Two-Dimensional Evaluation of ATHAM-Fluidity, a Nonhydrostatic Atmospheric Model Using Mixed Continuous/Discontinuous Finite Elements and Anisotropic Grid Optimization

JULIEN SAVRE

Department of Geography, University of Cambridge, Cambridge, United Kingdom

JAMES PERCIVAL

Department of Earth Science and Engineering, Royal School of Mines, Imperial College London, London, United Kingdom

MICHAEL HERZOG

Department of Geography, University of Cambridge, Cambridge, United Kingdom

CHRIS PAIN

Department of Earth Science and Engineering, Royal School of Mines, Imperial College London, London, United Kingdom

(Manuscript received 15 November 2015, in final form 1 June 2016)

ABSTRACT

This paper presents the first attempt to apply the compressible nonhydrostatic Active Tracer High-Resolution Atmospheric Model–Fluidity (ATHAM-Fluidity) solver to a series of idealized atmospheric test cases. ATHAM-Fluidity uses a hybrid finite-element discretization where pressure is solved on a continuous second-order grid while momentum and scalars are computed on a first-order discontinuous grid (also known as $\mathbb{P}_{1\mathrm{DG}} - \mathbb{P}_2$). ATHAM-Fluidity operates on two- and three-dimensional unstructured meshes, using triangular or tetrahedral elements, respectively, with the possibility to employ an anisotropic mesh optimization algorithm for automatic grid refinement and coarsening during run time. The solver is evaluated using two-dimensional-only dry idealized test cases covering a wide range of atmospheric applications. The first three cases, representative of atmospheric convection, reveal the ability of ATHAM-Fluidity to accurately simulate the evolution of large-scale flow features in neutral atmospheres at rest. Grid convergence without adaptivity as well as the performances of the Hermite–Weighted Essentially Non-oscillatory (Hermite-WENO) slope limiter are discussed. These cases are also used to test the grid optimization algorithm implemented in ATHAM-Fluidity. Adaptivity can result in up to a sixfold decrease in computational time and a fivefold decrease in total element number for the same finest resolution. However, substantial discrepancies are found between the uniform and adapted grid results, thus suggesting the necessity to improve the reliability of the approach. In the last three cases, corresponding to atmospheric gravity waves with and without orography, the model ability to capture the amplitude and propagation of weak stationary waves is demonstrated. This work constitutes the first step toward the development of a new comprehensive limited area atmospheric model.

1. Introduction

Despite the development of highly scalable massively parallel codes for atmospheric modeling [including general circulation models (GCMs) and limited area models such as cloud-resolving models (CRMs)], we are still not able to accurately resolve all physical scales involved in the climate and weather systems. For example, whereas large cloud systems such as tropical or midlatitude cyclones operate on scales of several hundreds to thousands of kilometers, cloud-resolving simulations, designed to follow the evolution of individual clouds, require spatial resolutions of only a few tens of meters. While the emergence of extremely powerful high-performance

Corresponding author address: Julien Savre, Department of Geography, University of Cambridge, Downing Place, Cambridge CB2 3EN, United Kingdom.
E-mail: js2176@cam.ac.uk

DOI: 10.1175/MWR-D-15-0398.1

© 2016 American Meteorological Society

computing resources allows for an increase of the typical affordable grid resolution for weather forecasting and climate predictions, the simultaneous increase of the complexity and the subsequent increase in CPU demand of the necessary physical parameterizations tend to slow down the performance improvements we could normally expect. It now appears clear that sustaining the trend toward an increase of the affordable spatial resolution will not be possible without completely rethinking our existing models.

Recent reviews have pointed out the necessity to develop new-generation atmospheric models, using state-of-the-art numerical methods, to adequately capture all the physical processes needed for a complete representation of our climate system (Slingo et al. 2009; Marras et al. 2016). In particular, it has been argued that increasing the flexibility offered by the numerical grids used in atmospheric models will be crucial to improve the representation of the various spatial and physical scales involved (Williamson 2007; Slingo et al. 2009; Staniforth and Thuburn 2012). This implies the development of highly scalable models supporting irregular grids [e.g., global icosahedral (Giraldo and Warburton 2005; Satoh et al. 2008; Skamarock et al. 2012) or cubed-sphere grids (Nair et al. 2005; Chen et al. 2011; Harris and Lin 2013; Ullrich 2014; Staniforth and Thuburn 2012)] or adaptive remeshing techniques (Behrens et al. 2005; Jablonowski et al. 2006; St Cyr et al. 2008; Weller 2009; Müller et al. 2013; Yelash et al. 2014; McCorquodale et al. 2015) (and possibly both). Adaptive remeshing methods in particular allow for the focus of computational efforts on areas where physical processes occur at small spatial scales and have been considered as a viable approach for operational weather prediction models for more than three decades.

Implementing non-Cartesian unstructured grids and adaptive remeshing techniques in atmospheric models is, however, not a trivial task. Advanced numerical methods are required for which the discretized equations can be formulated in a general framework while preserving important stability and accuracy properties. The numerical methods used to solve the basic flow equations are still often based on finite-difference methods, but if one wants to efficiently take advantage of advanced meshing techniques, numerical discretizations have to be redeveloped in consequence (Marras et al. 2016). In this context, both finite-volume and finite-element methods (FEMs) emerge as good candidates to solve atmospheric flows on irregular, adaptive grids because of their overall flexibility, high scalability, and excellent conservation properties. In particular, FEM with inexact integration or mass lumping (including Galerkin and high-order spectral element methods) is becoming increasingly popular among atmospheric and climate modelers, and some of the most recent atmospheric solvers (both global and limited area) rely on such techniques (Giraldo et al. 2002; Nair et al. 2005; Thomas et Loft 2005; Giraldo and Restelli 2008; Nair et al. 2009; Kelly and Giraldo 2012; Müller et al. 2013; Kopera and Giraldo 2014; Marras et al. 2015). However, Galerkin methods with exact integration and nondiagonal mass matrix have not previously been adopted, perhaps because of the extensive computational cost of solving the implicit system associated with them. With the increased availability of high-performance computing facilities, Galerkin methods with exact integration are just starting to receive interest from the atmospheric modeling community (Brdar et al. 2013; Schuster et al. 2014; Choi et al. 2014; Thuburn and Cotter 2015).

In the following, we introduce a new modeling framework for the simulation of atmospheric processes based on a mixed continuous/discontinuous Galerkin (CG/DG) finite-element discretization. The model also includes anisotropic adaptive remeshing that allows for modifications of both the connectivities between grid cells and the position of grid vertices (*hr* adaptivity). The new model, Active Tracer High-Resolution Atmospheric Model–Fluidity (ATHAM-Fluidity), combines the mixed FEM dynamical core from Fluidity (Ford et al. 2004; Piggott et al. 2009) with the physical package and active tracer concept from the ATHAM (Oberhuber et al. 1998). The numerical discretization follows the approach suggested by Cotter and Ham (2011) and Cotter et al. (2009a,b), where governing equations are solved on triangular meshes using linear discontinuous elements for momentum and scalars (denoted $\mathbb{P}_{1\mathrm{DG}}$) but continuous quadratic elements for pressure and density ($\mathbb{P}_2$). A semi-implicit compressible pressure projection method is then used to diagnose the pressure at each time-level by inverting a general Helmholtz equation devised to satisfy the continuity equation. This class of methods is known to perform well in both low-Mach and high-Mach number regimes (although this latter may not be relevant for atmospheric applications), while relaxing the severe time step restriction typically fixed by fast-propagating acoustic waves in density-based solvers [for which a continuity equation is solved explicitly, as in Giraldo and Restelli (2008) and Kopera and Giraldo (2014)]. In this context, using a mixed FEM formulation with the velocity possessing more or equal degrees of freedom than the pressure is essential to prevent the formation of spurious numerical pressure modes (Cotter et al. 2009a; Botti and Di Pietro 2011).

So far, the $\mathbb{P}_{1\mathrm{DG}} - \mathbb{P}_2$ element pair has mostly been used in the geophysical fluid dynamics community to solve the shallow water equations (Cotter et al. 2009a,b;

Cotter and Ham 2011; Düben et al. 2012). However, Cotter and Shipton (2012) recently drew the comparison between the $\mathbb{P}_{1DG} - \mathbb{P}_2$ pair and the popular Arakawa C grid and concluded on the suitability of such mixed FEM methods for numerical weather prediction and atmospheric modeling in general. The present study therefore builds on Cotter and Shipton (2012) and proposes an extension of the $\mathbb{P}_{1DG} - \mathbb{P}_2$ discretization to solve generalized compressible governing equations for atmospheric flows over limited area domains.

In the following, the dynamical core employed in ATHAM-Fluidity is introduced in section 2. The grid adaptivity procedure is then briefly discussed in section 3. In section 4 we evaluate the model's performances without grid adaptivity based on six elementary test cases (three atmospheric bubble-like cases and three gravity wave cases), commonly used to assess the numerics of new atmospheric models. We have restricted our study to dry atmosphere simulations only. Preliminary results obtained with grid optimization under atmospheric convection conditions (bubble-like test cases) are then shown in section 5. Conclusions are given in section 6.

2. The dynamical core

a. Governing equations

The Fluidity dynamical core solves a general set of fully compressible governing equations, including equations for the density (continuity equation),

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0; \tag{1}$$

momentum,

$$\rho \frac{\partial \mathbf{u}}{\partial t} + \rho \mathbf{u} \cdot \nabla \mathbf{u} = -2\boldsymbol{\Omega} \times (\mathbf{u} - \mathbf{u}_g) - \rho g \mathbf{k} + \nabla \cdot \boldsymbol{\sigma}; \tag{2}$$

and potential temperature (typically used in atmospheric applications as a proxy for energy),

$$\rho \frac{\partial \Theta}{\partial t} + \rho \mathbf{u} \cdot \nabla \Theta = -\nabla Q + S_\Theta - \nabla \cdot \boldsymbol{\tau}_\Theta. \tag{3}$$

The above set of equations is formulated in nonconservative form, where ρ is the density, p the pressure, $\mathbf{u}$ the velocity vector, and Θ the potential temperature. In the momentum equation, g is the gravitational constant; $\mathbf{k} = (0, 0, 1)^T$ is the unit vector defining the vertical direction; $\boldsymbol{\Omega}$ is the Earth's angular velocity vector; $\mathbf{u}_g$ represents the geostrophic wind vector; and $\boldsymbol{\sigma} = -p\mathbf{I} + \boldsymbol{\tau}$ is the stress tensor where p is the pressure, $\mathbf{I}$ is the identity matrix, and $\boldsymbol{\tau}$ is the deviatoric part of the tensor (viscous term). The term $\boldsymbol{\tau}_\Theta$ can be expressed in analogy with the stress tensor as $\mu/\mathrm{Pr}(\nabla \cdot \boldsymbol{\Theta})\mathbf{I}$, with Pr being the Prandtl number and μ being the dynamic viscosity. Additional terms related to subgrid-scale turbulence modeling have been omitted here (none of the cases considered in section 4 requires such parameterization). Finally, ∇Q represents external heating/cooling rates while S_Θ represents additional sources and sinks, including, for example, microphysical processes in clouds. In the above set of equations, no external force has been considered and molecular diffusion has been omitted.

The thermodynamic quantities are related via the equation of state for ideal gas:

$$p = p_0 \left(\frac{\rho R \Theta}{p_0} \right)^{c_p/c_v}, \tag{4}$$

where $p_0 = 1000$ hPa is a reference pressure, $R = c_p - c_v$ is the specific gas constant, and c_p and c_v are the specific heat capacities for dry air at constant pressure and volume, respectively. The potential temperature is defined in terms of primitive variables following $\Theta = T(p/p_0)^{-R/c_p}$.

The system is similar to equation set 2 in Giraldo and Restelli (2008), but written in nonconservative form. During the solution procedure, however, total mass is conserved, as Eq. (1) is used to devise the pressure projection method employed to diagnose the pressure field (see section 2d). In contrast, the potential temperature (used as a proxy for internal energy) is not conserved locally when solving Eq. (3). Errors related to energy conservation, however, are not expected to severely affect the solution because of the relatively short integration times typical of large-eddy simulation applications (less than a day).

b. Spatial discretization

We first recall here that ATHAM-Fluidity employs a mixed discontinuous/continuous FEM discretization where the momentum and scalar equations are solved on first-order linear discontinuous elements while the pressure and density are discretized on second-order parabolic elements ($\mathbb{P}_{1DG} - \mathbb{P}_2$). This approach conserves the high accuracy of DG methods (Comblen et al. 2010), while being more computationally efficient than fully discontinuous schemes and preventing the development of spurious pressure modes generated on collocated grids with a pressure projection procedure (Cotter and Shipton 2012). According to Cotter et al. (2009a), Cotter and Ham (2011), and Cotter and Shipton (2012), this latter property of the $\mathbb{P}_{1DG} - \mathbb{P}_2$ discretization is similar in principle to conventional staggered

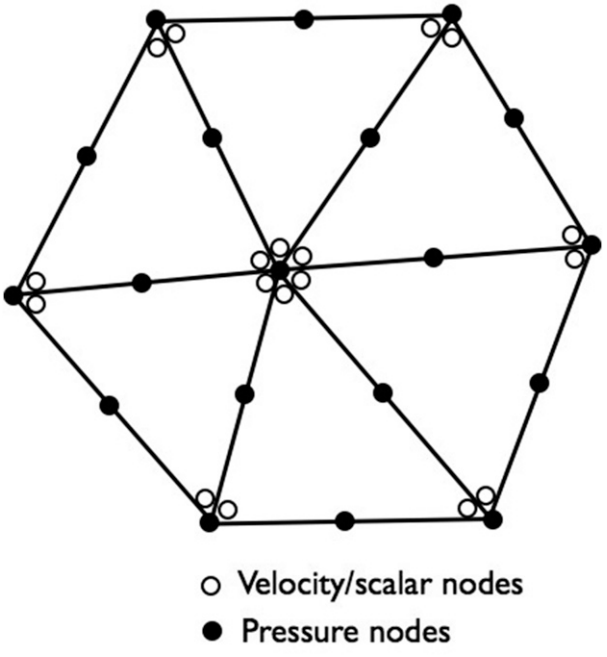


FIG. 1. Velocity/scalar and pressure nodes repartition on the triangular mesh.

Arakawa C grids used in many global and regional atmospheric models. More specifically, the use of unbalanced numbers of degrees of freedom (DOF) per grid cell between the pressure and velocity discretizations, with $N_{\mathrm{DOF}}(\mathbf{u}) \geq N_{\mathrm{DOF}}(p)$, was proved to be an important condition, although not sufficient, to prevent the propagation of spurious pressure modes polluting the numerical solution (Cotter et al. 2009a).

