

Eric Lenning\* and Ron Przybylinski

*National Weather Service, St. Louis, Missouri*

James A. Nelson, Jr.

*National Weather Service, Salt Lake City, Utah*

## 1. INTRODUCTION

As desktop computing power continues to grow, local mesoscale modeling is becoming increasingly practical and popular in National Weather Service (NWS) forecast offices. Several NWS offices across the country already are running a mesoscale model for use in their daily forecast operations. Furthermore, the flexibility of AWIPS means that viewing the local model data operationally is relatively easy. This is true both for data generated within a forecast office, or for data obtained from an outside institution such as a nearby university.

This paper discusses the general process for configuring, running, and viewing output from a mesoscale model entirely within an AWIPS environment. Although the paper focuses primarily on the Fifth-Generation NCAR / Penn State Mesoscale Model (MM5, Anthes and Warner, 1978) and how these tasks are accomplished at the St. Louis forecast office, the information is applicable to any office that may be considering its own mesoscale modeling program.

## 2. HARDWARE AND SOFTWARE CONFIGURATION

Determining where the model is going to run and obtaining the software needed to run the model are the first steps toward utilizing MM5 data operationally. While the latter task is well documented by the MM5 Tutorial Notes and Users Guide (Dudhia et al., 2001), choosing the platform on which the model will run is a decision for each individual user. The AWIPS hardware at NWS St. Louis includes four servers and five workstations, in addition to a PC used for applications development (the Application PC) and one used for displaying AWIPS case studies in displaced real-time (DRT, also known as the Warning Event Simulator: WES). Of course, NWS users

are given little or no flexibility in choosing hardware or software for their AWIPS system and have little control over when these components may change. Even so, changes to any of the AWIPS components can affect the process of setting up and running the MM5.

We will first discuss the AWIPS hardware and software configuration at NWS St. Louis, which should be similar if not identical to other offices. We then explain the various software packages required for running the MM5 and mention any possible incompatibilities with current or future AWIPS hardware and software.

### 2.1 Hardware Configuration

Until recently, computer hardware that could produce operational mesoscale model guidance was beyond the reach of NWS field offices. Even today, NWS field offices must use the most powerful computer they have for the guidance to be at all timely. For many, including St. Louis, this is the Linux-based DRT (or WES) PC.

As of late 2001, the DRT/WES machine at St. Louis was a Gateway E-5400 containing dual Pentium III 800-megahertz processors, one gigabyte of RAM, and two 20-gigabyte disk drives. The machine was delivered with Red Hat Linux 6.2 but was upgraded to Red Hat 7.0 in the summer of 2001 to support the Warning Event Simulator software provided by COMET. Despite its two processors and large amount of RAM, performance is a major consideration when choosing the size, number, and resolution of domains to be used in the model, as well as the length of forecast to produce. See *Section 6: Operational Considerations*, for more information.

Another important reason for putting the MM5 on the DRT machine is that none of the normal AWIPS tasks are affected. In fact, the other AWIPS servers and workstations are completely unaware that the model is even running. Furthermore, since the DRT machine is not used operationally, the extra CPU load while running the MM5 has no effect on mission-critical work.

---

\* Corresponding author address: Eric Lenning, NWS  
12 Research Park Drive, St. Charles, MO 63304;  
e-mail: eric.lenning@noaa.gov

## 2.2 Software Configuration

Building the MM5 executables on a Linux PC requires Fortran 77 and 90 compilers that support the Cray pointer data type. The Portland Group High-Performance Fortran (PGHPF) Workstation, available for purchase from the [www.pgroup.com](http://www.pgroup.com) website, is the only such compiler recommended by the MM5 User's Guide, although others might theoretically work. Unfortunately, only a few Linux-compatible compilers support Cray pointers, and to our knowledge none of them are freely available. Furthermore, versions of the PGHPF compiler earlier than 3.2-4 are incompatible with Red Hat 7.0 – a fact not realized at St. Louis until after the summer 2001 operating system upgrade. It is possible, even likely, that future AWIPS requirements will call for additional operating system upgrades. If so, the compatibility with Portland Group compilers may continue to be an issue.

