Accelerating Simulations of Tropical Cyclone Rapid Intensification Using Adaptive Mesh Refinement

YASSINE TISSAOUI,[a] STEPHEN R. GUIMOND,[b] FRANCIS X. GIRALDO,[c] AND SIMONE MARRAS[d]

[a] *Department of Mechanical and Industrial Engineering, New Jersey Institute of Technology, Newark, New Jersey*
[b] *Department of Atmospheric and Planetary Sciences, Hampton University, Hampton, Virginia*
[c] *Department of Applied Mathematics, Naval Postgraduate School, Monterey, California*
[d] *Department of Mechanical and Industrial Engineering, and Center for Applied Mathematics and Statistics, New Jersey Institute of Technology, Newark, New Jersey*

(Manuscript received 14 November 2024, in final form 17 October 2025, accepted 28 October 2025)

ABSTRACT: Tropical cyclones (TCs) are powerful, natural phenomena that can severely impact populations and infrastructure. Enhancing our understanding of the mechanisms driving their intensification is crucial for mitigating these impacts. To this end, researchers are pushing the boundaries of TC simulation resolution down to scales of just a few meters. However, high-resolution simulations come with significant computational challenges, increasing both time and energy costs. Adaptive mesh refinement (AMR) is a technique widely used in computational fluid dynamics but has seen limited application in atmospheric simulations. This study explores the use of *h*-adaptive spectral elements to accelerate simulations of TC rapid intensification while allowing very high resolutions in certain parts of the domain. By applying AMR to a rapidly intensifying TC test case, we demonstrate that AMR can replicate the results of uniform grid simulations in terms of mean and local wind speed maxima while dramatically reducing computational costs. We show that AMR can speed up dry simulations of rapidly intensifying TCs forced by heating observations by a factor of 2–13 for the set of tested refinement criteria. Additionally, we show that TC intensity changes as resolution is increased and that AMR can deliver high-resolution simulations at the cost of coarser static simulations. Our findings indicate that AMR and spectral element methods are promising tools for enhancing TC simulations.

SIGNIFICANCE STATEMENT: Each year, tropical cyclones are among the most deadly and costliest of all natural disasters. These storms intensify through the consolidation of the energy released inside thunderstorms, which contain a hierarchy of small spatial scales. To simulate the intensification process accurately with mathematical models, it is necessary to account for all these small spatial scales, which can be difficult even on modern computing systems. In this work, we present a numerical algorithm called adaptive mesh refinement (AMR) that can alleviate some of this computational burden while still capturing the essence of the thunderstorms. The AMR algorithm shown here produces faster simulations by up to a factor of 10, which will significantly enhance the ability to capture tropical cyclone (TC) intensification.

KEYWORDS: Hurricanes/typhoons; Adaptive models; Idealized models; Nonhydrostatic models; Numerical weather prediction/forecasting

1. Introduction

Atmospheric phenomena can be challenging to simulate with mathematical models, given the large range of scales that is present, which span several orders of magnitude, and the strong nonlinear interactions between those scales (Klein et al. 2010; Klein 2010). Tropical cyclones (TCs) are at the top of the list of challenging phenomena, given their extremely large Reynolds number ($\sim 10^{10}$) which requires large domains with high-resolution and high-order discretization to accurately capture the vast array of interactions controlling the system intensity. Physical processes such as (i) vortex Rossby waves (Guimond et al. 2010; Guinn and Schubert 1993; Wang 2002; Montgomery and Kallenbach 1997), (ii) mesovortices at the eye/eyewall interface (Schubert et al. 1999; Kossin and Schubert 2001; Hendricks et al. 2009; Guimond et al. 2016b; Cram et al. 2007), (iii) boundary layer coherent turbulent structures (Guimond et al. 2018; Foster 2005), (iv) air–sea turbulent fluxes (Wu et al. 2005; Emanuel 1986), and (v) deep convective bursts (Rogers et al. 2013; Molinari and Vollaro 2010; Hendricks et al. 2004; Guimond et al. 2010) all play a role in TC intensification and require high fidelity numerical treatment to accurately capture their energetics (Hasan et al. 2022). Increasing the model resolution and numerical accuracy are the primary means to reduce both the model error from parameterization and the error from excessive diffusion (Johnson et al. 2024). The computational burden of increasing the model resolution and numerical accuracy for explicitly simulating large turbulent eddies in TCs is substantial, even with the large core counts and optimized hardware of modern computing systems.

Scientists and engineers have developed various techniques to alleviate the computational burden of high-resolution simulations of atmospheric phenomena including 1) grid nesting and low-order discretizations (e.g., Kurihara et al. 1979, 1998;

Corresponding author: Yassine Tissaoui, tissaoui@wisc.edu

DOI: 10.1175/JAS-D-24-0242.1

Hodur 1997; Skamarock et al. 2005; Doyle et al. 2014), 2) static grid stretching (e.g., Guimond and Reisner 2012), and 3) adaptive mesh refinement (AMR) (e.g., Skamarock et al. 1989; Kopera and Giraldo 2014b; Abdi et al. 2024). Each of these techniques has positive and negative attributes. For example, static grid stretching is a relatively simple method to implement in a code, but it does not always reduce the computational cost effectively. Grid nesting has been effective for zooming into a smaller region with higher resolution while leaving surrounding areas with coarser resolution, but noise at the nest boundaries and some difficulty following features of interest are apparent (Ferguson et al. 2016). The AMR method works by dynamically adjusting the grid resolution in a series of refinements according to a set of predefined criteria such as the wind speed or vorticity. In this case, the grid becomes adaptive by refining and coarsening as targeted flow features appear and vanish over time.

The early works of Berger and Oliger (1984), Skamarock et al. (1989), and Skamarock and Klemp (1993) showed that adaptive refinement is useful for solving hyperbolic equations and problems in numerical weather prediction. Bacon et al. (2000) developed the first operational model making use of AMR to generate horizontally adaptive grids. This model was successful in simulating hurricane storm tracks. AMR has been shown to be effective for simulations of the shallow-water equations (Berger et al. 2009; LeVeque et al. 2011; McCorquodale et al. 2015), and it has been tested with high-order Galerkin methods for atmospheric applications (Mueller et al. 2013; Kopera and Giraldo 2014a; Chen et al. 2011). Recent applications of AMR led to the improvement of modeling atmospheric flows around topography (Orlando et al. 2024; Yamazaki et al. 2022; Li et al. 2021).

In terms of TC simulations, Hendricks et al. (2016) demonstrated that AMR can speed up highly idealized simulations of TCs by a factor of 4–15 using a shallow-water spectral element model. The goal of the current paper is to extend these results to realistic, three-dimensional simulations of TC rapid intensification using a similar spectral element model that provides high-order accurate discretizations of the governing equations. The authors will quantify the computational advantage of AMR over uniform resolution simulations. These numerical methods are applied to the problem of TC rapid intensification, which is characterized by a strong increase (approximately 15 m s^{-1} or more) in the maximum sustained winds over a period of 24 h. Rapid intensification is often associated with the pulsing of deep convective towers for ~6–12-h time periods (Guimond et al. 2010, 2016a). The thermal forcing associated with these convective towers generates a spectrum of smaller-scale motions that serve as an excellent test case for the AMR algorithm.

2. Model equations

We consider a dry atmosphere of density ρ, pressure p, and potential temperature $\theta = T/\pi$, where T is the sensible temperature and π is the exner pressure. Taking the ideal gas constant of dry air to be R_d and its specific heat capacity at constant pressure to be c_p, the exner pressure is defined as $\pi = (p/p_0)^{R_d/c_p}$, where p_0 is the sea surface pressure. Let $\mathbf{u}$ be the wind velocity vector. We consider a fixed spatial domain Ω and a time interval of interest $(0, t_f]$. We write the compressible Euler equations in nonconservative form as follows:

$$\frac{\partial \rho}{\partial t} + \nabla\cdot(\rho\mathbf{u}) = 0 \qquad \text{in } \Omega \times (0, t_f], \tag{1}$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u} = -\frac{1}{\rho}\nabla p + g\mathbf{k} - f\mathbf{u}_h + \nabla\cdot\boldsymbol{\tau} + S_u \qquad \text{in } \Omega \times (0, t_f], \tag{2}$$

$$\frac{\partial \theta}{\partial t} + \mathbf{u}\cdot\nabla\theta = \kappa\nabla^2\theta + S_\theta \qquad \text{in } \Omega \times (0, t_f], \tag{3}$$

where $g = 9.81$ m s^{-2} is the magnitude of the acceleration due to gravity; $\mathbf{k} = [0, 0, -1]$ is the unit vector aligned with the vertical axis; f is the Coriolis parameter; $\mathbf{u}_h$ is the horizontal velocity; S_u and S_θ represent the sources and sinks for velocity and potential temperature, respectively; τ is the horizontal turbulent stress tensor which is built using the Smagorinsky–Lilly model described in section 5b(3); and κ is a turbulent diffusivity defined using the Smagorinsky–Lilly model. The source and sink terms account for the observational heating, the vertical turbulent subgrid diffusion, and the Rayleigh damping which are all described in section 5. The ideal gas law applied to dry air is used to close the system:

$$p = \rho R_d T. \tag{4}$$

To facilitate the numerical solution of the system (1)–(3), we write density, pressure, and potential temperature as the sum of their mean hydrostatic values and fluctuations:

$$\rho(x, y, z, t) = \rho_0(z) + \rho'(x, y, z, t), \tag{5}$$

$$\theta(x, y, z, t) = \theta_0(z) + \theta'(x, y, z, t), \tag{6}$$

$$p(x, y, z, t) = p_0(z) + p'(x, y, z, t). \tag{7}$$

Note that the hydrostatic reference states are the functions of the vertical coordinate z only. Hydrostatic balance relates from p_0 to ρ_0 as follows:

$$\frac{dp_0}{dz} = -\rho_0 g. \tag{8}$$

Plugging (5)–(7) into (1)–(3) and accounting for (8) leads to

$$\frac{\partial \rho'}{\partial t} + \nabla\cdot[(\rho_0 + \rho')\mathbf{u}] = 0 \qquad \text{in } \Omega \times (0, t_f], \tag{9}$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u} = -\frac{1}{\rho_0 + \rho'}\nabla p' + \frac{\rho'}{\rho_0 + \rho'}g\mathbf{k} + \nabla\cdot\boldsymbol{\tau} \qquad \text{in } \Omega \times (0, t_f], \tag{10}$$

$$\frac{\partial \theta'}{\partial t} + \mathbf{u}\cdot\nabla\theta_0 + \mathbf{u}\cdot\nabla\theta' = \kappa\nabla^2\theta \qquad \text{in } \Omega \times (0, t_f]. \tag{11}$$

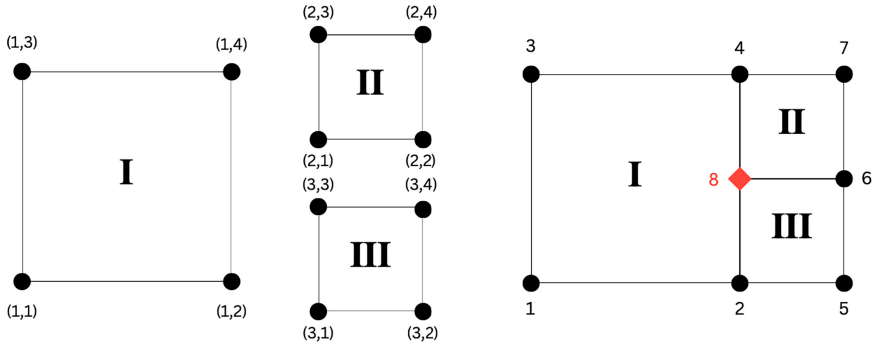

FIG. 1. Illustration of nonconforming linear elements in (left) local view (right) and global view. While node number 8 is shared by elements II and III corresponding to local nodes (2, 1) and (3, 3), respectively, it does not have a corresponding local node on element I. Node 8 is thus a hanging node.

3. The spectral element method

The continuous Galerkin (CG) spectral element method is used to discretize the governing equations in space. The three-dimensional (3D) basis functions for these elements are constructed using tensor products of the one-dimensional Lagrange-interpolating polynomials h of order N:

$$\psi_l(\mathbf{x}) = h_i[\xi(\mathbf{x})] \otimes h_j[\eta(\mathbf{x})] \otimes h_k[\zeta(\mathbf{x})],$$
$$l = i + (N+1)[j + k(N+1)], \tag{12}$$

where $\mathbf{x} = (x, y, z)$ and ξ, η, and ζ are the mappings from the physical coordinate $\mathbf{x}$ onto each coordinate of the reference element. The Legendre–Gauss–Lobatto points are used as both the interpolation and integration points. We use spectral elements of order 4. This is a sufficiently high order for inexact integration to be accurate and allows for important computational cost savings. For details on the CG method, we refer the interested reader to Kopriva (2008) and Giraldo (2020).

Considerations for adaptive mesh refinement

The use of AMR leads to the presence of nonconforming elements. That is to say that it is possible to obtain neighboring elements that do not necessarily have matching nodes along the face or edge that they share. We use linear elements to illustrate this in Fig. 1 where node 8 in the global view (right) corresponds to the local nodes (2, 1) and (3, 3) of elements II and III, respectively. However, this node has no corresponding local node on element I and is called a hanging node.

The global solution of the CG method is constructed via direct stiffness summation (DSS) which constructs the solution at each global node based on the values of the solution at the corresponding local nodes of each element that shares the global node. A modified DSS approach must be used to appropriately construct the global solution when there are hanging nodes. Essentially, the values of the local solution at hanging nodes contribute to the global solution through interpolation onto nonhanging nodes. The mass matrix is also constructed through this approach which insures the correctness of the DSS operation. Once a global solution is obtained, it is then interpolated onto the hanging nodes to construct the correct local solutions. We refer the reader to section d of Kopera and Giraldo (2014b) for algorithmic details.

4. AMR

We rely on the "P4est" (Burstedde et al. 2011) library to handle the grid adaptivity, domain partitioning, and load balancing. This library also allows for refining and coarsening grids, and we make use of this capability to perform adaptive mesh refinement during our simulations.

a. The AMR procedure

Thanks to P4est, the AMR procedure is straightforward to perform. The process only requires that P4est knows if an element should be refined, coarsened, or left untouched. For this to take place, an array adapt of size N_e, the total number of elements, is sent to P4est. The values of this array for each element e are as follows:

$$\text{adapt}(e) = \begin{cases} -1 & \text{If the element is to be marked for coarsening,} \\ 0 & \text{If the element is to be remain as is,} \\ 1 & \text{If the element is to be marked for refinement.} \end{cases} \tag{13}$$

1) REFINING AN ELEMENT

In this work, we are only considering adaptive mesh refinement in the horizontal directions. As such, if a hexahedral element is marked for refinement, it is split into four elements of equal size in the horizontal plane. Since we depend on column data structures for the planetary boundary layer package being used in our simulations, if an element belonging to a column is refined, the entire column must be refined with it to maintain this structure.

2) COARSENING AN ELEMENT

If the refinement process involves splitting an element into smaller elements, then the coarsening process is the opposite. Four neighboring elements that are marked for coarsening are merged together to form a single larger element. As such,

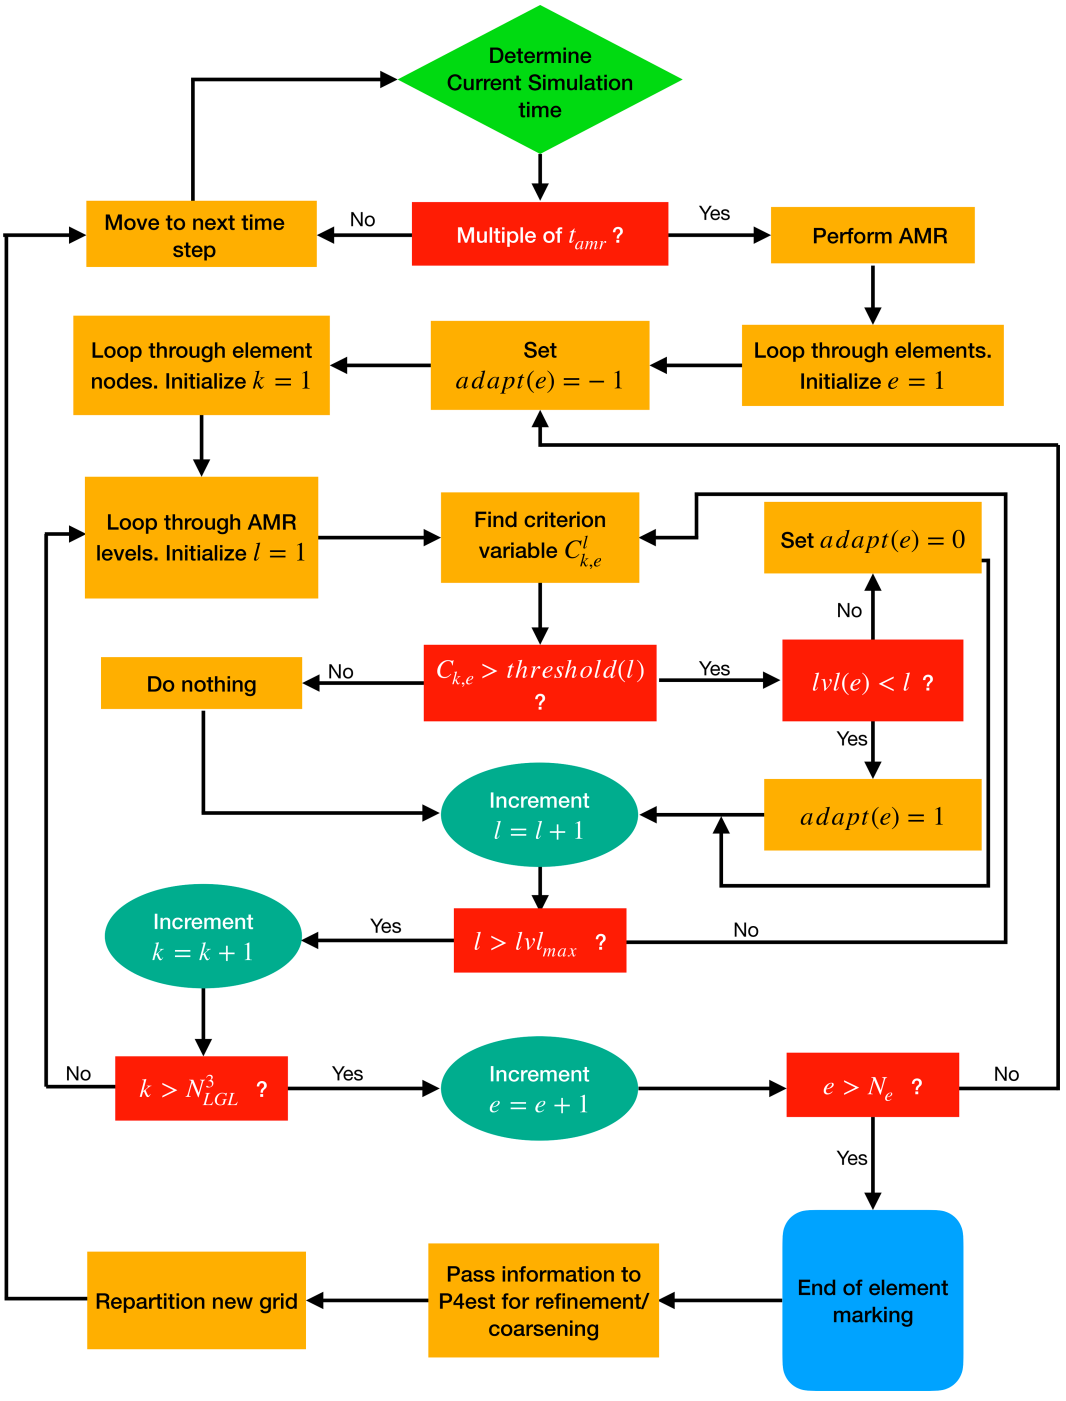

FIG. 2. Workflow of the AMR procedure. The start of the flow diagram is shown in the green shape labeled "determine current simulation time."

coarsening cannot take place unless there are four adjacent elements that all share a corner.

