On the performance of exponential integrators for problems in magnetohydrodynamics

Lukas Einkemmer^{a,*}, Mayya Tokman^b, John Loffeld^c

^a*University of Innsbruck, Innsbruck, Austria*

^b*University of California, Merced, United States of America*

^c*Lawrence Livermore National Laboratory, Livermore, United States of America*

Abstract

Exponential integrators have been introduced as an efficient alternative to explicit and implicit methods for integrating large stiff systems of differential equations. Over the past decades these methods have been studied theoretically and their performance was evaluated using a range of test problems. While the results of these investigations showed that exponential integrators can provide significant computational savings, the research on validating this hypothesis for large scale systems and understanding what classes of problems can particularly benefit from the use of the new techniques is in its initial stages. Resistive magnetohydrodynamic (MHD) modeling is widely used in studying large scale behavior of laboratory and astrophysical plasmas. In many problems numerical solution of MHD equations is a challenging task due to the temporal stiffness of this system in the parameter regimes of interest. In this paper we evaluate the performance of exponential integrators on large MHD problems and compare them to a state-of-the-art implicit time integrator. Both the variable and constant time step exponential methods of EpiRK-type are used to simulate magnetic reconnection and the Kelvin–Helmholtz instability in plasma. Performance of these methods, which are part of the EPIC software package, is compared to the variable time step variable order BDF scheme included in the CVODE (part of SUNDIALS) library. We study performance of the methods on parallel architectures and with respect to magnitudes of important parameters such as Reynolds, Lundquist, and Prandtl numbers. We find that the exponential integrators provide superior or equal performance in most circumstances and conclude that further development of exponential methods for MHD problems is warranted and can lead to significant computational advantages for large scale stiff systems of differential equations

such as MHD.

1. Introduction

The focus of this article lies at the intersection of computational plasma physics and numerical analysis. The system of resistive magnetohydrodynamics (MHD) equations is used extensively in both laboratory and astrophysical plasma physics [19]. MHD simulations are employed to study large scale plasma dynamics in applications ranging from the design of the next generation of fusion devices to modeling astrophysical jets (e.g. [29, 12, 57]). Since experiments in this field usually impose a high cost and the ability to conduct observations can be very limited, computational modeling in general, and numerical MHD in particular, have become very important components of research in plasma physics. While MHD theory is widely used to study plasma dynamics, the complexity of the MHD equations makes their numerical solution a challenging task. In particular, the MHD system describes a wide range of temporal scales. In applications such as large scale modeling of tokamaks or modeling of solar eruptions, the physical processes of interest evolve on the time scale of the shear Alfvén waves, which is much slower than the fastest mode in the system - the fast magnetosonic wave - and much faster than the slow magnetosonic wave and the resistive processes. Such temporal stiffness coupled with the complexity of the equations poses a double-edged challenge to developing a time integrator for MHD. From one perspective, the stiffness implies that explicit time integrators will suffer from prohibitively severe restrictions on the time step size due to stability considerations. On the other hand the inherent three-dimensional nature of the processes MHD describes and the complexity of spatial operators describing the dynamics makes evaluation of forcing terms computationally expensive. The latter difficulty served as one of the major reasons most large scale MHD codes came to rely extensively on explicit time integrators even for problems where resistivity plays a major role. However, the temporal stiffness of the equations forced researchers working on these software packages to turn to at least some form of implicitness to overcome the small time steps sizes of explicit integrators [25]. Production codes used

*Corresponding author

Email addresses: `lukas.einkemmer@uibk.ac.at` (Lukas Einkemmer),
`mtokman@ucmerced.edu` (Mayya Tokman), `loffeld1@llnl.gov` (John Loffeld)

Preprint submitted to Elsevier

November 16, 2016

in fusion research and astrophysics [48, 4, 18] started embedding some form of implicitness into their time integration schemes. The NIMROD code, for instance, employs an implicitly integrated term in the momentum equation derived by linearizing equations around a static magnetic field and using it to stabilize stiff wave components [49, 48]. In addition to this stabilizing term the overall time integration strategy also uses operator splitting and a staggered temporal grid. The choice of the implicit terms in such semi-implicit schemes is also influenced by the availability of efficient algebraic solvers used to approximate solutions to the associated linear systems. The complexity of the resulting overall time integration method makes it very difficult to use rigorous analysis to study stability and convergence. Such an approach also precludes or makes it challenging to derive error estimators necessary for employing adaptive time stepping strategies or using the time integration as part of the uncertainty quantification methods. Other partially implicit approaches such as alternating direction implicit (ADI) methods [30, 46] or implicit-explicit techniques [27] are also limited as they try to maintain a balance between stabilizing the time integration while maintaining the accuracy of the solution [9]

The difficulties associated with partially implicit methods and the ever pressing need for more efficient and accurate algorithms for high-performance computing platforms gave rise to the efforts to develop fully implicit time integrators for resistive MHD [42, 41, 9, 33]. A typical approach to constructing a fully implicit integrator involves using an implicit scheme such as, for instance, BDF and coupling it with a Newton–Krylov iteration to solve the implicit equations within each time step. In [42, 41], for example, the resistive MHD equations are solved using a variable order variable time step BDF method from the CVODE portion [10] of the SUNDIALS package [20]. While implicit methods possess better stability properties compared to explicit techniques, they are also affected by the stiffness of the equations. In an implicit integrator temporal stiffness manifests itself in the stiff linear systems that have to be solved as part of each Newton iteration. A key to obtaining an efficient fully implicit Newton–Krylov solver is in developing effective scalable preconditioners for these stiff linear systems. For a complex anisotropic system such as MHD construction of a sufficiently efficient preconditioner is simultaneously a central and a challenging task that requires a significant investment of time and effort. If the parameter regime is changed or an extra term is added to the system, it is possible that a new preconditioner has to be developed. Often the structure of the equations

is exploited to build a physics based preconditioner [9] which makes the overall integrator highly dependent on the equations, thus precluding one from using off-the-shelf time integration software packages.

In recent years exponential integrators have emerged as a promising alternative to explicit and implicit methods for solving evolution equations. A significant body of research has been accumulated in which these methods are investigated both from a theoretical (see e.g. [24] for a review article) as well as from a numerical point of view (see e.g. [55] and [31]). This research includes the construction of exponential integrators that have been tailored to a given differential equation (see e.g. [22]) as well as such integrators that can be applied to a class of problems (see e.g. [23, 51]). There is strong evidence to suggest that exponential integrators can offer computational advantages compared to standard integrators particularly for large scale complex problems where no efficient preconditioner is available. Exponential integration can also be used as part of an Implicit-Exponential (IMEXP) scheme [53] for problems where an efficient preconditioner has been constructed for some terms in the equations.

Exponential integrators were first shown to be a promising approach in the context of the MHD equations in [50, 54]. Since then a number of more efficient and sophisticated exponential time integrators have been developed. Exponential propagation iterative methods of Runge-Kutta type (EpiRK) were first introduced in [51] for general nonlinear systems of ordinary differential equations. The EpiRK framework [51, 55, 53] for constructing exponential time integrators has been designed to reduce computational complexity in each time step and provide a flexible ansatz that allows construction of particularly efficient exponential schemes [55, 52, 39, 53]. The most efficient EpiRK methods have been implemented as part of the Exponential Integrator Collection (EPIC) package, a C++/MPI software developed for serial and parallel computational platforms (available for download by request to mtokman@ucmerced.edu). Performance of EPIC integrators has been tested on a suit of numerical examples and it was shown that the EpiRK schemes can achieve superior or equal performance compared to the BDF integrator found in the CVODE library [31, 32].

However, up to now most numerical tests of exponential methods found in the literature considered problems of small to medium complexity (see, for example, [31], [28], or [14]) where in many instances efficient preconditioners can be constructed. The goal of this paper is to apply the EpiRK integrators found in the EPIC library to the problem of solving the resis-

tive magnetohydrodynamics (MHD) equations. General preconditioners for such systems have, for example, been investigated in [41] and it was found that constructing an efficient preconditioner for such system is challenging and the performance gain from standard preconditioning techniques for several dimensional examples is rather modest under most circumstances. Clearly a problem specific preconditioner or a tailor-made time integrator for a specific system of equations will be the most efficient approach in the majority of cases, but developing such a scheme requires extensive effort and time. The goal of this paper is to investigate what computational savings can be achieved with a general exponential integrator compared to an unpreconditioned implicit scheme. Thus we will perform a comparison of EPIC integrators with the BDF methods from the CVODE library without a preconditioner. We aim to demonstrate that exponential integration is a viable technique to use in integrating MHD equations.

