Bridging Scales in 2- and 3-Dimensional Atmospheric Modeling with Adaptive Mesh Refinement

by

Jared Ferguson

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Applied Physics)
in The University of Michigan
2018

Doctoral Committee:

Associate Professor Christiane Jablonowski, Chair
Associate Professor Mark G. Flanner
Hans Johansen, Lawrence Berkeley National Laboratory
Professor Richard B. Rood

Jared Ferguson

joferg@umich.edu

ORCID iD: 0000-0002-7741-0070

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my thesis advisor and committee chair, Christiane Jablonowski. She has provided superb mentorship and has contributed in a significant manner to my future career. I would like to thank the remainder of my committee, Hans Johansen, Mark Flanner, and Richard Rood. I appreciate all your advice and support in this endeavor.

Further thanks to Hans Johansen for all the guidance, mentorship, and countless weeks spent debugging to get the model to work. I would also like to thank the others at Lawerence Berkeley National Laboratory who assisted in the various facets of working on the model that this thesis is centered on, including Phil Colella, Peter McCorquodale, Eli Goodfriend, and Paul Ullrich. Thank you for the support and helpfully answering my numerous questions.

Also, thank you Cyndi McNabb, Lauren Segall, Chuck Sutton, Sandra Pytlinksi, and Allison Lyons for guiding me through the bureaucracy of being a graduate student in two departments. Sometimes it seemed like more than twice the work, so thanks for making it easy.

Last but far from least, I thank all my friends and family without whom my success and the completion of my degree would not have been possible. Thank you to my friends and officemates in the Jablonowksi Data Lab Kevin, Paul, James, Weiye, Colin, Allison, and Diana for the advice and helpful work-related conversation. However, I am especially grateful for the countless hours we spent joking in the office, hanging out at conferences, and venturing outside the SRB for activities around Ann

Arbor. To my Applied Physics compatriots, thanks for the six years of adventures that only physicists can have. Our game nights and random late night shenanigans made getting though Jackson a breeze. I also thank my second family in AOSS (now CLASP) for welcoming me into the department and teaching me how the atmosphere and weather actually work. You also provided much needed mental breaks from the day to day stresses of graduate school. Of course, thank you to my long time roommate and good friend Carrie as well as my numerous IM sports compatriots including. Though we established an illustrious Red Zeppelin dynasty in co-rec softball, we failed to win any other championships in the half dozen other IM sports we played. The fun we had playing and then recovering from our injuries afterwards at The Arena more than made up for those losses. I would be remiss to forget my science policy and student government friends and mentors who helped turn an enjoyable side interest into a future career path.

I am forever thankful for my parents who provided me with countless opportunities to succeed in life as well as in academics. They have always supported me in every decision I have made. I would also begrudgingly thank my sister Brennan for the best sibling rivalry a brother could ask for. Most importantly, I'd like to thank the soon to be doctor Katie whose companionship has made my time here a joy. Her love and support, as well as some key proofreading, allowed me to get through these stressful times.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

ABSTRACT

Complex multi-scale atmospheric phenomena, like tropical cyclones, challenge conventional weather and climate models, which use relatively coarse uniform-grid resolutions to cope with computational costs. Adaptive Mesh Refinement (AMR) techniques mitigate these challenges by dynamically and transiently placing high-resolution grids over salient features, thus providing sufficient local resolution while limiting the computational burden.

This thesis explores the development of AMR, a technique that has been featured only sporadically in the atmospheric science literature, within a new nonhydrostatic, finite-volume dynamical core and demonstrates AMR's effectiveness in improving model accuracy and ability to resolve multi-scale features. 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 developed with the Chombo AMR library. The AMR dynamical core is implemented in a hierarchy of models of increasing complexity, from an idealized 2D shallow water configuration to the non-hydrostatic 3D equation set with subgrid-scale parameterizations schemes. AMR's numerical accuracy, computational efficiency, and ability to track and resolve multifaceted and evolving features are assessed with a variety of existing and new test cases, implemented within each model iteration.

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 vortic-

ity, or minimum pressure based thresholds, in order to achieve the best accuracy for cost. Simulations show that the model's AMR can accurately resolve key local features in both shallow water and 3D test cases without requiring global high-resolution grids, as 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 degradation of the large-scale smooth flows to a minimum. 2D and 3D physics parameterizations are able to function effectively over multiple levels of refinement, though the parameterizations are sensitive to grid resolution.

AMR is most effective when refinement is triggered early or when the base uniform resolution can partially resolve the features of interests. Very coarse base resolutions lead to large initial errors that cannot be overcome by AMR. However, the addition of refinement later in the simulation still results in significant improvements, especially in resolving small-scale features. The research showed that flow properties, such as strong gradients or rainbands, can be sensitive to small changes in AMR criteria. These may delay the onset of the refinement or alter the shape of the refined area, which impacts the evolution of the flow. With coarse base resolutions, the tagging criteria must therefore be uniquely tailored to capture the early growth phases of the feature of interest. A promising refinement technique is a combination of some initial refinement and AMR. The initial refinement limits error growth at the base resolution and ensures that the model can resolve the feature of interest. Overall, AMR is shown to be a powerful modeling approach that bridges the resolution gap for extreme weather events.

CHAPTER I

Introduction: Dynamic refinement in climate and weather models

1.1 Motivation

Tropical cyclones are some of the most powerful and destructive weather systems on earth having the potential to cause major devastation and significant loss of life. Yet traditional global weather and climate models typically have coarse grid resolutions that cannot resolve the extreme pressure gradients, high wind speeds, and other small-scale (1-10 km) processes that drive these cyclones. The development of Adaptive Mesh Refinement (AMR) techniques offers a transformative new approach to incorporate these scales within atmospheric General Circulation Models (GCMs). GCMs with AMR capabilities will be able to identify traveling salient features and selectively enhance the grid resolution over them while keeping less active regions at coarser mesh spacings, thereby reducing computational costs. This has the potential to improve the representation of TCs and other extreme weather and climate events by allowing next-generation models to achieve unprecedented high resolutions in local areas.

1.2 General circulation models

General circulation models (GCMs) are the backbone of today's global climate modeling research; these models use numerical approximations to simulate the circulation of the Earth's atmosphere. GCMs are comprised of two components. The dynamical core component is considered the engine of a GCM. The atmosphere is spatially discretized onto a grid and the dynamical core (dycore) uses numerical methods to resolve the fluid motion and thermodynamic quantities on each grid. The second component contains the physics parameterizations. These approximate the atmospheric features including radiation and sub-grid scale process like convection, clouds, and turbulence that are not resolvable by the dycore. GCMs are used in wide ranging temporal scales, from hours and days for short term weather prediction to decades and centuries for long-term climate assessments.

1.3 High resolution in climate and weather models

Atmospheric phenomena which have an outsized effect on humans and society, such as tropical cyclones, atmospheric rivers, and severe thunderstorms, exist on relative small spatial scales. Even large scale global phenomena, like the Madden-Julian Oscillation, rely on key processes like cloud formation and convection that need to be parameterized for resolutions higher than a few kilometers. Decreasing the distance between grid elements improves the ability of the model to directly resolve atmospheric features at smaller spatial scales. Increased resolution results in broad improvements for resolving key climate and weather features as models can better capture physical processes and interactions between the atmosphere and land or ocean (Prodhomme et al., 2016). Increased resolution, however, is computationally intensive. The processing power and memory constraints of computers used to run the models are the limiting factors in increasing horizontal resolution. Increasing

resolution not only increases the number of grid cells of the model but also requires a shorter time step due to a more restrictive Courant-Friedrichs-Lewy (CFL) condition. As a result, a doubling of the horizontal resolution generally results in eight-fold more calculations.

The global climate models used in the Fifth Assessment Report by the Intergovernmental Panel on Climate Change (IPCC) used grid resolutions between 50 and 300 km (Flato et al., 2013). For comparison the average tropical cyclone has a diameter of ~ 600 km with the average cyclone eye ~ 30 km across. The latest high-resolution global climate models can implement 10 - 50 km grid spacing (Kinter et al., 2013; Hayhoe et al., 2017). However, they are computationally expensive to run extensively, and they are still unable to explicitly represent key processes, such as clouds. The global weather models used by the European Centre for Medium-Range Weather Forecasts (ECMWF) and NOAA have grid spacings of 9 km and 13 km respectively (Haiden et al., 2016; Dias et al., 2018) and are only run for spans of a few weeks. A few non-hydrostatic, partly cloud-resolving global simulations by Putman and Suarez (2011) and Miyamoto et al. (2013) were run for short, multi-day time periods with grid spacings in the 0.93.5 km range. These simulations were single runs demonstrating the concept of such models and are not feasible for climate simulations or even numerical weather predictions. With global high-resolution GCMs burdened by computational costs, alternative methods are needed to increase resolution over specific areas of interests.

Limited area models (LAM) are one possible solution. They can operate with increased resolution down to $10-50$ km grid spacing for climate simulations and 1 to 3 km grid spacing for short-term weather forecasting. LAMs eliminate computational expense by numerically simulating only a small specific region (e.g the continental United States), but this requires the lateral boundaries to be externally forced. These boundary conditions are typically derived from coarser GCMs, which use different nu-

merical schemes and physical parameterizations, thereby introducing possible biases or numerical discrepancies. Since LAMs are not global in scale, key conservation properties are not obeyed. Due to these boundary conditions, LAMs might not be able to effectively capture the large-scale, global climate feedbacks triggered by localized phenomena such as tropical cyclones.

1.4 Variable Resolution models

We are interested in combining the localized high-resolution and computational economy of LAMs with the numerical and physical consistency of a global model. Variable resolution GCMs (VRGCMs) place additional grid elements only where high-resolution is required. They maintain global connectivity between the areas of coarse and fine grids cells, eliminating the need for forced lateral boundary conditions. In addition, they maintain two-way interactions, unlike one-way nested LAMs, which permit features within the refined nest to affect the global solution. VRGCM is a classification that encompasses any global model which implements multiple grid resolutions with two-way interactions. The two primary techniques used to implement variable resolution are grid stretching and grid nesting. In either case, the grid can be fixed in place, *static refinement*, or be designed to adapt in time based on some preset criteria, *dynamic refinement*.

1.4.1 Static Refinement

VRGCMs with static refinement have rapidly grown in prominence for climate and weather research over the last decade. Though stretched grids have historically been the more popular method of achieving variable resolution with GCMs, nested grids have become more of a fixture in the last ten years.

1.4.1.1 Stretched Grids

Grid stretching techniques deform the global grid with a fixed number of grid elements so that more grid elements are concentrated in the region of interest, reducing the number of elements over the rest of the domain. This grid alteration results in a single global variable-resolution grid with a smooth, gradual transition between resolutions. This technique was historically attractive for VRGCMs because it required only minor modifications to existing numerical schemes and grid structures.

Two general stretching methods are a physically stretched spherical coordinate system developed for grid point models by Staniforth and Mitchell (1978) and a conformal coordinate transformation method developed for spectral models by Schmidt (1977). Early stretched spherical coordinate system models used two band-like grid structures perpendicular to each other with the highest resolution in the area where the two bands intersect. Stretched grid techniques have been added to several short-term forecasting models beginning in the mid-1990s (Déqué and Piedelievre, 1995; Côté et al., 1998; Fox-Rabinovitz et al., 1997). Later developments include a conformal-cubic stretched grid, which focuses stretching on a single area of interest, avoiding the extra computational overhead created by the traditional global banded grid structure (McGregor, 1996). A stretching technique that permits refinements over multiple regions of interest was introduced by Fox-Rabinovitz et al. (2002). More detailed assessments of stretched grids in atmospheric research are presented in Fox-Rabinovitz et al. (2006) and in McGregor (2013).

Several newer GCMs employ stretched grids for climate and weather simulations. The Conformal-Cubic Atmospheric Model (CCAM), which employs grid stretching on the cubed-sphere grid (McGregor and Dix, 2008), has beed used in high-resolution climate projections (McGregor et al., 2016). The Non-hydrostatic Icosahedral Atmospheric Model (NICAM) model has a stretched grid version (Tomita, 2008), which has been used to study African easterly waves (Satoh et al., 2013) and aerosols over Japan

(Goto et al., 2015). Harris and Lin (2013) have also implemented grid stretching in the Geophysical Fluid Dynamics Laboratory (GFDL) cubed-sphere finite volume model (FV3) and evaluated tropical cyclone forecasts.

1.4.1.2 Nested Grids

In a nested grid setup, a higher resolution grid is embedded on or in the coarse grid. Unlike grid stretching, additional grid elements are added to the mesh in regions of interest. Regions away from this remain at the same resolution. While LAMs and their parent global models are a form of nested grids, the traditional LAM setup only permits one-way transfer of information, from the global model to the LAM at the boundary of the inner LAM grid domain. The focus of this section is models that operate on the high and coarse resolution grids concurrently and allow two-way information flow between the grids.

Two-way nesting has been used for multiple resolution levels within LAMs since the 1970s (Koch and McQueen, 1987). Several global model and LAM pairings have implemented two-way nesting (e.g. Dudhia and Bresch (2002) and Lorenz and Jacob (2005)). In these models, the solution from the coarse global grid is interpolated on the regional model grid before each time step and, after each advance, the fine grid solution is remapped to the global coarse grid. An improvement can be made by using the same model (i.e. same numerical discretization, computation grid and, physical parameterizations) for both the global coarse grid and localized refined grid as described for the shallow water equations in Ruge et al. (1995). More recently a two-way nesting refinement was implemented with GFDL's FV3 model (Harris and Lin, 2013, 2014).

A second nested approach is multiscale grids which span multiple resolutions within a single mesh rather than overlaying grids of different resolutions. This method removes the difficulties of interpolating and communicating between multiple sepa-

rate grids, but simultaneously it requires a more complex computational grid. It also requires the entire global grid to be numerical integrated with a global time step restricted by the size of the small grid cells. Examples of GCMs using this single grid multiscale technique include the Model for Prediction Across Scales (MPAS) (Skamarock et al., 2012), the Ocean-Land-Atmosphere Model (OLAM) (Walko and Avissar, 2011), and the Community Atmosphere Model's Spectral Element dynamical core (CAM-SE) (Dennis et al., 2012; Zarzycki et al., 2014a). OLAM has been used to investigate teleconnections between Amazon deforestation and the snow pack in the western United States (Medvigy et al., 2013). Variable resolution in MPAS was employed for Madden-Julian oscillation simulations (Rauscher and Ringler, 2014) and for sensitivity tests of atmospheric rivers to model resolution (Hagos et al., 2015). CAM-SE was used in regional climate studies (Huang et al., 2016; Rhoades et al., 2018) and experimental tropical cyclone forecasting (Zarzycki and Jablonowski, 2015).

1.4.2 Dynamic Refinement

Dynamic refinement permits the grid resolution to change as the model simulation progresses, allowing high-resolution meshes to track transient features or processes in the simulation. Dynamic refinement is more fully adaptive than a simple moving fixed-sized nested grid. In dynamic refinement, the grid is refined over important physical processes or atmospheric features that need additional resolution and then coarsened over the same space once higher resolution is no longer needed. Dynamic refinement techniques do not need to know *a priori* where to place high-resolution, an added advantage over static refinement. They can create, migrate, and remove high-resolution patches as the simulation requires. Adaptive refinement methods are well established in many areas of computational fluid dynamics, such as aerospace or space weather modeling (see e.g. Tóth et al. (2012)). In atmospheric science, however, the costs and benefits of AMR methods have only been evaluated in idealized simulations

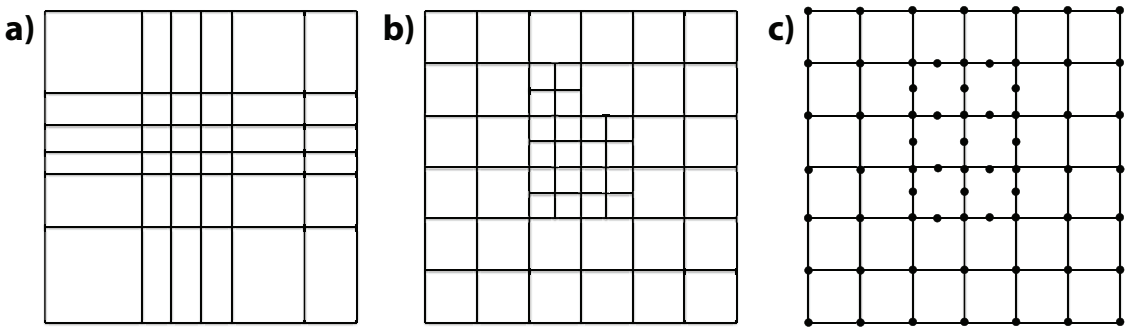

Figure 1.1: Simplified examples of adaptive refinement techniques: (a) r-refinement — The number of grid cells is held constant, but the spacing between them changed, (b) h-refinement — the number of grid cells is increased by subdividing the grid, and (c) p-refinement — the polynomial order on the sub-grid scale is increased over the area of interest. The black dots represent the nodes in each grid cell.

and simplified models so far.

TC prediction researchers made the first foray into dynamic refinement within atmospheric science (Ley and Elsberry, 1976; Kurihara et al., 1979; Jones, 1977). These moving nested-grid models, though not quite fully adaptive, would shift the entire nested grid to keep the cyclone within the refined mesh area. The grid's movement was rigid and predetermined, demanding some prior knowledge of the cyclone's trajectory, and it required the total number of grid points to remain constant. Some of the first truly adaptive models for atmospheric flows, which added and moved grid elements dynamically, were developed by Skamarock et al. (1989), Skamarock and Klemp (1993), and Dietachmayer and Droegemeier (1992). Adaptive schemes can be categorized into three grid refinement strategies: r-refinement (or r-adaptivity), h-refinement, and p-refinement. Figure 1.1 shows simplified examples of the three types.

1.4.2.1 r-refinement

In r-refinement, the grid topology and number of grid elements remain unchanged; instead, the elements are dynamically redistributed to increase resolution in parts of the grid while decreasing it everywhere else (Fig. 1.1a). This dynamic mesh redistribution creates smoother transition regions between resolutions but requires a complex global remapping of the mesh to move the location of the high resolution. The number of refined areas that can be formed are limited, and if the remapping transformation algorithms are not carefully constructed, grid tangling problems and convergence issues can arise (Budd et al., 2009). r-refinement is the dynamic version of static stretched grids.

In an early application of r-refinement, Dietachmayer and Droegemeier (1992) and Dietachmayer (1992) used this global grid redistribution technique to increase resolution in areas where the estimated solution error is high. Giraldo (2000) and Iselin et al. (2002) applied this type of dynamic adaptation for the 2D shallow water equations and advection problems, respectively. More recently, Walsh (2010) studied two-dimensional baroclinic instability with r-refinement and 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.

1.4.2.2 h-refinement: Adaptive Mesh Refinement

Adaptive mesh refinement (AMR), also known as 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 (Fig. 1.1b). Nested grids are the comparable static refinement technique.

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 trun-

cation error estimates. The grid-cell solutions and boundary conditions between the finer-resolution meshes and the base grid are continually updated. In a more recent example, Chen et al. (2011a) use AMR to overlay high-resolution meshes in areas of interest in a shallow water model on a cubed-sphere grid.

Examples of AMR techniques which 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 in a non-conforming grid, cells can share an edge with more than one neighboring element (Fig. 1.2). Conformal grids can maintain conservation properties easily, avoid possible discontinuities at refinement boundaries, and provide smoother transitions. Non-conformal grids can provide more flexible refinement capabilities and, with simpler computational grid structures, run more efficiently on large parallel supercomputers (Jablonowski et al., 2009a). 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 non-conformal grids.

Use of AMR in a regional model paired with a physical parameterization package was presented in Bacon et al. (2000). OMEGA has been tested in forecasting hurricane storm tracks (Bacon et al., 2007). 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 non-hydrostatic 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

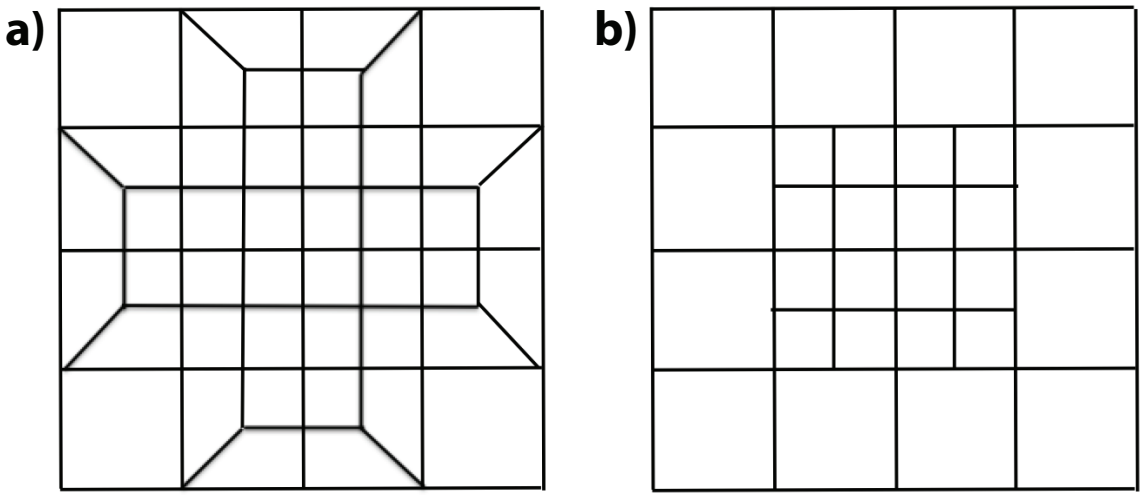

Figure 1.2: (a) Conformal grid — each cell shares an edge with exactly one other cell and (b) Non-conformal grid — cells can share an edge with more than one neighbor, which results in hanging nodes.

model on an f-plane with a constant Coriolis parameter f.

1.4.2.3 p-refinement: Adaptive Order Refinement

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). For smooth problems, p-refinement offers a high rate of convergence with significantly less computational expense than other refinement types, but it has difficulty resolving sharp gradients in the flow field. To ameliorate these problems, a hybrid refinement method that combines the h- and p-refinements methods was demonstrated by Eskilsson (2011). Blaise and St-Cyr (2012) used an hp-adaptive DG method to model the shallow water equations on a sphere for global tsunami simulations and Kopera and Giraldo (2014) implemented an DG method with hp-refinement in a two dimensional xz-projection of the compressible Euler equations.

1.5 Challenges associated with dynamic refinement

It has been nearly 30 years since the first adaptive atmospheric models were published (Skamarock et al., 1989; Dietachmayer and Droegemeier, 1992). Static refinement has been used for decades in regional weather models and VRGCMs have become a growing fixture in global climate modeling. Adaptive refinement for atmospheric modeling has remained mostly an academic exercise, with recent development centered on testing various refinement techniques in idealized one and two dimensional systems. Modeling the atmosphere is a far more complex problem than current applications of dynamic refinement as features of interest are more diverse and less well defined. To that end, there are several ongoing challenges that need to be addressed for AMR to become a tool for climate and weather modeling.

The key issue that most recent dynamical refinement work has been seeking to address is the development of robust numerical methods that run efficiently on today's supercomputers. A continuously changing grid structure sharply increases the difficulty in developing numerical schemes that are efficient, accurate, and stable. Dynamic refinement models still need to be locally conservative, preserve monotonicity, and minimize and manage numerical artifacts that arise from the grid refinement structure. Furthermore, any method must preserve consistency in the full equation sets. For example, at around 10 km resolution the hydrostatic approximation is no longer valid and models need to solve the more complex full non-hydrostatic equations.

Adaptive methods present more challenges for running simulations on current massively parallel supercomputers than traditional static models face. The processor and memory requirements of a dynamic refinement model change as grid elements are added and removed. Dynamic refinement techniques also use complex data structures requiring more computational overhead to organize the grid elements and communicate the grid changes. For dynamic refinement methods to become more mainstream,

they need to be more computationally efficient than current uniform models yet generate the same accuracy and skill. Both the block-structured (Jablonowski et al., 2009a; Chen et al., 2011a) and unstructured quad-tree refinements (Läuter et al., 2007; Blaise and St-Cyr, 2012) show promise as methods that ensure computationally efficient refinement on high performance computing platforms.

Another issue facing both dynamic and static VRGCMs is scale-aware physics parameterizations. In current uniform climate and weather models, the physics schemes that approximate the unresolved and sub-grid scale processes are specifically tuned to the model's resolution. Increasing resolution without re-tuning the physics parameterizations can adversely affect results (Rauscher et al., 2013; Zarzycki et al., 2014b). So, in adaptive schemes, the sub-grid parameterizations need to be able to adjust for changes in scale. A model would need to be able to phase out certain sub-grid processes, like deep convection, as resolution is increased and these processes are resolved on the actual mesh. Several parameterization schemes (e.g. the newly developed cloud scheme Cloud Layers Unified By Binormals (CLUBB) in the National Center for Atmospheric Research's Community Atmosphere Model) are being designed to work with varying vertical and horizontal resolutions.

A final issue in dynamic refinement is the selection of refinement criteria that determine where refinement is added. If a refinement criterion is too selective, key features will not be refined; if a criterion is too relaxed, refinement areas will become too large, making the model inefficient. Furthermore, the multiscale nature of the atmosphere and its interconnectivity creates a plethora of refinement criteria possibilities such that one set of criteria will not work effectively on all possible features of interest. Thus finding criteria that detect the areas that need to be refined as accurately and efficiently as possible is vital for effective implementation of adaptive meshes. A variety of criteria have been proposed and tested in the literature including local error estimates (Skamarock et al., 1989; Läuter et al., 2007; Bauer et al., 2014)

and physical variables like pressure or wind (Bacon et al., 2000; St-Cyr et al., 2008; Kopera and Giraldo, 2014).

Global dynamic refinement models retain potential as powerful tools for solving complex, nonlinear, multiscaled weather and climate systems. Dynamic refinement could benefit many applications involving transient localized phenomena including tropical cyclones, atmospheric blocking events, orographic precipitation, and pollution dispersion. Even long-term climate simulations could benefit as localized flow structures may be important in determining climate signals, such as modeling atmospheric rivers on the West coast of the United States.

1.6 Outline of Thesis

This thesis presents ongoing work to implement AMR within a new nonhydrostatic, finite-volume dycore and to demonstrate AMR's effectiveness in improving model accuracy and resolving multiscale features. These efforts are accomplished by implementing a hierarchy of new and existing test cases of increasing complexity for the 2D shallow water and 3D non-hydrostatic equations. The thesis is organized as follows. In Chapter II we implement existing shallow water test cases and develop several new ones to assess the capabilities of AMR in the shallow water version of the McCorquodale et al. (2015) finite volume Chombo AMR model. The results of this research were published in Ferguson et al. (2016). Chapter III presents two different forced shallow water systems with moisture variables designed to represent convective and precipitation processes as more challenging testbeds to use with the finite volume Chombo AMR model. In Chapter IV, we investigate the effects of AMR using the 3D non-hydrostatic version of the finite-volume Chombo AMR model in two 3D test cases. The first consists of a simple dry flow without forcing while the second adds a simplified physics parameterization scheme. Conclusions and future directions are presented in Chapter V.

CHAPTER II

Adaptive mesh refinement in 2D advection and shallow-water simulations

2.1 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 multi-scale 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, e.g., 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, multi-day time periods with grid spacings in the 0.9–3.5 kilometer 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 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 (non-moving) mesh adaptations are provided in, for example, Zarzycki et al. (2014a), Rauscher and Ringler (2014), Zarzycki and Jablonowski (2015) and Huang et al. (2016) (see also further references therein). This chapter 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), and Blaise and St-Cyr (2012)) or flow characteristics (e.g. Hubbard and Nikiforakis (2003), Jablonowski et al. (2006, 2009b) and 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 (see 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) uses 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 grid-cell 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 non-conformal 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 non-hydrostatic 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 non-conformal 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 wave-like reflections at AMR boundaries. The 2D shallow-water equations serve as a useful testbed 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, that 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. Lastly, 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 chapter is organized into three main sections. A brief description of the model and a discussion of the multiblock AMR techniques are provided in Section 2.2.

Section 2.3 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 lastly a test that assesses the interaction of idealized binary vortices. Conclusions are provided in Section 2.4.

2.2 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 shallow water equations on the sphere in coordinate invariant form are given by

$$\frac{\partial h\mathbf{v}}{\partial t} + \nabla \cdot \left(h\mathbf{v}\mathbf{v} + \mathcal{I}\frac{gh^2}{2} \right) = -gh\nabla z_{\mathrm{s}} - f\mathbf{k} \times (h\mathbf{v}) \tag{2.1}$$

$$\frac{\partial H}{\partial t} + \nabla \cdot (h\mathbf{v}) = 0 \tag{2.2}$$

where $\mathbf{v}$ is the velocity vector, $\mathbf{v}\mathbf{v}$ denotes the outer product of the velocity vector, $\mathcal{I}$ is the identity matrix, $f = 2\Omega \sin\phi$ is the Coriolis parameter in terms of the angular rotation $\Omega = 7.292 \times 10^{-5}$ s^{-1}, and $g = 9.80616$ m s^{-2} is the acceleration due to gravity. H denotes the total height of the fluid surface. Bottom topography z_s and h, the fluid depth above bottom topography, relates to the total height by $H = h + z_s$.

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

Table 2.1: 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_{min}/A_{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 RLL_{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_{min}/A_{max}	Eq. Res.	RLL_{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°

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 2.1 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 which 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, 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 (Sec. 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.

2.2.1 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 two. The finer levels are overlaid on top of the coarser levels and are organized in the block-structure described in the previous section. Figure 2.1a provides an

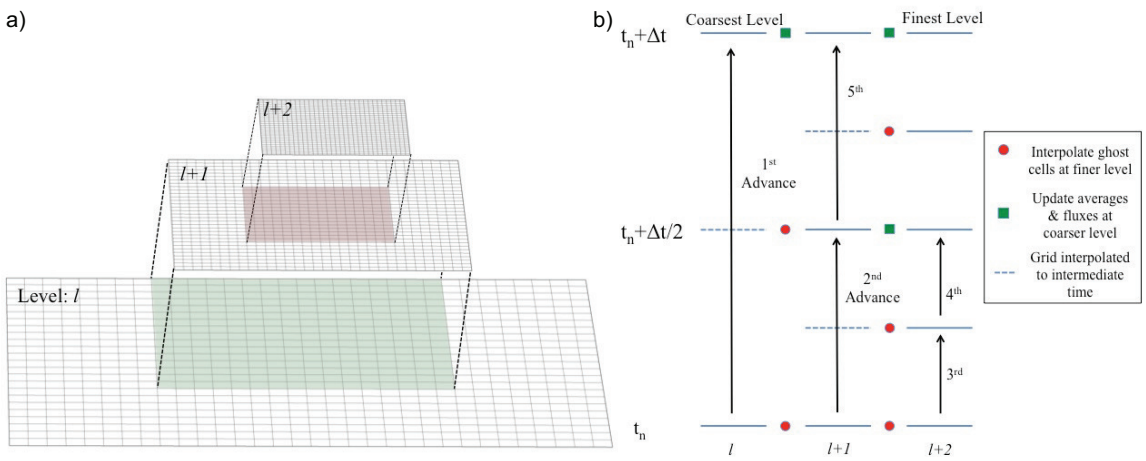

Figure 2.1: (a) Example of a grid with two levels of AMR using a x2 refinement ratio between levels. The additional levels do not replace the coarse level cells underneath, they are overlaid on top. (b) Diagram depicting the sub-cycling of AMR levels in time. The coarse level is advanced first from time 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.

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 calculated 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 2.1b depicts an overview of the time-stepping and sub-cycling process. Rather than having a uniform time step dependent on the smallest grid-cell size, Chombo-AMR sub-cycles 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 non-linear 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 which 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 Sec. 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 2.2 depicts the location of ghost cells for two of the three layers and the stencil of coarse grid cells used for their interpolation.

4. Perform previous steps for level $l+1$. Level $l+1$ is advanced using refined time steps (sub-cycling) as depicted in Fig. 2.1b. 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 sub-cycling 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).

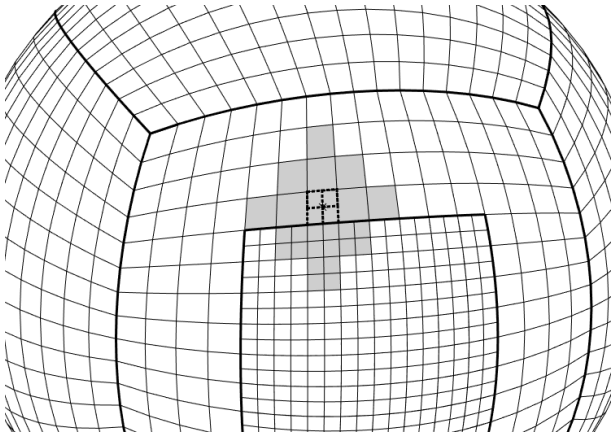

Figure 2.2: Example of a 1-level AMR approach with a refinement ratio of two on the cubed sphere grid. The four dashed cells are schematic examples of ghost cells for the finer grid, which are interpolated from the values in the coarse cells shaded in gray. Our actual AMR applications need three ghost cells.

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 chapter, 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 ten), and the maximum number of levels is set for each simulation. In our simulations, we explore the effects of adding up to two refinement levels (called 2-level AMR). In addition, we explore the grid-resolution refinement ratios between successive levels, and present selected results for three powers of two: x2, x4 and x8. As an example, a grid with a base level of c32 (2.8°) resolution and 2 levels of x4 refinement has a maximum resolution of c512 (0.18°).

2.3 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, non-linear, localized flows. The first category is comprised of the following two shallow-water tests:

- a steady-state geostrophic flow (test case 2 in Williamson et al. (1992))
- an unsteady solid body rotation (example 3 from Läuter 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.
- a binary-vortices test case in which two vortices interact.

The bottom two 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.

The model results are presented using normalized l_2 and l_∞ error norms as defined in Williamson et al. (1992). These errors for computed variable field h are calculated via the usual global error norms,

$$l_2(h) = \sqrt{\frac{I\left[(h - h_\tau)^2\right]}{I\left[h_\tau^2\right]}} \tag{2.3}$$

$$l_\infty(h) = \frac{\max|h - h_\tau|}{\max|h_\tau|} \tag{2.4}$$

where h_τ is the reference field and I is a discrete approximation to the global integral

given by

$$I[x] = \sum_{\text{all cells } k} x_k A_k \tag{2.5}$$

where A_k denotes the area of grid cell k.

2.3.1 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 which 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, which 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 $\pm 45°$). Figures 2.3a–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 $|\nabla q| > 1.5 \times 10^{-7}$ 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

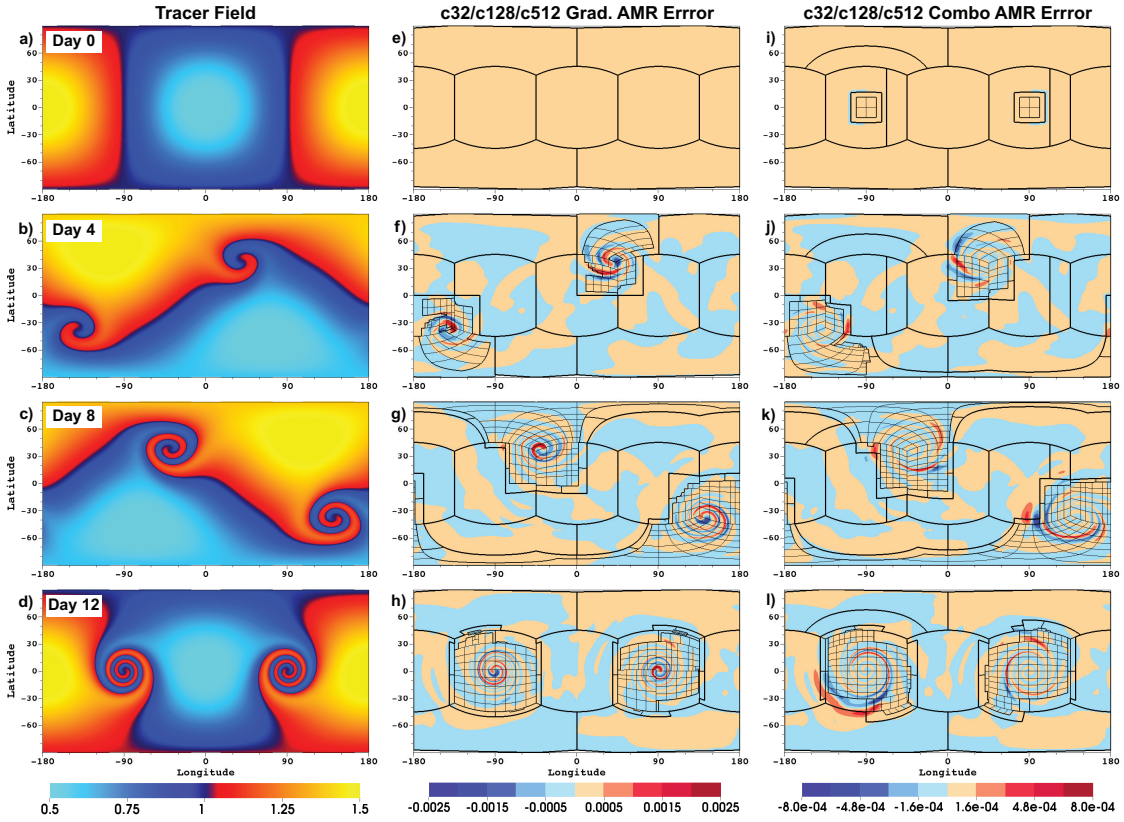

Figure 2.3: (a)–(d): Analytic tracer field at days 0, 4, 8, and 12 for the moving-vortices advection test of Section 2.3.1. (e)–(h): Tracer error at the select days for a 2-level AMR run with a c32 base resolution using x4 refinement (c32/c128/c512) and the tracer-gradient refinement tagging criterion. (i)–(l): same as (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 (e)–(h) have different contour scale than (i)–(l).

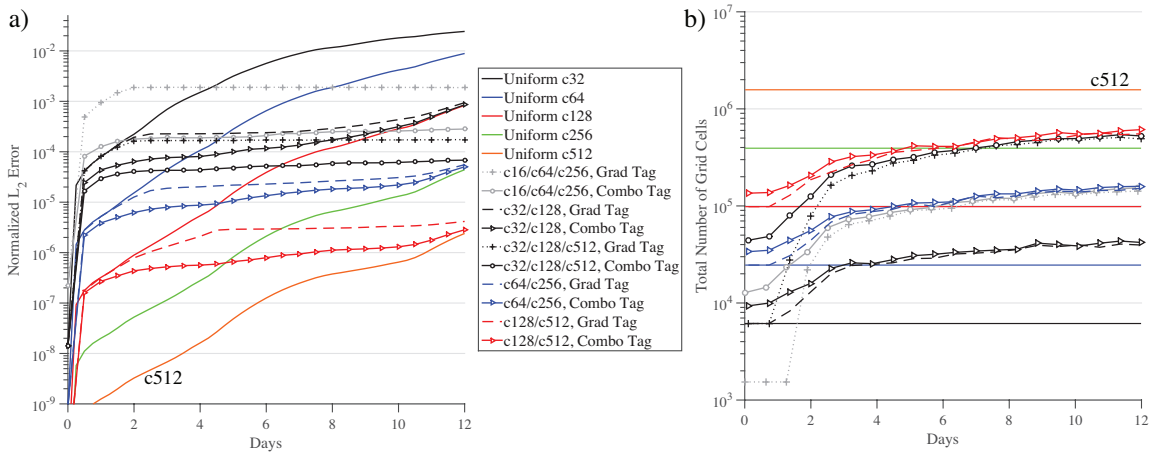

Figure 2.4: For the moving-vortices advection test of Section 2.3.1, growth over time of: (a) normalized l_2 tracer error, and (b) total number of grid cells. The plots provide a comparison of uniform runs (solid lines, no markers) and AMR runs using gradient tagging (broken lines, with and without markers) and combination of relative-vorticity and tracer-gradient tagging (solid lines with markers). All AMR runs use the x4 refinement ratio between resolution levels.

initial refinement around the vorticity patches without being too low as to refine on the background vorticity. With both criteria, a x4 refinement ratio between levels was used. The evolution of the grids can be seen in Fig. 2.3, which shows the tracer error for the two AMR runs. Figures 2.3e–h depict the 2-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 2.3i–l depict the same 2-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. 2.3. Each patch contains a user-selected maximum number of grid cells. However, patches might also be split into smaller blocks by the Chombo library. 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. 2.3. For the tracer-gradient tagging scheme (Figs. 2.3e–h), the error is largest near the center of the spiral. The combination tagging scheme (Figs. 2.3i–l) significantly reduces the error near the cen-

ter, and the largest error is now towards 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 one, 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. 2.3i), which limits the error growth. This effect can also be seen in Fig. 2.4a 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. 2.4a compares the l_2 tracer errors of the five uniform-resolution simulations c32, c64, c128, c256, and c512 (see Table 2.1 for the associated grid spacings) to various 1-level and 2-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. 2.4b. 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 grid-cell count for only the first few days (see Fig. 2.4b). The errors in all single-level AMR runs in Fig. 2.4a 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 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 2-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. 2.5. AMR runs that

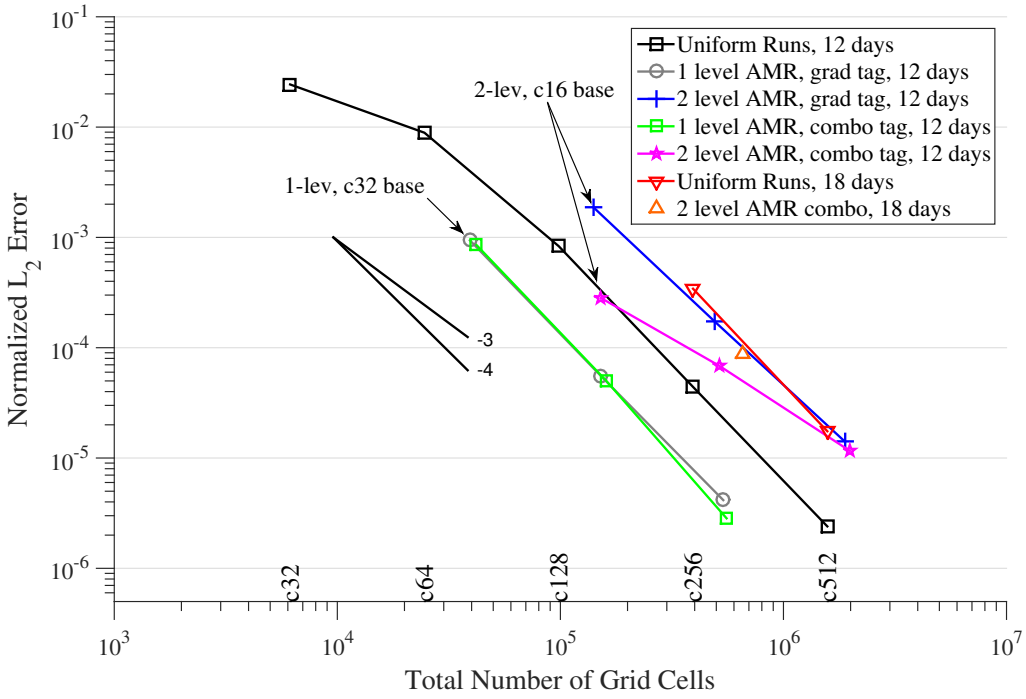

Figure 2.5: Normalized l_2 gradient error as a function of total number of grid cells for the moving-vortices advection test case of Section 2.3.1 at day 12. The grid-resolution labels along the bottom axis note the number of grid cells in the uniform grids with those resolutions. The black square markers are for the uniform runs c32 through c512. The gray circles and green squares represent the 1-level x4 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 2-level x4 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 2-level x4 refinement runs with a c32 base resolution, using the combination tagging at day 18. Solid black lines depict convergence rates.

