6.02

Utilizing Local Terrain to Determine Targeted Weather Observation Locations*

Paul E. Bieringer*

Weather Sensing Group
Massachusetts Institute of Technology Lincoln Laboratory
244 Wood Street
Lexington, MA 02420

1. Introduction

Many of the recent conflicts where the United States (US) military forces have been deployed are regions that contain complex terrain (i.e. Korea, Kosovo, Afghanistan, and northern Iraq). Accurate weather forecasts are critical to the success of operations in these regions and are typically supplied by numerical weather prediction (NWP) models like the US Navy NOGAPS, CAOMPS, and US Airforce MM5. Unfortunately the weather observations required to generate accurate initial conditions needed by these models are often not available. In these cases it is desirable to deploy additional weather sensors. The question then becomes: Where should the military planners deploy their sensor resources? This study demonstrates that knowledge of just the terrain within the model domain may be a useful factor for military planners to consider.

For NWP, model forecast errors in mountainous areas are typically thought to be due to poorly resolved terrain, or model physics not suited for use in a complex terrain environment. Recent advances in computational technology are making it possible to run these models at resolutions where many of the significant terrain features are now being well resolved. While terrain can be accurately specified, often the gradients in wind, temperature, and moisture fields associated with the higher resolution terrain are not. As a result, initial conditions in complex terrain environments are not be adequately specified.

Since not all initial condition errors contribute significantly to model forecast error, knowledge of terrain induced NWP model forecast sensitivity may be important when developing and deploying a weather sensor network to support a regional scale NWP model. The terrain induced model sensitivity can provide an indication of which variables in the initial conditions have a significant influence on the forecast and where initial conditions need to be most accurate to minimize model forecast error. A sensor

---

* This work was sponsored by the Department of the Air Force under Air Force Contract No. F19628-00-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the U.S. Government.

* Corresponding author address: Paul Bieringer, MIT Lincoln Laboratory, 244 Wood Street, Lexington, MA 02420-9108; e-mail: paulb@ll.mit.edu

network can then be designed to minimize these errors by deploying critical sensors in sensitive locations, thereby reducing relevant initial condition error without the costly deployment of a high-density sensor network. This is similar to the targeted observation technique first suggested by Emanuel *et al.* (1995), except that in this example the targeted observations would be designed to reduce initial condition error associated with poorly resolved atmospheric features created by the terrain.

This paper is organized as follows. Section 2 contains a brief description of the data collection effort designed to support this study. The experimental design and the specifics of the case used in this study are described in section 3. The analysis and results from both the forward and adjoint simulations are presented in section 4. Section 5 contains a summary of the results, and a brief discussion of their implications.

2. Data collection

Many of the targeted observational studies utilizing real observations have been associated with large field research campaigns. Several examples of this kind of study are the Fronts and Atlantic Storm Track Experiment (FASTEX) Joly et al. (1997), the North Pacific Experiment (NORPEX) Langland et al. (1999), and work conducted by the Tropical Prediction Center (TPC) Aberson and Frankland (1999). In all of these experiments, targeted observations are taken in regions where initial condition errors are suspected to grow rapidly into significant forecast errors. While these studies indicate that this concept shows promise as a practical technique for reducing forecast errors in global, synoptic, and hurricane forecast models (Emanuel and Langland 1998; Aberson and Franklin 1999; Szunyogh et al. 1999), it was often logistically difficult to deploy observational resources in a timely manner to illustrate the benefit of the targeted observations (Langland et al. 1999).

Here the logistical problems associated with collecting suitable data for a targeted observation study were minimized by reducing the overall size of the experiment and conducting it on a mesoscale domain in which a multi-month intensive data collection effort was more manageable. Data were collected over several months, making it possible to assemble a relatively diverse set of spring, summer and fall weather events. This experimental design made it possible to wait for a suitable event to move into the domain instead requiring that observational platforms be deployed prior to the weather event. Consequently it was easier to select a case in which the targeted observation technique can be tested with real observations.

a. Experimental domain and observations

This study is conducted using the terrain, and observational data from the Berkshire Mountain region of western Massachusetts, eastern New York State and southern Vermont. This region was chosen to capitalize on a small network of weather instruments located there and its proximity to the Albany, New York (KENX) Weather Service Radar – 1988 Doppler (WSR-88D) Doppler weather radar. Figure 1 illustrates the terrain elevations in this region, existing observational resources, and target forecast region. The small rectangle in the right center of the image indicates the target forecast region located over the Greylock Valley. The ⊗ symbol in Fig. 1 denotes the location of the KENX WSR-88D radar, the triangles denote the locations of the National Weather Service (NWS) rawinsonde launch sites, and the dots designate the surface observation locations. Although located in western Massachusetts, the Berkshire Mountains provide a complex topographic environment with valley to peak elevation variations of over 2500 feet (760 meters).

