



Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam



On the stability of exponential integrators for non-diffusive equations[☆]



Tommaso Buvoli ^{a,*}, Michael L. Minion ^b

^a University of California, Merced, CA 95343, USA

^b Lawrence Berkeley National Lab, Berkeley, CA 94720, USA

ARTICLE INFO

Article history:

Received 12 July 2021

Received in revised form 16 January 2022

MSC:
65L04
65L05
65L06
65L07

Keywords:
Exponential integrators
Linear stability analysis
Non-diffusive equations
Repartitioning
Hyperviscosity

ABSTRACT

Exponential integrators are a well-known class of time integration methods that have been the subject of many studies and developments in the past two decades. Surprisingly, there have been limited efforts to analyze their stability and efficiency on non-diffusive equations to date. In this paper we apply linear stability analysis to showcase the poor stability properties of exponential integrators on non-diffusive problems. We then propose a simple repartitioning approach that stabilizes the integrators and enables the efficient solution of stiff, non-diffusive equations. To validate the effectiveness of our approach, we perform several numerical experiments that compare partitioned exponential integrators to unmodified ones. We also compare repartitioning to the well-known approach of adding hyperviscosity to the equation right-hand-side. Overall, we find that the repartitioning restores convergence at large timesteps and, unlike hyperviscosity, it does not require the use of high-order spatial derivatives.

© 2022 Elsevier B.V. All rights reserved.

1. Introduction

In the past two decades, exponential integrators [1] have emerged as an attractive alternative to both fully implicit and semi-implicit methods [2–5]. The growing interest in exponential integrators is jointly driven by new method families [1,3,6–14] and by advances in algorithms for computing the associated exponential functions [15–25].

The purpose of this paper is to study the stability limitations of exponential integrators on problems with no diffusion, and to propose a repartitioning strategy that improves stability without damaging accuracy. At first glance, investigating the stability of exponential integrators may seem unfruitful since the methods have already been used to successfully solve a range of stiff partial differential equations including purely dispersive ones [2,3,5]. However, as was recently described in [26], exponential integrators possess linear stability regions for non-diffusive problems that are on par with those of explicit methods. These instabilities are often very small in magnitude, which explains why they can sometimes go unnoticed. In this work we will present two simple equations that excite instabilities in exponential integrators, and then demonstrate how a simple repartitioning approach can be used to stabilize them.

Our motivation for pursuing this work goes beyond simply stabilizing integrators. In particular, we believe that incorporating exponential integrators into parallel-in-time frameworks has the potential to produce new and efficient parallel methods for solving large scale application problems. These methods will be of particular interest if exponential integrators can also be used to achieve meaningful parallel speedup for equations with little to no diffusion. Unfortunately,

[☆] This work was funded by the National Science Foundation, USA, Computational Mathematics Program DMS-2012875.

* Corresponding author.

E-mail addresses: tbuvoli@ucmerced.edu (T. Buvoli), mlminion@lbl.gov (M.L. Minion).

we have found that the effectiveness of exponential integrators in parallel frameworks is significantly hindered by the stability issues discussed in this paper. Therefore, in this preliminary work, we propose a strategy for overcoming this problem.

In light of this goal we will focus on three integrator families that can be immediately used within parallel frameworks. In particular, we will analyze: (1) exponential Runge–Kutta (ERK) methods [7,8], which can be used to construct exponential Parareal methods, (2) exponential spectral deferred correction methods (ESDC) [13], which are the exponential equivalents of the classical SDC methods used in the Parallel Full Approximation Scheme in Space and Time (PFASST) framework [27], and (3) exponential polynomial block methods (EPBMs) [14] that allow for parallel function evaluations. Lastly, we note that the stability issues we describe affect many other families of exponential integrators, and therefore the developments from this paper can also be used to improve the efficiency of any serial exponential integrators for hyperbolic and non-diffusive problems.

The outline of this paper is as follows. In Section 2 we provide a brief overview of exponential integrators and the three families of methods that will be the subject of this work. In Section 3 we then introduce a simple one-dimensional, nonlinear partial differential equation that leads to severe stepsize restrictions for exponential integrators. Next, in Section 4, we use linear stability analysis to showcase the poor stability properties of exponential integrators on non-diffusive equations. Lastly, in Section 5, we propose a simple repartitioning approach to stabilize exponential integrators, and compare it against the commonly applied practice of adding hyperviscosity.

2. Exponential integrators

Exponential integrators [1] are a family of time integration methods with coefficients that can be written in terms of exponential functions of a linear operator. They have proven to be highly suitable for solving stiff systems and, in certain situations, they can outperform both fully-implicit and linearly-implicit methods [2–5,20]. Nearly all exponential integrators can be classified into two subfamilies: unpartitioned and partitioned. Unpartitioned exponential integrators [13,14,28–31] can be used to solve the generic differential equation $\mathbf{y}' = F(\mathbf{y})$ and require the exponentiation of the local Jacobian of $F(\mathbf{y})$ at each timestep. In contrast, partitioned exponential integrators [6–8,10,13,14,32] are designed to solve the semilinear initial value problem

$$\mathbf{y}' = \mathbf{Ly} + N(t, \mathbf{y}), \quad \mathbf{y}(t_0) = \mathbf{y}_0, \quad (1)$$

and only require exponential functions involving the autonomous linear operator \mathbf{L} .

In this work we will focus exclusively on partitioned exponential integrators. All partitioned exponential integrators can be derived by applying the variation of constants formula to (1), yielding

$$\mathbf{y}(t_0 + h) = e^{h\mathbf{L}}\mathbf{y}_0 + \int_{t_0}^{t_0+h} e^{(t_0+h-\tau)\mathbf{L}}N(\tau, \mathbf{y}(\tau))d\tau, \quad (2)$$

and then approximating the nonlinear term using a polynomial. For example, to obtain the first-order partitioned exponential Euler method with stepsize h ,

$$\mathbf{y}_{n+1} = e^{h\mathbf{L}}\mathbf{y}_n + h \int_0^1 e^{(1-s)h\mathbf{L}}N(t_n, \mathbf{y}_n)ds, \quad (3)$$

we take (2), replace t_0 with $t_n = t_0 + hn$, approximate the nonlinear term $N(\tau, \mathbf{y}(\tau))$ with a zeroth-order polynomial $p(\tau) = N(t_n, \mathbf{y}_n)$ and then rewrite the integrand in local coordinates s using $\tau = t_n + hs$. Higher order exponential multistep, Runge–Kutta, and general linear methods can be constructed by approximating N with a higher-order polynomial that respectively uses previous solution values, previous stage values, or any combination of the two.

In this work we will be concerned with three specific families of exponential integrators.

