

Time-parallelisation techniques with applications to plasma simulation



Federico Danieli
Trinity College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity term 2020/2021

Statement of Originality

Unless otherwise indicated, all the work in this thesis has been done by the author. Chapter 3 is currently under review for publication in the *SIAM Journal on Scientific Computing*. A manuscript illustrating the results in Chapter 4 is in preparation. Chapter 5 is under review for publication in the *Electronic Transaction on Numerical Analysis* journal. Chapter 6 has been published in the *Numerical Linear Algebra with Applications* journal [32].

Abstract

Parallel-in-time (PinT) methods have become an increasingly popular tool in the numerical solution of time-dependent Partial Differential Equations (PDEs), in light of their potential for extracting additional concurrency when spatial parallelism saturates. Particularly complex applications, such as the ones involving plasma simulations for fusion research, would benefit greatly from the implementation of such methods, given the large computational time usually necessary for the numerical solution of their governing equations.

In the first part of this thesis, we develop the concept of Space-Time Block Preconditioning (STBP), a technique for simplifying the task of time-parallelising the solution of a whole system of PDEs by reducing it to that of inverting space-time operators acting on single variables only, which are more amenable to be tackled by PinT algorithms. This technique borrows from block preconditioning strategies developed for spatial operators in the framework of time-stepping algorithms, and extends them to the whole space-time setting. After illustrating the guiding principles behind STBP, we measure its effectiveness with applications in fluid dynamics and magnetohydrodynamics. In both cases, the resulting preconditioners show excellent scaling properties with respect to mesh refinement, and their applications come at little-to-no overhead cost in terms of iterations to convergence with respect to their time-stepping counterpart, thus providing evidence for their potential parallel efficiency.

Ultimately, however, the parallel performance of a space-time block preconditioner depends on that of its internal components, which still rely on already-existing PinT methods. This might become an issue when dealing with advection-dominated or hyperbolic equations, as classical PinT algorithms have shown to struggle in this case. For this reason, the second part of this thesis is dedicated to investigating the causes behind this degradation in performance, as well as to study more recent PinT approaches for overcoming this limitation.

Acknowledgements

My gratitude goes first and foremost to my academic supervisor, **Andrew Wathen**: even after working with Andy for so long, I have barely managed to scratch the surface of his intuition of linear algebra, and I can only hope to make the best out of his teaching. Close second is my industrial supervisor, **Debasmita Samaddar**: Debbie's enthusiasm, support, and knack for organisation are qualities I will always aspire to.

My Viva experience was made particularly enjoyable by the examiners, **Matthias Bolten** and **Patrick Farrell**. The discussion ensued during the final exam and their valuable input contributed greatly towards the quality of this thesis, as well as in directing future work.

Also deserving my sincerest thanks are my collaborators: **Ben Southworth**, who has been an inexhaustible source of interesting ideas, and **Scott MacLachlan**, whose guidance and supervision have been priceless. Special thanks go to the people who made InFoMM possible, too: **Chris, Colin, Jonathan, Laura, Amanda, and Sarah**.

And then there are all the people who accompanied me in this experience, making it one of the most memorable. First of all, Cohort 3: a big hug goes to **Mel**, who will always be my favourite. Another one to **Ali, Ana**, and **Attila**, (and I am quite lucky their names come first alphabetically, so I can pretend that is the real reason why I put them there, and **Raquel** will not get offended), together with **Clint, Jonathan, Joe, Kris, Ollie, Scott**, and our undercover agent **Helen**. I want to thank **Matteo** for all the coffee, and apologise to the rest of my officemates, **Florian, Gonzalo** and **Nicolas**, for all the singing. Special thanks go to **Yifi**, whose support and affection have accompanied me throughout the hardest times of my DPhil. A big thank you also to **Anna, Giuseppe, Arno, Sam, Matilde, Lili** and **Diana**, who made my stay in Oxford and in Italy more light-hearted (especially during the pandemic), and to all the amazing people

I have engaged in interesting discussions with during my DPhil, and from whom I have learnt so much.

Finally, my parents, **Patrizia e Lucio**, and my brother, **Matteo**: they might never understand what this thesis is about, but for sure they understood all the effort that went into it.

Contents

1	Introduction	1
1.1	Importance of parallel algorithms	1
1.2	Relevance of parallel computing in plasma simulations	3
1.3	Motivation, scope, and structure of the thesis	4
2	Time parallelisation	7
2.1	Parallelisation of time-integration schemes	8
2.1.1	Parareal	8
2.1.2	PFASST	14
2.1.3	ParaExp	17
2.2	All-at-once methods	20
2.2.1	ParaDiag	22
2.2.2	Space-time circulant preconditioning	23
2.2.3	MGRIT	25
2.2.4	AIR	27
3	Space-time block preconditioning	29
3.1	Problem definition	30
3.1.1	Space-time discretisation	32
3.2	Space-time block preconditioning for incompressible flow	35
3.2.1	Small commutator approach	37
3.2.2	Green’s function approach	42
3.3	Eigenvalues clustering	45
3.4	Results	49
3.4.1	Model problems	50
3.4.2	Performance of preconditioner	53
3.4.3	The nonlinear setting: incompressible Navier-Stokes	57
3.4.4	Comparison with sequential time-stepping	59

3.5	Remarks on STBP for incompressible flow	60
4	Applications to incompressible magnetohydrodynamics	62
4.1	Problem definition	64
4.1.1	Space-time discretisation	69
4.2	Space-time block preconditioning for IMHD	73
4.2.1	Velocity-pressure coupling	74
4.2.2	Velocity-magnetic field coupling	75
4.2.3	Definition of the space-time block preconditioner	80
4.2.4	Comparison with single time-step block preconditioner	82
4.3	Eigenvalues clustering	84
4.4	Results	87
4.4.1	Model problems	88
4.4.2	Performance of preconditioner	94
4.4.3	Comparison with sequential time-stepping	96
4.5	Limitations of STBP for compressible MHD	98
5	Multigrid Reduction in Time for nonlinear hyperbolic equations	100
5.1	Model problems	101
5.1.1	Space discretisation	102
5.1.2	Time discretisation	108
5.2	Time parallelisation with MGRIT	109
5.2.1	Choices for the coarse integrator	114
5.3	Numerical results	116
5.3.1	Fine integrator matching	117
5.3.2	High-order schemes	119
5.3.3	Changing accuracy per coarsening level	122
5.4	Remarks on MGRIT for hyperbolic equations	122
6	Space-time circulant preconditioning for linear wave equations	124
6.1	Circulant preconditioning for Toeplitz systems	125
6.1.1	ω -circulant preconditioner	125
6.1.2	Convergence of circulant-preconditioned Krylov methods	127
6.2	All-at-once solution of linear ODEs	129
6.2.1	Discretisation aspects	129
6.2.2	Results	132
6.3	All-at-once solution of linear PDEs	133

6.3.1	Discretisation aspects	134
6.3.2	Block circulant preconditioning	137
6.3.3	Cost analysis of circulant preconditioning <i>versus</i> time-stepping	138
6.3.4	Results	140
7	Conclusion	145
7.1	Summary of findings	145
7.2	Future work	148
7.3	Final remarks	150
	Bibliography	151

Chapter 1

Introduction

In the framework of this project, we examine strategies to effectively employ parallel algorithms for the numerical solution of the equations describing the behaviour of plasma and electrically charged fluids in general. The particularity of the algorithms investigated lies in the way concurrency is achieved: the evolution of the solution is tracked simultaneously at different instants, owing to a process called *time parallelisation*.

We begin by providing a general introduction to the two main topics treated in this thesis, namely time parallelisation and plasma simulation, and then proceed to define the principal goal of this research work.

1.1 Importance of parallel algorithms

The growing complexity of systems arising from the solution of *Partial Differential Equations* (PDEs) demands ever increasing computational power in order to tackle such problems numerically in a reasonable amount of time. Unfortunately, the frequency of computations that can be carried out on a single processor unit is limited (see Fig. 1.1), and represents a fundamental upper bound on the efficiency of serial algorithms. To overcome this limit, much attention has been directed towards the implementation of parallel algorithms, which provide an alternative to sheer computational power to speed up computations.

Some standard parallelisation methods involve the decomposition of the spatial domain of the target PDE into smaller subdomains, where some computations are carried out independently [147]. Within this framework, a certain amount of information must be exchanged between the subdomains, so that the solution could be matched at the interfaces, usually in an iterative fashion. However, this causes additional overhead with respect to solving the PDE serially, so that increasing the number of

subdomains above a certain limit becomes detrimental to the purpose of speeding up the computation, and spatial parallelisation is said to *saturate*. Moreover, systems of *Initial Value Problems* (IVPs) for *Ordinary Differential Equations* (ODEs) do not present a spatial domain to subdivide in the first place, preventing the use of such techniques. Indeed, the temporal evolution of IVPs (for both ODEs and PDEs) is classically resolved via *time-stepping*, where the solution at each instant is recovered using information from the previous ones, in an inherently sequential fashion.

Nonetheless, the temporal domain is still available for parallelisation, and one can think of applying similar methodologies as for the spatial domain in order to concurrently recover the solution at various instants. Possibly the first person behind this intuition is Nievergelt, who speculated on such opportunity as early as in 1964 [110] and proposed a scheme that eventually gave birth to a whole new category of techniques for the solution of IVPs, namely *multiple shooting methods*. The additional challenge of parallelising over the temporal domain (rather than over the spatial one), stems from the causality principle: information flows forward in time, so it might seem counterintuitive to apply a parallel algorithm to a phenomenon which is sequential in nature. This notwithstanding, in recent years a renewed interest sparked the research on techniques for time parallelisation [55], brought forth both by the




Figure 1.1: Evolution of microprocessors performance. In recent years, while the clock frequency of a single processor (green squares) has plateaued, the number of cores per processors (black diamonds) has been steadily increasing, which reflects the focus on parallelisation over power. Source: [122, Creative Commons Attribution 4.0 International Public License]

increased availability of parallel computational power (with exascale computing being the next target at the time of writing [36]), and by the development of *Parareal* [92], which represents the first modern parallel-in-time (PinT) algorithm.

1.2 Relevance of parallel computing in plasma simulations

Plasma is one of the fundamental states of matter, and can be considered as an *ionised* fluid, whose charges are relatively free to move independently [66]. As a consequence, its dynamics can internally generate electromagnetic fields, and is subject to interactions with externally imposed ones. Plasma plays a fundamental role in fusion energy research, since the conditions at which fusion occurs require the nuclear fuel to be in this particular state, at least for the standard methods in which fusion is achieved nowadays. Designing a stable, efficient fusion reactor would represent a major step towards solving the problem of generating clean energy from a durable and sustainable source. However, one of the greatest challenges in this endeavour consists in controlling the behaviour of the plasma within the reactor in an effective manner.

One way to achieve this is by imposing strong magnetic fields, which force the plasma into specific trajectories, in what is called *magnetic confinement*. Unfortunately, confined plasma often shows turbulent behaviour, characterised by the build-up of instabilities [106, Chap. 6-7 and 12-13]. It is precisely the growth of these instabilities that poses the main obstacle towards effective confinement, as these eventually resolve in sudden bursts occurring at the outer region of the plasma (namely *Edge-Localised Modes*, or ELMs). Therefore, controlling them is of utmost importance: if left unchecked, in fact, they can cause damage to the interior wall of the reactor, with consequent release of impurities within the plasma, and a drastic reduction in the temperature of the core (see Fig. 1.2), which might end up stopping the fusion reaction altogether.

In order to design an effective magnetic confinement, fusion researchers resort in large measure to numerical simulations. Given the complexity of the PDEs describing plasma behaviour, however, recovering their approximate solutions is a computationally demanding task, which needs parallel computing to be completed in a timely fashion. For this purpose, time parallelisation represents an appealing technique to further reduce computational time once more classical, spatial parallelisation methods have reached saturation and capped their potential for speedup.




Figure 1.2: Picture of the *Mega Ampère Spherical Tokamak* (MAST) fusion reactor in Culham, UK. Incandescent plasma is filling the vessel, but it can only be seen at the bottom due to heat loss towards the reactor walls, which makes it radiate in the visible spectrum of light.

Research on PinT methods for plasma simulation remains however scarce. To make matters worse, in the most extreme regimes the relevant equations describing plasma behaviour are dominated by transport phenomena and, in the compressible limit, become hyperbolic in nature. This represents an additional complication to parallelisation in time, as most classical PinT algorithms tend to suffer from performance degradation when applied to this specific class of PDEs. Regardless, the recent successes from the applications of time parallelisation techniques to plasma simulation [119, 128, 129] justify the strong interest in further investigating the properties of such algorithms, on top of possible extensions and improvements, fine-tuned for the applications of interest.

1.3 Motivation, scope, and structure of the thesis

The motivation behind our work comes from the industrial partner and sponsor for this project, the *Culham Centre for Fusion Energy* (CCFE). CCFE is one of the leading research centres involved in the development of the technology necessary to attain a sustained fusion reaction, and the host to the *Joint European Torus* (JET) and the *Mega Ampère Spherical Tokamak* (MAST), two of the most advanced experimental fusion reactors in the world [151]. The interest of CCFE in time-parallelisation tech-

niques stems from the discussion outlined in Sec. 1.2, and in the scope of our work we investigate the applicability of PinT methods to the acceleration of the numerical solution of plasma simulations. This defines our research purpose.

The work conducted in this thesis aims at developing effective time-parallelisation techniques with the potential of speeding-up the numerical solution of complex systems of PDEs, specifically of models which are of relevance to the field of magnetohydrodynamics predicting plasma behaviour. In addition, it investigates the causes behind the loss of performance of available PinT algorithms when applied to the solution of hyperbolic PDEs, together with possible alternatives to overcome these limitations.

The remainder of the thesis is structured as follows.

Chapter 2 acts as an introduction to the field of parallelisation-in-time. There, we list some of the most renowned time-parallel algorithms developed in recent years, with the goal of providing the reader with an overview of state-of-the-art PinT methods. In addition to a general description of the schemes selected, we illustrate their applicability and potential for parallel efficiency by reporting some results available in the literature. In particular, we point out the difficulties some of these methods experience in their application for the solution of hyperbolic PDEs.

In Chap. 3, we introduce the novel concept of *Space-Time Block Preconditioning* (STBP), which represents the main technique developed in the framework of our work. This considers existing block preconditioning strategies, which have been studied for the acceleration of the solution of systems arising from the *spatial* discretisation of a target multi-variable PDEs, and extends them to tackle its simultaneous *space-time* discretisation. The main advantage of this approach stems from the fact that it renders the target system more prone to be time-parallelised using already-existing PinT methods for single-variable PDEs. To illustrate the guiding principles behind this strategy, we consider applications to incompressible fluid dynamics, and measure the effectiveness of our approach for several test-cases, with the goal of testing its potential for parallelisation.

The target application in our research is introduced in Chap. 4, where we expand upon the work conducted in Chap. 3, to consider model problems in magnetohydrodynamics (MHD). After a brief introduction to the relevant system of PDEs, we showcase the flexibility of our STBP strategy by designing a space-time block preconditioner for the more complex MHD system, following similar principles as in

the previous chapter. Even for this problem, we investigate the performance of the preconditioner when used to accelerate the solution of a few test-cases, aiming at evaluating its applicability for the solution of plasma simulations.

As the success of our STBP approach is ultimately reliant on the parallel effectiveness of PinT algorithms in accelerating the solution of single-variable PDEs, with Chap. 5 the focus of this thesis shifts to investigating the application of time-parallelisation methods to the challenging class of hyperbolic PDEs. The time-parallel scheme we investigate in this chapter is the MGRIT algorithm, whose performance we measure in its application to some nonlinear conservation laws. The main purpose is to shed some light on the causes behind the degradation in convergence speed observed when MGRIT is used for this class of problems.

An alternative approach for the time-parallel solution of hyperbolic problems is investigated in Chap. 6. There, we consider the use of a *circulant preconditioner* for accelerating the solution of monolithic space-time systems arising from the discretisation of linear, constant-coefficient wave equations. We provide theoretical and experimental justifications for the effectiveness of this preconditioner, which supports its applicability for the solution of this type of PDEs.

Our results are finally summarised in Chap. 7, where we also provide a list of possibilities to further expand the work in this thesis, as well as some open-ended questions which might represent interesting topics for future research.

Chapter 2

Time parallelisation

In this chapter, we present some of the most renowned modern algorithms for time parallelisation available in the literature. The methods described here have been chosen with the aim to provide a reasonable sample of the variety of guiding concepts that can inspire the design of PinT methods.

In order to better organise this exposition, we choose to subdivide the algorithms illustrated here into two groups, loosely based on the framework in which they operate. The schemes belonging to the first group, described in Sec. 2.1, consider a given time integration algorithm, and attempt to achieve parallelisation in time by introducing concurrency into its application; on the other hand, the schemes belonging to the second group, described in Sec. 2.2, consider the system arising from the simultaneous (space-)time discretisation of a differential equation, and focus on providing a parallel method for the solution of the whole system at once. We point out that this distinction is, to some extent, arbitrary, and indeed some of the methods listed here have often been categorised differently in the literature; nonetheless, we find it useful to operate this separation, in that it highlights two substantially different frameworks in which to attempt to achieve parallelisation in time.

As a caveat, the list presented here is by no means exhaustive, and it reports only a fraction of the different manners in which time parallelism can be achieved. For a broader overview, we refer to the excellent chapter on the topic by Gander, [55], which describes the evolution of PinT methods since their inception. A more recent review by Ong and Schroder, collecting performance results for time-parallelisation algorithms, can be found in [111].

2.1 Parallelisation of time-integration schemes

In aiding our presentation of this first class of PinT schemes, we introduce the generic target ODE,

$$\begin{cases} u'(t) = f(u(t), t) & t \in (T_0, T] \\ u(T_0) = \bar{u}^0, \end{cases} \quad (2.1)$$

whose integration in time we are interested in conducting in parallel. Here, \bar{u}^0 represents the initial condition, while the right-hand side $f : \mathbb{R} \times (T_0, T] \rightarrow \mathbb{R}$ is a given Lipschitz continuous function. We point out that in this introduction we focus on a scalar ODE purely for the purpose of simplifying our description, but the methods listed here can be seamlessly extended to systems thereof, including (but not limited to) PDEs discretised using a *method-of-lines* approach [130].

2.1.1 Parareal

The *Parareal* algorithm is one of the most famous PinT methods, also owing to the simplicity of its implementation. Its original formulation, which we follow here to provide a description of the scheme, was originally presented in [92] by Lions, Maday, and Turinici.

To apply the algorithm, first the temporal domain of (2.1) is split into N_p subdomains (or *time chunks*), $[T_p, T_{p+1}]$, for $p = 0, \dots, N_p - 1$, of size $\Delta T_p = T_{p+1} - T_p$ and with $T_{N_p} = T$. Parallelisation is then achieved by solving independently in each subdomain the Cauchy problems

$$\begin{cases} u'_p(t) = f(u_p(t), t) & t \in (T_p, T_{p+1}] \\ u_p(T_p) = \lambda_p \end{cases} \quad p = 0, \dots, N_p - 1. \quad (2.2)$$

We immediately realise the main difficulty we need to overcome for this approach to be effective: the initial condition is only available for the first time chunk, $\lambda_0 = \bar{u}^0$, and to collect the others exactly we would need to integrate sequentially each of the (2.2) and impose $\lambda_p = u_{p-1}(T_p)$, thus reducing the problem to the serial case. To obviate this inconvenience, Parareal works in an iterative fashion: a guess is provided for the initial conditions λ_p^k necessary to solve the problems (2.2), and this guess is refined at each iteration k .

This is achieved by introducing two numerical solvers (or *propagators*). Classically, these come in the form of time-stepping routines, but they can more generally be identified as functions that map an initial condition λ to an approximate solution of (2.2) at the end of the corresponding time chunk. The two solvers have different

Algorithm 1: Parareal

```

  /* Initialisation: */
1 Set  $\lambda_0^0 = \bar{u}^0$ ;
2 for  $p = 0$  to  $N_p - 2$  do
3   | Set  $\lambda_{p+1}^0 = \mathcal{G}_p(\lambda_p^0)$ ; //  $\mathcal{G}$  is used to get an initial guess for  $\lambda_p^0$ 

  /* Main Parareal iterations: */
4 for  $k = 0$  to  $It_{\max}$  do
   |
   | /* Parallel step: */
   | 5 parfor  $p = k$  to  $N_p - 1$  do
   | 6   | Compute  $\mathcal{F}_p(\lambda_p^k)$ ;
   |
   | /* Serial (update) step: */
   | 7 Set  $\lambda_{k+1}^{k+1} = \mathcal{F}_k(\lambda_k^k)$ ; // Update (2.3a)
   | 8 for  $p = k + 1$  to  $N_p - 2$  do
   | 9   | Compute  $\mathcal{G}_p(\lambda_p^{k+1})$ ;
   | 10  | Update  $\lambda_{p+1}^{k+1} = \mathcal{F}_p(\lambda_p^k) + \mathcal{G}_p(\lambda_p^{k+1}) - \mathcal{G}_p(\lambda_p^k)$ ; // Update (2.3b)
  
```

properties: the *fine* one, \mathcal{F}_p , is more precise and expensive, while the *coarse* one, \mathcal{G}_p , is cheaper to compute but potentially less accurate¹. The fine solver is employed to solve the N_p Cauchy problems (2.2) in parallel at the desired level of accuracy; the task of the coarse solver is instead to quickly propagate forward in time corrections to the eventual discrepancies between the guessed initial conditions λ_p^k and the actual results from the integration on the previous time chunk, $\mathcal{F}_{p-1}(\lambda_{p-1}^k)$. This is achieved by applying the following update formula for the initial conditions:

$$\begin{cases} \lambda_{k+1}^{k+1} = \mathcal{F}_k(\lambda_k^k) & (2.3a) \\ \lambda_{p+1}^{k+1} = \mathcal{F}_p(\lambda_p^k) + \mathcal{G}_p(\lambda_p^{k+1}) - \mathcal{G}_p(\lambda_p^k), & p = k + 1, \dots, N_p - 2, \end{cases} \quad (2.3b)$$

where an initial guess is given, *e.g.*, by a first serial application of the coarse solver $\lambda_p^0 = \mathcal{G}_p(\lambda_{p-1}^0)$, starting from the available initial condition $\lambda_0^0 = \bar{u}^0$. Pseudo-code for the corresponding algorithm is provided in Alg. 1, while a sketch of its behaviour is given in Fig. 2.1.

From (2.3a), we can see that at each iteration k , the initial condition for the $(k + 1)$ -th subdomain λ_{k+1} becomes exact and does not need to be updated any more. This reflects the fact that information is propagating forward in time and, by that iteration, the applications of the fine solver have caught up with that time chunk.

¹The subscripts p remind us that the action of each solver can in general depend on the time chunk p we are integrating on, in particular if the right-hand side of (2.1) is time-dependent.




Figure 2.1: Example of the first iteration of the Parareal algorithm. To parallelise the computation of the solution $u(t)$ to (2.1) (black curve), we begin by splitting the domain into time chunks. An initial guess for the values of the solution at the beginning of each time chunk, λ_p^0 , is collected using a first serial sweep of the coarse solver \mathcal{G} (light blue), starting from the known value u_0 . Using \mathcal{F} , integration is then conducted in parallel within each time subdomain (dark blue). At each iteration k , the k -th time chunk has achieved convergence. To update the other initial conditions, we perform a serial sweep of the coarse solver and sequentially apply formula (2.3b).

Indeed, after N_p iterations, the algorithm always converges to the solution one would obtain by serially applying the fine solver; for there to be any gain in employing Parareal, though, we should request the number of iterations to be much smaller than N_p . Moreover, (2.3b) states that the update needs to be carried out serially, since λ_{p+1}^{k+1} depends explicitly on λ_p^{k+1} ; this, together with the fact that at least one fine integration needs to be performed, effectively caps the maximum parallel efficiency of the algorithm (intended as the ratio between time-to-solution in the serial and parallel case, divided by number of processors used). An alternative approach stemming from the attempt of improving upon this limitation is described in Sec. 2.1.2.

Since the inception of the algorithm, a number of interpretations have been given to the update formula (2.3). For example, following [54], consider the system

$$\begin{bmatrix} \lambda_0 \\ \lambda_1 - \mathcal{F}_0(\lambda_0) \\ \lambda_2 - \mathcal{F}_1(\lambda_1) \\ \vdots \\ \lambda_{N_p-1} - \mathcal{F}_{N_p-2}(\lambda_{N_p-2}) \end{bmatrix} = \begin{bmatrix} \bar{u}^0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \quad (2.4)$$

Solving this system means finding the correct initial conditions λ_p so to match the result from the integration over the previous time chunk, $\mathcal{F}_p(\lambda_{p-1})$, which is the goal of the Parareal algorithm. Applying Newton iterations for solving (2.4) would give

rise to the following iterative update:

$$\begin{cases} \lambda_{k+1}^{k+1} = \mathcal{F}_k(\lambda_k^k) & (2.5a) \\ \lambda_{p+1}^{k+1} = \mathcal{F}_p(\lambda_p^k) + \frac{\partial \mathcal{F}_p(\lambda)}{\partial \lambda} \Big|_{\lambda=\lambda_p^k} (\lambda_p^{k+1} - \lambda_p^k), \quad p = k+1, \dots, N_p-2, & (2.5b) \end{cases}$$

where $\partial \mathcal{F}_p(\lambda)/\partial \lambda$ represents the sensitivity of the result from the p -th integrator with respect to variations in initial conditions. Comparing (2.5) with (2.3), we can view the application of the coarse solver as a way to provide an approximation for the sensitivities, which would classify Parareal as a quasi-Newton method [34, Chap. 6]. Similar in spirit is the interpretation provided in [96], which casts Parareal as a gradient descent algorithm applied to an optimal control problem, whose objective function directly relates to the norm of the residual in (2.4), or the one in [62], which sees the algorithm as a multiple shooting technique. Alternatively, in [96], as well as in [124], the algorithm is instead described as a preconditioned iterative method. All these different views on Parareal provide a varied framework for analysing its characteristics, which we proceed to describe in the following section.

2.1.1.1 Stability and issues with hyperbolic equations

Parareal is a relatively mature algorithm, whose properties have been studied extensively over the course of the last two decades. In this section we focus on reporting some results from the literature regarding the stability of the method, as well as on discussing the performance of the algorithm when applied to different classes of PDEs.

To this purpose, we introduce a convergence result presented in [62], recovered for the Dahlquist equation² ($f(u(t), t) = zu(t)$ in (2.1), with $z \in \mathbb{C}$), and which relies on one-step time-stepping algorithms as fine and coarse solvers to be used in the update formula (2.3). For simplicity, we consider time chunks of uniform size $\Delta T_p \equiv \Delta T$, so that $\mathcal{F}_p \equiv \mathcal{F}$ and $\mathcal{G}_p \equiv \mathcal{G}$ for each $p = 0, \dots, N_p-1$. Notice that in this simplified framework we can express the application of the solvers as multiplication by their corresponding *stability functions* [71, Chap. IV.3]: $\mathcal{G}u = \Phi_{\mathcal{G}}(z\Delta t/N_{\mathcal{G}})^{N_{\mathcal{G}}}u$, and $\mathcal{F}u = \Phi_{\mathcal{F}}(z\Delta t/N_{\mathcal{F}})^{N_{\mathcal{F}}}u$, where $N_{\mathcal{F}}$ and $N_{\mathcal{G}}$ are the number of time-steps per time chunk for fine and coarse solver, respectively. We also require that both solvers act on their region of absolute stability, given by $|\mathcal{F}| < 1$ and $|\mathcal{G}| < 1$. The bound is on the infinity norm of the error of the k -th Parareal iteration with respect to the result

²A somewhat more general framework for the convergence of Parareal is provided in [5], but we found the bound presented here better suited for describing the characteristics of Parareal outlined in this section.

from the serial application of \mathcal{F} , $\|e^k\|_\infty = \max\{\mathcal{F}\bar{u}^0 - \lambda_1^k, \dots, \mathcal{F}^{N_p-1}\bar{u}^0 - \lambda_{N_p-1}^k\}$, and it reads

$$\|e^k\|_\infty \leq |\mathcal{G} - \mathcal{F}|^k \min \left\{ \binom{N_p - 1}{k}, \left(\frac{1 - |\mathcal{G}|^{N_p-1}}{1 - |\mathcal{G}|} \right)^k \right\} \|e^0\|_\infty. \quad (2.6)$$

The binomial coefficient reminds us once again that the algorithm always converges in a finite number of steps; moreover, it does so superlinearly as k increases. However, since our focus is on the regime $N_p \gg 1$ and $k \ll N_p$, to achieve fast convergence we need to ensure that

$$\frac{|\mathcal{G} - \mathcal{F}|}{1 - |\mathcal{G}|} < 1, \quad (2.7)$$

as to guarantee that the error decreases at each iteration. Loosely speaking, this condition suggests that the coarse solver should be at the same time dissipative enough (in order to keep the denominator reasonably large), and a good approximation of the fine solver (so that the numerator remains reasonably small). This is still feasible when the dynamics described by the target differential equation incorporate some form of smoothing, and indeed Parareal has proven to be effective in accelerating the solution of parabolic problems. Examples of successful applications for this class of PDEs are reported in [98] with a theoretical speedup with respect to sequential time-stepping of $22\times$ using 50 cores, in [86] ($7\times$ speedup for 32 cores), in [20] ($300\times$ speedup with 1600 cores), and in [6] ($6.25\times$ speedup for 50 cores). When moving to transport-dominated equations, however, one can expect that ensuring (2.7) might become harder. In the case of pure transport, in fact, accurate solvers should preserve the energy within the system: as a consequence, the conditions on the coarse solver become even tighter, since this cannot be too dissipative.

It is possible to illustrate the difference in performance of Parareal when applied to a parabolic or hyperbolic problem with the example PDE

$$\begin{cases} \frac{\partial u}{\partial t}(x, t) = -v \frac{\partial u}{\partial x}(x, t) + \mu \frac{\partial^2 u}{\partial x^2}(x, t) & (x, t) \in [0, L] \times (0, T] \\ u(x, 0) = \frac{e^{-\frac{x^2}{2}}}{\sqrt{2\pi}} & x \in [0, L] \\ u(0, t) = u(L, t) & t \in (0, T] \end{cases}, \quad (2.8)$$

where, by varying its parameters, one can tweak the nature of the system from purely advective ($\mu = 0$) to purely diffusive ($v = 0$). The resulting convergence behaviour of the algorithm is compared in Fig. 2.2. While for the parabolic case the error diminishes relatively quickly, we can see how it remains substantially unchanged for the hyperbolic case, up until the very last iteration, when it must converge. Even

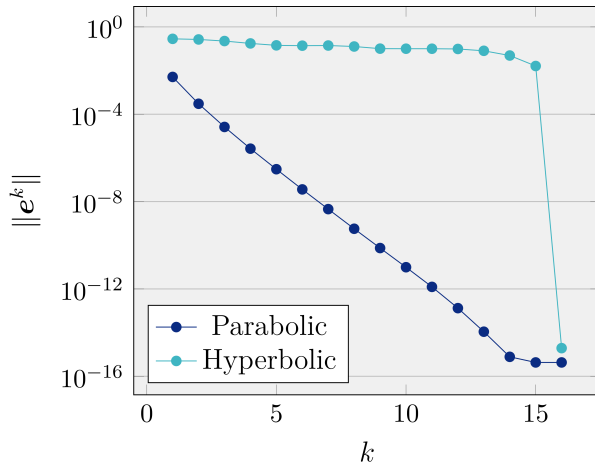


Figure 2.2: Evolution of the L^∞ norm of the error in both space and time for Parareal applied to (2.8), in the case $v = 0$, $\mu = 0.1$ (light blue), and $v = 1$, $\mu = 0$ (dark blue). Simulation parameters for both cases: *Implicit Euler* as coarse solver, exact fine solver, $N_p = T = 16$, $L = 2\pi$. Spectral discretisation in space with $N_x = 61$ modes.

worse scenarios, where the error spikes during the first $\sim N_p/2$ iterations to start decreasing only afterwards, are far from uncommon (see for example [62, Fig. 5.1]). Already in [143, Fig. 2 and 3] and [62, Fig. 4.1] it is illustrated how eigenvalues with large imaginary part tend to fall outside the stability region of Parareal for numerous choices of coarse/fine solvers, providing evidence for undermining the applicability of the vanilla version of this algorithm to systems where transport phenomena are dominant.

Investigating the causes for this instability, as well as identifying effective countermeasures, is in large part still a matter of research, and we address some of these issues in Chap. 5. The work conducted in [123, 124] provides a useful intuition behind this undesirable behaviour, and connects it to the *local* nature of formula (2.3): when dealing with a vector unknown, (as it might arise from the spatial discretisation of a PDE), the update acts component-wise, so Parareal has no way to counteract errors stemming from potential relative shifts in the solution of the coarse solver with respect to that of the fine one. In solving hyperbolic problems, these occur whenever there is a mismatch in the wave propagation speed registered by the two solvers.

This notwithstanding, the literature presents a number of examples where Parareal is successful, even in the unfavourable transport-dominated case: for example, in [26], convergence of Parareal applied to Burgers' equation is preserved by projecting the solution back to an energy manifold after every iteration; in [38] the authors show the potential for Parareal in speeding up the solution to a hydrodynamic problem char-

acterised by a Reynolds number of $5 \cdot 10^4$; in [109], choosing a Roe-average Riemann solver seems to be key for fast convergence when solving the shallow-water equations, reaching a parallel efficiency of $\sim 35\%$. On top of this, numerous attempts at stabilising the algorithm are also reported, for example in [18, 19, 63], and are at the base of the development of alternative methods, such as the one presented in Sec. 2.1.3.

2.1.2 PFASST

The *Parallel Full Approximation Scheme in Space and Time* (PFASST) was first introduced by Emmett and Minion in [43]. Similarly to Parareal, this is an iterative method which employs solvers of different levels of accuracy in order to refine and propagate the solution along the time domain. In contrast to Parareal, however, the exchange of information between the levels is performed using a *Full Approximation Scheme* (FAS), a concept originally developed for multigrid techniques [12, Chap. 6]. Moreover, in its original formulation, the numerical integrators used at each level belong to the class of *Spectral Deferred Correction* (SDC) iterative methods [37], rather than being time-stepping schemes. The combination of SDC and FAS represents the core of the PFASST algorithm: the method in fact makes use of both techniques in order to propagate information about the evolution of the solution forward in time, as well as to refine it in parallel to an increased level of accuracy.

More in detail, as with the Parareal algorithm described in Sec. 2.1.1, the temporal domain of (2.1) is separated into N_p chunks (2.2), each assigned to a different processor. Within each chunk $p = 0, \dots, N_p - 1$, two different discretisations of the time domain are introduced, $\{t_{m,p}^G\}_{m=0}^{N_p^G-1} \subset [T_p, T_{p+1}] \supset \{t_{m,p}^F\}_{m=0}^{N_p^F-1}$, where $N_p^G < N_p^F$, so that the superscript G corresponds to the coarser level, and F to the finer one. We further request the first and last point of each time chunk to be included in these discretisations, *i.e.*, $t_{0,p}^G = t_{0,p}^F = T_p$ and $t_{N_p^G-1,p}^G = t_{N_p^F-1,p}^F = T_{p+1}$, $\forall p$. The two levels communicate between each other using a pair of interpolation/restriction operators, respectively denoted as I_G^F and R_F^G , which map solutions from coarse to fine level and vice-versa; an SDC solver is also associated to each level. For each main iteration k of PFASST, we define the approximate fine solution within the p -th time chunk at the i -th SDC sub-iteration as:

$$\tilde{\mathbf{u}}_p^{k,i} = \left[\tilde{u}_{0,p}^{k,i}, \dots, \tilde{u}_{N_p^F-1,p}^{k,i} \right]^T \approx \mathbf{u}_p = \left[u(t_{0,p}^F), \dots, u(t_{N_p^F-1,p}^F) \right]^T. \quad (2.9)$$

In complete analogy, the coarse-level approximate solution is denoted as:

$$\tilde{\mathbf{y}}_p^{k,i} = \left[\tilde{y}_{0,p}^{k,i}, \dots, \tilde{y}_{N_p^G-1,p}^{k,i} \right]^T \approx \mathbf{y}_p^k = \left[y^k(t_{0,p}^G), \dots, y^k(t_{N_p^G-1,p}^G) \right]^T. \quad (2.10)$$

Upon initialisation of the PFASST algorithm, the initial condition \bar{u}^0 is broadcast to all processors p and restricted to the coarse level, in order to build

$$\tilde{\mathbf{u}}_{p,IN} = [\bar{u}^0, \dots, \bar{u}^0]^T \quad \text{and} \quad \tilde{\mathbf{y}}_{p,IN}^0 = R_F^G(\tilde{\mathbf{u}}_{p,IN}), \quad (2.11)$$

so to provide the initial guess for the coarse SDC solver. Here, the subscript $(\cdot)_{,IN}$ identifies quantities computed during the initialisation phase. Then, in a similar fashion to the Parareal algorithm, the initial conditions are serially updated using the coarse SDC solver, and propagated forward. For each processor, this amounts to iteratively applying the SDC update a total of N_{SDC}^G times, and sending the last component of the recovered solution to the following processor, which will then substitute it to the initial component of its own solution before applying the coarse solver in its turn. Contrarily to Parareal, however, in PFASST the processors $p > 0$ do not need to remain idle while waiting for the updated initial conditions to reach them, and can instead start applying SDC sweeps straight away. Thus, at the end of the initialisation phase, each processor p has completed a total of $(p+1)N_{SDC}^G$ SDC iterations and hence recovered $\tilde{\mathbf{y}}_{p,IN}^{(p+1)N_{SDC}^G}$. The initialisation terminates by interpolating the coarse solution back to the fine level, providing the initial conditions for the fine SDC sweep:

$$\tilde{\mathbf{u}}_p^{0,0} = \tilde{\mathbf{u}}_{p,IN} + I_G^F \left(\tilde{\mathbf{y}}_p^{(p+1)N_{SDC}^G} - \tilde{\mathbf{y}}_p^0 \right). \quad (2.12)$$

At each main iteration k of the PFASST algorithm, each processor independently performs a number N_{SDC}^F of SDC sweeps at the fine level, in order to refine the approximate solution to the target ODE. Updated information from the previous processor is then collected and restricted to the coarse level, where it is quickly propagated in time using the coarse SDC sweeps. The initial guess on the coarse solution is given by the restriction on the coarse grid of the current value of the approximate fine solution, but with the first component modified. This component is in fact recovered from the previous time-chunk:

$$\tilde{\mathbf{y}}_p^{k,0} = R_F^G \left(\left[\tilde{u}_{N_{p-1}^F-1,p-1}^{k+1,0}, \tilde{u}_{1,p}^{k,N_{SDC}^F}, \dots, \tilde{u}_{N_{p-1}^F,p}^{k,N_{SDC}^F} \right]^T \right). \quad (2.13)$$

Upon termination of the coarse SDC sweeps, the resulting correction is interpolated back to fine level and used to retrieve the starting point for the next PFASST iteration:

$$\tilde{\mathbf{u}}_p^{k+1,0} = \left[\tilde{u}_{N_{p-1}^F-1,p-1}^{k+1,0}, \tilde{u}_{1,p}^{k,N_{SDC}^F}, \dots, \tilde{u}_{N_{p-1}^F,p}^{k,N_{SDC}^F} \right]^T + I_G^F \left(\tilde{\mathbf{y}}_p^{k,N_{SDC}^F} - \tilde{\mathbf{y}}_p^{k,0} \right). \quad (2.14)$$

The iteration ends by sending the last component $\tilde{u}_{M_{p-1}^F,p}^{k+1,0}$ to the following time-chunk, before starting the new SDC sweeps at the fine level. The complete pseudo-code for the algorithm described is given in Alg. 2.

While this description is limited to the use of a pair of fine/coarse solvers, this framework can be extended to a complete hierarchy of discretisation grids. Even more generally, if we are solving PDEs rather than ODEs, the refinement level could vary in the space domain as well as in the time domain. In that case, some adjustments can be applied to the description of the algorithm in order to improve its speed, giving priority to the forward propagation of information, before performing the remaining tasks: that is to say, trying to reduce as much as possible the computational time between the `receive` and `send` instructions (lines 16 and 22 of Alg. 2 respectively).

Algorithm 2: PFASST

```

/* Initialisation: */
1 parfor  $p = 0$  to  $N_p - 1$  do
2   Broadcast initial conditions  $\bar{u}^0$ ;
3   Restrict to coarse level and recover  $\tilde{\mathbf{y}}_{p,IN}^0$ ; // Formula (2.11)
4   Apply  $pN_{SDC}^G$  coarse SDC sweeps to recover  $\tilde{\mathbf{y}}_{p,IN}^{pN_{SDC}^G}$ ;
5   if  $p > 0$  then
6     Receive the last component of  $\tilde{\mathbf{y}}_{p-1,IN}^{pN_{SDC}^G}$  from processor  $p - 1$ ;
7     Substitute it to first component of  $\tilde{\mathbf{y}}_{p,IN}^{pN_{SDC}^G}$ ;
8   Apply  $N_{SDC}^G$  additional SDC sweeps to recover  $\tilde{\mathbf{y}}_{p,IN}^{(p+1)N_{SDC}^G}$ ;
9   if  $p < N_p - 1$  then
10    Send last component of  $\tilde{\mathbf{y}}_{p,IN}^{(p+1)N_{SDC}^G}$  to processor  $p + 1$ ;
11  Interpolate correction at fine level to recover  $\tilde{\mathbf{u}}_p^{0,0}$ ; // Formula (2.14)

/* Main PFASST iterations: */
12 parfor  $k = 0$  to  $It_{\max}$  do
    /* Parallel step: */
13  for  $p = 0$  to  $N_p - 1$  do
14    Apply  $N_{SDC}^F$  fine SDC sweeps to recover  $\tilde{\mathbf{u}}_p^{k,N_{SDC}^F}$ ;
15  if  $p > 0$  then
16    Receive the last component of  $\tilde{\mathbf{u}}_{p-1}^{k+1,0}$  from processor  $p - 1$ ;
17    Substitute it to first component of  $\tilde{\mathbf{u}}_p^{k,N_{SDC}^F}$ ;
18  Restrict fine solution to coarse and recover  $\tilde{\mathbf{y}}_p^{k,0}$ ; // Formula (2.13)
19  Apply  $N_{SDC}^G$  coarse SDC sweeps to recover  $\tilde{\mathbf{y}}_p^{k,N_{SDC}^G}$ ;
20  Interpolate correction at fine level to recover  $\tilde{\mathbf{u}}_p^{k+1,0}$ ; // Formula (2.14)
21  if  $p < N_p - 1$  then
22    Send last component of  $\tilde{\mathbf{u}}_p^{k+1,0}$  to processor  $p + 1$ ;

```

For example, rather than interpolating the *whole* coarse correction to the fine level before sending the final conditions to the following processor, we can focus on recovering the sole final component (in time) of the fine solution, propagating that one forward, and only afterwards completing the interpolation of the remaining components, as presented in the original paper [43]. Furthermore, there is no need to limit the propagation of information to the fine grid only: each level can dialogue with the corresponding level of the following processor, sending updated information forward as soon as it becomes available [10, 142]. While this increases the amount of communication between processors, it might help reducing computational time, depending on the architecture of the system. Apart from these small but relevant adjustments, however, the core of the algorithm remains substantially unchanged.

Similarly to Parareal, also PFASST has proven to be effective in applications to parabolic PDEs: for example, in [104] a speedup of 6-9× with 24 cores is reported, while 8× is achieved in [43] with 16 cores. Unfortunately, though, not even PFASST is immune to instabilities in applications to transport-dominated problems, and often incurs in similar issues as those presented in Sec. 2.1.1.1. However, a recent result on linear convergence [11], (owing to an interpretation of the algorithm in a multi-grid framework [10]), indicates that opportunely choosing certain parameters in the algorithm can help dampen these undesirable effects.

2.1.3 ParaExp

The *Parallel Exponential Propagation* (ParaExp) scheme represents another efficient alternative for parallel-in-time integration of linear differential equations. Introduced in [56] by Gander and Güttel, the algorithm has proven to be particularly effective on systems featuring a forcing term which is challenging to integrate numerically. Based on the principle of superposition of effects, it splits the system into an autonomous and a non-autonomous part, and employs exponential integration in order to rapidly solve the former. Further details are presented next, following the description originally given in [56].

The target equation for the ParaExp scheme is similar to (2.1), but the right-hand side is simplified to

$$f(u(t), t) = au(t) + g(t), \quad (2.15)$$

where a is a scalar³. Within this framework, the function $g(t)$ represents the main difficulty towards integration: this could be because it is rapidly oscillating, render-

³The description still holds in the multi-dimensional case, substituting a by a linear operator \mathcal{A} .

ing the system stiff, or simply because it is particularly expensive to evaluate. In complete analogy to the other time-parallel algorithms presented in Sec. 2.1.1 and in Sec. 2.1.2, the ParaExp algorithm achieves a speedup by dividing the time domain of the differential equation into chunks, each assigned to a different processor. Each of the processors solves then two types of problem. The first one is defined as:

$$\begin{cases} v_p'(t) = av_p(t) + g(t), & t \in (T_p, T_{p+1}] \\ v_p(T_p) = \delta_{p,0} \bar{u}^0 \end{cases} \quad p = 0, \dots, N_p - 1, \quad (2.16)$$

where the Kronecker delta $\delta_{p,0}$ indicates that homogeneous initial conditions are picked for each subdomain, except the very first one, where the only known initial condition for the solution, $u(T_0) = \bar{u}^0$, is imposed. Due to the presence of $g(t)$, integrating these problems is expected to be time consuming; nonetheless, all of (2.16) are independent of each other, which hints at parallelisation as a valuable option to save computational time. The second problem is given by:

$$\begin{cases} w_p'(t) = aw_p(t), & t \in (T_p, T] \\ w_p(T_p) = v_{p-1}(T_p) \end{cases} \quad p = 1, \dots, N_p - 1. \quad (2.17)$$

Similarly, all of (2.17) are also independent. However, in order to be solved, they require $v_{p-1}(T_p)$, *i.e.* the result from the integration of the corresponding sub-problem (2.16). Moreover, notice that this integration does not stop at the end of a time chunk, but continues until the end of the time domain T , which causes an imbalance in the amount of computation that each processor needs to perform. This issue can be addressed by adjusting the sizes of each time chunk, if an estimate of the amount of computations necessary to solve both (2.16) and (2.17) is provided. Still, this might be seen as a large overhead with respect to the serial solution; however, each of the Cauchy problems in (2.17) is prone to be solved efficiently using exponential integration, as described in [76].

Finally, once the problems (2.16) and (2.17) are solved, the complete solution at a given instant $t \in [T_p, T_{p+1}]$ is promptly recovered by superposition of effects:

$$u(t) = v_p(t) + \sum_{i=1}^p w_i(t). \quad (2.18)$$

An illustration of the actions performed by each processor is sketched in Fig. 2.3, while the pseudo-code for the algorithm is provided in Alg. 3.




Figure 2.3: Schematic of the of the action of the ParaExp algorithm. Each processor p starts solving (2.16) on the corresponding time chunk, and recovers v_p (dark blue lines). This provides the initial condition (light blue dots), to proceed integrating (2.17) all the way until the end of the time domain (light blue dotted lines). Once $w_{p+1}(t)$ is also recovered, the complete solution can be attained simply summing all contributions.

2.1.3.1 Extension to nonlinear equations

If the linearity assumption breaks down, in general the principle of superposition of effects does not hold any more, and the procedure outlined in Alg. 3 fails to provide the correct solution to the target equation. However, one can hope to achieve convergence to the actual solution by iterating. In [57], this process is described and analysed in

Algorithm 3: Paraexp

```

/* Initialisation: */
1 Set  $v_0(T_0) = u_0$ ;
2 parfor  $p = 1$  to  $N_p - 1$  do
3   | Set  $v_p(T_p) = 0$ ;

/* Parallel steps: */
4 parfor  $p = 0$  to  $N_p - 1$  do
5   | Integrate  $v_p(t)$  in  $[T_p, T_{p+1}]$ ; // Problem (2.16)
6   | if  $p < N_p - 1$  then
7     | | /* Fast exponential integration: */
8     | | Set  $w_{p+1}(T_{p+1}) = v_p(T_{p+1})$ ;
9     | | Integrate  $w_{p+1}(t)$  in  $[T_{p+1}, T]$ ; // Problem (2.17)
9 Recover the complete solution by superposition; // Formula (2.18)

```

detail, giving rise to a ParaExp algorithm for nonlinear equations.

A different approach is instead pursued in [73], where the *Asymptotic Parareal* algorithm is introduced. This lies somewhat at the interface between the ParaExp and Parareal algorithms. It is intended for applications where the differential operator involved can be split into a linear and nonlinear part,

$$\frac{\partial u(t)}{\partial t} = \frac{1}{\varepsilon} \mathcal{L}u(t) + \mathcal{N}(u(t)), \quad (2.19)$$

whose linear component \mathcal{L} encapsulates the effect of a transport phenomenon, while the nonlinear one \mathcal{N} corresponds to diffusion. Moreover, it is assumed that advection occurs over a much smaller timescale (of order $\varepsilon \ll 1$) than diffusion. The main idea proposed in [73] is to employ a framework similar to that of Parareal, (alternating the use of a parallel fine integrator to recover an accurate solution and a serial coarse integrator to run an update), but one in which the coarse solver is chosen in such a way to exploit the separation of timescales. In particular, this consists of an exponential integrator that tracks exactly the evolution imposed by \mathcal{L} , while taking an appropriately averaged version of \mathcal{N} . It is possibly thanks to the very use of exponential integration that this method manages to circumvent many of the shortcomings of Parareal illustrated in Sec. 2.1.1.1: in fact, it allows to safely use a relatively large time-step for the coarse solver while still preserving an adequate level of accuracy, thus manufacturing a considerably more efficient algorithm: results in [73] show a speedup of $10\times$ with respect to vanilla Parareal.

2.2 All-at-once methods

The schemes introduced in the previous sections employ an already existing algorithm for the integration of a time-dependent differential equation (be it a generic time-stepping routine for Parareal, SDC for PFASST, or exponential integration for ParaExp), which is left substantially unchanged, while focusing on implementing an overarching framework which allows for the introduction of concurrency in the recovery of the solution along the temporal domain. An alternative approach to time parallelisation consists of considering the monolithic system arising from the simultaneous space-time discretisation of a time-dependent partial differential equation, and developing techniques which tackle its solution concurrently as a whole, rather than proceeding instant-by-instant. In light of this, the class of algorithms following this strategy is often referred to as *all-at-once methods* (AAO)⁴.

⁴We point out, however, that this term has been used with slightly different meanings throughout the literature.

All the techniques developed and studied in this thesis belong to this second class of PinT methods. In this section, we introduce the main ideas behind this approach, and we briefly illustrate a choice of AAO algorithms. As with the previous section, we introduce a model problem to help us in their description, and consider a generic time-dependent, linear, homogeneous PDE with constant coefficients, discretised in time using *Implicit Euler*. Temporal integration is performed by solving the following recurrence formula:

$$\frac{\mathbf{u}^k - \mathbf{u}^{k-1}}{\Delta t_k} + \mathcal{L}\mathbf{u}^k = \mathbf{0}, \quad k = 1, \dots, N_t, \quad (2.20)$$

where the instants t_k identify the temporal grid $\{t_k\}_{k=0}^{N_t}$, with $\Delta t_k = t_k - t_{k-1}$, and starting from a given initial condition $\mathbf{u}^0 = \bar{\mathbf{u}}^0$. In (2.20), the term $\mathcal{L} \in \mathbb{R}^{N_x \times N_x}$ represents a discretisation of a spatial operator, where N_x is the number of unknowns in the spatial discretisation. The set of equations (2.20) can be collected into a single discrete space-time system:

$$\underbrace{\begin{bmatrix} \frac{I}{\Delta t_1} + \mathcal{L} & & & & \\ -\frac{I}{\Delta t_2} & \frac{I}{\Delta t_2} + \mathcal{L} & & & \\ & \ddots & \ddots & & \\ & & & -\frac{I}{\Delta t_{N_t}} & \frac{I}{\Delta t_{N_t}} + \mathcal{L} \end{bmatrix}}_{=:A} \underbrace{\begin{bmatrix} \mathbf{u}^1 \\ \mathbf{u}^2 \\ \vdots \\ \mathbf{u}^{N_t} \end{bmatrix}}_{=:u} = \underbrace{\begin{bmatrix} \frac{\bar{u}^0}{\Delta t_1} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix}}_{=:g}. \quad (2.21)$$

All-at-once methods attempt to solve the whole system (2.21) at once, either in a direct or an iterative fashion, and with the goal of exposing concurrency in its solution. Some examples of how this can be achieved are described in the following section.

Remark 1. We expect that the type of temporal discretisation chosen will play a major role in defining the final structure of the resulting space-time system. Particularly, (2.21) presents the distinctive block lower-triangular structure stemming from employing classical time-stepping schemes⁵. In a sense, time-stepping algorithms *must* produce such a structure, since they recover the solution at a given instant by using information only coming from the previous instants. Classical time-stepping can then be interpreted as solving (2.21) via block forward substitution, making use of the fact that the space-time matrix stemming from this temporal discretisation retains a block lower-triangular form.

⁵More in detail, the number of steps employed in the time-stepping routine determines the number of block sub-diagonals in the monolithic system. Since we use a *one*-step method for our example, the resulting space-time matrix contains only *one* block sub-diagonal. More complex examples are treated in Chap. 6.

2.2.1 ParaDiag

First introduced in [97], with its properties being analysed in [59], the ParaDiag algorithm aims at recovering the solution of the space-time system (2.21) via block diagonalisation. As a consequence, each of the blocks composing the resulting block diagonal matrix represents an independent system, whose solution can be trivially parallelised. Following [97], this procedure can be better illustrated by considering the tensor product representation of (2.21). To this purpose, introducing the discrete time-differentiation operator

$$\Phi := \begin{bmatrix} \frac{1}{\Delta t_1} & & & & \\ -\frac{1}{\Delta t_2} & \frac{1}{\Delta t_2} & & & \\ & \ddots & \ddots & & \\ & & & -\frac{1}{\Delta t_{N_t}} & \frac{1}{\Delta t_{N_t}} \end{bmatrix} \in \mathbb{R}^{N_t \times N_t}, \quad (2.22)$$

we can express the space-time system (2.21) as

$$A = \Phi \otimes I_{N_x} + I_{N_t} \otimes \mathcal{L}, \quad (2.23)$$

where \otimes represents the Kronecker product, while I_{N_t} and I_{N_x} are identity matrices of size the number of temporal (N_t) and spatial (N_x) nodes, respectively. Given how Φ in this case is lower triangular, its eigenvalues are given by the elements in its diagonal, $\lambda_k(\Phi) = 1/\Delta t_k$; the matrix is hence diagonalisable if and only if all the time-steps δt_k are different from one another:

$$\Delta t_i \neq \Delta t_j \quad \forall i, j \in 1, \dots, N_t \iff \Phi = U \Lambda_\Phi U^{-1}, \quad \text{with } \Lambda_\Phi := \text{diag}(\lambda_k(\Phi)),_{k=1, \dots, N_t} \quad (2.24)$$

where $U \in \mathbb{R}^{N_t \times N_t}$ denotes the matrix of eigenvectors of Φ . Under this assumption, the eigendecomposition of (2.22) can be recovered explicitly; moreover, U presents a lower-triangular structure [59, Thm. 3], making it relatively simple to invert via backward substitution. Enforcing diagonalisability, we can finally rewrite the space-time matrix A as

$$\begin{aligned} A &= (U \Lambda_\Phi U^{-1}) \otimes I_{N_x} + U U^{-1} \otimes \mathcal{L} \\ &= (U \otimes I_{N_x})(\Lambda_\Phi \otimes I_{N_x})(U^{-1} \otimes I_{N_x}) + (U \otimes I_{N_x})(I_{N_t} \otimes \mathcal{L})(U^{-1} \otimes I_{N_x}) \\ &= (U \otimes I_{N_x}) \underbrace{(\Lambda_\Phi \otimes I_{N_x} + I_{N_t} \otimes \mathcal{L})}_{=\text{diag}_{k=1, \dots, N_t}(\lambda_k(\Phi)I_{N_x} + \mathcal{L})} (U^{-1} \otimes I_{N_x}), \end{aligned} \quad (2.25)$$

where we used the property of the Kronecker product $(ac) \otimes (bd) = (a \otimes b)(c \otimes d)$. Looking at the central factor in the formula above, we can see that the solution of

(2.21) reduces to the inversion of N_t independent spatial systems involving operators $\lambda_k(\Phi)I_{N_x} + \mathcal{L}$, which is what gives concurrency to this method. This procedure has been extended to applications with time-dependent parameters (*i.e.*, where the main operator \mathcal{L} itself might vary with time), as well as to the nonlinear setting [58]. Interestingly, this method can be successfully applied also to the solution of wave equations: some results on the parallel capabilities of this method are reported in [60], where it is often reported a parallel efficiency of up to $\sim 90\%$.

