Assessing Portability of MOM6 to GPUs Using OpenACC

End of Summer Progress Report

G. Dylan Dickerson, Briley James, Matthew Stack

July 28, 2020

Team Members

NCAR-Wyoming Interns

G. Dylan Dickerson
*Graduate Student**

Briley James
*Undergraduate Student**

NCAR Visiting Scholar

Matthew Stack
Undegraduate Student

University of Wyoming

Sumathi Lakshmiranganatha
*Graduate Student**

Zephaniah Connell
*Graduate Student**

Madison Shippy
*Undergraduate Student**

Oreoluwa Babatunde
*Undergraduate Student**

*Department of Electrical and Computer Engineering

Project Goals

Overarching Goal

- Achieve an operational, single source ocean model capable of running on heterogeneous computer architectures

MOM6 Goal

- Evaluate the potential for achieving CPU-GPU performance portability for MOM6

Summer Objectives

- Port the barotropic portion of the dynamical core of MOM6 onto hybrid architecture
- Determine a method of validating results of ported code against CPU versions

Porting Work Cycle

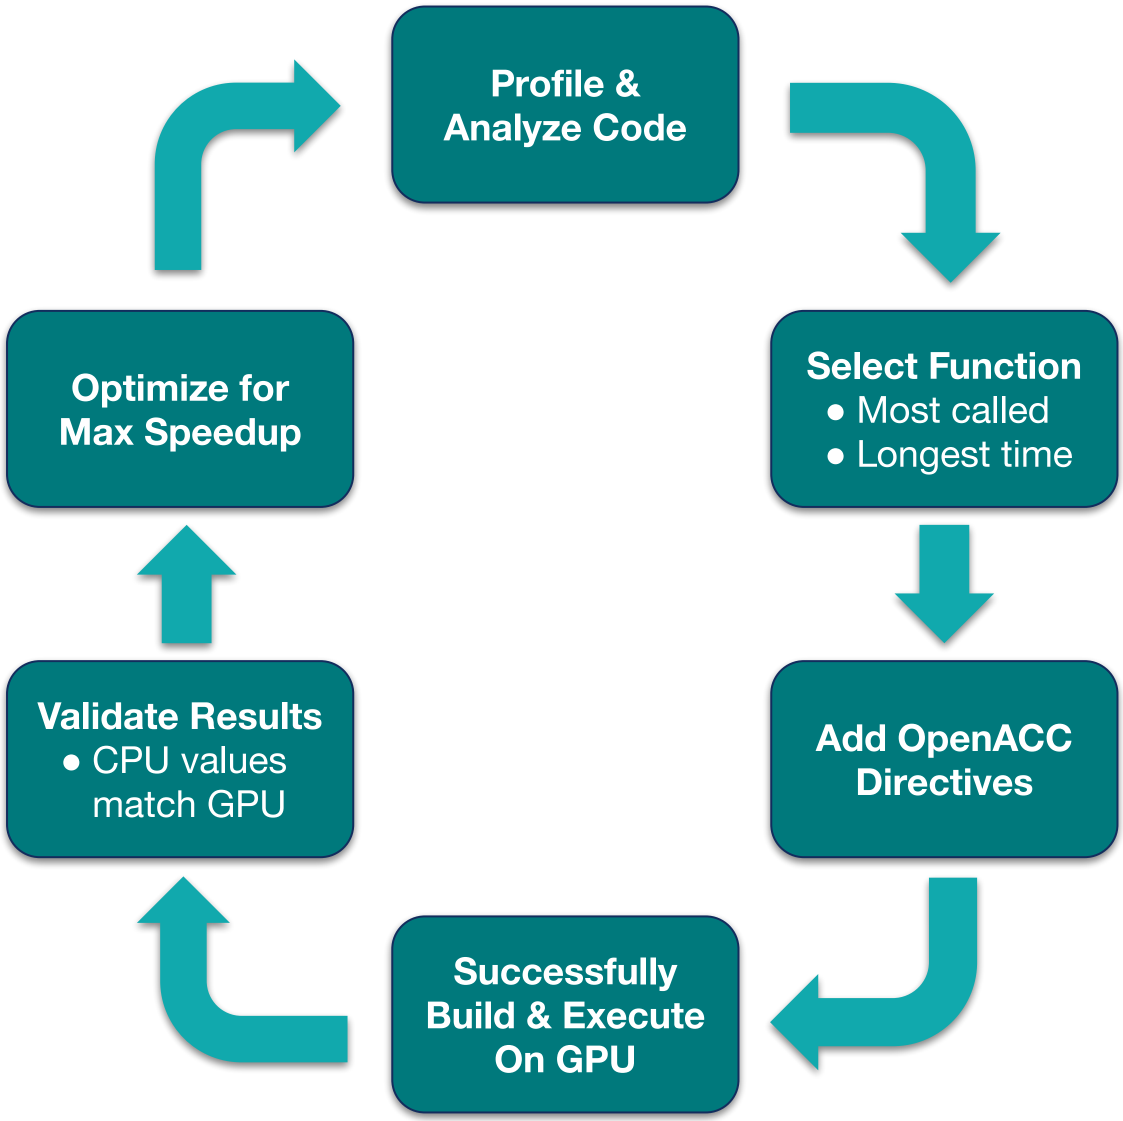

Project Timeline

Initial Work

- Modify CIME files to build/run CESM with CMOM compset on Casper
- Profile & identify important functions in CESM/CMOM

Post-Summer 2019

- Grow group/ Enlist help
- Strong scaling on CPUs
- Begin porting Barotropic Step

Project Timeline

Start of Summer 2020

- Recollect CESM profile
- CESM Strong Scaling
- Swap to Standalone MOM6
- Create workaround for Structure issue

Future Milestones

- Address Memory Issues
- Port the Remaining Functions
- Validate GPU Results
- Optimize Ported Functions
- Calculate Speedup

Summer 2020 | Beyond 2020

End of Summer 2020

- Standalone Profile
- Create new CPU strong/weak scaling plots
- Developed Automation Scripts
- Verify ported code results