The paper is organized as follows. In section 2 we provide an introduction to exponential integrators in general as well as to the specific EpiRK methods we use in the numerical examples and their implementation in EPIC. In section 3 we describe the MHD equations which are solved numerically in section 4 for a reconnection problem, and in section 5 for the Kelvin–Helmholtz instability. Implementation details are discussed in section 6. Finally, our conclusions are presented in section 7.

2. Exponential integrators

Let us consider the following initial value problem

$$\begin{aligned} y' &= F(y), \\ y(t_0) &= y_0 \end{aligned} \tag{1}$$

which is assumed to be large ($y, F(y) \in \mathbf{R}^N$ with $N \gg 1$) and stiff. Let us define by $y(t_n)$ and $y(t_n + h)$ solutions to (1) at times t_n and $t_n + h$ respectively. Assuming that the first order Taylor expansion of the solution $y(t_n + h)$ around solution $y(t_n)$ exists we can write $y(t_n + h)$ in an integral form as follows. Let us denote the remainder function of the first order Taylor expansion of $F(y)$ at $\bar{y}$ as follows

$$R(y, \bar{y}) = F(y) - F(\bar{y}) - F'(\bar{y})(y - \bar{y}). \tag{2}$$

Then (1) can be written as

$$y' = F(y) = F(y(t_n)) + F'(y(t_n))(y - y(t_n)) + R(y, y(t_n)). \quad (3)$$

To simplify the notation we define $J(t_n) = F'(y(t_n))$ and $F(t_n) = F(y(t_n))$. Multiplying equation (3) by a factor $e^{tJ(t_n)}$ and integrating the resulting formula from t_n to $t_n + h$, we obtain the variation-of-constants formula

$$\begin{aligned} y(t_n + h) = y(t_n) &+ (e^{hJ(t_n)} - I) (hJ(t_n))^{-1} hF(t_n) \\ &+ \int_{t_n}^{t_n+h} e^{(t_n+h-t)J(t_n)} R(y(t), y(t_n)) dt. \end{aligned} \quad (4)$$

Changing the integration variable in (4) to s with $t = t_n + sh$, $0 < s < 1$ we obtain

$$\begin{aligned} y(t_n + h) = y(t_n) &+ (e^{hJ(t_n)} - I) (hJ(t_n))^{-1} hF(t_n) \\ &+ \int_0^1 e^{h(1-s)J(t_n)} hR(y(t_n + sh), y(t_n)) ds. \end{aligned} \quad (5)$$

If t_n and $t_n + h$ represent discretization nodes on the time interval $[t_0, t_{end}]$ over which we want to approximate $y(t)$ ($t \in [t_0, t_{end}]$) we can use the integral representation of the solution $y(t_n + h)$ in (5) to construct an exponential integrator. In the following we will denote the numerical approximation at time t_n by y_n , i.e. $y_n \approx y(t_n)$ and our goal is thus to compute $y_{n+1} \approx y(t_{n+1})$ given y_n .

To do that equation (5) has to be approximated in a suitable manner. This task is done in two steps. First, since we do not know $y(t_n)$, $F(t_n)$, and $J(t_n)$ we replace them by the numerical approximations y_n , $F_n = F(y_n)$, and $J_n = J(y_n)$. Then we use a polynomial approximation of the function $R(\cdot, y_n)$ on some nodes s_k . The integral in (5) can then be written as a linear combination of terms of type $\varphi_k(s_k h J_n) v_k$, where the entire functions $\varphi_k(z)$ are defined by

$$\begin{aligned} \varphi_0(z) &= e^z \\ \varphi_k(z) &= \int_0^1 e^{(1-s)z} \frac{s^{k-1}}{(k-1)!} ds. \end{aligned} \quad (6)$$

Given definitions (6) the second term on the right-hand-side of (5) can be written as $\varphi_1(hJ_n)hF(t_n)$.

Note that all terms of type $\varphi_k(chJ)v$, where c and h are constants, J is a matrix and v is a vector, are products of matrix functions and vectors. Since J is the Jacobian of the large stiff system (1), it is unlikely that diagonalizing the matrix J is a computationally cheap, or even a feasible task. Thus the terms $\varphi_k(chJ)v$ have to be approximated in a different way. Computing these terms constitutes the main computational cost of an exponential method. Therefore choosing the quadrature for the integral in (5) in order to minimize the required number of these calculations as well as selecting an appropriate algorithm to evaluate terms of type $\varphi_k(chJ)v$ is essential to the construction of an efficient exponential integrator. Note that construction of a particular integrator involves choosing the parameter c . Small values of c can reduce computational cost as explained below. To explicitly make this point we leave the argument of the φ_k function as chJ rather than use a simplified notation with a single matrix variable.

A number of methods have been proposed in the literature to approximate the exponential-like matrix functions $\varphi_k(chJ)$ or their product with a given vector (e.g. [38, 1, 26, 8, 36, 56] or see the review in [24]). However, many of these techniques are applicable only to small matrices J [35, 47]. In this paper we are concerned with large scale problems such as MHD, where the dimensionality of the system (1) is high (e.g. number of unknowns in the MHD equations multiplied by the number of grid points in two- or three-dimensions), and thus standard techniques for approximating a product of a matrix function φ_k with a vector are either computationally infeasible or prohibitively expensive.

In this paper we use an approach based on the Krylov projection algorithm [2, 11]. Krylov methods have been extensively used to approximate the inverse of a matrix [45] and in 1987 were employed by Van der Vorst to approximate a general function of a matrix and a vector product [11]. The algorithm works by iteratively computing a projection of $\varphi_k(chJ)v$ onto a Krylov subspace

$$K_m(J, v) = \text{span} \{v, Jv, \dots, J^{m-1}v\}$$

using the Arnoldi algorithm [2]. Note that the Krylov subspace does not depend on c and h but since we are computing $\varphi_k(chJ)v$ the number of Krylov vectors computed m depends on the parameters c and h . In fact, the value of m is determined dynamically by estimating the residuals in the course of the iteration [44, 21]. For each m a matrix V_m is formed that consists of the orthonormal vectors that form the basis of $K_m(chJ, v)$. A

side product of the Arnoldi iteration is the $m \times m$ matrix $H_m = chV_m^T J V_m$ which represents the projection of chJ onto the Krylov subspace. A general product of the matrix function f with a vector is then approximated as

$$f(chJ)b \approx V_m (V_m^T \varphi_k(chJ) V_m) V_m^T v \approx V_m \varphi_k(H_m) V_m^T v.$$

H_m is a matrix of dimension $m \times m$ and a small m is usually sufficient to achieve the desired accuracy. Then the application of $\varphi_k(H_m)$ to a vector can be cheaply computed by using any of the standard methods (such as Padé approximation or methods based on polynomial interpolation [35, 24]). The residuals are used to determine the stopping value of m that meets the provided tolerance. Note that the Arnoldi iteration portion of the algorithm is a scale invariant procedure. Thus several terms $\varphi_k(c_i h J)v$ ($i = 1, 2, \dots, K$) can be computed using a single Arnoldi iteration, i.e. using the same Krylov basis. To meet the desired accuracy requirements we just have to ensure that the dimensionality of the basis m is chosen with respect to computing $\varphi_k(c_i h J)v$ with the largest c_i . This can be accomplished using appropriate residuals as described for example in [55].

Another possibility is to use a direct polynomial approximation of φ_k . To that end we have to choose m interpolation nodes on some compact set $K \subset \mathbb{C}$. Chebyshev nodes are an obvious choice, however, they suffer from the disadvantage that to compute the interpolation for $m+1$ points we have to reevaluate all the matrix-vector products already computed. To remedy this, Leja points have been proposed which share many of the favorable properties of the Chebyshev nodes but can be generated in sequence [8, 34, 7, 3, 5, 6]. Another advantage of interpolation at Leja points is its modest memory requirements. This makes it an attractive alternative for computer systems where memory is limited (such as graphic processing units, see [14]). The main disadvantage of the method, however, is that (at least) some approximation of the spectrum of J has to be available.