Having to request irregular time-steps, as per (2.24), can be beneficial in specific frameworks⁶, but is somewhat limiting in general, especially when considering that the condition number of the eigenvector matrix U is directly affected by the distribution of temporal nodes. However, work is under way in order to overcome this difficulty [60]. An alternative approach when uniform time-steps are employed instead is described in the following.

2.2.2 Space-time circulant preconditioning

Rather than using a direct solver to tackle (2.21), as done in Sec. 2.2.1, the method introduced in this section relies on Krylov-subspace iterative schemes [126], opportunely preconditioned with an operator whose application can be parallelised in time. To describe how this is achieved, consider again the operator in (2.21): if a uniform grid is employed for the temporal discretisation, this simplifies to

$$A \left(\begin{array}{c} = \\ \text{if } \Delta t_k \equiv \Delta t \\ \forall k = 1, \dots, N_t \end{array} \right) \left[\begin{array}{ccc} \frac{I}{\Delta t} + \mathcal{L} & & \\ -\frac{I}{\Delta t} & \frac{I}{\Delta t} + \mathcal{L} & \\ & \ddots & \ddots \end{array} \right]. \quad (2.26)$$

We notice that its structure further specialises: its blocks are constant down each diagonal, rendering A a block Toeplitz matrix. The authors in [101] suggest then to use a block circulant preconditioner in order to accelerate the convergence of iterative methods applied to the whole (2.26).

The idea of using circulant matrices as preconditioners for Toeplitz systems dates back to 1986 [145], and has proven to be remarkably effective: see for example [17, Chap. 2.5], or [108]. Circulant matrices can be seen as “wrapped around” Toeplitz matrices, where elements on a diagonal are also repeated N diagonals away (with N

⁶One such example is the classical heat equation: when tending to steady state, the rate of change of the solution decreases with time, so it makes sense to coarsen the temporal grid as time progresses.

being the size of the matrix), in the form

$$C = \begin{bmatrix} c_0 & c_{N-1} & \cdots & c_2 & c_1 \\ c_1 & c_0 & c_{N-1} & & c_2 \\ \vdots & c_1 & \ddots & \ddots & \vdots \\ c_{N-2} & & \ddots & & c_{N-1} \\ c_{N-1} & c_{N-2} & \cdots & c_1 & c_0 \end{bmatrix}. \quad (2.27)$$

A particularly favourable property of circulant matrices, is that they can be diagonalised via Fourier transform: introducing the *discrete Fourier transform* operator,

$$[F_N]_{m,n} := \frac{1}{\sqrt{N}} e^{-\frac{2\pi i}{N} mn}, \quad m, n = 0, \dots, N-1, \quad (2.28)$$

we have that

$$C = F_N \Lambda_C F_N^*, \quad \text{with} \quad \Lambda_C := \text{diag}_{k=0, \dots, N-1} ([F_N[C]_{-,0}]_k), \quad (2.29)$$

where Λ_C is the diagonal matrix containing the Fourier transform of the first column of C on its diagonal (which also identifies the eigenvalues of C). Systems involving C can then be effectively solved via *Fast Fourier Transform* (FFT).

This framework can be seamlessly extended to block circulant matrices, and applied to accelerate the solution of block Toeplitz systems such as (2.26): the preconditioner then assumes a form similar to that of (2.25), and its solution can be once again trivially parallelised in time⁷. We refer to Chap. 6, where we extensively make use of this method, for a more in-depth description. On top of this, the effectiveness of circulant preconditioning applied to space-time systems has been studied in detail in [101] and [65], and more recently in [93]. There, strong and weak scaling properties of the algorithm are discussed. In particular, the wall-clock time required by the method for some test-cases remains substantially constant as the number of processors used increases to up to 64, providing evidence for the optimal scalability of the algorithm.

Finally, we point out that, for constant-coefficient, linear problems which prescribe periodicity in time, the resulting space-time matrix is often already block circulant, namely

$$A_C = \begin{bmatrix} \frac{I}{\Delta t} + \mathcal{L} & & -\frac{I}{\Delta t} \\ -\frac{I}{\Delta t} & \frac{I}{\Delta t} + \mathcal{L} & \\ & \ddots & \ddots \end{bmatrix}, \quad (2.30)$$

⁷Indeed the approaches illustrated in Sec. 2.2.1 and 2.2.2 present some similarities, to the point that circulant preconditioning for AAO systems is described as a *variant* of ParaDiag in [60]. Given how the focus of the latter is more on the *preconditioning* aspect, however, we believe it is worth it to draw a distinction.

and the considerations on its diagonalisability can be applied directly to the solution of A : this has been exploited in [87] to design an efficient time-parallel solver for time-periodic problems, reporting a potential speedup of $500\times$ when using 50 processors.

2.2.3 MGRIT

The *Multigrid Reduction in Time* (MGRIT) algorithm also aims at parallelising the solution of the monolithic system (2.21), but, as suggested by its name, it follows a *Multigrid Reduction* approach (MGR, [51, 120]), with the particularity that its coarsening procedure acts along the temporal domain. First proposed in [47], MGRIT is not the only multigrid-inspired algorithm for time parallelisation (see for example [48, 61, 77]), but is one of the most mature and widely studied. To introduce the idea behind this algorithm, we follow much of the original description in [47].

In accordance with the MGR approach, the N_t nodes corresponding to the temporal discretisation are split into two sets: C , consisting of *coarse* nodes, and F , consisting of *fine* ones. Both the components of the monolithic space-time discretisation \mathbf{u} and the coefficients of A in (2.21), are rearranged accordingly, so that the matrix system can be factorised in the following way:

$$A = \begin{bmatrix} A_{FF} & A_{FC} \\ A_{CF} & A_{CC} \end{bmatrix} = \begin{bmatrix} I_{N_F} & 0 \\ A_{CF}A_{FF}^{-1} & I_{N_C} \end{bmatrix} \begin{bmatrix} A_{FF} & 0 \\ 0 & S \end{bmatrix} \begin{bmatrix} I_{N_F} & A_{FF}^{-1}A_{FC} \\ 0 & I_{N_C} \end{bmatrix} \quad (2.31)$$

provided A_{FF} is invertible. Here, $I_{(*)}$ denotes the identity matrix of size $(*)$, with N_C and N_F identifying the number of coarse and fine nodes, respectively, so that $N_C + N_F = N_t$. The block S in the block diagonal matrix represents the coarse-nodes Schur complement,

$$S := A_{CC} - A_{CF}A_{FF}^{-1}A_{FC}, \quad (2.32)$$

i.e., the ideal Petrov-Galerkin coarse grid operator [12, Chap. 7.7.3]. This factorisation allows us to separate the solution of the system over the F and C nodes: we have in fact

$$A^{-1} = \underbrace{\begin{bmatrix} -A_{FF}^{-1}A_{FC} \\ I_{N_C} \end{bmatrix}}_{=:I_F^C} S^{-1} \underbrace{\begin{bmatrix} A_{CF}A_{FF}^{-1} & I_{N_C} \end{bmatrix}}_{=:R_C^F} + \underbrace{\begin{bmatrix} A_{FF}^{-1} & 0 \\ 0 & 0 \end{bmatrix}}_{=:A_{FF,0}^{-1}}, \quad (2.33)$$

where we introduced the *ideal* restriction/interpolation operators R_C^F/I_F^C between the coarse and fine level, for which $S = R_C^F A I_F^C$ holds. To conclude the MGR description,

we use (2.33) to more conveniently express the error propagation operator associated with an *exact* 2-level multigrid method:

$$E := I - A^{-1}A = \underbrace{(I - I_F^C S^{-1} R_C^F A)}_{=: E_C} \underbrace{(I - A_{FF,0}^{-1} A)}_{=: E_F}, \quad (2.34)$$

where we highlighted the factors responsible for error relaxation at fine and coarse level, E_F and E_C , respectively. This suggests that an iterative multigrid method can be built by substituting the components in (2.34), (and specifically the restriction/interpolation operators and the Schur complement, R_C^F , I_F^C and S), with approximate ones (\tilde{R}_C^F , \tilde{I}_F^C and \tilde{S}).

Interestingly, Parareal itself can be interpreted as a special case of MGRIT, as shown in [47]. In fact, by choosing the injection operator $\tilde{R}_C^F := \begin{bmatrix} 0 & I_F^C \end{bmatrix}$ instead of ideal restriction, and by substituting the Schur complement S with a space-time operator \tilde{A} which represents time-stepping over the coarse nodes only, it can be proven that Parareal corresponds to a 2-level MGR scheme which applies a semi-coarsening procedure along the temporal domain only. Its associated error propagator is given by

$$E_{Parareal} := \left(I - I_F^C \tilde{A}^{-1} \tilde{R}_C^F A \right) \left(I - A_{FF,0}^{-1} A \right). \quad (2.35)$$

This way MGRIT provides the great advantage of establishing a framework for interpreting Parareal as a MGR method, thus allowing the possibility to extend the algorithm using concepts borrowed from the multigrid technology⁸. Particularly, this framework enables the use of a whole hierarchy of coarse grids (in a true *multi-grid* spirit), as well as different approaches to error relaxation. For example, in [47] it is recommended to substitute the F-relaxation factor E_F in (2.34) with the FCF-relaxation term

$$E_{FCF} := \left(I - A_{FF,0}^{-1} A \right) \left(I - I_F^C S^{-1} R_C^F A \right) \left(I - A_{FF,0}^{-1} A \right). \quad (2.36)$$

To explore these details further, we refer to Chap. 5, where we consider applications of MGRIT, and particularly to Sec. 5.2, for a more in-depth discussion on the actual choice of its approximated components \tilde{R}_C^F , \tilde{I}_F^C and \tilde{S} .

To get an idea of the performance of MGRIT, an example on PDE-constrained optimisation reports a $19\times$ speedup for 256 cores [69]; a more advanced application in fluid-structure interaction [74] shows a speedup of $18\times$ with 32 cores.

⁸A different multigrid interpretation of Parareal is given in [62], and does share some similarities with the approach presented in MGRIT, albeit the latter uses more standard operators in its description.

The characteristic of applying coarsening along the time domain only, specific of the original MGRIT algorithm, has the advantage of producing a non-intrusive, *black-box* solver which often squeezes additional (temporal) parallelisation out of an already-developed time-marching scheme, provided the latter allows for some flexibility in the time-step chosen: this makes it closer in spirit to Parareal and the other algorithms in Sec. 2.1. Indeed, thanks to their similarity, one can exploit some of the convergence results derived for Parareal also in the MGRIT framework (and vice-versa: theory developed for MGRIT can be applied to Parareal as well), see for example [35, 139, 140]. Regrettably, this also suggests that some of the *limitations* of Parareal, and specifically those described in Sec. 2.1.1.1 in applications to hyperbolic systems, carry over to MGRIT as well. We address some of these issues in Chap. 5, and an attempt to overcome them is presented in [33]. In practice, however, once the connection with MGR has been established, one has the freedom of applying its machinery to produce coarsening strategies involving both spatial and temporal domains at once. The advantages this ensues are highlighted for instance in [78]. An algorithm which heavily relies on this feature is AIR, discussed in the next section.

2.2.4 AIR

Rooted on the same MGR concept as MGRIT, the *Approximation to Ideal Restriction* (AIR) technology was originally developed in an attempt to design an effective method for applying the multigrid framework to nonsymmetric systems, and quickly established itself as an effective PinT algorithm.

When applying multigrid to elliptic operators, the focus generally lies on the interpolation operator: this must be designed so that it can effectively map the correction from the coarse grid into the range of the fine-level relaxation (in other words, \tilde{I}_F^C should be chosen so that it does not perturb the residual at the coarse nodes after coarse grid correction, since F-relaxation can usually take care of dampening the residual on the fine nodes); once this is achieved, if the target operator of the multigrid routine is self-adjoint, then restriction is usually taken as the adjoint of interpolation: $\tilde{R}_C^F = (\tilde{I}_F^C)^T$. On the other hand, with AIR the focus shifts to the restriction operator (as the name of the algorithm suggests). The role of restriction is complementary to that of interpolation, in that it should help defining a coarse-grid correction procedure which dampens error components falling outside the range of \tilde{I}_F^C . In [99] the case is made that for nonsymmetric systems, choosing \tilde{R}_C^F in a way that it covers this role effectively (thus approximating the action of R_C^F in (2.33)), is particularly relevant in order to ensure convergence.

To achieve this, an MGR method is introduced in [99] for (block) lower-triangular matrices, as they might arise from the discretisation of advection equations, which relies on approximating the action of R_C^F , (and particularly the A_{FF}^{-1} term appearing in its C block (2.33)), via *Neumann series*: this represents the n AIR version of this algorithm. Unfortunately, the effectiveness of the approximation heavily relies on the lower-triangular structure of the target matrix, so that if this is disrupted (for example by introducing a diffusive term), the performance of the method degrades quickly. To counteract this, an alternative approach based on a *local* approximation of the effect of R_C^F (aptly dubbed *l*AIR) has been proposed in [100], which shows much more robustness when applied to the whole spectrum from advection- to diffusion-dominated problems.

Remarkably, AIR can be successfully applied to the solution of space-time systems, once we realise that a d -dimensional time-dependent advection(-diffusion) equation can be interpreted as a $d + 1$ dimensional one, where time is yet another direction for advection. Moreover, the AIR procedure has the additional advantage of not relying on the target space-time system having the typical structure outlined in (2.21), which stems from a time-stepping scheme. As such, it can be successfully applied also to alternative approaches to temporal discretisation, such as space-time Finite Elements. This is showcased in [138], which also demonstrates the optimal parallel efficiency of the method. We make use of this algorithm in Chap. 3.

In the continuation of this thesis, we decide to focus on algorithms belonging to the AAO class described in Sec. 2.2. This preference is guided both by the potential they have shown over most algorithms in Sec. 2.1 when dealing with hyperbolic equations, and by the fact that they introduce a flexible framework where to operate our analysis, namely that of multigrid methods and/or preconditioning. Indeed, the technique developed in the following chapter takes inspiration precisely from block preconditioning strategies, which we extend to the whole space-time setting.

Chapter 3

Space-time block preconditioning

As mentioned in Sec. 2.2.3 and 2.2.4, the application of multigrid technologies to the simultaneous solution of space-time systems has been a guiding principle in designing successful parallel-in-time techniques [48,61,77], often providing significantly more parallelism and faster time-to-solution than algorithms listed in Sec. 2.1, not least because this approach naturally allows for adaptive mesh refinement in both space and time [138]. To our knowledge, however, most of the development of space-time MG techniques has been limited to applications involving single-variable (advection-)diffusion equations, and no multigrid methods have been developed and proven effective on space-time systems of PDEs. Indeed, effectively handling simultaneously both the complications introduced by inter-variable coupling, as well as the purely advective term in one dimension representing a time derivative, is beyond the capabilities of the current state-of-the-art.

In this chapter, we take a novel approach and introduce the concept of *space-time block preconditioning* (STBP) for the all-at-once solution of block-structured systems of PDEs. The main idea behind this method relies on considering well-developed spatial block preconditioning techniques, already used in accelerating the solution of sequential time-stepping or steady-state problems, and extending their application to the whole space-time setting. The approach described in [113,144], where the authors propose a preconditioner for a time-dependent optimal control problem, is similar in spirit to our work, as it exploits the block structure of the system at the space-time level. However, to our knowledge, similar concepts have not been applied before for the purpose of time parallelisation.

To better deliver and crystallise our description, we focus on the Oseen equations for illustrating the guiding principle behind the assembly of our space-time block preconditioner. This target problem is chosen both for the relevance of developing efficient time-parallelisation techniques for applications to fluid dynamics, due to the

typically high cost associated with their solution (see for example [152] for a discussion on the need for time parallelism in this field for emerging architectures), and because it serves as a building block for the more complex magnetohydrodynamics problem analysed in Chap. 4. Indeed, we believe that the underlying principle of space-time block preconditioning can be extended to other systems of PDEs for which effective spatial block preconditioners have been developed for the single time-step case (or for its stationary counterpart).

The block separation introduced by our preconditioner allows for great freedom in the choice of the solver for the velocity block, making it straightforward to apply, *e.g.*, space-time multigrid for single-variable equations as in [48, 77, 138], thus overcoming the issue described in the first paragraph. The proposed method builds on the *Pressure Convection-Diffusion* (PCD) preconditioner described in [40], and we show that similar principles to the ones that justified the validity of PCD can be generalised to the space-time setting, resulting in a space-time block preconditioner whose application is highly parallelisable in time.

The remainder of this chapter is organised as follows. Section 3.1 introduces our target problem, together with notation that will be used throughout. A derivation of the space-time block preconditioner representing our main contribution is given in Sec. 3.2, together with multiple theoretical justifications for its formulation. In Sec. 3.4 we investigate the effectiveness of the proposed preconditioner, testing its performance on four classical models in incompressible flow. Results indicate perfect scalability in refinement of spatial and temporal mesh spacing, perfect scalability in nonlinear Picard iteration count when applied to a nonlinear Navier-Stokes problem, and minimal overhead in terms of number of preconditioner applications compared with sequential time-stepping.

The work in this chapter stems from a collaboration with Ben S. Southworth: the findings have been collected into a manuscript, which at the time of writing is under review for publication in the *SIAM Journal on Scientific Computing*. A pre-print version of this work can be found in [31].

3.1 Problem definition

This section introduces notation that is used throughout this chapter. Here we describe a time-dependent version of the Oseen equations, as well as the discrete operators arising from its finite element discretisation.

Our target system is

$$\left\{ \begin{array}{l} \frac{\partial \mathbf{u}}{\partial t}(\mathbf{x}, t) + (\mathbf{w}(\mathbf{x}, t) \cdot \nabla) \mathbf{u}(\mathbf{x}, t) \\ -\mu \nabla^2 \mathbf{u}(\mathbf{x}, t) + \nabla p(\mathbf{x}, t) = \mathbf{f}(\mathbf{x}, t) \\ \nabla \cdot \mathbf{u}(\mathbf{x}, t) = 0 \end{array} \right. \quad \text{in } \Omega \times (T_0, T], \quad (3.1)$$

defined on a spatial domain $\Omega \subset \mathbb{R}^d$. Its unknowns are the d -dimensional vector field $\mathbf{u}(\mathbf{x}, t)$, and the scalar field $p(\mathbf{x}, t)$, representing velocity and pressure, respectively; μ is the viscosity coefficient, which we consider constant on Ω . The vector function $\mathbf{f}(\mathbf{x}, t)$ is a given forcing term, while $\mathbf{w}(\mathbf{x}, t)$ represents a given d -dimensional advection field. The special case where $\mathbf{w}(\mathbf{x}, t) \equiv 0$ gives rise to a time-dependent version of Stokes equations:

$$\left\{ \begin{array}{l} \frac{\partial \mathbf{u}}{\partial t}(\mathbf{x}, t) - \mu \nabla^2 \mathbf{u}(\mathbf{x}, t) + \nabla p(\mathbf{x}, t) = \mathbf{f}(\mathbf{x}, t) \\ \nabla \cdot \mathbf{u}(\mathbf{x}, t) = 0 \end{array} \right. \quad \text{in } \Omega \times (T_0, T]; \quad (3.2)$$

While most of the analysis conducted in this chapter considers the more general problem (3.1), we will at times refer to (3.2) for some simplifications. Conversely, by introducing the nonlinearity $\mathbf{w}(\mathbf{x}, t) \equiv \mathbf{u}(\mathbf{x}, t)$ we produce the Navier-Stokes equations for incompressible flow

$$\left\{ \begin{array}{l} \frac{\partial \mathbf{u}}{\partial t}(\mathbf{x}, t) + (\mathbf{u}(\mathbf{x}, t) \cdot \nabla) \mathbf{u}(\mathbf{x}, t) \\ -\mu \nabla^2 \mathbf{u}(\mathbf{x}, t) + \nabla p(\mathbf{x}, t) = \mathbf{f}(\mathbf{x}, t) \\ \nabla \cdot \mathbf{u}(\mathbf{x}, t) = 0 \end{array} \right. \quad \text{in } \Omega \times (T_0, T]. \quad (3.3)$$

The Oseen system can then be interpreted as a linearisation of Navier-Stokes, implying that the analysis conducted here can be extended to more general frameworks: this is addressed more in detail in Sec. 3.4.3.

In order to provide closure to the definition of the initial-boundary value problem, we also equip (3.1) with an initial condition

$$\mathbf{u}(\mathbf{x}, T_0) = \bar{\mathbf{u}}^0(\mathbf{x}) \quad \text{on } \Omega, \quad (3.4)$$

and the boundary conditions¹ (BC)

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{u}_D(\mathbf{s}, t) \quad \text{on } \Gamma_D^{\mathbf{u}} \times (T_0, T], \quad \text{and} \quad (3.5a)$$

$$\mu \nabla \mathbf{u}(\mathbf{s}, t) \mathbf{n}(\mathbf{s}) - p(\mathbf{s}, t) \mathbf{n}(\mathbf{s}) = \mathbf{g}(\mathbf{s}, t) \quad \text{on } \Gamma_N^{\mathbf{u}} \times (T_0, T], \quad (3.5b)$$

¹Choosing the Neumann boundary conditions in the form (3.5b) ignores the contribution from the $\mu(\nabla \mathbf{u})^T$ term in the fluid stress tensor. This simplification is dangerous, since it results in a violation of the *objectivity principle* of continuum mechanics [91]. However, this choice of BC remains unfortunately common when dealing with the Laplace form of the unsteady Stokes equations (3.2), as it results in a system where the velocity components are decoupled (and indeed it has been used in [40] as well).

where we split the boundary $\partial\Omega$ into its Dirichlet and Neumann parts, respectively, with $\partial\Omega = \Gamma_D^{\mathbf{u}} \cup \Gamma_N^{\mathbf{u}}$ and $\Gamma_D^{\mathbf{u}} \cap \Gamma_N^{\mathbf{u}} = \emptyset$. In (3.5), $\mathbf{n}(\mathbf{s}) \in \mathbb{R}^d$ represents the outward-pointing normal to the boundary, while $\mathbf{u}_D(\mathbf{s}, t)$, $\mathbf{g}(\mathbf{s}, t)$, and $\bar{\mathbf{u}}^0(\mathbf{x})$ are some known functions.

At any given instant $t \in [T_0, T]$, the velocity unknown lives in the functional space [118, Chap. 2.4]

$$\mathbf{u}(\cdot, t) \in V(\Omega) \equiv \left(H_{\Gamma_D^{\mathbf{u}}}^1(\Omega) \right)^d, \quad (3.6)$$

while for the pressure we have

$$p(\cdot, t) \in Q(\Omega) \equiv L^2(\Omega), \quad \text{or} \quad Q(\Omega) \equiv L_0^2(\Omega) = \left\{ q(\mathbf{x}) \in L^2(\Omega) : \int_{\Omega} q(\mathbf{x}) d\mathbf{x} = 0 \right\}, \quad (3.7)$$

where the zero-mean valued counterpart, $L_0^2(\Omega)$, is chosen if Dirichlet boundary conditions are imposed everywhere on the boundary of the domain, $\Gamma_D \equiv \partial\Omega$. In either case, having introduced eqs. (3.6) and (3.7), we can provide the weak formulation [118, Chap. 17.2] of (3.1): we seek $\mathbf{u}(\mathbf{x}, t)$ and $p(\mathbf{x}, t)$ such that

$$\begin{cases} \frac{\partial}{\partial t} \langle \mathbf{u}, \mathbf{v} \rangle + \mu \langle \nabla \mathbf{u}, \nabla \mathbf{v} \rangle \\ + \langle (\mathbf{w} \cdot \nabla) \mathbf{u}, \mathbf{v} \rangle - \langle p, \nabla \cdot \mathbf{v} \rangle = \langle \mathbf{f}, \mathbf{v} \rangle + \int_{\Gamma_N^{\mathbf{u}}} \mathbf{g} \cdot \mathbf{v}, & \forall \mathbf{v}(\mathbf{x}) \in V(\Omega), \\ - \langle q, \nabla \cdot \mathbf{u} \rangle = 0, & \forall q(\mathbf{x}) \in Q(\Omega) \end{cases} \quad (3.8)$$

for each instant $t \in (T_0, T]$, where $\langle \mathbf{a}(\mathbf{x}), \mathbf{b}(\mathbf{x}) \rangle := \int_{\Omega} \mathbf{a}(\mathbf{x}) \cdot \mathbf{b}(\mathbf{x}) d\mathbf{x}$ (or $\int_{\Omega} a(\mathbf{x}) b(\mathbf{x}) d\mathbf{x}$ when scalar quantities are involved) denotes the spatial scalar product. Notice we dropped the variables' dependence on \mathbf{x} and t in order to contract notation.

3.1.1 Space-time discretisation

For the sake of simplicity, we discretise the temporal derivative in (3.8) using implicit Euler. To this end, we introduce a uniformly-spaced grid, $\{t_k\}$, over the temporal domain, where

$$t_k = T_0 + k\Delta t, \quad \text{with} \quad \Delta t = (T - T_0)/N_t \quad \text{and} \quad k = 0, \dots, N_t, \quad (3.9)$$

at which we seek to recover our numerical solution. Following a finite element approach [118, Chap. 4] [40, Chap. 10], we discretise (3.8) in space by introducing a triangulation of Ω of characteristic length Δx , on which we build finite-dimensional

approximations to both $V(\Omega)$ and $Q(\Omega)$, named $V_h(\Omega)$ and $Q_h(\Omega)$. These are identified by their basis functions

$$V_h(\Omega) := \text{span}\{\phi_0(\mathbf{x}), \dots, \phi_{N_u-1}(\mathbf{x})\}, \quad \text{and} \quad (3.10a)$$

$$Q_h(\Omega) := \text{span}\{\psi_0(\mathbf{x}), \dots, \psi_{N_p-1}(\mathbf{x})\}, \quad (3.10b)$$

with N_u and N_p being the number of degrees of freedom associated with the velocity and pressure variables, respectively, so that vectors $\mathbf{v} \in \mathbb{R}^{N_u}$ and $\mathbf{q} \in \mathbb{R}^{N_p}$ correspond to functions in $V_h(\Omega)$ and $Q_h(\Omega)$ via

$$\mathbf{v}(\mathbf{x}) = \sum_{j=0}^{N_u-1} [\mathbf{v}]_j \phi_j(\mathbf{x}) \quad \text{and} \quad q(\mathbf{x}) = \sum_{j=0}^{N_p-1} [\mathbf{q}]_j \psi_j(\mathbf{x}), \quad (3.11)$$

respectively. (In the remainder of the chapter, we will refer somewhat interchangeably to functions in our discrete spaces and their vector representations according to the bases (3.10).) With this, we assemble the Galerkin matrices, namely the discrete counterparts of the bilinear forms appearing in (3.8). In particular, we define the mass and stiffness matrices for the velocity variable, respectively,

$$[\mathcal{M}_u]_{m,n=0}^{N_u-1} := \int_{\Omega} \phi_m(\mathbf{x}) \cdot \phi_n(\mathbf{x}) \, d\mathbf{x} \quad \text{and} \quad (3.12a)$$

$$[\mathcal{K}_u]_{m,n=0}^{N_u-1} := \int_{\Omega} \nabla \phi_m(\mathbf{x}) : \nabla \phi_n(\mathbf{x}) \, d\mathbf{x}, \quad (3.12b)$$

the discrete (negative) divergence matrix, responsible for the coupling between velocity and pressure,

$$[\mathcal{B}]_{m,n=0}^{N_p-1, N_u-1} := - \int_{\Omega} \psi_m(\mathbf{x}) \nabla \cdot \phi_n(\mathbf{x}) \, d\mathbf{x}, \quad (3.13)$$

and discretisations of the advection operator for each temporal instant t_k in (3.9):

$$[\mathcal{W}_{u,k}]_{m,n=0}^{N_u-1} := \int_{\Omega} ((\mathbf{w}(\mathbf{x}, t_k) \cdot \nabla) \phi_n(\mathbf{x})) \cdot \phi_m(\mathbf{x}) \, d\mathbf{x}. \quad (3.14)$$

Combining spatial and temporal discretisations gives rise to the following discrete linear system approximating (3.8):

$$\begin{cases} \frac{1}{\Delta t} \mathcal{M}_u (\mathbf{u}^k - \mathbf{u}^{k-1}) + \mu \mathcal{K}_u \mathbf{u}^k + \mathcal{W}_{u,k} \mathbf{u}^k + \mathcal{B}^T \mathbf{p}^k = \mathbf{f}^k \\ \mathcal{B} \mathbf{u}^k = \mathbf{0} \end{cases}, \quad k = 1, \dots, N_t, \quad (3.15)$$

where $\mathbf{0}$ is an all-zero vector of size N_p . Here, $[\mathbf{u}^k]_j$ and $[\mathbf{p}^k]_j$ represent the coefficients corresponding to the j -th basis function for the velocity and pressure variables

in eqs. (3.10a) and (3.10b), respectively, evaluated at the k -th temporal instant; analogously, $[\mathbf{f}^k]_j$ identifies the L^2 -projection of the right-hand side of (3.8) on the j -th basis function for $V_h(\Omega)$, at the instant t_k :

$$[\mathbf{f}^k]_j := \int_{\Omega} \mathbf{f}(\mathbf{x}, t_k) \cdot \phi_j(\mathbf{x}) d\mathbf{x} + \int_{\Gamma_N^u} \mathbf{g}(\mathbf{s}, t_k) \cdot \phi_j(\mathbf{s}) ds. \quad (3.16)$$

We now have all the ingredients necessary to assemble the monolithic space-time system corresponding to (3.15). To simplify notation, we introduce the operators

$$\mathcal{F}_{u,k} := \frac{\mathcal{M}_u}{\Delta t} + \mathcal{W}_{u,k} + \mu \mathcal{K}_u, \quad k = 1, \dots, N_t. \quad (3.17)$$

We then arrange all the equations of (3.15) in a single block matrix, as

$$\begin{bmatrix} \boxed{\begin{matrix} \mathcal{F}_{u,1} & \mathcal{B}^T \\ \mathcal{B} \end{matrix}} & & & \\ & \ddots & & \\ & & \boxed{\begin{matrix} -\frac{\mathcal{M}_u}{\Delta t} \end{matrix}} & \ddots \\ & & & \ddots \\ & & & & \boxed{\begin{matrix} \mathcal{F}_{u,N_t} & \mathcal{B}^T \\ \mathcal{B} \end{matrix}} \end{bmatrix} \begin{bmatrix} \mathbf{u}^1 \\ \mathbf{p}^1 \\ \vdots \\ \mathbf{u}^{N_t} \\ \mathbf{p}^{N_t} \end{bmatrix} = \begin{bmatrix} \mathbf{f}^1 \\ \mathbf{0} \\ \vdots \\ \mathbf{f}^{N_t} \\ \mathbf{0} \end{bmatrix}, \quad (3.18)$$

where, with a slight abuse of notation, we re-define \mathbf{f}^1 from (3.16) in order to include the initial condition (3.4),

$$[\mathbf{f}^1]_j := \int_{\Omega} \left(\mathbf{f}(\mathbf{x}, t_1) + \frac{1}{\Delta t} \bar{\mathbf{u}}^0(\mathbf{x}) \right) \cdot \phi_j(\mathbf{x}) d\mathbf{x} + \int_{\Gamma_N^u} \mathbf{g}(\mathbf{s}, t_1) \cdot \phi_j(\mathbf{s}) ds. \quad (3.19)$$

Notice that the system in (3.18) is of the block lower-triangular structure typical of space-time discretisations [47, 62]; moreover, since we use a *one*-step method to approximate the time derivative, the system is made of *two* block diagonals, with the 2×2 blocks composing them highlighted in (3.18). Usually, the solution to (3.15) is recovered via a time-stepping routine, which corresponds to solving (3.18) directly via (block) forward substitution. There are several methods to accelerate this procedure, mostly focusing on accelerating the inversion of the main diagonal blocks

$$\mathcal{A}_k := \begin{bmatrix} \mathcal{F}_{u,k} & \mathcal{B}^T \\ \mathcal{B} \end{bmatrix}, \quad k = 1, \dots, N_t, \quad (3.20)$$

by designing an optimal preconditioner (see for example [118, Chap. 17.8], or [40, Chap. 4] for an overview). The *pressure convection-diffusion* (PCD) preconditioner is one of the most successful among these [85, 137], preconditioning the block structure

of (3.20) by finding an accurate and computable approximation to the pressure Schur complement. Our work stems from the same principles as PCD, but considers the whole space-time system (3.18), rather than just the spatial diagonal blocks (3.20). Motivation and derivation of our preconditioner are described in the following section.

3.2 Space-time block preconditioning for incompressible flow

Rather than inverting (3.18) using sequential time-stepping (i.e., forward substitution), here we tackle the solution of the full space-time system (3.18) all at once by using preconditioned GMRES [127], with an appropriate space-time block preconditioner. In particular, we focus on developing a preconditioner that treats space and time together and in parallel, which may yield faster time-to-solution than block forward substitution when sufficient processors are available. In the remainder of this section, we show how we can effectively build on concepts developed in the study of the single time-step matrix (3.20) to design preconditioners that are effective in the full space-time setting.

To develop our preconditioner, we start by reordering the space-time operator in (3.18) to take the form

$$\left[\begin{array}{ccc|ccc} \mathcal{F}_{\mathbf{u},1} & & & \mathcal{B}^T & & \\ -\frac{\mathcal{M}_{\mathbf{u}}}{\Delta t} & \ddots & & & \ddots & \\ & \ddots & \mathcal{F}_{\mathbf{u},N_t} & & & \mathcal{B}^T \\ \hline \mathcal{B} & & & & & \\ & \ddots & & & & \\ & & & \mathcal{B} & & \end{array} \right] \begin{bmatrix} \mathbf{u}^1 \\ \vdots \\ \mathbf{u}^{N_t} \\ \mathbf{p}^1 \\ \vdots \\ \mathbf{p}^{N_t} \end{bmatrix} = \begin{bmatrix} \mathbf{f}^1 \\ \vdots \\ \mathbf{f}^{N_t} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix}. \quad (3.21)$$

This rearrangement favours a separation of variables based on the physical unknown they refer to, rather than on the individual time-steps. We further denote the blocks composing (3.21) by

$$F_{\mathbf{u}} := \begin{bmatrix} \mathcal{F}_{\mathbf{u},1} & & \\ -\frac{\mathcal{M}_{\mathbf{u}}}{\Delta t} & \ddots & \\ & \ddots & \mathcal{F}_{\mathbf{u},N_t} \end{bmatrix}, \quad \text{and} \quad B := \text{diag}_{N_t}(\mathcal{B}), \quad (3.22)$$

where $\text{diag}_N(*)$ identifies a block diagonal matrix with N blocks containing $(*)$. As in the single time-step case [40, Chap. 9.2], we exploit the block structure of (3.21), and

look for a right preconditioner for GMRES in the following block upper-triangular form:

$$P_T := \begin{bmatrix} F_{\mathbf{u}} & B^T \\ & \tilde{S}_p \end{bmatrix} \iff P_T^{-1} := \begin{bmatrix} F_{\mathbf{u}}^{-1} & -F_{\mathbf{u}}^{-1}B^T\tilde{S}_p^{-1} \\ & \tilde{S}_p^{-1} \end{bmatrix}. \quad (3.23)$$

This is justified in [107], and stems from considering the block LU factorisation of matrix (3.21),

$$\begin{bmatrix} F_{\mathbf{u}} & B^T \\ B & S_p \end{bmatrix} = \begin{bmatrix} I & \\ BF_{\mathbf{u}}^{-1} & I \end{bmatrix} \begin{bmatrix} F_{\mathbf{u}} & B^T \\ & S_p \end{bmatrix}, \quad (3.24)$$

and choosing P_T as an approximation of the rightmost factor: particularly, taking \tilde{S}_p in (3.23) as an approximation of the space-time pressure Schur complement S_p

$$\tilde{S}_p \approx S_p := -BF_{\mathbf{u}}^{-1}B^T. \quad (3.25)$$

Convergence of GMRES applied to the full space-time system and preconditioned by P_T^{-1} is then exactly the same as the convergence of GMRES applied to the Schur complement problem, preconditioned by \tilde{S}_p^{-1} [141].

The key advantage of the reordering in (3.21) and the block preconditioning discussed above is that we have effectively decoupled the imposition of the Lagrangian constraint from the solution of the PDE. Applying the preconditioner then requires inverting $F_{\mathbf{u}}$ (that is, time-integrating the velocity field), and applying an approximate Schur complement inverse (3.25). The former represents the discretisation of a single-variable parabolic PDE, which is much simpler to solve in parallel in space-time than the fully-coupled system (3.21); indeed parallel-in-(space-)time algorithms have shown to behave well on such problems: see [48, 138], and more generally the discussion in Chap. 2. Moreover, we claim that the pressure Schur complement (3.25) can naturally be approximated in a time-parallel fashion. To derive such an approximation, we follow two approaches which have been an inspiration for designing efficient preconditioners for the single time-step case.

Remark 2. One could also consider a block diagonal preconditioner, but due to the nonsymmetry of eqs. (3.22) and (3.25), MINRES [112] cannot be applied, thus negating one of the primary benefits of using a preconditioner in this form. With GMRES, block diagonal preconditioners are typically expected to require twice as many iterations [141] as block triangular ones, and in this case offer only a marginal reduction in computational cost. Block lower-triangular preconditioning is another option, but numerical tests (not included here) have favoured using an upper-triangular preconditioner. The latter is also more suited for right-preconditioning [141], which is required

by flexible GMRES [125] (and flexible Krylov methods are necessary if we use an inner Krylov method to approximately invert F_u). For these reasons, P_T (3.23) remains our preconditioner of choice.

3.2.1 Small commutator approach

In [137], an effective preconditioning strategy for (3.20) is recovered by starting from an assumption on the commuting of certain operators. Although our Schur complement (3.25) differs from the case considered there, as it involves a time-derivative, we can still follow a similar line of reasoning. Let us begin by characterising our space-time Schur complement in more detail. Firstly, we can explicitly write F_u^{-1} as²

$$F_u^{-1} = \begin{bmatrix} \mathcal{F}_{u,1}^{-1} & & & \\ \mathcal{F}_{u,2}^{-1} \left(\frac{\mathcal{M}_u}{\Delta t} \mathcal{F}_{u,1}^{-1} \right) & \mathcal{F}_{u,2}^{-1} & & \\ \mathcal{F}_{u,3}^{-1} \left(\frac{\mathcal{M}_u}{\Delta t} \mathcal{F}_{u,2}^{-1} \frac{\mathcal{M}_u}{\Delta t} \mathcal{F}_{u,1}^{-1} \right) & \mathcal{F}_{u,3}^{-1} \left(\frac{\mathcal{M}_u}{\Delta t} \mathcal{F}_{u,2}^{-1} \right) & \ddots & \\ \vdots & \ddots & \ddots & \ddots \end{bmatrix}. \quad (3.26)$$

Multiplying on the left and right by the block diagonal matrices B and B^T , respectively, results in inner operators (*i.e.*, blocks) of the form

$$\mathcal{B} \mathcal{F}_{u,k}^{-1} \left(\prod_{i=1}^{k-j} \left(\frac{\mathcal{M}_u}{\Delta t} \mathcal{F}_{u,k-i}^{-1} \right) \right) \mathcal{B}^T, \quad \text{with} \quad \begin{cases} k = 1, \dots, N_t \\ j = 1, \dots, k \end{cases}, \quad (3.27)$$

where the product is assumed to expand towards the right, that is $\prod_{i=0}^k a_i = a_0 \prod_{i=1}^{k-1} a_i$, and returns identity if the subscript is larger than the superscript. We seek to adequately approximate the operators appearing in (3.27): to do so, we follow [137] as well as [40, Chap. 9.2], and make the assumption that the reaction-advection-diffusion operator and the gradient operator approximately commute. For each $\mathcal{F}_{u,k}$ in (3.17), let us define $\mathcal{F}_{p,k}$, an equivalent reaction-advection-diffusion operator acting on the pressure field:

$$\mathcal{F}_{p,k} := \frac{\mathcal{M}_p}{\Delta t} + \mathcal{W}_{p,k} + \mu \mathcal{K}_p, \quad k = 1, \dots, N_t, \quad (3.28)$$

²If we were solving the time-dependent Stokes equations (that is, if $\mathbf{w}(\mathbf{x}, t) \equiv 0$), the structure of (3.26) would be further simplified, as F_u would be block Toeplitz; the presence of a time-dependent advection field, instead, causes the main diagonal block to vary down the diagonal. A similar disruption would be caused by using a non-uniform temporal grid when discretising Stokes: in that context, the difference between the blocks down the same diagonal is given by the factor Δt_k by which we divide the mass matrix. A similar analysis as the one reported here can be applied to that case as well.

where \mathcal{M}_p , \mathcal{K}_p , and $\mathcal{W}_{p,k}$ play the same role as the mass, stiffness, and advection matrices ((3.12a), (3.12b), and (3.14)), but for the pressure variable, namely

$$[\mathcal{M}_p]_{m,n=0}^{N_p-1} := \int_{\Omega} \psi_m(\mathbf{x}) \psi_n(\mathbf{x}) d\mathbf{x}, \quad (3.29a)$$

$$[\mathcal{K}_p]_{m,n=0}^{N_p-1} := \int_{\Omega} \nabla \psi_m(\mathbf{x}) \cdot \nabla \psi_n(\mathbf{x}) d\mathbf{x}, \quad \text{and} \quad (3.29b)$$

$$[\mathcal{W}_{p,k}]_{m,n=0}^{N_p-1} := \int_{\Omega} (\mathbf{w}(\mathbf{x}, t_k) \cdot \nabla) \psi_n(\mathbf{x}) \psi_m(\mathbf{x}) d\mathbf{x}, \quad \text{for } k = 1, \dots, N_t. \quad (3.29c)$$

More details on how to assemble these operators are provided in Sec. 3.2.1.2. Translating the commutation assumption to the discrete form introduced in Sec. 3.1, we have

$$(\mathcal{M}_u^{-1} \mathcal{B}^T) (\mathcal{M}_p^{-1} \mathcal{F}_{p,j}) - (\mathcal{M}_u^{-1} \mathcal{F}_{u,j}) (\mathcal{M}_u^{-1} \mathcal{B}^T) \approx 0 \quad (3.30)$$

for each $j = 1, \dots, N_t$. Here, rescaling by the corresponding mass matrices is necessary, since the Galerkin operators have the effect of mapping functions to *functionals*, integrating over the whole spatial domain.

To get closer to our desired operators (3.27), we left-multiply (3.30) by $\mathcal{F}_{u,j}^{-1} \mathcal{M}_u$ and right-multiply by $\mathcal{F}_{p,j}^{-1} \mathcal{M}_p$, thus obtaining

$$\mathcal{F}_{u,j}^{-1} \mathcal{B}^T \approx \mathcal{M}_u^{-1} \mathcal{B}^T \mathcal{F}_{p,j}^{-1} \mathcal{M}_p. \quad (3.31)$$

Continuing to left-multiply by $\mathcal{F}_{u,j+i}^{-1} \frac{\mathcal{M}_u}{\Delta t}$, increasing i until we reach $j+i=k$, using (3.31) for each $j+i$, and finally left-multiplying by \mathcal{B} , gives us

$$\begin{aligned} \mathcal{F}_{u,j+1}^{-1} \left(\frac{\mathcal{M}_u}{\Delta t} \mathcal{F}_{u,j}^{-1} \right) \mathcal{B}^T &\approx \underbrace{\frac{\mathcal{F}_{u,j+1}^{-1} \mathcal{B}^T}{\Delta t}}_{\approx \mathcal{M}_u^{-1} \mathcal{B}^T \mathcal{F}_{p,j+1}^{-1} \frac{\mathcal{M}_p}{\Delta t}} \mathcal{F}_{p,j}^{-1} \mathcal{M}_p \\ \implies \mathcal{B} \mathcal{F}_{u,k}^{-1} \left(\prod_{i=1}^{k-j} \left(\frac{\mathcal{M}_u}{\Delta t} \mathcal{F}_{u,k-i}^{-1} \right) \right) \mathcal{B}^T &\approx \mathcal{B} \mathcal{M}_u^{-1} \mathcal{B}^T \mathcal{F}_{p,k}^{-1} \left(\prod_{i=1}^{k-j} \left(\frac{\mathcal{M}_p}{\Delta t} \mathcal{F}_{p,k-i}^{-1} \right) \right) \mathcal{M}_p, \end{aligned} \quad (3.32)$$

for any desired index $k = 1, \dots, N_t$. This gives us an approximation for each of the operators in (3.27), and consequently identifies the following candidate for a reasonable approximation of the space-time pressure Schur complement (3.25):

$$\tilde{S}_p := - \text{diag}_{N_t} (\mathcal{B} \mathcal{M}_u^{-1} \mathcal{B}^T) \begin{bmatrix} \mathcal{F}_{p,1}^{-1} \\ \mathcal{F}_{p,2}^{-1} \left(\frac{\mathcal{M}_p}{\Delta t} \mathcal{F}_{p,1}^{-1} \right) & \mathcal{F}_{p,2}^{-1} \\ \mathcal{F}_{p,3}^{-1} \left(\frac{\mathcal{M}_p}{\Delta t} \mathcal{F}_{p,2}^{-1} \frac{\mathcal{M}_p}{\Delta t} \mathcal{F}_{p,1}^{-1} \right) & \mathcal{F}_{p,3}^{-1} \left(\frac{\mathcal{M}_p}{\Delta t} \mathcal{F}_{p,2}^{-1} \right) & \ddots \\ \vdots & \ddots & \ddots \end{bmatrix} \text{diag}_{N_t} (\mathcal{M}_p). \quad (3.33)$$

The expression above can be simplified by noticing that the central matrix shares the same structure as (3.26). In fact, it is exactly the inverse of the block bi-diagonal matrix

$$F_p := \begin{bmatrix} \mathcal{F}_{p,1} & & & \\ -\frac{\mathcal{M}_p}{\Delta t} & \ddots & & \\ & & \ddots & \mathcal{F}_{p,N_t} \\ & & & \end{bmatrix}, \quad (3.34)$$

which can be interpreted as a system stemming from the implicit Euler discretisation of a time-dependent PDE, just like its velocity counterpart (3.22). Moreover, by making use of the discrete inf-sup condition, it has been shown [40, Chap. 3.5] that the operator $\mathcal{B}\mathcal{M}_u^{-1}\mathcal{B}^T$ in (3.33) is spectrally equivalent to a discrete Laplacian operator acting on the pressure field. In principle, this operator can differ from the one appearing in (3.28), hence we denote it as

$$\tilde{\mathcal{K}}_p \approx \mathcal{B}\mathcal{M}_u^{-1}\mathcal{B}^T, \quad (3.35)$$

and delay a more in-depth discussion on its definition to Sec. 3.2.1.2. Furthermore, to simplify notation, we define the following block diagonal matrices

$$M_p := \text{diag}_{N_t}(\mathcal{M}_p) \quad \text{and} \quad K_p := \text{diag}_{N_t}(\tilde{\mathcal{K}}_p), \quad (3.36)$$

which ultimately allows us to write our approximation to the space-time pressure Schur complement (3.33) as

$$\tilde{\mathcal{S}}_p := -K_p F_p^{-1} M_p \quad \iff \quad \tilde{\mathcal{S}}_p^{-1} := -M_p^{-1} F_p K_p^{-1}. \quad (3.37)$$

3.2.1.1 Comparison with single time-step block preconditioner

As briefly remarked at the end of Sec. 3.1, the PCD preconditioner [137] which is used for the acceleration of the solution for the single time-step system (3.20), and which largely motivates our work, is given by

$$\mathcal{P}_{T,k} := \begin{bmatrix} \mathcal{F}_{u,k} & \mathcal{B}^T \\ & \tilde{\mathcal{S}}_{p,k} \end{bmatrix} \quad \iff \quad \mathcal{P}_{T,k}^{-1} := \begin{bmatrix} \mathcal{F}_{u,k}^{-1} & -\mathcal{F}_{u,k}^{-1}\mathcal{B}^T\tilde{\mathcal{S}}_{p,k}^{-1} \\ & \tilde{\mathcal{S}}_{p,k}^{-1} \end{bmatrix}, \quad (3.38)$$

for each time-step $k = 1, \dots, N_t$, where the single time-step pressure Schur complement, $\mathcal{S}_{p,k} := -\mathcal{B}\mathcal{F}_{u,k}^{-1}\mathcal{B}^T$, is approximated as

$$\mathcal{S}_{p,k} \approx \tilde{\mathcal{S}}_{p,k} := -\tilde{\mathcal{K}}_p \mathcal{F}_{p,k}^{-1} \mathcal{M}_p \quad \iff \quad \tilde{\mathcal{S}}_{p,k}^{-1} := -\mathcal{M}_p^{-1} \mathcal{F}_{p,k} \tilde{\mathcal{K}}_p^{-1}, \quad (3.39)$$

rendering it strikingly similar to (3.37). The sum of the computational costs associated with applying the pressure Schur complement approximation (3.39) at each

time-step is analogous to that of applying its space-time counterpart once; moreover, the application of (3.37) can be conducted in parallel over time naturally. We can convince ourselves of this by considering the block diagonal structure of the matrices in (3.37): applying the inverses of K_p and M_p requires solving N_t independent systems involving $\tilde{\mathcal{K}}_p$ and \mathcal{M}_p , akin to applying (3.39) for each of the N_t time-steps. The sole overhead with respect to sequential time-stepping comes from the *bi*-diagonal structure of F_p : the off-diagonal coupling between one time-step and the next must be accounted for for *each* application of the space-time preconditioner, rather than once per time-step. Nonetheless, this extra cost remains negligible compared to the other operations.

In light of these considerations, we conclude that the difference in performance between using (3.38) at each time-step, and tackling the whole space-time system at once with (3.23), is given by two factors: (i) the difference in the total number of GMRES iterations to convergence using the space-time block preconditioner (3.23), which we denote with N_{it}^{ST} , and the average number of GMRES iterations to convergence per time-step using (3.38), identified by N_{it}^0 ; and (ii) the computational time associated with inverting the velocity spatial operator $\mathcal{F}_{u,k}$ at each time-step versus inverting the whole space-time velocity block F_u . We denote the latter as C_{F_u} , while for the former we consider an average of $C_{\mathcal{F}_u}$ per time-step, so that the total time is given by $C_{\mathcal{F}_u} \cdot N_t$. We expect our approach to be competitive if

$$\frac{N_{it}^{ST}}{N_{it}^0} C_{F_u} < C_{\mathcal{F}_u} N_t. \quad (3.40)$$

The ratio N_{it}^{ST}/N_{it}^0 defines the overhead cost of the all-at-one approach in terms of the total number of preconditioner applications. Numerical results in Sec. 3.4.4 indicate that this ratio is $\in (1, 2)$ for most cases, and only increasing to 3 in some extreme settings. Thus, the overhead in performing all-at-once block preconditioning in the space-time setting is quite marginal compared to standard spatial preconditioning in sequential time-stepping. The efficiency of our approach then directly depends on the speedup that can be obtained when solving a vector-based time-dependent (advection-)diffusion problem using PinT or all-at-once techniques. With a sufficiently large number of MPI processes available, however, PinT and all-at-once solutions to space-time (advection-)diffusion problems have shown significant speedups over sequential time-stepping, as outlined in Chap. 2.

We point out that in general our space-time approach is associated with a larger memory footprint than sequential time-stepping. In order to ensure that the parallel

application of (3.37) can be conducted in a time comparable to that of applying (3.39), in fact, in our implementation we assign one processor per time-step, and have each processor assemble all the relevant spatial operators. As a consequence, the required memory grows linearly with the number of processors employed. Even though this effectively caps the maximum system size of our simulations in Sec. 3.4, such characteristic is common among PinT methods, and we do not consider it as a particularly burdensome limitation. As outlined in Chap. 1, in fact, time-parallelisation is generally employed within scenarios where computational resources are not a limiting factor.

3.2.1.2 Defining the pressure operators

One issue associated with the definition of the pressure reaction-advection-diffusion operator (3.28) and the pressure Laplacian (3.35) involves the choice of BC to impose on the operators themselves. In [41] and [40, Chap. 9.2.2], it is suggested that the various $\mathcal{F}_{p,k}$ should be assembled as if Robin BC,

$$\mu \nabla p(\mathbf{s}, t) \cdot \mathbf{n}(\mathbf{s}) - (\mathbf{w}(\mathbf{s}, t) \cdot \mathbf{n}(\mathbf{s})) p(\mathbf{s}, t) = 0, \quad (3.41)$$

were imposed everywhere, while $\tilde{\mathcal{K}}_p$ should have Dirichlet BC on the outflow boundary (that is, on $\mathbf{s} \in \partial\Omega : \mathbf{u}(\mathbf{s}) \cdot \mathbf{n}(\mathbf{s}) > 0$), and Neumann BC everywhere else. In particular, this choice implies that \mathcal{K}_p in (3.28) might differ from $\tilde{\mathcal{K}}_p$. However, prescribing the same BC for the assembly of both the pressure reaction-advection-diffusion operator and the pressure Laplacian has the advantage of further simplifying (3.37) for the Stokes equations (3.2). Under the assumption $\mathcal{K}_p \equiv \tilde{\mathcal{K}}_p$, in fact, this becomes

$$\tilde{\mathcal{S}}_p^{-1} \begin{pmatrix} = \\ \text{(if } \mathcal{K}_p \equiv \tilde{\mathcal{K}}_p) \\ \mathbf{w} \equiv \mathbf{0} \end{pmatrix} - \begin{bmatrix} \frac{\mathcal{K}_p^{-1}}{\Delta t} + \mu \mathcal{M}_p^{-1} & & \\ -\frac{\mathcal{K}_p^{-1}}{\Delta t} & \ddots & \\ & & \ddots \end{bmatrix}. \quad (3.42)$$

In comparison to (3.37), we need to multiply by one fewer operator, and inverting \mathcal{M}_p and \mathcal{K}_p can be done independently. (Notice that a preconditioner in this form is reminiscent of the Cahouet-Chabard preconditioner [14].) Hence, keeping performance in mind, in our code we rather choose to assemble both the pressure Laplacian and the pressure reaction-diffusion operator considering homogeneous Dirichlet BC at the outflow of the domain, and homogeneous Neumann BC everywhere else, thus mimicking the implementation in the IFISS package [42] (available at [82]). Moreover, early experiments showed that the two approaches provide comparable results in terms of acceleration of convergence for the test-cases considered.

We also remark that, using the BC prescribed above, (3.29b) becomes singular when dealing with fully enclosed Stokes flows, for which $\mathbf{u}(\mathbf{s}) \cdot \mathbf{n}(\mathbf{s}) = 0$, $\forall \mathbf{s} \in \partial\Omega$. In this context, in fact, the operator mimics the action of a Laplacian with Neumann BC everywhere on the boundary, whose kernel contains all constant functions. The GMRES procedure remains robust in this case too [40, Chap. 9.3.5], but we need to ensure that the chosen solver for $\tilde{\mathcal{K}}_p$ can deal with singular matrices.

Finally, since with (3.7) we only require L^2 -regularity in space for the pressure variable, the operator (3.29b) might be ill-defined, depending on the discretisation chosen. A solution to this issue when considering a discontinuous pressure field is detailed in [40, Chap. 9.2.1].

3.2.1.3 Small commutator approach focusing on the *divergence* operator

The most recent edition of [40] proposes to consider the commutator between the discrete *divergence* (rather than *gradient*) operator, and the reaction-diffusion operator. This provides a better framework for justifying the kind of BC to assign to the discrete operators defined over the pressure field (see also [41]). In our case, choosing this approach would translate to considering the commutator

$$(\mathcal{M}_p^{-1} \mathcal{F}_{p,k}) (\mathcal{M}_p^{-1} \mathcal{B}) - (\mathcal{M}_p^{-1} \mathcal{B}) (\mathcal{M}_u^{-1} \mathcal{F}_{u,k}) \approx 0, \quad (3.43)$$

as opposed to (3.30). Following steps similar to the derivation in (3.32), this would eventually lead to a definition of the pressure Schur complement approximation where the block diagonal operators containing the pressure stiffness and mass matrix are swapped with respect to (3.37), that is:

$$\tilde{S}_p := -M_p F_p^{-1} K_p \iff \tilde{S}_p^{-1} := -K_p^{-1} F_p M_p^{-1}. \quad (3.44)$$

The operations necessary for its application are virtually the same as those for (3.37). However, numerical experiments have indicated that the definition (3.44) does not provide significant improvements over (3.37) in terms of convergence. Additionally, the form (3.37) is more prone to simplifications similar to the one illustrated in (3.42) when solving Oseen equations. Therefore, for the remainder of this chapter, we consider the gradient commutator (3.37).

