



A Reynolds-robust preconditioner for the Scott–Vogelius discretization of the stationary incompressible Navier–Stokes equations

PATRICK E. FARRELL¹
LAWRENCE MITCHELL²
L. RIDGWAY SCOTT³
FLORIAN WECHSUNG⁴

¹ Mathematical Institute, University of Oxford, Oxford, UK

Email address: patrick.farrell@maths.ox.ac.uk

² Department of Computer Science, Durham University, Durham, UK

Email address: lawrence.mitchell@durham.ac.uk

³ Department of Computer Science, University of Chicago, Chicago, USA

Email address: ridg@uchicago.edu

⁴ Courant Institute of Mathematical Sciences, New York University, New York, USA

Email address: wechsung@nyu.edu

Abstract. Augmented Lagrangian preconditioners have successfully yielded Reynolds-robust preconditioners for the stationary incompressible Navier–Stokes equations, but only for specific discretizations. The discretizations for which these preconditioners have been designed possess error estimates which depend on the Reynolds number, with the discretization error deteriorating as the Reynolds number is increased. In this paper we present an augmented Lagrangian preconditioner for the Scott–Vogelius discretization on barycentrically-refined meshes. This achieves both Reynolds-robust performance and Reynolds-robust error estimates. A key consideration is the design of a suitable space decomposition that captures the kernel of the grad-div term added to control the Schur complement; the same barycentric refinement that guarantees inf-sup stability also provides a local decomposition of the kernel of the divergence. The robustness of the scheme is confirmed by numerical experiments in two and three dimensions.

Navier–Stokes, Scott–Vogelius element, exactly divergence-free, multigrid, preconditioning, Reynolds-robust solvers

65N55, 65F08, 65N30

1. Introduction

The stationary Navier–Stokes equations for the flow of a viscous, isothermal, incompressible, Newtonian fluid on a bounded Lipschitz domain $\Omega \subset \mathbb{R}^d$, $d \in \{2, 3\}$, are given by: find $(u, p) \in H^1(\Omega; \mathbb{R}^d) \times Q$

This research is supported by the Engineering and Physical Sciences Research Council [grant numbers EP/R029423/1 and EP/V001493/1], and by the EPSRC Centre For Doctoral Training in Industrially Focused Mathematical Modelling [grant number EP/L015803/1] in collaboration with London Computational Solutions. LM also acknowledges support from the UK Fluids Network [EPSRC grant number EP/N032861/1] for funding a visit to Oxford. This work used the ARCHER UK National Supercomputing Service (<http://www.archer.ac.uk>). The authors would like to acknowledge insightful comments and advice from PD Dr Alexander Linke.

such that

$$-\nabla \cdot 2\nu\varepsilon(u) + (u \cdot \nabla)u + \nabla p = f \quad \text{in } \Omega, \quad (1.1a)$$

$$\nabla \cdot u = 0 \quad \text{in } \Omega, \quad (1.1b)$$

$$u = g \quad \text{on } \Gamma_D, \quad (1.1c)$$

$$2\nu\varepsilon(u) \cdot n = pn \quad \text{on } \Gamma_N, \quad (1.1d)$$

where $\varepsilon(u) = \frac{1}{2}(\nabla u + \nabla u^\top)$, $\nu > 0$ is the kinematic viscosity, $f \in L^2(\Omega; \mathbb{R}^d)$, n is the outward-facing unit normal to $\partial\Omega$, Γ_D and Γ_N are disjoint with $\Gamma_D \cup \Gamma_N = \partial\Omega$, and $g \in H^{1/2}(\Gamma_D; \mathbb{R}^d)$. If $|\Gamma_N| > 0$, then a suitable trial space for the pressure is $Q := L^2(\Omega)$; if $|\Gamma_N| = 0$, then the pressure is only defined up to an additive constant and $Q := L_0^2(\Omega)$ is used instead [17]. The Reynolds number UL/ν , where U and L are the characteristic velocity and length scale of the flow, is a dimensionless number governing the nature of the system [36].

For high-Reynolds number flows, it is important that the error estimates should not degrade as the Reynolds number increases. For most discretizations, the velocity error estimates are not robust, as they are polluted by the pressure approximation scaled by the inverse viscosity. For the Stokes equations, the discrete solution $u_h \in V_h$ can be shown to satisfy [33, (3.5)]

$$\|\nabla(u - u_h)\|_{L^2(\Omega)} \leq 2 \inf_{\tilde{u}_h \in \mathcal{N}_h} \|\nabla(u - \tilde{u}_h)\|_{L^2(\Omega)} + \nu^{-1} \inf_{q_h \in Q_h} \|p - q_h\|_{L^2(\Omega)}, \quad (1.2)$$

where $V_h \subset H^1(\Omega; \mathbb{R}^d)$ is the velocity trial space, $Q_h \subset Q$ is the pressure trial space, and

$$\mathcal{N}_h = \{v_h \in V_h : \int_{\Omega} q_h \nabla \cdot v_h \, dx = 0 \, \forall q_h \in Q_h\} \quad (1.3)$$

is the space of discretely divergence-free velocity trial functions.

One way to achieve robustness is to choose a pair $V_h \times Q_h$ such that $\nabla \cdot V_h \subseteq Q_h$, so that $\nabla \cdot u_h = 0$ holds pointwise [33]. If this choice is made, the second term on the right-hand side of the error estimate (1.2) for the Stokes equations can be removed and the velocity error is then independent of both pressure and viscosity. The analysis is more complicated for the Navier–Stokes equations, but progress has recently been made [2]. We also mention that similar results are available for the time-dependent case [42]. Consequently, it is highly advantageous to choose such a discretization, especially for high Reynolds number flows. Such discretizations are termed pressure robust or Reynolds-robust.

There are several choices of element pairs that satisfy this requirement. They include $H(\text{div})$ - $L^2(\Omega)$ discretizations, such as the Raviart–Thomas and Brezzi–Douglas–Marini families [16, 60, 32]; hybrid discontinuous Galerkin schemes [39]; or H^1 conforming approaches such as the Scott–Vogelius \mathbb{P}_k - $\mathbb{P}_{k-1}^{\text{disc}}$ pair [55, 56]. The former two options allow for arbitrary order approximations, but are nonconforming: the discretization of the momentum equation requires penalty terms. In contrast, the Scott–Vogelius pair is straightforward to implement, but is only inf-sup stable on certain types of meshes and for certain polynomial degrees. In this work, we employ the Scott–Vogelius element. In [45], Olshanskii & Rebholz demonstrate the accuracy of this discretization and investigate the numerical performance of direct sparse solvers. While these solvers perform very well for problems with sizes of the order of millions of degrees of freedom, they do not scale and a different strategy is required to solve larger problems.

In this work, we build on the insights of [7, 32, 25, 24] to develop a Reynolds-robust block preconditioner for the Reynolds-robust Scott–Vogelius discretization of the Navier–Stokes equations. When building block preconditioners, the usual difficulty is developing a good approximation for the inverse of the Schur complement. We employ an augmented Lagrangian term to control the Schur complement of the system; this simplifies the approximation of the Schur complement, at the cost of making the momentum equation significantly more difficult to solve. We then develop a specialized geometric multigrid scheme for the resulting augmented momentum equation. This approach is based on the

work of Schöberl [53]. Its application to the Navier–Stokes equations was pioneered in two dimensions by Benzi & Olshanskii [7], and has recently been extended to three dimensions [25]. A similar strategy has proven successful for a $H(\text{div})$ - $L^2(\Omega)$ discretization of the Stokes equations [32]. While effective at controlling iteration counts as the Reynolds number is varied, the schemes presented in [7, 25] heavily rely on the use of piecewise constant pressure functions and are not effective for Scott–Vogelius discretizations.

Many alternative approaches to solving (1.1) have been considered in the literature. These include the pressure convection-diffusion (PCD), least-squares commutator (LSC) and SIMPLE block preconditioners [46, 34, 19, 20], monolithic multigrid approaches [59, 58], and a modified augmented Lagrangian strategy which trades off control of the Schur complement against ease of solving the augmented momentum block [8]. These preconditioners do not generally enjoy Reynolds-robustness.

The remainder of this manuscript is structured as follows. In section 2 we discuss the discretization and the conditions under which inf-sup stability of the Scott–Vogelius element is known. In section 3 we recall the augmented Lagrangian preconditioning strategy and the difficulties it introduces for solving the momentum equations. A Reynolds-robust multigrid cycle for the augmented Scott–Vogelius momentum operator is introduced in section 4. Finally, numerical examples in two and three dimensions are presented in section 5.

2. Discretization

We begin by recalling that finite dimensional subspaces $V_h \subset V$ and $Q_h \subset Q$ are said to satisfy the inf-sup condition if there exists a $\gamma > 0$ such that

$$\inf_{\substack{q_h \in Q_h \\ q_h \neq 0}} \sup_{\substack{v_h \in V_h \\ v_h \neq 0}} \frac{(\nabla \cdot v_h, q_h)}{\|v_h\| \|q_h\|} \geq \gamma \quad (2.1)$$

for all mesh sizes $h > 0$. Intuitively, the inf-sup condition encourages large velocity and small pressure spaces, but the opposite is true for the condition $\nabla \cdot V_h \subseteq Q_h$. For this reason it is difficult to construct discretizations that satisfy both of these properties. Most of the popular inf-sup stable finite element discretizations of the Stokes and Navier–Stokes equations (such as the Taylor–Hood, the MINI or the $[\mathbb{P}_2]^2 - \mathbb{P}_0$ elements) do not satisfy $\nabla \cdot V_h \subseteq Q_h$.

The Scott–Vogelius element is given by choosing continuous piecewise polynomials of degree k for the velocity and discontinuous piecewise polynomials of degree $k-1$ for the pressure. While this clearly implies that $\nabla \cdot V_h \subseteq Q_h$, inf-sup stability of the Scott–Vogelius element is more delicate, and is a topic of ongoing research. In two dimensions, Scott & Vogelius proved [56] that the element is inf-sup stable for $k \geq 4$ if the mesh does not have nearly singular vertices. In three dimensions, it was proven more recently in [67] that the element is stable for $k \geq 6$ on uniform meshes. The stability on general tetrahedral meshes continues to be an open question [44].

