Manuscript prepared for Geosci. Model Dev.
with version 2014/07/29 7.12 Copernicus papers of the LaTeX class copernicus.cls.
Date: 7 October 2014

The spectral element method on variable resolution grids: evaluating grid sensitivity and resolution-aware numerical viscosity

O. Guba[1], M. A. Taylor[2], P. A. Ullrich[3], J. R. Overfelt[2], and M. N. Levy[4]

[1]University of New Mexico, Albuquerque, New Mexico, USA
[2]Sandia National Laboratories, Albuquerque, New Mexico, USA
[3]University of California, Davis, California, USA
[4]National Center for Atmospheric Research, Boulder, Colorado, USA

Correspondence to: O. Guba (oksana@unm.edu)

Abstract. We evaluate the performance of the Community Atmosphere Model's (CAM) spectral element method on variable resolution grids using the shallow water equations in spherical geometry. We configure the method as it is used in CAM, with dissipation of grid scale variance implemented using hyperviscosity. Hyperviscosity is highly scale selective and grid independent, but does require a resolution dependent coefficient. For the spectral element method with variable resolution grids and highly distorted elements, we obtain the best results if we introduce a tensor-based hyperviscosity with tensor coefficients tied to the eigenvalues of the local element metric tensor. The tensor hyperviscosity is constructed so that for regions of uniform resolution it matches the traditional constant coefficient hyperviscsosity. With the tensor hyperviscosity the large scale solution is almost completely unaffected by the presence of grid refinement. This later point is important for climate applications where long term climatological averages can be *imprinted* by stationary inhomogeneities in the truncation error. We also evaluate the robustness of the approach with respect to grid quality by considering unstructured conforming quadrilateral grids generated with a well-known grid-generating toolkit and grids generated by SQuadGen, a new open source alternative which produces lower valence nodes.

1 Introduction

In climate and weather forecast applications there is an increased demand for variable-resolution capabilities. This demand is motivated by the need to resolve various temporal and spatial scales in forecast and regional climate studies with limited computational resources. Several approaches can be employed to this end including nesting techniques, multiphysics modeling, and multiresolution simulations, recently overviewed in Ringler et al. (2011).

Here we focus on the multiresoluton approach made possible by the spectral element method (SEM). We use global spherical grids, constructed to have uniform high resolution over a region of interest, uniform low resolution over the rest of the globe, and a transition region between them. The SEM has a long history of running on highly unstructured grids, including spherical geometry with climate applications (Fournier et al., 2004; St.-Cyr et al., 2008; Baer et al., 2006; Marras et al., 2014). Here our goal is to evaluate and improve the multiresolution capabilities of the SEM formulation from the High-Order Method Modeling Environment (HOMME) recently adopted as the default dynamical core by the Community Atmosphere Model (CAM) (Dennis et al., 2012). HOMME uses a locally conservative/mimetic formulation from Taylor and Fournier (2010) and relies on a constant coefficient hyperviscosity term to both dissipate energy near the grid scale and to damp grid scale modes with spurious propagation (Ainsworth and Wajid, 2009). This hyperviscosity operator is not suitable for variable resolution grids, and thus we consider two resolution-aware extensions. The first is the straightforward extension of allowing the hyperviscsoity coefficient ν to depend on the local element length scale. This approach was used in Zarzycki et al. (2014a, b) for realistic CAM simulations. Here, we use the shallow water equations to show some deficiencies with this approach, and then show that better results are obtained with a tensor-based hyperviscosity operator which can better represent both length scales within non-square spectral elements. We evaluate this approach using the shallow water equations on the sphere with the two-dimensional version of HOMME's spectral ele-

ment dynamical core. There have been other modifications of the viscosity and hyperviscosity operators. For example, Yu et al. (2014) uses a flow-dependent coefficient for the Laplacian. In Dobrev et al. (2012) several tensor coefficient viscosity operators were considered including a formulation where directional viscosity coefficients were chosen to depend on directional length scales of a deformed Lagrangian volume. Here we follow a similar approach, only the length scales come from the deformation of the elements, especially those elements in the grid transition region.

The mimetic SEM requires conforming quadrilateral grids. To generate variable-resolution grids, we employ a well-known grid-generating toolkit, CUBIT (https://cubit.sandia.gov). In the grid-transition region, the CUBIT grids have many valence 6 nodes (corner nodes shared by 6 elements). For grids in spherical geometry, it is possible to construct transition regions with mostly valence 5 nodes, and such elements will have less acute angles. To generate grids with low valence nodes, we use the Spherical Quadrilateral Grid Generator (SQuadGen, http://climate.ucdavis.edu/squadgen.php). This toolkit uses a paving technique (Blacker and Stephenson, 1991) in combination with a set of low-valence tiles to generate smooth quadrilateral grids based on cubed-sphere geometry. Regions of enhancement are determined via a user-specified image file which is mapped onto a cubed-sphere grid. Grid smoothing is performed via straightforward application of spring dynamics in 3-D geometry (Persson and Strang, 2004). Grids obtained via this technique exhibit several improved characteristics, including greater uniformity in the transition region, and elements with angles that are closer to 90 °.

In this study we use multiresolution grids with a single region of quasi-uniform high resolution, Δx_{high}, which transitions to a quasi-uniform grid of low resolution, Δx_{low}, covering most of the globe.

While evaluating the model's performance, it is natural to compare the multiresolution simulation with the corresponding Δx_{low} and Δx_{high} uniform grid simulations. Motivated by climate applications, and following Weller et al. (2009); Ringler et al. (2011) we evaluate $\Delta x_{\text{low}}/\Delta x_{\text{high}}$ variable-resolution simulations with two criteria:

1. Refinement does no harm to the global scales. For the shallow water equation initial value problem, at short times, it is reasonable to expect many features in the refined region would not be contaminated by information from the low resolution region. In contrast, for longer times we expect all features to be influenced by information from both the low and high resolution regions. So, in general, global errors should depend on Δx_{low} and may not be improved by the presence of a Δx_{high} resolution region. However, they should not be worsened by the presence of refinement and we thus expect global errors in a multiresolution simulation to be as low as those obtained by a uniform Δx_{low} simulation.
2. Local scales are resolved in the refined region. The purpose of the multiresolution approach in climate modeling is not to reduce the initial value problem error, but to resolve features of interest such as hurricanes, eddies, or topographically driven features in select regions at a lower cost. We thus expect that in the refined region, the multiresolution simulation can resolve the same scales as the uniform Δx_{high} resolution simulation. With hyperviscosity, this requires a resolution aware formulation which locally matches what would be used in a uniform resolution simulation.

In the study we use the popular set of shallow water test cases on the sphere compiled by Williamson et al. (1992) to show that the SEM satisfies both requirements when tensor hyperviscosity is used. We show that tensor hyperviscosity is both more accurate and more robust with respect to grid quality. The rest of the paper is organized as follows: In Sect. 2 we introduce two dissipation mechanisms, scalar and tensor hyperviscosity, in Sect. 3 we discuss grid refinement techniques, in Sect. 4 we describe shallow water test cases, and in Sect. 5 we present numerical results.

2 Hyperviscosity formulations

In many climate models a hyperviscosity term is added to the right hand side of the dynamical equations for both physical and numerical reasons. Hyperviscosity is preferred over regular viscosity because it is more scale selective. It strongly damps grid scale modes while having less of an impact on the resolved large scale modes. To introduce the various types of hyperviscosity that will be considered here, we first consider a model equation containing only the hyperviscosity operator

$$Q_t = -\nu\Delta^2 Q, \quad \nu > 0 \tag{1}$$

acting on scalar fields with a constant-coefficient ν. Here subscript t denotes differentiation w.r.t. time. The tensor formulation replaces $\nu\Delta^2$ by $(\nabla\cdot\boldsymbol{\tau}\nabla)\Delta$ for a symmetric positive definite matrix $\boldsymbol{\tau}$. That is, instead of (1) we consider

$$Q_t = -(\nabla\cdot\boldsymbol{\tau}\nabla)\Delta Q. \tag{2}$$

Our intention is to derive a formulation for $\boldsymbol{\tau}$ that is suitable for uniform and quasi-uniform grids and can be extended to non-uniform grids.

We start by writing equation (2) in weak form. Since (2) is equivalent to the set of equations $Q_t = -(\nabla\cdot\boldsymbol{\tau}\nabla)q$, $q = \Delta Q$, we rewrite the set as a system of integral equations,

$$\int_\Omega \phi_1 Q_t = \int_\Omega \nabla\phi_1\cdot\boldsymbol{\tau}\nabla q, \tag{3}$$

$$\int_\Omega \phi_2 q = -\int_\Omega \nabla\phi_2\cdot\nabla Q. \tag{4}$$

Here Ω is the problem domain, which in our case will be the sphere of radius R. This system of equations is discretized by the standard SEM and solved for all SEM test functions ϕ_1 and ϕ_2. We first decompose the domain Ω into a set of quadrilateral elements on the surface of the sphere, Ω_m, $m = 1, \ldots, M$, such as in Fig. 1, and then write

$$\sum_m \int_{\Omega_m} \phi_1 Q_t = \sum_m \int_{\Omega_m} \nabla\phi_1 \cdot \boldsymbol{\tau} \nabla q, \tag{5}$$

$$\sum_m \int_{\Omega_m} \phi_2 q = -\sum_m \int_{\Omega_m} \nabla\phi_2 \cdot \nabla Q. \tag{6}$$

Each term in this sum is then written as an integral over the reference element $\Omega_{\text{ref}} = [-1,1]\times[-1,1]$. We define $\boldsymbol{r}(\boldsymbol{x}; m)$ as the map from Ω_{ref} to Ω_m, with $\boldsymbol{r} \in \Omega_m$ a point on the sphere and $\boldsymbol{x} = (x_1, x_2) \in \Omega_{\text{ref}}$. We require this map to be differentiable and invertible, and further define

$$\mathbf{D} = \partial \boldsymbol{r} / \partial \boldsymbol{x} \tag{7}$$

where $\mathbf{D}$ is a two by two matrix whose columns are the covariant basis vectors expressed in spherical coordinates. The map and analytic expressions for $\mathbf{D}$ are given in the appendix. The integral over each spherical element Ω_m can then be written with respect to Ω_{ref}, using derivatives with respect to the reference element coordinates,

$$\int_{\Omega_m} \nabla\phi \cdot \boldsymbol{\tau} \nabla q = \int_{\Omega_{\text{ref}}} \begin{pmatrix} \frac{\partial\phi}{\partial x_1} \\ \frac{\partial\phi}{\partial x_2} \end{pmatrix}^{\mathrm{T}} \mathbf{D}^{-1} \boldsymbol{\tau} \mathbf{D}^{-T} \begin{pmatrix} \frac{\partial q}{\partial x_1} \\ \frac{\partial q}{\partial x_2} \end{pmatrix} \det(\mathbf{D})\, \mathrm{d}x_1 \mathrm{d}x_2. \tag{8}$$

where $\det(\mathbf{D})\,\mathrm{d}x_1\mathrm{d}x_2$ is the transformed area measure, $\boldsymbol{\tau}$ is the tensor expressed in spherical coordinates.

