Science and Technology Facilities Council

Preprint
RAL-P-2020-002

An element-based preconditioner for mixed finite element problems

T Rees, M Wathen

May 2020

Submitted for publication in SIAM Journal on Scientific Computing

Enquiries concerning this report should be addressed to:

RAL Library
STFC Rutherford Appleton Laboratory
Harwell Oxford
Didcot
OX11 0QX

Tel: +44(0)1235 445384
Fax: +44(0)1235 446677
email: libraryral@stfc.ac.uk

Science and Technology Facilities Council reports are available online at:
https://epubs.stfc.ac.uk

ISSN 1361-4762

An element-based preconditioner for mixed finite element problems *

Tyrone Rees† and Michael Wathen†

Abstract. We introduce a new and generic approximation to Schur complements arising from inf-sup stable mixed finite element discretizations of self-adjoint multi-physics problems. The approximation exploits the discretization mesh by forming local, or element, Schur complements and projecting them back to the global degrees of freedom. The resulting Schur complement approximation is sparse, has low construction cost (with the same order of operations as assembling a general finite element matrix), and can be solved using off-the-shelf techniques, such as multigrid. Using results from saddle point theory, we give conditions such that this approximation is spectrally equivalent to the global Schur complement. We present several numerical results to demonstrate the viability of this approach on a range of applications. Interestingly, numerical results show that the method gives an effective approximation to the non-symmetric Schur complement from the steady state Navier-Stokes equations.

Key words. saddle-point linear systems, preconditioners, Krylov subspace methods, finite element methods, Schur complements

AMS subject classifications. 65F08, 65F10, 65F15, 65F50, 65N22, 74S05

1. Introduction.

We seek $(u,p) \in \mathcal{X} \times \mathcal{M}$ that solves the saddle point problem

$$\begin{aligned} a(u,v) + b(v,p) &= \langle f, v\rangle \quad \forall v \in \mathcal{X}, \\ b(u,q) &= \langle g, q\rangle \quad \forall q \in \mathcal{M}, \end{aligned} \tag{1.1}$$

where $\mathcal{X}$ and $\mathcal{M}$ are two Hilbert spaces, $a(\cdot,\cdot): \mathcal{X}\times\mathcal{X} \to \mathbb{R}$ and $b(\cdot,\cdot): \mathcal{X}\times\mathcal{M}\to\mathbb{R}$ are two bounded bilinear forms, and $(f,g) \in \mathcal{X}' \times \mathcal{M}'$, where $\mathcal{X}'$ and $\mathcal{M}'$ are the dual spaces of $\mathcal{X}$ and $\mathcal{M}$, respectively. This is a well-posed problem provided that it satisfies certain conditions, given in Section 4.1. Such problems arise in fields as diverse as constrained optimization, constrained least-squares, fluid dynamics, electomagnetics, elasticity, optimal control, and many others (see, for example, the references within the survey of Benzi, Golub and Liesen [8]).

We focus on the case where (1.1) comes from the solution of a partial differential equation (PDE). We discretize using the Finite Element Method, and in particular, by choosing finite dimensional spaces $\mathcal{X}_h \subset \mathcal{X}$ and $\mathcal{M}_h \subset \mathcal{M}$. This entails overlaying the domain with a grid and use this to define local elements, which gives an elemental structure. We select basis functions that have support only on a small number of neighbouring elements, which are usually chosen to be unity at one mesh point and vanish at the others.

Suppose we have such a pair of spaces defined by the basis functions $\phi_i \in \mathcal{X}_h$, $\psi_i \in \mathcal{M}_h$; we assume they share the same elements, but the basis functions may differ. We look for an approximation $(u_h, p_h) = \left(\sum_i \mathbf{u}_i\phi_i, \sum_j \mathbf{p}_j\psi_j\right)$, where we find the coefficents of each basis

*This work was supported by EPSRC Grant EP/M025179/1

†STFC Rutherford Appleton Laboratory, Chilton, Didcot, UK (tyrone.rees@stfc.ac.uk, michael.wathen@stfc.ac.uk)

element by solving a linear system of the form

$$\underbrace{\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix}}_{\mathcal{A}} \underbrace{\begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix}}_{\mathbf{x}} = \underbrace{\begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix}}_{\mathbf{b}}, \tag{1.2}$$

where $A_{i,j} = a(\phi_i, \phi_j) \in \mathbb{R}^{n\times n}$ and $B_{i,j} = b(\phi_j, \psi_i) \in \mathbb{R}^{m\times n}$ (see, e.g., [10, 16] for more detail).

We can decompose the matrices in (1.2) into a sum of σ elemental matrices. On the eth element, let n_e and m_e be the number of local degrees of freedom in $\mathcal{X}_h$ and $\mathcal{M}_h$, respectively (see Figure 1 for an example). Let the small dense matrices $A_e \in \mathbb{R}^{n_e\times n_e}$ and $B_e \in \mathbb{R}^{m_e\times n_e}$ be the element equivalents of A and B.

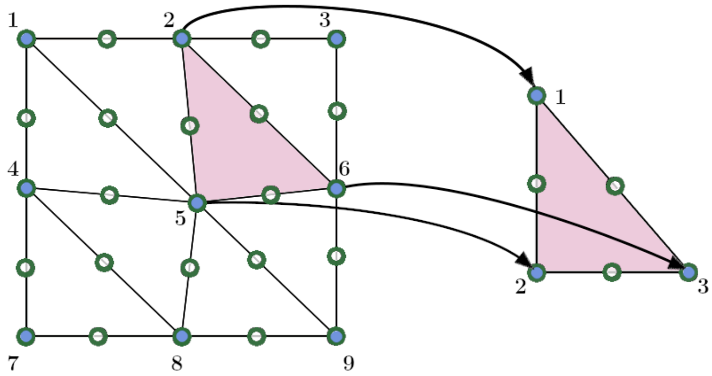

Figure 1: Global (left) and local (right) meshes for a P2-P1 triangulation of a domain. Here $n = 25$, $m = 9$, $n_e = 6$, $m_e = 3$, and $\sigma = 8$.

We introduce Boolean matrices $L_e \in \mathbb{R}^{n_e\times n}$ and $N_e \in \mathbb{R}^{m_e\times m}$ that map the local orderings to the global orderings. Thus, we write

$$A = \sum_e L_e^T A_e L_e := L^T \widehat{A_e} L, \quad B = \sum_e N_e^T B_e L_e := N^T \widehat{B_e} L. \tag{1.3}$$

where $L^T = [L_1^T, \dots L_\sigma^T]$, $N^T = [N_1^T, \dots N_\sigma^T]$, $\widehat{A_e} = \texttt{blkdiag}(A_e)$, and $\widehat{B_e} = \texttt{blkdiag}(B_e)$.

As we will outline in Section 3.1, preconditioners based on Schur complement approximations work well for systems of the form (1.2). In this paper we describe a sparse approximation to the Schur complement based on local contributions by the element matrices:

$$\hat{S}_{\text{dual}} = \sum_e N_e^T B_e Y_e^{-1} B_e{}^T N_e \quad \text{and} \quad \hat{S}_{\text{primal}} = \sum_e L_e^T (A_e + B_e{}^T W_e{}^{-1} B_e) L_e, \tag{1.4}$$

where Y_e and W_e are local element matrices that, if assembled, would form symmetric positive definite weighting matrices Y and W, respectively. In (1.4), we form the Schur complement on an element and then map the local matrix into a global matrix by using the standard assembly process. We call the matrices $B_e Y_e^{-1} B_e^T$ and $A_e + B_e{}^T W_e{}^{-1} B_e$ the *dual* and *primal* element Schur complement, respectively. The idea of using elements to define preconditioners is not new, and we outline their development in the literature in Section 2.

In Section 3, we review preconditioning for saddle point systems, in particular looking at (ideal) Schur complement preconditioners and natural norm, or Reisz-map, preconditioners. We further explore the link between these methods in Section 4, where we extend existing results in the literature. In Section 5, we prove our main result that, under certain conditions on Y and W, $\hat{S}_{\text{dual}}$ and $\hat{S}_{\text{primal}}$ are spectrally equivalent to the respective global Schur complements.

We apply this preconditioner to three well-known examples of saddle-point systems, Stokes equation (Section 6.1), Maxwell's equation (Section 6.2) and the Navier-Stokes equation (Section 6.3). Finally, we conclude in Section 7.

2. Related work.

Many people have proposed methods that use individual elements to accelerate the solution of systems of the form (1.2). Hughes, Levit and Winget [27] and Ortiz, Pinksy and Taylor [40] introduced element-by-element (EBE) methods for solving self-adjoint PDEs in the early eighties. Such methods approximate the sum in (1.3) by a product. Nour-Omid and Parlett [39] show that we may apply the EBE method as a preconditioner for conjugate gradients [26], and Wathen [52] and van Gijzen [51] give an analysis of such methods for symmetric and non-symmetric problems, respectively. Gustafsson and Lindskog [25] also exploit local contributions by developing a preconditioner based on global structure preserving Cholesky factorizations on elements.

