Monolithic multigrid for implicit Runge–Kutta discretizations of incompressible fluid flow[⋆]

Razan Abu-Labdeh[a,∗], Scott MacLachlan[a], Patrick E. Farrell[b]

[a]*Department of Mathematics and Statistics, Memorial University of Newfoundland,
St. John's, NL, Canada*
[b]*Mathematical Institute, University of Oxford, Oxford, UK*

Abstract

Most research on preconditioners for time-dependent PDEs has focused on implicit multi-step or diagonally-implicit multi-stage temporal discretizations. In this paper, we consider monolithic multigrid preconditioners for fully-implicit multi-stage Runge–Kutta (RK) time integration methods. These temporal discretizations have very attractive accuracy and stability properties, but they couple the spatial degrees of freedom across multiple time levels, requiring the solution of very large linear systems. We extend the classical Vanka relaxation scheme to implicit RK discretizations of saddle point problems. We present numerical results for the incompressible Stokes, Navier–Stokes, and resistive magnetohydrodynamics equations, in two and three dimensions, confirming that these relaxation schemes lead to robust and scalable monolithic multigrid methods for a challenging range of incompressible fluid-flow models.

Keywords: Implicit Runge–Kutta time integration, Monolithic multigrid, Newton–Krylov–multigrid Methods

1. Introduction

Among the many applications of advanced computer simulation, models of fluid flow have been a persistent and common driving force in research and practice. The history of spatial discretization of fluid problems dates back at least to the 1960's (e.g., the MAC-scheme discretization of the Navier–Stokes equations [1–3]), but continues to this day with investigation of higher-order mixed finite-element discretizations for both Newtonian and complex fluids [4–11]. Alongside this thrust to higher-order spatial discretizations comes a need for stable higher-order temporal discretizations, for which implicit Runge-Kutta methods are a natural choice. In this paper, we investigate the development of efficient Newton–Krylov–multigrid strategies for implicit Runge–Kutta discretizations of incompressible fluid-flow problems.

Effective solver strategies for both stationary problems and time-dependent flow models discretized via either multi-step schemes or diagonally implicit Runge-Kutta (DIRK) schemes have been studied for many years. For time-dependent Newtonian flows, both fully and semi-implicit pressure-correction schemes (e.g., [2, 3, 12–14]) have been proposed, based primarily on multigrid solution of the pressure-Poisson equation, but the construction and analysis of general high-order schemes is non-trivial [15]. Monolithic multigrid schemes (both linear and nonlinear) have also been broadly considered, first arising in the late 1970's and early 1980's [16, 17]. More approaches have been proposed since these early works, including techniques for Newtonian flows based on Vanka [18] and Braess–Sarazin [19] relaxation, and generalizations of these techniques to more complex flow settings and discretizations [20–22]. Simultaneously, block preconditioning strategies have also been developed, for a variety of discretizations and flow settings [23–28]. Despite this substantial body of work on multi-step methods, there are (to our knowledge) few comparable publications on solution strategies for multi-stage implicit Runge–Kutta (IRK) discretizations of flow models [29, 30].

A small body of work exists on solvers for IRK discretizations for parabolic

PDEs [30–38]. Much of this work focuses on block-structured preconditioners for the tensor-product systems generated by IRK discretization [34–37] where, for example, standard multigrid methods can be used to solve the diagonal blocks. The recent method of Southworth et al. [30, 38] appears to be very effective, again leveraging standard preconditioners for linear systems corresponding to BDF-type discretizations. On the other hand, the work of Vandewalle and others [31–33] applies monolithic multigrid methods to these discretizations, using block-Gauss–Seidel type relaxation for parabolic equations and a block-extension of the Hiptmair relaxation [39] for the eddy-current form of the curl-curl equation. Similar block-Jacobi relaxation was used for both the heat and Gross–Pitaevskii equations in [40]. Here, we investigate extensions of Vanka relaxation for IRK discretizations of fluid flow problems.

In this paper, we consider standard mixed finite-element (spatial) discretizations of Stokes, Navier–Stokes, and magnetohydrodynamic (MHD) flows, coupled with IRK discretizations in time. We focus on the development of monolithic geometric multigrid preconditioners for the coupled systems of equations to be solved at each timestep. For nonlinear problems, we use these preconditioners in a standard Newton–Krylov–multigrid setting, using Newton's method to linearize the coupled nonlinear systems at each timestep. We expect the same techniques would apply to the various simplifications of Newton's method that are applicable in the IRK context [41, 42]. Numerical results are presented for standard benchmarks in two and three spatial dimensions, showing that this solution approach is equally effective for IRK discretizations as it is for BDF discretizations.

The remainder of this paper is organized as follows. In Section 2, we review the Runge–Kutta discretization approach for systems of ODEs. For fluid-flow models, this is typically used in a method-of-lines approach with some spatial discretization, and Section 3 reviews mixed finite-element discretization of the Stokes, Navier–Stokes, and MHD models considered here. In Section 4, we present the constituent parts of the monolithic multigrid algorithm that we propose for solution of the resulting linear(ized) systems of equations. Numerical

results that confirm the effectiveness of this approach are given in Section 5. Finally, conclusions and directions for future work are given in Section 6.

2. Runge–Kutta temporal discretizations

While BDF (and other linear multi-step) schemes can achieve higher-order convergence, they do so at a cost to their stability, with the widely known result that no linear multi-step scheme with order greater than two can be A-stable (the so-called Second Dahlquist Barrier) [43]. Because of this (and other reasons), Runge–Kutta integrators are widely used when we seek higher-order time integration methods. In contrast to multi-step schemes (where solutions at past time-steps are used in the approximation), Runge–Kutta methods are *multi-stage* schemes, where a number of intermediate stage values are used to achieve the approximation. In general, an r-stage Runge–Kutta method applied to the system of ordinary differential equations $u'(t) = f(u(t), t)$ is given by

$$
k_i = f\left(u^n + \Delta t \sum_{j=1}^{r} a_{ij} k_j, t^n + c_i \Delta t\right), \text{ for } i = 1, 2, \dots, r,
$$

$$
u^{n+1} = u^n + \Delta t \sum_{j=1}^{r} b_j k_j.
$$

(1)

The coefficients in the scheme are the stage times (or nodes) c_i, the weights b_j, and the Runge–Kutta matrix $A = [a_{ij}]$. Taken together, these form the Butcher tableau for a given scheme [41, 44]. For consistency, we require that $\sum_{j=1}^{r} b_j = 1$ and $\sum_{j=1}^{r} a_{ij} = c_i$, for all $i = 1, 2, \dots, r$. The r stage values are represented by the set $\{k_i\}_{i=1}^{r}$ and the approximation at time $t^n = t^0 + n\Delta t$ is denoted by u^n.

Runge–Kutta methods are generally classified by the non-zero pattern of the matrix A. Methods can be explicit, with $a_{ij} = 0 \ \forall j \geq i$, or implicit, when $\exists j \geq i$ with $a_{ij} \neq 0$. The implicit methods can further be classified into diagonally implicit, with $a_{ij} = 0 \ \forall j > i$, or fully implicit, when $\exists j > i$ such that $a_{ij} \neq 0$. Further specialization is also possible, such as singly diagonally implicit Runge–Kutta (SDIRK) methods, which are diagonally implicit (DIRK)

4

methods with the added property that $a_{ii} = a_{jj}$ for all i and j, and explicit singly diagonally implicit (ESDIRK) methods, which have an all-zero first row of A, followed by SDIRK structure on lower rows (of which the Crank–Nicolson scheme is a well-known example).

There are three main points to consider when choosing a Runge–Kutta method, regarding its stability, accuracy, and computational cost. For any scheme, we define the function $r(z)$ as the map produced when applying the scheme to the (scalar, linear) Dahlquist test problem, $u' = \lambda u$ for $\lambda \in \mathbb{C}$, with $u^{n+1} = r(\lambda \Delta t) u^n$. The domain of stability of the scheme is defined as the region in the complex plane where $|r(z)| < 1$. RK methods are said to be *A-stable* if their domain of stability includes the entire left-half of the complex plane. If, additionally, we have that $\lim_{z \to -\infty} |r(z)| = 0$, we say that the scheme is *L-stable*. For many applications, L-stability is the preferred property, since an L-stable scheme generally damps non-physical high-frequency oscillations that may pollute a numerical solution. As is typical, explicit Runge–Kutta (ERK) methods have finite regions of stability, and only implicit Runge–Kutta (IRK) schemes can be A- or L-stable.

The local truncation error of an RK scheme is defined as the error made in a single step of the scheme, starting with the analytical solution of the differential equation as u^n, compared to $u(t^{n+1})$, while the global error is the accumulated error in the approximate solution over the timesteps needed to reach a fixed time. We typically discuss such errors by their order, meaning that we bound the error by a constant (depending on $f(u, t)$ and the analytical solution, $u(t)$) times $(\Delta t)^p$ to establish that a scheme has order p. Typically (e.g., when $f(u, t)$ is continuous in t and Lipschitz continuous in u), the global error is one order less than the local truncation error. A well-known result is that the order of global error of an ERK method cannot be greater than its number of stages (and, to achieve order $p > 5$, an ERK scheme must have at least $p + 1$ stages) [45, Section 324]. In contrast, the maximum order of global error for an IRK discretization can be as much as twice the number of stages in the scheme. While higher-order global error is attractive, for both stiff DEs and systems of

differential-algebraic equations (DAEs), the so-called *stage order* of a Runge–Kutta method is more important [45, Section 362]. Here, in addition to the truncation error, the accuracy of a scheme is determined by also bounding the approximation of stage i to $u(t^n + c_i \Delta t)$ by some constant (depending on $f(u, t)$ and $u(t)$) times $(\Delta t)^{q+1}$, thus defining the stage order as $\min\{q, p\}$. For index-2 DAEs (as are considered here), the order of accuracy of a scheme is limited by its stage order, due to perturbation bounds on the solution of the constrained system [43, Section VII.4]. This greatly limits our choice of schemes that allow higher-order accuracy. While DIRK methods can have reasonable global order, their stage order is typically limited to 1. We note that ESDIRK methods are an exception to this, with stage order limited to 2, due to the structure of their Butcher tableau. In contrast, the stage order of fully IRK schemes can be as large as the number of stages, making these the preferred schemes for integrating DAEs.

The downside of IRK schemes is their computational cost. ERK methods can be implemented at the cost of one evaluation of $f(u, t)$ for each stage in the method. In contrast, IRK methods require solution of a system of equations for each timestep (that may be large when u represents a spatially discretized approximation to the solution of a PDE). Herein lies the attraction of DIRK, SDIRK, and ESDIRK schemes. In these approaches, rather than having to solve for the stages in a coupled manner, each stage can be solved for sequentially, allowing the reuse of standard linear and nonlinear solvers from backward-Euler type schemes. SDIRK and ESDIRK afford even more of an advantage, particularly in the linear case, as the same solver architecture can be directly reused in the solution process for each stage. General IRK methods, in contrast, do not allow this simplification. While block-preconditioning strategies can be used to again leverage existing solver architectures from the multistep case [30, 34–38], these theoretical results tend to be limited to simple cases, excluding (for example) nonlinear systems of DAEs, as arise in standard models of computational fluid dynamics.

In this paper, we consider a standard Newton–Krylov–multigrid framework

6

for the solution of the nonlinear systems of equations that arise from using general IRK discretizations for the Navier–Stokes equations and the equations of magnetohydrodynamics. Because the details of these solvers depend directly on the spatial discretization, we next discuss the mixed finite-element discretization of these models.

3. Discretization of fluid models

In this section, we consider the interplay of mixed finite-element spatial discretization for incompressible models of fluid flow with temporal discretization by IRK methods. We consider three models: the linear Stokes model, the (nonlinear) Navier–Stokes equations, and the equations of single-fluid visco-resistive incompressible magnetohydrodynamics (MHD). In Section 4, we will focus on the development of a monolithic multigrid methodology for the linearized systems that result from applying Newton's method to the nonlinear problems. Both here and in that exposition, we will focus on the details of the algorithm for the simplest case of the linear Stokes model.

