30 April 2009 2009 Regional Smoke Management Meeting 1

For over a decade scientists and forecasters in the Pacific Northwest have been favored by access to high resolution meteorological models. Many in the community have become quite dependent on the products available from the twice-daily realtime forecast runs made at the University of Washington.

In April 2008, after a year of testing, the Consortium that oversees the operation of the forecast system decided to switch from the older MM5 model to the newer, more modular, Weather Research and Forecasting (WRF) model. As you will learn during this presentation, that change was not without its share of glitches.

I will give a brief overview of the two models, highlighting where they differ significantly and show how similar they are. I will conclude with some good news about the future (after April 2009).

Moving From MM5 to WRF

- Overview of MM5
- Overview of WRF
- Differences between MM5 and WRF
- Performance comparison

The MM5 model has its roots in model development at Penn State and subsequently at the National Center for Atmospheric Research (NCAR). As an early community model, its development was somewhat haphazard and the code had pitfalls that could trap the unwary contributor. However, it demonstrated adequate skill, especially in areas where topography plays a significant role in weather patterns.

Because the operations at the UW are set up to provide a timely forecast, the computer power exists to run multiple models during the day to test different models or individual modules within a model for their performance and skill. After a year of comparing the WRF model with the MM5, confidence increased that the WRF model would provide forecasts with equal or better accuracy of the MM5 model.

Moving From MM5 to WRF

- MM5 Modeling System Overview
- Minimum Requirements for Running the Modeling System Software

In this presentation think of yourself as an investor in a large piece of machinery, perhaps a ship, that you wish to use but really don't need to understand all of the intricacies of its operation. However, it helps to obtain some understanding so that you may judge whether the new ship will perform better.

MM5 Modeling System Overview

- Limited-area
- Terrain-following
- MM5 is a regional model,
  - it requires
    - an initial condition as well as
    - lateral boundary condition

The MM5 model forecasts weather over a limited area, although that area could extend from the Aleutians to El Paso. As with all gridded models it can resolve topography only down to a scale that is related to the grid spacing. In the case of MM5 that means that the minimum distance across a terrain feature, such as the distance from one ridge across a valley to the ridge on the other side, must be at least five times the grid spacing before the model will recognize that the valley exists. Although a fellow named Nyquist showed that the minimum distance required at least twice the spacing, practical implementations involve heavy doses of filtering and smoothing to remove noise that accumulates during the model run. Hence the much larger minimum distance.

The fact that the MM5 model uses a terrain following vertical coordinate system is not so important by itself for the user but will become more apparent as we look more closely at the WRF model. However, as a regional model, MM5 cannot begin to forecast weather without some initial conditions much the way that a game of checkers or chess cannot proceed without the pieces initially placed on the board.

If you restrict your attention to just a small portion of the board during a chess match, you can quickly understand how important boundary conditions are. Without being able to move a piece into the squares in the area you would not be able to follow the play. Similarly, a mesoscale model must have changeable boundaries to reflect the weather systems moving into its domain.

