

Lawrence Berkeley National Laboratory

LBL Publications

Title

Assessing Adaptive Mesh Refinement (AMR) in a Forced Shallow-Water Model with Moisture

Permalink

<https://escholarship.org/uc/item/83x4z0m6>

Journal

Monthly Weather Review, 147(10)

ISSN

0027-0644

Authors

Ferguson, Jared O
Jablonowski, Christiane
Johansen, Hans

Publication Date

2019-10-01

DOI

10.1175/mwr-d-18-0392.1

Peer reviewed

Assessing Adaptive Mesh Refinement (AMR) in a Forced Shallow-Water Model with Moisture

JARED O. FERGUSON AND CHRISTIANE JABLONOWSKI

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

HANS JOHANSEN

Lawrence Berkeley National Laboratory, Berkeley, California

(Manuscript received 9 November 2018, in final form 26 July 2019)

ABSTRACT

Two forced shallow-water flow scenarios are explored in a 2D fourth-order finite-volume dynamical core with adaptive mesh refinement (AMR) to investigate AMR's ability to track and resolve complex evolving features. Traditional shallow-water test cases are mainly characterized by large-scale smooth flows that do not effectively test the multiscale abilities of variable-resolution and AMR models to resolve sharp gradients and small-scale flow filaments. Therefore, adding forcing mechanisms to the shallow-water system to model key atmospheric processes adds complexity and creates small-scale phenomena. These can serve as foci for dynamic grid refinement while remaining simple enough to study the numerical design of a model's dynamical core. The first shallow-water flow scenario represents a strengthening, tropical cyclone-like, vortex that is driven by a Betts–Miller-like convection scheme. The second shallow-water test is built upon a barotropically unstable jet with an added Kessler-like warm rain scheme that leads to precipitating frontal zones. The key feature of both tests is that there is significant sensitivity to the model grid while converging (structurally) at high resolution. Both test cases are investigated for a series of uniform resolutions and a variety of AMR tagging criteria. The AMR simulations demonstrate that grid refinement can resolve local features without requiring global high-resolution meshes. However, the results are sensitive to the refinement criteria. Criteria that trigger refinement early in a simulation reproduce the uniform-resolution reference solutions most reliably. In contrast, AMR criteria that delay refinement for several days require careful tuning of the AMR thresholds to improve results compared with uniform-resolution simulations.

1. Introduction

The spherical shallow-water equations serve as an effective test bed for assessing numerical methods for atmospheric general circulation models (GCMs). They exhibit many of the dynamical characteristics of the full 3D equations with the advantage of being two-dimensional and thus less computationally intensive. Complex 3D models pair the GCM dynamical core with a collection of subgrid-scale parameterization schemes for unresolved physical processes. However, shallow-water models and the unforced test cases traditionally associated with them (Williamson et al. 1992) miss these subgrid-scale physical interactions (e.g., condensation and latent heat release), which play key roles in atmospheric and climatological phenomena.

Including simplified forcing mechanisms to represent 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 mimic some of the dynamical complexities of full 3D models and retain the nonlinearity of the physical processes. However, the forced shallow-water equations are still simple enough to effectively study key components of the dynamical core such as the numerical algorithm, computational grid, the physics-dynamics coupling and grid-scale interactions, and, for variable-resolution and adaptive mesh refinement (AMR) models, grid refinement strategies and efficacy. An advantage is that shallow-water models are computationally cheap and can be run at high resolutions down to a few kilometers.

A variety of studies have implemented forcing mechanisms in shallow-water models to study the fundamental

Corresponding author: Jared O. Ferguson, joferg@umich.edu

DOI: 10.1175/MWR-D-18-0392.1

© 2019 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](https://www.ametsoc.org/PUBSReuseLicenses) (www.ametsoc.org/PUBSReuseLicenses).

dynamical aspects of large synoptic-scale climatological features, including flows resembling key aspects of the Madden–Julian oscillation (MJO) (Ferreira et al. 1996; Yang and Ingersoll 2013), as well as intense, small-scale features including the evolution of tropical cyclone (TC)-like vortices, cumulus convection, and frontal propagation (Enagonio and Montgomery 2001; Hendricks et al. 2014). 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) implemented the moist-convective parameterizations in a fully nonlinear rotating shallow-water model although the Coriolis parameter was set to zero in their tests. The latter model was also used by Lambaerts et al. (2011) for dry and moist barotropic instability analyses, and by Lahaye and Zeitlin (2016) (with an added evaporation mechanism) to study the dynamical role of moisture in tropical cyclone instabilities. Furthermore, Rostami and Zeitlin (2017) implemented both a one-layer barotropic version and a two-layer shallow-water baroclinic version of the Lambaerts et al. (2011) model to investigate large-scale small-Rossby-number vortices.

Other recent frameworks for simulating precipitation and convection in the shallow-water system include the models by Würsch and Craig (2014) and Zerroukat and Allen (2015, hereafter ZA). Würsch and Craig (2014) developed a simplified model of cumulus convection which incorporated representations of updrafts, downdrafts, and idealized precipitation effects in a 1D nonrotating shallow-water model. ZA heuristically derived the rotating 2D shallow-water system from the 3D 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. ZA also implemented a three-state moisture model consisting of water vapor, cloud water and rainwater. The latter resembles the warm rain scheme by Kessler (1969) and is therefore characterized as “Kessler-like” here.

In this paper, we explore two forcing frameworks that mimic moisture interactions in a rotating shallow-water model on the sphere with adaptive mesh refinement

capabilities. The first framework extends the work by Bouchut et al. (2009) and Lahaye and Zeitlin (2016) to introduce forcing and precipitation mechanisms that resemble a Betts–Miller-like (Betts and Miller 1986) relaxation scheme. This framework is used to generate and strengthen TC-like vortices. In the second framework, the barotropic instability shallow-water test case of Galewsky et al. (2004) is implemented using the Kessler-like forcing as described in ZA. As the barotropic jet becomes unstable and collapses, frontal zones containing cloud water and precipitation develop. Using these “intermediate-complexity” frameworks as a test hierarchy, we investigate the distinctive dynamical features produced by the nonlinear physical processes in the shallow-water system. The goals of this paper are to 1) fully describe the forcing frameworks and initial conditions, 2) explore the sensitivity of the flow scenarios to the horizontal resolution for uniform-resolution configurations, and 3) use these as reference solutions to explore the transient flow scenarios with AMR. The overarching question is how well the AMR configuration of the model can capture and resolve these complex, moving, and growing flow features.

Though AMR methods for atmospheric flows were first explored several decades ago, they are still in a developmental stage and have not been used for real weather and climate model simulations yet. Jablonowski (2004), Ferguson et al. (2016) and Ferguson (2018) provide detailed reviews of the AMR approaches for atmospheric models and their AMR refinement criteria. We refer to these references for an in-depth overview of the atmospheric AMR literature. Here, we only briefly note that AMR refinement criteria can either be based on flow-based characteristics, like a vorticity or gradient threshold, or estimates of the local truncation error. This paper only assesses flow-based AMR criteria that are either guided by the dynamical fluid flow or physical attributes like cloud moisture. Our AMR assessments utilize the fourth-order finite-volume Chombo AMR model presented in McCorquodale et al. (2015), Ferguson et al. (2016) and Ferguson (2018). This shallow-water model implements dynamic grid refinement using a mapped-multiblock AMR technique which overlays the base grid with 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. In particular, we seek to quantify improvements gained from AMR grids and determine effective refinement criteria.

This paper is organized as follows. Section 2 provides a brief overview of the finite-volume model and the Chombo multiblock AMR techniques. Section 3 describes

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

Resolution (N_c)	No. of cells	Δx (km)	A_{avg} (km ²)	$A_{\text{min}}/A_{\text{max}}$	Eq. Res.	$\text{RLL}_{\text{equiv}}$
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°

the shallow-water equations with moist Betts–Miller-like forcing and the design of the TC-like vortex test case. Sections 4 and 5 compare numerical results of the strengthening vortex test case for uniform and AMR runs of varying resolution, respectively. The moist shallow-water system with Kessler-like forcing and the results of the barotropically unstable flow for both uniform and adaptive grids are presented in section 6. Section 7 summarizes the findings.

2. High-order finite-volume Chombo AMR model

We employ an unstaggered finite-volume (FV) mapped-multiblock dynamical core in a shallow-water configuration that is fourth-order accurate and adaptive in both space and time. Its AMR technique is based on the Chombo AMR library (Adams et al. 2015). An in-depth description of the model design on an equiangular cubed-sphere grid can be found in McCorquodale et al. (2015), Ferguson et al. (2016) and Ferguson (2018), and we only provide a brief summary here. Additional details about the cubed-sphere geometry and its metric terms are also provided in Ullrich et al. (2010). The model uses a classical fourth-order Runge–Kutta (RK4) time discretization scheme. In the horizontal domain, a fourth-order accurate finite-volume discretization is implemented to compute flux averages on the faces of each grid cell. Additionally, a sixth-order diffusive operator is applied to smooth the flux calculations while still maintaining the scheme’s fourth-order accuracy. The model is mass conserving and conserves energy up to the temporal truncation error, when limiters or explicit dissipation are not applied.

The model’s cubed-sphere grid consists of a cube whose six separate panels are projected onto the surface of a sphere. The cubed-sphere grid, which was originally suggested by Sadourny (1972), eliminates the two strong polar singularities found in spherical latitude–longitude grids and replaces them with weaker singularities along the edges of the cubed sphere and at the eight corners.

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 1 for the resolutions used in this paper. They range from about 320 to 5 km.

Our mapped-multiblock AMR approach implements a hierarchy of nested grid levels of increasing resolution. The grid resolution of an AMR level is defined by its refinement ratio to the grid resolution of the coarser level below it. In our investigations, we typically set the refinement ratio to $X = 4$, but $X = 2$ has also been tested. Finer levels are placed over regions where coarse cells have been marked (tagged) by the model as meeting the refinement criterion. We note that each refined block contains a collection of additional grid cells (between a minimum of 8×8 and a maximum of 32×32) in the horizontal direction. Therefore, a refined region always contains a collection of grid cells and is not refined on a cell-by-cell basis. This also provides somewhat of a buffer zone around the feature of interest. The block structure of the grid is schematically depicted in Fig. 1 of Ferguson et al. (2016).

Cell values at finer levels 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.

To ensure stable integration the Courant number is kept approximately constant, which requires that finer cells be substepped in time. Once the coarse cells have been advanced in time, the finer levels can be advanced by the required number of substeps using ghost cells, interpolated from the course cells, as boundary data.

After the substepping is complete, the values on the coarse grid are updated from the solution on the finer grid.

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 dynamical core can incorporate multiple levels of refinement, with the maximum number of levels refinement that could be activated in a run set by the user for each simulation, and tagging criteria can be uniformly enforced across all levels or required to scale with increasing resolution.

As an aside, the Chombo AMR model version used for the simulations here does not preserve monotonicity or apply filters to transported tracers. Therefore, small negative undershoots can occur in tracer fields, which needs to be remedied in future model versions. However, negative tracer values never affect the actual physical forcing calculations presented here. This is due to the design of the Betts–Miller-like and Kessler-like physics forcings that ignore or filter out negative moisture quantities.