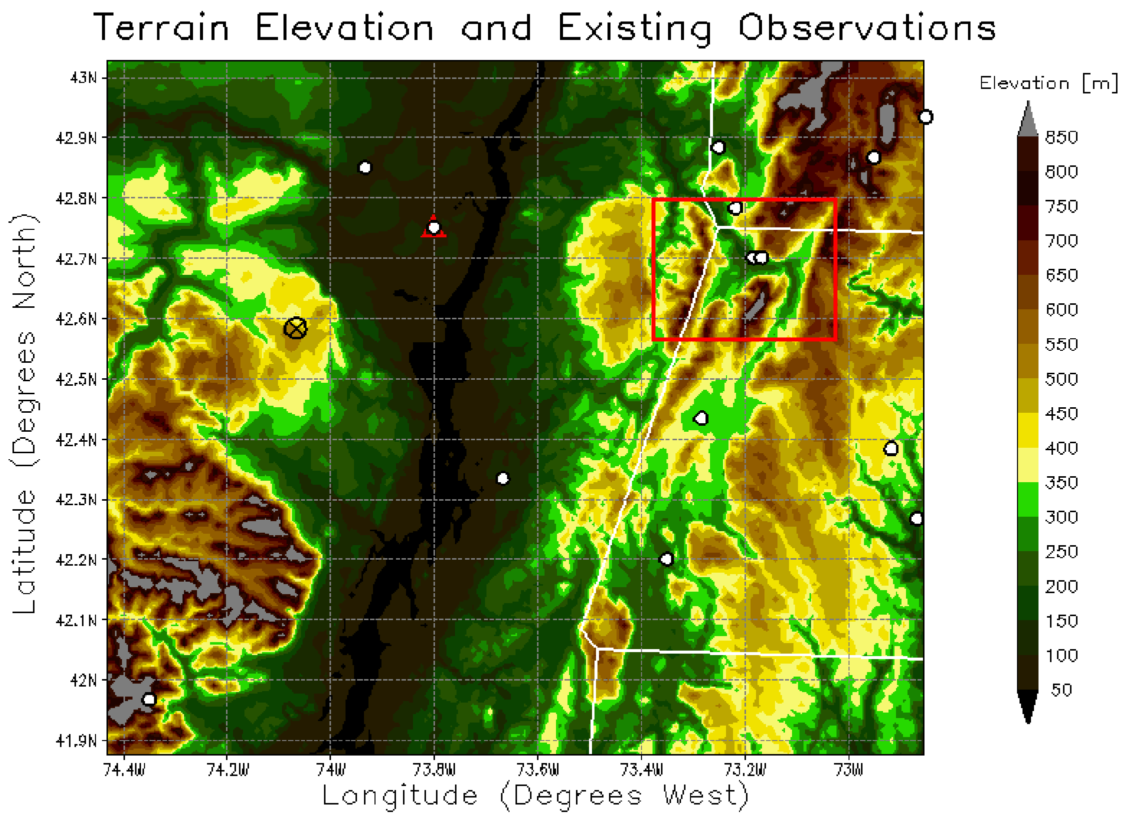

Fig. 1 A horizontal depiction of the terrain heights in western Massachusetts, southern Vermont, and eastern New York State. Terrain is given in meters above sea level. The ⊗ denotes the location of the KENX WSR-88D Doppler weather radar. The dots represent the locations of NWS or Federal Aviation Administration (FAA) surface observation sites, the triangles denote the location of the NWS upper air sites. The small red rectangle indicates the target forecast region.

Figure 2 depicts the region inside the target forecast region. The locations of the surface mesonet sensors deployed in the Greylock Valley of western Massachusetts are overlaid on a zoomed image of the local terrain. This sensor network will be referred to as the Berkshire mesonet. The Berkshire mesonet was established as a complex terrain and mountain weather test bed to support Department of Defense (DoD) weather information system technology development. This observation network was designed to support both a real-time data acquisition/display effort in addition to mesoscale model verification activities and serves as the verification data for this experiment. (Clark and Matthews 2000).

Concurrent with the mesonet data collection, upper air soundings, automated surface observation sites (ASOS) observations, radar data, and initial analysis fields from the Rapid Update Cycle (RUC) and ETA models were collected and archived. The upper air and ASOS observations were collected via the internet from the Albany NWS office.

Archive level II, radar data from the KENX WSR-88D were obtained via a dedicated communication line between Massachusetts Institute of Technology (MIT) Lincoln Laboratory (LL) and the WSR-88D radar products generator. Archive level II radar data provide access to radial velocity, spectrum width, radar reflectivity, and all of the derived products at the radar sampling resolution. The WSR-88D data serves as the targeted observation data for the experiment. Data were collected from June through November of 2001 and although some sensor outages occurred, the network provided a nearly continuous record of the atmospheric conditions in the experimental domain during this period.

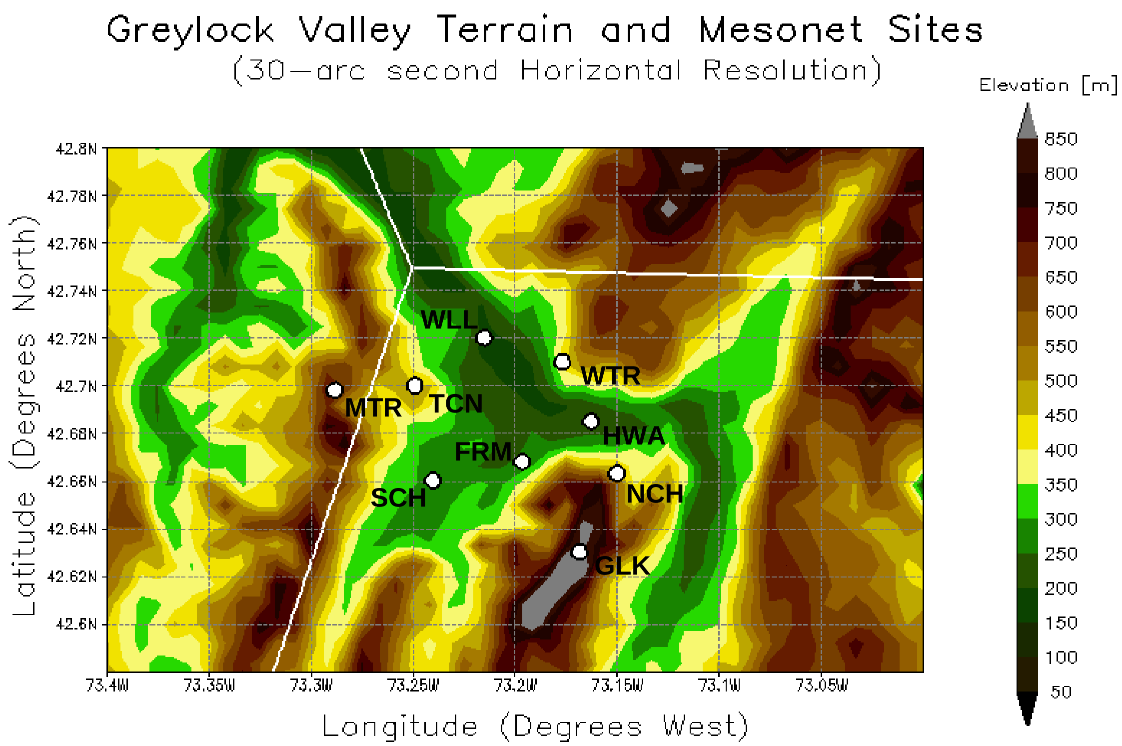

Fig. 2 A magnified map depicting the terrain heights near Williamstown in western Massachusetts, southern Vermont, and eastern New York State. Terrain is in meters above sea level. The circles denote the locations of the surface mesonet observation sites.

3. Background and experiment design

In Bieringer (2003), (hereafter referred to as B-2003) an idealized terrain environment where a lone mountain, surrounded by homogeneous "flat" terrain is used to characterize the impact of the terrain induced initial condition sensitivity. Gradient computations from the fifth generation Pennsylvania State University (PSU) National Center for Atmospheric Research (NCAR) Mesoscale Model (MM5) adjoint model were used to identify the locations and relative magnitudes of initial condition sensitivity. The adjoint sensitivity analysis results give a preliminary indication of initial condition sensitivity and provide information regarding the locations where the surface wind forecast is sensitive to adjustments in the initial conditions. When compared to adjoint simulations where the terrain was removed, the results from simulations using the idealized terrain indicate an increase in adjoint initial condition sensitivity over the elevated terrain.