Note that the discrete operator will conserve mass, $\int Q$, since the change in mass is obtained by taking test function $\phi_1 = 1$ and then the right hand side of Eq. (5) will be zero if we ensure $\nabla 1 = 0$.

We now consider the eigenvalues of the metric tensor and its inverse,

$$\mathbf{D}^{\mathrm{T}}\mathbf{D} = \mathbf{E} \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \mathbf{E}^{\mathrm{T}},\ \mathbf{D}^{-1}\mathbf{D}^{-T} = \mathbf{E} \begin{pmatrix} \lambda_1^{-1} & 0 \\ 0 & \lambda_2^{-1} \end{pmatrix} \mathbf{E}^{\mathrm{T}}, \tag{9}$$

with orthonormal matrix $\mathbf{E}$ whose columns are the basis vectors which diagonalize the Laplace operator. For any practical grid, both $\mathbf{D}$ and $\mathbf{D}^{-1}$ are well defined and hence the symmetric metric tensor is guaranteed to have such an eigenvalue decomposition with positive eigenvalues. We note that these two eigenvalues can be used to define the two length scales associated with each element Ω_m. For the special case of a grid of rectangular elements in Cartesian geometry of size $l_x \times l_y$, we have that $\mathbf{E}$ is the identity matrix and

$$\lambda_1 = (l_x/2)^2 \qquad \lambda_2 = (l_y/2)^2. \tag{10}$$

For general possibly distorted element, we define its two length scales by $2\sqrt{\lambda_1}$ and $2\sqrt{\lambda_2}$.

2.1 Constant-coefficient hyperviscosity

The traditional constant coefficient hyperviscosity is obtained by taking $\boldsymbol{\tau} = \nu\mathbf{I}$, with the identify matrix $\mathbf{I}$. For uniform resolutions with an average grid spacing of Δx, often $\nu = c_0(\Delta x)^{3.2}$, for some constant $c_0 > 0$. This scaling is obtained by experimentation and is found to be effective for several different dynamical cores over a wide range of resolutions (Boville, 1991; Takahashi et al., 2006; Dennis et al., 2012). We take a slightly more general form and allow $\nu = c_0(\Delta x)^s$ for a scaling parameter s. The constant coefficient hyperviscosity is used for quasi-uniform grids, where we follow the convention of defining Δx by the average number of degrees of freedom on the equator. For square elements, $l_x = p\Delta x$ where p is the polynomial order of the basis functions in the SEM.

In order to motivate how we generalize this operator to a full tensor, we first express $\boldsymbol{\tau} = \nu\mathbf{I}$ in the basis which diagonalizes the Laplace operator (the local element basis defined by $\mathbf{E}$). Some algebra shows that

$$\mathbf{E}^{-1}\mathbf{D}^{-1}\boldsymbol{\tau}\mathbf{D}^{-T}\mathbf{E}^{-T} = \begin{pmatrix} \nu\lambda_1^{-1} & 0 \\ 0 & \nu\lambda_2^{-1} \end{pmatrix}.$$

Below, we will ensure that the more general tensor formulations retain this scaling in regions where the grid has uniform resolution $\lambda_1 \simeq \lambda_2$.

2.2 Scalar hyperviscosity

For scalar hyperviscosity, we again take $\boldsymbol{\tau} = \nu\mathbf{I}$ and now allow ν to vary in space. The natural choice is to use the same scaling as with the constant-coefficient operator, $\nu = c_0(\Delta x)^s$, but with Δx now chosen locally for each element. To preserve the scaling for the constant-coefficient operator, but to also ensure that the coefficient does not become too small (and thus provide insufficient dissipation), we use Eq. (10) and approximate the resolution locally by taking $l_x = 2(\max\{\lambda_1, \lambda_2\})^{1/2}$, $\Delta x = l_x/p$. For scalar hyperviscosity, this tensor scales according to

$$\mathbf{E}^{-1}\mathbf{D}^{-1}\boldsymbol{\tau}\mathbf{D}^{-T}\mathbf{E}^{-T} = \begin{pmatrix} \nu\lambda_1^{-1} & 0 \\ 0 & \nu\lambda_2^{-1} \end{pmatrix}, \quad \nu = c_0(\Delta x)^s.$$

On a Cartesian grid with square elements, $\lambda_1 = \lambda_2 = (l_x/2)^2$ and so the scalar and constant-coefficient operators are identical. For a quasi-uniform grid, where $\lambda_1 \simeq \lambda_2 \simeq (l_x/2)^2$, the scalar and constant-coefficient operators will have the same scaling with resolution.

2.3 Tensor hyperviscosity

Now consider a grid with only rectangles of size $l_x \gg l_y$. Based on our expected scaling hyperviscosity with resolution, we note that scalar hyperviscosity above would give us the desired amount of dissipation in the x direction, but will have excessive dissipation acting in the y direction. The natural choice for a grid of rectangles is a tensor coefficient,

$$\boldsymbol{\tau} = \begin{pmatrix} \nu_1 & 0 \\ 0 & \nu_2 \end{pmatrix} \qquad \nu_1 = c_0(\Delta x)^s, \qquad \nu_2 = c_0(\Delta y)^s. \tag{11}$$

For a grid of pure rectangles, $\mathbf{D}$ is diagonal and $\mathbf{E} = \mathbf{I}$, so $\boldsymbol{\tau}$ expressed in the $\mathbf{E}$ basis is given by

$$\mathbf{E}^{-1}\mathbf{D}^{-1}\boldsymbol{\tau}\mathbf{D}^{-T}\mathbf{E}^{-T} = \begin{pmatrix} \nu_1\lambda_1^{-1} & 0 \\ 0 & \nu_2\lambda_2^{-1} \end{pmatrix}. \tag{12}$$

We use this same formulation for unstructured grids by defining the two locally varying element length scales as in Eq. (10) and taking $\nu_1 = c_0(2\sqrt{\lambda_1}/p)^s$ and $\nu_2 = c_0(2\sqrt{\lambda_2}/p)^s$. On a Cartesian grid of pure rectangles, the scalar and tensor operators are identical. For a quasi-uniform grid of rectangles where $\lambda_1 \simeq (l_x/2)^2$ and $\lambda_2 \simeq (l_y/2)^2$, the scalar and tensor operators will have the same scaling with resolution.

For direct comparison, we summarize the three different approaches:

- Constant-coefficient: For quasi-uniform grids with average grid spacing Δx

$$\boldsymbol{\tau} = \nu\mathbf{I} = \mathbf{DE}\begin{pmatrix} \nu\lambda_1^{-1} & 0 \\ 0 & \nu\lambda_2^{-1} \end{pmatrix}(\mathbf{DE})^{\mathrm{T}}$$
$$\nu = c_0(\Delta x)^s$$

- Scalar: $\nu = \nu(\boldsymbol{r})$ depends on local element length scales

$$\boldsymbol{\tau} = \nu\mathbf{I} = \mathbf{DE}\begin{pmatrix} \nu\lambda_1^{-1} & 0 \\ 0 & \nu\lambda_2^{-1} \end{pmatrix}(\mathbf{DE})^{\mathrm{T}}$$
$$\nu = c_0(\Delta x)^s$$
$$\Delta x = 2\sqrt{\max\{\lambda_1, \lambda_2\}}/p$$

- Tensor: $\boldsymbol{\tau}$ depends on local element length scales

$$\boldsymbol{\tau} = \mathbf{DE}\begin{pmatrix} \nu_1\lambda_1^{-1} & 0 \\ 0 & \nu_2\lambda_2^{-1} \end{pmatrix}(\mathbf{DE})^{\mathrm{T}}$$
$$\nu_1 = c_0(\Delta x)^s,$$
$$\Delta x = 2\sqrt{\lambda_1}/p$$
$$\nu_2 = c_0(\Delta y)^s,$$
$$\Delta y = 2\sqrt{\lambda_2}/p$$

For a smoothly deformed grid, the matrix entries of $\boldsymbol{\tau}$ will be smooth functions over the domain Ω_m. For our discrete grids, we ensure $\boldsymbol{\tau}$ is continuous across element edges by applying the standard SEM projecting operation to each entry of $\boldsymbol{\tau}$. We further reduce variations in $\boldsymbol{\tau}$ by computing it only at element corner points, forming a bilinear fit to these corner values, and using this bilinear approximation at all nodes within the element.

2.4 Hyperviscosity acting on vector fields

In our method, we represent vector fields $\boldsymbol{u}$ in spherical coordinates, but care must be taken not to differentiate individual vector components when represented in spherical coordinates since they are multiply valued at the poles. Instead we transform $\boldsymbol{u}$ from spherical to Cartesian coordinates and solve Eqs. (5–6) for each Cartesian component of velocity field, and then transform the result back to spherical coordinates.

