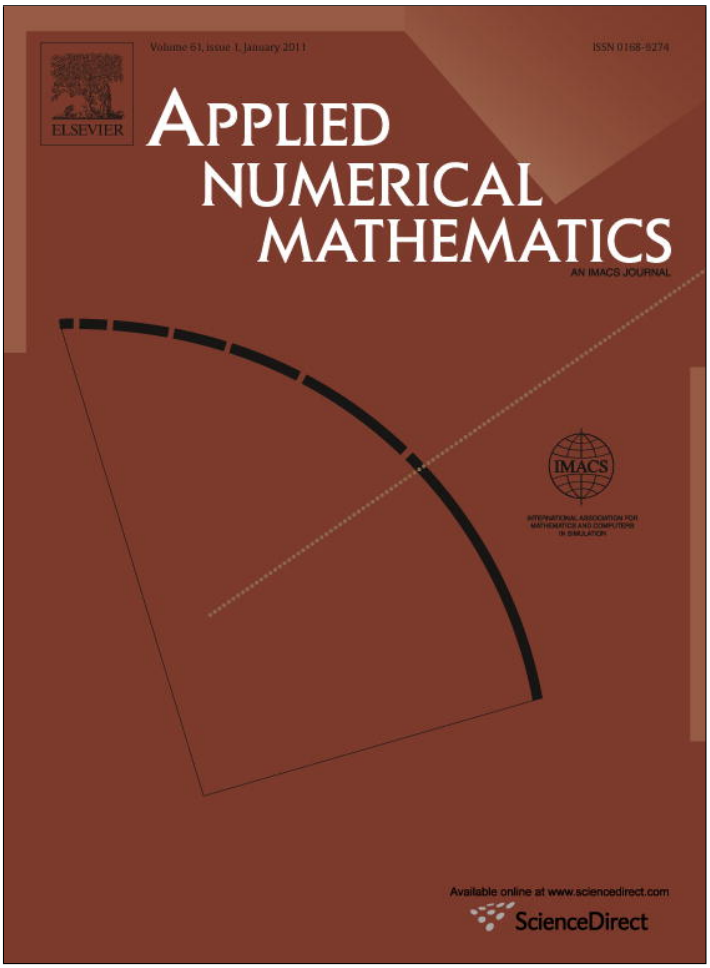

Applied Numerical Mathematics 61 (2011) 66–76

Contents lists available at ScienceDirect

Applied Numerical Mathematics

www.elsevier.com/locate/apnum

A dimensional split preconditioner for Stokes and linearized Navier–Stokes equations

Michele Benzi [a,*,1], Xue-Ping Guo [b,2]

[a] Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA
[b] Department of Mathematics, East China Normal University, Shanghai, 200062, PR China

ARTICLE INFO

Article history:
Received 23 March 2009
Received in revised form 2 August 2010
Accepted 9 August 2010
Available online 19 August 2010

Keywords:
Saddle point problems
Matrix splittings
Iterative methods
Preconditioning
Stokes problem
Oseen problem
Stretched grids

ABSTRACT

In this paper we introduce a new preconditioner for linear systems of saddle point type arising from the numerical solution of the Navier–Stokes equations. Our approach is based on a dimensional splitting of the problem along the components of the velocity field, resulting in a convergent fixed-point iteration. The basic iteration is accelerated by a Krylov subspace method like restarted GMRES. The corresponding preconditioner requires at each iteration the solution of a set of discrete scalar elliptic equations, one for each component of the velocity field. Numerical experiments illustrating the convergence behavior for different finite element discretizations of Stokes and Oseen problems are included.

1. Introduction

We consider the solution of the incompressible Navier–Stokes equations governing the flow of viscous Newtonian fluids. For an open bounded domain $\Omega \subset \mathbb{R}^d$ $(d = 2, 3)$ with boundary $\partial\Omega$, time interval $[0, T]$, and data $\mathbf{f}$, $\mathbf{g}$ and $\mathbf{u}_0$, the goal is to find a velocity field $\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$ and pressure field $p = p(\mathbf{x}, t)$ such that

$$\frac{\partial \mathbf{u}}{\partial t} - \nu\Delta\mathbf{u} + (\mathbf{u}\cdot\nabla)\mathbf{u} + \nabla p = \mathbf{f} \quad \text{on } \Omega \times (0, T], \tag{1}$$

$$\operatorname{div}\mathbf{u} = 0 \quad \text{on } \Omega \times [0, T], \tag{2}$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega \times [0, T], \tag{3}$$

$$\mathbf{u}(\mathbf{x}, 0) = \mathbf{u}_0(\mathbf{x}) \quad \text{on } \Omega, \tag{4}$$

where ν is the kinematic viscosity, Δ is the vector Laplacian, ∇ is the gradient and div the divergence. Implicit time discretization and linearization of the Navier–Stokes system by Picard fixed-point iteration result in a sequence of (generalized) Oseen problems of the form

* Corresponding author.
E-mail addresses: benzi@mathcs.emory.edu (M. Benzi), xpguo@math.ecnu.edu.cn (X.-P. Guo).

[1] The work of this author was supported in part by the National Science Foundation Grants DMS-0511336 and by a grant by the University Research Committee at Emory University.
[2] Part of this work was performed while visiting CERFACS, 42 Ave. Gaspard Coriolis, 31057 Toulouse Cedex, France.

0168-9274/$30.00 © 2010 IMACS. Published by Elsevier B.V. All rights reserved.
doi:10.1016/j.apnum.2010.08.005

$$\sigma \mathbf{u} - \nu \Delta \mathbf{u} + (\mathbf{v} \cdot \nabla)\mathbf{u} + \nabla p = \mathbf{f} \quad \text{in } \Omega, \tag{5}$$

$$\operatorname{div} \mathbf{u} = 0 \quad \text{in } \Omega, \tag{6}$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega, \tag{7}$$

where $\mathbf{v}$ is a known velocity field from a previous iteration or time step (the 'wind') and σ is proportional to the reciprocal of the time step ($\sigma = 0$ for a steady problem). When $\mathbf{v} = \mathbf{0}$ we have a (generalized) Stokes problem.

Spatial discretization of the Stokes or Oseen problem using LBB-stable finite elements (cf. [18,19]) results in large, sparse linear systems in *saddle point form*:

$$\begin{bmatrix} A & B^T \\ -B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ -g \end{bmatrix}, \quad \text{or} \quad \mathcal{A}\mathbf{x} = \mathbf{b}, \tag{8}$$

where now $\mathbf{u}$ and p represent the discrete velocity and pressure, respectively, A is the discretization of the diffusion, convection, and time-dependent terms, B^T is the discrete gradient, B the (negative) discrete divergence, and $\mathbf{f}$ and g contain forcing and boundary terms.

The efficient solution of (8) calls for rapidly convergent iterative methods. Much work has been done in developing efficient preconditioners for Krylov subspace methods applied to this problem; see, e.g., [4,5,8,14,15,17,18,22]. The ultimate goal is to develop robust solvers with optimal complexity. In particular, the rate of convergence should be independent of the mesh size h. For the Oseen problem, the rate of convergence should also depend only weakly on the kinematic viscosity ν (equivalently, on the Reynolds number $Re = O(\nu^{-1})$), although this goal is difficult to achieve in practice.

