

Semi-implicit time-integrators for a scalable spectral element atmospheric model

By FRANCIS X. GIRALDO*

Naval Research Laboratory, Monterey, USA

(Received 10 December 2003; revised 7 March 2005)

SUMMARY

The Naval Research Laboratory's spectral element atmospheric model (NSEAM) for scalable computer architectures is presented. This new dynamical core is based on a high-order spectral element (SE) method in space and uses semi-implicit methods in time based on either the traditional second-order leapfrog (LF2) or second-order backward difference formulas (BDF2). The novelties of NSEAM are: it is geometrically flexible and thereby can accommodate any type of grid; LF2 or BDF2 are used to construct the semi-implicit method; and the horizontal operators are written, discretized, and solved in three-dimensional Cartesian space. The semi-implicit NSEAM is validated using: five baroclinic test cases; direct comparisons to the explicit version of NSEAM which has been extensively tested and the results previously reported in the literature; and comparisons with operational weather prediction and well-established climate models. A comparison with the US Navy's spectral transform global forecast model illustrates that NSEAM is 60% faster on an IBM SP4 using 96 processors for the current operational resolution of T239 L30. However, NSEAM can accommodate many more processors while continuing to scale efficiently even at higher grid resolutions. In fact, we show that at T498 L60, NSEAM scales linearly up to 384 processors.

KEYWORDS: Backward difference formula Hexahedral Hydrostatic Icosahedral Leapfrog Primitive equations

1. INTRODUCTION

The current trend in high performance computing has shifted to the development of systems having tens of thousands of processors; the two fastest computers in the world reported at Supercomputing 2004 have 32 000 (IBM, BlueGene/L) and 10 000 (SGI) processors. In fact BlueGene/L is expected to reach its full capacity of 131 000 processors by June 2005. Therefore, to fully exploit this type of architecture requires utilizing numerical methods that rely on a decomposition of the global domain into a multitude of smaller subdomains. Methods that rely on domain decomposition are known as local methods whereas those that do not are referred to as global methods because they require the information of the entire global domain in order to operate on a specific subdomain. The best example of a global method is the spectral transform (ST) method. Examples of local methods include the finite difference (FD), finite element (FE), and finite volume (FV) methods. However, the biggest disadvantage of local methods is that they have not been able to compete, in terms of accuracy, with ST methods which have been used traditionally in operational numerical weather prediction (NWP) and climate models.

Spectral element (SE) methods combine the local domain decomposition property of FE methods with the high-order accuracy and weak numerical dispersion of ST methods. SE methods have shown promise in many areas of the geosciences including: seismic wave propagation (Komatitsch and Tromp 1999), deep earth flows (Fournier *et al.* 2004a), climate (Thomas *et al.* 2002; Fournier *et al.* 2004b), ocean (Iskandarani *et al.* 2002), and NWP (Giraldo and Rosmond 2004) modelling. These methods are high-order FE methods where the grid points are chosen to be the Legendre–Gauss–Lobatto points. In Giraldo and Rosmond (2004) we introduced an SE atmospheric model with an explicit leapfrog time-integrator that was shown to scale linearly while achieving

* Corresponding address: Naval Research Laboratory, Monterey, CA 93943-5502, USA.

e-mail: giraldo@nrlmry.navy.mil

© Royal Meteorological Society, 2005.

accuracies similar to those obtained with ST models. However, in that paper it was shown that the explicit SE model would only outperform ST models at resolutions beyond T406. In order to surpass ST models at resolutions below T406 requires upgrading the time-integrator from explicit to semi-implicit. Explicit time-integrators are too inefficient for NWP applications because the fast-moving gravity waves require the use of small time steps to maintain stability. In order to ameliorate this rather stringent time-step restriction, researchers have discretized the gravity-wave terms implicitly in time and the Rossby wave terms explicitly; this is the idea behind the semi-implicit method (Kwizak and Robert 1971).

In this paper we describe a new semi-implicit Eulerian atmospheric model based on: the SE method in space where the horizontal operators are written, discretized, and solved in three-dimensional (3-D) Cartesian space; and a discretization in time by the second-order leapfrog (LF2) and backward difference formulas (BDF2). Eulerian atmospheric models typically use LF2 for their semi-implicit method; examples include the US Navy (Hogan and Rosmond 1991) and the National Center for Environmental Prediction (Trémolet and Sela 1999) models. The reasons for experimenting with BDF2 are:

- (i) BDF2 are absolutely stable in the region of interest for these equations (for real and distinct eigenvalues);
- (ii) the resulting computational modes are damped and, therefore, no time-filter is required (time-filters diminish the order of accuracy);
- (iii) the resulting pseudo-Helmholtz matrix has a smaller condition number than the one obtained with the LF2 method which translates into fewer iterative solves per time step and results in a more efficient model.

However, BDF2 is by no means ideal for this class of equations (Hamiltonians) and their weaknesses will be discussed in section 2.

The advantages of using Cartesian coordinates are: the pole singularity which plagues the equations in spherical coordinates disappears and the numerical model is completely independent from the grid. Because the numerical method is constructed independently from the grid, this then permits any grid to be used including: icosahedral, hexahedral, telescoping, and adaptive unstructured grids. This independence from the grid is not shared by any of the existing and newly proposed global atmospheric models including the FD model in Davies *et al.* (2005), the FE model in Côté *et al.* (1998), the FV model in Lin and Rood (1996), the icosahedral models in Randall *et al.* (2002) and Majewski *et al.* (2002), the SE model in Thomas *et al.* (2002), and the ST models in Hack *et al.* (1992), Hogan and Rosmond (1991), Temperton *et al.* (2001), and Trémolet and Sela (1999). In fact, the formulations of all these models are restricted to a specific grid geometry.

Element-based Galerkin (EBG) methods, such as the SE method, offer many more benefits in addition to permitting the use of any grid. For example, to switch from quadrilateral to triangular elements merely requires changing the basis functions, and the associated quadrature points and weights as is done in Giraldo and Warburton (2005) which are being considered for future implementation into NSEAM. Changing from globally conservative to locally conservative methods only requires changing the element boundary conditions to account for fluxes as is done in Giraldo *et al.* (2002); this then simplifies the construction of adaptive solutions in addition to giving a fully conservative method (the development of non-hydrostatic ocean and atmospheric models using this approach is currently underway). In short, whenever a new contribution from approximation theory emerges, the new basis functions (and associated quadrature) can

be easily implemented into the existing EBG model; an example is the current work on non-polynomial expansions derived from the prolate spheroidal wave functions (Boyd 2004) which we are currently testing. This flexibility in EBG methods allows an existing model to adapt to the changing needs in science and computing which justifies the further development of SE models and should ensure their longevity.

The objective of the present work is to introduce a new grid-point semi-implicit atmospheric model which:

- (i) is spectrally accurate;
- (ii) is highly scalable on distributed-memory computers;
- (iii) allows for the use of any type of grid;
- (iv) facilitates its continuing augmentation in accuracy and efficiency due to its element-based construction.

The present work essentially extends the explicit time-integrator of NSEAM to semi-implicit. For convergence rates of the discrete horizontal operators the reader is referred to the article by Giraldo and Rosmond (2004).

The remainder of the paper is organized as follows. Section 2 describes the construction of the semi-implicit time-integrators. Section 3 contains: a description of the implementation of the model on distributed-memory computers using the Message-Passing Interface standard; a scalability comparison between NSEAM and the US Navy's Operational Global Atmospheric Prediction System (NOGAPS); and a performance comparison between the LF2 and BDF2 semi-implicit time-integrators of NSEAM. In section 4 the results for the five test cases used to validate the model are presented. Finally, in section 5 we summarize the key findings of this research. For completeness, appendix A contains a description of the semi-implicit method applied to the hydrostatic primitive equations discretized in time by a general second-order time-integrator and in space by the SE method in a Cartesian coordinate system.

2. SEMI-IMPLICIT TIME-INTEGRATORS

The governing equations solved in the present work are the hydrostatic primitive equations (HPE). We assume an adiabatic atmosphere (i.e. no diabatic forcing) and thus only take into account the dynamical processes. Let the HPE be written in the following vector form

$$\frac{\partial \mathbf{q}}{\partial t} = \mathbf{S}(\mathbf{q}), \quad (1)$$

where $\mathbf{q} = (\pi, \mathbf{u}^T, \theta)^T$ is the state vector containing the prognostic variables, $\mathbf{u} = (u, v, w)^T$, T is the transpose operator, and

$$\mathbf{S}(\mathbf{q}) = - \left\{ \begin{array}{c} \nabla \cdot (\pi \mathbf{u}) + \frac{\partial}{\partial \sigma} (\pi \dot{\sigma}) \\ \mathbf{u} \cdot \nabla \mathbf{u} + \dot{\sigma} \frac{\partial \mathbf{u}}{\partial \sigma} + \frac{2\Omega z}{a^2} (\mathbf{x} \times \mathbf{u}) + \nabla \phi + c_p \theta \frac{\partial P}{\partial \pi} \nabla \pi + \mu x \\ \mathbf{u} \cdot \nabla \theta + \dot{\sigma} \frac{\partial \theta}{\partial \sigma} \end{array} \right\} \quad (2)$$

is the source vector function. For closure we require the hydrostatic equation

$$\frac{\partial \phi}{\partial P} = -c_p \theta. \quad (3)$$

For completeness, we define the terms contained in these equations. The terms (a, Ω) are the radius and angular rotation of the earth, respectively. The prognostic variables are: $\pi = p_S - p_T$, where p_S is the surface pressure and p_T is the pressure at the top of the model; the wind velocities \mathbf{u} ; the potential temperature θ . The diagnostic variables are: the vertical velocity $\dot{\sigma}$; the geopotential height ϕ ; the pressure p . Other variables requiring definition are: the Cartesian coordinate of the grid points, x ; the vertical coordinate, $0 \leq \sigma \leq 1$, defined from the top of the atmosphere to the surface of the planet; the Exner function P ; the coefficient of specific heat for constant air pressure, c_p . Finally, the term μ is a Lagrange multiplier required only because we use a 3-D momentum equation in Cartesian coordinates to represent the corresponding 2-D momentum equation in spherical coordinates (see Côté 1988). With the equations defined we can now proceed to the description of the semi-implicit time-integrators.