The second component of the B-2003 study used these adjoint sensitivity results to direct the perturbation of the initial condition fields in a series of forward simulations using the MM5. All perturbations were of the same magnitude and based on the results of the relative sensitivity computations. Again, simulations using the idealized lone mountain terrain were compared to the forecasts from the simulations where the terrain had been removed. The impact of the initial condition perturbations, were measured in terms of maximum differences (impact) and root mean square (RMS) forecast impact on the surface horizontal wind forecast in the target region. The results of B-2003 indicate that an initial perturbation made over elevated terrain will have a larger impact on the forecast than comparable perturbations made to the initial analysis over flat terrain. This study provides a clear example of model sensitivity to initial conditions that can be used to help interpret the more complex results of the present study involving the use of real observational data.

The experimental design of this study is similar to B-2003 and uses both forward and adjoint model simulations. Both the adjoint and forward model components of the experiment use a domain with 19 vertical sigma levels, and a 125 by 125 point, 1 km horizontal resolution model domain centered over the Hudson River Valley in New York. The data used in this case were from 16:00 – 19:00 universal coordinated time (UTC) on October 4th, 2001. This was a case with negligible synoptic scale forcing and relatively uniform environmental flow. On this day, skies were mostly clear, winds were primarily out of the west-southwest at 10-15 knots, and there was no significant precipitation in the region. The vertical profile of horizontal winds was backing with height and varied from southerly at 5 knots at the surface, to westerly at 20 knots at the top of the mountain.

Adjoint sensitivity calculations are used to identify the locations where the surface horizontal wind forecast is sensitive to changes in the initial horizontal wind analysis as a function of simulation length. When using an adjoint model for a sensitivity study, a response function representing the forecast aspect of interest must be specified Errico (1997), Giering and Kaminski (1998), and Zou *et al.* (1997). This study uses vorticity at the lowest model level as the response function since it effectively captures both components of the surface horizontal winds. It is convenient to think of the gradient computations as an indication of initial condition sensitivity for each of the state variables in the model (Errico 1997). These results make it possible to correlate model initial condition sensitivity to the underlying terrain for a given simulation length. Since the locations of initial condition sensitivity pass over both elevated and flat terrain during the simulation, the efficacy of the initial analysis adjustments made in the forward model simulations over both areas can be compared.

The forward model provides surface wind forecasts that are compared to surface mesonet wind observations collected every five minutes. A control simulation, based solely on the initial analysis derived from the 20 km RUC, is used as the baseline forecast. A second "experimental" simulation demonstrating the effects of adjusting the model initial conditions is created by initializing the model with a 20 km RUC analysis adjusted through an objective analysis that incorporated the radar observations. The adjoint sensitivity analysis results link an upstream location where the model is sensitive to initial conditions to the forecast accuracy of a forward simulation of comparable length. As a result, it possible to determine the degree to which terrain variability influences the impact that initial analysis adjustments have on the surface wind forecast accuracy.

In this study, the surface horizontal wind forecast is used as the metric by which this influence is measured. Surface winds are used for several reasons: 1) the *u* and *v* wind components are state variables directly forecasted by the model, 2) the idealized adjoint sensitivity results presented in B-2003 indicate, that local terrain can significantly influence the relative impact that including additional observations can have on surface horizontal wind forecasts, 3) wind speed and direction were among the more reliable measurements made by the Berkshire mesonet and therefore serve as a suitable observational truth against which model forecasts can be evaluated.

4. Analysis and results

The results of the idealized initial condition sensitivity analysis discussed in B-2003 suggest that improvements made to the initial analysis over elevated terrain will have a larger impact (presumably positive) on forecast accuracy than comparable improvements made over flat terrain. This study builds on this work by examining the relationships between elevated terrain and the impact that the addition of radar derived wind observations had on forecast accuracy.

a. Doppler radar data analysis

WSR-88D Doppler radar data from KENX radar in Albany, NY are used to measure a vertical profile of the horizontal winds upwind of the forecast verification region. These data are then used to adjust the initial wind analysis for the experimental simulation. The radar derived vertical wind profile from 16:01 UTC was used to adjust the initial wind analysis used by the experimental MM5 simulation. The MM5 Little_R objective analysis preprocessing software was used to create the new initial wind analysis. Little_R starts with a 20 km RUC background analysis and then objectively incorporates the vertical wind profile. No variables other than the *u* and *v* winds were modified by Little_R. The objective analysis was configured such that the radar wind observations would adjust the initial winds over the Hudson River Valley and the mountains on the New York, Massachusetts border (Fig. 3). Since the adjusted analysis is based on background RUC analysis, the new wind analysis is modified only at the altitudes and locations where the radar derived winds differ from the background (Figs. 3 and 4). For additional details regarding the Little_R objective analysis scheme used in this study, the reader is encouraged to consult Dudhia *et al.* (2000).

b. Mesonet data analysis

Surface wind observations collected from the Berkshire mesonet serve as ground truth against which the accuracy of the surface wind forecasts can be measured. These data are used in a separate objective analysis that creates a gridded surface wind analysis. Again, the MM5 Little_R data preprocessing software is used to produce this analysis. Here Little_R uses the wind field from the RUC model, interpolated to the 1 km MM5 grid, as the background field. Ground truth wind analyses based on the surface wind observations and background winds were generated from 16:05 – 21:00 UTC at 5 minute intervals. A visual inspection of the objective analyses indicate that they provide a reasonable depiction of the surface winds when compared to the station observations (Fig. 5).