3.2.2 Green's function approach

Alternatively to the small-commutator approach described in Sec. 3.2.1, another heuristic for recovering a good candidate preconditioner for the Oseen equations has

been proposed in [85]. It is based on mapping the operators appearing in (3.1) to Fourier space, finding the associated Green's tensor in Fourier domain, and taking inspiration from its form to recover an approximation of the Schur complement. Here we demonstrate that such principles can also be extended to the space-time setting, and, in fact, arrive at the same Schur-complement approximation as derived in (3.37), providing additional validation for our choice of preconditioner.

We start by arranging the operators appearing in (3.1) in tensor form,

$$\left[\begin{array}{c|c} \mathcal{D}_\varphi & \nabla \\ \hline -\nabla \cdot & 0 \end{array} \right] \left[\begin{array}{c} \mathbf{u}(\mathbf{x}, t) \\ p(\mathbf{x}, t) \end{array} \right] = \left[\begin{array}{c} \mathbf{f}(\mathbf{x}, t) \\ 0 \end{array} \right], \quad (3.45)$$

where $\mathcal{D}_{(*)}$ denotes the block diagonal operator responsible for applying $(*)$ to each component of a vector in \mathbb{R}^d , while φ is defined as

$$\varphi := \frac{\partial}{\partial t} + \mathbf{w} \cdot \nabla - \mu \nabla^2. \quad (3.46)$$

The (tensor) operator associated to (3.1) is then given by

$$\mathcal{L} := \left[\begin{array}{c|c} \mathcal{D}_\varphi & \nabla \\ \hline -\nabla \cdot & 0 \end{array} \right]. \quad (3.47)$$

We seek to recover a candidate approximation for the pressure Schur complement appearing in our preconditioner (3.23), by following the form of the Green's tensor \mathcal{G} associated with (3.47), under the assumption of an advection field constant over the spatial domain (but not necessarily in time), $\mathbf{w}(\mathbf{x}, t) \equiv \mathbf{w}(t)$. To this purpose, we can exploit the fact that the Fourier transform of the Green's tensor $\hat{\mathcal{G}}$ is given by the inverse of the Fourier transform of (3.47), that is, $\hat{\mathcal{G}} = \hat{\mathcal{L}}^{-1}$. The operator equivalent to (3.46) in Fourier space is given by

$$\hat{\varphi} = \frac{\partial}{\partial t} + i\mathbf{w} \cdot \mathbf{k} + \mu|\mathbf{k}|^2, \quad (3.48)$$

where $\mathbf{k} = [k_0, \dots, k_{d-1}]^T$ is the frequency vector, and i represents the imaginary unit; consequently, the Fourier equivalent of (3.47) is given by

$$\hat{\mathcal{L}} = \left[\begin{array}{c|c} \mathcal{D}_{\hat{\varphi}} & i\mathbf{k} \\ \hline -i\mathbf{k} \cdot & 0 \end{array} \right]. \quad (3.49)$$

Let $\hat{\sigma}_p$ denote the (2,2) Schur complement in Fourier space,

$$\hat{\sigma}_p := i\mathbf{k}^T \mathcal{D}_{\hat{\varphi}^{-1}} i\mathbf{k} = -|\mathbf{k}|^2 \hat{\varphi}^{-1}. \quad (3.50)$$

Then, using the classic formula for inversion of a 2×2 block operator based on a block LDU decomposition, we recover

$$\begin{aligned}\hat{\mathcal{G}} = \hat{\mathcal{L}}^{-1} &= \left[\begin{array}{c|c} \mathcal{D}_{\hat{\varphi}^{-1}} - \mathcal{D}_{\hat{\varphi}^{-1}} \mathbf{k} \hat{\sigma}_p^{-1} \mathbf{k}^T \mathcal{D}_{\hat{\varphi}^{-1}} & -\mathcal{D}_{\hat{\varphi}^{-1}} \mathbf{i} \mathbf{k} \hat{\sigma}_p^{-1} \\ \hline \hat{\sigma}_p^{-1} \mathbf{i} \mathbf{k}^T \mathcal{D}_{\hat{\varphi}^{-1}} & \hat{\sigma}_p^{-1} \end{array} \right] \\ &= (\hat{\varphi} |\mathbf{k}|^2)^{-1} \left[\begin{array}{c|c} \mathcal{D}_{|\mathbf{k}|^2} - \mathbf{k} \mathbf{k}^T & \mathbf{i} \mathbf{k} \hat{\varphi} \\ \hline -\mathbf{i} \mathbf{k}^T \hat{\varphi} & -\hat{\varphi}^2 \end{array} \right].\end{aligned}\quad (3.51)$$

We focus our attention on the bottom-right block of (3.51). Transforming back from Fourier space to physical space, and using the properties of Fourier transforms, namely the correspondence of the operators $|\mathbf{k}|^2 \leftrightarrow -\nabla^2$ and $-\mathbf{k} \mathbf{k}^T \leftrightarrow \nabla(\nabla \cdot)$, as well as the identity $\nabla \times (\nabla \times) = \nabla(\nabla \cdot) - \nabla^2$, we get an expression for the Green's tensor:

$$\mathcal{G} = \left[\begin{array}{c|c} (\nabla \times (\nabla \times)) \mathcal{D}_{\mathcal{G}_{\hat{\varphi} \nabla^2}} & \nabla \mathcal{G}_{-\nabla^2} \\ \hline (-\nabla \cdot) \mathcal{D}_{\mathcal{G}_{-\nabla^2}} & -\left(\frac{\partial}{\partial t} + \mathbf{w} \cdot \nabla - \mu \nabla^2\right) \mathcal{G}_{-\nabla^2} \end{array} \right], \quad (3.52)$$

where $\mathcal{G}_{-\nabla^2}$ is the fundamental solution for the negative Laplacian operator, while $\mathcal{G}_{\varphi \nabla^2}$ is the fundamental solution for the operator $\nabla^2 \left(\frac{\partial}{\partial t} + \mathbf{w} \cdot \nabla - \mu \nabla^2\right)$. Keeping only the bottom-right block and discarding the rest, we get that the space-time pressure Schur complement operator σ_p satisfies

$$\sigma_p^{-1} q \approx - \left(\left(\frac{\partial}{\partial t} + \mathbf{w} \cdot \nabla - \mu \nabla^2 \right) \mathcal{G}_{-\nabla^2} \right) * q, \quad \forall q \in Q(\Omega), \quad (3.53)$$

on which we base its approximation. Considering that $\mathcal{G}_{-\nabla^2} * q$ should represent the solution to a (negative) Laplacian defined on the pressure space, we can discretise it using $\tilde{\mathcal{K}}_p^{-1} q$ defined in (3.35). Analogously, the space-time operator acting on the pressure space can be discretised via $F_p \approx \frac{\partial}{\partial t} + \mathbf{w} \cdot \nabla - \mu \nabla^2$, introduced in (3.34).

Altogether, this analysis suggests that a discrete approximation to the space-time pressure Schur complement S_p should satisfy

$$S_p^{-1} \approx \tilde{S}_p^{-1} := -M_p^{-1} F_p K_p^{-1}, \quad (3.54)$$

where M_p and K_p are defined as in (3.36). Note that here we rescaled the operator by the pressure mass matrix for similar reasons as discussed for (3.30); this rescaling was also applied to the single time-step case in [85]. We see that the Green's function heuristic followed in this section yields the same approximation to the space-time pressure Schur complement as found via the small-commutator approach in Sec. 3.2.1. This is not necessarily surprising, as both analyses in some sense rely on neglecting boundary conditions, but the fact that two approaches provide the same approximation (3.54) gives confidence in the validity of the proposed preconditioner.

3.3 Eigenvalues clustering

To provide an indication of the effectiveness of (3.23) as a preconditioner, we investigate how the eigenvalues of the resulting preconditioned system spread in the complex plane. Although eigenvalues are not necessarily indicative of fast convergence for GMRES, particularly for highly nonsymmetric systems [68], poor clustering of eigenvalues is likely to yield poor convergence, while nicely bounded and clustered eigenvalues are still often indicative of an effective preconditioning.

We are interested in solving the following generalised eigenvalue problem:

$$\begin{bmatrix} F_{\mathbf{u}} & B^T \\ B & \tilde{S}_p \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \lambda \begin{bmatrix} F_{\mathbf{u}} & B^T \\ & \tilde{S}_p \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix}. \quad (3.55)$$

The system presents $\lambda = 1$ as a solution with multiplicity $N_{\mathbf{u}}$, corresponding to a choice of $\mathbf{p} = \mathbf{0}$; the remaining eigenvalues are given by satisfying

$$\left(\tilde{S}_p^{-1} B F_{\mathbf{u}}^{-1} B^T + \lambda I \right) \mathbf{p} = 0. \quad (3.56)$$

From this, we see that the general eigenvalues of (3.55) directly correspond to the eigenvalues of the matrix $-\tilde{S}_p^{-1} B F_{\mathbf{u}}^{-1} B^T$. Substituting our approximation to the space-time pressure Schur complement (3.37), we have that the operator of interest becomes

$$-\tilde{S}_p^{-1} B F_{\mathbf{u}}^{-1} B^T = M_p^{-1} F_p K_p^{-1} B F_{\mathbf{u}}^{-1} B^T. \quad (3.57)$$

This matrix is block lower-triangular, because F_p and $F_{\mathbf{u}}^{-1}$ are both block lower-triangular and the remaining factors are block diagonal. Thus, to find its eigenvalues, we just need to recover the eigenvalues of the blocks on the main diagonal. These blocks consist of the operators

$$\mathcal{M}_p^{-1} \mathcal{F}_{p,k} \tilde{\mathcal{K}}_p^{-1} \mathcal{B} \mathcal{F}_{\mathbf{u},k}^{-1} \mathcal{B}^T, \quad \text{with} \quad k = 1, \dots, N_t, \quad (3.58)$$

whose factors can be separated as

$$\left(\tilde{\mathcal{K}}_p \mathcal{F}_{p,k}^{-1} \mathcal{M}_p \right)^{-1} \quad \text{and} \quad \mathcal{B} \mathcal{F}_{\mathbf{u},k}^{-1} \mathcal{B}^T, \quad \text{with} \quad k = 1, \dots, N_t. \quad (3.59)$$

By comparing these two factors with the definitions of the single time-step pressure Schur complement $\mathcal{S}_{p,k}$ and of its PCD approximation $\tilde{\mathcal{S}}_{p,k}$, both provided in Sec. 3.2.1.1, we can convince ourselves that the eigenvalue distribution of the space-time preconditioned system is directly dependent on that of the single time-step preconditioned system, and in particular on the spectral equivalence of $\mathcal{S}_{p,k}$ and $\tilde{\mathcal{S}}_{p,k}$, which we investigate next.

In order to simplify the upcoming analysis, we limit ourselves to the Stokes case, and we substitute $\tilde{\mathcal{K}}_p$ with \mathcal{K}_p , following Sec. 3.2.1.2. This renders the operator on the left in (3.59) *Symmetric Positive Definite* (SPD), since

$$\tilde{\mathcal{K}}_p \mathcal{F}_p^{-1} \mathcal{M}_p \begin{matrix} = \\ \left(\begin{array}{l} \text{if } \mathcal{K}_p \equiv \tilde{\mathcal{K}}_p \\ \mathbf{w} \equiv \mathbf{0} \end{array} \right) \end{matrix} \left(\frac{\mathcal{K}_p^{-1}}{\Delta t} + \mu \mathcal{M}_p^{-1} \right)^{-1}, \quad (3.60)$$

where both \mathcal{K}_p and \mathcal{M}_p are SPD. Notice that since $\mathbf{w} \equiv \mathbf{0}$, we can drop the dependence on the index k and write $\mathcal{F}_p \equiv \mathcal{F}_{p,k}$ and $\mathcal{F}_u \equiv \mathcal{F}_{u,k}$, for each k .

Under these assumptions, we can investigate the spectral equivalence of the factors in (3.59) by considering the *generalised singular value problem*³

$$\begin{bmatrix} & \mathcal{B}^T \\ \mathcal{B} & \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \sigma \begin{bmatrix} \mathcal{F}_u & \\ & \mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix}. \quad (3.62)$$

In fact, we have that the following relationship holds between the generalised singular values and the Rayleigh quotient of the matrices in (3.59):

$$\frac{\langle \mathcal{B} \mathcal{F}_u^{-1} \mathcal{B}^T \mathbf{q}, \mathbf{q} \rangle}{\langle \mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle} = \sigma^2 = \frac{\langle \mathcal{B}^T \mathcal{M}_p^{-1} \mathcal{F}_p \mathcal{K}_p^{-1} \mathcal{B} \mathbf{v}, \mathbf{v} \rangle}{\langle \mathcal{F}_u \mathbf{v}, \mathbf{v} \rangle}, \quad (3.63)$$

for all $(\mathbf{v}, \mathbf{q}) \in V_h(\Omega) \times Q_h(\Omega)$ such that $\langle \mathcal{F}_u \mathbf{v}, \mathbf{v} \rangle = \langle \mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle$. Proving their spectral equivalence is then reduced to bounding σ^2 , *i.e.*, finding two scalars $\underline{\sigma}^2$ and $\bar{\sigma}^2$ such that

$$\underline{\sigma}^2 \leq \frac{\langle \mathcal{B} \mathcal{F}_u^{-1} \mathcal{B}^T \mathbf{q}, \mathbf{q} \rangle}{\langle \mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle} = \sigma^2 = \frac{\langle \mathcal{B}^T \mathcal{M}_p^{-1} \mathcal{F}_p \mathcal{K}_p^{-1} \mathcal{B} \mathbf{v}, \mathbf{v} \rangle}{\langle \mathcal{F}_u \mathbf{v}, \mathbf{v} \rangle} \leq \bar{\sigma}^2, \quad (3.64)$$

either $\forall \mathbf{q} \in Q_h(\Omega)$ or $\forall \mathbf{v} \in V_h(\Omega)$. To simplify notation in the following derivation, we introduce these norms for our discrete pressure space:

$$\|\mathbf{q}\|_p^2 = \langle \mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle \quad \text{and} \quad \|\mathbf{q}\|_0^2 = \langle \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle. \quad (3.65)$$

Notice $\|\cdot\|_0$ is simply the L^2 -norm. For the discrete velocity space, instead, we define

$$\|\mathbf{v}\|_u^2 = \langle \mathcal{F}_u \mathbf{v}, \mathbf{v} \rangle \quad \text{and} \quad \|\mathbf{v}\|_1^2 = \langle \mathcal{K}_u \mathbf{v}, \mathbf{v} \rangle, \quad (3.66)$$

where $\|\mathbf{v}\|_1 \equiv \|\nabla \mathbf{v}(\mathbf{x})\|_0$, the H_1 -semi-norm.

³This procedure is akin to the analysis performed in [40, Chap. 3.5] for the Stokes equations. In that case, the aim was to prove the spectral equivalence between the chosen preconditioner \mathcal{M}_p and the pressure Schur complement $\mathcal{B} \mathcal{K}_u^{-1} \mathcal{B}^T$: the generalised singular value problem considered for this purpose was

$$\begin{bmatrix} & \mathcal{B}^T \\ \mathcal{B} & \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \sigma \begin{bmatrix} \mathcal{K}_u & \\ & \mathcal{M}_p \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix}. \quad (3.61)$$

In our case the structure is identical, but the operators involved take into account the extra terms coming from the temporal discretisation: the single time step pressure Schur complement is $\mathcal{B} \mathcal{F}_u \mathcal{B}^T$, and the proposed approximation is $\mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p$.

Upper bound. We first find $\bar{\sigma}^2$: starting from the Rayleigh quotient on the left of (3.64), we obtain

$$\frac{\langle \mathcal{B}\mathcal{F}_u^{-1}\mathcal{B}^T\mathbf{q}, \mathbf{q} \rangle^{\frac{1}{2}}}{\langle \mathcal{K}_p\mathcal{F}_p^{-1}\mathcal{M}_p\mathbf{q}, \mathbf{q} \rangle^{\frac{1}{2}}} = \frac{\langle \mathcal{F}_u^{-\frac{1}{2}}\mathcal{B}^T\mathbf{q}, \mathcal{F}_u^{-\frac{1}{2}}\mathcal{B}^T\mathbf{q} \rangle^{\frac{1}{2}}}{\|\mathbf{q}\|_p} = \frac{|\langle \mathbf{q}, \mathcal{B}\mathbf{v} \rangle|}{\|\mathbf{q}\|_p\|\mathbf{v}\|_u}, \quad (3.67)$$

where we have substituted $\mathbf{v} = \mathcal{F}_u^{-1}\mathcal{B}^T\mathbf{q}$. We then focus on the right-most term, and notice its numerator equals $|\langle q(\mathbf{x}), \nabla \cdot \mathbf{v}(\mathbf{x}) \rangle|$, where $q(\mathbf{x})$ and $\mathbf{v}(\mathbf{x})$ are the functions corresponding to the vectors \mathbf{q} and \mathbf{v} , according to (3.11). Using Cauchy-Schwarz,

$$|\langle q(\mathbf{x}), \nabla \cdot \mathbf{v}(\mathbf{x}) \rangle| \leq \|q(\mathbf{x})\|_0 \|\nabla \cdot \mathbf{v}(\mathbf{x})\|_0, \quad (3.68)$$

and moreover, we have

$$\|\nabla \cdot \mathbf{v}(\mathbf{x})\|_0^2 \leq \|\nabla \cdot \mathbf{v}(\mathbf{x})\|_0^2 + \|\nabla \times \mathbf{v}(\mathbf{x})\|_0^2 = \|\nabla \mathbf{v}(\mathbf{x})\|_0^2 = \|\mathbf{v}\|_1^2, \quad (3.69)$$

which allows us to write

$$\frac{|\langle \mathbf{q}, \mathcal{B}\mathbf{v} \rangle|}{\|\mathbf{q}\|_p\|\mathbf{v}\|_u} \leq \max_{\mathbf{q} \in Q_h(\Omega)} \frac{\|\mathbf{q}\|_0}{\|\mathbf{q}\|_p} \max_{\mathbf{v} \in V_h(\Omega)} \frac{\|\mathbf{v}\|_1}{\|\mathbf{v}\|_u}, \quad (3.70)$$

at which point finding the upper bound for (3.64) becomes an exercise in proving equivalence of norms. For the velocity field, we have:

$$\|\mathbf{v}\|_1^2 = \langle \mathcal{K}_v \mathbf{v}, \mathbf{v} \rangle \leq \max\{1, \mu^{-1}\} \left\langle \left(\frac{\mathcal{M}_u}{\Delta t} + \mu \mathcal{K}_u \right) \mathbf{v}, \mathbf{v} \right\rangle = \max\{1, \mu^{-1}\} \|\mathbf{v}\|_u^2. \quad (3.71)$$

For the pressure field, the analysis requires more care. Firstly, we can write

$$\frac{\|\mathbf{q}\|_0^2}{\|\mathbf{q}\|_p^2} = \frac{\langle \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle}{\langle \mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle} \leq \frac{\bar{\lambda}(\mathcal{M}_p)}{\underline{\lambda}(\mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p)} = \bar{\lambda}(\mathcal{M}_p) \bar{\lambda}(\mathcal{K}_p^{-1} \mathcal{F}_p \mathcal{M}_p^{-1}), \quad (3.72)$$

where $\underline{\lambda}(\ast)$ and $\bar{\lambda}(\ast)$ respectively denote the smallest and largest eigenvalues of matrix (\ast) , and we once again exploit the fact that the matrices involved are SPD. We also notice that

$$\lambda(\mathcal{K}_p^{-1} \mathcal{F}_p \mathcal{M}_p^{-1}) = \lambda\left(\frac{\mathcal{K}_p^{-1}}{\Delta t} + \mu \mathcal{M}_p^{-1}\right) \subseteq \left[\frac{\Delta t^{-1}}{\bar{\lambda}(\mathcal{K}_p)} + \frac{\mu}{\bar{\lambda}(\mathcal{M}_p)}, \frac{\Delta t^{-1}}{\underline{\lambda}(\mathcal{K}_p)} + \frac{\mu}{\underline{\lambda}(\mathcal{M}_p)} \right]. \quad (3.73)$$

The final ingredient we need to recover an expression for $\bar{\sigma}^2$, is a bound on the eigenvalues of the Galerkin mass and stiffness matrices: for $\Omega \in \mathbb{R}^2$, these are given by

$$\begin{aligned} \underline{c}_{\mathcal{M}}^{(\ast)} \Delta x^2 &\leq \lambda(\mathcal{M}_{(\ast)}) \leq \bar{c}_{\mathcal{M}}^{(\ast)} \Delta x^2 && [40, \text{Prop.1.30}], \\ \underline{c}_{\mathcal{K}}^{(\ast)} \Delta x^2 &\leq \lambda(\mathcal{K}_{(\ast)}) \leq \bar{c}_{\mathcal{K}}^{(\ast)} && [40, \text{Thm.1.33}], \end{aligned} \quad (3.74)$$

where $\underline{c}_{\mathcal{M}}$, $\bar{c}_{\mathcal{M}}$, $\underline{c}_{\mathcal{K}}$, and $\bar{c}_{\mathcal{K}}$ are constants depending on the domain and the type of discretisation; these bounds hold both for pressure and velocity, $(*) = p$ or $(*) = \mathbf{u}$, and are valid for a quasi-uniform discretisation defined on a regular 2D domain, with characteristic mesh size Δx . Substituting (3.74) for the pressure variable in (3.72) gives the upper bound for the equivalence of the norms introduced in (3.65):

$$\frac{\|\mathbf{q}\|_0^2}{\|\mathbf{q}\|_p^2} \leq \bar{c}_{\mathcal{M}} \Delta x^2 \left(\frac{1}{\Delta t} \frac{1}{\underline{c}_{\mathcal{K}}^p \Delta x^2} + \frac{\mu}{\underline{c}_{\mathcal{M}}^p \Delta x^2} \right) = \frac{1}{\Delta t} \frac{\bar{c}_{\mathcal{M}}^p}{\underline{c}_{\mathcal{K}}^p} + \frac{\mu \bar{c}_{\mathcal{M}}^p}{\underline{c}_{\mathcal{M}}^p}. \quad (3.75)$$

The right bound on (3.64) is finally recovered by combining eqs. (3.71) and (3.75)

$$\frac{\langle \mathcal{B}\mathcal{F}_u^{-1}\mathcal{B}^T \mathbf{q}, \mathbf{q} \rangle}{\langle \mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle} \leq \max\{\mu, 1\} \left(\frac{1}{\mu \Delta t} \frac{\bar{c}_{\mathcal{M}}^p}{\underline{c}_{\mathcal{K}}^p} + \frac{\bar{c}_{\mathcal{M}}^p}{\underline{c}_{\mathcal{M}}^p} \right) = \bar{\sigma}^2. \quad (3.76)$$

Lower bound. To find an expression for the lower bound $\underline{\sigma}^2$, we make use of the discrete inf-sup condition

$$0 < \gamma \leq \min_{\mathbf{q} \neq \mathbf{1}} \max_{\mathbf{v} \neq \mathbf{0}} \frac{|\langle \mathbf{q}, \mathcal{B}\mathbf{v} \rangle|}{\|\mathbf{v}\|_1 \|\mathbf{q}\|_0}, \quad (3.77)$$

where $\mathbf{1}$ denotes a vector filled with 1 (representing a FE function constant over the whole domain). Starting from (3.67) and using (3.77) we obtain

$$\frac{\langle \mathcal{B}\mathcal{F}_u^{-1}\mathcal{B}^T \mathbf{q}, \mathbf{q} \rangle^{\frac{1}{2}}}{\langle \mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle^{\frac{1}{2}}} \stackrel{=}{=} \frac{|\langle \mathbf{q}, \mathcal{B}\mathbf{v} \rangle|}{\|\mathbf{q}\|_p \|\mathbf{v}\|_u} \geq \min_{\mathbf{q} \neq \mathbf{1}} \max_{\mathbf{v} \neq \mathbf{0}} \frac{|\langle \mathbf{q}, \mathcal{B}\mathbf{v} \rangle|}{\|\mathbf{q}\|_p \|\mathbf{v}\|_u} \geq \gamma \frac{\|\mathbf{q}\|_0}{\|\mathbf{q}\|_p} \frac{\|\mathbf{v}\|_1}{\|\mathbf{v}\|_u}, \quad (3.78)$$

which once again leaves us with the task of proving equivalence of norms. For the velocity norms, we exploit Poincaré's inequality:

$$\begin{aligned} \|\mathbf{v}\|_1^2 &\geq \frac{1}{2c_p^2} \|\mathbf{v}\|_0^2 + \frac{1}{2} \|\mathbf{v}\|_1^2 = \frac{\Delta t}{2c_p^2} \frac{\langle \mathcal{M}_u \mathbf{v}, \mathbf{v} \rangle}{\Delta t} + \frac{1}{2\mu} \mu \langle \mathcal{K}_u \mathbf{v}, \mathbf{v} \rangle \\ &\geq \frac{1}{2} \min \left\{ \frac{\Delta t}{c_p^2}, \frac{1}{\mu} \right\} \|\mathbf{v}\|_u^2, \end{aligned} \quad (3.79)$$

where c_p is Poincaré's constant, such that $\|\mathbf{v}\|_0 \leq c_p \|\mathbf{v}\|_1$. For the pressure norms we follow a similar derivation as for the upper bound (3.75), and recover

$$\frac{\|\mathbf{q}\|_0^2}{\|\mathbf{q}\|_p^2} \geq \underline{c}_{\mathcal{M}} \Delta x^2 \left(\frac{1}{\bar{c}_{\mathcal{K}}^p \Delta t} + \frac{\mu}{\bar{c}_{\mathcal{M}}^p \Delta x^2} \right) = \frac{\Delta x^2}{\Delta t} \frac{\underline{c}_{\mathcal{M}}^p}{\bar{c}_{\mathcal{K}}^p} + \mu \frac{\underline{c}_{\mathcal{M}}^p}{\bar{c}_{\mathcal{M}}^p}. \quad (3.80)$$

Combining the two bounds eqs. (3.79) and (3.80), and substituting them in (3.78) provides a lower bound for (3.64):

$$\frac{\langle \mathcal{B}\mathcal{F}_u^{-1}\mathcal{B}^T \mathbf{q}, \mathbf{q} \rangle}{\langle \mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle} \geq \frac{\gamma^2}{2} \min \left\{ 1, \frac{\mu \Delta t}{c_p^2} \right\} \left(\frac{\Delta x^2}{\mu \Delta t} \frac{\underline{c}_{\mathcal{M}}^p}{\bar{c}_{\mathcal{K}}^p} + \frac{\underline{c}_{\mathcal{M}}^p}{\bar{c}_{\mathcal{M}}^p} \right) = \underline{\sigma}^2. \quad (3.81)$$

Collecting the results from eqs. (3.76) and (3.81), and assuming small μ and Δt , we get

$$\frac{\gamma^2}{2c_p^2} \left(\Delta x^2 \frac{c_M^p}{c_K^p} + \mu \Delta t \frac{c_M^p}{c_M^p} \right) \leq \frac{\langle \mathcal{B} \mathcal{F}_u^{-1} \mathcal{B}^T \mathbf{q}, \mathbf{q} \rangle}{\langle \mathcal{K}_p \mathcal{F}_p^{-1} \mathcal{M}_p \mathbf{q}, \mathbf{q} \rangle} \leq \frac{1}{\mu \Delta t} \frac{c_M^p}{c_K^p} + \frac{c_M^p}{c_M^p}. \quad (3.82)$$

Unfortunately, formula (3.82) fails to provide a useful bound as $\Delta x, \Delta t \rightarrow 0$, since the lower side approaches 0 and the upper bound diverges to ∞ . Even so, we can still provide a numerical indication of the clustering of the eigenvalues of the preconditioned system. Plots of the eigenvalues of (3.58) in the complex plane are shown in Fig. 3.1, for varying values of Δt , Δx , and intensity of the advection field $\mathbf{w}(\mathbf{x}, t)$ in (3.1) (the latter is regulated by the Péclet number, Pe: see (3.86) for its definition). We see that in all cases the eigenvalues are nicely bounded, they remain away from the origin, and are largely independent of variations in the discretisation parameters, providing semi-rigorous theoretical support for the proposed preconditioner.

It is worth pointing out that increasing the Péclet number has the effect of pushing some eigenvalues towards 2, and spreading more of them away from the real axis. It also makes the distributions more sensitive with respect to the level of refinement in the spatial discretisation, as can be seen comparing top and bottom plots in Fig. 3.1: the dark blue tokens distribute themselves more neatly in an arc as Δx decreases. These considerations make us wary of a possible degradation in the performance of the preconditioner as advection becomes predominant, an issue that is observed and discussed in Sec. 3.4.2.2. In general, however, most of the eigenvalues remain clustered inside $[0.1, 2] \times [-0.6, 0.6]$ in the complex plane.

3.4 Results

To demonstrate the performance of the preconditioner introduced in Sec. 3.2, we consider its application for the solution of a variety of flow configurations. For simplicity, we focus on problems defined on $\Omega \subset \mathbb{R}^2$, although we point out that the applicability of (3.23) is by no means limited to 2D simulations. The test-cases hereby proposed have been adapted from problems used extensively in the literature (see for example [40, Chap. 3.1 and Chap. 6.1]), and are briefly described in Sec. 3.4.1. The experiments conducted provide evidence for the optimal scalability of (3.23), as well as its potential for speedup via time parallelisation. Additionally, we investigate simple extensions to more general frameworks, such as the solution of nonlinear incompressible flow, and its applicability in advection-dominated regimes.

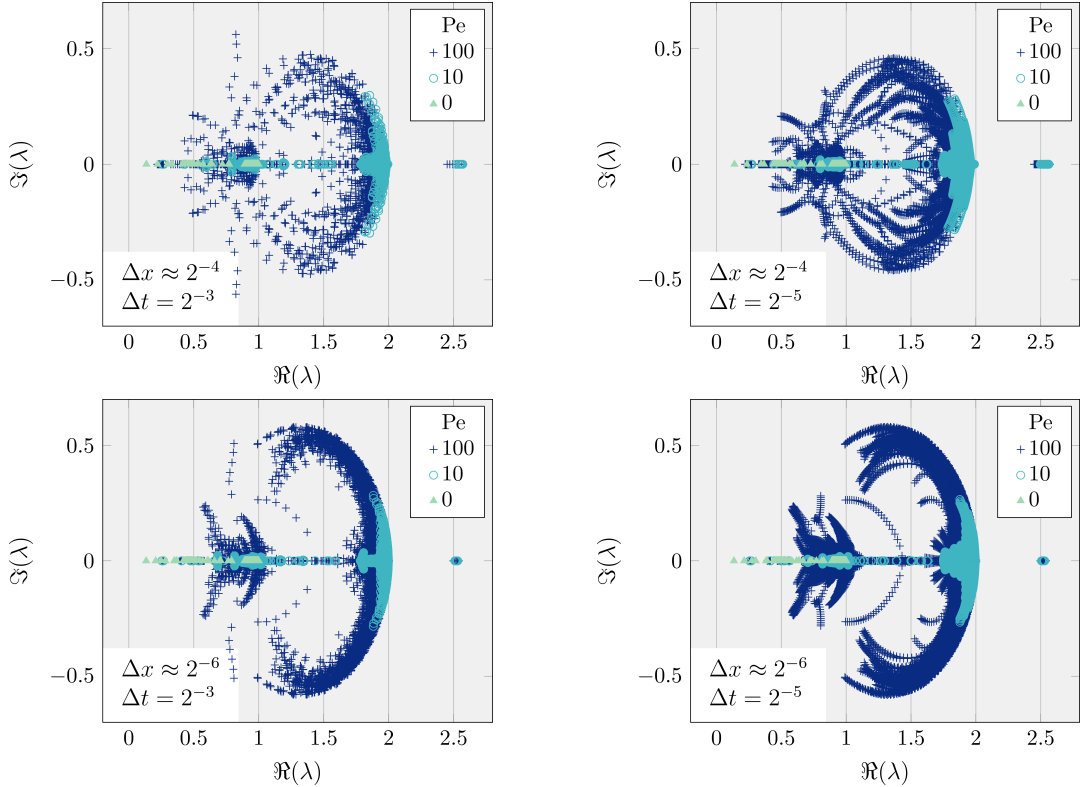


Figure 3.1: Eigenvalue distribution of the operators (3.58), corresponding to problem 4, for varying values of Δt , Δx , and Péclet number, Pe , which represents the intensity of the advection field. Here, Pe denotes the continuous Péclet number (i.e., not scaled by mesh spacing); see Sec. 3.4.1 for its definition (notice choosing $Pe = 0$ corresponds to problem 1). The eigenvalues are computed numerically using the MATLAB function `eigs`.

3.4.1 Model problems

The first three problems described in this section are test-cases for the time-dependent Stokes equations (3.2), while the last one is designed for Oseen (3.1). They are defined as follows.

Problem 1 *Driven cavity flow*. This represents a fully enclosed flow, defined on a square spatial domain $\mathbf{x} = (x, y)^T \in [0, 1]^2 =: \Omega_{\square}$. No forcing term is considered, and homogeneous Dirichlet BC are imposed on the bottom, left and right sides of the domain. To accelerate the flow, we prescribe the velocity on the top side

$$\mathbf{u}_C(\mathbf{x}|_{y=1}, t) = 8tx(1-x)(2x^2 - 2x + 1). \quad (3.83)$$

This profile is regularised so as to smoothly match the BC on the left and right sides. Notice that, in order to introduce time-dependency, the velocity is ramped up linearly




Figure 3.2: Velocity magnitude and streamlines (left), and pressure contour plot (right) for an example solution of the driven cavity flow problem 1.

from a quiet state: this is done in the following test-cases as well. An example solution for this problem is shown in Fig. 3.2.

Problem 2 *Poiseuille flow*. Also defined on Ω_{\square} , this set-up models laminar flow along a channel, for which the following analytical solution can be recovered:

$$\mathbf{u}_P(\mathbf{x}, t) = 4t \begin{bmatrix} y(1-y) \\ 0 \end{bmatrix}, \quad p_P(\mathbf{x}, t) = 8(1-x). \quad (3.84)$$

The forcing term is defined in such a way to counterbalance the time-derivative: $\mathbf{f}_P = 4[y(1-y), 0]^T$. No-slip conditions are included in the top and bottom sides of the channel, and a parabolic inflow profile is prescribed at $x = 0$. At the outflow




Figure 3.3: Velocity magnitude and streamlines (left), and pressure contour plot (right) for the analytical solution of the Poiseuille flow problem 2.




Figure 3.4: Velocity magnitude and streamlines (top), and pressure contour plot (bottom) for an example solution of the flow over backward-facing step problem 3.

$x = 1$, homogeneous Neumann BC are imposed. The analytical solution is plotted in Fig. 3.3.

Problem 3 *Flow over backward-facing step.* This is another example of flow through a channel, except the fluid undergoes a sudden expansion. The domain is a stretched L-shape, $\Omega_L = [0, 8] \times [0, 1] \cup [1, 8] \times [0, -1]$, and the BC imposed are similar to those for Poiseuille (the inflow, where the parabolic velocity profile is prescribed, is on the side $x = 0$). As for problem 1, no forcing term is included; an example solution is provided in Fig. 3.4.

Problem 4 *Double-glazing.* The set-up for this problem is analogous to that of the driven cavity flow, except now the effect of a recirculating wind is included in the system:

$$\mathbf{w}_G(\mathbf{x}, t) = 2t \begin{bmatrix} -(2y - 1)(4x^2 - 4x + 1) \\ (2x - 1)(4y^2 - 4y + 1) \end{bmatrix} \mu \text{Pe}, \quad (3.85)$$

whose intensity can be adjusted by tweaking the value of the Péclet number Pe [118, Chap. 13.2]. This parameter describes the relative intensity of advective transport with respect to diffusive transport, and for a generic advection field $\mathbf{w}(\mathbf{x}, t)$ it can be defined as

$$\text{Pe} := \frac{LU}{\mu} = \frac{L}{\mu} \left(\frac{1}{T - T_0} \int_{T_0}^T \max_{\mathbf{x} \in \Omega} \mathbf{w}(\mathbf{x}, t) dt \right), \quad (3.86)$$




Figure 3.5: Velocity magnitude and streamlines (left), and pressure contour plot (right) for an example solution of the double-glazing flow problem 4 with $Pe = 10$.

where L is a characteristic length of the spatial domain ($L = 1$ for this problem), and U a characteristic speed (considered averaged over the temporal domain). The impact that the added advection field has on the solution is illustrated in Fig. 3.5.

For each test-case the temporal domain is given by $t \in [0, 1]$, and the viscosity parameter is fixed to $\mu = 1$. The discrete functional spaces are approximated using Taylor-Hood finite elements: P_2 for velocity and P_1 for pressure [118, Chap. 17.4], both defined on a triangular mesh. The code developed for our experiments is publicly available at the repository [29]. It makes use of the MFEM finite element package [102] for the assembly of the finite element matrices, and we rely on the PETSc [7] and *hypr* [81] implementations of the various solvers used. Unless otherwise specified, we prescribe a null space-time initial guess for the GMRES iterations.

3.4.2 Performance of preconditioner

In this section, we conduct a series of experiments aimed at measuring the effectiveness of the proposed preconditioner (3.23) with GMRES. We chose the total number of iterations to convergence as the main measure of performance. The reason behind this preference (rather than, for example, computational time), lies in the fact that we believe it provides a more indicative metric for assessing the efficacy of our preconditioning strategy. This is in fact less dependent on the details of the actual implementation of the underlying solvers, and particularly of the one for the space-time velocity block (3.22), the optimal design of which is beyond the purpose of this

work: indeed a key point of our approach is the flexibility it provides in choosing such a solver.

3.4.2.1 Exact *versus* approximate solvers

As a first experiment, we directly test the application of our preconditioner to the solution of the model problems introduced in Sec. 3.4.1. We provide two different set-ups for the solvers involved in the application of (3.23): an *ideal* one, in which we make use of exact solvers for the relevant blocks, in order to provide a best-case scenario for the performance of such preconditioner; and an *approximate* case, which instead gives us a measure of the performance degradation we can expect by applying iterative solvers instead of exact solvers to approximate the solution of the relevant systems, which is more realistic in practice.

In more detail, the application of the inverse of (3.23) involves inverting three different operators: two of these are associated with the application of the approximate space-time Schur complement (3.37), and require inverting the pressure mass matrix \mathcal{M}_p (3.29a), and the pressure “Laplacian” $\tilde{\mathcal{K}}_p$ (3.35); the last one must deal with the monolithic space-time matrix for the velocity variable $F_{\mathbf{u}}$ (3.22). In the ideal case, we exploit LU solvers for both (3.29a) and (3.35), and a sequential time-stepping procedure for (3.22) (the latter further requires inverting the spatial velocity operator $\mathcal{F}_{\mathbf{u},k}$ (3.17) at each time-step k , for which we also use LU).

However, resorting to direct methods becomes infeasible when solving systems of a large size, both in terms of memory and of computations required. Moreover, our final goal is to actually parallelise in time the solution of the monolithic system. In order to do so, we must substitute the time-stepping procedure for the space-time velocity block $F_{\mathbf{u}}$ with a parallel alternative. Here, we use GMRES preconditioned by *hypre*’s implementation of nonsymmetric AMG based on approximate ideal restriction (AIR) [99,100], which we introduced in Sec. 2.2.4. Although likely not as effective as space-time geometric multigrid techniques (e.g., [77]) for diffusive problems, AIR is likely more robust for advective regimes and also less intrusive to apply, as we simply need to form the space-time velocity matrix and pass it to *hypre*. Due to the use of an inner Krylov method, we swap the outer GMRES solver with a *flexible*-GMRES (FGMRES) solver [125]. The inversions of the pressure mass and stiffness matrices are also approximated: for \mathcal{M}_p , we apply a fixed number of Chebyshev iterations [154], making use of the bounds on its eigenvalues provided by [155]; for $\tilde{\mathcal{K}}_p$, we apply *hypre*’s BoomerAMG [121].

Table 3.1: Number of GMRES iterations to 10^{-10} relative residual tolerance, when applied to the monolithic space-time system (3.21) and right-preconditioned with P_T^{-1} (3.23). Different levels of refinement are considered both for the spatial (rows) and temporal meshes (columns). The four sets of results refer to the four test problems introduced in Sec. 3.4.1 (problem 4 has Péclet number $Pe = 10$). Results on the left of each column are recovered using exact solvers, and results on the right of each column (in brackets) are recovered using approximate iterative solvers (inversion of \mathcal{M}_p is approximated with 8 Jacobi-preconditioned Chebyshev iterations, inversion of \mathcal{K}_p is approximated using 15 iterations of AMG, and the solution of the space-time velocity block F_u is recovered applying 15 iterations of GMRES preconditioned by AIR). The outer solver in the latter case is FGMRES. Crosses identify simulations that could not be completed due to memory requirements becoming too severe.

Pb	$\frac{\Delta t \rightarrow}{\Delta x \downarrow}$	2^{-1}		2^{-2}		2^{-3}		2^{-4}		2^{-5}		2^{-6}		2^{-7}	
1	2^{-2}	23	(23)	24	(24)	26	(26)	27	(27)	29	(29)	30	(30)	31	(31)
	2^{-3}	22	(22)	23	(23)	25	(25)	26	(29)	27	(29)	29	(29)	30	(30)
	2^{-4}	22	(22)	23	(23)	24	(24)	25	(25)	26	(26)	28	(38)	29	(36)
	2^{-5}	20	(20)	21	(21)	22	(22)	23	(23)	25	(25)	26	(36)	28	(30)
	2^{-6}	19	(19)	20	(20)	21	(21)	22	(21)	22	(22)	23	(24)	25	(26)
	2^{-7}	18	(18)	19	(19)	20	(20)	21	(20)	21	(21)	22	(24)	24	(25)
	2^{-8}	17	(17)	18	(18)	19	(19)	19	(19)	20	(20)	22	(23)	23	(25)
	2	2^{-2}	28	(28)	32	(32)	34	(34)	36	(36)	40	(40)	43	(43)	49
2^{-3}		31	(31)	34	(34)	35	(35)	36	(41)	38	(42)	39	(38)	38	(39)
2^{-4}		30	(30)	32	(32)	33	(33)	34	(34)	34	(36)	35	(49)	35	(44)
2^{-5}		29	(29)	31	(30)	32	(32)	33	(33)	34	(35)	34	(49)	34	(37)
2^{-6}		28	(28)	30	(30)	31	(31)	32	(32)	32	(32)	33	(34)	31	(32)
2^{-7}		27	(27)	29	(28)	30	(30)	30	(30)	30	(30)	29	(30)	29	(30)
2^{-8}		25	(25)	26	(26)	27	(26)	28	(27)	26	(26)	26	(27)	26	(27)
3		2^{-2}	33	(34)	35	(40)	37	(42)	38	(38)	42	(42)	46	(46)	53
	2^{-3}	33	(33)	35	(35)	36	(37)	37	(49)	39	(47)	39	(39)	40	(40)
	2^{-4}	31	(31)	33	(33)	34	(34)	35	(47)	35	(38)	36	(53)	37	(48)
	2^{-5}	30	(30)	32	(32)	33	(33)	34	(34)	34	(35)	36	(53)	35	(39)
	2^{-6}	28	(28)	31	(31)	32	(32)	32	(32)	33	(33)	33	(34)	33	(35)
	2^{-7}	27	(27)	29	(29)	30	(30)	31	(31)	31	(31)	30	(30)	30	(30)
	2^{-8}	25	(26)	27	(27)	28	(28)	28	(28)	28	(28)	26	(28)	×	(×)
	4	2^{-2}	25	(25)	28	(28)	30	(30)	31	(31)	32	(32)	31	(31)	32
2^{-3}		24	(24)	26	(26)	27	(27)	29	(32)	30	(32)	30	(30)	31	(31)
2^{-4}		24	(24)	25	(25)	26	(26)	27	(29)	28	(29)	29	(38)	31	(38)
2^{-5}		23	(24)	24	(26)	25	(27)	26	(28)	27	(30)	28	(40)	29	(32)
2^{-6}		21	(23)	22	(25)	23	(27)	24	(29)	25	(31)	26	(33)	26	(36)
2^{-7}		20	(24)	21	(27)	22	(29)	23	(32)	24	(35)	25	(38)	25	(42)
2^{-8}		19	(24)	20	(27)	21	(30)	22	(33)	23	(36)	23	(39)	24	(44)

Results from using direct or approximate solvers are reported side-by-side in Tab. 3.1, for ease of comparison. The two cases show a remarkably similar convergence behaviour, with the total number of iterations to convergence remaining substantially unchanged. The only notable disparities are observed for the most com-




Figure 3.6: Details on residual evolution of GMRES applied to problem 1, using P_T^{-1} with ideal components as a right-preconditioner. Results on the left refer to simulations with fixed $\Delta t = 2^{-7}$ and varying Δx ; results on the right are obtained by fixing $\Delta x \approx 2^{-8}$ and varying Δt . The set-up for the solvers is similar to the one used for Tab. 3.1.

plex problem 4, likely because the 15 AIR iterations used as an approximate solver for the space-time velocity block do not provide as accurate of an approximation as for problems 1 to 3. Even here, the increase in iterations compared with exact inner solves is nicely bounded, remaining below a factor of 2 with respect to the ideal case. This shows that the approximate solvers employed can be as effective as their direct counterparts, and gives an example of an effective time-parallel procedure which can work well in tandem with the preconditioner proposed.

We point out that, for the largest problems considered, the relevant systems have sizes $N_u = 526, 338$, $N_p = 66, 049$, and $N_t = 128$, for a total number of unknowns of $(N_u + N_p) \cdot N_t \approx 76 \cdot 10^6$; this further increases to $N_u = 5, 775, 362$, $N_p = 722, 945$, and $N_t = 128$ (totalling $\approx 832 \cdot 10^6$) for problem 3. Compared to the noticeable size of the system, the total number of iterations to convergence reported in Tab. 3.1 remains small, regardless of mesh size and problem type, providing evidence of the optimal scalability of the preconditioners proposed. This is further confirmed by tracking in more detail the evolution of the residual as the GMRES iterations progress: an example of this is provided in Fig. 3.6, which shows reasonably uniform convergence profiles, independently of the choice of both Δx and Δt .

3.4.2.2 Dependence on Péclet number

When applied to the solution of Oseen equations, we observe that the performance of the preconditioner degrades as we increase the intensity of the advection field

Table 3.2: Number of iterations to convergence for problem 4 with different values of Pe. Dashes represent simulations that did not converge in the maximum prescribed number of iterations (100). Same set-up as for the ideal case in Tab. 3.1.

$\frac{\Delta t \rightarrow}{\Delta x \downarrow}$	2^{-4}					2^{-5}					2^{-6}					2^{-7}				
Pe \rightarrow	2^4	2^5	2^6	2^7	2^8	2^4	2^5	2^6	2^7	2^8	2^4	2^5	2^6	2^7	2^8	2^4	2^5	2^6	2^7	2^8
2^{-2}	32	60	//	//	//	32	57	//	//	//	31	55	//	//	//	31	54	//	//	//
2^{-3}	28	35	54	//	//	28	35	53	//	//	28	35	53	//	//	27	35	52	//	//
2^{-4}	27	32	40	55	//	27	32	39	55	//	25	31	38	53	//	25	31	38	52	//
2^{-5}	25	30	35	45	66	25	28	35	44	66	24	28	34	44	64	23	27	34	43	63
2^{-6}	24	27	32	39	49	23	27	32	38	49	23	27	32	37	48	22	26	30	37	47
2^{-7}	23	26	31	35	40	22	26	30	35	41	22	26	29	33	41	21	25	29	33	40
2^{-8}	22	25	29	32	37	21	25	29	32	37	21	25	28	32	37	20	24	27	32	37

imposed, as foreshadowed by the analysis in Sec. 3.3. This is shown in Tab. 3.2, where we progressively ramp up the value of Pe in problem 4, solving for various Δx : convergence fails for coarser spatial meshes and largest Péclet numbers. This hints at the necessity of either employing appropriately refined meshes, or relying on stabilisation techniques [118, Chap. 13.8] (which we did not apply in our analysis), whenever advection-dominated flows need to be resolved. The preconditioner for the single time-step case (3.38) was shown to suffer from a similar lack of robustness [40, Table 9.3], so it is somewhat expected that this weakness carries over to the whole space-time case. However, we point out that the performance of the preconditioner does not vary significantly as long as the *grid* Péclet number, $\tilde{Pe} := \Delta x Pe / L$ [118, Chap. 13.2], is kept constant. Unlike its *global* counterpart (3.86), this parameter provides a *local* measure of the dominance of advection over diffusion, and is useful in identifying conditions for which instabilities and oscillations might arise in the numerical solution. This parameter is kept fixed down each diagonal of Tab. 3.2, which shows only slow growth in the number of iterations to convergence as \tilde{Pe} increases.

3.4.3 The nonlinear setting: incompressible Navier-Stokes

As stated in Sec. 3.1, Oseen equations (3.1) can be interpreted as a linearisation of the Navier-Stokes system (3.3). If we then introduce an *outer* solver to tackle the nonlinear advection term in Navier-Stokes, we can use our preconditioner to accelerate the solution of the corresponding *inner* (linearised) problems, much like what is done for the single time-step case in [40, Chap. 10]. We experiment on the performance of our preconditioner P_T (3.23) in this situation, and resolve the nonlinearity in (3.3)

via *Picard iterations* [115]. That is, we iteratively solve the linearised system

$$\left[\begin{array}{c|c} F_{\mathbf{u}}((\mathbf{u})^j) & B^T \\ \hline B & \mathbf{0} \end{array} \right] \begin{bmatrix} (\mathbf{u})^{j+1} \\ (\mathbf{p})^{j+1} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{0} \end{bmatrix}, \quad (3.87)$$

until convergence is reached. Here, $(\mathbf{u})^{j+1}$ and $(\mathbf{p})^{j+1}$ denote the whole space-time velocity and pressure solutions at the $j + 1$ -th Picard iteration, while we have made explicit the dependence of the space-time velocity block $F_{\mathbf{u}} = F_{\mathbf{u}}(\mathbf{w})$ on the advection field, which we substitute with the approximate solution at the previous iteration $(\mathbf{u})^j$.

Linear and nonlinear convergence obtained by applying a Picard linearisation (3.87) to the problems in Sec. 3.4.1 are reported in Tab. 3.3. Especially for problem 1, both the inner and outer solve iterations remain $\mathcal{O}(1)$, independently on the spatial and temporal mesh spacing. Similar results hold for the inflow-outflow problem problem 3, but here the average number of inner GMRES iterations is larger. In the specific case of a very large time-step relative to the spatial mesh, the number of inner iterations can be a fair amount larger. However, this is not necessarily surprising, as problem 3 offers more complex dynamical behaviour and we are resolving

Table 3.3: Convergence results for Navier-Stokes with the set-up of problems 1 and 3 (Reynolds number = 1 in both cases). Picard iterations were used as the outer nonlinear solver, which achieves convergence with a tolerance of 10^{-9} . Inside each column, the number of outer iterations to convergence is reported on the left, while the average number of inner GMRES iterations per outer Picard iteration appears on the right (in brackets). The inner GMRES solver uses the same set-up used for recovering the results in the left column of Tab. 3.1. Notice problem 2 is missing: this is because the nonlinearity does not affect the analytical solution for Poiseuille, and hence results are not particularly informative in this case (convergence occurs at the very first Picard iteration in almost all of the experiments). Also problem 4 is missing, since it already represents a linearisation of Navier-Stokes applied to problem 1.

Pb	$\frac{\Delta t \rightarrow}{\Delta x \downarrow}$	2 ⁻¹		2 ⁻²		2 ⁻³		2 ⁻⁴		2 ⁻⁵		2 ⁻⁶		2 ⁻⁷	
1	2 ⁻²	5	(11.6)	4	(14.75)	4	(14.25)	4	(14.50)	4	(13.50)	4	(13.25)	4	(13.0)
	2 ⁻³	4	(13.5)	4	(13.25)	4	(13.25)	4	(13.00)	4	(12.50)	4	(12.25)	5	(9.40)
	2 ⁻⁴	4	(12.5)	4	(12.50)	4	(12.25)	4	(12.00)	4	(11.50)	5	(8.800)	5	(8.60)
	2 ⁻⁵	4	(11.5)	4	(11.00)	4	(11.50)	4	(10.75)	5	(8.000)	5	(8.200)	4	(9.25)
	2 ⁻⁶	4	(10.5)	4	(10.00)	4	(10.00)	4	(9.500)	4	(9.250)	4	(8.500)	4	(8.75)
	2 ⁻⁷	4	(9.75)	4	(9.000)	4	(9.250)	5	(6.800)	4	(8.250)	4	(8.000)	4	(7.75)
	2 ⁻⁸	3	(11.0)	3	(11.00)	3	(10.67)	3	(10.33)	3	(10.33)	4	(7.250)	4	(6.75)
	3	2 ⁻²	5	(19.80)	5	(18.20)	5	(18.00)	5	(18.00)	5	(18.40)	5	(19.00)	5
2 ⁻³		5	(21.80)	5	(19.60)	5	(17.80)	5	(17.00)	5	(16.80)	5	(16.40)	5	(16.20)
2 ⁻⁴		5	(22.60)	5	(20.60)	5	(18.80)	5	(16.80)	5	(15.60)	5	(14.80)	5	(14.20)
2 ⁻⁵		5	(26.60)	5	(22.80)	5	(20.60)	5	(17.60)	5	(16.20)	5	(15.80)	5	(14.20)
2 ⁻⁶		4	(36.00)	4	(31.75)	4	(28.25)	5	(18.00)	5	(17.20)	5	(16.00)	5	(14.80)
2 ⁻⁷		4	(43.00)	4	(37.00)	4	(31.75)	4	(24.75)	5	(18.40)	5	(18.60)	5	(15.00)
2 ⁻⁸		4	(52.50)	4	(46.25)	4	(33.50)	4	(28.25)	5	(22.40)	×	(×)	×	(×)

the nonlinearity over a very large time-step. This likely results in a very stiff and ill-conditioned linear system to solve for each nonlinear iteration. Nevertheless, the outer space-time Picard linearisation remains perfectly scalable.

3.4.4 Comparison with sequential time-stepping

In Sec. 3.2.1.1 we explain how the two main factors in analysing the effectiveness of our preconditioning strategy compared with classical time-stepping are given by the ratios $C_{F_u}/(C_{F_u} N_t)$ and N_{it}^{ST}/N_{it}^0 . The first ratio directly refers to the efficiency of the method chosen for the parallel-in-time integration of the velocity block F_u : in this section we do not investigate the effectiveness of the various PinT methods available in the literature, but we refer to Chap. 2 for some examples on their typical performance, which is generally satisfactory for a parabolic problem such as the one considered here. Instead, we focus on the second ratio, since it is intrinsic to the preconditioning strategy proposed: in fact, it describes the overhead produced by tackling the whole space-time system at once, rather than via sequential time-stepping.

To give an indication of how N_{it}^0 and N_{it}^{ST} compare with each other when ideal solvers are employed, we collect convergence results from the application of (3.38) within the time-stepping procedure, for problems equivalent to those analysed in

Table 3.4: Ratio N_{it}^{ST}/N_{it}^0 , where N_{it}^{ST} is taken from the ideal case in Tab. 3.1, while N_{it}^0 represents the average number of iterations to convergence per time-step, solving for (3.18) via forward substitution. For the latter, GMRES is used at each time-step, right-preconditioned with the single time-step counterpart of P_T , (3.38) with (3.39). Convergence is reached with a tolerance of $10^{-10}/\sqrt{N_t}$. The remaining options for the solvers involved are the same ones used to fill the left column in Tab. 3.1.

