



RESEARCH ARTICLE

10.1029/2019MS001848

Special Section:

The Energy Exascale Earth System Model (E3SM)

MPAS-Ocean Simulation Quality for Variable-Resolution North American Coastal Meshes

Kristin E. Hoch¹, Mark R. Petersen¹, Steven R. Brus², Darren Engwirda^{3,4}, Andrew F. Roberts², Kevin L. Rosa^{1,5}, and Phillip J. Wolfram²

¹Computational Physics and Methods (CCS-2), Los Alamos National Laboratory, Los Alamos, NM, USA, ²Fluid Dynamics and Solid Mechanics (T-3), Los Alamos National Laboratory, Los Alamos, NM, USA, ³NASA Goddard Institute for Space Studies, New York, NY, USA, ⁴Center for Climate Systems Research, Columbia University, New York, NY, USA, ⁵Graduate School of Oceanography, University of Rhode Island, Narragansett, RI, USA

Key Points:

- Regionally refined MPAS-Ocean simulations are comparable to global high-resolution simulations for numerous metrics
- Variable-resolution unstructured Voronoi meshes created using JIGSAW are evaluated for quality
- Simulation quality remains high for steep resolution transitions and intentionally degraded meshes

Correspondence to:

K. E. Hoch, kehoch@lanl.gov

Citation:

Hoch, K. E., Petersen, M. R., Brus, S. R., Engwirda, D., Roberts, A. F., Rosa, K. L., & Wolfram, P. J. (2020). MPAS-Ocean simulation quality for variable-resolution North American coastal meshes. *Journal of Advances in Modeling Earth Systems*, 12, e2019MS001848. <https://doi.org/10.1029/2019MS001848>

Received 7 AUG 2019

Accepted 19 FEB 2020

Accepted article online 26 FEB 2020

Abstract Climate model components utilizing unstructured meshes enable variable resolution, regionally enhanced simulations within global domains. Here we investigate the relationship between mesh quality and simulation statistics using the JIGSAW unstructured meshing library and the Model for Prediction Across Scales-Ocean (MPAS-Ocean) with a focus on Gulf Stream dynamics. In the base configuration, the refined region employs 8 km cells that extend 400 km from the coast of North America. This coastal-refined region is embedded within a low-resolution global domain, with cell size varying latitudinally between 30 and 60 km. The resolution transition region between the refined region and background mesh is 600 km wide. Three sensitivity tests are conducted: (a) The quality of meshes is intentionally degraded so that horizontal cells are progressively more distorted; (b) the transition region from high to low resolution is steepened; and (c) resolution of the coastal refinement region is varied from 30 to 8 km. Overall, the ocean simulations are shown to be robust to mesh resolution and quality alterations. Meshes that are substantially degraded still produce realistic currents, with Southern Ocean transports within 0.4% and Gulf Stream transports within 12% of high-quality mesh results. The narrowest transition case of 100 km did not produce any spurious effects. Refined regions with high-resolution produce eddy kinetic energy and sea surface height variability that are similar to the high-resolution reference simulation. These results provide heuristics for the design criteria of variable-resolution climate model domains.

Plain Language Summary Computer simulations used to study the ocean use grids that cover the ocean's surface, and computations are conducted in each grid cell. The smaller these cells are, the more detailed the simulation is, but simulations with more cells are more expensive to run. We experiment with adding small cells in the region of interest, in this case the North American coast, and larger cells in the rest of the ocean. We conducted three series of tests and looked at the effects on the Gulf Stream, an ocean current off the East Coast of North America. (a) We wanted to know how much adding these small cells improved the simulation. We changed the size of the coastal cells from 30 km wide (less detailed) to 8 km wide (more detailed). Smaller cells improved the results along the North American coast. (b) We cannot go straight from the small to large cells and must have intermediate-sized cells in between. We experiment with different numbers of these intermediate transition cells. The more intermediate cells we added, the better the results were. (c) We wanted to know whether the cells have to be a regular shape in order to get good results. We experimented with irregular cell shapes. The irregular cells produced results that were very similar to the regular cells.

1. Introduction

Climate models based on unstructured horizontal meshes have matured in recent years. Unstructured global simulations of historical periods compare well when validated against observations and against other future climate projections (Golaz et al., 2019; Petersen et al., 2019; Scholz et al., 2019). Unstructured meshes offer great freedom in placing resolution in the areas of interest for regionally refined simulations and also suggest the possibility of improving global simulation quality with targeted areas of high resolution. However, modelers now have a dizzying array of choices to make in designing their meshes, compared to the limited

©2020. The Authors.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

variations of stretched aspect ratio in latitude-longitude-type quadrilateral grids. Furthermore, the role of regional refinement strategies on simulation quality is currently largely unknown.

There is a pressing need for constraints on mesh design and model configuration criteria that are informed by how local resolution affects simulation quality. However, time constraints and available computational resources generally allow only a limited number of configurations to be rigorously tested. In this study, we explore the role of mesh design and quality on various ocean simulation metrics using the Model for Prediction Across Scales (MPAS) (Ringler et al., 2013) with the goal of providing guidance on the design of meshes for variable-resolution climate models.

The generation of high-quality unstructured meshes for General Circulation Models (GCMs) is a challenging problem, and a new generation of mesh creation tools has been developed to satisfy the needs of high-resolution unstructured mesh models. This paper documents the use of JIGSAW (Engwirda, 2017) to produce optimized spherical Voronoi/Delaunay meshes for use with MPAS. MPAS-Ocean and MPAS-Sea Ice are components of the Department of Energy's Energy Exascale Earth System Model (E3SM) (<https://e3sm.org>) (Golaz et al., 2019; Petersen et al., 2019).

An ensemble of horizontal meshes was investigated using the Coastal United States "Plus" (CUSP) configuration, which is designed to enhance the resolution of coastal regions of North and Central America plus Hawaii. Three case studies were performed: one where global mesh quality is intentionally degraded; a second where the resolution transition width is varied in the CUSP mesh; and a third where the coastal-refined region is tested at a number of resolutions. In each case, a family of meshes was generated, and the results of a 10-year simulation were analyzed, allowing for the convergence of model metrics to be assessed with respect to perturbations in the underlying grid and model configuration. Analysis was focused on the Gulf Stream. Accurately resolving the Gulf Stream had been a persistent challenge in the MPAS low-resolution model. The Gulf Stream also crosses the transition zone, allowing the effect of changing resolution on ocean currents to be tested. Using these data, modelers can assess which mesh characteristics are most important for the needs of their application and inform their choices for the design of future configurations.

We aim to highlight the impact of various mesh characteristics on simulation quality and to document how different choices in mesh design feed back onto the simulated state. We focus on the geometric "quality" of a mesh, its rate of transition from regions of low to high resolution, and the placement of high-resolution near-energetic boundary currents and areas of interest. The configurations used in this paper enhance resolution of the North American coastal region, but the aim is to provide general guidelines that may be applied to the design of any variable-resolution mesh.

This paper is structured as follows. Section 2 reviews the state of variable-resolution meshes on ocean modeling. Section 3 introduces MPAS-Ocean, JIGSAW, and the details of the meshes created for this work. Section 4 presents the analysis of global simulations for the three sensitivity studies. Based on this evidence, the paper concludes with recommendations for mesh generation criteria in section 5.

2. Background

There now exists a growing selection of unstructured mesh models that are used for various global and regionally focused forecasts and analyses. This includes MPAS (Ringler et al., 2013), FESOM (Danilov et al., 2017; Wang et al., 2014), ICON (Korn, 2017), FVCOM (Chen et al., 2003), SCHISM (Zhang et al., 2016), SLIM (Kärnä et al., 2013), and Fluidity (Davies et al., 2011). Mesh creation tools such as Shingle 2.0 have been developed to produce high-quality reproducible meshes efficiently (Candy & Pietrzak, 2018). Models differ in the arrangement of variables on the underlying computational grid and in the numerical techniques employed, with both unstructured triangle- and polygon-based finite-volume and finite-element type discretization schemes adopted in various frameworks. As such, different approaches to the construction and optimization of the models' underlying unstructured meshes have been explored, including techniques based on Centroidal Voronoi Tessellation (CVTs) (Jacobsen et al., 2013; Yang et al., 2018), optimization via optimal transport (McRae et al., 2018; Weller et al., 2016), as well as triangulation-based refinement schemes (Lambrechts et al., 2008a; Remacle & Lambrechts, 2018). A mesh generation tool developed by Lambrechts et al. (2008b) refines according to bathymetry, bathymetry gradients, and distances from coasts. In the context of MPAS-Ocean, the numerical scheme requires that the mesh define a highly regular, orthogonal

tessellation, constraining grid generation choice to algorithms that can generate optimized Voronoi-type meshes (Jacobsen et al., 2013).

Variable resolution is advantageous in situations where highlighting a region may help to correct a bias or resolve a dynamic condition. In many cases, the resolved region will also be the focus of the investigation, but resolution can also be placed to correct a bias that is impacting a global simulation. Variable resolution can serve as a replacement for nested grids, with the advantage that variable resolution can be applied in more complex configurations and is more integrated with the global simulation (Biastoch et al., 2018; Hagos et al., 2013). Nested grids have the advantages of more easily implemented variable time stepping and simplified grid geometry. However, nesting introduces challenges with conservation, coupling, interpolation, and noise control (Debreu & Blayo, 2008).

Meshes in which the resolution varies as a function of latitude have been used to compensate for the changing Rossby radius with latitude. This approach is used in the standard high-resolution MPAS mesh (Petersen et al., 2019). Variable-resolution meshes are designed to improve the dynamics of a particular region or process and also to provide good global dynamics. Variable resolution meshes may refine particular regions, for example, the Arctic Ocean (Wang et al., 2018) or a coastal region (Androsov et al., 2019). Variable resolution has been applied in regional ocean models to capture a wide range of scales, from tens of kilometers to tens of meters. SCHISM and FVCOM have been applied in variable-resolution cases to place high resolution in estuaries and straits, where narrow channels and complex bathymetry must be properly represented, for example, the Chesapeake Bay (Ye et al., 2018) and the Canadian Archipelago (Zhang et al., 2016). It has also been applied for a variety of coastal processes, for example, storm surge (Fernández-Montblanc et al., 2019; Wang et al., 2014) and nutrient distribution (Tian et al., 2014).

Resolution can also be placed based on a particular parameter. For example, FESOM uses meshes that refine to the local Rossby radius (Sein et al., 2017), a more sophisticated approach than refining based on latitude alone. FESOM also uses meshes that refine according to eddy variability. This approach is much less computationally expensive than refining based on Rossby radius but has been shown to improve deep ocean biases and Gulf Stream separation (Rackow et al., 2019). FESOM's meshes which refine based on sea surface height (SSH) variability are useful for capturing boundary currents (Biastoch et al., 2018). A configuration which used high resolution over areas of high SSH variability, areas upstream of the separation of mid-latitude jets, and in the Nordic Seas improved Gulf Stream separation and biases in the Northwest Corner (Sein et al., 2016).

Because of the computational cost and complexity of global simulations, the majority of variable-resolution tests have been performed on idealized or simplified domains. For example, in order to eliminate the effects of continental geography, many tests have used aquaplanet configurations (Abiodun et al., 2008; Hagos et al., 2013; Lorant & Royer, 2001; Rauscher & Ringler, 2014; Rauscher et al., 2012; Zhao et al., 2016). Others have used two-dimensional domains (Düben & Korn, 2014). These simplified domains can demonstrate the effects of mesh resolution independent of other variables. Additionally, atmospheric variable-resolution simulations can inform choices in ocean domains (Abiodun et al., 2008; Düben & Korn, 2014; Park et al., 2014; Rauscher & Ringler, 2014; Zarzycki et al., 2015; Zhao et al., 2016). However, mesh resolution and design consequences on more realistic simulations are still largely unknown, even though the use of variable resolution in realistic simulations is becoming more widespread.

While mesh design is still a developing field, the literature points to several important considerations. In the past, parameter values for sub-grid scale physics were typically tuned for each resolution. Now, for variable-resolution meshes, parameterization schemes must work well across the span of grid cell sizes. Another consideration is that variable-resolution results compared against uniform high-resolution simulations may not necessarily be comparable near mesh transition regions. For example, a current flowing from a non-eddy permitting to an eddy-permitting region may not immediately develop eddies. Instead, eddies will develop downstream of the beginning of the high-resolution region once perturbations have time to evolve (Danilov & Wang, 2015). A similar result was found in atmospheric variable-resolution aquaplanet simulations, in which precipitation error was decreased in the eastern (downstream) section of the high-resolution region but not in the western (upstream) section (Hagos et al., 2013).

