14A.1

DYNAMICALLY ADAPTIVE NUMERICAL WEATHER PREDICTION: MODELS, OBSERVATIONS AND CYBERINFRASTRUCTURE RESPONDING TO THE ATMOSPHERE

*[1,2]Kelvin K. Droegemeier and [2]Yunheng Wang
[1]Center for Analysis and Prediction of Storms and [2]School of Meteorology
University of Oklahoma
Norman, Oklahoma

1. MOTIVATION AND GOAL

Since the dawn in 1955 of operational numerical weather prediction (NWP) in the United States (Shuman 1989), forecast models have, for the most part[1], been run on fixed time schedules, and in fixed configurations (e.g., geographic region, grid spacing), using dedicated computational resources. This strategy has numerous advantages for forecasters, particularly because it affords an ability, via repeated examination of results from the same model configuration over long periods of time, to identify and account for situation-dependent model idiosyncrasies.

Ironically, such a static framework that is independent of the weather occurring or being predicted is incongruent with the dynamic nature of the atmosphere, especially on the mesoscale, where events such as tornadoes, flash floods, hail, lightning, intense straight-line winds and localized winter storms are characterized by a high degree of locality as well as rapid onset and evolution. In light of operational models now boasting near-cloud-resolving grid spacing (e.g., Kain et al. 2008), a potentially more effective and, indeed, natural approach to operational NWP involves having forecast models and their associated infrastructure (data acquisition, storage and assimilation systems, IT networks, high performance computers, and data output frameworks) change their configuration automatically, or with minimal human intervention, to produce the "best" prediction possible for a given weather situation and available computational resources (Droegemeier et al. 2005, 2009; Plale et al. 2006).

This so-called *dynamically adaptive NWP* strategy is the focus of the present paper and the principal foundation of Linked Environments for Atmospheric Discovery (LEAD), a five-year Large Information Technology Research (ITR) grant funded by the National Science Foundation (NSF). LEAD has created an integrated, scalable framework in which meteorological analysis tools, forecast models, and data repositories can operate as dynamically adaptive, on-demand, grid-enabled systems that a) change configuration rapidly and automatically in response to weather; b) respond to decision-driven inputs from users; c) initiate other processes automatically; and d) steer remote observing technologies to optimize data collection for the problem at hand (e.g., Droegemeier 2009). Although mesoscale meteorology is the particular science domain to which these concepts have been applied in LEAD, the methodologies and infrastructures developed are extensible to other domains including medicine, ecology, hydrology, geology, oceanography and biology.

We take in this paper a first step toward answering the following fundamental question regarding dynamically adaptive NWP: For a given weather event or situation, what configuration of a cloud-resolving numerical prediction system yields the "best" (defined appropriately) solution under specified computational resource and other constraints? Stated another way, what is the most effective way to utilize a given set of computational resources (e.g., as measured by total CPU time) among choices such as a single fine-grid forecast in a small domain, a single medium-grid forecast in a larger domain, nested grids, ensembles, etc?

2. DYNAMIC ADAPTATION TO WEATHER

a. Objective and Strategies

Dynamic adaptation can take many forms but in all cases, the objective of dynamically adaptive

---

* Corresponding Author Address: Prof. Kelvin K. Droegemeier, University of Oklahoma, 120 David L. Boren Boulevard, Suite 5900, Norman, OK 73072. Email: kkd@ou.edu. Visit LEAD on the web at http://portal.leadproject.org.

[1] The exception in the US is hurricane models, run in domains that follow hurricanes with time (e.g., Davis et al. 2008); the high resolution window of the National Weather Service Weather Research and Forecast (WRF) model (Skamarock et al. 2008); and the UK Meteorological Office model, windowed domains of which can be configured by forecasters to run on demand.)

systems is to improve upon their static counterparts in some manner, ideally one that formally optimizes or at least quantitatively improves upon certain aspect(s) of performance (Droegemeier et al. 2007). In the case of NWP, systems or components may adapt in time (e.g., rapid update cycling, where new forecasts are launched in rapid sequence, perhaps overlapping in time), space (e.g., adaptive nesting as a means of reducing local error or increasing fidelity e.g., Dietachmayer and Droegemeier 1992; Skamarock and Klemp 1993), modality (e.g., ensembles of a particular number or configuration), or via adaptive observations (e.g., Morss et al. 2001, Brotzge et al. 2006; Plale et al. 2006).

Adaptation can be automated, manual, objective, or heuristic and can occur in a variety of locations within the system, at multiple levels and in highly connected, nonlinear ways. Finally, cyberinfrastructure is an important but often overlooked component of dynamically adaptive systems. For example, networks, high performance computers, data bases, and other cyber elements must be able to change configuration quickly, and in a coordinated, fault-tolerant manner, in response to a given situation (e.g., Marru et al. 2008) in order to accommodate the severe quality of service demands associated with operational NWP. It is this latter aspect of dynamic adaptation that presents among the greatest challenges, especially given the long history of batch-operated computing systems, which are not designed for interrupt-driven, on-demand operations.

b. Key Questions

Because dynamic adaptation is a complex, multi-disciplinary problem, associated with it are a number of challenging questions, including but not limited to the following:

- When is adaptation useful and can the costs and benefits of adaptation be quantified?
- What types of adaptation are possible and most effective and how can they be chosen and combined?
- How is adaptivity triggered/controlled?
- What elements of the system can or should adapt?
- How can one deal with loss of resources or less than ideal availability to achieve the required adaptation?
- What metrics can be used to measure the effectiveness of adaptation and can “optimal” adaptivity be defined?
- What negative consequences exist to adaptation?
- What are the time scales of adaptivity and what controls them?
- Do adaptivity and on-demand functionality need to be pre-scheduled to any extent?
- What triggers the decision to adapt and how is the decision communicated across the system?
- What does quality of service mean in an adaptive system?
- In the context of NWP, how might adaptation impact predictability?

c. Canonical Example of Adaptive NWP

A canonical example of adaptation is shown in Figure 1 from Droegemeier (2009). The far left side of the figure depicts observations transmitted by a variety of observing systems, including the NEXRAD operational and CASA (Brotzge et al. 2006) experimental Doppler radar networks, as well as forecast model output. These data can be processed, separately or in parallel, by a data mining system (marker 1), within which exists a persistent agent that searches for user-defined atmospheric conditions associated with the development of deep convection (e.g. instability, precipitation on radar of a certain intensity or vertical extent). Alternatively, the observations and model output can be assimilated using algorithms, such as ensemble Kalman filtering (e.g. Tong & Xue 2005) or four-dimensional variational methods (e.g. Rihan et al. 2005), to produce gridded fields of all relevant atmospheric quantities.

