R: In Situ Eddy Analysis in MPAS-Ocean

Jonathan Woodring, Mark Petersen, Andre Schmeißer,
John Patchett, James Ahrens, Hans Hagen

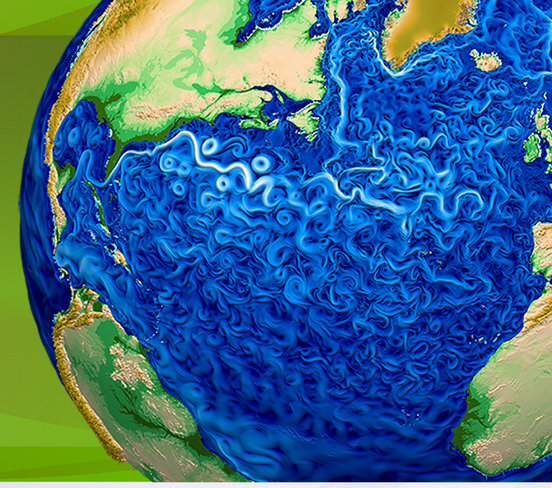

Objective

The study of eddies in global ocean-climate models requires large-scale, high-resolution simulations. MPAS-Ocean needs to write (store) a massive amount of data needed for eddy detection, which poses a problem for frequent (daily) censuses.

Shown in the right figure is the total end-to-end time for both simulation and simulation + parallel eddy census, if it is analyzed in a post-processing manner. At eddying resolution (15 km grid-cell size or smaller), the time to completion at 8192 processors takes 10 hours: i.e., the simulation only takes 2 hours, while daily eddy censuses adds 8 hours to the total.

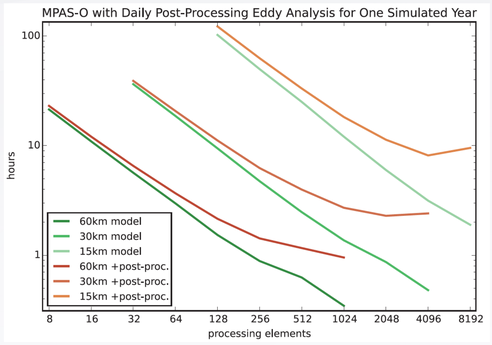

Why? We can see that the simulation continues to scale, while the I/O performance flattens out (Lustre filesystem on Wolf at LANL). Read/write bandwidth does not scale with the number of processors: it is a fixed, shared resource. This is a best case scenario, as it does not take into account copying data to and from archive (HPSS).

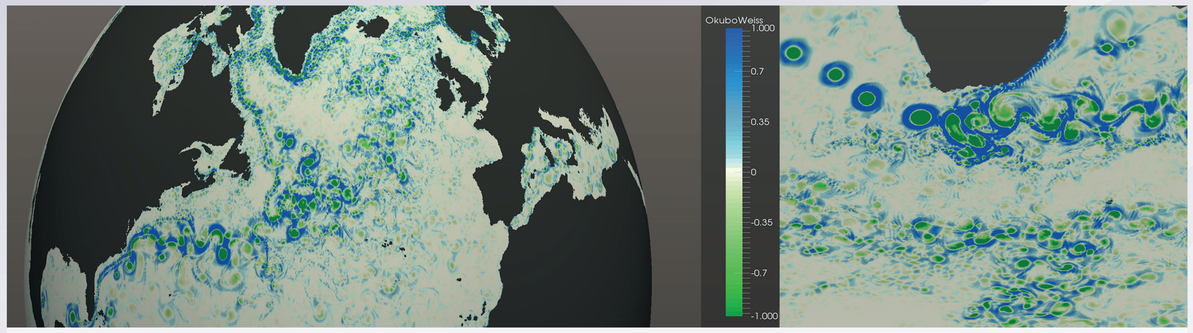

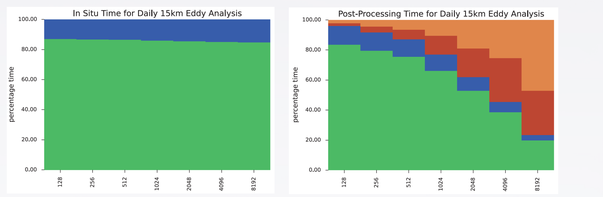

Relative time taken for both *in situ* and post-processing eddy analysis in MPAS-Ocean for daily eddy censuses. Green is model time, blue is eddy analysis time, red is model write time, and orange is analysis read time. As processor count increases, read and write time does not scale out.

