UC Berkeley

UC Berkeley Previously Published Works

Title

Analyzing the adaptive mesh refinement (AMR) characteristics of a high-order 2D cubed-sphere shallow-water model

Permalink

<https://escholarship.org/uc/item/7kr8c2qs>

Journal

Monthly Weather Review, 144(12)

ISSN

0027-0644

Authors

Ferguson, JO
Jablonowski, C
Johansen, H
et al.

Publication Date

2016

DOI

10.1175/MWR-D-16-0197.1

Peer reviewed

Analyzing the Adaptive Mesh Refinement (AMR) Characteristics of a High-Order 2D Cubed-Sphere Shallow-Water Model

JARED O. FERGUSON AND CHRISTIANE JABLONOWSKI

Department of Climate and Space Sciences and Engineering, University of Michigan, Ann Arbor, Michigan

HANS JOHANSEN, PETER MCCORQUODALE, AND PHILLIP COLELLA

Lawrence Berkeley National Laboratory, Berkeley, California

PAUL A. ULLRICH

Department of Land, Air and Water Resources, University of California, Davis, Davis, California

(Manuscript received 23 May 2016, in final form 21 August 2016)

ABSTRACT

Adaptive mesh refinement (AMR) is a technique that has been featured only sporadically in atmospheric science literature. This paper aims to demonstrate the utility of AMR for simulating atmospheric flows. Several test cases are implemented in a 2D shallow-water model on the sphere using the Chombo-AMR dynamical core. This high-order finite-volume model implements adaptive refinement in both space and time on a cubed-sphere grid using a mapped-multiblock mesh technique. The tests consist of the passive advection of a tracer around moving vortices, a steady-state geostrophic flow, an unsteady solid-body rotation, a gravity wave impinging on a mountain, and the interaction of binary vortices. Both static and dynamic refinements are analyzed to determine the strengths and weaknesses of AMR in both complex flows with small-scale features and large-scale smooth flows. The different test cases required different AMR criteria, such as vorticity or height-gradient based thresholds, in order to achieve the best accuracy for cost. The simulations show that the model can accurately resolve key local features without requiring global high-resolution grids. The adaptive grids are able to track features of interest reliably without inducing noise or visible distortions at the coarse-fine interfaces. Furthermore, the AMR grids keep any degradations of the large-scale smooth flows to a minimum.

1. Introduction

Global climate models have become vital tools for simulating the present and future climate and for predicting important climate trends. However, current global models are limited in their ability to represent many multiscale aspects of atmospheric flows. Their resolutions, limited by computational costs, are too coarse to accurately represent key processes that span a wide range of temporal and spatial scales. High-resolution simulations are essential for capturing these scale interactions and for accurately representing local and regional phenomena such as convection, orographically induced precipitation,

mesoscale storm systems, and tropical cyclones. Such events have large regional impacts as well as broader feedbacks onto the large-scale climate system. Today, high-resolution general circulation models (GCMs) used for global climate simulations can utilize uniform grid spacings down to 10 km as documented by [Manganello et al. \(2012\)](#) or [Kinter et al. \(2013\)](#). However, they are computationally very expensive and still unable to represent key processes such as clouds explicitly. Exceptions are the cloud-permitting, partly cloud-resolving, global simulations by [Miura et al. \(2007\)](#), [Putman and Suarez \(2011\)](#) or [Miyamoto et al. \(2013\)](#) that were run for short, multiday time periods with grid spacings in the 0.9–3.5-km range. To employ such high resolutions over longer climate time periods, modelers typically use limited-area models (LAMs) that focus the computational resources in areas of interest. A major drawback is that LAMs require their lateral boundaries to be externally forced. These

Corresponding author address: Jared O. Ferguson, Department of Climate and Space Sciences and Engineering, University of Michigan, 2455 Hayward, Ann Arbor, MI 48109.
E-mail: joferg@umich.edu

boundary conditions are typically derived from much coarser GCMs, which use different numerical schemes and physical parameterizations, thereby introducing possible biases or numerical discrepancies. In addition, it is an open question how well LAMs can capture teleconnections of global large-scale dynamics and localized features particularly for tropical cyclones and other phenomena that have feedbacks onto the larger climate system.

Variable-resolution GCMs can utilize static or dynamic grid refinements, which are promising options to bridge the gap between global and regional climate modeling. Application examples for static (nonmoving) mesh adaptations are provided in [Zarzycki et al. \(2014\)](#), [Rauscher and Ringler \(2014\)](#), [Zarzycki and Jablonowski \(2015\)](#), and [Huang et al. \(2016\)](#) (see also further references therein). Our paper focuses on dynamically adaptive grids, which track features of interest during the model simulation by locally adding or removing grid points as needed. Adaptation criteria based on error estimates (e.g., [Skamarock et al. 1989](#); [Behrens 1998](#); [Blaise and St-Cyr 2012](#)) or flow characteristics (e.g., [Hubbard and Nikiforakis 2003](#); [Jablonowski et al. 2006, 2009](#); [St-Cyr et al. 2008](#)) can be used to determine where the high-resolution mesh should be placed. By increasing resolution only locally, dynamic refinement significantly decreases the total number of degrees of freedom for the simulation. However, since dynamic refinement is used within a global model, it also eliminates the need for forced boundary conditions and solves the local high-resolution area and global flow using the same dynamical core and physics package. Global models allow for a better representation of global and synoptic-scale phenomena and permit them to interact better with meso- and small-scale features that can be resolved in the model. Additionally, a key advantage of dynamic refinement compared to a static refinement setup is that the location of the refined area does not have to be determined a priori.

Adaptive refinement methods are well established in many areas of computational fluid dynamics like aerospace engineering or space weather modeling (e.g., [Tóth et al. 2012](#)). In atmospheric science though, the costs and benefits of AMR methods have mostly been evaluated in idealized simulations and simplified models so far. Some of the first adaptive atmospheric models were developed by [Skamarock et al. \(1989\)](#), [Skamarock and Klemp \(1993\)](#), and [Dietachmayer and Droegemeier \(1992\)](#). In general, the grid refinement strategies can be categorized into three overarching types: *r* refinement (or *r* adaptivity), *h* refinement, and *p* refinement.

In *r* refinement, the number of grid cells remains unchanged; instead, the cells are dynamically redistributed

to increase resolution in parts of the grid while decreasing it everywhere else. This dynamic grid adaptation creates smoother transition regions between resolutions but requires a complex global remapping of the mesh to move the location of the high resolution. [Dietachmayer and Droegemeier \(1992\)](#) use this global grid redistribution technique to increase resolution in areas where the estimated solution error is high. [Giraldo \(2000\)](#) and [Iselin et al. \(2002\)](#) have also applied this type of dynamic adaptation for the 2D shallow-water equations and advection problems, respectively. More recently, [Kühnlein et al. \(2012\)](#) implemented *r* adaptivity within a 3D Cartesian framework, [Bauer et al. \(2014\)](#) used *r*-refinement grids guided by error estimates in a shallow-water model, and [Weller et al. \(2016\)](#) demonstrated *r*-refinement use on the sphere.

Adaptive mesh refinement (AMR), another term for *h* refinement, increases resolution locally either by adding cells within the grid structure or by overlaying additional cells of finer resolution on top of the grid without changing the base grid structure. [Skamarock et al. \(1989\)](#) and [Skamarock and Klemp \(1993\)](#) implemented AMR by placing finer-resolution meshes over the coarse grid in areas which had large truncation error estimates. The gridcell solutions and boundary conditions between the higher-resolution meshes and the base grid are continually updated. In a more recent example, [Chen et al. \(2011a\)](#) use AMR that overlays high-resolution meshes in areas of interest in a shallow-water model on a cubed-sphere grid. Examples of AMR techniques that locally add and remove cells to the base grid for the shallow-water equations on the sphere have been presented by [Behrens et al. \(2005\)](#), [Läuter et al. \(2007\)](#), [St-Cyr et al. \(2008\)](#), and [Marras et al. \(2015\)](#). Both [Behrens et al. \(2005\)](#) and [Läuter et al. \(2007\)](#) use conformal unstructured finite-element meshes. In conformal grids, each cell shares an edge with exactly one other cell, while on a nonconforming grid, cells can share an edge with more than one neighboring element. The two AMR models described in [St-Cyr et al. \(2008\)](#)—a block-structured finite-volume method on a latitude–longitude grid and a spectral-element method on a cubed-sphere grid—use nonconforming meshes and a quad-tree based AMR method with gradient- or vorticity-based refinement criteria. [Marras et al. \(2015\)](#) compared the use of an AMR approach on several structured and unstructured nonconformal grids. Use of AMR in a regional model paired with a physical parameterization package was presented in [Bacon et al. \(2000\)](#). Furthermore, AMR methods have also been investigated for 2D flow fields in Cartesian geometry. Recent examples include [Müller et al. \(2013\)](#) and [Kopera and Giraldo \(2014\)](#) who analyzed a tree-structured AMR

algorithm for nonhydrostatic dynamical cores in the x - z plane, and [Hendricks et al. \(2016\)](#) who explored static and dynamic AMR for tropical cyclone-like vortices in a shallow-water model on an f plane with a constant Coriolis parameter f .

The third type of dynamic refinement, p refinement, holds the grid spacing fixed but changes the order of the polynomial approximation within each grid element to increase local resolution. Use of p refinement with discontinuous Galerkin (DG) methods for the shallow-water equations were described by [Kubatko et al. \(2009\)](#) and [Tumolo and Bonaventura \(2015\)](#). A hybrid refinement method that combines the h - and p -refinements methods was demonstrated by [Eskilsson \(2011\)](#), and [Blaise and St-Cyr \(2012\)](#) used an hp -adaptive DG method to model the shallow-water equations on the sphere for global tsunami simulations. Recently, [Aechtner et al. \(2015\)](#) implemented a new adaptive wavelet approach for local dynamic refinement with the 2D shallow-water equations on the sphere.

The purpose of this study is to demonstrate the pros and cons of using AMR, h refinement, for simulating atmospheric flows. It assesses the effectiveness of AMR, employing a nonconformal grid, in achieving similar results as uniform-grid simulations while reducing computational cost. Furthermore, it is revealed that AMR can be implemented without harming or degrading the large-scale smooth flows or inducing numerical noise and wavelike reflections at AMR boundaries. The 2D shallow-water equations serve as a useful test bed for an AMR model as they exhibit many of the complexities present in a full 3D model. We utilize the cubed-sphere fourth-order finite-volume AMR model presented in [McCorquodale et al. \(2015\)](#) for the 2D shallow-water equations on the sphere. The model implements a mapped-multiblock AMR technique that overlays refined patches on the coarser grid. Our work tests the model's ability to track and refine over dynamic small-scale features of interest and to evaluate refinement criteria. We investigate various refinement criteria, such as thresholds for the height gradient or relative vorticity, which guide the locations of refinement patches. In addition, we shed light on factors that may limit AMR applications including the size of the refinement ratios between grid levels and the total number of levels. Last, we examine the effect of AMR on large relatively smooth flows where extra refinement is unnecessary. Specifically we focus on how the coarse-fine interfaces of the AMR patches influence the overall flow and error.

The paper is organized into three main sections. A brief description of the model and a discussion of the multiblock AMR techniques are provided in [section 2](#). [Section 3](#) discusses the results of the numerical tests.

One advection test and four shallow-water tests are presented. The tests are the moving vortices advection test by [Nair and Jablonowski \(2008\)](#), a steady-state geostrophic flow [test case 2 from [Williamson et al. \(1992\)](#)], the unsteady solid-body rotation test of [Läuter et al. \(2005\)](#), a shallow-water test consisting of a gravity wave impinging on an isolated mountain, and last a test that assesses the interaction of idealized binary vortices. Conclusions are provided in [section 4](#).

