

Parallelization of a Subgrid Orographic Precipitation Scheme in an MM5-Based Regional Climate Model

L. Ruby Leung¹, John G. Michalakes², and Xindi Bian¹

¹ Atmospheric Science and Global Change Resource, Pacific Northwest National Laboratory, Richland, WA 99352, USA

{ruby.leung, xindi.bian}@pnnl.gov

² Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA

michalak@mcs.anl.gov

Abstract. Regional Climate Models (RCMs) are practical downscaling tools to yield regional climate information for assessing the impacts of climate variability and change. The Pacific Northwest National Laboratory (PNNL) RCM, based on the Penn State/NCAR Mesoscale Model (MM5), features a novel subgrid treatment of orographic precipitation for coupling climate, hydrologic, and ecologic processes at the watershed scale. The parameterization aggregates subgrid variations of surface topography into a finite number of surface elevation bands. An airflow model and a thermodynamic model are used to parameterize the orographic uplift/descent as air parcels cross over mountain barriers or valleys. The parameterization has significant performance advantages over nesting to achieve comparable resolution of climate information; however, previous implementations of the subgrid scheme required significant modification to the host MM5 model, prohibiting its incorporation within the NCAR-supported community version of MM5. With this effort, software engineering challenges have been addressed to incorporate, parallelize, and load-balance the PNNL subgrid scheme with minimum changes to MM5. The result is an efficient, maintainable tool for regional climate simulation and a step forward in the development of an MM5-based community regional climate model.

1 Introduction

In areas with heterogeneous surface elevation and vegetation, high spatial resolution is required to accurately simulate precipitation and surface hydrology. Since computational cost increases approximately as the cube of resolution in atmospheric models, techniques such as nesting are often used to focus costly high-resolution computation where it is needed. However, the use of nesting to resolve topography in climate simulations has a number of disadvantages. First, topography may be highly spatially variable from cell to cell in a model domain so that even within the limited area of a high-resolution nest, computation is wasted in areas of the nested domain that do not require it. Second, major climate processes that require additional topographical refinement involve

only column physics (e.g., cloud and precipitation processes), not model dynamics; therefore, increased temporal resolution—that is, reducing the time step—is unnecessary. A subgrid parameterization of orographic precipitation [4,3,5] has been developed as an alternative to the use of nesting to efficiently perform long-term integration that yields high spatial resolution climate information, which is important for hydrological applications and climate impact assessments.

The subgrid parameterization was first implemented in the prototype version of the Penn State/NCAR Mesoscale Model (MM5) [1] model, producing the PNNL Regional Climate Model, a separate version. Subsequently, the community mesoscale model, now MM5 version 3, underwent additional development including nonhydrostatic dynamics, new options relating to land surface processes, and adaptation to distributed-memory scalable computing systems [6]. The demonstrated effectiveness of the same-source parallelization approach and the need for scalable performance in the PNNL RCM suggested the feasibility and appropriateness of integrating the PNNL-developed climate parameterizations into the NCAR-supported community version of the model. Ultimately, this effort is intended to form the basis for an MM5-based Community Regional Climate Model (CRCM). To date, we have developed a new parallel version of the PNNL RCM and subgrid scheme based on the current (at this writing) MM5 Version 3.4. Section 2 of this paper describes the host MM5 model. Section 3 describes subgrid scheme and details of parallelization and load balancing. Section 4 provides preliminary performance and load balancing results. Section 5 summarizes the issues addressed in this community effort.

2 Penn State/NCAR MM5

MM5 is a limited-area, nonhydrostatic, terrain-following sigma-coordinate model designed to simulate mesoscale and regional-scale atmospheric circulation. Features of the models include (i) a multiple-nest capability, (ii) nonhydrostatic dynamics, which allows the model to be used at a few-kilometer scale, (iii) multitasking capability on shared- and distributed-memory machines, (iv) a four-dimensional data-assimilation capability, and (v) numerous physics options. The latest versions of MM5 include several features that are important for climate applications such as regular updating of the lower boundary conditions (sea surface temperature and sea ice), two options for land surface modeling, and physics options such as radiative transfer that are more accurate for long-term integrations.

