

On the use of exponential time integration methods in atmospheric models

By COLM CLANCY* and JANUSZ A. PUDYKIEWICZ, *Atmospheric Numerical Weather Prediction Research, Environment Canada, 2121 Route Transcanadienne, Dorval, Québec, Canada H9P 1J3*

(Manuscript received 19 March 2013; in final form 6 November 2013)

ABSTRACT

Exponential integration methods offer a highly accurate approach to the time integration of large systems of differential equations. In recent years, they have attracted increased attention in a number of diverse fields due to advances in their computational efficiency. This has been as a result of the use of Krylov subspace methods for the approximation of the matrix exponentials which typically arise. In this work, we investigate the potential of exponential integration methods for use in atmospheric models. Two schemes are implemented in a shallow water model and tested against reference explicit and semi-implicit methods. In a number of experiments with standard test cases, the exponential methods are found to yield very accurate solutions with time-steps far longer than even the semi-implicit method allows. The relative efficiency of the exponential integrators, which depends mainly on the choice of the specific algorithm used for the calculation of the matrix exponent, is also discussed. The future work aimed at further improvements of the proposed methodology is outlined.

Keywords: numerical weather prediction, time integration, exponential methods, shallow water equations, advection schemes, icosahedral grid

1. Introduction

The choice of the time integration scheme is an important consideration in the design of an atmospheric model. Accuracy is an obvious requirement for any method, but this has to be balanced by stability and efficiency concerns. Particularly for numerical weather prediction (NWP) in an operational context, we would like to be able to use time-steps which are long enough to reduce computational expense, while maintaining sufficient accuracy.

The existence of vastly differing time-scales in atmospheric phenomena, from slow large-scale modes to very fast gravity and sound waves, poses a challenge for time integration techniques. In particular, purely explicit schemes are restricted by severe time-step limitations. A number of approaches, such as semi-implicit and split-explicit schemes, have been used throughout the history of NWP to circumvent these constraints. Details of the various methods may be found in Durran (1999).

An alternative strategy for solving systems of time-dependent differential equations, which has attracted attention in a number of fields, is the use of exponential

integration methods. The extensive review of Minchev and Wright (2005) traces their use back to 1960 and provides the general definition of an exponential integrator as ‘a numerical method which involves an exponential function (or a related function) of the Jacobian or an approximation to it’. These methods are exact for linear systems of ordinary differential equations (ODEs) with constant coefficients and are thus trivially A-stable.

While exponential methods offer the potential for high accuracy, the appearance of matrix exponentials in the solution can be very computationally expensive. This is of particular concern in the solution of partial differential equations (PDEs), where the spatial discretisation results in very large, albeit sparse, matrices. This rendered the methods somewhat unattractive in the past. However, a number of researchers began to investigate the use of Krylov subspace methods as an efficient solution to this problem. These are commonly used in conjunction with implicit time integration methods, where we typically have to approximate expressions of the form $(I - \tau M)^{-1}v$, for a given matrix M and scalar $\tau > 0$. Results in the literature have shown, in fact, that exponential methods have an advantage over the traditional implicit approach, due to faster Krylov convergence of exponential functions,

*Corresponding author.
email: Colm.Clancy@ec.gc.ca

compared with the rational functions arising from the implicit approach (Hochbruck and Lubich, 1997; Loffeld and Tokman, 2013). This has provided strong motivation for the further development of these methods; *c.f.* Hochbruck et al. (1998), Tokman (2006) and the numerous references therein.

In this paper, we investigate the potential use of exponential integrators in NWP models. A preliminary testing with a shallow water model was outlined in Bonaventura (2010). Here we provide a comprehensive examination of the performance of exponential schemes in the shallow water framework by comparing them to alternative time integrators, both explicit and semi-implicit.

In the next section, we provide a general introduction to the theory of exponential integration methods, along with some details of specific schemes and the Krylov methods commonly used. In particular, we discuss the adaptive Krylov subspace algorithm of Niesen and Wright (2012), which we will use for the efficient handling of matrix exponentials. In Section 3, we introduce the shallow water model and describe the implementation of the exponential methods. Numerical results are presented in Section 4 followed by discussion and conclusions in Section 5.

2. Exponential integration methods

2.1. General theory

We consider the following system which may arise after the spatial discretisation of a system of PDEs:

$$\frac{dU}{dt} = F(U) \quad (1)$$

where $U \in \mathbf{R}^N$ is the state vector and N denotes the number of degrees of freedom. There are two main approaches. Firstly, we assume that we can split the right-hand side forcing into linear and nonlinear parts; that is, we rewrite the above as:

$$\frac{dU}{dt} = LU + \mathcal{N}(U) \quad (2)$$

We wish to step the solution from $U_n = U(n\Delta t)$ to U_{n+1} . By multiplying through by an integrating factor e^{-Lt} we arrive at

$$U_{n+1} = e^{L\Delta t} U_n + \int_0^{\Delta t} e^{L(\Delta t-\tau)} \mathcal{N}(U(n\Delta t + \tau)) d\tau \quad (3)$$

A number of exponential-based schemes using this idea of separating the linear terms were considered in Beylkin et al. (1998) and Cox and Matthews (2002).

An alternative to this separation is described in detail in Tokman (2006). The forcing in eq. (1) is expanded about time $t_n = n\Delta t$ to give

$$\frac{dU(t)}{dt} = F(U_n) + A_n(U(t) - U_n) + R(U(t)) \quad (4)$$

where A_n is the Jacobi matrix $\frac{\partial F}{\partial U}$ evaluated at t_n and $R(U(t))$ represents the remainder terms:

$$R(U(t)) = F(U(t)) - F(U_n) - A_n(U(t) - U_n) \quad (5)$$

After multiplying eq. (4) by the integrating factor $e^{-A_n t}$, we obtain

$$U_{n+1} = U_n + (e^{A_n \Delta t} - I) A_n^{-1} F_n + \int_0^{\Delta t} e^{A_n(\Delta t-\tau)} R(U(n\Delta t + \tau)) d\tau \quad (6)$$

where I is the identity matrix.

Note that eqs. (3) and (6) represent exact forms of the solution of their respective ODE systems. The method of evaluation of the integrals determines a particular numerical scheme.

As will be seen from the details of specific schemes in the coming sections, a particular family of functions arise in the development of exponential integrators. These are known as the φ -functions and may be defined for a scalar z as

$$\begin{aligned} \varphi_0(z) &= e^z \\ \varphi_\ell(z) &= \frac{1}{(\ell-1)!} \int_0^1 e^{(1-\theta)z} \theta^{\ell-1} d\theta, \quad \ell = 1, 2, \dots \end{aligned}$$

and the functions satisfy the recurrence relation

$$\varphi_\ell(z) = z \varphi_{\ell+1}(z) + \frac{1}{\ell!} \quad (7)$$

With matrix arguments, the first three φ -functions are given by $\varphi_0(A) = e^A$, $\varphi_1(A) = (e^A - I)A^{-1}$, and $\varphi_2(A) = (e^A - I - A)(A^{-1})^2$.

