

NEPTUNE Technical Report 2060049-TN-01

Deliverables D1.1 and D2.1

A Review of Time Stepping Techniques and Preconditioning for Hyperbolic and Anisotropic Elliptic Problems

Hussam Al Daas¹, Tyrone Rees¹, and H. Sue Thorne²

¹*Scientific Computing Department, UKRI-STFC*

²*Hartree Centre, UKRI-STFC*

March 25, 2022

1 Introduction

Exascale targeted plasma modelling will require the efficient and accurate solution of systems of hyperbolic partial differential equations (PDEs), and the corresponding elliptic problems, in the presence of highly anisotropic dynamics. Furthermore, the techniques used must scale with the computing power available, and be robust enough to be portable to any emerging hardware that arises in the future. We are anticipating that high order finite elements/spectral elements will be used to discretize the equations, and we focus on methods amenable to such problems.

In Section 2 we investigate the current state of the art in time advance techniques, while in Section 3 we turn our attention to the state of the art for preconditioners of elliptic systems; these fulfill deliverables 1.1 and 2.1 of ExCALIBUR project NEPTUNE: Mathematical Support for Software Implementation. We give a high level summary of our findings at the end of each section.

2 State-of-the-art time stepping techniques for hyperbolic PDEs

We consider the solution of hyperbolic systems of the form

$$\frac{\partial u}{\partial t} = \mathcal{A}(u, t),$$

together with appropriate initial and boundary conditions. Stability requirements mean that explicit methods (such as forward Euler) would require an unfeasibly small temporal step size. This problem which is compounded by the fact that the step size for the *entire domain* is restricted by the finest mesh patch or wave velocity, which generally make explicit methods unsuitable for anisotropic hyperbolic equations.

We can make use of larger time steps with an implicit (or semi-implicit) method. While such schemes may be unconditionally stable, even in this case we must still restrict the size of the time step with a CFL-like condition to ensure that the solution is sufficiently accurate. Also, with any implicit method, there is the requirement to solve a large system of equations at each time step, and we must do this carefully to ensure performance on modern HPC systems.

2.1 Fully Implicit Methods

The Method of Lines [44] is a technique to transform a PDE into system of ordinary differential equations (ODEs) by applying a pre-determined discretization strategy to the spatial dimensions of the PDE. We may then solve the resulting ODE with an appropriate temporal scheme to the required accuracy.

In the linear case, and without algebraic constraints, the ODE system takes the form

$$Mu'(t) = \mathcal{L}u + \hat{f}(t) \quad \text{in } (0, T], \quad u(0) = 0, \quad (1)$$

where $M \in \mathbb{R}^{n \times n}$ is a mass matrix, $\mathcal{L} \in \mathbb{R}^{n \times n}$ is a discrete linear operator, and $\hat{f}(t)$ a time-dependent forcing function.

One may solve the ODE (1) by applying an s -stage Runge-Kutta scheme:

$$u_{n+1} = u_n + \delta t \sum_{i=1}^s b_i k_i, \quad (2)$$

$$Mk_i = \mathcal{L} \left(u_n + \delta t \sum_{j=1}^s a_{ij} k_j \right) + f(t_n + \delta t c_i). \quad (3)$$

Such schemes are commonly expressed in terms of a Runge-Kutta matrix, $A = \{a_{ij}\} \in \mathbb{R}^{s \times s}$, weight vector, $b^T = (b_1, \dots, b_s)^T$, and quadrature nodes $c = (c_1, \dots, c_s)$, often presented in a Butcher tableau:

$$\begin{array}{c|c} c & A \\ \hline & b^T \end{array}.$$

The stage vectors $\{k_i\}$ are the solution of the block linear system,

$$\left(\begin{bmatrix} M & & \\ & \ddots & \\ & & M \end{bmatrix} - \delta t \begin{bmatrix} a_{11}\mathcal{L} & \cdots & a_{1s}\mathcal{L} \\ \vdots & \ddots & \vdots \\ a_{s1}\mathcal{L} & \cdots & a_{ss}\mathcal{L} \end{bmatrix} \right) \begin{bmatrix} k_1 \\ \vdots \\ k_s \end{bmatrix} = \begin{bmatrix} f_1 \\ \vdots \\ f_s \end{bmatrix}, \quad (4)$$

where $f_i := \hat{f}(t_n + \delta t c_i) + \mathcal{L}(t_n + \delta t c_i)u_n$. The difficulty in applying fully implicit Runge–Kutta methods lies in solving the $ns \times ns$ block linear system (4).

The solution of nonlinear systems requires solves with the form (4), but with \mathcal{L} now being a linearized operator as determined by, say, a Newton or Picard iteration [43, 76].

We may also incorporate constraints (such as incompressibility); if the Method of Lines applied to a PDE gives the Differential Algebraic Equation

$$\begin{aligned} Mu'(t) &= \mathcal{N}(u, w, t) \\ 0 &= \mathcal{G}(u, w, t) \end{aligned}$$

then an s –stage Runge–Kutta method applied to this gives iterates of the form

$$\begin{bmatrix} u_{n+1} \\ w_{n+1} \end{bmatrix} = \begin{bmatrix} u_n \\ w_n \end{bmatrix} + \delta t \sum_{i=1}^s b_i \begin{bmatrix} k_i \\ \ell_i \end{bmatrix}.$$

In the linear case we obtain the stage vectors k_i, ℓ_i by solving the system of equations

$$\left(\begin{bmatrix} \begin{bmatrix} M & \\ & 0 \end{bmatrix} & & \\ & \ddots & \\ & & \begin{bmatrix} M & \\ & 0 \end{bmatrix} \end{bmatrix} - \delta t \begin{bmatrix} a_{11} \begin{bmatrix} \mathcal{N}_u & \mathcal{N}_w \\ \mathcal{G}_u & \mathcal{G}_w \end{bmatrix} & \cdots & a_{1s} \begin{bmatrix} \mathcal{N}_u & \mathcal{N}_w \\ \mathcal{G}_u & \mathcal{G}_w \end{bmatrix} \\ \vdots & \ddots & \vdots \\ a_{s1} \begin{bmatrix} \mathcal{N}_u & \mathcal{N}_w \\ \mathcal{G}_u & \mathcal{G}_w \end{bmatrix} & \cdots & a_{ss} \begin{bmatrix} \mathcal{N}_u & \mathcal{N}_w \\ \mathcal{G}_u & \mathcal{G}_w \end{bmatrix} \end{bmatrix} \right) \begin{bmatrix} k_1 \\ \ell_1 \\ \vdots \\ k_s \\ \ell_s \end{bmatrix} = \begin{bmatrix} f_1 \\ g_1 \\ \vdots \\ f_s \\ g_s \end{bmatrix}. \quad (5)$$

Again, the nonlinear case is also possible, and results in a series of linearized systems of the form (5). For more details see e.g., [8, Section 10.1.3], [76, Section 6]. In the following, for simplicity of exposition, we consider only the linear case without constraints (unless we state otherwise), but the ideas can also be applied in the more general case.