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 1-level AMR runs show a decrease in error while using fewer grid cells, the 2-level AMR runs generally result in a higher error for a comparable number of grid cells than the uniform runs. Only the 2-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 (c128/c512, for example). Extending the run time to 18 days, the l_2 error from the c32 base 2-level AMR run with the combination tagging (orange triangle in Fig. 2.5) is now slightly lower than the error from the uniform runs at day 18 (red downward-pointing triangles). Figure 2.5 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 $\alpha = 0°$ setup (not $\alpha = 45°$ as we present here) in an FV-AMR model using a coarse 5° base grid and one to three levels of x2 refinement which were guided by a 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 3-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 ($\sim 1.4°$) run with $\alpha = 45°$ has a normalized l_2 error of 8.9×10^{-3} after twelve days, which is comparable to the results from the 1.25° grid using a multi-moment 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

Table 2.2: Run times (wall-clock time in s and as % of c512 run time) for 12-day moving-vortices advection simulations (Section 2.3.1). 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 Res.	AMR Levels	Run Time (s)	Run Time (%)	c32 cells	c128 cells	c512 cells	Total cells (%)
c512	-	24152	100%	0	0	1.6×10^6	100%
c128	1	10116	42%	0	9.8×10^4	4.5×10^5	34%
c32	2	10468	43%	6.1×10^3	4.2×10^4	4.6×10^5	32%

grid cells in Table 2.2, for some of the 12-day runs in Fig. 2.5. 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 2-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.

2.3.2 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. The steady state background velocity field in latitude θ and longitude λ coordinates is

$$u = u_0 \left(\cos\theta \cos\alpha + \cos\lambda \sin\theta \sin\alpha \right) \tag{2.6}$$

$$v = -u_0 \sin\lambda \sin\alpha \tag{2.7}$$

where the background velocity $u_0 = \pi/6$ Earth radii day^{-1} = 38.61 m s^{-1}. The analytic height field is given by

$$h = h_0 - \frac{1}{g}\left(\Omega u_0 + \frac{u_0^2}{2}\right)(-\cos\lambda\cos\theta\sin\alpha + \sin\theta\cos\alpha)^2 \tag{2.8}$$

with background height $h_0 = 2998$ m and planetary constants $a = 6.37122 \times 10^6$ m, $\Omega = 7.292 \times 10^{-5}$ s^{-1}, and $g = 9.80616$ m s^{-2}.

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. 2.6. The first configuration (Fig. 2.6b) consists of a statically-refined patch centered at zero degrees 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. 2.6c), we place static patches over the locations of high relative vorticity, refining where $|\zeta| > 1.18 \times 10^{-5}$ 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 x2, x4, and x8 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

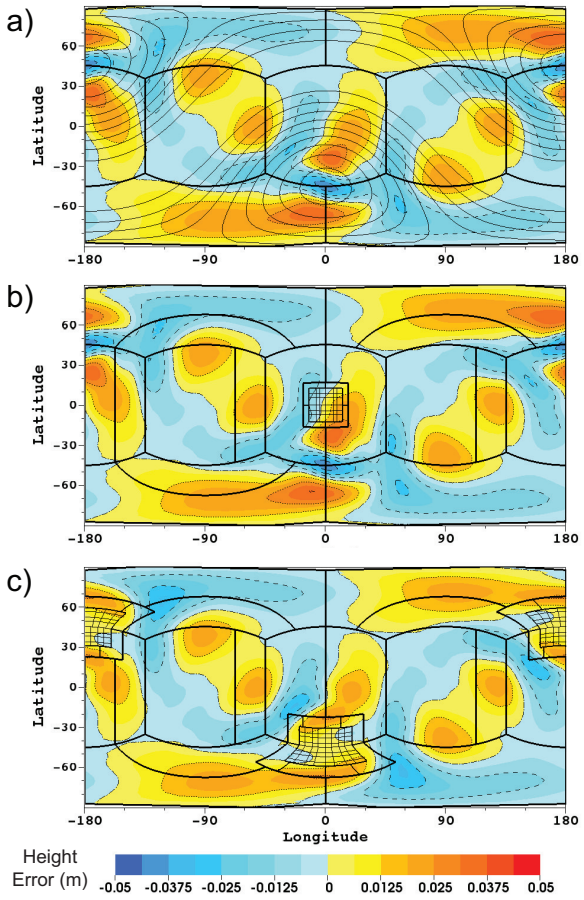

Figure 2.6: Height error (in meters) at day 5 for the steady-state geostrophic flow test case of Section 2.3.2. The configurations are: (a) a c32 uniform-resolution run, (b) a c32 grid with a static equatorial patch using two levels of x4 refinement, and (c) a c32 grid with static midlatitudinal patches using two levels of x4 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.

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 2.3 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 five 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 x8 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 x8 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 x4 additional refinement (Fig. 2.6a-c) depict a similar result. The equatorial patch run (Fig. 2.6b) has essentially the same height error profile as the uniform run (Fig. 2.6a). The height error for a base c32 run with the midlatitudinal refinement patches (Fig. 2.6c) 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 2.7 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

Table 2.3: Global steady-state geostrophic flow test of Section 2.3.2: 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 $l_\infty = 1.4505 \times 10^{-5}$.

		Equatorial Refinement		Midlatitudinal Refinement	
AMR Levels	Ref. Ratio	l_2 error	l_∞ error	l_2 error	l_∞ error
1	x2	5.4848×10^{-6}	1.4519×10^{-5}	4.7472×10^{-6}	1.0077×10^{-5}
1	x4	5.4837×10^{-6}	1.4542×10^{-5}	4.7835×10^{-6}	9.5043×10^{-6}
1	x8	5.4875×10^{-6}	1.4592×10^{-5}	4.8233×10^{-6}	9.6988×10^{-6}
2	x2	5.4849×10^{-6}	1.4520×10^{-5}	4.5341×10^{-6}	1.0559×10^{-5}
2	x4	5.4836×10^{-6}	1.4542×10^{-5}	4.7983×10^{-6}	1.1019×10^{-5}

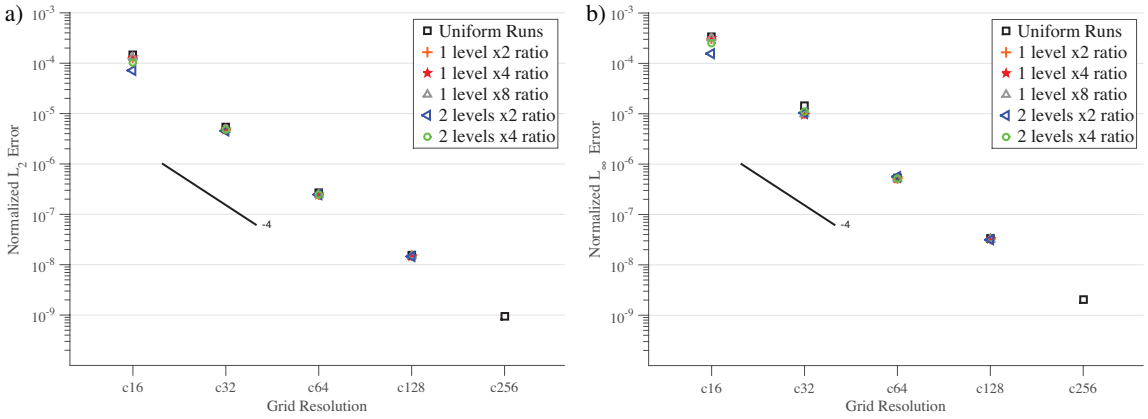

Figure 2.7: (a) Normalized l_2 and (b) l_∞ height errors at day 5 as a function of base grid resolution for the steady-state geostrophic flow test case of Section 2.3.2. Uniform runs and runs using the static midlatitudinal refinement patches are depicted with x2, x4, or x8 refinement ratios. The fourth-order convergence rate is shown by the black line.

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 2.3 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 x2 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 x3 refinement ratio. After five 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.

2.3.3 Unsteady Solid-Body Rotation

The time-dependent zonal flow test proposed in Läuter 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 the 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 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), the unsteady analytic solution can be written in latitude-longitude (θ, λ) coordinates as

$$u = u_0 \left(\sin\alpha \sin\theta \left(\cos\lambda \cos\Omega t - \sin\lambda \sin\Omega t\right) + \cos\alpha \cos\theta\right) \tag{2.9}$$

$$v = -u_0 \sin\alpha \left(\sin\lambda \cos\Omega t + \cos\lambda \sin\Omega t\right) \tag{2.10}$$

$$\begin{aligned} h = &-\frac{1}{2g}[u_0 \left(\sin\alpha \cos\theta \left(-\cos\lambda \cos\Omega t + \sin\lambda \sin\Omega t\right) + \cos\alpha \sin\theta\right) \\ &+ a\Omega \sin\theta]^2 + \frac{1}{2g}\left(a\Omega \sin\theta\right)^2 + k_1 \end{aligned} \tag{2.11}$$

$$h_b = \frac{1}{2g}\left(a\Omega \sin\theta\right)^2 + k_2 \tag{2.12}$$

The solutions for the zonal u and meridional v velocities are dependent on time t, Earth's angular velocity Ω, the velocity constant $u_0 = 2\pi a/12$ m day^{-1} = 38.61 m s^{-1}, and the the solid body axis of rotation angle α. The height field h is also dependent on the Earth radius $a = 6.37122 \times 10^6$ m and constant $k_1 = 1.362 \times 10^4$ m. The surface topography h_b has the constant offset k_2 which is set to zero for these simulations. 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 one day for better comparison of the results, the Earth's angular velocity is slightly modified to be based on a solar day instead of sidereal day, so that the angular velocity is $\Omega = 2\pi/86400$ s^{-1} $\approx 7.2722 \cdot 10^{-5}$ 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 Sec. 3b now centered at 0°N, 90°E.

2. Dynamic AMR refinement with height tag, the threshold for the free surface height is $h > 1.38 \times 10^4$ m.

3. Dynamic AMR refinement with relative vorticity tag, the threshold is $|\zeta| > 1.18 \times 10^{-5}$ s^{-1}.

The three grid configurations can be seen in Figs. 2.8b-d, which show the patches at the identical initial and final (day 5) positions. The second and third configurations (Figs. 2.8c,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 x2, x4, and x8 refinement ratio simulations, so we only discuss runs with the x4 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 five days. The normalized global l_2 and l_∞ height errors after five days are shown for simulations with a c32 base grid in Table 2.4. Errors are calculated by comparing runs to the analytic solution of the test case. The uniform grid results are compared with 1- and 2-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 2-level height-tag AMR run, we observe that the l_2 height error increases by about 12%, while the 2-level vorticity-tag AMR run decreases the error by roughly 23%. Figure 2.8 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. 2.8b), the height errors remain nearly the same as for the uniform run (Fig. 2.8a), with only very slight increases in the large error areas on the polar panels. Along the

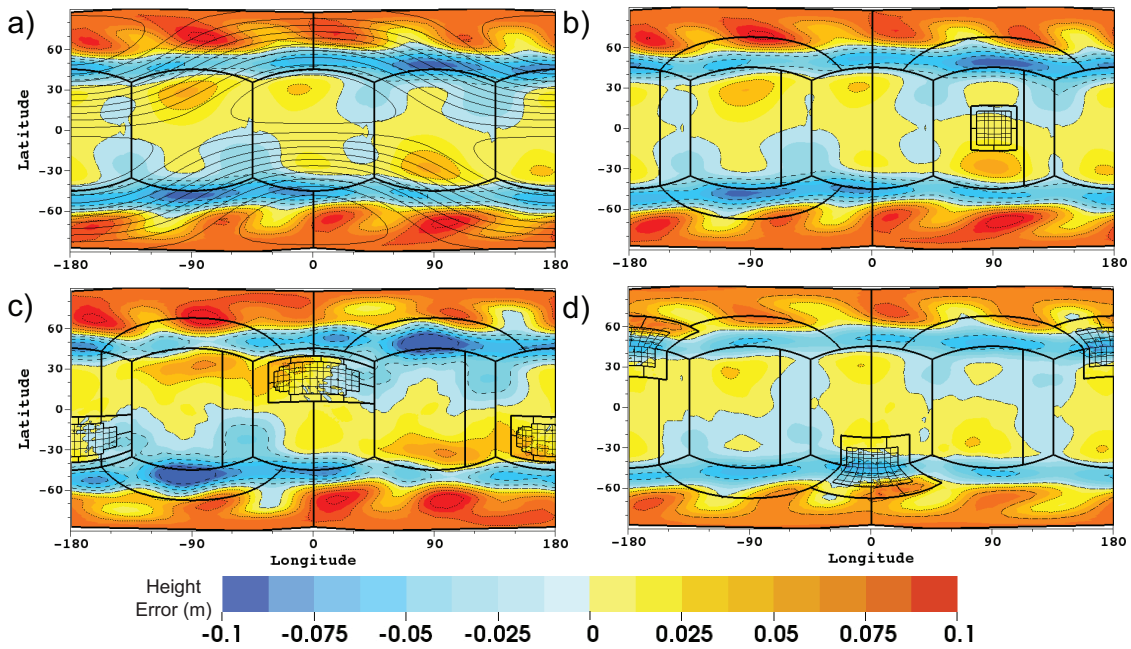

Figure 2.8: Height field errors (in meters) at day 5 for four simulations of the unsteady solid-body rotation test case of Section 2.3.3: (a) a c32 uniform-resolution run, (b) a c32 base grid with a static equatorial patch using two levels of x4 refinement, (c) a c32 base grid with two levels of dynamic x4 refinement using the height-tag criterion, and (d) a c32 base grid with two levels of dynamic x4 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.

coarse-fine boundary, no spurious grid-induced error is observed. In the height-tag AMR run (Fig. 2.8c) the errors along the polar-equatorial panel boundaries increase, while in the vorticity-tag AMR run (Fig. 2.8d), 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 ($\sim$ 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

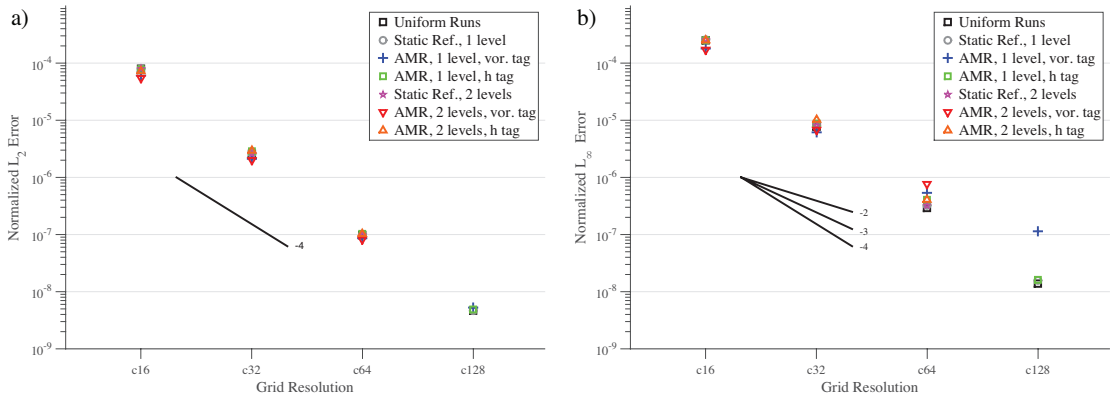

Figure 2.9: (a) Normalized l_2 and (b) l_∞ height errors at day 5 as a function of base grid resolution for the USBR test case of Section 2.3.3. Uniform runs and AMR runs with one and two refinement levels using the height-tag or vorticity-tag criteria. All AMR runs use x4 refinement ratio between levels. Solid black lines depict convergence rates. Errors are determined with respect to the analytic solution.

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 hrs, after which the flow features have progressed only half way around the sphere. Pudykiewicz (2011) showed a normalized l_2 height error of $\sim 6 \times 10^{-6}$ after 12 hours using a second-order icosahedral geodesic model with a 2° ($\sim$ 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 hrs. These results are comparable with those obtained by the fourth-order multi-moment 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 hours. Additionally, the third-order multi-moment 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 of $\sim 6.5 \times 10^{-7}$ at 12 hrs. 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 ($\sim$ 78 km spacing) with the two dynamic grid configurations. The normalized l_2 and l_∞ height errors after five days are plotted in Figs. 2.9a and 2.9b, 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. 2.9a). The fourth-order convergence is maintained for all grid configurations. In the l_∞ height error plot (Fig. 2.9b), 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. 2.8c-d. 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. 2.8c. 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. 2.9b 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.

2.3.4 Isolated Mountain Gravity Wave

This shallow-water test was developed to assess AMR when topography is present. In the test a gravity wave, that 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

$$z_s = \frac{z_0}{4}\left(1 + \cos\left(\frac{\pi r}{R}\right)\right)^2 \tag{2.13}$$

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

$$h = h_0 - h_{max} \exp\left(-\left(\frac{\tau}{\beta}\right)^2\right). \tag{2.14}$$

The maximum depth of the perturbation is set to $h_{max} = 100$ m, $\beta = \pi/36$ is a width parameter, and τ is the great-circle distance from point (λ, θ) to the dip's center (λ_d, θ_d) such that

$$\tau = a \arccos\left(\sin\theta_d \sin\theta + \cos\theta_d \cos\theta \cos\left(\lambda - \lambda_d\right)\right) \tag{2.15}$$

where $a = 6.37122 \times 10^6$ m is the average radius of the 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 2.10 depicts the initial height field and its distance from the mountain. The simulation is run for a period of twelve hours so that the gravity wave

Table 2.4: Day 5 normalized l_2 and l_∞ height error norms for the unsteady solid body rotation test of Section 2.3.3. 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 x4 refinement ratio.

Grid Config.	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}

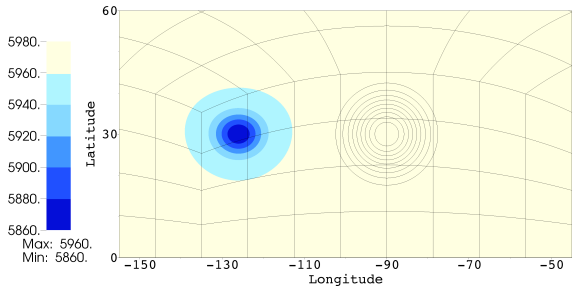

Figure 2.10: Initial free surface height field (colored, in m) for the gravity wave over an idealized mountain test of Section 2.3.4. The free surface height field (above the reference sphere at sea level) is uniform everywhere except for the 100 m deep Gaussian depression. The black contour lines represent the location of the mountain with 200 m contour spacing and a peak mountain height of 2000 m.

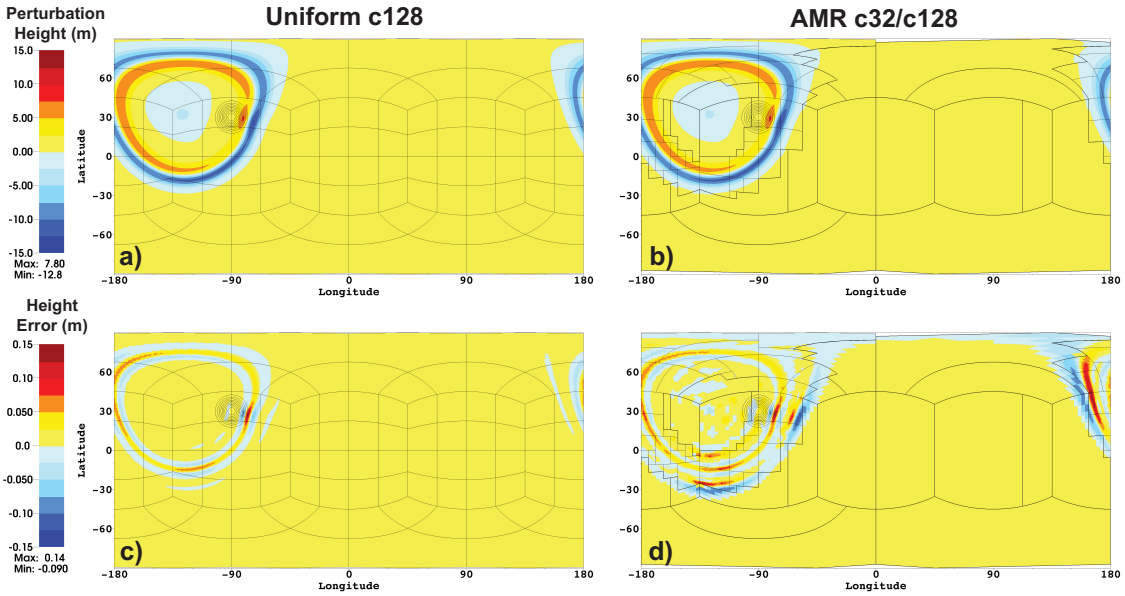

Figure 2.11: Mountain gravity-wave test of Section 2.3.4, at hour 6. (a) and (b) depict the perturbation height of the gravity wave as it passes over the mountain for a uniform c128 run and a c32/c128 AMR run with the height gradient tag $|\nabla h| > 7.5 \times 10^{-6}$, (c) and (d) depict the height error of each run after six hours compared to a reference uniform c1024 run. The block structure of the grid and the mountain contours are overlaid with thin black lines.

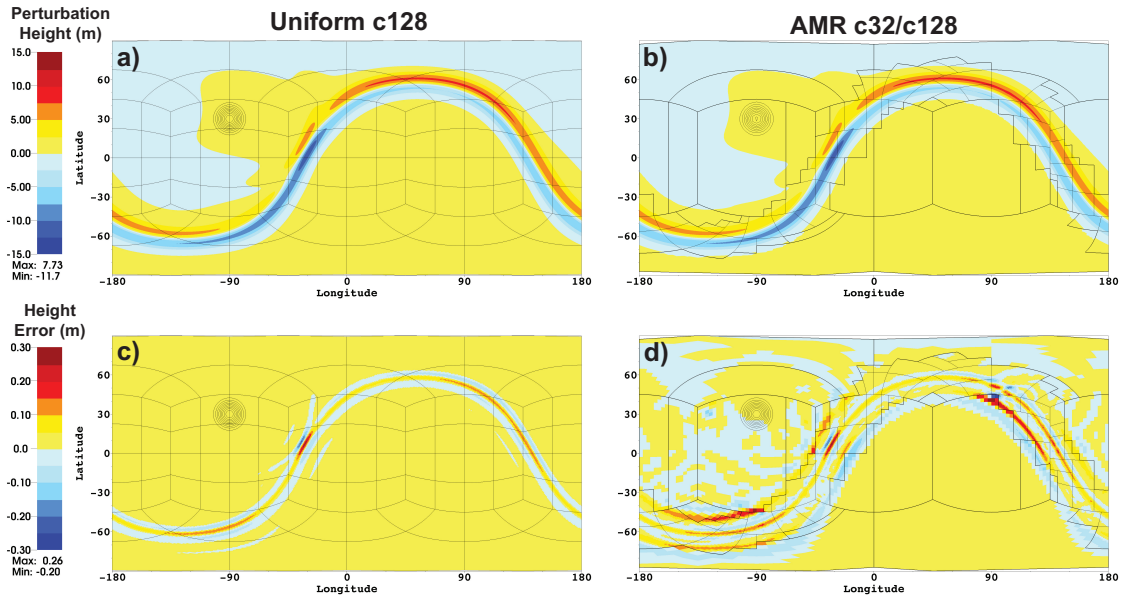

Figure 2.12: Same runs and plots as in Fig. 2.11 except for hour 12 as the gravity wave has moved halfway around the sphere.

has moved halfway around the sphere. Figures 2.11 and 2.12 show the perturbation height (defined as deviations from h_0, top panels) and the height difference from a uniform c1024 ($\sim$ 10 km) reference solution (bottom panels) at hour 6 and hour 12, respectively, for a uniform c128 run and a c32 base 1-level AMR run (c32/c128) tagged on a height-gradient threshold of $|\nabla h| > 7.5 \times 10^{-6}$. After six hours (Fig. 2.11), 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. 2.11b, d. The height differences for the c32/c128 AMR run (plot (d) in Figs. 2.11 and 2.12) are similar in position and magnitude to the uniform c128 run errors (plot (c)) within the refined AMR domain. The areas of larger error at the borders of the AMR region seen in Figs. 2.11d and 2.12d are located over the leading and trailing edges of the gravity wave, which are not fully covered by the AMR refinement criterion. The

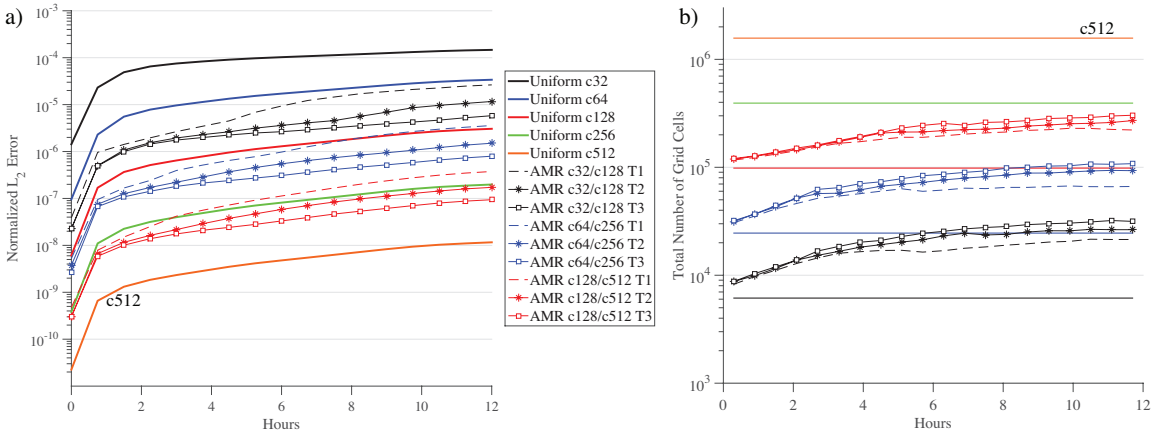

Figure 2.13: For the mountain gravity-wave test of Section 2.3.4, growth over the twelve-hour period of: (a) normalized l_2 height error with respect to the uniform c1024 run, and (b) total number of grid cells. Error and number of grid cells for uniform-resolution runs and 1- and 2- level AMR runs with height-gradient tagging. The thresholds are $|\nabla h| = 1.5\times10^{-5}$ (T1), 1.0×10^{-5} (T2), and 7.5×10^{-6} (T3).

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 again covered by only the coarse grid. The AMR reinitializes the topography (using Eq. (2.13)) whenever adaptations are triggered. No spurious errors appear as the AMR refines and coarsens over the topography region (Figs. 2.11d and 2.12d). Results in Fig. 2.13 show the normalized l_2 height error (Fig. 2.13a) and the total number of grid cells (Fig. 2.13b) as a function of forecast hour for uniform runs from c32 to c512 and 1-level AMR runs using height-gradient tagging with thresholds of $|\nabla h| > 1.5\times10^{-5}$, 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 1-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 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.

2.3.5 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 post-interaction 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 multi-scale 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

$$\phi = \bar{\phi} - \phi' \tag{2.16}$$

$$\phi' = \phi_c \left(1 - \exp\left(-\left(\frac{r_m}{r}\right)^b\right)\right) \tag{2.17}$$

where $\bar{\phi} = g\bar{h}$ is the background geopotential with the background height $\bar{h} = 4200$ m and the Earth's gravity $g = 9.80616$ 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. (2.15)). 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

$$v_T = -\frac{rf}{2} \pm \sqrt{\frac{r^2 f^2}{4} + r\frac{\partial \phi}{\partial r}}. \tag{2.18}$$

represents the initial axisymmetric flow in gradient wind balance. Using $\partial\phi/\partial r$ derived from Eq. (2.16), we get the corresponding tangential velocity for a cyclonic (+ sign in Eq. (2.18)) circulation

$$v_T = \sqrt{\phi_c b \left(\frac{r_m}{r}\right)^b \exp\left(-\left(\frac{r_m}{r}\right)^b\right) + \frac{r^2 f^2}{4}} - \frac{rf}{2} \tag{2.19}$$

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

$$u = v_T \frac{d_1}{d} \quad \text{and} \quad v = v_T \frac{d_2}{d} \tag{2.20}$$

where

$$d_1 = \sin\theta_c \cos\theta - \cos\theta_c \sin\theta \cos(\lambda - \lambda_c) \quad (2.21)$$

$$d_2 = \cos\theta_c \sin(\lambda - \lambda_c) \quad (2.22)$$

$$d = \max\left(\epsilon, \sqrt{{d_1}^2 + {d_2}^2}\right). \quad (2.23)$$

The threshold value $\epsilon = 10^{-25}$ 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. 2.14a. Here, all magnitudes are normalized to one 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. 2.14b. It shows that the tangential wind profile from Eq. (2.19) 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. (2.17). Here, we use the sum of the u, v and ϕ' values of both vortices to initialize our shallow-water system.

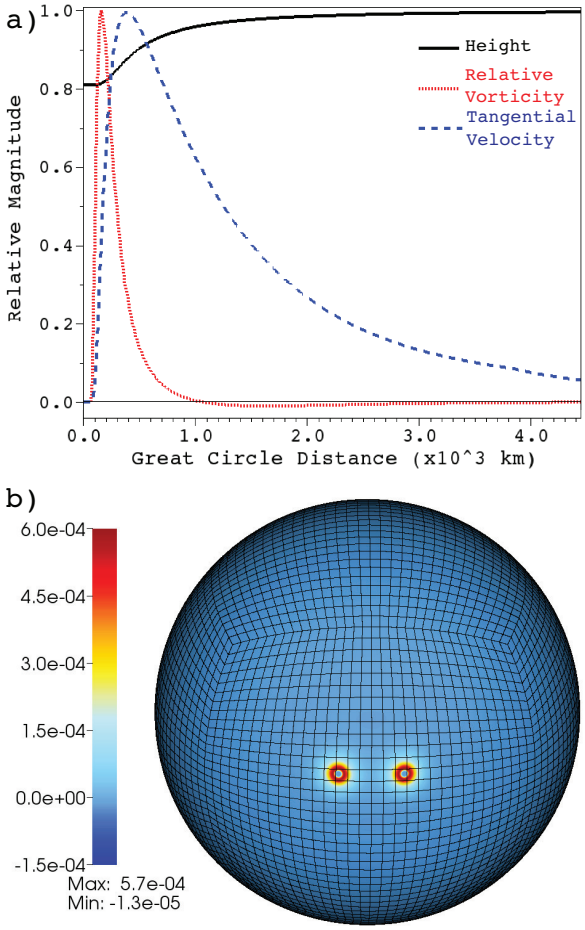

Figure 2.14: Initial conditions for the binary-vortices test case of Section 2.3.5. (a) Radial profiles of relative vorticity (red), tangential wind (blue) and height field (black) as a function of great-circle distance from the center for a single vortex. Profiles are scaled to the maximum of each value (see text). (b) Initial profile of relative vorticity (s^{-1}) of the two vortices on a cubed-sphere grid (merger test case).

2.3.5.1 Separation Case

In the separation case, the two vortices are centered at $\theta_c = \pi/36 = 5°\text{N}$ which defines the constant Coriolis parameter $f = 2\Omega \sin\theta_c$ with the Earth's angular velocity $\Omega = 7.292 \times 10^{-5}$ s^{-1}. The maximum geopotential perturbation is set to $\phi_c = gh_c$ with the maximum height perturbation $h_c = 800$ m. The radius of maximum wind is set to $r_m = 250$ km. This results in a maximum tangential wind of 64 m s^{-1} and a maximum relative vorticity of 9.4×10^{-4} s^{-1}. The centers of the two vortices are 13.5° apart ($\sim$ 1500 km) in the longitudinal direction, six times the radius of maximum wind, so that their negative vorticity regions still overlap. In particular, the two vortex centers are located at $\lambda_{c_1} = (3\pi/2 - 6.75\pi/180)$ and $\lambda_{c_2} = (3\pi/2 + 6.75\pi/180)$ 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 three days of the simulation, the two vortices make one complete orbit around each other as the beta-drift steers them towards the northwest, after which the two vortices then 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 anti-cyclone. 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. 2.15, which shows the vorticity field of the cyclone pair at day 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 through c1024, of which the c128, c256, and c1024 runs are depicted in Fig. 2.15. 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. 2.15k–o), the vortices start separating earlier and at the end of the run are in markedly different positions. The

c256 simulation (Figs. 2.15f–j) more closely resembles the highest resolution c1024 run (Figs. 2.15a–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 fine-scale vorticity filaments. A comparison of the uniform c512 run and c1024 run at day 6 can be seen in Figs. 2.16e and 2.16i. To assess the AMR performance, we ran the model using relative vorticity refinement criteria with one and two levels of AMR and x4 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. 2.16. They are divided into columns that share the same highest resolutions, e.g. in the leftmost column, Fig. 2.16a is the uniform c256 run, while Figs. 2.16b-d are for 1- or 2-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 anti-cyclonic filaments and patches around the cores, and the wave train developing to the lee of the vortices with only minor differences in the anti-cyclone 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.

2.3.5.2 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

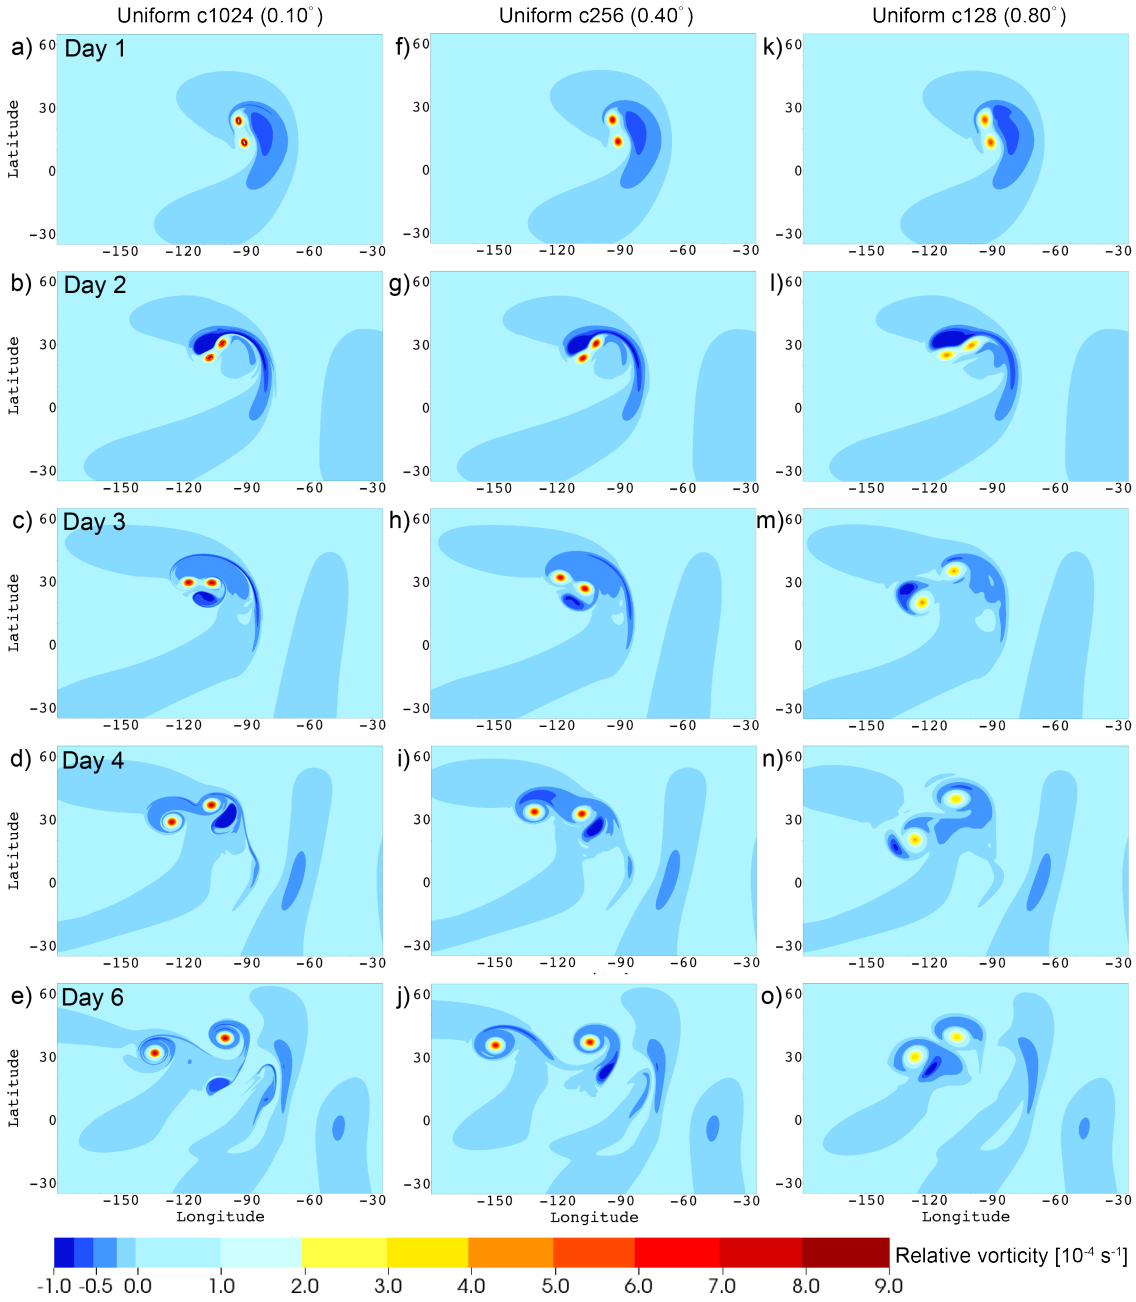

Figure 2.15: Evolution of the relative vorticity of uniform-resolution runs for the binary-vortices test (separation case) of Section 2.3.5.1. (a)–(e) Results for the uniform c1024 run for days 1, 2, 3, 4, and 6. (f)–(j) The uniform c256 run and (k)–(o) uniform c128 run.

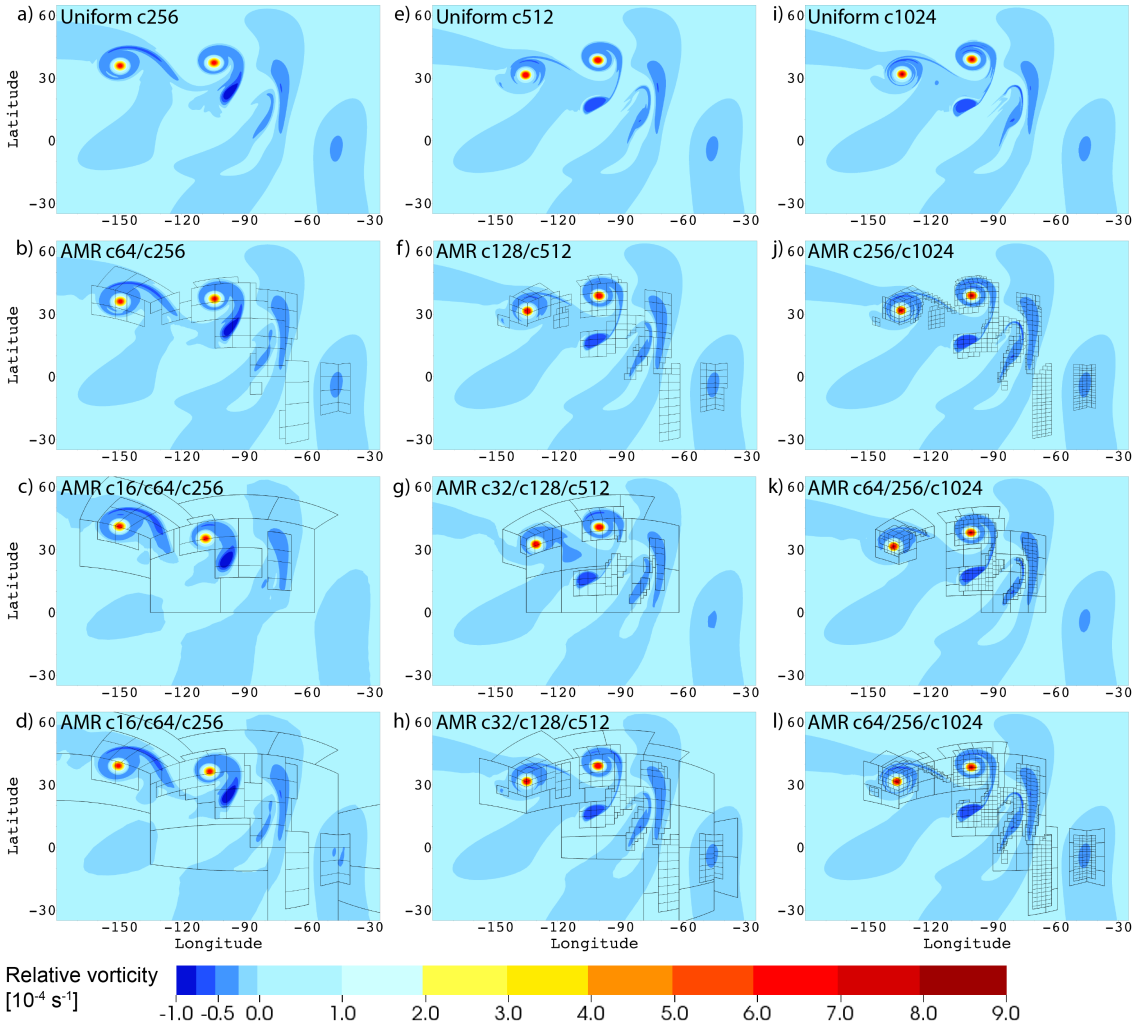

Figure 2.16: Relative vorticity fields at day 6 for several runs of the vortex separation case of Section 2.3.5.1 using AMR. (a) uniform c256 run, (e) uniform c512 run, and (i) uniform c1024 run. (b), (c), and (d) are AMR runs whose highest refinement level is c256. (f) through (h) and (j) through (l) are AMR runs with a highest refinement level of c512 and c1024, respectively. (b), (f), and (j) depict 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), and (k) depict AMR runs using two levels of refinement using a relative-vorticity threshold of $|\zeta| > 3.5 \times 10^{-5}\text{s}^{-1}$. The last row, (d), (h), and (l) depict AMR runs using the same refinement criteria as the second row but with two levels of refinement. All AMR runs use the x4 refinement ratio. The block structure of refinement levels 1 and 2 are outlined in black.

