



Parallel-in-time integration of the shallow water equations on the rotating sphere using Parareal and MGRIT

João Guilherme Caldas Steinstraesser, Pedro da Silva Peixoto, Martin Schreiber

► To cite this version:

João Guilherme Caldas Steinstraesser, Pedro da Silva Peixoto, Martin Schreiber. Parallel-in-time integration of the shallow water equations on the rotating sphere using Parareal and MGRIT. 2023. hal-04165113

HAL Id: hal-04165113

<https://hal.science/hal-04165113>

Preprint submitted on 27 Oct 2023

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

Parallel-in-time integration of the shallow water equations on the rotating sphere using Parareal and MGRIT

João Guilherme Caldas Steinstraesser^{a,*}, Pedro da Silva Peixoto^a, Martin Schreiber^{b,c,d}

^a*Universidade de São Paulo, Rua do Matão, 1010, São Paulo, 05508-090, Brazil*

^b*Université Grenoble Alpes, 621 avenue Centrale, Saint-Martin-d'Hères, 38400, France*

^c*Inria AIRSEA team, 700 Avenue Centrale, Grenoble, 38058, France*

^d*Technical University of Munich, Boltzmannstrasse 3, Garching b. Muenchen, 85748, Germany*

Abstract

Despite the growing interest in parallel-in-time methods as an approach to accelerate numerical simulations in atmospheric modeling, improving their stability and convergence remains a substantial challenge for their application to operational models. In this work, we study the temporal parallelization of the shallow water equations on the rotating sphere combined with time-stepping schemes commonly used in atmospheric modeling due to their stability properties, namely an Eulerian implicit-explicit (IMEX) method and a semi-Lagrangian semi-implicit method (SL-SI-SETTLS). The main goal is to investigate the performance of parallel-in-time methods, namely Parareal and Multigrid Reduction in Time (MGRIT) when these well-established schemes are used on the coarse discretization levels and provide insights on how they can be improved for better performance. We begin by performing an analytical stability study of Parareal and MGRIT applied to a linearized ordinary differential equation depending on some temporal parallelization parameters, including the choice of a coarse scheme. Next, we perform numerical simulations of two standard tests in atmospheric modeling to evaluate the stability, convergence, and speedup provided by the parallel-in-time methods compared to a fine reference solution computed serially. We also conduct a detailed investigation on the influence of artificial viscosity and hyperviscosity approaches, applied on the coarse discretization levels, on the performance of the temporal parallelization. Both the analytical stability study and the numerical simulations indicate a poorer stability behavior when SL-SI-SETTLS is used on the coarse levels, compared to the IMEX scheme. With the IMEX scheme, a better trade-off between convergence, stability, and speedup compared to serial simulations can be obtained under proper parameters and artificial viscosity choices, opening the perspective of the potential competitiveness for realistic models.

Keywords: Parallel-in-time methods, shallow water equations on the sphere, Parareal method, MGRIT method, atmospheric modeling

1. Introduction

The numerical simulation of atmospheric circulation models, in the context of climate modeling and numerical weather prediction, is a challenge that motivates constant research efforts. Since the middle of the last century, several spatial and temporal discretization schemes have been proposed, seeking a balance between accuracy, numerical stability, and computational costs for simulations in large domains both in space and time [67]. Over the decades, advances in terms of computing technology, with increasing processing and memory resources available, as well as the advent of massively parallel high-performance computing (HPC) systems, opened new possibilities and motivated the development of new approaches in atmospheric modeling [66].

*Corresponding author

Among the several temporal discretization approaches, parallel-in-time (PinT) methods have recently raised an increasing research interest. The term “PinT” refers to a large variety of numerical methods [22], of which the most popular, *e.g.*, Parareal [42], Multigrid Reduction in Time (MGRIT) [19] and Parallel Full Approximation Scheme in Space and Time (PFASST) [17], are iterative algorithms that, by using a fine, computationally expensive and one or more coarser, less expensive discretizations of the problem, allow to compute several time steps simultaneously, thus replacing the classical approach of serial time-stepping.

The interest in PinT in various application domains has grown mainly in the past two decades, with these methods being seen as an alternative for overcoming the saturation of already well-established spatial parallelism approaches and taking more advantage of HPC systems. This increasing popularity is also explained by the non-intrusive character of several PinT methods, allowing the user to combine them with arbitrary temporal and spatial discretizations, and the availability of open-source libraries implementing them, *e.g.*, Xbraid (MGRIT) [1] and LibPFASST (PFASST) [44], which makes the temporal parallelization of an operational code a relatively easy task. However, PinT is still making its first steps in atmospheric circulation models. Indeed, examples of efficient application of PinT methods, in the sense of effectively allowing to reduce the time-to-solution when compared to serial time-stepping, mainly include problems of parabolic and diffusive nature, *e.g.*, [65, 27, 3, 29]. Also, temporal parallelization suffers from stability and convergence issues when applied to simple hyperbolic problems, *e.g.*, the linear advection equation. This issue has been identified and studied by several works, *e.g.*, [5, 24, 14, 62], and has in some way discouraged further application and investigation of PinT methods to more complex advection-dominated problems, such as those arising in atmospheric modeling. Notably, it is known that the lack of convergence and stability arise on high wavenumbers of the solution due to the mismatch of phase representations on the discretization levels [51], and several approaches trying to overcome it have been proposed, *e.g.*, [52, 12, 34], but most of them have limited application or introduce algorithmic complications that reduce the possible gains in terms of time-to-solution and diminish the non-intrusive character of the PinT methods, making them less attractive to complex applications such as atmospheric circulation problems.

However, even if most studies of PinT methods for hyperbolic problems still focus on relatively simple one-dimensional models in order to develop a better understanding, some recent works already propose studies towards the temporal parallelization of more complete atmospheric models. These works focus on the shallow water equations (SWE) on the rotating sphere, a two-dimensional model commonly used as a starting point in atmospheric modeling since it contains most of the numerical and implementation challenges related to the horizontal discretization found in more complete, tridimensional models [69]. [2] studies an asymptotic MGRIT method based on the asymptotic Parareal proposed and applied to the one-dimensional rotating SWE by [34]; [32] implemented the PFASST method combined with a spherical harmonics discretization of the rotating SWE; and [55] developed a parallel-in-time method based on a temporal discretization using rational approximation for exponential integrators. It should also be noted that research efforts have been made to apply PinT methods to solve more complex and operational models. For instance, in the context of ocean circulation and sea-ice model, Parareal has been applied to the Finite-volume Sea ice-Ocean circulation Model (FESOM2) by [49], presenting slow convergence and stability issues that prevented theoretical speedups, with better convergence results, in the same context, being obtained by [48] with a variant of Parareal using spatial coarsening, but with still present instabilities in long-term simulations and lack of real wall-clock speedup due to parallel overheads. In the context of magnetically confined plasma simulation in tokamaks, in the Integrated Plasma Simulator (IPS) [8] and the nuclear fusion research project ITER [53], an event-based modification of Parareal allows for well-measured speedup, parallel efficiency, and parallel scaling results.

In this work, we also study the application of PinT methods to the SWE on the rotating sphere. Here, we focus on two- and multilevel temporal parallelization using Parareal and MGRIT. We combine them with popular temporal and spatial discretization schemes for atmospheric circulation, used operationally and/or for research purposes in weather and climate forecast models, to evaluate if well-established numerical methods in the atmospheric modeling community would suit temporal parallelization. Namely, we consider a spectral discretization in space using spherical harmonics; in time, we consider two schemes, an Eulerian implicit-explicit (IMEX) one and the semi-Lagrangian semi-implicit SL-SI-SETTLS proposed by [35]. We highlight that there is a recent interest in semi-Lagrangian methods in the PinT framework, *e.g.*, with studies

on the Burgers [54] and advection [61, 60] equations indicating stability and convergence improvements of Parareal and/or MGRIT by using semi-Lagrangian coarse discretizations; thus, we seek to study if this behavior is also verified in the context of a more complex problem.

This investigation is conducted following two approaches. First, we perform an analytical stability study of Parareal and MGRIT applied to a linearized ODE as a function of some parameters for the temporal parallelization, including the choice of a coarse time-stepping scheme (IMEX or SL-SI-SETTLS). This analysis is based on the work developed in the Parareal framework by [58], and we extend it to the two-level MGRIT with arbitrary relaxation. Second, we perform numerical simulations to evaluate the PinT methods in terms of stability, convergence to a reference solution, and computational time compared to the reference simulation. We consider two test cases commonly used for studying the numerical simulation of the SWE on the rotating sphere. Moreover, we investigate the use of artificial viscosity and hyperviscosity on the coarse discretization levels and how they influence the stability and convergence of the temporal parallelization.

The proposed study is conducted having in sight possible practical applications in atmospheric modeling. Indeed, the high complexity of operational models makes the non-intrusive character of PinT methods such as Parareal and MGRIT a very attractive feature. By combining these methods with time-stepping schemes, which are effectively used in these models, and by considering discretization and parametric choices which are coherent with practical applications, we may provide indications of the feasibility of their temporal parallelization and how the PinT performance can be improved by properly parametrizing it.

This paper is organized as follows: in Section 2, we present the shallow water equations on the rotating sphere and their discretization using IMEX and SL-SI-SETTLS in time and spherical harmonics in space; in Section 3 we describe the Parareal and MGRIT methods; the analytical stability study is developed in Section 4; the numerical simulations for evaluating the performance of the methods and the influence of the artificial viscosity and hyperviscosity are presented in Section 5; finally, conclusions are presented in Section 6.

2. The shallow water equations on the rotating sphere

2.1. Governing equations

The SWE on the rotating sphere read

$$\begin{aligned} \frac{\partial}{\partial t} \mathbf{U} &= \mathbf{L}_G(\mathbf{U}) + \mathbf{L}_C(\mathbf{U}) + \mathbf{N}_A(\mathbf{U}) + \mathbf{N}_R(\mathbf{U}) \\ &= \mathbf{L}\mathbf{U} + \mathbf{N}(\mathbf{U}) \end{aligned} \quad (1)$$

where $\mathbf{U} = (\Phi, \xi, \delta)^T$, $\Phi = \bar{\Phi} + \Phi' = gh$ is the geopotential field (with g the gravitational acceleration, h the fluid depth, and $\bar{\Phi}$ and Φ' the mean geopotential and the geopotential perturbation, respectively), $\xi := \mathbf{z} \cdot (\nabla \times \mathbf{V})$ is the vorticity (\mathbf{z} is the unit vector in the vertical direction), $\delta := \nabla \cdot \mathbf{V}$ is the divergence, $\mathbf{V} := (u, v)^T$ is the horizontal velocity, $\mathbf{L} := \mathbf{L}_G + \mathbf{L}_C$, $\mathbf{N} := \mathbf{N}_A + \mathbf{N}_R$ and \mathbf{L}_G , \mathbf{L}_C , \mathbf{N}_A and \mathbf{N}_R are respectively the linear gravity, linear Coriolis, nonlinear advection and nonlinear rest terms, given respectively by

$$\begin{aligned} \mathbf{L}_G(\mathbf{U}) &= \begin{pmatrix} -\bar{\Phi}\delta \\ 0 \\ -\nabla^2 \Phi \end{pmatrix}, \quad \mathbf{L}_C(\mathbf{U}) = \begin{pmatrix} 0 \\ -\nabla \cdot (f\mathbf{V}) \\ \mathbf{z} \cdot \nabla \times (f\mathbf{V}) \end{pmatrix} \\ \mathbf{N}_A(\mathbf{U}) &= \begin{pmatrix} -\mathbf{V} \cdot \nabla \Phi \\ -\nabla \cdot (\xi \mathbf{V}) \\ -\nabla^2 \left(\frac{\mathbf{V} \cdot \mathbf{V}}{2} \right) + \mathbf{z} \cdot \nabla \times (\xi \mathbf{V}) \end{pmatrix}, \quad \mathbf{N}_R(\mathbf{U}) = \begin{pmatrix} -\Phi'\delta \\ 0 \\ 0 \end{pmatrix} \end{aligned}$$

Moreover, in the numerical simulations performed in this work, we consider an artificial (hyper)viscosity approach [40]. We then include the linear term

$$\mathbf{L}_\nu(\mathbf{U}) = (-1)^{\frac{q}{2}+1}\nu \begin{pmatrix} \nabla^q \Phi' \\ \nabla^q \xi \\ \nabla^q \delta \end{pmatrix}$$

where $\nu \geq 0$ is the viscosity coefficient and $q \geq 2$ is the viscosity order, with q even.

2.2. Spatial discretization

In this work, the SWE equations on the rotating sphere are discretized in space using a spectral approach based on spherical harmonics, which is briefly presented below. We refer the reader to [15] for details. We remark that spherical harmonics discretization is used in important operational applications in atmospheric modeling, *e.g.*, by the Integrated Forecasting System (IFS) at the European Centre for Medium-Range Weather Forecast (ECMWF) [16], the Global Spectral Model (GSM) at the National Centers for Environmental Prediction (NCEP) of the U.S. National Oceanic and Atmospheric Administration (NOAA) [45] and the Global Spectral Model (GSM) of the Japan Meteorological Agency (JMA) [37].

A given time-dependent smooth field $\psi(\lambda, \mu, t)$ defined on the sphere, where $\mu := \sin(\theta)$, λ is the longitude, θ is the latitude, and t is the time, can be written as a spherical harmonics expansion

$$\psi(\lambda, \mu, t) = \sum_{m=-\infty}^{\infty} \sum_{n=|m|}^{\infty} \psi_{m,n}(t) Y_{m,n}(\lambda, \mu) \quad (2)$$

where $Y_{m,n}$ is the spherical harmonic function of zonal and total wavenumbers m and n , respectively, and $\psi_{m,n}$ is the respective expansion coefficient. The spherical harmonics are defined as the product of associated Legendre functions and Fourier modes, so the direct and inverse spherical transforms can be performed via Fast Fourier and Fast Legendre transforms, respectively, in the zonal and meridional directions. In the implementation considered here, we consider a triangular truncation for the expansion (2):

$$\psi(\lambda, \mu, t) = \sum_{m=-M}^M \sum_{n=|m|}^M \psi_{m,n}(t) Y_{m,n}(\lambda, \mu)$$

Two main reasons motivate using spectral discretization for the SWE. First, spherical harmonics are eigenfunctions of the spherical Laplacian:

$$\nabla^2 Y_{m,n} = \frac{-n(n+1)}{a^2} Y_{m,n}$$

where a is the sphere's radius. This property is interesting since the spherical Laplacian arises on the temporal discretization of the SWE. We remark that a pseudospectral approach is adopted, relying on collocated (A-grid) gridpoint variables, with the nonlinear terms as well as the Coriolis term being computed in the physical space (the discretization size of the physical grid being determined by the well-known anti-aliasing “3/2-rule” [15]). Second, the spherical harmonics expansions and their properties related to the spherical Laplacian allow to easily formulate semi-implicit methods, thus avoiding the so-called “pole problem” found in physical grid discretizations used *e.g.*, in explicit finite difference methods: when uniform latitude-longitude meshes are used, cells near the poles have small longitudinal length, requiring the use of too small time step size to fulfill CFL stability constraints [30].

2.3. Temporal discretization

In this work, we consider two time-stepping schemes for discretizing (1), namely an Eulerian, Strang-splitting implicit-explicit (IMEX) scheme and the semi-Lagrangian semi-implicit SL-SI-SETTLS proposed by [35], which are briefly presented below. In both schemes, using a backward Euler method, the (hyper)viscosity term \mathbf{L}_ν is solved at the end of each time step.

2.3.1. IMEX

In the implicit-explicit scheme considered here, the stiff, linear terms of the governing equations are treated implicitly, and the nonlinear ones are treated explicitly, allowing to overcome stability constraints imposed by the former, which makes it a popular time-stepping approach in fluid dynamics problems [11]. Although not as popular in operational numerical weather prediction models as semi-Lagrangian (SL) methods, IMEX schemes are also a relevant class of methods in atmospheric modeling, being used in research models such as the Model for Prediction Across Scales (MPAS) at the U.S. National Center for Atmospheric Research (NCAR) [57] and the Nonhydrostatic Unified Model of the Atmosphere (NUMA) at the Naval Research Laboratory (NPS/NRL) [38], including implicit-explicit approaches for the horizontal terms of the governing equations, as considered here, and also explicit and implicit treatments respectively for the horizontal and vertical terms [43].

We consider a second-order Strang-splitting IMEX scheme, with a half timestep of the implicit solver, followed by a full explicit timestep and a second implicit half timestep:

$$\mathbf{U}^{n+1} = \mathbf{F}_I^{\Delta t/2} \left(\mathbf{F}_E^{\Delta t} \left(\mathbf{F}_I^{\Delta t/2} (\mathbf{U}^n) \right) \right) \quad (3)$$

The implicit term $\mathbf{F}_I^{\Delta t/2}$ consists of a Crank-Nicolson discretization:

$$\mathbf{U}^* - \frac{\Delta t}{4} \mathbf{L}(\mathbf{U}^*) = \mathbf{U}^n + \frac{\Delta t}{4} \mathbf{L}(\mathbf{U}^n) \quad (4)$$

which is solved for \mathbf{U}^* using the “semi-implicit treatment of the Coriolis term” described by [64]. The explicit term $\mathbf{F}_E^{\Delta t}$ consists of a second-order Runge-Kutta scheme:

$$\begin{aligned} \bar{\mathbf{U}} &= \Delta t \mathbf{N}(\mathbf{U}^*) \\ \overline{\bar{\mathbf{U}}} &= \Delta t \mathbf{N}(\mathbf{U}^* + \bar{\mathbf{U}}) \\ \mathbf{U}^{**} &= \frac{1}{2} (\bar{\mathbf{U}} + \overline{\bar{\mathbf{U}}}) \end{aligned} \quad (5)$$

2.3.2. SL-SI-SETTLS

Semi-Lagrangian semi-implicit schemes are popular numerical methods in atmospheric circulation modeling, being used by several operational numerical weather prediction models, *e.g.*, the IFS-ECMWF, the GSM-JMA and the Global Forecast System (GFS) at NCEP/NOAA, to cite only a few [43]. They use a semi-implicit (Crank-Nicolson) discretization of the linear terms and a semi-Lagrangian (SL) treatment of the nonlinear advection term. The principle of SL schemes is to combine the Eulerian and Lagrangian approaches for spatiotemporal PDEs, the former relying on a fixed spatial grid (*e.g.*, the IMEX method presented above), being conceptually simpler but usually restricted by Courant-Friedrichs-Lewy (CFL) stability conditions limiting the time step size. In contrast, the latter follows the trajectories of the fluid particles along time instead of using a fixed grid, which allows larger time steps but is much more complex in implementation. SL methods avoid these issues by using a Lagrangian approach at each time step $[t_n, t_{n+1}]$: it traces the trajectories of the particles arriving at each point of a fixed spatial grid at time t_{n+1} , performing a spatial interpolation at time t_n to retrieve the departure values. It simplifies SL compared to purely Lagrangian schemes but still allows the use of larger time steps than Eulerian methods [59].