In ATHAM-Fluidity, the domains are discretized using triangular (in 2D, tetrahedra in 3D) elements in a way that the velocity/scalar and pressure nodes are distributed as shown in Fig. 1. At each node, the discontinuous quantities can possess up to six different values (according to Fig. 1).

In the following, we focus on the description of the DG discretization used for momentum and scalar advection. Although its use in atmospheric sciences is still relatively new, with pioneering works published in the early 2000s by Giraldo et al. (2002) and Nair et al. (2005), an extensive body of literature exists presenting the basics of the DG method to solve the Euler or Navier–Stokes equations (e.g., Bassi and Rebay 1997; Cockburn and Shu 2001). Since the dynamical core implemented in ATHAM-Fluidity mainly follows commonly described DG methods, we restrict our presentation to the essential aspects of DG as well as to the specifics of the present solver induced by the mixed finite element and pressure projection methods.

1) THE DG DISCRETIZATION

For simplicity, we consider here only the advection–diffusion equation for an arbitrary scalar q written in the following form:

$$\frac{\partial q}{\partial t} + \mathbf{u} \cdot \nabla q - \nabla \cdot (\boldsymbol{\mu} \cdot \nabla q) = 0. \tag{5}$$

Here, we have omitted additional sources and sinks and expended the deviatoric stress tensor with $\boldsymbol{\mu} = \mu/\mathrm{Pr}\mathbf{I}$. Equation (5) can then be rewritten in its weak form by multiplying with a predefined test function φ and integrating over the whole domain Ω:

$$\int_{\Omega} \varphi \left[\frac{\partial q}{\partial t} + \mathbf{u} \cdot \nabla q - \nabla \cdot (\boldsymbol{\mu} \cdot \nabla q) \right] = 0. \tag{6}$$

In DG methods, the test function φ is continuous over each element but can be discontinuous at the interface between two elements. The domain Ω is then decomposed into N_e nonoverlapping elements of arbitrary shape (triangles in the following 2D examples) able to cover the entire domain. Over each element, the quantity q is approximated by a linear combination of a finite number of functions:

$$q(\mathbf{x}; t) = \sum_{i=1}^{N_t} \underline{q}^i(t) \psi^i(\mathbf{x}), \tag{7}$$

where the $\underline{q}^i$ coefficients are defined locally at the nodes; ψ^i are the trial functions, defined continuously over each element; and N_t is the dimension of the trial function space. The test φ and trial ψ^i functions have not been defined yet. In Galerkin methods, they are selected in the same basis function space. They will both be denoted φ for simplicity.

Equation (6) can be integrated by parts over each element e, with test functions possibly discontinuous at the cell interface, yielding

$$\int_e \left[\varphi \frac{\partial q}{\partial t} + q \nabla(\varphi \cdot \mathbf{u}) - (\boldsymbol{\mu} \cdot \nabla q) \cdot \nabla \varphi \right] + \int_{\partial e} (\varphi \widehat{\mathbf{n} \cdot \mathbf{u} q} - \varphi \mathbf{n} \cdot \widehat{\boldsymbol{\mu} \cdot \nabla q}) = 0. \tag{8}$$

The second term in the equation represents integration over the element boundary ∂e and the hatted terms represent fluxes across this boundary, with $\mathbf{n}$ being the outward-pointing normal unit vector at the interface. This term can be decomposed into contributions from boundaries between internal elements and element boundaries belonging to the domain boundaries. Considering, for example, the imposition of Dirichlet conditions at

the domain boundary denoted $\partial\Omega^D$, the interface term in Eq. (8) can be recast into (dropping the viscous contribution)

$$\int_{\partial e} \varphi \widehat{\mathbf{n}\cdot\mathbf{u}q} = \int_{\partial e/\partial\Omega} \varphi \widehat{\mathbf{n}\cdot\mathbf{u}q}_{\text{int}} + \int_{\partial e\cap\partial\Omega^D} \varphi \widehat{\mathbf{n}\cdot\mathbf{u}q}_D, \quad (9)$$

where q_{int} represents internal values of q and q_D represents the values imposed at the domain boundary. The fully discretized version of Eq. (8) is obtained by replacing q using Eq. (7) with summation over the basis function space.

A critical aspect of DG methods resides in the evaluation of the flux terms, that is, $\widehat{\mathbf{n}\cdot\mathbf{u}q}$, representing mass and energy exchanges between two adjacent elements. Because quantities might by definition not be continuous and therefore uniquely defined at the interface, evaluating the numerical fluxes requires properly defined quantities at the interface. Because a continuous discretization is used for the pressure and density fields, a corresponding continuous velocity field $\hat{\mathbf{u}}$ can be obtained by projection on the pressure mesh. Scalar fluxes are then evaluated using a classical upwind flux formulation using the projected velocity:

$$\widehat{\mathbf{n}\cdot\mathbf{u}q} = \begin{cases} \mathbf{n}\cdot\hat{\mathbf{u}}q^- & \text{for} \quad \mathbf{n}\cdot\hat{\mathbf{u}} > 0 \\ \mathbf{n}\cdot\hat{\mathbf{u}}q^+ & \text{for} \quad \mathbf{n}\cdot\hat{\mathbf{u}} < 0 \end{cases}, \quad (10)$$

where q^- and q^+ represent scalar element values on the left and right sides of the element boundary, respectively. The upwind flux is the simplest approach satisfying the minimal requirements for DG fluxes, that is, consistency and conservation. Monotonicity is enforced by the slope limiter described in section 3.

The basis functions are typically chosen in the space of polynomials of degree $n \leq N$, continuous over each element e. This approach is usually referred to as a $\mathbb{P}_N$ (for CG) or $\mathbb{P}_{\text{NDG}}$ (for DG) discretization on triangular elements. In DG methods, no additional constraint on the basis function space is required, and it is not necessarily the same for all the prognostic variables. In the examples presented in section 4, the governing equations are discretized using discontinuous piecewise linear basis functions $N = 1$.

2) Discretization of the diffusion operator

In the scalar discretization case, the diffusion operator appearing in Eq. (8), namely, $\boldsymbol{\mu}\cdot\nabla q\cdot\nabla\varphi$, involves the integration of a second-order derivative. The problem is circumvented in the local DG approach by defining a vector field $\boldsymbol{\xi} = \nabla q$, which allows us to rewrite Eq. (5) as follows:

$$\boldsymbol{\xi} - \nabla q = 0$$

$$\frac{\partial q}{\partial t} + \mathbf{u}\cdot\boldsymbol{\xi} - \nabla\cdot(\boldsymbol{\mu}\cdot\boldsymbol{\xi}) = 0. \quad (11)$$

The diffusion operator appearing in Eq. (8) is then discretized on element e following

$$\int_e \nabla\varphi\cdot\boldsymbol{\xi} + \int_{\partial e} \varphi\mathbf{n}\cdot\hat{\boldsymbol{\xi}} \quad (12)$$

(the tensor field $\boldsymbol{\mu}$ has been dropped for simplicity). Evaluating properly the face values $\hat{\boldsymbol{\xi}}$ is crucial to solve the problem: the compact DG method proposed by Peraire and Persson (2008) is used in ATHAM-Fluidity to minimize the discretization stencil while preserving the stability and high-order accuracy of the solution.

Extending this procedure to discretize the viscous stress in the momentum equation is straightforward and involves the auxiliary tensor field $\boldsymbol{\zeta} = \nabla\mathbf{u}$. Depending on the form taken by the viscous term, $\boldsymbol{\zeta}$ can alternatively be defined using the stress tensor $\boldsymbol{\tau}$ instead of the velocity gradient. However, this option seems less computationally efficient (Persson and Peraire 2006), so that $\boldsymbol{\zeta}$ in practice is always related to the velocity gradient only.

3) The Hermite-WENO slope limiter

By virtue of Godunov's theorem, we may expect any numerical advection scheme of order >1 to generate new numerical extrema in the presence of steep gradients. If no particular care is taken during time integration, we can expect the method to produce spurious numerical artifacts and yield unbounded solutions at element interfaces. In ATHAM-Fluidity, the Hermite–Weighted Essentially Nonoscillatory (Hermite-WENO) slope limiter (Qiu and Shu 2005) is employed as a postprocessing filter after the main time integration stage (see section 2c) to smooth numerical oscillations generated by the advection step.

The Hermite-WENO limiter makes use of the WENO interpolation method, originally used to obtain high-order nonoscillatory fluxes for finite volume methods, to reconstruct the solution in elements where unbounded solutions are produced. The Hermite-WENO limiter constructs a series of high-order polynomials based on the cell-averaged scalar values in the neighboring elements and produces a smooth interpolated solution at the cell interface based on weighted averages of the solutions given by each polynomial. The weights are defined based on a nonoscillatory criterion ensuring that the largest weights are assigned to the polynomials producing the smoothest solutions at the interface

[see Jiang and Shu (1996) for further details on the WENO method].

The polynomial reconstruction can become very tedious at high orders when an increasing number of neighboring elements must be considered. The reconstruction step can, however, be greatly simplified by using element gradients in addition to the averaged element values from the direct neighbors only (Hermite interpolation).

Among other possible limiters, the Hermite-WENO limiter has been retained, as it was found to give the best compromise between spurious numerical diffusion and shape preservation of the solution in the numerical examples presented in section 4.

c. Time discretization

Using Eqs. (6) and (7), and choosing the test and trial functions in the same basis function space, Eq. (8) reduces to the following matrix equation for each $\underline{q}$ coefficient:

$$\mathbf{M}\frac{\partial \underline{q}}{\partial t} + \mathbf{A}(\mathbf{u})\underline{q} + \mathbf{T}\underline{q} = \mathbf{R}, \tag{13}$$

where $\mathbf{M}$ is the mass matrix; $\mathbf{A}$ is the advection matrix depending on the velocity vector; $\mathbf{T}$ is the diffusion operator; and $\mathbf{R}$ represents all the residual terms, including boundary terms (sponge layers), buoyancy, and other sources. In practice, a sequential first-order splitting is employed where advection along with sources are first solved together using the classical trapezoidal rule, while the diffusion operator is solved separately. After the first stage, the updated scalar, denoted $\underline{q}^*$, is filtered using the monotonicity preserving slope limiter described in section 3 to yield bounded solutions and improve stability of the diffusion stage.

The full time integration sequence reads

$$\begin{aligned}
\mathbf{M}\frac{\underline{q}^* - \underline{q}^n}{\Delta t} + \mathbf{A}(\mathbf{u}^{n+\theta_u})\underline{q}^{n+\theta_q} &= \mathbf{R}^{n+\theta_q},\\
\underline{\tilde{q}} &= \mathbf{F}\underline{q}^*,\\
\mathbf{M}\frac{\underline{q}^{n+1} - \underline{\tilde{q}}}{\Delta t} + \mathbf{T}\underline{\tilde{q}}^{n+\theta_q} &= 0,
\end{aligned} \tag{14}$$

where θ_q (θ_u) is the integration parameter for q (u) varying between 0 (forward Euler explicit) and 1 (backward Euler implicit) with

$$\underline{q}^{n+\theta_q} = \theta_q\underline{q}^* + (1-\theta_q)\underline{q}^n, \tag{15}$$

$$\underline{\tilde{q}}^{n+\theta_q} = \theta_q\underline{q}^{n+1} + (1-\theta_q)\underline{\tilde{q}}. \tag{16}$$

In the above, $\underline{q}^*$ represents the intermediate solution after the first integration stage, $\underline{\tilde{q}}$ is the filtered intermediate solution, and $\mathbf{F}$ is the slope limiting operator described in section 3.

In DG methods, the mass and advection matrices are generally sparse, which enables the use of efficient linear solvers. In all the simulations presented in section 4, we chose $\theta_q = \theta_u = 0.55$. Allowing the scheme to be slightly skewed backward can substantially improve the stability of the numerical solution (the use of off-centered time integration was notably found to be necessary to stabilize the finest warm bubble simulations shown in section 4a). Alternatively, in case explicit time integration is used, the first advection step can be subcycled, with the number of subcycles determined based on an appropriate target Courant–Friedrichs–Levy (CFL) number.

d. The pressure projection method

A compressible pressure projection (or pressure correction) method is currently used in ATHAM-Fluidity to determine the pressure–velocity coupling. The algorithm is based on the classical Semi-Implicit Method for Pressure Linked Equations–Consistent (SIMPLEC) iterative scheme initially designed for incompressible flows but extended to the compressible case and to all Mach numbers by Karki and Patankar (1989). Despite requiring more efforts than density-based solvers, the class of SIMPLEC projection methods still provides nonnegligible advantages: 1) mass conservation is strongly enforced at all flow speeds, 2) the solution procedure is robust at all flow speeds, and 3) the semi-implicit formulation of the pressure correction equation alleviates the stringent CFL criterion typically imposed by fast moving acoustic waves. Note that in the present formulation, only acoustic waves are treated semi-implicitly, with gravity waves being treated explicitly [semi-implicit treatment of gravity waves has been adopted, for example, in Smolarkiewicz et al. (2014)]. The reader is referred to Botti and Di Pietro (2011) for the implementation and stability of pressure-correction methods combined with DG.

In compressible pressure projection methods, the pressure and density fields are diagnosed to satisfy the full continuity equation at each time level [unlike projection methods used in the context of ''soundproof'' equation systems (Smolarkiewicz et al. 2014)]. As a preliminary step, the momentum equation is solved using an approximate pressure gradient and the density is computed using the equation of state at the previous time level. At this stage, the predicted velocity vector and density fields do not satisfy Eq. (1), but it is possible to define momentum and

pressure corrections for which an elliptic Helmholtz equation can be devised to force the system toward continuity. After the correction is applied, the density is updated via the equation of state and the whole procedure is then repeated until convergence. The potential temperature must also be updated after each pressure subiteration for consistency and energy conservation. This iterative procedure represents an efficient way to solve the coupled momentum–continuity equations and guarantee mass conservation with great precision.

For a given pressure iteration, we note that q^n is the value taken by any scalar q at time t^n, q_i^{n+1} is the q estimate at time t^{n+1} for the ith pressure iteration, and q_{i+1}^{n+1} is the estimate at time t^{n+1} for the next iteration. At the end of the iterative procedure, we set $q^{n+1} = q_{i+1}^{n+1}$. For clarity reasons, the projection scheme is described in the following for the Euler implicit time-stepping method only [that is, $\theta = 1$ in Eq. (14)].

