Exploring a Multi-Resolution Modeling Approach within the Shallow-Water Equations

Todd D. Ringler *

Los Alamos National Laboratory

Doug Jacobsen and Max Gunzburger

Florida State University

Lili Ju

University of South Carolina

Michael Duda and William Skamarock

National Center for Atmospheric Research

* *Corresponding author address:* Todd D. Ringler, Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

E-mail: ringler@lanl.gov

ABSTRACT

The ability to solve the global shallow-water equations with a conforming, variable-resolution mesh is evaluated using standard shallow-water test cases. While our long-term motivation is the creation of a global climate modeling framework capable of resolving different spatial and temporal scales in different regions, we begin with an analysis of the shallow-water system in order to better understand the strengths and weaknesses of our approach. The multi-resolution meshes are spherical centroidal Voronoi tessellations where a single, user-supplied density function determines the region(s) of fine- and coarse-mesh resolution. We explore the shallow-water system with a suite of meshes ranging from quasi-uniform resolution meshes, where grid spacing is globally uniform, to highly-variable resolution meshes, where grid spacing varies by a factor of 16 between the fine and coarse regions. We find that potential vorticity is conserved to within machine precision and total available energy is conserved to within time-truncation error. This finding holds for the full suite of meshes, ranging from quasi-uniform resolution and highly-variable resolution meshes. Using shallow-water test cases 2 and 5, *we find that solution error is controlled primarily by the grid resolution in the coarsest part of the model domain*. This finding is consistent with results obtained by others. When these variable resolution meshes are used for the simulation of an unstable zonal jet, we find that the core features of the growing instability are largely unchanged as the variation in mesh resolution increases. The main differences between the simulations occur outside the region of mesh refinement and these differences are attributed to the additional truncation error that accompanies increases in grid spacing. Overall, the results demonstrate support for this approach as a path toward multi-resolution climate system modeling.

1. Introduction

A defining feature of the global atmosphere and ocean circulations is their broad range of temporal and spatial scales. The climate of the atmosphere is determined by both global patterns of motion, $O(10^4)$ km, as well as, for example, boundary layer processes with $O(10^{-1})$ km characteristic scales (Klein and Hartmann 1993). Similarly the climate of the ocean is controlled by both basin scales of motions, $O(10^4)$ km, and sub-mesoscale processes with $O(10^{-1})$ km scales. (Boccaletti et al. 2007). As is typical of nonlinear systems, the broad range of climate-relevant spatial scales in the atmosphere and ocean are highly interacting; the $O(10^4)$ km global scales modify and are modified by the $O(10^{-1})$ km local scales. In terms of simulating the atmosphere and ocean climate systems, the strong interaction across scales implies that an accurate representation of the smallest scales is a prerequisite for the robust simulation of the largest scales.

As a result of the broad scale interaction, the numerical simulation of the climate system is particularly challenging. For example, we do not presently have the computational resources to globally resolve all the scales associated with fundamental processes in the atmosphere and ocean, e.g. clouds and ocean eddies (Randall and Bony 2007). This unfortunate reality will remain true for decades to come. The corollary is that the numerical simulation of the global climate system is, and will likely always remain, an under-resolved endeavor.

Given the importance of small-scale processes such as clouds and ocean eddies in the climate system, numerical models are obligated, either through direct simulation or parameterization, to account for how these processes modify and are modified by the larger scales. Due to the constraint presented by today's computational resources, climate models are al-

most always relegated to the latter option of parameterization. Parameterizing a process is significantly more challenging than directly simulating that same process. When conducting a direct simulation, the interaction across scales is naturally accommodated. When parameterizing a process, we need to know *a priori* how the larger (resolved) scales act to regulate the smaller (unresolved) parameterized process and, in turn, how the parameterized process acts in an aggregate sense to modify the largest scales. In effect, an accurate parameterization requires significantly greater understanding of the underlying physics than does the direct simulation of that same process.

The pitfall of parameterization has led to what might be considered the defining tenet of global climate modeling: increasing model resolution allows for less parameterization and, thereby, a more accurate simulation of the observed climate system. Faced with the daunting challenges posed by global climate modeling, the community has embarked on at least three research paths to address this challenge. The first approach is that of global ultra high-resolution climate system modeling (McClean et al. 2010, accepted). In this approach, high-resolution climate system models are paired with the world's most advanced high performance computing systems to conduct climate simulations at unprecedented resolution. The underlying premise is that as the model resolves more and more of the scales of interests, less of the system is left unresolved and, thus, less of the systems requires parameterization. This approach is very much in the theme of traditional climate modeling but at very high resolution and, as a result, benefits from the decades of experience that this activity has obtained. The main disadvantage of this approach is that the presently unresolved parts of the spectrum are resolved at a painfully slow rate. Reducing the horizontal grid spacing by a factor of two typically requires a factor of 2^3 increase in computing resources, where lon-

gitude, latitude and time account individually for a factor of 2. Thus moving from a global 50 km mesh like those presently used for high-resolution IPCC atmosphere simulations to a global 4 km mesh that would be required for convection-permitting atmosphere simulations will entail an increase in computing resources of approximately 2^{12}, or about 4000 times the present-day computing capacity. And this, of course, neglects the substantial increase in vertical resolution that will also be required.

In order to circumvent the tyranny of global, high-resolution climate modeling, a second approach based on limited-area climate modeling has been explored over the last two decades (Giorgi and Mearns (1991), McGregor (1997) and Wang et al (2004)). This approach employs a high-resolution mesh placed only over the area of interest. Since the area of interest generally spans only a small portion of the sphere, e.g. the continental United States, the computational demands are significantly reduced as compared to global high-resolution modeling. As a result of being more computationally accessible, it is much easier to explore physical processes that might be relevant to regional climate dynamics and regional climate change (e.g. Diffenbaugh et al. (2005)). The disadvantage of the limited-area approach is the requirement to impose one-way, non-interactive lateral boundary conditions. These lateral boundary conditions can be obtained from reanalysis data or coarse-grain global climate simulations. The imposition of lateral-boundary conditions can lead to inconsistencies in the physics and dynamics of the limited-area models (see Wang et al (2004) for a review). Physical inconsistencies can arise when the global and regional models use different physical parameterizations (e.g. McGregor (1997) Fig. 4). Dynamical inconsistencies can arise from the lack of well-posedness of the lateral boundary conditions (Oliger and Sundström (1978), Staniforth (1997)) and mismatch between the solution of the global and regional models in

the nesting region (Davies (1976), Marbaix et al. (2003), Harris and Durran (2010). These inconsistencies can result in the regional and global simulations diverging toward different climate states (Jones et al. 1995). Physical inconsistencies can be ameliorated by using the same physical parameterizations in both the global and regional model (McGregor (1997), Lorenz and Jacob (2005)). Dynamical inconsistencies can be mitigated by overwriting the global model solution with the regional model solution after every time step (Lorenz and Jacob (2005), Inatsu and Kimoto (2009).

The third option being pursued is that of multi-scale modeling. While this method has been investigated primarily in the atmosphere modeling (Grabowski 2010), a preliminary exploration of this approach in ocean modeling is underway (Campin et al. 2010, revised). As the name suggested, this approach couples models at different scales to create a full simulation. Efforts to date have focused primarily on coupling global, coarse-grain models of atmosphere dynamics with embedded high-resolution models of cloud and radiation processes. As a result, the multi-scale approach significantly reduces the need for physical parameterizations by resolving those processes directly via truncated large-eddy simulations (Khairoutdinov and Randall 2001). Multi-scale approaches are constructed on the premise that there exists a scale separation that can be exploited in the modeling of the physical system. Essentially, this approach assumes that the fine-scale processes act on temporal and spatial scales that are sufficiently far away from the coarse-grain processes such that the fine and coarse scales can be coupled without a representation of the intervening scales. The extent to which this assumption is valid for the atmosphere and ocean systems remains unclear.

In this contribution, we start what we hope will be a fourth line of research to address the

computational challenges in modeling the climate system. This approach, that we informally refer to as a multi-resolution approach, is essentially a merging of the traditional global climate modeling approach with the regional limited-area approach. As will be discussed below, in our multi-resolution approach we maintain a global modeling framework in the sense that we simulate the entire spatial extent of the atmosphere and/or ocean systems within a single model, yet we allow for arbitrary regions of local mesh refinement.

In the sense that this method maintains a global, conforming mesh, it is similar to the stretch-grid or conformal mapping approaches that have been explored over the last two decades (Fox-Rabinovitz et al. (1997), Déqué et al (2005) and Fox-Rabinovitz et al. (2006)). Since the stretched-grid approaches require the mesh to be deformed through a continuous mapping (i.e. the mesh is topologically unchanged as resolution is changed), the increase in resolution in one region must necessarily come at the expense of decreasing resolution in another region. Stretched-grid approaches are also limited in their ability to place enhanced resolution in more than one region. The multi-resolution approach developed below is not based on a continuous deformation of a mesh, does not require that the increase in resolution in any region come at the expense of resolution elsewhere and is not limited to resolution increases in only one region. The stretched-grid approach does highlight a primary challenge of any method that includes a wide range of spatial scales, namely the lack of access to scale-aware physical parameterizations. We revisit this challenge, along with the other challenges that multi-resolution approaches must address, in the last section of the paper.

As illustrated in Figure 1, the multi-resolution approach allows for grid-resolution in one or more regions to be significantly higher than the grid resolution in other regions. This can be acomplished in one of two ways. First, a variable density function could be employed

to redistribute a fixed number of grid points, causing the same effect as a stretched-grid approach. Second, using a set of grid points, an arbitrary number of refinement nodes can be added into the grid, causing refinement only in the area of interest, without hindering the results in other areas. We have the ability to directly simulate processes, such as clouds and ocean eddies, in the region(s) of high resolution while parameterizing those same processes in the region(s) of low resolution. This multi-resolution approach is built upon two key components: a conforming, variable-resolution mesh with exceptional mesh-quality characteristics and a finite-volume method that maintains all of its conservation properties even when implemented on a highly non-uniform grid.

The first of the two pillars on which this multi-resolution approach is build on is Spherical Centroidal Voronoi Tessellations (SCVTs). Voronoi tessellations have a long history in the sciences, probably because Voronoi-like polygons are commonly found in nature (Barlow 1974). In climate modeling, Voronoi-like tessellations were introduced by Sadourny et al. (1968) and Williamson (1968) due to their uniformity and isotropy in tiling the surface of the sphere.[1] Neither Sadourny nor Williamson refer to their grids as Voronoi tessellations, but both appeared to use Voronoi tessellations as their base mesh. Even over the last decade there has been much ambiguity with respect to the terminology used to describe these meshes (see Ju et al. (2010)). More recently, Voronoi-like meshing of the sphere has found significant success in global atmosphere modeling (Heikes and Randall (1995), Thuburn

[1]Voronoi tessellations have been reinvented many times over in the past hundred and fifty years. The first systematic treatment of what we now call Voronoi tessellations was given by Dirichlet (1850). Voronoi (1908) generalized the work of Dirichlet to arbitrary dimensions. These tessellations have been given many different names by their reinventors (see Ju et al. (2010))

(1997), Ringler et al. (2000), Ringler and Randall (2002), Randall et al. (2002), Tomita et al. (2005), Weller and Weller (2008)). In each of these examples, the use of Voronoi-like tessellations is motivated through their ability to produce high-quality meshes of uniform resolution while at the same time eliminating problematic grid singularities associated with other meshing approaches. While we certainly agree and appreciate this motivation, recent work suggests that Voronoi diagrams are equally valuable for the generation of variable resolution meshes. As will be discussed fully in Section 2, by adding the centroidal constraint to the construction of Voronoi tessellations we can produce a very regular, high-quality, variable-resolution meshing of the sphere. A *centroidal* Voronoi tessellation differs from the generic Voronoi tessellation by requiring that the generating points (grid points) are the centroids (centers of mass) of the corresponding Voronoi regions. This seemingly minor requirement that grid points be the centers of mass of the Voronoi grid cell results in meshes of remarkably high quality even when the mesh resolution changes (Gersho (1979), Du et al. (1999)).

The second pillar of this approach is the finite-volume scheme that we pair with the variable-resolution SCVTs to produce robust simulations of rotationally-dominated geophysical flows. A hallmark of robust finite-volume techniques used in global atmosphere and ocean models has been their ability to constrain the spurious growth of nonlinear quantities, such as potential enstrophy and total energy (Arakawa 1966). While a more nuanced view of the importance of conserving nonlinear quantities has emerged over the last decade (Thuburn 2008), anecdotal evidence has continually shown that there is value in developing numerical schemes that respect certain underlying constraints imposed by the continuous system. This is a particularly challenging task when the underlying mesh is not uniform

(e.g. see Perot (2000), Bonaventura and Ringler (2005), Stuhne and Peltier (2006), Ham et al. (2007), Kleptsova et al. (2009)). The recent contributions from Thuburn et al. (2009) and Ringler et al. (2010) detail a finite-volume approach that allows for the conservation of nonlinear quantities, even when the underlying mesh is highly variable. One purpose of this contribution is a full characterization of this scheme's performance on variable resolution meshes.

Our goals for this contribution are modest in the sense that we only wish to characterize the ability of this approach to simulate the shallow-water system with multi-resolution meshes. Such a characterization is, in our view, a prerequisite to performing variable-resolution simulations of the full atmosphere and ocean systems. We choose to begin with the analysis of the shallow-water system due to its proven usefulness as a simplified proxy of the 3D primitive equations. To this end, in Section 2 we provide a brief overview of the SCVTs, their properties and how these meshes are generated. In Section 3 we provide a brief summary of the underlying numerical method used in our multi-resolution approach with special attention toward the method's properties when the mesh is non-uniform. Results from a few of the standard shallow-water test cases are shown in Section 4 where the focus is on geostrophic balance, conservation properties and solution error as a function of mesh size and mesh refinement. In Section 5 we compare the results obtained herein with previously published results. The multi-resolution approach that we begin to develop here is not without its own set of challenges. In Section 6 we highlight the challenges that will have to be overcome if this approach is to make substantive contributions to the field of global and regional climate modeling.

