

## DEVELOPMENT OF A NEXT-GENERATION REGIONAL WEATHER RESEARCH AND FORECAST MODEL

J. MICHALAKES,<sup>2</sup> S. CHEN,<sup>1</sup> J. DUDHIA,<sup>1</sup> L. HART,<sup>3</sup>  
J. KLEMP,<sup>1</sup> J. MIDDLECOFF,<sup>3</sup> W. SKAMAROCK<sup>1</sup>

<sup>1</sup>*Mesoscale and Microscale Meteorology Division, National Center for Atmospheric  
Research, Boulder, Colorado 80307 U.S.A.*

<sup>2</sup>*Mathematics and Computer Science Division, Argonne National Laboratory, Argonne,  
Illinois 60439 U.S.A, michalak@mcs.anl.gov, +1 303 497-8199*

*NOAA Forecast Systems Laboratory, Boulder, Colorado 80303 U.S.A*

The Weather Research and Forecast (WRF) project is a multi-institutional effort to develop an advanced mesoscale forecast and data assimilation system that is accurate, efficient, and scalable across a range of scales and over a host of computer platforms. The first release, WRF 1.0, was November 30, 2000, with operational deployment targeted for the 2004-05 time frame. This paper provides an overview of the project and current status of the WRF development effort in the areas of numerics and physics, software and data architecture, and single-source parallelism and performance portability.

### 1 Introduction

The WRF project is developing a next-generation mesoscale forecast model and assimilation system that will advance both the understanding and the prediction of mesoscale precipitation systems and will promote closer ties between the research and operational forecasting communities. The model will incorporate advanced numerics and data assimilation techniques, a multiple relocatable nesting capability, and improved physics, particularly for treatment of convection and mesoscale precipitation. It is intended for a wide range of applications, from idealized research to operational forecasting, with priority emphasis on horizontal grids of 1–10 kilometers. A prototype

---

<sup>1</sup> This work was supported under National Science Foundation Cooperative Agreement ATM-9732665.

<sup>2</sup> This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-Eng-38.

has been released and is being supported as a community model.<sup>1</sup> Based on its merits, it will be a candidate to replace existing forecast models such as the PSU/NCAR Mesoscale Model (MM5), the ETA model at NCEP, and the RUC system at FSL. The first release of the model, WRF 1.0, was November 30, 2000. This paper reports on progress since our first ECMWF workshop paper on the WRF design two years ago [3]. Section 2 provides an overview of the WRF project and model, Section 3 the software design and implementation, and Section 4 preliminary performance results for the WRF 1.0 code.

## 2 WRF Overview

A large number of organizations are participating in the WRF project. The principal organizations are the Mesoscale and Microscale Meteorology Division of the National Center for Atmospheric Research (NCAR/MMM), the National Centers for Environmental Prediction (NOAA/NCEP), the Forecast Systems Laboratory (NOAA/FSL), the University of Oklahoma Center for the Analysis and Prediction of Storms (CAPS), and the U.S. Air Force Weather Agency (AFWA). Additional participants include the Geophysical Fluid Dynamics Laboratory (NOAA/GFDL), the National Severe Storms Laboratory (NOAA/NSSL), the Atmospheric Sciences Division of the NASA Goddard Space Flight Center, the U.S. Naval Research Laboratory Marine Meteorology Division, the U.S. Environmental Protection Agency Atmospheric Modeling Division, and a larger number of university

Conservative variables:		$U = \rho u, \quad V = \rho v, \quad W = \rho w, \quad \Theta = \rho \theta$
Inviscid, 2-D equations in Cartesian coordinates	$\frac{\partial U}{\partial t} + \gamma R \pi \frac{\partial \Theta}{\partial x} - fV = -\frac{\partial Uu}{\partial x} - \frac{\partial Wu}{\partial z}$	
	$\frac{\partial W}{\partial t} + \gamma R \pi \frac{\partial \Theta}{\partial z} + g\rho = -\frac{\partial Uw}{\partial x} - \frac{\partial Ww}{\partial z}$	
	$\frac{\partial \Theta}{\partial t} + \frac{\partial U\theta}{\partial x} + \frac{\partial W\theta}{\partial z} = \rho Q$	
	$\frac{\partial \rho}{\partial t} + \frac{\partial U}{\partial x} + \frac{\partial W}{\partial z} = 0$	
Pressure terms directly related to $\Theta$ :	$\gamma R \pi \nabla \Theta = c_p \Theta \nabla \pi = \nabla p$	

Figure 1 Flux form equations in height coordinates.