The mining agent can then be applied to the resulting quantities to search for specified values or patterns indicative of convection. If such conditions or features are found, LEAD automatically triggers a WRF numerical forecast (marker 2), the specific grid spacing, domain size, forecast duration and allowable wall clock time of which are communicated by a brokering agent to the TeraGrid (marker 3; www.teragrid.org). If resources available within the TeraGrid at that time are insufficient, LEAD adjusts model parameters based upon user-assigned priorities until the job can be run and output returned sufficiently quickly. Such an output (marker 4) is analysed by the same data mining engine used previously in an attempt to identify regions in which targeted observations might improve forecast quality (this operation could be performed on sensitivity fields as well; e.g. Errico & Vukicevic 1992; Park & Droegemeier 2000). If such regions are found, an agent communicates with an adaptive observing system, such as the radars being developed by the US National Science Foundation (NSF) Center for the Collaborative Adaptive Sensing of the

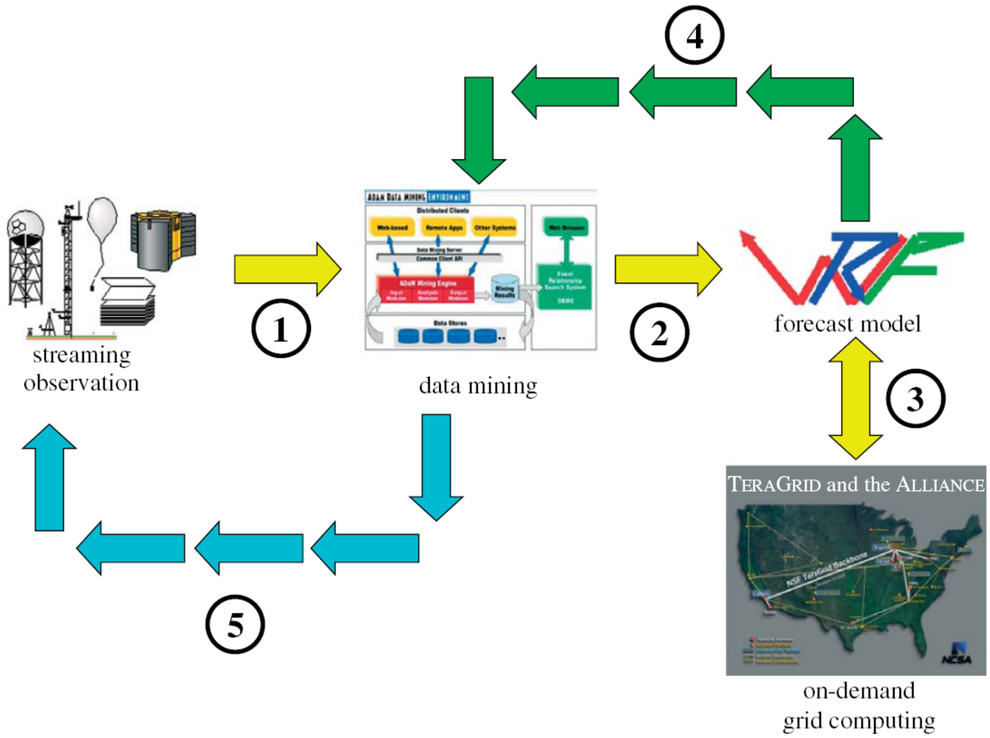

Figure 1. Sample closed-loop dynamically adaptive weather analysis and prediction scenario enabled by LEAD. See text for further details. (From Droegemeier 2009)

Atmosphere (Brotzge et al. 2006; Plale et al. 2006), and new targeted observations are collected. The process then repeats or is modified automatically if other specified criteria are met.

d. Potential Benefits and Drawbacks

The potential benefits of a dynamically adaptive approach to NWP are many. First and foremost is that the prediction system has the ability to differentiate among weather events or atmospheric scenarios, thus potentially providing a "best" solution in a given situation rather than "good" solutions across all situations. Second, the adaptive approach potentially makes optimal use of available data and computational resources, adjusting model grid spacing and other parameters, for example, to provide necessary capability when and where needed. Third, as shown by Brewster et al. (2008), a dynamic approach allows operational forecasters to launch predictions themselves, in anticipation of specific events such as deep convective storms, thus providing direct control in generating the specific information needed to make informed decisions.

Likewise, the dynamically adaptive approach also has a number of potential drawbacks. For example, in operational forecasting, the evaluation of a model operated daily in the same configuration allows forecasters to identify idosyncracies that can be factored into output analysis. If, as is the case in a dynamically adaptive system, the forecast configuration changes frequently in response to changing weather, the skill obtained through repeated examination of output in a static framework would be diminished or perhaps lost. The extent to which this drawback is overwhelmed by benefits of increased accuracy and reliability owing to dynamic adaptation is a dual research-operations question that is being addressed in part by the NOAA Hazardous Weather Test Bed (Brewster et al. 2008).

3. EXPERIMENT DESIGN

a. Model Description and Parameter Settings

We utilize Version 3.0 of the Weather Research and Forecast (WRF; http://wrf-model.org) model to begin assessing the question posed in §1, namely, for a given weather event or situation, what configuration of a cloud-resolving numerical prediction model yields the "best" solution under specified computational resource and other constraints? To develop an initial understanding of the many complex components of dynamically adaptive NWP, we keep these initial experiments as simple as possible. Specifically, we simulate the evolution, over three hours, of an isolated supercell storm in an idealized, initially horizontally uniform, conditionally unstable environment using only basic cloud and precipitation physics. The case chosen is the default supercell test scenario in WRF for which the environment is characterized by a wind hodograph that turns through a quarter-circle. The storm is initiated by a warm thermal impulse placed within the boundary-layer at the center of the model domain (see §3b).

