

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/fgiraldo/projects/NUMA>

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 and graphical processing units or GPUs. The numerical methods used on these sub-domains are local, high-order accurate, fully conservative, highly efficient, and geometrically flexible. Using these ideas we are developing global and mesoscale nonhydrostatic atmospheric models that will improve 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 improve the accuracy of the dynamics for the current global spherical harmonics model (NOGAPS) and the finite difference local-area model (COAMPS). It is conjectured that improving the accuracy of the dynamics (including tracers) will increase the overall accuracy of the forecast; however, such improvements will also have to be made to both the physics and the data assimilation systems but these two topics are currently beyond the scope of this project. The objective of this project is to improve the accuracy of the dynamics while increasing the geometric flexibility in order 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 models.

APPROACH

To meet our objectives we explore:

1. Unified high-order continuous Galerkin (CG) and discontinuous Galerkin (DG) spatial discretization methods;
2. Unified high-order implicit-explicit (IMEX) time-integrators with adaptive time-stepping for vastly improved efficiency;
3. unified global and mesoscale models based on the nonhydrostatic equations;
4. fully unstructured and adaptive grids; and
5. scalable parallel implementations.

The power of CG 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, may reside on separate processors. 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 CG 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, but has been at the forefront of research efforts in this project.

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 few 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. Two of the papers developed within this project are now becoming the standard papers for identifying the proper equation sets and test cases for nonhydrostatic atmospheric modeling (Giraldo-Restelli JCP 2008 and Giraldo et al. SISC 2010). Accordingly to Google Scholar, Giraldo-Restelli JCP 2008 already has 70 citations (14 citations per year).

One final area that needs to be explored is the concept of adaptive mesh refinement (AMR). In the past few years, adaptive grids have gained considerable momentum in the atmospheric modeling community. From August through December 2012, I participated in a four month program on this topic at the Newton Institute in Cambridge University. During this program we showed that the form of AMR that we are developing (called nonconforming) is by far the most efficient for constructing

variable resolution grids. I participated in a study on AMR and variable grids at the Newton Institute where we will publish our finding this year.

WORK COMPLETED

In this section, we describe the work completed this fiscal year. The work can be categorized into four sections: unified IMEX time-integration methods, unified spatial discretization methods, adaptive mesh refinement, and the inclusion of moist processes into NUMA.

Unified IMEX Time-Integrators. In studies performed in 2010-2012, we compared IMEX 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 explicit models. Our results show that if the Schur form is used, then the IMEX methods are always faster than explicit models. Furthermore, we have analyzed the types of problems that are amenable to an IMEX approach in all directions (which we call 3D-IMEX) versus IMEX approaches only in the vertical (which is known as HEVI, Horizontal Explicit – Vertically Implicit, or 1D-IMEX). This study is contained in our publication in SIAM Journal of Scientific Computing (see Ref. [1]).

One of the main results of that paper is that we have been able to extract the radial direction of our model in a general way in order to build the 1D-IMEX methods for both multi-stage and multi-step methods. This general 1D-IMEX is the clear winner for most global modeling problems that we have considered due to the disparate horizontal versus vertical scales typically used in global modeling (e.g., the horizontal resolution is typically around 50km whereas the vertical resolution is 1km). An additional advantage of 1D-IMEX is that it scales as perfectly on parallel computers as a fully explicit method if the domain decomposition is done across the horizontal direction (which is explicit) while the points along the vertical column are on processor and so the IMEX method incurs no MPI communication per solve. On the other hand, having the capacity to use the 3D-IMEX approach will allow us to use far more processors which we will have access to in the near future. The difficulty with this idea is that much work has to be invested in iterative solvers and preconditioners. We have developed a strategy that we believe will scale for very large processor counts and NUMA now contains both the 1D-IMEX as well as the 3D-IMEX; the results are shown in Ref. [1].

Adaptive Mesh Refinement. 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. Last year, we began a study of adaptive non-conforming quadrilateral grids (we published our conforming triangular grid results in Ref. [2]). One of the attractions of NUMA is that it uses either CG or DG methods. In other words, the CG and DG spatial discretizations have been written in a unified way which allows the model to use either method and either method can now be run with adaptive mesh refinement. Our AMR algorithm only takes 1% of the total runtime and works well in unison with our IMEX methods, iterative solvers, and preconditioners (see Ref. [3]). A publication detailing the approach is in preparation.