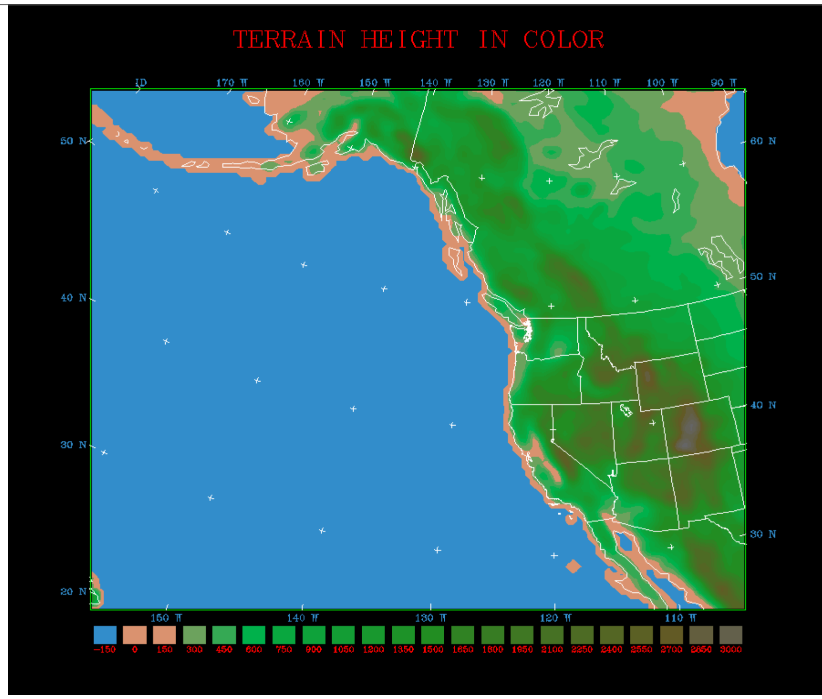

30 April 2009 2009 Regional Smoke Management Meeting 5

The MM5 model at the UW forecasts weather over this domain with a 36 km grid spacing. Within this domain there is a grid with 12 km spacing which covers part of the Eastern Pacific and most of the western US and southwestern Canada. Further, within the 12 km grid is a smaller grid with 4 km spacing covering all of Washington, Oregon, and Idaho and adjacent areas of California, Nevada, Utah, Wyoming, Montana, Alberta, and British Columbia.

Terrain Following Coordinates

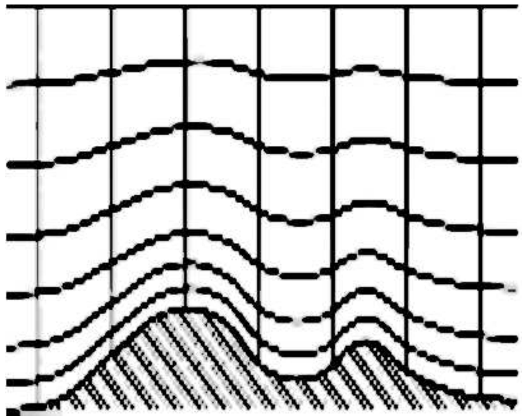

30 April 2009 2009 Regional Smoke Management Meeting 6

So that you can explain terrain following coordinates, I show a view of the vertical coordinate system as it crosses a couple of ridges.

Minimum Requirements for Running MM5

- Unix workstation with 128 Mb memory and 1-2 GB disk
- Fortran compiler
- Able to acquire gridded meteorological dataset, and observations

Ten years ago running the MM5 model was pretty much restricted to super computers but today's desktop computers have increased in capability so that you could likely run simulations at work or home. It's not likely that you could produce timely forecasts at small grid spacings (the execution times would be too long), but you could certainly make good wind and temperature fields adequate for further analysis and study.

Probably the biggest impediment to running MM5 is acquiring the necessary gridded meteorological data to initialize your forecast and the time varying boundary conditions needed to keep your forecast synchronized with the actual weather.