3 High and low connectivity conforming quadrilateral grids on the sphere

The mimetic formulation of the SEM we are using requires conforming quadrilateral grids. The cubed-sphere is a popular way to construct these grids on the sphere with quasi-uniform resolution. An inscribed cube is projected onto the surface of the sphere and each panel is further subdivided into a grid of elements, as shown in Fig. 1.

For multiresolution, we consider grids with a single refined region over an area of interest. We define a coarse resolution Δx_{low} and fine resolution Δx_{high}. We restrict ourselves to choices $\Delta x_{\text{high}} = \Delta x_{\text{low}}/N$, $N = 2, 4$ and 8. Starting from a cubed-sphere grid with resolution Δx_{low}, the region under refinement is substituted by uniform elements with Δx_{high}, as shown on Fig. 2. The approximate placement of the transition region is colored grey. For each N, we generate a family of grids with different low resolution regions Δx_{low}. Following Ringler et al. (2011), we refer to these family of grids as $\times 1$, $\times 2$, $\times 4$ and $\times 8$. The $\times 1$ family is the set of uniform cubed-sphere grids with resolutions ranging from 3 to 0.5°. The $\times 2$ family (shown in Fig. 3) is similar, but each $\times 2$ grid has a refined region with twice the resolution ($N = 2$). The $\times 4$ family has a refined region with 4 times the resolution ($N = 4$) and the $\times 8$ family has $N = 8$.

It is nontrivial to construct the transition region. We need to avoid hanging nodes and prefer the elements to be as close to squares as possible. In Fig. 4, we provide two example $\times 8$ grids with the same $\Delta x_{\text{low}} = 3^\circ$. Here we represent two approaches to construct the transition region. Both are based on periodic templates, as seen in Fig. 5. The transition region in Fig. 5a is constructed by CUBIT, a grid-generating software for complex geometries in two and three dimensions (https://cubit.sandia.gov). Figure 5b contains the tran-

sition generated by SQuadGen. SQuadGen was developed to generate two-dimensional refined spherical grids based on a cubed-sphere (http://climate.uccdavis.edu/squadgen.php).

As seen in Fig. 5, the transition region in Fig. 5a contains nodes of higher valence comparing to the similar region in Fig. 5b. In this context, valence of a node is a number of edges it is connected too. Node valence greater than 4 results in quadrilaterals with more acute angles and more distorted elements, and thus lower valence grids are usually preferred. In Fig. 5a, most nodes are of valence 3–6, with a few of valence 7. In Fig. 5b, most nodes are of valence 3–5 with a few nodes of valence 6. For the approach used in Fig. 5a it is possible to avoid valence-7 nodes altogether with less automated, more user-dependent procedure, but not valence-6 nodes.

After the transition has been constructed, it is standard procedure to apply a smoothing algorithm to further improve grid quality. SQuadGen employs the algorithm of Persson and Strang (2004) with a uniform spring force function. Moreover, the user can choose a halo size around the inner and outer boundaries of the transition region, in terms of graph distances. This halo is then used to define a region where smoothing will be applied. To investigate the performance of our resolution aware hyperviscosity operators, we take two extremes: Non-smooth grids with higher valence nodes generated by CUBIT, and smoothed grids with lower valence nodes generated by SQuadGen. We call the former *high-connectivity (or highly distorted) grids* and latter *low-connectivity grids*.

4 Shallow-water test cases

The shallow water equations on a rotating sphere are given by

$$\frac{\partial \boldsymbol{u}}{\partial t} + (\zeta + f)\boldsymbol{k} \times \boldsymbol{u} + \nabla\left(\frac{1}{2}\boldsymbol{u}^2 + g(h+b)\right) = -\nabla \cdot \boldsymbol{\tau}\nabla\Delta\boldsymbol{u} \tag{13}$$

$$\frac{\partial h}{\partial t} + \nabla \cdot (h\boldsymbol{u}) = -\nabla \cdot \boldsymbol{\tau}\nabla\Delta h. \tag{14}$$

Here h is the fluid thickness, $\boldsymbol{u}$ represents velocity, $\zeta = \boldsymbol{k} \cdot \nabla \times \boldsymbol{u}$ the vorticity, f is the Coriolis parameter, g is gravity, and b denotes bottom topography. The equations are discretized following Taylor and Fournier (2010) with the hyperviscosity operator discretized as per Sect. 2. We take $p = 3$ for a fouth order accurate spatial discretization and use the 2nd order accurate Leapfrog–Trapezoidal timestepping method.

For our studies, we choose two standard shallow water test cases from Williamson et al. (1992), test case 2 (TC2) and test case 5 (TC5). TC2 represents a global steady state zonal geostrophic flow. Since the analytical solution is known, TC2 is often used to investigate convergence rates. The solution is very smooth, resulting in small errors, but the errors are very sensitive to local fluctuations in truncation error such as caused by grid irregularities. Following Ringler et al. (2011) we run TC2 for 12 simulation days instead of the originally proposed 5 days, in order to allow longer time for the error growth to disrupt the steady state solution.

TC5 consists of a more realistic zonal flow over an isolated mountain run for 15 days. Error measures are obtained from a high resolution reference simulation. Establishing convergence rates is difficult as the rate decreases to zero as the errors approach the uncertainty in the reference solution. Instead, global errors are used primarily to measure the impact of the refined region. TC5 has much larger errors than TC2 which are less sensitive to small fluctuations in the local truncation error. In TC5 we examine the vorticity which contains small scale structures that are only captured at high resolution. We use the vorticity results to ensure that these structures can also be captured within the high resolution region of a variable resolution grid.

One of the purposes of this study is to confirm that in the SEM with hyperviscosity, the large scale errors are not harmed by the presence of refinement and thus are primarily controlled by coarse resolution Δx_{low}. For this, we collect a series of grids with $\times 2$, $\times 4$, and $\times 8$ refinements. For both TC2 and TC5, the refined region covers a circle with coordinates $\lambda = 3\pi/2$, $\theta = \pi/6$ and radius $\pi/9$. This placement is centered over the TC5 mountain. We summarize characteristics of the grids in Table 1. In Table 2 we summarize some parameters for simulations in Figs. 6–10. Resolutions Δx_{low} and Δx_{high} are computed according to the formula for an equatorial uniform resolution, considering that the whole sphere is covered by corresponding large or small elements.

5 Numerical results

5.1 Grid and hyperviscosity sensitivity in TC2

We present error plots for the TC2 height field h after 12 days in Fig. 6. The error

$$\Delta h = h_{\text{numerical}} - h_{\text{analytic}} \tag{15}$$

is contoured for several different meshes all with a relatively low $\Delta x_{\text{low}} = 3^\circ$ resolution. Part a contains a plot for a uniform resolution and constant-coefficient hyperviscosity. Plots b and d are simulations with the scalar hyperviscosity on $\times 8$ grids. Plots c and e are simulations with the tensor hyperviscosity on $\times 8$ grids.

We first note that the errors are quite small relative to the height field (which ranges from $1000\,\mathrm{m}$ to $3000\,\mathrm{m}$). The height field is not plotted since it would be identical to the analytic solution in Williamson et al. (1992). For the uniform grid with constant coefficient hyperviscosity, the error is quite uniform with no indication of any grid sensitivity. There is no visible $m = 4$ mode that might be expected because of the cubed-sphere grid.

To investigate performance and robustness of two dissipation mechanisms, we chose the two $\times 8$ grids shown in Fig. 4: a highly distorted unsmoothed grid generated by CUBIT, *deg3-*$\times 8$*-highconn*, and a grid with low-connectivity nodes and selectively applied smoothing generated by SQuadGen, *deg3-*$\times 8$*-lowconn*. Both the $\times 8$ grids have the same low resolution region as the uniform grid, and thus the errors in the $\times 8$ grids should be equal or lower than the errors in panel a. This is obviously not the case when scalar hyperviscosity is used, as seen in panels b and d. Those results are contaminated with significant numerical noise with larger errors than panel a. This is even true in the Southern Hemisphere, away from the region of local refinement. Comparing the distorted grid (panel b) with the low connectivity grid (panel d), we see that the scalar hyperviscosity has some grid sensitivity, as the better quality grid in panel d shows more zonal contours similar to panel a and somewhat less noise in the refined mesh region, although panel d does have larger minimum and maximum errors (given in the figure caption) than panel b. Both panels b and d have minimum and maximum errors significantly larger than panel a.

Contrary to this, results using the tensor coefficient hyperviscosity are very close to panel a and much less sensitive to the different types of refinement. The minimum and maximum errors with either the distorted grid (panel c) or the low connectivity grid (panel e) are slightly less than the values obtained on the uniform grid (panel a). In both panels c and e, the error contours in the Southern Hemisphere are almost identical to panel a. In the Northern Hemisphere, the errors are sensitive to the presence of refinement, but are actually lower in this region than with panel a. Thus with the tensor hyperviscosity, the presence of mesh refinement does no harm to the solution and actually results in a minor local reduction in the error.

5.2 Vorticity in TC5

We now examine the vorticity field for TC5. The vorticity after 15 days is plotted in Fig. 7. A closeup of the region over the TC5 mountain is shown in Fig. 8. In both figures, panel a shows a low resolution uniform mesh result, and panel b shows the reference solution that was computed on a uniform high resolution grid. Note the sharp gradient in the flow that is well resolved in our reference solution but not present in the low resolution result.

We show results computed using locally refined grids in Figs. 9 and 10. Based on the results presented for TC2, here we compare the worst and best performing extremes: scalar hyperviscosity running on the highly distorted $\times 8$ mesh in Fig. 4a and tensor hyperviscosity on the low connectivity $\times 8$ mesh in Fig. 4b.