3. Forced TC-like shallow-water vortices

The first “moist” framework of our shallow-water system simulates the growth and development of TC-like vortices using a Betts–Miller-like precipitation forcing mechanism. The flow field is initialized with weak vortices on a background field of uniform height at a state of rest. Evaporation and precipitation then cause these vortices to strengthen. After several days of strengthening, the vortices collapse and a more chaotic system evolves, characterized by several smaller vortices and a jetlike background flow. We first provide a description of the moist shallow-water system and an overview of the initial conditions. We then present the evolution of an isolated vortex at a uniform high resolution that serves as a reference solution.

a. “Moist” shallow-water equations

The rotating shallow-water equations on the sphere are modified to include the transport of a moisture variable and the effects of moisture, precipitation, and evaporation. These modifications extend the work by Bouchut et al. (2009) and Lahaye and Zeitlin (2016). As in Lahaye and Zeitlin (2016) our framework includes a newly added moisture equation [Eq. (3)] with a precipitation sink P and an evaporation source E . Precipitation is triggered when moisture levels exceed a saturation

threshold. A corresponding mass sink is added to the right-hand side (rhs) of the continuity equation [Eq. (2)]. The vector momentum, continuity, and moisture equations for this modified shallow-water system in conservation form are as follows:

$$\frac{\partial h\mathbf{v}}{\partial t} + \nabla \cdot (h\mathbf{v} \otimes \mathbf{v}) + f\hat{\mathbf{k}} \times (h\mathbf{v}) + gh\nabla H = 0, \quad (1)$$

$$\frac{\partial h}{\partial t} + \nabla \cdot (h\mathbf{v}) = -\beta P, \quad (2)$$

$$\frac{\partial hQ}{\partial t} + \nabla \cdot (hQ\mathbf{v}) = h(E - P). \quad (3)$$

Here \mathbf{v} is the horizontal velocity vector, $\mathbf{v} \otimes \mathbf{v}$ denotes the outer product of the velocity vector, $\hat{\mathbf{k}}$ is the unit vector in the vertical direction, f is the Coriolis parameter, g is the acceleration due to gravity, h is the height of the fluid, $H = h + h_b$ is the total height including the bottom topography h_b (here set to zero), ∇ is the horizontal gradient operator, $\nabla \cdot$ represents the horizontal divergence, and the dimensionless moisture variable Q represents bulk humidity. The latent heat release from precipitation cannot be explicitly modeled in this equation set. It is therefore indirectly represented as a mass sink in Eq. (2) that is governed by an adjustable constant β (defined later). More detailed motivations for similar forcing mechanisms are presented in Bouchut et al. (2009) and Lahaye and Zeitlin (2016). Note that the conservation form necessitates the use of $h(E - P)$ on the rhs of Eq. (3). If a model is formulated in advective form, the following moisture equation needs to be used:

$$\frac{\partial Q}{\partial t} + \mathbf{v} \cdot \nabla Q = E - P, \quad (4)$$

which mimics Eq. (1) of Lahaye and Zeitlin (2016). The forcings on the rhs of the vector momentum equation (rhs = 0) and continuity equation (rhs = $-\beta P$) are independent of the model formulation and identical for models in either advective or conservation form. This mimics the physics–dynamics coupling strategies of real weather and climate models. The latter ignore physics forcings in the dynamical cores and add the physical forcings to the rhs of the equations after an adiabatic dynamical core integration step.

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

$$P = \frac{Q - Q_s}{\tau} \mathbf{H}(Q - Q_s), \quad (5)$$

with a relaxation time of $\tau = 1$ day; $\mathbf{H}(\cdot)$ is the Heaviside function so that $P = 0$ whenever $Q \leq Q_s$. In contrast to

the uncapped, wind speed dependent, E formulation in [Lahaye and Zeitlin \(2016\)](#) we link the evaporation rate E in Eq. (3) to an intermediate evaporation rate E_r and the concept of a “moisture reservoir” C_r to limit the evaporative forcing effect. This was motivated by the fact that the [Lahaye and Zeitlin \(2016\)](#) formulation can lead to runaway supersaturation and very large height forcing which causes exceedingly high wind velocities and negative height values in longer simulations. Thus, our evaporation rate E is capped and given by

$$E = \frac{1}{\Delta t} \min(E_r \Delta t, C_r), \quad (6)$$

where Δt is the model time step. This formulation utilizes the intermediate evaporation rate E_r :

$$\begin{aligned} E_r &= \alpha_e |\mathbf{v}| \quad \text{for } |\mathbf{v}| < v_{\max} \\ E_r &= \alpha_e (v_{\max}) \quad \text{for } |\mathbf{v}| \geq v_{\max}, \end{aligned} \quad (7)$$

where the evaporated moisture scales with wind speed $|\mathbf{v}|$ and depends on the tuning coefficient α_e as in [Lahaye and Zeitlin \(2016\)](#). In the simulations presented below, we select $\alpha_e = 0.055 \text{ m}^{-1}$ and $v_{\max} = 30 \text{ m s}^{-1}$. The new aspect is that we cap E_r for velocity magnitudes greater than v_{\max} .

The second new aspect is that E is also limited by the available moisture reservoir. This reservoir is somewhat analogous to the heat content of an ocean surface and its limiting effects on tropical cyclone intensity (the so-called cold-wake effect of tropical cyclones). Equation (6) demonstrates that the evaporation rate E cannot exceed the amount of moisture remaining in the reservoir $C_r(\lambda, \phi)$ for any given longitude λ and latitude ϕ . The reservoir is affected by the evaporation rate E , which removes moisture from the reservoir, and a Newtonian relaxation that slowly returns the reservoir to its initial state. This time-dependent moisture reservoir is therefore defined as

$$\frac{\partial C_r(\lambda, \phi)}{\partial t} = -E + \frac{1}{\tau_c} [C_0(\lambda, \phi) - C_r(\lambda, \phi)], \quad (8)$$

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

$$C_0(\lambda, \phi) = C_{\max} \cos^4(\phi), \quad (9)$$

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. (6), ensures that C_r cannot become negative, and E tends toward zero for a declining moisture reservoir.

Another difference to the [Bouchut et al. \(2009\)](#) and [Lahaye and Zeitlin \(2016\)](#) forcing mechanisms is that they implemented a constant β in Eq. (2) to represent the precipitation mass sink. We implement a variable β that scales with fluid height. As the fluid height decreases below the initial value, β is also decreased, 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 forcing and limits the vortex strength. Thus the coefficient β is set as

$$\beta = \beta_0 \tanh \left[\max \left(0, \lambda_b \frac{h - h_t}{h_0 - h_t} \right) \right]. \quad (10)$$

We set the constants to $\beta_0 = 0.01 \text{ m}$ and $\lambda_b = 10$. The parameter h_0 is the initial background height and h_t is the cutoff height set to 75% of h_0 . As observed in the TC-like test case, the β term does not limit the vortex strengthening until the vortex is well established.

b. Vortex initialization technique

The initial conditions for the TC-like vortex test case consist of one or more small, weak vortices that are seeded onto constant-height and zero-velocity background conditions. In addition, each vortex is initialized with some supersaturation to ensure that the forcing is triggered immediately, creating convergence and allowing the vortices to strengthen.

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

$$h = h_0 - h_f \exp \left[- \left(\frac{r}{r_w} \right)^2 \right]. \quad (11)$$

Here, $h_0 = 4000 \text{ m}$ is the uniform initial height, h_f is the depth of the Gaussian depression, and r_w is the radius of maximum wind. The values for h_f and r_w are designated later and, if desired, multiple vortices with varying characteristics can be placed into the domain. The great circle distance r from point (λ, ϕ) to the vortex center (λ_d, ϕ_d) is given by

$$r = a \arccos[\sin \phi_d \sin \phi + \cos \phi_d \cos \phi \cos(\lambda - \lambda_d)], \quad (12)$$

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

$$u = -\frac{g}{af} \frac{\partial h}{\partial \phi} \quad \text{and} \quad v = \frac{g}{af} \frac{\partial h}{\cos \phi \partial \lambda}. \quad (13)$$

The Coriolis parameter $f = 2\Omega \sin\phi$ utilizes the constant rotation rate $\Omega = 7.292 \times 10^{-5} \text{ s}^{-1}$. If desired, a steering background wind field could be overlaid. However, this is not done here. The full equations for u and v in Eq. (13) are also provided in Chapter 3.3.2 of Ferguson (2018). 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 the forcing triggered immediately by an initial supersaturation. The associated mass sink in Eq. (2) creates significantly larger gravity waves which quickly overwhelm the effects 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

$$Q_0 = \min(Q_{\text{sat}} - Q_{\text{off}}, \cos\phi + Q_{\text{min}}), \quad (14)$$

with the saturation value $Q_{\text{sat}} = 0.9$, $Q_{\text{off}} = 0.01$ sets the maximum background field just below saturation, and $Q_{\text{min}} = 0.05$ 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

$$Q = Q_0 + Q_f \exp\left[-\left(\frac{r}{r_w}\right)^2\right], \quad (15)$$

where Q_f is the maximum moisture perturbation value at the center of the vortex stated later. 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 vortices are summed up and overlaid on the background fields. If the vortices are close together, their fields may significantly overlap.

c. Reference solution: Evolution of a strengthening vortex

We implement the test case by initializing one isolated vortex at a uniform high resolution and observe its evolution over 12 days. This c2048 resolution (~ 5 km) run serves as a reference solution for AMR and coarser

resolution runs. The initial weak vortex is centered at ($\lambda_d = 0^\circ$, $\phi_d = 10^\circ$) with the 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. 1a. In addition, we initialize a non-symmetric binary pair of vortices 90°E of the isolated vortex. The two additional vortices have little effect on the evolution of the isolated main vortex during the first 12 days of simulation. They are added to the test case to help form a complex chaotic global flow approximately 14 days into the simulation. The initialization details for the two vortices and the extended-time (up to 16 days) results focusing on all three vortices are discussed in detail in Ferguson (2018).

The evolution of the vortex's relative vorticity profile over a period of 12 days is depicted in Fig. 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. At day 8 the maximum wind has reached 31.2 m s^{-1} , while by day 10 it has increased to 69.0 m s^{-1} . During its evolution, the vortex develops a symmetric ring of maximum vorticity around day 4. This ring can be clearly seen at days 6 and 7 (Figs. 1d,e). As the vortex rapidly intensifies, this symmetric ring becomes elongated, as seen at day 8 in Fig. 1f and collapses. The filaments of large positive vorticity then begin to collate, creating a concentrated area of maximum vorticity (day 9, Fig. 1g). A small section of the vorticity filament is not reincorporated into the center spirals of the main vortex (day 10, Fig. 1h), becoming a separate, smaller secondary vortex pair seen by day 12 (Fig. 1i).

Figure 2 provides a snapshot at day 8 of the height field (Fig. 2a), wind magnitude (Fig. 2b), precipitation rate (Fig. 2c), and ocean-like reservoir of available moisture for evaporation (Fig. 2d), corresponding to the relative vorticity profile in Fig. 1f. Similar to the vorticity profile at day 8, an elongated ring of strongest winds and heaviest precipitation is present around the vortex center. In addition, a Rossby wave train forms toward the southeast is visible in the wind and precipitation fields (Figs. 2b,c). The area of low moisture reservoir levels to the southeast of the vortex in Fig. 2d shows where evaporation has been the strongest and reflects the path of the vortex. This lack of moisture for evaporation suppresses precipitation and vortex strengthening. The main vortex strengthens more slowly after day 10, reaching a peak wind speed of 176 m s^{-1} before day 13. It then starts to weaken