2. Dimensional splitting

For simplicity, in this paper we limit ourselves to the 2D case. Extension to the 3D case is possible (see Section 7), but will not be described here. The system matrix $\mathcal{A}$ admits the following splitting:

$$\mathcal{A} = \begin{bmatrix} A_1 & 0 & B_1^T \\ 0 & A_2 & B_2^T \\ -B_1 & -B_2 & 0 \end{bmatrix} = \begin{bmatrix} A_1 & 0 & B_1^T \\ 0 & 0 & 0 \\ -B_1 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & A_2 & B_2^T \\ 0 & -B_2 & 0 \end{bmatrix} = \mathcal{A}_1 + \mathcal{A}_2. \tag{9}$$

Here each diagonal submatrix A_i is a scalar discrete convection-diffusion-reaction operator:

$$A_i = \sigma M + \nu L + N_i \quad (i = 1, 2), \tag{10}$$

and B_1^T, B_2^T are discretizations of the partial derivatives $\frac{\partial}{\partial x}$, $\frac{\partial}{\partial y}$, respectively. Note that A_1 and A_2 act, respectively, on u (the x-component of the velocity field $\mathbf{u}$) and on v (the y-component of $\mathbf{u}$). Denoting by n_1, n_2 and m the number of degrees of freedom of u, v and p, respectively, then $A_1 \in \mathbb{R}^{n_1 \times n_1}$, $A_2 \in \mathbb{R}^{n_2 \times n_2}$, $B_1 \in \mathbb{R}^{m \times n_1}$ and $B_2 \in \mathbb{R}^{m \times n_2}$. Thus, $\mathcal{A} \in \mathbb{R}^{(n+m)\times(n+m)}$ with $n = n_1 + n_2$. In (10), M denotes the velocity mass matrix, L the discrete (negative) Laplacian, and N_i the convective terms. Note that for the discrete Stokes problem the convective terms are absent ($N_i = 0$) so that each A_i is symmetric and positive definite. For the discrete Oseen problem $A_i \neq A_i^T$, but $A_i + A_i^T$ is positive definite ($i = 1, 2$). As a consequence, $\mathcal{A}_1$ and $\mathcal{A}_2$ in (9) are nonsymmetric but positive semidefinite, in the sense that $\mathcal{A}_1 + \mathcal{A}_1^T$ and $\mathcal{A}_2 + \mathcal{A}_2^T$ are both symmetric positive semidefinite. In particular, $\mathcal{A}_1$ and $\mathcal{A}_2$ are singular. We refer to (9) as to a *dimensional splitting*, since $\mathcal{A}_1$ contains terms that correspond to the x-component of the solution and $\mathcal{A}_2$ contains terms that correspond to the y-component of the solution. Although this splitting is somewhat reminiscent of ADI (alternating direction implicit) methods [29], it is actually quite different since we do not split the operators A_i into their constituent components. To distinguish it from ADI splitting, we refer to (9) as to *dimensional splitting* (or DS for short). We further mention that our approach is also different from previous ADI schemes for saddle point problems, such as those described in [11] and [13].

Let now $\alpha > 0$ be a parameter, and denote by $\mathcal{I}$ the identity matrix of order $n_1 + n_2 + m$. Then $\mathcal{A}_1 + \alpha\mathcal{I}$ and $\mathcal{A}_2 + \alpha\mathcal{I}$ are both nonsingular, nonsymmetric, and positive definite. Consider the two splittings of $\mathcal{A}$,

$$\mathcal{A} = (\mathcal{A}_1 + \alpha\mathcal{I}) - (\alpha\mathcal{I} - \mathcal{A}_2) \quad \text{and} \quad \mathcal{A} = (\mathcal{A}_2 + \alpha\mathcal{I}) - (\alpha\mathcal{I} - \mathcal{A}_1).$$

Associated to these splittings is the alternating iteration

$$(\mathcal{A}_1 + \alpha\mathcal{I})\mathbf{x}^{k+\frac{1}{2}} = (\alpha\mathcal{I} - \mathcal{A}_2)\mathbf{x}^k + \mathbf{b}, \tag{11}$$

$$(\mathcal{A}_2 + \alpha\mathcal{I})\mathbf{x}^{k+1} = (\alpha\mathcal{I} - \mathcal{A}_1)\mathbf{x}^{k+\frac{1}{2}} + \mathbf{b} \tag{12}$$

($k = 0, 1, \ldots$). Eliminating $\mathbf{x}^{k+\frac{1}{2}}$ from these, we can rewrite (11)–(12) as the stationary scheme

$$\mathbf{x}^{k+1} = \mathcal{T}_\alpha \mathbf{x}^k + \mathbf{c}, \quad k = 0, 1, \ldots$$

where $\mathcal{T}_\alpha$ is the iteration matrix

$$\mathcal{T}_\alpha = (\mathcal{A}_2 + \alpha\mathcal{I})^{-1}(\alpha\mathcal{I} - \mathcal{A}_1)(\mathcal{A}_1 + \alpha\mathcal{I})^{-1}(\alpha\mathcal{I} - \mathcal{A}_2) \tag{13}$$

and $\mathbf{c}$ is a certain vector. As in [9], one can show that there is a unique splitting $\mathcal{A} = \mathcal{P}_\alpha - \mathcal{Q}_\alpha$ with $\mathcal{P}_\alpha$ nonsingular such that the iteration matrix $\mathcal{T}_\alpha$ is the matrix induced by that splitting, i.e., $\mathcal{T}_\alpha = \mathcal{P}_\alpha^{-1}\mathcal{Q}_\alpha = \mathcal{I} - \mathcal{P}_\alpha^{-1}\mathcal{A}$. Furthermore, $\mathbf{c} = \mathcal{P}_\alpha^{-1}\mathbf{b}$. Matrices $\mathcal{P}_\alpha$ and $\mathcal{Q}_\alpha$ are given by

$$\mathcal{P}_\alpha = \frac{1}{2\alpha}(\mathcal{A}_1 + \alpha\mathcal{I})(\mathcal{A}_2 + \alpha\mathcal{I}), \qquad \mathcal{Q}_\alpha = \frac{1}{2\alpha}(\alpha\mathcal{I} - \mathcal{A}_1)(\alpha\mathcal{I} - \mathcal{A}_2). \tag{14}$$

We refer to iteration (11)–(12) as the *DS iteration*, and to $\mathcal{P}_\alpha$ as the *DS preconditioner*. Besides the already mentioned resemblance to the Peaceman–Rachford ADI method, the DS iteration bears some resemblance to another alternating method, the Hermitian and skew-Hermitian splitting (HSS) iteration [2–4,6]. While the HSS method has proved quite successful in solving such problems as the generalized Stokes problem and the rotation form of the Navier–Stokes equations (see [6]), it is not well-suited for the standard (convection) form (1)–(2). This limitation of HSS was one of the main motivations for introducing the DS approach.

3. Convergence of the fixed-point iteration

We now prove that, under standard assumptions on the saddle point problem (8), the alternating iteration (11)–(12) converges to the solution of (8) for any choice of $\alpha > 0$ and for all initial guesses. First we state two auxiliary results. The first one is classical, and is known as *Kellogg's Lemma*; see [21]. In the following, a (not necessarily Hermitian) matrix $A \in \mathbb{C}^{n\times n}$ is said to be *positive definite* (*semidefinite*) if the Hermitian matrix $A + A^*$ is positive definite (resp., semidefinite) in the usual sense.

Lemma 1. *Let $A \in \mathbb{C}^{n\times n}$ be positive semidefinite. Then*

$$\left\|(\alpha I_n + A)^{-1}(\alpha I_n - A)\right\|_2 \leqslant 1$$

for all $\alpha > 0$. Furthermore, if A is positive definite then

$$\left\|(\alpha I_n + A)^{-1}(\alpha I_n - A)\right\|_2 < 1$$

for all $\alpha > 0$.

Lemma 2. *Assume that the $(1,1)$ block A in (8) has positive definite symmetric part and that B has full row rank. Then the following are equivalent*:

(i) *The matrix*

$$\mathcal{C}_\alpha := \begin{bmatrix} A_1 & \frac{1}{\alpha^2}B_1^T B_2 A_2 & B_1^T + \frac{1}{\alpha^2}B_1^T B_2 B_2^T \\ 0 & A_2 & B_2^T \\ -B_1 & -B_2 & 0 \end{bmatrix} \tag{15}$$

has no purely imaginary eigenvalues.