A high-resolution region will also have effects on the rest of the domain. Most obviously, a high-resolution region will have an effect immediately downstream, as the increased variability of the high-resolution

region is carried into the low-resolution region (Danilov & Wang, 2015). Changes to dynamics within the high-resolution region can propagate to other global processes (Hagos et al., 2013; Lorant & Royer, 2001; Sakaguchi et al., 2016; Sein et al., 2017). Conversely, the impact of the global domain on the high-resolution region is also important. A high-resolution region can decrease local error but will have a limited impact on processes that are due to causes outside the high-resolution region (Zarzycki et al., 2015).

3. Methods

3.1. The Model for Prediction Across Scales-Ocean (MPAS-Ocean)

The Model for Prediction Across Scales (MPAS) is an open source framework that provides common functionality for climate model components on unstructured meshes. This includes a mesh specification, decomposition of variables across processors, parallel input and output specified in a run time streams file, timers, and error handling. Finite volume operators were developed for Voronoi tessellations in Ringler et al. (2010) for the shallow water equations using mimetic methods to guarantee that mass, velocity, and potential vorticity evolve in a consistent and compatible manner.

MPAS-Ocean solves prognostic equations for momentum, thickness (volume), and tracers using these operators (Ringler et al., 2013) and can be run using both regular and unstructured meshes on Cartesian and spherical domains. The time stepping is split-explicit, where the 2D barotropic equations are sub-cycled within 3D baroclinic time steps. Both parts use a second-order predictor-corrector method based on Higdon (2005), as detailed in Appendix A5 of Ringler et al. (2013). Advection uses the flux-corrected transport scheme (Skamarock & Gassmann, 2011), which blends high- and low-order fluxes to preserve monotonicity and is second-order accurate on variable-resolution meshes (Ringler et al., 2013). The simulations presented here use a z-star vertical coordinate, which is the standard choice for global simulations. The MPAS-Ocean vertical coordinate is designed within an Arbitrary Lagrangian-Eulerian (ALE) framework (Petersen et al., 2015; Reckinger et al., 2015). Simulations typically include 60, 80, or 100 vertical layers, which vary from 2 m thick at the surface to 150 m thick at a depth of 5,000 m.

The vertical mixing scheme is the K-Profile Parameterization (KPP) (Van Roekel et al., 2018), calculated in the CVMix library (<https://github.com/CVMix/CVMix-src>, <https://doi.org/10.5281/zenodo.1000800>) and applied implicitly. The horizontal mesoscale eddy parameterization is Gent-McWilliams thickness advection (Gent & McWilliams, 1990), applied to variable-resolution meshes with a coefficient of $600 \text{ m}^2 \text{ s}^{-1}$ at grid cells larger than 30 km and tapering linearly to zero between 30 and 20 km. Viscosity (del-2) and hyperviscosity (del-4) are applied to the momentum equation with coefficients that depend on the grid cell size as

$$\nu_2 = 1000 [m^2 s^{-1}] \frac{\Delta x}{30 [km]}, \quad (1)$$

$$\nu_4 = 1.2e11 [m^4 s^{-2}] \left(\frac{\Delta x}{30 [km]} \right)^3, \quad (2)$$

respectively, where Δx is the horizontal grid cell width. The coefficients were tuned in Petersen et al. (2019) to be as small as possible, while ensuring that dissipation is sufficient for stability and to prevent grid scale noise. The addition of hyperviscosity removes energy more strongly at the highest wave numbers and allows a smaller viscosity coefficient. No horizontal diffusion is explicitly applied to the tracers.

For this study MPAS-Ocean was run with the same choice of parameters as typical global simulations, such as those presented in Petersen et al. (2019). One exception is that this study used the stand-alone version of MPAS-Ocean, rather than the coupled E3SM code. Stand-alone mode applies idealized, constant atmospheric forcing, where wind forcing is averaged over a 65-year CORE (coordinated ocean-ice reference experiments) cycle (Griffies et al., 2009). The choice to use stand-alone mode was made for two reasons. First, it considerably simplified the required setup, streamlining the work required for a large parameter study. Achieving realistic climatological results would require a lengthy spin-up process and longer run time, neither of which were possible for this number of simulations. Second, the idealized forcing simplified the conditions of the simulations, making it easier to evaluate any numerical effects of the meshes. Because this study used both new variable-resolution meshes and a new mesh creation tool, it was important to test simplified domains before running the meshes with the complexities inherent in fully coupled E3SM simulations.

The simulation is spun up for 1 year from an initial climatology of Polar Science Center Hydrographic Climatology, version 3 (PHC3.0, Steele et al., 2001). Surface salinity and temperature restoring to yearly averaged PHC3.0 is conducted with a piston velocity of 1.37 m day^{-1} to represent surface fluxes. Sea-ice is not included in these simulations. Simulations with more realistic atmospheric forcing (6-hourly CORE winds and surface fluxes) and active sea ice have been run within E3SM using the coastal-refined mesh (CUSP8) are currently underway and will be presented in a future publication.

3.2. JIGSAW Mesh Generation

JIGSAW is an unstructured meshing library designed to generate high-quality grids for computational simulation, with a focus on constructing optimized Voronoi-type grids for unstructured mesh GCM's. JIGSAW is a hybrid algorithm that combines both Delaunay-refinement and Voronoi optimization type approaches to enable the rapid generation of very high quality, high-resolution Voronoi/Delaunay meshes on the sphere. The Delaunay mesh is composed of triangular cells, and the Voronoi mesh is the dual composed of polygonal cells. A key advantage of this combined strategy is efficiency and guaranteed mesh quality. Previous mesh generation methods used in MPAS (Jacobsen et al., 2013) used an iterative Lloyd's method and were extremely slow.

With JIGSAW, highly optimized, large-scale variable-resolution Voronoi-type meshes can be generated in the order of minutes, allowing model users to easily create and explore a range of alternative configurations, investigate mesh quality and resolution dependence, and tailor the overall mesh and model configuration to their simulation needs. This capability was exploited in the present study to design and assess a range of coastal-enhanced MPAS-Ocean configurations and to explore various model/mesh feedbacks.

Meshes can be generated in local two-dimensional domains and over general spheroidal surfaces. Mesh resolution can be adapted to follow complex user-defined metrics, including topographic contours, solution profiles, and/or coastal features. This flexibility enables the construction of complex, variable-resolution model configurations, offering enhanced simulation fidelity in regions of interest or importance.

Given a particular geometry definition and resolution specification, JIGSAW proceeds to assemble the unstructured mesh incrementally—first creating a conforming Delaunay triangulation of the domain using a “frontal” Delaunay-refinement strategy (Engwirda & Ivers, 2016), before optimizing the resulting Voronoi/Delaunay tessellation using Optimal Delaunay Tessellation (ODT) type techniques (Chen & Holst, 2011; Engwirda, 2017). The final mesh is guaranteed to consist of high-quality triangular and polygonal cells that facilitate a locally orthogonal unstructured C-grid staggering. The final meshes are heavily optimized, typically satisfying the stringent mesh quality requirements imposed by the TRiSK (Thuburn Ringler Skamarock Klemp) discretization scheme (Ringler et al., 2010) used in MPAS-Ocean.

For TRiSK-based schemes, a complex array of geometrical and topological constraints must be satisfied (Engwirda, 2018), requiring tessellations to be orthogonal, centroidal, well-centered, and smoothly varying. These criteria require that the vertices of the triangular and polygonal grid cells lie close to the centroids of their enclosing control-volumes, that the staggered Voronoi and Delaunay edges intersect near their midpoints, that the Delaunay triangles contain their own circumcenters, and that the cell angles and edge-lengths be “nicely” distributed with respect to the desired mesh resolution constraints. Satisfying such criteria is nontrivial, and failure to do so has been shown to impact on the asymptotic accuracy and stability of the underlying numerical scheme (Peixoto, 2016) in idealized cases.

The expected accuracy of the TRiSK formulation is thus a function of both the geometry and topology of the mesh and can be quantified by considering the nature of the discrete gradient, divergence, curl, and interpolation operators used to discretize the continuous PDE's (Engwirda, 2018; Ringler et al., 2010). Based on theoretical analysis, it is expected that the accuracy of TRiSK is maximized (achieving quasi second-order scaling) only for “perfect” tessellations consisting of regular hexagons and equilateral triangles. For general unstructured meshes incorporating irregular and/or deformed polygonal and triangular cells, numerical accuracy is expected to degrade—leading to quasi first-order behavior in many practical configurations (Peixoto, 2016). The goal of mesh optimization is to construct a tessellation that serves to minimize these numerical errors, thus maximizing the quality of the resulting simulation.

A key question in the current study is to assess what impact mesh quality has on practical MPAS-Ocean simulations and to define an associated set of “best practice” guidelines for mesh generation. To this end, an




Figure 1. Cell quality of the degraded meshes. A small region of the mesh is shown. Cell quality is the ratio of the smallest to largest sides of a cell, 1.0 being a perfect polygon.

“ensemble” of meshes was considered in the current work—exploring the impact of different mesh quality perturbations and variable-resolution designs on the characteristics of spun up ocean simulations.

3.3. Meshes and Simulations

All the meshes used are based on two base configurations, a global low-resolution mesh and a mesh with refinement along the coast of North America. The global low-resolution mesh, EC60to30, varies from 30 km resolution at the equator and poles to 60 km resolution at the mid-latitudes and uses 100 vertical layers.

The base EC60to30 mesh created using JIGSAW was compared against the EC60to30-E3SM-V1 mesh created using a parallel Lloyd’s algorithm (Jacobsen et al., 2013), which was used in previously published E3SM simulations (Golaz et al., 2019; Petersen et al., 2019). Images of the two EC60to30 meshes can be seen in the first two panels of Figures 1 and 2, which show two different metrics for measuring cell quality. Figure 2 shows the percent change between the size of neighboring cells, and Figure 1 shows close-up images of the mesh and the ratio of the smallest to largest sides of the cells. These metrics show the different strategies used by each of the mesh creation methods. In order to cover the sphere, the mesh must deviate from regular hexagons. E3SM-V1 spreads these imperfections between large numbers of cells, resulting in smooth regions of lower-quality cells. JIGSAW concentrates the imperfections into “seams” of low-quality cells separating regions of very high quality cells.

The second base mesh is the North American refined mesh, created to investigate processes affecting North American coastal regions at high resolution while avoiding the cost of running a global high-resolution model. In addition to the improvements in the dynamics of the Gulf Stream investigated in this study, using the CUSP8 mesh will allow improved simulation of a variety of coastal processes around North America. The CUSP8 mesh (Coastal United States “Plus” with 8 km coastal resolution) has high resolution along the Atlantic and Pacific coasts from Central America to the Arctic, with additional high resolution in the Caribbean and around Greenland, Hawaii, and the Bering Strait (see Figure 3). The CUSP8 mesh is built on top of a background low-resolution EC60to30 mesh. It uses 80 vertical layers.

In the CUSP8 mesh, the transition between the high-resolution region and the background mesh begins 400 km off the coast. This was chosen so that the high-resolution region encompassed the flow of the Gulf Stream along the coast and other important coastal processes. The transition region is 600 km wide and




Figure 2. Percent change in grid cell area between neighboring cells.




Figure 3. The Coastal United States “Plus” mesh (CUSP8) on the left and the North Atlantic refined mesh (NA8) on the right. The white areas show the 8 km high-resolution regions. Blue shows the background EC60to30 low-resolution mesh, with 30 km resolution at the tropics and poles (light blue) and 60 km resolution in between (dark blue).

follows the following functions:

$$W = 0.5 \left(\tanh \frac{D - D_{start} - 0.5D_{width}}{0.2D_{width}} + 1 \right), \quad (3)$$

$$C = C_{coast} (1.0 - W) + C_{back} W, \quad (4)$$

where W is the weight, D is the distance from the coast, D_{start} is the distance from the coast where the transition region begins, and D_{width} is the transition width. The final cell width, C , shown in Figures 3 and 4, is simply a linear combination of the coastal and background cell widths, C_{coast} and C_{back} .

In addition to these two base meshes, a mesh with 8 km resolution spanning the full North Atlantic basin (NA8) was created. Like the CUSP8 mesh, it was built on a background EC60to30 mesh (see Figure 3). A global high-resolution simulation was not feasible for this study, but the NA8 mesh provides high resolution within the region of interest in the North Atlantic, providing a benchmark for the performance of the CUSP8 mesh.

All resolutions are first created on a full sphere, and then continental and island land cells are culled if the cell center is within a high-resolution coastline defined by connected points (<http://www.naturalearthdata.com>). The bathymetry is obtained by interpolation of the ETOPO2 2-Minute Gridded Global Relief Dataset available from the National Geophysical Data Center (section 4.1; Ringler et al., 2013). Partial bottom cells are used for a better representation of the bathymetry. All domains presented in this study use this standard method of initializing coastlines and bathymetry, with the highest resolution data available. This means that regions with higher mesh resolution also have finer coastal and depth features.