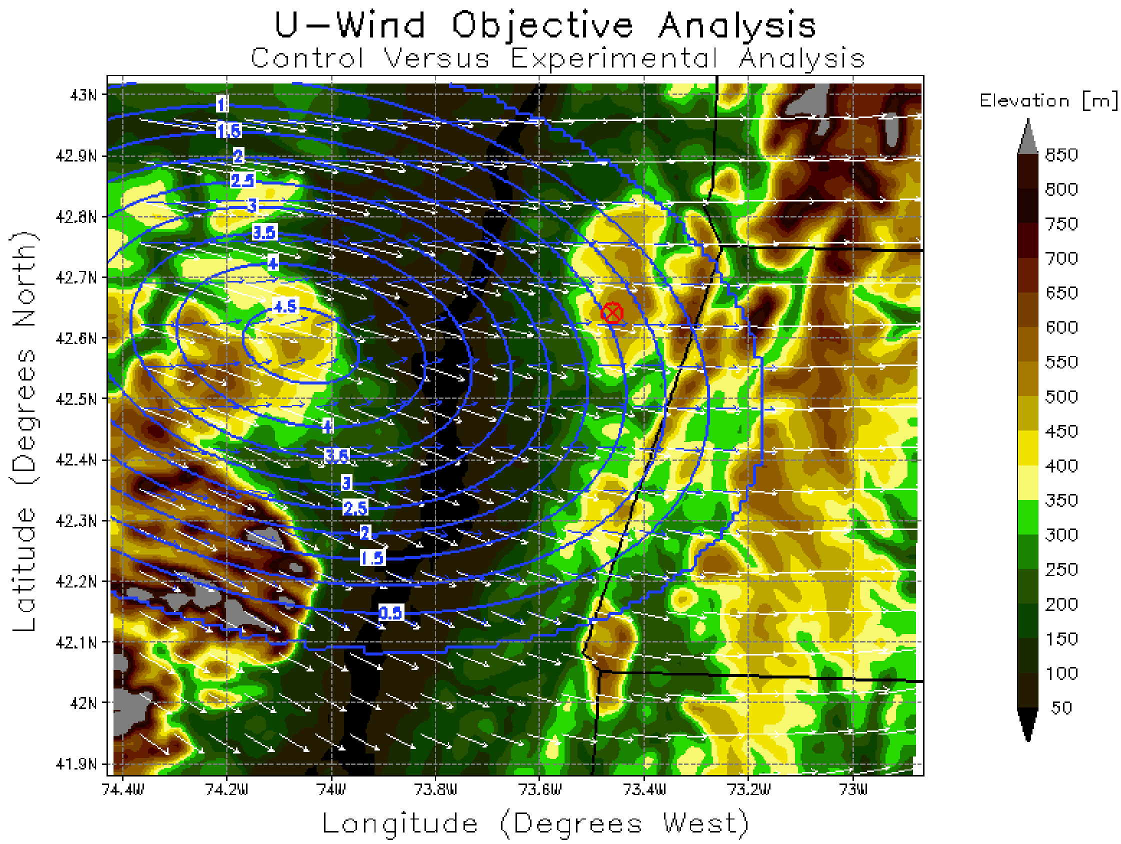

Fig. 3 A cross-section at 875 mb of the horizontal wind analyses used to initialize the control and experimental simulations. The white arrows represent the background wind analysis used in the control simulations. The blue arrows represent the wind analysis after the addition of the radar derived horizontal wind observations. The blue contours represent the difference between the two u-wind analyses in m/s, and illustrate the influence radius of the Cressman objective analysis. The ⊗ denotes the location where the vertical wind profile shown in Fig. 4 was taken.

c. Adjoint sensitivity analysis

The MM5 adjoint model is used to determine locations of initial condition sensitivity for a series of simulation lengths extending from 10 to 210 minutes in 10-minute intervals. The adjoint simulations use the same model domain and initial conditions used by the forward model control runs. This analysis uses a vorticity response function defined at the lowest model level in the forecast verification region. The *u* wind relative sensitivity calculations from the adjoint model provide the geographic locations where the upstream winds will influence the surface wind forecast in the Greylock Valley. A series of 21 adjoint sensitivity simulations provided initial condition sensitivity results resembling those shown in Fig. 6. The center locations for the regions of initial condition sensitivity are subjectively determined by reviewing the adjoint sensitivity results. Because this analysis uses real data to adjust the background wind analyses, it is important to also consider the locations of adjoint sensitivity in context with where the analysis is adjusted. In this case, the incorporation of the radar data results in an adjustment of the background winds between 900 and 700 mb (Fig. 4). Since all of the forecast modification is due to these adjustments, it was critical to examine the adjoint sensitivity patterns at these levels. The product of this analysis is a backward trajectory of points illustrating the approximate centers of regions of adjoint initial condition sensitivity of the surface wind forecast in the Greylock Valley that correspond to where the analysis was adjusted (Fig. 7). This analysis becomes difficult for adjoint simulations longer than 200 minutes because the regions of sensitivity become more diffuse and begin to encounter the domain boundary. This problem begins to manifests itself in the two points furthest to the southwest in analysis, and therefore should be viewed with caution (Fig. 6).

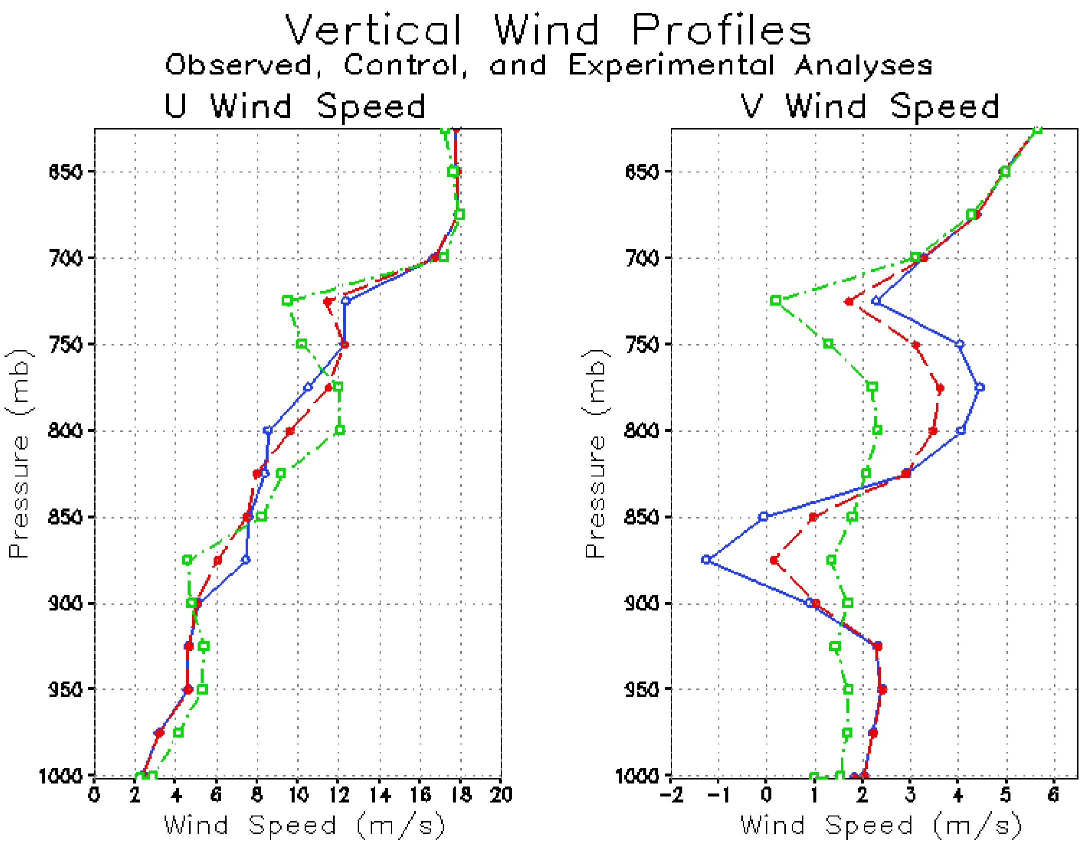

Fig. 4 Vertical wind profiles from the 16:00 UTC initial wind analyses used in both the control and experimental simulations. The wind profiles are taken at the radar location and over the elevated terrain southwest of the Greylock valley at 42.64º N, 73.46º W. The blue line represents the winds without the addition of any observations. The red dashed line represents the wind profile after it was adjusted using the radar observations. The green dot-dashed line represents the observed wind profile from the radar VAD algorithm.