2. Properties and Generation of SCVTs

A full review of SCVTs and their potential benefit in global climate system modeling is provided in Ju et al. (2010) and Ringler et al. (2008). Our discussion here is restricted to the most salient aspects of SCVTs with a focus on the practical aspect of the meshes. The analysis that yields these practical results is not discussed, but is referenced for those interested in better understanding the mathematical underpinning of this mesh generation technique.

Voronoi diagrams can be specified as follows: We are given a bounded domain $\Omega \in \mathbb{R}^d$ and a set of distinct points $\{\mathbf{x}_i\}_{i=1}^n \subset \Omega$. For each point $\mathbf{x}_i$, $i = 1, \ldots, n$, the corresponding *Voronoi region*, V_i, $i = 1, \ldots, n$, is defined by

$$V_i = \big\{\mathbf{x} \in \Omega \;\mid\; \|\mathbf{x} - \mathbf{x}_i\| < \|\mathbf{x} - \mathbf{x}_j\| \quad \text{for} \;\; j = 1, \cdots, n \quad \text{and} \quad j \neq i\big\}, \tag{1}$$

where $\|\cdot\|$ denotes the geodesic distance measured along the surface of the sphere. Clearly $V_i \cap V_j = \emptyset$ for $i \neq j$, and $\cup_{i=1}^n \overline{V}_i = \overline{\Omega}$ so that $\{V_i\}_{i=1}^n$ is a *tessellation* of Ω, i.e. $\cup_{i=1}^n \overline{V}_i$ spans $\overline{\Omega}$ with a non-overlapping mesh. We refer to $\{V_i\}_{i=1}^n$ as the *Voronoi tessellation* or *Voronoi diagram* of Ω associated with the point set $\{\mathbf{x}_i\}_{i=1}^n$. In the nomenclature of Voronoi diagrams, a point $\mathbf{x}_i$ is called a *generator* and a subdomain V_i is referred to as the *Voronoi region* or *Voronoi cell*. Each generator is uniquely associated with a single Voronoi region. For our purposes, generator points are equivalent to grid points and Voronoi regions are equivalent to grid cells. If the domain $\Omega \in \mathbb{R}^d$ spans all or part of the surface of the sphere, then we refer to the mesh as a spherical Voronoi tessellation.

A spherical Voronoi tessellation becomes a spherical *centroidal* Voronoi tessellation when

the *generators* are also *centers of mass* of the corresponding Voronoi region. Given a density function $\rho(\mathbf{x}) \geq 0$ defined on Ω, for any region $V \subset \Omega$, the standard *mass center* $\mathbf{x}^*$ of V is given by

$$\mathbf{x}^* = \frac{\int_V \mathbf{x}\rho(\mathbf{x})\, \mathrm{d}\mathbf{x}}{\int_V \rho(\mathbf{x})\, \mathrm{d}\mathbf{x}}. \tag{2}$$

This center-of-mass calculation will always result in an $\mathbf{x}^*$ that lies inside the surface of the sphere. In order to constrain the generator points to lie on the unit sphere, $\mathbf{x}^*$ is radially-projected on to the surface of the unit sphere. In general, the $\mathbf{x}^*$ for each grid cell does not correspond to grid point $\mathbf{x}_i$ of that cell. Only when $\mathbf{x}^* \equiv \mathbf{x}_i$ is the spherical Voronoi tessellation also a spherical *centroidal* Voronoi tessellation.

In practice, finding a SCVT given any SVT is a relatively straightforward, iterative process based on Lloyd's algorithm (Du et al. 1999). Given a set of $\mathbf{x}_i$, we first find the corresponding V_i and compute $\mathbf{x}_i^*$ for each V_i. In general, $\mathbf{x}_i^* \neq \mathbf{x}_i$, so we simply move generators to be the centroids with $\mathbf{x}_i = \mathbf{x}_i^*$ and repeat the process. The iterative procedure continues until $\mathbf{x}_i^*$ and $\mathbf{x}_i$ are deemed sufficiently close based on, say, the L_2 or L_{inf} norms. For a more detailed discussion of this iterative procedure, restrictions on ρ, the guarantees related to convergence and the optimality of the resulting mesh see, for example, Du et al. (1999), Du et al. (2003) or Ringler et al. (2008). While we only interested here in the extension of CVT to SCVT, the CVT-approach can be generalized to any manifold or surface, see Du et al. (2003).

The power of an approach based on SCVTs resides in the freedom to specify $\rho(\mathbf{x})$ and, thereby, control the local grid resolution and local grid variation with a high degree of

precision. If we pick any two Voronoi regions and arbitrarily index them with i and j, then the conjecture is

$$\frac{dx_i}{dx_j} \approx \left[\frac{\rho(\mathbf{x}_j)}{\rho(\mathbf{x}_i)}\right]^{1/(d'+2)} \tag{3}$$

where d' is the dimension of the manifold on which the tessellation is constructed, $\rho(\mathbf{x}_i)$ is the density function evaluated at $\mathbf{x}_i$ and dx_i is a measure of the local mesh spacing in the vicinity of the $\mathbf{x}_i$. Similarly for $\rho(\mathbf{x}_j)$ and dx_j. While (3) remains an open conjecture for $d' \geq 2$, its validity has been supported through many numerical studies. In our grid generation examples below, we demonstrate the accuracy of (3) and provide evidence for our assertion that we have precise control on the relative mesh spacing in different parts of Ω through the choice of ρ. Equation (3) becomes even more powerful when paired with Gersho's conjecture. Asymptotically and for a fixed density function, as the number of generators becomes larger and larger, Gersho's conjecture (Gersho 1979) states that the tessellation becomes more and more regular in the sense that, locally, the tessellation converges to a replication of a polytope. In other words, Gersho's conjecture states that if the number of generators n is large enough and one focuses on a small enough region, then a centroidal Voronoi tessellation appears to be a uniform mesh involving congruent polytopes. The regular hexagon provides a confirmation of the conjecture in two dimensions for the constant density case (Newman 1982).

The rigorous application of Gersho's conjecture to tessellating the surface of the sphere fails since we know that no regular single polytope can be used to tessellate the sphere (Saff and Kuijlaars 1997). Yet the spirit of Gersho's conjecture does carry over to the sphere;

for a given density function as the number of generators is increased the resulting meshes are composed, proportionally, of more hexagons that converge uniformly toward regular hexagons. Both Ju et al. (2010) and Ringler et al. (2008) demonstrate this in a variety of settings.

In summary, the utility of SCVTs resides in the ability to precisely control grid resolution through the specification of the density function as described in (3) and the guarantee that the meshes will become more regular as the number of grid points is increased.

3. Example SCVTs

The simulations discussed below will employ meshes sampled from a three-parameter density function expressed as

$$\rho(\mathbf{x}_i) = \frac{1}{2\,(1-\gamma)} \left[\tanh\left(\frac{\beta - \|\mathbf{x}_c - \mathbf{x}_i\|}{\alpha} \right) + 1 \right] + \gamma \tag{4}$$

where $\mathbf{x}_i$ is constrained to lie on the surface of the unit sphere. This function results in relatively large value of ρ within a distance β of the point $\mathbf{x}_c$ where β is measured in radians and $\mathbf{x}_c$ is also constrained to lie on the surface of the sphere. The function transitions to relatively small values of ρ across a radian distance of α. The distance between $\mathbf{x}_c$ and $\mathbf{x}_i$ is computed as $\|\mathbf{x}_c - \mathbf{x}_i\| = \cos^{-1}(\mathbf{x}_c \cdot \mathbf{x}_i)$ with a range from 0 to π.

The density function is constructed such that it has a maximum value of 1 and a minimum value of γ, where $\gamma > 0$. Based on (3) we know that the mesh spacing in the high resolution

region, dx_f, and the mesh spacing in the low resolution region, dx_c, will be related as

$$\frac{dx_f}{dx_c} \approx \gamma^{\frac{1}{4}}. \tag{5}$$

For this study we fix $\beta = \pi/6$ and $\alpha = \pi/20$. For reasons that will be clear below, we specify the location of $\mathbf{x}_c$ to coincide with the center of the orographic feature present in Shallow-Water Test Case 5 (Williamson et al. 1992). Our focus will be on the impact of γ, i.e. the impact of the relative resolution between the fine-mesh region and the coarse-mesh region. Figure 1 shows meshes that were generated with 2562 grid points based on γ values of $(1)^4$, $(1/2)^4$, $(1/4)^4$ and $(1/16)^4$. We refer to these meshes as the X1, X2, X4 and X16 meshes since fine-mesh and coarse-mesh resolutions vary by ratios of 1, 2, 4 and 16, respectively. The simulations discussed below will also use a X8 mesh that is not shown in Figure 1. The X1 through X16 meshes are generated with 2562, 10242, 40962, 163842 and 655362 grid points. As a result of this choice of grid points, the X1 meshes are very similar to other Voronoi-like meshes that are derived from the recursive bisection of the icosahedron. We made this choice in order to facilitate comparison of the error norms computed below to error norms already found in the published literature. Table 1 summarizes the resolutions of all of the meshes used in this study.

Figure 2 shows the distribution of mesh resolution measured in the vicinity of each grid cell as a function of geodesic distance from $\mathbf{x}_c$. At each grid cell we define the local grid resolution, dx_i as

$$dx_i = \frac{1}{n_i} \sum_{j=1}^{n_i} \|\mathbf{x}_j - \mathbf{x}_i\| \tag{6}$$

where $\mathbf{x}_j$ are the across-edge neighbors of grid cell i (see Figure 3). dx_i represents the average

distance between grid point x_i and all of its nearest neighbors. Also shown in Figure 2 is the theoretical estimate of local mesh resolution for the X1, X2, X4, X8 and X16 meshes based on (3).

Figure 2 confirms that the theoretical estimate of local grid resolution is remarkably accurate; the mesh spacing as computed from the meshes essentially falls on top of the theoretical estimate.

4. Summary of Numerical Method

This study focuses on the nonlinear shallow-water equations expressed as

$$\frac{\partial h}{\partial t} + \nabla \cdot (h\mathbf{u}) = 0, \tag{7}$$

$$\frac{\partial \mathbf{u}}{\partial t} + \eta \mathbf{k} \times \mathbf{u} = -g\nabla (h + b) - \nabla K, \tag{8}$$

where h represents the fluid layer thickness and $\mathbf{u}$ represents the fluid velocity along the surface of the sphere. The absolute vorticity, η, is defined as $\mathbf{k} \cdot (\nabla \times \mathbf{u}) + f$ and the kinetic energy, K, is defined as $|\mathbf{u}|^2 /2$. At all points on the surface of the sphere the vector $\mathbf{k}$ points in the local vertical direction and we require $\mathbf{k} \cdot \mathbf{u} = 0$ at all points. The three parameters in the system are gravity, g, Coriolis parameter, f, and bottom topography, b.

For our application, a more appropriate form of the continuous equations is expressed as

$$\frac{\partial h}{\partial t} + \nabla \cdot \mathbf{F} = 0, \tag{9}$$

$$\frac{\partial \mathbf{u}}{\partial t} + q\mathbf{F}^{\perp} = -g\nabla (h + b) - \nabla K, \tag{10}$$

where $\mathbf{F} = h\mathbf{u}$, $\mathbf{F}^{\perp} = \mathbf{k} \times h\mathbf{u}$ and $\eta = hq$ where q is the total potential vorticity.

A numerical method used to model the shallow-water system is discussed at length in Thuburn et al. (2009) (T09 hereafter) and Ringler et al. (2010) (R10 hereafter). In T09 an analysis of the linearized version of (7) and (8) is conducted in order to derive a numerical method that is able to reproduce stationary geostrophic modes found in the continuous system, even when the numerical method is implemented on variable resolution meshes such as those shown in Figure 1. In R10, the analysis is extended to the nonlinear shallow-water equations shown in (9) and (10) in order to derive a method that conserves total energy and potential vorticity while allowing for a physically-appropriate amount of potential enstrophy dissipation. While the analyses and derivations in both T09 and R10 are for any mesh that is a Vororoi tessellation, the numerical simulations presented in both of those papers only evaluate the method when implemented on a quasi-uniform mesh.

The numerical scheme is a standard finite-volume method with a C-grid staggering as shown in Figure 3. The thickness field is defined on the Voronoi cells while all vorticity-related fields, such as relative vorticity, absolute vorticity and potential vorticity, are defined on the Delaunay triangles. The discrete thickness equation is obtained by simply supplying a discrete approximation to the divergence operator (See R10 Figure 3). As with all C-grid methods, only the component of velocity in the direction normal to the thickness finite-volume cell is prognosed. To derive this normal-component velocity equation, the inner product of $\mathbf{n}_e$ (shown in Figure 3) and (10) is computed at each edge location. The resulting discrete system of equations is then expressed as

$$\frac{\partial h_i}{\partial t} = -\left[\nabla \cdot F_e\right]_i , \tag{11}$$

$$\frac{\partial u_e}{\partial t} + F_e^{\perp}\widehat{q}_e = -\left[\nabla \left(g(h_i + b_i) + K_i\right)\right]_e \tag{12}$$

where $F_e = h_e u_e$ represents the mass flux across the edge of a Voronoi cell and $F_e^{\perp}$ represents the mass flux across the edge of each Delaunay cell. The discrete approximations of the divergence and gradient operator are shown in R10 Figure 3. In (11) and (12) the yet-to-be-defined fields are K_i, h_e, $\widehat{q}_e$ and $F_e^{\perp}$. These fields are defined following R10 without exception. Also following R10, we use the anticipated potential vorticity method (Sadourny and Basdevant 1985) to dissipate potential enstrophy.

