Accelerated Prediction of the Polar Ice and Global Ocean (APPIGO)

Eric Chassignet
Center for Ocean-Atmosphere Prediction Studies (COAPS)
Florida State University, PO Box 3062840
Tallahassee, FL 32306-2840
phone: (850) 645-7288 fax: (850) 644-4841 email: echassignet@fsu.edu

Award Number: N00014-13-1-0861
https://www.earthsystemcog.org/projects/espc-appigo/

Other investigators:

Phil Jones (co-PI)	Los Alamos National Laboratory (LANL)
Rob Aulwes	Los Alamos National Laboratory
Tim Campbell	Naval Research Lab, Stennis Space Center (NRL-SSC)
Mohamed Iskandarani	University of Miami
Elizabeth Hunke	Los Alamos National Laboratory
Ben Kirtman	University of Miami
Alan Wallcraft	Naval Research Lab, Stennis Space Center

LONG-TERM GOALS

Arctic change and reductions in sea ice are impacting Arctic communities and are leading to increased commercial activity in the Arctic. Improved forecasts will be needed at a variety of timescales to support Arctic operations and infrastructure decisions. Increased resolution and ensemble forecasts will require significant computational capability. At the same time, high performance computing architectures are changing in response to power and cooling limitations, adding more cores per chip and using Graphics Processing Units (GPUs) as computational accelerators. This project will improve Arctic forecast capability by modifying component models to better utilize new computational architectures. Specifically, we will focus on the Los Alamos Sea Ice Model (CICE), the HYbrid Coordinate Ocean Model (HYCOM) and the Wavewatch III models and optimize each model on both GPU-accelerated and MIC-based architectures. These codes form the ocean and sea ice components of the Navy's Arctic Cap Nowcast/Forecast System (ACNFS) and the Navy Global Ocean Forecasting System (GOFS), with the latter scheduled to include a coupled Wavewatch III by 2016. This work will contribute to improved Arctic forecasts and the Arctic ice prediction demonstration project for the Earth System Prediction Capability (ESPC).

OBJECTIVES

The objective of this effort is to create versions of the Los Alamos Sea Ice Model (CICE), the HYbrid Coordinate Ocean Model (HYCOM) and the Wavewatch III models that can perform optimally on both GPU-accelerated and MIC-based computer architectures. These codes form the ocean and sea ice components of the Navy's Arctic Cap Nowcast/Forecast System (ACNFS) and the Navy Global Ocean Forecasting System (GOFS), with the latter scheduled to include a coupled Wavewatch III by 2016.

This work will contribute to improved Arctic forecasts and the Arctic ice prediction demonstration project for the Earth System Prediction Capability (ESPC).

APPROACH

We will utilize an incremental acceleration approach to ensure we maintain code fidelity while improving performance. We will begin by improving the performance of selected sections of each code and expanding those regions until we have accelerated the three application codes. Acceleration may start with directive-based mechanisms like OpenACC and OpenMP, but may also include targeted kernels written in CUDA or other lower-level accelerator libraries. This approach provides early successes and opportunities to test the changes as they are made. A second approach will redesign code infrastructure to incorporate a multi-level parallelism by design. The modified codes will be validated both on a single component basis and within the forecast systems.

WORK COMPLETED

While we continued attempts at GPU acceleration, this year included a reassessment of progress and the start of a change in focus and target architectures for the project. In particular, given the difficulties in obtaining significant performance increases on GPU-accelerated architectures, we have decided to transition to activities that are needed for performance on many-core (Knights Landing or KNL) and that will likely also improve performance on GPU architectures. These include an increased focus on threading and vectorization. We have continued science simulations to prepare a validation test for the improved model using a fully coupled configuration in collaboration with other related projects.

CICE Performance Improvements

During the year, we explored a number of performance improvements for the CICE model. We continued work on GPU acceleration of the model by porting the transport and dynamics to the GPU. Previously, we had attempted to use OpenACC directives and were able to match CPU performance, but had difficulty improving performance. OpenACC provides an easier-to-use programming model, but has limitations that impedes performance tuning. In particular, since OpenACC is a directive-based programming model, additional code is injected over which developers do not have control. We noticed that synchronization calls were added even for code we identified as asynchronous. With these considerations, we decided to explore using CUDA Fortran and investigate whether we could gain better performance, since using CUDA would allow for finer control over GPU code. Both the transport (horizontal remapping advection) and the elastic-viscous-plastic rheology (evp) computations were re-written in CUDA Fortran. Unfortunately, even these efforts resulted in parity (remapping) or worse (evp) performance relative to CPU-only code.