(a) *General form of second-order semi-implicit time-integrators*

Before describing the implementation of the semi-implicit (SI) method it is crucial to understand which terms must be discretized implicitly. The maximum characteristic wave speed of the HPE, Eqs. (1) and (2), is given by $U + \sqrt{\phi}$ where $U = \mathbf{u} \cdot \mathbf{n}$ is the wind speed along the direction \mathbf{n} and $\sqrt{\phi}$ is the speed of the gravity waves. The fastest gravity waves may travel up to six times faster than the fastest wind velocities. In the SI method, the terms responsible for the propagation of the gravity waves are treated implicitly and the remaining terms explicitly. This essentially slows down the gravity waves which does not adversely affect the medium-range forecast skill because they only carry a small amount of energy. It should be mentioned that there exist alternatives to the SI method with the leading contender perhaps being the Jacobian-free Newton–Krylov method (see Knoll and Keyes (2004), for example) in which the entire set of equations are solved fully implicitly in time. However, in the present work we shall only consider the SI approach.

In Eq. (2) the source terms contributing to the propagation of gravity waves (G) are

$$\mathbf{S}^G(\mathbf{q}) = - \left\{ \begin{array}{c} \nabla \cdot (\pi \mathbf{u}) + \frac{\partial}{\partial \sigma} (\pi \dot{\sigma}) \\ \nabla \phi + c_p \theta \frac{\partial P}{\partial \pi} \nabla \pi \\ \dot{\sigma} \frac{\partial \theta}{\partial \sigma} \end{array} \right\}. \tag{4}$$

We then seek a solution to the equations recast in the following form

$$\frac{\partial \mathbf{q}}{\partial t} = \{\mathbf{S}(\mathbf{q}) - \delta \mathbf{L}^G(\mathbf{q})\} + \delta [\mathbf{L}^G(\mathbf{q})], \tag{5}$$

where the terms inside the curly brackets are time-integrated explicitly, those inside the square brackets implicitly, \mathbf{L}^G (defined in Eq. (A.2)) represents the linearization of \mathbf{S}^G , and $\delta = 0$ or 1 depending on whether the method is purely explicit or semi-implicit.

To simplify the following discussion, let us write the integration of Eq. (5) in the following general form

$$\mathbf{q}^{n+1} = \sum_{m=0}^1 \alpha_m \mathbf{q}^{n-m} + \gamma \Delta t \sum_{m=0}^2 \beta_m \mathbf{S}(\mathbf{q})^{n-m} + \delta \gamma \Delta t \sum_{m=-1}^2 \rho_m \mathbf{L}(\mathbf{q})^{n-m}, \tag{6}$$

where Table 1 lists the associated coefficients corresponding to the BDF2 and LF2 methods and ϑ is the explicit/implicit weighting with $\vartheta = 0.5$ yielding the trapezoidal

TABLE 1. BACKWARD DIFFERENCE FORMULAS (BDF2) AND LEAPFROG (LF2) TIME-INTEGRATORS AND THEIR ASSOCIATED COEFFICIENTS CORRESPONDING TO EQ. (6)

Method	α_0	α_1	γ	β_0	β_1	β_2	ρ_{-1}	ρ_0	ρ_1	ρ_2
BDF2A	4/3	-1/3	2/3	2	-1	0	1	-2	1	0
BDF2B	4/3	-1/3	2/3	8/3	-7/3	2/3	1	-8/3	7/3	-2/3
LF2	0	1	2	1	0	0	ϑ	-1	$1 - \vartheta$	0

rule (also known as Crank–Nicholson). The BDF2A method shown in Table 1 is the method proposed by Karniadakis *et al.* (1991) for the incompressible Navier–Stokes equations and used by Shen and Wang (1999) for the HPE, while the BDF2B method was proposed by Hulstén (1996, personal communication) but has not been used or further studied. Note that, unlike the BDF2 methods, the LF2 method requires the application of the following time filter (Asselin 1972)

$$\tilde{\mathbf{q}}^n = \mathbf{q}^n + \epsilon(\mathbf{q}^{n+1} - 2\mathbf{q}^n + \tilde{\mathbf{q}}^{n-1}), \quad (7)$$

where $\tilde{\mathbf{q}}$ denotes the time-filtered variable with the time-filter weight ϵ .

By exploiting the linearity of the \mathbf{L} operator we can rewrite Eq. (6) as

$$\mathbf{q}^{n+1} = \sum_{m=0}^1 \alpha_m \mathbf{q}^{n-m} + \gamma \Delta t \sum_{m=0}^2 \beta_m \mathbf{S}(\mathbf{q})^{n-m} + \delta \gamma \Delta t \mathbf{L} \left(\sum_{m=-1}^2 \rho_m \mathbf{q}^{n-m} \right). \quad (8)$$

Furthermore, we can simplify this equation by extracting the fully explicit solution from Eq. (8) as follows

$$\mathbf{q}^{\text{explicit}} = \sum_{m=0}^1 \alpha_m \mathbf{q}^{n-m} + \gamma \Delta t \sum_{m=0}^2 \beta_m \mathbf{S}(\mathbf{q})^{n-m}. \quad (9)$$

Multiplying Eq. (8) by ρ_{-1} , and adding $\sum_{m=0}^2 \rho_m \mathbf{q}^{n-m}$ yields

$$\mathbf{q}_{tt} = \hat{\mathbf{q}} + \delta \gamma \Delta t \rho_{-1} \mathbf{L}(\mathbf{q}_{tt}), \quad (10)$$

where

$$\hat{\mathbf{q}} = \rho_{-1} \mathbf{q}^{\text{explicit}} + \sum_{m=0}^2 \rho_m \mathbf{q}^{n-m} \quad (11)$$

and

$$\mathbf{q}_{tt} = \sum_{m=-1}^2 \rho_m \mathbf{q}^{n-m} \equiv \rho_{-1} \mathbf{q}^{n+1} + \sum_{m=0}^2 \rho_m \mathbf{q}^{n-m}. \quad (12)$$

Equation (10) is the form that we use for the construction of the semi-implicit method along with the definitions in Eqs. (9), (11), and (12). The subscript tt in the semi-implicit state vector \mathbf{q} is meant to emphasize the similarity between the semi-implicit correction and a temporal second-order derivative. This is quite evident for BDF2A where $\mathbf{q}_{tt} = \mathbf{q}^{n+1} - 2\mathbf{q}^n + \mathbf{q}^{n-1}$. Note that the form given in Eq. (10) is only possible by using the linearization in Eq. (5) which allows the semi-implicit method to be written as a correction to the explicit method. This can be seen by taking $\delta = 0$ and equating Eqs. (12) and (11) which results in $\mathbf{q}^{n+1} = \mathbf{q}^{\text{explicit}}$. It should be mentioned that constructing the semi-implicit method as a correction to an explicit method as shown in Eq. (10) has been adopted universally by many, if not most, of the operational NWP

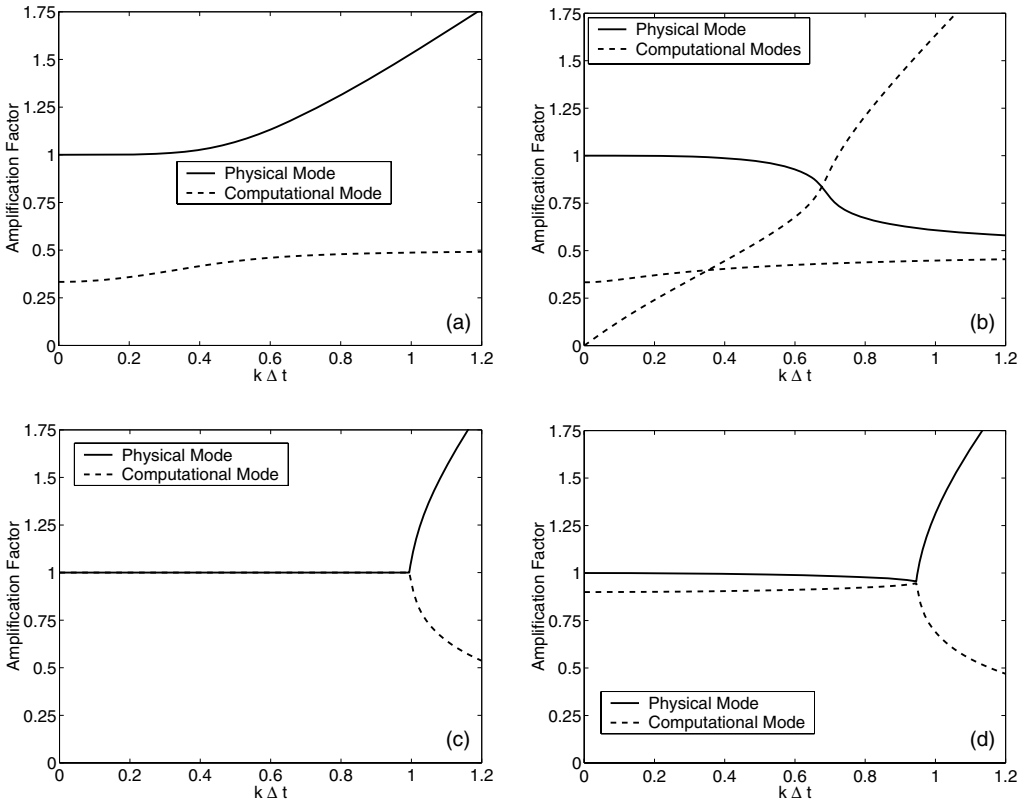


Figure 1. The stability of the explicit versions of (a) BDF2A, (b) BDF2B, (c) LF2 with $\epsilon = 0$, and (d) LF2 with $\epsilon = 0.05$. The solid lines represent the physical solutions and the dashed lines are the computational modes. See text for explanation.