A software package not required for running (but extremely useful for configuring) the model is NCAR Graphics. This software was formerly quite expensive but now is freely available from the [ngwww.ucar.edu](http://ngwww.ucar.edu) website. This software is used with the MM5 TERRAIN program when positioning model domains, and with the optional RIP and GRAPH programs that provide a way to look at model output. Source code and binary executables for the latest version of NCAR Graphics (4.2.3 as of September 2001) are available for both Red Hat Linux 6 and 7. MM5 users with the PGHPF compiler should download the binaries built with this compiler.

A third software package, netCDF, is used to convert MM5 output files into a format that AWIPS can read, although this package is not required for running the MM5. The latest version of this software (3.5 as of March 2001) is available from Unidata at the [www.unidata.ucar.edu](http://www.unidata.ucar.edu) website.

## 3. MM5 INSTALLATION AND CONFIGURATION

A special user account controls all of the MM5 tasks on the DRT machine at St. Louis. Having this separate account for running the model means anyone can administer the process without accessing the files of any specific user. Of course, given space limitations, it is not a good idea to install the model directly under the home directory of this user. A better place is under the /usr1 partition, which provides much more space (~18 gigabytes) for the large initialization model grids, terrain data files, and model output grids.

Some special steps are required for installing the MM5 software on the DRT machine. Mainly this is because the AWIPS LDAD firewall isolates the AWIPS network from the rest of the world (Fig. 1), and because the DRT machine at St. Louis is not set up to access machines beyond the Application PC. The best way around these limitations is to use the Application PC itself to download all the files necessary for running the MM5.

There are two ways of going through the AWIPS firewall to get software to the Application PC and eventually to the DRT PC. The first is to use the gateway

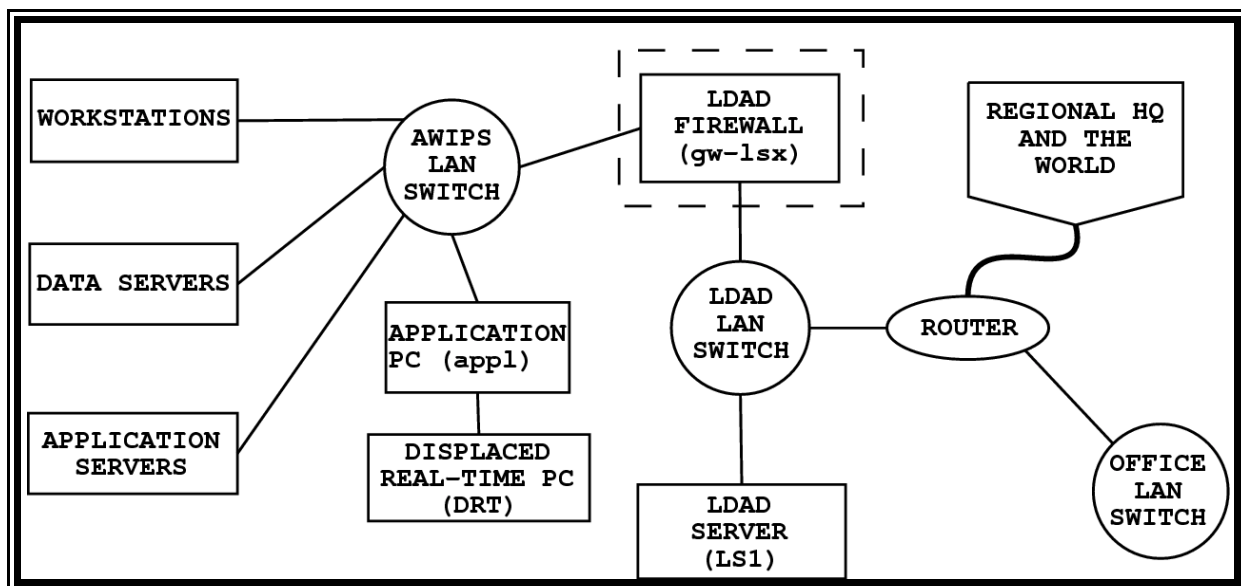


Figure 1. Network configuration at WFO St. Louis (LSX).