For general properties of such methods, we refer the reader to Ascher and Petzold [8, Section 4]. In the specific case that interests us, the algebraic block is very large and ill conditioned, being the spatial discretization of a partial differential equation. We can expect standard implementations of algorithms for solving ODEs to fail, as direct methods would struggle to solve even a single system with \mathcal{L} . Below we describe the state of the art.

2.1.1 Diagonally-implicit Runge–Kutta (DIRK) methods

Diagonally-implicit Runge–Kutta (DIRK) methods have a lower triangular Runge–Kutta matrix, A . This simplifies the solution of (4) considerably, as the implicit solve requires only a series of $n \times n$ systems, rather than one $ns \times ns$ system solve. Kennedy and Carpenter [40] performed a comprehensive survey of DIRK methods for NASA, followed up by the development of several new methods [41]. We recommend these papers, and the references therein, for more detail on this class of methods.

DIRK methods can be further divided into a series of subclasses: SDIRK methods are DIRK methods where A has a constant diagonal; EDIRK methods are DIRK methods with an explicit first stage (so $a_{1,1} = 0$); ESDIRK methods [45] are EDIRK methods where the non-zeros on the diagonal are constant. Nektar++ implements SDIRK with two and three stages, and an ESDIRK method with six stages; Yan et al. [84] gives a preliminary comparison of the performance of these methods against each other, and against explicit methods.

Pan et al. [61] point out that one must be careful when choosing the time step, as making a naive choice may lead to extra work without any gain in accuracy (see, e.g., [57]). They observe that, when solving Navier–Stokes equations using a Discontinuous Galerkin discretization in the spatial domain, an ESDIRK method in temporal domain, and solving the nonlinear system using a Jacobian Free-Newton Krylov method (JFNK) [43], then the error in one time step is the sum of the local temporal error (from ESDIRK), and the averaged spatial error (from DG) and the averaged algebraic JFNK error (from preconditioned GMRES). They therefore propose a system, implemented in Nektar++ [84], which uses explicit formulae for the spatial truncation error and temporal error to estimate the errors in these components. These are then used to choose *a priori* a time step and a Newton tolerance such that the temporal and iterative errors are smaller than the spatial errors, thus ensuring that the accuracy

is enough to ensure sufficient convergence, but not too much to waste CPU time. Since this method is based on local errors, there is no guarantee on the global errors; however, it suggests a reasonable heuristic which has been shown to be effective on the isentropic vortex, Taylor-Green vortex, flat plate boundary layer, and turbulent flow over a cylinder model problems. While this approach provides an upper bound for the time step, this may not be enough to maintain stability for challenging problems, yet the method is a promising approach for automatic simulation pipelines.

2.1.2 High order implicit Runge–Kutta (IRK) method

While DIRK schemes are attractive, since they simplify the structure of the linear systems, they come with a number of drawbacks, as outlined in [40]. For a DIRK method with formal integration order p and stage-order q , the order of accuracy observed in practice stiff nonlinear PDEs or DAEs is approximately $\min\{p, q + 1\}$. Stable DIRK methods have a maximum order of $p = s + 1$, and the stage-order, q , is usually 1, although can be 2 for DIRK methods with an explicit first stage. However, DIRK methods cannot have a stage-order larger than 2. Symplectic DIRK methods can be at most 4th order, and have the additional restriction that documented methods above order two have Runge–Kutta matrices with negatives on the diagonal, which usually leads to more difficult linear systems to solve. Overall, DIRK methods have limited accuracy: often this is good enough for the application, but can be an issue for difficult modelling problems.

Fully implicit Runge–Kutta (IRK) methods, in contrast, can have high-order accuracy, since they may have any stage-order: an s –stage IRK method may have accuracy of order $2s$. However, the linear system solved at each step (4) is formidable, and is intractable for realistic sized problems without careful handling of the linear algebra. Common IRK methods include the Gauss-Legendre, Gauss-Lobatto and Gauss-Radau families, which contain classical techniques such as backward Euler (RadauIIA), and Crank-Nicolson (LobattoIIIA).

Farrell et al. [26] have developed a high-level library called *Irksome* for manipulating UFL (Unified Form Language) expressions of semidiscrete variational forms to obtain UFL expressions for the coupled Runge–Kutta stage equations (3) at each time step. *Irksome* works with the *Firedrake* package to enable the efficient solution of the resulting coupled algebraic systems, which are solved matrix-free. *Irksome* solves the matrix (4) using a Krylov method with preconditioner

$$\text{blkdiag}(M - a_{1,1}\delta t\mathcal{L}, \dots, M - a_{s,s}\delta t\mathcal{L}),$$

which has been shown [50] to be a good preconditioner for parabolic problems.

The system (4) can be re-written in Kronecker product form as

$$(I \otimes M - \delta t A \otimes \mathcal{L})k = f, \tag{6}$$

which is the Sylvester matrix equation. Such systems arise commonly in model order reduction and control applications, and we refer the reader to Simoncini’s survey [75] on techniques for solving matrix equations. This connection has led to a number of interesting approaches for the solution of (4) using fully-implicit IRK methods over the past few years [37, 51, 67, 76, 77].

Southworth et al. introduce a theoretical and algorithmic preconditioning framework for solving (4) in the linear [77] and nonlinear [76] setting. This framework also naturally applies to discontinuous Galerkin discretizations in time. Under quite general assumptions on the spatial discretization that yield stable time integration, they prove that the preconditioned operator has a condition number bounded by a small, order-one constant, independent of the spatial mesh and time-step size, and with only weak dependence on number of stages/polynomial order; for example, the preconditioned operator for 10th-order Gauss IRK has condition number less than two, independent of the spatial discretization and time step.

The proposed method can be used with arbitrary existing preconditioners for backward Euler-type time-stepping schemes and is amenable to the use of three-term recursion Krylov methods when the underlying spatial discretization is symmetric. In [77], the authors apply their method to various high-order finite-difference and finite element discretizations of linear parabolic and hyperbolic problems, demonstrating fast, scalable solution of up to 10th-order accuracy. The proposed method, in several cases, can achieve 4th-order accuracy using Gauss integration with roughly half the number

of preconditioner applications and wallclock time as required using standard diagonally IRK methods. In [76] the authors treat the nonlinear case and the existence of differential algebraic equations, applying the method to the nonlinear Navier–Stokes equation. Again, the method only requires an efficient preconditioner for matrices arising from classic backward Euler scheme.