On barycentrically refined meshes, however, the pair is known to be stable for polynomial order $k = d$, see [48, Section 4.6] for the 2D case and [65] for the 3D case. If one is willing to consider the more complicated Powell–Sabin split, the order can be reduced further to $k = d-1$ [66, 68]. The two refinement patterns are shown for the two dimensional case in Figure 1. In this work we will consider the case of $k \geq d$ on barycentrically refined meshes, but the arguments apply *mutatis mutandis* to the Powell–Sabin split.

In the context of the multigrid scheme that we will develop in section 4, the requirement for barycentrically refined elements has some implications for our mesh hierarchy. First, note that repeatedly barycentrically refining a mesh leads to degenerate elements. Furthermore, it is not known whether regularly refining a mesh, on which an element pair with $k = d$ is stable, always results in a refined mesh for which stability is maintained. Consequently, we build the multigrid hierarchy in a different

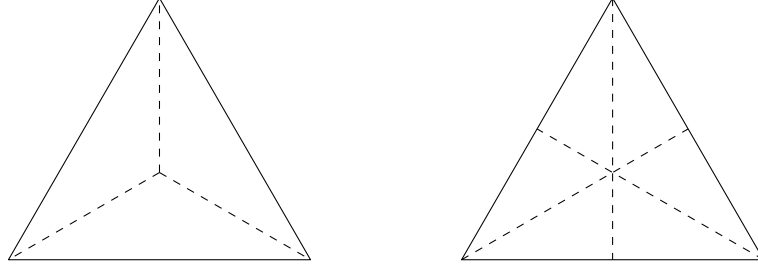


FIGURE 1. Barycentrically refined triangle (also known as Alfeld split) on the left, and Powell–Sabin split on the right.

way. Given a domain Ω , we consider a given simplicial mesh $\mathcal{M}_h = \{K^h\}$ with $\cup_{K^h \in \mathcal{M}_h} K^h = \overline{\Omega}$ and $(K_i^h)^\circ \cap (K_j^h)^\circ = \emptyset$ for $i \neq j$. The elements $K^h \in \mathcal{M}_h$ will be referred to as the *macro cells*. For each level h , we obtain the mesh $\hat{\mathcal{M}}_h$ by barycentric refinement: that is, for each macro cell $K^h \in \mathcal{M}_h$ we obtain $d + 1$ many cells \hat{K}_i^h , $0 \leq i \leq d$ and

$$\hat{\mathcal{M}}_h = \{\hat{K}_i^h : 0 \leq i \leq d, K^h \in \mathcal{M}_h\}. \quad (2.2)$$

The function spaces on $\hat{\mathcal{M}}_h$ are then given by

$$V_h := \{\mathbf{v} \in H^1(\Omega; \mathbb{R}^d) : \mathbf{v}|_{\hat{K}} \in [\mathbb{P}^d(\hat{K})]^d \ \forall \hat{K} \in \hat{\mathcal{M}}_h\}, \quad (2.3)$$

$$Q_h := \{q \in L^2(\Omega) : q|_{\hat{K}} \in \mathbb{P}^{d-1}(\hat{K}) \ \forall \hat{K} \in \hat{\mathcal{M}}_h\}. \quad (2.4)$$

We construct the hierarchy as follows. We start with an initial coarse triangulation of the domain, given by \mathcal{M}_H . We obtain \mathcal{M}_h , $h = \frac{1}{2}H$, by uniform refinement of the initial mesh. Both \mathcal{M}_H and \mathcal{M}_h are then refined barycentrically to obtain $\hat{\mathcal{M}}_H$ and $\hat{\mathcal{M}}_h$. Note that though \mathcal{M}_H and \mathcal{M}_h form a nested hierarchy, this is not true for $\hat{\mathcal{M}}_H$ and $\hat{\mathcal{M}}_h$. This two-level approach canonically extends to many levels; a hierarchy of three levels is shown in Figure 2.

We will see in Section 4.1 that this macro element structure not only guarantees inf-sup stability, but is also crucial in defining a robust relaxation method.

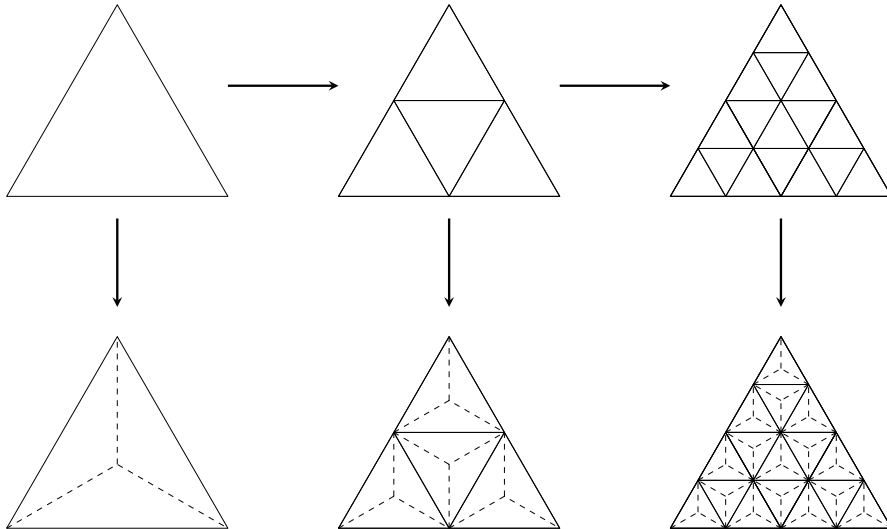


FIGURE 2. A three level barycentrically refined multigrid hierarchy.

3. Variational formulation and augmented Lagrangian strategy

For boundary data $g \in H^{1/2}(\Gamma_D)$, let

$$V_g = \{v \in H^1(\Omega; \mathbb{R}^d) : v|_{\Gamma_D} = g\}. \quad (3.1)$$

The weak form of (1.1) is: find $(u, p) \in V_g \times Q$ such that

$$\int_{\Omega} 2\nu \varepsilon(u) : \varepsilon(v) \, dx + \int_{\Omega} (u \cdot \nabla) u \cdot v \, dx - \int_{\Omega} p \nabla \cdot v \, dx - \int_{\Omega} q \nabla \cdot u \, dx = \int_{\Omega} f \cdot v \, dx, \quad (3.2)$$

for all $(v, q) \in V_0 \times Q$.

Given finite dimensional subspaces $V_h \subset H^1(\Omega; \mathbb{R}^d)$ and $Q_h \subset Q$ and after applying Newton's method to the nonlinear equations, at each Newton step we must solve a nonsymmetric linear system of generalized saddle point structure:

$$\begin{pmatrix} A & B^\top \\ B & 0 \end{pmatrix} \begin{pmatrix} \delta u \\ \delta p \end{pmatrix} = \begin{pmatrix} b \\ c \end{pmatrix}, \quad (3.3)$$

where A is the discrete linearized momentum operator, B^\top the discrete gradient, B the discrete divergence, and δu and δp are the updates to the velocity and pressure solutions respectively. This system becomes increasingly difficult to solve as the Reynolds number is increased.

There are two key ingredients when building block preconditioners for (3.3): an effective solver for A and an effective solver for the Schur complement $S = -BA^{-1}B^\top$. Since S is usually dense, tractable approximations to S^{-1} must be developed on a PDE-specific basis. The main issue with PCD, LSC and SIMPLE is that their choice for the approximate Schur complement becomes a poor approximation to the true Schur complement as the Reynolds number is increased, which in turn results in significant growth of the iteration counts. In [7], Benzi & Olshanskii proposed an augmented Lagrangian strategy that significantly simplifies the approximation of the Schur complement: for $\gamma > 0$, the linear system (3.3) is augmented by adding a term to the top-left block and adjusting the residual accordingly:

$$\begin{bmatrix} A + \gamma B^\top M_p^{-1} B & B^\top \\ B & 0 \end{bmatrix} \begin{bmatrix} \delta u \\ \delta p \end{bmatrix} = \begin{bmatrix} b + \gamma B^\top M_p^{-1} c \\ c \end{bmatrix}, \quad (3.4)$$

where M_p is the mass matrix for the pressure space. It is immediately clear that this modification does not change the solution of the linear system. Furthermore, one can show that the Schur complement of the augmented system, \tilde{S} , satisfies

$$\tilde{S}^{-1} = S^{-1} - \gamma M_p^{-1}, \quad (3.5)$$

where S is the Schur complement of the original system. The advantage is clear: as $\gamma \rightarrow \infty$, \tilde{S}^{-1} can be approximated by a scaled inverse pressure mass matrix, which is easy to solve.

In general, a triple matrix product as it occurs in the augmented Lagrangian term is both expensive to compute and store. However, it is straightforward to check that adding $\gamma B^\top M_p^{-1} B$ to the linear system corresponds to augmenting the weak form with a term

$$\gamma \int_{\Omega} \Pi_{Q_h}(\nabla \cdot u) \Pi_{Q_h}(\nabla \cdot v) \, dx, \quad (3.6)$$

where Π_{Q_h} is the projection onto Q_h .

The same augmentation, without the projection onto Q_h , is known as grad-div stabilization as it corresponds to the weak form of $\gamma \nabla \nabla \cdot u$. We remark that this form does not lead to any additional non-zeros in the assembled linear system. As we focus on discretizations that satisfy $\nabla \cdot V_h \subseteq Q_h$, in this work the projection onto the pressure space is always the identity operator and hence the augmented Lagrangian and grad-div stabilization coincide.

To summarize, the augmented problem in weak form reads: find $(u, p) \in (V_h \cap V_g) \times Q_h$ such that

$$\begin{aligned} \int_{\Omega} 2\nu \varepsilon(u) : \varepsilon(v) \, dx + \int_{\Omega} (u \cdot \nabla) u \cdot v \, dx - \gamma \int_{\Omega} \nabla \cdot u \nabla \cdot v \, dx \\ - \int_{\Omega} p \nabla \cdot v \, dx - \int_{\Omega} q \nabla \cdot u \, dx = \int_{\Omega} f \cdot v \, dx, \end{aligned} \quad (3.7)$$

for all $(v, q) \in (V_h \cap V_0) \times Q_h$.

Remark 3.1. A discretization for which the augmented Lagrangian and the grad-div stabilization *do not* coincide is the classical $[\mathbb{P}_k]^d - \mathbb{P}_{k-1}$ Taylor–Hood element. While the Taylor–Hood element does not enforce the divergence constraint exactly, it can be shown that when using grad-div stabilization, in the limit of $\gamma \rightarrow \infty$ the solution converges to that obtained when using the Scott–Vogelius discretization, so long as the Scott–Vogelius element is stable [14, Theorem 3.1].