FTP service on the firewall. The second and more flexible option is to temporarily configure the Netscape web browser on the Application PC to use the LDAD firewall as a proxy for accessing the Internet. Accessing the web from the Application PC is extremely useful during the process of downloading the Portland Group compilers, the MM5 software, the terrain data files, and the various support packages such as NCAR Graphics and netCDF. Since all of these eventually are transferred to the DRT machine, it also is convenient to have the /usr1 disk of the DRT machine NFS mounted on the Application PC.

The minimal list of programs needed to run the model includes TERRAIN, REGRID, INTERPF, and MM5, all of which are available from [www.mmm.ucar.edu/mm5](http://www.mmm.ucar.edu/mm5), the MM5 Home Page. Although the MM5 User's Guide discusses these in great detail, it is useful to mention aspects of each that are unique to the AWIPS environment.

The first program, TERRAIN, configures the mesoscale domains used in the model and interpolates terrain and land-use data onto these grids from a master set of input files. Because terrain information is static, this program is not re-executed each time the model runs. The TERRAIN program gives you the option of downloading the master set of input files automatically, but since the DRT machine is behind the AWIPS firewall these should be downloaded and decompressed manually before the TERRAIN program is executed. Compressed tar files containing the necessary data are available via anonymous FTP from [ftp.ucar.edu](ftp://ftp.ucar.edu) under the /mesouser/Data/TERRAIN directory.

The REGRID programs prepare initial and boundary conditions for the model by interpolating archived meteorological analyses to the grid and map projection defined by the TERRAIN program. This is necessary because a regional model requires a parent model to feed it information from beyond its limited domain throughout its forecast cycle. The 00-hour analysis from the latest run of the NCEP Eta model initializes the MM5 at St. Louis, while the six-hourly forecasts provide boundary values. These boundary values completely specify the behavior of the outer row and column of the MM5 grid and also influence the next four rows and columns in from the boundary. The Eta forecasts are stored in GRIB format on the 40 km Lambert Conformal grid 212 known as AWIP3D. They are available from the NCEP FTP server at [ftp.ncep.noaa.gov](ftp://ftp.ncep.noaa.gov) under the /pub/data1/eta/erl.yymmdd directory, where yymmdd represents the current 2-digit year, month, and

day. The AWIP3D grid covers much of North America and provides data on pressure levels in 25-mb increments and 3-hour forecast intervals. The file for each forecast projection presently is around 14 MB in size.

The INTERPF program handles the data transformation from the analysis programs (in pressure coordinates) to the mesoscale model (in sigma coordinates) and generates initial, lateral, and lower boundary conditions. No special steps are required to run this program in an AWIPS environment.

Finally, the MM5 program is the numerical weather prediction part of the modeling system. By editing the configure.user file in the MM5 directory and appending "-Mconcur" to the FCFLAGS and LDOPTIONS variables, the MM5 code can be compiled to use both processors on the DRT machine. Also, the NCPUS environment variable must be set to 2 so the MM5 program will know to use two processors at run-time. Another item to consider when setting up the MM5 program is the time resolution for the model output (set using the TAPFRQ variable in the mm5.deck script). While this does not greatly affect the time it takes the model to run, saving data too frequently will produce an unnecessarily large output file. Normally there is no need to save data more frequently than what you wish to see in AWIPS.

## **4. RUNNING THE MODEL**

After the model is compiled and configured, the next step is to run the model. In reality, this involves three separate steps: pre-processing, execution, and post-processing. Pre-processing means acquiring and preparing data for model initialization. Execution means running the model and producing output files. Post-processing means converting MM5 output into a format that AWIPS can display. Each of these steps will be discussed in detail.