The numerical experiments in [76, 77] used the software MFEM [7], and employs multigrid preconditioners that are available in Hypr (see 3). We highlight that Masud et al. [51] presented a similar approach to that of Southworth et al., treating the parabolic case.

Common to all the approaches described here are fast methods for the solution of a sequence of linear systems of equations of the form

$$(\alpha_i M - \mathcal{L})x = f_i \text{ or } \begin{bmatrix} \alpha_i M - \mathcal{L} & B^T \\ B & C \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} f_{i1} \\ f_{i2} \end{bmatrix}, \quad (7)$$

depending on if constraints are present or not, where M is a mass matrix; \mathcal{L} a linear differential operator; and B , B^T and C represent the differential algebraic constraints. The α_i 's depend on the values in the Runge–Kutta matrix A , and can be assumed to be positive. The symmetry and definiteness of (7) therefore depends on the linear differential operator \mathcal{L} . We point to Section 3 for a description of recent advances in preconditioning (7) for both the symmetric and nonsymmetric case.

2.2 Linear Multistep Methods

Linear multistep time integration techniques applied to (1) take the form

$$\sum_{j=-k}^0 \alpha_{k+j} M u_{i+j} = \delta t \sum_{j=-k}^0 \beta_{k+j} (\mathcal{L}(u_{i+j}) + f(t_{i+j})),$$

for a given set of α_i, β_j ; methods with $\beta_k = 0$ are explicit (e.g., the Adams–Bashforth family), and non-zero β_k gives implicit methods (e.g., the Adams–Moulton (AM) and Backward Differentiation Formulae (BDF) families).

There are, however, a number of drawbacks which limit their application to hyperbolic PDEs. Although they don't suffer from the accuracy issues of DIRK methods, there are no implicit multistep methods which are A-stable and of order greater than two. There are also no generally symplectic multistep methods. By the nature of multistep methods, which build approximations to the solution using solutions at previous points in time, multistep methods are more memory intensive than Runge–Kutta methods, and so may not be so applicable for large simulations at exascale.

2.3 Implicit-Explicit (IMEX) methods

It is often the case that partial differential equations take the form

$$\frac{\partial u}{\partial t} = \mathcal{B}(u, t) + \mathcal{C}(u, t),$$

where the operators $\mathcal{B}(u, t)$ and $\mathcal{C}(u, t)$ model phenomena on different time-scales; examples include a non-stationary convection-diffusion equation where we follow the convected time scale, where $\mathcal{B}(u, t)$ and $\mathcal{C}(u, t)$ may represent the diffusion and convection terms, respectively, or coupled multiphysics problems, where mixed fluids may have different behaviours.

The presence of the fast term precludes the use of fully explicit integrators, as we would require a prohibitively small time-step to ensure convergence. On the other hand, using a fully implicit method may ask too much of the linear (or non-linear) solvers. Implicit-Explicit (IMEX) time integration schemes exploit the structure of the problem by solving for the slow term explicitly, where we may safely take larger time steps, and solving for the fast term implicitly. The split between the two different regimes is not always clear, and schemes may leak stiffness. IMEX methods also may suffer from accuracy issues, and fully implicit methods may need to be used if a tight tolerance is required.

A number of IMEX schemes have recently been proposed for general hyperbolic problems [16, 33, 62], for multiphysics problems (see [42, Section 3.2.2] and the references therein), and for plasma modelling [54]. Nektar++ implements a scheme by Karndarkis et al. [38] for thermal convection problems [36]; we note that Nektar++ uses an equivalence between IMEX method and general linear methods to ease the

switch between various implicit, explicit and IMEX methods [82]. PETSc’s TS package [11, Section 6] provides interfaces to the IMEX methods proposed in [9, 16, 32, 39, 62] (mostly using an SDIRK implicit solver). We would like to highlight that IMEX methods (together with bespoke preconditioners for the implicit solves) have recently been successfully applied at scale as the integrator in the Unified Model [15, 52].

2.4 Parallel-in-time

As the parallel performance of spatial solvers has become saturated, there has been an increase in interest in parallel in time (PinT) methods; see the survey by Gander [27]. The development of the parareal method by Lions, Maday and Turinici in 2001 [46] was the catalyst for much of this activity. The parareal algorithm combines two solvers; a cheap, global, ‘coarse’ solver, and a more accurate ‘fine’ solver. The basic concept is simple: we employ a fine solver in parallel, which constitutes the bulk of the computational effort, and then combine the results into a global solution using the coarse solver. While most results in the literature apply this method to model problems and, in particular, those of a diffusive nature, there has been some success in applying the parareal paradigm to real scientific applications with a hyperbolic PDE, most notably in the work of Samaddar et al. on tokamak edge plasma simulations [14, 23, 69–74]. Key in this work is the selection of an appropriate coarse grid solver, and unfortunately there is currently little (if any) theoretical guidance on how which schemes would prove successful for other problems. Nevertheless, these studies report a speed up of roughly a factor of ten by using parareal over conventional time-stepping techniques.

An alternative approach to parareal is Multigrid Reduction in Time (MGRIT) [24], a method built on the equivalence between a traditional time integration method and a block lower triangular system of equations. The parareal method is equivalent to two grid (in time) multigrid method [29] and MGRIT is, in some ways, the natural extension, considering a hierarchy of grids. XBraid software [1] is an implementation of MGRIT, and has been shown to give speedups of up to a factor of fifty over sequential time-stepping.

Both parareal and MGRIT struggle on hyperbolic problems. Southworth et al. [78] recently provided a convergence analysis detailing whether or not parareal or two-level MGRIT will converge on a given problem, as well as what spatial or temporal discretizations are applicable with MGRIT. Recent work by Wathen and collaborators [20, 53] and Gander et al [28], based on the development of block preconditioners for the large block lower-triangular time-stepping matrix, has potential to overcome this limitation.

We highlight there is an ExCALIBUR project to compare the performance of parallel in time methods for exascale use, the findings of which will be relevant here. <https://excalibur.ac.uk/projects/exposing-parallelism-parallel-in-time/>

Summary

- DIRK time-stepping methods, where applicable, are well developed, and [40] is a nice summary.
- Fully implicit Runge–Kutta methods were largely dismissed as impractical for hyperbolic problems, but recent advances in preconditioning techniques may mean this is no-longer the case; Southworth et al. [76, 77] describes (and advances) the state of the art.
- The theory of IMEX methods is also well developed, and should be used where appropriate.
- Recent work in parallel-in-time methods may help break the inherently sequential nature of traditional time integration algorithms, which rely exclusively on the solution of the spatial discretization for parallelism.
- All approaches for solving implicit time-stepping problems require good methods for solving (variations of) the stationary system; see Section 3.