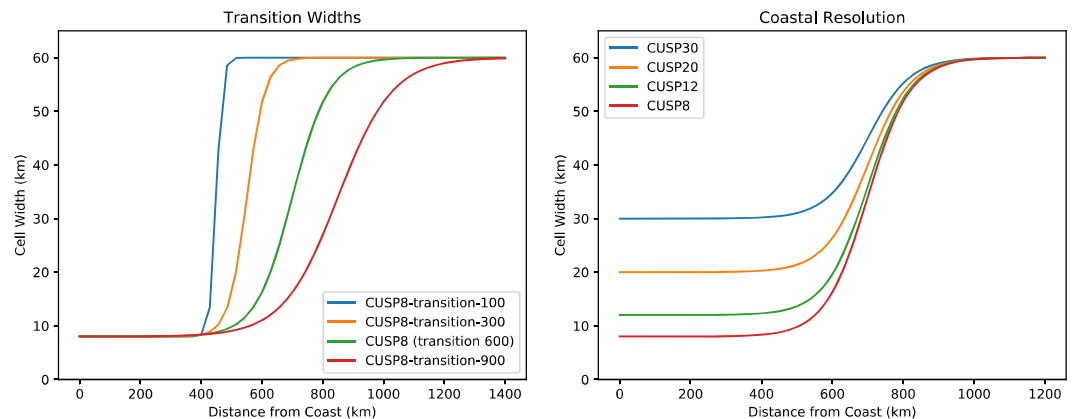


Figure 4. Plots of the transition function (equation (4)) for the transition width study (left) and the convergence study (right). The background resolution plotted is 60 km; however, the background resolution varies from 30 to 60 km depending on latitude.




Figure 5. A view of the East Coast showing the different transition widths used. The transition begins at 400 km off the coast for all transition widths. Note that the size of the higher-resolution region is expanded with a wider transition.

In order to ensure that all the meshes could be compared, EC60to30 simulations were run in each vertical configuration: 60, 80, and 100 layers. All three EC60to30 meshes performed similarly in terms of kinetic energy (KE), sea surface height (SSH), eddy kinetic energy (EKE), and sea surface height root mean squared (SSH RMS) (see Figure A2).

Three studies were performed to investigate mesh features and their effects on simulation quality. All the simulations were run for 10 years, with analysis performed on the last 9 years.

The first study uses the EC60to30 mesh to examine the effect of poor mesh quality on simulations. Meshes were intentionally degraded, producing poor quality cells. Variable-resolution meshes by necessity contain distorted cells within the transition regions. This is a particular concern when designing complex meshes such as the CUSP8 mesh that have large variations in resolution and relatively narrow transition regions. Because of the difficulty of decoupling the effects of poor cell quality from the effects of a change in resolution, the effect of poor cell quality on simulations was investigated using EC60to30 meshes with cell quality degraded globally.

Table 1
Simulation Parameters

Study	Mesh name	Refined resolution (km)	Number of cells (thousands)	Vertical layers	Transition width (km)	Degradation factor	Time step (min:s)	Barotropic step (min:s)
Reference meshes	EC60to30	None	236	100	None	None	30:00	1:00
	CUSP8	8	649	80	600	None	7:30	00:15
	NA8	8	842	80	600	None	7:30	00:15
	EC60to30-E3SM-V1	None	235	100	None	None	20:00	1:00
Degraded meshes	EC60to30 (not degraded)	None	236	100	None	None	30:00	1:00
	EC60to30-degraded-0.25	None	237	100	None	0.25	30:00	1:00
	EC60to30-degraded-0.50	None	248	100	None	0.50	20:00	0:40
	EC60to30-degraded-0.75	None	338	100	None	0.75	2:00	0:06
Transition width	CUSP8-transition-900	8	700	80	900	None	7:30	00:15
	CUSP8 (transition 600)	8	649	80	600	None	7:30	00:15
	CUSP8-transition-300	8	603	80	300	None	7:30	00:15
	CUSP8-transition-100	8	574	80	100	None	7:30	00:15
Coastal resolution	CUSP8	8	649	80	600	None	7:30	00:15
	CUSP12	12	414	80	600	None	12:00	00:24
	CUSP20	20	295	80	600	None	20:00	00:40
	CUSP30	30	256	80	600	None	30:00	1:00

Note. The reference simulations, EC60to30 and CUSP8, are bold. The varied parameter for each study is in italics. Time steps are proportional to the smallest resolution present in the mesh.



Figure 6. Shaded region indicates the area in which SSH RMS and EKE averages were computed. Yellow sections show the locations of transects along the Gulf Stream.

A mesh degradation heuristic was developed to systematically reduce the quality of meshes, perturbing the position of vertices and updating topology to effectively “de-optimize” the overall structure of a given mesh and degrade the shape of its cells. Care was taken to ensure that degraded meshes inherited the large-scale properties of their parent grids, adhering to variations in resolution and matching cell counts exactly. The kernel of the degradation operation consisted of randomly perturbing a subset of vertices toward the centroid of their largest neighboring triangle. By controlling the magnitude of the average relative vertex perturbation, the notion of a “ β -degraded” mesh was introduced—a 0.5-degraded mesh would re-position vertices (on average) halfway between their current position and the neighboring centroid location. Mesh topology was updated following the re-positioning of vertices to ensure that the orthogonality of the mesh was preserved. Starting from a fully optimized initial mesh, several iterations of this process were repeated to ensure that degraded grids were sufficiently randomized.

Three degraded meshes were created, EC60to30-degraded-0.25, EC60to30-degraded-0.50, and EC60to30-degraded-0.75, with larger degradation fractions indicating a more degraded mesh. Figure 1 shows the mesh quality of the standard EC60to30 mesh and the degraded meshes.

The second study investigates the effects of the steepness of the transition function in the CUSP8 mesh (equation (3)) by varying the transition width from 100 to 900 km (Figure 4). A 10 km transition was attempted as

well but failed early in the spin-up process. This study was designed to investigate how steep the transition function could be without negatively affecting the simulation quality. In addition to exploring the steepness of the transition function, this study also investigates the impact of the size of the higher-resolution region. Because the beginning of the transition region was kept fixed, the center of the transition region and the beginning of the low-resolution region were closer to the coast for steeper transitions, effectively shrinking the higher-resolution region (see Figure 5).

The third study investigates different coastal resolutions ranging from 8 km (CUSP8) to 30 km (CUSP30) in order to explore the improvements in dynamics with increased resolution. Resolutions were chosen to span the range between the highest resolution in the MPAS high-resolution model, which varies from 8 to 16 km, and low-resolution model, which varies from 30 to 60 km. The computational performance of the meshes was also examined in order to give a better sense of the trade-off between higher resolution and higher simulation costs. These meshes were compared against the EC60to30 and NA8 meshes. Ideally, the CUSP8 mesh would show dynamics comparable to the NA8 mesh within the high-resolution region with a much lower cost than a global high-resolution mesh. This study is designed to examine the degree to which we can produce the correct dynamics on variable-resolution meshes. It is likely that the CUSP8 mesh will not be able to fully recover the dynamics of a high-resolution simulation and that further modifications to the mesh will be required.

Table 1 shows the parameter values used for each simulation. These values were chosen based on the highest resolution region of the simulation. The EC60to30-degraded-0.50 and EC60to30-degraded-0.75 meshes had to be run at a smaller time step than the standard EC60to30 meshes due to the smaller cell sizes introduced by the degradation process. All simulations were started from a short initialization. This initialization process maintains stability using Rayleigh damping and small time steps after the run is started from rest. It is not intended to produce an equilibrium state. All runs used a 7 day initializations except the EC60to30-E3SM-V1, EC60to30-degraded-0.50, and EC60to30-degraded-0.75 meshes. The EC60to30-E3SM-V1 mesh used a 21 day initialization. The EC60to30-degraded-0.50 and EC60to30-degraded-0.75 meshes required longer initializations and were spun up to a different point because of the smaller time step required.

4. Results and Discussion

The analysis focuses on the Gulf Stream because it is the most prominent feature within the high-resolution region of the CUSP simulations. The Gulf Stream also crosses out of the high-resolution region, allowing

Table 2

The Average Transport in Sverdrups Through Transects for Years 2–10, Followed by Standard Deviation for Simulations and Error for Observations

	Florida-Cuba	Florida-Bahamas	Cape Hatteras	New Jersey	Drake Passage	Tasmania-Ant	Africa-Ant
Observation	31.0 ± 1.5	31.5 ± 1.5	87.8 ± 17.3	94.5	173.0 ± 10	157 ± 10	150.0 ± 30
CUSP8	16.43 ± 1.21	19.17 ± 1.13	47.52 ± 16.67	22.87 ± 30.44	174.42 ± 2.11	190.51 ± 2.87	174.49 ± 2.04
NA8	17.46 ± 1.28	20.83 ± 1.35	37.66 ± 12.33	46.99 ± 15.45	174.29 ± 1.60	188.58 ± 2.37	173.74 ± 1.50
CUSP8-transition-100	16.45 ± 0.93	19.04 ± 0.85	56.86 ± 12.56	8.45 ± 29.96	174.12 ± 1.56	189.67 ± 2.24	174.79 ± 1.44
CUSP8-transition-300	16.95 ± 1.19	19.50 ± 0.99	49.63 ± 18.68	26.51 ± 34.42	172.50 ± 2.27	187.14 ± 3.10	172.30 ± 2.17
CUSP8-transition-900	17.20 ± 1.12	19.66 ± 1.11	45.48 ± 13.72	49.51 ± 23.45	176.22 ± 1.31	191.36 ± 1.98	176.41 ± 1.23
CUSP12	15.66 ± 0.99	17.89 ± 0.99	61.07 ± 19.56	10.70 ± 33.87	171.01 ± 2.14	186.81 ± 2.93	171.25 ± 2.06
CUSP20	14.76 ± 0.72	16.37 ± 0.74	64.01 ± 14.37	-1.37 ± 17.20	170. ± 1.97	186.66 ± 2.91	170.82 ± 1.99
CUSP30	13.62 ± 0.41	14.89 ± 0.47	34.91 ± 1.19	25.47 ± 1.50	173. ± 1.88	187.82 ± 2.64	171.87 ± 1.80
EC60to30-degraded-0.25	12.31 ± 0.50	14.84 ± 0.52	40.49 ± 0.70	23.89 ± 1.00	175.20 ± 1.95	190. ± 2.69	175.58 ± 1.90
EC60to30-degraded-0.50	10.75 ± 0.39	11.56 ± 0.39	44.56 ± 0.61	24.92 ± 0.69	173.97 ± 1.68	189.82 ± 2.45	173.73 ± 1.58
EC60to30-degraded-0.75	10.91 ± 0.37	11.25 ± 0.36	44.58 ± 0.65	25.43 ± 0.83	172.18 ± 1.41	187.95 ± 2.16	172.99 ± 1.33
EC60to30	10.13 ± 0.40	12.55 ± 0.53	40.84 ± 0.73	22.93 ± 1.14	172.81 ± 2.37	188.21 ± 3.10	173.52 ± 2.36
EC60to30-E3SM-V1	10.51 ± 0.51	10.57 ± 0.45	43.07 ± 0.68	23.22 ± 0.53	173.18 ± 1.79	188.85 ± 2.55	172.30 ± 1.83

Note. See Figure 7 for plots of the data and Figure 6 for a map of the Gulf Stream transects. Observational references: Florida-Cuba (Johns et al., 2002), Florida-Bahamas (Johns et al., 2002), Cape Hatteras (Halkin & Rossby, 1985), New Jersey (Rossby et al., 2014), Drake Passage (Donohue et al., 2016), Tasmania-Ant (Ganachaud & Wunsch, 2000), and Africa-Ant (Ganachaud & Wunsch, 2000).

the effect of the transition in resolution to be investigated. The sea surface height, kinetic energy, sea surface height root mean squared, and eddy kinetic energy were analyzed for all simulations. Transport through transects along the Gulf Stream was calculated (see Figure 6 for a map of the Gulf Stream transects). Transport through Southern Ocean transects were also calculated in order to see if the high-resolution region had an impact on global dynamics (see Table 2 and Figure 7 for the transect results). SSH RMS and EKE were averaged along the Gulf Stream region (see Figure 6). These results are not expected to closely match observations, both because of the idealized forcing used and because of the differences between the sampling techniques used to calculate observational estimates and those used in our calculations. Global analysis was also run looking at global temperature, salinity, SSH, and EKE. However, because of the extremely similar results for all simulations, this paper focuses only on analysis of the areas within and around the high-resolution region. Preliminary results from simulations with realistic climatological forcing are also used to give an indication of how CUSP meshes perform in realistic climate simulations. Further results will follow in subsequent papers.

The comparison of the JIGSAW EC60to30 mesh and the EC60to30-E3SM-V1 mesh showed that they performed very similarly, confirming that the meshes created using JIGSAW produce comparable results to those used in previous MPAS studies (see Figure A1).

