

# Performance assessment of exponential Rosenbrock methods for large systems of ODE

E. J. Carr<sup>1</sup>T. J. Moroney<sup>2</sup>I. W. Turner<sup>3</sup>

(Received 1 November 2012; revised 27 March 2013)

## Abstract

This article studies time integration methods for stiff systems of ordinary differential equations of large dimension. For such problems, implicit methods generally outperform explicit methods because the step size is usually less restricted by stability constraints. Recently, however, a family of explicit methods, called exponential integrators, have become popular for large stiff problems due to their favourable stability properties and the rapid convergence of non-preconditioned Krylov subspace methods for computing matrix-vector products involving exponential-like functions of the Jacobian matrix. In this article, we implement the so-called exponential Rosenbrock methods using Krylov subspaces. Numerical experiments on a challenging real-world test problem reveal that these methods are a promising preconditioner-free alternative to well-established approaches based on preconditioned

---

<http://journal.austms.org.au/ojs/index.php/ANZIAMJ/article/view/6331>

gives this article, © Austral. Mathematical Soc. 2013. Published May 11, 2013, as part of the Proceedings of the 16th Biennial Computational Techniques and Applications Conference. ISSN 1446-8735. (Print two pages per sheet of paper.) Copies of this article must not be made otherwise available on the internet; instead link directly to this URL for this article.

Newton–Krylov implementations of the backward differentiation formulas.

*Subject class:* 65L04

*Keywords:* exponential integrators, backward differentiation formulas, stiff ordinary differential equations, Krylov subspace methods, porous media

# Contents

<b>1</b>	<b>Introduction</b>	<b>C103</b>
<b>2</b>	<b>Exponential Rosenbrock methods</b>	<b>C107</b>
2.1	General framework . . . . .	C107
2.2	Embedded methods . . . . .	C108
2.3	Krylov subspace methods . . . . .	C109
<b>3</b>	<b>Results</b>	<b>C111</b>
<b>4</b>	<b>Conclusions</b>	<b>C116</b>
	<b>References</b>	<b>C117</b>

## 1 Introduction

We study time integration methods for initial value problems involving large stiff systems of nonlinear ordinary differential equations (ODEs) in autonomous form:

$$\mathbf{u}'(t) = \mathbf{g}(\mathbf{u}), \quad \mathbf{u}(0) = \mathbf{u}_0, \quad 0 < t \leq T, \tag{1}$$

where  $\mathbf{u} \in \mathbb{R}^N$  with  $N$  large,  $\mathbf{g} : \mathbb{R}^N \rightarrow \mathbb{R}^N$  is a nonlinear vector-valued function of  $\mathbf{u}$ ,  $\mathbf{u}'$  is the time derivative of  $\mathbf{u}$ , and  $\mathbf{u}_0$  is the known initial solution vector. Such problems frequently arise from spatial discretisations of

nonlinear time-dependent partial differential equations, where the entries of the vector  $\mathbf{u}$  are the discrete solution values and the stiffness of the system is characterised by the presence of one or more eigenvalues of the Jacobian matrix  $\mathbf{J} = \partial \mathbf{g} / \partial \mathbf{u}$  with large negative real parts.

A time integration method applied to (1) produces an approximation to the continuous solution  $\mathbf{u}(t)$  taking the form of individual approximations  $\mathbf{u}_n \approx \mathbf{u}(t_n)$  at a discrete set of times  $\{t_n, n \in \mathbb{N}\}$  with step sizes  $\tau_n = t_{n+1} - t_n$ . Such methods are often classified as being explicit or implicit. In terms of the computation required per time step, explicit methods are generally more attractive because they do not require the solution of a linear or nonlinear system of equations at each time step. However, for stiff problems implicit methods generally outperform explicit methods. Roughly speaking, the reason for this is due to the eigenvalue distribution described above. Since implicit methods have a much larger region of absolute stability<sup>1</sup>, the step size is less restricted by stability constraints.

Perhaps the most widely used implicit methods for stiff problems are the backward differentiation formulas (BDFs). The idea is to evaluate (1) at  $t = t_{n+1}$ , and then use a  $q$ th order backward difference approximation<sup>2</sup> to approximate  $\mathbf{u}'(t_{n+1})$ :

$$\mathbf{u}'(t_{n+1}) \approx \sum_{i=0}^q \alpha_{n,i} \mathbf{u}_{n+1-i},$$

where the coefficients  $\alpha_{n,i}$  depend on the order  $q$ , the recent step size history and the current step size  $\tau_n$ . The  $q$ th order scheme<sup>3</sup> takes the form of a

---

<sup>1</sup>For a fixed step size  $\tau$ , the region of absolute stability specifies any constraints on  $\tau$  needed to ensure the approximate solution to the linear problem  $\mathbf{y}' = \lambda \mathbf{y}$ , where  $\lambda \in \mathbb{C}$  with  $\Re(\lambda) < 0$ , matches the asymptotic behaviour of the exact solution, that is,  $\mathbf{y}(t) \rightarrow 0$  as  $t \rightarrow \infty$ . A scheme is said to be A-stable if it is stable for all  $\Re(\lambda) < 0$ .