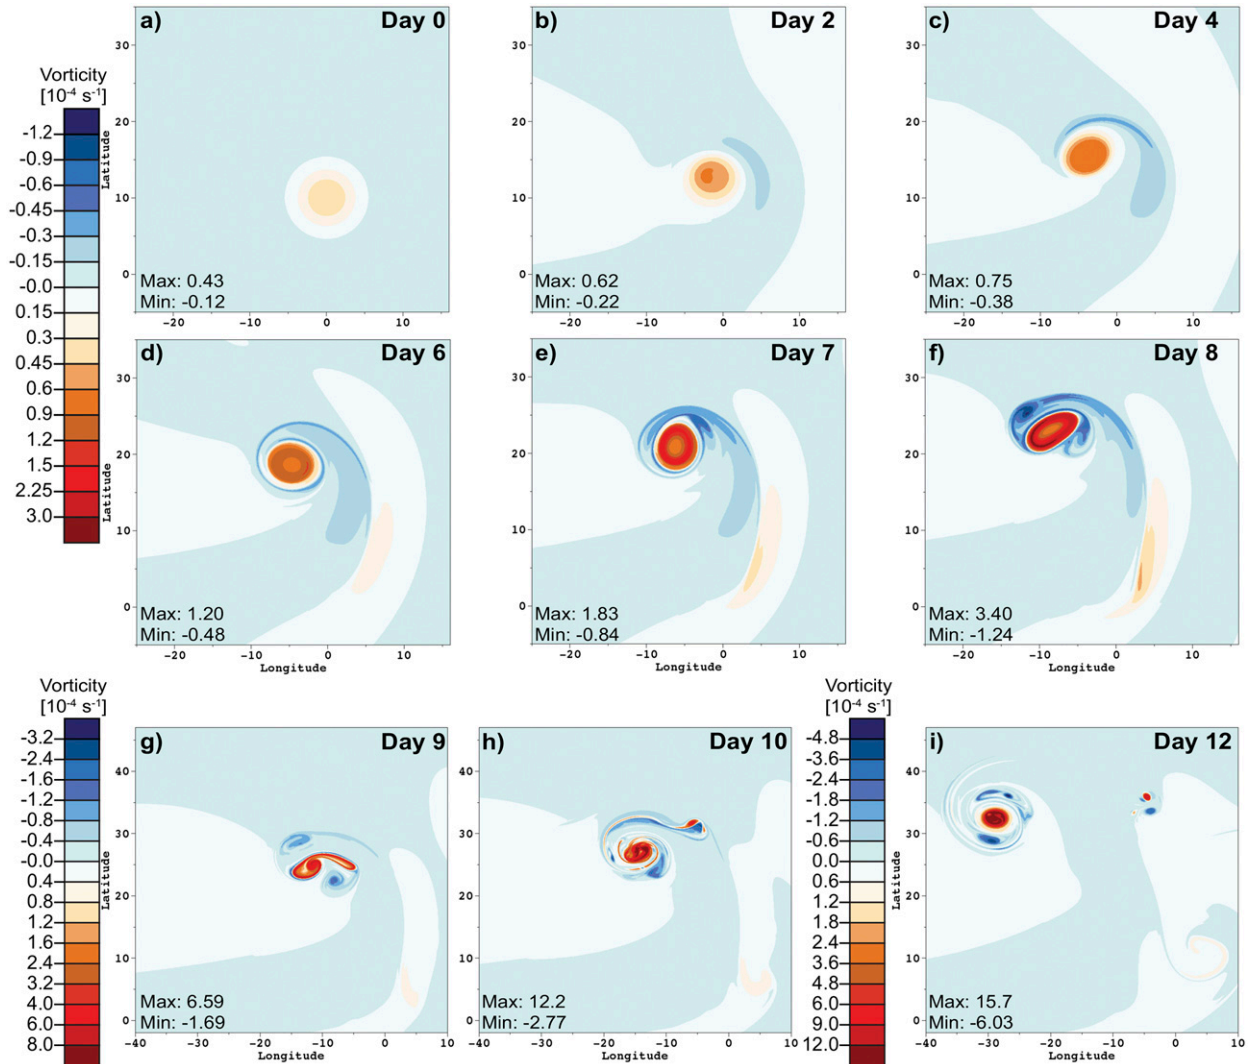


FIG. 1. The evolution of the relative vorticity for an isolated strengthening vortex in a c2048 uniform run. (a)–(f) Relative vorticity plots for the initial condition, day 0, and days 2, 4, 6, 7, and 8 with color contour range of -1.2×10^{-4} to $3.0 \times 10^{-4} \text{ s}^{-1}$. (g), (h) Relative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

having reached higher latitudes where less moisture is available in the reservoir.

4. TC-like vortices: Impact of the resolution in uniform-resolution simulations

We implement the TC-like vortex test case with a series of uniform resolution runs (as listed in Table 1) which can readily be compared to the high-resolution c2048 ($\approx 5 \text{ km}$) simulation. The time step changes in proportion to the grid resolution. Figure 3 depicts the growth in the maximum magnitude of relative vorticity for the main vortex in uniform resolution runs with grid spacings between $\approx 160 \text{ km}$ (c64) and $\approx 5 \text{ km}$ (c2048).

In general, vortices in higher-resolution runs have larger maximum vorticity over the first 10 days, while the c64 resolution is too coarse to properly resolve the vortex and it slowly weakens.

In Fig. 3 we see that the c256 run strengthens more rapidly after day 9 in comparison to higher resolution runs. The c256 run’s peak vorticity is higher than the c512 and c1024 runs by day 12, which peak between days 11 and 12 before weakening. This is in contrast to a continued increase in strength with increasing resolution that would be expected in a pure dynamics simulation. This difference in vorticity strength can be seen in Fig. 4, which depicts the relative vorticity field for uniform runs c256 ($\approx 40 \text{ km}$), c512 ($\approx 20 \text{ km}$), and c1024 ($\approx 10 \text{ km}$) at

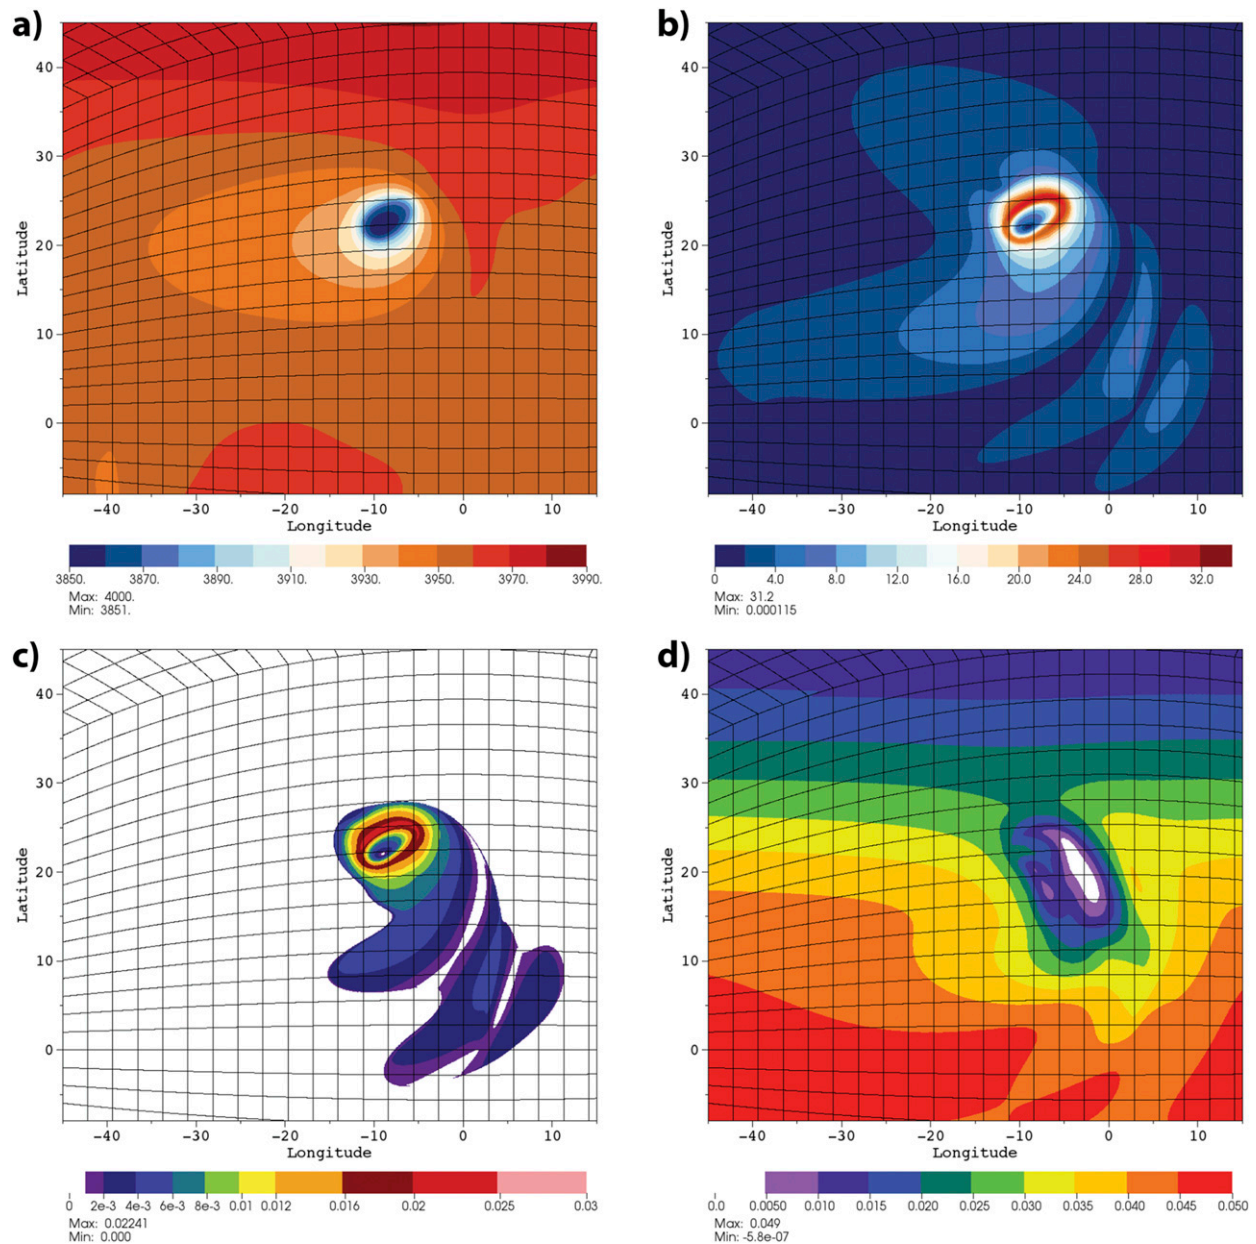


FIG. 2. Day 8 plots for the uniform c2048 run of the isolated strengthening vortex for several variables: (a) Height field (m), (b) wind magnitude (m s^{-1}), (c) instantaneous precipitation rate (moisture value per day), and (d) reservoir moisture content (moisture value). These plots correspond to the day 8 vorticity plot in Fig. 1f, though note the larger latitude–longitude domain in these plots.

day 9 in Figs. 4a–c and day 12 in Figs. 4d–f respectively. The collapse and roll-up of the main vortex at day 9 and the secondary vortex observed at day 12 are clearly resolved in the c1024. In contrast, the c256 run neither develops the distinct vortex ring structure nor its collapse. The c256 and coarser resolutions are unable to resolve the ring-like structure that the physics forcing induces. Instead, they dissipate that finescale feature into a broader area of vorticity that

continues to strengthen rather than collapsing and weakening.

5. TC-like vortices: Impact of the resolution in AMR simulations

The goal of this section is to assess the ability of different AMR refinement criteria to achieve comparable results to the uniform-resolution simulations of the




FIG. 3. The daily maximum relative vorticity of the strengthening vortex over a period of 13 days for uniform runs with resolutions from c64 to c2048. This vorticity value is only for the central vortex; it excludes any maximum that may occur in the secondary vortices spunoff.

TC-like vortices. First, the chosen tagging criteria for refinement are discussed.

a. AMR tagging criteria

Our AMR analyses focus on the relative vorticity field as the tagging variable since it reliably picks out a combination of large-scale and finescale flow structures, and is sensitive to changes in resolution. We conduct AMR runs with three different tagging criteria:

- Tag 1: a scaled vorticity threshold with a base threshold of $|\zeta| > 2 \text{ days}^{-1} = 2.3 \times 10^{-5} \text{ s}^{-1}$,
- Tag 2: a scaled vorticity threshold with a base threshold of $|\zeta| > 3 \text{ days}^{-1} = 3.5 \times 10^{-5} \text{ s}^{-1}$,
- Tag 3: a constant vorticity threshold of $|\zeta| > 5 \text{ days}^{-1} = 5.8 \times 10^{-5} \text{ s}^{-1}$.

The constant threshold triggers all levels of refinement simultaneously when the threshold is surpassed. The two scaled refinement criteria increase the threshold value for triggering each additional level of refinement beyond the first AMR level proportionally, in a one-to-one ratio, with increasing resolution. For example, in a c32-base two-level AMR run with a $X = 4$ refinement ratio using Tag 1, the first level of AMR, the c128 level, is triggered when relative vorticity exceeds $2.3 \times 10^{-5} \text{ s}^{-1}$, but the second layer of AMR, the c512 level, is triggered when relative vorticity exceeds $9.2 \times 10^{-5} \text{ s}^{-1}$. This delay in triggering higher refinement results in some loss of

detail on features of interest in their initial stages, but significantly fewer grid cells, and the corresponding computational cost, are used.