A total of 51 vertical levels are used with exponential stretching, and an average storm speed is subtracted from the environment to keep the principal right-moving storm of interest near the center of the domain. Cloud and precipitation processes are represented using the Kessler (1969) warm-rain parameterization and no radiation, surface physics or terrain effects are included. All other parameters represent WRF default values unless otherwise specified.

b. Domain Configuration

To mimic a routine forecast from a marginally cloud-resolving operational model, we create a baseline forecast (we use the terms forecast and simulation synonymously) at 5 km grid spacing in a domain having an area 1000 x 1000 km^2 (outer, black box in Figure 2). It is upon this forecast we wish to improve via the application of simple adaptation strategies. A "truth" or "nature" simulation, against which all other experiments are compared, is created in the same domain as the baseline forecast (Figure 2) though using a uniform horizontal grid spacing of 250 m.

Each of our adaptation experiments consists of a *single two-way nested grid placed within the baseline grid*. The size and grid spacing of the particular nest used (Figure 2) is determined, for the principal set of experiments, by its nesting ratio (i.e., ratio of grid spacing between the baseline forecast and the single nest) under the constraint that all nests must utilize a nearly fixed amount of computing time[2]. Consequently, as shown in Table 1, grid spacing varies in proportion to domain size, and because only a single nested grid is used in each experiment, the nesting ratio varies inversely with grid spacing. Shaded diagonals of Tables 1 and 2 show experiments in which computing time is constrained. For experiments below (above) the diagonal, less (more) computer time is required for a fixed grid spacing because smaller (larger) domains are used.

[2] Computing time as used here is the sum of dedicated CPU and shared I/O time.

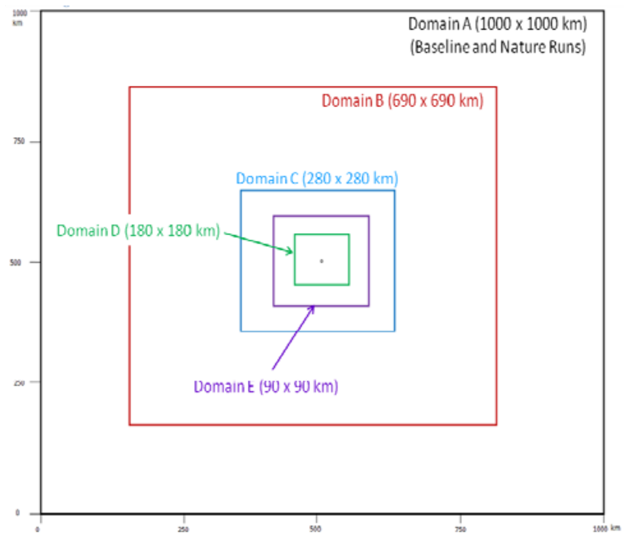

Figure 2. Model domains used in the experiments. See the text for further details.

In reality, adaptation may be triggered by an atmospheric event (e.g., first echo on radar or first towering cumulus on satellite), or well in advance based upon expected regions of convective initiation (Brewster et al. 2008). In the present experiments, all forecasts are run for three hours and all nests are initiated at t = 0. In the future, nests will be initiated later in time and also will be repositioned as the forecast proceeds.

4. RESULTS

a. Methodology

Owing to the spatial structure of deep convective storms, traditional measures of forecast skill (e.g., equitable threat, RMS error) designed for fields having global structure, and that measure broad overlap between prediction and verification, are less effective (e.g., Brown et al. 2004; Ebert 2008) because small temporal or spatial errors in storms can yield poor scores when, in reality, the forecast is reasonably "similar" to the actual event. Consequently, we utilize here a relatively simple measure of agreement between forecast and "truth," namely, the mean square error, which can be expressed as the sum of dissipation (amplitude) and dispersion (phase) error (Takacs 1985). This approach was used by Hou et al. (2001) to verify ensemble forecasts, and errors in the present case are computed on each forecast grid by interpolating to it values from the nature experiment. Subsequent work will improve upon this admittedly limited approach by applying other quantitative measures of skill.

Table 1. Relationship among model domain size (left column), grid spacing and nesting ratio (top row), and computing time (table cell entries, in seconds) for all experiments. Domain A is the baseline forecast (see Figure 2), and the single nested grid domains are indicated by the letters B-E (see also Figure 2). The shaded diagonal cells show computing time associated with each experiment which, for the nested runs, is constrained to be nearly constant.

Nest Ratio / Domain	Baseline (Δx = 5 km)	3 (Δx = 1.67 km)	5 (Δx = 1 km)	7 (Δx= 714 m)	9 (Δx = 555 m)
A (1000 x 1000 km^2)	840 s				
B (690 x 690 km^2)		5333 s	23,517 s	46,766 s	60,416 s
C (280 x 280 km^2)		1799 s	5306 s	10,565 s	31,387 s
D (180 x 180 km^2)		1440 s	2957 s	5359 s	14,752 s
E (90 x 90 km^2)		1177 s	1520 s	2269 s	5240 s

Table 2. As in Table 1 except showing the number of horizontal grid points for each domain.

Nest Ratio / Domain	Baseline (Δx = 5 km)	3 (Δx = 1.67 km)	5 (Δx = 1 km)	7 (Δx= 714 m)	9 (Δx = 555 m)
A (1000 x 1000 km^2)	201 x 201				
B (690 x 690 km^2)		415 x 415	691 x 691	967 x 967	1243 x 1243
C (280 x 280 km^2)		169 x 169	281 x 281	393 x 393	505 x 505
D (180 x 180 km^2)		109 x 109	181 x 181	253 x 253	325 x 325
E (90 x 90 km^2)		55 x 55	91 x 91	127 x 127	163 x 163

b. Baseline (Parent) Grid Forecast Errors

Because a two-way nested gridding approach is used for all adaptive forecasts, the solution on the 5 km parent or baseline grid also is impacted. Considering surface rainwater mixing ratio, total error for the baseline forecast (no nesting) is 0.025 g kg^{-1}. The use of nesting is expected to yield a smaller error, and Table 3 shows such is the case. In general, for a given nested domain size, total error on the parent grid decreases as grid spacing decreases and nesting ratio increases. Changes generally are small because the parent domain is extremely large relative to the size of the storm system being represented.

Table 3. Total error at 3 hours for surface rainwater mixing ratio (g kg^{-1}) for the baseline (parent) grid.