(ii) *The spectral radius of the iteration matrix $\mathcal{T}_\alpha$ in (13) is strictly less than unity.*

Proof. First we note that under the assumptions made on A and B, the matrix $\mathcal{A}$ in (8) is nonsingular; see [4, Lemma 1.1]. Let λ be an eigenvalue of the iteration matrix $\mathcal{T}_\alpha = \mathcal{I} - \mathcal{P}_\alpha^{-1}\mathcal{A}$. Then $\lambda = 1 - \mu$ where μ is a generalized eigenvalue of the matrix pencil $(\mathcal{A}, \mathcal{P}_\alpha)$; that is, there exists a vector $\mathbf{x} \neq \mathbf{0}$ such that $\mathcal{A}\mathbf{x} = \mu\mathcal{P}_\alpha\mathbf{x}$. Expanding the right-hand side, we get

$$\mathcal{A}\mathbf{x} = \frac{\mu}{2\alpha}(\mathcal{A}_1\mathcal{A}_2 + \alpha\mathcal{A} + \alpha^2\mathcal{I})\mathbf{x}.$$

Collecting terms in $\mathcal{A}$, we rewrite this as

$$\left(1 - \frac{1}{2}\mu\right)\mathcal{A}\mathbf{x} = \frac{\mu\alpha}{2}\left(\mathcal{I} + \frac{1}{\alpha^2}\mathcal{A}_1\mathcal{A}_2\right)\mathbf{x}. \tag{16}$$

Since both $\mathcal{A}$ and $\mathcal{P}_\alpha$ are nonsingular, it must be $\mu \neq 0$. Also, it must be $1 - \frac{1}{2}\mu \neq 0$ for otherwise (16) implies that $(\mathcal{I} + \frac{1}{\alpha^2}\mathcal{A}_1\mathcal{A}_2)\mathbf{x} = \mathbf{0}$ has a nonzero solution, but this is impossible since

$$\mathcal{G} := \mathcal{I} + \frac{1}{\alpha^2}\mathcal{A}_1\mathcal{A}_2 = \begin{bmatrix} I_{n_1} & -\frac{1}{\alpha^2}B_1^T B_2 & 0 \\ 0 & I_{n_2} & 0 \\ 0 & 0 & I_m \end{bmatrix}$$

is clearly nonsingular. Hence, $\mu \neq 2$ and we can set (as in [25, p. 378])

$$\theta := \frac{\mu\alpha}{2 - \mu}, \quad \text{from which } \mu = 2 - \frac{2\alpha}{\theta + \alpha} = \frac{2\theta}{\theta + \alpha}.$$

Hence, the generalized eigenproblem (16) can be reformulated as

$$\mathcal{A}\mathbf{x} = \theta\mathcal{G}\mathbf{x},$$

that is,

$$\begin{bmatrix} A_1 & 0 & B_1^T \\ 0 & A_2 & B_2^T \\ -B_1 & -B_2 & 0 \end{bmatrix} \begin{bmatrix} u \\ v \\ p \end{bmatrix} = \theta \begin{bmatrix} I_{n_1} & -\frac{1}{\alpha^2} B_1^T B_2 & 0 \\ 0 & I_{n_2} & 0 \\ 0 & 0 & I_m \end{bmatrix} \begin{bmatrix} u \\ v \\ p \end{bmatrix},$$

or

$$\mathcal{G}^{-1}\mathcal{A}\mathbf{x} = \theta\mathbf{x}, \quad \text{where } \mathcal{G}^{-1} = \begin{bmatrix} I_{n_1} & \frac{1}{\alpha^2} B_1^T B_2 & 0 \\ 0 & I_{n_2} & 0 \\ 0 & 0 & I_m \end{bmatrix}.$$

This eigenproblem is precisely $\mathcal{C}_\alpha\mathbf{x} = \theta\mathbf{x}$, where $\mathcal{C}_\alpha = \mathcal{G}^{-1}\mathcal{A}$ is the matrix in (15). Note that $\mathcal{C}_\alpha$ is necessarily nonsingular. Recall now that the eigenvalues of $\mathcal{P}_\alpha^{-1}\mathcal{A}$ are of the form $\mu = 2\theta/(\theta+\alpha)$. It must be $|1-\mu| \leqslant 1$, since $\lambda = 1-\mu$ is an eigenvalue of the iteration matrix

$$\mathcal{T}_\alpha = \mathcal{I} - \mathcal{P}_\alpha^{-1}\mathcal{A} = (\alpha\mathcal{I} + \mathcal{A}_2)^{-1}(\alpha\mathcal{I} - \mathcal{A}_1)(\alpha\mathcal{I} + \mathcal{A}_1)^{-1}(\alpha\mathcal{I} - \mathcal{A}_2),$$

which is similar to

$$\widehat{\mathcal{T}}_\alpha = (\alpha\mathcal{I} - \mathcal{A}_1)(\alpha\mathcal{I} + \mathcal{A}_1)^{-1}(\alpha\mathcal{I} - \mathcal{A}_2)(\alpha\mathcal{I} + \mathcal{A}_2)^{-1},$$

hence $\varrho(\mathcal{T}_\alpha)$, the spectral radius of $\mathcal{T}_\alpha$, satisfies

$$\varrho(\mathcal{T}_\alpha) = \varrho(\widehat{\mathcal{T}}_\alpha) \leqslant \big\|(\alpha\mathcal{I} - \mathcal{A}_1)(\alpha\mathcal{I} + \mathcal{A}_1)^{-1}\big\|_2 \big\|(\alpha\mathcal{I} - \mathcal{A}_2)(\alpha\mathcal{I} + \mathcal{A}_2)^{-1}\big\|_2 \leqslant 1. \tag{17}$$

The last inequality is an immediate consequence of Kellogg's Lemma.

Denoting by $\Re(\theta)$ and $\Im(\theta)$ the real and imaginary parts of θ, respectively, we claim that $\varrho(\mathcal{T}_\alpha) < 1$ if and only if $\Re(\theta) \neq 0$ for every eigenvalue θ of $\mathcal{C}_\alpha$; equivalently, $\varrho(\mathcal{T}_\alpha) = 1$ if and only if there exists at least one θ with $\Re(\theta) = 0$. Indeed, we have

$$|1-\mu| = 1 \quad \Leftrightarrow \quad \left|\frac{2\theta}{\theta+\alpha} - 1\right| = 1 \quad \Leftrightarrow \quad \left|\frac{\theta-\alpha}{\theta+\alpha}\right| = 1 \quad \Leftrightarrow \quad |\theta-\alpha| = |\theta+\alpha|.$$

The last equality can be rewritten as $(\Re(\theta)-\alpha)^2 + \Im(\theta)^2 = (\Re(\theta)+\alpha)^2 + \Im(\theta)^2$, or

$$\big(\Re(\theta)-\alpha\big)^2 = \big(\Re(\theta)+\alpha\big)^2 \quad \Leftrightarrow \quad 4\alpha\Re(\theta) = 0 \quad \Leftrightarrow \quad \Re(\theta) = 0.$$

Therefore, $\varrho(\mathcal{T}_\alpha) = 1$ if and only if $\mathcal{C}_\alpha$ has at least one purely imaginary eigenvalue. The proof is complete. □

We are now in a position to prove the following convergence result.

Theorem 3. *Under the assumptions of Lemma* 2*, the iteration* (11)–(12) *is unconditionally convergent*; *that is,* $\varrho(\mathcal{T}_\alpha) < 1$ *for all* $\alpha > 0$.