SL methods require to perform an estimation of the Lagrangian trajectories $(t, \mathbf{x}(t))$ along each time step $[t_n, t_{n+1}]$, with several approaches being proposed in the literature. We consider here the Stable Extrapolation Two-Time-Level Scheme (SETTLS) proposed by [35], in which the trajectories are computed iteratively based on an estimation of the velocity field on the intermediate timestep:

$$\mathbf{x}_d^{k+1} = \mathbf{x}_a - \frac{\Delta t}{2} \mathbf{V} \left(t_n + \frac{\Delta t}{2} \right) \approx \mathbf{x}_a - \frac{\Delta t}{2} ([2\mathbf{V}(t_n, \mathbf{x}_d) - \mathbf{V}(t_{n-1}, \mathbf{x}_d) + \mathbf{V}(t_n, \mathbf{x}_a)]) \quad (6)$$

where $\Delta t := t_{n+1} - t_n$, \mathbf{x}_a denotes the arrival point of the trajectories $(t, \mathbf{x}(t))$ along $[t_n, t_{n+1}]$, corresponding to the fixed spatial grid, and \mathbf{x}_d denotes the departure point. The SL-SI-SETTLS thus reads

$$\frac{\mathbf{U}^{n+1} - \mathbf{U}_*^n}{\Delta t} = \frac{1}{2} (\mathbf{L}\mathbf{U}^{n+1} + \mathbf{L}\mathbf{U}_*^n) + \frac{1}{2} ([2\mathbf{N}_R(\mathbf{U}^n) - \mathbf{N}_R(\mathbf{U}^{n-1})]_* + \mathbf{N}_R(\mathbf{U}^n)) \quad (7)$$

where the subscript $*$ denotes interpolation to the departure point \mathbf{x}_d . In (6) and (7), the terms in brackets are linear extrapolations to t_{n+1} , which are then averaged with the known values at t_n to obtain an estimation at $t_{n+1/2} := t_n + \Delta t/2$.

We notice that this two-step time-stepping scheme, with the solution at time t_{n+1} depending on the two preview times t_n and t_{n-1} , poses additional challenges in terms of practical implementation of the parallel-in-time methods. Therefore, we use a modified version of the method in the PinT framework as discussed in Section 3.2.4.

3. Parallel-in-time methods

Following the classification proposed by [7], PinT methods comprise iterative schemes allowing to compute several time steps simultaneously, algorithms based on spatial domain decomposition methods, and direct methods parallelizing the scheme used for advancing each time step. Detailed reviews on PinT can be found in [22, 46].

In this work, we focus on two of the most popular PinT methods, Parareal and MGRIT. Both can be interpreted as iterative, predictor-corrector algorithms, based on the simultaneous use of coarse (low expensive) and fine (expensive) time-stepping methods, the former being computed sequentially along the entire temporal domain, whereas the latter is computed in parallel, *i.e.*, with several time steps being computed simultaneously. Parareal is a two-level scheme, using fine and coarse discretizations of the problem; on the other hand, MGRIT is a multilevel scheme, using more than two discretization levels, besides other generalizations w.r.t. Parareal, as explained below. The popularity of these methods can be explained by their non-intrusive character, allowing the user to implement their own time-stepping scheme as coarse and fine methods, and by their simple implementation and formulation, mainly in the case of Parareal. In this section, we briefly describe these methods and introduce some notation used in this paper. For this description, we consider the time-dependent system of ODEs

$$\frac{\partial \mathbf{U}}{\partial t}(t) = \mathbf{f}(t, \mathbf{U}(t)), \quad \mathbf{U}(0) = \mathbf{U}_0, \quad t \in [0, T] \quad (8)$$

which can be obtained *e.g.*, via a spatial discretization of a PDE.

3.1. Parareal

First developed by [42] and presented as a predictor-corrector algorithm by [4, 6], the Parareal method iteratively computes approximations to the solution of (8) using simultaneously two numerical schemes, ϕ_f and ϕ_c , named respectively *fine* and *coarse* propagators. The method aims to provide these approximations with a smaller computational cost compared to the serial simulation of the accurate (thus expensive) fine propagator. Let $T_0 = 0, T_1, \dots, T_N = T$ be a discretization of the temporal domain $[0, T]$. We denote by $\phi_f(\mathbf{U}, T_n, T_{n+1})$ and $\phi_c(\mathbf{U}, T_n, T_{n+1})$ the propagation of \mathbf{U} from T_n to T_{n+1} , using respectively ϕ_f and ϕ_c . The initial guess (solution at iteration $k = 0$) is provided by the serial simulation of the coarse propagator along the entire temporal domain:

$$\mathbf{U}_{n+1}^0 = \phi_c(\mathbf{U}_n^0, T_n, T_{n+1}), \quad n = 0, \dots, N-1$$

In the following iterations, Parareal computes

$$\mathbf{U}_{n+1}^{k+1} = \phi_c(\mathbf{U}_n^{k+1}, T_n, T_{n+1}) + \phi_f(\mathbf{U}_n^k, T_n, T_{n+1}) - \phi_c(\mathbf{U}_n^k, T_n, T_{n+1}), \quad n = 0, \dots, N-1, \quad k \geq 0 \quad (9)$$

where \mathbf{U}_n^k is an approximation to the solution of (8) at time T_n and iteration k . Note that the only term on the right-hand side of (9) that needs to be computed serially at each iteration is the first one, using the coarse

propagator (which is supposed to be relatively cheap). The remaining terms, including the expensive, fine one, can be computed in parallel (*i.e.*, the propagations along each time slice $[T_n, T_{n+1}]$ can be distributed to different parallel processors) since they depend only on the solution of the previous iteration k , which has already been computed for every $n = 0, \dots, N$.

In general, the coarse propagator is defined with a larger timestep compared to the fine one. This, however, is not necessary. The fine and coarse schemes can be defined by different spatial resolutions, numerical schemes, or integration orders, for example. Finally, it can be easily shown that the Parareal solution converges exactly to the fine solution (*i.e.*, the solution obtained via a serial simulation of the fine propagator ϕ_f) in at most N iterations [21]; however, a much faster convergence is required in practice.

3.2. MGRIT

The MGRIT algorithm is a multilevel, predictor-corrector iterative parallel-in-time method introduced by [19] and based on spatial multigrid methods [50]. In the following paragraphs, we describe the main ideas behind MGRIT and its parameters that are relevant to the study proposed in this work; we refer the reader to [18, 1] for details that are omitted here.

3.2.1. Basic definitions

Let N_{levels} be the number of temporal discretization levels, indexed from 0 to $N_{\text{levels}} - 1$, the former being the finest one and the latter the coarsest one. The time-stepping scheme ϕ_l in each level is defined by a timestep Δt_l . As in Parareal, in general (but not necessarily), the coarse levels are defined such as to have a coarser temporal discretization w.r.t. the fine levels. We then define a coarsening factor $m_c \geq 1$ such that $\Delta t_{l+1} = m_c \Delta t_l, l = 0, \dots, N_{\text{levels}} - 2$. For the sake of simplicity, we consider the same coarsening factor m_c for all pairs of consecutive levels and also that the temporal discretization in each level is homogeneous.

In the following, we consider a pair of consecutive levels $(l, l+1)$, called respectively “fine” and “coarse” levels, to introduce some notation and vocabulary of MGRIT, which can be easily extended for any pair of levels.

The temporal domain is divided into $N_{l+1} := T/\Delta t_{l+1}$ coarse time steps, defining the coarse time instances $T_i = i\Delta t_{l+1}, i = 0, \dots, N_{l+1}$; and $N_l := T/\Delta t_l = m_c N_{l+1}$ fine time steps, defining the time instances $t_j = j\Delta t_l, j = 0, \dots, N_l$. Note that $T_i = t_{m_c i}, i = 0, \dots, N_{l+1}$. The times $T_i, i = 0, \dots, N_{l+1}$ of the coarse temporal discretization are named *C-points* and the fine time instances not present in the coarse discretization, *i.e.*, $t_j = 1, \dots, N_l - 1, j \neq m_c i, i = 0, \dots, N_{l+1}$, are named *F-points*. These definitions are illustrated in Figure 1.




Figure 1: Basic definitions for the temporal discretization adopted in MGRIT. Bullets and squares represent, respectively, F-points and C-points. The sequence of solid arrows represents an F-relaxation. The dashed arrow represents a C-relaxation.

The numerical approximations for the solution of (8) are defined on the fine temporal grid, such that $U_j^{(l)}$ and $U_{m_c i}^{(l)}$ denote approximations at the F-points and C-points, respectively, with $i = 0, \dots, N_{l+1}$ and $j = 0, \dots, N_l, j \neq m_c i$.

We call *F-relaxation* (or *F-relaxation sweep*) the update of the solution at all F-points between two C-points (*i.e.*, $t_j \in]T_i, T_{i+1}[$), by using ϕ_l with the C-point value $U_{m_c i}^{(l)}$ as an initial solution. Analogously, we call *C-relaxation* the update of each C-point value $U_{m_c i}^{(l)}$ by advancing one fine time step using ϕ_l with initial

solution $\mathbf{U}_{m_c i-1}^{(l)}$, see Figure 1. An update of the approximations on the fine level can be performed using a combination of F- and C-relaxations; in general, MGRIT is set to use *FCF-relaxation*, *i.e.*, an F-relaxation followed by a C- and a second F- one, but one can define any $\text{F(CF)}^{N_{\text{relax}}}$ -relaxation, $N_{\text{relax}} \geq 0$, including the simplest F-relaxation (which is used *e.g.*, by Parareal, corresponding to $N_{\text{relax}} = 0$). In general, we can expect a faster MGRIT convergence with a more complex relaxation scheme, but with a higher numerical cost.

3.2.2. Full approximation scheme

The idea of the MGRIT algorithm is to compute error corrections on the coarse grid that improve the solution obtained by relaxation on the fine grid. It is performed via the Full Approximation Scheme (FAS) [10], which can be briefly described in five steps:

1. Perform an $\text{F(CF)}^{N_{\text{relax}}}$ -relaxation on the fine temporal grid;
2. Restrict the solution and its residual from the fine to the coarse grid at the C-points;
3. Solve the so-called coarse grid equation;
4. Use the solution of the coarse grid equation to compute a coarse grid error approximation;
5. Correct the solution on the fine grid by injecting the computed error approximation from the coarse to the fine grid and performing a further F-relaxation sweep.

3.2.3. Multilevel scheme

In the two-level algorithm presented above, the third step is the only one that needs to be computed sequentially along the entire temporal domain. However, in the case of multilevel runs, this step is performed recursively, thus adding additional degrees of parallelism. In this case, the entire algorithm described above is applied for the pair of levels $(l+1, l+2)$ and so on. The order of recursive calls in each level is another parameter to be set in an MGRIT simulation, influencing the convergence speed and computational cost of the method. In this work, we consider F-cycles followed by one post-V-cycle for the order of recursive calls through levels, as defined in [1].

Finally, as in Parareal, an exact convergence towards the serial solution of the finest temporal discretization is obtained in a finite number of iterations. In a MGRIT simulation with N_0 time points on the finest level, coarsening factor m_c and using an $\text{F(CF)}^{N_{\text{relax}}}$ -relaxation strategy, exact convergence is obtained in at most $N_0/((N_{\text{relax}} + 1)m_c)$ iterations [18].

3.2.4. Implementation

The numerical results presented in Section 5 were obtained using the XBraid library [1], which implements MGRIT combined with user-defined time-stepping methods. The Parareal simulations were also performed using XBraid, by setting the parameters $N_{\text{levels}} = 2$ and $N_{\text{relax}} = 0$.

As a major feature, XBraid allows us to easily incorporate MGRIT into an operational code. Notably, every parallelization is made in a “black-box” fashion inside the library itself, with the user only needing to define *what* is communicated between processors, but not *when* or *how*. This, however, may not be convenient in the cases of two-step time-stepping schemes (such as SL-SI-SETTLS, described in Section 2.3.2), since two consecutive time steps $[t_{n-1}, t_n]$ and $[t_n, t_{n+1}]$, both contributing to the solution at t_{n+1} , may be treated by different parallel processors. In [54], a modified version of Parareal is proposed in order to deal with it. In XBraid, this issue can be solved on the coarsest level by including the solution at t_{n-1} in the solution to be communicated, since the simulation is sequential at this level. In all other levels, however, time steps are computed in parallel, with no predefined order, making it impracticable to define this extra communication in a coherent way. An alternative approach would be to replace the two-step scheme with a one-step one, with the solution at t_{n-1} replaced by the one at t_n . Preliminary studies using the Gaussian bumps test case introduced in Section 5 were conducted comparing the MGRIT solution under two scenarios: (i) with the alternative one-step scheme applied in all time steps of all levels, except for the coarsest one, in which the original two-step scheme is applied; (ii) with the one-step scheme being applied in all time steps of all levels, including the coarsest one. These results show that the chosen procedure has little influence

on the convergence and stability of MGRIT. Therefore, and to avoid extra communications on the coarsest level, the following procedure is considered for adapting SL-SI-SETTLS to MGRIT:

- On the coarsest level, the original two-step SETTLS scheme is used in all time steps, except for the *temporal boundaries* between processors, *i.e.*, if $[t_{n-1}, t_n]$ and $[t_n, t_{n+1}]$ are treated by different processors, in which case the one-step alternative is adopted;
- On all other levels using SL-SI-SETTLS, the one-step alternative is used in all time steps.

We remark that, except for the simulations performed for evaluating wall-clock times, all the simulations presented in Section 5 were run using a single processor in time, which implies that the original SL-SI-SETTLS scheme is applied on the entire coarsest level. Notably, in the two-level configurations using SL-SI-SETTLS on the coarse level and another scheme on the fine one, the PinT simulation is entirely conducted without modification of SL-SI-SETTLS. Finally, we notice that the analytical stability study presented in the next section, which indicates stability issues of Parareal and MGRIT using SL-SI-SETTLS on the coarse level (which are confirmed in the numerical simulations presented in Section 5), are conducted considering the original SL-SI-SETTLS.

4. Stability study

We perform in this section a stability study of Parareal and MGRIT depending on given parameters and given fine and coarse time-stepping methods. This study follows the procedure presented in [13]. We consider the ODE

$$\dot{u} = \lambda_{\mathbf{L}} u + \mathbf{N}(u) \quad (10)$$

which is obtained *e.g.*, by solving (1) with a spectral method, in which case $\lambda_{\mathbf{L}} \in \mathbb{C}$ is a wavenumber mode of \mathbf{L} and u is a spectral coefficient of \mathbf{U} . Problem (10) is linearized by considering $\mathbf{N}(u) = \lambda_{\mathbf{N}} u$, with $\lambda_{\mathbf{N}} \in \mathbb{C}$. Thus,

$$\dot{u} = \lambda_{\mathbf{L}} u + \lambda_{\mathbf{N}} u \quad (11)$$

As pointed out by [13], a full stability study of a numerical scheme to (11) should be performed in a four-dimensional space, *i.e.*, as a function of the real and imaginary parts of $\xi_{\mathbf{L}} := \lambda_{\mathbf{L}} \Delta t$ and $\xi_{\mathbf{N}} := \lambda_{\mathbf{N}} \Delta t$. Since the SWE equations on the rotating sphere are characterized by the propagation of purely imaginary wavenumber modes, we perform this study in the function of $Re(\xi_{\mathbf{N}})$ and $Im(\xi_{\mathbf{N}})$ for fixed values of $\xi_{\mathbf{L}}$. Namely, we are interested in the cases $\xi_{\mathbf{L}} \approx 0$ (Rossby modes) and positive and negative purely imaginary values of $\xi_{\mathbf{L}}$ (inertia-gravity wave modes).