3.1. Time-dependent Stokes equations

In the viscous limit of incompressible flow, inertial forces in the model can be neglected, leading to the time-steady Stokes equations. We consider here the time-dependent analogue of the Stokes equations on a bounded Lipschitz domain $\Omega \subset \mathbb{R}^d$, $d \in \{2, 3\}$:

$$\rho \mathbf{u}_t - \mu \Delta \mathbf{u} + \nabla p = \mathbf{f} \qquad \text{in } \Omega \times (0, T_f) \tag{2a}$$

$$-\nabla \cdot \mathbf{u} = 0 \qquad \text{in } \Omega \times (0, T_f), \tag{2b}$$

$$\mathbf{u} = 0 \qquad \text{on } \partial\Omega \times (0, T_f), \tag{2c}$$

$$\mathbf{u}(\mathbf{x}, 0) = \mathbf{g}(\mathbf{x}) \text{ on } \Omega \times \{t = 0\}, \tag{2d}$$

where $\mathbf{u}(\mathbf{x}, t)$ is the velocity, $p(\mathbf{x}, t)$ is the pressure, and $\mathbf{f}(\mathbf{x}, t)$ is a suitably smooth forcing term. Here, ρ denotes the fluid density and μ denotes the fluid viscosity; we set both to 1 for simplicity. The final time is denoted by T_f. Since

no time derivative of the pressure appears in the system, it is a DAE. The *index* of a DAE is defined as the number of analytical differentiations needed (along with algebraic manipulations) to convert the DAE into an explicit system of ODEs [43, Section VII.1]. Here, since the constraint equation is of the form $-\nabla \cdot \mathbf{u} = 0$, this is an index-two DAE, since one differentiation of the constraint (and applying the divergence to (2a)) allows us to explicitly solve for p in terms of $\mathbf{u}$, and a second gives an ODE for p. For index-two DAEs, the order of accuracy of a Runge–Kutta time-discretization is limited to the stage order of the scheme.

For the spatial discretization of (2), we use the mixed finite-element framework, considering the stable Taylor–Hood discretization on simplices [24]. Let $\mathcal{V} = \mathbf{H}_0^1(\Omega)$, where $\mathbf{H}_0^1(\Omega) = \{\mathbf{v} \in \mathbf{H}^1(\Omega) : \mathbf{u} = 0 \text{ on } \partial\Omega\}$, and $\mathcal{W} = L_0^2(\Omega)$ (the space of zero-mean functions in $L^2(\Omega)$), and consider a weak solution of (2) that is (at least) once continuously differentiable in time and such that for every $t \in (0, T_f)$, $\mathbf{u}(\cdot, t) \in \mathcal{V}$ and $p(\cdot, t) \in \mathcal{W}$. Multiplying the time-dependent equation by $\mathbf{v} \in \mathcal{V}$ and the divergence constraint by $q \in \mathcal{W}$ and integrating by parts, we get the weak form

$$\langle \mathbf{u}_t, \mathbf{v} \rangle + \langle \nabla \mathbf{u}, \nabla \mathbf{v} \rangle - \langle p, \nabla \cdot \mathbf{v} \rangle = \langle \mathbf{f}, \mathbf{v} \rangle, \quad \forall \mathbf{v} \in \mathcal{V},$$

$$-\langle q, \nabla \cdot \mathbf{u} \rangle = 0, \qquad \forall q \in \mathcal{W},$$

where the inner-product notation, $\langle \cdot, \cdot \rangle$, denotes integration in space but not time. The finite-element discretization is realized by constructing a triangulation, τ_h, of Ω, and approximating $\mathbf{u}$ and p in piecewise polynomial spaces defined over τ_h. Here, we use standard continuous Lagrange finite-element spaces, defining

$$P_k(\Omega, \tau_h) = \left\{ u \in C^0(\Omega) : \forall T \in \tau_h, \, u|_T(\mathbf{x}) \text{ is a polynomial of degree no more than } k \right\}.$$

We consider the standard stable Taylor–Hood discretization, with $\mathcal{V}_h = (P_2(\Omega, \tau_h))^d \cap \mathcal{V}$ and $\mathcal{W}_h = P_1(\Omega, \tau_h) \cap \mathcal{W}$ [24, 46]. This leads to the semi-discretized weak

form of finding $(\mathbf{u}(\cdot, t), p(\cdot, t)) \in \mathcal{V}_h \times \mathcal{W}_h$ such that

$$\langle \mathbf{u}_t, \mathbf{v} \rangle + \langle \nabla \mathbf{u}, \nabla \mathbf{v} \rangle - \langle p, \nabla \cdot \mathbf{v} \rangle = \langle \mathbf{f}, \mathbf{v} \rangle, \quad \forall \mathbf{v} \in \mathcal{V}_h,$$

$$-\langle q, \nabla \cdot \mathbf{u} \rangle = 0, \qquad \forall q \in \mathcal{W}_h.$$

Now writing $\vec{\mathbf{u}}(t)$ and $\vec{p}(t)$ for the (time-dependent) coefficients of $\mathbf{u}(\mathbf{x}, t)$ and $p(\mathbf{x}, t)$ in the finite-element basis, we can write this as a coupled linear system of DAEs, as

$$\begin{bmatrix} M\vec{\mathbf{u}}_t \\ 0 \end{bmatrix} + \begin{bmatrix} K & B \\ B^T & 0 \end{bmatrix} \begin{bmatrix} \vec{\mathbf{u}} \\ \vec{p} \end{bmatrix} = \begin{bmatrix} M\vec{\mathbf{f}} \\ 0 \end{bmatrix},$$

where $\vec{\mathbf{f}}$ is the vector of coefficients of the interpolant of $\mathbf{f}$ in $\mathcal{V}_h$. Here, M and K are the $(P_2(\Omega, \tau_h))^d$ mass and stiffness matrices, respectively, while B is the weak gradient operator mapping from $\mathcal{W}_h$ into $\mathcal{V}_h$.

It is this system of equations that we discretize using Runge–Kutta methods. As the system is a set of DAEs, and not ODEs, we cannot directly apply (1), but use its DAE analogue [43], writing

$$\vec{\mathbf{u}}_i^n = \vec{\mathbf{u}}^n + \Delta t \sum_{j=1}^{r} a_{ij} \vec{k}_j^{(\mathbf{u})}, \qquad \vec{p}_i^n = \vec{p}^n + \Delta t \sum_{j=1}^{r} a_{ij} \vec{k}_j^{(p)},$$

$$M\vec{k}_i^{(\mathbf{u})} + K\vec{\mathbf{u}}_i^n + B\vec{p}_i^n = M\vec{\mathbf{f}}_i^n, \qquad B^T \vec{\mathbf{u}}_i^n = 0,$$

$$\vec{\mathbf{u}}^{n+1} = \vec{\mathbf{u}}^n + \Delta t \sum_{j=1}^{r} b_j \vec{k}_j^{(\mathbf{u})}, \qquad \vec{p}^{n+1} = \vec{p}^n + \Delta t \sum_{j=1}^{r} b_j \vec{k}_j^p,$$

where $\vec{\mathbf{f}}_i^n$ is the interpolant of $\mathbf{f}$ in $\mathcal{V}_h$ at time $t^n + c_i \Delta t$, $\vec{\mathbf{u}}_i^n$ and $\vec{p}_i^n$ are the approximations of $\vec{\mathbf{u}}$ and $\vec{p}$ at time $t^n + c_i \Delta t$, and $\vec{k}_i^{(\mathbf{u})}$ and $\vec{k}_i^{(p)}$ are the RK stages for which we solve. Rewriting the equations for $\vec{k}_i^{(\mathbf{u})}$ and $\vec{k}_i^{(p)}$, we have

$$M\vec{k}_i^{(\mathbf{u})} + K \left(\vec{\mathbf{u}}^n + \Delta t \sum_{j=1}^{r} a_{ij} \vec{k}_j^{(\mathbf{u})} \right) + B \left(\vec{p}^n + \Delta t \sum_{j=1}^{r} a_{ij} \vec{k}_j^{(p)} \right) = M\vec{\mathbf{f}}_i^n$$

$$B^T \left(\vec{\mathbf{u}}^n + \Delta t \sum_{j=1}^{r} a_{ij} \vec{k}_j^{(\mathbf{u})} \right) = 0$$

or

$$M\vec{k}_i^{(\mathbf{u})} + \Delta t \sum_{j=1}^{r} a_{ij} \left(K\vec{k}_j^{(\mathbf{u})} + B\vec{k}_j^{(p)} \right) = M\vec{\mathbf{f}}_i^n - K\vec{\mathbf{u}}^n - B\vec{p}^n$$

$$\Delta t \sum_{j=1}^{r} a_{ij} B^T \vec{k}_j^{(\mathbf{u})} = -B^T \vec{\mathbf{u}}^n$$

for $1 \leq i \leq r$. The matrix on the left can easily be written in tensor-product form, leading to a concise description of the scheme as

$$\left(\mathbf{I}_r \otimes \begin{bmatrix} M & 0 \\ 0 & 0 \end{bmatrix} + \Delta t A \otimes \begin{bmatrix} K & B \\ B^T & 0 \end{bmatrix} \right) \vec{\mathbf{k}} = \vec{\mathbf{F}}, \tag{3}$$

where $\vec{\mathbf{k}}$ is the vector of stages, ordered consecutively by stage index i, keeping the ordering of $(\vec{k}_i^{(\mathbf{u})}, \vec{k}_i^{(p)})$ pairs together, and $\vec{\mathbf{F}}$ is the corresponding vector of right-hand sides (including terms from timestep n).

3.2. Navier–Stokes Equations

We next include the full inertial term, leading to the nonlinear incompressible Navier–Stokes equations,

$$\rho \left(\mathbf{u}_t + \mathbf{u} \cdot \nabla \mathbf{u} \right) - \mu \Delta \mathbf{u} + \nabla p = \mathbf{f} \qquad \text{in } \Omega \times (0, T_f), \tag{4a}$$

$$-\nabla \cdot \mathbf{u} = 0 \qquad \text{in } \Omega \times (0, T_f), \tag{4b}$$

$$\mathbf{u} = 0 \qquad \text{on } \partial\Omega \times (0, T_f), \tag{4c}$$

$$\mathbf{u}(\mathbf{x}, 0) = \mathbf{g}(\mathbf{x}) \text{ on } \Omega \times \{t = 0\}. \tag{4d}$$

We again take the density to be 1, but will allow the viscosity μ to be chosen differently, to consider problems at different Reynolds numbers. The additional term passes directly to the weak form, which we again discretize using a Taylor–Hood mixed finite-element discretization. The semi-discretized weak variational form of (4) is to find $(\mathbf{u}(\cdot, t), p(\cdot, t)) \in \mathcal{V}_h \times \mathcal{W}_h$ such that

$$\langle \mathbf{u}_t, \mathbf{v} \rangle + \langle \mathbf{u} \cdot \nabla \mathbf{u}, \mathbf{v} \rangle + \mu \langle \nabla \mathbf{u}, \nabla \mathbf{v} \rangle - \langle p, \nabla \cdot \mathbf{v} \rangle = 0,$$
$$-\langle \nabla \cdot \mathbf{u}, q \rangle = 0, \tag{5}$$

for all test functions $(\mathbf{v}, q) \in \mathcal{V}_h \times \mathcal{W}_h$. Again writing $\vec{\mathbf{u}}(t)$ and $\vec{p}(t)$ for the coefficients of $\mathbf{u}$ and p in the finite-element basis, this leads to a nonlinear coupled system of DAEs, as

$$\begin{bmatrix} M\vec{\mathbf{u}}_t \\ 0 \end{bmatrix} + \begin{bmatrix} N(\vec{\mathbf{u}}) \\ 0 \end{bmatrix} + \begin{bmatrix} K & B \\ B^T & 0 \end{bmatrix} \begin{bmatrix} \vec{\mathbf{u}} \\ \vec{p} \end{bmatrix} = \begin{bmatrix} M\vec{\mathbf{f}} \\ 0 \end{bmatrix},$$

where $N(\vec{\mathbf{u}})$ represents the discretization of $\langle \mathbf{u} \cdot \nabla \mathbf{u}, \mathbf{v} \rangle$. As above, accounting for this term in the RK stage equations leads to the nonlinear coupled system

$$M\vec{k}_i^{(\mathbf{u})} + N\left(\vec{\mathbf{u}}^n + \Delta t \sum_{j=1}^{r} a_{ij} \vec{k}_j^{(\mathbf{u})} \right) + \Delta t \sum_{j=1}^{r} a_{ij} \left(K\vec{k}_j^{(\mathbf{u})} + B\vec{k}_j^{(p)} \right) = M\vec{\mathbf{f}}_i^n - K\vec{\mathbf{u}}^n - B\vec{p}^n,$$

$$\Delta t \sum_{j=1}^{r} a_{ij} B^T \vec{k}_j^{(\mathbf{u})} = -B^T \vec{\mathbf{u}}^n,$$

for $1 \le i \le r$. This system is solved using Newton's method.