Background

Real World

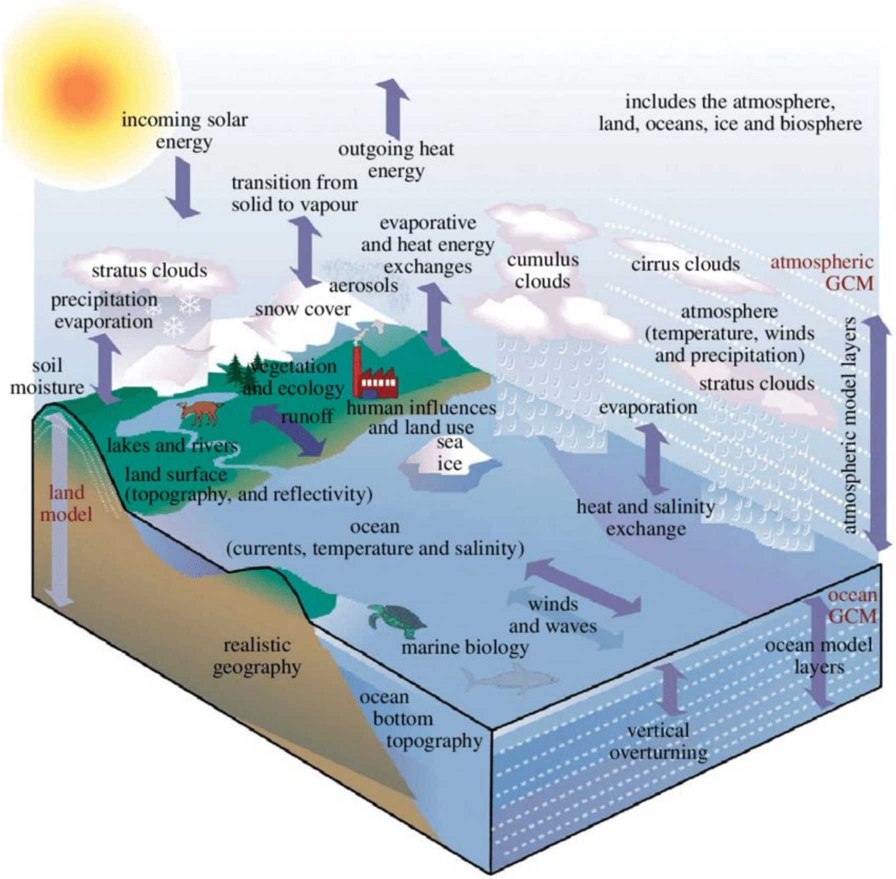

Community Earth System Model

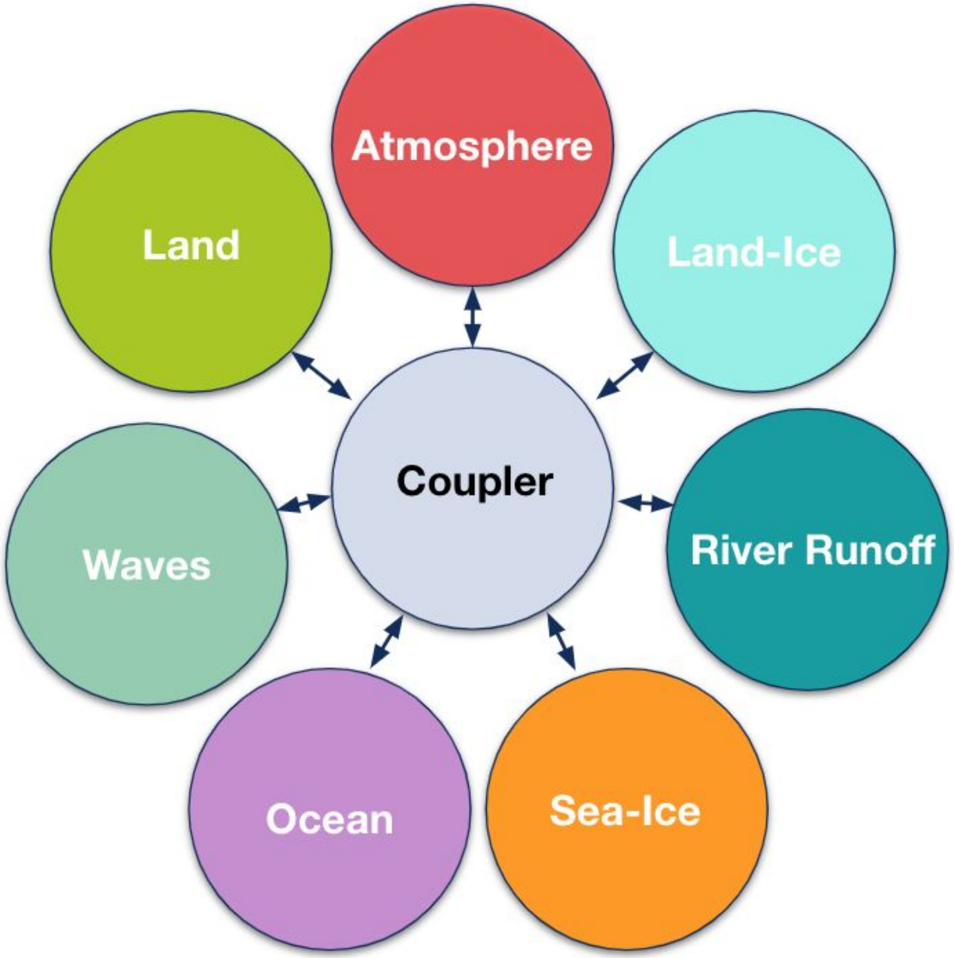

Per Nyberg (2006), Climate Model
https://www.csm.ornl.gov/PR/PR2006/hpc-07-21-06.html

