A primal-dual mimetic finite element scheme for the rotating shallow water equations on polygonal spherical meshes

John Thuburn[a], Colin J Cotter[b]

[a]*College of Engineering, Mathematics and Physical Sciences, University of Exeter, Exeter, UK*
[b]*Department of Mathematics, Imperial College, London, UK*

Abstract

A new numerical method is presented for solving the shallow water equations on a rotating sphere using quasi-uniform polygonal meshes. The method uses special families of finite element function spaces to mimic key mathematical properties of the continuous equations and thereby capture several desirable physical properties related to balance and conservation. The method relies on two novel features. The first is the use of *compound finite elements* to provide suitable finite element spaces on general polygonal meshes. The second is the use of *dual finite element spaces* on the dual of the original mesh, along with suitably defined discrete Hodge star operators to map between the primal and dual meshes, enabling the use of a finite volume scheme on the dual mesh to compute potential vorticity fluxes. The resulting method has the same mimetic properties as a finite volume method presented previously, but is more accurate on a number of standard test cases.

Keywords: compound finite element, dual finite element, mimetic, shallow water

1. Introduction

In order to exploit the new generation of massively parallel supercomputers that are becoming available, weather and climate models will require good parallel scalability. This requirement has driven the development of numerical methods that do not depend on the orthogonal coordinate system and quadrilateral structure of the longitude-latitude grid, whose polar resolution clustering is predicted to lead to a scalability bottleneck. A significant challenge is to obtain good scalability without sacrificing accuracy; in particular conservation, balance, and wave propagation are important for accurate modelling of the atmosphere (Staniforth and Thuburn, 2012).

Building on earlier work (Ringler et al., 2010; Thuburn and Cotter, 2012), Thuburn et al. (2014) presented a finite volume scheme for the shallow water equations on polygonal meshes. They start from the continuous shallow water equations in the so-called vector invariant form:

$$\phi_t + \nabla \cdot \mathbf{f} = 0, \qquad (1)$$

$$\mathbf{u}_t + \mathbf{q}^\perp + \nabla(\phi_T + k) = 0, \qquad (2)$$

where ϕ, the geopotential, is equal to the fluid depth times the gravitational acceleration, $\phi_T = \phi + \phi_{orog}$ is the total geopotential at the fluid's upper surface including the contribution from orography, $\mathbf{u}$ is the velocity, $\mathbf{f} = \mathbf{u}\phi$ is the mass flux, and $k = |\mathbf{u}|^2/2$. The $\perp$ symbol is defined by $\mathbf{u}^\perp = \mathbf{k} \times \mathbf{u}$ where $\mathbf{k}$ is the unit vertical vector. Finally, $\mathbf{q} = \mathbf{f}\pi$ is the flux of potential vorticity (PV), where $\pi = \zeta/\phi$ is the PV and $\zeta = f + \xi$ is the absolute vorticity, with f the Coriolis parameter and $\xi = \mathbf{k} \cdot \nabla \times \mathbf{u}$ the relative vorticity. (The notation used in the paper is summarized in Appendix B.) By the use of a C-grid placement of prognostic variables, and by ensuring that the numerical method mimics key mathematical properties of the continuous governing equations (hence the term 'mimetic'), the scheme was designed to have good conservation and balance properties. These good properties were verified in numerical tests on hexagonal and cubed sphere spherical meshes. However, their scheme has a number of drawbacks. Most seriously, the Coriolis operator, whose discrete form is essential to obtaining good geostrophic balance, is numerically inconsistent and fails to converge in the L_∞ norm (Weller, 2014; Thuburn et al., 2014). Also, although the gradient and divergence operators are consistent, their combination to form the discrete Laplacian operator also fails to converge in the L_∞ norm in some cases. These inaccuracies are clearly visible in idealized convergence tests, and give rise to marked 'grid imprinting' for initially symmetrical flows. Although they are less conspicuous in more complex flows, they are clearly undesirable.

Cotter and Shipton (2012) (see also McRae and Cotter, 2014; Cotter and Thuburn, 2014) showed that the same mimetic properties can be obtained using a certain class of mixed finite element method. The mimetic properties follow from the choice of an appropriate hierarchy of function spaces for the prognostic and diagnostic variables (e.g. section 3 below), which also provides a finite element analogue of the C-grid placement of variables, or a higher-order generalization. (The use of such a hierarchy goes by various names in the literature, including 'mimetic finite elements', 'compatible finite elements', and 'finite element exterior calculus'; see Cotter and Thuburn (2014) for a discussion of the shallow water equation case in the language of exterior calculus.) Importantly, the resulting schemes are numerically consistent.

While the mimetic finite element approach appears very attractive, it is not yet clear which particular choice of mesh and function spaces is most suitable. Standard finite element methods use triangular or quadrilateral elements. For the lowest-order mimetic finite element scheme on triangles, the dispersion relation for the linearized shallow water equations suffers from extra branches of inertio-gravity waves, which are badly behaved numerical artefacts (Le Roux et al., 2007), analogous to the problem that occurs on the triangular C-grid (Danilov, 2010). Higher-order finite element methods also typically exhibit anomalous features in their wave dispersion relations, such as extra branches, frequency gaps, or zero group velocity modes. Some progress has been made in reducing these problems, at least on quadrilateral meshes, through the inclusion of dissipation or modification of the mass matrix (e.g. Ullrich, 2013; Melvin et al., 2014), though the remedies are somewhat heuristic except in the most idealized cases. Finally, coupling to subgrid models of physical processes such as cumulus convection or cloud microphysics may be less straightforward with higher-order elements (P. Lauritzen, pers. comm.)[1]. These factors suggest that it may still be worthwhile investigating lowest-order schemes on quadrilateral and hexagonal meshes.

---

[1] The issues include, but are not limited to, (i) uneven spacing of the collocation points, especially in the vertical; (ii) whether the degrees of freedom should be treated as pointwise values, and, for discontinuous Galerkin schemes in particular, how to handle the case in which the data are discontinuous at an element boundary node; (iii) how to ensure positivity of fields like specific humidity; (iv) consistency of operators like divergence between the subgrid model and the dynamical core.

The above arguments raise two related questions. Can the mimetic finite element method inspire a development to fix the inconsistency of the mimetic finite volume method? Alternatively, can the mimetic finite element method at lowest order be adapted to work on polygonal meshes such as hexagons? Below we answer the second question by showing that the mimetic finite element method can indeed be adapted. In fact, from a certain viewpoint the mimetic finite volume and mimetic finite element schemes have very similar mathematical structure. The notation below is chosen to emphasize this similarity[2]. Moreover, the similarity is sufficiently strong that much of the code of the mimetic finite volume model of Thuburn et al. (2014) could be re-used in the model presented below. This, in turn, facilitates the cleanest possible comparison of the two approaches.

The adaptation of the mimetic finite element method employs two novel features. The first is the definition of a suitable hierarchy of finite element function spaces on polygonal meshes. Following Christiansen (2008), this is achieved by defining compound elements built out of triangular subelements, and is described in section 3. The second ingredient is the introduction of a dual family of function spaces that are defined on the dual of the original mesh. This permits the definition of a spatially averaged mass field that lives in the same function space as the vorticity and potential vorticity fields; this, in turn, enables the use of an accurate finite volume scheme on the dual mesh for advection of potential vorticity, and keeps the formulation of the finite element model as close as possible to that of the finite volume model.

2. Meshes and dual meshes

The scheme described here is suitable for arbitrary two-dimensional polygonal meshes on flat domains or, as used here, curved surfaces approximated by planar facets. Two particular meshes are used to obtain the results in section 5, namely the same variants of the hexagonal-icosahedral mesh and the cubed sphere mesh used by Thuburn et al. (2014), in order to facilitate comparison with their results. The hexagonal-icosahedral mesh uses the Heikes and Randall (1995b) optimization, which greatly improves the accuracy of the Laplacian for a finite volume scheme but has little effect for the present finite element scheme. The cubed sphere mesh is based on an equiangular cubed sphere, modified so that vertices are located at the barycenters of the surrounding cell centers. Again, this modification significantly improves the accuracy of a finite volume scheme but has little effect for the present finite element scheme. Coarse-resolutions versions of the two meshes are shown in Fig. 1.

Any polygonal mesh has a corresponding dual mesh. (We will refer to the original mesh as the 'primal' mesh where necessary to distinguish it from the dual.) Each primal cell contains one dual vertex; each dual cell contains one primal vertex; each primal edge corresponds to one dual edge and these usually cross each other. Figure 1 shows both primal and dual edges for the two meshes. For the hexagonal mesh, the dual edges are orthogonal to the primal edges. However, this is not true for the cubed sphere mesh.

3. Function spaces and compound finite elements

The mimetic properties of the scheme arise from the relationships between the finite element function spaces. Three function spaces are used on the primal mesh ($\mathbf{V}_0$, $\mathbf{V}_1$, and $\mathbf{V}_2$), and three

---

[2]Readers wishing to compare the two formulations should note that a different sign convention is used for the expansion coefficients of $\mathbf{k}\times$ any vector, such as $U^\perp$ in (48).
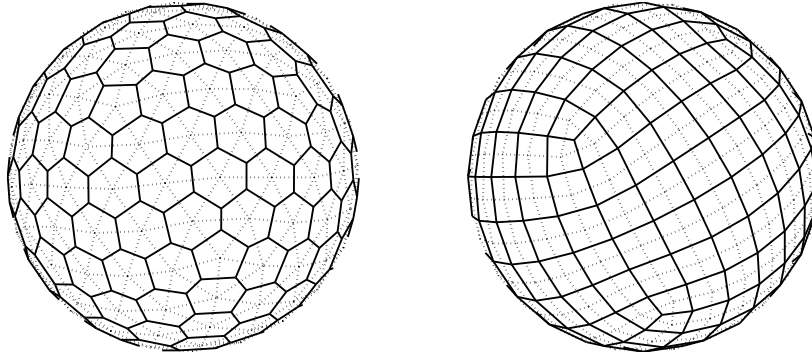
Figure 1: Left: a hexagonal–icosahedral mesh with 162 cells and 642 degrees of freedom. Right: a cubed-sphere mesh with 216 cells and 648 degrees of freedom. Continuous lines are primal mesh edges, dotted lines are dual mesh edges.
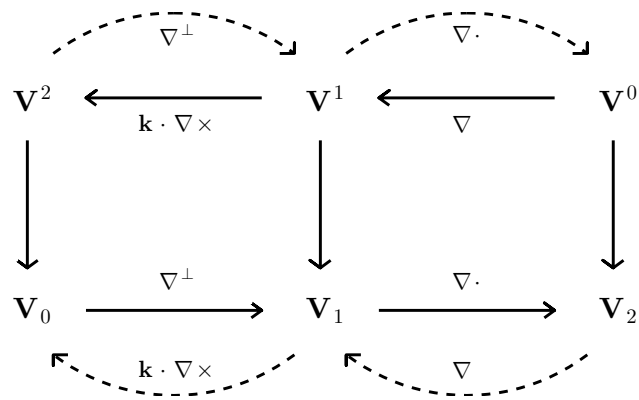

Figure 2: Schematic showing the function spaces used in the scheme and the relationships between them. Primal function spaces are on the bottom row and dual function spaces are on the top row.

on the dual mesh ($\mathbf{V}^2$, $\mathbf{V}^1$, and $\mathbf{V}^0$). Figure 2 indicates that $\nabla^\perp$ (i.e. $\mathbf{k} \times \nabla$) maps from $\mathbf{V}_0$ to $\mathbf{V}_1$ and $\nabla\cdot$ maps from $\mathbf{V}_1$ to $\mathbf{V}_2$[3]. More precisely, the primal function spaces satisfy the following properties.

Property List 1

- $\mathbf{u} \in \mathbf{V}_1 \;\; \Rightarrow \;\; \nabla \cdot \mathbf{u} \in \mathbf{V}_2$.

- $\phi \in \mathbf{V}_2$ with $\int \phi \, dA = 0 \;\; \Rightarrow \;\; \exists \mathbf{u} \in \mathbf{V}_1$ st $\nabla \cdot \mathbf{u} = \phi$.

- $\psi \in \mathbf{V}_0 \;\; \Rightarrow \;\; \nabla^\perp \psi \in \mathbf{V}_1$.

- $\forall \psi \in \mathbf{V}_0$, $\nabla \cdot \nabla^\perp \psi = 0$, and $\forall \mathbf{u} \in \mathbf{V}_1$ st $\nabla \cdot \mathbf{u} = 0$, $\exists \psi \in \mathbf{V}_0$ st $\mathbf{u} = \nabla^\perp \psi$. That is, $\nabla^\perp$ maps onto the kernel of $\nabla\cdot$.

The second condition assumes spherical geometry so that there are no lateral boundaries. The same assumption will be made throughout this paper; in particular, no boundary terms will arise when integrating by parts[4].