Moist Processes. Most of the work on NUMA conducted by the group at NPS has been focused on the dynamical core and new features that are related to it (e.g., time-integrators, positivity-preservation, adaptive mesh refinement) while the subgrid-scale parameterization has been the focus of the NRL

group. This year, however, the NPS group has focused on including moist processes within NUMA to test the capacity for the model to handle physical parameterization on non-uniform vertical grids as well as how to handle the time-integration of these processes. Through a collaboration with NRL we assisted in the publication of a paper on moist physics with the CG code in 2012 (see Ref. [5]). In that work, moist physics was implemented within a research version of NUMA. Now, NUMA is a far more sophisticated and professional software package that has been re-designed to simplify the inclusion of new variables and physical processes. This year, we have re-implemented the moist processes described in Ref. [5] and have obtained the same solutions with the NUMA model. The significance of this is that we can now reproduce the exact same 2D results described in Ref. [5] by our fully 3D highly-parallel NUMA code. With this in place, we are now in a position to run far more realistic moist simulations and are now a few steps closer to begin running hurricane simulations.

RESULTS

The main results for this year can be divided into the following sections: 1) development of a new IMEX time-integrator; 2) adaptive mesh refinement for both CG and DG; and 3) the inclusion of moist processes within 2D and 3D NUMA-CG codes and into the 2D NUMA-DG code.

New IMEX Time-Integrator. In Ref. [1] we were able to derive all Implicit-Explicit methods to be written in the exact same way using the usual semi-implicit formulation used by all operational centers. The advantage of this is that any IMEX method (multistep, multistage, etc.) can be incorporated easily into an existing model using our approach. NUMA currently contains the following time-integrators: explicit Runge-Kutta, Adams-Bashforth, Backward Difference Formulas, and Leapfrog and also implicit-explicit Runge-Kutta, Adams-Moulton, Backward Difference Formulas, and the Trapezoidal rule. All these methods are defined in a similar manner that the standard Leapfrog-Trapezoidal rule is defined. This allows NUMA to add new time-integrators quite easily. In Ref. [1] we compared all of these time-integrators and have found the IMEX Runge-Kutta methods to be the clear winners in terms of accuracy versus cost. We introduced a new 2nd order IMEX-RK method that has been shown to be the most accurate 2nd order method across many metrics. Colleagues at ECMWF, the University of Reading, and the UK Met Office have confirmed this as well in a series of papers published on IMEX methods. The stability analysis of various 2nd order methods considered can be found in Fig. 1.

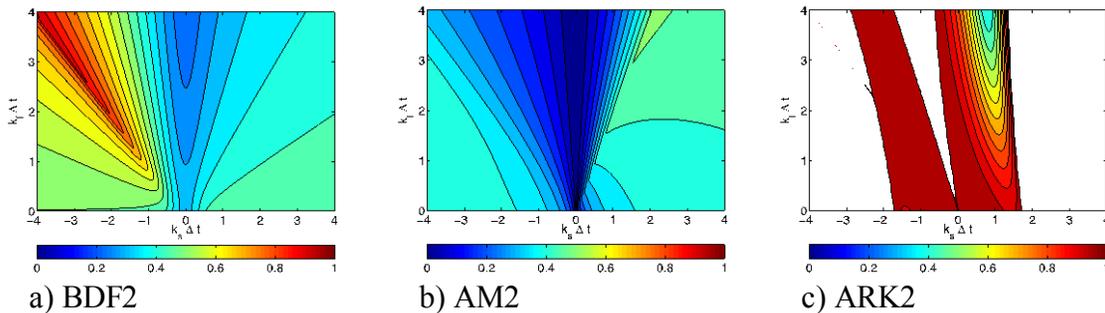


Figure 1: The stability analysis for three second-order IMEX methods a) backward difference formula (BDF2), b) Adams-Moulton method (AM2), and c) the new Runge-Kutta method (ARK2). The red region in ARK2 shows that the method is not dissipative in the region of stability whereas the blue regions in BDF2 and AM2 show that these methods are stable within a larger region but are dissipative.

The analyses and results in Ref. [1] show that ARK2 is almost an order of magnitude more accurate than the other two second order methods. This new IMEX-ARK2 method has been tested under various NWP-like scenarios and has become the default time-integrator in NUMA. It is very robust, has automatic time-stepping built into it, is effective with AMR and is designed with dense output in mind, meaning that the method has the capacity to evaluate the solution at any time (not just at the end of a time-step) without loss of accuracy since a built-in interpolation function comes with the method. In Fig. 2 we show the results of all of our time-integrators in NUMA.

