

GungHo! A new dynamical core for the Unified Model*

Andrew Staniforth, Thomas Melvin and Nigel Wood

Met Office, FitzRoy Road, Exeter

EX1 3PB, United Kingdom

Andrew.Staniforth@metoffice.gov.uk, Thomas.Melvin@metoffice.gov.uk,

Nigel.Wood@metoffice.gov.uk

Abstract

The collaborative GungHo! Project has been established to develop a new dynamical core for the Met Office's Unified Model. The rationale for this project, together with associated scientific issues and progress to date made on them, are reviewed.

1 Introduction

The Unified Model (hereafter UM) is a single model that underpins the Met Office's capabilities for operational weather and environment forecasting, and for climate simulation (Davies *et al.*, 2005). It is run in a wide range of configurations with resolutions spanning a very broad range of length and time scales. Horizontal resolutions currently vary from O(100 m) (for physical process studies), through O(1 km) (for mesoscale forecasts over a limited area) and O(25 km) (for global weather forecasts), to O(100 km) (for global climate simulation). Integration periods for climate simulation can be a century or longer.

At the heart of the UM is its dynamical core, and essential to its performance are the choice and form of the continuous governing equations, and the numerical formulae used to approximate them. Improving the efficiency of the dynamical core whilst maintaining, or even improving, accuracy allows maximisation of the model's resolution for given computing resources. This aspect is crucial to the Met Office strategy of moving to ever-higher-resolution models (Brown *et al.*, 2012).

During the past decade or so, the speed of individual computational processors for supercomputers has plateaued, and increased supercomputer performance has primarily been achieved by using many processors working in parallel - see e.g. www.top500.org. Thus an increasingly important aspect of the UM is how its performance scales when run with ever more computational processors. Due to horizontal connectivity, the dynamical core is a key element for this. Currently, the UM's dynamical core uses a lat-lon grid, as do many other atmospheric models. However, it is recognised that there are communication issues around the two poles for models with a lat-lon grid, and this has motivated renewed interest, at the Met Office and elsewhere, in developing models on quasi-uniform spherical grids.

This paper is organised as follows. The choice of a suitable set of governing equations is discussed in section 2. Scalability of the UM's New Dynamics and ENDGame dynamical cores, together with the motivation for, and birth of, the GungHo!¹ Project, are presented in section 3. This then leads to the review in section 4 of various issues associated with developing a new dynamical core on quasi-uniform spherical grids, and progress made to date on these. Finally, some conclusions are drawn in section 5. Note that due to space limitations, the references herein to the literature are necessarily quite parochial.

*© British Crown copyright, the Met Office

¹Although "gung-ho" also has a pejorative meaning, its original meaning (introduced into American slang during World War II) is, appropriately, "working harmoniously together".

2 Unified Model

2.1 Governing equations

An important step in designing a dynamical core is to identify an appropriate set of governing dynamical equations for subsequent discretisation. For simplicity, only the dry unforced equations of the UM's dynamical core are discussed herein. They are (see [Davies et al., 2005](#), for notation)

$$\frac{D_r u}{Dt} - \frac{uv \tan \phi}{r} - 2\Omega \sin \phi v + \frac{c_p \theta}{r \cos \phi} \frac{\partial \pi}{\partial \lambda} + \underbrace{\frac{uw}{r} + 2\Omega \cos \phi w}_{\text{deep}} = 0, \quad (1)$$

$$\frac{D_r v}{Dt} + \frac{u^2 \tan \phi}{r} + 2\Omega \sin \phi u + \frac{c_p \theta}{r} \frac{\partial \pi}{\partial \phi} + \underbrace{\frac{vw}{r}}_{\text{deep}} = 0, \quad (2)$$

$$\underbrace{\frac{D_r w}{Dt}}_{\text{non-hydro}} + c_p \theta \frac{\partial \pi}{\partial r} + g - \underbrace{\left[\left(\frac{u^2 + v^2}{r} \right) + 2\Omega \cos \phi u \right]}_{\text{deep}} = 0, \quad (3)$$

$$\frac{D_r \rho}{Dt} + \rho \left\{ \frac{1}{r \cos \phi} \left[\frac{\partial u}{\partial \lambda} + \frac{\partial}{\partial \phi} (v \cos \phi) \right] + \frac{1}{r^2} \frac{\partial (r^2 w)}{\partial r} \right\} = 0, \quad (4)$$

$$\frac{D_r \theta}{Dt} = 0, \quad (5)$$

$$\pi^{\frac{1-\kappa}{\kappa}} = \frac{R}{p_0} \rho \theta, \quad (6)$$

where

$$\frac{D_r}{Dt} \equiv \frac{\partial}{\partial t} + \frac{u}{r \cos \phi} \frac{\partial}{\partial \lambda} + \frac{v}{r} \frac{\partial}{\partial \phi} + w \frac{\partial}{\partial r}. \quad (7)$$

Equations (1)-(3) are the three components of the momentum equation, (4)-(6) are the continuity, thermodynamic and state equations, respectively, and (7) defines the material derivative. Collectively they are termed the fully-compressible non-hydrostatic equations for a deep atmosphere, and the spherical geopotential approximation is assumed. The UM's current operational dynamical core (colloquially known as New Dynamics) and its new alternative core (known as ENDGame - see below) are both based on (1)-(7).

2.2 Approximate equation sets

A quartet of dynamically-consistent approximate models of the global atmosphere was examined in [White et al. \(2005\)](#), where dynamical consistency is defined to be formal respect of conservation principles for axial angular momentum, energy and potential vorticity. This quartet can be obtained from (1)-(7) as the following four combinations of deep/shallow with non-hydrostatic/(quasi-)hydrostatic:

- fully-compressible non-hydrostatic deep equations - retain all terms;
- quasi-hydrostatic deep equations - omit the blue “non-hydro” term;
- non-hydrostatic shallow equations - omit the red “deep” terms, and replace r by Earth's radius a in the remaining terms; and

- hydrostatic shallow (or hydrostatic primitive) equations - omit the blue “*non-hydro*” and red “*deep*” terms, and replace r by Earth’s radius a in the remaining terms.

The most general of these four sets is the first (used in the UM), and the most approximated the last (the hydrostatic primitive equations, used in most operational global weather prediction and climate models). However, as computers become ever more powerful, there is an emerging need to develop global non-hydrostatic deep-atmosphere models that can be run at horizontal resolutions beyond the limit of validity (resolved scales larger than about 10 km) of the hydrostatic primitive equations, and for other applications (such as: space weather, with an atmospheric depth as large as 600 km; and other planetary atmospheres).