d. Forecast sensitivity analysis

The forward model forecast sensitivity analysis utilizes two model simulations. The control simulation uses only the RUC analyses between 16:00 – 21:00 UTC on October 4th, 2001 to provide the initial and boundary conditions. This simulation provides the surface wind forecasts that serves as the control forecast against which improvements in all experimental forecast are measured. The initial and boundary conditions for the experimental simulation use same RUC analyses as a background field; but as described earlier, are combined with the radar derived vertical wind profile to provide a more accurate representation of the upstream wind fields at the time of model initiation. The surface wind forecasts from the experimental simulation are compared to forecasts from the control simulation to measure the forecast improvement or degradation that results from the inclusion of the radar observations in the initial analysis.

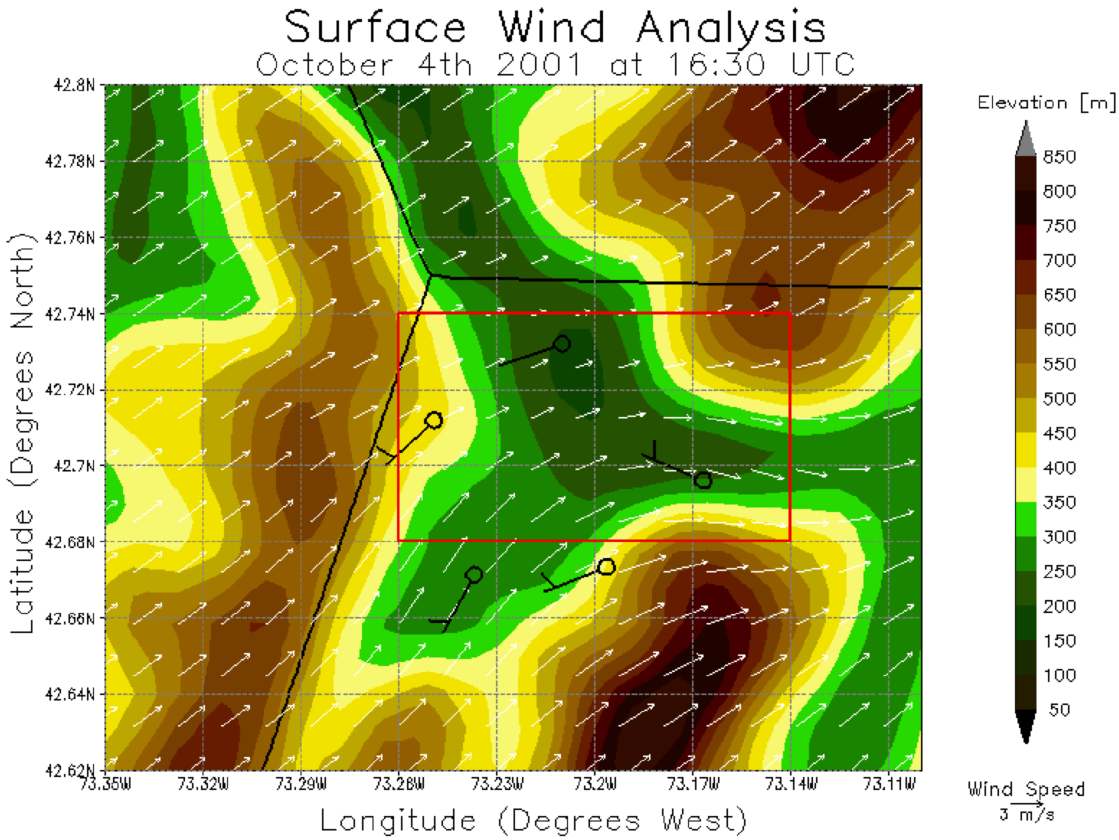

Fig. 5 The gridded surface wind analysis and surface wind observations from the Berkshire mesonet valid at 16:30 UTC on October 4th, 2001. This gridded wind analysis is used as the ground truth against which the accuracy of the MM5 surface wind forecasts are measured. This image is characteristic of the surface wind analyses used. The red rectangle illustrates the forecast verification region used in the study.

Surface wind observations were available every 5 minutes from 16:00 – 21:00 UTC and forecast accuracy was characterized at each observation time by the root mean square errors in the wind forecast at the lowest model level compared to the gridded surface wind analysis. The gridded analysis of these observations (similar to the one shown in Fig. 5) were compared to corresponding surface wind forecasts from the control and experimental simulations. Both components of the wind were examined. As in the idealized initial condition sensitivity experiments discussed in B-2003, the *u*-wind forecast indicates a positive response to the addition of upstream observations, while the *v*-wind forecast shows little or no response. Throughout the majority of the simulation, the inclusion of the radar observations in the initial analysis results in a decrease in *u*-wind RMS forecast error. The RMS error in the *v*-wind forecasts from the control and experimental simulations were very similar in most of the forecasts (Fig. 8).

