



Fast solution of incompressible flow problems with two-level pressure approximation

Jennifer Pestana¹ · David J. Silvester²

Received: 17 March 2023 / Revised: 23 May 2024 / Accepted: 25 May 2024 /

Published online: 17 June 2024

© The Author(s) 2024

Abstract

This paper develops efficient preconditioned iterative solvers for incompressible flow problems discretised by an enriched Taylor–Hood mixed approximation, in which the usual pressure space is augmented by a piecewise constant pressure to ensure local mass conservation. This enrichment process causes over-specification of the pressure when the pressure space is defined by the union of standard Taylor–Hood basis functions and piecewise constant pressure basis functions, which complicates the design and implementation of efficient solvers for the resulting linear systems. We first describe the impact of this choice of pressure space specification on the matrices involved. Next, we show how to recover effective solvers for Stokes problems, with preconditioners based on the singular pressure mass matrix, and for Oseen systems arising from linearised Navier–Stokes equations, by using a two-stage pressure convection–diffusion strategy. The codes used to generate the numerical results are available online.

Mathematics Subject Classification 65F08 · 65F10 · 65N30 · 35Q30

1 Introduction

Reliable and efficient iterative solvers for models of steady incompressible flow emerged in the early 1990s. Strategies based on block preconditioning of the underlying matrix operators using (algebraic or geometric) multigrid components have proved to be the key to realising mesh independent convergence (and optimal complexity) without the need for tuning parameters, particularly in the context of classical mixed finite element approximation, see Elman et al. [8, chap. 9]. The focus of this contribu-

✉ Jennifer Pestana
jennifer.pestana@strath.ac.uk

David J. Silvester
d.silvester@manchester.ac.uk

¹ Department of Mathematics and Statistics, University of Strathclyde, Glasgow G1 1XH, UK

² Department of Mathematics, University of Manchester, Oxford Road, Manchester M13 9PL, UK

tion is on efficient solver strategies in cases where (an inf–sup) stable Taylor–Hood mixed approximation is augmented by a piecewise constant pressure in order to guarantee local conservation of mass. The augmentation leads to over-specification of the pressure solution when the pressure space is defined in the natural way via the union of Taylor–Hood pressure basis functions and basis functions for the piecewise constant pressure space, requiring a redesign of the established solver technology.

The idea of adding a piecewise constant pressure to the standard rectangular biquadratic velocity, bilinear pressure ($\mathbb{Q}_2\text{--}\mathbb{Q}_1$) approximation was originally suggested during discussion around a blackboard at a conference on finite elements in fluids held in Banff in 1980; see Gresho et al. [9]. The need for local mass conservation was motivated by competition from finite volume methods (such as the MAC scheme) in the design of effective strategies for modelling buoyancy-driven flow in the atmosphere.

The extension of the augmentation idea to Taylor–Hood ($\mathbb{P}_2\text{--}\mathbb{P}_1$) triangular approximation was proposed in a paper presented at a conference in Reading in 1982; see Griffiths [10]. A proof of stability of the augmented $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ approximation on triangular meshes was constructed by Thatcher and Silvester [24] in 1987. An extended version of this manuscript included discussion of $\mathbb{Q}_2\text{--}\mathbb{Q}_1^*$ hexahedral elements [23]. A rigorous assessment of the augmentation strategy was undertaken by Boffi et al. two decades later [4]. The strategy of augmenting a continuous pressure approximation to give a locally mass-conserving strategy can also be generalised to higher-order $\mathbb{Q}_k\text{--}\mathbb{Q}_{k-1}$ and $\mathbb{P}_k\text{--}\mathbb{P}_{k-1}$ Taylor–Hood approximations. Inf–sup stability is assured for $k \geq 2$ in two dimensions and for $\mathbb{P}_k\text{--}\mathbb{P}_{k-1}$ with $k \geq 3$ in three dimensions; see [3, p. 506].

A closely related, yet fundamentally different Stokes approximation strategy is developed in [27]. Their idea is to stabilise the lowest-order $\mathbb{P}_1\text{--}\mathbb{P}_0$ approximation space defined on triangles or tetrahedra by augmenting the continuous velocity approximation by a piecewise constant approximation. The extended mixed approximation is stable but nonconforming. A similar approach is taken in [6], but an alternative weak form is used. Note that the design of linear solvers for the resulting discrete equations is relatively straightforward. A special feature of the extended approximation is that it can be augmented by postprocessing to give a pressure-robust approximation; see [13].

More generally, such an extended Galerkin (EG) approximation can be viewed as an intermediate between continuous Galerkin (CG) and discontinuous Galerkin (DG). This interpretation of piecewise constant augmentation was originally put forward by Sun and Liu [22] and is motivated by the fact that it gives local flux conservation when modelling transport in porous media flow problems, but with fewer degrees of freedom compared to vanilla DG.

The novel contribution in this work lies in the linear algebra aspects of two-field pressure approximation. The immediate issue that needs to be dealt with is the fact that, when the pressure space is specified by the frame formed by adding usual finite element basis functions for the underlying Taylor–Hood space to basis functions for the discontinuous pressure space, the mass matrix that determines the *stability* of the resulting mixed approximation is *singular*. This aspect is noted in the context of EG approximation in [15, Remark 4.1] and is discussed in Sect. 2. The main issue

that has to be dealt with in practical flow simulation is the over-specification and associated ill-conditioning of the discrete operators that arise in the preconditioning of the linearised Navier–Stokes operator. This is the focus of the discussion in Sect. 3. The conclusion of this study is that *optimal* complexity convergence rates can be recovered when using two-field pressure approximation, but only after a careful redesign of the preconditioning components.

2 Two-field pressure mass matrix

In this section we consider the Stokes problem

$$\begin{aligned} -\Delta \vec{u} + \nabla p &= \vec{0} & \text{in } \Omega, \\ \nabla \cdot \vec{u} &= 0 & \text{in } \Omega, \\ \vec{u} &= \vec{g} & \text{on } \partial\Omega, \end{aligned}$$

where \vec{u} and p are the fluid velocity and pressure, respectively, and $\Omega \subset \mathbb{R}^d$, $d \in \{2, 3\}$, is a polygonal or polyhedral domain.

Throughout this section, we assume that $V_h \subset H_0^1(\Omega)^d$ and $Q_h \subset L_0^2(\Omega) := \{q \in L_2 : \int_{\Omega} q = 0\}$ are an inf–sup stable pair of finite element spaces. Then the corresponding finite element approximation problem is to find $(\vec{u}_h, p_h) \in V_h \times Q_h$ such that

$$a(\vec{u}_h, \vec{v}_h) + b(\vec{v}_h, p_h) = (\vec{f}_h, \vec{v}_h) \quad \text{for all } \vec{v}_h \in V_h, \quad (1a)$$

$$b(\vec{u}_h, q_h) = (g_h, q_h) \quad \text{for all } q_h \in Q_h, \quad (1b)$$

where (\cdot, \cdot) denotes the usual L^2 inner product, right-hand side functions \vec{f}_h and g_h are associated with the specified boundary velocity (g_h is zero for enclosed flow), and

$$a(\vec{u}, \vec{v}) = \int_{\Omega} \nabla \vec{u} : \nabla \vec{v}, \quad b(\vec{u}, p) = - \int_{\Omega} p \nabla \cdot \vec{u}.$$

Solving the finite element problem (1) is then equivalent to solving the linear system

$$\underbrace{\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix}}_{\mathcal{A}} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix}, \quad (2)$$

where $A \in \mathbb{R}^{n_u \times n_u}$ is symmetric positive definite and $B \in \mathbb{R}^{n_p \times n_u}$ [8, chap. 3]. The matrix \mathcal{A} is of well-known saddle point type, and solvers for this sort of system have been extensively studied, see, e.g., Benzi et al. [1]. Since \mathcal{A} is large and sparse, the system is typically solved by an iterative method, with preconditioned MINRES [18] a popular choice.