4.1. Study 1: Degraded Meshes

Though the degradation factor for the degraded mesh study and the transition widths for the transition width study were chosen independently, the degraded meshes were found to be a good proxy for the transition regions (see Figure 8). The cell quality in the transition region for the 100 km transition width is comparable to the cell quality in the 0.75 degraded mesh, and the cell quality in the transition region for the 900 km transition is comparable to the cell quality in the 0.25 degraded mesh. Thus, the results of the degraded mesh study should also be considered when interpreting the results within the transition regions of the CUSP meshes.

The results of the degraded mesh study are summarized in Figure 9, showing snapshots and averaged distributions of sea surface height and kinetic energy in the CUSP region. Overall, it was found that mesh degradation did not significantly effect the quality of the simulations, with the pattern and magnitude of sea surface height and kinetic energy for the set of degraded meshes and the optimized EC60to30 mesh visually near identical. The more degraded meshes were found to have slightly higher average sea surface height variability and eddy kinetic energy (see Figure 10). Transport through all transects measured showed no significant variation between the degraded meshes (see Figure 7) and the reference EC60to30 configuration.

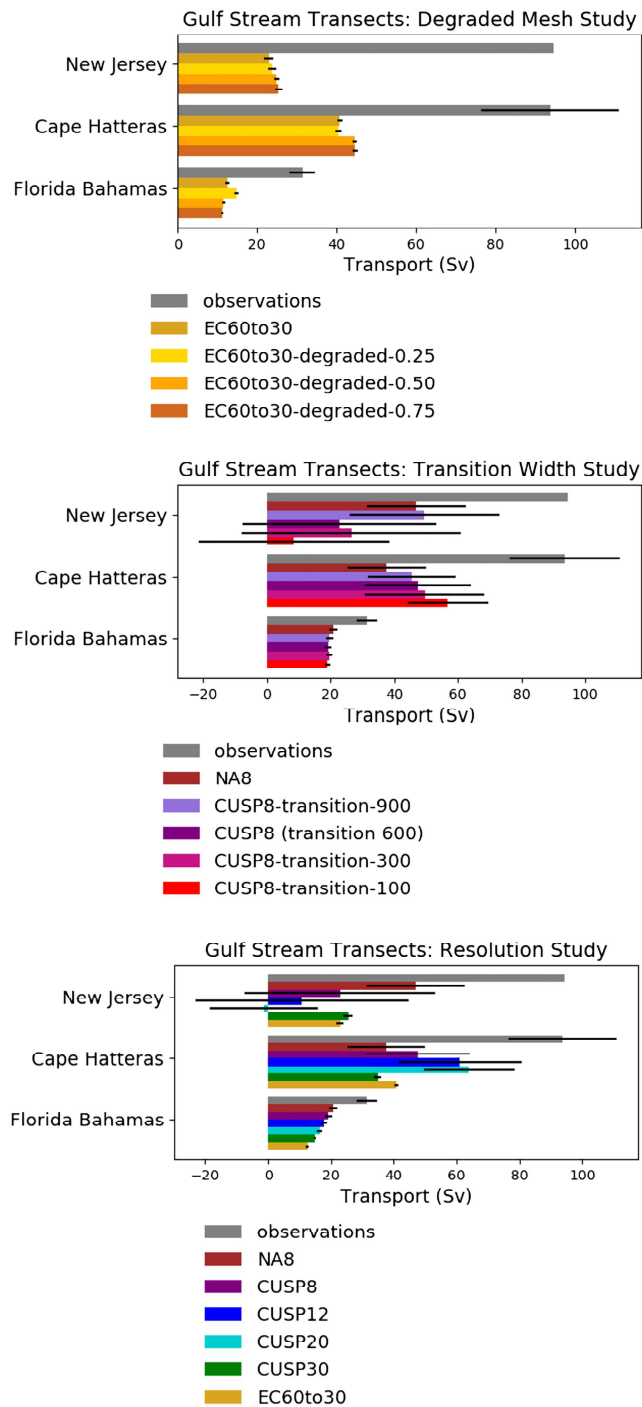


Figure 7. Transport through transects along the Gulf Stream. Table 2 shows the data and Figure 6 shows the locations of the transects.

Overall, it was not found that a reduction in mesh quality had notable adverse effects on the simulations, beyond the need for the use of smaller time steps in highly degraded cases, due to the presence of smaller grid cells. While increasing computation time, the use of smaller time steps does not impact the quality of results. An EC60to30 mesh was run with a 2 min time step (the time step used for the most degraded mesh), and the results were compared against the standard EC60to30 case. There was no apparent effect of the smaller time step.




Figure 8. Plot of cell quality (the ratio of the largest to smallest sides of a cell) in the transition region and, for comparison, the global cell quality for the global low-resolution mesh and the degraded meshes. The degraded meshes can serve as a proxy for the impact of cell quality in the transition region. Notice that the cell quality in the transition region of the CUPS8-transition-900 mesh is comparable to that of the EC60to30-degraded-0.25 mesh and that of the CUPS8-transition-100 mesh is comparable to the EC60to30-degraded-0.75 mesh.

While these results are encouraging—showing that the TRiSK-based numerical formulation employed by MPAS-Ocean appears to be relatively insensitive to mesh distortion—these conclusions should be tempered by the nature of simulations run. Specifically, our analysis is restricted to relatively low-resolution, eddy parameterized configurations, where it may be expected that the dissipation due to viscous mixing acts to damp down any noise and/or oscillations generated at the grid scale. It is further noted that at low resolution and with constant forcing, energy is primarily injected into the system at relatively long, well-resolved wavelengths. Future studies may expand on the results presented here, using a set of high-resolution, eddy-resolving configurations to study interactions between the discretization, mesh quality, and grid scale response in the absence of explicit viscous damping.

4.2. Study 2: Transition Width

The analysis of the transition width study can be found in Figure 11. As the transition width increases, the dynamics of the simulations improve. The simulations with wider transition regions show greater SSH RMS and EKE (see Figure 10). This is to be expected, both because the transition is less steep, leading to higher quality cells in the transition region, and because the higher-resolution area is effectively larger with a greater transition width (note the locations of the center of the transition region in Figure 11). The three widest transitions (900, 600, and 300 km) have closer average values. The 100 km transition, where the Gulf Stream is meandering into the low-resolution region, shows a more significant decline in average SSH RMS and EKE. Even within 400 km of the coast, where the resolution is 8 km in all the simulations, the dynamics were improved by a wider transition region. With a narrower transition, meanders and eddies from the Gulf Stream cross into regions of lower resolution. It appears that these features are then smoothed out and do not have time to recover even when returning to the high-resolution region. This result is consistent with that found by Danilov and Wang (2015), in which eddies did not develop at the beginning of the eddy-permitting region but instead developed only once perturbations had developed further downstream. This is clearly seen in Figure 11. It also appears possible that the transition region is affecting the path of the Gulf Stream, “trapping” it within the high-resolution region. However, there is not a wide enough spread of transition widths in this study to say anything definitive about this effect. The transport through the Gulf Stream transects increases with wider transition widths, with the exception of the Cape Hatteras transect, which shows the opposite pattern.

In addition to the results examined here, the results of the degraded mesh study should be considered as a proxy for the transition regions. Although the cell quality within the CUPS8-transition-100 transition region

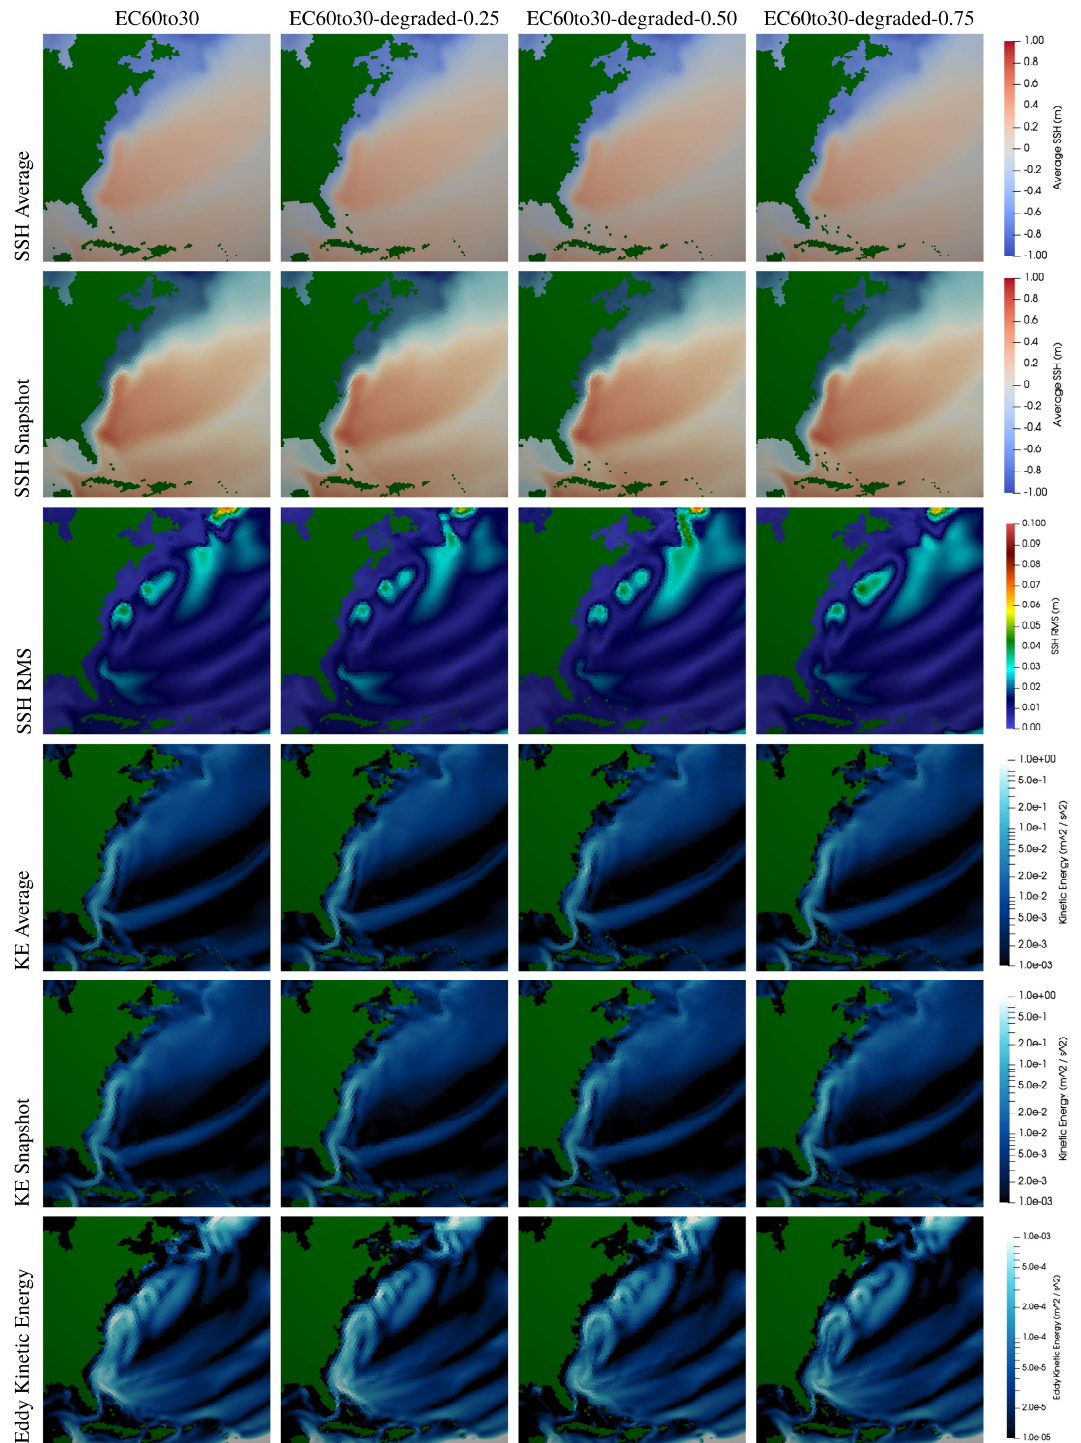


Figure 9. Degraded Mesh Study: Averages are taken from years 2–10 of the simulation, snapshots from 0002–06–01. The degraded meshes have a minimal impact on simulation quality.

is comparable to that of the EC60to30-degraded-0.75 mesh, the CUSP8-transition-100 mesh did not require the smaller time steps that the EC60to30-degraded-0.50 and EC60to30-degraded-0.75 meshes did. If the reduction in time step in the case of the degraded meshes is due to the presence of smaller cells, it would make sense that it is not required here, as the degraded cells in the transition region are larger than the cells in the high-resolution region. The results of the degraded mesh study indicate that mesh quality does




Figure 10. Plot of the average surface SSH RMS and EKE over the region shown in Figure 6 for years 2–10. The CUSP meshes have significantly higher average SSH RMS and EKE than the EC60to30 meshes. The variability increased as the mesh degradation increased. As the transition width was narrowed, the variability decreased, though this effect was small between CUSP8-transition-900 and CUSP8-transition-300. As the resolution decreased, the variability decreased, reaching the same values as the EC60to30 mesh for CUSP30, as would be expected.

not have a large impact on simulation results. The variation between the CUSP meshes is probably due primarily to other effects, such as the smaller region of higher resolution, rather than the cell quality within the transition region.