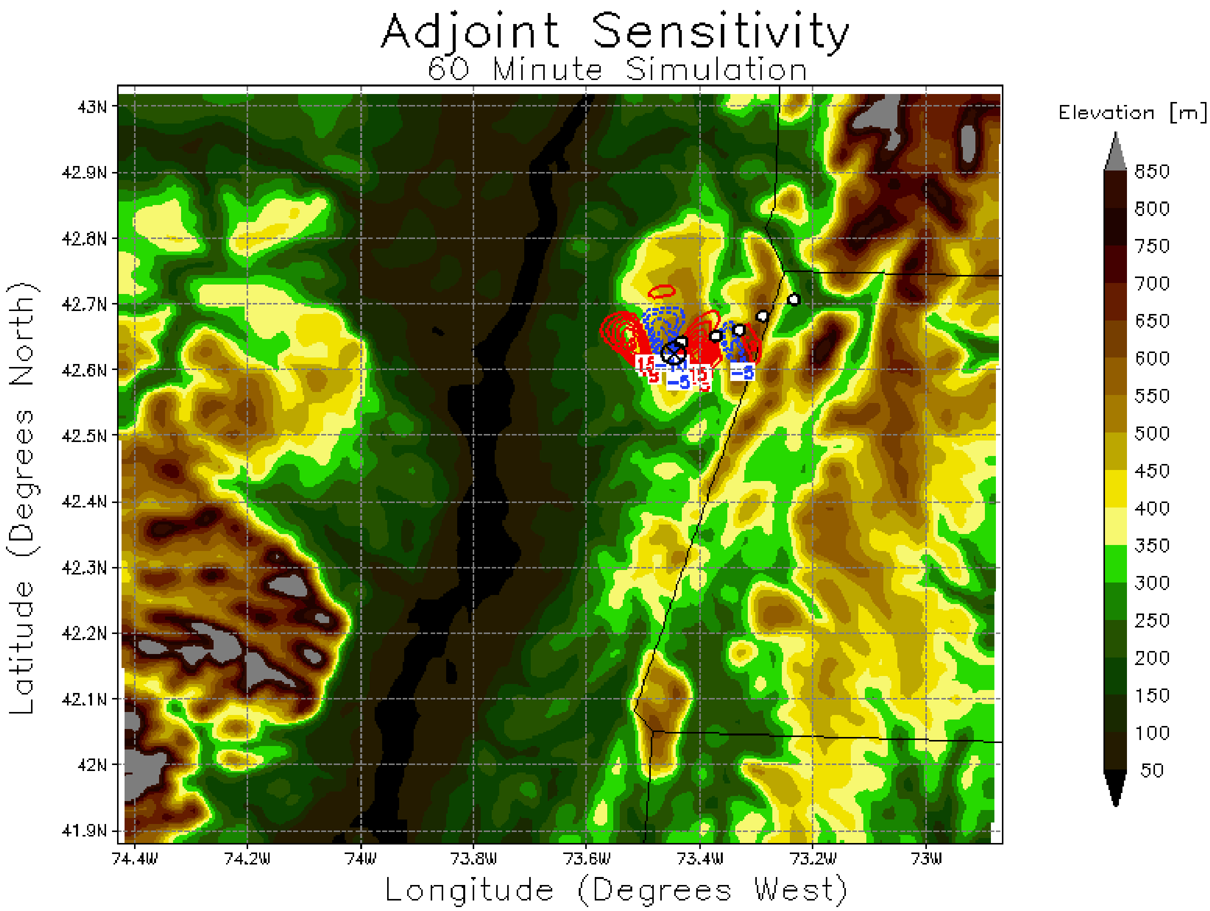

Fig. 6 A horizontal cross-section of adjoint sensitivity for a 60 minute simulation. Terrain is given in meters above sea level. The red and blue contours represent adjoint sensitivity of the u-wind component. The ⊗ denotes the location of the center of the adjoint sensitivity for the 60 minute simulation and the ↄ denote the center locations of adjoint sensitivity for simulations ranging from 10 – 50 minutes in 10 minute intervals.

Figure 8 demonstrates that the addition of the upstream radar observations decreases the error in the short-term surface wind forecasts in the Greylock Valley. The absolute value of the forecast improvements are relatively small, a likely consequence of the light environmental winds on this day; however, in relative terms, the accuracy improvements were on the order of 10%. Forecast improvements of this magnitude are consistent with other targeted observation studies which report forecast accuracy improvements that range from 10% – 20% (Szunyogh et. al. 2002). Figure 8 also illustrates that the amount of forecast error improvement may be linked to presence or lack of elevated terrain in the region of initial condition sensitivity. This is more evident when examining the differences between the RMS forecast errors of the control and experimental simulations plotted versus the simulation length (Fig. 9). Forecast improvement is defined as: (RMS error in the control forecast – the RMS error in the experimental forecast). The pattern of the forecast improvement roughly coincides with areas of high terrain variability below the center of the adjoint sensitivity regions. The peak forecast improvement coincides with simulations where the adjoint initial condition sensitivity was present over areas where the terrain was variable. From 160 to 210 minutes, the forecast error improvements again increase slightly. The terrain corresponding to simulations of this length again shows more variation. Beyond 210 minutes, the initial condition sensitivity centers are too diffuse to identify or are influenced by the boundary conditions and therefore not suitable for this analysis.

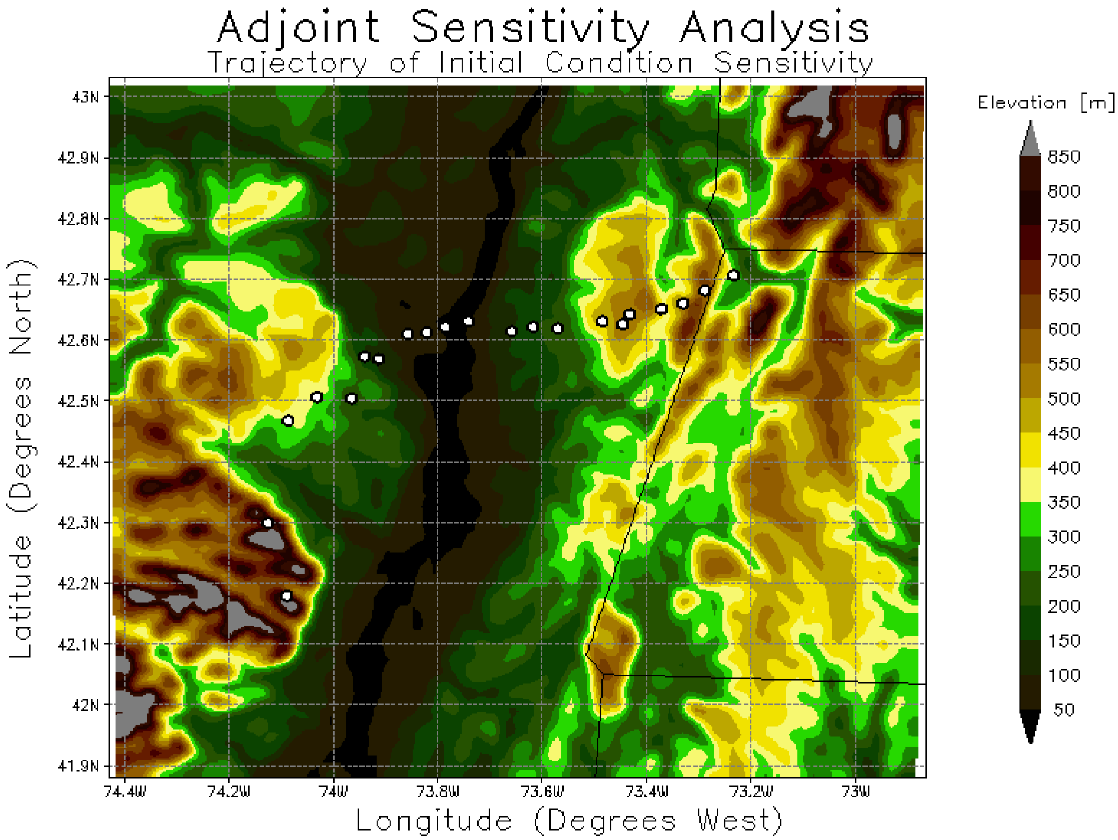

Fig. 7 Center points of the adjoint initial condition sensitivity regions. The dots designate the center locations for the regions of wind adjoint sensitivity in the simulations ranging from 10 to 210 minutes in 10-minute intervals. As the length of the adjoint simulation increases, the initial condition sensitivity tends to be located further west of the Greylock Valley.