In a similar way, Fig. 2 indicates that ∇ maps from $\mathbf{V}^0$ to $\mathbf{V}^1$ and $\mathbf{k} \cdot \nabla\times$ maps from $\mathbf{V}^1$ to $\mathbf{V}^2$. More precisely, the dual function spaces satisfy the following properties.

Property List 2

- $\hat{\mathbf{u}} \in \mathbf{V}^1 \;\; \Rightarrow \;\; \mathbf{k} \cdot \nabla \times \hat{\mathbf{u}} \in \mathbf{V}^2$.

- $\hat{\xi} \in \mathbf{V}^2$ with $\int \hat{\xi} \, dA = 0 \;\; \Rightarrow \;\; \exists \hat{\mathbf{u}} \in \mathbf{V}^1$ st $\mathbf{k} \cdot \nabla \times \hat{\mathbf{u}} = \hat{\xi}$.

- $\hat{\chi} \in \mathbf{V}^0 \;\; \Rightarrow \;\; \nabla \hat{\chi} \in \mathbf{V}^1$.

- $\forall \hat{\chi} \in \mathbf{V}^0$, $\mathbf{k} \cdot \nabla \times \nabla \hat{\chi} = 0$, and $\forall \hat{\mathbf{u}} \in \mathbf{V}^1$ st $\mathbf{k} \cdot \nabla \times \hat{\mathbf{u}} = 0$, $\exists \hat{\chi} \in \mathbf{V}^0$ st $\hat{\mathbf{u}} = \nabla \hat{\chi}$. That is, ∇ maps onto the kernel of $\mathbf{k} \cdot \nabla\times$.

As noted earlier, standard finite element schemes in two dimensions typically use triangular or quadrilateral elements. Several families of mixed finite elements that satisfy Property List 1 on such meshes are known. However, in order to apply our scheme on more general polygonal meshes we will need families of mixed finite elements satisfying Property List 1 on those meshes. Building on the ideas of Buffa and Christiansen (2007), Christiansen (2008) showed how suitable

---

[3]Note $\nabla^\perp$ and $\mathbf{k} \cdot \nabla\times$ (like ∇ and $\nabla\cdot$) can both be defined as intrinsic operations on a curved surface, without reference to $\mathbf{k}$ or a third dimension.

[4]The most general form of the Helmholtz-Hodge decomposition of a vector field $\mathbf{u}$ in 2D is

$$\mathbf{u} = \nabla \phi + \nabla^\perp \psi + \mathbf{h},$$

where ϕ is a potential, ψ is a stream function, and $\mathbf{h}$ is a harmonic vector field, i.e. one satisfying

$$\nabla \nabla \cdot \mathbf{h} + \nabla^\perp \mathbf{k} \cdot \nabla \times \mathbf{h} = \mathbf{0}.$$

The fourth condition in Property List 1 implies that all nondivergent fields $\mathbf{u}$ can be written as $\nabla^\perp \psi$, which rules out the possibility of harmonic vector fields. This is appropriate for spherical geometry, since there exist no non-zero harmonic vector fields on the sphere. However, for a doubly period plane, for example, for which a constant vector field is harmonic, we would have to extend the fourth condition to allow for harmonic vector fields. This issue does not affect any of the discussion below except for the discrete Helmholtz decomposition (section 4.3), which would only need to be extended in the obvious way to allow for harmonic vector fields.

families of finite elements could be constructed by a method of *harmonic extension*; the key ideas are summarized below. The desire to use a dual mesh increases the need for finite element spaces on polygons. Only in special cases (such as the cubed sphere, Fig. 1) can both the primal and dual meshes be built of triangles and quadrilaterals; other cases require higher degree polygons for either the primal or dual mesh (or both).

Christiansen (2008) formally and rigorously presents the method of harmonic extension in the language of differential forms. Here we summarize the key ideas for the two-dimensional case in language that may be more familiar to some readers. Let γ_j be a basis function for $\mathbf{V}_0$ associated with primal vertex j. Define γ_j to equal 1 at vertex j and zero at all other vertices. Next extend γ_j harmonically along primal mesh edges; that is, its second derivative should vanish so that its gradient is constant along each edge. Then extend γ_j harmonically into the interior of each element; that is, solve

$$\nabla^2 \gamma_j = 0 \tag{3}$$

subject to the Dirichlet boundary conditions given by the known values of γ_j on element edges. In a similar way, let $\mathbf{v}_e$ be a basis function for $\mathbf{V}_1$ associated with edge e. Define the normal component of $\mathbf{v}_e$ to be a nonzero constant along edge e (some arbitrary sign convention must be chosen to define the positive direction) and zero at all other edges. Then extend $\mathbf{v}_e$ harmonically into the interior of each element; that is, solve

$$\nabla (\nabla \cdot \mathbf{v}_e) = 0 \tag{4}$$

and

$$\mathbf{k} \cdot \nabla \times \mathbf{v}_e = 0 \tag{5}$$

subject to the known values of the normal component at element edges. For the last function space $\mathbf{V}_2$ the basis function associated with cell i is defined to be a nonzero constant in cell i and zero in all other cells. It may then be verified that the properties in List 1 do indeed hold for the spaces spanned by these basis functions.

The difficulty with the method as just described is that, except for the simplest element shapes, the basis functions cannot be found analytically. However, it is possible to proceed as follows (Christiansen, 2008). Subdivide the mesh into triangular subelements (i.e., form a 'simplicial refinement'). On triangles, the lowest order element family satisfying List 1 is known and comprises continuous piecewise linear elements ($\mathbf{V}_0$), lowest order Raviart-Thomas elements ($\mathbf{V}_1$), and piecewise constant elements ($\mathbf{V}_2$); P1-RT0-P0^{DG} in standard shorthand. The basis elements on the original polygonal mesh are then constructed by solving discrete versions of (3), (4), and (5) on the simplicial refinement using P1-RT0-P0^{DG} element spaces. Although only discrete versions of (3), (4), and (5) are solved, it may be verified that the properties in List 1 hold exactly.

The resulting polygonal element basis functions are linear combinations of the triangular subelement basis functions, and provide a generalization to polygonal meshes of the P1-RT0-P0^{DG} hierarchy of spaces; thus we will refer to them as compound P1-RT0-P0^{DG} elements. The basis functions on the triangular subelements are known analytically; therefore, integrals of products of compound basis functions, for example to compute entries of a mass matrix, can all be computed exactly, either analytically or by numerical quadrature. The expansion coefficients for $\mathbf{V}_0$ correspond to mesh vertices, for $\mathbf{V}_1$ to edges, and for $\mathbf{V}_2$ to cells. Thus, this hierarchy provides a finite element analogue of the polygonal C-grid if we choose to represent velocity in $\mathbf{V}_1$ and the mass variable in $\mathbf{V}_2$.
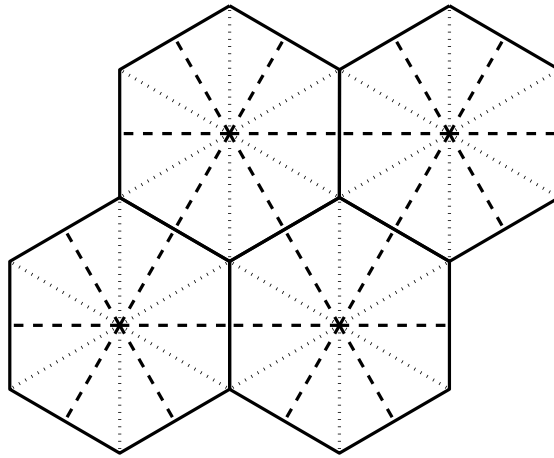
Figure 3: Example of part of a hexagonal primal mesh (solid lines) with its triangular dual mesh (dashed lines) and the supermesh of triangular subelements (all lines) used to construct the compound elements.

The construction of basis elements for the dual spaces proceeds in a very similar way, except that the basis function for $\mathbf{V}^1$ is given by $\mathbf{k}\times$ the solution of (4) and (5). This gives rise to a compound P1-N0-P0$^{\text{DG}}$ hierarchy of spaces, where N0 refers to the lowest order two-dimensional Nédélec elements (Nédélec, 1980).

An important detail concerns the number of subelements needed. It may appear natural to subdivide an n-gon cell into n triangular subelements. However, it will be necessary to calculate integrals on the overlap between primal and dual elements (section 4.2). In order to be able to do this when the domain is a curved surface approximated by plane triangular subelement facets, both the primal and dual compound element meshes must be built from triangular subelements of the the same supermesh. To achieve this we divide n-gon cells (whether primal or dual) into $2n$ subelements (Fig. 3).

It is convenient to normalize the basis functions as follows:

$$\alpha_i \in \mathbf{V}_2 : \qquad \int_{\text{cell } i'} \alpha_i \, dA = \delta_{i\,i'}; \tag{6}$$

$$\mathbf{v}_e \in \mathbf{V}_1 : \qquad \int_{\text{edge } e'} \mathbf{v}_e \cdot \mathbf{n} \, dl = \delta_{e\,e'}; \tag{7}$$

$$\gamma_j \in \mathbf{V}_0 : \qquad \gamma_j\big|_{\text{vertex } j'} = \delta_{j\,j'}; \tag{8}$$

$$\beta_j \in \mathbf{V}^2 : \qquad \int_{\text{dualcell } j'} \beta_j \, dA = \delta_{j\,j'}; \tag{9}$$

$$\mathbf{w}_e \in \mathbf{V}^1 : \qquad \int_{\text{dualedge } e'} \mathbf{w}_e \cdot \mathbf{m} \, dl = \delta_{e\,e'}; \tag{10}$$

$$\chi_i \in \mathbf{V}^0 : \qquad \chi_i\big|_{\text{dual vertex } i'} = \delta_{i\,i'}. \tag{11}$$

Here $\mathbf{n}$ is the unit normal vector to primal edge e and $\mathbf{m}$ is the unit tangent vector to dual edge e, with $\mathbf{m}$ and $\mathbf{n}$ pointing in the same sense (i.e. $\mathbf{n} \cdot \mathbf{m} > 0$, though they need not be parallel if the dual edges are not orthogonal to the primal edges), as in Thuburn and Cotter (2012). The normalization is chosen so that degrees of freedom for fields in $\mathbf{V}_2$ and $\mathbf{V}^2$ correspond to area integrals of scalars over primal cells and dual cells, respectively, degrees of freedom for a field in $\mathbf{V}_1$ correspond to normal fluxes integrated along primal edges, degrees of freedom in $\mathbf{V}^1$ correspond to circulations integrated along dual edges, and degrees of freedom for fields in $\mathbf{V}_0$ and $\mathbf{V}^0$ correspond to nodal values of scalars at primal vertices and dual vertices respectively. Again, this corresponds closely to the framework of Thuburn and Cotter (2012).

Melvin and Thuburn (2015) have analyzed the wave dispersion properties for finite element discretizations of the linear shallow water equations using these compound elements. That paper gives explicit expressions for the $\mathbf{V}_1$ and $\mathbf{V}_2$ compound element basis functions for the cases of a square mesh and a regular hexagonal mesh on a plane. On both meshes the compound elements give a noticeable improvement in accuracy compared to a C-grid finite difference scheme, while on a square mesh the compound element results are very similar to those obtained using standard quadrilateral P1-RT0-P0$^{\text{DG}}$ elements. The numbers of wave modes of each type are exactly the same as on a finite difference C-grid. For more general meshes it is straightforward and convenient to construct the compound element basis functions numerically rather than analytically. Figure 4 shows typical basis elements for the three primal mesh function spaces for quadrilateral and hexagonal cells.

