

# Modeling the Weather on a Cray X1

Tony Meys  
Army HPC Research Center / Network Computing Services, Inc.  
Minneapolis, Minnesota  
May 15, 2003

**ABSTRACT:** *The Army High Performance Computing Research Center (AHPCRC) installed two early-production Cray X1 computers in September 2002. In March 2003, the Center received a production level Cray X1 LC. This paper will discuss code migration, optimization, testing, and use of two numerical weather prediction codes on the Cray X1: the Fifth-Generation NCAR / Penn State Mesoscale Model (MM5) and the Weather Research and Forecasting (WRF) Model. Early experiences working with these models on the Cray X1 will be discussed along with a sampling of results.*

## 1. Introduction

### *NCSI and the AHPCRC Mission*

The Army High Performance Computing Research Center (AHPCRC) is a collaborative effort between the United States Army, university partners, and Network Computing Services, Inc. (NCSI). NCSI provides infrastructure and support for high performance machines used for AHPCRC research. As part of its role within the AHPCRC, NCSI installs, evaluates, and maintains HPC resources. Additionally, NCSI provides highly trained staff scientists and user support specialists to assist the Army and its university partners.

### *Weather and the Battlespace Environment [1]*

*“Weather significantly impacts on the feasibility of using military force and on ensuing operations. It impacts differently on various types of forces and, in some cases, dictates the types of forces that can be employed effectively. Weather data is part of the intelligence information required by commanders and staffs to plan and conduct combat operations. The answers achieved by analysing weather data, identifying weather effects, and assessing the impact of weather on systems, tactics, and operations provide vital information for commanders to optimally employ their forces.”* Army Field Manual 34-81 (Weather Support for Army Tactical Operations)

Over the millennia, military commanders have recognized the important role of weather in battle. Despite this recognition, the history of warfare is replete with battles lost because of the failure to adequately predict the weather or to properly understand the potential impact of weather on combat operations. Napoleon and Hitler lost their eastern

front wars because of a failure to prepare for cold weather. Other battles have been lost by the inability to predict storms or the effects of rain, heat and wind-blown sand. Cloud cover or man-made obscurants can hide military forces, a sudden change in weather conditions can make them visible and vulnerable. Because of this, the Army, Air Force, Navy and the Marine Corps all have substantial operational activities for gathering, sharing and predicting weather and for the dissemination of weather information to battlefield commanders. In addition, the Armed Services rely on other Government agencies, such as the National Oceanographic and Atmospheric Administration (NOAA) for the development of state-of-the-art weather forecasting models and detailed weather information.

Each of the Armed Services also has unique military related weather requirements. As explained in Army Field Manual 3-100/MCWP 3-3.7.1, the Army must evaluate and address the threat of enemy nuclear, biological and chemical weapons, smoke, non-lethal and flame weapons – weapons whose local effectiveness is substantially affected by weather conditions. In addition, local weather conditions have a significant effect on observability – both the ability to see an adversary as well as the ability to remain undetected. Because of this, the Army has particularly stringent requirements for localized weather forecasts and for analyses of the impact of weather on “ground terrain”. For the Army, “local” means under 5,000 feet above the ground and within a one-to-two kilometer range. Ground terrain could include desert, open fields, mountainous regions, or urban areas.

Weather forecasts require accurate initialization data (i.e., to predict tomorrow’s weather in a given region, it is necessary to know today’s weather at the region boundary),

good weather prediction models, and powerful computers with large memory capable of running the prediction models.

Within the United States, the initialization data usually comes from the National Weather Service/NOAA and the National Aeronautics and Space Administration (NASA) and includes both ground- and satellite-based data collection systems. NOAA and NASA provide relatively complete weather coverage for the U.S., but detailed ground weather station data is likely to be unavailable in foreign battlespace areas. For this reason, the Army has developed a deployable ground weather station. But still, complete coverage of a hostile battlespace environment is doubtful so initialization data is unlikely to be complete. Providing for accurate predictions in the absence of complete initialization data is one of the challenges confronting the Army.