Pb	$\frac{\Delta t \rightarrow}{\Delta x \downarrow}$	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	2^{-7}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	2^{-7}	Pb
1	2^{-2}	1.18	1.37	1.39	1.61	1.82	2.13	1.39	1.54	1.75	2.13	2.44	2.94	3
	2^{-3}	1.08	1.16	1.33	1.64	1.92	2.08	1.28	1.43	1.64	1.82	2.00	2.27	
	2^{-4}	1.08	1.22	1.41	1.56	1.79	2.00	1.20	1.32	1.43	1.59	1.79	2.08	
	2^{-5}	1.06	1.23	1.32	1.52	1.64	1.89	1.18	1.28	1.43	1.56	1.82	2.08	
	2^{-6}	1.08	1.18	1.25	1.32	1.47	1.67	1.18	1.27	1.37	1.56	1.85	2.27	
	2^{-7}	1.04	1.14	1.20	1.28	1.41	1.67	1.15	1.27	1.45	1.69	1.96	2.33	
	2^{-8}	1.03	1.11	1.15	1.28	1.56	1.82	1.16	1.30	1.52	1.75	1.92	×	
2	2^{-2}	1.43	1.61	1.82	2.13	2.50	3.13	1.28	1.43	1.54	1.69	1.82	2.13	4
	2^{-3}	1.35	1.43	1.61	1.82	2.08	2.27	1.14	1.20	1.39	1.69	1.85	2.08	
	2^{-4}	1.19	1.25	1.35	1.49	1.72	1.96	1.12	1.20	1.41	1.56	1.72	1.89	
	2^{-5}	1.15	1.19	1.28	1.47	1.67	2.00	1.19	1.27	1.37	1.49	1.59	1.69	
	2^{-6}	1.10	1.18	1.28	1.45	1.82	2.17	1.10	1.15	1.27	1.39	1.45	1.59	
	2^{-7}	1.10	1.20	1.30	1.56	1.89	2.44	1.06	1.10	1.22	1.35	1.49	1.59	
	2^{-8}	1.08	1.19	1.43	1.59	1.96	2.56	1.05	1.10	1.22	1.37	1.47	1.72	

Tab. 3.1. To make the comparison more fair, for the single time-step we ask for a stricter tolerance on convergence, scaled by a factor $\sqrt{N_t}$ with respect to the whole space-time case: in this way, the Euclidean norm of the final space-time residual is similar in the two cases. We report the resulting ratio N_{it}^{ST}/N_{it}^0 in Tab. 3.4. We see that this sits between roughly 1 and 3, in most cases closer to one, but becoming larger as $\Delta t \rightarrow 0$, particularly $\Delta t \ll \Delta x$. This observation can be explained by pointing out that a large advantage of time-stepping is given by the fact that, at each iteration, a good initial guess is provided by the value of the solution at the previous time-step, which we exploit in our experiments: with a reduced step-size, it is safe to expect that also the solution varies less between one time-step and the next, improving the quality of such initial guess. Nevertheless, results demonstrate little-to-no degradation in convergence speed between the two approaches and, thus, minimal overhead in terms of iteration counts when applying block preconditioning techniques in an all-at-once fashion, when compared with sequential time-stepping.

3.5 Remarks on STBP for incompressible flow

The space-time preconditioning approach pursued in this chapter exploits the block structure of the target system *at the space-time level* in order to simplify the type of problem we need to (parallel-)integrate in time. The most demanding task of the resulting procedure reduces in fact to the solution of a single-variable, time-dependent advection-diffusion equation. This approach provides a different take on time parallelisation for incompressible flows, compared with what has typically been done in the literature [23, 50, 103, 148, 149]. To our knowledge, most of the latter relies on algorithms such as Parareal, which leave the underlying time-stepping procedure for solving (3.18) largely unchanged, and rather focus on building an overarching framework which breaks its innate sequentiality.

A number of arguments can be made for supporting the latter approach: most noticeably, time parallelisation is a relatively new concept, largely seen as an extra feature to include when one desires to gain some additional concurrency. In this sense, relying on a non-intrusive PinT algorithm has the appealing advantage of allowing to reuse already-existing time-integration schemes, out of the box. This is especially true in the case of complex systems of PDEs, for which time-stepping in itself can be a challenging task, often requiring the design of solvers *ad hoc*. Having to reinvent such solvers from scratch in order to achieve time parallelisation would then be detrimental.

Fortunately, the analysis conducted in this chapter provides evidence for stating that this needs not be the case. In fact, in Sec. 3.2.1 and 3.2.2 we have shown that, at least for the model problem considered, the very same concepts that guided the design of a block preconditioner for the single time-step case can be seamlessly extended to the space-time framework, achieving similar advantages and with a marginal overhead, as detailed in Sec. 3.4 and more specifically in Sec. 3.4.4. Indeed an interesting analogy can be drawn between the role of the preconditioner for the single time-step case (3.38) and our space-time preconditioner (3.23): the first provides an effective method for solving Oseen via sequential time stepping, so long as a fast solver for an elliptic operator is available; the second provides an effective method for solving Oseen all-at-once, so long as a fast (time-parallel) solver for a *parabolic* operator is available (and research on PinT schemes for this class of problems is already relatively mature: see for example [48, 55, 77, 138], as well as Chap. 2).

In the following chapter, we investigate how the concept of space-time block preconditioning can be extended to a more complicated problem, representing the target application of our project: magnetohydrodynamics.

Chapter 4

Applications to incompressible magnetohydrodynamics

Lying at the interface between the fields of electromagnetism and fluid dynamics, *Magnetohydrodynamics* (MHD) is responsible of modelling the behaviour of electrically charged fluids when subject to electromagnetic interactions. MHD models are used to describe a variety of fascinating phenomena, such as galaxy formation, the generation of solar flares, as well as their interaction with the magnetosphere of the Earth; our interest lies in more technical applications, linked to the development of effective magnetic confinement for the plasma within a fusion reactor. More specifically, our target model considers *resistive* plasmas: in this framework, it is possible for the topology of the magnetic field lines to change, and phenomena in which these break or coalescence can occur [66]. This represents a problem in the development of fusion reactors, as during these events some matter might escape the magnetic confinement, and end up colliding with the interior walls of the vessel, thus damaging them, losing heat in the process (see also Fig. 1.2), and ultimately undermining the possibility of sustaining a reaction for long periods of time. Due to the complexity and cost associated with setting up large-scale experiments, fusion researchers at our industrial partner institution, CCFE, rely extensively on the results from numerical simulations in guiding their design choices for properly controlling the reaction.

The corresponding system of PDEs is tightly coupled, non-symmetric, and presents nonlinearities which give rise to turbulent behaviour and spawn a range of phenomena interacting at various scales, both temporal and spatial: these are all characteristics which make its approximate solution particularly challenging. Nonetheless, there is great interest in developing strategies for tackling these problems effectively, as testified by the amount of research dedicated to the design of stable and efficient schemes and preconditioners for accelerating the solution of the discretised MHD system: see

for example [1, 15, 21, 88, 114, 134, 135], and the work this chapter mostly relies on, [25]. Given the difficulties already posed by having to solve this problem in a time-stepping framework, perhaps it is not surprising to see that attempts to include time parallelisation into this task are rare and far-between. To our knowledge, only [8, 119, 129] and more recently [128], have tested the application of a parallel-in-time algorithm (specifically, Parareal from Sec. 2.1.1) to MHD problems, showing modest speedups of 8-10 \times using hundreds of processors. It is precisely the complexity of the equations involved, however, that makes the development of effective parallelisation techniques to reduce time-to-solution particularly appealing.

In this chapter we describe how the concept of space-time block preconditioning developed in Chap. 3 can be extended to the framework of resistive, incompressible MHD, thus building evidence supporting the flexibility and broad spectrum of applicability of this technology, and showcasing an alternative approach to time parallelisation for plasma simulation. In designing our space-time block preconditioner for MHD, we follow similar steps as in the previous chapter, and start by identifying a suitable single time-step counterpart. The choice falls on the *Continuous Schur Complement preconditioner* introduced in [25, (3.32)]: its particularity is that it relies on the very same PCD operator (3.38) which is the foundation underpinning our STBP preconditioner for incompressible flow derived in Sec. 3.2.1. As such, when considering its extension to the whole space-time framework, we can directly reuse most of the technology developed in the previous chapter. We measure the performance of the preconditioner when applied to the space-time solution of some MHD model problems. The convergence results are in large part still preliminary, and additional work is required in order to perfect the final definition of the space-time preconditioner, but they already hint at its scalability and show that it compares well with its single time-step counterpart.

As the work described here represents a direct extension of the technique developed in Chap. 3, the structure of this chapter closely resembles that of the previous one. In Sec. 4.1 we introduce the resistive incompressible magnetohydrodynamic system used for our experiments, together with a brief derivation of the model and a justification for the choice of the specific formulation employed, and cover details regarding its space-time discretisation. Our preconditioning strategy is illustrated in Sec. 4.2, where we show how to extend the work done in [25] from the single time-step to the whole space-time setting. Finally, some results on the effectiveness of the preconditioner, both from a theoretical and an experimental point of view, are presented and discussed in Sec. 4.3 and Sec. 4.4.

4.1 Problem definition

The system we are interested in solving can be interpreted as an extension of the Navier-Stokes equations introduced in (3.3), which takes into account the interaction between flow and a magnetic field [66, Chap. 6]. The model is provided by

$$\left\{ \begin{array}{l} \frac{\partial \mathbf{u}}{\partial t}(\mathbf{x}, t) + (\mathbf{u}(\mathbf{x}, t) \cdot \nabla) \mathbf{u}(\mathbf{x}, t) - \mu \nabla^2 \mathbf{u}(\mathbf{x}, t) \\ \quad + \nabla p(\mathbf{x}, t) - \nabla \cdot \mathcal{T}_M(\mathbf{B}(\mathbf{x}, t)) = \mathbf{f}(\mathbf{x}, t) \\ \quad \quad \quad - \nabla \cdot \mathbf{u}(\mathbf{x}, t) = 0 \\ \frac{\partial \mathbf{B}(\mathbf{x}, t)}{\partial t} - \nabla \times (\mathbf{u}(\mathbf{x}, t) \times \mathbf{B}(\mathbf{x}, t)) \\ \quad + \frac{\eta}{\mu_0} \nabla \times (\nabla \times \mathbf{B}(\mathbf{x}, t)) = \mathbf{0} \end{array} \right. \quad \text{in } \Omega \times (T_0, T], \quad (4.1)$$

where η is the electric resistivity, and μ_0 is the magnetic permeability in the void, both considered constant in $\Omega \subset \mathbb{R}^2$. Compared to (3.3), the momentum equation here contains an extra term,

$$\mathcal{T}_M(\mathbf{B}) = \frac{1}{\mu_0} \mathbf{B} \otimes \mathbf{B} - \frac{1}{2\mu_0} \|\mathbf{B}\|^2 I, \quad (4.2)$$

representing the *magnetic stress tensor*, which models the action of the force exerted by the magnetic field onto the flow, due to the charged nature of the medium (I is the 2×2 identity matrix). On top of this, the additional equation in the system models the behaviour of the magnetic field $\mathbf{B}(\mathbf{x}, t)$: in particular, notice how the flow is responsible for its advection, thus rendering the system fully coupled. This last PDE is derived using some well-known formulas for electromagnetism: see for example [83, Chap. 5-6] for their description. Specifically, we start from Maxwell-Faraday equation,

$$\frac{\partial \mathbf{B}(\mathbf{x}, t)}{\partial t} = -\nabla \times \mathbf{E}(\mathbf{x}, t), \quad (4.3)$$

which links the magnetic field to the electric field $\mathbf{E}(\mathbf{x}, t)$. We then substitute the latter by relying on Ohm's law,

$$\mathbf{E}(\mathbf{x}, t) = -\mathbf{u}(\mathbf{x}, t) \times \mathbf{B}(\mathbf{x}, t) + \eta \mathbf{j}(\mathbf{x}, t), \quad (4.4)$$

where we introduce the current $\mathbf{j}(\mathbf{x}, t)$. This in turn is related back to the magnetic field via Ampère's law

$$\mathbf{j}(\mathbf{x}, t) = \frac{1}{\mu_0} \nabla \times \mathbf{B}(\mathbf{x}, t), \quad (4.5)$$

and we can see that following this chain of substitutions results directly in the last equation of (4.1). The *resistive* nature of the plasma is dictated by the presence of

the current term in (4.4); by contrast, *ideal* plasmas are perfect conductors, and for those we would have $\eta = 0$: see also [66, Chap. 6.5] and [83, Chap. 10.2].

Notice that, by taking the divergence of the magnetic field equation, we recover

$$\frac{\partial(\nabla \cdot \mathbf{B}(\mathbf{x}, t))}{\partial t} = 0, \quad (4.6)$$

as the curl terms cancel out, since $\nabla \cdot (\nabla \times (*)) = 0$. Relationship (4.6) implies that the solenoidality of the magnetic field is maintained throughout the temporal evolution, provided $\mathbf{B}(\mathbf{x}, t)$ is initially solenoidal. We can use this property to rewrite the divergence of the magnetic stress tensor in a more familiar way:

$$-\nabla \cdot \mathcal{T}_M(\mathbf{B}) = -\frac{1}{\mu_0} \mathbf{B}(\nabla \cdot \mathbf{B}) - (\mathbf{B} \cdot \nabla) \mathbf{B} + \frac{1}{2\mu_0} \nabla(\|\mathbf{B}\|^2) = \frac{1}{\mu_0} \mathbf{B} \times (\nabla \times \mathbf{B}), \quad (4.7)$$

which represents the action of the *Lorentz force*. In practice, however, ensuring that the magnetic field remains solenoidal when solving (4.1) numerically might pose some difficulties. One possible workaround consists in enforcing the constraint $\nabla \cdot \mathbf{B} = 0$ weakly using a Lagrangian multiplier [133], thus augmenting (4.1) with an additional variable which covers a similar role for the magnetic field as the pressure does for the flow. An alternative way to enforce solenoidality of $\mathbf{B}(\mathbf{x}, t)$ consists in imposing

$$\mathbf{B}(\mathbf{x}, t) = \nabla \times \mathbf{A}(\mathbf{x}, t), \quad (4.8)$$

thus introducing the vector potential $\mathbf{A}(\mathbf{x}, t)$ for the magnetic field. In our analysis we pursue this latter approach, both because this follows the work done in [25], which our analysis heavily relies upon, and because it allows for some simplifications when dealing with 2D problems, as explained next.

In order to reformulate the system (4.1) in terms of \mathbf{A} , we employ the following relationship between the vector potential and the electric field

$$\mathbf{E}(\mathbf{x}, t) = -\nabla\phi(\mathbf{x}, t) - \frac{\partial\mathbf{A}(\mathbf{x}, t)}{\partial t}, \quad (4.9)$$

which holds for a given scalar electric potential $\phi(\mathbf{x}, t)$. Substituting eqs. (4.5), (4.8) and (4.9) into (4.4), provides the following PDE:

$$\frac{\partial\mathbf{A}(\mathbf{x}, t)}{\partial t} - \mathbf{u}(\mathbf{x}, t) \times (\nabla \times \mathbf{A}(\mathbf{x}, t)) + \frac{\eta}{\mu_0} \nabla \times (\nabla \times \mathbf{A}(\mathbf{x}, t)) + \nabla\phi(\mathbf{x}, t) = \mathbf{0}. \quad (4.10)$$

Given its definition in (4.8), we notice that adding a curl-free perturbation to $\mathbf{A}(\mathbf{x}, t)$ does not modify the resulting magnetic field $\mathbf{B}(\mathbf{x}, t)$. We can exploit this fact to get rid of the additional variable $\phi(\mathbf{x}, t)$: we do so by re-defining

$$\tilde{\mathbf{A}}(\mathbf{x}, t) = \mathbf{A}(\mathbf{x}, t) + \int_{T_0}^T \nabla\phi(\mathbf{x}, t) dt; \quad (4.11)$$

notice $\mathbf{B}(\mathbf{x}, t) = \nabla \times \tilde{\mathbf{A}}(\mathbf{x}, t) \equiv \nabla \times \mathbf{A}(\mathbf{x}, t)$, since $\nabla \times (\nabla \cdot *) = \mathbf{0}$. From here on, with a slight abuse of notation, we drop the tilde on $\tilde{\mathbf{A}}(\mathbf{x}, t)$ and refer to this modified vector potential instead. Substituting the new definition (4.11) in (4.10) simplifies the latter equation to

$$\frac{\partial \mathbf{A}(\mathbf{x}, t)}{\partial t} - \mathbf{u}(\mathbf{x}, t) \times (\nabla \times \mathbf{A}(\mathbf{x}, t)) + \frac{\eta}{\mu_0} \nabla \times (\nabla \times \mathbf{A}(\mathbf{x}, t)) = \mathbf{0}. \quad (4.12)$$

Still, $\mathbf{A}(\mathbf{x}, t)$ remains defined up to a curl-free function: to fix it once and for all¹, we impose the extra condition $\nabla \cdot \mathbf{A} = 0$. Using the identity $\nabla \times (\nabla \times *) = \nabla(\nabla \cdot *) - \nabla^2(*)$, this allows us to rewrite

$$\frac{\partial \mathbf{A}(\mathbf{x}, t)}{\partial t} - \mathbf{u}(\mathbf{x}, t) \times (\nabla \times \mathbf{A}(\mathbf{x}, t)) - \frac{\eta}{\mu_0} \nabla^2 \mathbf{A}(\mathbf{x}, t) = \mathbf{0}. \quad (4.13)$$

When limiting ourselves to 2D geometries, it suffices to track the z -component of the vector potential,

$$\mathbf{A}(\mathbf{x}, t) \underset{(\text{if } \mathbf{x} \in \mathbb{R}^2)}{=} \begin{bmatrix} 0 \\ 0 \\ A(\mathbf{x}, t) \end{bmatrix} = A(\mathbf{x}, t) \hat{\mathbf{k}} \implies \mathbf{B}(\mathbf{x}, t) = \begin{bmatrix} \frac{\partial A}{\partial y}(\mathbf{x}, t) \\ -\frac{\partial A}{\partial x}(\mathbf{x}, t) \\ 0 \end{bmatrix}, \quad (4.14)$$

($\hat{\mathbf{k}}$ represents the unit vector in the z -direction), which allows us to effectively turn equation (4.13) into a scalar one:

$$\frac{\partial A(\mathbf{x}, t)}{\partial t} + \mathbf{u}(\mathbf{x}, t) \cdot \nabla A(\mathbf{x}, t) - \frac{\eta}{\mu_0} \nabla^2 A(\mathbf{x}, t) = 0. \quad (4.15)$$

To fully express system (4.1) in terms of $A(\mathbf{x}, t)$, we also need to rewrite the Lorentz force term (4.7). This is done as follows:

$$\begin{aligned} \frac{1}{\mu_0} \mathbf{B}(\mathbf{x}, t) \times (\nabla \times \mathbf{B}(\mathbf{x}, t)) &= \frac{1}{\mu_0} (\nabla \times (A(\mathbf{x}, t) \hat{\mathbf{k}})) \times (\nabla \times \nabla \times (A(\mathbf{x}, t) \hat{\mathbf{k}})) \\ &= \frac{1}{\mu_0} (\nabla A(\mathbf{x}, t) \times \hat{\mathbf{k}}) \times (-\hat{\mathbf{k}} \nabla^2 A(\mathbf{x}, t)) \\ &= \frac{1}{\mu_0} \nabla A(\mathbf{x}, t) \nabla^2 A(\mathbf{x}, t). \end{aligned} \quad (4.16)$$

Changing (4.7) to (4.16), and substituting the magnetic field equation with (4.15),

¹ The procedure for fixing the choice of the vector potential goes under the name of *gauge fixing*. There is no unique way to achieve this: our approach, where we request $\mathbf{A}(\mathbf{x}, t)$ to have null divergence, corresponds to imposing the *Coulomb gauge fixing condition* [83, Chap. 6.5].

we modify (4.1) to

$$\left\{ \begin{array}{l} \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial t} + (\mathbf{u}(\mathbf{x}, t) \cdot \nabla) \mathbf{u}(\mathbf{x}, t) - \mu \nabla^2 \mathbf{u}(\mathbf{x}, t) \\ \quad + \nabla p(\mathbf{x}, t) + \frac{1}{\mu_0} \nabla A(\mathbf{x}, t) \nabla^2 A(\mathbf{x}, t) = \mathbf{f}(\mathbf{x}, t) \\ \quad - \nabla \cdot \mathbf{u}(\mathbf{x}, t) = 0 \\ \frac{\partial A(\mathbf{x}, t)}{\partial t} + \mathbf{u}(\mathbf{x}, t) \cdot \nabla A(\mathbf{x}, t) - \frac{\eta}{\mu_0} \nabla^2 A(\mathbf{x}, t) = 0 \end{array} \right. \quad \text{in } \Omega \times (T_0, T], \quad (4.17)$$

Under this formulation, the Lorentz term presents some high-order derivatives (on top of the nonlinearity), which might make its treatment difficult depending on the chosen discretisation. We tackle this issue by introducing the current as an auxiliary variable: in the 2D case, this reduces to

$$\begin{aligned} \mathbf{j}(\mathbf{x}, t) &= \frac{1}{\mu_0} \nabla \times (\nabla \times \mathbf{A}(\mathbf{x}, t)) \quad \left(\begin{array}{l} = j(\mathbf{x}, t) \hat{\mathbf{k}}, \\ \text{(if } \mathbf{x} \in \mathbb{R}^2) \\ \nabla \cdot \mathbf{A} \equiv 0 \end{array} \right) \quad \text{with} \\ j(\mathbf{x}, t) &= \frac{1}{\mu_0} \nabla^2 A(\mathbf{x}, t); \end{aligned} \quad (4.18)$$

(notice we used (4.5) once again). We are now ready to write our target system of equations:

$$\left\{ \begin{array}{l} \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial t} + (\mathbf{u}(\mathbf{x}, t) \cdot \nabla) \mathbf{u}(\mathbf{x}, t) - \mu \nabla^2 \mathbf{u}(\mathbf{x}, t) \\ \quad + \nabla p(\mathbf{x}, t) + \nabla A(\mathbf{x}, t) j(\mathbf{x}, t) = \mathbf{f}(\mathbf{x}, t) \\ \quad - \nabla \cdot \mathbf{u}(\mathbf{x}, t) = 0 \\ \quad j(\mathbf{x}, t) - \frac{1}{\mu_0} \nabla^2 A(\mathbf{x}, t) = 0 \\ \frac{\partial A(\mathbf{x}, t)}{\partial t} + \mathbf{u}(\mathbf{x}, t) \cdot \nabla A(\mathbf{x}, t) - \frac{\eta}{\mu_0} \nabla^2 A(\mathbf{x}, t) = -E(\mathbf{x}, t) \end{array} \right. \quad \text{in } \Omega \times (T_0, T], \quad (4.19)$$

where we also introduced a right-hand side into the vector potential equation²: $E(\mathbf{x}, t)$ represents the intensity of an applied external electric field (along the z -direction).

Next, we proceed to close the system with opportune initial and boundary conditions. For those on velocity, we reuse much of the notation introduced in Sec. 3.1, and particularly (3.5). Similarly, the initial condition prescribed for the vector potential is in the form

$$A(\mathbf{x}, T_0) = \bar{A}^0(\mathbf{x}) \quad \text{on } \Omega, \quad (4.20)$$

²Notice we could have used (4.18) to substitute $j(\mathbf{x}, t)$ into the equation for the vector potential. Not applying this substitution, however, has the advantage of giving the 2×2 block in the bottom-right of (4.26) an upper-triangular structure, with a standard parabolic operator in its diagonal, which makes it simpler to treat.

while its boundary conditions are denoted as

$$\begin{aligned} A(\mathbf{x}, t) &= A_D(\mathbf{s}, t) \quad \text{on } \Gamma_D^A \times (T_0, T], \quad \text{and} \\ \frac{\eta}{\mu_0} \nabla A(\mathbf{s}, t) \cdot \mathbf{n}(\mathbf{s}) &= h(\mathbf{s}, t) \quad \text{on } \Gamma_N^A \times (T_0, T], \end{aligned} \quad (4.21)$$

with $A_D(\mathbf{s}, t)$, $h(\mathbf{s}, t)$ and $\bar{A}^0(\mathbf{x})$ being some known functions, while Γ_D^A and Γ_N^A identify the Dirichlet and Neumann parts of the boundary for $A(\mathbf{x}, t)$, respectively, (which might in general differ from Γ_D^u and Γ_N^u).

In the same fashion, we inherit the notation in (3.6) and (3.7) to describe the functional spaces where the velocity and pressure variables live in, and we introduce two more for the vector potential and current intensity variables, namely:

$$A(\cdot, t) \in C(\Omega) \equiv H_{\Gamma_D^A}^1(\Omega), \quad \text{and} \quad j(\cdot, t) \in D(\Omega) \equiv L^2(\Omega). \quad (4.22)$$

With this, we can provide the weak formulation for (4.19): we seek $[\mathbf{u}, p, j, A]^T \in Y(\Omega) = V(\Omega) \times Q(\Omega) \times D(\Omega) \times C(\Omega)$ such that

$$\left\{ \begin{array}{l} \frac{\partial}{\partial t} \langle \mathbf{u}, \mathbf{v} \rangle + \langle (\mathbf{u} \cdot \nabla) \mathbf{u}, \mathbf{v} \rangle + \mu \langle \nabla \mathbf{u}, \nabla \mathbf{v} \rangle \\ \quad - \langle p, \nabla \cdot \mathbf{v} \rangle + \langle j \nabla A, \mathbf{v} \rangle = \langle \mathbf{f}, \mathbf{v} \rangle + \int_{\Gamma_N^u} \mathbf{g} \cdot \mathbf{v} \\ \quad - \langle q, \nabla \cdot \mathbf{u} \rangle = 0 \\ \quad \langle j, l \rangle + \langle \nabla A, \nabla l \rangle - \int_{\Gamma_D^A} \nabla A \cdot \mathbf{n} l = \int_{\Gamma_N^A} h l \\ \frac{\partial}{\partial t} \langle A, c \rangle + \langle \mathbf{u} \cdot \nabla A, c \rangle + \frac{\eta}{\mu_0} \langle \nabla A, \nabla c \rangle = - \langle E, c \rangle + \int_{\Gamma_N^A} h c \end{array} \right. , \quad (4.23)$$

for each $[\mathbf{v}(\mathbf{x}), q(\mathbf{x}), l(\mathbf{x}), c(\mathbf{x})]^T \in Y(\Omega)$, and for each instant $t \in (T_0, T]$. (In order to reduce notation, in (4.23) we don't report explicitly the dependence of the variables on \mathbf{x} and t .) We express the weak formulation more compactly by introducing the nonlinear operator $\mathcal{N}(\mathbf{z})$, which maps $\mathbf{z} \in Y(\Omega)$ to the residual of (4.19), so that (4.23) can be rewritten as: find $\mathbf{z} \in Y(\Omega)$ such that

$$\langle \mathcal{N}(\mathbf{z}), \mathbf{y} \rangle = 0 \quad \forall \mathbf{y} \in Y(\Omega). \quad (4.24)$$

To resolve the nonlinearity in (4.23), we use Newton's method: that is, we iteratively solve the linearised variational problem

$$\begin{aligned} \langle \mathcal{J}_{\mathcal{N}}(\mathbf{z}^k) \delta \mathbf{z}^k, \mathbf{y} \rangle &= - \langle \mathcal{N}(\mathbf{z}^k), \mathbf{y} \rangle, \quad \forall \mathbf{y} \in Y(\Omega) \\ \mathbf{z}^{k+1} &= \mathbf{z}^k + \delta \mathbf{z}^k \end{aligned} \quad (4.25)$$

for the perturbation $\delta \mathbf{z}^k$, at each iteration k . Here, $\mathcal{J}_{\mathcal{N}}(\mathbf{z})$ denotes the Jacobian of the nonlinear form \mathcal{N} , evaluated at \mathbf{z} . This is a linear tensor operator in the form

$$\mathcal{J}_{\mathcal{N}} \left(\begin{bmatrix} \mathbf{u} \\ p \\ j \\ A \end{bmatrix} \right) := \left[\begin{array}{c|c|c|c} \frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla + \nabla \mathbf{u} - \mu \nabla^2 & \nabla & \nabla A & j \nabla \\ \hline & -\nabla \cdot & & \\ \hline & & \mathcal{I} & -\frac{1}{\mu_0} \nabla^2 \\ \hline & \nabla A \cdot & & \frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla - \frac{\eta}{\mu_0} \nabla^2 \end{array} \right], \quad (4.26)$$

whose discretisation is described in the following section.

We point out that resorting to Newton's method as a nonlinear solver is by no means necessary, and that our choice mostly stems from the desire of testing the efficacy of our space-time block preconditioners with solvers alternative to Picard (already used in Sec. 3.4.3), as well as the interest in following the setup presented in [25]. Although not investigated in our work, an appealing alternative in this sense could be provided by the *minimal decoupling Picard* iterations outlined in [88]: by dropping the $\nabla A \cdot$ term in the bottom-left block of (4.26), this scheme renders the target operator block upper-triangular, thus making it more straightforward to invert, and in particular simplifying most of the discussion in Sec. 4.2.

4.1.1 Space-time discretisation

For discretising the system (4.19), we follow the same approach used for (3.1): temporal derivatives are approximated using implicit Euler on a uniform grid, while for the spatial discretisation we use a Finite Element approach. In addition to the discrete velocity and pressure function spaces introduced in (3.10), we also consider

$$C_h(\Omega) := \text{span}\{\zeta_0(\mathbf{x}), \dots, \zeta_{N_A-1}(\mathbf{x})\}, \quad \text{and} \quad (4.27a)$$

$$D_h(\Omega) := \text{span}\{\xi_0(\mathbf{x}), \dots, \xi_{N_j-1}(\mathbf{x})\}, \quad (4.27b)$$

approximations of $C(\Omega)$ and $D(\Omega)$, with N_A and N_j degrees of freedom, respectively, so that functions in (4.27) can be represented by vectors $\mathbf{A} \in \mathbb{R}^{N_A}$ and $\mathbf{j} \in \mathbb{R}^{N_j}$ according to

$$A(\mathbf{x}) = \sum_{i=0}^{N_A-1} [\mathbf{A}]_i \zeta_i(\mathbf{x}) \quad \text{and} \quad j(\mathbf{x}) = \sum_{i=0}^{N_j-1} [\mathbf{j}]_i \xi_i(\mathbf{x}). \quad (4.28)$$

With these, we are ready to describe the relevant Galerkin matrices composing the discretised version of the linearised system (4.25). We begin with the definition of the spatial operators of interest, proceeding then to opportunistically combine them in order to create a discretisation of the space-time operators appearing in (4.26).

Starting from the momentum equation, we have that the linearisation of the advection term $(\mathbf{u} \cdot \nabla)\mathbf{u}$ produces the operator

$$\left[\tilde{\mathcal{W}}_{\mathbf{u}}(\mathbf{u}(\mathbf{x}))\right]_{m,n=0}^{N_{\mathbf{u}}-1} := \int_{\Omega} \boldsymbol{\phi}_m(\mathbf{x}) \cdot [(\mathbf{u}(\mathbf{x}) \cdot \nabla) \boldsymbol{\phi}_n(\mathbf{x}) + (\boldsymbol{\phi}_n(\mathbf{x}) \cdot \nabla) \mathbf{u}(\mathbf{x})] d\mathbf{x}; \quad (4.29)$$

while linearising the Lorentz force term with respect to the vector potential and the current gives instead

$$[\mathcal{Z}_j(A(\mathbf{x}))]_{m,n=0}^{N_{\mathbf{u}}-1, N_j-1} := \int_{\Omega} \nabla A(\mathbf{x}) \cdot \boldsymbol{\phi}_m(\mathbf{x}) \xi_n(\mathbf{x}) d\mathbf{x}, \quad \text{and} \quad (4.30a)$$

$$[\mathcal{Z}_A(j(\mathbf{x}))]_{m,n=0}^{N_{\mathbf{u}}-1, N_A-1} := \int_{\Omega} j(\mathbf{x}) \boldsymbol{\phi}_m(\mathbf{x}) \cdot \nabla \zeta_n(\mathbf{x}) d\mathbf{x}. \quad (4.30b)$$

For the definition of $j(\mathbf{x}, t)$, the relevant operators appearing are the current mass matrix,

$$[\mathcal{M}_j]_{m,n=0}^{N_j-1} := \int_{\Omega} \xi_m(\mathbf{x}) \xi_n(\mathbf{x}) d\mathbf{x}, \quad (4.31)$$

and the mixed current-vector potential stiffness matrix (scaled by μ_0),

$$[\mathcal{K}_{jA}]_{m,n=0}^{N_j-1, N_A-1} := \frac{1}{\mu_0} \int_{\Omega} \nabla \xi_m(\mathbf{x}) \cdot \nabla \zeta_n(\mathbf{x}) d\mathbf{x} - \frac{1}{\mu_0} \int_{\Gamma_D^A} \xi_m(\mathbf{s}) \nabla \zeta_n(\mathbf{s}) \cdot \mathbf{n}(\mathbf{s}) d\mathbf{s}. \quad (4.32)$$

Lastly, for the vector potential equation, we introduce its mass, stiffness, and advection matrices

$$[\mathcal{M}_A]_{m,n=0}^{N_A-1} := \int_{\Omega} \zeta_m(\mathbf{x}) \cdot \zeta_n(\mathbf{x}) d\mathbf{x}, \quad (4.33a)$$

$$[\mathcal{K}_A]_{m,n=0}^{N_A-1} := \int_{\Omega} \nabla \zeta_m(\mathbf{x}) \cdot \nabla \zeta_n(\mathbf{x}) d\mathbf{x}, \quad \text{and} \quad (4.33b)$$

$$[\mathcal{W}_A(\mathbf{u}(\mathbf{x}))]_{m,n=0}^{N_A-1} := \int_{\Omega} \zeta_m(\mathbf{x}) \mathbf{u}(\mathbf{x}) \cdot \nabla \zeta_n(\mathbf{x}) d\mathbf{x}. \quad (4.33c)$$

The latter is one of the two terms stemming from the linearisation of $\mathbf{u} \cdot \nabla A$; the other is given by

$$[\mathcal{Y}(A(\mathbf{x}))]_{m,n=0}^{N_A-1, N_{\mathbf{u}}-1} := \int_{\Omega} \zeta_m(\mathbf{x}) \nabla A(\mathbf{x}) \cdot \boldsymbol{\phi}_n(\mathbf{x}) d\mathbf{x}. \quad (4.34)$$

The relevant discrete space-time operators are then assembled using the matrices defined above as building blocks. In particular, combining the linearised advection operator (4.29) with the velocity mass and stiffness matrices (3.12), we produce

$$\mathcal{F}_{\mathbf{u}}(\mathbf{u}(\mathbf{x})) := \frac{\mathcal{M}_{\mathbf{u}}}{\Delta t} + \tilde{\mathcal{W}}_{\mathbf{u}}(\mathbf{u}(\mathbf{x})) + \mu \mathcal{K}_{\mathbf{u}}. \quad (4.35)$$

This appears in the discrete space-time velocity matrix

$$F_{\mathbf{u}}(\mathbf{u}(\mathbf{x}, t)) := \begin{bmatrix} \mathcal{F}_{\mathbf{u}}(\mathbf{u}(\mathbf{x}, t_1)) & & \\ -\frac{\mathcal{M}_{\mathbf{u}}}{\Delta t} & \ddots & \\ & \ddots & \mathcal{F}_{\mathbf{u}}(\mathbf{u}(\mathbf{x}, t_{N_t})) \end{bmatrix}, \quad (4.36)$$

which shares the same block bi-diagonal structure as (3.22). Similarly, we combine (4.33) to create

$$\mathcal{F}_A(\mathbf{u}(\mathbf{x})) := \frac{\mathcal{M}_A}{\Delta t} + \mathcal{W}_A(\mathbf{u}(\mathbf{x})) + \frac{\eta}{\mu_0} \mathcal{K}_A, \quad (4.37)$$

which we use in the definition of the discrete space-time vector potential operator

$$F_A(\mathbf{u}(\mathbf{x}, t)) := \begin{bmatrix} \mathcal{F}_A(\mathbf{u}(\mathbf{x}, t_1)) & & \\ -\frac{\mathcal{M}_A}{\Delta t} & \ddots & \\ & \ddots & \mathcal{F}_A(\mathbf{u}(\mathbf{x}, t_{N_t})) \end{bmatrix}. \quad (4.38)$$

The remaining operators do not involve any temporal derivative, and as a consequence their space-time counterparts preserve a block diagonal structure, much like operator B in (3.22). These are given by:

$$\begin{aligned} Z_j(A(\mathbf{x}, t)) &:= \text{diag}_{k=1, \dots, N_t} (Z_A(A(\mathbf{x}, t_k))), \\ Z_A(j(\mathbf{x}, t)) &:= \text{diag}_{k=1, \dots, N_t} (Z_j(j(\mathbf{x}, t_k))), \\ M_j &:= \text{diag}_{N_t}(\mathcal{M}_j), \\ K_{jA} &:= \text{diag}_{N_t}(\mathcal{K}_{jA}), \quad \text{and} \\ Y(A(\mathbf{x}, t)) &:= \text{diag}_{k=1, \dots, N_t} (\mathcal{Y}(A(\mathbf{x}, t_k))), \end{aligned} \quad (4.39)$$

where $\text{diag}_{k=1, \dots, N}((*)_k)$ denotes a $N \times N$ block diagonal matrix, whose diagonal blocks are taken in order from the indexed operators $\{(*)_k\}_{k=1}^N$.

In turn, the space-time operators eqs. (4.36), (4.38) and (4.39) compose the discrete version of the Jacobian (4.26),

$$J_{\mathcal{N}} \left(\begin{bmatrix} \mathbf{u} \\ \mathbf{p} \\ \mathbf{j} \\ \mathbf{A} \end{bmatrix} \right) := \begin{bmatrix} F_{\mathbf{u}}(\mathbf{u}) & B^T & Z_j(\mathbf{A}) & Z_A(\mathbf{j}) \\ B & & & \\ & & M_j & K_{jA} \\ Y(\mathbf{A}) & & & F_A(\mathbf{u}) \end{bmatrix}, \quad (4.40)$$

where \mathbf{u} , \mathbf{p} , \mathbf{j} , and \mathbf{A} represent the whole space-time approximations to the unknowns of (4.19), so that, for example, $\mathbf{A} = [(\mathbf{A}^1)^T, (\mathbf{A}^2)^T, \dots, (\mathbf{A}^{N_t})^T]^T$, with $\mathbf{A}^i \approx A(\mathbf{x}, t_i)$ according to (4.28).

To complete the description of the discretisation of the linearised system (4.25), we also need to define the discrete counterpart of the residual $\mathcal{N}(\mathbf{z}(\mathbf{x}, t))$, which we denote as $\mathbf{N}(\mathbf{z})$, with $\mathbf{z} = [\mathbf{u}^T, \mathbf{p}^T, \mathbf{j}^T, \mathbf{A}^T]^T$. This is also expressed in a space-time block fashion:

$$\mathbf{N}(\mathbf{z}) = [\mathbf{N}_u(\mathbf{z})^T, \mathbf{N}_p(\mathbf{z})^T, \mathbf{N}_j(\mathbf{z})^T, \mathbf{N}_A(\mathbf{z})^T]^T, \quad (4.41)$$

where for each variable $(*)$ in the system, $\mathbf{N}_{(*)}(\mathbf{z})$ is in itself a space-time vector, created by collating the spatial vectors $\mathbf{N}_{(*)}^k(\mathbf{z})$ for each instant $k = 1, \dots, N_t$. These are given by

$$\begin{aligned} \mathbf{N}_u^k(\mathbf{z}) := & \frac{\mathcal{M}_u}{\Delta t} (\mathbf{u}^k - \mathbf{u}^{k-1}) + \mu \mathcal{K}_u \mathbf{u}^k + \mathcal{B} \mathbf{p}^k - \mathbf{f}^k \\ & + \left[\int_{\Omega} ((\mathbf{u}^k \cdot \nabla) \mathbf{u}^k + \mathbf{j}^k \nabla \mathbf{A}^k) \cdot \boldsymbol{\phi}_i \right]_{i=0}^{N_u-1}, \end{aligned} \quad (4.42a)$$

$$\mathbf{N}_p^k(\mathbf{z}) := \mathcal{B} \mathbf{u}^k, \quad (4.42b)$$

$$\mathbf{N}_j^k(\mathbf{z}) := \mathcal{M}_j \mathbf{j}^k + \mathcal{K}_{jA} \mathbf{A}^k - \mathbf{h}^k, \quad \text{and} \quad (4.42c)$$

$$\mathbf{N}_A^k(\mathbf{z}) := \frac{\mathcal{M}_A}{\Delta t} (\mathbf{A}^k - \mathbf{A}^{k-1}) + \frac{\eta}{\mu_0} \mathcal{K}_A \mathbf{A}^k + \mathbf{E}^k + \left[\int_{\Omega} \mathbf{u}^k \cdot \nabla \mathbf{A}^k \zeta_i \right]_{i=0}^{N_A-1}, \quad (4.42d)$$

with \mathbf{f}^k in (4.42a) defined as in (3.16), while for (4.42c) and (4.42d) we take

$$\begin{aligned} [\mathbf{h}^k]_{i=0}^{N_j-1} &:= \int_{\Gamma_N^A} h(\mathbf{s}, t_k) \xi_i(\mathbf{s}) d\mathbf{s}, \quad \text{and} \\ [\mathbf{E}^k]_{i=0}^{N_A-1} &:= \int_{\Omega} E(\mathbf{x}, t_k) \zeta_i(\mathbf{x}) d\mathbf{x} - \int_{\Gamma_N^A} h(\mathbf{s}, t_k) \zeta_i(\mathbf{s}) d\mathbf{s}. \end{aligned} \quad (4.43)$$

Notice that, similarly to (3.19), for $k = 1$ we modify $\mathbf{N}_u^1(\mathbf{z})$ and $\mathbf{N}_A^1(\mathbf{z})$ by considering the discretised initial conditions $\bar{\mathbf{u}}^0$ and $\bar{\mathbf{A}}^0$ instead of \mathbf{u}^0 and \mathbf{A}^0 , respectively. With the relevant quantities defined, we can finally write the target discrete system for the k -th Newton iteration:

$$J_{\mathcal{N}}(\mathbf{z}^k) \delta \mathbf{z}^k = -\mathbf{N}(\mathbf{z}^k). \quad (4.44)$$

Solving systems in the form of (4.44) is the main target of the study in this chapter. To do so effectively, and in a time-parallel fashion, we follow a strategy similar to the one pursued in Chap. 3. That is, we identify a preconditioner which has shown its efficacy when used within a time-stepping procedure, and extend it to cover the whole space-time case. The preconditioner of choice was initially developed by Cyr *et al.* in [25] for the single time-step system arising from the discretisation of (4.17). The main motivation behind this particular choice stems from the specific design of this

an *approximate* LDU factorisation of (4.40): adapted to our formulation (4.19), and extended to the whole space-time case, this reads

$$\begin{aligned}
J_N &\approx \begin{bmatrix} F_{\mathbf{u}} & B^T & Z_j & Z_A \\ B & & & \\ & & M_j & K_{jA} \\ Y & \boxed{YF_{\mathbf{u}}^{-1}B^T} & & F_A \end{bmatrix} \\
&= \underbrace{\begin{bmatrix} F_{\mathbf{u}} & & Z_j & Z_A \\ & I & & \\ & & M_j & K_{jA} \\ Y & & & F_A \end{bmatrix}}_{=:P_{\mathbf{u}jA}} \underbrace{\begin{bmatrix} F_{\mathbf{u}}^{-1} & & & \\ & I & & \\ & & I & \\ & & & I \end{bmatrix}}_{=:F_{\mathbf{u}I}^{-1}} \underbrace{\begin{bmatrix} F_{\mathbf{u}} & B^T & & \\ B & & & \\ & & I & \\ & & & I \end{bmatrix}}_{=:P_{\mathbf{u}p}}. \tag{4.48}
\end{aligned}$$

Such approximation consists in including an additional term in the vector potential equation, identified by the highlighted block in (4.48). Although a valid physical intuition about the effect of this perturbation is somewhat hard to grasp, the analyses conducted in [25] and Sec. 4.3 hint at the fact that the perturbed matrix has a spectrum similar to that of the original Jacobian, making this a valid approximation for our purposes. Moreover, the resulting factorisation has the remarkable advantage of neatly splitting the original fully-coupled system into two simpler subsystems, denoted as $P_{\mathbf{u}jA}$ and $P_{\mathbf{u}p}$, thus removing the complexity of dealing with nested Schur complements. We can convince ourselves of this by noticing how the factors in (4.48) present some block rows/columns which contain only an identity on the main diagonal: this effectively isolates the influence of some variables in the system. In the following sections we analyse more in detail the single factors composing this factorisation, providing approximations to their respective Schur complements, with the goal of simplifying the application of their inverse. It is these simplifications that our space-time preconditioner ultimately depends upon.

4.2.1 Velocity-pressure coupling

The rightmost matrix $P_{\mathbf{u}p}$ in (4.48) considers only the coupling between the velocity and pressure variables: applying the inverse of this factor is equivalent to solving an incompressible flow problem, without considering the effect of the magnetic field. This system corresponds to a similar problem to the one we tackled in Chap. 3, the sole difference lying in the velocity advection operator $\mathcal{W}_{\mathbf{u}}$. In fact, comparing the definition of (4.29) with that of (3.14), we notice the former presents an extra term,

$$[\Delta\mathcal{W}_{\mathbf{u}}(\mathbf{u}(\mathbf{x}))]_{m,n=0}^{N_{\mathbf{u}}-1} := [(\tilde{\mathcal{W}}_{\mathbf{u}} - \mathcal{W}_{\mathbf{u}})(\mathbf{u}(\mathbf{x}))]_{m,n=0}^{N_{\mathbf{u}}-1} = \int_{\Omega} ((\phi_n(\mathbf{x}) \cdot \nabla) \mathbf{u}(\mathbf{x})) \cdot \phi_m(\mathbf{x}) d\mathbf{x}, \tag{4.49}$$

due to the Newton linearisation. Since in Sec. 3.2.1 we derived our space-time PCD preconditioner from a commutation argument between \mathcal{F}_u (which contains \mathcal{W}_u) and an equivalent operator \mathcal{F}_p acting on the pressure field, one would expect that modifications to \mathcal{F}_u should push us to change \mathcal{F}_p as well, thus modifying the final form of the preconditioner. In practice, the extra term in (4.49) includes some interplay between different components of the velocity variable, and it is not straightforward to construct its equivalent on the pressure variable, as the latter is only a scalar field. For this reason, to our knowledge the common approach in the single time-step case [24, 39] is to just ignore this extra contribution on the definition of \mathcal{F}_p .

We pursue the same approach here, limiting ourselves to reusing the same preconditioner introduced in Chap. 3. Specifically, to simplify the solution of systems involving P_{up} , we consider its block LU factorisation,

$$P_{up} = \underbrace{\begin{bmatrix} I & & & \\ BF_u^{-1} & I & & \\ & & I & \\ & & & I \end{bmatrix}}_{=:L_{up}} \underbrace{\begin{bmatrix} F_u & B^T & & \\ & S_p & & \\ & & I & \\ & & & I \end{bmatrix}}_{=:U_{up}}, \quad (4.50)$$

and substitute its block upper-triangular term U_{up} with the approximation

$$\tilde{U}_{up} := \begin{bmatrix} F_u & B^T & & \\ & \tilde{S}_p & & \\ & & I & \\ & & & I \end{bmatrix}, \quad (4.51)$$

where the space-time pressure Schur complement S_p is swapped for the operator $\tilde{S}_p := -K_p F_p^{-1} M_p$ introduced in (3.37).

4.2.2 Velocity-magnetic field coupling

On the other hand, the leftmost matrix P_{ujA} in (4.48) considers only the coupling between velocity, current, and vector potential, ignoring the incompressibility constraint associated with the pressure variable. To find an effective way to approximately solve systems involving P_{ujA} , we once again consider its block LU decomposition:

$$P_{ujA} = \underbrace{\begin{bmatrix} I & & & & \\ & I & & & \\ & & I & & \\ YF_u^{-1} & & -YF_u^{-1}Z_jM_j^{-1} & & I \end{bmatrix}}_{=:L_{ujA}} \underbrace{\begin{bmatrix} F_u & Z_j & Z_A & & \\ & I & & & \\ & & M_j & K_{jA} & \\ & & & S_A & \end{bmatrix}}_{=:U_{ujA}}. \quad (4.52)$$

Here, S_A represents another space-time magnetic field Schur complement³, different from (4.47), and specific to this sub-problem. This is defined as:

$$S_A := F_A - Y F_{\mathbf{u}}^{-1} \underbrace{(Z_A - Z_j M_j^{-1} K_{jA})}_{=: Z}, \quad (4.53)$$

where we compact the whole linearisation of the Lorentz force term⁴ into

$$\begin{aligned} Z(j(\mathbf{x}, t), A(\mathbf{x}, t)) &:= Z_A(j(\mathbf{x}, t)) - Z_j(A(\mathbf{x}, t)) M_j^{-1} K_{jA} \\ &= \text{diag}_{k=1, \dots, N_t} (Z(j(\mathbf{x}, t_k), A(\mathbf{x}, t_k))), \end{aligned} \quad (4.55)$$

whose individual blocks are defined as

$$\mathcal{Z}(j(\mathbf{x}), A(\mathbf{x})) := \mathcal{Z}_A(j(\mathbf{x})) - \mathcal{Z}_j(A(\mathbf{x})) \mathcal{M}_j^{-1} \mathcal{K}_{jA}. \quad (4.56)$$

The main goal of this section is to provide an adequate approximation to the space-time operator (4.53), and in particular its term $Y F_{\mathbf{u}}^{-1} Z$. Analogously to (3.27), this term is composed of blocks in the form

$$\mathcal{Y}_k \mathcal{F}_{\mathbf{u}, k}^{-1} \left(\prod_{l=i}^{k-1} \left(\frac{\mathcal{M}_{\mathbf{u}}}{\Delta t} \mathcal{F}_{\mathbf{u}, l}^{-1} \right) \right) \mathcal{Z}_i^T, \quad \text{with} \quad \begin{cases} k = 1, \dots, N_t \\ i = 1, \dots, k \end{cases}, \quad (4.57)$$

where \mathcal{Z}_i , \mathcal{Y}_i and $\mathcal{F}_{\mathbf{u}, i}$ is shorthand notation for $\mathcal{Z}(j(\mathbf{x}, t_i), A(\mathbf{x}, t_i))$, $\mathcal{Y}(A(\mathbf{x}, t_i))$ and $\mathcal{F}(\mathbf{u}(\mathbf{x}, t_i))$, for some fixed fields $\mathbf{u}(\mathbf{x}, t)$, $j(\mathbf{x}, t)$, and $A(\mathbf{x}, t)$. In [25] the authors make use of a small-commutator argument akin to the one exploited in Sec. 3.2.1, and assume that

$$\mathcal{M}_A^{-1} \mathcal{F}_{A, i} \mathcal{M}_A^{-1} \mathcal{Y}_k \approx \mathcal{M}_A^{-1} \mathcal{Y}_k \mathcal{M}_{\mathbf{u}}^{-1} \mathcal{F}_{\mathbf{u}, i}, \quad \forall i, k = 1, \dots, N_t. \quad (4.58)$$

Notice that the operator $\mathcal{F}_{\mathbf{u}, i}$ contains the extra term (4.49) which does not find an equivalent in $\mathcal{F}_{A, i}$, but we accept this discrepancy under similar considerations as for

³More specifically, S_A simplifies \bar{S}_A , by dropping the contribution from the whole term $B F_{\mathbf{u}}^{-1} Z$. Considering the role covered by the single operators appearing in that term, one can interpret this simplification as assuming that the divergence (B) of the velocity field resulting from ($F_{\mathbf{u}}^{-1}$) the Lorentz force (Z) is null.

⁴Notice that in [25, (3.7)], the single time-step magnetic Schur complement assumes the form

$$S_A = \mathcal{F}_A - \mathcal{Y} \mathcal{F}_{\mathbf{u}}^{-1} \mathcal{Z}, \quad (4.54)$$

where \mathcal{Z} comes from the linearisation of (4.16). In virtue of our formulation (4.19), which explicitly considers the current as an additional variable, this term appears in a different form in (4.53), *i.e.*, it is further split into its two components \mathcal{Z}_A and $-\mathcal{Z}_j \mathcal{M}_j^{-1} \mathcal{K}_{jA}$. Nonetheless, as it represents the same operator, we can still apply also to our context similar considerations used in [25] for approximating (4.54).