An ideal preconditioner for \mathcal{A} is [14, 17]

$$\mathcal{P}_{\text{ideal}} = \begin{bmatrix} \mathbf{A} & \\ & \mathbf{B}\mathbf{A}^{-1}\mathbf{B}^T \end{bmatrix},$$

since the eigenvalues of the preconditioned matrix are 0, 1, and $(1 \pm \sqrt{5})/2$, with the zero eigenvalue appearing only if the preconditioned matrix is singular. This preconditioner is usually too costly to apply, and so efficient block diagonal preconditioners for \mathcal{A} are typically based on spectrally equivalent preconditioners for \mathbf{A} and the negative Schur complement $\mathbf{S} = \mathbf{B}\mathbf{A}^{-1}\mathbf{B}^T$. For Stokes problems, the solve with \mathbf{A} can be replaced by, e.g., an algebraic or geometric multigrid solver, while for stable discretisations \mathbf{S} is spectrally equivalent to M_Q , the pressure mass matrix [8, chap. 3], i.e., there exist constants γ and Γ , independent of h , such that

$$\gamma^2 \leq \frac{\mathbf{q}^T \mathbf{B}\mathbf{A}^{-1}\mathbf{B}^T \mathbf{q}}{\mathbf{q}^T M_Q \mathbf{q}} \leq \Gamma^2 \quad (3)$$

for all vectors \mathbf{q} except those corresponding to the function $q_h \equiv 1$ on Ω . For certain element pairs, M_Q itself is easily inverted, e.g., for discontinuous \mathbb{P}_0 or \mathbb{Q}_0 pressures, the mass matrix is diagonal. Otherwise, M_Q can be replaced by its diagonal or by a fixed number of steps of Chebyshev semi-iteration acceleration of a Jacobi iteration when, as is common, $\text{diag}(M_Q)^{-1}M_Q$ has eigenvalues within a small interval and this interval lies away from the origin [25, 26].

To summarise, for inf-sup stable finite element pairs, an effective preconditioner for (2) is

$$\mathcal{P} = \begin{bmatrix} \mathbf{M} & \\ & M_S \end{bmatrix}, \quad (4)$$

where $\mathbf{M} \in \mathbb{R}^{n_u \times n_u}$ is \mathbf{A} or an approximation, and $M_S \in \mathbb{R}^{n_p \times n_p}$ is M_Q or an approximation. (See, e.g., Elman et al. [8, chap. 4] for results with \mathbb{Q}_2 – \mathbb{Q}_1 approximation on quadrilaterals.)

In this section, our aim is to determine effective preconditioners for the enriched Taylor–Hood element. As we will see, although enriching the pressure space results in better mass conservation properties, it can also present challenges for solving (2).

2.1 Augmented Taylor–Hood elements

We see from (1) that the mass conservation condition $\nabla \cdot \vec{u} = 0$ is imposed only in a weak sense, and if Taylor–Hood elements are employed we can only guarantee that (1b) will hold. However, by augmenting the pressure space by piecewise constant pressures it is possible to obtain local mass conservation, so that the average of the divergence is zero on each individual element. The natural choice is to define this space by a frame consisting of the union of the continuous Taylor–Hood pressure basis functions and the discontinuous pressure basis functions. In this case, as we

will discuss later in this section, constant pressures have multiple representations, and this results in certain challenges for solving (2). Although it is certainly possible to instead compute a basis for this augmented pressure space, we show that the linear algebra challenges associated with using a frame can be overcome, which simplifies the implementation of the method.

Let us now describe the enriched Taylor–Hood finite element spaces. We first introduce a shape-regular family of simplicial, quadrilateral (in 2D) or hexahedral (in 3D) decompositions of the domain Ω . We assume that any two elements have at most a common face, edge, or vertex and denote by h the maximum diameter of any element. The total number of elements in the resulting mesh is n_{el} .

We denote the usual Taylor–Hood finite element space by $V_h^{TH} \times Q_h^{TH}$, so that $V_h^{TH} \times Q_h^{TH} = (\mathbb{Q}_{k+1})^d \times \mathbb{Q}_k$ or $V_h^{TH} \times Q_h^{TH} = (\mathbb{P}_{k+1})^d \times \mathbb{P}_k$ with $d \geq 2$. In the latter case, we additionally assume that the polynomial degree, k , satisfies $k \geq d - 1$. The corresponding enriched Taylor–Hood space is $V_h^{TH} \times Q_h^*$ where

$$Q_h^* = \left\{ q = q_k + q_0, q_k \in Q_h^{TH}, q_0 \in Q_h^0 \right\} \tag{5}$$

and Q_h^0 is the space of discontinuous pressures that are constant on each element. Thus, we see that the velocity approximation space is identical to that of the corresponding Taylor–Hood element, while Q_h^* is Q_h^{TH} augmented with piecewise constant pressures. It follows that functions in Q_h^* may be discontinuous across inter-element boundaries.

We stress that the enriched Taylor–Hood space is well defined, and inf–sup stable. From a linear algebra perspective, however, a critical point is the representation of Q_h^* . Ideally, we would like to represent functions in Q_h^* as linear combinations of basis functions. In this case the resulting linear system would be nonsingular, and it is likely that the approaches described at the start of this section could be applied directly. However, the most natural choice is define Q_h^* by a frame, i.e., to let

$$Q_h^* = \text{span}\{\phi_1, \dots, \phi_{n_k}, \phi_{n_k+1}, \dots, \phi_{n_p}\}, \tag{6}$$

where $\{\phi_k\}_{k=1}^{n_k}$ and $\{\phi_k\}_{k=n_k+1}^{n_p}$ are Lagrange bases of Q_h^{TH} and Q_h^0 (see (5)). This approach to specifying the pressure space has been used in, e.g., [4, 20, 23].

Although this makes specification of the pressure straightforward, with this choice any constant function on Ω can be represented by a function in Q_h^{TH} or a function in Q_h^0 . This has profound consequences for the linear algebra: the pressure mass matrix M_Q that determines M_S in (4) becomes singular, and the rank of the matrix B^T is reduced. In the rest of this section we first establish these properties, before showing that it is still possible to solve (2) by preconditioned MINRES in this case.

Using (6), we can relate vectors $\mathbf{p} \in \mathbb{R}^{n_p}$ and functions $p = p_k + p_0 \in Q_h^*$, where $p_k \in Q_h^{TH}$, $p_0 \in Q_h^0$. Specifically,

$$p = \sum_{i=1}^{n_p} \mathbf{p}_i \phi_i, \quad p_k = \sum_{i=1}^{n_k} \mathbf{p}_i \phi_i \quad \text{and} \quad p_0 = \sum_{i=n_k+1}^{n_p} \mathbf{p}_i \phi_i. \tag{7}$$

As mentioned above, any constant function has multiple representations when specified via (6). In particular, $p = p_k + p_0 \equiv 0$ with $p_k = \alpha$ and $p_0 = -\alpha$, for any $\alpha \in \mathbb{R}$. We see from (7) that this representation of the zero function corresponds to the vector $\alpha \mathbf{k}$, where

$$\mathbf{k} = \begin{bmatrix} \mathbf{1}_{n_k} \\ -\mathbf{1}_{n_0} \end{bmatrix} \tag{8}$$

and $n_0 = n_p - n_k$. A direct consequence of the correspondence between \mathbf{k} and the zero function on Ω is that $M_Q \mathbf{k} = \mathbf{0}$ and $B^T \mathbf{k} = \mathbf{0}$, as we now show.

Proposition 1 *Let*

$$M_Q = [q_{ij}]_{i,j=1}^{n_p}, \quad q_{ij} = \int_{\Omega} \phi_i \phi_j,$$

be the pressure mass matrix for the enriched Taylor–Hood pressure space Q_h^ in (5), specified by the frame (6). Then, $\text{null}(M_Q) = \text{span}\{\mathbf{k}\}$, where \mathbf{k} is given in (8).*

Proof Let $\mathbf{p} \in \mathbb{R}^{n_p}$, $\mathbf{p} \neq \mathbf{0}$. Then, using (7), we find that

$$\begin{aligned} M_Q \mathbf{p} = 0 &\Leftrightarrow \mathbf{p}^T M_Q \mathbf{p} = 0 \Leftrightarrow \int_{\Omega} \left(\sum_{i=1}^{n_p} p_i \phi_i \right) \left(\sum_{j=1}^{n_p} p_j \phi_j \right) = 0 \\ &\Leftrightarrow \int_{\Omega} p^2 = 0, \end{aligned}$$

which implies that $p \equiv 0$ in Ω since p is continuous on each element. Since $p \equiv 0$ corresponds to vectors of the form $\alpha \mathbf{k}$, we find that $\text{span}\{\mathbf{k}\} \subseteq \text{null}(M_Q)$.

We now show that there are no other vectors in the nullspace. Since $p = p_k + p_0 \equiv 0$, $p_k \in Q_h^{TH}$, $p_0 \in Q_h^0$, we must have $p_k = -p_0$ everywhere in Ω . But p_0 is piecewise constant, and p_k is continuous on $\overline{\Omega}$. It follows that $p_0 \equiv \alpha$ and $p_k \equiv -\alpha$ on Ω for some constant $\alpha \in \mathbb{R}$. Such functions correspond to vectors of the form $\alpha \mathbf{k}$, which shows that $\text{null}(M_Q) = \text{span}\{\mathbf{k}\}$, as required. \square

Remark 2 A very similar argument shows that $\mathbf{k} \in \text{null}(B^T)$.