The Armed Services also rely heavily on the weather prediction codes developed by organizations such as NOAA and the National Science Foundation (NSF). One of the major codes currently used is the Fifth-Generation National Center for Atmospheric Research (NCAR)/Pennsylvania State University Mesoscale Model (MM5). As an example, one recent AHPCRC project has been resource support for MM5 forecast runs for the National Test Center, Fort Irwin, CA. Robert Dumais, ARL/CISD, is the scientist leading that work. (See also, Figure 8.) The Army needs a finer resolution forecast capability particularly in the areas of nuclear, biological and chemical defense.

The AHPCRC in conjunction with the Battlefield Environment Division (BED), Army Research Laboratory (ARL), led by Jon Mercurio, is working to address each of these three areas. The goal of this effort is to create a forecasting capability for various battlespace environment phenomena all the way from the area of operation (kilometres in resolution), to the fine scale (meters in resolution), to the microscale (atmospheric turbulence modeling, molecular chemical interactions and obscuration effects).

### ***The Early Production X1 Experience***

In September 2002, NCSI took delivery of and began integrating into the AHPCRC infrastructure two early production Cray X1 computers [2]. These machines each contain 16 multistreaming processors (MSPs). There are 4 MSPs per Cray X1 node. MSPs share memory on a node, which for the early production X1 computers was 16 gigabytes. A 16 MSP machine, therefore, has 4 node boards installed in it. The early production (EP) machines differed from the production version of the machine in the following ways: 1) half clock rate - the EP X1 computers have a 200 Mhz scalar, and a 400 Mhz vector clock frequency; 2) early revision chip set - as an early production machine, the EP X1 had the early version of the set of chips which make up the MSPs and other on node logic; 3) early version of UNICOS/mp - these two machines experienced frequent OS

upgrading until the production level version of UNICOS/mp became available in early 2003.



Figure 1. The two AHPCRC Cray X1 EP machines in the machine room of NCSI. These computers were used for early development and training.

Although very useful tools for migrating codes, learning about the particulars of the Cray X1 architecture, and getting a rough idea of performance of codes on the production follow-on, working on early production hardware is not an activity for the impatient or faint of heart. Cray Inc. was, however, quite responsive to problems known by them and discovered by NCSI, and NCSI was well aware that the first three or four months on an early production system would include these kinds of issues. Most common among the early issues with the X1 were slow compiler times, a few common tools not being available (for example, the m4 preprocessor) issues with make, occasional problems with the compiler, and the need to rewrite code that provided an interface between Fortran and C. The latter issue was a result of the Cray X1 adopting a more industry standard approach to data sizes.

### ***A Hybrid Machine: MPP and PVP***

Perhaps the greatest value of the early exposure we had to the X1 through the early production systems was the opportunity to experiment with and better understand the X1 architecture. Through my experience with the machine, I very much came to regard the X1 as a hybrid of the Cray T3E and the Cray SV1. With OpenMP or other shared memory style parallel models not yet supported, MPI, shmem, or one of the other supported distributed memory programming models is a must to make use of more than one MSP. Working with distributed memory applications on the X1 has many similarities to using a T3E with two major exceptions: 1) the MSP processor, and 2) the availability of much more memory to a single processor. The availability of more memory to a single processor (or MSP in X1 jargon) is a very big plus. The T3E was limited in its ability to host large single processor, memory intensive jobs that often are required in the pre-processing

of data for weather models. The X1 can provide a node's worth of memory for such jobs.

The multi-streaming processor (MSP) makes the Cray X1 a lot like predecessor PVP machines. The inclusion of vector registers, pipes, and high bandwidth low latency access to local memory are the hallmarks of Cray PVP machines. Differing from that tradition is the use of four sub-processors (SSPs) ganged together in hardware and software to appear to the user as a single processor (the MSP). Inner loops with long vector content and an outer loop that would have been autotasked or microtasked in the past perform well on an MSP in most cases. As discussed below, however, there are situations where an analyst must be keenly aware of all the different levels of parallelism. Vectors, multistreaming loops, and distributed parallel memory with inter-MSP communication can all be happening simultaneously in a code, and it may be necessary to help the compiler find all of these opportunities to enhance performance.