Sec. 4.2.1. Even ignoring such a discrepancy, the formula would not hold in general; nonetheless, similar assumptions are not uncommon when applying this argument. Following a procedure similar to the one described in Sec. 3.2.1, we recursively use assumption (4.58) to recover reasonable approximations to the blocks (4.57), and thus to the whole space-time operator (4.53). Left-multiplying (4.58) by $\mathcal{M}_A \mathcal{F}_{A,i}^{-1} \mathcal{M}_A$, and right-multiplying it by $\mathcal{F}_{u,i}^{-1}$, we recover

$$\mathcal{Y}_k \mathcal{F}_{u,i}^{-1} \approx \mathcal{M}_A \mathcal{F}_{A,i}^{-1} \mathcal{Y}_k \mathcal{M}_u^{-1}, \quad \forall i, k = 1, \dots, N_t. \quad (4.59)$$

Starting from $i = k$, continuing to right-multiply by $\frac{\mathcal{M}_u}{\Delta t} \mathcal{F}_{u,k-j}^{-1}$, increasing the index j from 0 until we reach $k - j = l$, repeatedly using (4.59) for each $i = k - j$, and finally right-multiplying by \mathcal{Z}_l , gives us the following approximation

$$\begin{aligned} \mathcal{Y}_k \mathcal{F}_{u,k}^{-1} \frac{\mathcal{M}_u}{\Delta t} \mathcal{F}_{u,k-1}^{-1} &\approx \mathcal{M}_A \mathcal{F}_{A,k}^{-1} \underbrace{\frac{\mathcal{Y}_k \mathcal{F}_{u,k-1}^{-1}}{\Delta t}}_{\approx \frac{\mathcal{M}_A}{\Delta t} \mathcal{F}_{A,k-1}^{-1} \mathcal{Y}_k \mathcal{M}_u^{-1}} \\ \implies \mathcal{Y}_k \mathcal{F}_{u,k}^{-1} \left(\prod_{j=1}^{k-l} \left(\frac{\mathcal{M}_u}{\Delta t} \mathcal{F}_{u,k-j}^{-1} \right) \right) \mathcal{Z}_l &\approx \mathcal{M}_A \mathcal{F}_{A,k}^{-1} \left(\prod_{j=1}^{k-l} \left(\frac{\mathcal{M}_A}{\Delta t} \mathcal{F}_{A,k-j}^{-1} \right) \right) \mathcal{Y}_k \mathcal{M}_u^{-1} \mathcal{Z}_l, \end{aligned} \quad (4.60)$$

which we can apply to each of the blocks in (4.57). In turn, this identifies a first approximation for the space-time magnetic Schur complement (4.53):

$$\begin{aligned} Y F_u^{-1} Z &\approx M_A F_A^{-1} Y M_u^{-1} Z \\ \implies S_A &\approx F_A - M_A F_A^{-1} Y M_u^{-1} Z = M_A F_A^{-1} (F_A M_A^{-1} F_A - Y M_u^{-1} Z), \end{aligned} \quad (4.61)$$

which is a direct space-time equivalent of the one proposed in [25, (3.24)] for the single time-step magnetic Schur complement at instant k :

$$\mathcal{S}_{A,k} \approx \mathcal{M}_A \mathcal{F}_{A,k}^{-1} (\mathcal{F}_{A,k} \mathcal{M}_A^{-1} \mathcal{F}_{A,k} - \mathcal{Y}_k \mathcal{M}_u^{-1} \mathcal{Z}_k), \quad \forall k = 1, \dots, N_t. \quad (4.62)$$

An additional approximation is performed in [25] in order to further simplify the term $\mathcal{Y}_k \mathcal{M}_u^{-1} \mathcal{Z}_k$ in (4.62): this term has proven difficult to treat there, as it heavily disrupts the sparsity pattern of the first term $\mathcal{F}_{A,k} \mathcal{M}_A^{-1} \mathcal{F}_{A,k}$, rendering its solution via multigrid more challenging⁵. The justification behind the simplification proposed in [25] comes from analysing the continuous problem associated with inverting $P_{u_j A}$, and can be seamlessly extended to the whole space-time framework. In particular,

⁵In our formulation, an additional complication stems from the presence in \mathcal{Z}_k of the inverse of the mass matrix \mathcal{M}_j , which would also need to be properly approximated (together with \mathcal{M}_u).

solving for the velocity-magnetic field coupling (4.52) corresponds to finding a solution to the PDE

$$\begin{cases} \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} - \mu \nabla^2 \mathbf{u} + \frac{1}{\mu_0} \mathbf{B} \times (\nabla \times \mathbf{B}) = \mathbf{0} \\ \frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{u} \times \mathbf{B}) + \frac{\eta}{\mu_0} \nabla \times (\nabla \times \mathbf{B}) = \mathbf{0} \end{cases} \quad \text{in } \Omega \times (T_0, T], \quad (4.63)$$

obtained by removing the pressure variable from (4.1). Solving for the magnetic Schur complement is the discrete equivalent of using the first equation to express \mathbf{u} in terms of \mathbf{B} , substituting it in the second equation, and recovering its solution \mathbf{B} . To find an approximation to (4.62), the authors in [25] perform a similar procedure, but at the continuous level, and consider a linearised version of (4.63). Using perturbation theory, they expand the solution

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{u}_0(\mathbf{x}, t) + \mathbf{u}_1(\mathbf{x}, t) \quad \text{and} \quad \mathbf{B}(\mathbf{x}, t) = \mathbf{B}_0(\mathbf{x}, t) + \mathbf{B}_1(\mathbf{x}, t) \quad (4.64)$$

around a quiet state $\mathbf{u}_0(\mathbf{x}, t) \equiv \mathbf{0}$ with a constant background magnetic field $\mathbf{B}_0(\mathbf{x}, t) \equiv \mathbf{B}_0$. Positioning themselves in the hyperbolic limit, for which the second-order dissipative terms in (4.63) are neglected, and dropping high-order perturbations, they recover

$$\begin{cases} \frac{\partial \mathbf{u}_1}{\partial t} + \frac{1}{\mu_0} \mathbf{B}_0 \times (\nabla \times \mathbf{B}_1) = \mathbf{0} \\ \frac{\partial \mathbf{B}_1}{\partial t} - \nabla \times (\mathbf{u}_1 \times \mathbf{B}_0) = \mathbf{0} \end{cases} \quad \text{in } \Omega \times (T_0, T]. \quad (4.65)$$

To find the solution \mathbf{B}_1 , they proceed to take an extra time derivative on the second equation,

$$\frac{\partial^2 \mathbf{B}_1}{\partial t^2} - \nabla \times \left(\frac{\partial \mathbf{u}_1}{\partial t} \times \mathbf{B}_0 \right) - \nabla \times \left(\mathbf{u}_1 \times \frac{\partial \mathbf{B}_0}{\partial t} \right) = 0, \quad (4.66)$$

and substitute back in the first equation, to recover

$$\begin{aligned} & \frac{\partial^2 \mathbf{B}_1}{\partial t^2} - \nabla \times \left(\frac{1}{\mu_0} (\nabla \times \mathbf{B}_1) \times \mathbf{B}_0 \times \mathbf{B}_0 \right) = 0 \\ \implies & \frac{\partial^2 \mathbf{B}_1}{\partial t^2} - \frac{\|\mathbf{B}_0\|^2}{\mu_0} \nabla^2 \mathbf{B}_1 = 0. \end{aligned} \quad (4.67)$$

This equation can be analogously expressed in terms of the vector potential,

$$\frac{\partial^2 A}{\partial t^2} - \frac{\|\mathbf{B}_0\|^2}{\mu_0} \nabla^2 A = 0, \quad (4.68)$$

since the discussion in Sec. 4.1 ensures that there is unique correspondence between A and its curl. The magnetic Schur complement (4.62) should then approximate at least to some extent the action of the operator in (4.68), that is, it should represent a

wave equation with propagation speed⁶ $\sqrt{\|\mathbf{B}_0\|^2/\mu_0}$. In [25], this fact is exploited to justify substituting the term $\mathcal{Y}\mathcal{M}_u^{-1}\mathcal{Z}$ in (4.62) with a discrete Laplacian operator, opportunely scaled by the wave speed. In the space-time case, this translates to the following approximation of the magnetic Schur complement:

$$S_A \approx M_A F_A^{-1} \underbrace{(F_A M_A^{-1} F_A + K_A)}_{=: C_A}, \quad (4.69)$$

where we define

$$K_A := \text{diag}_{k=1, \dots, N_t} \left(\frac{\|\bar{\mathbf{B}}_0^k\|^2}{\mu_0} \mathcal{K}_A \right), \quad \text{with} \quad \bar{\mathbf{B}}^k := \frac{1}{|\Omega|} \int_{\Omega} \nabla \times (A(\mathbf{x}, t_k) \hat{\mathbf{k}}) d\mathbf{x}, \quad (4.70)$$

that is, $\bar{\mathbf{B}}_0^k$ is taken as the space average of the magnetic field at instant t_k . The task of explicitly assembling C_A in (4.69) is complicated by the presence of \mathcal{M}_A^{-1} . As a last simplification, we consider instead its approximation \tilde{C}_A , where we pick only the diagonal of \mathcal{M}_A , denoted as $\mathcal{D}_{\mathcal{M}_A}$, as a surrogate for the whole mass matrix. Overall, the space-time operator \tilde{C}_A has the following block tri-diagonal structure:

$$\tilde{C}_A := \text{diag}_{k=1, \dots, N_t} \left(\tilde{\mathcal{C}}_{A,k} \right) + \frac{1}{\Delta t} \text{diag}_{k=1, \dots, N_t-1} \left(\tilde{\mathcal{C}}_{A,-1,k}, -1 \right) + \frac{1}{\Delta t^2} \text{diag}_{N_t-2} \left(\tilde{\mathcal{C}}_{A,-2}, -2 \right), \quad (4.71)$$

where the notation $\text{diag}_{k=1, \dots, N-i} ((*)_k, -i)$ indicates a $N \times N$ block diagonal matrix whose i -th block sub-diagonal is filled by taking in order the indexed operators $\{(*)_k\}_{k=1}^{N-i}$. For each block diagonal, these are given by

$$\tilde{\mathcal{C}}_{A,k} := \mathcal{F}_{A,k} \mathcal{D}_{\mathcal{M}_A}^{-1} \mathcal{F}_{A,k} + \frac{\|\bar{\mathbf{B}}_0^k\|^2}{\mu_0} \mathcal{K}_A \quad k = 1, \dots, N_t \quad (4.72a)$$

$$\tilde{\mathcal{C}}_{A,-1,k} := \mathcal{M}_A \mathcal{D}_{\mathcal{M}_A}^{-1} \mathcal{F}_{A,k} + \mathcal{F}_{A,k+1} \mathcal{D}_{\mathcal{M}_A}^{-1} \mathcal{M}_A \quad k = 1, \dots, N_t - 1 \quad (4.72b)$$

$$\tilde{\mathcal{C}}_{A,-2} := \mathcal{M}_A \mathcal{D}_{\mathcal{M}_A}^{-1} \mathcal{M}_A. \quad (4.72c)$$

This allows us to define the final form for our approximate space-time magnetic Schur complement:

$$S_A \approx \tilde{S}_A := M_A F_A^{-1} \tilde{C}_A \quad \iff \quad \tilde{S}_A^{-1} := \tilde{C}_A^{-1} F_A M_A^{-1}. \quad (4.73)$$

⁶This corresponds to the speed at which *Alfvén waves* travel: this is a type of magnetohydrodynamic wave generated in response to a perturbation to the magnetic field induced by the presence of currents in the plasma [2].

We thus proceed to simplify the solution of systems involving (4.52), in complete analogy with Sec. 4.2.1, and substitute its block upper-triangular factor $U_{\mathbf{u}jA}$ with its approximation

$$\tilde{U}_{\mathbf{u}jA} := \begin{bmatrix} F_{\mathbf{u}} & Z_j & Z_A \\ & I & \\ & & M_j & K_{jA} \\ & & & \tilde{S}_A \end{bmatrix}. \quad (4.74)$$

4.2.3 Definition of the space-time block preconditioner

Incorporating the approximate Schur complements introduced in Sec. 4.2.1 and 4.2.2 into the approximate factorisation (4.48), we are ready to define our space-time block preconditioner for (4.40). This appears in the form

$$P := L_{\mathbf{u}jA} \tilde{U}_{\mathbf{u}jA} F_{\mathbf{u}I}^{-1} L_{\mathbf{u}p} \tilde{U}_{\mathbf{u}p} \iff P^{-1} := \tilde{U}_{\mathbf{u}p}^{-1} L_{\mathbf{u}p}^{-1} F_{\mathbf{u}I} \tilde{U}_{\mathbf{u}jA}^{-1} L_{\mathbf{u}jA}^{-1}, \quad (4.75)$$

so that applying its inverse involves the following sequence of operations:

Step 1 *Invert $L_{\mathbf{u}jA}$.* From its definition in (4.52), after some modifications we can show that applying its inverse to a generic block vector $[\mathbf{x}_{\mathbf{u}}^T, \mathbf{x}_p^T, \mathbf{x}_j^T, \mathbf{x}_A^T]^T$ is equivalent to computing

$$L_{\mathbf{u}jA}^{-1} \begin{bmatrix} \mathbf{x}_{\mathbf{u}} \\ \mathbf{x}_p \\ \mathbf{x}_j \\ \mathbf{x}_A \end{bmatrix} = \begin{bmatrix} \mathbf{x}_{\mathbf{u}} \\ \mathbf{x}_p \\ \mathbf{x}_j \\ \mathbf{x}_A + Y F_{\mathbf{u}}^{-1} (Z_j M_j^{-1} \mathbf{x}_j - \mathbf{x}_{\mathbf{u}}) \end{bmatrix}. \quad (4.76)$$

This operation requires solving N_t independent systems involving the current intensity mass matrix \mathcal{M}_j , (which can be trivially parallelised over the various time-steps), and more importantly inverting the velocity space-time matrix $F_{\mathbf{u}}$ (4.36). We faced a similar⁷ problem in Chap. 3, so we refer to Sec. 3.4 for strategies on how to effectively tackle its parallel-in-time solution.

Step 2 *Apply $F_{\mathbf{u}I} \tilde{U}_{\mathbf{u}jA}^{-1}$.* The two operators can be combined, rendering the task in this step equivalent to inverting the system

$$\bar{U}_{\mathbf{u}jA} := \tilde{U}_{\mathbf{u}jA} F_{\mathbf{u}I}^{-1} = \begin{bmatrix} I & Z_j & Z_A \\ & I & \\ & & M_j & K_{jA} \\ & & & \tilde{S}_A \end{bmatrix}. \quad (4.77)$$

⁷The sole difference being the definition of the diagonal block $\mathcal{F}_{\mathbf{u}}$, which here contains the extra term (4.49), as mentioned in Sec. 4.2.1.

This step is possibly the most challenging in the application of the preconditioner. In fact, apart from an additional inversion of M_j , which can be dealt with in the same way as in step 1, this requires the inversion of the space-time operator \tilde{S}_A , and particularly its factor \tilde{C}_A . From its definition in (4.71), we see that this is a fundamentally different operator from $F_{\mathbf{u}}$: while the latter essentially represents the space-time discretisation of a standard parabolic PDE, the former rather corresponds to a space-time operator containing high-order mixed spatial and temporal derivatives. It is hence not clear what would be the best strategy to time-parallelise its solution. To keep the scope of this chapter contained, then, we limit ourselves to considering direct solvers (*i.e.*, time-stepping) for inverting (4.73), thus momentarily suspending considerations on the parallel performance of the preconditioner.

Step 3 *Invert $L_{\mathbf{u}p}$.* Once again, the main task of this step lies in the inversion of the space-time velocity matrix $F_{\mathbf{u}}$, and is similar in complexity to step 1.

Step 4 *Invert $\tilde{U}_{\mathbf{u}p}$.* Comparing the definition of (4.51) with that of (3.23), we can convince ourselves that this task is equivalent to applying the block upper-triangular space-time preconditioner P_T proposed in Sec. 3.2, which involves once again inverting the space-time velocity matrix $F_{\mathbf{u}}$, as well as the approximate space-time pressure Schur complement \tilde{S}_p (the latter task is trivially parallelisable, as discussed in detail in Sec. 3.2.1.1).

The heaviest cost in the application of the preconditioner is associated with the solution of the four space-time systems illustrated above: three involving $F_{\mathbf{u}}$, and one \tilde{C}_A . We reduce this number by dropping some of the block lower-triangular terms in (4.75): this has the result of rendering the preconditioner closer to a block upper-triangular form. In particular, by dropping its $L_{\mathbf{u}jA}$ factor, we obtain

$$\tilde{P} := \bar{U}_{\mathbf{u}jA} L_{\mathbf{u}p} \tilde{U}_{\mathbf{u}p} = \begin{bmatrix} F_{\mathbf{u}} & B^T & Z_j & Z_A \\ B & \Delta S_p & & \\ & & M_j & K_{jA} \\ & & & \tilde{S}_A \end{bmatrix}, \quad (4.78)$$

where $\Delta S_p := \tilde{S}_p - S_p$. Notice that applying this simplification amounts to skipping step 1, so that one fewer space-time solve for $F_{\mathbf{u}}$ and one fewer diagonal solve for M_j are necessary for applying the preconditioner. A similar simplification in the single time-step framework is investigated in [25] as well. In addition, we perform one final approximation, and drop also the remaining block lower-triangular factor $L_{\mathbf{u}p}$ (*i.e.*, skip step 3), further reducing by one the total number of space-time solve

necessary for $F_{\mathbf{u}}$. With this, we end up with the block upper-triangular space-time preconditioner

$$P_T := \bar{U}_{\mathbf{u}jA} \tilde{U}_{\mathbf{u}p} = \begin{bmatrix} F_{\mathbf{u}} & B^T & Z_j & Z_A \\ & \tilde{S}_p & & \\ & & M_j & K_{jA} \\ & & & \tilde{S}_A \end{bmatrix}. \quad (4.79)$$

Considerations on the spectral properties of the preconditioner, together with some partial indications on how these are affected by dropping its block lower triangular terms, are provided in Sec. 4.3. Before addressing this issue, in the following section we briefly discuss on the cost associated with the application of the space-time block preconditioner introduced here, and on how it compares with its single time-step counterpart.

4.2.4 Comparison with single time-step block preconditioner

Ultimately, in order to judge whether the parallel application of the space-time block preconditioners introduced in Sec. 4.2.3 has the potential of reducing time-to-solution with respect to the sequential application of their single time-step equivalent, we need to compare the total cost (in computational time) associated with either strategy. To do so, we follow similar steps as in Sec. 3.2.1.1, and first identify the main parameters which determine the relative efficiency of one approach over the other. An estimate of these parameter is then provided in Sec. 4.4.3. To simplify our exposition, we focus our analysis on the block upper-triangular preconditioner P_T (4.79), but similar considerations hold for (4.75) and (4.78) as well.

To this purpose, let us begin by explicitly defining the single time-step equivalent of P_T . At the k -th time-step, this is given by

$$\mathcal{P}_{T,k} := \underbrace{\begin{bmatrix} I & \mathcal{Z}_{j,k} & \mathcal{Z}_{A,k} \\ & I & \\ & \mathcal{M}_j & \mathcal{K}_{jA} \\ & & \tilde{\mathcal{S}}_{A,k} \end{bmatrix}}_{=: \tilde{U}_{\mathbf{u}jA,k}} \underbrace{\begin{bmatrix} \mathcal{F}_{\mathbf{u},k} & \mathcal{B}^T & & \\ & \tilde{\mathcal{S}}_{p,k} & & \\ & & I & \\ & & & I \end{bmatrix}}_{=: \tilde{U}_{\mathbf{u}p,k}}, \quad (4.80)$$

where $\tilde{\mathcal{S}}_{p,k}$ is the single time-step pressure Schur complement approximation introduced in (3.39), while $\tilde{\mathcal{S}}_{A,k}$ is the single time-step magnetic field Schur complement approximation from [25, (3.32)], defined as

$$\tilde{\mathcal{S}}_{A,k} := \mathcal{M}_A \mathcal{F}_{A,k}^{-1} \tilde{\mathcal{C}}_{A,k}, \quad (4.81)$$

with $\tilde{\mathcal{C}}_{A,k}$ as in (4.72a). The two factors in (4.80) mimic the velocity-magnetic field and velocity-pressure couplings discussed in Sec. 4.2.1 and 4.2.2. In Sec. 3.2.1.1 we have already covered how the cost of inverting the rightmost factor in (4.80) in the single time-step setting compares with that of the whole space-time one, and we have reached the conclusion that, if one processor per time-step is available, the competitive advantage of the parallel method ultimately depends on two ratios: $C_{\mathcal{F}_u} \cdot N_t / C_{F_u}$, which provides an indication of the parallel efficiency at which we can solve the whole velocity space-time system F_u ; and N_{it}^{ST} / N_{it}^0 , which measures the overhead in terms of number iterations to convergence associated with preconditioning in space-time. We proceed by breaking down the cost comparison for the leftmost factor in (4.80) following a similar analysis. The total cost associated with the application of (4.80) is then given by summing these two contributions.

Inverting $\tilde{\mathcal{U}}_{ujA,k}$ requires solving for each of its diagonal blocks, as well as applying operators \mathcal{K}_{jA} , $\mathcal{Z}_{A,k}$ and $\mathcal{Z}_{j,k}$. When considering its space-time counterpart \bar{U}_{ujA} (4.77), a similar set of operations must be performed, except we need to deal with space-time operators. We have that M_j , K_{jA} , Z_A and Z_j are all block-diagonal, so their application (or that of their inverse, in the case of M_j) is trivially parallelisable over the time-steps: if N_t processors are available, we then expect this procedure to require a computational time comparable to that for the single time-step equivalent. The main difference rather lies in the block lower-triangular space-time operator \tilde{S}_A (4.73). By looking at the single factors composing it, we have that the application of its inverse requires: (i) inverting the space-time mass matrix M_A (which is block diagonal, and hence again trivially parallelisable); (ii) applying the space-time vector potential operator F_A (which is block *bi*-diagonal, and hence requires some overhead communication between the processors, but of marginal impact in terms of total cost, following considerations similar to those discussed for F_p (3.34) in Sec. 3.2.1.1); finally, (iii) inverting the space-time operator \tilde{C}_A (4.71), which represents the biggest obstacle in determining the parallel efficiency of the application of the space-time magnetic Schur complement approximation. We denote with $C_{\tilde{C}_A}$ the cost associated with operation (iii), and with $C_{\tilde{C}_A}$ the average cost per time-step associated with inverting its single time-step counterpart, $\tilde{\mathcal{C}}_{A,k}$ in (4.81). Combining this with the considerations above regarding $\tilde{\mathcal{U}}_{up,k}$, we have that the computational times necessary to apply P_T once, and $\mathcal{P}_{T,k}$ once *per time-step* are, respectively, $C_{F_u} + C_{\tilde{C}_A}$ and $(C_{F_u} + C_{\tilde{C}_A})N_t$. Therefore, one way to measure whether the space-time approach is competitive over

the time-stepping one consists in determining if

$$\frac{N_{it,L}^{ST}}{N_{it,L}^0} (C_{F_u} + C_{\tilde{C}_A}) < (C_{F_u} + C_{\tilde{C}_A}) N_t, \quad (4.82)$$

$N_{it,L}^{ST}$ and $N_{it,L}^0$ being the number of iterations to convergence required by the solver for the linearised system in the two cases, respectively (compare this to (3.40)).

However, an additional complication with respect to Sec. 3.2.1.1, is given by the nonlinearities present in the IMHD target system of this chapter (4.23). As a consequence, another relevant metric in determining the effectiveness of the space-time approach over the time-stepping one is given by the relative cost and number of iterations associated with the Newton solver. For each outer iteration, the main task consists in re-assembling the operators composing the discrete Jacobian (4.40). Most of these are either constant (and hence can be assembled once and for all during the start-up phase of the solver), or depend on values on a single time-step (and hence their assembly can be trivially parallelised, at no overhead cost with respect to time-stepping). One exception to this, however, is given by the space-time operator \tilde{C}_A . In fact, the k -th block in its first block subdiagonal (defined in (4.72b)) depends on values both at instant k and $k + 1$. Therefore, if the solver chosen for this space-time operator requires explicit assembly of this block, its construction might require significant additional computational time with respect to inverting its single time-step equivalent (even if it does not, additional communication between processors might be required, likely increasing the cost for inverting \tilde{C}_A). To keep the analysis in this chapter contained, we do not explore this issue in detail, and limit ourselves to track also the relative number of Newton iterations for the space-time and the time-stepping approaches, denoted as $N_{it,NL}^{ST}$ and $N_{it,NL}^0$ respectively.

Finally, also in terms of memory usage the situation is analogous to the one described in Sec. 3.2.1.1: even in this case, in fact, in our implementation we associate each processor with a different time-step, and have it assemble its corresponding spatial operators. As in the previous chapter, this results in the total memory required scaling linearly with the number of processors employed, which is in line with other PinT approaches.

4.3 Eigenvalues clustering

In this section we analyse the eigenvalue distribution for the discrete operator (4.40), right-preconditioned with the approximate block LU factorisation introduced in (4.45).

We do not consider here the effect of the approximations to the space-time Schur complements discussed in Sec. 4.2.1 and 4.2.2, and rather stick to the ideal case where the exact S_p and S_A operators ((4.46) and (4.53)) are used, which gives us a best-case scenario of the performance of the preconditioner. The goal is once again to provide some evidence for the effectiveness of this choice of preconditioner, although similar caveats as those made for Sec. 3.3 hold here as well: eigenvalues do not tell the whole picture when it comes to convergence of GMRES [68], especially when dealing with nonsymmetric matrices the likes of which compose our target system (4.44). Similarly to Sec. 3.3, we show that the eigenvalues distribution of the space-time preconditioned system is tightly connected to that of its single time-step counterpart.

Using the definitions of (4.40) and (4.45), we see that the preconditioned system is given by

$$J_{\mathcal{N}}P^{-1} = \left[\begin{array}{cc|cc} F_{\mathbf{u}} & B^T & Z_j & Z_A \\ B & & & \\ \hline & & M_j & K_{jA} \\ Y & & & F_A \end{array} \right] \left[\begin{array}{cc|cc} F_{\mathbf{u}} & B^T & Z_j & Z_A \\ B & & & \\ \hline & & M_j & K_{jA} \\ Y & YF_{\mathbf{u}}^{-1}B^T & & F_A \end{array} \right]^{-1}. \quad (4.83)$$

Notice we highlighted the 2×2 block structure of the two factors: they can in fact be more compactly expressed as

$$\left[\begin{array}{cc} A & B \\ C & D \end{array} \right] \left[\begin{array}{cc} A & B \\ \tilde{C} & D \end{array} \right]^{-1}, \quad (4.84)$$

whose common composing blocks are given by

$$A = \left[\begin{array}{cc} F_{\mathbf{u}} & B^T \\ B & \end{array} \right], \quad B = \left[\begin{array}{cc} Z_j & Z_A \end{array} \right], \quad \text{and} \quad D = \left[\begin{array}{cc} M_j & K_{jA} \\ & F_A \end{array} \right], \quad (4.85)$$

while the bottom-left corner differs, being

$$C = \left[\begin{array}{c} Y \end{array} \right], \quad \text{and} \quad \tilde{C} = \left[\begin{array}{cc} Y & YF_{\mathbf{u}}^{-1}B^T \end{array} \right], \quad (4.86)$$

for the matrix on the left and on the right, respectively. After some manipulations, we see that explicitly multiplying the factors in (4.83) results in the following block lower-triangular matrix:

$$\left[\begin{array}{cc} I_{N_t(N_{\mathbf{u}}+N_p)} & \\ CA^{-1} - S_D S_{\tilde{D}}^{-1} \tilde{C} A^{-1} & S_D S_{\tilde{D}}^{-1} \end{array} \right]. \quad (4.87)$$

This already identifies a number $N_t(N_{\mathbf{u}} + N_p)$ of unitary eigenvalues: due to the triangular nature of the matrix, the remaining eigenvalues are given by those of the

bottom-right block. The two factors composing it are provided, respectively, by

$$\begin{aligned}
S_D &= D - CA^{-1}B = \begin{bmatrix} M_j & K_{jA} \\ & F_A \end{bmatrix} - \begin{bmatrix} Y & \\ & \end{bmatrix} \begin{bmatrix} F_u & B^T \\ & B \end{bmatrix}^{-1} \begin{bmatrix} Z_j & Z_A \end{bmatrix} \\
&= \begin{bmatrix} M_j & K_{jA} \\ \underbrace{-YF_u^{-1}(I + B^T S_p^{-1} B F_u^{-1}) Z_j}_{=:E} & \underbrace{F_A - YF_u^{-1}(I + B^T S_p^{-1} B F_u^{-1}) Z_A}_{=:F} \end{bmatrix}, \tag{4.88}
\end{aligned}$$

where S_p is the familiar space-time pressure Schur complement (4.46), and

$$\begin{aligned}
S_{\tilde{D}} &= D - \tilde{C}A^{-1}B = \begin{bmatrix} M_j & K_{jA} \\ & F_A \end{bmatrix} - \begin{bmatrix} Y & YF_u^{-1}B^T \\ & \end{bmatrix} \begin{bmatrix} F_u & B^T \\ & B \end{bmatrix}^{-1} \begin{bmatrix} Z_j & Z_A \end{bmatrix} \\
&= \begin{bmatrix} M_j & K_{jA} \\ \underbrace{E + YF_u^{-1}B^T S_p^{-1} B F_u^{-1} Z_j}_{\tilde{E}} & \underbrace{F + YF_u^{-1}B^T S_p^{-1} B F_u^{-1} Z_A}_{\tilde{F}} \end{bmatrix}. \tag{4.89}
\end{aligned}$$

These too have a 2×2 block structure, and their product is in the form

$$\begin{bmatrix} M_j & K_{jA} \\ E & F \end{bmatrix} \begin{bmatrix} M_j & K_{jA} \\ \tilde{E} & \tilde{F} \end{bmatrix}^{-1}. \tag{4.90}$$

After some more calculations, their multiplication produces yet another block lower-triangular matrix,

$$\begin{bmatrix} I_{N_t N_j} & \\ EM_j^{-1} - S_F S_{\tilde{F}}^{-1} \tilde{E} M_j^{-1} & S_F S_{\tilde{F}}^{-1} \end{bmatrix}, \tag{4.91}$$

which allows us to rule out some additional $N_t N_j$ unitary eigenvalues. Just like for (4.87), even in this case the remaining eigenvalues coincide with those of the bottom-right block of the matrix. Another round of algebraic manipulations gives us an explicit formula for this operator:

$$\begin{aligned}
S_F S_{\tilde{F}}^{-1} &= (F - EM_j^{-1} K_{jA})(\tilde{F} - \tilde{E} M_j^{-1} K_{jA})^{-1} \\
&= I_{N_t N_A} - YF_u^{-1} B^T S_p^{-1} B F_u^{-1} (Z_A - Z_j M_j^{-1} K_{jA}) S_A^{-1}, \tag{4.92}
\end{aligned}$$

where S_A is the space-time magnetic Schur complement introduced in (4.53). Each of the space-time operators composing (4.92) has either a block diagonal or a block lower-triangular structure, implying that in order to recover its eigenvalues it suffices to identify those of the operators appearing in its diagonal blocks. These in turn are given by

$$I_{N_A} - \mathcal{Y}_k \mathcal{F}_{u,k}^{-1} \mathcal{B}^T \mathcal{S}_{p,k}^{-1} \mathcal{B} \mathcal{F}_{u,k}^{-1} (\mathcal{Z}_{A,k} - \mathcal{Z}_{j,k}, \mathcal{M}_j^{-1} \mathcal{K}_{jA}) \mathcal{S}_{A,k}^{-1}, \quad \forall k = 1, \dots, N_t, \tag{4.93}$$

where $\mathcal{Z}_{j,k}$ and $\mathcal{Z}_{A,k}$ are the Lorentz force linearisation terms, following (4.30). By considering how we compact these two terms into (4.56), and adjusting for the notation used in their article, we can convince ourselves that operator (4.93) corresponds to the one described in [25, (3.10)], which ultimately dictates the clustering properties of their preconditioner. As the spectrum of the single time-step preconditioned system is analysed in detail in [25, Fig. 3.1], we do not repeat their analysis here, and rather limit ourselves to pointing out how even in the application to IMHD considered in this chapter, the eigenvalues distribution of the whole space-time preconditioned system ends up depending directly on that of its single time-step counterpart, similarly to what described in Sec. 3.3.

This remains valid if we decide to simplify (4.75) by dropping its first block lower-triangular term L_{ujA} , as discussed at the end of Sec. 4.2.3. In fact, if we consider \tilde{P} from (4.78) as a simplification to P , the preconditioned system is modified to

$$\begin{aligned} J_{\mathcal{N}}\tilde{P}^{-1} &= J_{\mathcal{N}}\tilde{U}_{up}^{-1}L_{up}^{-1}F_{uI}\tilde{U}_{ujA}^{-1} = J_{\mathcal{N}}(\tilde{U}_{up}^{-1}L_{up}^{-1}F_{uI}\tilde{U}_{ujA}^{-1}L_{ujA}^{-1})L_{ujA} \\ &= J_{\mathcal{N}}P^{-1} \begin{bmatrix} I & & & & \\ & I & & & \\ & & I & & \\ YF_u^{-1} & & -YF_u^{-1}Z_jM_j^{-1} & & I \end{bmatrix}. \end{aligned} \quad (4.94)$$

We have already established how, in the ideal case where we pick the exact pressure and magnetic Schur complements, $J_{\mathcal{N}}P^{-1}$ gives rise to the block lower-triangular operator (4.91). The matrix we are now further multiplying by, L_{ujA} , is also block lower-triangular, and presents identities along its whole diagonal: from this, we safely conclude that excluding this factor has no effect on the eigenvalue distribution of the original preconditioned system.

Unfortunately, dropping the second block lower-triangular term $L_{u,p}$ as well has a less clear effect on the spectral properties of the resulting preconditioned system, and the corresponding analysis becomes more challenging. In the scope of this thesis, we do not investigate this further, and rather showcase the effectiveness of this additional simplification with some numerical experiments, which are presented in the following section.

4.4 Results

In this section, we experiment with the effectiveness of the preconditioner introduced in Sec. 4.2 by employing it for the acceleration of the solution of some model problems describing the arising and evolution of instabilities in resistive plasma. In the

framework of resistive plasmas, the instabilities described here result in modifications to the topology of the magnetic field lines, giving rise to separation and reconnection phenomena, usually associated with the generation of current sheets [66, Chap. 9.6]. These phenomena have been a fascinating subject of study over the last decades, and research on this area has undergone considerable effort: providing a complete overview of the topic is beyond the scope of this thesis, and we rather restrict ourselves to giving a general description of the dynamics triggered within the problems considered in Sec. 4.4.1. Nonetheless, we refer the interested reader to [66, Chap. 9 and Chap. 20], as well as to the excellent review on the topic by Dieter Biskamp, particularly [9, Chap. 3-4].

We point out that most of the results in this section are still preliminary: additional effort needs to be put into the design of the space-time preconditioner before it can reach the maturity necessary for the application to real-world scenarios. Particularly, we do not discuss its performance for small values of μ and η , which require *ad-hoc* stabilisation, and will be a matter for future research. Nonetheless, the experiments conducted here serve as proof of concept to validate the effectiveness of the preconditioner developed in Sec. 4.2, and more generally to sustain the claim for the possibility to extend the framework of space-time block preconditioning to more challenging systems of PDEs.

4.4.1 Model problems

The problems introduced in this section have been often used in the literature in order to test algorithms for MHD: see for example [15, 114] for problem 1, and [25, 114, 134] for problem 2, whose examples we follow in our set-up. We proceed to describe these in detail next.

Problem 1 *Tearing mode.* This is one of the classic instabilities in resistive plasmas. It arises from the interaction of magnetic fields with opposite orientations: in the presence of resistivity, the opposite fields are allowed to “diffuse” into each other, thus progressively weakening until annihilation occurs. At this stage, the magnetic field lines running in one direction reconnect with those in the opposite direction, and produce a so-called *x-point*, where the lines intersect (as opposed to *o-points*, around which the lines rotate). This evolution is dictated by nonlinear interactions, and it occurs on a relatively large time-scale; nonetheless, the solution eventually stabilises, at which point it is said to have reached *saturation*. An example solution

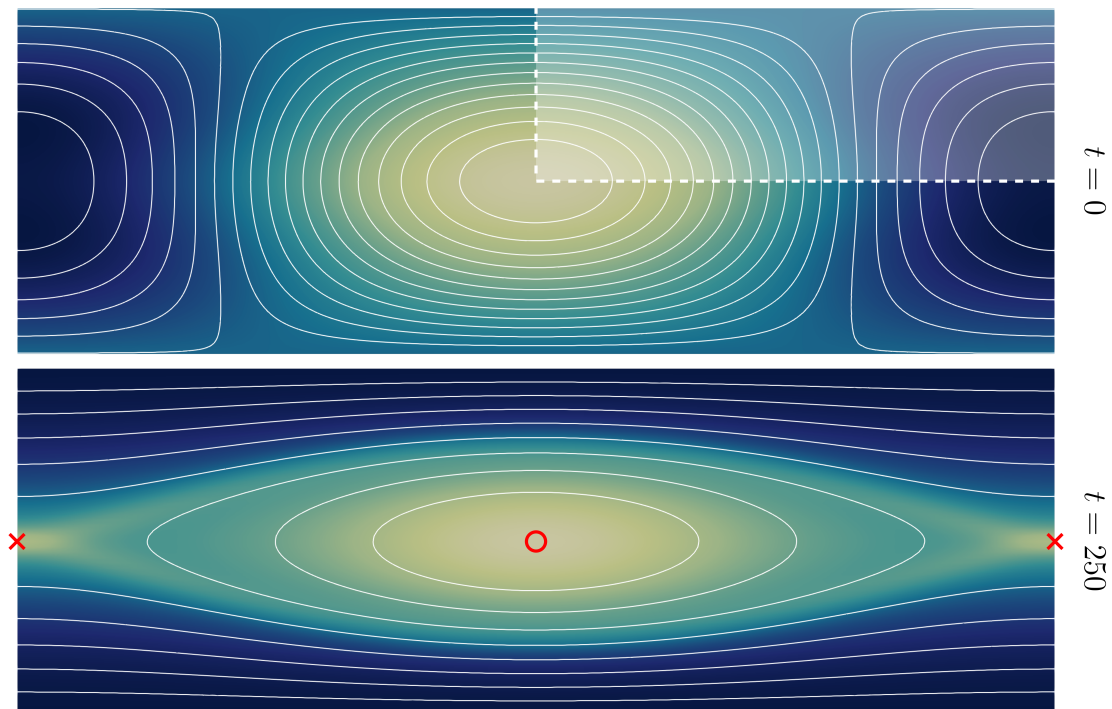


Figure 4.1: Evolution of contour plots of current intensity with superimposed isolines of the magnetic vector potential A (*i.e.*, streamlines of the magnetic field), for an example solution of problem 1, with $\mu = \eta = 10^{-3}$ and $\mu_0 = 1$. The top plot shows the initial conditions, with the current flowing alternately outside the domain (light area) and inside the domain (dark areas) as we move along the x -axis, and the resulting magnetic field lines rotating counter- and clock-wise, respectively; the shaded area highlights the actual domain of the simulation. As time proceeds, the magnetic field lines at the sides of the domain get progressively squashed and eroded, and eventually annihilated. Notice that this causes the lines to intersect at the middle of the left and right sides of the domain (x -points, highlighted in red together with the o -points). The solution at saturation is presented in the bottom plot.

of this problem, illustrating the dynamics described above, is plotted in Fig. 4.1; the corresponding velocity and pressure fields are instead reported in Fig. 4.2.

The system configuration for this first test-case is as follows. The flow is initially at quiet state,

$$\bar{\mathbf{u}}_{TM}^0(\mathbf{x}) = 0, \quad (4.95)$$

and the vector potential is at the *Harris sheet equilibrium* [9, Chap. 4.1], that is

$$A_{TM}^{eq}(\mathbf{x}) = \frac{1}{\lambda} \ln(\cosh(\lambda y)), \quad (4.96)$$

for a given λ . This must be sustained by the external electric field

$$E_{TM}^{eq}(\mathbf{x}) = \frac{\eta}{\mu_0} \nabla^2 A_{TM}^{eq}(\mathbf{x}) = \frac{\lambda \eta}{\mu_0} \cosh(\lambda y)^{-2}, \quad (4.97)$$




Figure 4.2: Velocity magnitude and streamlines (top), and pressure contour plot (bottom) for an example solution of the flow in the tearing mode problem 1 at saturation $t = 250$, with $\mu = \eta = 10^{-3}$ and $\mu_0 = 1$. The flow is accelerated in a vortex by Lorentz force: the vorticity of the field has largest intensity at the *separatrix* (the border of the central “eye” in Fig. 4.1, where the current gradient is largest). Only the simulation domain (top-right corner in Fig. 4.1) is shown.

which is imposed as a right-hand side for the vector potential equation in (4.19). The action of $A_{TM}^{eq}(\mathbf{x})$ on the flow (via Lorentz force) is counterbalanced by the pressure term

$$p_{TM}^{eq}(\mathbf{x}) = \frac{1}{2\mu_0} \cosh(\lambda y)^{-2} - \bar{p}_{TM}^{eq}, \quad (4.98)$$

where the constant \bar{p}_{TM}^{eq} is selected to ensure that the pressure has zero mean at equilibrium. In order to excite the tearing mode, we perturb the equilibrium by considering the following initial conditions for the vector potential:

$$\bar{A}_{TM}^0(\mathbf{x}) = A_{TM}^{eq}(\mathbf{x}) - \varepsilon \cos(\pi y) \cos\left(\frac{2\pi x}{L}\right), \quad (4.99)$$

with L being the x -length of the domain. The resulting solution is periodic in $x \in [-L, L]$, and shows alternating x - and o -points at saturation; the “eye” around the o -point is referred to as a *magnetic island*. The solution also presents axial symmetries, so instead of considering the whole spatial domain $[-L, L] \times [-1/2, 1/2]$, we can simulate only its top-right corner $\Omega = [0, L] \times [0, 1/2]$, and flip its solution around

both the x and y axes. As such, we prescribe symmetric boundary conditions on sides $x = 0$, $x = L$, and $y = 0$; on the top side $y = 1/2$, instead, we prescribe Dirichlet BC for A (fixed at the equilibrium value (4.96)), and slip conditions for \mathbf{u} . Notice the resulting flow is fully enclosed, so that the considerations detailed in Sec. 3.2.1.2 on the assembly of the approximate space-time pressure Schur complement \tilde{S}_p hold in this case as well. The parameters used in our problem are $\lambda = 5$, $\varepsilon = 10^{-3}$, and $L = 3$.

Problem 2 *Island coalescence.* This problem describes the evolution of two magnetic islands in a plasma, as they could stem from the tearing mode perturbation introduced in problem 1: notice the final condition in Fig. 4.1 roughly matches the initial condition in half of the domain in Fig. 4.3. The two σ -points are pulled together by Lorentz force, until eventually the x -point in between them separates and the magnetic islands merge into one (that is, they *coalesce*), generating an electric discharge; the specific dynamics depends on the value of the magnetic resistivity, but generally reconnection occurs in a time-scale much smaller than that for problem 1; an example of the evolution of this system is provided in Fig. 4.3 and 4.4.

As in problem 1, we start from a quiet state,

$$\bar{\mathbf{u}}_{IC}^0(\mathbf{x}) = 0. \quad (4.100)$$

with the vector potential given by a modification to the Harris sheet equilibrium in (4.96). This is at times referred to as *Fadeev* or *Schmid-Burgk equilibrium*, since they pioneered its study in [45, 131], or also as *corrugated sheet pinch equilibrium* [9, Chap. 4.3], and is given by

$$A_{IC}^{eq}(\mathbf{x}) = \frac{1}{2\pi} \ln [\cosh(2\pi y) + \beta \cos(2\pi x)], \quad (4.101)$$

for a given parameter β . This too must be sustained by an external electric field, in this case given by

$$E_{IC}^{eq}(\mathbf{x}) = \frac{\eta}{\mu_0} \nabla^2 A_{IC}^{eq}(\mathbf{x}) = \frac{\eta}{\mu_0} \frac{2\pi(1 - \beta^2)}{(\cosh(2\pi y) + \beta \cos(2\pi x))^2}. \quad (4.102)$$

The flow equilibrium is maintained by a counterbalancing pressure,

$$p_{IC}^{eq}(\mathbf{x}) = \frac{1}{2\mu_0} \frac{1 - \beta^2}{(\cosh(2\pi y) + \beta \cos(2\pi x))^2} - \bar{p}_{IC}^{eq}, \quad (4.103)$$

where again the quantity \bar{p}_{IC}^{eq} is chosen to ensure the equilibrium pressure is zero-mean valued. This equilibrium is perturbed by considering as initial conditions for




Figure 4.3: Evolution of contour plots of current intensity with superimposed isolines of the magnetic vector potential A (*i.e.*, streamlines of the magnetic field), for an example solution of problem 2, with $\mu = \eta = 10^{-3}$ and $\mu_0 = 1$. The top plot shows the initial conditions, with the magnetic islands induced by the perturbation; the shaded area highlights the actual domain of the simulation (notice the y -axis has been cropped to better zoom on the islands). As time proceeds, the islands start attracting each other until reconnection occurs in the middle plot, where we see the associated electric discharge. In the last plot the reconnection has completed, and the two islands have coalesced into one.

the vector potential

$$\bar{A}_{IC}^0(\mathbf{x}) = A_{IC}^{eq}(\mathbf{x}) + \varepsilon \cos\left(\frac{\pi}{2}y\right) \cos(\pi x). \quad (4.104)$$

Similarly to problem 1, in this case too the resulting solution is periodic in $x \in [0, 2]$ and presents a number of symmetries, which we exploit to simplify our simulation and reduce the domain from $[0, 2] \times [-1, 1]$ to its top-left corner $\Omega = \Omega_{\square} = [0, 1]^2$. We impose symmetry on the $x = 0$, $x = 1$ and $y = 0$ axes, while on $y = 1$, we prescribe Dirichlet BC for A (again fixed to the equilibrium (4.101)), and slip conditions for \mathbf{u} ,




Figure 4.4: Velocity magnitude and streamlines (left), and pressure contour plot (right) for an example solution of the flow in the island coalescence problem 2 at reconnection $t = 7$, with $\mu = \eta = 10^{-3}$ and $\mu_0 = 1$. Only the simulation domain (top-left corner in Fig. 4.3) is shown.

rendering the flow fully enclosed also in this case. The parameters considered for this problem are $\beta = 0.2$ and $\varepsilon = 10^{-3}$.

In our experiments, the choice of finite-element approximation for the discretisation of the functional spaces (3.10) and (4.27) falls on the Taylor-Hood $P_3 - P_2$ pair for velocity and pressure [118, Chap. 17.4], while we consider P_1 for both the vector potential and the current intensity variables. This stems from considerations on *Finite Element Exterior Calculus* [4, 79], and from the necessity of ensuring that the two forcing terms in the momentum equation, ∇p and $j\nabla A$, can balance each other out at equilibrium; this choice is otherwise rather empirical, and an additional theoretical justification is needed to show whether these FE spaces indeed satisfy the required stability conditions.

Regarding the parameters for the problems, we put ourselves in a somewhat favourable position, by considering only $\mu = \eta = \mu_0 = 1$. While this choice has the effect of rendering some of the dynamics described in problems 1 and 2 less evident, it allows us to neglect matters regarding stability under advection-dominated regimes, which is beyond the purpose of this analysis, and rather focus on providing a general indication of what kind of performance one might expect from the space-time preconditioner.

The actual code used for the numerical simulations is an expansion of the one developed for Sec. 3.4: as such, it still relies on MFEM and PETSc [7, 102] for the

linear solvers necessary to tackle the linearised system (4.44) at each Newton iteration. As an initial guess for the discrete space-time velocity, pressure, and vector potential, we pick the projection of the equilibrium conditions, (4.95), (4.98), (4.96) and (4.100), (4.103), (4.101), onto their respective discrete finite element spaces (3.10) and (4.27); the additional variable j is instead initialised as the solution to its associated discrete system (4.42c):

$$\mathcal{M}_j \mathbf{j}^k = -\mathcal{K}_{jA} \mathbf{A}^k + \mathbf{h}^k, \quad \forall k = 1, \dots, N_t, \quad (4.105)$$

where \mathbf{A}^k is the initial guess on the vector potential at instant k , so that the initial residual associated with this variable is effectively 0.

In the following, we report some preliminary results from the application of the space-time preconditioner.

4.4.2 Performance of preconditioner

This section is dedicated to measuring the effectiveness of the application of the space-time preconditioner introduced in Sec. 4.2.3 for accelerating the iterative solution of the linearised IMHD system (4.44). In analogy with the previous chapter, in order to evaluate the preconditioner performance, we focus on tracking the number of iterations necessary to achieve convergence, for both the outer Newton solver and the inner GMRES solver. This preference is motivated by similar considerations as those made for Sec. 3.4: for the purpose of this study, this parameter represents a more indicative measure of the behaviour of the preconditioner than, *e.g.*, time-to-solution, as the latter strongly depends on implementation choices. This becomes even more significant given the possible challenges represented by step 2, and particularly by the inversion of \tilde{C}_A in (4.71), whose effective time parallelisation we do not investigate further in this chapter.

For similar reasons, in this preliminary study we consider only the use of *exact* solvers for inverting the relevant block operators appearing in P_T . That is to say, we employ sequential time-stepping for the solution of the space-time operators $F_{\mathbf{u}}$ and \tilde{C}_A , and use LU factorisation for inverting the single blocks in their block diagonals $\mathcal{F}_{\mathbf{u},k}$ and $\tilde{C}_{A,k}$ (see their definitions in (4.35) and (4.71)), as well as the vector potential mass matrix \mathcal{M}_A (appearing within M_A in (4.73)), the current intensity mass matrix \mathcal{M}_j (present in M_j in (4.74)), and the pressure mass and stiffness operators $\mathcal{M}_p, \tilde{\mathcal{K}}_p$ (included in (3.37) and discussed in Sec. 3.2).

Notice moreover that the results included in this section only consider the application of the *simplified* block upper-triangular preconditioner P_T in (4.79), since

Table 4.1: Iterations to convergence for Newton method applied for the solution of problems 1 and 2, for various refinement levels of the spatial and temporal grids (the temporal domain is fixed at $T = 1$). The outer solver achieves convergence with a tolerance of 10^{-10} . The inner GMRES solver is right-preconditioned with P_T (4.79), and converges with relative tolerance 10^{-2} or absolute tolerance of 10^{-14} . To invert the relevant operators appearing in P_T , we use exact solvers (time-stepping and LU). Inside each column, the number of Newton iterations to convergence is reported on the left, the average number of inner GMRES iterations per outer Newton iteration appears on the right (in brackets).

Pb	$\frac{\Delta t \rightarrow}{\Delta x \downarrow}$	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	2^{-7}
1	2^{-2}	5 (7.80)	5 (9.40)	5 (11.60)	5 (13.20)	5 (13.60)	5 (16.40)
	2^{-3}	5 (8.40)	5 (10.20)	5 (13.00)	5 (14.80)	5 (15.00)	5 (16.60)
	2^{-4}	5 (8.00)	5 (10.40)	5 (12.20)	5 (13.80)	5 (15.40)	5 (17.20)
	2^{-5}	5 (7.40)	5 (9.40)	5 (12.00)	5 (13.40)	5 (14.00)	5 (18.40)
	2^{-6}	5 (7.60)	5 (9.40)	5 (12.00)	5 (13.40)	5 (13.80)	5 (17.00)
	2^{-7}	5 (7.60)	5 (9.80)	5 (11.80)	5 (13.20)	5 (14.40)	5 (16.80)
2	2^{-2}	4 (8.75)	5 (11.00)	5 (11.80)	5 (14.20)	4 (15.75)	4 (17.00)
	2^{-3}	5 (9.20)	4 (11.50)	4 (13.00)	4 (14.75)	4 (16.75)	4 (18.75)
	2^{-4}	4 (9.75)	4 (11.25)	4 (13.50)	4 (15.50)	4 (17.25)	4 (19.25)
	2^{-5}	4 (8.00)	4 (11.00)	4 (13.75)	4 (15.25)	3 (15.67)	3 (20.33)
	2^{-6}	3 (8.00)	3 (10.00)	3 (12.33)	3 (14.00)	3 (15.67)	3 (17.33)
	2^{-7}	3 (7.67)	3 (10.67)	3 (12.33)	3 (14.00)	3 (16.00)	3 (17.67)

early experiments showed negligible difference in convergence behaviour with respect to using its *full* version P from (4.75) (which makes P_T a much more appealing choice, given how its associated computational cost is significantly reduced, as per the discussion in Sec. 4.2.3).

In Tab. 4.1 we report convergence results for both inner and outer solvers, collected for various choices of spatial and temporal grid sizes Δx and Δt , on a fixed temporal domain of size $T = 1$. There, we see how the number of Newton iterations necessary for convergence remains substantially unchanged, regardless of the mesh chosen. The number of GMRES iterations seems to be increasing slightly with temporal refinement, in a somewhat more evident manner than for the fluid dynamics problems in Tab. 3.1, but the rate at which this occurs remains reasonably bounded (increasing the number of temporal nodes by a factor of 2^5 results in slightly more than ~ 2 times as many iterations). To put these numbers into perspective, notice that for these experiments the spatial unknowns for the most refined discretisation of problem 1 amount to $N_u = 222,722$, $N_p = 49,665$, $N_j = N_A = 12,545$, while those for problem 2 reach $N_u = 296,450$, $N_p = 66,049$, $N_j = N_A = 16,641$. Considering that the temporal grids count up to $N_t = 128$ nodes, this results in space-time sys-

Table 4.2: Convergence results for Newton iterations applied for the solution of problems 1 and 2, for various refinement levels of the spatial grid. We fix $\Delta t = 0.5$, and progressively enlarge the temporal domain to accommodate for more processors. The solvers used are as in Tab. 4.1. Inside each column, the number of Newton iterations to convergence is reported on the left, the average number of inner GMRES iterations per outer Newton iteration appears on the right (in brackets).

Pb	$\frac{T \rightarrow}{\Delta x \downarrow}$	2^2		2^3		2^4		2^5		2^6		2^7	
1	2^{-2}	5	(5.80)	5	(6.20)	5	(6.20)	5	(6.60)	5	(6.40)	5	(6.40)
	2^{-3}	5	(6.80)	5	(6.80)	5	(6.80)	5	(6.80)	5	(6.80)	5	(6.80)
	2^{-4}	5	(6.60)	5	(6.60)	5	(6.60)	5	(6.60)	5	(6.60)	5	(6.80)
	2^{-5}	4	(6.00)	5	(6.20)	4	(6.25)	4	(6.25)	5	(6.40)	5	(6.40)
	2^{-6}	4	(6.00)	4	(6.00)	4	(6.00)	4	(6.00)	4	(6.25)	4	(6.25)
	2^{-7}	4	(5.75)	4	(6.00)	4	(6.00)	4	(6.00)	4	(6.00)	4	(6.00)
	2	2^{-2}	5	(6.40)	5	(6.40)	5	(6.40)	5	(6.40)	5	(6.40)	5
2^{-3}		4	(6.75)	5	(7.00)	5	(7.00)	5	(7.00)	5	(7.00)	5	(7.00)
2^{-4}		4	(6.25)	4	(6.25)	4	(6.50)	4	(6.75)	5	(6.80)	5	(6.80)
2^{-5}		4	(6.50)	4	(6.50)	4	(6.75)	4	(6.50)	4	(6.50)	4	(6.50)
2^{-6}		3	(6.33)	4	(6.75)	4	(6.75)	4	(6.75)	4	(6.75)	4	(6.75)
2^{-7}		3	(6.33)	3	(6.33)	3	(6.33)	3	(6.33)	4	(6.75)	3	(6.33)

tems of total size $(N_{\mathbf{u}} + N_p + N_j + N_A) \cdot N_t \approx 38 \cdot 10^6$ and $\approx 52 \cdot 10^6$ for the two problems, respectively.

The temporal domain $T = 1$ considered for the experiments in Tab. 4.1 is however too small to allow for the relevant nonlinear behaviours described in Sec. 4.4.1 to manifest themselves into the solution (notice the time labels on the side of Fig. 4.1 and 4.3). To get an idea of how the preconditioner performs when these become evident, we consider another set of experiments, where we fix the time-step and progressively increase the size of the time domain. The corresponding convergence results are reported in Tab. 4.2. Overall, they draw a similar picture as those in Tab. 4.1, with the number of both outer and inner iterations kept bounded in all cases. In this test, though, we do not observe the slight increase in iterations count present in Tab. 4.1 as we increase the domain size (indeed, iterations number barely changes at all), providing evidence that the performance of the preconditioner is affected by the actual value of Δt , rather than by an increase in the number of processors available.

4.4.3 Comparison with sequential time-stepping

In this section, we expand on the discussion in Sec. 3.2.1.1 and provide an indication of how the number of inner and outer iterations to convergence for the space-time

Table 4.3: Ratios $N_{it,NL}^{ST}/N_{it,NL}^0$ and $N_{it,L}^{ST}/N_{it,L}^0$, where $N_{it,NL}^{ST}$ and $N_{it,L}^{ST}$ are the total number of outer and inner iterations to convergence for the space-time approach (resulting from the same experiments used to fill Tab. 4.1), while $N_{it,NL}^0$ and $N_{it,L}^0$ represent the average number of outer and inner iterations to convergence per time-step, using time-stepping. For the latter, the inner solver for each time-step is GMRES, right-preconditioned with the single time-step counterpart of the block upper-triangular preconditioner, (4.80). The outer solver for time-stepping reaches convergence with absolute tolerance $10^{-10}/\sqrt{N_t}$, while the inner solver has the same tolerances used for Tab. 4.1. For some of the most extreme parameters combinations, we noticed that the solution from time-stepping “freezes” from a certain (reasonably late) instant onwards: that is, the residual after the Newton update becomes smaller than the prescribed tolerance, thus preventing any other update from being carried out (notice we do not implement mesh adaptivity). To account for this, we compute the averages $N_{it,NL}^0$ and $N_{it,L}^0$ by considering the *effective* time-steps only.

Pb	$\frac{\Delta t \rightarrow}{\Delta x \downarrow}$	2^{-2}		2^{-3}		2^{-4}		2^{-5}		2^{-6}		2^{-7}	
1	2^{-2}	1.05	(1.05)	1.14	(1.18)	1.31	(1.45)	1.38	(1.54)	1.63	(1.97)	1.66	(2.41)
	2^{-3}	1.11	(1.17)	1.18	(1.29)	1.27	(1.54)	1.45	(1.80)	1.54	(2.15)	1.65	(2.40)
	2^{-4}	1.11	(1.21)	1.21	(1.40)	1.40	(1.65)	1.47	(1.82)	1.50	(2.35)	1.65	(2.73)
	2^{-5}	1.11	(1.13)	1.25	(1.38)	1.45	(1.79)	1.55	(1.93)	1.66	(2.51)	1.86	(3.50)
	2^{-6}	1.11	(1.21)	1.33	(1.50)	1.48	(1.89)	1.60	(2.06)	1.72	(2.71)	1.91	(3.52)
	2^{-7}	1.25	(1.42)	1.29	(1.56)	1.51	(1.94)	1.68	(2.21)	1.77	(3.06)	2.01	(3.91)
	2	2^{-2}	1.07	(1.12)	1.54	(1.73)	1.57	(1.78)	1.72	(2.25)	1.48	(2.13)	1.59
2^{-3}		1.54	(1.79)	1.39	(1.76)	1.60	(2.05)	1.60	(2.41)	1.70	(3.25)	1.79	(3.70)
2^{-4}		1.23	(1.53)	1.45	(1.90)	1.60	(2.23)	1.78	(2.68)	1.95	(4.36)	2.18	(5.26)
2^{-5}		1.45	(1.60)	1.45	(1.96)	1.78	(2.67)	2.03	(3.19)	1.60	(3.40)	1.90	(5.36)
2^{-6}		1.09	(1.33)	1.26	(1.63)	1.55	(2.20)	1.57	(2.38)	1.79	(4.21)	1.98	(5.20)
2^{-7}		1.20	(1.39)	1.41	(2.00)	1.60	(2.38)	1.75	(2.83)	1.88	(5.01)	2.04	(5.99)

preconditioner P_T , ($N_{it,L}^{ST}$ and $N_{it,NL}^{ST}$), compare with those of its single time-step counterpart $\mathcal{P}_{T,k}$, ($N_{it,L}^0$ and $N_{it,NL}^0$), for the test cases introduced in Sec. 4.4.1.

In order to conduct a balanced comparison between the single time-step and the whole space-time approaches, in Sec. 3.4.4 we adjusted the absolute tolerance of the linear solver for the former by scaling it by a factor $\sqrt{N_t}$. In this section, we do the same for the *outer* solver, but leave that of the inner solver untouched (this might slightly favour the time-stepping approach, but convergence of GMRES occurred due to the *relative* tolerance being met in almost the totality of cases). Moreover, to more faithfully represent a real-case scenario, when time-stepping we reuse the value of the solution at the previous instant as initial guess for the Newton solver at the following time-step, analogously to what has been done in Sec. 3.4.4.