Denoting the nonlinear system as $F\left(\vec{\mathbf{k}}^n \right) = 0$, a standard Newton approximation would be to solve

$$F\left(\vec{\mathbf{k}}^{n,\ell+1} \right) \approx F\left(\vec{\mathbf{k}}^{n,\ell} \right) + J\left(\vec{\mathbf{k}}^{n,\ell} \right) \delta\vec{\mathbf{k}}^{n,\ell} = 0,$$

where $J\left(\vec{\mathbf{k}}^{n,\ell} \right)$ is the Jacobian of the system at the current approximation, $\vec{\mathbf{k}}^{n,\ell}$ and $\delta\vec{\mathbf{k}}^{n,\ell} := \vec{\mathbf{k}}^{n,\ell+1} - \vec{\mathbf{k}}^{n,\ell}$ is the Newton search direction. Since we are timestepping, we use the computed solution at the previous time-step, $\vec{\mathbf{k}}^{n-1}$, for the initial guess for the stage values at step n, $\vec{\mathbf{k}}^{n,0}$. In this work, we use the Eisenstat–Walker stopping criterion for the Krylov iteration to solve for $\delta\vec{\mathbf{k}}^{n,\ell}$ [47], requiring that

$$\left\| F\left(\vec{\mathbf{k}}^{n,\ell} \right) + J\left(\vec{\mathbf{k}}^{n,\ell} \right) \delta\vec{\mathbf{k}}^{n,\ell} \right\| \le \eta_\ell \left\| F\left(\vec{\mathbf{k}}^{n,\ell} \right) \right\|,$$

for every step, ℓ, where $\eta_\ell \in [0, 1)$ is updated for each nonlinear iteration based on convergence of the method.

3.3. Magnetohydrodynamics

Finally, we consider the equations of single-fluid viscoresistive magnetohydrodynamics (MHD). In general, MHD models the flow of conducting fluids in

the presence of an electromagnetic field. These models are nonlinear and contain strong coupling between the fluid velocity and the electromagnetic variables. We follow the MHD formulation presented in [22],

$$\mathbf{u}_t + (\mathbf{u} \cdot \nabla)\mathbf{u} - \nabla \cdot (\frac{2}{\text{Re}}\epsilon(\mathbf{u})) + \nabla p - (\nabla \times \mathbf{B}) \times \mathbf{B} = \mathbf{f_u} \text{ in } \Omega \times (0, T_f), \quad (6a)$$

$$\mathbf{B}_t + \frac{1}{\text{Re}_m}\nabla \times \nabla \times \mathbf{B} - \nabla \times (\mathbf{u} \times \mathbf{B}) - \nabla \gamma = \mathbf{f_B} \text{ in } \Omega \times (0, T_f), \quad (6b)$$

$$-\nabla \cdot \mathbf{u} = 0 \text{ in } \Omega \times (0, T_f), \quad (6c)$$

$$\nabla \cdot \mathbf{B} = 0 \text{ in } \Omega \times (0, T_f), \quad (6d)$$

$$\mathbf{u} = 0 \text{ on } \partial\Omega \times (0, T_f), \quad (6e)$$

$$\mathbf{B} \times \mathbf{n} = 0 \text{ on } \partial\Omega \times (0, T_f), \quad (6f)$$

$$\mathbf{u}(\mathbf{x}, 0) = \mathbf{g_u}(\mathbf{x}) \text{ on } \Omega \times \{t = 0\}, \quad (6g)$$

$$\mathbf{B}(\mathbf{x}, 0) = \mathbf{g_B}(\mathbf{x}) \text{ on } \Omega \times \{t = 0\}, \quad (6h)$$

where the four unknowns are the velocity vector, $\mathbf{u}$, the pressure, p, the magnetic field, $\mathbf{B}$, and the Lagrange multiplier, γ. The Lagrange multiplier is used to enforce the solenoidal condition (6d), while the pressure is used to enforce the incompressibility condition (6c). The strain-rate tensor is $\epsilon(\mathbf{u}) = \frac{1}{2}(\nabla\mathbf{u} + \nabla\mathbf{u}^T)$, and the two dimensionless constants, Re and Re_m, are the hydrodynamic Reynolds number and magnetic Reynolds number, respectively. We consider this equation in both two- and three-dimensional domains, Ω; in 2D, the curl and cross-product are defined by the natural extensions from two-dimensional vector fields to three-dimensional fields.

Here, for $\Omega \subset \mathbb{R}^d$, we take

$$(\mathbf{u}(\cdot, t), \mathbf{B}(\cdot, t), p(\cdot, t), \gamma(\cdot, t)) \in \mathbf{H}_0^1(\Omega) \times \mathbf{H}_0(\text{curl}, \Omega) \times L_0^2(\Omega) \times H_0^1(\Omega),$$

where

$$\mathbf{H}_0^1(\Omega) = \{\mathbf{v} \in \mathbf{H}^1(\Omega) : \mathbf{u} = 0 \text{ on } \partial\Omega\},$$

$$\mathbf{H}_0(\text{curl}, \Omega) = \{\mathbf{c} \in \mathbf{L}^2(\Omega) : \nabla \times \mathbf{c} \in \mathbf{L}^2(\Omega), \mathbf{n} \times \mathbf{c} = 0 \text{ on } \partial\Omega\},$$

$$L_0^2(\Omega) = \{q \in L^2(\Omega) : \int_\Omega q \, d\mathbf{x} = 0\},$$

$$H_0^1(\Omega) = \{s \in H^1(\Omega) : s = 0 \text{ on } \partial\Omega\},$$

and $\mathbf{n}$ is the outward unit normal vector on $\partial\Omega$ [22, 48]. We discretize the fluid variables again with the Taylor–Hood discretization $\mathcal{V}_h \times \mathcal{W}_h \subset \mathbf{H}_0^1(\Omega) \times L_0^2(\Omega)$, and use lowest-order Nédélec elements for $\mathcal{C}_h \subset \mathbf{H}_0(\text{curl}, \Omega)$ and $\mathcal{S}_h = P_1(\Omega, \tau_h) \cap H_0^1(\Omega)$ for the Lagrange multiplier. Well-posedness (under small-data assumptions) of both the continuous and discrete formulations is shown in [48].

Multiplying (6) by the test functions $(\mathbf{v}, \mathbf{c}, q, s) \in \mathcal{V}_h \times \mathcal{C}_h \times \mathcal{W}_h \times \mathcal{S}_h$ and integrating by parts, we get the semi-discretized weak variational form of finding $(\mathbf{u}(\cdot, t), \mathbf{B}(\cdot, t), p(\cdot, t), \gamma(\cdot, t)) \in \mathcal{V}_h \times \mathcal{C}_h \times \mathcal{W}_h \times \mathcal{S}_h$ such that

$$\int_\Omega \mathbf{u}_t \cdot \mathbf{v} + ((\mathbf{u} \cdot \nabla)\mathbf{u}) \cdot \mathbf{v} - \frac{2}{\text{Re}} \left(\epsilon(\mathbf{u}) : \epsilon(\mathbf{v})\right) - p\nabla \cdot \mathbf{v} - ((\nabla \times \mathbf{B}) \times \mathbf{B}) \cdot \mathbf{v} \, d\mathbf{x}$$

$$= \int_\Omega \mathbf{f_u} \cdot \mathbf{v} \, d\mathbf{x},$$

$$\int_\Omega \mathbf{B}_t \cdot \mathbf{c} + \frac{1}{\text{Re}_m}(\nabla \times \mathbf{B}) \cdot (\nabla \times \mathbf{c}) - (\mathbf{u} \times \mathbf{B}) \cdot (\nabla \times \mathbf{c}) - (\nabla\gamma) \cdot \mathbf{c} \, d\mathbf{x} = \int_\Omega \mathbf{f_B} \cdot \mathbf{c} \, d\mathbf{x},$$

$$- \int_\Omega (\nabla \cdot \mathbf{u}) q \, d\mathbf{x} = 0,$$

$$- \int_\Omega \mathbf{B} \cdot (\nabla s) \, d\mathbf{x} = 0,$$

$$\tag{7}$$

for all $(\mathbf{v}, \mathbf{c}, q, s) \in \mathcal{V}_h \times \mathcal{C}_h \times \mathcal{W}_h \times \mathcal{S}_h$. The corresponding IRK discretization is derived from this semi-discretized form as described above, and solved in the same Newton–Krylov–multigrid manner, using the Eisenstat–Walker stopping criterion for the Krylov iteration. We note that linear solvers for this spatial discretization using BDF2 in time was the subject of [22]; given the Dahlquist barrier, and the driving need for L-stability, multistage schemes, such as IRK, are needed to achieve higher-order time integration for this problem.

4. Monolithic multigrid for fluid problems

As described above, we use a Newton–Krylov–multigrid framework for solving the non-linear systems of equations resulting from spatial and temporal discretization of the models in Section 3 (noting that the Newton linearization is trivial in the case of the linear Stokes equations). Since the systems are nonsymmetric, we use FGMRES [49] as the Krylov method, and seek to effectively precondition it. For IRK discretizations of scalar PDEs, such as the heat equation, block-diagonal preconditioning of the stage-coupled linear systems is known to be effective [34]. While block-diagonal preconditioning has also been developed for fluid models discretized using BDF-like methods [24, 25], we leave extension of this approach to IRK discretizations for future work. Instead, we follow the approach of [18, 31], and develop a monolithic multigrid preconditioner that makes use of an overlapping additive Schwarz relaxation that can be viewed as the extension of Vanka relaxation to the IRK case. Compared to the use of block-structured preconditioners, this offers the advantage of not needing to explicitly approximate Schur complements in the stage-coupled IRK linearizations (which may depend on properties of the specific IRK scheme chosen, for example). We note that FGMRES and classical right-preconditioned GMRES solve the same underlying optimization problem for the approximation in the same Krylov space, but that the underlying algorithms have important differences, with FGMRES requiring extra vector storage (to store both the Arnoldi vectors and their preconditioned counterparts) but right-preconditioned GMRES requiring an extra application of the preconditioner once the solution to the underlying Hessenberg system has been found. Thus, we choose to use FMGRES, instead of classical right-preconditioned GMRES, primarily because the cost of application of our preconditioners is non-trivial, but we are not memory-bound on the parallel machine used in the numerical results. Thus, the extra vector storage of FGMRES is an attractive trade-off over the extra preconditioner application required by standard right-preconditioned GMRES. An auxiliary advantage is that FGMRES allows the use of GMRES inside inner

iterations (such as the relaxation).