Furthermore, as models become more accurate, due to higher resolution with improved numerics and physical parametrisations, it can be argued that even the fully-compressible non-hydrostatic deep-atmosphere equations may be insufficiently accurate, due to the assumed spherical geopotential approximation, and particularly so for planetary atmospheres (e.g. [Mayne et al., 2013](#)). With this in mind, [White et al. \(2008\)](#) examined relaxing this spherical approximation to a spheroidal one, and proposed a spheroidal coordinate system for the fully-compressible non-hydrostatic deep-atmosphere equations. An analogous quartet to the spherical one discussed in [White et al. \(2005\)](#), but with relaxation of the spherical geopotential assumption, has recently been derived and discussed in [White and Wood \(2012\)](#) and [Staniforth \(2014\)](#).

2.3 ENDGame

A new dynamical core ([Wood et al., 2013](#)) has recently been developed at the Met Office, known as ENDGame (**E**ven **N**ewer **D**ynamics for **G**eneral **a**tmospheric **m**odelling of the **e**nvironment). It has many similarities with the current (New Dynamics) dynamical core, but also a number of important differences. Some of these are now summarised.

Both cores are based on the fully-compressible deep-atmosphere equations. However ENDGame is designed to not only be optionally re-configured to use any of the other three members of the above-discussed quartet of approximate equation sets in spherical geometry, but also in ellipsoidal geometry (the New Dynamics formulation is limited to the fully-compressible deep-atmosphere equations in spherical geometry). This opens the way to performing controlled experiments with ENDGame to assess the importance, as a function of application, of the various assumptions (non-hydrostatic vs. (quasi-) hydrostatic, deep vs. shallow, spherical vs. spheroidal) of simpler equation sets.

Both cores use an off-centred semi-implicit semi-Lagrangian time scheme, an Arakawa C grid in the horizontal, and a Charney-Phillips grid in the vertical. However the (off-centred) semi-implicit semi-Lagrangian scheme in ENDGame is both more stable and more accurate than that in New Dynamics, with a much-reduced need for temporal off-centring. Furthermore, although both cores use an Arakawa C grid in the horizontal, ENDGame uses a different placement of variables with respect to the pole than New Dynamics: this influences the derived Helmholtz problem and contributes to its better scaling on parallel-processing computers. ENDGame also has a more accurate discretisation of the Coriolis terms, leading to improved numerical dispersion and energy properties ([Thuburn and Staniforth, 2004](#)).

2.4 Some history

The development timeframe for Met Office dynamical cores is summarised in Table 1 (Terry Davies, personal communication). It can be seen that the development time for a new dynamical core is typically a decade or more. Part of the reason for this is the ever-increasing complexity of the suite of applications.

Year	Equation Set	# of levels	Δx (km)	Notes
1959	Quasi-geostrophic	2/3	320/300	
1972	Shallow hydrostatic	10	300	
1982	Shallow hydrostatic	15	150	Global
1991	Deep quasi-hydrostatic (Unified Model)	20	90	1st global deep model
2002	Deep non-hydrostatic ("New Dynamics")	38/50/70/85	60/40/25	1st global deep non-hydrostatic model
2014	Deep non-hydrostatic ("ENDGame")	70/85	17	Switchable equation sets and geometries

Table 1: Development timeframes for Met Office dynamical cores.

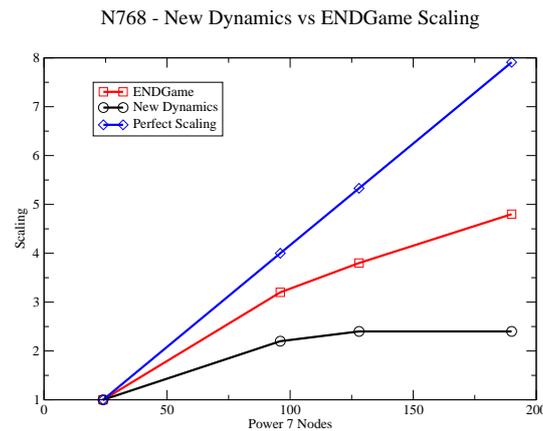


Figure 1: Scaling of the UM (including physical parametrisations) using New Dynamics and ENDGame dynamical cores, as a function of number of nodes (from 24 to 192 with 32 processors/ node) on an IBM Power 7 computer at fixed N768 resolution (approximately 17 km meshlength at mid-latitudes).

3 All change and GungHo!

3.1 Scalability

The (strong) scaling of the UM using the New Dynamics and ENDGame dynamical cores is displayed in Fig. 1 as a function of the number of nodes (each node has 32 processors) on an IBM Power 7 computer at fixed N768 resolution (approximately 17 km meshlength at mid-latitudes, and 70 vertical levels). It is seen that the current (New Dynamics) dynamical core scales poorly (there is virtually no performance gain when doubling the number of nodes from 96 to 192), whereas the new (ENDGame) core scales much better, albeit still significantly less well than perfect scaling. Although further performance gains can be expected from optimisation of ENDGame, it is thought that significant issues will nevertheless remain with further increases in the number of processors.

3.2 Global grids

As noted in section 2.3, the New Dynamics and ENDGame dynamical cores both use semi-implicit semi-Lagrangian integration schemes on a traditional lat-lon grid - see Fig. 2 (a) for a low-resolution example of a lat-lon grid. These schemes then require significant data communication among the grid points clustered around the two poles, but communication between processors is generally slow, leaving processors waiting for data. This then adversely affects scalability, so that an increased number of processors does

not lead to a commensurate reduction in computing time. This problem is exacerbated when horizontal resolution is increased, due to the resolution clustering around the poles. For example, at 25-km nominal (mid-latitude) resolution, the grid spacing for the nearest latitude to the poles is 75 m, but at 10-km resolution, it is only 12 m: this spacing reduces quadratically, rather than linearly, as a function of increased nominal resolution. As discussed in Staniforth and Thuburn (2012), this has motivated renewed interest, worldwide, in investigating alternative, quasi-uniform spherical grids for dynamical cores - see Fig. 2 (b)-(d) for examples of three such global grids.

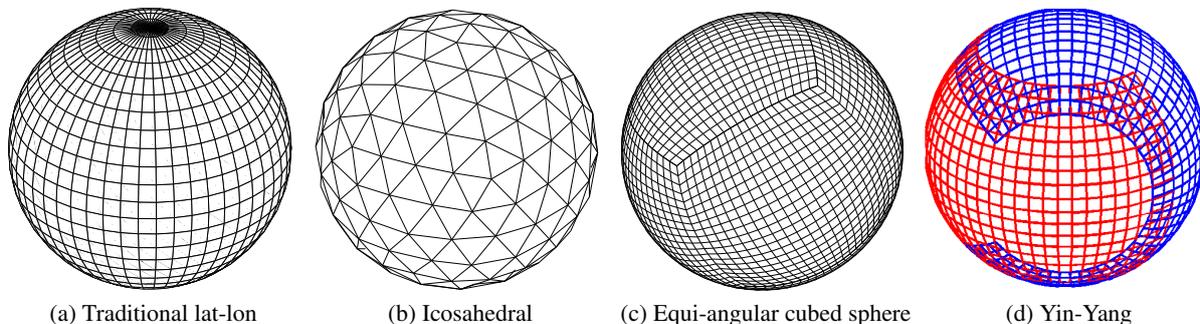


Figure 2: Four kinds of global grid, displayed at low resolution.