2.2. Specific schemes

By using a polynomial approximation for the integral in eq. (3), we can recover multistep numerical schemes of arbitrary order. Alternatively, we can use a multistage Runge–Kutta type approach. Many examples may be found in the literature; for example, Beylkin et al. (1998), Cox and Matthews (2002), Minchev and Wright (2005).

In relation to atmospheric applications, we note in passing the ‘slave-frog’ scheme of Beylkin et al. (1998), which is derived in a similar manner from eq. (2) over three time-levels and may be written as:

$$U_{n+1} = e^{L2\Delta t} U_{n-1} + (e^{L2\Delta t} - I) L^{-1} \mathcal{N}_n \quad (8)$$

In a shallow water model based on a normal mode expansion, Daley (1980) used this method for the integration of the fast modes of the system, pointing out that this exact treatment of the linear terms does not result in a slowing of high-frequency gravity waves, unlike a semi-implicit treatment. In this case, of course, the system is diagonalised through the normal mode decomposition and thus the use of exponentials is not a computational concern.

Exact treatment of the linear terms with the resulting phase accuracy may also be achieved using a Laplace transform approach, whereby the system is solved in transform space and the solution then inverted to physical space. Clancy and Lynch (2011) describe such an algorithm in which the inversion is carried out using numerical quadrature. The scheme in eq. (8) is an analogous form of this, where inversion is exact.

In our current work, we focus on methods which use the Jacobian of the system, rather than splitting the forcing into linear and nonlinear parts, that is, approach (6) instead of eq. (3). A number of such schemes are evaluated in Tokman (2006); we consider two of these in particular. The first is the second-order scheme denoted EPI2:

$$U_{n+1} = U_n + (e^{A_n \Delta t} - I) A_n^{-1} F_n \quad (9)$$

This method also commonly appears in the literature as the ‘exponentially fitted Euler’ method, for example, Hochbruck et al. (1998), Newman (2003).

The second scheme we choose is the third-order EPI3:

$$U_{n+1} = U_n + (e^{A_n \Delta t} - I) A_n^{-1} F_n + \frac{2}{3} \Delta t (e^{A_n \Delta t} - I - A_n \Delta t) (A_n \Delta t)^{-2} R_{n-1} \quad (10)$$

where $R_{n-1} = F_{n-1} - F_n - A_n(U_{n-1} - U_n)$.

The motivation for choosing these two particular schemes will be discussed in the coming sections.

2.3. Krylov subspace approximation

For a given matrix $A \in \mathbf{R}^{N \times N}$ and vector $b \in \mathbf{R}^N$, the m -dimensional Krylov subspace is given by

$$K_m(A, b) = \text{span}\{b, Ab, A^2b, \dots, A^{m-1}b\} \quad (11)$$

Krylov subspace methods are widely used in the numerical solution of linear systems; see, for example, Chapters 6 and 7 of Saad (2003) for more details. As mentioned in the introduction, they have also been used to great effect with exponential integrators. All of these discretisations involve functions of matrix exponentials, for example, eq. (10). We do not need these explicitly, rather their action on a given vector. We can therefore seek approximations in a Krylov subspace, whose dimension is $m \ll N$. This approach to

approximation of functions of matrices was first proposed by Van der Vorst (1987).

An Arnoldi algorithm may be used to construct an orthonormal basis $\{v_1, v_2, \dots, v_m\}$ for the Krylov subspace $K_m(A, b)$. The algorithm also produces an upper Hessenberg matrix $H_m = V_m^T A V_m$, where V_m is an $N \times m$ matrix whose columns are the basis vectors (Saad, 2003). For a given function of a matrix $f(A)$, we may approximate the action on the given vector, $f(A)b$, in $K_m(A, b)$ as

$$f(A)b \approx \|b\|_2 V_m f(H_m) e_1$$

where e_1 is the first standard unit vector and $\|\cdot\|_2$ denotes the vector 2-norm. Further details may be found in Tokman (2006). Since H_m has size $m \times m$, the key to successful exponential integration is to be able to use m sufficiently small so that $f(H_m)$ may be efficiently calculated using traditional methods, such as Padé approximation for exponential functions (Moler and Van Loan, 2003).

2.4. The phipm algorithm

In terms of the φ -functions introduced earlier, we may rewrite the schemes (9) and (10) as

$$U_{n+1} = U_n + \varphi_1(\Delta t A_n) \Delta t F_n \quad (12)$$

and

$$U_{n+1} = U_n + \varphi_1(\Delta t A_n) \Delta t F_n + \frac{2}{3} \varphi_2(\Delta t A_n) \Delta t R_{n-1} \quad (13)$$

The two-stage, third-order EPIRK3 scheme of Tokman (2006) may be written

$$r_1 = U_n + \varphi_1\left(\frac{1}{2} \Delta t A_n\right) \Delta t F_n \quad (14)$$

$$U_{n+1} = U_n + \varphi_1(\Delta t A_n) \Delta t F_n + \frac{2}{3} \varphi_2(\Delta t A_n) \Delta t R(r_1)$$

All of the methods presented require the evaluation of expressions of the form

$$\varphi_0(A)b_0 + \varphi_1(A)b_1 + \dots + \varphi_p(A)b_p \quad (15)$$

for some set of vectors $b_k \in \mathbf{R}^N$, $k = 0, 1, \dots, p$. Niesen and Wright (2012) have developed the **phipm** algorithm to efficiently carry out this task. This is achieved by using an adaptive time-stepping on the non-autonomous linear ODE

$$u'(t) = Au(t) + b_1 + tb_2 + \dots + \frac{t^{p-1}}{(p-1)!} b_p, \quad u(0) = b_0 \quad (16)$$

which has solution

$$u(t) = \varphi_0(tA)b_0 + t\varphi_1(tA)b_1 + \dots + t^p\varphi_p(tA)b_p \quad (17)$$

By taking $t=1$, we recover the expression (15). The algorithm steps through this pseudo-time interval. Due to the recurrence relations between the φ -functions, at each step just one exponential-type expression of the form $\varphi_p(\tau A)v$ is needed, where τ is the pseudo-time-step. This is done using a Krylov projection as outlined in the previous section. The full description of the algorithm is given in the paper of Niesen and Wright (2012).

For a specified accuracy, the efficiency is determined both by the size of the Krylov subspaces used and the number of steps taken. When approximating $\varphi_p(\tau A)v$ with $\tau < 1$, the smaller τ is, the fewer Krylov basis vectors will be needed from the Arnoldi iteration. However, this will then require more steps over the interval. The algorithm is fully adaptive and determines, after each step, which is the more cost-effective: more steps or larger subspace dimension m . Tokman et al. (2012) use the **phimp** algorithm in conjunction with Runge–Kutta exponential integrators and report a definite computational advantage to this adaptive approach.

