

Real-Time High-Resolution Weather Modelling in “Rugged” Regions

Don Morton, *Arctic Region Supercomputing Center / University of Montana*, **Eugene Petrescu**, *National Weather Service, Missoula, Montana*, and **Bryce Nordgren**, *United States Forest Service, Missoula, Montana*

ABSTRACT: *Numerical weather prediction in mountainous regions presents unique challenges for capturing the important small-scale dynamics of these complex environments. The stakes are high, as expeditious, localized forecasts are required for activities such as fire-fighting and aviation. This work presents a performance evaluation of the Cray XD1 (in comparison with an IBM p655+) execution of the Weather Research and Forecasting (WRF) model applied to “rugged” regions with resolution of 7.5km, 2.5km and 830m.*

KEYWORDS: XD1, WRF, numerical weather prediction

1. Introduction

Weather in Northern Rockies

Weather forecasting in mountainous regions presents challenges not encountered in the “flatlands.” Generally smooth atmospheric processes, upon encountering mountains, are often drastically modified by the topography and the localized weather that results from the rough terrain. It is well known that mountains can generate their own weather, and many of these processes are well understood. For example, mountains have a tendency to generate downslope winds at night due to heavier, cool air flowing down the drainages, and this is often realized by strong winds at drainage outlets. Residents of Missoula, Montana are familiar with the Hellgate Winds which tend to drain the mountains to the east of their cold air, with the flow constricting and accelerating through the narrow Hellgate Canyon on the east side of town (Figure 1). These winds frequently appear on an otherwise calm, cool morning when winds aloft are negligible. Small aircraft departing to the east encounter these stiff winds up to 2,000-3,000 feet AGL, then experience a region of turbulence due to shear between the easterly surface winds and negligible winds aloft, followed by very smooth conditions at higher altitudes. This pattern is also prevalent in the vicinity of other area drainages, and although local pilots are well

aware of them, they do not show up in numerical weather forecasts.




Figure 1: Easterly Hellgate Wind, Missoula, Montana

Another weather event familiar to Missoulians is the valley inversion. Several large-scale drainages terminate in the “Missoula bowl” (Figure 2), often resulting in an accumulation of cold, dense air in the bowl. These inversions occur throughout the year and, again, are not captured effectively in numerical weather forecasts. In the winter, the inversions are often associated with poor air quality and poor visibility that halts aviation activities. In the summer, the inversions frequently contrast with moderate winds aloft – morning in the valley may be cool

and calm, and as the air mixes due to surface heating, the inversion dissipates and allows the winds aloft to reach down to the surface. The ability to track the inversion's daily modification leads to more accurate wind forecasts.




Figure 2: Missoula Bowl

The Missoula area has been described here because of its role as a population center, and for its weather reporting capabilities, but it should be noted that similar patterns exist in complex interactions throughout the Missoula County Warning Area (CWA). Figure 3 shows this CWA, which is the area of interest to the Missoula Weather Forecast Office.

The Missoula CWA has experienced numerous wildfires in recent years, and accurate, localized forecasts are of paramount importance in battling these dynamic systems. The complex weather patterns found in the region, coupled with weather generated by the fires, makes this a daunting task. Evolving fire prediction models of interest to the Missoula Fire Sciences Lab (FiSL) are geared towards running CFD code on a laptop and generating detailed wind information in the mountains and valleys given upper level trends based on gridded wind data. Gridded wind data is obtained from course-resolution models and the ability to introduce more realistic wind fields as input would be desirable.

Hence, our motivation is to ultimately produce high-resolution forecasts that capture the numerous small-scale processes found in “rugged” regions and effectively couple such processes with the larger atmospheric system. Many in the weather modelling community will argue that high-resolution simulations suffer due to unavailability of

high-resolution initial conditions, and this point is acknowledged by the authors. However, at the same time, we recognize that at some point these small-scale processes need to be dealt with in the absence of parameterisations, and it is our intent to pursue these high-resolution simulations in order to begin determining where improvements need to be made.



Figure 3: Missoula CWA

WRF

The Weather Research and Forecast (WRF) modelling system, maintained and supported by the Mesoscale and Microscale Meteorology Division of NCAR, is advertised as providing a framework suitable for use in a broad range of applications across scales ranging from meters to thousands of meters. WRF is the product of community collaborations involving NCAR and several government agencies including the National Weather Service. It builds on widely-used models and is slated to replace regional versions of the popular Eta model. WRF is portable across a number of Unix and HPC architectures, making it attractive for research and development on platforms ranging from home Linux machines to high-end supercomputers.

