

Analysis of preconditioned iterative solvers for incompressible flow problems

S. A. Melchior^{1,*}, V. Legat¹, P. Van Dooren¹ and A. J. Wathen²

¹*CESAME, Batiment Euler, Avenue Georges Lemaître, 4, Louvain-la-Neuve B-1348, Belgium*

²*Computing Laboratory, Wolfson Building, Parks Road, Oxford OX1 3QD, U.K.*

SUMMARY

Solving efficiently the incompressible Navier–Stokes equations is a major challenge, especially in the three-dimensional case. The approach investigated by Elman *et al.* (*Finite Elements and Fast Iterative Solvers*, Oxford University Press: Oxford, 2005) consists in applying a preconditioned GMRES method to the linearized problem at each iteration of a nonlinear scheme. The preconditioner is built as an approximation of an ideal block-preconditioner that guarantees convergence in 2 or 3 iterations. In this paper, we investigate the numerical behavior for the three-dimensional lid-driven cavity problem with wedge elements; the ultimate motivation of this analysis is indeed the development of a preconditioned Krylov solver for stratified oceanic flows which can be efficiently tackled using such meshes. Numerical results for steady-state solutions of both the Stokes and the Navier–Stokes problems are presented. Theoretical bounds on the spectrum and the rate of convergence appear to be in agreement with the numerical experiments. Sensitivity analysis on different aspects of the structure of the preconditioner and the block decomposition strategies are also discussed. Copyright © 2011 John Wiley & Sons, Ltd.

Received 19 December 2008; Revised 28 September 2010; Accepted 14 November 2010

KEY WORDS: preconditioned iterative solvers; Navier–Stokes equations; spectral bounds; block preconditioner; pressure convection–diffusion preconditioner; algebraic multigrid

1. INTRODUCTION

In this paper, we analyze the performance and the theoretical behavior of a preconditioning methodology used with Krylov subspace iteration to obtain a numerical solution of the steady-state incompressible Navier–Stokes boundary value problem

$$\begin{aligned} \rho(\underline{u} \cdot \nabla)\underline{u} &= -\nabla p + \mu \nabla^2 \underline{u} + \underline{f} \quad \text{in } \Omega, \\ \nabla \cdot \underline{u} &= 0 \quad \text{in } \Omega, \\ \underline{u} &= \underline{w} \quad \text{in } \partial\Omega_D, \\ -p\underline{n} + \mu \frac{\partial \underline{u}}{\partial n} &= \underline{0} \quad \text{in } \partial\Omega_N, \end{aligned} \tag{1}$$

where the velocity \underline{u} and the pressure p are the unknown fields. The viscosity, the mass density of the fluid and the external forces are denoted by μ , ρ and f , respectively. The computational domain Ω is a subset of \mathbb{R}^3 bounded by $\partial\Omega = \partial\Omega_D \cup \partial\Omega_N$ with outward-pointing normal \underline{n} .

*Correspondence to: S. A. Melchior, CESAME, Batiment Euler, Avenue Georges Lemaître, 4, Louvain-la-Neuve B-1348, Belgium.

†E-mail: samuel.melchior@uclouvain.be

A large number of papers are devoted to iterative solutions of the Navier–Stokes equations. The incompressibility constraint implies that the linearized matrix of the discrete problem has the structure of a *saddle-point* problem. For this class of equations, special iterative solvers must be designed because of their indefiniteness and poor spectral properties [1–3]. An extensive overview of these methods is given by Benzi *et al.* [4], Elman *et al.* [5] and Turek [6]. Some recent works on numerical solution of saddle-point problems include [7–10].

We investigate the approach of Elman *et al.* [5] in order to solve the Stokes equations, which model creeping flows, and the Navier–Stokes equations. This analysis is developed in the framework of the development of new unstructured grid oceanic models [11–14]. A specific feature appears, as the ocean flows are strongly stratified. In other words, the horizontal scale is far bigger than the vertical one and the discretization of in-plane components of the velocities are distinguished from the out-of-plane velocity. In order to better address this vertical stratification of the flows, the three-dimensional meshes used for oceanic models are obtained by extruding triangular grids along the vertical direction. Unstructured meshes composed of several layers of wedges are attractive in view of their geometrical flexibility to represent coastlines. As the problem in oceanography are typically large scale, we investigate the use of the preconditioned Krylov solvers on such meshes.

We decide to restrict ourselves to a simple three-dimensional lid-driven cavity problem [15, 16], i.e. $\underline{f}=0$ and homogeneous Dirichlet boundary conditions everywhere except on the upper boundary where a constant velocity aligned with an horizontal axis is imposed. The trajectories of particles in this flow are drawn in Figure 1. Such a problem exhibits many relevant effects and it can be considered as a first real benchmark testcase for iterative preconditioned solvers.

One motivation for this analysis is to address time-varying Navier–Stokes equations with an implicit scheme. In such an approach, the discrete linear systems to be solved contains a combination of a mass matrix divided by the time step with a discrete convection–diffusion operator. The inclusion of this mass matrix term generates extra features that must be taken into account by the preconditioner [17–20]. Nevertheless, the time step must be limited for accuracy reason and it renders the contribution of the transport operator relatively small with respect to the steady-state problem.

The meshes used for analyzing the convergence are given in Figure 2. We consider a MINRES solver [21] with a block diagonal preconditioner. The velocity is preconditioned by applying one V-cycle of *algebraic multigrid* AMG. Such a method was introduced in the 1980s by Brandt [22] and Ruge and Stüben [23]. The main idea is to reduce the low-frequency components of the error on coarser operators while the high-frequency components are reduced by a smoother on the fine




Figure 1. The trajectories of particles of the three-dimensional lid-driven cavity problem.