First estimates of the velocity vector and potential temperature are obtained at the end of a given iteration step, noted $\tilde{\mathbf{u}}_{i+1}^{n+1}$ and $\tilde{\Theta}_{i+1}^{n+1}$, by solving the system of discretized equations Eq. (14), but using pressure and density estimates from the previous iteration, namely, p_i^{n+1} and ρ_i^{n+1}. Using the updated potential temperature but keeping the density estimate from the previous iteration, one can evaluate an intermediate pressure estimate $\tilde{p}_{i+1}^{n+1}$ using the equation of state:

$$\tilde{p}_{i+1}^{n+1} = p_0 \left(\frac{\rho_i^{n+1} R \tilde{\Theta}_{i+1}^{n+1}}{p_0} \right)^{c_p/c_v}. \tag{17}$$

Small momentum and pressure corrections are then defined following

$$\tilde{\rho}_{i+1}^{n+1} \mathbf{u}_{i+1}^{n+1} = \rho_i^{n+1} \tilde{\mathbf{u}}_{i+1}^{n+1} + \delta(\rho \mathbf{u}) \tag{18}$$

and

$$p_{i+1}^{n+1} = p_i^{n+1} + \delta p, \tag{19}$$

where $\tilde{\rho}_{i+1}^{n+1}$ is still unknown. The pressure and momentum corrections must satisfy the linearized momentum equation so that

$$\delta(\rho \mathbf{u}) = -\Delta t \theta_p \nabla(\delta p), \tag{20}$$

with θ_p being a pressure relaxation coefficient. The intermediate density estimate $\tilde{\rho}_{i+1}^{n+1}$ is obtained using the following Taylor expansion truncated to first order:

$$\tilde{\rho}_{i+1}^{n+1} = \rho_i^{n+1} + \left. \frac{\partial \rho}{\partial p} \right|_{\rho_i^{n+1}, \tilde{p}_{i+1}^{n+1}} (p_{i+1}^{n+1} - \tilde{p}_{i+1}^{n+1}), \tag{21}$$

where $\partial\rho/\partial p$ is known from the equation of state. The discretized continuity equation for the present subiteration

$$\tilde{\rho}_{i+1}^{n+1} + \Delta t \nabla \cdot (\tilde{\rho}_{i+1}^{n+1} \mathbf{u}_{i+1}^{n+1}) = \rho^n \tag{22}$$

can thus be simplified using Eq. (21) and Eq. (18) with

$$\tilde{\rho}_{i+1}^{n+1} = \rho_i^{n+1} + \left. \frac{\partial \rho}{\partial p} \right|_{\rho_i^{n+1}, \tilde{p}_{i+1}^{n+1}} (p_i^{n+1} - \tilde{p}_{i+1}^{n+1} + \delta p) \tag{23}$$

and

$$\nabla \cdot (\tilde{\rho}_{i+1}^{n+1} \mathbf{u}_{i+1}^{n+1}) = \nabla \cdot (\rho_i^{n+1} \tilde{\mathbf{u}}_{i+1}^{n+1}) - \Delta t \theta_p \nabla^2 \delta p. \tag{24}$$

Rearranging these equations, the pressure correction equation can be written as

$$-(\Delta t)^2 \theta_p \nabla^2 \delta p + \left. \frac{\partial \rho}{\partial p} \right|_{\rho_i^{n+1}, \tilde{p}_{i+1}^{n+1}} \delta p = \left. \frac{\partial \rho}{\partial p} \right|_{\rho_i^{n+1}, \tilde{p}_{i+1}^{n+1}} (\tilde{p}_{i+1}^{n+1} - p_i^{n+1}) + \rho^n - \rho_i^{n+1} - \Delta t \nabla \cdot (\rho_i^{n+1} \tilde{\mathbf{u}}_{i+1}^{n+1}). \tag{25}$$

The right-hand side of this equation is readily known after having solved the uncorrected momentum equation for the $(i + 1)$th iteration. The Helmoltz operator appearing on the left-hand side of the equation must be inverted to yield the pressure correction term δp for the ith nonlinear iteration. Both the pressure and momentum can then be updated using Eqs. (18)–(21), and the new corresponding density value can be derived from the equation of state. The procedure is then iterated until convergence using the new estimates of velocity, pressure, and density.

e. *Boundary conditions*

In the test cases presented hereafter, we make use of three different types of boundary conditions: open boundaries (inflow or outflow), surface (free slip) boundaries, and nonreflecting boundaries.

Open boundaries consist of Dirichlet-type conditions. Discretizing the advection operator weakly yields a boundary term $\mathscr{B}$ of the form (for an arbitrary scalar q)

$$\mathscr{B} = \int_{\partial\Omega^D} \varphi \widehat{\mathbf{n} \cdot \mathbf{u} q}, \tag{26}$$

where $\partial\Omega^D$ represents the section of the domain boundary where a Dirichlet condition is imposed. Considering an outflow condition, Eq. (26) can be added to Eq. (8) and solved as such. For inflow condition, we replace q by q_D, the imposed value at the boundary.

We want the normal component of the velocity at a surface boundary to vanish while preserving the tangential

component (free-slip condition for the tangential wind). This reduces to the following condition

$$\mathbf{n} \cdot \mathbf{u} = 0, \tag{27}$$

which can be easily integrated in the weak form of the equations.

Finally, in most cases, nonreflecting boundary conditions must be applied at the inflow and outflow of the domain to prevent the spurious reflection of physical or numerical waves propagating inside the domain. Unlike the other two boundary types, nonreflecting boundaries are defined as layers extending inside the domain where specified prognostic quantities, typically the velocity vector and the potential temperature, are relaxed toward a prescribed state with a given time scale. The relaxation is performed by adding a source term to the considered equations:

$$S_{\mathrm{nr}} = -C_{\mathrm{nr}}(x, y, z)\frac{q - q_0}{\tau_{\mathrm{nr}}}, \tag{28}$$

where $q = \{\mathbf{u}, \theta\}$, q_0 is the prescribed reference value of q, τ_{nr} is the relaxation time scale set equal to the time step Δt, and $C_{\mathrm{nr}}(x, y, z)$ is a function allowing a smooth increase of the relaxation strength toward the boundaries. The functional dependence of C_{nr} on the position follows Klemp and Lilly (1978).

3. Grid adaptivity

The grid adaptivity (or equivalently grid optimization) algorithm implemented in ATHAM-Fluidity belongs to the *hr*-adaptivity family, meaning that both connectivities of the grid elements (*h* adaptive) and the location of the grid vertices (*r* adaptive) may change over the course of a simulation. The algorithm can provide both a refinement of the grid in targeted regions of the numerical domain with strong flow inhomogeneities and a coarsening of the mesh in homogeneous parts of the flow. Adaptivity is performed in an anisotropic way so that grid refinement can follow preferential directions (for instance, vertical refinement only in cases of strong flow stratification), hence possibly resulting in high-aspect-ratio grid cells. Although highly flexible, adaptivity may require a long trial-and-error procedure to yield the optimal mesh based on error-bound criteria while controlling the shape and properties of the target mesh.

All grids handled by ATHAM-Fluidity are by default treated as unstructured. In contrast to other adaptivity algorithms that preserve the overall structure of the mesh while applying local refinements (see, e.g., Behrens et al. 2005; Müller et al. 2013), the optimized grids produced by ATHAM-Fluidity are therefore unstructured by construction and may not preserve key properties of the original mesh such as symmetry.

The overall adaptation procedure can be divided into three main steps described below: 1) based on certain predefined criteria, a metric tensor that will be used to guide grid optimization is computed on the original mesh; 2) the grid is iteratively modified until it satisfies the conditions given by the metric tensor; and 3) the solution field on the original mesh is projected onto the target mesh. A more detailed description of the adaptive algorithm implemented in ATHAM-Fluidity can be found in Pain et al. (2001) and Piggott et al. (2009).

a. *Definition of the metric tensor*

The metric tensor is a symmetric positive-definite tensor including information on the mesh size in all spatial directions (this representation is adequate for anisotropic, fully unstructured meshes). The symmetry and positive definiteness of the tensor are essential properties to allow the definition of a norm characterizing the distance between points on the mesh. We note that $\mathbf{M}$ is the metric tensor and $\mathbf{v}_l$ is a vector directed in the direction parallel to a given edge, with magnitude equal to the edge length. The edge length d_l with respect to a known metric tensor is given by the norm

$$d_l = \sqrt{\mathbf{v}_l^{\mathrm{T}}\mathbf{M}\mathbf{v}_l}. \tag{29}$$

The adapted mesh can thus be computed by defining a metric tensor containing information yielding appropriately adjusted edge lengths. The adaptive metric tensor is typically defined to satisfy a target interpolation error for a selected scalar field q, so that (Chen et al. 2007),

$$\mathbf{M} = \det|\mathbf{H}|^{-(1/2p+n)}\frac{|\mathbf{H}|}{\varepsilon} = \gamma\frac{|\mathbf{H}|}{\varepsilon}, \tag{30}$$

where $\mathbf{H} = \nabla^{\mathrm{T}}\nabla q$ is the Hessian matrix of the field we seek to optimize, ε is the target absolute interpolation error (user defined for a selected scalar), $p = 2$ is the order of the interpolation error norm, and n is the space dimension. The operator $|\mathbf{H}|$ provides a majorant of the Hessian $\mathbf{H}$ defined by $|\mathbf{H}| = \mathbf{V}_{\mathbf{H}}\mathrm{diag}(|\lambda_{Hi}|)\mathbf{V}_{\mathbf{H}}^{-1}$, where the columns of $\mathbf{V}_{\mathbf{H}}$ are the eigenvectors of $\mathbf{H}$ and $\mathrm{diag}(|\lambda_{Hi}|)$ is the diagonal matrix formed by the absolute values of $\mathbf{H}$'s eigenvalues. For simplicity, in very anisotropic conditions, an approximate Hessian matrix can be used instead (Pain et al. 2001).

Rearranging Eq. (30), the target edge length that satisfies the prescribed interpolation error is given by

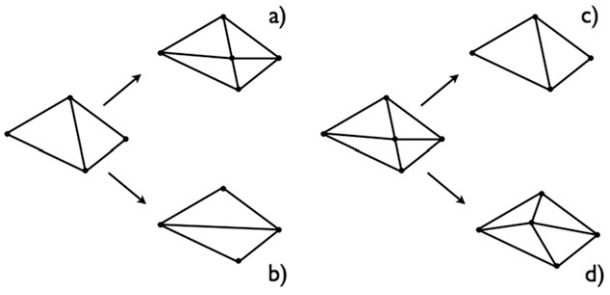


FIG. 2. Examples of elementary operations performed on the grid by the optimization algorithm: (a) element splitting (addition of one node), (b) edge swapping, (c) element merging (removal of one node), and (d) translation of one node.

$$d_l = \sqrt{\gamma \frac{\mathbf{v}_l^{\mathrm{T}} |\mathbf{H}| \mathbf{v}_l}{\varepsilon}}. \tag{31}$$

After adaptation, d_l therefore locally depends on the original edge length before adaptation (through $\mathbf{v_l}$), on the Hessian matrix $\mathbf{H}$ of the considered scalar, and on a user-defined target interpolation error ε.

The definition of the metric tensor may be further modified to account for minimum or maximum edge lengths (Pain et al. 2001). This may be necessary to control the aspect and size of the target mesh and therefore guarantee a certain quality of the adapted grid.

b. The adaptation step

The adaptivity algorithm returns an optimized mesh for a precomputed metric tensor. Various methods exist to generate the new mesh. ATHAM-Fluidity's algorithm is based on the iterative optimization of the original mesh until the new metrics is satisfied [see Vasilevskii and Lipnikov (1999) for the 2D case]. The mesh is progressively deformed following precise operations, including merging two adjacent elements, splitting an existing element, moving nodes, or swapping edges (Fig. 2). If the deformation results in a grid that satisfies the target metric tensor, that is, if the new edge length is consistent with Eq. (29) for a given $\mathbf{M}$, as well as potential additional constraints on the mesh quality (e.g., limiting element aspect ratios or adjacent edge length gradients), the modification is applied. Otherwise, new operations are performed on the original mesh until the entire grid has been successfully modified.

The adaptive step can be further constrained by additional requirements on the minimum and maximum element size allowed, the size gradient between two adjacent elements (gradation), or the total number of elements created. Overall, creating an optimized adapted mesh requires adjusting several such parameters, all having important effects on the appearance of the final grid, therefore making the determination of an optimal set of parameters a nontrivial process.

c. Conservative interpolation

Once the target mesh has been successfully created, the solution fields must be projected onto the new mesh in a consistent and conservative manner. Traditional consistent interpolations are unsuitable for this task as they are typically not designed to handle discontinuous discretizations. Galerkin projection methods provide an excellent alternative as they are optimally accurate to minimize the L_2 norm and conservative. However, their practical implementation appears to be very challenging. In ATHAM-Fluidity, the implementation of Galerkin projection described by Farrell and Maddison (2011) is used.

The projection method is designed to minimize the L_2 norm of the error between the original and interpolated fields. Denoting q_o, a scalar field defined on the original mesh, and q_t, the interpolated field on the target mesh, the Galerkin interpolation procedure reduces to solving the following equation:

$$\int_\Omega q_o \varphi_t = \int_\Omega q_t \varphi_t \tag{32}$$

for each basis function φ_t defined on the target mesh. Variables q_o and q_t can be further expanded using Eq. (7), and Eq. (32) can be recast into the form of a matrix equation:

$$\mathbf{M}_t \underline{q}_t = \mathbf{M}_{to} \underline{q}_o, \tag{33}$$

where

$$(M_t)_{ij} = \int_\Omega \varphi_t^i \varphi_t^j, \tag{34}$$

$$(M_{to})_{ij} = \int_\Omega \varphi_t^i \varphi_o^j, \tag{35}$$

with indices i and j denoting each basis function on the original and target meshes. Evaluating the matrix $\mathbf{M}_{to}$ requires the nontrivial computation of a product between (discontinuous) basis functions defined on two different meshes. To improve the method's accuracy, Farrell and Maddison (2011) suggested the creation of a "supermesh," meshing the superposition of elements from the original and target meshes. Because the basis functions are defined and continuous inside each element of both meshes, the product defining $\mathbf{M}_{to}$ can be computed on the supermesh.

Despite being very accurate, this method remains computationally expensive. Besides, although the present method can be shown to be conservative, it is not bounded so that mesh-to-mesh interpolations do not prevent the generation of local extrema in the projected field. Note that the same method can also be employed to project a discontinuous field onto a continuous grid, as required, for example, for the computation of advective fluxes (see section 1).

4. Model evaluation using uniform grids

The test cases introduced below constitute a standard suite of benchmark simulations used to test and evaluate new numerical methods for nonhydrostatic atmospheric models (Giraldo and Restelli 2008; Choi et al. 2014; Schuster et al. 2014). Six different test cases were selected that can be divided into two main categories. In the first three cases, we seek to evaluate the numerical stability of the solver, grid convergence, and conservation properties. These tests are based on buoyancy-driven dry bubble simulations that are of particular interest for atmospheric convection processes such as convective cloud systems. The last three cases are dedicated to the simulation of physical atmospheric waves (gravity waves), often characterized by weak thermodynamic perturbations. Note that all six cases presented only consider dry atmospheric processes (no water vapor nor liquid water was included) in two dimensions.

