DISTRIBUTION STATEMENT A: Approved for public release; distribution is unlimited.

Next-Generation Global and Mesoscale Atmospheric Models

Francis X. Giraldo
Department of Applied Mathematics
Naval Postgraduate School
833 Dyer Road
Monterey, CA 93943-5216
phone: (831) 656-2293 fax: (831) 656-2355 e-mail: fxgirald@nps.edu

Award Number: N0001408WR20136
http://faculty.nps.edu/fxgirald/projects/mesoscale

LONG-TERM GOALS

The long-term goal of this research is to construct a unified global and mesoscale nonhydrostatic numerical weather prediction (NWP) models for the U.S. Navy using new numerical methods specifically designed for modern computer architectures; this unified model is called the Non-hydrostatic Unified Model of the Atmosphere or NUMA. To take full advantage of distributed-memory computers, the global domains of these new models are partitioned into local sub-domains, or elements, that can then be solved independently on multiple processors. The numerical methods used on these sub-domains are local, high-order accurate, fully conservative, and highly efficient. Using these ideas we are developing global and mesoscale nonhydrostatic atmospheric models that will improve upon the operational models currently used by all U.S. agencies including the U.S. Navy.

OBJECTIVES

The objective of this project is to construct new high-order local methods for the Navy's next-generation global and mesoscale nonhydrostatic NWP models. The high-order accuracy of these methods will ensure that the new model yields better forecasts than the current global spherical harmonics model (NOGAPS) and better accuracy than the current mesoscale finite difference model (COAMPS). The objective is to achieve this accuracy while increasing the geometric flexibility to use any grid as well as to increase the efficiency of these models on large processor-count distributed-memory computers. Higher efficiency means that the new models will require less computing time that then allows for increasing the number of ensemble members and/or increasing the resolution of the NWP models. The methods that we propose to use for these models are state-of-the-art and are not being used by either current or newly emerging NWP and climate models.

APPROACH

To meet our objectives we explore:

1. spectral element (SE) and discontinuous Galerkin (DG) spatial discretization methods;

2. high-order semi-implicit (SI) time-integrators with adaptive time-stepping for vastly improved efficiency;
3. unified hydrostatic, nonhydrostatic, and pseudo-incompressible formulations of the equations;
4. fully unstructured and adaptive grids; and
5. scalable parallel implementation of the models.

The power of SE and DG methods is that they are high-order accurate yet are completely local in nature – meaning that the equations are solved independently within each individual element and, on parallel computers, processor. Furthermore, high-order methods have minimal dispersion error – this is an important property for capturing fine-scale atmospheric phenomena (e.g., tropical cyclones, Kelvin and Rossby waves). The theoretical development of SE and DG methods are now well established and these methods are currently the two most successful methods found in the literature for fluid flow problems.

Semi-implicit (SI) and fully-implicit (FI) time-integrators offer vast improvements in efficiency due to the longer time steps that they permit; semi-implicit methods can be classified under the heading of implicit-explicit (IMEX) methods that has garnered much attention in the computational mathematics literature. Furthermore, in order to reap the full benefits of the high-order spatial discretization methods requires increasing the order of accuracy of the time-integration methods as well; this is a topic that too often has been ignored by most scientific computing communities, including the NWP community.

Before committing resources towards the development of new NWP models, it is important to identify the form of the governing equations that is most capable of conserving all quantities deemed important. We have been performing studies on this topic for the past three years – that is, to identify the form of the governing equations capable of representing conservation of either mass, energy, or both. In addition, we have analyzed various forms of the governing equations with respect to robustness, flexibility, and efficiency in the context of implicit-explicit (IMEX) time-integration methods. Within this work we also explored hybrid models that solve either the hydrostatic, nonhydrostatic, or pseudo-incompressible equations. This feature allows the models to be used for research purposes by Navy scientists in order to test the importance of multi-scale phenomena at specific resolutions.

One final area that needs to be explored is the concept of adaptive grids. In the past few years, adaptive grids have gained considerable momentum in the atmospheric modeling community – I will participate in a four month program on this topic in 2012 at the Newton Institute in Cambridge University, England.