<sup>1</sup> <http://www.wrf-model.org>

Hydrostatic pressure coordinate:	$\eta = (\pi - \pi_t) / \mu, \quad \mu = \pi_s - \pi_t$
Conservative variables:	$U = \mu u, \quad W = \mu w, \quad \Theta = \mu \theta, \quad \Omega = \mu \dot{\eta}$
Inviscid, 2-D equations without rotation:	$\frac{\partial U}{\partial t} + \mu \alpha \frac{\partial p}{\partial x} + \frac{\partial p}{\partial \eta} \frac{\partial \phi}{\partial x} = - \frac{\partial U u}{\partial x} - \frac{\partial \Omega u}{\partial \eta}$ $\frac{\partial W}{\partial t} + g \left( \mu - \frac{\partial p}{\partial \eta} \right) = - \frac{\partial U w}{\partial x} - \frac{\partial \Omega w}{\partial \eta}$ $\frac{\partial \Theta}{\partial t} + \frac{\partial U \theta}{\partial x} + \frac{\partial \Omega \theta}{\partial \eta} = \mu Q$ $\frac{\partial \mu}{\partial t} + \frac{\partial U}{\partial x} + \frac{\partial \Omega}{\partial \eta} = 0$ $\frac{d\phi}{dt} = gw, \quad \frac{\partial \phi}{\partial \eta} = -\mu \alpha, \quad p = \left( \frac{R\theta}{p_0 \alpha} \right)^{\gamma}$
Figure 2 Flux form equations in mass coordinates.	

researchers. The project is organized under a WRF Oversight Board, who appoint and work with a WRF Coordinator (Klemp) and a WRF Science Board. Five WRF development teams -- Numerics and Software, Data Assimilation, Analysis and Validation, Community Involvement, and Operational Implementation -- are further divided into a number of working groups, which include Dynamic Model Numerics; Software Architecture, Standards, and Implementation; Analysis and Visualization; and Data Handling and Archiving.

The WRF development timeline is in two main phases: full implementation as a research model, to be completed by the end of 2002 (with the exception of 4DVAR, which is planned for 2003), and then full implementation as an operational forecasting system at NCEP and AFWA, to be largely completed by the end of 2004 but with additional effort for 4DVAR implementation and diagnosis and operational performance refinement stretching into the 2005-08 time frame.

The WRF design allows for multiple dynamical cores. The dynamical core in this first release uses Eulerian finite-differencing to integrate the fully compressible nonhydrostatic equations in height coordinates (Figure 1) in scalar-conserving flux form using a time-split small step for acoustic modes. A mass-coordinate prototype is also being implemented (Figure 2). Large time steps utilize a third-order Runge-Kutta technique, and second- to sixth-order advection operators can be specified. The horizontal staggering is an Arakawa-C grid. Future releases may include an implicit option to time splitting for the Eulerian dynamical cores. A semi-implicit semi-Lagrangian prototype is also under development at NCEP [2,5]; this will use an Arakawa-A grid and allow WRF and the new global model under development at NCEP to operate as unified models.

WRF physics is intended to be plug compatible [1] to allow easy exchange of physics packages with other models. In actual fact, the physics packages that are

included in this first release of WRF are from other modeling systems, adapted to operate within the WRF software framework. At least one representative package for each of the physical processes necessary for real-data mesoscale forecasts are included:

- Longwave radiation: RRTM
- Shortwave radiation: NASA/GSFC, MM5 (Dudhia)
- Cumulus: Kain-Fritsch, Betts-Miller-Janjic
- Explicit microphysics: Kessler, Lin et al, NCEP 3-class (Hong)
- PBL: MRF, MM5 (Slab)

Work is continuing to add additional parameterizations -- notably in the area of land surface processes, for which a new working group is being added to the WRF development effort. Parameterizations and packages of concern for atmospheric chemistry and regional climate are also being considered.

The WRF 1.0 distribution includes idealized cases for testing: two-dimensional hill case, baroclinic wave, squall line, and supercell thunderstorm. Model initialization for real data is handled by the WRF Standard Initialization (SI) package, developed largely at NOAA/FSL. The SI package is also distributed with the model. WRF is package-independent by design; datasets are currently stored in NetCDF format, and other formats such as GriB and HDF will be supported. Effort is under way within the WRF collaboration to design and implement a three-dimensional variational assimilation (3DVAR) system to initialize WRF in the 2001-02 timeframe; this will be followed by a full four-dimensional variational (4DVAR) system in the 2004-05 time frame.

The current release of WRF 1.0 supports a single domain. The WRF model will support nesting, with 2-way interacting nests that are dynamically instantiable and that may move and overlap. Control of nesting will be via a priori scripting. Eventually, adaptive mesh refinement -- that is, nesting based on evolving conditions within a running simulation -- will be supported.