For symmetric positive definite systems, Kraus [30] described the use of local (dual) Schur complements to build a preconditioner, and Axelsson, Baheta and Neytcheva [5] give descriptive eigenvalue bounds for $\hat{S}_{\text{dual}}^{-1} S_{\text{dual}}$, where $Y_e = A_e$, which rely only on the Cauchy-Bunyakowski-Schwarz constant. Neytcheva and co-authors [36, 37, 38] apply this idea to non-positive definite systems, and Neytcheva [36] gives a bound on $\|\hat{S}_{\text{dual}}^{-1} S_{\text{dual}}\|$ (again for $Y_e = A_e$) based on the norms of relations involving the constituent blocks. In that work, Neytcheva also shows numerically that this is a viable method for solving Stokes equation and the Oseen equation.

More recently, there has been interest in multilevel preconditioners based on local subdomains or elements. For example, in the context of overlapping domain decomposition, GENEO [47, 48] is designed specifically for heterogeneities within the variational form by incorporating local generalized eigenvalue problems on overlaps to define the coarsening. Also, for multigrid methods, PCPATCH [19, 20] derives effective relaxation methods based on specific "patches", or local contributions from elements.

3. Iterative solution of saddle point systems.

The convergence of a Krylov subspace method applied to (1.2) is generally unsatisfactory unless paired with a suitable preconditioner, $\mathcal{P}$, which has two, competing, properties: a solve with $\mathcal{P}$ must be cheap, and the Krylov method applied to $\mathcal{P}^{-1}\mathcal{A}$ must converge more quickly. We might satisfy the latter condition if, for example, the eigenvalues of $\mathcal{P}^{-1}\mathcal{A}$ are in clusters away from the origin. We recommend the surveys by Wathen [54] and Benzi, Golub, and Leisen [8] for a comprehensive overview of the state-of-the-art in preconditioning techniques in general (in the former), and for saddle point matrices in particular (in the latter).

In the case where the leading matrix A in (1.2) is symmetric, the saddle point system

$\mathcal{A}$ is symmetric but indefinite, and MINRES [41] is the method of choice. MINRES requires a symmetric positive definite preconditioner, and convergence is exactly described by the eigenvalues of the preconditioned system $\mathcal{P}^{-1}\mathcal{A}$ (see, e.g., [16, Section 4.1]). For nonsymmetric systems, GMRES [45] is often used, but in this case, as shown by Greenbaum, Pták and Strakoš [21], clustered eigenvalues are not necessarily enough to guarantee rapid convergence.

3.1. Schur complement preconditioners. An idea that has proved successful for solving systems of the form (1.2) is to build a preconditioner that exploits an approximation to the Schur complement. The ideal preconditioners are

$$\mathcal{P}_{\text{dual}} = \begin{bmatrix} A & 0 \\ 0 & BA^{-1}B^T \end{bmatrix}, \mathcal{P}_{\text{primal}} = \begin{bmatrix} A + B^T W^{-1} B & 0 \\ 0 & W \end{bmatrix}.$$

We can interpret $\mathcal{P}_{\text{dual}}$ as taking the weighting matrix $Y = A$. These involve the negative of the *dual Schur complement*, $BA^{-1}B^T$, and the *primal Schur complement*, $A + B^T W^{-1} B$. Often the choice between using $\mathcal{P}_{\text{dual}}$ and $\mathcal{P}_{\text{primal}}$ depends on the invertibility of the leading block, A. Murphy, Golub and Wathen [35] and Ipsen [28] show that the preconditioned matrix, $\mathcal{P}_{\text{dual}}^{-1}\mathcal{A}$, has three distinct eigenvalues. Thus, MINRES would converge in precisely three iterations. For non-symmetric A, GMRES usually converges rapidly, but this is not guaranteed. For the primal Schur complement, Greif and Schötzau [22] give bounds for the eigenvalues of $\mathcal{P}_{\text{primal}}^{-1}\mathcal{A}$. For the specific case where the dimension of the null space of leading block A is m (the dimension of the dual variables), Greif and Schötzau [23] show that the preconditioned matrix has two distinct eigenvalues.

Although in one sense these preconditioners are ideal, as convergence is rapid, forming and solving with these is not practical as the primal and dual Schur complements are typically dense. However, they give us something to aim for in a practical preconditioner.

3.2. Operator preconditioning for saddle point problems. Another framework for solving systems of the form (1.2) is operator preconditioning. Given a Hilbert space $\mathcal{V}$, the Riesz map is a mapping $\tau : \mathcal{V}' \to \mathcal{V}$ such that, for any $r \in \mathcal{V}'$,

$$(\tau r, v)_{\mathcal{V}} := \langle r, v \rangle \quad \forall v \in \mathcal{V}.$$

The continuous saddle point system (1.1) is an operator $\mathcal{L} : \mathcal{X} \times \mathcal{M} \to \mathcal{X}' \times \mathcal{M}'$, which returns functions outside of the space $(\mathcal{X}, \mathcal{M})$. Consider the Reisz map

$$\tau = \begin{bmatrix} \chi & 0 \\ 0 & \mu \end{bmatrix}, \tag{3.1}$$

where $\chi : \mathcal{X}' \to \mathcal{X}$ and $\mu : \mathcal{M}' \to \mathcal{M}$ are the Reisz maps associated with the spaces $\mathcal{X}$ and $\mathcal{M}$, respectively. Applying τ to (1.1) gives the transformed equation

$$\tau \mathcal{L}\,(u, p) = \tau\,(f, g)\,,$$

where $\tau\,(f, g) \in \mathcal{X} \times \mathcal{M}$ and $\tau\mathcal{K} : \mathcal{X} \times \mathcal{M} \to \mathcal{X} \times \mathcal{M}$. We now have a reformuation of (1.1) posed entirely in the space $\mathcal{X} \times \mathcal{M}$, and the map τ is therefore analogous to applying a preconditioner to the operator $\mathcal{L}$. Furthermore, we can build a practical preconditioner by choosing an inner product which is sufficiently close to the bilinear forms associated with these spaces, yet is numerically tractable. We refer the reader to the monograph by Malek and Strakos [32], for example, for more detail.

In some sense, therefore, block diagonal preconditioners are natural for saddle point problems of the form (1.1). We refer the reader to the article by Mardal and Winther [33], which describes in detail such preconditioners. However, their derivation requires detailed knowledge of the problem, and may be non-trivial, especially in real-world applications.

It is well known that there is a strong link between operator preconditioning and Schur complement preconditioning. Preconditioners for a range of problems from mixed finite elements [23, 43, 46, 53] were developed as approximations to the primal or dual Schur complement, but can be thought of as a finite dimensional analogue to τ.

4. The relationship between natural norm and Schur complement preconditioners.

Pestana and Wathen [42] describe the link between the dual Schur complement ($BY^{-1}B^T$) and the Riesz map, μ, on the secondary variables for a specific choice of Y, and we give an alternative proof of this result below. We also derive an analogous link between the primal Schur complement ($A + B^T W^{-1} B$) and Riesz map, χ, on the primary variables for a specific choice of W. These relationships are central to the theory that we develop for practical element preconditioners in Section 5.

4.1. Saddle point theory. We assume that the operators $a(\cdot,\cdot)$ and $b(\cdot,\cdot)$ are bounded, satisfying

$$|a(u,v)| \leq \Gamma_a^* \|u\|_{\mathcal{X}} \|v\|_{\mathcal{X}} \quad \forall u,v \in \mathcal{X} \tag{4.1}$$
$$|b(u,p)| \leq \Gamma_b^* \|u\|_{\mathcal{X}} \|p\|_{\mathcal{M}} \quad \forall u \in \mathcal{X}, p \in \mathcal{M}, \tag{4.2}$$

for some positive constants Γ_a^*, Γ_b^*. We define the space

$$\mathcal{V} := \{v \in \mathcal{X} : b(v,p) = 0 \ \forall \ p \in \mathcal{M}\}.$$

Brezzi's splitting theorem [11] tells us that the mapping in (1.1) defines an isomorphism if and only if the bilinear forms satisfy the following conditions:

1. the bilinear form $a(\cdot,\cdot)$ is $\mathcal{V}-$elliptic, i.e.

$$\exists \alpha_* > 0 \text{ s.t. } a(v,v) \geq \alpha_* \|v\|_{\mathcal{X}}^2 \quad \forall v \in \mathcal{V}. \tag{4.3}$$

2. The bilinear form $b(\cdot,\cdot)$ satisfies the inf-sup condition:

$$\exists \beta_* > 0 \text{ s.t. } \inf_{p \in \mathcal{M}} \sup_{u \in \mathcal{X}} \frac{b(u,p)}{\|u\|_{\mathcal{X}} \|p\|_{\mathcal{M}}} \geq \beta_*. \tag{4.4}$$

When we discretize problems of the form (1.1), we need to be careful in the choice of approximation spaces. It is not necessarily true that the finite dimensional problem will satisfy the equivalent inf-sup condition. In particular, we cannot choose the spaces $\mathcal{X}_h$ and $\mathcal{M}_h$ independently; see, for example, Brezzi and Fortin [12, Chapter 2]. If the spaces are complimentary, we say they satisfy the Ladyshenskaja-Babuška-Brezzi (LLB) condition, or that they are inf-sup stable. The discrete analogue to (4.4) is therefore

$$\inf_{p_h \in \mathcal{M}_h} \sup_{u_h \in \mathcal{X}_h} \frac{b(u_h, p_h)}{\|u_h\|_{\mathcal{X}_h} \|p_h\|_{\mathcal{M}_h}} \geq \beta. \tag{4.5}$$

For a stable discretization, the inf-sup constant, β, is independent of the discretization parameter (mesh size).