The culmination of the derivations in T09 and R10 is a numerical method that conserves total energy to within time-truncation error, conserves total potential vorticity to within machine round-off error and dissipates potential enstrophy at a rate that depends on a single parameter. This derivation was carried out for a general Voronoi mesh; the results in Section 5 are intended to confirm this analysis.

5. Results

Through the use of three shallow-water test cases, we confirm the derivations in T09 and R10 related to system energetics, geostrophic balance and potential vorticity dynamics. Shallow-Water Test Case 5 (SWTC5) and Shallow Water Test Case 2 (SWTC2) from Williamson et al. (1992) (hereafter W92) are used primarily to confirm the numerical methods ability to mimic conservation properties and maintain geostrophic balance, respectively.

A final test case, the Barotropic Instability Test Case, is used to illustrate the method's ability to allow prototypical structures of the atmosphere and ocean to enter and exit mesh transition zones (Galewsky et al. 2004).

Along the way we compute L_2 error norms of the thickness field, h_i, in order to better understand how the solution error varies with the amount of mesh variation. The L_2 norm is computed as

$$L_2 = \frac{\left\{ S\left[(h_i - h_i^r)^2 \right] \right\}^{\frac{1}{2}}}{\left\{ S\left[(h_i^r)^2 \right] \right\}^{\frac{1}{2}}}. \tag{13}$$

The field h_i^r is the reference solution that has been calculated at or interpolated to $\mathbf{x}_i$ positions. The reference solution represents either an analytic solution or, if an analytic solution is not available, a high-resolution solution. The function $S[f]$ computes the area-weighted average of f over the entire sphere.

Twenty-five simulations are conducted for each test case, thus filling the [grid points $\times$ mesh variation] matrix shown in Table 1. Every simulation in every test case is conducted with the exact same executable with the exact same parameter settings. The spatial discretization discussed above is paired with a $4^{th}-$order Runge-Kutta time stepping method using a time step of $dt = 25\,s$. Each simulation employs the anticipated potential vorticity method with the upwind-bias parameter θ set to $dt/2$ (see Sadourny and Basdevant (1985) Eq. 8). All simulations are conducted with 64-bit floating point arithmetic.

a. SWTC5

As explained in the Introduction, our long-term goal is the creation of full-physics, multi-resolution models of the global atmosphere and ocean systems. Our motivation for evaluating this approach in the shallow-water system is to identify, to the extent possible, the strengths and weaknesses in an idealized setting. We begin the analysis with SWTC5 because it offers an analog to what we hope to accomplish in more realistic settings. SWTC5 contains a single feature (orography) that is completely responsible for the transient evolution of the system. While the orography is large-scale, it is localized and, in that sense, is conducive to local mesh refinement. To greater and lesser extents, all of the meshes depicted in Figure 1 and Table 1 enhance resolution in the vicinity of the orography.

SWTC5 prescribes an analytic initial condition of large-scale geostrophic flow that would be in steady state, if not for the presence of an orographic feature. The orographic feature is centered at $\mathbf{x}_c$ and extends $\pi/9$ radians in latitude and longitude. Recall that the variable resolution meshes developed in Section 3 are also centered at $\mathbf{x}_c$ and extend the fine-mesh region a distance of $\pi/6$ radians; the fine-mesh region includes all of the orography.

The analytic initial condition is mapped to the discrete model by sampling Eq. (95), with the appropriate constants for SWTC5, from W92 at Voronoi grid points (i.e. $\mathbf{x}_i$ locations) to determine the initial thickness fields. The initial u_e field is obtained by determining the streamfunction via Eq. (92) from W92 at Delaunay grid points (i.e. $\mathbf{x}_v$ locations), then computing u_e as $\mathbf{k} \times \nabla\psi$. Even though errors in u_e are present at $t = 0$, this approach guarantees that the discrete divergence is identically zero at $t = 0$.

As a result of the orography, the geostrophically-balanced zonal flow impinges on the

mountain at $t = 0$, resulting in the radiation of gravity and Rossby waves as the flow adjusts to the presence of the orographic feature. The interaction between the zonal flow and the orography leads to strong nonlinearity, which is why this test case is chosen to assess the numerical method's conservation properties.

We begin with a qualitative assessment of SWTC5 by showing in Figure 4 the fluid height field, $h_i + b_i$, at day 15 for the X1, X2, X4 and X16 meshes using 40962 cells. Broadly, the simulations are identical as depicted by Figure 4. Since the flow is characterized by large-scale Rossby waves that are well resolved by the full suite of meshes using 40962 cells, we would expect that the simulations to be qualitatively similar. The coarse grid resolution in regions far removed from the orography is clearly seen in Figure 4. Note that while the simulation with the 40962/X16 mesh ranges in resolution from 40 km in the vicinity of the orography to 611 km elsewhere, there is not hint of noise in the mass field, even through the mesh transition zone.

Two quantities are conserved to round-off error in every simulation: the area-weighted global sum of thickness and the volume-weighted potential vorticity. Specifically we find

$$\frac{\partial}{\partial t} V = \frac{\partial}{\partial t} \sum_{i=1}^{N_i} h_i A_i = 0, \tag{14}$$

$$\frac{\partial}{\partial t} \sum_{v=1}^{N_v} q_v h_v A_v = 0, \tag{15}$$

to within round-off error in all simulations, where the quantity V represents the total fluid volume.

In order to evaluate the energetics of the system, the total energy is computed following

R10 Eq. (70) as

$$E = \sum_e A_e \left[\frac{\widehat{h}_e u_e^2}{2} \right] + \sum_i A_i \left[g h_i \left(\frac{1}{2} h_i + b_i \right) \right] - E_r. \tag{16}$$

In the computation of total energy, the unavailable potential energy, E_r, with the form

$$E_r = \sum_i g \, \bar{H}_i \, A_i \left[\frac{\bar{H}_i}{2} + b_i \right] \tag{17}$$

where

$$\bar{H}_i = \frac{\sum_i A_i \, (h_i + b_i)}{\sum_i A_i} - b_i \tag{18}$$

has been subtracted; hereafter references to "total energy" imply "total available energy". E_r represents the potential energy of the fluid at rest. Figure 5 demonstrates the degree to which total energy is conserved in the simulations. The figures show $\log_{10} \frac{|(E(t)-E(0))|}{E(0)}$ over the 15 day integration for the X1, X2, X4, X8 and X16 meshes with 40962 grid points. Figure 5 measures the extent to which the sum of available potential energy and kinetic energy is conserved. At day 15, all solutions conserve total energy to within 1.0×10^{-8} relative to total energy present at $t = 0$; this is orders of magnitude better than is required when considering the dissipation mechanisms present in the real atmosphere and ocean (Thuburn 2008).

The total energy is conserved in the physically-appropriate manner; the nonlinear Coriolis force neither creates nor destroys kinetic energy and the exchange of energy between its potential and kinetic forms is equal and opposite. We evaluate the degree to which the nonlinear Coriolis force is energetically-neutral by computing the time it would take for the nonlinear Coriolis force to double the kinetic energy in the system. With 40962 grid

points, the time required for the nonlinear Coriolis force to double the kinetic energy is approximately 10^4 years for all meshes. This finding is consistent with Figure 4 of R10.

The other important component in the total energy budget is the conservative exchange of energy between its potential and kinetic forms. The potential and kinetic energy equations each have a source term. These source terms are equal and opposite (see, for example, Eqs. (15) and (16) of R10). We evaluate the source term for kinetic and potential energy following Eqs. (65) and (67), respectively, from R10. Since these RHS sources are algebraically equivalent in the discrete system, we expect a very high degree of cancellation between the sources. All 25 simulations show that the time scale for doubling the kinetic energy of the system due to the imperfect cancellation of KE and PE sources terms to be approximately 10^{10} years. This is essentially machine precision round-off error.

In regards to conservation, the final quantity of interest is potential enstrophy. Figure 6 shows $\log_{10} \frac{|(R(t)-R(0))|}{R(0)}$ where R is the global-integrated potential enstrophy defined as

$$R = \frac{1}{V} \sum_{v=1}^{N_v} q_v^2 h_v A_v - R_r. \tag{19}$$

Just as energy has an unavailable reservoir, potential enstrophy has a unavailable reservoir that is equal to the amount of potential entrophy that exists when the fluid is at rest. This unavailable reservoir, R_r is removed from the computation in order to obtain a more representative evaluation of potential enstrophy conservation.

Figure 6 shows the relative change in globally-averaged potential enstrophy for the X1, X2, X4, X8 and X16 meshes with 40962 nodes. At day 15, the relative changes in globally-averaged potential enstrophy vary between 10^{-4} and $10^{-2.5}$ for the X1 and X16 meshes,

respectively. In these simulations, the X1 and X2 simulations show a monotonic decrease on globally-averaged potential enstrophy, while the X8 and X16 simulations show a monotonic increase in globally-averaged potential enstrophy. The X4 simulation fluctuates about its initial globally-averaged value. Clearly the amount of potential enstrophy dissipation provided by the anticipated potential vorticity method needs to vary with mesh resolution; this is discussed further in Section 7.

In terms of formal L_2 global error norms, previous work using local mesh refinement with the shallow-water system all find that the solution error is relatively unchanged when adding resolution in a specific region (e.g. Weller et al. (2009), St-Cyr et al (2008) and Chen et al (2011); see next section for a full discussion). Stated alternatively, previous work has found that the solution error is primarily controlled by the coarse region of the mesh when using static mesh refinement. In order to test if this is the case in our simulations, we plot the global L_2 error norm for each of the 25 simulations *as a function of coarse-mesh resolution* in Figure 7. Since SWTC5 does not have a known analytic solution, error norms are computed with respect to a T511 global spectral model (Swarztrauber 1996). For TC5 at T511, the global spectral model requires a scale-selective ∇^4 dissipation of $8.0 \times 10^{12}\, m^4/s$ in order to prevent the accumulation of energy and potential enstrophy at the grid scale.

Figure 7 shows that the solution error is controlled by the mesh resolution in the coarse region. All of the simulations show the same convergence rate of approximately 1.5. Note that we have plotted these errors norms on a $log - log$ scale to emphasis the primary finding that the L_2 error is controlled by the coarse-mesh resolution. If we parse the results more closely, we find that the variable resolution meshes provide a small, but measurable, improvement in solution error, i.e. adding degrees of freedom in the vicinity of the orography,

while holding the coarse-mesh resolution fixed, results in a small reduction in the error norm.

b. SWTC2

Having confirmed the ability of the numerical model to simulate transient flows in a robust manner with SWTC5, we now use SWTC2 to measure the method's ability to maintain large-scale geostrophic balance. SWTC2 prescribes an analytic initial condition that is an exact, steady-state solution to (9) and (10). The analytic initial condition is mapped to the discrete model by sampling Eq. (95) from W92 at Voronoi grid points (i.e. $\mathbf{x}_i$ locations) to determine the initial thickness fields. As with SWTC5, the initial u_e field is obtained by determining the streamfunction via Eq. (92) from W92 at Delaunay grid points (i.e. $\mathbf{x}_v$ locations), then computing u_e as $\mathbf{k} \times \nabla\psi$. Any deviation of the numerical solution from its initial condition is considered to be numerical error.

While SWTC5 offers a plausible reason for mesh refinement, no comparable reason is present in SWTC2. The motivation for evaluating our multi-resolution method using SWTC2 is not to demonstrate the approaches utility, but rather to measure the cost of mesh refinement. Maintaining large-scale balance is an important property of any numerical model of the atmosphere or ocean. SWTC2 provides the opportunity to precisely measure, through the L_2 error norm, the impact of mesh refinement on maintaining geostrophic balance.

Following our finding in SWTC5 that global error is controlled by the coarse mesh resolution, Figure 8 plots the global L_2 error for all 25 simulations against the resolution in the coarse-mesh region. As found with SWTC5, essentially all of the variation in the L_2 error in the simulations is controlled by the coarse resolution grid spacing. For a given coarse

resolution, solution error increases by approximately a factor of 2 between the X2 and X16 meshes. In contrast, the solution error for the X1 mesh is approximately a factor of 10 smaller, regardless of the coarse mesh resolution.

Each grid point in the X1 mesh is uniquely associated with a node produced when generating a mesh through the recursive bisection-projection of an inscribed icosahedron[2] (Heikes and Randall 1995). This method results in a particularly uniform distribution of grid points resulting in a relatively small solution error. This special distribution of nodes is lost when producing the variable-resolution meshes. As a result, we incur a relatively large cost, in terms of global error, by choosing to move away from the special quasi-uniforms meshes, but incur very little additional cost by increasing the extent of the mesh variation.

The rate of convergence for SWTC2 is not uniform. Meshes with minimum grid resolutions above 100 km show a convergence rate of approximately 1.9 with respect to the coarse mesh resolution. As the minimum resolution of the mesh becomes smaller and smaller, the rate of convergence becomes smaller. The likely culprits for this reduction in convergence rate are the following: deficiencies in the structure of the grids, deficiencies in the manner in which we compute the error norms, deficiencies in the numerical model. We have been unable to definitely rule out any of these possibilities and continue to seek the underlying cause of this issue. We full expect that 2nd-order convergence rate to continue indefinitely as resolution is increase.

[2]While the X1 meshed is topologically equivalent to a mesh produced through the recursive bisection-projection of an inscribed icosahedron, the actual positions of the nodes on the unit sphere differ because in our system we move the nodes so that the resulting mesh is a SCVT.

c. Barotropically-Unstable Zonal Jet

The final test case to be discussed is the growth a barotropic instability on a zonally-symmetic zonal jet (Galewsky et al. 2004) (G04 hereafter). In order to generate the initial conditions for this test case, we derive a streamfunction from G04 Eq. 2. This streamfunction is sampled at vertex locations and the initial u_e field is computed analogous to SWTC2 and SWTC5. The initial thickness field is computed based on G04 Eq. 3 and we include the height perturbation shown in G04 Eq. 4.