3) CONDITIONS FOR REFINEMENT AND COARSENING

P4est also stores the current level of refinement of each element in an array we will call lvl. Initially, all elements have the level $\text{lvl}(e) = 0$, indicating no refinement has taken place. This is also the maximum size of a given element, meaning that it cannot be coarsened if $\text{lvl}(e)$ is not strictly positive. In other words, an element cannot be coarsened if it is currently at the maximum size. An element also cannot be further refined if it is at the maximum allowed level of refinement lvl_{max}.

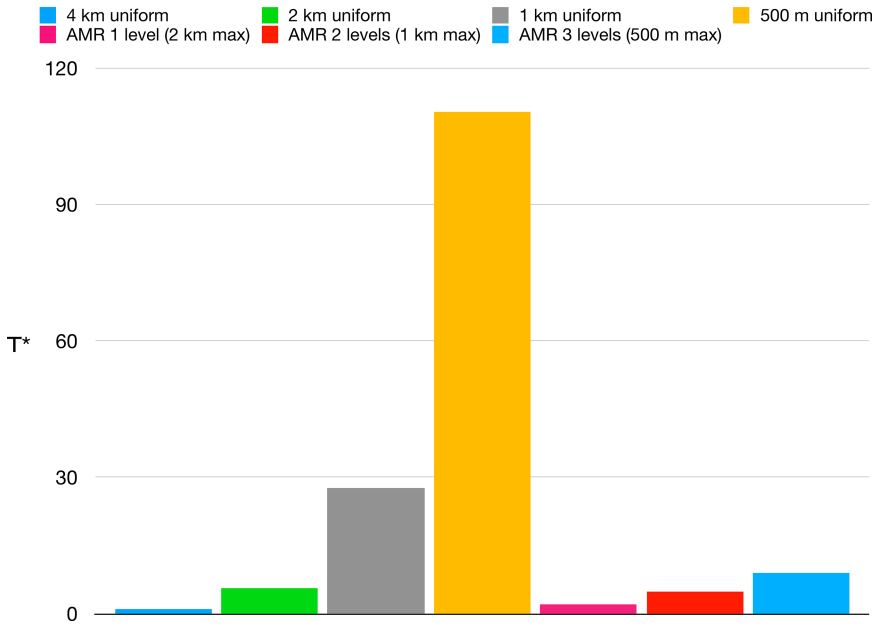

FIG. 3. Time-to-solution comparison of uniform and AMR simulations in normalized time for the first set of tests. The time to solution for the 500-m uniform case is an estimate based on the time to solution for the 1-km uniform case.

A criterion is set for each level of refinement. If one of the nodes belonging to an element satisfies the refinement criterion and $\mathrm{lvl}(e) < \mathrm{lvl}_{\max}$, then it is marked for refinement. If an element no longer satisfies the refinement criterion for its current level of refinement, then it is marked for coarsening.

4) THE REFINEMENT/COARSENING CRITERION

Consider that the refinement level l depends on the value of a flow variable C^l, and consider $C^l_{k,e}$ the value of this variable at the kth node of element e. We consider a threshold-type criterion for refinement. This means that if any node k belonging to element e satisfies $C_{k,e} > \mathrm{threshold}[\mathrm{lvl}(e) + 1]$, then the element is marked for refinement. It also means that if $\forall k\ C_{k,e} \leq \mathrm{threshold}[\mathrm{lvl}(e)]$, then the element is marked for coarsening.

5) FREQUENCY OF THE ADAPTIVE MESH REFINEMENT PROCEDURE

Because AMR is not without cost, it should not be performed at every simulation time step. We define a time interval t_{amr} that is a multiple of the time step Δt and allow the AMR procedure to take place at every instance that the current simulation time is a multiple of t_{amr}. This interval should be large enough to avoid needlessly executing the AMR procedure (a needless AMR procedure would be one where no elements are coarsened or refined) and small enough to be able to adapt to substantial changes in the flow.

Remark 4.1: An element can only be refined once per AMR iteration, even if it satisfies the criterion for a higher level of refinement. The same applies for coarsening. The level of a given element can only increase or decrease by a value of one at the most every time the AMR procedure is executed.

b. *The AMR algorithm and workflow*

We present the reader with graphical representation of the adaptive mesh refinement process in Fig. 2. This figure is a workflow diagram that illustrates the sequence of operations taking place, where t_{final} is the final time of a given simulation, threshold is an array storing the threshold criteria for each level of refinement, N_{LGL} is the number of Legendre–Gauss–Lobatto (LGL) points, N_e is the number of elements, and mod is the remainder operator. We also provide a pseudocode for the AMR procedure in the appendix.

5. Simulations and results

a. *Summary of results*

This section presents the results of simulations of idealized dry rapidly intensifying tropical cyclones on static and adaptive grids. The simulations performed with AMR compare well both qualitatively and quantitatively with simulations performed on static grids, while delivering speedups of up to a factor of 13. The speed and accuracy of AMR simulations are also shown to depend on the choice of refinement criterion and the duration of the simulation, with AMR being able to reproduce the results of shorter simulations with higher fidelity.

b. *Model setup*

The Nonhydrostatic Unified Model of the Atmosphere (NUMA) model (Giraldo et al. 2013) is used for all simulations. The initial conditions for this test are similar to those in Guimond et al. (2016a). A dry tropical storm–like vortex is

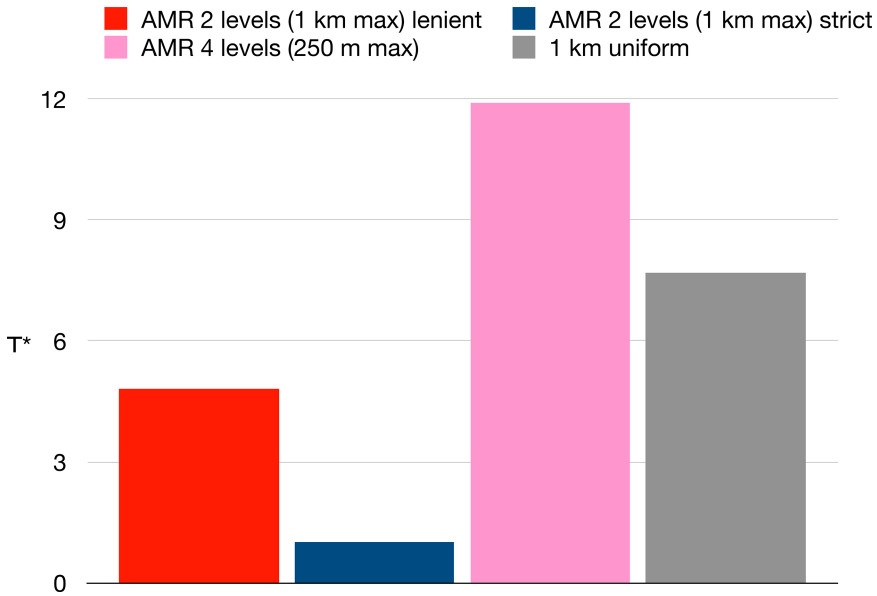

FIG. 4. Time-to-solution comparison of uniform (static) and AMR simulations in normalized time for the second set of tests. The strict two-level AMR uses threshold$_{st}$, while the lenient two-level AMR uses threshold$_{1n}$.

initialized using the following profile for azimuthal-mean tangential velocity:

$$\overline{v}(r, z) = V(r)\exp\left[-\frac{z^{\sigma}}{\sigma D_1^{\sigma}}\right]\exp\left[-\left(\frac{r}{D_2}\right)^6\right], \tag{14}$$

where V is the surface tangential velocity, $\sigma = 2$, $D_1 = 5823$ m, and $D_2 = 200$ km. The surface tangential velocity can be found by following the procedure described by Nolan and Grasso (2003) and Nolan et al. (2007) and integrating a specified Gaussian distribution with a vorticity peak of 1.5×10^{-3} s^{-1} and maximum winds of 21.5 m s^{-1} at a radius of 50 km from the center. The vertical velocity is initially taken to be zero everywhere.

The density and potential temperature are initialized by an iterative procedure that oscillates between satisfying the gradient wind balance and the hydrostatic balance until a specified criterion is met. This procedure is described by Nolan (2011). The background state is defined by vertically interpolating the Jordan (1958) mean hurricane-season sounding onto the spectral element grid.

1) OBSERVATIONAL HEATING

The time evolution of the vortex is driven by a four-dimensional source term in the energy equation. This source term represents latent heating/cooling rates in convective clouds derived from Doppler radar measurements in Hurricane Guillermo (1997). The algorithm and observations of latent heating/cooling rates are described by Guimond et al. (2011), and its implementation in NUMA is described by Hasan et al. (2022). Hurricane Guillermo (1997) was a rapidly intensifying TC, and the observational heating data should provide an excellent testing ground for how AMR responds to the small-scale variability inherent in this complex system. As described by Guimond et al. (2011), the heating is computed on a grid covering the inner core of the system out to a radius of $r = 60$ km from the domain center. This grid has a resolution of 2 km in the horizontal direction and 0.5 km in the vertical direction. The heating observations are split into 10 snapshots covering a 5.7-h period in intervals of 34 min. These numbers are based on the aircraft revisit times and overall sampling time period. The largest heating rates are present at a radius of 25–30 km from the domain center, well within the radius of maximum winds for the initial conditions. Outside of $r = 60$-km radius, the heating term is zero. The heating rates are interpolated in space onto the spectral element grid. After initialization, the first heating snapshot is gradually introduced over the first 30 min of the simulation, by the way of a hyperbolic tangent function. The snapshots are then linearly interpolated to the next observation time over the course of the remainder of the simulation. Past $t = 5.7$ h, the heating is maintained at the last snapshot until the simulation finishes at $t = 6$ h.

2) BOUNDARY CONDITIONS

At the lower boundary, the surface layer model by Jiménez et al. (2012) is active, while the sea surface temperature is maintained constant at its reference value. The lateral boundaries are doubly periodic, and a Rayleigh absorbing layer is used to damp gravity waves at the domain top and is described by

$$R(z) = \min\left(1, \max\left\{0, \gamma\sin^2\left[\frac{\pi}{2}\left(1 - \frac{z_{\text{top}-z}}{z_d}\right)\right]\right\}\right), \tag{15}$$

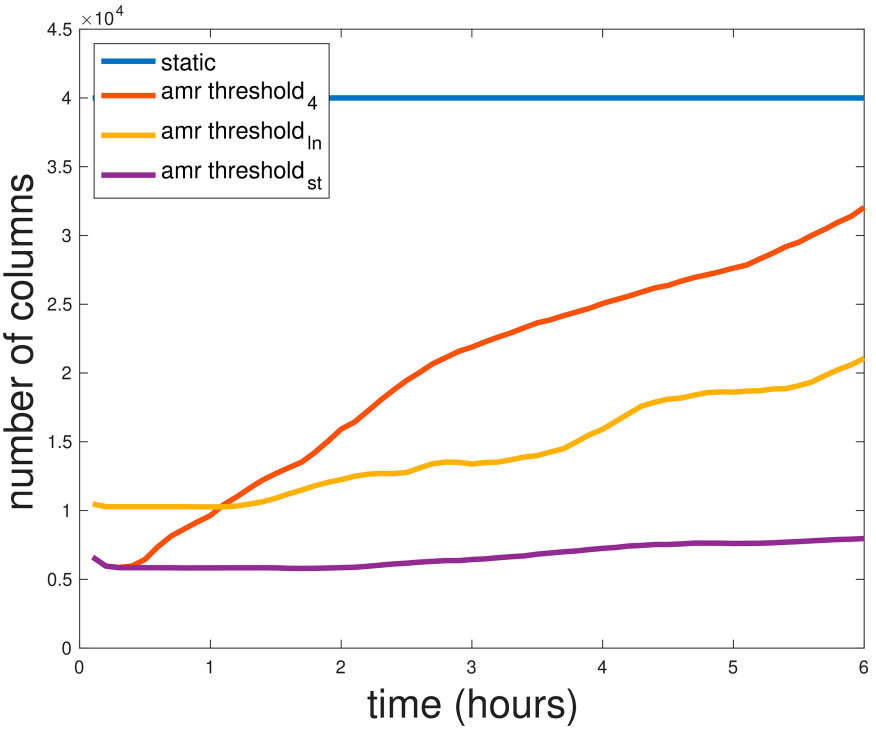

FIG. 5. Number of element columns over time for the uniform (static) grid simulation (blue), the four-level AMR simulation (red) with threshold$_4$ = [7.5, 15, 22.5, 30] m s^{-1}, the two-level AMR simulation with threshold$_{1n}$ = [2, 5] m s^{-1} (yellow), and the two-level AMR simulation with threshold$_{st}$ = [7.5, 15] m s^{-1} (purple).

where $\gamma = 1.0$, $z_{top} = 20$ km, and $z_d = 4$ km. The term $q'R(z)$ is added to the right-hand side of each equation where $q' = \rho'$ for (9), $q' = \mathbf{u}$ for (10), and $q' = \theta'$ for (6).

The lower boundary uses the surface layer scheme described by Jiménez et al. (2012) along with a free-slip boundary condition and a fixed sea surface temperature $T_{surf} = 299.95$ K.

3) SUBGRID MODELS

The Smagorinsky–Lilly model (Smagorinsky 1963; Lilly 1962) is used to model subgrid-scale turbulence in the horizontal direction. This model is a turbulent viscosity model and contributes to the right-hand side through the divergence of the turbulent stress tensor $\nabla \boldsymbol{\tau}$. The turbulent stress tensor is defined as

$$\boldsymbol{\tau} = -2\nu_t \mathbf{S}, \tag{16}$$

where $\mathbf{S} = (1/2)[\nabla\mathbf{u} + (\nabla\mathbf{u})^{\mathrm{T}}]$ is the strain rate tensor, and ν_t is the turbulent viscosity defined as

$$\nu_t = (\Delta C_s)^2 \sqrt{2|\mathbf{S}|^2}, \tag{17}$$

where C_s is the constant Smagorinsky coefficient and is taken as $C_s = 0.21$ for our simulations. The Δ symbol is the filter width of the Smagorinsky model and is taken to be the mean horizontal resolution of a given element, meaning that larger values of viscosity will be present for coarser elements. The turbulent viscosity and the turbulent Prandtl number are used to define the turbulent diffusivity $\kappa = \mathrm{Pr}\nu_t$ which is present in the potential temperature equation.

The vertical turbulent subgrid diffusion is modeled by the planetary boundary layer (PBL) scheme of Hong (2010) and Hong et al. (2006), commonly used for TC modeling. The scheme is described by

$$\frac{\partial C}{\partial t} = \frac{\partial}{\partial z}\left[K_c\left(\frac{\partial C}{\partial z} - \gamma_c\right) - \overline{(w'c')}_h\left(\frac{z}{h^3}\right)\right], \tag{18}$$

where C is a prognostic variable, K_c is the eddy diffusivity coefficient, $\overline{w'c'}_h$ is the flux at the inversion layer, and γ_c is a correction to the local gradient. For a more detailed description of the PBL scheme, we refer the reader to Hong (2010) and Hong et al. (2006).

4) COMPUTATIONAL DOMAIN

All of the simulations are run on a domain with size [−400, 400] km × [−400, 400] km × [0, 20] km. Continuous Galerkin spectral elements of order 4 are used. Note that all figures in the paper are only showing a subset of the full model domain to make details more visible.

