The Modular Ocean Model (MOM6)

Youwei Ma
Stony Brook University
03/23/2023

Tropical cyclone (TC) and climate interaction

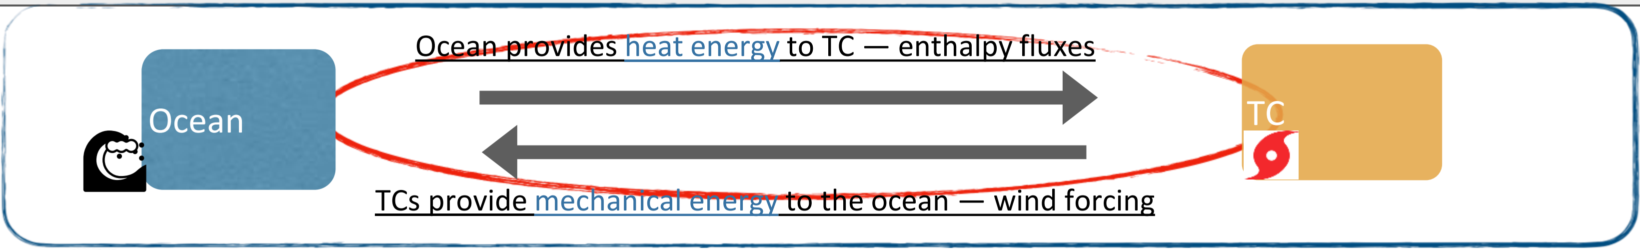

From Li's slides presented on CESM tutorial 2022

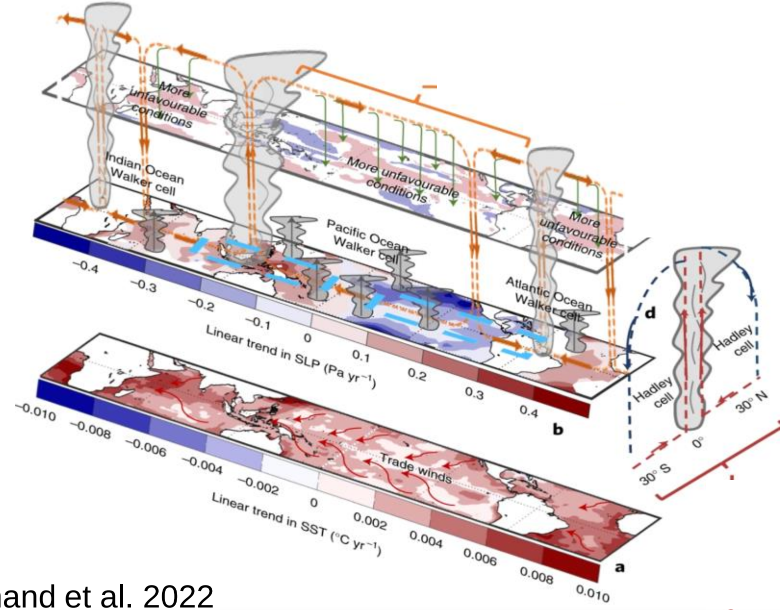

Chand et al. 2022

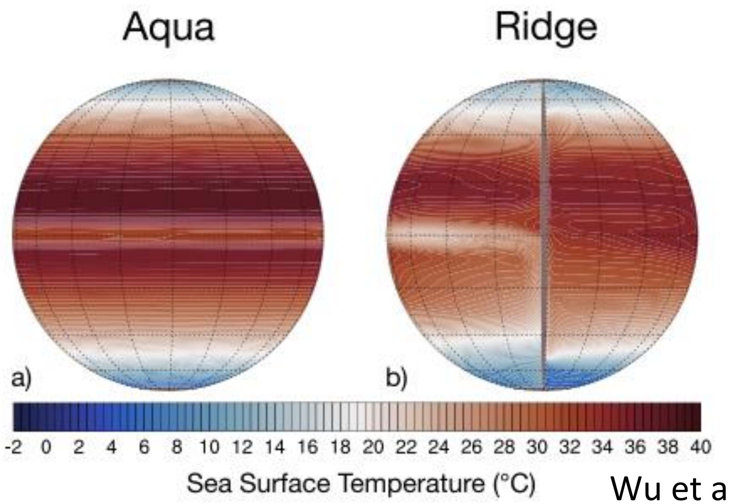

Wu et al. 2021

The Modular Ocean Model (MOM6)

- MOM6 is a numerical model simulating fluid properties and circulations based on the Navier–Stokes equations on the rotating sphere with thermodynamic terms.
- MOM6 is the ocean component in Community Earth System Model (CESM).
- Fortran based model.