The vorticity is plotted over the mesh refinement region for these two simulations in Fig. 9. The computation with scalar hyperviscosity develops unphysical oscillations which are not present in the tensor hyperviscosity result or reference solution. Figure 9b shows a very smooth field across the highly non-uniform transition region and the sharp gradient that is present in Fig. 8b is resolved without numerical noise. Note that exact matching of Figs. 8b and 9b is not expected because the reference solution used a grid with three times finer resolution then the finest resolution used in the $\times 8$ grids.

To quantify these observations, we plot the error in the vorticity,

$$\Delta\zeta = \zeta_{\text{numerical}} - \zeta_{\text{reference}} \tag{16}$$

field in Fig. 10. We show the error for the two $\times 8$ simulations as well as a uniform low resolution simulation. The error is computed using our high resolution reference solution as an approximation to the exact solution. The noise seen in the scalar hyperviscosity vorticity field (Fig. 9) is more evident throughout the refined region in the vorticity error plot (Fig. 10b). With the tensor viscosity on the low connectivity grid, Fig. 10c shows very little noise in the refinement region and the mesh transition region. In addition, the error is substantially reduced in the refinement region as compared to the low resolution uniform grid solution (Fig. 10b). The fact that the error after 15 days in this region can be reduced by local mesh refinement suggests that the solution contains standing features induced by the mountain which benefit from mesh refinement and are not sensitive to the solution in the rest of the domain where both grids have the same Δx_{low} resolution. In fact, the improved resolution of these standing features leads to slightly less error than is obtained by the global Δx_{low} resolution grid.

To investigate effects of tensor hyperviscosity, one can compare panels Fig. 10b (scalar hyperviscosity and the highly distorted grid) and Fig. 10d (tensor hyperviscosity and the highly distorted grid). The numerical noise in the transition region present in the simulation with scalar hyperviscosity is practically eliminated when tensor formulation is used. If the simulation Fig. 10c for a better quality grid (the low-connectivity grid) is considered optimal, then one can conclude that the tensor hyperviscosity simulation provides a very close to optimal result even if a low quality grid is used.

5.3 Convergence under grid refinement

We now present mesh convergence results for several choices of local refinement. We use our best configuration, tensor hyperviscosity running with low-connectivity grids. We compare the convergence properties of the method with uniform grids using constant coefficient hyperviscosity, uniform grids using tensor coefficient hyperviscosity, the $\times 2$ family of grids, the $\times 4$ family of grids and the $\times 8$ family of grids. The $\times 2$, $\times 4$ and $\times 8$ simulations all use tensor hyperviscosity.

We want to focus only on the spatial error, and in all cases use timesteps where the time truncation error is negligible as

compared to the spatial error. While running simulations for convergence we made sure that temporal errors did not dominate. For TC2, we obtain 4th order convergence when using timesteps near the CFL limit. For TC5 we reduced time steps so that time truncation errors are of 4th order. For example, for the resolution with spatial scales Δx the time is Δt. For the refined grid with spatial scales $\Delta x/2$ we execute a simulation with $\Delta t/4$.

The global l_2 errors for the TC2 height field for all these families of grids are shown in Fig. 11. We use the relative l_2 defined in Williamson et al. (1992). As noted in Sect. 2.1, the hyperviscosity scaling with resolution is typically chosen as $s = 3.2$. For TC2, we instead choose $s = 4$ so that the hyperviscosity term goes to zero at a 4th order rate and we can confirm the 4th order accuracy of our $p = 3$ SEM spatial discretization. For TC2 mesh refinement adds no value to the simulation, and as we saw in Sect. 5.1 it can lead to a slight reduction in the local error but no reduction in the error away from the refinement region. We thus examine convergence with respect to Δx_{low} for each family of grids. As expected, the uniform resolution simulations with 4th order constant-coefficient hyperviscosity demonstrate 4th order convergence, with the tensor and scalar hyperviscosity results nearly identical. Similar results are obtained for the $\times 2$, $\times 4$ and $\times 8$ family of grids. The error is completely determined by Δx_{low} for all grids, and all grids show 4th order convergence under mesh refinement with respect to Δx_{low}. Thus the presence of mesh refinement, with refinements as much as 8x, does no harm to the global errors.

The global l_2 errors for the TC5 height field for all these families of grids are shown in Fig. 12, again plotted as a function of Δx_{low} and normalized as in Williamson et al. (1992). For TC5, we return to the conventional hyperviscosity resolution scaling of $s = 3.2$. For TC5, we compute the l_2 errors from our high resolution reference solution. The convergence rates are lower in this case due to the fact that the mountain is not smooth, limiting the convergence to the 1st order in the max norm and second order in the l_2 norm. We first note that for uniform resolution grids, the constant coefficient hyperviscosity performs nearly identical to the tensor coefficient hyperviscosity. For TC5 we also see that for grids with the same Δx_{low}, the global l_2 error is slightly reduced by the presence of mesh refinement, as conjectured in Sect. 5.2. The effect is small, and fully captured by the $\times 2$ grid with twice the resolution over the TC5 mountain. Further local refinement in the $\times 4$ and $\times 8$ grids does not further improve the error. Thus in this case, the presence of mesh refinement does no harm to the global errors and in some special cases can decrease global error.

6 Conclusions

We compared two resolution aware hyperviscosity operators for the SEM running on unstructured grids with a region of local mesh refinement: a conventional scalar approach based on a single length scale for each element, and a tensor approach that respects the resolution scaling of both length scales within each element. In both shallow water test cases 2 and 5, the scalar approach had noticeable noise and oscillations near regions of local mesh refinement which was not present with the tensor formulation. Results for both formulations were sensitive to the grid quality, as shown by comparing results on a highly distorted grid with sharp mesh transitions and a smooth grid with less acute angles due to its lower valence nodes. But in TC2, the tensor formulation showed less grid sensitivity and obtained excellent results on both grids.

When running with tensor hyperviscosity in the SEM, the presence of local mesh refinement in TC2 had no impact on the global errors. The SEM obtained its formal order of accuracy for all grids tested (up to 8x regional refinement). In TC5, with refinement over the mountain, the presence of refinement again did no harm to the global errors and actually resulted in a small improvement. Asymptotically, the global errors were controlled by the coarse resolution and the locally refined meshes obtained the same convergence rates as the global uniform meshes.

The tensor and scalar hyperviscosity operators were constructed to be resolution aware and to preserve the resolution scaling often used with uniform grids. In a high resolution region of quasi-uniform elements, the resolution aware operators are very similar to the SEM's well proven constant coefficient operator. Hence we expect that within the high resolution region of a locally refined mesh, the SEM can locally resolve the same types of structures that the uniform high resolution grid can resolve. This was verified by looking at steep gradients in TC5 and comparing to coarse and high uniform resolution solutions.

Appendix A: An element-local map for quadrilaterals on a sphere

Numerical methods for the sphere based on cubed-sphere grids need to define a map $\boldsymbol{r}(\boldsymbol{x})$ from the reference element to the sphere (in the case of finite element methods) or from each cube face to the sphere for finite difference or finite element methods. Most approaches use the equidistant central projection (Sadourny, 1972), the equiangular central projection (Rancic et al., 1996), or their combination (Fournier et al., 2004). All three of these approaches were compared in Nair et al. (2005), where the equiangular mapping was found to be the most accurate. However, all three aforementioned projections are based on an inscribed cube and cannot correctly treat elements lying across cube edges. In particular, for an edge of such an element, the reference element maps for the two elements which share this edge may not agree, resulting in a loss of the SEM's mimetic properties. Here we present a map which avoids this issue by using a map local to

each element. The map uses a bilinear transformation based on elements' physical coordinates and does not make use of an inscribed cube. It is similar to the map used for triangular elements in Lauter et al. (2008).

For each quadrilateral element Ω_m on the surface of the unit sphere, we denote the map and its inverse by $\boldsymbol{r}(\boldsymbol{x};m)$ and $\boldsymbol{x}(\boldsymbol{r};m)$, where $\boldsymbol{x}=(x_1,x_2)$. To construct $\boldsymbol{r}(\boldsymbol{x};m)$, let $\boldsymbol{c}_1$, $\boldsymbol{c}_2$, $\boldsymbol{c}_3$, and $\boldsymbol{c}_4$ be Cartesian coordinates of the vertices of Ω_m with $\boldsymbol{c}_1=(c_1^x,c_1^y,c_3^z)^{\mathrm{T}}$, etc. and define

$$\boldsymbol{r}=\tilde{\boldsymbol{r}}/\|\tilde{\boldsymbol{r}}\|_2$$

with

$$\tilde{\boldsymbol{r}}=\frac{1}{4}\left((1-x_1)(1-x_2)\boldsymbol{c}_1+(1+x_1)(1-x_2)\boldsymbol{c}_2+(1+x_1)(1+x_2)\boldsymbol{c}_3+(1-x_1)(1+x_2)\boldsymbol{c}_4\right). \tag{A1}$$

We now give an analytical expression for $\mathbf{D}=\partial\boldsymbol{r}/\partial\boldsymbol{x}$ needed in Sect. 2. We use both Cartesian and longitude-latitude coordinates:

$$\boldsymbol{r}=\begin{pmatrix}\cos\lambda\cos\theta\\ \sin\lambda\cos\theta\\ \sin\theta\end{pmatrix}\text{ where }\lambda\in[0,2\pi],\ \theta\in[-\pi/2,\pi/2]. \tag{A2}$$

Since $\mathrm{d}\boldsymbol{r}=\cos\theta\,\mathrm{d}\lambda\,\boldsymbol{e}_\lambda+\mathrm{d}\theta\,\boldsymbol{e}_\theta$, it follows that

$$\mathbf{D}=\begin{pmatrix}\cos\theta & 0\\ 0 & 1\end{pmatrix}\frac{\partial(\lambda,\theta)}{\partial\boldsymbol{x}}=\begin{pmatrix}\cos\theta & 0\\ 0 & 1\end{pmatrix}\frac{\partial(\lambda,\theta)}{\partial\boldsymbol{r}}\frac{\partial\boldsymbol{r}}{\partial\boldsymbol{x}}. \tag{A3}$$