2. Model description

The Chombo-AMR dynamical core (dycore) is a new model that is built upon an unstaggered high-order finite-volume (FV) multiblock approach with a classical fourth-order Runge-Kutta (RK4) time discretization scheme. A detailed description of the model setup to solve the shallow-water equations in conservative flux form can be found in [McCorquodale et al. \(2015\)](#). In addition, the Chombo-AMR library is described in [Adams et al. \(2015\)](#). The finite-volume approach is implemented on an equiangular cubed-sphere grid. This grid consists of six separate panels that are projected onto the surface of the sphere. The mesh thereby eliminates the singularities due to converging meridians at the poles found in spherical latitude-longitude grids. Additionally, the equiangular cubed sphere leads to a quasi-uniform spherical grid with grid cells of similar size across the sphere. The discrete resolution of the cubed-sphere grid is represented as $c\{N_c\}$, where N_c denotes the number of grid cells in each direction on the six panels. A list of properties of the equiangular cubed-sphere grid, including the approximate grid spacings, is given in [Table 1](#) for several resolutions. The finite-volume method for the spatial discretization uses a fourth-order accurate discretization to compute flux averages on the faces. The central difference operators used to obtain the fluxes are smoothed by an explicitly added sixth-order diffusive operator that maintains the fourth-order accuracy of the scheme [see [McCorquodale et al. \(2015\)](#) for details]. No additional limiters or filters are implemented. The numerical scheme is mass conserving to machine precision and energy conserving up to the temporal truncation order, when used without limiters or explicit dissipation. The total-energy conservation properties for the model with added dissipation are demonstrated in [Fig. 11 of McCorquodale et al. \(2015\)](#).

Since high-order FV schemes make use of neighboring elements, a mapped-multiblock approach is used to coordinate the remapping of element values that are needed for the flux calculations across panel boundaries on the cubed sphere. Though the cells at panel edges are conformal with neighboring cells across panel boundaries,

TABLE 1. Properties for several cubed-sphere grid resolutions where N_c is the number of cells along an edge of a cubed-sphere panel. Here the number of cells is the total number of grid cells ($N_c^2 \times 6$), Δx is the approximate grid spacing, A_{avg} is the average area of a grid cell, $A_{\text{min}}/A_{\text{max}}$ is the ratio between the minimum and maximum cell areas, “Eq. res.” is the grid resolution in degrees given by $90^\circ/N_c$, and $\text{RLL}_{\text{equiv}}$ is the equivalent grid spacing on a regular latitude–longitude grid with the same total number of cells.

Resolution (N_c)	No. of cells	Δx (km)	A_{avg} (km ²)	$A_{\text{min}}/A_{\text{max}}$	Eq. res.	$\text{RLL}_{\text{equiv}}$
c16	1.54×10^3	625	3.321×10^5	0.7434	5.63°	6.50°
c32	6.14×10^3	313	8.302×10^4	0.7249	2.81°	3.25°
c64	2.46×10^4	156	2.076×10^4	0.7159	1.41°	1.62°
c128	9.83×10^4	78.2	5.189×10^3	0.7115	0.70°	0.82°
c256	3.93×10^5	39.1	1.297×10^3	0.7093	0.35°	0.41°
c512	1.57×10^6	19.5	3.243×10^2	0.7082	0.18°	0.20°
c1024	6.29×10^6	9.77	8.107×10^1	0.7076	0.09°	0.10°

the transition between the panels is not smooth due to the separate mapping on each panel. To preserve the order of accuracy of the fluxes, the domain is expanded at the panel edges with the addition of three layers of ghost cells to perform the FV calculations on each panel. As a result of different mappings, ghost cells of one panel will not have the same shape as cells on the neighboring panel. Therefore, the values in the ghost cells are set by least squares interpolation from a stencil of surrounding cells that are within the domain of the ghost cell’s panel and on neighboring panels [section 3.4 in [McCorquodale et al. \(2015\)](#) describes in detail the interpolation process]. Additionally, flux values for the cell faces that lie on a panel edge are calculated separately for each panel, and the mean of the two fluxes is taken as the value for that face to ensure conservation. Thus, communication between the separate domains for each panel is limited to the fluxes at the domain boundary and the neighboring cell values needed to interpolate the solution to the ghost cell regions. The block-structured AMR method allows for further subdivision of the computational domain of each panel into rectangular regions of grid cells called patches, which allow the calculation to be distributed efficiently on parallel computing platforms.

Adaptive mesh refinement

Our 2D AMR shallow-water model uses the strategies within the Chombo library ([Adams et al. 2015](#)) that has been designed for parallel computing architectures. AMR calculations are performed on a hierarchy of nested meshes, called levels, which have a defined refinement ratio between them. This refinement ratio must be a power of 2. The finer levels are overlaid on top of the coarser levels and are organized in the block structure described in the previous section. [Figure 1a](#) provides an example of the AMR grid structure with two refinement levels. Whenever new cells are created, they are initialized via interpolations from the coarser level and ghost cells are used to calculate the fluxes at patch boundaries. At these coarse–fine interfaces, as at panel

boundaries, the space–time accuracy drops from fourth order to third order due to a lack of error cancellation that would normally occur with the FV method. Intermediate levels must have a sufficient number of cells separating a finer level and a coarser one. This ensures that the finer level is properly nested so that the interpolation to fill ghost cells on the finer level can be performed using cells from only one level.

[Figure 1b](#) depicts an overview of the time-stepping and subcycling process. Rather than having a uniform time step dependent on the smallest gridcell size, Chombo-AMR subcycles the refined levels in time maintaining a constant Courant number. As noted by [Ullrich \(2014\)](#) the single-wave-mode characteristics of a numerical method often have an unexpected and nonlinear dependence on the Courant number. Therefore, a constant Courant number helps ensure consistent dispersive properties across the refinement levels. The typical work flow for advancing an AMR grid level l in time is as follows:

- 1) Regrid levels finer than l if required:
Evaluate the refinement criterion and mark (tag) all cells that should be included in finer levels. In these regions, new blocks of cells at levels $l + 1$ are overlaid. The new cell values are interpolated from the coarser level using a fourth-order least squares algorithm that maintains conservation as described in section 4.1 of [McCorquodale et al. \(2015\)](#).
- 2) Advance level l one time step using the RK4 time-stepping method.
- 3) Interpolate values to the ghost cells that surround level $l + 1$ using the least squares algorithm also implemented for ghost cells at panel boundaries. Three layers of ghost cells are required. The interpolation does not need to be conservative as the ghost cell values are only used to reconstruct the flux on the faces of the level $l + 1$ cells. [Figure 2](#) depicts the location of ghost cells for two of the three layers and the stencil of coarse grid cells used for their interpolation.

FIG. 1. (a) Example of a grid with two levels of AMR using a $\times 2$ refinement ratio between levels. The additional levels do not replace the coarse level cells underneath, they are overlaid on top. (b) Diagram depicting the sub-cycling of AMR levels in time. The coarse level is advanced first from time t_n to time $t_n + \Delta t$, where Δt is the time step for the coarsest level. Then the finer levels are advanced with a smaller time step and periodic updates of fluxes from the coarser grids.

- 4) Perform previous steps for level $l + 1$. Level $l + 1$ is advanced using refined time steps (subcycling) as depicted in Fig. 1b. A temporal interpolation closely related to the RK4 method is used to update the values in the level $l + 1$ ghost cells from cells on level l at the intermediate time steps (McCorquodale et al. 2015).
- 5) After the subcycling in time is completed, average the solution from level $l + 1$ and sum up the fluxes to update the values on the coarse grid. Corrections are applied to the fluxes at coarse–fine interfaces to ensure conservation (Berger and Colella 1989).

As mentioned in the description of the first step, the additional levels are placed in locations that have been tagged by the refinement criterion of the model. The size of the refined grid that is added over tagged cells is determined by three aspects: 1) the need to ensure proper nesting of finer levels, 2) parameters from the Chombo library designed for efficient parallelization, and 3) a user-defined buffer parameter. Refinement (tagging) criteria are based on thresholds for user-selected flow properties, like relative vorticity, that indicate where refinement should be placed. The tagging strategies can be based on a variety of properties including tracer values, vorticity thresholds, gradients, or a combination of several criteria. The type of refinement criteria and the threshold values are set independently for each simulation. The threshold values can be uniform across all refinement levels or designed to scale with increasing refinement. For example, the relative vorticity threshold can be set so that it increases with increasing resolution. With the idealized test cases presented in this paper, the selection of refinement criteria and

thresholds was problem dependent, with the simplicity of the tests offering only a few possible options. A range of refinement thresholds for the AMR test cases was explored. Here, we present various thresholds for the gravity wave and binary vortices tests to demonstrate how changes in refinement affect the solutions and grids. For the moving-vortices advection test we only present a single threshold value that represents a compromise between reducing the error and increasing the computational run time.

The Chombo-AMR dycore is designed to have multiple refinement levels (up to 10), and the maximum number of levels is set for each simulation. In our simulations, we explore the effects of adding up to two

FIG. 2. Example of a one-level AMR approach with a refinement ratio of two on the cubed sphere grid. The four dashed cells are schematic examples of ghost cells for the finer grid, which are interpolated from the values in the coarse cells shaded in gray. Our actual AMR applications need three ghost cells.

refinement levels (called two-level AMR). In addition, we explore the grid-resolution refinement ratios between successive levels, and present selected results for three powers of 2: $\times 2$, $\times 4$, and $\times 8$. As an example, a grid with a base level of c32 (2.8°) resolution and two levels of $\times 4$ refinement has a maximum resolution of c512 (0.18°).

3. Results of the numerical experiments

For our assessment of the Chombo-AMR shallow-water model, we select five test cases: one advection test and four shallow-water tests. The test cases are divided into two categories: large-scale smooth flows and simulations with either sharp gradients or strong, nonlinear, localized flows. The first category comprises the following two shallow-water tests:

- a steady-state geostrophic flow [test case 2 in Williamson et al. (1992)], and
- an unsteady solid-body rotation [example 3 from Luter et al. (2005)].

These large-scale flows, which have no realistic use for AMR, serve as “do no harm” tests. They are used to check the model’s ability to preserve the characteristics of smooth flows as they cross the AMR patches. We measure the impact that the refinement ratios, the number of AMR levels, and the location of the refinement patches have on the solution. Convergence tests are also performed with these test cases that both have analytical solutions.

The second test category with localized flow features consists of three tests for which AMR could improve the solution, and we seek to evaluate how effectively it is able to do so. These tests are

- the moving vortices advection test by Nair and Jablonowski (2008) (with analytical solution),
- a gravity wave impinging on an idealized mountain shallow-water test, and
- a binary vortices test case in which two vortices interact.

The bottom two tests do not have analytical solutions and the evaluations rely on high-resolution reference solutions. A variety of refinement criteria are used with these test cases to demonstrate the AMR’s ability to track, adapt to, and resolve these localized features accurately. The model results are presented using normalized l_2 and l_∞ error norms as defined in Williamson et al. (1992). Additionally, the total number of grid cells quoted for AMR runs include the sum of all valid grid cells from all refinement levels (not including ghost cells) since the finer levels overlay the coarser grids beneath them. The total number of grid cells can serve

as a rough benchmark of computational cost when comparing AMR runs to uniform runs.

a. Moving-vortices advection test

The moving-vortices test is a challenging deformational-flow advection test proposed by Nair and Jablonowski (2008). The test represents the roll-up of an initially smooth tracer into tight spiral bands. The roll-up creates steep gradients that provide a good test for the AMR. In this test, a pair of vortices is generated on diametrically opposite sides of the sphere. The wind field is the summation of a solid-body rotation and a deformational flow such that the two vortices move along a great circle and an exact solution is known at all times [see Nair and Jablonowski (2008) for details]. A 12-day time period is simulated that advects the spiraling vortices once around the sphere. The background flow is prescribed with a rotation angle of $\alpha = \pi/4$ so that the two vortices are advected through the corners of the cubed sphere (located at $\pm 45^\circ$). Figures 3a–d depict the analytical solution for the roll-up of the tracer at days 0, 4, 8, and 12.

Numerical tests were carried out with uniform grids and AMR grids with two different tagging criteria: a tracer-gradient tagging and a combination tagging. In tracer-gradient tagging, the model refines in regions where $|\nabla q| > 1.5 \times 10^{-7} \text{ m}^{-1}$, where q is the unitless value of the passive tracer. Combination tagging combines the tracer-gradient tagging criterion with a relative-vorticity tagging that refines in areas where $|\zeta| > 1.45 \times 10^{-5} \text{ s}^{-1}$, so that the model refines over regions in which either threshold is reached. The gradient threshold value demonstrates a balance between reducing error and limiting computational costs. A lower threshold increases run time without significantly reducing error and a higher one results in a large increase in error. The vorticity threshold is set to maximize initial refinement around the vorticity patches without being too low as to refine on the background vorticity. With both criteria, a $\times 4$ refinement ratio between levels was used. The evolution of the grids can be seen in Fig. 3, which shows the tracer error for the two AMR runs. Figures 3e–h depict the two-level AMR run (c32/c128/c512), which has a base grid with c32 resolution, a c128 intermediate AMR level, and a finest c512 AMR level, using the tracer-gradient tagging. Figures 3i–l depict the same two-level AMR setup but with the combination tagging. The adaptive block structure is clearly successful at tracking the evolving vortices. The block structure of patches, not the individual grids cells, is outlined by the black lines in Fig. 3. Each patch contains a user-selected maximum number of grid cells. However, patches might also be split into smaller blocks by the Chombo library.