We also mention another algorithm for approximating products of type $\varphi_k(chJ)v$ – the Taylor polynomial approximation-based method introduced in [1] (in fact, this is a direct polynomial interpolation method, as described above, where all interpolation nodes coincide in a single point). This technique combines ideas of scaling-and-squaring, Taylor expansion and certain error approximations based on the 1-norm of the matrix to estimate $\varphi_k(chJ)v$ or sequences of $\varphi_k(\tau_i chJ)$ at equally spaced intervals τ_i , $\tau_i = i\tau$. As indicated in the experiments conducted in [1] (e.g. Experiment 4), the adaptive Krylov algorithm of [38] can outperform the Taylor method. Note

that not all numerical experiments in [1], particularly Experiment 9 which involves the largest matrix (dimension 250,000), include comparisons with the adaptive Krylov method of Niesen and Wright [38]. In our numerical tests we found that for the large scale matrices we encountered in MHD and other test problems the adaptive Krylov method was more efficient than the Taylor-based algorithm. We have not encountered the instability in the adaptive Krylov algorithm found in Experiment 10 in [1] because such instability can only occur with a Krylov-based method if p and c are increased without bound in the linear combination $\sum_{k=1}^p \varphi_k(chJ)v_k$ (typically $p > 10$). This case is not relevant to most of the high-order exponential integrators proposed in the literature including EpiRK methods which use linear combination of only a few φ_k ($p \leq 5$) with appropriately (and adaptively) scaled c and h . Additionally, as described above, the ability of the adaptive Krylov methods to yield an approximation for any discretization τ_i (not only the equally spaced τ_i 's) allows taking advantage of the EpiRK framework to reduce the number of Krylov projections and to make each projection more efficient by optimizing the coefficients c . Unlike the Taylor-based method, the adaptive Krylov algorithm does not require any estimates of the matrix norm of chJ which reduces computational cost and is preferable for matrix-free implementations. **We note, however, that for problems where the matrix is stored in memory Taylor methods or methods based on Leja interpolation can be beneficial (for example, a backward error analysis is available for these methods [1, 6]).** Since approximating the action of matrix functions is an active area of research, it is likely that more efficient algorithms, including Taylor-based methods, will emerge in the future that allow us to further reduce the computational cost of exponential integrators. At present we focus on the EpiRK methods based on an adaptive Krylov approach.

Since evaluations of $\varphi_k(chJ)v$ constitute the most computationally expensive part of the exponential integrator, special care has to be taken to ensure that each time step requires as few of these evaluations as possible and the most efficient algorithm is chosen for these approximations. The EpiRK framework [51, 52] is designed to construct integrators with a reduced number of $\varphi_k(chJ)v$ evaluations and with φ_k functions chosen to speed up the Krylov iteration. In addition, the EpiRK framework allows optimizing the coefficients c to gain more efficiency. Recall that the Krylov iteration converges faster if the spectrum of the matrix chJ is more clustered. Minimizing the coefficients c increases clustering of the spectrum of

chJ and consequently reduces the dimensionality m of the Krylov subspace for a given tolerance.

A particularly efficient class of EpiRK methods [55] also takes advantage of an adaptive version of the Krylov iteration introduced in [38]. The adaptive Krylov based EpiRK schemes link the construction of a time integrator with the choice of the method to approximate $\varphi_k(chJ)v$ to build a more efficient overall time stepping scheme in the following way.

The computational cost of the Krylov algorithm to approximate terms of type $\varphi_k(chJ)v$ scales as $O(m^2)$ with the size of the Krylov basis m . Clearly, m depends on the eigenvalues of J , the vector v as well as the values of c and h . As c or h is increased, so is the size of the basis. The EpiRK framework allows one to reduce the coefficients c and choose the better suited functions φ_k in the derivation of a time integrator. The adaptive Krylov technique allows replacing each computation of $\varphi_k(chJ)v$ with several evaluations of $\varphi_k(\tau_i chJ)v$ where $0 < \tau_i < 1$ for $i = 0, 1, 2, \dots, P$. Since the complexity of each of these evaluations now scales as $O(m_i^2)$, in most cases increased efficiency is observed as $O(m^2)$ can be larger than $O(m_1^2) + O(m_2^2) + \dots + O(m_P^2)$.

A general EpiRK method can be written as

$$Y_i = y_n + a_{i1}\psi_{i1}(g_{i1}hJ_{i1})hF_n + \sum_{j=2}^{i-1} a_{ij}\psi_{ij}(g_{ij}hJ_{ij})h\Delta^{(j-1)}R_n, \quad i = 1, \dots, s-1$$

$$y_{n+1} = y_n + b_1\psi_{s1}(g_{s1}hJ_{s1})hF_n + \sum_{j=2}^s b_j\psi_{sj}(g_{sj}hJ_{sj})h\Delta^{(j-1)}R_n$$

where $R_n(\cdot) = R(\cdot, y_n)$ and the divided differences $\Delta^{(j-1)}R_n$ are computed by using the nodes $y_n, Y_1, \dots, Y_{s-1}$ (note that since $R_n(y_n) = 0$ the node y_n does not actually appear in the divided differences). Different choices of functions ψ_{ij} and matrices J_{ij} result in different classes of EpiRK methods such as unpartitioned, partitioned or hybrid exponential or implicit-exponential (IMEXP) integrators [53]. In this paper we employ the schemes with the choice of $J_{ij} = J_n$ and ψ_{ij} set as linear combinations of φ_k functions

$$\psi_{ij}(z) = \sum_{k=1}^s p_{ijk}\varphi_k(z).$$

As in the Runge–Kutta case, to construct an efficient integrator the

coefficients a_{ij}, g_{ij}, b_j and p_{ijk} have to be chosen subject to the appropriate order conditions. In this paper, we use the variable time step fifth order method EpiRK5P1

$$\begin{aligned} Y_1 &= y_n + a_{11}\varphi_1(g_{11}hJ_n)hF_n \\ Y_2 &= y_n + a_{21}\varphi_1(g_{21}hJ_n)hF_n + a_{22}\varphi_1(g_{22}hJ_n)hR(Y_1) \\ y_1 &= y_n + b_1\varphi_1(g_{31}hJ_n)hF_n + b_2\varphi_1(g_{32}hJ_n)hR(Y_1) + b_3\varphi_3(g_{33}hJ_n)h(-2R(Y_1) + R(Y_2)), \end{aligned} \quad (7)$$

that has been derived in [55] (see Table 1 for the coefficients; this method is also employed in the performance comparison given in [31]) and the constant time step fourth order EpiRK4 method proposed in [40] given by

$$\begin{aligned} Y_1 &= y_n + \frac{1}{2}\varphi_1(\frac{1}{2}h_nJ_n)h_nF_n \\ Y_2 &= y_n + \frac{2}{3}\varphi_1(\frac{2}{3}h_nJ_n)h_nF_n \\ y_{n+1} &= y_n + \varphi_1(h_nJ_n)h_nF_n + (32\varphi_3(h_nJ_n) - 144\varphi_4(h_nJ_n))h_nR(Y_1) \\ &\quad + \left(-\frac{27}{2}\varphi_3(h_nJ_n) + 81\varphi_4(h_nJ_n)\right)h_nR(Y_2) \end{aligned} \quad (8)$$

Table 1: Coefficients of the fifth order method EPIRK5P1.

$\begin{bmatrix} a_{11} & & \\ a_{21} & a_{22} & \\ b_1 & b_2 & b_3 \end{bmatrix}$	=	$\begin{bmatrix} 0.35129592695058193092 & & \\ 0.84405472011657126298 & 1.6905891609568963624 & \\ & 1.0 & 1.2727127317356892397 & 2.2714599265422622275 \end{bmatrix}$
$\begin{bmatrix} g_{11} & & \\ g_{21} & g_{22} & \\ g_{31} & g_{32} & g_{33} \end{bmatrix}$	=	$\begin{bmatrix} 0.35129592695058193092 & & \\ 0.84405472011657126298 & 0.5 & \\ & 1.0 & 0.71111095364366870359 & 0.62378111953371494809 \end{bmatrix}$

Both of these EpiRK schemes utilize the adaptive Krylov technique for the evaluations of $\varphi_k(chJ)v$ in order to improve performance. It is important to note that while both of these are three-stage methods, EpiRK4 requires only two Krylov projections to be executed per time step since the scale invariance of the Arnoldi iteration allows the simultaneous calculation of the Krylov bases for Y_1 and Y_2 and one more Krylov projection to evaluate y_{n+1} (e.g. see [40] for details). EPIRK5P1 requires three Krylov projections to be performed each time step [55]. In addition EPIRK4 is a so-called stiffly accurate method while EPIRK5P1 is derived using classical order conditions. The advantages of the EPIRK5P1 method, however, include the higher order and the ability to design embedded methods for the automatic

time step control that do not require an additional Krylov projection. The adaptive time step control version of the methods of type EPIRK4 is still under development.