There is an affinity between the continuous operators and their matrix representations, and we can write the relations above in terms of matrices. In finite dimensions, (4.1) and (4.2) become

$$|\mathbf{u}^T A \mathbf{v}| \leq \Gamma_a \|\mathbf{u}\|_X \|\mathbf{v}\|_X \quad \forall \mathbf{u}, \mathbf{v} \in \mathbb{R}^n \tag{4.6}$$

$$|\mathbf{p}^T B \mathbf{u}| \leq \Gamma_b \|\mathbf{u}\|_X \|\mathbf{p}\|_M \quad \forall \mathbf{u} \in \mathbb{R}^n, \mathbf{p} \in \mathbb{R}^m \tag{4.7}$$

for positive constants Γ_a and Γ_b, where $X_{i,j} = (\phi_i, \phi_j)_{\mathcal{X}_h}$ and $M_{i,j} = (\psi_i, \psi_j)_{\mathcal{M}_h}$ are the matrices that define the natural inner products of the primary and secondary spaces, respectively. Similarly, the matrix representation of the inf-sup condition (4.5) is

$$\sup_{\mathbf{u}} \frac{(B\mathbf{u}, \mathbf{p})}{\|\mathbf{u}\|_X} \geq \beta \|\mathbf{p}\|_M \quad \forall \mathbf{p} \in \mathbb{R}^m. \tag{4.8}$$

Following the theory in Braess [10, Chapter 4], we associate the mapping $\mathbb{B} : \mathcal{X}_h \to \mathcal{M}_h'$ with $(\mathbb{B}u_h, p_h)_{\mathcal{M}_h} = b(u_h, p_h)$ for all $p_h \in \mathcal{M}_h$, and its adjoint mapping $\mathbb{B}' : \mathcal{M}_h \to \mathcal{X}_h'$ with $(\mathbb{B}'q_h, v_h)_{\mathcal{M}_h} = b(v_h, q_h)$ for all $v_h \in \mathcal{X}_h$. The following lemma gives five alternative statements of the inf-sup condition (4.8).

Lemma 4.1. *Let $\mathcal{V}_h$ be the finite dimensional analogue of $\mathcal{V}$. The following statements are equivalent:*

1. *There exits a constant β that satisfies (4.8).*
2. *The operator $\mathbb{B} : \mathcal{V}_h^\perp \to \mathcal{M}_h'$ is an isomorphism, and*

$$\|\mathbb{B}v_h\|_{\mathcal{M}_h} \geq \beta \|v_h\|_{\mathcal{X}_h} \qquad \forall v_h \in \mathcal{V}_h^\perp.$$

3. *For all $\mathbf{v} \in \mathbb{R}^n$, $\mathbf{v} \in \text{null}(B)^\perp$,*

$$\|B\mathbf{v}\|_{M^{-1}} \geq \beta \|\mathbf{v}\|_X ,$$

4. *The operator $\mathbb{B}' : \mathcal{M}_h \to \mathcal{V}_h^0 \subset \mathcal{X}_h'$ is an isomorphism, and*

$$\|\mathbb{B}'p_h\|_{\mathcal{X}_h} \geq \beta \|p_h\|_{\mathcal{M}_h} \qquad \forall p_h \in \mathcal{M}_h.$$

5. *For all $\mathbf{p} \in \mathbb{R}^m$,*

$$\|B^T\mathbf{p}\|_{X^{-1}} \geq \beta \|\mathbf{p}\|_M .$$

Proof. For the equivalence of 1, 2 and 4, see Braess [10, Lemma 4.2].

We now show 2 $\iff$ 3. First note that $\|v_h\|^2_{\mathcal{X}_h} = \mathbf{v}^T X \mathbf{v}$. Since $\mathbb{B}v_h \in \mathcal{M}_h$, $\mathbb{B}v_h = \sum_j \mathbf{q}_j \psi_j$ for some coefficients $\mathbf{q} \neq \mathbf{0}$. Then

$$\|\mathbb{B}v_h\|_{\mathcal{M}_h} = \sum_i \sum_j \mathbf{v}_i \mathbf{q}_j \left(\mathbb{B}\phi_i, \psi_j\right)_{\mathcal{M}_h} = \sum_i \sum_j \mathbf{v}_i \mathbf{q}_j b(\phi_i, \psi_j) = \mathbf{q}^T B \mathbf{v}.$$

Furthermore, we have

$$(B\mathbf{v})_j = \sum_i \mathbf{v}_i (\mathbb{B}\phi_i, \psi_j)_{\mathcal{M}_h} = (\mathbb{B}v_h, \psi_j)_{\mathcal{M}_h} = \sum_k \mathbf{q}_k \left(\psi_k, \psi_j\right)_{\mathcal{M}_h} = (M\mathbf{q})_j,$$

and so $\mathbf{q} = M^{-1}B\mathbf{v}$, which shows the equivalence of 2 and 3.

A similar argument gives 4 $\iff$ 5. ■

We will use these results below to show relationships between a Schur complement and the natural norm for such problems, which will inform our choice of the weighting matrices Y and W.

4.2. Dual Schur complements and the natural norm. For Stokes equations (see Section 6.1), we take the weighting matrix $Y = A$, which can also be identified with the matrix defining the natural norm. Figure 2a shows a plot of the entries of the resulting dual Schur complement, where the magnitude of the entries is shown in the colour bar. The matrix is dense, however the largest entries have a clear structure. Indeed, as we see in Figure 2b, in this case this structure resembles the sparsity pattern of the matrix M, the mass matrix for this example.

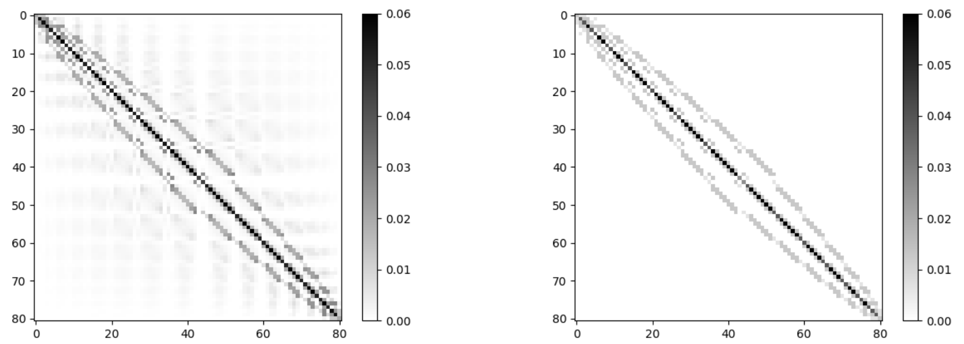

(a) $BA^{-1}B^T$ – the Stokes equations

(b) The P1 mass matrix

Figure 2: Plots of the entries of the Schur complements and mass matrix, P2-P1 discretization, $h = 2^{-3}$, with the Reverse Cuthill-McKee ordering.

The following theorem formalises this intuition for saddle point systems in general. We refer the reader to Pestana and Wathen [42, Section 3] for an alternative proof of this result.

Theorem 4.2. *Suppose we have an inf-sup stable discretization of a saddle point problem of the form* (1.1) *where* $a(\cdot,\cdot)$ *is* $\mathcal{V}_h-$*elliptic. Let* λ *satisfy the generalized eigenvalue problem*

$$BX^{-1}B^T\mathbf{p} = \lambda M\mathbf{p}.$$

Then λ *is independent of* h*, and lies in the range* $[\beta^2, \Gamma_b^2]$.

Proof. Consider the generalized Rayleigh quotient

$$\frac{\mathbf{p}^T BX^{-1}B^T\mathbf{p}}{\mathbf{p}^T M\mathbf{p}}.$$

We can bound this quantity from below by β^2 using Condition 5 of Lemma 4.1, and from above by Γ_b^2 using (4.7) with $\mathbf{u} = X^{-1}B^T\mathbf{p}$. ■

Theorem 4.2 suggests that the natural choice of the weighting matrix Y is the matrix that defines the natural inner product of the primary space, X. We will use this in subsequent sections.

Remark 4.3. *From Theorem 4.2, if* A *can be identified with* X*, then the dual Schur complement of the leading block of* $\mathcal{A}$ *is spectrally equivalent to the matrix* M*.*

4.3. Primal Schur complements and the natural norm.

We are unaware of a result in the literature that links the primal Schur complement to the natural norm in a way analogous to Theorem 4.2. Below we prove results for the cases when A is symmetric positive definite (Theorem 4.4) and maximally rank deficient (Theorem 4.5).