3 State-of-the-art approaches for preconditioning elliptic problems discretized with high-order methods

Existing and emerging computing architectures suffer from an exponentially growing gap between the time necessary to perform a floating point operation and the time to move data across the network on a distributed memory environment. High-order numerical methods provide higher accuracy with fewer degrees of freedom, at the cost of more arithmetic operations performed per degree of freedom. Because of their high arithmetic intensity, high-order methods are promising candidates to get the best performance on exascale systems.

A finite element or discontinuous Galerkin method of polynomial degree p in d spatial dimensions will have $\mathcal{O}(p^d)$ degrees of freedom per element, and hence the system matrix will have $\mathcal{O}(p^{2d})$ non-zeros entries. For this reason the coefficient matrix is often not assembled, and operations with it are carried out in a matrix-free fashion. A matrix-free method reduces the memory requirements to $\mathcal{O}(p^d)$, and, making use of sum factorization techniques, the cost of a matrix-vector product can be reduced to $\mathcal{O}(dp^{d+1})$ [60].

The condition number of the discretized matrix increases quadratically with the polynomial degree. Due to the large bandwidth, direct methods for solving linear systems are not an option. We therefore must turn to iterative solvers, and it is vital that we pair the Krylov subspace solver with an appropriate preconditioner to get good performance.

In this section, we summarize the state-of-the-art in techniques proposed to precondition linear systems arising from the discretization of elliptic PDEs by using high-order methods. We point the reader to the NEPTUNE report [6] for an overview of general-purpose preconditioners, as well as the excellent review by Wathen [83].

3.1 Preconditioning high order finite element matrices

There is a good body of work on preconditioners for low to moderate order finite elements, and there are some excellent implementations of algorithms which have been used in a wide variety of practical situations; see Section 3.2. However, it has been observed that such methods often give less than desirable performance when applied to high order discretization (however, see Heys et al. [35], who show that AMG can be applied successfully – albeit with a mild p –dependence – with minor modifications).

One technique that has proved successful, first proposed by Orszag in 1980 [60], is to exploit the spectral equivalence between low-order and high-order operators, known as FEM-SEM equivalence; see the review by Canuto, Gervasio and Quarteroni [18] and the references therein. This approach allows us to use a low order discretization to precondition the high order problem, allowing the use of the methods outlined in Section 3.2 as a fast approximation to this ‘ideal’, low-order, preconditioner.

In practice, the tensor-product of the Gauss-Lobatto-Legendre points used to generate the grid for spectral element methods result in anisotropic finite element meshes, which challenge basic multigrid and domain decomposition methods [47]. Nevertheless, there has been considerable success in the development of such methods in recent years; Pazner and Kolev [64,65] develop a multigrid preconditioner for high-order continuous and discontinuous Galerkin methods with hp –refinement; Chalmers and Warburton [19] and Olson [59] apply this idea to simplexes in two and three dimensions; Bello-Maldonado and Fischer [13] create a scheme that uses a preconditioner based on $P1$ -elements, using a meshing technique for rectangular and hexahedral elements; Pazner, Dohrmann and Kolev [21,66] show this can be applied to finite element problems on $H(\text{curl})$ and $H(\text{div})$ spaces (using Nédélec and Raivart-Thomas elements, respectively). These methods are typically independent of the polynomial degree, mesh size and, for DG methods, the penalty parameter. In a comparison by Sundar, Stadler and Biros [80] conclude that this approach is more advantageous than h – or p – multigrid.

We particularly highlight the work of Pazner, Dohrmann and Kolev [66], who present numerical experiments using the finite element library MFEM, with which the construction of such preconditioners requires only one or two lines of code. These tests corroborate the theoretical properties of the proposed preconditioners, and demonstrate the flexibility and scalability of the method on a range of challenging three-dimensional problems. These new solvers are flexible and easy to use; any black-box preconditioner for low-order problems can be used to create an effective and efficient preconditioner

for the corresponding high-order problem. The `lor_solvers` miniapp, and its parallel counterpart `plor_solvers`, illustrate the construction of low-order-refined discretizations and solvers, and come distributed with MFEM’s source code, available at <https://github.com/mfem/mfem>.

An alternative approach is to use the observation of Pavarino [63] that additive Schwarz, together with an additive coarse space of order one and a vertex-centered space decomposition, is robust with respect to both p and h for symmetric and coercive problems. The local solves for such problems are dense, and solved with a direct method, and also the coarse-grid operator is also fairly dense and can quickly become expensive. However, in recent years good methods for solving this system have become available; see, e.g., Lottes and Fischer [47].

Recently Brubeck and Farrell [17] introduced a p -robust preconditioner which uses the additive Schwarz method with vertex patches combined with a low-order coarse space as a solver for symmetric and coercive problems. By constructing a tensor product basis that diagonalizes the blocks in the stiffness matrix for the internal degrees of freedom of each individual cell, they show that the patch problem is as sparse as a low-order finite difference discretization and having a sparse Cholesky factorization allows them to scale to large polynomial degree p and afford the assembly and factorization of the matrices in the vertex-patch problems. They successfully apply their method to the Poisson equation and the mixed formulation of linear elasticity both with constant coefficient problems and claim that the theory of [10] suggests that it would remain effective for spatially varying coefficients. In their conclusion, they state that the downside of their approach: (1) its narrow applicability; it will not be effective on more general problems (tested on Poisson and mixed formulation of linear elasticity), especially for those where the dominant terms include mixed derivatives and mixed vector components. (2) their method relies on having a good quality mesh, with its performance depending on the minimal angle.

Finally, we highlight that even explicit time integration methods require the solution of a mass matrix at each time step; while this is trivial in low-order, it is less obviously the case for high-order problems, as the condition number of the mass matrix grows algebraically with the polynomial order. Ainsworth and Jiang [2] describe an efficient preconditioner for the mass matrix, independent of h and p , based on a specific choice of hierarchical basis, which involves only diagonal solves.

3.2 Preconditioners for low order problems

The methods described above depend on a robust and efficient linear solver for low-order finite element systems. Outside of basic PDEs on uniform grids, where, where Fast Multipole Methods and Fast Fourier Transforms may be successfully applied [31], the choice is between multigrid methods and domain decomposition methods.

3.2.1 Multigrid methods

Multigrid methods [81] are split into algebraic (AMG) and geometric (GMG) variants. AMG methods use the algebraic properties of the assembled matrix to restrict to coarser grids, whereas GMG solvers use information about the underlying mesh connectivity. GMG solvers can be over an order of magnitude faster than the best AMG solvers [31], since they can be used matrix-free, and operators can be modified on different levels, which can be advantageous when solving, e.g., advection diffusion problems [68], or when needing to accommodate non-standard boundary conditions. However, there are excellent robust AMG implementations available which work well off-the-shelf for a wide range of problems, which we describe below, and AMG solvers and preconditioners are some of the fastest black-box numerical methods to solve linear systems.