centres including the European Centre for Medium-Range Weather Forecasting (see Ritchie *et al.* 1995), the National Center for Environmental Prediction, and the US Navy. In the following discussion we use the equation

$$\frac{\partial \mathbf{q}}{\partial t} = i k \mathbf{q}$$

to evaluate the stability properties of BDF2 and LF2. In Fig. 1 we show the stability region for the explicit formulations of BDF2 and LF2. Because BDF2 and LF2 both yield multiple numerical solutions to a first-order equation this then means that only one solution is physical while the others are purely computational. BDF2A clearly is inferior to the other two methods because it becomes unstable quite early near $k\Delta t = 0.15$ (Fig. 1(a)); however, the computational solution remains damped and thereby does not require the use of a filter. The physical solution of the BDF2B method is completely stable but becomes quite damped for increasing time step; however, one of the computational solutions eventually becomes unstable near $k\Delta t = 0.73$ (Fig. 1(b)). In contrast, since the original LF2 scheme is symplectic (i.e. it exactly conserves all Lagrangian integral invariants), it is completely undamped (Fig. 1(c)). This is true for both the physical and computational modes. For the physical mode this is a highly desirable property but not for the computational mode because it can become excited through nonlinear interactions with the physical mode and eventually become unstable

(see Sanz-Serna 1985 and Aoyagi 1995). For this reason, in the geophysical fluid dynamics community LF2 is typically used in conjunction with a time filter (Asselin 1972) which damps the computational mode while selectively modifying the physical mode (Fig. 1(d)). This is by no means the only choice available for eradicating the computational mode. In the astrophysics community, which places much importance on symplecticness (which is ideal for Hamiltonian systems such as the first-order HPE) the approach has been to introduce second-order Runge–Kutta smoothers (see Aoyagi 1995; New *et al.* 1998) which obviate the need for time-filters; however, in the present work we only consider LF2 with the Asselin time-filter.

Based on the above discussion it is safe to conclude that BDF2B and LF2 with $\epsilon = 0.05$ are the most stable schemes shown in Fig. 1. Let us now compare these two methods. Clearly, the ideal method would be one whereby the physical solution is undamped while the computational solution is damped. BDF2B and LF2 with $\epsilon = 0.05$ damp both the physical and computational solutions. The question we now try to answer is how much unwanted dissipation do these two methods introduce? Assuming that 1% numerical dissipation is an acceptable level (i.e. the amplification factor approaches 0.990) we find that BDF2B reaches this value at $k\Delta t = 0.38$ while LF2 with $\epsilon = 0.05$ reaches it at $k\Delta t = 0.58$. Thus BDF2B is much more dissipative than LF2 with $\epsilon = 0.05$. However, for $\epsilon = 0.1$, LF2 reaches this level at $k\Delta t = 0.41$ which is as dissipative as BDF2B. In the present work we use $\epsilon = 0.05$ for LF2 for most of the simulations but for very long time-integrations (such as the Held–Suarez test case) we use $\epsilon = 0.1$.

In Fig. 2 we show the stability region for the implicit formulations of BDF2 and LF2. For the fully implicit formulation ($\delta = 1$) BDF2A and BDF2B collapse to the classical BDF2 method found in numerical analysis text books (e.g. Gear 1971, p. 217). The reason there is no computational mode for BDF2 in Fig. 2(a) is because the amplification factor of this mode remains below 0.32 and hence is not visible in this amplified plot. Clearly BDF2 and LF2 are unconditionally stable. However, this stability is achieved at the price of damping the physical solution. Once again, if we take 1% dissipation as an acceptable level we find that BDF2 reaches it at $k\Delta t = 0.51$ (Fig. 2(a)) while LF2 with $\epsilon = 0.05$ and $\vartheta = 0.5$ at $k\Delta t = 0.73$ (Fig. 2(b)). For $\epsilon = 0.1$ and $\vartheta = 0.5$, LF2 reaches this level at $k\Delta t = 0.46$ (Fig. 2(c)) which is now below BDF2. Changing the implicit weight of LF2 has even more drastic consequences. For example, for $\epsilon = 0.05$ and $\vartheta = 0.6$, LF2 reaches 1% dissipation at $k\Delta t = 0.22$ (Fig. 2(d)). In fact, for this choice of ϵ and ϑ , LF2 is more dissipative than BDF2 for the entire range of $k\Delta t$ values. Values of $\vartheta > 0.5$ are not uncommon in NWP models and in fact are often used to off-centre semi-Lagrangian schemes in order to eliminate unwanted noise near steep topography (see Rivest *et al.* 1994; Côté *et al.* 1995; Davies *et al.* 2005). BDF2 has the advantage of being naturally off-centred and therefore able to avoid the spurious resonant response induced by orographic forcing.

In summary, LF2 with $\epsilon = 0$ and $\vartheta = 0.5$ is a better method than BDF2 because it is non-dissipative; however, as shown by Sanz-Serna (1985), LF2 can become unstable due to the nonlinear interaction between the physical and computational modes even when linear stability analysis shows otherwise. By choosing $\epsilon > 0$ and/or $\vartheta > 0.5$, LF2 becomes more stable but at the price of becoming dissipative and losing its second-order accuracy. Thus, if one is willing to accept a small amount of dissipation, have off-centring built into the time-integrator, but would like to retain second-order accuracy, then BDF2 is a reasonable choice. The stability analysis also shows that, due to its larger explicit stability region, LF2 admits a larger time step than BDF2, at least for the Rossby waves. However, as we shall see in section 3(b), a larger time step does not necessarily translate into a faster model—at least not for pure dynamics (i.e. no diabatic forcing).

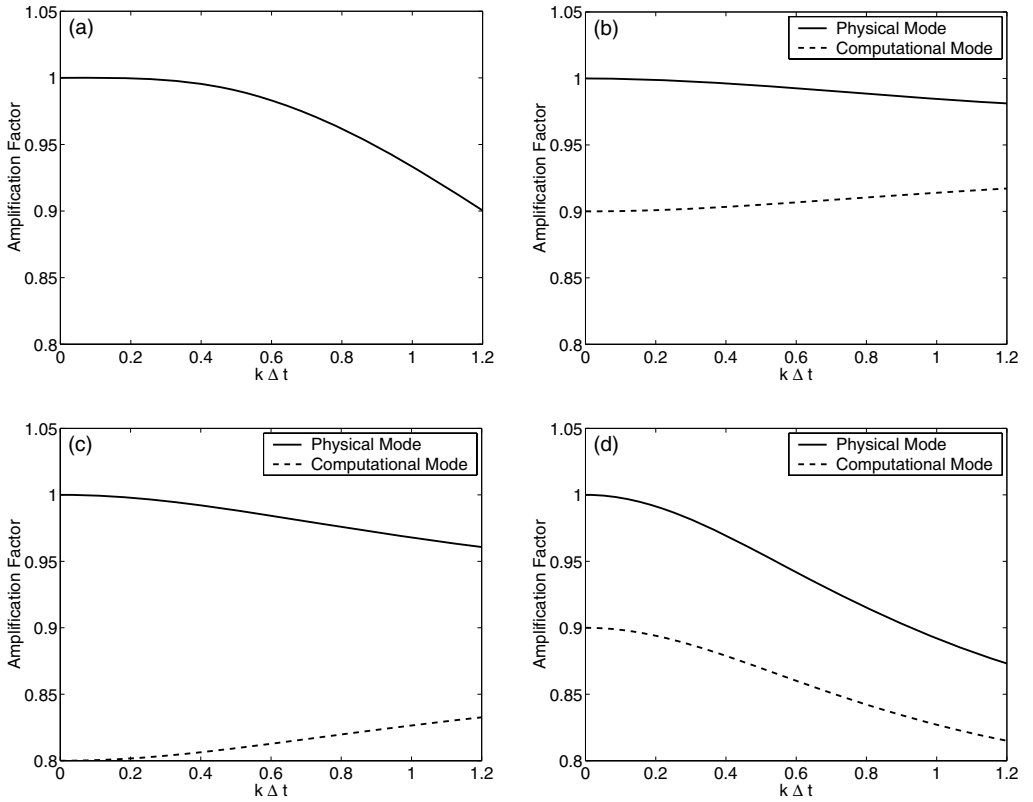


Figure 2. The stability of the implicit versions of (a) BDF2, (b) LF2 with $\epsilon = 0.05$ and $\vartheta = 0.5$, (c) LF2 with $\epsilon = 0.1$ and $\vartheta = 0.5$, and (d) LF2 with $\epsilon = 0.05$ and $\vartheta = 0.6$. See text for explanation.

The success of BDF2 for the shallow-water equations using a spectral element semi-Lagrangian method (Giraldo *et al.* 2003), the incompressible Navier–Stokes equations using Eulerian time-integrators (Karniadakis *et al.* 1991), semi-Lagrangian methods (Xiu and Karniadakis 2001), and operator-integration-factor splitting methods (Maday *et al.* 1990) has been the main motivation for considering BDF2 for the semi-implicit version of NSEAM. We now turn to the construction of the Helmholtz operator resulting from the generalized semi-implicit time-integration.

(b) The 3-D pseudo-Helmholtz operator

After application of the semi-implicit method, as outlined in appendix A for the generalized second-order form given in Eq. (8), we obtain the following 3-D pseudo-Helmholtz equation

$$\Phi_l - \lambda^2 V_{l,m} (\mathbf{M}^{-1} \mathbf{D}^T \mathcal{P} \mathbf{M}^{-1} \mathbf{D} \Phi)_m = \hat{\Phi}_l - \lambda V_{l,m} (\mathbf{M}^{-1} \mathbf{D}^T \mathcal{P} \hat{\mathbf{u}})_m \quad (13)$$

for the variable Φ which is a linear combination of the potential temperature and surface pressure semi-implicit corrections (see Eq. (A.12)). To properly explain the solution strategy of this 3-D pseudo-Helmholtz problem, a detailed description of Eq. (13) is first in order.

