

ESPC Computational Efficiency of Earth System Models

Timothy R. Whitcomb
Marine Meteorology Division, Code 7532
Naval Research Laboratory
Monterey, CA 93943
phone: (831) 656-4812 fax: (831) 656-4769 email: timothy.whitcomb@nrlmry.navy.mil

Alan J. Wallcraft
Oceanography Division, Code 7323
Naval Research Laboratory
Stennis Space Center, MS 39571
phone: (228) 688-4813 fax: (228) 688-4759 email: alan.wallcraft@nrlssc.navy.mil

Award Number: N0001414WX20355
<http://www.nrlmry.navy.mil>

LONG-TERM GOALS

Successful operational implementation of the global ESPC fully coupled system in 2018 will require that the coupled system and the constituent models are able to run efficiently and in a timely manner on Navy operational computer systems to ensure that products produced by the systems are available for fleet user consumption and downstream dependencies. This product will analyze and modify the atmosphere model (NAVGEM) and the ocean model (HYCOM) to ensure that the ESPC coupled system is able to take advantage of modern computational platforms and increase computational efficiency and scalability.

OBJECTIVES

Instrument, analyze, and modify the Navy's global atmosphere model (NAVGEM), global ocean model (HYCOM), and global atmosphere data assimilation system (NAVDAS-AR) to increase the scalability of the component systems and thereby improve the computational efficiency of the ESPC system as a whole.

APPROACH

Greater efficiency of a coupled system begins with identifying opportunities for improvement. This will require enhanced instrumentation and analysis of model and coupling components to identify the particular sections of code with the largest impact on performance. This instrumentation will show both communication patterns (typically within component models) and places in the code where the model spends a lot of time.

Once identified, we will use an incremental approach to improve the efficiency of those sections. Typically these enhancements will be in the form of adjusting the control flow, altering

communication patterns, changing data structures, or taking advantage of improved algorithms that make better use of modern architectures.

As each improvement is complete, the full system will be re-instrumented and re-evaluated to ensure that the changes made provided improved efficiency and did not alter the results.

For FY14, the focus was on the atmosphere model alone: key performers for that system are Timothy Whitcomb (NRLMRY, focus on global modeling) along with Steven Lowder (NRLMRY/SAIC, computer scientist focused on I/O framework).

WORK COMPLETED

- Identify external contractor for NAVGEM instrumentation and optimization, established memorandum of understanding and statement of work
- Initial full-model scaling test for high-resolution NAVGEM
- Begin analysis and refactoring of NAVGEM I/O

RESULTS

As a first step in evaluating model scalability, we tested a full-model scaling (i.e. with no features disabled) with a high-resolution NAVGEM integration on the NRL Cray XE6m supercomputer. Traditional scaling tests are focused purely on computational limitations (and so disable things like I/O) but our approach was to initially assess model performance as it would be run in an operational system. Results of this test are shown in Figure 1 – the key is the increasingly poor scaling of the model as the core count increases. A large portion of this growing discrepancy is from the communication required to write out the model state so we began our efforts focusing on the I/O portion of the model. There are definite opportunities for optimization in this system.




Figure 1 – Plot showing seconds per forecast day wallclock time for a T639L64 (~21 km at the equator) NAVGEM integration. Dashed slanted lines show “perfect” scaling (i.e. double the number of processors, cut wallclock time in half). The growing discrepancy between ideal scaling and actual scaling includes contributions from computational limitations (e.g. the spherical harmonic transforms) as well as system limitations, primarily input/output.

One persistent bottleneck in the NAVGEM is the I/O subsystem for handling input output of model history files, which are used for restarts, data assimilation, and post-processing for output and downstream products. The current architecture for output uses MPI gather operations to bring data from each core to a single process that uses unformatted Fortran I/O calls to write a custom binary formatted file while all other processors wait. This simple setup works well for lower model resolutions and systems with fast disks, but represents serious performance issues with filesystems like Lustre that are geared toward parallel file operations. The percentage of model runtime for I/O can reach over 30% on some platforms which poses a serious impediment to scalability.

We analyzed the file format used by NAVGEM and began refactoring into routines that will be shared in a library between the data assimilation system and forecast model. This understanding of the storage order of spectral coefficients and how they are distributed across processors will allow for modification of the backend storage once the I/O library is completely extracted from the model. This refactoring will allow other output options to better take advantage of parallel I/O (such as MPIIO or

parallel HDF5) as well as future capabilities of asynchronous ESMF components for performing input and output.

The use of MPIIO for NAVGEM in a draft implementation developed under a previous project showed significant improvement in performance and we have begun to incorporate that draft implementation into what will become the standard I/O library for NAVGEM while attempting to maintain backward compatibility. We have two candidate solutions that allow for no format change (a shortcoming of the draft implementation) that are currently being tested. The understanding of the particular file layouts and ongoing abstraction of the I/O layer are critical pieces to addressing the negative influences to scalability shown in Figure 1.

IMPACT/APPLICATIONS

The future impact of this project is to ensure operational efficiency of the Naval operational environmental prediction system that is targeted for IOC in 2018.

RELATED PROJECTS

This work is part of a larger ESPC project. Other related projects are 6.2 NOPP: Accelerated Prediction of the Polar Ice and Global Ocean (APPIGO) with a strong focus on accelerators for the ocean and ice model, as well as extensions to and validation of the atmosphere model through 6.4 NAVGEM and extensions to and validation of the ocean model through 6.4 Large Scale Ocean modeling. A funded PETTT proposal for ESMF asynchronous I/O component will be leveraged in future years of this project.