The fields used in the computation are represented as expansions in terms of these basis elements. For example,

$$\phi = \sum_i \phi_i \alpha_i \qquad \in \mathbf{V}_2, \tag{12}$$

$$\mathbf{u} = \sum_e u_e \mathbf{v}_e \qquad \in \mathbf{V}_1 \tag{13}$$

for the prognostic geopotential and velocity fields, and

$$\xi = \sum_j \xi_j \gamma_j \qquad \in \mathbf{V}_0 \tag{14}$$

for the relative vorticity field. Here the sums are global sums over all basis elements in the relevant spaces. In some cases it will be useful to introduce dual space representations of fields; these
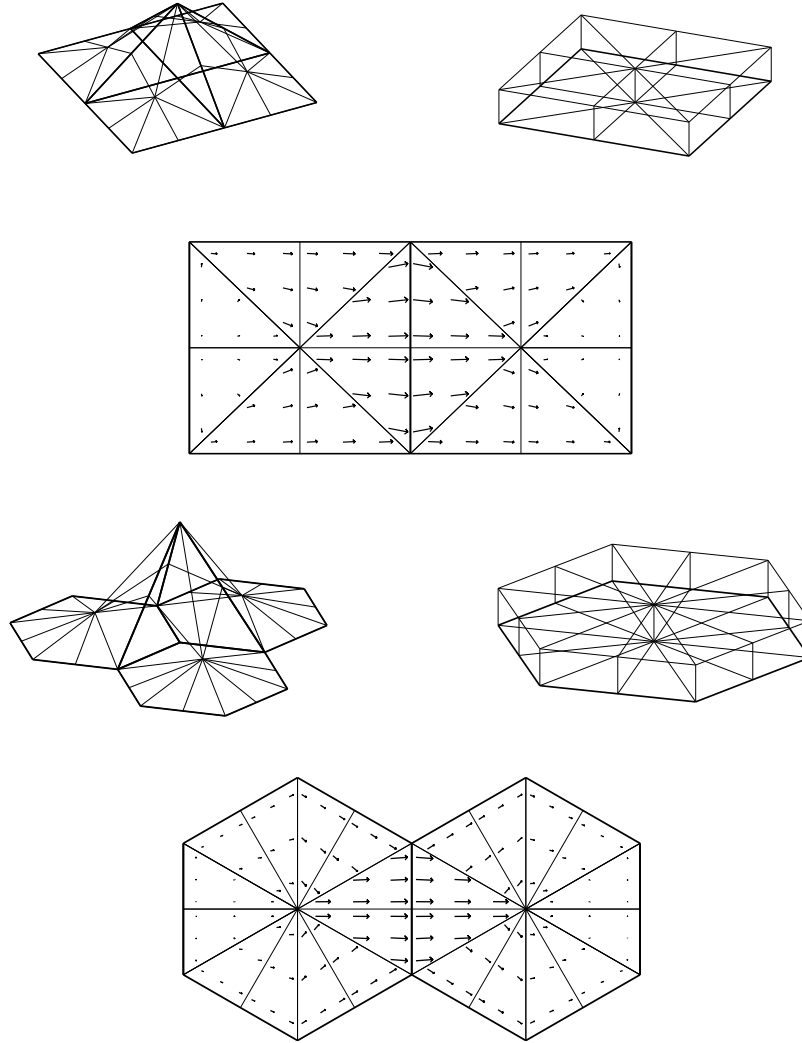
8

Figure 4: Typical compound basis elements. First row: $\gamma_j \in \mathbf{V}_0$ and $\alpha_i \in \mathbf{V}_2$, square mesh; second row: $\mathbf{v}_e \in \mathbf{V}_1$, square mesh; third row: $\gamma_j \in \mathbf{V}_0$ and $\alpha_i \in \mathbf{V}_2$, hexagonal mesh; fourth row: $\mathbf{v}_e \in \mathbf{V}_1$, hexagonal mesh.

will be indicated by a hat symbol where necessary to distinguish them from the corresponding primal space representations. For example,

$$\hat{\phi} = \sum_i \hat{\phi}_i \chi_i \qquad \in \mathbf{V}^0, \tag{15}$$

$$\hat{\mathbf{u}} = \sum_e \hat{u}_e \mathbf{w}_e \qquad \in \mathbf{V}^1, \tag{16}$$

$$\hat{\xi} = \sum_j \hat{\xi}_j \beta_j \qquad \in \mathbf{V}^2. \tag{17}$$

The fields ϕ and $\hat{\phi}$ have the same number of degrees of freedom, and it is possible construct a map between them by demanding that they agree when integrated against any test function in the primal space $\mathbf{V}_2$. Similarly, the fields $\mathbf{u}$ and $\hat{\mathbf{u}}$ have the same number of degrees of freedom, and it is possible construct a map between them by demanding that they agree when integrated against any test function in the primal space $\mathbf{V}_1$. On the meshes we have used we have found these maps to be reversible and well-conditioned. It will also be useful to introduce spatially averaged versions of some fields. For example,

$$\widetilde{\phi} = \sum_j \widetilde{\phi}_j \gamma_j \qquad \in \mathbf{V}_0, \tag{18}$$

$$\overline{\phi} = \sum_j \overline{\phi}_j \beta_j \qquad \in \mathbf{V}^2. \tag{19}$$

Here, $\widetilde{\phi}$ and $\overline{\phi}$ have the same number of degrees of freedom, and it is possible construct a map between them by demanding that they agree when integrated against any test function in the primal space $\mathbf{V}_0$. $\widetilde{\phi}$ or $\overline{\phi}$ can be obtained from ϕ by demanding that they agree when integrated against any test function in $\mathbf{V}_0$; in effect this provides an averaging operation from $\mathbf{V}_2$ to $\mathbf{V}_0$ or $\mathbf{V}^2$. (However, we should not expect to be able to obtain ϕ from $\widetilde{\phi}$ or $\overline{\phi}$, as this would require an un-averaging operation, which will be ill-conditioned if it exists at all.)

It will be convenient to be able to refer to the vector of degrees of freedom for any field. To do this, we will use the same letter (with hat, tilde or bar if needed) but in upper case. Thus, for example, Φ will be the vector of values $(\phi_1, \phi_2, \dots)^T$, $\hat{U}$ will be the vector of values $(\hat{u}_1, \hat{u}_2, \dots)^T$, etc.

4. Finite element scheme

Finite element schemes solve the governing equations by approximating the solution in the chosen function spaces, written as expansions in terms of basis functions (e.g. (12), (13)), and demanding that the equations be satisfied in weak form, that is, when multiplied by any test function in the appropriate space and integrated over the domain. In this approach (1) becomes

$$\int \alpha_i \left(\phi_t + \nabla \cdot \mathbf{f} \right) dA = 0 \quad \forall \alpha_i \in \mathbf{V}_2, \tag{20}$$

or, regarding the integral as an inner product for which we introduce angle backet notation,

$$\langle \alpha_i, \phi_t \rangle + \langle \alpha_i, \nabla \cdot \mathbf{f} \rangle = 0 \quad \forall \alpha_i \in \mathbf{V}_2. \tag{21}$$

10

Similarly, (2) becomes

$$\int \mathbf{v}_e \cdot \left\{ \mathbf{u}_t + \mathbf{q}^\perp + \nabla(\phi_T + k) \right\} dA = 0 \quad \forall \mathbf{v}_e \in \mathbf{V}_1, \tag{22}$$

or

$$\langle \mathbf{v}_e, \mathbf{u}_t \rangle + \langle \mathbf{v}_e, \mathbf{q}^\perp \rangle + \langle \mathbf{v}_e, \nabla(\phi_T + k) \rangle = 0. \tag{23}$$

(The construction of the nonlinear terms $\mathbf{f}$, $\mathbf{q}$ and k is discussed in section 4.6 below.) The method generally leads to a system of algebraic equations for the unknown coefficients in the expansion of the solution.

The following subsections show how the mimetic finite element method can be re-expressed in terms of certain matrix operators acting on the coefficient vectors Φ, U, etc. The notation is chosen to highlight the similarity to the finite volume scheme of Thuburn et al. (2014).

4.1. Matrix representation of derivatives – strong derivatives

The velocity basis elements are constructed and normalized so as to have constant divergence over the cell upwind of the edge where the degree of freedom resides, with area integral equal to 1, and constant divergence over the cell downwind of this edge, with area integral equal to -1, with zero velocity and hence zero divergence in all other cells. Thus

$$\nabla \cdot \mathbf{v}_e = \sum_i n_{e\,i} \alpha_i \quad \in \mathbf{V}_2, \tag{24}$$

where $n_{e\,i}$ is equal to 1 when the normal at edge e points out of cell i, equal to -1 when the normal at edge e points into cell i, and is zero otherwise. We will write D_2 for the matrix whose transpose has components $n_{e\,i}$. D_2 is called an incidence matrix because it describes some aspects of the grid topology. Hence, the divergence δ of an arbitrary velocity field $\mathbf{u}$ is

$$\sum_i \delta_i \alpha_i = \delta = \nabla \cdot \mathbf{u} = \sum_e u_e \nabla \cdot \mathbf{v}_e = \sum_{e\,i} u_e n_{e\,i} \alpha_i. \tag{25}$$

Equating coefficients of α_i gives

$$\delta_i = \sum_e n_{e\,i} u_e, \tag{26}$$

or, in matrix-vector notation

$$\Delta = D_2 U. \tag{27}$$

Note we could have demanded that (25) should hold when integrated against any test function in $\mathbf{V}_2$, to obtain the same result. However, this would obscure the fact that (25) actually holds at every point in the domain (except on cell edges where all terms are discontinuous), not just when integrated against a test function. In this sense, $\nabla \cdot : \mathbf{V}_1 \to \mathbf{V}_2$ is a *strong* derivative operator.

Similarly, the basis elements in $\mathbf{V}_0$ are constructed so that

$$\nabla^\perp \gamma_j = \sum_e -t_{e\,j} \mathbf{v}_e, \tag{28}$$

where $t_{e\,j}$ is defined to equal 1 if edge e is incident on vertex j and the unit tangent vector $\mathbf{t}$ at edge e points towards vertex j, -1 if it points away from vertex j, and zero otherwise. The unit

normal and unit tangent at any edge are related by $\mathbf{t} = \mathbf{k} \times \mathbf{n}$. Hence, a stream function ψ is related to the corresponding rotational velocity field $\mathbf{u}$ by

$$\sum_e u_e \mathbf{v}_e = \mathbf{u} = \nabla^\perp \psi = \sum_j \psi_j \nabla^\perp \gamma_j = \sum_{je} -\psi_j t_{ej} \mathbf{v}_e. \tag{29}$$

Equating coefficients of $\mathbf{v}_e$ and defining D_1 to be the matrix whose entries are t_{ej} gives the matrix-vector form

$$U = -D_1 \Psi. \tag{30}$$

Equation (29) holds pointwise (again with the exception of discontinuities), so $\nabla^\perp : \mathbf{V}_0 \to \mathbf{V}_1$ is a strong derivative operator.

The matrices D_1 and D_2 are exactly the same as in the finite volume framework of Thuburn and Cotter (2012). In particular, they have the property that

$$D_2 D_1 \equiv 0, \tag{31}$$

giving a discrete analogue of the continuous property $\nabla \cdot \nabla^\perp \equiv 0$.

Analogous relations hold on the dual spaces.

$$\nabla \chi_i = -\sum_e n_{ei} \mathbf{w}_e \tag{32}$$

implies that the discrete analogue of

$$\hat{\mathbf{u}} = \nabla \hat{p} \tag{33}$$

is

$$\hat{U} = \overline{D}_1 \hat{P}, \tag{34}$$

where $\overline{D}_1 = -D_2^T$. Similarly

$$\mathbf{k} \cdot \nabla \times \mathbf{w}_e = \sum_j t_{ej} \beta_j \tag{35}$$

implies that the discrete analogue of

$$\hat{\xi} = \mathbf{k} \cdot \nabla \times \hat{\mathbf{u}} \tag{36}$$

is

$$\hat{\Xi} = \overline{D}_2 \hat{U}, \tag{37}$$

where $\overline{D}_2 = D_1^T$. Again, these are strong derivative operators.

The matrices $\overline{D}_1$ and $\overline{D}_2$ have the property

$$\overline{D}_2 \overline{D}_1 \equiv 0, \tag{38}$$

giving a discrete analogue in the dual space of the continuous relation $\mathbf{k} \cdot \nabla \times \nabla \equiv 0$.

4.2. Mass matrices and other operators

Define the following mass matrices for the primal function spaces:

$$L_{i\,i'} = \langle \alpha_i, \alpha_{i'} \rangle = \int \alpha_i \alpha_{i'} \, dA, \qquad (\mathbf{V}_2 \to \mathbf{V}_2), \tag{39}$$

$$M_{e\,e'} = \langle \mathbf{v}_e, \mathbf{v}_{e'} \rangle = \int \mathbf{v}_e \cdot \mathbf{v}_{e'} \, dA, \qquad (\mathbf{V}_1 \to \mathbf{V}_1), \tag{40}$$

$$N_{j\,j'} = \langle \gamma_j, \gamma_{j'} \rangle = \int \gamma_j \gamma_{j'} \, dA, \qquad (\mathbf{V}_0 \to \mathbf{V}_0). \tag{41}$$

The expressions in parentheses indicate that L maps $\mathbf{V}_2$ to itself, etc. (Analogous mass matrices may be defined for the dual spaces; however, they will not be needed here.)

The following matrices are also needed.

$$R_{j\,i} = \langle \gamma_j, \alpha_i \rangle, \qquad (\mathbf{V}_2 \to \mathbf{V}_0), \tag{42}$$

$$W_{e\,e'} = -\langle \mathbf{v}_e, \mathbf{v}_{e'}^{\perp} \rangle = -W_{e'\,e}, \qquad (\mathbf{V}_1 \to \mathbf{V}_1), \tag{43}$$

$$H_{e\,e'} = \langle \mathbf{v}_e, \mathbf{w}_{e'} \rangle, \qquad (\mathbf{V}^1 \to \mathbf{V}_1), \tag{44}$$

$$J_{j\,j'} = \langle \gamma_j, \beta_{j'} \rangle, \qquad (\mathbf{V}^2 \to \mathbf{V}_0). \tag{45}$$

For completeness we may also define

$$I_{i\,i'} = \langle \alpha_i, \chi_{i'} \rangle, \qquad (\mathbf{V}^0 \to \mathbf{V}_2), \tag{46}$$

though we will not need to employ this matrix in the shallow water scheme.