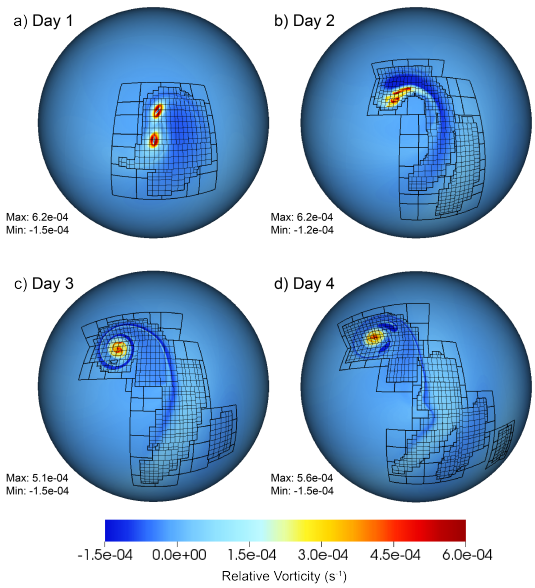

Figure 2.17: Evolution of the merging vortices in the test of Section 2.3.5.2 in a 2-level AMR run (c64/c256/c1024) using a relative-vorticity threshold refinement criterion. Refinement occurs when the absolute value of the relative vorticity is greater than $2.3 \times 10^{-5}\text{s}^{-1}$. Snapshots of relative vorticity at day (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.

relative vorticity is around 5.7×10^{-4} s^{-1}. The vortices are centered at $\theta_c = \pi/18 = 10°$N where the constant Coriolis parameter f is evaluated for this test case. The vortex centers centers are 15.65° apart ($\sim$ 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 2.17 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. 2.17 has a c64 base resolution with two levels of x4 refinement so that its finest level has a c1024 resolution. The AMR is triggered when the absolute value of the relative vorticity exceeds 2.3×10^{-5} s^{-1}. With that relatively low refinement threshold, the AMR captures not only the main vortex cores, but also the small-scale anti-cyclonic filament that extends far south of the merged vortex, and the small-magnitude wave train that develops by day 4. Figure 2.18 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 three 2-level AMR runs with x4 refinement on c64 base resolution that have a finest resolution of c1024. The second column has the c512 uniform run and three 2-level AMR runs with a c32 base. The last column has the uniform c256 run and three 2-level AMR runs with a c16 base. The three AMR refinement criteria in Fig. 2.18 are

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. 2.18)

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 $|\zeta| > 3.5 \times 10^{-5}$ 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. 2.18b and 2.18f, 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 anti-cyclonic filament wrapping around the new vortex and extending down from it as well as the development of the secondary lee side 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. 2.19. 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. 2.19a) and total number of grid cells (Fig. 2.19b) as a function of forecast days for uniform-resolution runs and 1- and 2- level AMR runs using the large-height-gradient tag or the vorticity tag with x4 refinement ratios. The vorticity-tag AMR simulations (both 1- and 2-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. 2.18b, f, and j). Although the large-scale shape and locations of the two merging vortices and the post-merger 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 non-linear 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 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 2.5. The table thereby represents some of the runs from Fig. 2.19b. 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

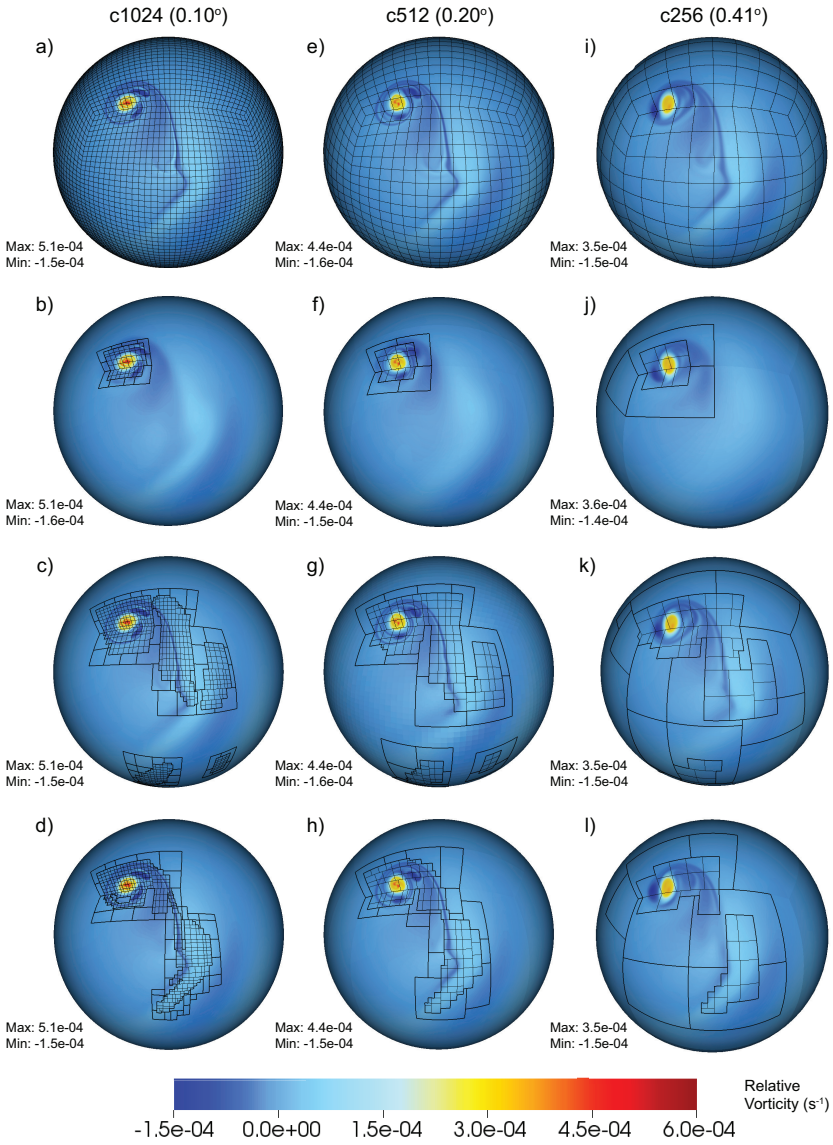

Figure 2.18: Relative vorticity field at day 4 on the cubed-sphere grid for several runs of the merging-vortices test of Section 2.3.5.2 using AMR. (a) Uniform c1024 run, (e) uniform c512 run, and (i) uniform c256 run. (b), (c), and (d) are AMR runs whose highest refinement level is c1024. (f), (g), (h) AMR runs have a maximum refinement level of c512 while (j), (i), and (l) AMR runs have a c256 maximum resolution. (b), (f), and (j) depict AMR runs with two level of refinement using the large-height-gradient tag. (c) , (g), and (k) depict 2-level AMR runs using the small-height-gradient tag. In the last row, (d), (h), and (l) the 2-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. All AMR runs use a x4 refinement ratio.

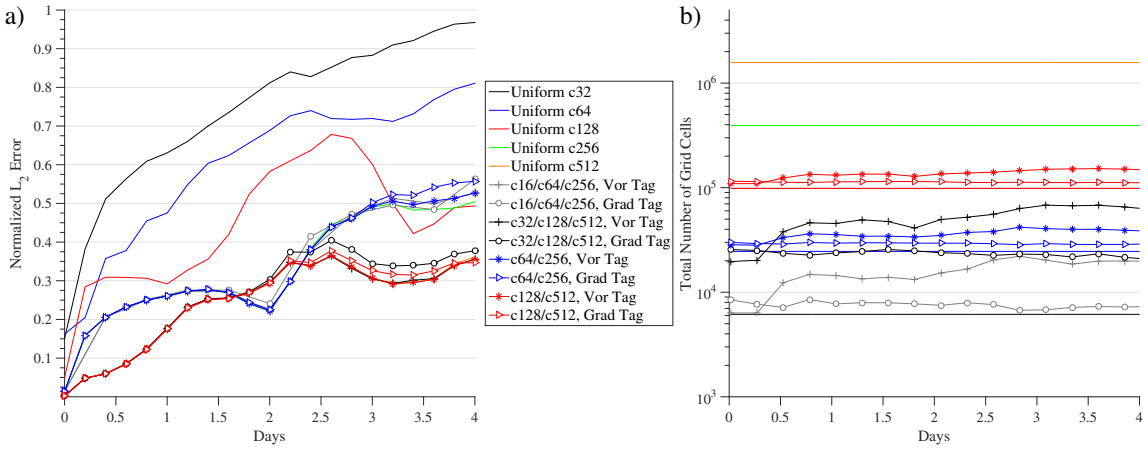

Figure 2.19: For the merging-vortices test of Section 2.3.5.2, growth over the four-day period of: (a) normalized l_2 error for relative vorticity calculated with respect to the uniform c1024 run, and (b) total number of grid cells. Error and number of grid cells are for uniform runs and 1- and 2- level AMR runs using the large-height-gradient tag or the relative vorticity tag with x4 refinement ratios.

Table 2.5: Run times (wall-clock time in s and as % of the c512 time) for 4-day merging-vortices simulations (Section 2.3.5.2) 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 Res.	AMR Levels	Run Time (s)	Run Time (%)	Total cells	Total cells (%)
c512	-	24127	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%

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 for the coarsest 2-level AMR runs (c32/c128/c512 and c16/c64/c256).

2.4 Conclusion

In this chapter, 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 x4 or x8 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 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 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 re-

finement 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 non-hydrostatic 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.

CHAPTER III

Assessing AMR in forced shallow water systems with moisture

3.1 Introduction

The spherical shallow water equations serve as an effective test bed for assessing numerical methods for general circulation models (GCMs). They exhibit many of the dynamics and complexities of the full 3D equations with the advantage of being two-dimensional and thus less computationally intensive. Full 3D models use the dycore of the GCM with a plethora of sub-grid parameterization schemes for unresolved physical and thermodynamical processes, further adding to their complexity. However, shallow water-type models and the unforced test cases traditionally associated with them (Williamson et al., 1992) are missing features (e.g. water vapor, convection, latent heat release, thermal forcing) which play key roles in atmospheric and climatological phenomena. Including simplified forcing mechanisms to represent such moisture and heating processes in the shallow water system narrows the gap between idealized unforced studies and full-physics models. These forced shallow water models include the dynamical complexities of the full models and can retain the non-linearity of the physical processes. However, the shallow water equations are still simple enough to effectively study the key components of the dycore such as the numerical algorithms,

computational grid, and, for variable resolution and AMR models, grid refinement strategies and efficacy. Shallow water models have the added advantage of being computationally cheap to run at high resolutions.

A variety of studies have implemented forcing in shallow water models to study the fundamental dynamical aspects of large synoptic scale climatological features, such as the Madden-Julian Oscillation (MJO), as well as intense, small-scale features including tropical cyclone evolution, cumulus convection, and frontal propagation. Ferreira et al. (1996) and Yang and Ingersoll (2013) implement spatially and temporally varying mass sinks and sources within nonlinear shallow-water models to simulate MJO convection in studies of MJO induced twin tropical cyclones and MJO convective envelope development, respectively. In Enagonio and Montgomery (2001), the authors examined tropical cyclogenesis by forcing a shallow water system with periodic pulses of vorticity to represent multi-burst convection. Hendricks et al. (2014) mimicked diabatic heating in hurricane eyewalls within a Cartesian shallow water model by using a prescribed annular axisymmetric mass sink.

A framework to study the specific dynamical role of moist processes in a shallow water system was proposed in the seminal work by Gill (1982). In this system, a moisture equation with nonlinear precipitation thresholds was added to the linearized shallow-water equations to model the effects of latent heat release on the propagation of large scale disturbances. Similar models incorporating this framework for parameterizing moisture were analyzed by Goswami and Goswami (1991) in the context of large-scale equatorial wave propagation and by Frierson et al. (2004), Stechmann and Majda (2006), and Bouchut et al. (2009) in studies of tropical precipitation fronts. Unlike the other studies mentioned, Bouchut et al. (2009) implements the moist-convective parameterizations in a fully nonlinear rotating shallow water model, also used by Lambaerts et al. (2011) to compare the differences in barotropic instability evolution between moist and dry conditions. To study the dynamical role of mois-

ture in tropical cyclone instabilities, Lahaye and Zeitlin (2016) also used this moist rotating shallow-water model with an added evaporation mechanism. Rostami and Zeitlin (2017) implemented both a 1-level barotropic version and a 2-level shallow water baroclinic version of the model introduced in Lambaerts et al. (2011) to study large-scale small Rossby number vortices.

Other frameworks for simulating precipitation and convection in the shallow water system that have been implemented recently included models by Würsch and Craig (2014) and Zerroukat and Allen (2015). Würsch and Craig (2014) developed a simplified model of cumulus convection incorporating representation of updrafts, downdrafts, and idealized precipitation effects in a 1D shallow water model. Once the fluid exceeds a certain threshold height, signaling the onset of convection, mechanisms serving as simplified representations of cumulus convection modify the geopotential height to create conditional instability and mimic updrafts. This method was extended to a rotating, 2D shallow water model in Kent et al. (2017). In Zerroukat and Allen (2015), the authors re-derived the two-dimensional shallow water system from the three-dimensional moist Boussinesq approximation. Density was permitted to vary with temperature, resulting in additional buoyancy related terms in the momentum equations and permitting a dynamics-moisture feedback. Zerroukat and Allen (2015) also implemented a three-state moisture model consisting of vapor, cloud, and precipitated species.

In this paper, we implement two forcing frameworks that seek to mimic moisture and convection in a spherical, 2D shallow water system. In the first, moist shallow water equations derived in Bouchut et al. (2009) and Lahaye and Zeitlin (2016) cause convective forcing and precipitation, which induce TC-like vortices to develop and strengthen. In the second framework, the barotropic instability shallow water test case of Galewsky et al. (2004) is implemented within a different moist shallow water system developed from Zerroukat and Allen (2015). As the jet becomes unstable

and collapses, front-like systems containing clouds and precipitation develop. Using these frameworks, we investigate the distinctive dynamics produced by the non-linear physics processes within the traditional shallow water system. A main goal of this paper is to introduce test cases with more challenging complex and transient features to test the ability of adaptive mesh refinement (AMR) to track and resolve these moving and growing features.

We use the fourth-order finite-volume Chombo AMR model presented in McCorquodale et al. (2015) and Ferguson et al. (2016) for the 2D shallow water equations. This model implements dynamic refinement using a mapped-multiblock AMR technique which overlays the base grid with more refined patches. Using AMR, we observe how features in the test cases evolve due to the forcing processes and how those forcing processes are affected by the AMR refinement. Ideally, the physics forcing schemes should be able to effectively span the multiple levels of refinement and changing resolutions created by AMR. In addition, we seek to quantify improvements gained from AMR grids and determine effective refinement criteria.

This chapter is organized as follows. Section 3.2 provides a brief overview of the finite volume model and the Chombo multiblock AMR techniques. Section 3.3 describes the shallow water equations with moist convective forcing and the strengthening vortex test case. Section 3.4 compares numerical results of the strengthening vortex test case for uniform and AMR runs of varying resolution. The Zerroukat and Allen (2015) moist shallow water system with barotropic instability test case and its numerical results are presented in Section 3.5. Section 3.6 summarizes our conclusions from the two test cases.

3.2 High-Order Finite-Volume Chombo AMR Model

For this study we employ an unstaggered finite-volume (FV) mapped-multiblock dynamical core (dycore) that is fourth-order accurate and adaptive in both space

Table 3.1: 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_{min}/A_{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 RLL_{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_{min}/A_{max}	Eq. Res.	RLL_{equiv}
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°
c2048	2.52×10^7	4.89	2.027×10^1	0.7074	0.04°	0.05°

and time. The model's AMR is based on the Chombo AMR library (Adams et al., 2015) and an in-depth description of the model dynamical core for the shallow-water equations on an equiangular cubed-sphere grid can be found in McCorquodale et al. (2015). The model is mass conserving and conserves energy up to the temporal truncation error, when limiters or explicit dissipation are not applied.

The cubed-sphere grid consists of a cube whose six separate panels are projected onto the surface of a sphere. In addition to eliminating the polar singularities found in spherical latitude-longitude grids, the equiangular cubed-sphere also leads to a quasi-uniform mesh with similarly sized grid cells across the sphere. The discrete resolution of the cubed-sphere grid is denoted by $c\{N_c\}$ where N_c is the number of grid cells in each direction on a panel. Several properties of the equiangular cubed-sphere grid, including approximate grid spacings, are given in Table 3.1 for resolutions used in this chapter.

The model uses a classical fourth-order Runge-Kutta (RK4) time discretization scheme. in the spatial domain, fourth-order accurate finite-volume discretization is implemented to compute flux averages on the faces of each cell. In addition, a high-order least-squares interpolation is used to compute stencil operations near panel

or block boundaries. The stencil operations create three layers of ghost cells at the panel edges to preserve the order of accuracy of the fluxes. At a panel edge, the fluxes for cell faces along the panel edge are calculated separately for each panel and are then averaged together. This mean value is taken as the flux for each face to ensure conservation. Additionally, a sixth-order diffusive operator is applied to smooth the flux calculations while still maintaining the scheme's fourth-order accuracy.

Our mapped-multiblock AMR approach implements a hierarchy of nested grid levels of increasing resolution, using the numerical tools developed in the Chombo AMR library (Adams et al., 2015). The grid resolution of an AMR level is defined by its refinement ratio to the grid resolution of the coarser level below it. Finer levels are placed over regions where coarse cells have been marked (tagged) by the model as meeting the refinement criterion. The cell values in the finer level are initialized via interpolations from the coarser level. Ghost cells are used to calculate fluxes at the level boundaries in the same manner as is done at the cubed-sphere panel boundaries. If multiple levels are used, intermediate levels must cover enough area to ensure that the finer level is nested within the intermediate level; it is required that the ghost cells for the finer level are only interpolated from cells within the intermediate level.

Finer levels are sub-cycled in time to maintain a constant Courant number across all resolutions. The sub-cycling routine can be summarized as follows: after interpolating values to the refined level, the coarser level is advanced by one time step. The finer level is then advanced in time using a smaller time step determined by dividing the coarse time step by the refinement ratio between the two levels. At each finer time step, the ghost cells used to calculate the boundary fluxes at the finer level's edges are updated via a temporal interpolation from the RK4 method. After the sub-cycling is complete, the values on the coarse grid are updated from the solution on the finer grid. McCorquodale et al. (2015) and Ferguson et al. (2016) both provide more detailed descriptions of the AMR technique.

The refinement criteria determine the regions over which additional grid levels are placed based on user-selected threshold values for flow properties. The thresholds are set independently for each simulation and their criteria can be based on a variety of properties, such as tracer values, gradients, relative vorticity, or a combination of these. The AMR dycore can incorporate multiple levels of refinement, preset for each simulation, and tagging criteria can be uniformly enforced across all levels or required to scale with increasing resolution.

3.3 Convectively forced shallow water vortices

This test cases simulates the growth and development of TC-like vortices in a 2D shallow water framework using a moist convective forcing mechanism. Weak vortices are initialized on a calm background field of uniform height. Evaporation and precipitation cause these vortices to strengthen. After several days of strengthening, the orderly vortices collapse and a more chaotic system evolves, characterized by several smaller vortices and jet-like background flow.

We first provide a description of the moist convective shallow water system and an overview of the initial conditions. Next we present the evolution of an isolated vortex in this test case at a uniform high resolution. Then we demonstrate the use of AMR in this test case comparing the affects of different resolutions and various tagging criteria.

3.3.1 "Moist-convective" shallow water equations

The shallow water equations are modified to include the transport of a moisture variable and the effects of moist convection, precipitation, and evaporation. We extend the forcing schemes developed by Bouchut et al. (2009) and Lahaye and Zeitlin (2016) to the 2D shallow water equations on the sphere. Bouchut et al. (2009) developed a moist convective scheme for a rotating shallow water model to study precip-

itation fronts. It only included an advected moisture variable, precipitation, and a convection mimicking process, while Lahaye and Zeitlin (2016) added an evaporation forcing to the system.

In this moist convective system, a relaxation sink is added to the moisture equation (Eq. 3.3), representing precipitation when moisture levels exceed a saturation value. A corresponding convective mass sink is added to the continuity equations and a moisture source is added to the moisture equation to represent evaporation. The equations for this modified shallow water system are

$$\frac{\partial h\mathbf{v}}{\partial t} + \nabla \cdot (h\mathbf{v}\mathbf{v}) + f\hat{\mathbf{k}} \times (h\mathbf{v}) + gh\nabla H = 0 \tag{3.1}$$

$$\frac{\partial h}{\partial t} + \nabla \cdot (h\mathbf{v}) = -\beta P \tag{3.2}$$

$$\frac{\partial hQ}{\partial t} + \nabla \cdot (hQ\mathbf{v}) = h(E - P). \tag{3.3}$$

Here $\mathbf{v}$ is the velocity vector, $\mathbf{v}\mathbf{v}$ denotes the outer product of the velocity vector, f is the Coriolis parameter, g is the acceleration due to gravity, h is the height of the fluid, H is the total height including topography, and the dimensionless moisture variable Q represents bulk humidity. E is the evaporation in the moisture budget and P is the precipitation. The latent heat release from precipitation does not directly influence the horizontal momentum. It is instead linked to convective vertical velocity directly proportional to P at the upper surface of the fluid. Since this system is designed to represent only the lower part of the troposphere, the convective updraft can be constructed as a mass exchange from this surface layer. The mass exchange is then represented as a mass sink in Equation 3.2 governed by an adjustable constant β. A detailed explanation for the implementation is presented by Bouchut et al. (2009).

The precipitation sink is calculated in terms of Q and the saturation value Q_s,

$$P = \frac{Q - Q_s}{\tau}\mathbf{H}(Q - Q_s). \tag{3.4}$$

with a relaxation time of τ. $\mathbf{H}(\cdot)$ is the Heaviside function so that $P = 0$ whenever $Q \leq Q_s$. A common parameterization used in Lahaye and Zeitlin (2016) is modified to include an upper velocity limit on the evaporation rate, thus the evaporation rate E_r becomes

$$\begin{aligned} E_r &= \alpha_e|\vec{v}| && \text{for } |\vec{v}| < v_{\max} \\ E_r &= \alpha_e(v_{\max}) && \text{for } |\vec{v}| > v_{\max} \end{aligned} \tag{3.5}$$

where the evaporated moisture is dependent on the magnitude of the velocity $\vec{v}$ and the adjustable coefficient a_e. This evaporation scales with wind speed until $v_{\max}$. In the simulations presented below, $a_e = 0.055$ m^{-1} and $v_{\max} = 30$ m s^{-1}. At higher velocity magnitudes, E_r is constant.

As established, evaporation is unlimited, but such a setup can lead to runaway supersaturation and very large height forcing which causes exceedingly high wind velocities and negative height values in longer simulations. To limit these sources of instability, we refine the forcing mechanics with the following additional mechanisms. We add a moisture reservoir from which the evaporation is drawn. This reservoir is a simplistic representation of ocean surface heat content and its limiting effects on tropical convection and tropical cyclone intensity. In this forcing scheme, the evaporation rate E_r defined in Eq. 3.5 draws moisture from the finite reservoir $C_r(\phi, \lambda)$. However, evaporation cannot exceed the amount of moisture remaining in C_R for any given latitude-longitude (ϕ, λ) point in space and time. Thus, the evaporation rate

E applied in Eq. 3.3 is capped so that

$$E = \frac{1}{\Delta t} \min\left(E_r \Delta t, C_r\right) \tag{3.6}$$

where Δt is the model time step. So the evaporation rate becomes zero if there is no longer moisture to be drawn from C_r.

The amount of moisture in the reservoir is affected by evaporation which removes moisture from the reservoir, and a Newtonian relaxation slowly returns the reservoir to its initial state. The rate of change of moisture in the reservoir is

$$\frac{\partial C_r(\phi, \lambda)}{\partial t} = -E + \frac{1}{\tau_c}(C_0(\phi, \lambda) - C_r(\phi, \lambda)) \tag{3.7}$$

where $C_0(\phi, \lambda)$ is the initial value and $\tau_c = 10$ days is the relaxation parameter. The initial moisture value C_0 is a zonally symmetric field where

$$C_0 = C_{max} \cos^4(\phi) \tag{3.8}$$

with the constant $C_{max} = 0.05$. The reservoir is largest near the equator and declines to zero at the poles. The evaporation rate E as calculated in Eq. 3.6 ensures that C_r will not go negative.

Bouchut et al. (2009) and Lahaye and Zeitlin (2016) implement a constant β to represent the precipitation mass sink. We implement a variable β which scales with fluid height. As the fluid height decreases below the initial value, β is minimized, reducing precipitation forcing and even removing it completely after h has fallen below a certain height. This setup prevents negative fluid height caused by the convective forcing and limits the vortex strength. Thus the coefficient β is set as

$$\beta = \beta_0 \tanh\left(\max\left[0, \lambda_b \frac{h - h_t}{h_0 - h_t}\right]\right). \tag{3.9}$$

The constants β_0 and λ_b are set for the tests below to 0.01 and 10.0 respectively, while h_0 is the initial background height and h_t is the cutoff height set to 75% of h_0. As implemented in the vortices test case, the β term does not limit the vortex strengthening until the vortex is well established.

3.3.2 Initialization of vortices

The initial conditions of this test case consist of one or more weak vortices seeded onto a constant height zero velocity background. In the moisture field, each small vortex is initialized with some supersaturation to ensure the convective forcing is triggered immediately, creating convergence and allowing the vortices to strengthen.

Each small initial vortex is a small Gaussian depression in the uniform height field such that the height field for one initial vortex is

$$h = h_0 - h_f \exp\left(-\left(\frac{r}{r_w}\right)^2\right). \tag{3.10}$$

Here, $h_0 = 4000$m is the uniform initial height, h_f is the depth of the Gaussian depression, and r_w is the radius of maximum winds. The values for h_f and r_w are designated later for each initial vortex. The great circle distance r from point (ϕ, λ) to the vortex center (ϕ_d, λ_d) is given by , and such that

$$r = a \arccos\left(\sin\phi_d \sin\phi + \cos\phi_d \cos\phi \cos\left(\lambda - \lambda_d\right)\right) \tag{3.11}$$

with the Earth's mean radius $a = 6.37122 \times 10^6$ m. The chosen (ϕ_d, λ_d) vortex center points are explained later. The corresponding perturbations in the zonal u and meridional v wind components for each vortex are derived from the geostrophic wind balance so that

$$u = -\frac{g}{af}\frac{\partial h}{\partial \phi} = \frac{2gh_f}{ar_w^2 f} r \exp\left(-\left(\frac{r}{r_w}\right)^2\right)\frac{\partial r}{\partial \phi} \tag{3.12}$$

$$v = \frac{g}{af\cos\phi}\frac{\partial h}{\partial \lambda} = \frac{-2gh_f}{ar_w^2 f\cos\phi} r \exp\left(-\left(\frac{r}{r_w}\right)^2\right)\frac{\partial r}{\partial \lambda}. \tag{3.13}$$

The Coriolis parameter $f = 2\Omega\sin\phi$ has the constant rotation rate $\Omega = 7.292 \times 10^{-5}$ s^{-1}. The partial derivatives of the great circle distance with respect to spherical coordinates ϕ and λ are

$$\frac{\partial r}{\partial \phi} = \frac{\sin\phi_d\cos\phi - \cos\phi_d\sin\phi\cos(\lambda - \lambda_d)}{\sqrt{1 - (\sin\phi_d\sin\phi + \cos\phi_d\cos\phi\cos(\lambda - \lambda_d))^2}} \tag{3.14}$$

and

$$\frac{\partial r}{\partial \lambda} = \frac{\cos\phi_c\cos\phi\sin(\lambda_d - \lambda)}{\sqrt{1 - (\sin\phi_d\sin\phi + \cos\phi_d\cos\phi\cos(\lambda - \lambda_d))^2}}. \tag{3.15}$$

This initialization of the weak vortices is not perfectly balanced on the sphere and will cause the creation of some gravity waves. However, there is no need to further balance these initial conditions due to the effects of convective forcing triggered immediately by the initial supersaturation. The associated mass sink in Equation 3.2 creates significantly larger gravity waves which quickly overwhelm the affects of the initial field imbalances.

The initial dimensionless moisture profile Q consists of a background profile Q_0 and a Gaussian hill leading to supersaturation overlaying each vortex. The background profile has a minimum value at the poles and increases to a maximum just below the saturation point in the equatorial region, and is given by

$$Q_0 = \min(Q_{sat} - Q_{off}, \cos\phi + Q_{min}) \tag{3.16}$$

with the saturation value $Q_{sat} = 0.9$, $Q_{off} = 0.01$ sets the maximum background field just below saturation, and $Q_{min} = 0.05$ is the minimum initial moisture value at the poles. On top of the background profile is a Gaussian hill which raises the initial moisture level in each vortex above saturation. With this perturbation, the

total initial moisture profile is

$$Q = Q_0 + Q_f \exp\left(-\left(\frac{r}{r_w}\right)^2\right) \tag{3.17}$$

where Q_f is the maximum moisture perturbation value at the center of the vortex. This value is stated later for the individual vortices. The initial supersaturation allows the vortex to begin strengthening immediately. It provides the initial convergence which pulls additional moisture toward the vortex center, sustaining the growth of the vortex. Without supersaturation, the growth of the vortex is less consistent and requires more time to initiate. For runs with more than one vortex, the height, velocity, and moisture perturbations for all initial the vortices are summed up and overlaid on the background fields. If the vortices are close together, their fields may significantly overlap.

3.3.3 Evolution of a strengthening vortex

As a demonstration of the convective forcing mechanisms, we implement the test case by initializing one isolated vortex at a uniform high resolution and observe its evolution over twelve days. This c2048 resolution ($\sim$ 5 km) run serves as a reference solution for AMR and coarser resolution runs. The initial weak vortex is centered at ($\phi_d = 10°, \lambda_d = 0°$) with maximum height perturbation $h_f = 10$ m, radius of maximum wind $r_w = 600$ km, and maximum moisture perturbation $Q_f = 0.0175$. These values result in an initial wind magnitude peak of 5.7 m s^{-1} and a peak moisture value of 0.83% above the saturation value Q_s. The initial vorticity profile is depicted in Fig. 3.1a. In addition, we initialize a non-symmetric binary pair of vortices 90° east of the isolated vortex. These vortices interact with each other as they strengthen but do not merge. Each has the same size and strength as the first vortex and but a lower level of supersaturation. The second vortex is centered at ($\phi_d = 10°, \lambda_d = 90°$)

with $Q_f = 0.015$ and the third is at $(10°, 105°)$ with $Q_f = 0.0125$. The two additional vortices have little effect on the evolution of the isolated main vortex during the first twelve days of simulation. They are added to the test case to help form a complex chaotic global flow approximately 14 days into the simulation. This analysis and discussion is focused on the single isolated vortex during the initial twelve days; the extended-time (up to 16 days) results focusing on all three vortices are discussed in detail in Section 3.4.1.

Initial supersaturation triggers convection immediately, creating convergence at the vortex center and driving vortex strengthening. The evolution of the vortex's relative vorticity profile over a period of twelve days is depicted in Fig. 3.1. As the vortex drifts toward the northwest due to beta drift, it undergoes a steady increase in strength over the first six days. At day 6, the maximum wind magnitude has increased to 16.7 m s^{-1} and the vortex strengthens more rapidly from this point. By the eighth day of the simulation the maximum winds are 31.2 m s^{-1} and by day 10 they are 69.0 m s^{-1}.

A key feature that develops around day 4 is a symmetric ring of maximum vorticity. This ring can be clearly seen at days 6, 7, and 8 (Figs. 3.1d-f). As the vortex intensifies more rapidly, this symmetric ring becomes elongated, as seen at day 8 in Fig. 3.1f. The unstable elongated ring collapses and the filaments of large positive vorticity begin to collate, creating a concentrated area of maximum vorticity (day 9, Fig. 3.1g). A small section of the vorticity filament is not reincorporated into the center spirals of the main vortex (day 10, Fig. 3.1h), becoming a separate, smaller secondary vortex with a vorticity dipole feature by day 12 (Fig. 3.1i).

Figure 3.2 provides a snapshot of the height field, wind magnitude, precipitation rate, and moisture reservoir at day 8, corresponding to the relative vorticity profile in Fig. 3.1f. After eight days, the minimum height shown in Fig. 3.2a at the center of the vortex has decreased by 140 m. The velocity magnitude profile depicted in Fig.

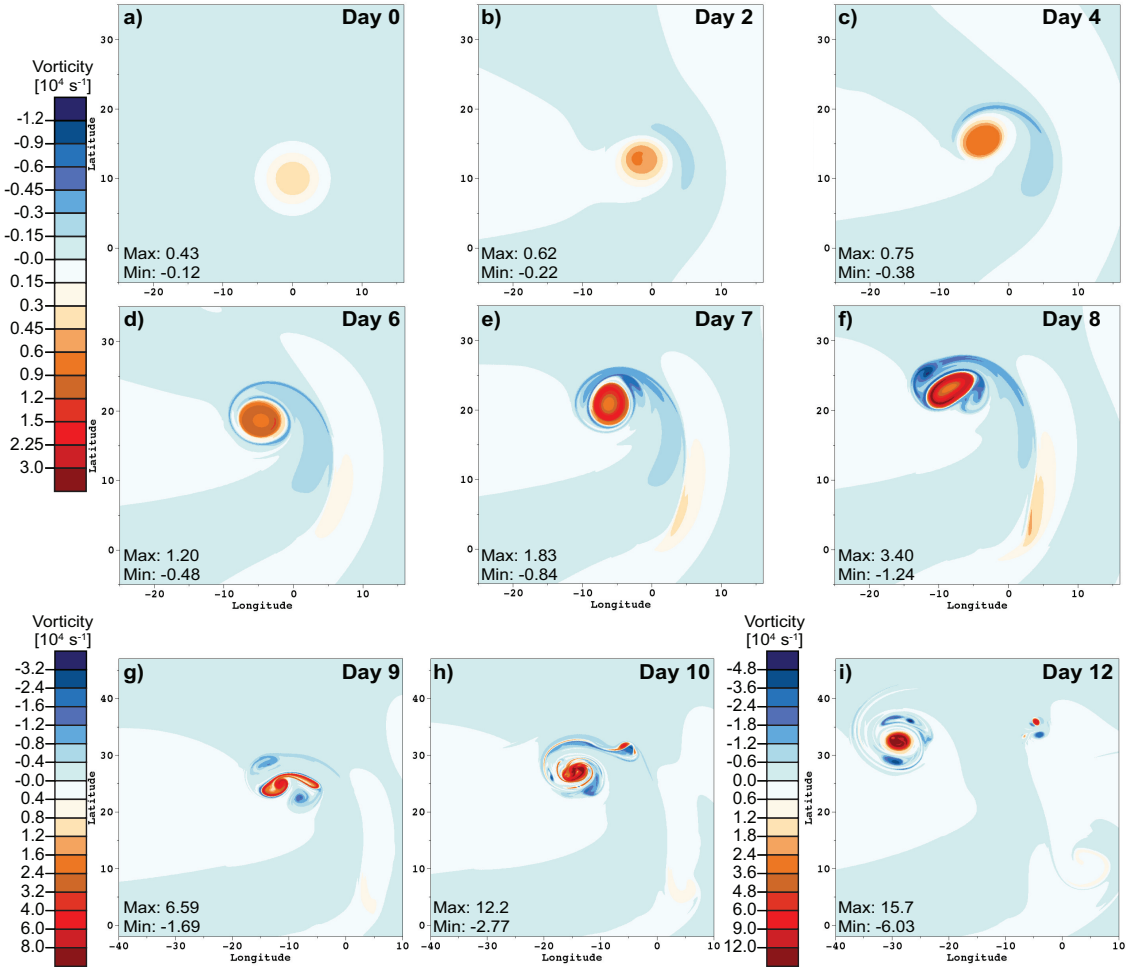

Figure 3.1: The evolution of the relative vorticity for an isolated strengthening vortex in a c2048 uniform run. (a)-(f) Relative vorticity plots for the initial condition, day 0, and days 2, 4, 6, 7, and 8 with color contour range of -1.2×10^{-4} s^{-1} to 3.0×10^{-4} s^{-1}. (g) and (h) Relative vorticity for days 9 and 10 with the color contour ranged increased to between -3.2×10^{-4} s^{-1} to 8.0×10^{-4} s^{-1}. (i) Relative vorticity for day 12 with a contour range of -4.8×10^{-4} s^{-1} to 12.0×10^{-4} s^{-1}. Note that (g)-(i) have an expanded latitude-longitude domain.

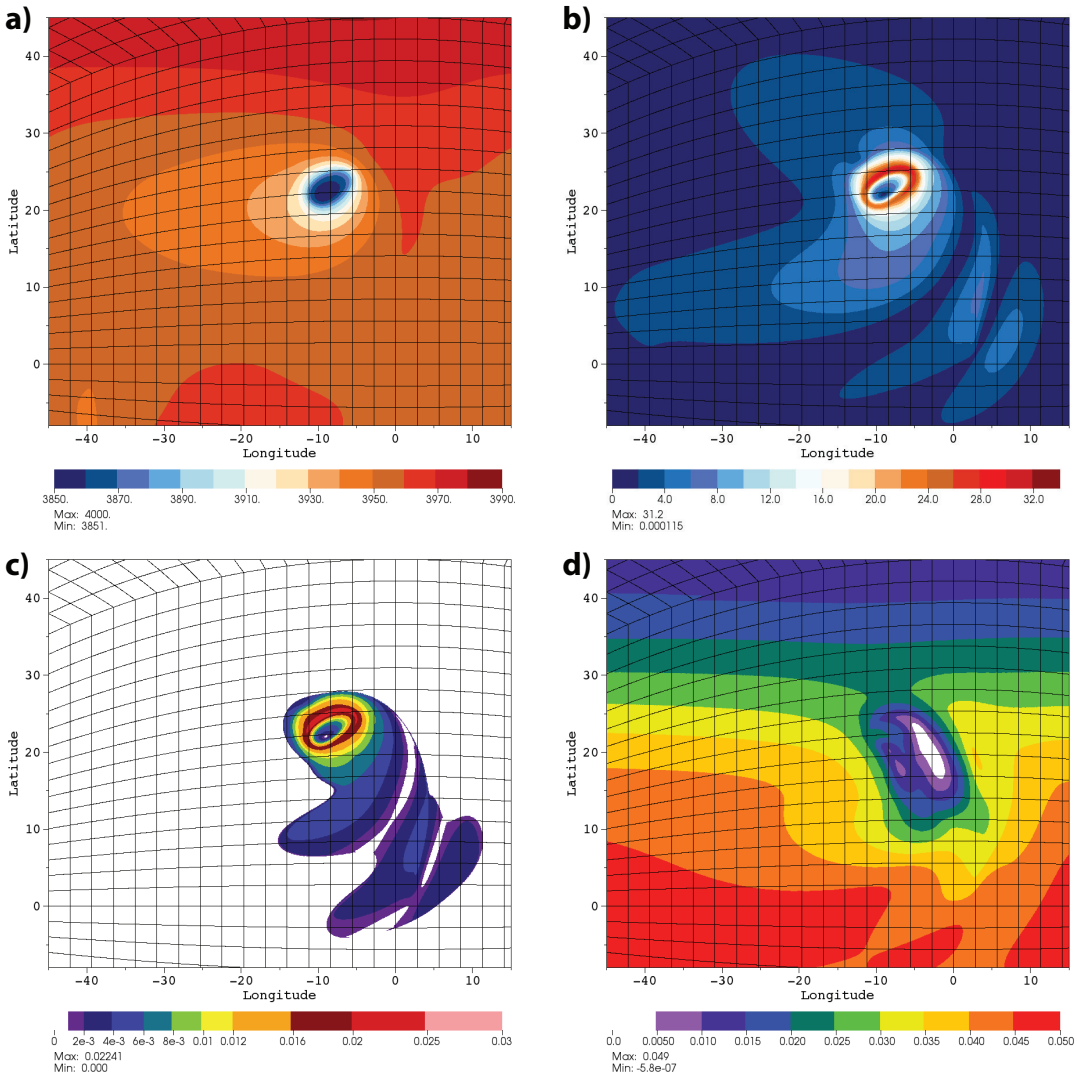

Figure 3.2: Day 8 plots for the uniform c2048 run of the isolated strengthening vortex for several variables: (a) Height field (m), (b) Wind magnitude (m s^{-1}), (c) Instantaneous precipitation rate (moisture value per day), and (d) Reservoir moisture content (moisture value). These plot correspond to the day 8 vorticity plot in Fig. 3.1(f), though note the larger latitude-longitude domain in these plots.

3.2b and the instantaneous precipitation rate in Fig. 3.2c contain, as in the vorticity profile at day 8, an elongated ring of strongest winds and heaviest precipitation. to the southeast of the vortex center, a Rossby wave train forms as visible in the wind and precipitation fields (Figs. 3.2b,c). The ocean-like reservoir of available moisture for evaporation is plotted in Fig. 3.2d. The area of low reservoir levels to the southeast of the vortex shows where evaporation has been the strongest and reflects the path of the vortex.