Theorem 4.4. *Suppose that the bilinear form* $a(\cdot,\cdot)$ *is elliptic with ellipticity constant* α_x*, and the associated matrix* A *is symmetric. For an inf-sup stable discretization, the generalized eigenvalues satisfying*

$$(A + B^T M^{-1}B)\mathbf{x} = \lambda X\mathbf{x}$$

are such that $\lambda \in [\alpha_x, \Gamma]$*, where* $\Gamma = \Gamma_a + \Gamma_b^2$.

Proof. For the lower bound, we have

$$\mathbf{x}^T(A + B^T M^{-1}B)\mathbf{x} = \mathbf{x}^T A\mathbf{x} + \mathbf{x}^T B^T M^{-1}B\mathbf{x} \geq \mathbf{x}^T A\mathbf{x} \geq \alpha_x \mathbf{x}^T X\mathbf{x}.$$

For the upper bound, first note that

$$\mathbf{x}^T B^T M^{-1}B\mathbf{x} \leq \Gamma_b\|\mathbf{x}\|_X\|M^{-1}B\mathbf{x}\|_M = \Gamma_b(\mathbf{x}^T X\mathbf{x})^{1/2}(\mathbf{x}^T B^T M^{-1}B\mathbf{x})^{1/2},$$

where we have used (4.7). Therefore using this result, together with (4.6), we obtain

$$\mathbf{x}^T(A + B^T M^{-1}B)\mathbf{x} = \mathbf{x}^T A\mathbf{x} + \mathbf{x}^T B^T M^{-1}B\mathbf{x} \leq \Gamma_a \mathbf{x}^T X\mathbf{x} + \Gamma_b^2 \mathbf{x}^T X\mathbf{x},$$

which gives the required result. ■

Theorem 4.5. *Suppose that the bilinear form is* $\mathcal{V}_h$*-elliptic with ellipticity constant* α*, and the matrix* A *associated with the bilinear form* $a(\cdot,\cdot)$ *is symmetric and positive semi-definite with nullity* $n - m$*. For an inf-sup stable discretization, the generalized eigenvalues satisfying*

$$(A + B^T M^{-1}B)\mathbf{x} = \lambda X\mathbf{x}$$

are such that $\lambda \in [\gamma, \Gamma]$*, where* $\gamma = \frac{1}{2}\min(\alpha, \beta^2)$ *and* $\Gamma = \Gamma_a + \Gamma_b^2$.

Proof. We can decompose $\mathbf{x} = \mathbf{y} + \mathbf{z}$ where $\mathbf{y} \in \text{null}(B)$ and $\mathbf{z} \in \text{null}(B)^{\perp}$. Since A has nullity $n - m$, and $a(\cdot,\cdot)$ is $\mathcal{V}_h$−elliptic, we must have that $A\mathbf{z} = \mathbf{0}$.

We have that

$$\begin{aligned}\mathbf{x}^T(A + B^T M^{-1} B)\mathbf{x} &= \mathbf{y}^T A\mathbf{y} + \mathbf{z}^T B^T M^{-1} B\mathbf{z} \\ &\geq \alpha \mathbf{y}^T X\mathbf{y} + \beta^2 \mathbf{z}^T X\mathbf{z} \\ &\geq \min(\alpha, \beta^2)(\mathbf{y}^T X\mathbf{y} + \mathbf{z}^T X\mathbf{z})\end{aligned}$$

Since $(\mathbf{y} - \mathbf{z})^T X(\mathbf{y} - \mathbf{z}) \geq 0$, we have that $\mathbf{y}^T X\mathbf{y} + \mathbf{z}^T X\mathbf{z} \geq 2\mathbf{y}^T X\mathbf{z}$. Therefore

$$\begin{aligned}\mathbf{x}^T(A + B^T M^{-1} B)\mathbf{x} &\geq \min(\alpha, \beta^2)\left(\frac{1}{2}\left(\mathbf{y}^T X\mathbf{y} + \mathbf{z}^T X\mathbf{z}\right) + \frac{1}{2}\left(\mathbf{y}^T X\mathbf{y} + \mathbf{z}^T X\mathbf{z}\right)\right) \\ &\geq \min(\alpha, \beta^2)\frac{1}{2}\left(\mathbf{y}^T X\mathbf{y} + 2\mathbf{y}^T X\mathbf{z} + \mathbf{z}^T X\mathbf{z}\right) \\ &= \gamma\, \mathbf{x}^T X\mathbf{x},\end{aligned}$$

from which we obtain the lower bound. The upper bound follows the proof of Theorem 4.4. ■

Theorems 4.4 and 4.5 provide the motivation for the "natural" choice for the weighting matrix $W = M$. From this point onwards we consider the *primal* Schur complement to be $A + B^T M^{-1} B$.

Remark 4.6. *If* $\text{null}(B)$ *and* $\text{null}(B)^{\perp}$ *are* X−*orthogonal subspaces, then it is straightforward to adapt the proof of Theorem 4.5 to show that minimum eigenvalue is* $\min(\alpha, \beta^2)$. *This is the case in, for example, Maxwell's equations [23, Section 2.2].*

5. Element Schur complement preconditioners.

We now turn our attention to the element Schur complement approximations defined in (1.4). Computing these matrices requires a dense calculation on each element. Since the matrices involved are small in comparison to the mesh size (e.g., for P2 elements in 2D, A_e will be in $\mathbb{R}^{6\times 6}$) the cost of forming the element Schur complement is asymptotically the same as standard assembly.

Let X_e and M_e be the local Gram matrices defined by the norm on $\mathcal{X}_h$ and $\mathcal{M}_h$, respectively. Then we have the following results.

Lemma 5.1. *Suppose we have a saddle point problem that satisfies the conditions of Theorem 4.2. The generalized eigenvalues satisfying:*

$$B_e X_e^{-1} B_e^T \mathbf{x} = \lambda M_e \mathbf{x}$$

are bounded within some finite region $[\widehat{\gamma}_d^2, \widehat{\Gamma}_d^2]$ *independently of* h.

Lemma 5.2. *Suppose we have a saddle point problem that satisfies the conditions of Theorem 4.4 or Theorem 4.5. The generalized eigenvalues satisfying:*

$$(A_e + B_e^T M_e^{-1} B_e)\mathbf{x} = \lambda X_e \mathbf{x}$$

are bounded within some finite region $[\widehat{\gamma}_p^2, \widehat{\Gamma}_p^2]$ *independently of* h.

Proof of Lemma 5.1 and Lemma 5.2. The discretization is inf-sup stable on any domain, so must be so on a single element. Thus, by Theorem 4.2 (for Lemma 5.1) and Theorem 4.4 or Theorem 4.5 (for Lemma 5.2), we get the required result, with the eigenvalues bounded above and below by the bounds in those theorems. ■

Theorem 5.3. *Suppose we have a discretization of a saddle point problem that satisfies the conditions of Theorem 4.2, and consider $\hat{S}_{\text{dual}}$ as defined in (1.4) with $Y_e = X_e$. The generalized eigenvalues satisfying*

$$BX^{-1}B^T\mathbf{x} = \lambda \hat{S}_{\text{dual}}\mathbf{x}$$

are bounded independently of the mesh size.

Proof. For the upper bound, we have that

$$\begin{aligned}
\lambda &\leq \max_{\mathbf{x}\notin \text{null}(B)} \frac{\mathbf{x}^T BX^{-1}B^T\mathbf{x}}{\mathbf{x}^T\hat{S}_{\text{dual}}\mathbf{x}} \\
&= \max_{\mathbf{x}\notin \text{null}(B)} \frac{\mathbf{x}^T BX^{-1}B^T\mathbf{x}}{\mathbf{x}^T M\mathbf{x}} \cdot \frac{\mathbf{x}^T M\mathbf{x}}{\mathbf{x}^T\hat{S}_{\text{dual}}\mathbf{x}} \\
&\leq \Gamma_b^2 \max_{\mathbf{x}\notin \text{null}(B)} \frac{\mathbf{x}^T N^T M_e N\mathbf{x}}{\mathbf{x}^T N^T \hat{B}_e \hat{X}_e{}^{-1} \hat{B}_e N\mathbf{x}} \\
&\leq \Gamma_b^2 \max_e \max_{\mathbf{z}\notin \text{null}(B_e)} \frac{\mathbf{z}^T M_e\mathbf{z}}{\mathbf{z}^T B_e X_e^{-1} B_e^T\mathbf{z}}, \\
&= \Gamma_b^2/\widehat{\gamma}_d^2.
\end{aligned}$$

An analogous argument holds for the lower bound, which is given by $\beta^2/\widehat{\Gamma}_d^2$. ■

Theorem 5.4. *Suppose we have a discretization of a saddle point problem that satisfies the conditions of either Theorem 4.4 or Theorem 4.5, and consider $\hat{S}_{\text{primal}}$ as defined in (1.4) with $W_e = M_e$. Then the generalized eigenvalues satisfying*

$$(A + BM^{-1}B^T)\mathbf{x} = \lambda \hat{S}_{\text{primal}}\mathbf{x}$$

are bounded away from zero.