3.3 GungHo! - the miracle of birth

Following on from the above discussion, work is well underway to investigate future options to enable the UM to continue to make efficient use of emerging computer architectures. The principal mechanism for this is a significant, funded, collaboration between the Met Office and researchers within both the Natural Environment Research Council and the Science & Technology Facilities Council. This dynamical core project is informally known as the GungHo! (**G**lobally **U**niform, **N**ext **G**eneration, **H**ighly **O**ptimised) project. It started in early 2011 and is funded until early 2016. The focus of the collaboration is the development and evaluation of the options for a new dynamical core that maintains the accuracy of the UM's ENDGame dynamical core, whilst improving its scalability for use on future supercomputer architectures. Some key issues, and recent progress on these, are discussed below. The reader is also referred to the associated papers in this volume by Gung-Ho! collaborators Colin Cotter, Sarah-Jane Lock and John Thuburn.

4 GungHo! to not so gungho

4.1 Top ten properties

In their review paper, Staniforth and Thuburn (2012) identified the following top ten “*essential and desirable properties of a dynamical core*”:

1. Mass conservation
2. Accurate representation of balanced flow and adjustment
3. Computational modes should be absent or well controlled
4. Geopotential and pressure gradients should produce no unphysical source of vorticity
 $\Rightarrow \nabla \times \nabla p = \nabla \times \nabla \Phi = 0$

5. Terms involving the pressure should be energy conserving $\Rightarrow \mathbf{u} \cdot \nabla p + p \nabla \cdot \mathbf{u} = \nabla \cdot (p\mathbf{u})$
6. Coriolis terms should be energy conserving $\Rightarrow \mathbf{u} \cdot (\boldsymbol{\Omega} \times \mathbf{u}) = 0$
7. There should be no spurious fast propagation of Rossby modes; geostrophic balance should not spontaneously break down
8. Axial angular momentum should be conserved
9. Accuracy approaching second order
10. Minimal grid imprinting

One of the reasons for the sustained popularity of the lat–lon grid is that its logically rectangular structure, orthogonality and symmetry can be exploited to obtain the first eight properties. Obtaining these properties on alternative grid structures is a challenging current area of research, and to achieve property 8 on a grid other than a lat–lon grid is particularly challenging. Properties 9–10 are particularly challenging for grids with special points or regions, and likely to require higher-order schemes.

4.2 Balancing degrees of freedom and spurious modes

For good representations of geostrophic adjustment and Rossby-mode propagation, the number of degrees of freedom between momentum and pressure for 2D horizontal flow needs to be in the ratio 2:1, and failure to respect this necessary (but not sufficient) constraint leads to spurious discrete (computational) modes (Le Roux *et al.*, 1998; Thuburn *et al.*, 2009; Weller *et al.*, 2012). Spurious modes are potentially very serious, since they may be excited in realistic applications by boundary conditions, nonlinearity, physical forcing, and data assimilation - see John Thuburn’s paper in the present volume. The geometric properties of polyhedra provide valuable insight for this issue (Staniforth and Thuburn, 2012).

For example, assuming the popular C-grid placement of variables, with pressure at the centres of the faces and the normal component of momentum at the midpoints of the edges, the required balance is obtained when there are twice as many edges (E) as faces (F), i.e. $E = 2F$. The cube is the only regular polyhedron that respects this balance: every quadrilateral-based grid on the sphere (e.g. cubed sphere - Fig. 2(c)) also does so. For triangular regular polyhedra (tetrahedron, octahedron and icosahedron), as well as other grids built entirely from triangles (e.g. icosahedral - Fig. 2(b)), $E = 3F/2$: there are then too many pressure nodes, and two sets of spurious gravity modes. For a pentagonal-hexagonal grid composed primarily of hexagons, $E \approx 3F$: there are then too many momentum nodes, and a spurious set of Rossby modes.

Motivated by these observations, Cotter and Shipton (2012) have recently proposed the use of some particular families of finite elements as a means of improving the horizontal discretisation of atmospheric dynamical cores on quasi-uniform global grids. The finite-element framework: (a) intrinsically facilitates the derivation of discretisations having good conservation properties, which are a direct consequence of an appropriate choice of finite-element spaces; and (b) straightforwardly leads to discretisations that are valid with irregular resolution. Two families of mixed finite elements are specifically recommended: the Raviart–Thomas (RT) element on quadrilaterals, and the Brezzi–Douglas–Fortin–Marini (BDFM) element on triangles.

The (lowest-degree) RT0 element pair on quadrilaterals is the finite-element analogue of the C grid: it, and the (next degree) RT1 pair, are both candidates for discretisation on a cubed sphere, with the RT1 pair offering the promise of improved accuracy. The BDFM element on triangles offers promise for discretisations on triangularly subdivided icosahedral grids. However, the lowest possible degree for triangular elements that allows a 2:1 ratio is the BDFM1 element pair (piecewise linear and discontinuous