Like its predecessor, MM5, WRF supports nested domains to provide high-resolution simulation over “interesting” regions and coarser resolution – reducing the

computational load – over outer regions still necessary to drive the forecast. Many in the weather modelling community would like to see more regionally based forecasting performed in weather forecast offices for a localized, customized model for a particular part of the country. With the positioning of WRF as a well-supported, widely-used weather model, our goal is to initiate research to investigate and improve upon its utility in providing the high-resolution forecasts that we need in our unique region.

Computing Platforms

To date, we have been executing a daily 72-hour forecast with WRF at 7.5km resolution over the Missoula CWA (a 73x86 grid) on a cluster of two dual-CPU Linux workstations with 2.8 GHz Intel Xeon processors and 2 Gbytes memory per node. With a timestep of 40 seconds, simulation walltime is approximately twelve hours.

In Autumn 2004 a 24-processor Cray XD1 was acquired by the USDA FS Fire Sciences Laboratory (FiSL) in Missoula, Montana and housed in the National Weather Service machine room. The primary purpose for the XD1 is to run WRF-Chem for smoke plume simulations, and we were graciously invited to move some of our work to this new machine. The FiSL's XD1 is built with 2.2 GHz AMD Optron processors running Linux. The primary focus of this paper is to describe our experiences in attempting high-resolution WRF simulations on the XD1 and determine what degree of resolution is reasonable for us. The compiler used was PGI version 5.2. Compilation was performed using default options in the WRF configuration, with MPI as the IPC protocol.

Additionally, through the generosity of the Arctic Region Supercomputing Center we have been provided with the opportunity to use a collection of 92 tightly-coupled IBM p655+ servers. Each server possesses eight 1.5 GHz Power4 processors and is supported by 16GB of memory. Initial use of this resource gives us something to compare our XD1 performance with, and also, due to the large number of processors, might provide some indication of the performance to expect if the XD1 is upgraded with more processors. Compilation was performed with the xlf90 compiler using default options in the WRF configuration, with MPI used for communication.

2. The tests on uniform domains

The first tests were performed on a uniform 7.5km grid covering the Missoula CWA. The 73 x 86 grid (6,278 points per level) has 75 vertical levels. Timings, displayed in Figure 4, are for a 6-hour forecast using a timestep of 20 seconds.




Figure 4: Timings for 7.5 km grid, 6-hour simulation

During the testing, it was discovered that choice of rows and columns for a particular processor partitioning had a significant effect on performance, and in general, a column-major partitioning (number of columns less than number of rows) provided better performance. This is illustrated in Figure 5.




Figure 5: Performance vs. partitioning, 7.5km

A disturbing observation is that there is significant difference in simulation output between the XD1 and the p655+. Figures 6 and 7 show surface pressure (isobars) and temperature (filled, colored contours) after three simulation hours. Although temperatures are not radically different, the isobars in the northeast region differ significantly and may result in further model divergence in subsequent timesteps.




Figure 6: XD1, surface pressure and temperature at 3 hours, 7.5km resolution




Figure 7: p655+, surface pressure and temperature at 3 hours, 7.5km resolution

The next tests were performed on a uniform 2.5km grid covering the same region as the previous example. This 216 x 264 grid (57,024 points per level) had 75 levels, using a timestep of 8 seconds. Timings are displayed in Figure 8. With this simulation of respectable size, the higher speeds of the XD1 processors become evident. On the XD1, the nearly one hour of walltime required for a three hour forecast is probably too slow for producing regular real-time forecasts, but the trend showed by the p655+ suggests that an increase in the number of XD1 CPU's might put this problem into one that's manageable for our purposes.

Again, as seen in Figure 9, both machines exhibit a preference for a particular processor partitioning, and it is advisable that users pay attention to this, especially for long simulations.




Figure 8: Timings for 2.5km, 3-hour simulation




Figure 9: Performance vs. partitioning, 2.5km, 3-hour simulation

Close observation of the p655+ performance on 64 and 128 processors began to reveal an issue in tackling large problems. The WRF modelling system actually consists of three primary steps

1. Retrieval and interpolation of input and boundary conditions. This step essentially takes gridded data from various sources and maps it to the domain and resolution chosen for simulation. In the problems discussed in this paper, this process typically took half a minute up to five minutes for the largest problems. Longer forecast times, of course, require more of this pre-processing.
2. Real data initialization, using *real.exe*. This step uses the pre-processed input data and creates the files to be used as initial and boundary conditions in the simulation. For

most of the problems *real.exe* took a small amount of time, but in cases of deep nesting, large datasets, and/or a large number of processors, the time taken to run *real.exe* was excessive, and this will be discussed shortly.