Proof. The proof follows a similar format to Theorem 5.3. The eigenvalues are bounded above by $\Gamma/\widehat{\gamma}_p^2$. The lower bound is given by $\alpha_x/\widehat{\Gamma}_p^2$ in the elliptic case (as in Theorem 4.4), and $\gamma/\widehat{\Gamma}_p^2$ in the maximally rank deficient case (as in Theorem 4.5). ■

Theorems 5.3 and 5.4 show spectral equivalence between the element and global Schur complements for both the dual and primal cases. The existence of a natural norm preconditioner is vital in the proofs, but explicit knowledge of it is not required to form a practical preconditioner. We provide numerical experiments in the following section to demonstrate the performance of the element based Schur complement approximation against well-known and state-of-the-art preconditioners on standard test problems.

6. Numerical results. In this section, we present numerical results for both symmetric and nonsymmetric problems, using `Firedrake` [44] with `PETSc` [6, 7] and `PETSc4PY` [13] as the solver interface. We have also used `FEniCS` [1, 31] to generate the eigenvalue plots that depict the theoretical bounds produced in Section 5. For systems involving simple H^1 elliptic type operators, we use an algebraic multigrid solver from `HYPRE` [18], and for more complex operators, arising from $H(\text{curl})$ discretizations, we use the sparse direct solver `MUMPS` [2, 3]. We set the absolute and relative tolerance of the Krylov subspace solver to be 10^{-6} and 10^{-8}, respectively. Table 1 introduces the notation used in the results tables. We note that the timings are for the total time to solve the model, including the assembly of both the linear system and the preconditioners, as well as the linear solve.

Column label	Linear model	Nonlinear model
DoF	Total degrees of freedom (system size)	Total degrees of freedom (system size)
Time	Total time to solve the linear system (including assembly of the linear system and preconditioner, as well as the linear solve time)	Total time to solve the nonlinear system (including assembly of all linear systems and preconditioners, as well as linear solve time at each nonlinear iteration)
Iteration	Total number of linear iterations	Total number of nonlinear iterations/average number of linear iterations

Table 1: Notation used in the results tables

6.1. Stokes Flow. The Stokes equations describe viscous incompressible flow over some bounded, sufficiently regular, domain $\Omega \subset \mathbb{R}^d$:

$$-\frac{1}{\text{Re}}\nabla^2\vec{u} + \nabla p = \vec{f} \qquad \text{in } \Omega \tag{6.1}$$

$$\nabla \cdot \vec{u} = 0 \qquad \text{in } \Omega, \tag{6.2}$$

where $\vec{u}$, p and Re are the fluid velocity, fluid pressure and the Reynolds number, respectively. For a detailed description of the problem, see, e.g., Temam [50, Chapter 1] or Elman, Silvester and Wathen [16, Chapter 3]. We seek a weak solution $(\vec{u}, p) \in H^1_E(\Omega)^d \times L^2(\Omega)$ such that

$$\begin{aligned} \tfrac{1}{\text{Re}}(\nabla\vec{u}, \nabla\vec{v})_\Omega - (p, \nabla\cdot\vec{v})_\Omega &= (\vec{f}, \vec{v})_\Omega, \qquad \forall \vec{v} \in H^1_{E_0}(\Omega)^d, \\ -(q, \nabla\cdot\vec{u})_\Omega &= 0, \qquad \forall q \in L^2(\Omega), \end{aligned}$$

where $H^1_E(\Omega)^d$ and $H^1_{E_0}(\Omega)^d$ are subsets of $H^1(\Omega)^d$ that satisfy the required boundary conditions. We consider the classical test problem known as the *leaky* cavity driven flow, see Elman, Silvester and Wathen [16, Example 3.1.3] for full details.

In the following we use the (inf-sup stable) P2-P1 (Taylor-Hood) element [49], although the methods can be applied to any stable element pair. Upon discretization, we obtain a saddle point system of the form (1.2), where A is a discrete vector Laplacian scaled by the Reynolds number, and B is a fluid divergence operator.

The natural norm preconditioner for the Stokes problem is analogous to the approximate Schur complement preconditioner of Silvester and Wathen [46, 53] and is given by

$$\mathcal{P} = \begin{bmatrix} \frac{1}{\mathrm{Re}}K & 0 \\ 0 & \mathrm{Re}\,Q_p \end{bmatrix},$$

where Q_p is the pressure mass matrix and K is the discrete vector Laplacian; see Mardal and Winther [33, Example 7.1] for full details.

Since the leading block of the Stokes problem defines the natural inner product on the primary space, then from Theorem 5.3 the element and global dual Schur complements are spectrally equivalent. However, in the construction of the element Schur complement, it is necessary to invert the local element matrices of the Laplacian. These local matrices are singular so we shift the matrix to form the element Schur complements as follows:

$$\hat{S}_{\mathrm{dual}} = \sum_e N_e^T B_e \left(\frac{1}{\mathrm{Re}} (K_e + \epsilon Q_e) \right)^{-1} {B_e}^T N_e \quad \text{where} \quad \epsilon = 10^{-6},$$

$$\hat{S}_{\mathrm{primal}} = \frac{1}{\mathrm{Re}} \sum_e L_e^T \left(K_e + {B_e}^T [Q_p]_e^{-1} B_e \right) L_e.$$

Here B_e and K_e are defined as before, Q_e is the velocity element mass matrix and $[Q_p]_e$ is the pressure element mass matrix. The shift in the dual Schur complement corresponds to using a weighting matrix $Y = \frac{1}{\mathrm{Re}}(K + \epsilon Q)$.

We first numerically examine the eigenvalue bounds derived in Theorems 5.3 and 5.4. Figure 3 depicts the computed (blue) eigenvalues of

$$\mathrm{Re}\, B K^{-1} B^T \mathbf{x} = \lambda \hat{S}_{\mathrm{dual}} \mathbf{x} \quad \text{and} \quad \frac{1}{\mathrm{Re}} (K + B^T Q_p^{-1} B) \mathbf{x} = \lambda \hat{S}_{\mathrm{primal}} \mathbf{x},$$

where the pressure mass matrix, Q_p, is the matrix associated with the norm on the pressure space, together with the bounds. The computed bounds are shown with square and regular parentheses for the primal and dual Schur complement, respectively. Finally, the y-axis shows the order of the mixed discretization used, see *Periodic Table of the Finite Elements* [4] for the key.

From the figure, we see that the computed eigenvalues for both Schur complements remain constant, and within the bounds, as the order of the discretization is increased. The bounds are not tight in this example, but are similar in value to equivalent bounds obtained for the natural norm preconditioners (see, for example, Elman, Silvester and Wathen [16, Section 3.5.1]), which our estimates rely on. We observe that the both the computed eigenvalues and the bounds are well away from zero. The worst case in Theorems 5.3 and 5.4 would give a small condition number, and in practice we see the eigenvalues are much more tightly clustered than the worst case predicts.

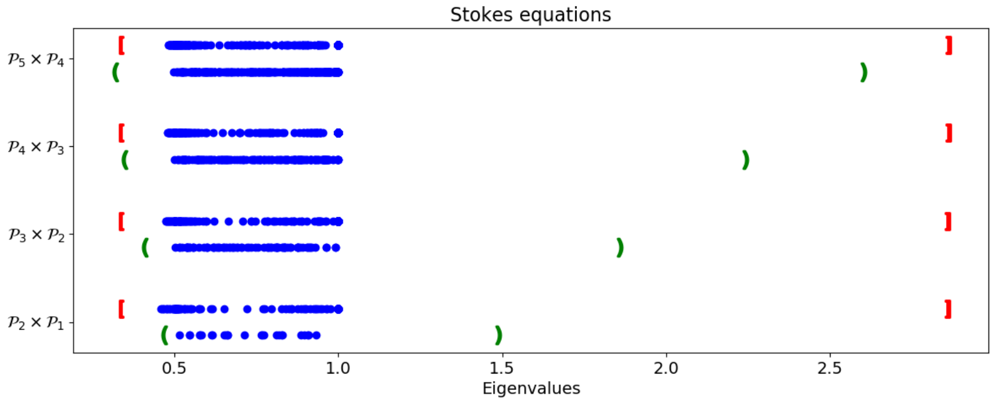

Figure 3: Eigenvalue bounds for different discretization orders for the Taylor-hood mixed element for the Stokes example in Section 6.1.

Tables 2 and 3 give the two and three dimensional timing and iteration results for a constant Reynolds number, Re = 1000. For the two dimensional results in Table 2, we see that the iterations for both element Schur complement preconditioners and the natural norm preconditioner do not grow significantly as we refine the mesh, and the performance of all preconditioners is comparable. Also, for all preconditioners, the time scales linearly with the degrees of freedom. For the three dimensional results in Table 3, we see that the iterations for the dual element Schur complement appear to increase mildy, however, the iterations for the natural norm and the primal element Schur complement preconditioner are almost identical and are scalable with respect to the mesh.

	Dual element Schur complement		Primal element Schur complement		Natural norm preconditioner	
DoF	Time	Iteration	Time	Iteration	Time	Iteration
2,467	0.2	45	0.2	40	0.2	38
9,539	0.4	43	0.3	42	0.3	41
37,507	0.8	45	1.1	45	0.7	41
148,739	3.0	50	5.6	45	3.6	43
592,387	12.1	52	18.3	48	13.0	51
2,364,419	53.8	58	56.4	45	60.2	55
9,447,427	197.9	50	239.0	51	209.3	53

Table 2: Iteration and timing results for the two-dimensional Stokes *leaky* cavity driven flow problem with Re = 1000.