### **4.1 Model Pre-Processing**

The main pre-processing task involves obtaining gridded model data for boundary conditions. Although there are various ways to approach this task, at St. Louis it is handled by a set of Unix shell and Perl scripts. Namely, a Unix shell script is started twice a day from a cron job on the DRT machine. This script uses the Unix remote-shell command (rsh) to execute a Perl script on the Application PC. Since the DRT machine cannot

access any machine other than the Application PC, the Perl script handles the acquisition of the latest Eta model grids from the NCEP FTP server. This process can take between 20 and 40 minutes once the data files become available (presently around 2 hours past synoptic time). When the Perl script finishes, the original Unix shell script continues by calling the REGRID programs on the DRT to begin processing the Eta grids (Section 4.2).

Early plans for running the MM5 at St. Louis considered using Eta netCDF files already present on the AWIPS system to save both bandwidth and time, since the model grids arrive automatically over the AWIPS Satellite Broadcast Network. Unfortunately, the AWIPS Eta grids do not contain sea-surface temperature or snow cover fields for initializing the MM5, nor is the 80-km grid scale a small enough horizontal resolution. In addition, software for converting the netCDF files into a format the MM5 programs could use was not readily available.

There are some minor limitations with the NCEP Eta model grids. For example, sea-surface temperature is not archived directly, so skin temperature is used instead. Although ground temperatures and sea-surface temperatures at coastlines can be unrealistic, this is more of a problem for locations near a large body of water than for someplace like St. Louis. Also, snow cover is available in the Eta as water-equivalent snow depth rather than actual depth of snow.

#### **4.2 Model Execution**

The task of actually running the model in an operational setting involves setting up scripts to execute the various parts of the MM5 system (i.e., REGRID, INTERPF, and MM5). Using scripts to run the model is no more difficult than running each part of the system manually. The only difference is that the scripts can automatically feed information about the current date and time to the various MM5 programs.

#### **4.3 Model Post-Processing**

The post-processing task prepares the MM5 output for use by AWIPS or some other display program. To prepare the data for AWIPS, St. Louis uses software developed by James Nelson and Jason Burks of the NWS Western Region Headquarters Scientific Services Division and described in AWIPS Technical Note (ATN) 4.3-01 (Nelson and Burks, 2000). The software converts individual forecast hours of raw MM5 output into netCDF

files, then concatenates selected forecast hours into one large netCDF file that can be used by AWIPS.

If the model is running behind the LDAD firewall as described in this paper, post-processing also takes place behind the firewall. Of course, offices that do not run the MM5 behind the LDAD firewall (or do not run the MM5 at all) still can obtain raw MM5 output files and post-process them behind the firewall. In both of these cases, the final netCDF file is copied directly to the AWIPS Data Server. If the post-processing is performed outside the firewall, the final netCDF file is placed on the LDAD Server machine (Is1 in Fig. 1) and pulled through the firewall into AWIPS by the Data Server. The directory in which this file is stored at St. Louis is called /data/fxa/Grid/LOCAL/netCDF/MM5LOCAL.

### **5. VIEWING MODEL OUTPUT IN AWIPS**

Once the netCDF files are placed in the proper directory on the AWIPS Data Server, there are five steps that must be taken before the D2D program on the AWIPS workstations can display the MM5 data. These steps are detailed in the ATN 4.3-01 mentioned above, but the additional background information that follows may provide a greater understanding of this process.

The first step toward displaying MM5 data in D2D is to give AWIPS a description of the netCDF file created during the post-processing task. A human-readable description of this netCDF data set is provided in the mm5local.cdl file that comes with the Nelson and Burks software. The only local modifications to this file involve setting the number of MM5 grid points and forecast valid times. Once these changes are made, this file should be copied to each workstation.

The next step prepares a supplementary file that D2D uses for navigation purposes when displaying the MM5 data. This task is handled in part by the AWIPS maksuparg program, which must be given the latitude and longitude of the center point, lower left point, and upper right point of the MM5 domain. One source for this information is the terrain.print.out file created after running the MM5 TERRAIN program (when IPRINT=1 in terrain.deck). See the TERRAIN chapter of the MM5 User's Guide for more information. Running the maksuparg command on one workstation produces a file called mm5local.sup, which then can be copied to all the other workstations.