Initial MOM6 Scaling Under CESM

May 19, 2020

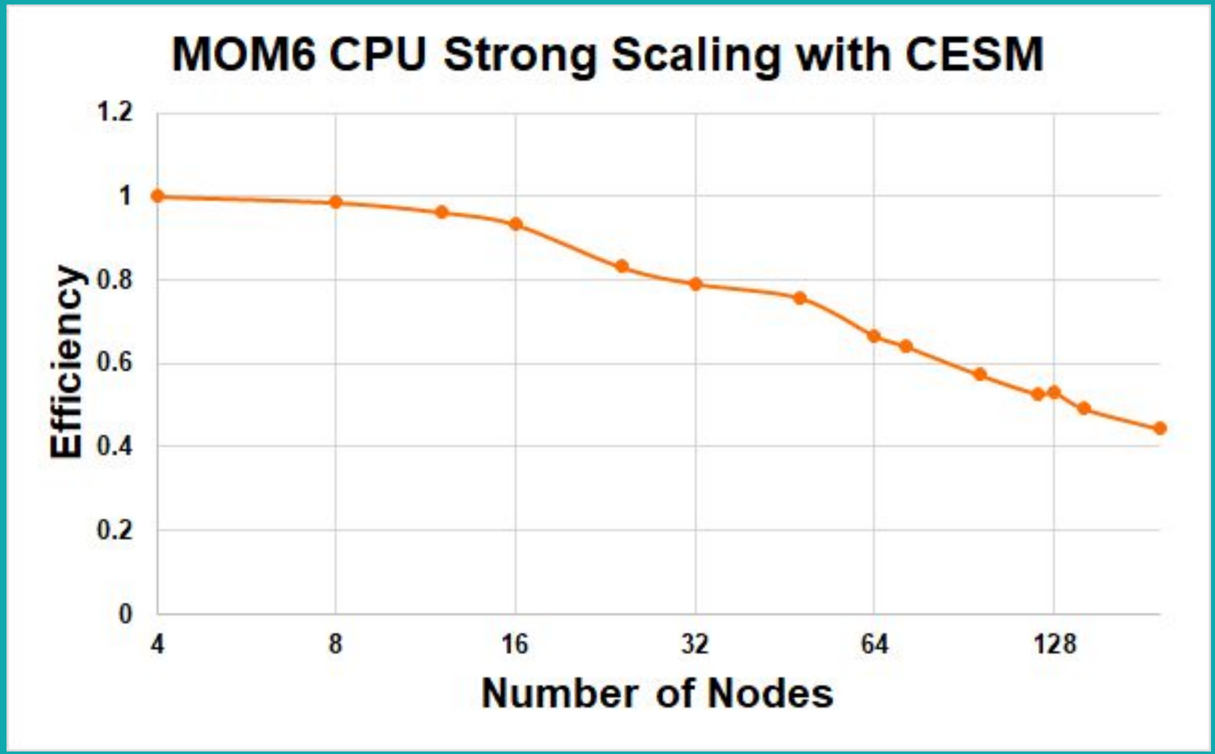

Machine:	NCAR Cheyenne
Compiler:	Intel 19.0.5
MPI Library:	MPT 2.19

Test Case:	CMOM
Simulated Days:	5
Resolution:	T62_t061

MOM6 Profile Under CESM Framework

May 19, 2020

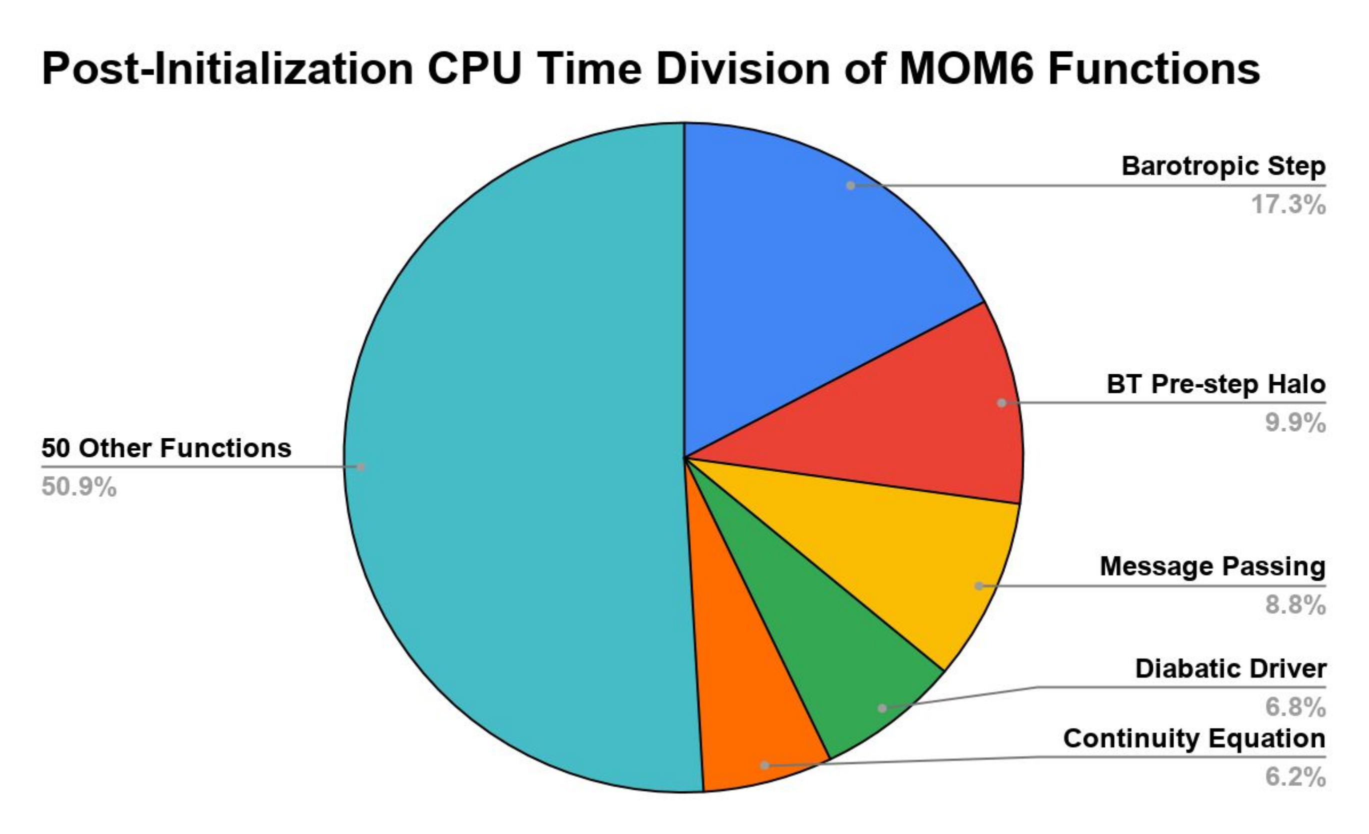

Machine: NCAR Cheyenne
Compiler: Intel 19.0.5
MPI Library: MPT 2.22
# of Nodes: 21

Test Case: CMOM
Simulated Days: 30
Resolution: T62_061
Ranks per node: 36

Swap to Standalone

CESM

Standalone

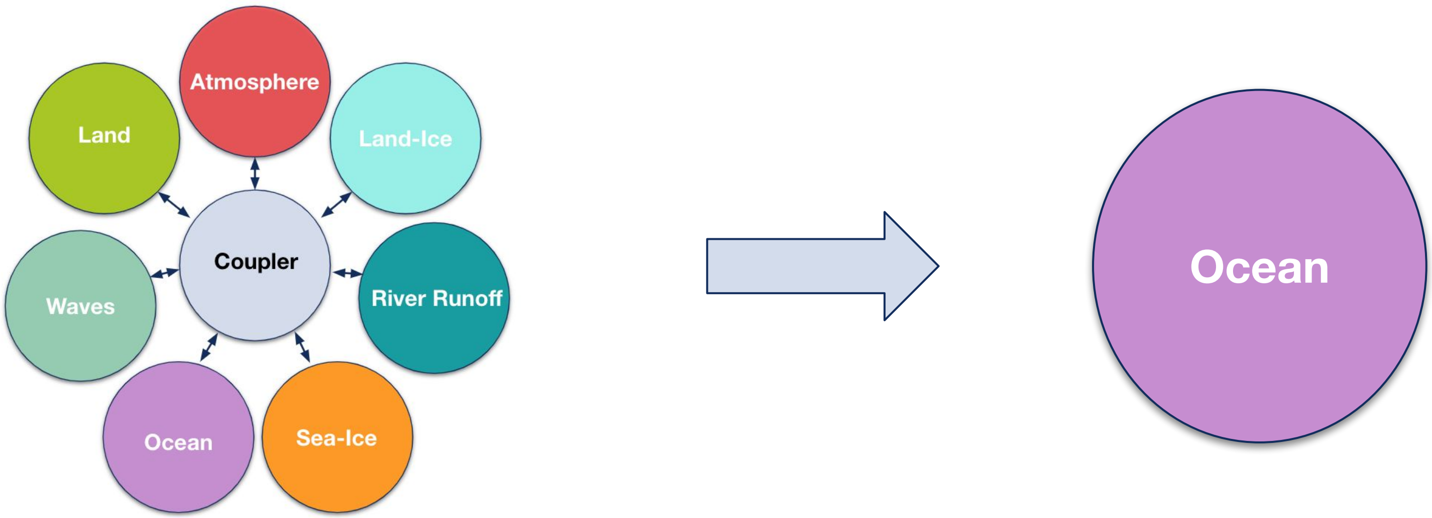

Framework	CESM2+MOM6	Standalone Channel Benchmark
Compilation Time (s)	593	510
Execution Time (s)	675	9.5

Machine: NCAR Cheyenne
Compiler: Intel 19.0.5
MPI Library: MPT 2.22
Number of Threads: 8

Test Case: CMOM Compset
Simulated Days: 5
Resolution: T62_t061
# of Nodes: 1
Ranks per node: 36

Test Case: Channel Benchmark
Simulated Days: 1
Resolution: 360x180
# of Nodes: 1
Ranks per node: 36