	Dual element Schur complement		Primal element Schur complement		Natural norm preconditioner	
DoF	Time	Iteration	Time	Iteration	Time	Iteration
402	0.2	45	0.2	52	0.2	41
2,312	0.3	60	0.3	56	0.2	54
15,468	1.8	69	2.5	62	1.8	61
112,724	20.5	75	27.3	65	18.1	66
859,812	219.2	83	279.6	69	186.1	68

Table 3: Iteration and timing results for the three-dimensional Stokes *leaky* cavity driven flow problem with $\mathrm{Re} = 1000$.

6.2. Mixed formulation of Maxwell's equations. We now consider the time-harmonic Maxwell equation in mixed form [23, 24, 34]. The continuous problem is given as follows:

$$\begin{aligned} \frac{1}{\mathrm{Re}_m}\nabla \times \nabla \times \mathbf{b} + \nabla r &= \mathbf{f} \quad \text{in } \Omega \\ \nabla \cdot \mathbf{b} &= 0 \quad \text{in } \Omega, \end{aligned} \tag{6.3}$$

where $\mathbf{b}$ is the magnetic field, r is the Lagrange multiplier associated with the divergence constraint on the magnetic field and Re_m is the magnetic Reynolds number. For this problem, we set the forcing terms and Dirichlet boundary conditions corresponding to the exact solution

$$\vec{u} = \begin{bmatrix} \exp(x)\cos(y) \\ \exp(x)\sin(y) \end{bmatrix} \quad \text{and} \quad p = xy.$$

The standard weak form is: find $(\vec{b}, r) \in H(\mathrm{curl}, \Omega) \times H^1_E(\Omega)$ such that

$$\begin{aligned} \frac{1}{\mathrm{Re}_m}(\nabla \times \vec{b}, \nabla \times \vec{c})_\Omega + (\vec{c}, \nabla r)_\Omega &= (\mathbf{f}, \vec{c})_\Omega, \qquad \forall \vec{c} \in H(\mathrm{curl}, \Omega), \\ (\vec{b}, \nabla s)_\Omega &= 0, \qquad \forall s \in H^1_{E_0}(\Omega). \end{aligned}$$

In the mixed Maxwell case, A in (1.2) is the discrete curl-curl operator, $\frac{1}{\mathrm{Re}_m}K$.

The primal element Schur complement is formed using the matrix associated with the natural inner product for $H^1_0(\Omega)$, the Laplacian operator, L, and so we take this to be the weighting matrix, W. The natural norm preconditioner is $\mathtt{blkdiag}\left(\frac{1}{\mathrm{Re}_m}K + Q, L\right)$, where Q is a vector mass matrix; see Mardal and Winther [33, Example 7.4]. Greif and Schötzau [23] derive the same preconditioner using a Schur complement argument. They show directly that $B^T L^{-1} B$ is spectrally equivalent to the vector mass matrix, Q. Thus, their proposed preconditioner is identical (under mild parameter assumptions) to the natural norm preconditioner. We also consider the dual element Schur complement preconditioner, taking Y as $\frac{1}{\mathrm{Re}_m}K + Q$.

Figure 4 shows the eigenvalue distribution, following the same conventions introduced for Figure 3. The computed eigenvalues and corresponding bounds for the dual Schur complement are precisely 1, since the dual element and global Schur complement are identical; this can be seen from identities established by Greif and collaborators [17, 23, 24]. We see that the

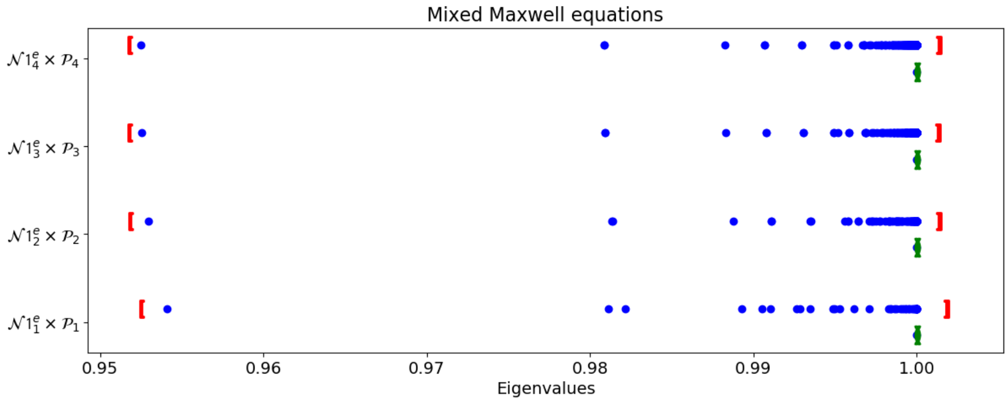

Figure 4

bounds on the primal element Schur complement are looser, however, the bounds seem to be fairly descriptive of the eigenvalues and range between 0.95 and just over 1.

Table 4 gives the timing and iteration results for this example. In general, the number of iterations for all three preconditioning techniques do not grow significantly as the system gets larger. The timings do not scale linearly in this case, which is due to our use of a direct solver to apply the preconditioner. The iteration counts for the dual element Schur complement and the natural norm preconditioner are identical. Since we are applying the preconditioner with a direct solver we do not provide three dimensional results for this example.

	Dual element Schur complement		Primal element Schur complement		Natural norm preconditioner	
DoF	Time	Iteration	Time	Iteration	Time	Iteration
4,225	0.4	33	0.4	33	0.4	33
16,641	0.6	30	0.7	30	0.6	30
66,049	1.8	33	2.9	33	1.9	33
263,169	6.6	33	9.4	33	8.4	33
1,050,625	27.4	33	31.0	33	30.2	33
4,198,401	113.9	34	116.3	35	121.9	34
16,785,409	596.9	35	638.3	40	590.8	35

Table 4: Timing and iterations results for the mixed Maxwell equations with $\mathrm{Re}_m = 100$.

6.3. The Navier-Stokes equations. Finally, we consider the Navier-Stokes equations, given by

$$-\frac{1}{\mathrm{Re}}\nabla^2\vec{u} + \vec{u}\cdot\nabla\vec{u} + \nabla p = \vec{f}$$
$$\nabla\cdot\vec{u} = 0.$$

The presence of the advection term $(\vec{u} \cdot \nabla \vec{u})$ provides both non-linearity and non-symmetry in the model, which are two key differences between this example and the previous ones. The nonlinear scheme we use is Newton's method with absolute and relative tolerances of 10^{-6} and 10^{-8}, respectively.

A standard mixed formulation of this model is: find $(\mathbf{u}, p) \in H_0^1(\Omega)^d \times L^2(\Omega)$ such that

$$\begin{array}{rcll} \frac{1}{\mathrm{Re}}(\nabla \vec{u}, \nabla \vec{v})_\Omega + (\vec{u} \cdot \nabla \vec{u}, \vec{v})_\Omega + (p, \nabla \cdot \vec{v})_\Omega & = & (\vec{f}, \vec{v})_\Omega, & \forall \vec{v} \in H_0^1(\Omega)^d, \\ (q, \nabla \cdot \vec{u})_\Omega & = & 0, & \forall q \in L^2(\Omega). \end{array}.$$

Upon discretization and linearization, we obtain a saddle point system of the form (1.2) with $A = \frac{1}{\mathrm{Re}} K + C$, where K is the vector Laplacian and C is the advection diffusion matrix, and B is the divergence matrix.

Due to their non-symmetric nature, the discrete Navier Stokes equations do not fit into the framework for natural norm preconditioners, thus we cannot apply the theory in Sections 4 and 5. However, we can still form local Schur complements using $Y = \frac{1}{\mathrm{Re}} K + C$, and use this as a heuristic preconditioner. We compare this approach with the Pressure Convection Diffusion (PCD) preconditioner of Kay, Loghin and Wathen [29].

For this example we will use the full $L^T DL$ decomposition:

$$P^{-1} = \begin{bmatrix} I & -\hat{A}^{-1} B^T \\ 0 & I \end{bmatrix} \begin{bmatrix} \hat{A}^{-1} & 0 \\ 0 & \hat{S}^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -B\hat{A}^{-1} & I \end{bmatrix}$$

where the action of $\hat{A}^{-1}$ is performed by a multigrid V-cycle approximating the inverse of A. For $\hat{S}$ we either use the dual element Schur complement or the PCD approximation, $\hat{S} = A_p F_p^{-1} Q_p$ (where A_p, F_p and Q_p are the pressure space Laplacian, convection-diffusion operator and mass matrix, respectively). We use the implementation of PCD distributed with `Firedrake`.

Tables 5 and 7 show timing and iteration results for a two-dimensional cavity driven flow (described in Elman, Silvester and Wathen [16, Example 8.1.2]) and flow over a two-dimensional backwards facing step (described in Elman, Silvester and Wathen [16, Example 8.1.3]), respectively. The backwards facing step is more complicated due to both the nonconvex domain and the natural outflow conditions on the rightmost boundary. Table 6 shows timing and iterations results for a three-dimensional cavity driven flow problem.