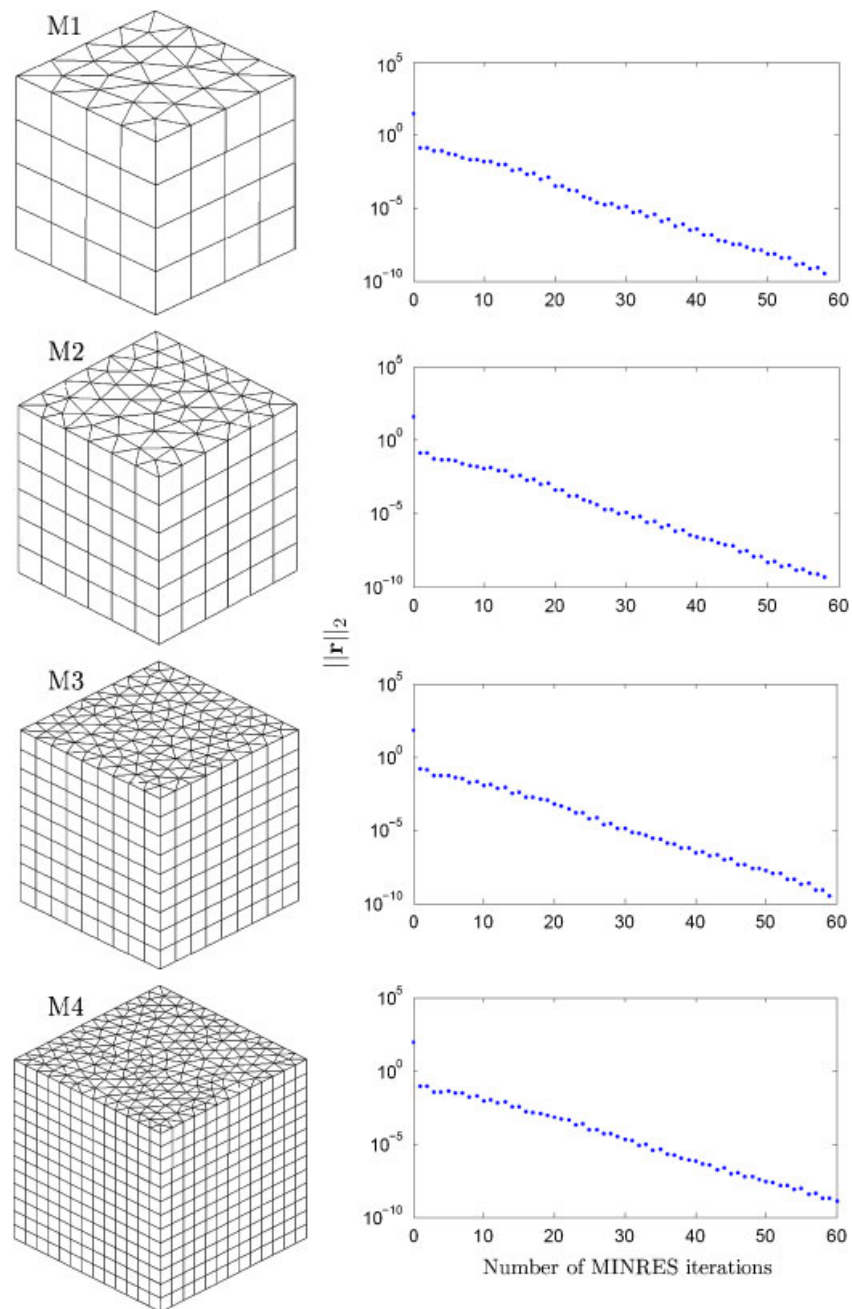


Figure 2. Convergence of the MINRES solver preconditioned with one AMG V-cycle and the mass matrix, for the Stokes equations for meshes M1, M2, M3 and M4, respectively. The rate of convergence is 0.73 on each mesh.

grid. Instead of using several meshes, the AMG variant automatically builds a series of coarser operators by only using the matrix given by the discretization on one fine grid [24]. An iteration where the error is smoothed successively on each grid from the finest to the coarsest and then from the coarsest to the finest is called a V-cycle.

As far as the pressure is concerned, its preconditioner is obtained from the mass matrix Q . For cubic meshes composed of horizontal layers of wedges extruded along the vertical direction, the l_2 norm of the residual converges linearly at the rate of approximately 0.73, independently of the mesh refinement. The norm of the residual is decreased to 10^{-10} in approximately 60 iterations of

Table I. Characteristics of four finite element meshes and CPU time for solving the Stokes equations with both direct and iterative methods on each mesh.

Mesh	# Wedges	# Unknowns	γ^2	Direct solver (s)	Iterative solver (s)
M1	192	3216	0.126286	1	4
M2	588	9053	0.124565	18	10
M3	1998	28717	0.123363	210	42
M4	4602	64497	0.121416	898	114

this preconditioned MINRES solver. The characteristics of the meshes and the typical CPU times required by both the direct sparse and the iterative solvers used are given in Table I. Moreover, numerical estimations based on the relation (4) of the LBB constant γ is also provided [25–27].

This paper is organized as follows. The Stokes equations are first investigated as tight theoretical bounds can be derived in that case. Then, several choices in the extension of this iterative solver to the Navier–Stokes equations are discussed from numerical experiments.

2. PRECONDITIONED MINRES SOLVER FOR THE STOKES EQUATIONS

The finite element discretization of the Stokes equations leads to a linear system

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

where A is the stiffness matrix, i.e. a discrete vectorial Laplacian operator; B and B^T are the discrete negative divergence and gradient operators, respectively. This mixed formulation is the first-order optimality condition for finding the saddle-point to the Lagrangian of the Stokes problem [4]. Hence, \mathcal{A} is indefinite even though the stiffness matrix A is symmetric and positive definite. To solve such a system, an iterative MINRES solver is in general considered because the coefficient matrix \mathcal{A} is very large, ill-conditioned and indefinite. Moreover, this system must be properly preconditioned in order to avoid stagnation in the convergence of the norm of the residual.

An ideal preconditioner would be

$$\mathcal{M}_* = \begin{bmatrix} A & 0 \\ 0 & BA^{-1}B^T \end{bmatrix}. \quad (2)$$

As shown by Murphy *et al.* [28], the minimal polynomial of the preconditioned matrix $\mathcal{T}_* = \mathcal{M}_*^{-1}\mathcal{A}$ has degree 3; this is exactly the number of iterations required for the MINRES scheme to compute the solution with this preconditioner. However, it is not useful in practice. In addition to the inversion of Laplacian operator A , the major drawback of such a preconditioner \mathcal{M}_* is that solving the negative Schur complement $BA^{-1}B^T$ is required. Nevertheless, if M_A and $M_{BA^{-1}B^T}$ are matrix representations of linear approximations of A and $BA^{-1}B^T$, respectively, then this ideal preconditioner should for practical computations be substituted by

$$\mathcal{M} = \begin{bmatrix} M_A & 0 \\ 0 & M_{BA^{-1}B^T} \end{bmatrix}.$$

In order to quantify the quality of both those approximations, we use the concept of *spectral equivalence*, introduced by Axelsson and Barker [29], for symmetric and positive matrices. The matrices A and M_A are said to be spectrally equivalent if there exists some constants $\delta, \Delta > 0$ independent of the mesh size h , such that $\forall \mathbf{v} \neq \mathbf{0}$ the following condition holds:

$$\delta \leq \frac{\langle A\mathbf{v}, \mathbf{v} \rangle}{\langle M_A\mathbf{v}, \mathbf{v} \rangle} \leq \Delta. \quad (3)$$

This relation implies that the real interval $[\delta, \Delta]$ contains all eigenvalues of the preconditioned matrix $M_A^{-\frac{1}{2}} A M_A^{-\frac{1}{2}}$. In the special case where M_A is a preconditioner that yields convergence of the Richardson iteration at a rate independent of h , such inequalities are satisfied (see [5], p. 294). In this case, $\delta = 1 - \rho$ and $\Delta = 1 + \rho$, where ρ is the rate of convergence of this simple iterative scheme. This is the case, for instance, for one multigrid V-cycle.

The Schur complement is approximated by the pressure mass matrix Q , as shown for example in [5] $\forall \mathbf{q} \notin \ker(B^T)$

$$\gamma^2 \leq \frac{\langle B A^{-1} B^T \mathbf{q}, \mathbf{q} \rangle}{\langle Q \mathbf{q}, \mathbf{q} \rangle} \leq \Gamma^2, \quad (4)$$

where $\gamma > 0$ is the LBB constant. The estimated values of γ^2 given in Table I for several meshes are computed as the smallest eigenvalue of $Q^{-1} S$ as indicated by the eigenvalue interpretation of this relation. On the other hand, the ratio between the inner products with the Schur complement and the mass matrix is bounded above by Γ which is also independent of h . For enclosed flow, this boundedness constant satisfies $\Gamma = 1$; otherwise, $\Gamma \leq \sqrt{d}$, where d is the dimension of the space. Note that the kernel of B^T corresponds to the hydrostatic pressure, i.e. $\mathbf{q} = q_0 \mathbf{1}$.