For the first three cases, total mass conservation is assessed using the time dependent variable M defined by

$$M(t) = \frac{\left| \sum_{i=1}^{N_n} \rho^i(t) - \sum_{i=1}^{N_n} \rho_0^i \right|}{\sum_{i=1}^{N_n} \rho_0^i}, \tag{36}$$

where ρ_0^i is the initial density at node i, $\rho^i(t)$ is the density at node i and time t, and N_n is the total number of nodes in the numerical domain. Energy conservation is assessed using a similar criterion E:

$$E(t) = \frac{\left| \sum_{i=1}^{N_n} (\rho\Theta)^i(t) - \sum_{i=1}^{N_n} (\rho\Theta)_0^i \right|}{\sum_{i=1}^{N_n} (\rho\Theta)_0^i}, \tag{37}$$

where we employ $\rho\Theta$ as a proxy for internal energy in the absence of heat sources and sinks (no radiation nor microphysical processes are included).

a. Case 1: Rising smooth warm bubble

1) Case description

The rising smooth warm bubble configuration follows (Robert 1993). A Gaussian potential temperature perturbation is initially imposed in an otherwise neutral and static environment in hydrostatic balance following

$$\Theta' = \begin{cases} \dfrac{\Theta_0'}{2}[1 + \cos(\pi r)] & \text{for} \quad r > 1 \\ 0 & \text{otherwise} \end{cases}, \tag{38}$$

with $\Theta_0' = 0.5\,\text{K}$, $\Theta = \Theta_0 + \Theta'$ (the subscript 0 denotes the hydrostatic base state with $\Theta_0 = 303.15\,\text{K}$), and

$$r = \sqrt{\left(\frac{x - x_c}{x_r}\right)^2 + \left(\frac{z - z_c}{z_r}\right)^2}, \tag{39}$$

where $x_c = 500\,\text{m}$, $z_c = 260\,\text{m}$, and $x_r = z_r = 250\,\text{m}$. The initial pressure is found by solving the hydrostatic balance equation

$$\frac{dp}{dz} = -\frac{gp}{\pi R\Theta}. \tag{40}$$

The 2D numerical domain extends between 0 and 1000 m in the horizontal direction and from 0 to 1500 m in the vertical direction. No-flux conditions are imposed on all four boundaries.

Stabilization of the potential temperature field is achieved using the Hermite-WENO slope limiter in all cases shown. In addition, the use of artificial viscosity was found to be necessary to improve grid convergence and preserve the bubble's shape as it rises. A fixed and homogeneous viscosity value of $\mu = 0.1\,\text{m}^2\,\text{s}^{-1}$ was used as in Yelash et al. (2014). No artificial scalar diffusion was otherwise added.

Three different meshes were tested with spatial resolutions of 20, 10, and 5 m. Note that on the 20-m grid, a potential temperature anomaly initially exists along the lower boundary because of the lower part of the bubble being underresolved. All three meshes are symmetric about a vertical axis passing through the center of the bubble. The model time step is determined as the minimum between a CFL limited time step with $\text{CFL}_{\max} = 0.75$ and $\Delta t = 5\,\text{s}$. An effective "velocity based" CFL number is defined on triangular elements:

$$\text{CFL} = \frac{\max|U|\Delta t}{\mathscr{D}}, \tag{41}$$

where $\max|U|$ is the maximum velocity magnitude within each element and $\mathscr{D}$ is the diameter of the circle

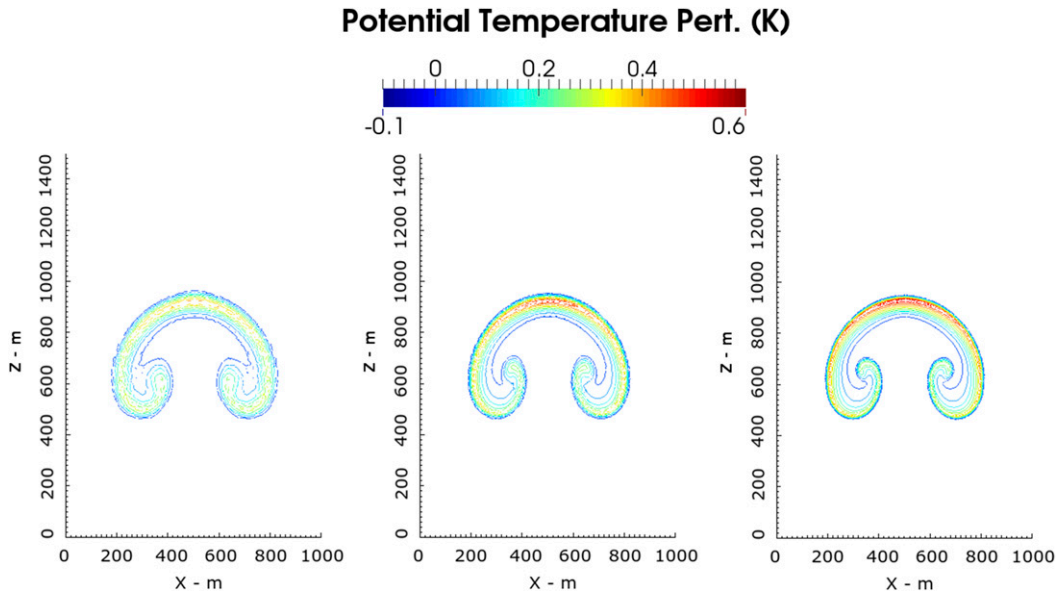


FIG. 3. Perturbation potential temperature contours after 720 s for the rising warm bubble test case. Contours are drawn between −0.1 and 0.6 K at an interval of 0.05 K for (left) 20-, (center) 10-, and (right) 5-m resolution.

inscribed in the triangular element. In the present simulations, minimum time steps of ~ 4, ~ 2, and ~ 1 s are found at 20-, 10-, and 5-m resolution, respectively.

2) RESULTS

Results for the rising warm bubble at the three tested resolutions and after 720 s are shown in Fig. 3. As time progresses and the bubble rises, two Kelvin–Helmoltz rotors develop on each side of the bubble. A clear improvement of the solution is obtained as the resolution is increased. In particular, the two rotors have a more distinct shape and roll further inside at 5-m resolution. At later times and high grid resolution, oscillations start to develop along the bubble interface that quickly evolve into turbulent-like perturbations (Robert 1993; Giraldo and Restelli 2008).

More details on the results at 720 s are provided in Table 1. While the minimum potential temperature perturbation remains only slightly affected by the grid resolution, $\Theta'_{\max}$, $w'_{\max}$, and $w'_{\min}$ are seen to increase with increasing resolution. In an ideal situation, Θ' should not exceed 0.5 K or drop below 0 K, but the use of slope limitation and artificial viscosity does not prevent the development of potential temperature overshoots. Additional simulations have been performed using other, more diffusive slope limitation methods (not shown here). In some instances, excessive anisotropic diffusion provided by certain slope limiters can actually deteriorate the shape of the rising thermal, and the WENO-based limiter has always been found to yield the best overall results despite the fact that it could not completely suppress the Θ' overshoots. The position of the tip of the rising bubble is, however, seen to converge at high resolution toward ~950 m.

The Θ' extreme values shown in Table 1 at all three resolutions appear to be very similar to those reported by Yu et al. (2015) on a very similar configuration with low-order polynomials and selective artificial viscosity. Yu et al. (2015) show that keeping the perturbation potential temperature signal within its expected bounds requires the use of high-resolution grids ($\Delta x \leq 5$ m) and at least eighth-order polynomials. Increasing the level of numerical diffusion may also help convergence, but this would be done at the expense of the overall quality of the results [see Yelash et al. (2014) and discussion above].

Figure 4 displays mass and energy conservation properties for the rising warm bubble test case. Total mass is conserved to machine precision at all resolutions thanks to the use of the efficient iterative semi-implicit pressure-correction procedure. Not surprisingly, energy is not as accurately conserved as total mass, the potential temperature equation being solved in nonconservative form. Note that the conservation error for energy is

TABLE 1. Comparison of maximum and minimum perturbations for the rising warm bubble test case after 720 s and for the three different resolutions; z_{tip} is the altitude reached by the leading Θ' contour.

Resolution (m)	z_{tip}	$\Theta'_{\min}$	$\Theta'_{\max}$	$w'_{\min}$	$w'_{\max}$
20	954.4	−0.076	0.446	−1.6276	2.1775
10	951.6	−0.076	0.526	−1.7092	2.3817
5	950.8	−0.072	0.569	−1.7322	2.4955

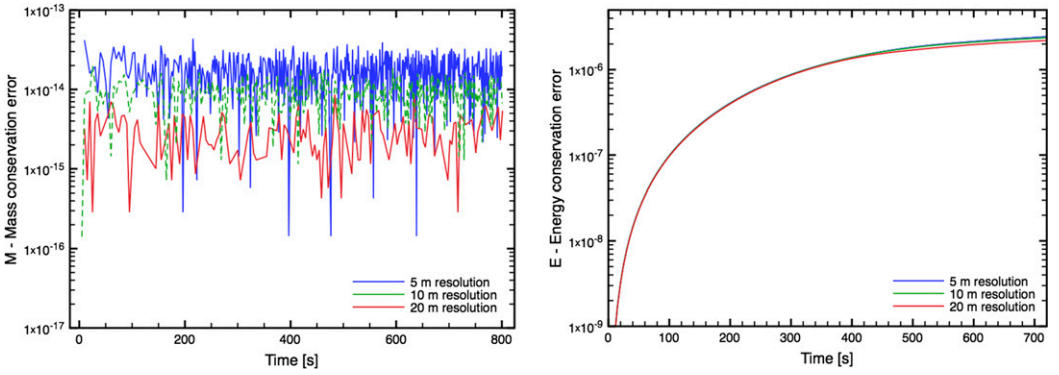


FIG. 4. Time evolution of mass and energy conservation criteria (left) M and (right) E for the warm bubble simulations. Bold lines and dashed lines are for positive and negative values, respectively.

systematically biased positively, indicating an accumulation of internal energy within the numerical domain.

b. Case 2: Density current

1) CASE DESCRIPTION

The density current case consists of an initial cold potential temperature Gaussian perturbation introduced in the domain, which rapidly sinks and hits the surface, therefore creating a cold density current progressing horizontally along the surface. The setup follows Straka et al. (1993) and Giraldo and Restelli (2008).

The initial sounding is computed as in Case 1 using Eqs. (38) and (39). The initial cold bubble corresponds to a potential temperature perturbation Θ_0' of -15 K introduced in a neutrally stratified atmosphere at $\Theta_0 = 300$ K in hydrostatic balance. The initial perturbation is centered at $x_c = 0$ and $z_c = 3000$ m, with dimensions defined by $x_r = 4000$ and $z_r = 2000$ m. The numerical domain is symmetric with respect to the vertical axis at $x = 0$ m and extends to $x = 21\,600$ m. The vertical extent of the domains is set to 9000 m. No-flux conditions are imposed on all boundaries.

Four different grid resolutions were tested: 400, 200, 100, and 50 m. Again, the Hermite-WENO slope limiter has been used in the simulations presented below. To remain consistent with the original configuration proposed by Straka et al. (1993), a fixed artificial viscosity of 75 m s^{-2} is used. Similar to the warm bubble test case, the time step is determined as the minimum between a CFL limited time step with $\mathrm{CFL}_{\max} = 0.75$ and $\Delta t = 5$ s. The minimum time steps reported for the four grids range between 5 (400-m resolution) and 1 s (50-m resolution).

2) RESULTS

Potential temperature perturbation contours after 900 s at each grid resolution are displayed in Fig. 5. At

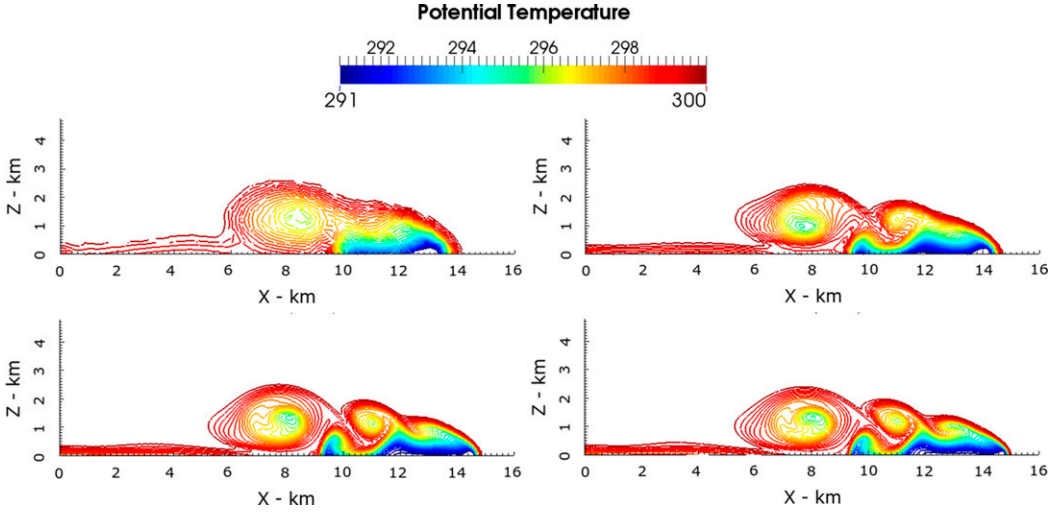


FIG. 5. Potential temperature contours after 900 s for the density current test case. Contours are drawn between 291 and 300 K at an interval of 0.25 K at (top left) 400-, (top right) 200-, (bottom left) 100-, and (bottom right) 50-m resolution. The x and y axes are expressed in kilometers.

TABLE 2. Comparison of maximum and minimum perturbations for the density current test case after 900 s for the fixed meshes at four different resolutions and for the adaptive grid simulation; x_{tip} is the location of the density current front along the x axis, defined by the $\Theta' = -1$°C contour.

Resolution (m)	x_{tip}	Θ'_{min}	Θ'_{max}	w'_{min}	w'_{max}
400	13 989	−10.93	1.18	−9.2307	7.4171
200	14 638	−11.863	0.92	−11.3593	10.7678
100	14 823	−11.94	0.67	−13.3	12.66
50	14 991	−10.49	0.37	−13.14	11.03

400-m resolution, only the largest of the three rotors is clearly visible. At 200 m, two well-developed rotors can be distinguished as in Giraldo and Restelli (2008). When the resolution is increased to 100 and 50 m, the two main eddy structures are clearly visible and a third rotor is seen to develop. Qualitatively, the results at the two highest resolutions are very similar to Giraldo and Restelli (2008) [despite Giraldo and Restelli (2008) using tenth-order polynomials] and no clear improvement between the 100-m grid and the 50-m grid can be seen from Fig. 5.