5. Summary and Conclusions

The surface wind forecast in a small mountain valley has been used to evaluate the connection between the forecast impact resulting from the addition of observations over varied terrain. Forecast error was measured by comparing the model forecast with a surface objective wind analysis based on the wind observations from mesonet stations in the Greylock Valley. Forecast improvements were then determined by comparing the error analysis from the experimental simulation containing additional observations to the results of the error analysis from the control forecast which used no additional measurements. The geographic locations of initial condition sensitivity for a corresponding simulation length (illustrated in Fig. 7) were then used to determine if the forecast responses were due to the initial wind analysis adjustments made over the elevated terrain or over the flat terrain. This made it possible to characterize the influence that terrain had on the efficacy of the adjusted initial analysis in the sensitive regions.

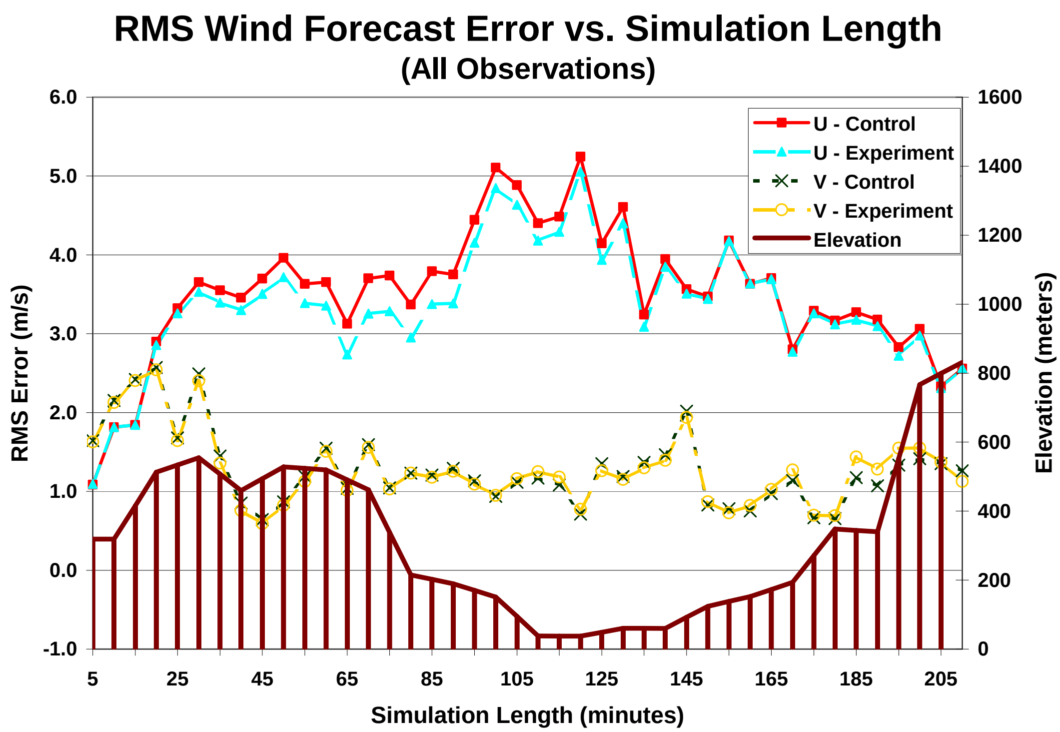

Fig. 8 RMS forecast error in the *u* and *v* winds as a function of simulation length. The control and experimental simulation forecast errors are red and blue respectively for the *u*-wind forecast and green and yellow for the *v*-wind forecasts. Terrain elevations are based on values that correspond to the centers of adjoint initial condition sensitivity illustrated in Fig. 7. Throughout the simulation, the experimental forecast, based on an adjustment of the wind field at a single time, typically indicates improved *u*-wind forecast accuracy, while little or no improvement is evident in the *v*-wind forecast.

This study has confirmed the findings of B-2003 and demonstrated with real observations that terrain variability can be a significant element of initial condition sensitivity. The art of observing the atmosphere primarily consists of capturing gradients in the state variables that drive its processes. In models that use a terrain following vertical coordinate, sharp gradients in the state variables often occur in the vicinity of terrain variations. While this is a realistic representation of the atmosphere within the confines of the model, this study suggests that the forecast is more sensitive to variations in the initial atmospheric variables in these areas, than in regions with less variable terrain. One explanation is that in these areas more information is being represented by fewer grid points. When the terrain elevation varies sharply, a set of grid points on the single terrain following model level will be influenced by the atmospheric conditions both adjacent and above its location. In contrast, a set of grid points on a similar vertical model level in a region with little terrain variability will be primarily influenced only by conditions horizontally adjacent to them. In essence, more atmospheric information is packed into each grid point when the terrain elevation varies than when it is homogeneous. Consequently, improvements to the initial analysis in these regions appear to have a larger impact on forecast accuracy than comparable improvements made elsewhere.

This has potential implications with regards to the deployment of weather sensors and weather sensor networks to support US military operations. In many cases it is difficult to get suitable weather observations to initialize the high-resolution NWP models used to support military operations. When additional observational resources can be deployed, the question then becomes: Where should a sensor be deployed that will provide a significant improvement in forecast accuracy? Clearly, in any scenario there are always locations which will exhibit more sensitivity to initial conditions than others, and often these areas change as the environmental conditions change. This study indicates that initial condition sensitivity is also linked to stationary elements within the model, such as terrain. Consequently, just a simple knowledge of the local terrain may be a good starting point when military planners are making sensor deployment decisions.