One further operator will be needed to construct the kinetic energy per unit mass. It is

$$T_{i\,e\,e'} = \int_{\text{cell } i} \mathbf{v}_e \cdot \mathbf{v}_{e'} \, dA = A_i \langle \alpha_i, \mathbf{v}_e \cdot \mathbf{v}_{e'} \rangle \quad (\mathbf{V}_1 \otimes \mathbf{V}_1 \to \mathbf{V}_2). \tag{47}$$

where $A_i = (L_{ii})^{-1}$ is the area of primal cell i.

All of these matrices can be precomputed, so that no quadrature needs to be done at run time. Moreover, they are all sparse, so they can be efficiently stored as lists of stencils and coefficients.

Let $U^{\perp}$ be the coefficients of the expansion of the projection of $\mathbf{u}^{\perp}$ into $\mathbf{V}_1$:

$$\langle \mathbf{v}_e, \sum_{e'} U_{e'}^{\perp} \mathbf{v}_{e'} \rangle = \langle \mathbf{v}_e, \sum_{e'} U_{e'} \mathbf{v}_{e'}^{\perp} \rangle = \langle \mathbf{v}_e, \mathbf{u}^{\perp} \rangle \qquad \forall \mathbf{v}_e \in \mathbf{V}_1. \tag{48}$$

Using (40) and (43) gives the discrete version of the $\perp$ operator:

$$MU^{\perp} = -WU. \tag{49}$$

Demanding agreement between (12) and (15) when integrated against any test function in $\mathbf{V}_2$ leads to

$$L\Phi = I\hat{\Phi}. \tag{50}$$

Similarly, demanding agreement between (13) and (16) when integrated against any test function in $\mathbf{V}_1$ gives

$$MU = H\hat{U}, \tag{51}$$

13

while demanding agreement between (14) and (17) when integrated against any test function in $\mathbf{V}_0$ gives

$$N\Xi = J\hat{\Xi}. \tag{52}$$

The relations (50), (51), (52) provide invertible maps between the primal and dual function spaces. Thus, they are examples of discrete Hodge star operators (e.g. Hiptmair, 2001). They may be contrasted with the analogous relations employed by Thuburn and Cotter (2012) and Thuburn et al. (2014) for the finite volume case, which do not involve mass matrices.

Demanding agreement between (12), (18), and (19) when integrated against any test function in $\mathbf{V}_0$ leads to

$$N\widetilde{\Phi} = J\overline{\Phi} = R\Phi. \tag{53}$$

This is the matrix representation of the averaging operator discussed in section 3.

4.3. Matrix representation of derivatives – weak derivatives

A field in $\mathbf{V}_2$ is discontinuous, so its gradient in $\mathbf{V}_1$ can only be defined in a *weak* sense, by integrating against all test functions in $\mathbf{V}_1$. For example,

$$\mathbf{g} = \nabla\phi \tag{54}$$

must be approximated as

$$\langle \mathbf{v}_e, \mathbf{g} \rangle = \langle \mathbf{v}_e, \nabla\phi \rangle \quad \forall \mathbf{v}_e \in \mathbf{V}_1, \tag{55}$$

where $\phi \in \mathbf{V}_2$, $\mathbf{g} \in \mathbf{V}_1$. Expanding both ϕ and $\mathbf{g}$ in terms of basis elements and integrating by parts then leads to the matrix form

$$MG = \overline{D}_1 L\Phi. \tag{56}$$

Similarly, the curl of a vector field in $\mathbf{V}_1$ must be defined by integration against all test functions in $\mathbf{V}_0$. For example, the discrete analogue of

$$\xi = \mathbf{k} \cdot \nabla \times \mathbf{u}, \tag{57}$$

after expanding in basis functions and integrating by parts, is

$$N\Xi = \overline{D}_2 MU. \tag{58}$$

Combining these two results, the discrete analogue of

$$z = \mathbf{k} \cdot \nabla \times \nabla\phi \tag{59}$$

is

$$NZ = \overline{D}_2\overline{D}_1 L\Phi, \tag{60}$$

which is identically zero.

These derivative operators can be combined to obtain the Laplacian of a scalar. For a scalar $\phi \in \mathbf{V}_2$, the discrete Laplacian is $D_2 M^{-1}\overline{D}_1 L\Phi$. For a scalar $\psi \in \mathbf{V}_0$, the discrete Laplacian is $-N^{-1}\overline{D}_2 MD_1\Psi$. The operators introduced above lead to a discrete version of the Helmholtz decomposition, in which an arbitrary vector field is decomposed into its divergent and rotational parts:

$$U = M^{-1}\overline{D}_1 L\Phi - D_1\Psi. \tag{61}$$

Figure 5 summarizes how the operators introduced here map between the different function spaces.
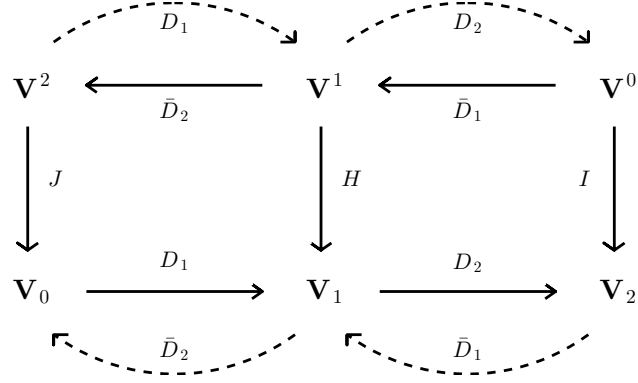
Figure 5: Schematic showing the function spaces used in the scheme and how the various matrices introduced above map between them.

4.4. Some operator identities

The operators defined above satisfy some key relations that underpin the mimetic properties of the scheme. We have already seen that

$$D_2 D_1 \equiv 0 \tag{62}$$

and

$$\overline{D}_2 \overline{D}_1 \equiv 0, \tag{63}$$

leading to discrete analogues of $\nabla \cdot \nabla^\perp \equiv 0$ and $\mathbf{k} \cdot \nabla \times \nabla \equiv 0$.

Next, note that the basis elements γ_j give a partition of unity, that is

$$\sum_j \gamma_j = 1 \tag{64}$$

at every point in the domain. Consequently

$$\sum_j R_{ji} = \langle 1, \alpha_i \rangle = 1 \tag{65}$$

and

$$\sum_j J_{jj'} = \langle 1, \beta_{j'} \rangle = 1. \tag{66}$$

Now let

$$\psi = \sum_j \psi_j \gamma_j \quad \in \mathbf{V}_0. \tag{67}$$

By considering the projection of $\nabla \psi$ into $\mathbf{V}_1$

$$\langle \mathbf{v}_e, (\nabla^\perp \psi)^\perp \rangle = -\langle \mathbf{v}_e, \nabla \psi \rangle \tag{68}$$

15

and integrating by parts and using the matrices defined in sections 4.1 and 4.2, we obtain

$$-\overline{D}_2 W = R D_2. \tag{69}$$

This identity is key to obtaining the steady geostrophic mode property (section 4.5.3 below). A rough interpretation is that averaging velocities to construct the Coriolis terms (W) then taking their divergence ($\overline{D}_2$) gives the same result as computing the velocity divergence (D_2) followed by averaging to $\mathbf{V}_0$ (R). One consequence is $\overline{D}_2 W D_1 = -R D_2 D_1 \equiv 0$.

An identical formula to (69) relating R and W was obtained by Thuburn and Cotter (2012) for the finite volume case. The result was originally derived for the construction of the Coriolis terms on orthogonal grids by Thuburn et al. (2009), and Thuburn and Cotter (2012) showed that it could be embedded in a more general framework applicable to nonorthogonal grids. Moreover, Thuburn et al. (2009) showed that, for any given R with the appropriate stencil (which we have here) and satisfying (65), there is a unique antisymmetric W satisfying (69), and gave an explicit construction for W in terms of R. Thus, although the context and interpretation are slightly different here, we can, nevertheless, use the Thuburn et al. construction in implementing the mixed finite-element version of the W operator!

Now consider the two representations of any vector field $\mathbf{u} \in \mathbf{V}_1$, $\hat{\mathbf{u}} \in \mathbf{V}^1$ related by

$$\langle \mathbf{v}, \mathbf{u} \rangle = \langle \mathbf{v}, \hat{\mathbf{u}} \rangle \quad \forall \mathbf{v} \in \mathbf{V}_1, \tag{70}$$

so that

$$MU = H\hat{U}. \tag{71}$$

Since $\nabla^\perp \gamma_j \in \mathbf{V}_1$,

$$\langle \nabla^\perp \gamma_j, \mathbf{u} \rangle = \langle \nabla^\perp \gamma_j, \hat{\mathbf{u}} \rangle \quad \forall \gamma_j \in \mathbf{V}_0, \tag{72}$$

and integrating by parts gives

$$-\langle \gamma_j, \mathbf{k} \cdot \nabla \times \mathbf{u} \rangle = -\langle \gamma_j, \mathbf{k} \cdot \nabla \times \hat{\mathbf{u}} \rangle \quad \forall \gamma_j \in \mathbf{V}_0. \tag{73}$$

Hence

$$-\overline{D}_2 M U = -J \overline{D}_2 \hat{U}. \tag{74}$$

Finally, substituting from (71) and noting that $\hat{U}$ is arbitrary gives

$$\overline{D}_2 H = J \overline{D}_2. \tag{75}$$

The interpretation of this identity is that, for a velocity field in $\mathbf{V}^1$, taking the curl followed by mapping to the primal space is equivalent to mapping the velocity field to the primal space then taking its curl. One consequence is that $\overline{D}_2 H \overline{D}_1 = J \overline{D}_2 \overline{D}_1 \equiv 0$.

Finally, let $\chi \in \mathbf{V}_2$ and $\hat{\chi} \in \mathbf{V}^0$ be two discrete representations of a scalar field related by

$$\langle \alpha, \chi \rangle = \langle \alpha, \hat{\chi} \rangle \quad \forall \alpha \in \mathbf{V}_2, \tag{76}$$

so that

$$LX = I\hat{X}. \tag{77}$$

Since $\nabla \cdot \mathbf{v}_e \in \mathbf{V}_2$ for any $\mathbf{v}_e \in \mathbf{V}_1$, we have

$$\begin{aligned}
\langle \nabla \cdot \mathbf{v}_e, \chi \rangle &= \langle \nabla \cdot \mathbf{v}_e, \hat{\chi} \rangle, \\
\langle \mathbf{v}_e, \nabla \chi \rangle &= \langle \mathbf{v}_e, \nabla \hat{\chi} \rangle, \\
\overline{D}_1 L X &= H \overline{D}_1 \hat{X},
\end{aligned} \tag{78}$$

16

or, using (77) and noting that $\hat{X}$ is arbitrary,

$$\overline{D}_1 I = H\overline{D}_1. \tag{79}$$

Using these identities and the Hodge star operators, it can be seen that taking a weak derivative in the primal space is equivalent to applying a Hodge star to map to the dual space, taking a strong derivative in the dual space, and applying another Hodge star to map back to the primal space:

$$M^{-1}\overline{D}_1 L = (M^{-1}H)\overline{D}_1(I^{-1}L); \tag{80}$$
$$N^{-1}\overline{D}_2 M = (N^{-1}J)\overline{D}_2(H^{-1}M) \tag{81}$$

(Cotter and Thuburn, 2014). Thus, certain paths in Fig. 5 commute. Weak derivative operators in the dual space can be defined by demanding a similar equivalence with primal space strong derivatives; however, the resulting formulas are less elegant and, in any case, will not be needed.

4.5. Linear shallow water equations

We first examine the spatial discretization of the linear shallow water equations to illustrate how some key conservation and balance properties arise. The rotating shallow water equations (1), (2) when linearized about a resting basic state with constant geopotential ϕ_0 and with constant Coriolis parameter f become

$$\phi_t + \nabla \cdot (\phi_0 \mathbf{u}) = 0, \tag{82}$$
$$\mathbf{u}_t + f\mathbf{u}^\perp + \nabla\phi = 0. \tag{83}$$

By writing these in weak form (analogous to (21) and (23)), expanding ϕ and $\mathbf{u}$ in terms of basis functions, and using the notation and operators defined above, we obtain

$$\dot{\Phi} + \phi_0 D_2 U = 0, \tag{84}$$
$$M\dot{U} - fWU + \overline{D}_1 L\Phi = 0. \tag{85}$$

4.5.1. Mass conservation

Mass conservation is trivially satisfied (for both the linear and nonlinear equations) because the discrete divergence is a strong operator, so the domain integral of the discrete divergence of any vector field vanishes.