We consider the use of **phimp** with the EPI2 and EPI3 methods given by eqs. (12) and (13), respectively. Taking eq. (13), for example, we see from eq. (15) that the second and third terms on the right-hand side may be evaluated by one call to the algorithm, with b_0 the zero vector, $b_1 = \Delta t F_n$, $b_2 = \frac{2}{3} \Delta t R_{n-1}$ and $A = \Delta t A_n$. In contrast, the two-stage EPIRK3 scheme in eq. (14) requires two calls to **phimp** at each time-step; as pointed out in Tokman et al. (2012), the number of stages in the scheme equals the number of required executions of the algorithm.

Tokman (2006) discusses the benefits of Runge–Kutta-type schemes, such as EPIRK3, when used in the context of adaptive time-stepping. In our present work, we are interested in methods with a constant time-step for atmospheric models. For this reason, we concentrate here on the single-stage EPI2 and EPI3 methods. As will be seen in Section 4, the main computational cost at each time-step is the call to the **phimp** subroutine, and so we wish to minimise the number of required calls overall.

3. Shallow water implementation

The shallow water equations, describing a shallow rotating layer of incompressible and homogeneous fluid, are a commonly used testbed for numerical methods for atmospheric models. Here we use the model developed in Pudykiewicz (2011). This uses a finite volume discretisation of the governing equations on an icosahedral grid. The time integration was carried out using an explicit Runge–Kutta fourth-order scheme (RK4). Clancy and Pudykiewicz (2013) subsequently proposed a number of semi-implicit predictor–corrector time schemes and tested these against the RK4 using a number of standard experiments. These

methods used the Krylov-based generalised conjugate residual method (Eisenstat et al., 1983) to solve the linear systems arising from the semi-implicit formulation and were found to allow considerably longer time-steps while retaining sufficient levels of accuracy. We now seek to test the utility of the EPI2 and EPI3 schemes within this experimental framework.

The shallow water model consists of four prognostic variables: the three Cartesian wind components and the depth of the fluid layer. After spatial discretisation, we have a system of ODEs which can be written in the form (1). Details are given in the appendix. The state vector, U , contains the averages of the prognostic variables in each of the control volumes on the icosahedral grid. Thus, for a grid with N vertices, we have a system with a state vector of length $4N$ and a $4N \times 4N$ sparse Jacobi matrix.

To solve the ODE system using the EPI2 or EPI3 schemes given by eqs. (12) and (13), respectively, we proceed as follows. At any given time-level n , representing time $t = n\Delta t$, $n = 0, 1, \dots$, we first compute the forcing terms at this time, F_n , along with the Jacobi matrix A_n . An expression for this is given in the appendix. Note that EPI3 uses three time-levels, with the older $(n-1)$ information needed in the R_{n-1} term. To start the integration, we take a single EPI2 step to get U_1 from U_0 .

Each scheme then requires a single call to the **phimp** subroutine. As discussed in the previous section, the input vectors b_k from eq. (15) are $[b_0; b_1] = [0; \Delta t F_n]$ for EPI2 and $[b_0; b_1; b_2] = [0; \Delta t F_n; \frac{2}{3} \Delta t R_{n-1}]$ for EPI3. The vector (15) calculated by **phimp** is then used to update the solution from U_n to U_{n+1} .

The **phimp** subroutine has a number of additional input arguments. As mentioned earlier, we set the parameter t in eq. (15) to 1. In both EPI2 and EPI3, the matrix used in the routine is $\Delta t A_n$; in our case this is specified as non-symmetric (the logical flag *symm* allows for efficient Krylov projection in the symmetric case).

The first guess of the Krylov subspace dimension is given by the argument m . The adaptive algorithm subsequently adjusts the dimension. The remaining input is the tolerance parameter *tol*. This appears in eqs. (17) and (18) of Niesen and Wright (2012) and relates to the initial choice of step size in the algorithm and the adaptive procedure for this.

4. Numerical tests

We now test the exponential schemes in the icosahedral model using a number of standard test cases from the literature. In all of the following, we use the icosahedral grid number 6, corresponding to six successive divisions of the faces of the original icosahedron. This grid has $N = 40,962$ vertices with an average spatial resolution of 112 km (Pudykiewicz, 2011).

In order to analyse the simulation results, we will use the standard normalised ℓ_2 and ℓ_∞ height error measures, as defined in Williamson et al. (1992):

$$\ell_2(h) = \sqrt{\frac{\mathcal{I}\{(h - h_a)^2\}}{\mathcal{I}\{h_a^2\}}}, \quad \ell_\infty(h) = \frac{\max|h - h_a|}{\max|h_a|} \quad (18)$$

where h_a is the analytic or reference height field and $\mathcal{I}\{\}$ denotes integration of the field over the whole spherical domain.

Before tests with the full shallow water system, we will first consider the advection of a scalar field by a prescribed velocity field in Section 4.1 below. This involves just the continuity equation. As we will see, it is a special case for exponential methods and is therefore considered separately. In Section 4.2, we return to the full system and discuss the various model configurations for the experiments. A description of the shallow water test cases and the simulation results follow this.

4.1. Passive advection

In this case, we deal with just one of the discretised equations from the shallow water set given in the appendix, namely the continuity equation:

$$\frac{d}{dt}\{h\} = -\mathbf{D}_x\{u_x h\} - \mathbf{D}_y\{u_y h\} - \mathbf{D}_z\{u_z h\} \quad (19)$$

We note that with the winds constant (or alternatively prescribed at each time-step), the above can be written as a linear system

$$\frac{d}{dt}\{h\} = M\{h\}$$

where the (i, j) -th component of the matrix M is

$$M^{ij} = -D_x^j u_x^i - D_y^j u_y^i - D_z^j u_z^i$$

No summation convention is used in the above definition.

As noted in previous sections, exponential methods are A-stable and exact for linear constant-coefficient systems; the Jacobi matrix is simply M . The only limits in terms of accuracy and stability will be from the computation of the exponential $e^{M\Delta t}$.

We test this with the deformational flow case of Nair and Machenhauer (2002), with the same parameter settings as in Pudykiewicz (2011). This consists of the advection of the scalar field by two static vortices centred at the poles. In Fig. 1, the analytic solutions (looking down on the north pole) after 6, 20 and 60 simulated days are plotted in panels (a), (c) and (e), respectively. Table 1 shows the normalised ℓ_∞ errors for a number of simulations. First is the RK4 scheme with a time-step of 2160s, also used by Pudykiewicz (2011). The scheme is stable for Courant numbers of

around 2.8 (Durran, 1999) and the simulation may in fact be carried out with as much as a 4-h time-step. The results, however, are indistinguishable from those even with a much lower time-step, and so the solution shown is dominated by spatial truncation errors.

The other results in Table 1 are for the EPI2 and EPI3 schemes with a time-step of 12 h. The results are essentially identical to those with the RK4. In Fig. 1b,d,f, we show the solutions for the EPI2 run, which very closely match the analytic solutions.