Nest Ratio (Δx) / Domain	3 (1.67 km)	5 (1 km)	7 (714 m)	9 (555 m)
B (690^2 km^2)	0.017	0.014	0.014	0.014
C (280^2 km^2)	0.017	0.014	0.017	0.014
D (180^2 km^2)	0.019	0.018	0.017	0.017
E (90^2 km^2)	0.024	0.013	0.012	0.011

The behavior of total error as a function of domain size for a fixed grid spacing (nesting ratio) is less regular. One would expect errors to be smaller for smaller nested domains because feedback from the fine to coarse grid would be occurring over a much smaller fraction of the total grid. This generally appears to be the case for grid E, except for a nesting ratio of three (Table 3).

From the perspective of a fixed amount of computing resource, it is difficult to assess which of the four model configurations represented by diagonal entries in the table – all of which use approximately the same computing time – yields the “best” solution for rainwater mixing ratio. Because the practical utility of nested grids lies in the increased fidelity of the solution on them rather than the parent grid, improvements on the latter likely are less relevant in an operational context.

Table 4 shows the same information as Table 3 except for vertical velocity at 4 km altitude. For comparison, the total error for the baseline forecast (no nesting) is 0.31 m s^{-1}. Perhaps not surprisingly, given the small size of storm updrafts, no terribly clear pattern emerges from the errors other than an overall reduction relative to the no-nest baseline case.

Table 4. Total error at 3 hours for vertical velocity ($m\ s^{-1}$) at 4 km altitude for the baseline (parent) grid.

Nest Ratio (Δx) / Domain	3 (1.67 km)	5 (1 km)	7 (714 m)	9 (555 m)
B (690^2 km^2)	0.21	0.19	0.19	0.20
C (280^2 km^2)	0.21	0.19	0.19	0.20
D (180^2 km^2)	0.23	0.21	0.22	0.22
E (90^2 km^2)	0.25	0.19	0.19	0.18

c. Nested Grid Forecast Errors

Tables 5 and 6 show, for the complete area of each nested grid, total error at 3 hours (end of the forecast) for surface rainwater mixing ratio and vertical velocity at 4 km altitude, respectively. In all experiments, error increases monotonically for a given grid spacing (nesting ratio) as domain size decreases, indicating both an impact of the nested lateral boundaries (at which interpolation occurs every time step in the two-way nesting procedure) as well as, for the two smallest domains (D and E), a partial exit of a portion of the left-moving split storm.

For a given domain size, however, as the grid becomes successively finer, total error decreases, reaching a minimum at 714 m spacing (nesting ratio of 7) for rainwater (Table 5) and 555 m spacing for vertical velocity (Table 6, nesting ratio of 9). The error in all cases is overwhelmingly dominated by dispersion, indicating that phase error decreases as the grid becomes finer, up to the point where the nesting ratio is so large that errors associated with interpolation in the two-way nesting procedure begin to dominate.

An interesting and somewhat anticipated result is that, *for a fixed amount of computing time* (diagonal entries), total error is a minimum for the *largest* domain and *coarsest* grid spacing. As noted previously, the increase in error for the smallest two domains as grid spacing decreases is to some extent associated with the partial exit of the left-moving storm; however, the entire storm system is contained within the rather large domains B and C, suggesting that nesting ratio is the dominant factor in those error trends. [To eliminate this effect, we examine in §4c errors computed within the one domain common to all experiments, i.e., domain E.]

It is important to recognize that, in practice, multiply nested grids likely would be used to mitigate this behavior, and indeed we are studying such configurations now. However, establishing a foundation using a single nest was an important first step before attempting to understand more complicated scenarios.

Table 5. Total error at 3 hours for surface rainwater mixing ratio ($g\ kg^{-1}$) computed over the complete domain of the nested grids shown.

Nest Ratio (Δx) / Domain	3 (1.67 km)	5 (1 km)	7 (714 m)	9 (555 m)
B (690^2 km^2)	0.052	0.049	0.048	0.052
C (280^2 km^2)	0.315	0.296	0.293	0.315
D (180^2 km^2)	0.489	0.434	0.356	0.365
E (90^2 km^2)	1.156	1.020	0.786	0.871

Table 6. Total error at 3 hours for vertical velocity ($m\ s^{-1}$) at 4 km altitude computed over the complete domain of the nested grids shown.

Nest Ratio (Δx) / Domain	3 (1.67 km)	5 (1 km)	7 (714 m)	9 (555 m)
B (690^2 km^2)	0.56	0.57	0.60	0.60
C (280^2 km^2)	3.36	3.44	3.64	3.61
D (180^2 km^2)	5.76	5.62	6.05	6.00
E (90^2 km^2)	16.01	15.47	14.40	13.78

d. Common Grid Forecast Errors

To deal with the shortcoming, noted above, regarding error interpretation given that part of the left-moving storm exits the smallest two domains, we present statistics computed only over the domain (E) that is common to all nested runs. Tables 7 and 8 show, *over domain E*, total error at 3 hours (end of the forecast) for surface rainwater mixing ratio and vertical velocity at 4 km altitude, respectively. Corresponding horizontal cross sections of forecasts over the same area, arrayed identical to the cells in the tables, are shown in Figures 3 and 4, respectively, and in Figures 5 and 6 for the nature or “truth” run, also interpolated to domain E.

In contrast to the previous discussion, error behavior is somewhat more variable as a function of domain size and grid spacing (nesting ratio). In particular, error increases as domain size decreases for a spacing of 1.67 km (Table 7), but for finer grids and other nesting ratios, smaller domains sometimes produce smaller errors. A significant decrease in error occurs in rainwater mixing ratio for all nested grids (Table 7) when the nesting ratio increases from five to seven (compare Figures 3 and 5), contrary to the behavior noted above when error was computed over the entire nested domain. Interestingly, a nesting ratio of seven still appears to be optimal for rainwater, and for a fixed amount of computing resource (diagonal cells), this ratio also produces

Table 7. Total error at 3 hours for surface rainwater mixing ratio (g kg^{-1}) computed over Domain E for all nested grids.

Nest Ratio (Δx) / Domain	3 (1.67 km)	5 (1 km)	7 (714 m)	9 (555 m)
B (690^2 km^2)	1.121	1.130	0.767	0.878
C (280^2 km^2)	1.122	1.130	0.763	0.883
D (180^2 km^2)	1.124	1.125	0.753	0.871
E (90^2 km^2)	1.156	1.020	0.786	0.871

Table 8. Total error at 3 hours for vertical velocity (m s^{-1}) at 4 km altitude computed over Domain E for all nested grids.