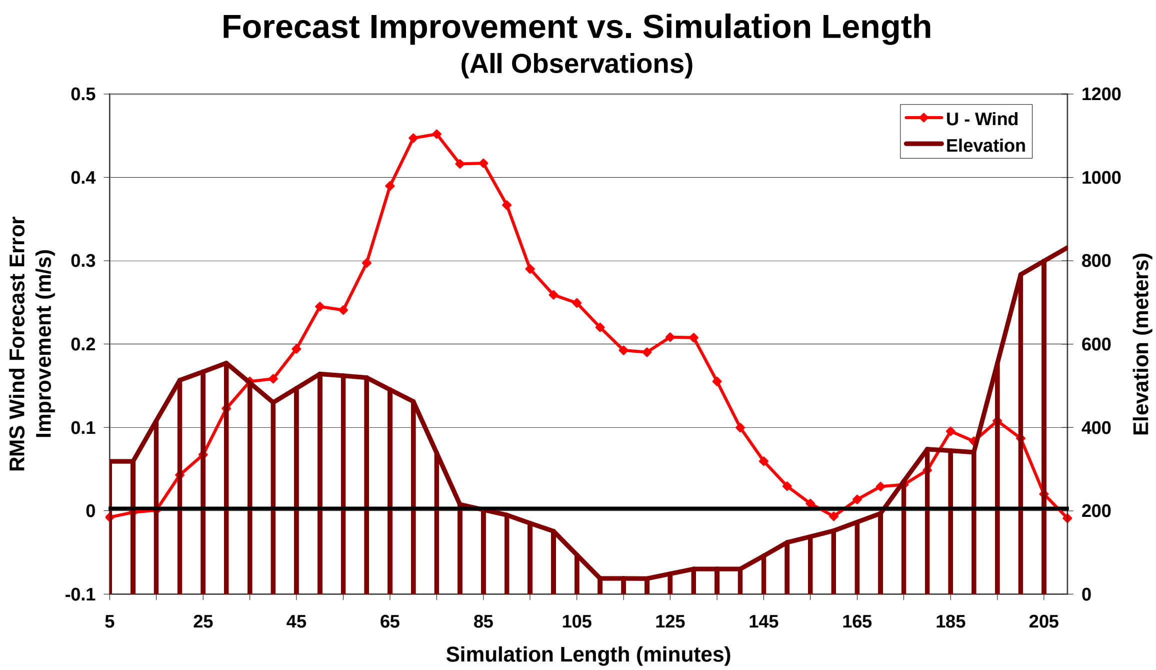

Fig. 9 RMS *u*-wind forecast improvement and terrain versus simulation length. RMS forecast improvement, defined as: (control forecast error – experimental forecast error) illustrates the differences between the *u*-wind RMS error lines in Fig. 8. Terrain elevations are based on values corresponding to the locations identified as the centers of the regions of adjoint initial condition sensitivity illustrated in Fig. 7. The forecast improvement peaks and decreases, based on the adjustment of the winds at a single time, suggest that terrain influences the impact of the additional wind observations included in the experimental simulations.

6. References

Aberson, S. D., and J. L. Franklin, 1999: The impact on hurricane track and intensity forecasts of GPS dropwindsonde observations from the first-season flights of the NOAA Gulfstream-IV jet aircraft. *Bull. Amer. Met. Soc.*, **80**, 421–428.

Bieringer, P. E., 2003: The effect of topography on the initial condition sensitivity of a mesoscale model. *10th Conference on Mesoscale Meteorology, 23-27 June, 2003, Portland, Oregon, American Meteorological Society, Boston MA.*

Clark, D. A., and M. P. Matthews, 2000: An integrated weather sensor testbed for support of theater operations in areas of complex terrain. *Battlespace*

Atmospheric and Cloud Impacts on Military Operations, 25-27 April, 2000, Fort Colins, CO.

Dudhia, J., D. Gill, Y-R. Guo, K. Manning, W. Wang, 2000: *PSU/NCAR Mesoscale modeling system tutorial class notes and user's guide: (MM5 Modeling system version 3).* Available from the National Center for Atmospheric Research, Boulder, CO, 336 pp.

Emanuel, K., D. Raymond, A. Betts, L. Bosart, C. Bretherton, K. Droegemeir, B. Farrell, J. M. Fritsch, R. Houze, M. L. Mone, D. Lilly, R. Rotunno, M. Shapiro, R. Smith, and A. Thorpe, 1995: Report of the first prospectus development Team of the U.S. weather research program to NOAA and the NSF. *Bulletin of the American Meteorological Society,* **76,** 1194 – 1208.

Emanuel, K. A., and R. Langland, 1998: FASTEX adaptive observations workshop. *Bulletin of the American Meteorological Society,* **79,** 1915 – 1919.

Errico, R. M., 1997: What is an adjoint model. *Bulletin of the American Meteorological Society,* **78,** 2577 – 2591.

Giering, R., T. Kaminski, 1998: Recipes for adjoint code construction. ACM transactions on Mathematical Software, 24, No. 4, 437-474.

Joly, A., D. Jorgensen, M. A. Shapiro, A. Thorpe, P. Bessemoulin, K. A. Browning, J-P. Cammas, J-P. Chalon, S.A. Clough, K. A. Emanuel, L. Eymard, R. Gall, P. H. Hildebrand, R. H. Langland, Y. Lemaître, P. Lynch, J. A. Moore, P. O. G. Persson, C. Snyder, and R. M. Wakimoto, 1997: The Fronts and Atlantic Strom-Track EXperiment (FASTEX): Scientific objectives and experimental design. *Bulletin of the American Meteorological Society,* **78,** 1917-1940.

Langland, R. H., Z. Toth, R. Gelaro, I. Szunyogh, M. A. Shapiro, S. J. Majumdar, R. E. Morss, G. D. Rohaly, C. Velden, N. Bond, and C. H. Bishop, 1999: The north pacific experiment (NORPEX-98): Targeted observations for improved North American weather forecasts. *Bulletin of the American Meteorological Society,* **80,** 1363-1384.

Szunyogh, I. Z., Toth, Z. Majumdar, R. Morss, C. Bishop, and S. Lord, 1999: Ensemble-based targeted observations during NORPEX. Preprints: *3rd Symposium on Integrated Observing Systems*, 10 – 15 January 1999, Dallas, TX, American Meteorological Society, Boston MA, 74 – 77.

Szunyogh, I. Z., Toth, Z., Zimin, A. V., Majumdar, S. J., Persson, A., 2002: Propagation of the effect of targeted observations: the 2000 winter storm reconnaissance program. *Monthly Weather Review,* **130,** 1144-1165.

Zou, X., F. Vandenberghe, M. Pondeca, and Y.-H. Kuo, 1997: Introduction to adjoint techniques and the MM5 adjoint modeling system. *NCAR Technical Note.* NCAR/TN-435-STR, 107 pp.