The same fields are depicted for day 10 in Fig. 3.3. The height has decreased to a minimum of 3400 m. As in the vorticity profile in Fig. 3.1h, the wind magnitude profile in Fig. 3.3b has begun to re-form a cyclone-like profile with a central calm eye surrounded by a ring of maximum winds. However, the precipitation rate (Fig. 3.3c) does not reform a symmetric ring, and precipitation is limited in the southeastern sector of the vortex due to a lack of additional moisture being evaporated. As seen in Fig. 3.3d, the available moisture for evaporation in that area has been depleted. This lack of moisture for evaporation suppresses convection and vortex strengthening. A comparable real world effect is the cold wake reduction in ocean surface heat content that forms behind the path of tropical cyclones due to the cyclone wind induced surface layer mixing. The main vortex strengthen more slowly after day 10, as the vortex has reached higher latitudes that have less moisture available in the reservoir. The maximum relative vorticity peaks after eleven days before slowly declining in magnitude, though the small secondary vortex continues to strengthen. The main vortex wind speed follows the same trend, reaching a maximum between twelve and thirteens days with an unrealistic speed of 176 m s^{-1} before slowly declining.

3.4 Numerical results for uniform and AMR grids

Next we examine the effects of grid resolution on the evolution of the vortex and assess the ability of different AMR refinement criteria to achieve comparable results

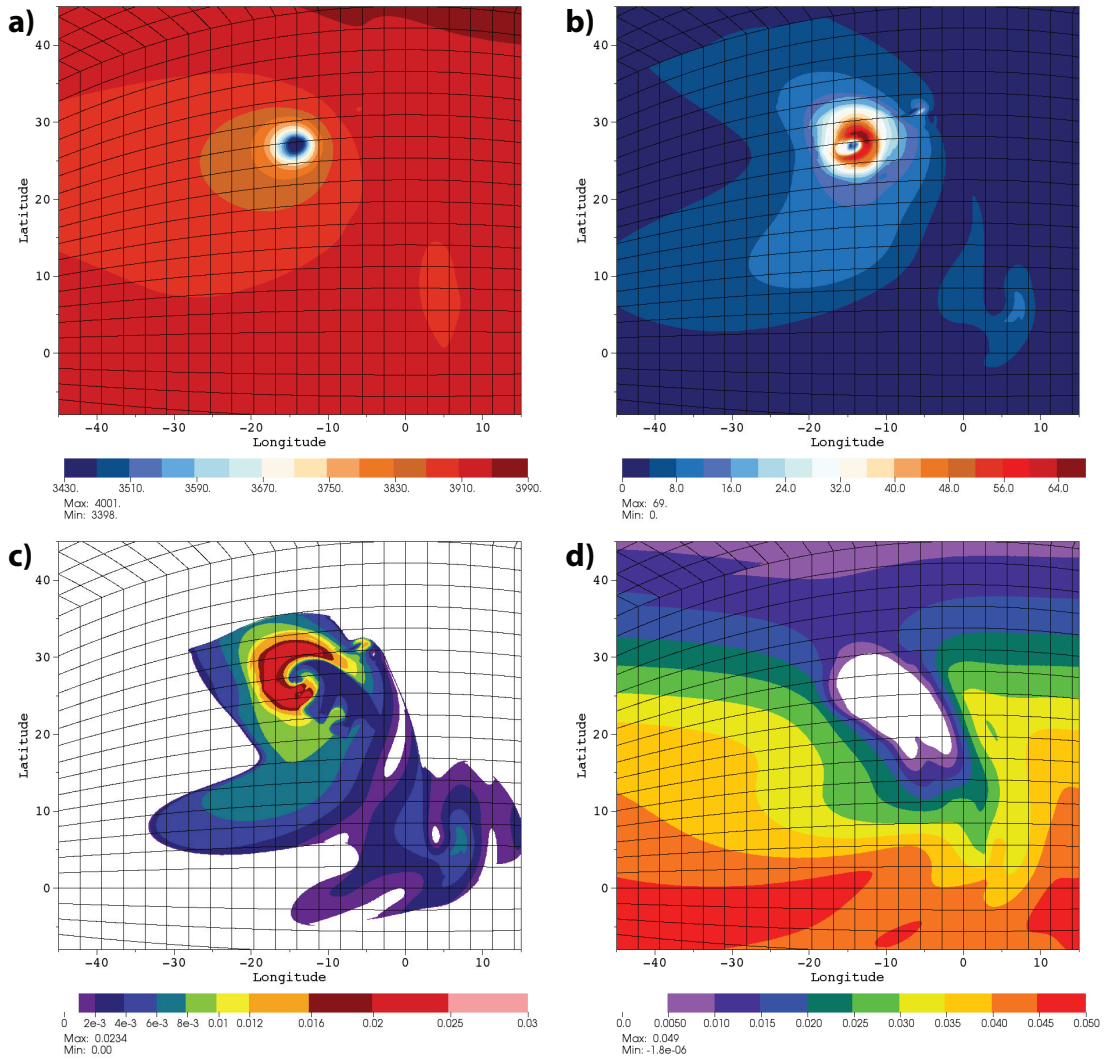

Figure 3.3: Day 10 plots for the uniform c2048 run of the isolated strengthening vortex for several variables: (a) Height field (m), (b) Wind magnitude (m s^{-1}), (c) Instantaneous precipitation rate (moisture value per day), and (d) Reservoir moisture content (moisture value). These plot correspond to the day 10 vorticity plot in Fig. 3.1(h), though note the different latitude-longitude domain in these plots.

to uniform resolution runs. With the high resolution c2048 run serving as a reference, we implement the test case with the resolutions listed in Table 3.1. In addition, we conduct a variety of runs with AMR based on 4 different tagging criteria sets to determine the optimal tagging criterion. The differences between these sets are (a) the magnitude of the absolute value of the relative vorticity threshold, on which the criteria are refined, and (b) the change in that threshold between multiple levels of AMR. We implement two types of thresholds. The first is a constant threshold in which the required vorticity value is the same across all resolutions and, when met, all the AMR levels are added simultaneously. The second version is a scaled threshold that increases with the level's resolution. In this setup, one level of refinement is implemented when the threshold value is met, but the next level of AMR is not added until a new, higher threshold value scaled with the increasing resolution is reached. For example, a three level AMR run with x4 refinement, a base resolution of $c32$, and a vorticity tag threshold of $|\zeta| > 2$ days$^{-1} = 2.3 \times 10^{-5}$ s^{-1} will implement the first AMR level at a c128 resolution once that threshold has been exceeded. With this new resolution of c128, the threshold is four times higher $|\zeta| > 8$ days$^{-1} = 9.3 \times 10^{-5}$ s^{-1} in regions covered by the first level of AMR. This causes the second level of refinement, with a c512 resolution, to not implement until the vorticity has exceeded this new higher threshold. This delay in triggering higher refinement results in some loss of detail on features of interest in their initial stages, but fewer grid cells, and thus computational cost, are added until absolutely necessary. For AMR with only one level of refinement, there is no difference between the two threshold scalings. The four tagging criteria are:

- Tag 1: a vorticity threshold that scales with resolution with a base level threshold of $|\zeta| > 2$ days$^{-1} = 2.3 \times 10^{-5}$ s^{-1}
- Tag 2: a vorticity threshold that scales with resolution with a base level threshold of $|\zeta| > 3$ days$^{-1} = 3.5 \times 10^{-5}$ s^{-1}

- Tag 3: a vorticity threshold that scales with resolution with a base level threshold of $|\zeta| > 8$ days $^{-1} = 9.3 \times 10^{-5}$ s^{-1}
- Tag 4: a constant threshold across all AMR levels for vorticity of $|\zeta| > 5$ days $^{-1} = 5.8 \times 10^{-5}$ s^{-1}.

The higher vorticity threshold in Tag 4 compared to the scaled thresholds in Tags 1 and 2 attempts to reduce the computational cost of the constant threshold method and make the model more comparable to that of the scaled threshold methods. Tag 3 simply increases the threshold by four times that of Tag 1 so that a c32 base level AMR with x4 refinement using Tag 1 will have the same threshold of refinement for its second level of AMR as a c128 base level AMR using Tag 3 has for its first level of AMR.

We focus on the vorticity field in this analysis, as it contains a combination of large-scale and fine-scale structures and is sensitive to changes in resolution. The height and velocity fields have fewer small-scale features and the minimum height and maximum wind associated with the main vortex correlate well to the vorticity strength. The precipitation field is also larger in scale and has little variability with resolution.

Focusing on the single isolated vortex, Fig. 3.4 depicts the growth of the vortex with a daily maximum magnitude of relative vorticity for uniform resolution runs (Fig. 3.4a), AMR runs using tags 1 and 3 (Fig. 3.4b), AMR runs with tag 2 (Fig. 3.4c), and AMR runs with tag 4 (Fig. 3.4d). The line color for each AMR run corresponds to base level resolution while the line style corresponds to the resolution of the highest AMR level. The uniform runs in Fig. 3.4 serve as the reference for the color and style of each resolution.

Overall vortices with higher resolution have larger maximum vorticity over the first nine days. The c64 resolution has not strengthened, nor has the c32 resolution, which is not plotted in Fig. 3.4a. In these cases, the grid resolution is too coarse and

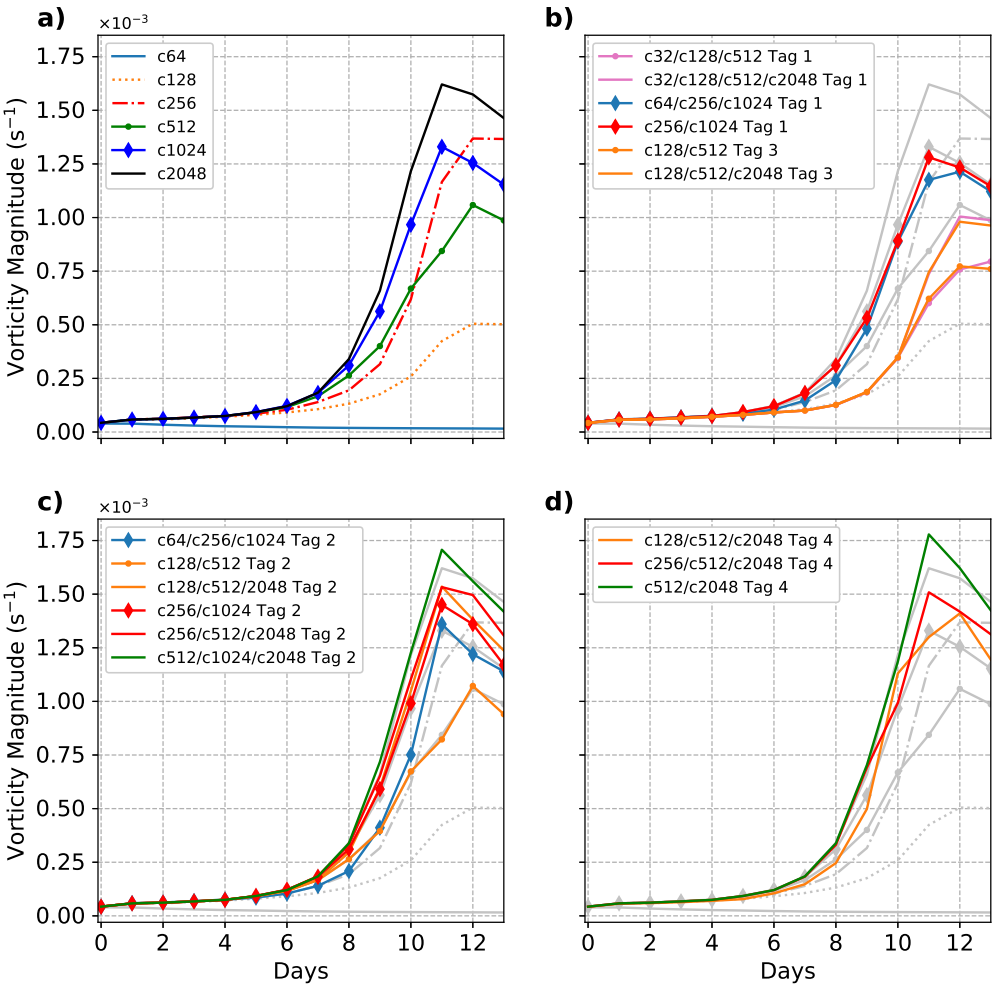

Figure 3.4: Maximum relative vorticity of the strengthening vortex over a period of 13 days for (a) uniform runs, (b) AMR runs using the Tag 1 or Tag 3 refinement criteria, (c) AMR runs using the Tag 2 criteria, and (d) AMR runs using the Tag 4 criteria. All the plots follow a line color and marker system dependent on resolution. The line color denotes the run's base resolution while the line style denotes the run's highest AMR resolution. The line color and style for each resolution is as follows: c2048 (black, plane), c1024 (blue, solid diamond markers), c512 (green, small circle markers), c256 (red, dot-dash line), and c128 (orange, dotted line). The coarse resolutions c64 and c32 have only a line color, light blue and pink respectively, and no line style as none of the AMR runs have a highest refinement at these resolutions given that the vortex does not develop on such a coarse grid (as seen by the c64 uniform run in (a)). For comparison purposes the uniform run lines from (a) have been imposed in light grey on the other three plots.

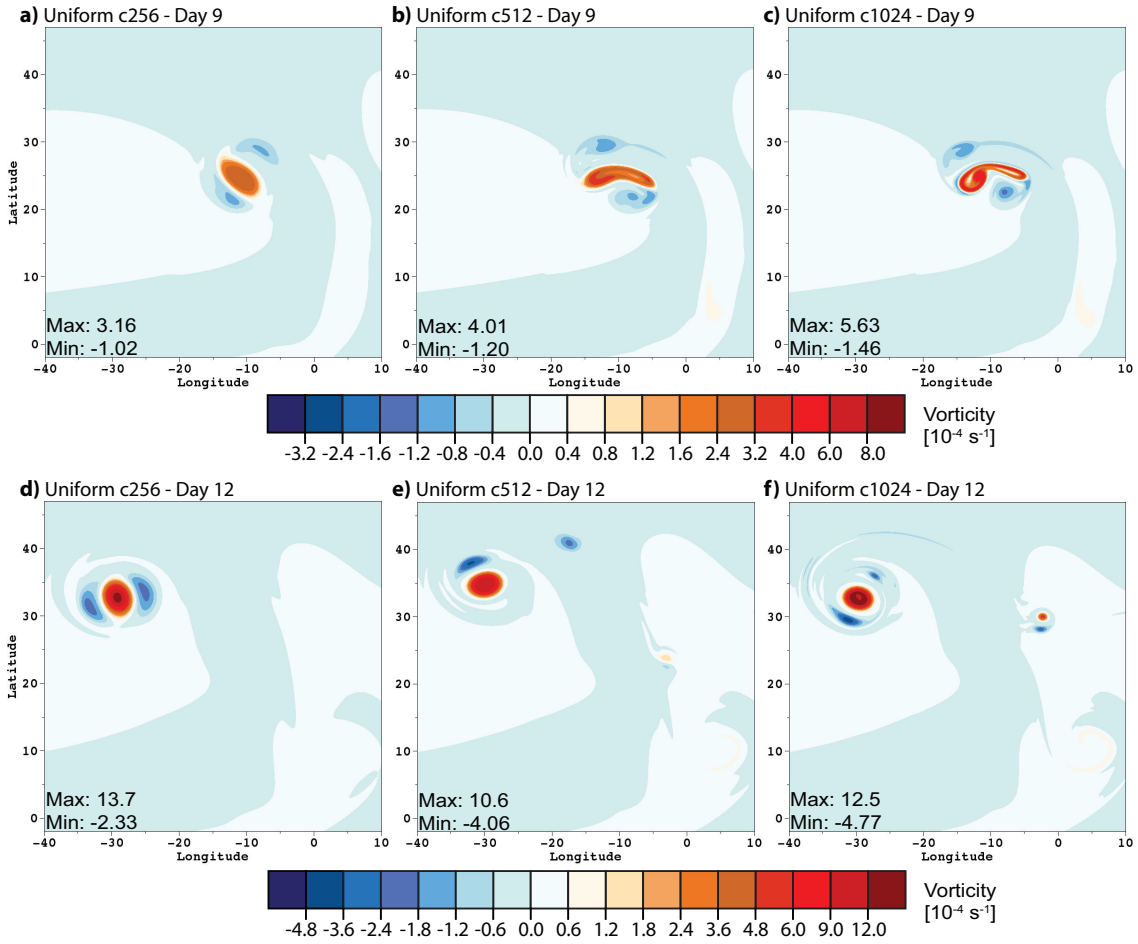

Figure 3.5: Relative vorticity field of the strengthening vortex case at day 9 (a) - (c) and day 12 (d)-(f) for uniform runs c256 resolution (a) and (d), c512 resolution (b) and (e), and c1024 resolution (c) and (f). These plots correspond to the day 9 uniform c2048 plot, Fig. 3.1g, and day 12 plot Fig. 3.1i.

the initialized vortices slowly weaken. In Fig. 3.4 we see the c256 run strengthening more rapidly after day 9 than higher resolution runs such that its maximum vorticity is higher than the c512 and c1024 runs by day 12. Over the same time period we see that the c512 and higher resolution runs reach a peak strength between day 11 and day 12 before slowly decreasing, observing a regime change that the coarser resolution c256 and c128 are not able to resolve. This regime change can be seen in Fig. 3.5, which depicts the relative vorticity field for uniform runs c256, c512, and c1024 at day 9 in Figs. 3.5 (a)-(c) and day 12 in Figs. 3.5(d)-(f).

Comparing day 9 in Fig. 3.5 and Fig. 3.1g, the strength of the vortex is clearly dependent on resolution, as the maximum vorticity increases with each resolution. The structure of the vortex, however, visually converges. The ring collapse and roll-up seen at day 9 in the uniform c2048 run in Fig. 3.1g is also observed in the c1024 run (Fig. 3.5c) despite the slightly weaker maximum vorticity. The weaker vortex of the uniform c512 run results in a delay of the stretching and collapse of the vortex ring. In Fig. 3.5b, the vortex is stretched, but the roll-up has not begun. The ring structure is still visible but less distinct as in the higher resolution runs. In the uniform c256 run the vortex does not form a distinct vortex ring structure and continues to strengthen at day 9 without a significant change to its structure. (Fig. 3.5a).

By day 12, the c1024 vorticity filaments have rolled back up into a single vortex and spun off a smaller secondary vortex with a vorticity dipole feature. The field is visually similar to the c2048 run at day 12 in Fig. 3.1i, albeit slightly weaker overall and thus less northward by day 12. Since the c256 run did not undergo the collapse and roll-up, its maximum vorticity at day 12 in Fig. 3.5d is now stronger than that of the c512 and c1024 vortices. In addition, no secondary vortex develops. The uniform c128 run follows a similar, though weaker evolution, so it is not shown here. For the c512 uniform run at day 12 in Fig. 3.5e, a secondary vortex has spun off (centered around $(25°N, -3°W)$ but is significantly weaker than the secondary vortices in the c1024 and c2048 runs. In addition, one of the two anticyclonic regions that abut the main vortex in the c1024 and c2048 runs has been spun off as well. Figure 3.5 shows these two different regimes are dependent on resolution. At resolutions c256 and below the development of a distinct vortex ring and its collapse cannot be resolved. The c512 resolution appears to be the start of the high resolution regime, which begins to converge between the c1024 and c2048 resolution runs.

The AMR runs shown in Figs. 3.4(b)-(d) do an effective job of following the

growth trajectory of uniform run with the same resolution as the finest AMR level with a few exceptions. However, several AMR maximum vorticity values remain lower than the corresponding uniform run. This occurs because the higher resolution refinement is not implemented early enough in the simulation, resulting in the lower strength observed in the uniform c256 and lower resolutions runs. This can be seen in the c128/c512/c2048 AMR with Tag 4 in Fig. 3.4d. The refinement levels in that run are triggered at day 2, and the resulting delay causes the maximum strength to remain below the uniform c2048 level throughout the thirteen days of the simulation. The largest divergence occurs for the c32 AMR runs with Tag 1 and c128 AMR runs with Tag 3 in Fig. 3.4b. The two runs with the highest AMR level of c2048 have a maximum vorticity nearly 40% weaker at day 12 than the uniform c2048 run. The other two AMR runs, with a highest AMR level of c512, are approximately 25% weaker than the c512 uniform vortex. These runs also follow the low resolution regime more closely than the regime expected by their highest AMR level. All four of these runs begin with a highest resolution of c128, as the Tag 1 threshold for c32 AMR runs triggers the c128 level of refinement immediately at initialization. However, for both the c32 AMR Tag 1 runs and the c128 Tag 3 runs, the c512 AMR level is not triggered until six and a half days into the simulation, and the two runs which reach the c2048 level resolution do not trigger it until day 10. As a result, the model benefits of instituting this higher resolution occur too late in the simulation. As Fig. 3.4b shows, these four runs do not diverge from the c128 uniform run until day 9. In other AMR setups, the delayed implementation of the highest AMR resolution does not degrade the growth of the vortex strength. For the c64/c256/c1024 AMR runs with Tag 1 and Tag 2 refinement, the first refinement level of c256 resolution is triggered initially. The second c1024 level, however, is not triggered until after five days or seven days, for Tag 1 and Tag 2, respectively. As a result, we observe that the c64 AMR run with Tag 1 in Fig. 3.4b does not diverge from the uniform c256

run until after day 7, and in Fig. 3.4c the c64 Tag 2 AMR run does not diverge until after day 8. Both, however, follow the c1024 maximum vorticity closely by day 10. Though delayed, the refinement still occurs before the rapid intensification period during which the developing vortex becomes unstable and collapses. Thus these two runs are able to more successfully match the c1024 uniform results than the c32 Tag 1 AMR runs. In the cases of the c512/c2048 AMR runs, the AMR results in a slightly higher maximum vorticity than the corresponding uniform run between day 10 and day 12 using Tag 2 (Fig. 3.4c) and Tag 4 (Fig. 3.4d) as well as with the c256/c1024 AMR run using Tag 2 (Fig. 3.4c).

Figures 3.6 and 3.7 depict the relative vorticity field for day 9 and day 12 respectively for nine AMR runs. They provide a more detailed comparison of the overall vortex and the small-scale features in the vorticity field between the AMR runs and the uniform resolution runs in Fig. 3.1 and Fig. 3.5. At day 9, the selected AMR runs in Fig. 3.6 are at various stages of the evolution process, depending on their resolutions and tagging criteria. As in Fig. 3.4b, the c32 base 4-level AMR with Tag 1 run (Fig. 3.6a) and the c128 2-level AMR with Tag 3 run (Fig. 3.6b) are significantly weaker than the uniform c2048 run at day 9 (Fig. 3.1g). The highest level of AMR for both runs, c2048, has not yet been triggered, and both runs use c512 resolution over the vortex. The evolution of the vortex corresponds to the high resolution regime, albeit delayed, with the clearly visible vortex ring. Figures 3.6a and b are more comparable to the c2048 uniform run at day 6 (Fig. 3.1d). Though not as pronounced, the two c64 based AMR runs are also delayed in their vortex evolution. The c64 2-level AMR runs with Tag 1 (3.6d) and Tag 2 (3.6d) both have a clearly defined vortex ring that has begun to stretch and deform. However, they are more similar in strength and appearance to the c512 uniform run at day 9 in Fig. 3.5b. In addition, the vortex in the c64 Tag 1 run is slightly stronger and more deformed, reflecting the two day advantage provided by the Tag 1 criteria in which

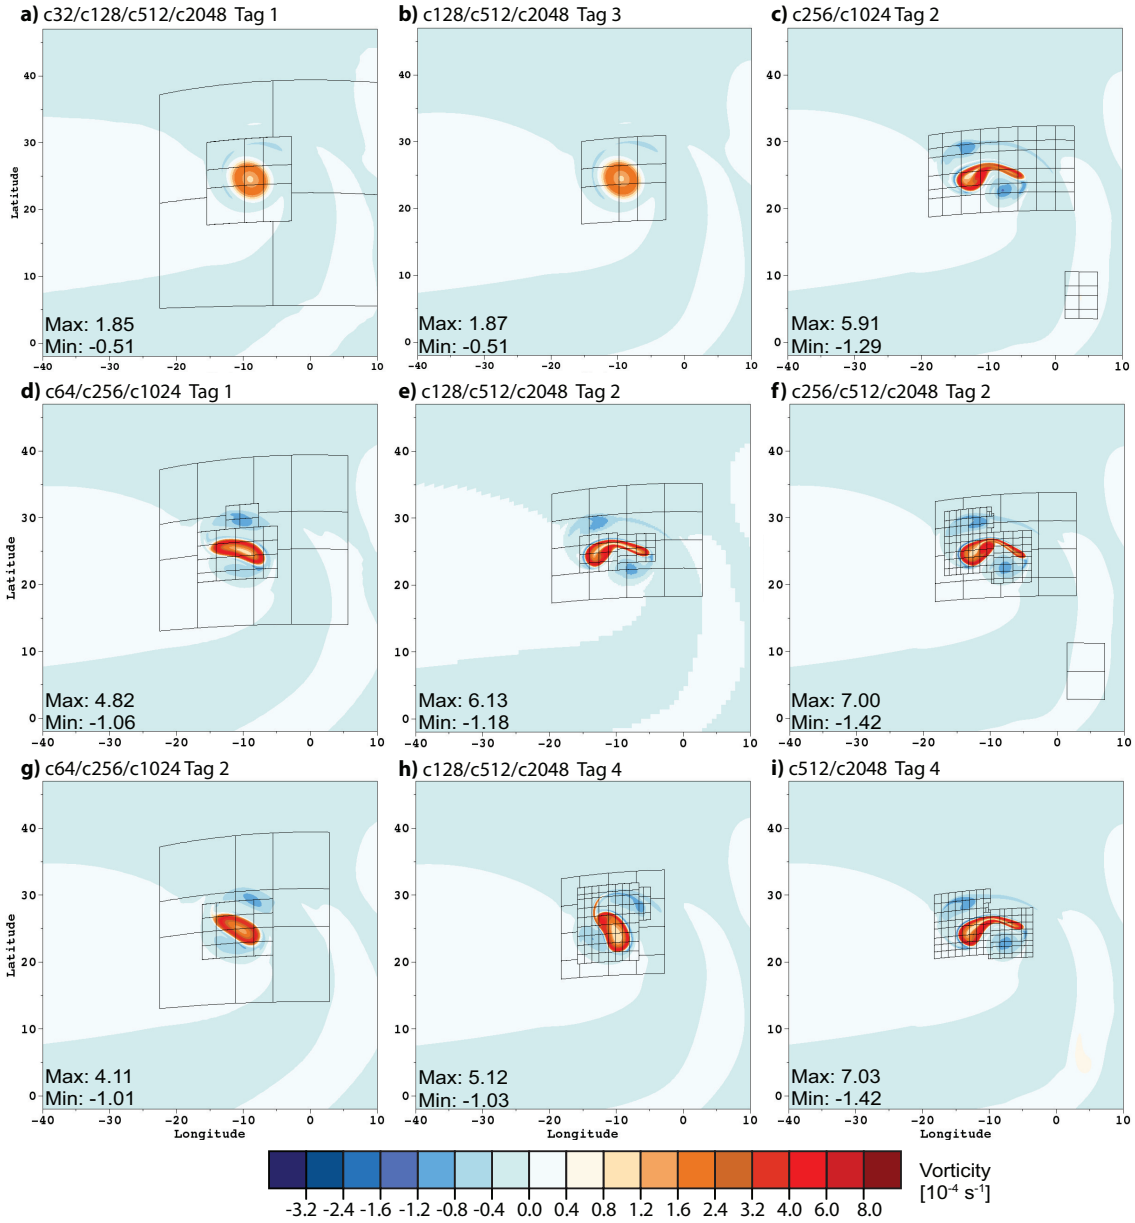

Figure 3.6: Relative vorticity fields at day 9 for several AMR runs of the strengthening vortex case. (a) c32 base 4-level AMR run with x4 refinement using Tag 1. (b) c128 base 2-level AMR run with x4 refinement using Tag 3. (c) c256 base 1-level AMR run with x4 refinement using Tag 2. (d) c64 base 2-level AMR run with x4 refinement using Tag 1. (e) c128 base 2-level AMR run with x4 refinement using Tag 2. (f) c256 base 2-level AMR run with one level of x2 refinement and one of x4 refinement using Tag 2. (g) c64 base 2-level AMR run with x4 refinement using Tag 2. (h) c128 base 2-level AMR run with x4 refinement using Tag 4. (i) a c512 base 1-level AMR run with x4 refinement using Tag 4. These plots correspond to the day 9 uniform plots in Fig. 3.1g and 3.5(a)-(c). The block structures of the multiple refinement levels are outlined in black

the c1024 AMR level is triggered at a lower threshold.

The c128 and higher resolution base level AMR runs are able to capture the vortex roll-up effectively at day 9. The three AMR runs using Tag 2 refinement criteria, c256/c1024 (Fig. 3.6c), c128 2-level AMR (Fig. 3.6e), and c256 2-level AMR (Fig. 3.6f, as well as the c512 1-level AMR Tag 4 refinement criteria run (Fig. 3.6i) capture the collapse and roll-up of the vortex and closely match the day 9 vorticity magnitudes seen in the c1024 and c2048 uniform runs. The AMR runs are also able to resolve the two areas of negative vorticity on either side of the main vortex. However, the roll-up process in the AMR runs is slightly delayed compared to the high resolution uniform runs. The distinct comma-like positive vorticity feature of the main vortex ring seen in the uniform c2048 (Fig. 3.1g) and c1024 (Fig. 3.5c) runs is not as prominently developed in these AMR runs. An exception to this is observed in the c128 base 2-level AMR run with Tag 4. Seen in Fig. 3.6h, the vortex ring in this Tag 4 run (Fig. 3.4d) is approximately 25% weaker than the c2048 uniform run and the vortex roll-up is slowed by half a day. Both the slight ring deformation and the northward extension of the positive vorticity filament from the northwest sector of the vortex ring correspond well to the structure of the vortices in the uniform c1024 and c2048 runs at 8.5 days.

Figure 3.7 depicts the vorticity fields of the AMR runs after vortex ring collapse and spin-off of a smaller secondary vortex. The c256 and c512 base resolution AMR runs (Figs. 3.7c, f, and i) exhibit vorticity maximums roughly 10% higher than their corresponding high resolution uniform runs reflecting Fig. 3.4. These three AMR runs most closely resemble the formation and location of the secondary vortex and capture the anticyclonic filaments wrapping aground the main vortex. The six AMR runs with coarser base resolutions have lower peak magnitudes. The vorticity fields for the c32 Tag 1 AMR run in Fig. 3.7a and the c128 Tag 3 AMR run in Fig. 3.7 show delayed spin-offs of the secondary vortex and less symmetric main vortices,

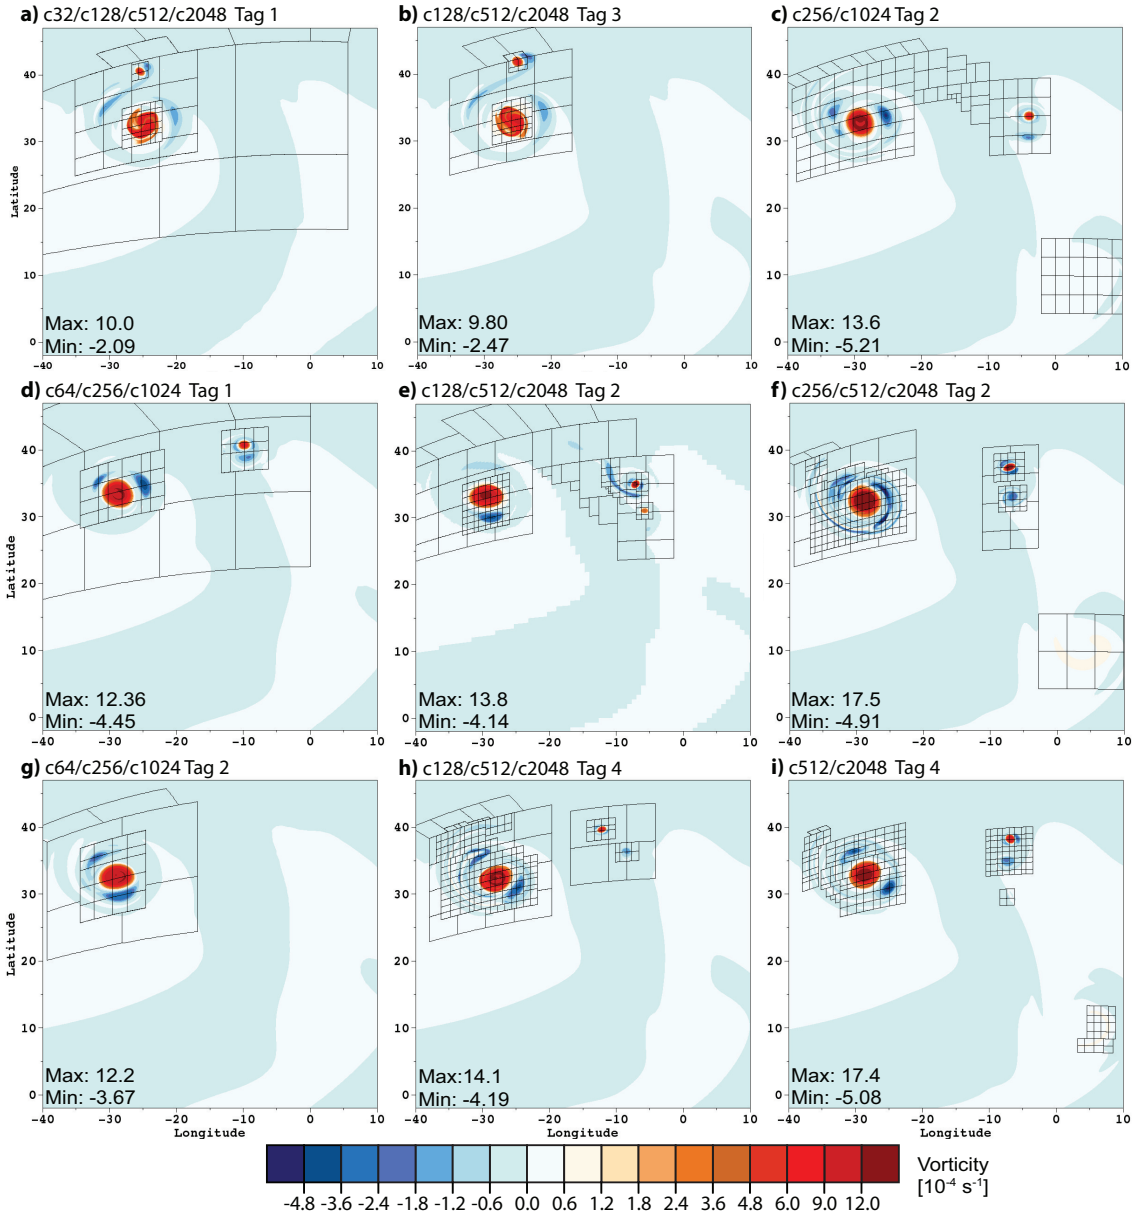

Figure 3.7: Same as Fig. 3.6, but for day 12 after the small secondary vortex has spun off. These plots correspond to the day 12 uniform plots in Fig. 3.1i and 3.5(d)-(f). The block structures of the multiple refinement levels are outlined in black

comparable to day 10 of the c2048 uniform run in Fig. 3.1h. The c64 2-level AMR Tag 1 run in Fig. 3.7d has a vorticity maximum similar to the c1024 uniform reference. This AMR run also resolves the secondary vortex, though its position is further to the northwest than that of the c1024 uniform run shown in Fig. 3.5 at day 12. In contrast, the slightly higher threshold of Tag 2 refinement criteria combined with the c64 2-level AMR run (Fig. 3.7g) is unable to reproduce the secondary vortex spin-off. The c128 AMR runs depicted in Fig. 3.7e and h both demonstrate development of a secondary vortex, though the Tag 2 run is more chaotic. This is shown by the large negative voricity filament area around the Tag 2 run secondary vortex and a weak third cyclonic vortex. The anticyclonic filaments are also more defined for the c128 2-level AMR Tag 4 run (Fig. 3.7h) than the Tag 2 run in Fig. 3.7e.

Figure 3.8 shows the number of grid cells at each AMR level as a function of time for both c128 2-level AMR runs. While the lower base level refinement threshold in Tag 2 results in the c512 level AMR being applied initially, refinement does not occur for two days with the Tag 4 criteria, at which time both levels of refinement are implemented due to the constant threshold value. Though both levels are implemented, the c2048 refinement level covers only the center core of the vortex and the outer edges of the vortex remain at c512 or c128 resolution for another day. After day 3, the Tag 4 run has almost 10x more cells than the Tag 2 run. That difference narrows significantly once the c2048 level is triggered for the Tag 2 run after day 6. Through the entire 16 day simulation, the Tag 2 run has more c512 grid cells while the Tag 4 run has more c2048 cells. Therefore, though the Tag 4 run has more grid cells and has triggered the c2048 refinement level for four days more than the c128 2-level AMR Tag 2 run, the Tag 2 run better captures the vortex roll-up at day 9 (Figs. 3.6e and h). The c512 resolution gained earlier, in the first two days, appears to outweigh the benefits of more c2048 resolution throughout the majority of the run time. The benefit of that early advantage fades as the run continues, and the benefit

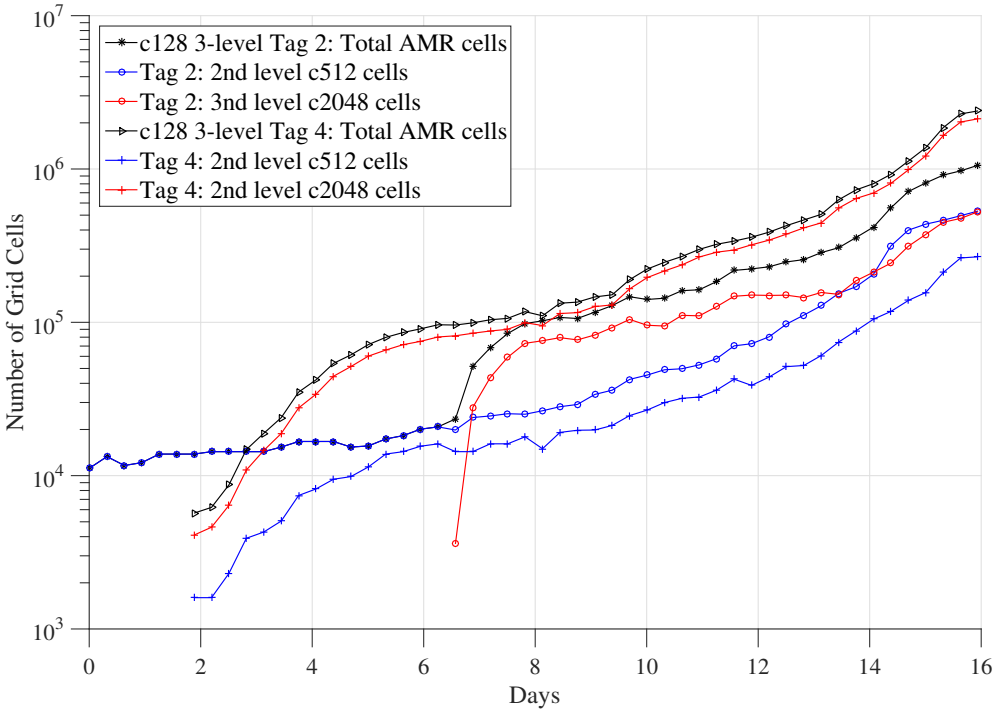

Figure 3.8: The growth of AMR grid cells over time for the two c128 2-level AMR runs with tag 2 and tag 4 refinement criteria. The base level c128 grid cells are excluded while the 2nd level c512 cells are plotted in blue and the third level c2048 cells are plotted in red with a plus marker to denote the tag 4 run and a circle to denote the tag 2 run. The sum of the two levels are plotted in black with a triangle marker denoting the tag 4 run and an asterisk marking the tag 2 run.

of more coverage by the c2048 refinement level becomes apparent. By day 12 the c128 Tag 2 run lacks the c2048 resolution around the core of the main vortex needed to resolve the thin negative vorticity filaments (Fig. 3.7e), which the Tag 4 run is able to capture (Fig. 3.7h). The secondary vortex in the Tag 4 run also more closely resembles the uniform c2048 run.

A key delineation between these AMR runs is apparent when c512 resolution or higher is implemented. At these levels of refinement, the vortex undergoes the high resolution evolution regime. The AMR runs with tagging criteria that triggered refinement levels of at least c512 initially, or within the first day, exhibited vortex growth most similar to the uniform c2048. The subset of these runs that do not trigger the c2048 refinement level until well into the simulation (six days or later) outperform AMR runs which have coarser than c512 resolutions initially but trigger c2048 resolution much earlier. Refinement, no matter what time it is applied still improves the results. Once c512 or higher refinement is triggered, rapid strengthening occurs and the vortex eventually transition to the high resolution evolution regime. The critical vortex collapse merely occurs later in time and we see some of those AMR runs can catch-up to the reference solution run by day 10 or 12.

3.4.1 Extended Run Results

Extending the simulation time leads to development of an active and complex global flow pattern. The isolated vortex discussed in detail in the previous sections, as well as the binary pair of weaker vortices, evolve independently of each other for twelve days. Beyond twelve days, the Rossby wave trains, though small in magnitude, spread and trigger evaporation and convection over a significant area. As a result, by day 16 of the simulation, a global chaotic flow has developed with multiple unseeded cyclonic and anticyclonic areas. Additionally, jet-like features develop around the 30°N and S latitudes which correspond to the areas of transition between high and

low moisture as determined by the initial moisture field and reservoir. Figure 3.9 depicts vorticity fields of all three vortices at day 9 and day 12 as well as the resulting global vorticity field at day 16 for uniform c256 (Fig. 3.9a), c512 (Fig. 3.9b), c1024 (Fig. 3.9c), and c2048 (Fig. 3.9d) resolution runs. The plots for day 9 and 12 are restricted to the three vortices as these are the primary features during that time frame. Day 16 plots depict the global relative vorticity field. The overall flow, number of large-scale vortices, and vortex locations remain fairly consistent across the increasing uniform resolutions. The higher resolutions naturally resolve more fine-scale structures in this chaotic system, though the uniform c1024 and c2048 runs also resolve several small-scale vortices as can be observed around 15°N and 30°W and nearby anticyclonic patches The high maximum vorticity in the c2048 run at day 16 is from the cyclonic vortex centered near 15°N and 30°W.