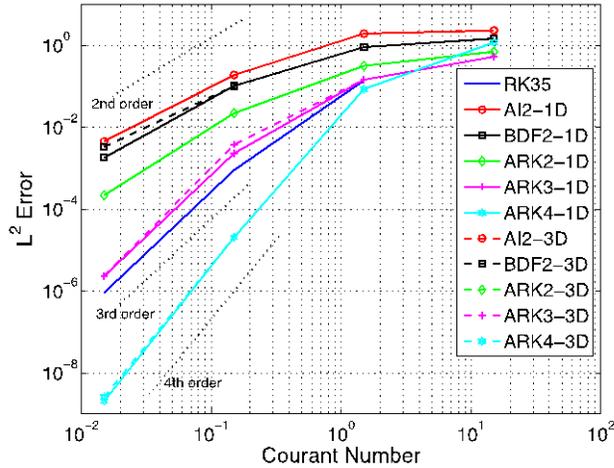


Figure 2: The accuracy of NUMA using various time-integrators for a global atmospheric problem (propagation of an acoustic wave).

Figure 2 shows the rates of convergence revealing that each method yields the expected theoretical rate of convergence (the dashed line show these ideal rates). We can see that the ARK2 result is about an order of magnitude more accurate than the BDF2 and AM2 results. This figure shows the results for a global atmospheric problem but numerous other simulations have been shown to give the same behavior (see Ref. [1] for more simulations). We have not discussed preconditioners and iterative solvers but these are discussed in detail in Ref. [3] where we derive a new class of preconditioners specifically designed for our IMEX methods.

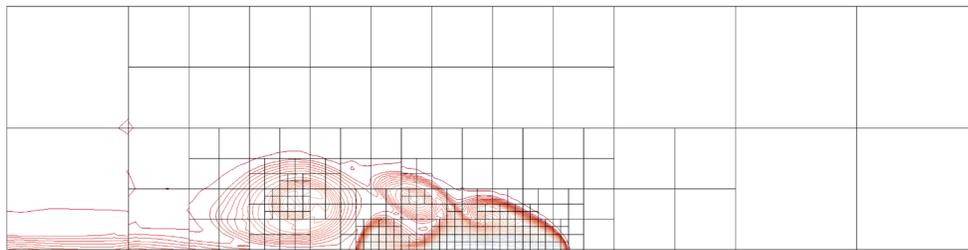


Figure 3: The grid (black lines) and the density current contours (red, blues, etc.) for potential temperature during the simulation. The grid has followed the rotors throughout the simulation.

Adaptive Mesh Refinement. One of the distinguishing features of NUMA is that it is the only model from all the ESPC-considered models that has the capacity for adaptive mesh refinement (AMR).

Currently, AMR is only available in the 2D versions of NUMA but the extension to 3D-NUMA will begin this year. Although it is not clear if AMR can be used for all types of simulations, it is certainly a good candidate for high-resolution hurricane simulations. To explain AMR let's consider Fig. 3 where a density current simulation is illustrated.

Figure 3 shows the potential temperature contours plotted with the grid. Note that some grid boxes are smaller than others. The AMR algorithm automatically figures out which grid box needs to be small (where something is happening) and which boxes can be large (where nothing is happening). This way one can follow an interesting feature with high precision yet use far fewer degrees of freedom compared to a uniform high-resolution simulation where all the grid boxes are small.

The question that one might ask is: is AMR faster than using a uniform simulation? In other words, what is the overhead in using the AMR algorithm? In Fig. 4 we show the speed-up (cost of the uniform simulation divided by the cost of the AMR simulation in CPU time).