In Eq. (13) \mathbf{M} and \mathbf{D} are the mass and differentiation matrices resulting from the spectral element discretization, \mathcal{P} is the projection matrix which constrains the 3-D

Cartesian velocities, and \mathbf{V} is the matrix containing the vertical contribution of the semi-implicit method (see Eq. (A.14)). At this point we have only used subscripts for the matrices corresponding to the vertical discretization. In Eq. (13) \mathbf{V} is an $N_{\text{lev}} \times N_{\text{lev}}$ matrix with $l, m = 1, \dots, N_{\text{lev}}$ which is defined for every horizontal grid point $i = 1, \dots, N_p$ where N_{lev} are the number of vertical layers in the model and N_p are the total number of horizontal grid points. Thus the matrix \mathbf{V} couples every vertical layer at each horizontal grid point. The mass matrix \mathbf{M} and the differentiation matrix \mathbf{D} are $N_p \times N_p$ matrices with $i, j = 1, \dots, N_p$, and thus only involve the horizontal direction. Furthermore, the differentiation matrix \mathbf{D} is a vector of matrices such that $\mathbf{D} = (D^x \mathbf{i} + D^y \mathbf{j} + D^z \mathbf{k})$ where (i, j, k) are the Cartesian directional unit vectors, and D^s denotes the differentiation matrix along the s direction.

To further simplify the discussion let us define two additional horizontal matrices: let

$$H_{i,j}^L = \mathbf{M}^{-1} \mathbf{D}^T \mathcal{P} \mathbf{M}^{-1} \mathbf{D} \quad (14)$$

represent the discrete pseudo-Laplacian operator and

$$H_{i,j}^D = \mathbf{M}^{-1} \mathbf{D}^T \mathcal{P} \quad (15)$$

the discrete divergence operator. Substituting Eqs. (14) and (15) into Eq. (13), factoring terms, and including subscripts for the matrices corresponding to the horizontal and vertical discretizations yields

$$[I_{l,m} \otimes I_{i,j} - \lambda^2 (V_{l,m} \otimes H_{i,j}^L)] \Phi_{j,m} = \widehat{\Phi}_{l,i} - \lambda (V_{l,m} \otimes H_{i,j}^D) \widehat{u}_{j,m}, \quad (16)$$

where \otimes denotes the tensor product; $I_{l,m}$ and $I_{i,j}$ are identity matrices associated with the vertical and horizontal discretizations, respectively. Let us now describe some important details of the matrix problem defined in Eq. (16). The tensor product $I_{l,m} \otimes I_{i,j}$ results in the identity matrix $I_{l,i,m,j}$ which is an $N_{\text{lev}}^2 \times N_p^2$ matrix; the same is true for the tensor products of \mathbf{V} with \mathbf{H}^L and \mathbf{H}^D . Thus if we replaced the left-hand-side matrix of Eq. (16), which is in the square brackets, by \mathbf{H}^{3D} then it becomes evident that the product of $\mathbf{H}_{l,i,m,j}^{3D}$ with the solution matrix $\Phi_{j,m}$ results in

$$G_{l,i} = H_{l,i,m,j}^{3D} \Phi_{j,m},$$

where \mathbf{G} is an $N_{\text{lev}} \times N_p$ matrix that spans the entire 3-D domain because $l = 1, \dots, N_{\text{lev}}$ spans the vertical and $i = 1, \dots, N_p$ spans the horizontal directions.

At this point, the Helmholtz operator of the semi-implicit formulation given in Eq. (16) yields a fully 3-D matrix. Below we describe the vertical mode decomposition which transforms the 3-D HPE into a series of 2-D shallow-water equations which are then solved much more efficiently.

(c) Vertical-mode decomposition

Note that the matrix \mathbf{H}^L in Eq. (16) is completely independent of the vertical direction while the matrix \mathbf{V} is independent of the horizontal direction; however, the tensor product of \mathbf{V} and \mathbf{H}^L results in a fully 3-D system. Thus to reduce this intimidatingly large 3-D matrix problem we apply a mode decomposition of the matrix \mathbf{V} in the vertical direction. Upon completion of this decomposition the full 3-D Helmholtz matrix problem will be converted into a series of N_{lev} 2-D Helmholtz problems.

In order to perform this vertical decomposition we write the matrix \mathbf{V} in the canonical form

$$\mathbf{V} = \mathbf{R} \mathbf{\Lambda} \mathbf{R}^{-1},$$

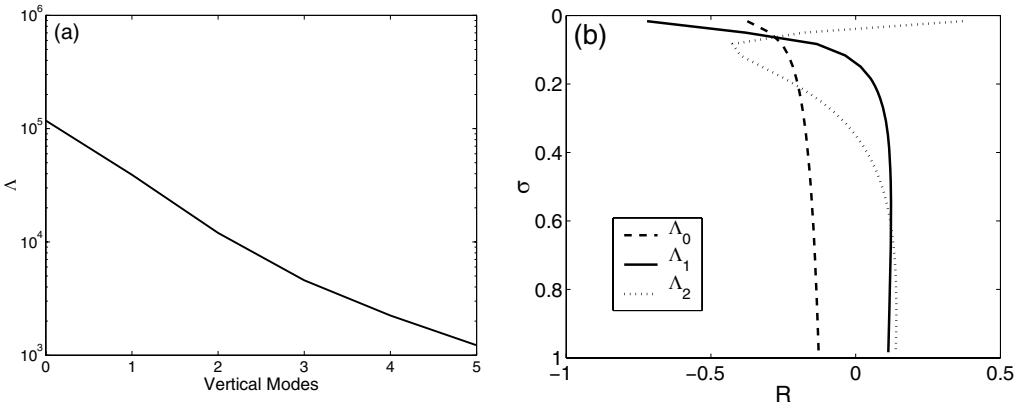


Figure 3. (a) The eigenvalues Λ of the six fastest vertical modes, and (b) the right eigenvectors \mathbf{R} of the three fastest vertical modes. See text for explanation.

where Λ are the eigenvalues of \mathbf{V} and \mathbf{R} are its associated right eigenvectors which satisfy

$$\mathbf{V}\mathbf{R} = \mathbf{R}\Lambda.$$

Note that this eigenvalue problem is not computationally expensive because \mathbf{V} is an $N_{\text{lev}} \times N_{\text{lev}}$ matrix which only needs to be decomposed once. Upon obtaining this vertical mode decomposition we simply left-multiply Eq. (16) by \mathbf{R}^{-1} to yield

$$(I_{i,j} - \lambda^2 \Lambda_l H_{i,j}^L)(R_{l,k}^{-1} \Phi_{k,j})^T = [R_{l,k}^{-1} (\hat{\Phi}_{k,i} - \lambda (V_{k,m} \otimes H_{i,j}^D) \hat{u}_{j,m})]^T. \tag{17}$$

Letting

$$\Phi_{j,l}^R = (R_{l,k}^{-1} \Phi_{k,j})^T,$$

we then obtain the following series of 2-D Helmholtz problems

$$(I_{i,j} - \lambda^2 \Lambda_l H_{i,j}^L) \Phi_{j,l}^R = [R_{l,k}^{-1} (\hat{\Phi}_{k,i} - \lambda (V_{k,m} \otimes H_{i,j}^D) \hat{u}_{j,m})]^T \tag{18}$$

for the variable Φ^R . Note that, for each of the $l = 1, \dots, N_{\text{lev}}$ modes, Eq. (18) represents an individual 2-D Helmholtz problem of size $N_p \times N_p$, which is in fact analogous to a shallow-water model scaled by the eigenvalue Λ_l .

Figure 3(a) shows that the eigenvalues Λ of the six fastest vertical modes decrease exponentially. In fact, the third eigenvalue Λ_2 has a value of 10^4 , the square root of which is already in the range of the highest horizontal wind velocities encountered in the atmosphere ($\sim 100 \text{ m s}^{-1}$). Thus the semi-implicit method is only required for the vertical modes which have gravity-wave speeds beyond the highest horizontal wind velocities. For this reason we only solve Eq. (18) for the first four vertical modes in much the same way first proposed by Burridge (1975). Figure 3(b) shows the right eigenvectors \mathbf{R} , as a function of σ for the three fastest vertical modes. Let us now turn our attention to the solution of the 2-D Helmholtz problems.

(d) *Solution of the 2-D pseudo-Helmholtz operator*

From the vertical mode analysis it was determined that only the Helmholtz operator corresponding to the first four vertical modes need to be solved implicitly in time; the remainder can be computed explicitly. This coupling between explicit and implicit

methods can be seamlessly included into the numerical approach by using the explicit solution $\hat{\Phi}$ as the initial guess for the implicit solution Φ . This is why it is so beneficial to write the semi-implicit method as a correction to the explicit method. First, the explicit solution is obtained and then only for the first four vertical modes are the semi-implicit corrections included. The semi-implicit correction is obtained by solving Eq. (18) using the generalized minimum residual method (GMRES) with point Jacobi preconditioning, L_2 projection for the next iterate, and restarts every 10 time steps (see Fischer 1998). There are numerous other Krylov subspace methods to choose from but we have decided on GMRES based on previous experiences (see Giraldo *et al.* 2003). In addition, there are also more elaborate preconditioners and we shall report on our experiences with methods such as overlapping Schwarz (see Pavarino 2002) in future work.

The number of GMRES iteration required for convergence are dependent on the stopping criterion, ϵ_{stop} , and the size of $\lambda^2\Lambda$ in Eq. (18); the smaller ϵ_{stop} or the larger $\lambda^2\Lambda$ the more iterations required. The value of ϵ_{stop} is defined by the user and there are numerous strategies for choosing this value which are beyond the scope of the present work. It turns out that the value of $\lambda^2\Lambda$ is dependent on the eigenmode (more iterations are required for the external mode), the coefficients γ and ρ_{-1} of the semi-implicit method (see Eq. (A.6)), and of course the time-step. It is shown in section 3(b) that increasing the time step by a given factor does not mean that an efficiency gain of this size will be achieved. Similarly, one time-integrator running with a larger time step may not be more efficient than another time-integrator running with a smaller time step. The overall efficiency of the model is determined by the scaling $\lambda^2\Lambda$ which affects the condition number of the resulting Helmholtz problem.