From Table 5, we see that for both PCD and dual element Schur complement preconditioners exhibit mesh independent iterations. However, for both Reynolds numbers the dual element Schur complement preconditioner converges in noticeably fewer Krylov iterations. For the timing results, we see that both preconditioners scale linearly with the system size. We highlight that the time to solve the system with the element Schur complement preconditioner is about 75% of the equivalent solve for the PCD preconditioner. As well as needing to take fewer iterations, the element Schur complement requires only one sparse system solve per application, as opposed to two solves and a matrix-vector multiply for the PCD preconditioner.

We can also apply the element Schur complement preconditioner in three dimensions, as shown in Table 6. Here the difference between the element Schur complement preconditioner and the PCD preconditioner is more pronounced. The iterations for the element preconditioner

	Re = 1				Re = 100			
	Dual element Schur complement		PCD preconditioner		Dual element Schur complement		PCD preconditioner	
DoF	Time	Iteration	Time	Iteration	Time	Iteration	Time	Iteration
2,467	0.3	2/24.0	0.4	2/34.5	0.7	4/59.8	1.0	4/99.5
9,539	0.5	2/26.0	0.6	2/36.0	1.6	4/63.2	2.6	4/104.5
37,507	2.1	2/27.0	4.0	2/38.5	8.8	4/65.0	14.1	4/103.0
148,739	7.5	2/27.5	14.0	2/39.0	33.7	4/68.2	49.7	4/99.2
592,387	26.6	2/28.0	39.5	2/40.5	128.1	4/73.8	173.0	4/99.5
2,364,419	108.2	2/29.0	188.1	3/33.3	520.6	4/77.5	741.9	4/107.5
9,447,427	439.9	2/29.0	676.6	2/41.5	1651.1	3/79.0	2162.2	3/103.3

Table 5: Results for preconditioning GMRES for the Navier-Stokes equations with Re = 1 and Re = 100 for the 2D cavity driven flow problem.

	Dual element Schur complement		PCD preconditioner	
DoF	Time	Iteration	Time	Iteration
402	-	-/-	3.0	5/103.0
2,312	1.7	4/80.8	4.3	4/219.2
15,468	18.4	4/63.8	51.2	4/190.2
112,724	170.7	4/71.5	488.4	4/216.2
859,812	1810.9	4/82.5	4986.4	4/238.2

Table 6: Results for preconditioning GMRES for the Navier-Stokes equations with Re = 100 for the 3D cavity driven flow problem.

are similar to those for the two-dimensional problem in Table 5, whereas the linear iteration count for the PCD preconditioner approximately doubled.

For the backwards facing step, Table 7, the `Firedrake` implementation of the PCD preconditioner does not converge, and so we only report the element Schur complement results. While the iterations grow with the mesh size, this growth appears to be stabilizing. Again, the timings scale linearly with the problem size.

We remark that by carefully considering the boundary conditions, the PCD preconditioner can be modified to yield scalable iterations for outflow problems (see, for example, Elman, Silvester and Wathen [16, Section 9.2.2] and Bootland [9, Chapter 5]), but these are not yet implemented in `Firedrake`. We speculate that similar considerations would also improve the performance of the element Schur complement approach considered here; this is beyond the scope of this paper.

The element Schur complement preconditioner was robust with respect to different mesh tessellations (both triangles and quadrilaterals) in our tests. This is also the case for the PCD preconditioner. However, for other preconditioners for Navier-Stokes, specifically the Least-

DoF	Dual element Schur complement Time	Iteration
24,498	7.1	3/79.3
61,659	20.1	3/89.0
133,092	39.2	3/87.3
271,286	93.8	3/103.3
599,892	234.4	3/117.3
2,455,263	972.5	3/119.0

Table 7: Results for preconditioning GMRES for the Navier-Stokes equations with Re = 10 for the backwards facing step problem.

squares commutator preconditioner [15], it has been observed that the performance degrades for triangular cells; see Bootland [9, Chapter 6.3.1] and the references within.

Finally, in our numerical experiments for the Navier-Stokes equations we saw that, for an initial guess that interpolates the Dirichlet boundary conditions, reassembling the element Schur complement preconditioner at each nonlinear iteration did not significantly reduce the number of Krylov iterations. However, it seemed necessary to reassemble the PCD preconditioner to obtain a scalable solver.

7. Conclusion and outlook. In this work, we presented effective sparse Schur complement approximations for inf-sup stable mixed finite element discretizations of self-adjoint problems. Our new preconditioner is based on forming local, or element, Schur complements and projecting them back to the global degrees of freedom. Utilizing the existence of matrices associated with the natural inner products of the underlying function spaces, we show spectral equivalence between global and local Schur complements under mild conditions. However, given the algebraic nature of these preconditioners, they can also be used as a “black-box”.

From the numerical results, the element based Schur complement preconditioners perform similarly to the natural norm preconditioners for the self-adjoint multi-physics problems we have considered. Interestingly, for the nonsymmetric and nonlinear Navier-Stokes example, where the theory breaks down, our new Schur complement approximation outperforms the well-known PCD preconditioner in our tests. One possible area of future work would be to extend the theory to cover nonsymmetric systems.

We only consider single element contributions, but grouping multiple local elements together to form slightly larger local Schur complements may lead to a better approximation, requiring only marginally more storage.

We focus on problems from mixed finite elements, but the ideas should work for any partially separable functions which can be represented as a sum of element functions (see Daydé, L’Excellent and Gould [14]). This may allow us use similar ideas to build fast solvers for problems that do not come from PDEs, and where the natural norms are not obvious. The extension of element preconditioners to these problems is left for future work.

Code availability. Together with the written manuscript, we provide the code which was used to generate the results [55].

References.

[1] M. S. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells, *The FEniCS Project Version 1.5*, Archive of Numerical Software, 3 (2015), https://doi.org/10.11588/ans.2015.100.20553.

[2] P. R. Amestoy, I. S. Duff, J. Koster, and J.-Y. L'Excellent, *A Fully Asynchronous Multifrontal Solver Using Distributed Dynamic Scheduling*, SIAM Journal on Matrix Analysis and Applications, 23 (2001), pp. 15–41, https://doi.org/10.1137/S0895479899358194.

[3] P. R. Amestoy, A. Guermouche, J.-Y. L'Excellent, and S. Pralet, *Hybrid scheduling for the parallel solution of linear systems*, Parallel Computing, 32 (2006), pp. 136 – 156, https://doi.org/https://doi.org/10.1016/j.parco.2005.07.004.

[4] D. N. Arnold and A. Logg, *Periodic table of the finite elements*, SIAM News, 47 (2014), p. 212, http://femtable.org/.

[5] O. Axelsson, R. Blaheta, and M. Neytcheva, *Preconditioning of boundary value problems using elementwise schur complements*, SIAM Journal on Matrix Analysis and Applications, 31 (2009), pp. 767–789, https://doi.org/https://doi.org/10.1137/070679673.

[6] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, and H. Zhang, *PETSc users manual*, Tech. Report ANL-95/11 - Revision 3.13, Argonne National Laboratory, 2020, https://www.mcs.anl.gov/petsc.

[7] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, *Efficient management of parallelism in object oriented numerical software libraries*, in Modern Software Tools in Scientific Computing, E. Arge, A. M. Bruaset, and H. P. Langtangen, eds., Birkhäuser Press, 1997, pp. 163–202, https://doi.org/https://doi.org/10.1007/978-1-4612-1986-6_8.

[8] M. Benzi, G. H. Golub, and J. Liesen, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), p. 1137, https://doi.org/10.1017/S0962492904000212.

[9] N. Bootland, *Scalable two-phase flow solvers*, PhD thesis, University of Oxford, 2018, https://ora.ox.ac.uk/objects/uuid:d315682a-4308-4b2a-8deb-0b2f7cb4bfc3.

[10] D. Braess, *Finite Elements: Theory, Fast Solvers, and Applications in Solid Mechanics*, Cambridge University Press, 3rd ed., 2007, https://doi.org/10.1017/CBO9780511618635.

[11] F. Brezzi, *On the existence, uniqueness and approximation of saddle-point problems arising from lagrangian multipliers*, R.A.I.R.O. Analyse Numérique, 8 (1974), pp. 129–151, http://www.numdam.org/item/M2AN_1974__8_2_129_0.

[12] F. Brezzi and M. Fortin, *Mixed and Hybrid Finite Element Methods*, Springer-Verlag, Berlin, Heidelberg, 1991, https://doi.org/10.1007/978-1-4612-3172-1.

[13] L. D. Dalcin, R. R. Paz, P. A. Kler, and A. Cosimo, *Parallel distributed computing using Python*, Advances in Water Resources, 34 (2011), pp. 1124–1139, https://doi.org/http://dx.doi.org/10.1016/j.advwatres.2011.04.013. New Computational Methods and Software Tools.