4.3. Study 3: Coastal Resolution

The analysis of the coastal resolution study can be found in Figure 12. The meshes with higher coastal resolution showed significantly improved dynamics, particularly in eddy kinetic energy and sea surface height variability, which were much smaller in CUSP30 (see Figure 12). The Gulf Stream within the high-resolution region in CUSP8 is similar to that of NA8. However, as noted in the transition width study, features that cross into the lower resolution transition region and then back into the high-resolution region, such as meanders and eddies, are less well resolved in the CUSP8 simulation. Figure 10 shows the effect of this as well. The integration region is within the high-resolution region in all of the simulation. However, the average SSH RMS and EKE decreases with decreasing transition width. This indicates that the variability within the high-resolution region is affected by the adjacent low-resolution region, presumably by current passing from low to high resolution. This effect was described by Danilov and Wang (2015).

Figure 13 shows the path of the Gulf Stream in the coastal resolution study. The NA8 simulation shows very little variability in the path of the Gulf Stream, while the CUSP8, CUSP12, and CUSP20 simulations show much more. The CUSP30 simulation also does not show much variability in the Gulf Stream path, but this is expected as the resolution is too low to be eddy permitting. The variation in the Gulf Stream path is also apparent in the transport through the transects along the Gulf Stream. In the southernmost transects (Florida-Cuba and Florida-Bahamas) where the flow is geographically constrained, the transport increases with increased resolution. The Cape Hatteras and New Jersey transects do not show this pattern. Figure 13 and Table 2 show that in the CUSP8, CUSP12, and CUSP20 simulations, there is significant variability in the path of the flow in the region of the New Jersey transect. Periods of very low or negative transport are probably due to North-South flow through the transect as the Gulf Stream separates from the coast. This can be seen in some of the monthly Gulf Stream paths seen in Figure 13.

This high variability is not due only to high resolution, as the NA8 simulation, which has the same coastal resolution, shows very little variability in the path of the Gulf Stream. The very low variability in the NA8 simulation is probably due largely to the idealized forcing used, as this effect did not show up in global high-resolution simulations with realistic climatological forcing. Initial results from global high-resolution simulations with climatological forcing show that the Gulf Stream had a realistic path and variability (see Figure A3). It appears that the lack of variation in the forcing or in the mesh itself prevents the NA8 mesh from developing meandering features. Variable forcing appears to resolve this problem. The variability in cell size and quality in the CUSP meshes may allow these features to develop. For all the transects, the

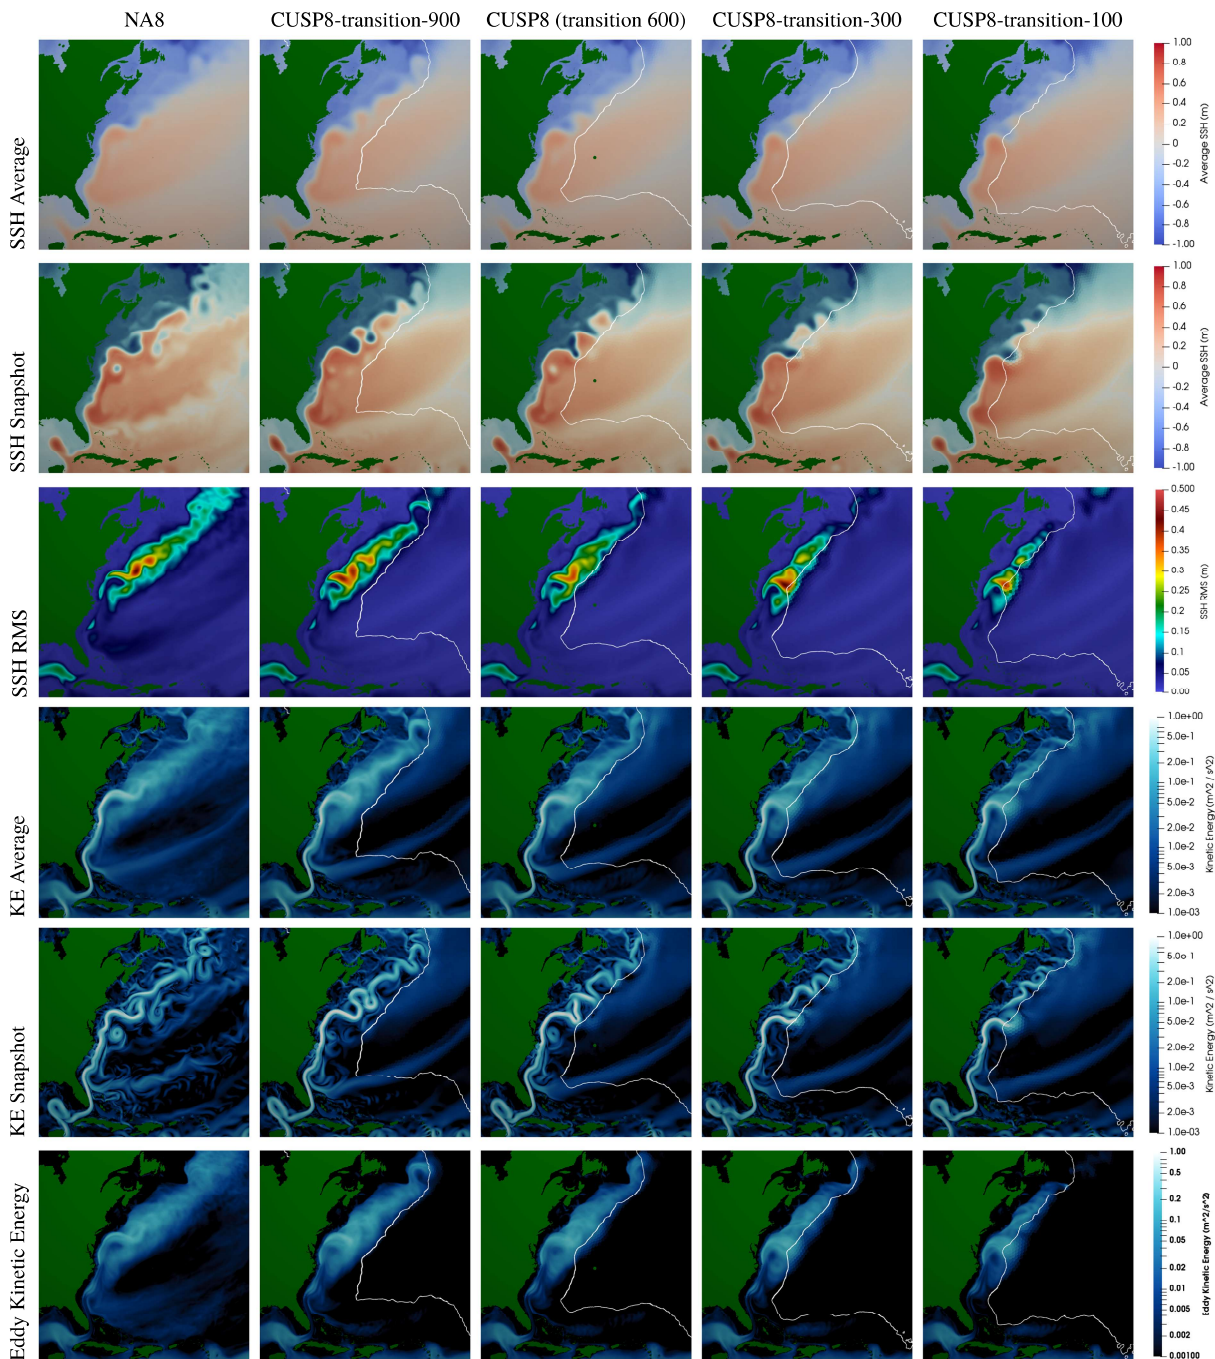


Figure 11. Transition Width Study: A wider transition improves simulation quality and increases variability within the coastal region. This appears to be less a function of the transition itself and rather a function of the size of the higher-resolution region (see position of the center of the transition region). The white line shows the center of the transition region. See Figure 9 for details.

variability increased significantly between CUSP30 and the higher-resolution meshes, as would be expected when transitioning to an eddy-permitting resolution.

The CUSP simulations are also a significant improvement on global high resolution in terms of cost (see Figure 14). CUSP8 offers an order of magnitude improvement in speed when compared to global high-resolution simulations with cell sizes ranging from 18 to 6 km. EC60to30, while an order of magnitude faster than CUSP8, lacks the improvements in coastal dynamics that motivated the creation of the CUSP8 mesh. Performance tests were run on Grizzly at Los Alamos National Laboratory. Grizzly is an Institutional

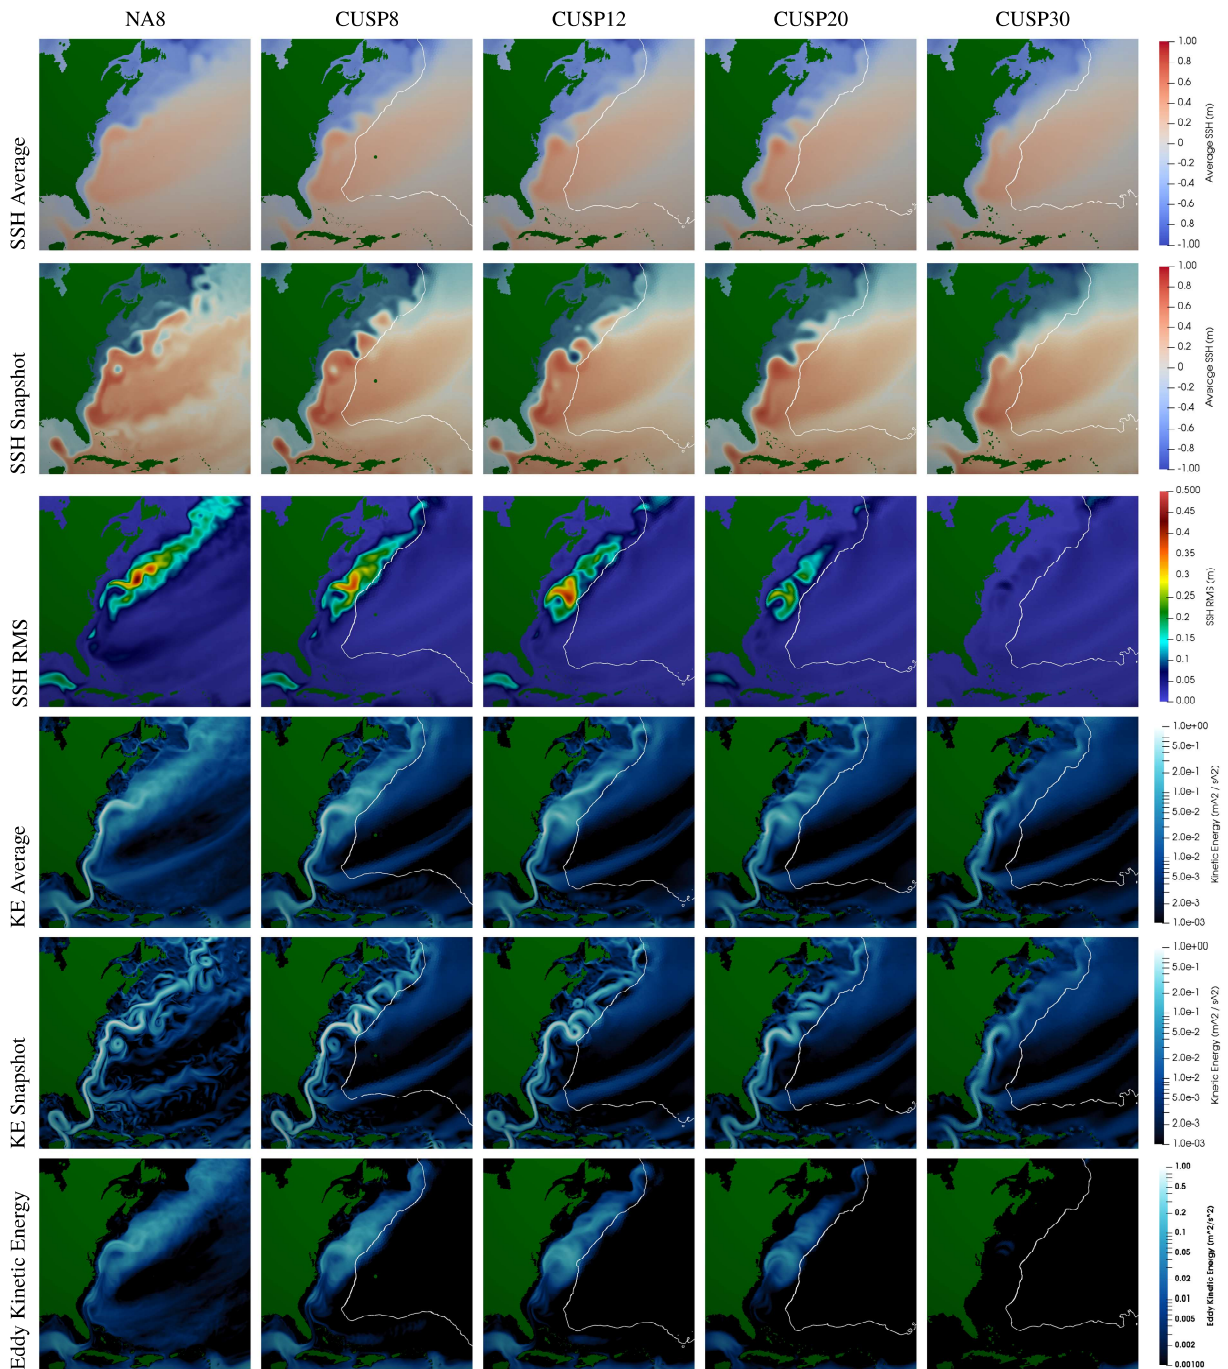


Figure 12. Resolution Study: Same as Figure 11. A higher-resolution improves the dynamics of the Gulf Stream significantly, with CUSP8 approaching the dynamics of NA8 along the coast. Variability increases with increasing resolution.