Let us first determine some physically meaningful values of $\xi_{\mathbf{L}}$ for this stability study, which can be easily obtained via a simplification of the linear operator \mathbf{L} defined in (1). By considering the Coriolis parameter f to be constant (which corresponds to the so-called f -plane approximation), we can write

$$\mathbf{L}(\mathbf{U}) = \mathbf{L}\mathbf{U} = \begin{pmatrix} 0 & 0 & -\bar{\Phi} \\ 0 & 0 & -f \\ -\nabla^2 & f & 0 \end{pmatrix} \begin{pmatrix} \Phi \\ \xi \\ \delta \end{pmatrix}$$

which, in the spherical harmonics spectral space, reads

$$\mathcal{L}\mathbf{U}_{m,n} = \begin{pmatrix} 0 & 0 & -\bar{\Phi} \\ 0 & 0 & -f \\ -\frac{n(n+1)}{a^2} & f & 0 \end{pmatrix} \begin{pmatrix} \Phi_{m,n} \\ \xi_{m,n} \\ \delta_{m,n} \end{pmatrix} \quad (12)$$

The three eigenvalues of \mathcal{L} are

$$\lambda_{\mathbf{L}} = 0, \quad \lambda_{\mathbf{L}} = \pm i \sqrt{f^2 + \frac{n(n+1)}{a^2} \bar{\Phi}}$$

In the numerical simulations presented in Section 5, the mean geopotential field has an order of magnitude of $10^5 \text{m}^2 \text{s}^{-2}$. We then consider here $\bar{\Phi} = 10^5 \text{m}^2 \text{s}^{-2}$. The radius of the Earth is $a \approx 6371.22 \times 10^3 \text{m}$ and the Coriolis parameter in a f -plane approximation reads $f = 2\Omega \approx 2 \times 7.292 \times 10^{-5} \text{s}^{-1}$, where Ω is the Earth's angular velocity. Since $\bar{\Phi}/a^2 \gg f^2$, we neglect this term in the approximation of $\lambda_{\mathbf{L}}$ and we plot the stability regions for fixed integer multiples of $\tilde{\xi}_{\mathbf{L}} := i\sqrt{\bar{\Phi}/a^2} \approx 2.5 \times 10^{-4}i$, where the scaling of this quantity can be interpreted as a consequence of the choice of wavenumber n and/or the choice of timestep Δt .

For a given numerical scheme for (10), we compute the amplification factor $A := |u^{n+1}|/|u^n|$ and plot its stability region (*i.e.*, where $|A| \leq 1$) on the $Re(\xi_{\mathbf{N}}) - Im(\xi_{\mathbf{N}})$ plane. We are mainly interested in the overlapping between the stability region and the imaginary axis.

We begin by presenting the stability plots for IMEX and SL-SI-SETTLS individually (*i.e.*, in a serial framework), where we introduce some specific details of the stability analysis of each method. Then, we present the stability study of Parareal presented by [58] and we extend it to the MGRIT framework in a two-level configuration with an arbitrary relaxation strategy.

4.1. Stability of the time-stepping schemes individually

4.1.1. Stability of IMEX

Applying (3) to the linearized problem (10), the implicit step (4) with a half timestep $\Delta t/2$ has the amplification factor

$$A_{\text{IMEX-I}} = \frac{4 + \xi_{\mathbf{L}}}{4 - \xi_{\mathbf{L}}}$$

and, for the explicit step with a full timestep,

$$A_{\text{IMEX-E}} = \frac{\xi_{\mathbf{N}} + \xi_{\mathbf{N}}(1 + \xi_{\mathbf{N}})}{2} = \frac{\xi_{\mathbf{N}}(2 + \xi_{\mathbf{N}})}{2}$$

Thus, the amplification factor of IMEX reads

$$A_{\text{IMEX}} = A_{\text{IMEX-I}} A_{\text{IMEX-E}} A_{\text{IMEX-I}} = \left(\frac{4 + \xi_{\mathbf{L}}}{4 - \xi_{\mathbf{L}}} \right)^2 \frac{\xi_{\mathbf{N}}(2 + \xi_{\mathbf{N}})}{2}$$

4.1.2. Stability of SL-SI-SETTLS

In the stability study of semi-Lagrangian schemes, the amplification factors also depend on the spatial wavenumber (see *e.g.*, [15]). Indeed, the amplification factor A is more rigorously defined by considering a solution under the form $u(x, t_n) = A^n e^{i\kappa x}$, where κ is the wavenumber. In the case of Eulerian schemes (*e.g.*, IMEX), all terms are evaluated in the same spatial point and the exponential term vanishes, leading to $A = |u^{n+1}|/|u^n|$, which is not the case in the semi-Lagrangian framework. We define $s := x_j - x_*$ as the distance between the grid point x_j and the departure point x_* (at t_n) of the trajectory arriving at x_j at t_{n+1} . Replacing in (7), we obtain

$$A^{n+1} e^{i\kappa x_j} - A^n e^{i\kappa x_*} = \frac{\xi_{\mathbf{L}}}{2} (A^{n+1} e^{i\kappa x_j} + A^n e^{i\kappa x_*}) + \frac{\xi_{\mathbf{N}}}{2} ([2A^n - A^{n-1}] e^{i\kappa x_*} + A^n e^{i\kappa x_j})$$

which can be simplified and rearranged to

$$A_{\text{SL-SI-SETTLS}}^2 \left(1 - \frac{\xi_{\mathbf{L}}}{2} \right) - A_{\text{SL-SI-SETTLS}} \left(e^{-i\kappa s} \left(1 + \frac{\xi_{\mathbf{L}}}{2} + \xi_{\mathbf{N}} \right) + \frac{\xi_{\mathbf{N}}}{2} \right) + \frac{\xi_{\mathbf{N}}}{2} e^{-i\kappa s} = 0 \quad (13)$$

Eq. (13) has two roots $A_{\text{SL-SI-SETTLS}}^+$ and $A_{\text{SL-SI-SETTLS}}^-$, and the stability region of SL-SI-SETTLS is the intersection of their stability regions. Moreover, the stability region depends on the spatial wavenumber κs . As shown by [35], SL-SI-SETTLS has improved stability properties since the intersection of the stability regions for all κs is not empty, which is not the case of previous SL-schemes used in IFS-EMCWF. In this work, we compute the stability region of SL-SI-SETTLS by intersecting the stability regions of $\kappa s \in [0, 2\pi]$ with a step $\pi/10$.

Figure 2 compares the stability regions of both schemes for $\xi_L \in \{0, 5 \times 10^3 \tilde{\xi}_L, 10^4 \tilde{\xi}_L, 2.5 \times 10^4 \tilde{\xi}_L\}$. For negative ξ_L , the plots are symmetric w.r.t. the $Re(\xi_N)$ -axis. We note that the stability regions of IMEX do not vary with $\xi_L \in i\mathbb{R}$ since, in this case, $|(4 + \xi_L)/(4 - \xi_L)| = 1$, indicating that $A_{\text{IMEX}} = \xi_N(2 + \xi_N)/2$. On the other hand, the stability region of SL-SI-SETTLS depends on ξ_L , and for larger values of the linear mode, the method outperforms IMEX, notably with a larger intersection between the stability region and the imaginary axis.




Figure 2: Stability regions of IMEX and SL-SI-SETTLS applied to (10) as a function of ξ_N for fixed values of ξ_L . The contours for IMEX visually coincide.

4.2. Stability of Parareal

We briefly present the stability analysis developed by [58] for the parareal iteration

$$u_n^k = \phi_c(u_{n-1}^k) + \phi_f(u_{n-1}^{k-1}) - \phi_c(u_{n-1}^{k-1}) \quad (14)$$

using the fine and coarse timestepping methods ϕ_c and ϕ_f with respective timesteps Δt_f and $\Delta t_c = m_c \Delta t_f$. The parareal solution u_n^k at iteration k is computed at time $t_n := n\Delta T$, with $\Delta t/\Delta t_f =: N_f$ and $\Delta t/\Delta t_c =: N_c$.

Let $A_f := A_f(\Delta t_f)$ and $A_c := A_c(\Delta t_c)$ be the stability functions of ϕ_f and ϕ_c , respectively. Then, (14) can be written as

$$u_n^k = A_c^{N_c} u_{n-1}^k + A_f^{N_f} u_{n-1}^{k-1} - A_c^{N_c} u_{n-1}^{k-1} = A_c^{N_c} u_{n-1}^k + (A_f^{N_f} - A_c^{N_c}) u_{n-1}^{k-1} \quad (15)$$

Eq. (15) is a recurrence relation leading to

$$u_n^k = \left(\sum_{i=0}^k \binom{n}{i} (A_f^{N_f} - A_c^{N_c})^i (A_c^{N_c})^{n-i} \right) u_0 \quad (16)$$

where u_0 is the initial solution at $t = 0$. The coefficients in (16) are the binomial coefficients of Pascal's triangle, which can be easily identified in Figure 3a, adapted from [58].

From (16), we define the Parareal stability function




Figure 3: Illustration of the recurrence relations of: (a) the Parareal iteration $u_n^k = Ru_{n-1}^k + Su_{n-1}^{k-1}$, with $R := A_c^{N_c}$ and $S := A_f^{N_f} - A_c^{N_c}$; and (b) the two-level MGRIT iteration with $N_{\text{relax}} = 1$, $u_n^k = Ru_{n-1}^k + Su_{n-2}^{k-1}$, with $R := A_c^{N_c}$ and $S := A_f^{2N_f} - A_c^{2N_c}$.

$$A_{\text{parareal}}(n, k, N_f, N_c, \phi_f, \phi_c) := \sum_{i=0}^k \binom{n}{i} (A_f^{N_f} - A_c^{N_c})^i (A_c^{N_c})^{n-i}, \quad n \geq k \quad (17)$$

Let us study the behavior of A_{parareal} along iterations and as a function of the coarse time-stepping method. We fix $n = 100$, $m_c = 2$, $N_c = 1$, and IMEX as a fine scheme and we plot the stability contours at iterations $k = 0, 1, 5, 10$ in Figure 4. In all cases (for both choices of coarse time-stepping method and for all chosen values of ξ_L), we observe a decrease of the stability region along iterations. Also, for both coarse schemes, the stability region is not larger than the ones using the schemes alone (which corresponds to $k = 0$). A much more notable stability loss is observed in the configurations using SL-SI-SETTLS as a coarse method, whose stability contours no longer intersect the $\text{Im}(\xi_N)$ -axis for $k \geq 1$. Also, in the case $\xi_L = 10^4 \xi_{\text{L}}$, in which the serial SL-SI-SETTLS outperforms IMEX, the stability regions vanish (or almost) in the Parareal framework after few iterations. On the other hand, the configurations using IMEX as a coarse method are able to preserve quite well their shape and only suffer a small reduction in size.

4.3. Stability of MGRIT

We generalize the Parareal stability analysis presented above to the MGRIT framework. We consider here the two-level case with $\text{F(CF)}^{N_{\text{relax}}}$ -relaxation and we study the influence of N_{relax} on the stability contours using the fact that the two-level MGRIT can be seen as an overlapping variant of Parareal, with the parallel steps performed at each iteration being computed along $N_{\text{relax}} + 1$ time slices [23]:

$$u_n^k = \phi_c(u_{n-1}^k) + \phi_f^{N_{\text{relax}}+1} \left(u_{n-(N_{\text{relax}}+1)}^{k-1} \right) - \phi_c \phi_f^{N_{\text{relax}}} \left(u_{n-(N_{\text{relax}}+1)}^{k-1} \right) \quad (18)$$

In terms of the fine and coarse stability functions:

$$u_n^k = \left[A_c^{N_c} u_{n-1}^k + \left(A_f^{(N_{\text{relax}}+1)N_f} - A_c^{N_c} A_f^{N_{\text{relax}}N_f} \right) u_{n-(N_{\text{relax}}+1)}^{k-1} \right] \quad (19)$$

which can be rewritten as

$$u_n^k = A_{\text{MGRIT}}(n, k, \phi_f, \phi_c, N_{\text{levels}} = 2, N_{\text{relax}}) u_0$$

with

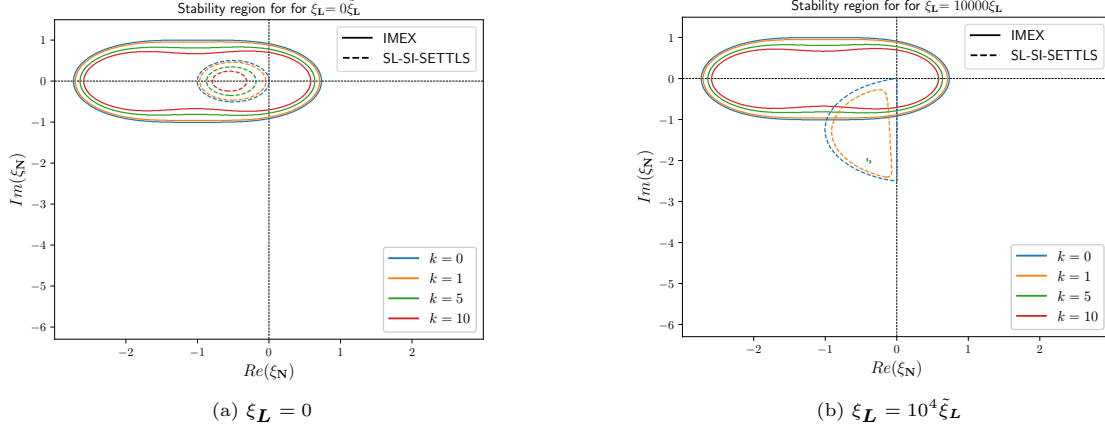


Figure 4: Stability regions of Parareal (with $m_c = 2$, $n = 100$, $k \in \{0, 1, 5, 10\}$) applied to (10) as a function of the coarse timestepping method and ξ_N for fixed values of ξ_L . In all configurations, the fine time-stepping method is IMEX. In the case $\xi_L = 10^4 \tilde{\xi}_L$, the stability region of SL-SI-SETTLS is empty for $k = 10$.

$$A_{\text{MGRIT}}(n, k, N_f, N_c, \phi_f, \phi_c, N_{\text{levels}} = 2, N_{\text{relax}}) := \sum_{i=0}^{\lfloor k/(N_{\text{relax}}+1) \rfloor} \binom{n - iN_{\text{relax}}}{i} \left(A_f^{N_f N_{\text{relax}}} (A_f^{N_f} - A_c^{N_c}) \right)^i (A_c^{N_c})^{n-i(N_{\text{relax}}+1)}, \quad n \geq \frac{k}{N_{\text{relax}} + 1} \quad (20)$$

being the stability function of MGRIT, where $\lfloor \cdot \rfloor$ is the floor function. It is easy to check that (20) reduces to (17) in the case $N_{\text{relax}} = 0$ (Parareal).

The derivation of (20) is analogous to the Parareal case. The two-level MGRIT iteration (19) can be identified with a Pascal's triangle, but with a “jump” of N_{relax} rows in all “diagonal relations”. Figure 3b illustrates it in the case $N_{\text{relax}} = 1$.

In Figure 5, we plot the stability regions for fixed MGRIT parameters ($k = 5$, $n = 100$, $m_c = 2$) and $N_{\text{relax}} \in \{0, 1, 2, 3\}$. We observe that the stability regions increase when N_{relax} increases, mainly from $N_{\text{relax}} = 0$ to $N_{\text{relax}} = 1$. In most cases, the stability regions of the serial schemes are almost or totally recovered with $N_{\text{relax}} \geq 2$; however, for large ξ_L ($\xi_L = 10^4 \tilde{\xi}_L$), it is not enough to improve the stability region of MGRIT using SL-SI-SETTLS as coarse method, remaining smaller than the stability region of the serial scheme.

We highlight that the small stability regions in the configurations using SL-SI-SETTLS may be linked to the fact that the stability region is defined as the intersection between the stability regions of all wavenumbers κ . In the Parareal and MGRIT framework, each of these regions is reduced depending on the parametric choice; thus, their intersection is drastically reduced compared to the stability region of the serial scheme. It does not imply that SL schemes will necessarily suffer from this issue; if the amplification factor can be written in the form $A(\xi_L, \xi_N, e^{-i\kappa s}) = e^{-i\kappa s} \tilde{A}(\xi_L, \xi_N)$, then the region where $|\tilde{A}| \leq 1$ will be the same for all wavenumbers, and its reduction in the PinT framework may be less remarkable. This is the case, for instance, of semi-Lagrangian exponential methods [47], which also present a better stability behavior in PinT simulations, as will be presented in future work.

5. Numerical tests

We consider here two test cases for evaluating the application of Parareal and MGRIT combined with IMEX or SL-SI-SETTLS for the parallel-in-time integration of the SWE on the rotating sphere. For each test, we perform a convergence study, in which we study the convergence speed and stability as a function of




Figure 5: Stability regions of MGRIT (with $k = 5$, $m_c = 2$, $n = 100$ and $N_{\text{relax}} \in \{0, 1, 2, 3\}$) applied to (10) as a function of the coarse timestepping method and $\xi_{\mathbf{N}}$ for fixed values of $\xi_{\mathbf{L}}$. In all configurations, the fine time-stepping method is IMEX. In almost all cases, the contour plots for $N_{\text{relax}} = 2$ and $N_{\text{relax}} = 3$ visually coincide.

selected Parareal and MGRIT parameters (namely the number of levels N_{levels} , the coarsening factor m_c , the relaxation strategy N_{relax} and the spectral resolution M_{coarse} of the coarse levels), and we evaluate the computational times and speedups w.r.t. the parallelization. Moreover, in the first test case, we study the influence of the artificial (hyper)viscosity on the convergence and stability of the temporal parallelization, and we use the obtained conclusions for setting up the viscosity parameters in the second test case.

We are also interested in comparing the performance of the PinT methods using IMEX or SL-SI-SETTLS. Since the finest level is commonly discretized with a small time step, it provides an accurate solution to the problem and is not critical in terms of stability, which is not the case for the coarser levels, in which larger time steps are used. Therefore, we simplify this study by fixing the time-stepping scheme on the finest level ($l = 0$), namely IMEX, and considering the same time-stepping scheme (IMEX or SL-SI-SETTLS) on all coarse levels ($l > 0$).

The choices of parameters and test cases presented in this section represent an effort towards the operational application of PinT methods in atmospheric modeling, even though their application to more complex and realistic problems would evidently require more detailed studies. As done in the analytical stability study presented in Section 4, in which the modes $\xi_{\mathbf{L}}$ were chosen based on realistic physical parameters, we consider here, for instance, meaningful temporal and spatial discretization sizes and viscosity coefficient values. Moreover, the two test cases represent quite complex dynamics, mainly the second one, which is a standard and challenging benchmark in atmospheric modeling research.