4.1. Coarse-grid correction and transfer operators

In the numerical results that follow, we consider hierarchies of grids generated by taking uniform refinements of a given coarsest grid. To map functions from a coarse mesh to its refinement, we use canonical finite-element interpolation operators for each field in the discretization. For the single-stage case, for both the time-dependent Stokes and Navier–Stokes problems, interpolation takes the form

$$
P = \begin{bmatrix} P_{\mathbf{u}} & \\ & P_p \end{bmatrix},
$$

where $P_{\mathbf{u}}$ and P_p represent the interpolation operators for the P_2 and P_1 finite-element spaces, respectively. In the MHD case, we introduce finite-element interpolation operators $P_{\mathbf{B}}$ for the lowest-order Nédélec space and P_γ for the P_1 space (with suitable boundary conditions for γ), following [22], making the interpolation operator for the single-stage case

$$
P = \begin{bmatrix} P_{\mathbf{u}} & & & \\ & P_{\mathbf{B}} & & \\ & & P_p & \\ & & & P_\gamma \end{bmatrix}.
$$

For multistage IRK discretizations, the interpolation operator is defined as $I_r \otimes P$, creating a block-diagonal interpolation operator that applies the finite-element interpolation in P to each stage independently. We use the transpose of interpolation as the restriction operator.

We use rediscretization to define the coarse-grid operators, noting that this is equivalent to Galerkin coarsening in the finite-element case (if compatible quadrature rules are used to assemble on the fine and coarse grids). The coarsest-grid systems are solved directly, using MUMPS [50].

4.2. Vanka Relaxation

Vanka relaxation was first introduced for the time-steady MAC-scheme discretization of the Navier–Stokes equations [51], but it has been extensively used

15

in many more general settings in recent years [20–22, 52, 53]. Broadly defined, Vanka relaxation schemes are overlapping Schwarz (domain decomposition) methods used as relaxation for a multigrid algorithm. In order to achieve the expected cost of a multigrid relaxation scheme, the subdomain problems are generally quite small, on the order of 10s-100s of DoFs. Historically, the most common approaches were multiplicative in nature; however, we follow the recent trend towards additive schemes [22, 52, 53] that are naturally parallelizable.

To specify the details of relaxation, we now describe how the Schwarz subdomains (commonly referred to as the Vanka "blocks" or "patches") are constructed from the underlying mesh on any given level of the multigrid hierarchy. In this work, we follow the topological construction described in [53]. In particular, we form a Vanka patch for each vertex in the mesh, which consists of all degrees of freedom associated with the closure of the cells adjacent to the vertex. As is typical in Vanka relaxation, we exclude all degrees of freedom associated with P_1 constraints (the pressure and Lagrange multiplier in our systems) from the patch, except for those located at the vertex around which the patch is formed. For the models considered here, this results in patches like those shown in Figure 1 for regular two-dimensional grids, with a single pressure degree of freedom and all velocity DoFs on all elements adjacent to the node. When used in an IRK discretization, these patches include all stage degrees of freedom. For MHD, we note that the patch shown at right of Figure 1 coincides topologically with the *coupled Vanka* approach for the BDF2 discretization considered in [22], but the patches used here contain more degrees of freedom than those used in [22], due to inclusion of all stages in the IRK discretization.

Denoting the set of DoFs in the i^{th} Vanka patch by $\mathcal{S}_i$, we have (by construction) that every degree of freedom in the domain is contained in at least one patch: $\mathcal{S} = \bigcup_{i=1}^{N} \mathcal{S}_i$, where N is the total number of patches and $\mathcal{S}$ is the complete set of DoFs for the problem. Denoting R_i as a "restriction" operator that maps global DoFs to those in patch $\mathcal{S}_i$, we can write a single iteration of a
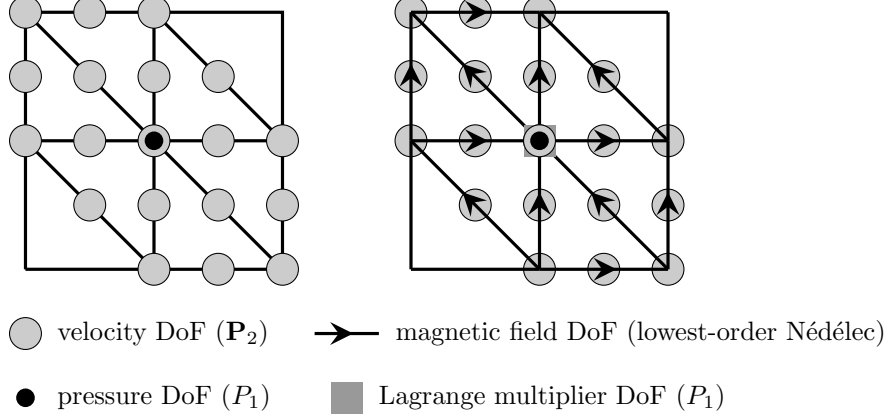
Figure 1: Left: Vanka patch for the Stokes and Navier–Stokes equations, consisting of P_2 velocity DoFs, and one P_1 pressure DoF. Right: Vanka patch for the MHD equations, consisting of P_2 velocity DoFs, lowest-order Nédélec DoFs for the magnetic field, one P_1 Lagrange multiplier DoF and one P_1 pressure DoF.

weighted stationary iteration as

$$\vec{\mathbf{k}} \leftarrow \vec{\mathbf{k}} + \omega \sum_{i=1}^{N} R_i^T (R_i J R_i^T)^{-1} R_i (\vec{\mathbf{F}} - J\vec{\mathbf{k}}),$$

where $J\vec{\mathbf{k}} = \vec{\mathbf{F}}$ is the linear system to be solved, and $R_i J R_i^T$ is the restriction of J to the DoFs in patch $\mathcal{S}_i$. In practice, we use several steps of a Vanka-preconditioned Chebyshev or GMRES iteration as the relaxation scheme for our problems, with the endpoints of the interval defining the associated Chebyshev polynomials tuned by hand.

4.3. Implementation

The numerical results below are produced using Firedrake [54] for the spatial finite-element discretization and Irksome [40] for the temporal discretization. Linear and nonlinear solvers are implemented in PETSc [55], taking advantage of the close integration between discretizations and solvers provided by this combination [56]. The Vanka relaxation is implemented through PCPATCH [53]. For reproducibility, the codes used to generate the numerical results and the major components of Firedrake, Irksome, and PETSc needed, have been archived

on Zenodo [57]. We emphasize that all aspects of the discretization and solver software are chosen to be naturally parallelizable. The coarsest mesh in each hierarchy is distributed across the available parallel cores, and then refined in parallel. For the two-dimensional problems below, after each refinement, the mesh is redistributed to better balance parallel work, but this was not done for the 3D example, due to software limitations. While not rebalancing the meshes leads to a small load imbalance on the finer grids in the hierarchies, this was not seen to lead to significant loss of performance in the weak scaling tests reported below. To account for the need to compute residuals for each DoF in each Vanka patch, a node-distance-2 halo is included in the parallel mesh distribution, to allow the relaxation scheme to be performed in parallel without additional communication [58].

5. Numerical Results

For the numerical results in this paper, we focus on 3 families of IRK methods: Gauss (also known as Gauss–Legendre), LobattoIIIC, and RadauIIA. We note that LobattoIIIC and RadauIIA are both L-stable and A-stable, while Gauss is A-stable, but not L-stable. All are fully implicit schemes, with stage order equal to the number of stages. For an r-stage method, Gauss schemes have order $2r$, while RadauIIA have order $2r - 1$ and LobattoIIIC have order $2r - 2$. We consider both 2- and 3-stage schemes here, with standard Butcher tableaux [43, 59].

We present results for four separate test cases: a simple two-dimensional time-dependent Stokes model, two-dimensional Navier–Stokes flow past a cylinder, and two MHD examples, a two-dimensional island-coalescence problem and a three-dimensional lid-driven cavity model. All results presented in this paper were computed on the Compute Canada cluster, *Niagara*, consisting of 2,024 nodes, each with 40 2.4 GHz Intel Skylake cores and 202GB of RAM, connected using a 100Gb/s EDR Dragonfly+ network.

5.1. Two-dimensional time-dependent Stokes

We consider a method of manufactured solutions test case, solving (2) on the two-dimensional unit square, $\Omega = (0,1)^2$. The forcing function $\mathbf{f}$ and boundary conditions are chosen so that the exact solution is

$$\mathbf{u} = \begin{bmatrix} \sin(\pi x)\cos(\pi y)e^{-2t\pi^2} \\ -\cos(\pi x)\sin(\pi y)e^{-2t\pi^2} \end{bmatrix}, \text{ and } p = 0.$$

For this example, we construct a coarsest grid by creating a uniform 8×8 quadrilateral mesh of the unit square, then cut each quadrilateral cell into 4 triangles, adding a vertex at the center of the quadrilateral. This mesh is then uniformly refined ℓ times; below, we present results for $\ell = 5, 6, 7$, where the Taylor–Hood discretization of the Stokes equations results in about 1.1 million DoFs per stage for $\ell = 5$ up to about 19 million DoFs per stage for $\ell = 7$. The initial condition is chosen by interpolating the exact solution into the finite-element space at $t = 0$, and we integrate up to time $T_f = 0.5$, with timestep $\Delta t = T_f/N$ for $N = 2^{\ell+3}$. To our knowledge, there are no rigorous stopping tolerances that guarantee discretization-error level accuracy for these systems; we use a hand-tuned stopping tolerance, where we require the absolute value of the ℓ_2 norm of the residual of the system to be reduced below $10^{-2} \times N^{-3}$ at each timestep, or a relative reduction in this norm by 10^{-8}. Based on preliminary experiments, we accelerate the relaxation process using Chebyshev polynomials of the first kind on the interval $[2, 8]$ and employ 2 pre- and post-relaxation sweeps. Proper choice of the Chebyshev interval is critical to achieving scalable performance. For two-dimensional problems with geometric coarsening by a factor of two (as used here), a reasonable strategy is to estimate the largest eigenvalue, λ, of the relaxation-preconditioned matrix (e.g., using Ritz values from preconditioned GMRES) and choose the interval to be $[\lambda/4, \lambda]$. Here, we started from similar estimates, but hand-tuned the intervals to optimize performance.

Table 1 presents a weak scaling study for this problem, for both two- and three-stage methods. For the two-stage methods, we use 10 cores on 1 node for

$\ell = 5$, 40 cores on 1 node for $\ell = 6$ and 160 cores on 4 nodes for $\ell = 7$. For the three-stage methods, we increase core counts by 50%, to account for the increased number of degrees of freedom in the resulting linear systems, using 15 cores on 1 node for $\ell = 5$, 60 cores on 2 nodes for $\ell = 6$ and 240 cores on 6 nodes for $\ell = 7$. We report the relative L_2 error in the velocity and the absolute L_2 error in the pressure approximation at the final time, as well as the average number of linear iterations to achieve convergence over all timesteps and the total computational time needed in minutes. In the final column of Table 1, we report the average wall-clock time per Krylov iteration (t/K) in seconds.

Table 2 summarizes rates of convergence for the results shown in Table 1. We observe at least second-order convergence in the velocity error for all three IRK schemes; however, we notice much larger errors for the Gauss results than for the other two schemes. We note that the stopping tolerance decreases by a factor of 8 with each refinement, so that the slight increase in averaged iterations to convergence is not overly surprising, and seems to remain bounded at reasonable levels. Nonetheless, the factor four increase in the number of cores with each refinement is insufficient to lead to ideal time scaling (which would be to double with each refinement, due to the doubling of the number of time-steps).