Figure 9 shows the relative vorticity field at day 6 for the X1, X2, X4, X8 and X16 meshes with 655362 cells. The fine-mesh region is coincident with the center of each panel. In addition, the envelope of the growing barotropic instability is roughly coincident with the fine mesh region at day 6, with parts of the wave system entering and exiting the fine-mesh region at this point in time.

Conducting test cases based on instabilities that grow on a zonally-symmetric base state is particularly challenging for our modeling system. Specification of the test case is zonally symmetric and the instability is triggered by a small amplitude perturbation. The meshes used in this study are not zonally-symmetric and, as a result, lead to truncation error projecting onto non-zero zonal wave numbers. This truncation error serves as an additional trigger for the instability and can lead to wave growth that is either too fast or not in the correct location. As the resolution is increased, the amplitude of the spurious forcing by truncation error diminishes and the instability is solely controlled by the perturbation contained in the initial conditions.

In addition, the growth of the unstable waves depends strongly on the type and strength

of the sub-grid scales closures that are either implicit in the underlying numerical formulation or explicitly added to the numerical models. For example, the X1 panel in Figure 8 agrees very closely with panel D in Figure 17 of Ii and Xiao (2010), but is significantly different than panel D in Figure 9 of G04. This is because the simulations presented here and in Ii and Xiao (2010) do not use any explicit closure, whereas G04 uses hyper-diffusion on the RHS of the momentum equation.

The strong correspondence of our X1 simulations with panel D in Figure 17 of Ii and Xiao (2010) indicates that the X1 simulation is broadly representative of the instability when simulated in a minimally or undamped system. Our primary purpose here is to understand how the use of variable resolution meshes alters the growth of the barotropic instability.

First, if we focus on the deep, tilted trough just right of center in each panel along with the ridge-trough-ridge system just upstream to the west we find that these dominate features are present in all simulations with the same amplitude and phase. The X2 simulation is qualitatively equivalent to the X1 simulation in all respects. In addition, the X8 simulation is qualitatively equivalent to the X4 simulation in all respects. The X4 simulation differs from the X2 simulation only along the edges of the panels that corresponds to the center of the coarse-mesh regions. The primary difference between these two groups of simulations is that the X4/X8 simulations produce an additional ridge in the upstream wave. The X16 simulation is qualitatively different from the other simulations in all regions other than the fine-mesh region. The X16 simulation produces a relatively strong ridge-trough systems in the coarse-mesh region that are not present in the other simulations. It is important to note that the fine-mesh resolution of the X8 and X16 simulations is essentially the same at approximately 10 km, yet the coarse-mesh resolution between these same two simulations differ by

a factor of two (see Table 1). The X16/655362 simulation is more similar to the X1/40962 simulation (not shown) than any of the other simulations with 655362 nodes. Since the coarse resolution of the X16/655362 simulation is comparable to the X1/40962 simulation, this finding is consistent with Figures 7 and 8 which demonstrate that the accuracy of the simulation is controlled primarily by the resolution in the coarse-mesh region.

6. Comparison to Previous Results

Our introduction emphasized that there are several approaches to regional climate simulation that are being actively explored. Given the diversity of existing approaches and the novelty of the approach discussed herein, an obvious question is how the results obtained in the previous section compare to other published results.

Unfortunately, the literature is sparse with respect to the evaluation of regional modeling approaches using the standard shallow-water test cases. For example, while full-physics, 3D, regional climate simulations employing the limited-area modeling approach have been conducted over the last two decades, we have been unable to find any results where the limited-area method has been evaluated using of the standard shallow-water test cases. With regard to the stretched-grid and conformally-mapped grid approaches, we are also unable to find evaluations of the methods within the context of the shallow-water modeling system. We note that the situation is exactly the opposite with respect to numerical methods evaluated using global, quasi-uniform meshes in the shallow water system; in this case the literature is extremely rich. But the comparison of the numerical scheme proposed here, when paired with quasi-uniform meshes, has already appeared in R10. Furthermore, the

literature that does exist is primarily focused on dynamic adaptivity, whereas our focus is on static adaptivity. In what follows, we compare our multi-resolution simulations to previously published findings presented in three manuscripts: Weller et al. (2009), St-Cyr et al (2008) and Chen et al (2011).

The cleanest and most useful comparison of our results is with Weller et al. (2009) (W09, hereafter). W09 focuses on static mesh refinement, employs a finite-volume approach based on AtmosFOAM (Weller and Weller, 2008) and utilizes variable-resolution meshes based on Voronoi tessellations, Delaunay triangles and quadilateral polygons. In addition, the spatial location of the mesh refinement used in W09 and herein is the same; both place mesh refinement in the vicinity of the orographic feature present in SWTC5 (e.g. see W09 Fig. 4). One difference between our work and W09 is the extent of mesh refinement; we employ meshes that vary in resolution by a factor of 16, whereas W09 use meshes that vary in resolution by a factor of 2. We also explore meshes with approximately 5×10^5 cells, whereas W09 uses meshes of significantly lower resolution with 1×10^4 cells. In terms of accuracy, the results presented in Figures 7 and 8 show error norms that are approximately a factor of five more accurate than W09 for SWTC5 and SWCT2, respectively. While Figure 7 shows that the error is controlled almost entirely by the coarse mesh resolution with a small gain received for adding more degrees of freedom in the fine mesh region, W09 find that errors increase slightly for all meshes when extra resolution is added around the mountain.

St-Cyr et al (2008) (S08, hereafter) evaluate two numerical methods in the context of dynamically-adaptive mesh refinement. One method uses a high-order, spectral element method while the other uses a standard finite-volume method. S08 conduct SWTC2 with static mesh refinement resulting a mesh that varies in resolution by a factor of 4 in grid

spacing, with a coarse mesh grid spacing of approximate 250 km. The extent of the refined region is 30 degrees in latitude and 45 degrees in longitude and covers approximately 3% of the surface of the sphere. This region of grid refinement is placed at two latitudes (30N and 45N) and the simulation error norms are compared to the errors from global simulations with no mesh refinement. The two numerical methods perform markedly different in SWTC2 with mesh refinement. Mesh refinement with the spectral element method reduces the global error by 30% regardless of where the refined region is positioned, whereas mesh refinement with the finite volume method increases the global error by between 60% and 300% with the amount of increase sensitive to the location of the refined region. S08 contains no discussion with regards to how a refinement over an arbitrary 3% of the sphere can lead to a 30% reduction in global error in SWTC2. Our multi-resolution simulations of SWTC2 fall in between the results in S08. In terms of absolute accuracy, the global error norms that we present for SWTC2 are marginally lower than the errors produced by the finite-volume method in S08, but are nearly a factor of ten larger than the errors produced by the spectral element method in S08. When the flow is infinitely differentiable, as is SWTC2, spectral element methods are hard to match in terms of global error. The advantage that spectral element methods have on infinitely smooth flows is largely lost when discontinuities in the flow or forcing are present, such as in SWTC5. Since S08 only evaluate SWTC5 with static, quasi-uniform meshes and dynamically adapting meshes, it is not possible to make a close comparison to our results. We do note that our results are very much consistent with S08 when making a comparison of the global errors based on quasi-uniform meshes. With a uniform grid resolution of approximately 240 km we obtain a normalized global error of approximately 1.0e-3, whereas S08 show normalized errors of approximately 7.5e-4 and 2.0e-3 when using

the spectral element and finite-volume method, respectively.

The recent results of Chen et al (2011) (hereafter, C10) are also focused on dynamic adaptivity. Similar to S08, C10 evaluates SWTC2 with static mesh refinement and SWTC5 with dynamic mesh refinement. The numerical method used in C10 is a multi-moment method that utilizes both a cell average equation (similar to finite-volume methods) and a large number of point values (similar to spectral element methods). For SWTC2, C10 statically refines over a region that spans 22.5 degrees in longitude and 15 degrees in latitude. The grid spacing in the coarse and fine mesh zone is approximately 120 km and 15 km, respectively. Similar to S08, C10 places the refinement in arbitrary regions. C10 finds that using a refined mesh leads to an increase in error norms by between 5% and 35% as compared to the unrefined mesh. Qualitatively this result is consistent with our finding that the global error is controlled by the coarse mesh resolution. C10 evaluate SWTC5 with static, quasi-uniform meshes and dynamically adapting meshes, thus making a close comparison of the results difficult. We do note that when comparing errors based on the uniform meshes, our results are consistent with C10; on a mesh with a resolution of approximately 240 km we obtain a normalized global error of 1.0e-3 whereas C10 obtains a normalized global error of 5.0e-3.

The above comparison to W09, S08 and C10 focuses on each methods ability to minimize the global error in the shallow-water test case suite. In this comparison, the results obtained herein compare respectably to previously published results. At the same time, in our opinion the global error tells only a part of the story. The fact remains that long-term, robust solutions that are analogous to climate simulations are far more sensitive to conservation properties of the numerical scheme than to absolute accuracy. In terms of conservation of

mass and tracers, the finite-volume schemes presented in W09, S08 and C10 are all conservative. The spectral element scheme in S08 does not conserve mass or tracer substance. In terms of conserving potential vorticity, potential enstrophy or total energy, none of the schemes presented in W09, S08 or C10 have formal guarantees on conservation or boundedness. Furthermore, no anecdotal evidence comparable to Fig 5 and Fig 6 are presented in W09, S08 or C10 that would better illuminate each of the numerical methods character with respect to conservation. In this respect, the numerical scheme presented in T09, R10 and evaluated herein appears to be unique.

7. Discussion

Using a suite of shallow-water test cases, we evaluate the numerical scheme presented in T09 and R10 when implemented on variable resolution meshes. We produce a set of variable resolution meshes (see Figure 1 and Table 1) with grid-resolution spacing varying from quasi-uniform (X1) to highly-variable (X16). The simulations are conducted over a range of mesh sizes from 2562 to 655362 nodes.

The analysis included in T09 indicates that the numerical scheme evaluated herein supports geostrophic balance, even on variable resolution meshes. Since SWTC2 provides an initial condition in exact, nonlinear geostrophic balance, it provides an excellent means for evaluating the analysis in T09. We find that regardless of the mesh variation, geostrophic balance is maintained in the numerical simulations.

The analysis included in R10 indicates that the numerical scheme should maintain all its conservation properties on variable resolution meshes. We use SWTC5 with its large

transient forcing at $t = 0$ to measure conservation of mass, energy, potential vorticity and potential enstrophy. We find that both mass and potential vorticity are conserved to machine precision. Normalized total available energy is conserved to within 1.0×10^{-8} over the standard 15 day integration. We evaluate the spurious sources of energy stemming from the nonlinear Coriolis force and exchanges of energy between its kinetic and potential forms by measuring the time required for these spurious sources to double the globally-averaged kinetic energy. Consistent with the finding from R10 using quasi-uniform meshes, we find doubling times to be on the order of 10^4 years, regardless of the variation in mesh resolution.

The numerical scheme uses the anticipated potential vorticity method developed in Sadourny and Basdevant (1985) and explored further in R10. This numerical technique allows for the generation of physically-appropriate levels of potential enstrophy dissipation without dissipating kinetic energy. The simulations with SWTC5 show changes in globally-averaged potential enstrophy betweeen 10^{-4} and $10^{-2.5}$ for the X1 and X16 meshes, respectively. In some of those simulations (X1 and X2) the globally-averaged potential enstrophy decreased over time. In other simulations (X8 and X16) the globally-averaged potential enstrophy increased over time. We conducted all simulations with the same parameter setting $\theta = dt/2$ (see Sadourny and Basdevant (1985) Eq. 8). This parameter was chosen arbitrarily and, in retrospect, somewhat naively. We have confirmed that different choices for θ can lead to monotonically decreasing values of globally-averaged potential enstrophy in any of the simulations presented here. Instead of engaging in an *ad hoc* tuning exercise for θ, we plan to implement the scale-aware formulation of the anticipated potential vorticity method developed in Chen et al. (2011, accepted).

The rate of convergence for SWTC5 is approximately 1.5 *with respect to the coarse mesh*

resolution (see Figure 7). This rate of convergence is consistent across all meshes used in this study, regardless of the ratio between the minimum and maximum resolution. This rate of convergence is consistent with that found in T10 using quasi-uniform meshes. The rate of convergence for SWTC2 is not uniform. Meshes with minimum grid resolutions above 100 km show a convergence rate of approximately 1.9 with respect to the coarse mesh resolution (see Figure 8). Meshes with minimum grid resolutions less than 100 km show a continual reduction of convergence rate as the minimum grid resolution deceases. We have analyzed the mesh quality, the manner in which we compute the error norms and the numerical algorithm in an attempt to identify this shortcoming. We are uncomfortable with this reduction in convergence rate and will continue to seek its source.

We have carefully compared the results obtained herein to the works of Weller et al. (2009), St-Cyr et al (2008) and Chen et al (2011) (See Section 6). We find that the conservation properties demonstrated herein have not been demonstrated elsewhere. In this sense, the results produced in Section 5 are notable. In terms of the global normalized L_2 error norms obtained from SWTC2 and SWTC5, we find that our results are competitive in the sense that we obtain error norms that are both smaller and larger than those found in these other works.

We find that the mesh resolution in the coarse-mesh region is the primary factor controlling solution error. Figures 7 and 8 show that for SWTC5 and SWTC2, respectively, nearly all of the variation in the global L_2 error norm can be explained by the coarse-mesh resolution. This should not be surprising because in terms of reducing solution error, grid refinement is most advantageous when the solution in one part of the domain contains structures with relatively large derivatives and the solution in another part of the domain contains

structures with relatively small derivatives. Under this circumstance, it is plausible to reduce the solution error by a judicious rearrangement of a fixed number of grid cells. This situation is certainly not present in SWTC2 and, at least for this numerical scheme, is not sufficiently strong in SWTC5. As a result, the increase in solution error that accompanies the coarsening of the mesh in the coarse-mesh region exceeds any reduction in solution error that accompanies the refinement in the fine-mesh region. The larger error in the coarse-mesh region is propagated to all other regions, including the fine-mesh region, via advection and wave phenomena.

