Nesting in WRF 2.0

David Gill
John Michalakes
Jimy Dudhia
William Skamarock
National Center for Atmospheric Research[1], Boulder, Colorado

S.G. Gopalakrishnan
SAIC and NOAA/NWS/NCEP, Camp Spring, Maryland

1. Introduction

A horizontal mesh refinement capability is available to users with the release of WRF version 2. The capability, referred to as *nesting,* supports 1-way and 2-way interaction between a lower-resolution parent domain and one or more higher-resolution nests to arbitrary (i.e. telescoping) nesting levels. The nesting capability runs efficiently on single-processor, shared-memory parallel, and distributed-memory systems. Development is continuing towards allowing overlapping and moving nests. Effort is ongoing to provide additional smoothing and feedback options, and alternative input files.

In addition to changes in the WRF model itself (relative to version 1.3) users will notice nesting related improvements to the Standard Initialization (SI) program. The SI generates fine grid domain information for ingest onto nested WRF domains through the WRF IO API. There is also a new program, called ndown, used for 1-way, non-interactive nesting.

In the presentation, examples of real-data cases (both 1-way and 2-way forcing) and ideal simulations are shown. Preliminary performance results are also presented.

2. User-Level Specifications

The WRF system supports both 1-way and 2-way nesting. The 2-way interactive nesting refers to the forcing that the parent and child domains provide to the each other. The coarse grid provides time-dependent lateral boundary conditions to the fine grid, and the child domain feeds back the higher resolution forecast data to the parent in the area of the nest. The 2-way nesting option in WRF initializes a new child domain by interpolating from the parent domain; in addition, a fine-grid input file may be used to initialize the child domain (specifically the orography, land/sea mask, and various masked surface fields). By selecting the "no feedback" option in the namelist file, a 2-way nested forecast essentially becomes a 1-way simulation. More typically, however, 1-way nesting involves use of the ndown program to generate the fine grid initial and boundary conditions from a previous coarse-grid forecast. This form of 1-way nesting may be used if a lower temporal frequency of coarse-to-fine boundary forcing is acceptable/preferred or if nested forecasts are to be run as separate WRF jobs.

Most settings for user control of nesting are run-time options set in the namelist.input file. The namelists are set up with two kinds of entries: those with a single value, and list-valued entries with one value per domain. The listed values are ordered with the first entry corresponding to the first domain, the second entry associated with the second domain, and so on. Domain dimensions, grid-distance and time step ratio, nesting hierarchical structure ("who is my mother"), the number of active domains, and when domains start/stop are all controlled by various namelist options. The default limit is 11 total domains, which can be changed at compile time. Otherwise, available machine memory is the only limit to the number of domains.

Users and developers interested in changing or extending nesting in WRF have control over nesting functionality through the WRF Registry, an editable text file from which much of the nesting code is automatically generated at compile-time. The Registry controls the following compile-time options for nesting:

[1] The National Center for Atmospheric Research is supported by the National Science Foundation.

- variables to horizontally interpolate
- horizontal interpolation method
- variables to feedback
- feedback method
- variables ingested with ndown
- variables to smooth
- variables requiring lateral boundary forcing
- I/O

Compile-time changes to the Registry allow the user to modify nesting without introducing source code modifications. However, most users will simply modify the WRF namelist file to generate a multi-domain forecast and leave the Registry file alone.

Prior to running the WRF model for real-data cases, the Standard Initialization (SI) may be run for a nested forecast to generate initial data for the nested domains. The SI is a system of Fortran code and perl scripts with a graphical user interface (GUI) that generates WRF input from external GriB data. The release of the SI software package (tightly coordinated with the WRF version 2 release) provides coarse and fine grid data suitable for later ingest by the WRF system. The SI produces three types of input files for WRF: 1) multiple coarse-grid files that each contain a single time period of meteorologically complete data, 2) a time-independent, static-data only file for each domain requested, and 3) the initial time of each fine grid with a complete meteorological data stream.

3. Software Technical Details

As with the single-domain version of WRF released in March 2003, much of the code for mesh refinement is automatically generated in the form of Fortran include files. The Registry is used as an input file and interpreted for use in all phases of nesting: I/O, horizontal interpolation of the fine grid, lateral boundary forcing from the coarse grid to the fine grid, feedback to the coarse grid, and smoothing on the coarse grid in the area of the nest. Other automatically generated files handle more infrastructure related tasks, such as movement of data between domains.

At the software infrastructure level, nesting is based on built-in capabilities for defining, allocating, initializing, integrating over, and exchanging forcing data between multiple nested domains in the WRF Advanced Software Framework (ASF) (Michalakes *et al*, 2002, Michalakes *et al*, 1998) and in underlying libraries such as RSL (Michalakes, 2000) that are distributed with WRF. Currently, in WRF 2.0, RSL must be compiled and linked into the model whenever nesting is required. Compile-time settings for WRF are provided that allow WRF and RSL to be built without MPI (a parallel message passing library). This avoids the need to install MPI on machines that will only ever run WRF as a single processor or shared-memory parallel program.