In addition to the uniform runs, we extend four AMR runs to 16 days of simulation. The global vorticity fields at day 16 for these runs are plotted in Fig. 3.10. The first column depicts the relative vorticity, while the second shows the block structure for the multiple AMR levels. All four AMR runs capture the overall jet-like flow of the system. The c64 base 2-level AMR run in Fig. 3.10a captures all the features with its intermediate AMR level. However, the highest c1024 level is limited to the main vortices and a few of the larger filaments. Thus much of the secondary flow and small scale features in this chaotic system diverge from the high resolution uniform runs. The c128 base 2-level AMR run in Fig. 3.10b, the c256/c1024 AMR run in Fig. 3.10c, and the 2-level c256/c512/c2048 AMR run in Fig. 3.10d all have a high resolution level over much of the critical areas that is better able to capture the small scale filaments and features. The one key metric in which the latter three AMR runs differ from their corresponding highest resolution uniform run is in the maximum vorticity. In the AMR runs, as in the c1024 and c2048 uniform runs, the vorticity maxima and minima are located in the small scale vortices that develop. However,

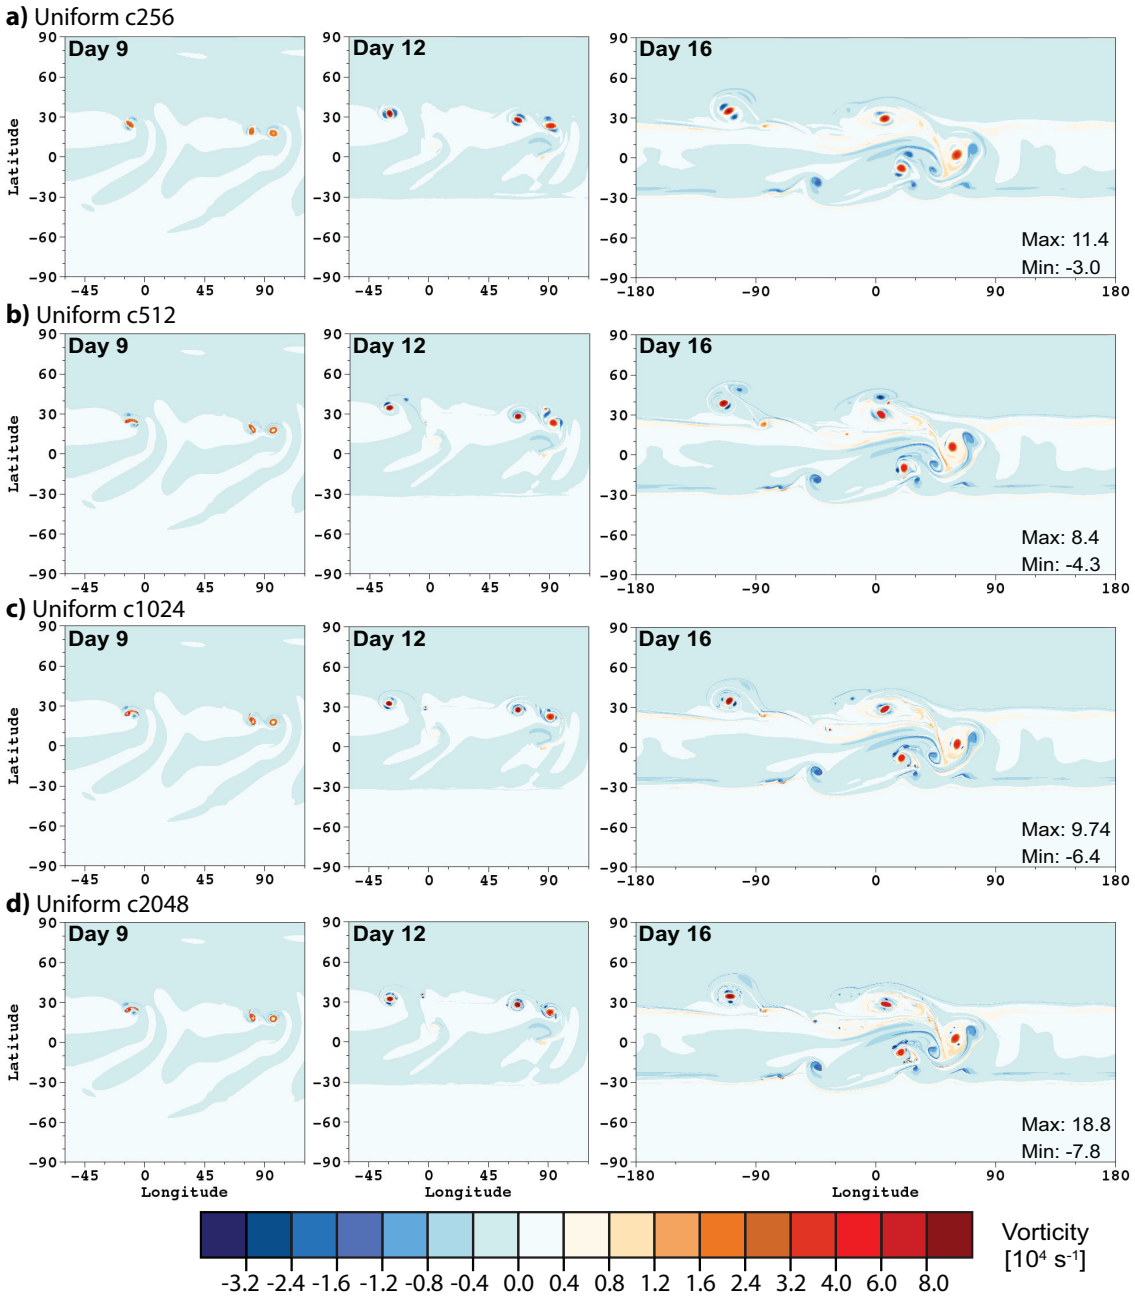

Figure 3.9: Late run evolution of the relative vorticity field for the strengthening vortex case with three initialized vortices. These plots show the growth of a global chaotic regime by day 16 in for uniform resolution runs. Relative vorticity snapshots at days 9 (left column), 12 (middle column), and 16 (right column) are given for (a) uniform c256 run, (b) uniform c512 run, (c) uniform c1024 run, and (d) uniform c2048 run. The leftmost vortex in the days 9 and 12 plots located around $(30°\text{N}, 15°\text{W})$ is the isolated vortex discussed in previous sections. The two vortices centered around $(20°\text{N}, 90°\text{E})$ are the binary pair. Note: the vorticity extrema occur in the isolated vortex in all cases for days 9 and 12 so they are not displayed.

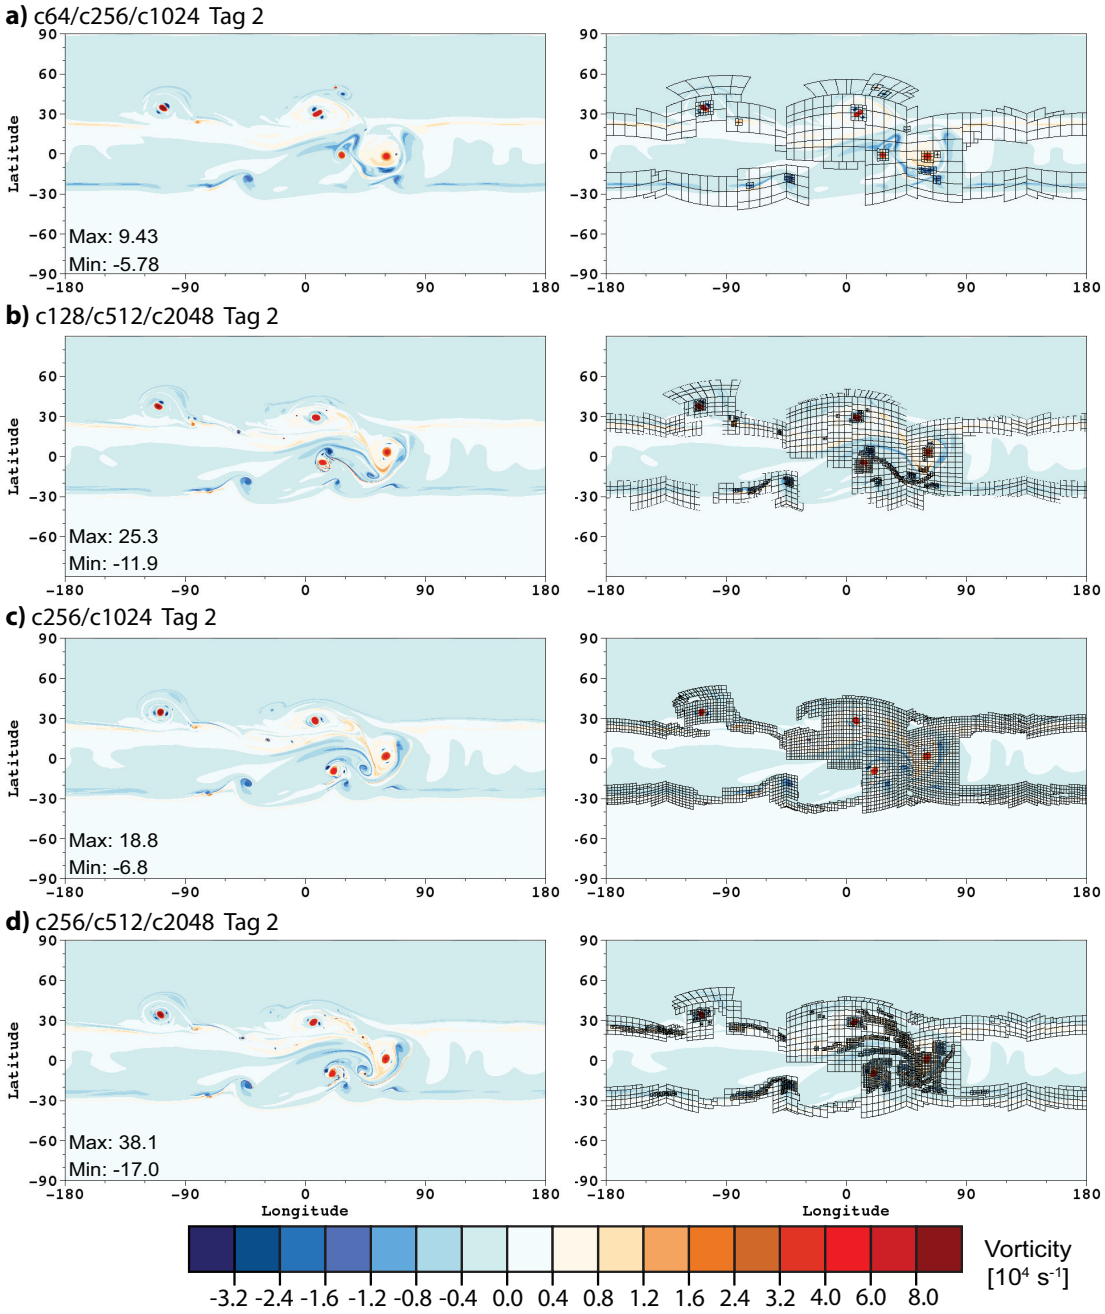

Figure 3.10: Relative vorticity fields at day 16 for four AMR runs of the strengthening vortex case with three initialized vortices: (a) c64 base 2-level AMR run with x4 refinement using Tag 2, (b) c128 base 2-level AMR run with x4 refinement using Tag 2, (c) c256 base 1-level AMR run with x4 refinement using Tag 2, and (d) c256 base 2-level AMR run with one level of x2 refinement and one of x4 refinement using Tag 2. The left column depicts the vorticity field at day 16 while the right column overlays the block structures of the refinement levels in black. These plots are comparable to the day 16 plots in Fig. 3.9.

these extremes for the AMR runs in Figs. 3.10b, c, and d are more than double the values seen in the uniform runs. The $2.53 \times 10^{-3}\text{s}^{-1}$ maximum in the c128 three-level run and the $3.81 \times 10^{-3}\text{s}^{-1}$ maximum in the c256 three-level run arise from small-scale vortices covered by small patches of the highest level of refinement. The maximum in the c256/c1024 AMR run is equally large compared to its uniform run counterpart ($1.8 \times 10^{-3}\text{s}^{-1}$) and is centered in a broad area of refinement that had been in place for several days. These vortices are well resolved with more than ten grid cells across and a significant buffer of refinement between the vortices and coarse-fine boundaries. Therefore it seems unlikely that these large vorticity values are being formed at AMR level boundaries. The cause is more likely to be the chaotic nature of this system at day 16.

3.5 Kessler-like Physics for Shallow Water Equations

An alternative setup for a moist forced shallow water system presented in Zerroukat and Allen (2015) derives the 2D shallow water equations from the moist Boussinesq equations. It includes a three-state moist physics model that simulates water vapor, cloud vapor, and liquid water, similar to the simplified 3D physics scheme presented in Kessler (1969). The forcing setup is comparable to the generalized shallow water equations of Ripa's model (Ripa, 1993, 1995) used in ocean modeling. In this model, latent heat released due to precipitation increases the average potential temperature of the fluid, which is coupled to the momentum equations. The model can be viewed as a diabatic non-convective model in contrast to the system used in Section 3.3 which can be considered an adiabatic convective model. It assumes energy released during precipitation has no effect on the temperature and instead produces convection represented as a mass flux. A brief discussion comparing the two models is presented in Appendix A of Bouchut et al. (2009). We implement this physics forcing here for both uniform and AMR runs of the barotropic instability test case of

Galewsky et al. (2004).

3.5.1 The Shallow Water and Physics Equations

Zerroukat and Allen (2015) dimensionally reduce the Boussinesq equations to obtain a derivation of the shallow water equations that retain some buoyancy terms. These augmented shallow water equations are presented below in flux-form:

$$\frac{\partial h\mathbf{v}}{\partial t} + \nabla \cdot (h\mathbf{v}\mathbf{v}) + f\hat{\mathbf{k}} \times (h\mathbf{v}) + gh\nabla H = gh\theta\nabla h_b + g\nabla(\frac{1}{2}h^2\theta) \tag{3.18}$$

$$\frac{\partial h}{\partial t} + \nabla \cdot (h\mathbf{v}) = 0 \tag{3.19}$$

$$\frac{\partial h\theta}{\partial t} + \nabla \cdot (h\theta\mathbf{v}) = hS_\theta \tag{3.20}$$

$$\frac{\partial hq^{(k)}}{\partial t} + \nabla \cdot (hq_{(k)}\mathbf{v}) = hS_{q_{(k)}} \tag{3.21}$$

where H is the total height, h_b is the bottom topography height, θ is the temperature based quantity, $q^{(k)}$ represents the moist physics tracers, and S_θ and $S_q^{(k)}$ are the depth averaged moisture and temperature sources, respectively.

The Zerroukat and Allen (2015) physics scheme consists of three forms of water vapor, q_v, cloud q_c, and rain q_r. When the local value of q_v exceeds a prescribed function for the saturation $q_{sat}(h, \theta)$, a fraction of the oversaturation is condensed into cloud, represented by Δq_v, with a corresponding latent heat release that increases the local temperature θ. In the same manner, a fraction of a cloud present in unsaturated air evaporates, Δq_c, with a corresponding cooling effect. In both cases, only a fraction of the water is converted to avoid a two time step oscillation between oversaturated and sub-saturated air induced by the changing temperature. Cloud can also be converted to rain when q_c exceeds a prescribed threshold q_{precip} and a fraction of the excess cloud is then converted to rain, Δq_r. It is important to note that the moisture q_v, q_c, and q_r and temperature θ variables as well the related constants like q_{sat} are

not associated with the typical physical units or value ranges used for such physical quantities.

In this setup, only q_v and q_c are advected as tracers as $q_{(k)}$ in Eq. 3.21. Once cloud moisture is transformed into rain, the rain water is removed from the system. Processes such as rain evaporation and accretion are neglected in this simplified model. The equations for the processes and source terms for the moisture variables ($S_q^{(1)} = S_{q_v},\ S_q^{(2)} = S_{q_c}$) from Eq. 3.21 and temperature S_θ are from Eq. 3.20.

$$\begin{aligned}
\Delta q_v &= \frac{1}{\Delta t} \max\left[0, \gamma_v \left(q_v - q_{sat}\right)\right] && (3.22)\\
\Delta q_c &= \frac{1}{\Delta t} \min\left[q_c, \max\left[0, \gamma_v \left(q_{sat} - q_v\right)\right]\right] && (3.23)\\
\Delta q_r &= \frac{1}{\Delta t} \max\left[0, \gamma_r \left(q_c - q_{precip}\right)\right] && (3.24)\\
S_{q_v} &= \Delta q_c - \Delta q_v && (3.25)\\
S_{q_c} &= \Delta q_v - \Delta q_c - \Delta q_r && (3.26)\\
S_\theta &= L\left(\Delta q_v - \Delta q_c\right) . && (3.27)
\end{aligned}$$

The constant γ_r is the rain conversion rate and L is a pseudo-latent heat constant for the θ variable. As derived in Zerroukat and Allen (2015), the saturation threshold $q_{sat}(h, \theta)$ and γ_v, the θ-dependent conversion rate between vapor and cloud moisture, are:

$$q_{sat} = \frac{q_0}{gH} \exp\left(20\theta\right) \tag{3.28}$$

where q_0 is a test case dependent constant set to have the initial max $q_v = 0.02$ and

$$\gamma_v = \left(1 + L\frac{\partial q_{sat}}{\partial \theta}\right)^{-1} . \tag{3.29}$$

Given the simplicity of the model, the constants can be arbitrarily chosen. However, we use the values described in Zerroukat and Allen (2015), where $L = 10$, $\gamma_r = 10^{-3}\mathrm{s}^{-1}$, and $q_{precip} = 10^{-4}$. The model does not have a sink for q_r, and q_r can instead

be viewed as water molecules suspended and advected in the gas phase. Additionally, there is no evaporation of q_r in this setup, though the equations can be easily modified to incorporate this phase-change. The Chombo-AMR model version used for these simulations does not preserve monotonicity or apply filters to tracers. So, negative undershoots can occur in the tracer fields, which needs to be remedied in future model versions.

3.5.2 Barotropic Instability Test Case Initialization

The barotropic instability test case of Galewsky et al. (2004) consists of a balanced zonal jet centered at 45° N to which a small height perturbation is added to initiate the roll-up of the jet. The initial velocity and height fields, along with the height perturbation, are defined in Galewsky et al. (2004). We add to that initialization the θ and q_v profiles, and set q_c and q_r to zero over the entire model space. The initial θ profile is a quadratic function with a north-south variation taken from Zerroukat and Allen (2015) so that

$$\theta(\phi,\lambda) = \theta^{SP}\left(\phi - \frac{\pi}{2}\right)\phi - \left(1 - \mu_1\theta^{EQ}\right)\left(\phi + \frac{\pi}{2}\right)\left(\phi - \frac{\pi}{2}\right) + \theta^{NP}\left(\phi + \frac{\pi}{2}\right)\phi. \quad (3.30)$$

The constants used for this test case are $\mu_1 = 2 \times 10^{-5}$, $\theta^{SP} = -40\epsilon$, $\theta^{EQ} = 30\epsilon$, and $\theta^{NP} = -20\epsilon$ where $\epsilon = 1/300$. The initial moisture profile is set just below the saturation level so $q_v(\phi,\lambda) = 0.98 q_{sat}(h,\theta)$, where $q_{sat}(h,\theta)$ is established from Eq. 3.28, where $q_0 = 0.0492238$.

3.5.3 Effects of resolution in the moist barotropic instability test case

The development of instability in the jet and evolution of the initial vorticity roll-ups into sharp gradients are consistent with the results in Galewsky et al. (2004). Significant cloud formation q_c does not begin until after four days, and that q_c does

not precipitate until five days into the simulation. By day six, the barotropic wave has created distinct vortices and thin vorticity filaments. Within these front and cutoff low-like features, areas of cloud and rain have formed. Figure 3.11 shows several variable fields for the barotropic wave at day 6 for a uniform c2048 ($\sim$ 5 km). The temperature θ in Fig. 3.11a and water vapor q_v in Fig. 3.11b echo the structure of the vorticity field (shown in Figs. 3.11c and d as black solid and dashed contour lines). The protrusions of colder and drier areas within the vorticity troughs mimic frontal systems. The q_c field is depicted in Fig. 3.11c, while Fig. 3.11d shows the total amount of water precipitated, q_r, in the preceding twelve hours. The highest areas of cloud and rain are within these vorticity troughs with smaller values of q_c located around the cutoff lows.

The effects of resolution and variable refinement on the barotropic instability's vorticity field has been well covered (see St-Cyr et al. (2008), Weller et al. (2009), and Scott et al. (2016)). So, we focus our investigation on how the cloud q_c and precipitation of the physics scheme are affected by changing resolution and AMR. The q_c fields at day 6 for four other uniform resolutions, c128, c256, c512, and c1024, are depicted in Fig. 3.12 for comparison with the c2048 run q_c plot in Fig. 3.11c. The accumulation of precipitated water q_r over a half day period before day 6 for the four uniform resolutions is plotted in Fig. 3.13 and corresponds to Fig. 3.11d for the uniform c2048 run.

Cloud cover area is fairly consistent across all resolutions in Fig. 3.12. The structure of the cloud field and smaller scale features centered on the cutoff lows and the leftmost wave converge at resolutions of c512 and higher. A key difference in cloud cover structure is in the c128 run in Fig. 3.12a. The c128 run has two extra areas of cloud cover between 80° and 170° longitude, which we interpret as artifacts of the cubed-sphere grid. As observed by St-Cyr et al. (2008) and Ullrich et al. (2010), the barotropic test case provides some difficulties for the cubed sphere. As the jet

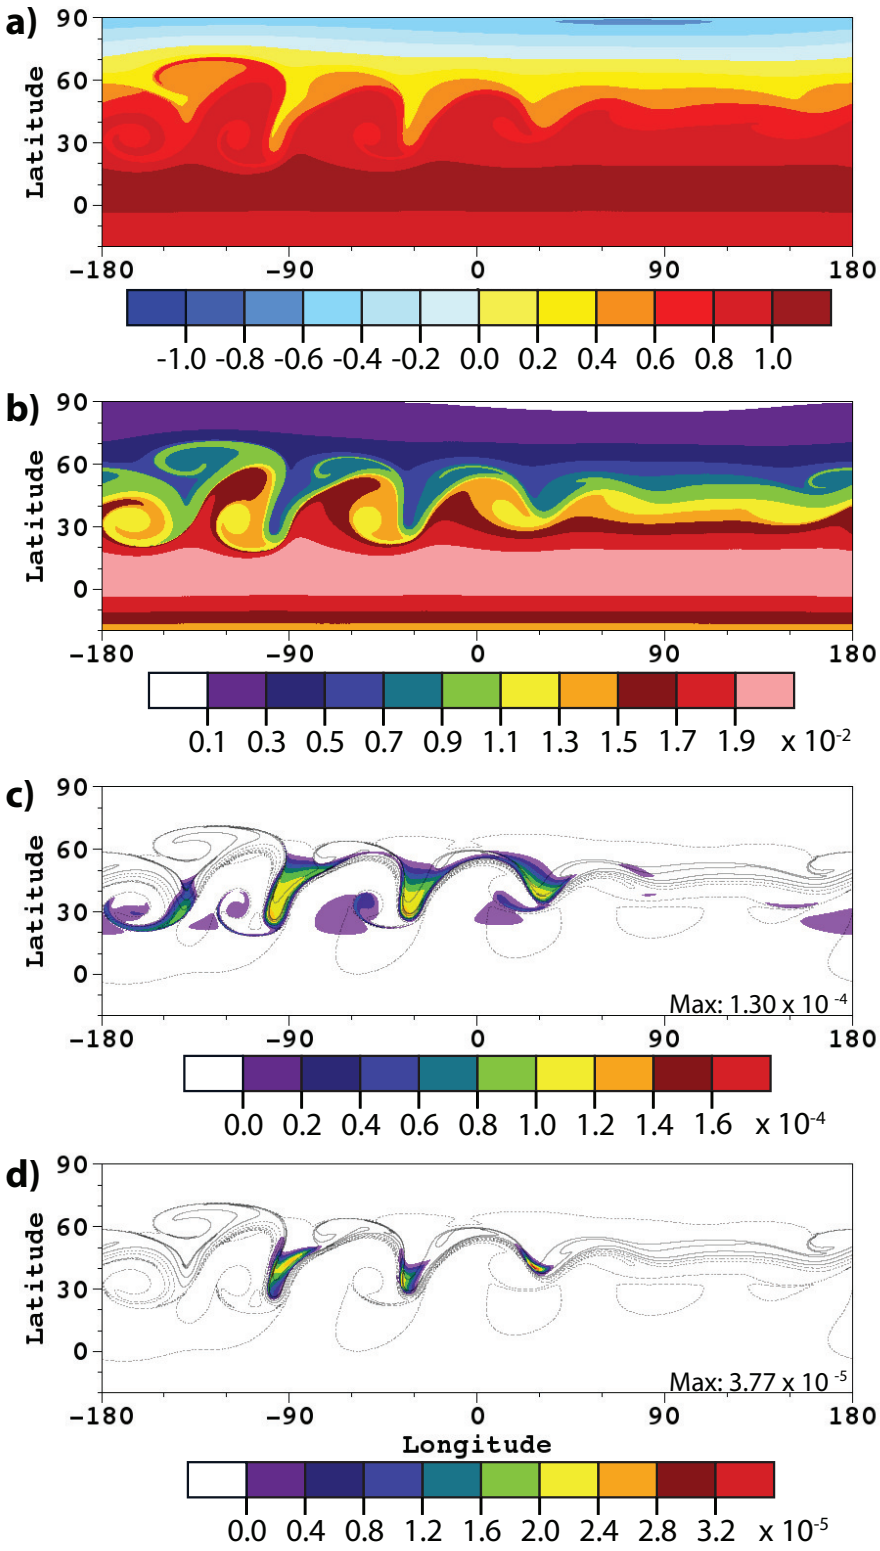

Figure 3.11: Day 6 snapshots of the evolving barotropic wave for the c2048 uniform run's (a) temperature field, (b) q_v moisture field, (c) q_c cloud field, (d) past 12-hour accumulation of the q_r precipitated water field. The solid and dashed black contour lines in (c) and (d) represent the positive and negative relative vorticity respectively. The spacing between contour lines is 5×10^{-5} s^{-1}.

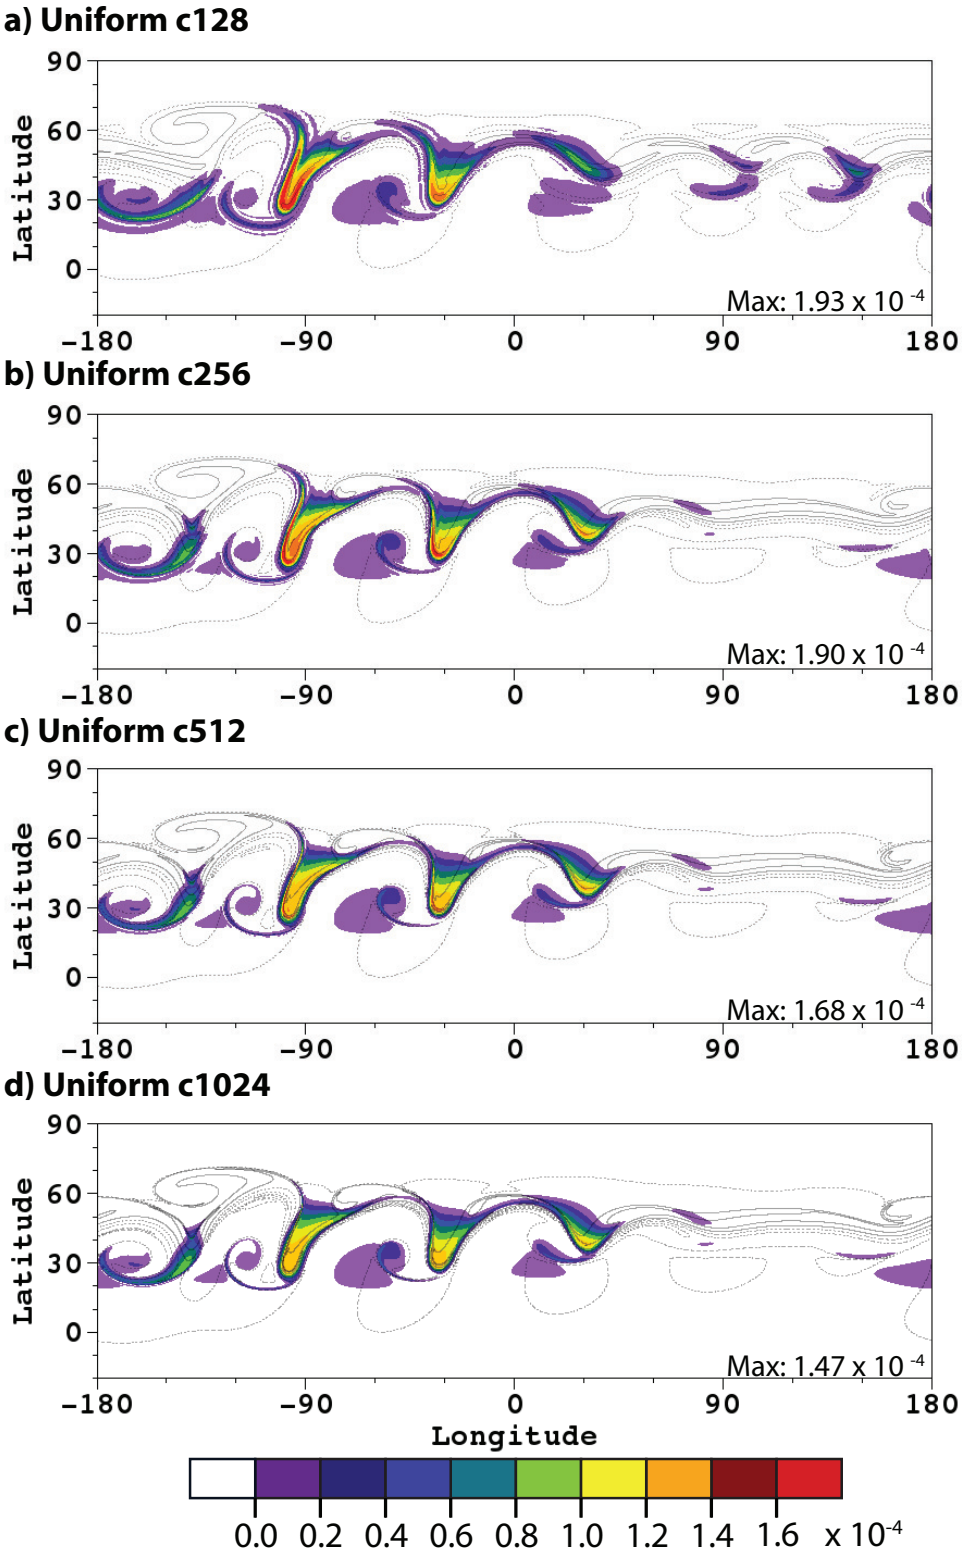

Figure 3.12: Plots of the q_c cloud field at day 6 for several uniform resolutions: (a) c128, (b) c256, (c) c512, and (d) c1024.The c2048 uniform run plot of the same field in Fig. 3.11c serves as a reference. The solid and dashed black contour lines represent the positive and negative relative vorticity respectively using the same contour spacing as in Fig. 3.11.

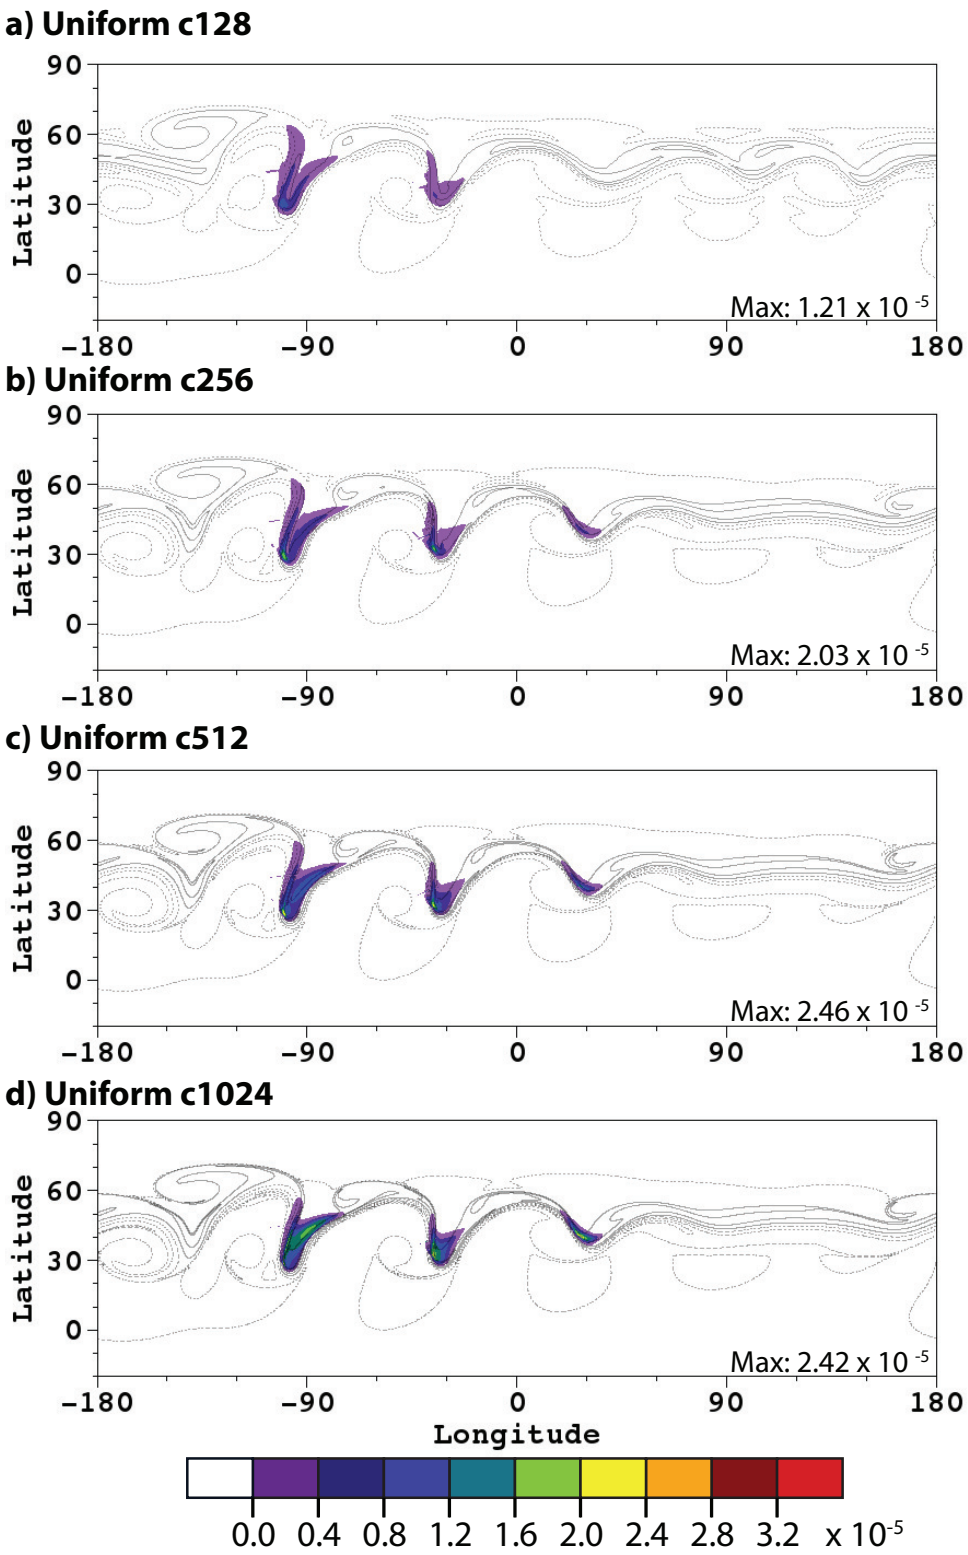

Figure 3.13: Plots depicting the 12-hour accumulation in the q_r precipitated water field for (a) c128, (b) c256, (c) c512, and (d) c1024 uniform runs. The c2048 uniform run plot of the same field in Fig. 3.11d serves as a reference. The solid and dashed black contour lines represent the positive and negative relative vorticity respectively using the same contour spacing as in Fig. 3.11.

moves over the four corners of the cubed-sphere, coarser resolution runs generate a wave number four forcing which affects the development of the solution. In the c256 run (Fig. 3.12b) the artifact has disappeared.

While the overall shape and area of the cloud field converges, the concentration of the q_c field continues to decrease with increasing resolution. The ribbon-like area of peak cloud concentrations on the edges of the two main troughs seen in Figs. 3.12a and b are not present in the higher resolution c1024 (Fig. 3.12d) and c2048 (Fig. 3.11c) runs. In the plots of 12-hour precipitation accumulation in Fig. 3.13, we observe the opposite trend with the maximum amount of precipitation accumulation nearly doubling between the c256 run (Fig. 3.13b) and the c2048 run (Fig. 3.11d). Like the q_c field, the overall coverage and structure of the rain field converges well with resolution, with only the area of heavier precipitation expanding as resolution increases. One key feature that shifts with increased resolution is the area of highest precipitation in the front-like system centered around $-90°$ longitude. In the day 6 precipitation plots for the c256 (Fig. 3.13b) and c512 (Fig. 3.13c) runs, the highest levels of accumulation are concentrated in a small area at the western edge of the bottom of the trough. In the higher resolution c1204 (Fig. 3.13d) and c2048 runs (Fig. 3.11d), the areas of most accumulation change to a broad area along the leading (eastern) edge and a secondary long and narrow area along the western edge.

3.5.4 The moist barotropic instability with AMR

For our implementation of AMR in this moist shallow water system, we created three refinement tagging criteria. In addition to the previously used relative vorticity thresholds, two criteria based on the physics variable q_c are used. The first q_c based criterion is a simple threshold of $q_c > 3.0 \times 10^{-5}$. The second cloud based criterion is designed to track the leading edges of cloudy areas and is based on the relative gradient of q_c. Its threshold is $|\nabla q_c| > 7.5 \times 10^{-8}$ m^{-1}. The simple relative vorticity

threshold is $|\zeta| > 2.3 \times 10^{-5}$ s^{-1}.

Figure 3.14 shows the q_c field at day 6 for six AMR runs using the three AMR tagging criteria stated above. The left column in Fig. 3.14 shows the cloud field overlaid with the block structure of the AMR levels, while the right removes them for easier viewing of the cloud field. At coarser resolutions, AMR runs using the q_c tagging criteria have only small areas of refinement in place after day 4. This still allows the grid imprinting to develop as seen in Fig. 3.14b, the c128 base resolution 1-level AMR run with the q_c threshold tag. Additionally, while the two heaviest cloud areas compare well their counterparts in the uniform c512 run (Fig. 3.12c), the two weaker areas of clouds centered around $-170°$ longitude and $30°$ longitude more closely resemble the uniform c128 run (Fig. 3.12a) The vorticity tagging criterion used for the c64 base resolution AMR run in Fig. 3.14a places higher resolution at the start of the simulation over the entire jet, which prevents grid imprinting. As a result of the larger area of refinement, this c64 AMR run achieves a similar result to the uniform c256 run (Fig. 3.12b).

For the c256 base resolution AMR runs with one-level of x4 refinement in Figs. 3.14c, d, and e, as well as the c128 base resolution two-level x4 refinement AMR run in Fig. 3.14f, the resolution over the jet is high enough to avoid the wavenumber four grid imprinting. The q_c fields for the three c256 base resolution AMR runs visually converge to the uniform c1024 run, albeit with some poorly resolved rings of low concentration along the edges of the main q_c areas in the two AMR runs using the q_c value and q_c gradient tags. This structure corresponds to the coarse-fine grid boundary, though it is caused by a response to the increased resolution rather than a computational artifact of the grid itself. In the c128 and c256 uniform runs, the main filaments of cloud in Figs. 3.12a and b are buttressed by thin parallel low density cloud filaments. These secondary filaments are either reduced in size or no longer present in the higher resolution uniform runs. In the AMR runs, the coarse-fine

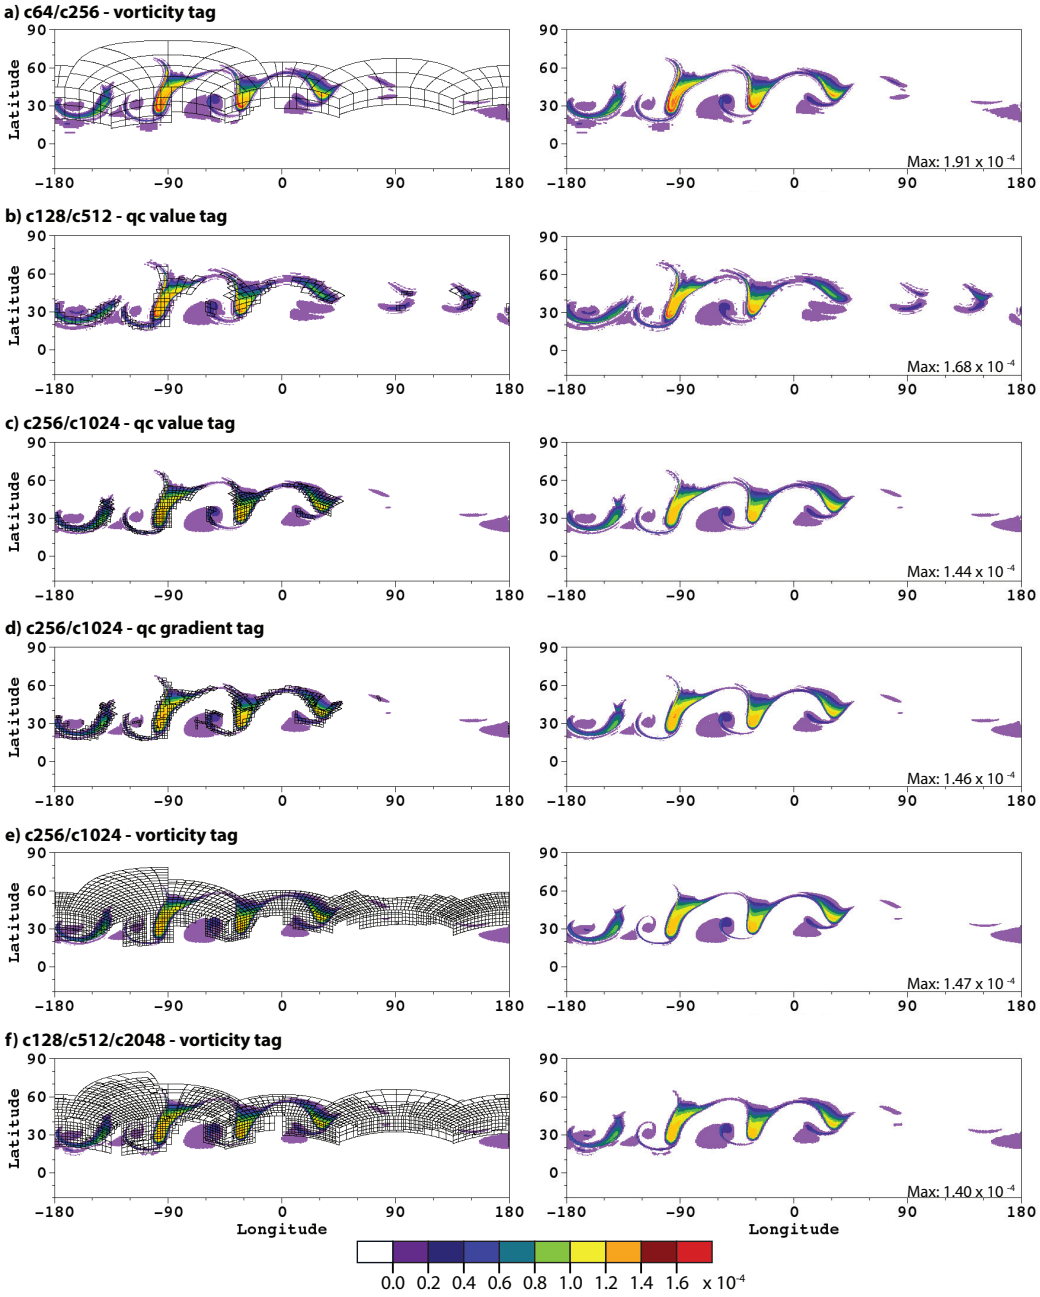

Figure 3.14: The cloud q_c field profile at day 6 for several AMR runs. The left column overlays the q_c variable with the block structures of the refinement levels in black, while the right columns removes these AMR blocks so that the q_c field can be viewed more clearly. (a) - (e) depict AMR runs with one level of x4 refinement while (f) depicts an AMR run with two levels of x4 refinement. The tagging criterion for (a), (e), and (f) is a relative vorticity threshold of $|\zeta| > 2.3 \times 10^{-5}$ s^{-1}. The criterion for (b) and (c) is $q_c > 3.0 \times 10^{-5}$, and the criterions for the AMR run in (d) is $|\nabla q_c| > 7.5 \times 10^{-8}$ km^{-1}.

boundary intersects these secondary filaments, causing the patchwork pattern seen in Fig. 3.14b, c, and d. It is less pronounced in Fig. 3.14e because the vorticity tagging threshold triggers refinement over a broader area, providing a buffer-like effect. The c128 two-level AMR run with vorticity tagging has a highest refinement level of c2048 over most of the jet for the duration of the model run time. Its q_c field at day 6, seen in Fig. 3.14f, depicts levels of cloud concentration (denoted by the orange and yellow contours) which are more comparable to the uniform c1024 run.