The Tag 1 vorticity threshold results in the initial precipitation region being covered by the first level of refinement, while the threshold in Tag 2 results in refinement over only part of the initial precipitation area. Tag 3's value is set above the initial maximum vorticity so that it is triggered within the first few days of the simulation. The higher vorticity threshold reduces some of the computational cost associated with the Tag 3 constant threshold triggering at once all levels of refinement. Other thresholds were explored but those were either too restrictive, not providing enough resolution to properly resolve the vortex, or triggered too much refinement, significantly increasing the computational cost. The selected thresholds are a middle ground and offer some examples of both cases.

b. Evolution of maximum vortex strength

The growth of the peak relative vorticity is shown in Fig. 5 for the various choices of refinement criteria, base resolution, and number of AMR levels. The plotted AMR runs represent several combinations of starting base resolution, number of AMR levels, refinement ratios between levels, and tagging criteria to explore effective combinations of these attributes. In total, 12 of the 13 configurations have a maximum of only 1 or 2 levels of refinement. We implement one c32 base-level AMR with three levels of x4 refinement with Tag 1 to observe what improvement is provided by the triggering of a c2048 high-resolution level late in the simulation compared to the c32 base-level run with only two levels of AMR. Several runs have a refinement ratio of x2 instead of x4, and two c256 base-level AMR runs with Tag 2 and Tag 3 have an initial level of x2 refinement and a second level of x4 refinement. This small-ratio attribute was implemented to observe what effects more gradual transitions have on the vortex. Since Tag 3 does not trigger initial AMR refinement, only base resolutions of c128 or higher could be used since the vortex does not strengthen in lower-resolution runs.

Figure 5 shows that the AMR runs are able to follow the growth trajectory of uniform runs with the same resolution as the finest AMR level. AMR runs, where refinement occurs within the first day, are able to match the maximum vorticity of the uniform run with the same resolution as the highest AMR level. The c256-base 1-level AMR runs with Tag 1 and Tag 2 refinement in Figs. 5a and 5b and the c512 base-level AMR runs in Figs. 5b and 5c fit this category. The maximum relative vorticity for several other AMR runs remain slightly lower than the corresponding uniform run because the

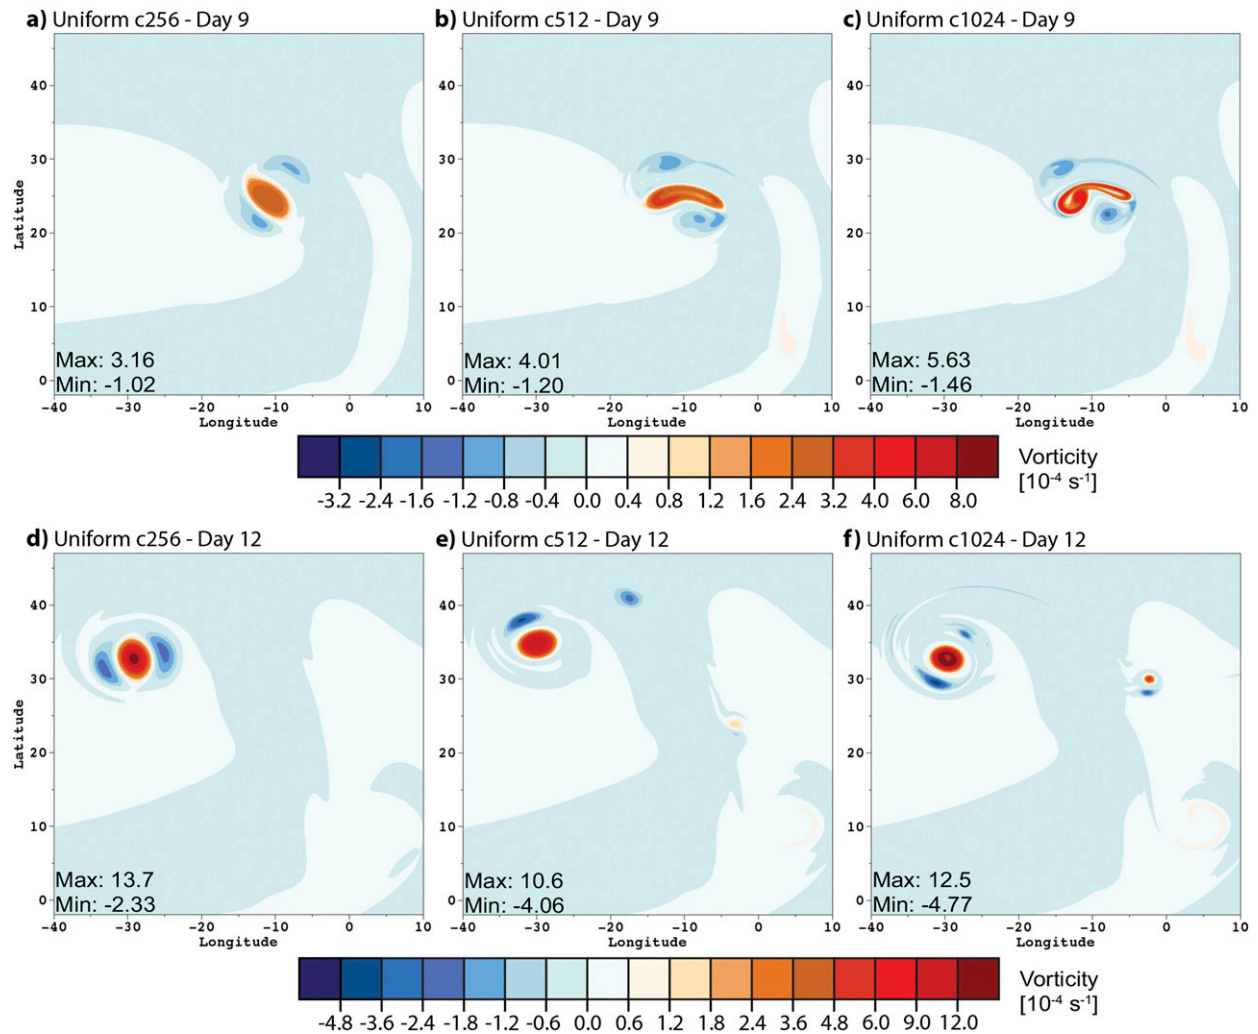


FIG. 4. Relative vorticity field of the strengthening vortex case at (a)–(c) day 9 and (d)–(f) day 12 for uniform runs (a),(d) c256 resolution, (b),(e) c512 resolution, and (c),(f) c1024 resolution. These plots correspond to the day 9 uniform c2048 plot Fig. 1g, and day 12 plot Fig. 1i.

higher-resolution refinement is triggered several days into the simulation. This delay can be seen in the c128/c512/c2048 AMR run with Tag 3 in Fig. 5c. The c512 and c2048 refinement levels in that run are not triggered until after day 2.

The key exceptions to this trend are the two c32 Tag 1 AMR runs in Fig. 5a. The three-level (c32/c128/c512/c2048) and the two-level (c32/c128/c512) AMR runs have maximum vorticities nearly 40% and 25% weaker, respectively, at day 12 in comparison to their uniform reference runs. These two runs had the c128 refinement level triggered at initialization. However, the c512 AMR level is not triggered until after six and half days, and c2048 AMR level after 10 days, so these runs follow the low-resolution regime more closely. In contrast, the c64-base AMR runs with Tag 1 (Fig. 5a) and Tag 2 (Fig. 5b) do resolve the vortex's rapid strengthening,

albeit delayed, and result in comparable vortex strength by the end of the situation. Both runs have initial c256 level refinement, but the c1024 level is not triggered until after day 7 for the Tag 1 run and day 8 for the Tag 2 one. Though the trigger times for the second AMR level are comparable to the c32-base AMR runs, the c64 AMR runs do follow the c1024 maximum vorticity by day 10. The higher refinement occurs earlier in the c64 AMR runs before the rapid intensification and vortex collapse period.

Figure 6 depicts the mass loss ratio between select uniform and AMR runs and the c2048 reference run. Given that the mass sink forcing is key to the vortex strengthening, the mass loss for the select AMR runs aligns well with the maximum vorticity comparisons

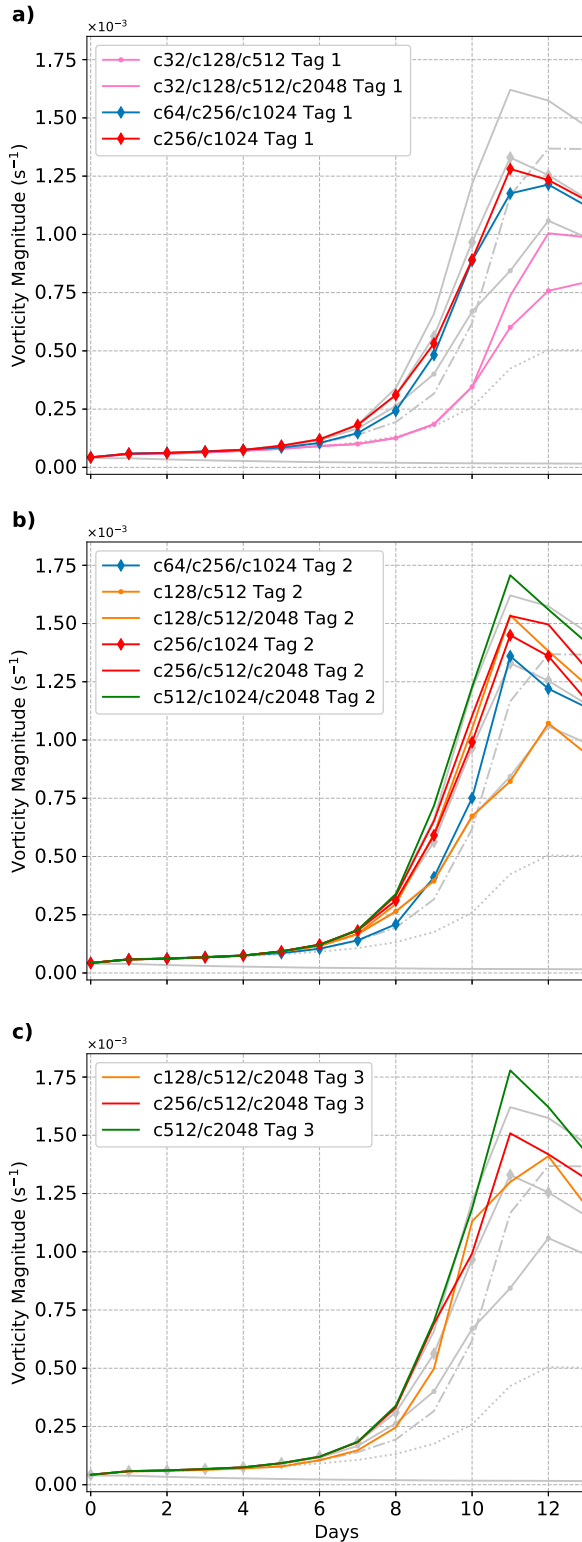


FIG. 5. Maximum relative vorticity of the strengthening vortex over a period of 13 days for (a) AMR runs using the Tag 1 refinement criteria, (b) AMR runs using the Tag 2 criteria, and (c) AMR runs using the Tag 3 criteria. For comparison purposes

seen in Figs. 3 and 5. The higher resolution uniform runs have comparable mass losses, while the coarser c128 run and the c32-base Tag 1 AMR run retain more mass as expected. The other Tag 1 and Tag 2 AMR runs begin with slightly slower mass loss rates than the c2048 reference run but as more refinement is triggered, the mass loss rate increases until it is comparable to the reference run. The mass loss in the c128-base Tag 3 AMR run follows the uniform c128 run more closely even though its maximum vorticity and other features are similar to the higher resolution runs. However, the run has a comparable mass loss rate to the reference run once refinement is triggered around day 3. After day 9 its mass loss rate decreases and once again tracks the c128 run. This may be the result of the higher AMR threshold that prevents refinement over areas south and west of the main vortex which become the main areas of precipitation (mass loss) toward the end of the simulation. As an aside, convection schemes in real weather and climate models are not a source or sink for dry air mass. The mass loss here is a special characteristic of the chosen Betts–Miller-like forcing mechanism for the shallow-water equations.