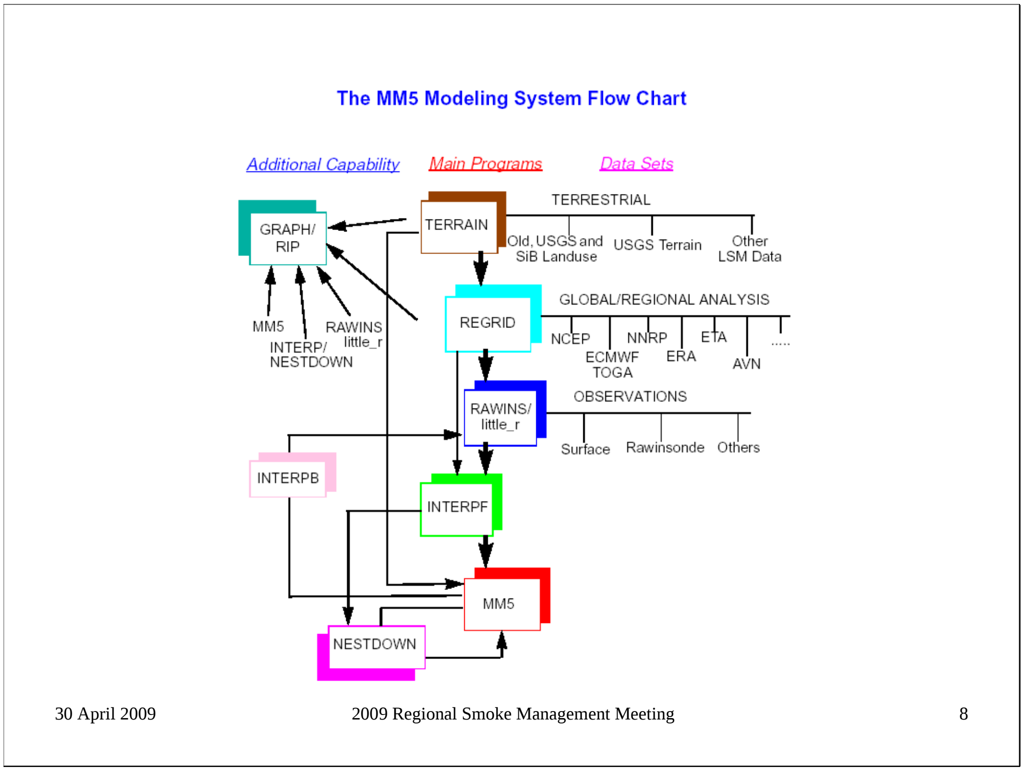

Let's open up the door to the machine room to see the internals of MM5. Hmm, there are a lot of parts down there. You can see that you need terrain, output from global models, perhaps some observations. It's likely that the observations and the global model won't lie on your grid points and so you need a way to interpolate the input to your grid before running your scenario. There are also numerous programs to help with visualization of the input and output.

The User Should Have

- Experience with numerical modeling of the atmosphere.
- Understanding of atmospheric science at the MS level.
- Basic UNIX/Linux knowledge.
- Basic FORTRAN 77 and 90 knowledge.
- <u>No previous experience in the MM5 model is required.</u>

Learn MM5 modeling system program from the on line tutorial

Except for the first two bullets, there are few specialized requirements to run MM5 that an advanced high school student couldn't have. The first two are required to avoid producing and believing an impossible (or improbable) forecast.

Overview of WRF

- Limited-area
- Terrain-following
- WRF is a regional model,
  - it requires
    - an initial condition as well as
    - lateral boundary condition

WRF is also a limited area model with a terrain-following vertical coordinate system. As a regional model is requires both initial and boundary conditions and, as such, doesn't differ much from MM5.

Overview of WRF

- Improved variable staggering
- Improved accuracy of approximations
- Improved isolation of different physics choices
  - Microphysics, cumulus parameterization, planetary boundary layer, surface layer, land-surface, longwave radiation, shortwave radiation, sub-grid turbulence

Now let's look at WRF. The WRF code is pretty much all new. It incorporates numerous improvements which translate to less numerical noise (less smoothing of the output), more accurate formulation of the equations, and a more modular approach which will allow different schemes to be used with less concern about incompatible combinations.

The Arakawa staggered grids

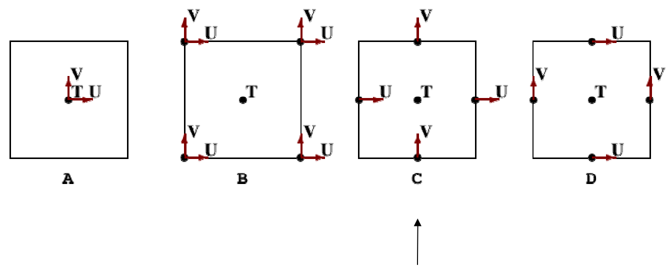

WRF

Forty years ago Akio Arakawa published his analysis of the stability of various ways to store the important variables used to predict weather. The Arakawa A grid is not particularly suited for long term runs. MM5 uses the B grid which can produce stable solutions but when combined with the truncation error of second order differencing schemes requires heavy smoothing to filter out the noise generated during long term runs.

WRF uses the Arakawa C grid and has much higher order differencing schemes to produce a solution that requires much less smoothing. Accordingly WRF has a much higher resolution for a given grid spacing than MM5. The coding is also much better at running on multiple CPUs and therefore the WRF runs actually complete is less time than MM5.