Computing (IC) cluster, running on the TOSS operating system (Tri-Lab Operating System Stack) and using the Intel OmniPath interconnect. Each core is a 2.1 GHz Broadwell with a 45 MB cache, with 36 processors per node.

5. Conclusion

Overall, this mesh resolution case study indicates that simulations are robust to changes in the mesh. Changes to mesh quality were found to have little impact on the simulation quality and statistics. Problems with the stability of the simulations at large time steps occurred in spinning-up the two most degraded

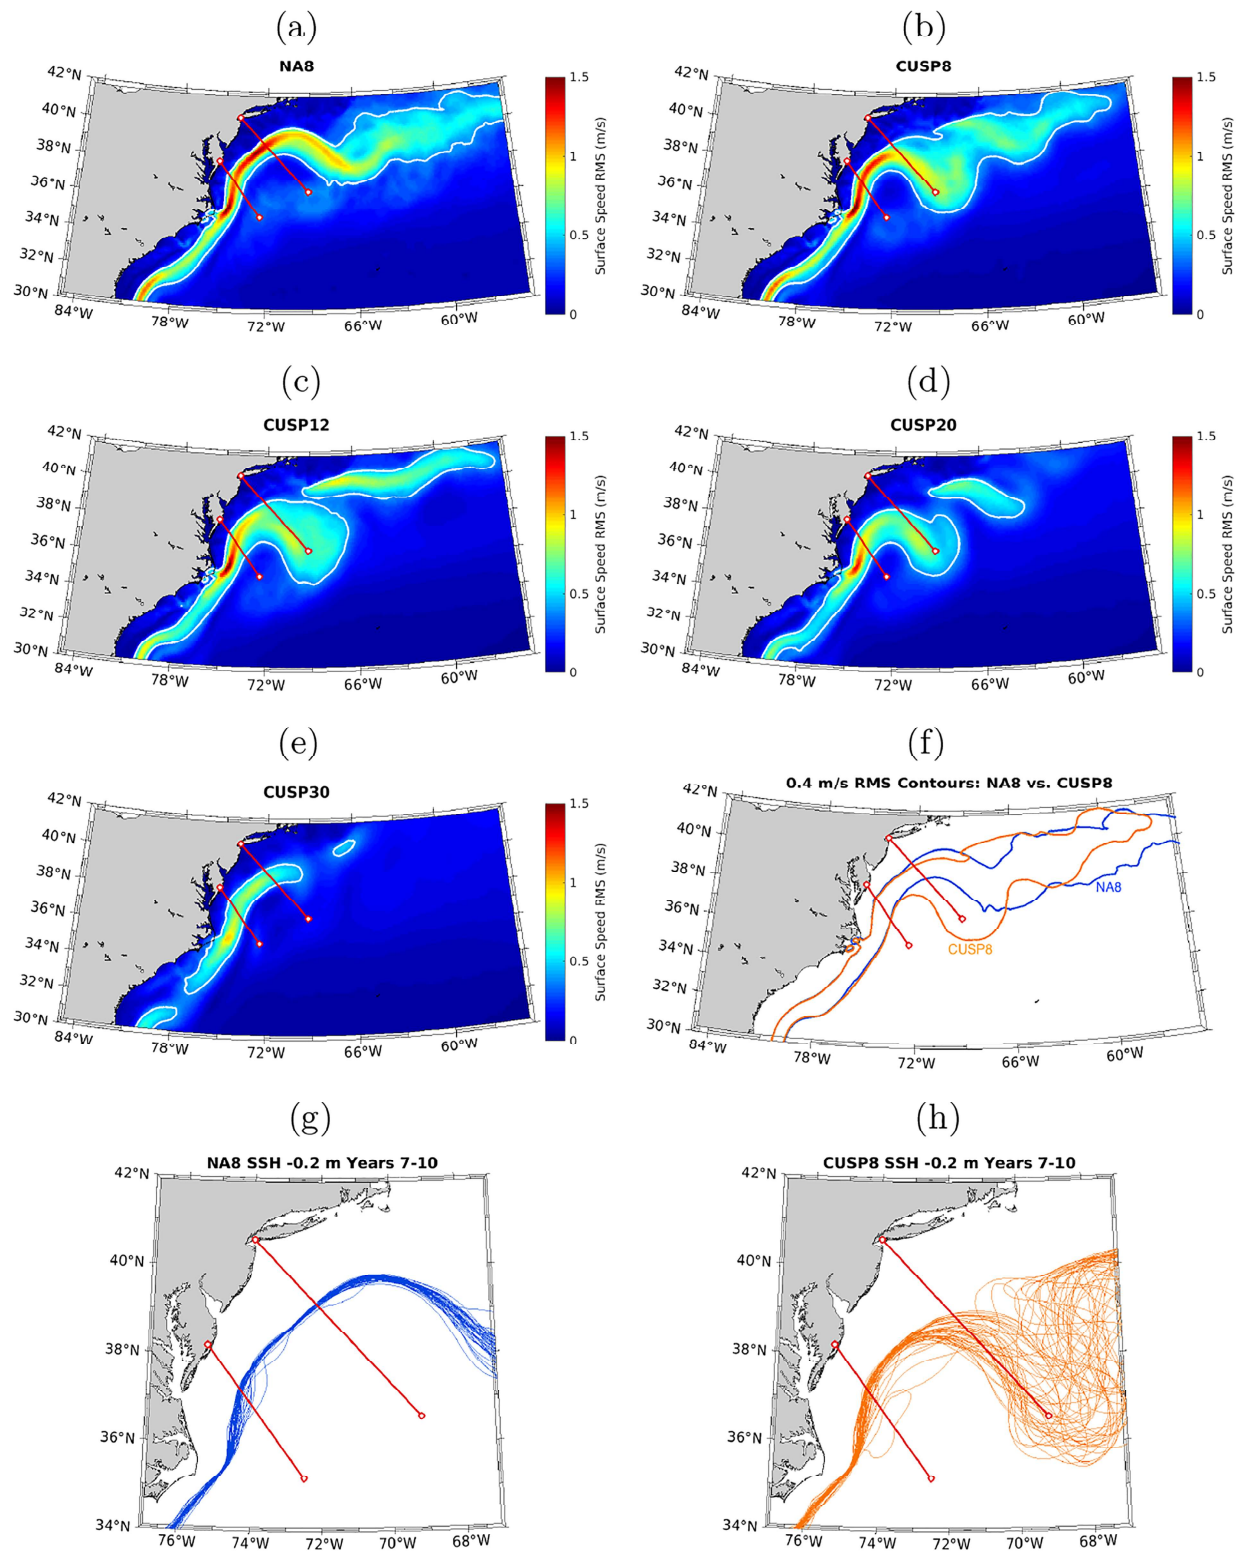


Figure 13. Root mean square (RMS) of surface speed for resolution convergence runs. Surface speed is taken from the first 5-day snapshot of each month for years 2 through 10. White contour on (a) to (e) indicates 0.4 m s^{-1} surface speed RMS contour. These contours are seen in (f) for NA8 (blue) and CUSP8 (orange). The northern and southern red lines on each plot indicate the New Jersey and Cape Hatteras transects, respectively. Snapshots of the Gulf Stream path are shown for NA8 (g) and CUSP8 (h). These pathlines follow the -0.2 m SSH contour. Paths are taken from the first 5-day snapshot of each month for years 7 through 10. Notice how the path of CUSP8 (h) frequently loops back through the New Jersey transect. This is the cause of the low transport through this transect for CUSP8 relative to NA8 and observations.




Figure 14. Performance for resolution study, showing simulated years per wall clock day (SYPD). Black dotted lines show perfect scaling. RRS18to6 is the global high-resolution mesh, with resolution varying from 18 to 6 km. The RRS18to6 mesh has 3.7 million cells, as compared to 236,000 cells in the EC60to30 mesh and 649,000 in the CUSP8 mesh. The SYPD values for 1,024 processors are CUSP8: 2.0, CUSP12: 5.1, CUSP20: 11.7, CUSP30: 20.0, EC60to30: 32.5, and RRS18to6: 0.38. CUSP8 is 16 times slower than EC60to30, but 5.3 times faster than global high resolution with cell sizes ranging from 18 to 6 km (RRS18to6). All simulations use 80 layers, except the EC60to30, which is 60 layers.

configurations, but with a modified time step, these simulations were found to perform similarly to the non-degraded cases. Such behavior is consistent with the expected reduction of the time step due to CFL limits associated with heavily degraded meshes that incorporate small grid cells. Despite previous theoretical analysis suggesting a strong link between mesh quality and numerical discretization error in inviscid settings (Peixoto, 2016), it was found that simulation quality was not obviously diminished with increasing levels of mesh degradation. In this sense, it appears that the TRISK formulation used in MPAS-Ocean may outperform its theoretical bounds in many practical cases, when using the typical suite of parameters for global simulations with non-zero dissipation. Changes to the transition width were also found to have relatively little impact on the quality of the simulations.

It is likely that much of the variation in the transition width study was due to the change in the size of the higher-resolution portion of the transition region rather than the transition itself. The difference between the steady Gulf Stream path in the NA8 simulation and the variable paths in the CUSP simulations shows that the transition region has some impact on variability. It is not clear if this is due to mesh quality or to the effect of changing resolution. In this case, the CUSP meshes had more realistic variability, but it is not clear that this added variability would be desirable in a simulation with realistic forcing. This study also demonstrated that higher coastal resolution improved the dynamics of the Gulf Stream at a much lower cost than a high-resolution global model.

When designing a mesh, the effect of processes outside of the high-resolution region is essential. The transition width study showed that processes within the high-resolution region cannot be properly resolved if they interact with processes in the low-resolution region. For example, in the CUSP8-transition-100 simulation, meanders and eddies crossing into the low-resolution region had a strong impact on the dynamics present along the Gulf Stream within the high-resolution region. More broadly, it is important to evaluate the dependence of the coastal dynamics on basin scale or global dynamics. A coastal high-resolution model may be of limited use if the ultimate drivers of the coastal dynamics are not modeled accurately. For instance, flooding during a hurricane requires that offshore storm surges are modeled at appropriate resolution in order to predict accurate coastal surges.

Physical dynamics considerations appear to be much more important than mesh metrics considerations in these stand-alone ocean simulations. The cases presented here were limited to idealized surface forcing in order to conduct a large parameter study. Future studies will look in more detail at CUSP8 simulations with realistic atmospheric forcing and in coupled configurations, which may have more stringent mesh quality requirements due to cross-component feedbacks. Upcoming papers will examine the CUSP8 mesh in simulations with realistic forcing and active sea ice. Subsequent research has been done to explore the causes of

the weak Gulf Stream in the CUSP meshes and has revealed that it may be linked to biases in the Labrador Sea. Work to resolve these issues in both low and variable-resolution meshes is ongoing. Based on the results of this study and further research, modifications have been made to the CUSP8 mesh. The high-resolution region has been extended to encompass the Gulf Stream extension, which will prevent the current being steered by the transition. High resolution will also extend into the Labrador Sea and the coastlines of Greenland and the Canadian Arctic, resolving other processes that are essential to the Gulf Stream, including downwelling in the Labrador sea and the Labrador current.