It turns out that not only are the solutions obtained from these two elements related, but also that (on the particular meshes considered here) the solver developed for the Scott–Vogelius element can be used for the Taylor–Hood element. First observe that since the pressure space is continuous, the projection onto Q_h in the augmented Lagrangian term introduces additional coupling and adding (3.6) increases the number of non-zeros in the assembled matrix. However, it was shown in [31, Lemma 2] that in the limit of $h \rightarrow 0$ grad-div stabilization converges to the discrete augmented Lagrangian approach. Hence one can simply use grad-div stabilization and still use the Schur complement approximation in (3.5). Clearly, the obtained top-left block is then identical to that obtained from the Scott–Vogelius element and hence the multigrid scheme that we propose in the next section applies in the same way. From a solver point of view, a minor disadvantage of this approach is that the pressure mass matrix corresponding to a continuous finite element space is more expensive to invert than that of a discontinuous space.

4. Solving the top-left block

Although the augmented Lagrangian approach is appealing since it provides for an excellent Schur complement approximation, it introduces some new challenges to developing fast solvers for the equations as a whole. In contrast to A , the performance of standard multigrid schemes deteriorates quickly for $A_\gamma = A + \gamma B^\top M_p^{-1} B$ as γ is increased. The additional term has a large kernel consisting of solenoidal vector-fields, implying that the problem becomes nearly singular as γ increases. If a scalable solver for the nearly singular top-left block can be developed that is robust to ν and γ , the overall solver will be Reynolds-robust.

Multigrid schemes for nearly singular problems have been studied previously by Schöberl [52, 53, 54] and Lee et al. [38, 37], and a similar analysis was carried out for overlapping Schwarz methods by Ewing and Wang [21]. While these works only consider the symmetric case, they clearly demonstrate that a necessary condition for a γ -robust scheme is a good understanding and characterization of the kernel of the semidefinite term, which in this case means understanding the kernel of the divergence operator. This has implications for both multigrid relaxation and prolongation, as we now consider.

4.1. Relaxation

The core requirement for obtaining a parameter robust relaxation is that the space decomposition defining the relaxation needs to provide a decomposition of the kernel of the singular operator. Classical point relaxation methods such as point-block Jacobi do not satisfy this property and are ineffective for this problem; their smoothing strength degrades as γ increases.

REYNOLDS-ROBUST PRECONDITIONERS FOR THE NSE

Many smoothers can be expressed as so-called subspace correction methods [63]. We consider a decomposition

$$V_h = \sum_i V_i, \quad (4.1)$$

where the sum is not necessarily a direct sum. This decomposition naturally defines an associated additive Schwarz relaxation method. For each subspace i we then denote the natural inclusion by $I_i : V_i \rightarrow V_h$ and we define the restriction A_i of A onto V_i as

$$(A_i u_i, v_i) = (A I_i u_i, I_i v_i) \quad \text{for all } u_i, v_i \in V_i. \quad (4.2)$$

Now denoting

$$D_h^{-1} = \sum_i I_i A_i^{-1} I_i^\top \quad (4.3)$$

and introducing a damping parameter $\tau > 0$, we can express one update of the additive Schwarz method as

$$u_{k+1} = u_k + \tau D^{-1}(f - A u_k). \quad (4.4)$$

The method is also known as the parallel subspace correction method [63]. For the choice $V_i = \text{span}(\{\varphi_i\})$, we recover the classical Jacobi iteration. Usually each subspace is described by an index set J_i and $V_i = \text{span}(\{\varphi_j : j \in J_i\})$. In that case we also speak of a block Jacobi method.

Both Schöberl (for additive relaxation) and later Lee et al. (for multiplicative relaxation) recognize that the key condition for a γ -robust relaxation is that the subspaces V_i need to satisfy a *kernel decomposition property*:

$$\mathcal{N}_h = \sum_i (V_i \cap \mathcal{N}_h), \quad (4.5)$$

where we recall that \mathcal{N}_h is the space of discretely divergence-free vector-fields. In essence, we require that any kernel function can be written as the sum of kernel functions drawn from the subspaces.

When considering the $[\mathbb{P}_2]^2 - \mathbb{P}_0$ element on a 2D triangulation \mathcal{M}_h , that is

$$\begin{aligned} V_h &= \{\mathbf{v} \in H^1(\Omega; \mathbb{R}^d) : \mathbf{v}|_K \in [\mathbb{P}^2(K)]^d \ \forall K \in \mathcal{M}_h\}, \\ Q_h &= \{q \in L^2(\Omega) : q|_K \equiv \text{const} \ \forall K \in \mathcal{M}_h\}, \end{aligned} \quad (4.6)$$

Schöberl proved that the space decomposition defined by

$$V_i := \{\mathbf{u}_h \in V_h : \text{supp}(\mathbf{u}_h) \subset \text{star}(v_i)\}, \quad (4.7)$$

where

$$\text{star}(v_i) := \bigcup_{K \in \mathcal{M}_h : v_i \in K} K, \quad (4.8)$$

satisfies the kernel decomposition property (4.5). The subspaces are illustrated in Figure 3. This




FIGURE 3. The star patch satisfies the kernel decomposition property (4.5) for the $[\mathbb{P}_2]^2 - \mathbb{P}_0$ element.

decomposition was subsequently used in [7] for the Navier–Stokes equations in two dimensions and in [25] in three dimensions. It was also used for $H(\text{div})$ - $L^2(\Omega)$ discretizations of the Stokes and linear elasticity equations in [32], and provides a robust relaxation method for the $H(\text{div})$ and $H(\text{curl})$ Riesz maps [4, 5].

However, the proof that this element pair satisfies the kernel decomposition property depends on the pressure space being piecewise constant, and does not generalise to either Taylor–Hood or conforming divergence-free finite elements. In fact, the same choice of space decomposition applied to a barycentrically refined mesh does not result in a γ -robust smoother for the Scott–Vogelius element considered here.

To find a space decomposition that decomposes the kernel of the divergence operator for the Scott–Vogelius element, we consider the following Hilbert complexes in two

$$\mathbb{R} \xrightarrow{\text{id}} H^2 \xrightarrow{\text{curl}} H^1 \xrightarrow{\text{div}} L^2 \xrightarrow{\text{null}} 0 \quad (4.9)$$

and three dimensions

$$\mathbb{R} \xrightarrow{\text{id}} H^2 \xrightarrow{\text{grad}} H^1(\text{curl}) \xrightarrow{\text{curl}} H^1 \xrightarrow{\text{div}} L^2 \xrightarrow{\text{null}} 0, \quad (4.10)$$

where $H^1(\text{curl}) = \{u \in (H^1)^3, \text{curl } u \in (H^1)^3\}$. These sequences are exact on simply connected domains, which implies that every field u in the kernel of the divergence can be represented as the curl of a potential Φ . If we are given a discrete subcomplex of the form

$$\dots \rightarrow \Sigma_h \xrightarrow{\text{curl}} V_h \xrightarrow{\text{div}} Q_h \xrightarrow{\text{null}} 0, \quad (4.11)$$

then for a divergence-free discrete vector-field $u_h \in V_h$, we can write it as the curl of a potential $\Phi_h \in \Sigma_h$. Writing $\Phi_h = \sum_j \Phi_j$ in terms of basis functions of Σ_h , we can then define a divergence-free decomposition of u_h as $u_h = \sum_j \nabla \times \Phi_j$. Hence, a space decomposition such that $\nabla \times \Phi_j$ is contained in some V_i for any basis function Φ_j decomposes the kernel and satisfies (4.5).

In two dimensions, the Scott–Vogelius velocity and pressure spaces V_h and Q_h form an exact sequence with Σ_h chosen as the HCT finite element space [15, §6.1]. The three elements are displayed in Figure 4. For a given vertex v_i in the macro mesh \mathcal{M}_h , we define the macrostar(v_i) of the vertex




FIGURE 4. 2D exact Stokes complex.

as the union of all macro elements touching the vertex. We then see that for every basis function Φ_j there exists a vertex v_i such that $\text{supp}(\Phi_j) \subset \text{macrostar}(v_i)$. Hence also $\text{supp}(\nabla \times \Phi_j) \subset \text{macrostar}(v_i)$ and if we define

$$V_i = \{v \in V_h : \text{supp}(v) \subset \text{macrostar}(v_i)\} \quad (4.12)$$

then these subspaces decompose the kernel. The macrostar is shown in Figure 5.

A similar argument can be made in three dimensions. Here the existence of an exact sequence is given by the recent work of Fu, Guzmán, and Neilan [27].

We note that the proof for robustness of the multigrid scheme in [53] actually has a stricter requirement than simply the kernel decomposition property (4.5). The kernel decomposition must be *stable*: given a function $u \in \mathcal{N}_h$ with $u = \sum u_i$ and $u_i \in V_i \cap \mathcal{N}_h$, one needs to be able to estimate the norm of




FIGURE 5. Left: The star patch applied to the barycentrically refined mesh does not yield a robust relaxation method for the Scott–Vogelius element. Right: The macrostar patch satisfies the kernel decomposition property (4.5) for the $[\mathbb{P}_k]^d - \mathbb{P}_{k-1}^{\text{disc}}$ element for $k \geq d$ (here shown for $k = d = 2$).

the u_i in terms of u . In general this estimate does not follow purely from the exactness of the discrete sequence, but it was shown in [24] that one can use the existence of a particular Fortin operator to obtain the required bounds.

4.2. Relaxation in the presence of stabilization

It is well known that straightforward Galerkin discretizations of advection-dominated problems are oscillatory [9, 58, 49, 20]. Several approaches have been developed to address these issues, for example by adding a small amount of artificial viscosity as in the case of Streamline Upwind/Petrov Galerkin (SUPG) or Galerkin Least Squares (GLS) or by enriching the space with bubble functions. SUPG stabilization was used in previous works on Reynolds-robust preconditioners [7, 25]. As both of these works consider piecewise constant pressures, the pressure gradient on each cell vanishes and hence only the top-left block is modified when adding SUPG stabilization. For the Scott–Vogelius element this is not the case, and SUPG stabilization results in a modification of the top-right block, introducing additional coupling between the velocity and the pressure [10, p. 1249], [29, p. 250]. This makes the nonlinear problem much more difficult to solve. In fact, in numerical experiments we observe that the outer Newton solver already fails to converge at $\text{Re} = 50$ for a two dimensional lid-driven cavity, even with the use of a direct solver. Furthermore, the modified top-right block needs to be taken into account when adding the augmented Lagrangian term and in the multigrid scheme for the top-left block.