After discussions amongst the team, we decided to discontinue our efforts using GPU acceleration and instead focus on many-core architectures using threading. With the installation of Trinity Phase 2 with the advanced Xeon Phi Knights Landing (KNL) architecture, we concurred that we should focus on the Xeon Phi architecture and improving CICE threading and vector performance. We have built CICE on LANL test cluster of KNL, Trinitite and are in the process of fixing some threading bugs and gathering performance profiling. Our first step is identifying code sections that are not vectorizing; vectorization will be critical to KNL performance.

For any architecture, one of the largest performance bottlenecks in the CICE model are the halo exchanges that communicate domain boundary information between domains distributed across nodes.

During the year, we restructured the halo exchanges in an attempt to overlap communication and aggregate messages. The new update method works for both GPUDirect and regular CPU updates. The original halo updates are synchronous calls for each physical field being updated. The new scheme breaks the update into multiple steps in an attempt to make the updates asynchronous and allow computation to continue while the updates occur. We replaced the halo updates in the remapping transport routines with the new update and tested it with the p4 test problem using 30 MPI ranks. Disappointingly, we did not see performance improvements. We still believe this new approach will offer improvements with halo updates, and we will expand it to CPU halo updates in addition to GPU halo updates. We will also experiment with the MVAPICH MPI implementation that also supports GPUDirect and check if it gives better GPU performance than OpenMPI.

Task parallelism may provide another means for exposing more parallelism in CICE as well as mapping tasks to architectural elements. We started exploration of task parallelism, generating an initial dataflow analysis. Our strategy is to identify tasks that could be run concurrently, not only on the GPU, but also on the CPU. Instead of porting all code to the GPU, we are exploring schemes to use CPU resources while kernels are executing on the GPUs. We have begun implementing this strategy with the `horizontal_remap` routine, which was shown as the top hotspot in the performance profiler. Intial results showed our CUDA version is within 3% of the normal CPU code for the tp1 test problem using 80 MPI ranks. The test problem was run on Nvidia's PSG cluster using four nodes with each node containing six Kepler K80 GPUs.

HYCOM/Wavewatch

The proposed work plan had a switch to wave model work in the final year of the base 3-year plan, so performance work on HYCOM has paused to examine potential improvements to Wavewatch.

HYCOM in CESM

In order to validate changes in the model, we are configuring and running HYCOM in CESM in several different configurations. Initial work was done using two low-resolution grids: the POP gx1v6 (~1º) grid, that was already available in the CESM set-up, and the HYCOM glbt0.72 (~0.72º) tripolar grid that has been especially implemented in the CESM set-up.

As a first step, the CESM passive ice component DICE was added to the standalone HYCOM. This new option allows HYCOM to evolve with a prescribed ice cover derived from SSMR/SSMI NSIDC climatology (Cavalieri et al., 1997). In addition to those changes, options to use a spatially varying sea surface salinity and/or temperature relaxation as well as a correction of the precipitation based on the global salinity have been introduced to comply with the POP simulation parameters used for comparison. Several 30-year runs of HYCOM have been performed with the CORE normal year atmospheric fields (Large and Yeager, 2009) to assess the sensitivity of the model to several parameters (reference pressure, thermobaricity, and isopycnal smoothing) and find the optimal parameters for HYCOM.

As second step, HYCOM was fully implemented in CESM as an alternative ocean component to POP. Starting with version 2.2.35 of HYCOM, we first focused on the proper implementation of the routine responsible for the exchange of fluxes between HYCOM and the ice and atmospheric components. We then updated HYCOM to the latest version available (2.2.98) and connected the river transport component. Also, we implemented the latest HYCOM-CESM together with the latest version of ESMF (7.0.0beta60) on the NAVY DSRC machines, Kilrain (IBM iDataPlex), Shepard (Cray XC30) and

Gordon (Cray XC40). A 20-year experiment forced the CORE normal year climatological fields was performed to ensure that HYCOM-CESM performed as expected when coupled with CICE only (G compset). When comparing HYCOM-CESM to POP-CESM simulations, we find that differences arisedover the polar regions (Fig. 1). While the ice cover are in good agreement between the two models, HYCOM-CESM presents a thicker ice sheet on the Beaufort gyre than POP-CESM.