While this linear system is a trivial case for the exponential methods, it is still noteworthy that they converge so quickly even with such unrealistically long time-steps. The **phipm** routine at each step requires fewer than 10 Krylov basis vectors, on average. Furthermore we see no obvious errors originating from the computation of the matrix exponentials.

The method is also mass conserving down to the round-off error level (Qaddouri et al., 2012). Combination of this property with the stability of the algorithm for very large Courant numbers is particularly important for simulation of cloud variables and atmospheric transport of chemically reactive species. In the past, these fields were advected in the NWP and climate models using semi-Lagrangian advection which suffers from the inherent mass conservation errors. The use of exponential integrators for advecting the crucial microphysical fields will be particularly important to increase the efficiency and accuracy of many aspects of atmospheric chemistry.

4.2. Model configurations

We now return to the full shallow water system. As mentioned, the icosahedral grid number 6 has $N=40,962$ vertices. The state vector of the system, comprised of the four prognostic variables, is therefore of size 163 848. The construction of the sparse Jacobi matrix, of size $(163,848 \times 163,848)$, is discussed in the appendix.

We compare the integrations with EPI2 and EPI3 against those with two reference time schemes. The first is the explicit RK4 mentioned in the previous section. For this, we use a time-step of $\Delta t = 240$ s, as in Pudykiewicz (2011).

In addition, we consider the semi-implicit T-ABT scheme of Clancy and Pudykiewicz (2013). Briefly, this uses a split form of the equations, as in eq. (2). The non-linear terms are discretised with an Adams-Bashforth-trapezoidal predictor-corrector method (Kar, 2012), while a trapezoidal averaging is applied to the linear terms at each stage. With the system written in the form (2), the scheme is given by

$$\frac{U_* - U_n}{\Delta t} = \frac{1}{2}[LU_* + LU_n] + \frac{1}{2}[3\mathcal{N}(U_n) - \mathcal{N}(U_{n-1})] \quad (20)$$




Fig. 1. Deformational flow case: exact solution at (a) day 6, (c) day 20 and (e) day 60. The corresponding numerical solutions with EPI2 with a 12-h time-step are shown in (b), (d) and (f).

Table 1. Deformational flow case: normalised ℓ_∞ errors

	Day 6	Day 20	Day 60
RK4 $\Delta t = 2160$ s	1.386977×10^{-3}	3.475878×10^{-3}	1.894418×10^{-1}
EPI2 $\Delta t = 12$ h	1.386983×10^{-3}	3.476770×10^{-3}	1.894432×10^{-1}
EPI3 $\Delta t = 12$ h	1.387031×10^{-3}	3.477584×10^{-3}	1.894430×10^{-1}

$$\frac{U_{n+1} - U_n}{\Delta t} = \frac{1}{2} [LU_{n+1} + LU_n] + \frac{1}{2} [\mathcal{N}(U_*) + \mathcal{N}(U_n)] \quad (21)$$

Details of the linear splitting are given in the appendix. With this T-ABT scheme, we take $\Delta t = 900$ s, to balance accuracy and efficiency concerns (Clancy and Pudykiewicz, 2013).

For the exponential integration schemes, we select time-steps of 3600 and 7200s. These are considerably longer than those for the two reference methods. As mentioned in the introduction, the exponential schemes can be shown to be trivially A-stable, since they are exact for linear equations. The choice of time-step is then governed by accuracy and also by the convergence of the Krylov projections. Tests showed that taking $\Delta t = 10800$ s was too long; the algorithm was slow to converge and the results were noticeably less accurate. For the arguments of **phipm** already discussed, we use $m = 30$ and $tol = 1 \times 10^{-4}$. Reducing tol to 1×10^{-5} was found to have negligible effect on the solutions.

We now discuss the results from four test cases, in terms of accuracy and efficiency. We note that, unless specified, no explicit dissipation is added by default in these experiments.

4.3. Shallow water test cases

We first consider one of the time-dependent flow cases of Lauter et al. (2005). This unsteady case is a useful test because it has a known analytic solution, allowing us to measure the numerical accuracy. In terms of the zonal (u) and meridional (v) winds, the height of the fluid layer (h) and the height of the orography (h_s), the flow is given at time t as follows:

$$\begin{aligned} u(\lambda, \phi, t) &= u_0 (\sin \theta \sin \phi (\cos \lambda \cos \Omega t - \sin \lambda \sin \Omega t) \\ &\quad + \cos \theta \cos \phi) \\ v(\lambda, \phi, t) &= -u_0 \sin \theta (\sin \lambda \cos \Omega t + \cos \lambda \sin \Omega t) \\ gh(\lambda, \phi, t) &= -\frac{1}{2} [u_0 (\sin \theta \cos \phi (-\cos \lambda \cos \Omega t + \sin \lambda \sin \Omega t) \\ &\quad + \cos \theta \sin \phi) + a\Omega \sin \phi]^2 + k_1 \\ gh_s(\lambda, \phi, t) &= \frac{1}{2} (a\Omega \sin \phi)^2 + k_2 \end{aligned}$$

where λ and ϕ are longitude and latitude, respectively, a is the Earth's radius and Ω is the Earth's angular velocity.

As in Pudykiewicz (2011), we use the parameter values $u_0 = 2\pi a/12\text{m/d}$, $k_1 = 133681\text{m}^2\text{ s}^{-2}$, $k_2 = 10\text{ m}^2\text{s}^{-2}$ and $\theta = \pi/4$.

The next two cases are from the standard suite of Williamson et al. (1992). We use case 5, consisting of an initially balanced zonal flow interacting with an isolated mountain, and the Rossby-Haurwitz wave case 6. Neither of these cases has a known analytic solution. In order to compute error measures for comparison, we use a high-resolution reference solution obtained from the icosahedral model at grid 7 (163 842 vertices, with an average grid resolution of 56 km) with the RK4 scheme at $\Delta t = 90$ s. This reference was also used in Clancy and Pudykiewicz (2013).

The final case to be considered is the barotropic instability case described in Galewsky et al. (2004). The initial conditions consist of a mid-latitude jet which is destabilised by the addition of a small height perturbation, resulting in a more complex, multiscale flow when compared with other test cases.

4.4. Accuracy

In Fig. 2, we plot the errors for a 10-d simulation of the Lauter et al. (2005) unsteady case. Here, for interest, we also include forecasts with the exponential methods at $\Delta t = 900$ s, matching the semi-implicit T-ABT.

The explicit RK4 with the lowest time-step is the most accurate. However, both exponential methods show very similar accuracy with the much longer $\Delta t = 900$ s, despite both having lower order of accuracy than the RK4. Increasing the time-step by a factor of 4 to 1 h, the EPI2 and EPI3 are both still considerably more accurate than the T-ABT at 900s. Even at a 2-h time-step, they generally outperform the T-ABT, although the EPI3 begins to lose accuracy towards the end of the simulation (in fact, this particular simulation begins to destabilise after day 10).