c. Vortex structure at days 9 and 12

Figures 7 and 8 depict the relative vorticity field for day 9 and day 12, respectively, for six of the 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 Figs. 1 and 4. At day 9 the effect of AMR is relatively straight forward and expected. The earlier in the simulation higher resolutions are applied the more closely the run resembles the uniform c2048 run at day 9. The three-level c32-base AMR run with Tag 1 (Fig. 7a) only has c128 and c512 AMR levels triggered. Thus its vortex evolution is delayed several days and more comparable to the uniform c2048 vortex at day 7 (Fig. 1d). The vortices in the two c64-based AMR runs with Tag 1 (Fig. 7b) and Tag 2 (Fig. 7e) are also delayed, though by less than a day. They are similar to the c512 uniform run at day 9 (Fig. 4b). The vortex in the c64-base AMR Tag 1 run is slightly stronger and more deformed than the Tag 2 run, reflecting the earlier refinement triggered by the lower threshold. The c128 two-level AMR (Fig. 7c) and

←
the uniform run lines from Fig. 3 have been imposed in light gray on the other three plots. The labeling of each run shows the base resolution and the resolution of each AMR level that can be triggered (e.g., the c256/c1024 Tag 1 run has a base resolution of c256 and one level of AMR at c1024 resolution).

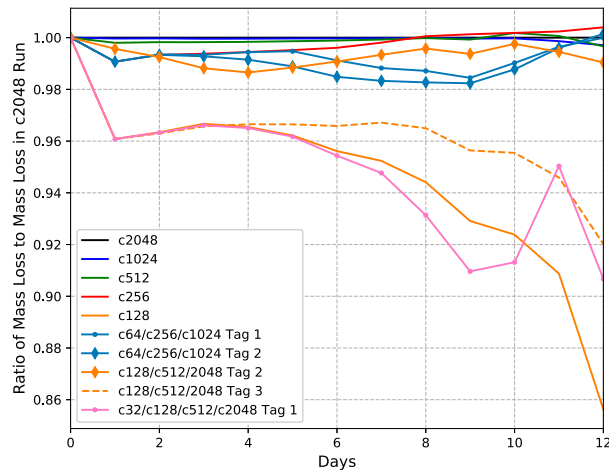


FIG. 6. The ratios of mass loss in the uniform and AMR runs to the mass loss in the c2048 uniform reference run as a function of time. Thus the c2048 run line (solid black) is constantly one, and runs that have lost less mass than the uniform c2048 reference, have ratios of less than one. The c64 uniform run is not included in the plot to show more clearly the variation in the runs that do have a strengthening vortex.

c256 two-level AMR (Fig. 7d) runs with Tag 2 criterion, closely match the day 9 vorticity magnitudes of the c2048 uniform run. However, a small delay is evident by the less developed comma-like positive vorticity feature of the main vortex ring. In c128-base two-level AMR run with the Tag 3 criteria in Fig. 7f the vortex structure is comparable to the c2048 uniform run at 8.5 days.

Figure 8 shows that at day 12 most of the AMR runs, even with delayed vortex evolution, now have comparable strength and vortex structure, including the spinoff of a smaller secondary vortex, as observed in the corresponding high-resolution uniform runs. The main exception is the vorticity field for the c32-base Tag 1 AMR run (Fig. 8a), which is still delayed with a vorticity field more comparable to day 10 of the c2048 uniform run (Fig. 1h). The c64-base Tag 2 AMR run (Fig. 8e), unlike the Tag 1 run (Fig. 8b), fails to reproduce the secondary vortex spinoff. The c128-base AMR Tag 3 run (Fig. 8f) and the c256-base AMR Tag 2 run (Fig. 8d) effectively capture the anticyclonic filaments around the main vortex observed in the uniform c2048 run at day 12 (Fig. 1i).

d. Comparison of growth and coverage of refinement levels for two AMR runs

Table 2 compares the number of grid cells at each AMR level at several days for both the Tag 2 and Tag 3 c128 two-level AMR runs. The Tag 2 run's initial c512 refinement appears to outweigh the benefits of the extra c2048 resolution that the Tag 3 run has after day 2. However, by the end of the run the additional c2048

refinement in Tag 3 provides better coverage of finescale structures around the main vortex as seen in Figs. 7 and 8. A scaled tagging criteria similar to Tag 2 but with less than one-to-one scaling with resolution (e.g., the threshold only increases by half with doubling of resolution) may strike a balance by providing both early refinement and more coverage of key features with the highest resolution levels, without significantly increasing computational costs.

Computationally, both AMR runs had significantly fewer grid cells than the c2048 uniform run. At day 12, the c128 two-level Tag 2 and Tag 3 AMR runs have c512 resolution over approximately 5% and 2.7% of the cubed sphere, respectively. The c2048 resolution covered 0.6% of the cubed-sphere surface for the Tag 2 run and 1.3% for the Tag 3 run. For a 12-day simulation the number of calculations for the c128 two-level AMR Tag 2 run is approximately 0.2% the number of calculations needed for the c2048 uniform run, and for the c128 two-level AMR Tag 3 run it is roughly 0.4% that of the c2048 run. Though the regridding process and interpolation between coarse-fine levels do add some additional overhead, AMR costs are dominated by the evolution of the finest grids. Table 5 in Ferguson et al. (2016) shows that the number of grid cells is a proxy for the total computational cost for the Chombo-AMR model.

e. Summary assessment of the AMR simulations

A key delineation between all 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.

6. Forcing of the shallow-water equations via a Kessler-like warm rain scheme

An alternative setup for a moist, forced shallow-water system can be heuristically derived from the

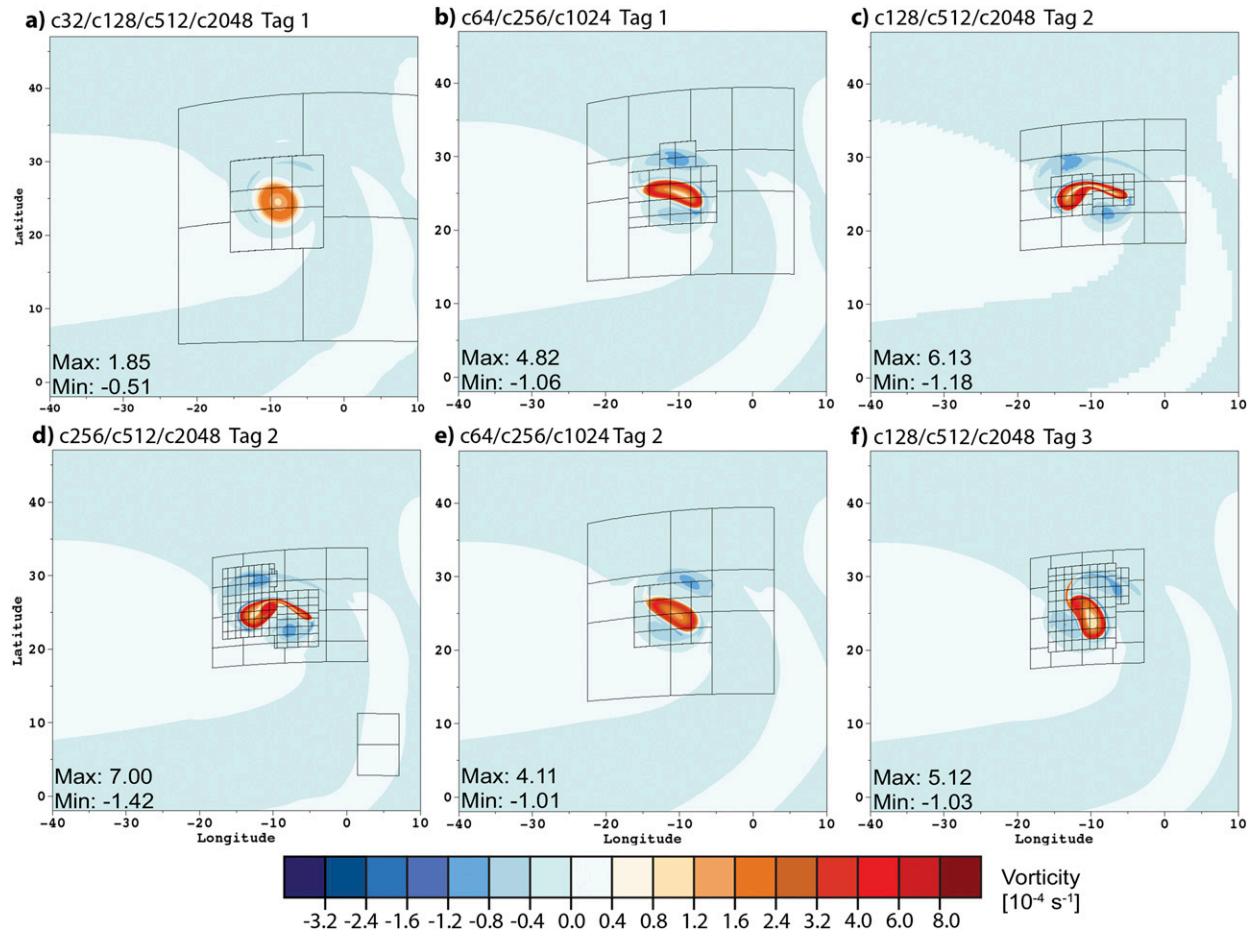


FIG. 7. Relative vorticity fields at day 9 for six AMR runs of the strengthening vortex case. These plots correspond to the day 9 uniform plots in Figs. 1g and 4a–c. The block structures of the multiple refinement levels are outlined in black.

moist 3D Boussinesq equations. This is done via depth averaging in *ZA*. This framework includes a three-state moist physics model that simulates water vapor, cloud water and rainwater, similar to the warm-rain parameterization developed by 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 release due to precipitation increases the average temperature of the fluid, which is coupled to the momentum equations. This is in contrast to our earlier moist Betts–Miller-like framework, in which the precipitation impacts the continuity equation. The latter somewhat mimics a “convective mass flux”. A brief discussion comparing the two models is presented in appendix A of Bouchut et al. (2009). We implement the Kessler-like physics forcing here for both uniform and AMR runs, and use the barotropic instability test case of Galewsky et al. (2004) for the model initialization.

a. The shallow-water and physics equations

The forced shallow-water equations implemented in this test case are presented in section 2.1 of *ZA*. The *ZA* physics scheme consists of three forms of moisture which are the water vapor q_v , cloud water q_c , and rainwater q_r , and a depth-averaged temperature variable θ . When the local value of q_v exceeds a prescribed function for the saturation a fraction of the oversaturation is condensed into cloud water with a corresponding latent heat release that increases the local temperature θ . In the same manner, a fraction of cloud water (present in unsaturated air) evaporates 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 subsaturated air induced by the changing temperature. Cloud water can also be converted to rainwater when q_c exceeds a prescribed threshold and a fraction of the excess cloud water is then converted to rainwater. The equation set for this physics scheme is