Parallelism in MM5 is implemented by using the “same-source” approach, described in [6]. This involves a traditional two-dimensional horizontal data domain decomposition, but with minimal changes to the original source, allowing the parallel code to be maintained as part of the official MM5 for use on a range of both shared- and distributed-memory parallel computers, including the IBM SP, Cray T3E, Fujitsu VPP, Compaq ES40, SGI Origin2000, PC and Alpha-based Beowulf clusters, and workstations. Interprocessor communication to update subdomain halos, exchange forcing and feedback data between nested

domains, and implement distributed I/O is supplied by the RSL library [7]. RSL also provides support for automatic domain decomposition with unequally sized subdomains for load balancing. The FLIC source translator automates changes to MM5 loops and indices for parallel computation [8]. More recently, FLIC has been extended to collapse MM5 physics loops over the two horizontal dimensions into a single loop for improved performance on vector machines, and this translation is identical to that required for the PNNL subgrid scheme. More details are provided in the following section.

3 Subgrid Orographic Precipitation Scheme

The subgrid parameterization of orographic precipitation aggregates subgrid variations of surface topography into a finite number of surface elevation bands. A dominant vegetation cover is defined for each elevation class of each model grid cell to account for the subgrid heterogeneity in vegetation and land-water contrast. An airflow model and a thermodynamic model are used to parameterize the orographic uplift/descent as air parcels cross over mountain barriers or valleys. Physical processes such as cloud microphysics, convection, turbulence transfer, radiative transfer, and land-atmosphere transfer are all calculated for each subgrid elevation class based on its surface elevation, vegetation cover, and atmospheric conditions. The result is separate predictions of precipitation, temperature, snow water equivalent, soil moisture, and surface runoff for a selected number of surface elevation classes within each grid cell.

Figure 1 shows a schematic of the subgrid parameterization as applied to a grid cell 50 km by 50 km in the western United States. During postprocessing, the simulated fields can be distributed according to the spatial distribution of surface elevation within each grid cell to yield predictions at the scale of the surface elevation data. Hence, the RCM can operate at a coarser spatial resolution (typically 50–100 km) while still accounting for subgrid spatial heterogeneity in surface topography and vegetation, but at a reduced computational cost. The subgrid method significantly improves the simulation of surface temperature, precipitation, and snowpack over mountainous areas (see, e.g., [2]).

3.1 Structure and Decomposition of Subgrid Variables

When running with the subgrid parameterization of orographic precipitation, MM5 computes two solutions of the physically forced prognostic equations, one for the grid cell mean variables, and one for the subgrid variables. A separate set of arrays stores the subgrid variables. These include prognostic variables (e.g., temperature, wind, moisture, the various forms of cloud water, and the pressure perturbation) and their tendencies, and diagnostic variables such as precipitation and ground temperature. Several new arrays are added to store information for mapping between the grid cells and subgrid elevation classes.

SOLVE is the MM5 subroutine that computes the main physics and dynamics at each model time step. It includes calls to all advection, diffusion, time-split