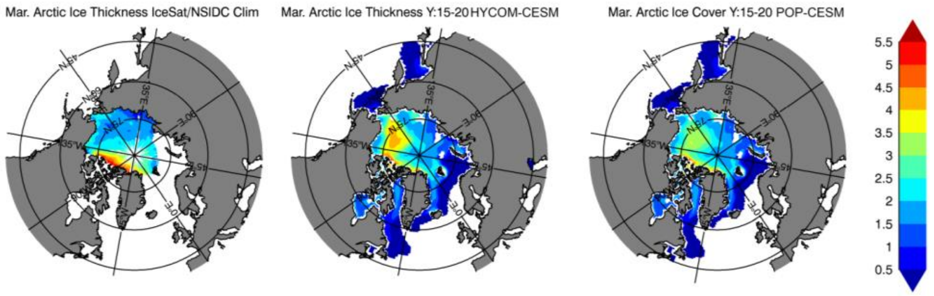

Fig. 1: Ice thickness climatology from IceSat (left) and averaged over the last 5 years of simulation for HYCOM-CESM (middle) and POP-CESM (right).

After performing sensitivity tests on different CICE constants and parameterizations that were set differently between the standalone HYCOM-CICE and HYCOM-CESM (i.e. snowpatch, ustar_scale, shortwave parameterizations), it was established that HYCOM-CESM produced more ice than POP-CESM when using the Delta-Eddington shortwave parameterization under the CORE normal year forcing. To ensure that these differences did not come from a technical mistake in the coupling process, a standalone HYCOM-CICE simulation was performed and compared to a HYCOM-CESM simulation using the same set of CICE parameters in both. Results presented almost identical ice thickness between HYCOM-CICE and HYCOM-CESM (not shown). Then, in order to properly compared both ocean models in more realistic settings, a simulation of 60 years was run using the atmospheric CORE interannual forcing for each ocean model within the CESM (GIAF compset). Unlike with the CORE normal year forcing, the ice thickness averaged over the last 10 years of simulation in HYCOM_CESM was very similar to the ice thickness in POP-CESM (Fig. 3) with even slightly less ice in the middle of the Beaufort Gyre in HYCOM-CESM compared to POP-CESM. The ice extent was however slightly higher in HYCOM-CESM, especially in the Labrador Sea Region and the Nordic Sea (Fig. 2).

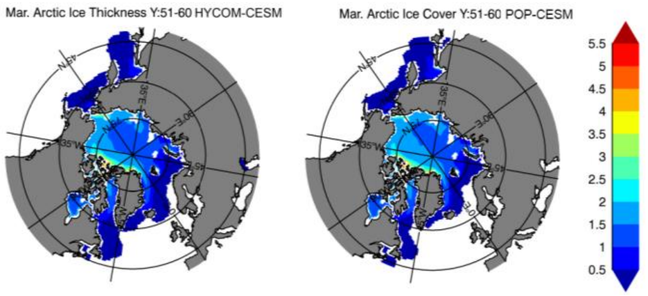

Fig. 2: Ice thickness averaged over the last 10 years for HYCOM-CESM (left) and POP-CESM (right).

On the ocean part, HYCOM-CESM's global temperature increases by 0.1ºC in 60 years whereas POP-CESM's global temperature stays almost constant over the whole time period (not shown). At the surface, HYCOM-CESM presents a saltier bias compared with POP-CESM but a similar bias in SST except in the North Pacific where HYCOM-CESM is warmer and the North Atlantic subpolar gyre where POP-CESM is warmer (Fig. 3). The maximum of Atlantic Meridional Overturning Circulation

(AMOC) at 26.5ºN averaged over the last 10 years of simulations is 17.5Sv for POP-CESM and 9Sv for HYCOM-CESM, in agreement with the CORE-II HYCOM experiments (Danabasoglu et al., 2015) (not shown). In parallel to this study with the gx1v6 POP grid, the HYCOM GLBt0.72 tripolar grid (500x382) was also implemented and tested within CESM and showed overall comparable results with the gx1v6 grid HYCOM simulations.