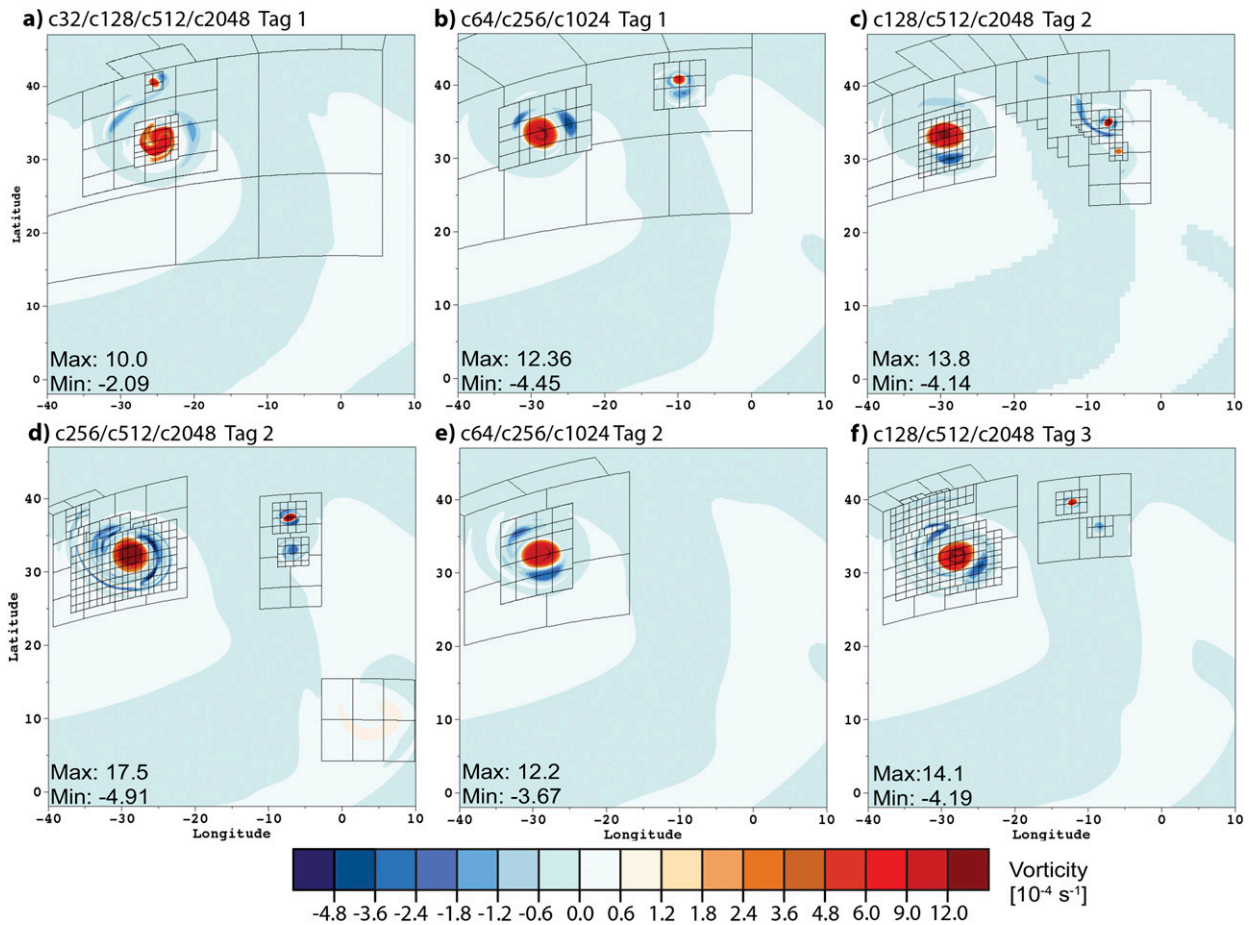


FIG. 8. As in Fig. 7, but for day 12 after the small secondary vortex has spun off. These plots correspond to the day 12 uniform plots in Figs. 1i and 4d–f. The block structures of the multiple refinement levels are outlined in black. Note the vorticity maximum in (d) is located in the secondary vortex.

presented in section 3 of ZA. In this setup, the rainwater q_r is removed from the system and counted as precipitation, whereas in ZA the q_r variable is advected around. In addition, our physics forcing functions for temperature and the moisture quantities are multiplied with h due to the conservation form of our equation set. In both setups, the rainwater does not feed back onto the dynamics so processes such as rain evaporation and accretion are neglected. It is important to note that the moisture q_v , q_c , and q_r and temperature θ variables as well as the related constants from ZA, though derived from realistic values are not suitably scaled to commonly used physical units or value ranges. Given the simplicity of the forcing, its constants can be, in principle, arbitrarily chosen. However, we use the identical parameter values described in ZA.

b. 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 rollup of the jet. The initial velocity and height fields, along with the height perturbation, are defined in Galewsky et al. (2004). We add the θ and q_v profiles to this initialization, and set the initial q_c and q_r fields to zero. The initial θ profile is a quadratic

TABLE 2. A snapshot of the number of grid cells per refinement level every three days for the two c128 2-level AMR runs using Tag 2 and Tag 3 refinement criteria.

Time	Number of cells per AMR level			
	c128 2-level AMR Tag 2 run		c128 2-level AMR Tag 3 run	
	c512 level	c2048 level	c512 level	c2048 level
Day 0	1.03×10^4	—	—	—
Day 3	1.43×10^4	—	3.90×10^3	1.18×10^4
Day 6	2.14×10^4	—	1.56×10^4	7.51×10^4
Day 9	3.07×10^4	8.10×10^4	1.87×10^4	1.24×10^5
Day 12	7.67×10^4	1.46×10^5	4.31×10^4	3.22×10^5

function with a north–south variation taken from ZA so that

$$\theta(\phi, \lambda) = \theta^{\text{SP}} \left(\phi - \frac{\pi}{2} \right) \phi - (1 - \mu_1 \theta^{\text{EO}}) \left(\phi + \frac{\pi}{2} \right) \left(\phi - \frac{\pi}{2} \right) + \theta^{\text{NP}} \left(\phi + \frac{\pi}{2} \right) \phi. \quad (16)$$

The constants used for this test case are $\mu_1 = 2 \times 10^{-5}$, $\theta^{\text{SP}} = -40\varepsilon$, $\theta^{\text{EO}} = 30\varepsilon$, and $\theta^{\text{NP}} = -20\varepsilon$ where $\varepsilon = 1/300$. The initial moisture profile is set just below the saturation level so $q_v(\phi, \lambda) = 0.98q_{\text{sat}}(h, \theta)$, where $q_{\text{sat}}(h, \theta)$ is established from Eq. (A.11) in ZA, where $q_0 = 0.0492238$.

c. Uniform-resolution results: Moist barotropic instability test

Figure 9 depicts several fields for the barotropic wave at day 6 for the uniform high-resolution reference solution c2048 (≈ 5 km). This simulation shows that the development of the jet instability and the evolution of the initial vorticity rollups into sharp gradients (indicated as black solid and dashed contour lines in Figs. 9c and 9d) are consistent with the dry simulation results in Galewsky et al. (2004). The new aspects are that significant cloud formation q_c does not begin before day 4, and q_c does not precipitate (via the conversion to q_r) until five days into the simulation. By day six, the barotropic wave has created distinct vortices and thin vorticity filaments. Within these frontal and cutoff low-like features, areas of cloud and rain have formed. The temperature θ in Fig. 9a and water vapor q_v in Fig. 9b echo the structure of the contoured relative vorticity field (in Figs. 9c and 9d). The protrusions of colder and drier areas within the vorticity troughs mimic frontal systems in the midlatitudes. The q_c field is depicted in Fig. 9c, while Fig. 9d shows the total amount of rainwater q_r that has precipitated out over the preceding 12 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 mesh refinement on the barotropic instability’s vorticity field has been well covered for dry simulations [see St-Cyr et al. (2008), Weller et al. (2009), and Scott et al. (2016)]. Therefore, we focus our investigation on how the cloud q_c and precipitation of the physics scheme are affected by changing resolutions which serve as a reference for AMR runs. The q_c fields at day 6 for four other uniform resolutions, c128, c256, c512, and c1024, are depicted in Fig. 10 for comparison with the c2048 run q_c plot in Fig. 9c. Cloud cover area is fairly consistent across all resolutions in Fig. 10, with the exception of the c128 run. We interpret the two extra areas of cloud cover between 80° and 170° longitude

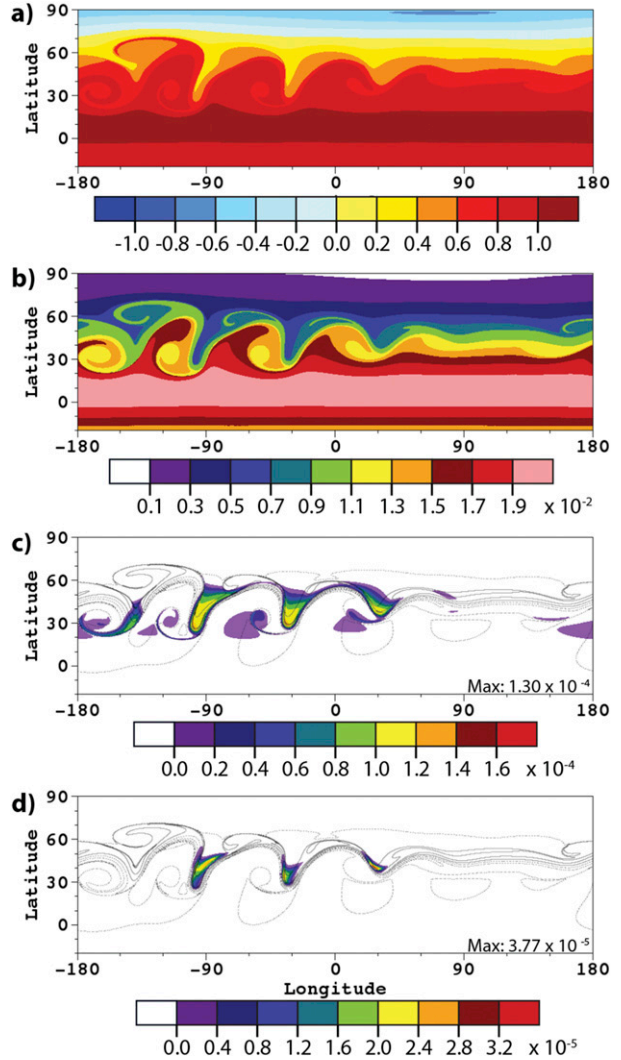


FIG. 9. Day 6 snapshots of the evolving barotropic wave for the c2048 uniform run’s (a) temperature field, (b) q_v moisture field, (c) q_c cloud field, and (d) past 12-h accumulation of the q_r precipitated water field. The solid and dashed black contour lines in (c),(d) represent the positive and negative relative vorticity, respectively. The spacing between contour lines is $5 \times 10^{-5} \text{ s}^{-1}$.

in Fig. 10a as artifacts of a wavenumber four created as the jet moves over the four corners of the cubed sphere. Similar difficulties for coarser resolutions runs on the cubed sphere were observed by St-Cyr et al. (2008) and Ullrich et al. (2010). While the overall shape and area of the cloud field converge at resolutions of c512 and higher, the concentration of the q_c field decreases with increasing resolution.