Table 2 presents a quantitative comparison between the four simulations. Both the potential temperature perturbation and vertical velocity extrema are reduced when increasing the resolution from 100 to 50 m. The 100-m-resolution grid produces the most extreme results of all cases (except for Θ'_{max}). At 200-m resolution, our results are comparable to the solutions presented by Straka et al. (1993) using spectral models that indicate low numerical dissipation. At the lowest resolution tested, the model is still seen to produce potential temperature overshoots ($\Theta'_{max} = 0.282$ K instead of 0). Unlike the warm bubble test case presented in section 4a, the maximum potential temperature perturbation is seen to converge as the resolution is increased. Note also that the very high resolution (25 m) simulations discussed by Giraldo and Restelli (2008) converge toward $\Theta'_{max} \sim 0$ K and $\Theta'_{min} \sim -9$ K, which is consistent with our results.

The leading edge of the current propagates faster with increasing grid resolution (x_{tip} in Table 2), similar to Giraldo and Restelli (2008). The position of the tip at 200-m resolution is consistent with the upwind biased high-order finite difference results shown by Straka et al. (1993). At high resolution, the front propagates faster than in the high-resolution test presented by Giraldo and Restelli (2008). Overall, a 50-m-resolution grid does not seem to be sufficient to accurately capture all the features of the density current solution.

The total mass and energy conservation indicators M and E are shown in Fig. 6 for the density current case. Unlike the other two bubble cases, mass does not appear to be conserved to machine precision. In the density current case, the flow is indeed driven by large Θ' anomalies that produce strong pressure oscillations in the numerical domain. The amount of nonlinear iterations needed by the projection method has been restricted to three in all cases shown to limit computational costs, but more subcycles appear to be necessary in this case for the iterative procedure to converge. Total energy is again not as accurately conserved as mass, with conservation errors consistently biased positively. In general, parameter E exhibits a larger error here compared to the warm bubble test case.

c. *Case 3: Interacting warm and cold bubbles*

1) CASE DESCRIPTION

Besides the single rising warm bubbles, Robert (1993) also proposed an alternative configuration consisting of one large smooth warm bubble similar to case 1 and a smaller cold bubble located above the first one. As time progresses, both bubbles collide and mix. This test case was designed to evaluate the accuracy of the numerical discretization in presence of smaller, underresolved physical features (the small bubble is only 50 m wide, which is equivalent to $\sim 5 - 10\Delta x$). This case has already been used to evaluate the performances of DG methods on adaptive grids (Müller et al. 2013; Yelash et al. 2014).

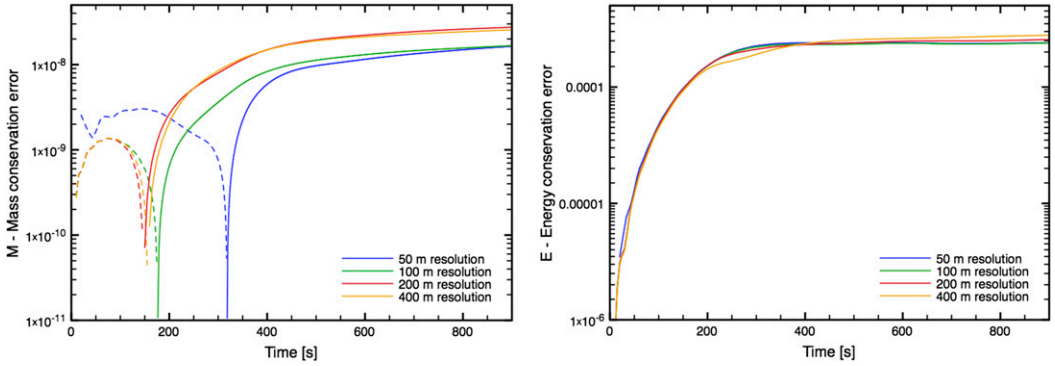


FIG. 6. As in Fig. 4, but for the density current simulations.

The simulations are carried out with the same initial base state in hydrostatic balance as in section 4a and using the same numerical domain. The initial potential temperature perturbations of both bubbles are described by

$$\Theta' = \begin{cases} \Theta'_0 \exp\left[-\dfrac{(r-r_c)^2}{s^2}\right] & \text{for} \quad r > r_c \\ \Theta'_0 & \text{otherwise} \end{cases}. \tag{42}$$

The large warm perturbation is centered at $x = 500$ and $z = 300$ m, with $\Theta'_0 = 0.5$ K, $r_c = 150$ m, and $s = 50$ m (Robert 1993). The smaller cold perturbation is centered at $x = 560$ and $z = 640$ m, with $\Theta'_0 = -0.15$ K, $r_c = 0$ m, and $s = 50$ m.

Simulations were performed on two fixed uniform grids with spatial resolutions of 10 and 5 m, respectively. The Hermite-WENO slope limiter was used in all the simulations and the time step was determined following the same procedure as in the warm bubble case.

2) Results

As time progresses, the smaller cold bubble collides with the right side of the larger rising thermal and significantly alters its shape. The solution after 600 s shown on Fig. 7 (right) resembles the results shown by Robert (1993) or Müller et al. (2013). Only minor differences can be seen between the 10- and 5-m-resolution grids after 600 s. In particular, the roll-ups developing on the right branch of the warm bubble are more clearly defined at higher resolution. No substantial differences between the two mesh resolutions can be seen at 400 s. Overall, the solution field does not seem to be particularly degraded in the low-resolution case, although the smaller cold bubble is only discretized by five elements.

Figure 8, displaying the mass and energy conservation parameters M and E for the interacting bubbles simulations, shows total mass conservation to machine precision, similar to the single warm bubble case. Again, $\rho\Theta$ does not appear to be conserved as well as total mass because of the use of a nonconservative equation, resulting in the accumulation of internal energy within the domain (positive error).

d. Case 4: Inertia–gravity waves

1) Case description

Following Skamarock and Klemp (1994), the inertia–gravity wave test case involves the horizontal propagation of a nonhydrostatic gravity wave in a channel. An initial potential temperature perturbation is imposed in an otherwise uniformly stratified atmosphere, producing gravity waves propagating symmetrically toward the left and right sides of the domain. Because a homogeneous flow is added from left to right, the leftward propagating part of the wave appears to be quasi stationary. This configuration was originally designed to test the performances of time integration schemes in nonhydrostatic models.

The initial sounding is taken to be in hydrostatic balance with a uniform stable potential temperature stratification. The potential temperature profile is based on the definition of the Brunt–Väisälä frequency:

$$\mathcal{N}^2 = g\frac{d\ln\Theta}{dz}, \tag{43}$$

which yields, after integration:

$$\Theta = \Theta_s \exp\left(-\frac{\mathcal{N}^2}{g}z\right), \tag{44}$$

with $\mathcal{N}$ initially set to 0.01 s^{-1} and $\Theta_s = 300$ K the surface potential temperature. Equation (44) combined with the hydrostatic balance Eq. (40) defines the initial hydrostatic pressure sounding. The initial potential temperature perturbation is given by

$$\Theta' = \Theta_c \frac{\sin(\pi z/h_c)}{1 + [(x - x_c)/a_c]^2}, \tag{45}$$

with $h_c = 10\,000$ m, $x_c = 100\,000$ m, $a_c = 5000$ m, and $\Theta_c = 0.01$ K. A uniform flow with constant velocity $u = 20$ m s^{-1} is imposed.

The mesh is 300 000 m long and 10 000 m high, with a 500-m resolution in both directions (grid aspect ratio of 1). Free-slip conditions are set for the top and bottom boundaries, with the vertical velocity component set to 0 at both surfaces. In contrast to the original configuration, the left and right boundaries are defined as an inflow and outflow, respectively. Damping layers 20 000 m wide have been added on each side of the domain to avoid spurious wave reflection. These layers allow fast relaxation of the prognostic velocity and potential temperature fields toward the initial state. Compared to the dry bubble test cases presented previously, the inertia–gravity wave simulations (as well as the topographically forced waves in the following sections) were performed without slope limiter or grid adaptation.

Two different simulations were performed for this configuration. In the first case, the time step was automatically computed based on a fixed CFL number of 0.25 [corresponding to a time step of ~6.25 s, similar to Wicker and Skamarock (1998)], while in the second case, a fixed time step of 1 s was selected (corresponding

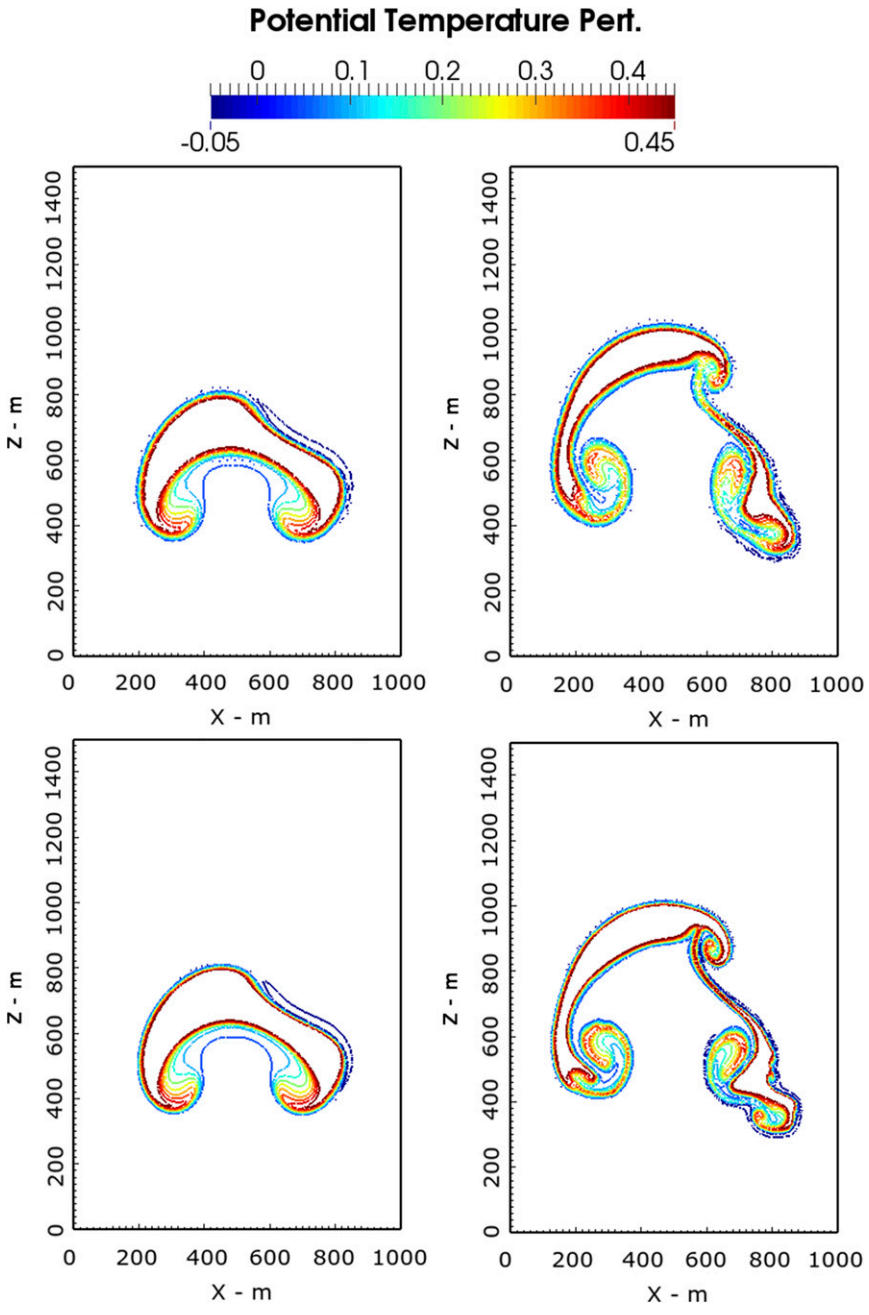


FIG. 7. Perturbation potential temperature contours after (left) 400 and (right) 600 s for the interacting warm and cold bubbles test case. Contours are drawn between −0.05 and 0.45 K with a 0.05-K interval. Uniform grids at (top) 10- and (bottom) 5-m resolution.

to a CFL number of ~0.04). At $\Delta t = 1$ s, the maximum acoustic CFL number based on the sound speed is ~0.7 and sound waves are therefore explicitly resolved.

2) RESULTS

The solution of the inertia–gravity wave propagation after 3000 s is shown in Fig. 9 for the two different configurations. A clear asymmetry of the solution is found when using a relatively large time step (CFL = 0.25), with the right side of the wave propagating downstream being significantly damped compared to the leftward propagating part of the wave. The difference between the left and right potential temperature perturbation extrema is ~0.0001 K, that is, 4% of the maximum value reached within the entire domain. This result is not consistent with the perfectly symmetric analytical solution produced by

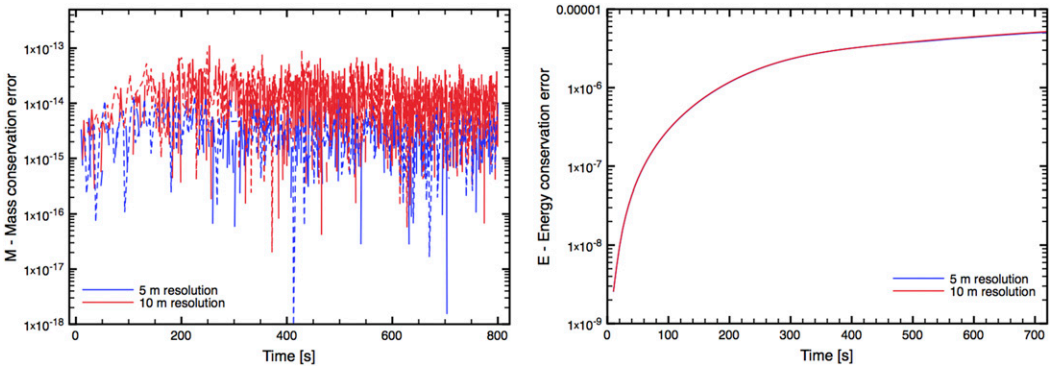


FIG. 8. As in Fig. 4, but for the interacting warm and cold bubbles simulations.

Skamarock and Klemp (1994) using the linearized Boussinesq equations. Note also that in this configuration, the central part of the wave appears to be significantly degraded with the development of small-scale numerical artifacts.

Increasing the grid resolution from 500 to 250 m (not shown here) does not improve the accuracy and symmetry of the solution. However, using a smaller time step of 1 s perfectly preserves the symmetry of the solution. Asymmetric solutions for this case have already been reported in the literature under similar conditions, including in the original work by Skamarock and Klemp (1994). Phase-speed errors introduced by low-order time integrators, especially at high CFL number, were identified as the sources of these asymmetries (Wicker and Skamarock 1998).