4.5.2. Energy conservation

For the linearized equations the total energy is given by

$$
\begin{aligned}
E &= \frac{1}{2}\int \phi^2 + \phi_0 \mathbf{u} \cdot \mathbf{u}\, dA \\
&= \frac{1}{2}\Phi^T L\Phi + \frac{1}{2}\phi_0 U^T MU.
\end{aligned} \tag{86}
$$

Hence, the rate of change of total energy is

$$
\begin{aligned}
\frac{dE}{dt} &= \Phi^T L\dot{\Phi} + \phi_0 U^T M\dot{U} \\
&= -\phi_0\Phi^T LD_2 U + \phi_0 U^T(fWU - \overline{D}_1 L\Phi) \\
&= 0, \tag{87}
\end{aligned}
$$

where we have used the fact that L and M are symmetric, W is antisymmetric, and $D_2^T = -\overline{D}_1$.

17

4.5.3. Steady geostrophic modes

The linear shallow water equations support steady non-divergent flows in geostrophic balance. A numerical method must respect this property in order to be able to represent geostrophic balance. However, it is non-trivial to achieve this property because several ingredients must fall into place.

- The geopotential ϕ must be steady. The steadiness of ϕ follows immediately from the assumption that $\nabla \cdot \mathbf{u} = 0$.

- The relative vorticity ξ must be steady; neither the pressure gradient nor the Coriolis term should generate vorticity. First note that, from Property List 1, $U = -D_1\Psi$ for some Ψ. Taking the curl of the momentum equation then gives

$$N\dot{\Xi} = \overline{D_2}M\dot{U} = \overline{D_2}(-fWD_1\Psi - \overline{D_1}L\Phi). \tag{88}$$

The pressure gradient term does not contribute because $\overline{D_2}\,\overline{D_1} \equiv 0$, and the Coriolis term does not contribute because $\overline{D_2}WD_1 \equiv 0$.

- There must exist a geopotential ϕ that balances the Coriolis term so that the divergence is steady. Taking the divergence of the momentum equation gives the divergence tendency

$$\dot{\Delta} = D_2 M^{-1}(-fWD_1\Psi - \overline{D_1}\Phi). \tag{89}$$

If we define $\Phi = fL^{-1}R^T\Psi$ and use the transpose of (69) we find that $\dot{\Delta}$ does indeed vanish; thus the required ϕ does exist.

Consequently, the scheme does support steady geostrophic modes for the linearized equations. (Note, it is not necessarily true that any given ϕ field can be balanced by some non-divergent velocity field. On some meshes, particularly those with triangular primal cells, there might not be enough velocity degrees of freedom to balance all possible ϕ fields.)

4.5.4. Linear PV equation

A generalization of the steady geostrophic mode property is that the scheme should have a suitable PV equation. In this section we consider the linear case; the nonlinear case is dealt with in section 4.6.3.

The mass field ϕ and the vorticity field ξ live in different spaces. To construct a suitable discrete PV we need an averaged mass field $\widetilde{\phi}$ that lives in the same space as ξ. The linearized PV should be independent of time. For this to hold, ξ and $\widetilde{\phi}$ must see the same divergence field.

For a general (possibly divergent) velocity field $\mathbf{u}$, the vorticity equation (88) becomes

$$N\dot{\Xi} = f\overline{D_2}WU. \tag{90}$$

Define $\widetilde{\phi}$ using (53). Then the evolution of $\widetilde{\phi}$ is given by

$$
\begin{aligned}
N\dot{\widetilde{\Phi}} &= R\dot{\Phi} \\
&= -\phi_0 RD_2 U \\
&= \phi_0\overline{D_2}WU, \tag{91}
\end{aligned}
$$

(using (69)). Thus, Ξ and $\widetilde{\Phi}$ see the same divergence $-\overline{D_2}WU$; consequently the linearized PV $\Xi/\phi_0 - f\widetilde{\Phi}/\phi_0^2$ is independent of time.

4.6. Nonlinear shallow water equations

The nonlinear rotating shallow water equations are (1) and (2). Writing these in weak form (21) and (23), and letting F, Q, and K be the vectors of coefficients for the discrete representations of the mass flux $\mathbf{f} = \mathbf{u}\phi$, the PV flux $\mathbf{q} = \mathbf{f}\pi$, and the kinetic energy per unit mass $k = \mathbf{u} \cdot \mathbf{u}/2$, the nonlinear discretization becomes

$$\dot{\Phi} + D_2 F \quad = \quad 0, \tag{92}$$

$$M\dot{U} + MQ^{\perp} + \overline{D}_1 L(\Phi_T + K) \quad = \quad 0. \tag{93}$$

The remaining issue is how to construct suitable values of the three nonlinear terms K, F, and $Q^{\perp}$.

4.6.1. Constructing K

The discretization of k follows the standard finite element construction, which is to project ∇k into $\mathbf{V}_1$. It may easily be verified that this is equivalent to projecting k into $\mathbf{V}_2$ before taking the weak gradient. Using the T operator defined in section 4.2, the expansion coefficients K of the projected k are given by

$$k_i = \frac{1}{2} \sum_{e\,e'} T_{i\,e\,e'} u_e u_{e'}. \tag{94}$$

4.6.2. Constructing F

Because the ϕ field is approximated as piecewise constant, its degrees of freedom can be interpreted as primal cell integrals. Similarly, the degrees of freedom of the $\mathbf{u}$ field are the integrals of the normal velocity fluxes across primal cell edges, and the D_2 operator looks exactly like a finite volume divergence operator. Thus, it is straightforward to use a finite volume advection scheme for advection of ϕ. The mass flux is constructed using a forward in time advection scheme, identical to that used by Thuburn et al. (2014), using the fluxes U and the mass field Φ as input. We write this symbolically as

$$F = \mathrm{adv}_1(U, \Phi). \tag{95}$$

The subscript 1 indicates that this version of the advection scheme operates on the primal mesh and works with densities or concentrations.

4.6.3. Constructing $Q^{\perp}$

So far we have not needed to use the dual mesh representation of any field. However, in order to use the same finite volume advection scheme as Thuburn et al. (2014) to compute PV fluxes, we need a piecewise constant representation of the PV field on dual cells, and a representation of the mass flux field in terms of components normal to dual cell edges. These are naturally given by the dual function spaces:

$$\pi = (f + \hat{\xi})/\bar{\phi} \qquad \in \mathbf{V}^2, \tag{96}$$

and

$$\widehat{\mathbf{f}^{\perp}} \qquad \in \mathbf{V}^1. \tag{97}$$

Applying (49) followed by (51) to the mass flux gives

$$H\widehat{F^{\perp}} = MF^{\perp} = -WF. \tag{98}$$

19

Now consider the evolution of the dual mass field $\bar{\phi}$.

$$J\dot{\bar{\Phi}} = N\dot{\bar{\Phi}} = \overline{D}_2 WF = -\overline{D}_2 H\widehat{F^\perp} = -J\overline{D}_2\widehat{F^\perp}, \tag{99}$$

i.e.

$$\dot{\bar{\Phi}} + \overline{D}_2\widehat{F^\perp} = 0. \tag{100}$$

Since $\overline{D}_2$ acts exactly like a finite volume divergence operator on the dual mesh, $\bar{\phi}$ behaves exactly as if it were evolving according to a finite volume advection scheme.

Next, in order for PV to evolve in a way consistent with the mass field $\bar{\phi}$, we construct PV fluxes in $\mathbf{V}^1$ using the dual mesh finite volume advection scheme:

$$\widehat{Q^\perp} = \mathrm{adv}_2(\widehat{F^\perp}, \Pi) \tag{101}$$

The subscript 2 indicates that this version of the advection scheme operates on the dual mesh and works with quantities analogous to mixing ratios (such as PV π). Finally, these dual mesh PV fluxes are mapped to the primal mesh for use in the momentum equation:

$$H\widehat{Q^\perp} = MQ^\perp. \tag{102}$$

It may be verified that the resulting vorticity equation for $\dot{\hat{\Xi}}$ is indeed analogous to (100), involving the potential vorticity flux $\widehat{Q^\perp}$. Using (52), (58), (93) and (75), we have

$$
\begin{aligned}
J\dot{\hat{\Xi}} &= & N\dot{\Xi} &= & \overline{D}_2 M\dot{U} &= & -\overline{D}_2 MQ^\perp \\
&= & -\overline{D}_2 H\widehat{Q^\perp} &= & -J\overline{D}_2\widehat{Q^\perp}.
\end{aligned}
$$

Hence,

$$\dot{\hat{\Xi}} + \overline{D}_2\widehat{Q^\perp} = 0, \tag{103}$$

which is of the desired form. The similarity of (103) and (100) means that it is possible to construct PV fluxes from the dual mass fluxes $\widehat{F^\perp}$ such that the evolution of the PV is consistent with the evolution of $\bar{\Phi}$.

4.7. Time integation scheme

The same time integration scheme as in Thuburn et al. (2014) is used.

$$\Phi^{n+1} - \Phi^n + D_2\widetilde{F} = 0, \tag{104}$$

$$MU^{n+1} - MU^n + H\widetilde{\widehat{Q^\perp}} + \overline{D_1 L(\Phi + K)}^t = 0. \tag{105}$$

Here, $\overline{()}^t$ indicates the usual (possibly off-centred) Crank-Nicolson approximation to the integral over one time interval:

$$\overline{\psi}^t = (\alpha\psi^{n+1} + \beta\psi^n)\Delta t \tag{106}$$

(for any field ψ) where $\alpha + \beta = 1$. All results presented below use $\alpha = \beta = 0.5$.

$\widetilde{F}$ is an approximation to the time integral of the mass flux across primal cell edges computed using the advection scheme. The velocity field used for the advection is $\overline{U}^t$. We write this symbolically as

$$\widetilde{F} = \mathrm{Adv}_1(\overline{U}^t, \Phi^n). \tag{107}$$

20

(The notation Adv_1, as distinct from adv_1 in (95), indicates that here we are working with time integrals $\overline{U}^t$ and $\widetilde{F}$.)

Finally, $\widetilde{Q^\perp}$ is an approximation to the time integral of the PV flux across dual edges computed using the advection scheme. Dual grid time integrated mass fluxes are calculated from the primal grid time integrated mass fluxes as

$$H\widetilde{\overline{F^\perp}} = -W\widetilde{F}. \tag{108}$$

These are then used in the dual grid advection scheme to compute the time integrated PV fluxes:

$$\widetilde{Q^\perp} = \mathrm{Adv}_2(\widetilde{\overline{F^\perp}}, \Pi^n). \tag{109}$$

4.8. Incremental iterative solver

The system (104), (105) is nonlinear in the unknowns Φ^{n+1}, U^{n+1}. It can be solved efficiently using an incremental method; this may be viewed as a Newton method with an approximate Jacobian. After l iterations (104) and (105) will not be satisfied exactly but will have some residuals R_Φ, R_U defined by:

$$R_\Phi = \Phi^{(l)} - \Phi^n + D_2\widetilde{F}, \tag{110}$$

$$R_U = MU^{(l)} - MU^n + H\widetilde{Q^\perp} + \overline{\overline{D}_1 L(\Phi + K)}^t. \tag{111}$$

Here $\Phi^{(l)}$ and $U^{(l)}$ are the approximations after l iterations to Φ^{n+1} and U^{n+1} and it is understood that these have been used in evaluating $\widetilde{F}$, $\widetilde{Q^\perp}$, and $\overline{\overline{D}_1 L(\Phi + K)}^t$. We then seek updated values

$$\Phi^{(l+1)} = \Phi^{(l)} + \Phi', \qquad U^{(l+1)} = U^{(l)} + U', \tag{112}$$

that will reduce the residuals, where the increments Φ', U' satisfy

$$\Phi' + \alpha\Delta t D_2\phi^* U' = -R_\Phi, \tag{113}$$

$$U' + \alpha\Delta t \mathcal{M}^{-1}\overline{D}_1 L\Phi' = -\mathcal{M}^{-1}R_U. \tag{114}$$

Here, ϕ^* is a reference value of ϕ; in the current implementation it is given by ϕ^n interpolated to cell edges. To avoid the appearance of the non-sparse matrix M^{-1} in the Helmholtz problem below, a sparse approximation $\mathcal{M}^{-1}$ has been introduced. The construction of $\mathcal{M}^{-1}$ is briefly discussed in Appendix A.

Eliminating U' leaves a Helmholtz problem for Φ':

$$\alpha^2\Delta t^2 D_2\phi^*\mathcal{M}^{-1}\overline{D}_1 L\Phi' - \Phi' = R_\Phi - \alpha\Delta t D_2\phi^*\mathcal{M}^{-1}R_U. \tag{115}$$

In the current implementation, the Helmholtz problem is solved using a single sweep of a full multigrid algorithm. This gives sufficient accuracy to avoid harming the convergence rate of the Newton iteration. Once Φ' is found, U' is obtained by backsubstitution in (114). Finally, (112) is used to obtain improved estimates for the unknowns.