Nest Ratio (Δx) / Domain	3 (1.67 km)	5 (1 km)	7 (714 m)	9 (555 m)
B (690^2 km^2)	2.52	2.16	2.20	2.05
C (280^2 km^2)	2.52	2.18	2.18	2.06
D (180^2 km^2)	2.51	2.15	2.20	2.10
E (90^2 km^2)	2.55	2.09	2.01	1.91

the best overall solution. Consequently, in a dynamically adaptive NWP scenario when available computing time is limited to ~5300 s and a single nested grid is inserted within the 5 km operational forecast, the best model configuration with respect to rainwater at 3 hours would be a nest having an area of 180 x 180 km^2 and using 714 m grid spacing. The manner in which this choice is actually made for a given weather scenario is a component of our ongoing research (see §5) given that the LEAD infrastructure can accommodate exactly this sort of capability.

Interestingly, vertical velocity (Table 8, Figures 4 and 6) does not show a similar marked change in error between nesting ratios of five and seven, and possible reasons for this behavior continue to be evaluated. Similar abrupt changes have been noted in previous studies (e.g., Adlerman and Droegemeier 2002; Petch 2006) and may be related to differences in dynamical and microphysical processes at kilometer to sub-kilometer grid spacing. Of the many other fields and altitudes examined for the present simulations (not shown), only surface rainwater mixing ratio shows this abrupt behavior.

Recalling Table 1, which shows a wide range of computing times across the study parameter space, it is clear that substantially large nested domains at fine grid spacing do not necessarily yield a commensurately more accurate solution (or return on investment) for the simple storm case and single nested grid configuration used here. This likely would not be the case if, for example, a substantial portion of the baseline domain were filled with convection, in which case a single nested grid, while perhaps "good" locally, would for some of the configurations shown here be fully inadequate by virtue of excluding potentially important weather. Further, such a situation would lead to potentially smaller domains relative to available computer time because physics components of the code would consume a proportionally large amount of time.

5. SUMMARY AND FUTURE WORK

We used the WRF model to examine a very simple scenario of dynamically adaptive numerical weather prediction in which, under the constraint of fixed computing time, a single nested grid of variable spacing and size, but fixed location, was placed within a baseline or pseudo-operational forecast run at 5 km grid spacing. The case examined was an idealized, isolated supercell storm with only warm rain microphysics and no terrain, radiation or surface physics.

We noted that phase error (dispersion) was the dominant error in all experiments, and that the mean square or total error tended to increase for a fixed grid spacing (nesting ratio) as domain size decreased. Not all variables exhibited this behavior, and for some the "best" forecast was not produced by the largest domain and finest grid spacing owing to the influence of the nesting ratio, which increased as grid spacing decreased.

We are pursuing a variety of additional, much more sophisticated experiments to build upon this admittedly simple example. These include considerably greater aerial coverage and complexity of storms within the "truth" experiment, the use of multiple and moving nests, the launching of nests at different times, and different ways of specifying computational constraints (e.g., total processing time, forecast turnaround time).

In these experiments, we did not concern ourselves with the decision by which any given nested grid or other change in model configuration might be implemented. However, the LEAD infrastructure has the capability to automatically launch multiple forecasts on demand using a variety of "triggers," including but not limited to features detected in observations such as radar data, regions of the atmosphere identified as potentially threatening by humans (e.g., mesoscale discussions, watches) or data (e.g., an exceedance threshold for CAPE), features detected in previous forecasts that might suggest areas of refinement, errors established by adjoint sensitivity analysis, or manually by humans.

As we answer more questions regarding dynamically adaptive systems (§2b), we will utilize the LEAD infrastructure to test the concept of forecast model solution optimization under a variety of physical and computational constraints.