There are several contributing factors to the less-than-perfect scaling, beyond the simple increase in total number of Krylov iterations with refinement. When going from $\ell = 5$ to $\ell = 6$ with the two-stage schemes, we increase the number of cores used for the calculation, but those cores remain on one physical node, leading to a saturation of the memory bandwidth available. The same limitation occurs when going from $\ell = 6$ to $\ell = 7$ with the three-stage schemes, where we go from using 30 cores on each of 2 nodes to all 40 cores on 6 nodes. While this could be avoided by using the same number of cores on more nodes of the parallel machine, such usage is impractical when a single node has sufficient memory for the $\ell = 6$ problem with 2 IRK stages. Furthermore, when going from $\ell = 6$ to $\ell = 7$, the (direct) coarsest-grid solve goes from being dominated by its computation to being dominated by its communication. Here, we clearly see another increase in the cost per linear iteration, especially in the three-stage

	ℓ	velocity error	pressure error	iterations	time	t/K
	5	1.786×10^{-2}	2.328×10^{-2}	9.85	57.47	1.401
Gauss(2)	6	3.155×10^{-3}	1.162×10^{-2}	13.39	202.98	1.779
	7	5.575×10^{-4}	6.271×10^{-3}	19.14	725.02	2.285
	5	3.380×10^{-6}	6.327×10^{-9}	8.70	56.13	1.570
RadauIIA(2)	6	4.297×10^{-7}	2.795×10^{-9}	10.99	171.08	1.837
	7	4.971×10^{-8}	1.028×10^{-9}	14.22	575.77	2.391
	5	9.823×10^{-4}	4.594×10^{-7}	9.23	54.13	1.912
LobattoIIIC(2)	6	2.495×10^{-4}	1.479×10^{-7}	11.71	181.88	1.833
	7	6.289×10^{-5}	4.823×10^{-8}	15.25	618.77	2.393
	5	9.098×10^{-5}	1.481×10^{-4}	13.38	109.39	1.929
Gauss(3)	6	1.839×10^{-5}	6.802×10^{-4}	24.51	483.75	2.323
	7	3.293×10^{-6}	3.235×10^{-4}	25.97	1332.40	3.249
	5	6.151×10^{-7}	1.298×10^{-9}	9.32	82.29	2.083
RadauIIA(3)	6	1.122×10^{-7}	1.310×10^{-9}	12.40	267.78	2.566
	7	1.300×10^{-8}	1.718×10^{-10}	13.91	877.99	3.786
	5	6.019×10^{-7}	2.728×10^{-9}	9.61	85.25	2.120
LobattoIIIC(3)	6	1.083×10^{-7}	1.966×10^{-9}	12.91	248.95	2.293
	7	1.271×10^{-8}	3.797×10^{-10}	14.91	938.06	3.770

Table 1: Numerical results for two-dimensional Stokes model problem with two- and three-stage IRK schemes. Relative L_2 errors in velocity and absolute L_2 errors for pressure are reported, along with average number of linear solver iterations per time-step, total wall-clock time-to-solution in minutes, and time per Krylov iteration in seconds, for refinement levels $\ell = 5, 6, 7$.

		Gauss		RadauIIA		LobattoIIIC	
		u	p	u	p	u	p
two-stage	$\log_2 e_5/e_6$	2.5	1.0	3.0	1.2	2.0	1.6
	$\log_2 e_6/e_7$	2.5	0.9	3.1	1.4	2.0	1.6
three-stage	$\log_2 e_5/e_6$	2.4	-2.3	2.5	0.0	2.5	0.5
	$\log_2 e_6/e_7$	2.5	1.1	3.1	2.9	3.1	2.4

Table 2: Rates of convergence in velocity and pressure for data in Table 1 with two- and three-stage IRK schemes for refinement levels $\ell = 5, 6, 7$. Here, e_ℓ denotes the error in a quantity on refinement level ℓ.

methods where the time required increases by about 50% for $\ell = 7$. Improved performance would almost certainly be seen by duplicating the coarse-grid solve on each node, as considered in [60–62]. We leave these performance enhancements for future work.

The experiment in Table 1 highlights convergence as we change both the spatial and temporal discretizations. Here, however, we note that the temporal discretizations are higher order than the spatial, particularly for the 3-stage discretizations. Thus, for comparison with this data, Table 3 presents results with the same setup as Table 1, but using a fixed timestep of $\Delta t = 0.5/2^8$, to match results with $\ell = 5$ (noting that these results were run independently, so small differences in timings for $\ell = 5$ naturally arise). We see that while quite reasonable convergence is observed in Table 1, we observe significant stagnation in convergence here. Thus, even though the temporal discretizations are higher order, we still see substantial benefits to varying the timestep simultaneously with refinement of the spatial mesh.

Finally, Table 4 presents comparison results for diagonal IRK schemes. Here, we consider the two-stage second-order Pareschi-Russo (with parameter $1 - \sqrt{2}/2$) [63] and three-stage third-order Alexander [64] integrators, which are both L-stable. While these results show some outperformance of the theoretical

	ℓ	velocity error	pressure error	iterations	time	t/K
	5	3.380×10^{-6}	6.327×10^{-9}	8.70	56.37	1.560
RadauIIA(2)	6	3.119×10^{-6}	1.896×10^{-9}	12.18	96.77	0.940
	7	3.101×10^{-6}	2.793×10^{-9}	18.04	165.42	0.540
	5	6.151×10^{-7}	1.298×10^{-9}	9.32	82.07	2.087
RadauIIA(3)	6	3.677×10^{-8}	3.288×10^{-9}	14.22	116.27	0.964
	7	3.566×10^{-8}	8.063×10^{-10}	17.30	238.18	0.809
	5	6.019×10^{-7}	2.728×10^{-9}	9.61	75.87	1.850
LobattoIIIC(3)	6	6.066×10^{-8}	3.076×10^{-9}	14.60	107.87	0.866
	7	5.611×10^{-8}	1.136×10^{-9}	17.47	239.03	0.805

Table 3: Results analogous to Table 1, but with $\Delta t = 0.5/2^8$. Relative L_2 errors in velocity and absolute L_2 errors for pressure are reported, along with average number of linear solver iterations per time-step, total wall-clock time-to-solution in minutes, and time per Krylov iteration in seconds, for refinement levels $\ell = 5, 6, 7$.

guarantees given by their stage order of one, they are also quite poor in comparison to the RadauIIA integrators of the same number of stages. In particular, comparing with the results in Table 1, we see that the errors achieved using DIRK(3) with $\ell = 6$ are comparable to those achieved when using RadauIIA(2) with $\ell = 5$, but that the latter calculation was achieved in about 60% of the wall-clock time and on one-eighth of the number of cores (10 for RadauIIA(2) with $\ell = 5$ vs. 80 for DIRK(3) with $\ell = 6$). Similarly, the errors for DIRK(3) with $\ell = 7$ are slightly better than those achieved with RadauIIA(2) and $\ell = 6$, and slightly worse than those achieved with RadauIIA(3) and $\ell = 6$. The two-stage Radau results, however, are achieved in just over 40% of the wall-clock time, and on one-sixth the cores, while the three-stage Radau results are achieved in about two-thirds the wall-clock time, on one-fourth the cores. These results highlight the added accuracy that can be gained using fully implicit RK methods over DIRK methods, and the added efficiency possible when using state-of-the-art

	ℓ	velocity error	pressure error	iterations	time	t/K
	5	2.480×10^{-5}	9.266×10^{-8}	6.42	35.20	1.312
DIRK(2)	6	6.198×10^{-5}	3.254×10^{-8}	8.23	110.43	1.590
	7	1.546×10^{-5}	1.173×10^{-8}	10.07	213.76	1.835
	5	1.106×10^{-5}	4.161×10^{-8}	6.99	35.28	1.240
DIRK(3)	6	2.325×10^{-6}	1.857×10^{-9}	9.51	94.98*	1.12
	7	2.607×10^{-7}	4.578×10^{-9}	11.17	415.26	2.36

Table 4: Numerical results for two-dimensional Stokes model problem with two- and three-stage DIRK schemes. Relative L_2 errors in velocity and absolute L_2 errors for pressure are reported, along with average number of linear solver iterations per time-step, total wall-clock time-to-solution in minutes, and time per Krylov iteration in seconds, for refinement levels $\ell = 5, 6, 7$. Due to a change in the configuration of the machine, results for DIRK(3) at $\ell = 6$ were run on 80 cores, instead of 60; all other results were run with same parallelism as in Table 1.

linear solvers to achieve that accuracy. For this reason, we focus on only the fully implicit RK schemes in the remainder of the paper.

5.2. Two-dimensional Navier–Stokes

We next consider two-dimensional Navier–Stokes flow past a cylinder, following the example given in [40, 65, 66]. Here, we consider the spatial domain $\Omega = (0, 2.2) \times (0, 0.41) \setminus B_r(0.2, 0.2)$, where $B_r(0.2, 0.2)$ is the disc of radius $r = 0.05$ centred at $(0.2, 0.2)$, shown in Figure 2. No-slip (zero-velocity) boundary conditions are imposed on the top and bottom boundaries of the rectangle and along the surface of the cylinder. Time-dependent inflow conditions are given on the left edge, prescribing

$$\mathbf{u}(0, y, t) = \begin{bmatrix} \frac{4U(t)y(0.41-y)}{0.41^2} \\ 0 \end{bmatrix},$$

where $U(t) = 1.5 \sin\left(\frac{\pi t}{8}\right)$ is the mean inflow velocity. No-stress outflow is prescribed on the right boundary. The viscosity is set as $\mu = 10^{-3}$, resulting
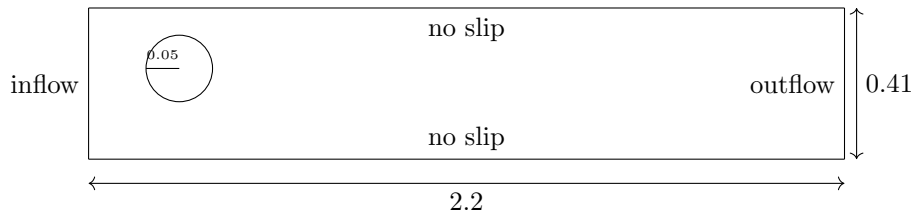
Figure 2: Domain for Navier–Stokes flow past a cylinder.

in a Reynolds number of 100. The time step for these experiments is fixed as $\Delta t = \frac{1}{400}$, and we consider the final time $T_f = 8$. As above, we discretize using Taylor–Hood elements in space and IRK in time.

For this problem, an unstructured coarsest grid with 972 triangular elements is used, chosen to refine the representation around the included cylinder. Below, we report results for $3 \leq \ell \leq 6$, with discrete problem sizes for the Taylor–Hood discretization ranging from about 245 thousand DoFs per stage for $\ell = 3$ to about 15.5 million DoFs per stage for $\ell = 6$. Details of the parallelization are provided in Table 5, where we again note that we have increased the number of cores for the 3-stage IRK methods by about 60% over those for the 2-stage methods. For this problem, we use a nonlinear stopping tolerance requiring the absolute ℓ_2 norm of the nonlinear residual be below $1/N^3$ with $N = 2^{\ell+3}$, and use an Eisenstat–Walker inexact Newton scheme to determine the linear stopping tolerances for each nonlinear iteration. Here, again 2 pre- and post-relaxation sweeps are used, with Chebyshev polynomials for relaxation taken over the interval $[1.5, 8]$.

As no analytical solution is available in this case, we instead record the maximum drag and lift values computed over the simulations, for comparison with reference data [40, 65]. Figure 3 presents time-histories of these quantities for one simulation, showing excellent agreement with reference data. Results for other simulations with both RadauIIA and LobattoIIIC are visually similar. Table 6 presents these values for both of these integrators, along with the average wall-clock time in minutes per time-step, and average nonlinear and linear

25

stages	ℓ	total DoFs	nodes	cores
2	3	489,656	1	6
	4	1,948,144	1	25
	5	7,771,616	4	100
	6	31,044,544	10	400
3	3	734,484	1	10
	4	2,922,216	1	40
	5	11,657,424	4	160
	6	46,566,816	16	640

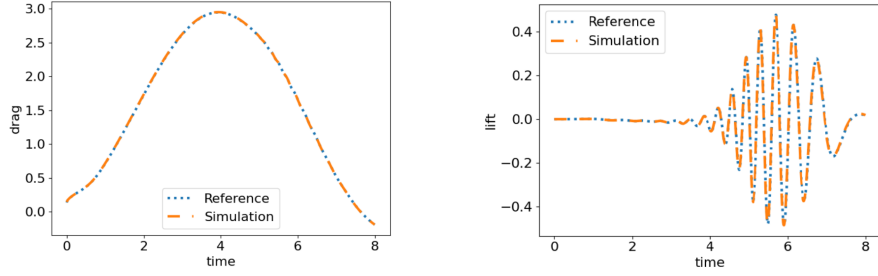Table 5: Total number of DoFs and number of nodes and cores used for the Navier–Stokes test problem with two- and three-stage IRK discretizations.