All of the results presented use 4 Newton iterations, and we have not found any case in which more iterations were needed. Monitoring of the residuals R_Φ and R_U indicates that convergence of the Newton iterations is very good, and it might be possible to use fewer iterations in practice. The convergence rate of the Newton iterations is limited by the approximations made to the Jacobian, namely the neglect of the Coriolis and nonlinear terms, and the approximation of the velocity mass matrix inverse; the first two of these approximations are quantified by the dimensionless parameters $\alpha\Delta t f$ and $\alpha\Delta t|\nabla\mathbf{u}|$. The algorithm requires the inversion of several of the linear operators represented as matrices above. Appendix A describes how this is done.

5. Results

The same tests were applied to the finite element shallow water model as were applied to the finite volume model of Thuburn et al. (2014). Only a subset of results are shown here to emphasize the differences between the two models. Other aspects are the following.

- **Stability.** All experimentation to date suggests the two models have the same stability limit: with no temporal off-centring ($\alpha = \beta = 0.5$) the models are stable for large gravity wave Courant numbers and advective Courant numbers less than 1. (For the polygonal meshes used, a natural definition of the gravity wave Courant number is $\phi^{1/2} \Delta t / d$, where d is the distance between neighbouring cell centers, while the natural definition of the advective Courant number is $(\nabla \cdot \mathbf{u}^+) \Delta t$, where the superscript + indicates that the divergence in a cell should be calculated using only the velocities at outflow edges.)

- **Advection.** The same advection scheme is used in the two models to compute mass, PV, and tracer fluxes on primal and dual meshes. In particular, the models share the consistency between mass and PV, between mass and tracers, and between primal mass and dual mass discussed by Thuburn et al. (2014).

- **Balance.** The balance test discussed in section 6.8 of Thuburn et al. (2014) was repeated for the finite element shallow water model. The results on both the hexagonal and cubed sphere meshes were very similar to those for the finite volume model and the ENDGame semi-implicit semi-Lagrangian model (Zerroukat et al., 2009), implying that any spurious numerical generation of imbalance is extremely weak.

- **Computational Rossby modes.** The experiment to test the ability of the scheme on hexagonal meshes to handle grid-scale vorticity features was not repeated here. However, given the general arguments in Thuburn et al. (2014) (see also Weller, 2012), and the similarities between the numerics of the finite volume and finite element models, the results are expected to be very similar for the finite element model.

For the remaining tests discussed below, the same mesh resolutions and time steps were used as in Thuburn et al. (2014).

5.1. Convergence of the Laplacian

The discrete Laplacian defined in section 4.3 was applied to the $\mathbf{V}_2$ representation of the field $\cos \varphi \sin \lambda$ on the unit sphere, where φ is latitude and λ is longitude, and the L_∞ and L_2 errors computed on different resolution meshes. The results are shown in table 1.

On both the hexagonal and cubed sphere meshes the L_∞ errors converge at first order. On the hexagonal mesh the L_2 errors converge at close to second order, while on the cubed sphere mesh the convergence rate is between first and second order. For the cubed sphere mesh the convergence of the discrete scalar Laplacian is significantly better than for the finite volume scheme of Thuburn et al. (2014) (their table 4).

5.2. Convergence of the Coriolis operator

The convergence of the Coriolis operator was investigated as follows. A stream function equal to $\cos \varphi \sin \lambda$ was sampled at dual vertices ($\hat{\Psi}$), enabling exact dual edge normal fluxes $\widehat{U^\perp} = -\overline{D}_1 \hat{\Psi}$ to be computed. The same stream function was also sampled at primal vertices (Ψ),

Table 1: Convergence of the scalar Laplacian on hexagonal and cubed sphere grids.

Hex Ncells	L_∞ err	L_2 err	Cube Ncells	L_∞ err	L_2 err
42	0.14	0.074	54	0.12	0.064
162	0.033	0.019	216	0.030	0.016
642	0.0090	0.0049	864	0.0077	0.0043
2562	0.0026	0.0012	3456	0.0038	0.0012
10242	0.00082	0.00031	13824	0.0022	0.00037
40962	0.00036	0.000081	55296	0.0012	0.00012
163842	0.00018	0.000022	221184	0.00062	0.000039

Table 2: Convergence of the Coriolis operator on hexagonal and cubed sphere grids.

Hex Ncells	L_∞ err	L_2 err	Cube Ncells	L_∞ err	L_2 err
42	0.018	0.0092	54	0.0079	0.0049
162	0.0049	0.0026	216	0.0055	0.0021
642	0.0018	0.00066	864	0.0039	0.00092
2562	0.00078	0.00017	3456	0.0022	0.00037
10242	0.00036	0.000042	13824	0.0012	0.00014
40962	0.00017	0.000011	55296	0.00060	0.000050

enabling exact primal edge normal fluxes $U = -D_1\Psi$ to be calculated; approximate dual edge normal fluxes are then given by the Coriolis operator: $H\widehat{U^\perp}_{\text{approx}} = MU^\perp_{\text{approx}} = -WU = WD_1\Psi$. The difference between the two estimates $WD_1\Psi + H\overline{D}_1\hat{\Psi}$ gives a measure of the error in the Coriolis operator.

Values of the error at different resolutions on the two meshes are shown in table 2. On both meshes the L_∞ errors converge at first order. The L_2 errors converge at second order on the hexagonal mesh and between first and second order on the cubed sphere mesh. This consistency of the Coriolis operator, in contrast to the finite volume scheme of Thuburn et al. (2014), was one of the primary motivations for investigating the finite element approach.

5.3. Balanced solid body rotation

Test case 2 of Williamson et al. (1992) tests the ability of models to represent large-scale steady balanced flow. The exact solution is known, allowing errors in ϕ and u to be computed. The errors on the two meshes after 5 days are given in table 3, along with the time steps used at different resolutions.

On the hexagonal mesh the convergence rate is close to second order or better. On the cubed sphere mesh it is between first and second order for $L_\infty(u)$ and close to second order for the other error measures. The errors are considerably smaller than for the finite volume scheme of Thuburn et al. (2014) (their table 6).

Figure 6 shows the pattern of geopotential errors after 5 days at the second highest resolution in the table. The errors clearly reflect the mesh structure, showing a zonal wavenumber 5 pattern on the hexagonal mesh and a zonal wavenumber 4 pattern on the cubed sphere mesh. However, in contrast to the finite volume model, which shows errors concentrated along certain features of the mesh, the error pattern here is large scale and almost smooth.

Table 3: Geopotential errors (m²s⁻²) and velocity errors (ms⁻¹) for the solid body rotation test case.

Ncells	Δt (s)	$L_2(\phi)$	$L_\infty(\phi)$	$L_2(u)$	$L_\infty(u)$
Hex					
642	7200	19.62	43.40	0.290	0.774
2562	3600	8.59	14.52	0.0940	0.217
10242	1800	2.27	4.01	0.0244	0.0551
40962	900	0.584	1.13	0.00609	0.0144
Cube					
864	7200	35.04	87.48	0.212	0.569
3456	3600	10.16	18.06	0.0754	0.235
13824	1800	2.57	4.65	0.0194	0.0692
55296	900	0.639	1.17	0.00484	0.0257

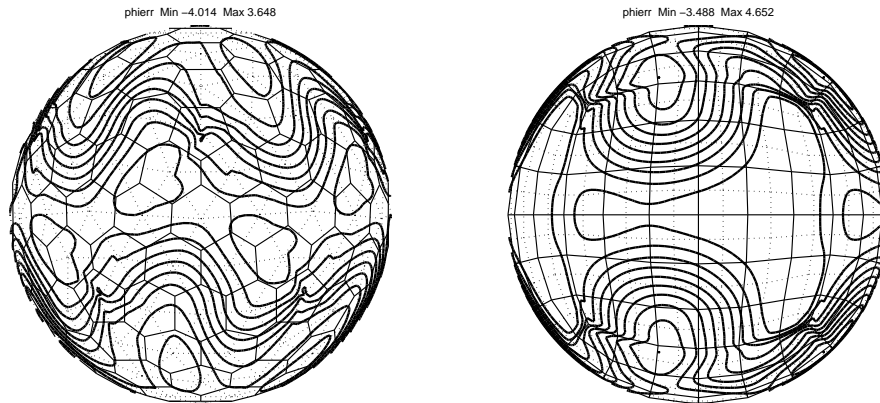phierr Min −4.014 Max 3.648   phierr Min −3.488 Max 4.652

Figure 6: Geopotential error (m²s⁻²) after 5 days for the solid body rotation test case. Left: hexagonal mesh, 10242 cells. Right: cubed sphere mesh, 13824 cells. In each case 11 evenly space contours (i.e. 10 intervals) are used between the minimum and maximum values. (The coarse resolution meshes shown as background are for orientation only.)

Table 4: Height errors (m) for test case 5.

Ncells	Δt (s)	$L_1(h)$	$L_2(h)$	$L_\infty(h)$
Hex				
642	1800	36.37	50.91	191.47
2562	900	11.62	15.83	66.84
10242	450	3.12	4.11	15.06
40962	225	1.27	1.82	9.28
Cube				
864	1800	44.11	64.93	291.35
3456	900	17.57	25.14	100.66
13824	450	3.75	5.25	21.42
55296	225	1.08	1.46	6.47

5.4. *Flow over an isolated mountain*

Test case 5 of Williamson et al. (1992) involves an initial solid body rotation flow impinging on a conical mid-latitude mountain, leading to the generation of gravity and Rossby waves and, eventually, a complex nonlinear flow. There is no analytical solution for this test case, so a high-resolution reference solution was generated using the semi-implict, semi-Lagrangian ENDGame shallow water model (Zerroukat et al., 2009). The finite element model runs stably with the time steps given in table 3, but, as discussed by Thuburn et al. (2014) for the finite volume model and for ENDGame itself, the errors are then dominated by the semi-implicit treatment of the large amplitude gravity waves present in the solution. At any given resolution, the errors look almost identical for all combinations of model and mesh tested. The test was therefore repeated with the time steps reduced by a factor 4. The resulting height errors at day 15 are shown in table 4. The errors on the two meshes are generally very similar, and in most cases are a little smaller than those produced by the finite volume model (Thuburn et al., 2014, table 7).

Maps of height error at day 15 are shown in Fig. 7. The errors produced by the finite element model are of comparable size to those from ENDGame, though the error patterns are different in the three cases. Comparison with figure 6 of Thuburn et al. (2014) confirms that the errors in the finite element model are somewhat smaller than those in the finite volume model.

This test case was also run to 50 days at the highest resolutions in table 4 and several diagnostics relevant to the mimetic properties of the scheme were calculated. The results are very similar to those shown in figure 8 of Thuburn et al. (2014). They confirm that mass is conserved to within roundoff error, and that changes in the total available energy (available potential energy plus kinetic energy) are much smaller than the conversions between available potential energy and kinetic energy. The dissipation of available energy and potential enstrophy is associated almost entirely with the inherent scale-selective dissipation in the advection scheme; it is very small, of order 1 part per thousand, during the first 20 days, but increases subsequently as PV contours begin to wrap up and nonlinear cascades become significant, implying that the inherent dissipation adapts automatically to the flow complexity in a reasonable way. A dual-mass-like tracer remains consistent with the diagnosed dual mass field $\bar{\phi}$ to within 2 parts in 10^4, and a PV-like tracer remains consistent with the diagnosed PV field, to within 3 parts in 10^3. The small errors result from imperfect convergence of various iterative aspects of the solver, and can be reduced by taking more iterations.

h error time 0001296000 Min −15.06 Max 14.4

h error time 0001296000 Min −18.55 Max 21.42
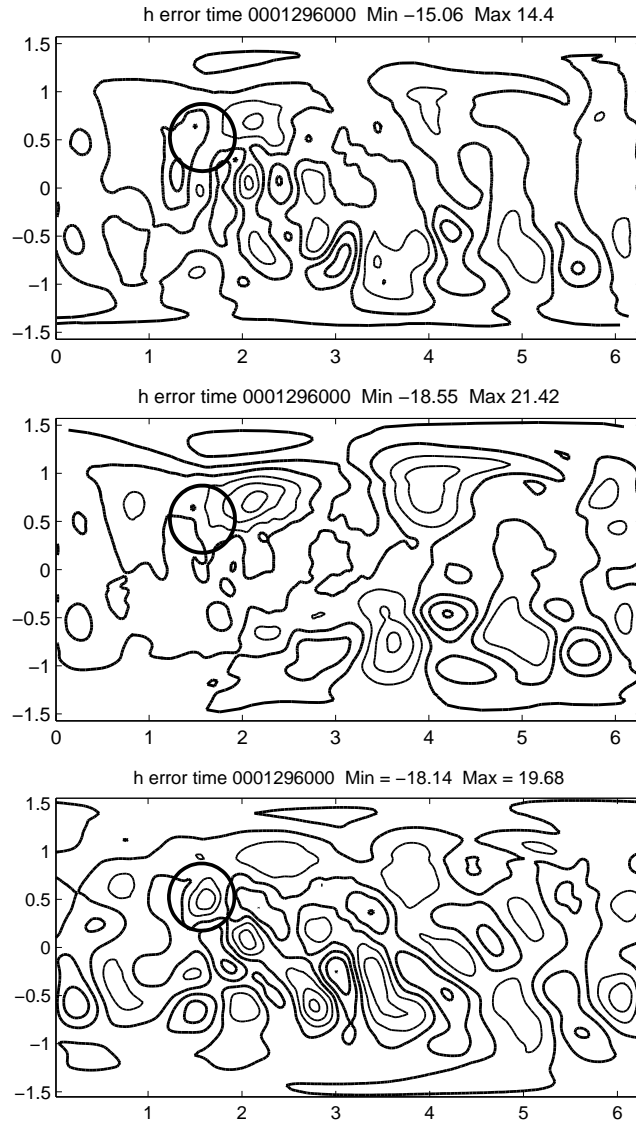
h error time 0001296000 Min = −18.14 Max = 19.68

Figure 7: Height errors (m) at day 15 for the isolated mountain test case. Top: hexagonal mesh, 10242 cells. Middle: cubed sphere mesh, 13824 cells. Bottom: ENDGame on a regular longitude-latitude mesh, 160×80 cells. The contour interval is 6 m, and zero and negative contours are bold. The bold circle indicates the position of the mountain.