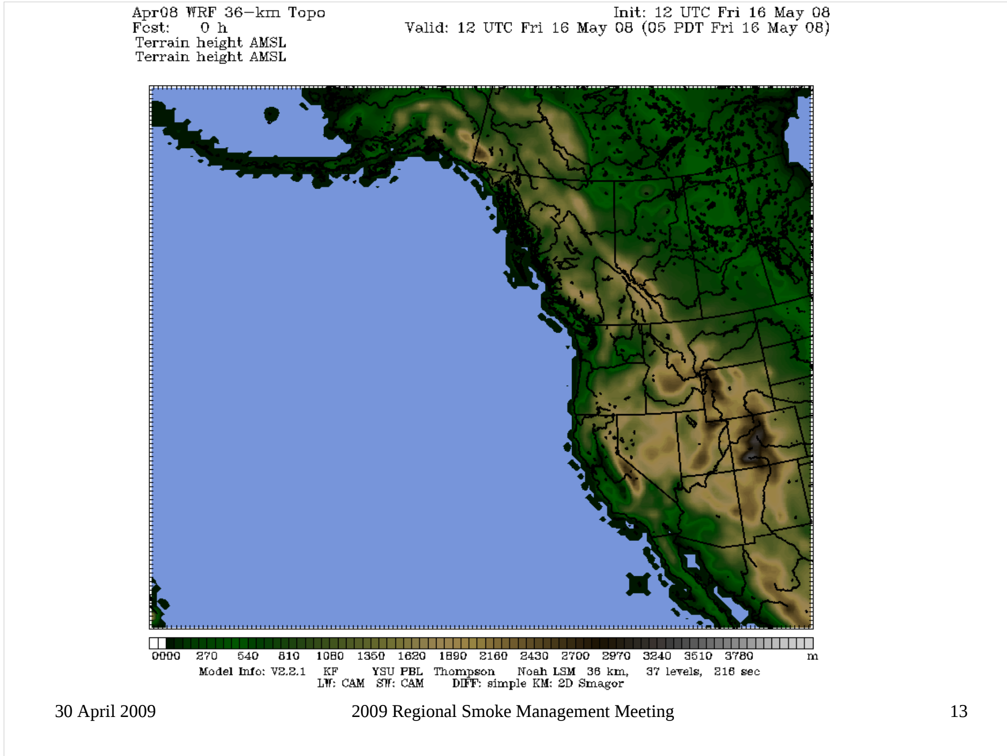

Except that the map is a bit fancier and covers slightly more area, the outer WRF domain is nearly the same as the MM5 we saw earlier. WRF is also run with nearly the same12 km and 4 km nested domains used in MM5.

Overview of WRF

- Platforms it runs on:
  - IBM SP systems (e.g. NCAR blackvista/blueice, Power5-based systems)
  - IBM Blue Gene
  - SGI: Origin 2000 and Altix
  - Pentium 4 cluster iJet system at NOAA FSL
  - IA64/Linux MPP (SGI Altix)
  - IA64/Linux MPP (HP Superdome at PNNL)
  - IA64/Linux MPP (NCSA)
  - Sun (single and SMP)
  - Cray: Unicos/X1, X1e (vector), and XD1 (Opteron) series
  - HP-UX
  - NEC: SX/8
  - Fujitsu: VPP 5000
  - Linux: PGI, Intel ifort, Pathscale, gfortran and g95 compilers
  - Pentium 3/4 clusters and SMPs
  - Intel Xeon IA32
  - IA64/Linux SMP (local)
  - AMD Opteron
  - Mac Intel/PPC, PGI/ifort/g95

WRF runs on a wide range of computers—I've grouped the super-computers on the left side in fine print and emphasized the smaller computers that might be found in a typical office or home in the right column. WRF runs on some pretty average computers with large memories and disk drives.