5) TIME INTEGRATION

All simulations use an explicit third-order Runge–Kutta scheme for time integration. The time step is maintained and

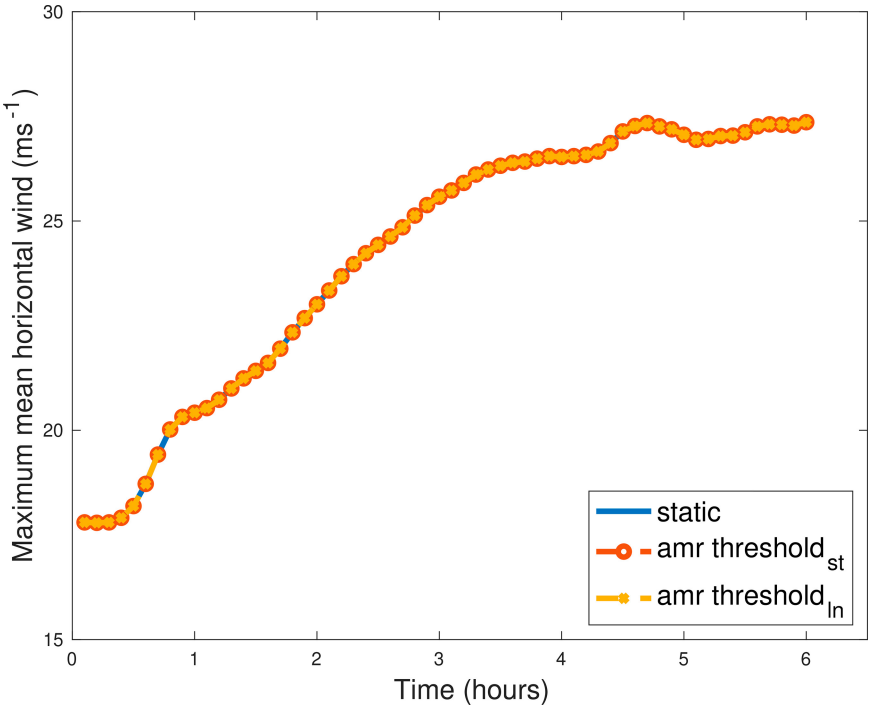

FIG. 6. Maximum of horizontally averaged horizontal velocity within the $[-50, 50]$ km^2 subdomain over the course of the 6-h simulation period. These values are found by calculating the average value of horizontal velocity at each vertical level and then finding the maximum among these values. The blue solid line (uniform/static simulation) overlaps perfectly with the dashed line with red circles (stricter criterion AMR simulation) and the dashed line with yellow crosses (more lenient criterion AMR simulation). The largest differences between the static simulation and the AMR simulations are ~0.002 m s^{-1}.

is the same in all of the tests that we perform with the exception of a single AMR test with four levels of refinement; because of a higher vertical resolution for this test, a smaller time step is required. The value of Δt will be clearly indicated in the test descriptions below.

c. *Time-to-solution comparison*

A set of 1-h-long simulations is performed to compare the time to solution of simulations with and without AMR at varying resolutions. The vertical resolution is maintained constant at $\Delta z = 312$ m. A simulation using a uniform horizontal resolution of $\Delta x = \Delta y = 4$ km is used as the baseline, and its time to solution $T_{4\text{km}}$ is used to obtain a normalized time to solution $T^* = T/T_{4\text{km}}$, where T is the time to solution for a given simulation.

Six simulations in total are performed for this comparison. Three constant grid spacing simulations are performed with $\Delta x = \Delta y = 4, 2, 1$ km, respectively. Three simulations using AMR are performed with respective minimum horizontal grid spacings of $\Delta x = \Delta y = 2, 1, 0, 5$ km, corresponding to one, two, and three levels of refinement. The criterion for refinement in the AMR simulations is that the velocity magnitude must pass a set of predefined thresholds threshold$_k$, where $k = 1$, lvl$_{\text{max}}$, and lvl$_{\text{max}}$ is the maximum refinement level. For these tests, we perform tests with the following thresholds: threshold$_1$ = [7.5] m s^{-1}, threshold$_2$ = [7.5, 15] m s^{-1}, and threshold$_3$ = [7.5, 15, 22.5] m s^{-1}. All these simulations are performed on the same machine with 1536 cores, and we will refer to these tests as the first set in the remainder of the paper. The normalized time to solution for these tests is presented in Fig. 3. For the static uniform grids, the increase in horizontal resolution results in ~4 times increase in time to solution from 4 to 2 km and then similarly from 2 to 1 km. This is expected as the number of grid cells quadruples when doubling the horizontal resolution in a 3D simulation. We use a scaling factor of 4 to present an estimated time to solution for a uniform grid with a horizontal resolution $\Delta x = 500$ m on the same figure (yellow histogram). With the set of criteria being used, the time to solution for the AMR simulations with up to 1-km horizontal resolution remains lower than a uniform 2-km horizontal resolution simulation. Additional refinement levels result in a ~2 times increase in cost for the AMR simulations with the criteria being used in these tests. The AMR simulation with a horizontal resolution of up to 500 m is still much cheaper than a uniform 1-km resolution simulation and barely more expensive than a uniform 2-km resolution simulation. Compared to the estimated cost of a uniform $\Delta x = 500$-m simulation, the AMR simulation with a resolution of up to $\Delta x = 500$ m is nearly 13 times faster. The speedup obtained for these simulations is obtained while using a single explicit time stepper for all processes with the

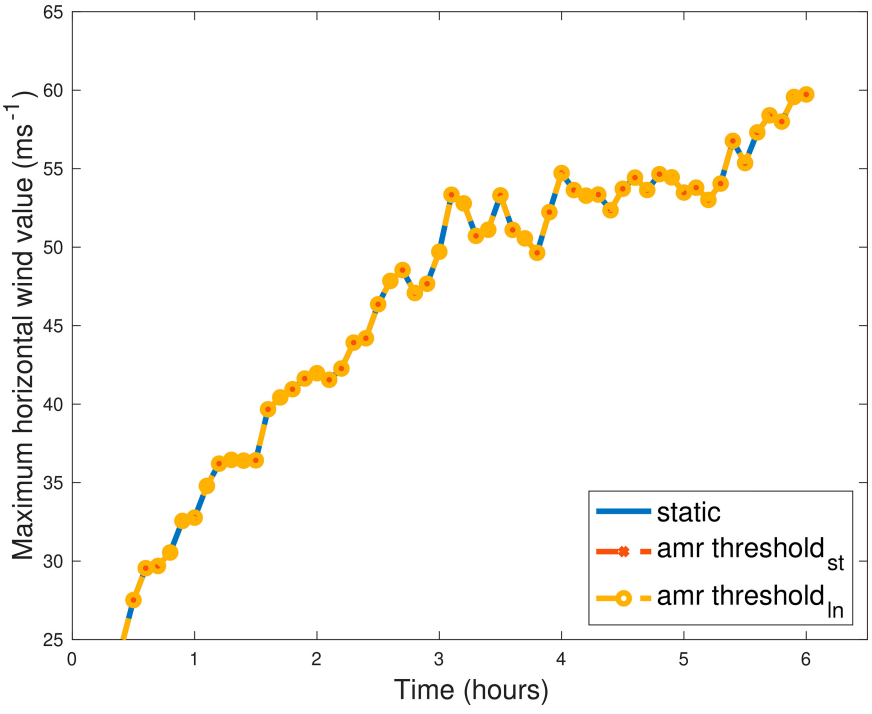

FIG. 7. Maximum value of horizontal velocity within the [−50, 50] km × [−50, 50] km × [0, 20] km subdomain over the course of the 6-h simulation period. The blue solid line (uniform/static simulation) overlaps perfectly with the dashed line with red crosses (stricter criterion AMR simulation) and the dashed line with yellow circles (more lenient criterion AMR simulation). The largest differences between the static simulation and the AMR simulations are ~0.001 m s^{-1}.

time step constrained by the vertical resolution. A different time integration scheme might allow adaptive time stepping, where the time step depends on the horizontal resolution. With the use of such a scheme, further speedup may be possible.

A second set of tests is performed to observe the effect of changing the refinement criterion on the time to solution and to see how AMR performs over the full 6-h simulation period. Like the previous tests, these are also performed using 1536 cores, and with the exception of the test with four levels of AMR, all use a vertical resolution $\Delta z = 312$ m. This set of tests includes four simulations:

- A simulation using a uniform grid with a horizontal resolution $\Delta x = \Delta y = 1$ km.
- A simulation with two levels of refinement (up to 1-km horizontal resolution) with the same refinement criterion $\text{threshold}_{\text{st}} = \text{threshold}_2$ as the previous two-level simulation; this will be referred to as the strict two-level simulation.
- A simulation with two levels of refinement but with a more lenient refinement criterion $\text{threshold}_{\text{ln}} = [2, 5]$ m s^{-1}; this will be referred to as the lenient two-level simulation.
- A simulation with four levels of refinement (up to 250-m horizontal resolution) and a vertical resolution $\Delta z = 250$ m. The thresholds for the simulation are $\text{threshold}_4 = [7.5, 15, 22.5, 30]$ m s^{-1}.

Figure 4 presents the time to solution for the second set of tests in normalized time. For these tests, the baseline is the strict two-level AMR simulation as it is the cheapest to perform. We can see that while the two simulations using two levels of AMR remain cheaper than the uniform alternative, the more lenient criterion incurs a substantial increase in the cost of the simulation.

Figure 5 shows the number of vertical columns over time for this set of simulations. We can see that with the stricter AMR criteria and two levels of refinement (purple line), the number of columns does not vary substantially over the course of the simulation. With the more lenient criteria and two levels of refinement (yellow line), the number of columns at the end of the simulation is double what it was at the beginning which explains the significant increase in cost. The simulation using four levels of AMR shows a substantial increase in the number of columns over the course of the simulation; the number of columns at the end is 5 times what it was at the beginning. As the hurricane intensifies, more areas of high velocity magnitude appear and the mesh adapts by refining around these areas and generating more columns. While the total number of columns for the four-level AMR simulation is less than the total number of columns for the uniform (static) simulation, it still requires a longer time to complete. The main reason for this is the increased vertical resolution which increases the total number of degrees of freedom and further restricts the time step required for stability. For the remainder of this section, we will continue referring to this set of simulations.

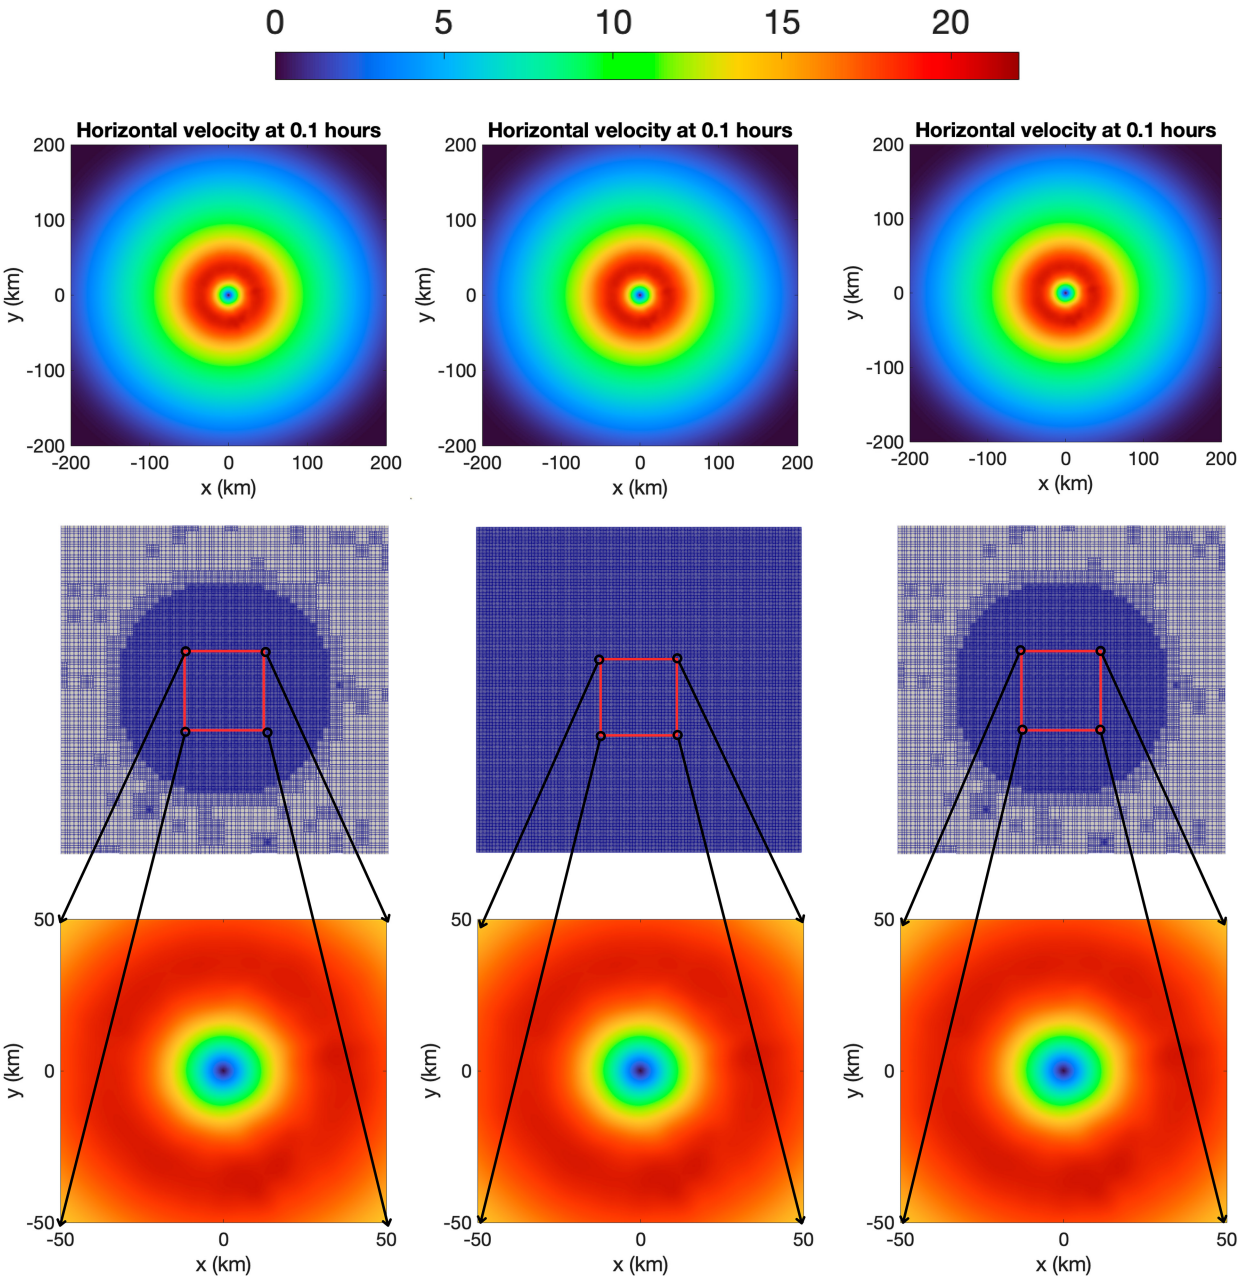

FIG. 8. Horizontal winds taken at height $z = 1000$ m above sea level at $t = 0.1$ h. (left) Results for a simulation with two levels of AMR and a maximum horizontal resolution of 1 km. (center) Results for a uniform grid with 1-km constant horizontal resolution. (right) Results for a simulation with a maximum of four levels of AMR. At this early stage of the simulation, only two levels of refinement can be seen even on the four-level capable simulation as the criteria for higher refinement are not met yet.

d. Accuracy of AMR simulations

Figure 6 shows the maximum values of the horizontally averaged horizontal velocity over time for $(x, y) \in [-50, 50]$ km $\times$ $[-50, 50]$ km. These values are found by finding the average value of horizontal velocity at each vertical level and then finding the maximum among these.

This horizontal subdomain is chosen as it should allow for focusing on where the observational heating takes place. Figure 7 shows the values of maximum horizontal velocity over time. The height of maximum horizontal velocity varies over time but is usually found between $z = 500$ m and $z = 4000$ m. Both figures compare these values for the uniform simulation and the two two-level AMR. Both figures show an essentially perfect overlap for these quantities and demonstrate the AMR's ability to capture the intensification of the storm even with relatively strict criteria. In fact, the largest differences