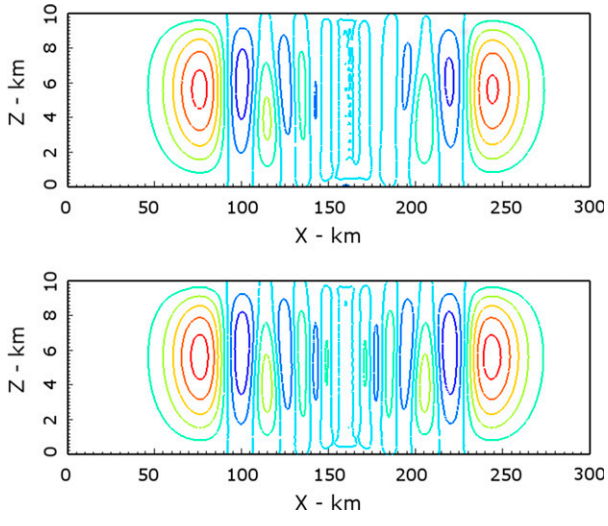


FIG. 9. Inertia–gravity wave solution after 3000 s: (top) fixed CFL number simulation (CFL = 0.25) and (bottom) fixed time step simulation ($dt = 1$ s). Perturbation potential temperature contours were drawn between −0.0015 and 0.003 K with an interval of 0.0005 K.

The center of the propagating wave is located at 160 000 m (Fig. 10) in both cases, indicating accurate advection of the initial perturbation at a constant 20 m s^{-1} velocity, even for relatively large time steps.

e. *Case 5: Schär mountain case, hydrostatic flow*

1) CASE DESCRIPTION

In this case, a dry atmospheric flow is forced over a five-peak mountain range with constant horizontal velocity, therefore producing steady-state gravity waves. The configuration follows Schär et al. (2002). The initial state is computed using Eq. (44), with $\mathcal{N} = 0.01\ \mathrm{s}^{-1}$, and the surface potential temperature is set to 280 K.

The domain extends between −25 000 and 25 000 m in the horizontal direction and is 22 000 m high. The vertical and horizontal resolutions are set to 250 and 500 m, respectively. The five-peak mountain profile is defined by

$$h = h_c \exp\left[-\left(\frac{x}{x_c}\right)^2\right] \cos^2\left(\frac{\pi x}{\lambda_c}\right), \tag{46}$$

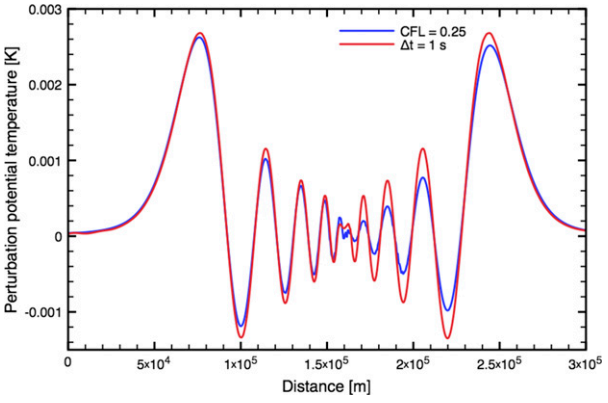


FIG. 10. Perturbation potential temperature profiles after 3000 s along the x axis at a 5000-m height for the inertia–gravity wave.

where h is the terrain elevation, $h_c = 250\,\text{m}$ corresponds to the maximum elevation at the center of the domain, $x_c = 5000\,\text{m}$, and $\lambda_c = 4000\,\text{m}$. The left boundary is defined as an inflow, forced with a uniform horizontal velocity $u = 10\,\text{m}\,\text{s}^{-1}$. The right and top boundaries are defined as outflows. Damping layers 7000 m deep have been added near all open boundaries (lateral and top) following the description given in section 2e. The surface is treated as a free-slip condition with no vertical velocity. A constant time step of 1.9 s has been used, corresponding to an effective CFL number of ~0.1. At this resolution, we can expect both the gravity and sound waves to be accurately resolved as this was shown to be of particular importance in section 4d.

Defining a Froude number as $\mathcal{N} x_c/u$, we find Fr = 5, so that this case stays in the hydrostatic regime. The model results are presented after a simulation time of 8 h.

2) RESULTS

A reference solution for the linearized problem was computed and used for comparisons. The linearized steady-state pseudocompressible equations from Durran (1989) are solved using Fourier decomposition in the vertical direction and finite differences in the horizontal direction. The spatial discretization was set to be the same as in the model simulation.

The model results are qualitatively in good agreement with the reference solution (Fig. 11). The results are particularly similar in a region centered around the main mountain peak and close to the surface. On the sides and higher in the atmosphere, the model solution might be already affected by the presence of the damping layers. Such topographically forced atmospheric flows are typically very sensitive to the imposed boundary conditions and damping layers, which must provide just the appropriate level of numerical damping. The results shown could possibly be improved by optimizing the imposed damping layers, but this remains outside of the scope of the present work.

Let's recall here that the reference solution was obtained by solving linearized steady-state hydrostatic equations. Nonlinearities and nonhydrostatic effects in ATHAM-Fluidity as well as truncation errors introduced when solving the fully discrete time-dependent equations are likely to cause the observed discrepancies between both sets of results. Note also that the $\mathbb{P}_{1\text{DG}} - \mathbb{P}_2$ discretization, while completely inhibiting the development of spurious pressure modes, may support spurious stationary inertial oscillations (Cotter and Ham 2011), which could introduce weak additional biases in the numerical solution.

f. Case 6: Linear single-mountain case, nonhydrostatic flow

1) CASE DESCRIPTION

In this case, we consider the steady-state solution produced by a dry atmospheric flow forced over a single linear mountain profile. The initial conditions can be found in Giraldo and Restelli (2008). As in case 5, the initial state is in hydrostatic balance with a potential temperature stratification corresponding to a buoyancy frequency of $\mathcal{N} = 0.01\,\text{s}^{-1}$, with the surface potential temperature set to $\Theta_0 = 280\,\text{K}$. The flow is initially assigned a uniform velocity of $u = 10\,\text{m}\,\text{s}^{-1}$.

The linear mountain profile is given by

$$h = h_c \Big/ \left[1 + \left(\frac{x - x_c}{a_c}\right)^2\right], \tag{47}$$

with $h_c = 1$, $x_c = 0$, and $a_c = 1000\,\text{m}$. The numerical domain extends between $-72\,000$ and $72\,000\,\text{m}$ in the horizontal direction and is $22\,000\,\text{m}$ high. The horizontal and vertical resolutions are set to 600 and 300 m, respectively. The numerical domain is defined as an inflow/outflow configuration with damping layers at both lateral boundaries and starting $12\,000\,\text{m}$ away from each boundary. A $50\,000$-m-deep damping layer is also imposed near the top boundary. As in case 5, the surface is defined as a free-slip condition with no vertical velocity. The time step was set to 2.3 s, again corresponding to an effective CFL number of ~0.1.

While the five-peak mountain case considers the solution of a hydrostatic flow, the single mountain case is defined by a Froude number $\mathcal{N} a_c/u = 1$ and therefore belongs to the nonhydrostatic regime. The simulation is continued for 5 h before analyzing the results.

2) RESULTS

A reference solution based on the linearized steady-state pseudocompressible equations has been computed as exposed in section 4e and is used here for comparison. Figure 12 presents the modeled steady-state perturbation velocities and potential temperature contours. Again, all results compare well with the reference solutions close to the surface but are quickly degraded when moving higher in the atmosphere. Similarly, the model results agree better at the center of the domain than away from it, indicating the influence of the damping layers at the boundaries. As already noted by Giraldo and Restelli (2008), the simulated horizontal velocity perturbation differs more strongly from the reference solution compared to the vertical velocity or perturbation potential temperature.

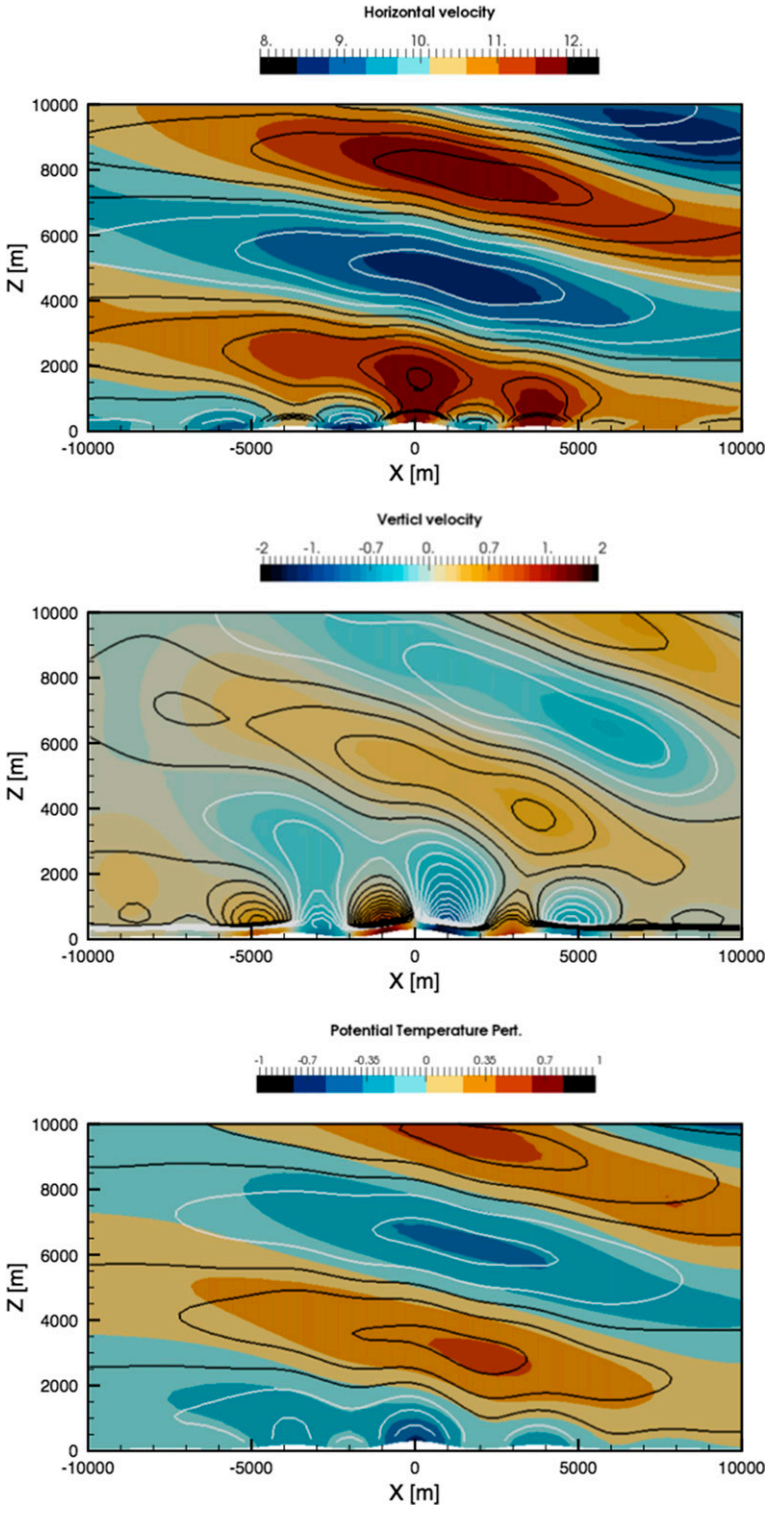


FIG. 11. Contours of (top) horizontal velocity u, (middle) vertical velocity v, and (bottom) potential temperature perturbations Θ' for the five-peak mountain test case. The color scale ranges between 8 and 12.5 m s^{-1} with 10 levels for horizontal velocity, between -2 and 2 m s^{-1} with 40 levels for vertical velocity, and between -1 and 1 K for the potential temperature perturbation with 10 levels. The contour lines correspond to the analytical solution and were defined similarly to the color scales, with black and white lines representing positive and negative values, respectively.

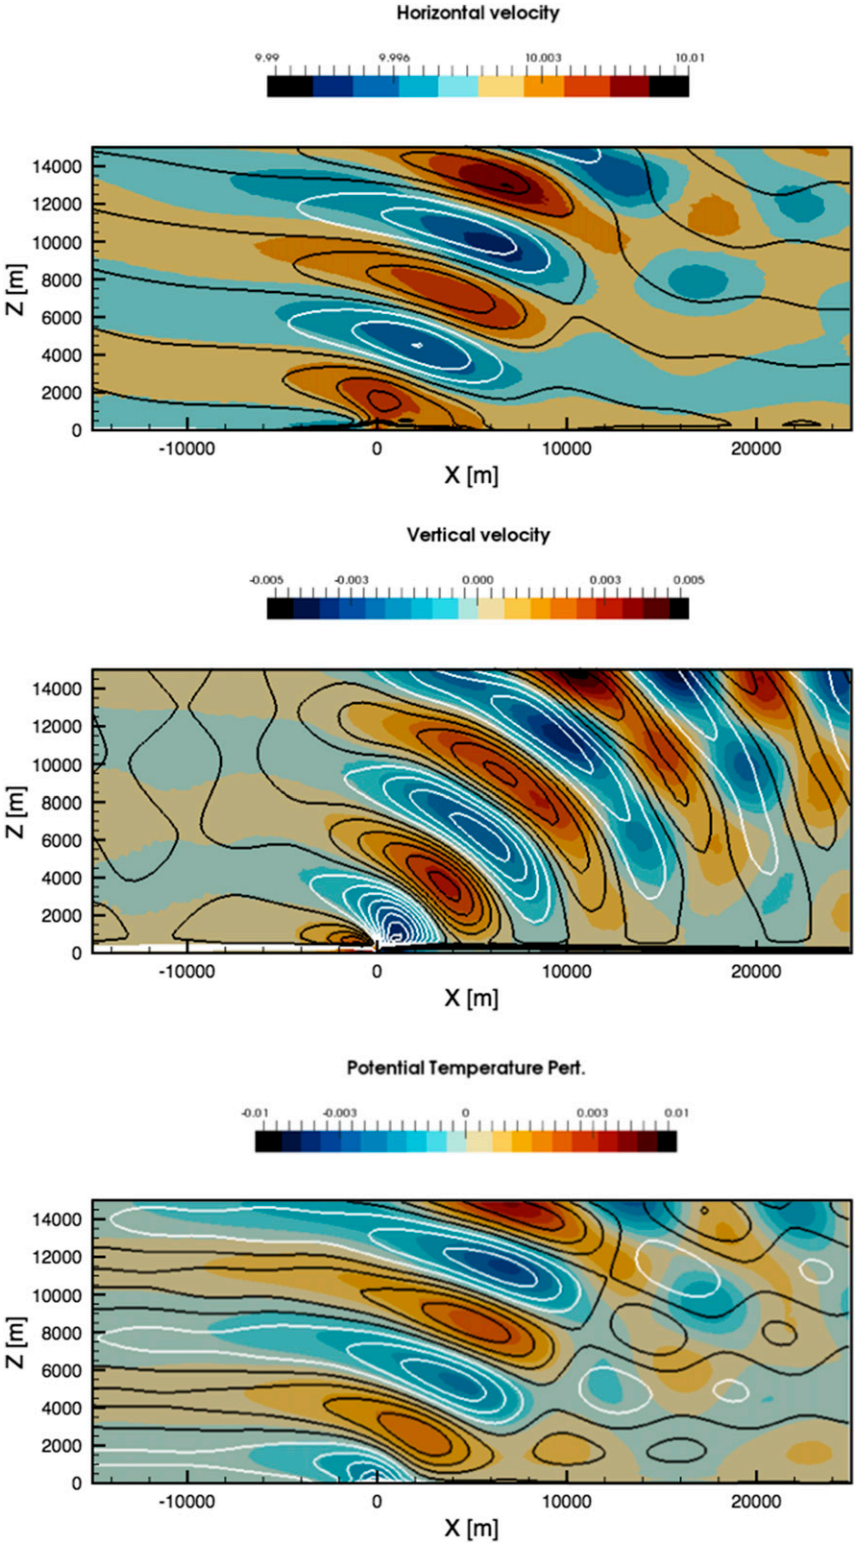


FIG. 12. As in Fig. 11, but for the nonhydrostatic single mountain test case. The color scales (and contours) now contain 11, 21, and 21 levels for u, v, and Θ', respectively.