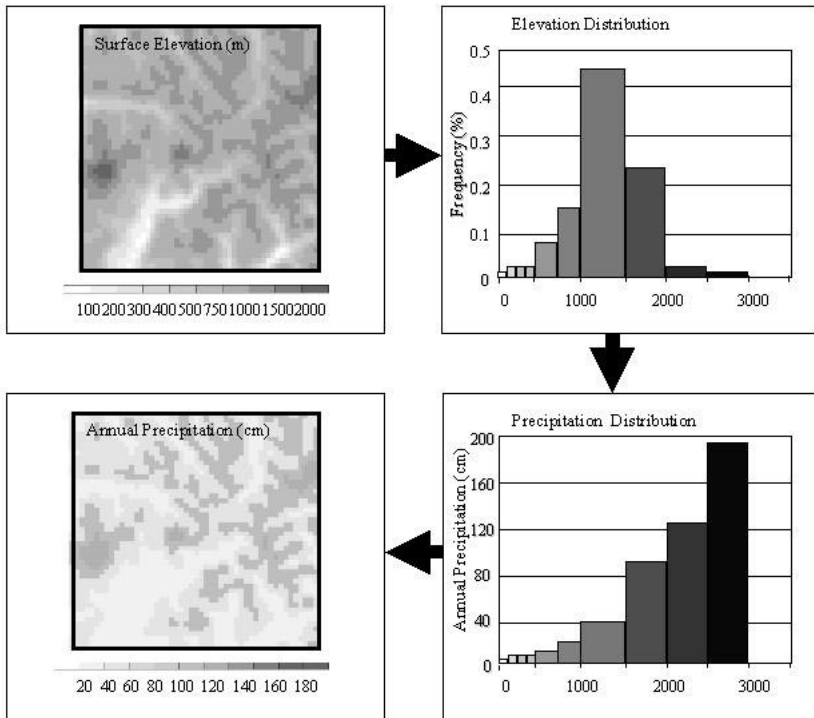


Fig. 1. Schematic illustration of the subgrid parameterization of orographic precipitation applied to a grid cell 50 km by 50 km in the western U.S. Upper left: surface topography within the grid cell at 1 km spatial resolution. Upper right: subgrid elevation classification. Lower right: simulations of precipitation at each subgrid elevation class. Lower left: mapping of precipitation to the geographical area based on elevation to yield high spatial resolution distribution of climate conditions for driving hydrology models

integration, and model physics routines. Most of these routines are called from SOLVE within loops over the J (east-west) dimension and compute one sweep of the I (north-south) dimension each time they are called. With the subgrid scheme, the enclosing J loops in the SOLVE routine are removed. Within the subroutines, the horizontal indices (I, J) are collapsed so that the iteration sweeps over the single index that runs from 1 to NHT , where NHT is the total number of subgrid elevation classes of all grid cells. This collapsing of indices is identical to a translation that FLIC performs for performance improvements on vector machines. Thus, the approach to integrating the PNNL subgrid parameterization easily leverages the overall same-source infrastructure already employed in MM5.

Parallelization involves decomposing subgrid arrays so that subgrid classes are on the same processor as the corresponding cells in the regular model grid. The NHT_{global} -sized elevation class arrays are decomposed so that the NHT_{local} -sized arrays on each processor contain variables for the elevation

classes corresponding to the grid cells in the local processor's subdomain. Because the distribution of elevation classes over grid cells is nonuniform, NHT_{local} may vary considerably from processor to processor in a simple equal-area decomposition. This is the basic source of load imbalance associated with the subgrid scheme.

3.2 Load Balancing

Load imbalance results from unequal distribution of subgrid elevation classes (NHT_{local}) when a domain is decomposed over processors. This load imbalance is static because NHT depends only on the spatial heterogeneity in surface topography that is determined once the domain is selected. Since a load-balancing mechanism already exists in the parallel MM5, a simple and effective approach to balancing the number of elevation classes is to redistribute the grid cells with which these classes are associated. We redistribute the cells in the regular MM5 grid to maximize $\frac{NHT_{global}}{p \max_p(NHT_{local})}$, where p is the number of processors and \max_p is the maximum over p processors. A load-balanced decomposition is computed at the beginning and remains in force for the duration of the model run.

The algorithm used to compute the decomposition is only a slight modification of the MM5 algorithm, which weights grid cells according to whether they are interior domain points (higher computational cost associated with physics calculations) or boundary points and then computes a decomposition that yields subdomains having close to the same aggregate weights. The subgrid load-balancing algorithm includes an additional cell-weighting factor, called *band-influence*, which determines the influence of the number of elevation classes associated with a grid cell. Figure 2a shows the static imbalance associated with the subgrid scheme in a 70 by 70 cells domain covering the western United States. The subgrid parameterization is applied to the interior 50 by 40 grid cells. The distribution of elevation bands is determined by a terrain dataset at 1 km spatial resolution. Cells with higher numbers of elevation classes are shown in lighter colors; black represents cells with only a single elevation class. The latter are found mostly over the ocean and near the boundaries where the subgrid scheme is not applied.

Figure 2b shows a 64-processor decomposition computed with a band-influence of zero—grid cells distributed more or less equally to different processors. The cost for an average time step on this domain is 513 milliseconds on an IBM SP. Figure 2c shows the domain decomposition that is computed with band-influence of one. The sizes of the local processor subdomains are varied to achieve a more uniform distribution of subgrid elevation bands to each processor. Here, the cost for an average time step is 269 milliseconds, a significant improvement.