In Figs. 3 and 4, we show the results for the Williamson et al. (1992) mountain and Rossby-Haurwitz wave cases, respectively. Both of these sets of experiments show similar behaviour. The exponential methods with time-steps of 1 or 2 h are as accurate as the explicit RK4 with only $\Delta t = 240$ s. They are more accurate than the T-ABT forecasts, although the difference is less in the Rossby-Haurwitz case, particularly at the later stages (Fig. 4).




Fig. 2. Normalised ℓ_2 (left) and ℓ_∞ (right) error measures for the height field in the unsteady flow case. Exponential methods EPI2 (top) and EPI3 (bottom) using various time-steps are compared against the reference RK4 and T-ABT. Note the logarithmic scale on the vertical axis.

Finally we consider the barotropic instability case. In Fig. 5, we plot the relative vorticity fields after 6 d of integration. The performance of the icosahedral model with the RK4 scheme in this case has been compared with simulations from other models (Qaddouri et al., 2012).

From the top of Fig. 5, we have the RK4 followed by the T-ABT, both giving comparable results. The bottom two panels are for the EPI2 and EPI3, respectively, with a time-step of 1 h. Looking in particular at the region between $\lambda = \pi/2$ and $\lambda = \pi$, the T-ABT is, perhaps, the smoothest; each of the other schemes produces some more noise in this area. Given the length of the time-steps for the two exponential schemes, it is impressive that they can compete with the RK4 in this quite challenging test case, even without the addition of explicit diffusion. In Fig. 6, we plot similar fields for the exponential methods using a 2-h time-step. For these values, however, we see that there is a noticeable distortion of the solutions when compared with Fig. 5.

In general, some form of artificial dissipation is needed in simulations with atmospheric models to control the smallest grid-scales, even in the case of fully monotonic, non-oscillatory schemes; for a further discussion see Jablonowski and Williamson (2011). We now add the same fourth-order diffusion described in Pudykiewicz (2011); *c.f.* Fig. 21 in that paper. In Fig. 7, we again plot the relative vorticity fields at day 6 for the RK4 at $\Delta t = 240$ s along with the exponential schemes with a 2-h time-step. In contrast to the results in Figs. 5 and 6, there is now close agreement between the models.

This test case also allows us to examine the effect of a discretisation method on the phase speed of fast-moving waves. Galewsky et al. (2004) examine the instantaneous height and divergence fields (their Fig. 2) during the early hours of the simulation. At this adjustment stage, the flow evolves with the propagation of gravity waves.




Fig. 3. Normalised ℓ_2 (left) and ℓ_∞ (right) error measures for the height field in the Williamson et al. (1992) mountain case 5. Exponential methods EPI2 (top) and EPI3 (bottom) using various time-steps are compared against the reference RK4 and T-ABT. Note the logarithmic scale on the vertical axis.

In Fig. 8, we plot the divergence fields after 12 h for the six simulations depicted in Figs. 5 and 6, that is, no dissipation included. We would expect a phase error with the semi-implicit T-ABT, due to the use of trapezoidal averaging (Durran, 1999). Indeed, all of the other models show comparable performance, whereas for the T-ABT (top right panel) there is a noticeable slowing and distortion of the gravity wave. This fact is a key motivation for work with exponential integrators.

4.5. Efficiency

In the previous section, we saw how exponential integration can offer a high level of accuracy at rather long time-steps. We now consider issues of comparative efficiency between the various methods. We expect that the Krylov projections involved in the exponential methods will add

substantially to the computational cost per time-step. Ideally, this would be compensated by fewer steps.

The model is implemented in MATLAB and run on a desktop computer with an AMD Athlon 64 X2 Dual Core 4800+ processor. In Table 2, we list some timings for the first three of the test cases used. The timings given, in minutes, are the average execution time per simulated day of the experiment. The semi-implicit T-ABT simulations are consistently the fastest. The exponential methods are comparable with the RK4, particularly with the 2-h time-step.

In the tests of comparative performance, Loffeld and Tokman (2013) demonstrate that exponential methods can be superior to traditional implicit schemes. In our case, however, the discretisation of T-ABT reduces to a linear system in the height field only (details are given in Appendix A.2). The exponential methods, on the other




Fig. 4. Normalised ℓ_2 (left) and ℓ_∞ (right) error measures for the height field in the Williamson et al. (1992) Rossby–Haurwitz wave case 6. Exponential methods EPI2 (top) and EPI3 (bottom) using various time-steps are compared against the reference RK4 and T–ABT. Note the logarithmic scale on the vertical axis.

hand, work with the entire system of four prognostic variables. The difference in execution times is therefore to be expected.

A partial breakdown of the cost of the exponential methods is given in Table 3 for the 14-d Rossby–Haurwitz wave simulation. This shows the average time, in seconds, taken by three different components of the discretisations at each step: the computation of the Jacobi matrix, the calculation of the forcing terms and the execution time of the **phipm** routine. As expected, the **phipm** routine is the dominant cost. The final column shows the average number of Krylov subspace basis vectors required at each step. Looking at Tables 2 and 3 together, we see why doubling the time-step for EPI2 and EPI3 does not significantly reduce the overall execution time. With the longer time-step, a larger Krylov subspace is necessary to maintain the same level of accuracy, and this results in the longer timings for **phipm**.

The tolerance parameter in **phipm**, *tol*, is another consideration when discussing the efficiency of the exponential integrators. The default value used by Niesen and Wright (2012) is 10^{-7} . This parameter determines the accuracy of the evaluation of the φ -functions by controlling the dimension of the Krylov subspace and the length of the internal pseudo-time-step. As mentioned earlier, we have used a value of $tol = 10^{-4}$ throughout this work, with time-steps of 1 and 2 h. In experiments with less accurate calculation of the φ -functions using a tolerance of 10^{-3} , it was possible to keep a stable integration even with the longer time-step, with only slight deterioration of the conservation properties. More complete analysis and experimentation with the balancing of stability, efficiency and accuracy is an ongoing effort.

Further gains in efficiency may also be achieved by exploring other recently described techniques for evaluating the solution of eq. (6). One promising direction is




Fig. 5. Relative vorticity field (northern hemisphere) at day 6 for the Galewsky et al. (2004) case. From top to bottom: RK4 $\Delta t = 240$ s, T-ABT $\Delta t = 900$ s, EPI2 $\Delta t = 3600$ s and EPI3 $\Delta t = 3600$ s.

discussed by Al-Mohy (2010), who presents an alternate algorithm for exponential integrating schemes which is shown to be more efficient than **phipm** in a number of

experiments. The examination of this new scheme will be the subject of forthcoming research in the context of NWP.




Fig. 6. Relative vorticity field (northern hemisphere) at day 6 for the Galewsky et al. (2004) case with $\Delta t = 7200$ s for EPI2 (top) and EPI3 (bottom).




Fig. 7. Relative vorticity field (northern hemisphere) at day 6 for the Galewsky et al. (2004) case with explicit dissipation added. From top: RK4 $\Delta t = 240s$, EPI2 $\Delta t = 7200s$ and EPI3 $\Delta t = 7200s$.