Similar variable-resolution meshes are in development for investigating other regions of interest, including the Arctic and Southern Oceans. Our results suggest robust capabilities inherent in the MPAS-Ocean discretization and mesh generation approaches. These provide the capability to create a diverse range of variable-resolution configurations, which will allow modelers to accurately resolve additional physical processes at lower computational costs.

Appendix a

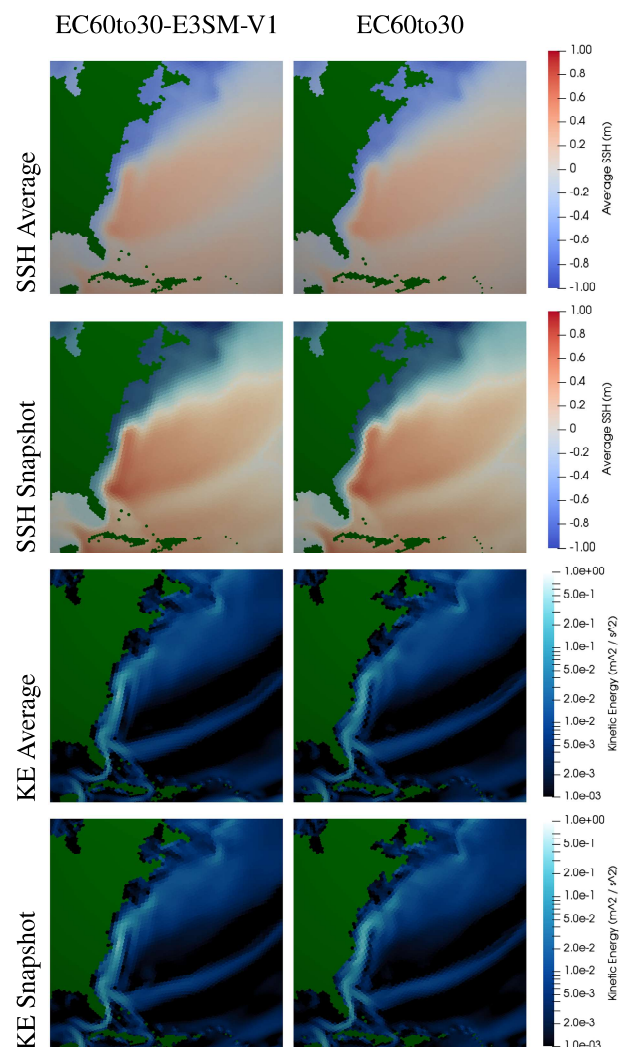


Figure A1. EC60to30-E3SM-V1 vs EC60to30: Averages are taken from years 2–10 of the simulation, snapshots from 0002-06-01.




Figure A2. EC60to30 Layers: A comparison of EC60to30 meshes with different numbers of vertical layers. The mesh used in this paper was the 100 layer mesh. The CUSP meshes used an 80 layer mesh.

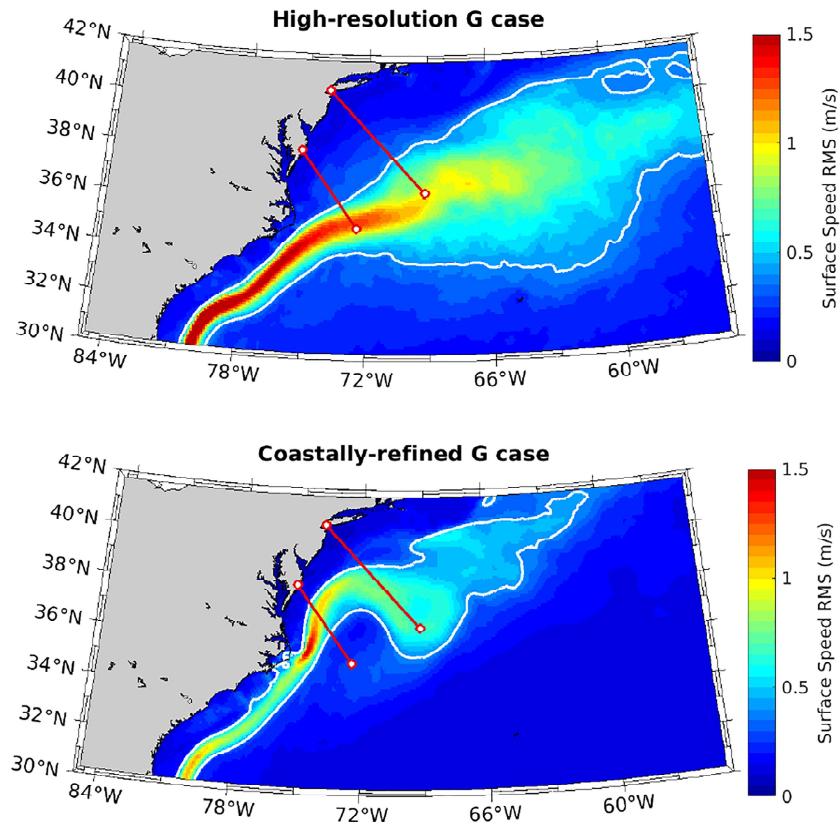


Figure A3. Root mean square (RMS) of surface speed for runs forced with CORE realistic atmosphere. High-resolution 18–6 km eddy-permitting run (a) and coastal-refined 8 km run (b). White contour indicates 0.4 m s^{-1} surface speed RMS contour. In the North Atlantic, the high-resolution grid in (a) is similar to the NA8 grid. The Gulf Stream separation, variability, and transport are much more realistic in the CORE-forced high-resolution run (a) than in any of the climatology-forced runs (Figure 13).

Acknowledgments

We thank the editor and reviewers for their careful reading of the manuscript and helpful comments. In particular, we appreciate the suggestions made by Sergey Danilov, whose guidance was extremely useful. This research was supported as part of the Energy Exascale Earth System Model (E3SM) project, funded by the U.S. Department of Energy, Office of Science, Office of Biological and Environmental Research as well as the U.S. Department of Energy Advanced Research Projects Agency - Energy (ARPA-E) Macroalgae Research Inspiring Novel Energy Resources (MARINER) program (Funding Opportunity DE-FOA-0001726, MARINER Award 17/CJ000/09/01, Pacific Northwest National Laboratory, prime recipient). E3SM simulations are conducted at Argonne Leadership Computing Facility (contract DE-AC02-06CH11357); National Energy Research Scientific Computing Center (DE-AC05-00OR22725); Oak Ridge Leadership Computing Facility (DE-AC05-00OR22725); Argonne Nat. Lab. high-performance computing cluster, provided by BER Earth System Modeling; and Los Alamos Nat. Lab. Institutional Computing, US DOE NNSA (DE-AC52-06NA25396). E3SM data are freely available through the Earth System Grid Federation (ESGF) distributed archives. See details online (<https://e3sm.org/data>). JIGSAW can be obtained freely online (<https://github.com/dengwirda/jigsaw-geo-matlab>).

References

- Abiodun, B. J., Prusa, J. M., & Gutowski, W. J. (2008). Implementation of a non-hydrostatic, adaptive-grid dynamics core in CAM3. Part I: Comparison of dynamics cores in aqua-planet simulations. *Climate Dynamics*, *31*(7), 795–810. <https://doi.org/10.1007/s00382-008-0381-y>
- Androsov, A., Fofonova, V., Kuznetsov, I., Danilov, S., Rakowsky, N., Harig, S., et al. (2019). FESOM-C v.2: Coastal dynamics on hybrid unstructured meshes. *Geoscientific Model Development*, *12*(3), 1009–1028. <https://doi.org/10.5194/gmd-12-1009-2019>
- Biastoch, A., Sein, D., Durgadoo, J. V., Wang, Q., & Danilov, S. (2018). Simulating the Agulhas system in global ocean models—Nesting vs. multi-resolution unstructured meshes. *Ocean Modelling*, *121*, 117–131. <https://doi.org/10.1016/j.ocemod.2017.12.002>
- Candy, A. S., & Pietrzak, J. D. (2018). Shingle 2.0: Generalising self-consistent and automated domain discretisation for multi-scale geophysical models. *Geoscientific Model Development*, *11*(1), 213–234. <https://doi.org/10.5194/gmd-11-213-2018>
- Chen, L., & Holst, M. (2011). Efficient mesh optimization schemes based on optimal delaunay triangulations. *Computer Methods in Applied Mechanics and Engineering*, *200*(9), 967–984. <https://doi.org/10.1016/j.cma.2010.11.007>
- Chen, C., Liu, H., & Beardsley, R. C. (2003). An unstructured grid, finite-volume, three-dimensional, primitive equations ocean model: Application to coastal ocean and estuaries. *Journal of Atmospheric and Oceanic Technology*, *20*(1), 159–186. [https://doi.org/10.1175/1520-0426\(2003\)020<0159:AUGFVT>2.0.CO;2](https://doi.org/10.1175/1520-0426(2003)020<0159:AUGFVT>2.0.CO;2)
- Danilov, S., Sidorenko, D., Wang, Q., & Jung, T. (2017). The finite-volume sea ice–ocean model (FESOM2). *Geoscientific Model Development*, *10*(2), 765–789. <https://doi.org/10.5194/gmd-10-765-2017>
- Danilov, S., & Wang, Q. (2015). Resolving eddies by local mesh refinement. *Ocean Modelling*, *93*, 75–83. <https://doi.org/10.1016/j.ocemod.2015.07.006>
- Davies, D. R., Wilson, C. R., & Kramer, S. C. (2011). Fluidity: A fully unstructured anisotropic adaptive mesh computational modeling framework for geodynamics. *Geochemistry, Geophysics, Geosystems*, *12*, Q06001. <https://doi.org/10.1029/2011GC003551>
- Debreu, L., & Blayo, E. (2008). Two-way embedding algorithms: A review. *Ocean Dynamics*, *58*, 415–428.
- Donohue, K. A., Tracey, K. L., Watts, D. R., Chidichimo, M. P., & Chereskin, T. K. (2016). Mean Antarctic circumpolar current transport measured in drake passage. *Geophysical Research Letters*, *43*, 11,760–11,767. <https://doi.org/10.1002/2016GL070319>
- Düben, P. D., & Korn, P. (2014). Atmosphere and ocean modeling on grids of variable resolution—A 2D case study. *Monthly Weather Review*, *142*(5), 1997–2017. <https://doi.org/10.1175/MWR-D-13-00217.1>
- Engwirda, D. (2017). JIGSAW-GEO (1.0): Locally orthogonal staggered unstructured grid generation for general circulation modelling on the sphere. *Geoscientific Model Development*, *10*(6), 2117–2140. <https://doi.org/10.5194/gmd-10-2117-2017>