<sup>2</sup>Such an approximation is derived by differentiating the interpolating polynomial (of degree at most  $q$ ) that passes through the  $q + 1$  points  $(t_{n+1-i}, \mathbf{u}_{n+1-i})$ ,  $i = 0, \dots, q$ .

<sup>3</sup>For  $q = 1$  the well-known backward Euler method, a first order one-step method, is

system of nonlinear equations,

$$\mathbf{f}(\mathbf{u}_{n+1}) \equiv \mathbf{u}_{n+1} - \gamma_n \mathbf{g}(\mathbf{u}_{n+1}) + \mathbf{a}_n = \mathbf{0}, \quad (2)$$

where  $\mathbf{a}_n = \sum_{i=1}^q (\alpha_{n,i}/\alpha_{n,0}) \mathbf{u}_{n+1-i}$  and  $\gamma_n = \tau_n/\alpha_{n,0}$ , which must be solved at every time step for  $\mathbf{u}_{n+1}$ .

The computational cost of solving (2) is dominated by the solution of a system of linear equations at each Newton iteration:

$$[\mathbf{I} - \gamma_n \mathbf{J}(\mathbf{u}_{n+1,k})] \mathbf{x}_k = -\mathbf{f}(\mathbf{u}_{n+1,k}), \quad (3)$$

where  $\mathbf{x}_k = \mathbf{u}_{n+1,k+1} - \mathbf{u}_{n+1,k}$  and  $\mathbf{u}_{n+1,k}$  denotes the  $k$ th Newton iterate.

For large problems, Krylov subspace methods are favoured because they require only matrix-vector products with the Jacobian matrix. This is attractive because such products are approximated cheaply using finite difference approximations [7, §3.2.1, e.g.]. Such methods for solving systems of nonlinear equations are called *Jacobian-free Newton-Krylov methods* [8]. The preferred Krylov method is usually the generalised minimal residual method (GMRES), since it minimises the 2-norm of the residual vector at each iteration. The catch is that convergence is usually unsatisfactory without some form of preconditioning. Using *right* preconditioning, the linear system (3) is solved in two stages by introducing a preconditioner matrix  $\mathbf{M}$ :

- solve  $[\mathbf{I} - \gamma_n \mathbf{J}(\mathbf{u}_{n+1,k})] \mathbf{M}^{-1} \tilde{\mathbf{x}}_k = -\mathbf{f}(\mathbf{u}_{n+1,k})$  for  $\tilde{\mathbf{x}}_k$  using GMRES;
- solve  $\mathbf{M} \mathbf{x}_k = \tilde{\mathbf{x}}_k$  for  $\mathbf{x}_k$ .

The choice of  $\mathbf{M}$  is highly problem-dependent: a good preconditioner for one problem is not necessarily a good preconditioner for another [9]. A standard approach, which works well in many cases, is to choose  $\mathbf{M}$  equal to the coefficient matrix of (3) but ‘freeze’ the matrix over multiple time steps [8]. Note that for this particular choice of  $\mathbf{M}$ , the approach is not *truly*

---

obtained, whereas for  $q > 1$ , the backward differentiation formulas are all higher order multistep methods as  $\mathbf{u}_{n+1}$  depends on not only  $\mathbf{u}_n$  but also  $\mathbf{u}_{n-1}, \dots, \mathbf{u}_{n+1-q}$ .

Jacobian-free since the Jacobian matrix is formed and stored whenever  $\mathbf{M}$  is updated.

Over the past decade, a family of time integration methods known as exponential integrators have emerged as a serious contender to implicit methods for large, stiff problems. In the simplest of terms, an exponential integrator for (1) is any time integration method that involves the exponential of the Jacobian matrix. The prototype method (in essence, the backward Euler equivalent for exponential integrators) is known as the exponential Rosenbrock–Euler method [6]. It is summarised as follows. Given the approximate solution at  $t = t_n$ , denoted by  $\mathbf{u}_n$ , linearise the nonlinear function  $\mathbf{g}(\mathbf{u})$  to obtain