1. *Exponential Runge–Kutta (ERK)* methods accept a single input, compute s stage values, and produce a single output. The methods developed in [7,8] attempt to achieve a high order of accuracy in the fewest number of stages, and the coefficients have been derived by satisfying nonlinear order conditions. In this work, we will consider the fourth-order exponential Runge–Kutta method ETDRK4-B from [8] which will be referred to simply as ERK4.
2. *Exponential spectral deferred correction (ESDC)* methods [13] are a class of arbitrary order time integration schemes that iteratively improve a provisional solution by solving an integral equation that governs the error. All ESDC methods can be written as ERK methods with a large number of stages. In this paper we will consider a 6th-order ESDC method with four Gauss–Lobatto nodes that takes six correction iterations; we will refer to this method as ESDC6.
3. *Exponential polynomial block methods (EPBM)* [14] are multivalued, exponential general linear methods that advance a set of q different solution values at each timestep. The input values are approximations to the solution at different time points, and the outputs are computed by approximating the nonlinearity in (1) using a high-order polynomial that interpolates the nonlinearity at the input values. The methods can be constructed at any order of accuracy and allow for parallel function evaluations. Here we consider a fifth-order composite EPBM that is constructed using Legendre nodes and is run using an extrapolation factor of $\alpha = 1$.

The formulas for all three exponential integrators are contained in [Appendix A](#), and Matlab code for running all the numerical experiments can be downloaded from [33].

Table 1
Numerical parameters for the ZDS equation (4).

Domain:	$x \in [-4\pi, 4\pi]$, $t \in [0, 40]$
Boundary conditions:	periodic
Spatial discretization:	Fourier pseudo-spectral
Number of spatial grid points:	$N_x = 128$
Dealiasing:	$\frac{3}{2}$ rule (see for example [35, p. 84])

3. A motivating example

We begin by showcasing a simple one dimensional nonlinear partial differential equation that causes stability problems for exponential integrators. Specifically, we consider the zero-dispersion Schrödinger equation (ZDS)

$$\begin{aligned} iu_t + iu_{xxx} + 2u|u|^2 &= 0 \\ u(x, t=0) &= 1 + \frac{1}{100} \exp(3ix/4). \end{aligned} \quad (4)$$

This complex-valued, dispersive partial differential equation models optical pulses in zero dispersion fibers [34] and is obtained by replacing the second derivative in the canonical nonlinear Schrödinger equation with an iu_{xxx} term.

We equip the ZDS equation with periodic boundary conditions in x and integrate to time $t = 40$ using a Fourier pseudo-spectral method. To simplify the computation of the exponential functions, we solve the equation in Fourier space where the linear derivative operators are diagonal and (4) reduces to the semilinear Eq. (1) with

$$\mathbf{L} = \text{diag}(i\mathbf{k}^3) \quad \text{and} \quad N(\hat{u}) = 2\mathcal{F}(\mathcal{F}^{-1}(\hat{u}) \cdot \text{abs}(\mathcal{F}^{-1}(\hat{u}))^2)$$

where \mathbf{k} is a vector of Fourier wavenumbers and \mathcal{F} denotes the discrete Fourier transform. The full numerical parameters are summarized in Table 1.

In Fig. 1 we show the solution obtained by integrating in time using 2000 steps of the ERK4 method. For comparison we also compute the solution using the fourth-order implicit-explicit method from [36] named ARK4(3)6L2[2]SA, which we will simply refer to as IMRK4. Instabilities are clearly visible in the ERK4 solution while the IMRK4 solution is clean. Even more surprisingly, the problem is only mildly stiff. By using 2000 timesteps, or equivalently $h = 2 \times 10^{-2}$, the spectral radius of the linear operator $\rho(h\mathbf{L})$ is only 23.1525; note that we are computing the spectral radius for the lower 2/3 of modes where antialiasing is not applied.

To make matters worse, instabilities persist in the ERK4 integrator across a range of coarse timesteps and only disappear when the stepsize has been reduced by a factor of approximately ten. In Fig. 2 we show convergence diagrams for the ERK4, ESDC6, and EPBM5 integrators. For reference we also include IMRK4 and classical explicit RK4 [37, p. 194]. The reference solution was calculated using RK4 with 200,000 timesteps, and the relative error in all our plots is defined as $\|\mathbf{y}_{\text{ref}} - \mathbf{y}\|_{\infty}/\|\mathbf{y}_{\text{ref}}\|_{\infty}$.

At coarse timesteps the exponential integrators either generate completely inaccurate solutions or are outright unstable. Proper convergence is only achieved when one uses timesteps that are nearly on par with those required to keep the explicit RK4 method stable. Conversely, the IMRK4 method converges at slightly higher than fourth-order across the full range of timesteps. Lastly, we want to emphasize that this problem presents the same stability challenges for a wide range of exponential integrators. In Appendix B we show two additional plots with lower-order ERK, EPBM, and ESDC methods. Moreover, though we have not completed an exhaustive experiment with all known partitioned exponential integrators, all the families of integrators that we have tested exhibit similar behavior on this problem.

4. Linear stability

In this section we use linear stability analysis [38, IV.2] to study the stability properties of exponential integrators on non-diffusive problems. Since we are considering partitioned integrators, we analyze stability for the non-diffusive, partitioned Dahlquist equation

$$y' = i\lambda_1 y + i\lambda_2 y, \quad y(0) = y_0, \quad \lambda_1, \lambda_2 \in \mathbb{C}. \quad (5)$$

This equation is a special case of the partitioned Dahlquist equation with $\lambda_1, \lambda_2 \in \mathbb{C}$ that was used to study the general stability of both implicit-explicit [39–41] and exponential integrators [6,7] including ERK4 [8], ESDC [13] and EPBMs [14].

We can relate the partitioned Dahlquist equation to a general nonlinear system, by observing that (1) can be reduced to a decoupled set of partitioned Dahlquist equations provided that $N(t, \mathbf{y})$ is an autonomous, diagonalizable linear operator that shares all its eigenvectors with \mathbf{L} . Though this assumption does not hold true for general nonlinear systems, the partitioned Dahlquist equation has nevertheless proven invaluable for studying stability.

When solving (5) with a partitioned exponential integrator, we treat the term $i\lambda_1 y$ exponentially and the term $i\lambda_2 y$ explicitly. A one-step exponential integrator like ERK or ESDC then reduces to the scalar iteration,

$$y_{n+1} = R(ik_1, ik_2)y_n \quad \text{where} \quad k_1 = h\lambda_1, \quad k_2 = h\lambda_2, \quad (6)$$




Fig. 1. Solution to the ZDS equation in physical and Fourier space using the ERK4 and the IMRK4 integrators that are both run using 2000 timesteps. In physical space the squared absolute value of the solution $u(x, t)$ is plotted, while in Fourier space we show the log of the absolute value of the Fourier coefficients as a function of time t and wavenumber k .

while a multivalued integrator like EPBM5 reduces to a matrix iteration

$$\mathbf{y}^{[n+1]} = \mathbf{M}(ik_1, ik_2)\mathbf{y}^{[n]} \quad (7)$$

where $\mathbf{M}(ik_1, ik_2) \in \mathbb{C}^{q \times q}$ and q represents the number of input values. The stability region S includes all the (k_1, k_2) pairs that guaranteed a bounded iteration. For the scalar case we simply need a stability function R with magnitude less than one, so that

$$S = \{(k_1, k_2) \in \mathbb{R}^2 : |R(ik_1, ik_2)| \leq 1\}. \quad (8)$$

For the multivalued case, we require the matrix $\mathbf{M}(ik_1, ik_1)$ to be power bounded. Therefore, we redefine the stability function as $R(ik_1, k_2) = \rho(\mathbf{M}(ik_1, ik_1))$ where ρ denotes the spectral radius. On the boundary, we must take special care to ensure that any eigenvalues of magnitude one are non-defective.

