



HAL
open science

A Unification and Investigation of Rational Approximation of Exponential Integration Methods

Martin Schreiber, Jed Brown

► **To cite this version:**

Martin Schreiber, Jed Brown. A Unification and Investigation of Rational Approximation of Exponential Integration Methods. 2023. hal-04363335v1

HAL Id: hal-04363335

<https://hal.science/hal-04363335v1>

Preprint submitted on 24 Dec 2023 (v1), last revised 8 Oct 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

1 **A GENERALIZED RATIONAL APPROXIMATION OF**
2 **EXPONENTIAL INTEGRATION (REXI) FOR MASSIVELY**
3 **PARALLEL TIME INTEGRATION**

4 MARTIN SCHREIBER* AND JED BROWN†

5
6 **Abstract.** Solving partial differential equations (PDEs) is one of the most traditional tasks in
7 scientific computing. In this work, we consider numerical solutions of initial value problems (IVPs)
8 problems partly or entirely given by linear PDEs and how to compute solutions with a method
9 we refer to as rational approximation of exponential integration (REXI). REXI replaces a typically
10 sequential timestepping method with a sum of rational terms, leading to the possibility to parallelize
11 over this sum. Hence, this method can potentially exploit additional degrees of parallelization for
12 scaling problems limited in their spatial scalability to large-scale supercomputers.

13 The main contribution of this work lies in developing the “unified REXI” in which we show
14 algebraic equivalence to other methods developed up to five decades ago. Such methods cover, e.g.,
15 diagonalization of the Butcher table for implicit Runge-Kutta methods, Cauchy-contour integration-
16 based methods, and direct approximations. To our best knowledge, this is the first time of such a
17 comparison and deep investigation of all these methods.

18 Finally, we will show the applicability of REXI to the nonlinear shallow-water equations on the
19 rotating sphere, including HPC results. While previous REXI studies have focused on exposing more
20 parallelism to enable faster time to solution, we also consider efficiency at prescribed accuracy and
21 find that diagonalized Gauss Runge-Kutta methods (formulated as REXI) are compelling highly
22 efficient methods.

23 **Key words.** Exponential integrators, rational approximation, parallel-in-time, Cauchy contour,
24 Butcher table, diagonalization

25 **AMS subject classifications.**

26 **1. Introduction.** Time integration of IVPs is one of the most traditional tasks in
27 scientific computing, having seen two centuries of research. The IVPs we are interested
28 in are given entirely or partly by linear autonomous PDEs, which are ubiquitous in
29 applications ranging from daily weather forecasting [11] to full waveform inversion
30 [43]. Integration of such systems is sequential in time using conventional methods such
31 as explicit and diagonally implicit Runge-Kutta [29, 21]: Without special structure
32 [20], the state at each stage is necessary to compute the next stage, either explicitly
33 or implicitly. The time step size is typically limited by stability and/or accuracy
34 requirements and the method is purely sequential in the time dimension.

35 With the desire to solve PDEs with ever-higher resolutions, the demands on high-
36 performance computers (HPC) have increased. The steady and ongoing increase in

*Univ. Grenoble Alpes / Laboratoire Jean Kuntzmann / Inria, Grenoble, France (martin.schreiber@univ-grenoble-alpes.fr), Department of Informatics, Technical University of Munich, Germany (martin.schreiber@tum.de), <https://www.martin-schreiber.info>

†Department of Computer Science, University of Colorado, Boulder, USA (jed@jedbrown.org, <https://jedbrown.org>)

Submitted to the editors DATE.

Funding: This project has received funding from the Federal Ministry of Education and Research and the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955701. The JU receives support from the European Union’s Horizon 2020 research and innovation programme and Belgium, France, Germany, Switzerland.

Martin Schreiber gratefully acknowledges KONWIHR funding as part of the project “Parallel in Time Integration with Rational Approximations targeting Weather and Climate Simulations”.

Jed Brown acknowledges support from the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, applied mathematics program.

37 HPC performance is provided almost exclusively by increased parallelism; increasing
 38 resolution in space (spatial scalability) can be solved in the same amount of time per
 39 time step, but the wallclock time to simulate for a fixed physical duration increases due
 40 to the increasing number of time steps to satisfy the Courant-Friedrichs-Lewy (CFL)
 41 constraint [9] for transport phenomena. Consequently, refinement to increase accuracy
 42 on a transient physical problem is always a scaling challenge, and many applications
 43 are unable to increase spatial resolution without sacrificing external timelines such as
 44 IPCC assessment reports [6] or design/manufacturing timelines. Parallelism in the
 45 time dimension is seen as an opportunity to utilize greater parallelism to meet strin-
 46 gent simulation timelines. The Rational Approximation of Exponential Integration
 47 (REXI) family of methods, which we briefly explain next, are a promising candidate
 48 for hyperbolic PDEs. Consider a linear autonomous PDE given by $\frac{\partial U(t)}{\partial t} = \mathcal{L}U(t)$
 49 with $U(t)$ the current state and \mathcal{L} a linear differential operator. Discretizing state
 50 variable and operator in space leads to

$$51 \quad (1.1) \quad \frac{\partial U(t)}{\partial t} = LU(t)$$

52 with L the discrete linear operator and $U(t)$ the discrete state variables at time t .
 53 Solving such IVPs have been intensively studied over the last decades with various
 54 approaches, and one of the direct methods is the application of an exponential inte-
 55 gration

$$56 \quad (1.2) \quad U(t + \Delta t) = \exp(\Delta t L)U(t)$$

57 with the solution $U(t)$ at time t . We want to emphasize that no time discretization has
 58 been introduced and that the only approximations are related to space. The REXI
 59 method exploits the feature that $\exp(\Delta t L)$ only needs to be approximated within a
 60 spectrum related to time step size Δt and the spectrum of L . In the present work,
 61 we can express a variety of different time integration methods by what we refer to as
 62 the “unified REXI” formulation given by

$$63 \quad (1.3) \quad U(t + \Delta t) \approx \gamma U(t) + \underbrace{\sum_{n=1}^N \beta_n (\Delta t L - \alpha_n)^{-1} U(t)}_{\text{Parallelization}}$$

64 with the time step size given by Δt , and (typically) complex valued REXI coefficients
 65 α_n , β_n and real-valued γ is a new generalization that we will use in the following
 66 sections. The remainder of this work investigates different ways to infer these REXI
 67 coefficients and their relation to a class of Runge-Kutta methods. Based on this, we
 68 will study their numerical properties in the linear and nonlinear context.

69 **2. Related Work.**

70 **2.1. Exponential integration.** Exponential integration methods are formu-
 71 lated for nonlinear systems factored as

$$72 \quad (2.1) \quad \frac{\partial U(t)}{\partial t} = LU(t) + N(U(t)),$$

73 where the linear part L is intended to capture the “fast” dynamics and N is the
 74 remaining nonlinear part. An exact ansatz for advancing this split equation over a

75 finite time interval is given by

76 (2.2)
$$U(t + \Delta t) = \exp(\Delta t L)U(t) + \int_0^{\Delta t} \exp((\Delta t - \tau)L)N(U(t + \tau))d\tau.$$

77 In this form, the linear parts are integrated precisely by an exponential function,
 78 hence overcoming potential stiffness challenges caused by the linear parts. Due to
 79 this advantageous property, the interest in these exponential integrators has steadily
 80 increased over the last decades (see, e.g., [25, 19]) where various approaches have been
 81 taken to approximate the integral of the nonlinearities. One of the most commonly
 82 known approximations of the integral is, e.g., given by (see [10])

83 (2.3)
$$U(t + \Delta t) = \varphi_0(\Delta t L)U(t) + \Delta t \varphi_1(\Delta t L)N(U(t))$$

84 where we used the notations $\varphi_0(Z) = e^Z$ and $\varphi_1(Z) = \frac{e^Z - I}{Z}$. We skip further examples
 85 for discretized exponential integrator formulations and only like to point out the φ
 86 functions to be omnipresent in higher-order exponential integration methods, which
 87 is generally given, e.g., by

88 (2.4)
$$\varphi_{i+1}(Z) = (\varphi_i(Z) - \varphi_i(0)) Z^{-1} \quad \text{for } i \geq 0.$$

89 An investigation of all different varieties of discretizations of exponential integra-
 90 tors incorporating the nonlinearities is beyond the scope of this work, and we continue
 91 on the linear parts. These linear parts can be either given by full linear PDEs or by
 92 time integrating only a part of linear PDEs where the underlying requirement of time
 93 integration results in problems of the form $U(t+\Delta t) = \varphi_0(\Delta t L)U(t) = \exp(\Delta t L)U(t)$.
 94 In contrast to state-of-the-art time integration methods, which are used in operational
 95 codes, exponential integrators for linear operators avoid any time-discretization er-
 96 rors. However, the computational complexity can be tremendous and triggered the
 97 development of various ways to tackle this challenge [25]. We will briefly summarize
 98 the ones recently researched, namely based on Krylov subspaces and REXI.

99 The exponential can be approximated using Krylov subspace solvers (see [26,
 100 39, 40, 8]) where we see polynomial approximations (e.g., based on Chebyshev) as
 101 a subclass of them. The advantage of such methods is their simplicity – assuming
 102 the Krylov solver framework given – since only vector multiplications with the linear
 103 operator are required. However, the potential drawbacks of Krylov subspace solvers
 104 are their inherent property of sequential iterations over the Krylov subspace, hence
 105 not providing ways to exploit additional degrees of parallelization. An alternative is
 106 to use the REXI method, which will be discussed in the next section in further detail.

107 **2.2. REXI.** The particular way to evaluate the φ functions in the present work
 108 is strongly related to Padé approximations, which can be used as a first instance to
 109 approximate the φ functions. This approximation is most naturally related to how
 110 all Runge-Kutta formulations, e.g., based on the Butcher table, can be formulated for
 111 linear autonomous operators. However, higher-order polynomials in the denominator
 112 also make the development of solvers for such Padé approximations more challenging,
 113 and a partial fraction decomposition can be used. This well-known decomposition
 114 transforms a higher-order Padé approximation into a sum of lower-order terms, which
 115 can eventually be used to develop solvers parallelized over all terms. Although not
 116 being explained in the context of Padé approximations, REXI methods using the
 117 partial fraction decomposed form have been developed in different contexts. They

118 can be interpreted as a Padé approximation, which is why they are mentioned in this
 119 context.

120 In what follows, we will provide an overview of different methods, which can all
 121 be phrased in REXI form. One of the earliest REXI formulations for hyperbolic PDE
 122 time integrators is related to the Laplace transformation (cf. [23, 7]). Here, the PDE is
 123 transformed with the Laplace operator, where the backward transform is conducted
 124 with a Cauchy Contour integral. This transformation can be again related to an
 125 exponential integration scheme, namely to the Cauchy Contour method mentioned
 126 below, see also [41]. More recently, time integration based on Laplace transformations
 127 with a circle-based Cauchy contour integration have been more intensively studied
 128 in [28] with ODEs, in particular filtering properties. However, it needed a more
 129 extensive (community) effort to develop other, e.g., higher-order methods around
 130 them, as has been extensively the case for exponential integration methods. This lack
 131 of advancements with the REXI Laplace transform is also why we only concentrate
 132 on exponential integration-based formulations. We want to point out that the same
 133 approaches could also be taken from the Laplace transform perspective.