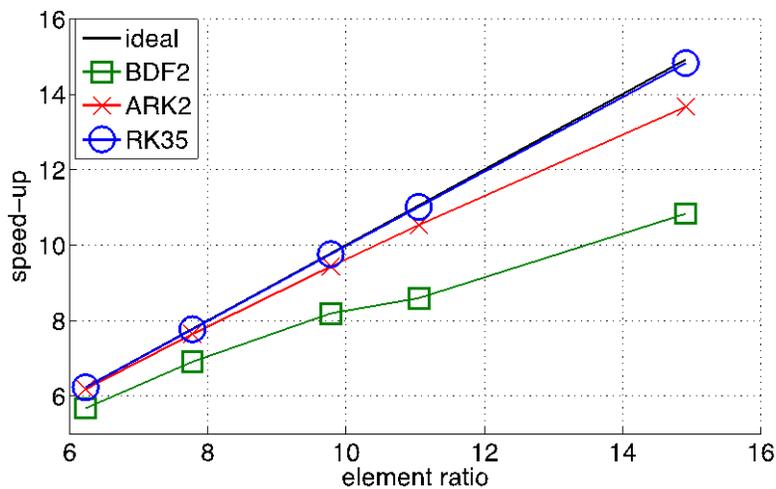


Figure 4: The solid black line below the blue circles is the ideal speed-up. I.e., if you use 10 times fewer elements (element ratio) the simulation should be 10 times faster (the speed-up).

Figure 4 shows that we get perfect speedup for the explicit RK35 method. This is a good result because it shows that the AMR algorithm adds very little cost to the total wallclock time of the simulation. On the other hand, we get less than ideal for both IMEX-BDF2 and IMEX-ARK2; note however that the IMEX-ARK2 is close to ideal. The reason why the IMEX methods do not achieve perfect speed-up is because every time a new grid is produced by the AMR algorithm, the IMEX time-integrators require the preconditioners to be recomputed. The results shown in Fig. 4 represent a worst case scenario since we apply AMR at every time-step; in simulations there is rarely significant changes from one time-step to the next and so we can reduce the effects of this overhead by calling the AMR algorithm every 10 time-steps or so.

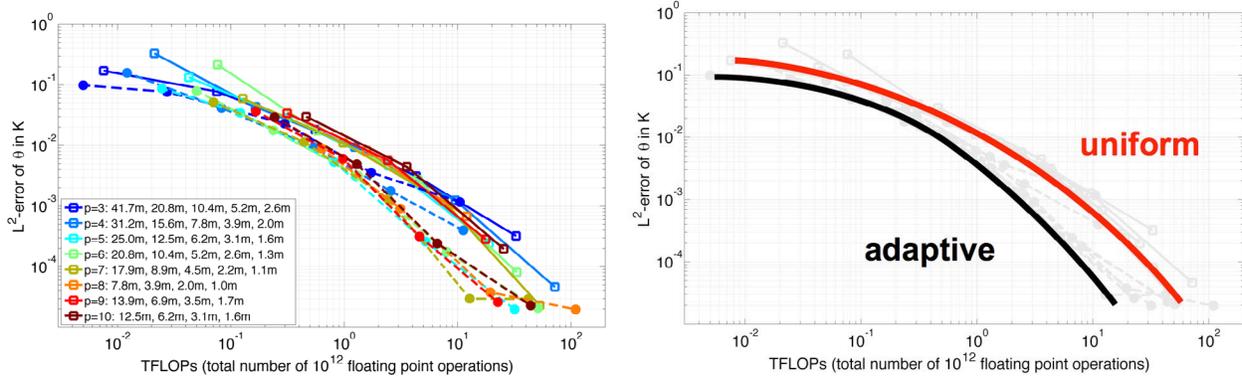


Figure 5: The L2 error versus total floating point operations for various simulations. The left plot shows the detailed simulations while the right plot shows the general trends of the simulations categorized by adaptive (AMR) and uniform grid.

Note that Fig. 4 says nothing about the error incurred by AMR only that it is much faster than a uniform simulation. The question one might ask now is: which is faster for a certain level of accuracy, AMR or a uniform simulation. To help answer this question we ran a detailed study that can be summarized by Fig. 5.

The left panel in Fig. 5 shows the details of various simulations using various AMR element ratios and different polynomial orders (orders of approximation) of our method. The right panel in Fig. 5 summarizes the results by collecting the AMR and uniform grid simulations. This panel shows that the AMR (adaptive) simulations can be as much as 5 to 10 times faster than the uniform simulations particularly when high-order polynomials are used (bottom right portions of the curves). To our knowledge, such a detailed study of the cost of AMR has not been conducted before.

Moist Simulations. In Ref. [4] we included moist processes to NUMA and showed that NUMA, although a high-order model, is more efficient than a 2nd order finite difference model (such as COAMPS) when one considers not only the wallclock time but also the quality of the solution obtained. This year we have added moist processes to the 3D code as well as to the 2D DG code.