To avoid the singularity at the poles in the term

$$\begin{pmatrix}\cos\theta & 0\\ 0 & 1\end{pmatrix}\frac{\partial(\lambda,\theta)}{\partial\boldsymbol{r}}=\begin{pmatrix}-\sin\lambda & \cos\lambda & 0\\ 0 & 0 & \frac{1}{\cos\theta}\end{pmatrix}$$

we further decompose the term $\frac{\partial\boldsymbol{r}}{\partial\boldsymbol{x}}=\frac{\partial\boldsymbol{r}}{\partial\tilde{\boldsymbol{r}}}\frac{\partial\tilde{\boldsymbol{r}}}{\partial\boldsymbol{x}}$ so that we can extract a factor of $\cos\theta$ from $\frac{\partial\boldsymbol{r}}{\partial\tilde{\boldsymbol{r}}}$. Some algebra shows

$$\frac{\partial\boldsymbol{r}}{\partial\tilde{\boldsymbol{r}}}=\frac{1}{||\tilde{\boldsymbol{r}}||_2}\begin{pmatrix}a_{11} & a_{12} & a_{13}\\ a_{21} & a_{22} & a_{23}\\ a_{31} & a_{32} & a_{33}\end{pmatrix}$$

where $a_{11}=\sin^2\lambda\cos^2\theta+\sin^2\theta$, $a_{12}=-\frac{1}{2}\sin2\lambda\cos^2\theta$, $a_{13}=-\frac{1}{2}\cos\lambda\sin2\theta$, $a_{21}=-\frac{1}{2}\sin2\lambda\cos^2\theta$, $a_{22}=\cos^2\lambda\cos^2\theta+\sin^2\theta$, $a_{23}=-\frac{1}{2}\sin\lambda\sin2\theta$, $a_{31}=-\frac{1}{2}\cos\lambda\sin2\theta$, $a_{32}=-\frac{1}{2}\sin\lambda\sin2\theta$, and $a_{33}=\cos^2\theta$.

Also,

$$\frac{\partial\tilde{\boldsymbol{r}}}{\partial\boldsymbol{x}}=\frac{1}{4}\begin{pmatrix}c_1^x & c_2^x & c_3^x & c_4^x\\ c_1^y & c_2^y & c_3^y & c_4^y\\ c_1^z & c_2^z & c_3^z & c_4^z\end{pmatrix}\begin{pmatrix}-1+x_2 & -1+x_1\\ 1-x_2 & -1-x_1\\ 1+x_2 & 1+x_1\\ -1-x_2 & 1-x_1\end{pmatrix}.$$

All together, we have

$$\mathbf{D}=\frac{1}{||\tilde{\boldsymbol{r}}||_2}\frac{1}{4}\begin{pmatrix}-\sin\lambda & \cos\lambda & 0\\ 0 & 0 & 1\end{pmatrix}\cdot\begin{pmatrix}a_{11} & a_{12} & a_{13}\\ a_{21} & a_{22} & a_{23}\\ a_{31}/\cos\theta & a_{32}/\cos\theta & a_{33}/\cos\theta\end{pmatrix}\cdot\begin{pmatrix}c_1^x & c_2^x & c_3^x & c_4^x\\ c_1^y & c_2^y & c_3^y & c_4^y\\ c_1^z & c_2^z & c_3^z & c_4^z\end{pmatrix}\begin{pmatrix}-1+x_2 & -1+x_1\\ 1-x_2 & -1-x_1\\ 1+x_2 & 1+x_1\\ -1-x_2 & 1-x_1\end{pmatrix},$$

where

$$\tilde{\boldsymbol{r}}=\frac{1}{4}\begin{pmatrix}c_1^x & c_2^x & c_3^x & c_4^x\\ c_1^y & c_2^y & c_3^y & c_4^y\\ c_1^z & c_2^z & c_3^z & c_4^z\end{pmatrix}\begin{pmatrix}(1-x_1)(1-x_2)\\ (1+x_1)(1-x_2)\\ (1+x_1)(1+x_2)\\ (1-x_1)(1+x_2)\end{pmatrix}.$$

Acknowledgements. This research was supported by the Department of Energy Office of Biological and Environmental Research, work package 12-015334 "Multiscale Methods for Accurate, Efficient, and Scale-Aware Models of the Earth System" and work package 11-014996 "Climate Science for a Sustainable Energy Future". Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the US Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

We would like to thank two anonymous reviewers for their suggestions which significantly improved the manuscript.

References

Ainsworth, M. and Wajid, H.: Dispersive and dissipative behavior of the spectral element method, SIAM J. Numer. Anal., 47, 3910–3937, 2009.

Baer, F., Wang, H., Tribbia, J., and Fournier, A.: Climate modeling with spectral elements, Mon. Weather Rev., 134, 3610–3624, 2006.

Blacker, T. D. and Stephenson, M. B.: Paving: a new approach to automated quadrilateral mesh generation, Int. J. Numer. Meth. Eng., 32, 811–847, 1991.

Boville, B.: Sensitivity of simulated climate to model resolution, J. Climate, 4, 469–485, 1991.

Dennis, J. M., Edwards, J., Evans, K. J., Guba, O., Lauritzen, P. H., Mirin, A. A., St.-Cyr, A., Taylor, M. A., and Worley, P. H.: CAM-SE: a scalable spectral element dynamical core for the Community Atmosphere Model, Int. J. High Perform. C., 26, 74–89, 2012.

Dobrev, V., Kolev, T., and Rieben, R.: High-order curvilinear finite element methods for Lagrangian hydrodynamics, SIAM J. Sci. Comput., 34, B606–B641, 2012.

Fournier, A., Taylor, M., and Tribbia, J.: The spectral element atmosphere model (SEAM): high-resolution parallel computation and localized resolution of regional dynamics, Mon. Weather Rev., 132, 726–748, 2004.

Lauter, M., Giraldo, F. X., Handorf, D., and Dethloff, K.: A discontinuous Galerkin method for the shallow water equations in spherical triangular coordinates, J. Comput. Phys., 227, 10226–10242, 2008.

Marras, S., Kopera, M., and Giraldo, F.: Grid Effect on Spherical Shallow Water Jets using Continuous and Discontinuous Galerkin Methods, Q. J. Roy. Meteorol. Soc., under review, 2014.

Nair, R. D., Thomas, S., and Loft, R.: A discontinuous Galerkin transport scheme on the cubed-sphere, Mon. Weather Rev., 133, 814–828, 2005.

Persson, P.-O. and Strang, G.: A simple mesh generator in MATLAB, SIAM Rev., 46, 329–345, 2004.

Rancic, M., Purser, R., and Mesinger, F.: A global shallow-water model using an expanded spherical cube: gnomonic versus conformal coordinates, Q. J. Roy. Meteorol. Soc., 122, 959–982, 1996.

Ringler, T. D., Jacobsen, D., Gunzburger, M., Ju, L., Duda, M., and Skamarock, W.: Exploring a multi-resolution modeling approach within the shallow-water equations, Mon. Weather Rev., 139, 3348–3368, 2011.

Sadourny, R.: Conservative finite-difference approximations of the primitive equations on quasi-uniform spherical grids, Mon. Weather Rev., 100, 136–144, 1972.

St.-Cyr, A., Jablonowski, C., Dennis, J. M., Tufo, H. M., and Thomas, S. J.: A comparison of two shallow water models with non-conforming adaptive grids, Mon. Weather Rev., 136, 1898–1922, 2008.

Takahashi, Y., Hamilton, K., and Ohfuchi, W.: Explicit global simulation of the mesoscale spectrum of atmospheric motions, Geophys. Res. Lett., 33, L12812, doi:10.1029/2006GL026429, 2006.

Taylor, M. A. and Fournier, A.: A compatible and conservative spectral element method on unstructured grids, J. Comput. Phys., 229, 5879–5895, 2010.

Weller, H., Weller, H. G., and Fournier, A.: Voronoi, Delaunay, and block-structured mesh refinement for solution of the shallow-water equations on the sphere, Mon. Weather Rev., 137, 4208–4224, 2009.

Williamson, D. L., Drake, J. B., Hack, J. J., Jakob, R., and Swarztrauber, P. N.: A standard test set for numerical approximations to the shallow water equations in spherical geometry, J. Comput. Phys., 102, 211–224, 1992.

Yu, M., Giraldo, F. X., Peng, M., and Wang, Z.: Localized artificial viscosity stabilization of discontinuous Galerkin methods for nonhydrostatic mesoscale atmospheric modeling, J. Comput. Phys., under review, 2014.

Zarzycki, C., Jablonowski, C., and Taylor, M.: Using variable-resolution meshes to model tropical cyclones in the Community Atmosphere Model, Mon. Weather Rev., 142, 1221–1239, 2014a.

Zarzycki, C., Levy, M., Jablonowski, C., Taylor, M., Overfelt, J., and Ullrich, P.: Aqua planet experiments using CAM's variable-resolution dynamical core, J. Climate, online first, doi:10.1175/JCLI-D-14-00004.1, 2014b.