Proof. By Lemma 2, it suffices to show that for all $\alpha > 0$, the matrix $\mathcal{C}_\alpha$ in (15) has no purely imaginary eigenvalues. We will argue by contradiction. Recall that $\mathcal{C}_\alpha$ is nonsingular. Thus, let $\theta \neq 0$ be an eigenvalue of $\mathcal{C}_\alpha$ corresponding to an eigenvector $\mathbf{x} = [u; v; p]$, where $u \in \mathbb{C}^{n_1}$, $v \in \mathbb{C}^{n_2}$ and $p \in \mathbb{C}^m$ are not all equal to zero. Expanding $\mathcal{C}_\alpha\mathbf{x} = \theta\mathbf{x}$ we obtain

$$A_1 u + \frac{1}{\alpha^2} B_1^T B_2 A_2 v + B_1^T p + \frac{1}{\alpha^2} B_1^T B_2 B_2^T p = \theta u, \tag{18}$$

$$A_2 v + B_2^T p = \theta v, \tag{19}$$

$$-B_1 u - B_2 v = \theta p. \tag{20}$$

Assuming that the eigenvector $\mathbf{x}$ has been normalized so that $\|\mathbf{x}\|_2 = 1$, we have

$$\theta = \mathbf{x}^*\mathcal{C}_\alpha\mathbf{x}, \qquad \bar{\theta} = \mathbf{x}^*\mathcal{C}_\alpha^T\mathbf{x} \quad \text{and} \quad \Re(\theta) = \frac{\theta+\bar{\theta}}{2} = \frac{1}{2}\mathbf{x}^*\big(\mathcal{C}_\alpha + \mathcal{C}_\alpha^T\big)\mathbf{x}.$$

Therefore, letting $H_1 = (A_1 + A_1^T)/2$ and $H_2 = (A_2 + A_2^T)/2$, we find after some easy algebraic manipulations

$$\Re(\theta) = u^* H_1 u + v^* H_2 v + \frac{1}{2\alpha^2}\big[u^* B_1^T B_2\big(A_2 v + B_2^T p\big) + \big(v^* A_2^T + p^* B_2\big)B_2^T B_1 u\big].$$

First we observe that at least one between u and v must be nonzero, for otherwise (20), together with the fact that $\theta \neq 0$, implies $p = 0$ and thus $\mathbf{x} = \mathbf{0}$, a contradiction. Since H_1 and H_2 are symmetric positive definite we have that $u^* H_1 u + v^* H_2 v > 0$, therefore

$$\Re(\theta) = 0 \quad \Rightarrow \quad u^* B_1^T B_2 \big(A_2 v + B_2^T p\big) + \big(v^* A_2^T + p^* B_2\big) B_2^T B_1 u < 0, \tag{21}$$

showing that if θ is purely imaginary, it must necessarily be $u \neq 0$. Next, consider the case where $v = 0$. Then Eqs. (18)–(20) reduce to

$$A_1 u + B_1^T p = \theta u, \tag{22}$$

$$B_2^T p = 0, \tag{23}$$

$$-B_1 u = \theta p, \tag{24}$$

that is, $\mathcal{A}\mathbf{x} = \theta \mathbf{x}$ where $\mathbf{x} = [u; 0; p] \neq \mathbf{0}$. Hence, θ is an eigenvalue of $\mathcal{A}$; but then $\Re(\theta) > 0$ by virtue of Lemma 1.1 in [4]. So it must be $v \neq 0$. Furthermore, if $p = 0$, then Eq. (19) becomes $A_2 v = \theta v$, hence $\Re(\theta) > 0$ since A_2 is positive definite (has positive definite symmetric part) by assumption, and therefore all its eigenvalues have positive real part. Thus, it must be $u \neq 0$, $v \neq 0$ and $p \neq 0$. Now, using (19) we rewrite the necessary condition in (21) in the form

$$\theta u^* B_1^T B_2 v + \bar{\theta} v^* B_2^T B_1 u = \theta (B_1 u)^* B_2 v + \bar{\theta} (B_2 v)^* B_1 u < 0. \tag{25}$$

Now, from (20) we obtain $B_2 v = -B_1 u - \theta p$ which substituted into (25) yields

$$\theta u^* B_1^T (-B_1 u - \theta p) + \bar{\theta}\big(-u^* B_1^T - \bar{\theta} p^*\big) B_1 u < 0,$$

or, equivalently,

$$-\theta u^* B_1^T B_1 u - \theta^2 u^* B_1^T p - \bar{\theta} u^* B_1^T B_1 u - \bar{\theta}^2 p^* B_1 u < 0.$$

Now, if $\Re(\theta) = 0$ then $\theta = \mathrm{i}\xi$ for some $\xi \in \mathbb{R}$, $\xi \neq 0$, where $\mathrm{i} = \sqrt{-1}$. After simplification, we find

$$-\theta u^* B_1^T B_1 u - \theta^2 u^* B_1^T p - \bar{\theta} u^* B_1^T B_1 u - \bar{\theta}^2 p^* B_1 u = \xi^2 \big(u^* B_1^T p + p^* B_1 u\big),$$

therefore condition (21) becomes

$$u^* B_1^T p + p^* B_1 u < 0$$

and since $u^* B_1^T p + p^* B_1 u = 2\Re(u^* B_1^T p)$, we conclude that

$$\Re(\theta) = 0 \quad \Rightarrow \quad \Re\big(u^* B_1^T p\big) < 0.$$

Likewise, from (20) we obtain $B_1 u = -B_2 v - \theta p$; substituting this into (25) and going through the same algebraic operations as before, we also find that

$$\Re(\theta) = 0 \quad \Rightarrow \quad \Re\big(v^* B_2^T p\big) < 0.$$

Therefore,

$$\Re\big((B_1 u + B_2 v)^* p\big) = \Re\big(\big(u^* B_1^T + v^* B_2^T\big) p\big) < 0,$$

but together with (20) this implies $\Re((-\theta p)^* p) < 0$, or $\Re(\mathrm{i}\xi \|p\|_2^2) < 0$, which is clearly absurd since $\mathrm{i}\xi \|p\|_2^2$ is imaginary. This proves that $\mathcal{C}_\alpha$ cannot have purely imaginary eigenvalues. □

The restriction in Theorem 3 that A have positive definite symmetric part is not essential. If $A + A^T$ is only positive semidefinite (and singular), the alternating iteration (11)–(12) is still well defined. Moreover, the spectral radius of the iteration matrix cannot exceed 1. Indeed, the iteration matrix $\mathcal{T}_\alpha$ still satisfies $\varrho(\mathcal{T}_\alpha) \leqslant 1$, and if the symmetric part of A and B have no null vectors in common, the coefficient matrix $\mathcal{A}$ in (8) is still nonsingular; see again Lemma 1.1 in [4]. Hence, 1 is not an eigenvalue of $\mathcal{T}_\alpha = \mathcal{I} - \mathcal{P}_\alpha^{-1}\mathcal{A}$. However, it may happen that $\varrho(\mathcal{T}_\alpha) = 1$ for some choices of $\alpha > 0$. A simple example is given by

$$\mathcal{A} = \left[\begin{array}{cc|c} I & 0 & 0 \\ 0 & 0 & I \\ \hline 0 & -I & 0 \end{array}\right].$$

Note that this matrix is nonsingular. It is easy to see that for $\alpha = 1$, the iteration matrix $\mathcal{T}_\alpha$ has only three distinct eigenvalues: $\lambda = 0$, $\lambda = \mathrm{i}$ and $\lambda = -\mathrm{i}$. Hence, the spectral radius is 1. Nevertheless, a simple modification of the basic algorithm yields a convergent iteration. To this end, recall that $\varrho(\mathcal{T}_\alpha) \leqslant 1$ for all $\alpha > 0$; see (17). Let $\gamma \in (0, 1)$ be a parameter, then the matrix $(1 - \gamma)\mathcal{I} + \gamma \mathcal{T}_\alpha$ has spectral radius less than 1 for all $\alpha > 0$. Indeed, the eigenvalues of $(1 - \gamma)\mathcal{I} + \gamma \mathcal{T}_\alpha$ are of the form $1 - \gamma + \gamma\lambda$, where λ is an eigenvalue of $\mathcal{T}_\alpha$. It is easy to see that since $|\lambda| \leqslant 1$ and $\lambda \neq 1$, all the quantities

$1-\gamma+\gamma\lambda$ have magnitude strictly less than 1. In practice, however, this modification is seldom used. In the next section we discuss Krylov subspace acceleration, which is much more effective and is applicable whether or not $\varrho(\mathcal{T}_\alpha) < 1$. Nevertheless, knowing that $\varrho(\mathcal{T}_\alpha) < 1$ is useful because it implies that the spectrum of the preconditioned matrix lies entirely in the right half-plane, a desirable property for Krylov subspace acceleration. Moreover, the smaller is $\varrho(\mathcal{T}_\alpha)$, the more clustered the spectrum of the preconditioned matrix is around 1.

