

LEAST SQUARES PRECONDITIONERS FOR STABILIZED DISCRETIZATIONS OF THE NAVIER–STOKES EQUATIONS*

HOWARD ELMAN[†], VICTORIA E. HOWLE[‡], JOHN SHADID[§], DAVID SILVESTER[¶],
AND RAY TUMINARO[‡]

Abstract. This paper introduces two stabilization schemes for the least squares commutator (LSC) preconditioner developed by Elman, Howle, Shadid, Shuttleworth, and Tuminaro [*SIAM J. Sci. Comput.*, 27 (2006), pp. 1651–1668] for the incompressible Navier–Stokes equations. This preconditioning methodology is one of several choices that are effective for Navier–Stokes equations, and it has the advantage of being defined from strictly algebraic considerations. It has previously been limited in its applicability to div-stable discretizations of the Navier–Stokes equations. This paper shows how to extend the same methodology to stabilized low-order mixed finite element approximation methods.

Key words. preconditioning, Navier–Stokes, iterative algorithms

AMS subject classifications. Primary, 65F10, 65N12, 65N30, 76D05; Secondary, 15A06, 35Q30

DOI. 10.1137/060655742

1. Introduction. Consider the Navier–Stokes equations

$$(1.1) \quad \begin{aligned} \eta \mathbf{u}_t - \nu \nabla^2 \mathbf{u} + (\mathbf{u} \cdot \text{grad}) \mathbf{u} + \text{grad } p &= \mathbf{f}, \\ -\text{div } \mathbf{u} &= 0, \end{aligned}$$

on $\Omega \subset \mathbb{R}^d$, $d = 2$ or 3 . Here, \mathbf{u} is the d -dimensional velocity field, which is assumed to satisfy suitable boundary conditions on $\partial\Omega$; p is the pressure; and ν is the kinematic viscosity, which is inversely proportional to the Reynolds number. The value $\eta = 0$ corresponds to the steady-state problem and $\eta = 1$ to the case of unsteady flow. Linearization and discretization of (1.1) by finite elements, finite differences, or finite volumes leads to a sequence of linear systems of equations of the form

$$(1.2) \quad \begin{bmatrix} F & B^T \\ B & -\frac{1}{\nu}C \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ g \end{bmatrix}.$$

*Received by the editors March 30, 2006; accepted for publication (in revised form) May 7, 2007; published electronically December 21, 2007. This work was partially supported by the DOE Office of Science MICS Program and by the ASC Program at Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000. The U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes. Copyright is owned by SIAM to the extent not limited by these rights.

<http://www.siam.org/journals/sisc/30-1/65574.html>

[†]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 (elman@cs.umd.edu). The work of this author was supported by the Department of Energy under grant DOEG0204ER25619.

[‡]Sandia National Laboratories, P.O. Box 969, MS 9159, Livermore, CA 94551 (vhowle@sandia.gov, rstumin@sandia.gov).

[§]Sandia National Laboratories, P.O. Box 5800, MS 1111, Albuquerque, NM 87185 (jnshadi@cs.sandia.gov).

[¶]School of Mathematics, University of Manchester, Manchester M601QD, UK (d.silvester@manchester.ac.uk). The work of this author was supported in part by the Engineering and Physical Sciences Research Council (EPSRC) via grant EP/C000528/1.

These systems, which are the focus of this paper, must be solved at each step of a nonlinear (Picard or Newton) iteration, or at each time step. Here, B and B^T are matrices corresponding to discrete divergence and gradient operators, respectively, and F operates on the discrete velocity space. For *div-stable* discretizations, $C = 0$. For mixed approximation methods that do not uniformly satisfy a discrete inf-sup condition, the matrix C is a nonzero *stabilization operator*. Examples of finite element methods that require stabilization are the mixed approximations using linear or bilinear velocities (trilinear in three dimensions) coupled with constant pressures, as well as any discretization in which equal-order discrete velocities and pressures are specified using a common set of nodes; see, for example, Brezzi and Fortin [4, p. 210].

In recent years, there has been considerable activity in the development of efficient iterative methods for the numerical solution of the stationary and fully implicit versions of this problem. These are based on new preconditioning methods derived from the structure of the linearized discrete problem given in (1.2). A complete overview of the ideas under consideration can be found in the monograph of Elman, Silvester, and Wathen [7]. The key to attaining fast convergence lies with the effective approximation of the Schur complement operator

$$(1.3) \quad S = BF^{-1}B^T + \frac{1}{\nu}C,$$

which is obtained by algebraically eliminating the velocities from the system.

One approach of interest is the *pressure convection-diffusion* preconditioner proposed by Kay, Loghin, and Wathen [10] and Silvester, Elman, Kay, and Wathen [13]. In this method, the Schur complement matrix is approximated as

$$(1.4) \quad S \approx M_S = A_p F_p^{-1} Q_p,$$

where Q_p is the pressure mass matrix associated with the pressure discretization (or a spectrally equivalent approximation to it), and A_p and F_p are discrete Laplace and convection-diffusion operators defined on the pressure space. Note that when we need to make the distinction, we will use a subscript p to indicate operators defined on the pressure space and a subscript v to indicate operators defined on the velocity subspace. Although effective for solving the system (1.2), this method has the drawback of requiring users to provide the discrete operators A_p and F_p . This means that integration of this idea into a code that models incompressible flow requires a sophisticated understanding of the discretization and other implementation issues, something often held only by the developers of the model.

An alternative approach is the *least squares commutator (LSC)* preconditioner developed by Elman, Howle, Shadid, Shuttleworth, and Tuminaro [6]. A description and derivation of this method will be given in section 2. In this case, the approximation to the Schur complement matrix is derived from purely algebraic considerations by solving a certain least squares problem. The resulting preconditioning methodology is competitive with the pressure convection-diffusion approach, and in some cases its performance is superior. However, so far this approach has only been shown to be applicable to the case where $C = 0$ in (1.2). The main goal of this paper is to show that the least squares commutator preconditioner can be extended to handle discretizations that require stabilization. This closes a gap in the derivation of these ideas, and a version of the new method can be also formulated from algebraic considerations.

An outline of the paper is as follows. Section 2 contains a derivation of the LSC approximation to the Schur complement operator for div-stable discretizations, and

section 3 gives a brief outline of stabilization of mixed finite element approximation for the Navier–Stokes equations. Section 4 presents the main contribution of the paper, the derivation of new LSC preconditioners applicable to stabilized approximation methods. Two new methods are derived. One approach is derived by appealing directly to the issue of stabilization, in which the LSC preconditioner is modified by adding stabilizing terms to certain unstable components. The other method is derived by treating “high frequency” and “low frequency” components of the pressure space separately in a manner reminiscent of the development of multigrid methods; loosely speaking, the high frequency terms are those that require stabilization. Finally, section 5 gives the results of numerical experiments that illustrate the effectiveness of the stabilized LSC preconditioners. Our focus in this work is on steady problems, although the ideas generalize in a straightforward manner to unsteady flow problems.

2. The least squares commutator preconditioner. In this section, we give a brief derivation of the LSC method for stable discretizations ($C = 0$ in (1.2)); additional details can be found in [6, 7]. The method is based on a notion of approximate algebraic commuting. Consider the linear convection-diffusion operator

$$\mathcal{L} = -\nu\nabla^2 + \mathbf{w} \cdot \nabla,$$

where the convection coefficient \mathbf{w} is a vector field in \mathbb{R}^d . This operator is defined on the velocity space and derived from linearization of the convection term in (1.1) via Picard iteration. We will also assume that there is an analogous operator

$$\mathcal{L}_p = (-\nu\nabla^2 + \mathbf{w} \cdot \nabla)_p$$

defined on the pressure space. These operators are used solely for the purpose of deriving certain matrix preconditioning operators, and we will not be precise about boundary conditions or function spaces.

Next, consider the commutator of the convection-diffusion operators with the gradient operator,

$$(2.1) \quad \mathcal{E} = \mathcal{L}\nabla - \nabla\mathcal{L}_p.$$

When \mathbf{w} is smooth or small, we expect this commutator to be small in some sense. If finite element methods are used to discretize the component operators of \mathcal{E} on both the velocity space and the pressure space, then a discrete version of the commutator takes the form

$$E = (Q_u^{-1}F)(Q_u^{-1}B^T) - (Q_u^{-1}B^T)(Q_p^{-1}F_p),$$

where Q_u and Q_p are the velocity mass matrix and pressure mass matrix, respectively. Assuming that this matrix version of the commutator is also small (i.e., $E \approx 0$), then a straightforward algebraic manipulation leads to the approximation

$$BF^{-1}B^T Q_p^{-1}F_p(BQ_u^{-1}B^T)^{-1} \approx I.$$

That is, we have the approximation

$$(2.2) \quad BF^{-1}B^T \approx BQ_u^{-1}B^T F_p^{-1}Q_p$$

to the Schur complement operator of (1.2) arising from stable discretizations. Note that the term $BQ_u^{-1}B^T$ can be viewed as a scaled discrete Laplacian operator; with this interpretation, the construction of (2.2) represents a particular example of (1.4).