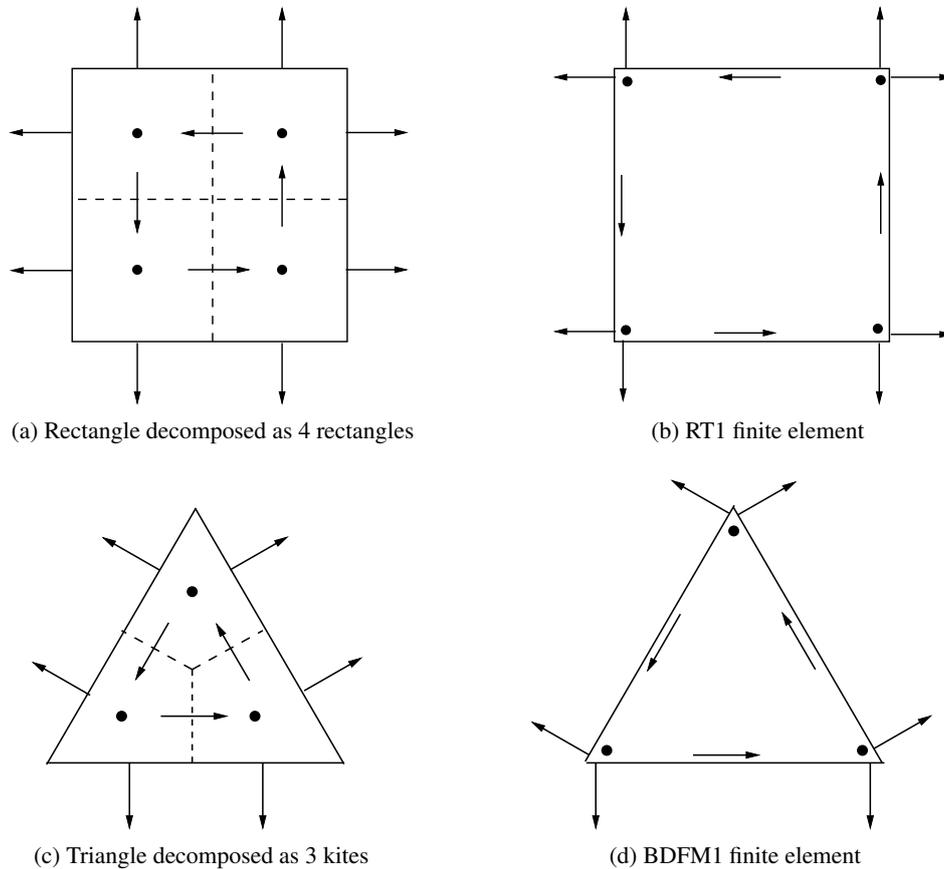


Figure 3: Placement of variables for various rectangular (panels a and b) and triangular (panels c and d) finite-element pairs. Pressure is held at circle points; and arrows denote components of momentum, with normal components continuous across edges, and tangential components discontinuous.

for pressure, piecewise quadratic with continuous normal components for velocity, with the constraint that normal components are linear on edges).

It is instructive to schematically examine the construction of these element pairs. For the RT0 element pair, consider any one of the four rectangles of Fig. 3(a). As for an Arakawa C grid, pressure is held at the rectangle centre (circle point), normal components of momentum at the midpoints of rectangle edges with cross-edge continuity, and the 2:1 ratio of degrees of freedom is exactly satisfied. Now consider a redistribution of the placement of variables of Fig. 3(a), such that: the four pressure points migrate to the corners of the rectangle encompassing the four displayed rectangles; the normal components of momentum across dotted edges migrate to the solid edges of the encompassing rectangle; and the normal components across the solid edges migrate to the corner points of the encompassing rectangle. This then gives Fig. 3(b), which is the placement of variables for the RT1 element pair. The 2:1 ratio is exactly satisfied since no degrees of freedom have been added or subtracted, only redistributed. The final step in constructing the element pair is to choose low-degree piecewise-defined polynomial basis functions over the encompassing rectangle with the appropriate number of degrees of freedom to match the placement of variables, and with the appropriate cross-edge continuity conditions. So, for example, whereas pressure for the RT0 element pair is piecewise constant, and discontinuous between elements, it is piecewise bilinear (i.e. a degree higher), and still discontinuous across elements for the RT1 pair.

A similar construction leads to the BDFM1 element pair for triangles, but with an added wrinkle. In contradistinction to quadrilaterals (see above), a triangular C-grid placement of variables does not satisfy the 2:1 ratio constraint. To address this, each triangle is first sub-divided into three quadrilaterals (kites)

with a C-grid placement of variables - see Fig. 3(c) - thereby generating a quadrilateral grid that does satisfy the 2:1 ratio constraint. The interior degrees of freedom are then migrated to the edges of the encompassing triangle in an analogous manner to that described above: this then gives Fig. 3(d). Low-degree piecewise-defined polynomial basis functions, with the appropriate number of degrees of freedom, are then chosen to match the placement of variables. See [Cotter and Shipton \(2012\)](#) for details.

4.3 Mimetics

Work (driven primarily by Colin Cotter and John Thuburn) over the last year or so has shown that most of [Staniforth and Thuburn \(2012\)](#)'s top-ten properties can be achieved, at least within the framework of the two-dimensional shallow-water equations, with the following approach.

First, the vector-invariant form of the equations (in which the self-advection term, $\mathbf{u} \cdot \nabla \mathbf{u}$, is replaced by $(\nabla \times \mathbf{u}) \times \mathbf{u} + \nabla(\mathbf{u} \cdot \mathbf{u}/2)$) is used. The advantage of this approach is that it makes explicit the appearance of a vorticity flux term: this permits materialisation of a vorticity, which is appropriately advected (the vorticity flux in the velocity equations is calculated by applying an advection scheme to the vorticity).

Second, and as discussed above, a finite-element method is used. ([Thuburn and Cotter \(2012\)](#) proposed a similar approach in a finite-volume context, but it appears difficult to obtain high-order accuracy, or even consistency, on irregular meshes.) The finite-element methodology replaces the more-or-less *ad hoc* finite-difference approach of arbitrarily mixing piecewise-constant or piecewise-linear (or higher degree) reconstructions (often of the same variable at different places in the same discretisation!) by a rigorous sub-element description of each variable. However, it is insufficient to apply a standard finite-element approach. Instead, a *mixed finite-element* method is used (originally developed in the field of electrodynamics) in which each (vector or scalar) variable is assigned to a very specific hierarchy of function spaces in order to give the required mimetic properties (in which the discrete equations *mimic* certain properties of the continuous ones). See [Cotter and Thuburn \(2013\)](#) for details of this approach and its link to exterior calculus and discrete differential forms; and also Colin Cotter's paper in this volume.

[Cotter and Shipton \(2012\)](#) demonstrate that their proposed mixed finite elements preserve the following desirable properties of the C-grid method when applied to linear barotropic wave propagation: mass conservation; energy conservation; steady geostrophic modes on an f plane; absence of spurious pressure modes; and the flexibility to adjust the balance of velocity degrees of freedom to pressure degrees of freedom, so as to achieve a 2:1 ratio.

Two approaches using this combination have been developed: a low-degree primal-dual one, and a family of primal-only ones. The low-degree primal-dual scheme has similar accuracy to that currently employed in the UM and achieves all of the required properties. By using a dual grid (a secondary grid that makes explicit the staggering of variables, akin to C-grid staggering), it additionally permits the use of any familiar finite-volume transport scheme. Extending this primal-dual finite-element method to higher degree (as might be required near the edges and corners of a cubed-sphere mesh for example) is an, as yet, unsolved problem. However primal-only mixed finite-element schemes have been developed ([Cotter and Shipton, 2012](#)) that also have all the required properties, but the higher-degree forms of these schemes have an isolated problem for the numerical dispersion of inertia-gravity waves. A method of addressing this problem has been developed - see section 4.4 - which works well in one dimension and for certain forms of two-dimensional grids. Work is ongoing to see how successful this is more generally, or whether an alternative solution can be found. (The ability to use a triangular mesh would also be an attractive option for an ocean model.)