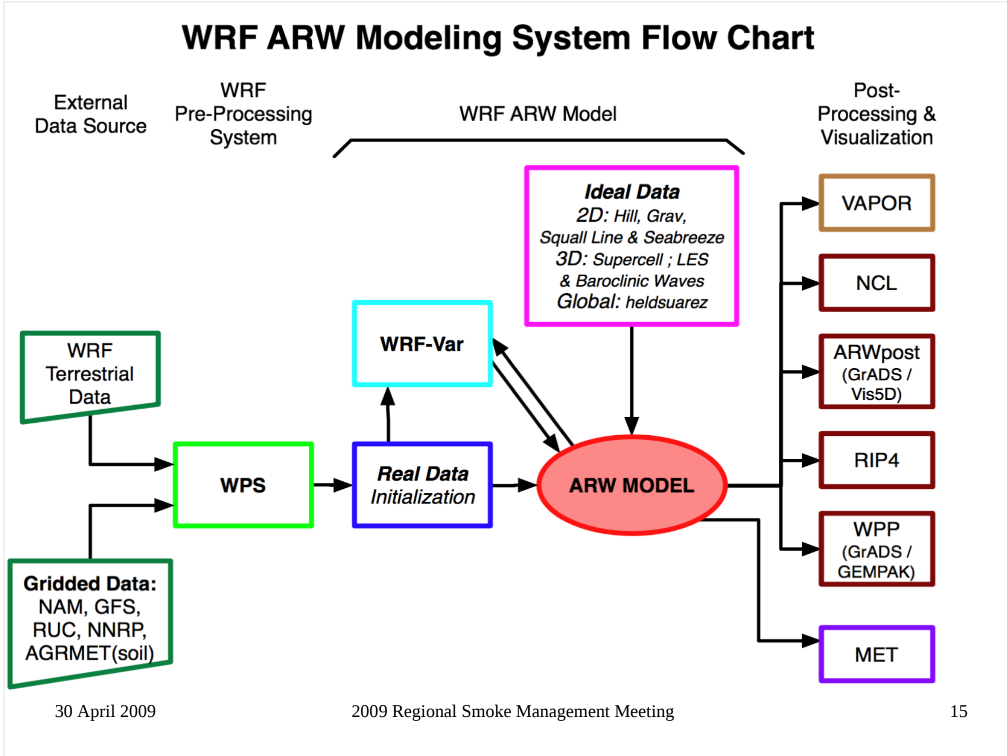

The machine room in WRF has many of the same components found in MM5. One change is the addition of a built-in capability to run idealized scenarios—most useful for testing new modules.

- Differences between MM5 and WRF
- WRF computer code is almost all new
  - Easier to modify, add new features
- Less spatial smoothing
- New bugs (see first bullet)

Most of the code in the WRF model is new and written to be more flexible and allow new modules to be added without the strong interaction encountered in MM5. The higher order schemes produce less noise which decreases the amount of filtering and smoothing required and allows smaller scale features to affect the solution for a given grid spacing.

However, because most of the code is new, there are new bugs that have been introduced some of which don't show up in the limited testing done before model release.

Comparison of Spatial Resolution

MM5

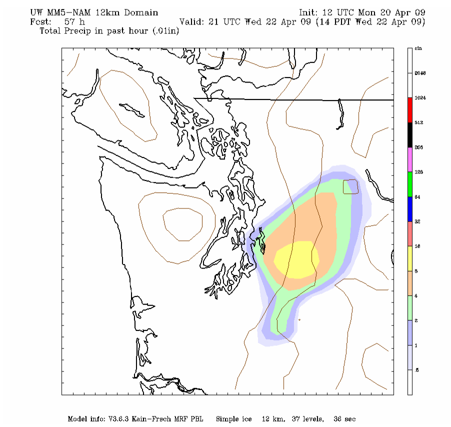

WRF

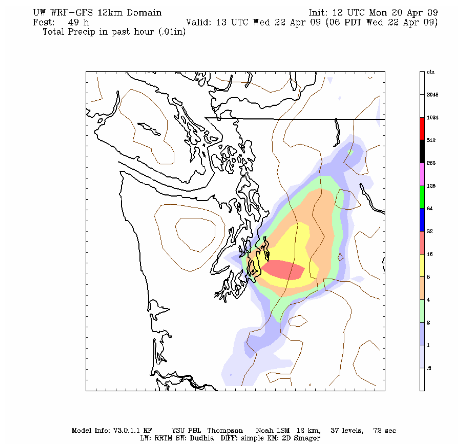

Here is a good example of the effects of the higher resolution. I compare an MM5 run with a WRF run. Because the two models were initialized by different global models, the time that a small system arrived in Puget Sound differed by eight hours between the two models. However the two panels show the system at about the same point in its passage.

The heavily smoothing in the MM5 run is apparent in the nearly featureless precipitation pattern on the left. The WRF model, with less smoothing and higher order approximations, produces higher precipitation rates with an apparent convergence zone just south of downtown Seattle. The periphery of the pattern also has much more structure reflecting the topographic influence that has been smoothed out in the MM5 run.