134 Another way to infer REXI coefficients originates from the REXI method based
 135 on Gaussian basis functions originally developed in [18], which only targets purely
 136 oscillatory problems (hence L has only imaginary or zero eigenvalues). This method
 137 also showed excellent properties regarding the wallclock-time vs. error for the linear
 138 shallow-water equations on the plane (see [33]) and on the rotating sphere [32, 34].

139 Although initially developed for analytical reasons, the Cauchy contour integra-
 140 tion method can indeed be used for REXI time integration. As pointed out above, one
 141 of the first times this has been used as a REXI-like method was with the Laplace trans-
 142 formations. However, exponential formulations (see Eq. 2.2) provide a more direct
 143 and substantial established way to integrate in time. Here, the property of φ_i being
 144 an analytical function plays a fundamental role in the Cauchy contour integration
 145 method as well as the exponentially fast converging trapezoidal rule to approximate
 146 the contour [41]. This method has already been used in different works: The ap-
 147 proximation of $\varphi_i(x)$ evaluations on scalar values has been used in various works to
 148 overcome singularities of $\varphi_{i>0}$ singularities at the origin, see, e.g., [5]. It has been
 149 used mainly for parabolic problems [35], also pointing out the potential of paralleliza-
 150 tion for the first time, as well as using a Carathéodory-Fejér method [31]. Regarding
 151 real applications, it was applied to nonlinear shallow-water equations on the rotating
 152 sphere [34], providing improved wallclock time-to-solution by using an enlarged and
 153 shifted contour to avoid numerical cancellation errors.

154 **2.3. Parallel-in-time.** Overcoming the wallclock time limitations of simula-
 155 tions, which cannot be accomplished by any further increase of parallelization in
 156 the spatial dimension, is the main focus of the parallel-in-time algorithms. Here, two
 157 different types of approaches exist: (a) minimally-invasive methods that take existing
 158 time integration methods and incorporate them into an iterative-in-time correction
 159 scheme (see, e.g., Parareal [22] and PFASST [24]); and (b) invasive methods that
 160 replace an existing time stepping with one that works entirely differently. Very often,
 161 one likes to use a combination of these approaches to enhance the convergence speed
 162 of the correction scheme in time. REXI is an invasive parallel-in-time algorithm (see
 163 [33]) since it requires efficient complex-valued solvers for each REXI term.

164 **3. Unified REXI formulation.** We start directly with the REXI formulation
 165 which will provide a standard fundament for the different variants to infer REXI
 166 coefficients. Given a discrete linear operator L , we can use an eigendecomposition

167 $L = Q\Lambda Q^{-1}$ with the eigenvectors stored in the columns of Q and the eigenvalues
 168 placed correspondingly on the diagonal of Λ .

169 (3.1)
$$\frac{\partial U(t)}{\partial t} = LU(t) = Q\Lambda Q^{-1}U(t)$$

170 where Q and Λ are the matrices with the eigenvectors and eigenvalues on the diagonal,
 171 respectively. In terms of the characteristic variable $u = Q^{-1}U$ and due to diagonal-
 172 only Λ , we get independent equations of the form

173
$$\frac{\partial u_i(t)}{\partial t} = \lambda_i u_i(t)$$

174 with λ_i the individual Eigenvalues on the diagonal of Λ . In characteristic variables,
 175 the unified REXI formulation (1.3) becomes

176 (3.2)
$$u_i(t + \Delta t) \approx \gamma u_i(t) + \sum_{n=1}^N \beta_n (\Delta t \lambda_i - I \alpha_n)^{-1} u_i(t).$$

177 Since each component u_i is decoupled, we can freely drop the subscript. For the
 178 purpose of time integration, the linear operator L is completely described by its
 179 eigenvalues λ , where imaginary components $\Im(\lambda)$ represent oscillation and negative
 180 real values $\Re(\lambda) < 0$ describe a diffusive/damping behavior. Note that substituting
 181 $\lambda = 1, t = 0, \Delta t = x, u(0) = 1$ in (3.2) yields $\exp(x) = \gamma + \sum_n \beta_n (x - \alpha_n)^{-1}$, which
 182 provides intuition as a sum of rational functions.

183 **3.1. Exploiting symmetry of coefficients.** We note that it is possible to
 184 reduce the workload by a factor of two for real-valued operators L when the poles
 185 α consist of complex conjugate pairs (see, e.g., [23, 18]). This optimization does not
 186 change the relative performance of the methods we consider here, so for simplicity,
 187 we do not apply it.

3.2. REXI-derived higher-order φ forms. Particular higher-order exponen-
 tial time integrators such as (2.3) require evaluations of higher-order $\varphi_{i|i>0}$. REXI
 coefficients for these functions are so far computed with methods tailored to them, see
 [18, 32]. We briefly present an new alternative way to compute them which is easily
 applicable. Given REXI coefficients for

$$\varphi_i(x) \approx \gamma + \sum_n \beta_n (x - \alpha_n)^{-1}$$

188 we can compute higher-order REXI approximations with

189
$$\varphi_{i+1}(x) = \frac{\varphi_i(x) - \varphi_i(0)}{x} = \frac{\gamma + \sum_n \frac{\beta_n}{x - \alpha_n} - \varphi_i(0)}{x}$$

 190
$$= \sum_n \left(\frac{\beta_n}{\alpha_n(x - \alpha_n)} \right) + \frac{1}{x} \underbrace{\left(\sum_n \left(-\frac{\beta_n}{\alpha_n} \right) + \gamma - \varphi_i(0) \right)}_{=0} = \sum_n \frac{\frac{\beta_n}{\alpha_n}}{x - \alpha_n}.$$

191 The cancellation of the terms is a consequence of the stationary modes which require
 192 $\sum_n \left(-\frac{\beta_n}{\alpha_n} \right) + \gamma = \varphi_i(0)$. Note that this leads to different coefficients compared to
 193 tailored computations.

194 **3.3. Linear solvers for REXI terms.** Efficient solvers are required for each
 195 REXI term. Over the last decades, this efficiency aspect turned out to be a very
 196 challenging task. E.g., in the context of shallow-water equations, this results in the
 197 original Helmholtz problem (rather than a backward Euler time step) where it is
 198 known that no off-the-shelf solvers such as GMRES and multigrid methods work in
 199 a highly-scalable way (see, e.g., [12]). This is ongoing research, and in the present
 200 work, we are using solvers developed in spherical harmonics formulations, hence a
 201 solver tailored particularly to one PDE problem.

202 **4. REXI methods.** These sections cover various ways to infer REXI coefficients,
 203 which represent, from our point of view, the most interesting cases. The goal is not
 204 to show all methods in great detail but their fundamental properties.

205 Although we present methods in characteristic form (3.2), the proposed methods
 206 also hold in system form (1.3). In the following, we will use the error

$$207 \quad (4.1) \quad e(z) = \left| \gamma + \sum_n \beta_n (z - \alpha_n)^{-1} - \exp(z) \right|$$

208 to compute the deviation from $\varphi_0(z) = \exp(z)$ with $z = \lambda\Delta t$ denoting the point on
 209 the complex plane to evaluate. Since approximating diffusive problems is relatively
 210 straightforward, we focus on purely oscillatory problems with $\lambda \in i\mathbb{R}$. The REXI
 211 methods we consider have complex-conjugate poles α , thus $e(z) = e(\bar{z})$ and so we
 212 only plot errors for $\Im[z] \geq 0$.

213 **4.1. B-REXI: Butcher/Bickart.** A Butcher table [2] provides a canonical rep-
 214 resentation of s -stage Runge-Kutta methods [29, 21] in terms of a matrix $A \in \mathbb{R}^{s \times s}$
 215 and completion vector $b \in \mathbb{R}^s$, with $c = A\mathbf{1}$ determining the abscissa (which we will
 216 see are related to REXI poles and $\mathbf{1}$ is a column vector of ones). The coefficients are
 217 selected to achieve the desired order of accuracy and stability properties as well as
 218 solution procedure, such as explicit, diagonally implicit, and fully implicit.

219 For fully nonlinear and non-autonomous ODEs $\frac{\partial u}{\partial t} = f(t, u)$, a Runge-Kutta
 220 method in Butcher form requires solving a system of stage equations

$$221 \quad (4.2) \quad y_s = u_n + \Delta t \sum_{j=1}^S A_{sj} f(t + c_j \Delta t, y_j), \quad i = 1, \dots, S$$

222 and evaluating the completion formula $u_{n+1} = u_n + \Delta t \sum_{j=1}^S b_j f(t + c_j \Delta t, y_j)$. Here,
 223 Δt is the time step size, and $\mathbf{y} = \{y_j\}_{j=1}^S$ is the vector of stage solutions. For
 224 linear autonomous equations, we can choose characteristic variables, in which case
 225 $f(t, u) = \lambda u$, and the stage equations (4.2) reduce to $\mathbf{y} = \mathbf{1}u + \Delta t \lambda A \mathbf{y}$ and

$$226 \quad (4.3) \quad u_{n+1} = \underbrace{[1 + \Delta t \lambda b^T (I - \Delta t \lambda A)^{-1} \mathbf{1}]}_{R(\Delta t \lambda)} u_n,$$

227 where we have identified the stability function $R(z) \approx \exp(z)$.

228 **4.1.1. Derivation.** We now show that unified REXI is algebraically equivalent
 229 to Runge-Kutta methods with a diagonal Butcher matrix A , starting with a decom-
 230 position inspired by the solution method developed independently by [3, 1]. Given an
 231 eigendecomposition $A = EDE^{-1}$ (which exists for the collocation methods we will
 232 consider [16]), we can rewrite (4.3) as

$$233 \quad (4.4) \quad u_{n+1} = [1 + \Delta t \lambda b^T E (I - \Delta t \lambda D)^{-1} E^{-1} \mathbf{1}] u_n.$$




FIG. 4.1. Error studies for the B-REXI method with (a) Gauss-Legendre and (b) Chebyshev quadrature points for the error given in Eq. (4.1). Each color refers to the same number of stages, marked to B-REXI or RK form. The non-diagonalized version provides significantly better results compared to the diagonalized version. In particular, results with B-REXI using 32 or more stages suffer from significant defects in the solution.