FIG. 3. (a)–(d) Analytic tracer field at days 0, 4, 8, and 12 for the moving-vortices advection test of section 3a. (e)–(h) Tracer error at the select days for a two-level AMR run with a c32 base resolution using $\times 4$ refinement (c32/c128/c512) and the tracer-gradient refinement tagging criterion. (i)–(l) As in (e)–(h), but with the combination of relative-vorticity magnitude and tracer-gradient criterion. The adaptive block structure is shown by black lines. The thickest lines are the base c32 grid, and thinner lines represent the c128 and c512 levels. Note that (e)–(h) have different contour scales than (i)–(l).

This optimizes the load balancing of the model on parallel computing architectures and increases the efficiency of the code. Therefore, various patch sizes are present in Fig. 3.

For the tracer-gradient tagging scheme (Figs. 3e–h), the error is largest near the center of the spiral. The combination tagging scheme (Figs. 3i–l) significantly reduces the error near the center, and the largest error is now toward the outer edges of the spiral just beyond the coarse–fine grid boundaries. When using only the tracer-gradient tagging, refinement does not begin until after day 1, so errors accumulate on the coarse grid until that time and remain higher even after being overlaid with finer levels. With the combination tagging, the central region is already refined (see Fig. 3i), which limits the error growth. This effect can also be seen in Fig. 4a where a time series of the normalized l_2 tracer errors with respect to the analytic solution are depicted for

15 different model configurations. In particular, Fig. 4a compares the l_2 tracer errors of the five uniform-resolution simulations c32, c64, c128, c256, and c512 (see Table 1 for the associated grid spacings) to various one- and two-level AMR experiments with either the tracer-gradient or combination tagging. The time evolution of the associated total number of grid cells, including the underlying coarse cells, is shown in Fig. 4b. In general, the error in the combination tagging simulations reaches the lower error values of the uniform runs (with the same resolution as the finest AMR level) more quickly than the tracer-gradient tagging ones, despite having a significant difference in gridcell count for only the first few days (see Fig. 4b). The errors in all single-level AMR runs in Fig. 4a using both tagging criteria converge to the error of uniform runs at the highest resolution (e.g., the c32/c128 AMR run and the uniform c128 run). This result is achieved using

FIG. 4. For the moving-vortices advection test of [section 3a](#), growth over time of (a) normalized l_2 tracer error and (b) total number of grid cells. The plots provide a comparison of uniform runs (solid lines, no markers) and AMR runs using gradient tagging (broken lines, with and without markers) and a combination of relative-vorticity and tracer-gradient tagging (solid lines with markers). All AMR runs use the $\times 4$ refinement ratio between resolution levels.

significantly fewer grid cells than the uniform runs. AMR limits the normalized l_2 error growth. AMR runs begin with global errors that are near the level of uniform runs with a grid resolution of their coarsest grid. However, their global error increases at a much slower rate than that of the uniform runs until the AMR error approaches the error level of the matching high-resolution uniform run. In the simulations presented here the two-level AMR runs see a diminishing effectiveness at decreasing the global error as the errors from the coarse section of the grid dominate.

The error at day 12 as a function of total number of grid cells for both uniform and AMR runs is plotted in [Fig. 5](#). AMR runs that have error values plotted to the left of the errors of the uniform runs (line of hollow black squares) achieve a lower l_2 error than a uniform run with a comparable number of grid cells. While the one-level AMR runs show a decrease in error while using fewer grid cells, the two-level AMR runs generally result in a higher error for a comparable number of grid cells than the uniform runs. Only the two-level AMR run with a c16 base using the combination tagging has a slightly improved error compared to the uniform runs (see the leftmost magenta star). These results demonstrate the reduced improvement to the global error from additional levels of AMR with very coarse base meshes. However, AMR still slows down the error growth over time in comparison to uniform runs, and even asymptotes to the finest uniform mesh errors (e.g., c128/c512). Extending the run time to 18 days, the l_2 error from the c32 base two-level AMR run with the combination tagging (orange triangle in [Fig. 5](#)) is now slightly lower than the error from the uniform runs at day 18 (red downward-pointing triangles). [Figure 5](#) also confirms

that the convergence rate for the uniform run is fourth order in the normalized l_2 error.

[Nair and Jablonowski \(2008\)](#) applied this test using an $\alpha = 0^\circ$ setup (not $\alpha = 45^\circ$ as we present here) in an FV-AMR model using a coarse 5° base grid and one to three levels of $\times 2$ refinement, which were guided by a

FIG. 5. Normalized l_2 gradient error as a function of total number of grid cells for the moving-vortices advection test case of [section 3a](#) at day 12. The grid-resolution labels along the bottom axis note the number of grid cells in the uniform grids with those resolutions. The black square markers are for the uniform runs c32 through c512. The gray circles and green squares represent the one-level $\times 4$ refinement AMR with base resolutions of c32, c64, and c128 using gradient tagging and combination tagging, respectively. The blue crosses and magenta stars represent the two-level $\times 4$ refinement runs with base resolutions of c16, c32, and c64 using gradient tagging and combination tagging, respectively. Finally, the red downward-pointing triangles represent the uniform c256 and c512 runs at day 18, while the orange triangle is the two-level $\times 4$ refinement runs with a c32 base resolution, using the combination tagging at day 18. Solid black lines depict convergence rates.

TABLE 2. Run times (wall-clock time in s and as % of c512 run time) for 12-day moving-vortices advection simulations (section 3a). These runs had a maximum resolution of c512, performed on two nodes of NCAR's Yellowstone computing platform with 32 processors total. Number of cells per refinement level is given at day 12 and as a percent of the c512 uniform run.

Base resolution	AMR levels	Run time (s)	Run time (%)	c32 cells	c128 cells	c512 cells	Total cells (%)
c512	—	24 152	100%	0	0	1.6×10^6	100%
c128	1	10 116	42%	0	9.8×10^4	4.5×10^5	34%
c32	2	10 468	43%	6.1×10^3	4.2×10^4	4.6×10^5	32%

tracer-gradient threshold. Their AMR errors were generally the same or lower than the errors of their comparable uniform runs for coarse grids with one or two levels of refinement. However, for their three-level AMR run (finest resolution 0.625°) errors were slightly higher than their uniform 0.625° run, agreeing with what we observe for multiple levels of AMR. Additionally, our error measures for uniform-resolution runs are comparable to other higher-order models. Our uniform c64 ($\sim 1.4^\circ$) run with $\alpha = 45^\circ$ has a normalized l_2 error of 8.9×10^{-3} after 12 days, which is comparable to the results from the 1.25° grid using a multimoment method in Chen et al. (2011b) and slightly higher than the c80 (1.125°) run in Lauritzen et al. (2010).

We also include the total run time versus number of grid cells in Table 2, for some of the 12-day runs in Fig. 5. In this case, the wall-clock time (as a % of the finest uniform run) is closely related to the total number of grid points, because most of the time steps and grid cells are at the finest level. For AMR runs, coarser levels must be completed before fine levels. There is, therefore, a slight performance penalty for the two-level AMR (c32/c128/c512) run, which actually increases the total number of finest-level points. This configuration only has 24 boxes (each 16×16 cells) at the c32 level to be distributed across 32 processors on the Yellowstone computing platform [operated by the National Center for Atmospheric Research (NCAR)]. This means that some processors run idle when the solution on the coarsest grid is computed, which slightly lessens the parallel performance of this AMR run. Overall, a good heuristic for this test is that the run time is approximately proportional to the total number of grid cells at the finest level.

b. Global steady-state geostrophic flow

This zonal steady-state flow is the second test case from the Williamson et al. (1992) test case suite for the shallow-water equations on the sphere. The test consists of a solid-body rotation along an axis that differs from the polar axis by angle α and a corresponding balanced height field. We use the more challenging $\alpha = \pi/4$ case, which means that the flow travels over the cubed-sphere corner points at a 45° angle. Since the flow is initialized

in a gradient-wind balance, any changes from the initial conditions (which serve as the analytical solution) are considered errors. No topography is present. The test was designed to measure how well the model can maintain a large-scale smooth balanced flow. Thus, we expect little benefit from AMR refinement. We use the test primarily to assess the sensitivity of the flow to the grid structure and abrupt changes in the grid resolution along coarse-fine mesh interfaces.

We implement two static refinement configurations, which can be seen in the bottom two panels of Fig. 6. The first configuration (Fig. 6b) consists of a statically refined patch centered at 0° latitude and longitude (fully contained within an equatorial cubed-sphere panel) over an area with strong gradients in the height field. In the second configuration (Fig. 6c), we place static patches over the locations of high relative vorticity, refining where $|\zeta| > 1.18 \times 10^{-5} \text{ s}^{-1}$. This criterion results in two midlatitudinal patches that transect polar-equatorial panel boundaries, a challenging location for the cubed-sphere grid. Using these static refinement configurations, we ran simulations that have one and two levels of refinement using $\times 2$, $\times 4$, and $\times 8$ refinement ratios. Increasing the refinement ratio permits us to test how abruptly resolution can increase without harming accuracy or causing spurious numerical noise at the boundary between the coarse and fine regions. We compare the height errors, characterized as the difference between the analytic initial condition and the numerical solution for the height at day 5, of uniform-resolution simulations and simulations using the equatorial and midlatitudinal patch configurations. The addition of refined patches should ideally result in no additional error if the flow is well resolved, or a small decrease in global height error if it is not.

Results in Table 3 show the normalized l_2 and l_∞ height errors for the uniform c32 (2.8°) run and c32 base runs with the two refinement configurations after 5 days. The errors for the runs with the equatorial patch are essentially unchanged from that of the uniform c32 run. Even the extreme cases of a $\times 8$ refinement ratio or multiple levels of refinement increase the l_2 height error by less than 0.25%. The results for midlatitudinal patch runs show a reduction in global error with even a $\times 8$

FIG. 6. Height error (in m) at day 5 for the steady-state geostrophic flow test case of section 3b. The configurations are as follows: (a) a c32 uniform-resolution run, (b) a c32 grid with a static equatorial patch using two levels of $\times 4$ refinement, and (c) a c32 grid with static midlatitudinal patches using two levels of $\times 4$ refinement tagging on the relative-vorticity extrema. The solid black contour lines in (a) represent the height field with a contour spacing of 200 m and a value range between 1200 and 2800 m with the minima encircled by the closed contours. The dotted and dashed contour lines correspond to the negative and positive values, respectively, of the height error tick marks in the label bar. The zero line is the dot-dashed line.

refinement ratio reducing the l_2 height error by more than 10% and the l_∞ error by at least 25% compared to the c32 uniform run. The height error plots at day 5 for the uniform c32 run and the runs for the two refinement configuration using two levels of $\times 4$ additional refinement (Figs. 6a–c) depict a similar result. The equatorial patch run (Fig. 6b) has essentially the same height error profile as the uniform run (Fig. 6a). The height error for a base c32 run with the midlatitudinal refinement patches (Fig. 6c) shows a clear improvement in error since the coarse base resolution does not fully resolve the flow and the refined patches cover areas of high

error. The refined patches do not create any spurious wave reflections or lead to an increase in error along the coarse–fine boundary.

Figure 7 depicts a comparison of the day-5 normalized l_2 and l_∞ height errors for runs with the midlatitudinal refinement patches and uniform runs for base resolutions of c16 (5.6°) to c256 (0.35°). At coarse resolutions, we see a slight improvement in the error for runs with refinement compared to the uniform runs. However, at higher resolutions when the flow is well resolved the change in error is indistinguishable. The figure also shows that fourth-order convergence is maintained in the runs with the static refinement patches. Results for runs with the equatorial refinement patch at other higher resolutions followed a similar pattern as the c32 base runs in Table 3 and also demonstrated fourth-order convergence (not shown).