What do these results mean for the solution of (2) by preconditioned MINRES, when the coefficient matrix is constructed using the frame in (6)? First, if $B^T \mathbf{k} = \mathbf{0}$ then it follows that \mathcal{A} is always singular, with $\mathcal{A} \mathbf{w} = \mathbf{0}$, where

$$\mathbf{w} = \begin{bmatrix} \mathbf{0}_{n_u} \\ \mathbf{k} \end{bmatrix}. \tag{9}$$

If the linear system (2) is consistent then this does not pose a problem for preconditioned MINRES. However, as we shall now see, the proposed block diagonal preconditioner (4), with $M_S = M_Q$, is also singular, and so effective preconditioning requires some care.

Algorithm 3 Preconditioned MINRES algorithm for solving $\mathcal{A}\mathbf{x} = \mathbf{b}$ with symmetric positive definite preconditioner \mathcal{P} [8, Algorithm 4.1].

- 1: $\mathbf{v}^{(0)} = \mathbf{0}$, $\mathbf{w}^{(0)} = \mathbf{0}$, $\mathbf{w}^{(1)} = \mathbf{0}$, $\gamma_0 = 0$
- 2: Choose $\mathbf{x}^{(0)}$, compute $\mathbf{v}^{(1)} = \mathbf{b} - \mathcal{A}\mathbf{x}^{(0)}$
- 3: Solve $\mathcal{P}\mathbf{z}^{(1)} = \mathbf{v}^{(1)}$, set $\gamma_1 = \sqrt{\langle \mathbf{z}^{(1)}, \mathbf{v}^{(1)} \rangle}$
- 4: Set $\eta = \gamma_1$, $s_0 = s_1 = 0$, $c_0 = c_1 = 1$
- 5: **for** $j = 1$ **until** convergence **do**
- 6: $\mathbf{z}^{(j)} = \mathbf{z}^{(j)}/\gamma_j$
- 7: $\delta_j = \langle \mathcal{A}\mathbf{z}^{(j)}, \mathbf{z}^{(j)} \rangle$
- 8: $\mathbf{v}^{(j+1)} = \mathcal{A}\mathbf{z}^{(j)} - (\delta_j/\gamma_j)\mathbf{v}^{(j)} - (\gamma_j/\gamma_{j-1})\mathbf{v}^{(j-1)}$
- 9: Solve $\mathcal{P}\mathbf{z}^{(j+1)} = \mathbf{v}^{(j+1)}$
- 10: $\gamma_{j+1} = \sqrt{\langle \mathbf{z}^{(j+1)}, \mathbf{v}^{(j+1)} \rangle}$
- 11: $\alpha_0 = c_j\delta_j - c_{j-1}s_j\gamma_j$
- 12: $\alpha_1 = \sqrt{\alpha_0^2 + \gamma_{j+1}^2}$
- 13: $\alpha_2 = s_j\delta_j + c_{j-1}c_j\gamma_j$
- 14: $\alpha_3 = s_{j-1}\gamma_j$
- 15: $c_{j+1} = \alpha_0/\alpha_1$; $s_{j+1} = \gamma_{j+1}/\alpha_1$
- 16: $\mathbf{w}^{(j+1)} = (\mathbf{z}^{(j)} - \alpha_3\mathbf{w}^{(j-1)} - \alpha_2\mathbf{w}^{(j)})/\alpha_1$
- 17: $\mathbf{x}^{(j)} = \mathbf{x}^{(j-1)} + c_{j+1}\eta\mathbf{w}^{(j+1)}$
- 18: $\eta = -s_{j+1}\eta$
- 19: <Test for convergence>
- 20: **end for**

2.2 Preconditioning considerations

Knowing that the enriched Taylor–Hood element is inf–sup stable implies that the preconditioner (4) will be effective when solving (2). However, with the common choice of frame (6), the matrix M_Q , and hence the preconditioner \mathcal{P} , are singular, since $\mathcal{P}\mathbf{w} = \mathbf{0}$, where \mathbf{w} is given in (9). (Note that this implies that \mathcal{A} and \mathcal{P} have a common nullspace.) Here, we will show that using a singular preconditioner causes no difficulty for the preconditioned MINRES method in Algorithm 3 in exact arithmetic. At the end of Sect. 2, we present numerical results with two different preconditioners to demonstrate our approach.

We start by noting that at each iteration step we need to solve a linear system of the form $\mathcal{P}\mathbf{z}^{(j)} = \mathbf{v}^{(j)}$. If this system is consistent there are infinitely many solutions, which take the form $\mathbf{z}^{(j)} = \bar{\mathbf{z}}^{(j)} + \zeta_j\mathbf{w}$, with $\bar{\mathbf{z}}^{(j)} \perp \mathbf{w}$ and $\zeta_j \in \mathbb{R}$. Our next step is to examine the effect of $|\zeta_j|$ on the scalars and vectors in Algorithm 3. Observe that if the linear system $\mathcal{A}\mathbf{x} = \mathbf{b}$ is consistent then $\mathbf{b} \perp \mathbf{w}$ and so $\mathbf{v}^{(1)} \perp \mathbf{w}$. It then follows by induction that $\mathbf{v}^{(j)} \perp \mathbf{w}$, $j = 1, 2, \dots$, and that the systems $\mathcal{P}\mathbf{z}^{(j)} = \mathbf{v}^{(j)}$ are all consistent. Furthermore, $\gamma_j = \langle \mathbf{z}^{(j)}, \mathbf{v}^{(j)} \rangle^{\frac{1}{2}} = \langle \bar{\mathbf{z}}^{(j)}, \mathbf{v}^{(j)} \rangle^{\frac{1}{2}}$ because of the orthogonality of $\mathbf{v}^{(j)}$ and \mathbf{w} , so that ζ_j does not affect γ_j . Similarly, $\delta_j = \langle \mathcal{A}\mathbf{z}^{(j)}, \mathbf{z}^{(j)} \rangle = \langle \mathcal{A}\bar{\mathbf{z}}^{(j)}, \bar{\mathbf{z}}^{(j)} \rangle$, which shows that δ_j is similarly unaffected by ζ_j . Indeed, the only quantities that are affected by the nullspace components $\zeta_j\mathbf{w}$, $j = 1, 2, \dots$, are the vectors $\mathbf{w}^{(j)}$ and $\mathbf{x}^{(j)}$. In exact arithmetic, this is not a problem, since solutions of $\mathcal{A}\mathbf{x} = \mathbf{b}$ may certainly

contain a component in the direction of \mathbf{w} . However, in unlucky cases, the size of the nullspace component of $\mathbf{x}^{(j)}$ may be so large as to dominate the approximate solution. Alternatively, in finite precision $\mathbf{v}^{(j)}$ and \mathbf{w} may not be exactly orthogonal. Hence, it may be wise to explicitly ensure that $\mathbf{z} \perp \mathbf{w}$. One option is to orthogonalise \mathbf{z} against \mathbf{w} after each preconditioner solve, but if $|\zeta_j|$ is large then the result may be inaccurate. A more robust approach is to note that solutions of $\mathcal{P}\mathbf{z} = \mathbf{r}$ are minimisers of the quadratic form

$$\frac{1}{2}\mathbf{z}^T \mathcal{P}\mathbf{z} - \mathbf{z}^T \mathbf{r},$$

since \mathcal{P} is positive semidefinite. Constraining \mathbf{z} to be orthogonal to \mathbf{w} is then equivalent to the following optimisation problem:

$$\min_{\mathbf{z}} \frac{1}{2}\mathbf{z}^T \mathcal{P}\mathbf{z} - \mathbf{z}^T \mathbf{r} \quad \text{s.t.} \quad \mathbf{w}^T \mathbf{z} = 0.$$

Applying a Lagrange multiplier approach results in the augmented system

$$\begin{bmatrix} \mathcal{P} & \mathbf{w} \\ \mathbf{w}^T & 0 \end{bmatrix} \begin{bmatrix} \mathbf{z} \\ \lambda \end{bmatrix} = \begin{bmatrix} \mathbf{r} \\ 0 \end{bmatrix},$$

where λ is the Lagrange multiplier, and solving this augmented system gives a vector \mathbf{z} that is orthogonal to \mathbf{w} . Moreover, since

$$\begin{bmatrix} \mathcal{P} & \mathbf{w} \\ \mathbf{w}^T & 0 \end{bmatrix} = \begin{bmatrix} \mathbf{M} & 0 & 0 \\ 0 & M_S & \mathbf{k} \\ 0 & \mathbf{k}^T & 0 \end{bmatrix}$$

we see that only the solve with M_S needs to be modified.