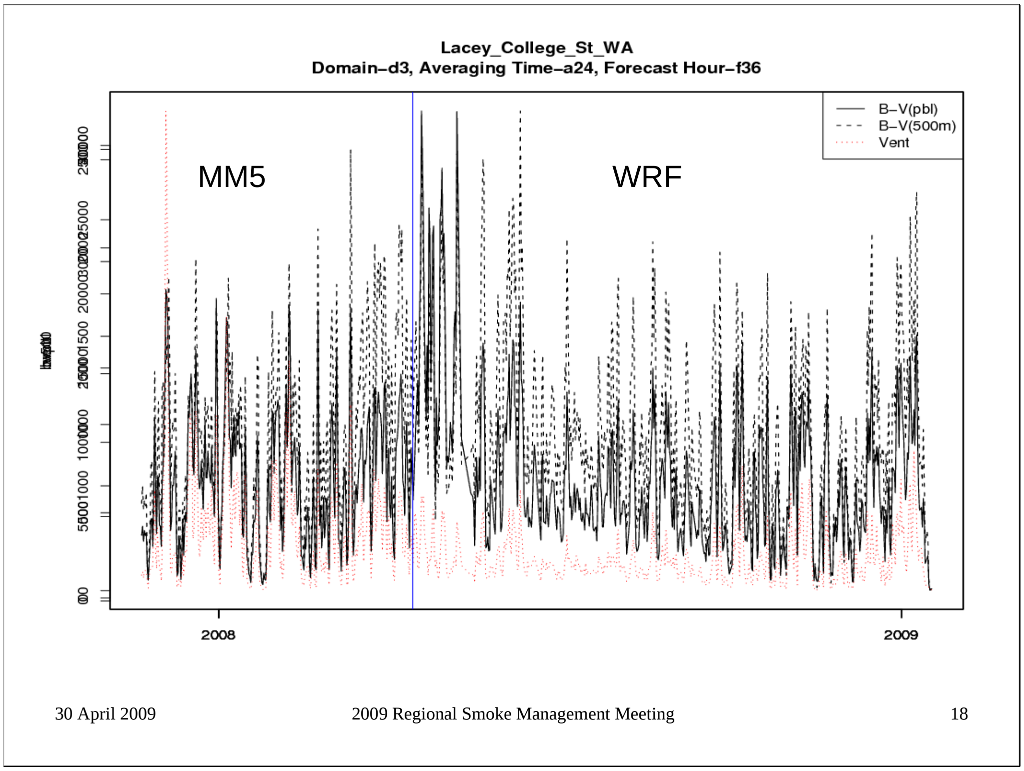

This figure shows the results of a bug in the WRF code and then its correction a few months later. The figure begins in late 2007 (on the left) with output from MM5. There are three variables being plotted in this time series—a classical ventilation index (the product of wind speed and mixing depth) in red dots, and two variations of a Brunt-Vaisala index which is based on the vertical temperature gradient plotted in black. I've intentionally allowed the ordinate axis to be overwritten to indicate that the three indices are self-scaled (they have differing values but here I'm emphasizing the difference in relative behavior).

During the five months of MM5 runs, the indices all have roughly the same behavior with the classical ventilation index consistently showing up just below the two B-V based indices. On 15 April 2008, marked by the vertical blue line, we switched to the WRF model and the ventilation index seemed to separate itself from the B-V based indices for about five months before the error in the WRF code was found and fixed. Afterwards, the relative behavior of the three ventilation indices returned to about the same as seen from the MM5 runs.