$$\mathbf{u}'(t) = \mathbf{g}_n + \mathbf{J}_n(\mathbf{u} - \mathbf{u}_n),$$

where  $\mathbf{g}_n = \mathbf{g}(\mathbf{u}_n)$  and  $\mathbf{J}_n = \mathbf{J}(\mathbf{u}_n)$ . Solving this system of linear ODEs exactly using the integrating factor:

$$\int_{t_n}^{t_{n+1}} \frac{d}{dt} (e^{-t\mathbf{J}_n} \mathbf{u}) dt = \int_{t_n}^{t_{n+1}} e^{-t\mathbf{J}_n} (\mathbf{g}_n - \mathbf{J}_n \mathbf{u}_n) dt,$$

produces the scheme

$$\mathbf{u}_{n+1} = \mathbf{u}_n + \tau_n \varphi_1(\tau_n \mathbf{J}_n) \mathbf{g}_n, \quad (4)$$

where

$$\varphi_1(z) = \frac{e^z - 1}{z}. \quad (5)$$

Note that (4) is explicit:  $\mathbf{u}_{n+1}$  is defined explicitly in term of  $\mathbf{u}_n$ , so there is no nonlinear or linear system to solve. Instead, computation of the matrix function vector product  $\varphi_1(\tau_n \mathbf{J}_n) \mathbf{g}_n$  is required per time step.

The exponential Rosenbrock–Euler method is a second order method. Higher order methods for (1) have been developed and include the exponential Rosenbrock methods [6] and the exponential propagation iterative methods [9]. Such methods require multiple matrix function vector products per time

step involving the so-called *phi functions*, which are defined by (5) and the recurrence relation:

$$\varphi_k(z) = \frac{\varphi_{k-1}(z) - \varphi_{k-1}(0)}{z}, \quad k = 2, 3, \dots$$

The attraction of using an exponential integrator is due to two main reasons. Firstly, they have excellent stability properties—they are usually exact for linear problems (that is, problems in the form of (1) with  $\mathbf{g}(\mathbf{u}) = \mathbf{A}\mathbf{u} + \mathbf{b}$ , where  $\mathbf{A}$  and  $\mathbf{b}$  have fixed entries), which means that they are  $\mathbf{A}$  stable. In contrast, only the first and second order backward differentiation formulas possess this property. The second reason is that Krylov subspace methods for computing matrix function vector products involving  $\varphi_k(\tau_n \mathbf{J}_n)$  typically converge rapidly, meaning that the integrators can be implemented efficiently without preconditioning.

The aim of this article is to compare the performance of higher order exponential integrators to the backward differentiation formulas. Similar comparisons for first and second order methods [3] and fixed step size implementations [9] have featured previously. We focus on the exponential Rosenbrock methods of Hochbruck, Ostermann and Schweitzer [6]. We briefly discuss their derivation, present methods of order three and four, and explain the implementation of the schemes using Krylov subspace methods (see Section 2). Numerical experiments comparing the performance of each method are carried out in Section 3.

## 2 Exponential Rosenbrock methods

### 2.1 General framework

We write the initial value problem (1) in the form

$$\mathbf{u}'(\mathbf{t}) = \mathbf{g}_n + \mathbf{J}_n(\mathbf{u} - \mathbf{u}_n) + \mathbf{d}(\mathbf{u}),$$

where  $\mathbf{d}(\mathbf{u}) = \mathbf{g}(\mathbf{u}) - \mathbf{g}_n - \mathbf{J}_n(\mathbf{u} - \mathbf{u}_n)$ . Using the integrating factor  $e^{-t\mathbf{J}_n}$ , the following expression is obtained for the exact solution at  $\mathbf{t} = \mathbf{t}_{n+1}$ ,

$$\mathbf{u}(\mathbf{t}_{n+1}) = \mathbf{u}_n + \tau_n \varphi_1[\tau_n \mathbf{J}(\mathbf{u}_n)] \mathbf{g}(\mathbf{u}_n) + \int_{\mathbf{t}_n}^{\mathbf{t}_{n+1}} e^{(\mathbf{t}_{n+1}-t)\mathbf{J}(\mathbf{u}_n)} \mathbf{d}[\mathbf{u}(t)] dt.$$