We first define the errors and computational time measures used for evaluating the PinT performance:

Error definition. As proposed by [33], we evaluate the PinT errors (w.r.t. a given reference solution) in the spectral space. Indeed, it is known that PinT methods suffer from stability and convergence issues for the higher wavenumber modes of the solution. Therefore, we evaluate the ability of Parareal and MGRIT to converge on different regions of the wavenumber spectrum. Let $\psi_{m,n}$ be the spectral coefficient of a given function ψ with modes $(m, n) \in [-M, \dots, M] \times [m], \dots, M]$ (corresponding to a triangular truncation of a spherical harmonics transform). We define a spectral resolution $0 < R_{\text{norm}} \leq M$ for evaluating the errors. Then, the error of a numerical approximation to ψ , compared to a given reference solution ψ_{ref} , reads

$$E_{\psi, R_{\text{norm}}} := \frac{\|\psi - \psi_{\text{ref}}\|_{\infty, R_{\text{norm}}}}{\|\psi_{\text{ref}}\|_{\infty, R_{\text{norm}}}}, \quad \|\psi\|_{\infty, R_{\text{norm}}} := \max_{\substack{m \in \{0, \dots, R_{\text{norm}}\} \\ n \in \{m, \dots, R_{\text{norm}}\}}} |\psi_{m,n}| \quad (21)$$

In some situations, we also evaluate the more classical L_2 error computed in the physical space:

$$E_{\psi, L_2} := \frac{\|\psi - \psi_{\text{ref}}\|_{L_2}}{\|\psi_{\text{ref}}\|_{L_2}}, \quad \|\psi\|_{L_2} := \sqrt{\frac{1}{M^{(x)}M^{(y)}} \sum_{i=1}^{M^{(x)}} \sum_{j=1}^{M^{(y)}} \psi_{i,j}^2}$$

where $\psi_{i,j}$ are the discrete values of ψ in a homogeneous physical grid with $M^{(x)}$ and $M^{(y)}$ points respectively along latitudes and longitudes.

Evaluation of computational time and speedup. The numerical simulations presented in the following paragraphs were executed in the GRICAD cluster from the University of Grenoble Alpes. The considered nodes are composed of 32 physical cores of Intel Xeon Gold 6130 @2.10GHz distributed between two sockets. The serial-in-time reference simulations were run considering a shared memory spatial parallelization using 16 OpenMP threads bound to cores in the same socket. The parallel-in-time simulations were run with a hybrid OpenMP-MPI parallelization, respectively in space and time, with $N_{\text{proc}}/2$ nodes being allocated to define N_{proc} MPI tasks (which we name hereafter processors), each one using 16 OpenMP threads bound to cores on the same socket. We remark that other space-time parallelization strategies could be adopted, *e.g.*, an MPI-MPI one relying on the splitting of the MPI communicator.

The computing times of the reference and the PinT simulations, the latter using N_{proc} processors, are denoted T_{ref} and $T_{\text{PinT}}(N_{\text{proc}})$. The speedup provided by the temporal parallelization is defined by

$$S(N_{\text{proc}}) := \frac{T_{\text{ref}}}{T_{\text{PinT}}(N_{\text{proc}})}$$

5.1. Gaussian bumps test case

In this first test case, adapted from the single Gaussian bump benchmark presented by [63], a geopotential field with mean value $\bar{\Phi} = g\bar{h}$, with $\bar{h} = 29400\text{m}$ and $g = 9.80616\text{ms}^{-2}$, is perturbed by three Gaussian bump centered at $(\lambda_1, \theta_1) = (\pi/5, \pi/3)$, $(\lambda_2, \theta_2) = (6\pi/5, \pi/5)$ and $(\lambda_3, \theta_3) = (8\pi/5, -\pi/4)$:

$$\Phi(\lambda, \theta) = \bar{\Phi} + A \sum_{i=1}^3 e^{-a_i d_i(\lambda, \theta)^2}$$

where $A = 6000\text{m}$, $a_1 = 20$, $a_2 = 80$, $a_3 = 360$ and

$$d_i(\lambda, \theta) = \arccos(\sin(\theta)\sin(\theta_i) + \cos(\theta)\cos(\theta_i)\cos(\lambda - \lambda_i))$$

The initial velocity field is identical to zero and the final time of simulation is $T = 36\text{h}$.

We begin by defining a discretization to be used on the fine level of the parallel-in-time simulations (and also as a reference solution for evaluating the errors). Since the solution provided by Parareal and MGRIT converges to the one obtained via a serial simulation on the fine level, we want the latter to be a good enough approximation to the exact solution of the problem (or a very refined solution); on the other hand, we do not want the fine level to be overresolved (*e.g.*, with too fine discretizations in time and/or space), since it may lead to unrealistically overestimated speedups provided by the temporal parallelization [28]. We fix the spectral discretization to $M_0 = 256$ and perform a serial simulation with time step sizes in $[2, 960]$, which we compare to a solution obtained with spectral resolution $M = 512$ and time step $\Delta t = 2$. In all cases, no artificial viscosity is used, and the solutions are obtained with IMEX. Figure 6 presents the L_2 errors, computed in the physical space, as well as the spectral errors for given R_{norm} values. Considering all the presented norms, no important error reduction is observed by using time steps smaller than approximately 100s, with the errors being dominated by spatial discretization. We then choose $(M_0 = 256, \Delta t_0 = 60)$ to be used on the fine levels of the temporal parallelization. The solution at $t = T$ obtained under this configuration is compared to the solution obtained with $(M = 512, \Delta t = 2)$ in Figure 7. We notice that for the chosen simulation length T , the turbulence regime is not yet fully developed and the $-5/3$ power law of the two-dimensional kinetic energy spectrum w.r.t. the wavenumber [41] is not yet verified, as illustrated in Figure 11, which compares the kinetic energy spectrum of the fine solution with those obtained with selected PinT configurations.




Figure 6: Gaussian bumps test case: relative L_2 error (computed in the physical space) and spectral errors (for various values of R_{norm}) between a solution obtained with spectral resolution $M = 512$ and time step $\Delta t = 2$ and solutions obtained with $M = M_0 = 256$ and various time steps. IMEX is used in all cases.




Figure 7: Gaussian bumps test case: solution at $t = T = 36\text{h}$ computed with IMEX.


5.1.1. Convergence study

The finest level in the PinT configuration is discretized with the same time step size $\Delta t_0 = \Delta t_{\text{ref}} = 60$, spectral resolution $M_0 = M_{\text{ref}} = 256$ and artificial viscosity coefficient $\nu_0 = \nu_{\text{ref}} = 0$ used in the reference simulation. We perform MGRIT simulations considering $N_{\text{levels}} \in \{2, 3\}$, with a coarsening factor $m_c \in \{2, 4\}$. The same spectral resolution is considered on all coarse levels, being set to $M_{\text{coarse}} \in \{51, 128\}$, which correspond respectively to spatial coarsening factors of 1/5 and 1/2 w.r.t. the finest level. The use of spatial coarsening on the coarse levels is motivated by stability issues observed in the PinT simulations, both when the coarse schemes are IMEX and SL-SI-SETTLS, if too large spectral resolutions are adopted since relatively large time step sizes are used on the coarse levels. We also consider several relaxation strategies, with $N_{\text{relax}} \in \{0, 1, 5\}$. We recall that MGRIT with $(N_{\text{levels}}, N_{\text{relax}}) = (2, 0)$ is equivalent to Parareal. Finally, all coarse levels $l > 0$ use an artificial second-order viscosity coefficient $\nu_{\text{coarse}} = 10^6 \text{m}^2 \text{s}^{-1}$, which is the same value adopted in the numerical simulations presented by [55] for integrating the SWE on the rotating sphere with PFASST. A more detailed study on the influence of the artificial viscosity order and coefficient applied in each discretization level is presented in Section 5.1.2.


Figure 8 presents the evolution, along iterations, of the relative spectral error between the parallel-in-time and fine solutions, evaluated at the final time of simulation when IMEX is used on the coarse levels. The errors are shown for two chosen spectral resolutions, namely $R_{\text{norm}} = 32$ and $R_{\text{norm}} = 128$. We observe that only a few simulations are stable and able to converge (and, when this is the case, with a relatively fast error decrease in the first iterations followed by a slower convergence in the next ones), namely when the spectral resolution on the coarse levels, the number of levels and/or the coarsening factor are sufficiently small; otherwise, the time step on the coarsest level is too large and leads to unstable behaviors. With $M_{\text{coarse}} = 51$, all configurations can converge, more or less rapidly, at the beginning of the wavenumber spectrum ($R_{\text{norm}} = 32$), except for the most aggressive configuration $((N_{\text{levels}}, m_c) = (3, 4))$; when more modes are considered in the error analysis ($R_{\text{norm}} = 128$), the simulations with $(N_{\text{levels}}, m_c) = (2, 4)$ also starts to develop instabilities that compromise convergence, with this unstable behavior being less remarkable when a more expensive relaxation strategy is adopted. On the other hand, when a finer spectral resolution is considered on the coarse levels ($M_{\text{coarse}} = 128$), almost none of the MGRIT configurations are stable, even under $R_{\text{norm}} = 32$, the only exception being the less aggressive one $((N_{\text{levels}}, m_c) = (2, 2))$. The configuration $(N_{\text{levels}}, m_c) = (3, 2)$ initially presents a convergent behavior; still, the error starts to increase after five iterations, which takes place sooner at the end of the wavenumber spectrum ($R_{\text{norm}} = 128$), and when the relaxation is less expensive; under other configurations, the instabilities start to develop even sooner. Finally, in all convergent configurations, we notice a slight improvement when the relaxation strategy is more expensive, but the error behaviors are essentially the same.

The same results are presented in Figure 9 for the simulations using SL-SI-SETTLS on the coarse levels. We notice a more important unstable behavior compared to the simulations using IMEX as a coarse scheme. Among all tested configurations, the only one presenting a relatively stable and convergent behavior is $(N_{\text{levels}}, m_c, M_{\text{coarse}}) = (2, 2, 51)$, *i.e.*, the one with the smallest time step and spectral resolution on the coarsest level (however, even in this case the convergence is less clear and not monotonic as observed in simulations using IMEX). In all other configurations, the error increases after a few iterations, independently of the chosen norm for evaluating the error. Moreover, contrary to the results with IMEX, we do not observe significant differences between the errors computed with $R_{\text{norm}} = 32$ and $R_{\text{norm}} = 128$, meaning that simulations fail to converge in the large spatial scales, and only the former case is presented. Finally, we observe only very small improvements when using more expensive relaxation strategies, but it is not enough to improve the stability of the simulations and lead to convergent behavior.


To illustrate the physical solutions obtained in the PinT simulations, we present in Figure 10 the relative difference between the PinT and the reference geopotential fields at iteration $k = 7$ for chosen configurations, compared to the solution obtained at the initial iteration. With IMEX, we choose $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (3, 2, 0, 51)$ and $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (3, 2, 5, 128)$. The former case presents a monotonic convergence behavior, and we observe small-scale oscillations of the error. On the other hand, the latter case diverges after a few iterations (which takes place sooner in the largest wavenumbers, *i.e.*, for larger R_{norm} values): in the seventh iteration, the PinT already diverged under $R_{\text{norm}} = 128$




(a) $M_{\text{coarse}} = 51, R_{\text{norm}} = 32$



(b) $M_{\text{coarse}} = 51, R_{\text{norm}} = 128$




(c) $M_{\text{coarse}} = 128, R_{\text{norm}} = 32$




(d) $M_{\text{coarse}} = 128, R_{\text{norm}} = 128$

Figure 8: Gaussian bumps test case: relative error $E_{\Phi, R_{\text{norm}}}$ between the PinT and fine solutions at $t = T$ along iterations, for $R_{\text{norm}} = 32$ (left) and $R_{\text{norm}} = 128$ (right); and $M_{\text{coarse}} = 51$ (top) and $M_{\text{coarse}} = 128$ (bottom), with IMEX used on the coarse levels. Simulations are identified by $(N_{\text{levels}}, m_c, N_{\text{relax}})$.



(a) $M_{\text{coarse}} = 51, R_{\text{norm}} = 32$



(b) $M_{\text{coarse}} = 128, R_{\text{norm}} = 32$

Figure 9: Gaussian bumps test case: relative error $E_{\Phi, R_{\text{norm}}}$ between the PinT and fine solutions at $t = T$ along iterations, for $R_{\text{norm}} = 32$ and $M_{\text{coarse}} = 51$ (left) and $M_{\text{coarse}} = 128$ (right), with SL-SI-SETTTLs used on the coarse levels. Results are identical under $R_{\text{norm}} = 128$. Simulations are identified by $(N_{\text{levels}}, m_c, N_{\text{relax}})$.

but not under $R_{\text{norm}} = 32$, and we indeed observe fine oscillations dominating the error plot. With SL-SI-SETTLS as coarse scheme, we present the results for $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (2, 2, 0, 51)$ and $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (2, 2, 0, 128)$. The former presents the best convergence behavior among the simulations depicted in Figure 9, but we observe an increasing magnitude of the fine-scale errors along iterations. The latter configuration starts to diverge after five iterations, under $R_{\text{norm}} = 32$ and $R_{\text{norm}} = 128$: indeed, the plot reveals small- and fine-scale errors at iteration 7.

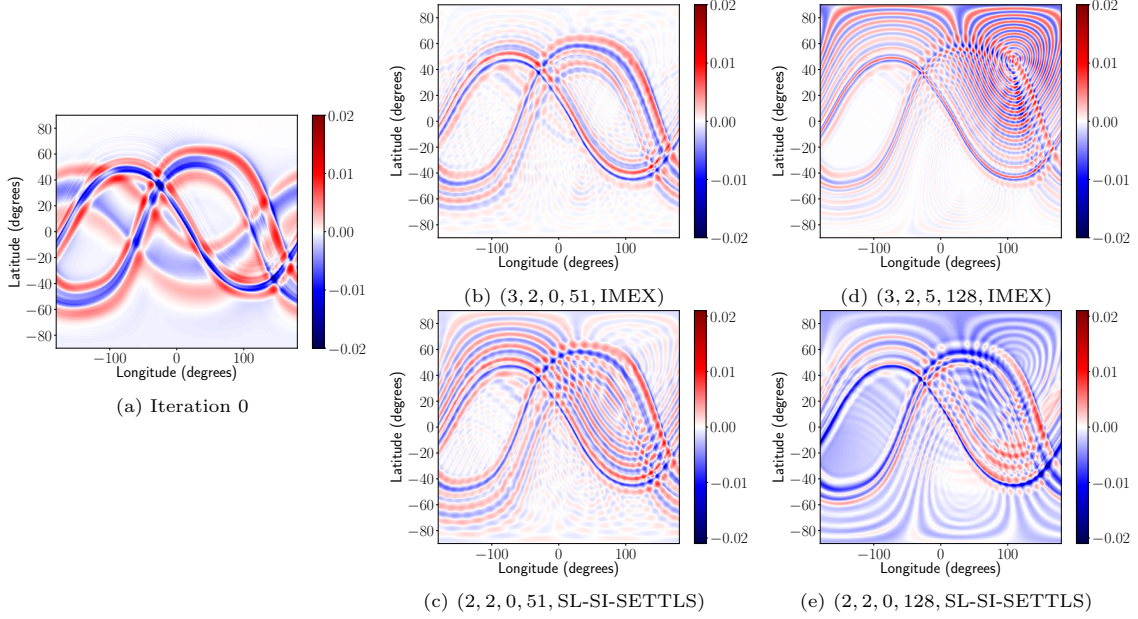


Figure 10: Gaussian bumps test case: the relative difference between the PinT and fine geopotential fields at $t = T$ after 7 iterations for chosen configurations $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}, \text{coarse time-stepping method})$. The left plot represents the solution at the initial iteration, which is visually similar for all configurations.

The stability and convergence behaviors of the simulations presented in Figure 10 are also illustrated in Figure 11, in which we compare the kinetic energy (KE) spectrum at $t = T$ at given iterations with those corresponding to the fine, reference solution. For a triangular spectral truncation, the KE spectrum is defined by [39]

$$E(n, t) = \frac{1}{4} \frac{a^2}{n(n+1)} \sum_{m=-n}^n (|\xi_{m,n}(t)|^2 + |\delta_{m,n}(t)|^2)$$

where $\xi_{m,n}(t)$ and $\delta_{m,n}(t)$ are the spherical harmonics coefficients of the vorticity and divergence, respectively. In Figure 11, we present the spectrum as a function of the wavelength $\mathcal{L}_n := 2\pi a/n$. Concerning the simulations using IMEX, the unstable behavior under configuration $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (3, 2, 5, 128)$ is clear in the spectrum, with over-amplifications of medium to large wavenumbers along iterations; on the other hand, the spectrum of the simulation using $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (3, 2, 0, 51)$ remains below the fine one, but slowly converging to it, indicating a more difficult convergence of large wavenumber modes. In the simulations using SL-SI-SETTLS, we observe some over-amplification in the middle region of the wavenumber spectrum in both PinT configurations.