Given the remarks in Sec. 4.2.4, it is unclear whether a relative increase in *outer* iterations is more detrimental in raising the total computational time with respect to an increase in *inner* iterations, or if the converse holds instead. Regardless, the

fact remains that both ratios $N_{it,L}^{ST}/N_{it,L}^0$ and $N_{it,NL}^{ST}/N_{it,NL}^0$ must be kept as close as possible to 1 for the space-time preconditioning approach to represent an appealing alternative to sequential time-stepping. The results in Tab. 4.3 show that this condition is satisfied reasonably well: the ratio of Newton iterations hits a maximum of ~ 2 , while that of GMRES iterations sits between 1 and 3 for most cases. As $\Delta t \rightarrow 0$, this ratio tends to increase, reaching almost 6 in certain instances. This is not surprising, as in this case the improved quality of the initial guess for Newton when using time-stepping (see the discussion in Sec. 3.4.4), and the slight increase in iterations count observed in Tab. 4.1 both play against the space-time approach.

Overall, the results in this section showcase the potential effectiveness of our STBP approach also in this more complex MHD case. At least when employing *exact* solvers, the space-time preconditioner requires limited overhead in terms of number of iterations to convergence with respect to its single time-step equivalent. This allows significant room for parallel efficiency, so long as effective PinT methods are developed to solve the single-variable space-time systems appearing within our preconditioner, namely involving operators $F_{\mathbf{u}}$ and \tilde{C}_A . Compared to having to parallelise in time the solution of the whole multi-variable system (4.44), this task is now simplified, by virtue of our space-time block preconditioner.

4.5 Limitations of STBP for compressible MHD

In Sec. 3.5, we remarked on how, rather than being a time-parallelisation method in and of itself, the space-time block preconditioning procedure outlined here and in Chap. 3 represents more of a general strategy for simplifying the application of already-existing parallel-in-time algorithms to systems of PDEs. As such, this implies that its performance is ultimately reliant on that of the aforementioned algorithms.

When dealing with the incompressible kind of flow analysed in the last two chapters, this dependence is not particularly burdensome: in fact, we mostly need to solve space-time systems stemming from discretisations of single-variable *parabolic* equations, which most traditional PinT methods have already proven to handle quite effectively, as shown in the introductory discussion in Chap. 2. However, this ceases to be the case when dealing with applications to compressible flow: within that framework, one can expect the resulting target single-variable PDEs to be *hyperbolic* in nature instead, thus posing a much greater challenge to classical time-parallelisation schemes (see Sec. 2.1, and more specifically Sec. 2.1.1.1).

Therefore, rather than considering magnetohydrodynamics problems in the compressible regime, the work described in the following chapters is mostly directed at understanding the mechanism behind the failure of PinT algorithms when applied to hyperbolic problems (Chap. 5), and in trying to determine the validity of alternative approaches (Chap. 6). This also dictates a shift from the tone used in Chap. 3 and 4, with most of the findings illustrated next revolving around the investigation of the performance of PinT methods available in the literature, rather than the description of novel techniques.

Chapter 5

Multigrid Reduction in Time for nonlinear hyperbolic equations

Recovering an accurate numerical solution to hyperbolic equations *per se* is challenging in many ways [89, Chap. 1.4], with one of the biggest difficulties lying in being able to properly capture shocks present in the solution, while retaining the desired level of accuracy and ensuring stability of the scheme used. These objectives are often at conflict with each other: better accuracy can be achieved using high-order approximations, but these approximations might generate spurious oscillations in proximity of discontinuities in the solution, undermining the stability of the algorithm. To counteract this behaviour, a variety of schemes have been proposed in the literature, some of which are described in this chapter.

When classical approaches to time parallelisation, such as Parareal and MGRIT, come into play, this problem becomes even more pronounced, as solutions from different integrators need to be combined together. Even if both fine and coarse solvers are individually stable, often their combined use triggers instabilities that, in turn, result in loss of accuracy and poor convergence, negating most of the advantages from parallelisation. Examples of successful applications, however, still remain even in this case, as outlined in the discussion in Sec. 2.1.1.1.

The aim of this chapter is to investigate in detail the performance of MGRIT when applied to some nonlinear hyperbolic problems frequently employed as test-cases. In particular, we discuss the behaviour of the algorithm when used in conjunction with high-order discretisations, which are frequently used for the solution of conservation laws. Example applications of parallel-in-time methods to high-order schemes are, to our knowledge, somewhat rare in the literature. Noticeable exceptions are: publications involving the PFASST algorithm [43], which achieves high-order accuracy

via repeated SDC corrections, but is limited to this very specific class of integrators; the works conducted in [48, 49], where k -th order multi-step BDF methods are analysed; [109], which shares a similar set-up to ours, but only shows results for a specific configuration of the integrators used; finally, [33], which is possibly the work whose goal is the closest to the one of this chapter. There, the efficacy of MGRIT applied to hyperbolic equations is studied in detail, with the aim of identifying suitable coarse integrators to achieve fast convergence, though the project only covers the linear case. The purpose of our study lies in investigating the applicability of MGRIT to the acceleration of solutions to nonlinear systems of conservation laws which are relevant in real-world scenarios. It focuses on determining the principal factors that cause degradation in the convergence speed of the algorithm, and identifies future directions for improvements.

The work in this chapter has been conducted while on a research visit at the Memorial University of Newfoundland, under the supervision of Scott MacLachlan. The corresponding findings have been submitted to the *Electronic Transaction on Numerical Analysis* journal, and are currently under review. A pre-print version can be found in [30]. The remainder of this chapter is structured as follows. In Sec. 5.1, we introduce the test problems considered here and the discretisations used in our experiments. In Sec. 5.2, we expand on the description of the MGRIT algorithm provided in Sec. 2.2.3, and discuss the important choice of the coarse integrator for a given fine integrator. In Sec. 5.3, we present and discuss numerical experiments.

5.1 Model problems

As test cases for our experiments, we pick examples of *conservation laws* that, due to their physical relevance, are widely used in the analysis and testing of numerical methods for hyperbolic equations. For simplicity, we limit ourselves to problems defined on a one-dimensional spatial domain, and with periodic boundary conditions. These equations can be written in the following general form:

$$\left\{ \begin{array}{ll} \frac{\partial \mathbf{u}}{\partial t} + \frac{\partial \mathbf{f}(\mathbf{u})}{\partial x} = \mathbf{0} & (x, t) \in [0, L] \times (0, T] \\ \mathbf{u}(x, 0) = \bar{\mathbf{u}}^0(x) & x \in [0, L] \\ \mathbf{u}(0, t) = \mathbf{u}(L, t) & t \in (0, T] \end{array} \right. , \quad (5.1)$$

where the variable $\mathbf{u}(x, t) \in \mathbb{R}^D$ contains the vector of conserved variables in the system, or *state*, with associated *flux* $\mathbf{f}(\mathbf{u})$. The parameters L and T define the size

of the spatial and temporal domains, respectively, while $\bar{\mathbf{u}}^0(x)$ identifies the initial condition. We consider three different systems, ubiquitous in the literature.

Problem 1 The *Burgers' equation*, considered here in its inviscid formulation, which is possibly the simplest example of a scalar conservation law including nonlinear effects. It is defined by

$$\mathbf{u} = u, \quad \mathbf{f}(\mathbf{u}) = \frac{u^2}{2}, \quad (5.2)$$

with u representing the flow velocity.

Problem 2 The *shallow-water equations*, also considered without viscosity, which compose a system describing a fluid flow in the regime where the vertical length scale is negligible with respect to the horizontal one. The associated state and flux are, respectively

$$\mathbf{u} = \begin{bmatrix} h \\ hu \end{bmatrix}, \quad \mathbf{f}(\mathbf{u}) = \begin{bmatrix} hu \\ hu^2 + \frac{1}{2}gh^2 \end{bmatrix}, \quad (5.3)$$

with g being the gravitational constant, u the velocity, and h the height of the fluid column.

Problem 3 The *Euler equations*, another system governing compressible fluid flow, with

$$\mathbf{u} = \begin{bmatrix} \rho \\ \rho u \\ E \end{bmatrix}, \quad \mathbf{f}(\mathbf{u}) = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ (E + p)u \end{bmatrix} \quad (5.4)$$

where ρ is the flow density, u is once again its velocity, E its total internal energy, and p its pressure. In our case, this system is closed by considering an ideal monoatomic gas, which gives the following relationship between pressure and energy:

$$p = (\gamma - 1) \left(E - \frac{1}{2}\rho u^2 \right), \quad \text{with } \gamma = \frac{5}{3}. \quad (5.5)$$

These three problems are presented in detail in [89, Chap. 3.2, Chap. 5.4, and Chap. 5.1], respectively.

In the following, we provide a description of the numerical schemes employed in our experiments in order to recover their approximate solution.

5.1.1 Space discretisation

To simplify our notation, in this section we consider the state of the conservation law as being a scalar, unless otherwise specified, although the treatment described here can be seamlessly extended to the systems listed above.

Equation (5.1) is discretised via a method of lines approach, using finite volumes in space. The spatial domain $[0, L]$ is subdivided into N_x cells of uniform length $\Delta x = L/N_x$. The semi-discretised unknown is approximated by a vector $\mathbf{u}(t) \in \mathbb{R}^{N_x} \approx u(x, t)$, whose i -th component represents the cell-average

$$u_i(t) \approx \frac{1}{\Delta x} \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} u(x, t) dx, \quad i = 0, \dots, N_x - 1. \quad (5.6)$$

Here, $x_{i\pm\frac{1}{2}}$ identifies the coordinate of the right (respectively, left) interface of the i -th cell,

$$x_{i\pm\frac{1}{2}} = \left(i + \frac{1 \pm 1}{2} \right) \Delta x, \quad \begin{array}{c} x_{N_x-\frac{1}{2}} = L \\ \text{-----} | \overset{i = N_x - 1}{\text{|||}} | \overset{i = 0}{\text{---}} | \text{-----} \\ x_{-\frac{1}{2}} = 0 \qquad x_{\frac{1}{2}} = \Delta x \end{array} \quad (5.7)$$

where the right-interface of cell $N_x - 1$, (on coordinate $x_{N_x-\frac{1}{2}} = L$), logically coincides with the left-interface of cell 0, (on coordinate $x_{-\frac{1}{2}} = 0$), due to the periodic boundary conditions being applied. Integrating equation (5.1) on cell i , we obtain

$$\frac{\partial u_i(t)}{\partial t} = - \frac{f_{i+\frac{1}{2}}(t) - f_{i-\frac{1}{2}}(t)}{\Delta x} =: [\mathcal{A}(\mathbf{u}(t))]_i, \quad (5.8)$$

where $f_{i+\frac{1}{2}}(t)$ defines an approximation of the flux at the cell interface at $x_{i+\frac{1}{2}}$:

$$f_{i+\frac{1}{2}}(t) \approx f \left(u \left(x_{i+\frac{1}{2}}, t \right) \right). \quad (5.9)$$

Typically, this approximation must be built starting from values of the numerical solution on the surrounding cells. There is, by far, no unique way to achieve this and, indeed, several schemes have been proposed in the literature, based on different methods for reconstructing the values in (5.9). Our experiments are based primarily on the WENO scheme, which is briefly introduced in Sec. 5.1.1.1. In Sec. 5.1.1.2 and Sec. 5.1.1.3, we illustrate the two procedures employed to recover the fluxes in (5.9) using WENO.

5.1.1.1 WENO reconstruction

The *Weighted Essentially Non-Oscillatory* scheme (or WENO) was originally developed by Liu, Chan and Osher in [94]. In this section, we aim at providing a short description of the main idea behind the method, but we refer to [75, Chap. 11.3] for a more complete treatment.

The aim of the WENO procedure is to provide a stable, high-order approximation of a certain value of interest (usually the state itself, or directly the flux) at the cell interface. We represent these generic recovered quantities as

$$q_{i+\frac{1}{2}}^{\pm} \approx q\left(x_{i+\frac{1}{2}}^{\pm}\right). \quad (5.10)$$

Since discontinuities might arise in the solution of hyperbolic equations, the approximations on the left ($-$) and right ($+$) of an interface can differ in general, hence the different superscripts in (5.10). These quantities are approximated via polynomial reconstruction, starting from the cell-values q_i in a stencil containing the interface. However, as pointed out in the introduction to this chapter, solutions to (5.1) can develop shocks: if the stencil used for the polynomial reconstruction happens to contain a sharp discontinuity, then the reconstructed polynomial can show oscillations due to Gibbs phenomenon, and the approximation recovered might be of poor quality. In order to counteract this, the WENO method collects approximations built using a number of different stencils, and opportunely combines them via a convex linear combination, with the weight of each approximation considered depending on an estimate of the smoothness of the particular reconstruction.

In more detail, in order to recover an approximation of the value to the left of the interface, $q_{i+\frac{1}{2}}^-$, using a polynomial $\tilde{q}_i^r(x)$ of degree k , one can choose between $k + 1$ different stencils each containing $k + 1$ cells (see Fig. 5.1). Each of these




Figure 5.1: Choice of stencils S_i^r for the polynomial reconstructions of degree $k = 2$ of the value left of the interface, $q_{i+\frac{1}{2}}^-$, starting from the cell values q_{i-k}, \dots, q_{i+k} (flat blue lines). Different stencils produce different approximating polynomials $\tilde{q}_i^r(x)$ (dotted blue lines) which, in turn, give different values at the interface (blue dots)

approximations can be expressed as a linear combination of the cell-averages,

$$q_{i+\frac{1}{2}}^{-,r} = \tilde{q}_i^r \left(x_{i+\frac{1}{2}} \right) = \sum_{j \in S_i^r} c_{r,j-(i+r-k)} q_j, \quad (5.11)$$

where here (and in the remainder of this section) r ranges from 0 to k ; the stencil S_i^r contains the cells with indices $S_i^r := \{i+r-k, \dots, i+r\}$. The quality of the r -th reconstruction $q_{i+\frac{1}{2}}^{-,r}$ on cell i is measured via a *smoothness indicator*, β_i^r , based on the (scaled) L^2 norm of the derivatives of the reconstructed polynomial $\tilde{q}_i^r(x)$ on S_i^r . This can be expressed as a bilinear form,

$$\beta_i^r := \sum_{l=1}^k \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \Delta x^{2l-1} \left(\frac{\partial^l \tilde{q}_i^r(x)}{\partial x^l} \right)^2 dx = (\mathbf{q}_i^r)^T B_r \mathbf{q}_i^r, \quad (5.12)$$

represented by $B_r \in \mathbb{R}^{k \times k}$, with $\mathbf{q}_i^r \in \mathbb{R}^k$ being the vector collecting the values of q_j on the stencil considered, $j \in S_i^r$. These smoothness indicators contribute to the definition of the weights ω_i^r as follows:

$$\omega_i^r := \frac{\alpha_i^r}{\sum_{r=0}^k \alpha_i^r}, \quad \text{and} \quad \alpha_i^r := \frac{d_r}{(\varepsilon + \beta_i^r)^2}, \quad (5.13)$$

where ε is a small parameter to prevent division by 0, (usually $\varepsilon = 10^{-6}$). The coefficients d_r are defined in such a way that, in the smooth case,

$$q_{i+\frac{1}{2}} = \sum_{r=0}^k d_r q_{i+\frac{1}{2}}^r = q \left(x_{i+\frac{1}{2}} \right) + \mathcal{O}(\Delta x^s) \quad (5.14)$$

holds with $s = 2k + 1$. This way, the choice (5.13) of the weights ω_i^r ensures that the recovered reconstruction is a high-order approximation of the underlying quantity. Finally, the actual reconstruction is computed as a convex combination of all the different approximations in (5.11),

$$q_{i+\frac{1}{2}}^- = \sum_{r=0}^k \omega_i^r q_{i+\frac{1}{2}}^{-,r}. \quad (5.15)$$

The specific values of the parameters $c_{r,j}$, d_r , and B_r described above depend directly on the degree k of the polynomials employed. For example, for $k = 2$, we have:

$$[c_{r,j}] := \frac{1}{6} \begin{bmatrix} 2 & 5 & -1 \\ -1 & 5 & 2 \\ 2 & -7 & 11 \end{bmatrix}, \quad [d_r] := \frac{1}{10} \begin{bmatrix} 3 \\ 6 \\ 1 \end{bmatrix}, \quad \text{and} \\ [B_0|B_1|B_2] := \frac{1}{6} \left[\begin{array}{ccc|ccc|ccc} 20 & -31 & 11 & 8 & -13 & 5 & 8 & -19 & 11 \\ -31 & 50 & -19 & -13 & 26 & -13 & -19 & 50 & -31 \\ 11 & -19 & 8 & 5 & -13 & 8 & 11 & -31 & 20 \end{array} \right]. \quad (5.16)$$

An analogous procedure is used to recover the value to the right of the interface, $q_{i+\frac{1}{2}}^+$. We also point out that, if piecewise constant polynomials are chosen ($k = 0$), then the reconstruction becomes trivial: $q_{i+\frac{1}{2}}^- = q_i$, and $q_{i+\frac{1}{2}}^+ = q_{i+1}$.

When dealing with a system of conservation laws, additional care is necessary to recover a proper reconstruction of the state. A naïve extension from the scalar case would consider reconstructing each component of the state independently, but this has been shown to cause some unphysical oscillations in the solution [116]. To limit this effect, a more stable approach rather consists of independently reconstructing the *characteristic* variables of the system, since these are more readily associated with the information carried by the characteristics [75, Chap. 11.3.4]. This comes at an additional cost, as it requires a local decomposition of the state in each of the cells considered for reconstruction on a given reference state: this procedure is described in [136, Procedure 2.8], where an averaged state at the interface is taken as a reference. Even though we broadly follow these guidelines in our implementation, we decide for simplicity to decompose the values $\mathbf{u}_{i-k}, \dots, \mathbf{u}_{i+k}$ using the central cell value \mathbf{u}_i as a reference state, in order to recover $\mathbf{u}_{i+\frac{1}{2}}^-$ and $\mathbf{u}_{i-\frac{1}{2}}^+$.

With the WENO procedure available to reconstruct the states to the left and right of each interface, we proceed to approximate the numerical flux (5.9) in two different ways. These are described next.

5.1.1.2 Lax-Friedrichs flux

We first consider the Lax-Friedrichs definition for the numerical flux [136, Chap. 2.3.1]. For a system of conservation laws, this is given by

$$\mathbf{f}_{i+\frac{1}{2}}^{LF} := \frac{1}{2} \left(\mathbf{f} \left(\mathbf{u}_{i+\frac{1}{2}}^+ \right) + \mathbf{f} \left(\mathbf{u}_{i+\frac{1}{2}}^- \right) \right) + \frac{1}{2} \alpha_{LF} \left(\mathbf{u}_{i+\frac{1}{2}}^+ - \mathbf{u}_{i+\frac{1}{2}}^- \right). \quad (5.17)$$

This flux is one of the simplest to prescribe; however, as a downside, it typically produces smeared out numerical solutions, since it adds artificial diffusion to the system. The amount of numerical diffusion added depends directly on the parameter α_{LF} : for our implementation, we choose the (somewhat loose) value of

$$\alpha_{LF} := \max_{i,k} \left\{ \left| \lambda_k \left(\mathbf{u}_{i+\frac{1}{2}}^\pm \right) \right| \right\}, \quad \begin{array}{l} i = 0, \dots, N_x - 1 \\ k = 0, \dots, D - 1 \end{array}, \quad (5.18)$$

where the λ_k 's are the D eigenvalues (if the system is composed of D equations) of the Jacobian of the flux:

$$J_{\mathbf{f}}(\mathbf{u}) \mathbf{r}_k = \lambda_k \mathbf{r}_k, \quad \text{with} \quad [J_{\mathbf{f}}(\mathbf{u})]_{m,n} := \frac{\partial [\mathbf{f}(\mathbf{u})]_m}{\partial [\mathbf{u}]_n}, \quad (5.19)$$

where the \mathbf{r}_k 's represent the corresponding eigenvectors.

5.1.1.3 Roe flux

The second scheme investigated in our experiments uses Roe's approximate Riemann solver [89, Chap. 14.2] in order to recover the numerical flux. For each interface, this solver targets a linearisation of the Riemann problem defined by (5.1) and the left and right values computed using the WENO procedure. This problem is written as:

$$\mathbf{u}_t + J_f(\hat{\mathbf{u}}) \mathbf{u}_x = 0, \quad (5.20)$$

where $\hat{\mathbf{u}}$ is an opportunely defined *Roe-averaged* state. All relevant variables in this section are to be intended as evaluated at an interface: we drop the subscripts $i + 1/2$ for ease of notation. The solution to the linear Riemann problem (5.20), as well as the associated flux at the interface, are both easily expressed in terms of the eigenvalues and eigenvectors of the Jacobian. These are given by

$$\begin{aligned} \lambda_0^S &= u - c_S & \text{and} & & \mathbf{r}_0^S &= \begin{bmatrix} 1 & \lambda_0^S \end{bmatrix}^T \\ \lambda_1^S &= u + c_S & & & \mathbf{r}_1^S &= \begin{bmatrix} 1 & \lambda_1^S \end{bmatrix}^T \end{aligned} \quad (5.21)$$

for the shallow-water system, with *speed of sound* $c_S := \sqrt{gh}$, and by

$$\begin{aligned} \lambda_0^E &= u - c_E & \mathbf{r}_0^E &= \begin{bmatrix} 1 & \lambda_0^E & H - uc_E \end{bmatrix}^T \\ \lambda_1^E &= u & \text{and} & & \mathbf{r}_1^E &= \begin{bmatrix} 1 & \lambda_1^E & u^2/2 \end{bmatrix}^T \\ \lambda_2^E &= u + c_E & & & \mathbf{r}_2^E &= \begin{bmatrix} 1 & \lambda_2^E & H + uc_E \end{bmatrix}^T \end{aligned} \quad (5.22)$$

for the Euler equations, with speed of sound $c_E := \sqrt{(\gamma - 1)(H - u^2/2)}$. Here, $H := (E + p)/\rho$ is the *specific enthalpy*. For Burgers' equation, the only eigenvalue is given by $\lambda_0^B = u$ itself. Since we are only interested in the eigenvalues and eigenvectors of the Jacobian $J_f(\hat{\mathbf{u}})$ in (5.20), (namely $\hat{\lambda}_k$ and $\hat{\mathbf{r}}_k$), it suffices to define the following Roe-averaged quantities: an average velocity for the Burgers' equation, given by

$$\hat{u} := \frac{u^+ + u^-}{2}; \quad (5.23)$$

average height and velocity for the Shallow-water equation, prescribed as

$$\hat{h} := \frac{h^+ + h^-}{2} \quad \text{and} \quad \hat{u} := \frac{u^+ \sqrt{h^+} + u^- \sqrt{h^-}}{\sqrt{h^+} + \sqrt{h^-}}; \quad (5.24)$$

finally, average velocity and specific enthalpy for the Euler equations:

$$\hat{u} := \frac{u^+ \sqrt{\rho^+} + u^- \sqrt{\rho^-}}{\sqrt{\rho^+} + \sqrt{\rho^-}} \quad \text{and} \quad \hat{H} := \frac{(E^+ + p^+) \sqrt{\rho^+} + (E^- + p^-) \sqrt{\rho^-}}{\sqrt{\rho^+} + \sqrt{\rho^-}}. \quad (5.25)$$

Since the target problem (5.20) is linear, this procedure is known to provide a non-entropic weak solution when transonic rarefaction waves arise [89, Chap. 14.2.2]. To

counteract this, we also apply the *entropy fix* proposed by Harten and Hyman in [72]. In the general case of a system of D conservation laws, this gives rise to the following formula for the numerical flux at the interface:

$$\mathbf{f}^R := \frac{1}{2} (\mathbf{f}(\mathbf{u}^-) + \mathbf{f}(\mathbf{u}^+)) - \sum_{k=0}^{D-1} q_H(\hat{\lambda}_k) \alpha_k \hat{\mathbf{r}}_k. \quad (5.26)$$

Here, α_k is the *shock strength*, that is the jump in the k -th characteristic variable between the states left and right of the interface: $\alpha_k := (\mathbf{u}^+ - \mathbf{u}^-) \cdot \hat{\mathbf{r}}_k$. Also, $q_H(\lambda_k)$ defines the entropy fix:

$$q_H(\hat{\lambda}_k) := \begin{cases} \frac{1}{2} \left(\frac{\hat{\lambda}_k}{\delta_k} + \delta_k \right) & \text{if } |\hat{\lambda}_k| < \delta_k \\ |\hat{\lambda}_k| & \text{else} \end{cases}, \quad (5.27)$$

where $\delta_k := \max \{0, \hat{\lambda}_k - \lambda_k^-, \lambda_k^+ - \hat{\lambda}_k\}$, and λ_k^\pm are the eigenvalues of $J_{\mathbf{f}}(\mathbf{u}^\pm)$.

Substituting (5.17) or (5.26) into (5.8) identifies the right-hand side $\mathcal{A}(\mathbf{u})$ of our semi-discretised equation. The details regarding the discretisation in time, as well as the definition of the time-steppers employed, are presented next.

5.1.2 Time discretisation

The time domain is also discretised using a uniform grid of N_t nodes, with time step $\Delta t = T/N_t$. The unknowns at each instant $t_n = n\Delta t$, $n = 0, \dots, N_t - 1$ are approximated by a vector denoted as $\mathbf{u}^n \approx \mathbf{u}(t_n)$.

Since we aim at employing high-order spatial reconstructions, it is sensible to request that our temporal discretisation matches this accuracy. The schemes chosen belong to the family of *Strong Stability-Preserving Runge Kutta* methods (SSPRK). As the name suggests, these have the remarkable property of being able to preserve strong stability and, in particular, to be *Total Variation Diminishing*, even while achieving high-order of accuracy; for this reason, they are often employed in conjunction with high-order spatial discretisation of hyperbolic equations. We refer to [67] for a complete review of these schemes, and we only report here the definition of the methods used in our experiments, which vary in the order of accuracy d recovered. They are: a third-order scheme (SSPRK3), whose stepping procedure applied to (5.8) is given by

$$\begin{aligned} \mathbf{u}^{n,(1)} &= \mathbf{u}^n + \Delta t \mathcal{A}(\mathbf{u}^n) \\ \mathbf{u}^{n,(2)} &= \frac{3}{4} \mathbf{u}^n + \frac{1}{4} \mathbf{u}^{n,(1)} + \frac{1}{4} \Delta t \mathcal{A}(\mathbf{u}^{n,(1)}) \\ \mathbf{u}^{n+1} &= \frac{1}{3} \mathbf{u}^n + \frac{2}{3} \mathbf{u}^{n,(2)} + \frac{2}{3} \Delta t \mathcal{A}(\mathbf{u}^{n,(2)}); \end{aligned} \quad (5.28)$$

a second-order scheme (SSPRK2), defined by

$$\begin{aligned}\mathbf{u}^{n,(1)} &= \mathbf{u}^n + \Delta t \mathcal{A}(\mathbf{u}^n) \\ \mathbf{u}^{n+1} &= \frac{1}{2}\mathbf{u}^n + \frac{1}{2}\mathbf{u}^{n,(1)} + \frac{1}{2}\Delta t \mathcal{A}(\mathbf{u}^{n,(1)}); \end{aligned} \quad (5.29)$$

and, finally, the first-order scheme, which reduces to the well-known forward Euler (FE) method,

$$\mathbf{u}^{n+1} = \mathbf{u}^n + \Delta t \mathcal{A}(\mathbf{u}^n). \quad (5.30)$$

5.2 Time parallelisation with MGRIT

In this section we expand on the introduction to MGRIT given in Sec. 2.2.3, so to provide a more in-depth description of the various components of the algorithm used in our experiments. To aid in our exposition, we follow much of [47] and consider the space-time discretisation of a constant-coefficient, linear, homogeneous system using an explicit single-step time integrator. Similarly to (2.26), this system takes the following block Toeplitz, block bi-diagonal form

$$A\mathbf{u} = \mathbf{g} \iff \begin{bmatrix} I & & & & \\ -\mathcal{F} & I & & & \\ & \ddots & \ddots & & \\ & & & -\mathcal{F} & I \end{bmatrix} \begin{bmatrix} \mathbf{u}^0 \\ \mathbf{u}^1 \\ \vdots \\ \mathbf{u}^{N_t-1} \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{u}}^0 \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix}, \quad (5.31)$$

where $\mathbf{u} \in \mathbb{R}^{N_x N_t}$ is the vector containing the values of the discrete solution at each node in the space-time grid, and $\bar{\mathbf{u}}^0$ contains the discretisation of the initial condition; since we do not consider forcing terms, the rest of \mathbf{g} is filled with zeros. The *fine integrator* \mathcal{F} represents the action of the time-stepper, so that, for each instant $n = 1, \dots, N_t - 1$, we have $\mathbf{u}^n = \mathcal{F}\mathbf{u}^{n-1}$. We proceed to splitting the temporal grid into fine and coarse nodes, and we subdivide the monolithic system (5.31) accordingly, as done in (2.31), and reported here for ease of reference:

$$A = \begin{bmatrix} A_{FF} & A_{FC} \\ A_{CF} & A_{CC} \end{bmatrix} = \begin{bmatrix} I & 0 \\ A_{CF}A_{FF}^{-1} & I \end{bmatrix} \begin{bmatrix} A_{FF} & 0 \\ 0 & S \end{bmatrix} \begin{bmatrix} I & A_{FF}^{-1}A_{FC} \\ 0 & I \end{bmatrix}. \quad (5.32)$$

For simplicity, we pick the coarse nodes to be spaced every m nodes, with m being the *coarsening factor* (see Fig. 5.2), so that the coarse nodes subdivide the time domain into $(N_t - 1)/m$ different *time chunks*, each of size $m\Delta t$.

It can be shown that in our case the A_{FF} sub-matrix of (5.32) presents a block diagonal structure,

$$A_{FF} = \begin{bmatrix} B_{\mathcal{F}} & & \\ & \ddots & \\ & & B_{\mathcal{F}} \end{bmatrix}, \quad \text{with} \quad B_{\mathcal{F}} = \begin{bmatrix} I & & & \\ -\mathcal{F} & I & & \\ & \ddots & \ddots & \\ & & -\mathcal{F} & I \end{bmatrix}, \quad (5.33)$$

which implies that systems involving it can be readily solved in parallel. In (5.33), we note that A_{FF} is overall an $N_x(N_t-1)(m-1)/m \times N_x(N_t-1)(m-1)/m$ matrix, written above as an $(N_t-1)/m \times (N_t-1)/m$ block diagonal system, with blocks $B_{\mathcal{F}}$ of size $(m-1)N_x \times (m-1)N_x$. Each solve with $B_{\mathcal{F}}$ involves “stepping” forward an approximate solution on the spatial mesh through $m-2$ applications of \mathcal{F} . The challenge lies in finding the solution to the system associated with the Schur complement, which in this case assumes the following shape:

$$S = A_{CC} - A_{CF}A_{FF}^{-1}A_{FC} = \begin{bmatrix} I & & & \\ -\mathcal{F}^m & I & & \\ & \ddots & \ddots & \\ & & -\mathcal{F}^m & I \end{bmatrix}. \quad (5.34)$$

Attempting to solve this directly is equivalent to applying forward block substitution to (5.31), (*i.e.*, time-stepping over all temporal nodes) and would thus nullify any advantage gained from parallelisation. Instead, in the scope of the MGRIT algorithm, we resort to solving a modified system, obtained by substituting another operator $\mathcal{G} \approx \mathcal{F}^m$ in (5.34), giving

$$S_{\mathcal{G}} = \begin{bmatrix} I & & & \\ -\mathcal{G} & I & & \\ & \ddots & \ddots & \\ & & -\mathcal{G} & I \end{bmatrix} \approx S. \quad (5.35)$$

The structure of this system is equivalent to that of (5.31), so that \mathcal{G} can be interpreted as yet another time-stepping routine, which acts only on the coarse nodes: we call this operator the *coarse integrator*. Notice that the splitting into fine and coarse




Figure 5.2: Sketch of the partitioning of the time discretisation: white dots represent coarse nodes, black dots are fine nodes. The last time chunk is also highlighted.

nodes can be applied in a recursive fashion, further extracting a hierarchy of N_l grids, together with their corresponding integrators $\mathcal{F}_{(l)}$. The hope is that, by opportunely alternating between solving for the fine nodes (inverting A_{FF}) and for the coarse nodes (inverting S_G) and iterating, we can quickly converge to a suitable approximation of the solution of the original system (5.31).

In more detail, at each iteration of the MGRIT algorithm we need to apply the operations described next

Relaxation *Update the solution at the fine nodes of the current level, given a guess at the coarse nodes.* This involves solving a system in the form

$$A_{FF}^{(l)} \mathbf{u}_{(l)}^F = \mathbf{g}_{(l)}^F - A_{FC}^{(l)} \mathbf{u}_{(l)}^C, \quad (5.36)$$

where $A_{FF}^{(l)}$ has the same structure as in (5.33), but whose subdiagonal blocks contain $\mathcal{F}_{(l)}$, while $\mathbf{u}_{(l)}^F$ and $\mathbf{u}_{(l)}^C$, respectively, represent the grouping of the unknowns at the fine and coarse nodes of the current level, l . The solution is hence updated by time-stepping using $\mathcal{F}_{(l)}$, starting from the given values at the coarse nodes. Here lies the parallel part of the algorithm, since the time-stepping procedure can be carried out independently on each chunk. In our work, we consider the two types of relaxation briefly introduced in Sec. 2.2.3: *F-relaxation*, where the time-stepping covers a single time chunk, updating the fine nodes within it, and *FCF-relaxation*, where the time-stepping carries on over the following coarse node (performing a *C-relaxation*) and continues on the following time chunk (adding another F-relaxation). The latter is an *overlapping* form of relaxation that requires more work per level of the hierarchy but can be implemented with the same parallel efficiency [47].

Restriction *Transfer information from the nodes of the current level l to the nodes at the coarser level $l + 1$.* Simple injection onto the coarse nodes is chosen as the restriction operator:

$$R_{(l+1)}^{(l)} \mathbf{u}_{(l)} = \mathbf{u}_{(l)}^C. \quad (5.37)$$

Since we are dealing with nonlinear equations, we also implement the *Full Approximation Scheme* (FAS) [150, Chap. 5.3.4]. This modifies the right-hand side of the system at the coarse level by adding a correction term:

$$S_{\mathcal{F}_{(l+1)}}(\mathbf{u}_{(l+1)}) = R_{(l+1)}^{(l)} \left(\mathbf{g}_{(l)} - S_{\mathcal{F}_{(l)}}(\mathbf{u}_{(l)}) \right) + \underbrace{S_{\mathcal{F}_{(l+1)}} \left(R_{(l+1)}^{(l)} \mathbf{u}_{(l)} \right)}_{\text{FAS correction}} = \mathbf{g}_{(l+1)}, \quad (5.38)$$

where $S_{\mathcal{F}_{(l+1)}}$ has the same structure as in (5.35), but with \mathcal{F}_{l+1} on the subdiagonal. We point out that the same restriction operator (5.37) is applied to both the residual and the state. Given the particular structure of the operators involved, formula (5.38) simplifies, so that the right-hand side of the coarse-level system, $\mathbf{g}_{(l+1)}$, reduces to

$$\mathbf{g}_{(l+1)}^i = \mathbf{g}_{(l)}^{mi} + \mathcal{F}_{(l)} \left(\mathbf{u}_{(l)}^{m^{i-1}} \right) - \mathcal{F}_{(l+1)} \left(\mathbf{u}_{(l)}^{m^{(i-1)}} \right), \quad (5.39)$$

for each temporal node $i = 1, \dots, (N_t - 1)/m^{l+1} + 1$ of the coarser level. We also need to provide an initial guess for the solution at the coarse level $\mathbf{u}_{(l+1)}$: this is usually chosen to be $\mathbf{u}_{(l)}^C$, however we have found that a better alternative is given by

$$\mathbf{u}_{(l+1)}^i = \mathcal{F}_{(l)} \left(\mathbf{u}_{(l)}^{m^{i-1}} \right) + \mathbf{g}_{(l)}^{mi}, \quad (5.40)$$

that is, by performing the last integration step for each chunk and injecting the resulting value. The rationale behind this choice is given by the fact that, if restriction follows relaxation, with (5.40) we are injecting into the coarse level information from the relaxation procedure as well, which should provide a more refined guess over $\mathbf{u}_{(l)}^C$. This operation comes at virtually no additional cost (since the quantity in (5.40) is already computed as part of (5.39)), and indeed in early experiments it was seen to improve convergence: in Fig. 5.5, we give an example of the effectiveness of (5.40) over simple injection. Notice that this procedure is *not* equivalent to performing an additional C-relaxation before injection, since the FAS correction (5.38) is still based on the non-updated values at the coarse nodes $\mathbf{u}_{(l)}^{m^{(i-1)}}$; notice moreover that this approach *does* retain the fixed-point property of the FAS algorithm, since (5.37) and (5.40) coincide when $\mathbf{u}_{(l)}$ is the exact solution on level l .

Coarse grid correction *Update solution at the coarsest level.* This involves solving the system

$$S_{\mathcal{F}_{(N_l-1)}} \mathbf{u}_{(N_l-1)} = \mathbf{g}_{(N_l-1)}. \quad (5.41)$$

This procedure is sequential but, by choosing a cheap coarse integrator, the overhead to the algorithm can be limited, and parallel efficiency can still be gained.

Interpolation *Transfer information from the nodes at the coarser level $l + 1$ to those at the current level l .* We choose *ideal interpolation*, which in our case consists of two steps: an injection from coarse to fine grid, followed by an F-relaxation on the

fine level, starting from the freshly updated values at the coarse nodes. Overall, we have the following definition for the interpolation operator:

$$\begin{aligned} \left[I_{(l)}^{(l+1)}(\mathbf{u}_{(l+1)}) \right]^i &= \bar{I}_{\lfloor \frac{i}{m} \rfloor, i \% m}^{(l)} \left(\mathbf{u}_{(l+1)}^{\lfloor \frac{i}{m} \rfloor} \right), \quad \text{with} \\ \bar{I}_{k,r}^{(l)}(\mathbf{u}) &= \begin{cases} \mathbf{u} & \text{if } r = 0, \\ \mathcal{F}_{(l)} \left(\bar{I}_{k,r-1}^{(l)}(\mathbf{u}) \right) + \mathbf{g}_{(l)}^{k+r} & \text{else,} \end{cases} \end{aligned} \quad (5.42)$$

where $\lfloor * \rfloor$ and $\%$ denote the floor and modulo operators, respectively. Notice that, if the interpolation step is followed by a relaxation step, then the F-relaxation within interpolation becomes redundant, since the results would get overwritten by the relaxation step. Also notice that here we consider only *pre*-relaxation within the multigrid cycles, without *post*-relaxation. This goes against what is commonly done when applying multigrid to symmetric operators, as it would break the self-adjointness of the multigrid procedure; it is however standard setup in the MGRIT framework, where the target operator is a space-time matrix: as this is non-symmetric to begin with, this choice is not particularly concerning: see the discussion in Sec. 2.2.4 and [153].

The order in which the algorithm moves between discretisation levels defines the sequence in which the operations above are prescribed, and identifies the type of *cycle* used in MGRIT. We experiment both with a *V-cycle* and an *F-cycle* [150, Chap. 3]. In Alg. 4, the pseudo-code outlining the (recursive) definition of a V-cycle with F-relaxation is provided; a sketch of the movement along the various levels for both V- and F-cycle is given in Fig. 5.3.

Algorithm 4: MGRIT, V-cycle iteration

```

1 if  $l < N_l - 1$  then
   | /* Parallel step */
2   Perform relaxation at current level and recover  $\mathbf{u}_{(l)}^F$  by solving (5.36);
3   Compute residual and FAS correction, and restrict to coarser level to
   | compute  $\mathbf{g}_{(l+1)}$  and  $\mathbf{u}_{(l+1)}$  following formulae (5.39) and (5.40);
   | /* Recursive step */
4   Invoke this function at coarser level and find  $\mathbf{u}_{(l+1)}$ ;
   | /* Parallel step */
5   Interpolate from coarser level to recover  $\mathbf{u}_{(l)}$  applying (5.42);
6 else
   | /* Serial step - base step of recursion */
7   Recover  $\mathbf{u}_{(N_l-1)}$  by solving system (5.41);

```




Figure 5.3: Sketch of the flow of the MGRIT algorithm if prescribing a V-cycle (left) or an F-cycle (right), on 5 levels. Restriction (solid lines) is applied when moving from a level to a coarser one (dashed lines: the size of the time-step at each level l is reported on the right). Relaxation (blue dots) is applied at each level, bar the bottom one, where the coarsest system is solved exactly (black dots). Interpolation (dotted lines) is applied when moving from a level to a finer one.

5.2.1 Choices for the coarse integrator

As hinted above, the key to the effective application of MGRIT lies in finding an adequate pair (or set) of fine and coarse integrators. In particular, the coarse integrator needs to do a “good enough” job of mimicking the action of the fine integrator, so that (5.35) represents a valid approximation to (5.34); at the same time, it should be cheap enough to compute, so that the solution to systems involving (5.35) can be promptly recovered. The necessity of overcoming this trade-off between cost and accuracy is common to other time-parallelisation algorithms (most notably, Parareal in Sec. 2.1.1), and a number of approaches have been proposed to address this. For MGRIT, the most straightforward is simple rediscritisation: the coarse integrator is none other than the fine integrator applied on a coarser time grid [46, 47]; an alternative consists in choosing integrators of varying levels of accuracy [109]. These solutions focus mostly on finding a coarse integrator which is cheap to compute, but it is unclear to what extent it matches the fine integrator. Indeed, in the hyperbolic framework, simple rediscritisation has been shown to fail in many cases (see Fig. 5.4), particularly if low-order time discretisations are employed: evidence of this is given in [33], as well as in the results of Fig. 5.6, at least for certain regimes. Therefore, when sticking to simple time-steppers, additional care is needed to ensure stability of MGRIT. To address this, we propose a novel approach which reverses the aforementioned principle: we start directly from the definition of the (explicit) fine integrator, and we progressively perform some simplifications which render it cheaper to apply,

without sacrificing “too much” accuracy. With this, we are able to “restore” MGRIT convergence for a low-order explicit time discretisation to be comparable to that of direct rediscratisation with high-order explicit time discretisations, although in neither case do we see CFL-robust convergence.

5.2.1.1 Fine integrator matching

To illustrate our approach, we consider our model equation (5.2). If a simple Lax-Friedrichs flux discretisation is applied in conjunction with a Forward Euler time discretisation, the resulting time-stepping formula reads

$$\begin{aligned} u_i^{n+1} &= \underbrace{\frac{u_{i+1}^n + u_{i-1}^n}{2}}_{Eu_i^n} - \Delta t \underbrace{\frac{f(u_{i+1}^n) - f(u_{i-1}^n)}{2\Delta x}}_{Df_i^n} \\ &= Eu_i^n - \Delta t Df_i^n, \end{aligned} \quad (5.43)$$

with $f_i^n = f(u_i^n)$, and where we introduced the *central difference* operator, $D((*)_i) := ((*)_{i+1} - (*))_{i-1})/2\Delta x$, and the *average* operator $E((*)_i) := ((*)_{i+1} + (*))_{i-1})/2$. If we were to simply rediscratisation at the coarse level and apply the same scheme to a coarser grid with time step $2\Delta t$, the resulting formula would similarly read

$$u_i^{n+1} = Eu_i^n - 2\Delta t Df_i^n. \quad (5.44)$$

However, we would like the result from the coarse integrator to be close to that of the fine integrator. The latter is given, in our case, by applying (5.43) twice (computing the nonlinear analogue of \mathcal{F}^2 , as in (5.34) with $m = 2$), which results in

$$\begin{aligned} u_i^{n+2} &= E(Eu_i^n - \Delta t Df_i^n) - \Delta t Df(Eu_i^n - \Delta t Df_i^n) \\ &\approx E^2 u_i^n - \Delta t (EDf_i^n + Df(Eu_i^n)) \\ &\quad + \Delta t^2 D(f'(Eu_i^n)Df_i^n) \\ &\quad - \Delta t^3 D(f''(Eu_i^n)(Df_i^n)^2). \end{aligned} \quad (5.45)$$

Here, we Taylor-expanded the flux in the second term around $E(u_i^n)$, exploiting the fact that Δt is a small parameter. In the case of (5.2), this formula is exact, since the flux f is a polynomial of degree 2. By directly comparing (5.45) and (5.44), we see that the formula for the coarse integrator is indeed much cheaper to compute, but it also significantly differs from that of the compounded fine integrator. A way to improve the accuracy of the coarse integrator without excessively increasing its cost consists in correcting (5.44) so that it matches (5.45) up to a certain order of Δt . This gives rise to the following integrators:

$$u_i^{n+1} = [\mathcal{G}^{(0)}(\mathbf{u}_i^n)]_i := E^2 u_i^n - 2\Delta t Df_i^n, \quad (5.46)$$

if we aim for a zeroth-order match, leaving the rest untouched, or

$$u_i^{n+1} = [\mathcal{G}^{(1)}(\mathbf{u}_i^n)]_i := E^2 u_i^n - \Delta t (EDf_i^n + Df(Eu_i^n)), \quad (5.47)$$

for a first-order match. Formula (5.45) refers to a coarsening factor $m = 2$, but this can be easily extended to any value of $m > 1$, as follows:

$$u_i^{n+m} = E^m u_i^n - \Delta t \sum_{i=1}^m E^{i-1} Df(E^{m-i} u_i^n) + \mathcal{O}(\Delta t^2), \quad (5.48)$$

and the corresponding (5.46) and (5.47) can be modified accordingly.

Regrettably, this approach presents a number of downsides. First of all, it carries an increased computational cost with respect to rediscritisation, as the stencil of the operators involved grows linearly with m . For zeroth-order matching, this simply translates into a larger number of vector additions, so that the cost is still contained, provided we refrain from using aggressive coarsening strategies. For first-order matching, though, the number of flux evaluations per time-step increases as well, which unfortunately makes the method far less appealing. Secondly, this approach is limited in its application, as it requires an explicit formula for the repeated application of the fine solver (5.48): this needs to be both readily computable, and prone to simplifications without having to resort to evaluating intermediate states. All these requirements drastically limit the type of problems that can be addressed, as well as the possible schemes we can employ. In fact, we have to resort to a low-order, highly dissipative method: notice that (5.43) corresponds to the Lax-Friedrichs flux (5.17) with a choice of an even larger parameter $\alpha_{LF} = \Delta x / \Delta t$. For these reasons, we still recommend higher-order schemes, which in our tests behave reasonably well in most regimes. Nonetheless, the results in Sec. 5.3 show the validity of this method, and highlight the importance of accurately mimicking the action of the fine solver at the coarser levels to achieve fast convergence.

5.3 Numerical results

In this section, we discuss how the convergence behaviour of MGRIT is impacted by factors such as the parameters in the multigrid algorithm, and the type of integrators chosen. The code used for the simulations is publicly available at [28].

5.3.1 Fine integrator matching

As a first test, we check the effectiveness of the time-stepper proposed in Sec. 5.2.1.1 applied to Burgers' equation: the convergence results for a number of runs are reported in Fig. 5.4. We choose two different initial conditions: first, the simple sinusoidal wave $\bar{u}_S^0 = \sin(2\pi x/L)$, for which the solution develops a *stationary shock* positioned at $x_S = L/2$, at the *breaking time* $t_S = -1/\min_x(\bar{u}_S^0(x))' = L/(2\pi)$; second, its translation $\bar{u}_M^0 = (1 + \sin(2\pi x/L))/2$, which instead develops a moving shock. These two choices allow us to investigate the impact of having to track a discontinuity which is not aligned with the grid.

First of all, we observe that the performance definitely improves if the coarse solver is chosen to match the fine solver up to a higher order: in fact, the blue curves in Fig. 5.4 (first-order matching, (5.47)) lie consistently below the green ones (zeroth-order matching, (5.46)). This is in line with expectations, since we are employing a more accurate solver. Still, the performance of zeroth-order matching remains competitive, achieving very effective reductions in error over the first iterations. Somewhat surprisingly, the zeroth-order matching gets slightly better initial error reductions in the moving-shock scenario than for the stationary shock. Unfortunately, though, this scheme clearly fails to effectively damp some components of the error, as we see from the fact that the corresponding error starts rising again after 5-10 iterations on finer meshes. This negative effect seems to be heightened as we refine the grid further, with the error starting to increase earlier. Nonetheless, in all the tests considered this approach behaves much better than naïve rediscritisation, for which the error blows up after the very first iteration and is hence not reported in the graph. Moreover, the convergence behaviour remains unchanged as we further refine our grids: the different lines superimpose almost perfectly. Modifying the type of multigrid cycle does not vary the nature of these observations, but as expected an F-cycle shows steeper convergence plots than a V-cycle. We have found that modifying the type of relaxation from an F- to an FCF-smoother, instead, does not improve convergence dramatically, so long as the initial guess for the solution at the coarse nodes is picked as in (5.40). An example of this is shown in Fig. 5.5, where a comparison of the error evolution of MGRIT is shown when using simple injection as a restriction operator instead of formula (5.40). The latter consistently shows better (or at worst comparable) results, particularly if MGRIT is equipped with simple multigrid components, such as V-cycle and F-relaxation as opposed to F-cycle and FCF-relaxation.

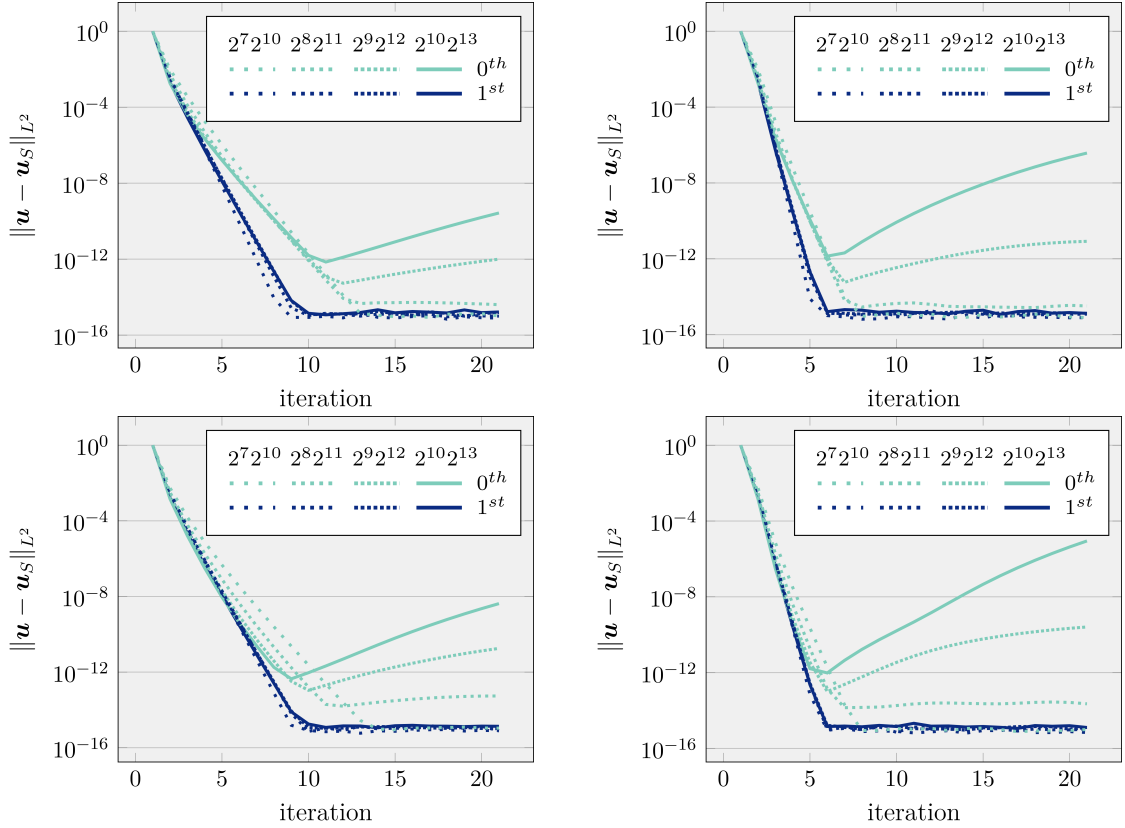


Figure 5.4: Impact of grid refinement and degree of fine-integrator matching on the evolution of the error between the MGRIT iterates and the serial solution \mathbf{u}_S . The relative L^2 norm in space-time is considered. MGRIT parameters: F-smoothing, 5 levels, coarsening factor $m = 2$, V-cycle (left), and F-cycle (right). Domain: $L = 1$, $T = 0.475$. More densely dotted graphs correspond to finer meshes: the numbers of nodes $N_x N_t$ at the finest level are reported in the legend. The time-steppers introduced in Sec. 5.2.1.1 are employed: FE with Lax-Friedrichs flux (5.43) for the finest level, and the corresponding zeroth- (5.46) and first-order matching (5.47), with varying coarsening factors, for the coarse levels (green and blue lines, respectively). Initial conditions: stationary shock (top, $u(x, 0) = \bar{u}_S^0$), and moving shock (bottom, $u(x, 0) = \bar{u}_M^0$). Results from pure rediscretisation are not reported, as the error diverges after the very first iteration.

The factor playing the main role in determining the performance of the algorithm, instead, lies in how closely the *Courant-Friedrichs-Lewy* CFL condition [89, Chap. 10.6] is met. In Fig. 5.4, the number of nodes N_t and N_x are chosen to guarantee that the CFL number on the *coarsest* temporal mesh is ~ 0.95 in all cases. The impact of relaxing the coarse-grid CFL condition can be clearly seen in the top-left graph of Fig. 5.6: convergence degrades noticeably, even using high-order matching. On the other hand, as we increase the number of levels in the multigrid scheme, this

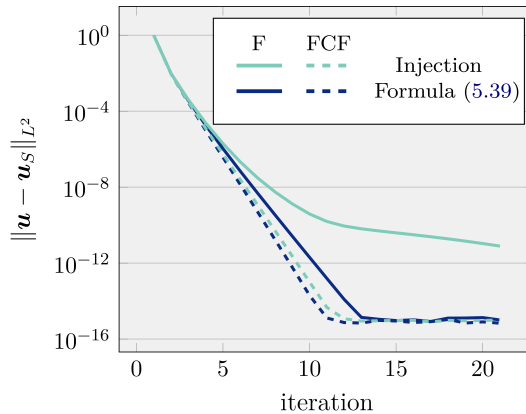


Figure 5.5: Effect of modifying the initial guess for the coarse solution on the error convergence of MGRIT: comparison of error convergence between classical injection (green) and formula (5.40) (blue). Relaxation used: F-relaxation (full lines) and FCF-relaxation (dashed lines); otherwise, same parameters as in the top-left graph in Fig. 5.4. Only results for $N_x N_t = 2^7 2^{10}$ (zeroth-order) are reported, but finer grids give comparable results.

condition becomes progressively more stringent, and translates into having to choose smaller and smaller Δt at the *finest* level, possibly over-resolving the solution.