Model Performance

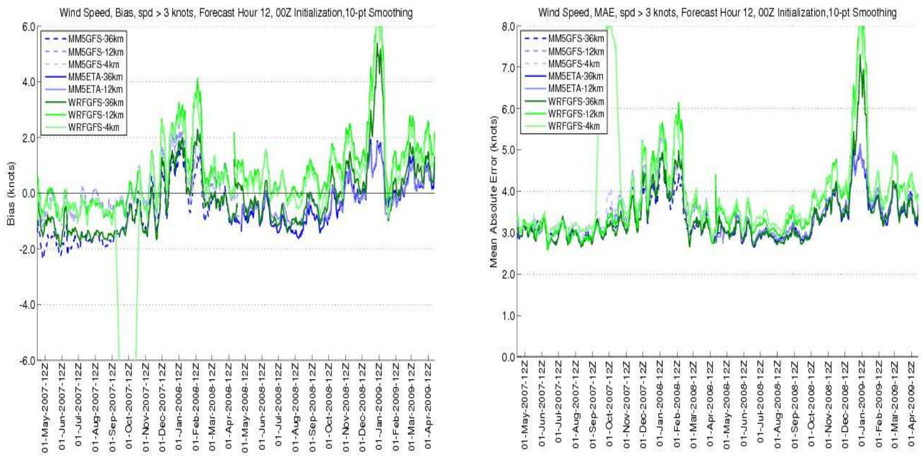

Here we see model performance of parallel runs of MM5 (blue lines) and WRF (green lines) for wind speed. Basically, there is little difference between WRF and MM5. Because WRF scales across processors better than MM5 (adding more processors decreases run time more with WRF than with MM5), the Consortium elected in April 2008 to change the operational runs to WRF.

Example of PBL Problems

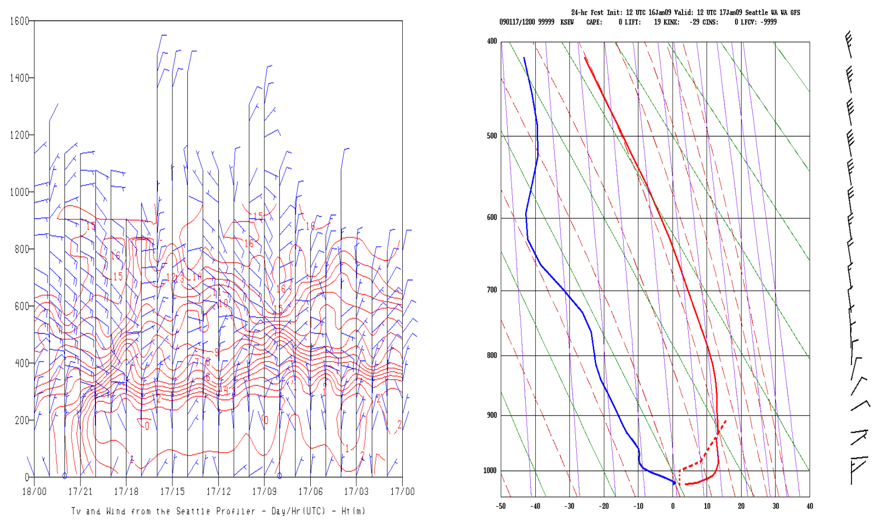

30 April 2009 2009 Regional Smoke Management Meeting 20

One of the problems first identified in MM5 runs continues after the switch to WRF—poor modeling of the lowest layers of the atmosphere. Both models, regardless of the module used to model the boundary layer (that layer of the atmosphere beginning at the earth surface and extending upward to a few hundred meters to perhaps a couple of thousand meters where the effects of the surface heating and friction influence the air motion) , have too much mixing. This increased mixing spreads out vertical temperature gradients and produces too high surface wind speeds. Both effects increase the mixing of emissions near the ground and result in low concentrations when used in an air quality model.

The left --hand panel in the figure above shows the temperature and wind speed and direction from the profiler at Sand Point on Lake Washington. At 12Z on the 17th the profiler shows a 200 m deep layer with a nearly isothermal temperature profile capped by a strong subsidence inversion. The right-hand panel shows the predicted sounding for the same time. Note that the predicted profile has the inversion beginning at the surface. I've approximated the profiler temperature by a dashed red line to show the difference between the predicted and observed temperatures.

Summary

- Same products available
- Better effective resolution
- Available sooner
- Similar or better performance
- PBL problems persist
  - Late breaking news—WRF 3.1 released, has new PBL schemes, being tested at UW.

In summary, the change from MM5 to WRF made no difference in the products available. The WRF output has less smoothing and shows increased structure and is available sooner.

The performance has not changed significantly and the pbl modeling suffers equally. The good news is that the latest release of WRF (version 3.1) has a new PBL scheme which testing at NCAR and other locations has shown is superior to any in the previous version. That release and the new PBL scheme, are being tested at the UW.