4. Krylov subspace acceleration

The basic method (11)–(12) although unconditionally convergent, is not competitive as a solver for problem (8), mainly due to the fact that convergence is generally slow. Fortunately, the rate of convergence can be greatly improved by Krylov subspace acceleration. In other words, $\mathcal{P}_\alpha$ can be used as a preconditioner for GMRES [24] or any other nonsymmetric Krylov method. It should be noted that when $\mathcal{P}_\alpha$ is used as a preconditioner, the pre-factor $\frac{1}{2\alpha}$ in (14) is irrelevant and can be neglected. In this paper we use the restarted GMRES algorithm with restart parameter m. Preconditioning is applied on the right.

The rate of convergence of nonsymmetric Krylov iterations (like GMRES) preconditioned by $\mathcal{P}_\alpha$ depends on the particular choice of α. Finding the value of α that optimizes the rate of convergence appears to be a difficult problem in general. Indeed, in practice the convergence rate depends to a large extent on the size, shape, and location of the entire spectrum of the preconditioned matrix $\mathcal{P}_\alpha^{-1}\mathcal{A}$, and not just on the spectral radius of $\mathcal{T}_\alpha = \mathcal{I} - \mathcal{P}_\alpha^{-1}\mathcal{A}$. (The rate of convergence may also be affected by the conditioning of the eigenbasis of the preconditioned matrix, but this is usually difficult to estimate; see, e.g., [26, p. 17].) Numerical experiments (see below) suggest that the value (or values) α_* of α for which the number of preconditioned iterations is minimized is a rather small number ($0 < \alpha_* \ll 1$). Moreover, the convergence rate is not overly sensitive to small relative changes in α.

5. Implementation aspects

For the proposed approach to be successful, it is imperative that the action of the DS preconditioner be computed efficiently within each GMRES iteration. Written out explicitly, system (11) reads

$$\begin{bmatrix} A_1+\alpha I_{n_1} & 0 & B_1^T \\ 0 & \alpha I_{n_2} & 0 \\ -B_1 & 0 & \alpha I_{n_3} \end{bmatrix} \begin{bmatrix} u^{k+\frac{1}{2}} \\ v^{k+\frac{1}{2}} \\ p^{k+\frac{1}{2}} \end{bmatrix} = \begin{bmatrix} \alpha u^k + f_1 \\ (\alpha I_{n_2} - A_2)v^k - B_2^T p^k + f_2 \\ B_2 v^k + \alpha p^k - g \end{bmatrix}, \tag{26}$$

while system (12) becomes

$$\begin{bmatrix} \alpha I_{n_1} & 0 & 0 \\ 0 & A_2+\alpha I_{n_2} & B_2^T \\ 0 & -B_2 & \alpha I_{n_3} \end{bmatrix} \begin{bmatrix} u^{k+1} \\ v^{k+1} \\ p^{k+1} \end{bmatrix} = \begin{bmatrix} (\alpha I_{n_1} - A_1)u^{k+\frac{1}{2}} - B_1^T p^{k+\frac{1}{2}} + f_1 \\ \alpha v^{k+\frac{1}{2}} + f_2 \\ B_1 u^{k+\frac{1}{2}} + \alpha p^{k+\frac{1}{2}} - g \end{bmatrix}. \tag{27}$$

Both systems (26) and (27) are highly reducible. Indeed, the second equation in (26) immediately yields $v^{k+\frac{1}{2}} = \frac{1}{\alpha}[(\alpha I_{n_2} - A_2)v^k - B_2^T p^k + f_2]$ together with the reduced system

$$\begin{bmatrix} A_1+\alpha I_{n_1} & B_1^T \\ -B_1 & \alpha I_{n_3} \end{bmatrix} \begin{bmatrix} u^{k+\frac{1}{2}} \\ p^{k+\frac{1}{2}} \end{bmatrix} = \begin{bmatrix} \alpha u^k + f_1 \\ B_2 v^k + \alpha p^k - g \end{bmatrix}. \tag{28}$$

Likewise, system (27) is equivalent to $u^{k+1} = \frac{1}{\alpha}[(\alpha I_{n_1} - A_1)u^{k+\frac{1}{2}} - B_1^T p^{k+\frac{1}{2}} + f_1]$ together with the reduced system

$$\begin{bmatrix} A_2+\alpha I_{n_2} & B_2^T \\ -B_2 & \alpha I_{n_3} \end{bmatrix} \begin{bmatrix} v^{k+1} \\ p^{k+1} \end{bmatrix} = \begin{bmatrix} \alpha v^{k+\frac{1}{2}} + f_2 \\ B_1 u^{k+\frac{1}{2}} + \alpha p^{k+\frac{1}{2}} - g \end{bmatrix}. \tag{29}$$

Both systems (28) and (29) can be further reduced. Let $c^k = \alpha u^k + f_1$ and $d^k = B_2 v^k + \alpha p^k - g$. For (28), the second equation yields

$$p^{k+\frac{1}{2}} = \frac{1}{\alpha}\left(d^k + B_1 u^{k+\frac{1}{2}}\right) \tag{30}$$

which, substituted in the first one, yields

$$\left(A_1 + \alpha I_{n_1} + \frac{1}{\alpha} B_1^T B_1\right) u^{k+\frac{1}{2}} = c^k - \frac{1}{\alpha} B_1^T d^k. \tag{31}$$

Once this equation has been solved for $u^{k+\frac{1}{2}}$, we recover $p^{k+\frac{1}{2}}$ from (30). Similarly, let $c^{k+\frac{1}{2}} = \alpha v^{k+\frac{1}{2}} + f_2$ and $d^{k+\frac{1}{2}} = B_1 u^{k+\frac{1}{2}} + \alpha p^{k+\frac{1}{2}} - g$, then the second equation of (29) yields

$$p^{k+1} = \frac{1}{\alpha}\left(d^{k+\frac{1}{2}} + B_2 v^{k+1}\right) \tag{32}$$

which, substituted in the first one, yields

$$\left(A_2 + \alpha I_{n_2} + \frac{1}{\alpha} B_2^T B_2\right) v^{k+1} = c^{k+\frac{1}{2}} - \frac{1}{\alpha} B_2^T d^{k+\frac{1}{2}}. \tag{33}$$

Once this equation has been solved, the new value p^{k+1} of the pressure can be obtained from (32).

Hence, the bulk of the work in applying the preconditioner is in the solution of the two reduced systems (31) and (33). Each of these is a discrete analogue of a scalar, second-order, elliptic, anisotropic convection–diffusion–reaction equation. The anisotropy is in a sense artificial, since it depends on the size of the algorithmic parameter α: the smaller α is, the stronger the anisotropy in the diffusion terms. Note that for the Stokes and generalized Stokes problems, the convection terms are missing in (31) and (33) and the coefficient matrices are symmetric and positive definite. Some remarks on the solution of these two subsystems are given in the next section.

On the basis of the foregoing discussion, it is clear that the DS approach can be regarded as a *dimensionally segregated method*, i.e., a method where the values of the velocity components u and v (or u, v and w in 3D) are updated separately through a decoupling process; the new value of the pressure p is obtained at very low cost from the new velocity values.