WORK COMPLETED

In this section, we describe the work completed this fiscal year. The work can be categorized into four sections: the governing equations, time-integration methods, spatial discretization methods, and parallel implementation.

Nonhydrostatic Governing Equations. In FY 2009, we analyzed 5 different forms of the governing equations of motion and determined that Set 2 in both conservation (2C) and non-conservation (2NC) form were optimal for building a non-hydrostatic atmospheric model.

Set 2C is defined as follows:

$$\frac{\partial \rho}{\partial t} + \nabla \bullet \mathbf{U} = 0; \quad \frac{\partial \mathbf{U}}{\partial t} + \nabla \bullet \left(\frac{\mathbf{U} \otimes \mathbf{U}}{\rho} + P\mathbf{I}_2 \right) = -f\left(\mathbf{k} \times \mathbf{U}\right) - \rho g \mathbf{k}; \quad \frac{\partial \Theta}{\partial t} + \nabla \bullet \left(\frac{\Theta \mathbf{U}}{\rho} \right) = 0; \quad P = P_A \left(\frac{R\Theta}{P_A} \right)^{\gamma}$$

where the solution vector is density, momentum, and density potential temperature. Set 2 is the form used in WRF. This form is very attractive because it conserves mass although it does not conserve energy. It does, however, conserve density potential temperature which is related to entropy.

This set is of interest because it can also be written in non-conservation form while still conserving mass. Set 2NC is defined as follows:

$$\frac{\partial \rho}{\partial t} + \nabla \bullet \left(\rho \mathbf{u}\right) = 0; \quad \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \bullet \nabla \mathbf{u} + \nabla P = -f\left(\mathbf{k} \times \mathbf{u}\right) - g\mathbf{k}; \quad \frac{\partial \theta}{\partial t} + \mathbf{u} \bullet \nabla \theta = 0; \quad P = P_A \left(\frac{\rho R \theta}{P_A} \right)^{\gamma}$$

The interest in equation Set 2NC is that it conserves mass exactly and can conserve energy if the advection term in the momentum equation is replaced by its rotation form; we are currently working on this aspect and will report on the success of this idea in the future.

The analysis of various forms of the governing equations can be found in our recent paper (see Ref. [5]). The summary of that paper is that set 2NC is the form that we should use in order to construct the optimal nonhydrostatic model, taking into account: accuracy, efficiency, and conservation measures. This year, we used this equation set (both SE2C and SE2NC) to build a 3D global/local nonhydrostatic model.

Time-Integrators. Directly connected with our choice of equation sets is the resulting semi-implicit operators. In studies performed last year, we compared semi-implicit time-integrators both in their Schur (i.e., pseudo-Helmholtz) and No Schur (i.e., full system) forms and compared them to the types of explicit time-integrators currently being used in split-explicit models. Our results show that if the Schur form is used, then the semi-implicit models are always faster than explicit models. However, as is shown in the Results section, much work still remains in order to scale implicit methods on massively parallel computers.

Spatial Discretization Methods. We have been arguing in the course of this project that the best next-generation models will be those based on element-based Galerkin (EBG) methods such as the spectral element (SE/CG) and discontinuous Galerkin (DG) methods. However, we have only partly showed the benefits of this approach such as: high parallel efficiency and high-order accuracy. We have also hinted in the past that another attribute of these methods is that they allow one to use unstructured and adaptive grids. We have teamed up with the University of Mainz (Germany) and the University of Hamburg (Germany) to combine the triangle-based discontinuous Galerkin method with adaptivity. Although only preliminary, the results show that indeed this combination will be quite formidable for tackling nonhydrostatic problems (such as tracking hurricanes, storm-surges, etc.).

One of the advantages of the NUMA code is that it is capable of using both spectral elements (that we call continuous Galerkin or CG methods) or discontinuous Galerkin methods. Even though the spectral element method is shown to scale linearly up to 16,000 processors, we expect the discontinuous Galerkin (DG) version of the code to scale for much larger processor counts.