This formula is the starting point for developing higher order one-step exponential integrators for (1). The basic idea is to approximate the integral using an appropriate quadrature rule. Since the integrand depends on the solution  $\mathbf{u}(t)$ , in general, a time integration method is obtained consisting of  $s$  internal stages  $\mathbf{u}_{n,i}$  that approximate the solution at  $\mathbf{t} = \mathbf{t}_n + \mathbf{c}_i \tau_n$ :

$$\begin{aligned} \mathbf{u}_{n,i} &= \mathbf{u}_n + \mathbf{c}_i \tau_n \varphi_1(\mathbf{c}_i \tau_n \mathbf{J}_n) \mathbf{g}_n + \tau_n \sum_{j=1}^{i-1} \mathbf{a}_{ij}(\mathbf{c}_i \tau_n \mathbf{J}_n) \mathbf{d}_{n,j}, \quad i = 1, \dots, s, \\ \mathbf{u}_{n+1} &= \mathbf{u}_n + \tau_n \varphi_1(\tau_n \mathbf{J}_n) \mathbf{g}_n + \tau_n \sum_{i=1}^s \mathbf{b}_i(\tau_n \mathbf{J}_n) \mathbf{d}_{n,i}, \end{aligned} \quad (6)$$

where  $\mathbf{d}_{n,i} = \mathbf{d}(\mathbf{u}_{n,i})$  and the scalars  $\mathbf{c}_i$  are called the nodes of the method. The quadrature weights  $\mathbf{b}_i(z)$  and coefficients  $\mathbf{a}_{ij}(z)$  are linear combinations of the functions  $\varphi_k(z)$  and  $\varphi_k(\mathbf{c}_i z)$ . Conditions on these functions determine the order of the method [6].

Note that the exponential Rosenbrock–Euler method (4) is an exponential Rosenbrock method with  $s = 0$ .

## 2.2 Embedded methods

A critical component of time integration methods for large, stiff problems is the use of adaptive step size control, whereby the local error is estimated at every time step and the step size adjusted to satisfy absolute and relative error tolerances. The idea is to construct a pair of embedded methods to obtain two approximate solutions:  $\mathbf{u}_{n+1}$  and an error estimator  $\hat{\mathbf{u}}_{n+1}$ . The

difference between these solutions is taken as an estimate of the local error at each time step. To reduce the computation introduced, the error estimator relies on the same internal stages  $\mathbf{u}_{n,i}$ ,  $i = 1, \dots, s$  and takes the general form

$$\hat{\mathbf{u}}_{n+1} = \mathbf{u}_n + \tau_n \varphi_1(\tau_n J_n) \mathbf{g}_n + \tau_n \sum_{i=1}^s \hat{\mathbf{b}}_i(\tau_n J_n) \mathbf{d}_{n,i}.$$

The following pairs of embedded methods were given by Hochbruck et al. [6].

The `exprb32` scheme is a one stage pair of embedded exponential Rosenbrock methods of orders two and three:

$$\begin{aligned} \mathbf{u}_{n,1} &= \mathbf{u}_n + \tau_n \varphi_1(\tau_n J_n) \mathbf{g}_n, \\ \mathbf{u}_{n+1} &= \mathbf{u}_n + \tau_n \varphi_1(\tau_n J_n) \mathbf{g}_n + \tau_n \mathbf{b}_1(\tau_n J_n) \mathbf{d}_{n,1}, \\ \hat{\mathbf{u}}_{n+1} &= \mathbf{u}_n + \tau_n \varphi_1(\tau_n J_n) \mathbf{g}_n, \end{aligned}$$

where  $\mathbf{b}_1(\mathbf{z}) = 2\varphi_3(\mathbf{z})$ . The `exprb43` scheme is a two stage pair of embedded exponential Rosenbrock methods of orders three and four,;

$$\begin{aligned} \mathbf{u}_{n,1} &= \mathbf{u}_n + \frac{1}{2} \tau_n \varphi_1\left(\frac{1}{2} \tau_n J_n\right) \mathbf{g}_n, \\ \mathbf{u}_{n,2} &= \mathbf{u}_n + \tau_n \varphi_1(\tau_n J_n) \mathbf{g}_n + \tau_n \mathbf{a}_{21}(\tau_n J_n) \mathbf{d}_{n,1}, \\ \mathbf{u}_{n+1} &= \mathbf{u}_n + \tau_n \varphi_1(\tau_n J_n) \mathbf{g}_n + \tau_n \mathbf{b}_1(\tau_n J_n) \mathbf{d}_{n,1} + \tau_n \mathbf{b}_2(\tau_n J_n) \mathbf{d}_{n,2}, \\ \hat{\mathbf{u}}_{n+1} &= \mathbf{u}_n + \tau_n \varphi_1(\tau_n J_n) \mathbf{g}_n + \tau_n \hat{\mathbf{b}}_1(\tau_n J_n) \mathbf{d}_{n,1} + \tau_n \hat{\mathbf{b}}_2(\tau_n J_n) \mathbf{d}_{n,2}, \end{aligned}$$

where  $\mathbf{a}_{21} = \varphi_1(\mathbf{z})$ ,  $\mathbf{b}_1(\mathbf{z}) = 16\varphi_3(\mathbf{z}) - 48\varphi_4(\mathbf{z})$ ,  $\mathbf{b}_2(\mathbf{z}) = -2\varphi_3(\mathbf{z}) + 12\varphi_4(\mathbf{z})$ ,  $\hat{\mathbf{b}}_1(\mathbf{z}) = 16\varphi_3(\mathbf{z})$  and  $\hat{\mathbf{b}}_2(\mathbf{z}) = -2\varphi_3(\mathbf{z})$ . In both schemes the higher order method is used to continue the integration.