Channel Benchmark Test Case

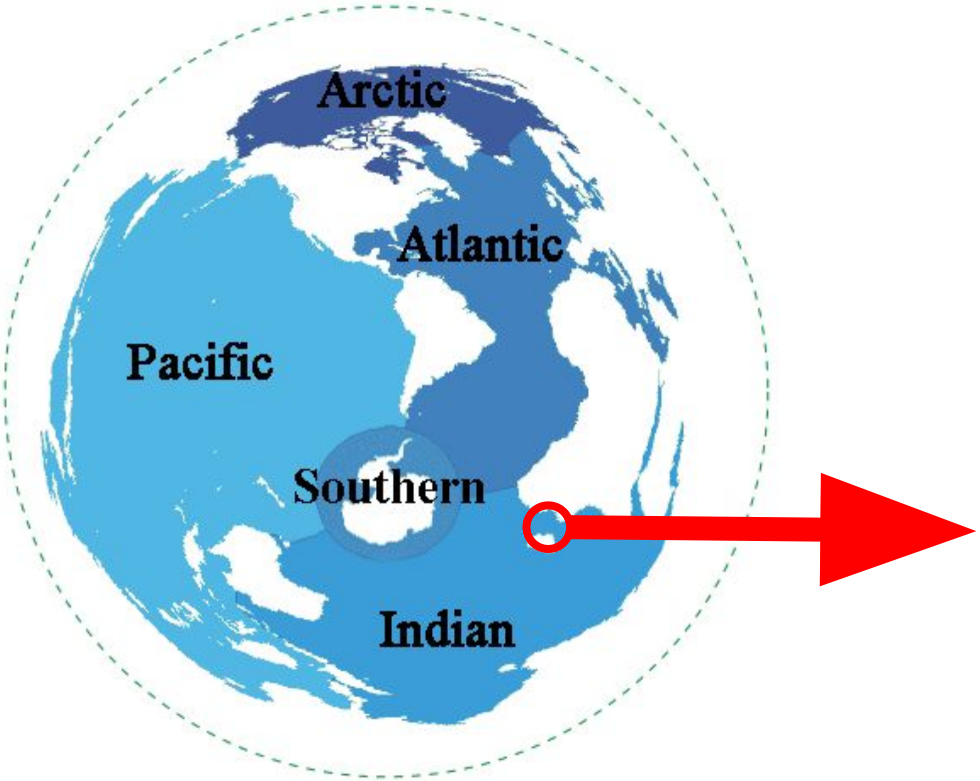

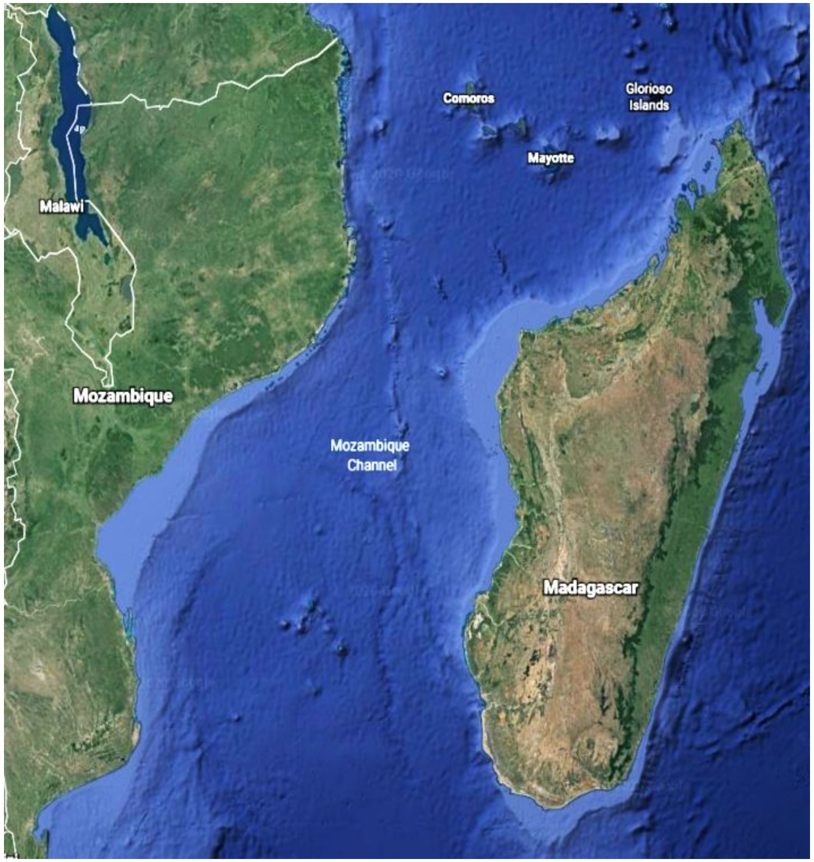

World Ocean Map 5 Oceans
https://commons.wikimedia.org/wiki/File:World_ocean_map_5_oceans.gif

Screenshot from Google Earth of Mozambique Channel
https://earth.google.com/web/

Standalone Build and Execution Configurations

CPU Only

- NCAR Cheyenne Supercomputer
- Intel Xeon E5-2697V4 (Broadwell) processors

Intel compiler
- Intel 19.0.5
- MPT 2.22
- NetCDF 4.7.3

PGI compiler
- PGI 20.4
- OpenMPI 4.0.3
- NetCDF 4.7.3

CPU + GPU

- NCAR Casper Supercomputer
- 4 V100 GPU's and 2 Intel Xeon Gold 6140 (Skylake) processors per node

- NVIDIA Prometheus Cluster
- 8 V100 GPU's and Broadwell Processors

PGI compiler
- PGI 20.4
- OpenMPI 4.0.3
- NetCDF 4.7.3

*See build scripts located in WorldShare directory for more details

Benchmarking Automation

- Written in Python
- Loops through based on mode and run-defining lists in the script