2.1. Defining a tight inclusion set that contains the preconditioned spectrum

Some theoretical bounds on the spectrum of the preconditioned matrix can be derived. Each eigenvalue $\lambda \neq 0$ of $\mathcal{T} = \mathcal{M}^{-1} \mathcal{A}$ lies in the inclusion set

$$\mathcal{S} = [-a, -b] \cup [c, d],$$

where a, b, c, d are the positive real numbers. As this matrix is indefinite, this set must include both positive and negative values. The eigenvalues of $\mathcal{M}^{-1} \mathcal{A}$ satisfy the generalized eigenvalue problem:

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \lambda \begin{bmatrix} M_A & 0 \\ 0 & Q \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix}. \quad (5)$$

Following the approach developed in [5], we first decouple this system by substituting one equation into the other

$$B(A - \lambda M_A)^{-1} B^T \mathbf{p} = -\lambda Q \mathbf{p}, \quad (6)$$

$$(A - \lambda M_A) \mathbf{u} = -B^T (\lambda Q)^{-1} B \mathbf{u} \quad (7)$$

for λ such that the inverse matrix $(A - \lambda M_A)^{-1}$ exists. The spectral equivalence is then used to obtain spectral bounds.

Let us first consider the case $\lambda < 0$. In this case, the matrix $A - \lambda M_A$ is certainly invertible and is spectrally equivalent to A

$$\left(1 - \frac{\lambda}{\Delta}\right) \langle A \mathbf{v}, \mathbf{v} \rangle \leq \langle (A - \lambda M_A) \mathbf{v}, \mathbf{v} \rangle \leq \left(1 - \frac{\lambda}{\delta}\right) \langle A \mathbf{v}, \mathbf{v} \rangle \quad (8)$$

as a result of the inequalities (3). The spectral equivalence between their inverse directly follows from the fact that $\langle M_1 \mathbf{v}, \mathbf{v} \rangle \leq \langle M_2 \mathbf{v}, \mathbf{v} \rangle$ is strictly equivalent to $\langle M_2^{-1} \mathbf{v}, \mathbf{v} \rangle \leq \langle M_1^{-1} \mathbf{v}, \mathbf{v} \rangle$, for symmetric and positive definite matrices. Finally, substituting \mathbf{v} by $B^T \mathbf{p}$, the relations (8) read:

$$\left(1 - \frac{\lambda}{\delta}\right)^{-1} \langle A^{-1} B^T \mathbf{p}, B^T \mathbf{p} \rangle \leq \langle (A - \lambda M_A)^{-1} B^T \mathbf{p}, B^T \mathbf{p} \rangle \leq \left(1 - \frac{\lambda}{\Delta}\right)^{-1} \langle A^{-1} B^T \mathbf{p}, B^T \mathbf{p} \rangle.$$

The latter gives after using the relations (7) and (4):

$$\left(1 - \frac{\lambda}{\delta}\right)^{-1} \gamma^2 \langle Q \mathbf{p}, \mathbf{p} \rangle \leq -\lambda \langle Q \mathbf{p}, \mathbf{p} \rangle \leq \left(1 - \frac{\lambda}{\Delta}\right)^{-1} \Gamma^2 \langle Q \mathbf{p}, \mathbf{p} \rangle.$$

The second inequality implies that $\lambda^2 - \lambda\Delta - \Delta\Gamma^2 \leq 0$. The negative eigenvalues satisfying the latter are bounded below as

$$\lambda \geq \frac{\Delta - \sqrt{\Delta^2 + 4\Delta\Gamma^2}}{2} = -a \geq -\Gamma^2, \quad (9)$$

where the lower bound on $-a$ is obtained after the addition of $4\Gamma^4$ inside the square root. Similarly, the first inequality implies that $\lambda^2 - \lambda\delta - \delta\gamma^2 \geq 0$. The negative solutions of this inequation are given by:

$$\lambda \leq \frac{\delta - \sqrt{\delta^2 + 4\delta\gamma^2}}{2} = -b < 0. \quad (10)$$

In the case $\lambda > 0$, the bounds are obtained similarly; it follows from Equation (6) and the relations (3) that

$$\left(1 - \frac{\lambda}{\delta}\right) \langle A\mathbf{u}, \mathbf{u} \rangle \leq \frac{-1}{\lambda} \langle B^T Q^{-1} B\mathbf{u}, \mathbf{u} \rangle \leq \left(1 - \frac{\lambda}{\Delta}\right) \langle A\mathbf{u}, \mathbf{u} \rangle.$$

From the first inequality, it can be derived that

$$\lambda \geq \delta = c > 0 \quad (11)$$

as Q and A are positive definite. From the second inequality, we take advantage of $\langle B^T Q^{-1} B\mathbf{u}, \mathbf{u} \rangle \leq \Gamma^2 \langle A\mathbf{u}, \mathbf{u} \rangle$, which corresponds to the upper bound in (4), in order to obtain $\lambda^2 - \lambda\Delta - \Delta\Gamma^2 \leq 0$. The latter's positive solutions satisfy:

$$\lambda \leq \frac{\Delta + \sqrt{\Delta^2 + 4\Delta\Gamma^2}}{2} = d \leq \Delta + \Gamma^2. \quad (12)$$

As a summary of the bounds (9), (10), (11) and (12), the non-zero eigenvalues of Equation (5) belong to the inclusion set \mathcal{S}

$$\lambda \in \mathcal{S} = \left[\frac{\Delta - \sqrt{\Delta^2 + 4\Delta\Gamma^2}}{2}, \frac{\delta - \sqrt{\delta^2 + 4\delta\gamma^2}}{2} \right] \cup \left[\delta, \frac{\Delta + \sqrt{\Delta^2 + 4\Delta\Gamma^2}}{2} \right]. \quad (13)$$

These bounds are tight. For problems with the Dirichlet boundary conditions, the constant Γ is equal to 1 while it is bounded by the square root of the dimension of the computational domain for Neumann boundary conditions [5]. If we restrict ourselves to the Dirichlet boundary condition, the set (13) for the ideal preconditioner \mathcal{M}_* ($\delta = \Delta = \gamma = \Gamma = 1$) reduces to $\{(1 - \sqrt{5})/2\} \cup [1, (1 + \sqrt{5})/2]$. The limits of this set are actually the three eigenvalues of \mathcal{T}_* that implies the three-iteration convergence. Hence, it appears that the approximation \mathcal{M} diffuses the discrete spectrum onto two intervals on each side of the origin.

It is interesting to compare the theoretical bounds (13) with the spectrum obtained on several meshes. In Figure 3, the spectrum obtained in the case $M_A = A$ (i.e. only $\gamma \neq 1$ as it can be shown that $\Gamma = 1$ for full Dirichlet problem) can be visualized. As it is required to compute each eigenvalue of the preconditioned matrix, it is not practical to perform computation with highly refined meshes. Therefore, we only consider a series of coarse regular cubic meshes with 1, 2 and 4 layers of wedges; the most refined mesh exhibits only 64 degrees of freedom. In order to visualize the density of the spectrum, a Gaussian with a small variance is centered on each eigenvalue. Then, this curve is obtained by taking the logarithm of the sum of all these Gaussians. Each curve is symmetric with respect to 0.5, except for the peaks around 0 and 1. Such a result is proved in [5].