Since we are only considering real-valued λ_1 and λ_2 in (5), we can visualize the stability regions for non-diffusive problems using a simple two-dimensional plot. Moreover, all the integrators we consider have stability functions that satisfy $R(ik_1, ik_2) = R(-ik_1, -ik_2)$. Therefore, we only need to consider $k_1 \geq 0$. This same approach was used in [42] to visualize the stability of parareal integrators based on implicit-explicit Runge-Kutta methods.

In Fig. 3 we show the stability regions S and the stability functions $R(ik_1, ik_2)$ for the ERK4, ESDC6, and EPBM5 methods. If we exclude the line $k_2 = 0$, which corresponds to $N(t, \mathbf{y}) = 0$ in (1), then the stability regions of all three exponential integrators are disjoint. This characteristic is extremely undesirable and explains the severe timestep restrictions we experienced when solving the ZDS equation.




Fig. 2. A convergence diagram for ERK4, ESDC6, and EPBM5 methods on the ZDS equation. We also include the IMRK4 method and classical explicit fourth-order RK4. At coarse timesteps, instabilities lead to a total loss of convergence for ERK4 and ESDC6 and cause EPBM5 to be completely unusable. These limitations make all three exponential integrators only marginally more stable than the fully explicit RK4. Conversely, the IMRK4 method converges across the full range of timesteps and exhibits more than fourth convergence for this problem (approximately 4.5).




Fig. 3. Stability regions \mathcal{S} (top row) and stability functions $R(ik_1, ik_2)$ (bottom row) of the ERK4, ESDC6, and EPBM5 methods for $0 \leq k_1 \leq 60$. To improve readability, the k_2 axis is scaled differently for each method. In the top row we show the proper stability region in dark gray and an extended stability region in light gray that allows for a small amount of instability.

Interestingly, the magnitude of the instabilities is very small, and allowing $R(ik_1, ik_2)$ to grow marginally larger than one leads to fully connected regions. However, this observation has little practical significance since we have already seen one example where unmodified exponential integrators produce unusable results. Nevertheless, this phenomenon suggests that exponential integrators will converge successfully on sufficiently small integration intervals that do not allow

instabilities to fully manifest. As we will see in Section 5.3, this likely explains the existence of successful convergence results found in the literature.

Another interesting feature of the stability region is the behavior near the line $k_2 = 0$. Since all exponential integrators treat the linear term exactly, the stability functions all satisfy $R(ik_1, 0) = 1$, which pins the stability function right on the border of stability along the $k_2 = 0$ line. Therefore, even small values of k_2 can easily perturb the stability function above one. However, unlike explicit methods, the instabilities decrease in magnitude as k_1 increases. This is due to the asymptotic behavior of the exponential integral that approximates the nonlinearity in (2); namely for any finite k_2 ,

$$\lim_{k_1 \rightarrow \pm\infty} \left| k_2 \int_0^1 e^{ik_1(s-1)} p(s) ds \right| \sim \frac{1}{|k_1|} \quad \text{for any polynomial } p(s). \quad (9)$$

To prove this result one simply needs to apply integration by parts.

5. Improving stability through repartitioning

We now introduce a simple repartitioning strategy to eliminate the stability limitations described in the previous two sections. In short, we propose to repartition the system (1) as

$$\mathbf{y}' = \hat{\mathbf{L}}\mathbf{y} + \hat{N}(t, \mathbf{y}) \quad (10)$$

where

$$\hat{\mathbf{L}} = \mathbf{L} + \epsilon \mathbf{D}, \quad \hat{N}(t, \mathbf{y}) = N(t, \mathbf{y}) - \epsilon \mathbf{D}\mathbf{y}, \quad (11)$$

and \mathbf{D} is a diffusive operator. The repartitioned system (10) is mathematically equivalent to (1), however a partitioned exponential integrator now exponentiates $\hat{\mathbf{L}}$ instead of the original matrix \mathbf{L} .

To improve stability, we seek a matrix \mathbf{D} so that the eigenvalues of $\hat{\mathbf{L}}$ have some small negative real-component. We start by first assuming that \mathbf{L} is diagonalizable so that $\mathbf{L} = \mathbf{U}\Lambda\mathbf{U}^{-1}$ and then select

$$\mathbf{D} = -\mathbf{U}\text{abs}(\Lambda)\mathbf{U}^{-1}. \quad (12)$$

Although this choice may be difficult to apply in practice, it is very convenient for analyzing the stability effects of repartitioning. If we apply (12) to the Dahlquist test problem (5), then $\mathbf{D} = -|\lambda_1|$ and we obtain the repartitioned non-diffusive Dahlquist equation

$$y' = \underbrace{(i\lambda_1 - \epsilon|\lambda_1|)}_{\hat{\mathbf{L}}} y + \underbrace{(i\lambda_2 + \epsilon|\lambda_1|)y}_{\hat{N}(t, y)}. \quad (13)$$

The associated stability region for an exponential integrator is

$$S = \left\{ (k_1, k_2) \in \mathbb{R}^2 : \left| \hat{R}(k_1, k_2) \right| \leq 1 \right\}, \quad (14)$$

$$\hat{R}(k_1, k_2) = R(ik_1 - \epsilon|k_1|, ik_2 + \epsilon|k_1|), \quad (15)$$

where R is the stability function of the unmodified integrator. By choosing

$$\epsilon = 1/\tan(\frac{\pi}{2} + \rho) \quad \text{for } \rho \in [0, \frac{\pi}{2}), \quad (16)$$

the single eigenvalue of the partitioned linear operator $\hat{\mathbf{L}}$ is now angled ρ radians off the imaginary axis into the left half-plane. Conversely, the “nonlinear” operator $\hat{N}(t, y)$ has been rotated and scaled into the right half-plane. Therefore, in order for the method to stay stable, the exponential functions of $\hat{\mathbf{L}}$ must damp the excitation that was introduced in the nonlinear component.

In Fig. 4 we show the stability regions for ERK4, ESDC6, and EPBM5 after repartitioning with different values of ρ . Rotating the linear operator by only $\rho = \pi/2048$ radians (≈ 0.088 degrees) already leads to large connected linear stability regions for all three methods. This occurs because the magnitude of the partitioned stability function $\hat{R}(k_1, k_2)$ along the line $k_2 = 0$ is now less than one for any $k_1 \neq 0$. In fact, by increasing ρ further, one introduces additional damping for large k_1 . Therefore, under repartitioning, high-frequency modes will be damped, while low frequency modes will be integrated in a nearly identical fashion to the unmodified integrator with $\rho = 0$. Excluding scenarios where energy conservation is critical, the damping of high-frequency modes is not a serious drawback since large phase errors in the unmodified integrator would still lead to inaccurate solutions (supposing the method is stable in the first place).