As noted above, construction of the discrete operator F_p requires information about the discretization or access to a code that may not be readily available. Moreover, a precise definition of F_p also entails some decisions about boundary conditions; see [7, 10].

An alternative to constructing F_p this way is to define it algebraically in a way that minimizes a measure of the discrete version of the commutator. Given two vectors \mathbf{u} and \mathbf{v} in the discrete velocity space, the form

$$(\mathbf{u}, \mathbf{v})_{Q_u} = (Q_u \mathbf{u}, \mathbf{v})$$

defines an inner product that represents a discrete analogue of the continuous L_2 inner product on the velocity space. Let $\|\cdot\|_{Q_u}$ denote the induced norm. We can define the j th column of the matrix F_p to solve the weighted least squares problem

$$\min \| [Q_u^{-1} F Q_u^{-1} B^T]_j - Q_u^{-1} B^T Q_p^{-1} [F_p]_j \|_{Q_u}.$$

The normal equations associated with this problem are

$$Q_p^{-1} B Q_u^{-1} B^T Q_p^{-1} [F_p]_j = [Q_p^{-1} B Q_u^{-1} F Q_u^{-1} B^T]_j.$$

This leads to the following definition of F_p :

$$F_p = Q_p (B Q_u^{-1} B^T)^{-1} (B Q_u^{-1} F Q_u^{-1} B^T).$$

Substitution of this expression into (2.2) then gives an approximation to the Schur complement matrix:

$$(2.3) \quad B F^{-1} B^T \approx (B Q_u^{-1} B^T) (B Q_u^{-1} F Q_u^{-1} B^T)^{-1} (B Q_u^{-1} B^T).$$

For most discretizations, the inverse of the velocity mass matrix Q_u^{-1} appearing in this expression is a dense matrix, so it is not practical to use the matrix on the right-hand side of (2.3). This difficulty can be resolved by replacing Q_u with a diagonal matrix \hat{Q}_u , such as the diagonal part of Q_u , that is spectrally equivalent to Q_u ; see Wathen [18]. Another approach to avoiding the dense matrix Q_u^{-1} is never to explicitly form Q_u^{-1} , but to approximate it using an iterative solver. However, this approach would involve four extra solves with Q_u for every application of M_S . The \hat{Q}_u approximation leads to the least squares commutator preconditioner for stable discretizations,

$$(2.4) \quad B F^{-1} B^T \approx M_S = (B \hat{Q}_u^{-1} B^T) (B \hat{Q}_u^{-1} F \hat{Q}_u^{-1} B^T)^{-1} (B \hat{Q}_u^{-1} B^T).$$

Use of this preconditioner entails application of the action of the inverse of M_S . This requires solution of two subsidiary problems of the form

$$(B \hat{Q}_u^{-1} B^T) p = q$$

for p , which is essentially a discrete Poisson equation.

3. Stabilization. Stabilized mixed approximation methods can be found in many incompressible flow codes. In particular, equal-order finite elements and non-staggered finite difference grids are computationally convenient and increasingly used in preference to stable alternatives. The problem with stabilized methods within a linear algebra setting is that “derived” operators like $B Q_u^{-1} B^T$ in (2.3) do not give a true representation of their continuous analogue. This problem needs to be addressed

if uniformly effective preconditioning methods with grid-independent convergence are to be constructed.

We first discuss the following Stokes problem: find eigenvalues and associated eigenvectors satisfying

$$(3.1) \quad \begin{aligned} -\nabla^2 \mathbf{u} + \text{grad } p &= -\lambda \nabla^2 \mathbf{u}, \\ -\text{div } \mathbf{u} &= \lambda p \quad \text{in } \Omega, \end{aligned}$$

with $\mathbf{u} = \mathbf{0}$ on $\partial\Omega$. Note that there is a zero eigenvalue associated with the null vector $(\mathbf{0}, 1)$.

Let \mathcal{T}_h denote a subdivision (or triangularization) of the domain Ω . The matrix analogue of (3.1) derived from finite element discretization on \mathcal{T}_h is: find pairs $\{\lambda, (\mathbf{u}, p)\}$ such that

$$(3.2) \quad \begin{bmatrix} A_u & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix} = \lambda \begin{bmatrix} A_u & 0 \\ 0 & Q_p \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix}.$$

By eliminating \mathbf{u} from (3.2), we see that the pressure component satisfies

$$(3.3) \quad BA_u^{-1}B^T p = \underbrace{\lambda(\lambda - 1)}_{\sigma} Q_p p.$$

The discrete gradient matrix B^T must be rank deficient; specifically, if $p = \text{constant}$, then $\sigma = \lambda = 0$. It is also known that $\sigma < 1$; see Stoyan [17]. An inf-sup stable approximation is characterized by the requirements that (3.2) and (3.3) each have only one zero eigenvalue, and that the nonzero eigenvalues of (3.3) be bounded away from zero independently of the discretization parameter h ; see Elman, Silvester, and Wathen [7, Theorem 5.22]. Thus, if n_p is the dimension of the discrete pressure space and the eigenvalues of (3.3) are ordered so that

$$0 = \sigma_1 \leq \sigma_2 \leq \sigma_3 \leq \cdots \leq \sigma_{n_p} < 1,$$

then inf-sup stability means that $\sigma_2 \geq \sigma_* > 0$, where σ_* is independent of h .

Another way of phrasing this property of boundedness of eigenvalues is to say that the Schur complement matrix $BA_u^{-1}B^T$ and the pressure mass matrix Q_p are spectrally equivalent. A glance at the operators producing these matrices reveals that from a heuristic point of view, this is not a surprise. The pressure mass matrix is a representation of the identity operator in the discrete pressure space, and the Schur complement $BA_u^{-1}B^T$ is a matrix representation of “ $-\text{div}(\text{div} \cdot \text{grad})^{-1} \text{grad}$,” which we might expect to behave like an identity operator. The problem with unstable discretizations (those not satisfying the inf-sup condition) is that this “expected” behavior is not manifested. Instability can arise in two different ways: “spurious” pressure modes $\{p_j\}_{j=2}^r$ with $0 = \sigma_2 = \sigma_3 = \cdots = \sigma_r$, and “pesky” pressure modes $\{p_k\}$ with $\sigma_k = O(h)$ for some set of indices $\{k\}$, where h is the characteristic mesh parameter. Pesky modes are insidious in the sense that they do not affect the solvability of the discrete system and can be detected only when computing solutions on a sequence of grids. Both types of instability occur in the case of Q_1-P_0 or Q_1-Q_1 mixed approximation in \mathbb{R}^2 and in \mathbb{R}^3 . Specifically, for enclosed flow in \mathbb{R}^2 it is known that $r = 2$ and $r = 8$ for Q_1-P_0 and Q_1-Q_1 , respectively. A quantitative description of the pesky modes associated with Q_1-P_0 can be found in Gresho and Sani [8,

pp. 686–691]. Fortunately, both types of instability are removed by the stabilization methods outlined below.

Unstable elements can be stabilized by relaxing the discrete incompressibility constraint in a consistent manner. For the Stokes problem, this leads to the introduction of a nonzero *stabilization matrix* C in (1.2), as well as the replacement of the zero matrix in the (2,2)-block in (3.2) by C . The stabilization matrix C should make the eigenvalues δ_k of the stabilized Schur complement eigenvalue problem

$$(3.4) \quad (BA_u^{-1}B^T + C)p = \delta Q_p p$$

satisfy $\delta_2 \geq \delta_* > 0$ and $\delta_{n_p} < \delta^* < \infty$, where δ_* and δ^* are independent of the discretization mesh. It is expected that both these constants have magnitudes on the order of 1. We discuss the cases of continuous and discontinuous pressure approximation separately. The stabilization methods that we consider are described in more detail in Elman, Silvester, and Wathen [7, section 5.3]. An important feature of these methods is that they are “parameter-free.” Alternative methods exist, but these involve one or more “user-defined” parameters that typically have to be well chosen in order to get good results.