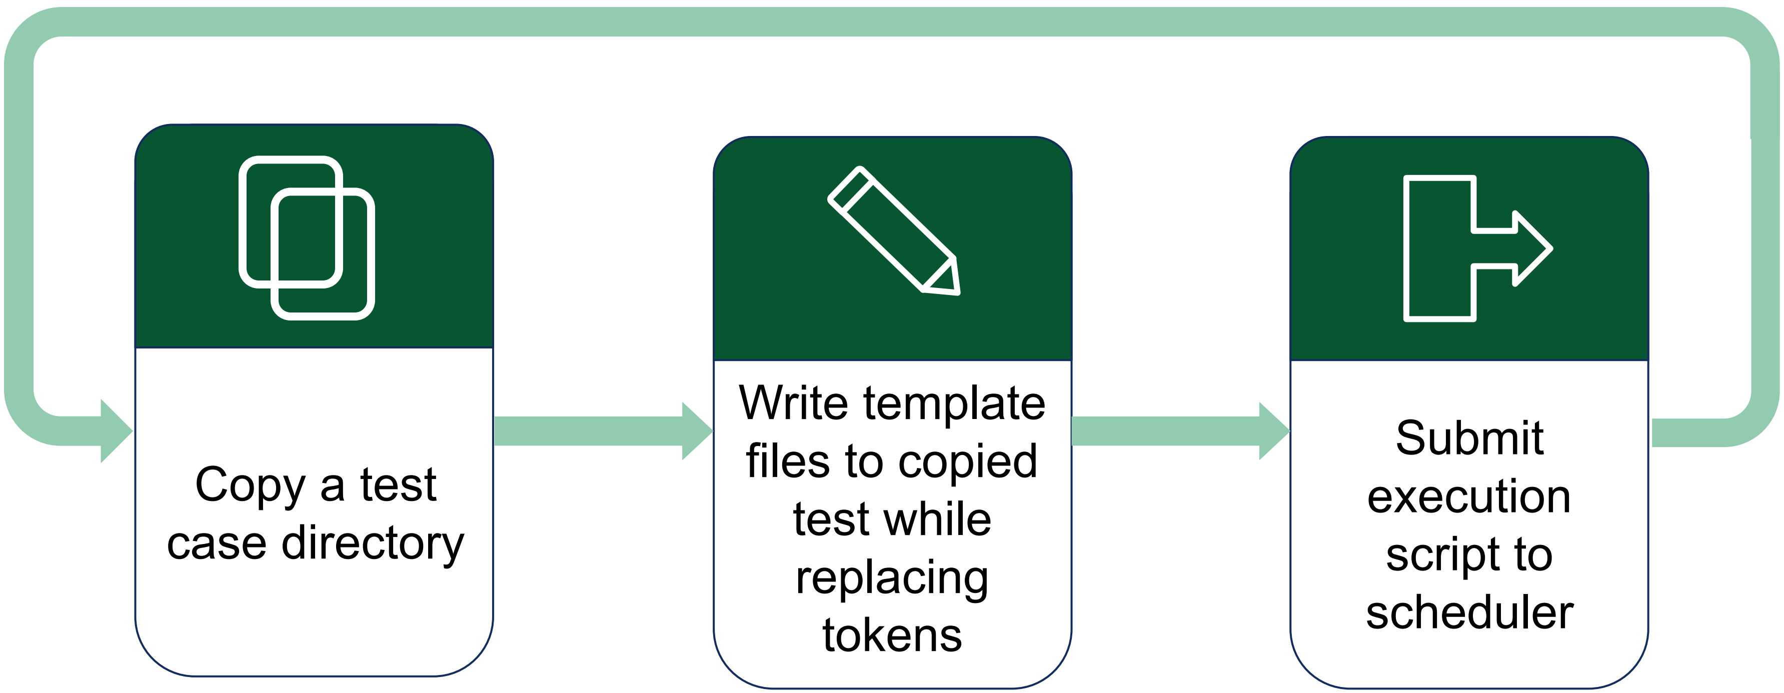

Weak Scaling Results

July 9, 2020

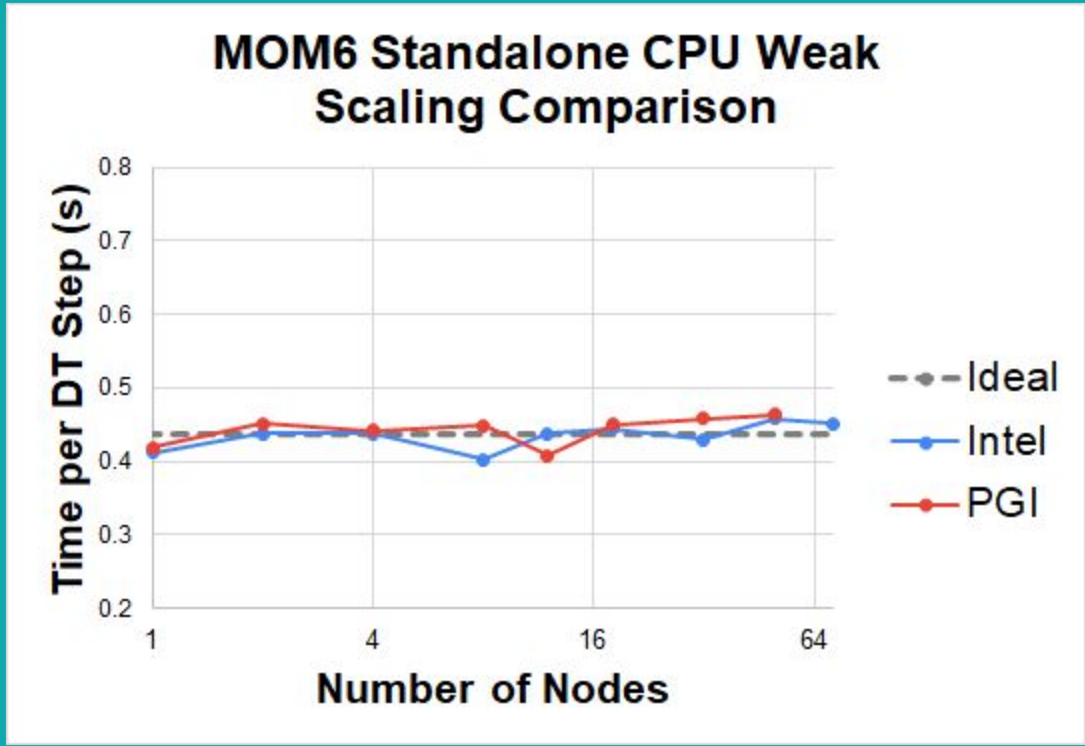

Machine: NCAR Cheyenne
Test Case: Channel Benchmark
Simulated Days: 5
MPI Ranks/node: 36

Intel Runs:
Compiler: Intel 19.0.5
MPI Library: MPT 2.22

PGI Runs:
Compiler: PGI 20.4
MPI Library: OpenMPI 4.0.3

Note: Resolution increases with node count for weak scaling

Strong Scaling Results

July 9, 2020

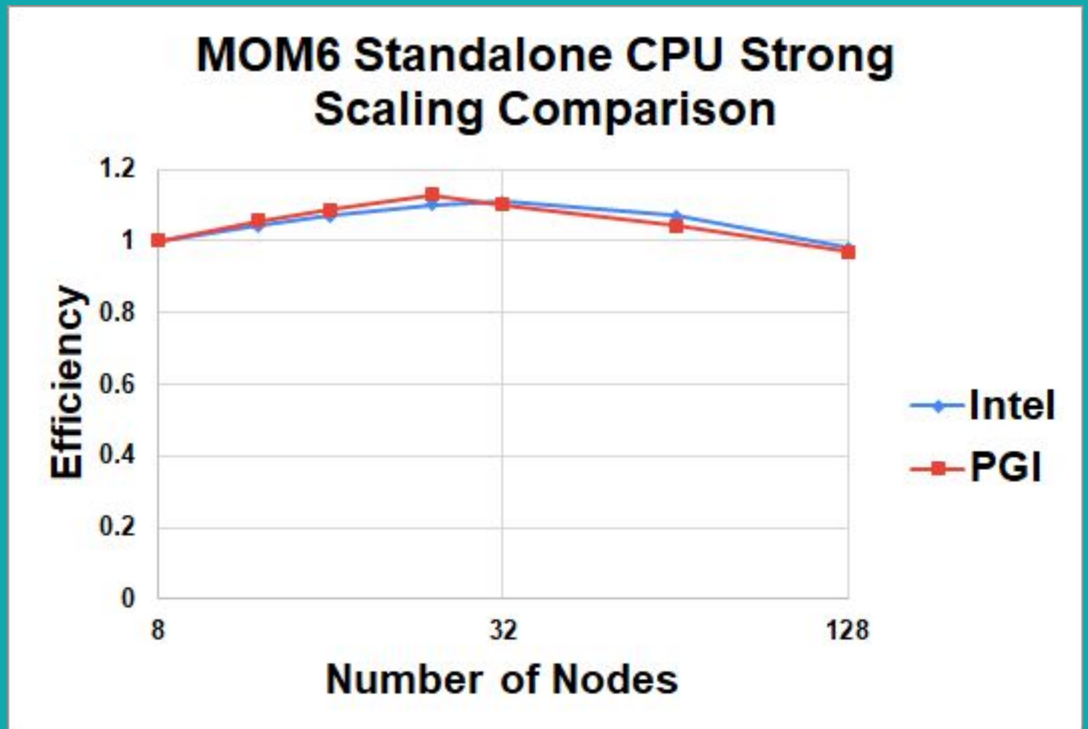

Machine: NCAR Cheyenne
Test Case: Channel Benchmark
MPI Ranks/node: 36
Simulated Days: 5
Resolution: 3072x1536

Intel Runs:
Compiler: Intel 19.0.5
MPI Library: MPT 2.22

PGI Runs:
Compiler: PGI 20.4
MPI Library: OpenMPI 4.0.3

Standalone Built-in Timers Profile

July 20, 2020

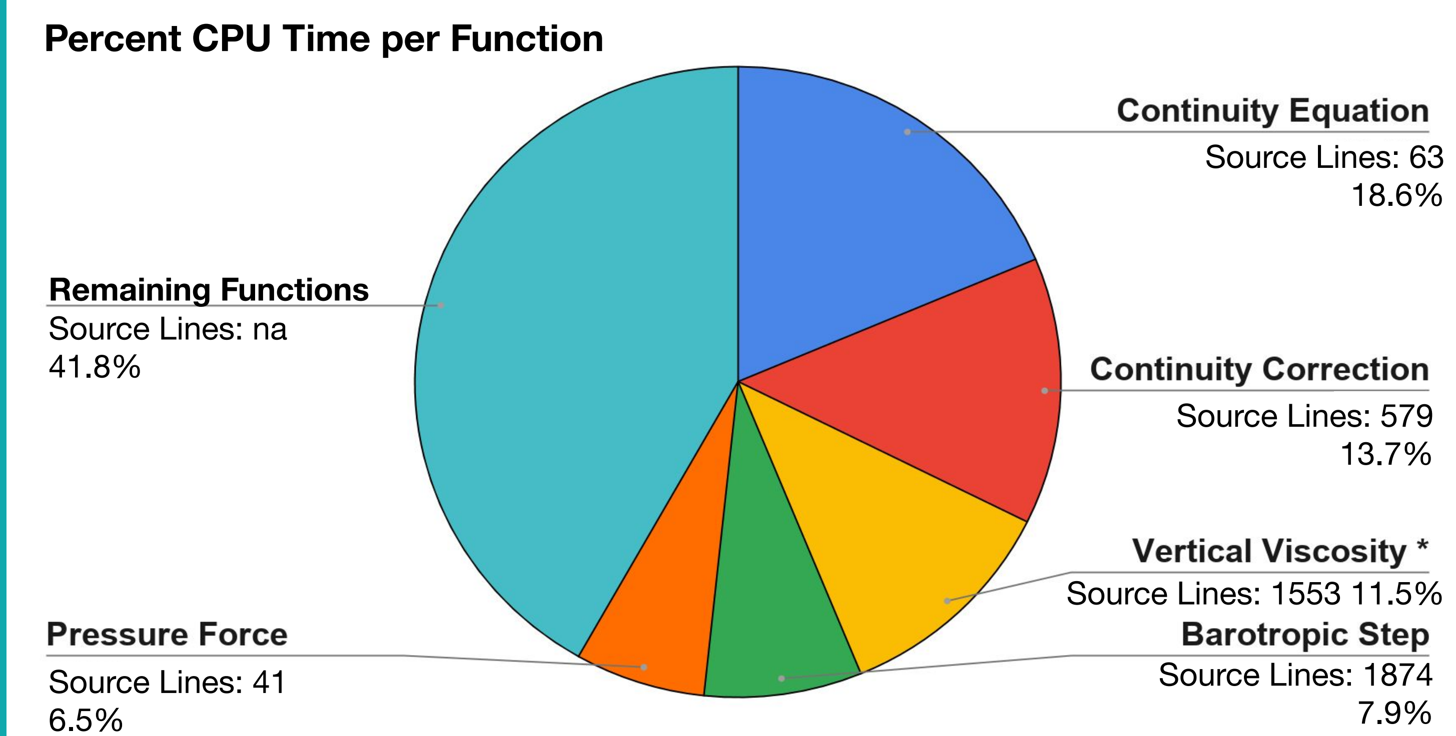

Machine: NCAR Cheyenne
Compiler: Intel 19.0.5
MPI Library: MPT 2.22
# of Nodes: 1