Figure 1. A cubed-sphere grid

Table 1. Summary of the $\times 2$, $\times 4$ and $\times 8$ family of grids.

name	Δx_{low}	Δx_{high}	refinement	connectivity
deg3-$\times$2-lowconn	3°	1.5°	$\times 2$	low-connectivity
deg1.5-$\times$2-lowconn	1.5°	0.75°	$\times 2$	low-connectivity
deg1-$\times$2-lowconn	1°	0.5°	$\times 2$	low-connectivity
deg0.75-$\times$2-lowconn	0.75°	0.375°	$\times 2$	low-connectivity
deg0.5-$\times$2-lowconn	0.5°	0.25°	$\times 2$	low-connectivity
deg3-$\times$4-lowconn	3°	0.75°	$\times 4$	low-connectivity
deg1.5-$\times$4-lowconn	1.5°	0.375°	$\times 4$	low-connectivity
deg1-$\times$4-lowconn	1°	0.25°	$\times 4$	low-connectivity
deg0.5-$\times$4-lowconn	0.5°	0.125°	$\times 4$	low-connectivity
deg3-$\times$8-lowconn	3°	0.375°	$\times 8$	low-connectivity
deg1.5-$\times$8-lowconn	1.5°	0.1875°	$\times 8$	low-connectivity
deg1-$\times$8-lowconn	1°	0.125°	$\times 8$	low-connectivity
deg3-$\times$8-highconn	3°	0.375°	$\times 8$	high-connectivity

Table 2. Summary for simulations. TC stands for a test case, numbers c_0 and s are parameters in the hyperviscosity coefficient $\nu = c_0(\Delta x)^s$.

TC	grid	HV method	c_0	s	Δt	figure
TC2	uniform, 3°	constant-coef.	6.12×10^{-6}	4	50 s	Fig. 6a
TC2	deg3-$\times$8-highconn	scalar	6.52×10^{-6}	4	30 s	Fig. 6b
TC2	deg3-$\times$8-highconn	tensor-based	6.12×10^{-6}	4	30 s	Fig. 6c
TC2	deg3-$\times$8-lowconn	scalar	6.52×10^{-6}	4	30 s	Fig. 6d
TC2	deg3-$\times$8-lowconn	tensor-based	6.12×10^{-6}	4	30 s	Fig. 6e
TC5	uniform, 3°	constant-coef.	7.18×10^{-2}	3.2	50 s	Fig. 10a
TC5	deg3-$\times$8-highconn	scalar	7.18×10^{-2}	3.2	20 s	Figs. 9b, 10b
TC5	deg3-$\times$8-lowconn	tensor-based	3.59×10^{-2}	3.2	50 s	Figs. 9c, 10c
TC5	deg3-$\times$8-highconn	tensor-based	3.59×10^{-2}	3.2	50 s	Fig. 10d

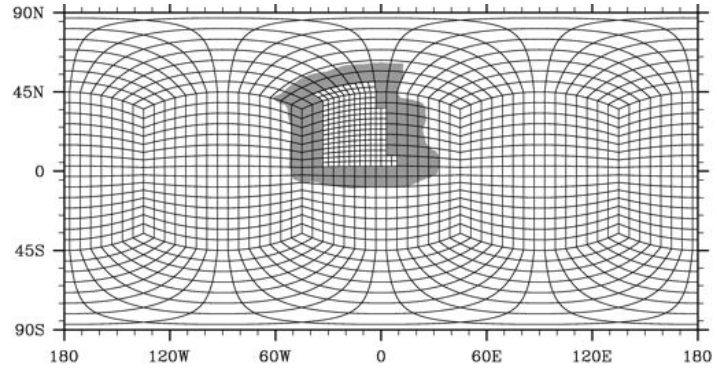

Figure 2. Schematic idea of constructing refined grids for conforming quadrilaterals on a sphere: We start from uniform grid with Δx_{low}. Next, the region of desired refinement is replaced with uniform elements of size Δx_{high}. The grey area approximately defines a transition region which is constructed by substituting quadrilaterals with Δx_{low} by transition templates. After transition region is assembled, spring dynamics can be used to smooth the grid.

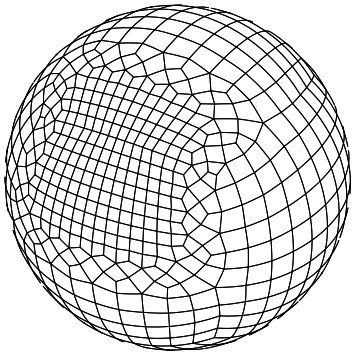

(a) $\Delta x_{\text{low}} = 3^\circ$

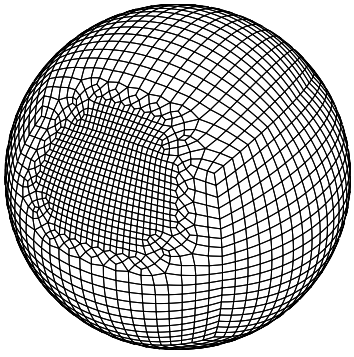

(b) $\Delta x_{\text{low}} = 1.5^\circ$

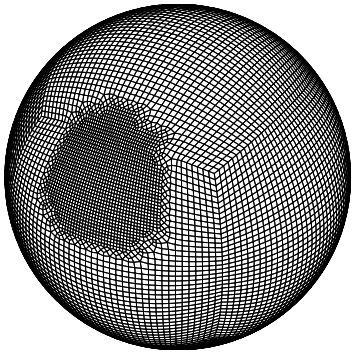

(c) $\Delta x_{\text{low}} = 0.75^\circ$

Figure 3. A family of $\times 2$ grids with, from left to right, $\Delta x_{\text{low}} = 3^\circ$, 1.5°, and 0.75°, $\Delta x_{\text{high}} = \Delta x_{\text{low}}/2$.

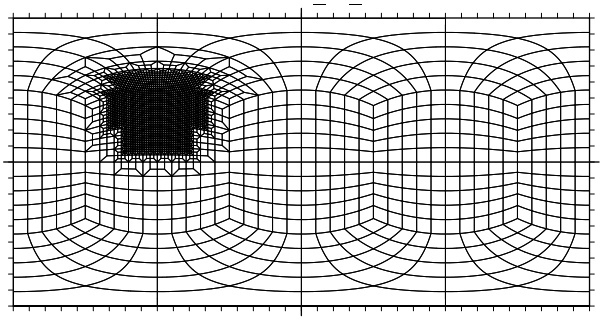

(a) Grid with a narrow, more distorted transition region from CUBIT mesh generation software with grid smoothing turned off

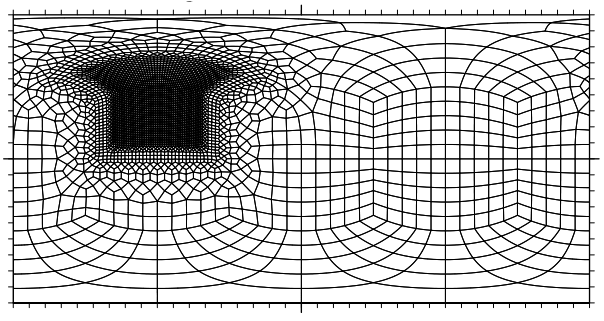

(b) Grid with a wider, more uniform transition region from SQuadGen and grid smoothing with spring dynamics

Figure 4. Example of $\times 8$ refined grids with $\Delta x_{\text{low}} = 3^\circ$, $\Delta x_{\text{high}} = \Delta x_{\text{low}}/8$

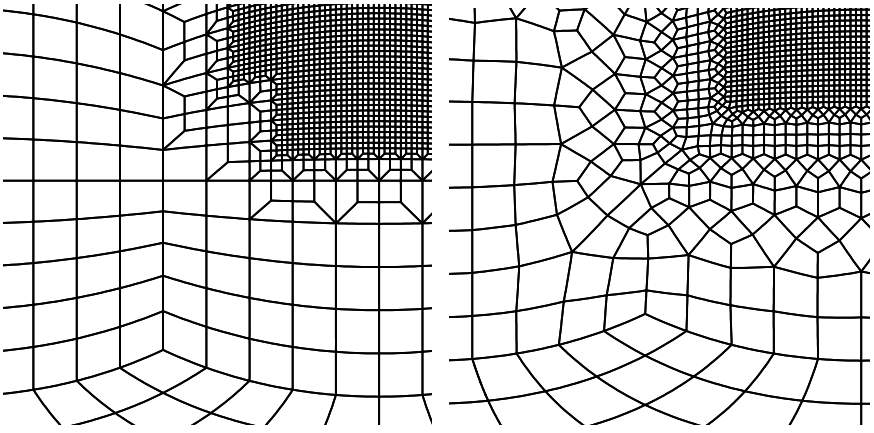

(a) CUBIT approach (b) SQuadGen approach

Figure 5. Different types of refined conforming quadrilateral grids. Plots **(a)** and **(b)** show closeups of the transition regions from plots Fig. 4(a) and Fig. 4(b) respectively

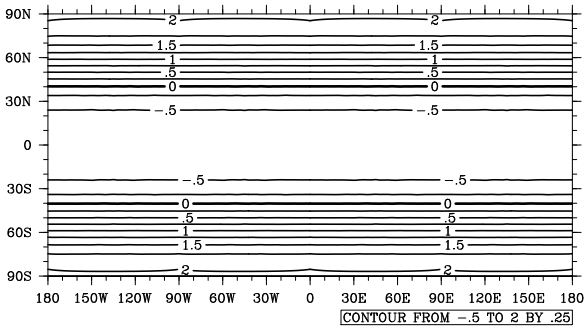

(a) Error plot for the uniform grid. The resolution is $\Delta x_{unit} = 3^{\circ}$. The error varies from -6.06×10^{-1} to 2.02×10^{0}.