Upon obtaining the solution for Φ^R we then left-multiply by \mathbf{R} to obtain Φ , i.e.

$$\Phi = \mathbf{R}\Phi_R.$$

Once Φ is obtained we then solve for \mathbf{u}_{tt} via Eq. (A.15). With this value of \mathbf{u}_{tt} known we can then solve for π_{tt} via Eq. (A.7) and θ_{tt} via Eq. (A.9). Finally, the prognostic variables are extracted from \mathbf{q}_{tt} via Eq. (12). This then concludes the solution strategy for each time step.

3. PARALLEL IMPLEMENTATION

(a) Model scalability

One of the main advantages of using SE methods over ST methods is that for an equivalent grid resolution the SE method allows the use of far more processors. Assuming that a 1-D decomposition along latitude rings is the most efficient decomposition for ST models (as in the US Navy's operational NWP model, NOGAPS), the maximum number of processors that an ST model can use is

$$N_{\text{proc}}^{\text{ST}} = N_{\text{lat}} \approx \frac{3}{2}T, \quad (19)$$

where N_{lat} denotes the number of latitude rings and T the resolution of the spectral triangular truncation. In contrast, on a hexahedral grid (see Fig. 4(b)) with $N_p = 6(n_H N)^2 + 2$ grid points and $N_e = 6n_H^2$ elements (where n_H and N are the number of elements in each of the x , y directions on each of the six faces of the hexahedron and the polynomial order, respectively) the maximum number of processors that a SE model can use is

$$N_{\text{proc}}^{\text{SE}} = N_e \equiv \frac{6}{N^2}H^2, \quad (20)$$




Figure 4. Three of the many possible grids that NSEAM can use. They include: (a) icosahedral, (b) hexahedral, and (c) telescoping grids. The three grids have approximately the same number of grid points as the spectral grid T239 where the 49 high-order grid points inside each quadrilateral ($N = 6$) have been omitted for clarity.

where $H = n_H N$ represents the hexahedral horizontal resolution. In other words an SE model can use as many processors as there are elements. Thus for fixed N the number of processors allowed by an SE model increases quadratically with resolution H , while only linearly for an ST model. Using the approximation based on equivalent number of grid points given by

$$H \sim \frac{\sqrt{3}}{2}(T + 1), \tag{21}$$

a spectral resolution of T239 translates to a hexahedral resolution of H208 which for simplicity we round up to H216 (e.g. $n_H = 27, N = 8$; $n_H = 36, N = 6$; or $n_H = 54, N = 4$). From Eq. (21) we see that the resolution H216 is equivalent to T249. However, due to the h - p nature of the SE method it can achieve a specified grid resolution by either increasing the number of elements h or the polynomial order p . Thus it is important to state the polynomial order used to obtain a specified grid resolution and so, for the resolution $n_H = 36, N = 6$, we denote it as T249 N6 in order to distinguish it from other possible resolutions such as $n_H = 27, N = 8$ (T249 N8), or $n_H = 54, N = 4$ (T249 N4).

At T239 (which is the current operational resolution used by NOGAPS) an ST model can use 360 processors. At T249 N6 an SE model can use 7700 processors—a twenty-fold increase in the number of processors. Equation (20) shows that, if we wish to further increase the number of processors of the SE model, we simply increase n_H while decreasing N accordingly in order to maintain the horizontal resolution fixed. Therefore, at T249 N4, NSEAM can accommodate well over 17 000 processors—a forty-fold increase in the number of processors. However, decreasing N will impact the solution accuracy and the issue of efficiency versus accuracy must be carefully weighed. The point here is that the SE method offers this flexibility to increase either the accuracy or efficiency—a luxury shared by neither the spectral transform nor finite-difference methods.

(b) Comparison of BDF2 and LF2 semi-implicit time-integrators

Efficiency is arguably one of the most important criteria for determining whether a specific algorithm will be included in an operational NWP model. In Fig. 5 we show the performance of NSEAM using the BDF2B and LF2 semi-implicit time-integrators using double precision on an IBM SP4 P690 with a clockspeed of 1.3 GHz; this P690 was reported as the 56th fastest computer in the world at SuperComputing 2004. The results for BDF2A are quite similar to those for BDF2B and so we shall refer to the BDF




Figure 5. Efficiency of BDF2 and LF2 semi-implicit methods for NSEAM T249 N6 L30 (see text) on an IBM SP4 (P690 with 1.3 GHz clockspeed).

methods simply as BDF2. The resolution is T249 N6 L30 with $\Delta t = 300$ s for BDF2 and LF2 and $\Delta t = 200$ s for BDF2.

The results in Fig. 5 show that the BDF2 semi-implicit formulation is more efficient than LF2 even when BDF2 uses a time step 50% smaller. The reason for this is quite simple: the 2-D Helmholtz problem, given in Eq. (18), which must be solved at each time step, only differs for the two methods by the parameter λ . For LF2 it is equal to $2\vartheta \Delta t$ and for BDF2 it is $(2/3)\Delta t$. Thus we can see that the matrix corresponding to BDF2 is more diagonally dominant than that for LF2 (since $\vartheta > 1/2$), which gives BDF2 a smaller condition number than LF2. A smaller condition number means that the iterative solver will require fewer iterations for a given convergence tolerance; in Fig. 5, on average, BDF2 required 10 GMRES iterations per time step while LF2 required 21. This difference in GMRES iterations allows BDF2 to use a smaller time step and still perform as efficiently as LF2. Another interesting result is the performance of BDF2. For this method using a time step 50% smaller the decrease in performance was not significant especially at high processor counts. BDF2 with $\Delta t = 200$ s required fewer GMRES iterations than with $\Delta t = 300$ s. Clearly, this has a dramatic effect on performance especially at high processor counts.

(c) Comparison of NSEAM and NOGAPS dynamical cores

Figure 6 shows a comparison between the dynamical cores (i.e. no diabatic forcing) of NSEAM and NOGAPS on the IBM SP4 P690 using single precision. For NOGAPS we use the current operational resolution T239 L30 while for NSEAM we use T249 N6 L30 with the BDF2B time-integrator. At T239 L30 the maximum time step that NOGAPS can use is 300 s. At this time step using 96 processors NSEAM is 60% faster than NOGAPS; this gap between the two models will continue to widen with increased resolution and/or processor count in favour of NSEAM as predicted by Eqs. (53) and (54) in Giraldo and Rosmond (2004).

It is expected that the resolution of NOGAPS will increase to T479 L60 in the near future. Thus in Fig. 7 we show the performance of NSEAM at the resolutions T249 N6 L30 (with $\Delta t = 300$ s) and T498 N6 L60 (with $\Delta t = 150$ s) on an IBM SP4 P655 with a clockspeed of 1.7 GHz (the P655 was reported as the 9th fastest computer in the world at SuperComputing 2004). Going from a resolution of T249 N6 L30 to T498 N6 L60 increases the problem size by a factor of 16; decreasing the spacing in all three directions by one half increase the number of grid points by a factor of 8 and decreasing




Figure 6. Simulation days per wallclock hours for the dynamical cores of NSEAM (T249 N6 L30 using BDF2B) and NOGAPS (T239 L30) on an IBM SP4 (P690 with 1.3 GHz clockspeed) using $\Delta t = 300$ s (see text). (The performance data for NOGAPS is courtesy of Tim Hogan.)




Figure 7. Simulation days per wallclock hours for the dynamical core of NSEAM (using BDF2B) on an IBM SP4 (P655 with 1.7 GHz clockspeed) using (a) T249 N6 L30 with $\Delta t = 300$ s, and (b) T498 N6 L60 with $\Delta t = 150$ s. See text for explanation.

the time step by half increases the number of time steps for one simulation day by another factor of 2. Comparing the performance of NSEAM T249 N6 L30 (Fig. 7(a)) with T498 N6 L60 (Fig. 7(b)) shows that the difference in performance is in fact less than the expected value of 16. This implies that NSEAM becomes more efficient for large problem sizes and high processor counts. This can be seen in Fig. 7(b) which shows the linear scalability of NSEAM T498 N6 L60. The largest cost incurred by NSEAM is in the solution of the 2-D Helmholtz problem. However, for 30, 60, or 90 vertical levels, NSEAM only solves a Helmholtz problem for the first four vertical modes thereby getting the remaining vertical modes virtually at no cost (the cost incurred by the remaining vertical modes comes in via the explicit solution). This behaviour is observed in Fig. 7 where one can see that at 384 processors the difference in performance between T249 N6 L30 and T498 N6 L60 is approaching a factor of 8; thus the vertical solution is achieved at almost no cost.

The scalability studies have shown that the NSEAM model performs extremely well; however, there are regions of the model which can still be improved to further increase its performance. Because NSEAM is currently only a research tool it has been




Figure 8. Rossby–Haurwitz wave number 4: the surface pressure contours (hPa) for NSEAM T185 N8 L24 using (a) the BDF2B and (b) the LF2 semi-implicit time-integrators.

designed with flexibility in mind. For example, the grids and the corresponding domain decomposition (DD) are generated automatically within the code. This consumes precious central processor unit time which could be avoided if the grid and its DD were generated off-line and then read in as input. This feature is currently being implemented into the next version of NSEAM which will have terrain, diabatic forcing, and non-reflecting boundary conditions. Another likely candidate for improvement is the solution of the external mode which requires far more GMRES iterations than the remaining modes. It is conjectured that the external mode can be solved much more efficiently by using a coarse grid in a multi-grid approach. However, this topic is reserved for future research.

4. RESULTS

The test cases used to validate NSEAM consist of five baroclinic tests. The difficulty with quantifying the error and/or accuracy of baroclinic models is that analytic solutions are difficult to obtain. Instead we view the following test cases as a means for qualitative comparisons to show that NSEAM gives similar results to existing models. In order to judge and compare the accuracy of NSEAM we plot normalized L_2 error norms defined as follows

$$\|\mathbf{q}\|_{L_2} = \sqrt{\frac{\int_A (\mathbf{q}_{\text{exact}} - \mathbf{q})^2 dA}{\int_A \mathbf{q}_{\text{exact}}^2 dA}}, \quad (22)$$

where \mathbf{q} is the computed solution vector, $\mathbf{q}_{\text{exact}}$ is the exact solution, and A represents the surface area of the earth.