2.3 Approximating the two-level pressure mass matrix

Now let us consider approximations of the matrix M_Q which, because of the two-level pressure approximation, has 2×2 block structure:

$$M_Q = \begin{bmatrix} Q_k & R^T \\ R & Q_0 \end{bmatrix},$$

where

$$\begin{aligned} Q_k &= [q_{k,ij}], \quad i, j = 1, \dots, n_k, & q_{k,ij} &= \int_{\Omega} \phi_j \phi_i, \\ R &= [r_{ij}], \quad i = n_k + 1, \dots, n_p, \quad j = 1, \dots, n_k, & r_{ij} &= \int_{\Omega} \phi_j \phi_i, \\ Q_0 &= [q_{0,ij}], \quad i, j = n_k + 1, \dots, n_p, & q_{0,ij} &= \int_{\Omega} \phi_j \phi_i. \end{aligned}$$

Note that Q_k is the standard Taylor–Hood pressure mass matrix, and Q_0 is the standard discontinuous pressure mass matrix, while R represents cross terms between the spaces Q_h^{TH} and Q_h^0 (see (5)).

Naively, one may wish to approximate Q_k by its diagonal, but this results in slow convergence rates; it can be shown, using the method described by Wathen [26], that $\text{diag}(Q_k)^{-1}Q_k$ has eigenvalues in an interval $[0, \mu]$ where, for example, for \mathbb{P}_1 elements on triangles $\mu = 3$. Accordingly, applying a fixed number of iterations of a Chebyshev semi-iteration based on a Jacobi iteration, as is standard for nonsingular Galerkin mass matrices [25], leads to large iteration counts here. Replacing Q_k or Q_0 in M_Q by their diagonals results in similarly poor performance.

However, it is possible to replace M_Q by an approximation designed for symmetric positive semidefinite matrices, see, e.g., [5, 7, 16] and the references therein. For example, symmetric Gauss–Seidel is convergent for such matrices [7]. Here, we have applied a fixed number of iterations of Chebyshev semi-iteration based on symmetric Gauss–Seidel [5]. The results are still mesh-dependent, because the largest non-unit eigenvalue of the underlying symmetric Gauss–Seidel iteration matrix approaches 1 as the mesh is refined but, for large problems, the cost per iteration is much lower than applying M_Q exactly.

2.4 Reliable computation of the discrete inf–sup constant

The inf–sup constant depends on the shape of the domain Ω . Thus the constant for a step domain with a long channel is much smaller than the constant for a square domain. Estimation of the inf–sup constant dynamically provides useful information regarding the connection between the velocity error and the pressure error. The inf–sup constant also features in the norm equivalence between the residual norm of the saddle-point system and the natural “energy” norm $\|\nabla u\|_{L^2(\Omega)} + \|p\|_{L_0^2(\Omega)}$, as discussed in the motivating paper [21].

A typical strategy to estimate the inf–sup constant for (1) is to find the largest γ that satisfies (3), i.e., to find the smallest nonzero eigenvalue of the generalised eigenvalue problem

$$BA^{-1}B^T v = \lambda M_Q v, \quad (10)$$

where A is prescribed by the velocity basis functions, M_Q by the pressure functions, and B by both the velocity and pressure functions.

Although the inf–sup constant is independent of the choice of these functions, in practice, its computation may be affected. For example, with the common choice to define Q_h^* by the frame (6), Proposition 1 and Remark 2 show that \mathbf{k} lies in the nullspaces of both B^T and M_Q , which means that this generalised eigenvalue problem is singular, i.e., any $\lambda \in \mathbb{R}$ satisfies (10) when $v = \mathbf{k}$. It is known that generalised eigenvalue problems with singular pencils are challenging to solve numerically [11], and additional checks must be performed to ensure that an estimate of γ is not associated with the eigenvector \mathbf{k} . In practice we find that standard methods for computing

eigenvalues of sparse matrices may struggle to accurately compute these eigenvalues, precisely because $BA^{-1}B^T$ and M_Q are singular.

An alternative is to define Q_h^* by a basis instead of the commonly-used frame (6), but it turns out that this is not necessary. The EST-MINRES approach proposed in [21] is much more robust for this singular eigenvalue problem and gives consistently reliable results for certain preconditioners. The intuition is that by ensuring that any solves with M_Q are orthogonal to k within the preconditioned MINRES method, (see Sect. 2.2), we instead solve (10) for $v \perp k$.

We illustrate the EST-MINRES inf-sup constant estimates using two representative test problems that we describe below. All numerical results were obtained using T-IFISS [2] (for triangular elements) and IFISS3D [19] (for cubic elements). The stopping criterion for preconditioned MINRES is a reduction of the norm of the preconditioned residual by eight orders of magnitude, i.e., $\|r_k\|_{\mathcal{P}^{-1}}/\|r_0\|_{\mathcal{P}^{-1}} < 10^{-8}$. We apply two different block diagonal preconditioners (4). The first is

$$\mathcal{P}_1 = \begin{bmatrix} A & \\ & M_Q \end{bmatrix}.$$

Although this preconditioner is expensive to apply, because it involves exact solves with A , we have used it here to more clearly illustrate the key findings of this section, namely, that (4) with $M_S = M_Q$ is an effective preconditioner for augmented Taylor–Hood problems, despite the singularity of M_Q , and that EST-MINRES provides reliable approximations to γ , the discrete inf-sup constant in (3). However, we note that it is possible to replace A by, e.g., an algebraic multigrid method such as the HSL code MI20 [12]. With this AMG approximation, we are able to solve a 3D problem with nearly 10^6 degrees of freedom in under 10s on a quad-core, 62 GB RAM, Intel i7-6700 CPU with 3.20GHz. Moreover, the effect on the inf-sup constant is fairly small.

Our second, cheaper, preconditioner is

$$\mathcal{P}_2 = \begin{bmatrix} A_{\text{AMG}} & \\ & M_{\text{cheb}} \end{bmatrix},$$

where A_{AMG} is one AMG V-cycle, with the default parameters in T-IFISS and IFISS3D and M_{cheb} is 20 iterations of a Chebyshev semi-iteration method based on symmetric Gauss–Seidel.

Test problem 1 (two-dimensional enclosed flow) Our first example is a classical driven-cavity flow in the square domain $D = [-1, 1]^2$. A Dirichlet no-flow condition is imposed on the bottom and side boundaries, while on the lid the nonzero tangential velocity is $u_y = 1 - x^4$. The domain is subdivided uniformly into n^2 bisected squares. We use the standard $\mathbb{P}_2\text{--}\mathbb{P}_1$ Taylor–Hood mixed approximation and the augmented $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ approximation. The two components of the pressure solution for the $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ approximation, computed for $n = 32$, are illustrated in Fig. 1. The centroid pressure field is concentrated in the two corners where the pressure is singular, and the centroid pressures are an order of magnitude smaller than the vertex pressure in all elements.




Fig. 1 Representative $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ pressure field solution for the cavity flow in Example 1 computed on a uniform mesh with 2048 right-angled triangles and 1089 vertices. Contours are equally spaced between the maximum and minimum values in both plots

Consequently, the overall pressure field is visually identical to the \mathbb{P}_1 pressure field shown in the left plot in Fig. 1.

Test problem 2 (three-dimensional enclosed flow) Our second problem is a three-dimensional version of driven-cavity flow. The domain is now $D = [-1, 1]^3$. As in the previous example, the flow is enclosed, but now the nonzero tangential velocity $u_y = (1 - x^4)(1 - z^4)$ is specified on the top of the cavity. The domain is subdivided uniformly into n^3 cubic elements, and we use $\mathbb{Q}_2\text{--}\mathbb{Q}_1$ and $\mathbb{Q}_2\text{--}\mathbb{Q}_1^*$ approximations.

Table 1 shows preconditioned MINRES iteration counts and EST-MINRES discrete inf-sup constant approximations for Example 1 with preconditioner \mathcal{P}_1 . We first note that for both the $\mathbb{P}_2\text{--}\mathbb{P}_1$ and $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ approximations the iteration counts are quite similar and are mesh independent. In both cases the inf-sup constant approximations also appear to be converging from above. The approximation for $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ elements appears to rapidly converge to four digits, indicating that even a relatively coarse grid is sufficient to obtain an approximation to the discrete inf-sup constant. However, the approximation for $\mathbb{P}_2\text{--}\mathbb{P}_1$ elements appears to converge more slowly. We also note that the two approaches give different inf-sup constant estimates, at least for the grids shown here. This is not so surprising as the matrices in (3) depend on the choice of finite element spaces.