Figure 3: Comparison of reference and drag (left) and lift (right) computed using LobattoI-IIC(2) and $\ell = 6$.

solver iterations per time-step. As before, we have decreasing solver tolerances as ℓ increases, so the small increases in iterations counts with refinement are expected.

Using both RadauIIA and LobattoIIIC IRK discretizations, with either two or three stages, results in computed lift and drag values that are consistent with those presented in [40, 65]. However, results computed with Gauss were not. Figure 4 shows results using the three-stage Gauss method with $\ell = 3$, computed with a stricter stopping tolerance (absolute nonlinear residual norm

26

	nref	drag max	lift max	time	nonlinear its	linear its
RadauIIA(2)	3	2.95	0.48	0.03	1.37	1.61
	4	2.95	0.48	0.04	1.76	2.39
	5	2.95	0.48	0.07	2.14	3.56
	6	2.95	0.48	0.13	2.52	5.09
LobattoIIIC(2)	3	2.95	0.48	0.03	1.35	1.80
	4	2.95	0.48	0.05	1.79	2.76
	5	2.95	0.48	0.07	2.15	4.30
	6	2.95	0.48	0.17	2.70	8.19
RadauIIA(3)	3	2.95	0.48	0.04	1.57	2.27
	4	2.95	0.48	0.07	1.90	2.78
	5	2.95	0.48	0.11	2.17	3.68
	6	2.95	0.48	0.17	2.56	5.01
LobattoIIIC(3)	3	2.95	0.48	0.04	1.38	1.74
	4	2.95	0.48	0.06	1.78	2.50
	5	2.95	0.48	0.10	2.14	3.60
	6	2.95	0.48	0.19	2.54	5.06

Table 6: Maximum drag and lift values, average wall-clock time per time-step (in minutes) and average numbers of nonlinear and linear iterations per time step for $3 \leq \ell \leq 6$ for Navier–Stokes flow past a cylinder.
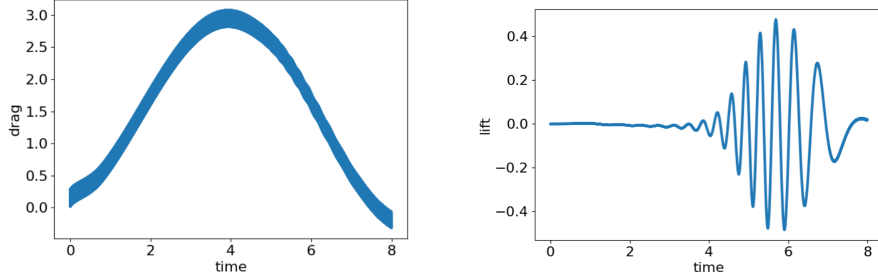
Figure 4: Drag and lift for $\ell = 4$ using Gauss(3). The "thick lines" indicate that the solutions are highly oscillatory in time, due to the lack of L-stability of the integrator.

below $1/N^4$) than used above for RadauIIA and LobattoIIIC methods[1]. The appearance of "thick lines" in these plots reflects highly oscillatory numerical solutions. We hypothesize that this is due to the lack of L-stability of the integrator, where large negative eigenvalues of the linearized spatial operator are not quickly damped but, rather, slowly decay and oscillate in time due to a stability function value close to -1. Refinement in time for fixed spatial grids should ameliorate the issue, but leads to increased computational costs to achieve similar accuracy to that given by RadauIIA and LobattoIIIC with these timesteps.

5.3. Two-Dimensional MHD Island Coalescence

We next consider a standard test model in MHD, of two-dimensional island coalescence. This model mimics flow in a large aspect ratio tokamak, considering a cross-section of flow of magnetically confined plasma. When a large external magnetic field is imposed in the "toroidal" direction of the tokamak, essentially two-dimensional dynamics result. This model geometry is then mapped and rescaled to a square domain, $\Omega = (-1, 1)^2$, with periodic boundary conditions on the left and right edges (see [22, 67, 68] for more details). In this geometry,

---

[1]Using the same stopping tolerance led to even more inconsistent data.

an equilibrium solution to the MHD equations is given by

$$\mathbf{u}_0(x, y) = \mathbf{0},$$

$$\mathbf{B}_0(x, y) = \frac{1}{\cosh(2\pi y) + k\cos(2\pi x)} \begin{pmatrix} \sinh(2\pi y) \\ k\sin(2\pi x) \end{pmatrix},$$

$$p(x, y) = \frac{1 - k^2}{2} \left(1 + \frac{1}{(\cosh(2\pi y) + k\cos(2\pi x))^2} \right),$$

$$\gamma(x, y) = 0,$$

where $k = 0.2$, when forcing terms of

$$\mathbf{f_u} = \mathbf{0},$$

$$\mathbf{f_B} = \frac{-8\pi^2(k^2 - 1)}{Re_m(\cosh(2\pi y) + k\cos(2\pi x))^3} \begin{pmatrix} \sinh(2\pi y) \\ k\sin(2\pi x) \end{pmatrix},$$

are imposed on the differential equation. To initialize a dynamic problem, these forcing terms are applied, but the initial condition is perturbed by adding

$$\delta\mathbf{B} = \frac{-0.01}{\pi} \begin{pmatrix} -\cos(\pi x)\sin(\frac{\pi y}{2})/2 \\ \cos(\frac{\pi y}{2})\sin(\pi x) \end{pmatrix}.$$

to the equilibrium solution at $t = 0$. The expected effect of this perturbation is to create two initially separated "islands" of current density that break the magnetic field lines, which then reconnect. At the reconnection point (or $\mathcal{X}$-point), a sudden sharp spike should be seen in the magnetic current density. At higher Reynolds numbers, a "sloshing" effect should occur before the islands of current density merge.

As above, no analytical solution is known for this problem. A key measure of the physical fidelity is the time-history of the *reconnection rate*, computed as the difference between the curl of $\mathbf{B}$ at the origin at the current time and its value at the origin, scaled by $1/\sqrt{Re_m}$. We compute this using the same methodology as in [22]. As Re and Re_m increase, the peak value of the reconnection rate should decrease, and the length of time for which this value is nonzero should increase. In this section, we consider only the two-stage LobattoIIIC integrator, and integrate until $T_f = 20$. Following [22], we "substep" for the first time-step,
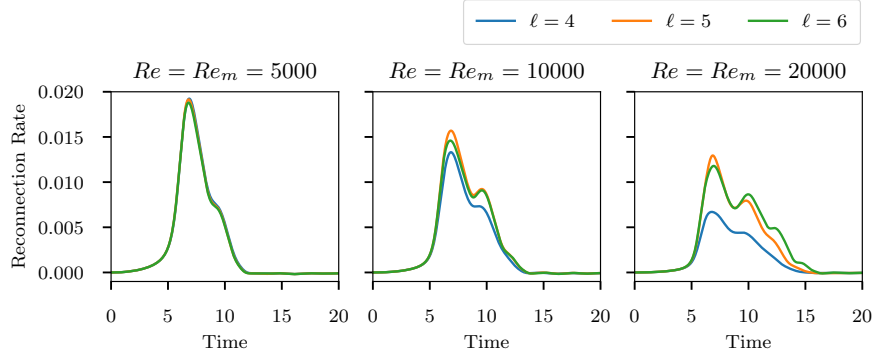
Figure 5: Reconnection rates recorded for the 2D MHD island coalescence model with varying Reynolds numbers of Re = Re$_m$ = 5000, 10000 and 20000 on 3 different levels of refinement using the LobattoIIIC(2) temporal discretization.

taking 10 substeps to initialize the simulation and avoid problems with nonlinear convergence. We consider a coarsest spatial mesh of 20×20 quadrilateral elements, again each cut into 4 triangles, and present results for $\ell = 4, 5, 6$ refinements. For $\ell = 4$, the resulting discretization has about 2.7 million DoFs per stage, while it has about 42.7 million DoFs per stage for $\ell = 6$. For $\ell = 4$, we use $\Delta t = 0.025$, which is halved with each spatial refinement. We test for 3 different pairs of Reynolds numbers, Re = Re$_m$: 5000, 1000 and 20000. Figure 5 shows the computed reconnection rates for these problems with varying Re = Re$_m$ and ℓ, properly reflecting the expected behaviour.

For this problem, we adjust the nonlinear and linear solver parameters as follows. Taking $N = 20 \times 2^\ell$ as a representative number of elements in one dimension on refinement level ℓ, we set both nonlinear and linear stopping tolerances to demand an absolute reduction of the ℓ_2 norm of the corresponding residual below $1/N^2$. We now use 3 pre- and post-relaxation sweeps, with the Chebyshev polynomials defining the relaxation taken over the interval [2, 10]. Figures 6 and 7 present results from a weak scaling study, using 40 cores on 1 node for $\ell = 4$, 160 cores on 4 nodes for $\ell = 5$, and 640 cores on 16 nodes for $\ell = 6$. We note that these are larger core counts than those used for the
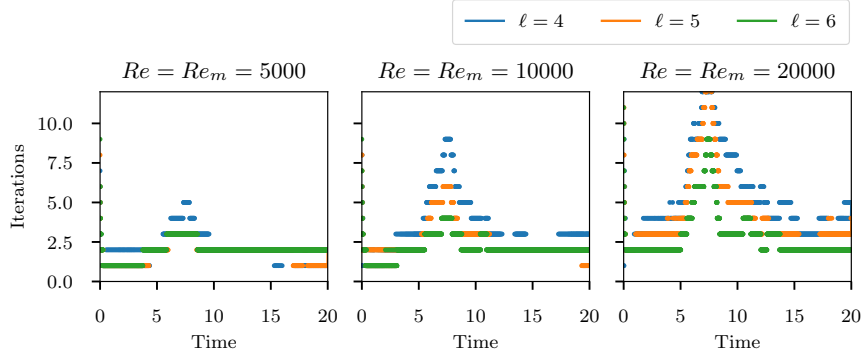
Figure 6: Number of linear iterations per timestep for the 2D MHD island coalescence model with varying Reynolds numbers of $Re = Re_m = 5000$, 10000 and 20000 on 3 different levels of refinements using the LobattoIIC(2) integrator.

same underlying meshes with a BDF2 discretization in [22]; however, this is due to the larger number of DoFs in the system using a 2-stage discretization. On average, our finest-grid problems have about 67 thousand DoFs per stage per core, which is a reasonable range for weak scaling. We note that the $\ell = 4$ problem takes about 2 hours of wall-clock time with these settings, with slightly better than doubling of wall-clock with each refinement (due to the halving of Δt with each refinement, but also improved solver performance).

Figure 6 shows the number of linear solver iterations recorded per timestep as we vary $Re = Re_m$ and ℓ. We note slight growth in iteration counts as $Re = Re_m$ increases (and the problem becomes less diffusive in nature), but also improving iteration counts at fixed values of $Re = Re_m$ as ℓ increases. Comparing with iteration counts from [22], we see slightly higher iteration counts here, with slightly worse dependence on $Re = Re_m$, but still reasonable performance overall. Wall-clock times per timestep, shown in Figure 7, generally reflect the linear iteration counts. In particular, we again see a general increase with $Re = Re_m$, and a general decrease with increasing ℓ. The most expensive solves in the test set are still achieved in under 1 minute per time-step, and the average time is about 0.2 minutes per time-step.
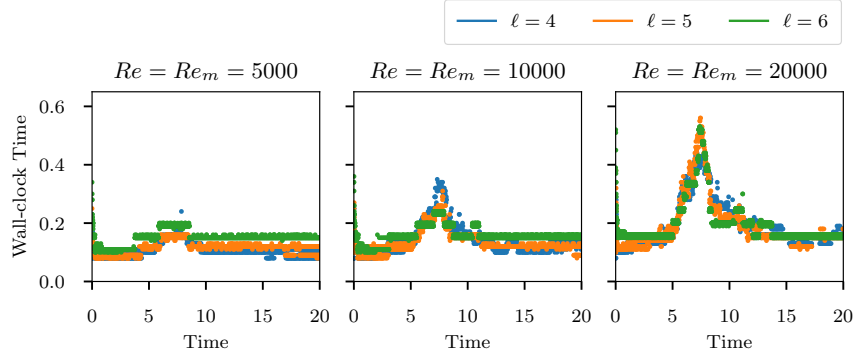
Figure 7: Wall-clock time (in minutes) for the nonlinear system solve at each time-step for the 2D MHD island coalescence model with varying Reynolds numbers of $Re = Re_m = 5000$, 10000 and 20000 on 3 different levels of refinements using the LobattoIIC(2) integrator.