(a) Rossby–Haurwitz wave number 4

In this test case (see Giraldo and Rosmond 2004) we track the propagation of Rossby–Haurwitz waves during a five-day period. Surface pressure contours after a five-day integration for NSEAM T185 N8 L24 with $\Delta t = 300$ s for BDF2B and LF2 semi-implicit formulations are shown in Fig. 8. The results between the two methods are virtually indistinguishable; however, BDF2B required 33% fewer iterations per time step to converge.

(b) Polvani et al. baroclinic instability with diffusion

For this test case (see Polvani *et al.* 2004), the atmosphere is initially balanced and a perturbation is added to the flow which begins the motion of the atmosphere.




Figure 9. Polvani *et al.* test case: the surface temperature of NSEAM T74 N8 L20 after 12 days using (a) the BDF2B and (b) the LF2 semi-implicit time-integrators.

This perturbation drives the atmosphere towards a singularity; however, in order to avoid this unpleasantness a diffusion operator is added to the momentum and thermodynamic equations with viscosity $7.0 \times 10^{-5} \text{ m}^2\text{s}^{-1}$. Figure 9 shows the surface temperature as a function of longitude and latitude at day 12 of the integration for the BDF2B and LF2 semi-implicit methods. The results are essentially identical between the two methods. In addition, these results compare extremely well with the Geophysical Fluid Dynamics Laboratory’s spectral transform model (see Polvani *et al.* 2004). In fact, at this resolution NSEAM has converged to the correct solution regardless of which semi-implicit method is used. Further increases in either horizontal or vertical resolution show no discernible difference in the solution.

(c) *Jablonowski–Williamson balanced initial state*

For this test case (see Jablonowski and Williams 2002), the atmosphere is initially balanced and should remain so indefinitely. Figure 10 shows the normalized π L_2 error norm as a function of time for a 30-day period for NSEAM with the semi-implicit BDF2A, semi-implicit BDF2B, semi-implicit LF2, and explicit LF2 methods for the resolution T185 N8 L26; the explicit method used in this figure is in fact the model described in Giraldo and Rosmond (2004). Note that all the methods use different time steps with the explicit using the smallest (42 s) and the semi-implicit LF2 the largest (270 s). Even though the semi-implicit BDF2 and LF2 methods use time steps much larger than the explicit method, the error norms are quite similar, which confirms that the semi-implicit solutions are as accurate as the explicit one. It is surprising, however, that both semi-implicit BDF2 methods yield virtually indistinguishable results from the explicit method but that the semi-implicit LF2 method does not. For this case LF2 uses a time step 50% larger than BDF2B but is only 210 s faster for the entire 30-day simulation. The advantages of the larger time step for LF2 are somewhat offset by requiring 50% more iterations per time step than BDF2B.

(d) *Jablonowski–Williamson baroclinic instability*

This case is similar to the balanced initial state except that now a perturbation is added to the initial zonal velocity. This perturbation grows until a baroclinic instability develops near day 9. Figure 11 shows the minimum surface pressure p_s , as a function of time for NSEAM (BDF2B) against various models including the National Centers for




Figure 10. Jablonowski–Williamson balanced initial state: the normalized π L_2 error as a function of days for NSEAM T185 N8 L26 using the semi-implicit BDF2A method with $\Delta t = 135$ s (solid line), semi-implicit BDF2B method with $\Delta t = 180$ s (dashed-dotted line), semi-implicit LF2 method with $\Delta t = 270$ s (dotted line), and the explicit LF2 method with $\Delta t = 42$ s (dashed line).




Figure 11. Jablonowski–Williamson baroclinic instability: the minimum surface pressure (hPa) as a function of days for the NASA (finite volume), GME (finite-difference), NCAR (spectral transform), and NSEAM (spectral element with $N = 8$ using BDF2B, see text) models using 26 vertical levels. (The data for the last three models are courtesy of Christiane Jablonowski.)

Atmospheric Research (NCAR) ST model (Hack *et al.* 1992), the National Aeronautics and Space Administration (NASA) Goddard FV model (Lin and Rood 1996), and the German Weather Service icosahedral FD model (Majewski *et al.* 2002) which we denote as GME. For the grid-point models, we use the definition of equivalent triangular truncation

$$T \sim \frac{\sqrt{2N_p}}{3}.$$

The results of this case are summarized as follows. Figure 11 shows that all four models are in complete agreement until day 8, at which point the two low-order models (NASA and GME) diverge from the NCAR and NSEAM models. The two low-order models, NASA and GME, show a similar pattern during the 14 day integration but do not match exactly. On the other hand the two high-order models, NCAR and NSEAM, behave

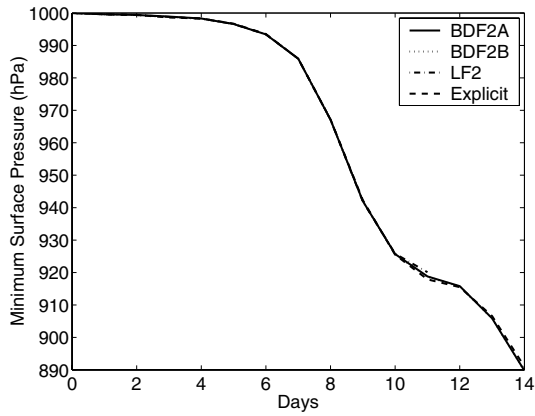


Figure 12. Jablonowski–Williamson baroclinic instability: The minimum surface pressure (hPa) as a function of days for NSEAM T185 N8 L26 using the semi-implicit BDF2A with $\Delta t = 135$ s (solid line), semi-implicit BDF2B method with $\Delta t = 180$ s (dotted line), semi-implicit LF2 method with $\Delta t = 270$ s (dashed-dotted line), and the explicit LF2 method with $\Delta t = 42$ s (dashed line). See text for explanation.




Figure 13. Held–Suarez mean planetary climate: the time and zonally averaged zonal velocity of the semi-implicit (a) BDF2B and (b) LF2 methods as functions of latitude and vertical coordinate for the resolution T74 N8 L20.

almost identically throughout the entire 14 day simulation. Having established that NSEAM using BDF2B behaves similarly to other well-established climate and NWP models, we compare various semi-implicit time-integrators of NSEAM.

In Fig. 12 we plot the minimum surface pressure as a function of days for NSEAM T185 N8 L26 using the semi-implicit BDF2A, semi-implicit BDF2B, semi-implicit LF2, and the explicit LF2. Even though all the models use different time steps they all agree rather well, confirming that the semi-implicit time-integrators are as accurate as the explicit method.

(e) Held–Suarez mean planetary climate

For this test case (see Held and Suarez 1994), the atmosphere is initially at rest and a perturbation is added to the flow which begins the motion of the atmosphere. A forcing function mimicking the radiation of the sun near the equator drives the model towards a realistic mean planetary climate. NSEAM was run for 1200 days with samples taken every 4 days beginning from day 200. The sample files from day 200 to day 1200

are then averaged to obtain a temporal mean. Figure 13 shows the time and zonally averaged zonal velocity for the semi-implicit BDF2B (Fig. 13(a)) and LF2 (Fig. 13(b)) as a function of latitude and vertical coordinate. Both models yield identical midlatitude jets in the upper atmosphere which agree with the results obtained with the explicit model in Giraldo and Rosmond (2004) and the spectral transform model in Held and Suarez (1994).

5. CONCLUSION

The Naval Research Laboratory's spectral element atmospheric model (NSEAM) for scalable computer architectures was presented. NSEAM is based on a Cartesian formulation of the equations, the spectral element method in space, and a general second-order semi-implicit method in time. Specifically, semi-implicit methods based on backward difference formulas (BDF2) and leapfrog (LF2) were compared. The stability analysis showed that, for reasonable values of the Asselin time-filter ϵ , LF2 is less dissipative than BDF2 as long as the implicit weight ϑ is equal to 0.5; if $\vartheta > 0.5$ then LF2 becomes more dissipative than BDF2. For pure dynamics BDF2 was shown to be as accurate as LF2 while being more efficient, even though it required a smaller time step. This difference in efficiency is due to the resulting Helmholtz matrix having a smaller condition number for BDF2 than LF2, which then translates into fewer iterations per time step to converge.

We showed that, regardless of time-integrator, NSEAM gives similar results to well-established climate and weather-prediction models while scaling quite efficiently on distributed-memory computers. At a resolution of T249 L30 using 96 processors, NSEAM was shown to be 60% faster than a spectral transform model and this gap will continue to grow in favour of NSEAM as the horizontal resolution and the number of processors are increased. At T498 L60 NSEAM was shown to scale linearly up to 384 processors—a feat impossible to achieve by a spectral transform model. Because SE models are constructed completely around basis functions this offers attractive flexibilities not shared by other methods. The shape, order, and characteristics of the grid, polynomials and model can be altered merely by changing the basis functions. This flexibility allows an SE model to adapt to the changing needs in science and computing throughout its lifetime. Finally, the advantage of using Cartesian coordinates with the SE method is that the model becomes completely independent from the grid. This means that any type of grid can be used with NSEAM. While we have only shown results on hexahedral grids, the extension to icosahedral, telescoping, and adaptive grids is immediately obvious. Various improvements in the accuracy, efficiency and flexibility of NSEAM are currently underway and the results of this research will be reported in the future.

ACKNOWLEDGEMENTS

This work was supported by the Office of Naval Research through programme element PE-0602435N. All computer simulations were performed on the IBM SP4 at the Department of Defense NAVO Major Shared Resource Center. I would like to thank Craig Bishop and João Teixeira for their helpful comments on an early draft of the manuscript. I also would like to thank Terry Davies and Lorenzo Polvani for allowing me access to their unpublished works. Finally, I would like to thank Andrew Staniforth, Clive Temperton, and the anonymous reviewer for their detailed review which helped to improved the manuscript.

