Characterizing and Optimizing Performance in MPAS-Ocean

Divya Jaganathan[a,b]
Mark R. Petersen[b], Luke P. Van Roekel[b]
Brown University[a], Los Alamos National Laboratory[b]

LA-UR-18-27173

Abstract

The **Model for Prediction Across Scales (MPAS)-Ocean** is an unstructured-mesh model with variable resolution capability that addresses the computational demand of high-resolution modeling of ocean in regional domains of interest.

- We improve computational **load-balance** across parallel processing units using **depth-weighted mesh-partitioning algorithm** and guide an optimal range for cell weight factor.
- We demonstrate **computational performance gains** in simulations using a **variable resolution mesh** that retain high-quality fine scale features regionally.

Background

- Ocean domain is decomposed into **partitions of equal number of surface cells** using Metis[2] for **distributed-memory parallel computations**.

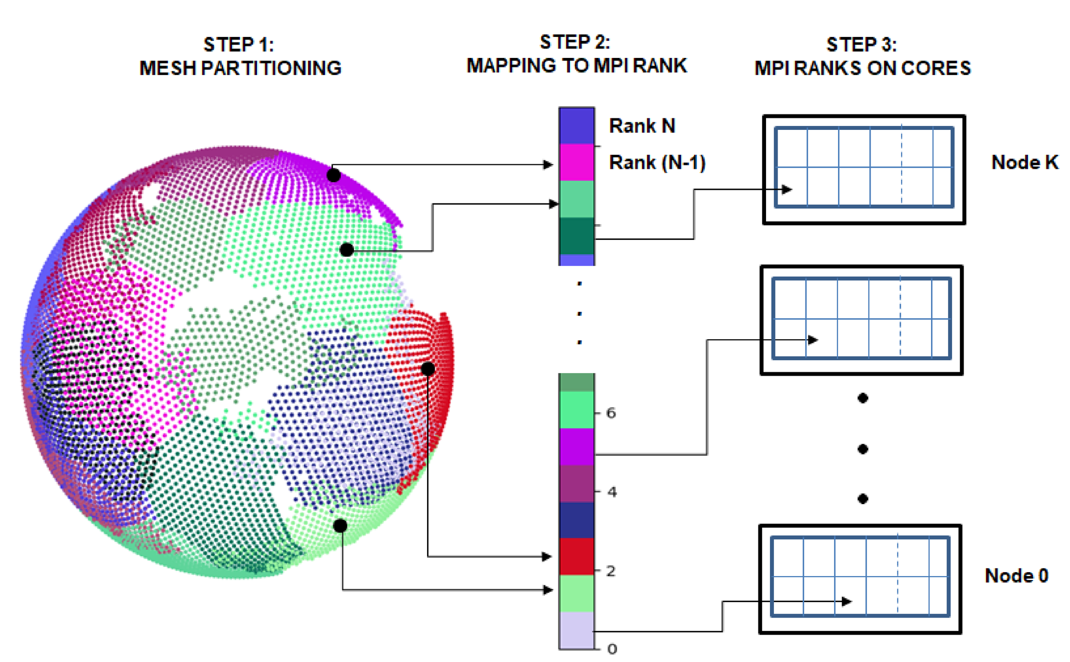

Figure: Representation of Parallelism in MPAS-Ocean Model

- Equal partitioning based on surface cells assumes uniform computational cost per cell despite an unequal number of vertical cells stacked under each surface cell.
- The computational cost is further complicated as different subroutines in the model scale differently with number of vertical cells, $\mathbf{N_{vcells}}$

Barotropic 2D Algorithms	Baroclinic 3D Algorithms	Vertical Mixing
O(1)	$O(N_{vcells})$	$O(N_{vcells}^2)$

Table: Sub-routines and their orders of dependence on N_{vcells}

Analysis of Weighted Domain Decomposition for Load Balancing

$$\text{Cell Weight factor} \propto (\text{Number of Vertical Cells})^{X}$$

X - global average dependence factor on number of vertical cells

$$\text{Parallel Efficiency} = \frac{\text{Average time taken by partitions}}{\text{Maximum time taken by a partition}}$$