### **3 Software Architecture**

The WRF prototype employs a layered software architecture that promotes modularity, portability, and software reuse. Information hiding and abstraction are employed so that parallelism, data management, and other issues are dealt with at specific levels of the hierarchy and transparent to other layers. The WRF software architecture (Figure 3) consists of three distinct model layers: a solver layer that is usually written by scientists, a driver layer that is responsible for allocating and deallocating space and controlling the integration sequence and I/O, and a mediation layer that communicates between the driver and model. In this manner, the user code is isolated from the concerns of parallelism.

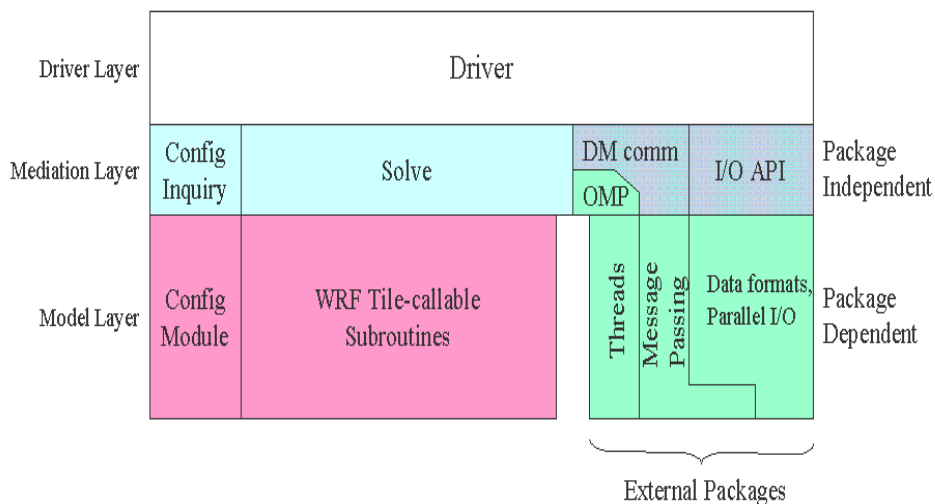


Figure 3 WRF software architecture schematic

The WRF prototype uses modern programming language features that have been standardized in Fortran90: modules, derived data-types, dynamic memory allocation, recursion, long variable names, and free format. Array syntax is avoided for performance reasons. A central object of a nested model is a *domain*, represented as an instance of a Fortran90 derived data type. The memory for fields within a domain are sized and allocated dynamically, facilitating run-time resizing of domain data structures to accommodate load balancing and dynamic nesting schemes.

External packages, such as MPI, OpenMP, single-sided message passing, and higher-level libraries for parallelization, data formatting, and I/O, will be employed as required by platform or application; however, these are considered external to the design, and their interfaces are being encapsulated within WRF-specific/package-independent Application Program Interfaces (API). For example, details of parallel I/O and data formatting are encapsulated within a WRF I/O API. A detailed specification of the WRF I/O API is distributed with the model so that implementers at particular sites can use an existing implementation or develop their own to accommodate site-specific requirements.

A flexible approach for parallelism is achieved through a two-level decomposition in which the model domain may be subdivided into patches that are assigned to distributed memory nodes, and then may be further subdivided into tiles that are allocated to shared-memory processors within a node. This approach addresses all current models for parallelism -- single-processor, shared-memory, distributed-memory, and hybrid -- and also provides adaptivity with respect to processor type: tiles may be

sized and shaped for cache-blocking or to preserve maximum vector length. Model layer subroutines are required to be *tile-callable*, that is, callable for an arbitrarily sized and shaped subdomain. All data must be passed through the argument list (state data) or defined locally within the subroutine. No COMMON or USE-associated state data is allowed. Domain, memory, and run dimensions for the subroutine are passed separately and unambiguously as the last eighteen integer arguments. Thus, the WRF software architecture and two-level decomposition strategy provides a flexible, modular approach to performance portability across a range of different platforms, as well as promoting software reuse. It will facilitate use of other framework packages at the WRF driver layer as well as the reverse, the integration of other models at the model layer within the WRF framework. A related project is under way at NCEP to adapt the nonhydrostatic Eta model to this framework as a proof of concept and to gauge the effectiveness of the design.