In Figure 4, the distributions of the eigenvalues for the mesh M1 and for an equivalent regular mesh are compared. In the case of the regular mesh, there are two distinct eigenvalues that are twice as close to 1 as the closest one in the irregular mesh case. This can be viewed as a direct consequence of a smaller LBB constant. As we will observe, it implies that the irregular mesh yields

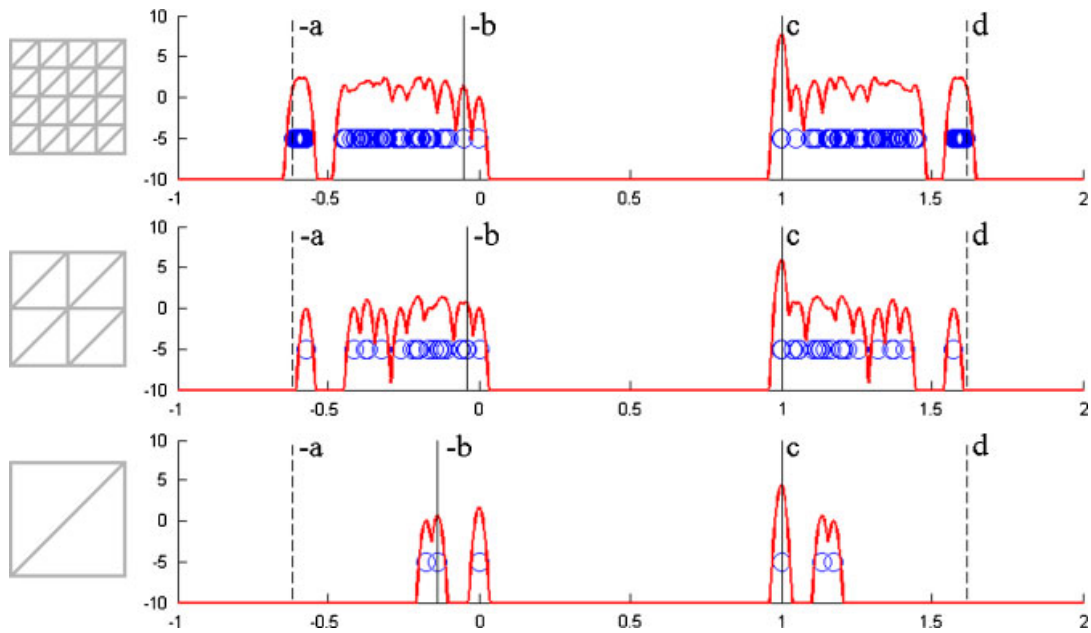


Figure 3. Spectrum of the preconditioned matrix $\mathcal{T} = \mathcal{M}^{-1}\mathcal{A}$ with $M_A = A$ and $M_{BA^{-1}B^T} = Q$ for a series of coarse regular cubic meshes with 1, 2 and 4 layers of wedges. The theoretical bounds are indicated with vertical lines.




Figure 4. Close-up view of the positive eigenvalues of the preconditioned matrix \mathcal{T} for the regular mesh with 4 layers of wedges and for the irregular mesh M1. Both meshes have approximately the same number of degrees of freedom.

better convergence, recalling that the spectrum is symmetric with respect to 0.5. Nevertheless, both curves are quite comparable, except the size of the gap around 1.5. In both cases, the upper theoretical bound is strongly tight.

Such an approach can be considered as an improved presentation of the eigenvalue bounds presented in Elman *et al.* [5]. Such bounds were established by Rusten and Winther [30] and independently for the Stokes problem by Wathen and Silvester [31, 32].

2.2. From the inclusion set to a theoretical bound on the rate of convergence

The preconditioned MINRES scheme is a member of the family of Krylov subspace methods, such as the *conjugate gradient method* (CG) [33] and the *generalized minimal residual method*

(GMRES) [34]. The common idea behind these methods is to minimize the preconditioned residual of the linear system in the Krylov subspace $\mathcal{K}_k(\mathcal{T}, \mathbf{r}_{(0)})$ with respect to a norm. Due to the structure of this Krylov subspace, the following equality holds:

$$\|\hat{\mathbf{r}}_{(k)}\| = \min_{p_k \in \Pi_k, p_k(0)=1} \|p_k(\mathcal{T})\hat{\mathbf{r}}_{(0)}\|, \tag{14}$$

where Π_k is the set of real polynomials of degree k , and $\hat{\mathbf{r}}_{(k)} = \mathcal{M}^{-1}\mathbf{r}_{(k)}$ is the preconditioned residual. Let us assume that \mathcal{T} is diagonalizable ($\mathcal{T} = \mathcal{V}\Lambda\mathcal{V}^{-1}$, where Λ is the diagonal matrix of the eigenvalues, and \mathcal{V} is the matrix of the eigenvectors). It follows that

$$\begin{aligned} \|p_k(\mathcal{T})\mathbf{r}_{(0)}\| &= \|\mathcal{V} p_k(\Lambda)\mathcal{V}^{-1}\mathbf{r}_{(0)}\| \\ &\leq \|\mathcal{V}\| \|p_k(\Lambda)\| \|\mathcal{V}^{-1}\| \|\mathbf{r}_{(0)}\| \\ &= \kappa(\mathcal{V}) \max_j |p_k(\lambda_j)| \|\mathbf{r}_{(0)}\|, \end{aligned}$$

where $\kappa(\mathcal{V}) \triangleq \|\mathcal{V}\| \|\mathcal{V}^{-1}\|$ is the condition number of \mathcal{V} . In the cases of CG and MINRES, \mathcal{T} is symmetric and \mathcal{V} is thus an orthogonal matrix with $\kappa(\mathcal{V}) = 1$.

The norm of the preconditioned residual at each iteration is then bounded by the solution of an optimization problem

$$\frac{\|\hat{\mathbf{r}}_{(k)}\|}{\|\hat{\mathbf{r}}_{(0)}\|} \leq \kappa(\mathcal{V}) \min_{p_k \in \Pi_k, p_k(0)=1} \max_j |p_k(\lambda_j)|. \tag{15}$$

In general, this problem is slightly relaxed by considering it over the inclusion region \mathcal{S} . For the Stokes equations, we can then easily estimate the rate of convergence. If both intervals have the same length, the optimal polynomial at the $2k$ th MINRES iteration is $\tau_k(s(\lambda))/\tau_k(s(0))$, where $\tau_k(s)$ is the k th Chebyshev polynomial and $s(\lambda)$ is the quadratic polynomial which maps $[-a, -b]$ as well as $[c, d]$ on $[-1, 1]$. Finally, the residual at iteration $2k$ satisfies:

$$\|\hat{\mathbf{r}}_{(2k)}\| \leq 2 \left(\frac{\sqrt{ad} - \sqrt{bc}}{\sqrt{ad} + \sqrt{bc}} \right)^k \|\hat{\mathbf{r}}_{(0)}\|. \tag{16}$$

In Figure 5, the rate of convergence of the preconditioned residual is compared with the theoretical bounds given by (16). We successively consider the cases where M_A^{-1} is A^{-1} and one algebraic multigrid (AMG) V-cycle [24], respectively. In the latter case, the values of δ and Δ are deduced from the observed rate of convergence $\rho_{AMG} = 0.185$ of the multigrid V-cycles. The only difference between the case $M_A = A$ and the ideal three-iterations preconditioner is that the mass matrix approximates the Schur complement. Nevertheless, the convergence behavior is really different: the number of iterations increases from 3 to more than 50.

Let us observe the staircase effect in the convergence behavior: the norm of the residual does not decrease at the even iterations. This can be deduced from the symmetry of the spectrum. Recalling the way the optimal polynomial is built from a Chebyshev polynomial, it is clear that the optimal polynomial of degree $2k + 1$ is very similar to the optimal polynomial of degree $2k$.