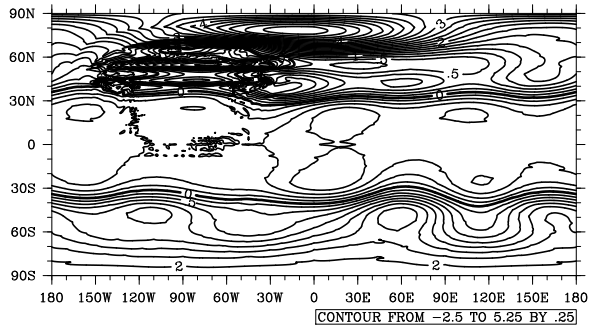

(b) Scalar hyperviscosity with the highly distorted $\times 8$ grid shown in Fig. 4(a), error varies from -2.61×10^{0} to 5.44×10^{0} meters

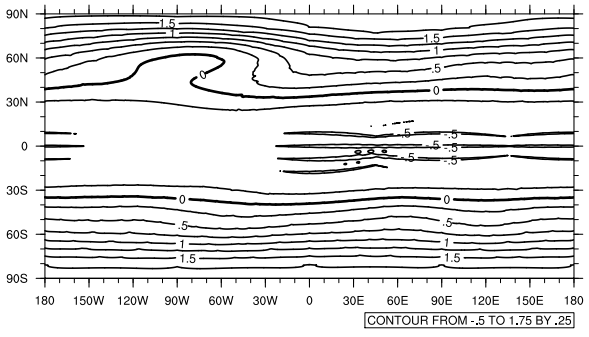

(c) Tensor hyperviscosity with the highly distorted $\times 8$ grid shown in Fig. 4(a), error varies from -5.32×10^{-1} to 1.90×10^{0} meters

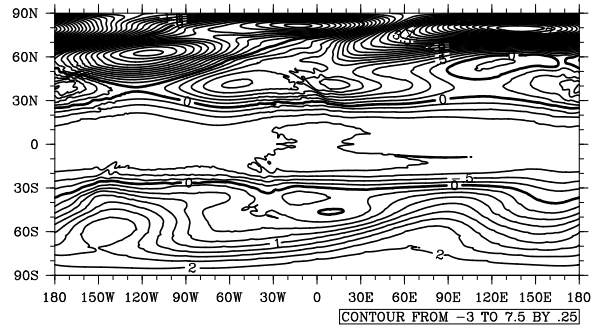

(d) Scalar hyperviscosity with the low-connectivity $\times 8$ grid shown in Fig. 4(b), error varies from -3.09×10^{0} to 7.73×10^{0} meters

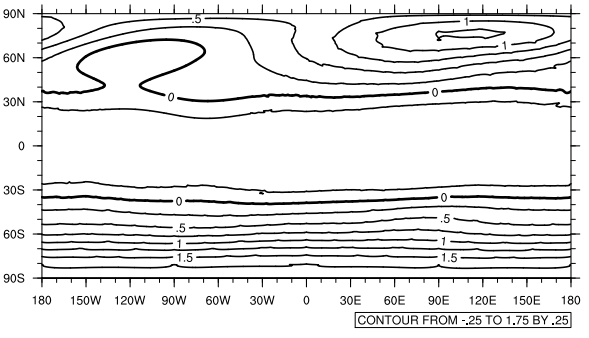

(e) Tensor hyperviscosity with the low-connectivity $\times 8$ grid shown in Fig. 4(b), error varies from -4.85×10^{-1} to 1.91×10^{0} meters

Figure 6. Error plots for TC2. A contour spacing of $0.25\,\mathrm{m}$ is the same for all plots. Maximum and minimum values of error Eq. (15) are given in captions. Tensor hyperviscosity produces smoother fields comparing to scalar hyperviscosity, as follows from comparing pairs **(b)**, **(c)** and **(d)**, **(e)**. In addition, the quality of the underlying grid significantly improves the outcome around the refined region when using scalar hyperviscosity, as seen by comparing **(b)** and **(d)**. Contrary to scalar hyperviscosity, tensor hyperviscosity is more robust with respect to mesh quality, as follows from comparing **(c)** and **(e)**. Simulations **(c)** and **(e)** with tensor hyperviscosity also exhibit a substantial error reduction in the vicinity of the refinement compared to the coarse uniform resolution in **(a)**.

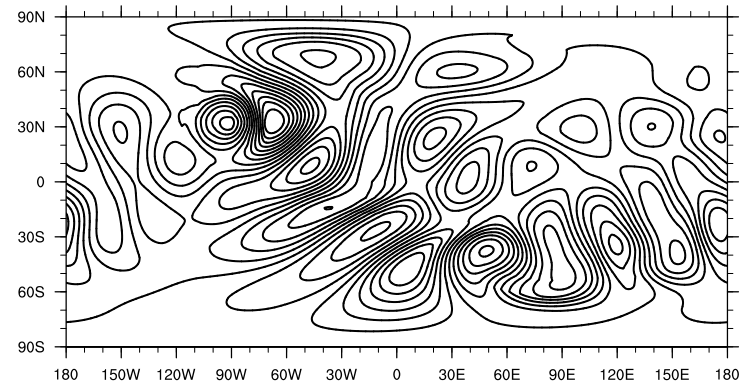

(a) Coarse resolution solution using a uniform grid with $\Delta x = 3^\circ$

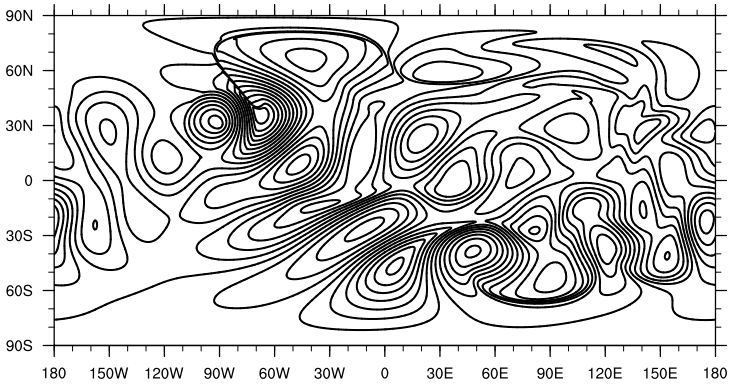

(b) High resolution reference solution, using a uniform grid with $\Delta x = 0.125^\circ$

Figure 7. TC5 vorticity contours for low and high uniform resolutions with constant coefficient hyperviscosity. The contour spacing is $5.0 \times 10^{-6}\,\mathrm{s}^{-1}$. A spherical mountain approximately 30° in diameter is centered at 30N, 90W.

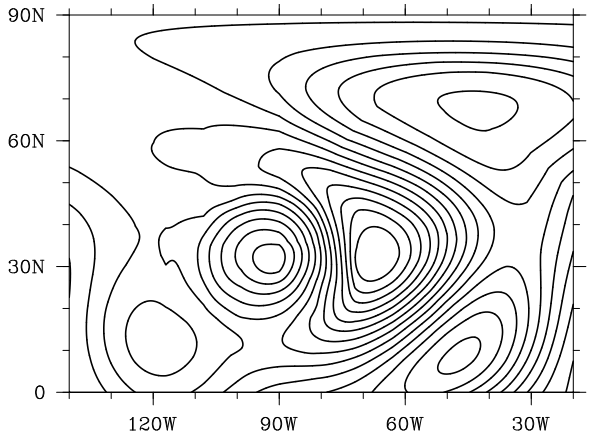

(a) Coarse resolution solution with $\Delta x = 3^\circ$.

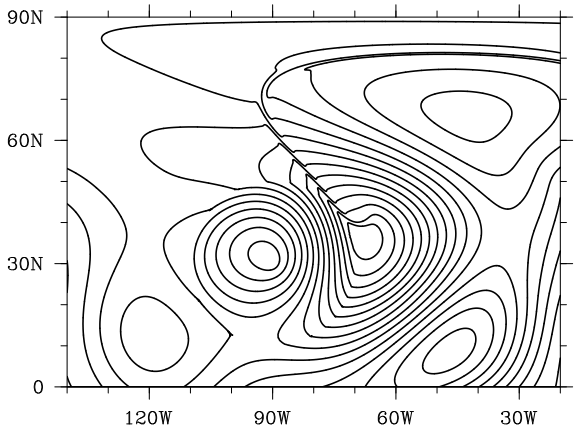

(b) High resolution reference solution, using a uniform grid with $\Delta x = 0.125^\circ$.

Figure 8. As in Fig. 7 but plotted in a subregion of the global domain.

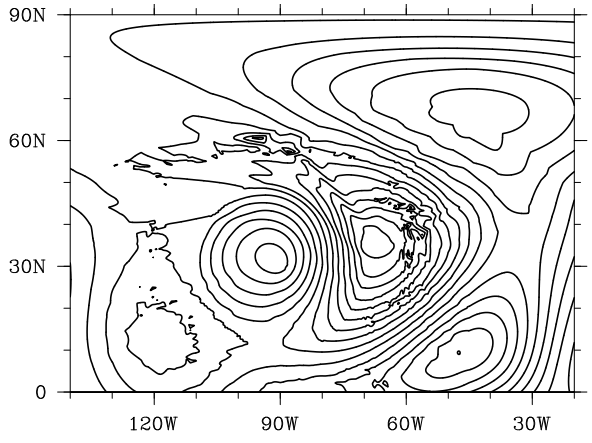

(a) Scalar hyperviscosity with the highly distorted $\times 8$ grid shown in Fig. 4(a).

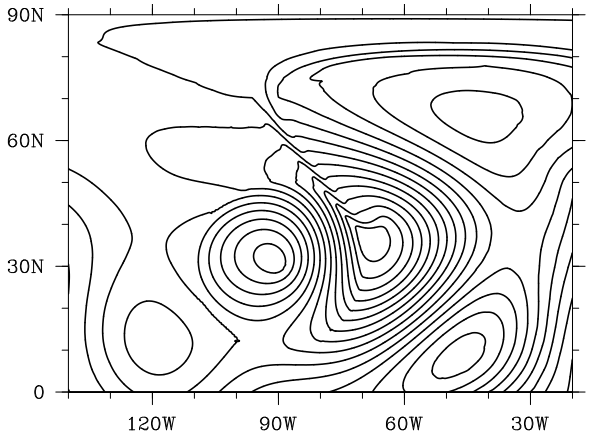

(b) Tensor hyperviscosity with the low connectivity $\times 8$ grid shown in Fig. 4(b).