3.1. Lowest-order approximation (P_1-P_0, Q_1-P_0). This methodology was introduced by Kechkar and Silvester [12]. It constructs a stabilization operator using a nonoverlapping macroelement partitioning \mathcal{M}_h of the subdivision \mathcal{T}_h . One simple way to achieve this in \mathbb{R}^2 is to construct \mathcal{T}_h by uniformly refining a coarse mesh $\mathcal{T}_{h/2}$, subdividing every element into four smaller triangles or rectangles. In this case, $\mathcal{M}_h = \mathcal{T}_{h/2}$, and the macroelements all consist of four adjoining elements. If this is done in \mathbb{R}^3 , then macroelements consist of twelve adjoining tetrahedra or eight adjoining bricks. A different element aggregation strategy must be adopted to generate \mathcal{M}_h in an adaptive refinement setting; see Kay and Silvester [11] for details.

Given a suitable partitioning \mathcal{M}_h , the pressure stabilization operator penalizes pressure jumps across interior element boundaries (edges in \mathbb{R}^2 , faces in \mathbb{R}^3) within each macroelement:

$$(3.5) \quad c^{(macro)}(p_h, q_h) = \frac{1}{4} |\mathcal{M}| \sum_{e \in \Gamma_{\mathcal{M}}} \langle \llbracket p_h \rrbracket_e, \llbracket q_h \rrbracket_e \rangle_{\bar{E}}.$$

Here, $|\mathcal{M}|$ is the mean element area within the macroelement, the set $\Gamma_{\mathcal{M}}$ consists of interior element edges/faces in \mathcal{M} , $\llbracket \cdot \rrbracket_e$ is the jump across edge/face e , and $\langle p, q \rangle_{\bar{E}} = \frac{1}{|\bar{E}|} \int_{\bar{E}} pq$. For example, if a uniform grid of rectangular $h_x \times h_y$ elements is decomposed into 2×2 macroelements, then the 4×4 macroelement contribution matrix is

$$(3.6) \quad C^{(macro)} = \frac{h_x h_y}{4} \begin{bmatrix} 2 & -1 & 0 & -1 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ -1 & 0 & -1 & 2 \end{bmatrix}.$$

Using (3.5), a block diagonal stabilization matrix C is then assembled that has the contributing macroelement stabilization matrices as its diagonal blocks. Note that the null space of $C^{(macro)}$ consists of constant vectors, and this means that the assembled stabilization matrix again has a null space that is consistent with (3.3). An important property of mixed methods based on discontinuous pressure is that of elementwise




FIG. 3.1. 2×2 Q_1 - P_0 subdivision of a square domain, with four pressure nodes (numbered) and one interior velocity node (solid).

mass conservation:

$$(3.7) \quad 0 = \int_k \nabla \cdot \mathbf{u}_h = \int_{\partial k} \mathbf{u}_h \cdot \vec{n}.$$

Although elementwise mass conservation is lost in the jump stabilized formulation, local incompressibility is retained at the macroelement level:

$$(3.8) \quad 0 = \int_{\mathcal{M}} \nabla \cdot \mathbf{u}_h = \int_{\partial \mathcal{M}} \mathbf{u}_h \cdot \vec{n}.$$

It is possible to develop a stabilization matrix that penalizes *all* the interelement pressure jumps rather than just those within macroelements [9]. This would eliminate the need for a data structure that keeps track of macroelements, but it would also sacrifice local mass conservation as given in (3.7) or (3.8).

The new versions of the LSC preconditioners will be developed by mimicking some properties of this strategy for stabilizing the discrete Stokes operator. To motivate the ideas used for preconditioning, we look more closely at the Stokes operator for the case of a square domain Ω that is divided into a 2×2 grid of square elements of size h^2 , with the Q_1 - P_0 approximation; see Figure 3.1. In this case, there are four pressure nodes, located at element centers, and one velocity node not on the boundary, located at the juncture where the four elements meet. For this simple example, with a single node for each velocity component, the discrete Laplacian on the velocity space is the 2×2 matrix

$$(3.9) \quad A_u = \frac{8}{3} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

and the discrete gradient operator is

$$(3.10) \quad B^T = \begin{bmatrix} B_x^T \\ B_y^T \end{bmatrix} = \begin{bmatrix} -h/2 & h/2 & h/2 & -h/2 \\ -h/2 & -h/2 & h/2 & h/2 \end{bmatrix}.$$

This results in the Schur complement matrix

$$(3.11) \quad S = BA_u^{-1}B^T = \frac{3h^2}{16} \begin{bmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \\ -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \end{bmatrix},$$

whose eigenvalues and eigenvectors are

$$(3.12) \quad \left\{ 0, \underbrace{\begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}}_{\mathbf{q}_1} \right\}, \quad \left\{ 0, \underbrace{\begin{bmatrix} 1 \\ -1 \\ 1 \\ -1 \end{bmatrix}}_{\mathbf{q}_2} \right\}, \quad \left\{ \frac{3h^2}{8}, \underbrace{\begin{bmatrix} 1 \\ 1 \\ -1 \\ -1 \end{bmatrix}}_{\mathbf{q}_3} \right\}, \quad \left\{ \frac{3h^2}{8}, \underbrace{\begin{bmatrix} 1 \\ -1 \\ -1 \\ 1 \end{bmatrix}}_{\mathbf{q}_4} \right\}.$$

There is a zero eigenvalue associated with the “checkerboard mode” q_2 , and q_2 is the source of instability in the system—if any vector \mathbf{p} is a pressure solution of a discrete Stokes system defined on Ω , then $\mathbf{p} + \alpha q_2$ is also a solution. The stabilization operator $C^{(macro)}$ of (3.6) has the same eigenvectors as S , with eigenvalues h^2 , $h^2/2$, and $h^2/2$ for q_2 , q_3 , and q_4 , respectively. The pressure mass matrix is $Q_p = h^2 I$. Hence, the nonzero eigenvalues of the generalized problem (3.4) are given by

$$\delta_2 = 1, \quad \delta_3 = 7/8, \quad \delta_4 = 7/8.$$

In particular, in the stabilized system, $\delta_2 = 1$ replaces $\sigma_2 = 0$, and stabilization makes the influence of q_2 commensurate with that of q_3 and q_4 . Each of these three modes contributes to the pressure solution in a manner consistent with their presence in the right-hand-side vector, which ensures stability at the macroelement level. Global stability on more general domains is established by demonstrating a *macroelement connectivity condition* whereby an inf-sup condition is shown to hold on the union of patches composed of macroelements; see Stenberg [16] or Boland and Nicolaides [3] for a rigorous statement.

3.2. Equal-order approximation (P_1 - P_1 , Q_1 - Q_1). This methodology was developed by Bochev, Dohrmann, and Gunzburger [2]. Motivated by the rigorous error analysis of mixed methods based on continuous pressure approximation, the deficiency of equal-order approximation can be associated with the mismatch between the discrete divergence of the velocity field (a subspace of the space P_0 of discontinuous piecewise constant functions) and the actual discrete pressure space P_1 or Q_1 . To get into the “right space,” a suitable pressure stabilization operator is needed, namely,

$$(3.13) \quad c^{(proj)}(p_h, q_h) = (p_h - \Pi_0 p_h, q_h - \Pi_0 q_h),$$

where Π_0 is the L^2 projection from the pressure approximation space into the space P_0 . Note that this projection is defined locally: $\Pi_0 p_h$ is a constant function in each element $k \in \mathcal{T}_h$. It is determined simply by local averaging,

$$(3.14) \quad \Pi_0 p_h|_k = \frac{1}{|k|} \int_k p_h \quad \text{for all } k \in \mathcal{T}_h.$$

It is clear that the associated grid stabilization matrix C can be assembled from element contribution matrices in the same way as a standard finite element stiffness matrix. For example, in the case of a rectangular grid in \mathbb{R}^2 the 4×4 contribution matrix $C^{(proj)}$ is given by

$$(3.15) \quad C^{(proj)} = Q - qq^T |\square_k|,$$

where $|\square_k|$ is the area of element k , Q is the 4×4 element mass matrix for the bilinear discretization, and $q = [1/4, 1/4, 1/4, 1/4]^T$ is the local averaging operator. Note that the null space of $C^{(proj)}$ consists of constant vectors, and this means that the assembled stabilization matrix has a null space that is consistent with (3.3); that is, the constant vector is a null vector in both cases.