Fortunately, our motivation for exploring grid refinement is *not* a formal reduction in solution error. Rather, our motivation is to employ multi-resolution meshes so that certain phenomena like clouds or ocean eddies can be resolved in certain regions of interest. In this respect, Figures 7 and 8 are very promising. These figures indicate that we can specify the resolution in the coarse-mesh region(s) by determining what is an acceptable level of accuracy. From that starting point, we can increase resolution in region(s) of interest in order to simulate new phenomena while knowing that we will not degrade the formal accuracy of the solution. In practice we expect that the resolution of the coarse-mesh region(s) will be chosen to match typical IPCC-class resolutions and the fine-mesh region(s) will be chosen based on the phenomena to be simulated and the availability of computational resources. While we recognize that conclusions based on the idealized simulations discussed above must be regarded as tentative, we see no reason not to pursue this multi-resolution technique in more realistic systems.

We also evaluate the method using a standard barotropic instability test case. Similar to SWTC2, this test case specifies a zonally-symmetic zonal jet that is in exact non-linear

geostrophic balance. Different from SWTC2, this jet is barotropically-unstable. The test case specifies a small perturbation in the height field at $t = 0$ that triggers the instability. None of the meshes used in this study are zonally-symmetic. As a result, truncation error projects onto non-zero zonal wavenumbers and acts as an additional trigger for the barotropic instability. As shown in Figure 9, the impact of the truncation on the growth and position of the instability increases with mesh variation. For the suite of meshes with 655362 nodes, we find the X1, X2, X4, and X8 simulations to be qualitatively similar. The outlier is the X16 simulation that compares more closely to a X1 simulation with 40962 nodes.

We only examine one parameter in our three parameter density function shown in (4). The suite of meshes shown in Figure 1 is produced by varying γ, the parameter that controls the relative mesh spacing between the fine and coarse regions. Another critical parameter that needs to be examined carefully is α, the parameter that controls the width of the transition zone between the fine and coarse regions. As α gets smaller the width of the transition zone is reduced, the mesh transition becomes more abrupt and the local mesh distortion is increased. This, in turn, leads to an increase in truncation error and a reduction in the accuracy of the simulation. We expect that future studies will identify an "optimal" rate of mesh variation that balances the conflicting desires to minimize α and maintain local accuracy.

While we motivate this work based on the challenges encountered in global climate modeling, the application of this approach extends beyond the domain of climate simulation. For example, numerical weather prediction faces most of the same daunting challenges as global climate modeling, especially with regard to our inability to directly simulate all of the important spatial and temporal scales in the system. With the gap between atmosphere

climate models and numerical weather models closing, we expect that the multi-resolution approach developed here will find applications in both arenas.

Given the tentative progress demonstrated above, it is appropriate to consider the overarching challenges that will need to be overcome before a robust multi-resolution approach to climate system modeling is successful. In our view, the creation of a robust approach to multi-resolution climate system modeling requires success on two fronts; an accurate simulation of resolved scales of motion on an underlying mesh that varies in resolution and the creation of scale-aware parameterizations.

While we demonstrated some ability with respect to model stability and formal accuracy of simulations on variable resolution meshes, substantial challenges remain on several fronts. In particular, we have not yet addressed issues related to transport and wave propagation through mesh transition zones. With respect to the transport of tracer constituents, we expect that the recent high-order transport schemes (Skamarock and Menchaca (2010, accepted), Skamarock and Gassmann (2011), Ii and Xiao (2010)) along with a new analysis of flux-limiters (Mittal and Skamarock 2010, accepted) should be sufficient to maintain tracer field structure and amplitude through highly variable mesh transition zones.

Issues related to wave propagation are likely to be more difficult to address. One of the main motivations for this approach is to allow phenomena, including wave dynamics, to be better resolved in certain portions of the domain. By construction, a part of the wavenumber spectrum resolved in the fine-mesh region will not be resolved in the coarse-mesh region. As these high-wavenumber waves propagate out of the fine-mesh region, special care will be required to insure that these waves exit into the coarse mesh region in a sensible manner. Since we view this as the major outstanding challenge within the context of accurately

simulating resolved scales, our efforts will be directed to this problem immediately.

Developing scale-aware parameterizations for the atmosphere and ocean will be a much harder endeavor. While the venerable closures for clouds in the atmosphere (Arakawa and Schubert 1974) and eddies in the ocean (Gent and McWilliams 1990) have been remarkable in their success over the last decades, neither has been generalized across spatial and/or temporal scales (Randall et al. (2003), Gent (2010, in press)).

Both limited-area domain and stretched-grid simulations have had to address the lack of access to scale-aware parameterizations, i.e. parameterizations that function appropriately across a wide range of spatial and temporal scales without *ad hoc* tuning. Those conducting full physics simulations on stretched-grids are more acutely aware of this problem simply due to the fact that these deficiencies are manifest in a single, global simulation. One remedy pursued by the stretched-grid community has been to compute all physical parameterizations on an quasi-uniform mesh of intermediate resolution (Fox-Rabinovitz et al. 2006). While this remedy certainly removes biases in parameterizations due to their lack of scaling, the approach is antithetical to our motivation. Our motivation for this multi-resolution approach is founded on the principle that there is scientific value in directly resolving (i.e. not parameterizing) certain processes in certain regions. As a result, remedies found in the stretched grid community only highlight the extent of the challenges ahead of us.

In the short term, say over the next three to five years, we expect that careful choices in the positioning of the mesh transition zone(s) along with *ad hoc* scaling of closure parameters across mesh transition regions will allow the approach developed here to produce scientifically valuable results. In turn, we expect that this modeling approach can be used as a *testbed* for the evaluation of proposed parameterizations that are intended to be scale-aware. Over the

long term, the broad success of this modeling approach depends upon the development of a full suite of scale-aware parameterizations.

This modeling approach could potentially benefit all physical components included in global climate and weather prediction system models, including the atmosphere, ocean, land ice, sea ice and land surface components. Given the broad applicability of this approach, we have codified the technique through the creation of the Model for Prediction Across Scales (MPAS) project. The purpose of the MPAS project is to produce a suite of models based on a common conceptual and algorithmic foundation. The project has already produced this shallow-water model as well as prototype global atmosphere and ocean models based on the primitive equations. Since the numerical method evaluated above forms the core for both the primitive equation atmosphere and ocean models, this contribution serves as a scoping exercise for the identification of the successes and challenges in developing global primitive equation models based on a multi-resolution approach.

Acknowledgments.

This work was supported by the Biological and Environmental Research Division in the US Department of Energy Office of Science through DE-FG02-07ER64431, DE-FG02-07ER64432 and DOE 07SCPF152. The National Center for Atmospheric Research is sponsored by the National Science Foundation.

REFERENCES

Arakawa, A., 1966: Computational design for long-term numerical integration of the equations of fluid motion: Two-dimensional incompressible flow. *J. Comput. Phys.*, **1**, 119–143.

Arakawa, A. and W. Schubert, 1974: Interaction of a cumulus cloud ensemble with the large-scale environment, part i. *Journal of Atmospheric Sciences*, **31**, 674–701.

Barlow, G., 1974: Hexagonal territories. *Animal Behaviour*, **22**, 876–878.

Boccaletti, G., R. Ferrari, and B. Fox-Kemper, 2007: Mixed layer instabilities and restratification. *Journal of Physical Oceanography*, **37 (9)**, 2228–2250.

Bonaventura, L. and T. Ringler, 2005: Analysis of discrete shallow-water models on geodesic delaunay grids with c-type staggering. *Monthly Weather Review*, 2351–2373.

Campin, J., C. Hill, H. Jones, and J. Marshall, 2010, revised: Super-parameterization in ocean modelling: Application to deep convection.*Ocean Modelling*.

Chen, Q., M. Gunzburger, and T. Ringler, 2011, accepted: A scale-invariant formulation of the anticipated potential vorticity method. *Monthly Weather Review*, 1–42.

Chen, C., F. Xiao and X. Li: An Adaptive Multi-Moment Global Model on Cubed Sphere. *Monthly Weather Review*, to appear.

Davies, H., 1976: A lateral boundary formulation for multi-level prediction models. *Royal Meteorological Society*, **102**, 405–418.

Déqué, M., R.G. Jones, M. Wild, F. Giorgi, J.H. Christensen, D.C. Hassell, P.L. Vidale, B. Rockel, D. Jacob, E. Kjellström, 2005: Global high resolution versus Limited Area Model climate change projections over Europe: quantifying confidence level from PRUDENCE results. *Climate Dynamics*, **25**, 653–670.

Diffenbaugh, N., J. Pal, R. Trapp, and F. Giorgi, 2005: Fine-scale processes regulate the response of extreme events to global climate change. *Proceedings of the National Academy of Sciences of the United States of America*, **102 (44)**, 15 774–15 778.

Dirichlet, G., 1850: Uber die Reduktion der positiven quadratischen Formen mit drei unbestimmten ganzen Zahlen. *J. Reine Angew. Math.*, **40**, 209–227.

Du, Q., V. Faber, and M. Gunzburger, 1999: Centroidal voronoi tessellations: applications and algorithms. *SIAM review*, **41 (4)**, 637–676.

Du, Q., M. Gunzburger, and L. Ju, 2003: Voronoi-based finite volume methods, optimal voronoi meshes, and pdes on the sphere. *Computer Methods in Applied Mechanics and Engineering*, **192 (35-36)**, 3933–3957.

Fox-Rabinovitz, M., G. Stenchikov, M. Suarez, and L. Takacs, 1997: A finite-difference gcm dynamical core with a variable-resolution stretched grid. *Monthly Weather Review*, **125 (11)**, 2943–2968.

Fox-Rabinovitz, M., J. Côté, B. Dugas, M. Déqué and J.L. McGregor, 2006: Variable resolution general circulation models: Stretched-grid model intercomparison project (SGMIP). *J. Geophys. Res*, **111**, D16104.

Galewsky, J., R. Scott, and L. Polvani, 2004: An initial-value problem for testing numerical models of the global shallow-water equations. *Tellus A*, **56**, 429–440.

Gent, P., 2010, in press: The gent-mcwilliams parameterization: 20/20 hindsight. *Ocean Modelling*.

Gent, P. and J. McWilliams, 1990: Isopycnal mixing in ocean circulation models. *Journal of Physical Oceanography*, **20 (1)**, 150–155.

Gersho, A., 1979: Asymptotically optimal block quantization. *Ieee T Inform Theory*, **25**, 373–380.

Giorgi, F. and L. Mearns, 1991: Approaches to the simulation of regional climate change: a review. *Reviews of Geophysics*, **29 (2)**, 191–216.

Grabowski, W., 2010: Coupling cloud processes with the large-scale dynamics using the cloud-resolving convection parameterization (CRCP). *Journal of Atmospheric Sciences*, **58**, 978–997.

Ham, D., S. Kramer, G. Stelling, and J. Pietrzak, 2007: The symmetry and stability of unstructured mesh c-grid shallow water models under the influence of coriolis. *Ocean Modelling*, **16**, 47–60.

Harris, L. and D. Durran, 2010: An idealized comparison of one-way and two-way grid nesting. *Monthly Weather Review*, **138**, 2174–2187.

Heikes, R. and D. Randall, 1995: Numerical integration of the shallow-water equations on a

twisted icosahedral grid. part i: Basic design and results of tests. *Monthly Weather Review*, **123**, 1862–1880.

Ii, S. and F. Xiao, 2010: A global shallow water model using high order multi-moment constrained finite volume method and icosahedral grid. *Journal of Computational Physics*, **229 (5)**, 1774–1796.

Inatsu, M. and M. Kimoto, 2009: A Scale Interaction Study on East Asian Cyclogenesis Using a General Circulation Model Coupled with an Interactively Nested Regional Model. *Monthly Weather Review*, **137**, 2851–2868.

Jones, R.G., J.M. Murphy and M. Noguer: Simulation of climate change over europe using a nested regionalclimate model. I: Assessment of control climate, including sensitivity to location of lateral boundaries. *Q. J. R. Meteorol. Soc.*, **121**, 1413-1449.

Lorenz, P. and D. Jacob, 2005: Influence of regional scale information on the global circulation: A two-way nesting climate simulation. *Geophys. Res. Lett.*, **32**, L18706.

Ju, L., T. Ringler, and M. Gunzburger, 2010: Voronoi tessellations and their application to climate and global modeling. *Numerical Techniques for Global Atmospheric Models, Springer Lecture Notes in Computational Science and Engineering*, Eds. P. H. Lauritzen, C. Jablonowski, M. A. Taylor and R. D. Nair, to appear, 2011.

Khairoutdinov, M. and D. Randall, 2001: A cloud resolving model as a cloud parameterization in the ncar community climate system model: Preliminary results. *Geophys. Res. Lett*, **28 (18)**, 36173 620.

Klein, S. and D. Hartmann, 1993: The seasonal cycle of low stratiform clouds. *Journal of Climate*, **6 (8)**, 1587–1606.

Kleptsova, O., J. Pietrzak, and G. Stelling, 2009: On the accurate and stable reconstruction of tangential velocities in c-grid ocean models. *Ocean Modelling*, 118–126.

Marbaix, P., H Gallée, O Brasseur and J.P van Ypersele, 2003: Lateral boundary conditions in regional climate models: a detailed study of the relaxation procedure. *Monthly Weather Review*, **131**, 461–479.