234 With $W = \text{diag}(E^{-1}\mathbf{1})^{-1}$, we may transform to

235 (4.5)
$$u_{n+1} = \left[1 + \Delta t \lambda \underbrace{b^T E W^{-1}}_{\tilde{b}^T} (I - \Delta t \lambda D)^{-1} \underbrace{W E^{-1} \mathbf{1}}_1 \right] u_n,$$

236 which is a diagonal Runge-Kutta method with A replaced by D and the original
 237 completion vector b replaced by \tilde{b} . Rewriting this to a REXI form leads to

238 (4.6)
$$\begin{aligned} u_{n+1} &= u_n + \Delta t \tilde{b}^T \left(-(\Delta t D)^{-1} \right) \left(I + (\Delta t \lambda D - I)^{-1} \right) \mathbf{1} u_n \\ &= \underbrace{\left(1 - \tilde{b}^T D^{-1} \mathbf{1} \right)}_{\gamma} u_n + \underbrace{\left(-\tilde{b}^T D^{-2} \right)}_{\beta^T} \left(\Delta t \lambda - \underbrace{D^{-1}}_{\text{diag}(\alpha)} \right)^{-1} \mathbf{1} u_n. \end{aligned}$$

239 Finally, we can write this in the unified REXI formulation (1.3) with

240 (4.7)
$$\gamma = 1 - \tilde{b}^T D^{-1} \mathbf{1} \quad \beta^T = -\tilde{b}^T D^{-2} \quad \alpha = \text{diag}(D^{-1}).$$

242 We have derived a transformation from implicit RK method with nonzero eigenval-
 243 ues to REXI form with the same stability function. Given a REXI method, one
 244 can construct an equivalent diagonal RK method (with complex coefficients) via
 245 $D = \text{diag}(\alpha)^{-1}$ and $\tilde{b}^T = -\beta^T D^2$. Note that a conventional Butcher table A, b^T
 246 is not uniquely determined by this procedure. We remark that standard techniques
 247 for analyzing Runge-Kutta methods can readily be applied to REXI methods. This
 248 includes barriers such as Theorem 4.3 of [20], which establishes that diagonal (parallel)
 249 RK methods can be no more than second order accurate for nonlinear problems.

250 **4.1.2. Error studies.** We choose the Gauss-Legendre and Chebyshev quadra-
 251 ture points for the error studies, with results given in Figure 4.1. We can observe
 252 that increasing the number of stages in the non-diagonalized version (using a dense
 253 Butcher table) always improves accuracy per stage. In contrast, B-REXI accuracy
 254 degrades when too many stages are used, becoming apparent beyond 8 stages. This
 255 effect is related to ill-conditioning that can be interpreted via the condition number
 256 of the eigenbasis E that effects diagonalization (4.5) or via the 1-norm of the com-
 257 pletion vector \tilde{b} , as shown in Figure 4.2. Note that completion vectors must sum
 258 to 1 so $\|\tilde{b}\|_1 = 1$ is optimal (and indeed holds for the original completion vector b);




FIG. 4.2. Condition number of the eigenbasis E for the B-REXI method on Gauss-Legendre collocation points (left) and 1-norm of completion vector $\tilde{\mathbf{b}}$ for the diagonalized method (right). The rounding errors incurred by the exponential growth precludes use of this approach for many stages.

259 a large 1-norm indicates the existence of large positive and negative entries, leading
 260 to cancellation errors. Despite this downside, the numerical experiments of §6 will
 261 show that these B-REXI (diagonalized Gauss Runge-Kutta) methods with lower stage
 262 counts are remarkably efficient compared to the other (better-conditioned) families.

263 **4.1.3. Relation to Crank-Nicolson.** Since this will be relevant for the results
 264 section, we would like to show the relation between the B-REXI approximation with
 265 a single pole using the Gauss-Legendre quadrature using just a simple pole (centered
 266 at the interval). This will lead to the terms $\gamma = -1$, $\alpha = 2$ and $\beta = -4$ which
 267 yields the REXI approximation $\exp(x) \approx -1 + \frac{-4}{x-2} = \frac{1+\frac{1}{2}x}{1-\frac{1}{2}x}$ with the equation on the
 268 right hand side matching the Crank-Nicolson formulation. This REXI approximation
 269 with a single term resembles the Crank-Nicolson formulation with a midpoint rule
 270 (forward Euler on nominator and backward Euler on denominator for $x = \Delta t L$ and a
 271 half-time step size). This will explain that later numerical results with B-REXI match
 272 the Crank-Nicolson method. Using more REXI poles will result in even higher-order
 273 approximations.

274 **4.2. T-REXI: Terry's Rational Approximation of the Exponential In-**
 275 **tegrator.** The approach which we will refer to as T-REXI was introduced in [18].
 276 Several steps are required to gain the α and β coefficients. Since these steps account
 277 for the computational workload and the properties, we briefly describe the derivation,
 278 including a discussion on the advantages and limitations of this method.

279 **4.2.1. Derivation.** The first step consists of an approximation of a Gaussian
 280 basis function $\psi_c(x)$ as follows:

$$281 \quad (4.8) \quad \psi_h(x) = (4\pi)^{-\frac{1}{2}} e^{-x^2/(4h^2)} \approx \operatorname{Re} \left(\sum_{k=-W}^W \frac{\omega_k}{i\frac{x}{h} + (\mu + ik)} \right)$$

282 Using $W = 11$, hence $L = 2W + 1 = 23$ terms in total, is sufficient for an accurate
 283 approximation up to numerical double precision (see [18]). The advantage of this
 284 representation is an efficient representation of the Gaussian basis function in Fourier
 285 space. The proxy with the Gaussian basis function allows for computing the coeffi-
 286 cients ν_k for an approximation of an oscillatory function within an approximate range

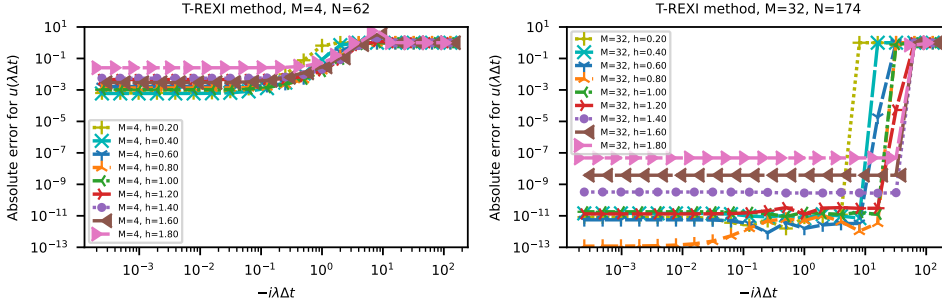


FIG. 4.3. Error studies for the T-REXI method for different M and h values. M relates to the number of Gaussian functions via $2M+1$ to approximate an oscillation with $N = 2(2M+L)$ number of REXI terms. Left image: We can observe a very high error for a low number of Gaussian basis bumps, which cannot be improved by changing h . Right image: Using significantly more Gaussian functions leads to significant improvements. In particular, we observe that optimal values for h influence the quality of the approximation. An optimum can be observed for $h \approx 0.8$.

287 $x \in [-Mh; Mh]$ in Fourier space, yielding

288 (4.9)
$$\exp(ix) \approx \sum_{k=-M}^M \nu_k \psi_h(x + kh).$$

289 Both steps are then combined, resulting in the approximation

290 (4.10)
$$\operatorname{Re}(\exp(ix)) \approx \sum_{n=-M-W}^{M+W} \operatorname{Re}(\beta_n^{\operatorname{Re}}(ix + \alpha_n)^{-1})$$

291 where we only showed the Re one. We combine Re and Im to the form

292 (4.11)
$$\exp(ix) \approx \sum_{n=-M-W}^{M+W} \beta_n(ix - \alpha_n)^{-1}$$

293 eventually leading to the REXI formulation with $\gamma = 0$, but x related directly to the
 294 imaginary value on the complex plane. So far, we only targeted the φ_0 function, and
 295 we like to point out that this method can also be used to approximate other φ_i terms
 296 (see [18]) or directly with the REXI coefficients (see §3.2). We want to emphasize
 297 that this approximation was derived only for purely oscillatory functions and, hence,
 298 does not include approximations with non-zero real eigenvalues components.

299 **4.2.2. Error studies.** We investigate the errors of the T-REXI method in Fig-
 300 ure 4.3. On the left image, we can observe that we need a minimum number of
 301 Gaussian basis functions to approximate the oscillations. The right image shows ex-
 302 ceptionally accurate results for $h \approx 0.8$ in the range $x \in [0; 10]$ and a rather large
 303 region of accuracy of about $e(x = 128) \leq 10^{-11}$. Other figures (not included) show
 304 that increasing M leads to a linear increase of the size of the region of high accuracy
 305 (see [18]) with an optimum value of $h \approx 0.8$. For the remainder of this work, we will
 306 use $h \approx 1.0$ as a compromise between accuracy and total workload.

307 **4.3. CI-REXI: Cauchy Contour Integral method.** Cauchy Contour Inte-
 308 gral (CI) methods offer an alternative way to infer the REXI coefficients (see e.g.
 309 [42, 4, 34]). We start with the general CI equation given by

310 (4.12)
$$g(x) = \frac{1}{2\pi i} \oint_{\Gamma} \frac{g(z)}{z - x} dz$$

311 where $g(x)$ is one of the analytic φ_i functions (2.4), Γ the contour enclosing the
 312 eigenvalue λ for ODEs and all eigenvalues on the diagonal of Λ for PDEs.

313 **4.3.1. REXI Derivation.** Regarding the contour, we can use different ap-
 314 proaches. In what follows, we used parametrized contours $\Gamma = \{\sigma(w)|w \in [0; 1]\}$
 315 with the contour function $\sigma(w) : \mathbb{R} \rightarrow \mathbb{C}$. Using integration by substitution and the
 316 contour function, we obtain

$$317 \quad (4.13) \quad g(x) = \frac{1}{2\pi i} \oint_0^1 \frac{g(\sigma(w))}{\sigma(w) - x} \sigma'(w) dw = \oint_0^1 \frac{i(2\pi)^{-1} g(\sigma(w)) \sigma'(w)}{x - \sigma(w)} dw.$$

318 Using the trapezoidal rule, which is exponentially fast converging on periodic bound-
 319 aries (see [41]) with N trapezoidal points in total, we obtain

$$320 \quad (4.14) \quad g(x) \approx \frac{1}{N} \sum_{n=1}^N \frac{i(2\pi)^{-1} g(\sigma(w_n)) \sigma'(w_n)}{x - \sigma(w_n)} \quad \text{with} \quad w_n = \frac{n}{N}.$$

321 Again, we can infer a unified REXI formulation (1.3) by setting

$$322 \quad (4.15) \quad \alpha_n = \sigma(w_n) \quad \beta_n = \frac{ig(\sigma(w_n))\sigma'(w_n)}{N2\pi} \quad \gamma = 0.$$

323 An ellipse contour is given by $\sigma(w) = R_x \cos(iw2\pi) + iR_y \sin(iw2\pi) - \mu$ with μ
 324 related to the center of the ellipse. This leads to the coefficients

$$325 \quad (4.16) \quad \alpha_n = R_x \cos(iw2\pi) + iR_y \sin(iw2\pi) - \mu$$

$$326 \quad (4.17) \quad \beta_n = \frac{i}{N} \exp(\sigma(w)) (-R_x \sin(iw2\pi) + iR_y \cos(iw2\pi))$$

327 A study of all kinds of contour shapes (rectangle, bean, polygonal shapes, etc.) is
 328 beyond the scope of this work. In the next section, we will mainly focus on the circle
 329 to show interesting characterizations and use the ellipse for numerical studies to show
 330 its superiority to another REXI method. In the following, we will refer to the special
 331 case of a circle as CI-REXI and to the ellipse case as CI-EL-REXI.