4. Stabilized LSC. The LSC preconditioner (2.4) does not handle stabilized approximations of the Navier–Stokes equations. As suggested in the previous section, the difficulty comes from the “derived” terms $B\hat{Q}_u^{-1}B^T$ and $B\hat{Q}_u^{-1}F\hat{Q}_u^{-1}B^T$ figuring in the definition of the preconditioning operator. In this section, we show how to modify the LSC preconditioner for stabilized approximations. We consider two approaches. First, we present stabilized versions of the two component terms, which are constructed by analogy with the stabilization matrix C of the previous section. Second, we show that it is possible to vary this approach in a manner that does not require knowledge of the underlying finite element discretization. This enables the construction of a stabilized preconditioner using essentially algebraic considerations, as in the derivation of the original operator of (2.4).

4.1. Element-based stabilized LSC. Stabilization strategies for the Stokes operator and the associated Schur complement $BA_u^{-1}B^T$ derive from the eigenvalue problem (3.1) and its discrete analogue (3.2). For the matrix $BQ_u^{-1}B^T$, the analogous continuous problem is the *potential flow* problem obtained by setting $\eta = 1$ and $\nu = 0$ in (1.1): find eigenvalues and associated eigenvectors satisfying

$$(4.1) \quad \begin{aligned} \mathbf{u} + \text{grad } p &= \lambda \mathbf{u}, \\ -\text{div } \mathbf{u} &= -\lambda \nabla^2 p \quad \text{in } \Omega, \end{aligned}$$

with $\mathbf{u} = \mathbf{0}$ on $\partial\Omega$. As for the analogous eigenvalue problem for the Stokes operator, (3.1), there is a zero eigenvalue with associated null vector $(\mathbf{0}, 1)$. The discrete analogue of (4.1) is: find pairs $\left\{ \lambda, \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix} \right\}$ such that

$$(4.2) \quad \begin{bmatrix} Q_u & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix} = \lambda \begin{bmatrix} Q_u & 0 \\ 0 & A_p \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix}.$$

Here, A_p is a discrete Laplacian defined on the pressure space; no boundary conditions are given for the pressure, so that A_p will have the form arising when a Neumann condition is specified. As in (3.3), the pressure components p are eigenvectors of an eigenvalue problem involving the (derived) Schur complement¹

$$(4.3) \quad BQ_u^{-1}B^T p = \sigma A_p p.$$

Note that the presence of the pressure Laplacian on the right-hand side of (4.1) and of A_p in (4.2) stems from the fact that for stable discretizations the eigenvalues of (4.3) satisfy

$$0 < \sigma_* \leq \sigma_2 \leq \sigma_3 \leq \cdots \leq \sigma_{n_p} \leq \sigma^* < \infty,$$

where σ_* and σ^* are independent of the subdivision mesh. That is, the Schur complement $BQ_u^{-1}B^T$ is spectrally equivalent to a discrete Laplacian on the pressure space. This is also related to an alternative inf-sup condition; see Elman, Silvester, and Wathen [7, pp. 272–273].

For unstable approximations like Q_1 – Q_0 and Q_1 – Q_1 , however, this alternative inf-sup condition will not be satisfied. In \mathbb{R}^2 , the spurious pressure modes $\{p_j\}_{j=2}^8$

¹Although this problem is more complicated than (3.3) in that both $BQ_u^{-1}B^T$ and A_p have a constant vector in their null spaces, this does not play any significant role in the derivation of a preconditioner.

cause (4.3) to have zero eigenvalues: $0 = \sigma_2 = \dots = \sigma_8$. Moreover, pesky pressure modes $\{p_k\}$ give rise to $O(h)$ eigenvalues in (4.3). What this means is that (4.2) is not an accurate approximation to (4.1), and $BQ_u^{-1}B^T$ as well as its approximation $B\hat{Q}_u^{-1}B^T$ suffer from exactly the same instabilities as those arising with the Stokes operator.

The way to fix this is to define a stabilization matrix C_1 in a manner analogous to what is done for the Stokes operator in section 3. That is, we construct a stabilized version of the eigenvalue problem (4.2) given by

$$(4.4) \quad \begin{bmatrix} Q_u & B^T \\ B & -C_1 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix} = \lambda \begin{bmatrix} Q_u & 0 \\ 0 & A_p \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix},$$

in such a way that the eigenvalues of the resulting Schur complement problem,

$$(4.5) \quad (BQ_u^{-1}B^T + C_1)p = \delta A_p p,$$

satisfy $\delta_2 \geq \delta_* > 0$ and $\delta_{n_p} < \delta^* < \infty$ for any conceivable grid.

To see how to construct such a matrix C_1 , first consider a 2×2 element grid as in section 3.1; see Figure 3.1. For Q_1 velocities, the discrete Laplacian appearing in the Stokes system is replaced in (4.4) by a velocity mass matrix

$$(4.6) \quad Q_u = \hat{Q}_u = \frac{4h^2}{9} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

For P_0 pressures (see (3.10)), the Schur complement is

$$BQ_u^{-1}B^T = \frac{9}{8} \begin{bmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \\ -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \end{bmatrix} = B\hat{Q}_u^{-1}B^T.$$

Up to scaling with respect to the element area, this is essentially the same as the analogous Schur complement (3.11) for the Stokes operator, and it has eigenvalues and eigenvectors given by

$$\{0, q_1\}, \quad \{0, q_2\}, \quad \{9/4, q_3\}, \quad \{9/4, q_4\}.$$

As above, the zero eigenvalue corresponding to the checkerboard mode q_2 is a source of instability. Adding the “unscaled” macroelement stabilization matrix

$$(4.7) \quad C_1 = \frac{1}{4} \begin{bmatrix} 2 & -1 & 0 & -1 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ -1 & 0 & -1 & 2 \end{bmatrix}$$

to $BQ_u^{-1}B^T$ will make the influence of the rogue eigenvector q_2 commensurate with q_3 and q_4 in the Schur complement $BQ_u^{-1}B^T + C_1$. In particular, $BQ_u^{-1}B^T + C_1$ has an eigenvalue of unity corresponding to q_2 .

Consider the impact of this stabilization operator on the generalized eigenvalue problem (4.5). The discrete Laplacian A_p for a P_0 approximation is constructed by

combining jump operators defined on interelement edges or faces. For details, see Elman, Silvester, and Wathen [7, pp. 352ff]. Specifically, for a uniform square grid, we have that

$$(4.8) \quad A_p = [a_{p,ij}], \quad a_{p,ij} = \sum_{E \in \mathcal{E}_{h,\Omega}} \langle \llbracket \psi_j \rrbracket, \llbracket \psi_i \rrbracket \rangle_E,$$

where $\{\psi_j\}$ is the basis of the P_0 approximation and $\mathcal{E}_{h,\Omega}$ is the set of interior edges. For the 2×2 element grid in Figure 3.1, this leads to

$$A_p = \begin{bmatrix} 2 & -1 & 0 & -1 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ -1 & 0 & -1 & 2 \end{bmatrix}.$$

That is, $A_p = 4C_1$, and it then follows that the generalized problem (4.5) has nonzero eigenvalues

$$\delta_2 = 1/4, \quad \delta_3 = 11/8, \quad \delta_4 = 11/8.$$

This discussion of the 2×2 macroelement shows what is needed in general to stabilize an unstable discretization of the potential operator and thereby identify the matrix C_1 needed to stabilize the factor $B\hat{Q}_u^{-1}B^T$ appearing in (2.4). The stabilization operator should be constructed in a manner analogous to what is done for Stokes problems, but scaled appropriately with respect to the mesh size. Thus, the analogue of (3.5) is given by omitting the scaling by mean element area, to give the form

$$(4.9) \quad c_1^{(macro)}(p_h, q_h) = \frac{1}{4} \sum_{e \in \Gamma_{\mathcal{M}}} \langle \llbracket p_h \rrbracket_e, \llbracket q_h \rrbracket_e \rangle_e.$$

For equal-order approximation (section 3.2), we rewrite $c^{(proj)}$ of (3.13) as a sum of integrals over the elements,

$$c^{(proj)}(p_h, q_h) = \sum_{k \in \mathcal{T}_h} (p_h - \Pi_0 p_h, q_h - \Pi_0 q_h)_k,$$

and then scale the local integrals by the inverse of the element area to give

$$(4.10) \quad c_1^{(proj)}(p_h, q_h) = \sum_{k \in \mathcal{T}_h} \frac{1}{|k|} (p_h - \Pi_0 p_h, q_h - \Pi_0 q_h)_k.$$