5.3.2 High-order schemes

While the integrator proposed in Sec. 5.2.1.1 seems effective, it remains very dissipative. On the one hand, it is well-documented (see for example [35]) that smoothing effects improve convergence of MGRIT; on the other, though, artificial diffusion is very undesirable in the model problems considered, as it results in a heavy smearing of the shocks, which are the characteristic feature of the solutions of conservation laws. The choice of high-order space reconstructions is instead preferable, as they offer the possibility to preserve such discontinuities sharply, and an equally high-order time-discretisation is consequently requested. Investigating the behaviour of the MGRIT algorithm when used in conjunction with these schemes is hence of relevance, as it would more faithfully represent the setup of a real-world application. For this reason, we proceed to applying MGRIT to the test cases introduced in Sec. 5.1, for discretisations using the Lax-Friedrichs and Roe definitions of numerical fluxes, (5.17) and (5.26), and employing WENO reconstructions at various orders of accuracy s . The initial conditions for the various problems are chosen to ensure that the solutions

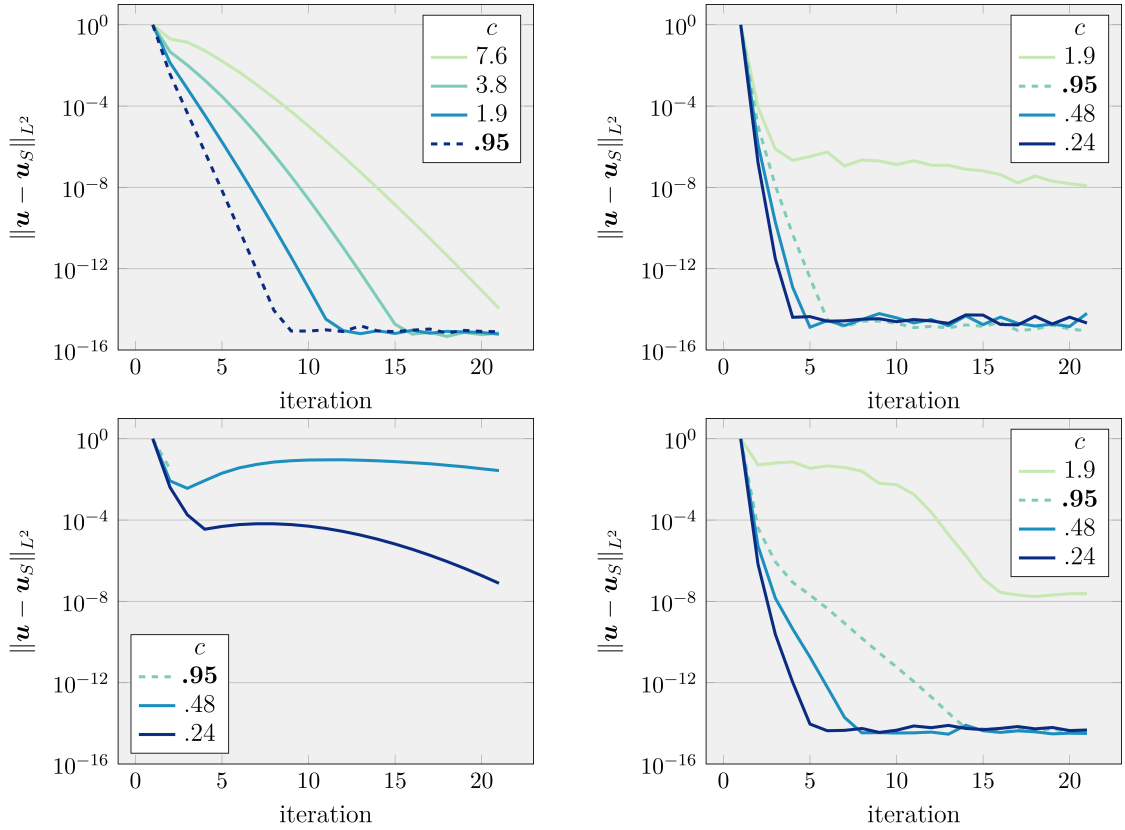


Figure 5.6: Influence of increasing the CFL number c at the coarsest level on the convergence behaviour of MGRIT applied to Burgers' equation. $N_x = 2^7$, while N_t is progressively halved to increase c . The other parameters are the same as in the top-left graph in Fig. 5.4, but different schemes are used in each plot. Top-left: first-order matching and FE; top-right: WENO reconstruction of order $s = 1$ and SSPRK3; bottom-left: WENO ($s = 5$) and FE; bottom-right: WENO ($s = 5$) and SSPRK3; the last three use a Roe flux and rediscritisation at all levels. Missing values are diverging (notice in particular the truncated $c = 0.95$ plot in the bottom-left graph).

achieve a similar maximum CFL number $c \sim 0.43$. These are:

$$\bar{u}_B^0 = \frac{4}{3} \sin\left(\frac{2\pi x}{L}\right), \quad \bar{u}_S^0 = \frac{1}{11} \begin{bmatrix} 1 + \frac{1}{2} \sin\left(\frac{2\pi x}{L}\right) \\ 0 \end{bmatrix}, \quad \text{and} \quad \bar{u}_E^0 = \begin{bmatrix} 1 \\ 0 \\ 1 + \frac{1}{2} \sin\left(\frac{2\pi x}{L}\right) \end{bmatrix}, \quad (5.49)$$

for Burgers', shallow-water, and Euler equations, respectively. From Fig. 5.7, we see that the error behaviour varies in quite an erratic way as we modify s . For both Lax-Friedrichs and Roe fluxes with the SSPRK3 time-stepper, the best performance is seen with an integrator which does not perform any sort of reconstruction, *i.e.*, $s = 1$. This is in line with expectations, as this choice of s corresponds to the most dissipative among the schemes considered. The tendency is for performance to

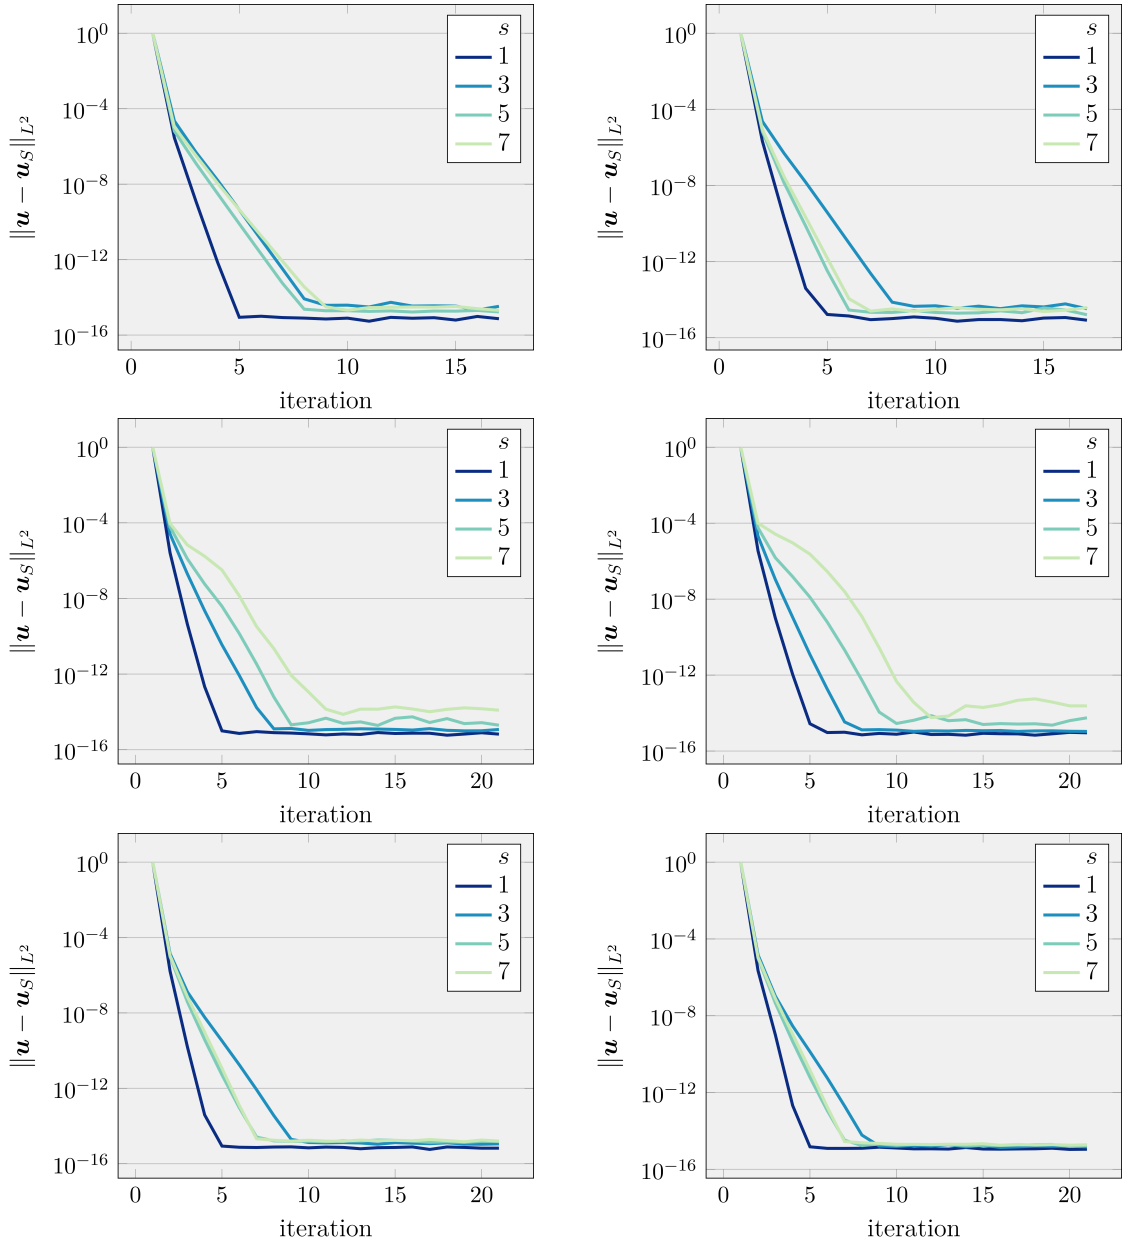


Figure 5.7: Evolution of the error in the MGRIT iterates, applied to Burgers' (top), shallow-water (middle) and Euler equations (bottom). MGRIT parameters: V-cycle, F-smoothing, 3 levels with $N_t = 100, 200, 400$, respectively, $N_x = 64$. Domain: $L = 1$, $T = 0.5$. Flux discretisation used: Lax-Friedrichs (left plots, (5.17)), and Roe (right plots, (5.26)), both employing WENO reconstruction of different orders of accuracy s . Time-stepper used: SSPRK3 (5.28) with rediscetisation at all levels.

worsen as we increase s , although $s = 3$ seems to provide an exception and behave particularly badly in some cases: the reason behind this remains unclear. Nonetheless, these higher-order methods still behave reasonably well, particularly in comparison to many results in the literature regarding more diffusive schemes. Evidence of this

is also given in Fig. 5.6: at least for $c \leq .95$, which corresponds to the dashed lines in the plots, SSPRK3 (top-right) outperforms FE (top-left), even if the latter is used in conjunction with the first order-matching procedure (5.47). Indeed it is accuracy in the *time* discretisation that helps convergence, while increasing the order of the *space* discretisation seems to produce the opposite effect, as already discussed before, and further testified by the poorer error behaviour in the bottom plots of Fig. 5.6. In particular, note how MGRIT (with rediscritisation on the coarse grid) fails to produce a converging solver for $c = .95$ (or even $c = 0.48$) in the bottom-left of Fig. 5.6, where it is used in combination with a high-order spatial discretisation and a low-order temporal discretisation.

Ultimately, however, even in this case the CFL condition has shown to be the true bottleneck in the effectiveness of such solvers. Indeed we can see from the plots in Fig. 5.6 that convergence of MGRIT using high-order schemes is much more sensitive to an increase in the CFL number c .

5.3.3 Changing accuracy per coarsening level

Experimenting with discretisations of different accuracy at different multigrid levels makes sense in the attempt to match the action of the fine solver: one can think of making up for the loss of precision due to coarsening by increasing the order of the reconstruction at a coarser level, and still obtain a result similar to what the fine integrator would provide. This might provide a feasible coarse solver, provided the additional computational cost associated with reconstruction is reasonably small. With this motivation in mind, we test applications of MGRIT where the order of the discretisations used, (both in space and time), varies across the levels. Unfortunately, the results reported in Fig. 5.8 pinpoint that this strategy is not beneficial to the algorithm performance and that, instead, simple rediscritisation is preferable. Among the alternatives considered, though, decreasing the accuracy as we descend to coarser levels has proven to be the most promising, which could make it attractive if one seeks to make the application of coarse integrators cheaper: this is also in line with the results shown in [109], where a lower-order WENO reconstruction was used at the coarse level.

5.4 Remarks on MGRIT for hyperbolic equations

From the experiments conducted in this chapter, it appears clear that the necessity of satisfying the CFL condition at each coarsening level represents a major hindrance




Figure 5.8: Impact of modifying the accuracy of the schemes used across the levels on the error convergence of MGRIT. Same parameters as in Fig. 5.7, but with the order of the WENO reconstruction s either increasing (green) or decreasing (blue) as we move to coarser levels. Dashed lines show error convergence if both s and the order accuracy of the time-stepper, d , are changed simultaneously, while solid lines show results if only s is varied. Left is for Lax-Friedrichs flux, right is for Roe flux. Only results for the application to Burgers’ equation are shown, but the ones collected for the other test-cases follow the same qualitative behaviour.

towards the stability of MGRIT, when applied to accelerate the solution of conservation laws discretised with explicit time-stepping schemes. This seems to indicate that, for a multigrid-in-time algorithm to be effective in this framework, its coarsening strategy needs to act on both the spatial and temporal domain, so that the CFL number is kept close to 1 at each level. This complicates the analysis even further, given how most of the currently available theoretical results on the convergence of MGRIT rely heavily on the algorithm implementing a purely temporal coarsening procedure. Nonetheless, some research in this sense has already been conducted [78, 138], and is providing further evidence of the effectiveness of employing a combined *space-time* multigrid method. Indeed the AIR algorithm introduced in Sec. 2.2.4 makes use of such strategies, on top of employing restriction techniques specific to advection problems [138].

In hindsight, AIR could have been a valuable method for the focus of our analysis in tackling time parallelisation of hyperbolic equations, but at the time the research reported in this and the following chapters was conducted, this scheme was still at its preliminary stages. For this reason, instead, we decided to work on the study of algorithms that were more mature, such as MGRIT in this chapter, or that had already started to prove their applicability in the hyperbolic case. The method investigated in the following chapter represents an alternative to MGRIT in this sense.

Chapter 6

Space-time circulant preconditioning for linear wave equations

The use of circulant preconditioners to accelerate the solution of Krylov subspace iterative methods for Toeplitz systems is a well-established practice, dating back to the work of Strang in 1986 [145], and employed in a range of frameworks ever since (see *e.g.* [108] for details on the iterative solution of Toeplitz systems, as well as a number of their applications). In more recent years, this practice has been investigated for accelerating the all-at-once solution of monolithic systems arising from the simultaneous space-time discretisation of time-evolving PDEs [3,64,65,87,93,101].

When dealing with constant-coefficient, linear PDEs discretised on a uniform temporal grid, in fact, the resulting space-time system presents a block Toeplitz structure, whose blocks are constant down each diagonal, making it amenable to circulant preconditioning. Systems involving circulant matrices are relatively simple to invert by means of the Fourier transform, which effectively diagonalise the operator, thus allowing for parallelisation in its solution. This is what makes circulant preconditioning for space-time matrices an actual parallel-in-time method. An appealing property of this strategy is that, when applied to hyperbolic PDEs, it does not seem to suffer from the same stability issues afflicting other PinT methods (see Sec. 2.1.1.1, as well as the analysis in Chap. 5), thus making it a useful alternative to achieve time parallelisation in this framework.

In this chapter, we show how a theoretical result from Chan, Potts, and Steidl [16] on the convergence of circulant-preconditioned GMRES applies to the solution of space-time systems. Under certain assumptions, this ensures that the iterative method will converge (terminate) in a number of steps which is independent of the

number of nodes in the temporal mesh, thus giving an indication of the effectiveness of the space-time circulant preconditioning strategy. We provide extensive experimental results that measure its performance when applied to linear wave equations, using high-order discretisations tailored for this type of problems, which confirm the theoretical results and provide further evidence for the applicability of space-time circulant preconditioning for hyperbolic PDEs.

In Sec. 6.1, we expand on the introduction to circulant preconditioners provided in Sec. 2.2.2, and describe the main theoretical convergence result exploited in this chapter. We proceed to test the effectiveness of the space-time circulant preconditioning strategy, first applying it to a toy ODE problem in Sec. 6.2, and then to a more interesting target PDE in Sec. 6.3, providing the relevant experimental results in the corresponding sections.

The analysis and results from this chapter have been published in the *Numerical Linear Algebra with Applications* journal, and are available at [32].

6.1 Circulant preconditioning for Toeplitz systems

In the literature, a variety of circulant preconditioners for Toeplitz matrices have been proposed: some examples of these are listed in [17, Chap. 2], and indeed some of these approaches coincide for the simple test cases considered here. In this chapter we exploit some of the results from R. Chan, D. Potts, and G. Steidl [16]: therefore, we follow most of the notation used in their paper, and describe the ω -circulant preconditioner therein defined.

6.1.1 ω -circulant preconditioner

Let us introduce a 2π -periodic *generating function* $f(x)$, with the associated Toeplitz matrix $A_N(f) \in \mathbb{R}^{N \times N}$, that is:

$$[A_N(f)]_{m,n} := a_{m-n}(f), \quad \text{with} \quad a_k(f) := \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) e^{-ikx} dx, \quad (6.1)$$

where the $a_k(f)$ are the *Fourier coefficients* of $f(x)$. In our work we are mostly concerned with *banded* Toeplitz matrices with a small band, which have *trigonometric polynomials* as generating functions, in the form

$$f(x) = \sum_{k=-s_2}^{s_1} a_k e^{ikx}, \quad \text{for some } a_k \in \mathbb{R}, \quad s_1, s_2 \in \mathbb{N}, \quad 0 < s_1 + s_2 \ll N. \quad (6.2)$$

that is, in order to recover $M_N(f)$, the off-diagonal elements of $A_N(f)$ are summed to those N diagonals away. Notice that in our particular setting, where the band of $A_N(f)$ is small (specifically, smaller than $N/2$), and $\omega = 1$, the ω -preconditioner coincides with the *Strang circulant* preconditioner [145].

Formula (6.7) implies that inverting a circulant matrix amounts to applying a *Fast Fourier Transform* (FFT), inverting a diagonal matrix, and applying the inverse FFT: a task of overall complexity $\mathcal{O}(N \log N)$. This argument can be similarly extended to block circulant matrices, in which case we would achieve block diagonalisation. Therein also lies the concurrency of the algorithm: solving a (block) diagonal system is a task which can be trivially parallelised. Being relatively simple to invert, (6.7) consequently satisfies the first requirement of a preconditioner. As a second requirement, we would also hope for it to effectively accelerate the convergence of iterative methods. We address this in the following.

6.1.2 Convergence of circulant-preconditioned Krylov methods

There is a choice of iterative methods one can use, paired with the preconditioner (6.7) (or slight modifications thereof), to solve a system of equations involving (6.3). For example, following [101], one could symmetrise the Toeplitz system via a Hankel matrix, and use MINRES with the absolute value of (6.7) as preconditioner. An alternative explored in [16] consists in solving the normal equation using *Conjugate Gradient*, preconditioned with $M_N(|f|^2)$. In both cases, strong theoretical bounds guarantee fast convergence of the methods used. We are mostly concerned with convergence of circulant preconditioners applied to GMRES. While convergence results are harder to come by in this case [68], some experiments [3, 64] hint at it being a more effective iterative solver with respect to other choices available.

In this section, we use a result from Chen, Potts and Steidl [16], which deals with the application of (6.7) as a right-preconditioner for GMRES iterations for a system involving (6.3). Adapted to our case, the theorem states the following:

Theorem 6.1.1. *Let f be a trigonometric polynomial such as (6.2), satisfying (6.4), and generating the Toeplitz matrix $A_N(f)$, with associated ω -circulant preconditioner $M_N(f)$. Then,*

$$A_N(f)M_N(f)^{-1} = I_N + R_N(s_1 + s_2), \quad (6.9)$$

where $I_N \in \mathbb{R}^{N \times N}$ is the identity matrix, while $R_N(s) \in \mathbb{R}^{N \times N}$ identifies a matrix of rank s .

Proof. In our situation, the proof simplifies significantly. We can split $A_N(f)$ as $M_N(f) - B_N(f)$, where

$$B_N(f) = \begin{bmatrix} & & & & a_{s_1} & a_{s_1-1} & \cdots & a_1 \\ & & & & & \ddots & \ddots & \vdots \\ & a_{-s_2} & & & & & \ddots & a_{s_1-1} \\ a_{-s_2+1} & \ddots & & & & & & a_{s_1} \\ \vdots & \ddots & \ddots & & & & & \\ a_{-1} & \cdots & a_{-s_2+1} & a_{-s_2} & & & & \end{bmatrix}, \quad (6.10)$$

so that

$$A_N(f)M_N(f)^{-1} = I_N(f) - B_N(f)M_N(f)^{-1}, \quad (6.11)$$

from which it follows that $B_N(f)M_N(f)^{-1} = R_N(s_1 + s_2)$, since $B_N(f)$ only has non-zero elements in the top s_1 and bottom s_2 rows. \square

Remark 3. Notice that the theorem allows for some flexibility in breaking the Toeplitz structure of $A_N(f)$. If we consider a perturbation $\tilde{A}_N(f)$ to $A_N(f)$, we can still decompose $A_N(f) + \tilde{A}_N(f) = M_N(f) - \tilde{B}_N(f)$, where $\tilde{B}_N(f) = B_N(f) - \tilde{A}_N(f)$. Provided that the entries in $\tilde{A}_N(f)$ which do not adhere to the Toeplitz structure are only limited to the first s_1 and last s_2 rows, then, Thm. 6.1.1 still holds. This is particularly useful if we consider that we are dealing with the solution of discrete space-time systems. In this framework, in fact, we are often forced to modify the equations corresponding to the first few temporal nodes, be it in order to accommodate for initial conditions on the derivative, and/or in order to kick-start high-order discretisations in time: this has the effect of perturbing the Toeplitz structure of the system, but not to an extent that would break the result of Thm. 6.1.1, as described more in detail in Sec. 6.2 and 6.3.

As a note, the original theorem is presented in [16] in a more general form, and deals rather with *rational* generating functions:

$$f(x) = \frac{p(x)}{q(x)} = \frac{\sum_{k=0}^{d_1} p_k e^{ikx}}{\sum_{k=0}^{d_2} q_k e^{ikx}}, \quad \text{with } d_1, d_2 \in \mathbb{N}, \quad d_1 d_2 \neq 0. \quad (6.12)$$

In that case, the original theorem states that the rank of the resulting perturbation of the identity is equal to $\max\{d_1, d_2\}$. It is possible to recover our version of Thm. 6.1.1 by noticing that we can rewrite (6.2) as $f(x) = (\sum_{k=0}^{s_1+s_2} a_{-s_2+k} e^{ikx}) / e^{is_2x}$.

Noticeably, Thm. 6.1.1 gives us a theoretical guarantee that GMRES with $M_N(f)$ as a right-preconditioner will converge to the solution after at most $s_1 + s_2 + 1$ iterations, at least in exact precision: if the band of $A_N(f)$ is small enough, then (6.7) provides an excellent preconditioner.

In the following sections, we illustrate how simple systems arising from the discretisation of differential equations are (close-to-) Toeplitz in nature, which makes them amenable to be preconditioned using (6.7). The goal is to provide examples of the effectiveness of this preconditioner, as well as to discuss cases in which this theorem fails to come to our aid. To this purpose, we introduce a few model problems, and the methods used in our experiments to recover their numerical solution.

6.2 All-at-once solution of linear ODEs

As a proof of concept, useful in introducing some of the notation employed, we propose the simplest linear second-order differential equation:

$$\begin{cases} \frac{d^2u}{dt^2} - \lambda u = 0 & t \in (0, T] \\ u(0) = \bar{u}^0 \\ \frac{du}{dt}(0) = \bar{u}^1 \end{cases}, \quad (6.13)$$

where \bar{u}^0 and $\bar{u}^1 \in \mathbb{R}$ represent the initial values of the solution itself and its derivative, respectively.

6.2.1 Discretisation aspects

To discretise our temporal domain, we choose a uniform grid with step size $\Delta t = T/(N_t - 1)$. Furthermore, to approximate (6.13) we consider *Störmer-Verlet* schemes (SV) of various orders: this is a family of multi-step methods specifically tailored for the solution of second-order differential equations [70, Chap. III-10]. The general formulation for a k -th order SV scheme is given by

$$u^n - 2u^{n-1} + u^{n-2} = \Delta t^2 \sum_{j=0}^{k-1} \alpha_j D_-^j f^n, \quad (6.14)$$

where the right-hand side of the system f^n simply corresponds to λu^n in our case; D_- identifies the backward difference operator

$$D_-^j f^n = D_-^{j-1} f^n - D_-^{j-1} f^{n-1}, \quad \text{with} \quad D_-^0 f^n = f^n, \quad (6.15)$$

while the coefficients α_j are picked in such a way to ensure the desired level of accuracy [70, Table.10.2]: their values for the schemes used in this chapter are

$$[\alpha_j]_{j=0}^8 = \left[1, -1, \frac{1}{12}, 0, -\frac{1}{240}, -\frac{1}{240}, -\frac{221}{60480}, -\frac{19}{60480}, -\frac{9829}{3628800} \right]. \quad (6.16)$$

Notice that for all cases but $k = 2$, (6.14) defines an implicit method. For ease of notation, we identify discretisation schemes of different accuracies by reporting their order as a subscript to their acronym. We can see that choosing SV_1 corresponds to approximating the second-order derivative using a first-order backward difference formula, which produces the following recurrence relation:

$$\begin{cases} \frac{u^n - 2u^{n-1} + u^{n-2}}{\Delta t^2} - \lambda u^n = 0 & n = 2, \dots, N_t - 1 \\ \frac{2u^1 - 2u^0}{\Delta t^2} - \lambda u^1 = \frac{2}{\Delta t} \bar{u}^1 \\ u^0 = \bar{u}^0 \end{cases}. \quad (6.17)$$

Notice that we use a *ghost node* approach to include the initial condition on the derivative at $t = 0$; that is, we introduce a fictitious node u^{-1} , approximate the first derivative using the *central difference* (CD) formula $u^1 - u^{-1} = 2\Delta t \bar{u}^1$, and substitute in the first equation of (6.17) with $n = 1$ to recover the second equation. We can gather the equations (6.17) together and build the system

$$\underbrace{(\Phi^{SV_1} - \lambda \Delta t^2 \Phi_r^{SV_1})}_{=: A^{SV_1}} \mathbf{u} = \mathbf{g}^{SV_1}, \quad (6.18)$$

where $\mathbf{u} = [u^0, u^1, \dots, u^{N_t-1}]^T \in \mathbb{R}_t^N$ collects the values of the solution at each temporal node. The operator $\Phi^{SV_1} \in \mathbb{R}^{N_t \times N_t}$ represents the discretisation of the time derivative (including the perturbation from the ghost node approach), and has the following structure:

$$\Phi^{SV_1} := \begin{bmatrix} 1 & & & & & & \\ -2 & 2 & & & & & \\ 1 & -2 & 1 & & & & \\ & \ddots & \ddots & \ddots & & & \\ & & & & 1 & -2 & 1 \end{bmatrix} \in \mathbb{R}^{N_t \times N_t}. \quad (6.19)$$

Operator $\Phi_r^{SV_1}$ stems from the evaluation of the right-hand side in (6.14), and it reduces to $\Phi_r^{SV_1} = I^{N_t} \in \mathbb{R}^{N_t \times N_t}$, the identity matrix, for SV_1 . The effect of the initial conditions is included in the right-hand side:

$$\mathbf{g}^{SV_1} = [(1 - \lambda \Delta t^2) \bar{u}^0, 2\Delta t \bar{u}^1, 0, \dots, 0]^T \in \mathbb{R}^{N_t}. \quad (6.20)$$

Notice how imposing the initial conditions breaks the Toeplitz structure of the matrix (6.19), and hence that of the system A^{SV_1} : the second element in the second

row is different from the others on the diagonal. However, we can still assemble the circulant preconditioner disregarding this perturbation, as such:

$$M^{\text{SV}_1} := \hat{\Phi}^{\text{SV}_1} - \lambda \Delta t^2 \Phi_r^{\text{SV}_1}, \quad (6.21)$$

where $\hat{\Phi}^{\text{SV}_1}$ is our circulant approximation to Φ^{SV_1} , built as follows:

$$\hat{\Phi}^{\text{SV}_1} := \begin{bmatrix} 1 & & & 1 & -2 \\ -2 & 1 & & & 1 \\ 1 & -2 & 1 & & \\ & \ddots & \ddots & \ddots & \\ & & & 1 & -2 & 1 \end{bmatrix} \in \mathbb{R}^{N_t \times N_t}. \quad (6.22)$$

As hinted in Rmk. 3, the statement of Thm. 6.1.1 remains valid using M^{SV_1} as a preconditioner for A^{SV_1} , even if the system is not strictly Toeplitz: in fact, we only perturbed the second row on a matrix with a band of size 2.

We also experiment using higher-order discretisations. For orders > 2 , this requires recovering the values of the solution at the first few time steps in some alternative way, in order to *kick-start* the multi-step method [90, Chap. 5.9.3]. This is due both to the stencil generally becoming larger as the order increases, and to the fact that the ghost-node approach used in (6.17) is only second-order accurate, and will end up polluting the global accuracy of higher-order schemes. There are a number of ways in which the kick-starting procedure might be conducted, (generally resorting to self-starting schemes to solve for the first few unknowns), which end up disrupting the Toeplitz structure of the space-time system. For simplicity, in our experiments we directly evaluate the analytical solution for the necessary nodes; however, in order to mimic the impact the perturbation stemming from a kick-starting procedure might have on the performance of the circulant preconditioner, we also modify the first few equations in the system by leaving only the contribution from the main diagonal. For example, using SV_5 gives rise to the space-time system:

$$\underbrace{(\Phi^{\text{SV}_5} - \lambda \Delta t^2 \Phi_r^{\text{SV}_5})}_{=: A^{\text{SV}_5}} \mathbf{u} = \mathbf{g}^{\text{SV}_5}, \quad (6.23)$$

where Φ^{SV_5} and $\Phi_r^{\text{SV}_5}$ play the same roles as Φ^{SV_1} and $\Phi_r^{\text{SV}_1}$ in (6.19), but both have

Table 6.1: Number of iterations to convergence for GMRES, right-preconditioned with (6.7), applied to the solution of systems arising from the discretisation of (6.13). Systems of different size N_t , and SV schemes of various orders of accuracy are considered. The top row reports the theoretical upper bound from Thm. 6.1.1. The two columns of values refer to different tolerances for convergence: $\varepsilon = 10^{-7}$ (left, in brackets), $\varepsilon = 10^{-10}$ (right). The initial guess for GMRES is randomly chosen (from a normal distribution). The problem parameters are: $\bar{u}^0 = 1$, $\bar{u}^1 = -1$, $\lambda = -1$, $T = 10^3$.

	SV ₁	SV ₂	SV ₃ (= SV ₄)	SV ₅	SV ₆	SV ₇	SV ₈
$s_1 + s_2 + 1$	3	3	3	5	6	7	8
$N_t = 10^2 \cdot 2^3$	(3) 3	(3) 3	(3) 3	(5) 5	(6) 6	(7) 7	(8) 8
$N_t = 10^2 \cdot 2^4$	(3) 3	(3) 3	(3) 3	(5) 5	(6) 6	(7) 7	(8) 8
$N_t = 10^2 \cdot 2^5$	(3) 3	(3) 3	(3) 3	(5) 5	(6) 6	(7) 7	(8) 8
$N_t = 10^2 \cdot 2^6$	(3) 3	(3) 3	(3) 3	(5) 5	(6) 6	(7) 7	(8) 8
$N_t = 10^2 \cdot 2^7$	(3) 3	(3) 3	(3) 3	(5) 5	(6) 6	(7) 8	(8) 9
$N_t = 10^2 \cdot 2^8$	(3) 3	(3) 3	(3) 4	(5) 6	(6) 7	(7) 8	(8) 9
$N_t = 10^2 \cdot 2^9$	(3) 3	(3) 4	(3) 4	(5) 6	(6) 7	(7) 8	(8) 9
$N_t = 10^2 \cdot 2^{10}$	(3) 3	(3) 4	(3) 4	(5) 6	(6) 7	(7) 8	(9) 9
$N_t = 10^2 \cdot 2^{11}$	(3) 4	(3) 4	(3) 4	(6) 6	(7) 7	(7) 8	(9) 9

refined grids. The ω -preconditioner is further tuned by picking the parameter w in (6.4) as $w = \pi/N_t$, which (at least for the lowest-order schemes) ensures that the roots of the corresponding generating function are the farthest away from 0; this in turns attains the desirable property of minimising the condition number of the preconditioner. A much more detailed analysis of the impact of this parameter on the effectiveness of the preconditioner, however, has been presented in [93].

6.3 All-at-once solution of linear PDEs

Ultimately, our interest resides in exploiting a (block) circulant preconditioner for the solution of hyperbolic partial differential equations. In presenting the test-case for this class of PDEs, we mostly follow the set-up provided in [65], and consider the *wave equation* with Dirichlet boundary conditions:

$$\begin{cases} \frac{\partial^2 u}{\partial t^2} - \frac{\partial^2 u}{\partial x^2} = 0 & (x, t) \in (0, L) \times (0, T] \\ u(x, 0) = \bar{u}^0(x) & x \in [0, L] \\ \frac{\partial u}{\partial t}(x, 0) = \bar{u}^1(x) & \\ u(0, t) = u(L, t) = 0 & t \in (0, T] \end{cases} \quad (6.27)$$

6.3.1 Discretisation aspects

Most of the notation used in Sec. 6.2 can be extended to the PDE case, but we need to introduce an appropriate discretisation of the spatial operator. For this purpose, we choose once again a uniform grid spanning the interior of the spatial domain, with spacing $\Delta x = L/(N_x + 1)$, and we consider a *central difference* scheme (CD) for the second-order spatial derivative. Also for spatial discretisation we make use of approximations of different orders of accuracy, to match the temporal discretisations. Furthermore, the Dirichlet boundary conditions in (6.27) are imposed naturally. This causes some inconvenience when using high order schemes: in that case, in fact, we cannot use symmetric stencils anymore as we approach the boundaries, since we would require information from nodes outside the domain. We circumvent this issue by using, for those few nodes, finite difference of the same order of accuracy, but defined on a slanted stencil [52]. For example, for a second-order accurate discretisation, we obtain the following matrix:

$$\mathcal{K}^{\text{CD}_2} := \begin{bmatrix} -2 & 1 & & & & & \\ & 1 & -2 & 1 & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & & 1 & -2 & 1 \\ & & & & & 1 & -2 \end{bmatrix} \in \mathbb{R}^{N_x \times N_x}; \quad (6.28)$$

while for a fourth-order one, we get

$$\mathcal{K}^{\text{CD}_4} := \frac{1}{12} \begin{bmatrix} -6 & 14 & -4 & -15 & 10 & & & & \\ 16 & -30 & 16 & 1 & & & & & \\ -1 & 16 & -30 & 16 & -1 & & & & \\ & \ddots & \ddots & \ddots & \ddots & \ddots & & & \\ & & -1 & 16 & -30 & 16 & -1 & & \\ & & & -1 & 16 & -30 & 16 & & \\ & & & 10 & -15 & -4 & 14 & -6 & \end{bmatrix} \in \mathbb{R}^{N_x \times N_x}. \quad (6.29)$$

We point out that the scheme stemming from choosing SV_2 for the temporal discretisation and CD_2 for the spatial one is also known as the *explicit Leapfrog* method [90, Chap. 10.2.2], which is *conditionally stable*. This can be shown via Von-Neumann analysis [90, Chap. 9.6]: assuming the solution is composed of a single Fourier mode $u^n = G^n e^{ikx}$, for a certain *growth factor* G and frequency k , and substituting this into the scheme, we obtain the recurrence relation

$$\frac{G^n - 2G^{n-1} + G^{n-2}}{\Delta t^2} e^{ikx} = \frac{(e^{ik\Delta x} - 2 + e^{-ik\Delta x})}{\Delta x^2} G^{n-1} e^{ikx}. \quad (6.30)$$

We are interested in the roots of the associated characteristic polynomial, given by the solutions of

$$g^2 - 2(1 + r^2 c_2)g + 1 = 0, \quad (6.31)$$

where we named $c_2 = \cos(k\Delta x) - 1 \in (-2, 0)$, and $r = \Delta t/\Delta x$, in order to reduce notation. Its roots, denoted as g_{\pm} , satisfy the relationships

$$g_+ g_- = 1, \quad (6.32a)$$

$$\text{and } g_+ + g_- = 2(1 + r^2 c_2). \quad (6.32b)$$

We want to ensure that $|g_{\pm}| \leq 1 \forall k$, and that roots falling on the unit circle are simple. This can only be guaranteed if these are given as a complex conjugate pair of modulus 1 (otherwise they would break (6.32a)). Consequently, we can translate (6.32b) into requesting $|\Re(g_{\pm})| = |1 + r^2 c_2| \leq 1$, which only holds if $r^2 \leq 1$, given the range of c_2 . For the fourth order discretisation, a similar reasoning can be applied. The roots of the characteristic polynomial associated with using SV_4 in time and CD_4 in space, are given by

$$\begin{aligned} & \left(1 - \frac{1}{12}r^2 c_4\right) g^2 - 2 \left(1 + \frac{5}{12}r^2 c_4\right) g + \left(1 - \frac{1}{12}r^2 c_4\right) = 0 \\ \implies & g^2 - 2 \left(\frac{12 + 5r^2 c_4}{12 - r^2 c_4}\right) g + 1 = 0, \end{aligned} \quad (6.33)$$

this time with $c_4 = (16 \cos(k\Delta x) - \cos(2k\Delta x) - 15)/12 \in (-8/3, 0)$. The conditions on its roots resemble those in (6.32); for stability, we then need to impose

$$|\Re(g_{\pm})| = \left|\frac{12 + 5r^2 c_4}{12 - r^2 c_4}\right| \leq 1 \implies r^2 \leq \frac{9}{4}. \quad (6.34)$$

For higher-order methods the analysis gets more complex, as increasing the size of the stencil increases the order of the associated characteristic polynomial, but they also require constraints on the size of r in order to ensure stability.

To avoid these restrictions, we also introduce another discretisation, namely an implicit version of the Leapfrog method [105], which we consider only to second-order accuracy, and which we denote by CI_2 . The only difference with respect to the explicit Leapfrog scheme lies in the fact that the spatial operator is averaged over three consecutive time steps, giving rise to the following recurrence relation:

$$\frac{u^n - u^{n-1} + u^{n-2}}{\Delta t^2} = \frac{1}{\Delta x^2} \mathcal{K}^{CD_2} \frac{u^n + 2u^{n-1} + u^{n-2}}{4}. \quad (6.35)$$

Even in this case, von Neumann analysis comes to our aid in establishing that the method is *unconditionally stable*. After some algebraic manipulations, we can show that the characteristic polynomial of interest this time is given by

$$g^2 - 2 \left(\frac{2 + r^2 c_2}{2 - r^2 c_2} \right) g + 1 = 0, \quad (6.36)$$

with $c_2 \in (-2, 0)$ defined as before. Similar considerations as for the case above show that its roots are a complex conjugate pair falling on the unit circle whenever $|2 + r^2 c_2| < |2 - r^2 c_2|$; given c_2 is always negative, this holds $\forall r^2 > 0$.

We are finally ready to assemble the monolithic space-time system, by opportunely combining the matrices introduced above. For CI_2 , or SV_k temporal discretisations of a given order k , we do so in the following block-wise fashion:

$$\begin{aligned} A^{\text{SV}_k} &:= \Phi^{\text{SV}_k} \otimes I^{N_x} - r^2 \left(\Phi_r^{\text{SV}_k} \otimes \mathcal{K}^{\text{CD}_2[\frac{k}{2}]} \right), \\ \text{and } A^{\text{CI}_2} &:= \Phi^{\text{SV}_2} \otimes I^{N_x} - \frac{r^2}{4} \left(\tilde{I}^{N_t} + 2I_{-1}^{N_t} + I_{-2}^{N_t} \right) \otimes \mathcal{K}^{\text{CD}_2}, \end{aligned} \quad (6.37)$$

where \otimes represents the *Kronecker product*, $[\ast]$ the *ceiling function*, and $I_{-i}^{N_t}$ is the *lower shift matrix*, containing only ones on the i -th sub-diagonal. Also, \tilde{I}^{N_t} is a slight modification of the identity matrix, containing a 2 in the second row, to account for the initial conditions on the derivative. Notice how the part for A^{CI_2} referring to the spatial discretisation (the term containing the Kronecker product with $\mathcal{K}^{\text{CD}_2}$) is made of three adjacent blocks per block row, corresponding to an averaging over three consecutive instants. The corresponding right-hand sides are assembled in a similar way as to what was done in Sec. 6.2:

$$\begin{aligned} \mathbf{g}^{\text{SV}_1} &= \left[(I^{N_x} - r^2 \mathcal{K}^{\text{CD}_2}) \bar{u}^0, 2\Delta t \bar{u}^1, 0, \dots, 0 \right]^T, \\ \mathbf{g}^{\text{SV}_k} &= \left[[A^{\text{SV}_k}]_{i,j=0}^{N_x-1} \bar{u}^0, \dots, [A^{\text{SV}_k}]_{i,j=(k-2)N_x}^{(k-1)N_x-1} \bar{u}^{k-2}, 0, \dots, 0 \right]^T \quad \text{for } k \geq 3, \\ \text{and } \mathbf{g}^{\text{CI}_2} &= \left[\left(I^{N_x} - \frac{r^2}{4} \mathcal{K}^{\text{CD}_2} \right) \bar{u}^0, 2 \left(\Delta t I^{N_x} - \frac{r^2}{4} \mathcal{K}^{\text{CD}_2} \right) \bar{u}^1, 0, \dots, 0 \right]^T. \end{aligned} \quad (6.38)$$

Notice that, since we defined the monolithic systems discretising (6.27) as a Kronecker product with the (quasi) Toeplitz matrices in Sec. 6.2, we still retain a (quasi) block Toeplitz structure even for the PDE case. We can then build our block circulant preconditioner.

6.3.2 Block circulant preconditioning

Analogously to what was done in Sec. 6.2, the following block circulant preconditioners are built for the systems (6.37),

$$\begin{aligned} M^{\text{SV}_k} &:= \hat{\Phi}^{\text{SV}_k} \otimes I^{N_x} - r^2 \left(\hat{\Phi}_r^{\text{SV}_k} \otimes \mathcal{K}^{\text{CD}_2 \lceil \frac{k}{2} \rceil} \right), \\ \text{and } M^{\text{Cl}_2} &:= \hat{\Phi}^{\text{SV}_2} \otimes I^{N_x} - \frac{r^2}{4} \left(I^{N_t} + 2\hat{I}_{-1}^{N_t} + \hat{I}_{-2}^{N_t} \right) \otimes \mathcal{K}^{\text{CD}_2}. \end{aligned} \quad (6.39)$$

where $\hat{I}_{-i}^{N_t}$ is the *circulant lower shift matrix*, that is the circulant counterpart to $I_{-i}^{N_t}$, containing ones both on the i -th sub-diagonal and the $(N_t - i)$ -th super-diagonal. As described in Sec. 6.1, these preconditioners can be (block) diagonalised using a (block) FFT. For instance, for SV_2 we have:

$$M^{\text{SV}_2} = (F_{N_t} \otimes I^{N_x}) \underbrace{\begin{pmatrix} I^{N_t} \otimes I^{N_x} \\ - \Lambda_1^{N_t} \otimes (2I^{N_x} + r^2 \mathcal{K}^{\text{CD}_2}) \\ + \Lambda_2^{N_t} \otimes I^{N_x} \end{pmatrix}}_{=: D^{\text{SV}_2}} (F_{N_t}^* \otimes I^{N_x}), \quad (6.40)$$

where $\Lambda_n^{N_t}$ are the diagonal matrices containing the N_t -th roots of unity, sampled at frequency n : more explicitly, we have $\Lambda_n^{N_t} = \text{diag}_{k=0, \dots, N_t-1} (e^{i2\pi nk/N_t}) = F_{N_t}^* \hat{I}_{-n}^{N_t}$. In order to apply the inverse of (6.40) to a space-time vector $\mathbf{b} \in \mathbb{R}^{N_x N_t}$, we then need to perform the following sequence of operations:

Step 1 Apply $F_{N_t}^* \otimes I^{N_x}$. This corresponds to computing a total of N_x different FFTs on the coefficients of \mathbf{b} . Each of these FFTs can be performed independently. The signals that need to be Fourier-transformed, denoted as $\mathbf{s}_n \in \mathbb{R}^{N_t}$, are each of length N_t and are recovered by collecting every N_x coefficients of \mathbf{b} , starting from the n -th coefficient:

$$[\mathbf{s}_n]_i = [\mathbf{b}]_{iN_x+n}, \quad \text{with } n = 0, \dots, N_x - 1, \quad \text{and } i = 0, \dots, N_t - 1. \quad (6.41)$$

Basically, we are tracking how the solution at each node evolves in time, and applying a Fourier-transform to each of these evolutions. The FFTs produce the transformed signals $\hat{\mathbf{s}}_n = F_{N_t}^* \mathbf{s}_n$. Their coefficients are then re-arranged in order to recover, as a final result, the vector

$$[\hat{\mathbf{s}}]_i = [\hat{\mathbf{s}}_{\lfloor i/N_x \rfloor}]_{i \% N_x}, \quad i = 0, \dots, N_x N_t - 1. \quad (6.42)$$

Step 2 Invert D^{SV_2} . As already pointed out, this matrix is block diagonal. As a consequence, this step requires solving for N_t independent systems (the blocks on

the diagonal), using $[\hat{\mathbf{s}}]_{nN_x}^{(n+1)N_x-1}$, $n = 0, \dots, N_t - 1$, as right-hand sides. Looking at (6.40) we can see how each of these solves requires inverting a complex-shifted discrete Laplacian operator.

Step 3 Apply $F_{N_t} \otimes I^{N_x}$. The last step is equivalent to the first, except N_x inverse FFTs (iFFT) must be performed.

6.3.3 Cost analysis of circulant preconditioning *versus* time-stepping

Canonically, approximate solutions to (6.27) are recovered via time-stepping, which corresponds to applying block forward substitution to space-time matrices such as (6.37). This is an inherently sequential procedure which, at least for implicit methods, requires inverting a shifted Laplacian at each time step. Assuming inverting the Laplacian operator has an associated cost of $C(N_x)$, the overall procedure has a complexity of $\mathcal{O}(N_t C(N_x))$. Given the availability of optimal solvers for this kind of operators [150], we can expect $C(N_x)$ to scale linearly with N_x , with spatial parallelisation eventually coming to our aid in reducing the required computational time.

Conversely, our procedure dramatically increases the global cost required for solving systems with the coefficient matrix (6.37). For each GMRES iteration, in fact, we need to multiply the space-time residual by the matrix (6.37), an operation which by itself has a complexity of $\mathcal{O}(N_t N_x)$. On top of this, applying the preconditioner involves N_x independent FFT and iFFT, of overall complexity $\mathcal{O}(N_x N_t \log(N_t))$, as well as inverting N_t systems involving a complex-shifted Laplacian. The latter can be safely assumed to have a cost comparable to $C(N_x)$ (in fact, even for this operator optimal solvers are available, for example in [44, 95]), which makes the application of the preconditioner a procedure with an overall complexity of $\mathcal{O}(N_x N_t \log(N_t) + N_t C(N_x))$. Unlike time-stepping, however, solution via GMRES with circulant preconditioning exposes parallelisation along the temporal domain, as stated in Sec. 6.1.1.

This can be exploited to reduce computational time at each step of the GMRES iteration. Firstly, when multiplying by the space-time matrix, if we have N_t processors available, each assigned to a specific instant (or block row in the matrix), then the computational time necessary is effectively reduced from $\mathcal{O}(N_t N_x)$ to $\mathcal{O}(N_x)$, if we assume that time-to-solution is directly proportional to complexity¹. Secondly, when

¹Here we are neglecting the fact that, if each processor contains information only pertaining to a specific instant, a certain amount of communication must occur in order to perform multiplication by the space-time matrix. This overload is however limited, since the space-time matrix has a small band of size $s_1 + s_2$, implying that this exchange of information only involves $s_1 + s_2 - 1$ processors

inverting the complex shifted Laplacians, time similarly reduces from $\mathcal{O}(N_t C(N_x))$ to $\mathcal{O}(C(N_x))$, since again each system can be inverted independently, (and no communication must occur among processors). Thirdly, a similar reasoning can also be followed when applying the Fourier transforms: not only the FFTs on each of the N_x signals are independent, meaning that any additional processor for spatial parallelisation can still be employed to attack the N_x factor in the $\mathcal{O}(N_x N_t \log(N_t))$ complexity, but also parallelisation over time is perfectly feasible. Most of the FFT algorithms available involve a step consisting of operations acting over the whole signal vector (see their detailed description in [13, 22]): the operations on each element are independent and can hence be parallelised (as is done in [84, 146]), ideally dropping computational time from $\mathcal{O}(N_x N_t \log(N_t))$ to $\mathcal{O}(N_x \log(N_t))$.

The actual performance of the parallel FFT depends heavily on how communication among processors is organised, and on the architecture of the system considered [53]: since FFT represents the bottleneck in the application of the circulant preconditioner, any efficient parallel implementation will need to be fine-tuned in order to take this into account. This is however beyond the scope of this work, which rather wishes to focus on the efficacy of the preconditioner when applied to space-time systems.

To summarise, we have on the one hand time-stepping: a sure-fire direct procedure which recovers the solution to a space-time system in a number of operations scaling as $\sim N_t C(N_x)$; on the other, we have an iterative procedure whose computational time *per iteration* scales as $\sim C(N_x) + N_x \log(N_t)$. The difference between the two regimes allows some room for speedup, provided that the number of necessary iterations remains bounded, and particularly that it does not scale with N_x , nor with N_t . Even in the PDE case, we can still invoke Thm. 6.1.1 to secure ourselves from the latter, since it confirms that the maximum theoretical number of iterations required for convergence is independent of the number of time steps taken. Unfortunately, though, Thm. 6.1.1 does not shield us from the former, since in principle the number of iterations could grow with the size of the blocks composing the matrices (6.37), *i.e.*, the number of spatial nodes N_x considered in the discretisation. In practice, however, and particularly under certain regimes, this seems not to be the case. Details on this are shown in the following.

at most, and in particular its cost per processor is independent on N_t .

6.3.4 Results

In this section, we report the results from the application of GMRES for the solution of systems involving (6.37) as coefficient matrix, using the corresponding (6.39) as right-preconditioners. In particular, and in light of the considerations made in the previous paragraph, we are interested in its performance in terms of number of iterations to convergence.

Table 6.2: Number of iterations to convergence for GMRES, right-preconditioned with (6.7), applied to the solution of systems arising from the discretisation of (6.27). Different number of spatial and temporal nodes considered (N_x and N_t , respectively). Time derivative approximated using SV, and space derivative approximated using CD of matching orders of accuracy. The rightmost column in each table reports the theoretical upper bound from Thm. 6.1.1; a cross identifies a simulation which did not converge due to memory requirements becoming too severe. Values for $N_x > N_t$ are not reported, as they violate the CFL condition, and give rise to unstable solutions: convergence results are very poor in those cases. The different colours in the cells background provide an indication of the accuracy of the solution for the test-case considered: the L^2 error in space-time with respect to the analytical solution decreases as we move to darker shades, becoming $< 10^{-4}$, $< 10^{-6}$, and $< 10^{-8}$, respectively.

Time: SV ₁ - Space: CD ₂								Time: SV ₂ - Space: CD ₂							
$N_t \rightarrow$ $N_x \downarrow$	120	240	480	960	1920	3840	(s_1+s_2) $\cdot N_x+1$	$N_t \rightarrow$ $N_x \downarrow$	120	240	480	960	1920	3840	(s_1+s_2) $\cdot N_x+1$
80	15	17	26	39	70	136	161	80	66	71	96	123	125	128	161
160	–	16	19	30	53	95	321	160	–	82	116	196	242	225	321
320	–	–	18	21	33	76	641	320	–	–	83	115	149	406	641
640	–	–	–	19	22	35	1281	640	–	–	–	103	153	144	1281
1280	–	–	–	–	21	23	2561	1280	–	–	–	–	130	204	2561
2560	–	–	–	–	–	24	5121	2560	–	–	–	–	–	159	5121
Time: SV ₃ = SV ₄ - Space: CD ₄								Time: SV ₅ - Space: CD ₆							
$N_t \rightarrow$ $N_x \downarrow$	120	240	480	960	1920	3840	(s_1+s_2) $\cdot N_x+1$	$N_t \rightarrow$ $N_x \downarrow$	120	240	480	960	1920	3840	(s_1+s_2) $\cdot N_x+1$
80	49	67	91	105	114	118	161	80	132	140	157	189	192	199	321
160	–	46	55	81	110	153	321	160	–	71	123	120	143	151	641
320	–	–	30	31	33	46	641	320	–	–	65	183	177	137	1281
640	–	–	–	25	25	27	1281	640	–	–	–	97	388	171	2561
1280	–	–	–	–	27	28	2561	1280	–	–	–	–	167	325	5121
2560	–	–	–	–	–	26	5121	2560	–	–	–	–	–	92	10241
Time: SV ₆ - Space: CD ₆								Time: SV ₇ - Space: CD ₈							
$N_t \rightarrow$ $N_x \downarrow$	120	240	480	960	1920	3840	(s_1+s_2) $\cdot N_x+1$	$N_t \rightarrow$ $N_x \downarrow$	120	240	480	960	1920	3840	(s_1+s_2) $\cdot N_x+1$
80	181	229	205	212	222	225	401	80	213	258	221	242	249	256	481
160	–	248	244	186	169	164	801	160	–	312	390	259	241	176	961
320	–	–	263	348	264	253	1601	320	–	–	480	506	406	270	1921
640	–	–	–	310	465	368	3201	640	–	–	–	844	822	395	3841
1280	–	–	–	–	335	907	6401	1280	–	–	–	–	1573	943	7681
2560	–	–	–	–	–	588	12801	2560	–	–	–	–	–	×	15261

For all of our experiments, GMRES is set to a tolerance of 10^{-10} , starting from a null initial guess. We consider the unit square, $T = L = 1$, as a domain for the PDE

Table 6.3: Number of iterations to convergence for GMRES, right-preconditioned with (6.7), applied to the solution of systems arising from the discretisation of (6.27). Different number of spatial and temporal nodes considered (N_x and N_t , respectively). Time derivatives approximated using CI_2 , space derivatives using CD with various accuracies. The rightmost column in each table reports the theoretical upper bound from Thm. 6.1.1. The different shades indicate the accuracy with respect to the analytical solution as in Tab. 6.2.

Time: CI_2 - Space: CD_2								Time: CI_2 - Space: CD_4							
$\begin{smallmatrix} N_t \rightarrow \\ N_x \downarrow \end{smallmatrix}$	120	240	480	960	1920	3840	$\begin{smallmatrix} (s_1+s_2) \\ \cdot N_x+1 \end{smallmatrix}$	$\begin{smallmatrix} N_t \rightarrow \\ N_x \downarrow \end{smallmatrix}$	120	240	480	960	1920	3840	$\begin{smallmatrix} (s_1+s_2) \\ \cdot N_x+1 \end{smallmatrix}$
80	61	75	111	124	125	128	161	80	64	70	74	90	109	112	161
160	83	92	136	203	243	226	321	160	82	99	87	70	70	94	321
320	116	119	140	146	224	372	641	320	124	117	100	86	70	74	641
640	242	188	152	157	183	137	1281	640	285	187	136	125	115	78	1281
1280	767	403	227	205	221	240	2561	1280	799	497	208	153	161	144	2561
2560	1266	1476	582	338	234	279	5121	2560	1844	1546	587	248	250	162	5121

(6.27). As an initial condition, we choose a shifted Gaussian:

$$\bar{u}^0(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2} \left(x - \frac{L}{2}\right)^2\right) - c, \quad (6.43)$$

where c is picked so that $\bar{u}^0(0) = \bar{u}^0(L) = 0$, and $\sigma^2 = 0.002$; for the derivative, we simply take $\bar{u}^1(x) = 0$. The solutions to the systems involving the complex-shifted Laplacian are recovered at each iteration using the *backslash* command in MATLAB. The code used for the experiments is publicly available in its Git repository [27]. We test the two different schemes presented above, for a variety of orders of accuracy. Results are presented in Tab. 6.2 for SV discretisations of the temporal derivative, and in Tab. 6.3 for CI_2 discretisations.

The actual number of iterations to convergence remains well below the limit indicated by Thm. 6.1.1 for all cases considered. We notice however that this number is still far from being independent from the spatial mesh size. The convergence behaviour when increasing N_t is somewhat hard to grasp, but tends to show an increase in iterations count as N_t gets higher. Increasing N_x has a similar impact, and also slows convergence. This is expected from the statement of Thm. 6.1.1: N_x determines the size of the blocks composing (6.7), and hence directly affects the rank of the perturbation (6.10). Also increasing the order of accuracy of the time discretisations used has a similar negative effect on performance. This is also expected from Thm. 6.1.1: in this case it is the actual number of blocks in (6.10), rather than their dimension, that is directly affected.

The SV_1 scheme (top-left in Tab. 6.2) deserves a separate discussion: this scheme produces a solution which presents a noticeable degree of numerical diffusion, as

reported in the literature [65] and shown in Fig. 6.1. There, we can see how the solution profile with SV_1 (at the bottom) is visibly smeared out, contrasting with the result from leapfrog (top and middle): the difference is clear by comparing the shapes of the final negative peaks in the three surface plots, or equivalently by looking at how the heatmap at the bottom appears more blurred. The fact that diffusion is present seems to aid with the convergence of the scheme, a feature shared with other parallel-in-time integration methods (see [5,35], and the discussion in Sec. 2.1.1.1). In fact, we can see from Tab. 6.2 that the number of iterations to convergence remains reasonably small, even when the meshes are refined. However, the recovered solution is of very poor quality: even with the finest meshes used, the error with respect to the analytical solution remains above 10^{-4} .