When M_A^{-1} is one AMG V-cycle, the convergence is slightly deteriorated. Nevertheless, the computational time decreases from 1025 (s) to 105 (s): using one algebraic multigrid iteration is thus an efficient strategy. The staircase effect is weaker because the spectrum is no longer exactly symmetric in this case, as illustrated by the computed limits of the inclusion set for the mesh M4:

$$\mathcal{S} = [-0.6469, -0.1073] \cup [0.8150, 1.8319].$$

To apply the theoretical result (16) to this case, we have to increase the first interval to $[-1.1242, -0.1073]$ in order to have both intervals of the same length. For both considered preconditioners, it appears that the theoretical bounds on the rate of convergence are fairly loose. This is partly due to the relaxation that gives the same length to both intervals, especially in the case with the AMG V-cycle.

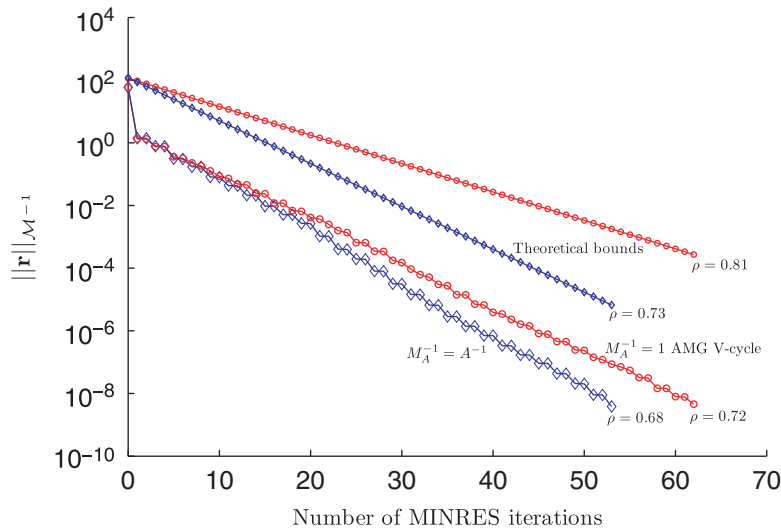


Figure 5. Convergence of the preconditioned MINRES solver and the corresponding theoretical bounds for two different M_A , for the Stokes lid-driven cavity with the mesh M4. The preconditioned residual norm $\|r\|_{M^{-1}}$ is shown. The Schur complement is approximated by the mass matrix.




Figure 6. Structure of the nested iterations in the solver of the Navier–Stokes equations.

3. PRECONDITIONED GMRES SOLVER FOR THE NAVIER–STOKES EQUATIONS

The Picard linearization of the discrete Navier–Stokes equations yield

$$\mathcal{A} \mathbf{x} = \begin{bmatrix} F & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix}, \tag{17}$$

where F is the linear combination $\mu A + \rho N$ of the Laplacian operator and the linearized convection term. As the coefficient matrix of this system is not symmetric, the MINRES method is replaced by the GMRES method. As illustrated in Figure 6, the computation of the solution of the Navier–Stokes equations with a preconditioned iterative solver requires a hierarchical interconnection of methods. At least four different levels can be distinguished. The outer loop consists of nonlinear iterations. Each of the latter requires to solve a linear system. Therefore, the inner loop is the preconditioned Krylov subspace method that computes its solution. Applying the inverse of the preconditioner requires the application of three operations. They can be either simple multiplication, one V-cycle of a multigrid scheme or the composition of three operators.

3.1. Block structure of the preconditioner

For the Stokes problem and the MINRES method, a symmetric and positive definite block diagonal preconditioner is used. Now, it is no longer required for the preconditioner to be symmetric and positive definite with the GMRES method applied to the Navier–Stokes equations. Note that the only extra cost between block triangular and block diagonal preconditioners is a simple sparse matrix multiplication.

On the one hand, an ideal block upper triangular preconditioner would be

$$\mathcal{M}_U = \begin{bmatrix} F & B^T \\ 0 & -S \end{bmatrix}, \quad (18)$$

where $S = BF^{-1}B^T$ is the negative Schur complement of the linearized matrix. It would achieve excellent convergence as was observed for the ideal Stokes preconditioner (2). The minimum polynomial of the preconditioned matrix has degree 2; this is exactly the number of iterations required for the GMRES scheme to compute the solution with this preconditioner. Indeed, the matrix

$$\mathcal{T}_U = \mathcal{M}_U^{-1}\mathcal{A} = \begin{bmatrix} I + F^{-1}B^TS^{-1}B & F^{-1}B^T \\ -S^{-1}B & 0 \end{bmatrix}$$

satisfies $(\mathcal{T}_U - \mathcal{I})^2 = 0$. The low-rank matrix $\mathcal{T}_U - \mathcal{I}$ can be written as $[F^{-1}B^T - I]^T [S^{-1}B \quad I]$ and it follows from the definition of S that

$$\begin{bmatrix} F^{-1}B^T \\ -I \end{bmatrix} \underbrace{[S^{-1}B \quad I] \begin{bmatrix} F^{-1}B^T \\ -I \end{bmatrix}}_{S^{-1}BF^{-1}B^T - I} [S^{-1}B \quad I] = 0.$$

On the other hand, an ideal block lower triangular preconditioner would be

$$\mathcal{M}_L = \begin{bmatrix} F & 0 \\ B & -S \end{bmatrix}. \quad (19)$$

It would also achieve excellent convergence as the matrix

$$\mathcal{T}_L = \mathcal{M}_L^{-1}\mathcal{A} = \begin{bmatrix} I & F^{-1}B^T \\ 0 & I \end{bmatrix}$$

satisfies $(\mathcal{T}_L - \mathcal{I})^2 = 0$. In Figure 7, almost the same convergence behavior is observed for both preconditioners. Intuitively, this might be understood as the same operations are applied but in a different order. Moreover, $\mathcal{M}_L^T = \mathcal{M}_U$ in the Stokes case.

The only eigenvalue of \mathcal{T}_L and \mathcal{T}_U is $\lambda = 1$, but its geometric multiplicity is smaller than its algebraic multiplicity. If the sign in front of the matrix S is changed, the spectrum now consists in a set of eigenvalues -1 and 1 whose geometric and algebraic multiplicity are equal. For any selection of sign, convergence would be achieved in two iterations for both choice of sign. However, it is no longer the case when approximations are used. The inclusion regions of the spectrum will evolve from a point to a non-degenerated area in the complex plane. If the minus sign is chosen, as in Equations (18) and (19), there will be only one region in the right half-plane, instead of one in each half-plane with the opposite sign. The bound (15), with the maximization relaxed over the inclusion region, indicates that the former choice might yield faster convergence.

Moreover, the preconditioned matrix has been written so far with the preconditioner applied on the left side $\mathcal{T} = \mathcal{M}^{-1}\mathcal{A}$. Instead of solving the equation $\mathcal{M}^{-1}\mathcal{A}\mathbf{x} = \mathcal{M}^{-1}\mathbf{f}$, right preconditioning consists in finding $\mathbf{y} = \mathcal{M}\mathbf{x}$ such that $\mathcal{A}\mathcal{M}^{-1}\mathbf{y} = \mathbf{f}$. One advantage of this alternative is that the preconditioned norm is not different from the norm of the true residual $\mathbf{f} - \mathcal{A}\mathbf{x}$. Therefore, the effect