When we replace the preconditioner by the cheaper \mathcal{P}_2 , the situation is largely unchanged for $\mathbb{P}_2\text{--}\mathbb{P}_1$ elements (see Table 2). However, for $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ elements, the preconditioned MINRES convergence rate and inf-sup constant estimate degrade as the mesh is refined. The effect on the convergence rate is caused by the largest non-unit eigenvalue of the symmetric Gauss-Seidel iteration matrix approaching 1 as the mesh is refined; this appears to also impact inf-sup constant approximation.

Figure 2 plots the inf-sup approximations at each iteration of preconditioned MINRES for Example 1 for preconditioner \mathcal{P}_1 . We see that a good approximation of the

Table 1 Preconditioned MINRES iterations and discrete inf–sup constant approximations for Example 1 and preconditioner \mathcal{P}_1

Grid	$\mathbb{P}_2\text{--}\mathbb{P}_1$				$\mathbb{P}_2\text{--}\mathbb{P}_1^*$			
	Velocity dof	Pressure dof	Iters	γ^2	Velocity dof	Pressure dof	Iters	γ^2
4	2178	289	37	0.1947	2178	801	42	0.1397
5	8450	1089	37	0.1926	8450	3137	42	0.1396
6	33,282	4225	39	0.1911	33,282	12417	40	0.1395
7	132,098	16,641	37	0.1898	132,098	49,409	40	0.1395
8	526,338	66,049	37	0.1888	526,338	197,121	40	0.1395

Table 2 Preconditioned MINRES iterations and discrete inf–sup constant approximations for Example 1 and preconditioner \mathcal{P}_2

Grid	$\mathbb{P}_2\text{--}\mathbb{P}_1$				$\mathbb{P}_2\text{--}\mathbb{P}_1^*$			
	Velocity dof	Pressure dof	Iters	γ^2	Velocity dof	Pressure dof	Iters	γ^2
4	2178	289	42	0.1961	2178	801	59	0.1404
5	8450	1089	42	0.1940	8450	3137	59	0.1443
6	33,282	4225	44	0.1925	33,282	12,417	117	0.0467
7	132,098	16,641	45	0.1917	132,098	49,409	152	0.0259
8	526,338	66,049	45	0.1912	526,338	197,121	197	0.0064




Fig. 2 EST-MINRES estimates of the discrete inf–sup constant γ^2 at each iteration for Example 1 and preconditioner \mathcal{P}_1 with $\mathbb{P}_2\text{--}\mathbb{P}_1$ (left) and $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ (right) elements for the grids specified in Table 1

inf–sup constant is obtained after 20–25 iterations. It is again clear that for the enriched Taylor–Hood approximations we obtain very similar approximations for all grids.

The results for Example 2 are given in Tables 3 and 4 and Fig. 3. Preconditioned MINRES iteration counts for \mathcal{P}_1 are again mesh independent and, for all but the coarsest mesh, are almost identical for the two element pairs. The discrete inf–sup constant approximations again appear to converge from above. Now, at least for the grids presented here, the estimates of γ^2 for Taylor–Hood elements seem to “track” the augmented Taylor–Hood estimates, i.e., the Taylor–Hood approximation on grid

Table 3 Preconditioned MINRES iterations and discrete inf–sup constant approximations for Example 2 and preconditioner \mathcal{P}_1

Grid	$\mathbb{Q}_2\text{--}\mathbb{Q}_1$				$\mathbb{Q}_2\text{--}\mathbb{Q}_1^*$			
	Velocity dof	Pressure dof	Iters	γ^2	Velocity dof	Pressure dof	Iters	γ^2
3	2187	125	45	0.1128	2187	189	48	0.1122
4	14,739	729	51	0.1122	14,739	1241	52	0.1115
5	107,811	4913	51	0.1116	107,811	9009	52	0.1110

Table 4 Preconditioned MINRES iterations and discrete inf–sup constant approximations for Example 2 and preconditioner \mathcal{P}_2

Grid	$\mathbb{Q}_2\text{--}\mathbb{Q}_1$				$\mathbb{Q}_2\text{--}\mathbb{Q}_1^*$			
	Velocity dof	Pressure dof	Iters	γ^2	Velocity dof	Pressure dof	Iters	γ^2
3	2187	125	50	0.1132	2187	189	53	0.1126
4	14,739	729	54	0.1128	14,739	1241	59	0.1124
5	107,811	4913	56	0.1123	107,811	9009	69	0.1094

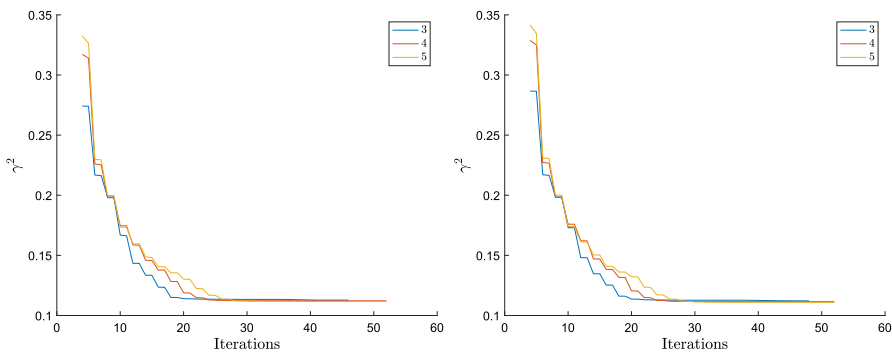


Fig. 3 EST-MINRES estimates of the discrete inf–sup constant γ^2 at each iteration for Example 2 and preconditioner \mathcal{P}_1 with $\mathbb{Q}_2\text{--}\mathbb{Q}_1$ (left) and $\mathbb{Q}_2\text{--}\mathbb{Q}_1^*$ (right) elements for the grids specified in Table 3

j is almost the same as the augmented Taylor–Hood estimate on grid $j - 1$. Table 4 shows that using the cheaper \mathcal{P}_2 also results in mesh-independent convergence for $\mathbb{Q}_2\text{--}\mathbb{Q}_1$ elements. Additionally, the inf–sup constant approximations seem to “track” the Taylor–Hood estimates in Table 3. Now, however, there is more modest growth in the iteration counts when this approximate preconditioner is used for $\mathbb{Q}_2\text{--}\mathbb{Q}_1^*$ elements and the inf–sup approximations are much closer to those in Table 3.

We also see from Fig. 2, which plots the inf–sup approximations at each iteration of preconditioned MINRES, for preconditioner \mathcal{P}_1 , that the discrete inf–sup constant is already reasonably well approximated after 25 iterations.

3 Two-field pressure preconditioning strategies for Oseen flow

In this section we consider the Oseen problem

$$\begin{aligned} -\nu \Delta \vec{u} + \vec{w} \cdot \nabla \vec{u} + \nabla p &= \vec{0} & \text{in } \Omega, \\ \nabla \cdot \vec{u} &= 0 & \text{in } \Omega, \\ \vec{u} &= \vec{g} & \text{on } \partial\Omega, \end{aligned}$$

that arises from applying fixed point iteration to the Navier–Stokes equations modelling steady flow in a channel domain $\Omega \subset \mathbb{R}^2$ with viscosity parameter ν . The Oseen problem is a linear elliptic system of PDEs and assuming a divergence-free convection field \vec{w} and a smooth (e.g., polygonal) domain Ω , has a unique weak solution for all positive values of the viscosity parameter.¹ The inf–sup stability of our two-field mixed approximation is a sufficient condition for convergence of the mixed approximation to the weak solution of the Oseen problem as $h \rightarrow 0$.

Our focus will be on the associated discrete matrix system,

$$\begin{bmatrix} \mathbf{F} & \mathbf{B}^T \\ \mathbf{B} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix}, \quad (11)$$

where the unknown coefficient vector involves the discrete velocity vector $\mathbf{u} \in \mathbb{R}^{n_u}$ and the two-field pressure vector $\mathbf{p} \in \mathbb{R}^{n_p}$. The right-hand side vectors \mathbf{f} and \mathbf{g} are associated with the specified boundary velocity (as in the Stokes flow case). References to the “residual norm” throughout this section will refer to the standard Euclidean norm for finite dimensional vectors. Thus, for the system (11) the convergence of our iterative solver strategies is assessed by monitoring

$$\|\mathbf{z}\| := \left\| \begin{bmatrix} \mathbf{f} - \mathbf{F}\mathbf{u} - \mathbf{B}^T\mathbf{p} \\ \mathbf{g} - \mathbf{B}\mathbf{u} \end{bmatrix} \right\|_{\ell_2}. \quad (12)$$

The nonsymmetry of the matrix \mathbf{F} means that the iterative solver of choice is GMRES (see [8, Sect.9.1]) together with a block preconditioning operator of the form

$$\mathcal{M} = \begin{bmatrix} \mathbf{M} & \mathbf{B}^T \\ 0 & -\mathbf{M}_S \end{bmatrix}, \quad (13)$$

where \mathbf{M} is an optimal complexity (multigrid) operator effecting the action of the inverse of the matrix \mathbf{F} and \mathbf{M}_S is an optimal complexity approximation of the Schur complement matrix $\mathbf{B}\mathbf{F}^{-1}\mathbf{B}^T$. We will discuss results for two representative flow problems herein.