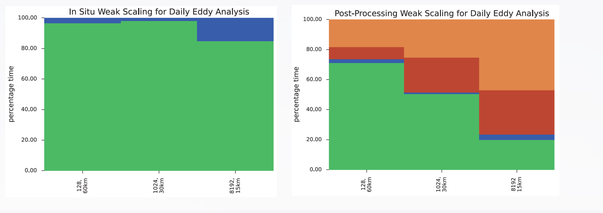

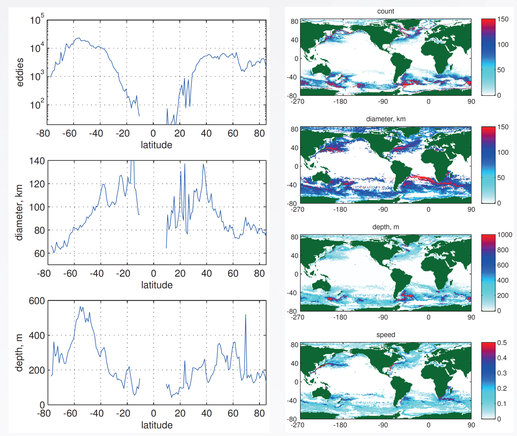

Early validation results showing the census distribution of eddies over the MPAS-Ocean global ocean-climate model.

Approach

To reduce the total time-to-completion (the time measured between when a simulation is started and when the analytical results are acquired), we have developed an *in situ* (an integrated computation along side the model time stepping) eddy census analysis that scales well to ten-thousand processing elements.

The time taken for generating an *in situ* census, as the same case as above (8192 processors), only increases total time to results by 15 minutes (from 2 hours to 2 ¼ hours) compared to 8 hours. This savings is due to to avoiding writing and reading the data (see the plots in the upper right corner of this poster).

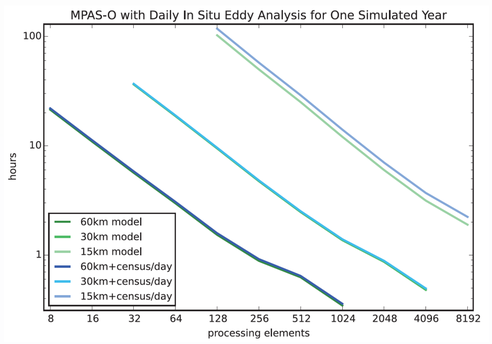

The In Situ Algorithms utilized for the Eddy Censuses:

- MPAS-Ocean analysis member framework that is used for *in situ* computations
- Parallel Okubo-Weiss for calculating rotating and shearing (similar to Lambda-2)
- Parallel connected components to extract eddies and generate the census

Impact

MPAS-Ocean was initially designed with an *in situ* analysis methodology as a core development principle, for both the model and the analysis of the model.

This study validates that design decision by quantifying that *in situ* computations can have many advantages that pertain to analysis tasks, as we move to higher spatial and temporal resolutions in climate models. It also shows early science results, validating eddies in the MPAS-Ocean global ocean-climate model.

Paper:
In Situ Eddy Analysis in a High-Resolution Ocean Climate Model
J Woodring, M Petersen, A Schmeißer, J Patchett, J Ahrens, H Hagen
Visualization and Computer Graphics, IEEE Transactions on (TVCG): 22 (1), 857-866

Source Code: MPAS-Ocean (available in ACME) in the analysis member framework
http://mpas-dev.github.io *http://github.com/MPAS-Dev*

Thanks to: Doug Jacobsen, Todd Ringler, and the MPAS-Ocean development team

ACME Accelerated Climate Modeling for Energy

For additional information, contact:
Jon Woodring
Research Scientist
Los Alamos National Laboratory
(505) 695-6161
woodring@lanl.gov

climatemodeling.science.energy.gov/acme

U.S. DEPARTMENT OF ENERGY