

NPS-NRL-Rice-UIUC Collaboration on Navy Atmosphere-Ocean Coupled Models on Many-Core Computer Architectures Annual Report

Lucas C. Wilcox
Department of Applied Mathematics
Naval Postgraduate School
833 Dyer Road
Monterey, CA 93943 USA
phone: (831) 656-3249 fax: (831) 656-2355 e-mail: lwilcox@nps.edu

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

Timothy Campbell
Naval Research Laboratory
Oceanography Division, Code 7322
Stennis Space Center, MS 39529 USA
phone: (228) 688-5284 e-mail: tim.campbell@nrlssc.navy.mil

Andreas Klöckner
Department of Computer Science
University of Illinois at Urbana-Champaign
201 North Goodwin Avenue
Urbana, IL 61801 USA
phone: (217) 244-6401 e-mail: andreak@illinois.edu

Timothy Warburton
Department of Computational and Applied Mathematics
Rice University
6100 Main Street, MS 134
Houston, TX 77005 USA
phone: (713) 348-5666 e-mail: tim.warburton@rice.edu

Timothy Whitcomb
Naval Research Laboratory
7 Grace Hopper Ave, MS2
Monterey, CA 93943 USA
phone: (831) 656-4812 fax: (831) 656-4769 email: tim.whitcomb@nrlmry.navy.mil

LONG-TERM GOALS

The ever increasing pace of improvement in the state-of-the-art high performance computing technology promises enhanced capabilities for the next-generation atmospheric models. In this project we primarily consider incorporating state of the art numerical methods and algorithms to enable the Nonhydrostatic Unified Model of the Atmosphere (<http://faculty.nps.edu/fxgiraldo/projects/NUMA>), also known as NUMA, to fully exploit the current and future generations of parallel many-core computers. This includes sharing the tools developed for NUMA (through open-source) with the U.S. community (building on NUOPC and ESPC) that can synergistically move the knowledge of accelerator-based computing to many of the climate, weather, and ocean laboratories around the country.

OBJECTIVES

The objective of this project is threefold. The first objective is to identify the bottlenecks of the NUMA and then circumvent these bottlenecks through the use of: 1) analytical tools to identify the most computationally intensive parts of both the dynamics and physics; 2) intelligent and performance portable use of heterogeneous accelerator-based many-core machines, such as General Purpose Graphics Processing Units (GPGPU or GPU, for short) or Intel's Many Integrated Core (MIC), for the dynamics; and 3) intelligent use of accelerators for the physics. The second objective is to implement Earth System Modeling Framework (ESMF) interfaces for the accelerator-based computational kernels of NUMA allowing the study of coupling many-core based components. We will investigate whether the ESMF data structures can be used to streamline the coupling of models in light of these new computer architectures which require memory access that has to be carefully orchestrated to maximize both cache hits and bus occupancy for out of cache requests. The third objective is to implement NUMA as an ESMF component allowing NUMA to be used as an atmospheric component in a coupled earth system application. A specific outcome of this objective will be a demonstration of a coupled air-ocean-wave-ice system involving NUMA, HYCOM, Wavewatch III, and CICE within the Navy ESPC. The understanding gained through this investigation will have a direct impact on the Navy ESPC that is currently under development. NUMA has already been shown to scale up to tens of thousands of processors on CPU-based distributed-memory platforms (Kelly and Giraldo 2012). This scalability has been achieved through the use of the Message Passing Interface (MPI) to exchange data between processors. The work planned here will further increase the performance of NUMA especially for the most costly operations that are currently taking place on-processor. Examples of such operations include the right-hand-side (RHS) vectors formed by the continuous/discontinuous Galerkin (CG/DG) high-order spatial operators, the implicit time integration strategy, and the sub-grid scale physics.