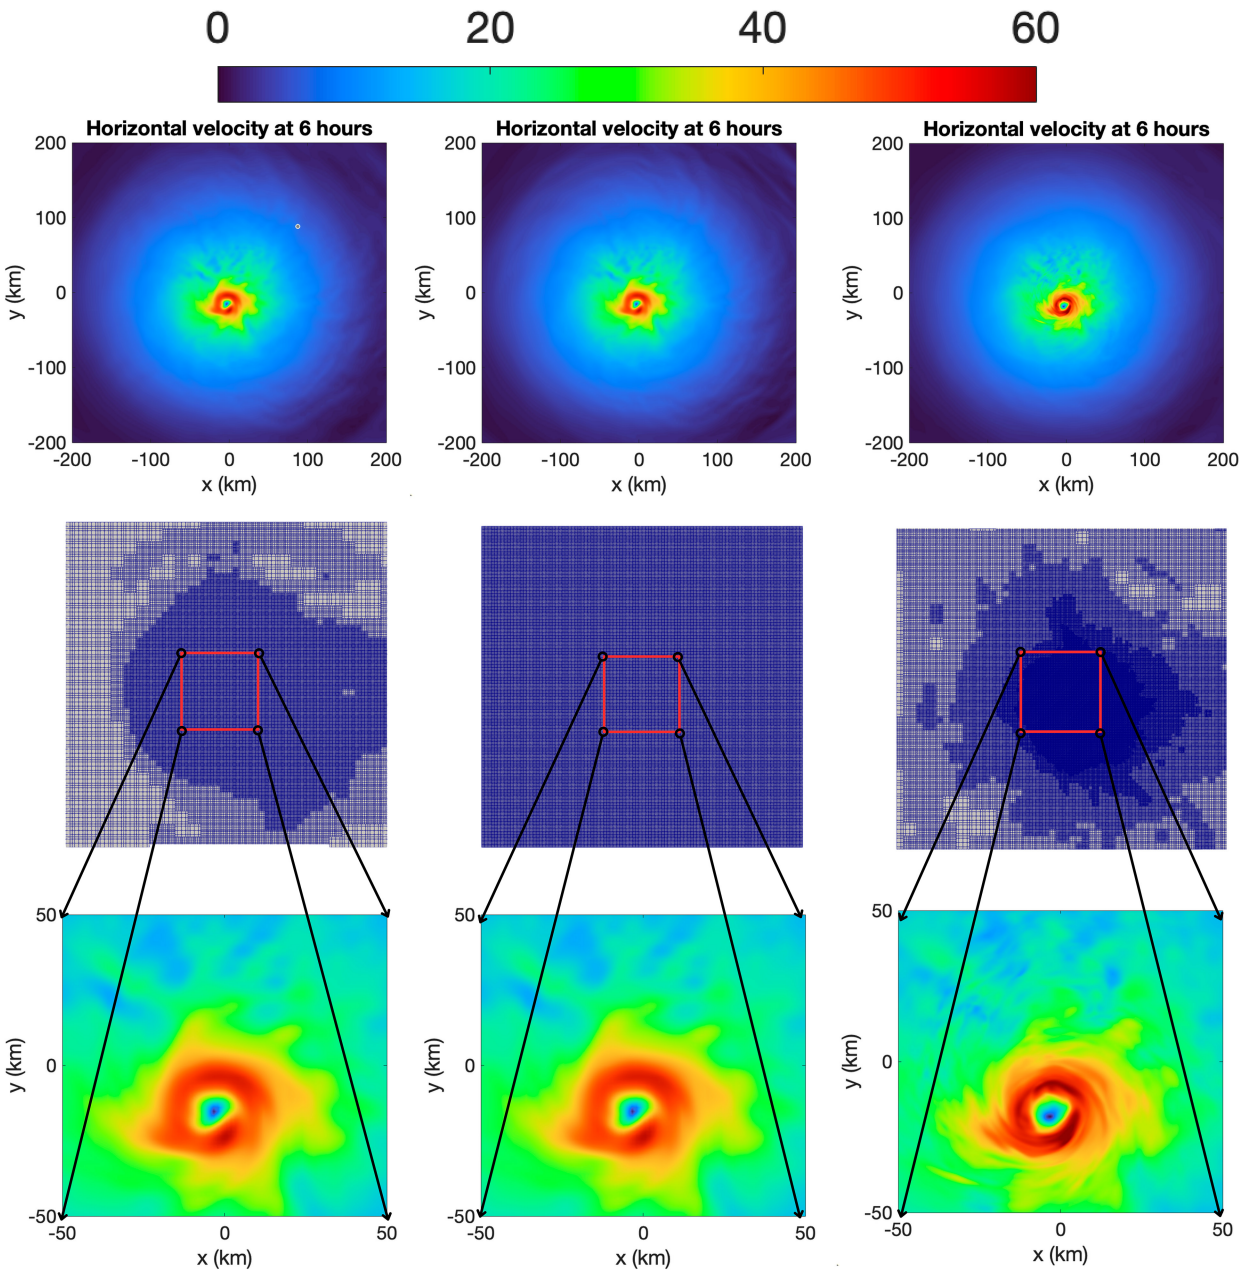

FIG. 9. Horizontal winds taken at height $z = 1000$ m above sea level at $t = 6$ h. (left) Results for a simulation with two levels of AMR and a maximum horizontal resolution of 1 km. (center) Results for a uniform grid with 1-km constant horizontal resolution. (right) Results for a simulation with four levels of AMR and a maximum horizontal resolution of 250 m. All four levels of refinement are visible on the grid in the right column.

between the static and AMR simulations are ~0.002 m s^{-1} for Fig. 6 and ~0.001 m s^{-1} for Fig. 7.

The intensification of the vortex is rather large over this short time period, but this is not without precedence in nature. Hurricane Patricia (2015) broke records for storm intensification rate with an increase in maximum sustained winds of 54 m s^{-1} in 24 h (Rogers et al. 2017). It is important to note that the simulations in this paper are idealized: 1) no environmental vertical wind shear, 2) no active moisture and associated entrainment of dry air into the system, and 3) no ocean coupling or cooling processes and associated effects on the system thermodynamics and dissipation characteristics. All of these idealizations will generally lead to weakening of the storm. In addition, the radius of maximum winds in our initial vortex is a bit large, which leads to more rapid contraction with the specified heating.

Despite the idealizations listed above, the inner-core heating perturbations are derived from airborne radar measurements in a real, rapidly intensifying hurricane with a sophisticated

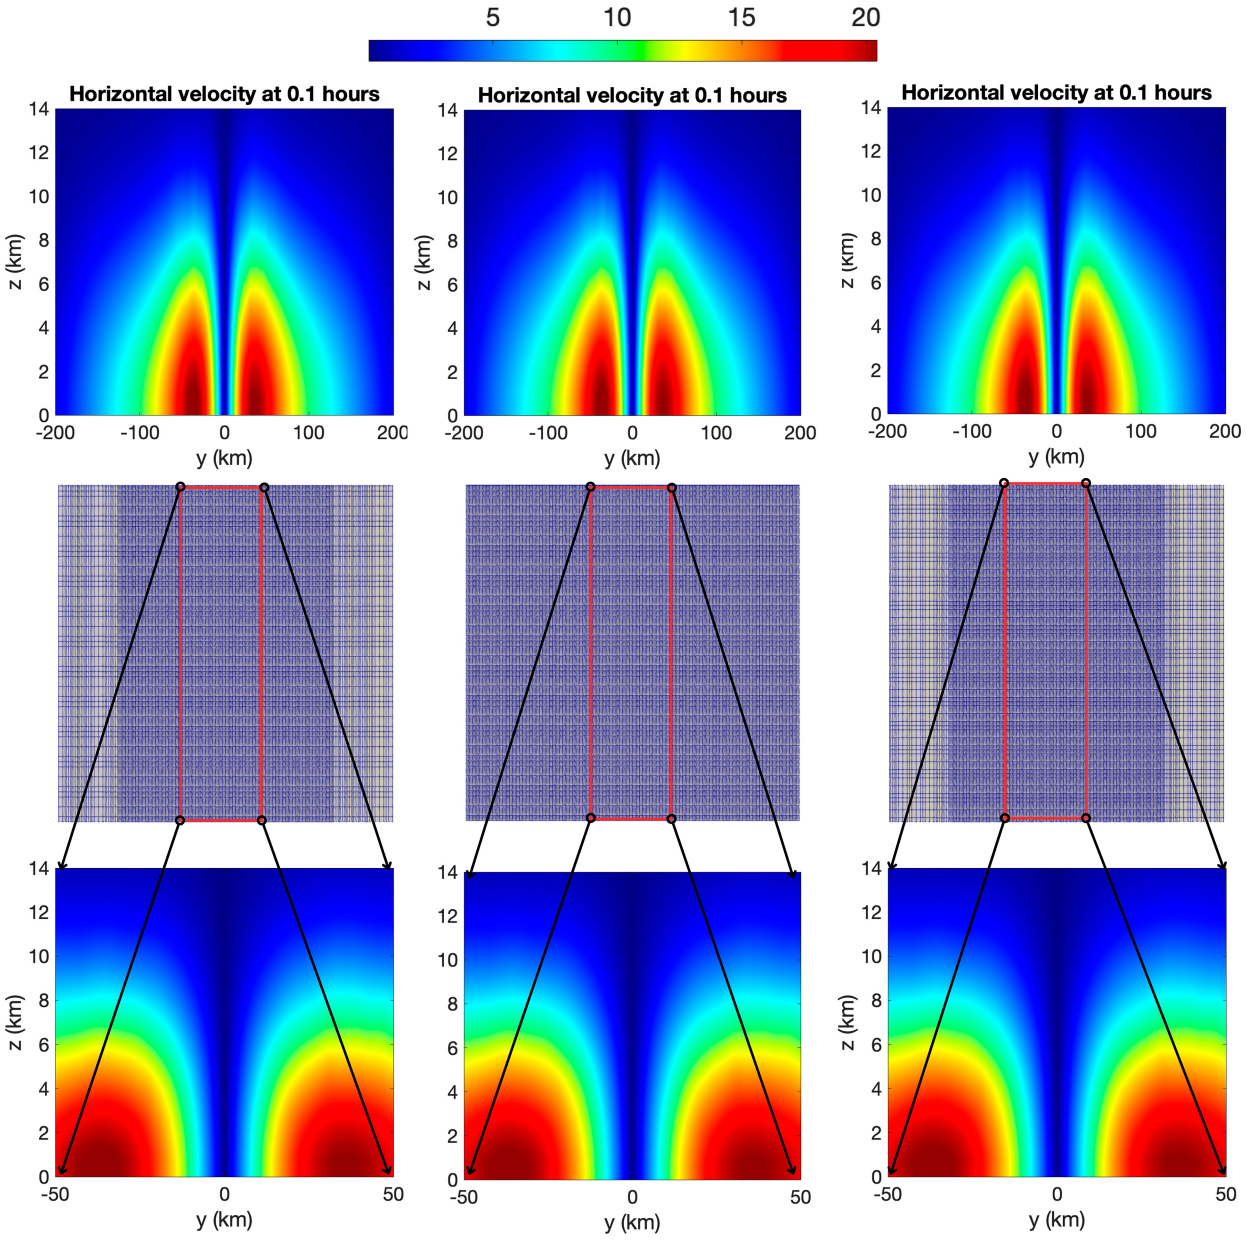

FIG. 10. Horizontal velocity taken at $x = 0$ m at $t = 0.1$ h. (left) Results for a simulation with two levels of AMR and a maximum horizontal resolution of 1 km. (center) Results for a uniform grid with 1-km constant horizontal resolution. (right) Results for a simulation with four levels of AMR and a maximum horizontal resolution of 250 m. At this early stage of the simulation, only two levels of refinement can be seen even on the four-level capable simulation as the criteria for higher refinement are not met yet.

retrieval algorithm. The heating perturbations are certainly not perfect [see Guimond et al. (2011) for a detailed uncertainty analysis], but they represent some of the highest quality fields available in the hurricane community. In addition, the fundamental physical process driving hurricane intensification is the inner-core heating, which is controlled largely by convective bursts. Thus, the simulations are representing the fundamental hurricane intensification dynamics with good quality forcing, which is a useful test case for our AMR algorithm. Extensions to more detailed simulations with the addition of wind shear and moisture are possible, but they are left for future work.

Figure 8 shows the horizontal velocity of the storm at $t = 0.1$ h for a horizontal slice taken at $z = 1000$ m. We can see that the left column (two-level AMR with strict criterion), middle column (uniform grid), and right column (four-level AMR) are identical. At this early stage of the simulation, the four-level AMR simulation behaves identically to the two-level AMR simulation because the criteria for the third and fourth levels of refinement are not met yet. Figure 9 is similar

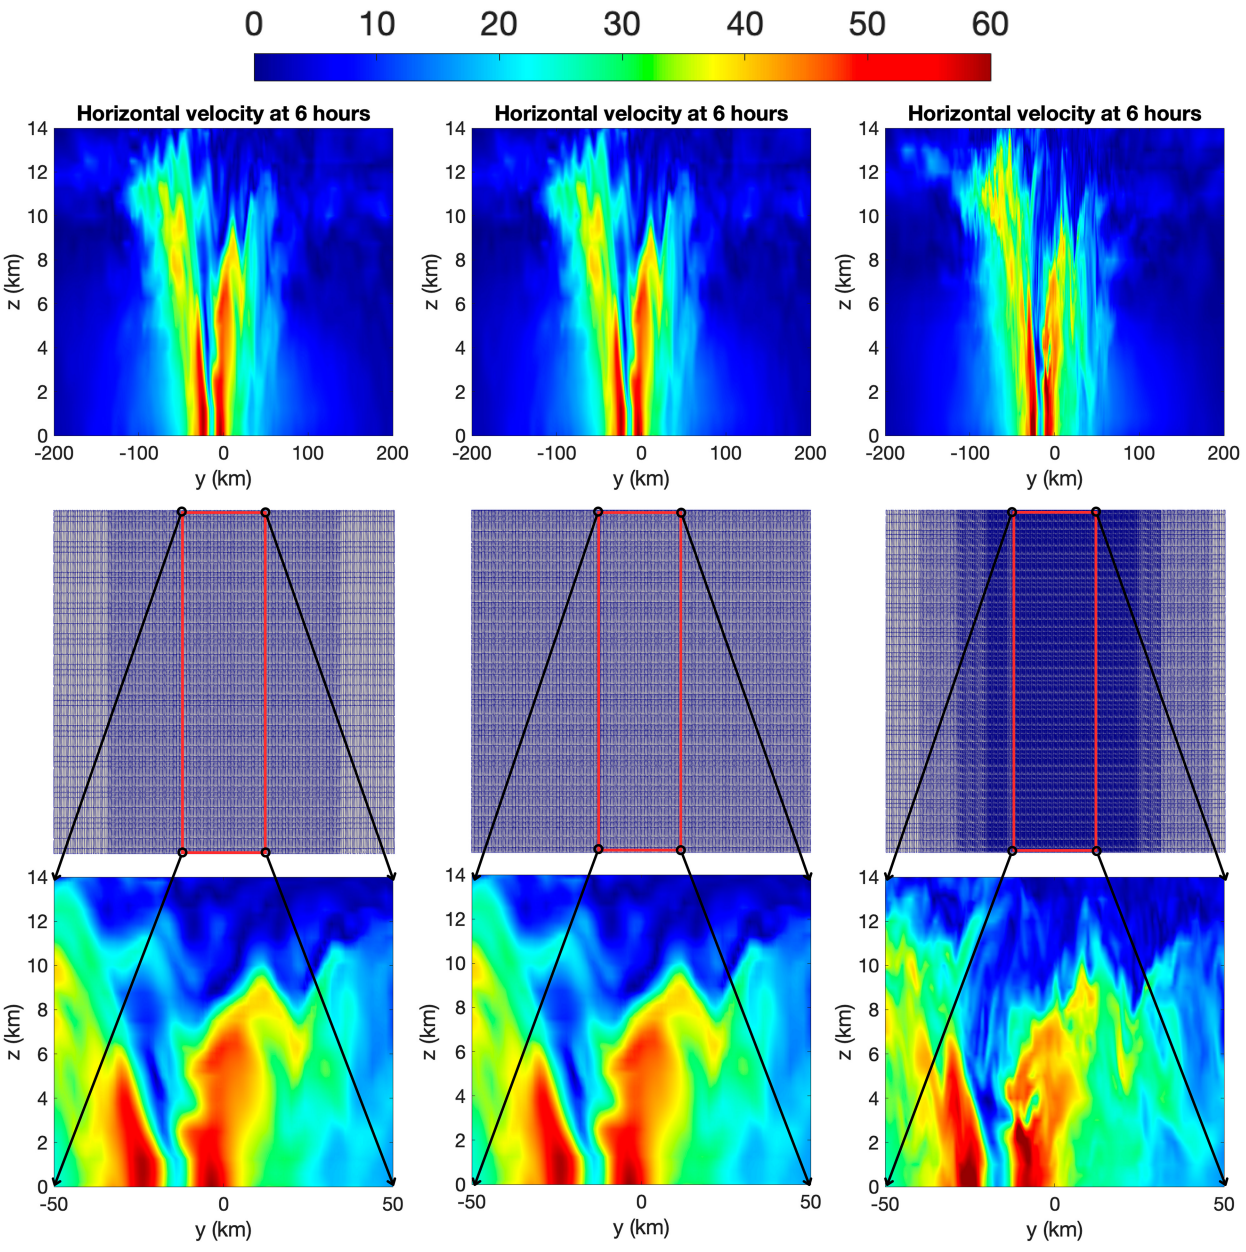

FIG. 11. Horizontal velocity taken at $x = 0$ m at $t = 6$ h. (left) Results for a simulation with two levels of AMR and a maximum horizontal resolution of 1 km. (center) Results for a uniform grid with 1-km constant horizontal resolution. (right) Results for a simulation with four levels of AMR and a maximum horizontal resolution of 250 m. All four levels of refinement are visible on the grid in the right column.

to Fig. 8 but corresponds to the end of the simulations at $t = 6$ h. The left and middle columns (two-level AMR and uniform grid, respectively) are still identical even at the end of the simulation. The right column (four-level AMR) allows us to see the additional details and structures that can be captured with additional refinement, thanks to criteria for all four levels of refinement being met in the domain center. Figure 10 shows the horizontal velocity of the storm at $t = 0.1$ h for a vertical slice taken at $x = 0$ m. Once again, all the columns are identical at this stage of the simulation. Figure 11 is similar to Fig. 10 but corresponds to the end of the simulations at $t = 6$ h. The left and middle columns are still identical and demonstrate the ability of AMR to obtain high fidelity results at a fraction of the cost of a simulation using a uniform grid. Figure 12 presents the vertical vorticity of the storm at $t = 3$ h (top two rows) and $t = 6$ h (bottom two rows) for a horizontal slice at $z = 1000$ m. Even for this derived quantity, the left and center columns are still identical. The right column showcases what additional resolution has to offer in terms of resolving turbulent structures. The simulations with resolutions of up to 250 m of horizontal resolution show an increase in both vorticity and velocity compared to their lower-resolution