The convergence of AMG solvers is well established for symmetric positive definite (SPD) linear systems resulting from the discretization of general elliptic PDEs or the spatial discretization of parabolic PDEs. Hyperbolic PDEs remain a challenge for AMG, as well as other fast linear solvers, in part because the resulting linear systems are often highly nonsymmetric. Nevertheless, modern AMG implementations can perform well on a wide range of problems.

The HYPRE library [25], originating from the Lawrence Livermore National Laboratories, contains a number of high quality implementations of multigrid preconditioners. BoomerAMG [34] is a parallel implementation of the classical Ruge-Stüben AMG method, and offers a range of coarsening, interpolation, and smoothing options.

Trilinos’s ML preconditioner [30], developed by Sandia National Laboratory, is another fully featured algebraic multigrid implementation, offering a range of multilevel schemes, smoothers, and coarse solvers.

PETSc’s GAMG preconditioner [11, Section 4.4.5] uses a smoothed aggregation coarsening strategy, and offers a reference implementation of the classical Ruge-Stüben method. It also provides unsmoothed aggregation, which can be useful for unsymmetric problems. Note that both HYPRE and ML are also available via the PETSc interface.

Notay’s AGMG method [55,56] is another aggregation-based algebraic multigrid method which can be applied to any system matrix that has positive diagonal entries. The implementation can use MPI, multithreading, or both.

It is well known that AMG methods don’t perform well ‘out-of-the-box’ for anisotropic problems, and may require some problem-specific tuning; the PETSc documentation [11, Section 4.4.5] gives some advice for this.

Manteuffel et al. [49] present a new variation on classical AMG for nonsymmetric matrices (denoted ℓ AIR), based on a local approximation to the ideal restriction operator, coupled with F-relaxation. They demonstrate the efficacy of the proposed preconditioner on systems arising from the discrete form of the advection-diffusion-reaction equation. ℓ AIR is shown to be a robust solver for various discretizations of the advection-diffusion-reaction equation, including time-dependent and steady-state, from purely advective to purely diffusive. Convergence is robust for discretizations on unstructured meshes and using higher-order finite elements, and is particularly effective on upwind discontinuous Galerkin discretizations. ℓ -AIR available in PyAMG [58], and an implementation through HYPRE is underway.

In a related, Manteuffel et al. [48] present a reduction-based AMG method developed for upwind discretizations of hyperbolic PDEs, based on the concept of a Neumann approximation to ideal restriction (n AIR). n AIR can be seen as a variation of local AIR (ℓ AIR) specifically targeting matrices with triangular structure. Although less versatile than ℓ AIR, setup times for n AIR can be substantially faster for problems with high connectivity. n AIR is shown to be an effective and scalable solver of steady state transport for discontinuous, upwind discretizations, with unstructured meshes, and up to 6th-order finite elements, offering a significant improvement over existing AMG methods. n AIR is also shown to be effective on several classes of nearly triangular matrices resulting from curvilinear finite elements and artificial diffusion.

3.2.2 Domain decomposition methods

Domain decomposition methods are among the most efficient for solving sparse linear systems of equations on massively parallel architectures; see, e.g., the book for Dolean, Jolivet and Nataf [22] for a recent overview of the field. Their effectiveness relies on a judiciously chosen coarse space. Spillane et al. introduced GenEO, a spectral coarse space, in [79]. GenEO proved to be very attractive as it can be constructed efficiently and adapts to the underlying problem and its difficulty with minimum amount of effort from the user perspective. This method is theoretically proved to be efficient and robust with respect to the problem size, mesh, discretization type and order and parameters for elliptic PDEs [3,12]. That is, the preconditioner can be constructed efficiently such that the condition number of the preconditioned matrix is bounded from above by a user defined number.

Al Daas et al. [4] have recently extended GenEO to be applicable as a black-box method to precondition general sparse linear system. They assessed the efficiency and scalability of the algebraic GenEO using a variety of very challenging problems arising from a wide range of applications and including highly non symmetric problems such as the advection dominated advection-diffusion equation. Comparisons against state-of-the-art multigrid preconditioners illustrated the robustness and efficiency of algebraic GenEO.

Al Daas et al. [5] recently introduced an algebraic extension of GenEO for the diagonally weighted normal equation matrix. Their motivation was to precondition iterative methods for solving linear least-squares problems. The weighted normal equation matrix also arises in block preconditioning where a preconditioner is required for the (approximate) Schur complement matrix. The preconditioner proposed in [5] can be directly employed within a block preconditioner resulting in a robust, efficient and scalable preconditioner for the (approximate) Schur complement matrix.

Implementations of the Algebraic GenEO methods proposed in [4, 5] are available through the PETSc preconditioner PCHPDDM, and only require the coefficient matrix.

Summary

- The most promising method solving matrices from high order elliptic PDEs is to precondition with a low order finite element operator; low-order solvers have become robust enough to solve the resulting (challenging) linear systems [66].
- Multigrid and Domain Decomposition based preconditioners are the leading contenders for performant parallel iterative solves on anisotropic elliptic problems of low order.
- There are several excellent algebraic multigrid packages available that are highly parallel. However, to get good performance it is vital to tune the parameters for a given problem. A well-implemented geometric multigrid method would give superior performance to an algebraic multigrid.
- The GenEO method [79] is the domain decomposition method with most promise, with a good trade off between performance and ease of setup; the recent development of fully algebraic variations [4, 5] make this even more the case.

References

- [1] XBraid: Parallel multigrid in time. <http://llnl.gov/casc/xbraid>.
- [2] Mark Ainsworth and Shuai Jiang. Preconditioning the Mass Matrix for High Order Finite Element Approximation on Tetrahedra. *SIAM Journal on Scientific Computing*, 43(1):A384–A414, January 2021.
- [3] H. Al Daas, L. Grigori, P. Jolivet, and P.-H. Tournier. A multilevel Schwarz preconditioner based on a hierarchy of robust coarse spaces. *SIAM Journal on Scientific Computing*, 43(3):A1907–A1928, 2021/07/01 2021.
- [4] H. Al Daas, P. Jolivet, and T. Rees. Efficient algebraic two-level schwarz preconditioner for sparse matrices, 2022.
- [5] H. Al Daas, P. Jolivet, and J. A. Scott. A robust algebraic domain decomposition preconditioner for sparse normal equations, 2022. In press. preprint:<https://arxiv.org/abs/2107.09006>.
- [6] V. Alexandrov, A. Lebedev, E. Sahin, and S. Thorne. Linear systems of equations and preconditioners relating to the NEPTUNE Programme. NEPTUNE Technical Report 2047353-TN-04, CCFE Culham, April 2021.
- [7] R. Anderson, J. Andrej, A. Barker, J. Bramwell, J.-S. Camier, J. Cerveny V. Dobrev, Y. Dudouit, A. Fisher, Tz. Kolev, W. Pazner, M. Stowell, V. Tomov, I. Akkerman, J. Dahm, D. Medina, and S. Zampini. MFEM: A modular finite element methods library. *Computers & Mathematics with Applications*, 81:42–74, 2021.
- [8] Uri M. Ascher and Linda R. Petzold. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. SIAM, August 1998.
- [9] Uri M. Ascher, Steven J. Ruuth, and Raymond J. Spiteri. Implicit-explicit Runge-Kutta methods for time-dependent partial differential equations. *Applied Numerical Mathematics*, 25(2):151–167, November 1997.
- [10] O. Axelsson and J. Karátson. Equivalent operator preconditioning for elliptic problems. *Numerical Algorithms*, 50(3):297–380, 2009.