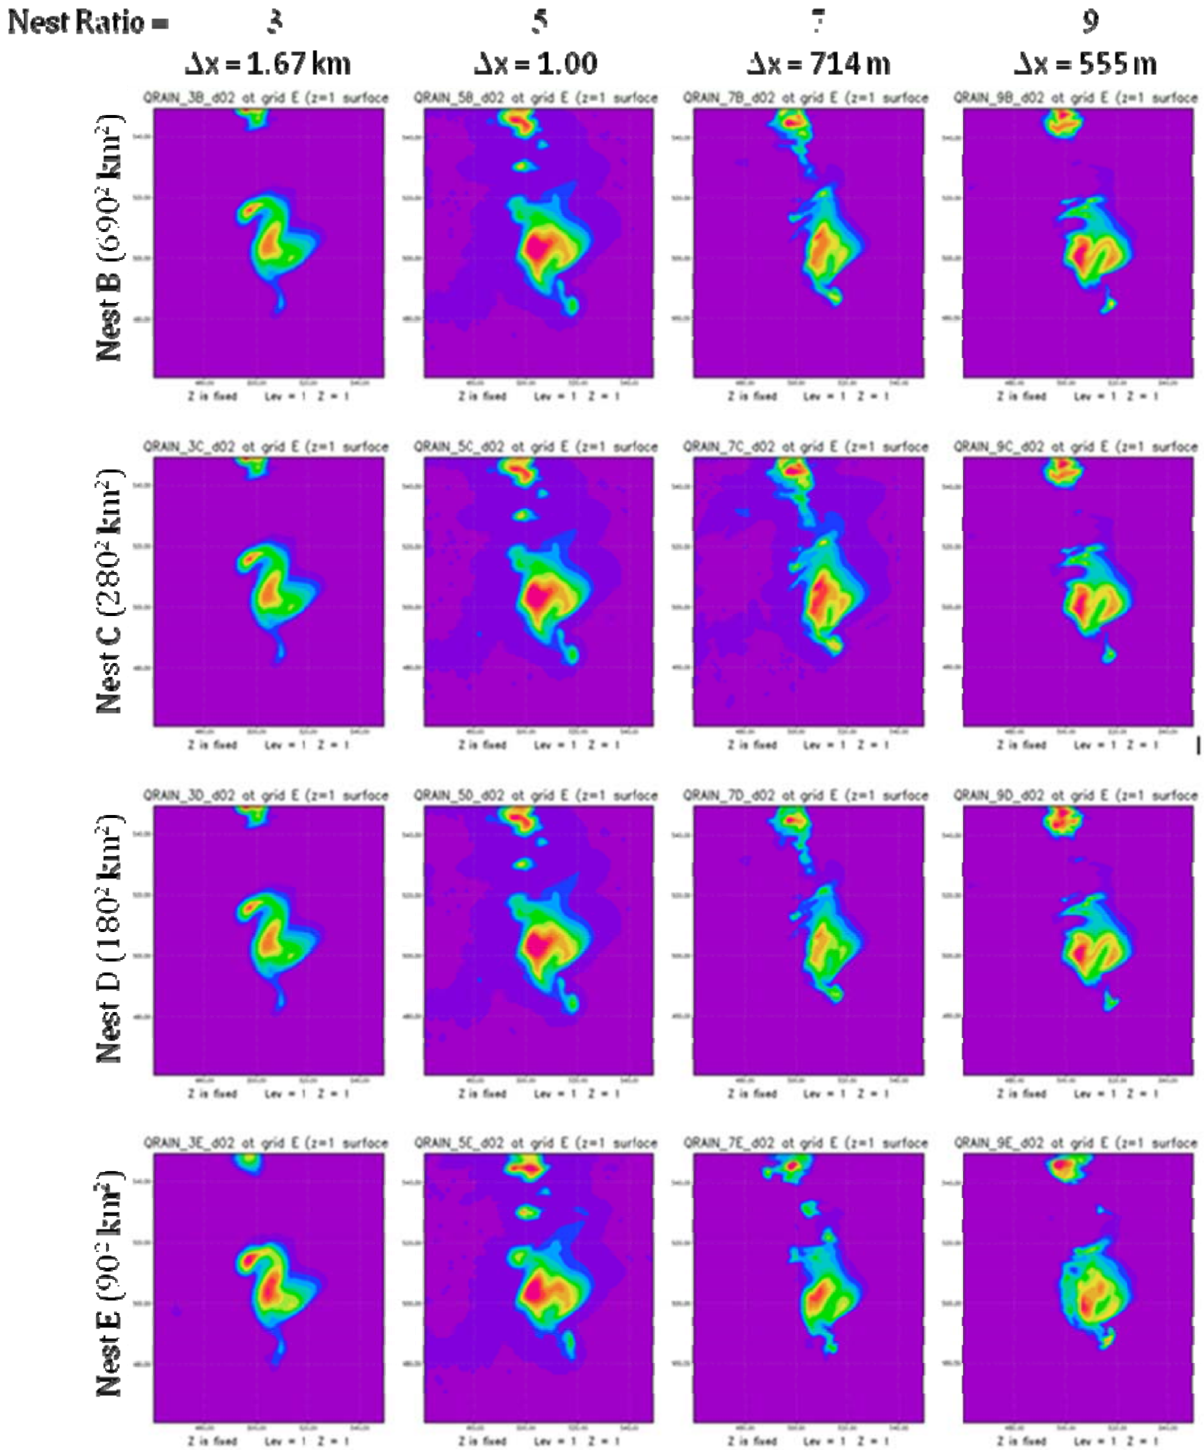

Figure 3. Horizontal cross sections, over the area of domain E, of surface rainwater mixing ratio for all nested grid experiments.