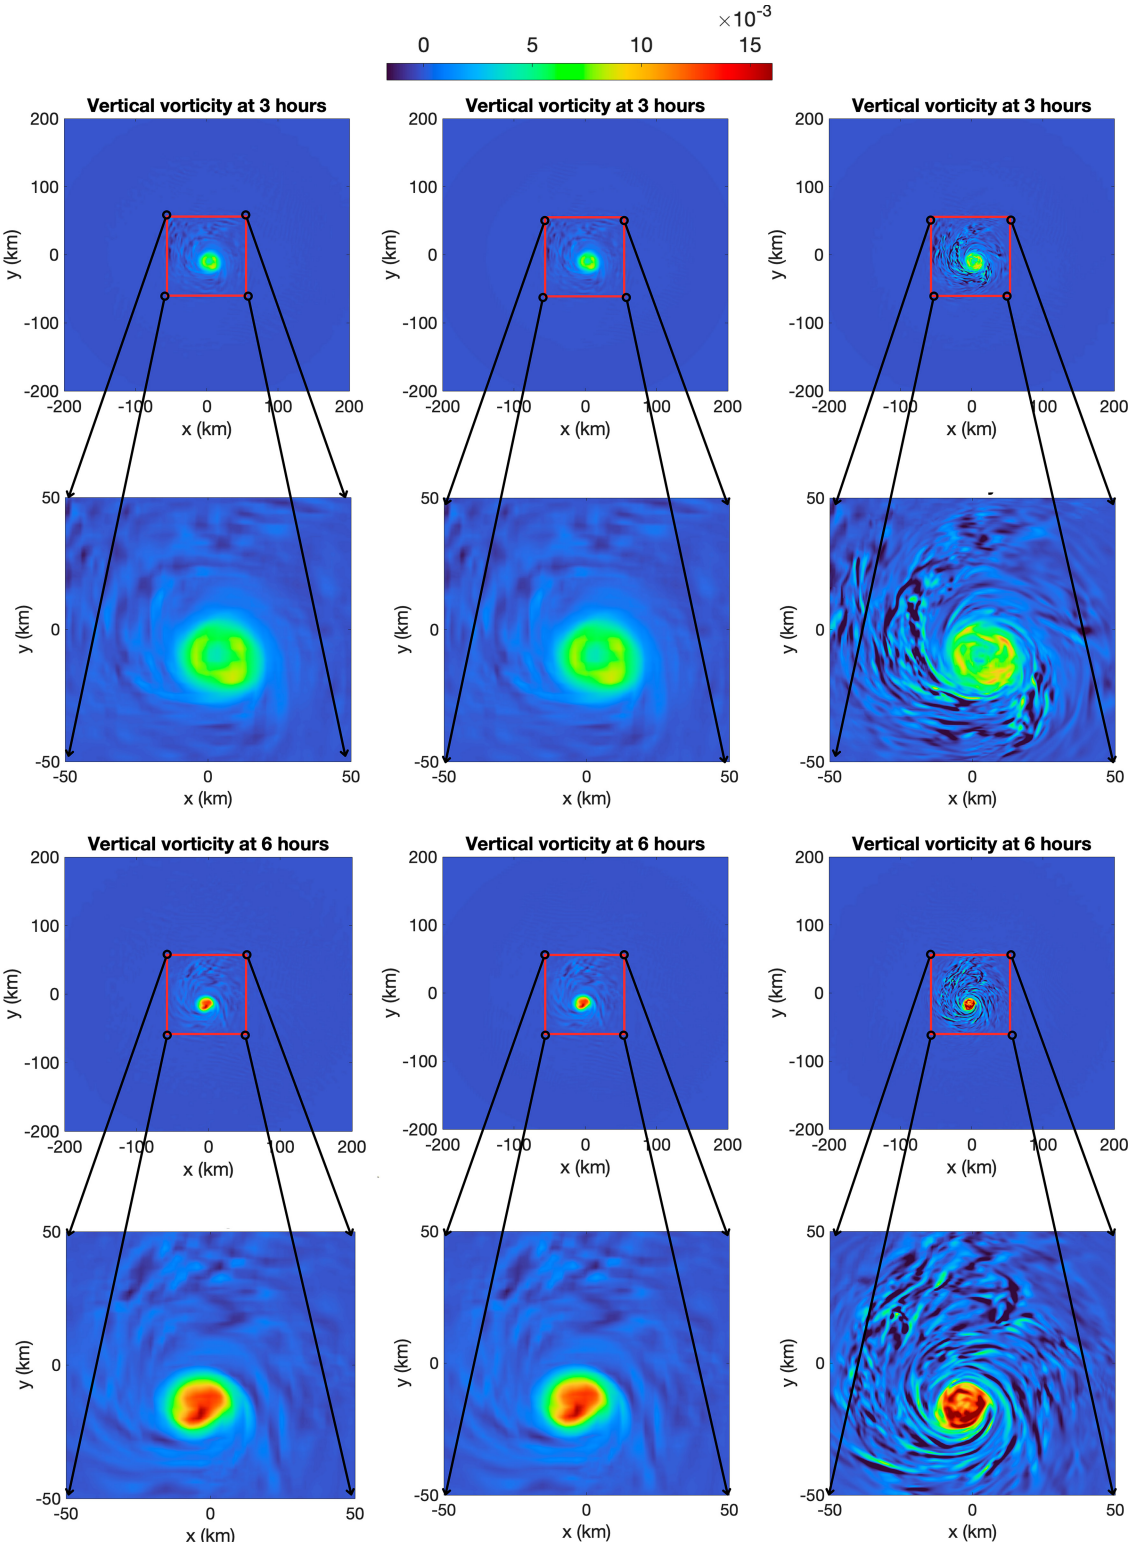

FIG. 12. Vertical vorticity taken at height $z = 1000$ m above sea level at $t = 3$ h in the top two rows and $t = 6$ h in the bottom two rows. (left) Results for a simulation with two levels of AMR and a maximum horizontal resolution of 1 km. (center) Results for a uniform grid with 1-km constant horizontal resolution. (right) Results for a simulation with four levels of AMR and a maximum horizontal resolution of 250 m.

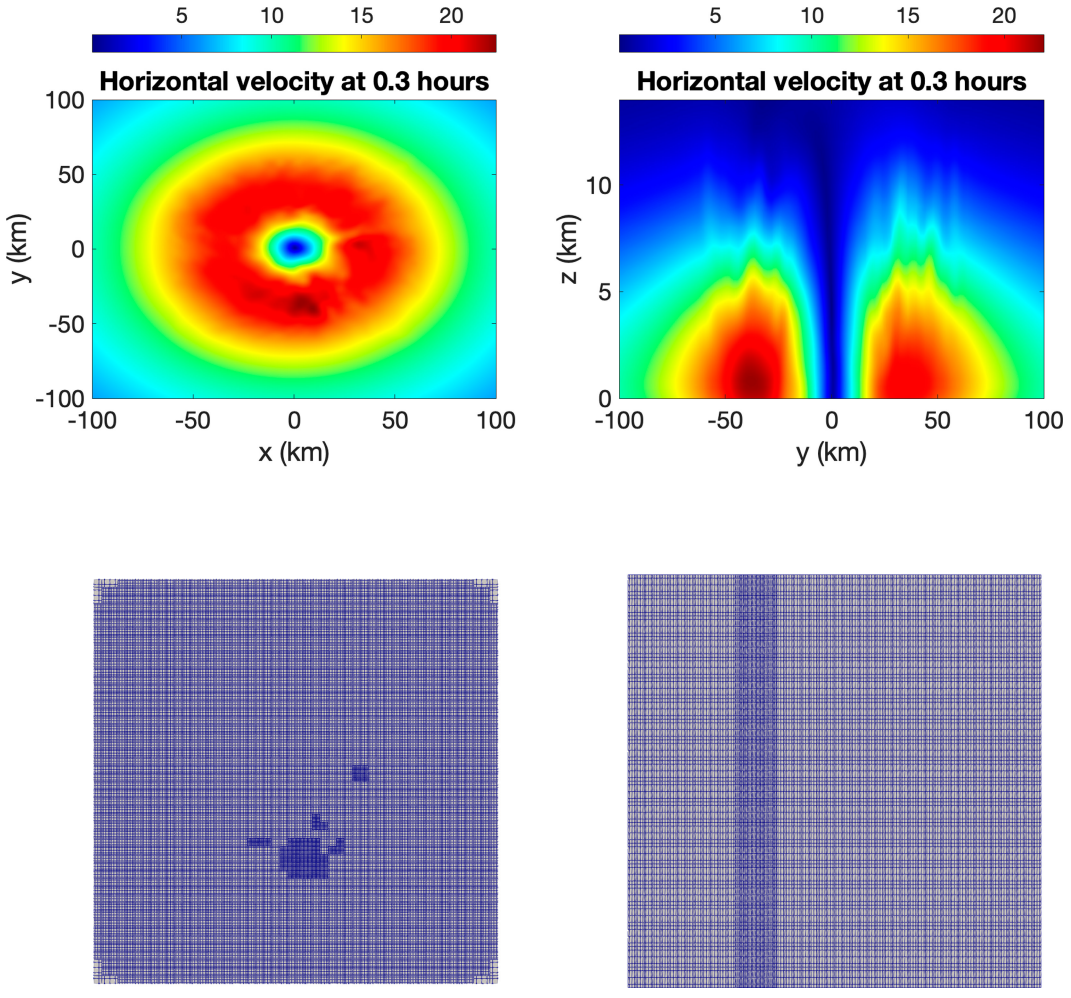

FIG. 13. Horizontal velocity and grid at the first instance of reaching the third level of refinement for the four-level AMR simulation. (top left) Horizontal velocity at $z = 1000$ m and $t = 0.3$ h. (top right) Horizontal velocity at $x = 0$ m and $t = 0.3$ h. (bottom left) Horizontal grid at $z = 1000$ m, $t = 0.3$ h, and for (x), $(y) \in [-100, 100]$ km $\times$ $[-100, 100]$ km. (bottom right) Vertical grid at $x = 0$ m, $t = 0.3$ h, and for (y), $(z) \in [-100, 100]$ km $\times$ $[-0, 20]$ km. Note that the lowest resolution visible in the bottom-left panel is 2 km (seen in the corners) and the highest is 500 m. The baseline 4-km resolution is not visible as it is present farther away from the domain center. In the bottom-right panel, the lowest visible horizontal resolution is 1 km and the highest is 500 m. The lower resolutions are present in other areas of the simulated domain not pictured here. Areas with a velocity magnitude larger than 22.5 m s^{-1} trigger the criterion for the third level of refinement and cause grid to reach 500 m of horizontal resolution wherever the criterion is met. As this refinement is done on a column basis, the entire vertical column is refined as shown in the bottom-right panel.

counterparts. This reflects the role of the additional resolution in capturing the intensification of the storm. Vorticity is especially important when considering features like strong updrafts and convective towers. Pushing this resolution even further could allow for studies of how these features affect TC intensification and RI through numerical experiments.

We now focus on the simulation using four levels of refinement. Figure 13 shows the first instance we detect of the grid refining to a horizontal resolution $\Delta x = \Delta y = 500$ m. This refinement takes place at $t \sim 0.3$ h, near the beginning of the simulation. The top two plots show areas of higher horizontal velocity (dark red shades) which corresponds to the criterion for refining to 500 m being met, and the bottom two plots of the figure show how the grid responds to the intensification. Figure 14 shows the first instance we detect of the grid refining to a horizontal resolution $\Delta x = \Delta y = 250$ m. This refinement occurs at $t \sim 0.4$ h, still near the beginning of the simulation. As shown in the top-middle plot, high values of velocity magnitude trigger the refinement process causing the grid to adapt (middle-bottom plot). At this time, the high value of velocity magnitude is associated with the occurrence of a strong updraft, as evidenced by the high vertical velocity values shown in the top-right plot of the figure. The refinement criterion being dependent on the magnitude of the velocity vector allows for refinement to happen even if the horizontal velocities are not yet strong enough to induce refinement on their own.

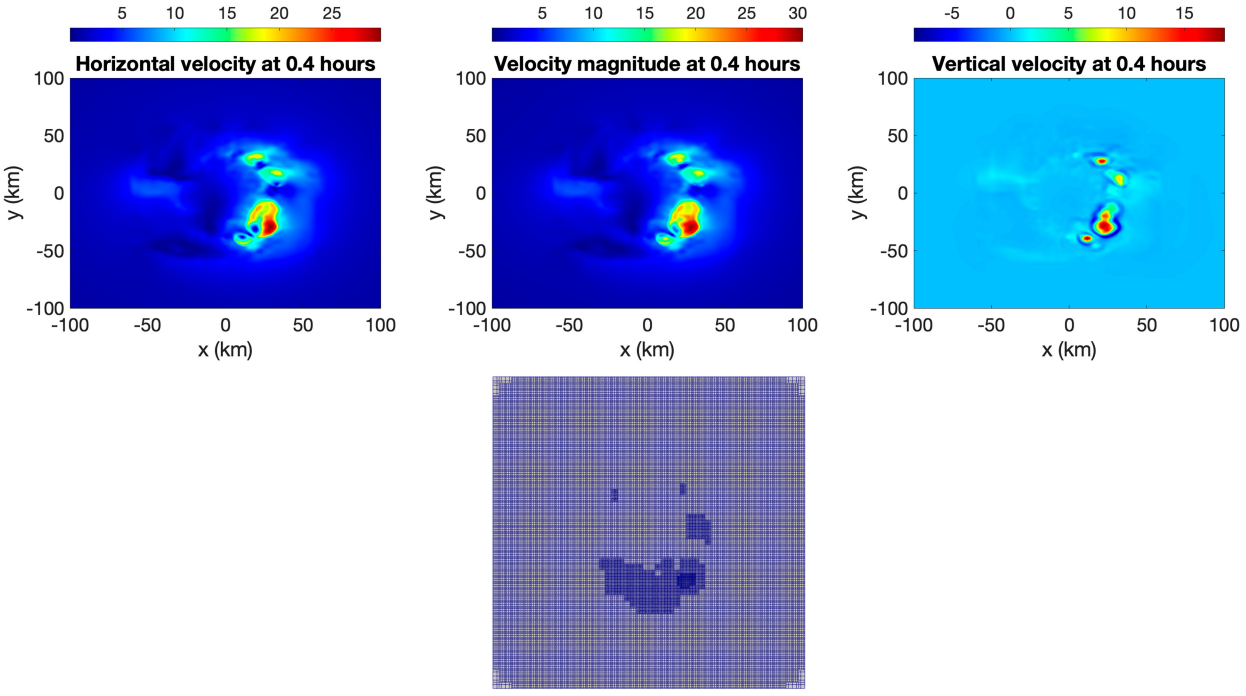

FIG. 14. Velocity and grid at the first instance of reaching the third level of refinement for the four-level AMR simulation. (top left) Horizontal velocity at $z = 12\,000$ m and $t = 0.4$ h. (top middle) Velocity magnitude at $z = 12\,000$ m and $t = 0.4$ h. (top right) Vertical velocity at $z = 12\,000$ m and $t = 0.4$ h. (bottom middle) Horizontal grid at $z = 12\,000$ m $t = 0.4$ h and for (x), $(y) \in [-100, 100]$ km $\times$ $[-100, 100]$ km. In the bottom-middle panel, the lowest visible horizontal resolution is 1 km and the highest is 250 m. A substantial updraft indicated by the elevated vertical velocity shown in the top-right panel allows the velocity magnitude (shown in the top-middle panel) to pass the 30 m s^{-1} threshold required for the fourth level of AMR. The horizontal velocity (shown in the top-left panel) would not have been sufficient to trigger refinement on its own, and the AMR algorithm is effectively refining around this updraft.

Figure 15 shows the four-level AMR simulation at $t = 6$ h and at the $z = 1000$-m horizontal plane within the boundaries of $(x, y) \in [-200, 200]$ km $\times$ $[-200, 200]$ km. All the levels of the AMR grid are visible, and as we would expect, the grid is finer near the storm center than it is farther away.

Figures 16 and 17 show the maximum horizontally averaged horizontal wind over time and the maximum value of horizontal wind over time, respectively. Both figures show a close overlap for the earliest stages of the respective simulations (up to $t = 0.3$ h) but begin to separate at $t = 0.4$ h for Fig. 16 and $t = 0.5$ h for Fig. 17. These times correspond to instants the AMR simulation first refines to $\Delta x = \Delta y = 500$ m ($t = 0.4$ h) and first refines to $\Delta x = \Delta y = 250$ m ($t = 0.5$ h). Both figures show that the four-level AMR simulation produces a more intense storm with the difference in intensity generally being between 5% and 10%. Both storms are responding to exactly the same forcing from the observational heating, and Figs. 6 and 7 show that for the same maximum resolutions, the AMR and uniform simulations are able to produce the same results. As such, the difference in storm intensity being observed here is a result of the difference in maximum resolution between the two simulations.

e. *Testing AMR for longer TC simulations*

We perform a set of two simulations to see how AMR performs over the course of an idealized, 18-h-long TC simulation. The first simulation uses a uniform fixed horizontal resolution $\Delta x = 1$ km, and the second simulation has an initial uniform resolution of $\Delta x = 4$ km and uses two levels of AMR to refine up to $\Delta x = 1$ km. The refinement criterion is $\text{threshold}_{\text{st}}$. These simulations are divided into two parts. During the first 12 h, we introduce an analytical heating function identical to that used in Hasan et al. (2022). This heating is obtained by adding potential temperature perturbations from the azimuthal wavenumbers 0, 2, 3, 4, and 5 with respective amplitudes 2.0, 1.5, 1.0, 1.0, and 1.0 K. The combined potential temperature perturbation is divided by its maximum value and multiplied by 100. The objective is to obtain a maximum heating rate of 100 K h^{-1}. Equations (19) and (20) describe the potential temperature perturbation corresponding to a given wavenumber:

$$\theta'_{\tilde{\nu}} = A_{\tilde{\nu}} \cos(\tilde{\nu}\phi) e^{-d}, \tag{19}$$

$$d = \left(\frac{r - r_b}{s_r}\right)^2 + \left(\frac{z - z_b}{s_z}\right)^2, \tag{20}$$

where $A_{\tilde{\nu}}$ is the amplitude corresponding to the wavenumber $\tilde{\nu}$, $\phi = \arctan(y, x)$, $r = \sqrt{x^2 + y^2}$, $r_b = 40\,000$ m, $z_b = 5000$ m, $s_r = 200\,000$ m, and $s_z = 2000$ m. The wavenumber-based heating is allowed to ramp up to the

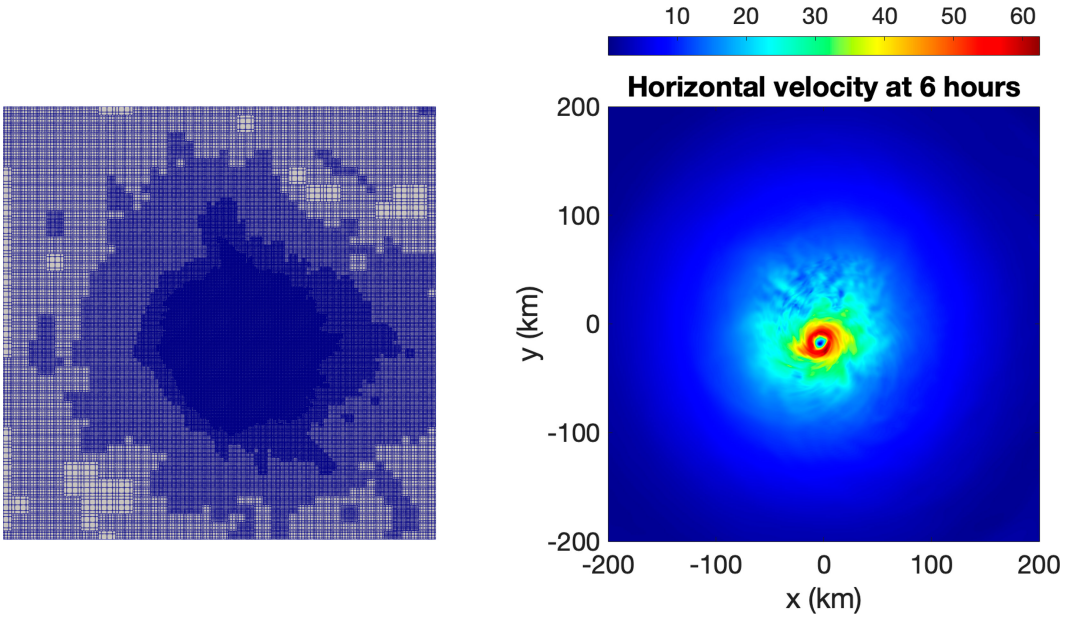

FIG. 15. Example of four-level refinement at $t = 6$ h. The criterion used is threshold$_4$ = [7.5, 15, 22.5, 30] m s^{-1} on velocity magnitude. (left) Horizontal cross section of the AMR grid. (right) Horizontal velocity at $t = 6$ h and $z = 1000$ m.

full effect during the first 30 min of the simulation. At 12 h, the wavenumber-based heating is ramped down while we introduce the observational heating in a similar manner to the 6-h simulation.