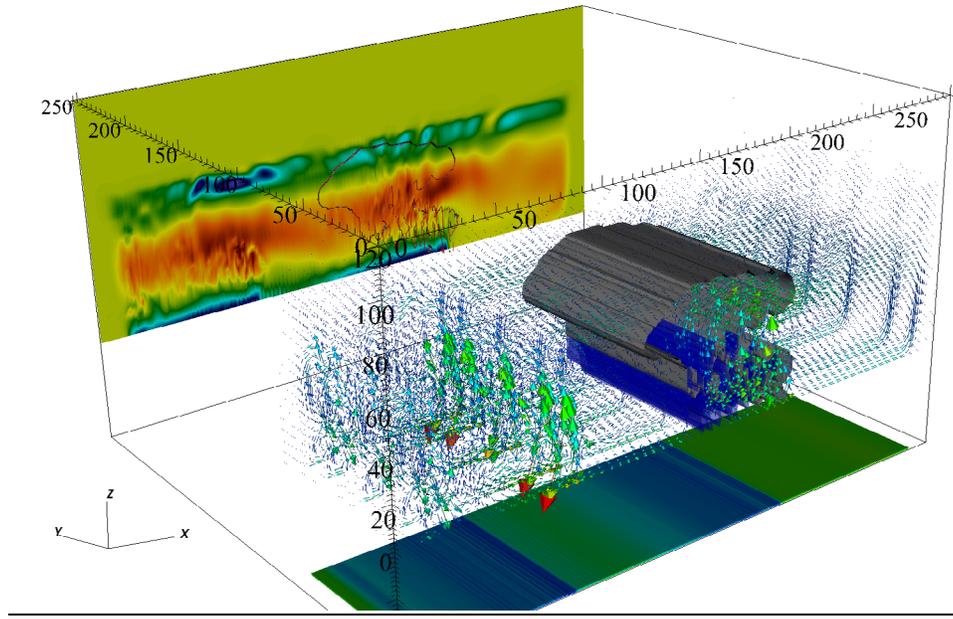


Figure 6: A squall line simulation with the 3D NUMA-CG model. Two cloud trains are present although only the one in the far end is shown. The one nearest the xyz axis is not shown in order to plot velocity vectors. The bottom panel shows the cold pools formed by the density current triggered by the evaporating rain.

Figure 6 shows the results of NUMA3dCG with rain. This 3D version of the model is able to reproduce exactly the 2D results published in Ref. [4]. With moisture now properly working we will now turn to implementing Large-Eddy Simulation and turbulence closure for NUMA in order to begin running hurricane simulations.

In addition to adding moisture to the 3D NUMA-CG code we have also added moisture to the unified NUMA2dCGDG-AMR. This code represents a new version of NUMA that contains unified CG and DG discretization in addition to AMR.

Figure 7 shows the results of, to our knowledge, the first simulation with moisture for a discontinuous Galerkin (DG) model. The significance of this simulation is that it was not well-understood how to handle the duplicate discontinuous degrees of freedom in the DG method within a column in order to simulate precipitation. Although this result is only preliminary, it does show that adding physical parameterizations to a DG model is not as difficult as once believed.

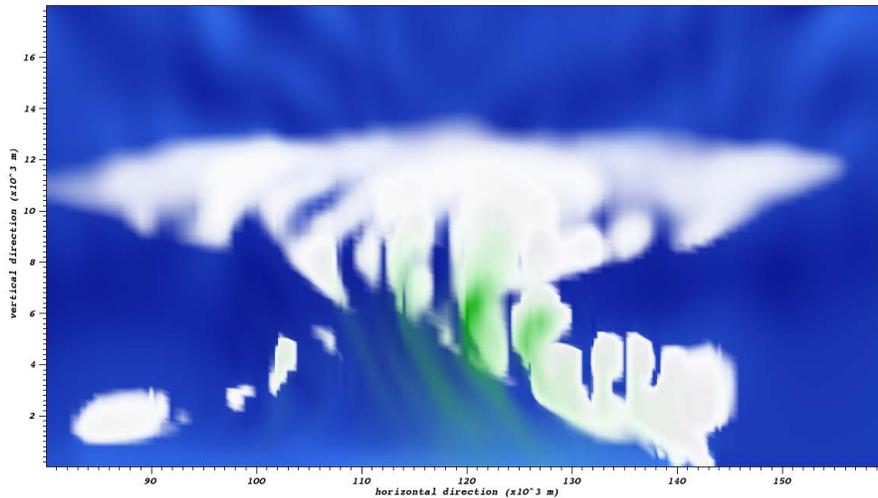


Figure 7: A squall line simulation with the 2D NUMA-DG model. Blue is water vapor, white is condensation, and green is precipitation.