3. The numerical weather simulation is performed by *wrf.exe*, using the input files created by *real.exe*.

All timings discussed in this paper, unless otherwise noted, were taken as a total time for the execution of *real.exe* and *wrf.exe*. Both of these routines can be run in parallel. The initial pre-processing, even for large problems, was deemed to take a negligible amount of time for execution relative to long-term forecasts. Additionally, the pre-processing is currently a serial task so wasn't included in the timings.

In the 7.5km-resolution problems and most of the 2.5km-resolution problems the cost of *real.exe* was negligible, on the order of one to two minutes. However, when running the 2.5km on 128 processors (see Figure 8), we noticed that almost half of the 39 minutes were being used by *real.exe*. The execution of *wrf.exe* required approximately 17 minutes. Thinking that the relative cost of *real.exe* might become more negligible with longer simulations, we also obtained timings on the p655+ system for 6-hour forecasts on 64 and 128 processors (Figure 10). In this test, we find that performance is slightly better with 128 processors (as opposed to the 3-hour simulation shown in Figure 8), and recognize that longer simulations might see more of a benefit. However, longer simulations also require more time in *real.exe* to create the initial and boundary condition files.




Figure 10: 2.5km, 6-hour simulation on p655+

We attempted further tests on a uniform grid of 833m. This was a 635 x 789 (501,015 points per level) grid with 75 levels. Unfortunately, the large size of the required input files made this prohibitive. On the XD1 we were able to pre-process the input data, but with only 1

Gbyte per processor available, *real.exe* took a prohibitive amount of time (we killed it after 20 hours). On the p655+ system we were unable to perform initial pre-processing due to array allocation issues. In both cases we believe there may be alternatives, and this is an area that warrants further exploration. An additional test that uses nesting to achieve 833m resolution in a select area is discussed in the next session.

Although we were unable to run a simulation on a uniform 833m grid over the entire Missoula CWA, we suspect that even with many processors, the costs would have been prohibitive. With a 2.5km uniform grid, the 22 processors available on the FiSL XD1 would be acceptable for research efforts, but probably not for real-time forecasts (it would take 24 hours to run a 72-hour forecast). However, results on the p655+ indicate that adding more processors to the XD1 will allow for real-time forecasts at this resolution. Figures 11 and 12 illustrate the difference in topography resolution we will realize, and Figures 13 and 14 illustrate the increased detail in fields we can expect to see by running 2.5km simulations on a regular basis.




Figure 11: Missoula CWA topography, 7.5km resolution




Figure 12: Missoula CWA topography, 2.5km resolution

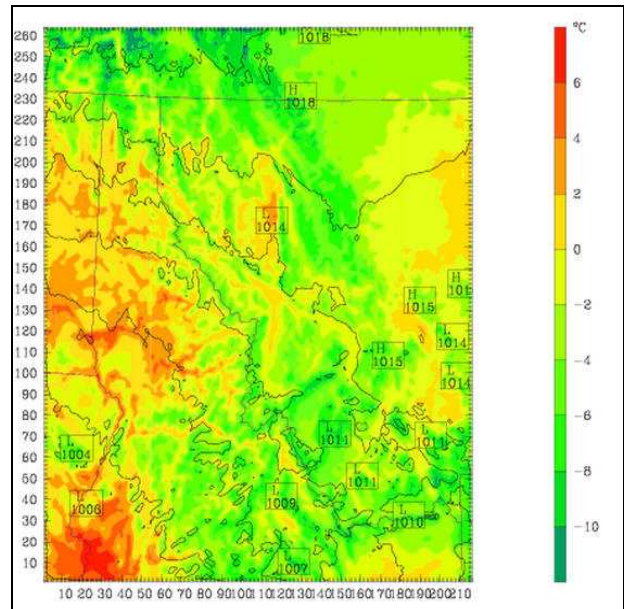


Figure 14: Surface pressure and temperature at 3 hours, 2.5km resolution




Figure 13: Surface pressure and temperature at 3 hours, 7.5km resolution

3. The tests on nested domains

WRF presents the ability to perform nested runs, allowing for high-resolution computations only in regions of interest rather than across the full forecast area, thus resulting in fewer computations. We tested the performance of a three-nest domain over the Missoula CWA, with highest resolution in the Missoula vicinity (Figure 15). Domain configuration was as follows:

- Outer domain – 73 x 86 points at 7.5km resolution
- Middle domain – 97 x 109 points at 2.5km resolution
- Inner domain – 124 x 139 points at 833m resolution

A small sample of timings for this configuration is presented in Figure 16. Although the 22-processor timings on the XD1 suggest that this isn't a usable configuration for real-time simulations, the timings for 32 and 64 processors on the p655+ indicate that there is hope for real-time simulations at this high degree of resolution if we have more processors available. The benefits of achieving this resolution are exciting, and depicted in Figures 17-19. The reader should verify that the topography illustrated in Figures 17 and 19 begin to look much like the graphic of the Missoula Bowl shown in Figure 2. Further analysis of Figure 19 reveals surface winds affected by the topography of narrow valleys, and

this is exactly the behaviour we hope to capture in our efforts to realistically model weather in “rugged” regions like the northern Rockies of the United States.




Figure 15: Missoula CWA, nested from 7.5km to 833m resolution

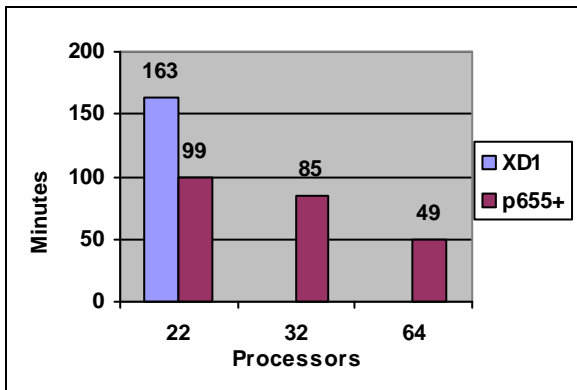


Figure 16: Timings for nested 7.5km to 833m, 3-hour simulation




Figure 17: Topography of innermost nest, 833m resolution. Missoula Bowl is just northeast of center.




Figure 18: Surface pressure and temperature, Missoula Bowl, 833m resolution.




Figure 19: Wind fields and terrain, Missoula Bowl, 833m resolution

4. Conclusions

Our group has a somewhat idealistic goal of producing numerical weather forecasts that realistically capture the complex behaviour of mountainous regions. Clearly, this goal requires advanced computational resources. With the advent of the affordable Cray XD1, we believe that the resources necessary for these computations may be just within reach. We only had access to a 22-processor XD1, but performance results suggested increased scalability, and we look forward to the utilization of larger XD1's for further testing.

Another feature of WRF not mentioned in this work, but critical for high-resolution modelling in remote regions such as the northern Rockies, is the assimilation of point observational data into the simulation. Although the Missoula CWA is sparsely populated, there are a number of remote weather stations whose observations can be utilized to “guide” simulations and provide better initial and boundary conditions than those obtained through gridded datasets, thus beginning to address concerns that high-resolution models lack the necessary initial conditions for a reasonable forecast.

The purpose of the work described in this paper was to determine the feasibility of pursuing high-resolution forecasts in “rugged” regions. The results of this study suggest that the available computational resources are sufficient to warrant continued effort, and are certainly

adequate to begin addressing issues of model verification at high resolutions.

Acknowledgments

The authors wish to thank Peter Johnsen and Frank Chism of Cray for their assistance in learning to use the XD1 and in porting WRF to it. Additionally, we would like to thank the Arctic Region Supercomputing Center not only for their use of the p655+ servers, but for their excellent User Services support. Finally, we thank the USDA FS Fire Sciences Laboratory for the extensive use of their new XD1, and the Missoula Weather Forecast Office of the National Weather Service for their part in housing and maintaining the system.

About the Authors

Don Morton is a long-time visiting scientist with ARSC since their beginning days and is currently with the University of Montana's Department of Computer Science; contact information is Department of Computer Science, The University of Montana, Missoula, MT 59812, Don.Morton@umontana.edu. Gene Petrescu is the Science Operations Officer at the National Weather Service Weather Forecast Office in Missoula, Montana; contact information is National Weather Service, PO Box 9350, 5765 Highway 10 West, Missoula, MT 59807, Eugene.M.Petrescu@noaa.gov. Bryce Nordgren is Physical Scientist for the Fire Chemistry Project at the USDA FS Fire Sciences Laboratory in Missoula, Montana; contact information is Fire Chemistry Project – FiSL, 5775 W Hwy 10, Missoula, MT 59808, bnordgren@fs.fed.us.