Parallel Implementation. This year, we were able to extend the 2D serial models that we developed for the compressible Navier-Stokes equations to 3D parallel. The first step was to extend the 2D serial models to 3D serial. This step represented a significant amount of work because the model was rewritten from scratch in order to plan ahead for its parallel implementation. We used modern software engineering approaches to achieve this task such as building the Fortran 95 model in terms of modules – this proved to be the reason for the straightforward extension of the serial model to MPI. The 3D MPI code is currently setup to use only CG and using Set 2NC; we will extend Set 2C that contains both the CG and DG methods to MPI in the next year.

RESULTS

For the sake of brevity, we shall only discuss the results of our 3D MPI model. We split this section into: validation of the model, domain decomposition approach, and scalability of the model.

Model Validation. The 3D serial model was first tested against our 2D serial models to ensure that we recovered the exact same results up to machine double precision. We now have a set of 3D serial models that can run all of the 2D test cases. At this point, we began implementing 3D test cases such as the 3D rising thermal bubble and 3D linear hydrostatic isolated mountain. The bubble problem does not have an analytic solution but the evolving dynamics is sufficiently simple to use as a good first test case (e.g., we can check the symmetry of the bubble in the horizontal as well as the conservation of mass and energy).

Next, we used the isolated mountain case which we can compare against semi-analytic solutions. The analytic solutions are only valid in the near-vicinity of the mountain surface. This test case is the most realistic idealized test case that we can run because it has almost all of the difficulties found in a real limited-area model run: such as orography (the mountain) and the necessity of using non-reflecting boundary conditions. Our models have all of these components.

In Figure 1 we show the grid geometry for the 3D isolated mountain test. On the left panel we show the orography (an isolated mountain) and on the right panel we show the hexahedral elements used to tile the domain. On the right panel we only show the vertices of the hexahedra but it should be understood that inside each hexahedra, an arbitrarily high-order polynomial can be used that gives the spectral element (SE) method its unprecedented high-order accuracy while maintaining the capacity to use completely unstructured grids.

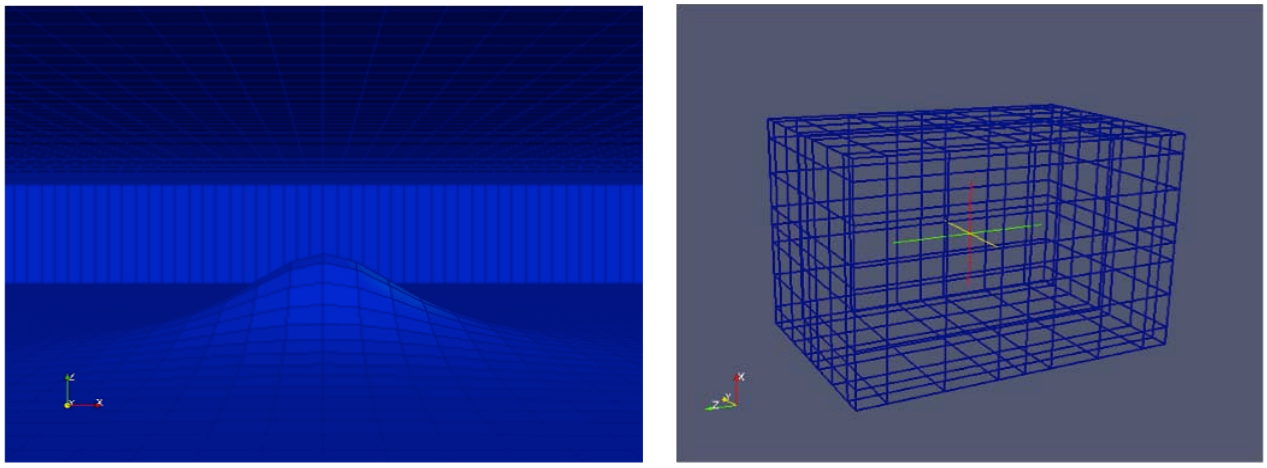

a) ***b)***

Figure 1: 3D Linear Hydrostatic Isolated Mountain with a) the orography and the structure of the 3D grid.