3. The resistive magnetohydrodynamics equations

The most fundamental theoretical description of a classical plasma comes from the kinetic equation. This so-called Vlasov equation (the collisionless case) or Boltzmann equation (when collisions are of physical significance) describes the time evolution of a particle-density in the $3 + 3$ dimensional phase space (the first three dimensions correspond to the space dependence while the remaining correspond to the velocity dependence of the particle-density). While a number of simulations of different plasma phenomena have been conducted with this approach, due to the high dimensionality of the phase space, a lower-dimensional approximation is usually used to render such simulations feasible. A similar discussion holds true for the gyrokinetic approximation that can be employed in plasmas where a strong external magnetic field along a given axis is present. In this case the phase space of the Vlasov equation is reduced to $3 + 2$ dimensions by averaging over the gyro motion. This procedure yields a good approximation under the assumption of low frequency as compared to the cyclotron frequency. For a more detailed discussion of gyrokinetic models, see e.g. [17] or [16].

However, in many applications (such as magnetic confinement fusion, spheromak experiments, and astrophysical plasmas) the timescales of interest are sufficiently long and/or the full three dimensional model is necessary to describe the physical phenomena. In that case the kinetic approach is usually not feasible (even on modern day supercomputers) and thus further simplifications have to be introduced. For the magnetohydrodynamics (MHD) equations, which we will describe in the remainder of this section, the assumption is made that, to a good approximation, the distribution in the velocity direction is Maxwellian; in other words, it is assumed that each sufficiently small volume in the plasma is in thermodynamic equilibrium. For a general overview of kinetic and MHD models (including a derivation of the MHD equations from the Vlasov equation) see e.g. [37]. Numerical computations in the context of the MHD model discussed in this paper are performed, for example, in [42], [43], and [41].

If the assumption is made that the plasma considered is in thermodynamic equilibrium, the equations of motion in a three dimensional phase

space (a so-called fluid model) can be derived. These are given by (in dimensionless units)

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \quad (9)$$

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho \mathbf{v} \cdot \nabla \mathbf{v} = \mathbf{J} \times \mathbf{B} - \nabla p \quad (10)$$

which we refer to as the continuity and momentum equation respectively. The equations are written in terms of the density ρ , the fluid velocity $\mathbf{v}$, the magnetic field $\mathbf{B}$, the electric current density $\mathbf{J}$, and the pressure p . The dynamics is determined by the Lorentz force (the $\mathbf{J} \times \mathbf{B}$ term) and the pressure gradient force (the ∇p term). These fluid equations have to be coupled to an appropriate model of the electric field. Note, however, that if (the ideal) Ohm's law is assumed to hold, i.e.

$$\mathbf{E} + \mathbf{v} \times \mathbf{B} = 0,$$

then the electric field $\mathbf{E}$ can be eliminated from Maxwell's equation. This leaves us with the following system

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{v} \times \mathbf{B}) \quad (11)$$

$$\mathbf{J} = \nabla \times \mathbf{B}$$

$$\nabla \cdot \mathbf{B} = 0.$$

The first of these equations yields the time evolution of $\mathbf{B}$, the second can be used to eliminate $\mathbf{J}$ from equation (10), and the third is the familiar solenoidal constraint imposed on the magnetic field.

A commonly employed approach to close these equations (see e.g. [43]) is to supplement them with the following equation of state

$$e = \frac{p}{\gamma - 1} + \frac{\rho}{2} v^2 + B^2,$$

where the time evolution of the energy density e is given by

$$\frac{\partial e}{\partial t} + \nabla \cdot \left((e + p + \frac{1}{2} B^2) \mathbf{v} - \mathbf{B}(\mathbf{B} \cdot \mathbf{v}) \right) = 0. \quad (12)$$

Collectively, the equations (9)-(12) in the variables $(\rho, \mathbf{v}, \mathbf{B}, e)$ yield a

first-order system of 8 differential equations in 8 variables and are called the ideal magnetohydrodynamics equations (or the ideal MHD equations).

For the purpose of performing the spatial discretization, these equations are often cast into the so-called divergence form. Then, the equations of motion read as (see e.g. [43])

$$\frac{\partial U}{\partial t} + \nabla \cdot F_h(U) = 0$$

with state vector

$$U = \begin{bmatrix} \rho \\ \rho \mathbf{v} \\ \mathbf{B} \\ e \end{bmatrix}$$

and

$$F_h(U) = \begin{bmatrix} \rho \mathbf{v} \\ \rho \mathbf{v} \otimes \mathbf{v} + (p + \frac{1}{2} B^2) \mathbf{I} - \mathbf{B} \otimes \mathbf{B} \\ \mathbf{v} \otimes \mathbf{B} - \mathbf{B} \otimes \mathbf{v} \\ (e + p + \frac{1}{2} B^2) \mathbf{v} - \mathbf{B} (\mathbf{B} \cdot \mathbf{v}) \end{bmatrix}, \quad (13)$$

where we have denoted the tensor product by using the $\otimes$ symbol.

In this paper, as in [42] and [43], we will consider a slightly more general class of equations which, in addition to the dynamics discussed so far, includes dissipative effects (due to particle collisions in the plasma). To that end the hyperbolic flux vector $F_h(U)$ is extended in [42] by a diffusive part given by

$$F_d(U) = \begin{bmatrix} 0 \\ Re^{-1} \boldsymbol{\tau} \\ S^{-1} (\eta \nabla \mathbf{B} - \eta (\nabla \mathbf{B})^T) \\ Re^{-1} \boldsymbol{\tau} \cdot \mathbf{v} + \frac{\gamma}{\gamma-1} Re^{-1} Pr^{-1} \nabla T + S^{-1} (\frac{1}{2} \nabla (\mathbf{B} \cdot \mathbf{B}) - \mathbf{B} (\nabla \mathbf{B})^T) \end{bmatrix} \quad (14)$$

with

$$\boldsymbol{\tau} = \nabla \mathbf{v} + (\nabla \mathbf{v})^T - \frac{2}{3} \nabla \cdot \mathbf{v} \mathbf{I},$$

where we have assumed a spatially homogeneous viscosity μ , resistivity η , and thermal conductivity κ . The non-dimensional parameters governing the solution of this system are the Reynolds number $Re = \rho_0 v_A l_0 / \mu_0$, the Lundquist number $S = \mu_0 v_A l_0 / \eta_0$ and the Prandtl number $Pr = c_p \mu_0 / \kappa_0$ for a characteristic density ρ_0 , a characteristic length scale l_0 , and the Alfvén velocity v_A (as usual the permeability of free space is denoted by μ_0 , $\kappa_0 =$

5/3, and c_p is the specific heat of the fluid).

Note that in all the simulations we use dimensionless units. Thus, we can rewrite (14) more conveniently in terms of the (dimensionless) viscosity $\mu = Re^{-1}$, the (dimensionless) resistivity $\eta = S^{-1}$, and the (dimensionless) thermal conductivity $\kappa = Pr^{-1}$. Then, the form of the resistive MHD equations used for the spatial discretization is

$$\frac{\partial U}{\partial t} + \nabla \cdot F_h(U) = \nabla \cdot F_d(U)$$

with $F_h(U)$ given by equation (13) and where $F_d(U)$ is given by

$$F_d(U) = \begin{bmatrix} 0 \\ \mu \boldsymbol{\tau} \\ \eta (\nabla B - (\nabla B)^T) \\ \mu \boldsymbol{\tau} \cdot \mathbf{v} + \frac{\gamma \mu \kappa}{\gamma - 1} \nabla T + \eta \left(\frac{1}{2} \nabla (\mathbf{B} \cdot \mathbf{B}) - B (\nabla \mathbf{B})^T \right) \end{bmatrix}.$$

For the numerical simulations conducted in this paper, we employ the MHD code developed in [13]. This C++ code is based on the Fortran code developed in [42] which has been used to conduct plasma physics simulations (see e.g. [42], [43], and [41]) as well as to construct more efficient preconditioners in the context of implicit time integrators (see [43]). The implementation of exponential integrators requires the evaluation of the action of the Jacobian applied to a vector. We compute this action by employing a simple forward difference stencil (this has the advantage that one function evaluation can be reused). Consequently, the present implementation is completely matrix-free. For more details we refer the reader to section 6.