7 equations and 7 unknowns:

3 velocity components; Potential temperature;

Salinity; Density;

Pressure.

Plus: 1 equation for each passive tracer, e.g. CFCs, Ideal Age.

Features of MOM6

Arakawa C-grid

Fig 1: Arakawa C-grid of variables around an h-cell with North-East indexing convention

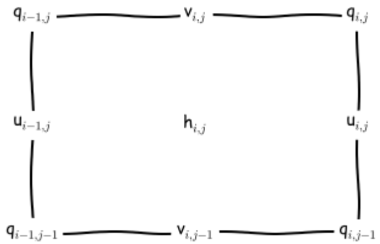

Advantages:

- Allow single-point channels
- Good for conservation

Disadvantages:

- Less accurate Coriolis term, inducing less accurate inertia gravity waves and Rossby waves

Vertical Lagrangian remapping: Arbitrary Lagrangian Eulerian (ALE) algorithm

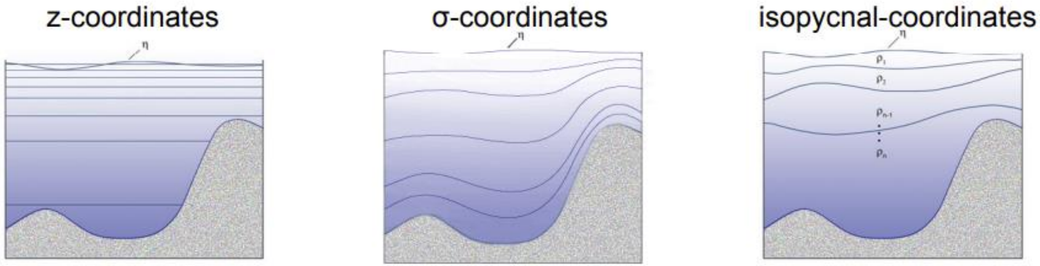

From: https://www.oc.nps.edu/nom/modeling/vertical_grids.html

- Enable the use of any vertical coordinate
- Removes the vertical advection CFL restriction on the time-step so that the model is unconditionally stable to thin (or even vanishing) layers.
- Good for representing the evolution of ice shelf grounding lines as well as coastal/tidal estuaries

Compiler Flags

Compiler+parallel implementation	Flags
GCC v12.1.0+Open MPI v4.1.4	-Ofast, -mcpu=a64fx,-fopenmp, -fallow-invalid-boz,-fallow-argument-mismatch
Arm v22.1+Open MPI v4.1.4	-Ofast, -mcpu=a64fx -armpl, -fopenmp, FPPFLAGS += -D'rank(X)=size(shape(X))
Cray v22.03+MVAPICH2 v2.3.5	-O3, -h vector3,-h omp
Fujitsu v4.8	-Kfast, -KSVE,-Kopenmp, -CcdRR8
Intel 2022.1 (Intel Broadwell) processors on Cheyenne)	-O3,-openmp

Reference makefile templates:
MOM6-examples/src/mkmf/templates/linux-GNU.mk
MOM6-examples/src/mkmf/templates/ncrc-cray.mk

MOM6 test case

MOM6 configuration for tests:

dev/gfdl MOM6-examples/ocean_only/global

- 1-degree global ocean

  Number of grids =360(longitude)x210(latitude)x63(depth) =4,762,800

  Depth 0.5m-6000m

- 30-model-day run

- Surface oundary conditions:
  prescribes the daily surface forcing including radiation flux, freshwater flux, and wind stress
- Bottom boundary conditions:
  realistic global bathymetry
- Initial conditions:
  realistic temperature and salinity.