332 **4.3.2. Characterization and numerical issues.** Next, we characterize the
 333 REXI terms, referred to as the β characterization, with an overview in Figure 4.4.
 334 We remind the reader that REXI approximates functions with a linear combination
 335 of rational basis functions. Depending on the placement of these functions (related
 336 to α_n) and the weighting of each basis (related to β_n), we have three different cases:

337 **a) Obsolete REXI terms:** Contours $Re(\sigma(x)) \rightarrow -\infty$ relating to areas of the
 338 contour in the distant negative real axis on the complex plane have exponentially
 339 fast decaying β coefficients, hence $\lim_{\sigma \rightarrow -\infty} \beta_n = \lim_{\sigma \rightarrow -\infty} \frac{i \exp(Re(\sigma(w_n))) \sigma'(w_n)}{N2\pi} = 0$.
 340 Once a particular β_n coefficient undershoots a threshold ϵ_β , the corresponding REXI
 341 term can be removed if $\beta_n < \bar{\epsilon}_\beta$ and $\bar{\epsilon}_\beta = \epsilon_\beta/N$. The last equation incorporates that
 342 a higher numerical resolution results in smaller values of the β weights.

343 **b) Regular REXI terms:** This characterization refers to those REXI terms
 344 that can be incorporated in the approximation in a useful way.

345 **c) Cancellation-prone REXI terms:** These terms are related to the contour
 346 $Re(\sigma(x)) \rightarrow +\infty$. Approximating the exp function in the far distance to the right
 347 of the origin leads to exponentially increasing the β values. An oscillatory function
 348 is also parallel to the imaginary axis, which is approximated. Both effects lead to




FIG. 4.4. Complex plane for the real (left image) and imaginary (right image) value of $\exp(x)$. We highlight the different areas related to the different β characterizations.




FIG. 4.5. Selection of contours used with the CI-REXI method.

349 very large positive and negative numbers, resulting in severe cancellation errors in
 350 this region. Consequently, this region should be avoided.

351 Examples of different contours are provided in Figure 4.5. Each contour targets a
 352 particular problem. The circle can be used for the approximation of a small spectral
 353 radius $\lambda\Delta t < 10$. Once requiring a larger approximation along the imaginary axis, the
 354 radius cannot be enlarged without sacrificing accuracy due to cancellation errors in
 355 β_n , see (c) above. This can be avoided by enlarging the radius and choosing the value
 356 μ , hence shifting the circle, to exclude a contour across areas with $Re(x) > 10$, which
 357 leads to the shifted circle. Other contours are, e.g., given by the ellipse or rectangle
 358 targeting the approximation of a spectrum on or close to the imaginary axis and
 359 the bean contour targeting an approximation of diffusive and oscillatory problems.
 360 Studies about these contours are beyond the scope of this work and we will focus on
 361 the (shifted) circle and ellipse throughout the remainder of this paper.

362 **4.3.3. Error studies.** We conduct error studies using the shifted circle CI-REXI
 363 method. The first study is based on a circle centered at the origin, with studies for
 364 different radii. The second is for a circle which is shifted to overcome problems related
 365 to the cancellation effects (see (c) above).

366 Results are given in Figure 4.6 with plots based on a fixed number of $N = 256$




FIG. 4.6. Error studies for the centered circle (left) and shifted circle (right) CI-REXI method. The centered circle suffers from cancellation effects for large radii, whereas the shifted circle limits these effects. In particular, for a larger imaginary spectrum to be approximated, adding more REXI poles leads to improved accuracy, which is not the case for the centered circle CI-REXI method.




FIG. 4.7. Error studies for the shifted circle with $N = 128$ REXI poles (left) and $N = 2048$ REXI poles (right) with different ϵ_β pruning values. We can observe significant reductions in the number of required REXI poles.

367 REXI poles. We can observe that the errors significantly increase for the centered
 368 circle once the radius exceeds a certain threshold. In particular, errors for a larger
 369 radius – including a larger spectrum on the imaginary axis – are outside the plotting
 370 range. The results for using a higher number of REXI poles do not significantly
 371 improve the results. Using a shifted and enlarged circle, we can gain improved results
 372 that overcome cancellation errors.

373 So far, we only investigated the error itself but neglected the total workload.
 374 Pruning β with ϵ_β (exploiting characterization (a)), we can reduce some workload
 375 significantly as depicted in Figure 4.7 for larger radii. For a moderate number of REXI
 376 poles $N = 128$ (left image), we observe a pruning close to the accuracy of REXI itself,
 377 hardly impacting the results. In contrast, larger radii already suffer from inaccuracies
 378 of the used quadrature, with errors outside the plotting range. For a larger number of
 379 REXI poles $N = 2048$ (right image), we observe very robust pruning, hardly affecting
 380 the accuracy of the REXI approximation quality but leading to a significant reduction
 381 of the workload.

382 **5. Stability, normalization & filtering.** So far, we have only studied errors
 383 in approximating the φ_0 function with REXI methods. However, once we use REXI
 384 methods for time integrating differential equations, additional properties such as sta-
 385 bility and convergence are assumed to be relevant. We will investigate these properties
 386 in this section for the ODE $\frac{du(t)}{dt} = \lambda u(t)$.




FIG. 5.1. Stability plots. For B-REXI (left): we observe an excellent stability behavior known for collocation methods. T-REXI (right): We observe instabilities at the imaginary axis for the boundaries of the approximation range.




FIG. 5.2. Stability plots for CI-REXI. The left image depicts the discrete contour points chosen so that one α pole lies on the imaginary axis. This leads to instabilities. The right image depicts a half-shifted variant of it.

387 **5.1. Stability.** The stability plots are generated based on the stability function
 388 $R(\lambda)$, which is defined by the execution of a single-time step $u(t + \Delta t) = R(\Delta t \lambda)u(t)$.
 389 We will plot the amplification factor $|R(\lambda)|$ of the solution $u(t)$ over a time step
 390 $\Delta t = 1$.

391 **B-REXI** (left image in Figure 5.1): The stability reflects the A-stability of these
 392 methods on the entire left half plane. In particular, stability is given for the entire
 393 imaginary axis, a known property of collocation methods.

394 **T-REXI** (right image in Figure 5.1): We can observe that T-REXI provides
 395 excellent stability for purely imaginary values. However, we can observe instabilities
 396 on the imaginary axis once we reach the boundaries of the approximation range. This
 397 can be avoided by an additional T-REXI filter, which could be applied to obtain
 398 stability also outside the approximation range (see [18]).

399 **CI-REXI:** Finally, we look at the CI-REXI method based on Cauchy contour
 400 integral methods in Figure 5.2. The left image shows an unstable region along the
 401 imaginary axis. This is caused by an α pole directly placed on the imaginary axis.
 402 We can avoid this by choosing the support points of the trapezoidal rule differently.
 403 The right image shows a solution to this by shifting them by a half interval, effectively
 404 avoiding this instability, and CI-REXI becomes unconditionally stable for oscillatory
 405 systems. To summarize, if using the CI-REXI method, one should avoid placing poles




FIG. 5.3. Contour lines for stability plots for CI-REXI with half-shifted intervals. We can observe an excellent stability region over this entire range.

406 near the eigenvalues of the linear operator.

407 A contour plot comparing CI methods with an increasing number of poles and
 408 approximation range is provided in Figure 5.3. We adopted the contour of the circle
 409 to pass through the points $\pm 10i$, $\pm 20i$, and $\pm 40i$ on the imaginary axis for increasing
 410 the number of poles while keeping the contour never exceeding 10 on the real axis.

411 **5.2. Normalization.** This section concerns particular problems for stationary
 412 or nearby modes requiring special treatment with the T-REXI method. So far, we
 413 only assessed errors for a single time step, and this section will investigate the accuracy
 414 and conservation properties of stationary modes concerning REXI methods. We will
 415 use Dahlquist's equation (5) with $\lambda = 10^{-3}i$, which is time-integrated until $t = 100$
 416 using different REXI methods. The particular choice of this low frequency is related
 417 to almost stationary modes of PDEs. Such modes play an important role, e.g., for
 418 geostrophic balance in atmospheric simulations, and not preserving them might lead
 419 to spurious/parasitic modes.

420 An investigation of the results at the absolute ODE errors at $t = 100$ is given in
 421 Figure 5.4. The left column shows REXI methods as they have been computed with
 422 the methods from before. We can observe that the CI-REXI method (top left image)
 423 has REXI coefficients preserving the stationary modes. However, the T-REXI suffers
 424 from significant defects in it. A normalization can be used to overcome this problem
 425 where stationary modes require $\sum_n \frac{\beta_n}{x - \alpha_n} = s = 1$ and we can ensure this by simply
 426 rescaling β_n so that $\beta_n^{\text{new}} = \frac{\beta_n}{s}$.

427 The results for this are given in the right column, where we observe relatively
 428 small improvements for the CI-REXI method (right top image). However, for the T-
 429 REXI method, the errors significantly drop from 10^{-8} to about 10^{-13} once applying
 430 this normalization. We also do not see any accuracy degradation for very large time
 431 step sizes. Hence, this normalization can be used without impacting the accuracy of
 432 other choices of λ , and we will use it throughout the remainder of this work.

433 We close with two side comments: First, studies for purely stationary modes
 434 ($\lambda = 0$, not shown here) showed that the errors are increasing for smaller time step
 435 sizes, but only due to round-off errors. Overall, these results still lie within the range
 436 of numerical precision; hence, we skipped them here. Second, we skipped the B-REXI
 437 method since it is not prone to this problem for a number of REXI terms usable as
 438 solvers.