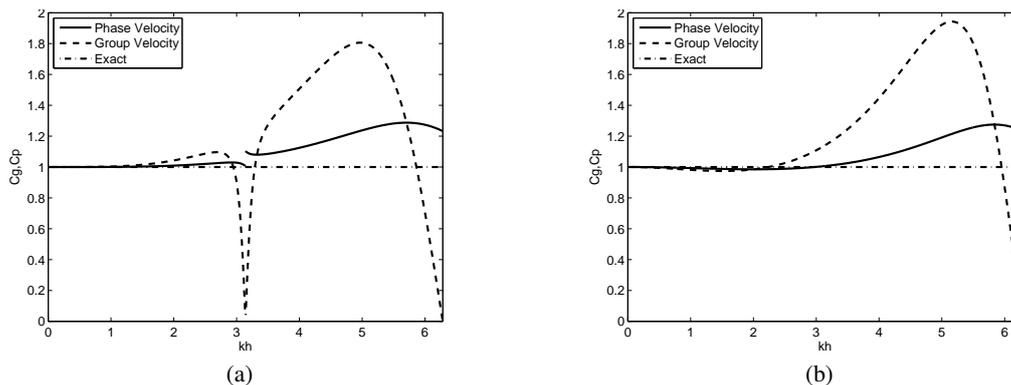


Figure 4: Normalised (with respect to exact) phase and group velocities vs. nondimensional wavenumber kh for two discretisations of the 1D gravity-wave equations: (a) original [Cotter and Shipton \(2012\)](#) RT1 scheme; and (b) [Staniforth et al. \(2013\)](#)'s partially-mass-lumped modification of this scheme.

4.4 Numerical dispersion

Even when everything theoretically looks good, including respect of the necessary 2:1 ratio of momentum and pressure degrees of freedom, and appropriate choice of finite-element function spaces for good representations of geostrophic adjustment and Rossby-mode propagation, things can still unfortunately go wrong. Numerical dispersion analysis is an important and powerful tool for diagnosing and understanding problems ([Staniforth and Thuburn, 2012](#)) and, in the present context, it has proven key for the assessment of two candidate methods: spectral element and finite element.

The spectral-element method has (with exact time integration) the desirable attribute of locally and globally conserving mass, energy and potential vorticity. It also respects the 2:1 ratio, scales well on massively parallel computers and, with appropriate spectral truncation, promises good accuracy. So it appears to be an almost ideal method. However, [Melvin et al. \(2012\)](#) show analytically that it has poor numerical dispersion properties. For the short-wavelength part of the spectrum, energy propagates in the wrong direction at all spectral truncations. At higher spectral truncations, gaps appear in the spectrum of frequencies that can be represented, and eigenmode structures are localised near element boundaries. Numerical integrations confirm that this leads to reversed group velocities and to grid imprinting at element boundaries.

As discussed in sections 4.2 and 4.3, [Cotter and Shipton \(2012\)](#)'s RT1 and BDFM1 mixed finite-element schemes also appear to be close to ideal for similar reasons. However [Staniforth et al. \(2013\)](#) show, both theoretically and experimentally, that there is a serious dispersion issue when discretising the 1-d gravity-wave equations using the Raviart-Thomas RT1 element pair, a particular member of [Cotter and Shipton \(2012\)](#)'s proposed families. For wavelengths close to twice the element width, there is a spectral gap in the dispersion relation - see Fig. 4(a) - such that the group velocity spuriously and strongly goes to zero, with serious implications for energy propagation. [Staniforth et al. \(2013\)](#) further argue that this problem can also be expected to occur for analogous quadrilateral and triangular element pairs in two dimensions. For the very simple model problem examined, [Staniforth et al. \(2013\)](#) show that, without losing any of the desirable properties of the RT1 element pair, this deficiency can be addressed by a partial lumping of the mass matrix associated with the time tendency term of the momentum equation - see Fig. 4(b) for the modified dispersion curves, with no spectral gap and much-improved phase and group velocities.

To practically illustrate the identified dispersion problem, [Staniforth et al. \(2013\)](#) perform various time integrations of the 1D gravity-wave equations using the [Cotter and Shipton \(2012\)](#) RT1 scheme, without and with [Staniforth et al. \(2013\)](#)'s partially-mass-lumped modification. All time integrations are performed to time $t = 0.75$ using Matlab software with a Runge-Kutta timestepping scheme, 120 elements

in a two-units-wide periodic domain, and Φ_0 (mean value of geopotential height) set to unity.

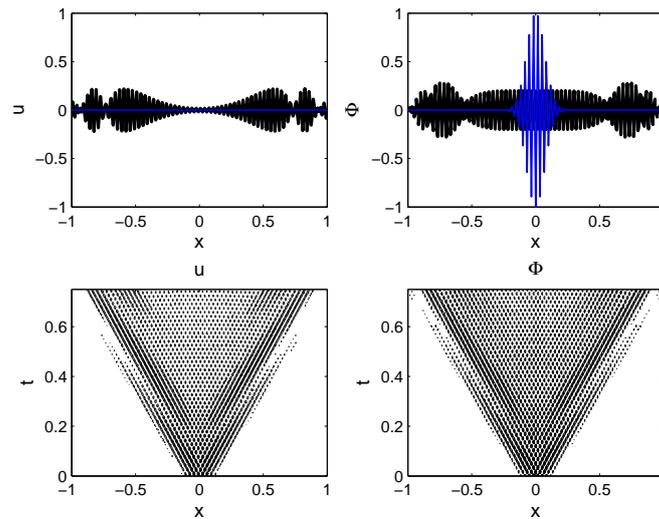


Figure 5: Results for integrations of 1D gravity-wave equations using the original [Cotter and Shipton \(2012\)](#) RT1 scheme. Upper panels: u (left) and Φ (right) at time $t = 0.75$, with the initial condition for Φ plotted in blue. Lower panels: time evolution of u (left) and Φ (right) as a function of x . Contour interval is 0.05 for all panels, and only positive contours are plotted on the lower panels.