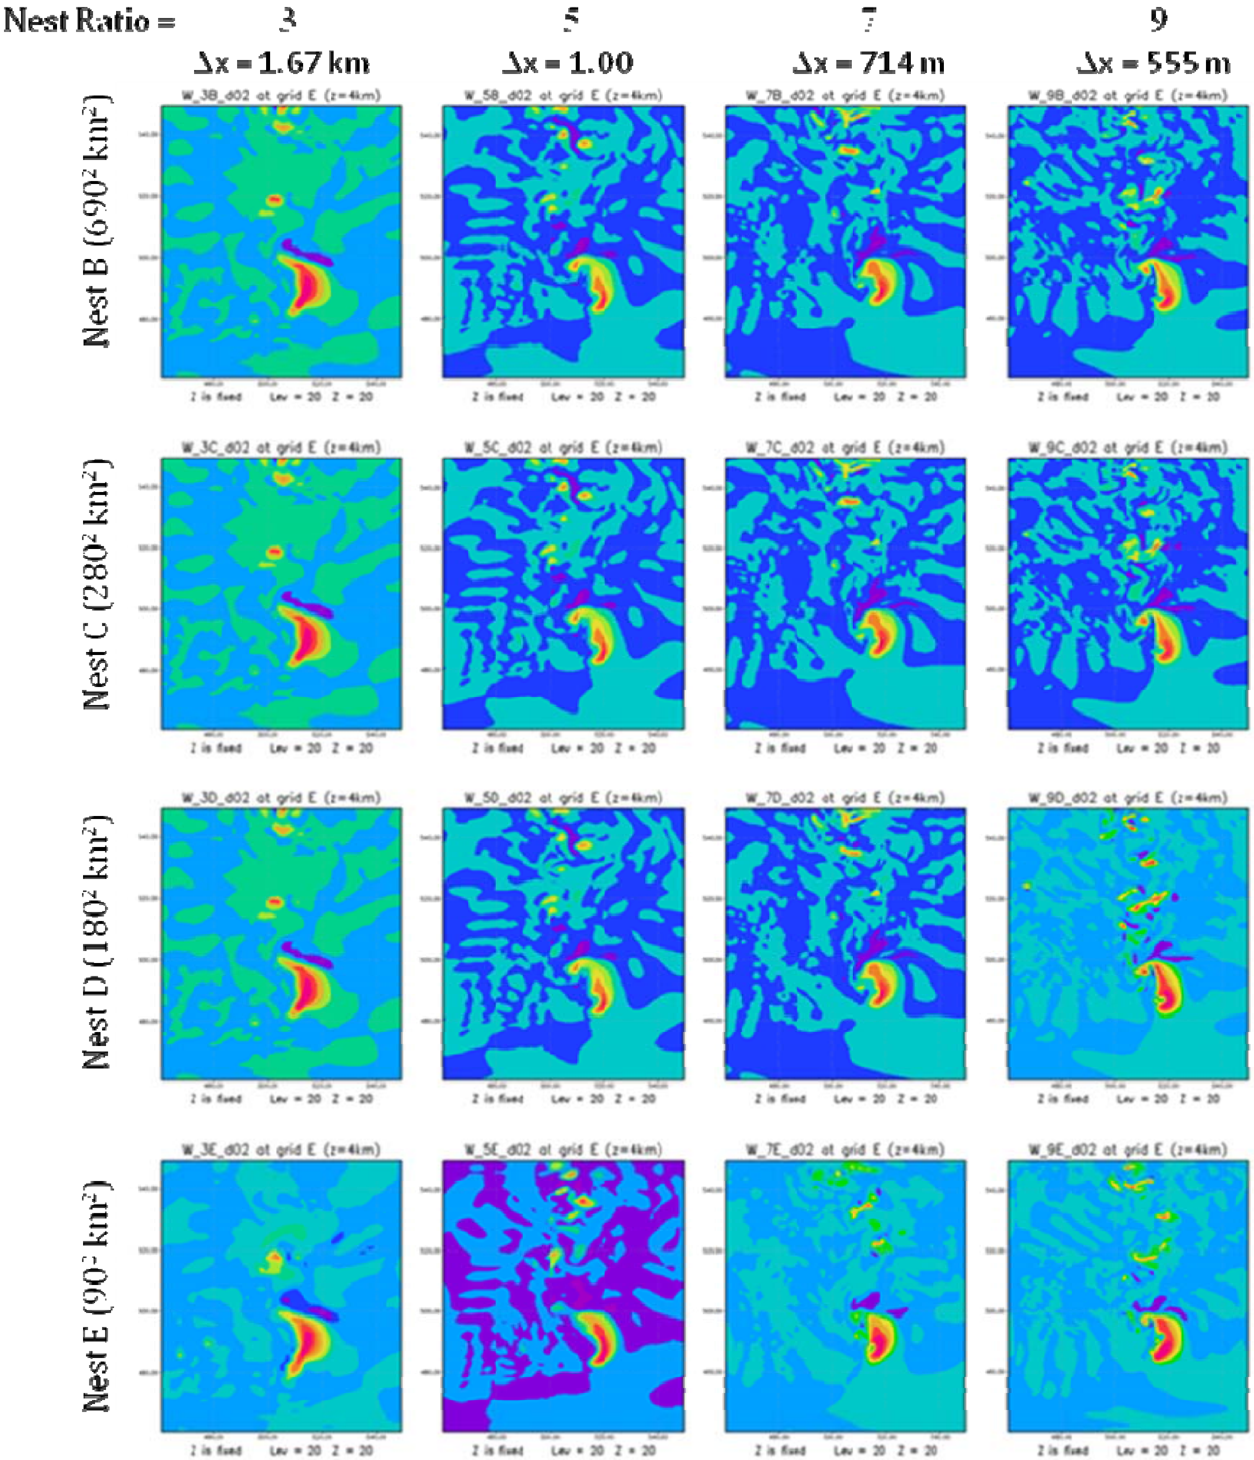

Figure 4. Horizontal cross sections, over the area of domain E, of vertical velocity at 4 km altitude for all nested grid experiments.

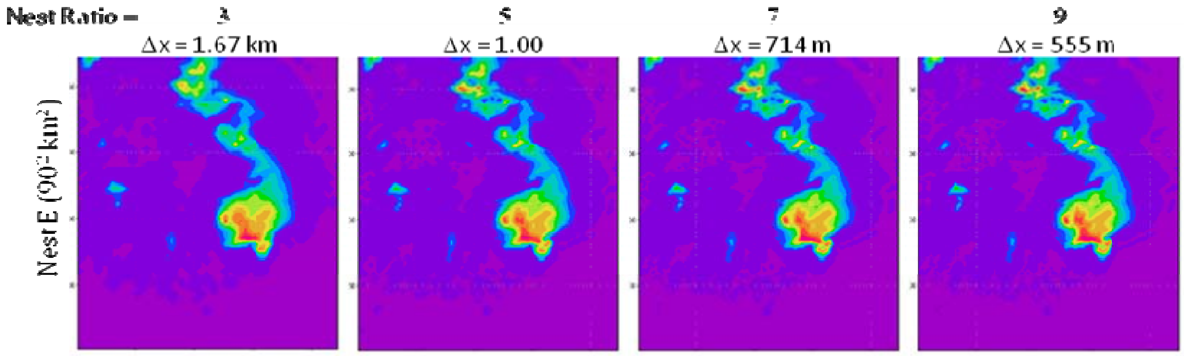

Figure 5. Horizontal cross sections of surface rainwater mixing ratio for the nature or "truth" experiment interpolated to each of the nested grids over the area of Grid E.

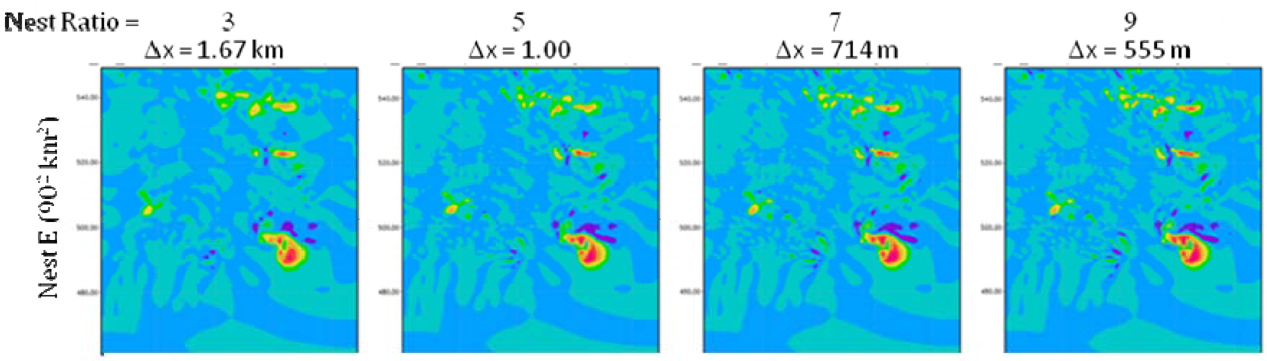

Figure 6. Horizontal cross sections of vertical velocity at 4 km altitude for the nature or "truth" experiment interpolated to each of the nested grids over the area of Grid E.