The 12-h accumulation of precipitated water q_r at day 6 for the four uniform resolutions is plotted in Fig. 11 and corresponds to Fig. 9d for the uniform c2048 run. We observe that the peak 12-h accumulation rate nearly doubles between the c256 run (Fig. 11b) and the c2048

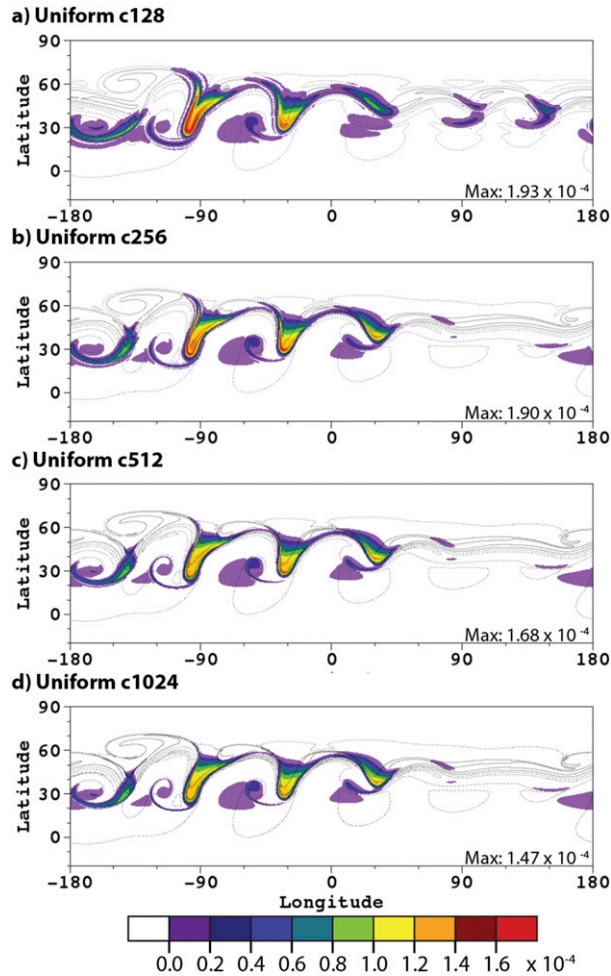


FIG. 10. Plots of the q_c cloud field at day 6 for several uniform resolutions: (a) c128, (b) c256, (c) c512, and (d) c1024. The c2048 uniform run plot of the same field in Fig. 9c serves as a reference. The solid and dashed black contour lines represent the positive and negative relative vorticity, respectively, using the same contour spacing as in Fig. 9.

run (Fig. 9d). 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. The location of the highest precipitation accumulation in the front-like system centered around -100° longitude shifts from a small area at the western edge of the bottom of the trough (as seen in Fig. 11c) to a broad area along the leading (eastern) edge and a secondary long and narrow area along the western edge (as seen in Figs. 11d and 9d).

d. AMR results: Moist barotropic instability test

For our implementation of AMR in this moist shallow-water system, we created two refinement tagging criteria: a relative vorticity threshold with

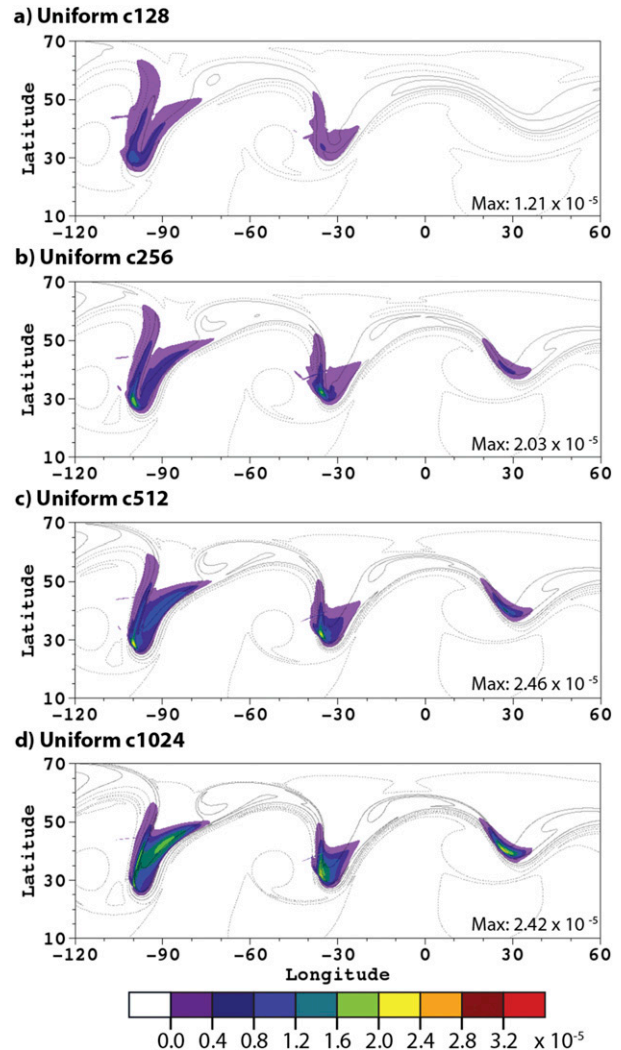


FIG. 11. Plots depicting the 12-h accumulation in the q_r precipitated water field for (a) c128, (b) c256, (c) c512, and (d) c1024 uniform runs. The c2048 uniform run plot of the same field in Fig. 9d serves as a reference. The solid and dashed black contour lines represent the positive and negative relative vorticity, respectively, using the same contour spacing as in Fig. 9.

$|\zeta| > 2.3 \times 10^{-5} \text{ s}^{-1}$ and a physics variable q_c -based threshold with $q_c > 3.0 \times 10^{-5}$. The magnitude of the vorticity threshold is roughly a quarter of the initial maximum absolute value of the relative vorticity. This ensures that the jet is sufficiently covered by refinement without adding excess refinement away from the main features. The magnitude of the cloud water threshold was selected to avoid refinement on low concentrations of cloud water that form around the equator throughout the simulation. The latter are triggered by gravity waves (due to the initial dynamical imbalance of the initial height perturbation, see the test case description) and having the initial water vapor values close to saturation.

Additionally, the vorticity tag triggers refinement from the start of the simulation over the entire jet, ensuring features of interest are well resolved throughout the simulation. The q_c tag is more stringent and does not trigger refinement until clouds form around day 4. It also only refines over more localized areas and uses significantly fewer computational resources. However, this leads to coarser resolutions over the developing instabilities. Four AMR runs with one level of $X = 4$ refinement are implemented. The vorticity tag is implemented within a c64 base-level run and a c256 run, while the q_c tag is implemented within a c128 base-level run and a c256 run.

Figure 12 depicts the q_c field at day 6 for four AMR runs while Fig. 13 depicts the 12-h q_r accumulation at day 6. The broader area of refinement created by the vorticity tagging criterion in the c64 base resolution AMR run reduces the grid imprinting typically seen with coarser resolutions. The cloud field (Fig. 12a) and 12-h accumulated precipitation (Fig. 13a) are comparable to those of the c256 uniform run (Figs. 10b and 11b). The c128 AMR run using the q_c tagging criterion demonstrates some of the drawbacks of the more stringent AMR criterion. Grid imprinting still occurs, with refinement actually being triggered over it (Fig. 12b). The three main areas of cloud concentration in Fig. 12b are, however, comparable to their counterparts in the uniform c512 run (Fig. 10c). The heaviest areas of precipitation in Fig. 11c match that of the c512 run as well, though the precipitation field has large areas of small-scale, but noisy, precipitation on the western sides of the two largest troughs.

The c64 vorticity tag AMR run and the c128 q_c tag AMR run have roughly the same computational cost. However, the c64 run has a higher computational burden for the first four days of the simulation, while the c128 run has a higher burden once refinement has been triggered after day 4. The c64 AMR run better captures the large global structure while the c128 AMR run better captures the localized structure of the cloud and rain distribution in the main troughs.

The higher base resolution in the two c256 AMR runs using the q_c tag in Fig. 12c and the vorticity tag in Fig. 12d eliminates the grid imprinting. For both runs, the q_c field visually converges to the uniform c1024 run and both runs capture the shifting of the precipitation maximum in the leftmost trough (Figs. 11c,d). The q_c tag run only has a base-resolution coverage of the rings of low concentration cloud along the edges of the main q_c areas in Fig. 12c. These features correspond to the thin weak cloud filaments that buttress the main filaments of clouds in the c128 and c256 uniform runs (Figs. 10a,b). They are present in the AMR runs

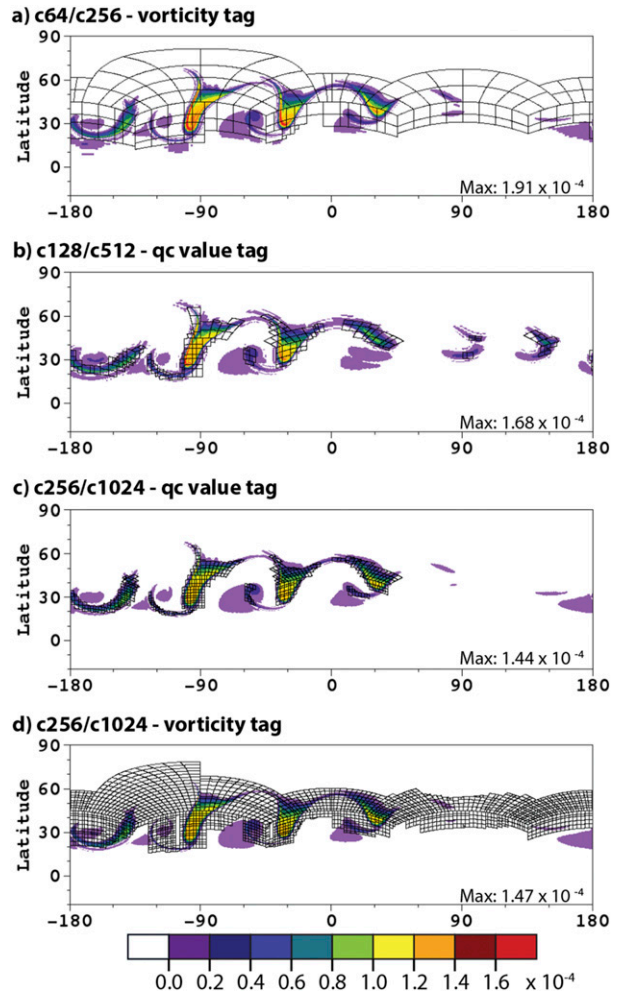


FIG. 12. The cloud q_c field profile at day 6 for several AMR runs with one level of $x4$ refinement. The tagging criterion for (a),(d) is a relative vorticity threshold of $|\zeta| > 2.3 \times 10^{-5} \text{ s}^{-1}$. The criterion for (b),(c) is $q_c > 3.0 \times 10^{-5}$. The block structures of the refinement levels are outlined in black.

because they lie along the coarse-fine grid boundary. The q_c tag runs also have areas of small-scale, noisy precipitation (Fig. 11c) seen in the c128 q_c tag run. Both these features are reduced significantly in the c256 vorticity tag AMR run.

These improvements come with significant computational cost. The c256 AMR vorticity tag run has roughly 4.5 times the computational cost as the q_c tag run, and by day 6 still has twice as many c1024 cells. These costs suggest that a higher base resolution that resolves the barotropic flow might make vorticity tagging unnecessary. The q_c tag is nearly as effective. Additionally, a lower cloud concentration threshold may reduce the noisy low-level edges by extending refinement out beyond the cloud formation areas. This can compensate the observed shortcomings of

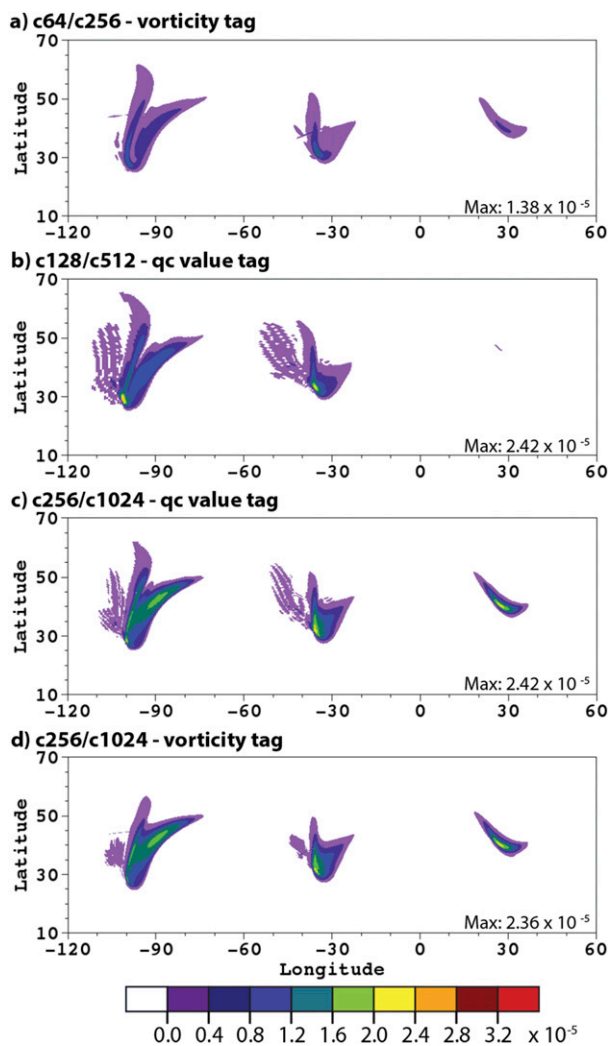


FIG. 13. Past 12-h accumulation of q_r at day 6 for the AMR runs depicted in Fig. 12.

the q_c tag without a significant increase in computational cost.