5. Discussion and conclusions

Two single-stage exponential time integration methods have been implemented and tested in a shallow water model. This represents the first step in a longer-term goal

of assessing the potential benefits of this approach for the time integration of more complete models of the atmosphere. The methods were compared with an explicit and semi-implicit time discretisation using a number of standard test cases.




Fig. 8. Divergence fields for the barotropic jet case after 12 h of simulation. Top: RK4 $\Delta t = 240s$ (left) and T-ABT $\Delta t = 900s$ (right). Middle: EPI2 (left) and EPI3 (right) with $\Delta t = 3600s$. Bottom: EPI2 (left) and EPI3 (right) with $\Delta t = 7200s$. Contour interval is $8 \times 10^{-7} s^{-1}$.

Table 2. Average execution time (minutes) per simulated day, for the cases of the unsteady flow (Läuter et al., 2005) and the mountain and Rossby–Haurwitz wave of Williamson et al. (1992)

	Unsteady flow	Mountain	RH wave
RK4 $\Delta t = 240$ s	2.6	2.6	2.4
T–ABT $\Delta t = 900$ s	0.8	0.9	0.9
EPI2 $\Delta t = 1$ h	3.2	3.3	3.6
EPI2 $\Delta t = 2$ h	2.7	2.3	2.8
EPI3 $\Delta t = 1$ h	3.2	3.3	3.4
EPI3 $\Delta t = 2$ h	2.7	2.3	2.8

The exponential methods were found to yield very accurate solutions, exceeding those of the semi-implicit method. The second-order EPI2 gave comparable results to the explicit fourth-order Runge–Kutta, even with dramatically longer time-steps. In spite of the caveat of Tokman (2006) that this method is not accurate enough for general nonlinear systems, due to its omission of the integral in eq. (6), it seems to be sufficient for simulating shallow water flow at least. Only in the challenging unstable jet case of Galewsky et al. (2004) do the two exponential methods struggle at the longer, 2-h time-step. However, once dissipation was incorporated into the simulation, the results again were consistent with those of the explicit scheme. In any practical applications with more complex models, some mechanism for dissipating smaller scales is necessary, and so this is not a major cause of concern.

The high accuracy of the exponential methods is to be somewhat expected, given their analytic treatment of the linearised part of the system. Indeed, in the case of the passive advection where we are dealing with a linear system with constant winds, the methods give an ‘exact’ solution with no formal time truncation errors and are subject only to the errors in the computation of the matrix exponentials.

The crucial issue to address is the efficiency of these methods, given the size of the systems involved and the potentially large computational overhead with the matrix

Table 3. Average time taken, in seconds, for the calculation of the Jacobi matrix, the forcing terms and the execution of the **phipm** routine at each step over the 14-d Rossby–Haurwitz wave simulation. The fourth column gives the average number of Krylov basis vectors per call to **phipm**

	Jacobi matrix	Forcing	phipm	Krylov vectors
EPI2 $\Delta t = 1$ h	2.3	0.1	6.3	34.6
EPI2 $\Delta t = 2$ h	2.2	0.1	11.8	59.5
EPI3 $\Delta t = 1$ h	2.2	0.1	5.9	33.1
EPI3 $\Delta t = 2$ h	2.2	0.1	11.6	58.3

exponentials. In this work, we used the **phipm** algorithm of Niesen and Wright (2012). This was found to give execution times roughly comparable with the existing RK4 method. On the other hand, the semi-implicit T–ABT was always faster. This is unsurprising since, with this scheme, we have to solve only a linear system (a discrete Helmholtz equation) for just one of the prognostic variables, as described in the Appendix A.2. The exponential methods, however, must deal with a system four times as large as a result of considering the full system of four prognostic variables. On the other hand, for the case of advection with prescribed winds, we have just one prognostic variable. With the high accuracy mentioned previously and the ability to take very long time-steps, exponential methods may prove to be of benefit in efficiently simulating the advection of various tracers and interacting species.

As discussed in the previous section, it may also be possible to improve on the efficiency of the exponential methods by experimenting with the tolerance parameter when computing the exponential functions. The resulting compromise in accuracy may be acceptable and should be measured and balanced against the performance of the semi-implicit schemes.

The next stage of this research is to investigate further the emerging exponential integrators and apply them to more complete atmospheric models. The increased stiffness and resulting time-step limitations of non-hydrostatic models, due to the presence of rapidly propagating acoustic waves, render purely explicit methods impractical. Exponential methods potentially offer a means of achieving high accuracy and stability, without the distortion of the fast waves typical of semi-implicit methods. It remains to be seen what level of efficiency can be maintained in the complete system.

6. Acknowledgements

The authors are grateful to Claude Girard for his helpful internal review and to the two anonymous reviewers whose comments were of great benefit. This research was partly funded by the Natural Sciences and Engineering Research Council of Canada.

7. Appendix

A. Details of the shallow water model

Here we provide details of the shallow water model relevant to the current work. A full description of the icosahedral grid and shallow water model may be found in Pudykiewicz (2006, 2011).

After the finite volume spatial discretisation, the shallow water equations are written as systems of ODEs as follows; *c.f.* eq. (63) of Pudykiewicz (2011).

$$\frac{d}{dt}\{u_x\} = -\mathbf{W}_x - \mathbf{GS}_x\{f\} \quad (\text{A1})$$

$$\frac{d}{dt}\{u_y\} = -\mathbf{W}_y - \mathbf{GS}_y\{f\}$$

$$\frac{d}{dt}\{u_z\} = -\mathbf{W}_z - \mathbf{GS}_z\{f\}$$

$$\frac{d}{dt}\{h\} = -\mathbf{D}_x\{u_x h\} - \mathbf{D}_y\{u_y h\} - \mathbf{D}_z\{u_z h\}$$

The prognostic variables are the Cartesian wind components, u_x , u_y and u_z , and the depth of the fluid layer, h . The $\{\}$ brackets represent column arrays consisting of the control volume averages of the variables. The $\mathbf{W}_x$, $\mathbf{W}_y$ and $\mathbf{W}_z$ arrays contain terms involving vorticity components and $f = g(h_s + h) + |\mathbf{u}|^2/2$, where h_s is the orography and $\mathbf{u}$ is the velocity vector. The sparse matrices ($\mathbf{GS}_x$, $\mathbf{GS}_y$, $\mathbf{GS}_z$) and ($\mathbf{D}_x$, $\mathbf{D}_y$, $\mathbf{D}_z$) are used to compute the Cartesian components of the gradient and divergence operators, respectively. The construction of the operators ensures that the wind is constrained to be tangent to the sphere. Full details are given in Pudykiewicz (2011).

We may also add explicit dissipation to the system in the form $\mathbf{Df}\{\psi\}$, for each prognostic variable ψ . The dissipation operator takes the form

$$\mathbf{Df} = -\nu \mathbf{L}^2 \quad (\text{A2})$$

where $\mathbf{L}$ is the sparse matrix representing the Laplace operator.