APPENDIX A

Semi-implicit formulation: Linearization of the nonlinear gravity-wave terms \mathbf{S}^G

The terms in Eq. (4) are linearized in the following manner. First, π and θ are linearized about the reference states $\pi^* = 1000$ hPa and

$$\theta^* = \frac{T^*}{P(\pi^*)},$$

with $T^* = 300$ K (see Temperton *et al.* 2001). With these reference states defined we can now proceed with the linearization of the terms in Eq. (4).

(i) *Surface pressure and thermodynamic equations.* Based on the previously defined reference states, the surface pressure and thermodynamic equations are linearized in the following simple form

$$\mathbf{S}^G(\pi) = -\pi^* \left(\nabla \cdot \mathbf{u} + \frac{\partial \dot{\sigma}}{\partial \sigma} \right) \quad \text{and} \quad \mathbf{S}^G(\theta) = -\dot{\sigma} \frac{\partial \theta^*}{\partial \sigma},$$

respectively.

(ii) *Momentum equation.* In order to linearize the pressure gradient the geopotential height itself must be linearized. Beginning with the finite-differenced equation for the geopotential height

$$\phi_l - \phi_{l+1} = c_p \theta_l (P_{l+1/2} - P_l) + c_p \theta_{l+1} (P_{l+1} - P_{l+1/2}),$$

where $l = 1, \dots, N_{\text{lev}}$ with N_{lev} representing the number of vertical levels, we can then take a Taylor series expansion about the reference states π^* and θ^* yielding

$$\begin{aligned} \phi_l - \phi_{l+1} = & c_p \theta_l (P_{l+1/2}^* - P_l^*) + c_p \theta_{l+1} (P_{l+1}^* - P_{l+1/2}^*) \\ & + c_p \theta_l^* \left(\frac{\partial P_{l+1/2}^*}{\partial \pi} - \frac{\partial P_l^*}{\partial \pi} \right) (\pi - \pi^*) \\ & + c_p \theta_{l+1}^* \left(\frac{\partial P_{l+1}^*}{\partial \pi} - \frac{\partial P_{l+1/2}^*}{\partial \pi} \right) (\pi - \pi^*). \end{aligned} \quad (\text{A.1})$$

Equation (A.1) can be written in matrix form as

$$A_{l,k} \phi_k = B_{l,k} \theta_k + C_l (\pi - \pi^*),$$

where

$$\begin{aligned} A_{l,k} = & \begin{cases} 1, & \text{if } k = l, \\ -1, & \text{if } k = l + 1, \\ 0, & \text{if } k < l \text{ or } k > l + 1, \end{cases} \\ B_{l,k} = & \begin{cases} c_p (P_{l+1/2}^* - P_l^*), & \text{if } k = l, \\ c_p (P_{l+1}^* - P_{l+1/2}^*), & \text{if } k = l + 1, \\ 0, & \text{if } k < l \text{ or } k > l + 1, \end{cases} \\ C_l = & \begin{cases} c_p \theta_l^* \left(\frac{\partial P_{l+1/2}^*}{\partial \pi} - \frac{\partial P_l^*}{\partial \pi} \right) + c_p \theta_{l+1}^* \left(\frac{\partial P_{l+1}^*}{\partial \pi} - \frac{\partial P_{l+1/2}^*}{\partial \pi} \right), & \text{if } l < N_{\text{lev}}, \\ c_p \theta_l^* \left(\frac{\partial P_{l+1/2}^*}{\partial \pi} - \frac{\partial P_l^*}{\partial \pi} \right), & \text{if } l = N_{\text{lev}}, \end{cases} \end{aligned}$$

and the range of the indices are $k, l, m = 1, \dots, N_{\text{lev}}$. This results in the following geopotential gradient

$$\nabla \phi_l = \nabla (A_{l,k}^{-1} B_{k,m} \theta_m + A_{l,k}^{-1} C_k (\pi - \pi^*)),$$

where

$$A_{l,k}^{-1} = \begin{cases} 1, & \text{if } k \geq l, \\ 0, & \text{if } k < l. \end{cases}$$

Linearizing the gradient due to surface pressure yields

$$\left(c_p \theta_l \frac{\partial P_l}{\partial \pi} \nabla \pi \right)^{n+1} = \left(c_p \theta_l^* \frac{\partial P_l^*}{\partial \pi} \nabla \pi^{n+1} \right).$$

To simplify the description of the semi-implicit method let us create the following definitions: let

$$E_{l,k} = A_{l,m}^{-1} B_{m,k} \quad \text{and} \quad F_l = A_{l,k}^{-1} C_k + c_p \theta_l^* \frac{\partial P_l^*}{\partial \pi},$$

which allows the gradient of geopotential and surface pressure to be written as

$$\nabla \phi_l + c_p \theta_l \frac{\partial P_l}{\partial \pi} \nabla \pi = \nabla (E_{l,k} \theta_k + F_l \pi).$$

We are now ready to construct the linear operator $\mathbf{L}^G(\mathbf{q})$ of Eq. (5).

Semi-implicit formulation: Linear operator and implicit correction

With the linearizations described above and dropping the subscripts, the linear operator \mathbf{L}^G in Eq. (5) can be written as follows

$$\mathbf{L}^G(q) = \begin{Bmatrix} \pi^* \left(\nabla \cdot \mathbf{u} + \frac{\partial \dot{\sigma}}{\partial \sigma} \right) \\ \nabla (\mathbf{E} \theta + \mathbf{F} \pi) \\ \dot{\sigma} \frac{\partial \theta^*}{\partial \sigma} \end{Bmatrix}. \quad (\text{A.2})$$

Using this linearization we can now write the equations in terms of the semi-implicit correction as follows

$$\pi_{tt} = \hat{\pi} - \lambda \pi^* \left(\nabla \cdot \mathbf{u}_{tt} + \frac{\partial}{\partial \sigma} (\dot{\sigma}_{tt}) \right), \quad (\text{A.3})$$

$$\mathbf{u}_{tt} = \hat{\mathbf{u}} - \lambda \nabla (\mathbf{E} \theta_{tt} + \mathbf{F} \pi_{tt}) - \mu x, \quad (\text{A.4})$$

$$\theta_{tt} = \hat{\theta} - \lambda \left(\dot{\sigma}_{tt} \frac{\partial \theta^*}{\partial \sigma} \right), \quad (\text{A.5})$$

where

$$\lambda = \delta \gamma \Delta t \rho_{-1} \quad (\text{A.6})$$

and $\hat{\mathbf{q}}$ and \mathbf{q}_{tt} are defined in Eqs. (11) and (12).

Construction of the 3-D pseudo-Helmholtz operator

Replacing the continuous spatial operators in Eqs. (A.3), (A.4) and (A.5) by their discrete spectral element counterparts results in

$$\pi_{tt} = \hat{\pi} - \frac{\lambda \pi^*}{\sigma_S - \sigma_T} \sum_{k=1}^{N_{lev}} \mathbf{M}^{-1} \mathbf{D}^T \mathbf{u}_{tt\,k} \Delta \sigma_k, \quad (\text{A.7})$$

$$\mathbf{u}_{tt\,l} = \mathcal{P} \hat{\mathbf{u}}_l - \lambda \mathcal{P} \mathbf{M}^{-1} \mathbf{D} (E_{l,k} \theta_{tt\,k} + F_l \pi_{tt}), \quad (\text{A.8})$$

$$\theta_{tt\,l} = \hat{\theta}_l - \lambda S_{l,k} (\mathbf{M}^{-1} \mathbf{D}^T \mathbf{u}_{tt})_k, \quad (\text{A.9})$$

where

$$\left. \begin{aligned} S_{l,k} &= \Theta_l \Delta_{N_{lev},k} - \Theta_{l+1/2} \Delta_{l,k} - \Theta_{l-1/2} \Delta_{l-1,k}, \\ \Theta_l &= \Theta_{l+1/2} \frac{\sigma_{l+1/2} - \sigma_T}{\sigma_S - \sigma_T} + \Theta_{l-1/2} \frac{\sigma_{l-1/2} - \sigma_T}{\sigma_S - \sigma_T}, \\ \Theta_{l+1/2} &= \frac{\theta_{l+1/2}^* - \theta_l^*}{\sigma_{l+1/2} - \sigma_l}, \quad \Theta_{l-1/2} = \frac{\theta_l^* - \theta_{l-1/2}^*}{\sigma_l - \sigma_{l-1/2}}, \\ \Delta_{l,k} &= \begin{cases} \Delta \sigma_k, & \text{if } k \leq l, \\ 0, & \text{if } k > l, \end{cases} \quad \Delta \sigma_k = \sigma_{k+1/2} - \sigma_{k-1/2}, \end{aligned} \right\} \quad (\text{A.10})$$

and \mathbf{M} is the mass matrix, \mathbf{D} is the differentiation matrix, and \mathcal{P} is the projection matrix which constrains the Cartesian velocity field to remain tangential to a sigma surface (for details on the SE discretization see Giraldo and Rosmond (2004)). To simplify the description of the semi-implicit formulation we only include subscripts for the matrices involving the vertical discretization because these are the terms to which the vertical mode decomposition is applied.