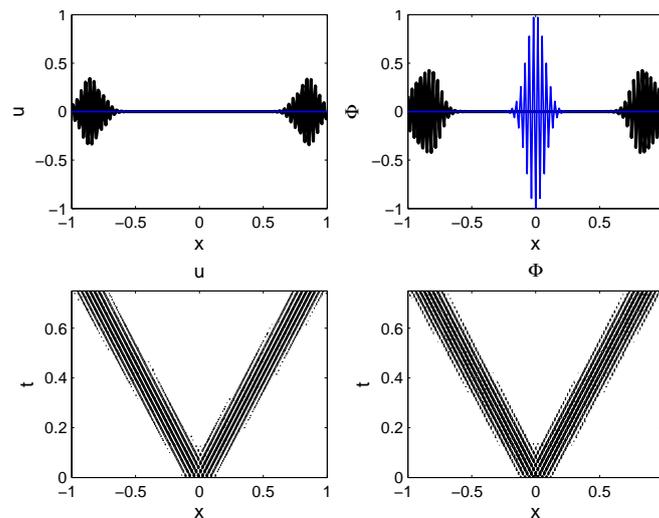


Figure 6: As in in Fig. 5, but for [Staniforth et al. \(2013\)](#)'s partially-mass-lumped modified scheme.

The initial structure for Φ corresponds to a pure mode of wavelength equal to twice the element width, modulated by a Gaussian envelope of unit amplitude. The initial wind field is set identically zero. These initial conditions (Fig. 5) correspond to a pair of leftward and rightward, symmetrically-propagating, gravity waves, each having half-unit amplitude, which reinforce one another at initial time to give unit amplitude. The exact solution after time $t = 0.75$ is comprised of two symmetrically-located wave packets of half-unit amplitude: these are centred on $x = \pm 0.75$ since the exact group velocity is $\pm\sqrt{\Phi_0} = \pm 1$.

Results of the integrations using the unmodified and modified schemes are displayed in Figs. 5-6, respectively. As predicted, the unmodified scheme behaves in an unphysical manner (Fig. 5). Although there are two symmetric wave packets approximately centred about their correct locations, they are of reduced amplitude. Importantly, there is also significant spurious wave activity between them, of similar

amplitude, where there should be none.

Staniforth *et al.* (2013)'s partially-mass-lumped scheme behaves in a much more physical manner (Fig. 6). The two wave packets propagate in a properly-separated manner, albeit slightly too quickly, and the half-unit amplitude for Φ is well maintained. This is consistent with analysis which predicts (Fig. 4(b)) that the group velocity is accelerated by approximately 15% for wavelengths around $kh = \pi$.

So far so good: but can Staniforth *et al.* (2013)'s partially-mass-lumped 1D scheme be extended to 2D? Whilst examining this question, it was realised that although their analysis is correct, it is incomplete. Melvin *et al.* (2013) first remedy this oversight by extending Staniforth *et al.* (2013)'s dispersion analysis of the 1D gravity-wave equations to the 1D linear shallow-water equations on an f plane. They show that for there to be no spectral gap, the presence of the Coriolis terms leads to an additional constraint on the discretisation parameters. Without the imposition of this further constraint, there remains a serious dispersion problem, including the special case of pure inertial oscillations.

Having addressed this oversight, Melvin *et al.* (2013) then show that the resulting 1D partial-mass-lumping approach can indeed be successfully extended to the discretisation of the 2D shallow-water equations on an f plane by a tensor-product approach using rectangular elements. Numerical simulations confirm that:

- the unmodified RT1 scheme leads to wavepackets of inertia-gravity modes at the two-element gridscale being unable to propagate away from the region of initial disturbance; and
- the partial-lumping scheme addresses this problem, allowing the wavepackets to properly disperse.

4.5 RT1 mixed finite elements on a cubed sphere

The modified RT1 scheme is of interest because of the concern that the fairly large variation of grid cell angles on the cubed sphere grid may lead to unsatisfactory grid imprinting for the (lowest-degree) RT0 scheme. To show that the RT1 scheme is viable for shallow-water wave propagation on the sphere, the partial lumping approach needs to be extended to non-affine quadrilaterals on a non-orthogonal, quasi-uniform, cubed-sphere grid. The hypothesis is that finite-element schemes can be constructed by mapping back to a reference square via the Piola transform, as proposed in Cotter and Shipton (2012). This approach has been implemented, and some preliminary results are now presented.

The experimental design is similar to that used in Melvin *et al.* (2013) for 2D integrations of the shallow-water equations on an f plane, but here for the gravity-wave equations on a cubed sphere, with a 20×20 equi-angular grid on each panel. The initial structure for $\hat{\Phi}$ (see Fig. 7(a)) corresponds to a pure two-element gridscale mode of amplitude $\hat{\Phi} \equiv \Phi_0/10$, superimposed on a mean state $\Phi_0 = 100 \text{ m}^2 \text{ s}^{-2}$: this is then modulated by a Gaussian envelope of unit amplitude, centred on a panel of the cubed sphere. The initial wind fields are set identically zero. These initial conditions are unbalanced and generate four outward-propagating gravity wave packets. Because Earth's rotation is suppressed in this experiment, these wave packets should symmetrically propagate towards the four panel edges until they reach them. They should then continue to propagate symmetrically across these four inter-panel edges.

Integrations have been performed without and with the partial mass lumping. Snapshots at time $t = 250$ h of the perturbation are displayed in Figs. 7(b) and (c), respectively. For the unmodified RT1 scheme, it is seen that it behaves in an unphysical manner, similar to that described in section 4.4 for integrations with the 1-d analogue of this scheme. Although there are four, albeit fuzzily defined, propagating wave packets (Fig. 7(b)), there is also significant spurious wave activity between them, of comparable amplitude. This illustrates the poor dispersion properties of the unmodified RT1 scheme for disturbances with wavelengths equal or close to twice the element width. The spurious wave activity between the four wave packets is however absent (Fig. 7(c)) from the integration using the partially-mass-lumped RT1 scheme,