4 Performance Results

A series of runs was performed on the IBM SP system at NCAR (Winterhawk-II, 4 375 MHz processors per node) using a 50 km resolution domain covering




Fig. 2. Distribution of subgrid elevation bands in the 70 by 70 grid cells domain in the western U.S. (a) lighter cells have more elevation bands (max. 12); black is single band per cell. Multiple elevation bands per cell create static load imbalance; (b) simple decomposition (band-influence=0.0); and (c) load-balanced decomposition (band-influence=1.0)

the western United States. The domain, as shown in Figure 2a, consists of 70 by 70 cells in the horizontal with 23 vertical layers. The time step was 150 seconds. Runs were conducted on 16, 36, 64, and 100 processors (4x4, 6x6, 8x8, and 10x10 decompositions, respectively). All runs were straight MPI, with four MPI processes per node. In each set of runs, the band-influence parameter in the modified MM5 load-balancing algorithm was varied from 0.0 (no influence of subgrid imbalance) to 1.1. Performance of each time step was measured using a millisecond timer and averaged over the last half-hour of a three-hour simulation, allowing sufficient time for spin-up of moisture fields. Initialization and I/O cost were ignored.

Figure 3 shows performance for these runs expressed as the number of model time steps executed per wall-clock second. All four sets showed improvement as band-influence was increased. Band-influence=1.0 was optimal for the 16, 36, and 64 processor runs; band-influence=0.9 was optimal for the 100 processor run. Load balancing provided a 96 percent improvement over the nonload-balanced

performance on 100 processors, 91 percent improvement on 64 processors, 88 percent on 36 processors, and 71 percent on 16 processors. The results indicate that for the processor counts tested, the benefit of load balancing increases with the number of processors. One expects this situation eventually to reverse with higher numbers of processors, however, since smaller subdomains will provide less opportunity for load balancing by redistributing grid points.

As expected, load balancing improves scaling efficiency, the speedup divided by the increase in the number of processors. Scaling with load balancing is 87 percent from 16 to 36 processors, 79 percent from 16 to 64 processors, and 64 percent from 16 to 100 processors. Scaling without load balancing is 79 percent from 16 to 36 processors, 70 percent from 16 to 64 processors, and 56 percent from 16 to 100 processors.

Performance of the subgrid scheme for the tested example compares quite favorably with traditional nesting. Employing the load-balanced subgrid scheme over half the area of the total domain required 2.05 times longer to run on 16 processors, 2.04 times on 36 processors, 1.7 times on 64 processors, and 1.98 times on 100 processors than without the subgrid scheme. A similarly sized nest would cost 14.5 times more than without a nest: an additional 4.5 times the number of cells time-stepping three times more frequently, plus the time for the original coarse domain.

5 Conclusions

Regional climate models are downscaling tools that enable the understanding and predictions of regional response to large-scale climate forcings. They can be used to provide spatially detailed seasonal climate forecasts and long-term climate projections useful for managing natural resources as well as serving as testbeds for developing physics parameterizations for global climate models.

Although the Penn State/NCAR Mesoscale Model MM5 was originally developed for short-term simulations of mesoscale weather phenomena, a community effort has been organized to add a capability for regional climate simulations. The following issues are being addressed to develop a Community Regional Climate Model (CRCM) based on MM5: (i) computational efficiency, (ii) stable numerics for long-term integration, (iii) lateral boundary condition formulation, (iv) a suite of physical parameterizations that provide accuracy and computational efficiency for long-term simulations, (v) model pre- and postprocessing, and (vi) a well-coordinated testing of different model components to ensure its suitability for long-term integration. As part of this community effort, the subgrid parameterization of orographic precipitation developed by Leung and Ghan has been implemented in the parallel MM5. This parameterization provides a computationally efficient alternative to the use of nesting for achieving simulations with high spatial resolution. With the example illustrated, model execution time increased approximately 2-fold using the subgrid scheme; the cost increase for a comparable refinement using a nested domain would have been 14-fold.