The steady-state geostrophic flow test case has been used in other AMR and static refinement studies. Similar refined grid locations were used with the finite-volume AMR models in Chen et al. (2011a) (on the cubed sphere) and St-Cyr et al. (2008) (on a latitude–longitude grid). In both models, the introduction of refined patches led to increases in error when compared with the uniform runs, with significantly larger increases in the height error for configurations in which the refinement patch was placed over strong height gradients. The error increased by $\sim 35\%$ in Chen et al. (2011a) and a factor of 2.5 in St-Cyr et al. (2008). However, for the higher-order spectral-element method (SEM) in St-Cyr et al. (2008), the error was considerably reduced with the addition of a refined patch. Weller et al. (2009) tested a number of grid geometries with variable resolutions using a $\times 2$ refinement ratio and found increases in error when a refinement patch was added. Harris and Lin (2013) used a nested-grid FV model with a $\times 3$ refinement ratio. After 5 days their l_2 height errors roughly doubled in comparison to their uniform run, though their l_∞ errors were nearly unchanged. Our results with static refinements are very competitive as they show almost no increase in error or even in some cases an improvement in the error. Thus, our model preserves the large-scale flow and limits the errors at the refinement patches very effectively.

c. Unsteady solid-body rotation

The time-dependent zonal flow test proposed in Luter et al. (2005) (example 3) consists of an unsteady solid-body rotation (USBR), which is forced by topography. It possesses an analytical solution. The large-scale flow and topography are smooth, zonally symmetric, and somewhat artificial. In particular, the topography is zero at the equator and rises to its maximum (around 11 km) at both poles. As with the previous test case, we expect little benefit from AMR given the smooth characteristics. The

TABLE 3. Global steady-state geostrophic flow test of [section 3b](#): normalized l_2 and l_∞ height errors at day 5 for a variety of refinement ratios and numbers of levels with the two refinement locations near the equator and in the midlatitudes. As a comparison, the normalized height errors of a uniform-resolution c32 run at day 5 are $l_2 = 5.4752 \times 10^{-6}$ and $l_\infty = 1.4505 \times 10^{-5}$.

AMR levels	Refinement ratio	Equatorial refinement		Midlatitudinal refinement	
		l_2 error	l_∞ error	l_2 error	l_∞ error
1	$\times 2$	5.4848×10^{-6}	1.4519×10^{-5}	4.7472×10^{-6}	1.0077×10^{-5}
1	$\times 4$	5.4837×10^{-6}	1.4542×10^{-5}	4.7835×10^{-6}	9.5043×10^{-6}
1	$\times 8$	5.4875×10^{-6}	1.4592×10^{-5}	4.8233×10^{-6}	9.6988×10^{-6}
2	$\times 2$	5.4849×10^{-6}	1.4520×10^{-5}	4.5341×10^{-6}	1.0559×10^{-5}
2	$\times 4$	5.4836×10^{-6}	1.4542×10^{-5}	4.7983×10^{-6}	1.1019×10^{-5}

benefit of the USBR test is that it has the added complication of moving features that can be tracked with AMR while still having an analytic solution to determine the error. One can observe how well the flow is maintained and whether numerical artifacts, if any, are created by the resolution change at grid boundaries and by the AMR regridding process. Using the setup described in [Läuter et al. \(2005\)](#), we have set the parameter $\alpha = \pi/4$ to let the flow field pass over the corners of the cubed sphere at a 45° angle. To force the initial condition to repeat itself after exactly 1 day for better comparison of the results, Earth's angular velocity is slightly modified to be based on a solar day instead of sidereal day, so that the angular velocity is $\Omega = 2\pi/86\,400 \text{ s}^{-1} \approx 7.2722 \times 10^{-5} \text{ s}^{-1}$.

We conduct a series of simulations over a range of base resolutions with either one or two levels of refinement using three refinement configurations:

- 1) Statically refined patch used in [section 3b](#) now centered at $0^\circ, 90^\circ\text{E}$.
- 2) Dynamic AMR refinement with height tag, the threshold for the free surface height is $h > 1.38 \times 10^4 \text{ m}$.
- 3) Dynamic AMR refinement with relative vorticity tag, the threshold is $|\zeta| > 1.18 \times 10^{-5} \text{ s}^{-1}$.

The three grid configurations can be seen in [Figs. 8b–d](#), which show the patches at the identical initial and final (day 5) positions. The second and third configurations ([Figs. 8c,d](#)) provide for the AMR tracking of a moving feature, so the effects of a moving mesh and regridding can be observed. The vorticity tag provides a more challenging test since it bisects a cubed-sphere panel edge. We see little deviation in error among the $\times 2$, $\times 4$, and $\times 8$ refinement ratio simulations, so we only discuss runs with the $\times 4$ refinement ratio. Whenever the dynamic AMR grids are moved, the underlying topography is reinitialized with the analytical formulation given in [Läuter et al. \(2005\)](#).

Simulations were run for 5 days. The normalized global l_2 and l_∞ height errors after 5 days are shown for simulations with a c32 base grid in [Table 4](#). Errors are calculated by comparing runs to the analytic solution of the test case. The uniform grid results are compared with one- and two-level refinement runs using the three grid configurations. For the static refinement simulations, the l_2 height error increases by approximately 0.75% in comparison to the uniform c32 run. In the two-level height-tag AMR run, we observe that the l_2 height error increases by about 12%, while the two-level vorticity-tag

FIG. 7. (a) Normalized l_2 and (b) l_∞ height errors at day 5 as a function of base grid resolution for the steady-state geostrophic flow test case of [section 3b](#). Uniform runs and runs using the static midlatitudinal refinement patches are depicted with $\times 2$, $\times 4$, or $\times 8$ refinement ratios. The fourth-order convergence rate is shown by the black line.

FIG. 8. Height field errors (in m) at day 5 for four simulations of the unsteady solid-body rotation test case of section 3c: (a) a c32 uniform-resolution run, (b) a c32 base grid with a static equatorial patch using two levels of $\times 4$ refinement, (c) a c32 base grid with two levels of dynamic $\times 4$ refinement using the height-tag criterion, and (d) a c32 base grid with two levels of dynamic $\times 4$ refinement tagging on the vorticity-tag criterion. The solid black contour lines in (a) represent the free surface height field of the USBR test (above sea level) with a 150-m contour spacing and a value range between 1.22×10^4 and 1.37×10^4 m with the minima in the polar regions. The dotted and dashed contour lines correspond to the negative and positive values, respectively, of the height error tick marks in the label bar. The zero line is the dot-dashed line.

AMR run decreases the error by roughly 23%. Figure 8 depicts the USBR height errors at day 5 for the c32 uniform run and c32 base grid runs with the three grid configurations. For the static equatorial patch run (Fig. 8b), the height errors remain nearly the same as for the uniform run (Fig. 8a), with only very slight increases in the large error areas on the polar panels. Along the coarse-fine boundary, no spurious grid-induced error is observed. In the height-tag AMR run (Fig. 8c) the errors along the polar-equatorial panel boundaries increase, while in the vorticity-tag AMR run (Fig. 8d), we observe that the large errors on the polar panels are reduced due to the addition of refinement over that area. With the height tagging, we do not see a similar improvement because the refined patches are over the low error areas on the equatorial panel.

After 5 days, the uniform c32 run (~ 313 -km grid) had a normalized l_2 height error of 2.6×10^{-6} . For comparison, the second-order icosahedral model by Düben et al. (2012) obtained a normalized l_2 height error of 1.05×10^{-4} at day 5 using a uniform grid with an average edge length of 240 km. Other investigations focused on results at 12 h, after which the flow features have progressed only halfway around the sphere. Pudykiewicz (2011) showed a normalized l_2 height error

of $\sim 6 \times 10^{-6}$ after 12 h using a second-order icosahedral geodesic model with a 2° (~ 220 km) grid resolution, while our uniform c32 (2.8°) run produced a normalized l_2 height error of 3.5×10^{-7} at 12 h. These results are comparable with those obtained by the fourth-order multimoment model on a Yin-Yang grid with an effective resolution of 1.875° in Li et al. (2015). They reported a normalized l_2 height error of $\sim 3 \times 10^{-7}$ after 12 h. Additionally, the third-order multimoment method on a cubed-sphere grid with a $N = 40$ (2.25°) grid resolution in Chen et al. (2014) had a normalized l_2 height error

TABLE 4. Day-5 normalized l_2 and l_∞ height error norms for the unsteady solid-body rotation test of section 3c. The c32 uniform-resolution run is compared to the static refinement runs and AMR runs tagging on relative vorticity and height with one and two refinement levels using the $\times 4$ refinement ratio.

Grid configuration	No. of levels	l_2 error	l_∞ error
Uniform	—	2.6052×10^{-6}	7.8216×10^{-6}
Static	1	2.6064×10^{-6}	8.1955×10^{-6}
Static	2	2.6064×10^{-6}	8.1977×10^{-6}
AMR, vorticity	1	2.1034×10^{-6}	6.1509×10^{-6}
AMR, vorticity	2	2.0017×10^{-6}	6.8035×10^{-6}
AMR, height	1	2.8837×10^{-6}	8.7651×10^{-6}
AMR, height	2	2.9205×10^{-6}	1.0255×10^{-5}

FIG. 9. (a) Normalized l_2 and (b) l_∞ height errors at day 5 as a function of base grid resolution for the USBR test case of section 3c. Uniform runs and AMR runs with one and two refinement levels using the height-tag or vorticity-tag criteria. All AMR runs use $\times 4$ refinement ratio between levels. Solid black lines depict convergence rates. Errors are determined with respect to the analytic solution.

of $\sim 6.5 \times 10^{-7}$ at 12 h. We are unaware of other results that use the USBR test with AMR applications.

We performed the USBR test in simulations with increasing base resolutions of up to c128 (~ 78 -km spacing) with the two dynamic grid configurations. The normalized l_2 and l_∞ height errors after 5 days are plotted in Figs. 9a and 9b, respectively. At higher base resolutions the slight improvements in the l_2 height error no longer occur for the vorticity-tag AMR as the large-scale flow features are well resolved (Fig. 9a). The fourth-order convergence is maintained for all grid configurations. In the l_∞ height error plot (Fig. 9b), we observe the same slight decrease in error for the coarse c16 and c32 base resolutions for the vorticity-tag AMR and the slight increase in the height-tag AMR errors as observed earlier in Figs. 8c and 8d. However, at higher base resolutions we find a large increase in the l_∞ error for the vorticity-tag AMR runs. While fourth-order convergence is maintained at all resolutions for the uniform, static refinement, and height-tag AMR configurations across all resolutions, the l_∞ convergence rate drops to between 3 and 2.5 for the vorticity-tag AMR runs at higher base resolutions. The increased error is due to the regridding of the AMR patches. The maximum errors occur in cells bordering the coarse-fine boundary of the AMR patch and the base grid when that boundary intersects an edge of the cubed sphere. This point-source-like artifact of the AMR grid occurs in both the height-tag and vorticity-tag AMR simulations. In the height-tag runs, the artifact is triggered only when the AMR grid passes over the corners of the cubed sphere, resulting in the slight l_∞ error increase observed in Fig. 8c. In the vorticity-tag AMR runs, the refined grid bisects the polar-equatorial panel edge during the entire run, thus this small error is compounded at each regridding step. This results in the sharp increase in the l_∞ error seen in Fig. 9b for the vorticity-tag runs with

c64 and c128 base resolutions. Overall, though, the error is small and very localized at the cells where the AMR patches intersect the cubed-sphere edge. It is, therefore, only obvious in the strict l_∞ error measure and only at high horizontal base resolutions. At lower base resolutions the magnitudes of other errors are bigger which then dominate the global l_∞ error measure.

d. Isolated mountain gravity wave

This shallow-water test was developed to assess AMR when topography is present. In the test a gravity wave, which is triggered by an unbalanced initial height perturbation in a quiescent background environment, passes over an idealized mountain. The change in topography deforms the structure of the gravity wave. The bottom topography z_s consists of a cosine mountain and is defined by

$$z_s = \frac{z_0}{4} \left[1 + \cos\left(\frac{\pi r}{R}\right) \right]^2, \quad (1)$$

where $R = \pi/9$ and $r^2 = \min[R^2, (\lambda - \lambda_c)^2 + (\theta - \theta_c)^2]$. Outside the radius R the topography is set to zero. The peak height of the mountain is $z_0 = 2000$ m, and it is centered at $(\lambda_c, \theta_c) = (3\pi/2, \pi/6)$ in the longitudinal and latitudinal direction, respectively. The initial velocity field is set to zero and the initial free surface height has a constant background value of $h_0 = 5960$ m with a local Gaussian dip perturbation. Thus, the initial free surface height field (above the reference sphere at sea level) is given as

$$h = h_0 - h_{\max} \exp\left[-\left(\frac{\tau}{\beta}\right)^2\right]. \quad (2)$$

The maximum depth of the perturbation is set to $h_{\max} = 100$ m, $\beta = a \times \pi/36$ is a width parameter, and τ is

FIG. 10. Initial free surface height field (colored, in m) for the gravity wave over an idealized mountain test of section 3d. The free surface height field (above the reference sphere at sea level) is uniform everywhere except for the 100-m-deep Gaussian depression. The black contour lines represent the location of the mountain with 200-m contour spacing and a peak mountain height of 2000 m.

the great-circle distance from point (λ, θ) to the dip's center (λ_d, θ_d) such that

$$\tau = a \arccos[\sin\theta_d \sin\theta + \cos\theta_d \cos\theta \cos(\lambda - \lambda_d)], \quad (3)$$

where $a = 6.37122 \times 10^6$ m is the average radius of Earth. The Gaussian dip is centered at $(\lambda_d, \theta_d) = (3\pi/2 - \pi/5, \pi/6)$ in the longitudinal and latitudinal direction, respectively. Figure 10 depicts the initial height field and its distance from the mountain.