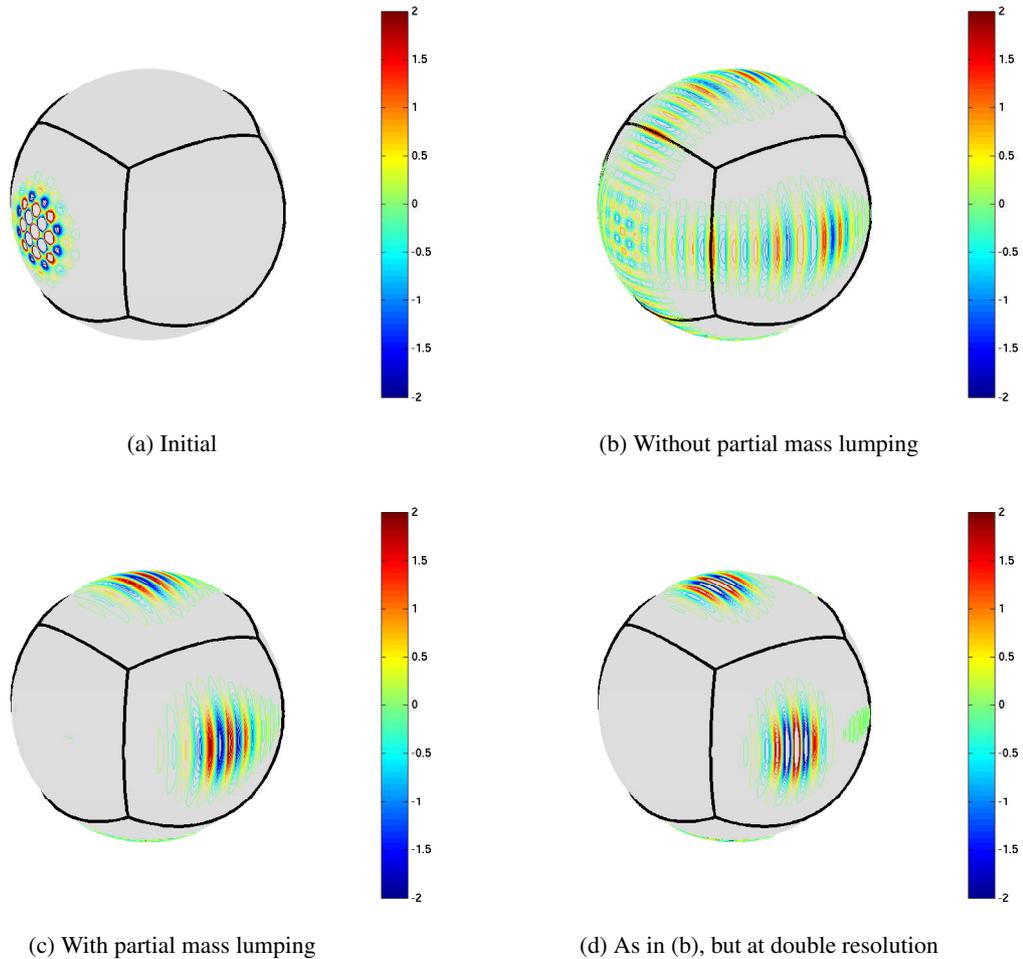


Figure 7: perturbation (m^2s^{-2}) at: (a) initial time; (b) time $t = 250$ h using the unmodified RT1 scheme; (c) as in (b), but using the partially-mass-lumped RT1 scheme; and (d) as in (b), but at double resolution. Contour interval = $0.2 \text{ m}^2\text{s}^{-2}$.

as desired, and the four wave packets have cleanly propagated across the inter-panel edges.

An integration without partial mass lumping has also been performed, but at double resolution. It is seen (Fig. 7(d)) that the spurious wave activity between the four wave packets observed in Fig. 7(b) for this scheme at standard resolution is absent. This is because the doubling of resolution means that most of the energy of the wave packets is no longer at, or near, the problematic scale $kh = \pi$ (Fig. 4(a)), but has been shifted to $kh = \pi/2$ (Fig. 4(b)) where the phase and group velocities behave almost perfectly.

4.6 Time schemes

Many dynamical cores use a semi-implicit timestepping scheme for stability and efficiency. Although this allows an order-of-magnitude larger timestep than would otherwise be possible, it introduces the overhead of solving a 3D elliptic boundary-value problem each timestep. For the semi-implicit scheme to be advantageous, this problem must be solved efficiently. It is widely believed that the overhead becomes prohibitively expensive on massively-parallel computer architectures due to interprocessor communication costs. However, [Mueller and Scheichl \(2013\)](#) have recently demonstrated very good scalability of Krylov subspace solvers and multigrid algorithms for a representative 3D model equation with more than 10^{10} unknowns on 65536 cores. This important finding suggests that the semi-implicit scheme is likely

to be viable for some years to come.

Should the semi-implicit scheme ultimately prove to be computationally too costly, then [Lock *et al.* \(2013\)](#) and [Weller *et al.* \(2013\)](#) have investigated various approaches to efficiently solve hyperbolic wave equations using HEVI (**H**orizontally-**E**xplicit **V**ertically-**I**mplicit) splitting schemes. See Sarah-Jane Lock's paper in the present volume.

5 Conclusion

At the halfway point, much progress has been made on the goals of the GungHo! Project: some is summarised above, and some in the companion papers in this volume of Colin Cotter, Sarah-Jane Lock and John Thuburn. The principal focus of the project thus far has been horizontal discretisation, particularly of the shallow-water equations. At the time of writing (July 2013), the most promising options for this are variants of the families of mixed finite-element schemes proposed in [Cotter and Shipton \(2012\)](#). There are however still some outstanding issues needing further work. Three mixed finite-element pairs are under active consideration. With our current assessment of their relative merits and outstanding issues, they are:

1. RT0 quadrilaterals on a cubed sphere

This is the lowest-degree finite-element pair with the targeted desirable properties, including absence of spurious modes. A concern, however, is that the fairly-large variation of grid cell angles on the cubed sphere grid may lead to unsatisfactory accuracy, and to grid imprinting where the truncation error is particularly large, *viz.* near the corners and edges of the six panels.

2. RT1 quadrilaterals on a cubed sphere

This finite-element pair, in principle, addresses the possible accuracy and grid-imprinting concern for the RT0 element pair. As originally formulated, the RT1 element pair suffers from a dispersion problem at the two-element gridscale, as discussed in section 4.4. However, the preliminary results presented in section 4.5 suggest that this may be addressed by partial mass lumping, but how robustly it does so in more general circumstances remains to be demonstrated. There could also be residual accuracy and grid imprinting issues near the corners and edges of the six panels of the cubed sphere where resolution varies most rapidly.

3. BDFM1 triangles on an icosahedral grid

The BDFM1 element pair is the lowest-order one of the BDFM family that respects the necessary 2:1 ratio of momentum and pressure degrees of freedom for avoidance of spurious modes. However, as discussed above, it almost certainly suffers from a serious dispersion problem at the two-element gridscale, similar to that of the RT1 element pair. A key, and as-yet unresolved, issue is whether or not the partial-mass-lumping remedy of the RT1 quadrilateral element pair (with tensor-product basis functions) can be adapted to BDFM1 triangular elements. Tensor-product basis functions are possible on quadrilaterals, but not on triangles, so if this is an essential (as opposed to convenient) ingredient for the success of partial mass-lumping for the RT1 element pair on quadrilaterals, then the approach will not generalise to triangles. However if partial mass lumping can be made to work for triangles, then this would lead to more-uniform resolution (and less risk of grid imprinting) on an icosahedral grid than for RT1 quadrilaterals on a cubed sphere. It would also be a promising method for ocean modelling, due to the geometric flexibility of triangles to represent complex coastlines.