Repartitioning cannot be applied indiscriminately and as ρ approaches $\frac{\pi}{2}$ one obtains an unstable integrator. To highlight this phenomenon more clearly, we show magnified stability regions in Fig. 5, in which we selected sufficiently large ρ values to cause stability region separation for each method. From these results we can see that the maximum amount of allowed repartitioning is integrator dependent, with ESDC6 allowing for the most repartition and EPBM5 the least.

Finally, we note that this repartitioning technique can also be applied to implicit-explicit methods. However, on all the methods we tried, we found that the repartitioning rapidly destabilizes the integrator. In Fig. 6 we present the stability




Fig. 4. Stability regions \mathcal{S} of the ERK4, ESDC6, and EPBM5 methods for various ρ values. On each plot, the x -axis represents k_1 and the y -axis represents k_2 . The color represents the $\log |\hat{R}(k_1, k_2)|$ from (15).




Fig. 5. Magnified stability regions that show the effects of adding too much repartitioning. Each integrator family has a different amount of maximal repartitioning before the stability regions split along the k_2 line. Amongst the three methods, ESDC6 was the most robust and EPBM5 was the least robust.




Fig. 6. Stability regions for the repartitioned IMRK4 methods with four values of ρ . For even small values of ρ , instabilities form along the $k_2 = 0$ line near the origin, and for large values of ρ , the stability regions separate. The effect appears for even smaller values of ρ , however we did not include these plots here, since the effect is not visible at this level of magnification.

regions for IMRK4 using different ρ values and show that stability along the $k_2 = 0$ line is lost even for small ρ values. The stability region corresponding to $\rho = 0$ can be compared with the stability regions of the exponential integrators in Fig. 4 to see how the damping properties of a repartitioned exponential methods compare to those of an IMEX method.

5.1. Solving ZDS with repartitioning

We now validate our stability results by solving the ZDS Eq. (4) using several different choices for the diffusive operator \mathbf{D} . Since we are solving in Fourier space where $\mathbf{L} = \text{diag}(ik^3)$, we can easily implement (12) by selecting $\mathbf{D} = -\text{diag}(|\mathbf{k}|^3)$. However, for more general problems, it will often not be possible to compute the eigenvalue decomposition of the operator \mathbf{L} . Therefore, to develop a more practical repartitioning, we consider a generic, repartitioned, semilinear partial differential equation

$$u_t = \underbrace{\hat{L}[u]}_{\hat{L}[u]} + \underbrace{\epsilon D[u]}_{\hat{D}[u]} + \underbrace{N(t, u) - \epsilon D[u]}_{\hat{N}(t, u)} \quad (17)$$


where the spatial discretizations of $\hat{L}[u]$ and $\hat{N}(t, u)$ are the $\hat{\mathbf{L}}\mathbf{y}$ and $\hat{N}(t, \mathbf{y})$ from (11), and the linear operator $D[u]$ is the continuous equivalent of the matrix \mathbf{D} . A natural choice for a diffusive operator in one spatial dimension is the even-powered derivative

$$D[u] = \begin{cases} -\frac{\partial^k u}{\partial x^k} & k \equiv 0 \pmod{4}, \\ \frac{\partial^k u}{\partial x^k} & k \equiv 2 \pmod{4}. \end{cases} \quad (18)$$


This operator can be easily implemented for different spatial discretizations and boundary conditions. Moreover, it can be generalized to higher dimensional PDEs by adding partial derivatives in other dimensions.

The only remaining question is how to choose k . To avoid increasing the number of required boundary conditions, it is advantageous if k is smaller than or equal to the highest derivative order found in $L[u]$. Therefore, in addition to (12), we also consider two additional repartitionings for the ZDS equation that are based on the zeroth and the second spatial derivative of $u(x, t)$. Below, we describe each repartitioning in detail, and in Fig. 7 we also plot the spectrum of the corresponding repartitioned linear operators $\hat{\mathbf{L}}$.

$$\mathbf{D} = -\text{diag}(|\mathbf{k}|^3)$$



$$\mathbf{D} = -\text{diag}(\mathbf{k}^2)$$



$$\mathbf{D} = -\mathbf{I}$$




Fig. 7. Spectrum of the repartitioned linear operator $\hat{\mathbf{L}} = \mathbf{L} + \epsilon \mathbf{D}$, for three choices of \mathbf{D} and multiple ϵ . In the first two plots ϵ is selected according to (16). The choices of ρ for the second-order repartitioning were selected to achieve comparable damping to third-order repartitioning.

- **Third-order repartitioning.** The diffusive operators are

$$D[u] = \mathcal{F}^{-1}(|k|^3) * u \quad (\text{Continuous - physical space}), \quad (19)$$

$$\mathbf{D} = -\text{diag}(|\mathbf{k}|^3) \quad (\text{Discrete - Fourier space}), \quad (20)$$

where \mathcal{F}^{-1} denotes the inverse Fourier transform and $*$ is a convolution. This choice is equivalent to (12). We choose ϵ according to (16), so that the eigenvalues of $\hat{\mathbf{L}} = \mathbf{L} + \epsilon \mathbf{D}$ lie on the curve

$$re^{i(\pi/2+\rho)} \cup re^{i(3\pi/2-\rho)} \quad (r \geq 0).$$

For this repartitioning we select ϵ using $\rho = \frac{\pi}{2048}, \frac{\pi}{512}, \frac{\pi}{128}$, and $\frac{\pi}{32}$.

- **Second-order repartitioning.** The diffusive operators are

$$D[u] = u_{xx} \quad (\text{Continuous - physical space}), \quad (21)$$

$$\mathbf{D} = -\text{diag}(\mathbf{k}^2) \quad (\text{Discrete - Fourier space}), \quad (22)$$

and we again choose ϵ according to (16). Compared to the previous choice, second-order repartitioning over-rotates eigenvalues with magnitude less than one, and under-rotates eigenvalues with magnitude larger than one. Therefore we require larger ρ values to achieve similar damping effects as third-order repartitioning (see Fig. 7); in particular we select ϵ using $\rho = \frac{\pi}{256}, \frac{\pi}{64}, \frac{\pi}{16}$, and $\frac{\pi}{4}$.

- **Zeroth-order repartitioning.** The diffusive operators are

$$D[u] = -u \quad (\text{Continuous - physical space}), \quad (23)$$

$$\mathbf{D} = -\mathbf{I} \quad (\text{Discrete - Fourier space}). \quad (24)$$

This choice translates every eigenvalue of the linear operator \mathbf{L} by a fixed amount ϵ into the left-hand plane. We consider $\epsilon = 1, 2, 4$ and 8 .

In Figs. 8 we present convergence diagrams for ERK4, ESDC6, and EPBM5 using each of the three repartitioning strategies. Overall, repartitioning resolved the stability issues and enabled the use of exponential integrators for efficiently solving the ZDS equation. We summarize the results as follows.

Third-order repartitioning. Adding even a small amount of third-order repartitioning immediately improves the convergence properties of the exponential integrators. For $\rho = \pi/128$, all integrators achieve proper convergence across the full range of stable timesteps. Moreover, adding additional repartitioning does not damage the accuracy so long as the underlying method remains stable.