6. ACKNOWLEDGMENTS

The simulations were performed on the Ranger computing system at the Texas Advanced Computing Center, which is funded in part by the National Science Foundation. The authors were funded by LEAD, which is supported in part by the National Science Foundation under Cooperative Agreements ATM-0331594 to the University of Oklahoma. They gratefully acknowledge the support of their many LEAD colleagues. Support for the NOAA Hazardous Weather Test Bed was provided by a C-STAR grant to the first author from NOAA under agreement NA17RJ1227/NAGO000070029 to the University of Oklahoma.

7. REFERENCES

Adlerman, E.J. and K.K. Droegemeier, 2002: The sensitivity of numerically-simulated cyclic mesocyclogenesis to variations in model physical and computational parameters. *Mon. Wea. Rev*., **130**, 2671-2691.

Brewster, K. and Co-Authors, 2008: Use of the LEAD portal for on-demand severe weather prediction. *Preprints*, 24th Conf. On Integrated Information and Processing, 21-24 January, New Orleans, LA, Amer. Meteor. Soc.

Brotzge, J., K.K. Droegemeier, and D.J. McLaughlin, 2006: Collaborative Adaptive Sensing of the Atmosphere (CASA): New radar system for improving analysis and forecasting of surface weather conditions. *J. Transport. Res. Board*, No. 1948, 145-151.

Brown, B.G. and co-authors, 2004: New verification approaches for convective weather forecasts. *Preprints*, 11th Conf. on Aviation, Range and Aerospace Meteorology, Amer. Meteor. Soc., 3-8 October, Hyannis, MA.

Davis, C. and co-authors, 2008: Prediction of landfalling hurricanes with the advanced hurricane WRF model. *Mon. Wea. Rev*., **136**, 1990-2005.

Dietachmayer, G. and K.K. Droegemeier, 1992: Application of continuous dynamic grid adaption techniques to meteorological modelling, Part I: Basic formulation and accuracy. *Mon. Wea. Rev*., **120**, 1675-1706.

Droegemeier, K.K. and Co-Authors, 2005: Service-oriented environments in research and education for dynamically interacting with mesoscale weather. *Computing in Science and Engineering*, **7**, 12-29.

Droegemeier, K.K., and Co-Authors, 2007: A new paradigm for mesoscale meteorology: Grid and web service-oriented research and education in LEAD. Preprints, *23rd Conf. On Integrated Information and Processing*, 15-18 January, San Antonio, TX, Amer. Meteor. Soc.

Droegemeier, K.K., 2009: Transforming the Sensing and Numerical Prediction of High Impact Local Weather Through Dynamic Adaptation. *Phil. Trans. of the Royal Soc*. A, 1-20.

Ebert, E.E., 2008: Fuzzy verification of high resolution gridded forecasts: A review and proposed framework. *Meteorol. Appls*., **15**, 51-64.

Errico, R.M. and T. Vukicevic, 1992: Sensitivity analysis using an adjoint of the PSU-NCAR mesoscale model. Mon. Wea. Rev., **120**, 1644-1600.

Hou, D., E. Kalnay, and K.K. Droegemeier, 2001: Objective verification of the SAMEX '98 ensemble forecasts. *Mon. Wea. Rev*., **129**, 73-91.

Kain, J.S., S.J. Weiss, D.R. Bright, M.E. Baldwin, J.J. Levit, G.W. Carbin, C.S. Schwartz, M. Weisman, K. Droegemeier, D. Weber, and K.W. Thomas, 2008: Some practical considerations for the first generation of operational convection-allowing NWP: How much resolution is enough? *Wea. and Forecasting*, **23**, 931-952.

Kessler, E., 1969: *On the Distribution and Continuity of Water Substance in Atmospheric Circulations. Meteor. Monogr.,* No. 32, Amer. Meteor. Soc., 84 pp.

Marru, S., D. Gannon, S. Nadella, P. Beckman, D.B. Weber, K.A. Brewster, K.K. Droegemeier, 2008: LEAD Cyberinfrastructure to Track Real-Time Storms Using SPRUCE Urgent Computing, CTWatch Quarterly. Volume 4, Number 1, March.

Morss. R. E., K.A. Emanuel and C. Snyder, 2001: Idealized adaptive observation strategies for improving numerical weather prediction. J. Atmos. Sci., **58**, 210-232.

Park, S.K. and K.K. Droegemeier, 2000: Sensitivity analysis of a 3-D convective storm: Implications for variational data assimilation and forecast error. *Mon. Wea. Rev*., **128**, 140-159.

Petch, J.C., 2006: Sensitivity studies of developing convection in a cloud-resolving model. Quart. *J. Roy. Meteor. Soc*., **132**, 345-398.

Plale, B., J. Alameda, R. Wilhelmson, D. Gannon, S. Hampton, A. Rossi, and K.K. Droegemeier, 2004: User-oriented active management of scientific data with my LEAD. *IEEE Internet Computing*, **9**, 27-34.

Plale, B., D. Gannon, J. Brotzge, K.K. Droegemeier and Co-Authors, 2006: CASA and LEAD: Adaptive cyberinfrastructure for real-time multiscale weather forecasting. *IEEE Computer*, **39**, 66-74.

Plale, B. and Co-Authors, 2007: Real time filtering and mining of NEXRAD streams for mesoscale forecast and prediction. Preprints, *23rd Conf. On Integrated Information and Processing*, 15-18 January, San Antonio, TX, Amer. Meteor. Soc.

Rihan, F. A., Collier, C. G. and I. Roulstone, 2005: Four-dimensional variational data assimilation for Doppler radar wind data. *J. Comput. Appl. Math*., **176**, 15–34.

Shuman, F.G., 1989: History of numerical weather prediction at the National Meteorological Center. *Wea. and Forecasting*, **3**, 286-296.

Skamarock, W.C. and J.B. Klemp, 1993: Adaptive grid refinement for two-dimensional and three-dimensional nonhydrostatic atmospheric flow. Mon. Wea. Rev., **121**, 788-804.

Skamarock, W.C. and co-authors, 2008: A description of the advanced research WRF version 3. NCAR Technical Note NCAR/TN-475+STR, 113 pp.

Takacs, L.L., 1985: A two-step scheme for the advection equation with minimized dissipation and dispersion errors. Mon. Wea. Rev., **113**, 1050-1065.

Tong, M. and M. Xue, 2005: Ensemble kalman filter assimilation of Doppler radar data with a compressible nonhydrostatic model: OSSE experiments. Mon. Wea. Rev., **133**, 1789-1807.