A.1. Jacobi matrix

When implementing the exponential integration schemes in Section 3, we write the system in terms of a single state vector consisting of the arrays of all the prognostic variables. We then also require the Jacobi matrix of the right-hand side forcing.

The system of eq. (A1), with the dissipation terms added, could be rewritten in the compact form

$$\frac{dU}{dt} = F(U) \quad (\text{A3})$$

where the state vector U contains components of the velocity and the height field

$$U = \{\{u_x^i\}, \{u_y^i\}, \{u_z^i\}, \{h^i\}\}^T$$

The lower index denotes Cartesian components whereas the upper index represents the control volumes. The elements of the right-hand side vector of eq. (A3)

$$F = \{\{F_x^i\}, \{F_y^i\}, \{F_z^i\}, \{H^i\}\}^T$$

are defined by the following relations:

$$F_x^i = -(V_x^{ij} u_x^j + V_y^{ij} u_y^j + V_z^{ij} u_z^j + \gamma^i) * (\mathcal{N}_y^i u_z^j - \mathcal{N}_z^i u_y^j) - GS_x^{ij} \left(\frac{1}{2} (u_x^2 + u_y^2 + u_z^2) + g(h^i + h_s^i) \right) + Df^{ij} u_x^j \quad (\text{A4})$$

$$F_y^i = -(V_x^{ij} u_x^j + V_y^{ij} u_y^j + V_z^{ij} u_z^j + \gamma^i) * (\mathcal{N}_z^i u_x^j - \mathcal{N}_x^i u_z^j) - GS_y^{ij} \left(\frac{1}{2} (u_x^2 + u_y^2 + u_z^2) + g(h^i + h_s^i) \right) + Df^{ij} u_y^j \quad (\text{A5})$$

$$F_z^i = -(V_x^{ij} u_x^j + V_y^{ij} u_y^j + V_z^{ij} u_z^j + \gamma^i) * (\mathcal{N}_x^i u_y^j - \mathcal{N}_y^i u_x^j) - GS_z^{ij} \left(\frac{1}{2} (u_x^2 + u_y^2 + u_z^2) + g(h^i + h_s^i) \right) + Df^{ij} u_z^j \quad (\text{A6})$$

$$H^i = -(D_x^{ij} h^j u_x^i + D_y^{ij} h^j u_y^i + D_z^{ij} h^j u_z^i) + Df^{ij} h^i \quad (\text{A7})$$

where GS_x^{ij} , GS_y^{ij} , GS_z^{ij} , V_x^{ij} , V_y^{ij} , V_z^{ij} , D_x^{ij} , D_y^{ij} , D_z^{ij} are the components of the sparse matrices used to calculate gradient, vorticity and divergence, $\mathcal{N}_x$, $\mathcal{N}_y$, $\mathcal{N}_z$ are column arrays with the x , y and z components of the vector normal to the surface of the sphere and the Coriolis parameter is given by γ . The symbol $*$ indicates that arrays are multiplied component-wise. The Einstein summation convention is applied in eqs. (A4)–(A7).

The Jacobi matrix of the system (A3)

$\mathcal{J} = \{\frac{\partial F}{\partial U^T}\}$ can be written in the explicit form

$$\mathcal{J} = \mathcal{J}_1 + \mathcal{J}_2 + \mathcal{J}_3 \quad (\text{A8})$$

where

$$\mathcal{J}_1 = - \begin{bmatrix} \{P_x^i V_x^{ij}\} & \{P_x^i V_y^{ij}\} & \{P_x^i V_z^{ij}\} & \{0\} \\ \{P_y^i V_x^{ij}\} & \{P_y^i V_y^{ij}\} & \{P_y^i V_z^{ij}\} & \{0\} \\ \{P_z^i V_x^{ij}\} & \{P_z^i V_y^{ij}\} & \{P_z^i V_z^{ij}\} & \{0\} \\ \{0\} & \{0\} & \{0\} & \{0\} \end{bmatrix}, \quad (\text{A9})$$

the arrays $\{P_x^i\}$, $\{P_y^i\}$ and $\{P_z^i\}$ contain the x , y and z components of the vector product of the velocity with the vector normal to the surface of the sphere, $\{0\}$ denotes the all-zero sparse matrix,

$$\mathcal{J}_2 = + \begin{bmatrix} \{Df^{ij}\} & \{z^i \mathcal{N}_y^j \delta^{ij}\} & -\{z^i \mathcal{N}_z^j \delta^{ij}\} & \{0\} \\ \{z^i \mathcal{N}_z^j \delta^{ij}\} & \{Df^{ij}\} & \{z^i \mathcal{N}_x^j \delta^{ij}\} & \{0\} \\ \{z^i \mathcal{N}_x^j \delta^{ij}\} & -\{z^i \mathcal{N}_y^j \delta^{ij}\} & \{Df^{ij}\} & \{0\} \\ \{0\} & \{0\} & \{0\} & \{Df^{ij}\} \end{bmatrix}, \quad (\text{A10})$$

and

$$\mathcal{J}_3 = - \begin{bmatrix} \{GS_x^{ij} u_x^j\} & \{GS_x^{ij} u_y^j\} & \{GS_x^{ij} u_z^j\} & \{g GS_x^{ij}\} \\ \{GS_y^{ij} u_x^j\} & \{GS_y^{ij} u_y^j\} & \{GS_y^{ij} u_z^j\} & \{g GS_y^{ij}\} \\ \{GS_z^{ij} u_x^j\} & \{GS_z^{ij} u_y^j\} & \{GS_z^{ij} u_z^j\} & \{g GS_z^{ij}\} \\ \{D_x^{ij} h^j\} & \{D_y^{ij} h^j\} & \{D_z^{ij} h^j\} & \{\mathcal{D}^{ij}\} \end{bmatrix} \quad (\text{A11})$$

with

$$\mathcal{D}^{ij} = D_x^{ij} u_x^j + D_y^{ij} u_y^j + D_z^{ij} u_z^j$$

The Kronecker delta is written above as δ^{ij} and the z^i components are defined in Pudykiewicz (2011). Please note that we do not use the summation convention in the bracketed terms in eqs. (A9)–(A11). Consequently, with this convention the expression $\{P_x^i V_x^{ij}\}$ denotes the matrix with ij -th element equal to $P_x^i V_x^{ij}$.

Due to the sparsity of this Jacobi matrix, we only need to store the relatively small number of non-zero elements. The product of the matrix and a vector can then be efficiently evaluated when required. An alternative to this approach would be to use a subroutine which evaluates this product without explicitly storing the constructed matrix.