Test Case: Channel Benchmark
Simulated Days: 1
Resolution: 360x180
Ranks per node: 36

*Remaining functions contains remaining 65 functions

Standalone VTune Profiler Results

July 20, 2020

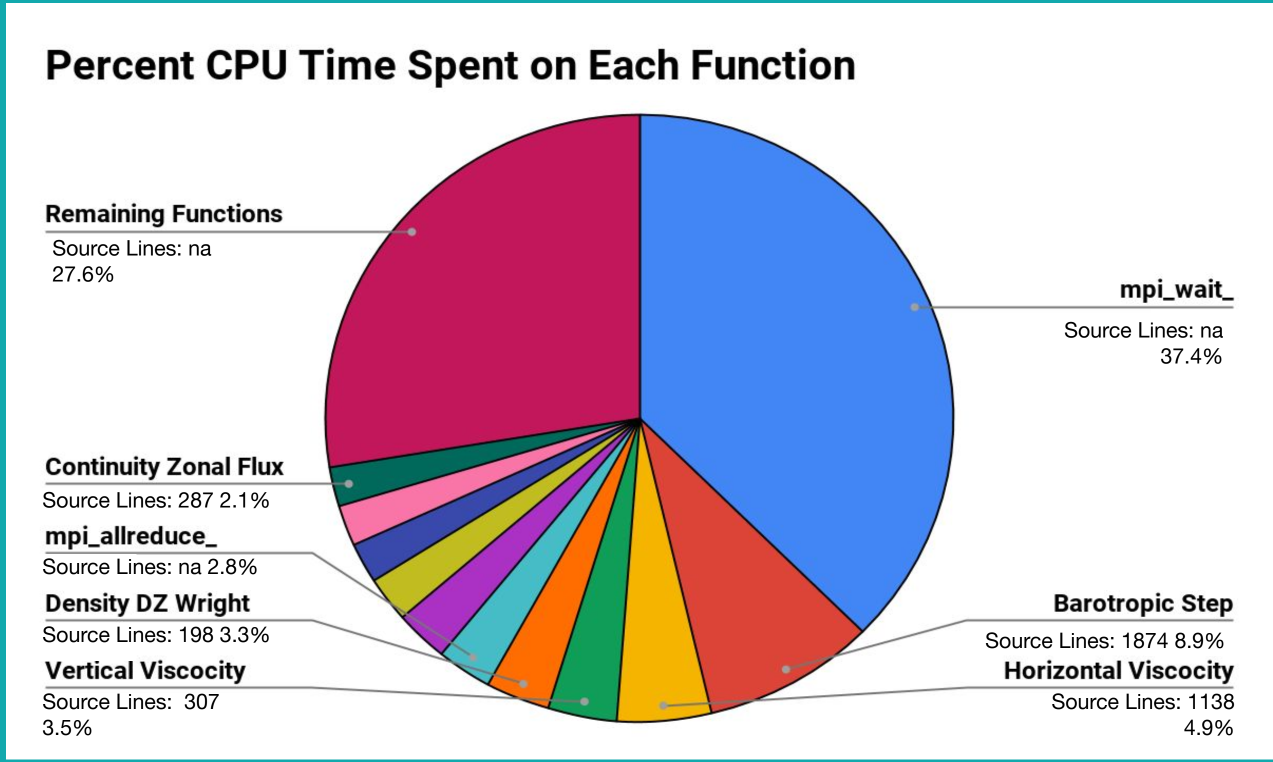

Machine: NCAR Cheyenne
Compiler: Intel 19.0.5
MPI Library: MPT 2.22
# of Nodes: 1

Test Case: Channel Benchmark
Simulated Days: 1
Resolution: 360x180
Ranks per node: 36

*Remaining contains 260 functions.
Collected by Sumathi L.

Issues Faced: Deep Copy

Deep Copy Structure Issue:

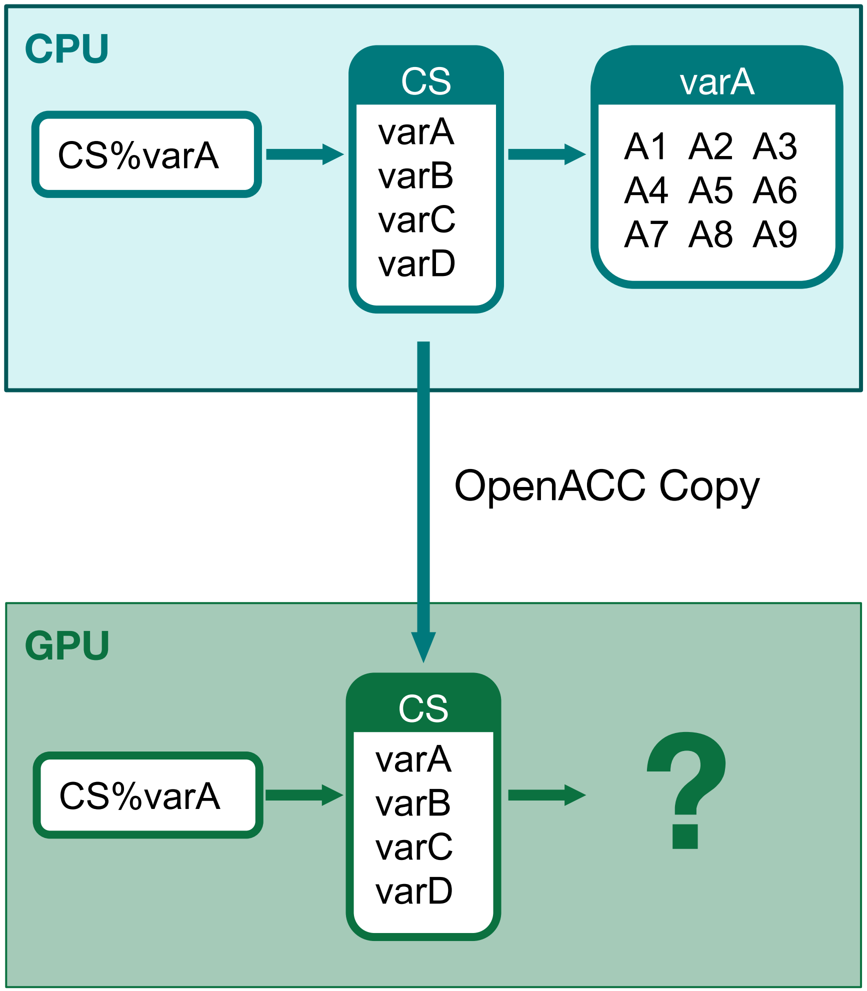

Workaround:

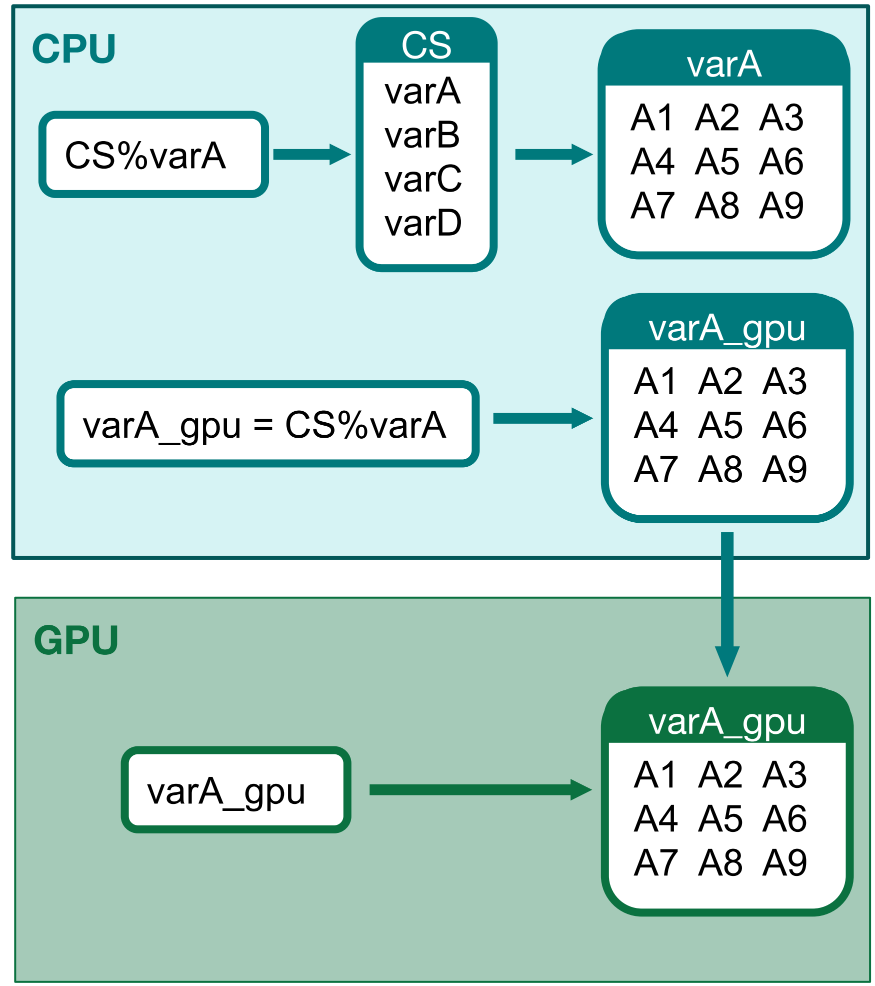

Issues Faced: Collapse Structures

Collapsing Dependencies: Scaling from 3D to 2D structure

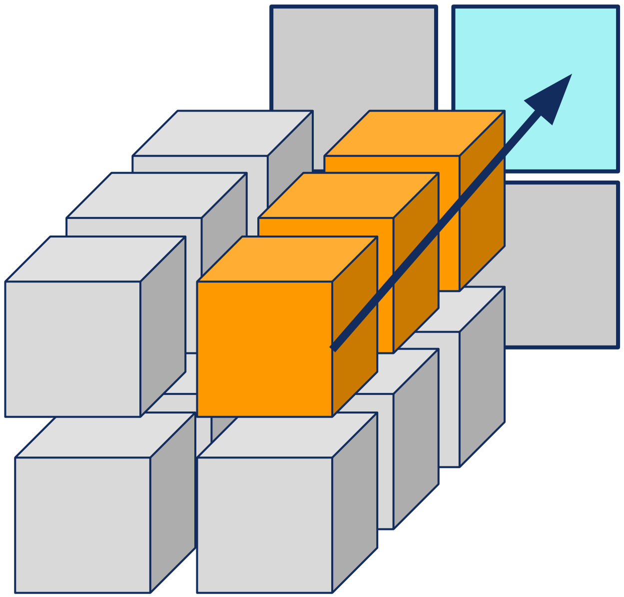

```
!Loop with dependency flag
!$acc parallel loop
do J=Jsq,Jeq ; do k=1,nz ; do i=is,ie
BT_force_v(i,J) = BT_force_v(i,J) + wt_v(i,J,k) * bc_accel_v(i,J,k)
enddo ; enddo ; enddo
!$acc end parallel
```

```
!$acc parallel
!$acc loop seq
do k=1,nz
!$acc loop collapse(2)
do J=Jsq,Jeq ; do i=is,ie
BT_force_v(i,J) = BT_force_v(i,J) + wt_v(i,J,k) * bc_accel_v(i,J,k)
enddo ; enddo ; enddo
!$acc end parallel
```

Issues Faced: Phantom Dependencies

Phantom Dependencies:

- Compiler saw false dependence, only optimized the outermost loop
- Compiler ignored overriding directives indicating independence
  - !$acc parallel loop private( )
  - !$acc parallel loop independent

Solution:

- Present compiler a 'single' loop for which it selects optimization
  - !$acc parallel loop collapse(2)

```
!$acc parallel loop
do j=jsv-1,jev+1 ; do i=isv-1,iev+1
   eta_pred(i,j) = (eta(i,j) + eta_src(i,j)) + (dtbt * CS_lareaT_gpu(i,j)) * &
             ((uhbt(I-1,j) - uhbt(I,j)) + (vhbt(i,J-1) - vhbt(i,J)))
enddo ; enddo
$acc end parallel
```

Issues Faced: Mystery Memory

Top command: PGI on Casper, No OpenACC

PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
20	0	858660	358028	19060	R	100.0	0.0	0:26.33	mom.exe
20	0	842480	350164	17676	R	100.0	0.0	0:26.45	mom.exe
20	0	842480	349996	17688	R	100.0	0.0	0:26.45	mom.exe
20	0	842480	349740	17732	R	100.0	0.0	0:26.45	mom.exe

Top command: PGI on Casper, OpenACC Enabled

PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
20	0	6040704	593104	216608	R	100.0	0.1	0:45.88	mom.exe
20	0	6028516	577696	215264	R	100.0	0.1	0:45.97	mom.exe
20	0	6028516	584008	215268	R	99.7	0.1	0:45.97	mom.exe
20	0	6028516	579264	215284	R	99.7	0.1	0:45.97	mom.exe

Top command: PGI on Cheyenne

PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
20	0	938996	362420	16788	R	100.33	0.553	0:05.76	mom.exe
20	0	943084	365528	16880	R	100.33	0.558	0:05.76	mom.exe
20	0	979620	378928	18268	R	100.00	0.578	0:05.63	mom.exe
20	0	943084	365260	16692	R	100.00	0.558	0:05.76	mom.exe

Virtual Memory

846,525 kB ≅ **.8 GB**

6,031,563 kB ≅ **6 GB**

951,196 kB ≅ **.9 GB**

Barotropic Step Subroutine Speedup

1 CPU node v. 1 GPU Prometheus node (RES 512x256)

2.23 x

CPU Intel v. GPU PGI

Machine / Hardware	Time(s)
Cheyenne 1 CPU	1.0432
Prometheus CPU+1 GPU	0.4669

4 CPU node v. 4 GPU Casper node (RES 1024x512)

0.8 x

CPU Intel v. GPU PGI

Machine / Hardware	Time(s)
Cheyenne 4 CPU	1.3099
Casper CPU+4 GPU	1.5346

Test Case: Channel Benchmark
Simulated Days: 0.1
Timesteps: 20

Intel CPU Runs:
Compiler: Intel 19.0.5
MPI Library: MPT 2.22

PGI GPU Runs:
Compiler: PGI 20.4
MPI Library: OpenMPI 4.0.3

Jupyter Lab Analysis

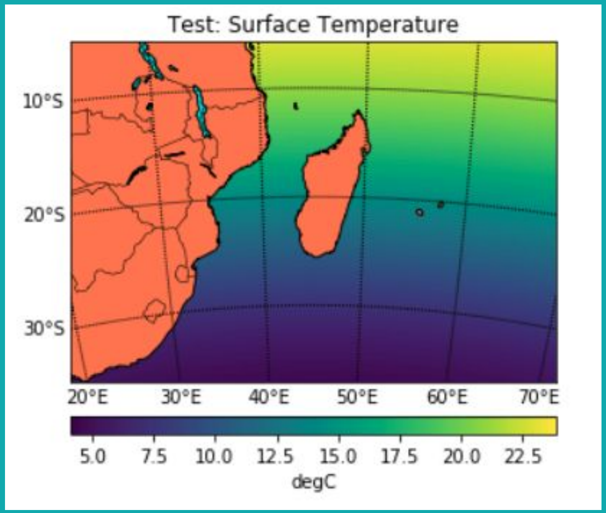

- Validates test values against reference data (generated on CPU, using Intel compiler)
- **Summarizes validation results** for selected variables
- **Visualizes surface variables** and their difference from the reference
- **Gathers timing information** for speedup comparison