¹ Uniquely defined solutions to the underlying steady-state Navier–Stokes problem are only guaranteed when the viscosity parameter is sufficiently large; see [8, Sect.8.2.1].

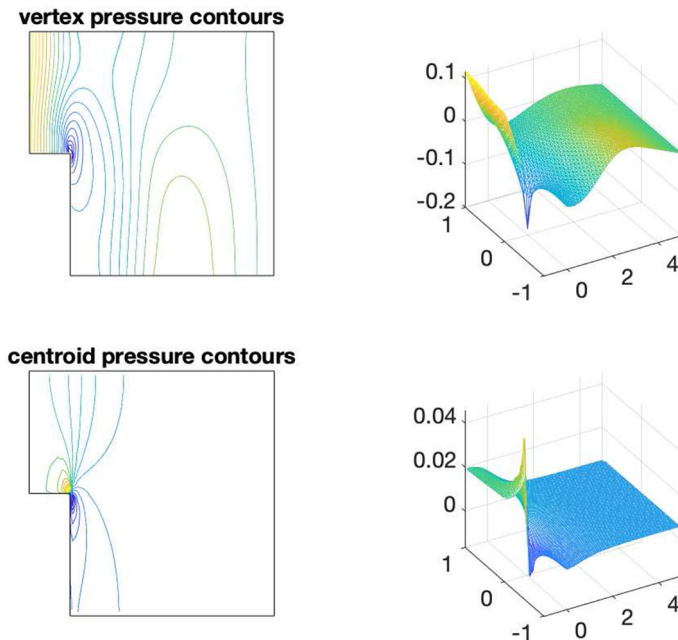


Fig. 4 Representative $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ pressure field solution for flow over a step ($\mathcal{R} = 100$) computed on a uniform mesh with 5632 right-angled triangles and 2945 vertices

Test problem 3 (flow over a step) We consider an inflow–outflow problem defined on the domain $D = [-1, 0) \times [0, 1] \cup [0, 5] \times [-1, 1]$ with the viscosity parameter ν set to $1/50$ and a parabolic velocity $u_x = 4y(1 - y)$ specified on the inflow boundary $x = -1$. The construction of the standard weak formulation (see [8, p. 127]) gives rise to a natural boundary condition that fixes the hydrostatic pressure level by weakly enforcing a zero mean pressure at the outflow boundary $x = 5$. (The other boundary conditions are associated with fixed walls.) We consider the discrete system (11) that results after 5 fixed-point iterations of the discretised Navier–Stokes system starting from the corresponding Stokes flow solution. We generate solutions using $\mathbb{Q}_2\text{--}\mathbb{Q}_1^*$ or $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ augmented Taylor–Hood mixed approximation with the domain subdivided uniformly into (bisected) squares. The resulting system (11) is singular with the one-dimensional pressure nullspace described in Sect. 2. The two components of the pressure solution computed on a representative $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ mesh are illustrated in Fig. 4. The centroid pressure values are an order of magnitude smaller than the vertex pressure in all the elements—they provide the “corrections” to the vertex pressures that are needed to ensure local (elementwise) conservation of mass.

Test problem 4 (two-dimensional enclosed flow) We consider the classical driven-cavity enclosed flow problem defined on the domain $D = [-1, 1]^2$ with the viscosity parameter ν set to $1/100$ and a nonzero tangential velocity $u_y = 1 - x^4$ specified on the top of the cavity. We take $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ augmented Taylor–Hood mixed approximation with the domain subdivided uniformly into n^2 bisected squares. We consider the discrete




Fig. 5 Comparison of divergence error and \mathbb{P}_2 - \mathbb{P}_1^* centroid pressure solution for cavity flow ($\mathcal{R} = 200$) computed on a uniform mesh with 2048 right-angled triangles and 1089 vertices

system (11) that arises after 5 fixed-point iterations starting from the corresponding Stokes flow solution. The discrete system is singular with a two-dimensional pressure nullspace, corresponding to a constant vertex pressure and a constant centroid pressure. The contour plot on the left in Fig. 5 shows the element divergence errors $\|\nabla \cdot \bar{u}_h\|_{L^2(\Delta)}$ computed on a coarse mesh ($n = 32$). The associated centroid “correction” pressure field contours shown on the plot in the right can be seen to provide a good indication of the regions where the divergence error is concentrated. As in Fig. 4 the centroid pressure values are an order of magnitude smaller than the vertex pressure in all elements, so the overall pressure field is visually identical to the \mathbb{P}_1 pressure field (not shown here).

3.1 Pressure convection-diffusion preconditioning for Oseen flow

There are two alternative ways of approximating the key matrix $B\mathbf{F}^{-1}B^T$ in the case that B^T is generated by a two-field pressure approximation so that $B^T = [B_1^T, B_0^T]$. The focus will be on pressure convection-diffusion (PCD) preconditioning in this section. Both of the Schur complement approximations can be motivated by starting with the Oseen matrix operator (11) and observing that the diagonal blocks of \mathbf{F} are discrete representations of the convection–diffusion operator

$$\mathcal{L} = -\nu\nabla^2 + \bar{w}_h \cdot \nabla, \quad (14)$$

defined on the velocity space. In practical calculations $\nu > 0$ is proportional to the inverse of the flow Reynolds number and \bar{w}_h is the discrete approximation to the flow velocity computed at the most recent nonlinear iteration. The PCD approximation supposes that there is an analogous operator to (14), namely

$$\mathcal{L}_p = (-\nu\nabla^2 + \bar{w}_h \cdot \nabla)_p \quad (15)$$

defined on the two components of the augmented pressure space.

To this end, defining $\{\phi_j\}_{j=1}^{n_1}$ to be the basis for the C^0 pressure discretisation, we construct matrices Q_1 and F_1 so that

$$Q_1 = [q_{1,ij}], \quad q_{1,ij} = \int_{\Omega} \phi_j \phi_i$$

$$F_1 = [f_{1,ij}], \quad f_{1,ij} = \nu \int_{\Omega} \nabla \phi_j \cdot \nabla \phi_i + \int_{\Omega} (\bar{w}_h \cdot \nabla \phi_j) \phi_i.$$

We then note that if the commutator with the divergence operator

$$\mathcal{E} = \nabla \cdot (-\nu \nabla^2 + \bar{w}_h \cdot \nabla) - (-\nu \nabla^2 + \bar{w}_h \cdot \nabla)_p \nabla \cdot$$

is small then we have the approximation

$$0 \approx (Q_1^{-1} B_1) (M^{-1} F) - (Q_1^{-1} F_1) (Q_1^{-1} B_1) \tag{16}$$

where M is the diagonal of the mass matrix associated with the basis representation of the velocity space.² Rearranging (16) gives the first Schur complement approximation

$$B_1 F^{-1} B^T \approx Q_1 F_1^{-1} (B_1 M^{-1} B^T). \tag{17}$$

A discrete version of \mathcal{L}_p for the piecewise constant pressure space can be generated by considering the jumps in pressure across inter-element boundaries; see [8, pp. 268–370]. To this end, defining $\{\varphi_j\}_{j=1}^{n_0}$ to be the (indicator function) basis for the discontinuous pressure, we construct matrices Q_0 and F_0 via

$$Q_0 = [q_{0,ij}], \quad q_{0,ij} = \int_{\Omega} \varphi_j \varphi_i = \begin{cases} |T_i| & \text{if } i = j, \\ 0 & \text{otherwise,} \end{cases}$$

$$F_0 = [f_{0,ij}], \quad f_{0,ij} = \nu \sum_{T \in \mathcal{T}_h} \int_T \nabla \varphi_j \cdot \nabla \varphi_i + \sum_{T \in \mathcal{T}_h} \int_T (\bar{w}_h \cdot \nabla \varphi_j) \varphi_i$$

and note that if the commutator with the divergence operator is small then we have

$$0 \approx (Q_0^{-1} B_0) (M^{-1} F) - (Q_0^{-1} F_0) (Q_0^{-1} B_0), \tag{18}$$

suggesting the second Schur complement approximation

$$B_0 F^{-1} B^T \approx Q_0 F_0^{-1} (B_0 M^{-1} B^T). \tag{19}$$

² The inverse of the mass matrix is a dense matrix. The diagonal of the mass matrix is a spectrally equivalent matrix operator with a sparse (diagonal) inverse.