In 1976 Douglas & Dupont [18] suggested an interior penalty scheme that penalises a jump of the derivative across facets:

$$S(u, v) = \sum_{K \in \mathcal{M}_h} \frac{1}{2} \int_{\partial K} \delta h_{\partial K}^2 [\![\nabla u]\!] : [\![\nabla v]\!] \, ds, \quad (4.13)$$

where $[\![\nabla u]\!]$ denotes the jump of the gradient, $h_{\partial K}$ is a function giving the facet size, and δ is a free parameter. The term vanishes when the velocity is C^1 -continuous. This scheme has received renewed attention and it was shown in [11] that it successfully stabilizes advection-dominated problems and has subsequently been used to stabilize the Stokes [12] and the Oseen equations [10, 13]. In addition to not introducing any additional coupling of the pressure and the velocity, an advantage of this scheme is that it is adjoint consistent. This means that for low Reynolds number one can prove that the order of convergence in the L^2 norm is preserved [13, Remark 12]. In fact, for the problem considered in

Section 5.2 we observe optimal convergence even at high Reynolds number, although this is not in general guaranteed.

We now consider the effect of adding (4.13) to the top-left block on the multigrid scheme. Since S vanishes for functions that have continuous gradients, we have added another bilinear form to our system that has a nontrivial kernel consisting of C^1 vector fields. As the weight $\delta h_{\partial K}^2$ is small, the impact is not as significant as that of the grad-div term, but for very high Reynolds number or coarse meshes, we still observe reduced performance of the multigrid scheme. As discussed in the previous section, we know that for the smoother to be robust the space decomposition must provide a decomposition of the kernel. In two dimensions, this is satisfied if $k \geq 3$, as the macrostar around vertices then captures the support of the HCT element. In three dimensions the lowest degree (that the authors are aware of) for a local basis for C^1 vector fields on barycentrically refined meshes is $k = 5$, see [3, 35].

This argument is heuristic, and a full analysis of problems with two different singular terms is out of the scope of this work. Nevertheless, in the numerical experiments we will see that the scheme is noticeably more robust for $k = 3$ in two dimensions and $k = 5$ in three dimensions in the presence of the stabilization (4.13).

Remark 4.1. We conjecture that similar robustness properties will carry over to the case of Powell–Sabin splits. In two dimensions, there is a local C^1 -conforming quadratic basis on Powell–Sabin splits [47], and in three dimensions there is a local C^1 -conforming cubic basis [62]. We therefore expect that the choice $k \geq d$ will provide more robust iteration counts in the presence of stabilization than the minimal $k = d - 1$ required for inf-sup stability on meshes with this macro structure.

4.3. Prolongation

The second key ingredient of the multigrid scheme is a robust prolongation operator. To keep notation simple we consider the case of a two-level scheme and denote the coarse-grid function space by V_H and the fine-grid space by V_h . We denote the standard prolongation operator induced by the finite element interpolation operator in V_h by $P_H : V_H \rightarrow V_h$. Let $u_H \in V_H$ be a coarse-grid function. For the multigrid scheme to be γ -robust, it was shown by Schöberl [53] that the prolongation operator must satisfy

$$\|P_H u_H\|_{A_{h,\gamma}} \leq C \|u_H\|_{A_{H,\gamma}} \quad (4.14)$$

with a constant C independent of γ . Calculating these norms for a divergence-free function $u_H \in V_H$, we observe that

$$\begin{aligned} \|u_H\|_{A_{H,\gamma}}^2 &= \|\nabla u_H\|_{L^2}^2 + \gamma \underbrace{\|\nabla \cdot u_H\|_{L^2}^2}_{=0}, \\ \|P_H u_H\|_{A_{h,\gamma}}^2 &= \|\nabla(P_H u_H)\|_{L^2}^2 + \gamma \|\nabla \cdot (P_H u_H)\|_{L^2}^2. \end{aligned} \quad (4.15)$$

However, since the multigrid hierarchy considered here is non-nested, the interpolation is not exact and u_H being divergence-free does not necessarily imply the same property for $P_H u_H$.

Remark 4.2. We note that this scenario differs from the situation in [53, 7]. There the function spaces are nested, but a discretely divergence-free function on the coarse-grid may not be discretely divergence-free on the fine grid, as Q_h is larger than Q_H .

Inspecting the mesh hierarchy in Figure 2, we notice that the interpolation is exact along the edges of the coarse-grid macro mesh. In turn, this means that the flux across these edges is preserved exactly and hence the interpolated vector-field is divergence-free with respect to pressure functions that are piecewise constant on the macro mesh (i.e. before barycentric refinement):

$$\tilde{Q}_H = \{q \in L^2(\Omega) : q|_K \equiv \text{const } \forall K \in \mathcal{M}_h\}. \quad (4.16)$$

REYNOLDS-ROBUST PRECONDITIONERS FOR THE NSE

For a robust prolongation, we therefore only need to modify the degrees of freedom within a coarse-grid macro cell to remove the divergence within the cell that was created by the interpolation. To this end we define the subspace $\tilde{V}_h \subset V_h$ of functions that vanish on the boundaries of macro cells

$$\tilde{V}_h = \{v \in V_h : v = 0 \text{ on } \partial K \text{ for all } K \in \mathcal{M}_h\}. \quad (4.17)$$

We then solve for $\tilde{u}_h \in \tilde{V}_h$ such that

$$\nu(\varepsilon(\tilde{u}_h), \varepsilon(\tilde{v}_h)) + \gamma(\nabla \cdot \tilde{u}_h, \nabla \cdot \tilde{v}_h) = \gamma(\Pi_{\tilde{Q}_h}(\nabla \cdot (P_H u_H)), \Pi_{\tilde{Q}_h}(\nabla \cdot \tilde{v}_h)) \quad \text{for all } \tilde{v}_h \in \tilde{V}_h. \quad (4.18)$$

It was shown in [24] that then the modified prolongation given by

$$\tilde{P}_H u_H = P_H u_H - \tilde{u}_h \quad (4.19)$$

is continuous in the energy norm with a continuity constant independent of γ . We emphasize that due to the nature of the space \tilde{V}_h , (4.18) decouples into many small, independent solves and hence can be solved efficiently, see Figure 6. In addition, since the problem solved on each coarse macro cell does not vary through the nonlinear iteration, the small matrices can be assembled and factorized once in an initialization step.




FIGURE 6. The robust prolongation operator solves local Stokes problems on each coarse macro cell.

5. Numerical examples

5.1. Details of the algorithm

A graphical depiction of the solver algorithm is given in Figure 7. The essential structure is the same as in [25], with different multigrid components for A_γ . The code is implemented in Firedrake [50] using PETSc [6] and PCPATCH [22]. Since the convergence behavior of Newton's method is not Reynolds-robust, we employ continuation in the Reynolds number to ensure its convergence. A flexible Krylov variant is required as we apply GMRES inside the multigrid relaxation, and hence we use flexible GMRES [51] as the outermost Krylov solver. We use the full block factorization preconditioner

$$P^{-1} = \begin{pmatrix} I & -\tilde{A}_\gamma^{-1} B^\top \\ 0 & I \end{pmatrix} \begin{pmatrix} \tilde{A}_\gamma^{-1} & 0 \\ 0 & \tilde{S}^{-1} \end{pmatrix} \begin{pmatrix} I & 0 \\ -B \tilde{A}_\gamma^{-1} & I \end{pmatrix} \quad (5.1)$$

with the scaled inverse of the (block diagonal) pressure mass matrix as \tilde{S}^{-1} , and one full multigrid cycle of the algorithm described in Section 4 as \tilde{A}_γ^{-1} . Each relaxation sweep conducts 6 (in 2D) or 10 (in 3D) GMRES iterations preconditioned by the additive macrostar iteration. As in [25], the problem

on the coarsest level is solved with the SuperLU_DIST sparse direct solver [41, 40] and uses PETSc's telescoping functionality [43] for improved parallel scalability.

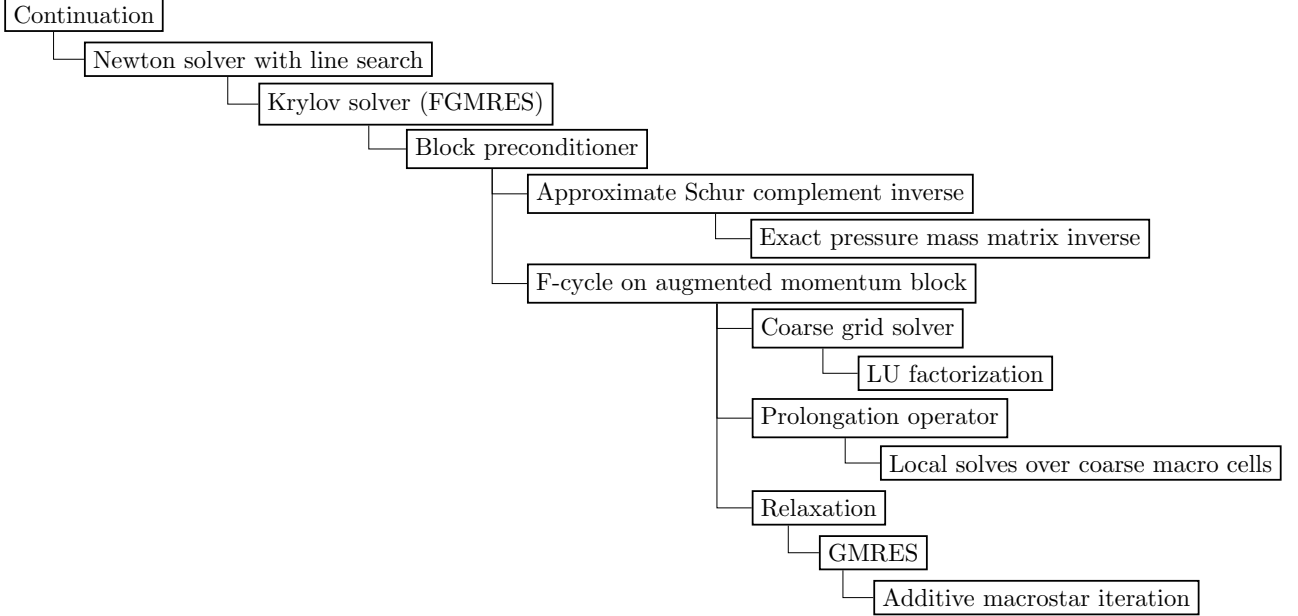


FIGURE 7. An outline of the algorithm for solving (1.1).