To better understand the linear and nonlinear solver performance shown above, we compute both the fluid and magnetic (Alfvén) CFL numbers for the flow. At each timestep, for the given solutions for $\mathbf{u}$ and $\mathbf{B}$, we approximate the maximum magnitude of the vector fields (by projecting $\mathbf{u} \cdot \mathbf{u}$ and $\mathbf{B} \cdot \mathbf{B}$ into the discontinuous piecewise-constant finite-element space on the finest mesh and computing the maximum values of these projections), $u_{\max}$ and $B_{\max}$, and then computing the fluid CFL value, $u_{\max}\frac{\Delta t}{h}$, and the Alfvén CFL value, $B_{\max}\frac{\Delta t}{h}$, where h is a representative edge length for the spatial mesh. Figure 8 shows both CFL values calculated at each timestep for the simulations considered, showing identical results to those obtained in the BDF2 case in [22]. We note that, aside from the initial substeps, the Alfvén CFL is roughly constant at a value around 6, while the fluid CFL peaks at the same time as the reconnection rate, and is above 1 for the largest values of $Re = Re_m$ considered.

5.4. Three-dimensional MHD lid-driven cavity

Finally, we present results for a three-dimensional lid-driven cavity MHD model on the unit cube, $\Omega = (0,1)^3$, following [69]. On the top face, $z = 1$, the flow is driven by imposed velocity $\mathbf{u} = (1,0,0)^T$, while $\mathbf{u} = (0,0,0)^T$ on all other faces. The tangential components of the magnetic field are set to
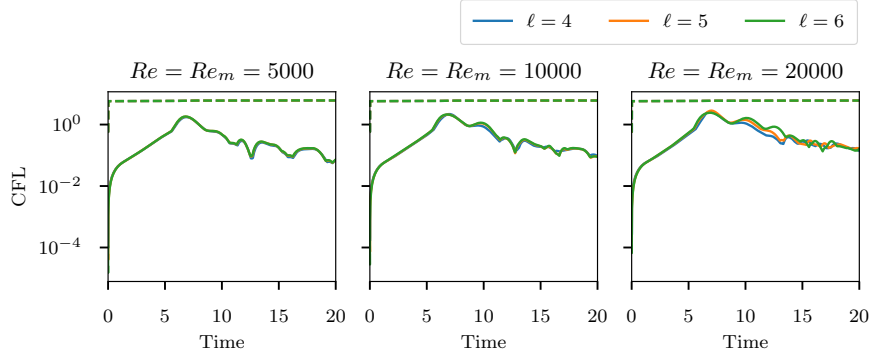
Figure 8: CFL values at each timestep for the 2D MHD island coalescence model with varying Reynolds numbers of Re = Re$_m$ = 5000, 10000 and 20000 on 3 different levels of refinements using the LobattoIIC(2) integrator. Solid lines represent fluid CFL, while dashed lines represent Alfvén CFL.

match those of $\mathbf{B} = (-1, 0, 0)^T$ on all faces of the cube. We set $\gamma = 0$ on all boundary faces and fix the pressure $p = 0$ at the origin. In this section, we consider 3-grid methods, refining a given coarsest grid twice for each test. This is driven by the consideration that, with increasing finest grid size, we require more cores over which to parallelize the computation; however, there is a software limitation within Firedrake that requires that the coarsest grid in the simulation must have at least 1 cell per core. While satisfying this requirement is not burdensome in 2D, it becomes problematic with increasing memory and computational requirements of 3D simulations. In all cases, we construct the coarsest grid by taking a uniform hexahedral mesh of the cube, then cutting each hexahedral element into 6 tetrahedra in the usual way. We still use ℓ to denote the levels of refinement, but now $\ell = 1$ denotes the smallest grid, created by refining a $2 \times 2 \times 2$ grid twice, while $\ell = 2$ denotes the grid created by refining a $4 \times 4 \times 4$ grid twice, and $\ell = 3$ denotes the grid created by refining a $8 \times 8 \times 8$ grid twice. With $\ell = 1$, our discretization has about 20 thousand DoFs per stage, increasing to about 1.1 million DoFs per stage for $\ell = 3$.

We employ the same spatial discretization and again use the LobattoIIIC(2) integrator. We integrate until $T_f = 2.5$. For $\ell = 1$, we take $\Delta t = 0.125$, and

halve Δt with each refinement. An all-zero initial condition is used. Table 7 presents average linear and nonlinear iterations per timestep, along with average wall-clock time per nonlinear solve for the three grids above and $\text{Re} = \text{Re}_m = 10^p$ for $1 \leq p \leq 3$. For $\ell = 1$, 10 cores on 1 node are used, increasing to 80 cores on 2 nodes for $\ell = 2$ and 640 cores on 16 nodes for $\ell = 3$. We use 3 pre- and post-relaxation sweeps, here accelerated using GMRES, as this was observed to result in better overall iteration counts and computation times than using Chebyshev acceleration, likely due to the convective nature of the problem at high Reynolds numbers. The nonlinear solve at each timestep requires the absolute value of the ℓ_2 norm of the residual to be reduced below 10^{-6}, and the same stopping criterion is used for the linear solves as well.

Several trends can be observed in these results. First, for fixed values of $\text{Re} = \text{Re}_m$, we generally observe improving solver performance as ℓ is increased, as expected. Similarly, we typically observe degrading solver performance as $\text{Re} = \text{Re}_m$ is increased for fixed ℓ. Overall, the iteration counts are quite reasonable, except for $\text{Re} = \text{Re}_m = 1000$ with $\ell = 1, 2$. Here, the problem is quite severely under-resolved, with a finest-grid mesh spacing of $h = 0.0625$ with $\ell = 2$, so it is not surprising that the solver suffers when the discretization is so poor. For smaller Reynolds numbers, $\text{Re} = \text{Re}_m = 1$ (not shown here), using Chebyshev acceleration gave significantly better results than using GMRES-accelerated relaxation, which failed to converge in some cases. Figure 9 presents representative solutions for $\ell = 3$ with $\text{Re} = \text{Re}_m = 10$ (where the solutions are well-resolved), showing streamlines of both the velocity field, $\mathbf{u}$, and the magnetic field, $\mathbf{B}$, at the final time at refinement $\ell = 3$.

6. Conclusion

In this paper, we have developed monolithic Vanka relaxation schemes for fully-implicit Runge–Kutta discretizations of saddle point problems arising in models of fluid flow. Within a Newton–Krylov–multigrid setting, our method is shown to be effective for both Newtonian and magnetohydrodynamic flows, in

ℓ		Re = 10	Re = 100	Re = 1000
	linear its.	6.86	8.34	31.55
1	nonlinear its.	3.10	2.76	3.52
	time	0.11	0.11	0.24
	linear its	6.61	5.41	17.59
2	nonlinear its	3.18	2.43	3.02
	time	0.21	0.17	0.30
	linear its	5.76	4.20	6.08
3	nonlinear its	2.57	2.33	2.16
	time	0.22	0.22	0.22

Table 7: Average number of linear and nonlinear iterations per time-step and wall-clock time per nonlinear iteration (in minutes) for the 3D MHD lid-driven cavity problem with various Reynolds numbers and grid refinements, using the LobattoIIIC(2) integrator.
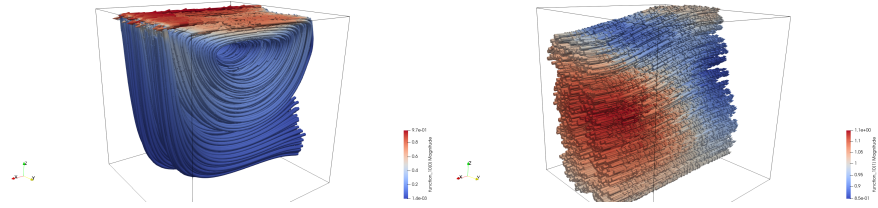

Figure 9: Streamlines of velocity (left) and magnetic field (right) for $\ell = 3$ with $\mathrm{Re} = \mathrm{Re}_m = 10$.

both two and three spatial dimensions. The algorithm is chosen with parallel implementation in mind, and weak scaling results are shown up to 640 cores.

There are many possibilities for future work. We note primarily that the current study uses relatively low-order spatial discretizations, based on classical Taylor–Hood elements for velocity and pressure. A next step in this research is to extend these solvers to more sophisticated finite-element discretizations that preserve the incompressibility and solenoidality constraints exactly, as in [9, 28]. An important question for future work is the extension of these techniques to higher-order discretizations, where the cost of classical sparse direct solvers for the patch problems becomes prohibitive.

References

[1] F. H. Harlow, J. E. Welch, Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface, Physics of Fluids 8 (12) (1965) 2182–2189.

[2] A. J. Chorin, Numerical solution of the Navier-Stokes equations, Math. Comp. 22 (1968) 745–762.

[3] R. Témam, Sur l'approximation de la solution des équations de Navier-Stokes par la méthode des pas fractionnaires. II, Arch. Rational Mech. Anal. 33 (1969) 377–385. `doi:10.1007/BF00247696`.

[4] L. R. Scott, M. Vogelius, Conforming finite element methods for incompressible and nearly incompressible continua, Lectures in Applied Mathematics 22 (2) (1985).

[5] S. Zhang, A new family of stable mixed finite elements for the 3D Stokes equations, Math. Comp. 74 (250) (2005) 543–554. `doi:10.1090/S0025-5718-04-01711-9`.

[6] V. John, A. Linke, C. Merdon, M. Neilan, L. G. Rebholz, On the divergence constraint in mixed finite element methods for incompressible flows, SIAM Review 59 (3) (2017) 492–544. `doi:10.1137/15M1047696`.

[7] N. R. Gauger, A. Linke, P. W. Schroeder, On high-order pressure-robust space discretisations, their advantages for incompressible high Reynolds number generalised Beltrami flows and beyond, The SMAI Journal of Computational Mathematics 5 (2019) 89–129. `doi:10.5802/smai-jcm.44`.

[8] D. Schötzau, Mixed finite element methods for stationary incompressible magneto-hydrodynamics, Numer. Math. 96 (4) (2004) 771–800. `doi:10.1007/s00211-003-0487-4`.

[9] K. Hu, Y. Ma, J. Xu, Stable finite element methods preserving $\nabla \cdot B = 0$ exactly for MHD models, Numer. Math. 135 (2) (2017) 371–396. `doi:10.1007/s00211-016-0803-4`.

[10] K. Hu, J. Xu, Structure-preserving finite element methods for stationary MHD models, Math. Comp. 88 (316) (2019) 553–581. `doi:10.1090/mcom/3341`.

[11] K. Hu, W. Qiu, K. Shi, Convergence of a B-E based finite element method for MHD models on Lipschitz domains, Journal of Computational and Applied Mathematics 368 (2020) 112477. `doi:10.1016/j.cam.2019.112477`.

[12] A. J. Chorin, On the convergence of discrete approximations to the Navier-Stokes equations, Math. Comp. 23 (1969) 341–353.

[13] J. van Kan, A second-order accurate pressure-correction scheme for viscous incompressible flow, SIAM J. Sci. Stat. Comput. 7 (3) (1986) 870–891. `doi:10.1137/0907059`.

[14] J. B. Bell, P. Colella, H. M. Glaz, A second-order projection method for the incompressible Navier-Stokes equations, J. Comput. Phys. 85 (2) (1989) 257–283.