5.5. Barotropically unstable jet

The test case proposed by Galewsky et al. (2004) follows the evolution of a perturbed barotropically unstable jet. The case tests the ability of models to handle the complex small scale vorticity features produced by the rapidly growing instability. The results are very sensitive to spurious triggering of the instability by error patterns related to the mesh structure.

Figure 8 shows the relative vorticity field at day 6 for the hexagonal mesh with 10242 cells and 163842 cells, the cubed sphere mesh with 13824 cells and 221184 cells, and, for comparison, from ENDGame on a 640×320 longitude-latitude mesh. In all cases the vorticity field is free of noise and spurious ripples. However, at coarse resolution the finite element model solutions show distinct 'grid imprinting', with a zonal wavenumber 5 pattern on the hexagonal mesh and a zonal wavenumber 4 pattern on the cubed sphere mesh. At finer resolution the solutions are more similar to the ENDGame solution, but still show significant development in the longitude range $\pi/2$ to π where the jet in the ENDGame solution remains quiescent. It should be noted that the rather artificial highly unstable initial state in this test case, while it is discriminating, may exaggerate the importance of grid imprinting errors. More appropriate test cases will be needed to determine whether the degree of grid imprinting is acceptable for practical applications.

The solutions on the hexagonal mesh, especially at the finer resolution, are remarkably similar to those from the finite volume model (Thuburn et al., 2014, figure 9). On the other hand, the solutions on the cubed sphere mesh show some noticable differences from the finite volume model. At the finer resolution, outside the region strongly affected by the spurious development, the structure of vorticity features is slightly more accurate in the finite element model.

6. Conclusions and discussion

A method of constructing low-order mimetic finite element spaces on arbitrary two-dimensional polygonal meshes using compound elements, along with corresponding discrete Hodge star operators for mapping between primal and dual function spaces, has been used as the basis of a numerical model to solve the shallow water equations on a rotating sphere. The model has the same mimetic properties, which underpin the ability to capture important physical properties, as the finite volume model of Thuburn et al. (2014), but with improved accuracy. Encouraged by these results and others, the GungHo project[5] has begun the development of three-dimensional schemes for the compressible Euler equations based on similar ideas.

The finite volume model of Thuburn et al. (2014) relies on certain properties of the mesh for accuracy, namely the Heikes and Randall (1995b) optimization on the hexagonal mesh and the placement of primal vertices relative to dual vertices on the cubed sphere mesh. Although identical meshes have been used here to ensure the cleanest possible comparison, the mimetic finite element scheme does not depend on such mesh properties for accuracy; thus it provides greater flexibility in the choice of mesh.

An alternative mimetic numerical method that has been proposed for atmospheric modelling is the spectral element method (Taylor and Fournier, 2010). It has the advantage of allowing arbitrarily high spatial order of accuracy, but has so far only been developed for quadrilateral and hexahedral meshes, albeit with almost arbitrary structure. Also, the same quadrature nodes are used for all fields, giving rise to wave dispersion issues analogous to those on an unstaggered

---

[5]GungHo: a collaboration between the Met Office and UK research councils to develop a scalable dynamical core for the Met Office weather and climate model.
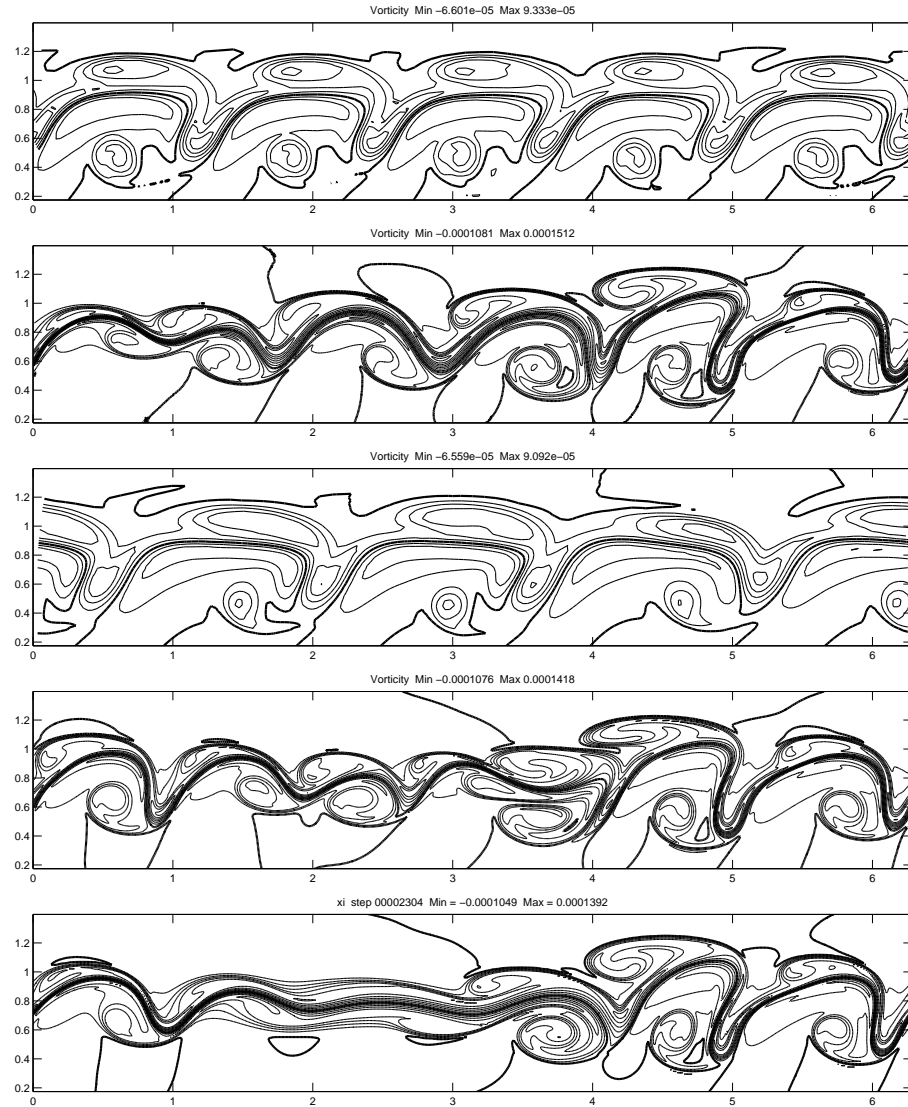
Figure 8: Relative vorticity field at day 6 for the barotropic instability test case. Row 1: hexagonal mesh, 10242 cells, $\Delta t = 900$ s. Row 2: hexagonal mesh, 163842 cells, $\Delta t = 225$ s. Row 3: cubed sphere mesh, 13824 cells, $\Delta t = 900$ s. Row 4: cubed sphere mesh, 221184 cells, $\Delta t = 225$ s. Row 5: ENDGame, 640×320 cells, $\Delta t = 225$ s. The plotted region is 0^o to 360^o longitude, 10^o to 80^o latitude. The contour interval is 2×10^{-5} s^{-1}.

28

A-grid (Melvin et al., 2012). An alternative approach to developing finite element methods on polygonal meshes is provided by the virtual element method (VEM) (Beirão da Veiga et al., 2013, 2014). The compound element method aims to produce a hierarchy of finite element spaces as illustrated in Fig. 2 on arbitrary polygonal meshes by subdividing elements into triangles to produce explicitly-defined finite element spaces. In contrast, the VEM constructs these spaces by harmonic extension directly, but avoids the specific construction of the finite element spaces. Instead, VEM defines the basis functions in terms of their projection onto a polynomial basis, which can be explicitly evaluated. In general VEM bilinear forms must be stabilised by additional terms. In our setting, it is not clear whether these extra terms can be designed to satisfy the steady geostrophic mode property or the PV conservation property that we discuss in this paper.

An important practical consideration is the computational cost of the method. As a rough guide, the cost of the mimetic finite element model on a single processor varied between 3.3 and 4.6 times the cost of ENDGame for the cubed sphere mesh and between 4.2 and 7.3 times the cost of ENDGame for the hexagonal mesh, at the resolutions tested[6]. (For comparison, the cost of the finite volume model varied between 2.7 and 3.7 times the cost of ENDGame for the cubed sphere mesh and between 3.3 and 4.9 times the cost of ENDGame for the hexagonal mesh.) The greater cost on the hexagonal mesh compared to the cubed sphere results from a combination of a greater stencil size for some operators and, in the current implementation, a less cache-friendly mesh numbering (the latter could straightfowardly be optimized). Given the potential to optimize the implementation and the expected gains in parallel scalability from the quasi-uniform mesh, these figures suggest that, despite the need for indirect addressing and the need to invert several linear operators, the finite element method need not be prohibitively expensive compared to methods currently used for operational forecasting, typified by ENDGame.

Although the finite element code has not yet been implemented in parallel, the algorithm presents no barriers to a scalable parallel implementation. The H, J and M operators are diagonally dominant, with a condition number that does not increase with increasing resolution; their inverses can be accurately computed with a small number of iterations, each involving only local communication. Similarly, the condition number of the Helmholtz problem does not increase with increasing resolution, assuming that Δt decreases in proportion to the grid spacing, and the multigrid solver requires only a small number of smoother iterations, each involving only local communication. Finally, the advection scheme and other operators in the model involve only local communication. In contrast, for a semi-implicit semi-Lagrangian scheme on a longitude-latitude grid, such as ENDGame, the Helmholtz problem becomes worse conditioned with increasing resolution, requiring ever more solver iterations, while the communication required for semi-Lagrangian advection becomes increasingly non-local near the poles. These effects will eventually limit the scalability of ENDGame.

A related issue is the complexity of the formulation, both conceptual and computational. Conceptually, the extra complexity compared to a standard triangular or quadrilateral finite element formulation comes from (a) the need to set up and solve a small linear system in order to construct each basis function, and (b) any integrals over elements must be broken down into a sum of integrals over subelements. (Formulating the scheme for an arbitrary unstructured mesh presents no new difficulties beyond those on any triangular or quadrilateral unstructured mesh; indeed it can often simply the code.) In the current implementation, all the operators L, M, H, J, R, W, and T are precomputed, thus avoiding the need for any run-time quadrature in the finite

---

[6]Martin Schreiber (pers. comm.) reports that the cost of the finite element model can be significantly reduced, by roughly a factor 2, by reordering the dimensions of a couple of key arrays to improve cache usage.

element parts of the calculations[7]. (Also, once these operators are computed, there is no need to retain the details of how the compound elements were built from subelements.) This precomputation is possible because all but one of these operators are linear; the only nonlinear term (other than advection) is a simple quadratic nonlinearity in the kinetic energy. In a system with more complex nonlinearities, such as the pressure gradient term in a compressible three-dimensional fluid, precomputation might not be possible and some run time quadrature would be unavoidable. Even so, in a high performance computing environment it is not clear whether precomputation or run-time quadrature would be most efficient, given the relative cost of memory access and computation (David Ham, pers. comm.).

The mathematical similarity of the finite element and finite volume formulations has been emphasized, the principal difference being the appearance of mass matrices in the finite element formulation. The similarity is made even clearer if we use (50), (51) and (79) to rewrite (93) in the equivalent dual space form

$$\dot{\hat{U}} + \widehat{Q^{\perp}} + \overline{D}_1(\hat{\Phi} + \hat{K}) = 0. \tag{116}$$

The velocity degrees of freedom then correspond to dual edge circulations, and $\hat{U}$ can be identified with the V of Thuburn et al. (2014). Equations (92) and (116) explicitly involve only the topological derivative operators D_2 and $\overline{D}_1$; the metric enters through the Hodge star operators needed to map between primal and dual function spaces. This approach of isolating the metric from the purely topological operators in order to construct numerical methods with mimetic properties on complex geometries or meshes has been advocated by several authors (e.g. Bossavit, 1998; Hiptmair, 2001; Palha et al., 2014, and references therein).