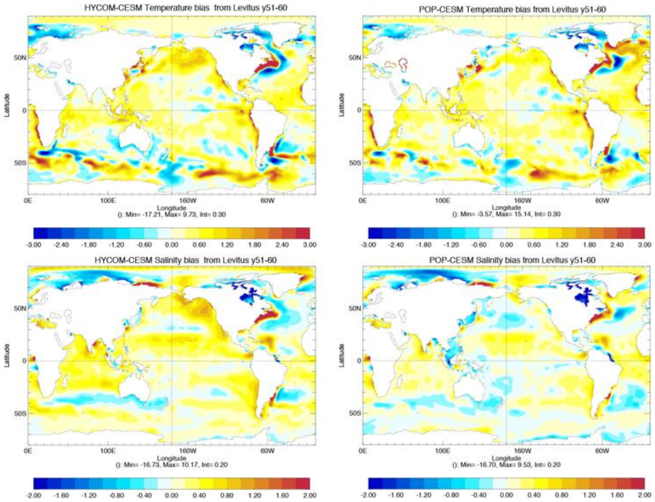

Fig. 3: SST (top) and SSS (bottom) bias from Levitus averaged over the last 10 years for HYCOM-CESM (left) and POP-CESM (right).

After validating HYCOM in an ocean-ice configuration, we are now moving into evaluating HYCOM into an ocean-ice-atmosphere configuration. With CAM as an atmospheric model on a 1.9ºx2.5º grid, the behaviors of HYCOM on a gx1v6 grid and a GLBt0.72 grid are compared with the behavior of POP on a gx1v6 grid. First results show a larger extent of the ice cover in the HYCOM simulations than in POP-CESM especially in the Labrador Sea (not showed). Consequently, HYCOM presents a colder and fresher bias over the North Atlantic subpolar region (not showed). Further analysis is needed to fully understand the behavior of HYCOM in a full coupled setting.

At the same time, a high-resolution configuration using the 1/10º tx01v2 POP tripolar grid and a 0.47ºx0.63º atmospheric grid has been set-up to be run with HYCOM. The physical parameters from a stand-alone 1/12º global HYCOM configuration have been used to test this new HYCOM-CICE-CAM configuration. The goals of this exercise are to compare 1) the HYCOM and POP ocean model in a high resolution framework, 2) the HYCOM behavior running on the 1/10º tripolar POP grid vs. running on the 1/12º HYCOM grid and finally 3) the HYCOM 1/12º behavior in CESM vs. in NAVGEM. While all the work involving the 1/12º grid will be done by COAPS (Alexandra Bozec and Eric Chassignet), the configuration involving the 1/10º tripolar POP grid will be prepared by COAPS but the simulations will be run by Ben Kirtman.

ESPC-HYCOM

A low-resolution configuration using a roughly 1º (T119) NAVGEM atmospheric grid and a 0.72º (GLBt0.72) ocean grid has been added to the sample runs of the ESPC system to allow for easier testing. This configuration will be run for 50 years and compared with the CESM-HYCOM using the same ocean grid. For a better comparison, the ESPC ice-ocean coupling now also allow for a full flux exchange instead of the usual partial coupling of the NRL. In addition, an automatic resubmission script has been added to the system to be able to run continuously month per month in long-term simulations.

Several performance tests have been done with output frequency and configuration similar to the CESM-HYCOM above (i.e. ocean IO every day, ice-atm IO every month and 32 layers instead of 41 layers for HYCOM). Results show that for a 900procs (ocn/ice) + 96procs (atm) configuration, we achieve 1 simulated day in 25-26 CPU minutes, and for a 1800procs + 96procs configuration, we achieve 1 simulated day in 19-20 min with a time step of 120 s/5 s. These results compare to 1 simulated day in 13-14 min for CESM-HYCOM with the same time step.

RESULTS

Performance improvements using GPU acceleration have been elusive and effort has now moved toward threading and vectorization for KNL architectures.

IMPACT/APPLICATIONS

Model performance improvements under this project will result in high-performance codes to enable improved future Arctic prediction, through improved resolution, increased realism or an ability to run ensembles.

RELATED PROJECTS

This project builds on the core model development activities taking place at the partner sites, including:

The Climate, Ocean and Sea Ice Modeling (COSIM) project that includes the primary development of the Los Alamos Sea Ice Model (CICE), funded by the US Department of Energy’s Office of Science.

The ongoing development of the Arctic Cap Nowcast-Forecast System (ACNFS) and Global Ocean Forecast System (GOFS) at the Naval Research Lab – Stennis, funded by the US Navy.

Continued development of the Hybrid Coordinate Ocean Model (HYCOM) at Florida State University, funded by the National Science Foundation, Department of Energy and US Navy.