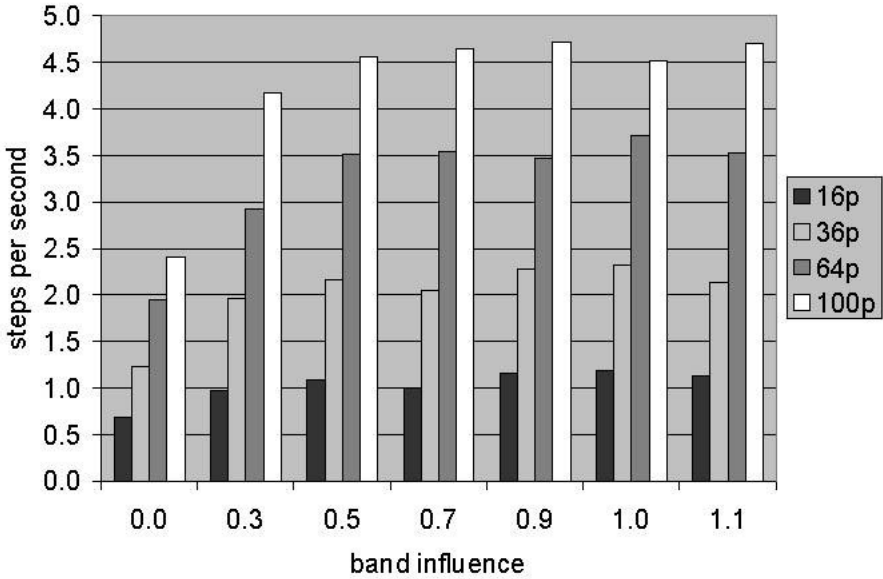


Fig. 3. Performance of a 70 by 70 cell domain on 20, 36, 64, and 100 IBM SP processors, varying the band-influence parameter in the MM5 load-balancing algorithm. Band-influence=0.0 disregards the load imbalance associated with the subgrid scheme; band-influence=1.0 gives each elevation class the full weight of an additional grid cell on the processor

Thus, the parallelized subgrid parameterization with load balancing represents a 7-fold savings over traditional nesting in MM5 for this scenario.

The implementation of the subgrid parameterization is consistent with the same-source approach to parallelization and vectorization adopted by the standard MM5. Changes are mostly transparent to users who opt not to use the parameterization. Load balancing is an important issue when the parameterization is applied to spatially diverse regions where the number of subgrid elevation classes vary strongly from one grid cell to another. The load-balancing algorithm in MM5 was modified to address this issue. In the near future, this parameterization will be implemented in the Weather Research and Forecast Model to improve its capability for regional climate simulations. Application of the subgrid scheme to enhance regional resolution in global climate simulations is also under way. The scheme has been implemented experimentally in the NCAR CCM3, and evaluation is being performed over different regions of the world.

References

1. Grell, G. A., J. Dudhia, and D. R. Stauffer: A Description of the Fifth-Generation Penn State/NCAR Mesoscale Model (MM5). Tech. Rep. NCAR/TN-398+STR, National Center for Atmospheric Research, Boulder, Colorado (1994).
2. Leung, L. R., and S. J. Ghan: Pacific Northwest Climate Sensitivity Simulated by a Regional Climate Model Driven by a GCM. Part I: Control Simulation. *J. Clim.* **12** (1999) 2010–2030.
3. Leung, L. R., and S. J. Ghan: Parameterizing Subgrid Orographic Precipitation and Surface Cover in Climate Models. *Mon. Wea. Rev.* **126** (1998) 3271–3291.
4. Leung, L. R., and S. J. Ghan: A Subgrid Parameterization of Orographic Precipitation. *Theor. Appl. Climatol.* **52** (1995) 2697–2717.
5. Leung, L. R., M. S. Wigmosta, S. J. Ghan, D. J. Epstein, and L. W. Vail: Application of a Subgrid Orographic Precipitation/Surface Hydrology Scheme to a Mountain Watershed. *J. Geophys. Res.* **101** (1996) 12803–12817.
6. Michalakes, J.: The Same-Source Parallel MM5. *J. Sci. Programming* **8** 2000 5–12.
7. 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.
8. Michalakes, J.: FLIC: A Translator for Same-source Parallel Implementation of Regular Grid Applications, Tech. Rep. ANL/MCS-TM-223, Argonne National Laboratory (1997).