Fig. 6 Absolute residual reduction for test problem 3 when computing $\mathbb{P}_2\text{-}\mathbb{P}_1^*$ solutions using preconditioners \mathcal{M}_1 (x) or \mathcal{M}_2 (o) on two nested meshes

Combining (17) with (19) then gives a two-field PCD approximation

$$BF^{-1}B^T \approx M_S := \begin{bmatrix} Q_1 & 0 \\ 0 & Q_0 \end{bmatrix} \begin{bmatrix} F_1^{-1} & 0 \\ 0 & F_0^{-1} \end{bmatrix} BM^{-1}B^T. \tag{20}$$

Two features of the PCD approximation (20) are worth noting. The first point is that the coupling between the pressure components is represented by the 2×2 block matrix $BM^{-1}B^T$ rather than by the pressure mass matrix M_Q . The second key point is that the matrices M_S and $BF^{-1}B^T$ have the same nullspace, independent of the nature of underlying flow problem that is being solved.

The PCD approximation in (20) is imperfect in practice. To illustrate this, representative convergence histories that arise in solving the inflow-outflow problem using $\mathbb{P}_2\text{-}\mathbb{P}_1^*$ approximation (shown in Fig. 4) are presented in Fig. 6. Taking the solution from the previous Picard iteration as the initial guess, we note that the initial residual norm of the target matrix system (11) is close to 10^{-4} independent of the spatial discretisation. Convergence plots are presented for two preconditioning strategies, namely

$$\mathcal{M}_1 = \begin{bmatrix} F & B^T \\ 0 & -\frac{1}{\nu}M_Q \end{bmatrix}, \quad \mathcal{M}_2 = \begin{bmatrix} F & B^T \\ 0 & -M_S \end{bmatrix}, \tag{21}$$

with M_S defined in (20). The first strategy is the block triangular extension of the Stokes preconditioning strategy discussed in Sect. 2. We note that the resulting convergence is very slow—around 50 iterations are required to reduce the residual norm by an order of magnitude—but is independent of the discretisation level. In contrast we see that the PCD preconditioning strategy has two distinctive phases of convergence behaviour. An initial phase of relatively fast convergence is followed by a secondary phase where GMRES stagnates. We hypothesise that this stagnation is a consequence




Fig. 7 Absolute residual reduction for test problem 3 when computing $\mathbb{Q}_2\text{-}\mathbb{Q}_1^*$ solutions using preconditioners \mathcal{M}_1 (x) or two-stage PCD (Step I: o; Step II: o) on two nested meshes

of the ill-conditioning of the eigenvectors of the matrix operator M_S . A notable feature is that the onset of the stagnation is delayed when solving the same problem on a finer grid. The two-phase convergence behaviour is ubiquitous—the same pattern is seen using rectangular elements and the stagnation does not go away when the viscosity parameter is increased from $1/50$ to $1/5$. An alternative strategy is clearly needed!

One way of designing a more robust PCD preconditioning strategy for a two-field pressure approximation is to exploit the fast convergence of the PCD approximation for the unaugmented Taylor-Hood approximation. The starting point for such a strategy is to rewrite the system (11) in the form

$$\begin{bmatrix} \mathbf{F} & B_1^T & B_0^T \\ B_1 & 0 & 0 \\ B_0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p}_1 \\ \mathbf{p}_0 \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}. \tag{22}$$

The proposed solution algorithm is then a two-stage process.

- **Input:** residual reduction tolerance η and reduced system residual factor c

Step I Generate a PCD solution to the *reduced system*

$$\begin{bmatrix} \mathbf{F} & B_1^T \\ B_1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{q}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{0} \end{bmatrix} \tag{23}$$

using the Schur complement approximation (16), stopping the GMRES iteration when the residual is reduced by a factor of $c\eta$.

Step II Generate a solution to the target system (22) with residual tolerance η using preconditioning strategy \mathcal{M}_1 in (21) with the refined initial guess $[\mathbf{u}_1^*, \mathbf{q}_1^*, \mathbf{0}]$.

- **Output:** refined solution $[\mathbf{u}^*, \mathbf{p}_1^*, \mathbf{p}_0^*]$




Fig. 8 Absolute residual reduction for test problem 4 when computing $\mathbb{P}_2\text{-}\mathbb{P}_1^*$ solutions using preconditioners \mathcal{M}_1 (x) or two-stage PCD (Step I: o; Step II: o) on two nested meshes

Sample results generated using this strategy with tolerance η set to 10^{-4} and c set to 10 are presented in Fig. 7. Sample results for the cavity test problem are presented in Fig. 8. The results in Figs. 7 and 8 are representative of the two-stage convergence profiles that are generated when solving these test problems at other Reynolds numbers. Our experience is that the level of residual reduction is perfectly robust with regards to the spatial discretisation—typically giving smaller iteration counts when the mesh resolution is increased (a known feature of PCD preconditioning). The convergence rates of both stages of the algorithm deteriorate slowly when the Reynolds number is increased. Our strategy for terminating the first stage of the iteration is motivated by the following result.

Proposition 4 *The residual error $\|\mathbf{z}^*\|$ associated with the intermediate solution $[\mathbf{u}_1^*, \mathbf{q}_1^*, \mathbf{0}]$ to the discrete system (22) satisfies the bound*

$$\|\mathbf{z}^*\|^2 \leq c^2 \eta^2 \|\mathbf{f}\|^2 + \|B_0 \mathbf{u}_1^*\|^2, \tag{24}$$

where $\|\mathbf{f}\|$ is the initial residual error associated with a zero initial vector.

Proof The vector \mathbf{z}^* associated with the intermediate solution is the three-component vector

$$\begin{bmatrix} \mathbf{r}^* \\ \mathbf{r}_1^* \\ \mathbf{r}_0^* \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathbf{F} & B_1^T & B_0^T \\ B_1 & 0 & 0 \\ B_0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u}_1^* \\ \mathbf{q}_1^* \\ \mathbf{0} \end{bmatrix}. \tag{25}$$

The stopping test for solving the reduced system ensures that

$$\|\mathbf{r}^*\|^2 + \|\mathbf{r}_1^*\|^2 \leq c^2 \eta^2 \|\mathbf{f}\|^2. \tag{26}$$

Thus we have

$$\|\mathbf{z}^*\|^2 = \|\mathbf{r}^*\|^2 + \|\mathbf{r}_1^*\|^2 + \|\mathbf{r}_0^*\|^2 \leq c^2 \eta^2 \|\mathbf{f}\|^2 + \|B_0 \mathbf{u}_1^*\|^2. \tag{27}$$

□

The bound (24) has two terms on the right-hand side. While the first term can be controlled by reducing η and/or c , the second term measures the local incompressibility of the intermediate Taylor–Hood solution $\|B_0 \mathbf{u}_1^*\|^2 = \sum_j (\int_{T_j} \nabla \cdot \bar{u}_h^*)^2$, where \bar{u}_h^* is the expansion of the coefficient vector \mathbf{u}_1^* in the basis of the velocity approximation space. Setting $\eta = 10^{-4}$ and $c = 10$ we see that the second term saturates the residual error and the residual error jumps up when the switch is made from the first to the second step of the algorithm. This “transition” jump in the residual norm is clearly evident in the convergence plots. The convergence in the second step is rapid initially but eventually mirrors the rate observed for \mathcal{M}_1 approximation with a standard starting guess.

3.2 Least square commutator approximation for Oseen flow

A second way of approximating the key matrix $B F^{-1} B^T$ in the case that B^T is generated by a two-field pressure approximation is given by the least-squares commutator (LSC) preconditioner

$$B F^{-1} B^T \approx M_S = (B H^{-1} B^T) (B M^{-1} F H^{-1} B^T)^{-1} (B M^{-1} B^T). \tag{28}$$

The attractive feature of LSC is that the construction of M_S is completely algebraic. The only technical issue is the need to make adjustments on rows and columns associated with tangential velocity degree of freedom adjacent to inflow and fixed wall boundaries. These adjustments are associated with a diagonal scaling matrix D so that $H = D^{-1/2} M D^{-1/2}$. Full details can be found in [8, pp. 376–379].

The LSC approximation (28) is also far from perfect when using triangular elements.³ To illustrate this, representative convergence histories that arise in solving the inflow-outflow problem using $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ approximation are presented in Fig. 9. Convergence plots are presented for two preconditioning strategies, namely the refined PCD from the previous section and

$$\mathcal{M}_3 = \begin{bmatrix} F & B^T \\ 0 & -M_S \end{bmatrix}, \tag{29}$$

with M_S defined in (28).