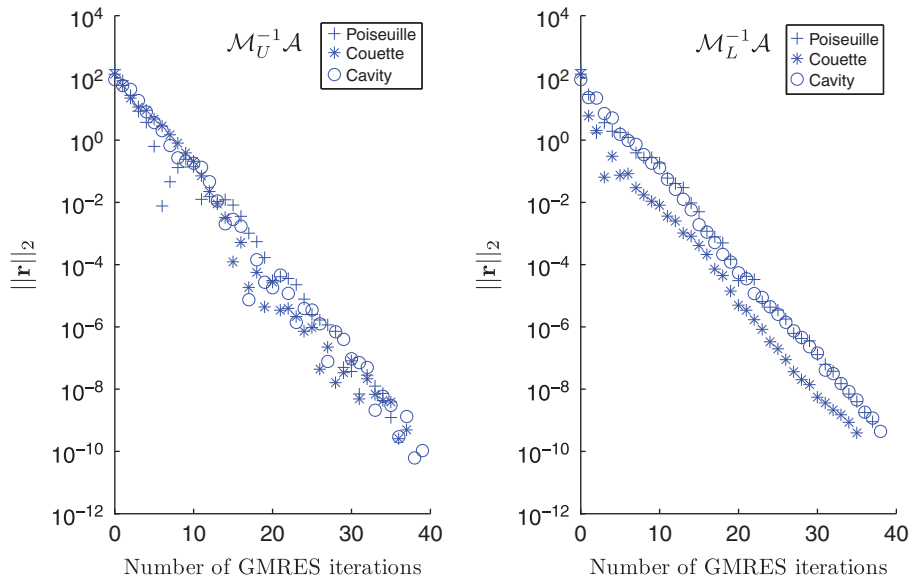


Figure 7. Convergence of the GMRES solver, with upper and lower triangular preconditioners, for several Stokes problems on the mesh M4. F and S are approximated in the preconditioner by 1 AMG V-cycle and the mass matrix, respectively. The true residual norm $\|\mathbf{r}\|_2$ is shown.

of distorted norm which could give a wrong impression of convergence does not exist in this case. In the previous strategy for the Stokes problem, the MINRES method actually solves the equation $\mathcal{H}^{-1}\mathcal{A}\mathcal{H}^{-T}(\mathcal{H}^T\mathbf{x}) = \mathcal{H}^{-1}\mathbf{f}$, where \mathcal{H} is obtained from a theoretical decomposition for positive definite matrix $\mathcal{M} = \mathcal{H}\mathcal{H}^T$, like, for instance, the Cholesky factorization. Such a factorization is not needed in practice. Hence, there is neither discussion about the side of preconditioning, nor about the sign of the Schur complement as this block must be positive definite.

As far as the spectrum is concerned, there is no difference between both possibilities of preconditioning, $\mathcal{M}^{-1}\mathcal{A}$ and $\mathcal{A}\mathcal{M}^{-1}$. Indeed, the eigenvalues are invariant under the similarity transformation, i.e the multiplication on the left by a matrix, and on the right by its inverse. However, the eigenvectors, which are the columns of the matrix \mathcal{V} in the diagonalization of the preconditioned matrix $\mathcal{T} = \mathcal{V}\Lambda\mathcal{V}^{-1}$, are different. As the condition number of \mathcal{V} appears in the bound (15) for the convergence of general Krylov methods, one of the sides might be more efficient.

One way to compare both possibilities would be to evaluate the non-normality of the preconditioned matrix, as normal matrices satisfy $\kappa(\mathcal{V}) = 1$. The Frobenius norm of the difference $\mathcal{N}(\mathcal{T}) = \mathcal{T}\mathcal{T}^T - \mathcal{T}^T\mathcal{T}$ is a possible measure of non-normality; its value is 0 for normal matrix according to their definition. This measure can be improved by normalizing it, as $\|\mathcal{N}(\mathcal{T})\|_F / \|\mathcal{T}\|_F^2$, so it is scale-invariant. This value is twice as small in the case of left preconditioning, which indicates that this choice would be more interesting. Nevertheless, this difference is not sufficiently large to draw such a conclusion. A good discussion about the choice of a measure of non-normality can be found in Trefethen and Embree [35]. Another possible measure is $\|\mathcal{N}(\mathcal{T})\|_F / \|\mathcal{T}^2\|_F$. This quantity is also scale-invariant and goes to infinity if $\mathcal{N}(\mathcal{T}) \rightarrow \infty$, unlike $\|\mathcal{N}(\mathcal{T})\|_F / \|\mathcal{T}\|_F^2$, which is ≤ 2 .

Both the questions of the sign in front of the Schur complement and of the preconditioning side are investigated in Figure 8. The latter shows for the four possibilities the convergence of the preconditioned GMRES solver when the approximation proposed for the Stokes problem is used instead of F and S in this preconditioner. This result tends to confirm that the minus sign in front of the Schur complement is more efficient, whatever the side of preconditioning. Because of the difference between the geometric and algebraic multiplicities, this change in sign does not modify the number of iterations to solve the ideally preconditioned system, even though the number of distinct eigenvalues is reduced from 2 to 1. However, it improves significantly the rate of convergence when approximations are used for the practical preconditioner.

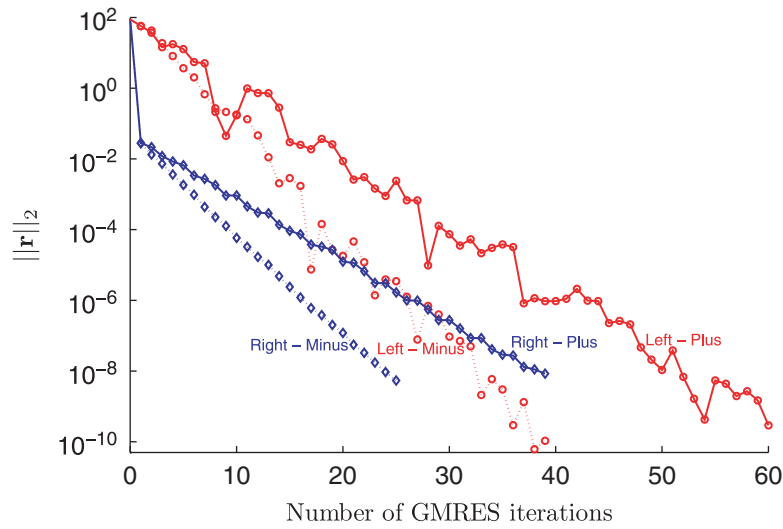


Figure 8. Convergence of the true residual norm $\|\mathbf{r}\|_2$ with the GMRES solver on the Stokes lid-driven cavity flow on the mesh M4. F and S are approximated in the preconditioner by 1 AMG V-cycle and the mass matrix, respectively.

Whatever preconditioning side is chosen, the observed rate of convergence is 0.53 in the minus case and 0.67 otherwise. Indeed, the relation (15) indicates that the eigenvalue spectrum solely determines the evolution of the upper bound on the norm of the residual at each iteration and by extension, an upper bound on the rate of convergence. However, nothing prevents the norm of the residual from being smaller than the upper bound; in other words, the number of iterations does not only depend on the eigenvalue spectrum. It appears that preconditioning on the right-hand side yields faster convergence thanks to the gap in the norm of the residual at the first iteration. This empirical observation does not confirm the trend indicated by the measure of non-normality.

The l_2 norm of the residual can increase through one iteration in the case of left preconditioning. The convergence is not monotonic because the GMRES method actually minimizes the norm of the preconditioned residual $\|\mathbf{r}\|_{\mathcal{M}^{-1}}$ which is thus distorted [36]. For the sake of completeness, the convergence of the same computations but observed in the preconditioned norm of the residual is shown in Figure 9. The termination criterion used is a relative tolerance on the preconditioned norm of 10^{-10} .

One may observe that the rate of convergence to solve the Stokes equation is faster for the GMRES solver with the block triangular preconditioner. This does not mean that the symmetric strategy with MINRES can be forgotten. The theoretical bounds set the latter in a reliable framework. Moreover, the MINRES algorithm is a short-term recurrence that works only with three vectors, whereas the GMRES method needs to keep in memory all vectors computed for one Gram–Schmidt orthogonalization per iteration.