- [11] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. Gropp, D. Karpeyev, D. Kaushik, M. Knepley, D. May, L. Curfman McInnes, R. Mills, T. Munson, K. Rupp, P. Sanan, B. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc Users Manual. 2019.
- [12] Peter Bastian, Robert Scheichl, Linus Seelinger, and Arne Strehlow. Multilevel spectral domain decomposition. *SIAM Journal on Scientific Computing*, pages S1–S26, 2022/03/16 2022.
- [13] Pedro D. Bello-Maldonado and Paul F. Fischer. Scalable Low-Order Finite Element Preconditioners for High-Order Spectral Element Poisson Solvers. *SIAM Journal on Scientific Computing*, 41(5):S2–S18, January 2019.
- [14] L. A. Berry, W. Elwasif, J. M. Reynolds-Barredo, D. Samaddar, R. Sanchez, and D. E. Newman. Event-based parareal: A data-flow based implementation of parareal. *Journal of Computational Physics*, 231(17):5945–5954, July 2012.
- [15] Jack Betteridge, Thomas H. Gibson, Ivan G. Graham, and Eike H. Müller. Multigrid preconditioners for the hybridised discontinuous galerkin discretisation of the shallow water equations. *Journal of Computational Physics*, 426:109948, 2021.
- [16] S. Boscarino, L. Pareschi, and G. Russo. Implicit-Explicit Runge–Kutta Schemes for Hyperbolic Systems and Kinetic Equations in the Diffusion Limit. *SIAM Journal on Scientific Computing*, 35(1):A22–A51, January 2013.
- [17] Pablo D Brubeck and Patrick E Farrell. A scalable and robust vertex-star relaxation for high-order fem, 2021.
- [18] Claudio Canuto, Paola Gervasio, and Alfio Quarteroni. Finite-Element Preconditioning of G-NI Spectral Methods. *SIAM Journal on Scientific Computing*, 31(6):4422–4451, January 2010.
- [19] Noel Chalmers and T. Warburton. Low-Order Preconditioning of High-Order Triangular Finite Elements. *SIAM Journal on Scientific Computing*, 40(6):A4040–A4059, January 2018.
- [20] Federico Danieli, Ben S. Southworth, and Andrew J. Wathen. Space-time block preconditioning for incompressible flow. *arXiv:2101.07003 [cs, math]*, January 2021.
- [21] Clark R. Dohrmann. Spectral Equivalence of Low-Order Discretizations for High-Order H(curl) and H(div) Spaces. *SIAM Journal on Scientific Computing*, 43(6):A3992–A4014, January 2021.
- [22] Victorita Dolean, Pierre Jolivet, and Frederic Nataf. *An Introduction to Domain Decomposition Methods: Algorithms, Theory, and Parallel Implementation*. SIAM, December 2015.
- [23] Wael R. Elwasif, Samantha S. Foley, David E. Bernholdt, Lee A. Berry, Debasmita Samaddar, David E. Newman, and Raul Sanchez. A dependency-driven formulation of parareal: Parallel-in-time solution of PDEs as a many-task application. In *Proceedings of the 2011 ACM International Workshop on Many Task Computing on Grids and Supercomputers*, pages 15–24. Association for Computing Machinery, New York, NY, USA, November 2011.
- [24] R. D. Falgout, S. Friedhoff, Tz. V. Kolev, S. P. MacLachlan, and J. B. Schroder. Parallel Time Integration with Multigrid. *SIAM Journal on Scientific Computing*, 36(6):C635–C661, January 2014.
- [25] Robert D. Falgout and Ulrike Meier Yang. Hypre: A Library of High Performance Preconditioners. In Peter M. A. Sloot, Alfons G. Hoekstra, C. J. Kenneth Tan, and Jack J. Dongarra, editors, *Computational Science — ICCS 2002*, pages 632–641, Berlin, Heidelberg, 2002. Springer.
- [26] Patrick E. Farrell, Robert C. Kirby, and Jorge Marchena-Menéndez. Irksome: Automating runge–kutta time-stepping for finite element methods. *ACM Trans. Math. Softw.*, 47(4), sep 2021.