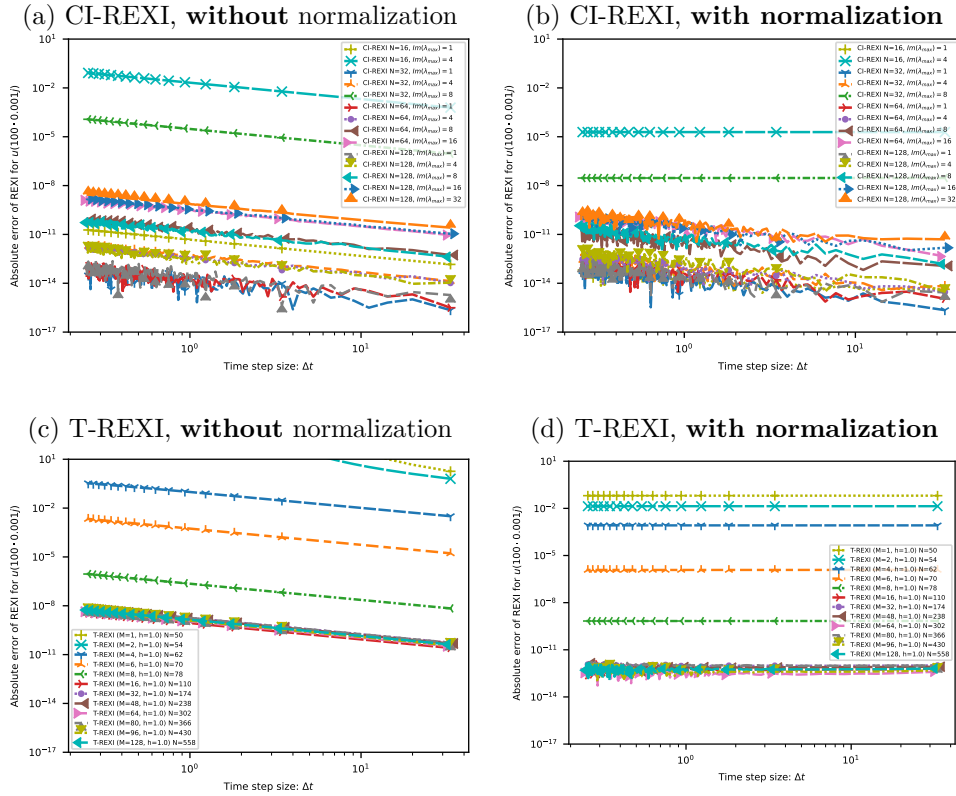


FIG. 5.4. Error studies for different REXI methods of time step size Δt vs. absolute error at $u(t = 100)$. The left column shows errors without normalization, and the right column shows errors with normalization for near-stationary modes. As we can observe in the first row, the normalization for near-stationary modes with the CI-REXI method does not lead to any significant improvements. In contrast, significant improvements can be observed for the T-REXI method. See the text for a detailed explanation of the results.

439 **5.3. Filtering.** This brief section points out the filtering capabilities of the dif-
 440 ferent REXI methods. We define a filter to apply a reduction of the amplitude of
 441 $\varphi_i(x)$ for a particular set of eigenvalues. It is, in particular, desirable to filter out
 442 (setting them close to zero) the so-called “fast modes” for x starting at a threshold
 443 and to have a smooth transition of the change in amplitude towards filtering out
 444 modes. Since diffusive problems already have a reduction of amplitude given natu-
 445 rally by their mathematical properties, we will again solely focus on purely oscillatory
 446 problems, with results also applicable to a mix of oscillatory/diffusive problems.

447 Using the **B-REXI** method, we can observe that the stability contour follows ex-
 448 actly the imaginary axis. Hence, there is no filtering at all. For the **T-REXI** method,
 449 we skip a discussion of filtering due to the inherent instability at the boundaries of
 450 the approximation range and point out to an additional filter proposed in [18]. The
 451 **CI-REXI** method has a natural filtering. This is due to the property that points
 452 outside the contour are rapidly approaching 0 as a property of the Cauchy contour
 453 integral.

454 **6. Comparison of REXI methods.** This section aims to provide guidance
 455 about which REXI method is best, and we will explore this in different ways. A

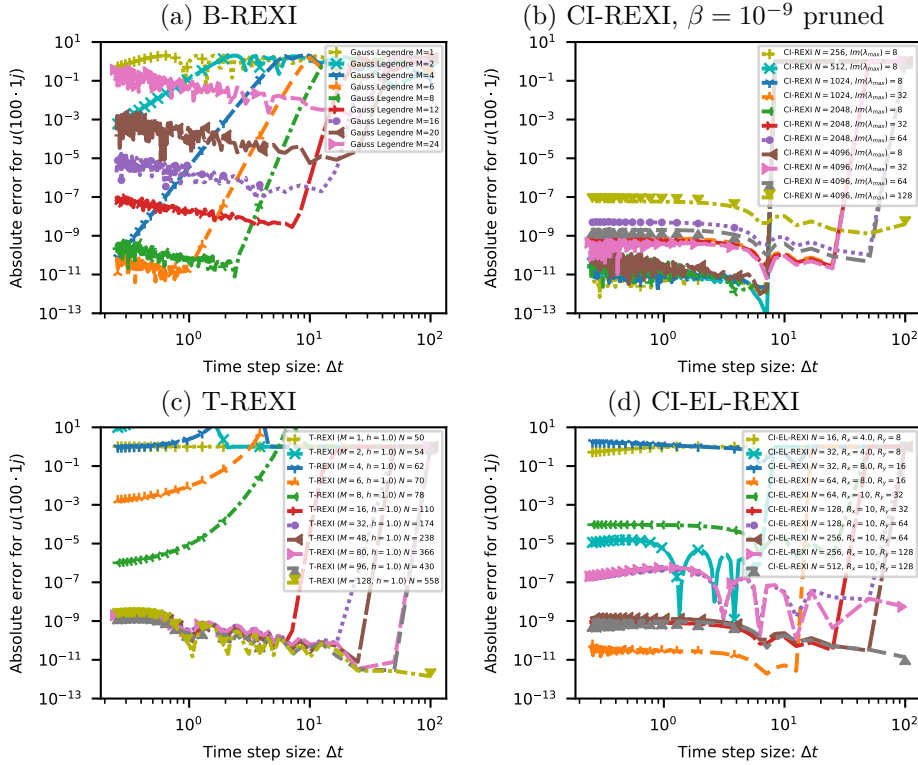


FIG. 6.1. Error studies for different REXI methods of time step size Δt vs. absolute error at $u(t = 100)$. *B-REXI* (left top) is suitable only for smaller timestep sizes. *CI-REXI* can be tuned to allow also very large time step sizes. *T-REXI* requires many poles for small time step sizes and allows also very large time step sizes. *CI-EL-REXI* allows also very large time step sizes and in addition requires the least number of poles for similar accuracy.

456 full exploration of all parameter combinations is obviously not possible. Hence, we
 457 focused on the ones that were most rational to us based on far more experiments than
 458 shown here. We first continue with concrete examples using a linear oscillatory ODE
 459 based on the Dahlquist equation followed by a PDE with the nonlinear shallow-water
 460 equations on the rotating sphere to gain insight into numerical properties once we
 461 apply this to more realistic test cases.

462 Based on the eigendecomposition, we classify linear operators as oscillatory, dif-
 463 fusive, or both. A purely oscillatory system [10] requires imaginary-only eigenvalues
 464 ($i\lambda \in \mathbb{R}$) whereas a diffusive behavior is based on negative real eigenvalues ($\lambda \in \mathbb{R}$ and
 465 $\lambda < 0$). Since oscillatory/hyperbolic systems belong to the most challenging problems
 466 for REXI methods, we will solely focus on them.

467 **6.1. ODE.** We investigate the ODE systems again with Dahlquist's equation
 468 (5) using $\lambda = 1$ and the simulation results at $t = 100$. We use $u(0) = (1 + i)/\sqrt{2}$
 469 as an initial condition. We compare various REXI methods in Figure 6.1. The total
 470 numbers of REXI coefficients are given by N .

471 The **B-REXI** method (left upper image) performs extremely well for small step
 472 sizes where only a few poles are required. For larger time step sizes of $\Delta t \approx 10$, using
 473 16 poles is sufficient to gain single precision accuracy.

474 The **CI-REXI** method (right top image) is tuned with a contour never exceeding

475 a real value of 10 and to include the points on the imaginary axis given by $Im(\lambda_{max})$.
 476 The CI-REXI method clearly outperforms the B-REXI method for medium-sized time
 477 step sizes and also allows taking very large time step sizes.

478 The **T-REXI** method (left bottom image) requires a significant number of REXI
 479 poles if only small time step sizes should be taken. This improves once larger step
 480 sizes are taken since the initial overheads of the large number of poles (due to the
 481 rational approximation of the Gaussian) have less relative impact on the number of
 482 REXI coefficients.

483 In addition, we also investigated the **CI-EL-REXI** (right bottom image) method,
 484 which is a natural choice for purely oscillatory problems. We chose the semi-major
 485 axis of the ellipse along the real axis in an empirical way and never exceeding 10 to
 486 avoid numerical issues. This method *outperforms both CI-REXI and T-REXI almost*
 487 *everywhere* regarding accuracy and number of terms required to solve it.

488 **6.2. PDE example.** In this final section, we will investigate different REXI
 489 methods with the shallow-water equations (SWE) on the rotating sphere. We decided
 490 not to investigate many different PDEs, but to go into depth of exponential integration
 491 for a single one which is of purely hyperbolic nature. We chose the SWE since they
 492 are frequently used to assess the quality and performance of discretizations in time
 493 and space concerning horizontal aspects of the full Euler equations solving the fluid
 494 dynamics equations related to the atmosphere. In velocity form, the nonlinear SWE
 495 are given by

496 (6.1)
$$\frac{\partial}{\partial t} \begin{pmatrix} \Phi \\ \vec{V} \end{pmatrix} = \underbrace{\begin{pmatrix} -\bar{\Phi}\nabla \cdot \vec{V} \\ -\nabla\Phi \end{pmatrix}}_{L_g: \text{ linear gravity}} + \underbrace{\begin{pmatrix} 0 \\ -f\vec{k} \times \vec{V} \end{pmatrix}}_{L_c: \text{ linear Coriolis}} + \underbrace{\begin{pmatrix} -\nabla \cdot (\Phi'\vec{V}) \\ -\vec{V} \cdot \nabla\vec{V} \end{pmatrix}}_N$$

497 with the horizontal velocity \vec{V} on the longitude/latitude field, geopotential $\Phi = g \cdot h$
 498 with height h , average geopotential $\bar{\Phi} = g \cdot \bar{h}$ with average height \bar{h} , a linearization
 499 around a state $\bar{h} = 10^5 m$, Coriolis effect $f = 2\Omega \sin(\phi)$ with latitude ϕ and angular
 500 rate of rotation Ω . We like to emphasize that no (hyper)viscosity is used in this PDE
 501 to avoid a simplification of the problem due to diffusive effects.

502 We use this PDE due to its particularly interesting features: The linear gravity
 503 term L_g is the stiffest one and can be solved with exponential integrators either
 504 analytically or with REXI. We want to point out that a comparison of some methods
 505 has already been under investigation in former work [34] but solely with the CI-
 506 REXI method and the geopotential field, which has also been identified to be the
 507 best Strang-split method. Anyhow, this study also lacked comparisons with other
 508 variables, particularly other REXI methods, which will lead to new revelations, as
 509 presented in the following sections. Since including the T-REXI method would not
 510 provide any beneficial insight, since the CI-REXI method is computationally much
 511 cheaper and provides additional benefits, we skip this method in the following studies.