This means that the analogue of (3.15) for the 4×4 element matrix is

$$C_1^{(proj)} = \frac{1}{|\square_k|} Q - qq^T.$$

We now turn to the second problematic component of the LSC preconditioner (2.4), the matrix $B\hat{Q}_u^{-1}F\hat{Q}_u^{-1}B^T$. We can see what is needed to stabilize this term by again considering the 2×2 macroelement case. In particular, on a square grid with Q_1 discrete velocities on elements of size h^2 ,

$$Q_u^{-1}FQ_u^{-1} = \frac{81}{16h^4} (\nu A_u),$$

where Q_u is given by (4.6), A_u is as in (3.9), and the viscosity parameter ν derives from the discrete linearized version of (1.1). Then

$$BQ_u^{-1}FQ_u^{-1}B^T = \frac{\nu}{h^4}36(BA_u^{-1}B^T).$$

The essential difference from the Schur complement (3.11) for the Stokes operator is the scaling ν/h^4 , that is, the viscosity parameter divided by the *square* of the element area. For the stabilization to be of commensurate influence, this leads to the scaled forms

$$(4.11) \quad c_2^{(macro)}(p_h, q_h) = \frac{\nu}{4|\mathcal{M}|} \sum_{e \in \Gamma_{\mathcal{M}}} \langle \llbracket p_h \rrbracket_e, \llbracket q_h \rrbracket_e \rangle_E$$

for the lowest-order case and

$$(4.12) \quad c_2^{(proj)}(p_h, q_h) = \sum_{k \in \mathcal{T}_h} \frac{\nu}{|k|^2} (p_h - \Pi_0 p_h, q_h - \Pi_0 q_h)_k$$

for the equal-order approximation case. Each of these forms generates an associated stabilization matrix C_2 .

In summary, the *stabilized LSC preconditioner* is

$$(4.13) \quad M_S^{-1} = (B\hat{Q}_u^{-1}B^T + C_1)^{-1}(B\hat{Q}_u^{-1}F\hat{Q}_u^{-1}B^T + C_2)(B\hat{Q}_u^{-1}B^T + C_1)^{-1},$$

where C_1 and C_2 are defined by (4.9)/(4.11) or (4.10)/(4.12). We hypothesise that this strategy would work for other unstable mixed approximations: if a stabilization operator C is constructed using some local element assembly process, then C_1 can be defined by scaling the local matrices by the inverse of the local element size, and C_2 can be defined by scaling by the viscosity times the square of the inverse of the local element size.

4.2. Algebraically stabilized LSC. A feature of the LSC preconditioner (2.4) for div-stable discretizations is that it is defined from purely algebraic considerations, using matrices that arise naturally from finite element discretization of the Navier–Stokes equations. In particular, its construction does not depend on knowledge of the underlying grid, which may be unavailable to the developer of a “solver” routine. This feature does not carry over to the preconditioner of (4.13): the stabilization matrices C_1 and C_2 are defined using local element information, and they are needed only for the preconditioner. In this section, we introduce a stabilized preconditioner that does not use local element information. For this, a modification is made to the “Laplacian component” of the preconditioner, $B\hat{Q}_u^{-1}B^T$, in a manner analogous to what was done in section 4.1 but using the existing stabilization matrix C . An extra term is added to compensate for the singularity of the operator $B\hat{Q}_u^{-1}F\hat{Q}_u^{-1}B^T$. The resulting preconditioner is a sum of two preconditioners, and it bears some resemblance to an additive Schwarz method; see Smith, Bjørstad, and Gropp [15]. Specifically,

$$(4.14) \quad M_S^{-1} = M_{S_\gamma}^{-1} + M_{S_\alpha}^{-1},$$

where the two parts of the preconditioner are given by

$$M_{S_\gamma}^{-1} = (B\hat{Q}_u^{-1}B^T + \gamma C)^{-1}(B\hat{Q}_u^{-1}F\hat{Q}_u^{-1}B^T)(B\hat{Q}_u^{-1}B^T + \gamma C)^{-1}$$

and

$$(4.15) \quad M_{S_\alpha}^{-1} = \alpha D^{-1},$$

where D is the diagonal of $(B \operatorname{diag}(F)^{-1} B^T + \frac{1}{\nu} C)$ and $\operatorname{diag}(F)$ is the diagonal part of F . This approach requires choices for the scalars γ and α .

These parameters will be determined using Fourier analysis. This requires a simplified statement of the problem. In particular, we will explore a steady version of (1.1) where the nonlinear term $(\mathbf{u} \cdot \operatorname{grad}) \mathbf{u}$ is replaced by a linear one $(\mathbf{w} \cdot \operatorname{grad}) \mathbf{u}$ (Oseen linearization), where \mathbf{w} is constant. The simplified problem is posed on a square in two dimensions or a cube in three dimensions with periodic boundary conditions, and the discretization is assumed to be on a uniform $(n \times n$ or $n \times n \times n)$ grid in which the number of velocity components in each coordinate direction is equal to the total number of pressure components. This would be the case, for example, in a Q_1 - Q_1 or Q_1 - P_0 discretization on a periodic mesh. We emphasize that these assumptions are made only to develop the stabilization by defining the parameters α and γ . These restrictions coincide with standard suppositions associated with Fourier stability analysis. Due to the local nature of stabilization, Fourier stability analysis often accurately describes more general problems even though it is formally valid only in limited situations. In section 5, the resulting preconditioner M_S^{-1} will be applied to problems that are not periodic and do not have constant coefficient stencils.

The Oseen linearization leads to the operator $\hat{Q}_u^{-1} F$ with block form in two dimensions²

$$(4.16) \quad \hat{Q}_u^{-1} F = \begin{bmatrix} \hat{F} & 0 \\ 0 & \hat{F} \end{bmatrix} = I_2 \otimes \hat{F},$$

where the block structure arises by grouping the velocities in each coordinate direction and I_2 refers to a 2×2 identity matrix. Notice that we have absorbed the blocks of \hat{Q}_u^{-1} into \hat{F} . Similarly, B has the block form $[B^x \ B^y]$, where $(B^x)^T$ and $(B^y)^T$ correspond to the derivative operator in the x and y coordinate directions. The assumptions of periodicity and constant coefficients imply that each of the matrices \hat{F} , B^x , B^y , C , and the product $B \hat{Q}_u^{-1} B^T$ commutes with one other and that they are all diagonalized by the Fourier transform. Specifically,

$$(4.17) \quad X \hat{F} X^H = \operatorname{diag}(f_{(\theta_x, \theta_y)}),$$

$$(4.18) \quad X B \hat{Q}_u^{-1} B^T X^H = \operatorname{diag}(a_{(\theta_x, \theta_y)}),$$

$$(4.19) \quad X C X^H = \operatorname{diag}(c_{(\theta_x, \theta_y)}),$$

where X is a matrix whose columns correspond to Fourier vectors and $\operatorname{diag}(g_{(\theta_x, \theta_y)})$ is the diagonal matrix with entries $g_{(\theta_x, \theta_y)}$ on the diagonal. The Fourier components that make up X have the form

$$(4.20) \quad \frac{1}{n} e^{2\pi i \theta_x j_x / n} e^{2\pi i \theta_y j_y / n},$$

where θ_x and θ_y are the frequencies in each coordinate direction, with $|\theta_x| = 1, \dots, n/2$ and $|\theta_y| = 1, \dots, n/2$, and j_x and j_y are the grid indices in each coordinate direction, with $j_x = 1, \dots, n$ and $j_y = 1, \dots, n$. For convenience, we will use the symbol θ to represent a generic frequency (θ_x, θ_y) ; i.e., we will write $c_{(\theta_x, \theta_y)}$ as c_θ .

²The analysis for three-dimensional problems is identical.

Under these assumptions, the Fourier-transformed Schur complement is

$$\begin{aligned} X S X^H &= X B F^{-1} B^T X^H + \frac{1}{\nu} X C X^H \\ &= X [B^x \ B^y] (I_2 \otimes \hat{F}^{-1}) \hat{Q}_u^{-1} B^T X^H + \frac{1}{\nu} X C X^H \\ &= X \hat{F}^{-1} [B^x \ B^y] \hat{Q}_u^{-1} B^T X^H + \frac{1}{\nu} X C X^H \\ &= (X \hat{F}^{-1} X^H) (X B \hat{Q}_u^{-1} B^T X^H) + \frac{1}{\nu} X C X^H, \end{aligned}$$

where we have used the commuting property of \hat{F}^{-1} , B^x , and B^y and the fact that X is a unitary matrix. Thus, the Schur complement is diagonalized by the Fourier transform, and the transformed Schur complement has diagonal elements given by

$$(4.21) \quad s_\theta = f_\theta^{-1} a_\theta + \frac{1}{\nu} c_\theta.$$

Here, f_θ is the symbol for a convection-diffusion operator, a_θ is from the (derived) Laplacian $B \hat{Q}_u^{-1} B^T$, and c_θ is from the stabilization operator C . A similar exercise leads to the Fourier symbol m_θ for the Schur complement preconditioner M_S^{-1} :

$$(4.22) \quad m_\theta = f_\theta a_\theta (a_\theta + \gamma c_\theta)^{-2} + \alpha d^{-1},$$

where the scalar d^{-1} corresponds to the eigenvalue of the constant coefficient matrix $D^{-1} (= d^{-1} I)$. Thus, the preconditioned system $S M_S^{-1}$ can be diagonalized by a Fourier transform, leading to an expression for the eigenvalues of $S M_S^{-1}$:

$$(4.23) \quad s_\theta m_\theta = a_\theta^2 (a_\theta + \gamma c_\theta)^{-2} + \frac{1}{\nu} f_\theta a_\theta c_\theta (a_\theta + \gamma c_\theta)^{-2} + \alpha s_\theta d^{-1}.$$