Second-order repartitioning. Second-order repartitioning is able to achieve nearly identical results to third-order repartitioning, provided that larger ρ -values are used. Overall, the results are not surprising since the spectrums of the corresponding linear operators shown in Fig. 7 look very similar. The main disadvantage of second-order repartitioning is that ρ needs to be tuned to ensure that the highest modes have been sufficiently rotated.

Zeroth-order repartitioning. Zeroth order repartitioning is extremely simple to implement, however it is also the least effective at improving convergence and preserving accuracy. A small ϵ does not introduce enough damping and the convergence curves are improved but not fully restored. On the other hand, large ϵ values stabilize stiff modes, however since all the eigenvalues are shifted by an equal amount, the repartitioning damages the accuracy of non-stiff modes. This leads to convergence curves that have been shifted to the left since we have effectively worsened the error constant of the exponential integrator. Zeroth-order repartitioning also negatively impacted the sensitivity of the integrator to roundoff errors, and we were unable to obtain the solution with a relative error of less than approximately 10^{-8} .

5.2. Comparing repartitioning to hyperviscosity

Time dependent equations without diffusion are known to cause stability issues for numerical methods, and a commonly applied strategy is to add a hyperviscous term to the equation right-hand-side [43–45]. If the magnitude of the hyperviscous term is constant, then the numerical solution will converge to the perturbed equation. To avoid affecting the truncation error (relative to the unmodified equation), the magnitude of the hyperviscosity can depend on the spatial grid spacing (e.g. see [44,46]) or the timestep (e.g. see [45]). In this work we are solely concerned with time convergence; therefore, we consider hyperviscous terms whose magnitude is proportional to the stepsize h raised to the power of $q+1$, where q is the order of the integrator. In the context of the continuous semilinear partial differential equation

$$u_t = L[u] + N(t, u) \quad (25)$$

this is equivalent to considering a new equation with a vanishing diffusive operator $\tilde{D}[u]$ added to the right-hand-side so that

$$u_t = L[u] + h^{q+1} \gamma \tilde{D}[u] + N(t, u), \quad (26)$$

where γ is a constant that controls the strength of the diffusion. One then approximates the solution to (25) by numerically integrating (26). The improvement to stability comes from the fact that we have replaced the original discretized linear operator \mathbf{L} with $\tilde{\mathbf{L}} = \mathbf{L} + h^{q+1} \gamma \mathbf{D}$.

Unlike repartitioning, we are no longer adding and subtracting the new operator. We must therefore ensure that $D[u]$ does not damage the accuracy of slow modes as they typically contain the majority of useful information. For this reason, $D[u]$ is generally chosen to be a high-order even derivative since these operators have a negligible effect on low frequencies while causing significant damping of high frequencies.

To compare the differences between repartitioning and hyperviscosity, we re-solve the ZDS equation using ERK4 with hyperviscosity of orders four, six, and eight. Since ERK4 is a fourth-order method, we take $q = 4$. In Fig. 9 we show




Fig. 8. Convergence diagrams for the repartitioned ERK4, ESDC6, and EPBM5 integrators on the ZDS equation using third, second, and zeroth order repartitioning. For reference we also include the unpartitioned IMRK4 integrator in each plot. The columns correspond to different choices for the matrix \mathbf{D} , and each row to different integrators. Color denotes the value of ρ from (16) for the first two columns and the value of ϵ for the third column.

convergence diagrams for these experiments. We immediately see that hyperviscosity is only effective when a sufficiently high-order spatial derivative is used. In particular, fourth-order hyperviscosity fails to improve stability for small γ and completely damages the accuracy of the integrator for larger γ . Sixth-order hyperviscosity offers a marginal improvement at coarse stepsizes, but also damages accuracy at fine stepsizes. Eighth-order hyperviscosity with $\gamma = 10^4$ is the only choice that achieves results comparable to repartitioning.

In summary, repartitioning offers two key advantages. First, it does not require the use of high-order spatial derivatives, and second, it is less sensitive to overdamping. These advantages are both due to the fact that repartitioning does not modify the underlying problem, while hyperviscosity is only effective if the modified problem (26) closely approximates the original problem (25). We discuss both points in more detail below.

Sensitivity to overdamping. When adding hyperviscosity, it is critical to select the smallest possible γ that suppresses instabilities. Selecting larger γ causes the solutions of (25) and (26) to grow unnecessarily far apart, and leads to a time integration scheme that converges more slowly to the solution of (26). Specifically, the convergence curves for methods with excessive hyperviscosity look like the convergence curves of a method with an increased error constant. This




Fig. 9. Convergence plots for ERK4 with hyperviscosity added to the linear operator; for convenience we write γ in terms of a new parameter ω . Low-order hyperviscosity is unable stabilize the integrator, while a properly selected 8th-order diffusion yields results that are nearly identical to repartitioning.

phenomenon appears prominently in Fig. 9, where ERK4 methods large amounts of hyperviscosity consistently performed worse than all other methods at fine timesteps (e.g. see graphs for $\omega = 10^6$).

In contrast, second-order and third-order repartitioning are significantly more flexible since they allow for a greater amount of over-partitioning without any significant damage to the accuracy or stability. Excessively large ϵ values can still cause the stability region separation shown in Fig. 5, however such values are unlikely to be used in practice, since they lead to a partitioned linear operator with eigenvalues that have a larger negative real part than imaginary part. Zeroth-order repartitioning is most similar to hyperviscosity since large values of ϵ also damage the error constant of the method; however, the effects are significantly less pronounced.

Importance of high-order spatial derivatives. When adding hyperviscosity we must ensure that the small eigenvalues of the modified linear operator $\hat{\mathbf{L}}$ closely approximate those of \mathbf{L} , or we risk altering the dynamics of slow modes. Therefore, we require a small γ for low-order hyperviscous terms. However, this creates a dilemma: choosing a small γ may not eliminate the instabilities while choosing a large γ damages accuracy. This is exactly why we were not able to efficiently stabilize the ZDS equation using fourth-order hyperviscosity.

In contrast, repartitioning does not require that the small eigenvalues of $\hat{\mathbf{L}}$ closely approximate those of \mathbf{L} , since the nonlinear term counteracts any changes. This is perhaps most easily explained by considering the Dahlquist Eq. (5). If $|\lambda| = |\lambda_1 + \lambda_2|$ is small (i.e. the mode is slow), then an exponential integrator will integrate the system accurately so long as $|\lambda_1|$ and $|\lambda_2|$ are also small. Hence, we can freely redistribute the wavenumber between λ_1 and λ_2 . This allows us to repartition using second-order diffusion without losing accuracy.

A comment on efficiency. Our numerical experiments involve equations with periodic domains that are solved in Fourier space where all linear derivative operators are diagonal. For these problems the computational cost of hyperviscosity and repartitioning are equivalent since all exponential functions can be cheaply precomputed before the first timestep. However, for more general boundary conditions, hyperviscosity can lead to non-diagonal linear operators with large negative eigenvalues. Computing the exponentials of such operators can be more expensive for both direct methods [15,16] (more scaling and squaring iterations), Krylov-based methods [22,23] (more basis vectors required), and polynomial or rational function based methods [17–19] (larger contours). Therefore, an additional benefit of repartitioning over hyperviscosity, is the potential improvement in computational efficiency for equations with more general boundary conditions and domains.