APPROACH

Following the lead of various DoE labs (Swaminarayan 2011), we will adapt NUMA to accelerator-based many-core machines in a step-by-step process. At each step we will develop *mini-apps* which are self-contained programs that capture the essential performance characteristics of different algorithms in NUMA. This plan to partition the development of the heterogeneous architecture version of NUMA into small chunks of work that can be handled somewhat independently will allow us to produce (at every

stage of the work pipeline) a result that is beneficial not only to the NUMA developers and user groups but also to the larger climate, weather, and ocean community. The many-core mini-apps that will be developed will include:

Dynamics

Explicit-in-time CG a continuous Galerkin discretization of the compressible Euler mini-app with explicit time integration;

Explicit-in-time DG a discontinuous Galerkin discretization of the compressible Euler mini-app with explicit time integration;

Vertically Semi-Implicit CG a continuous Galerkin discretization of the compressible Euler mini-app with vertically implicit semi-implicit time integration;

Vertically Semi-Implicit DG a discontinuous Galerkin discretization of the compressible Euler mini-app with vertically implicit semi-implicit time integration;

Physics

Moisture a Kessler parameterized moisture mini-app; and

Long-Wave Radiation a radiative transfer for inhomogeneous atmospheres (using for example the rapid radiation transfer scheme (Mlawer et al. 1997)) based mini-app.

Once the performance of a mini-app is accepted it will be considered for adoption into NUMA. We will also make these mini-apps available to the community to be imported into other codes. Warburton's group is developing a library, *occa*, that allows a single kernel to be compiled using many different threading frameworks, such as CUDA, OpenCL, OpenMP, and OpenACC. We will plan on using plain OpenCL or *occa* for the computational kernels in the mini-apps. The choice will be made weighing portability with usability. Parallel communication between devices will use the MPI standard to enable the mini-apps to run on large scale clusters. Using these community standards for parallel programming will allow our mini-apps to be portable to many platforms, however the performance may not be portable across devices. For performance portability, we will use *Loo.py* to develop OpenCL kernels which can be automatically tuned for current many-core devices along with future ones.

The second goal is to implement Earth System Modeling Framework (ESMF) interfaces for the accelerator-based computational kernels of NUMA allowing the study of coupling many-core based components. We will investigate whether the ESMF data structures can be used to streamline the coupling of models in light of these new computer architectures which require memory access that has to be carefully orchestrated to maximize both cache hits and bus occupancy for out of cache requests. We will coordinate with the "An Integration and Evaluation Framework for ESPC Coupled Models" team to develop and test ESMF based mini-apps within the proposed ESPC Coupling Testbed.

The third goal is to implement NUMA as an ESMF component allowing NUMA to be used as an atmospheric component in a coupled earth system application. A specific outcome of this goal will be a demonstration of a coupled air-ocean-wave-ice system involving NUMA, HYCOM, Wavewatch III, and CICE within the Navy ESPC. Optimized versions of HYCOM, Wavewatch III, and CICE will be obtained from the "Accelerated Prediction of the Polar Ice and Global Ocean" team. The understanding gained through this investigation will have a direct impact on the Navy ESPC that is currently under development.

Activity	Months:	0–12	13–24	25–36	37–48	49–60
Identifying Bottlenecks		•				
Explicit Dynamics		•	•			
Moisture		•	•			
Vertically-Implicit Dynamics			•	•	•	
More Physics Processes					•	•
Implement ESMF in mini-apps				•	•	•
Implement ESMF in NUMA		•	•			
Port many-core kernels into NUMA			•	•	•	•
Adapt Loo.py		•	•	•	•	
Develop source translation tool			•	•	•	•
Assess Performance		•	•	•	•	•
Create Working Group		•				
Disseminate work (Publications)			•	•	•	•

Table 1: Time-line of proposed activities for different months of the project

WORK COMPLETED

In the course of this project (the first three years plus the two year option), we plan to carry out the following work items:

1. identify current bottlenecks in the NUMA modeling system;
2. port the explicit time-integration portion of the dynamics onto many-core devices;
3. port the moisture schemes onto many-core devices;
4. port the implicit-in-the-vertical dynamics onto many-core devices;
5. port long-wave radiation and other costly physics onto many-core devices;
6. implement ESMF interfaces for many-core components;
7. implement NUMA as ESMF component;
8. transition many-core kernels into NUMA;
9. adapt the Loo.py code generator for the needs of the project;
10. develop a source-to-source translation capability built on Loo.py to facilitate the NUMA transition;
11. assess performance against current modeling suite;
12. foster the formation of a community working group on using accelerators for atmosphere-ocean modeling.

The management plan for these work items is shown in Table 1.

In preparation for the work to start we have done a couple of things. At the Naval Postgraduate School, Giraldo and Wilcox are in the process of hiring a postdoctoral student. We have selected a candidate and are awaiting his application to the National Research Council (NRC) Research Associateship Program.

Warburton and Wilcox are planning on proposing a minisymposium at the International Conference on Spectral and High Order Methods 2014 (<http://www.icosahom2014.org/>) to bring together senior and junior international researchers from the national labs, academia, and industry who are actively engaged in the development of high performance algorithms for high-order PDE discretizations on many-core architectures. This will be a venue to disseminate exchange ideas about running numerical codes like NUMA on many-core architectures.

At the kick-off meeting of the NOPP “Advancing Air-Ocean-Land-Ice Global Coupled Prediction on Emerging Computational Architectures” effort we will coordinate with the other projects to ensure that all efforts (whenever possible) will be in sync to avoid duplication. In addition to the NOPP kick-off meeting we are planning a kick-off meeting in Monterey for the NPS-NRL-Rice-UIUC collaboration and are currently in the process of determining the date.

RESULTS

As the work has not started there are no results to report.

IMPACT/APPLICATIONS

Ensuring that the U.S. gains and maintains a strategic advantage in medium-range weather forecasting requires pooling knowledge from across the disparate U.S. government agencies currently involved in both climate and weather prediction modeling. The new computer architectures currently coming into maturity have leveled the playing field because only those that embrace this technology and fully commit to harnessing its power will be able to push the frontiers of atmosphere-ocean modeling beyond its current state. The work in this project is critical to developing and distributing the knowledge of accelerator-based computing that will support the use of the new platforms in many of the climate, weather, and ocean laboratories around the country.

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

The Earth System Modeling Framework (ESMF) together with the NUOPC Interoperability Layer form the backbone of the Navy ESPC software coupling infrastructure. We will enable the many-core miniapps and NUMA to be used as components in the Navy ESPC by implementing them as a NUOPC compliant ESMF components. This will bring our work the ESPC community enabling coupling to codes from other projects such as HYCOM and Wavewatch III.

REFERENCES

Kelly, James F. and Francis X. Giraldo (2012). “Continuous and discontinuous Galerkin methods for a scalable three-dimensional nonhydrostatic atmospheric model: Limited-area mode”. English. In: *Journal of Computational Physics* 231.24, 7988–8008. DOI: [10.1016/j.jcp.2012.04.042](https://doi.org/10.1016/j.jcp.2012.04.042).

Mlawer, Eli J. et al. (1997). “Radiative transfer for inhomogeneous atmospheres: RRTM, a validated correlated-k model for the longwave”. In: *Journal of Geophysical Research: Atmospheres* 102.D14, pp. 16663–16682. DOI: [10.1029/97JD00237](https://doi.org/10.1029/97JD00237).

Swaminarayan, Sriram (2011). *Exaflops, Petabytes, and Gigathreads... Oh my!* DoE Advanced Scientific Computing Research (ASCR) 2011 Workshop on Exascale Programming Challenges. URL: http://science.energy.gov/~media/ascr/pdf/research/cs/Programming_Challenges_Workshop/Siriam%20Swaminarayan.pdf.