512 **6.2.1. Spatial discretization.** We solve these equations using the SWEET soft-
 513 ware¹ which utilizes spherical harmonics (SH) to solve these equations. Such a global
 514 spectral basis leads to a substantial reduction of spatial errors (besides a lack of non-
 515 linear interactions at the limit of resolution), hence allowing us to put the focus on
 516 time integration methods. We like to refer to [32, 15] for a detailed description of the
 517 spherical harmonics. In particular, we work with the vorticity-divergence formulation

¹<https://sweet.gitlabpages.inria.fr/sweet-www/>




FIG. 6.2. Vorticity field of barotropic instability benchmark after 8 days of inviscid shallow-water equations. We see the development of various large and small-scale vortices.

518 in spectral space to avoid spurious modes if one would convert the velocity to spectral
 519 space. The standard $\frac{2}{3}$ rule [27] is used for anti-aliasing to evaluate bi-non-linearities.

520 **6.2.2. Time stepping solvers.** We can find highly efficient solvers in spherical
 521 harmonics space for direct exponential integration (without numerical approxima-
 522 tions), REXI, and implicit Euler time integrators.

Regarding the direct exponential integration, we can straightforwardly find a direct solution using the vorticity divergence form, see also [38]. Due to orthogonality, each mode can be separately written as

$$\begin{bmatrix} \partial_t \Phi' \\ \partial_t \delta \end{bmatrix} = \begin{bmatrix} -\nabla^2 & -\bar{\Phi} \\ D & G \end{bmatrix} \begin{bmatrix} \Phi' \\ \delta \end{bmatrix} = \begin{bmatrix} D & G \\ D & G \end{bmatrix} \begin{bmatrix} \Phi' \\ \delta \end{bmatrix}$$

with the famous identity $\nabla^2 = -n(n+1)$ for this harmonic. Using $D = -\nabla^2$ and $G = -\bar{\Phi}$ for convenience, we find the eigenvectors Q and eigenvalues $diag(\Lambda)$

$$Q = \begin{bmatrix} -\sqrt{\frac{G}{D}} & +\sqrt{\frac{G}{D}} \\ 1 & 1 \end{bmatrix} \quad Q^{-1} = \begin{bmatrix} \frac{1}{2}\sqrt{\frac{D}{G}} & \frac{1}{2} \\ -\frac{1}{2}\sqrt{\frac{D}{G}} & \frac{1}{2} \end{bmatrix} \quad \Lambda = \begin{bmatrix} -\sqrt{DG} & \\ & \sqrt{DG} \end{bmatrix}.$$

523 We can then use $U(t + \Delta t) = Q \exp(\Delta t \Lambda) Q^{-1}$. From an algebraic perspective, this
 524 method matches the method in [17], which uses a rather cumbersome derivation using
 525 Laplace transforms, whereas our derivation is more elegant and short. This method
 526 is also used for investigating errors

527 For the exponential integration of the full linear terms $L = L_g + L_c$, this would
 528 relate to the Hough modes [44] and no direct exponential solution has been derived
 529 yet. Hence, it requires evaluations of the form $(\Delta t L - \alpha)^{-1}$ with complex-valued α .
 530 The first time this was solved for REXI using spherical harmonics was based on a
 531 method requiring transformations to grid space [32]. The present work is based on
 532 an implicit time stepper [37] of the form $(I - \Delta t L)^{-1}$ which has been transformed to
 533 solve a REXI term by simply using a complex-valued time step size. We also used it
 534 for the implicit time integration of L as it has been originally suggested.

535 **6.2.3. Benchmark.** Our benchmark is based on the barotropic instability test
 536 case (see [14]). This benchmark is initialized with a geostrophically balanced initial

Short notation	Description
$ERK(X, o = N)$	Explicit Runge-Kutta with order N
$IRK(X)$	Backward Euler using 2nd order Crank-Nicolson
$SS(X, Y)$	2nd order Strang-splitting as explained in the text with $F_1 = X$ and $F_2 = Y$
$EXP(X)$	Direct exponential integration on X
$REXI(X)$	A particular REXI method on X
$ETDRK(X, Y)$	2nd order ETDRK method with X being exponentially integrated and Y treated as the nonlinearity
$X + Y$	Time tendencies of terms X and Y are added

TABLE 6.1

Overview of time integration methods. Note that they can be composed together.

537 condition, which is perturbed by a small Gaussian bump (see reference for detailed
 538 initial conditions). We time integrate this system for 8 days with results in Figure 6.2.

539 **6.2.4. Time integration.** Regarding the particular Runge-Kutta (RK) based
 540 time integrators, we used 2nd order midpoint, 3rd order Heun, and classical 4th order
 541 RK. The reference solution to compute the errors is based on the 4th order RK with
 542 a time step size of $\Delta t = 5$.

543 Besides the methods already introduced, our investigation also includes the 2nd
 544 order Strang splitting (SS) method [36]. With SS, a PDE given by two terms $\frac{d}{dt}U =$
 545 $F_1(U) + F_2(U)$ can be integrated with 2nd order accuracy if a 2nd order accurate time
 546 integrator $R_{F_i}^{\Delta t}$ is provided for time step size Δt by $U(t + \Delta t) = R_{F_1}^{\frac{1}{2}\Delta t} \circ R_{F_2}^{\Delta t} \circ R_{F_1}^{\frac{1}{2}\Delta t}$.
 547 We use a function-like notation to refer to the particular time integration methods.
 548 An overview of this is given in Table 6.1 where we use X and Y as representatives for
 549 either term in the PDE such as L_g , L_c , and N or to refer to another time integrator.
 550 In the latter case, e.g., ERK , EXP , $REXI$, and IRK can both be used in the
 551 Strang-Splitting SS as arguments.

552 **6.2.5. Hardware, parallelization & batch configuration.** All results have
 553 been computed on the Thin Nodes of SUPERMUC-NG. Each node is equipped with
 554 two Intel SkylakeXeon Platinum 8174, resulting in two NUMA domains. For the
 555 spatial parallelization, we use solely OpenMP on one NUMA domain, resulting in a
 556 spatial scaling of up to 24 cores. Scalability for REXI is then based on MPI first
 557 by utilizing the 2nd NUMA domain, then other compute nodes. We gratefully ac-
 558 knowledge the usage of the SHTNS library [30] which is based on FFTW [13]. We
 559 precomputed transformation plans and reused them for all studies to ensure the uti-
 560 lization of the same ones over all studies. Each batch job is set to timeout after 1
 561 hour, which follows the idea that the simulations should be finished within a specific
 562 time frame.

563 **6.2.6. Performance comparison for splitting L_g and $L_c + N$.** We start
 564 with a comparison of standard methods in Figure 6.3 which we will use as a baseline
 565 for further comparisons with REXI-based methods. Plots are given for the three
 566 prognostic variables, which we define here as the variables required as input to one
 567 time step, since results differ for all of them.

568 First, the higher-order 3rd- and 4th-order RK method can outperform other lower-
 569 order methods for smaller time step sizes depending on the variable under study. This




FIG. 6.3. Studies *without* **REXI** methods (but using direct exponentiation) on all prognostic variables with error vs. time step size for the barotropic instability benchmark. We also include 2nd, 3rd, and 4th order Runge-Kutta based methods with gray lines.

570 is a known phenomenon for higher-order time integration methods, and we wanted
 571 to include it to also see its max. stable time step size. We are primarily interested in
 572 very large time step sizes while still having a moderately small error.

573 The best method concerning the geopotential and the divergence variable is the
 574 Strang-split $SS(ERK(L_c+N), EXP(L_g))$, which we account for by the more accurate
 575 treatment with the exponential treatment of both variables. Since the vorticity field is
 576 not treated exponentially (time tendency for this in L_g is null), there's also no benefit
 577 visible in the comparison of the vorticity field.

578 The ETD RK method itself – although assumed to be an excellent off-the-shelf
 579 method – does not provide the overall best results compared to the rather straight-
 580 forward Strang splitting. We can observe it to be the 2nd best for the geopotential
 581 and even lower ranked for the other variables. We account for that by the way a 2nd
 582 order accurate Strang-splitting is performed. This can be interpreted as a subcycling
 583 of time steps by executing two half-time steps for one of the terms (the time step size
 584 limiting one).

585 Next, we will continue with RE XI studies by comparing them with the best
 586 Strang-split exponential and implicit methods from the previous results in Figure 6.4.
 587 Overall, we can observe a 2nd order convergence even if using only a single pole for
 588 the B-RE XI methods.

589 Matching results for $SS(RE XI, ERK)$ B-RE XI $N=1$ and $SS(IRK, ERK)$ are ob-
 590 served which is explained in §4.1.3: This particular B-RE XI method resembles exactly
 591 the Crank-Nicolson method but uses one complex-valued pole to solve the system of
 592 equations.




FIG. 6.4. Studies with REXI methods on all prognostic variables with error vs. time step size for the barotropic instability benchmark.

593 The other B-REXI methods outperform all alternatives except for the direct exponential
594 integration EXP. We see one particularly interesting and highly important
595 effect: The B-REXI method does not provide any further advantages using more than
596 $N = 2$ poles. Even using $N = 4$ poles, the results are not further improved. A particularly
597 important point is the comparison of the CI-REXI method with B-REXI,
598 where absolutely no benefits are visible for $N = 128$ poles using CI-REXI compared
599 to $N = 2$ poles using B-REXI. This clearly indicates that significant computational
600 savings of a factor of 64 can be accomplished in this case compared to the former
601 work.

602 We close this section by HPC studies in Figure 6.5. For sake of better overview, we
603 only plotted the most promising candidates (ETDRK is worse than B-REXI methods,
604 the explicit RK order 3 and 4 methods are better for larger wallclock times (smaller
605 time steps), but unstable otherwise).

606 We start by comparing the performance of the direct exponential method EXP
607 with the REXI method, where we would expect that the direct method is faster, which
608 is not the case. We account for that by the direct method to be computationally more
609 intensive (square root, exponential, etc., see §6.2.2) in order to solve for this term,
610 whereas the B-REXI methods only require to evaluate two or 4 rational approxima-
611 tions. For the CI-REXI method, which requires $N = 128$ terms this is again different
612 due to the higher MPI overheads resulting in a lower performance than the others.

613 Although the Strang-splitting method with the implicit term is computationally
614 quite efficient to evaluate, its overall wallclock time performance is not optimal.

615 **6.2.7. Performance comparison for splitting into L and N .** Next, we in-
616 vestigate the performance of REXI methods using a splitting into the linear term




FIG. 6.5. Studies including REXI methods with wallclock time vs. time step size for the barotropic instability benchmark.

617 $L = L_g + L_c$ and the nonlinear term N . This leads to the situation that no direct
 618 computation of $\exp(\Delta t L)$ is possible, as previously explained. Plots are given in Fig-
 619 ure 6.6 where some data points of ETDRK are missing due to the 1h time out of the
 620 job (see discussion before).