5.3. Long-time integration of Korteweg-de Vries

The ZDS Eq. (4) causes considerable difficulties for unmodified exponential integrators even on short timescales when the dynamics are simple. However, this does not imply that repartitioning is always necessary for solving dispersive equations. In fact, there are many reported results in the literature where exponential integrators converge successfully without requiring any modifications; several examples include [3,5], and [2].

As discussed in Section 4, the instabilities of exponential integrators are very small in magnitude and can therefore go unnoticed for long periods of time. To explore this further, we present a final numerical experiment where we solve the




Fig. 10. Relative error versus time for the KDV Eq. (27) solved using ERK4, ESDC6, and EPBM5 with 56000 timesteps. Each integrator is shown using a different line marker. The blue curves denote unpartitioned integrators, while shades of gray denote various repartitionings. The differences in relative error between different repartitionings is very small, so the gray lines almost entirely overlap.

Korteweg–de Vries (KDV) equation from [47]

$$\frac{\partial u}{\partial t} = - \left[\delta \frac{\partial^3 u}{\partial x^3} + \frac{1}{2} \frac{\partial}{\partial x} (u^2) \right] \quad (27)$$

$$u(x, t = 0) = \cos(\pi x), \quad x \in [0, 2]$$

where $\delta = 0.022$. The boundary conditions are periodic, and we discretize in space using a 512 point Fourier spectral method. As with the ZDS equation, we solve in Fourier space where the linear operator $\mathbf{L} = \text{diag}(i\delta\mathbf{k}^3)$.

This exact problem was used in both [13,14] to validate the convergence of ERK, ESDC, and EPBM methods. In the original experiments the equation was integrated to time $3.6/\pi$. On these timescales ERK4, ESDC6, and EPBM5 all converge properly and show no signs of instability. We now make the problem more difficult to solve by extending the integration interval to $t = 160$. The longer time interval increases the complexity of the solution and allows for instabilities to fully manifest.

In Fig. 10 we show how the relative error of both unpartitioned and partitioned ERK4, ESDC6, and EPBM5 methods evolves in time. To produce the plot, we run all integrators using 56000 timesteps and compare their outputs to a numerically computed reference solution at 30 equispaced times between $t = 0$ and $t = 160$.

On short timescales all unmodified exponential integrators are stable and no repartitioning is required. However, on longer timescales, repartitioning becomes necessary. Moreover, the maximum time horizon before instabilities dominate differs amongst the integrators. The unmodified EPBM5 method is the first integrator to become unstable around $t = 20$. The unmodified ERK4 method is more robust and remains stable until approximately time $t = 55$, while the unmodified ESDC6 method remains stable across almost the entire time interval. Unlike the ZDS example, the time to instability is now correlated to the size of the methods stability region.

Adding zeroth-order, second-order, or third-order repartitioning stabilizes all the methods, and does not damage the accuracy in regions where the unmodified integrator converged. Furthermore, the accuracy differences between the three repartitioning strategies is effectively negligible. Lastly, the repartitioning parameters described in the legend of Fig. 10 allow us to compute the solution at even longer times; we tested all methods out to time $t = 1000$, after which we did not look further, and found that all partitioned method configurations remained stable.

6. Summary and conclusions

In this work we studied the linear stability properties of exponential integrators on stiff equations with no diffusion, and showed that without any modifications, exponential integrators are inherently unstable at large stepsizes. Moreover, we demonstrated that the severity of the instabilities varies across problems and integrator families.

The main contribution of this paper is the introduction of a simple repartitioning approach that resolves the stability issues and improves convergence for stiff non-diffusive equations. Compared to the common practice of adding hyperviscosity, repartitioning provides better accuracy and does not require introducing high-order spatial derivatives. Furthermore, there is no need to solve a modified equation; instead repartitioning introduces a pre-specified amount of numerical dissipation directly into an exponential integrator.

The applications of this work are two-fold. First, repartitioning is a useful tool for any practitioner experiencing stability issues when applying exponential integrators on hyperbolic or dispersive equations. The simplest way to apply repartitioning is to solve an equation using a range of different repartitioning terms until the desired result is achieved. Alternatively, one can study the linear stability properties of an integrator, and then compute the spectrum of the equation's linear operator to determine a suitable repartitioning term and constant ϵ .

The second application of this work relates to exponential integrators and parallel-in-time frameworks where effect of instabilities becomes much more significant. In followup works we will apply this repartitioning strategy within the Parareal and PFASST frameworks to solve non-diffusive equations.

Acknowledgments

The work of Buvoli was funded by the National Science Foundation, USA, Computational Mathematics Program DMS-2012875. The work of Minion was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under contract number DE-AC02005CH11231.

Appendix A. Method coefficients

- **ERK4** is a fourth-order exponential Runge–Kutta method given by

$$\begin{aligned}
 K_0 &= N(t_n, y_n) \\
 K_1 &= N(t_n + h/2, \varphi_0(h\mathbf{L}/2)y_n + \frac{1}{2}\varphi_1(h/2\mathbf{L})K_0) \\
 K_2 &= N(t_n + h/2, \varphi_0(h\mathbf{L}/2)y_n + (\frac{1}{2}\varphi_1(h/2\mathbf{L}) - \varphi_2(h/2\mathbf{L}))K_0 + \varphi_2(h/2\mathbf{L})K_1) \\
 K_3 &= N(t_n + h, \varphi_0(h\mathbf{L})y_n + (\varphi_1(h\mathbf{L}) - 2\varphi_2(h\mathbf{L}))K_0 + 2\varphi_2(h\mathbf{L})K_2) \\
 y_{n+1} &= \varphi_0(h\mathbf{L})y_n + (\varphi_1(h\mathbf{L}) - 3\varphi_2(h\mathbf{L}) + 4\varphi_3(h\mathbf{L}))K_1 \\
 &\quad + (2\varphi_2(h\mathbf{L}) - 4\varphi_3(h\mathbf{L}))(K_2 + K_3) + (-\varphi_2(h\mathbf{L}) - 4\varphi_3(h\mathbf{L}))K_3
 \end{aligned}$$

- **ESDC6** is a sixth-order exponential spectral deferred correction method. To write down the formula for ESDC6, we first define the quadrature nodes

$$t_{n,j} = t_n + h\eta_j, \quad j = 1, \dots, 4$$

for $\eta_j = \{0, \frac{1}{2} - \frac{\sqrt{5}}{10}, \frac{1}{2} + \frac{\sqrt{5}}{10}, 1\}$ and let $h_{n,j} = t_{n,j+1} - t_{n,j}$. The pseudocode for ESDC6 is:

```

 $Y_{n,1}^{[k]} = y_n, \quad N_{n,j}^{[k]} = N(t_{n,j}, Y_{n,j}^{[k]})$ 
for  $j = 1$  to  $3$ 
 $Y_{n,j+1}^{[1]} = \varphi_0(h_{n,j}\mathbf{L})Y_{n,j}^{[1]} + h_{n,j}\varphi_1(h_{n,j}\mathbf{L})N_{n,j+\gamma}^{[1]}$ 
for  $k = 1$  to  $6$ 
for  $j = 1$  to  $3$ 
 $Y_{n,j+1}^{[k+1]} = \varphi_0(h_{n,j}\mathbf{L})Y_{n,j}^{[k+1]} + h_{n,j}\varphi_1(h_{n,j}\mathbf{L})\left[N_{n,j+\gamma}^{[k+1]} - N_{n,j+\gamma}^{[k]}\right] + I_{n,j}^{[k]}$ 
 $y_{n+1} = Y_{n,p}^{[m+1]}$ 