We conclude this section with a discussion of diagonal scaling. In the alternating iteration (11)–(12), it is possible to replace the $(n+m)$-by-$(n+m)$ identity matrix $\mathcal{I}$ with an arbitrary symmetric positive definite matrix $\mathcal{D}$, leading to a preconditioner of the form

$$\widehat{\mathcal{P}}_\alpha = \frac{1}{2\alpha}(\mathcal{A}_1 + \alpha\mathcal{D})(\mathcal{A}_2 + \alpha\mathcal{D}).$$

It is easy to check that this is equivalent to applying the original DS preconditioner $\mathcal{P}_\alpha$ to the linear system $\widehat{\mathcal{A}}\hat{\mathbf{x}} = \hat{\mathbf{b}}$, where $\widehat{\mathcal{A}} := \mathcal{D}^{-\frac{1}{2}}\mathcal{A}\mathcal{D}^{-\frac{1}{2}}$, $\hat{\mathbf{x}} = \mathcal{D}^{-\frac{1}{2}}\mathbf{x}$, and $\hat{\mathbf{b}} = \mathcal{D}^{-\frac{1}{2}}\mathbf{b}$. A natural choice for finite element problems is to take $\mathcal{D}$ to be a block diagonal matrix with the velocity and pressure mass matrices on the main diagonal. In order to reduce the cost of applying the preconditioner, the mass matrices can be lumped or replaced by their main diagonals. In this paper we form $\mathcal{D}$ from the diagonals of the mass matrices. We found that this scaling, which is used in all numerical experiments in the next section, is quite effective in improving the convergence rate of DS preconditioning, especially for problems on stretched grids. Incidentally, it was noted in [4] that diagonal scaling is very beneficial for the HSS preconditioner as well.

6. Numerical experiments

In this section we report the results of numerical experiments on linear systems from Stokes and Oseen models of incompressible flow. We consider Q2–Q1 finite element discretizations of two standard model problems: the leaky-lid driven cavity problem, and the backward facing step problem (see [18]). For the driven cavity problem, we consider both uniform and stretched grids of increasing size. All test problems were generated under Matlab using the IFISS software package [16] (see also [18]). We use the DS preconditioner in conjunction with restarted GMRES with $m = 30$ as the restart. In all cases the initial guess was the zero vector, and the stopping criterion was a reduction of at least six orders of magnitude of the initial residual norm. We discuss experiments for both steady and unsteady cases.

6.1. Steady problems

Application of the DS preconditioner requires, at each iteration, the solution of the linear systems (31) and (33), where the general form of A_i is given in (10). For the steady Stokes problems, $\sigma = 0$ and $N_i = 0$; hence, the coefficient matrices in these two systems are symmetric positive definite. The systems can be solved very efficiently with a sparse Cholesky factorization with an approximate minimum degree (AMD) ordering; see [1,12]. The factorization is computed once and for all at the outset, and only forward and backward triangular solves need to be performed at each GMRES iteration. For the steady Oseen problem ($\sigma = 0$, $N_i \neq 0$ in (10)) the two systems (31) and (33) are nonsymmetric, although structurally symmetric. We compute sparse LU factorizations [12] using again an AMD reordering. These direct methods are much faster than iterative methods in the case of 2D problems; in the solution of large 3D problems iterative methods will have to be used instead, necessitating the use of a flexible Krylov method (like flexible GMRES [23]) for the outer iteration.

The first set of experiments is aimed at assessing the performance of the DS preconditioner on steady Stokes and Oseen problems, in particular to investigate the dependence on the discretization parameter h and on the viscosity ν. We begin with the leaky-lid driven cavity problem.

In Table 1 we show iteration counts for DS-preconditioned GMRES (30) applied to the steady Stokes problem on a sequence of uniform grids. We report results for the optimal choice of α, determined experimentally. We see that the iteration count is independent of mesh size. Our tests show that for the Stokes problem, small changes in the value of α do not have a dramatic effect on the number of iterations.

In Table 2 we report iteration counts for the steady Oseen problem on a sequence of uniform grids and for different values of ν, using optimal or near-optimal values of α. We found that for small values of the viscosity, scaling (as described

Table 1
Preconditioned GMRES iterations for Stokes problem with the optimal α.

Grid	Its	α_{opt}
16×16	11	0.006
32×32	12	0.001
64×64	12	0.0006
128×128	10	0.0002

Table 2
Preconditioned GMRES iterations for Oseen problem for different values of ν, optimal α (lid driven cavity, uniform grids).

Grid	$\nu = 0.1$	$\nu = 0.01$	$\nu = 0.001$
16×16	14	19	44
32×32	15	20	57
64×64	15	19	50
128×128	14	18	39

Table 3
Preconditioned GMRES iterations for Oseen problem for different values of ν, optimal α (backward facing step, uniform grids).

Grid	$\nu = 0.1$	$\nu = 0.01$	$\nu = 0.005$
16×48	19	21	22
32×96	19	22	27
64×192	20	24	30
128×384	19	25	32

Table 4
Preconditioned GMRES iterations for Stokes problem with the optimal α (lid driven cavity, stretched grids).

Grid	Its	α_{opt}
16×16	9	0.2
32×32	9	0.3
64×64	9	0.3
128×128	9	0.4

Table 5
Preconditioned GMRES iterations for Oseen problem for different values of ν, optimal α (lid driven cavity, stretched grids).

Grid	$\nu = 0.1$	$\nu = 0.01$	$\nu = 0.001$
16×16	14	30	137
32×32	14	37	166
64×64	14	39	177
128×128	14	39	189

in the previous section) dramatically improves the rate of convergence of preconditioned GMRES iteration. For example, for $\nu = 0.001$ the preconditioned iteration without scaling requires over 200 iterations on the fine grid. One can clearly see again that DS preconditioning results in h-independent convergence rates. There is a mild dependence of the rate of convergence on the viscosity. Note, however, that the convergence rate on the finest grid remains excellent.

Results for the backward facing step problem (Oseen only) are presented in Table 3. Note that for this problem, the number of cells is different in the horizontal and vertical directions. Here the smallest value of the viscosity we consider is $\nu = 0.005$, since the flow is unsteady for $\nu \approx 0.001$. The experiments show a fairly robust convergence behavior with respect to both h and ν.

Next, we present some results using stretched grids, using the default stretch factors provided by IFISS. These are 1.2712 for the 16×16 grid, 1.1669 for the 32×32 grid, 1.0977 for the 64×64 grid, and 1.056 for the 128×128 grid. The stretching is done in both the horizontal and vertical directions starting at the center of the domain, resulting in rather fine grids near the boundaries. In practice, stretched grids of this kind are often used to resolve boundary layers, if present. However, using stretched grids typically results in linear systems that are considerably more difficult to solve with iterative methods. Tables 4–5 report results on numerical experiments for the driven cavity Stokes and Oseen problems (respectively) discretized on a sequence of stretched grids. Clearly, the results for the Stokes problem are very good—indeed, even better than on the uniform grids. It is important to mention that without the diagonal scaling (using the diagonals of the mass matrices), the

Table 6
Preconditioned GMRES iterations for generalized Stokes driven cavity problem ($\nu = 1$, $\sigma = h^{-1}$) with the optimal α.

Grid	Its	α_{opt}
16×16	11	0.006
32×32	12	0.001
64×64	11	0.0005
128×128	12	0.0001

Table 7
Preconditioned GMRES iterations for generalized Oseen driven cavity problem ($\sigma = h^{-1}$) for different values of ν, optimal α.