The simulation is run for a period of 12 h so that the gravity wave has moved halfway around the sphere. Figures 11 and 12 show the perturbation height (defined as deviations from h_0 , top panels) and the height difference from a uniform c1024 (~ 10 km) reference solution (bottom panels) at hours 6 and 12, respectively, for a uniform c128 run and a c32 base one-level AMR run (c32/c128) tagged on a height-gradient threshold of $|\nabla h| > 7.5 \times 10^{-6}$. After 6 h (Fig. 11), the gravity wave has just passed over the mountain and the distortion to the wave from the mountain is clearly visible. The presence of the mountain breaks the symmetry of the circular, outward-propagating gravity wave. This propagation is captured by the AMR refinement criterion as indicated by the overlaid block structure in Figs. 11b and 11d. The height differences for the c32/c128 AMR run (Figs. 11d and 12d) are similar in position and magnitude to the uniform c128 run errors (Fig. 11c) within the refined AMR domain. The areas of larger error at the borders of the AMR region seen in Figs. 11d and 12d are located over the leading and trailing edges of the gravity wave, which are not fully covered by the AMR refinement criterion. The location and magnitude of these larger errors correlate with the error at the leading edge of the gravity wave observed in the uniform c32 run. At hour 6, the AMR refined grids are over the mountain and by hour 12, the mountain is once

FIG. 11. Mountain gravity wave test of section 3d, at hour 6. (a),(b) The perturbation height of the gravity wave as it passes over the mountain for a uniform c128 run and a c32/c128 AMR run with the height gradient tag $|\nabla h| > 7.5 \times 10^{-6}$. (c),(d) The height error of each run after 6 h compared to a reference uniform c1024 run. The block structure of the grid and the mountain contours are overlaid with thin black lines.

FIG. 12. As in Fig. 11, but for hour 12 as the gravity wave has moved halfway around the sphere.

again covered by only the coarse grid. The AMR re-initializes the topography [using Eq. (1)] whenever adaptations are triggered. No spurious errors appear as the AMR refines and coarsens over the topography region (Figs. 11d and 12d).

Results in Fig. 13 show the normalized l_2 height error (Fig. 13a) and the total number of grid cells (Fig. 13b) as a function of forecast hour for uniform runs from c32 to c512 and one-level AMR runs using height-gradient tagging with thresholds of $|\nabla h| > 1.5 \times 10^{-5}$, 1.0×10^{-5} , and 7.5×10^{-6} , labeled T1, T2, and T3, respectively. The

normalized error metrics are determined with respect to the uniform c1024 simulation, which serves as the reference solution. For uniform runs, the solution error converges to fourth order in both the normalized l_2 and l_∞ height error norms. The AMR runs have improved error but do not reach the error of the uniform run with the same resolution as the highest refinement level. The c32 base one-level AMR T3 run (c32/c128) has a maximum number of grid cells roughly equivalent to the uniform c64 run, but its error is nearly an order of magnitude smaller than the uniform c64 run. Reducing the gradient

FIG. 13. For the mountain gravity wave test of section 3d, growth over the 12-h period of (a) normalized l_2 height error with respect to the uniform c1024 run and (b) total number of grid cells. Error and number of grid cells for uniform-resolution runs and one- and two-level AMR runs with height-gradient tagging. The thresholds are $|\nabla h| = 1.5 \times 10^{-5}$ (T1), 1.0×10^{-5} (T2), and 7.5×10^{-6} (T3).

threshold in AMR runs so that more area around the gravity wave is covered by the AMR patches improves the solution and reduces error. A c32/c128 AMR run with a refinement threshold of $|\nabla h| > 4.5 \times 10^{-6}$, lower than the T3 criterion, results in the AMR grid covering 31% of the globe by hour 12 and a normalized l_2 height error of 3.957×10^{-6} . In comparison, the uniform c128 run has an error of 3.051×10^{-6} and the T3 run has an error of 5.806×10^{-6} with AMR blocks covering only 25.8% of the area. As more of the leading edge of the gravity wave is refined with lower tagging thresholds, the error is decreased further, though at a diminishing rate.

e. Binary–vortices interaction

The binary–vortices interaction test demonstrates the AMR benefits and its effectiveness in studying an important and more realistic problem. The interaction of two neighboring tropical cyclones (TCs) often alters the structures of the two, leads to complex tracks for the storms, and in some instances results in a merger of the two cyclones. These interactions were first studied by Fujiwhara (Fujiwhara 1921) and are commonly called the Fujiwhara effect. Idealized binary–vortex interactions have been extensively investigated using 2D idealized models by Melander et al. (1988), Waugh (1992), Ritchie and Holland (1993), Prieto et al. (2003), and Shin et al. (2006). A majority of the research has been conducted on two-dimensional Cartesian systems using a constant Coriolis force. These studies have used a variety of initial vortex profiles featuring discrete (Ritchie and Holland 1993) or continuous (Bauer et al. 2014) vortices in both symmetric and asymmetric pairs (Dritschel and Waugh 1992). They demonstrate that slight changes in initial conditions will cause widely diverging results. The vortices will either merge or repel each other depending on the strength, size, and separation distances of the vortices, and the postinteraction shapes of the vortices will be vastly different. Bauer et al. (2014) used an r -adaptive shallow-water model to demonstrate that the vortices' tracks are sensitive to initial conditions and to initial grid resolution. Given the sensitivity to resolution, binary–vortex interactions are a well-suited test of AMR. The steep gradients, localized areas of high vorticity, and complex flow fields around the vortices are transient and resolution-dependent, mimicking the multiscale nature of tropical cyclones. With this test, we can evaluate the AMR's ability to refine and track these features of interest and measure the AMR's accuracy in resolving the vortex interaction. It can assess how well the results and errors in AMR runs match the results of uniform high-resolution runs and can determine the sensitivity of the vortex tracks to the changing grid resolutions.

In our binary tropical cyclone–like vortices test, we use the full shallow-water equations on a spherical grid with a changing Coriolis parameter, whereas most other published studies use a nondivergent barotropic model on an f plane. We also restrict our study to only the symmetric case so that the two vortices are identical in size and strength. The two vortices are initialized near each other and are allowed to interact over a simulation period of several days. Two variations of this setup are presented. In the separation case, the two vortices orbit around each other and then slowly drift apart. In the merger case, the vortices merge. Our initializations of the continuous vortex profiles were inspired by the definition of the initial state in Holland and Dietachmayer (1993). The initial wind and height profiles are derived from the shallow-water equations in cylindrical coordinates using an f -plane approximation. The vortex structure is depicted as a radial perturbation in the geopotential field and is given by

$$\phi = \bar{\phi} - \phi', \quad (4)$$

$$\phi' = \phi_c \left\{ 1 - \exp \left[- \left(\frac{r_m}{r} \right)^b \right] \right\}, \quad (5)$$

where $\bar{\phi} = g\bar{h}$ is the background geopotential with the background height $\bar{h} = 4200$ m and Earth's gravity $g = 9.80616 \text{ m s}^{-2}$, ϕ' denotes the geopotential perturbation, ϕ_c symbolizes the maximum geopotential perturbation, r_m is the radius of maximum wind, b stands for a scaling parameter set to 1.5, and r is the great-circle distance from point (θ, λ) to the vortex center (θ_c, λ_c) [see also Eq. (3)]. The values of ϕ_c , r_m , θ_c , and λ_c are provided later.

A balanced tangential wind field is then found by using the steady-state shallow-water momentum equations in cylindrical coordinates. In particular, the tangential wind:

$$v_T = -\frac{rf}{2} \pm \sqrt{\frac{r^2 f^2}{4} + r \frac{\partial \phi}{\partial r}} \quad (6)$$

represents the initial axisymmetric flow in gradient wind balance. Using $\partial \phi / \partial r$ derived from Eq. (4), we get the corresponding tangential velocity for a cyclonic [the plus sign in Eq. (6)] circulation:

$$v_T = \sqrt{\phi_c b \left(\frac{r_m}{r} \right)^b \exp \left[- \left(\frac{r_m}{r} \right)^b \right] + \frac{r^2 f^2}{4}} - \frac{rf}{2}, \quad (7)$$

where f is the constant Coriolis parameter for an f -plane approximation at the latitude of the vortex center (specified later). The last initialization step is to project

the tangential velocity onto the sphere with the zonal u and meridional v spherical wind components given by

$$u = v_T \frac{d_1}{d} \quad \text{and} \quad v = v_T \frac{d_2}{d}, \quad (8)$$

where

$$d_1 = \sin\theta_c \cos\theta - \cos\theta_c \sin\theta \cos(\lambda - \lambda_c), \quad (9)$$

$$d_2 = \cos\theta_c \sin(\lambda - \lambda_c), \quad (10)$$

$$d = \max(\varepsilon, \sqrt{d_1^2 + d_2^2}). \quad (11)$$

The threshold value $\varepsilon = 10^{-25}$ prevents division by zero. The topography is set to zero.

This initialization technique represents a perfect balance for a single vortex in cylindrical coordinates and leads to a very good balance in spherical coordinates. Note that the perfect balance is broken on the sphere since f varies in the spherical domain and is held constant for the purpose of the initialization. In addition, an analytically derived balance is not fully balanced in a numerical (discrete) sense and the overlap region of two vortices is not strictly balanced either. However, the initial imbalances for our separation and merger test cases are very minor and the resulting small gravity waves do not interfere with our AMR analysis. The same initialization technique was also used for idealized tropical cyclone simulations in 3D GCMs (Reed and Jablonowski 2011).

The radial cross sections of the initial relative vorticity, height field, and tangential wind for a single vortex as a function of the great-circle distance from the center are depicted in Fig. 14a. Here, all magnitudes are normalized to 1 and are provided below for each test case. The initial relative vorticity for the merger test case with two initial vortices is depicted in Fig. 14b. It shows that the tangential wind profile from Eq. (7) results in a relative vorticity profile with a core of positive relative vorticity in the center of each storm surrounded by a broad ring of negative vorticity with relatively small magnitudes. The two vortices slightly overlap with very minor magnitudes of u , v , and ϕ' [Eq. (5)]. Here, we use the sum of the u , v , and ϕ' values of both vortices to initialize our shallow-water system.