5.1.2. Influence of artificial (hyper)viscosity

We now study the influence of the artificial (hyper)viscosity coefficients applied on each coarse discretization level on the convergence and stability performance of the PinT methods. Since different temporal and spatial discretization sizes are used on each level, their respective stability constraints are not the same; therefore, different viscosity approaches should probably be applied on different levels. A brief overview of




Figure 11: Gaussian bumps test case: kinetic energy spectra at $t = T$ of the reference (fine) and PinT solutions at iterations 0, 5 and 7 for given configurations. Left and right: IMEX and SL-SI-SETTTLs as coarse time-stepping schemes, respectively.

viscosity and hyperviscosity approaches in spectral methods and some guidelines for choosing the magnitude of the viscosity coefficients are presented in [Appendix A](#). For the study proposed here, we consider the configuration $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (3, 2, 0, 128)$, both with IMEX and SL-SI-SETTTLs as coarse time-stepping schemes. In the former case, this configuration leads to an initial convergent behavior followed by instabilities arising after the fifth iteration (Figure 8); in the latter, the unstable behavior is observed from the initial iteration (Figure 9). As before, no viscosity is used on the fine level ($\nu_0 = 0$); for the two coarse levels, we consider viscosity orders q_1 and q_2 from $\{2, 4, 6\}$, with the viscosity coefficients respectively in $\{10^5, 10^6, 10^7\}$ (m^2s^{-1}), $\{10^{15}, 10^{16}, 10^{17}\}$ (m^4s^{-1}) and $\{10^{25}, 10^{26}, 10^{27}\}$ (m^6s^{-1}). These ranges of coefficients are coherent with those documented in the literature of atmospheric modeling and used in operational models [36]; for instance, reported values of fourth-order coefficient varying between orders of $10^{12}\text{m}^4\text{s}^{-1}$ and $10^{16}\text{m}^4\text{s}^{-1}$, with larger coefficients being adopted, in general, when the spectral resolution is smaller.

It has been observed that the convergence and stability behaviors of the simulations, both in the cases with IMEX and SL-SI-SETTTLs as a coarse scheme, are determined mainly by the viscosity order and coefficient applied on the coarsest level, *i.e.*, q_2 and ν_2 , with only little influence of the values applied on the intermediate level. This observation indicates that it is critical to damp amplifications produced by the coarsest discretization, whose stability constraints are the most critical ones. Therefore, in order to avoid the presentation of a too large number of simulations (with both (q_1, ν_1) and (q_2, ν_2) varying), we fix $\nu_1 = 0$ (*i.e.*, no viscosity on the intermediate level) and present the results for each pair (q_2, ν_2) . Moreover, similar results were obtained in the cases $R_{\text{norm}} = 32$ and $R_{\text{norm}} = 128$, and we only present the former.

In the simulations using IMEX as a coarse scheme (Figure 12a), we observe that large second-, fourth- or six-order viscosities are required to ensure stability: indeed, zero or too small viscosities lead to strong unstable behaviors from the initial iteration, and with intermediate values the simulations initially converge but instabilities are triggered after few iterations. The results also reveal that a compromise between stability and accuracy needs to be fulfilled when choosing the viscosity order: among the stable simulations, a faster convergence is obtained when the viscosity order is higher, indicating that only the largest wavenumbers need to be damped. Notably, with a large second-order viscosity, which damps a large region of the wavenumber spectrum, the error is much larger compared to higher orders already in the initial iteration. On the other hand, in the case where SL-SI-SETTTLs is used as a coarse scheme (Figure 12b), a much more critical dependence of the stability on the viscosity approach is observed. Only the simulation using a large second-order viscosity presents a convergent behavior, indicating that a large region of the wavenumber spectrum needs to be damped. Finally, the results for both coarse schemes reveal that, compared to reported values in serial simulations, larger viscosity coefficients (by between one and three orders of magnitude) are required to ensure stability in the PinT context.

We illustrate in Figure 13 the evolution of the kinetic energy spectrum for two chosen viscosity con-

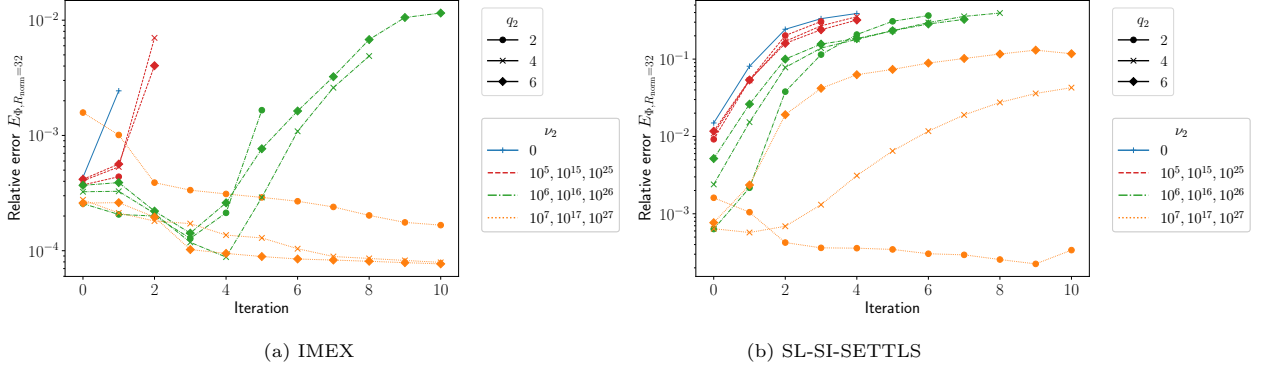


Figure 12: Gaussian bumps test case: relative error $E_{\Phi, R_{\text{norm}}}$ between the PinT and fine solutions at $t = T$ along iterations for $R_{\text{norm}} = 32$, with IMEX (left) or SL-SI-SETTLS (right) being used as a coarse time-stepping scheme, as a function of the viscosity coefficients applied on the coarsest level. All simulations use configuration $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (3, 2, 0, 128)$. No viscosity is applied on the finest and intermediate levels. Viscosity coefficients expressed in $\text{m}^q \text{s}^{-1}$.

figurations, namely $(q_1, \nu_1, q_2, \nu_2) = (-, 0, 2, 10^7)$ and $(q_1, \nu_1, q_2, \nu_2) = (-, 0, 6, 10^{27})$ (*i.e.*, with no viscosity applied on the intermediate level and high second- or sixth-order viscosity on the coarsest none), for both time-stepping schemes on the coarse level. As shown in Figure 12a, both viscosity configurations provide convergence in the simulations using IMEX, mainly the sixth-order one. We indeed observe that the second-order viscosity strongly damps the initial PinT solution almost in the entire wavenumber spectrum, and the sixth-order one provides a better approximation to the reference spectrum. On the other hand, in the simulations using SL-SI-SETTLS on the coarse levels, the large damping due to the second-order viscosity is required for keeping the simulation stable along iterations; a higher-order approach, even with a very large coefficient, damps only the largest wavenumbers, and we observe an overamplification of medium wavenumbers already at iteration $k = 0$, which propagates to the entire spectrum after a few iterations.

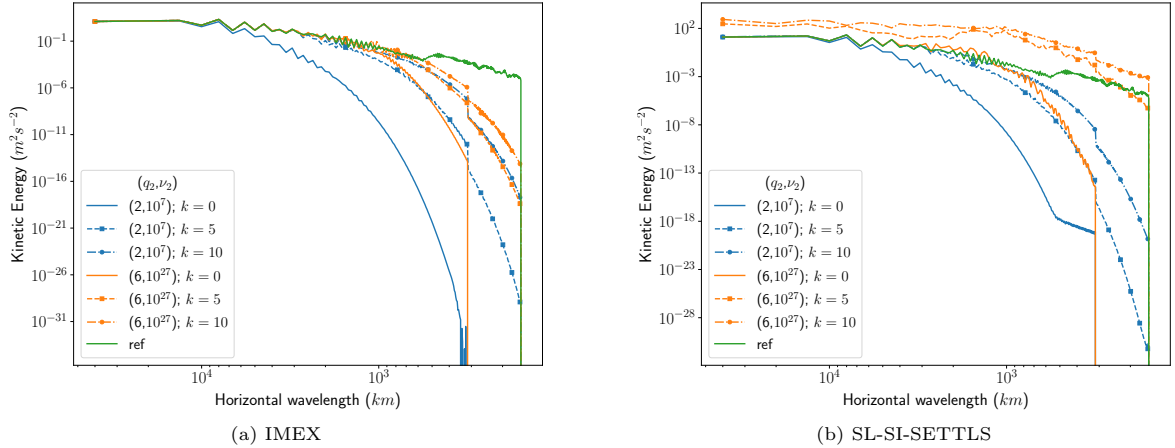


Figure 13: Gaussian bumps test case: kinetic energy spectra at $t = T$ of the reference (fine) and PinT solutions at iterations 0, 5 and 10 for chosen viscosity order and coefficient (q_2, ν_2) applied on the coarsest level. All simulations use configuration $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (3, 2, 0, 128)$, with no viscosity applied on the finest and intermediate levels. Left and right: IMEX and SL-SI-SETTLS as coarse time-stepping schemes, respectively.

5.1.3. Evaluation of computing times and speedups

We now evaluate the computing times and respective speedups of some chosen PinT configurations presenting relatively stable and convergent behavior. A larger number of configurations using IMEX on the coarse levels is chosen since it provides better behavior in general, which allows studying the speedups as a

function of various MGRIT parameters, such as the number of levels, the relaxation strategy, the spectral resolution on the coarse levels and the artificial viscosity approach. In the case of SL-SI-SETTLS, a smaller set of configurations is chosen.

We are mainly interested in the relation between the speedups and the errors provided by each PinT simulation. However, we first present some strong scaling results in Figure 14 for the configuration $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}, q_1, \nu_1) = (2, 2, 0, 51, 2, 10^6)$ using IMEX or SL-SI-SETTLS as coarse time-stepping schemes. The wall times for reaching given iterations are compared to the reference one ($T_{\text{ref}} \approx 137\text{s}$) for simulations using $N_{\text{proc}} \in [1, 64]$. In the initial iteration, we observe a good scaling in the entire range of N_{proc} ; in the following ones, we begin to notice a saturation of the speedup for $N_{\text{proc}} \geq 32$, but there seems to still exist some room for further improvements if more processors are used. In all results presented hereafter, we consider $N_{\text{proc}} = 64$. We notice that the number of processors considered here is smaller than the theoretical maximum value of N_{proc} , which would still provide speedups, which depends on the number of fine time steps, the coarsening factor and the relaxation strategy, as described in Section 3.2; for instance, in the simulation depicted in Figure 14, up to $N_{\text{proc}} = 1080$ could be used. This limit is smaller in other configurations considered in this work but remains larger than $N_{\text{proc}} = 64$ when $N_{\text{relax}} = 0$ is used. As a last remark, we notice in Figure 14 that, under the same PinT configurations, the simulations using SL-SI-SETTLS as a coarse scheme are more expensive than those using IMEX, which may be due to additional costs linked to the semi-Lagrangian approach (trajectory estimations and interpolations to departure points).

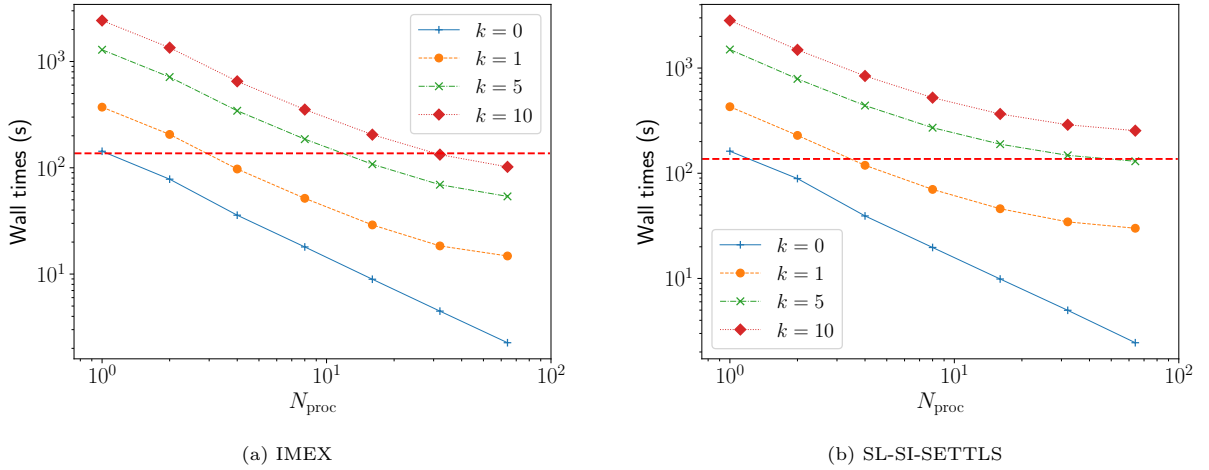


Figure 14: Gaussian bumps test case: wall times for completing k iterations of PinT simulations for the configuration $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}, q_1, \nu_1) = (2, 2, 0, 51, 2, 10^6)$ using IMEX (left) or SL-SI-SETTLS (right) on the coarse level as a function of the number of parallel processors in time. The horizontal, dashed line indicates the computing time of the reference solution, computed using $\Delta t_{\text{ref}} = \Delta t_0 = 60\text{s}$.

We now study the speedups w.r.t. the relative error in the geopotential field for simulations using IMEX in Figure 15a. The presented configurations are chosen to compare the influence of the PinT parameters individually. First, we observe that the use of a less refined resolution $M_{\text{coarse}} = 51$ reduces the cost of the time integration on the coarse level, leading to a speedup of approximately 5.5 for reducing the error by a factor close to 5; however, the convergence rapidly stagnates, with increasing computational times not leading to significant error reductions; if one wants to obtain smaller errors, it is necessary to use a larger coarse spectral resolution, which increases the computational cost of the temporal parallelization but leads to speedups still larger than one if enough processors are used. Concerning the relaxation strategy, it has been seen in Figure 8 that only slight improvements in convergence are obtained by increasing N_{relax} . A more expensive relaxation, however, strongly increases the computational cost of the MGRIT simulation since a larger number of time steps needs to be computed per iteration; therefore, increasing N_{relax} negatively impacts the speedup for obtaining approximately the same errors, and $N_{\text{relax}} = 0$ provides the better compromise. Finally, we compare the influence of the viscosity orders and coefficients applied on the coarse

levels: the worst compromise between speedup and convergence is obtained by using a too-large second-order viscosity, which overdamps the entire wavenumber spectrum and leads to larger errors despite the convergent behavior. Better results are obtained using a large higher-order viscosity, mainly in the beginning of the wavenumber spectrum, with a speedup of approximately 2.6 to obtain relative errors close to 10^{-4} under $R_{\text{norm}} = 32$. Still, the overall trade-off between convergence and computational cost is similar to moderate second-order viscosity configurations.

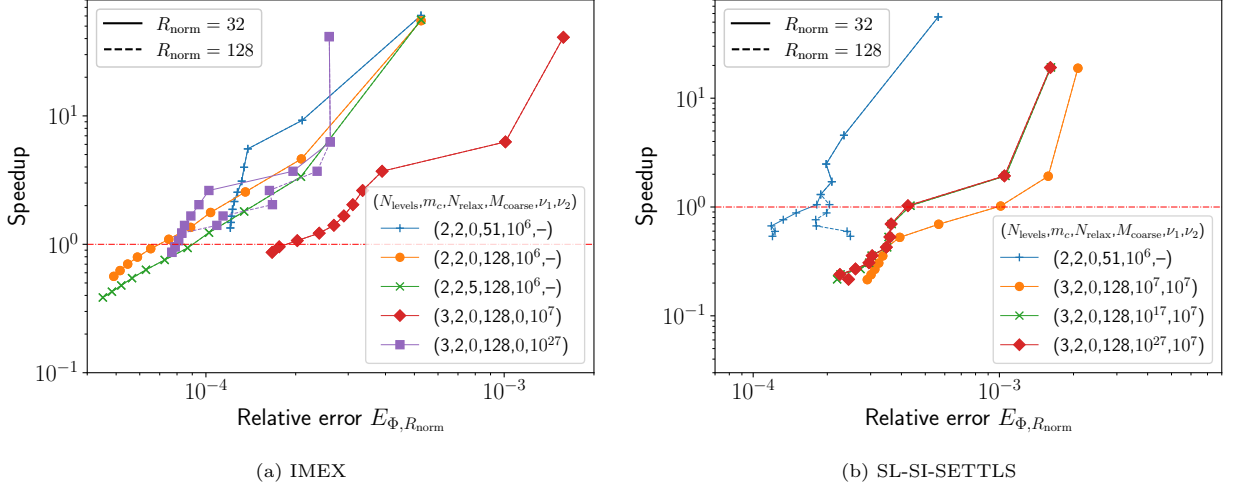


Figure 15: Gaussian bumps test case: speedup as a function of the relative geopotential error in the spectral space for chosen configurations using IMEX (left) or SL-SI-SETTLES (right) as coarse time-stepping scheme. Simulations are identified by $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}, \nu_1, \nu_2)$ and each data point corresponds to an iteration. The horizontal, dashed-dotted line indicates a unitary speedup. All simulations use $N_{\text{proc}} = 64$ parallel processors in time.

Figure 15b presents the speedups as a function of the errors of some of the few relatively stable simulations using SL-SI-SETTLES on the coarse levels, both with the default moderate second-order viscosity or other viscosity configurations. In the former case, only the less aggressive configuration $(N_{\text{levels}}, m_c) = (2, 2)$ remains stable and can improve the convergence with speedups larger than the unity. Still, after two iterations, the convergence deteriorates at the end of the wavenumber spectrum. In the other simulations, a better convergence behavior is obtained due to the larger second-order viscosity on the coarsest level; however, due to this same reason, mainly when a large second-order viscosity is also used on the intermediate level, the initial errors are larger than with $\nu_1 = 10^6 \text{m}^2 \text{s}^{-1}$ and even after ten iterations, they do not provide smaller errors than this case; moreover, the speedups are considerably smaller due to the larger spectral resolution on the coarse levels and go below the unity after two iterations.