- [27] Martin J. Gander. 50 Years of Time Parallel Time Integration. In Thomas Carraro, Michael Geiger, Stefan Körkel, and Rolf Rannacher, editors, *Multiple Shooting and Time Domain Decomposition Methods*, pages 69–113, Cham, 2015. Springer International Publishing.
- [28] Martin J. Gander, Jun Liu, Shu-Lin Wu, Xiaoqiang Yue, and Tao Zhou. ParaDiag: Parallel-in-time algorithms based on the diagonalization technique. *arXiv:2005.09158 [cs, math]*, April 2021.
- [29] Martin J. Gander and Stefan Vandewalle. Analysis of the Parareal Time-Parallel Time-Integration Method. *SIAM Journal on Scientific Computing*, 29(2):556–578, January 2007.
- [30] M.W. Gee, C.M. Siefert, J.J. Hu, R.S. Tuminaro, and M.G. Sala. Ml 5.0 smoothed aggregation user’s guide. Technical Report SAND2006-2649, Sandia National Laboratories, 2006.
- [31] Amir Gholami, Dhairya Malhotra, Hari Sundar, and George Biros. FFT, FMM, or Multigrid? A comparative Study of State-Of-the-Art Poisson Solvers for Uniform and Nonuniform Grids in the Unit Cube. *SIAM Journal on Scientific Computing*, 38(3):C280–C306, January 2016.
- [32] F. X. Giraldo, J. F. Kelly, and E. M. Constantinescu. Implicit-Explicit Formulations of a Three-Dimensional Nonhydrostatic Unified Model of the Atmosphere (NUMA). *SIAM Journal on Scientific Computing*, 35(5):B1162–B1194, January 2013.
- [33] Sigal Gottlieb, Zachary J. Grant, Jingwei Hu, and Ruiwen Shu. High Order Strong Stability Preserving MultiDerivative Implicit and IMEX Runge–Kutta Methods with Asymptotic Preserving Properties. *SIAM Journal on Numerical Analysis*, 60(1):423–449, February 2022.
- [34] Van Emden Henson and Ulrike Meier Yang. BoomerAMG: A parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics*, 41(1):155–177, April 2002.
- [35] J. J. Heys, T. A. Manteuffel, S. F. McCormick, and L. N. Olson. Algebraic multigrid for higher-order finite elements. *Journal of Computational Physics*, 204(2):520–532, April 2005.
- [36] Mohammad Z. Hossain, Chris D. Cantwell, and Spencer J. Sherwin. A spectral/hp element method for thermal convection. *International Journal for Numerical Methods in Fluids*, 93(7):2380–2395, 2021.
- [37] Xiangmin Jiao, Xuebin Wang, and Qiao Chen. Optimal and Low-Memory Near-Optimal Preconditioning of Fully Implicit Runge–Kutta Schemes for Parabolic PDEs. *SIAM Journal on Scientific Computing*, 43(5):A3527–A3551, January 2021.
- [38] George Em Karniadakis, Moshe Israeli, and Steven A Orszag. High-order splitting methods for the incompressible Navier-Stokes equations. *Journal of Computational Physics*, 97(2):414–443, December 1991.
- [39] Christopher A. Kennedy and Mark H. Carpenter. Additive Runge–Kutta schemes for convection–diffusion–reaction equations. *Applied Numerical Mathematics*, 44(1):139–181, January 2003.
- [40] Christopher A. Kennedy and Mark H. Carpenter. Diagonally Implicit Runge-Kutta Methods for Ordinary Differential Equations. A Review. Technical Report NF1676L-19716, March 2016.
- [41] Christopher A. Kennedy and Mark H. Carpenter. Diagonally implicit Runge–Kutta methods for stiff ODEs. *Applied Numerical Mathematics*, 146:221–244, December 2019.
- [42] David E Keyes, Lois C McInnes, Carol Woodward, William Gropp, Eric Myra, Michael Pernice, John Bell, Jed Brown, Alain Clo, Jeffrey Connors, Emil Constantinescu, Don Estep, Kate Evans, Charbel Farhat, Ammar Hakim, Glenn Hammond, Glen Hansen, Judith Hill, Tobin Isaac, Xiangmin Jiao, Kirk Jordan, Dinesh Kaushik, Efthimios Kaxiras, Alice Koniges, Kihwan Lee, Aaron Lott, Qiming Lu, John Magerlein, Reed Maxwell, Michael McCourt, Miriam Mehl, Roger Pawlowski, Amanda P Randles, Daniel Reynolds, Beatrice Rivière, Ulrich Rüde, Tim Scheibe,

John Shadid, Brendan Sheehan, Mark Shephard, Andrew Siegel, Barry Smith, Xianzhu Tang, Cian Wilson, and Barbara Wohlmuth. Multiphysics simulations: Challenges and opportunities. *The International Journal of High Performance Computing Applications*, 27(1):4–83, February 2013.

- [43] D. A. Knoll and D. E. Keyes. Jacobian-free Newton–Krylov methods: A survey of approaches and applications. *Journal of Computational Physics*, 193(2):357–397, January 2004.
- [44] Heinz-Otto Kreiss and Godela Scherer. Method of Lines for Hyperbolic Differential Equations. *SIAM Journal on Numerical Analysis*, 29(3):640–646, June 1992.
- [45] A. Kværnø. Singly Diagonally Implicit Runge–Kutta Methods with an Explicit First Stage. *BIT Numerical Mathematics*, 44(3):489–502, August 2004.
- [46] J.-L. Lions, Y. Maday, and G. Turinici. A ‘parareal’ in time discretization of pde’s. *Comptes Rendus de l’Académie des Sciences - Series I - Mathematics*, 332:661–668, 2001.
- [47] James W. Lottes and Paul F. Fischer. Hybrid Multigrid/Schwarz Algorithms for the Spectral Element Method. *Journal of Scientific Computing*, 24(1):45–78, July 2005.
- [48] T. A. Manteuffel, S. Münzenmaier, J. Ruge, and B. S. Southworth. Nonsymmetric reduction-based algebraic multigrid. *SIAM Journal on Scientific Computing*, 41(5):S242–S268, 2021/11/17 2019.
- [49] T. A. Manteuffel, J. Ruge, and B. S. Southworth. Nonsymmetric algebraic multigrid based on local approximate ideal restriction (*lAIR*). *SIAM Journal on Scientific Computing*, 40(6):A4105–A4130, 2021/11/17 2018.
- [50] K.-A. Mardal, T. K. Nilssen, and G. A. Staff. Order-Optimal Preconditioners for Implicit Runge–Kutta Schemes Applied to Parabolic PDEs. *SIAM Journal on Scientific Computing*, 29(1):361–375, January 2007.
- [51] Md. Masud Rana, Victoria E. Howle, Katharine Long, Ashley Meeke, and William Milestone. A New Block Preconditioner for Implicit Runge–Kutta Methods for Parabolic PDE Problems. *SIAM Journal on Scientific Computing*, 43(5):S475–S495, 2021.
- [52] Christopher Maynard, Thomas Melvin, and Eike Hermann Müller. Multigrid preconditioners for the mixed finite element dynamical core of the lfriac atmospheric model. *Quarterly Journal of the Royal Meteorological Society*, 146(733):3917–3936, 2022/03/16 2020.
- [53] Eleanor McDonald, Jennifer Pestana, and Andy Wathen. Preconditioning and Iterative Solution of All-at-Once Systems for Evolutionary Partial Differential Equations. *SIAM Journal on Scientific Computing*, 40(2):A1012–A1033, January 2018.
- [54] S. T. Miller, E. C. Cyr, J. N. Shadid, R. M. J. Kramer, E. G. Phillips, S. Conde, and R. P. Pawlowski. IMEX and exact sequence discretization of the multi-fluid plasma model. *Journal of Computational Physics*, 397:108806, November 2019.
- [55] Yvan Notay. An aggregation-based algebraic multigrid method. *Electronic Transactions on Numerical Analysis*, 37(6):123–146, 2010.
- [56] Yvan Notay. Aggregation-Based Algebraic Multigrid for Convection-Diffusion Equations. *SIAM Journal on Scientific Computing*, 34(4):A2288–A2316, January 2012.
- [57] G. Noventa, F. Massa, F. Bassi, A. Colombo, N. Franchina, and A. Ghidoni. A high-order Discontinuous Galerkin solver for unsteady incompressible turbulent flows. *Computers & Fluids*, 139:248–260, November 2016.
- [58] L. N. Olson and J. B. Schroder. PyAMG: Algebraic multigrid solvers in Python v4.0, 2018. Release 4.0.