### ***Moving to a Production Environment***

In late February 2003, NCS took delivery of and began integrating into the AHPARC infrastructure one of the first production level liquid cooled Cray X1 computers. This machine has an 800 Mhz vector clock, 400 Mhz scalar clock, and 8 nodes. Each node has 16 Gbytes of memory and is populated with 4 MSPs for a system total of 28 application MSPs and 4 command MSPs. Results presented in this report will be from the production machine. The two half-clock, air-cooled, early production machines will eventually be upgraded with production nodes and at least one will remain as part of the NCSI/AHPARC computer environment.



Figure 2. The production level Cray X1 LC (liquid cooled) machine following installation in the NCSI machine room for the AHPARC.

## **2. Evaluating NWP on the X1 with MM5**

### ***About MM5***

The Fifth-Generation NCAR/Penn State Mesoscale Model (MM5) [3,4] is the current version of a weather simulation system which has been in development and use for about thirty years. The model's vertical representation uses terrain-following sigma coordinates. In the horizontal, it uses a structured grid layout. Both the vertical and horizontal resolutions of the model can be adjusted to fit the needs of a given experiment. Additionally, MM5 supports the nesting of domains. This allows a user to define an outer mesh at a relatively coarse resolution so that an inner mesh can have boundary conditions fed into it at all sides. Nesting domains multiple times within the model is possible. Because a large community of researchers have used and contributed to MM5, it has a number of options for treatment of solar radiation, longwave radiation, cloud formation, interactions between air and land surfaces, and other physics.

### ***When to use HPC for Mesoscale Modeling***

A weather model such as MM5 can be run on almost any type of computer ranging from a laptop to an HPC system [4]. What determines the platform of choice is mainly based on the geographical size, resolution, and forecast length of the problem. For example, useful real-time simulations can be performed on 1 Ghz class Intel machines if the problem is limited to about 24 forecast hours, the domain sizes kept at 20 to 30 levels with fairly localized coverage (500 to 1000 km), and resolutions of inner and outer nests kept at near 15 km. Data analysis and a forecast run using MM5 can be done on this kind of problem with a desktop machine in a time window of about 2 hours. In this case, the number of grid points for either mesh in the horizontal will be on the order of 50x50. Increase resolution or coverage area for a forecast, however, and it is readily apparent that a much larger compute engine is needed to do operational forecasting or research in a timely manner. Doubling the horizontal resolution and the number of vertical levels of the above hypothetical domain increases the amount of memory required by a factor of 8. To do the same amount of work in a two hour period, the speed of the machine would have to also increase, but by a factor greater than 8 times since, in addition to more grid points to calculate, the time step of the problem must be decreased to maintain numerical stability at the higher resolution. This is why computers specifically designed to address HPC class problems still have a place in numerical weather prediction at all spatial and temporal scales.