1) SEPARATION CASE

In the separation case, the two vortices are centered at $\theta_c = \pi/36 = 5^\circ\text{N}$, which defines the constant Coriolis parameter $f = 2\Omega \sin\theta_c$ with Earth's angular velocity $\Omega = 7.292 \times 10^{-5} \text{ s}^{-1}$. The maximum geopotential perturbation is set to $\phi_c = gh_c$ with the maximum height perturbation $h_c = 800 \text{ m}$. The radius of maximum wind is

FIG. 14. Initial conditions for the binary vortices test case of section 3e. (a) Radial profiles of relative vorticity (red), tangential wind (blue), and height field (black) as a function of great-circle distance from the center for a single vortex. Profiles are scaled to the maximum of each value (see text). (b) Initial profile of relative vorticity (s^{-1}) of the two vortices on a cubed-sphere grid (merger test case).

set to $r_m = 250 \text{ km}$. This results in a maximum tangential wind of 64 ms^{-1} and a maximum relative vorticity of $9.4 \times 10^{-4} \text{ s}^{-1}$. The centers of the two vortices are 13.5° apart ($\sim 1500 \text{ km}$) in the longitudinal direction, 6 times the radius of maximum wind, so that their negative vorticity regions still overlap. In particular, the two vortex centers are located at $\lambda_{c1} = (3\pi/2 - 6.75\pi/180)$ and $\lambda_{c2} = (3\pi/2 + 6.75\pi/180)$ with the midway point between the two cyclones at 90°W .

This scenario is sensitive to variations in initial conditions, making it desirable for testing adaptive grids. Decreasing the separation distance by 20 km results in the merger of the two vortices. During the first 3 days of the simulation, the two vortices make one complete orbit around each other as the beta-drift steers them toward the northwest, after which the two vortices then

FIG. 15. Evolution of the relative vorticity of uniform-resolution runs for the binary vortices test (separation case) of [section 3e\(1\)](#). (a)–(e) Results for the uniform c1024 run for days 1, 2, 3, 4, and 6. (f)–(j) The uniform c256 run and (k)–(o) the uniform c128 run.

drift apart. In that time, the negative vorticity is stretched as it is advected around the pair of positive cores before being spun off behind the pair as an anticyclone. We also observed the growth of a large-scale wave train that forms in the lee of the orbiting pair. The time evolution of the flow

can be seen in [Fig. 15](#), which shows the vorticity field of the cyclone pair at days 1, 2, 3, 4, and 6 for several uniform-resolution runs. These serve as references for the AMR simulations. We ran uniform runs with resolutions from c32 to c1024, of which the c128, c256, and c1024 runs are

FIG. 16. Relative vorticity fields at day 6 for several runs of the vortex separation case of [section 3e\(1\)](#) using AMR: (a) uniform c256 run, (e) uniform c512 run, and (i) uniform c1024 run. (b)–(d) AMR runs whose highest refinement level is c256. (f)–(h) and (j)–(l) AMR runs with a highest refinement level of c512 and c1024, respectively. (b), (f), (j) AMR runs with one level of refinement using a tagging criterion based on a relative vorticity threshold of $|\zeta| > 2.3 \times 10^{-5} \text{ s}^{-1}$. (c), (g), (k) AMR runs using two levels of refinement using a relative-vorticity threshold of $|\zeta| > 3.5 \times 10^{-5} \text{ s}^{-1}$. (d), (h), (l) AMR runs using the same refinement criteria as in (b), (f), (j), but with two levels of refinement. All AMR runs use the $\times 4$ refinement ratio. The block structures of refinement levels 1 and 2 are outlined in black.

depicted in [Fig. 15](#). Results vary significantly with resolution, though results do converge with increasing resolution. Runs with coarser resolution than c128 (not shown) have very weak vortices that merge instead of drifting apart. In the c128 run ([Figs. 15k–o](#)), the vortices start separating earlier and at the end of the run are in markedly different positions. The c256 simulation ([Figs. 15f–j](#)) more closely resembles the highest-resolution c1024 run ([Figs. 15a–e](#)), but there are still significant differences. However, the c512 run (not shown as a time series) is nearly indistinguishable

from c1024 with only slight differences in the center of the vortex cores and in the finescale vorticity filaments. A comparison of the uniform c512 run and c1024 run at day 6 can be seen in [Figs. 16e and 16i](#).

To assess the AMR performance, we ran the model using relative vorticity refinement criteria with one and two levels of AMR and $\times 4$ refinement on base grid resolutions from c16 to c256. Samples of the resulting relative vorticity fields at day 6 using two different relative vorticity thresholds of $|\zeta| > 3.5 \times 10^{-5} \text{ s}^{-1}$ and $|\zeta| > 2.3 \times 10^{-5} \text{ s}^{-1}$

are displayed in Fig. 16. They are divided into columns that share the same highest resolutions; for example, in the leftmost column, Fig. 16a is the uniform c256 run, while Figs. 16b–d are for one- or two-level AMR runs that have a finest grid resolution of c256. The AMR runs agree well with the uniform solutions having the same resolution as the finest AMR level. The AMR blocks accurately capture the positions of the two vortices and the shape of their high-vorticity cores. The AMR runs also effectively reproduce the overall shapes of the anticyclonic filaments and patches around the cores, and the wave train developing to the lee of the vortices with only minor differences in the anticyclone filament between the two vortices in a few runs. Given the nonlinear sensitivity to initial conditions and grid resolution, unrefined areas and the coarse–fine boundaries near the vortices in AMR runs may cause divergent solutions. In none of our AMR simulations did this occur. The results of the AMR runs differed only slightly from the uniform-resolution runs.

2) MERGING VORTICES

In the merging-vortices case, the two vortices have the same maximum height perturbation with $h_c = 800$ m, but the radius of maximum wind is increased to $r_m = 400$ km so that the maximum tangential wind is 61 m s^{-1} and the maximum relative vorticity is around $5.7 \times 10^{-4} \text{ s}^{-1}$. The vortices are centered at $\theta_c = \pi/18 = 10^\circ\text{N}$ where the constant Coriolis parameter f is evaluated for this test case. The vortex centers are 15.65° apart (~ 1700 km) in the longitudinal direction. In particular, they are located at $\lambda_{c_1} = (3\pi/2 - 7.825\pi/180)$ and $\lambda_{c_2} = (3\pi/2 + 7.825\pi/180)$ with the same midway point as before at 90°W . This separation distance is about 4.3 times the radius of maximum wind, so that the edges of the positive vorticity cores slightly overlap.

Figure 17 depicts the evolution of these vortices as they merge over the course of four simulation days. As in the previous case, though small changes in initial condition lead to very different results, we do observe a slow convergence with increasing resolution. We ran the test case with several configurations using one and two levels of AMR with criteria based on a relative vorticity or a height gradient threshold. The AMR run shown in Fig. 17 has a c64 base resolution with two levels of $\times 4$ refinement so that its finest level has a c1024 resolution. The AMR is triggered when the absolute value of the relative vorticity exceeds $2.3 \times 10^{-5} \text{ s}^{-1}$. With that relatively low refinement threshold, the AMR captures not only the main vortex cores, but also the small-scale anticyclonic filament that extends far south of the merged vortex, and the small-magnitude wave train that develops by day 4. Figure 18 shows a column comparison of the vorticity field at day 4 for uniform-resolution runs and

AMR runs using three refinement criteria. The first column contains the uniform c1024 run and 3 two-level AMR runs with $\times 4$ refinement on c64 base resolution that have a finest resolution of c1024. The second column has the c512 uniform run and 3 two-level AMR runs with a c32 base. The last column has the uniform c256 run and 3 two-level AMR runs with a c16 base. The three AMR refinement criteria in Fig. 18 are as follows:

- 1) Large-height-gradient-tag AMR: tag where the absolute value of the height gradient is $|\nabla h| > 4 \times 10^{-4}$ (second row of Fig. 18);
- 2) Small-height-gradient-tag AMR: tag where $|\nabla h| > 1 \times 10^{-4}$ (third row);
- 3) Vorticity-tag AMR: tag where the absolute value of the relative vorticity is $|\zeta| > 3.5 \times 10^{-5} \text{ s}^{-1}$ (fourth row).

Locally refining the grid resolution with AMR effectively achieves a similar result in the refined areas as the corresponding high-resolution uniform runs. Even the large-height-gradient refinement threshold used in Figs. 18b and 18f, which results in very little refinement, is still able to produce a very similar vortex structure and position within the refined area demonstrating little to no negative effects from the coarse–fine boundaries surrounding the vortex. The lower refinement thresholds are further able to capture the anticyclonic filament wrapping around the new vortex and extending down from it as well as the development of the secondary leeside wave train.

The convergence to the error of the uniform-resolution runs can be observed in the normalized l_2 vorticity error seen in Fig. 19. The l_2 error norm is computed by comparing runs to the uniform c1024 run, which serves as a reference solution. The figure depicts the normalized l_2 relative vorticity error (Fig. 19a) and total number of grid cells (Fig. 19b) as a function of forecast days for uniform-resolution runs and one- and two-level AMR runs using the large-height-gradient tag or the vorticity tag with $\times 4$ refinement ratios. The vorticity-tag AMR simulations (both one- and two-level AMR) have nearly the same error as the uniform runs with the highest resolution while the gradient-tag runs have slightly higher error. This agrees with the fact that the gradient-tag runs use fewer grid cells and only cover the merged vortex core (see Figs. 18b,f,j). Although the large-scale shape and locations of the two merging vortices and the postmerger vortex appear visually to converge to a solution with increasing resolution, we do not observe a large reduction in the global errors with increasing resolution. The source of this error is the small differences that occur in the core of the vortices, caused by small-scale nonlinear features in the high-vorticity filaments, as well as slight variations in the beta drift created by small changes in the vorticity magnitude for the different resolution runs. These small

FIG. 17. Evolution of the merging vortices in the test of [section 3e\(2\)](#) in a two-level AMR run (c64/c256/c1024) using a relative-vorticity threshold refinement criterion. Refinement occurs when the absolute value of the relative vorticity is greater than $2.3 \times 10^{-5} \text{ s}^{-1}$. Snapshots of relative vorticity at days (a) 1, (b) 2, (c) 3, and (d) 4 are depicted. The block structures of the c256 and c1024 refinement levels are indicated by black contours.

differences in these features lead to localized large-magnitude errors in the vorticity. As in the separation case, the AMR improves the solution using fewer grid cells. Even when the AMR patch over the vortex is small and the coarse-fine boundary is near the high vorticity cores, the solution is not negatively distorted, showing the robustness of the model given the sensitivity of the test case to grid resolution and slight changes in initial conditions.

The computing run times versus number of grid cells for a 4-day simulation with vorticity tagging is presented

in [Table 5](#). The table thereby represents some of the runs from [Fig. 19b](#). Eight processors on one node of NCAR's Yellowstone computing platform are used. We see the approximate $8\times$ reduction in cost between the c512 and c256 uniform runs, as expected for a doubling of the horizontal resolution and a halving of the time step. For this test, the wall-clock run time for AMR runs is closer to $\approx 4\times$ for $2\times$ resolution changes, demonstrating some of the overhead of the AMR algorithm. The total wall-clock time roughly correlates with the total number of grid cells, as in the moving-vortices advection test, even

FIG. 18. Relative vorticity field at day 4 on the cubed-sphere grid for several runs of the merging-vortices test of [section 3e\(2\)](#) using AMR: (a) uniform c1024 run, (e) uniform c512 run, and (i) uniform c256 run. (b)–(d) AMR runs whose highest refinement level is c1024. (f)–(h) AMR runs that have a maximum refinement level of c512 and (j)–(l) AMR runs that have a c256 maximum resolution. (b),(f),(j) AMR runs with two level of refinement using the large-height-gradient tag. (c),(g),(k) Two-level AMR runs using the small-height-gradient tag. (d),(h),(l) The two-level AMR runs use the relative-vorticity tag. Thus, in the first column, all the AMR runs have a c64 base grid, a c32 base grid in the second column, and a c16 base grid in the third column. All AMR runs use a $\times 4$ refinement ratio.

FIG. 19. For the merging-vortices test of section 3e(2), growth over the 4-day period of (a) normalized L_2 error for relative vorticity calculated with respect to the uniform c1024 run and (b) total number of grid cells. Error and number of grid cells are for uniform runs and one- and two-level AMR runs using the large-height-gradient tag or the relative vorticity tag with $\times 4$ refinement ratios.

for the coarsest two-level AMR runs (c32/c128/c512 and c16/c64/c256).