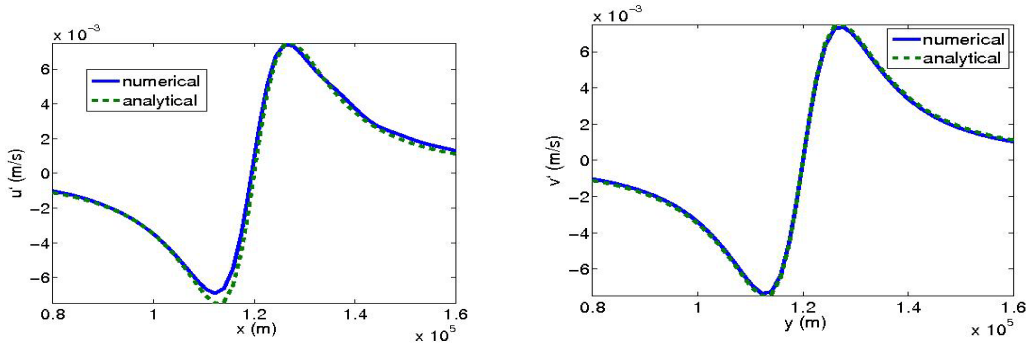

a) ***b)***

Figure 2: 3D Linear Hydrostatic Isolated Mountain with a) the u-velocity component and b) the v-velocity component.

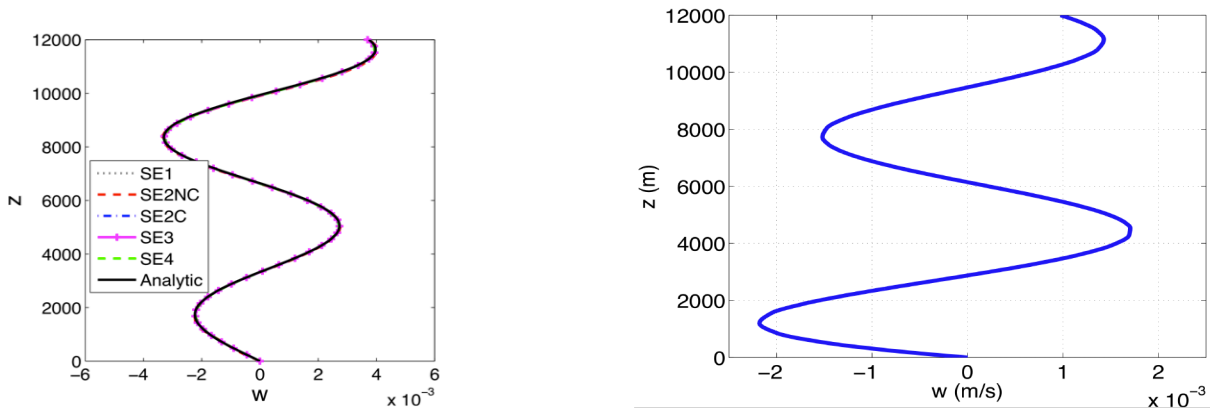

a) ***b)***

Figure 3: The vertical velocity profile for the a) 2D linear hydrostatic mountain ridge and b) 3D linear hydrostatic isolated mountain.

In Fig. 2, we compare the analytic and numerical solutions after 10 hours of simulation. However, there is no known analytic solution for the vertical velocity component. Instead, in Fig. 3 we compare the vertical velocity measured from the top of the mountain through the top of the atmosphere with the results from the linear hydrostatic mountain ridge; this problem is essentially the 2D version of the 3D isolated mountain. Note that the panel on the left and the panel on the right are not identical but quite similar especially near the surface and the very top. The main differences are near 8 kms but overall, the solutions look quite similar.