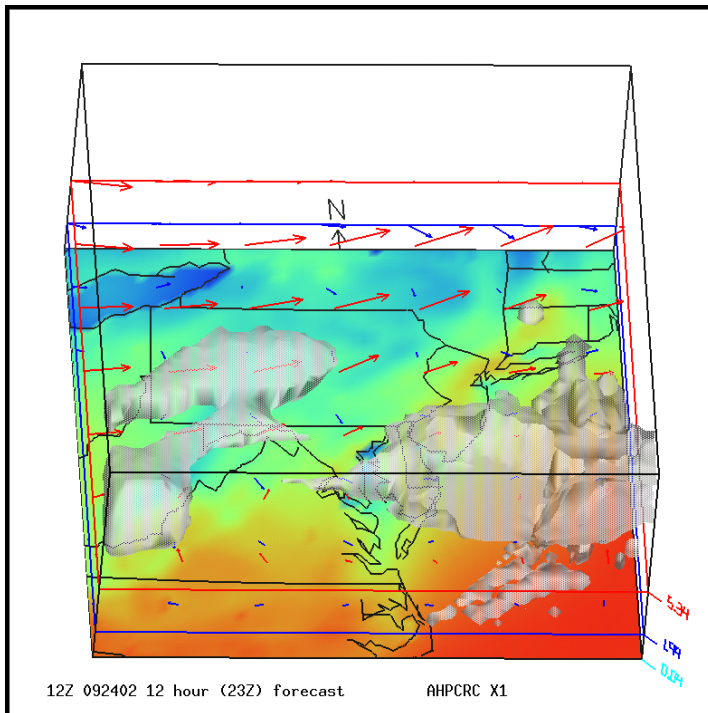


Figure 3. An example of visualized output from one of the first MM5 runs performed on an AHCRC Cray X1. In this case shades of color show surface temperature, arrows wind direction and magnitude, and isosurfaces clouds.

#### Initial Functionality and Performance Tests

Since the first X1 was installed a number of studies have been done to evaluate its functionality for numerical weather prediction (NWP) using MM5. Among the experiments have been:

- Performing a two nest MM5 run with outer mesh of 15 km spacing, 127x127 horizontal grid points, 47 levels, with inner 5 km mesh also at 127x127 by 47 levels, centered over Minneapolis, MN.
- A triply nested run similar to the above but with 23 levels in the vertical and resolutions of the three grids being 10 km, 5 km, and 2.5 km. This domain was centered over Atlanta. Leon Osborne (University of North Dakota) provided LAPS data analyses for the large configuration Atlanta experiments.
- Other smaller domain configurations and tests run in real-time with current weather data. These domains were centered on Minneapolis, Baltimore, and Atlanta.
- Tests using the NCAR Standard Benchmark. The information from these tests is discussed below.
- Large domain (continental United States) at 5 km resolution model runs. The information from these tests are in Section 4.

#### Results from the NCAR Standard Benchmark for MM5

One of the first tests performed on the Cray X1 LC was the NCAR Standard Benchmark for MM5 [5]. This test uses a data set with the following configuration:

- horizontal grid of 112 x 136
- 33 levels in the vertical
- 81 second time step
- 3 hour forecast

A standard count for floating point operations for this case was provided on the NCAR MM5 web site from which the Standard Benchmark is available. Since for our tests we were using the latest version of MM5, version 3.6.0, and it was unclear (and unlikely) that comparison information from earlier tests on other machines reported on that web site used 3.6.0, it was decided that floating point operations would be based on a known reliable count. The known reliable count chosen was the performance of the same single processor run on an SV1ex using hpm to count operations. Because of this methodology, the GF/sec in Table 1 are, very technically speaking, SV1ex equivalent floating point operations per second. The hpm count for the case on a Cray SV1ex was 290566 million floating point operations. Timings and performance metrics in Table 1 are for that part of the 3 hour benchmark run that includes all time steps, but does not include model initialization or shutdown time.

MSPs	time (seconds)	GF/sec
1	97.4	2.98
2	53.6	5.42
3	39.4	7.38
4	31.2	9.31
5	25.3	11.49
6	22.4	12.97
7	19.7	14.75
8	18.6	15.62
9	16.8	17.29
10	16.6	17.51
11	15.3	18.99
12	14.2	20.46
...	...	...
16	13.7	22.01
...	...	...
20	10.8	26.89
...	...	...
24	10.5	27.67
...	...	...
28	10.2	28.49

Table 1. NCAR Standard Benchmark on Cray X1 LC



### 3. Performance Tuning MM5 for the X1

#### Optimizations for Better X1 Performance

Migration of MM5 to the Cray X1 resulted in a functional implementation of MM5, but not an optimal one. The primary early performance issues with the model fell into the following categories: 64 versus 32 bit mode, areas of the code where multistreaming was inhibited due to the structure of MM5, and the requirement to decompose MM5 configurations such that long inner loops of vector work would be preserved. In the discussion below, each improvement is described along with performance comparisons. For these tests, version 3.6.0 of MM5 was used and the model configuration and data from the NCAR Standard Benchmark was used when executing tests.