The background colours of the cells in Tab. 6.2 draw a clearer picture of the performances of each scheme, by providing an indication of the accuracy of the recovered approximation: moving to darker shades, the numerical solution reaches an approximation error of 10^{-4} , 10^{-6} , and 10^{-8} , respectively, although at around 10^{-8} the convergence of high-order schemes tends to stall, likely because of round-off errors becoming relevant, which makes the results less informative for this last range. The first lesson we can extract from this is that there is no advantage in using a low-order scheme over a high-order one, so long as they share the same stencil: by comparing the tables relative to SV_2 and SV_4 , (two schemes both with a stencil of size 3), we see that SV_4 achieves a much smaller error with respect to SV_2 on the same meshes, and often even with a smaller number of iterations. The latter is possibly due to the negative impact of the extra perturbation introduced to the Toeplitz system when accounting for the initial conditions, since we used a ghost-node approach for SV_2 . As the order of the method increases, it becomes harder to make such strong statements, but if the stencil grows together with the order, then low-order, small-stencil schemes seem to be preferable. For example, SV_4 reaches an error of 10^{-6} with 25 iterations on a 640×960 mesh, while SV_5 needs 65 iterations (roughly 2.5 as many) on a mesh with $1/4$ as many degrees of freedom, and SV_7 uses 312, (or roughly 12.5 as many) iterations on a mesh with only $1/16$ as many degrees of freedom. Given how the cost per iteration scales *linearly* in N_x but *logarithmically* in N_t (as per the analysis conducted in the previous paragraph), SV_5 is of comparable cost with respect to SV_4 , but SV_7 far exceeds it, and is hence not worth employing for the same target accuracy.

The nature of the initial conditions also has an impact on the speed of convergence. We test this by running a series of experiments where we progressively narrow




Figure 6.1: Surface plots and heatmaps of examples of numerical solutions recovered from the experiments. Top: SV_2 discretisation in time, CD_2 in space (explicit Leapfrog); middle: implicit Leapfrog, still with CD_2 space discretisation; bottom: the pair $SV_1 - CD_2$ is used instead. $N_t = 240$, $N_x = 160$, and other parameters as in Tab. 6.2, and Tab. 6.3. In the surface plots, time progresses towards the reader.

the Gaussian used as an initial condition. The number of iterations to convergence increases as the Gaussian becomes steeper, regardless of the discretisation used, as is shown in Tab. 6.4. This is not entirely surprising: with steeper profiles, the range of relevant modes in the solution extends to higher frequencies, and we can expect to

Table 6.4: Number of iterations to convergence for GMRES, right-preconditioned with (6.7), applied to the solution of systems arising from the discretisation of (6.27), with different initial conditions. The first row reports results for a standing wave, while the others refer to Gaussian curves with different values of σ^2 in (6.43). System size: $N_t = 240$, $N_x = 160$. Various discretisations for the temporal derivatives considered; spatial derivatives approximated using CD of matching accuracy.

	SV ₁	SV ₂	SV ₃ = SV ₄	SV ₅	SV ₆	SV ₇	CI ₂
$\bar{u}^0 = \sin(3\pi x/L)$	2	2	47	16	22	49	2
$\sigma^2 = 10^{-3}2^1$	16	82	46	71	248	312	92
$\sigma^2 = 10^{-3}2^0$	17	110	80	113	267	347	107
$\sigma^2 = 10^{-3}2^{-1}$	17	123	103	169	288	373	125
$\sigma^2 = 10^{-3}2^{-2}$	17	135	120	208	302	388	139
$\sigma^2 = 10^{-3}2^{-3}$	17	142	134	228	310	397	146
$\sigma^2 = 10^{-3}2^{-4}$	18	146	139	277	338	400	149
$\sigma^2 = 10^{-3}2^{-5}$	18	148	141	285	395	401	151

require a richer Krylov subspace in order to properly approximate them. An extreme example of the converse can be shown by picking a single mode as an initial condition, of any frequency. This gives rise to a stationary wave, which is particularly simple to track, since the spatial and the temporal components of the solution factorise: the solution at each instant is nothing but the initial condition, opportunely rescaled. In this case, the circulant preconditioner is extremely effective, as it is capable of building a Krylov subspace which captures the degrees of freedom of the solution in very few iterations: only 2, for all schemes of order of accuracy 2 or smaller, which matches some observations in the literature [64,65]. This is however a peculiar feature of the circulant preconditioner, limited to the very specific combination of parameters employed: if we vary the type of spatial discretisation used, (or if we add a forcing term, or vary the initial guess for the GMRES algorithm), we cannot expect such a dramatic convergence anymore, as shown in the first row of Tab. 6.4.

Since the time this research was conducted, an account on the parallel performance of this algorithm has been provided in [93]. This, and the analysis conducted in this chapter, provide evidence for stating that circulant preconditioning is a feasible PinT method for hyperbolic PDEs. The biggest obstacle in its application, however, lies in the scarce generalisability of the scheme, which relies on the Toeplitz structure of the target space-time system. How to overcome this limitation is still largely a matter of research: some suggestions for future work in this regard, as well as for the other topics covered in the scope of this thesis, are reported in the next and final chapter.

Chapter 7

Conclusion

We conclude this thesis by providing a brief summary of the results of our work. This is done in the following section, where for each chapter we indicate the major findings stemming from the analyses conducted therein. Section 7.2 instead lists a number of issues which have not been addressed with our research, together with suggestions for possible solutions, which will be the matter of future research. As a final comment, in Sec. 7.3 we report our personal intuition behind the scope of space-time block preconditioning, the main technique developed in the course of this project.

7.1 Summary of findings

Space-time block preconditioning. In Chap. 3, we introduce a novel approach to block preconditioning for space-time systems. Its purpose is to reduce the solution of a space-time system of PDEs to that of a single-variable space-time PDE, which is a much more tractable problem for PinT solvers. We consider the discretised time-dependent Oseen equations as an example to describe our approach. In designing our space-time block preconditioner, we take inspiration from heuristics that have proven to be successful in the single time-step or steady-state setting, and extend similar principles to the whole space-time formulation. The resulting preconditioning procedure involves two operations: the space-time solution of a single-variable, time-dependent advection-diffusion equation, and a (spatial) pressure mass- and stiffness-matrix inverse at each time point. The latter operation is trivially parallelisable over the time-steps, so that the computational time it requires is comparable to its analogue for sequential time-stepping. For the former, efficient time-parallelisation techniques are readily available. We demonstrate the effectiveness of our preconditioner by applying it to a variety of model problems taken from the literature, which cover a range of flow configurations. The main measure for performance is given by the number

of iterations required for convergence, which we show to be scalable in spatial and temporal mesh size, and comparable to those required for a classical time-stepping (non-time-parallel) routine. We investigate a simple extension to a nonlinear Navier-Stokes problem as well: this is solved via Picard iterations, and even in this case we observe near perfect scalability, for both linear and nonlinear solvers.

Applications to incompressible magnetohydrodynamics. In Chap. 4, we extend the work from the previous chapter, and design a space-time block preconditioner for a more complicated, nonlinear problem in magnetohydrodynamics. This gives evidence for the flexibility of our STBP approach, paving the way for its application to other systems of PDEs. The system considered represents the main target application of our thesis, and is of particular relevance for our industrial partner, due to the fundamental role it plays in modelling the behaviour of plasma within a fusion reactor. Following the principles outlined in the previous chapter, we first identify a preconditioner which has shown its effectiveness in the time-stepping framework, and then proceed to extend it to the whole space-time setting. The resulting space-time block preconditioner acts by addressing two simplified subproblems in sequence: it solves first a system coupling magnetic field and velocity, followed by one coupling velocity and pressure. The latter is equivalent to the one tackled in Chap. 3, while the former requires inverting the discretisation of a single-variable, high-order space-time differential operator, analogous to the one arising in the sequential time-stepping setting. How to address the parallel-in-time solution of this operator remains a matter of future research, and to keep our work contained we focus on evaluating the performance of the space-time block preconditioner when using exact solvers. This is done by measuring iterations to convergence for both the linear and nonlinear solvers. Applying the space-time block preconditioner requires limited overhead with respect to its sequential time-stepping counterpart, and it shows excellent scalability with respect to spatial refinement and only marginal convergence degradation with temporal refinement.

Multigrid Reduction in Time for nonlinear hyperbolic equations. With Chap. 5, our focus moves from the time-parallel solution of systems of PDEs to that of single-variable PDEs. In particular, we consider the performance of PinT methods when applied to hyperbolic equations. In this chapter, the algorithm of choice is MGRIT, in light of its maturity and the flexibility it offers as a multigrid method, and we investigate its applicability to the solution of nonlinear conservation laws.

The aim is to understand how the choice of integrators used at each level of the multigrid algorithm impacts its convergence behaviour. To this purpose, we measure the performance of MGRIT applied using a combination of existing integrators, commonly used for the solution of conservation laws. From the results we can infer that dissipative schemes behave better in general, which is in line with the literature; however, higher-order methods for spatial reconstruction coupled with matching order time-discretisations still provide satisfactory convergence results, under suitable CFL limits. We note the importance of choosing coarse integrators that closely follow the action of the fine integrator, by showcasing the effectiveness of a new method proposed for the construction of accurate coarse solvers. This approach seeks to directly approximate the action of the fine solver: even though it comes at an increased cost with respect to simple rediscritisation, it offers superior performance. Its range of applicability remains, however, limited. In all cases, in fact, the performance is seen to degrade fairly quickly as the CFL number increases. This seems to be a strong limitation for the application of MGRIT to hyperbolic systems directly discretised with explicit time-steppers. The level of coarsening that can be applied to the temporal grid in this case is, thus, effectively capped. This suggests that MGRIT should be paired with a spatial coarsening strategy as well, in order to control the CFL number adequately across all levels.

Space-time circulant preconditioning for linear wave equations. Chapter 6 revolves around the investigation of a recently suggested method for the all-at-once solution of space-time systems arising from the discretisation of linear, constant-coefficient PDEs. We explore this approach as an alternative to MGRIT for the solution of hyperbolic equations, since preliminary results showed it not to suffer from the same stability limitations when applied to this type of PDEs. Under the assumptions considered, the resulting monolithic system presents a block Toeplitz structure. This is exploited to design a block circulant preconditioner for the acceleration of the iterative solution of the whole space-time system. Given that circulant operators are diagonalisable via FFT, the preconditioner can be inverted in parallel with relative ease, making this approach a *de facto* time-parallel method. We first prove the effectiveness of the resulting scheme theoretically, applying a result on the convergence of circulant-preconditioned GMRES, which guarantees that convergence is reached in a number of iterations independent on the number of temporal nodes. We then measure its performance with numerical experiments, where the preconditioner is applied to the solution of linear wave equations, discretised using high-order

schemes. The results support the theoretical findings, and testify on the effectiveness of this preconditioning strategy.

7.2 Future work

Measurement of parallel performance. The first extension of relevance to the work conducted in this thesis consists in designing an efficient parallel implementation of the space-time block preconditioners introduced in Chap. 3 and 4, in order to actually measure the parallel performance of the strategies proposed here. For the largest part, in our work we restrict ourselves to discussing the *ideal* performance of these preconditioners, focusing on giving theoretical (rather than experimental) justifications for their potential in providing parallel speedup. This choice is on the one hand dictated by the fact that the methods proposed are either relatively recent and still subject of research, or entirely novel, so it is reasonable to first test their effectiveness in a best-case scenario. On the other, including the implementation and design of efficient parallel code would have been hardly feasible in the timescale allocated for this project. Nonetheless, giving actual measurements of the time-to-solution required by these methods would be valuable in consolidating the theoretical results on their viability. In particular, for the space-time block preconditioners, this would help in gaining additional insight on how their performance depends on the choice of their internal PinT components, since in Tab. 3.1 we give only one such example. This is particularly true for the operator (4.71) appearing in the space-time magnetic Schur complement approximation (4.73), as its parallelisation in time, to our knowledge, is not covered in the literature.

Stability of STBP for advection-dominated regimes. In most of our experiments in Chap. 3 and 4, we choose the parameters for viscosity μ in (3.1) and (4.1), and of electric resistivity η and magnetic permeability μ_0 in (4.1), all to be $= 1$, thus steering away from the more challenging advection-dominated, turbulent regimes for these problems. This is justified by similar considerations as those for the previous paragraph, as well as by the fact that the single-time step PCD operator (3.38) our space-time block preconditioners rely on was not originally designed with robustness with respect to μ in mind (indeed we remark on this in Sec. 3.4.2.2). In light of the discussion in Sec. 1.2, though, the advection-dominated regime is of particular relevance for the development of fusion reactors. Fortunately, block preconditioners whose performance remains stable with respect to the problem parameters are

available in the literature: see for example [88]. This is originally designed in the single time-step framework, so the question remains open whether similar ideas can be extended to the whole space-time setting.

MGRIT with semi-implicit Lagrangian methods for hyperbolic equations.

In Chap. 5 we investigate the application of MGRIT for the time parallelisation of the solution of conservation laws discretised with high-order explicit time-integration methods, and in Sec. 5.4 we remark on how one of the biggest obstacles to the stability of the algorithm is given by the need to satisfy the CFL condition at all coarsening levels. One way to circumvent this, without renouncing the computational advantages of employing explicit methods, is given by using *semi-implicit Lagrangian integrators* [80, 117]. By approximately backtracking the characteristics in the system and including a correction term in the flux definition, these schemes have the advantage of easing the restriction imposed by the CFL condition. Employing these solvers at the lower coarsening levels would allow for more leeway in considering larger time-steps, thus possibly improving the overall stability of MGRIT. A similar setup, showing promising results, is investigated for Parareal in [132].

Circulant preconditioning for nonlinear PDEs. The most cumbersome limitation behind the use of the space-time circulant preconditioner introduced in Chap. 6 lies in its reduced scope for applicability. If the block Toeplitz structure of the target space-time matrix is severely broken (be it by choosing uneven temporal grids, by introducing some temporal variation into the PDE coefficients, or by considering nonlinearities), two problems arise: on the one hand, it is unclear how one should extract an appropriate circulant preconditioner from a non-Toeplitz matrix; on the other, as hinted in Rmk. 3, the theoretical result from Thm. 6.1.1 ensuring the effectiveness of the circulant preconditioner does not hold any more. A workaround for us to still make use of the circulant preconditioner, consists in designing an effective projection operator, which allows us to recover a block-Toeplitz approximation of the target space-time system. Possibly the most straightforward way to achieve this is given by diagonally averaging the entries in the monolithic system (which corresponds to a projection in the Frobenius norm), and indeed some preliminary results on a similar approach applied to equations with time-varying coefficients are reported in [93]. The next step would consist in investigating the effectiveness of this method in the framework of nonlinear equations.

7.3 Final remarks

In our opinion, the main contribution of this research work is represented by the space-time block preconditioning approach outlined in Chap. 3 and 4. In addition to its potential for parallel efficiency, which we showcase in applications to problems in fluid dynamics and magnetohydrodynamics in the corresponding chapters, we believe that its major strength lies in the flexibility that such approach provides in dealing with time-dependent systems of PDEs.

Indeed, STBP offers a fundamentally different take on time parallelisation for multi-variable PDEs: rather than focusing on designing *ad hoc* PinT solvers for the target system, this strategy allows to link the solution of such system to that of a (hopefully simpler to parallelise) single-variable, time-dependent PDE, represented by the approximation of a space-time Schur complement. As such, it takes what is accomplished by the block preconditioning strategy commonly used for *spatial* operators, and naturally extends it to the framework of *space-time* operators.

This space-time block preconditioning approach is by no means flawless, and its worth remains in large part to be proven: the work conducted in the scope of this thesis is, after all, introductory at most. However, if the comparison with its spatial counterpart holds, and the development of efficient PinT methods for single-variable PDEs continues to progress, it is not unreasonable to believe that it can become a valid tool in the arsenal of time-parallelisation techniques.

Bibliography

- [1] J. H. ADLER, T. BENSON, E. C. CYR, P. E. FARRELL, S. MACLACHLAN, AND R. TUMINARO, *Monolithic multigrid for magnetohydrodynamics*, ArXiv preprint, arXiv:2006.15700 [math.NA], (2020), <https://arxiv.org/abs/2006.15700>.
- [2] H. ALFVÉN, *Existence of electromagnetic-hydrodynamic waves*, *Nature*, 150 (1942), pp. 405–406, <https://doi.org/10.1038/150405d0>.
- [3] G. ANTONUCCI, *Iterative solution of evolutionary PDEs via all-at-once methods*, master’s thesis, University of Oxford, 2017.
- [4] D. N. ARNOLD, *Finite Element Exterior Calculus*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2018, <https://doi.org/10.1137/1.9781611975543>.
- [5] G. BAL, *On the convergence and the stability of the Parareal algorithm to solve partial differential equations*, in *Domain Decomposition Methods in Science and Engineering*, T. J. Barth, M. Griebel, D. E. Keyes, R. M. Nieminen, D. Roose, T. Schlick, R. Kornhuber, R. Hoppe, J. Périaux, O. Pironneau, O. Widlund, and J. Xu, eds., *Lecture Notes in Computational Science and Engineering*, Berlin, Heidelberg, 2005, Springer Berlin Heidelberg, pp. 425–432.
- [6] G. BAL AND Y. MADAY, *A “Parareal” time discretization for non-linear PDE’s with application to the pricing of an American put*, in *Recent Developments in Domain Decomposition Methods*, L. F. Pavarino and A. Toselli, eds., *Lecture Notes in Computational Science and Engineering*, Berlin, Heidelberg, 2002, Springer Berlin Heidelberg, pp. 189–202, https://doi.org/10.1007%2F978-3-642-56118-4_12.
- [7] S. BALAY, S. ABHYANKAR, M. F. ADAMS, J. BROWN, P. BRUNE, K. BUSCHELMAN, L. DALCIN, A. DENER, V. EIJKHOUT, W. D. GROPP,

- D. KARPEYEV, D. KAUSHIK, M. G. KNEPLEY, D. A. MAY, L. C. MCINNES, R. T. MILLS, T. MUNSON, K. RUPP, P. SANAN, B. F. SMITH, S. ZAMPINI, H. ZHANG, AND H. ZHANG, *PETSc Web page*. <https://www.mcs.anl.gov/petsc>, 2019, <https://www.mcs.anl.gov/petsc>. Software.
- [8] A.-M. BAUDRON, J.-J. LAUTARD, Y. MADAY, AND O. MULA, *The parareal in time algorithm applied to the kinetic neutron diffusion equation*, in Domain Decomposition Methods in Science and Engineering XXI, J. Erhel, M. J. Gander, L. Halpern, G. Pichot, T. Sassi, and O. Widlund, eds., Cham, 2014, Springer International Publishing, pp. 437–445.
- [9] D. BISKAMP, *Magnetic Reconnection in Plasmas*, Cambridge Monographs on Plasma Physics, Cambridge University Press, UK, 2000, <https://doi.org/10.1017/CB09780511599958>.
- [10] M. BOLTEN, D. MOSER, AND R. SPECK, *A multigrid perspective on the parallel full approximation scheme in space and time*, Numerical Linear Algebra with Applications, 24 (2016).
- [11] M. BOLTEN, D. MOSER, AND R. SPECK, *Asymptotic convergence of the parallel full approximation scheme in space and time for linear problems*, Numerical Linear Algebra with Applications, 25 (2018), p. e2208, <https://doi.org/10.1002/nla.2208>.
- [12] W. BRIGGS, V. HENSON, AND S. MCCORMICK, *A Multigrid Tutorial*, Other Titles in Applied Mathematics, Society for Industrial and Applied Mathematics, 2nd ed., 2000.
- [13] C. BURRUS, M. FRIGO, AND G. JOHNSON, *Fast Fourier Transforms*, Samurai Media Limited, 2018.
- [14] J. CAHOUE ET AND J.-P. CHABARD, *Some fast 3d finite element solvers for the generalized Stokes problem*, International Journal for Numerical Methods in Fluids, 8 (1988), pp. 869–895, <https://doi.org/10.1002/flid.1650080802>.
- [15] L. CHACÓN, D. KNOLL, AND J. FINN, *An implicit, nonlinear reduced resistive MHD solver*, Journal of Computational Physics, 178 (2002), pp. 15–36, <https://doi.org/10.1006/jcph.2002.7015>.

- [16] R. H. CHAN, D. POTTS, AND G. STEIDL, *Preconditioners for non-Hermitian Toeplitz systems*, Numerical Linear Algebra with Applications, 8 (2001), pp. 83–98, [https://doi.org/10.1002/1099-1506\(200103\)8:2<83::AID-NLA231>3.0.CO;2-X](https://doi.org/10.1002/1099-1506(200103)8:2<83::AID-NLA231>3.0.CO;2-X).
- [17] R. H.-F. CHAN AND X.-Q. JIN, *An Introduction to Iterative Toeplitz Solvers*, Society for Industrial and Applied Mathematics, 2007, <https://doi.org/10.1137/1.9780898718850>.
- [18] F. CHEN, J. S. HESTHAVEN, Y. MADAY, AND A. S. NIELSEN, *An adjoint approach for stabilizing the parareal method*. 2015, <http://infoscience.epfl.ch/record/211097>.
- [19] F. CHEN, J. S. HESTHAVEN, AND X. ZHU, *On the Use of Reduced Basis Methods to Accelerate and Stabilize the Parareal Method*, Springer International Publishing, Cham, 2014, pp. 187–214, https://doi.org/10.1007/978-3-319-02090-7_7.
- [20] A. T. CLARKE, C. J. DAVIES, D. RUPRECHT, AND S. M. TOBIAS, *Parallel-in-time integration of kinematic dynamos*, Journal of Computational Physics: X, 7 (2020), p. 100057, <https://doi.org/10.1016/j.jcpx.2020.100057>.
- [21] R. CODINA AND N. HERNÁNDEZ-SILVA, *Stabilized finite element approximation of the stationary magneto-hydrodynamics equations*, Computational Mechanics, 38 (2006), pp. 344–355, <https://doi.org/10.1007/s00466-006-0037-x>.
- [22] J. COOLEY AND J. W. TUKEY, *An algorithm for the machine calculation of complex Fourier series*, Mathematics of Computation, 19 (1965), pp. 297–301, <https://doi.org/10.1090/S0025-5718-1965-0178586-1>.
- [23] R. CROCE, D. RUPRECHT, AND R. KRAUSE, *Parallel-in-space-and-time simulation of the three-dimensional, unsteady Navier-Stokes equations for incompressible flow*, in Modeling, Simulation and Optimization of Complex Processes-HPSC 2012, Springer, 2014, pp. 13–23, https://doi.org/10.1007/978-3-319-09063-4_2.

- [24] E. C. CYR, J. N. SHADID, AND R. S. TUMINARO, *Stabilization and scalable block preconditioning for the Navier-Stokes equations*, Journal of Computational Physics, 231 (2012), pp. 345–363, <https://doi.org/10.1016/j.jcp.2011.09.001>.
- [25] E. C. CYR, J. N. SHADID, R. S. TUMINARO, R. P. PAWLOWSKI, AND L. CHACÓN, *A new approximate block factorization preconditioner for two-dimensional incompressible (reduced) resistive MHD*, SIAM Journal on Scientific Computing, 35 (2013), pp. B701–B730, <https://doi.org/10.1137/12088879X>.
- [26] X. DAI AND Y. MADAY, *Stable parareal in time method for first- and second-order hyperbolic systems*, SIAM Journal on Scientific Computing, 35 (2013), pp. A52–A78, <https://doi.org/10.1137/110861002>.
- [27] F. DANIELI, *All-at-once_code*. https://gitlab.com/fdanieli/all-at-once_code. Software repository.
- [28] F. DANIELI, *MGRIT_Non-linear_hyperbolic*. https://gitlab.com/fdanieli/mgrit_non-linear_hyperbolic. Software repository.
- [29] F. DANIELI, *STBP-incompressible-flow*. <https://gitlab.com/fdanieli/stbp-incompressible-flow>, 2020. Software repository.
- [30] F. DANIELI AND S. MACLACHLAN, *Multigrid reduction in time for non-linear hyperbolic equations*, ArXiv preprint [arXiv:2104.09404](https://arxiv.org/abs/2104.09404) [math.NA], (2021), <https://arxiv.org/abs/2104.09404>.
- [31] F. DANIELI, B. S. SOUTHWORTH, AND A. J. WATHEN, *Space-time block preconditioning for incompressible flow*, ArXiv preprint [arXiv:2101.07003](https://arxiv.org/abs/2101.07003) [math.NA], (2021), <https://arxiv.org/abs/2101.07003>.
- [32] F. DANIELI AND A. J. WATHEN, *All-at-once solution of linear wave equations*, Numerical Linear Algebra with Applications, p. e2386, <https://doi.org/10.1002/nla.2386>.
- [33] H. DE STERCK, R. D. FALGOUT, S. FRIEDHOFF, O. A. KRZYSIK, AND S. P. MACLACHLAN, *Optimizing multigrid reduction-in-time and Parareal coarse-grid operators for linear advection*, Numerical Linear Algebra with Applications, (2021), <https://doi.org/10.1002/nla.2367>.

- [34] J. DENNIS AND R. SCHNABEL, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Society for Industrial and Applied Mathematics, 1996, <https://doi.org/10.1137/1.9781611971200>.
- [35] V. DOBREV, T. KOLEV, N. PETERSSON, AND J. SCHRODER, *Two-level convergence theory for multigrid reduction in time (MGRIT)*, *SIAM Journal on Scientific Computing*, 39 (2017), pp. S501–S527, <https://doi.org/10.1137/16M1074096>.
- [36] J. DONGARRA AND P. LUSZCZEK, *TOP500*, Springer US, Boston, MA, 2011, pp. 2055–2057, https://doi.org/10.1007/978-0-387-09766-4_157.
- [37] A. DUTT, L. GREENGARD, AND V. ROKHLIN, *Spectral deferred correction methods for ordinary differential equations*, *BIT Numerical Mathematics*, 40 (2000), pp. 241–266, <https://doi.org/10.1023/A:1022338906936>.
- [38] A. EGHBAL, A. G. GERBER, AND E. AUBANEL, *Acceleration of unsteady hydrodynamic simulations using the parareal algorithm*, *Journal of Computational Science*, 19 (2017), pp. 57 – 76, <https://doi.org/10.1016/j.jocs.2016.12.006>.
- [39] H. ELMAN, V. HOWLE, J. SHADID, R. SHUTTLEWORTH, AND R. TUMINARO, *A taxonomy and comparison of parallel block multi-level preconditioners for the incompressible Navier-Stokes equations*, *Journal of Computational Physics*, 227 (2008), pp. 1790–1808, <https://doi.org/10.1016/j.jcp.2007.09.026>.
- [40] H. ELMAN, D. SILVESTER, AND A. WATHEN, *Finite Elements and Fast Iterative Solvers: with Applications in Incompressible Fluid Dynamics*, Numerical Mathematics and Scientific Computation, Oxford University Press, UK, 2014, <https://doi.org/10.1093/acprof:oso/9780199678792.001.0001>.
- [41] H. ELMAN AND R. TUMINARO, *Boundary conditions in approximate commutator preconditioners for the Navier-Stokes equations*, *Electronic Transactions on Numerical Analysis*, 35 (2009), pp. 257–280.
- [42] H. C. ELMAN, A. RAMAGE, AND D. J. SILVESTER, *IFISS: A computational laboratory for investigating incompressible flow problems*, *SIAM Review*, 56 (2014), pp. 261–273, <https://doi.org/10.1137/120891393>.

- [43] M. EMMETT AND M. L. MINION, *Toward an efficient parallel in time method for partial differential equations*, Communications in Applied Mathematics and Computational Science, 7 (2012), pp. 105–132, <https://doi.org/10.2140/camcos.2012.7.105>.
- [44] Y. A. ERLANGGA, C. W. OOSTERLEE, AND C. VUIK, *A novel multigrid based preconditioner for heterogeneous Helmholtz problems*, SIAM Journal on Scientific Computing, 27 (2006), pp. 1471–1492, <https://doi.org/10.1137/040615195>.
- [45] V. M. FADEEV, I. F. KVABTSKHAVA, AND N. N. KOMAROV, *Self-focusing of local plasma currents*, Nuclear Fusion, 5 (1965), pp. 202–209, <https://doi.org/10.1088/0029-5515/5/3/003>.
- [46] R. FALGOUT, T. MANTEUFFEL, B. O’NEILL, AND J. SCHRODER, *Multigrid reduction in time for nonlinear parabolic problems: A case study*, SIAM Journal on Scientific Computing, 39 (2017), pp. S298–S322, <https://doi.org/10.1137/16M1082330>.
- [47] R. D. FALGOUT, S. FRIEDHOFF, T. V. KOLEV, S. P. MACLACHLAN, AND J. B. SCHRODER, *Parallel time integration with multigrid*, SIAM Journal on Scientific Computing, 36 (2014), pp. C635–C661, <https://doi.org/10.1137/130944230>.
- [48] R. D. FALGOUT, S. FRIEDHOFF, T. V. KOLEV, S. P. MACLACHLAN, J. B. SCHRODER, AND S. VANDEWALLE, *Multigrid methods with space-time concurrency*, Computing and Visualization in Science, 18 (2017), pp. 123–143, <https://doi.org/10.1007/s00791-017-0283-9>.
- [49] R. D. FALGOUT, M. LECOUEZ, AND C. S. WOODWARD, *A parallel-in-time algorithm for variable step multistep methods*, in LLNL Technical Report, 2017.
- [50] P. FISCHER, F. HECHT, AND Y. MADAY, *A parareal in time semi-implicit approximation of the Navier-Stokes equations*, vol. 40, 01 2005, pp. 433–440, https://doi.org/10.1007/3-540-26825-1_44.
- [51] H. FOERSTER, K. STÜBEN, AND U. TROTTENBERG, *Non-standard multigrid techniques using checkered relaxation and intermediate grids*, in Elliptic Problem Solvers, M. H. Schultz, ed., Academic Press, 1981, pp. 285–300, <https://doi.org/10.1016/B978-0-12-632620-8.50027-9>.

- [52] B. FORNBERG, *Generation of finite difference formulas on arbitrarily spaced grids*, Mathematics of Computation, 51 (1988), pp. 699–706, <https://doi.org/10.2307/2008770>.
- [53] M. FRIGO AND S. G. JOHNSON, *The design and implementation of FFTW3*, Proceedings of the IEEE, 93 (2005), pp. 216–231, <https://doi.org/10.1109/JPROC.2004.840301>.
- [54] M. J. GANDER, *Analysis of the Parareal algorithm applied to hyperbolic problems using characteristics*, Boletín de la Sociedad Española de Matemática Aplicada, 42 (2008), pp. 21–35.
- [55] M. J. GANDER, *50 years of time parallel time integration*, in Multiple Shooting and Time Domain Decomposition Methods, T. Carraro, M. Geiger, S. Körkel, and R. Rannacher, eds., Cham, 2015, Springer International Publishing, pp. 69–113, https://doi.org/10.1007/978-3-319-23321-5_3.
- [56] M. J. GANDER AND S. GÜTTEL, *ParaExp: A parallel integrator for linear initial-value problems*, SIAM Journal on Scientific Computing, 35 (2013), pp. C123–C142, <https://doi.org/10.1137/110856137>.
- [57] M. J. GANDER, S. GÜTTEL, AND M. PETCU, *A nonlinear ParaExp algorithm*, in Domain Decomposition Methods in Science and Engineering XXIV, P. E. Bjørstad, S. C. Brenner, L. Halpern, H. H. Kim, R. Kornhuber, T. Rahman, and O. B. Widlund, eds., Cham, 2018, Springer International Publishing, pp. 261–270.
- [58] M. J. GANDER AND L. HALPERN, *Time parallelization for nonlinear problems based on diagonalization*, in Domain Decomposition Methods in Science and Engineering XXIII, C.-O. Lee, X.-C. Cai, D. E. Keyes, H. H. Kim, A. Klawonn, E.-J. Park, and O. B. Widlund, eds., Cham, 2017, Springer International Publishing, pp. 163–170.
- [59] M. J. GANDER, L. HALPERN, J. RYAN, AND T. T. B. TRAN, *A direct solver for time parallelization*, in Domain Decomposition Methods in Science and Engineering XXII, T. Dickopf, M. J. Gander, L. Halpern, R. Krause, and L. F. Pavarino, eds., Cham, 2016, Springer International Publishing, pp. 491–499.

- [60] M. J. GANDER, J. LIU, S.-L. WU, X. YUE, AND T. ZHOU, *Para-Diag: parallel-in-time algorithms based on the diagonalization technique*, ArXiv preprint arXiv:2005.09158 [math.NA], (2021), <https://arxiv.org/abs/2005.09158>.
- [61] M. J. GANDER AND M. NEUMÜLLER, *Analysis of a new space-time parallel multigrid algorithm for parabolic problems*, SIAM Journal on Scientific Computing, 38 (2016), pp. A2173–A2208, <https://doi.org/10.1137/15M1046605>.
- [62] M. J. GANDER AND S. VANDEWALLE, *Analysis of the Parareal time-parallel time-integration method*, SIAM Journal on Scientific Computing, 29 (2007), pp. 556–578, <https://doi.org/10.1137/05064607X>.
- [63] GANDER, M. AND PETCU, M., *Analysis of a Krylov subspace enhanced parareal algorithm for linear problems*, ESAIM Proceedings, 25 (2008), pp. 114–129, <https://doi.org/10.1051/proc:082508>.
- [64] A. GODDARD, *An implementation of the all-at-once method to evolutionary PDEs*, master’s thesis, University of Oxford, 2018.
- [65] A. GODDARD AND A. WATHEN, *A note on parallel preconditioning for all-at-once evolutionary PDEs*, Electronic Transactions on Numerical Analysis, 51 (2019), pp. 135–150.
- [66] R. J. GOLDSTON AND P. H. RUTHERFORD, *Introduction to plasma physics*, Institute of Physics Publishing, Bristol, 1995.
- [67] S. GOTTLIEB, C. SHU, AND E. TADMOR, *Strong stability-preserving high-order time discretization methods*, SIAM Review, 43 (2001), pp. 89–112, <https://doi.org/10.1137/S003614450036757X>.
- [68] A. GREENBAUM, V. PTÁK, AND Z. STRAKOŠ, *Any nonincreasing convergence curve is possible for GMRES*, SIAM Journal on Matrix Analysis and Applications, 17 (1996), pp. 465–469, <https://doi.org/10.1137/S0895479894275030>.
- [69] S. GÜNTHER, N. GAUGER, AND J. SCHODER, *A non-intrusive parallel-in-time adjoint solver with the XBraid library*, Computing and Visualization in Science, 19 (2018), <https://doi.org/10.1007/s00791-018-0300-7>.

- [70] E. HAIRER, S. NØRSETT, AND G. WANNER, *Solving Ordinary Differential Equations I: Nonstiff Problems*, Springer Series in Computational Mathematics, Springer Berlin Heidelberg, 2008, <https://doi.org/10.1007/978-3-540-78862-1>.
- [71] E. HAIRER AND G. WANNER, *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*, Springer Series in Computational Mathematics, Springer Berlin Heidelberg, 2010, <https://doi.org/10.1007/978-3-642-05221-7>.
- [72] A. HARTEN AND J. M. HYMAN, *Self adjusting grid methods for one-dimensional hyperbolic conservation laws*, Journal of Computational Physics, 50 (1983), pp. 235 – 269, [https://doi.org/10.1016/0021-9991\(83\)90066-9](https://doi.org/10.1016/0021-9991(83)90066-9).
- [73] T. HAUT AND B. WINGATE, *An asymptotic parallel-in-time method for highly oscillatory PDEs*, SIAM Journal on Scientific Computing, 36 (2014), pp. A693–A713, <https://doi.org/10.1137/130914577>.
- [74] A. HESSENTHALER, R. D. FALGOUT, J. B. SCHRODER, A. DE VECCHI, D. NORDSLETTEN, AND O. RÖHRLE, *Time-periodic steady-state solution of fluid-structure interaction and cardiac flow problems through multigrid-reduction-in-time*, ArXiv preprint, arXiv:2105.00305 [cs.CE], (2021), <https://arxiv.org/abs/2105.00305>.
- [75] J. HESTHAVEN, *Numerical Methods for Conservation Laws*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017, <https://doi.org/10.1137/1.9781611975109>.
- [76] M. HOCHBRUCK AND A. OSTERMANN, *Exponential integrators*, Acta Numerica, 19 (2010), p. 209–286, <https://doi.org/10.1017/S0962492910000048>.
- [77] G. HORTON AND S. VANDEWALLE, *A space-time multigrid method for parabolic partial differential equations*, SIAM Journal on Scientific Computing, 16 (1995), pp. 848–864, <https://doi.org/10.1137/0916050>.
- [78] A. HOWSE, H. STERCK, R. FALGOUT, S. MACLACHLAN, AND J. SCHRODER, *Parallel-in-time multigrid with adaptive spatial coarsening for the linear advection and inviscid burgers equations*, SIAM Journal on Scientific Computing, 41 (2019), pp. A538–A565, <https://doi.org/10.1137/17M1144982>.

- [79] K. HU. Private communication.
- [80] C.-S. HUANG, T. ARBOGAST, AND C.-H. HUNG, *A semi-Lagrangian finite difference WENO scheme for scalar nonlinear conservation laws*, Journal of Computational Physics, 322 (2016), pp. 559–585, <https://doi.org/10.1016/j.jcp.2016.06.027>.
- [81] *HYPRE: Scalable linear solvers and multigrid methods*. www.llnl.gov/casc/hypre/. Software.
- [82] *IFISS: Incompressible flow and iterative solver software*. <https://personalpages.manchester.ac.uk/staff/david.silvester/ifiss/default.htm>. Software.
- [83] J. D. JACKSON, *Classical electrodynamics*, Wiley, New York, NY, 3rd ed., 1999.
- [84] L. JAMIESON, P. T. MUELLER, AND H. SIEGEL, *FFT algorithms for SIMD parallel processing systems*, Journal of Parallel and Distributed Computing, 3 (1986), pp. 48–71, [https://doi.org/10.1016/0743-7315\(86\)90027-4](https://doi.org/10.1016/0743-7315(86)90027-4).
- [85] D. KAY, D. LOGHIN, AND A. WATHEN, *A preconditioner for the steady-state Navier-Stokes equations*, SIAM Journal on Scientific Computing, 24 (2002), pp. 237–256, <https://doi.org/10.1137/S106482759935808X>.
- [86] A. KREIENBUEHL, A. NAEGEL, D. RUPRECHT, R. SPECK, G. WITTUM, AND R. KRAUSE, *Numerical simulation of skin transport using parareal*, Computing and Visualization in Science, 17 (2015), pp. 99–108, <https://doi.org/10.1007/s00791-015-0246-y>.
- [87] I. KULCHYTSKA-RUCHKA AND S. SCHÖPS, *Efficient parallel-in-time solution of time-periodic problems using a multiharmonic coarse grid correction*, SIAM Journal on Scientific Computing, 43 (2021), pp. C61–C88, <https://doi.org/10.1137/20M1314756>.
- [88] F. LAAKMANN, P. E. FARRELL, AND L. MITCHELL, *An augmented lagrangian preconditioner for the magnetohydrodynamics equations at high Reynolds and coupling numbers*, ArXiv preprint, [arXiv:2104.14855 \[math.NA\]](https://arxiv.org/abs/2104.14855), (2021), <https://arxiv.org/abs/2104.14855>.

- [89] R. LEVEQUE, *Numerical Methods for Conservation Laws*, Lectures in Mathematics ETH Zürich, Department of Mathematics Research Institute of Mathematics, Springer, 1992.
- [90] R. J. LEVEQUE, *Finite Difference Methods for Ordinary and Partial Differential Equations*, Society for Industrial and Applied Mathematics, 2007, <https://doi.org/10.1137/1.9780898717839>.
- [91] A. LIMACHE, S. IDELSOHN, R. ROSSI, AND E. OÑATE, *The violation of objectivity in Laplace formulations of the Navier-Stokes equations*, International Journal for Numerical Methods in Fluids, 54 (2007), pp. 639–664, <https://doi.org/10.1002/flid.1480>.
- [92] J.-L. LIONS, Y. MADAY, AND G. TURINICI, *Résolution d'EDP par un schéma en temps $\langle\langle$ pararéel $\rangle\rangle$* , Comptes Rendus de l'Académie des Sciences - Series I - Mathematics, 332 (2001), pp. 661 – 668, [https://doi.org/10.1016/S0764-4442\(00\)01793-6](https://doi.org/10.1016/S0764-4442(00)01793-6).
- [93] J. LIU AND S.-L. WU, *A fast block α -circulant preconditioner for all-at-once system from wave equations*, SIAM Journal on Matrix Analysis and Applications, (2020), <https://doi.org/10.1137/19M1309869>.
- [94] X.-D. LIU, S. OSHER, AND T. CHAN, *Weighted essentially non-oscillatory schemes*, Journal of Computational Physics, 115 (1994), pp. 200 – 212, <https://doi.org/10.1006/jcph.1994.1187>.
- [95] S. MACLACHLAN AND K. OOSTERLEE, *Algebraic multigrid solvers for complex-valued matrices*, SIAM Journal on Scientific Computing, 30 (2008), pp. 1548–1571, <https://doi.org/10.1137/070687232>.
- [96] Y. MADAY, *The parareal in time algorithm*, in Substructuring Techniques and Domain Decomposition Methods, F. Magoulès, ed., Stirlingshire, UK, 2010, Saxe-Coburg Publications, pp. 19–44, <https://doi.org/10.4203/csets.24.2>.
- [97] Y. MADAY AND E. M. RØNQUIST, *Parallelization in time through tensor-product space-time solvers*, Comptes Rendus Mathématique, 346 (2008), pp. 113 – 118, <https://doi.org/10.1016/j.crma.2007.09.012>.

- [98] S. MAHMOUDKABER AND Y. MADAY, *Parareal in time approximation of the Korteweg-deVries-Burgers' equations*, PAMM, 7 (2007), pp. 1026403–1026404, <https://doi.org/10.1002/pamm.200700574>.
- [99] T. A. MANTEUFFEL, S. MÜNZENMAIER, J. RUGE, AND B. SOUTHWORTH, *Nonsymmetric reduction-based algebraic multigrid*, SIAM Journal on Scientific Computing, 41 (2019), pp. S242–S268, <https://doi.org/10.1137/18M1193761>.
- [100] T. A. MANTEUFFEL, J. RUGE, AND B. S. SOUTHWORTH, *Nonsymmetric algebraic multigrid based on local approximate ideal restriction (ℓ AIR)*, SIAM Journal on Scientific Computing, 40 (2018), pp. A4105–A4130, <https://doi.org/10.1137/17M1144350>.
- [101] E. McDONALD, J. PESTANA, AND A. WATHEN, *Preconditioning and iterative solution of all-at-once systems for evolutionary partial differential equations*, SIAM Journal on Scientific Computing, 40 (2018), pp. A1012–A1033, <https://doi.org/10.1137/16M1062016>.
- [102] *MFEM: Modular finite element methods [software]*. mfem.org, <https://doi.org/10.11578/dc.20171025.1248>.
- [103] Z. MIAO, Y.-L. JIANG, AND Y.-B. YANG, *Convergence analysis of a parareal-in-time algorithm for the incompressible non-isothermal flows*, International Journal of Computer Mathematics, 96 (2019), pp. 1398–1415, <https://doi.org/10.1080/00207160.2018.1498484>.
- [104] M. L. MINION, R. SPECK, M. BOLTEN, M. EMMETT, AND D. RUPRECHT, *Interweaving PFASST and parallel multigrid*, SIAM Journal on Scientific Computing, 37 (2015), pp. S244–S263, <https://doi.org/10.1137/14097536X>.
- [105] A. MITCHELL AND D. GRIFFITHS, *The Finite Difference Method in Partial Differential Equations*, John Wiley & Sons, Chichester-New York-Brisbane-Toronto, 1980.
- [106] K. MIYAMOTO, *Plasma physics for controlled fusion*, Springer-Verlag, Berlin, Germany, 2nd ed., 2016.
- [107] M. MURPHY, G. GOLUB, AND A. WATHEN, *A note on preconditioning for indefinite linear systems*, SIAM Journal on Scientific Computing, 21 (2000), pp. 1969–1972, <https://doi.org/10.1137/S1064827599355153>.

- [108] M. K. NG, *Iterative Methods for Toeplitz Systems*, Numerical Mathematics and Scientific Computation, Oxford University Press, UK, 2004.
- [109] A. S. NIELSEN, G. BRUNNER, AND J. S. HESTHAVEN, *Communication-aware adaptive Parareal with application to a nonlinear hyperbolic system of partial differential equations*, Journal of computational physics, 15 (2017), pp. 483–505, <https://doi.org/10.1016/j.jcp.2018.04.056>.
- [110] J. NIEVERGELT, *Parallel methods for integrating ordinary differential equations*, Communications of the ACM, 7 (1964), pp. 731–733, <https://doi.org/10.1145/355588.365137>.
- [111] B. W. ONG AND J. B. SCHRODER, *Applications of time parallelization*, Computing and Visualization in Science, 23 (2020), <https://doi.org/10.1007/s00791-020-00331-4>.
- [112] C. C. PAIGE AND M. A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, SIAM Journal on Numerical Analysis, 12 (1975), pp. 617–629, <https://doi.org/10.1137/0712047>.
- [113] J. W. PEARSON, M. STOLL, AND A. J. WATHEN, *Regularization-robust preconditioners for time-dependent PDE-constrained optimization problems*, SIAM Journal on Matrix Analysis and Applications, 33 (2012), pp. 1126–1152, <https://doi.org/10.1137/110847949>.
- [114] B. PHILIP, L. CHACÓN, AND M. PERNICE, *Implicit adaptive mesh refinement for 2d reduced resistive magnetohydrodynamics*, Journal of Computational Physics, 227 (2008), pp. 8855–8874, <https://doi.org/10.1016/j.jcp.2008.06.029>.
- [115] É. PICARD, *Sur l'application des méthodes d'approximations successives à l'étude de certaines équations différentielles ordinaires*, Journal de Mathématiques Pures et Appliquées, 9 (1893), pp. 217–272, <https://doi.org/10.2307/2369869>.
- [116] J. QIU AND C.-W. SHU, *On the construction, comparison, and local characteristic decomposition for high-order central WENO schemes*, Journal of Computational Physics, 183 (2002), pp. 187 – 209, <https://doi.org/doi.org/10.1006/jcph.2002.7191>.

- [117] J.-M. QIU AND C.-W. SHU, *Conservative high order semi-Lagrangian finite difference WENO methods for advection in incompressible flow*, Journal of Computational Physics, 230 (2011), pp. 863–889, <https://doi.org/10.1016/j.jcp.2010.04.037>.
- [118] A. QUARTERONI, *Numerical Models for Differential Problems*, MS&A, Springer International Publishing, 2017, <https://doi.org/10.1007/978-3-319-49316-9>.
- [119] J. REYNOLDS-BARREDO, D. NEWMAN, R. SÁNCHEZ, D. SAMADDAR, L. BERRY, AND W. ELWASIF, *Mechanisms for the convergence of time-parallelized, parareal turbulent plasma simulations*, Journal of Computational Physics, 231 (2012), pp. 7851–7867, <https://doi.org/10.1016/j.jcp.2012.07.028>.
- [120] M. RIES, U. TROTTENBERG, AND G. WINTER, *A note on MGR methods*, Linear Algebra and its Applications, 49 (1983), pp. 1–26, [https://doi.org/doi.org/10.1016/0024-3795\(83\)90091-5](https://doi.org/doi.org/10.1016/0024-3795(83)90091-5).
- [121] J. W. RUGE AND K. STÜBEN, *Algebraic Multigrid*, Society for Industrial and Applied Mathematics, 1987, ch. 4, pp. 73–130, <https://doi.org/10.1137/1.9781611971057.ch4>.
- [122] K. RUPP, *42 years of microprocessor trend data*. <https://www.karlrupp.net/2018/02/42-years-of-microprocessor-trend-data/>. Raw data available at <https://github.com/karlrupp/microprocessor-trend-data>. Accessed on: 2019/06/19.
- [123] D. RUPRECHT, *Parareal’s discrete dispersion relation*. <https://www.birs.ca/events/2016/5-day-workshops/16w5030/videos/watch/201611281030-Ruprecht.html>. Presentation at the *Fifth Parallel-in-Time Integration Workshop*. Accessed on: 2021/06/11.
- [124] D. RUPRECHT, *Wave propagation characteristics of parareal*, Computing and Visualization in Science, 19 (2018), pp. 1–17, <https://doi.org/10.1007/s00791-018-0296-z>.
- [125] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM Journal on Scientific Computing, 14 (1993), pp. 461–469, <https://doi.org/10.1137/0914028>.

- [126] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics, 2nd ed., 2003, <https://doi.org/10.1137/1.9780898718003>.
- [127] Y. SAAD AND M. SCHULTZ, *GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems*, *Siam Journal on Scientific and Statistical Computing*, 7 (1986), pp. 856–869, <https://doi.org/10.1137/0907058>.
- [128] D. SAMADDAR, D. COSTER, X. BONNIN, L. BERRY, W. ELWASIF, AND D. BATCHELOR, *Application of the parareal algorithm to simulations of elms in iter plasma*, *Computer Physics Communications*, 235 (2019), pp. 246–257, <https://doi.org/10.1016/j.cpc.2018.08.007>.
- [129] D. SAMADDAR, D. NEWMAN, AND R. SÁNCHEZ, *Parallelization in time of numerical simulations of fully-developed plasma turbulence using the parareal algorithm*, *Journal of Computational Physics*, 229 (2010), pp. 6558–6573, <https://doi.org/10.1016/j.jcp.2010.05.012>.
- [130] W. SCHIESSER, *The Numerical Method of Lines: Integration of Partial Differential Equations*, Academic Press, 1991.
- [131] SCHMID-BURGK, *Zweidimensionale selbstkonsistente Lösungen der stationären Wlassowgleichung für Zweikomponententplasma*, master’s thesis, Ludwig-Maximilians-Universität München, 1965.
- [132] 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 (2018), pp. 45–57.
- [133] D. SCHÖTZAU, *Mixed finite element methods for stationary incompressible magneto-hydrodynamics*, *Numerische Mathematik*, 96 (2004), pp. 771–800.
- [134] J. SHADID, R. PAWLOWSKI, J. BANKS, L. CHACÓN, P. LIN, AND R. TUMINARO, *Towards a scalable fully-implicit fully-coupled resistive MHD formulation with stabilized FE methods*, *Journal of Computational Physics*, 229 (2010), pp. 7649–7671, <https://doi.org/10.1016/j.jcp.2010.06.018>.
- [135] J. SHADID, R. PAWLOWSKI, E. CYR, R. TUMINARO, L. CHACÓN, AND P. WEBER, *Scalable implicit incompressible resistive MHD with stabilized fe*

- and fully-coupled Newton-Krylov-AMG, *Computer Methods in Applied Mechanics and Engineering*, 304 (2016), pp. 1–25, <https://doi.org/10.1016/j.cma.2016.01.019>.
- [136] C.-W. SHU, *Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws*, Springer Berlin Heidelberg, Berlin, Heidelberg, 1998, pp. 325–432, <https://doi.org/10.1007/BFb0096355>.
- [137] D. SILVESTER, H. ELMAN, D. KAY, AND A. WATHEN, *Efficient preconditioning of the linearized Navier-Stokes equations for incompressible flow*, *Journal of Computational and Applied Mathematics*, 128 (2001), pp. 261 – 279, [https://doi.org/10.1016/S0377-0427\(00\)00515-X](https://doi.org/10.1016/S0377-0427(00)00515-X). *Numerical Analysis 2000. Vol. VII: Partial Differential Equations*.
- [138] A. A. SIVAS, B. S. SOUTHWORTH, AND S. RHEBERGEN, *AIR algebraic multigrid for a space-time hybridizable discontinuous Galerkin discretization of advection (-diffusion)*, ArXiv preprint, [arXiv:2010.11130 \[math.NA\]](https://arxiv.org/abs/2010.11130), (2020), <https://arxiv.org/abs/2010.11130>.
- [139] B. SOUTHWORTH, *Necessary conditions and tight two-level convergence bounds for parareal and multigrid reduction in time*, *SIAM Journal on Matrix Analysis and Applications*, 40 (2019), pp. 564–608, <https://doi.org/10.1137/18M1226208>.
- [140] B. S. SOUTHWORTH, W. MITCHELL, A. HESSENTHALER, AND F. DANIELI, *Tight two-level convergence of linear parareal and MGRIT: Extensions and implications in practice*, ArXiv preprint, [arXiv:2010.11879 \[math.NA\]](https://arxiv.org/abs/2010.11879), (2020), <https://arxiv.org/abs/2010.11879>.
- [141] B. S. SOUTHWORTH, A. A. SIVAS, AND S. RHEBERGEN, *On fixed-point, Krylov, and 2×2 block preconditioners for nonsymmetric problems*, *SIAM Journal on Matrix Analysis and Applications*, 41 (2020), pp. 871–900, <https://doi.org/10.1137/19M1298317>.
- [142] R. SPECK AND D. RUPRECHT, *Toward fault-tolerant parallel-in-time integration with PFASST*, ArXiv preprint, [arXiv:1510.08334 \[cs.DC\]](https://arxiv.org/abs/1510.08334), (2015), <https://arxiv.org/abs/1510.08334>.

- [143] G. A. STAFF AND E. M. RØNQUIST, *Stability of the parareal algorithm*, in Domain Decomposition Methods in Science and Engineering, T. J. Barth, M. Griebel, D. E. Keyes, R. M. Nieminen, D. Roose, T. Schlick, R. Kornhuber, R. Hoppe, J. Périaux, O. Pironneau, O. Widlund, and J. Xu, eds., Berlin, Heidelberg, 2005, Springer Berlin Heidelberg, pp. 449–456, https://doi.org/10.1007/3-540-26825-1_46.
- [144] M. STOLL AND A. WATHEN, *All-at-once solution of time-dependent Stokes control*, Journal of Computational Physics, 232 (2013), <https://doi.org/10.1016/j.jcp.2012.08.039>.
- [145] G. STRANG, *A proposal for Toeplitz matrix calculations*, Studies in Applied Mathematics, 74 (1986), pp. 171–176, <https://doi.org/10.1002/sapm1986742171>.
- [146] P. N. SWARZTRAUBER, *Multiprocessor FFTs*, Parallel Computing, 5 (1987), pp. 197 – 210, [https://doi.org/10.1016/0167-8191\(87\)90018-4](https://doi.org/10.1016/0167-8191(87)90018-4). Proceedings of the International Conference on Vector and Parallel Computing - Issues in Applied Research and Development.
- [147] A. TOSELLI AND O. WIDLUND, *Domain Decomposition Methods - Algorithms and Theory*, Springer Series in Computational Mathematics, Springer Berlin Heidelberg, 2006.
- [148] J. TRINDADE AND J. PEREIRA, *Parallel-in-time simulation of the unsteady Navier–Stokes equations for incompressible flow*, International journal for numerical methods in fluids, 45 (2004), pp. 1123–1136, <https://doi.org/10.1002/flid.732>.
- [149] J. TRINDADE AND J. PEREIRA, *Parallel-in-time simulation of two-dimensional, unsteady, incompressible laminar flows*, Numerical Heat Transfer, Part B: Fundamentals, 50 (2006), pp. 25–40, <https://doi.org/10.1080/10407790500459379>.
- [150] U. TROTTENBERG, C. OOSTERLEE, AND A. SCHULLER, *Multigrid*, Elsevier Science, 2000.
- [151] UK ATOMIC ENERGY AUTHORITY, *Culham Centre for Fusion Energy*. <https://ccfe.ukaea.uk/>. Accessed on: 2021/06/01.

- [152] Q. WANG, S. A. GOMEZ, P. J. BLONIGAN, A. L. GREGORY, AND E. Y. QIAN, *Towards scalable parallel-in-time turbulent flow simulations*, *Physics of Fluids*, 25 (2013), p. 110818, <https://doi.org/10.1063/1.4819390>.
- [153] A. WATHEN, *Some observations on preconditioning for non-self-adjoint and time-dependent problems*, *Computers & Mathematics with Applications*, (2021), <https://doi.org/10.1016/j.camwa.2021.05.037>.
- [154] A. WATHEN AND T. REES, *Chebyshev semi-iteration in preconditioning for problems including the mass matrix*, *Electronic Transactions on Numerical Analysis*, 34 (2008-2009), pp. 125–135.
- [155] A. J. WATHEN, *Realistic eigenvalue bounds for the Galerkin mass matrix*, *IMA Journal of Numerical Analysis*, 7 (1987), pp. 449–457, <https://doi.org/10.1093/imanum/7.4.449>.