McClean, J. L., et al., 2010, accepted: A prototype two-decade fully-coupled fine-resolution ccsm simulation. *Ocean Modeling*.

McGregor, J.L, 1997: Regional climate modelling. *Meteorology and Atmospheric Physics*, **63**, 105–117.

Mittal, R. and W. Skamarock, 2010, accepted: Monotonic limiters for a second order finite volume advection scheme using icosahedral-hexagonal meshes. *Monthly Weather Review*.

Newman, D., 1982: The hexagon theorem. *IEEE Trans. Inform. Theo.*, **28**, 137–139.

Ničković, S., M. Gavrilov, and I. Tošić, 2002: Geostrophic adjustment on hexagonal grids. *Monthly Weather Review*, **130 (3)**, 668–683.

Oliger, J. and A. Sundström, 1978: Theoretical and practical aspects of some initial boundary value problems in fluid dynamics, *SIAM Journal on Applied Mathematics*, **35**, 419–446.

Perot, B., 2000: Conservation properties of unstructured staggered mesh schemes. *Journal of Computational Physics*, **159 (1)**, 58–89.

Randall, D. and S. Bony, 2007: Climate models and their evaluation. *IPCC WG1 Fourth Assessment Report*.

Randall, D., M. Khairoutdinov, A. Arakawa, and W. Grabowski, 2003: Breaking the cloud parameterization deadlock. *Bulletin of the american Meteorological society*, **84 (11)**, 1547–1564.

Randall, D., T. Ringler, R. Heikes, P. Jones, and J. Baumgardner, 2002: Climate modeling with spherical geodesic grids. *Computing in Science & Engineering*, 32–41.

Ringler, T., R. Heikes, and D. Randall, 2000: Modeling the atmospheric general circulation using a spherical geodesic grid: A new class of dynamical cores. *Monthly Weather Review*, **128 (7)**, 2471–2490.

Ringler, T., L. Ju, and M. Gunzburger, 2008: A multiresolution method for climate system modeling: application of spherical centroidal voronoi tessellations. *Ocean Dynamics*, **58 (5)**, 475–498.

Ringler, T. and D. Randall, 2002: A potential enstrophy and energy conserving numerical scheme for solution of the shallow-water equations on a geodesic grid. *Monthly Weather Review*, **130 (5)**, 1397–1410.

Ringler, T., J. Thuburn, J. Klemp, and W. Skamarock, 2010: A unified approach to energy conservation and potential vorticity dynamics for arbitrarily-structured c-grids. *Journal of Computational Physics*, **229 (9)**, 3065–3090.

Sadourny, R., A. Arakawa, and Y. Mintz, 1968: Integration of the nondivergent barotropic

vorticity equation with an icosahedral-hexagonal grid for the sphere. *Monthly Weather Review*, **96 (6)**, 351–356.

Sadourny, R. and C. Basdevant, 1985: Parameterization of subgrid scale barotropic and baroclinic eddies in quasi-geostrophic models: Anticipated potential vorticity method. *Journal of the Atmospheric Sciences*, **42 (13)**, 1353–1363.

Saff, E. and A. Kuijlaars, 1997: Distributing many points on a sphere. *The Mathematical Intelligencer*, **19 (1)**, 5–11.

Skamarock, W. and M. Menchaca, 2010, accepted: Conservative transport schemes for spherical geodesic grids: High-order reconstructions for forward-in-time schemes. *Monthly Weather Review*.

Skamarock, W. and A. Gassmann, 2011: Conservative Transport Schemes for Spherical Geodesic Grids: High-Order Flux Operators for ODE-Based Time Integration. *Monthly Weather Review*, revised.

Staniforth, A., 1997: Regional modeling: A theoretical discussion. *Meteorology and Atmospheric Physics*, **63**, 15–29.

Stuhne, G. and W. Peltier, 2006: A robust unstructured grid discretization for 3-dimensional hydrostatic flows in spherical geometry: A new numerical structure for ocean general circulation modeling. *Journal of Computational Physics*, **213 (2)**, 704–729.

St-Cyr, A., C. Jablonowski, J.M. Dennis, H.M. Tufo and S.J. Thomas, 2008: A Comparison of Two Shallow Water Models with Non-Conforming Adaptive Grids: classical tests. *Monthly Weather Review*, **136**, 1898–1922.

Swarztrauber, P., 1996:Spectral transform methods for solving the shallow water equations on the sphere. *Mon. Wea. Rev*, **124**, 730–744.

Thuburn, J., 1997: A pv-based shallow-water model on a hexagonal–icosahedral grid. *Monthly Weather Review*, **125 (9)**, 2328–2347.

Thuburn, J., 2008: Some conservation issues for the dynamical cores of nwp and climate models. *Journal of Computational Physics*, **227 (7)**, 3715–3730.

Thuburn, J., T. Ringler, W. Skamarock, and J. Klemp, 2009: Numerical representation of geostrophic modes on arbitrarily structured c-grids. *Journal of Computational Physics*, **228 (22)**, 8321–8335.

Tomita, H., H. Miura, S. Iga, T. Nasuno, and M. Satoh, 2005: A global cloud-resolving simulation: Preliminary results from an aqua planet experiment. *Geophys. Res. Lett.*, **32 (8)**, L08 805.

Voronoi, G., 1908: Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième Mémoire: Recherches usr les parallélloèdres primitifs,. *J. Reine Angew. Math.*, **134**, 198–287.

Wang, Y., L.R. Leung, J.L. McGregor, D.K.Lee, W.C. Wang, Y. Ding and F. Kimura, 2010: Regional climate modeling: progress, challenges, and prospects. *Journal of the Meteorological Society of Japan*, **82**, 1599–1628.

Weller, H. and H. Weller, 2008: A high-order arbitrarily unstructured finite-volume model of the global atmosphere: Tests solving the shallow-water equations. *International Journal for Numerical Methods in Fluids*, **56 (8)**, 1589–1596.

Weller, H., H.G. Weller and A. Fournier, 2009: Voronoi, Delaunay, and Block-Structured Mesh Refinement for Solution of the Shallow-Water Equations on the Sphere. *Monthly Weather Review*, **137**, 4208–4224.

Williamson, J. Drake, J. Hack, R. Jakob, and P. Swarztrauber, 1992: A standard test set for numerical approximations to the shallow water equations in spherical geometry. *Journal of Computational Physics*, **102**, 211–224.

Williamson, D., 1968: Integration of the barotropic vorticity equation on a spherical geodesic grid. *Tellus*, **20**, 642–653.

List of Tables

TABLE 1. Approximate mesh resolutions (km) of the fine-mesh (dx_f) and coarse-mesh (dx_c) regions of the global domain for the X1 through X16 meshes as a function of the number of grid points.

Grid Points	$X1(dx_f,\ dx_c)$	$X2(dx_f, dx_c)$	$X4(dx_f, dx_c)$	$X8(dx_f, dx_c)$	$X16(dx_f, dx_c)$
2562	(480, 480)	(282, 537)	(196, 737)	(169, 1293)	(163, 2419)
10242	(240, 240)	(141, 169)	(98, 368)	(85, 648)	(81, 1222)
40962	(120, 120)	(70, 134)	(49, 184)	(42, 324)	(40, 611)
163842	(60, 60)	(35, 67)	(25, 92)	(21, 162)	(20, 305)
655362	(30, 30)	(16, 32)	(12, 48)	(10, 78)	(9, 148)

List of Figures

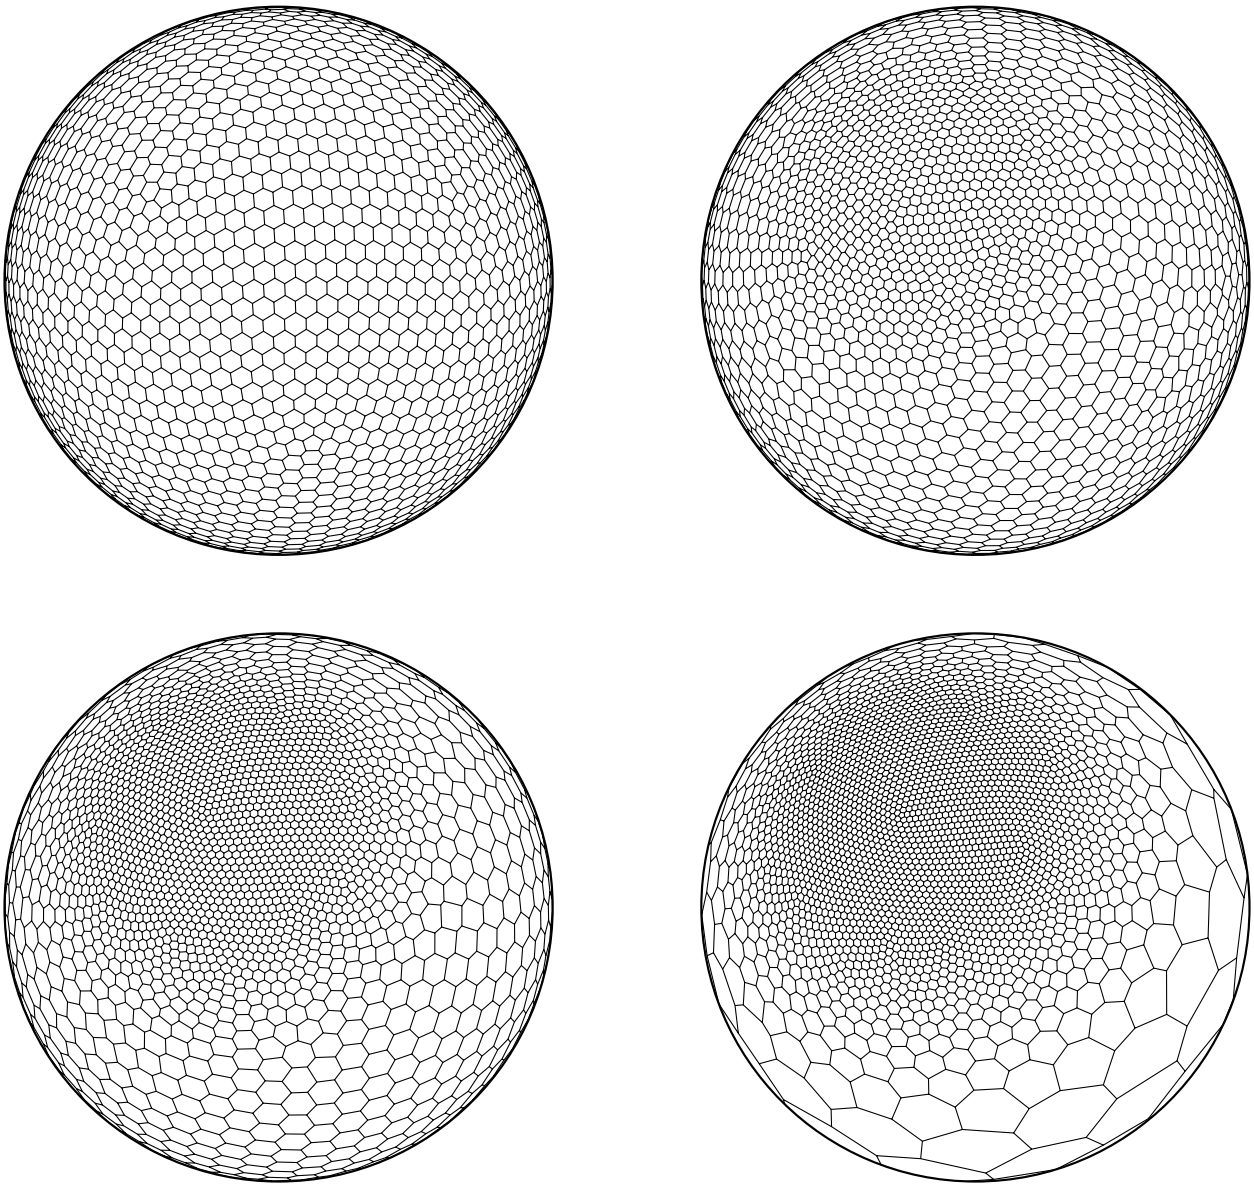

FIG. 1. Four members of a family of meshes constructed from Eq. (4). Each mesh uses 2562 grid points and only differ in the setting of the parameter γ to produce ratios in local grid resolution between the fine-mesh and coarse-mesh regions of 1, 2, 4 and 16 shown in the top-left, top-right, bottom-left and bottom-right, respectively.