The third step adds the MM5 to the list of AWIPS data sets known as the Grid Source Table. This task also

tells AWIPS where to find the MM5 data and the scales (Regional, State, etc.) on which it can be displayed.

The fourth step tells AWIPS the name used to describe the MM5 data. This name, which must match the name used in the Grid Source Table above, appears on the Volume menu of D2D and under the Source menu of the Volume Browser, as well as the product labels in the D2D display.

The final step for displaying MM5 data in D2D is to run a localization on each workstation. This localization makes use of information provided in the previous four steps. Once this localization is complete and D2D is restarted, the MM5 data should be available for display. Of course, any future changes to the MM5 parameters likely will require all five steps to be performed from the beginning.

## 6. OPERATIONAL CONSIDERATIONS

St. Louis continues to experiment with several domain configurations to determine how best to produce an accurate forecast in a reasonable amount of time. Early tests at NWS St. Louis suggested that synoptic scale features such as troughs and ridges were not able to enter and exit the MM5 domain properly if the domain covered too small an area. This problem is being addressed by running a relatively coarse outer domain covering the continental United States (CONUS), with a smaller but higher-resolution nest over the Midwest.

The present setup at NWS St. Louis uses a 45-km outer domain with 70 north-south and 110 east-west points over the CONUS, and a 15-km inner domain with 64 north-south and 61 east-west points to cover only the State scale on the D2D display at St. Louis. With a two-way nest (IFEED=3 in mm5.deck) and 23 sigma levels, a 24-hour forecast takes around 120 minutes. This is an acceptable amount of time for an operational setting and means MM5 model output is available to the forecaster around 5 hours past synoptic time.

While other model parameters can also affect the time required to produce a forecast, the model grid size and resolution certainly are two of the most important factors. Halving the distance between model grid points increases the workload necessary to produce a forecast by a factor of eight. Besides quadrupling the number of points in the model, this also doubles the number of time steps for which calculations must be made. For example, since the recommended time step is 3 seconds per km of grid resolution, a grid resolution of 60 km requires the

equations to be solved every 180 seconds (3 minutes). Cutting the distance between grid points to 30 km means the equations must be solved every 90 seconds, or twice as often.

## 7. CONCLUSION

We have discussed the general process of configuring, running, and viewing output from the NCAR/PSU MM5 in an AWIPS environment. Since local modeling is becoming increasingly common in the National Weather Service, we hope the information provided here will prove useful for those offices hoping to understand at least one way to run a local model and view the output in AWIPS.

It is highly unlikely that an NWS field office will ever be able to produce operational model guidance for large domains with resolutions (horizontal or vertical) comparable to those used at specialized centers like NCEP or the Forecast Systems Laboratory. Even so, a high-resolution regional model can add value to the local forecast process in two different ways. The first way is by supplementing the national products, possibly by addressing a specific problem like sea breeze convection, lake-effect snow, or quantitative precipitation forecasting. The second way is by providing a better understanding of numerical weather prediction in general. The process of configuring and running a local model is an excellent way to learn more about the national models, which should lead to a corresponding improvement in the forecasts produced when using these models.

## 8. REFERENCES

- Anthes, R.A., and T.T. Warner, 1978: Development of hydrodynamic models suitable for air pollution and other mesometeorological studies. *Mon. Wea. Rev.*, **106**, 1045-1078.
- Dudhia, J., D. Gill, Y.R. Guo, K. Manning, and W. Wang, 2001: PSU/NCAR Mesoscale Modeling System tutorial class notes and user's guide: MM5 Modeling System Version 3. Available online at <http://www.mmm.ucar.edu/mm5/doc.html>.
- Nelson, J.A., and J. Burks, 2000: AWIPS Technical Note 4.3-01: Converting MM5 raw model output for display on AWIPS D2D. NWS WRH SSD. Available electronically at [http://awips.wrh.noaa.gov/build4\\_2/atn42-65.htm](http://awips.wrh.noaa.gov/build4_2/atn42-65.htm).