To see if AMR is able to adapt to less intense flow features that might be present in the outer areas of a TC, we introduce a weaker wavenumber 6 heating perturbation that is present for the full 18-h duration. This heating is handled

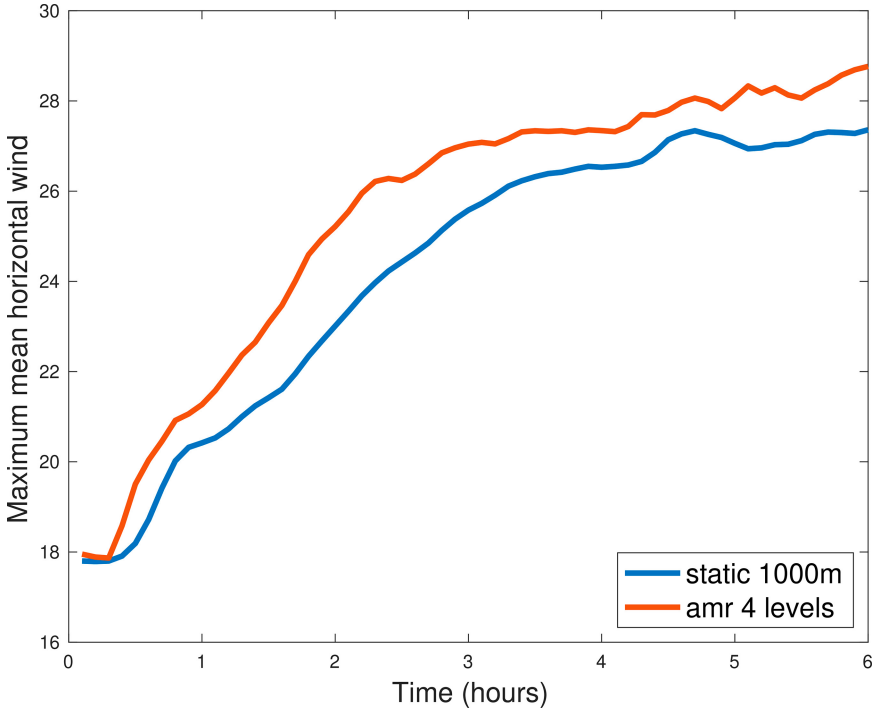

FIG. 16. Maximum of horizontally averaged horizontal wind within the [−50, 50] km^2 subdomain over the course of the 6-h simulation period for the uniform 1-km resolution simulation (blue) and the four-level AMR simulation (red).

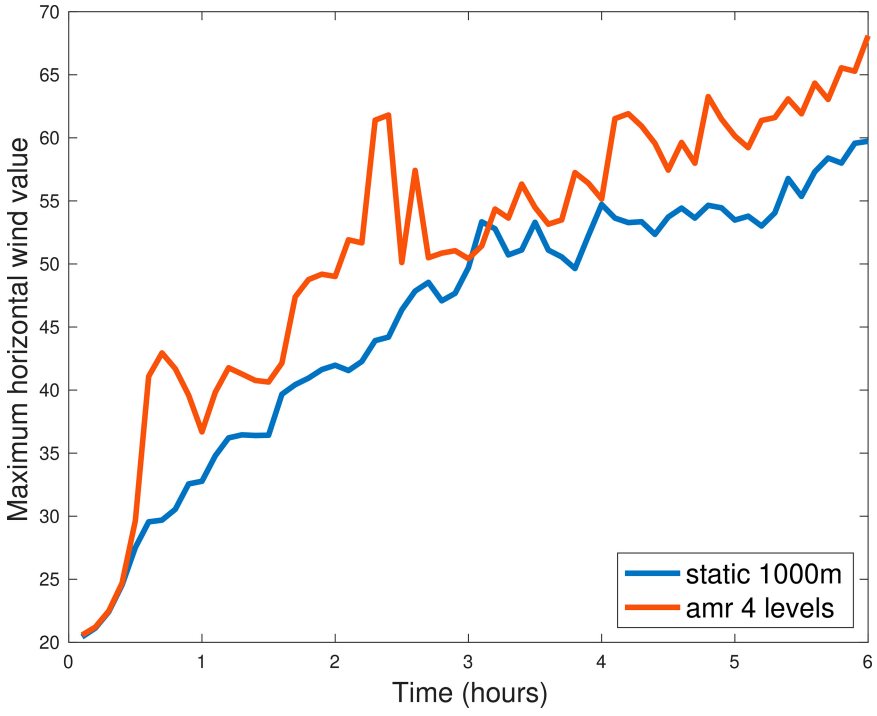

FIG. 17. Maximum value of horizontal wind within the $[-50, 50]$ km^2 subdomain over the course of the 6-h simulation period for the uniform 1-km resolution simulation (blue) and the four-level AMR simulation (red).

in a similar manner to the wavenumber heating introduced in the first 12 h of the simulations. However, the maximum heating rate is 20 K h^{-1}. We also choose to have this heating be effective in the outer regions of storm rather than in the core. As such, we adjust the corresponding perturbation as follows:

$$\theta_6' = \cos(6\phi)e^{-d_6}, \tag{21}$$

$$d_6 = \left(\frac{r - r_{b6}}{s_{r6}}\right)^2 + \left(\frac{z - z_{b_6}}{s_{z6}}\right)^2 + \left(\frac{180\,000 - r}{60\,000}\right)^2, \tag{22}$$

where $r_{b6} = 20\,000$ m, $z_{b6} = 7000$ m, $s_{r6} = 250\,000$ m, and $s_{z6} = 4000$ m.

1) TIME TO SOLUTION

Both tests are performed on 256 cores. The AMR simulation required just under 4 days of execution time, while the static simulation lasted over 15 days and required ~370 h of execution time. Although the AMR simulation is faster by a factor of 3.9, we have seen that in the previous simulation sets, we were able to obtain even more speedup using the same criterion. We think the reason for this is the wavenumber-6 heating that is present throughout the entire length of the simulation and is active even far from the core. The AMR responds to this additional activity by refining, and the extra grid points result in an increase in cost.

2) BEHAVIOR AND ACCURACY OF THE AMR ALGORITHM

Figures 18 and 19 compare the vertical vorticity and the horizontal velocity at the end of the two simulations. In both figures, it is clear that there is a substantial increase in activity in the outer regions of the storm compared to the previous simulation sets. A wavenumber-6 pattern is visible in both the vertical vorticity and the horizontal velocity fields. The wavenumber-6 heating has had the desired effect of increasing the amount of activity in the outer regions and that can be observed through the secondary flow features. Figure 19 even shows the presence of bands that extend out of the eyewall. As such, this test should let us analyze how AMR responds to these secondary flow features which are the results of a weaker but longer lasting heating.

3) AMR- AND CRITERION-ASSOCIATED LIMITATIONS

Figure 20 illustrates the spectra of the radially averaged velocity at height $z = 1000$ m for the two simulations for two different radial bands and at different times. Through this figure, we can see that at the end of the first 6 h, the spectra of the static simulation (red solid line) and that of the AMR simulation (purple solid line with stars) overlap very well for both choices of radial extension for $0 < r < 50$ km and $50 < r < 200$ km. While the spectra overlap very well until 12 h, they start to separate at 18 h. The behavior of the spectra over time indicated that AMR does very well in capturing the same behavior as the static grid, especially over shorter durations. The AMR seems to be able to

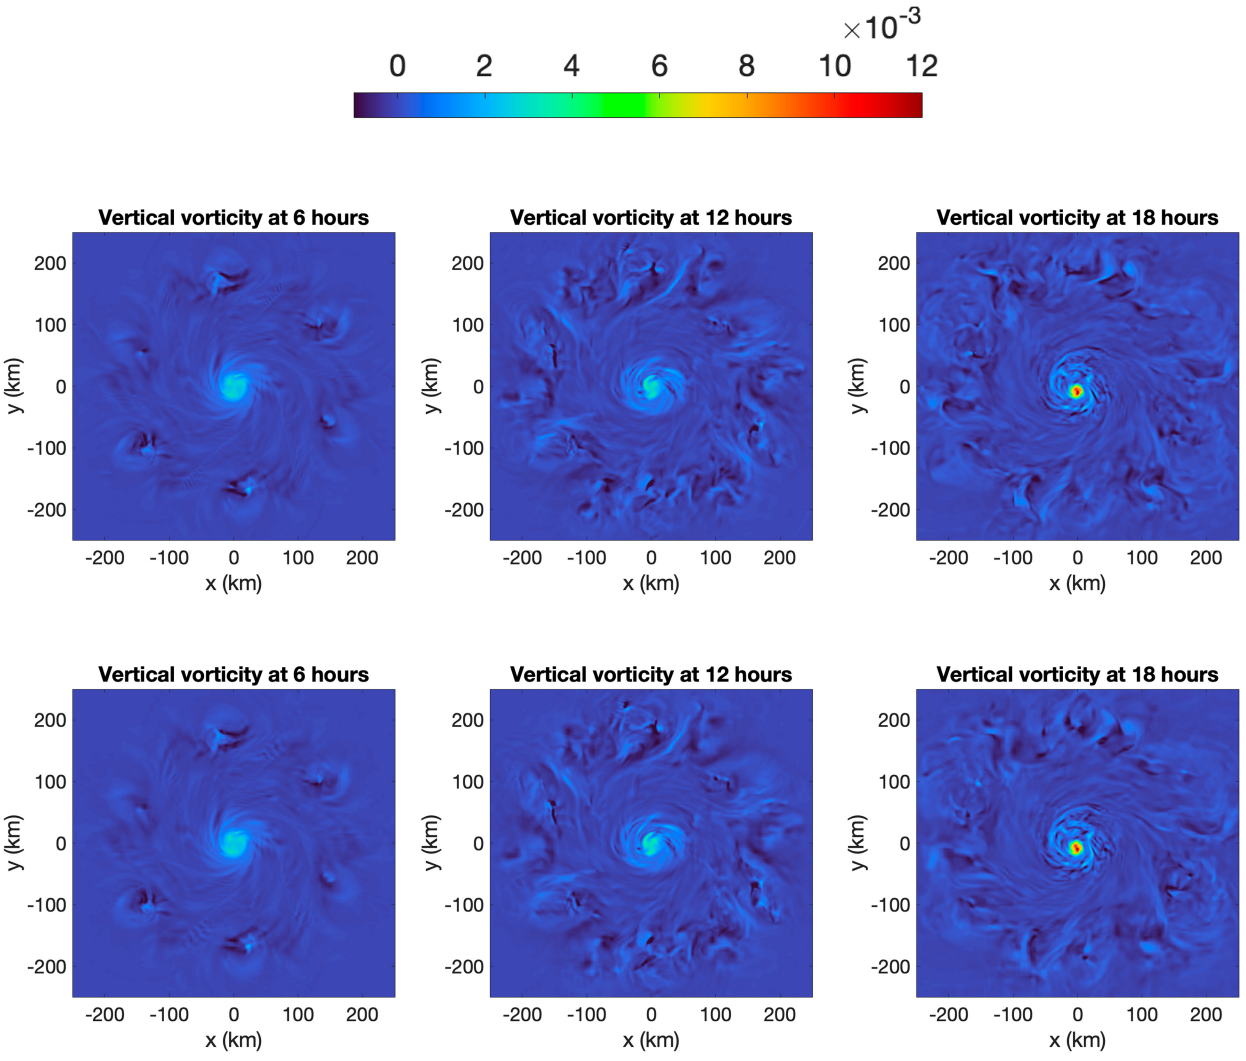

FIG. 18. Vertical vorticity taken at height $z = 2000$ m above sea level and (left) $t = 6$ h, (middle) $t = 12$ h, and (right) $t = 18$ h for the long duration simulations and for (x), $(y) \in [-250, 250]$ km $\times$ $[-250, 250]$ km. (top) Results of the static grid simulation with a constant horizontal resolution of 1 km. (bottom) Results of the simulation with two levels of AMR and a maximum horizontal resolution of 1 km.

capture the behavior of the storm and associated nonlinear interactions between scales, reasonably well. However, over longer durations, the accuracy of an AMR simulation with respect to a static grid seems to decrease a bit. We believe this is due to two main causes:

- The chosen refinement criterion might not be able to detect and refine around some flow features. These features could eventually become important to flow's development but remain poorly resolved because of an inadequate refinement criterion or threshold.
- Even if AMR eventually refines around an important flow feature, it might be too slow to react, and this might eventually result in differences with respect to a static mesh.

Given that the differences in the spectra are more pronounced for the larger maximum radius, we investigated the two simulations for any possible visible differences in the flow fields in the outer regions of the storm. Figure 21 shows the velocity magnitude of the storm region at $(x, y) \in [-250, -50]$ km $\times$ $[-250, -50]$ km and at a height of $z - 7000$ m. The left column of this figure shows a clear difference in the velocity magnitude fields of the static simulation (top right) and the AMR simulation (bottom right). Looking at the meshes (left column of Fig. 21) makes the source of these differences quite clear. The static grid simulation (top left) has a horizontal resolution of 1 km throughout the observed area, while the AMR simulation (bottom left) seems to have 1-km resolution in only a small portion of the area of concern with 2-km resolution in most of this area.

It is clear that differences in resolution are the primary sources of differences between the two simulations and their associated spectra. In fact, Fig. 22 tracks the behavior of

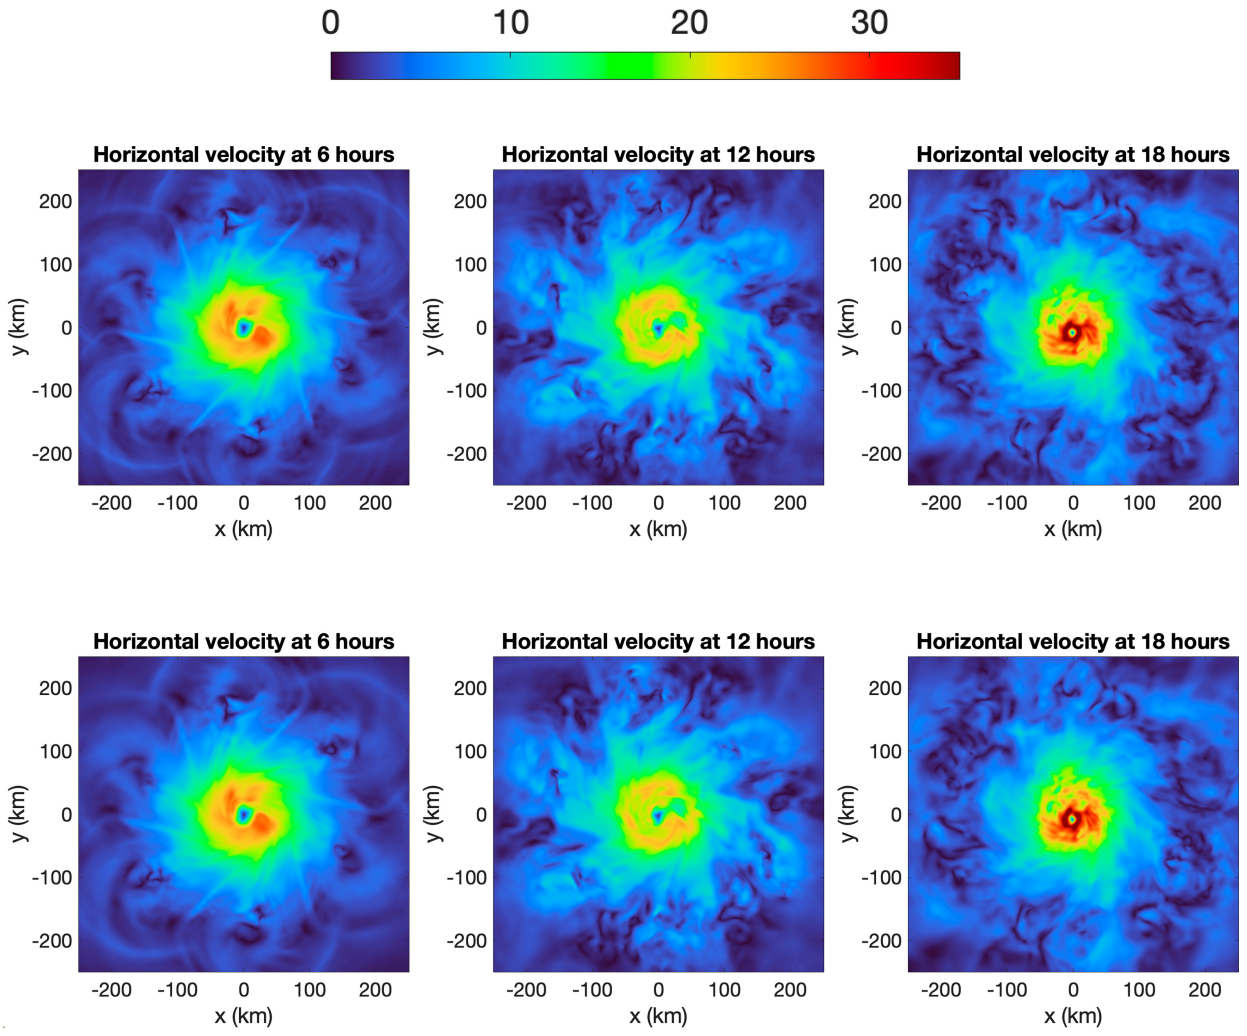

FIG. 19. Horizontal velocity taken at height $z = 2000$ m above sea level and (left) $t = 6$ h, (middle) $t = 12$ h, and (right) $t = 18$ h for the long duration simulations and for (x), $(y) \in [-250, 250]$ km $\times$ $[-250, 250]$ km. (top) Results of the static grid simulation with a constant horizontal resolution of 1 km. (bottom) Results of the simulation with two levels of AMR and a maximum horizontal resolution of 1 km.

the grid over time and highlights the area of focus that is presented in Fig. 21 (red box). Figure 22 shows that a full 30 min passes without any refinement taking place in the $(x, y) \in [-250, -50]$ km $\times$ $[-250, -50]$ km box. Even at $t = 1$ h, only about a quarter of the box has had level 1 refinement, and this is still the case at $t = 3$ h with only the addition of a few cells refined to level 2. Although eventually the entire box is refined to level 1, only a small portion of it is ever refined to level 2. The fact that refinement is limited in this area even though it is clearly an active area of the flow highlights the limitations of AMR and its reliance on the choice of criterion. Because the wavenumber 6 heating is comparatively weak, it takes longer to generate a powerful enough flow to trigger the AMR algorithm. This means that depending on the choice of criterion, some flow regions might remain poorly resolved for long periods of time and result in different behaviors compared to a static mesh. This emphasizes the importance of criterion selection for AMR simulations to appropriately resolve flow features of interest.