The hierarchy of nested domains in WRF is represented as a tree, with the top-level domain as the root. Integrating forward in time involves a depth-first traversal of the tree, using a recursive subroutine, *integrate(),* defined in frame/-module_integrate.F. Calling integrate for a domain, represented by a node in the tree, advances the domain and all of its children forward a specified time interval. Thus, at the highest level, a WRF run amounts to a single call to integrate(), passing the root domain and the simulation run-length as arguments (see main/wrf.F).

When a fine-grid domain is instantiated within WRF, the nest becomes a child domain of one and only one parent domain. The nest may itself be the parent for domains on a finer level. After the allocation process is complete then the state arrays that make up the domain are initialized. First, the coarse grid data is horizontally interpolated to the fine grid, where each η-surface for each 3D field is processed independently. Then the single time of refined-mesh data supplied by the SI is ingested. Any fields requested in the Registry that exist in the fine grid input file overwrite the horizontally interpolated data. The terrain elevation, reference surface pressure, and reference geopotential are linearly blended along the lateral boundary of the coarse and fine interface. This provides consistent coarse-domain information along the fine-grid boundaries. The fine grid is fed back to the coarse grid, and both the parent and child domains have their reference fields recomputed.

The primary horizontal interpolation used for most of the meteorological variables is a 3rd order, positive definite advection scheme (Smolarkiewicz and Grell, 1992). There are nearest-neighbor interpolators for fields that are categorical (e.g. land use and soil index) and special 4-point interpolators for fields that are stenciled based upon the domain's land/sea mask (e.g. soil temperature, soil moisture, sea ice, sea surface temperature).

The fine-grid forcing back to the parent domain, known as *feedback*, is computed as the mean of the entire parent-cell for mass points and a cell-face average for the horizontal momentum fields. For some masked variables, coincident points in the child and parent domain employ a single-point feedback technique (not suitable for even ratio refinement, yet). All feedback computations are separately performed on each η surface, as the parent and child computational surfaces always coincide.

For fields on the parent domain that have been updated via a feedback, an η-surface smoothing option in the area of the updated values may be used to filter the data. Both a 1-2-1 smoother and a smoother-desmoother data filter are available. All data fields selected for smoothing are processed with the same 2D smoother.

4. Summary

An efficient, parallel, full-function implementation of horizontal mesh refinement is available in WRF version 2. This nesting feature may be used for both real-data and ideal initialization cases. The horizontally-nested domains may be instantiated and stopped at any time during a forecast. A child domain exists within a single parent domain and may be recursively telescoped to any depth. Other than the outermost domain, which has no parent, these domains may have any number of children contained with the same parent domain and the children may have differing refinement ratios. The time-step ratio for a fine-grid domain need not be identical to the grid-distance refinement ratio.

In the short-term, there are plans to support even-ratio refinements, and to permit real-data initializations using only terrestrial/static fields. Long-term work includes a moving nest capability, overlapping domains, support of nesting under RSL_LITE for very large simulations, and documentation.

The possibility of extending the nesting capability to the E-grid-based WRF-NMM core is being explored for the hurricane problem. A bi-linear interpolation technique has been developed and tested within the WRF framework to horizontally interpolate most of the meteorological variables from the coarse domain on to the fine domain. The nearest-neighbor interpolation method may be adopted for most of the static fields. Development is continuing in the direction of lateral boundary forcing on the fine grid and feeding back to the coarse grid, specifically for communication between the domains. For future applications, there are proposals to use a 3D variational analysis to provide appropriate initial conditions.

References

Michalakes, J.G., M. McAtee, J. Wegiel (2002): Software Infrastructure for the Weather Research and Forecast Model, in proceedings of UGC 2002, June, Austin, Texas, 13 pp. http://www.hpcmo.hpc.mil/Htdocs/UGC/UGC02/paper/john_michalakes_paper.pdf

Michalakes, J.: RSL: A Parallel Runtime System Library for Regional Atmospheric Models with Nesting, in Structured Adaptive Mesh Refinement (SAMR) Grid Methods, IMA Volumes in Mathematics and its Applications (117), Springer, New York, 2000, pp. 59-74.

Michalakes, J., J. Dudhia, D. Gill, J. Klemp and W. Skamarock: Design of a next-generation regional weather research and forecast model: Towards Teracomputing, World Scientific, River Edge, New Jersey, 1998, pp. 117-124.

Smolarkiewicz, P. K., and G. A. Grell 1992: A class of monotone interpolation schemes. *J. Comp. Phys.*, **101**, 431-440.