#### 64 bit vs. 32 bit

Unlike prior machines carrying the Cray brand, the X1 has a default precision of 32 bits for both floating point and integer arithmetic. Although it is envisioned that most NWP work on the AHPARC Cray X1 machines will use the default 32 bit mode, MM5 was initially migrated as a 64 bit application to avoid some very early issues in the pre-production X1 environment. Those issues have long since been resolved. Given the tradition of 64 bit default precision Cray machines, however, it is still useful to start with a performance comparison with the new 32 bit default.

Code version (1 MSP tests)	Timing (seconds)	Speed-up
Base (no opt, 64 bit)	218	1
Base (no opt, 32 bit)	195	1.12

Table 2. 64 bit vs. 32 bit version of code

In the remainder of this review of performance improvements, the "Base Case" will refer to the 32 bit version of MM5 with no additional optimizations.

#### Multistreaming in MM5

As previously noted, the Cray X1 makes use of the concept of multistreaming to make 4 SSPs work together as one MSP. There are, however, structures in code which inhibit the ability of the current compiling system to make use of all parts of an MSP. In situations where the compiler can not determine if a section of code is safe, but an analyst knows it to be free of dependencies, "csd" directives can be used to force the compiler to stream a loop [5].

In the case of MM5, large loops in the modules solve.F and sound.F were found not to be streaming. The two most frequent reasons for this were loops which call functions that perform work on data and inner loops with the possibility of dependencies within the limited scope of the information the compiler had access to for a module. Many of these non-streaming areas were the same loops which had

been microtasked (e.g., "cmic" directives) previously in the code. Initial work on adding in csd directives for sound.F and solve.F was done by Frank Kampe, Cray, Inc. The following chart shows the effect of adding the csd directive changes to the 32-bit version of MM5.

Code version (1 MSP tests)	Timing (seconds)	Speed-up
Base (no opt, 32 bit)	195	1
Streaming mods on	96	2.03

Table 3. Multistreaming modifications

#### Impact of Model Distribution on Vectorization

Keeping vector length sufficiently long to make use of the vector registers of the X1 is critical for optimal performance. When working with small to medium sized MM5 test cases, it is important to ensure that the way a domain is decomposed to be spread over multiple MSPs does not adversely impact vector performance. Table 4 shows what was observed when the default pattern of decomposition of the NCAR Standard Benchmark domain was used. Table 5 shows the improved results when the domain was decomposed so that each MSP received a set of south-north strips. (See Figure 4).




Fig. 4(a) Default layout on 4 MSPs

Figure 4. Comparison of MM5 default distribution of work to MSPs and distribution used to maintain longer vector lengths.

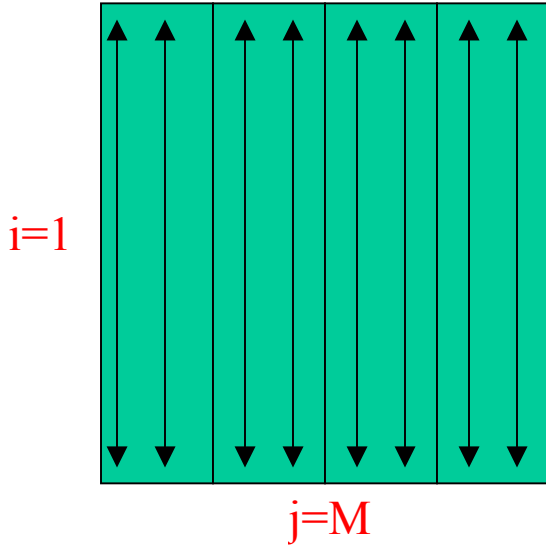


Fig. 4(b) Layout to maintain vector lengths.