The same comparison to uniform runs can be made for the 12 hour q_r accumulation at day 6 of the six AMR runs pictured in Fig. 3.15. The AMR grid structure is the same as in Fig. 3.14, so the block structure is not shown in these plots. The overall structure of the accumulation compares to uniform runs for the q_c field. The area and magnitude of heaviest accumulation correspond well to the uniform resolution runs, with the same resolution as the highest resolution AMR runs; the c128 1-level AMR run in Fig. 3.15f is again more similar to the uniform c1024 run rather than the c2048 run. One key difference is observed in the c256 AMR run tagging on q_c gradient in Fig. 3.15d. For this AMR run, the area of higher accumulation on the eastern side of the front system, centered near $-90°$ longitude, is disjointed and compressed compared to the larger and smoother structure of accumulation in the corresponding c1024 uniform run in Fig. 3.12d. With the q_c gradient tagging, the interior areas of the large troughs have no high resolution refinement. This results in this precipitation accumulation structure observed as the coarser resolution has lower levels of accumulation. Another contrast is seen in the c128 and c256 AMR runs, which have large areas of low-level accumulation noise on the western sides of the two largest troughs. The noise is most pronounced in the three q_c tagging runs (Figs. 3.15b,c, and d), though present over a smaller area in vorticity tagging runs, specifically Fig. 3.15e and f. A lower cloud concentration threshold may reduce the noisy low-level edges, extending refinement out beyond the cloud formation areas.

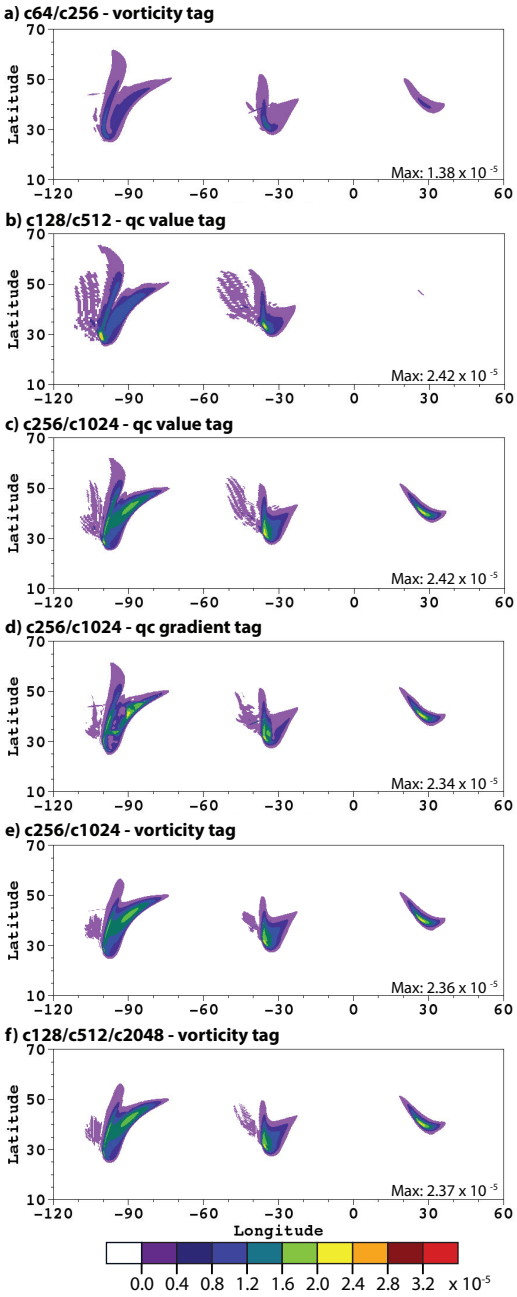

Figure 3.15: Past 12-hour accumulation of q_r at day 6 for the AMR runs depicted in Fig. 3.14.

3.6 Conclusion

We implemented two different forcing schemes designed to mimic the effects of moisture in the atmosphere within a 2D shallow water system in a fourth-order finite-volume model that is adaptive in both space and time. The first moist physics framework adds a water vapor variable and models convection as a mass sink triggered by saturation. We implemented a strengthening vortex test case with this setup. In the second forcing framework, a more complex moisture representation is used, consisting of vapor, cloud, and rain variables. The effects of moisture were coupled to the momentum equations through a potential temperature variable, linked to the moisture variables through latent heat. The model used this setup with the barotropic instability test case. With both forcing systems and test cases, we observe the evolution of features of interest at various resolutions and with differential refinement strategies. We also investigate AMR's effect on the physics forcing as grid resolutions changed. These simulations can aid the establishment of guidelines for effective AMR refinement criteria.

The resolution dependency of the physics forcing schemes were relatively mild. In the Kessler-like moisture scheme with the barotropic instability test, the overall structure of the moisture variables converged quickly with increasing resolution, though the concentration of clouds and precipitation accumulation continued to change with resolution. In the convective vortex forcing setup, we observed a resolution-dependent two-regime vortex evolution structure. For resolutions c512 and greater, we observed a convergence of vortex's shape and structure. Though the maximum wind speed and relative vorticity continued to slowly increase with resolution. The forcing in both cases functioned effectively across the varying resolutions and multiple levels of AMR. The coarse-fine interfaces did not induce noise or significant distortion

The sensitivity to resolution and AMR refinement criteria was much more pronounced in the strengthened vortex setup. The response of the moist barotropic wave

test case to AMR refinement criteria was fairly consistent; changes in criteria did not significantly alter the growth and structure of clouds and rain within the wave, so long as the initial refinement adequately resolved the wave to avoid computational grid artifacts. Any additional refinement was effective. In the strengthening vortex test case, the strength and evolution of the central vortex ring were quite sensitive to initial resolution and the time point at which AMR levels are triggered. Though the vortex does not strengthen significantly or undergo rapid structural changes during the first few days, we observed that AMR runs with solutions most similar to the uniform high resolution runs had some initial levels of refinement either initially (at least c256) or within the first day (at least c512). The vortices in these runs evolved at a comparable rate and strength to those in the uniform run with the same resolution as the highest AMR level, even if that level was not triggered until several days later. AMR not added initially was still beneficial. Once sufficient resolution was triggered, the vortex underwent the process of rapid strengthening, collapsing, and reforming. The high resolution vortex evolution process was merely delayed by lack of refinement. The application of refinement allowed the vortex to catch-up to the high resolution reference vortex. The time window in which AMR would trigger this process was limited. If high resolution AMR was not triggered until many days into the simulation, the AMR solution diverged from the high resolution runs.

Both sets of simulations show the starting grid resolution must be able to adequately resolve the features of interest to maximize AMR effectiveness. AMR cannot remove the errors caused before refinement begins. Additional refinement with AMR beyond that base level did improve the model, especially with regards to the small scale vorticity features in the strengthen vortex. To obtain this early refinement with AMR, the tagging criteria must be tailored to properties uniquely associated with the origins of the feature of interest, which is difficult even in these idealized shallow water systems, or must use a combination of initial static refinement and AMR.

For example, in tracking and resolving TCs in a realistic climate simulation, a static region of refinement could be placed over regions of cyclogenesis. Any storms that develop could be further refined with AMR tagging on surface pressure and followed as they traverse and exit the region of static refinement. Future work will consist of extending the analysis to AMR in the full 3D dycore, focusing on similarly simplified physics parameterization schemes.

CHAPTER IV

Implementing AMR in the 3D dynamical core with dry and simple physics test cases

4.1 Introduction

Variable resolution grids have been used in regional atmospheric models for several decades (Fox-Rabinovitz et al., 1997). The most common downscaling approach was the nested grid, in which a fixed-size refined grid is embedded within the coarse grid at a set location. It was implemented in several limited area models (LAM) for local weather forecasting and regional climate simulations (e.g. Phillips (1979), Pielke et al. (1992), Grell et al. (1994), and Caya and Laprise (1999)). Many of these models are regularly used today. For example, one of the most frequently used LAMs with embedded nested grids is the Weather Research and Forecasting Model (WRF, Skamarock et al. (2008)).

An alternative approach to downscaling designed for global models is the stretched grid that was developed around the same time (Schmidt, 1977; Staniforth and Mitchell, 1978). In these models, an originally uniform mesh is smoothly distorted such that more grid elements are concentrated over a localized region, providing higher resolution in that area and leaving fewer grid elements, and thus coarser resolution, in the remaining area. Stretched grids require few modifications to the numerical schemes

and computational grids to implement in general circulation models (GCMs). Thus they were an attractive technique to implemented in variable resolution modeling because GCMs avoid many drawbacks associated with LAMs such as lateral boundary conditions and their inability to capture upscaling effects. Several stretched grid models developed at that time include Paegle (1989), Déqué and Piedelievre (1995), Yessad and Bénard (1996), Fox-Rabinovitz et al. (1997), and Côté et al. (1998)).

Early GCMs using nested grids were developed by Ruge et al. (1995) and Dudhia and Bresch (2002). Nested grids within GCMs involve continuous two-way communication between resolutions and require more challenging numerical schemes and computational grids. This made them initially less attractive in comparison to stretched grids. A second form of nesting involves a single grid spanning multiple resolutions, where refined grid cells physically replace the coarse cells in the region of interests. An early example of this technique was implemented by Fournier et al. (2004).

With the advancement of large-scale parallel computing systems, variable resolution GCMs are a growing fixture in atmospheric and climate modeling. Variable resolution has been incorporated into several operational GCMs across major modeling centers (Skamarock et al., 2012; Harris and Lin, 2013; Zarzycki et al., 2014b). This combination has proven to be effective for assessing tropical cyclones (Zarzycki and Jablonowski, 2014, 2015; Huang et al., 2017), large-scale weather systems (Rauscher and Ringler, 2014), atmospheric rivers (Hagos et al., 2015), and regional climate (Medvigy et al., 2013; Huang et al., 2016; Gettelman et al., 2017; Rhoades et al., 2018). Specifically, variable resolution can bridge the difference in scale between global and regional climate modeling by overcoming many of the known issues with conventional downscaling methods of LAMs or high resolution GCMs. They provide high-resolution in a desired location while eliminating the need for forced lateral boundary conditions, capture small-to-large scale teleconnections, and demand fewer computational resources compared to standard uniform high-resolution GCMs.

The variable resolution models discussed thus far implement static grid refinement. The refined grid's location is determined *a priori* and remains fixed for the entirety of the simulation. Dynamic grid refinement, such as adaptive mesh refinement (AMR), is a more flexible variable resolution modeling technique. Adaptive grids track features of interest and refine the grid locally in advance of any important physical features or processes requiring additional resolution. When the additional resolution is no longer needed, that same region of the grid is coarsened. While dynamic refinement is frequently implemented in other areas of computational hydrodynamics such as aerospace and heliophysics, it has not been widely adapted for atmospheric and climate modeling. Since dynamically adaptive grids were first applied to atmospheric flows by Skamarock et al. (1989), Skamarock and Klemp (1993) and Dietachmayer and Droegemeier (1992), adaptive model work has been limited to algorithm development, simplified models, and idealized simulations. Dynamically adaptive shallow water models on the sphere have been discussed in Giraldo (2000), Behrens et al. (2005), Läuter et al. (2007), St-Cyr et al. (2008), Kubatko et al. (2009), and Chen et al. (2011a). More recent model designs have been presented in McCorquodale et al. (2015), Tumolo and Bonaventura (2015), Aechtner et al. (2015), and Weller et al. (2016). Several applications of AMR and dynamic refinement in the shallow water system include tsunami propagation in Blaise and St-Cyr (2012), Fujiwhara interactions in Bauer et al. (2014), and tropical cyclone eye-wall evolution in Hendricks et al. (2016). Ferguson et al. (2016) provides more detailed overview of these models and their dynamic refinement techniques. Though some LAM models like WRF (Shepherd and Walsh, 2017) have primitive moving nested grids, these cannot be considered true AMR models. The nests in these models are often fixed in size and grid elements cannot be added or removed. The only full 3D atmospheric model with dynamic refinement known is the OMEGA model by Bacon et al. (2000). This adaptive non-hydrostatic limited-area model implemented on rotated Cartesian coor-

dinates has been used as a regional hurricane forecasting system in Gopalakrishnan et al. (2002).

This chapter explores the use of mapped-multiblock AMR techniques for three-dimensional atmospheric flows. We use the nonhydrostatic finite-volume Chombo AMR model on the cubed-sphere extended to the full 3D atmospheric equations from the 2D shallow water model of McCorquodale et al. (2015). We wish to characterize the ability of this model and its multiblock refinement to simulate atmospheric flow in the global 3D system. To this end, we implement the Lin et al. (2017) modon test case in the dynamical core (dycore) without any forcing or subgrid physical parameterization. We run convergence tests and demonstrate the use of AMR with simple tagging criteria to observe how effectively AMR resolves and tracks the modons. We also implement a second test case, the idealized tropical cyclone from Reed and Jablonowski (2011), in which a single, idealized initial weak vortex evolves into a tropical cyclone. This test requires the simplified physics scheme from Reed and Jablonowski (2012) consisting of parameterizations for large-scale condensation, boundary layer diffusion, and surface fluxes for moisture, sensible heat, and momentum. With the idealized vortex, we investigate the AMR's ability to track and resolve the strengthening cyclone using various refinement criteria and thresholds to trigger refinement. The results presented are an preliminary assessment of the AMR capabilities of the 3D models. This model is still undergoing development and we are currently discovering, evaluating, and resolving instabilities. Several of these issues are noted in following sections.

This chapter proceeds as follows. Section 4.2 contains a brief description of the the nonhydrostatic 3D finite-volume model on the cubed-sphere. In section 4.3, we present the results of the colliding modon test, including convergence tests. The results of the idealized tropical cyclone test and the comparisons of various AMR refinement criteria are discussed in section 4.4. In Section 4.5 we summarize our

conclusions and comment on the direction of future research.

4.2 Experimental design

The non-hydrostatic finite-volume Chombo AMR model is a global 3D model built upon the shallow water version presented in McCorquodale et al. (2015). The dycore utilizes the full-non-hydrostatic moist fluid equations in a shallow-atmosphere approximation on a cubed-sphere grid.

4.2.1 Cubed-Sphere Grid

The cubed-sphere grid, first developed by Sadourny (1972), consists of a cube with six Cartesian panels inflated out to form a spherical shell. The grid avoids this pole-problem of traditional latitude-longitude grids by replacing the two strong singularities at the poles with eight weaker singularities at the corner points of the originally cube. The cubed-sphere grid also provides a near uniform tiling of the sphere as compared to large changes in grid spacing on the latitude-longitude mesh. There are multiple ways to map the grids of each panel to the sphere (see Putman and Lin (2007) for a review of several cubed sphere grids). The Chombo-AMR model uses gnomonic equiangular cubed-sphere grid, where the gridlines on each panel have equally spaced central angles relative to the center of the sphere. This projection yields a quasi-uniform spherical grid. The grid does not have perfectly uniform tile sizes, but as resolution increases the ratio between the largest and smallest grid cells converges to $\sqrt{2}$, the smallest ratio of cubed-sphere grids. Figure 4.1 depicts the equiangular cubed-sphere grid.

The physical domain of the model is a spherical shell $r \in [R_i, R_f]$ with a thickness of $H = (R_f - R_i)$, where R_i and R_f are the radii of the model bottom and top respectively, and the assumption that $H \ll R_i, R_f$. This shallow atmosphere assumption allows us to neglect r metric dependencies (e.g. no area increase with altitude). The

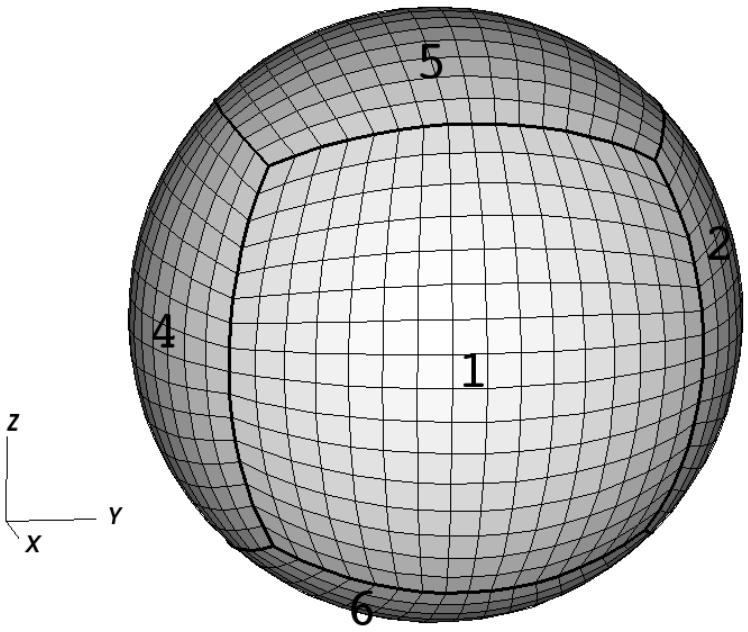

Figure 4.1: A cubed-sphere grid, shown with labels on panels. Panels 1 – 4 all straddle the equator ($z = 0$) of the unit sphere. Panel 5 is centered on the north pole ($z = +1$), panel 6 on the south pole ($z = -1$). On the cubed-sphere grid shown here, $N_c = 16$ (each panel contains 16×16 grid cells).

physical domain is mapped from the cubed-sphere grid.

The equiangular coordinates system for the cubed-sphere grid is given as $(\alpha, \beta, \xi, n_p)$, defined on six panels $n_p \in [1, 2, \ldots, 6]$, with central angles $\alpha, \beta \in [-\frac{\pi}{4}, \frac{\pi}{4}]$ and vertical coordinate $\xi \in [0, 1]$. By convention, panels 1-4 are along the equator and panels 5 and 6 are centered on the north and south poles, respectively, as seen in Fig. 4.1. Each panel is discretized into an $N_c \times N_c$ grid of quadrilaterals. The edges of each grid cell are non-orthogonal segments of the great circle in either the α or β direction. This property makes the computational grid sizes constant so that $d\alpha = d\beta = \frac{\pi}{2N_c}$. The discrete horizontal resolution of the cubed-sphere grid is represented as $c\{N_c\}$. A list of the horizontal properties of the equiangular cubed-sphere grid, including the approximate grid spacings, and comparable resolutions for other coordinate systems, is given in Table 2.1 for several resolutions.

The vertical direction is discretized into N_v layers with non-constant thickness but

constant thickness in the computational domain so that for each level $\Delta\xi = 1/N_v$. The mapping $r(\xi)$ between the physical and computational domains is such that $r(\xi = 0) = R_a$ and $r(\xi = 1) = R_a + H$ where R_a is the radius of the spherical shell and H is the height of the model top. The height-based vertical mapping is set by a user-created array consisting of the coordinate positions for the vertical level interfaces in physical space. For non-uniformly spaced vertical maps, the final positioning of the level interface points are smoothed via a cubic spline.

4.2.2 Fluid equations in cubed-sphere coordinates

The dycore utilizes the following state variables

$$\boldsymbol{S} = \begin{bmatrix} \langle J\rho u^{\alpha} \rangle \\ \langle J\rho u^{\beta} \rangle \\ [w]_f \\ [p]_c \\ \langle J\rho \rangle \\ \langle J\rho\theta_v \rangle \\ \langle J\rho q_v \rangle \end{bmatrix}. \tag{4.1}$$

Angled brackets $\langle\rangle$ indicate a cell averaged variable for the density ρ, the horizontal momentum density variables ρu^{α} and ρu^{β} in the α and β directions on the cubed-sphere, respectively. The virtual potential temperature is θ_v and the specific humidity q_v which will serves as a placeholder for other physics tracers. Vertical velocity w is face-centered variable indicated by $[]_f$ and pressure p is a cell-centered variable, indicated by $[]_c$. J is the metric Jacobian. The relationship between the physical vertical velocity w and the computational vertical velocity u^{ξ} is given by

$$w = r_{\xi} u^{\xi}. \tag{4.2}$$

The relationship between the derivates in the computational vertical coordinate ξ and physical vertical coordinate r is

$$\partial_r = r_\xi^{-1} \partial_\xi. \tag{4.3}$$

In both equations, r_ξ is the vertical coordinate transform term for converting between physical and computational spaces. The virtual potential temperature is defined as

$$\theta_v = (1 - 0.61 q_v) T \left(\frac{p_0}{p} \right)^\kappa, \tag{4.4}$$

where T denotes the temperature, $p_0 = 10^5$ Pa is the reference pressure and $\kappa = R_d/c_{pd}$. Here R_d is the ideal gas constant and c_{pd} is the specific heat at constant pressure, both in the case of dry air.

With these variables the equations of motion in conservation form on the cubed sphere grid are:

$$\frac{\partial J\rho u^\alpha}{\partial t} = -\sum_{i=\alpha,\beta} \partial_i (J\rho u^\alpha u^i + JG^{\alpha i} p) - \partial_\xi (J\rho u^\alpha u^\xi + JG^{\alpha\xi} p) + J\Psi_C^\alpha + J\Psi_M^\alpha \tag{4.5}$$

$$\frac{\partial J\rho u^\beta}{\partial t} = -\sum_{i=\alpha,\beta} \partial_i (J\rho u^\beta u^i + JG^{\beta i} p) - \partial_\xi (J\rho u^\beta u^\xi + JG^{\beta\xi} p) + J\Psi_C^\beta + J\Psi_M^\beta \tag{4.6}$$

$$\frac{\partial w}{\partial t} = -\sum_{i=\alpha,\beta} u^i \partial_i w - u^\xi \, \partial_\xi w - \frac{1}{\rho} \partial_r p - g \tag{4.7}$$

$$\frac{\partial p}{\partial t} = -\sum_{i=\alpha,\beta} \left(\gamma p \frac{1}{J} \partial_i (J u^i) + u^i \partial_i p \right) - u^\xi \partial_\xi p - \gamma p \, \partial_r w + \eta \left(P(\rho\theta_v) - p \right) \tag{4.8}$$

$$\frac{\partial J\rho}{\partial t} = -\sum_{i=\alpha,\beta} \partial_i (J\rho u^i) - \partial_\xi (J\rho u^\xi) \tag{4.9}$$

$$\frac{\partial J\rho\theta_v}{\partial t} = -\sum_{i=\alpha,\beta} \partial_i (J\rho\theta_v u^i) - \partial_\xi (J\rho\theta_v u^\xi) \tag{4.10}$$

$$\frac{\partial J\rho q_v}{\partial t} = -\sum_{i=\alpha,\beta} \partial_i (J\rho q_v u^i) - \partial_\xi (J\rho q_v u^\xi) \tag{4.11}$$

where G^{ij} represent the contravariant cubed-sphere metric term, g is the acceleration due to gravity, and $\gamma = 1/(1-\kappa)$. Ψ_C and Ψ_M represent source terms from the Coriolis

force and metric terms due to the cubed-sphere geometry, respectively. These source terms and coordinate transforms that arise from using cubed-sphere mapping may be found in Appendices A and B of Ullrich and Jablonowski (2012). Currently no topography is implemented. Finally the moist equation of state for the equations of motion is

$$P = p_0 \left(\frac{R_d \rho \theta_v}{p_0} \right)^{\gamma}. \tag{4.12}$$

Equation 4.8 includes a volume discrepancy term: $\eta \left(P(\rho\theta_v) - p \right)$ with relaxation parameter $\eta = 0.5$ to prevent the prognostic pressure p from drifting from the equation of state pressure P. The volume discrepancy approach is used in gas dynamics modeling to couple non-linear explicit flow solvers and stiff reactions and to maintain conservation in potential temperature density (Day and Bell, 2000). Because the equations of motion are closed with the equation of state Eq. 4.12, this leads to a non-linearity in θ_v which leads to pressure values drifting apart. Note that the redundant pressure equation makes it convenient to treat acoustic waves implicitly. The volume discrepancy term relaxes the prognostic pressure to the equation of state pressure.

4.2.3 Numerical Methods

The non-hydrostatic equations are spatially discretized with a finite-volume scheme that was implemented by McCorquodale et al. (2015) for the shallow water equations. The 3D scheme computes the horizontal fluxes with fourth-order accuracy, but is only accurate to second-order in the vertical. To perform the finite volume calculations on panel edges, the mapped-multiblock approach creates three layers of ghost cells and remaps the values from neighboring cells to these ghost cells. In addition, a sixth-order diffusive operator which maintains the fourth-order accuracy of the scheme is applied to the horizontal fluxes. There is no dissipation in the vertical. 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. Since topography is not yet mplemented in the dycore, simple height-based vertical coordinates are used.

Time integration is conducted by an implicit-explicit second-order, three stage ARS232 scheme by Ascher et al. (1997). Given the large aspect ratio in the vertical direction, vertical acoustic waves, which are supported by the shallow atmosphere equations of motion, would severely limit the global time step. Thus, the terms responsible for vertical acoustic waves in the w and p equations are separated out to be treated implicitly. All the horizontal terms and the remaining advection terms for the vertical variables are evaluated explicitly.

The mapped-multiblock AMR techniques developed from the Chombo library (Adams et al., 2015) and used for the shallow water model in McCorquodale et al. (2015) and Ferguson et al. (2016) are implemented in the full 3D dycore. AMR calculations are performed on a hierarchy of nested meshes, called levels, which have a defined refinement ratio between them. The finer levels are overlaid on top of the coarser levels. Information on these finer level is initialized via interpolations from the coarser level and ghost cells are used to calculated the fluxes at patch boundaries. Information from the finer level is averaged down to update values on the coarser level. These levels are sub-cycled in time to maintain a constant Courant number across all levels. A detailed overview of the AMR level strucutre and sub-cycling time integration is provided in Chapter II. Refinement is only done in the horizontal direction, vertical refinement is not currently implemented. When an AMR level is initialized, refinement is added uniformly throughout all vertical levels.

Several modifications have been made to achieve a working dycore to present these results. Refluxing across panel boundaries is currently not implemented so mass conservation is not guaranteed at panel edges or at coarse-fine boundaries. Additionally, the fourth-order horizontal discretization drops to second order at panel

boundaries, though coarse-fine boundaries are still fourth-order. A simple clipping limiter is used to prevent moisture variables from going negative. An optional ability to incorporate Rayleigh damping in the upper atmosphere has also been implemented. Rayleigh damping is added as a source term in the form

$$\Psi_R = -R_c(\alpha, \beta, \xi) \left(\rho \mathbf{u} - \rho \mathbf{u_0}\right), \tag{4.13}$$

where R_c denotes the strength of the damping, $\rho\mathbf{u}$ is the 3D momentum density vector, and $\rho\mathbf{u_0}$ denotes a reference state for the momentum. For these simulations, $\mathbf{u_0}$ is set to zero. The strength of the damping term is designed to smoothly transition from zero damping at lower levels to a maximum at the model top. We choose

$$R_c(\alpha, \beta, \xi) = 0 \qquad \text{if } \xi < \xi_r, \tag{4.14}$$

$$R_c(\alpha, \beta, \xi) = \frac{1}{\tau_R} \left(\frac{\xi - \xi_r}{1 - \xi_r}\right)^2 \qquad \text{if } \xi > \xi_r. \tag{4.15}$$

Here τ_R is the timescale of the damping and ξ_R is the starting height of the damping layer in ξ coordinates. We use $\tau_R = 1$ day and $\xi_R = \frac{2}{3}$, which places damping in the upper third of the atmosphere only. Rayleigh damping is currently not used in the modon simulations but is activated for the idealized tropical cyclone tests. No additional limiters or filters are implemented.

4.3 Modon test case implemented in the dry dynamical core

The Lin et al. (2017) colliding modon test consists of two anti-polar pairs of counter-rotating vortices, which are centered on the equator and overlay a calm, hydrostatic background. We implement it in a dry version of the dycore for use in some basic convergence studies and a first look at the efficacy of AMR. The modons are slow-moving, isolated features that can be easily tagged for refinement with a vorticity

threshold. These characteristics make the test case effective for discerning problems that might arise due to refinement, instabilities due to the numerical schemes, and noise or wavelike reflections at AMR boundaries.

4.3.1 Initialization and basic characteristics

Each modon is a dipole of positive and negative vorticity regions on either side of the equator. In this test case by Lin et al. (2017), two modons are initialized at the equator on opposite sides of the sphere. The two modons initially undergo rapid cyclostrophic adjustment due to unbalanced initial conditions. This creates a gravity wave which propagates around the globe but has no effect on the structure of the modons. The modons travel slowly (6-7 m/s) towards each other along the equator as there is no planetary rotation, until they collide after approximately 22 days. The collision creates a modon moving northward and another moving southward. In Lin et al. (2017) the simulation runs for 100 days, during which the two modons cross the poles and then collide at the equator on the opposite side, exchanging vorticity again. After 100 days they have passed through their initial positions and begin the loop again. Without diffusion, the modons should cycle indefinitely. However, in realistic models they will slowly decay. We only run the simulation to a maximum of twenty days due to instability issues that cause the modon structure to collapse shortly after this time. An additional issue affecting AMR runs is triggered when an AMR level intersects a polar panel edge, and the sources of these issues are under investigation.

The modons are initialized in an isothermal temperature profile of 300K and a pressure field that is hydrostatically balanced with an initial surface pressure of 1,000 hPa. Uniform vertical height coordinate mapping is utilized with 16 vertical levels, though one set of convergence tests are run with 32 vertical levels. The model top is set to 10 km. There is no moisture in this setup and the Coriolis rotational forcing is not included. The velocity is zero everywhere except for the initial zonal wind

perturbations of the modons, which is uniform throughout the vertical. The initial zonal wind is

$$u(\theta,\lambda) = U_0 \left(\exp\left(-\left(\frac{r_1}{r_0}\right)^2 \right) - \exp\left(-\left(\frac{r_2}{r_0}\right)^2 \right) \right) \tag{4.16}$$

where $U_0 = 40$ m s^{-1}, $r_0 = 500$ km, (θ, λ) are the latitude and longitude, and r_1 and r_2 are the great-circle distance from each modon's center. The modons are initially centered at $(\theta_1, \lambda_1) = (0, \pi/2)$ and $(\theta_2, \lambda_2) = (0, 3\pi/2)$. The vortex structure of the modons can be seen in Fig. 4.2a.

4.3.2 AMR with the modon test

We compare the results from a series of uniform and AMR runs to assess if the AMR correctly tags, refines, and follows the non-linear modons. We implement one refinement criterion for all AMR runs: a relative vorticity threshold of $|\zeta| > 2 \times 10^{-5}$ s^{-1}. Figure 4.2 shows the 5 km height vorticity profiles of the uniform c128 (left column) and c32/c128 AMR runs (right column) at days 0, 10, and 20. At this resolution, the modons lose more than a quarter of their strength over the first twenty days; at higher resolutions there is only minimal loss of strength. The vorticity tagged refinement in the c32/c128 run successfully refines the modon areas immediately and is able to track the propagating modons. The size and intensity of the modons in the AMR run are preserved in comparison to the c128 run. We can see this more concretely in Fig. 4.3 which plots the 10-day time evolution of the maximum vorticity of the modons for uniform runs from c64 to c512 and three AMR runs all using the same $|\zeta| > 2 \times 10^{-5}$ s^{-1} refinement criteria. The uniform runs show that for coarse resolutions, c128 and below, the modons rapidly decay. For c256 and higher the modons maintain their intensity. For the three AMR runs: c32/c128, c64/c256, and c128/c512, we see strong alignment in maximum vorticity to that of the uniform run

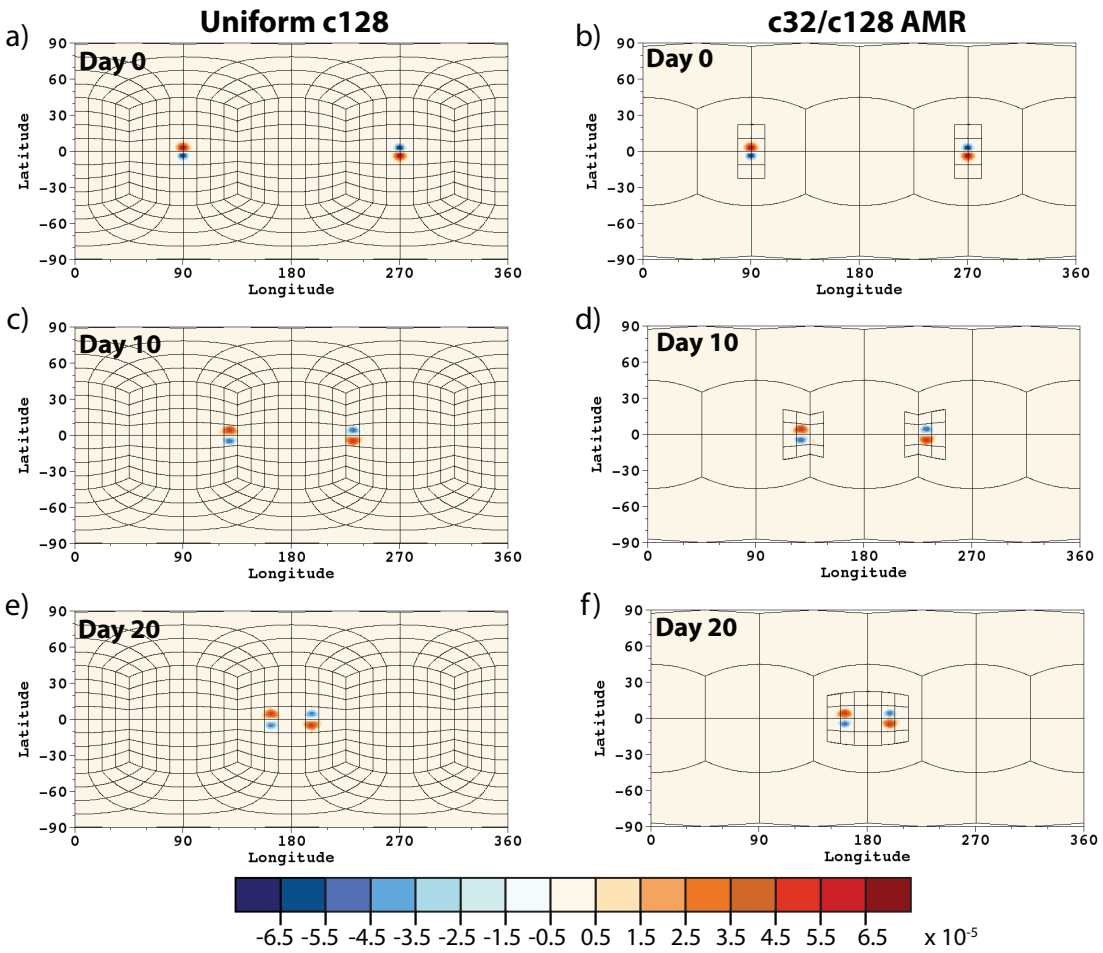

Figure 4.2: Snapshots of the vorticity field at day 0 (a)-(b), day 10 (c)-(d), and day 20 (e)-(f) for a uniform c128 run (left column) and a c32/c128 AMR run (right column).

with the same resolution as the AMR level. This result is not surprising as the AMR was in place over the modons for the entire run, but it does demonstrate that on the panels and across the equatorial panel edges AMR is functioning as expected.

4.3.3 Convergence tests

We also use the modons test for some basic convergence tests to assess the numerical methods. We compare a series of uniform runs of varying resolutions and AMRs to a uniform high-resolution reference solution. Since the modon test does not have a known analytical solution, the c1024 uniform run in each scenario serves as a

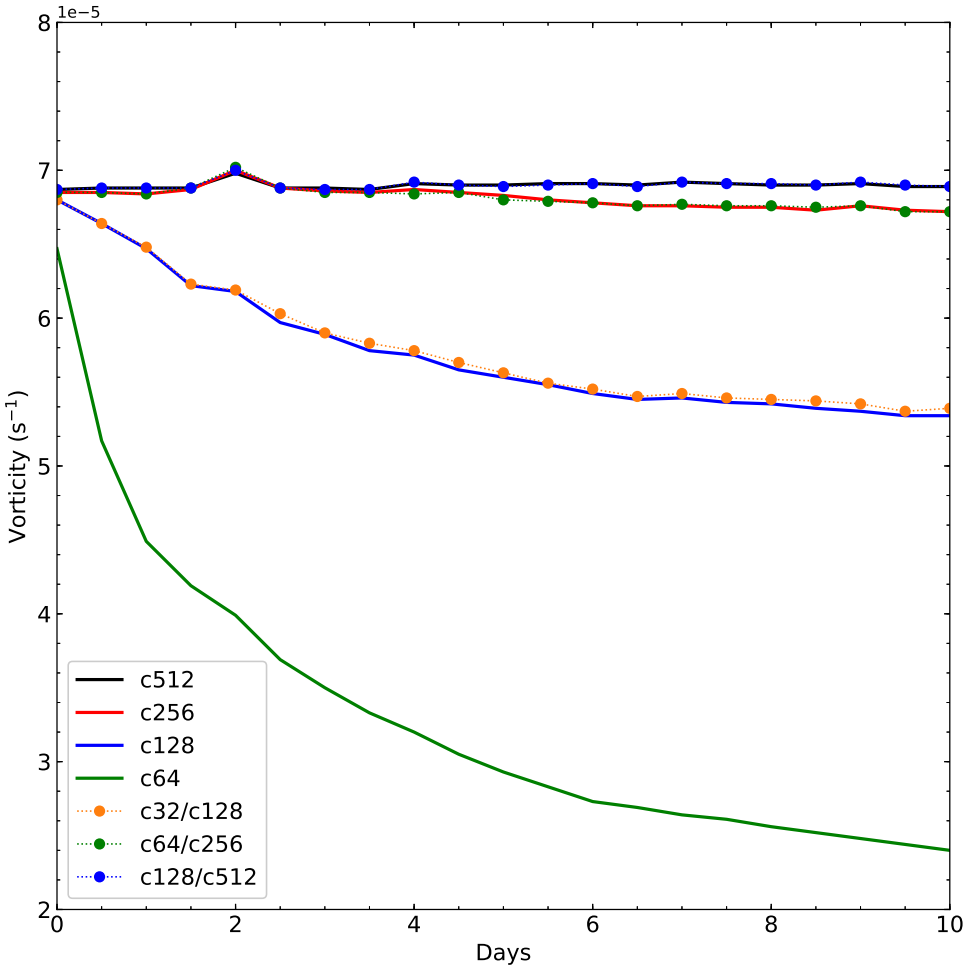

Figure 4.3: 10-day time evolution of the modons' maximum vorticity for uniform and AMR runs.

reference solution. The normalized error measures for each run are calculated via the traditional global error norms,

$$L_1(h) = \frac{I\left[|h - h_\tau|\right]}{I\left[|h_\tau|\right]} \tag{4.17}$$

$$L_2(h) = \sqrt{\frac{I\left[(h - h_\tau)^2\right]}{I\left[h_\tau^2\right]}} \tag{4.18}$$

$$L_\infty(h) = \frac{\max|h - h_\tau|}{\max|h_\tau|} \tag{4.19}$$

where h_τ is the reference field and I is a discrete approximation to the global 3D integral given by

$$I[x] = \sum_{\text{all cells } k} x_k V_k \tag{4.20}$$

where V_k denotes the volume of grid cell k. With the finite-volume model, we expect fourth-order convergence in the horizontal and second-order convergence in the vertical. However, the time scheme is only second-order so the overall error convergence for any variable tends towards second-order.