Multiplying Eq. (A.9) by E and Eq. (A.7) by F and adding gives

$$\Phi_l = \hat{\Phi}_l - \lambda E_{l,k} S_{k,m} (\mathbf{M}^{-1} \mathbf{D}^T \mathbf{u}_{tt})_m - \lambda F_l \frac{\pi^*}{\sigma_S - \sigma_T} \Delta_{N_{lev},m} (\mathbf{M}^{-1} \mathbf{D}^T \mathbf{u}_{tt})_m, \quad (\text{A.11})$$

where

$$\begin{aligned} \Phi_l &= E_{l,k} \theta_{tt\,k} + F_l \pi_{tt}, \\ \hat{\Phi}_l &= E_{l,k} \hat{\theta}_k + F_l \hat{\pi}, \end{aligned} \quad (\text{A.12})$$

and σ_S and σ_T are the sigma values at the surface and top of the model. Let us next factor $\mathbf{M}^{-1} \mathbf{D}^T \mathbf{u}_{tt}$ from Eq. (A.11) and rewrite the equation as

$$\Phi_l = \hat{\Phi}_l - \lambda V_{l,m} (\mathbf{M}^{-1} \mathbf{D}^T \mathbf{u}_{tt})_m \quad (\text{A.13})$$

where

$$V_{l,m} = E_{l,k} S_{k,m} + F_l \frac{\pi^*}{\sigma_S - \sigma_T} \Delta_{N_{lev},m} \quad (\text{A.14})$$

is the matrix containing the vertical contribution to the semi-implicit formulation. Note that Eq. (A.8) can now be written as

$$\mathbf{u}_{tt\,l} = \mathcal{P} (\hat{\mathbf{u}}_l - \lambda \mathbf{M}^{-1} \mathbf{D} \Phi_l). \quad (\text{A.15})$$

Substituting Eq. (A.15) into Eq. (A.13) yields

$$\Phi_l - \lambda^2 V_{l,m} (\mathbf{M}^{-1} \mathbf{D}^T \mathcal{P} \mathbf{M}^{-1} \mathbf{D} \Phi)_m = \hat{\Phi}_l - \lambda V_{l,m} (\mathbf{M}^{-1} \mathbf{D}^T \mathcal{P} \hat{\mathbf{u}})_m.$$

REFERENCES

- Aoyagi, A. 1995 Nonlinear leapfrog instability for Fornberg's pattern. *J. Comput. Phys.*, **120**, 316–322
- Asselin, R. 1972 Frequency filter for time integrations. *Mon. Weather Rev.*, **100**, 487–490
- Boyd, J. P. 2004 Prolate spheroidal wavefunctions as an alternative to Chebyshev or Legendre polynomials for spectral elements and pseudospectral algorithms. *J. Comput. Phys.*, **199**, 688–716
- Burridge, D. M. 1975 A split semi-implicit reformulation of the Bushby–Timpson 10-level model. *Q. J. R. Meteorol. Soc.*, **101**, 777–792
- Côté, J. 1988 A Lagrange multiplier approach for the metric terms of semi-Lagrangian models on the sphere. *Q. J. R. Meteorol. Soc.*, **114**, 1347–1352
- Côté, J., Gravel, S. and Staniforth, A. 1995 A generalized family of schemes that eliminate the spurious resonant response of semi-Lagrangian schemes to orographic forcing. *Mon. Weather Rev.*, **123**, 3605–3613
- Côté, J. C., Gravel, S., Methot, A., Patoine, A., Roch, M. and Staniforth, A. 1998 The operational CMC-MRB global environmental multiscale (GEM) model. Part I: design considerations and formulation. *Mon. Weather Rev.*, **126**, 1373–1395
- Davies, T., Cullen, M. J. P., Malcolm, A. J., Mawson, M. H., Staniforth, A., White, A. A. and Wood, N. 2005 A new dynamical core for the Met Office's global and regional modelling of the atmosphere. *Q. J. R. Meteorol. Soc.* (in press)
- Fischer, P. F. 1998 Projection techniques for iterative solutions of $Ax = b$ with successive right-hand sides. *Computer Methods in Applied Mechanics and Engineering*, **163**, 193–204
- Fournier, A., Bunge, H. P., Hollerbach, R. and Vilotte, J. P. 2004a Application of the spectral element method to the axisymmetric Navier–Stokes equations. *Geophys. J. Int.*, **156**, 682–700
- Fournier, A., Taylor, M. A. and Tribbia, J. J. 2004b The spectral element atmosphere model (SEAM): high-resolution parallel computation and localized resolution of regional dynamics. *Mon. Weather Rev.*, **132**, 726–748
- Gear, C. W. 1971 *Numerical initial value problems in ordinary differential equations*. Prentice-Hall, Englewood Cliffs, NJ, USA
- Giraldo, F. X. and Rosmond, T. E. 2004 A scalable spectral element Eulerian atmospheric model (SEE-AM) for NWP: Dynamical core tests. *Mon. Weather Rev.*, **132**, 133–153
- Giraldo, F. X. and Warburton, T. 2005 A nodal triangle-based spectral element method for the shallow water equations on the sphere. *J. Comput. Phys.*, **207**, 129–150
- Giraldo, F. X., Hesthaven, J. S. and Warburton, T. 2002 Nodal high-order discontinuous Galerkin methods for the spherical shallow water equations. *J. Comput. Phys.*, **181**, 499–525
- Giraldo, F. X., Perot, J. B. and Fischer, P. F. 2003 A spectral element semi-Lagrangian (SESL) method for the spherical shallow water equations. *J. Comput. Phys.*, **190**, 623–650
- Hack, J. J., Boville, B. A., Briegleb, B. P., Kiehl, J. T., Rasch, P. J. and Williamson, D. L. 1992 Description of the NCAR community climate model (CCM2). *NCAR Technical Note NCAR/TN-382+STR*, National Center for Atmospheric Research, Climate Modeling Section, PO Box 3000, Boulder CO 80307
- Held, I. M. and Suarez, M. J. 1994 A proposal for the intercomparison of the dynamical cores of atmospheric general circulation models. *Bull. Am. Meteorol. Soc.*, **75**, 1825–1830
- Hogan, T. F. and Rosmond, T. E. 1991 The description of the Navy Global Operational Prediction System's spectral forecast model. *Mon. Weather Rev.*, **119**, 1786–1815
- Iskandarani, M., Haidvogel, D. B., Levin, J. C., Curchitser E. and Edwards C. A. 2002 Multiscale geophysical modeling using the spectral element method. *Computing in Science and Engineering*, **4**, 42–48
- Jablonski, C. and Williamson, D. L. 2002 'Baroclinic instability test with two jets in the midlatitudes'. P. 10 in Abstracts of the 2002 workshop on the solution of partial differential equations on the sphere, Toronto, Canada
- Karniadakis, G. E., Israeli, M. and Orszag, S. A. 1991 High-order splitting methods for the incompressible Navier–Stokes equations. *J. Comput. Phys.*, **97**, 414–443
- Knoll, D. A. and Keyes, D. E. 2004 Jacobian-free Newton–Krylov methods: A survey of approaches and applications. *J. Comput. Phys.*, **193**, 357–397

- Komatitsch, D. and Tromp, J. 1999 Introduction to the spectral element method for three-dimensional seismic wave propagation. *Geophys. J. Int.*, **139**, 806–822
- Kwizak, M. and Robert, A. J. 1971 A semi-implicit scheme for grid point atmospheric models of the primitive equations. *Mon. Weather Rev.*, **99**, 32–36
- Lin, S. J. and Rood, R. B. 1996 Multidimensional flux-form semi-Lagrangian transport schemes. *Mon. Weather Rev.*, **124**, 2046–2070
- Maday, Y., Patera, A. T. and Ronquist, E. M. 1990 An operator-integration-factor splitting method for time-dependent problems: Application to incompressible fluid flow. *J. Sci. Computing*, **5**, 263–292
- Majewski, D., Liermann, D., Prohl, P., Ritter, R., Buchhold, M., Hanisch, T., Paul, G., Wergen, W. and Baumgardner, J. 2002 The operational global icosahedral-hexahedral gridpoint model GME: Description and high-resolution tests. *Mon. Weather Rev.*, **130**, 319–338
- New, K. C. B., Watt, K., Misner, C. W. and Centrella, J. M. 1998 Stable 3-level leapfrog integration in numerical relativity. *Phys. Rev. D*, **58**, 064022, 1–14
- Pavarino, L. F. 2002 Overlapping Schwarz preconditioners for spectral element discretizations of convection-diffusion problems. *International Journal for Numerical Methods in Engineering*, **53**, 1005–1023
- Polvani, L. M., Scott, R. K. and Thomas, S. J. 2004 Numerically converged solutions of the global primitive equations for testing the dynamical core of atmospheric GCMs. *Mon. Weather Rev.*, **132**, 2539–2552
- Randall, D. A., Ringler, T. D. and Heikes, R. P. 2002 Climate modeling with spherical geodesic grids. *Computing in Science and Engineering*, **4**, 32–41
- Ritchie, H., Temperton, C., Simmons, A., Hortal, M., Davies, T., Dent, D. and Hamrud, M. 1995 Implementation of the semi-Lagrangian method in a high-resolution version of the ECMWF forecast model. *Mon. Weather Rev.*, **123**, 489–514
- Rivest, C., Staniforth, A. and Robert, A. 1994 Spurious resonant response of semi-Lagrangian discretizations to orographic forcing—diagnosis and solution. *Mon. Weather Rev.*, **122**, 366–376
- Sanz-Serna, J. M. 1985 Studies in numerical nonlinear instability. 1: Why do leapfrog schemes go unstable. *SIAM J. Scientific and Statistical Computing*, **6**, 923–938
- Shen, J. and Wang, W. 1999 A fast and accurate numerical scheme for the primitive equations of the atmosphere. *SIAM J. Numer. Anal.*, **36**, 719–737
- Temperton, C., Hortal, M. and Simmons, A. 2001 A two-time-level semi-Lagrangian global spectral model. *Q. J. R. Meteorol. Soc.*, **127**, 111–127
- Thomas, S. J., Loft, R. D. and Dennis, J. M. 2002 Parallel implementation issues: Global versus local methods. *Computing in Science and Engineering*, **4**, 26–31
- Trémolet, Y. and Sela, J. 1999 ‘Parallelization of the NCEP global spectral model’. Pp. 109–116 in *Towards teracomputing*, Proceedings of the 8th ECMWF workshop on the use of parallel processors in meteorology. Eds. W. Zwiefelhofer and N. Kreitz, World Scientific, London
- Xiu, D. and Karniadakis, G. E. 2001 A semi-Lagrangian high-order method for the Navier–Stokes equations. *J. Comput. Phys.*, **172**, 658–684