Analytic trigonometric definitions of f_θ , c_θ , and a_θ can be used for particular constant coefficient stencils to obtain a function involving stencil coefficients, θ , γ , and α . For example, if the one-dimensional convection-diffusion operator is discretized using centered finite differences with mesh spacing h , we obtain

$$f_\theta = ih \sin(2\pi\theta h) + 4\nu \sin(2\pi\theta h/2)^2.$$

In principle, optimal algorithm parameters could be determined by finding the values of γ and α that minimize the maximum of $s_\theta m_\theta$ over all θ_x and θ_y . These expressions, however, are associated with particular constant coefficient stencils and are quite complex.

Instead, we further simplify the analysis to find computable expressions for α and γ that lead to a practical method, although the parameters may not be optimal. We use the following observations:

1. The first term of (4.23) corresponds to the main part of the preconditioned operator. For a stable discretization, this is the only term present, and $s_\theta m_\theta$ is identically 1 for all θ . (Thus, this is a perfect preconditioner for the periodic problem; see Elman [5].) For stabilized problems, any positive choice of γ ensures that this term is well defined and bounded above by one. Therefore, we need not be concerned about the first term of (4.23) becoming large.




FIG. 4.1. Magnitude of the eigenvalues of the preconditioned linear system as a function of z in one dimension with $\alpha = 0$, constant wind ($\nu = 0.1$), constant coefficient stencil, periodic boundary conditions, and $N = 100,000$. The left side of each plot corresponds to low frequencies and the right corresponds to higher frequencies.

2. The middle term of (4.23) comes from the interaction of the stabilizing term of the discrete Schur complement and the preconditioner. If an unstabilized preconditioner (with $\gamma = 0$) is applied to a stabilized problem, then this term has the form $\frac{1}{\nu} f_{\theta} c_{\theta} / a_{\theta}$. For spurious modes, standard stability arguments show that neither f_{θ} nor c_{θ} will vanish. An unstabilized preconditioning operator does not handle spurious modes, and the denominator a_{θ} will be zero in such cases. A positive γ is needed to remove this singularity and keep this term under control.
3. The first two terms of (4.23) are zero for spurious modes. The third term is the symbol for a scaled Schur complement. It will be bounded for all modes, and, moreover, we choose α so that this term is approximately one for spurious modes. This will ensure that the symbol $s_{\theta} m_{\theta}$ for the preconditioned operator is bounded below away from 0.

Using these observations, we can determine choices of γ and α independently, choosing γ to force the (stabilized) middle term to be bounded above, and choosing α to bound the third term below for modes where the first two terms are zero.

Let us introduce the variable $z_{\theta} = a_{\theta} / c_{\theta}$. For values of θ corresponding to spurious and pesky modes, $c_{\theta} \neq 0$ so that z_{θ} is well defined in this regime. To provide insight into how γ should be defined, in Figure 4.1 we plot the magnitude of $s_{\theta} m_{\theta}$ as a function of z for a one-dimensional problem with $\alpha = 0$. In this example, the matrix C is a standard central difference Laplacian scaled by one fourth, B is a standard central difference gradient operator, and F is a standard convection-diffusion operator. The “alternating pressure” function is the only spurious mode in this case. The left side of each plot corresponds to low frequencies. The ratio z_{θ} tends toward a constant as the lowest frequencies are approached. (In this case the constant is four, due to the scaling of C .) For $\gamma = 0$, the function is unbounded as the spurious mode is approached. As γ is increased, the peak moves to the left and becomes smaller in magnitude. In fact, for these computations it is easy to verify that the location of the peaks is approximately

equal to γ . For example, when $\gamma = 0.04$ in the “small γ ” picture on the left side of Figure 4.1, the peak occurs at $z = 0.04003932$, and when $\gamma = 0.6$ in the “intermediate γ ” picture on the right, the peak occurs at $z = 0.6000749$.

Since the first term (4.23) is bounded, we derive an explicit value of γ using the middle term. Rewriting this term as a function of f_θ and z_θ gives

$$(4.24) \quad \frac{f_\theta z_\theta / \nu}{(z_\theta + \gamma)^2}.$$

Consider an approximation obtained by replacing f_θ with a constant, for which we use

$$(4.25) \quad f_\theta = \rho(\hat{F}).$$

With this choice, we can view (4.24) as a function of the single parameter z_θ . Although the approximation (4.25) may appear crude, it is important to keep in mind that f_θ does not vary substantially compared to other terms in (4.24) near the small set of frequencies corresponding to the sharp peak in (4.24) (when γ is small). This approximation was implicitly made in the previous section by exclusively considering 2×2 macroelements.

Thus, we will use the maximal value of

$$(4.26) \quad g(z_\theta) = \frac{\rho(\hat{F})z_\theta / \nu}{(z_\theta + \gamma)^2}$$

to determine γ . Solving for $g'(z_\theta) = 0$ reveals that the maximum occurs at $z^* = \gamma$. Notice that $\gamma = 0$ gives $z^* = 0$, which is appropriate. To obtain an expression for γ , we evaluate (4.23) at $z^* = \gamma$ and set it equal to one. That is, we are choosing γ so that

$$(4.27) \quad \lambda(SM_S^{-1}(z^*)) \approx 1.$$

This forces the preconditioned linear system to be well behaved for Fourier modes corresponding to z^* and leads to the following formula:

$$(4.28) \quad \gamma = \frac{\rho(\hat{F})}{3\nu}.$$

Next we address the choice of the parameter α in (4.15). Notice that when a_θ is zero for a nonzero γ , (4.23) is zero (corresponding to the far right of the plots in Figure 4.1). We use α in the component $M_{S_\alpha}^{-1}$ of the preconditioner (see (4.14)–(4.15)) to fix this. As with the parameter γ , we want to find the smallest α that makes the eigenvalues of SM_S^{-1} approximately equal to one. To do this, we find the largest value of s_θ/d and set α equal to its reciprocal. Unfortunately, doing this in the most straightforward way involves the Schur complement, which we do not want to compute. In the high frequency case, we can approximate F with $\text{diag}(F)$ (an approximation similar to (4.25)). We therefore take α to be

$$(4.29) \quad \alpha = \frac{-1}{\rho(SD^{-1})} = \frac{-1}{\rho(B \text{diag}(F)^{-1}B^T D^{-1})},$$

where, as before, D is the diagonal of $(B \text{diag}(F)^{-1}B^T + \frac{1}{\nu}C)$ and $\text{diag}(F)$ is the diagonal part of F .

We end this section with a simple adaptation for nonuniform meshes, where Fourier analysis is not applicable. We note that some new analytic tools developed in [1] offer an alternative approach for developing and analyzing preconditioners for problems on nonuniform grids. The matrix C is defined with the intent of stabilizing $BF^{-1}B^T$. The method we have developed in this section for uniform meshes uses this same matrix, scaled by γ , to stabilize the (different) Schur complement operator $B\hat{Q}_u^{-1}B^T$. We have seen in section 4.1 that, for nonuniform meshes, stabilization of the latter operator takes into account the local element area, and such information cannot be incorporated into a single global scalar. To compensate for spatial variation, we define a diagonal matrix D_r whose elements are equal to the (componentwise) ratio of the diagonal entries of $B\hat{Q}_u^{-1}B^T$ to the diagonal entries of C . We then take M_S^{-1} to be

$$(4.30) \quad M_S^{-1} = (B\hat{Q}_u^{-1}B^T + \tilde{\gamma}D_r^{\frac{1}{2}}CD_r^{\frac{1}{2}})^{-1}(B\hat{Q}_u^{-1}F\hat{Q}_u^{-1}B^T)(B\hat{Q}_u^{-1}B^T + \tilde{\gamma}D_r^{\frac{1}{2}}CD_r^{\frac{1}{2}})^{-1} + \alpha D^{-1},$$

where $\tilde{\gamma} = \gamma/\|\text{diag}(D_r)\|_\infty$ and α is given by (4.29). In the case of uniform mesh and constant coefficients, D_r is just a scaled identity; its introduction has no effect on the preconditioner, and we recover (4.14). The idea behind D_r is that the amount of dissipation added to a row of $B\hat{Q}_u^{-1}B^T$ should be proportional to the size of the diagonal of $B\hat{Q}_u^{-1}B^T$. This can be roughly accomplished by scaling C such that its diagonal variations mirror those of $B\hat{Q}_u^{-1}B^T$.