4. Conclusions

In this paper, we utilized a fourth-order finite-volume model on a cubed-sphere grid, which is adaptive in both space and time, to demonstrate the effectiveness of the AMR in resolving and tracking chosen features of interest while maintaining large-scale smooth flows. Using selected shallow-water and advection test cases, we evaluated the AMR's ability to track and resolve features of interest without creating distortions or numerical noise in the large-scale smooth flows at the interfaces between meshes. A variety of static and dynamic refinement criteria and strategies are implemented to assess the strengths and weaknesses of the AMR method. With the large-scale smooth “do no harm” tests, one and two levels of static and adaptive refinement meshes with several refinement ratios were placed at several locations on the cubed-sphere grid. The results confirmed that multiple levels of refinement and abrupt $\times 4$ or $\times 8$ refinement ratios between levels still allowed flows to

move smoothly through the refined areas. There was little induced noise and numerical error at the refinement boundaries. For coarse resolutions, the refinement improved global errors slightly, and the errors remained nearly unchanged when refinement was added to a higher-resolution base grid for the two shallow-water tests. Only for high resolutions in the USBR tests when a moving AMR grid transected the cubed-sphere panel boundaries did we see a noticeable increase in error. This error, however, was very localized and only becomes apparent because the base global error is so low in the uniform-resolution simulations for this smooth idealized test. In the coarser runs for the USBR and in the other more complex shallow-water tests with larger expected global errors, this was not observed.

With the three AMR test cases, we demonstrated that AMR is able to track the features of interest and closely reproduces the results of uniform high-resolution runs using fewer grid cells. AMR was implemented using tracer and height-field gradients as well as relative-vorticity magnitude as tagging criteria with multiple refinement levels and a range of thresholds. The AMR grids are added and removed in time without creating significant

TABLE 5. Run times (wall-clock time in s and as % of the c512 time) for 4-day merging-vortices simulations [section 3e(2)] with uniform and AMR runs using only eight processors on one node of NCAR's Yellowstone computing platform. The total number of cells is counted at day 4.

AMR resolution	AMR levels	Run time (s)	Run time (%)	Total cells	Total cells (%)
c512	—	24 127	100%	1.6×10^6	100%
c128/c512	1	1913	7.9%	1.9×10^5	12%
c32/c128/c512	2	1225	5.0%	1.2×10^5	7.5%
c256	—	3596	15%	3.9×10^5	25%
c64/c256	1	394	1.6%	5.5×10^4	3.4%
c16/c64/c256	2	304	1.3%	3.6×10^4	2.3%

distortions or noise at the mesh interfaces. In the tracer advection test, fourth-order convergence was maintained while using AMR, and the error of AMR runs with one level of refinement was comparable to the error of uniform runs having the same fine-level resolution. The test showed the importance of refining early, as errors developed from the coarse grid propagate through the fine grids for the rest of the simulation. The gravity wave impinging on the mountain test demonstrated the use of AMR with topography. Refinement was added and removed from the areas with topography without creating additional negative impacts. In the binary vortices test, AMR improved accuracy of the position of the vortices as they interacted and their structures when compared with uniform-resolution runs. Even stringent criteria with high threshold values, which did not create a large buffer of high resolution around the vortices, still produced accurate results and improved the solution in the refined patches. Additional refinement, though, significantly improved the representation of the vorticity filaments that extended well away from the central vortices and the developing Rossby wave train.

All three test cases demonstrated that a variety of AMR criteria and thresholds lead to improvements in the results, though to maximize that improvement, the refinement criteria needed careful tailoring. Several conditions increased the effectiveness of AMR; however, there was no clear strategy for establishing the best general refinement criteria. Having initial refinement or refining early in the run before errors developed on the coarse grid was one of the key strategies for improving accuracy. When there is no initial refinement, the benefits of AMR are limited by the coarseness of the base mesh; AMR with two levels was ineffective due to large errors introduced by the coarsest base meshes early in the calculation. This speaks to the need for a sufficiently refined base mesh to avoid contaminating finer levels. Using more than one level of refinement and effective tagging strategies resulted in better-resolved features of interest, but at a diminishing rate of return of improvement. Our conclusion is that the benefit of AMR does not come automatically from the computational savings of a very coarse base mesh. However, there may still be benefits of two or more levels over uniform-resolution calculations that otherwise would not be computationally feasible without AMR. In a realistic climate simulation, tropical cyclones could, perhaps, be effectively captured early by criteria that place high resolution over the cyclogenesis region and then refine over and track emerging storms to ensure continued accuracy. For other, more complex or moist flows, more advanced criteria than just a simple relative-vorticity threshold

need to be investigated. They could be based, for example, on combinations of physics-based properties (like rainfall), thresholds of vorticity, or gradients. Future work will explore such refinement criteria in the 3D nonhydrostatic version of the Chombo-AMR model with and without a variety of physical parameterization schemes. The addition of physical parameterizations will also allow us to test the scale awareness of the physics schemes.

Acknowledgments. Support for this work has been provided by the Office of Science, U.S. Department of Energy, Award DE-SC0003990 and by the Director, Office of Science, Office of Advanced Scientific Computing Research of the U.S. Department of Energy under Contract DE-AC02-05CH11231 as part of their Mathematical, Computational, and Computer Sciences Research/Computational Partnerships Program. We would like to acknowledge high-performance computing support from Yellowstone provided by NCAR's Computational and Information Systems Laboratory, sponsored by the National Science Foundation. We thank the reviewers for their helpful comments and suggestions.

REFERENCES

- Adams, M., and Coauthors, 2015: Chombo software package for AMR applications—Design document. Tech. Rep. LBNL-6616E, Lawrence Berkeley National Laboratory, 204 pp.
- Aechtner, M., N.-R. Kevlahan, and T. Dubos, 2015: A conservative wavelet method for the shallow-water equations on the sphere. *Quart. J. Roy. Meteor. Soc.*, **141**, 1712–1726, doi:[10.1002/qj.2473](https://doi.org/10.1002/qj.2473).
- Bacon, D. P., and Coauthors, 2000: A dynamically adapting weather and dispersion model: The Operational Multiscale Environment Model with Grid Adaptivity (OMEGA). *Mon. Wea. Rev.*, **128**, 2044–2076, doi:[10.1175/1520-0493\(2000\)128<2044:ADAWAD>2.0.CO;2](https://doi.org/10.1175/1520-0493(2000)128<2044:ADAWAD>2.0.CO;2).
- Bauer, W., M. Baumann, L. Scheck, A. Gassmann, V. Heuveline, and S. C. Jones, 2014: Simulation of tropical-cyclone-like vortices in shallow-water ICON-hex using goal-oriented r-adaptivity. *Theor. Comput. Fluid Dyn.*, **28**, 107–128, doi:[10.1007/s00162-013-0303-4](https://doi.org/10.1007/s00162-013-0303-4).
- Behrens, J., 1998: Atmospheric and ocean modeling with an adaptive finite element solver for the shallow-water equations. *Appl. Numer. Math.*, **26**, 217–226, doi:[10.1016/S0168-9274\(97\)00090-1](https://doi.org/10.1016/S0168-9274(97)00090-1).
- , N. Rakowsky, W. Hiller, D. Handorf, M. Läuter, J. Päpke, and K. Dethloff, 2005: amatos: Parallel adaptive mesh generator for atmospheric and oceanic simulation. *Ocean Modell.*, **10**, 171–183, doi:[10.1016/j.ocemod.2004.06.003](https://doi.org/10.1016/j.ocemod.2004.06.003).
- Berger, M. J., and P. Colella, 1989: Local adaptive mesh refinement for shock hydrodynamics. *J. Comput. Phys.*, **82**, 64–84, doi:[10.1016/0021-9991\(89\)90035-1](https://doi.org/10.1016/0021-9991(89)90035-1).
- Blaise, S., and A. St-Cyr, 2012: A dynamic hp-adaptive discontinuous Galerkin method for shallow-water flows on the sphere with application to a global tsunami simulation. *Mon. Wea. Rev.*, **140**, 978–996, doi:[10.1175/MWR-D-11-00038.1](https://doi.org/10.1175/MWR-D-11-00038.1).

