

ACCESS-OM2

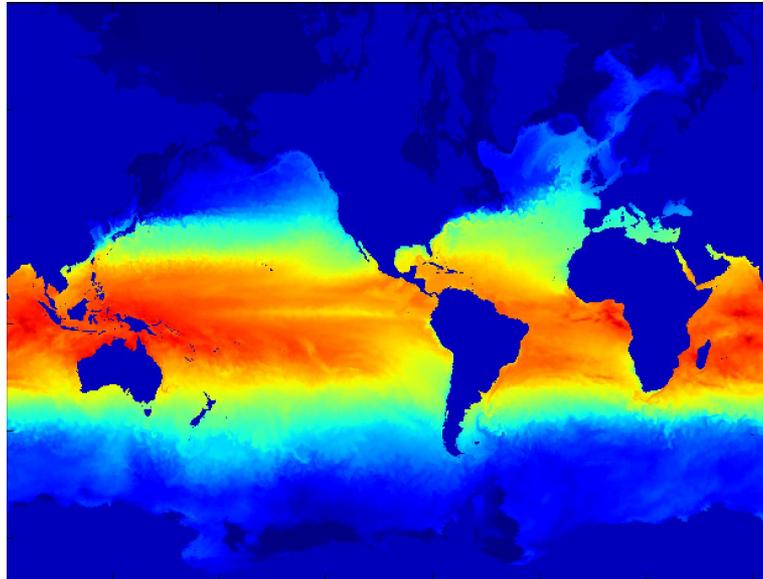
A coupled model renovation project for COSIMA

Nic Hannah
26/05/17



The goal

- Global 1/10th ocean and ice using: MOM5, CICE5, JRA55/CORE2, OASIS
- A framework - tools, process and documentation - useful for future model development



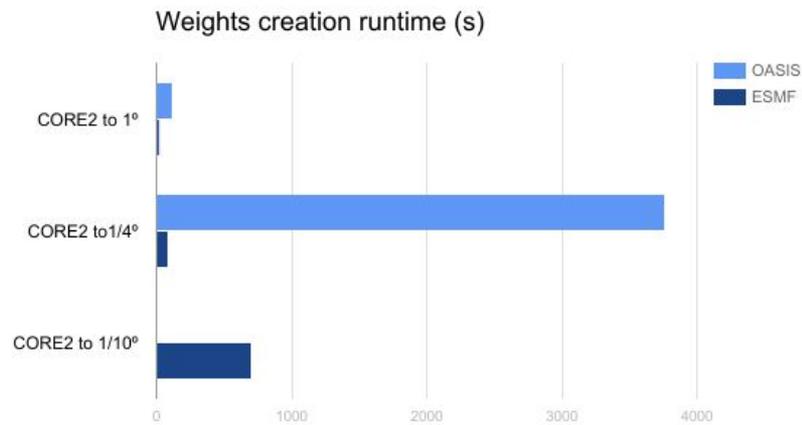
Foundations:

ACCESS-OM 1 and 1/4° MOM5, CICE4, OASIS3-MCT, CORE2

ACCESS-OM 1° as above with JRA55 (Fabio)

MOM+SIS 1/10°

Problem: interpolation weight generation

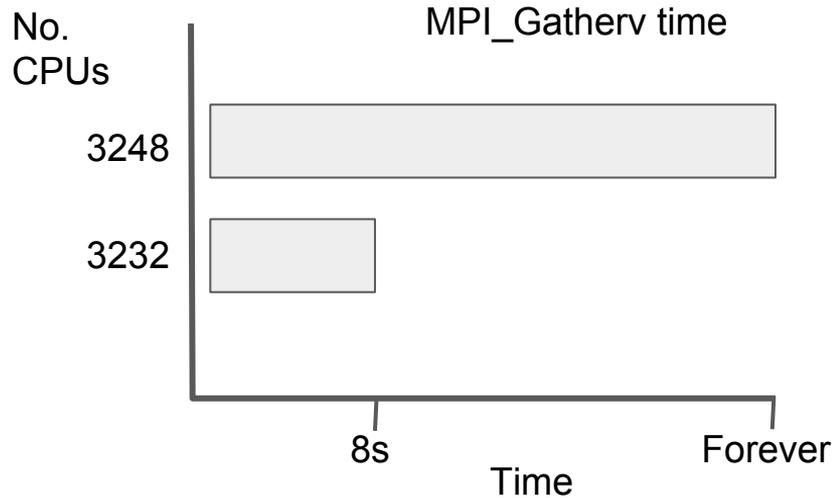


- Use ESMF_RegridWeightGen instead of SCRIP for big performance improvement
- Also gives us more sophisticated algorithms, e.g. 2nd order conservative
- OASIS3-MCT modifications were necessary

Hiccup: FMS MPI global operations

(this page is intentionally left blank)

Roadblock: OASIS startup won't scale (?)



Clue: “The compute nodes in the system utilise the latest Mellanox EDR interconnect (100Gb/s) in 2:1 blocking fat-tree topology...”
(<https://opus.nci.org.au/display/Help/Broadwell+Compute+Nodes>)

Realisation: we have serious performance problems

- 1/10th spin-up is looking intractable. E.g. 2 days per hour on 3k PEs! (dt = 150)
- Matching MOM+SIS is looking very unlikely.



**DON'T
PANIC**

Performance Makeover

- First step was to create equivalent 1 and $\frac{1}{4}$ degree configurations. These have been invaluable development tools.
- Add performance counters to the code, particularly to measure the difference between waiting for a coupling exchange and doing one.

Conclusion: reduce submodel wait times, i.e. maximise parallelism

Reducing wait times

1. Optimise CICE PE decomposition and balance ocean and ice. This is a long story.
2. Remove ice to atm coupling exchange.
3. Reorder CICE halo updates for incoming coupling fields from ocean and atmosphere.
4. Optimise OASIS3-MCT application of interpolation weights.

Timing for one month runs

	OM-025 (1168)	OM-025 sl-ice (2288)	OM-025 sq-ice (2416)	OM2-025 (2416)
Ocean wait time	33 sec		20 sec	8 sec
Total runtime	780 sec	418 sec	410 sec	354 sec

Mid-renovation status:

- MOM5, CICE5, JRA55 coupled with OASIS3-MCT / ESMF at three resolutions. 1 and $\frac{1}{4}^\circ$ need validation, $\frac{1}{10}^\circ$ is in spin-up.
- 1° performance: ~20 minutes per year. ~60 years per 24 hour submit. 304 PEs, dt = 3600s
- $\frac{1}{4}^\circ$ performance: ~72 minutes per year. 2416 PEs, dt = 1800s
- $\frac{1}{10}^\circ$ performance. Good news is that CICE5 seems to scale well so ocean wait times are low.

Conclusions

- Developed tools to help create model configurations.
- Introduced a new way to make interpolation weights.
- New configurations for three resolutions. Together these are more than the sum of parts.
- Overcame some tricky model and Raijin problems.
- Big improvements to performance. The coupling overhead is about 2-3% of MOM runtime. Can now compete with MOM+SIS.

Future work

- Unanswered performance questions. Should be faster than MOM+SIS.
- Improve efficiency. Need processor masking, also CICE waits a lot.
- Integrate 2nd order conservative remapping method.
- Lots, lots more.