5.2. Verification and pressure robustness

We consider the two-dimensional test case of [57] to verify the implementation and to confirm that the velocity errors are independent of the Reynolds number. The example is similar to the lid-driven cavity but with a known analytical solution. Using either the $[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$ Scott–Vogelius or the $[\mathbb{P}_2]^2 - \mathbb{P}_1$ Taylor–Hood element we observe the expected second order convergence of the velocity gradient and of the pressure as the mesh is refined, see Figure 8. In addition, we compare to the $[\mathbb{P}_2]^2 - \mathbb{P}_0$ element used in [7, 25] which converges at first order only.

As motivated in the introduction, since the Taylor–Hood and the $[\mathbb{P}_2]^2 - \mathbb{P}_0$ element do not enforce the divergence constraint exactly (see bottom right of Figure 8), the velocity error increases as the Reynolds number is increased. This is in contrast to the solutions obtained using the Scott–Vogelius element, which are divergence-free up to solver tolerances and exhibit Reynolds-robust errors.

5.3. Two dimensional examples

We begin by considering two classical benchmark problems: the regularised lid-driven cavity and backward-facing step problems, shown in Figures 9 and 10, and described in detail in [20, Examples 8.1.2 and 8.1.3]. We choose relative and absolute tolerances of 10^{-9} and 10^{-8} respectively for the nonlinear solver and 10^{-9} and 10^{-10} for the linear solver. The mesh hierarchy is constructed as described in Section 2. In each case, we start with a given coarse grid, perform the specified number of uniform refinements, and then barycentrically refine each level once.

We perform continuation in the Reynolds number: we start by solving the Stokes equations, and then solve for $\text{Re} = 1$, $\text{Re} = 10$, $\text{Re} = 100$, $\text{Re} = 200$, $\text{Re} = 300$, etc., up to $\text{Re} = 10\,000$, using the previous solution as initial guess for the subsequent Newton iteration. For the backward-facing step cases we add additional continuation steps at $\text{Re} \in \{50, 150, 250, 350\}$. The augmented Lagrangian

REYNOLDS-ROBUST PRECONDITIONERS FOR THE NSE

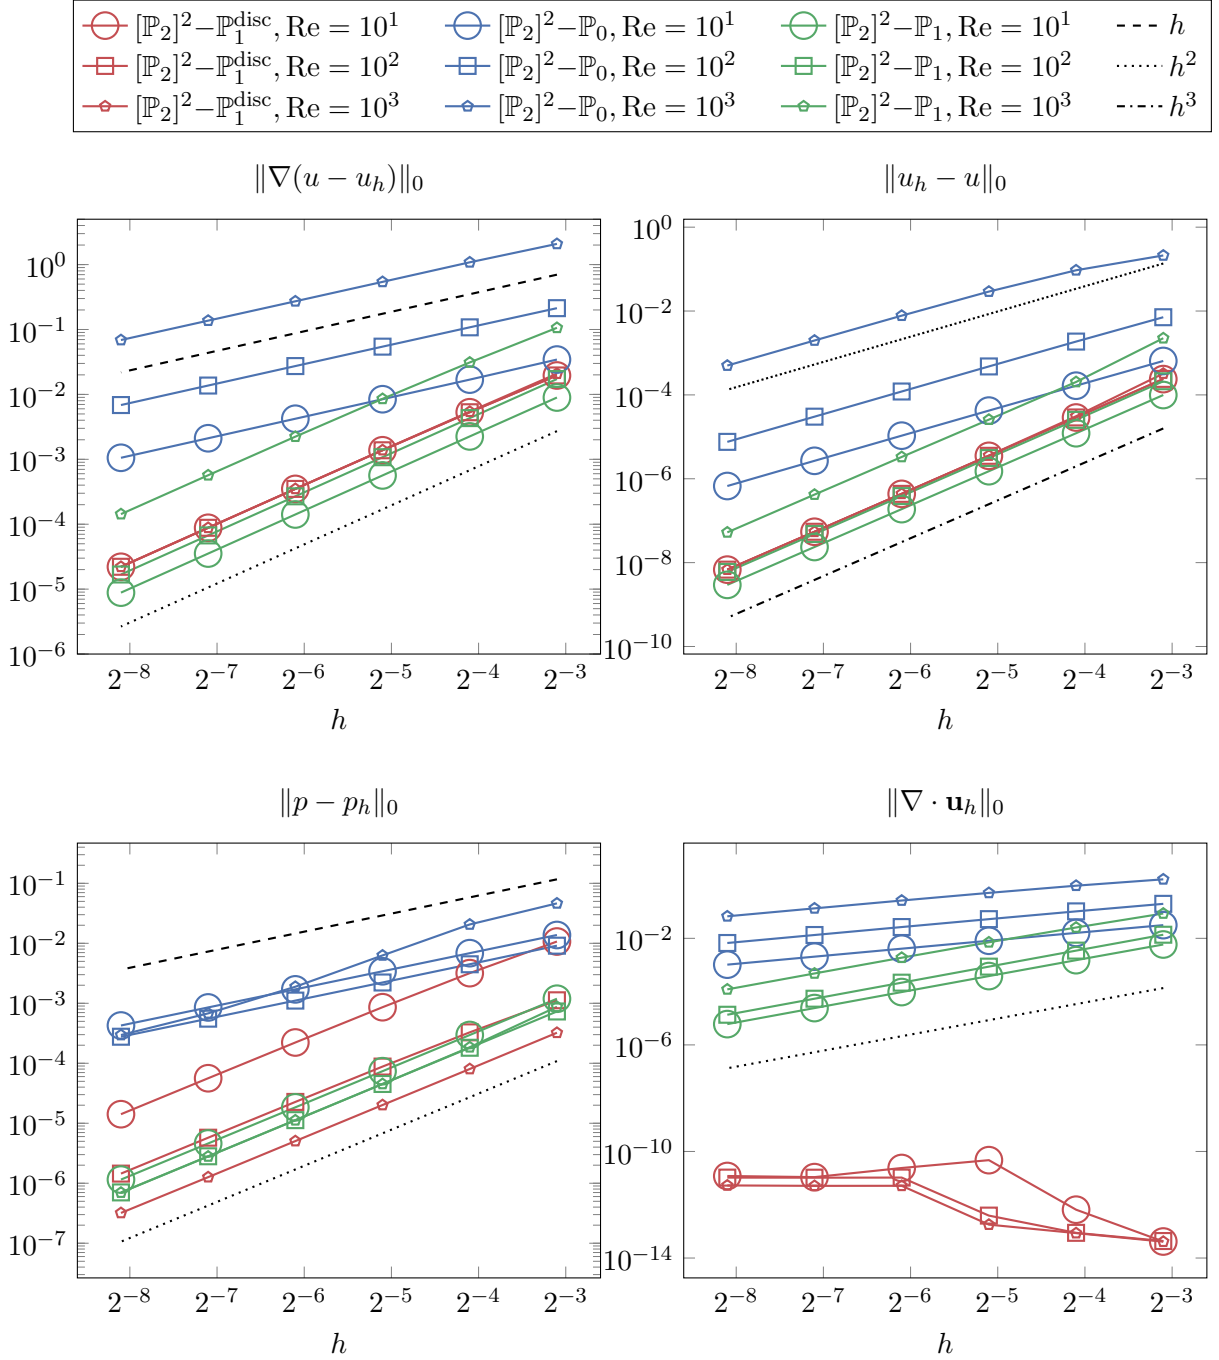


FIGURE 8. Velocity and pressure error as well as L^2 norm of the divergence for the $[\mathbb{P}_2]^2 - \mathbb{P}_0$, $[\mathbb{P}_2]^2 - \mathbb{P}_1$, and $[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$ elements for different Reynolds numbers. The $[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$ element yields velocity errors independent of the Reynolds number, while the other discretizations do not.




FIGURE 9. Streamlines for the two dimensional lid-driven cavity problem at $\text{Re} = 2500$. The domain is given by the $[0, 2] \times [0, 2]$ square. The boundary condition on the top is given by a horizontal velocity field $\mathbf{u}(x, y) = (x^2(2-x)^2, 0)$ and the other boundaries are equipped with a no-slip condition.




FIGURE 10. Streamlines for the two dimensional backwards-facing step problem at $\text{Re} = 200$. The domain is given by $([0, 10] \times [0, 2]) \setminus ([0, 1] \times [0, 1])$. The inflow condition at the top-left boundary is given by a horizontal velocity field $\mathbf{u}(x, y) = (4(2-y)(y-1), 0)$, a natural outflow condition (1.1d) is enforced on the right and the other boundaries are equipped with a no-slip condition.

parameter is set to $\gamma = 10^4$ and we run the solver with and without the stabilization (4.13). In the former case, the stabilization parameter is chosen as $\delta = 5 \cdot 10^{-3}$.

For the lid-driven cavity the coarse mesh is given by a regular 10×10 triangular mesh. Iteration counts are shown in Table 1.

Iteration counts for the backwards-facing step are shown in Table 2. For this case we consider a coarse mesh generated by the frontal meshing algorithm of Gmsh [30] consisting of 5996 triangles.

For both examples, we observe low and nearly flat iteration counts to $\text{Re} = 1000$ for both $k = 2$ and $k = 3$ both with and without stabilization. It is at larger Reynolds numbers that we see differences between the four different configurations. The combination of the $[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$ element and lack of stabilization is least robust. Adding stabilization and increasing the polynomial degree (to additionally capture the kernel of the stabilization, cf. Section 4.2) results in highly robust iteration counts on the finer meshes for both the lid-driven cavity and the backwards-facing step.