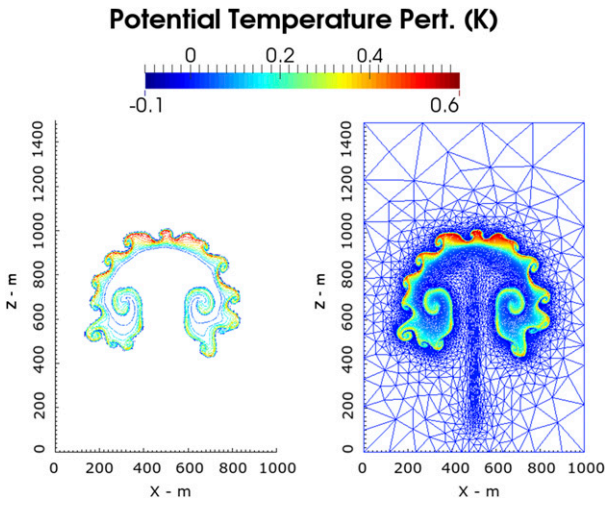


FIG. 13. Perturbation potential temperature contours and adapted mesh overview after 720 s for the rising warm bubble test case with grid optimization. Contours and color scale are similar to Fig. 3.

Again, it should be noted that the reference solution has been obtained for linearized steady-state hydrostatic equations so that nonlinear and nonhydrostatic effects in addition to truncation errors introduced by ATHAM-Fluidity are likely to cause the observed biases.

5. Preliminary results on adaptive grids

In the following, preliminary results obtained using the grid optimization algorithm are presented for the first three test cases introduced in section 4. These results constitute a very first attempt to use anisotropic *hr* adaptivity with ATHAM-Fluidity under atmospheric conditions. More dedicated investigations would be needed to better characterize the influence of various optimization parameters on the adaptive grid results and improve the overall reliability of the approach. In particular, as discussed in section 3b, finding the optimal set of parameters yielding the best target mesh possible for a given case is a very tedious process and is out of the scope of the present study.

a. Rising warm bubble

One additional simulation has been performed for the rising warm bubble case using mesh optimization, with a maximum resolution set to 5 m and a maximum edge length of 250 m. Optimization is computed based on a target absolute interpolation error for potential temperature of 0.005 K [Eq. (30)] and new adapted grids are generated every five time steps.

A snapshot of the mesh and potential temperature contours after 720 s is shown in Fig. 13. As expected, the optimization algorithm generates small elements in regions where steep potential temperature gradients are found. The highest grid resolution, corresponding to an effective resolution of 5 m, is reached at the boundaries of the rising thermal while very large elements are found in the upper part of the domain. Elements in the high-resolution parts of the grid appear to be rather uniform and isotropic, a result of the grid quality constraint imposed during adaptation.

Compared to the fixed grid simulations, the solution on the adapted grid quickly develops turbulent-like perturbations along the edge of the thermal (these perturbations are already visible after ~600 s while they develop only after 800 s on a uniform mesh). We speculate here that these spurious instabilities are caused by irregular element features (such as element orientation, aspect ratios, or asymmetry) in the adapted grid combined with errors introduced by the interpolation of the solution field onto the new mesh (we recall that the interpolation procedure is conservative but not bounded, as emphasized in section 3c). Besides, the asymmetry of the numerical solution directly ensues from the asymmetry of the optimized mesh as *hr* adaptivity naturally generates fully unstructured and asymmetric grids, even based on perfectly symmetric flow fields (see section 3).

No extra scalar diffusivity except for that provided by the WENO slope limiter has been employed here. However, using extra uniform artificial diffusion does not appear to be an appropriate solution to improve the results shown in Fig. 13. Although artificial scalar diffusivity

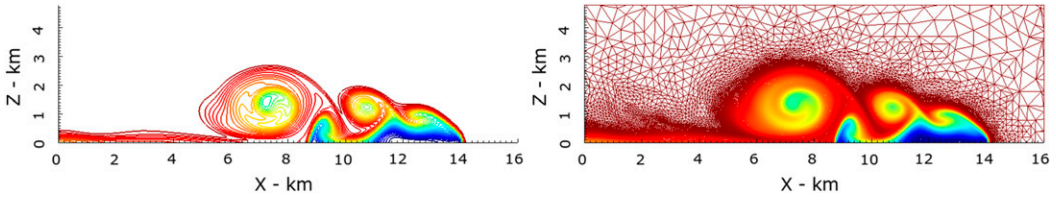


FIG. 14. Perturbation potential temperature contours after 900 s for the density current test case with adaptive mesh. (left) Contours are drawn as described in Fig. 5 (the same color scale is used). (right) An overview of the adapted grid after 900 s.

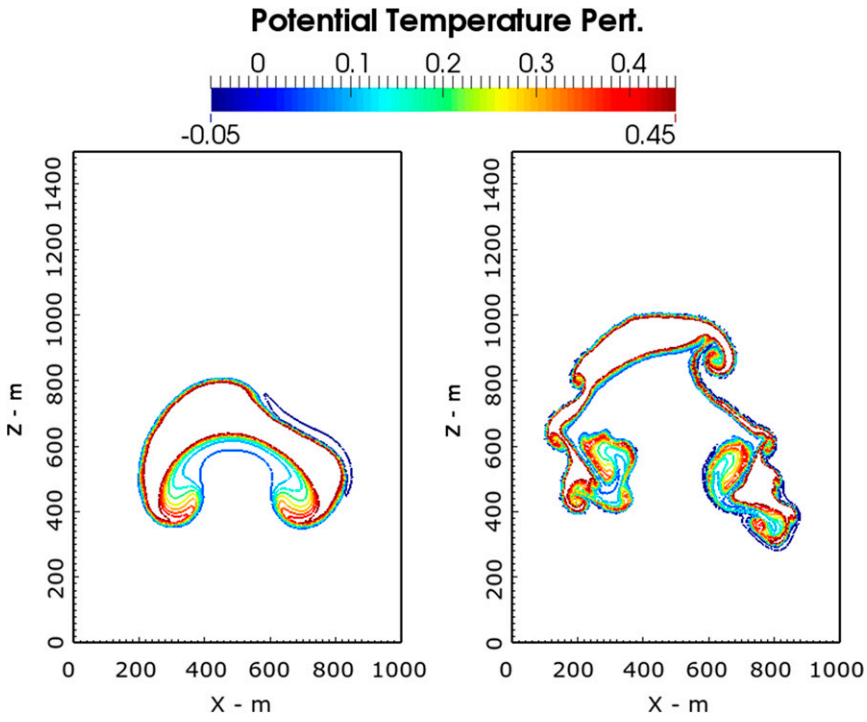


FIG. 15. As in Fig. 7, but with optimized mesh.

would indeed help prevent the development of numerical instabilities (as is the case for the density current shown in section 5b), the potential temperature gradients would be quickly damped, thereby modifying the bubble's ascent. In this situation, the adaptive results would no longer be comparable to the uniform grid solution.

Compared to the fixed 5-m-resolution mesh (at equivalent resolution), the number of elements has been reduced by a factor of ~6 (from 100 000 to 16 000). The computation time has been subsequently reduced by a factor of ~5.7.

b. Density current

The density current test case has been simulated using grid optimization with a maximum resolution equivalent to the highest-resolution uniform grid case (50 m). Optimization is performed every five time steps and is again configured to produce mesh refinement around strong potential temperature gradients. The target absolute interpolation error is set to 0.01 K.

In contrast to the rising thermal test case, the density current does not develop any turbulent-like perturbation when grid optimization is used (see Fig. 14). The artificial diffusivity added to help reach a converged solution (Straka et al. 1993) contributes to stabilizing the solution field even on an irregular adapted mesh. The optimized grid results are qualitatively similar to the uniform grid ones, and the essential features of the density current solution are well captured. A striking difference comes from the propagation of the leading front, which reaches only 14 166 m with the adaptive grid compared to over 14 900 m with a uniform grid. This notable speed reduction may to some extent stem from the propagation of the front through coarser grid elements between two optimization steps, as lower resolutions seem to slow down the current propagation (Table 2). The optimized mesh also limits the generation of potential temperature overshoots compared to a uniform grid, with $\Theta'_{\min} = -9.49$ and $\Theta'_{\max} = 0.282$ K.

Müller et al. (2013) also used a DG-based method to simulate the density current test case with mesh adaptivity. Their results show overall better agreements between the adapted mesh and fixed mesh simulations. In particular, the position of the leading front of the current varies by less than 0.1% between adapted and fixed grid simulations at different resolutions (the same can be noted for $\Theta'_{\max}$).

After 900 s, the number of elements in the adaptive grid setup is only 16 000, as compared to the 96 000 elements needed by the uniform high-resolution mesh (factor ~5.8 reduction). The CPU time has correspondingly been reduced by a factor ~5.5.

c. Interacting warm and cold bubbles

For the interacting bubbles case, one simulation using grid optimization with a minimum effective element size

of 5 m has been carried out. As in the previous two cases, grid adaptation is performed every five time steps. The remeshing algorithm uses an error estimate based on the potential temperature equal to 0.001 K and producing refinement in strong Θ gradient regions.

As in the warm rising thermal case, simulating the interacting warm and cold bubbles with grid optimization leads to the development of turbulent-like perturbations after 600 s (Fig. 15), most likely generated by irregular grid features (element aspect ratios, face orientations, etc.) and unbounded mesh to mesh interpolations. After 400 s, no substantial differences can be seen between the uniform and adaptive grid results. As in case 1, no substantial improvements can be expected by the addition of extra numerical diffusion.

Compared to the fixed 5-m-resolution mesh, the number of elements has been reduced by a factor of ~4.6 (from 100 000 to 20 250). The computation time has been reduced by a factor of ~4.3. These numbers are slightly lower than in case 1 because of the presence of the smaller cold bubble adding small-scale physical features that require additional meshing efforts.

6. Discussion and conclusions

In this paper, we introduce a new nonhydrostatic limited-area (cloud-resolving and large-eddy model scales) solver, ATHAM-Fluidity, employed for the first time here to simulate idealized atmospheric flows. ATHAM-Fluidity uses a mixed finite-element method (a DG discretization for momentum and scalars but a CG discretization for the pressure and density fields), along with a compressible pressure projection procedure, and operates on fully unstructured adaptive grids. The solver is tested using a series of standard benchmark cases designed to evaluate the performances of such models under various operational conditions. Among these test cases, the first three (dry bubble simulations) are selected to test the numerical methods, grid convergence, and the impact of the grid optimization algorithm while the last three provide a severe evaluation of the model accuracy through the simulation of weak gravity waves. Although all the cases presented were originally designed for two-dimensional simulations, the dynamical core has been developed as a three-dimensional solver and has already been extensively applied within this context.

The dry bubble simulations revealed the capacity of the fixed grid dynamical core to accurately capture dry convective processes in the atmosphere. This represents the first step toward the simulation of more realistic atmospheric processes leading to the formation of convective cloud systems and the associated weather phenomena. In particular, the model was found to provide relatively good results even at low spatial resolution and despite the use of low-order polynomials. A Hermite-WENO slope limiter has been used in all these simulations, as it was found to give a good compromise between numerical diffusion and overall accuracy.

The gravity waves (both in the hydrostatic and nonhydrostatic regimes) simulated using ATHAM-Fluidity were found to qualitatively and quantitatively reproduce reference solutions provided by a linear model. It should be emphasized that these reference solutions were obtained for a simplified set of equations (using the hydrostatic and pseudocompressible assumptions) in steady state, which naturally limits the possibility for accurate comparisons with our model (which was configured to solve the fully compressible Euler equations). Note also that these test cases are typically extremely sensitive to the lateral boundary conditions applied, as these gravity waves require appropriate damping layers to avoid wave reflection without perturbing the numerical solution in the region of interest. Overall, our results suggest that the new model, as introduced in the present paper, is able to accurately capture very weak perturbations in the background atmosphere from, for instance, topographical features as small as 1 m, and that the configuration used (including damping layers at the boundaries) can readily be used to simulate actual dry atmospheric flows.

Preliminary results using grid adaptivity were also presented based on the three dry bubble configurations. The optimization algorithm allowed for the decrease of the CPU cost of the three cases tested by a factor of ~5 for about 6 times fewer grid elements. We believe that these ratios could be further improved by optimizing the various parameters controlling the adaptive algorithm (error threshold, gradation factor, largest cell size, etc.). Simulations with adaptive grid consistently showed substantial discrepancies compared to the uniform grid simulations at equivalent resolutions. In particular, the development of turbulent-like perturbations was observed in two of the bubble cases, most likely caused by the irregularity and asymmetry of the optimized grids as well as small numerical errors introduced by the unbounded mesh-to-mesh interpolation procedure. Although these results did not compare favorably with the uniform grid simulations, the encouraging CPU cost ratios reported should serve as a motivation to pursue the development of the grid adaptivity technique and improve its reliability.

Future works will be dedicated to the implementation and testing of a comprehensive bulk cloud microphysics scheme. The model will ultimately be used for large-eddy simulations of atmospheric processes from the development of single clouds in idealized atmospheres to the evolution of large cloud systems forced by actual meteorological conditions.

Acknowledgments. This research has received funding from the European Union Seventh Framework Program (FP7/2007-2013) under Grant Agreement 603663 for the research project PEARL (Preparing for Extreme And Rare events in coastaL regions). C. Pain and J. Percival would also like to acknowledge the EPSRC multiphase program grant MEMPHIS. Finally, the authors thank the editor, Hilary Weller, as well as three anonymous reviewers for their valuable comments and suggestions. All data presented in this manuscript are freely available through the open-access online data repository Zenodo (https://zenodo.org).