Domain Decomposition. We now describe the domain decomposition strategy. One of the many nice features of NUMA is that it is capable of using any unstructured grid; currently, it only allows for grids based on hexahedra but triangular prisms will be added in FY 2011. The difficulty in using unstructured grids is in devising a domain decomposition strategy that allows for (near) perfect load balancing, meaning that each processor will have approximately the same number of unknowns. In Fig. 4 we show some possible grids that NUMA can use.

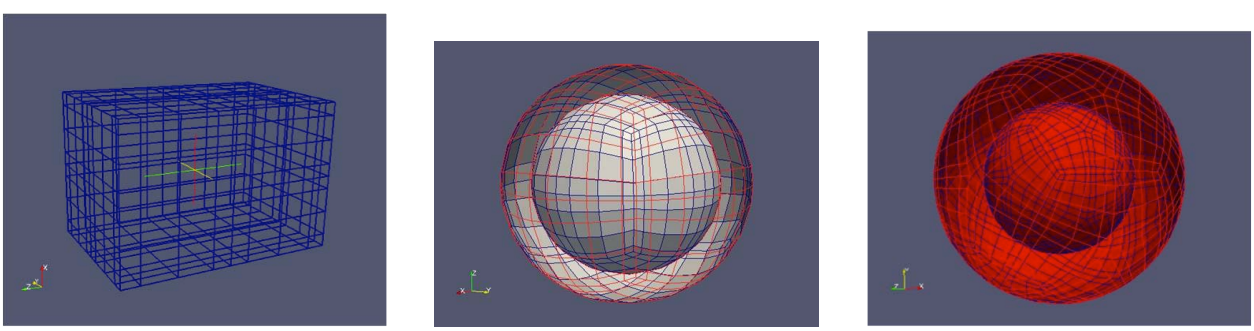

a) ***b)*** ***c)***

Figure 4: Some Possible Grids used by NUMA: a) a Cartesian grid, b) a cube-sphere grid and c) an icosahedral grid.

In Fig. 5, we show the domain decomposition of these grids. The different colors denote the different processors. In order to allow NUMA to use any unstructured grid, we had to rewrite NUMA in order to handle processors with variable number of spectral elements. That is, for a specific problem size, we can use any number of processors because we now partition the grid across processors using the METIS software. This software partitions the grids by minimizing the number of edge cuts (processor boundaries) and therefore results in imperfect load balancing but very efficient communication strategies.

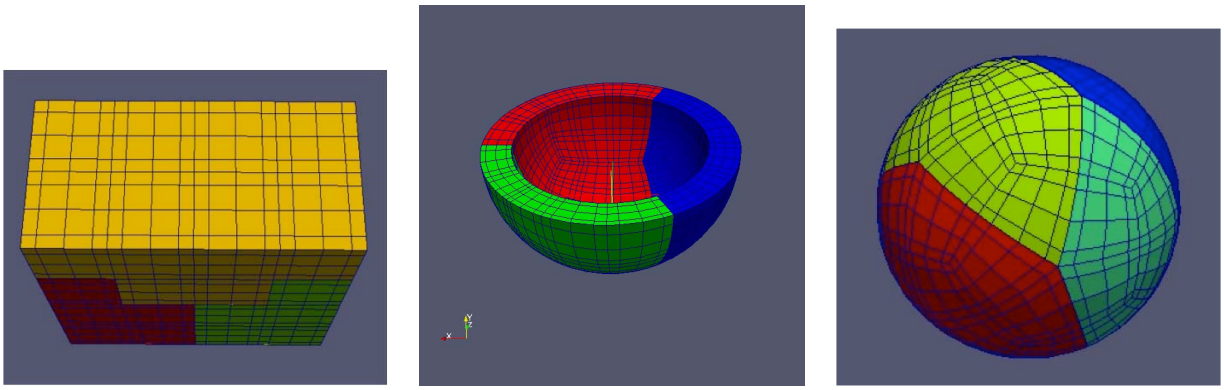

a) ***b)*** ***c)***

Figure 5: Domain Decomposition of the NUMA grids: a) a Cartesian grid, b) a cube-sphere grid and c) an icosahedral grid.