7. Conclusions

In this paper, we examined two different forcing frameworks to mimic the effects of atmospheric moisture within a 2D shallow-water system. These schemes add complexity to the shallow-water system and create dynamic features that can benefit from use of AMR. The first moist physics framework adds a water vapor variable and a mass sink triggered by saturation. We implemented a strengthening vortex test case with this setup. In the second forcing framework, the effects of moisture were coupled to the momentum equations through a temperature variable, linked to the moisture variables through latent heat. We used this setup with

the barotropic instability test case. The series of uniform simulations show that the evolution of the features of interest converges with increasing resolution for both test cases. Using the fourth-order finite-volume model Chombo AMR model, we implemented the two test cases in a series of uniform resolution and AMR simulations. These simulations can aid the establishment of guidelines for effective AMR refinement criteria.

With both forced shallow-water test cases, we investigated the effectiveness of different AMR strategies and AMR's effect on the physics forcing as grid resolutions changed. The forcing in both cases functioned effectively across the varying resolutions and multiple levels of AMR. With regards to refinement strategies, the sensitivity to 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. The changes in the 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.

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 did not strengthen significantly or underwent 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 or within the first day. AMR not added initially was still beneficial. The 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 test case simulations demonstrate that the starting grid resolution must be able to adequately resolve the features of interest to maximize the AMR effectiveness. AMR cannot correct errors caused before refinement begins. Additional refinement with AMR beyond that base level improved the model, particularly with regards to the small-scale vorticity features in the TC-like test case. To obtain 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. Refinement criteria also need to consider the computational cost. We used a higher vorticity threshold in Tag 3 for the vortex test case to

try to reduce number of grid cells but it still added significantly more refined blocks than the other simulations. The scaled refinement thresholds helped reduce computational cost but also delayed the use of the highest refinement levels until much later in the simulation. Adjusting the scale ratio, so that the vorticity threshold increases at less than the currently chosen one-to-one ratio could provide a compromise between adequate resolution and acceptable computational cost.

A more complex refinement option could include multiple criteria such as having the initial level of AMR be based on a large-scale feature like surface pressure or the wind field and have additional AMR levels trigger on more small-scale fields like sharp vorticity gradients or precipitation. An alternative technique is to use a combination of initial static refinement and AMR. For example, in tracking and resolving tropical cyclones in a realistic climate simulation, a static region of refinement could be placed over regions of cyclogenesis. The tropical 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 nonhydrostatic dynamical core, focusing on similar simplified physics parameterization schemes. This work has already started and initial results for a dry colliding modons test case (Lin et al. 2017) and the idealized moist tropical cyclone test case by Reed and Jablonowski (2011, 2012) are documented in Ferguson (2018).

Acknowledgments. Support for this work has been provided by the Office of Science, U.S. Department of Energy, Award DE-SC0003990, a University of Michigan (UM) Rackham Predoctoral Fellowship and a UM Catalyst grant from the Michigan Institute for Computational Discovery and Engineering (MICDE). We would like to acknowledge high-performance computing support from Cheyenne (<https://doi.org/10.5065/D6RX99HX>) provided by NCAR's Computational and Information Systems Laboratory, sponsored by the National Science Foundation. We thank the reviewers for their helpful comments and suggestions.

REFERENCES

- Adams, M., and Coauthors, 2015: Chombo software package for AMR applications—Design document. Tech. Rep. LBNL-6616E, Lawrence Berkeley National Laboratory, 204 pp.
- Betts, A. K., and M. J. Miller, 1986: A new convective adjustment scheme. Part II: Single column tests using GATE wave, BOMEX, and arctic air-mass data sets. *Quart. J. Roy. Meteor. Soc.*, **112**, 693–709, <https://doi.org/10.1002/qj.49711247308>.
- 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. *Phys. Fluids*, **21**, 116604, <https://doi.org/10.1063/1.3265970>.
- Enagonio, J., and M. T. Montgomery, 2001: Tropical cyclogenesis via convectively forced vortex Rossby waves in a shallow water primitive equation model. *J. Atmos. Sci.*, **58**, 685–706, [https://doi.org/10.1175/1520-0469\(2001\)058<0685:TCVCFV>2.0.CO;2](https://doi.org/10.1175/1520-0469(2001)058<0685:TCVCFV>2.0.CO;2).
- Ferguson, J. O., 2018: Bridging scales in 2- and 3-dimensional atmospheric modeling with adaptive mesh refinement. Ph.D. thesis, Department of Climate and Space Sciences and Engineering, University of Michigan, Ann Arbor, 168 pp.
- , 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. *Mon. Wea. Rev.*, **144**, 4641–4666, <https://doi.org/10.1175/MWR-D-16-0197.1>.
- Ferreira, R. N., W. H. Schubert, and J. J. Hack, 1996: Dynamical aspects of twin tropical cyclones associated with the Madden-Julian Oscillation. *J. Atmos. Sci.*, **53**, 929–945, [https://doi.org/10.1175/1520-0469\(1996\)053<0929:DAOTTC>2.0.CO;2](https://doi.org/10.1175/1520-0469(1996)053<0929:DAOTTC>2.0.CO;2).
- Frierson, D. M., A. J. Majda, and O. M. Pauluis, 2004: Large scale dynamics of precipitation fronts in the tropical atmosphere: A novel relaxation limit. *Commun. Math. Sci.*, **2**, 591–626, <https://doi.org/10.4310/CMS.2004.v2.n4.a3>.
- 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*, **56A**, 429–440, <https://doi.org/10.3402/tellusa.v56i5.14436>.
- Gill, A. E., 1982: Studies of moisture effects in simple atmospheric models: The stable case. *Geophys. Astrophys. Fluid Dyn.*, **19**, 119–152, <https://doi.org/10.1080/03091928208208950>.
- Goswami, P., and B. Goswami, 1991: Modification of $n = 0$ equatorial waves due to interaction between convection and dynamics. *J. Atmos. Sci.*, **48**, 2231–2244, [https://doi.org/10.1175/1520-0469\(1991\)048<2231:MOEWDT>2.0.CO;2](https://doi.org/10.1175/1520-0469(1991)048<2231:MOEWDT>2.0.CO;2).
- 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. *J. Atmos. Sci.*, **71**, 1623–1643, <https://doi.org/10.1175/JAS-D-13-0303.1>.
- Jablonowski, C., 2004: Adaptive grids in weather and climate modeling. Ph.D. thesis, Department of Atmospheric, Oceanic and Space Sciences, University of Michigan, Ann Arbor, 292 pp.
- Kessler, E., 1969: *On the Distribution and Continuity of Water Substance in Atmospheric Circulation*. Meteor. Monogr., No. 32, Amer. Meteor. Soc., 84 pp.
- Lahaye, N., and V. Zeitlin, 2016: Understanding instabilities of tropical cyclones and their evolution with a moist convective rotating shallow-water model. *J. Atmos. Sci.*, **73**, 505–523, <https://doi.org/10.1175/JAS-D-15-0115.1>.
- Lambaerts, J., G. Lapeyre, and V. Zeitlin, 2011: Moist versus dry barotropic instability in a shallow-water model of the atmosphere with moist convection. *J. Atmos. Sci.*, **68**, 1234–1252, <https://doi.org/10.1175/2011JAS3540.1>.
- 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. *J. Adv. Model. Earth Syst.*, **9**, 2483–2492, <https://doi.org/10.1002/2017MS000965>.
- McCorquodale, P., P. A. Ullrich, H. Johansen, and P. Colella, 2015: An adaptive multiblock high-order finite-volume method for solving the shallow-water equations on the sphere. *Commun. Appl. Math. Comput. Sci.*, **10**, 121–162, <https://doi.org/10.2140/camcos.2015.10.121>.

- Reed, K. A., and C. Jablonowski, 2011: An analytic vortex initialization technique for idealized tropical cyclone studies in AGCMs. *Mon. Wea. Rev.*, **139**, 689–710, <https://doi.org/10.1175/2010MWR3488.1>.
- , and —, 2012: Idealized tropical cyclone simulations of intermediate complexity: A test case for AGCMs. *J. Adv. Model. Earth Syst.*, **4**, M04001, <https://doi.org/10.1029/2011MS000099>.
- Ripa, P., 1993: Conservation laws for primitive equations models with inhomogeneous layers. *Geophys. Astrophys. Fluid Dyn.*, **70**, 85–111, <https://doi.org/10.1080/03091929308203588>.
- , 1995: On improving a one-layer ocean model with thermodynamics. *J. Fluid Mech.*, **303**, 169–201, <https://doi.org/10.1017/S0022112095004228>.
- 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. *Geophys. Astrophys. Fluid Dyn.*, **111**, 1–31, <https://doi.org/10.1080/03091929.2016.1269897>.
- Sadourny, R., 1972: Conservative finite-difference approximations of the primitive equations on quasi-uniform spherical grids. *Mon. Wea. Rev.*, **100**, 136–144, [https://doi.org/10.1175/1520-0493\(1972\)100<0136:CFAOTP>2.3.CO;2](https://doi.org/10.1175/1520-0493(1972)100<0136:CFAOTP>2.3.CO;2).
- Scott, R., L. Harris, and L. Polvani, 2016: A test case for the inviscid shallow-water equations on the sphere. *Quart. J. Roy. Meteor. Soc.*, **142**, 488–495, <https://doi.org/10.1002/qj.2667>.
- St-Cyr, A., C. Jablonowski, J. M. Dennis, H. M. Tufo, and S. J. Thomas, 2008: A comparison of two shallow water models with nonconforming adaptive grids: Classical tests. *Mon. Wea. Rev.*, **136**, 1898–1922, <https://doi.org/10.1175/2007MWR2108.1>.
- Stechmann, S. N., and A. J. Majda, 2006: The structure of precipitation fronts for finite relaxation time. *Theor. Comput. Fluid Dyn.*, **20**, 377–404, <https://doi.org/10.1007/s00162-006-0014-1>.
- Ullrich, P., C. Jablonowski, and B. van Leer, 2010: High-order finite-volume methods for the shallow-water equations on the sphere. *J. Comput. Phys.*, **229**, 6104–6134, <https://doi.org/10.1016/j.jcp.2010.04.044>.
- Weller, H., H. G. Weller, and A. Fournier, 2009: Voronoi, Delaunay, and block-structured mesh refinement for solution of the shallow-water equations on the sphere. *Mon. Wea. Rev.*, **137**, 4208–4224, <https://doi.org/10.1175/2009MWR2917.1>.
- Williamson, D. L., J. B. Drake, J. J. Hack, R. Jakob, and P. N. Swarztrauber, 1992: A standard test set for numerical approximations to the shallow water equations in spherical geometry. *J. Comput. Phys.*, **102**, 211–224, [https://doi.org/10.1016/S0021-9991\(05\)80016-6](https://doi.org/10.1016/S0021-9991(05)80016-6).
- Würsch, M., and G. C. Craig, 2014: A simple dynamical model of cumulus convection for data assimilation research. *Meteor. Z.*, **23**, 483–490, <https://doi.org/10.1127/0941-2948/2014/0492>.
- Yang, D., and A. P. Ingersoll, 2013: Triggered convection, gravity waves, and the MJO: A shallow-water model. *J. Atmos. Sci.*, **70**, 2476–2486, <https://doi.org/10.1175/JAS-D-12-0255.1>.
- Zerroukat, M., and T. Allen, 2015: A moist Boussinesq shallow water equations set for testing atmospheric models. *J. Comput. Phys.*, **290**, 55–72, <https://doi.org/10.1016/j.jcp.2015.02.011>.