MSPs	Time (seconds)	Speed-up
1	99.6	1
2	53.7	1.85
3	39.7	2.51
4	39.8	2.50
5	25.4	3.92
6	29.1	3.42
7	19.5	5.11
8	23.1	4.31
9	27.5	3.62
10	19.4	5.13
11	15.6	6.38
12	21.5	4.63
13	13.2	7.55

Table 4. Default decomposition

MSPs	Time (seconds)	Speed-up
1	99.6	1
2	53.7	1.85
3	39.7	2.51
4	31.5	3.16
5	25.4	3.92
6	22.6	4.41
7	19.5	5.11
8	19.1	5.21
9	16.6	6.00
10	16.2	6.15
11	15.6	6.38
12	14.0	7.11
13	13.2	7.55

Table 5. Decomposition to keep S-N inner loop vectors as long as S-N domain dimension

#### 4. Some Bigger MM5 Test Cases

##### *Minnesota with Higher Vertical Resolution*

During the initial porting of MM5, the number of levels in the vertical was kept at 23. One of the first larger tests configured after the smaller versions of the code were verified was a 127 by 127 horizontal, 47 level model run. Initialization data for this case was derived from Eta-coordinate output from the August 21, 2002, 12Z model run conducted by the National Center for Environmental Prediction (NCEP) [7].

This case, centered on Minneapolis, Minnesota used a fairly typical mesoscale forecast resolution of 15km in the horizontal. For testing purposes, the length of the model run was kept at six simulation hours. Performance for this case on a Cray X1 LC are shown in Table 6 on various combinations of MSPs. A graphic is also shown that illustrates the extent of the domain and the meteorological phenomena modelled in the run. See Figure 5.




Figure 5. Example output from a 127x127, 15 km spacing, 47 level MM5 simulation using August 21, 2002, 12Z NCEP Eta model data for initialization.

MSPs	time (seconds)	GF/sec
1	71.3	3.96
2	39.2	7.19
3	28.5	9.90
4	23.0	12.26
5	20.4	13.82
6	17.1	16.49
7	15.8	17.85
8	13.7	20.58
9	13.5	20.89
10	11.7	24.10
11	11.4	24.74
12	11.1	25.41
13	11.1	25.41
14	10.9	25.87
15	9.4	30.00

Table 6. Performance results for the 47 level case.

**Configuration for a National 5km Test**

Although configurations of MM5 like that used in the Standard NCAR Benchmark provided performance data points that assisted with early migration and optimization, larger configuration runs were needed to more fully test the Cray X1. To do this, a data set was constructed at 5 km resolution for a domain that covered the entire continental United States. Versions of the data set were constructed with 23, 33, and 47 levels in the vertical. All versions had north-south and east-west dimensions of 680 by 1000.

Reconfiguring the model for this large of a domain involved only re-dimensioning in the i,j,k dimensions. The RSL layer of MM5 did have one parameter that controls the maximum j-length of a dimension that required upward adjustment. After these changes, compilation of the model proceeded as it did for smaller test cases.

More challenging was the construction of the input data needed to do the forecast run. An initial attempt to build the input terrain file and MMINPUT\_DOMAIN file for the large configuration MM5 proved too memory intensive to be easily handled by a Linux server that had been used for the pre-processing of smaller input data sets. The work around for the problem was to migrate MM5 pre-processing routines TERRAIN, PREGRID, REGRID, and INTERPF to the Cray X1. Although the X1 had ample memory to generate the larger input data files, the execution time for pre-processing was not trivial. All of these routines have a substantial amount of input/output and scalar work, neither of which are the strong points of the Cray X1 at present.

**Performance Results**

A series of tests were run using the large case input data. Timings were made for model initialization, time per forward model step, model integration time (the sum of all forward model steps) and model shutdown time. All of the tests were run for 24 simulation hours.

One of the largest MM5 weather simulations tested to date on the Cray X1 has been a 33 level configuration with 5 km spacing over the entire continental United States. During a trial run of this model, the Cray X1 sustained 36.7 billion floating point operations per second on 16 MSPs while executing the forecast steps of the model. This allowed the model to simulate one hour of atmospheric physics and dynamics in 8.4 minutes on average, for a model integration time for 24 simulation hours of just under 3.5 wallclock hours. The problem required 20 billion bytes of memory. As a point of reference, current operational weather models that cover all of the United States are typically run at a resolution of about 10km. Doubling the resolution of MM5 from 10 km to 5 km increases its memory requirements by at least a factor of four and its compute requirement by a factor of approximately eight.