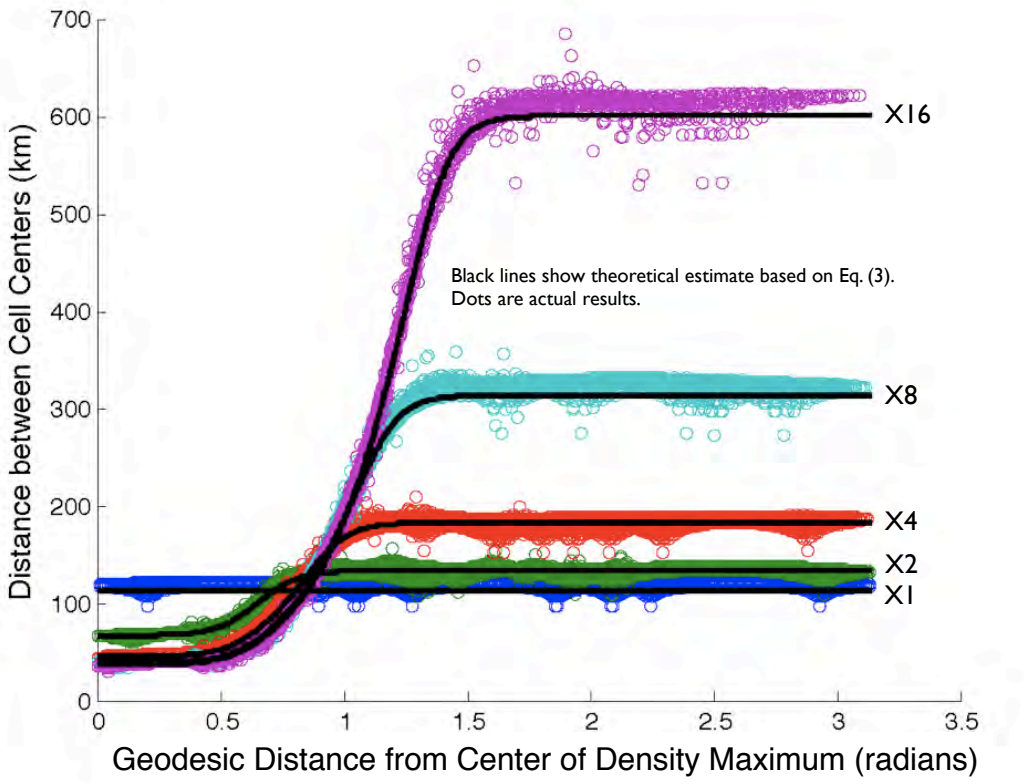

FIG. 2. Distribution of local mesh resolution as a function of geodesic distance from the center of the fine-mesh region. The x-axis measures the distance along the great circle arc between the center of the fine-mesh region, $\mathbf{x}_c$, and every grid point, $\mathbf{x}_i$. The y-axis measures the local mesh resolution in the vicinity of each $\mathbf{x}_i$ grid cell based on (6). Each open circle represents one cell on the X1, X2, X4, X8 or X16 meshs. Also shown for each mesh is the theoretical estimate of mesh resolution as a function of distance from $\mathbf{x}_c$ based on (4) with $\beta = \pi/6$, $\alpha = \pi/20$ and γ varies as described in (5).

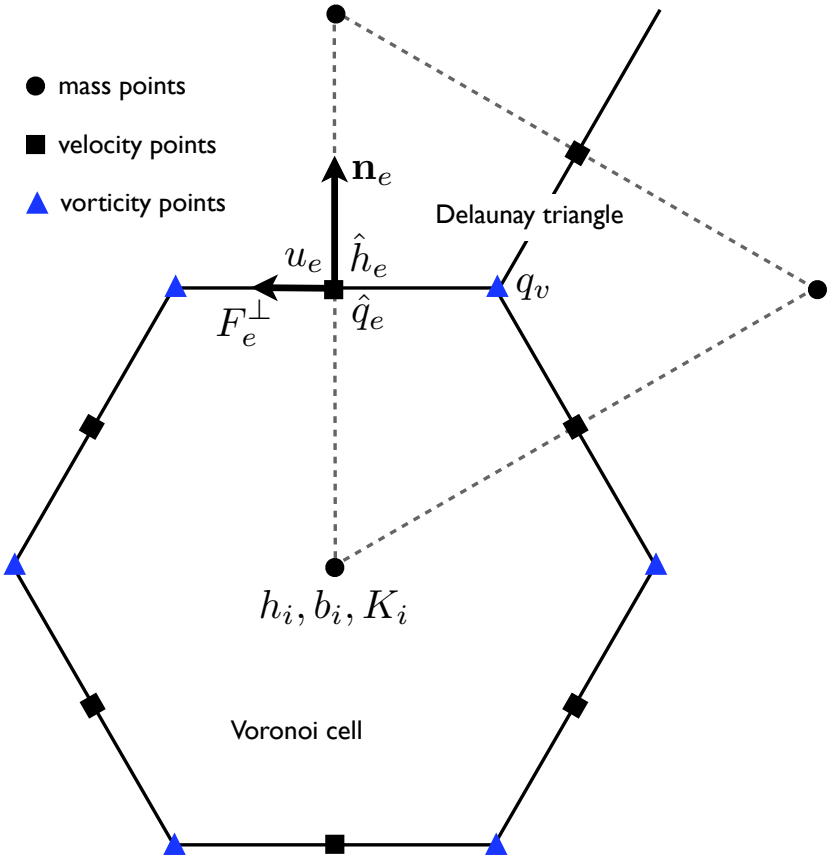

FIG. 3. Shown is the variable staggering for the finite-volume scheme. Mass, surface topography and kinetic energy are defined at the center of each Voronoi cell. The normal component of the velocity field, u_e, is defined at the mid-point of line segments connecting cell centers. All vorticity-related fields, such as the relative, absolute and potential vorticities are defined at the vertices of the Voronio cells. The derived fields, $\hat{h}_e$, $\hat{q}_e$ and $F_e^{\perp}$ must be reconstructed at each velocity point in order to evolve the velocity field forward in time.

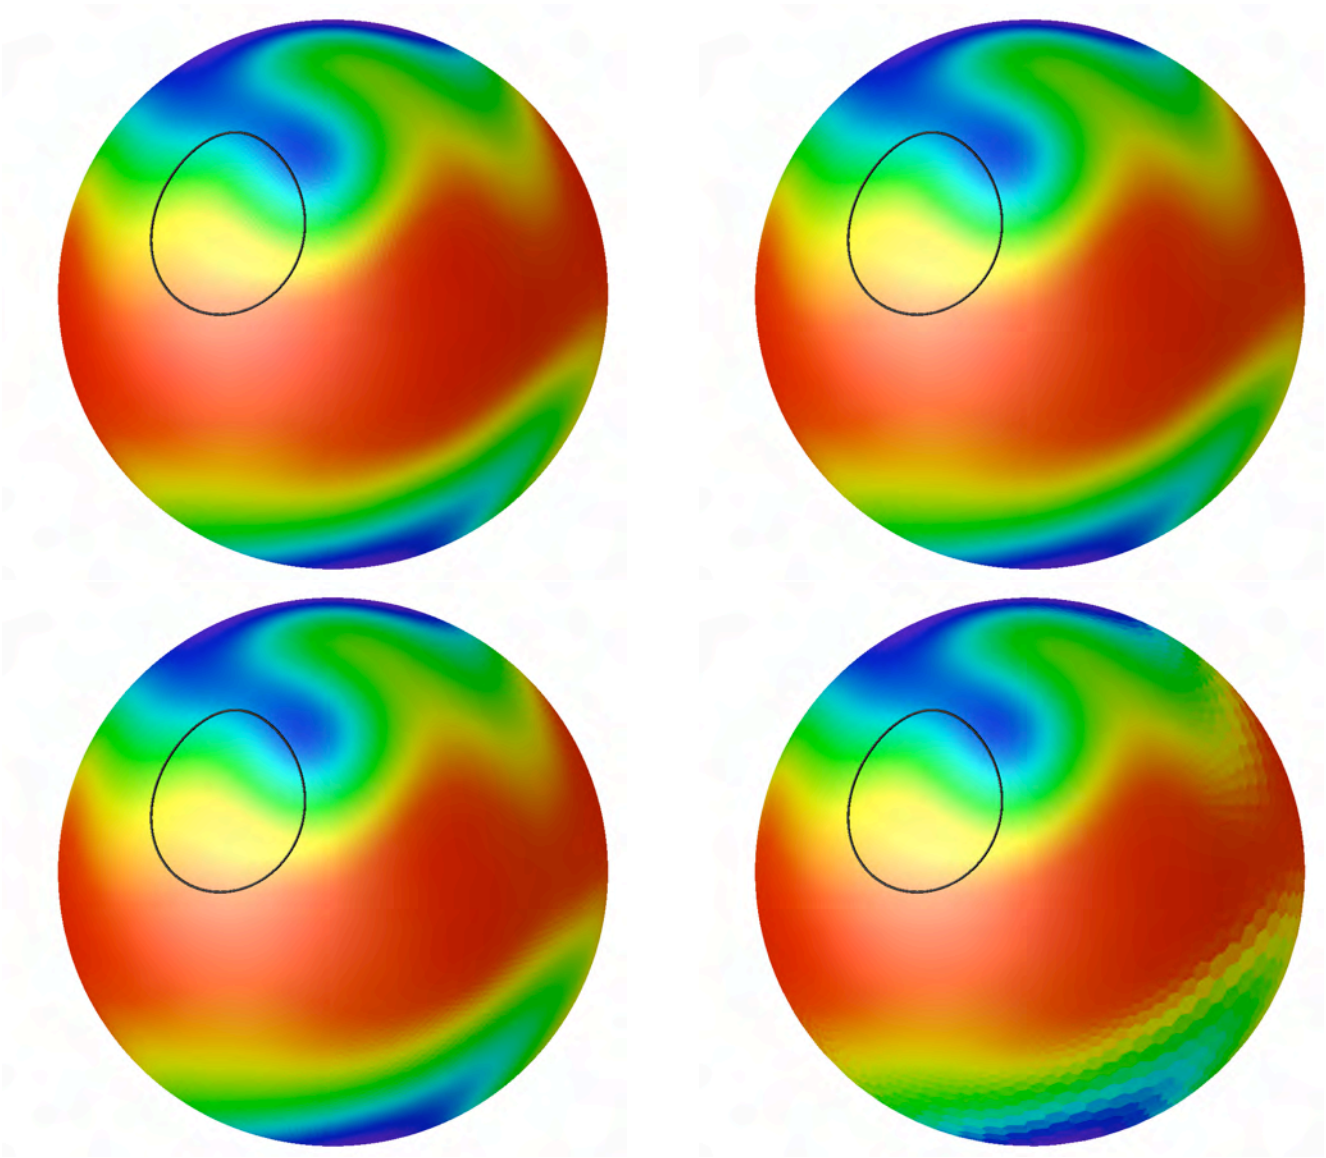

FIG. 4. The fluid height, $h_i + b_i$, at day 15 for SWTC5. Starting at the upper left and moving clockwise shows results from the X1, X2, X16 and X4 meshes using 40962 cells. The black oval denotes the location of the orography. The figures are generated by filling each Voronoi cells with a single color, i.e. there is no interpolation due to rendering. This allows the far-field grid resolution to be seen in the X4 and X16 simulations.

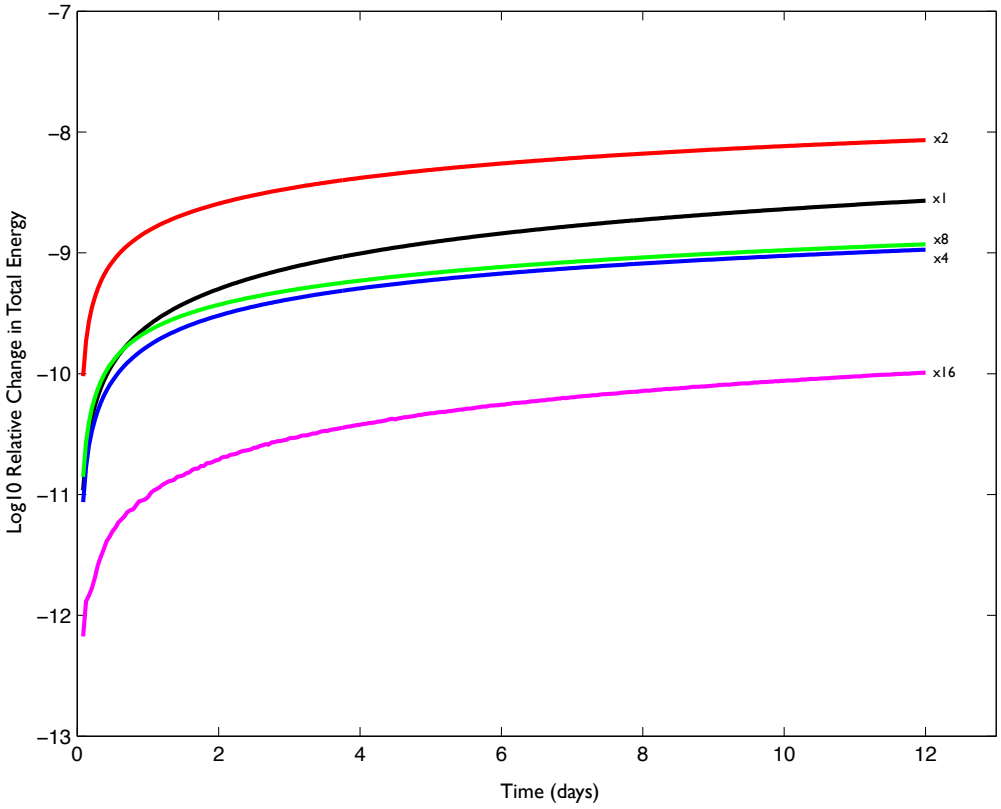

FIG. 5. Log_{10} of the relative change in available total energy for SWTC5 as a function of time for the X1, X2, X4, X8 and X16 meshes with 40962 grid points. All results are plotted with an identical color scheme with a maximum of 5975 m and a minimum of 5025 m.