3.2. Approximation to the Schur complement

For the Navier–Stokes equations, preconditioners of the Schur complement can be deduced from the assumption that a discrete version of the commutator between some convection–diffusion and mixed first-order differential operators is small. Choosing the latter as the gradient operator yield the formulation initially proposed by Kay *et al.* [37]

$$M_S = A_p F_p^{-1} Q \tag{20}$$

called *the pressure convection–diffusion preconditioner*. The matrix A_p is a discrete Laplacian operator built in the pressure space. In order to approximately solve a system with this sparse matrix, one AMG V-cycle can be efficient. The nonsymmetric matrix F_p is the analog, in the pressure space, of the operator F from the Picard linearization. Based on intuition, one can think

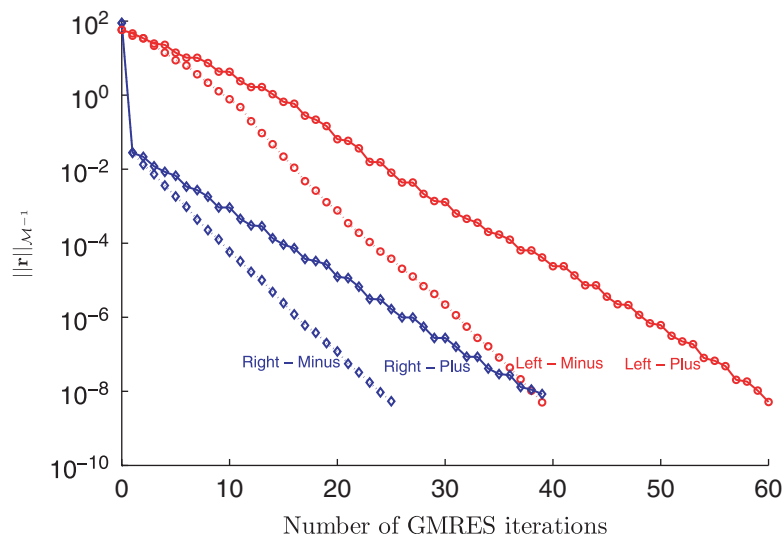


Figure 9. Convergence of the preconditioned residual norm $\|\mathbf{r}\|_{M^{-1}}$ with the GMRES solver on the Stokes lid-driven cavity flow on the mesh M4. F and S are approximated in the preconditioner by 1 AMG V-cycle and the mass matrix, respectively.

that both matrices somehow have the same non-normality; this part of the preconditioner is essential to take the convection term into account. Note that the matrix F_p is involved in one multiplication as only M_S^{-1} has to be applied.

Both these matrices are not involved in the discrete linearized system; one major issue is the boundary conditions used in their definition. As the three-dimensional lid-driven cavity problem is an enclosed flow, A_p and F_p corresponds to discretizations where the Neumann conditions are imposed and are thus singular; this seems reasonable as the mean value of the pressure is free in such a case. For problems with inflow and outflow boundary, it is shown in [38] that considering the commutator with the divergence operator and using, e.g. Robin conditions can help obtaining a better approximation.

Finally, Q is the mass matrix in the pressure space, whose inversion is, in fact, approximated. Only a few iterations of the Chebyshev Semi Iterative method (see Varga [39], pp. 149–156), preconditioned by its diagonal D_Q , are completed (see [40]). This scheme is an efficient Krylov subspace iteration that is constant and linear, which is needed to ensure convergence of GMRES. Otherwise, one should use FGMRES; this might be attractive even though it does not enjoy the same theoretical convergence properties as GMRES. Eigenvalue bounds, required as parameters of the Chebyshev Semi Iterative method, are given by the spectral bounds of the local mass matrix (see Wathen [41]). The computation of these bounds for the linear wedge elements shows that:

$$\frac{1}{4} \leq \frac{\langle Q\mathbf{p}, \mathbf{p} \rangle}{\langle D_Q\mathbf{p}, \mathbf{p} \rangle} \leq 3.$$

As this element is the Cartesian product of a triangle by an interval, the values of these bounds can be deduced from the bounds for these lower dimensional elements. In fact, the local mass matrix for the wedge element is equal to the Kronecker product of the equivalent matrices for the one-dimensional and the two-dimensional elements. Their spectrum and their eigenvectors are also linked by Kronecker products; each eigenvalue of the local mass matrix for the wedge element is the product of an eigenvalue for the triangle (either 0.5 or 2) with an eigenvalue for the interval (either 0.5 or 1.5).

The inverse of the mass matrix has a two-fold goal: scaling and unsmoothing. Indeed, the integration on each element bring an extra scaling proportional to its mesh size in the block of the discretized system. Moreover, applying the latter has the effect of averaging the values of the continuous differential operator between the nodes and their close neighbors. The mass matrix is

Table II. Comparison between the convergence rate using the scaled mass matrix and the convergence rate using the six possible orderings in the definition of the pressure convection–diffusion preconditioner, for solving the Stokes equations and for the inversion of its Schur complement.

Ordering	Convergence rate
$Q^{-1} F_p \hat{A}_p^{-1}$	Equivalent to scaled mass matrix
$\hat{A}_p^{-1} F_p Q^{-1}$	Nearly equivalent to scaled mass matrix
$Q^{-1} \hat{A}_p^{-1} F_p$	Slower than scaled mass matrix
$F_p \hat{A}_p^{-1} Q^{-1}$	Stagnation of the true residual norm
$F_p Q^{-1} \hat{A}_p^{-1}$	Stagnation of the true residual norm
$\hat{A}_p^{-1} Q^{-1} F_p$	No iteration

Table III. Comparison between the norm of the preconditioned residuals and the true residuals at each iteration using the six possible orderings in the definition of the pressure convection–diffusion preconditioner, for solving the Stokes equations, the Oseen problem and the inversion of their Schur complement.

Ordering	Preconditioned residual norm
$Q^{-1} F_p \hat{A}_p^{-1}$	Equivalent to true residual norm
$\hat{A}_p^{-1} F_p Q^{-1}$	Nearly equivalent to true residual norm
$Q^{-1} \hat{A}_p^{-1} F_p$	A bit distorted
$F_p \hat{A}_p^{-1} Q^{-1}$	Distorted norm
$F_p Q^{-1} \hat{A}_p^{-1}$	Distorted norm
$\hat{A}_p^{-1} Q^{-1} F_p$	Highly distorted norm

somehow a generalization of the identity matrix that includes this scaling and smoothing property. Hence, applying its inverse somehow improves the approximations used in the discrete version of the commutator. The way it is applied to derive M_S in [37, 38] can be interpreted in a least-square sense. Nevertheless, the other choices of ordering obtained by permuting the position of the mass matrix are also valid.

We studied the six different possibilities for solving the Stokes equations, the Oseen problem and for the inversion of their Schur complement. For the Stokes problem, the preconditioner M_S would be a scaled pressure mass matrix if the matrix A_p could be inverted exactly. However, the latter is singular and thus its inverse \hat{A}_p^{-1} is approximated by 1 AMG V-cycle. It may seem irrelevant to apply pressure convection–diffusion preconditioning on the Stokes problem, but it is nevertheless possible and the fact that the matrices do not commute necessarily makes the different choices distinct. A summary of the observation sorted from the best to the poorest is given in Tables II and III.

There are no clear explanations of these different results; nevertheless, intuition can be used to find some interpretations. The last three orderings are clearly inefficient. Even if the GMRES method manages to reduce the preconditioned norm of the residual, the norm of the true residual of the system remains approximately constant. One could have foreseen that the matrices F_p and A_p cannot be separated as, together, they form the nonsymmetric correction to the Schur complement approximation for the Stokes equations.