Common to all three of the above (and also other) methods is the important issue of how they respond in the presence of forcing, and whether this leads to grid imprinting or other problems. To date little work has been done on this due to a lack of time. A concern here is that the special vertices of cubed-sphere

and icosahedral grids occur on scales for which the baroclinic waves of Earth's atmosphere are most unstable. There is therefore a risk that any grid imprinting could spuriously magnify any errors in the representation of this crucially important phenomenon, and thereby unacceptably degrade accuracy.

Regarding temporal discretisation (section 4.6), work on efficiently solving 3D elliptic boundary-value problems on massively-parallel computer architectures suggests that the semi-implicit scheme is likely to be viable for some years to come. Should the semi-implicit scheme ultimately prove too costly, then a viable alternative has been investigated.

Work (not reported here) is now underway on vertical discretisation issues.

Whatever algorithmic schemes are developed and adopted, they need to be efficiently implemented on massively-parallel computer architectures, and significant effort has been devoted to this by GungHo!'s computational scientists. This has led to a code design to implement the emerging GungHo! dynamical core on exascale computers, and to enable it to be coupled to physics parametrisations, ocean models, sea ice models etc. (Ford *et al.*, 2013). The design accommodates the three above-described mixed finite-element schemes under consideration, as well as many other discretisations over the sphere.

References

- Brown A, Milton S, Cullen M, Golding B, Mitchell J, Shelly A. 2012. Unified modeling and prediction of weather and climate: a 25-year journey. *Bull. Amer. Meteor. Soc.* **93**: 1865–1877.
- Cotter CJ, Shipton J. 2012. Mixed finite elements for numerical weather prediction. *J. Comput. Phys.* **231**: 7076–7091.
- Cotter CJ, Thuburn J. 2013. A finite element exterior calculus framework for the rotating shallow-water equations. *J. Comp. Phys.* (submitted, available at <http://arxiv.org/abs/1207.3336>).
- Davies T, Cullen M, Malcolm A, Mawson M, Staniforth A, White A, Wood N. 2005. A new dynamical core for the Met Office's global and regional modelling of the atmosphere. *Q. J. R. Meteorol. Soc.* **131**: 1759–1782.
- Ford R, Glover MJ, Ham DA, Maynard CM, Pickles SM, Riley GD, Wood N. 2013. Gung Ho: A code design for weather and climate prediction on exascale machines. *Adv. Eng. Softw.* Submitted.
- Le Roux DY, Staniforth A, Lin CA. 1998. Finite elements for shallow-water equation ocean models. *Mon. Wea. Rev.* **126**: 1931–1951.
- Lock SJ, Wood N, Weller H. 2013. Numerical analyses of Runge-Kutta implicit-explicit schemes for horizontally-explicit vertically-implicit solutions of atmospheric models. *Q. J. R. Meteorol. Soc.* Submitted.
- Mayne NJ, Baraffe I, Acreman DM, Smith C, Browning MK, Amundsen DS, Wood N, Thuburn J, Jackson D. 2013. ENDGame for hot Jupiters I: HD 209458b test using the UK Met Office non-hydrostatic, deep-atmosphere Global Circulation Model. *Astron. Astrophys.* Submitted.
- Melvin T, Staniforth A, Cotter C. 2013. A two dimensional mixed finite element pair on rectangles. *Q. J. R. Meteorol. Soc.* doi:10.1002/qj.2189. (in press).
- Melvin T, Staniforth A, Thuburn J. 2012. Dispersion analysis of the Spectral Element Method. *Q. J. R. Meteorol. Soc.* **138**: 1934–1947.
- Mueller E, Scheichl R. 2013. Massively parallel solvers for elliptic PDEs in numerical weather- and climate prediction. *Q. J. R. Meteorol. Soc.* Submitted.

- Staniforth A. 2014. Deriving consistent approximate models of the global atmosphere using Hamilton's principle. *Q. J. R. Meteorol. Soc.* Submitted.
- Staniforth A, Melvin T, Cotter C. 2013. Analysis of a mixed finite-element pair proposed for an atmospheric dynamical core. *Q. J. R. Meteorol. Soc.* **139**: 1239–1254, doi:10.1002/qj.2028.
- Staniforth A, Thuburn J. 2012. Horizontal grids for global weather prediction and climate models: a review. *Q. J. R. Meteorol. Soc.* **138**: 1–26.
- Thuburn J, Cotter CJ. 2012. A framework for mimetic discretization of the rotating shallow-water equations on arbitrary polygonal grids. *SIAM J. Sci. Comp.* **34**(3): B203–B225.
- Thuburn J, Ringler TD, Skamarock WC, Klemp JB. 2009. Numerical representation of geostrophic modes on arbitrarily structured C-grids. *J. Comput. Phys.* **228**: 8321–8335.
- Thuburn J, Staniforth A. 2004. Conservation and linear Rossby-mode dispersion on the spherical C grid. *Mon. Wea. Rev.* **132**: 641–653.
- Weller H, Lock SJ, Wood N. 2013. Runge-Kutta IMEX schemes for the Horizontally Explicit/Vertically Implicit (HEVI) solution of wave equations. *J. Comp. Phys.* **252**: 365–381.
- Weller H, Thuburn J, Cotter CJ. 2012. Computational modes and grid imprinting on five quasi-uniform spherical C-grids. *Mon. Wea. Rev.* **140**: 2734–2755.
- White AA, Hoskins BJ, Roulstone I, Staniforth A. 2005. Consistent approximate models of the global atmosphere: shallow, deep, hydrostatic, quasi-hydrostatic and non-hydrostatic. *Q. J. R. Meteorol. Soc.* **131**: 2081–2107.
- White AA, Staniforth A, Wood N. 2008. Spheroidal coordinate systems for modelling global atmospheres. *Q. J. R. Meteorol. Soc.* **134**: 261–270.
- White AA, Wood N. 2012. Consistent approximate models of the global atmosphere in non-spherical geopotential coordinates. *Q. J. R. Meteorol. Soc.* **138**: 980–988.
- Wood N, Staniforth A, White A, Allen T, Diamantakis M, Gross M, Melvin T, Smith C, Vosper S, Zerroukat M, Thuburn J. 2013. An inherently mass-conserving semi-implicit semi-Lagrangian discretisation of the deep-atmosphere global nonhydrostatic equations. *Q.J.R. Meteorol. Soc.* (to appear).

Acknowledgments

The authors thank Paul Selwood for providing Fig. 1, and are happy to acknowledge the considerable input from, and many fruitful discussions with, all of the GungHo! team. GungHo! is jointly funded by the Met Office, the Natural Environment Research Council, and the Science & Technology Facilities Council.