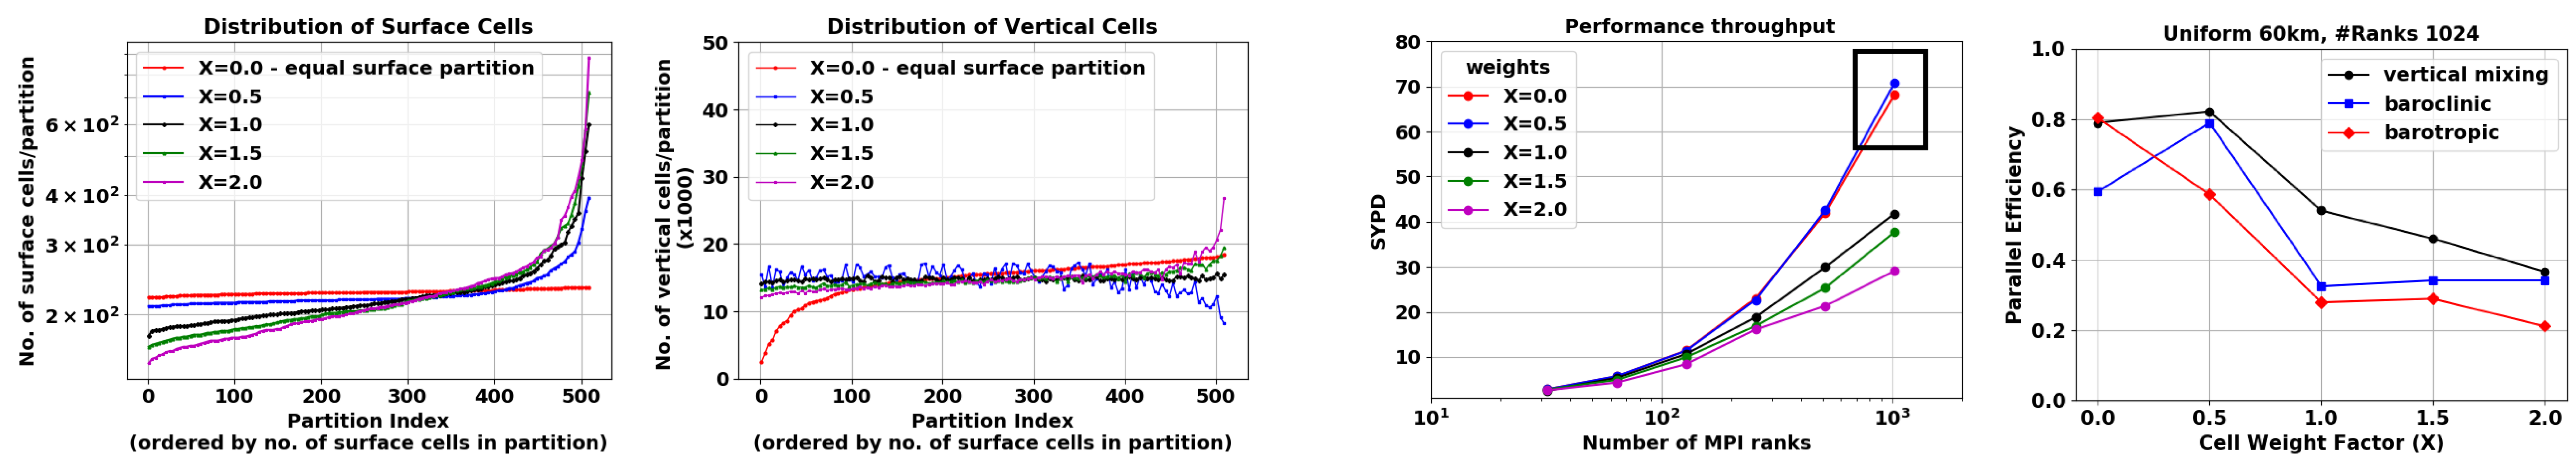

Figure: (From Left to Right) Distribution of number of surface cells and vertical cells shows that load balance in terms of equal number of cells per partition is achieved at X=0.5; Performance throughput improves for X<1; (Right-most) Sample sub-routines with different dependence on $\mathbf{N_{vcells}}$ attain maximum parallel efficiency at different cell weight factors
All tests shown above were done on **36-core Broadwell** at **60 km uniform resolution**.

Performance of Variable Resolution Mesh

$$\text{Simulated Years Per Day (SYPD)} = \frac{\text{Simulated Duration in Years}}{\text{Computational Time Taken in Days}}$$