Performance of pure-MPI tasks

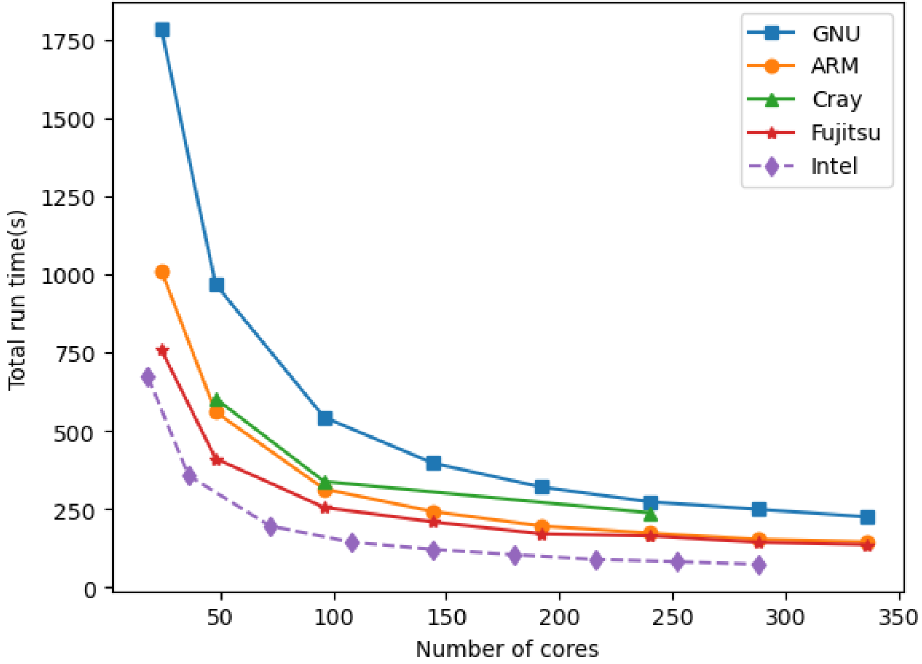

The total run time of 30-day simulations of MOM6 compiled by GCC, Arm, Cray, Fujitsu on Ookami, and Intel on Cheyenne.

Single node:
The Fujitsu compiler outperforms other compilers on Ookami
The Fujitsu compiler is two times faster than the GCC compiler

Multi-core performance:
Saturate after decrease
The Fujitsu compiler is the fastest

Cray simulations are aborted due to a Segmentation fault.

Intel model running on Cheyenne has a better performance than all the compilers on Ookami

Throughput

Throughput of model compiled by different compilers on 6 nodes

	Core hour/simulated year	Simulated year/wall clock day
GCC	239.0	28.9
ARM	146.9	47.0
Fujitsu	137.8	50.1

Power consumption ~117W per node

Mapping the MPI ranks by the NUMA regions

Number of tasks per NUMA	Total run time (seconds)
8	154.8(170.7)
9	147.4
10	143.5(164.2)
11	129.0
12	143.5

Total run time of Fujitsu-built MOM6 with respect to the number of tasks per NUMA running on 6 nodes. The numbers in parentheses indicate the run time using the same number of tasks but fully loaded NUMA in fewer nodes.

Fujitsu Profiling Tool

	12-task-per-NUMA case	11-task-per-NUMA case
The memory throughput peak ratio	38.2%	41.4%
The SVE operation rates	~85%	
mca_btl_vader_component_progress	29.5%	26.6%
MPI communication	36.3%	26.6%

MPI+OpenMP hybrid run

- 12 threads per NUMA, 4 MPI-tasks per node
- The OpenMP enabled MOM6 is successfully compiled by Arm, GCC, Fujitsu, and Intel compilers but fails to run
- Find an initialization bug in the source code

Summary and future work

- Fujitsu compiler is the fastest one on the Ookami.
- The throughput of Fujitsu-compiled MOM6 can satisfy our research requirements.
- 11 tasks per NUMA is the optimized setting for the future production run.
- MPI+OpenMP hybrid run will be tested in the future.
- Porting the fully coupled CESM on Ookami.

Thanks