Grid	$\nu = 0.1$	$\nu = 0.01$	$\nu = 0.001$
16×16	12	13	13
32×32	14	15	15
64×64	16	18	19
128×128	19	22	23

convergence behavior is much worse; in particular, the number of iterations increases as the mesh is refined. Also note that the optimal value of α is much larger now than for the uniform grid case. Table 5 shows that mesh independent convergence is observed also for the steady Oseen problem, although there is a noticeable dependence on the viscosity. We note that without scaling with the mass matrices, the rate of convergence deteriorates very rapidly for decreasing mesh size *and* viscosity. Hence, diagonal scaling helps mitigate the negative impact of stretched meshes and removes the dependence on mesh size; however, it does not appear to be enough to achieve robustness for very small values of ν. In our experience, this kind of degradation in solver performance is observed also with other preconditioners when stretched grids are used.

6.2. Unsteady problems

Next, we report on analogous experiments involving the generalized Stokes problem (with $\nu = 1$) and the generalized Oseen problem (for several values of ν). A sequence of linear systems of this type needs to be solved when the time-dependent Stokes or Navier–Stokes equations are integrated numerically using implicit time-stepping schemes. Now the matrices A_1 and A_2 in (31)–(33) are of the form (10) where $\sigma = h^{-1}$ and M is the velocity mass matrix; also, $N_i = 0$ for generalized Stokes and $N_i \neq 0$ for generalized Oseen. For brevity, we only consider the driven cavity problem on square meshes. As one can see from Table 6, for the generalized Stokes problem the results are virtually the same as those obtained in the steady case. The behavior, however, is somewhat different for the generalized Oseen problem. Indeed, we can see from the results in Table 7 that the rate of convergence for DS-preconditioned GMRES(20) is essentially independent of viscosity, while showing a mild dependence on h.

6.3. Choosing α

As with all parameter dependent preconditioners, some guidelines need to be provided for the choice of α. The analytic determination of the value of α which results in the fastest convergence of the preconditioned GMRES iteration appears to be quite difficult, especially in the case of the Oseen problem. Our numerical experiments indicate that for problems posed on uniform meshes, α should be taken very small; for the Q2–Q1 discretization used in our experiments, the best value of α is often of the order of 10^{-3} or even smaller. Moreover, the best value of α gets smaller as the mesh is refined. Note that taking too small a value of α can lead to excessive ill-conditioning in the subsystems (31)–(33) to be solved at each GMRES iteration; in our experiments, however, this was never a problem. For problems posed on stretched meshes, on the other hand, our experiments show that with proper diagonal scaling the optimal α is often of the order of 10^{-1}.

A possible rule of thumb, applicable in case of uniform meshes, is to tie α to the discretization parameter h. Taking $\alpha \approx h^2$ was found to give pretty good results in most cases, at least for h small enough.

In the case of a uniform mesh, the rate of convergence of DS-preconditioned GMRES does not appear to be overly sensitive to the choice of α, in the sense that small relative changes in the size of α do not usually cause the number of iterations to change too drastically. In Fig. 1 we show the total number of GMRES and GMRES(30) iterations for the solution of the steady Oseen problem (with $\nu = 0.01$) as a function of α, for two choices of h. Fig. 2 displays the corresponding data for the unsteady case. From these plots, it appears that slightly overestimating the parameter α does not lead to drastic changes in the number of iterations, especially in the steady case. Underestimating the optimal α can be more harmful, but this is easy to avoid. One way to do this is to find a (near-)optimal value of α on a coarse grid, and then use the same value of α on finer grids. Since the optimal α tends to decrease as the grid id refined, this strategy will generally overestimate the optimal α. This strategy, again, assumes uniform meshes are used.

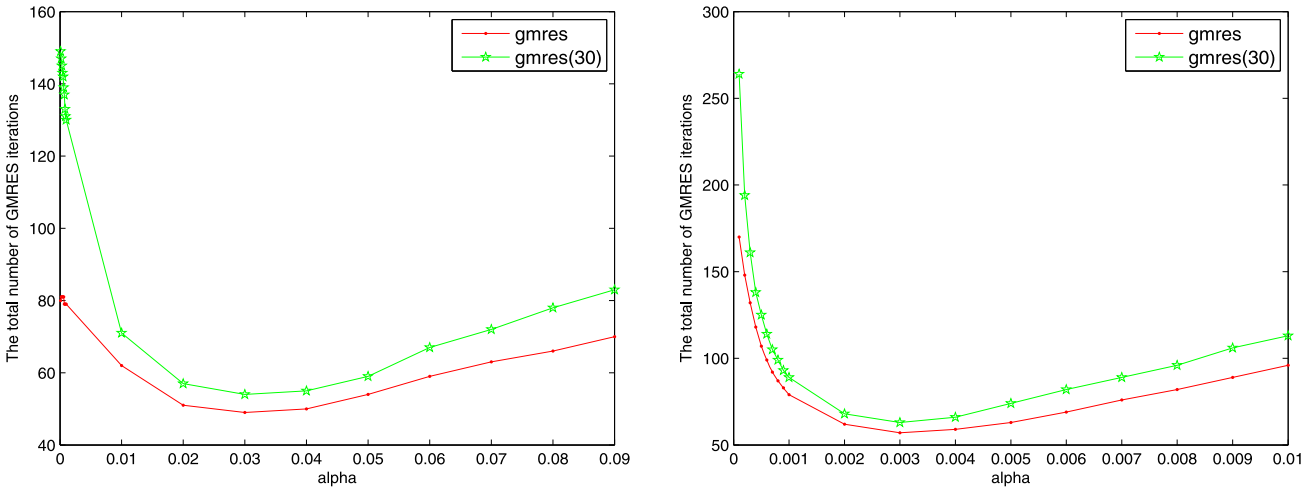

(a) 16×16 grid (b) 64×64 grid

Fig. 1. Number of GMRES iterations vs. the value of α, steady Oseen problem, $\nu = 0.01$. Uniform mesh.

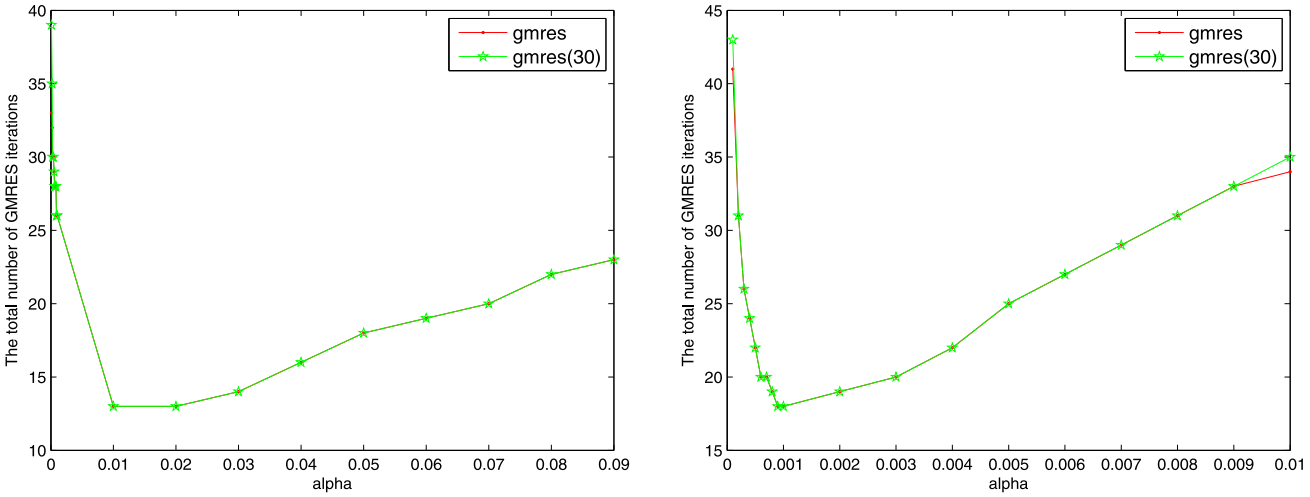

(a) 16×16 grid (b) 64×64 grid

Fig. 2. Number of GMRES iterations vs. the value of α, unsteady Oseen problem, $\nu = 0.01$. Uniform mesh.