Confirm Results

- Across Compilers: Intel vs PGI on CPU

Validation Results

Var Name	Max Error	
energyMass	2.2744359851056828e-07	PASS
maximumCFL	0.0	PASS
meanSeaLevel	1.3999999999999993e-10	PASS
totalMass	0.0	PASS
meanSalin	0.0	PASS
meanTemp	0.0	PASS
fracMassErr	2.0000000000000325e-18	PASS
salinErr	2.5999999999999995e-16	PASS
tempErr	5.0000000000000024e-17	PASS

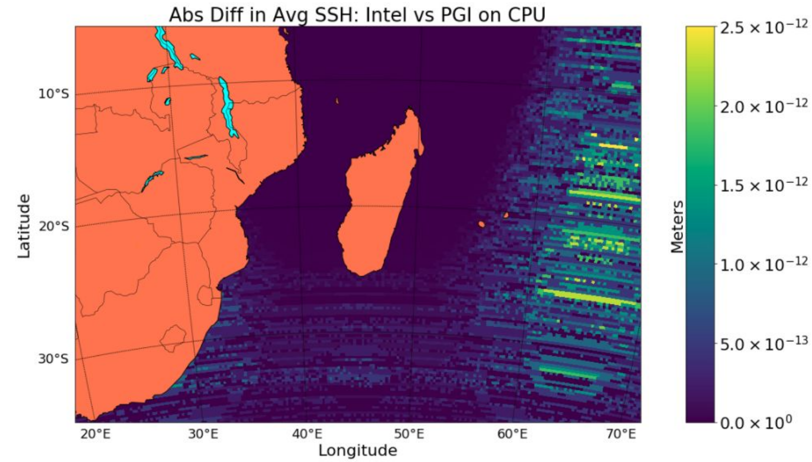

- CPU vs GPU: Intel vs PGI on GPU

Validation Results

Var Name	Max Error	
energyMass	2.7440251770816104e-07	PASS
maximumCFL	0.0	PASS
meanSeaLevel	1.5000000000001127e-10	PASS
totalMass	0.0	PASS
meanSalin	0.0	PASS
meanTemp	0.0	PASS
fracMassErr	2.0000000000000325e-18	PASS
salinErr	2.300000000000002e-16	PASS
tempErr	5.999999999999979e-17	PASS

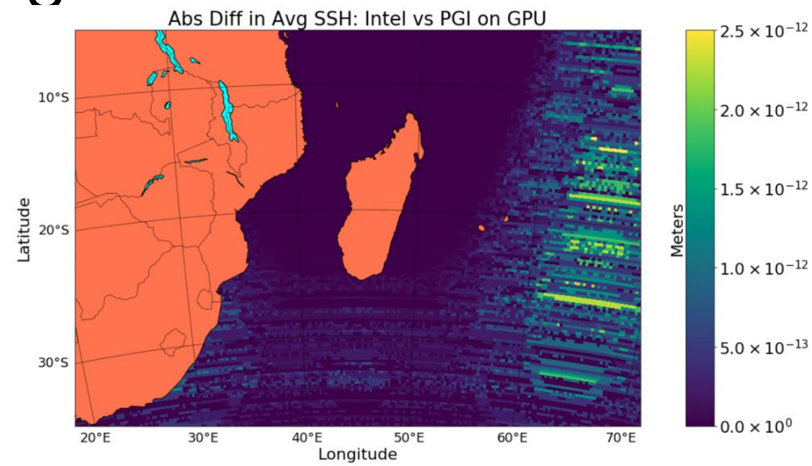

Current Project Status

Ported

Addressed the functions taking 19% of total time

- Barotropic Step
- Horizontal Viscosity
- Coriolis Advection Corad Calc
- Dynamics Split RK2

Up Next

Continue with functions taking the next 14% of total time

- Density DZ Wright
- Vertical Viscosity Coefficient
- Continuity Zonal Mass Flux
- Continuity Zonal Flux Adjust
- Continuity Meridional Mass Flux
- Continuity Zonal Flux Layer

Summary of Accomplishments

1. Changed to an Ideal Environment and Test Case
2. Identified Target Porting Functions
3. Ported **4369** lines in top 4 functions
4. Developed Automation Scripts for Scaling
5. Created Validation and Visualization Jupyter Notebooks

Future Milestones

End of Summer 2020

Future Milestones

Address Memory Issues

- Determine build settings for optimal code functionality
- Raise bug reports as necessary with compiler groups

Port the Remaining Functions

- Increase percent of code running on GPU's
- Update profile to select next target function(s)
- Calculate Speedup

Validate GPU Results

- Assess the correctness of a fully ported MOM6
- Run on Multiple GPU nodes to generate Scaling Plots

Optimize Ported Functions

- Reach a higher optimization standard for maximum performance
- Approach Roofline max

Code Sources

MOM6 Standalone

Framework: https://github.com/NOAA-GFDL/MOM6-examples

Github wiki for setup instructions:

https://github.com/gdicker1/MOM6/wiki/Casper-Standalone-Workflow

MOM6 under CESM

Framework: https://github.com/ESCOMP/CESM.git

Github wiki for setup instructions:

https://github.com/gdicker1/MOM6/wiki/Casper-Workflow

WorldShare (Build, Run, and Validation Scripts)

WorldShare: https://github.com/supreethms1809/MOM6_WorldShared.git

Validation Documentation:

https://github.com/supreethms1809/MOM6_WorldShared/blob/master/RegTestNBs/README.md

MOM6 port in Progress

See each of our GitHub repositories.

*See build scripts located in WorldShare directory for more details

References

1. Casper | Computational & Information Systems Lab. Retrieved July, 2020, from https://www2.cisl.ucar.edu/resources/computational-systems/casper
2. Cheyenne | Computational & Information Systems Lab. Retrieved July, 2020, from https://www2.cisl.ucar.edu/resources/computational-systems/cheyenne
3. Per Nyberg (2006).Climate model diagram [online image].HPCwire. https://www.csm.ornl.gov/PR/PR2006/hpc-07-21-06.html
4. Quizmodo (2020). World Ocean Map 5 Oceans [online image]. https://commons.wikimedia.org/wiki/File:World_ocean_map_5_oceans.gif
5. Mozambique Channel Screenshot [online image]. Retrieved July 20, 2020 from https://earth.google.com/web

- All logo Images were taken from respective organizations branding site

Acknowledgements

- Sumathi Lakshmiranganatha - University of Wyoming
- Zephaniah Connell - University of Wyoming
- Madison Shippy - University of Wyoming
- Oreoluwa Babatunde - University of Wyoming
- Supreeth Suresh - NCAR, project mentor
- Cena Miller - NCAR, project mentor
- Dr. Raghu Raj Kumar - NVIDIA
- Pranay Reddy Kommera - University of Wyoming
- Dr. Alper Altuntas - NCAR, CGD
- Dr. Gustavo Marques - NCAR, CGD
- Dr. Richard Loft - NCAR, CISL Deputy Director
- Mick Coady - NCAR, CISL
- Ben Matthews - NCAR, CISL
- Jared Baker - NCAR, CISL
- AJ Lauer - NCAR, CODE
- Virginia Do - NCAR, CODE
- Jerry Cyccone - NCAR, ASP

Thank You

G. Dylan Dickerson
gdicker1@uwyo.edu
github.com/gdicker1

Briley James
bjames10@uwyo.edu
github.com/brileyj

Matthew Stack
mattst@udel.edu
github.com/matt-stack