[14] M. J. Daydé, J.-Y. L'Excellent, and N. I. M. Gould, *Element-by-Element Pre-*

conditioners for Large Partially Separable Optimization Problems, SIAM J. Sci. Comput., 18 (1997), p. 1767?1787, https://doi.org/10.1137/S1064827594274796.

[15] 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.

[16] H. Elman, D. Silvester, and A. Wathen, *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, Numerical mathematics and scientific computation, Oxford University Press, 2014, https://doi.org/10.1093/acprof:oso/9780199678792.001.0001.

[17] R. Estrin and C. Greif, *On nonsingular saddle-point systems with a maximally rank-deficient leading block*, SIAM Journal on Matrix Analysis and Applications, 36 (2015), pp. 367–384, https://doi.org/10.1137/140989996.

[18] R. D. Falgout and U. M. Yang, *hypre: A Library of High Performance Preconditioners*, in Computational Science — ICCS 2002, P. M. A. Sloot, A. G. Hoekstra, C. J. K. Tan, and J. J. Dongarra, eds., Berlin, Heidelberg, 2002, Springer Berlin Heidelberg, pp. 632–641, https://doi.org/10.1007/3-540-47789-6_66.

[19] 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.

[20] 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), p. A3073?A3096, https://doi.org/10.1137/18m1219370.

[21] A. Greenbaum, V. Pták, and Z. Strakoš, *Any nonincreasing convergence curve is possible for GMRES*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 465–469, https://doi.org/10.1137/S0895479894275030.

[22] C. Greif and D. Schötzau, *Preconditioners for saddle point linear systems with highly singular (1, 1) blocks*, ETNA, Special Volume on Saddle Point Problems, 22 (2006), pp. 114–121, http://etna.mcs.kent.edu/volumes/2001-2010/vol22/abstract.php?vol=22&pages=114-121.

[23] C. Greif and D. Schötzau, *Preconditioners for the discretized time-harmonic maxwell equations in mixed form*, Numerical Linear Algebra with Applications, 14 (2007), pp. 281–297, https://doi.org/10.1002/nla.515.

[24] C. Greif and M. Wathen, *Conjugate gradient for nonsingular saddle-point systems with a maximally rank-deficient leading block*, Journal of Computational and Applied Mathematics, 358 (2019), pp. 1 – 11, https://doi.org/10.1016/j.cam.2019.02.016.

[25] I. Gustafsson and G. Lindskog, *A preconditioning technique based on element matrix factorizations*, Computer methods in applied mechanics and engineering, 55 (1986), pp. 201–220, https://doi.org/10.1016/0045-7825(86)90053-8.

[26] M. R. Hestenes and E. Stiefel, *Methods of conjugate gradients for solving linear systems*, J. Res. Nat. Bur. Stand., 49 (1952), pp. 409–436, https://doi.org/10.6028/jres.049.044.

[27] T. J. Hughes, I. Levit, and J. Winget, *An element-by-element solution algorithm for problems of structural and solid mechanics*, Computer Methods in Applied Mechanics

and Engineering, 36 (1983), pp. 241–254, https://doi.org/10.1016/0045-7825(83)90115-9.

[28] I. C. F. Ipsen, *A note on preconditioning nonsymmetric matrices*, SIAM Journal on Scientific Computing, 23 (2001), pp. 1050–1051, https://doi.org/10.1137/S1064827599355153.

[29] 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/doi={10.1137/S106482759935808X}.

[30] J. Kraus, *Algebraic multilevel preconditioning of finite element matrices using local schur complements*, Numerical linear algebra with applications, 13 (2006), pp. 49–70, https://doi.org/10.1002/nla.462.

[31] A. Logg, K.-A. Mardal, G. N. Wells, et al., *Automated Solution of Differential Equations by the Finite Element Method*, Springer, 2012, https://doi.org/10.1007/978-3-642-23099-8.

[32] J. Màlek and Z. Strakos, *Preconditioning and the Conjugate Gradient Method in the Context of Solving PDEs*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2014, https://my.siam.org/Store/Product/viewproduct/?ProductId=25965149.

[33] K.-A. Mardal and R. Winther, *Preconditioning discretizations of systems of partial differential equations*, Numerical Linear Algebra with Applications, 18 (2011), pp. 1–40, https://doi.org/10.1002/nla.716.

[34] P. Monk, *Finite Element Methods for Maxwell's Equations*, Numerical Mathematics and Scientific Computation, Clarendon Press, 2003, https://doi.org/10.1093/acprof:oso/9780198508885.001.0001.

[35] M. F. Murphy, G. H. Golub, and A. J. Wathen, *A note on preconditioning for indefinite linear systems*, SIAM J. Sci. Comput., 21 (2000), pp. 1969–1972, https://doi.org/10.1137/S1064827599355153.

[36] M. Neytcheva, *On element-by-element schur complement approximations*, Linear Algebra and Its Applications, 434 (2011), pp. 2308–2324, https://doi.org/10.1016/j.laa.2010.03.031.

[37] M. Neytcheva and E. Bängtsson, *Preconditioning of nonsymmetric saddle point systems as arising in modelling of viscoelastic problems*, Electronic Transactions on Numerical Analysis, 29 (2008), pp. 193–211, http://eudml.org/doc/130827.

[38] M. Neytcheva, M. Do-Quang, and H. Xin, *Element-by-element schur complement approximations for general nonsymmetric matrices of two-by-two block form*, in International Conference on Large-Scale Scientific Computing, Springer, 2009, pp. 108–115, https://doi.org/10.1007/978-3-642-12535-5_11.

[39] B. Nour-Omid and B. Parlett, *Element preconditioning using splitting techniques*, SIAM journal on scientific and statistical computing, 6 (1985), pp. 761–770, https://doi.org/10.1137/0906051.

[40] M. Ortiz, P. M. Pinsky, and R. L. Taylor, *Unconditionally stable element-by-element algorithms for dynamic problems*, Computer Methods in Applied Mechanics and Engineering, 36 (1983), pp. 223–239, https://doi.org/10.1016/0045-7825(83)90114-7.

[41] C. C. Paige and M. A. Saunders, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629, https://doi.org/10.1137/0712047,

http://link.aip.org/link/?SNA/12/617/1.
[42] J. Pestana and A. J. Wathen, *Natural preconditioning and iterative methods for saddle point systems*, SIAM Review, 57 (2015), pp. 71–91, https://doi.org/10.1137/130934921.
[43] C. E. Powell and D. Silvester, *Optimal preconditioning for raviart–thomas mixed formulation of second-order elliptic problems*, SIAM journal on matrix analysis and applications, 25 (2003), pp. 718–738, https://doi.org/10.1137/S0895479802404428.
[44] 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 Trans. Math. Softw., 43 (2016), pp. 24:1–24:27, https://doi.org/10.1145/2998441, https://arxiv.org/abs/1501.01809.
[45] Y. Saad and M. H. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869, https://doi.org/10.1137/0907058.
[46] D. Silvester and A. Wathen, *Fast iterative solution of stabilised Stokes systems Part II: Using general block preconditioners*, SIAM J. Numer. Anal., 31 (1994), pp. 1352–1367, https://doi.org/10.1137/0731070.
[47] N. Spillane, V. Dolean, P. Hauret, F. Nataf, C. Pechstein, and R. Scheichl, *A robust two-level domain decomposition preconditioner for systems of PDEs*, Comptes Rendus Mathematique, 349 (2011), pp. 1255 – 1259, https://doi.org/10.1016/j.crma.2011.10.021.
[48] N. Spillane and D. Rixen, *Automatic spectral coarse spaces for robust finite element tearing and interconnecting and balanced domain decomposition algorithms*, International Journal for Numerical Methods in Engineering, 95 (2013), pp. 953–990, https://doi.org/10.1002/nme.4534.
[49] C. Taylor and P. Hood, *A numerical solution of the navier-stokes equations using the finite element technique*, Computers & Fluids, 1 (1973), pp. 73 – 100, https://doi.org/10.1016/0045-7930(73)90027-3.
[50] R. Temam, *Navier–Stokes Equations*, American Mathematical Soc., 1984, https://www.elsevier.com/books/navier-stokes-equations/temam/978-0-444-85307-3.
[51] M. Van Gijzen, *An analysis of element-by-element preconditioners for nonsymmetric problems*, Computer methods in applied mechanics and engineering, 105 (1993), pp. 23–40, https://doi.org/10.1016/0045-7825(93)90114-D.
[52] A. Wathen, *An analysis of some element-by-element techniques*, Computer Methods in Applied Mechanics and Engineering, 74 (1989), pp. 271–287, https://doi.org/10.1016/0045-7825(89)90052-2.
[53] A. Wathen and D. Silvester, *Fast iterative solution of stabilised Stokes systems. part I: Using simple diagonal preconditioners*, SIAM J. Numer. Anal., 30 (1993), pp. 630–649, https://doi.org/10.1137/0730031.
[54] A. J. Wathen, *Preconditioning*, Acta Numerica, 24 (2015), pp. 329–376, https://doi.org/10.1017/S0962492915000021.
[55] M. Wathen and T. Rees, *ralna/elementschur: Version for tech report*, May 2020, https://doi.org/10.5281/zenodo.3801732.