```

To write the exponential integral terms $I_{j,k}$ we first define

$$\tau_{j,i} = \frac{t_{n,i} - t_{n,j}}{h_{n,j}}, \quad i, j = 1, \dots, 4,$$

$$\mathbf{V}(j)_{c,d} = (\tau_{j,c})^{d-1}, \quad \mathbf{V}(j) \in \mathbf{R}^{4 \times 4},$$

then,

$$I_{n,j}^{[k]} = h_{n,j} \sum_{v=1}^p \varphi_v(h_{n,j}\mathbf{L}) \mathbf{b}_v^{[k]} \quad \text{where} \quad \mathbf{b}_v^{[k]} = \sum_{l=1}^p a_{j,l}^{(v)} N_{n,l}^{[k]}. \quad (28)$$

where the weights $a_{j,l}^{(v)} = v! \mathbf{V}(j)_{v+1,l}^{-1}$.

- **EPBM5** is composite method based on an exponential polynomial block method that accepts five inputs $y_j^{[n]} \approx y(t_n + rz_j)$ and produces five outputs $y_j^{[n+1]} \approx y(t_n + rz_j + r\alpha)$ where the nodes $\{z_j\}_{j=1}^5$ are $\{-1, -\eta^+, -\eta^-, \eta^-, \eta^+\}$ for $\eta^\pm = \sqrt{\frac{3}{7} \pm \frac{2}{7}\sqrt{\frac{6}{5}}}$. The output computation is

$$y_j^{[n+1]} = \varphi_0(r\eta_j(\alpha)\mathbf{L})y_j^{[n]} + r \sum_{k=1}^4 \eta_j^k \varphi_k(r\eta_j(\alpha)\mathbf{L}) \mathbf{v}_k, \quad j = 1, \dots, 5, \quad (29)$$

where $\eta_j(\alpha) = z_j + \alpha + 1$ and the vectors \mathbf{v}_j are given by

$$\begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \mathbf{v}_3 \\ \mathbf{v}_4 \end{bmatrix} = \begin{bmatrix} (w_1^+ + u_1^+) & (-w_1^- + u_1^-) & (w_1^- + u_1^-) & (-w_1^+ + u_1^+) \\ (-w_2^- - u_2) & (w_2^+ + u_2) & (-w_2^+ + u_2) & (w_2^- - u_2) \\ (w_3^- + u_2) & (-w_3^+ - u_2) & (w_3^+ - u_2) & (-w_3^+ - u_2) \\ -w_4^+ & w_4^+ & w_4^+ & w_4^- \end{bmatrix} \begin{bmatrix} \mathbf{N}_2 \\ \mathbf{N}_3 \\ \mathbf{N}_4 \\ \mathbf{N}_5 \end{bmatrix}$$

where $\mathbf{N}_j = N(t_n + rz_j, y_j^{[n]})$ and the constants w_j^\pm and u_j^\pm are:

$$\begin{aligned} w_1^\pm &= \frac{\sqrt{75 \pm 4\sqrt{30}}}{12} & w_2^\pm &= \frac{\sqrt{10170 \pm 1104\sqrt{30}}}{24} \\ w_3^\pm &= \frac{7\sqrt{1350 \pm 180\sqrt{30}}}{24} & w_4^\pm &= \frac{7\sqrt{150 \pm 20\sqrt{30}}}{8} \\ u_1^\pm &= \frac{3 \pm \sqrt{30}}{12} & u_2 &= \frac{7\sqrt{30}}{24} \end{aligned}$$

EPBM5 first advances the solution using $\alpha = 1$, and then corrects the new solution using $\alpha = 0$. If we denote the right-hand-side of (29) as $M_j(\alpha, t_n, y^{[n]})$, then the composite method EPBM5 can be written as

$$\begin{aligned} \tilde{y}_j^{[n+1]} &= M_j(1, t_n, y^{[n]}) \\ y_j^{[n+1]} &= M_j(0, t_n + r, \tilde{y}_j^{[n+1]}) \quad j = 1, \dots, 5. \end{aligned}$$

The required nonlinear function evaluations $\mathbf{N}_j = N(t_n + rz_j, y_j^{[n]})$ and $\tilde{\mathbf{N}}_j = N(t_n + rz_j + r, \tilde{y}_j^{[n]})$ can each be evaluated in parallel.

Appendix B. Low-order exponential integrators on ZDS

In Fig. 11 we rerun the ZDS numerical experiment described in Table 1 with low-order exponential and IMEX integrators. These plots can be compared directly with those from Fig. 2 which show higher-order methods. The convergence plots demonstrate that low-order ERK, ESDC, and EPBM integrators all suffer from the same stability issues as their higher-order counterparts.




Fig. 11. Convergence diagrams for second-order and third-order ERK, ESDC, EPBM, and IMEXRK methods on the ZDS equation. In both plots we also include the classical explicit fourth-order RK4.