5.2. Unstable jet test case

We now consider the test case presented by [20]. In this test, a stationary zonal jet is perturbed by a Gaussian bump in the geopotential field, leading to the formation of vortices and a rapid energy transfer from low to high wavenumbers, which may be especially challenging in the context of PinT methods. This test case was used by [33] for studying the temporal parallelization of the SWE on the rotating sphere using PFASST, and we consider the same simulation length in time, namely $T = 144\text{h}$. As in the Gaussian bumps test case, we begin by choosing a temporal discretization size to be used as a reference and fine solution for the study performed here. We perform simulations using a spectral resolution $M_0 = 256$, timestep $\Delta t_0 \in [2, 960](\text{s})$ and no artificial viscosity ($\nu_0 = 0$) and evaluate the errors of the geopotential field both in physical and spectral spaces w.r.t. to a solution computed using $M = 512$ and $\Delta t = 2$, also computed without viscosity. In all cases, the integration is performed using IMEX. Figure 16 shows that the errors for $\Delta t_0 \leq 120\text{s}$ are mainly due to the spatial discretization, with no visible dependence of the spectral errors on Δt_0 , for all R_{norm} values, and a slight increase of the physical error between $\Delta t_0 = 60\text{s}$ and $\Delta t_0 = 120\text{s}$.

For all tested values $\Delta t_0 > 120$ s, instabilities develop and the simulations are not able to reach $t = T$. Thus, we choose $\Delta t_0 = 60$ s, allowing, in the context of Parareal and MGRIT, a stable simulation on the finest level and the use of not-too-large time steps on the coarse ones. Figure 17 presents the final vorticity field produced by the chosen fine discretization and by the simulation using $M = 512$ and $\Delta t = 2$.




Figure 16: Unstable jet test case: relative L_2 error (computed in the physical space) and spectral errors (for various values of R_{norm}) between a solution obtained with spectral resolution $M = 512$ and time step $\Delta t = 2$ and solutions obtained with $M = M_0 = 256$ and various time steps. IMEX is used in all cases.




Figure 17: Unstable jet test case: solution at $t = T = 144$ h computed with IMEX. Zoom on the north hemisphere, the vorticity field being negligible in the south one.

5.2.1. Convergence study

For the study of convergence and stability of Parareal and MGRIT, we consider the same set of parameters as in the Gaussian bumps test case, namely with $N_{\text{levels}} \in \{2, 3\}$, $m_c \in \{2, 4\}$, $N_{\text{relax}} \in \{0, 1, 5\}$ and $M_{\text{coarse}} \in \{51, 128\}$. However, instead of using a fixed second-order artificial viscosity approach, with the same coefficient applied on all coarse levels, we use the results obtained in Section 5.1.2 in order to improve the performance of the temporal parallelization. In the simulations using IMEX on the coarse levels, we consider a higher-order viscosity; however, due to the higher complexity of the unstable jet test case compared to the Gaussian bumps one, with a larger temporal domain and possibly more challenging stability, we choose here to use fourth instead of sixth-order viscosity; moreover, the viscosity coefficients are chosen as a function of the time step size used on each level, such that larger coefficients are applied when stability constraints are more restrictive. In the case of SL-SI-SETTLS, we apply large second-order viscosity coefficients on all coarse levels, independently of their discretizations, due to more critical stability behavior. Table 1 summarizes the adopted viscosity coefficients.

Figure 18 presents the evolution of the PinT errors along iterations when IMEX is used on the coarse levels. As in the Gaussian bumps test case, only slight improvements are observed by using more expensive

Coarse scheme	IMEX	SL-SI-SETTLS
Viscosity order q	4	2
$\Delta t = 2\Delta t_0$	10^{16}	10^7
$\Delta t = 4\Delta t_0$	10^{17}	10^7
$\Delta t = 16\Delta t_0$	10^{17}	10^7

Table 1: Unstable jet test case: viscosity coefficients (in m^qs^{-1}) applied on the coarse levels of the PinT simulations as a function of the time step and coarse time-stepping scheme.

relaxation strategies; therefore, we only present the results for $N_{\text{relax}} = 0$. Moreover, the observed stability and convergence behaviors are the same in all spatial scales, with the large ones dominating the error magnitudes, such that the convergence curves are identical under $R_{\text{norm}} = 32$ and $R_{\text{norm}} = 128$; we, therefore, present only the former case (these same remarks being valid for the results using SL-SI-SETTLS). Finally, in order to evaluate the proposed viscosity approach, we also present the results using the same viscosity order and coefficients initially adopted in the previous test case ($q = 2$ and $\nu = 10^6\text{m}^2\text{s}^{-1}$ on all coarse levels).

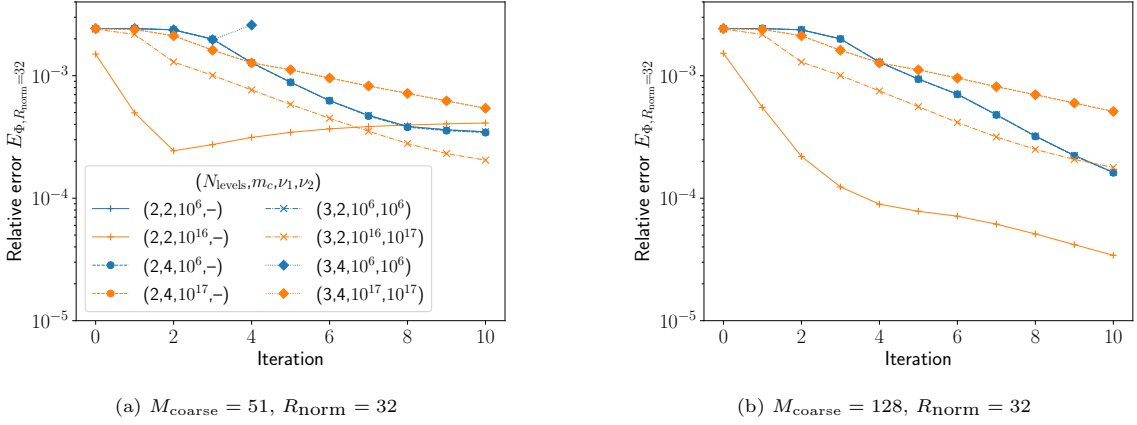


Figure 18: Unstable jet test case: relative error $E_{\Phi, R_{\text{norm}}=32}$ between the PinT and fine solutions at $t = T$ along iterations for $M_{\text{coarse}} = 51$ (left) and $M_{\text{coarse}} = 128$ (right), with IMEX used on the coarse levels. Results are identical under $R_{\text{norm}} = 128$. Simulations are identified by $(N_{\text{levels}}, m_c, \nu_1, \nu_2)$. The curves corresponding to the second-order viscosity visually coincide, as well as the curves $(N_{\text{levels}}, m_c) = (2, 4)$ and $(N_{\text{levels}}, m_c) = (3, 4)$ in the fourth-order viscosity case.

Under the second-order viscosity approach, convergence, and stability is effectively obtained, except for the most aggressive configuration $((N_{\text{levels}}, m_c) = (3, 4))$, which stops after three and one iteration, respectively in the cases $M_{\text{coarse}} = 51$ and $M_{\text{coarse}} = 128$, due to instabilities; however, all other simulations present an almost identical convergence behavior: it probably indicates that the second-order viscosity causes too large damping on the coarse levels and only a few contributions to the fine solution come from them. Indeed, proper choices of fourth-order viscosity coefficients provide better results: a much faster convergence is obtained in the configuration $(N_{\text{levels}}, m_c, M_{\text{coarse}}) = (2, 2, 128)$, in which a moderate viscosity coefficient ($\nu_1 = 10^{16}\text{m}^4\text{s}^{-1}$) is applied on the coarse level; a slower one (but still faster compared to the second-order viscosity) is obtained under $(N_{\text{levels}}, m_c) = (3, 2)$, in which a large viscosity ($\nu_2 = 10^{17}\text{m}^4\text{s}^{-1}$) is applied only on the coarsest level; however, all simulations using this large viscosity on all coarse levels $((N_{\text{levels}}, m_c) = (2, 4)$ and $(N_{\text{levels}}, m_c) = (3, 4))$ converge almost identically and slower than the second-order viscosity case, indicating an excessive damping. An unexpected result concerns the less aggressive configuration $((N_{\text{levels}}, m_c, M_{\text{coarse}}) = (2, 2, 51))$, which diverges after a very fast convergence in the first two iterations, probably indicating the triggering of numerical instabilities.

The importance of the choice of the viscosity approach is even clearer in the simulations using SL-SI-SETTLS (Figure 19). The second-order viscosity with a moderate coefficient ($\nu_{\text{coarse}} = 10^6\text{m}^2\text{s}^{-1}$) is not

able to ensure convergence and stability in any of the PinT configurations. For that, it is necessary to increase the viscosity coefficient to $\nu_{\text{coarse}} = 10^7 \text{m}^2 \text{s}^{-1}$. However, it leads to excessive damping on the coarse levels, and all configurations present the same convergence behavior, which stagnates around errors one to two orders of magnitude larger than those obtained using IMEX as a coarse scheme. Therefore, the viscosity setting does allow improving the PinT simulations using SL-SI-SETTLS, but it seems quite challenging to find a sweet spot between accuracy on the coarse levels and convergence.

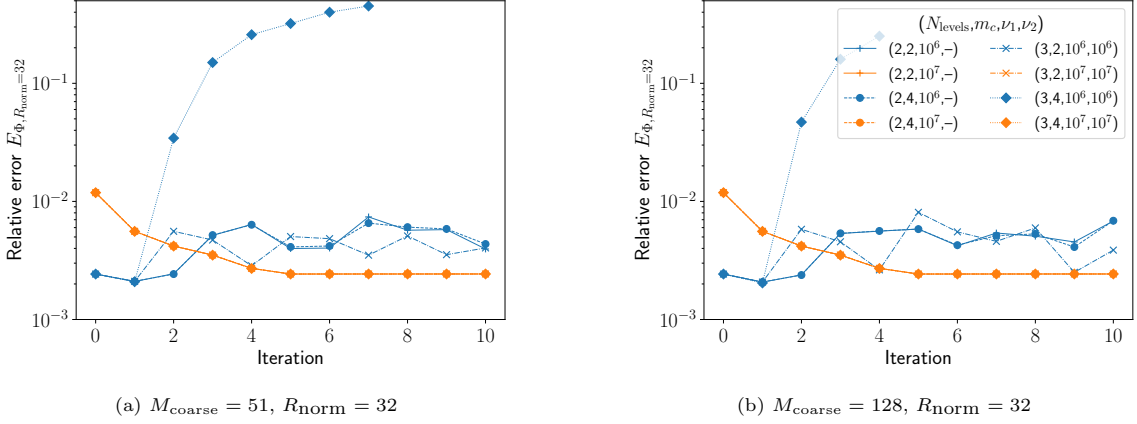


Figure 19: Unstable jet test case: relative error $E_{\Phi, R_{\text{norm}}=32}$ between the PinT and fine solutions at $t = T$ along iterations, for $M_{\text{coarse}} = 51$ (left) and $M_{\text{coarse}} = 128$ (right), with SL-SI-SETTLS used on the coarse levels. Results are identical under $R_{\text{norm}} = 128$. Simulations are identified by $(N_{\text{levels}}, m_c, \nu_1, \nu_2)$. The curves corresponding to $\nu_1 = 10^7 \text{m}^2 \text{s}^{-1}$ and $\nu_2 = 10^7 \text{m}^2 \text{s}^{-1}$ visually coincide.

Figure 21 presents the absolute difference between the PinT and reference vorticity fields at iterations $k = 0$ and $k = 5$ or $k = 10$ for chosen configurations, namely $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) \in \{(2, 2, 0, 51), (2, 2, 0, 128), (3, 2, 0, 51)\}$ and $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (2, 2, 0, 51)$, respectively with IMEX and SL-SI-SETTLS as coarse scheme, all of them using the viscosity configurations depicted in Table 1. The evolution of the respective kinetic energy spectra is illustrated in Figure 20. Concerning the configurations $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (2, 2, 0, 51)$ and $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (2, 2, 0, 128)$ using IMEX, which use a moderate viscosity coefficient on the coarse level, the kinetic energy spectrum is relatively close to the reference one already in the initial iteration; after five iterations, a slight amplification of medium wavenumbers is observed in the case $M_{\text{coarse}} = 51$, which translates to small-scale oscillations in the physical plot, despite a global reduction of the error magnitudes, whereas a much more important error decrease, without amplification of the spectrum, is obtained with $M_{\text{coarse}} = 128$. In the configuration $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}) = (3, 2, 0, 51)$ using IMEX, the larger fourth-order viscosity coefficient applied on the coarsest level leads to a larger damping of the spectrum at iteration $k = 0$, including small wavenumbers. It is clearly seen as large-scale oscillations in the respective physical plot. After five iterations, only smaller-scale errors are observed, but the spectrum is still outperformed by the initial one of the simulations using a smaller viscosity coefficient. Finally, in the simulation using SL-SI-SETTLS, in which a very aggressive second-order viscosity is applied, drastic damping is observed along the entire spectrum in the initial iteration, leading to very large-scale physical errors. The spectrum converges slowly: after ten iterations, it remains highly damped, and large-scale errors are still observed in the physical plot.

5.2.2. Evaluation of computing times and speedups

We now proceed to the evaluation of the computing times and speedups provided by the PinT simulations in the unstable jet test case, whose reference solution takes a computing time $T_{\text{ref}} \approx 546 \text{s}$. As in the Gaussian bumps test case, we consider $N_{\text{proc}} = 64$ in all simulations.

Figure 22 presents the speedups as a function of the relative error on the geopotential field for some chosen configurations using IMEX or SL-SI-SETTLS on the coarse levels. In the former case, we consider two

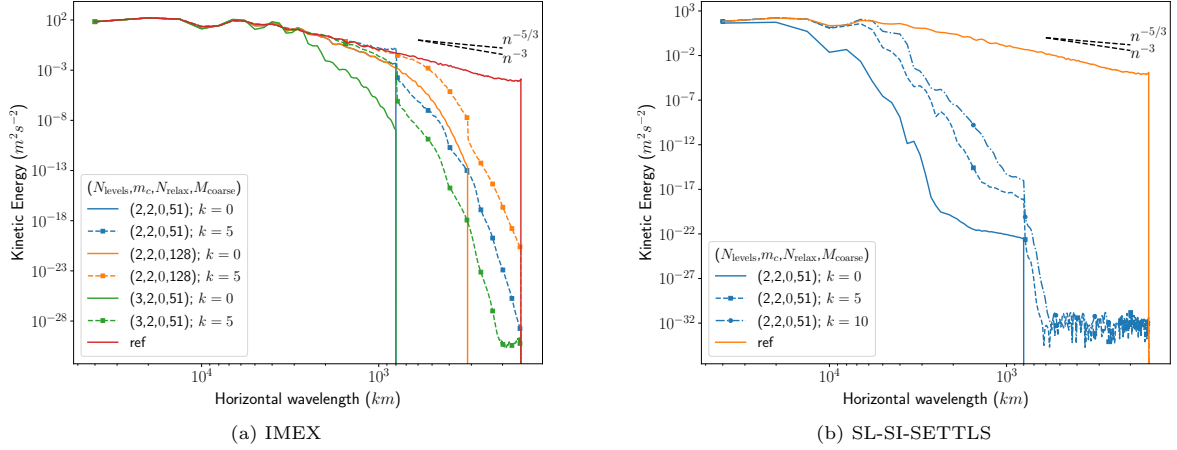


Figure 20: Unstable jet test case: kinetic energy spectra at $t = T$ of the reference (fine) and MGRIT solutions at given iterations k for the configurations illustrated in Figure 21. Left and right: IMEX and SL-SI-SETTLES as coarse time-stepping schemes.

pairs of configurations, namely $(N_{\text{levels}}, m_c, M_{\text{coarse}}) = (2, 2, 128)$ and $(N_{\text{levels}}, m_c, M_{\text{coarse}}) = (3, 2, 51)$; in each pair, one configuration uses a second-order viscosity and the other a fourth-order one. Contrary to the Gaussian bumps simulations, we now observe, in this more complex test case, characterized by more important interactions between wavenumber modes, clearer differences between the viscosity approaches in terms of compromise between convergence and numerical acceleration. We easily see that the simulations using the second-order approach provide less interesting results in terms of speedup since almost no convergence is observed in the first three iterations. Better results are obtained using the fourth-order viscosity, mainly the configuration $(N_{\text{levels}}, m_c, M_{\text{coarse}}) = (2, 2, 128)$, despite its larger spectral resolution on the coarse levels. Indeed, two iterations provide nearly the same error as ten iterations of the configuration $(N_{\text{levels}}, m_c, M_{\text{coarse}}) = (3, 2, 51)$ using fourth-order viscosity, with approximate speedup factors of 2.7 and 1.9, respectively. Finally, for these two configurations, the solutions whose error plots are depicted in Figure 21, corresponding to the iteration $k = 5$, are obtained with respective speedups of 1.1 and 3.7.

In the case where SL-SI-SETTLES is used as a coarse time-stepping scheme, we recall that all stable simulations present approximately the same convergence behavior, due to the overdamping induced by the large second-order viscosity approach. Therefore, the comparison of the speedups allows us to study the influence of the choice of Parareal and MGRIT parameters on the computational cost. In Figure 22b, we only present simulations using the coarse spectral resolution $M_{\text{coarse}} = 51$ and the errors under $R_{\text{norm}} = 32$. The largest and smallest speedups are provided respectively by $(N_{\text{levels}}, m_c) = (3, 4)$ and $(N_{\text{levels}}, m_c) = (2, 2)$, whose coarsest levels, on which the time integration is serial, use the largest and smallest time step sizes, with speedups of approximately 3.5 and 1.1 before the stagnation of the convergence. The simulations $(N_{\text{levels}}, m_c) = (3, 2)$ and $(N_{\text{levels}}, m_c) = (2, 4)$, whose coarsest levels use the same time step size, present intermediate and very similar speedup results.