- Chen, C., F. Xiao, and X. Li, 2011a: An adaptive multimoment global model on a cubed sphere. *Mon. Wea. Rev.*, **139**, 523–548, doi:[10.1175/2010MWR3365.1](https://doi.org/10.1175/2010MWR3365.1).
- , —, —, and Y. Yang, 2011b: A multi-moment transport model on cubed-sphere grid. *Int. J. Numer. Methods Fluids*, **67**, 1993–2014, doi:[10.1002/rd.2478](https://doi.org/10.1002/rd.2478).
- , X. Li, X. Shen, and F. Xiao, 2014: Global shallow water models based on multi-moment constrained finite volume method and three quasi-uniform spherical grids. *J. Comput. Phys.*, **271**, 191–223, doi:[10.1016/j.jcp.2013.10.026](https://doi.org/10.1016/j.jcp.2013.10.026).
- Dietachmayer, G. S., and K. K. Droegemeier, 1992: Application of continuous dynamic grid adaption techniques to meteorological modeling. Part I: Basic formulation and accuracy. *Mon. Wea. Rev.*, **120**, 1675–1706, doi:[10.1175/1520-0493\(1992\)120<1675:AOCDGA>2.0.CO;2](https://doi.org/10.1175/1520-0493(1992)120<1675:AOCDGA>2.0.CO;2).
- Dritschel, D., and D. Waugh, 1992: Quantification of the inelastic interaction of unequal vortices in two-dimensional vortex dynamics. *Phys. Fluids A*, **4**, 1737–1744, doi:[10.1063/1.858394](https://doi.org/10.1063/1.858394).
- Dübén, P. D., P. Korn, and V. Aizinger, 2012: A discontinuous/continuous low order finite element shallow water model on the sphere. *J. Comput. Phys.*, **231**, 2396–2413, doi:[10.1016/j.jcp.2011.11.018](https://doi.org/10.1016/j.jcp.2011.11.018).
- Eskilsson, C., 2011: An *hp*-adaptive discontinuous Galerkin method for shallow water flows. *Int. J. Numer. Methods Fluids*, **67**, 1605–1623, doi:[10.1002/rd.2434](https://doi.org/10.1002/rd.2434).
- Fujiwhara, S., 1921: The natural tendency towards symmetry of motion and its application as a principle in meteorology. *Quart. J. Roy. Meteor. Soc.*, **47**, 287–292, doi:[10.1002/qj.49704720010](https://doi.org/10.1002/qj.49704720010).
- Giraldo, F., 2000: The Lagrange–Galerkin method for the two-dimensional shallow water equations on adaptive grids. *Int. J. Numer. Methods Fluids*, **33**, 789–832, doi:[10.1002/1097-0363\(20000730\)33:6<789::AID-FLD29>3.0.CO;2-1](https://doi.org/10.1002/1097-0363(20000730)33:6<789::AID-FLD29>3.0.CO;2-1).
- Harris, L. M., and S.-J. Lin, 2013: A two-way nested global-regional dynamical core on the cubed-sphere grid. *Mon. Wea. Rev.*, **141**, 283–306, doi:[10.1175/MWR-D-11-00201.1](https://doi.org/10.1175/MWR-D-11-00201.1).
- Hendricks, E. A., M. A. Kopera, F. X. Giraldo, M. S. Peng, J. D. Doyle, and Q. Jiang, 2016: Evaluation of the utility of static and adaptive mesh refinement for idealized tropical cyclone problems in a spectral element shallow-water model. *Mon. Wea. Rev.*, **144**, 3697–3724, doi:[10.1175/MWR-D-15-0146.1](https://doi.org/10.1175/MWR-D-15-0146.1).
- Holland, G. J., and G. S. Dietachmayer, 1993: On the interaction of tropical-cyclone-scale vortices. III: Continuous barotropic vortices. *Quart. J. Roy. Meteor. Soc.*, **119**, 1381–1398, doi:[10.1002/qj.49711951408](https://doi.org/10.1002/qj.49711951408).
- Huang, X., A. M. Rhoades, P. A. Ullrich, and C. M. Zarzycki, 2016: An evaluation of the variable-resolution CESM for modeling California's climate. *J. Adv. Model. Earth Syst.*, **8**, 345–369, doi:[10.1002/2015MS000559](https://doi.org/10.1002/2015MS000559).
- Hubbard, M., and N. Nikiforakis, 2003: A three-dimensional, adaptive, Godunov-type model for global atmospheric flows. *Mon. Wea. Rev.*, **131**, 1848–1864, doi:[10.1175/2568.1](https://doi.org/10.1175/2568.1).
- Iselin, J. P., J. M. Prusa, and W. J. Gutowski, 2002: Dynamic grid adaptation using the MPDATA scheme. *Mon. Wea. Rev.*, **130**, 1026–1039, doi:[10.1175/1520-0493\(2002\)130<1026:DGAUTM>2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130<1026:DGAUTM>2.0.CO;2).
- Jablonowski, C., M. Herzog, J. E. Penner, R. C. Oehmke, Q. F. Stout, B. van Leer, and K. G. Powell, 2006: Block-structured adaptive grids on the sphere: Advection experiments. *Mon. Wea. Rev.*, **134**, 3691–3713, doi:[10.1175/MWR3223.1](https://doi.org/10.1175/MWR3223.1).
- , R. C. Oehmke, and Q. F. Stout, 2009: Block-structured adaptive meshes and reduced grids for atmospheric general circulation models. *Philos. Trans. Roy. Soc. London*, **367A**, 4497–4522, doi:[10.1098/rsta.2009.0150](https://doi.org/10.1098/rsta.2009.0150).
- Kinter, J. L., III, and Coauthors, 2013: Revolutionizing climate modeling with project Athena: A multi-institutional, international collaboration. *Bull. Amer. Meteor. Soc.*, **94**, 231–245, doi:[10.1175/BAMS-D-11-00043.1](https://doi.org/10.1175/BAMS-D-11-00043.1).
- Kopera, M. A., and F. X. Giraldo, 2014: Analysis of adaptive mesh refinement for IMEX discontinuous Galerkin solutions of the compressible Euler equations with application to atmospheric simulations. *J. Comput. Phys.*, **275**, 92–117, doi:[10.1016/j.jcp.2014.06.026](https://doi.org/10.1016/j.jcp.2014.06.026).
- Kubatko, E. J., S. Bunya, C. Dawson, and J. J. Westerink, 2009: Dynamic *p*-adaptive Runge–Kutta discontinuous Galerkin methods for the shallow water equations. *Comput. Methods Appl. Mech. Eng.*, **198**, 1766–1774, doi:[10.1016/j.cma.2009.01.007](https://doi.org/10.1016/j.cma.2009.01.007).
- Kühnlein, C., P. K. Smolarkiewicz, and A. Dörnbrack, 2012: Modeling atmospheric flows with adaptive moving meshes. *J. Comput. Phys.*, **231**, 2741–2763, doi:[10.1016/j.jcp.2011.12.012](https://doi.org/10.1016/j.jcp.2011.12.012).
- Lauritzen, P. H., R. D. Nair, and P. A. Ullrich, 2010: A conservative semi-Lagrangian multi-tracer transport scheme (CSLAM) on the cubed-sphere grid. *J. Comput. Phys.*, **229**, 1401–1424, doi:[10.1016/j.jcp.2009.10.036](https://doi.org/10.1016/j.jcp.2009.10.036).
- Läuter, M., D. Handorf, and K. Dethloff, 2005: Unsteady analytical solutions of the spherical shallow water equations. *J. Comput. Phys.*, **210**, 535–553, doi:[10.1016/j.jcp.2005.04.022](https://doi.org/10.1016/j.jcp.2005.04.022).
- , —, N. Rakowsky, J. Behrens, S. Frickenhaus, M. Best, K. Dethloff, and W. Hiller, 2007: A parallel adaptive barotropic model of the atmosphere. *J. Comput. Phys.*, **223**, 609–628, doi:[10.1016/j.jcp.2006.09.029](https://doi.org/10.1016/j.jcp.2006.09.029).
- Li, X., C. Chen, F. Xiao, and X. Shen, 2015: A high-order multi-moment constrained finite-volume global shallow-water model on the Yin-Yang grid. *Quart. J. Roy. Meteor. Soc.*, **141**, 2090–2102, doi:[10.1002/qj.2504](https://doi.org/10.1002/qj.2504).
- Manganello, J. V., and Coauthors, 2012: Tropical cyclone climatology in a 10-km global atmospheric GCM: Toward weather-resolving climate modeling. *J. Climate*, **25**, 3867–3893, doi:[10.1175/JCLI-D-11-00346.1](https://doi.org/10.1175/JCLI-D-11-00346.1).
- Marras, S., M. Kopera, and F. Giraldo, 2015: Simulation of shallow-water jets with a unified element-based continuous/discontinuous Galerkin model with grid flexibility on the sphere. *Quart. J. Roy. Meteor. Soc.*, **141**, 1727–1739, doi:[10.1002/qj.2474](https://doi.org/10.1002/qj.2474).
- McCorquodale, P., P. A. Ullrich, H. Johansen, and P. Colella, 2015: An adaptive multiblock high-order finite-volume method for solving the shallow-water equations on the sphere. *Commun. Appl. Math. Comput. Sci.*, **10**, 121–162, doi:[10.2140/camcos.2015.10.121](https://doi.org/10.2140/camcos.2015.10.121).
- Melander, M., N. Zabusky, and J. McWilliams, 1988: Symmetric vortex merger in two dimensions: Causes and conditions. *J. Fluid Mech.*, **195**, 303–340, doi:[10.1017/S0022112088002435](https://doi.org/10.1017/S0022112088002435).
- Miura, H., M. Satoh, H. Tomita, A. Noda, T. Nasuno, and S. Iga, 2007: A short-duration global cloud-resolving simulation with a realistic land and sea distribution. *Geophys. Res. Lett.*, **34**, L02804, doi:[10.1029/2006GL027448](https://doi.org/10.1029/2006GL027448).
- Miyamoto, Y., Y. Kajikawa, R. Yoshida, T. Yamaura, H. Yashiro, and H. Tomita, 2013: Deep moist atmospheric convection in a subkilometer global simulation. *Geophys. Res. Lett.*, **40**, 4922–4926, doi:[10.1002/grl.50944](https://doi.org/10.1002/grl.50944).
- Müller, A., J. Behrens, F. X. Giraldo, and V. Wirth, 2013: Comparison between adaptive and uniform discontinuous Galerkin simulations in dry 2D bubble experiments. *J. Comput. Phys.*, **235**, 371–393, doi:[10.1016/j.jcp.2012.10.038](https://doi.org/10.1016/j.jcp.2012.10.038).
- Nair, R. D., and C. Jablonowski, 2008: Moving vortices on the sphere: A test case for horizontal advection problems. *Mon. Wea. Rev.*, **136**, 699–711, doi:[10.1175/2007MWR2105.1](https://doi.org/10.1175/2007MWR2105.1).
- Prieto, R., B. D. McNoldy, S. R. Fulton, and W. H. Schubert, 2003: A classification of binary tropical cyclone-like vortex

- interactions. *Mon. Wea. Rev.*, **131**, 2656–2666, doi:[10.1175/1520-0493\(2003\)131<2656:ACOBTC>2.0.CO;2](https://doi.org/10.1175/1520-0493(2003)131<2656:ACOBTC>2.0.CO;2).
- Putykiewicz, J. A., 2011: On numerical solution of the shallow water equations with chemical reactions on icosahedral geodesic grid. *J. Comput. Phys.*, **230**, 1956–1991, doi:[10.1016/j.jcp.2010.11.045](https://doi.org/10.1016/j.jcp.2010.11.045).
- Putman, W. M., and M. Suarez, 2011: Cloud-system resolving simulations with the NASA Goddard Earth Observing System global atmospheric model (GEOS-5). *Geophys. Res. Lett.*, **38**, L16809, doi:[10.1029/2011GL048438](https://doi.org/10.1029/2011GL048438).
- Rauscher, S. A., and T. D. Ringler, 2014: Impact of variable-resolution meshes on midlatitude baroclinic eddies using CAM-MPAS-A. *Mon. Wea. Rev.*, **142**, 4256–4268, doi:[10.1175/MWR-D-13-00366.1](https://doi.org/10.1175/MWR-D-13-00366.1).
- Reed, K. A., and C. Jablonowski, 2011: An analytic vortex initialization technique for idealized tropical cyclone studies in AGCMs. *Mon. Wea. Rev.*, **139**, 689–710, doi:[10.1175/2010MWR3488.1](https://doi.org/10.1175/2010MWR3488.1).
- Ritchie, E. A., and G. J. Holland, 1993: On the interaction of tropical-cyclone-scale vortices. II: Discrete vortex patches. *Quart. J. Roy. Meteor. Soc.*, **119**, 1363–1379, doi:[10.1002/qj.49711951407](https://doi.org/10.1002/qj.49711951407).
- Shin, S.-E., J.-Y. Han, and J.-J. Baik, 2006: On the critical separation distance of binary vortices in a nondivergent barotropic atmosphere. *J. Meteor. Soc. Japan*, **84**, 853–869, doi:[10.2151/jmsj.84.853](https://doi.org/10.2151/jmsj.84.853).
- Skamarock, W. C., and J. B. Klemp, 1993: Adaptive grid refinement for two-dimensional and three-dimensional non-hydrostatic atmospheric flow. *Mon. Wea. Rev.*, **121**, 788–804, doi:[10.1175/1520-0493\(1993\)121<0788:AGRFTD>2.0.CO;2](https://doi.org/10.1175/1520-0493(1993)121<0788:AGRFTD>2.0.CO;2).
- , J. Oliger, and R. L. Street, 1989: Adaptive grid refinement for numerical weather prediction. *J. Comput. Phys.*, **80**, 27–60, doi:[10.1016/0021-9991\(89\)90089-2](https://doi.org/10.1016/0021-9991(89)90089-2).
- St-Cyr, A., C. Jablonowski, J. M. Dennis, H. M. Tufo, and S. J. Thomas, 2008: A comparison of two shallow water models with nonconforming adaptive grids: Classical tests. *Mon. Wea. Rev.*, **136**, 1898–1922, doi:[10.1175/2007MWR2108.1](https://doi.org/10.1175/2007MWR2108.1).
- Tóth, G., and Coauthors, 2012: Adaptive numerical algorithms in space weather modeling. *J. Comput. Phys.*, **231**, 870–903, doi:[10.1016/j.jcp.2011.02.006](https://doi.org/10.1016/j.jcp.2011.02.006).
- Tumolo, G., and L. Bonaventura, 2015: A semi-implicit, semi-Lagrangian discontinuous Galerkin framework for adaptive numerical weather prediction. *Quart. J. Roy. Meteor. Soc.*, **141**, 2582–2601, doi:[10.1002/qj.2544](https://doi.org/10.1002/qj.2544).
- Ullrich, P. A., 2014: Understanding the treatment of waves in atmospheric models. Part 1: The shortest resolved waves of the 1D linearized shallow-water equations. *Quart. J. Roy. Meteor. Soc.*, **140**, 1426–1440, doi:[10.1002/qj.2226](https://doi.org/10.1002/qj.2226).
- Waugh, D. W., 1992: The efficiency of symmetric vortex merger. *Phys. Fluids A*, **4**, 1745–1758, doi:[10.1063/1.858395](https://doi.org/10.1063/1.858395).
- 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. *Mon. Wea. Rev.*, **137**, 4208–4224, doi:[10.1175/2009MWR2917.1](https://doi.org/10.1175/2009MWR2917.1).
- , P. Browne, C. Budd, and M. Cullen, 2016: Mesh adaptation on the sphere using optimal transport and the numerical solution of a Monge–Ampère type equation. *J. Comput. Phys.*, **308**, 102–123, doi:[10.1016/j.jcp.2015.12.018](https://doi.org/10.1016/j.jcp.2015.12.018).
- Williamson, D. L., J. B. Drake, J. J. Hack, R. Jakob, and P. N. Swarztrauber, 1992: A standard test set for numerical approximations to the shallow water equations in spherical geometry. *J. Comput. Phys.*, **102**, 211–224, doi:[10.1016/S0021-9991\(05\)80016-6](https://doi.org/10.1016/S0021-9991(05)80016-6).
- Zarzycki, C. M., and C. Jablonowski, 2015: Experimental tropical cyclone forecasts using a variable-resolution global model. *Mon. Wea. Rev.*, **143**, 4012–4037, doi:[10.1175/MWR-D-15-0159.1](https://doi.org/10.1175/MWR-D-15-0159.1).
- , —, and M. A. Taylor, 2014: Using variable resolution meshes to model tropical cyclones in the Community Atmosphere Model. *Mon. Wea. Rev.*, **142**, 1221–1239, doi:[10.1175/MWR-D-13-00179.1](https://doi.org/10.1175/MWR-D-13-00179.1).