The PCD histories in Fig. 9 display mesh independent convergence and are comparable with those observed using square elements in Fig. 7. The associated cpu times for the solution on the intermediate grid with 2×11264 elements were 13 s for the

³ The strategy is designed for tensor-product approximation spaces. It gives good results in the case of $\mathbb{Q}_2\text{--}\mathbb{Q}_1^*$.

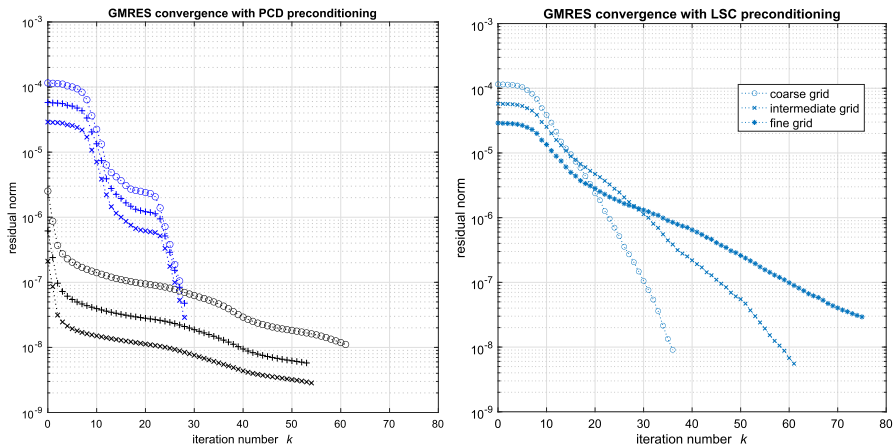


Fig. 9 Absolute residual reduction for test problem 3 when computing $\mathbb{P}_2\text{--}\mathbb{P}_1^*$ solutions using using refined PCD (left) or preconditioner \mathcal{M}_3 (right) on three nested meshes

first step (28 iterations) and 26 s for the second step (53 iterations). The LSC preconditioning strategy is not robust—the convergence rate deteriorates with increasing grid refinement. The cpu time for generating the LSC solution on the intermediate grid was over 100 s (61 iterations).

4 Conclusions

Two-level pressure approximation for incompressible flow problems offer the prospect of accurate approximation with minimal computational overhead. Derived quantities of practical importance such as the mean wall shear stress are likely to be computed much more precisely if incompressibility is enforced locally.⁴ However, the augmented pressure space causes some challenges for the linear algebra when the pressure space is defined by combining the usual Taylor–Hood pressure basis functions with basis functions for the piecewise continuous pressure space. This is because constant functions can be expressed using either the usual (continuous) Taylor–Hood pressure space, or the augmented piecewise constant pressure space. Specifically, with this common choice for specifying the pressure space, the pressure mass matrix becomes singular, and care should be taken to carefully construct and apply preconditioners that involve this matrix. Care should also be exercised when approximating the discrete inf–sup constant, and we find that naive approaches are not always reliable. On the other hand, the approximation implemented in EST-MINRES is robust. Our computational experimentation indicates that our two-stage PCD strategy could be the best way of iteratively solving two-level pressure discrete linear algebra systems in the sense of algorithmic reliability and computational efficiency.

⁴ See <https://personalpages.manchester.ac.uk/staff/david.silvester/lecture1.18.mp4> for a comparison of alternative strategies for computing the average shear stress for flow over a step at Reynolds number 200.

Acknowledgements This work was supported by EPSRC grant EP/W033801/1.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Benzi, M., Golub, G.H., Liesen, J.: Numerical solution of saddle point problems. *Acta Numer.* **14**, 1–137 (2005)
2. Bespalov, A., Rocchi, L., Silvester, D.: T-IFISS: a toolbox for adaptive FEM computation. *Comput. Math. Appl.* **81**, 373–390 (2021)
3. Boffi, D., Brezzi, F., Fortin, M.: *Mixed Finite Element Methods and Applications*, Springer Series in Computational Mathematics 44. Springer, Berlin (2013)
4. Boffi, D., Cavallini, N., Gardini, F., Gastaldi, L.: Local mass conservation of Stokes finite elements. *J. Sci. Comput.* **52**, 383–400 (2012)
5. Cao, Z.: Polynomial acceleration methods for solving singular systems of linear equations. *J. Comput. Math.* **9**, 378–387 (1991)
6. Chaabane, N., Girault, V., Riviere, B., Thompson, T.: A stable enriched Galerkin element for the Stokes problem. *Appl. Numer. Math.* **132**, 1–21 (2018)
7. Dax, A.: The convergence of linear stationary iterative processes for solving singular unstructured systems of linear equations. *SIAM Rev.* **32**, 611–635 (1990)
8. Elman, H., Silvester, D., Wathen, A.: *Finite Elements and Fast Iterative Solvers: with Applications in Incompressible Fluid Dynamics*, 2nd edn. Oxford University Press, Oxford (2014)
9. Gresho, P., Lee, R., Chan, S., Leone, J.: A new finite element for incompressible or Boussinesq fluids. In: *Third International Conference on Finite Elements in Flow Problems*, pp. 204–215. Wiley, New York (1981)
10. Griffiths, D.: The effect of pressure approximation on finite element calculations of incompressible flows. In: Morton, K.W., Baines, M.J. (eds.) *Numerical Methods for Fluid Dynamics*, pp. 359–374. Academic Press, San Diego (1982)
11. Hochstenbach, M.E., Mehl, C., Plestenjak, B.: Solving singular generalized eigenvalue problems by a rank-completing perturbation. *SIAM J. Matrix Anal. Appl.* **40**, 1022–1046 (2019)
12. HSL: A collection of Fortran codes for large scale scientific computation. <http://www.hsl.rl.ac.uk/>
13. Hu, X., Lee, S., Lin, M., Yi, S.: Pressure robust enriched Galerkin methods for the Stokes equations. *J. Comput. Appl. Math.* **436**, 115449 (2024)
14. Kuznetsov, Y.A.: Efficient iterative solvers for elliptic finite element problems on nonmatching grids. *Russ. J. Numer. Anal. Math. Modelling* **10**, 187–211 (1995)
15. Lee, S., Lee, Y.-J., Wheeler, M.: A locally conservative enriched Galerkin approximation and efficient solver for elliptic and parabolic problems. *SIAM J. Sci. Comput.* **38**, A1404–A1429 (2016)
16. Ludwig, E., Nabben, R., Tang, J.: Deflation and projection methods applied to symmetric positive semi-definite systems. *Linear Algebra Appl.* **489**, 253–273 (2016)
17. Murphy, M.F., Golub, G.H., Wathen, A.J.: A note on preconditioning for indefinite linear systems. *SIAM J. Sci. Comput.* **21**, 1969–1972 (2000)
18. Paige, C., Saunders, M.: Solution of sparse indefinite systems of linear equations. *SIAM J. Numer. Anal.* **12**, 617–629 (1975)
19. Papanikos, G., Powell, C.E., Silvester, D.J.: IFISS3D: a computational laboratory for investigating finite element approximation in three dimensions. *ACM Trans. Math. Softw.* **49**, 30 (2023)
20. Qin, J., Zhang, S.: Stability of the finite elements $9/(4c + 1)$ and $9/5c$ for stationary Stokes equations. *Comput. Struct.* **84**, 70–77 (2005)

21. Silvester, D.J., Simoncini, V.D.: An optimal iterative solver for symmetric indefinite systems stemming from mixed approximation. *ACM Trans. Math. Softw.* **37**, 42 (2011)
22. Sun, S., Liu, J.: A locally conservative finite element method based on piecewise constant enrichment of the continuous Galerkin method. *SIAM J. Sci. Comput.* **31**, 2528–2548 (2009)
23. Thatcher, R.: Locally mass-conserving Taylor-Hood elements for two- and three-dimensional flow. *Int. J. Numer. Methods Fluids* **11**, 341–353 (1990)
24. Thatcher, R., Silvester, D.: A locally mass conserving quadratic velocity, linear pressure element. Manchester Centre for Computational Mathematics report 147 1987. [arXiv:2001.11878](https://arxiv.org/abs/2001.11878) [math.NA]
25. Wathen, A., Rees, T.: Chebyshev semi-iteration in preconditioning for problems including the mass matrix. *Electron. Trans. Numer. Anal.* **34**, 125–135 (2009)
26. Wathen, A.J.: Realistic eigenvalue bounds for the Galerkin mass matrix. *IMA J. Numer. Anal.* **7**, 449–457 (1987)
27. Yi, S., Hu, X., Lee, S., Adler, J.: An enriched Galerkin method for the Stokes equations. *Comput. Math. Appl.* **120**, 115–131 (2022)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.