5.3. Discussion

The study presented in this section aimed to evaluate the performance, in terms of stability, convergence and computational cost, of Parareal and MGRIT applied to the numerical integration of the SWE on the rotating sphere. More specifically, the main goal was to study the influence of the choice of time stepping scheme used on the coarse levels, which is known to be a crucial factor for the PinT performance, and how the results can be improved by proper parametric and discretization choices. The obtained results provide indeed indications on which characteristics a coarse scheme should have in order to lead to a successful temporal parallelization, *i.e.*, with stable and relatively accurate solutions being provided within smaller computational times than serial integrations.

From the numerical simulations, and also from the analytical stability study developed in Section 4, it is clear that it is necessary but not sufficient for the coarse scheme to be stable: poor results are obtained when

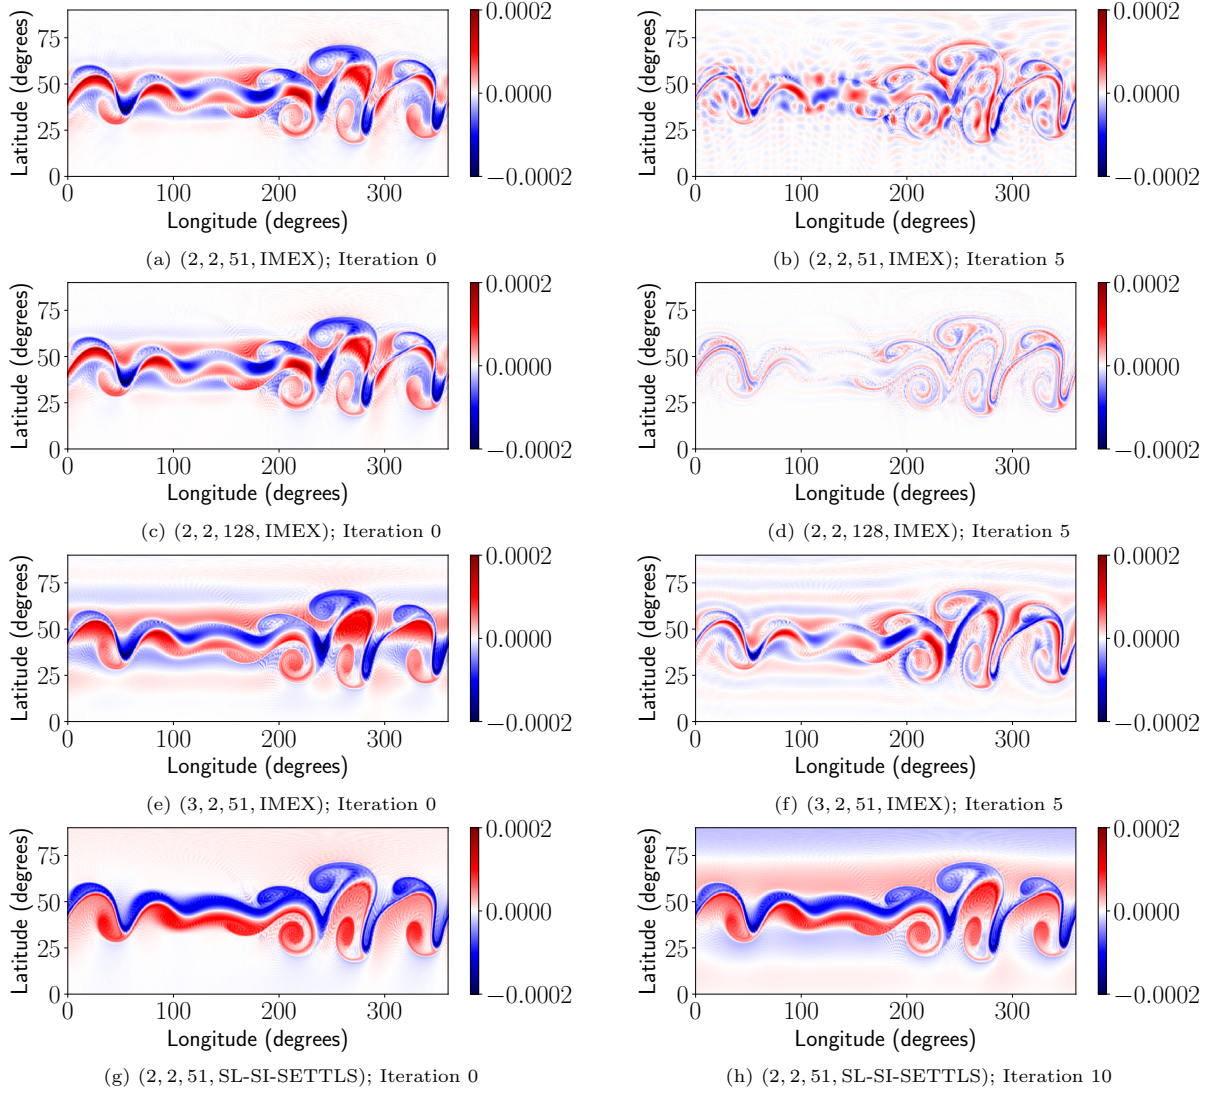


Figure 21: Unstable jet test case: the absolute difference between the PinT and the fine vorticity fields at $t = T$ under chosen configurations identified by $(N_{\text{levels}}, m_c, M_{\text{coarse}}, \text{coarse time-stepping scheme})$ at iterations 0 (left) and 5 or 10 (right). All simulations use $N_{\text{relax}} = 0$ and the viscosity configurations depicted in Table 1.

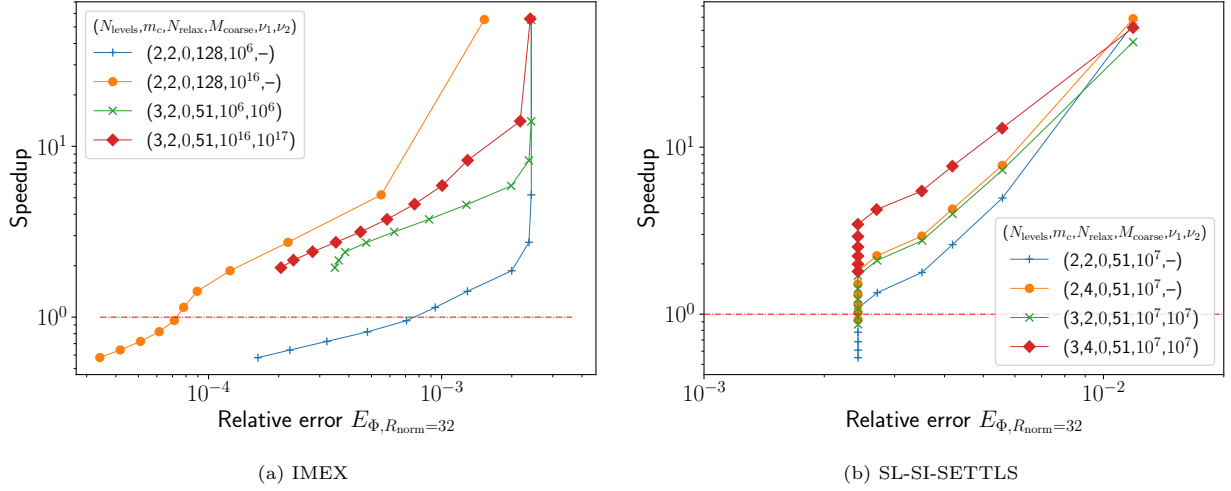


Figure 22: Unstable jet test case: speedup with respect to the relative geopotential error in spectral space for chosen configurations using IMEX or SL-SI-SETTLS as coarse time-stepping scheme. Simulations are identified by $(N_{\text{levels}}, m_c, N_{\text{relax}}, M_{\text{coarse}}, \nu_1, \nu_2)$ and each data point corresponds to an iteration. The horizontal, dashed-dotted line indicates a unitary speedup. All simulations use $N_{\text{proc}} = 64$ parallel processors in time.

the temporal parallelization uses SL-SI-SETTLS on the coarse levels, despite the stability and popularity of this scheme for the serial integration of atmospheric models and recent works indicating that the use of semi-Lagrangian approaches on the coarse levels improves the performance of PinT methods applied to simpler problems.

Moreover, the coarse discretization should be able to accurately represent large spatial scales and ensure stability on the fine ones. Indeed, the best results using IMEX as a coarse scheme are obtained with large higher-order viscosity coefficients on the coarse levels, damping the largest wavenumbers but preserving the intermediate and smallest ones. In the configurations using SL-SI-SETTLS on the coarse levels, the very restrictive stability constraints require the use of large second-order viscosities, which damps a large range of the spectrum, providing stability at the expense of a great accuracy loss.

The detailed and level-dependent study conducted on the influence of the adopted viscosity approach provides other insightful conclusions. The tests were performed considering viscosity values around typical values adopted on atmospheric modeling, and it was found that relatively large values are required for ensuring stability in the PinT framework. However, it was verified that the artificial viscosity is required only on the coarsest levels, on which the stability constraints are more severe due to the large time step sizes adopted. It implies that the fine dynamics do not need to be damped on the fine discretization levels, and more accurate and faster converging PinT solutions can be expected.

Finally, these conclusions imply that a compromise between accuracy and stability of the coarse discretization has to be found in order to achieve speedup using PinT methods. Our results indicate that, if the conditions of the coarse discretization described above are met, and by using enough parallel resources, Parareal and MGRIT using IMEX are able to provide relatively accurate solutions within shorter computational times compared to fine reference simulations.

6. Conclusion and perspectives

In this work, we have studied the temporal parallelization of the shallow water equations on the rotating sphere using Parareal and MGRIT. The development and application of PinT methods to hyperbolic problems still advance relatively slowly due to the well-known stability and convergence issues when applied to simple problems such as the advection equation, with the choice of coarse temporal discretization being a crucial aspect in this context. Therefore, the focus here was to investigate if popular and well-established

time-stepping schemes in the atmospheric modeling community can provide good performance results when used on the coarse levels of the PinT methods. The two considered schemes, IMEX and SL-SI-SETTLS, allow the use of relatively large time steps in serial integrations since they avoid stability issues linked to the stiff terms of the governing equations.

Two approaches were considered in this study. First, we conducted an analytical stability investigation of the PinT methods applied to a linearized ODE. A notable result is that Parareal and MGRIT using SL-SI-SETTLS as a coarse scheme present very poor stability properties compared to IMEX. Second, we performed numerical simulations of two test cases with increasing complexity. This confirmed the poor stability using SL-SI-SETTLS, with only very restrictive PinT configurations presenting a stable behavior. Better results were obtained with IMEX, but still relatively limited in terms of choice of several levels and coarsening factors in time. A further investigation indicated that better choices of artificial viscosity parameters only on the coarse levels allow for improved stability and convergence by using second- and higher-order viscosities for SL-SI-SETTLS and IMEX. In the former case, however, the stability is obtained at the cost of a strong damping of the entire wavenumber spectrum, which reduces the accuracy provided by the temporal parallelization. To the best of our knowledge, this is the first time that a level-selective viscosity has been investigated in PinT methods. Finally, studies on parallel performance indicated that the best compromises between accuracy and computational cost, with speedups larger than one, are obtained using IMEX as a coarse scheme, mainly due to its better stability and convergence properties in the PinT framework.

In summary, the study presented in this work indicates that inferior results are obtained when the temporal parallelization uses SL-SI-SETTLS on the coarse levels, despite the stability and popularity of this scheme for the serial integration of atmospheric model and recent works indicating that the use of semi-Lagrangian approaches on the coarse levels improves the performance of PinT methods. Better results are obtained using IMEX, but depending on the discretization applied on each level and, consequently, on appropriate choices of viscosity and hyperviscosity parameters, whose influence is especially remarkable in the more complex unstable jet test case. Under these conditions, and by using enough parallel resources, Parareal and MGRIT using IMEX can provide relatively accurate solutions within shorter computational times compared to fine reference simulations.

An important and open challenge concerns the gap to be filled between the results obtained in this work and the application of parallel-in-time methods for real weather and climate problems. Through parametrization and discretization choices coherent with practical applications, we tried, to the extent possible, to provide insights in this direction, *e.g.*, with the parameters considered in the analytical stability study, the spectral resolutions adopted, the ranges of tested viscosity and hyperviscosity coefficients and the simulation of quite challenging test cases, mainly the unstable jet one, which is a standard test in atmospheric modeling due to its complex dynamics. However, more refined investigations would need to be conducted to reach practical applications in operational models.

A natural future work consists of investigating if modifications of SL-SI-SETTLS and the use of other time-stepping schemes, eventually not still used operationally in atmospheric models or even similar schemes but with higher discretization orders, could provide better stability and convergence properties for Parareal and MGRIT. For instance, exponential integration methods, which can integrate precisely the linear terms of the governing equations, have aroused a growing interest in the context of atmospheric modeling, *e.g.*, in [26, 55, 25]. It includes a semi-Lagrangian variant of this family of method, which have been proposed and used for solving the SWE on the plane by [47], and to the SWE on the rotating sphere with a cubed sphere spatial discretization by [56]. Ongoing studies indicate that the use of exponential schemes and their semi-Lagrangian versions on the coarse discretization levels indeed provides better stability and convergence properties for Parareal and MGRIT, which will be presented in the following article.

Finally, the results presented in this work illustrate different behaviors of the wavenumber spectra of the solution along the PinT iterations, with the adopted artificial viscosity approach being crucial for ensuring stability and convergence, and notably being required only on the coarse discretization levels. Since the SWEs are a nonlinear model, with nonlinear interactions between modes transferring energy to the highest wavenumbers, future works should focus on understanding how each time-stepping scheme treats these interactions and how they influence the performance of the temporal parallelization.

CRediT authorship contribution statement

João G. Caldas Steinstraesser: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing - Original Draft, Visualization. **Pedro da Silva Peixoto:** Conceptualization, Methodology, Resources, Writing - Review & Editing, Supervision, Project administration, Funding acquisition. **Martin Schreiber:** Conceptualization, Methodology, Software, Validation, Formal analysis, Resources, Writing - Review & Editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The source code and numerical tests presented here are available for reproducibility in the repository

https://gitlab.inria.fr/sweet/sweet/-/tree/58a4c051fe1e928573284d9d94aaf571104d9e05/benchmarks_sphere/paper_jcp_pint_imex_slsi

Acknowledgements

Funding: This work was supported by the São Paulo Research Foundation (FAPESP) grants 2021/03777-2 and 2021/06176-0, as well as the Brazilian National Council for Scientific and Technological Development (CNPq), Grant 303436/2022-0. This project also received funding from the Federal Ministry of Education and Research and the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955701, Time-X. The JU receives support from the European Union’s Horizon 2020 research and innovation programme and Belgium, France, Germany, Switzerland.

Most of the computations presented in this paper were performed using the GRICAD infrastructure (<https://gricad.univ-grenoble-alpes.fr>), which is supported by Grenoble research communities.

Appendix A. Viscosity and hyperviscosity in spectral methods

We make a brief overview of the theory of viscosity approaches in spectral methods and some considerations about the orders of magnitude to be chosen for the viscosity coefficients, following the presentation in [36], to which we refer the reader for further details.

As mentioned in Section 2, the viscosity approach with even order $q \geq 2$ applied to a field ψ consists in

$$\frac{\partial \psi}{\partial t} = (-1)^{\frac{q}{2}+1} \nu \nabla^q \psi \quad (\text{A.1})$$

where the $(-1)^{\frac{q}{2}+1}$ allows to define $\nu \geq 0$. In the spectral space of spherical harmonics, (A.1) reads

$$\frac{\partial \psi_{m,n}}{\partial t} = (-1)^{\frac{q}{2}+1} \nu \left(\frac{-n(n+1)}{a^2} \right)^{\frac{q}{2}} \psi_{m,n} = -\nu \left(\frac{n(n+1)}{a^2} \right)^{\frac{q}{2}} \psi_{m,n}$$

whose exact solution is

$$\psi_{m,n}(t) = \psi_{m,n}(0) \exp \left(-\nu \left(\frac{n(n+1)}{a^2} \right)^{\frac{q}{2}} t \right)$$

The viscosity coefficient can be determined by setting a time τ in which a given wavenumber n_0 is damped to a fraction b_{τ, n_0} of $\psi_{m,n}(0)$. Typically, ν is set such that the largest wavenumber $n_0 = M$ damps to a fraction $b_{\tau, n_0} = 1/e$, i.e.,

$$\nu = \frac{1}{\tau} \left(\frac{M(M+1)}{a^2} \right)^{-\frac{q}{2}} \quad (\text{A.2})$$

With this viscosity parameter being applied to the entire spectrum, a given mode $\psi_{m,n}$ is damped after a time step Δt by a factor

$$\begin{aligned} b_{\Delta t, n} &:= \exp \left(-\nu \left(\frac{n(n+1)}{a^2} \right)^{\frac{q}{2}} \Delta t \right) = \exp \left(-\frac{\Delta t}{\tau} \left(\frac{n(n+1)}{M(M+1)} \right)^{\frac{q}{2}} \right) \\ &\approx \left[1 + \Delta t \nu \left(\frac{n(n+1)}{a^2} \right)^{\frac{q}{2}} \right]^{-1} =: \hat{b}_{\Delta t, n} \end{aligned}$$

where $\hat{b}_{\Delta t, n}$ corresponds to a backward Euler discretization of (A.1), as considered in this work.