6. Discussion and conclusions

The AMR simulations presented in this work are of idealized dry rapidly intensifying TCs. These simulations also do not allow for adaptive vertical refinement of the grid spacing. The study of vertical and horizontal AMRs for simulations of moist TCs over their full life cycle could be the subject of future work due to limitations imposed by the chosen planetary boundary layer scheme and its compatibility with the spectral element method used for the spatial discretization. The planetary boundary layer scheme's reliance on column-based grids makes it impractical to use with a vertically adaptive grid, and its reliance on finite differences results in numerical noise when used in our Galerkin spectral element model. This noise

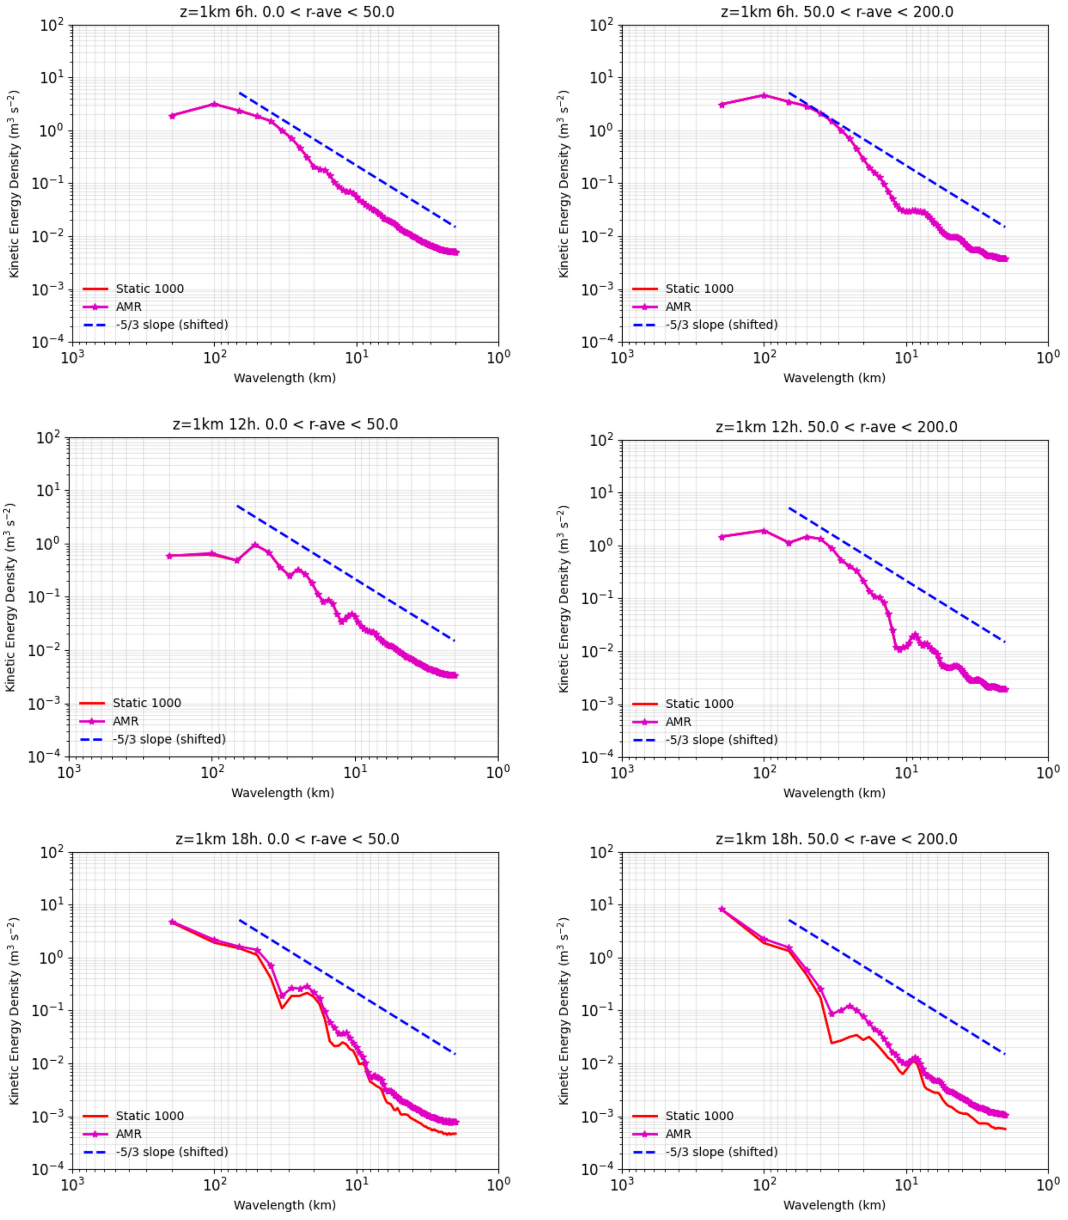

FIG. 20. One-dimensional spectra of radius-averaged velocity along the azimuth for the long duration simulation within a radial band (left) between the center and $r = 50$ km and (right) between $r = 50$ km and $r = 200$ km. The spectra are taken at $z = 1000$ m and (top) $t = 6$ h, (middle) $t = 12$ h, and (bottom) $t = 18$ h.

is manageable in dry simulations, but the sensitivity of microphysics models to noise makes it much more difficult to handle. To solve this issue would require fully rewriting the planetary boundary layer scheme for Galerkin numerical methods and is beyond the scope of this work. However, we want to emphasize that this model could be used to perform full-length moist simulations of TCs with adaptive mesh refinement (and even vertical adaptive mesh refinement), but it would require large-eddy simulation resolutions to bypass the need for a planetary boundary layer scheme. Additionally, the efficiency of AMR depends very closely on the choice of criterion chosen for refinement. A criterion that is too strict and hard to fulfill can result in insufficient refinement and lead to less accurate simulations. A criterion that is too lenient and easy to fulfill may result in less time savings and perhaps unnecessary costs. This work does not offer an in-depth study of different criteria for AMR, but we do demonstrate how it can affect the efficiency of a simulation. Deciding on a criterion requires knowledge of specific features that the AMR is meant to capture, and velocity magnitude was chosen for the simulations here as it allows for AMR to detect regions where intensification takes place. However, it could be

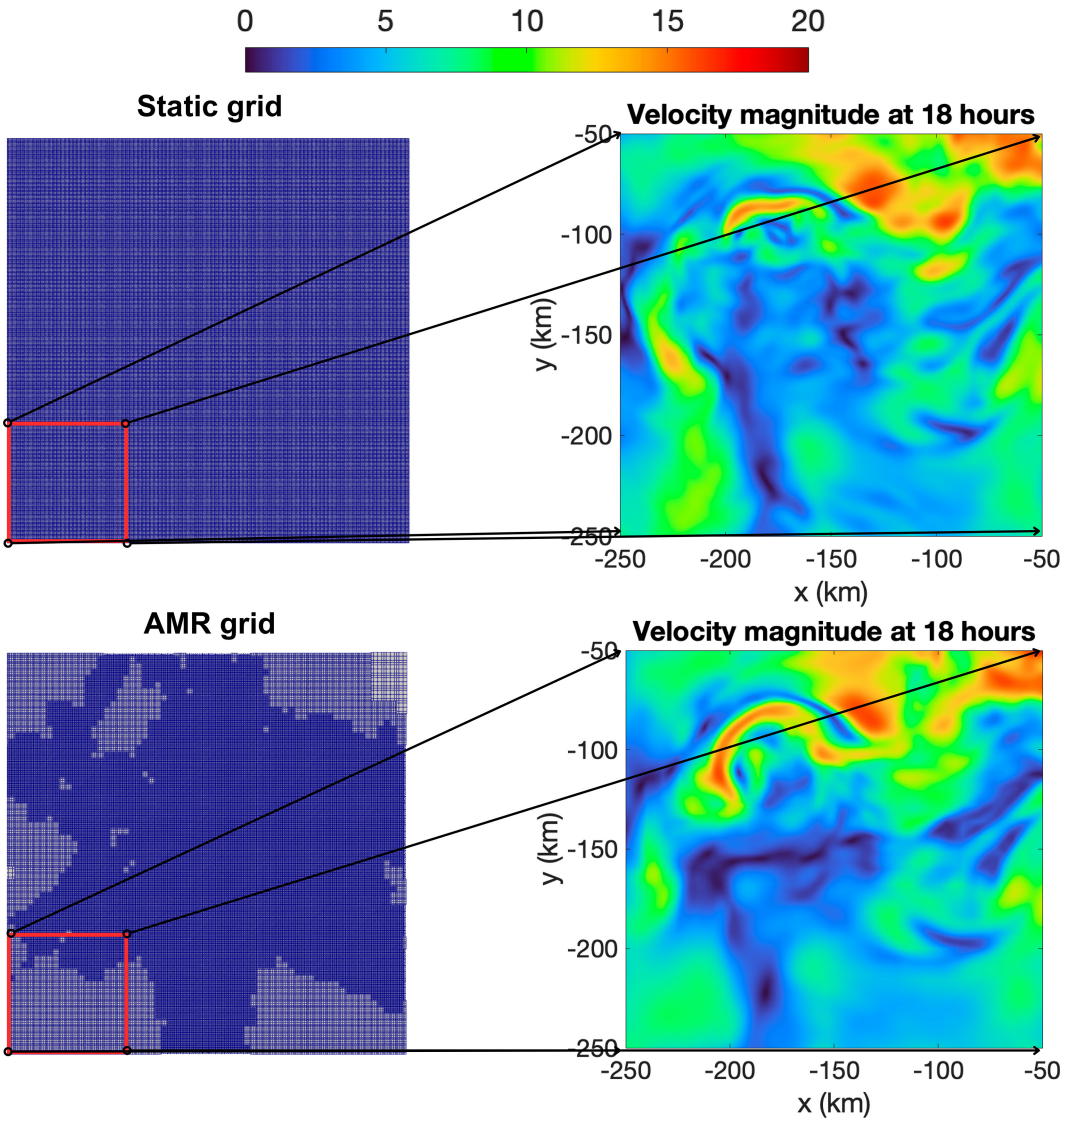

FIG. 21. Comparison of (top left) the static grid with (bottom left) the AMR grid at $t = 18$ h for $(x), (y) \in [-250, 250]$ km $\times$ $[-250, 250]$ km. The red boxes highlight $(x), (y) \in [-250, -50]$ km $\times$ $[-250, -50]$ km. Velocity magnitude of the (top right) static simulation and (bottom right) AMR simulation at height $z = 7000$ m above sea level and $t = 18$ h for $(x), (y) \in [-250, -50]$ km $\times$ $[-250, -50]$ km.

argued that this criterion can only be activated after intensification has already begun, and it might not be able to capture the onset of intensification. Potential vorticity has been studied by Hendricks et al. (2016) as a potential criterion for TC AMR, which showed promise in a simple shallow-water model. Perhaps, a criterion that combines multiple flow variables could be useful for future AMR simulations of TCs. This work presented AMR as a tool for accelerating the simulation of TC rapid intensification. The algorithm used to adapt the grid was presented and its effectiveness demonstrated through a suite of tests on a dry rapidly intensifying TC case. We have demonstrated that the results of simulations using uniform grids can be accurately replicated through the use of AMR. We have also shown that AMR allows for improving the time efficiency of rapidly intensifying TC simulations. Though dependent on the adaptivity criterion, we have shown that AMR simulations can be performed at a fraction of the time it takes to perform their uniform grid counterparts. We have also shown that it is possible to perform high-resolution AMR simulations in comparable and sometimes even less time than it would take to perform a lower-resolution uniform grid simulation. Furthermore, although specific considerations were taken to handle hanging nodes for the spectral elements method used in this work, a suitable AMR algorithm would be applicable to other types of spatial discretizations such as finite volumes

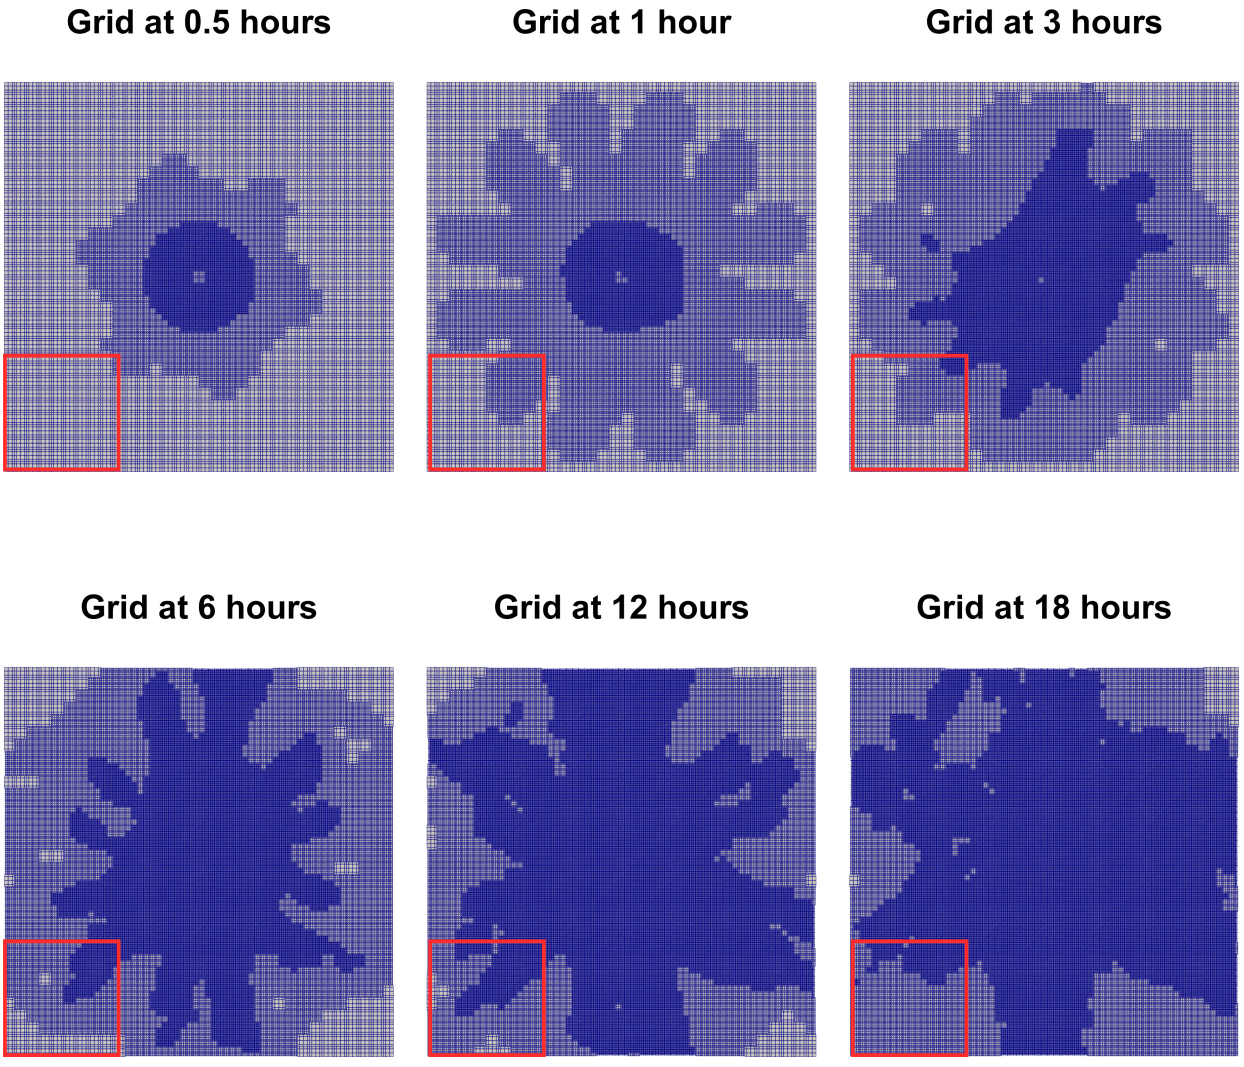

FIG. 22. Snapshots of the AMR grid for the long duration simulation at (top left) $t = 0.5$ h, (top middle) $t = 1$ h, (top right) $t = 3$ h, (bottom left) $t = 6$ h, (bottom middle) $t = 12$ h, and (bottom right) $t = 18$ h.

or finite differences and could yield similar performance improvements.