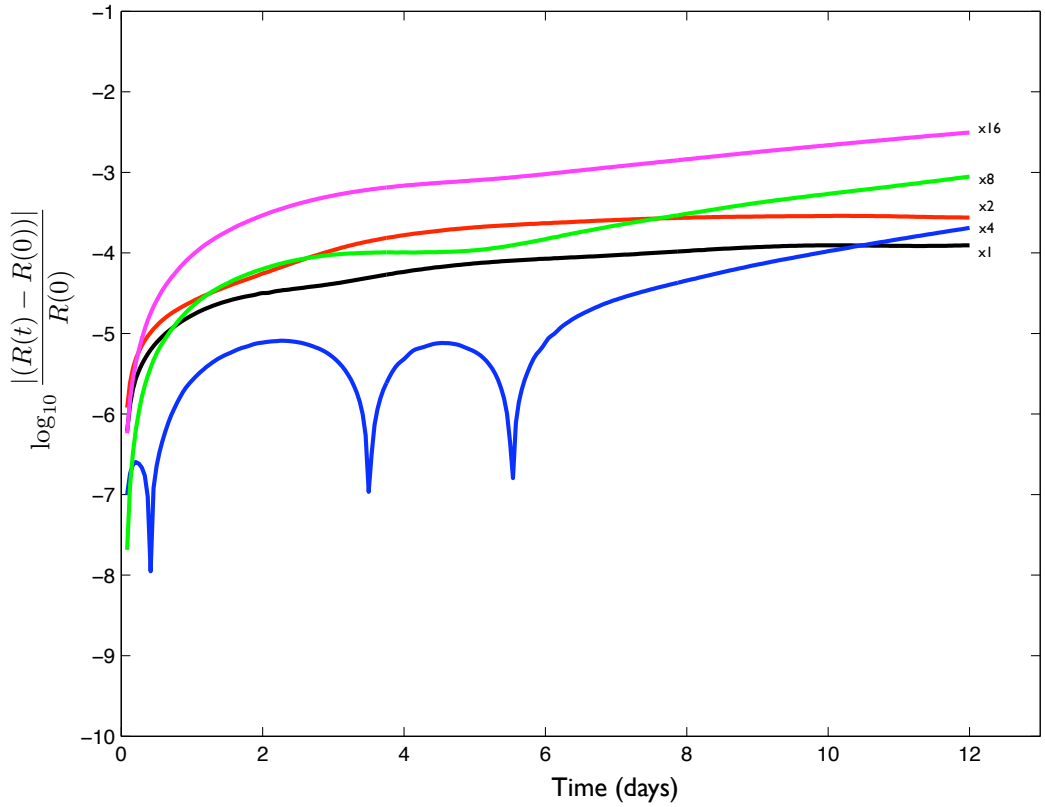

FIG. 6. Log_{10} of the relative change in available potential enstrophy for SWTC5 as a function of time for the X1, X2, X4, X8 and X16 meshes with 40962 grid points. In the X1 and X2 simulations the globally-averaged potential enstrophy is decreasing in time, while in the X8 and X16 simulations the the globally-averaged potential enstrophy is increasing in time. In the X4 simulation the globally-averaged potential enstrophy fluctuates about its initial value.

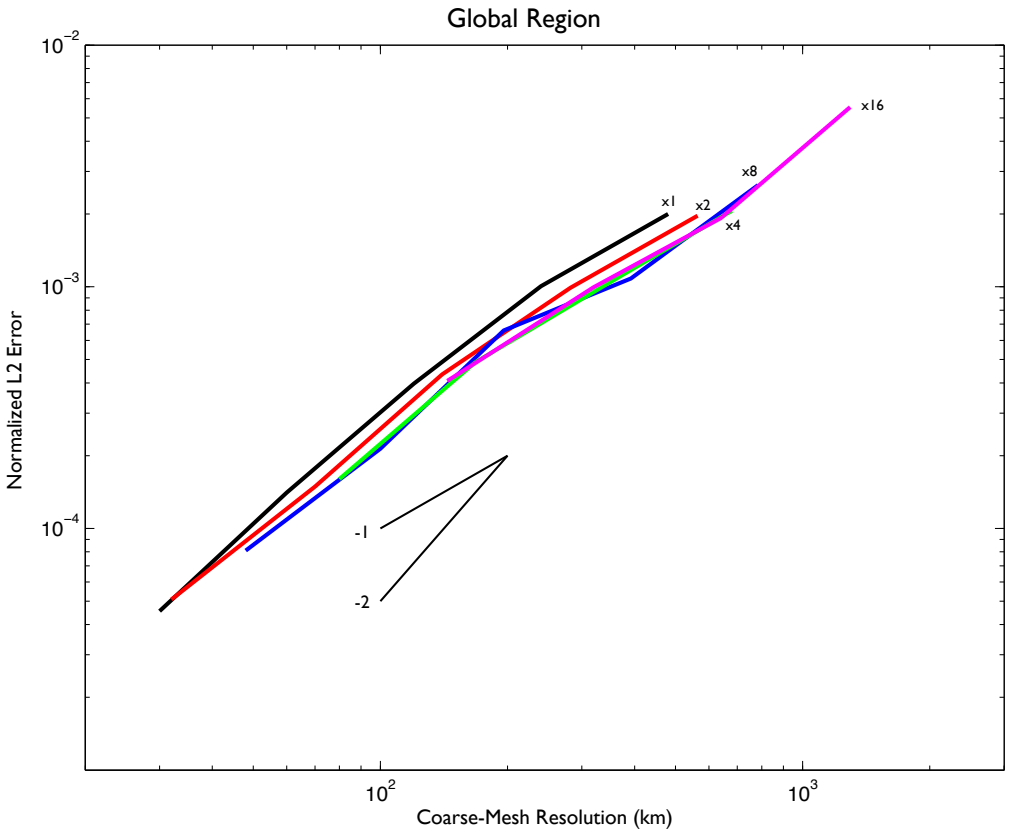

FIG. 7. The L_2 error of the thickness field at day 15 for SWTC5 shown for the X1, X2, X4, X8 and X16 meshes as a function of grid resolution found in the coarse-mesh region. Lines representing 1^{st}-order and 2^{nd}-order convergence rates are also shown.

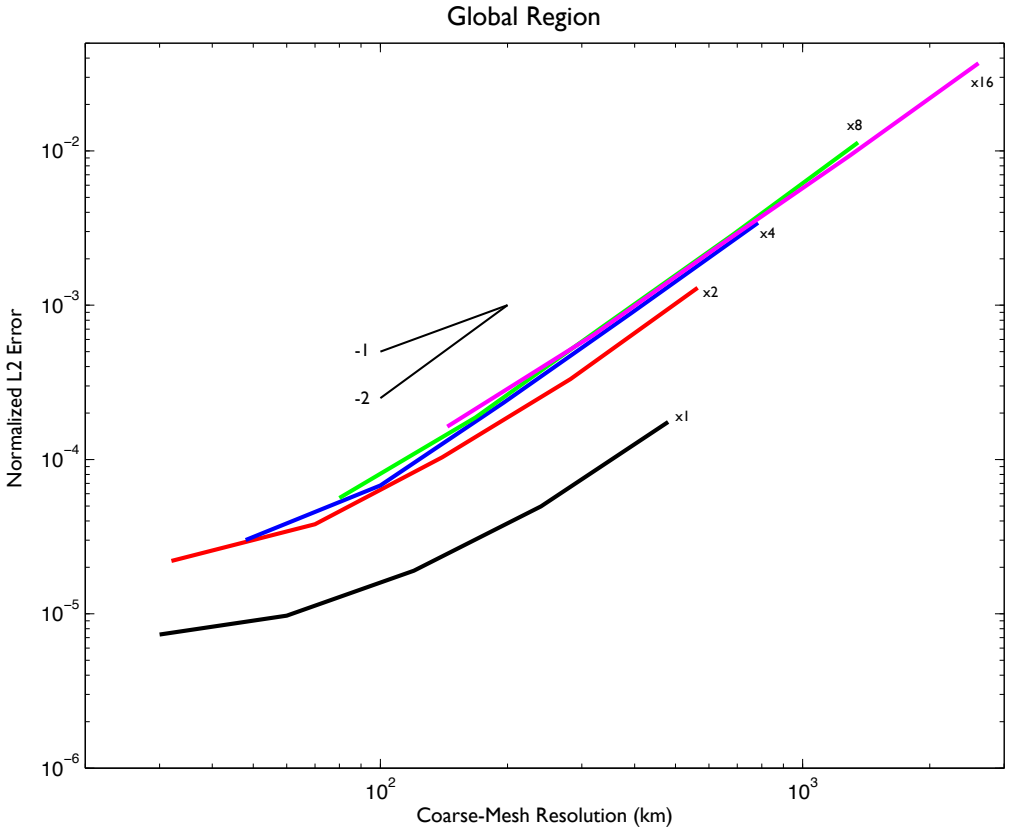

FIG. 8. The L_2 error of the thickness field at day 12 for SWTC2 shown for the X1, X2, X4, X8 and X16 meshes as a function of grid resolution found in the coarse-mesh region. Lines representing 1^{st}-order and 2^{nd}-order convergence rates are also shown.

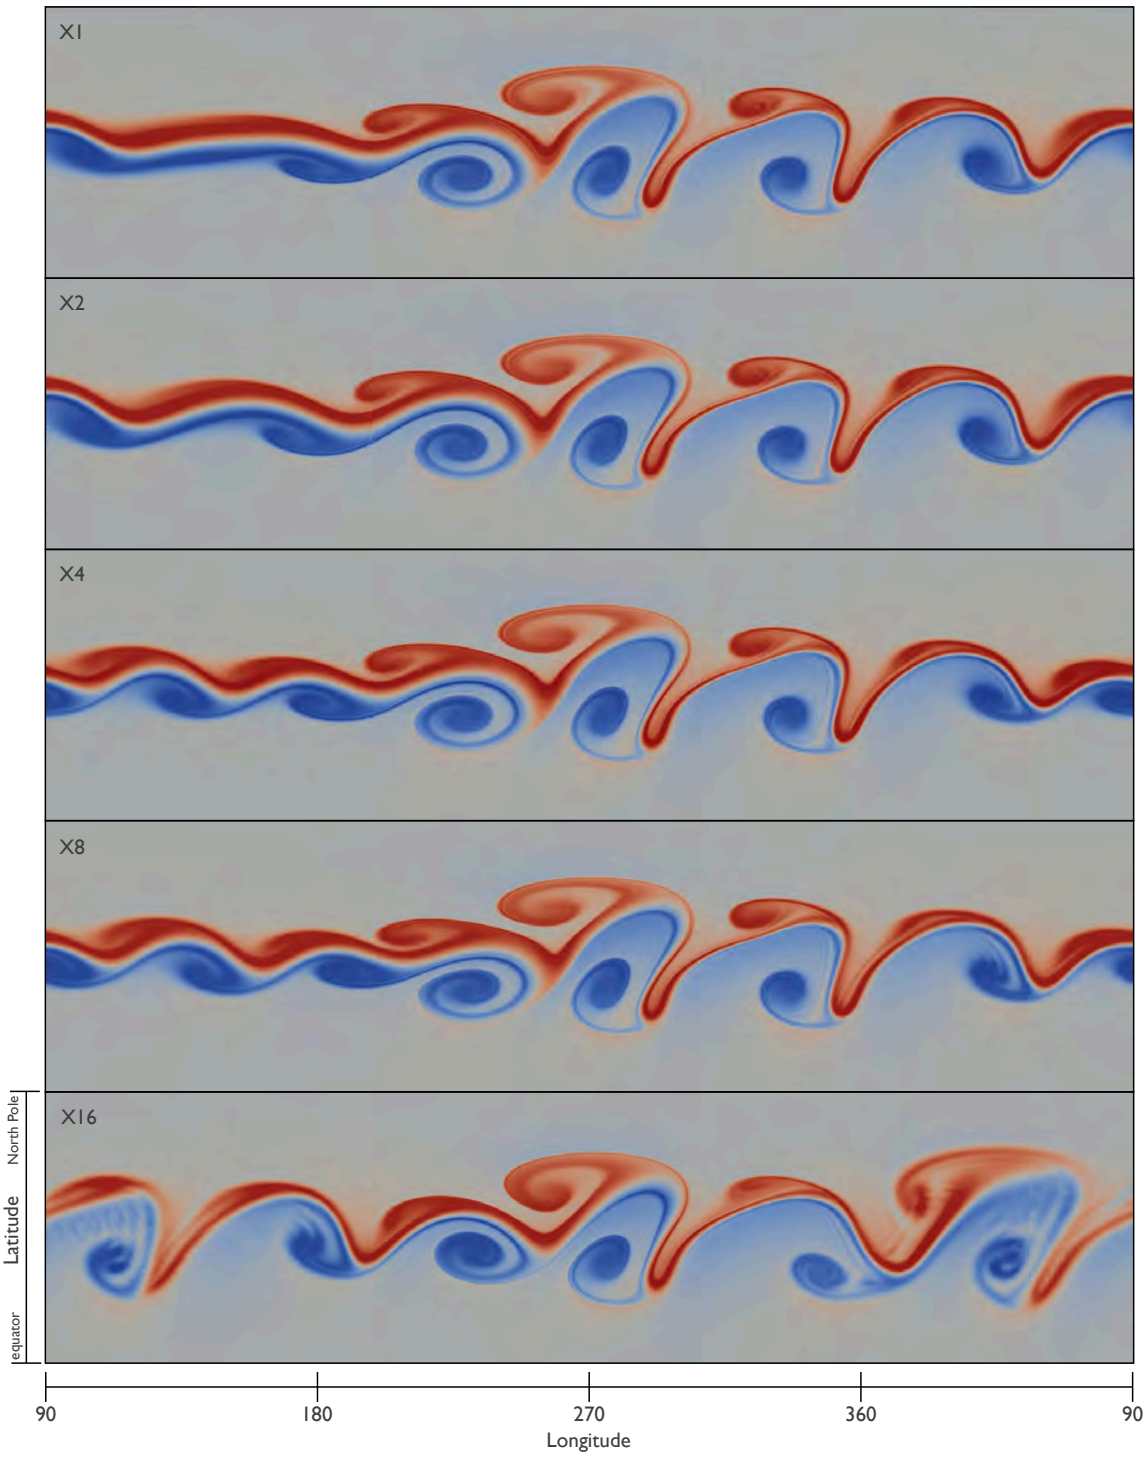

FIG. 9. Each panel depicts the relative vorticity field at day 6 for a barotropically-unstable jet using 655362 cells. The panels differ only in the mesh used in the simulation. The vertical extent of each panel covers the northern hemisphere. The horizontal extent covers all longitudes starting at -90 degrees such that the fine-mesh region is approximately centered on each panel. The color scales are identical for every panel and saturate at $\pm 1.0 \times 10^{-4}$.