Figure A.23a illustrates the orders of magnitude of the viscosity coefficient values for viscosity orders $q \in \{2, 4, 6\}$ as a function of the damping time τ and for the two spectral resolutions considered on the coarse discretization levels of the PinT numerical simulations performed here, namely $M \in \{51, 128\}$. Damping timescales reported and suggested in the literature for spectral models usually range around units or tenths of hours, depending on the spectral resolution (see *e.g.*, [9, 31, 68]). In Figure A.23b we plot the discrete damping factor $\hat{b}_{\Delta t, n}$ w.r.t. the wavenumber modes n for some viscosity coefficients around the approximate average values depicted in Figure A.23a and time step $\Delta t = 120\text{s}$. We observe that smaller-order viscosities lead to a faster decay of smaller wavenumbers and a larger sensitivity on the coefficient value (compare the spacing between the curves corresponding to the same viscosity order). It indicates that a higher-order viscosity produces a more refined stability filter, damping only the largest wavenumbers and better preserving the large-scale spatial features of the solution.

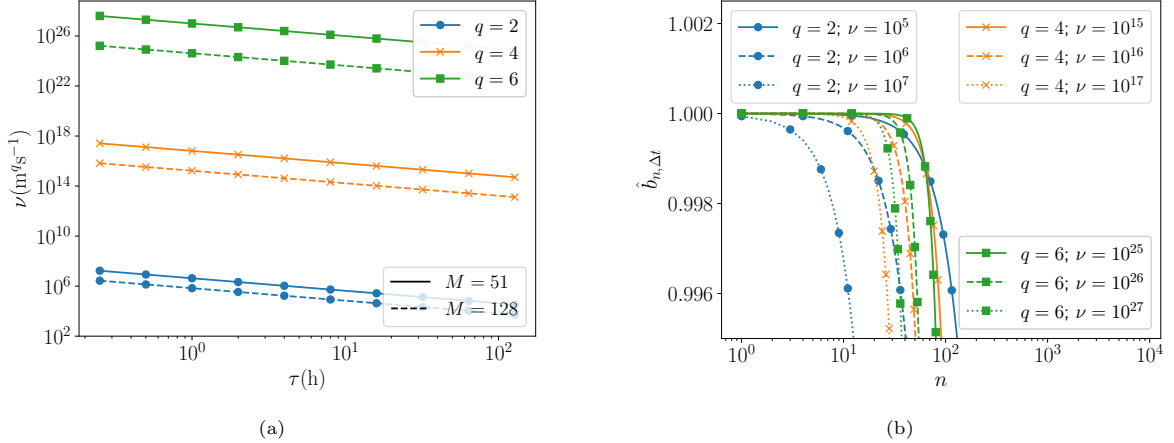


Figure A.23: (a): viscosity coefficient of orders $q \in \{2, 4, 6\}$ as a function of the damping time for chosen spectral resolutions M . Computed using (A.2). (b): discrete damping factor $\hat{b}_{\Delta t, n}$ as a function of the spectral mode for viscosity orders 2, 4 and 6 and chosen viscosity coefficients (in $\text{m}^q \text{s}^{-1}$). Computed considering $\Delta t = 120\text{s}$.

References

- [1] XBraid: Parallel multigrid in time. <http://llnl.gov/casc/xbraid>.
- [2] N. Abel, J. Chaudhry, R. D. Falgout, and J. Schroder. Multigrid-reduction-in-time for the rotating shallow water equations, 8 2020.
- [3] Matteo Astorino, Franz Chouly, and Alfio Quarteroni. Multiscale coupling of finite element and lattice Boltzmann methods for time dependent problems. Technical report, October 2012.

- [4] L. Baffico, S. Bernard, Y. Maday, G. Turinici, and G. Zérah. Parallel-in-time molecular-dynamics simulations. *Physical Review E*, 66(5), November 2002.
- [5] Guillaume Bal. On the convergence and the stability of the parareal algorithm to solve partial differential equations. In Timothy J. Barth, Michael Griebel, David E. Keyes, Risto M. Nieminen, Dirk Roose, Tamar Schlick, Ralf Kornhuber, Ronald Hoppe, Jacques Périaux, Olivier Pironneau, Olof Widlund, and Jinchao Xu, editors, *Domain Decomposition Methods in Science and Engineering*, pages 425–432, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [6] Guillaume Bal and Yvon Maday. A “parareal” time discretization for non-linear PDE’s with application to the pricing of an american put. In *Lecture Notes in Computational Science and Engineering*, pages 189–202. Springer Berlin Heidelberg, 2002.
- [7] Alfredo Bellen and Marino Zennaro. Parallel algorithms for initial-value problems for difference and differential equations. *Journal of Computational and Applied Mathematics*, 25(3):341–350, May 1989.
- [8] L.A. Berry, W. Elwasif, J.M. Reynolds-Barredo, D. Samaddar, R. Sanchez, and D.E. Newman. Event-based parareal: A data-flow based implementation of parareal. *Journal of Computational Physics*, 231(17):5945–5954, July 2012.
- [9] Byron A. Boville. Sensitivity of simulated climate to model resolution. *Journal of Climate*, 4(5):469–485, May 1991.
- [10] Achi Brandt. Multi-level adaptive solutions to boundary-value problems. *Mathematics of Computation*, 31(138):333–390, 1977.
- [11] M. H. Carpenter, C. A. Kennedy, Hester Bijl, S. A. Viken, and Veer N. Vatsa. Fourth-order runge–kutta schemes for fluid mechanics applications. *Journal of Scientific Computing*, 25(1):157–194, October 2005.
- [12] Feng Chen, Jan S. Hesthaven, and Xueyu Zhu. *On the Use of Reduced Basis Methods to Accelerate and Stabilize the Parareal Method*, pages 187–214. Springer International Publishing, Cham, 2014.
- [13] S. M. Cox and P. C. Matthews. Exponential time differencing for stiff systems. *Journal of Computational Physics*, 176(2):430–455, 2002.
- [14] Xiaoying Dai and Yvon Maday. Stable parareal in time method for first and second order hyperbolic system. Technical report, December 2011.
- [15] Dale Durran. *Numerical methods for fluid dynamics : with applications to geophysics*. Springer, New York, 2010.
- [16] ECMWF and P. White. IFS Documentation CY25R1 - Part III: Dynamics and Numerical Procedures. 2003.
- [17] Matthew Emmett and Michael L. Minion. Toward an Efficient Parallel in Time Method for Partial Differential Equations. *Communications in Applied Mathematics and Computational Science*, 7:105–132, 2012.
- [18] R. D. Falgout, S. Friedhoff, T. V. Kolev, Scott P. MacLachlan, and Jacob B. Schroder. Parallel time integration with multigrid. *SIAM Journal on Scientific Computing*, 36:C635–C661, 2014.
- [19] S. Friedhoff, R. D. Falgout, T. V. Kolev, Scott P. MacLachlan, and Jacob B. Schroder. A Multigrid-in-Time Algorithm for Solving Evolution Equations in Parallel. In *Presented at: Sixteenth Copper Mountain Conference on Multigrid Methods, Copper Mountain, CO, United States, Mar 17 - Mar 22, 2013*, 2013.
- [20] Joseph Galewsky, Richard K. Scott, and Lorenzo M. Polvani. An initial-value problem for testing numerical models of the global shallow-water equations. *Tellus A: Dynamic Meteorology and Oceanography*, 56(5):429–440, January 2004.
- [21] Martin J. Gander. Analysis of the Parareal Algorithm Applied to Hyperbolic Problems using Characteristics. *Bol. Soc. Esp. Mat. Apl.*, 42:21–35, 2008.
- [22] Martin J. Gander. 50 years of time parallel time integration. In *Contributions in Mathematical and Computational Sciences*, pages 69–113. Springer International Publishing, 2015.
- [23] Martin J. Gander, Felix Kwok, and Hui Zhang. Multigrid interpretations of the parareal algorithm leading to an overlapping variant and MGRIT. *Computing and Visualization in Science*, 19(3-4):59–74, June 2018.
- [24] Martin J. Gander and Stefan Vandewalle. Analysis of the parareal time-parallel time-integration method. *SIAM J. Scientific Computing*, 29:556–578, 01 2007.
- [25] Stéphane Gaudreault, Martin Charron, Valentin Dallerit, and Mayya Tokman. High-order numerical solutions to the shallow-water equations on the rotated cubed-sphere grid. *Journal of Computational Physics*, 449:110792, January 2022.
- [26] Stéphane Gaudreault and Janusz A. Pudykiewicz. An efficient exponential time integration method for the numerical solution of the shallow water equations on the sphere. *Journal of Computational Physics*, 322:827–848, October 2016.
- [27] Jürgen Geiser and Stefan Güttel. Coupling methods for heat transfer and heat flow: Operator splitting and the parareal algorithm. *Journal of Mathematical Analysis and Applications*, 388(2):873–887, April 2012.
- [28] Sebastian Götschel, Michael Minion, Daniel Ruprecht, and Robert Speck. Twelve ways to fool the masses when giving parallel-in-time results. In *Springer Proceedings in Mathematics & Statistics*, pages 81–94. Springer International Publishing, 2021.
- [29] Sebastian Götschel and Michael L. Minion. Parallel-in-time for parabolic optimal control problems using PFASST. In *Lecture Notes in Computational Science and Engineering*, pages 363–371. Springer International Publishing, 2018.
- [30] James Hack and Ruediger Jakob. Description of a global shallow water model based on the spectral transform method. Technical report, 1992.
- [31] Kevin Hamilton, Yoshiyuki O. Takahashi, and Wataru Ohfuchi. Mesoscale spectrum of atmospheric motions investigated in a very fine resolution global general circulation model. *Journal of Geophysical Research*, 113(D18), September 2008.
- [32] François P. Hamon, Martin Schreiber, and Michael L. Minion. Parallel-in-time multi-level integration of the shallow-water equations on the rotating sphere. *Journal of Computational Physics*, 407:109210, April 2020.
- [33] François P. Hamon, Martin Schreiber, and Michael L. Minion. Parallel-in-time multi-level integration of the shallow-water equations on the rotating sphere. *Journal of Computational Physics*, 407:109210, 2020.
- [34] Terry Haut and Beth Wingate. An asymptotic parallel-in-time method for highly oscillatory PDEs. *SIAM Journal on Scientific Computing*, 36(2):A693–A713, 2014.
- [35] Mariano Hortal. The development and testing of a new two-time-level semi-lagrangian scheme (SETTSL) in the ECMWF

- forecast model. *Quarterly Journal of the Royal Meteorological Society*, 128(583):1671–1687, July 2002.
- [36] Christiane Jablonowski and David L. Williamson. The pros and cons of diffusion, filters and fixers in atmospheric general circulation models. In Peter Lauritzen, Christiane Jablonowski, Mark Taylor, and Ramachandran Nair, editors, *Numerical Techniques for Global Atmospheric Models*, chapter 13, pages 381–493. Springer Berlin Heidelberg, 2011.
 - [37] M. Kanamitsu, K. Tada, T. Kudo, N. Sato, and S. Isa. Description of the JMA operational spectral model. *Journal of the Meteorological Society of Japan. Ser. II*, 61(6):812–828, 1983.
 - [38] James F. Kelly and Francis X. Giraldo. Continuous and discontinuous galerkin methods for a scalable three-dimensional nonhydrostatic atmospheric model: Limited-area mode. *Journal of Computational Physics*, 231(24):7988–8008, October 2012.
 - [39] John N. Koshyk and Kevin Hamilton. The horizontal kinetic energy spectrum and spectral budget simulated by a high-resolution troposphere–stratosphere–mesosphere GCM. *Journal of the Atmospheric Sciences*, 58(4):329–348, February 2001.
 - [40] Peter Lauritzen, Christiane Jablonowski, Mark Taylor, and Ramachandran Nair, editors. *Numerical Techniques for Global Atmospheric Models*. Springer Berlin Heidelberg, 2011.
 - [41] Erik Lindborg. Can the atmospheric kinetic energy spectrum be explained by two-dimensional turbulence? *Journal of Fluid Mechanics*, 388:259–288, June 1999.
 - [42] Jacques-Louis Lions, Yvon Maday, and Gabriel Turinici. Résolution d'EDP par un schéma en temps “pararéel”. *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, 332(7):661–668, April 2001.
 - [43] Gianmarco Mengaldo, Andrzej Wyszogrodzki, Michail Diamantakis, Sarah-Jane Lock, Francis X. Giraldo, and Nils P. Wedi. Current and emerging time-integration strategies in global numerical weather and climate prediction. *Archives of Computational Methods in Engineering*, 26(3):663–684, February 2018.
 - [44] Michael Minion, Brandon Krull, Mathew Emmett, and Sebastian Goetschel. libpfasst v1.0. [Computer Software] <https://doi.org/10.11578/dc.20180711.5>, jun 2018.
 - [45] U.S. National Oceanic and Atmospheric Administration (NOAA) / National Centers for Environmental Prediction (NCEP). Global Forecast System - Global Spectral Model (GSM) - V13.0.2. <https://vlab.noaa.gov/web/gfs/documentation>, May 2016.
 - [46] Benjamin W. Ong and Jacob B. Schroder. Applications of time parallelization. *Computing and Visualization in Science*, 23(1-4), September 2020.
 - [47] Pedro S. Peixoto and Martin Schreiber. Semi-lagrangian exponential integration with application to the rotating shallow water equations. *SIAM Journal on Scientific Computing*, 41(5):B903–B928, 2019.
 - [48] B. Philippi and T. Slawig. A Micro-Macro Parareal Implementation for the Ocean-Circulation Model FESOM2, 2023.
 - [49] Benedict Philippi and Thomas Slawig. The Parareal Algorithm Applied to the FESOM 2 Ocean Circulation Model, 2022.
 - [50] M. Ries and U. Trottenberg. Mgr-ein blitzschneller elliptischer löser. *Tech. Rep. Preprint*, 277(SFB 72), 1979.
 - [51] Daniel Ruprecht. Wave propagation characteristics of Parareal. *Computing and Visualization in Science*, 19:1–17, june 2018.
 - [52] Daniel Ruprecht and Rolf Krause. Explicit parallel-in-time integration of a linear acoustic-advection system. *Computers Fluids*, 59:72–83, Apr 2012.
 - [53] D. Samaddar, D.P. Coster, X. Bonnin, L.A. Berry, W.R. Elwasif, and D.B. Batchelor. Application of the parareal algorithm to simulations of ELMs in ITER plasma. *Computer Physics Communications*, 235:246–257, February 2019.
 - [54] A. Schmitt, M. Schreiber, P. Peixoto, and M. Schäfer. A numerical study of a semi-lagrangian parareal method applied to the viscous burgers equation. *Computing and Visualization in Science*, 19(1-2):45–57, June 2018.
 - [55] M. Schreiber, N. Schaeffer, and R. Loft. Exponential integrators with parallel-in-time rational approximations for shallow-water equations on the rotating sphere. *Parallel Computing*, 2019.
 - [56] Vladimir V. Shashkin and Gordey S. Goyman. Semi-lagrangian exponential time-integration method for the shallow water equations on the cubed sphere grid. *Russian Journal of Numerical Analysis and Mathematical Modelling*, 35(6):355–366, December 2020.
 - [57] William C. Skamarock, Joseph B. Klemp, Michael G. Duda, Laura D. Fowler, Sang-Hun Park, and Todd D. Ringler. A multiscale nonhydrostatic atmospheric model using centroidal voronoi tessellations and c-grid staggering. *Monthly Weather Review*, 140(9):3090–3105, September 2012.
 - [58] Gunnar Andreas Staff and Einar M. Rønquist. Stability of the parareal algorithm. In *Lecture Notes in Computational Science and Engineering*, pages 449–456. Springer-Verlag, 2005.
 - [59] Andrew Staniforth and Jean Côté. Semi-lagrangian integration schemes for atmospheric models—a review. *Monthly Weather Review*, 119(9):2206–2223, September 1991.
 - [60] H. De Sterck, R. D. Falgout, and O. A. Krzysik. Fast multigrid reduction-in-time for advection via modified semi-lagrangian coarse-grid operators. 2022.
 - [61] Hans De Sterck, Robert D. Falgout, Stephanie Friedhoff, Oliver A. Krzysik, and Scott P. MacLachlan. Optimizing multigrid reduction-in-time and parareal coarse-grid operators for linear advection. *Numerical Linear Algebra with Applications*, 28(4), March 2021.
 - [62] Hans De Sterck, Stephanie Friedhoff, Alexander J. M. Howse, and Scott P. MacLachlan. Convergence analysis for parallel-in-time solution of hyperbolic systems. *Numerical Linear Algebra with Applications*, 27(1), November 2019.
 - [63] Paul N. Swarztrauber. Shallow water flow on the sphere. *Monthly Weather Review*, 132(12):3010–3018, December 2004.
 - [64] Clive Temperton. Treatment of the coriolis terms in semi-lagrangian spectral models. *Atmosphere-Ocean*, 35(sup1):293–302, January 1997.
 - [65] Joana M. F. da Trindade and José F. Pereira. Parallel-in-time simulation of the unsteady Navier–Stokes equations for incompressible flow. *International Journal for Numerical Methods in Fluids*, 45(10):1123–1136, 2004.

- [66] Warren M Washington, Lawrence Buja, and Anthony Craig. The computational future for climate and earth system models: on the path to petaflop and beyond. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1890):833–846, December 2008.
- [67] David L. Williamson. The evolution of dynamical cores for global atmospheric models. *Journal of the Meteorological Society of Japan. Ser. II*, 85B:241–269, 2007.
- [68] David L. Williamson. Convergence of aqua-planet simulations with increasing resolution in the community atmospheric model, version 3. *Tellus A: Dynamic Meteorology and Oceanography*, 60(5):848, January 2008.
- [69] David L. Williamson, John B. Drake, James J. Hack, Rüdiger Jakob, and Paul N. Swarztrauber. A standard test set for numerical approximations to the shallow water equations in spherical geometry. *Journal of Computational Physics*, 102(1):211–224, September 1992.