5. Numerical results. In this section, we present the results of numerical experiments with the stabilized versions of the LSC preconditioner described in section 4.³ We consider two benchmark problems on domains depicted in Figure 5.1:

1. *Flow over a backward facing step.* The steady version of (1.1) ($\eta = 0$) is posed on a step-shaped domain with boundary conditions consisting of a parabolic velocity profile




FIG. 5.1. Domains and samples of streamlines for the backward facing step with $Re = 100$ (left) and driven cavity problem with $Re = 200$ (right).

³It is not our intention in this study to compare the LSC preconditioner with the pressure convection-diffusion method of (1.4). See [6], [7, Chapter 8] for such comparisons, where it is observed that each of these has some advantages with respect to the other.

$$u_1 = 4y(1 - y), \quad u_2 = 0$$

at the inflow $x = -1$, no-flow (zero velocity) conditions on the horizontal walls, and a Neumann condition

$$\nu \frac{\partial u_1}{\partial x} - p = 0, \quad \frac{\partial u_2}{\partial x} = 0$$

at the outflow boundary $x = 5$.

2. *Driven cavity flow.* This enclosed flow problem in a square has no-flow conditions on the bottom, left, and right boundaries, and the regularized condition

$$u_1 = 1 - x^4, \quad u_2 = 0$$

modeling the moving cavity lid.

The backward facing step is discretized on a uniform grid of square finite elements, where the grid is such that the rectangle enclosing the step would contain a uniform $n \times 3n$ grid of elements of width $h = 2/n$. The cavity is discretized on an $n \times n$ rectangular element grid; this problem is studied using both uniform grids (of width $h = 2/n$) and nonuniform stretched grids with more points concentrated near the domain boundaries. The discrete problems were generated using the IFISS software package developed by Silvester, Elman, and Ramage [14].

Details of the computations are as follows. For each problem, the nonlinear algebraic system derived from discretization of the steady Navier–Stokes equations was solved by Picard iteration [7] so that the nonlinear residual satisfies

$$\left\| \begin{bmatrix} \mathbf{f} - (F(\mathbf{u})\mathbf{u} + B^T p) \\ g - (B\mathbf{u} - Cp) \end{bmatrix} \right\| \leq 10^{-5} \left\| \begin{bmatrix} \mathbf{f} \\ g \end{bmatrix} \right\|.$$

The linear system for the correction that arises from the next Picard step is

$$(5.1) \quad \begin{bmatrix} F(\mathbf{u}) & B^T \\ B & -C \end{bmatrix} \begin{bmatrix} \Delta \mathbf{u} \\ \Delta p \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{f}} \\ \hat{g} \end{bmatrix},$$

where the right-hand side is the nonlinear residual. The results reported in this section correspond to the solution of such systems. We denote the system (5.1) as $\mathcal{A}x = b$ and report iteration counts for preconditioned GMRES to satisfy the tolerance

$$\|b - \mathcal{A}x_k\|_2 \leq 10^{-6} \|b\|_2,$$

with zero initial iterate.

Tables 5.1, 5.2, and 5.3 show the results for the element-based stabilized LSC preconditioner, for various Reynolds numbers and grid sizes. Here $Re = 2/\nu$. Each of these tables shows on the left the iteration counts required when the stabilized LSC preconditioner (4.13) is applied to problems arising from the (stabilized) Q_1 – Q_1 and Q_1 – P_0 discretizations described in section 3. On the right, they show the analogous counts when the LSC preconditioner (2.4) is used with a stable discretization, Q_2 – Q_1 , which consists of biquadratic velocities and bilinear pressures. The main point to observe here is that for each of the problems, the trends in iteration counts are the same for the three discretizations. In particular, the stabilization strategies introduced in section 4 lead to methods whose performance is the same as when no stabilization

TABLE 5.1

Iteration counts for element-based LSC-preconditioned GMRES applied to the backward facing step on a uniform $2^n \times 3 \cdot 2^n$ grid, for three choices of elements.

n	Stabilized								Stable			
	$Q_1 - Q_1$				$Q_1 - P_0$				$Q_2 - Q_1$			
	4	5	6	7	4	5	6	7	4	5	6	7
$Re=10$	12	15	20	28	11	15	21	31	11	15	19	23
$Re=100$	19	17	20	29	17	16	22	32	18	17	21	29
$Re=200$	32	27	20	27	26	18	21	30	30	24	22	29

TABLE 5.2

Iteration counts for element-based LSC-preconditioned GMRES applied to the driven cavity problem on a uniform $2^n \times 2^n$ grid, for three choices of elements.

n	Stabilized						Stable		
	$Q_1 - Q_1$			$Q_1 - P_0$			$Q_2 - Q_1$		
	5	6	7	5	6	7	5	6	7
$Re=10$	11	15	20	12	16	22	11	16	18
$Re=100$	17	19	25	17	21	28	16	21	27
$Re=500$	42	34	31	36	30	32	36	34	37
$Re=1000$	63	60	42	51	50	36	62	55	45

TABLE 5.3

Iteration counts for element-based LSC-preconditioned GMRES applied to the driven cavity problem on a nonuniform $2^n \times 2^n$ grid, for three choices of elements.

n	Stabilized								Stable			
	$Q_1 - Q_1$				$Q_1 - P_0$				$Q_2 - Q_1$			
	4	5	6	7	4	5	6	7	4	5	6	7
$Re=10$	14	18	25	34	13	17	25	37	12	18	24	31
$Re=100$	21	31	41	53	19	26	37	51	17	26	37	53
$Re=500$	35	43	68	101	31	36	54	80	29	45	67	97
$Re=1000$	47	61	77	118	41	49	60	90	41	63	92	132

TABLE 5.4

Iteration counts for algebraic LSC-preconditioned GMRES applied to the backward facing step on a uniform $2^n \times 3 \cdot 2^n$ grid, two stabilized elements.

n	Stabilized							
	$Q_1 - Q_1$				$Q_1 - P_0$			
	4	5	6	7	4	5	6	7
$Re=10$	22	22	25	29	16	18	26	37
$Re=100$	28	30	32	36	22	24	27	38
$Re=200$	30	30	32	35	26	28	29	37

is required. This conclusion is further substantiated in Figure 5.2, which plots the residual norm through the course of the GMRES iteration for two of these examples, corresponding to the backward step with $n = 5$ and $Re = 100$ and the cavity on a nonuniform grid with $n = 5$ and $Re = 500$. In particular, these results indicate that the weights given to the stabilization operators C_1 and C_2 are appropriate.

Tables 5.4, 5.5, and 5.6 show the iteration counts required when the same set of problems are solved with preconditioned GMRES using the “algebraic” versions of the stabilized LSC preconditioner described in section 4.2. These results should be compared with those in Tables 5.1–5.3. They indicate that the algebraic approach is highly competitive with the element-based method: it tends to be somewhat slower for the finest meshes but faster in cases with coarse mesh and large Reynolds numbers.

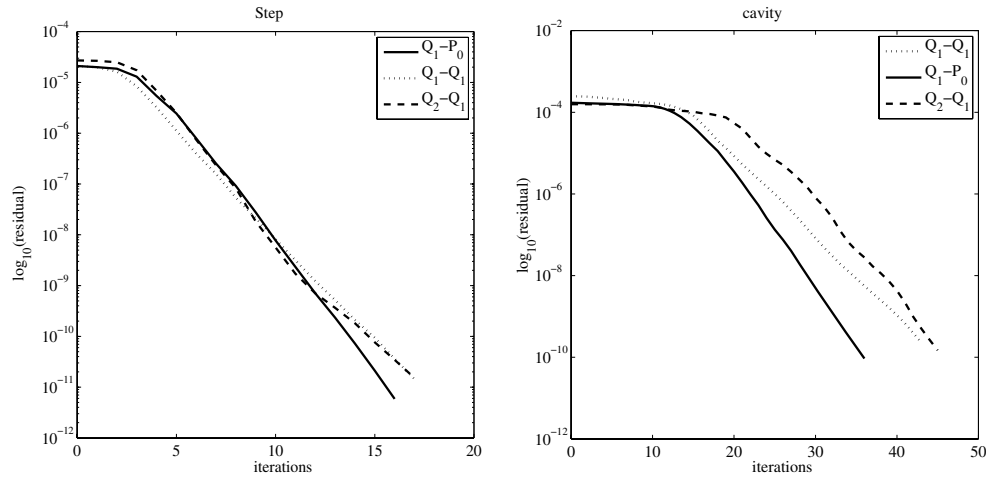


FIG. 5.2. Performance of element-based LSC-preconditioned GMRES on two benchmark problems. Left: backward facing step, $Re = 100$, $n = 5$. Right: driven cavity, $Re = 500$, $n = 5$, nonuniform grid.