The 33 level, 5 km resolutions national domain case was also run on 8 MSPs. It was found that it took 6.4 hours to do the forecast calculations with a sustained wallclock performance of 19.253 Gflops/sec. Comparing this with the approximately 3.5 hours model integration time on 16 MSPs results in a speed-up of about 1.9 times from 8 to 16 MSPs, which is consistent with what can be expected from MM5 on this case.

As a final test, a 47 level version of the 5 km national test case was configured. In this instance, a total of 12 MSPs were utilized to perform a 24 simulation hour model integration. The performance observed for the model integration time was 29.5 Gflops/second. Performing the model integration took 5 hours, 50 min. On 12 MSPs, this run would probably not be timely for operational forecasting. Further evaluation needs to be done to determine how many MSPs could do this kind of higher resolution vertical and horizontal case within an operational forecast window. Still, the test did confirm that on even as few as 12 MSPs, higher resolution research runs are possible with quite timely turn around on the machine.

Vertical levels	MSPs	Time (minutes)	Gflops/sec
33	8	384	19.25
33	16	207	36.71
47	12	350	29.12

Table 7. Summary table for larger case 24 hour simulations.




Figure 6. Subsection of 5km national domain showing low level winds at each gridpoint.

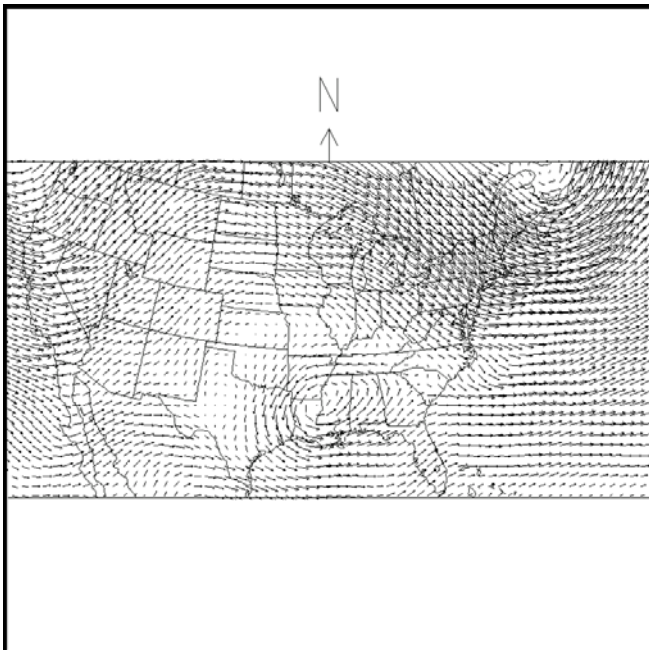


Figure 7. Continental United States domain showing mid-level winds from test case. Note that wind barbs are shown at a density of one-tenth the actual number of grid points.

## 5. Weather Research and Forecast Model

### *Introduction to WRF*

Although MM5 remains the workhorse mesoscale numerical weather prediction model for a large part of the operational and research community, including the Air Force and Army, there is considerable interest in the Weather Research and Forecasting (WRF) Model, a regional and mesoscale model which is being jointly developed by the research community. It is anticipated that AHPCRC users doing weather and environmental work will have an interest in using this code on the Cray X1. As a result, work has begun to provide users with a working version of WRF on the Cray X1[4,8].

### *Migration of WRF to the X1*

Migration of WRF to the X1 has proven to be more challenging than the initial MM5 port. The difficulties have resulted from a combination of WRF complexity and several issues with the early Cray X1 programming environment.

On the WRF side of things, a complex system of makefiles and configuration tools have required time to understand and modify. One of the advantages of the WRF system is the ability to specify a target machine at the time of model set-up, and then have the makefiles and configure utilities automatically take care of most of the build. This works well for architectures already built into the WRF configuration system. Adding a new machine, however, is not a simple process. Using configuration information from similar machines (e.g., 32-bit machines with MPI that will use RSL for parallelism) as a template, most of the build configuration for the X1 has been added. There are, however, a few parts of a WRF compile that need to be done by hand at this point.