## 2.3 Krylov subspace methods

In general, an  $s$  stage exponential Rosenbrock method contains (at most)  $(s+1)(s+2)/2$  unique products involving a matrix function and a vector. However,

since all products involve the matrix  $J_n$  and either of the vectors  $\mathbf{g}_n$  or  $\mathbf{d}_{n,k}$ ,  $k = 1, \dots, s$ , only  $s + 1$  Krylov subspaces are required. For example, the scheme `exprb43` requires seven unique products but only three Krylov subspaces per time step.

In what follows, we describe the computation of  $f(\tau_n J_n) \mathbf{b}_k$ , where

$$\mathbf{b}_k = \begin{cases} \mathbf{g}_n, & k = 0, \\ \mathbf{d}_{n,k}, & k = 1, \dots, s, \end{cases}$$

and  $f(\mathbf{z})$  is any one of the functions  $\varphi_1(\mathbf{z})$ ,  $\mathbf{b}_i(\mathbf{z})$ ,  $\mathbf{a}_{ij}(\mathbf{z})$  or  $\hat{\mathbf{b}}_i(\mathbf{z})$ . Approximations are extracted from the small dimensional Krylov subspaces  $\mathcal{K}_m(J_n, \mathbf{b}_k)$ ,  $k = 0, \dots, s$  using  $s + 1$  applications of Arnoldi's method, yielding the relations

$$J_n \mathbf{V}_m^{(k)} = \mathbf{V}_m^{(k)} \mathbf{H}_m^{(k)} + \beta_m^{(k)} \mathbf{v}_{m+1}^{(k)} \mathbf{e}_m^T, \quad k = 0, \dots, s,$$

where  $\mathbf{v}_1 = \mathbf{b}_k / \|\mathbf{b}_k\|_2$ , the columns of  $\mathbf{V}_m^{(k)} = [\mathbf{v}_1^{(k)}, \mathbf{v}_2^{(k)}, \dots, \mathbf{v}_m^{(k)}] \in \mathbb{R}^{N \times m}$  form an orthonormal basis for  $\mathcal{K}_m(J_n, \mathbf{b}_k)$ ,  $\mathbf{H}_m^{(k)} = \mathbf{V}_m^{(k)T} J_n \mathbf{V}_m^{(k)} \in \mathbb{R}^{m \times m}$  is an upper Hessenberg matrix and  $\mathbf{e}_m$  is the  $m$ th canonical basis vector in  $\mathbb{R}^m$  [1]. Since  $\mathbf{d}_{n,1}$  depends on  $\mathbf{g}_n$  and  $\mathbf{d}_{n,k}$  depends on  $\mathbf{d}_{n,k-1}$  for  $k = 2, \dots, s$  all applications of Arnoldi's method must be performed serially.

The Krylov approximation is then [2, 5]

$$f(\tau_n J_n) \mathbf{b}_k \approx \|\mathbf{b}_k\|_2 \mathbf{V}_m^{(k)} f(\mathbf{H}_m^{(k)}) \mathbf{e}_1, \quad k = 0, \dots, s, \quad (7)$$

which reduces the evaluation of  $f$  to a small  $m \times m$  matrix that is computed using standard approaches such as Padé approximation [5]. Beginning at  $m = 1$  the subspace is expanded by one additional vector at each iteration producing an updated Arnoldi relation and approximation (7). Termination of the iterative procedure is based on the *generalized residual vector* [2, 5].

### 3 Results

Our interest in initial value problems in the form of (1) lies in the field of simulating transport in porous media and therefore we use the well-known two dimensional test problem, Richards' equation. We solve the two-dimensional Richards' equation,

$$\frac{\partial \theta(\mathbf{h})}{\partial t} = \frac{\partial}{\partial x} \left[ \mathbf{K}(\mathbf{h}) \frac{\partial \mathbf{h}}{\partial x} \right] + \frac{\partial}{\partial z} \left[ \mathbf{K}(\mathbf{h}) \left( \frac{\partial \mathbf{h}}{\partial z} + 1 \right) \right],$$

on  $\Omega = \{(x, z) \mid 0 < x < 5, 0 < z < 3\}$  and  $0 < t < 12.5$ , subject to the initial condition and boundary conditions:

$$\begin{aligned} \mathbf{h} &= -500, & t &= 0, \\ \mathbf{K}(\mathbf{h}) \frac{\partial \mathbf{h}}{\partial x} &= 0, & x &= 0, x = 5, \\ \mathbf{K}(\mathbf{h}) \frac{\partial \mathbf{h}}{\partial z} + \mathbf{K}(\mathbf{h}) &= 0, & z &= 0, \\ \mathbf{K}(\mathbf{h}) \frac{\partial \mathbf{h}}{\partial z} + \mathbf{K}(\mathbf{h}) &= 0, & z &= 3, 0 < x < 2, 3 < x < 5, \\ \mathbf{K}(\mathbf{h}) \left( \frac{\partial \mathbf{h}}{\partial x} + 1 \right) &= 0.05, & z &= 3, 2 < x < 3. \end{aligned}$$

The moisture content  $\theta$ , hydraulic conductivity  $\mathbf{K}$  and effective saturation  $S_e$  are all defined in terms of the pressure head  $\mathbf{h} < 0$  as given by the van Genuchten relations [1]:

$$\begin{aligned} \theta(\mathbf{h}) &= \theta_{\text{res}} + (\theta_{\text{sat}} - \theta_{\text{res}}) S_e(\mathbf{h}), \\ \mathbf{K}(\mathbf{h}) &= \mathbf{K}_{\text{sat}} \sqrt{S_e(\mathbf{h})} \left\{ 1 - [1 - S_e(\mathbf{h})^{1/m}]^m \right\}^2, \\ S_e(\mathbf{h}) &= [1 + (-\alpha \mathbf{h})^n]^{-m}, \end{aligned} \tag{8}$$

where  $m = 1 - 1/n$ .

The specific test problem concerns the flow of water into a very dry rectangular region divided into nine alternating blocks of two heterogeneous soils: sand and clay (see Figure 1 and Table 3). All boundaries are zero flux apart from a 1 m strip in the centre of the top block of sand, where a constant influx of water at 0.05 m per day is applied.

A spatial discretisation was performed using the finite volume method on a rectangular grid consisting of 129 nodes in both the  $x$  and  $z$  directions. Horizontal symmetry of the problem is not exploited. This produces an initial value problem in the form of (1) with  $N = 129^2 = 16641$ , where the vector  $\mathbf{u}$  contains the unknown values of the pressure head  $h$  at each of the nodes. Carr, Moroney and Turner [1] give full details.

The schemes `exprb32` and `exprb43` were implemented in Matlab with Krylov subspace methods. In addition, we implemented the exponential Rosenbrock–Euler method using the step size control strategy featured in our previous article [1]. We refer to this method as `exprem22` henceforth. Performance comparisons were made with the Matlab interface to the CVODE module of the suite of nonlinear/differential-algebraic solvers (SundialsTB v2.5.0). For stiff problems, CVODE employs a variable order implementation of the backward differentiation formulas, meaning that the order is adjusted throughout the time integration with the goal of maximising the step size. Implementations with upper bounds of two, three and four on the order are denoted by `bdf2`, `bdf3` and `bdf4`. CVODE provides several choices for the solution of the linear systems at each Newton iteration. In this work, we used the in-built GMRES solver with the ‘freezing’ preconditioning strategy outlined in Section 1. The specific solver options are given in Table 1. Full details are provided in the Sundials documentation [4].

All methods tested were implemented in a Jacobian-free framework, that is, Jacobian matrix-vector products were approximated using

$$\mathbf{J}(\mathbf{u})\mathbf{v} \approx [\mathbf{g}(\mathbf{u} + \varepsilon\mathbf{v}) - \mathbf{g}(\mathbf{u})]/\varepsilon, \quad (9)$$

where  $\varepsilon$  is a suitably-defined perturbation [1, 4]. For CVODE, when updating

Table 1: Solver options used in CVODE.

Option	Value
LinearSolver	GMRES
PrecType	Right
PrecModule	BandPre
LowerBwidth	29
UpperBwidth	129

Figure 1: Numerical solution of the test problem. Contour plot of the saturation field ( $\theta/\theta_{\text{sat}}$ ). Animation over varying time viewable in at least Adobe Reader version 9.

the preconditioner, the full Jacobian was computed numerically taking advantage of its banded structure by using (9) with appropriate choices for the vector  $\mathbf{v}$  [7, §2.3].

A benchmark solution was obtained using CVODE (with maximum order five) with very small absolute and relative error tolerances. For each of the six methods tested, simulations were performed in Matlab 2012b with absolute

Table 2: Hydraulic properties of sand and clay used in the van Genuchten relations in equation (8).