Acknowledgments. Stephen R. Guimond and Simone Marras acknowledge the partial support by the National Science Foundation through Grants PD-2121366 and PD-2121367 which supported the work by graduate student Yassine Tissaoui. Francis X. Giraldo gratefully acknowledges the support of the Office of Naval Research under Grant N0001419WX00721 and National Science Foundation under Grant AGS-1835881. The authors gratefully acknowledge the contributions of several persons who, over many years, contributed to the AMR capability in NUMA. A brief list includes Andreas Müller, Michal Kopera, and Jeremy Kozdon. Discussions with Lucas Wilcox on P4est also helped to get this work over the finish line.

Data availability statement. The data used to generate all of the figures in this work are publicly available at the following link: https://doi.org/10.5281/zenodo.15291298.

APPENDIX

The Adaptive Mesh Refinement Algorithm

We provide the reader with a pseudocode for the adaptive mesh refinement procedure. Here, t_{final} is the final time of a given simulation, threshold is an array storing the threshold

criteria for each level of refinement, N_{LGL} is the number of LGL points, N_e is the number of elements, and mod is the remainder operator.

Algorithm 1: Algorithm of the adaptive mesh refinement procedure.

for $t = 0, t_{\mathrm{final}}, \Delta t$ **do**
  Check if current time is a multiple of t_{amr}.
  if mod$(t, t_{\mathrm{amr}}) = = 0$ **then**
    for $e = 1, N_e$ **do**
      adapt$(e) = -1$
      elements are marked for coarsening unless they pass at least one AMR threshold.
      for $k = 1, N^3_{\mathrm{LGL}}$ **do**
        Compute local AMR criterion $C_{k,e}$.
        for $l = 1, \mathrm{lvl}_{\mathrm{max}}$ **do**
          if $C^l_{k,e} >$ threshold(kv) **then**
            if lvl$(e) < l$ **then**
              adapt$(e) = 1$
            else
              adapt$(e) = 0$
            end if
          end if
        end for
      end for
    end for
  end if
  Refine and coarsen marked elements through P4est
  Repartition new mesh
end for

REFERENCES

Abdi, D. S., A. Almgren, F. X. Giraldo, and I. Jankov, 2024: Comparison of adaptive mesh refinement techniques for numerical weather prediction. arXiv, 2404.16648v1, https://doi.org/10.48550/arXiv.2404.16648.

Bacon, D. P., and Coauthors, 2000: A dynamically adapting weather and dispersion model: The Operational Multiscale Environment Model with Grid Adaptivity (OMEGA). *Mon. Wea. Rev.*, **128**, 2044–2076, https://doi.org/10.1175/1520-0493(2000)128<2044:ADAWAD>2.0.CO;2.

Berger, M. J., and J. Oliger, 1984: Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.*, **53**, 484–512, https://doi.org/10.1016/0021-9991(84)90073-1.

——, D. A. Calhoun, C. Helzel, and R. J. Leveque, 2009: Logically rectangular finite volume methods with adaptive refinement on the sphere. *Philos. Trans. Roy. Soc.*, **A367**, 4483–4496, https://doi.org/10.1098/rsta.2009.0168.

Burstedde, C., L. C. Wilcox, and O. Ghattas, 2011: p4est: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees. *SIAM J. Sci. Comput.*, **33**, 1103–1133, https://doi.org/10.1137/100791634.

Chen, C., F. Xiao, and X. Li, 2011: An adaptive multimoment global model on a cubed sphere. *Mon. Wea. Rev.*, **139**, 523–548, https://doi.org/10.1175/2010MWR3365.1.

Cram, T. A., J. Persing, M. T. Montgomery, and S. A. Braun, 2007: A Lagrangian trajectory view on transport and mixing processes between the eye, eyewall, and environment using a high-resolution simulation of Hurricane Bonnie (1998). *J. Atmos. Sci.*, **64**, 1835–1856, https://doi.org/10.1175/JAS3921.1.

Doyle, J. D., and Coauthors, 2014: Tropical cyclone prediction using COAMPS-TC. *Oceanography*, **27** (3), 104–115, https://doi.org/10.5670/oceanog.2014.72.

Emanuel, K. A., 1986: An air-sea interaction theory for tropical cyclones. Part I: Steady-state maintenance. *J. Atmos. Sci.*, **43**, 585–605, https://doi.org/10.1175/1520-0469(1986)043<0585:AASITF>2.0.CO;2.

Ferguson, J. O., C. Jablonowski, H. Johansen, P. McCorquodale, P. Collela, 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.

Foster, R. C., 2005: Why rolls are prevalent in the hurricane boundary layer. *J. Atmos. Sci.*, **62**, 2647–2661, https://doi.org/10.1175/JAS3475.1.

Giraldo, F. X., 2020: *An Introduction to Element-based Galerkin Methods on Tensor-Product Bases - Analysis, Algorithms, and Applications.* Vol. 24. Springer, 559 pp.

——, J. F. Kelly, and E. M. Constantinescu, 2013: Implicit-explicit formulations of a three-dimensional Nonhydrostatic Unified Model of the Atmosphere (NUMA). *SIAM J. Sci. Comput.*, **35**, B1162–B1194, https://doi.org/10.1137/120876034.

Guimond, S. R., and J. M. Reisner, 2012: A latent heat retrieval and its effects on the intensity and structure change of Hurricane Guillermo (1997). Part II: Numerical simulations. *J. Atmos. Sci.*, **69**, 3128–3146, https://doi.org/10.1175/JAS-D-11-0201.1.

——, G. M. Heymsfield, and F. J. Turk, 2010: Multiscale observations of Hurricane Dennis (2005): The effects of hot towers on rapid intensification. *J. Atmos. Sci.*, **67**, 633–654, https://doi.org/10.1175/2009JAS3119.1.

——, M. Bourassa, and P. D. Reasor, 2011: A latent heat retrieval and its effects on the intensity and structure change of Hurricane Guillermo (1997). Part I: The algorithm and observations. *J. Atmos. Sci.*, **68**, 1549–1567, https://doi.org/10.1175/2011JAS3700.1.

——, J. M. Reisner, S. Marras, and F. X. Giraldo, 2016a: The impacts of dry dynamic cores on asymmetric hurricane intensification. *J. Atmos. Sci.*, **73**, 4661–4684, https://doi.org/10.1175/JAS-D-16-0055.1.

——, G. M. Heymsfield, P. D. Reasor, and A. C. Didlake Jr., 2016b: The rapid intensification of Hurricane Karl (2010): New remote sensing observations of convective bursts from the Global Hawk platform. *J. Atmos. Sci.*, **73**, 3617–3639, https://doi.org/10.1175/jas-d-16-0026.1.

——, J. A. Zhang, J. W. Sapp, and S. J. Frasier, 2018: Coherent turbulence in the boundary layer of Hurricane Rita (2005) during an eyewall replacement cycle. *J. Atmos. Sci.*, **75**, 3071–3093, https://doi.org/10.1175/JAS-D-17-0347.1.

Guinn, T. A., and W. H. Schubert, 1993: Hurricane spiral bands. *J. Atmos. Sci.*, **50**, 3380–3403, https://doi.org/10.1175/1520-0469(1993)050<3380:HSB>2.0.CO;2.

Hasan, M. B., S. R. Guimond, M. L. Yu, S. Reddy, and F. X. Giraldo, 2022: The effects of numerical dissipation on hurricane rapid intensification with observational heating. *J. Adv. Model. Earth Syst.*, **14**, e2021MS002897, https://doi.org/10.1029/2021MS002897.

Hendricks, E. A., M. T. Montgomery, and C. A. Davis, 2004: The role of "vortical" hot towers in the formation of Tropical Cyclone Diana (1984). *J. Atmos. Sci.*, **61**, 1209–1232, https://

doi.org/10.1175/1520-0469(2004)061%3C1209:TROVHT%3E2.0.CO;2.

——, W. H. Schubert, R. K. Taft, H. Wang, and J. P. Kossin, 2009: Life cycles of hurricane-like vorticity rings. *J. Atmos. Sci.*, **66**, 705–722, https://doi.org/10.1175/2008JAS2820.1.

——, M. A. Kopera, F. X. Giraldo, M. S. Peng, J. D. Doyle, and Q. Jiang, 2016: Evaluation of the utility of static and adaptive mesh refinement for idealized tropical cyclone problems in a spectral element shallow-water model. *Mon. Wea. Rev.*, **144**, 3697–3724, https://doi.org/10.1175/MWR-D-15-0146.1.

Hodur, R., 1997: The Naval Research Laboratory's Coupled Ocean/Atmosphere Mesoscale Prediction System (COAMPS). *Mon. Wea. Rev.*, **125**, 1414–1430, https://doi.org/10.1175/1520-0493(1997)125<1414:TNRLSC>2.0.CO;2.

Hong, S.-Y., 2010: A new stable boundary-layer mixing scheme and its impact on the simulated East Asian summer monsoon. *Quart. J. Roy. Meteor. Soc.*, **136**, 1481–1496, https://doi.org/10.1002/qj.665.

——, Y. Noh, and J. Dudhia, 2006: A new vertical diffusion package with an explicit treatment of entrainment processes. *Mon. Wea. Rev.*, **134**, 2318–2341, https://doi.org/10.1175/MWR3199.1.

Jiménez, P. A., J. Dudhia, J. F. González-Rouco, J. Navarro, J. P. Montávez, and E. García-Bustamante, 2012: A revised scheme for the WRF surface layer formulation. *Mon. Wea. Rev.*, **140**, 898–918, https://doi.org/10.1175/MWR-D-11-00056.1.

Johnson, A. A., J. Schwendike, A. N. Ross, A. Lock, J. M. Edwards, and J. D. Kepert, 2024: Impacts of free tropospheric turbulence parametrisation on a sheared tropical cyclone. *Quart. J. Roy. Meteor. Soc.*, **150**, 4437–4456, https://doi.org/10.1002/qj.4823.

Jordan, C. L., 1958: Mean soundings for the West Indies area. *J. Meteor.*, **15**, 91–97, https://doi.org/10.1175/1520-0469(1958)015<0091:MSFTWI>2.0.CO;2.

Klein, R., 2010: Scale-dependent models for atmospheric flows. *Annu. Rev. Fluid Mech.*, **42**, 249–274, https://doi.org/10.1146/annurev-fluid-121108-145537.

——, U. Achatz, D. Bresch, O. M. Knio, and P. K. Smolarkiewicz, 2010: Regime of validity of soundproof atmospheric flow models. *J. Atmos. Sci.*, **67**, 3226–3237, https://doi.org/10.1175/2010JAS3490.1.

Kopera, M. A., and F. X. Giraldo, 2014a: Analysis of adaptive mesh refinement for IMEX discontinuous Galerkin solutions of the compressible Euler equations with application to atmospheric simulations. *J. Comput. Phys.*, **275**, 92–117, https://doi.org/10.1016/j.jcp.2014.06.026.

——, and ——, 2014b: Mass conservation of the unified continuous and discontinuous element-based Galerkin methods on dynamically adaptive grids with application to atmospheric simulations. *J. Comput. Phys.*, **297**, 90–103, https://doi.org/10.1016/j.jcp.2015.05.010.

Kopriva, D., 2008: *Implementing Spectral Methods for Partial Differential Equations.* Springer, 394 pp.

Kossin, J. P., and W. H. Schubert, 2001: Mesovortices, polygonal flow patterns, and rapid pressure falls in hurricane-like vortices. *J. Atmos. Sci.*, **58**, 2196–2209, https://doi.org/10.1175/1520-0469(2001)058<2196:MPFPAR>2.0.CO;2.

Kurihara, Y., G. J. Tripoli, and M. A. Bender, 1979: Design of a movable nested-mesh primitive equation model. *Mon. Wea. Rev.*, **107**, 239–249, https://doi.org/10.1175/1520-0493(1979)107<0239:DOAMNM>2.0.CO;2.

——, R. E. Tuleya, and M. A. Bender, 1998: The GFDL Hurricane Prediction System and its performance in the 1995 hurricane season. *Mon. Wea. Rev.*, **126**, 1306–1322, https://doi.org/10.1175/1520-0493(1998)126<1306:TGHPSA>2.0.CO;2.

LeVeque, R. J., D. L. George, and M. J. Berger, 2011: Tsunami modelling with adaptively refined finite volume methods. *Acta Numer.*, **20**, 211–289, https://doi.org/10.1017/S0962492911000043.

Li, J., F. Fang, J. Steppeler, J. Zhu, Y. Cheng, and X. Wu, 2021: Demonstration of a three-dimensional dynamically adaptive atmospheric dynamic framework for the simulation of mountain waves. *Meteor. Atmos. Phys.*, **133**, 1627–1645, https://doi.org/10.1007/s00703-021-00828-8.

Lilly, D. K., 1962: On the numerical simulation of buoyant convection. *Tellus*, **14** (2), 148–172, https://doi.org/10.3402/tellusa.v14i2.9537.

McCorquodale, P., P. A. Ulrich, 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.

Molinari, J., and D. Vollaro, 2010: Rapid intensification of a sheared tropical storm. *Mon. Wea. Rev.*, **138**, 3869–3885, https://doi.org/10.1175/2010MWR3378.1.

Montgomery, M. T., and R. J. Kallenbach, 1997: A theory for vortex Rossby-waves and its application to spiral bands and intensity changes in hurricanes. *Quart. J. Roy. Meteor. Soc.*, **123**, 435–465, https://doi.org/10.1002/qj.49712353810.

Mueller, A., J. Behrens, F. X. Giraldo, and V. Wirth, 2013: Comparison between adaptive and uniform deiscontinuous Galerkin simulations in 2D dry bubble experiments. *J. Comput. Phys.*, **235**, 371–393, https://doi.org/10.1016/j.jcp.2012.10.038.

Nolan, D. S., 2011: Evaluating environmental favorableness for tropical cyclone development with the method of point-downscaling. *J. Adv. Model. Earth Syst.*, **3**, M08001, https://doi.org/10.1029/2011MS000063.

——, and L. D. Grasso, 2003: Nonhydrostatic, three-dimensional perturbations to balanced, hurricane-like vortices. Part II: Symmetric response and nonlinear simulations. *J. Atmos. Sci.*, **60**, 2717–2745, https://doi.org/10.1175/1520-0469(2003)060<2717:NTPTBH>2.0.CO;2.

——, Y. Moon, and D. P. Stern, 2007: Tropical cyclone intensification from asymmetric convection: Energetics and efficiency. *J. Atmos. Sci.*, **64**, 3377–3405, https://doi.org/10.1175/JAS3988.1.

Orlando, G., T. Benacchio, and L. Bonaventura, 2024: Robust and accurate simulations of flows over orography using nonconforming meshes. arXiv, 2402.07759v6, https://doi.org/10.48550/arXiv.2402.07759.

Rogers, R., P. Reasor, and S. Lorsolo, 2013: Airborne Doppler observations of the inner-core structural differences between intensifying and steady-state tropical cyclones. *Mon. Wea. Rev.*, **141**, 2970–2991, https://doi.org/10.1175/MWR-D-12-00357.1.

Rogers, R. F., and Coauthors, 2017: Rewriting the tropical record books: The extraordinary intensification of Hurricane Patricia (2015). *Bull. Amer. Meteor. Soc.*, **98**, 2091–2112, https://doi.org/10.1175/BAMS-D-16-0039.1.

Schubert, W. H., M. T. Montgomery, R. K. Taft, T. A. Guinn, S. R. Fulton, J. P. Kossin, and J. P. Edwards, 1999: Polygonal eyewalls, asymmetric eye contraction, and potential vorticity mixing in hurricanes. *J. Atmos. Sci.*, **56**, 1197–1223, https://doi.org/10.1175/1520-0469(1999)056<1197:PEAECA>2.0.CO;2.

Skamarock, W., J. Oliger, and R. L. Street, 1989: Adaptive grid refinement for numerical weather prediction. *J. Comput. Phys.*, **80**, 27–60, https://doi.org/10.1016/0021-9991(89)90089-2.

Skamarock, W. C., and J. B. Klemp, 1993: Adaptive grid refinement for two-dimensional and three-dimensional nonhydrostatic atmospheric flow. *Mon. Wea. Rev.*, **121**, 788–804, https://doi.org/10.1175/1520-0493(1993)121%3C0788:AGRFTD%3E2.0.CO;2.

——, ——, J. Dudhia, D. O. Gill, D. M. Barker, W. Wang, and J. G. Powers, 2005: A description of the Advanced Research WRF version 2. NCAR Tech. Note NCAR/TN-468+STR, 88 pp., https://doi.org/10.5065/D6DZ069T.

Smagorinsky, J., 1963: General circulation experiments with the primitive equations: I. The basic experiement. *Mon. Wea. Rev.*, **91**, 99–164, https://doi.org/10.1175/1520-0493(1963)091%3C0099:GCEWTP%3E2.3.CO;2.

Wang, Y., 2002: Vortex Rossby waves in a numerically simulated tropical cyclone. Part I: Overall structure, potential vorticity, and kinetic energy budgets. *J. Atmos. Sci.*, **59**, 1213–1238, https://doi.org/10.1175/1520-0469(2002)059<1213:VRWIAN>2.0.CO;2.

Wu, L., B. Wang, and S. A. Braun, 2005: Impacts of air–sea interaction on tropical cyclone track and intensity. *Mon. Wea. Rev.*, **133**, 3299–3314, https://doi.org/10.1175/MWR3030.1.

Yamazaki, H., H. Weller, C. J. Cotter, and P. A. Browne, 2022: Conservation with moving meshes over orography. *J. Comput. Phys.*, **461**, 111217, https://doi.org/10.1016/j.jcp.2022.111217.