It is also worth emphasizing that the roles of primal and dual function spaces are not symmetrical here. Although any given field may be represented in both the primal and dual spaces, with a reversible Hodge star map between them, only primal space test functions are ever used, and so only primal space mass matrices appear, and dual space weak derivatives are never needed. (An interesting alternative would be to use only dual space test functions; then the prognostic equations remain (92) and (116), but (50)-(52) are replaced by

$$I^T \Phi = \hat{L}\hat{\Phi}, \tag{117}$$

$$H^T U = \hat{M}\hat{U}, \tag{118}$$

$$J^T \Xi = \hat{N}\hat{\Xi}, \tag{119}$$

where $\hat{L}$, $\hat{M}$ and $\hat{N}$ are the mass matrices for the spaces $\mathbf{V}^0$, $\mathbf{V}^1$ and $\mathbf{V}^2$, respectively.)

Only the lowest order polygonal finite element spaces are used here: compound P1-RT0-P0$^{\mathrm{DG}}$. Christiansen (2010) has shown that the idea of discrete harmonic extension also applies to the construction of higher order compatible finite element spaces. However, it is not known whether higher order dual finite element spaces can be constructed.

Another, more straightforward, extension of the compound element approach is to three dimensions. The compound elements used here can be extruded into polygonal prisms; we have made some initial progress in working out the details of using such a scheme for the compressible Euler equations. (In atmosphere and ocean models it is desirable, for several reasons, to use a columnar mesh.) Fully three-dimensional compound elements can also be constructed using

---

[7]Some run-time quadrature is done in the advection scheme to compute swept area integrals.

Table A.5: Relaxation parameters used for Jacobi iteration for operator inverses.

Grid	J^{-1}	M^{-1}	H^{-1}
Hex	1.4	1.4	1.4
Cube	1.4	0.9	1.4

the discrete harmonic extension approach. These might be useful, for example, to implement a finite element version of the cut cell method for handling bottom topography (e.g. Lock et al., 2012, and references therein) while retaining a columnar mesh.

Besides their ability to use arbitrary polygonal meshes, another potentially useful property of the compound elements used here is that the function spaces are built directly in physical space, without the need for Piola transforms. Thus, for example, globally constant functions are always contained in $\mathbf{V}_2$. In this way, the compound elements avoid the reduced convergence rate, and even loss of consistency, discussed by Arnold et al. (2014), and so provide an alternative to the *rehabilitation* technique of Bochev and Ridzal (2008).

Appendix A. Operator inverses and sparse approximate M inverse

Inverses of the H and J operators are needed at the beginning of every time step, and inverses of H and M are needed at every Newton iteration. These are computed by (under- or over-relaxed) Jacobi iteration based on a diagonal approximation of the relevant operator. E.g., to solve $Ax = R$, define

$$x^{(1)} = (A^*)^{-1}R \tag{A.1}$$

where A^* is a diagonal approximation to A, then iterate:

$$x^{(l+1)} = x^{(l)} + \mu(A^*)^{-1}(R - Ax^{(l)}). \tag{A.2}$$

A diagonal approximation J^* to the operator J is defined by demanding that, for every dual cell j, J^* and J should give the same result in dual cell j when acting on the $\mathbf{V}^2$ representation of a constant scalar field. A diagonal approximation M^* to the velocity mass matrix M is defined by demanding that, for every edge e, M^* and M should give the same result at edge e when acting on the $\mathbf{V}_1$ representation of a solid body rotation velocity field whose maximum velocity is normal to primal edge e. A diagonal approximation H^* to the operator H is defined by demanding that, for every edge e, H^* and H should give the same result at edge e when acting on the $\mathbf{V}^1$ representation of a solid body rotation velocity field whose maximum velocity is tangential to dual edge e.

Optimal values of the relaxation parameter μ were found to depend on the operator and mesh structure. The values used are given in table A.5. For the inverses that occur once per time step,

Table B.6: Summary of Notation.

	Term	Meaning	Discussed in
Continuous equations	ϕ	Geopotential	Equations (1), (2)
	$\mathbf{u}$	Velocity	
	$\mathbf{f}$	Mass flux	
	k	Kinetic energy per unit mass	
	ξ	Relative vorticity	
	ζ	Absolute vorticity	
	π	Potential vorticity	
	$\mathbf{q}$	Potential vorticity flux	
Function spaces	V_0, V_1, V_2	Primal function spaces	Section 3
	V^0, V^1, V^2	Dual function spaces	
Basis functions	$\alpha_i, \mathbf{v}_e, \gamma_j$	for spaces V_0, V_1, V_2	Equations (6)-(11)
	$\beta_j, \mathbf{w}_e, \chi_i$	for spaces V^0, V^1, V^2	
Operators in matrix form	$D_2, D_1, \overline{D}_2, \overline{D}_1$	Derivatives	Section 4.1
	L, M, N	Mass matrices	Section 4.2
	R, W	Averaging operators	Section 4.2
	H, J, I	Hodge star operators	Section 4.2
	T	Kinetic energy operator	Section 4.2
	$\mathcal{M}^{-1}$	Sparse approximation to M^{-1}	Section 4.8

10 Jacobi iterations are used. For those that occur at every Newton iteration, 2 Jacobi iterations are used taking the solution obtained at the previous Newton iteration as the first guess (or (A.1) on the first Newton iteration).

A sparse approximate inverse $\mathcal{M}^{-1}$ of the $\mathbf{V}_1$ mass matrix is needed for the Helmholtz problem. On the hexagonal mesh it is sufficient to use a diagonal approximation

$$\mathcal{M}^{-1} = (M^*)^{-1}. \tag{A.3}$$

However, on the cubed sphere mesh, whose dual and primal edges are not mutually orthogonal, such a diagonal approximation is less accurate and limits the convergence of the Newton iterations. Therefore we use instead an approximation based on a single Jacobi iteration towards the inverse of M:

$$\mathcal{M}^{-1} = (M^*)^{-1} \left\{ (1 + \mu)\text{Id} - \mu M (M^*)^{-1} \right\}, \tag{A.4}$$

where Id is the identity matrix. This approximate inverse is not diagonal but has the same stencil as M itself.

Appendix B. Notation

For convenience, the notation used in the paper is summarized here in table B.6.

Any field represented in a primal function space V_0, V_1, or V_2 has another representation in the corresponding dual function space V^2, V^1, or V^0, respectively. A ˆ above the variable name is used to indicate the dual space representation. For a field represented in V_2, an averaged

version of that field may be represented in V_0 or in V^2; these averaged fields are indicated by $\overline{}$ and $\widetilde{}$ above the variable name, respectively. See section 3.

References

Arnold, D. N., Boffi, D., and Bonizzoni, F.: Finite element differential forms on curvilinear cubic meshes and their approximation properties, Numer. Math., DOI:10.1007/s00211-014-0631-3, 2014.

Beirão da Veiga, L., Brezzi, F., Cangiani, A., Manzini, G., Marini, L. D., and Russo, A.: Basic principles of virtual element methods, Math Models and Methods in Appl. Sci., 23, 199–214, 2013.

Beirão da Veiga, L., Brezzi, F., Marini, L. D., and Russo, A.: Hitchhiker's guide to the virtual element method, Math Models and Methods in Appl. Sci., 24, 1541–1573, 2014.

Bochev, P. B., and Ridzal, D.: Rehabilitation of the lowest-order Raviart-Thomas element on quadrilateral grids, SIAM J. Numer. Anal., 47, 487–507, 2008.

Bossavit, A.: Computational Electromagnetism: Variational Formulation, Complementarity, Edge Elements. Academic Press Electromagnetism Series, Number 2 San Diego: Academic Press, 1998.

Buffa, A., and Christiansen, S. H.: A dual finite element complex on the barycentric refinement, Math. of Comput., 76, 1743–1769, 2007.

Christiansen, S. H.: A construction of spaces of compatible differential forms on cellular complexes, Math. Models and Methods in Appl. Sci., 18, 739–757, 2008.

Christiansen, S. H.: Minimal mixed finite elements on polyhedra, C. R. Acad. Sci. Paris, 348, 217–221, 2010.

Cotter, C. J., and Shipton, J.: Mixed finite elements for numerical weather prediction, J. Comput. Phys., 231, 7076–7091, 2012.

Cotter, C. J. and Thuburn, J.: A finite element exterior calculus framework for the rotating shallow-water equations, J. Comp. Phys., 257, 1506–1526, 2014.

Danilov, S.: On the utility of triangular C-grid type discretization for numerical modeling of large-scale ocean flows, Ocean Dyn., 60, 1361–1369, 2010.

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

Heikes, R., and Randall, D.: Numerical integration of the shallow-water equations on a twisted icosahedral grid, Part II: A detailed description of the grid and analysis of numerical accuracy, Mon. Weather Rev., 123, 1881–1997, 1995b.

Hiptmair, R.: Discrete Hodge operators, Numer. Math., 90, 265–289, 2001.

Le Roux, D. Y., Rostand, V., and Pouliot, B.: Analysis of numerically induced oscillations in 2D finite-element shallow-water models. Part I: Inertia-gravity waves, SIAM J. Sci. Comp., 29, 331–360, 2007.

Lock, S.-J., Bitzer, H.-W., Coals, A., Gadian, A., and Mobbs, S.: Demonstration of a cut-cell representation of 3D orography for studies of atmospheric flows over very steep hills, Mon. Weather Rev., 140, 411–424.

McRae, A. T. T., and Cotter, C.: Energy- and enstrophy-conserving schemes for the shallow-water equations, based on mimetic finite elements, Submitted to Q. J. R. Meteorol. Soc. (available at http://dx.doi.org/10.1002/qj.2291).

Melvin, T., Staniforth, A., and Cotter, C.: A two-dimensional mixed finite-element pair on rectangles, Quart. J. Roy. Meteorol. Soc., 140, 930–942, 2014.

Melvin, T., Staniforth, A., and Thuburn, J.: Dispersion analysis of the spectral element method, Quart. J. Roy. Meteorol. Soc., 138, 1934–1947, 2012.

Melvin, T., and Thuburn, J.: Wave dispersion properties of compound finite elements, submitted to J. Comput. Phys.

Nédélec, J.-C.: Mixed finite elements in $\mathbb{R}^3$, Numer. Math., 35, 315–341, 1980.

Palha, A., Rebelo, P. P., Hiemstra, R., Kreeft, J., and Gerritsma, M.: Physics-compatible discretization techniques on single and dual grids, with application to the Poisson equation of volume forms. J. Comput. Phys., 257, 1394–1422, 2014.

Ringler, T. D., Thuburn, J., Klemp, J. B., and Skamarock, W. C.: A unified approach to energy conservation and potential vorticity dynamics for arbitrarily structured C-grids, J. Comput. Phys., 229, 3065–3090, 2010.

Staniforth, A., and Thuburn, J.: Horizontal grids for global weather and climate prediction models: a review, Q. J. Roy. Meteor. Soc., 138, 1–26, 2012.

Taylor, M. A., and Fournier, A,: A compatible and conservative spectral element method on unstructured grids, J. Comput. Phys, 229, 5879–5895, 2010.

Thuburn, J., and Cotter, C. J.: A framework for mimetic discretization of the rotating shallow water equations on arbitrary polygonal grids, SIAM J. Sci. Comput., 34, 203–225, 2012.

Thuburn J., Cotter, C. J., and Dubos, T.: A mimetic, semi-implicit, forward-in-time, finite volume shallow water model: comparison of hexagonal–icosahedral and cubed sphere grids, Geosci. Model Dev., 7, 909–929, 2014.

Thuburn, J., Ringler, T. D., Skamarock, W. C., and Klemp, J. B.: Numerical representation of geostrophic modes on arbitrarily structured C-grids, J. Comput. Phys., 228, 8321–8335, 2009.

Ullrich, P.A.: Understanding the treatment of waves in atmospheric models. Part 1: The shortest resolved waves of the 1D linearized shallow-water equations, Quart. J. Roy. Meteorol. Soc., 140, 1426–1440, 2014.

Weller, H.: Controlling the computational modes of the arbitrarily structured C-grid, Mon. Weather Rev., 140, 3220–3234, 2012.

Weller, H.: Non-orthogonal version of the arbitrary polygonal C-grid and a new diamond grid, Geosci. Model Dev., 7, 779–797, 2014.

Williamson, D. L., Drake, J. B., Hack, J. J., Jakob, R., and Swartztrauber, P. N.: A standard test set for numerical approximations to the shallow water equations in spherical geometry, J. Comput. Phys., 102, 211–224, 1992.

Zerroukat, M., Wood, N., Staniforth, A., White, A. A., and Thuburn, J.: An inherently mass-conserving semi-implicit semi-Lagrangian discretization of the shallow water equations on the sphere, Quart. J. Roy. Meteorol. Soc., 135, 1104–1116, 2009.