The 12-hour normalized errors as a function of grid resolution for density (left column) and vorticity (right column) are plotted in Fig. 4.4. The runs depicted are uniform runs with resolution from c64 to c512 and three 1-level x4 refinement AMR runs with c64, c128, and c256 base levels using the vorticity refinement criteria noted earlier. These runs are implemented with the model's standard time step which fixes the CFL number at roughly 0.5. With this, the uniform c512 model has a time step of $\Delta t = 20$ s, while the c64/c256 AMR run has a coarse fine step of $\Delta t = 160$ s and a sub-cycling time step of $\Delta t = 40$ s for the c256 AMR level. In the left column of Fig. 4.4, we observe second-order convergence in the L_1 (Fig. 4.4a) and L_2 (Fig. 4.4c) errors. It is further reduced to first order for the L_{max} in Fig. 4.4e. Additionally, the AMR runs have roughly the same errors as their base resolution

uniform runs. No noticeable improvement is gained from the AMR; the vorticity tagging and refinement of only two localized areas is not expected to significantly improve global density errors.

We do observe improvements in the vorticity errors (Figs. 4.4b, d, and f) as a result of AMR. The c64/c256 AMR has higher errors than the uniform c256, but lower errors than the uniform c128 run in all three norms despite using roughly one-third the total number of grid cells in the c128 run. The error improvement is not as significant in the c256/c1024 AMR. This could be explained by the convergence of the solution at higher resolutions resulting in diminishing returns for the benefits of AMR in improving global error (Lin et al., 2017). Surprisingly, vorticity maintains a fourth-order convergence in the L_1 and L_2 error norm (Figs. 4.4b and d) and approximately third-order convergence in the L_{max} norm (Fig. 4.4f) even though it is not a prognostic variable, but derived from the momentum variables. By 24 hours, convergence does drop to second order for vorticity, and several days into simulation errors for all variables are first-order at best.

To assess if the horizontal discretization is achieving fourth-order accuracy, we can work to minimize the temporal error source, so that the error is dominated by the spatial discretization. To this end, we perform a second convergence study comparing c32 through c512 resolution uniforms with a c1024 reference run. The time step in these runs is set to have CFL number of approximately 1/160 in contrast to the 0.5 in the previous study. We also only perform the error analysis at $t = 160$ s so the temporal error does not have time to build up. Note these runs were run with 32 vertical levels (all previous runs had 16 levels) to create a smoother vertical transition of variables. The L_1, L_2, and L_{max} errors are plotted in Fig. 4.5 for density 4.5a, momentum in the α direction 4.5b, and vorticity 4.5c. The uniform grid resolutions c64 and higher display fourth-order convergence in all three error norms for both ρu_α momentum in Fig. 4.5b and vorticity in Fig. 4.5c. The convergence rates

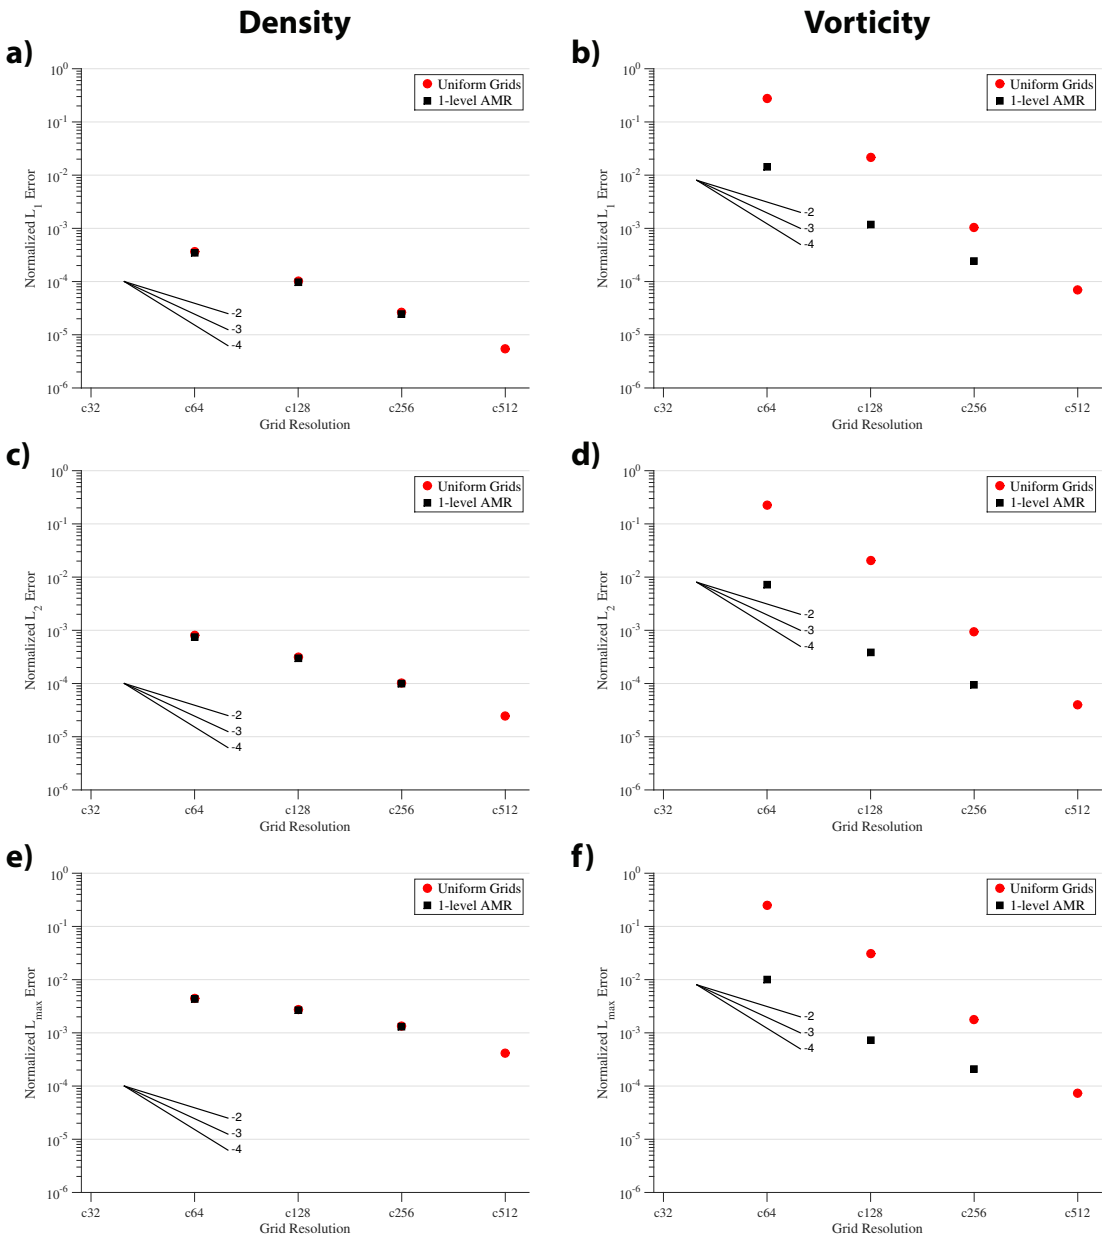

Figure 4.4: The 12 hour L_1 (a)-(b), L_2 (c)-(d), and L_{max} (e)-(f) normalized errors for ρ density (left column) and vorticity (right column) with respect to the uniform c1024 run for uniform and 1-level x4 refinement AMR runs. The AMR runs are plotted with respect to their base resolutions. These runs used the standard time steps which results in a CFL number of around 0.5 for all resolutions.

of the three error norms for the density variable diverge (Fig. 4.5a). The second-order convergence of the L_1 error reduces to first-order at higher resolutions. The L_2 error is approximately third-order convergent at resolutions c128 and below, but also reduces to first-order at the highest resolutions. In contrast, the L_{max} density error exhibits fourth-order convergence except at the coarsest c32 and c64 resolutions. The variations in density convergence could be attributed to the test case is not fully balanced. The models has to rapidly adjust to cyclostrophic balance. The error distortions could rise from the adjustment processes affecting the density field, though this question warrants further investigation to confirm it is not an artifact of an issue with the model numerics.

4.4 Idealized tropical cyclone

In the idealized TC test case by Reed and Jablonowski (2012), a simple physics parameterization suit with the important driving mechanism of TCs is added to the dycore. This causes an initially weak vortex to rapidly intensify into a TC and is a dynamic test case with real world applications. We use the idealized TC to demonstrate AMR's effectiveness in tracking and resolving the TC and asses how the physics forcing interacts with the AMR grid levels. Its use is analogous to the 2D moist shallow water strengthening vortex test in Chapter III. A key focus area is the ability of the physics scheme to handle changes in resolution, either at coarse-fine boundaries or when AMR is added or removed. In addition, we are interested in assessing refinement criteria and AMR implementation timing effects on the evolution of the TC.

4.4.1 Simple physics parameterization

The simple-physics parameterization suite of Reed and Jablonowski (2012) consists of three main mechanisms that drive tropical cyclone development. It implements

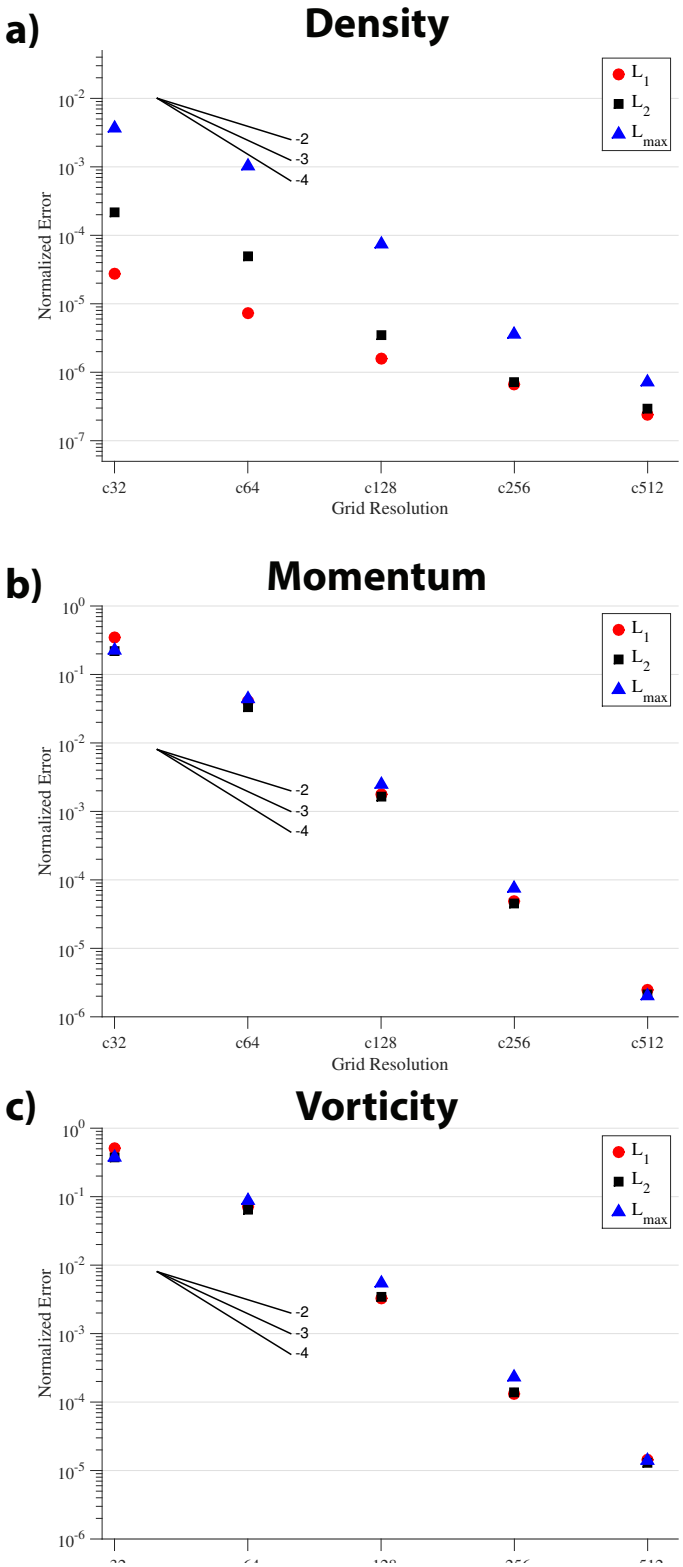

Figure 4.5: L_1, L_2, and L_{max} normalized errors after 160s for (a) ρ density, (b) ρu_α momentum density, and (c) vorticity with respect to the uniform c1024 run. Small time steps are used which result in a CFL number of around 1/160 for all resolutions.

a large-scale condensation mechanism that gets triggered when the atmosphere becomes saturated. The condensation scheme does not include a cloud stage, and it instantaneously removes condensed moisture without any reevaporation at lower model levels. The second component is a representation of the surface fluxes that describe the interactions between the lower atmosphere and the ocean surface. Specifically, it sets fluxes of horizontal momentum, sensible heat, and latent heat. The package assumes an aqua-planet surface with no topography and uniform sea surface temperatures. Finally, this scheme parameterizes boundary layer turbulence. For the Chombo AMR model, the physics time step is set to 16 minutes for all resolutions, comparable to the average physics time step used by models in Reed and Jablonowski (2012). This time step is longer than the dynamic time steps (40s to 6 mins, depending on resolution) used in this chapter. The frequency at which physics is called affects the evolution of the TC and is an area of further interest for AMR and the broad atmospheric modeling community.

4.4.2 Idealized tropical cyclone

The initialization of the starting analytic vortex is described in detail by Reed and Jablonowski (2011). The prescribed pressure, temperature, moisture, and velocity fields establish an initial vortex with initial maximum winds of 20 m s^{-1} located near the surface. The radius of maximum wind of roughly 250 km. The vortex is embedded in a background designed to mimic favorable tropical conditions. The background surface pressure is set to $p_0 = 1015$ hPa and the background wind at all vertical levels is zero. The model top height is set to 30 km and a stretched mapping is implemented in the height-based vertical coordinates. The mapping smoothly concentrates vertical levels closer to the surface to improve the resolution of the tropical cyclone and the physics processes that force it.

Figure 4.6 depicts the evolution of the TC in a c256 (0.35°) uniform resolution

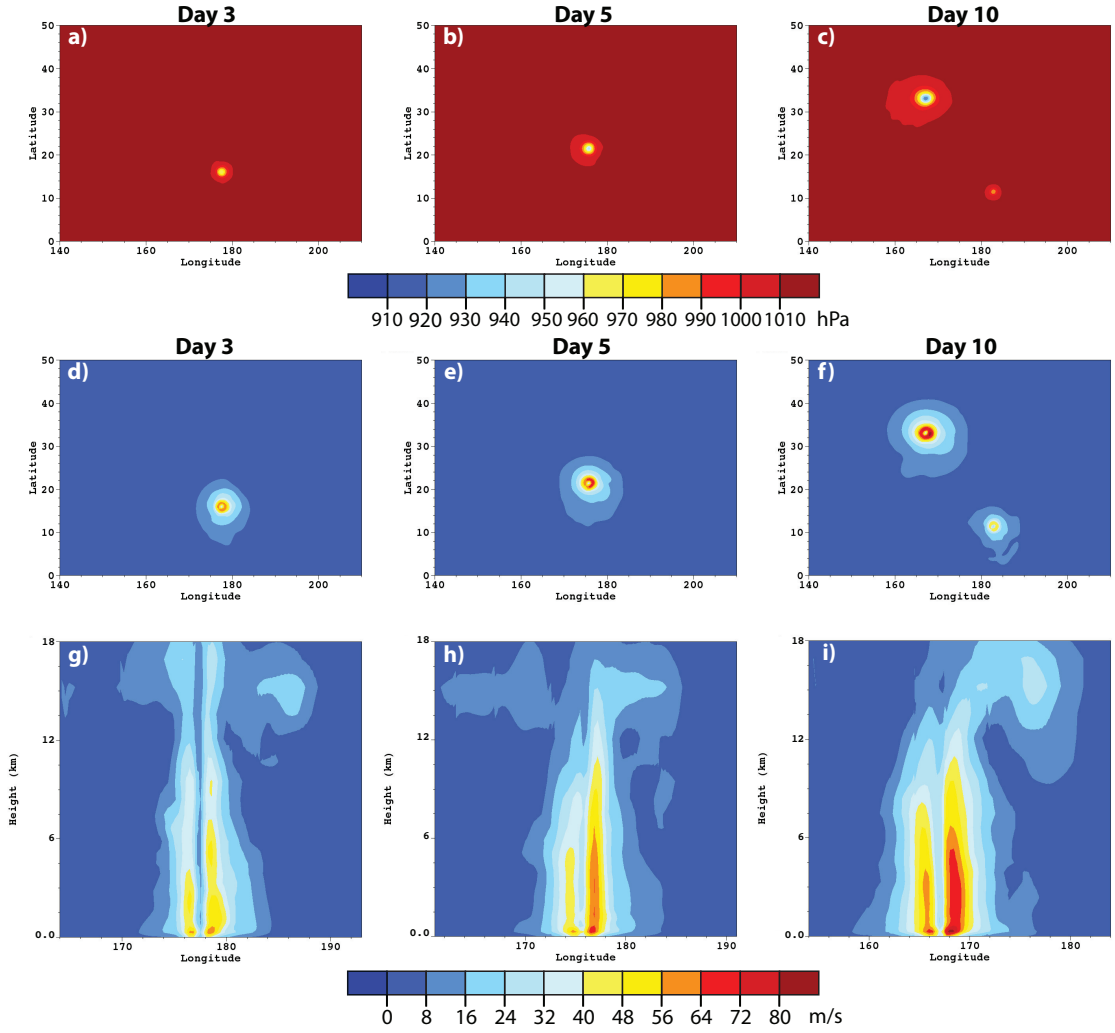

Figure 4.6: Snapshots of the tropical cyclone at day 3 (left), day 5 (middle), and day 10 (right) for the uniform c256 model run. (a)-(c) Surface pressure. (d)-(f) Wind speed at a height of 250 m. (g)-(i) Longitude-height cross-section of the wind speed through the center latitude of the vortex.

run as it beta-drifts northwestward due to the Coriolis force over a period of ten days. The leftmost column shows the surface pressure (Fig. 4.6a), the wind speed at 250m (Fig. 4.6d), and the vertical cross-section of the wind speed (Fig. 4.6g). The middle column (Figs. 4.6b, e, and h) are the same plots for day 5, and the rightmost column (Figs. 4.6c, f, and i) depicts day 10.

The c256 resolution is the highest resolution used for these simulations. As resolution increases, certain wave modes in the vertical components grow stronger and become unstable between the first and second day of the simulation. This time corresponds to the start of significant TC strengthening which drives a rapid increase in vertical velocity at the TC's center. The unstable modes add to this growth, eventually triggering CFL violations that stop the simulation. This instability can be mitigated by further decreasing the time step at higher resolutions. However, such methods are not sustainable, as resolutions c512 or beyond require time steps that are too small, making the simulation prohibitively costly to run. Determining the causes of these unstable modes and possible solutions is an ongoing effort. For this reason, the highest resolution implemented for the TC test in uniform or AMR mode is c256.

4.4.3 AMR with the idealized TC

In addition, to the uniform c256 run, we also ran uniform c64 and c128 runs. We implemented four different refinement criteria for use in the AMR runs.

- Tag 1 is a surface pressure threshold tag based on the absolute difference from the initial surface pressure $p_0 = 1015$ hPa. Refinement is triggered where $|p_s - p_0| > 4$ hPa.
- Tag 2 is a relative vorticity threshold that refines where $|\zeta| > 2 \times 10^{-5}\text{s}^{-1}$.
- Tag 3 is a relative vorticity threshold that refines where $|\zeta| > 1 \times 10^{-5}\text{s}^{-1}$.

- Tag 4 is a surface pressure tag that refines where $|p_s - p_0| > 9$ hPa.

An overview of all simulations performed for this section, including the run resolutions and AMR refinement criteria used, is presented in Table 4.1.

Table 4.1: Brief description of all uniform and AMR runs performed for Sec. 4.4. For each run the model resolutions, the refinement criteria (Tags 1 through 4), and the tagging variable are presented. Some of the AMR simulations have prescribed delays that prevent refinement from occurring until after a set time. That delay, in hours, is shown in the right most column. The model resolutions (left most column) are presented in cubed-sphere coordinates (cN) with N being the number of cells along each panel edge. For the AMR runs, the resolutions for the multiple levels are given in the form c32/c128/c512, where the left most resolution is the base level's resolution and the subsequent resolutions are for each level of AMR implemented.

Run Resolution	AMR Criteria	Tagging Variable	Delay (hrs)
c16/c64/c256	Tag 3	Vorticity	0
c32/c128	Tag 1	Surface Pressure	0
c32/c128	Tag 2	Vorticity	0
c64	-	-	-
c64/c256	Tag 1	Surface Pressure	0
c64/c256	Tag 2	Vorticity	0
c64/c256	Tag 3	Vorticity	0
c64/c256	Tag 4	Surface Pressure	0
c64/c256	Tag 1	Surface Pressure	48
c64/c256	Tag 3	Vorticity	24
c64/c256	Tag 3	Vorticity	48
c128	-	-	-
c256	-	-	-

Figure 4.7 depicts the time evolution of the minimum surface pressure (Fig. 4.7a) and maximum wind speed at a height of 250 m (Fig. 4.7b) for the three uniform runs and six AMR runs using the first three tagging criteria. We observe that the c256 run has a minimum surface pressure of around 920 hPa by day 10. This is comparable to the strongest TC produced in Reed and Jablonowski (2012) using the simple physics scheme. Wind speeds in Reed and Jablonowski (2012) were calculated at 100 m, making direct comparisons difficult. In the uniform c64 run, the TC slowly weakens initially. After day 3 it starts to strengthen but at a much slower rate than

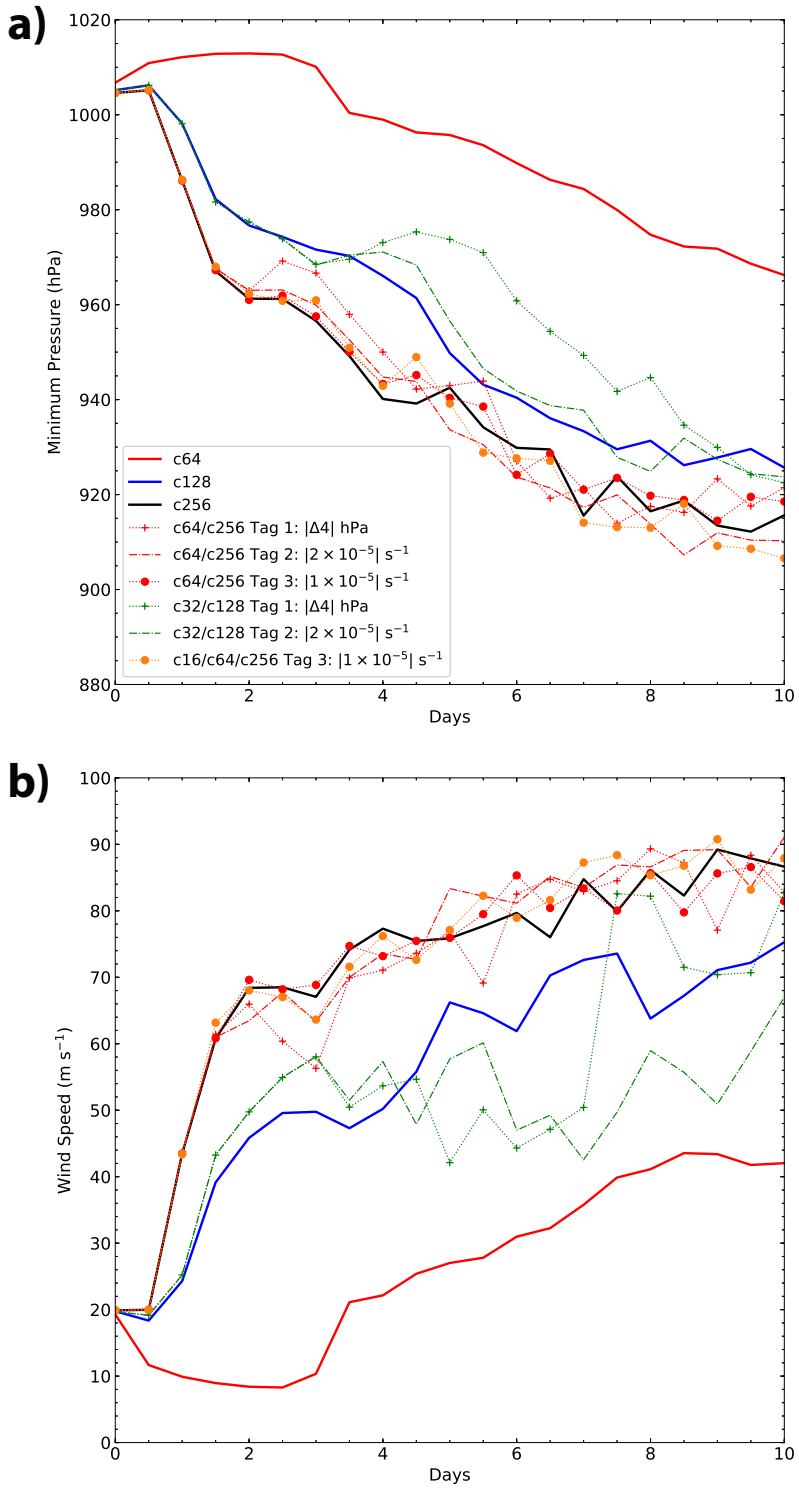

Figure 4.7: Time evolution of the (a) minimum surface pressure and (b) maximum wind speed at 250m for three uniform runs and six AMR runs using various refinement criteria.

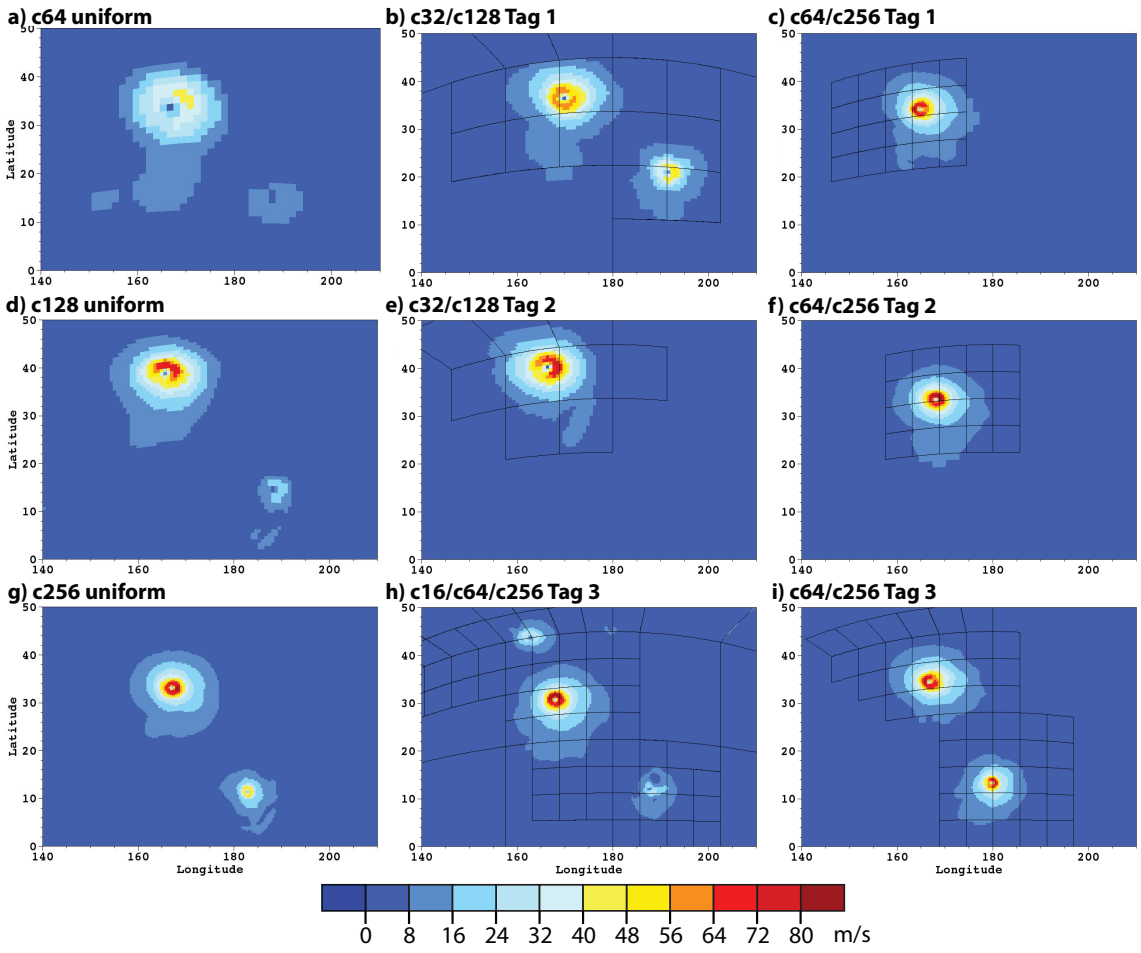

Figure 4.8: Day 10 snapshots of the horizontal wind speed at 250m for three uniform runs, (a) c64, (d) c128, and (g) c256, and the six AMR runs also depicted in Fig. 4.7.

the higher resolution runs. By day 10, its maximum wind speed is less than half the c256 uniform run. In 4.7, the four AMR runs that have a c256 level of refinement track the evolution of the uniform c256 TC closely. The two c32/c128 AMR runs also follow the uniform c128 run's progression but not as closely. Both the Tag 1 and Tag 2 c32 base AMR runs diverge markedly at several points during the ten days. In all six AMR runs, the refinement threshold is met initially so that the TC is resolved by the highest level of AMR for the entire simulation.

The wind fields at a height of 250 m for all nine runs are shown in Fig. 4.8. The vertical cross-section of the wind fields centered around the strongest vortex are

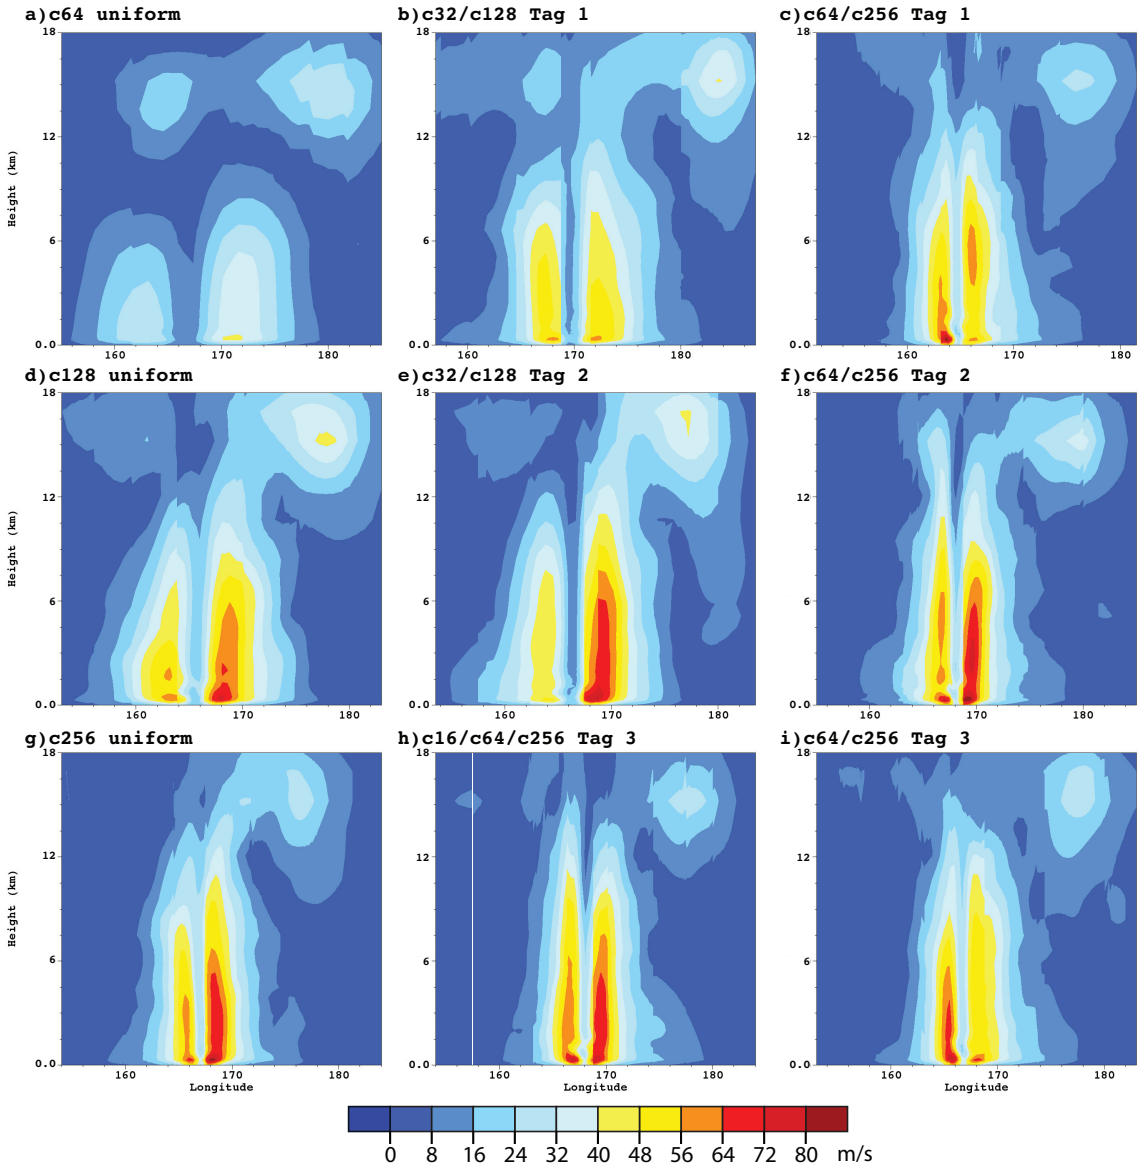

Figure 4.9: Day 10 snapshots of the longitude-height cross-section of the wind speed through the center latitude of the vortex, for the same uniform and AMR runs as in Fig. 4.8.

plotted in Fig. 4.9. From the uniform runs we observe increasing strength with increasing resolution as well increasing compactness, a recognized sensitivity of simple physics (Reed and Jablonowski, 2012). The TC strength and structure in the AMR runs compare well to the corresponding uniform runs. The main exception to this similarity is the secondary vortex which spins up at the main TC's origin point. This TC is another resolution-sensitive feature, that is only produced in some models in Reed and Jablonowski (2012). In our simulation all three uniform runs have the secondary vortex developing, though it is only clearly defined in the c256 run. However, only half of the AMR runs develop this secondary vortex. It develops in both Tag 3 AMR runs which have the broadest refinement, though it is fairly weak in Fig. 4.8h and much stronger than the c256 uniform run in the c64/c256 AMR run (Fig. 4.8i). Neither Tag 2 AMR runs develop the secondary vortex, and only the lower resolution c32/c128 surface pressure based AMR run (Fig. 4.8b) develops it. In Fig. 4.8b. the main TC is also significantly weaker than either the c128 uniform run (Fig. 4.8d) or the other c32/c128 AMR run (Fig. 4.8e) This main vortex is weakened by the secondary vortex interfering with the main vortex's source of heat and moisture. This also explains the weakened eastern side of the c64/c256 Tag 3's main TC, seen in Fig. 4.9i. Reed and Jablonowski (2012) observed that small perturbations in the initial vortex structure led to noticeable spread in the evolution of the TC. So it is not unexpected that differences in AMR grid locations can affect the evolution of the secondary vortex. The tertiary vortex that appears in the c16/c64/c256 AMR run (Fig. 4.8h) on the polar panel edge north of the main vortex is a numerical artifact from the same AMR/panel edge issue observed in the modon tests.

Initiating refinement at the start of the simulation minimizes the growth of early coarse grid errors that carry over to the finer grids. In more realistic scenarios, a distinct initial condition to tag on may not always be present. We are therefore interested in studying how the idealized TC and the simple physics mechanisms will

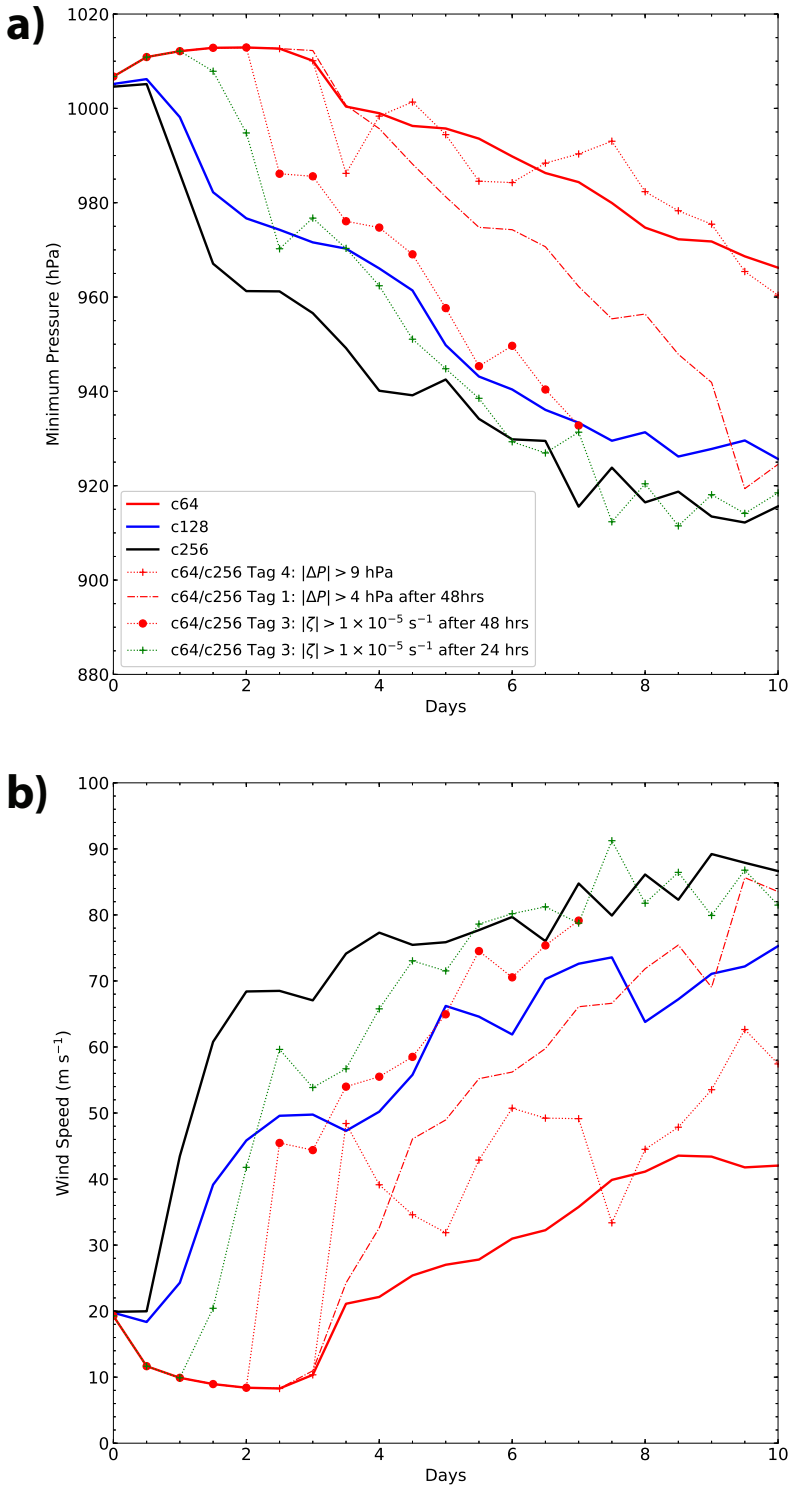

Figure 4.10: Time evolution of (a) minimum surface pressure and (b) maximum wind speed at 250m for three uniform runs and four AMR runs that do not have initial refinement over the vortex.

adjust to the addition of resolution when AMR is triggered a few days into the simulation rather than initially. The tag 4 criterion has a refinement threshold that is higher than the initial maximum surface pressure spread, so refinement is delayed until the TC undergoes some strengthening on the coarse grid. Due to the high threshold though, the newly triggered AMR grid fails to sufficiently cover the whole TC for several additional days. To explore more refinement options, we implement an artificial delay to refinement. For the Tag 1 and Tag 3 criteria, we manually shut off tagging in the model for either 24 or 48 hours. After that cut off, a broad area of refinement is applied over the TC. Figure 4.10 shows the time evolution of the minimum surface pressure (Fig. 4.7a) and maximum wind speed at a height of 250 m (Fig. 4.7b) for the four AMR runs. The uniform runs are again plotted for comparison purposes. Figures 4.11 and 4.12 depict the horizontal and vertical cross-sections of the winds of the three AMR runs, respectively, with the uniform c256 run serving as a reference. The c64/c256 Tag 3 AMR run with a 48-hour delay was stopped at day 7 by the vertical velocity instability discussed at the beginning of this section, so it is not pictured in these plots.