We have been able to successfully build and run a single MSP test for WRF. Getting WRF to run on multiple MSPs using the RSL/MPI distributed memory model is still a work in progress. One stumbling block that has been encountered in early WRF porting work resulted from early compiler problems in the pre-4.3 UNICOS/mp programming environment. We have recently upgraded to a new version of UNICOS/mp which should address the problem.

### *Status of Multiple MSP Testing*

At the time this paper was prepared, all of WRF had been compiled for the distributed memory version of the model with the exception of a required input and output routine. Any progress made toward getting WRF to successfully run on more than a single MSP will be reported in the presentation associated with this paper during the May CUG.






Figure 8. Shown above is a weather forecast for 19 December (2200Z) computed for the National Test Center (NTC), Ft. Irwin, CA. The view is from the south looking north and shows the predicted cloud formations and that precipitation (light blue) is predicted over the NTC. The dark lines represent California county boundaries. The forecast was computed on the high performance computing systems at the AHPARC. [1]

## 6. Comments and Conclusions

Even with the progress made to date, there are significant improvements remaining to be implemented to improve MM5 performance on the Cray X1. Initialization time, which involves reading in boundary conditions, the initial state of the atmosphere, distributing the data to multiple processes, and the set-up of memory currently takes a long time on the Cray X1. Even for a smaller configuration run the initialization process can take several minutes. Initial reports from analysis of the code indicate opportunities for better multistreaming in some of the functions used at model start-up by RSL.

Further analysis and performance improvement work is also necessary for MM5 cases that have a nested domain within them. Analysis has shown that feedback routines between the outer and inner nest are not multistreaming well and are running at scalar speeds. This bottleneck needs to be removed to improve performance on the Cray X1.

We do not think we can over emphasize the necessity of having vector code when running on the Cray X1. Even modest amounts of scalar work can greatly limit the

performance of applications on the machine. MM5 is a vector computer “friendly” code, but it does contain a fair amount of scalar work.

Time to compile and link MM5 in the Cray X1 programming environment is slow when compared to compile times on other machines. This issue has been discussed numerous times with Cray, Inc., and NCS has been told that ideas for improving the compile time are being investigated.

## References

1. Paul Muzio, “*Weather and the Battlespace Environment*”, AHPARC Bulletin Fall 2002 – Vol. 12 No. 4
2. Paul Muzio, “*The Cray X1 Has Arrived!*”, AHPARC Bulletin Fall 2002 – Vol. 12 No. 3
3. MM5 page: <http://www.mmm.ucar.edu/mm5>
4. Tony Meys, “*Modeling the Weather on the AHPARC Cray X1*”, AHPARC Bulletin Winter 2002 – Vol. 12 No.4
5. Parallel MM5 benchmarks page: <http://www.mmm.ucar.edu/mm5/mpp/helpdesk/20020218.html>
6. Cray, Inc., “*Cray X1 Application Programming and Optimization*”, TR-SV2PO-0
7. NCEP Products page: <http://www.nco.ncep.noaa.gov/pmb/products/>
8. Weather Research and Forecasting (WRF) Model page: <http://wrf-model.org>

## Acknowledgments

The author wishes to thank Frank Kampe and John Salinas of Cray, Inc., for their many contributions to the early evaluation work discussed in this paper. In particular, the work of Mr. Kampe in multistreaming optimizations and the many questions answered about the system by Mr. Salinas were greatly appreciated.

The research reported in this document/presentation was performed in connection with contract/instrument DAAD19-03-D-0001 with the U.S. Army Research Laboratory. The views and conclusions contained in this document/presentation are those of the authors and should not be interpreted as presenting the official policies or positions, either expressed or implied, of the U.S. Army Research Laboratory or the U.S. Government unless so designated by other authorized documents. Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

## **About the Author**

Tony Meys is a User Support Specialist, Atmospheric Science, with Network Computing Services, Inc. (NCS). Mr. Meys has worked with atmospheric and other environmental models on Cray Research and other HPC machines over the past decade. He can be contacted at Network Computing Services, Inc., 1200 Washington Avenue South, Minneapolis, MN 55415 USA. E-mail contact is [tmeys@ahpcrc.org](mailto:tmeys@ahpcrc.org).