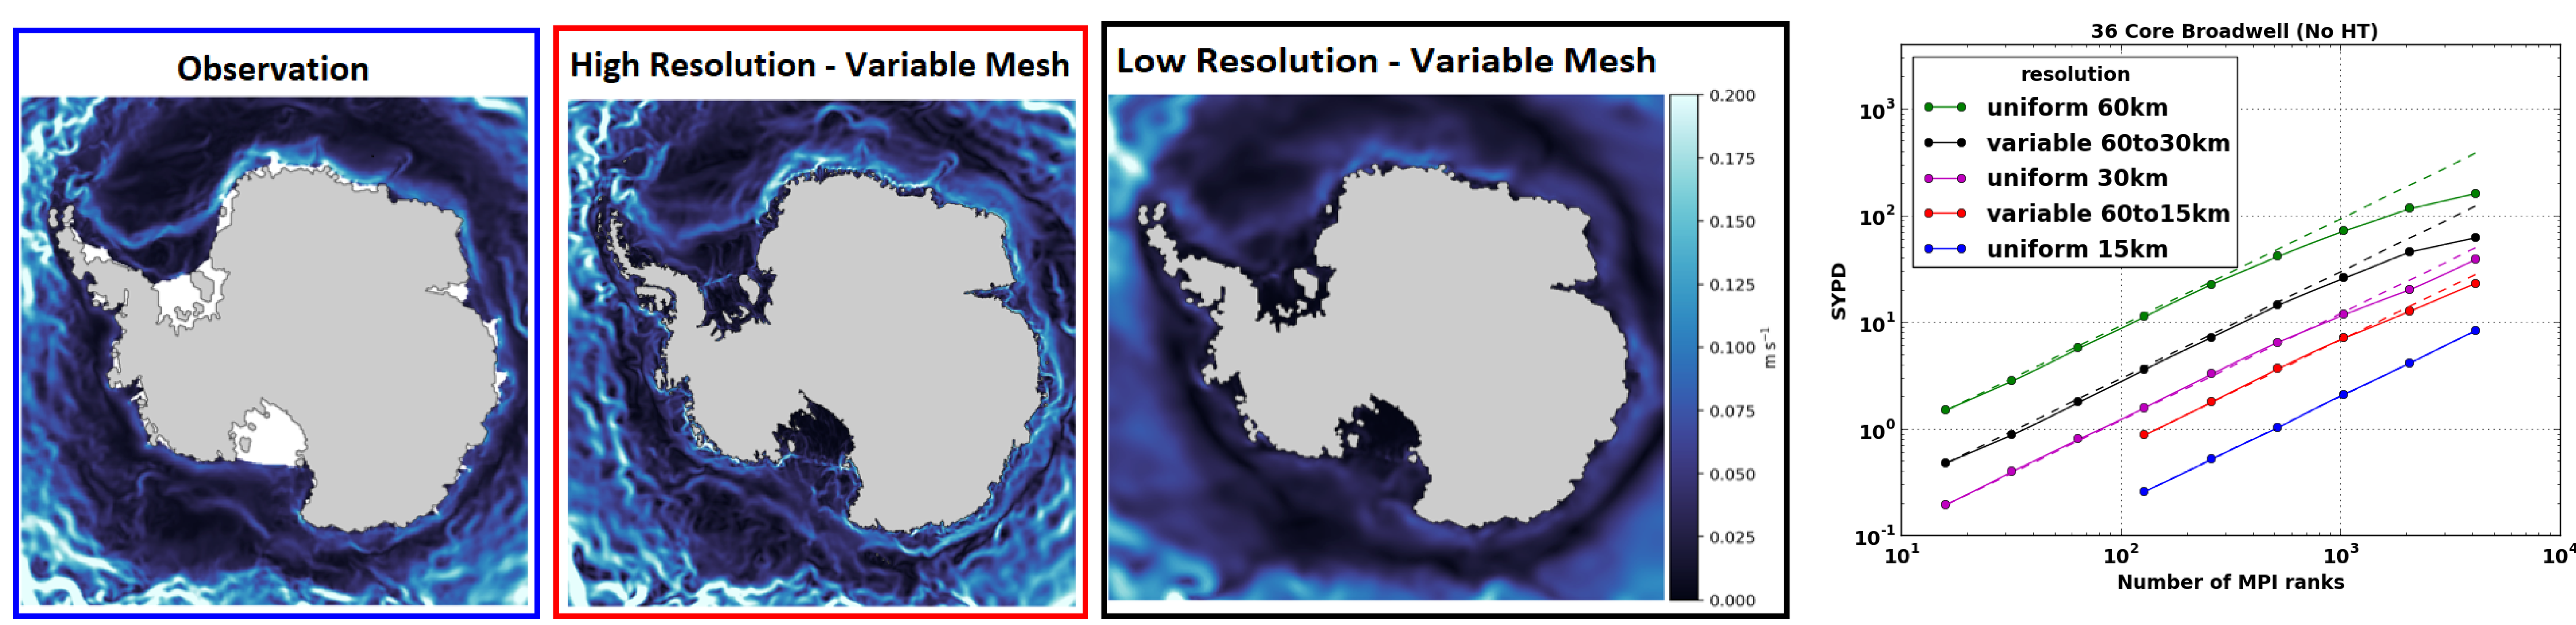

Figure: Sea Surface Velocity Magnitude: Southern Ocean enhanced high-resolution simulation (red) has higher performance throughput and similar simulation quality compared to the global high resolution observation (blue). Eddy activity and currents are absent in Southern Ocean enhanced low-resolution simulation (black)

Conclusion

- Cell weight factor, proportional to N_{vcells}^{X} with **X** in the range **0–0.5** ensures load balance on depth-dependent routines
- Variable resolution mesh simulates regional fine scale features at good quality at a fraction of the computational cost compared to global high resolution mesh.

Acknowledgements

- This research was supported as part of the Energy Exascale Earth System Model (E3SM) project, funded by the U.S. Department of Energy
- This work used **Grizzly*** at Los Alamos National Laboratory* (**DE-AC52-06NA25396**), and **Cori**+, **Edison**+ at National Energy Research Scientific Computing Center+ (**DE-AC02-05CH11231**)
- We thank the mentors at Parallel Computing Summer Research Internship: Hai Ah Nam, Bob Robey, Kris Garrett

References

[1] Ringler, T., Petersen, M., Higdon, R. L., Jacobsen, D., Jones, P. W., & Maltrud, M. A multi-resolution approach to global ocean modeling Ocean Modelling (2013).
[2] A Fast and Highly Quality Multilevel Scheme for Partitioning Irregular Graphs. George Karypis and Vipin Kumar. SIAM Journal on Scientific Computing (1999).