Remark 5.1. As discussed in Remark 3.1, we expect that for the Taylor–Hood element on barycentrically-refined meshes the combination of grad-div stabilization and the presented multigrid scheme for the top-left block results in an effective preconditioner. To confirm this we solved the lid-driven cavity using the $[\mathbb{P}_2]^2 - \mathbb{P}_1$ and $[\mathbb{P}_3]^2 - \mathbb{P}_2$ Taylor–Hood elements and obtained essentially identical iteration counts (up to ± 1) to those shown in Table 1.

Clearly the patches in the multigrid relaxation considered here are significantly larger than the star patches considered in [7, 25]. To investigate the impact this has on performance, in Table 3

REYNOLDS-ROBUST PRECONDITIONERS FOR THE NSE

Refinements	Degrees of freedom	Reynolds number									
		10		100		1 000		5 000		10 000	
$[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$											
1	1.70×10^4	4.50 4.50	5.00 5.00	6.67 5.67	18.00 25.00	25.50 >100					
2	6.75×10^4	4.00 4.00	4.33 4.33	5.67 4.33	16.00 16.50	29.00 56.50					
3	2.69×10^5	4.00 4.00	4.00 4.00	4.00 3.00	10.00 9.50	18.00 27.00					
4	1.08×10^6	3.00 3.00	3.33 3.33	2.67 2.67	6.50 6.50	11.50 17.50					
$[\mathbb{P}_3]^2 - \mathbb{P}_2^{\text{disc}}$											
1	3.62×10^4	2.50 2.50	2.67 2.67	3.33 2.67	8.00 8.00	11.50 15.00					
2	1.44×10^5	2.50 2.50	2.67 2.67	2.33 2.00	5.50 4.50	9.50 8.50					
3	5.77×10^5	2.00 2.00	2.67 2.67	2.00 1.67	4.00 3.00	5.50 5.00					
4	2.31×10^6	2.00 2.00	2.67 2.67	2.00 2.00	2.50 2.00	5.00 4.00					

TABLE 1. Average number of outer Krylov iterations per Newton step for the 2D regularised lid-driven cavity problem using the Scott–Vogelius discretization with and without stabilization for $k = 2$ and $k = 3$.

Refinements	Degrees of freedom	Reynolds number					
		10	100	1 000	5 000	10 000	
$[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$							
1	4.79×10^5	3.67 3.67	3.25 3.25	5.00 6.50	14.50 >100	19.50 >100	
2	1.91×10^6	3.67 3.67	3.25 3.25	4.00 4.00	8.50 25.00	13.50 >100	
3	7.64×10^6	4.50 4.50	4.33 4.00	3.00 3.00	5.00 12.50	9.50 >100	
$[\mathbb{P}_3]^2 - \mathbb{P}_2^{\text{disc}}$							
1	1.02×10^6	2.00 2.00	2.00 2.00	2.50 2.50	4.50 11.00	6.00 >100	
2	4.10×10^6	2.50 2.50	2.33 2.33	1.50 1.50	2.00 3.50	3.00 10.50	
3	1.64×10^7	2.50 2.50	3.33 3.33	2.00 2.00	1.50 2.00	2.00 4.50	

TABLE 2. Average number of outer Krylov iterations per Newton step for the 2D backwards-facing step problem using the Scott–Vogelius discretization with and without stabilization for $k = 2$ and $k = 3$.

we compare the runtime of the implementation in [25] using the $[\mathbb{P}_2]^2 - \mathbb{P}_0$ with the $[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$ and $[\mathbb{P}_3]^2 - \mathbb{P}_2^{\text{disc}}$ discretization presented here. Since the mesh for the $[\mathbb{P}_2]^2 - \mathbb{P}_0$ does not require barycentric refinement, we choose a finer 16×16 grid, to obtain a problem of comparable size. As expected, the solver is more expensive for the Scott–Vogelius discretization, as it requires the larger macrostar. However, we emphasize that the improved robustness of the $[\mathbb{P}_3]^2 - \mathbb{P}_2^{\text{disc}}$ element implies that this discretization is very attractive, especially at higher Reynolds number: it converges at higher order but the computational cost is only $\sim 50\%$ bigger than that of the $[\mathbb{P}_2]^2 - \mathbb{P}_0$ and $[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$ discretization on the same mesh.

We study one final two dimensional problem. As discussed in [28], the nature of flows at high Reynolds number is dependent on the Helmholtz decomposition of the convection term. Ahmed et

Discretization	Degrees of freedom	Reynolds number				
		10	100	1 000	5 000	10 000
$[\mathbb{P}_2]^2 - \mathbb{P}_0$	6.57×10^5	3.38	3.74	4.56	8.85	9.64
$[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$	1.08×10^6	10.39	10.84	9.06	18.63	32.31
$[\mathbb{P}_3]^2 - \mathbb{P}_2^{\text{disc}}$	2.31×10^6	25.83	30.28	24.73	28.43	48.07

TABLE 3. Runtime per Newton step (in seconds) for the 2D regularised lid-driven cavity using the $[\mathbb{P}_2]^2 - \mathbb{P}_0$, $[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$, and $[\mathbb{P}_3]^2 - \mathbb{P}_2^{\text{disc}}$ element pairs (with stabilization). Measured on two Intel(R) Xeon(R) Gold 5118 CPUs running 12 MPI processes.

al. [1] introduce the seminorm

$$|f| = \sup_{0 \neq v \in (H_0^1)^d, \nabla \cdot v = 0} \frac{|\langle f, v \rangle|}{\|\nabla v\|_{L^2}}. \quad (5.2)$$

This seminorm vanishes for gradient fields, i.e. when $\langle f, v \rangle = (\nabla \varphi, v)_{L^2}$ for some $\varphi \in H_0^1$. In the case when $|(u \cdot \nabla)u| = 0$ the non-linear convective term and the pressure gradient approximately balance. These types of flows are referred to as *generalized Beltrami flows* and it is particularly for these types of flows that exactly divergence-free methods were shown to outperform standard discretizations [28]. In particular, both of the previously studied examples are close to generalized Beltrami flows at high Reynolds numbers. However, in general, flows may contain a significant divergence-free component in the non-linear term. To demonstrate that the proposed preconditioner is efficient for this type of flow, we consider the Oseen problem proposed in [1, §4.4] with exact velocity $u(x, y) = (\sin(2\pi x) \sin(2\pi y), \cos(2\pi x) \cos(2\pi y))^\top$, pressure $p = \frac{1}{4}(\cos(4\pi x) - \cos(4\pi y))$ and wind $\beta = u + (0, 1)^\top$ on a domain $\Omega = (0, 1)^2$. Since this problem is linear we no longer perform continuation in the Reynolds number, and instead start the linear solver with a zero initial guess for each Reynolds number.

Iteration counts are shown in Table 4 — the results are qualitatively similar to those for the backwards-facing step: robust iteration counts both with and without stabilization to $\text{Re} = 1000$. Robust iteration counts to $\text{Re} = 10000$ are obtained when using stabilization and a cubic velocity space.

5.4. Three dimensional examples

We now study three dimensional variants of the lid-driven cavity and backwards-facing step problems; these are described in detail in [25, §5.5]. The solver tolerances are all relaxed to 10^{-8} in three dimensions. We study iteration counts both with and without adding the stabilization terms in (4.13).

Results for the lid-driven cavity are shown in Table 5. Both with and without stabilization we observe iteration counts that approximately double as the Reynolds number is increased from $\text{Re} = 10$ to $\text{Re} = 5000$.

We show results for the backward-facing step in Table 6. Without stabilization, we observe iteration counts that approximately triple over the same range of Reynolds numbers. However, when adding stabilization, iteration counts increase significantly and blow up for very high Reynolds number. We attribute this to the issue raised in Section 4.2: the stabilization term itself has a large nullspace (consisting of C^1 vector fields) that is not captured by the relaxation induced by the macrostar around vertices. If we choose $k = 5$ we know that a local basis for C^1 functions exists. Indeed, we see in Table 6 that iteration counts for the $[\mathbb{P}_5]^3 - \mathbb{P}_4^{\text{disc}}$ element are significantly more robust.

The results in Tables 5 and 6 were obtained on the ARCHER supercomputer. To provide an impression of computational performance, the results on the twice refined mesh of Table 6 were run on

REYNOLDS-ROBUST PRECONDITIONERS FOR THE NSE

Refinements	Degrees of freedom	Reynolds number				
		10	100	1 000	5 000	10 000
$[\mathbb{P}_2]^2 - \mathbb{P}_1^{\text{disc}}$						
1	1.70×10^4	3 3	3 3	6 13	10 >100	10 >100
2	6.75×10^4	3 3	3 2	6 11	15 >100	18 >100
3	2.69×10^5	3 3	2 2	4 7	15 >100	19 >100
4	1.08×10^6	3 3	2 2	3 5	12 >100	19 >100
$[\mathbb{P}_3]^2 - \mathbb{P}_2^{\text{disc}}$						
1	3.62×10^4	2 2	2 2	3 3	5 28	5 >100
2	1.44×10^5	2 2	2 2	3 3	4 21	6 >100
3	5.77×10^5	2 2	2 2	2 2	4 11	5 70
4	2.31×10^6	2 2	2 2	2 2	3 6	4 25

TABLE 4. Number of Krylov iterations at each Reynolds number for the example in [1, §4.4] with and without stabilization for $k = 2$ and $k = 3$.

Refinements	Degrees of freedom	Reynolds number				
		10	100	1 000	2 500	5 000
$k = 3$ without stabilization						
1	1.03×10^6	3.00	3.67	3.50	4.00	5.00
2	8.22×10^6	3.50	3.67	4.00	4.00	4.00
3	6.55×10^7	3.00	3.33	3.50	3.50	4.00
$k = 3$ with stabilization						
1	1.03×10^6	3.00	4.00	4.50	5.00	6.00
2	8.22×10^6	3.50	4.00	5.50	6.00	6.50
3	6.55×10^7	3.00	3.33	5.00	6.00	7.50

TABLE 5. Average number of outer Krylov iterations per Newton step for the 3D regularised lid-driven cavity problem.

512 cores (for $k = 3$, $\sim 60\,000$ dofs/core) and 960 cores (for $k = 5$, $\sim 31\,000$ dofs/core). Without stabilization, each Krylov iteration takes ~ 16 seconds. Adding stabilization increases this to ~ 30 seconds (due to the larger number of non-zeros in the matrix) and increasing the order to $k = 5$ increases the runtime further to ~ 450 seconds. Though we would not recommend the case $k = 5$ as a practical discretization due to its high computational cost and memory requirements, we include the results to demonstrate the necessity of capturing the nullspace of all singular operators in order to obtain a fully robust scheme.