7. Conclusions and future work

In this paper we have introduced a "dimensional splitting" approach for the solution of saddle point systems in which the $(1,1)$ block can be partitioned into a two-by-two block diagonal structure. Saddle point systems of this kind arise in a number of applications: in this paper we focused on linear systems arising from the discretization of two-dimensional incompressible flow problems. We have established the convergence of the fixed-point iteration, and investigated experimentally its use as a preconditioner for restarted GMRES on a set of Stokes and Oseen problems (both steady and unsteady) discretized by Q2–Q1 finite elements on both uniform meshes and stretched meshes. The numerical experiments indicate that for uniform meshes, the preconditioner results in fast convergence for a wide range of mesh sizes and Reynolds numbers. Diagonal scaling was found to be essential to achieve robustness for small values of the viscosity. For stretched meshes, on the other hand, we found that diagonal scaling (using the diagonals of the velocity and pressure mas matrices) is absolutely necessary to retain mesh-independent convergence for steady Stokes and Oseen problems. On stretched meshes, however, the performance of the preconditioner deteriorates for small values of the viscosity. This problem occurs with other preconditioners as well.

Future work should include further analysis of the preconditioned iteration, including using Local Fourier Analysis [28,3] for estimating the optimal value of the relaxation parameter α, and extension to the 3D case. We observe here that the basic alternating iteration (11)–(12) is of *Peaceman–Rachford type* and cannot be directly extended to the case of three splittings. Extension to the 3D case requires the alternating iteration to be of *Douglas–Rachford type*; see [27, pp. 244–245]. From the viewpoint of implementation, the 3D case necessitates using (inexact) inner iterative solves for the subproblems that occur in the application of the preconditioner. The effect of inexact solves on the performance of the DS preconditioner needs to be investigated. We mention that a promising approach for solving systems of the type (31)–(33) was recently described

in [10]. Another possibility would be to use the scalable algebraic multilevel solvers for scalar elliptic PDEs provided in the state-of-the-art Trilinos software package [20].

Finally, we mention that the method presented in this paper forms the starting point for a new preconditioning scheme, called *Relaxed Dimensional Factorization*, which is currently under development; see [7].

Acknowledgements

We would like to thank Zhen Wang and Qiang Niu for useful comments and for help with some of the numerical experiments. We are also indebted to Daniel Szyld and two anonymous referees for constructive criticism that resulted in several improvements.

References

[1] P.R. Amestoy, T.A. Davis, I.S. Duff, An approximate minimum degree ordering algorithm, SIAM J. Matrix Anal. Appl. 17 (1996) 886–905.
[2] Z.Z. Bai, G.H. Golub, M.K. Ng, Hermitian and skew-Hermitian splitting methods for non-Hermitian positive definite linear systems, SIAM J. Matrix Anal. Appl. 24 (2003) 603–626.
[3] M. Benzi, M.J. Gander, G.H. Golub, Optimization of the Hermitian and skew-Hermitian splitting iteration for saddle-point problems, BIT Numer. Math. 43 (2003) 881–900.
[4] M. Benzi, G.H. Golub, A preconditioner for generalized saddle point problems, SIAM J. Matrix Anal. Appl. 26 (2004) 20–41.
[5] M. Benzi, G.H. Golub, J. Liesen, Numerical solution of saddle point problems, Acta Numer. 14 (2005) 1–137.
[6] M. Benzi, J. Liu, An efficient solver for the incompressible Navier–Stokes equations in rotation form, SIAM J. Sci. Comput. 29 (2007) 1959–1981.
[7] M. Benzi, M.K. Ng, Q. Niu, Z. Wang, A relaxed dimensional factorization preconditioner for the incompressible Navier–Stokes equations, Technical Report TR-2010-010, Department of Mathematics and Computer Science, Emory University, 2010.
[8] M. Benzi, M.A. Olshanskii, An augmented Lagrangian-based approach to the Oseen problem, SIAM J. Sci. Comput. 28 (2006) 2095–2113.
[9] M. Benzi, D.B. Szyld, Existence and uniqueness of splittings for stationary iterative methods with applications to alternating methods, Numer. Math. 76 (1997) 309–321.
[10] D. Bertaccini, G.H. Golub, S. Serra-Capizzano, Spectral analysis of a preconditioned iterative method for the convection-diffusion equation, SIAM J. Matrix Anal. Appl. 29 (2007) 260–278.
[11] D.C. Brown, Alternating-direction iterative schemes for mixed finite element methods for second order elliptic problems, PhD thesis, Department of Mathematics, University of Chicago, Chicago, IL, 1982.
[12] T.A. Davis, Direct Methods for Sparse Linear Systems, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2006.
[13] J. Douglas Jr., R. Durán, P. Pietra, Alternating-direction iteration for mixed finite element methods, in: Computing Methods in Applied Sciences and Engineering VII, Versailles, France, 1985, Springer-Verlag, New York, 1986, pp. 181–196.
[14] H.C. Elman, Preconditioning for the steady-state Navier–Stokes equations with low viscosity, SIAM J. Sci. Comput. 20 (1999) 1299–1316.
[15] H.C. Elman, Preconditioners for saddle point problems arising in computational fluid dynamics, Appl. Numer. Math. 43 (2002) 75–89.
[16] H.C. Elman, A. Ramage, D.J. Silvester, IFISS: A Matlab toolbox for modelling incompressible flow, ACM Trans. Math. Software 33 (2007), Article 14.
[17] H.C. Elman, D.J. Silvester, A.J. Wathen, Performance and analysis of saddle point preconditioners for the discrete steady-state Navier–Stokes equations, Numer. Math. 90 (2002) 665–688.
[18] H. Elman, D. Silvester, A. Wathen, Finite Elements and Fast Iterative Solvers with Applications in Incompressible Fluid Dynamics, Numer. Math. Sci. Comput., Oxford University Press, Oxford, UK, 2005.
[19] R. Glowinski, Numerical Methods for Fluids (Part 3). Finite Element Methods for Incompressible Viscous Flow, in: P.G. Ciarlet, J.L. Lions (Eds.), Handb. Numer. Anal., vol. 9, Elsevier/North-Holland, Amsterdam, 2003.
[20] M.A. Heroux, R.A. Bartlett, V.E. Howle, R.J. Hoekstra, J.J. Hu, T.G. Kolda, R.B. Lehoucq, K.R. Long, R.P. Pawlowski, E.T. Phipps, A.G. Salinger, H.K. Thornquist, R.S. Tuminaro, J.S. Willenbring, A. Williams, K.S. Stanley, An overview of the Trilinos project, ACM Trans. Math. Software 31 (2005) 397–423.
[21] R.B. Kellogg, Another alternating-direction-implicit method, J. Soc. Ind. Appl. Math. 11 (1963) 976–979.
[22] M.A. Olshanskii, Y.V. Vassilevski, Pressure Schur complement preconditioners for the discrete Oseen problem, SIAM J. Sci. Comput. 29 (2007) 2686–2704.
[23] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, SIAM J. Sci. Comput. 14 (1993) 461–469.
[24] Y. Saad, M.H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM J. Sci. Stat. Comput. 7 (1986) 856–869.
[25] V. Simoncini, M. Benzi, Spectral properties of the Hermitian and skew-Hermitian splitting preconditioner for saddle point problems, SIAM J. Matrix Anal. Appl. 26 (2004) 377–389.
[26] V. Simoncini, D.B. Szyld, Recent computational developments in Krylov subspace methods for linear systems, Numer. Linear Algebra Appl. 14 (2007) 1–59.
[27] R.S. Varga, Matrix Iterative Analysis, Prentice-Hall, Englewood Cliffs, NJ, 1962.
[28] R. Wienands, W. Joppich, Practical Fourier Analysis for Multigrid Methods, Chapman & Hall, New York, 2005.
[29] D.M. Young, Iterative Solution of Large Linear Systems, Academic Press, New York, NY, 1971.