[15] J. Guermond, P. Minev, J. Shen, An overview of projection methods for incompressible flows, Computer Methods in Applied Mechanics and Engineering 195 (44) (2006) 6011–6045. `doi:10.1016/j.cma.2005.10.010`.

[16] A. Brandt, N. Dinar, Multigrid solutions to elliptic flow problems, in: S. Parter (Ed.), Numerical Methods for Partial Differential Equations, Academic Press, New York, 1979, pp. 53–147.

[17] A. Brandt, Multigrid techniques: 1984 guide with applications to fluid dynamics, GMD–Studien Nr. 85, Gesellschaft für Mathematik und Datenverarbeitung, St. Augustin, 1984.

[18] S. P. Vanka, Block-implicit multigrid solution of Navier-Stokes equations in primitive variables, Journal of Computational Physics 65 (1) (1986) 138–158.

[19] D. Braess, R. Sarazin, An efficient smoother for the Stokes problem, Applied Numerical Mathematics 23 (1) (1997) 3–19.

[20] J. H. Adler, T. R. Benson, S. P. MacLachlan, Preconditioning a mass-conserving discontinuous Galerkin discretization of the Stokes equations, Numerical Linear Algebra with Applications 24 (3) (2017) e2047.

[21] J. H. Adler, T. R. Benson, E. C. Cyr, S. P. MacLachlan, R. S. Tuminaro, Monolithic multigrid methods for two-dimensional resistive magnetohydrodynamics, SIAM Journal on Scientific Computing 38 (1) (2016) B1–B24.

[22] J. H. Adler, T. Benson, E. C. Cyr, P. E. Farrell, S. MacLachlan, R. Tuminaro, Monolithic multigrid for magnetohydrodynamics, SIAM J. Sci. Comput. 43 (5) (2021) S70–S91.

[23] D. Kay, D. Loghin, A. Wathen, A preconditioner for the steady-state Navier–Stokes equations, SIAM Journal on Scientific Computing 24 (1) (2002) 237–256. `doi:10.1137/S106482759935808X`.

[24] H. C. Elman, D. J. Silvester, A. J. Wathen, Finite elements and fast iterative solvers: with applications in incompressible fluid dynamics, Oxford University Press, USA, 2014.

[25] M. Wathen, C. Greif, D. Schötzau, Preconditioners for mixed finite element discretizations of incompressible MHD equations, SIAM Journal on Scientific Computing 39 (6) (2017) A2993–A3013. `doi:10.1137/16M1098991`.

[26] P. E. Farrell, L. Mitchell, F. Wechsung, An augmented Lagrangian preconditioner for the 3D stationary incompressible Navier–Stokes equations at high Reynolds number, SIAM Journal on Scientific Computing 41 (5) (2019) A3073–A3096. `doi:10.1137/18M1219370`.

[27] P. E. Farrell, L. Mitchell, L. R. Scott, F. Wechsung, A Reynolds-robust preconditioner for the Scott-Vogelius discretization of the stationary incompressible Navier-Stokes equations, The SMAI Journal of Computational Mathematics 7 (2021) 75–96. `doi:10.5802/smai-jcm.72`.

[28] F. Laakmann, P. E. Farrell, L. Mitchell, An augmented Lagrangian preconditioner for the magnetohydrodynamics equations at high Reynolds and coupling numbers (2021). `arXiv:2104.14855`.

[29] W. Pazner, P.-O. Persson, Stage-parallel fully implicit Runge–Kutta solvers for discontinuous Galerkin fluid simulations, Journal of Computational Physics 335 (2017) 700–717.

[30] B. S. Southworth, O. Krzysik, W. Pazner, H. D. Sterck, Fast parallel solution of fully implicit Runge-Kutta and discontinuous Galerkin in time for numerical PDEs, Part I: the linear setting (2021). `arXiv:2101.00512`.

[31] J. Van Lent, S. Vandewalle, Multigrid methods for implicit Runge–Kutta and boundary value method discretizations of parabolic PDEs, SIAM Journal on Scientific Computing 27 (1) (2005) 67–92.

[32] E. Rosseel, T. Boonen, S. Vandewalle, Algebraic multigrid for stationary and time-dependent partial differential equations with stochastic coefficients, Numer. Linear Algebra Appl. 15 (2-3) (2008) 141–163. `doi:10.1002/nla.568`.

[33] T. Boonen, J. Van lent, S. Vandewalle, An algebraic multigrid method for high order time-discretizations of the div-grad and the curl-curl equations, Applied Numerical Mathematics 59 (3) (2009) 507–521. `doi:10.1016/j.apnum.2008.03.004`.

[34] K.-A. Mardal, T. K. Nilssen, G. A. Staff, Order-optimal preconditioners for implicit Runge-Kutta schemes applied to parabolic PDEs, SIAM Journal on Scientific Computing 29 (1) (2007) 361–375.

[35] L. O. Jay, Inexact simplified Newton iterations for implicit Runge-Kutta methods, SIAM Journal on Numerical Analysis 38 (4) (2000) 1369–1388. `doi:10.1137/S0036142999360573`.

[36] H. Chen, A splitting preconditioner for the iterative solution of implicit Runge-Kutta and boundary value methods, BIT 54 (3) (2014) 607–621. `doi:10.1007/s10543-014-0467-3`.

[37] M. M. Rana, V. E. Howle, K. Long, A. Meek, W. Milestone, A new block preconditioner for implicit Runge-Kutta methods for parabolic PDE, SIAM J. Sci. Comp. 43 (5) (2021) S475–S495.

[38] B. S. Southworth, O. Krzysik, W. Pazner, Fast parallel solution of fully implicit Runge-Kutta and discontinuous Galerkin in time for numerical PDEs, Part II: nonlinearities and DAEs (2021). `arXiv:2101.01776`.

[39] R. Hiptmair, Multigrid method for Maxwell's equations, SIAM J. Numer. Anal. 36 (1) (1999) 204–225.

[40] P. E. Farrell, R. C. Kirby, J. Marchena-Menéndez, Irksome: Automating Runge–Kutta time-stepping for finite element methods, ACM Trans. Math. Softw. 47 (4) (2021). `doi:10.1145/3466168`.

[41] J. C. Butcher, On the implementation of implicit Runge-Kutta methods, BIT Numerical Mathematics 16 (3) (1976) 237–240.

[42] T. A. Bickart, An efficient solution process for implicit Runge–Kutta methods, SIAM Journal on Numerical Analysis 14 (6) (1977) 1022–1027.

[43] G. Wanner, E. Hairer, Solving ordinary differential equations II, Vol. 375, Springer Berlin Heidelberg, 1996.

[44] J. C. Butcher, General linear methods, Acta Numerica 15 (2006) 157–256. `doi:10.1017/S0962492906220014`.

[45] J. C. Butcher, Numerical methods for ordinary differential equations, John Wiley & Sons, 2016.

[46] C. Taylor, P. Hood, A numerical solution of the Navier-Stokes equations using the finite element technique, Computers & Fluids 1 (1) (1973) 73–100.

[47] S. C. Eisenstat, H. F. Walker, Choosing the forcing terms in an inexact Newton method, SIAM Journal on Scientific Computing 17 (1) (1996) 16–32.

[48] A. Schneebeli, D. Schötzau, Mixed finite elements for incompressible magneto-hydrodynamics, Comptes Rendus Mathematique 337 (1) (2003) 71–74.

[49] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, SIAM Journal on Scientific Computing 14 (2) (1993) 461–469.

[50] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, J. Koster, A fully asynchronous multifrontal solver using distributed dynamic scheduling, SIAM Journal on Matrix Analysis and Applications 23 (1) (2001) 15–41.

[51] S. P. Vanka, Block-implicit calculation of steady turbulent recirculating flows, International Journal of Heat and Mass Transfer 28 (11) (1985) 2093–2103.

[52] P. E. Farrell, Y. He, S. P. MacLachlan, A local Fourier analysis of additive Vanka relaxation for the Stokes equations, Numerical Linear Algebra with Applications (2020) e2306.

[53] P. E. Farrell, M. G. Knepley, L. Mitchell, F. Wechsung, PCPATCH: Software for the topological construction of multigrid relaxation methods, ACM Trans. Math. Softw. 47 (3) (2021). `doi:10.1145/3445791`.

[54] F. Rathgeber, D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. T. McRae, G.-T. Bercea, G. R. Markall, P. H. J. Kelly, Firedrake: automating the finite element method by composing abstractions, ACM Transactions on Mathematical Software (TOMS) 43 (3) (2016) 1–27.

[55] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. Gropp, et al., PETSc users manual: Revision 3.10, Tech. rep., Argonne National Lab.(ANL), Argonne, IL (United States) (2018).

[56] R. C. Kirby, L. Mitchell, Solver composition across the PDE/linear algebra barrier, SIAM J. Sci. Comput. 40 (1) (2018) C76–C98.

[57] Software used in 'Monolithic multigrid for Implicit Runge–Kutta discretizations of incompressible fluid flow' (feb 2022). `doi:10.5281/zenodo.6036268`.
URL `https://doi.org/10.5281/zenodo.6036268`

[58] M. Lange, L. Mitchell, M. G. Knepley, G. J. Gorman, Efficient mesh management in firedrake using PETSc DMPLEX, SIAM Journal on Scientific Computing 38 (5) (2016) S143–S155.

[59] Wikipedia, List of Runge-Kutta methods, `https://en.wikipedia.org/wiki/List_of_Runge-Kutta_methods`, [Online; accessed 7-January-2021] (2021).

[60] J. D. Betteridge, P. E. Farrell, D. A. Ham, Code generation for productive, portable, and scalable finite element simulation in firedrake, Computing in Science and Engineering 23 (4) (2021) 8–17. `doi:10.1109/MCSE.2021.3085102`.

[61] D. A. May, P. Sanan, K. Rupp, M. G. Knepley, B. F. Smith, Extreme-scale multigrid components within PETSc, in: Proceedings of the Platform for Advanced Scientific Computing Conference, PASC '16, Association for Computing Machinery, New York, NY, USA, 2016. `doi:10.1145/2929908.2929913`.

[62] A. Reisner, L. N. Olson, J. D. Moulton, Scaling structured multigrid to 500k+ cores through coarse-grid redistribution, SIAM Journal on Scientific Computing 40 (4) (2018) C581–C604.

[63] L. Pareschi, G. Russo, Implicit-Explicit Runge-Kutta schemes and applications to hyperbolic systems with relaxation, J. Sci. Comput. 25 (1-2) (2005) 129–155. `doi:10.1007/s10915-004-4636-4`.

[64] R. Alexander, Diagonally implicit Runge-Kutta methods for stiff O.D.E.'s, SIAM J. Numer. Anal. 14 (6) (1977) 1006–1021. `doi:10.1137/0714068`.

[65] V. John, Reference values for drag and lift of a two-dimensional time-dependent flow around a cylinder, International Journal for Numerical Methods in Fluids 44 (7) (2004) 777–788.

[66] M. Schäfer, S. Turek, F. Durst, E. Krause, R. Rannacher, Benchmark computations of laminar flow around a cylinder, in: Flow simulation with high-performance computers II, Springer, 1996, pp. 547–566.

[67] D. A. Knoll, L. Chacón, Coalescence of magnetic islands, sloshing, and the pressure problem, Physics of Plasmas 13 (3) (2006) 032307.

[68] J. H. Adler, M. Brezina, T. A. Manteuffel, S. F. McCormick, J. W. Ruge, L. Tang, Island coalescence using parallel first-order system least squares on incompressible resistive magnetohydrodynamics, SIAM Journal on Scientific Computing 35 (5) (2013) S171–S191.

[69] E. G. Phillips, J. N. Shadid, E. C. Cyr, H. C. Elman, R. P. Pawlowski, Block preconditioners for stable mixed nodal and edge finite element rep-

resentations of incompressible resistive MHD, SIAM Journal on Scientific Computing 38 (6) (2016) B1009–B1031.