6. Summary

The goal of this work was to develop a scalable solver for the stationary incompressible Navier–Stokes equations that exhibits both Reynolds-robust iteration counts and errors. To achieve this goal we apply the augmented Lagrangian approach to the exactly incompressible Scott–Vogelius discretization on

Refinements	Degrees of freedom	Reynolds number				
		10	100	1 000	2 500	5 000
$k = 3$ without stabilization						
1	3.85×10^6	4.50	4.33	5.33	9.00	15.00
2	3.06×10^7	5.00	5.33	5.33	10.00	12.00
$k = 3$ with stabilization						
1	3.85×10^6	4.50	5.33	7.33	11.50	13.50
2	3.06×10^7	5.00	6.33	12.50	14.00	154.00
$k = 5$ with stabilization						
1	3.81×10^6	1.50	1.67	2.00	3.50	4.00
2	3.03×10^7	2.00	1.67	2.00	2.50	5.00

TABLE 6. Average number of outer Krylov iterations per Newton step for the 3D backwards-facing step problem. The results for $k = 5$ were obtained on a coarser mesh to have a comparable number of degrees of freedom to the case of $k = 3$.

barycentric grids, and solve the augmented momentum block with a specialized multigrid method that exploits the barycentric structure in both relaxation and prolongation.

Dictated by inf-sup stability, the minimal polynomial degree that we require is $k = d$. For this degree we observe robust iteration counts to $\text{Re} \sim 1000$ but increase for higher Reynolds numbers. We attribute this behaviour to the singular nature of the stabilization term employed and show that when using a sufficiently high order discretization, full robustness can be achieved.

Code availability

The code for the Navier–Stokes solver and the numerical experiments in this paper can be found at <https://github.com/florianwechsung/alfi/>. For reproducibility, all major Firedrake components as well as the code used to run these examples have been archived on Zenodo. The results in Tables 5, and 6 were obtained on ARCHER, the UK national supercomputer using the code archived at [26, 61]. The results in Figure 8 and Tables 1, 2, 3, and 4 were obtained using the code archived at [64, 23].

References

- [1] N. AHMED, G. R. BARRENECHEA, E. BURMAN, J. GUZMÁN, A. LINKE, AND C. MERDON, *A pressure-robust discretization of Oseen’s equation using stabilization in the vorticity equation*, arXiv:2007.04012 [cs, math], (2020). arXiv: 2007.04012.
- [2] N. AHMED, A. LINKE, AND C. MERDON, *Towards pressure-robust mixed methods for the incompressible Navier–Stokes equations*, Computational Methods in Applied Mathematics, 18 (2018), pp. 353–372, <https://doi.org/10.1515/cmam-2017-0047>.
- [3] P. ALFELD, *A trivariate Clough–Tocher scheme for tetrahedral data*, Computer Aided Geometric Design, 1 (1984), pp. 169–181, [https://doi.org/10.1016/0167-8396\(84\)90029-3](https://doi.org/10.1016/0167-8396(84)90029-3).

REYNOLDS-ROBUST PRECONDITIONERS FOR THE NSE

- [4] D. N. ARNOLD, R. FALK, AND R. WINTHER, *Preconditioning in $H(\text{div})$ and applications*, Mathematics of Computation, 66 (1997), pp. 957–984, <https://doi.org/10.1090/S0025-5718-97-00826-0>.
- [5] D. N. ARNOLD, R. S. FALK, AND R. WINTHER, *Multigrid in $H(\text{div})$ and $H(\text{curl})$* , Numerische Mathematik, 85 (2000), pp. 197–217, <https://doi.org/10.1007/pl00005386>.
- [6] S. BALAY, S. ABHYANKAR, M. F. ADAMS, J. BROWN, P. BRUNE, K. BUSCHELMAN, L. DALCIN, V. ELJKHOUT, W. D. GROPP, D. KARPEYEV, D. KAUSHIK, M. G. KNEPLEY, D. A. MAY, L. C. MCINNES, R. T. MILLS, T. MUNSON, K. RUPP, P. SANAN, B. F. SMITH, S. ZAMPINI, H. ZHANG, AND H. ZHANG, *PETSc users manual*, Tech. Report ANL-95/11 - Revision 3.11, Argonne National Laboratory, 2019.
- [7] M. BENZI AND M. A. OLSHANSKII, *An augmented Lagrangian-based approach to the Oseen problem*, SIAM Journal on Scientific Computing, 28 (2006), pp. 2095–2113, <https://doi.org/10.1137/050646421>.
- [8] M. BENZI, M. A. OLSHANSKII, AND Z. WANG, *Modified augmented Lagrangian preconditioners for the incompressible Navier-Stokes equations*, International Journal for Numerical Methods in Fluids, 66 (2011), pp. 486–508, <https://doi.org/10.1002/fld.2267>.
- [9] A. N. BROOKS AND T. J. R. HUGHES, *Streamline upwind/Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations*, Computer Methods in Applied Mechanics and Engineering, 32 (1982), pp. 199–259, [https://doi.org/10.1016/0045-7825\(82\)90071-8](https://doi.org/10.1016/0045-7825(82)90071-8).
- [10] E. BURMAN, M. A. FERNÁNDEZ, AND P. HANSBO, *Continuous interior penalty finite element method for Oseen’s equations*, SIAM Journal on Numerical Analysis, 44 (2006), pp. 1248–1274, <https://doi.org/10.1137/040617686>.
- [11] E. BURMAN AND P. HANSBO, *Edge stabilization for Galerkin approximations of convection–diffusion–reaction problems*, Computer Methods in Applied Mechanics and Engineering, 193 (2004), pp. 1437–1453, <https://doi.org/10.1016/j.cma.2003.12.032>.
- [12] E. BURMAN AND P. HANSBO, *Edge stabilization for the generalized Stokes problem: A continuous interior penalty method*, Computer Methods in Applied Mechanics and Engineering, 195 (2006), pp. 2393–2410, <https://doi.org/10.1016/j.cma.2005.05.009>.
- [13] E. BURMAN AND A. LINKE, *Stabilized finite element schemes for incompressible flow using Scott–Vogelius elements*, Applied Numerical Mathematics, 58 (2008), pp. 1704–1719, <https://doi.org/10.1016/j.apnum.2007.11.001>.
- [14] M. A. CASE, V. J. ERVIN, A. LINKE, AND L. G. REBHOLZ, *A connection between Scott–Vogelius and grad-div stabilized Taylor–Hood FE approximations of the Navier–Stokes equations*, SIAM Journal on Numerical Analysis, 49 (2011), pp. 1461–1481, <https://doi.org/10.1137/100794250>.
- [15] P. G. CIARLET, *The Finite Element Method for Elliptic Problems*, North-Holland, 1978. Reprinted by SIAM in 2002.

- [16] B. COCKBURN, G. KANSCHAT, AND D. SCHÖTZAU, *A note on discontinuous Galerkin divergence-free solutions of the Navier–Stokes equations*, Journal of Scientific Computing, 31 (2007), pp. 61–73, <https://doi.org/10.1007/s10915-006-9107-7>.
- [17] C. CONCA, F. MURAT, AND O. PIRONNEAU, *The Stokes and Navier-Stokes equations with boundary conditions involving the pressure*, Japanese Journal of Mathematics. New series, 20 (1994), pp. 279–318.
- [18] J. DOUGLAS AND T. DUPONT, *Interior penalty procedures for elliptic and parabolic Galerkin methods*, in Computing Methods in Applied Sciences, vol. 58, Springer Berlin Heidelberg, 1976, pp. 207–216, <https://doi.org/10.1007/BFb0120591>.
- [19] H. ELMAN, V. E. HOWLE, J. SHADID, R. SHUTTLEWORTH, AND R. TUMINARO, *Block preconditioners based on approximate commutators*, SIAM Journal on Scientific Computing, 27 (2006), pp. 1651–1668, <https://doi.org/10.1137/040608817>.
- [20] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers: with applications in incompressible fluid dynamics*, Oxford University Press, 2014.
- [21] R. E. EWING AND J. WANG, *Analysis of the Schwarz algorithm for mixed finite elements methods*, ESAIM: Mathematical Modelling and Numerical Analysis, 26 (1992), pp. 739–756, <https://doi.org/10.1051/m2an/1992260607391>.
- [22] P. E. FARRELL, M. G. KNEPLEY, L. MITCHELL, AND F. WECHSUNG, *PCPATCH: Software for the topological construction of multigrid relaxation methods.*, 2019, <https://arxiv.org/abs/1912.08516>.
- [23] P. E. FARRELL, L. MITCHELL, L. R. SCOTT, AND F. WECHSUNG, *florianwechsung/alfi: Reynolds-Robust Preconditioners for the NSE*, 2020, <https://doi.org/10.5281/zenodo.4311759>, <https://github.com/florianwechsung/alfi/tree/fw/smaipaperresponse>.
- [24] P. E. FARRELL, L. MITCHELL, L. R. SCOTT, AND F. WECHSUNG, *Robust multigrid methods for nearly incompressible elasticity using macro elements*, 2020, <https://arxiv.org/abs/2020.02051>.
- [25] P. E. FARRELL, L. MITCHELL, AND F. WECHSUNG, *An augmented Lagrangian preconditioner for the 3D stationary incompressible Navier–Stokes equations at high Reynolds number*, SIAM Journal on Scientific Computing, 41 (2019), pp. A3073–A3096, <https://doi.org/10.1137/18M1219370>, <https://arxiv.org/abs/1810.03315>.
- [26] FIREDRAKE-ZENODO, *Software used in ‘Shape optimisation examples for “Shape Optimisation and Robust Solvers for Incompressible Flow”’*, 2019, <https://doi.org/10.5281/zenodo.3369183>, <https://zenodo.org/record/3369183>.
- [27] G. FU, J. GUZMÁN, AND M. NEILAN, *Exact smooth piecewise polynomial sequences on Alfeld splits*, Mathematics of Computation, 89 (2020), pp. 1059–1091, <https://doi.org/10.1090/mcom/3520>.
- [28] N. R. GAUGER, A. LINKE, AND 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), pp. 89–129, <https://doi.org/10.5802/smai-jcm.44>.