As with any large code development project, software management in the WRF project is a concern. For source code revision control, the project is relying on the CVS package. In addition, a rudimentary computer-aided software engineering (CASE) tool called a Registry has been developed for use within the WRF project. The Registry comprises a database of tables pertaining to the state data structures in WRF and their associated attributes: type, dimensionality, number of time levels, staggering, interprocessor communication (for distributed memory parallelism), association with specific physics packages, and attributes pertaining to their inclusion within WRF initial, restart, and history data sets. The Registry, currently implemented as a text file and Perl scripts, is invoked at compile time to automatically generate interfaces between the driver and model layers in the software architecture, calls to allocate state fields within the derived data type representing WRF domains, communicators for various halo exchanges used in the different dynamical cores, and calls to the WRF I/O API for initial, restart, and history I/O. The Registry mechanism has proved invaluable in the WRF software development to date, by allowing developers to add or modify state information for the model by modifying a single file, the database of Registry tables.

#### **4 Preliminary Results**

Initial testing of the WRF prototype has focused on idealized cases such as simple flow over terrain, channel models, squall lines, and idealized supercell simulations. Subsequent testing has involved real data cases and comparison with output from existing models such as MM5. An automated real-time forecasting system similar to the MM5 real-time system at NCAR<sup>1</sup> is under construction and will provide a means for real-time testing and verification of the WRF system over extended periods.

---

<sup>1</sup> <http://rain.mmm.ucar.edu>




Figure 4 WRF performance compared with MM5

Benchmarking to evaluate and improve computational performance is also under way. The model has run in shared-memory, distributed-memory, and hybrid parallel modes on the IBM SP, Compaq ES40, Alpha and PC Linux clusters, and other systems. WRF is neutral with respect to packages and will use a number of different communication layers for message passing and multithreading. The current prototype uses MPI message passing through the RSL library [4] and OpenMP. Typically, straight distributed-memory parallelism (multiple patches with one tile per patch) has been the fastest, most scalable mode of operation to date, suggesting that additional effort may be required to realize the full potential from the shared-memory implementations of WRF. Using straight distributed-memory parallelism and benchmarking an idealized baroclinic wave case for which a floating-point operations count is known (measured by using the Perfex tool on the SGI Origin), WRF ran at 467 Mflop/s on 4 processors and 6,032 Mflop/s on 64 processors of the IBM SP Winterhawk-II system at NCAR.<sup>1</sup> This is approximately 81 percent efficient relative to 4 processors (6032/467/16). On an SGI Origin2000 with 400 MHz processors, performance was 2497 Mflop/s on 16 processors and 8,914 Mflop/s on 64 processors, or 89 percent efficient relative to 16 processors.

A comparison of WRF performance with performance of an existing model, MM5, is shown in Figure 4. This shows both MM5 and WRF running a standard MM5 benchmark case, the AFWA T3A (Europe) scenario<sup>2</sup>: 36 km resolution on a

<sup>1</sup> <http://www.scd.ucar.edu/main/computers.html>

<sup>2</sup> <http://www.mmm.ucar.edu/mm5/mpp/helpdesk>

136 by 112 by 33 grid. The WRF is using comparable to or more sophisticated physical options than the MM5. A comparable operations count for WRF on this scenario is not known, so instead time-to-solution is compared. In terms of time-per-timestep, the WRF simulation is considerably more costly, however, the 2 time-level Runge-Kutta solver in WRF allows for a considerably longer time step: 200 seconds versus 81 seconds for MM5. Thus, the time-to-solution performance for WRF is slightly better than MM5 and should improve with tuning and optimization.

## 5 Summary

With the release of WRF 1.0, an important milestone has been reached in the effort to develop an advanced mesoscale forecast and data-assimilation system designed to provide good performance over a range of diverse parallel computer architectures. Ongoing work involves design and implementation of nesting, expanding physics and dynamical options, development of a parallel 3DVAR system based on the same WRF software architecture, performance optimization, and testing and verification over a range of applications including research and operational forecasting.

## References

1. Kálnay E., M. Kanamitsu, J. Pfaendtner, J. Sela, M. Suarez, J. Stackpole, J. Tuccillo, L. Umscheid, and D. Williamson: "Rules for interchange of physical parameterization," *Bull. Amer. Meteor. Soc.* **70** (1989) 620-622.
2. Leslie, L. M., and R. J. Purser: "Three-dimensional mass-conserving semi-Lagrangian scheme employing forward trajectories," *Mon. Wea. Rev.*, **123** (1995) 2551-2566.
3. 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 (1999), pp. 117-124.
4. 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.
5. Purser, R. J., and L. M. Leslie: "An efficient semi-Lagrangian scheme using third-order semi-implicit time integration and forward trajectories." *Mon. Wea. Rev.* **122** (1994) 745-756.