Material	$\theta_{\text{res}}$	$\theta_{\text{sat}}$	$K_{\text{sat}} [\text{ms}^{-1}]$	$\alpha [\text{m}^{-1}]$	$n$
Clay	0.106	0.469	$1.52 \times 10^{-6}$	1.04	1.40
Sand	0.029	0.366	$6.26 \times 10^{-5}$	2.80	2.23

and relative error tolerances (`tol`) equal to  $10^{-4}, 10^{-5}, \dots, 10^{-8}$ , respectively. Table 3 lists the relative error of the solution at  $t = 12.5$  days, the total number of function calls to  $g(\mathbf{u})$  (for `bdf2`, `bdf3` and `bdf4` this includes the function evaluations required to form the Jacobian matrix when updating the preconditioner), the total number of accepted time steps and the CPU time for each simulation. Although we understand that CPU times strongly depend on the available computer architecture, we feel that the comparisons are of interest.

Comparing the second order methods, we conclude that `exprem22` comprehensively outperforms `bdf2` because it requires fewer function calls, fewer time steps and less CPU time to compute a solution of equivalent accuracy. For the third order methods, `exprem32` and `bdf3`, we observe that both schemes are comparable in terms of computation required to obtain a solution of desired accuracy. Finally, the fourth order backward differentiation formula, `bdf4`, performs better than the fourth order exponential Rosenbrock method, `exprb43`.

An important observation is the increase in computation per step when increasing the order of the exponential Rosenbrock methods. For a comparable number of time steps, the number of function calls increases as evidenced by the simulations: `exprem22` (2437 steps, 18335 calls), `exprb32` (2783 steps, 25217 calls) and `exprb43` (2692 steps, 36200 calls). Such a trend is not observable for the the backward differentiation formulas: `bdf2` (2860 steps, 20183 calls), `bdf3` (2926 steps, 20603 calls), `bdf4` (2386 steps, 16191 calls). The main reason for this is that the higher order exponential Rosenbrock

Table 3: Simulation statistics for the test problem. For each of the six methods, approximate solutions are obtained using absolute and relative error tolerances ( $\text{tol}$ ). Statistics listed are the relative error of the approximate solution at  $t = 12.5$  days, the total number of function calls to  $g(\mathbf{u})$  (for `bdf2`, `bdf3` and `bdf4` this includes the function evaluations required to form the Jacobian matrix when updating the preconditioner), the total number of accepted time steps and the CPU time.

$\text{tol}$	<code>exprem22</code>	<code>bdf2</code>	<code>exprb32</code>	<code>bdf3</code>	<code>exprb43</code>	<code>bdf4</code>
Relative error						
$10^{-4}$	8.7E-3	1.4E-3	1.6E-2	3.3E-4	5.2E-2	2.1E-5
$10^{-5}$	4.1E-4	3.2E-4	3.4E-4	6.0E-5	5.2E-3	1.5E-6
$10^{-6}$	7.4E-6	5.8E-5	3.9E-6	8.8E-6	1.6E-4	4.2E-7
$10^{-7}$	3.4E-6	1.4E-5	3.5E-6	1.9E-6	6.7E-7	1.0E-7
$10^{-8}$	8.7E-7	3.2E-6	8.0E-7	3.3E-7	2.8E-7	2.9E-8
Total number of function calls						
$10^{-4}$	11688	20183	15408	12267	20945	11134
$10^{-5}$	18335	40286	25217	20603	25721	16191
$10^{-6}$	31062	85000	45404	36026	36200	24292
$10^{-7}$	56778	177813	89813	62152	55758	38410
$10^{-8}$	112346	382927	186361	107069	92538	59556
Total number of accepted time steps						
$10^{-4}$	1234	2860	1360	1707	1014	1561
$10^{-5}$	2437	6039	2783	2926	1598	2386
$10^{-6}$	5156	13365	5906	5484	2692	3660
$10^{-7}$	10660	28022	12652	9717	4664	5960
$10^{-8}$	22233	60273	27221	16826	8214	9329
CPU time [mins]						
$10^{-4}$	2.6	8.2	3.4	5.4	5.5	4.6
$10^{-5}$	4.0	15.3	5.3	8.4	6.2	6.5
$10^{-6}$	6.6	31.3	9.5	14.2	8.2	9.7
$10^{-7}$	11.9	64.6	18.5	24.2	12.7	14.2
$10^{-8}$	23.6	142.0	38.4	41.2	21.0	22.1

methods require computation of additional matrix function vector products per time step, whereas, for the backward differentiation formulas, the solution of one nonlinear system is required per step, regardless of the order used.