The c64/c256 Tag 4 AMR run triggers refinement after three days. Figure 4.10 shows the drop in surface pressure and sharp jump in wind speed, however the AMR does not sufficiently cover the vortex resulting in the TC's strength fluctuating around the uniform c64 TC's level despite refinement. In contrast, by day 10 the TC in the c64/c256 Tag 1 AMR run with a 48-hour delay has strengthened to c256 TC comparable wind speeds and surface pressure levels. Its AMR is triggered after two and a half days but refines a broader area, capturing the whole TC. Once the refinement is triggered, the TC strengthens consistently until day 10, but the increase is not as rapid as the other delayed AMR runs. This is also true in comparison to the initial strengthening of the c128 and c256 uniform runs. The TCs in the two delayed Tag 3 AMR runs follow similar trajectories. After the delay, refinement is immediately

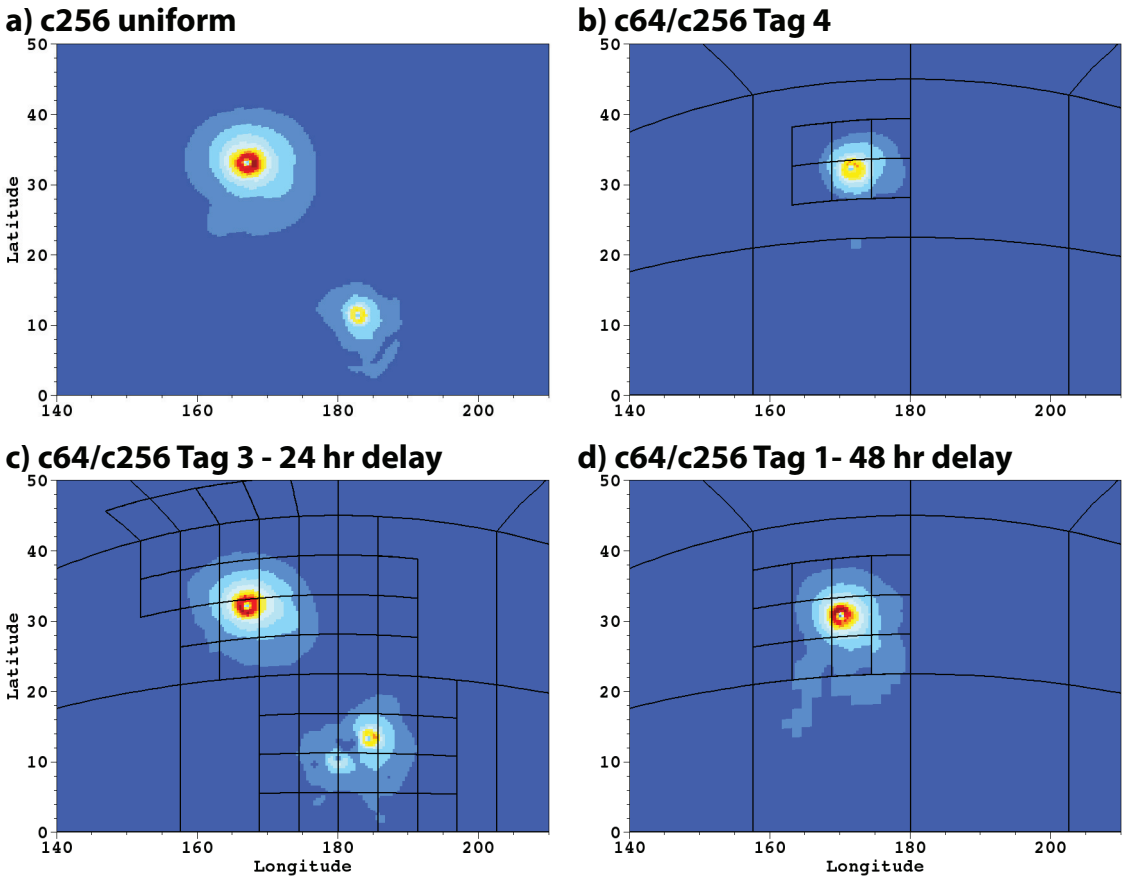

Figure 4.11: Day 10 snapshots of the horizontal wind speed at 250m for the three AMR runs that do not have initial refinement over the vortex. (a) Uniform c256, for reference. (b) c64/c256 using a tagging criterion of $|\Delta p| > 9$hPa. (c) c64/c256 using Tag 3 with a 24-hour delay. (d) c64/c256 using Tag 1 with a 48-hour delay.

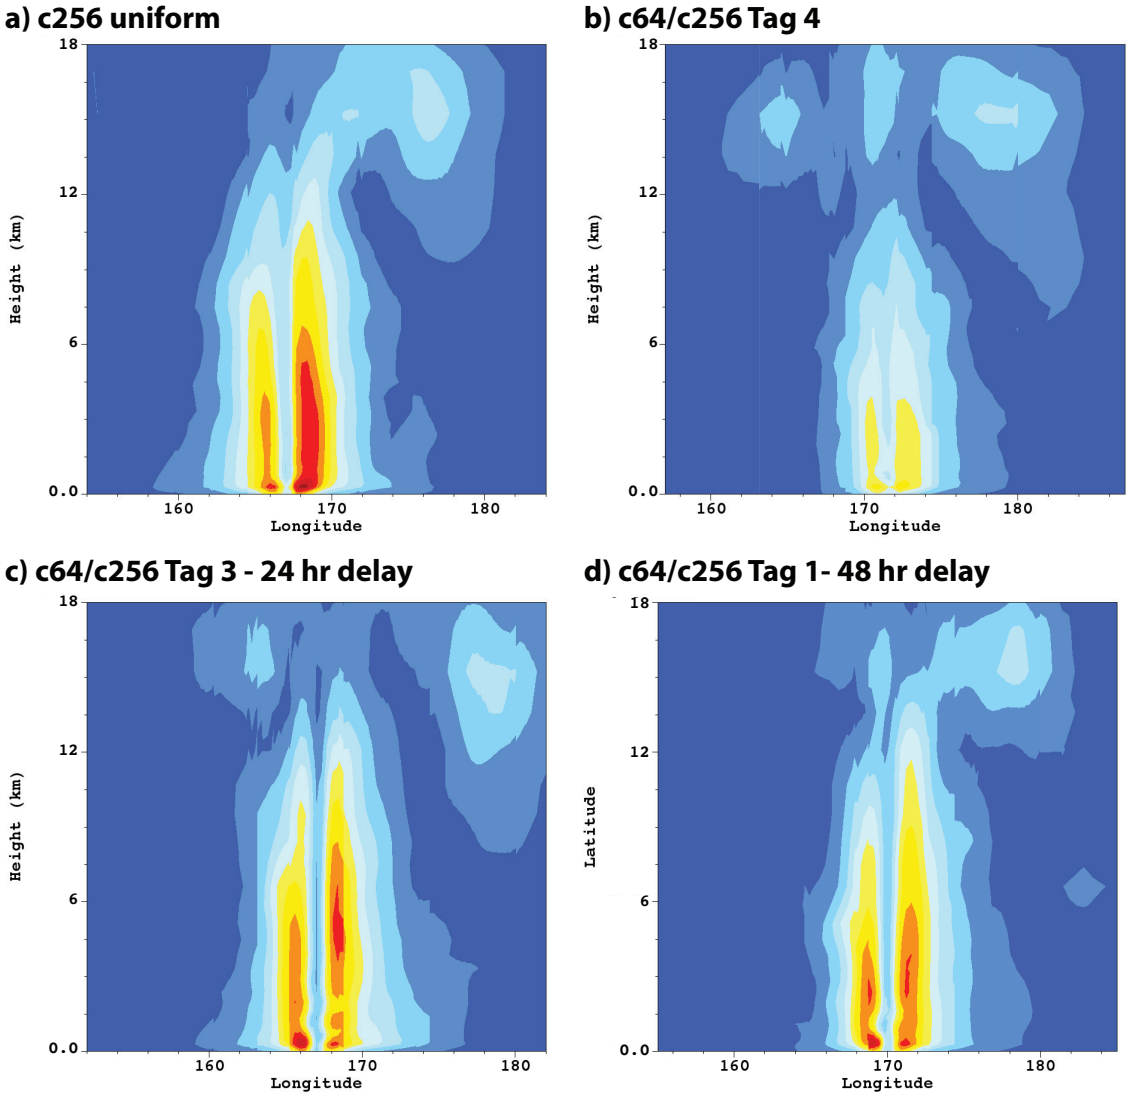

Figure 4.12: Day 10 snapshots of the jongitude-height cross-section of the wind speed through the center latitude of the vortex, for the same uniform and AMR runs as in Fig. 4.11.

triggered over the whole TC and the surrounding area. The TC undergoes an abrupt transition that temporarily distorts the vortex structure for approximately 24 hours and causes it to rapidly strengthen. The evolution of the c64/c256 Tag 3 AMR run with the 24-hours delay mimics the rapid strengthening of the uniform c256 TC, merely a day behind. By day 6, the minimum surface pressure and wind speed in the Tag 3 AMR and the c256 uniform runs are approximately the same. At day 10, Figs. 4.11 and 4.12 shows that both the Tag 1 and Tag 3 delayed runs have TCs comparable to the uniform c256 TC. The c64/256 Tag 3 delayed run even captures the development of the secondary vortex (Fig. 4.7c). The processes in the evolution of the high resolution TC can still be triggered in AMR runs that do not have high levels of refinement initially. These AMR runs demonstrate a time window during which the TC can undergo rapid strengthening and evolve into a vortex comparable to the TC from the uniform high resolution run.

4.5 Conclusions

In this chapter, we used the non-hydrostatic finite-volume Chombo AMR dycore and demonstrated its AMR characteristics for idealized 3D atmospheric flows on the sphere. We implemented two test cases: the colliding modons test, in the dry dycore and the idealized TC test with an added simple moist physics parameterization package. In the modon test, AMR functioned as expected. It was able to tag, refine, and follow the modons and it effectively reduced the global vorticity errors. The error convergence properties met our expectations given the implemented numerical schemes and the current status of the model. However, a few notable results trigger the need for further investigations. Several stability problems were observed in both test cases that need to be better understood and corrected.

In the idealized TC test, the AMR runs were able to effectively reproduce results from uniform high resolution runs. The physics scheme was able to function effectively

over multiple levels of refinement. No noise or instabilities were observed at coarse-fine boundaries, though some numerical artifacts were observed on polar panel edges. Several aspects of the TC's evolution were sensitive to the tagging criteria and AMR coverage such as the generation of the secondary vortex. Having AMR levels applied initially is the most effective way of improving results. However, the delayed AMR tests demonstrated that the TC can still undergo the same strengthening processes as high resolution runs once refinement is triggered. There is a narrow window of flexibility, in which the triggering of AMR allows the TC in the AMR run to catch up to the results in the high resolution uniform run. The most promising refinement technique is a combination of some initial refinement and AMR. The initial refinement limits error growth at the base resolution and ensures that the model can resolve the feature of interest. Additional AMR level then enable the feature to then be fully resolved.

CHAPTER V

Conclusion

Complex multi-scale atmospheric phenomena, like tropical cyclones, challenge conventional weather and climate models, which use relatively coarse uniform-grid resolutions to cope with computational costs. Adaptive Mesh Refinement (AMR) techniques mitigate these challenges by dynamically and transiently placing high-resolution grids over salient features, thus providing sufficient local resolution while limiting the computational burden. This thesis explores the development of AMR within a new non-hydrostatic, finite-volume dynamical core and demonstrates its effectiveness in improving model accuracy and its ability to resolve multi-scale features. The AMR dynamical core is implemented in a hierarchy of models of increasing complexity, from an idealized 2D shallow water configuration to the non-hydrostatic 3D equation set with subgrid-scale physics parameterization schemes. The research explores effective refinement practices and assesses the benefits achieved with increased dynamic refinement. It is shown that AMR is a powerful modeling approach that bridges the resolution gap for extreme weather events.

In Chapter II 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. Static and dynamic refinements were 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, height, or gradient-based thresholds, in order to achieve the best accuracy. The simulations showed 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.

In Chapter III we implemented two different forcing schemes designed to mimic the effects of moisture in the atmosphere within the 2D AMR shallow water system. The first moist physics framework added water vapor variable and models convection as a mass sink triggered by saturation. We implemented a strengthened vortex test case with this setup. In the second forcing framework, a more complex moisture representation is used, consisting of vapor, cloud, and rain variables. The effects of moisture were coupled to the momentum equations through a potential temperature variable, linked to the moisture variables through latent heat. This physics scheme was used with a barotropic instability test case. With both forcing systems and test cases, we observe the evolution of features of interest at various resolutions and with different refinement strategies. We also investigated the AMR's effect on the physics forcing as grid resolutions changed. These simulations can help develop AMR tagging strategies and refinement criteria. Both sets of simulations showed that the starting resolution must be able to adequately resolve the feature of interest to maximize the effectiveness of AMR. AMR cannot remove the errors caused by coarse grids before refinement begins. Additional refinement with AMR beyond the base grid level did

improve the model, especially with regards to the small-scale vorticity features in the strengthening vortex test case. To obtain such early refinement with AMR, the tagging criteria must be tailored to properties that are uniquely associated with the origins of the feature of interest. This is difficult even in these idealized shallow water systems.

In Chapter IV, we implemented AMR in the non-hydrostatic finite-volume Chombo AMR dycore and tested it for idealized 3D atmospheric flows on the sphere. We used two test cases: the dry dycore colliding modons test and an idealized TC test with a simple moist physics parameterization scheme. In the modon test, AMR functioned as expected. It was able to tag, refine, and follow the modons, and the added resolution effectively reduced global the vorticity errors. The error convergence properties met our expectations given the numerical schemes used and the current status of the model. However, a few results triggered the need for additional in-depth examinations. Several stability problems were observed in both test cases that need to be better understood and addressed.

In the idealized TC test, the AMR runs were able to effectively reproduce results from uniform high resolution runs. The physics scheme was able to function effectively over multiple levels of refinement. No noise or instabilities were observed at coarse-fine boundaries, though some numerical artifacts were observed on polar panel edges. Several aspects of the TC's evolution were sensitive to the tagging criteria and AMR coverage such as the generation of the secondary vortex. Having AMR levels applied initially is the most effective way of improving results. However, the delayed AMR tests demonstrated that the TC can still undergo the same strengthening processes as high resolution runs once refinement is triggered. There is a narrow window of flexibility, in which the triggering of AMR allows the TC in the AMR run to catch up to the results in the high resolution uniform run. The most promising refinement technique is a combination of some initial refinement and AMR. The initial refinement

limits error growth at the base resolution and ensures that the model can resolve the feature of interest. Additional AMR level then enable the feature to then be fully resolved. For example, in tracking and resolving TCs in a realistic climate simulation, a static region of refinement could be placed over regions of cyclogenesis. Any storms that develop could be further refined and followed with AMR as they traverse and exit the region of static refinement.

5.0.1 Collaborations

This research has built collaborations between the University of Michigan (UofM), the Applied Numerical Algorithms Group (ANAG) at the Lawrence Berkeley National Laboratory (LBNL), and the University of California, Davis (UC Davis). In particular, direct contributors to the research contained in this thesis include Christiane Jablonowski (UofM), Hans Johansen (LBNL), Peter McCorquodale (LBNL), Phillip Colella (LBNL), and Paul Ullrich (UC Davis).

5.0.2 Future Work

AMR dycore development: Additional effort is required in implementing, validating, and improving AMR within GCMs. Use of AMR in a 3D non-hydrostatic dycore is a relatively recent application. Further development includes the creation of more complex models with realistic physics schemes, with the goal of running aqua-planet simulations with full-physics schemes. Implementing vertical refinement with AMR is another algorithm development focus. With an AMR non-hydrostatic model, grid resolutions on the order of 1km or finer are feasible. At these resolutions, the aspect ratios between the horizontal and vertical lengths of a grid cell can be come highly skewed. The addition of more vertical cells would alleviate this issue, and increase vertical resolution.

Scale aware subgrid physics parameterizations: The effectiveness of VRGCMs

is limited by the sub-grid physical parameterizations used to approximate atmospheric features including radiation and sub-grid scale process like convection, clouds, and turbulence that cannot be resolved by the dycore. In adaptive schemes, the subgrid-scale parameterizations need to be able to adjust for changes in scale. A model will need to be able to phase out certain subgrid-scale processes, like deep convection, as resolution is increased and these processes are resolved on the actual mesh. One novel research avenue is the application of artificial intelligence to aid in the development of these scale-aware physics schemes.

AMR and topography: Choosing an effective method to handle topography is a challenging endeavor for any model. The question of smoothing always arises as very steep topological gradients can cause issues for the vertical coordinates and create numerical noise. AMR faces the additional hurdle of how to deal with topography when the grid is refined over it. Literature on this subject is rather sparse, but there are several possible methods of approaching this issue. One traditional method is to merely interpolate the existing coarse-grid topography to the new grid cells. This method, though preserving monotonicity, will not necessarily improve orographic representation as it cannot make mountains taller, valleys deeper, or gradients steeper. A second method would be to actually alter the topographic features making them more pronounced and even steeper or taller when the grid is refined. The topography can then be smoothed down for coarse grids and merely updated back to the high-resolution version when AMR refines it. Such a method provides a more realistic orographic representation but could be numerically unsound, causing air mass conservation issues and creating large gravity waves. A third method is to merely keep a static refinement mesh over large topographic features, though this could be computationally inefficient.

BIBLIOGRAPHY

BIBLIOGRAPHY

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 adaptive wavelet method for the shallow-water equations on the sphere. *Quarterly Journal of the Royal Meteorological Society*, **141 (690)**, 1712–1726.

Ascher, U. M., S. J. Ruuth, and R. J. Spiteri, 1997: Implicit-explicit Runge-Kutta methods for time-dependent partial differential equations. *Applied Numerical Mathematics*, **25 (2-3)**, 151–167.

Bacon, D. P., T. J. Dunn, S. Gopalakrishnan, M. S. Hall, A. Sarma, and Coauthors, 2007: Hurricane track forecasting with OMEGA. *Natural Hazards*, **41 (3)**, 457–470.

Bacon, D. P., and Coauthors, 2000: A dynamically adapting weather and dispersion model: the operational multiscale environment model with grid adaptivity (OMEGA). *Monthly Weather Review*, **128 (7)**, 2044–2076.

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. *Theoretical and Computational Fluid Dynamics*, **28 (1)**, 107–128.

Behrens, J., 1998: Atmospheric and ocean modeling with an adaptive finite element solver for the shallow-water equations. *Applied Numerical Mathematics*, **26 (1)**, 217–226.

Behrens, J., 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 Modelling*, **10 (1)**, 171–183.

Berger, M. J., and P. Colella, 1989: Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, **82 (1)**, 64–84.

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. *Monthly Weather Review*, **140 (3)**, 978–996.

Bouchut, F., J. Lambaerts, G. Lapeyre, and V. Zeitlin, 2009: Fronts and nonlinear waves in a simplified shallow-water model of the atmosphere with moisture and convection. *Physics of Fluids*, **21 (11)**, 116 604.

Budd, C. J., W. Huang, and R. D. Russell, 2009: Adaptivity with moving grids. *Acta Numerica*, **18**, 111–241.

Caya, D., and R. Laprise, 1999: A semi-implicit semi-Lagrangian regional climate model: The Canadian RCM. *Monthly Weather Review*, **127 (3)**, 341–362.

Chen, C., 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. *Journal of Computational Physics*, **271**, 191–223.

Chen, C., F. Xiao, and X. Li, 2011a: An adaptive multimoment global model on a cubed sphere. *Monthly Weather Review*, **139**, 523–548.

Chen, C., F. Xiao, X. Li, and Y. Yang, 2011b: A multi-moment transport model on cubed-sphere grid. *International Journal for Numerical Methods in Fluids*, **67 (12)**, 1993–2014.

Côté, J., S. Gravel, A. Méthot, A. Patoine, M. Roch, and A. Staniforth, 1998: The operational CMC–MRB global environmental multiscale (GEM) model. part i: Design considerations and formulation. *Monthly Weather Review*, **126 (6)**, 1373–1395.

Day, M. S., and J. B. Bell, 2000: Numerical simulation of laminar reacting flows with complex chemistry. *Combustion Theory and Modelling*, **4 (4)**, 535–556.

Dennis, J. M., and Coauthors, 2012: CAM-SE: A scalable spectral element dynamical core for the Community Atmosphere Model. *The International Journal of High Performance Computing Applications*, **26 (1)**, 74–89.

Déqué, M., and J. P. Piedelievre, 1995: High resolution climate simulation over europe. *Climate Dynamics*, **11 (6)**, 321–339.

Dias, J., M. Gehne, G. N. Kiladis, N. Sakaeda, P. Bechtold, and T. Haiden, 2018: Equatorial waves and the skill of NCEP and ECMWF numerical weather prediction systems. *Monthly Weather Review*, **146 (6)**, 1763–1784.

Dietachmayer, G. S., 1992: Application of continuous dynamic grid adaption techniques to meteorological modeling. Part II: Efficiency. *Monthly Weather Review*, **120 (8)**, 1707–1722.

Dietachmayer, G. S., and K. K. Droegemeier, 1992: Application of continuous dynamic grid adaption techniques to meteorological modeling. Part I: Basic formulation and accuracy. *Monthly Weather Review*, **120 (8)**, 1675–1706.

Dritschel, D., and D. Waugh, 1992: Quantification of the inelastic interaction of unequal vortices in two-dimensional vortex dynamics. *Physics of Fluids A: Fluid Dynamics (1989-1993)*, **4 (8)**, 1737–1744.

Düben, P. D., P. Korn, and V. Aizinger, 2012: A discontinuous/continuous low order finite element shallow water model on the sphere. *Journal of Computational Physics*, **231 (6)**, 2396–2413.

Dudhia, J., and J. F. Bresch, 2002: A global version of the PSU–NCAR mesoscale model. *Monthly Weather Review*, **130 (12)**, 2989–3007.

Enagonio, J., and M. T. Montgomery, 2001: Tropical cyclogenesis via convectively forced vortex Rossby waves in a shallow water primitive equation model. *Journal of the Atmospheric Sciences*, **58 (7)**, 685–706.

Eskilsson, C., 2011: An hp-adaptive discontinuous Galerkin method for shallow water flows. *International Journal for Numerical Methods in Fluids*, **67 (11)**, 1605–1623.

Ferguson, J. O., C. Jablonowski, H. Johansen, P. McCorquodale, P. Colella, and P. A. Ullrich, 2016: Analyzing the adaptive mesh refinement (AMR) characteristics of a high-order 2d cubed-sphere shallow-water model. *Monthly Weather Review*, **144 (12)**, 4641–4666.

Ferreira, R. N., W. H. Schubert, and J. J. Hack, 1996: Dynamical aspects of twin tropical cyclones associated with the Madden–Julian oscillation. *Journal of the Atmospheric Sciences*, **53 (7)**, 929–945.

Flato, G., and Coauthors, 2013: Evaluation of climate models. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, T. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. Midgley, Eds., Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, book section 9, 741–866.

Fournier, A., M. A. Taylor, and J. J. Tribbia, 2004: The spectral element atmosphere model (SEAM): High-resolution parallel computation and localized resolution of regional dynamics. *Monthly Weather Review*, **132 (3)**, 726–748.

Fox-Rabinovitz, M., J. Côté, B. Dugas, M. Déqué, and J. L. McGregor, 2006: Variable resolution general circulation models: Stretched-grid model intercomparison project (SGMIP). *Journal of Geophysical Research: Atmospheres*, **111 (D16)**.

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

Fox-Rabinovitz, M. S., L. L. Takacs, and R. C. Govindaraju, 2002: A variable-resolution stretched-grid general circulation model and data assimilation system with multiple areas of interest: Studying the anomalous regional climate events of 1998. *Journal of Geophysical Research: Atmospheres*, **107 (D24)**, ACL 12–1–ACL 12–25.

Frierson, D. M., A. J. Majda, O. M. Pauluis, and Coauthors, 2004: Large scale dynamics of precipitation fronts in the tropical atmosphere: A novel relaxation limit. *Communications in Mathematical Sciences*, **2 (4)**, 591–626.

Fujiwhara, S., 1921: The natural tendency towards symmetry of motion and its application as a principle in meteorology. *Quarterly Journal of the Royal Meteorological Society*, **47 (200)**, 287–292.

Galewsky, J., R. K. Scott, and L. M. Polvani, 2004: An initial-value problem for testing numerical models of the global shallow-water equations. *Tellus A*, **56 (5)**, 429–440.

Gettelman, A., P. Callaghan, V. Larson, C. Zarzycki, J. Bacmeister, P. Lauritzen, P. Bogenschutz, and R. Neale, 2017: Regional climate simulations with the Community Earth System Model. *Journal of Advances in Modeling Earth Systems*, **10 (6)**, 1357–1380.

Gill, A. E., 1982: Studies of moisture effects in simple atmospheric models: The stable case. *Geophysical & Astrophysical Fluid Dynamics*, **19 (1-2)**, 119–152.

Giraldo, F., 2000: The Lagrange-Galerkin method for the two-dimensional shallow water equations on. *Int. J. Numer. Meth. Fluids*, **33**, 789–832.

Gopalakrishnan, S., and Coauthors, 2002: An operational multiscale hurricane forecasting system. *Monthly Weather Weview*, **130 (7)**, 1830–1847.

Goswami, P., and B. Goswami, 1991: Modification of n = 0 equatorial waves due to interaction between convection and dynamics. *Journal of the Atmospheric Sciences*, **48 (20)**, 2231–2244.

Goto, D., and Coauthors, 2015: Application of a global nonhydrostatic model with a stretched-grid system to regional aerosol simulations around Japan. *Geoscientific Model Development*, **8 (2)**, 235–259.

Grell, G. A., J. Dudhia, D. R. Stauffer, and Coauthors, 1994: A description of the fifth-generation Penn State/NCAR mesoscale model (MM5). Tech. Rep. NCAR/TN-398+STR, NCAR, 117 pp.

Hagos, S., L. R. Leung, Q. Yang, C. Zhao, and J. Lu, 2015: Resolution and dynamical core dependence of atmospheric river frequency in global model simulations. *Journal of Climate*, **28 (7)**, 2764–2776.

Haiden, T., M. Janousek, J. Bidlot, L. Ferranti, F. Prates, F. Vitart, P. Bauer, and D. Richardson, 2016: Evaluation of ECMWF forecasts, including the 2016 resolution upgrade. Technical Memo 792, European Centre for Medium Range Weather Forecasts, 55 pp.

Harris, L. M., and S.-J. Lin, 2013: A two-way nested global-regional dynamical core on the cubed-sphere grid. *Monthly Weather Review*, **141 (1)**, 283–306.

Harris, L. M., and S.-J. Lin, 2014: Global-to-regional nested grid climate simulations in the GFDL high resolution atmospheric model. *Journal of Climate*, **27 (13)**, 4890–4910.

Hayhoe, K., J. Edmonds, R. Kopp, A. LeGrande, B. Sanderson, M. Wehner, and D. Wuebbles, 2017: Climate models, scenarios, and projections. in: *Climate Science Special Report: Fourth National Climate Assessment*, volume i. DC: US Global Change Research Program, 133–160 pp.

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. *Monthly Weather Review*, **144 (10)**, 3697–3724.

Hendricks, E. A., W. H. Schubert, Y.-H. Chen, H.-C. Kuo, and M. S. Peng, 2014: Hurricane eyewall evolution in a forced shallow-water model. *Journal of the Atmospheric Sciences*, **71 (5)**, 1623–1643.

Holland, G. J., and G. S. Dietachmayer, 1993: On the interaction of tropical-cyclone-scale vortices. III: Continuous barotropic vortices. *Quarterly Journal of the Royal Meteorological Society*, **119 (514)**, 1381–1398.

Huang, C.-Y., Y. Zhang, W. C. Skamarock, and L.-H. Hsu, 2017: Influences of large-scale flow variations on the track evolution of typhoons Morakot (2009) and Megi (2010): Simulations with a global variable-resolution model. *Monthly Weather Review*, **145 (5)**, 1691–1716.

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. *Journal of Advances in Modeling Earth Systems*, **8**, 345–369.

Hubbard, M., and N. Nikiforakis, 2003: A three-dimensional, adaptive, Godunov-type model for global atmospheric flows. *Monthly Weather Review*, **131 (8)**, 1848–1864.

Iselin, J. P., J. M. Prusa, and W. J. Gutowski, 2002: Dynamic grid adaptation using the MPDATA scheme. *Monthly Weather Review*, **130 (4)**, 1026–1039.

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. *Monthly Weather Review*, **134**, 3691–3713.

Jablonowski, C., R. C. Oehmke, and Q. F. Stout, 2009a: Block-structured adaptive meshes and reduced grids for atmospheric general circulation models. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **367 (1907)**, 4497–4522.

Jablonowski, C., R. C. Oehmke, and Q. F. Stout, 2009b: Block-structured adaptive meshes and reduced grids for atmospheric general circulation models. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **367**, 4497–4522.

Jones, R. W., 1977: A nested grid for a three-dimensional model of a tropical cyclone. *Journal of the Atmospheric Sciences*, **34 (10)**, 1528–1553.

Kent, T., O. Bokhove, and S. Tobias, 2017: Dynamics of an idealized fluid model for investigating convective-scale data assimilation. *Tellus A: Dynamic Meteorology and Oceanography*, **69 (1)**, 1369 332.

Kessler, E., 1969: On the distribution and continuity of water substance in atmospheric circulation. *Meterological Monographs*, **10 (32)**, 1–84.

Kinter, J. L., III, and Coauthors, 2013: Revolutionizing climate modeling with project Athena: A multi-institutional, international collaboration. *Bulletin of the American Meteorological Society*, **94**, 231–245.

Koch, S. E., and J. T. McQueen, 1987: A survey of nested grid techniques and their potential for use within the MASS weather prediction model. Tech. Rep. 87808, NASA, 28 pp.

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. *Journal of Computational Physics*, **275**, 92–117.

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. *Computer Methods in Applied Mechanics and Engineering*, **198 (21)**, 1766–1774.

Kühnlein, C., P. K. Smolarkiewicz, and A. Dörnbrack, 2012: Modelling atmospheric flows with adaptive moving meshes. *Journal of Computational Physics*, **231 (7)**, 2741–2763.

Kurihara, Y., G. J. Tripoli, and M. A. Bender, 1979: Design of a movable nested-mesh primitive equation model. *Monthly Weather Review*, **107 (3)**, 239–249.

Lahaye, N., and V. Zeitlin, 2016: Understanding instabilities of tropical cyclones and their evolution with a moist convective rotating shallow-water model. *Journal of the Atmospheric Sciences*, **73 (2)**, 505–523.

Lambaerts, J., G. Lapeyre, and V. Zeitlin, 2011: Moist versus dry barotropic instability in a shallow-water model of the atmosphere with moist convection. *Journal of the Atmospheric Sciences*, **68 (6)**, 1234–1252.

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. *Journal of Computational Physics*, **229 (5)**, 1401–1424.

Läuter, M., D. Handorf, and K. Dethloff, 2005: Unsteady analytical solutions of the spherical shallow water equations. *Journal of Computiational Physics*, **210**, 535–553.

Läuter, M., D. Handorf, N. Rakowsky, J. Behrens, S. Frickenhaus, M. Best, K. Dethloff, and W. Hiller, 2007: A parallel adaptive barotropic model of the atmosphere. *Journal of Computational Physics*, **223 (2)**, 609–628.

Ley, G. W., and R. L. Elsberry, 1976: Forecasts of typhoon Irma using a nested-grid model. *Monthly Weather Review*, **104 (9)**, 1154–1161.

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. *Quarterly Journal of the Royal Meteorological Society*, **141 (691)**, 2090–2102.

Lin, S.-J., L. Harris, X. Chen, W. Yao, and J. Chai, 2017: Colliding modons: A nonlinear test for the evaluation of global dynamical cores. *Journal of Advances in Modeling Earth Systems*, **9 (7)**, 2483–2492.

Lorenz, P., and D. Jacob, 2005: Influence of regional scale information on the global circulation: A two-way nesting climate simulation. *Geophysical Research Letters*, **32 (18)**.

Manganello, J. V., and Coauthors, 2012: Tropical cyclone climatology in a 10-km global atmospheric GCM: Toward weather-resolving climate modeling. *Journal of Climate*, **25**, 3867–3893.

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. *Quarterly Journal of the Royal Meteorological Society*, **141 (690)**, 1727–1739.

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. *Communications in Applied Mathematics and Computational Science*, **10 (2)**, 121–162.

McGregor, J. L., 1996: Semi-Lagrangian advection on conformal-cubic grids. *Monthly Weather Review*, **124 (6)**, 1311–1322.

McGregor, J. L., 2013: Recent developments in variable-resolution global climate modelling. *Climatic Change*, **119**, 1–12.

McGregor, J. L., and M. R. Dix, 2008: An updated description of the conformal-cubic atmospheric model. *High Resolution Numerical Modelling of the Atmosphere and Ocean*, Springer-Verlag, New York, 51–75.

McGregor, J. L., K. C. Nguyen, D. G. Kirono, and J. J. Katzfey, 2016: High-resolution climate projections for the islands of Lombok and Sumbawa, Nusa Tenggara Barat Province, Indonesia: Challenges and implications. *Climate Risk Management*, **12**, 32 – 44.

Medvigy, D., R. L. Walko, M. J. Otte, and R. Avissar, 2013: Simulated changes in northwest US climate in response to Amazon deforestation. *Journal of Climate*, **26 (22)**, 9115–9136.

Melander, M., N. Zabusky, and J. McWilliams, 1988: Symmetric vortex merger in two dimensions: causes and conditions. *Journal of Fluid Mechanics*, **195**, 303–340.

Miura, H., M. Satoh, H. Tomita, A. Noda, T. Nasuno, S. Iga, M. Satoh, and T. Matsuno, 2007: A short-duration global cloud-resolving simulation with a realistic land and sea distribution. *Geophysical Research Letters*, **34**, L02 804.

Miyamoto, Y., Y. Kajikawa, R. Yoshida, T. Yamaura, H. Yashiro, and H. Tomita, 2013: Deep moist atmospheric convection in a subkilometer global simulation. *Geophysical Research Letters*, **40 (18)**, 4922–4926.

Müller, A., J. Behrens, F. X. Giraldo, and V. Wirth, 2013: Comparison between adaptive and uniform discontinuous Galerkin simulations in dry 2D bubble experiments. *Journal of Computational Physics*, **235**, 371–393.

Nair, R. D., and C. Jablonowski, 2008: Moving vortices on the sphere: A test case for horizontal advection problems. *Monthly Weather Review*, **136 (2)**, 699–711.

Paegle, J., 1989: A variable resolution global model based upon Fourier and finite element representation. *Monthly Weather Review*, **117 (3)**, 583–606.

Phillips, N. A., 1979: The nested grid model. Tech. Rep. NWS22, NOAA, 80 pp.

Pielke, R. A., and Coauthors, 1992: A comprehensive meteorological modeling system-RAMS. *Meteorology and Atmospheric Physics*, **49 (1-4)**, 69–91.

Prieto, R., B. D. McNoldy, S. R. Fulton, and W. H. Schubert, 2003: A classification of binary tropical cyclone-like vortex interactions. *Monthly Weather Review*, **131 (11)**, 2656–2666.

Prodhomme, C., L. Batté, F. Massonnet, P. Davini, O. Bellprat, V. Guemas, and F. J. Doblas-Reyes, 2016: Benefits of increasing the model resolution for the seasonal forecast quality in EC-Earth. *Journal of Climate*, **29 (24)**, 9141–9162.

Pudykiewicz, J. A., 2011: On numerical solution of the shallow water equations with chemical reactions on icosahedral geodesic grid. *Journal of Computational Physics*, **230 (5)**, 1956–1991.

Putman, W. M., and S.-J. Lin, 2007: Finite-volume transport on various cubed-sphere grids. *Journal of Computational Physics*, **227 (1)**, 55–78.

Putman, W. M., and M. Suarez, 2011: Cloud-system resolving simulations with the NASA Goddard Earth Observing System global atmospheric model (GEOS-5). *Geophysical Research Letters*, **38**, L16 809.

Rauscher, S. A., and T. D. Ringler, 2014: Impact of variable-resolution meshes on midlatitude baroclinic eddies using CAM-MPAS-A. *Monthly Weather Review*, **142**, 4256–4268.

Rauscher, S. A., T. D. Ringler, W. C. Skamarock, and A. A. Mirin, 2013: Exploring a global multiresolution modeling approach using aquaplanet simulations. *Journal of Climate*, **26 (8)**, 2432–2452.

Reed, K. A., and C. Jablonowski, 2011: An analytic vortex initialization technique for idealized tropical cyclone studies in AGCMs. *Monthly Weather Review*, **139**, 689–710.

Reed, K. A., and C. Jablonowski, 2012: Idealized tropical cyclone simulations of intermediate complexity: A test case for AGCMs. *Journal of Advances in Modeling Earth Systems*, **4 (2)**, 1942–2466.

Rhoades, A. M., P. A. Ullrich, and C. M. Zarzycki, 2018: Projecting 21st century snowpack trends in western USA mountains using variable-resolution CESM. *Climate Dynamics*, **50 (1-2)**, 261–288.

Ripa, P., 1993: Conservation laws for primitive equations models with inhomogeneous layers. *Geophysical & Astrophysical Fluid Dynamics*, **70 (1-4)**, 85–111.

Ripa, P., 1995: On improving a one-layer ocean model with thermodynamics. *Journal of Fluid Mechanics*, **303**, 169–201.

Ritchie, E. A., and G. J. Holland, 1993: On the interaction of tropical-cyclone-scale vortices. II: Discrete vortex patches. *Quarterly Journal of the Royal Meteorological Society*, **119**, 1363–1379.

Rostami, M., and V. Zeitlin, 2017: Influence of condensation and latent heat release upon barotropic and baroclinic instabilities of vortices in a rotating shallow water f-plane model. *Geophysical & Astrophysical Fluid Dynamics*, **111 (1)**, 1–31.

Ruge, J. W., S. F. McCormick, and S. Y. K. Yee, 1995: Multilevel adaptive methods for semi-implicit solution of shallow-water equations on a sphere. *Monthly Weather Review*, **123 (7)**, 2197–2205.

Sadourny, R., 1972: Conservative finite-difference approximations of the primitive equations on quasi-uniform spherical grids. *Monthly Weather Review*, **100 (2)**, 136–144.

Satoh, M., R. Nihonmatsu, and H. Kubokawa, 2013: Environmental conditions for tropical cyclogenesis associated with African Easterly Waves. *Scientific Online Letters on the Atmosphere*, **9**, 120–124.

Schmidt, F., 1977: Variable fine mesh in spectral global models. *Contributions to Atmospheric Physics, Oxford*, **50 (1)**, 211–217.

Scott, R., L. Harris, and L. Polvani, 2016: A test case for the inviscid shallow-water equations on the sphere. *Quarterly Journal of the Royal Meteorological Society*, **142 (694)**, 488–495.

Shepherd, T. J., and K. J. Walsh, 2017: Sensitivity of hurricane track to cumulus parameterization schemes in the WRF model for three intense tropical cyclones: impact of convective asymmetry. *Meteorology and Atmospheric Physics*, **129 (4)**, 345–374.

Shin, S.-E., J.-Y. Han, and J.-J. Baik, 2006: On the critical separation distance of binary vortices in a nondivergent barotropic atmosphere. *Journal of the Meteorological Society of Japan*, **84 (5)**, 853–869.

Skamarock, W., J. Oliger, and R. L. Street, 1989: Adaptive grid refinement for numerical weather prediction. *Journal of Computational Physics*, **80 (1)**, 27–60.

Skamarock, W. C., and J. B. Klemp, 1993: Adaptive grid refinement for two-dimensional and three-dimensional nonhydrostatic atmospheric flow. *Monthly Weather Review*, **121 (3)**, 788–804.

Skamarock, W. C., J. B. Klemp, M. G. Duda, L. D. Fowler, S.-H. Park, and T. D. Ringler, 2012: A multiscale nonhydrostatic atmospheric model using centroidal Voronoi tesselations and C-grid staggering. *Monthly Weather Review*, **140 (9)**, 3090–3105.

Skamarock, W. C., and Coauthors, 2008: A description of the advanced research WRF version 3. Tech. Rep. NCAR/TN-475+STR, NCAR, 113 pp.

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

Staniforth, A. N., and H. L. Mitchell, 1978: A variable-resolution finite-element technique for regional forecasting with the primitive equations. *Monthly Weather Review*, **106 (4)**, 439–447.

Stechmann, S. N., and A. J. Majda, 2006: The structure of precipitation fronts for finite relaxation time. *Theoretical and Computational Fluid Dynamics*, **20 (5-6)**, 377–404.

Tomita, H., 2008: A stretched icosahedral grid by a new grid transformation. *Journal of the Meteorological Society of Japan. Ser. II*, **86**, 107–119.

Tóth, G., and Coauthors, 2012: Adaptive numerical algorithms in space weather modeling. *Journal of Computational Physics*, **231 (3)**, 870–903.

Tumolo, G., and L. Bonaventura, 2015: A semi-implicit, semi-Lagrangian discontinuous Galerkin framework for adaptive numerical weather prediction. *Quarterly Journal of the Royal Meteorological Society*, **141 (692)**, 2582–2601.

Ullrich, P., C. Jablonowski, and B. van Leer, 2010: High-order finite-volume methods for the shallow-water equations on the sphere. *Journal of Computiational Physics*, **229**, 6104–6134.

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. *Quarterly Journal of the Royal Meteorological Society*, **140 (682)**, 1426–1440.

Ullrich, P. A., and C. Jablonowski, 2012: MCore: A non-hydrostatic atmospheric dynamical core utilizing high-order finite-volume methods. *Journal of Computational Physics*, **231 (15)**, 5078–5108.

Walko, R. L., and R. Avissar, 2011: A direct method for constructing refined regions in unstructured conforming triangular-hexagonal computational grids: Application to OLAM. *Monthly Weather Review*, **139 (12)**, 3923–3937.

Walsh, E. J., 2010: Moving mesh methods for problems in meteorology. Ph.D. thesis, Univeristy of Bath, 225 pp.

Waugh, D. W., 1992: The efficiency of symmetric vortex merger. *Physics of Fluids A: Fluid Dynamics (1989-1993)*, **4 (8)**, 1745–1758.

Weller, H., 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. *Journal of Computational Physics*, **308**, 102–123.

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

Williamson, 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. *Journal of Computational Physics*, **102**, 211–224.

Würsch, M., and G. C. Craig, 2014: A simple dynamical model of cumulus convection for data assimilation research. *Meteorologische Zeitschrift*, **5**, 483–490.

Yang, D., and A. P. Ingersoll, 2013: Triggered convection, gravity waves, and the MJO: A shallow-water model. *Journal of the Atmospheric Sciences*, **70 (8)**, 2476–2486.

Yessad, K., and P. Bénard, 1996: Introduction of a local mapping factor in the spectral part of the Météo-France global variable mesh numerical forecast model. *Quarterly Journal of the Royal Meteorological Society*, **122 (535)**, 1701–1719.

Zarzycki, C. M., and C. Jablonowski, 2014: A multidecadal simulation of Atlantic tropical cyclones using a variable-resolution global atmospheric general circulation model. *Journal of Advances in Modeling Earth Systems*, **6 (3)**, 805–828.

Zarzycki, C. M., and C. Jablonowski, 2015: Experimental tropical cyclone forecasts using a variable-resolution global model. *Monthly Weather Review*, **143**, 4012–4037.

Zarzycki, C. M., C. Jablonowski, and M. A. Taylor, 2014a: Using variable-resolution meshes to model tropical cyclones in the Community Atmosphere Model. *Monthly Weather Review*, **142 (3)**, 1221–1239.

Zarzycki, C. M., M. N. Levy, C. Jablonowski, J. R. Overfelt, M. A. Taylor, and P. A. Ullrich, 2014b: Aquaplanet experiments using CAM's variable resolution dynamical core. *Journal of Climate*, **27**, 5481–5503.

Zerroukat, M., and T. Allen, 2015: A moist Boussinesq shallow water equations set for testing atmospheric models. *Journal of Computational Physics*, **290**, 55–72.