- Engwirda, D. (2018). Generalised primal-dual grids for unstructured co-volume schemes. *Journal of Computational Physics*, 375, 155–176. <https://doi.org/10.1016/j.jcp.2018.07.025>
- Engwirda, D., & Ivers, D. (2016). Off-centre steiner points for delaunay-refinement on curved surfaces. *Computer-Aided Design*, 72, 157–171. 23rd International Meshing Roundtable Special Issue: Advances in Mesh Generation <https://doi.org/10.1016/j.cad.2015.10.007>
- Fernández-Montblanc, T., Vousdoukas, M. I., Ciavola, P., Voukouvalas, E., Mentaschi, L., Breyiannis, G., et al. (2019). Towards robust pan-European storm surge forecasting. *Ocean Modelling*, 133, 129–144. <https://doi.org/10.1016/j.ocemod.2018.12.001>
- Ganachaud, A., & Wunsch, C. (2000). Improved estimates of global ocean circulation, heat transport and mixing from hydrographic data. *Nature*, 408(6811), 453–457. <https://doi.org/10.1038/35044048>
- Gent, P. R., & McWilliams, J. C. (1990). Isopycnal mixing in ocean circulation models. *Journal of Physical Oceanography*, 20(1), 150–155. [https://doi.org/10.1175/1520-0485\(1990\)020<0150:IMIOCM>2.0.CO;2](https://doi.org/10.1175/1520-0485(1990)020<0150:IMIOCM>2.0.CO;2)
- Golaz, J.-C., Caldwell, P. M., Van Roekel, L. P., Petersen, M. R., Tang, Q., Wolfe, J. D., et al. (2019). The DOE E3SM coupled model version 1: Overview and evaluation at standard resolution. *Journal of Advances in Modeling Earth Systems*, 11, 2089–2129. <https://doi.org/10.1029/2018MS001603>
- Griffies, S. M., Biastoch, A., Böning, C., Bryan, F., Danabasoglu, G., Chassignet, E. P., et al. (2009). Coordinated ocean-ice reference experiments (COREs). *Ocean Modelling*, 26(1-2), 1–46. <https://doi.org/10.1016/j.ocemod.2008.08.007>
- Hagos, S., Leung, R., Rauscher, S. A., & Ringler, T. (2013). Error characteristics of two grid refinement approaches in aquaplanet simulations: MPAS-A and WRF. *Monthly Weather Review*, 141(9), 3022–3036. <https://doi.org/10.1175/MWR-D-12-00338.1>
- Halkin, D., & Rossby, T. (1985). The structure and transport of the Gulf Stream at 73°W. *Journal of Physical Oceanography*, 15(11), 1439–1452. [https://doi.org/10.1175/1520-0485\(1985\)015<1439:TSATOT>2.0.CO;2](https://doi.org/10.1175/1520-0485(1985)015<1439:TSATOT>2.0.CO;2)
- Higdon, R. L. (2005). A two-level time-stepping method for layered ocean circulation models: Further development and testing. *Journal of Computational Physics*, 206(2), 463–504. <https://doi.org/10.1016/j.jcp.2004.12.011>
- Jacobsen, D. W., Gunzburger, M., Ringler, T., Burkardt, J., & Peterson, J. (2013). Parallel algorithms for planar and spherical Delaunay construction with an application to centroidal Voronoi tessellations. *Geoscientific Model Development*, 6(4), 1353–1365. <https://doi.org/10.5194/gmd-6-1353-2013>
- Johns, W. E., Townsend, T. L., Fratantoni, D. M., & Wilson, W. D. (2002). On the Atlantic inflow to the Caribbean Sea. *Deep Sea Research Part I: Oceanographic Research Papers*, 49(2), 211–243. [https://doi.org/10.1016/S0967-0637\(01\)00041-3](https://doi.org/10.1016/S0967-0637(01)00041-3)
- Kärnä, T., Legat, V., & Deleersnijder, E. (2013). A baroclinic discontinuous Galerkin finite element model for coastal flows. *Ocean Modelling*, 61, 1–20. <https://doi.org/10.1016/j.ocemod.2012.09.009>
- Korn, P. (2017). Formulation of an unstructured grid model for global ocean dynamics. *Journal of Computational Physics*, 339, 525–552. <https://doi.org/10.1016/j.jcp.2017.03.009>
- Lambrechts, J., Comblen, R., Legat, V., Geuzaine, C., & Remacle, J.-F. (2008a). Multiscale mesh generation on the sphere. *Ocean Dynamics*, 58(5), 461–473. <https://doi.org/10.1007/s10236-008-0148-3>
- Lambrechts, J., Comblen, R., Legat, V., Geuzaine, C., & Remacle, J.-F. (2008b). Multiscale mesh generation on the sphere. *Ocean Dynamics*, 58, 461–473. <https://doi.org/10.1007/s10236-008-0148-3>
- Lorant, V., & Royer, J.-F. (2001). Sensitivity of equatorial convection to horizontal resolution in aquaplanet simulations with a variable-resolution GCM. *Monthly Weather Review*, 129(11), 2730–2745. [https://doi.org/10.1175/1520-0493\(2001\)129<2730:SOECTH>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<2730:SOECTH>2.0.CO;2)
- McRae, A., Cotter, C., & Budd, C. (2018). Optimal-transport-based mesh adaptivity on the plane and sphere using finite elements. *SIAM Journal on Scientific Computing*, 40(2), A1121–A1148. <https://doi.org/10.1137/16M1109515>
- Park, S.-H., Klemp, J. B., & Skamarock, W. C. (2014). A comparison of mesh refinement in the global MPAS-A and WRF models using an idealized normal-mode baroclinic wave simulation. *Monthly Weather Review*, 142(10), 3614–3634. <https://doi.org/10.1175/MWR-D-14-00004.1>
- Peixoto, P. S. (2016). Accuracy analysis of mimetic finite volume operators on geodesic grids and a consistent alternative. *Journal of Computational Physics*, 310, 127–160. <https://doi.org/10.1016/j.jcp.2015.12.058>
- Petersen, M. R., Asay-Davis, X. S., Berres, A. S., Chen, Q., Feige, N., Hoffman, M. J., et al. (2019). An evaluation of the ocean and sea ice climate of E3SM using MPAS and interannual CORE-II forcing. *Journal of Advances in Modeling Earth Systems*, 11, 1438–1458. <https://doi.org/10.1029/2018MS001373>
- Petersen, M. R., Jacobsen, D. W., Ringler, T. D., Hecht, M. W., & Maltrud, M. E. (2015). Evaluation of the arbitrary Lagrangian–Eulerian vertical coordinate method in the MPAS-Ocean model. *Ocean Modelling*, 86, 93–113. <https://doi.org/10.1016/j.ocemod.2014.12.004>
- Rackow, T., Sein, D. V., Semmler, T., Danilov, S., Koldunov, N. V., Sidorenko, D., et al. (2019). Sensitivity of deep ocean biases to horizontal resolution in prototype CMIP6 simulations with AWI-CM1.0. *Geoscientific Model Development*, 12(7), 2635–2656. <https://doi.org/10.5194/gmd-12-2635-2019>
- Rauscher, S. A., & Ringler, T. D. (2014). Impact of variable-resolution meshes on midlatitude baroclinic eddies using CAM-MPAS-A. *Monthly Weather Review*, 142(11), 4256–4268. <https://doi.org/10.1175/MWR-D-13-00366.1>
- Rauscher, S. A., Ringler, T. D., Skamarock, W. C., & Mirin, A. A. (2012). Exploring a global multiresolution modeling approach using aquaplanet simulations. *Journal of Climate*, 26(8), 2432–2452. <https://doi.org/10.1175/JCLI-D-12-00154.1>
- Reckinger, S. M., Petersen, M. R., & Reckinger, S. J. (2015). A study of overflow simulations using MPAS-Ocean: Vertical grids, resolution, and viscosity. *Ocean Modelling*, 96, 291–313. <https://doi.org/10.1016/j.ocemod.2015.09.006>
- Remacle, J.-F., & Lambrechts, J. (2018). Fast and robust mesh generation on the sphere—Application to coastal domains. *Computer-Aided Design*, 103, 14–23. 25th International Meshing Roundtable Special Issue: Advances in Mesh Generation <https://doi.org/10.1016/j.cad.2018.03.002>
- Ringler, T., Petersen, M., Higdon, R. L., Jacobsen, D., Jones, P. W., & Maltrud, M. (2013). A multi-resolution approach to global ocean modeling. *Ocean Modelling*, 69, 211–232. <https://doi.org/10.1016/j.ocemod.2013.04.010>
- Ringler, T. D., Thuburn, J., Klemp, J. B., & Skamarock, W. C. (2010). A unified approach to energy conservation and potential vorticity dynamics for arbitrarily-structured C-grids. *Journal of Computational Physics*, 229(9), 3065–3090. <https://doi.org/10.1016/j.jcp.2009.12.007>
- Rosby, T., Flagg, C. N., Donohue, K., Sanchez-Franks, A., & Lillibridge, J. (2014). On the long-term stability of Gulf Stream transport based on 20 years of direct measurements. *Geophysical Research Letters*, 41, 114–120. <https://doi.org/10.1002/2013GL058636>
- Sakaguchi, K., Lu, J., Leung, L. R., Zhao, C., Li, Y., & Hagos, S. (2016). Sources and pathways of the upscale effects on the Southern Hemisphere jet in MPAS-CAM4 variable-resolution simulations. *Journal of Advances in Modeling Earth Systems*, 8, 1786–1805. <https://doi.org/10.1002/2016MS000743>

- Scholz, P., Sidorenko, D., Gurses, O., Danilov, S., Koldunov, N., Wang, Q., et al. (2019). Assessment of the finite volume sea ice ocean model (FESOM2.0), part I: Description of selected key model elements and comparison to its predecessor version. *Geoscientific Model Development Discussions*, 2019, 1–42. <https://doi.org/10.5194/gmd-2018-329>
- Sein, D. V., Danilov, S., Biastoch, A., Durgadoo, J. V., Sidorenko, D., Harig, S., & Wang, Q. (2016). Designing variable ocean model resolution based on the observed ocean variability. *Journal of Advances in Modeling Earth Systems*, 8, 904–916. <https://doi.org/10.1002/2016MS000650>
- Sein, D. V., Koldunov, N. V., Danilov, S., Wang, Q., Sidorenko, D., Fast, I., et al. (2017). Ocean modeling on a mesh with resolution following the local rossby radius. *Journal of Advances in Modeling Earth Systems*, 9, 2601–2614. <https://doi.org/10.1002/2017MS001099>
- Skamarock, W. C., & Gassmann, A. (2011). Conservative transport schemes for spherical geodesic grids: High-order flux operators for ODE-based time integration. *Monthly Weather Review*, 139(9), 2962–2975. <https://doi.org/10.1175/MWR-D-10-05056.1>
- Steele, M., Morley, R., & Ermold, W. (2001). PHC: A global ocean hydrography with a high-quality arctic ocean. *Journal of Climate*, 14(9), 2079–2087. [https://doi.org/10.1175/1520-0442\(2001\)014<2079:PAGOHW>2.0.CO;2](https://doi.org/10.1175/1520-0442(2001)014<2079:PAGOHW>2.0.CO;2)
- Tian, R., Chen, C., Qi, J., Ji, R., Beardsley, R. C., & Davis, C. (2014). Model study of nutrient and phytoplankton dynamics in the Gulf of Maine: Patterns and drivers for seasonal and interannual variability. *ICES Journal of Marine Science*, 72(2), 388–402. <https://doi.org/10.1093/icesjms/fsu090>
- Van Roekel, L., Adcroft, A. J., Danabasoglu, G., Griffies, S. M., Kauffman, B., Large, W., et al. (2018). The KPP boundary layer scheme for the ocean: Revisiting its formulation and benchmarking one-dimensional simulations relative to les. *Journal of Advances in Modeling Earth Systems*, 10, 2647–2685. <https://doi.org/10.1029/2018MS001336>
- Wang, Q., Danilov, S., Sidorenko, D., Timmermann, R., Wekerle, C., Wang, X., et al. (2014). The finite element sea ice-ocean model (FESOM) v.1.4: Formulation of an ocean general circulation model. *Geoscientific Model Development*, 7(2), 663–693. <https://doi.org/10.5194/gmd-7-663-2014>
- Wang, H., Loftis, J., Liu, Z., Forrest, D., & Zhang, J. (2014). The storm surge and sub-grid inundation modeling in New York City during Hurricane Sandy. *Journal of Marine Science and Engineering*, 2, 226–246. <https://doi.org/10.3390/jmse2010226>
- Wang, Q., Wekerle, C., Danilov, S., Wang, X., & Jung, T. (2018). A 4.5 km resolution Arctic Ocean simulation with the global multi-resolution model FESOM1.4. *Geoscientific Model Development*, 11, 1229–1255. <https://doi.org/10.5194/gmd-11-1229-2018>
- Weller, H., Browne, P., Budd, C., & Cullen, M. (2016). Mesh adaptation on the sphere using optimal transport and the numerical solution of a Monge–Ampère type equation. *Journal of Computational Physics*, 308, 102–123. <https://doi.org/10.1016/j.jcp.2015.12.018>
- Yang, H., Gunzburger, M., & Ju, L. (2018). Fast spherical centroidal Voronoi mesh generation: A Lloyd-preconditioned LBFGS method in parallel. *Journal of Computational Physics*, 367, 235–252. <https://doi.org/10.1016/j.jcp.2018.04.034>
- Ye, F., Zhang, Y. J., Wang, H. V., Friedrichs, M. A. M., Irby, I. D., Alteljevich, E., et al. (2018). A 3D unstructured-grid model for Chesapeake bay: Importance of bathymetry. *Ocean Modelling*, 127, 16–39. <https://doi.org/10.1016/j.ocemod.2018.05.002>
- Zarzycki, C. M., Jablonowski, C., Thatcher, D. R., & Taylor, M. A. (2015). Effects of localized grid refinement on the general circulation and climatology in the community atmosphere model. *Journal of Climate*, 28(7), 2777–2803. <https://doi.org/10.1175/JCLI-D-14-00599.1>
- Zhang, Y., Chen, C., Beardsley, R. C., Gao, G., Lai, Z., Curry, B., et al. (2016). Studies of the Canadian Arctic Archipelago water transport and its relationship to basin-local forcings: Results from AO-FVCOM. *Journal of Geophysical Research: Oceans*, 121, 4392–4415. <https://doi.org/10.1002/2016JC011634>
- Zhang, Y. J., Ye, F., Stanev, E. V., & Grashorn, S. (2016). Seamless cross-scale modeling with SCHISM. *Ocean Modelling*, 102, 64–81. <https://doi.org/10.1016/j.ocemod.2016.05.002>
- Zhao, C., Leung, L. R., Park, S.-H., Hagos, S., Lu, J., Sakaguchi, K., et al. (2016). Exploring the impacts of physics and resolution on aqua-planet simulations from a nonhydrostatic global variable-resolution modeling framework. *Journal of Advances in Modeling Earth Systems*, 8, 1751–1768. <https://doi.org/10.1002/2016MS000727>