A.2. Linear terms for semi-implicit methods

For the T-ABT integration scheme given in eq. (20), we apply a trapezoidal averaging to the linear terms governing the fast gravity waves in the shallow water system. As was done in Clancy and Pudykiewicz (2013), we define constant reference and perturbation depths such that $h = \bar{h} + h'$ and redefine $f = g h_s + |\mathbf{u}|^2/2$. We can then rewrite the system (A1) as

$$\frac{d}{dt} \{u_x\} = -g \mathbf{GS}_x \{h'\} + \mathbf{N}_x \quad (\text{A12})$$

$$\frac{d}{dt} \{u_y\} = -g \mathbf{GS}_y \{h'\} + \mathbf{N}_y$$

$$\frac{d}{dt} \{u_z\} = -g \mathbf{GS}_z \{h'\} + \mathbf{N}_z$$

$$\frac{d}{dt} \{h'\} = -\bar{h} [\mathbf{D}_x \{u_x\} + \mathbf{D}_y \{u_y\} + \mathbf{D}_z \{u_z\}] + \mathbf{N}_h$$

where the $\mathbf{N}_x$, $\mathbf{N}_y$, $\mathbf{N}_z$, and $\mathbf{N}_h$ represent those terms remaining from the forcing in eq. (A1). With the system now written in the form of eq. (2), we can apply the discretisation of eq. (20). After some substitution, we arrive at a discrete Helmholtz equation for $\{h'\}$ at the new time-level. Full details may be found in Clancy and Pudykiewicz (2013).

References

- Al-Mohy, A. H. 2010. *Algorithms for the matrix exponential and its Fréchet derivative*. PhD Thesis. Faculty of Engineering and Physical Sciences, University of Manchester, Manchester, UK, 155 pp.
- Beylkin, G., Keiser, J. M. and Vozovoi, L. 1998. A new class of time discretization schemes for the solution of nonlinear PDEs. *J. Comput. Phys.* **147**, 362–387.
- Bonaventura, L. 2010. Exponential integrators for multiple time scale problems of environmental fluid dynamics. *Innovative Integrators Workshop 27–30 October 2010*. Universität Innsbruck, Austria.
- Clancy, C. and Lynch, P. 2011. Laplace transform integration of the shallow-water equations. Part I: Eulerian formulation and Kelvin waves. *Q. J. Roy. Meteorol. Soc.* **137**, 792–799.
- Clancy, C. and Pudykiewicz, J. A. 2013. A class of semi-implicit predictor–corrector schemes for the time integration of atmospheric models. *J. Comput. Phys.* **250**, 665–684.
- Cox, S. M. and Matthews, P. C. 2002. Exponential time differencing for stiff systems. *J. Comput. Phys.* **176**, 430–455.
- Daley, R. 1980. The development of efficient time integration schemes using model normal modes. *Mon. Weather Rev.* **108**, 100–110.
- Durrant, D. R. 1999. *Numerical Methods for Wave Equations in Geophysical Fluid Dynamics*. Springer, New York.
- Eisenstat, S. C., Elman, H. C. and Schultz, M. H. 1983. Variational iterative methods for nonsymmetric systems of linear equations. *SIAM J. Numer. Anal.* **20**, 345–357.
- Galewsky, J., Scott, R. K. and Polvani, L. M. 2004. An initial-value problem for testing numerical models of the global shallow-water equations. *Tellus A*. **56**, 429–440.
- Hochbruck, M. and Lubich, C. 1997. On Krylov subspace approximations to the matrix exponential operator. *SIAM J. Numer. Anal.* **34**, 1911–1925.
- Hochbruck, M., Lubich, C. and Selhofer, H. 1998. Exponential integrators for large systems of differential equations. *SIAM J. Sci. Comput.* **19**, 1552–1574.
- Jablonski, J. and Williamson, D. L. 2011. The pros and cons of diffusion, filters and fixers in atmospheric general circulation models. In: *Numerical Techniques for Global Atmospheric Models* (eds P. H. Lauritzen, C. Jablonowski, Taylor M. A. and R. D. Nair), Springer–Verlag, Berlin Heidelberg, 381–493.
- Kar, S. K. 2012. An explicit time-difference scheme with an Adams–Bashforth predictor and a trapezoidal corrector. *Mon. Weather Rev.* **140**, 307–322.
- Läuter, M., Handorf, D. and Dethloff, K. 2005. Unsteady analytical solutions of the spherical shallow water equations. *J. Comput. Phys.* **210**, 535–553.
- Löffeld, J. and Tokman, M. 2013. Comparative performance of exponential, implicit, and explicit integrators for stiff systems of ODEs. *J. Comput. Appl. Math.* **241**, 45–67. DOI: 10.1016/j.cam.2012.09.038.
- Minchev, B. V. and Wright, W. M. 2005. *A Review of Exponential Integrators for First Order Semi-Linear Problems*. Technical Report 2/2005, Department of Mathematical Sciences, Norwegian

- University of Science and Technology, N-7491 Trondheim, Norway.
- Moler, C. and Van Loan, C. 2003. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Rev.* **45**, 3–49.
- Nair, R. D. and Machenhauer, B. 2002. The mass-conservative cell-integrated semi-Lagrangian advection scheme on the sphere. *Mon. Weather Rev.* **130**, 649–667.
- Newman, C. K. 2003. *Exponential Integrators for the Incompressible Navier–Stokes Equations*. PhD Thesis. Virginia Polytechnic Institute and State University, Blacksburg, VA, 101 pp.
- Niesen, J. and Wright, W. M. 2012. Algorithm 919: a Krylov subspace algorithm for evaluating the φ -functions appearing in exponential integrators. *ACM Trans. Math. Software* **38**(3), 1–19. DOI: 10.1145/2168773.2168781.
- Pudykiewicz, J. A. 2006. Numerical solution of the reaction-advection-diffusion equation on the sphere. *J. Comput. Phys.* **213**, 358–390.
- Pudykiewicz, J.A. 2011. On numerical solution of the shallow water equations with chemical reactions on icosahedral geodesic grid. *J. Comput. Phys.* **230**, 1956–1991.
- Qaddouri, A., Pudykiewicz, J., Tanguay, M., Girard, C. and Côté, J. 2012. Experiments with different discretizations for the shallow-water equations on a sphere. *Q. J. Roy. Meteorol. Soc.* **138**, 989–1003.
- Saad, Y. 2003. *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia.
- Tokman, M. 2006. Efficient integration of large stiff systems of ODEs with exponential propagation iterative (EPI) methods. *J. Comput. Phys.* **213**, 748–776.
- Tokman, M., Loffeld, J. and Tranquilli, P. 2012. New adaptive exponential propagation iterative methods of Runge–Kutta type. *SIAM J. Sci. Comput.* **34**, A2650–A2669.
- Van der Vorst, H. A. 1987. An iterative solution method for solving $f(a)x=b$ using Krylov subspace information obtained for the symmetric positive definite matrix a . *J. Comput. Appl. Math.* **18**, 249–263.
- Williamson, D. L., Drake, J. B., Hack, J. J., Jakob, R. and Swarztrauber, P. N. 1992. A standard test set for numerical approximation to the Shallow Water equations in spherical geometry. *J. Comp. Phys.* **102**, 211–224.