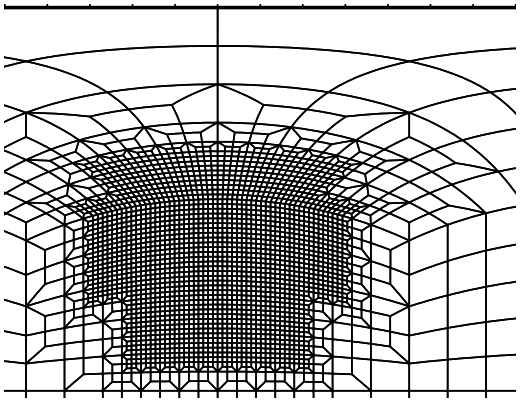

(c) The highly distorted grid used in (a)

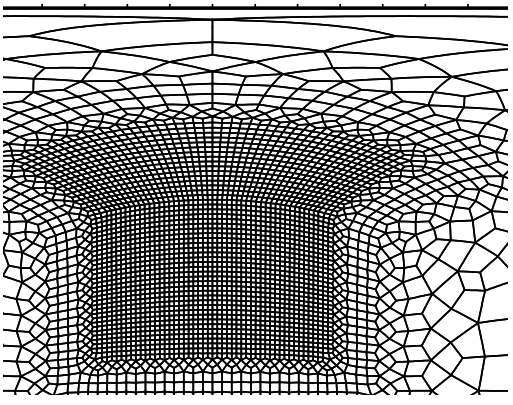

(d) The low connectivity grid used in (b)

Figure 9. Contour plots of TC5 vorticity are shown in the top panels while the grid is shown in the bottom panel. The contour interval is $5.0 \times 10^{-6}\,\mathrm{s}^{-1}$. All panels show a subregion of the global domain which contains most of the refined region and is identical to the subregion used in Fig. 8. Panel **(a)** shows results using scalar hyperviscosity on the highly distorted grid shown in **(c)**. Panel **(b)** shows results using tensor hyperviscosity on the low connectivity grid shown in **(d)**. The improved hyperviscosity and mesh quality result in significantly improved results.

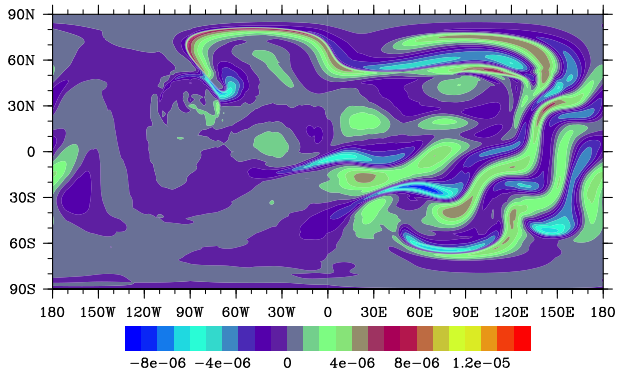

(a) Uniform grid with $\Delta x = 3^\circ$ and constant hyperviscosity. Normalized l_2 error is 1.16×10^{-3}. Error (16) varies from -9.0×10^{-6} to 6.2×10^{-6}.

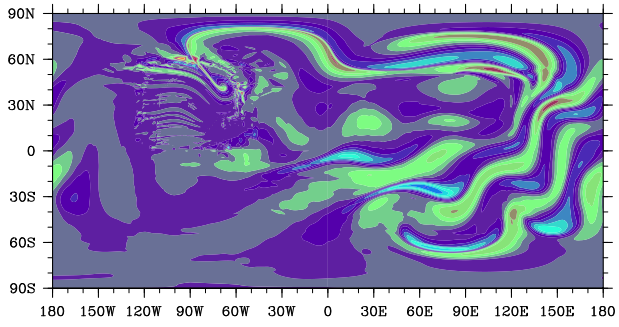

(b) Scalar hyperviscosity using the highly distorted $\times 8$ shown in Fig. 4(a). Normalized l_2 error is 8.44×10^{-4}. Error (16) varies from -8.1×10^{-6} to 1.4×10^{-5}.

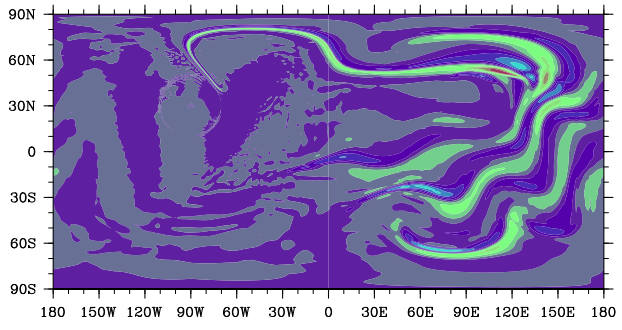

(c) Tensor-based hyperviscosity using the low-connectivity $\times 8$ grid shown in Fig. 4(b). Normalized l_2 error is 2.85×10^{-4}. Error (16) varies from -4.7×10^{-6} to 6.0×10^{-6}.

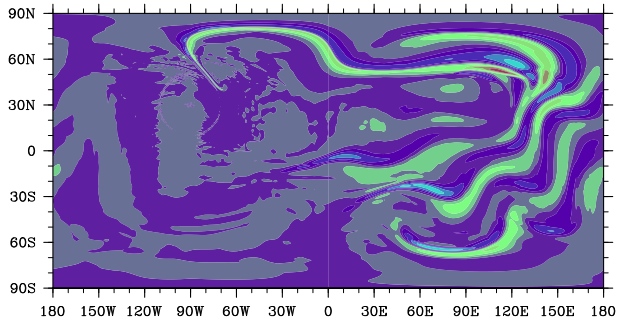

(d) Tensor-based hyperviscosity using the highly distorted $\times 8$ grid shown in Fig. 4(b). Normalized l_2 error is 3.36×10^{-4}. Error (16) varies from -4.7×10^{-6} to 5.75×10^{-6}.

Figure 10. The error in the TC5 vorticity field is plotted for the global domain. The color scheme given in **(a)** is the same for all plots. The tensor hyperviscosity again produces the best results with very little noise. From comparing panel **(c)** (tensor hyperviscosity and the low-connectivity grid) and panel **(d)** (tensor hyperviscosity and the highly distorted grid), we conclude that tensor hyperviscosity is not sensitive to grid quality.

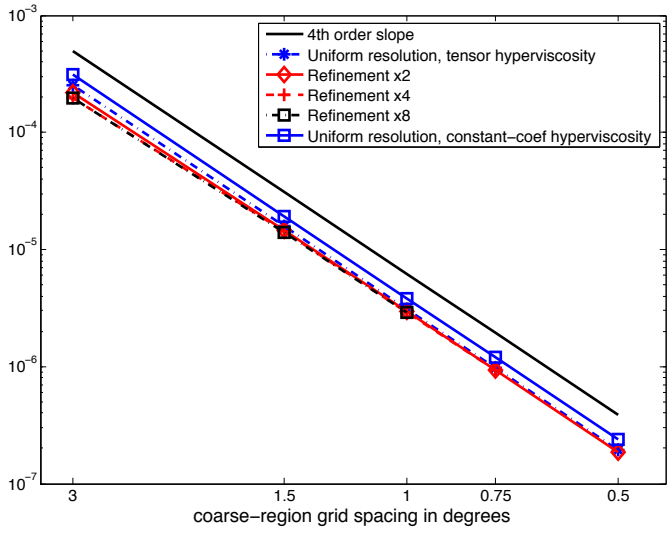

Figure 11. TC2 l_2 errors for uniform and low-connectivity grids plotted as a function of Δx_{low}. The solid line shows 4th order convergence. The error is controlled by the coarse region Δx_{low} and the SEM obtains its formal order of accuracy (4th order for $p = 3$) with respect to Δx_{low}. This is true for both uniform grids (blue squares and blue stars) and grids containing mesh refined regions with $\Delta x_{\text{high}} = \Delta x_{\text{low}}/2$ ($\times 2$ family shown as red diamonds), $\Delta x_{\text{high}} = \Delta x_{\text{low}}/4$ ($\times 4$ family, shown as red plus marks) and $\Delta x_{\text{high}} = \Delta x_{\text{low}}/8$ ($\times 8$ family, shown as black squares). All curves are practically indistinguishable.

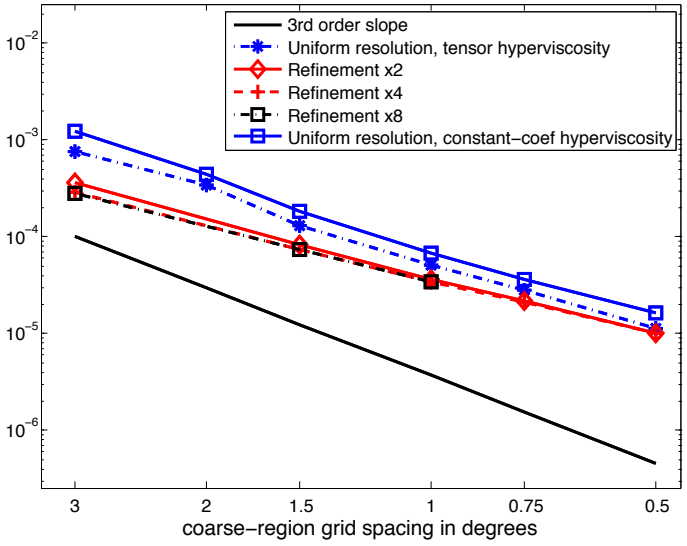

Figure 12. As in Fig. 11, except for TC5. In case of tensor-based hyperviscosity, three error curves for grid refinement with $\times 2$, $\times 4$ or $\times 8$ local refinement are practically indistinguishable and obtain the same convergence rates as seen with uniform grids. For a given Δx_{low}, the presence of some local refinement ($\times 2$) does decrease the global error, but $\times 4$ or $\times 8$ levels of local refinement produce no further benefit.