OOKAMI PROJECT APPLICATION

Date: June 15, 2021

Project Title: Vertical Structure of Oceanic Mesoscale Mixing

Usage:

☒ Testbed

☐ Production

Principal Investigator:

University/Company/Institute: Stony Brook University

Mailing address including country: Z=5000 (USA)

Phone number: 631-632-3152

Email: christopher.wolfe@stonybrook.edu

Names & Email of initial project users:

Christopher Wolfe (christopher.wolfe@stonybrook.edu)
Wenda Zhang (wenda.zhang@stonybrook.edu)

Usage Description:

The scientific goal of this project is to validate, invalidate, and/or constrain theories for the vertical structure of oceanic mesoscale tracer diffusivity using the diffusivity determined from a pseudo-inversion method as the "ground truth" for comparison to theory. The pseudo-inversion method obtains the diffusivity from the fluxes and gradients of a large number of numerical tracers advected by a high-resolution ocean circulation model. The ocean circulation model is a configuration of the MIT general circulation model (MITgcm: mitgcm.org) which achieves parallelism via domain decomposition of the horizontal domain into tiles. The tiles themselves have relatively few grid points, so each process typically has a fairly modest memory footprint. The tiles primarily communicate through one-to-one exchanges of edge values. The equation for the ocean's surface is solved implicitly to avoid having to take very small time steps to resolve surface waves—the resulting elliptic problem is solved via conjugate gradients (CG), necessitating global sums. At high resolution, the CG solver begins to dominate the time-to-solution for the flow itself. Adding many tracers increases number of computations-per-tile as well as the amount of edge data that must be exchanged. Under these conditions, tracer advection dominates the time-to-solution rather than the CG solver.

Computational Resources

We will to test the performance of our configuration both with and without advected tracers—these will be referred to as the **no tracer** and **tracer** configurations, respectively.

The **no tracer** experiment will be run for approx. 500 simulated years to generate a stable circulation to use as input for the **tracer** experiment which will run for approx. 20 simulated years. Computational requirements were estimated from performance on *Seawulf* assuming a 3.5 times speed-up factor between *Seawulf* and *Ookami*.

Total node hours per year: approx. 5100

Size (nodes) and duration (hours) for a typical batch job:
Either 8 nodes for 2 days (long queue) or 24 nodes for 8 hours (large queue) depending which configuration is mode efficient.

Disk space (home, project, scratch):
Home: a few GB for compiling and staging
Scratch: < 1 TB for temporary files
Project: 2 TB for sharing model output between project members, post-processing, and medium-term (~year) archival.

Personnel Resources (assistance in porting/tuning, or training for your users**):**

The PI is well-aquatinted with porting and compiling the MITgcm, which depends only on the presence of standard MPI and OpenMP libraries. On the other hand, it has been some time since the PI has optimized code for vector machines. It is therefore difficult to make precise predictions of personnel requirements. Informal channels (KB, Slack channel, webinars, etc.) will probably suffice.

Required software:

The only software required are modern Fortran and C compilers as well as MPI and OpenMP libraries. Access to a Scientific Python stack would be helpful for pre- and post-processing, but these can also be done offline at lower convenience.

If your research is supported by US federal agencies:

Agency: NSF

Grant number(s): OCE-2048826