- [29] T. GELHARD, G. LUBE, M. A. OLSHANSKII, AND J.-H. STARCKE, *Stabilized finite element schemes with LBB-stable elements for incompressible flows*, Journal of Computational and Applied Mathematics, 177 (2005), pp. 243–267, <https://doi.org/10.1016/j.cam.2004.09.017>.
- [30] C. GEUZAIN AND J.-F. REMACLE, *Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities*, International Journal for Numerical Methods in Engineering, 79 (2009), pp. 1309–1331, <https://doi.org/10.1002/nme.2579>.
- [31] T. HEISTER AND G. RAPIN, *Efficient augmented Lagrangian-type preconditioning for the Oseen problem using grad-div stabilization*, International Journal for Numerical Methods in Fluids, 71 (2012), pp. 118–134, <https://doi.org/10.1002/flid.3654>.
- [32] Q. HONG, J. KRAUS, J. XU, AND L. ZIKATANOV, *A robust multigrid method for discontinuous Galerkin discretizations of Stokes and linear elasticity equations*, Numerische Mathematik, 132 (2015), pp. 23–49, <https://doi.org/10.1007/s00211-015-0712-y>.
- [33] V. JOHN, A. LINKE, C. MERDON, M. NEILAN, AND L. G. REBHOLZ, *On the divergence constraint in mixed finite element methods for incompressible flows*, SIAM Review, 59 (2017), pp. 492–544, <https://doi.org/10.1137/15m1047696>.
- [34] D. KAY, D. LOGHIN, AND A. WATHEN, *A preconditioner for the steady-state Navier–Stokes equations*, SIAM Journal on Scientific Computing, 24 (2002), pp. 237–256, <https://doi.org/10.1137/S106482759935808X>.
- [35] A. KOLESNIKOV AND T. SOROKINA, *Multivariate C^1 -Continuous Splines on the Alfeld Split of a Simplex*, in Approximation Theory XIV: San Antonio 2013, G. E. Fasshauer and L. L. Schumaker, eds., vol. 83, Springer International Publishing, Cham, 2014, pp. 283–294, https://doi.org/10.1007/978-3-319-06404-8_16. Series Title: Springer Proceedings in Mathematics & Statistics.
- [36] L. D. LANDAU AND E. M. LIFSHITZ, *Fluid mechanics*, no. 6 in Course of Theoretical Physics, Elsevier/Butterworth Heinemann, Amsterdam, 2nd ed., 1987.
- [37] Y.-J. LEE, J. WU, AND J. CHEN, *Robust multigrid method for the planar linear elasticity problems*, Numerische Mathematik, 113 (2009), pp. 473–496, <https://doi.org/10.1007/s00211-009-0232-8>.
- [38] Y.-J. LEE, J. WU, J. XU, AND L. ZIKATANOV, *Robust subspace correction methods for nearly singular systems*, Mathematical Models and Methods in Applied Sciences, 17 (2007), pp. 1937–1963, <https://doi.org/10.1142/s0218202507002522>.
- [39] C. LEHRENFELD AND J. SCHÖBERL, *High order exactly divergence-free hybrid discontinuous Galerkin methods for unsteady incompressible flows*, Computer Methods in Applied Mechanics and Engineering, 307 (2016), pp. 339–361, <https://doi.org/10.1016/j.cma.2016.04.025>.
- [40] X. S. LI AND J. W. DEMMEL, *SuperLU_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems*, ACM Transactions on Mathematical Software, 29 (2003), pp. 110–140, <https://doi.org/10.1145/779359.779361>.
- [41] X. S. LI, J. W. DEMMEL, J. R. GILBERT, I. L. GRIGORI, M. SHAO, AND I. YAMAZAKI, *SuperLU Users’ Guide*, Tech. Report LBNL-44289, Lawrence Berkeley National Laboratory, September 1999.

- [42] A. LINKE AND L. G. REBHOLZ, *Pressure-induced locking in mixed methods for time-dependent (Navier–)Stokes equations*, Journal of Computational Physics, 388 (2019), pp. 350–356, <https://doi.org/10.1016/j.jcp.2019.03.010>.
- [43] D. A. MAY, P. SANAN, K. RUPP, M. G. KNEPLEY, AND B. F. SMITH, *Extreme-scale multigrid components within PETSc*, in Proceedings of the Platform for Advanced Scientific Computing Conference, 2016, <https://doi.org/10.1145/2929908.2929913>.
- [44] M. NEILAN, *The Stokes complex: a review of exactly divergence-free finite element pairs for incompressible flows*, in 75 Years of Mathematics of Computation, S. C. Brenner, I. Shparlinksi, C. Shu, and D. B. Szyld, eds., vol. 754 of Contemporary Mathematics, American Mathematical Society, 2020, pp. 141–158, <https://doi.org/10.1090/conm/754/15142>.
- [45] M. A. OLSHANSKII AND L. G. REBHOLZ, *Application of barycenter refined meshes in linear elasticity and incompressible fluid dynamics*, Electronic Transactions on Numerical Analysis, 38 (2011), pp. 258–274.
- [46] S. PATANKAR, *Numerical Heat Transfer and Fluid Flow*, Hemisphere Series on Computational Methods in Mechanics and Thermal Science, Taylor & Francis, 1 ed., 1980.
- [47] M. J. D. POWELL AND M. A. SABIN, *Piecewise quadratic approximations on triangles*, ACM Transactions on Mathematical Software, 3 (1977), pp. 316–325, <https://doi.org/10.1145/355759.355761>.
- [48] J. QIN, *On the convergence of some low order mixed finite elements for incompressible fluids*, PhD thesis, Pennsylvania State University, 1994.
- [49] A. QUARTERONI AND A. VALLI, *Numerical Approximation of Partial Differential Equations*, vol. 23 of Springer Series in Computational Mathematics, Springer, 2008.
- [50] F. RATHGEBER, D. A. HAM, L. MITCHELL, M. LANGE, F. LUPORINI, A. T. T. MCRAE, G.-T. BERCEA, G. R. MARKALL, AND P. H. J. KELLY, *Firedrake: automating the finite element method by composing abstractions*, ACM Transactions on Mathematical Software, 43 (2016), pp. 1–27, <https://doi.org/10.1145/2998441>.
- [51] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM Journal on Scientific Computing, 14 (1993), pp. 461–469, <https://doi.org/10.1137/0914028>.
- [52] J. SCHÖBERL, *Multigrid methods for a parameter dependent problem in primal variables*, Numerische Mathematik, 84 (1999), pp. 97–119, <https://doi.org/10.1007/s002110050465>.
- [53] J. SCHÖBERL, *Robust Multigrid Methods for Parameter Dependent Problems*, PhD thesis, Johannes Kepler Universität Linz, Linz, Austria, 1999.
- [54] J. SCHÖBERL, *Robust Multigrid Preconditioning for Parameter-Dependent Problems I: The Stokes-Type Case*, in Multigrid Methods V, W. Hackbusch and G. Wittum, eds., Lecture Notes in Computational Science and Engineering, Springer Berlin Heidelberg, 1998, pp. 260–275, https://doi.org/10.1007/978-3-642-58734-4_15.
- [55] L. R. SCOTT AND M. VOGELIUS, *Conforming finite element methods for incompressible and nearly incompressible continua*, in Large Scale Computations in Fluid Mechanics, B. E. Engquist, et al., eds., vol. 22 (Part 2), Providence: AMS, 1985, pp. 221–244.

- [56] L. R. SCOTT AND M. VOGELIUS, *Norm estimates for a maximal right inverse of the divergence operator in spaces of piecewise polynomials*, ESAIM: Mathematical Modelling and Numerical Analysis, 19 (1985), pp. 111–143, <https://doi.org/10.1051/m2an/1985190101111>.
- [57] T. SHIH, C. TAN, AND B. HWANG, *Effects of grid staggering on numerical schemes*, International Journal for Numerical Methods in Fluids, 9 (1989), pp. 193–212, <https://doi.org/10.1002/flid.1650090206>.
- [58] S. TUREK, *Efficient Solvers for Incompressible Flow Problems: An Algorithmic and Computational Approach*, Lecture Notes in Computational Science and Engineering, Springer, 1999.
- [59] S. P. VANKA, *Block-implicit multigrid solution of Navier-Stokes equations in primitive variables*, Journal of Computational Physics, 65 (1986), pp. 138–158, [https://doi.org/10.1016/0021-9991\(86\)90008-2](https://doi.org/10.1016/0021-9991(86)90008-2).
- [60] J. WANG AND X. YE, *New finite element methods in computational fluid dynamics by $H(\text{div})$ elements*, SIAM Journal on Numerical Analysis, 45 (2007), pp. 1269–1286, <https://doi.org/10.1137/060649227>.
- [61] F. WECHSUNG, *florianwechsung/alfi: Thesis submission*, 2019, <https://doi.org/10.5281/zenodo.3369207>.
- [62] A. J. WORSEY AND G. FARIN, *An n -dimensional Clough–Tocher interpolant*, Constructive Approximation, 3 (1987), pp. 99–110, <https://doi.org/10.1007/BF01890556>.
- [63] J. XU, *Iterative methods by space decomposition and subspace correction*, SIAM Review, 34 (1992), pp. 581–613, <https://doi.org/10.1137/1034116>.
- [64] *Software used in 'A Reynolds-robust preconditioner for the Reynolds-robust Scott–Vogelius discretization of the stationary incompressible Navier–Stokes equations'*, dec 2020, <https://doi.org/10.5281/zenodo.4311750>, <https://doi.org/10.5281/zenodo.4311750>.
- [65] S. ZHANG, *A new family of stable mixed finite elements for the 3D Stokes equations*, Mathematics of Computation, 74 (2004), pp. 543–555, <https://doi.org/10.1090/s0025-5718-04-01711-9>.
- [66] S. ZHANG, *On the P_1 Powell–Sabin divergence free finite element for the Stokes equations*, Journal of Computational Mathematics, 26 (2008), pp. 456–470.
- [67] S. ZHANG, *Divergence-free finite elements on tetrahedral grids for $k \geq 6$* , Mathematics of Computation, 80 (2011), pp. 669–695, <https://doi.org/10.1090/S0025-5718-2010-02412-3>.
- [68] S. ZHANG, *Quadratic divergence-free finite elements on Powell–Sabin tetrahedral grids*, Calcolo, 48 (2011), pp. 211–244, <https://doi.org/10.1007/s10092-010-0035-4>.