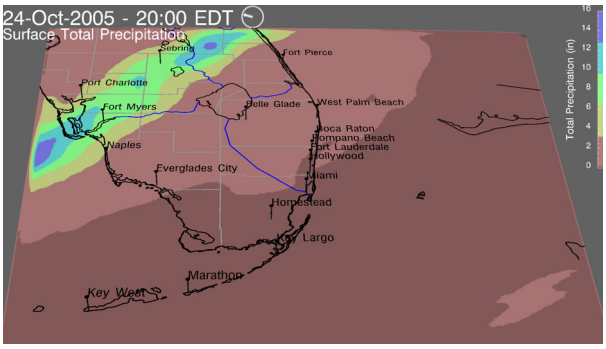

Figure 16. Forecasted Total Precipitation in the 6 km Nest for Hurricane Wilma.

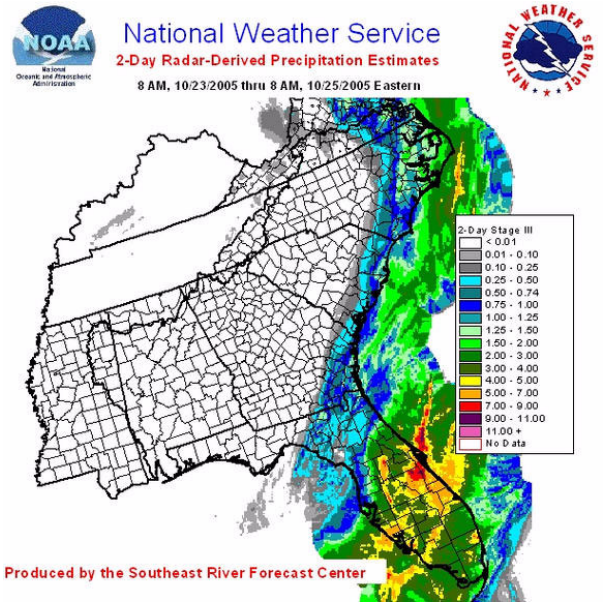

Figure 17. Total Estimated Radar Precipitation for Hurricane Wilma.

To aid in the improvement of overall forecast quality, several steps are planned. While the ability to leverage the availability of full-resolution 12 km NAM results on the AWIPS 218 grid has been implemented within the current pre-processing tools, additional data such as daily global sea surface temperature will be used to

improve initial conditions. In parallel, a modest mesonet is under construction to provide better temporal and spatial sampling in surface observations to aid in forecast verification for the New York City metropolitan area.

Long-term, the viability of the current processing component is limited. In addition to gaps in available modelling capabilities, it is unlikely that further optimization of the underlying code for newer computing platforms will be feasible. Therefore, it is expected that the customized version of RAMS will be replaced with the Weather Research and Forecast Model (WRF). The WRF model has reached sufficient maturity in recent months to address some of capabilities of the current *Deep Thunder* system (Michalakes et al, 2004). Hence, the other components of the system will be adapted to utilize WRF to enable the same class of automated, integrated operations.

The next step after enabling parallel operations with WRF with comparable results, will then consider the utilization of more sophisticated physics and parameterization as well as assimilation of available observations to improve initial conditions. Although all of these changes will result in up to an order of magnitude increase in data production and processing, the current hardware environment does have the capacity to support it.

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

7. ACKNOWLEDGEMENTS

This work is supported by the Deep Computing Systems Department at the IBM Thomas J. Watson Research Center.

8. 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.

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.

de Elía, R. and R. Laprise. *Distribution-Oriented Verification of Limited-Area Model Forecasts in a Perfect Model Framework*. **Monthly Weather Review**, **131**, n 10., pp. 2492-2509.

Edwards, J., J. S. Snook and Z. Christidis. *Forecasting for the 1996 Summer Olympic Games with the SMS-RAMS Parallel Model.* **Proceedings of the 13th 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.

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.

Michalakes, J., J. Dudhia, D. Gill, T. Henderson, J. Klemp, W. Skamarock, and W. Wang. *The Weather Reseach and Forecast Model: Software Architecture and Performance*. **Proceedings of the 11th ECMWF Workshop on the Use of High Performance Computing In Meteorology**, October 2004, Reading, England.

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.

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 1th 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 an Operational Mesoscale Numerical Modelling System in the Northeast U.S. -- Winter 2002-2003.* **Proceedings of the 20th Conference on Weather Analysis and Forecasting/16th Conference on Numerical Weather Prediction**, January 2004, Seattle, WA.

Praino, A. P. and L. A. Treinish. *Convective forecast performance of an operational mesoscale modelling system*. **Proceedings of the 21st International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, January 2005, San Diego, CA.

Praino, A. P. and L. A. Treinish. *Forecast Performance of an Operational Meso-γ-Scale Modelling System for Extratropical Systems*. **Proceedings of the 22nd International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, January 2006, Atlanta, GA

Roebber, P. J., D. M. Schultz, B. A. Colle and D. J. Stensrud. *Toward Improved Prediction: High-Resolution and Ensemble Modeling Systems in Operations*. **Weather and Forecasting**, **19**, no. 5, pp. 936–949, 2004.

Roohr, P.B., L. A. Treinish and A. P. Praino. *Evaluation and Utilization of Meso-γ-Scale Numerical Weather Prediction for Logistical and Transportation Applications*. **Proceedings of the 22nd International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, January 2006, Atlanta, GA

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.

Treinish, L. *Multi-Resolution Visualization Techniques for Nested Weather Models*. **Proceedings of the IEEE Visualization 2000 Conference**, October 2000, Salt Lake City, UT, pp. 513-516, 602.

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 18th International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, January 2002, Orlando, FL, pp. J159-161.

Treinish, L. A. and A. P. Praino. *Customization of a*

Mesoscale Numerical Weather Prediction System for Transportation Applications. **Proceedings of the 20th International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, January 2004, Seattle, WA.

Treinish, L. A. and Praino, A. P. *The Potential Role for Cloud-Scale Numerical Weather Prediction for Terminal Area Planning and Scheduling*. **Proceedings of the 21st International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, January 2005, San Diego, CA.

Treinish, L. A, Praino, A. P and C. Tashman. *Reconstruction of Gridded Model Data Received via NOAAport*. **Proceedings of the 21st International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology**, January 2005, San Diego, CA.

Zhong, S., H.-J. In, X. Bian, J. Charney, W. Heilman, and B. Potter. *Evaluation of Real-Time High-Resolution MM5 Predictions over the Great Lakes Region*. **Weather and Forecasting**, **20**, no. 1, pp. 63–81, 2005.