Another key finding is the relatively low accuracy of `exprb32` and `exprb43` compared with `bdf3` and `bdf4` for a given value of the tolerance (`tol`). This indicates that the exponential Rosenbrock methods are under estimating the local error at each time step. The reason for this is almost certainly that the embedded methods, used to estimate the local error, are of lower order.

Finally, interestingly, the computational cost required to compute a solution with relative error less than  $10^{-5}$ . Surprisingly, `exprem22` (6.6 mins, 31062 calls) places second in this category, just behind `bdf4` (6.5 mins, 16191 calls) and well in front of `exprb32` (9.5 mins, 45404 calls), `exprb43` (12.7 mins, 55758 calls), `bdf3` (14.2 mins, 36026 calls) and `bdf2` (142 mins, 382927 calls).

## 4 Conclusions

Exponential Rosenbrock methods are competitive with backward differentiation formulas when applied to stiff problems arising from spatial discretisations of time-dependent partial differential equations. For very large problems, these methods are particularly attractive because they do not require preconditioning. Our numerical experiments indicate that, although the third and fourth order exponential Rosenbrock methods are competitive with the third and fourth order BDFs, they suffer from an increase in the number of function evaluations due to the multiple matrix function vector products (and therefore multiple Krylov subspaces) required per time step. A key point is that all of these subspaces involve the same Jacobian matrix. Future work will therefore focus on reducing the function evaluation count by recycling some of the spectral information from one subspace to the next.

**Acknowledgements** This work was supported by the Australian Research Council Discovery Project grant DP120103770. We thank the two anonymous reviewers for their detailed analysis and suggestions which improved the quality of the final version of this article.

## References

- [1] E. J. Carr, T. J. Moroney and I. W. Turner. Efficient simulation of unsaturated flow using exponential time integration. *Appl. Math. Comput.*, 217(14): 6587–6596, 2011. doi:[10.1016/j.amc.2011.01.041](https://doi.org/10.1016/j.amc.2011.01.041) C110, C111, C112
- [2] E. J. Carr, I. W. Turner and M. Ilic. Krylov subspace approximations for the exponential Euler method: error estimates and the harmonic Ritz approximant. *ANZIAM Journal*, 52: C612–C627, 2011. <http://journal.austms.org.au/ojs/index.php/ANZIAMJ/article/view/3938> C110
- [3] E. J. Carr, I. W. Turner and P. Perré. A variable step size Jacobian-free exponential integrator for simulating transport in heterogeneous porous media: application to wood drying. *J. Comput. Phys.*, 233: 66–82, 2013. doi:[10.1016/j.jcp.2012.07.024](https://doi.org/10.1016/j.jcp.2012.07.024) C107
- [4] Sundials user documentation. <https://computation.llnl.gov/casc/sundials/documentation/documentation.html> C112
- [5] M. Hochbruck, C. Lubich and H. Selhofer. Exponential integrators for large systems of differential equations. *SIAM J. Sci. Comput.*, 19(5): 1552–1574, 1998. doi:[10.1137/S1064827595295337](https://doi.org/10.1137/S1064827595295337) C110
- [6] M. Hochbruck, A. Ostermann and J. Schweitzer. Exponential Rosenbrock-type methods. *SIAM J. Numer. Anal.*, 47(1): 786–803, 2009. doi:[10.1137/080717717](https://doi.org/10.1137/080717717) C106, C107, C108, C109

- [7] C. T. Kelley. *Solving nonlinear equations with Newton's method*. SIAM, USA, 2003. [C105](#), [C113](#)
- [8] D. A. Knoll and D. E. Keyes. Jacobian-free Newton–Krylov methods: a survey of approaches and applications. *J. Comput. Phys.*, 193(2): 357–397, 2004. doi:[10.1016/j.jcp.2003.08.010](#) [C105](#)
- [9] M. Tokman. Efficient integration of large stiff systems of ODEs with exponential propagation iteration (EPI) methods. *J. Comput. Phys.*, 213: 748–776, 2006. doi:[10.1016/j.jcp.2005.08.032](#) [C105](#), [C106](#), [C107](#)

## Author addresses

1. **E. J. Carr**, Mathematical Sciences School, Queensland University of Technology, Brisbane, Queensland 4000, Australia.  
<mailto:elliott.carr@qut.edu.au>
2. **T. J. Moroney**, Mathematical Sciences School, Queensland University of Technology, Brisbane, Queensland 4000, Australia.  
<mailto:t.moroney@qut.edu.au>
3. **I. W. Turner**, Mathematical Sciences School, Queensland University of Technology, Brisbane, Queensland 4000, Australia.  
<mailto:i.turner@qut.edu.au>