With respect to space discretization, our implementation assumes a finite difference or finite volume method where a single value is stored in each cell. Thus, in each cell we store the value of the density ρ , the fluid velocity $\mathbf{v}$, the magnetic field $\mathbf{B}$, the energy e (but not the pressure p). For the numerical simulations conducted in this paper, we have implemented the 2.5-dimensional case (that is, the state variables only depend on the x - and y -direction but both the velocity $\mathbf{v}(x, y)$ and the magnetic field $\mathbf{B}(x, y)$ are three-dimensional vectors) using a classic centered stencil for the divergence; that is, for each vector G , corresponding to the flux of a given (scalar) state variable, we compute the following approximation of the divergence

$$\nabla \cdot G(U) \approx \frac{G(U_{i+1,j}) - G(U_{i-1,j})}{2h} + \frac{G(U_{i,j+1}) - G(U_{i,j-1})}{2h},$$

where i and j are the cell indices in the x - and y -directions, respectively and h is the cell size. To evaluate the spatial derivatives present in the diffusion vector F_d , we also employ the classic centered stencil. Thus, the code has order two accuracy in space.

It is well known that for MHD problems preserving the divergence free property of the magnetic field is important. The space discretization used in our implementation ensures that if the solenoidal property $\nabla \cdot \mathbf{B} = 0$ is satisfied for the initial value, then this is true for all later times. In [43] it is shown that this property also holds true if we couple the space discretization with an implicit method that is based on a matrix-free inexact Newton–Krylov algorithm.

We will now show that an equivalent result holds for the exponential integrators considered in this paper. That is, if the discrete solenoidal property

$$\frac{B_{i+1,j}(t) - B_{i-1,j}(t)}{2h} + \frac{B_{i,j+1}(t) - B_{i,j-1}(t)}{2h} = 0$$

is satisfied for the initial value this is also true for all later times (independent of the step size and the tolerance specified for the computation of the φ_k functions). Now, it is well known that $F(U)$ satisfies the discrete solenoidal condition if U does (see, for example, [43]). In addition, since the Jacobian J is computed using a finite difference approximation, the solenoidal property also holds for J and for $F(U) - J$. This together with the fact that the solenoidal property is a linear invariant (i.e. the set of all states that satisfy the discrete solenoidal constraint forms a vector space which we denote by $\mathcal{V}$) allows us to apply Theorem 2 in [15] which implies that any exponential Runge–Kutta method satisfies the solenoidal property. The only issue here is that the cited theorem assumes that the computation of $\varphi_k(J)v$ preserves the linear constraint in question. However, since we use a Krylov approach the term $\varphi_k(J)v$ is approximated by $p(J)v$, where p is a polynomial. Since we already know that $\mathcal{V}$ is an invariant subspace with respect to J , the desired result follows immediately.

4. The reconnection problem

The examples in this and the next section are drawn from [42] and [43], respectively. We start with a reconnection problem for which the initial value of the magnetic field is given by

$$B_0(x, y, z) = \begin{bmatrix} \tanh(2y) - \psi_0 k_y \cos(k_x x) \sin(k_y y) \\ \psi_0 k_x \sin(k_x x) \cos(k_y y) \\ 0 \end{bmatrix},$$

where as in [42] we have chosen $k_x = \pi/x_r$, $k_y = \pi/(2y_r)$, $\psi_0 = 0.1$, and the computational domain is given by $[-x_r, x_r] \times [-y_r, y_r]$ for $x_r = 12.8$ and $y_r = 6.4$. This implies that the magnetic field reverses direction from pointing along $\mathbf{e}_x$ to pointing along $-\mathbf{e}_x$ abruptly at $y = 0$. Furthermore, at time $t = 0$, we impose a density that is given by

$$\rho = 1.2 - \tanh^2(2y)$$

and a pressure that is proportional to the density; to be more precise $p = 0.5\rho$ (from which the energy is determined). A vanishing velocity in both space directions is prescribed.

The time evolution for the current is shown in Figure 1. In all the simulations conducted the relative tolerance, the absolute tolerance, and the tolerance for the Krylov iteration are equal. We have investigated the effect of varying the tolerance for the Krylov iteration and found no significant difference in accuracy or run time as long as we choose a tolerance for the Krylov iteration that is at least as small as the tolerance for the time integration. In addition, Figure 1 shows that the value of $\nabla \cdot B$ is (almost) independent of the tolerance chosen and therefore that the error made is only due to numerical round-off.

First, we consider the configuration that is investigated in [42]; that is, we employ 256 grid points in the x -direction and 128 grid points in the y -direction. The results obtained are shown in Figure 2. We observe that the performance of the EPIRK5P1 method is superior to CVODE if the desired accuracy is equal or less than 10^{-4} . In the low precision regime that is often of interest in applications, EPIRK5P1 outperforms CVODE by almost a factor of 2. However, for more stringent tolerances the CVODE implementation is significantly more efficient. This is due to the fact that the run time for the CVODE method is almost independent of the accuracy achieved. What is perhaps even more interesting is that for a specified tolerance of 10^{-4} or above the CVODE method does not converge (in the situation described a final step size on the order of 10^{-16} is reported by the application). Let us also note that the constant time step method EPIRK4 is not competitive for this problem.

Next, we investigate the effect of increasing the number of grid points.

Figure 1: The current $J = \nabla \times B$ for the reconnection problem is shown at times $t = 0$ (top-left), $t = 100$ (top-right), $t = 150$ (middle-left), and $t = 200$ (middle-right). We employ 256 grid points in the x -direction and 128 grid points in the y -direction. The dimensionless parameters are chosen as follows: $\mu = 10^{-2}$, $\eta = 10^{-3}$, and $\kappa = 10^{-2}$. The absolute and relative tolerance for EPIRK5P1 is set to 10^{-6} . On the bottom graph $\nabla \cdot B$ is shown as a function of time for a variety of tolerances.

Figure 2: The maximum error of the components of the numerical solution (measured using the l^2 norm in space at the final time) is shown as a function of the run time. The reference solution for the CVODE/EpiRK is computed using the EpiRK/CVODE method with a tolerance of 10^{-14} . The simulation is conducted up to a final time of $T = 100$. We have used 256 grid points in the x -direction and 128 grid points in the y -direction (this corresponds to the configuration investigated in [42]). The viscosity is given by $\mu = 5 \cdot 10^{-2}$, the resistivity by $\eta = 5 \cdot 10^{-3}$, and the thermal conductivity by $\kappa = 4 \cdot 10^{-2}$. CVODE tolerance: 10^{-4} , 10^{-5} , 10^{-7} , 10^{-8} , 10^{-9} . EPIRK5P1 tolerance: 10^{-1} , $2 \cdot 10^{-2}$, 10^{-2} , 10^{-3} , 10^{-5} , 10^{-7} , 10^{-9} . EPIRK4 step size: 2, 1, $4 \cdot 10^{-1}$, $2 \cdot 10^{-1}$, 10^{-1} .

Figure 3: The maximum error of the components of the numerical solution (measured using the l^2 norm in space at the final time) is shown as a function of the run time. The reference solution for the CVODE/EpiRK is computed using the EpiRK/CVODE method with a tolerance of 10^{-11} . The simulation is conducted up to a final time of $T = 20$. We have used 512 grid points in the x -direction and 256 grid points in the y -direction. The viscosity is given by $\mu = 5 \cdot 10^{-2}$, the resistivity by $\eta = 5 \cdot 10^{-3}$, and the thermal conductivity by $\kappa = 4 \cdot 10^{-2}$. CVODE tolerance: $10^{-3}, 10^{-4}, \dots, 10^{-9}$. EPIRK5P1 tolerance: $10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-8}, 10^{-9}$. EPIRK4 step size: 4, 2, 1, $4 \cdot 10^{-1}, 2 \cdot 10^{-1}, 10^{-1}$.

The results obtained with 512 grid points in the x -direction and 256 grid points in the y -direction are shown Figure 3. In this case almost identical conclusions can be drawn. Once again the performance of EPIRK5P1 is superior to CVODE library if the desired accuracy is equal or less than 10^{-4} . Furthermore, the constant step size method EPIRK4 is not competitive except for very stringent tolerance requirements.