TABLE 5.5

Iteration counts for algebraic LSC-preconditioned GMRES applied to the driven cavity problem on a uniform $2^n \times 2^n$ grid, for two stabilized elements.

n	Stabilized					
	$Q_1 - Q_1$			$Q_1 - P_0$		
	5	6	7	5	6	7
$Re=10$	17	18	22	18	23	33
$Re=100$	27	30	32	22	28	40
$Re=500$	37	42	45	32	36	42
$Re=1000$	43	49	56	41	44	52

TABLE 5.6

Iteration counts for algebraic LSC-preconditioned GMRES applied to the driven cavity problem on a nonuniform $2^n \times 2^n$ grid, for two stabilized elements.

n	Stabilized							
	$Q_1 - Q_1$				$Q_1 - P_0$			
	4	5	6	7	4	5	6	7
$Re=10$	20	25	34	50	20	27	40	68
$Re=100$	26	34	46	67	24	30	48	79
$Re=500$	34	38	59	99	34	36	57	100
$Re=1000$	35	47	63	107	41	46	61	107

Finally, we investigate the sensitivity of the linear solve to our choices of the α and γ parameters and the matrix D in the “algebraic” preconditioner. There is little sensitivity to the use of $\text{diag}(F)$ versus F in (4.29) or in the matrix D in (4.15) and (4.29). In practice, we have found that replacing $\text{diag}(F)$ with the full F matrix in the algebraic preconditioner has little effect on iteration counts. However, there is a more noticeable effect from other approximations in our choices of α and γ . This can be seen in Tables 5.7 and 5.8, where we ran a simple optimization loop over α and γ with initial values taken to be those determined by the Fourier analysis, minimizing with respect to iteration count. Table 5.7 displays the lowest

TABLE 5.7

Iteration counts for optimized algebraic LSC-preconditioned GMRES applied to the backward facing step on a uniform $2^n \times 3 \cdot 2^n$ grid, two stabilized elements.

n	Stabilized							
	$Q_1 - Q_1$				$Q_1 - P_0$			
	4	5	6	7	4	5	6	7
$Re=10$	20	19	18	19	15	16	16	17
$Re=100$	28	30	32	35	22	23	25	28
$Re=200$	29	30	32	34	26	26	28	30

TABLE 5.8

Iteration counts for optimized algebraic LSC-preconditioned GMRES applied to the driven cavity problem on a uniform $2^n \times 2^n$ grid, for two stabilized elements.

n	Stabilized					
	$Q_1 - Q_1$			$Q_1 - P_0$		
	5	6	7	5	6	7
$Re=10$	15	13	13	14	14	14
$Re=100$	26	29	29	22	24	27
$Re=500$	37	40	42	32	36	37
$Re=1000$	41	46	52	40	44	46

iteration count determined by this simple optimization on the backward facing step problem. Comparison of these results with Table 5.4 indicates that our analysis has given reasonable values for the parameters for the higher Reynolds numbers, with the optimized parameters sometimes reducing the iteration count by a small amount, but not dramatically. There is a more noticeable reduction in iteration count with the optimized parameters for lower Reynolds numbers. Similarly, Table 5.8 displays the optimized iteration count on the lid-driven cavity problem. Comparison with Table 5.5 again shows that our analysis has given reasonable values for the parameters in the higher Reynolds number case, but that they are less optimal for lower Reynolds numbers. One explanation for the parameters being less optimal for lower Reynolds numbers is the effect of our approximating \hat{F} by $\rho(\hat{F})$ in (4.25) and (4.29). This approximation is chosen because it does not have significant impact on the high frequency modes. The high frequency modes are more dominant, and thus our approximation is better, for higher Reynolds number problems where the boundary layers in the solution become thinner, that is, where the solution becomes less smooth.

6. Concluding remarks. In this study we have shown that the *least squares commutator* preconditioner for the discrete linearized Navier–Stokes equations can be extended for use in solving the systems of equations that arise from stabilized finite element discretizations. The new preconditioner is developed by stabilizing the components of the LSC operator in a manner analogous to what is done to stabilize the discretization operator, and as in the case for stable problems, a version of the preconditioner can also be constructed from purely algebraic considerations. Note that the costs of implementing the element-based and algebraic versions of this preconditioner are virtually identical; the difference lies in the technique used for stabilization. The performance of the new method for stabilized discrete problems is virtually identical to the performance of the LSC preconditioner for stable problems.

REFERENCES

- [1] B. BECKERMANN AND S. SERRA-CAPIZZANO, *On the asymptotic spectrum of finite element matrix sequences*, SIAM J. Numer. Anal., 45 (2007), pp. 746–769.
- [2] P. B. BOCHEV, C. R. DOHRMANN, AND M. D. GUNZBURGER, *Stabilization of low-order mixed finite elements for the Stokes equations*, SIAM J. Numer. Anal., 44 (2006), pp. 82–101.
- [3] J. M. BOLAND AND R. A. NICOLAIDES, *Stability of finite elements under divergence constraints*, SIAM J. Numer. Anal., 20 (1983), pp. 722–731.
- [4] F. BREZZI AND M. FORTIN, *Mixed and Hybrid Finite Element Methods*, Springer-Verlag, New York, 1991.
- [5] H. C. ELMAN, *Preconditioning for the steady-state Navier–Stokes equations with low viscosity*, SIAM J. Sci. Comput., 20 (1999), pp. 1299–1316.
- [6] H. ELMAN, V. E. HOWLE, J. SHADID, R. SHUTTLEWORTH, AND R. TUMINARO, *Block preconditioners based on approximate commutators*, SIAM J. Sci. Comput., 27 (2006), pp. 1651–1668.
- [7] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers*, Oxford University Press, Oxford, UK, 2005.
- [8] P. M. GRESHO AND R. SANI, *Incompressible Flow and the Finite Element Method: Isothermal Laminar Flow*, Wiley, Chichester, UK, 1998.
- [9] T. J. HUGHES AND L. P. FRANCA, *A new finite element formulation for computational fluid dynamics: VII. The Stokes problem with various well-posed boundary conditions: Symmetric formulations that converge for all velocity/pressure spaces*, Comput. Methods Appl. Mech. Engrg., 65 (1987), pp. 85–96.
- [10] D. KAY, D. LOGHIN, AND A. WATHEN, *A preconditioner for the steady-state Navier–Stokes equations*, SIAM J. Sci. Comput., 24 (2002), pp. 237–256.
- [11] D. KAY AND D. SILVESTER, *A posteriori error estimation for stabilized mixed approximations of the Stokes equations*, SIAM J. Sci. Comput., 21 (1999), pp. 1321–1336.
- [12] N. KECHKAR AND D. SILVESTER, *Analysis of locally stabilized mixed finite element methods for the Stokes problem*, Math. Comp., 58 (1992), pp. 1–10.
- [13] D. SILVESTER, H. ELMAN, D. KAY, AND A. WATHEN, *Efficient preconditioning of the linearized Navier–Stokes equations for incompressible flow*, J. Comput. Appl. Math., 128 (2001), pp. 261–279.
- [14] D. J. SILVESTER, H. C. ELMAN, AND A. RAMAGE, *IFISS: Incompressible Flow Iterative Solution Software*, available online at <http://www.manchester.ac.uk/ifiss>.
- [15] B. SMITH, P. BJØRSTAD, AND W. GROPP, *Domain Decomposition*, Cambridge University Press, Cambridge, UK, 1996.
- [16] R. STENBERG, *Analysis of mixed finite element methods for the Stokes problem: A unified approach*, Math. Comp., 42 (1984), pp. 9–23.
- [17] G. STOYAN, *Towards discrete Veltre decompositions and narrow bounds for inf-sup constants*, Comput. Math. Appl., 38 (1999), pp. 243–261.
- [18] A. J. WATHEN, *Realistic eigenvalue bounds for the Galerkin mass matrix*, IMA J. Numer. Anal., 7 (1987), pp. 449–457.