621 For the geopotential Φ' , we can observe significant improvements in terms of
 622 accuracy. In particular, we can take very large time step sizes and still observe a con-
 623 vergence, whereas the 2nd order IRK-like methods already stagnate. With respect to
 624 the ETDRK scheme, its performance is worse compared to the best (straightforward)
 625 Strang-split methods.

626 For the vorticity η we can observe that the ETDRK method does not lead to
 627 any improvement. The best methods are the Strang-split IRK-based ones and some
 628 REXI-based methods. Hence, we do not see any improvement in the accuracy of the
 629 vorticity field by using exponential integration methods. This is kind of surprising
 630 at first glimpse since we expected a better treatment of the vorticity due to the
 631 exponential integration of the Coriolis effect. However, the errors in the nonlinear
 632 parts dominate the overall errors. Hence, this does not provide any better results.

633 The divergence δ study shows REXI methods to be the best ones. Again, the
 634 accuracy cannot be improved by using more than $N = 2$ poles. Everything beyond
 635 that would be an additional computational burden. The ETDRK methods again show
 636 a poorer performance than the more straightforward approach.

637 Finally, we investigate the wallclock time vs. errors with results given in Figure 6.7.
 638 We can observe that fully explicit ERK methods actually provide excellent results due

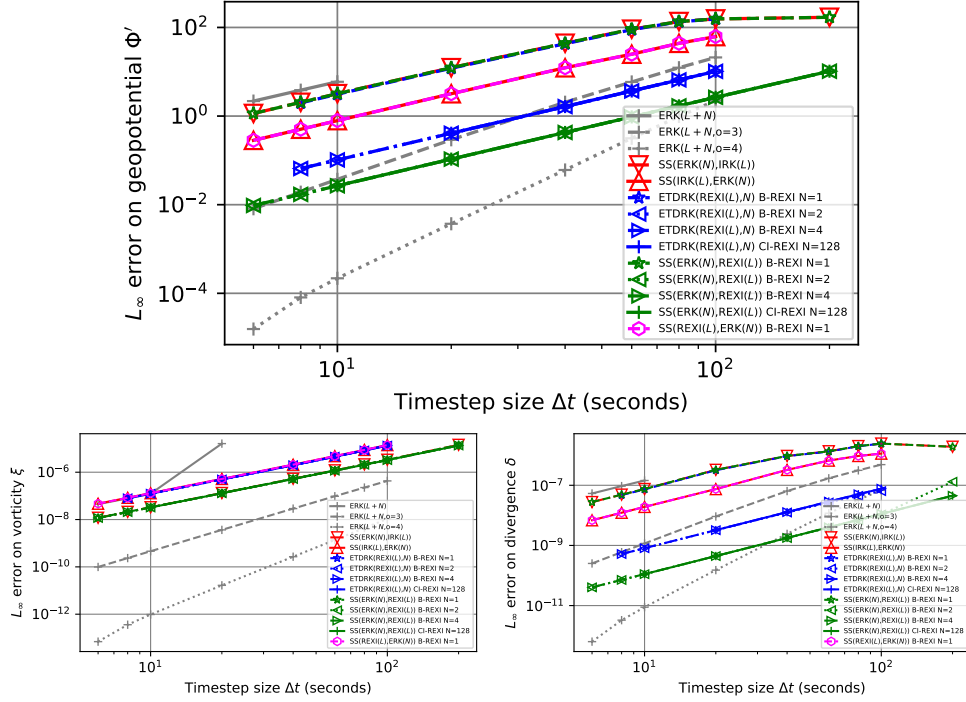


FIG. 6.6. Studies using non-REXI methods (using direct exponentiation) on all prognostic variables with error vs. time step size for the barotropic instability benchmark. (ETDRK data points are missing due to 1h timeouts of the job.)

639 to their computationally efficient way. In particular, the classical 4th-order accurate
 640 ERK method provides excellent results across all prognostic variables.

641 A closer look at the geopotential Φ' errors shows that the B-REXI-based methods
 642 with $N=2$ poles are to be preferred compared to all other methods. Again, the ETDRK
 643 method shows no real benefits.

644 Investigating the vorticity η leads to a different interpretation: Now, the implicit
 645 Strang-split method provides the best results which can be easily explained by the
 646 situation that the exponential treatment of the L_c term did not lead to any beneficial
 647 results already in the error vs. time step size plots and additional computational time
 648 is required here. Finally, ETDRK are literally the worst in here, not paying off at all.

649 The errors on the divergence δ show similar results compared to the geopotential,
 650 which is why we skip a detailed discussion here.

651 **6.2.8. Summary of PDE results.** The CI-REXI method with $N = 128$ poles
 652 is not beneficial at all compared to B-REXI with $N = 2$ poles. Using only $N = 2$ poles
 653 with the B-REXI method already provides the best results, and no improvement can
 654 be gained by adding more poles. This is actually quite surprising, with expectations of
 655 exponential integration methods to always provide significantly better results. How-
 656 ever, using such a higher-order approximation seems to provide sufficient accuracy so
 657 that the errors from the splitting approach dominate the overall errors.

658 We would like to emphasize that all the statements are specific to the SWE on
 659 the rotating sphere PDE and should not be generalized.

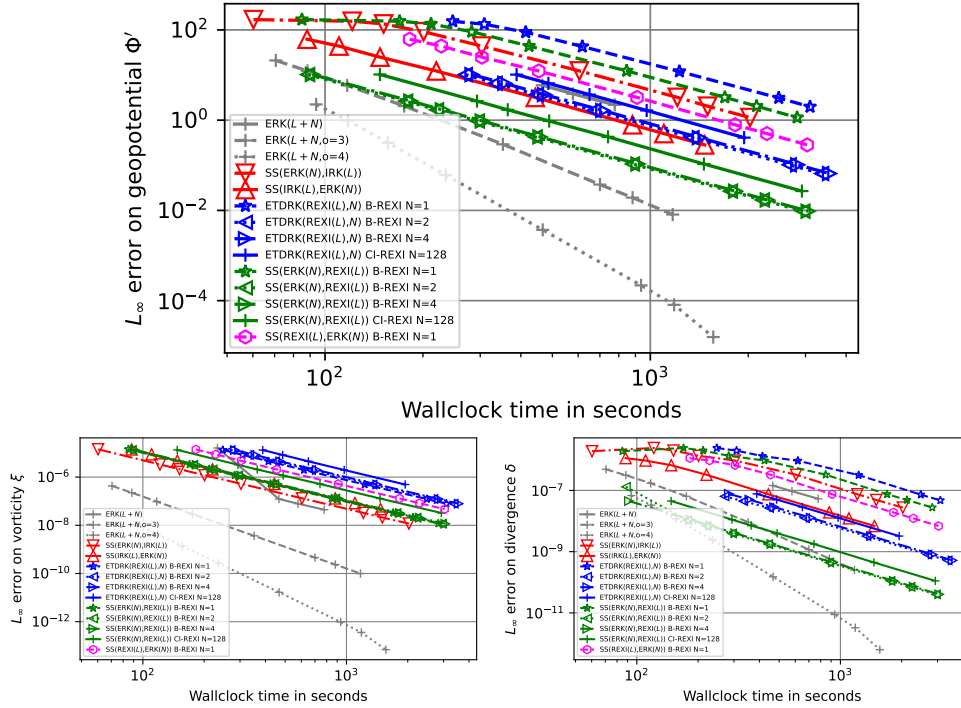


FIG. 6.7. Studies using non-REXI methods (using direct exponentiation) on all prognostic variables with error vs. time step size for the barotropic instability benchmark. (ETDRK data points are missing due to 1h timeouts of the job.)

660 **7. Summary and Conclusions.** Exponential integration methods are consid-
 661 ered to be a way to integrate with high efficiency. As part of that, φ functions need
 662 to be solved, which turn out to be computationally rather challenging.

663 This paper investigated different ways to approximate φ functions with rational
 664 approximations of exponential integration (REXI). The coefficients of REXI meth-
 665 ods can be derived in many ways and we introduced a generalized REXI approach,
 666 finally allowing to express many different methods in this way. We showed this
 667 for the Butcher/Bickard-based REXI, Cauchy Contour integration based REXI and
 668 T(erry)-REXI method. All methods have been introduced in a way making its ca-
 669 pabilities and limitations easily graspable. With respect to physical properties, the
 670 T-REXI method requires special treatment for (quasi-)stationary modes and became
 671 obsolete with CI-REXI. In addition, we derived an elegant way to compute higher-
 672 order φ functions based on REXI coefficients for lower-order φ .

673 An in-depth investigation of the approximation quality of each REXI method has
 674 been conducted including an explanation of numerical issues for all of the methods.
 675 Next, we put it into the context of time integration methods. We first used linear
 676 ODEs where we studied and discussed properties of stability, convergence and also
 677 the filtering capabilities. Second, we performed in-depth studies using the nonlinear
 678 shallow-water equations on the rotating sphere. Surprisingly, the best REXI method
 679 turned out B-REXI with only $N = 2$ poles, leading to a significant reduction of
 680 computational effort compared to former REXI methods in this context using $N = 128$
 681 poles. Consequently, regarding demands on computational resources, B-REXI showed
 682 a reduction of a factor of 64 compared to previous work. This also means that a

683 higher-order implicit Runge-Kutta method is competitive to traditional exponential
684 integration methods for this PDE.

685 **Acknowledgements.** Both authors like to thank Pedro S. Peixoto for pointing
686 out the potential relation of exponential integration methods to Laplace transforms
687 and Peter Lynch’s work in this context. Martin Schreiber is grateful to NCAR for
688 providing financial support and a very inspiring office space with a splendid view to the
689 flatirons, which strongly supported this work. Both authors thank Matthew Normile
690 for preliminary work as well as Finn Capelle and Raphael Schilling who indirectly
691 contributed to this work with the REXInsight software.

692 The authors gratefully acknowledge the Gauss Centre for SC e.V. (www.gauss-
693 centre.eu) for funding this project by providing computing time on the GCS Super-
694 computer SUPERMUC-NG at Leibniz Supercomputing Centre (www.lrz.de).