Up to this point we have only considered relatively large values for the viscosity, resistivity, and the thermal conductivity. Thus, we proceed by decreasing these dimensionless quantities by a factor of five. The corresponding results are shown in Figure 4. In this case the EPIRK5P1 implementation outperforms the CVODE implementation for all the tolerances studied here (although the difference in performance between the two methods in the low precision regime is smaller than was the case for the previous two examples). For accuracies between 10^{-6} and 10^{-8} the constant time step method EPIRK4 manages to outperform both the EPIRK5P1 as well

Figure 4: The maximum error of the components of the numerical solution (measured using the l^2 norm in space at the final time) is shown as a function of the run time. The reference solution for the CVODE/EpiRK is computed using the EpiRK/CVODE method with a tolerance of 10^{-14} . The simulation is conducted up to a final time $T = 100$. We have used 256 grid points in the x -direction and 128 grid points in the y -direction. The viscosity is given by $\mu = 10^{-2}$, the resistivity by $\eta = 10^{-3}$, and the thermal conductivity by $\kappa = 10^{-2}$. CVODE tolerance: $10^{-4}, 10^{-5}, \dots, 10^{-9}$. EPIRK5P1 tolerance: $10^{-1}, 10^{-2}, 10^{-3}, 10^{-5}, 10^{-7}, 10^{-8}$. EPIRK4 step size: $2, 1, 2 \cdot 10^{-1}, 10^{-1}$. Note that CVODE does not converge for tolerances that are less tight than 10^{-4} .

as CVODE implementation.

We observe that for larger Reynolds and Lundquist numbers the EpiRK method shows better or identical performance for the range of tolerances studied here. In addition, the EpiRK method shows a clear advantage over the CVODE implementation if a precision equal to or lower than 10^{-5} is required.

In summary, it can be said that for the reconnection problem the EPIRK5P1 implementation shows superior performance in the loose and medium tolerance range that is of interest for the majority of practical applications. However, for some of the examples considered here the CVODE implementation has a significant advantage for more stringent tolerance requirements.

To conclude this section let us investigate the strong scaling (i.e. the problem size is kept fixed while the number of cores is increased) for both our time integration routines and the CVODE implementation. This allows us to ascertain how well the current implementation is parallelized. In the

Figure 5: The run time (normalized to the single core run time for both the CVODE and EpiRK implementation) is shown as a function of the number of cores. The simulation is conducted up to a final time $T = 1$. We have used 512 grid points in both the x - and y -direction. The viscosity is given by $\mu = 5 \cdot 10^{-2}$, the resistivity by $\eta = 5 \cdot 10^{-3}$, and the thermal conductivity by $\kappa = 4 \cdot 10^{-2}$. These performance measurements were done on the VSC-2 (Vienna Scientific Cluster 2; see <http://vsc.ac.at/systems/vsc-2/>) which uses two 8 core AMD Opteron Magny Cours 6132HE and 32 GB of DDR3 memory per node. The nodes are connected by Infiniband QDR and the Intel MPI library (version 4.1) is employed. The line of ideal scaling is shown in black.

ideal case we would observe a linear decrease in the run time as a function of the number of cores.

The strong scaling results on the VSC-2 are shown in Figure 5 and are almost identical for both implementations. We conclude that the algorithms scale well up to approximately 256 cores. No further gain in performance can be observed for 1024 or more cores. This is to be expected as in the case of strong scaling the amount of work available to a single core eventually decreases sufficiently such that communication between the different MPI processes limits the performance of the application. At this point no further increase in performance can be expected.

parameter	value
ϵ_x, ϵ_y	0.1
v_0	0.5
ω_x, ω_y	2
pressure p	0.25
B_z	10
B_x	0.1

Table 2: Parameters for the initial value of the Kelvin–Helmholtz instability.

5. The Kelvin–Helmholtz instability

As a second example we consider the Kelvin–Helmholtz (KH) instability. The KH instability is triggered by superimposing the following perturbation

$$\epsilon_x \cos\left(\frac{2\pi\omega_x}{L_x}x\right) + \epsilon_y \sin\left(\frac{\pi(2\omega_y - 1)}{L_y}y\right)$$

in the x -direction on the velocity field

$$\mathbf{v} = \begin{bmatrix} v_0 \tanh(\frac{y}{\lambda}) \\ 0 \\ 0 \end{bmatrix}.$$

The density is initialized to unity and a uniform pressure is chosen. The magnetic field is initialized to be uniform in both the x - and z -direction with strength B_x and B_z , respectively, and is assumed to vanish in the y -direction. All the parameters used to determine the numerical value of the initial value are listed in Table 2.

The time evolution of the numerical solution is shown in Figure 6. In all the simulations conducted the relative tolerance, the absolute tolerance, and the tolerance for the Krylov iteration are equal. We have investigated the effect of varying the tolerance for the Krylov iteration and found no significant difference in accuracy or run time as long as we choose a tolerance for the Krylov iteration that is at least as small as the tolerance for the time integration.

First, we perform a numerical simulation of the Kelvin–Helmholtz instability for 128 grid points in both space directions. The results obtained are shown in Figure 7. Both exponential integrators considered here (EPIRK5P1 and EPIRK4) outperform the CVODE implementation by a

Figure 6: For the Kelvin–Helmholtz instability the velocity in the x -direction is shown at time $t = 0$ (top-left), $t = 1.5$ (top-right), $t = 2$ (bottom-left), and at $t = 2.5$ (bottom-right). In both space directions 256 grid points are employed. The dimensionless parameters are chosen as follows: $\mu = 10^{-4}$, $\eta = 10^{-4}$, and $\kappa = 10^{-4}$. The absolute and relative tolerance for EPIRK5P1 is set to 10^{-6} .

significant margin (approximately a factor of two). Furthermore, CVODE is very sensitive to even small perturbations in the chosen tolerance, while the exponential integrators are much more robust in this regard. Thus, a user who would naively (i.e. without measuring the run time for a range of tolerances) choose a specific tolerance would increase the run time of the application by significantly more than the factor of two stated earlier. What is also interesting to see is that the constant step size method EPIRK4 is equal in performance or outperforms EPIRK5P1 for all tolerances studied here (although the difference between these two methods is at most 30%).

Now, let us increase the number of grid points to 256 (in both space directions). The corresponding results are shown in Figure 8. As before, we observe that both exponential integrators significantly outperform the CVODE implementation. In addition, the constant step size method EPIRK4 outperforms the EPIRK5P1 implementation by up to 30%. This is also consistent with the results on the coarse grid considered before. Let us emphasize here again that the major advantage of EPIRK4 is that it uses fewer Krylov projections (two vs. three for EPIRK5P1) per time step. However, at present this advantage over EPIRK5P1 is limited to constant time step implementations. A third projection is needed to construct automatic time stepping control schemes for EPIRK4. The variable time stepping version of the integrators is more efficient for EPIRK5P1 since it is a higher order method. For the Kelvin–Helmholtz instability, considered in this section, these differences work in favor of the EPIRK4 method while for the reconnection problem, considered in the previous section, the EPIRK5P1 scheme is superior (in the case of [loose](#) and medium tolerances).

In summary, both exponential integrators considered in this paper show significantly improved performance compared to CVODE for the Kelvin–Helmholtz instability. Let us also mention that for both the exponential integrators as well as CVODE it is not entirely straightforward to choose a good tolerance. This is due to the fact that if the tolerance is chosen too [loose](#), the run time can actually increase. This is due to the deficiency of the step size controller which assumes that the computational cost of each step is independent of the step size (a reasonable assumption for explicit time integrators; however, this is not true for the implicit and exponential methods considered here). We emphasize, however, that this behavior is significantly more serious for the CVODE implementation.

Figure 7: The maximum error of the components of the numerical solution (measured using the l^2 norm in space at the final time) is shown as a function of the run time. The reference solution for the CVODE/EpiRK is computed using the EpiRK/CVODE method with a tolerance of 10^{-11} . The simulation is conducted up to a final time $T = 1$ (top) and $T = 2$ (bottom) and 128 grid points are used in both the x - and y -direction. The (dimensionless) viscosity is given by $\mu = 10^{-4}$, the resistivity by $\eta = 10^{-4}$, and the thermal conductivity by $\kappa = 10^{-4}$. CVODE tolerance ($T = 1$): 10^{-3} , 10^{-6} , 10^{-7} , 10^{-8} , 10^{-9} . CVODE tolerance ($T = 2$): 10^{-6} , 10^{-7} , 10^{-8} , 10^{-9} . EPIRK5P1 tolerance ($T = 1$): 10^{-3} , 10^{-4} , $\dots$, 10^{-9} . EPIRK5P1 tolerance ($T = 2$): 10^{-4} , 10^{-5} , $\dots$, 10^{-9} . EPIRK4 step size ($T = 1$): 10^{-1} , $4 \cdot 10^{-2}$, $2 \cdot 10^{-2}$, 10^{-2} , $4 \cdot 10^{-3}$, $2 \cdot 10^{-3}$, 10^{-3} . EPIRK4 step size ($T = 2$): $4 \cdot 10^{-2}$, $2 \cdot 10^{-2}$, 10^{-2} , $4 \cdot 10^{-3}$, $2 \cdot 10^{-3}$.