References

- [1] M. Hochbruck, A. Ostermann, Exponential integrators, *Acta Numer.* 19 (2010) 209–286.
- [2] I. Grooms, K. Julien, Linearly implicit methods for nonlinear PDEs with linear dispersion and dissipation, *J. Comput. Phys.* 230 (9) (2011) 3630–3650.
- [3] A. Kassam, L. Trefethen, Fourth-order time stepping for stiff PDEs, *SIAM J. Sci. Comput.* 26 (2005) 1214–1233.
- [4] J. Loffeld, M. Tokman, Comparative performance of exponential, implicit, and explicit integrators for stiff systems of ODEs, *J. Comput. Appl. Math.* 241 (2013) 45–67.
- [5] H. Montanelli, N. Bootland, Solving periodic semilinear stiff PDEs in 1D, 2D and 3D with exponential integrators, 2016, arXiv preprint [arXiv:1604.08900](https://arxiv.org/abs/1604.08900).
- [6] G. Beylkin, J.M. Keiser, L. Vozovoi, A new class of time discretization schemes for the solution of nonlinear PDEs, *J. Comput. Phys.* 147 (2) (1998) 362–387.
- [7] S.M. Cox, P.C. Matthews, Exponential time differencing for stiff systems, *J. Comput. Phys.* 176 (2) (2002) 430–455.
- [8] S. Krogstad, Generalized integrating factor methods for stiff PDEs, *J. Comput. Phys.* 203 (1) (2005) 72–88.
- [9] M. Hochbruck, A. Ostermann, Explicit exponential Runge–Kutta methods for semilinear parabolic problems, *SIAM J. Numer. Anal.* 43 (3) (2005) 1069–1090.
- [10] M. Hochbruck, A. Ostermann, Exponential Runge–Kutta methods for parabolic problems, *Appl. Numer. Math.* 53 (2) (2005) 323–339.
- [11] S. Koikari, Rooted tree analysis of Runge–Kutta methods with exact treatment of linear terms, *J. Comput. Appl. Math.* 177 (2) (2005) 427–453.
- [12] A. Ostermann, M. Thalhammer, W.M. Wright, A class of explicit exponential general linear methods, *BIT Numer. Math.* 46 (2) (2006) 409–431.
- [13] T. Buvoli, A class of exponential integrators based on spectral deferred correction, *SIAM J. Sci. Comput.* 42 (1) (2020) A1–A27.
- [14] T. Buvoli, Exponential polynomial block methods, *SIAM J. Sci. Comput.* 43 (3) (2021) A1692–A1722.
- [15] A.H. Al-Mohy, N.J. Higham, A new scaling and squaring algorithm for the matrix exponential, *SIAM J. Matrix Anal. Appl.* 31 (3) (2010) 970–989.
- [16] A.H. Al-Mohy, N.J. Higham, Computing the action of the matrix exponential, with an application to exponential integrators, *SIAM J. Sci. Comput.* 33 (2) (2011) 488–511.
- [17] M. Caliari, L. Einkemmer, A. Moriggi, A. Ostermann, An accurate and time-parallel rational exponential integrator for hyperbolic and oscillatory PDEs, *J. Comput. Phys.* (2021) 110289.
- [18] T.S. Haut, T. Babb, P.G. Martinsson, B.A. Wingate, A high-order time-parallel scheme for solving wave propagation problems via the direct construction of an approximate time-evolution operator, *IMA J. Numer. Anal.* 36 (2) (2015) 688–716.
- [19] M. Caliari, P. Kandolf, A. Ostermann, The Leja method revisited: Backward error analysis for the matrix exponential, *SIAM J. Sci. Comput.* 38 (3) (2016) A1639–A1661.
- [20] M. Hochbruck, C. Lubich, On Krylov subspace approximations to the matrix exponential operator, *SIAM J. Numer. Anal.* 34 (5) (1997) 1911–1925.
- [21] J. Niesen, W.M. Wright, Algorithm 919: A Krylov subspace algorithm for evaluating the φ -functions appearing in exponential integrators, *ACM Trans. Math. Software* 38 (3) (2012) 22:1–22:19.
- [22] J. Niesen, W.M. Wright, A Krylov subspace method for option pricing, 2011.
- [23] S. Gaudreault, G. Rainwater, M. Tokman, KIOPS: A fast adaptive Krylov subspace solver for exponential integrators, *J. Comput. Phys.* 372 (2018) 236–255.
- [24] M. Caliari, P. Kandolf, A. Ostermann, S. Rainer, Comparison of software for computing the action of the matrix exponential, *BIT Numer. Math.* 54 (1) (2014) 113–128.
- [25] N.J. Higham, E. Hopkins, A catalogue of software for matrix functions. Version 3.0, 2020.
- [26] N. Crouseilles, L. Einkemmer, J. Massot, Exponential methods for solving hyperbolic problems with application to collisionless kinetic equations, *J. Comput. Phys.* 420 (2020) 109688.
- [27] M. Emmett, M. Minion, Toward an efficient parallel in time method for partial differential equations, *Commun. Appl. Math. Comput. Sci.* 7 (1) (2012) 105–132.
- [28] M. Hochbruck, A. Ostermann, J. Schweitzer, Exponential Rosenbrock-type methods, *SIAM J. Numer. Anal.* 47 (1) (2009) 786–803.
- [29] M. Tokman, Efficient integration of large stiff systems of ODEs with exponential propagation iterative (EPI) methods, *J. Comput. Phys.* 213 (2) (2006) 748–776.
- [30] M. Tokman, A new class of exponential propagation iterative methods of Runge–Kutta type (EPIRK), *J. Comput. Phys.* 230 (24) (2011) 8762–8778.
- [31] G. Rainwater, 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.
- [32] V.T. Luan, Efficient exponential Runge–Kutta methods of high order: construction and implementation, *BIT Numer. Math.* 61 (2) (2021) 535–560.

- [33] T. Buvoli, Codebase for “exponential integrators for non-diffusive problems”, 2021, URL <https://github.com/buvoli/exp-non-diffusive>.
- [34] G.P. Agrawal, Nonlinear fiber optics, in: Nonlinear Science at the Dawn of the 21st Century, Springer, 2000, pp. 195–211.
- [35] C. Canuto, M.Y. Hussaini, A. Quarteroni, A. Thomas Jr., et al., Spectral Methods in Fluid Dynamics, Springer-Verlag, 1988.
- [36] C.A. Kennedy, M.H. Carpenter, Additive Runge–Kutta schemes for convection–diffusion–reaction equations, *Appl. Numer. Math.* 44 (1–2) (2003) 139–181.
- [37] J.C. Butcher, Numerical Methods for Ordinary Differential Equations, John Wiley & Sons, 2016.
- [38] G. Wanner, E. Hairer, Solving Ordinary Differential Equations II, Springer Berlin Heidelberg, 1996.
- [39] U.M. Ascher, S.J. Ruuth, B.T. Wetton, Implicit-explicit methods for time-dependent partial differential equations, *SIAM J. Numer. Anal.* 32 (3) (1995) 797–823.
- [40] A. Sandu, M. Günther, A generalized-structure approach to additive Runge–Kutta methods, *SIAM J. Numer. Anal.* 53 (1) (2015) 17–42.
- [41] G. Izzo, Z. Jackiewicz, Highly stable implicit–explicit Runge–Kutta methods, *Appl. Numer. Math.* 113 (2017) 71–92.
- [42] T. Buvoli, M. Minion, IMEX Runge–Kutta Parareal for Non-diffusive Equations, in: B. Ong, J. Schroder, J. Shipton, S. Friedhoff (Eds.), Parallel-in-Time Integration Methods, Springer International Publishing, Cham, ISBN: 978-3-030-75933-9, 2021, pp. 95–127.
- [43] C. Jablonowski, D.L. Williamson, The pros and cons of diffusion, filters and fixers in atmospheric general circulation models, in: Numerical Techniques for Global Atmospheric Models, Springer, 2011, pp. 381–493.
- [44] P.A. Ullrich, D.R. Reynolds, J.E. Guerra, M.A. Taylor, Impact and importance of hyperdiffusion on the spectral element method: A linear dispersion analysis, *J. Comput. Phys.* 375 (2018) 427–446.
- [45] C.J. Vogl, A. Steyer, D.R. Reynolds, P.A. Ullrich, C.S. Woodward, Evaluation of implicit-explicit additive Runge–Kutta integrators for the homme-nh dynamical core, *J. Adv. Model. Earth Syst.* 11 (12) (2019) 4228–4244.
- [46] B. Fornberg, A Practical Guide to Pseudospectral Methods, Vol. 1, Cambridge Univ. Press, Cambridge, UK, 1998.
- [47] N.J. Zabusky, M.D. Kruskal, Interaction of solitons in a collisionless plasma and the recurrence of initial states, *Phys. Rev. Lett.* 15 (6) (1965) 240–243.