REFERENCES

Bassi, F., and S. Rebay, 1997: A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier-Stokes equations. *J. Comput. Phys.*, **131**, 267–279, doi:10.1006/jcph.1996.5572.

Behrens, J., N. Rakowsky, W. Hiller, D. Handorf, M. Läutter, J. Päpke, and K. Dethloff, 2005: Amatos: Parallel adaptive mesh generator for atmospheric and oceanic simulation. *Ocean Modell.*, **10**, 171–183, doi:10.1016/j.ocemod.2004.06.003.

Botti, L., and D. A. Di Pietro, 2011: A pressure-correction scheme for convection-dominated incompressible flows with discontinuous velocity and continuous pressure. *J. Comput. Phys.*, **230**, 572–585, doi:10.1016/j.jcp.2010.10.004.

Brdar, S., M. Baldauf, A. Dedner, and R. Kloöfkorn, 2013: Comparison of dynamical cores for NWP models: Comparison of COSMO and DUNE. *Theor. Comput. Fluid Dyn.*, **27**, 453–472, doi:10.1007/s00162-012-0264-z.

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

Chen, L., P. Sun, and J. Xu, 2007: Optimal anisotropic meshes for minimizing interpolation errors in Lp-norm. *Math. Comput.*, **76**, 179–204, doi:10.1090/S0025-5718-06-01896-5.

Choi, S.-J., F. X. Giraldo, J. Kim, and J. Shin, 2014: Verification of a non-hydrostatic dynamical core using a horizontal spectral element method and vertical finite difference method: 2-D aspects. *Geosci. Model Dev.*, **7**, 2717–2731, doi:10.5194/gmd-7-2717-2014.

Cockburn, B., and C.-W. Shu, 2001: Runge-Kutta discontinuous Galerkin methods for convection dominated problems. *SIAM J. Sci. Comput.*, **16**, 173–261, doi:10.1023/A:1012873910884.

Comblen, R., J. Lambrechts, J.-F. Remacle, and V. Legat, 2010: Practical evaluation of five partly discontinuous finite element pairs for the non-conservative shallow water equations. *Int. J. Numer. Methods Fluids*, **63**, 701–724, doi:10.1002/fld.2094.

Cotter, C. J., and D. A. Ham, 2011: Numerical wave propagation for the triangular $P1_{DG}$–P2 finite element pair. *J. Comput. Phys.*, **230**, 2806–2820, doi:10.1016/j.jcp.2010.12.024.

——, and J. Shipton, 2012: Mixed finite elements for numerical weather prediction. *J. Comput. Phys.*, **231**, 7076–7091, doi:10.1016/j.jcp.2012.05.020.

——, D. A. Ham, C. C. Pain, and S. Reich, 2009a: LBB stability of a mixed Galerkin finite element pair for fluid flow simulations. *J. Comput. Phys.*, **228**, 336–348, doi:10.1016/j.jcp.2008.09.014.

——, ——, and ——, 2009b: A mixed discontinuous/continuous finite element pair for shallow-water ocean modelling. *Ocean Modell.*, **26**, 86–90, doi:10.1016/j.ocemod.2008.09.002.

Düben, P. D., P. Korn, and V. Aizinger, 2012: A discontinuous/continuous low order finite element shallow water model on the sphere. *J. Comput. Phys.*, **231**, 2396–2413, doi:10.1016/j.jcp.2011.11.018.

Durran, D. R., 1989: Improving the anelastic approximation. *J. Atmos. Sci.*, **46**, 1453–1461, doi:10.1175/1520-0469(1989)046<1453:ITAA>2.0.CO;2.

Farrell, P. E., and J. R. Maddison, 2011: Conservative interpolation between volume meshes by local Galerkin projection. *Comput. Methods Appl. Mech. Eng.*, **200**, 89–100, doi:10.1016/j.cma.2010.07.015.

Ford, R., C. C. Pain, M. D. Piggott, A. J. H. Goddard, C. R. E. De Oliveira, and A. P. Umpleby, 2004: A nonhydrostatic finite-element model for three-dimensional stratified oceanic flows. Part I: Model formulation. *Mon. Wea. Rev.*, **132**, 2816–2831, doi:10.1175/MWR2824.1.

Giraldo, F. X., and T. Warburton, 2005: A nodal triangle-based spectral element method for the shallow water equations on the sphere. *J. Comput. Phys.*, **207**, 129–150, doi:10.1016/j.jcp.2005.01.004.

——, and M. Restelli, 2008: A study of spectral elements and discontinuous Galerkin methods for the Navier-Stokes equations in nonhydrostatic mesoscale atmospheric modeling: Equations sets and tests cases. *J. Comput. Phys.*, **227**, 3849–3877, doi:10.1016/j.jcp.2007.12.009.

——, J. S. Hesthaven, and T. Warburton, 2002: Nodal high-order discontinuous Galerkin methods for the spherical shallow water equations. *J. Comput. Phys.*, **181**, 499–525, doi:10.1006/jcph.2002.7139.

Harris, L. M., and S.-J. Lin, 2013: A two-way nested global-regional dynamical core on the cubed-sphere grid. *Mon. Wea. Rev.*, **141**, 283–306, doi:10.1175/MWR-D-11-00201.1.

Jablonowski, C., M. Herzog, J. Penner, R. Oehmke, Q. Stout, B. van Leer, and K. Powell, 2006: Block-structured adaptive grids on the sphere: Advection experiments. *Mon. Wea. Rev.*, **134**, 3691–3713, doi:10.1175/MWR3223.1.

Jiang, G. S., and C. W. Shu, 1996: Efficient implementation of weighted ENO schemes. *J. Comput. Phys.*, **126**, 202–228, doi:10.1006/jcph.1996.0130.

Karki, K. C., and S. V. Patankar, 1989: Pressure based calculation procedure for viscous flows at all speeds in arbitrary configurations. *AIAA J.*, **27**, 1167–1174, doi:10.2514/3.10242.

Kelly, J. F., and F. X. Giraldo, 2012: Continuous and discontinuous Galerkin methods for a scalable three-dimensional nonhydrostatic atmospheric model: Limited-area mode. *J. Comput. Phys.*, **231**, 7988–8008, doi:10.1016/j.jcp.2012.04.042.

Klemp, J. B., and D. K. Lilly, 1978: Numerical simulation of hydrostatic mountain waves. *J. Atmos. Sci.*, **35**, 78–107, doi:10.1175/1520-0469(1978)035<0078:NSOHMW>2.0.CO;2.

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

Marras, S., M. A. Kopera, and F. X. Giraldo, 2015: Simulation of shallow-water jets with a unified element-based continuous/discontinuous Galerkin model with grid flexibility on the sphere. *Quart. J. Roy. Meteor. Soc.*, **141**, 1727–1739, doi:10.1002/qj.2474.

——, and Coauthors, 2016: A review of element-based Galerkin methods for numerical weather prediction: Finite elements, spectral elements and discontinuous Galerkin. *Arch. Comput. Methods Eng.*, doi:10.1007/s11831-015-9152-1, in press.

McCorquodale, P., P. A. Ullrich, H. Johansen, and P. Colella, 2015: An adaptive multiblock high-order finite-volume method for solving the shallow-water equations on the sphere.

Commun. Appl. Math. Comput. Sci., **10**, 121–162, doi:10.2140/camcos.2015.10.121.

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

Nair, R. D., S. J. Thomas, and R. D. Loft, 2005: A discontinuous Galerkin transport scheme on the cubed sphere. *Mon. Wea. Rev.*, **133**, 814–828, doi:10.1175/MWR2890.1.

——, H.-W. Choi, and H. M. Tufo, 2009: Computational aspects of a scalable high-order discontinuous Galerkin atmospheric dynamical core. *Comput. Fluids*, **38**, 309–319, doi:10.1016/j.compfluid.2008.04.006.

Oberhuber, J. M., M. Herzog, H.-F. Graf, and K. Schwanke, 1998: Volcanic plume simulation on large scales. *J. Volcanol. Geotherm. Res.*, **87**, 29–53, doi:10.1016/S0377-0273(98)00099-7.

Pain, C. C., A. P. Umpleby, C. R. E. de Oliveira, and A. J. H. Goddard, 2001: Tetrahedral mesh optimisation and adaptivity for steady-state and transient finite element calculations. *Comput. Methods Appl. Mech. Eng.*, **190**, 3771–3796, doi:10.1016/S0045-7825(00)00294-2.

Peraire, J., and P. O. Persson, 2008: The compact discontinuous Galerkin (CDG) method for elliptic problems. *SIAM J. Sci. Comput.*, **30**, 1806–1824, doi:10.1137/070685518.

Persson, P. O., and J. Peraire, 2006: Sub-cell shock capturing for discontinuous Galerkin methods. *44th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, AIAA, 2006-112, doi:10.2514/6.2006-112.

Piggott, M. D., P. E. Farrell, C. R. Wilson, G. J. Gorman, and C. C. Pain, 2009: Anisotropic mesh adaptivity for multi-scale ocean modelling. *Philos. Trans. Roy. Soc. London*, **A367**, 4591–4611, doi:10.1098/rsta.2009.0155.

Qiu, J., and C.-W. Shu, 2005: Runge-Kutta discontinuous Galerkin method using WENO limiters. *SIAM J. Sci. Comput.*, **26**, 907–929, doi:10.1137/S1064827503425298.

Robert, A., 1993: Bubble convection experiments with a semi-implicit formulation of the Euler equations. *J. Atmos. Sci.*, **50**, 1865–1873, doi:10.1175/1520-0469(1993)050<1865:BCEWAS>2.0.CO;2.

Satoh, M., T. Matsuno, H. Tomita, H. Miura, T. Nasuno, and S. Iga, 2008: Nonhydrostatic icosahedral atmospheric model (NICAM) for global cloud resolving simulations. *J. Comput. Phys.*, **227**, 3486–3514, doi:10.1016/j.jcp.2007.02.006.

Schär, C., D. Leuenberger, O. Fuhrer, D. Lüthi, and C. Girard, 2002: A new terrain-following vertical coordinate formulation for atmospheric prediction models. *Mon. Wea. Rev.*, **130**, 2459–2480, doi:10.1175/1520-0493(2002)130<2459:ANTFVC>2.0.CO;2.

Schuster, D., S. Brdar, M. Baldauf, A. Dedner, R. Klöfkorn, and D. Kröher, 2014: On discontinuous Galerkin approach for atmospheric flow in the mesoscale with and without moisture. *Meteor. Z.*, **23**, 449–464, doi:10.1127/0941-2948/2014/0565.

Skamarock, W. C., and J. B. Klemp, 1994: Efficiency and accuracy of the Klemp–Wilhelmson time-splitting technique. *Mon. Wea. Rev.*, **122**, 2623–2630, doi:10.1175/1520-0493(1994)122<2623:EAAOTK>2.0.CO;2.

——, ——, M. G. Duda, L. D. Fowler, and S.-H. Park, 2012: A multi scale nonhydrostatic atmospheric model using centrical Voronoi tesselations and C-grid staggering. *Mon. Wea. Rev.*, **140**, 3090–3105, doi:10.1175/MWR-D-11-00215.1.

Slingo, J., and Coauthors, 2009: Developing the next-generation climate system models: challenges and achievements. *Philos. Trans. Roy. Soc. London*, **A367**, 815–831, doi:10.1098/rsta.2008.0207.

Smolarkiewicz, P. K., C. Kühnlein, and N. Wedi, 2014: A consistent framework for for discrete integrations of soundproof and compressible PDEs of atmospheric dynamics. *J. Comput. Phys.*, **263**, 185–205, doi:10.1016/j.jcp.2014.01.031.

St-Cyr, A., C. Jablonowski, J. M. Dennis, H. M. Tufo, and S. J. Thomas, 2008: A Comparison of two shallow water models with non-conforming adaptive grids. *Mon. Wea. Rev.*, **136**, 1898–1922, doi:10.1175/2007MWR2108.1.

Staniforth, A., and J. Thuburn, 2012: Horizontal girds for global weather and climate prediction models: A review. *Quart. J. Roy. Meteor. Soc.*, **138**, 1–26, doi:10.1002/qj.958.

Straka, J. M., R. B. Wilhelmson, L. J. Wicker, J. R. Anderson, and K. K. Droegemeier, 1993: Numerical solutions of a nonlinear density current: A benchmark solution and comparisons. *Int. J. Numer. Methods Fluids*, **17**, 1–22, doi:10.1002/fld.1650170103.

Thomas, S. J., and R. Loft, 2005: The NCAR spectral element climate dynamical core: semi-implicit Eulerian formulation. *J. Sci. Comput.*, **25**, 307–322, doi:10.1007/s10915-004-4646-2.

Thuburn, J., and C. J. Cotter, 2015: A primal-dual mimetic finite element scheme for the rotating shallow water equations on polygonal spherical meshes. *J. Comput. Phys.*, **290**, 274–297, doi:10.1016/j.jcp.2015.02.045.

Ullrich, P. A., 2014: A global finite-element shallow-water model supporting continuous and discontinuous elements. *Geosci. Model Dev.*, **7**, 3017–3035, doi:10.5194/gmd-7-3017-2014.

Vasilevskii, Y., and K. Lipnikov, 1999: An adaptive algorithm for quasi optimal mesh generation. *Comput. Math. Math. Phys.*, **39**, 1468–1486.

Weller, H., 2009: Predicting mesh density for adaptive modelling of the global atmosphere. *Philos. Trans. Roy. Soc. London*, **A367**, 4523–4542, doi:10.1098/rsta.2009.0151.

Wicker, L. J., and W. C. Skamarock, 1998: A time-splitting scheme for the elastic equations incorporating second-order Runge–Kutta time differencing. *Mon. Wea. Rev.*, **126**, 1992–1999, doi:10.1175/1520-0493(1998)126<1992:ATSSFT>2.0.CO;2.

Williamson, D. L., 2007: The evolution of dynamical cores for global atmospheric models. *J. Meteor. Soc. Japan*, **85B**, 241–269, doi:10.2151/jmsj.85B.241.

Yelash, L., A. Müller, M. Lukáčová-Medvid'ová, F. X. Giraldo, and V. Wirth, 2014: Adaptive discontinuous evolution Galerkin method for dry atmospheric flow. *J. Comput. Phys.*, **268**, 106–133, doi:10.1016/j.jcp.2014.02.034.

Yu, M., F. X. Giraldo, M. Peng, and Z. J. Wang, 2015: Localized artificial viscosity stabilization of discontinuous Galerkin methods for nonhydrostatic mesoscale atmospheric modeling. *Mon. Wea. Rev.*, **143**, 4823–4845, doi:10.1175/MWR-D-15-0134.1.