Figure 8: The maximum error of the components of the numerical solution (measured using the l^2 norm in space at the final time) is shown as a function of the run time. The reference solution for the CVODE/EpiRK is computed using the EpiRK/CVODE method with a tolerance of 10^{-11} . The simulation is conducted up to a final time $T = 1$ and 256 grid points are used in both the x - and y -direction. The (dimensionless) viscosity is given by $\mu = 10^{-4}$, the resistivity by $\eta = 10^{-4}$, and the thermal conductivity by $\kappa = 10^{-4}$. CVODE tolerance: $10^{-5}, 10^{-6}, \dots, 10^{-9}$. EPIRK5P1 tolerance: $10^{-4}, 10^{-5}, \dots, 10^{-9}$. EPIRK4 step size: $4 \cdot 10^{-2}, 2 \cdot 10^{-2}, 10^{-2}, 4 \cdot 10^{-3}, 2 \cdot 10^{-3}, 10^{-3}$.

6. Implementation

We have chosen to base our implementation on the 2.5 dimensional¹ MHD code developed by Daniel R. Reynolds et al. A large number of numerical simulations have been conducted (see e.g. [42] and [43]) and the performance of preconditioners for Newton–Krylov based implicit methods has been investigated (see e.g. [43]). Furthermore, the extension of the code to three dimensional problems as well as to non-square geometries (such as a Tokamak geometry) have been investigated in [41].

Our code is implemented in C++ and is designed to easily accommodate different time integrators, space discretization schemes, as well as different MHD models. The design considerations and a detailed description of the computer code can be found in [13].

The second component of our code is the EPIC library (see [31]) which implements the fifth order EpiRK5P1 with an adaptive time step control

¹All quantities depend on only two spatial variables; however, the direction of the velocity as well as the magnetic fields are three dimensional vectors.

(*EpiRK5P1VerticalVar* function in EPIC library), the fourth order EpiRK4 method (*EpiRK4MixedConst* function in EPIC library) with constant time steps and the adaptive Krylov algorithm to evaluate the φ_k functions as described in [38]. This library is written in the C++ programming language with MPI, and uses routines from the BLAS and LAPACK libraries. In the experiments, the BLAS and LAPACK implementations were supplied by the Intel Math Kernel Library (Intel MKL).

One detail warrants further discussion: in the original Fortran program the CVODE interface is used to approximate the Jacobian by a simple forward difference stencil. In our implementation, however, we provide a custom function to compute an approximation to the Jacobian. This is necessary for the EPIC library, which before was only used in the context of problems where an exact Jacobian is available (as in [31]). Computing an exact Jacobian for the significantly more complex MHD problem considered in this paper is infeasible. Therefore, we have to approximate $J(a)v$, i.e. the application of the Jacobian J at position a to a vector v , where both a and v depend on the specific numerical algorithm as well as the initial values. We note here that there is a qualitative difference in the norm of the vector v that depends on the numerical time integration scheme used. In the EpiRK the norm of v is close to unity whereas in the BDF method the value is often significantly below $\sqrt{\epsilon}$ (where we use $\epsilon \approx 10^{-16}$ to denote machine precision). Therefore, care is taken to scale vectors to the norm $\sqrt{\epsilon}$ only if the initial norm is above $\sqrt{\epsilon}$ in magnitude. Then the same implementation can be used for both methods and a difference in performance due to an internally optimized function to compute the Jacobian (as is provided by the CVODE library) is precluded.

In all of the computations conducted in this paper, we have used the infinity norm to scale v to its appropriate size. In general, we have found that the performance and accuracy of the computation is not significantly altered if the scaling is performed to some value that is reasonably close to $\sqrt{\epsilon}$.

7. Conclusion

Using several examples of MHD problems, we have shown that exponential integrators constitute a viable alternative to the more commonly employed BDF method (as implemented in the CVODE library). In almost all instances equal or superior performance has been observed for the

adaptive and variable time stepping fifth order EpiRK method for low to medium accuracy requirements.

Note that while the first version of the CVODE package was released several decades ago and the package has been optimized and refined over extended period of time, the EPIC package is the first C++/MPI implementation of the exponential integrators. As indicated in [32] there are a number of additional optimizations possible both from an algorithmic and from a computer science perspective. For example, significant performance gains are expected from improvements in the adaptivity algorithms in the Krylov projections and the derivation of the time integrators specifically adapted to a given system of equations. In addition, the newly introduced implicit-exponential (IMEXP) methods [40] demonstrate that exponential integration can also be used as a component of the overall integration scheme and can be combined with the preconditioned implicit integrators as well as time splitting strategies. It is expected that further development will lead to significant performance improvements.

Furthermore, we note that the results presented here are somewhat different from the simpler models considered in [32]. In that case order of magnitude speedups are observed for the EpiRK method as compared to the CVODE implementation. Further investigations are needed to determine whether the degradation in comparative performance is due to the specifics of the PDE system, the details of the integrator or the particulars of the implementation. Also it is not clear to what degree the approximate computation of the Jacobian contributes to the observed differences (in all the problems investigated in [32] an analytical form of the Jacobian has been used).

Acknowledgements

We would like to take the opportunity to thank D. R. Reynolds for providing the code of the Fortran MHD solver and for the helpful discussion. Dr. Einkemmer was supported by the Marshall plan scholarship of the Austrian Marshall plan foundation (<http://www.marshallplan.at/>) and by the Fonds zur Förderung der Wissenschaften (FWF) – project id: P25346. Dr. Tokman was supported by a grant from the National Science Foundation, Computational Mathematics Program, under Grant No. 1115978.

The simulations were conducted using the Vienna Scientific Cluster (VSC).

References

- [1] A.H. Al-Mohy and N.J. Higham, *Computing the action of the matrix exponential, with an application to exponential integrators*, SIAM J. Sci. Comput. **33** (2011), no. 2, 488–511.
- [2] W.E. Arnoldi, *The principle of minimized iteration in the solution of the matrix eigenvalue problem*, Quart. Appl. Math. **9** (1951), no. 1, 17–29.
- [3] L. Bergamaschi, M. Caliari, and M. Vianello, *The ReLPM Exponential Integrator for FE Discretizations of Advection-Diffusion Equations*, Computational Science – ICCS 2004: 4th International Conference, Kraków, Poland, Springer, 2004, pp. 434–442.
- [4] J. Breslau, N. Ferraro, and S. Jardin, *Some properties of the M3D-C1 form of the three-dimensional magnetohydrodynamics equations*, Phys. Plasmas **16** (2009), no. 9, 822–838.
- [5] M. Caliari, P. Kandolf, A. Ostermann, and S. Rainer, *Comparison of software for computing the action of the matrix exponential*, BIT **54** (2014), no. 1, 113–128.
- [6] ———, *The Leja method revisited: backward error analysis for the matrix exponential*, SIAM J. Sci. Comput. **38** (2016), no. 3, A1639–A1661.
- [7] M. Caliari and A. Ostermann, *Implementation of exponential Rosenbrock-type integrators*, Appl. Numer. Math. **59** (2009), no. 3–4, 568–581.
- [8] M. Caliari, M. Vianello, and L. Bergamaschi, *Interpolating discrete advection–diffusion propagators at Leja sequences*, J. Comput. Appl. Math. **172** (2004), no. 1, 79 – 99.
- [9] L. Chacón, *An optimal, parallel, fully implicit Newton–Krylov solver for three-dimensional viscoresistive magnetohydrodynamics*, Phys. Plasmas **15** (2008), no. 5, 056103.
- [10] S.D. Cohen and A.C. Hindmarsh, *CVODE, a stiff/nonstiff ODE solver in C*, Comput. Phys. **10** (1996), no. 2, 138–143.
- [11] H. A. Van der Vorst, *An iterative solution method for solving $f(A) = b$ using Krylov subspace information obtained for the symmetric positive definite matrix A* , J. Comput. Appl. Math. **18** (1987), no. 2, 249–263.
- [12] F. Ebrahimi and R. Raman, *Plasmoids formation during simulations of coaxial helicity injection in the national spherical torus experiment*, Phys. Rev. Lett. **114** (2015), no. 20, 205003.
- [13] L. Einkemmer, *A modern resistive magnetohydrodynamics solver using C++ and the Boost library*, Comput. Phys. Commun. **206** (2016), 69–77.
- [14] L. Einkemmer and A. Ostermann, *Exponential integrators on graphic processing units*, High Performance Computing and Simulation HPCS, 2013 International Conference on High Performance Computing and Simulation, arXiv preprint arXiv:1603.07008 (2013).
- [15] ———, *A splitting approach for the Kadomtsev–Petviashvili equation*, J. Comput. Phys. **299** (2015), 716–730.
- [16] M. R. Fahey and J. Candy, *Gyro: A 5-D Gyrokinetic-Maxwell solver*, Proceedings of the 2004 ACM/IEEE Conference on Supercomputing (Washington, DC, USA), SC ’04, IEEE Computer Society, 2004, pp. 26–.
- [17] E. Frieman and L. Chen, *Nonlinear gyrokinetic equations for low-frequency electromagnetic waves in general plasma equilibria*, Phys. Fluids **25** (1982), 502.