- [59] Luke Olson. Algebraic Multigrid Preconditioning of High-Order Spectral Elements for Elliptic Problems on a Simplicial Mesh. *SIAM Journal on Scientific Computing*, 29(5):2189–2209, January 2007.
- [60] Steven A Orszag. Spectral methods for problems in complex geometries. *Journal of Computational Physics*, 37(1):70–92, August 1980.
- [61] Yu Pan, Zhen-Guo Yan, Joaquim Peiró, and Spencer J. Sherwin. Development of a Balanced Adaptive Time-Stepping Strategy Based on an Implicit JFNK-DG Compressible Flow Solver. *Communications on Applied Mathematics and Computation*, August 2021.
- [62] Lorenzo Pareschi and Giovanni Russo. Implicit-explicit runge-kutta schemes and applications to hyperbolic systems with relaxation. *Journal of Scientific Computing*, 25(1):129–155, November 2005.
- [63] Luca F. Pavarino. Additive schwarz methods for thep-version finite element method. *Numerische Mathematik*, 66(1):493–515, 1993.
- [64] Will Pazner. Efficient low-order refined preconditioners for high-order matrix-free continuous and discontinuous galerkin methods. *SIAM Journal on Scientific Computing*, 42(5):A3055–A3083, 2022/03/17 2020.
- [65] Will Pazner and Tzanio Kolev. Uniform Subspace Correction Preconditioners for Discontinuous Galerkin Methods with hp-Refinement. *Communications on Applied Mathematics and Computation*, July 2021.
- [66] Will Pazner, Tzanio Kolev, and Clark Dohrmann. Low-order preconditioning for the high-order finite element de rham complex, 2022.
- [67] Will Pazner and Per-Olof Persson. Stage-parallel fully implicit Runge–Kutta solvers for discontinuous Galerkin fluid simulations. *Journal of Computational Physics*, 335:700–717, April 2017.
- [68] A. Ramage. A multigrid preconditioner for stabilised discretisations of advection–diffusion problems. *Journal of Computational and Applied Mathematics*, 110(1):187–203, October 1999.
- [69] J. M. Reynolds-Barredo, D. E. Newman, R. Sanchez, D. Samaddar, L. A. Berry, and W. R. Elwasif. Mechanisms for the convergence of time-parallelized, parareal turbulent plasma simulations. *Journal of Computational Physics*, 231(23):7851–7867, October 2012.
- [70] Jose M. Reynolds Barredo, David E. Newman, Raul Sanchez, Debasmita Samaddar, Lee A. Berry, and Wael Elwasif. Toward the application of the Parareal algorithm to 5D gyrokinetic plasma turbulence. page TP9.053, October 2011.
- [71] D. Samaddar, T. A. Casper, S. H. Kim, L. A. Berry, W. R. Elwasif, D. Batchelor, and W. A. Houlberg. Time parallelization of advanced operation scenario simulations of ITER plasma. *Journal of Physics: Conference Series*, 410:012032, February 2013.
- [72] D. Samaddar, D. P. Coster, X. Bonnin, C. Bergmeister, E. Havlíčková, L. A. Berry, W. R. Elwasif, and D. B. Batchelor. Temporal parallelization of edge plasma simulations using the parareal algorithm and the SOLPS code. *Computer Physics Communications*, 221:19–27, December 2017.
- [73] D. Samaddar, D. P. Coster, X. Bonnin, L. A. Berry, W. R. Elwasif, and D. B. Batchelor. Application of the parareal algorithm to simulations of ELMs in ITER plasma. *Computer Physics Communications*, 235:246–257, February 2019.
- [74] D. Samaddar, D. E. Newman, and R. Sánchez. Parallelization in time of numerical simulations of fully-developed plasma turbulence using the parareal algorithm. *Journal of Computational Physics*, 229(18):6558–6573, September 2010.

- [75] V. Simoncini. Computational Methods for Linear Matrix Equations. *SIAM Review*, 58(3):377–441, 2016.
- [76] Ben S. Southworth, Oliver Krzysik, and Will Pazner. Fast solution of fully implicit runge–kutta and discontinuous galerkin in time for numerical pdes, part ii: Nonlinearities and daes. *SIAM Journal on Scientific Computing*, 44(2):A636–A663, 2022.
- [77] Ben S. Southworth, Oliver Krzysik, Will Pazner, and Hans De Sterck. Fast Solution of Fully Implicit Runge–Kutta and Discontinuous Galerkin in Time for Numerical PDEs, Part I: the Linear Setting. *SIAM Journal on Scientific Computing*, 44(1):A416–A443, 2022.
- [78] Ben S. Southworth, Wayne Mitchell, Andreas Hestenthaler, and Federico Danieli. Tight Two-Level Convergence of Linear Parareal and MGRIT: Extensions and Implications in Practice. In Benjamin Ong, Jacob Schroder, Jemma Shipton, and Stephanie Friedhoff, editors, *Parallel-in-Time Integration Methods*, pages 1–31, Cham, 2021. Springer International Publishing.
- [79] N. Spillane, V. Dolean, P. Hauret, F. Nataf, C. Pechstein, and R. Scheichl. Abstract robust coarse spaces for systems of PDEs via generalized eigenproblems in the overlaps. *Numerische Mathematik*, 126(4):741–770, 2014.
- [80] Hari Sundar, Georg Stadler, and George Biros. Comparison of multigrid algorithms for high-order continuous finite element discretizations. *Numerical Linear Algebra with Applications*, 22(4):664–680, 2015.
- [81] U Trottenberg, C Oosterlee, and A Schüller. *Multigrid*. Academic Press, 2000.
- [82] Peter E.J. Vos, Claes Eskilsson, Alessandro Bolis, Sehun Chun, Robert M. Kirby, and Spencer J. Sherwin. A generic framework for time-stepping partial differential equations (PDEs): General linear methods, object-oriented implementation and application to fluid problems. *International Journal of Computational Fluid Dynamics*, 25(3):107–125, March 2011.
- [83] A. J. Wathen. Preconditioning. *Acta Numerica*, 24:329–376, May 2015.
- [84] Zhen-Guo Yan, Yu Pan, Giacomo Castiglioni, Koen Hillewaert, Joaquim Peiró, David Moxey, and Spencer J. Sherwin. Nektar++: Design and implementation of an implicit, spectral/hp element, compressible flow solver using a Jacobian-free Newton Krylov approach. *Computers & Mathematics with Applications*, 81:351–372, January 2021.