695

REFERENCES

- 696 [1] T. A. BICKART, *An Efficient Solution Process for Implicit Runge-Kutta Methods*, SIAM Journal
697 on Numerical Analysis, 14 (1977), pp. 1022–1027, <https://doi.org/10.1137/0714069>.
- 698 [2] J. C. BUTCHER, *Implicit Runge-Kutta Processes*, AMS, 18 (1964), pp. 50–64.
- 699 [3] J. C. BUTCHER, *On the implementation of implicit Runge-Kutta methods*, BIT, 16 (1976),
700 pp. 237–240, <https://doi.org/10.1007/BF01932265>.
- 701 [4] T. BUVOLI, *A Class of Exponential Integrators Based on Spectral Deferred Correction*, (2015),
702 pp. 1–22, <http://arxiv.org/abs/1504.05543>.
- 703 [5] T. BUVOLI, *A class of exponential integrators based on spectral deferred correction*, SIAM Jour-
704 nal on Scientific Computing, 42 (2020), pp. A1–A27, <https://doi.org/10.1137/19M1256166>.
- 705 [6] K. E. A. CALVIN, *IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Work-*
706 *ing Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on*
707 *Climate Change. IPCC, Geneva, Switzerland.*, tech. report, Intergovernmental Panel on
708 Climate Change (IPCC), July 2023, <https://doi.org/10.59327/IPCC/AR6-9789291691647>,
709 <https://www.ipcc.ch/report/ar6/syr/> (accessed 2023-11-09). Edition: First.
- 710 [7] C. CLANCY AND P. LYNCH, *Laplace transform integration of the shallow-water equations. Part*
711 *I: Eulerian formulation and Kelvin waves*, Quarterly Journal of the Royal Met. Society,
712 137 (2011), pp. 792–799, <https://doi.org/10.1002/qj.793>.
- 713 [8] C. CLANCY AND J. A. PUDYKIEWICZ, *On the use of exponential time integration methods in*
714 *atmospheric models*, Tellus, Series A: Dynamic Meteorology and Oceanography, 65 (2013),
715 <https://doi.org/10.3402/tellusa.v65i0.20898>.
- 716 [9] R. COURANT, H. LEWY, AND K. FRIEDRICHS, *Über die partiellen Differenzgleichungen der*
717 *mathematischen Physik*, Mathematische Annalen, (1932).
- 718 [10] S. M. COX AND P. C. MATTHEWS, *Exponential time differencing for stiff systems*, Journal of
719 Computational Physics, 176 (2002), pp. 430–455, <https://doi.org/10.1006/jcph.2002.6995>.
- 720 [11] ECMWF, *The Strength of a Common Goal: A Roadmap To 2025*, (2016), [https://www.ecmwf](https://www.ecmwf.int/sites/default/files/ECMWF_Roadmap_to_2025.pdf)
721 [int/sites/default/files/ECMWF_Roadmap_to_2025.pdf](https://www.ecmwf.int/sites/default/files/ECMWF_Roadmap_to_2025.pdf).
- 722 [12] O. G. ERNST AND M. J. GANDER, *Why it is Difficult to Solve Helmholtz Problems with Classical*
723 *Iterative Methods*, vol. 83, 2012, <https://doi.org/10.1007/978-3-642-22061-6>.
- 724 [13] M. FRIGO, *A Fast Fourier Transform Compiler*, 1999.
- 725 [14] J. GALEWSKY, R. K. SCOTT, AND L. M. POLVANI, *An initial-value problem for testing numerical*
726 *models of the global shallow-water equations*, Tellus, Series A: Dynamic Meteorology and
727 Oceanography, 56 (2004), pp. 429–440, <https://doi.org/10.1111/j.1600-0870.2004.00071.x>.
- 728 [15] J. HACK AND R. JAKOB, *Description of a global shallow water model based on the spectral*
729 *transform method*, vol. NCAR/TN-34, 1992. Publication Title: NCAR Technical Note.
- 730 [16] E. HAIRER, S. NORSETT, AND G. WANNER, *Solving ordinary differential equations I: Nonstiff*
731 *problems*, 1987.
- 732 [17] E. HARNEY AND P. LYNCH, *Laplace transform integration of a baroclinic model*, Quarterly
733 Journal of the Royal Met. Society, (2019), pp. 347–355, <https://doi.org/10.1002/qj.3435>.
- 734 [18] T. S. HAUT, T. BABE, P. G. MARTINSSON, AND B. A. WINGATE, *A high-order time-parallel*
735 *scheme for solving wave propagation problems via the direct construction of an approximate*
736 *time-evolution operator*, IMA Journal of Numerical Analysis, 36 (2016), pp. 688–716, <https://doi.org/10.1093/imanum/drv021>.
- 737 [19] M. HOCHBRUCK AND A. OSTERMANN, *Exponential integrators*, vol. 19, 2010, <https://doi.org/>
738

- 739 [10.1017/S0962492910000048](https://doi.org/10.1017/S0962492910000048). ISSN: 09624929 Publication Title: Acta Numerica.
- 740 [20] K. R. JACKSON AND S. P. NØRSETT, *The potential for parallelism in RK methods. Part 1: RK*
741 *formulas in standard form*, SIAM Journal on Numerical Analysis, 32 (1995), pp. 49–82.
- 742 [21] W. KUTTA, *Beitrag zur näherungsweise integration totaler Differentialgleichungen*, Z. Math.
743 Phys., (1901), pp. 435–453.
- 744 [22] J. L. LIONS, Y. MADAY, AND G. TURINICI, *Résolution d'EDP par un schéma en temps*
745 *"pararéel"*, Comptes Rendus de l'Académie des Sciences - Series I: Mathematics, 332 (2001),
746 pp. 661–668, [https://doi.org/10.1016/S0764-4442\(00\)01793-6](https://doi.org/10.1016/S0764-4442(00)01793-6). ISBN: 0764-4442.
- 747 [23] P. LYNCH, *Filtering integration schemes based on the Laplace and Z transforms*, 1991. ISSN:
748 00270644 Issue: 3 Pages: 653–666 Publication Title: Monthly Weather Review Volume:
749 119.
- 750 [24] M. MINION, R. SPECK, M. BOLTEN, M. EMMETT, AND D. RUPRECHT, *Interweaving PFASST*
751 *and Parallel Multigrid*, (2014), pp. 1–20, <https://doi.org/10.1137/14097536X>.
- 752 [25] C. MOLER AND C. VAN LOAN, *Nineteen Dubious Ways to Compute the Exponential of a*
753 *Matrix, Twenty-Five Years Later*, SIAM Review, 45 (2003), pp. 3–49, [https://doi.org/10.](https://doi.org/10.1137/S00361445024180)
754 [1137/S00361445024180](https://doi.org/10.1137/S00361445024180).
- 755 [26] J. NIESEN AND W. M. WRIGHT, *A Krylov subspace algorithm for evaluating the phi-functions*
756 *appearing in exponential integrators*, ACM Transactions on Mathematical Software, 38
757 (2012), p. 20, <https://doi.org/10.1145/2168773.2168781>.
- 758 [27] S. A. ORSZAG, *On the Elimination of Aliasing in Finite-Difference Schemes by Filtering High-*
759 *Wavenumber Components*, (1971), p. 1. 2/3 rule.
- 760 [28] A. C. ROJAS MENDOZA AND P. D. S. PEIXOTO, *Numerical Solution of Ordinary Differential*
761 *equations using Laplace transform integration*, master's thesis, Univ. de São Paulo, 2020.
- 762 [29] C. RUNGE, *Über die numerische Auflösung von Differentialgleichungen*, Mathematische An-
763 nalen, 46 (1895).
- 764 [30] N. SCHAEFFER, *Efficient spherical harmonic transforms aimed at pseudospectral numerical*
765 *simulations*, Geochemistry, Geophysics, Geosystems, 14 (2013), pp. 751–758.
- 766 [31] T. SCHMELZER AND L. N. TREFETHEN, *Evaluating matrix functions for exponential integrators*
767 *via carathéodory-fejér approximation and contour integrals*, Electronic Transactions on
768 Numerical Analysis, 29 (2007), pp. 1–18, <https://doi.org/10.1007/s00586-009-1106-6>.
- 769 [32] M. SCHREIBER AND R. LOFT, *A parallel time integrator for solving the linearized shallow water*
770 *equations on the rotating sphere*, Numerical Linear Algebra with Applications, 26 (2018).
- 771 [33] M. SCHREIBER, P. S. PEIXOTO, T. HAUT, AND B. WINGATE, *Beyond spatial scalability limita-*
772 *tions with a massively parallel method for linear oscillatory problems*, International Journal
773 of High Performance Computing Applications, 32 (2017), pp. 913–933.
- 774 [34] M. SCHREIBER, N. SCHAEFFER, AND R. LOFT, *Exponential integrators with parallel-in-time rat.*
775 *approx. for the shallow-water equations on the rotating sphere*, Parallel Computing, 85
776 (2019), pp. 56–65, <https://doi.org/10.1016/j.parco.2019.01.005>. Publisher: Elsevier B.V.
- 777 [35] D. SHEEN, I. H. SLOAN, AND V. THOMÉE, *A parallel method for time-discretization of para-*
778 *bolic problems based on contour integral representation and quadrature*, Mathematics of
779 Computation, 69 (1999), pp. 177–196.
- 780 [36] STRANG, GILBERT, *On the Construction and Comparison of Difference Schemes*, SIAM Journal
781 on Numerical Analysis, 5 (1968), pp. 506–517, <https://doi.org/10.1137/0705041>.
- 782 [37] C. TEMPERTON, *Treatment of the Coriolis Terms in Semi-Lagrangian Spectral Models*, (1995).
- 783 [38] J. THUBURN, T. RINGLER, W. SKAMAROCK, AND J. KLEMP, *Numerical representation of geo-*
784 *strophic modes on arbitrarily structured C-grids*, Journal of Computational Physics, 228
785 (2009), pp. 8321–8335, <https://doi.org/10.1016/j.jcp.2009.08.006>.
- 786 [39] M. TOKMAN, *Efficient integration of large stiff systems of ODEs with exponential propagation*
787 *iterative (EPI) methods*, Journal of Computational Physics, 213 (2006), pp. 748–776.
- 788 [40] M. TOKMAN, *A new class of exponential propagation iterative methods of Runge-Kutta type*
789 *(EPIRK)*, Journal of Computational Physics, 230 (2011), pp. 8762–8778, [https://doi.org/](https://doi.org/10.1016/j.jcp.2011.08.023)
790 [10.1016/j.jcp.2011.08.023](https://doi.org/10.1016/j.jcp.2011.08.023). Publisher: Elsevier Inc.
- 791 [41] L. N. TREFETHEN AND J. A. C. WEIDEMAN, *The Exponentially Convergent Trapezoidal Rule*,
792 56 (2014), pp. 385–458.
- 793 [42] L. N. TREFETHEN, J. A. C. WEIDEMAN, AND T. SCHMELZER, *Talbot quadratures and rational*
794 *approximations*, BIT Numerical Mathematics, 46 (2006), pp. 653–670.
- 795 [43] J. VIRIEUX, A. ASNAASHARI, R. BROSSIER, L. MÉTIVIER, A. RIBODETTI, AND W. ZHOU, *An*
796 *introduction to full waveform inversion*, 2014, [https://doi.org/10.1190/1.9781560803027.](https://doi.org/10.1190/1.9781560803027.entry6)
797 [entry6](https://doi.org/10.1190/1.9781560803027.entry6). Publication Title: Encyclopedia of Exploration Geophysics.
- 798 [44] H. WANG, J. P. BOYD, AND R. A. AKMAEV, *On computation of Hough functions*, Geoscientific
799 Model Development, 9 (2016), pp. 1477–1488, <https://doi.org/10.5194/gmd-9-1477-2016>.
800 ISBN: 1471239314.