- [18] B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, and H. Tufo, *FLASH: an adaptive mesh hydrodynamics code for modeling astrophysical thermonuclear flashes*, *Astrophys. J. Suppl. S.* **131** (2000), no. 1, 273.
- [19] H. Goedbloed and S. Poedts, *Principles of magnetohydrodynamics with applications to laboratory and astrophysical plasmas.*, Cambridge University Press, 2004.
- [20] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, and C. S. Woodward, *SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers*, *ACM Trans. Math. Softw.* **31** (2005), no. 3, 363–396.
- [21] M. Hochbruck and Ch. Lubich, *On Krylov subspace approximations to the matrix exponential operator*, *SIAM J. Numer. Anal.* **34** (1997), 1911–1925.
- [22] ———, *Exponential integrators for quantum-classical molecular dynamics*, *BIT Num. Math.* **39** (1999), no. 4, 620–645.
- [23] M. Hochbruck and A. Ostermann, *Explicit exponential Runge–Kutta methods for semilinear parabolic problems*, *SIAM J. Numer. Anal.* **43** (2005), no. 3, 1069–1090.
- [24] M. Hochbruck and A. Ostermann, *Exponential integrators*, *Acta Numerica* **19** (2010), 209–286.
- [25] S.C. Jardin, *Review of implicit methods for the magnetohydrodynamic description of magnetically confined plasmas*, *J. Comput. Phys.* **231** (2012), no. 3, 822 – 838.
- [26] A.K. Kassam and L.N. Trefethen, *Fourth-order time stepping for stiff PDEs*, *SIAM J. Sci. Comput.* **26** (2005), no. 4, 1214–1233.
- [27] R. Keppens, G. Toth, M.A. Botchev, and A. Van Der Ploeg, *Implicit and semi-implicit schemes: Algorithms*, *Int. J. Numer. Meth. Fl.* **30** (1999), no. 3, 335–352.
- [28] C. Klein and K. Roidot, *Fourth order time-stepping for Kadomtsev-Petviashvili and Davey-Stewartson equations*, *SIAM J. Sci. Comput.* **33** (2011), no. 6, 3333–3356.
- [29] T. Kuwabara, K. Shibata, T. Kudoh, and R. Matsumoto, *The acceleration mechanism of resistive magnetohydrodynamic jets launched from accretion disks*, *Astrophys. J.* **621** (2005), no. 2, 921.
- [30] I. Lindemuth and J. Killeen, *Alternating direction implicit techniques for two-dimensional magnetohydrodynamic calculations*, *J. Comput. Phys.* **13** (1973), no. 2, 181–208.
- [31] J. Loffeld and M. Tokman, *Comparative performance of exponential, implicit, and explicit integrators for stiff systems of ODEs*, *J. Comput. Appl. Math.* **241** (2012), 45–67.
- [32] ———, *Implementation of parallel adaptive-Krylov exponential solvers for stiff problems*, *SIAM J. Sci. Comput.* **36** (2014), no. 5, C591–C616.
- [33] H. Lütjens and J. Luciani, *XTOR-2F: a fully implicit Newton-Krylov solver applied to nonlinear 3D extended MHD in tokamaks*, *J. Comput. Phys.* **229** (2010), no. 21, 8130–8143.
- [34] A. Martinez, L. Bergamaschi, M. Caliari, and M. Vianello, *A massively parallel exponential integrator for advection-diffusion models*, *J. Comput. Appl. Math.* **231** (2009), no. 1, 82–91.
- [35] C.B. Moler and C.F. Van Loan, *Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later*, *SIAM Rev.* **45** (2003), no. 1, 3–49.
- [36] I. Moret and P. Novati, *RD-rational approximation of the matrix exponential operator*, *BIT* **44** (2014), 595–615.

- [37] D.R. Nicholson, *Introduction to plasma theory*, Cambridge University Press, 1983.
- [38] J. Niesen and W.M. Wright, *Algorithm 919: a Krylov subspace algorithm for evaluating the φ -functions appearing in exponential integrators*, ACM T. Math. Software **38** (2012), no. 3, 22.
- [39] G. Rainwater and M. Tokman, *A new class of split exponential propagation iterative methods of Runge–Kutta type (sEPIRK) for semilinear systems of ODEs*, J. Comput. Phys. **269** (2014), 40–60.
- [40] ———, *A new approach to constructing efficient stiffly accurate EPIRK methods*, to appear in J. Comput. Phys. (2016).
- [41] D. R. Reynolds, R. Samtaney, and H. C. Tiedeman, *A fully implicit Newton–Krylov–Schwarz method for tokamak magnetohydrodynamics: Jacobian construction and preconditioner formulation*, Comput. Sci. Discov. **5** (2012), no. 1, 014003.
- [42] D. R. Reynolds, R. Samtaney, and C. S. Woodward, *A fully implicit numerical method for single-fluid resistive magnetohydrodynamics*, J. Comput. Phys. **219** (2006), no. 1, 144 – 162.
- [43] D.R. Reynolds, R. Samtaney, and C.S. Woodward, *Operator-based preconditioning of stiff hyperbolic systems*, SIAM J. Sci. Comput. **32** (2010), no. 1, 150–170.
- [44] Y. Saad, *Analysis of some Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal. **29** (1992), 209–228.
- [45] ———, *Iterative methods for sparse linear systems*, PWS Publishing Company, 1996.
- [46] D. Schnack and J. Killeen, *Nonlinear, two-dimensional magnetohydrodynamic calculations*, J. Comput. Phys. **35** (1980), no. 1, 110 – 145.
- [47] R.B. Sidje, *Expokit: A software package for computing matrix exponentials*, ACM Trans. Math. Softw. **24** (1998), no. 1, 130–156.
- [48] C.R. Sovinec, A.H. Glasser, T.A. Gianakon, D.C. Barnes, R.A. Nebel, S.E. Kruger, S.J. Plimpton, A. Tarditi, M.S. Chu, and the NIMROD Team, *Nonlinear magnetohydrodynamics simulation using high-order finite elements*, J. Comp. Phys. **195** (2004), 355–386.
- [49] C.R. Sovinec and J.R. King, *Analysis of a mixed semi-implicit/implicit algorithm for low-frequency two-fluid plasma modeling*, J. Comput. Phys. **229** (2010), no. 16, 5803 – 5819.
- [50] M. Tokman, *Magnetohydrodynamic modeling of solar magnetic arcades using exponential propagation methods*, Ph.D. thesis, California Institute of Technology, 2001.
- [51] M. Tokman, *Efficient integration of large stiff systems of ODEs with exponential propagation iterative (EPI) methods*, J. Comput. Phys. **213** (2006), no. 2, 748–776.
- [52] ———, *A new class of exponential propagation iterative methods of Runge–Kutta type (EPIRK)*, J. Comput. Phys. **230** (2011), no. 24, 8762–8778.
- [53] ———, *Four classes of exponential EPIRK integrators*, Oberwolfach Reports **14** (2014), 855–858.
- [54] M. Tokman and P. M. Bellan, *Three-dimensional model of the structure and evolution of coronal mass ejections*, Astrophys. J. **567** (2002), no. 2, 1202.
- [55] M. Tokman, J. Loffeld, and P. Tranquilli, *New adaptive exponential propagation iterative methods of Runge–Kutta type (EPIRK)*, SIAM J. Sci. Comput. **34** (2012), no. 5, A2650–A2669.
- [56] J. van den Eshof and M. Hochbruck, *Preconditioning Lanczos approximations to the*

- matrix exponential*, SIAM J. Sci. Comp. **27** (2006), no. 4, 1438–1457.
- [57] C. Zanni, A. Ferrari, R. Rosner, G. Bodo, and S. Massaglia, *MHD simulations of jet acceleration from Keplerian accretion disks - The effects of disk resistivity*, Astron. Astrophys. **469** (2007), no. 3, 811–828.