<u>*Model Scalability*</u>. Using the domain decomposition strategy described above, we have developed a very general MPI (message-passing interface) implementation of the 3D NUMA-CG; that is, the 3D NUMA using continuous Galerkin methods with Set 2NC. We have also developed a 3D NUMA-CG/DG for Set 2C but this code is currently only serial; we will extend this code to MPI in FY 2011 as well as develop a 3D NUMA-CG/DG for Set 3 which is the form based on density and total energy.

For 3D NUMA-CG-2NC we have run preliminary scalability studies of the model. In Fig. 6 we show the results for various time-integrators. On the left panel (Fig. 6a) we show the results for a simulation with 2 million grid points comparing an explicit Runge-Kutta (RK) method and the semi-implicit backward difference formula (BDF2-SI) using both the Schur and No Schur forms. This figure shows that both BDF2-SI results yield more efficient results than the explicit method at least up to 128 processors. However, although subtle, the results at 128 processors reveal that the perfect scalability of the semi-implicit methods is beginning to breakdown whereas the explicit code is able to sustain this scalability. In Fig. 6b (right panel) we show results for the explicit version of NUMA but now using 16 million grid points. It is quite clear that this code is able to scale linearly up to the 12,000 processors that we have used. We expect this code to scale for processor counts reaching 100,000 processors.

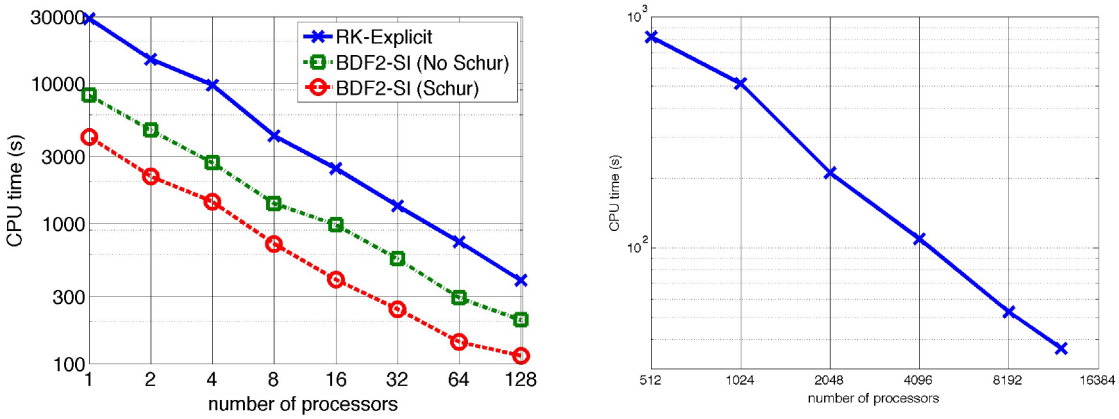

a) ***b)***

Figure 6: Scalability of NUMA on Ranger: a) three time-integrators with 2 million grid points and b) explicit time-integrator with 16 million grid points.

We believe that we can redesign the semi-implicit time-integrators to scale linearly up to tens of thousands of processors. This of course will require a new class of time-integrators, iterative solvers and preconditioners. We hope to report on the implementation of such methods into NUMA next year.

IMPACT/APPLICATIONS

NOGAPS and COAMPS are run operationally by FNMOC and is the heart of the Navy's operational support to nearly all DOD users worldwide. This work targets the next-generation of these systems for massively parallel computer architectures. NUMA has been designed specifically for these types of computer architectures while offering more flexibility, robustness, and accuracy than the current operational systems. Additionally, the new models are expected to conserve all quantities such as mass and energy and use state-of-the-art time-integration methods that will greatly improve the capabilities of the Navy's forecast systems.

TRANSITIONS

Improved algorithms for model processes will be transitioned to 6.4 as they are ready, and will ultimately be transitioned to FNMOC.

RELATED PROJECTS

Some of the technology developed for this project could be used to improve NOGAPS in other NRL projects. The work on the mesoscale models will help improve COAMPS. An example is the time-integration methods that we are exploring for the new models may well be incorporated into the current operational version of COAMPS. The insight gained on grids in the current project could also be leveraged to develop a global version of COAMPS by virtue of the cubed-sphere grid (hexahedral) and data structures developed for NUMA.