The next phase in this research is to compare and contrast the moist processes simulated by both CG and DG. In addition, we plan to run CG and DG moisture with AMR. With these steps worked out and the addition of turbulence, we are one step closer to running an adaptive high-order simulation of a hurricane.

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 many important quantities such as mass 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. In a separate Department of Energy (DoE) project, the Mathematics and Computational Science group at Argonne National Laboratory is working on interfacing NUMA with their highly scalable software PETSc (Portable Extensive Toolkit for

Scientific Computing). This union will vastly increase the capability of NUMA through the access to the large suite of time-integrators, preconditioners, and solvers contained in PETSc. Furthermore, having NUMA being scrutinized by software and parallel computing scientists will allow us to improve NUMA.

PUBLICATIONS

Journals

1. Giraldo, F.X., Kelly, J.F., and Constantinescu, E.M., 2013: Implicit-Explicit Formulations for a 3D Nonhydrostatic Unified Model of the Atmosphere (NUMA). *SIAM Sci. Comp.* (in press).
2. Mueller, A., Behrens, J., Giraldo, F.X., and Wirth, V., 2013: Adaptive Discontinuous Galerkin Nonhydrostatic Model for Studying Convection. *Journal of Computational Physics*, Vol. 235, 371-393.
3. Carr, L.E., Borges, C.F., and Giraldo, F.X., 2012: An element-based spectrally-optimized approximate inverse preconditioner for the Euler equations. *SIAM J. Sci. Comp.* Vol. 34, B392-B420.
4. Kelly, J.F., Giraldo, F.X., 2012: Continuous and Discontinuous Galerkin Methods for a Scalable 3D Nonhydrostatic Atmospheric Model: limited-area mode. *J. Comp. Phys.*, Vol. 231, 7988–8008.
5. Gabersek, S., Giraldo, F.X., and Doyle, J.D., 2012: Simple Microphysics Experiments with a Spectral Element Model. *Monthly Weather Review* Vol. 140, 3163-3182.
6. S. Marras, J.F. Kelly, F.X. Giraldo, and M. Vazquez, 2012: Variational Multiscale Stabilization and p-adaptivity of High-Order Spectral Elements for the Convection-Diffusion Equation, *J. Comp. Phys.* Vol. 231, 7187-7213.

Plenary Talks

1. Giraldo, F.X.: Continuous and Discontinuous Galerkin Methods for Atmospheric Modeling, European Center for Medium-Range Weather Forecasting, Reading UK (September 2013).

Invited Talks

1. Giraldo, F.X.: Development of the Nonhydrostatic Unified Model of the Atmosphere: challenges in Adaptive Mesh Refinement. Plenary talk at the Newton Institute, Cambridge University (August 20 2012).
2. Giraldo, F.X.: Development of the Nonhydrostatic Unified Model of the Atmosphere (NUMA) on Multi-core and Many-core Architectures, The Met Office, Exeter UK, Weather and climate prediction on next generation supercomputers (October 22-25 2012).
3. Giraldo, F.X.: Challenges in Designing a Multi-purpose Nonhydrostatic Unified Model of the Atmosphere, European Center for Medium-Range Weather Forecasting, Reading UK (November 22 2012).
4. Giraldo, F.X.: The Nonhydrostatic Unified Model of the Atmosphere for Weather Prediction, Lawrence Berkeley Laboratory, Berkeley CA (May 31 2013).

Conference Abstracts

1. Kopera, M.A., Giraldo, F.X.: Non-conforming Adaptive Mesh Refinement for Discontinuous Galerkin Atmospheric Model. *PDES on the Sphere 2012, Cambridge UK* (September 2012).
2. Mueller, A., Giraldo, F.X., Behrens, J., Wirth, V.: Accuracy of Adaptive Discontinuous Galerkin Simulations. *PDES on the Sphere 2012, Cambridge UK* (September 2012).
3. Giraldo, F.X: Derivation of IMEX Formulations for CG and DG Methods. *SIAM Geosciences 2013, Padova Italy* (June 2013).