PUBLICATIONS

Journals

1. Restelli, M., and Giraldo F.X., 2009: A Conservative Semi-Implicit Discontinuous Galerkin Formulation for the Navier-Stokes Equations in Nonhydrostatic Mesoscale Modeling. *SIAM Journal of Scientific Computing*, Vol. 31, pp. 2231-2257.
2. Dea, J., Giraldo F. X., and Neta B., 2009: High-Order Non-Reflecting Boundary Conditions for the Linearized 2D Euler Equations. *Wave Motion*, Vol. 46, pp. 210-220.
3. Giraldo, F.X., and Restelli M., 2010: High-Order Semi-Implicit Time-Integrators for Discontinuous Galerkin Methods. *International Journal for Numerical Methods in Fluids*, Vol. 63, pp. 1077-1102.
4. Lindquist, J., Neta, B., and Giraldo, F.X., 2010: Spectral Element Solution of the Klein-Gordon Equation with High-Order Boundary Treatment and Time. *Wave Motion* Vol. 47, pp. 289-298.
5. Giraldo, F.X., Restelli M., and Läuter, M., 2010: Semi-Implicit Time-Integrators for the Navier-Stokes Equations: Applications to Nonhydrostatic Atmospheric Modeling. *SIAM Journal on Scientific Computing*, in press.
6. Lindquist, J., Giraldo, F.X., and Beta, N., 2010: Klein-Gordon Equation with Advection on Unbounded Domains using Spectral Elements and High-Order Non-Reflecting Boundary Conditions. *Applied Mathematics and Computation,* in press.
7. Gabersek, S., Giraldo, F.X., and Doyle, J.D., 2010: Simple Microphysics Experiments with a Spectral Element Model. *Monthly Weather Review* [in preparation].

Theses

1. Dea J., 2009: *High-Order Non-Reflecting Boundary Conditions for Linearized Euler Equations*. PhD Thesis, Naval Postgraduate School [published].
2. Gibbons, S.L., 2009: *Impact of Sigma Coordinates on Continuous and Discontinuous Galerkin Navier-Stokes Models for Nonhydrostatic Atmospheric Modeling*. Master's Thesis, Naval Postgraduate School [published].
3. Lindquist, J. 2010: *A Spectral Element Solution of the Klein-Gordon Equations with High-order Treatment in Time and Non-Reflecting Boundary Conditions*. PhD Thesis, Naval Postgraduate School [published].

Conference Abstracts

1. Kelly, J.F., Giraldo, F.X., and Carr, L.E., 2010: Parallel Implementation of a Spectral Element Navier-Stokes Model. *2010 Partial Differential Equations on the Sphere*, Potsdam Germany (August 2010).

Invited Talks

1. Giraldo, F.X., 2010: A Research Framework for the Construction of Next-Generation Unified Atmospheric Models, Institute for Pure and Applied Mathematics, UCLA (April 2010).
2. Giraldo, F.X., 2010: Mathematical Derivation of the Non-hydrostatic Unified Atmospheric Model (NUMA), Oberwolfach, Mathematische Forschung Institute, Oberwolfach Germany (August 2010).
3. Giraldo, F.X., 2010: Optimal Implicit Time-Integrators for the Compressible Navier-Stokes Equations, CWI, Amsterdam Netherlands (August 2010).
4. Giraldo, F.X., 2010: Development of the Non-hydrostatic Unified Atmospheric Model (NUMA), ESPC Meeting, Boulder CO (September 2010).
5. Giraldo, F.X., 2010: Algorithmic and Mathematical Formulation of the Non-hydrostatic Unified Atmospheric Model (NUMA), Brown University, Providence RI (September 2010).
6. Giraldo, F.X., 2010: Developing a continuous/discontinuous Galerkin formulation in the Non-hydrostatic Unified Atmospheric Model (NUMA), Cal State Northridge, Northridge CA (October 2010).