The first three orderings are all acceptable. The best one is the ordering proposed initially which fits the scaled mass matrix perfectly. The second one is also efficient; the common feature of these two is that the matrix F_p is in the middle. It seems reasonable that both are somehow equivalent as, in the Stokes case, one is the transpose of the other.

3.3. Numerical results

In Figure 10, the convergence of the GMRES solver to compute the second Picard iterate is shown for the three-dimensional lid-driven cavity problem at different $Re \leq 500$. The preconditioner uses one AMG V-cycle and the pressure convection–diffusion preconditioner $A_p F_p^{-1} Q$ as approximation of the convection–diffusion operator and of the negative Schur complement, respectively. The evolution with the Reynolds numbers of the number of iterations to reach convergence can be visualized in Figure 11. If the matrix A_p could be inverted exactly, the product $F_p^{-1} A_p$ would yield a scaled identity matrix for the Stokes problem. Therefore, it can be seen as a nonsymmetric correction to this approximation of the Schur complement for the Navier–Stokes equations. It appears that the number of iterations increases rather linearly with the mass matrix while a curve

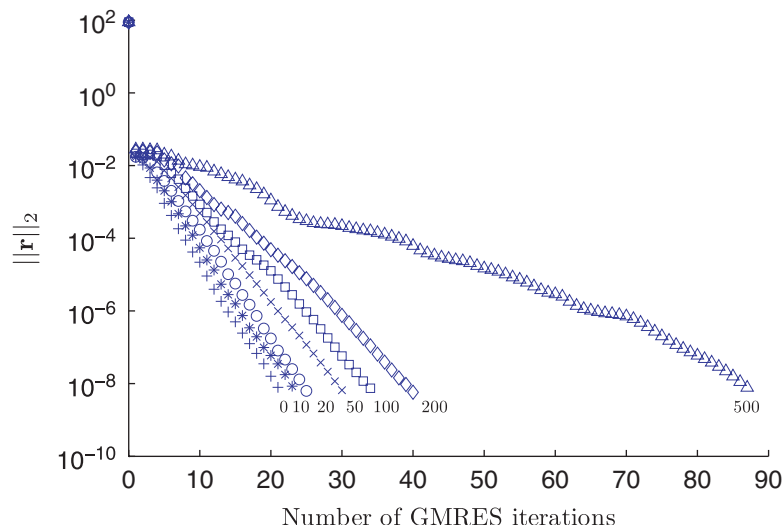


Figure 10. Convergence of the GMRES solver to compute the second Picard iterate with the preconditioner using one AMG V-cycle for M_F and $M_S = A_p F_p^{-1} Q$, for several Reynolds numbers on the mesh M4.




Figure 11. Number of GMRES iterations in order to reach a relative tolerance of 10^{-10} , in function of the Reynolds number for both approximations of the Schur complement on the mesh M4. A curve fitting of Δn , the increase in the number of iterations compared with the case $Re=0$, is given for the pressure convection–diffusion preconditioner.

fitting with a power law gives an experimental exponent of approximately $\frac{3}{5}$ for the pressure convection–diffusion preconditioner. This shows that this modification of M_S significantly improves the convergence of the GMRES solver.

The value for $Re = 500$ is ignored in the curve fitting because the behavior of the iterative solver is very different. Around $Re = 300$, new physics begins to emerge in the form of detachment from the upstream side-wall of the separation line between the primary core and the secondary eddies. It is interesting to observe that this critical value for the Reynolds number has a similar impact on the iterative scheme and on the physical observations from the transient flows obtained numerically [15]. The flow appears to remain laminar in the range $1 < Re < 1200$ with vortices forming around $Re = 1300$.

4. CONCLUSIONS

A Krylov solver for the discrete Stokes equations is presented. The MINRES method preconditioned using the mass matrix and the algebraic multigrid yields convergence independent of the mesh size. Hence, the complexity for solving this linear system is linear with respect to the number of unknowns. Tight theoretical bounds on the spectrum of the preconditioned matrix are derived and provide a better understanding of the behavior of the technique. A new symmetrical and simpler proof of such bounds is given. Convergence analysis exhibits some specific features: in particular, better convergence is observed for irregular meshes that seems to better transfer local information into the iterative scheme.

The generalization of this iterative method to the linearized Navier–Stokes equations is discussed on the basis of numerical experiments. First, the MINRES method must be substituted by another Krylov solver that is suitable for any matrices as the coefficient matrix is nonsymmetric. We choose the GMRES scheme as it computes the optimal solution in each successive Krylov subspaces. Second, a symmetric positive definite preconditioned is not required by GMRES. Hence, the block diagonal preconditioner can be replaced by a block triangular one that improves significantly the rate of convergence with a simple sparse matrix multiplication as extra cost. As this preconditioned is less constrained, new degrees of freedom appear in its definition. In this paper, the following issues have been analyzed for solving the linearized Navier–Stokes equations with a preconditioned GMRES scheme:

- First, the issue of selecting an upper or a lower triangular preconditioner is investigated. Surprisingly, both structures exhibit almost the same behavior.
- Right preconditioning seems to be more efficient; moreover, it avoids the ambiguity on the norm to measure the convergence.
- A minus sign in front of the negative Schur complement tends to yield a better rate of convergence.
- A sensitivity analysis is performed on the ordering for the pressure convection–diffusion preconditioner. All approaches are not equivalent and only the $Q^{-1}F_p\hat{A}_p^{-1}$ order is able to yield the same convergence as the scaled mass matrix in the Stokes case.

Several numerical experiments support those observations, based on an original implementation of the solver. However, no rigorous theoretical explanations are provided in this paper.

ACKNOWLEDGEMENTS

Samuel Melchior is a Research fellow with the Belgian National Fund for Scientific Research (FNRS). This study was carried out within the scope of the project ‘A second-generation model of the ocean system’, which is funded by the Communauté Française de Belgique, as Actions de Recherche Concertées, under Contract ARC 04/09-316. This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office. The scientific responsibility rests with its author(s).

REFERENCES

1. Lipnikov K, Vassilevski Yu. Parallel adaptive solution of the Stokes and Oseen problems on unstructured 3D meshes. In *Parallel Computational Fluid Dynamics 2003*, Ecer A, Satofuka N, Periaux J, Fox P (eds). Elsevier: Amsterdam, 2004; 153–161. ISBN: 978-0-44-451612-1, DOI: 10.1016/B978-044451612-1/50020-2.
2. Olshanskii MA. A low order Galerkin finite element method for the Navier–Stokes equations of steady incompressible flow: a stabilization issue and iterative methods. *Computer Methods in Applied Mechanics and Engineering* 2002; **191**:5515–5536.
3. Quarteroni A, Saleri F, Veneziani A. Factorization methods for the numerical approximation of Navier–Stokes equations. *Computer Methods in Applied Mechanics and Engineering* 2000; **188**:505–526.
4. Benzi M, Golub GH, Liesen J. Numerical solution of saddle point problems. *Acta Numerica* 2005; **14**:1–137.
5. Elman HC, Silvester DJ, Wathen AJ. *Finite Elements and Fast Iterative Solvers*. Oxford University Press: Oxford, U.K., 2005.
6. Turek S. *Efficient Solvers for Incompressible Flow Problems*. Springer: Berlin, 1999.
7. Benzi M, Olshanskii MA. An augmented Lagrangian-based approach to the Oseen problem. *SIAM Journal on Scientific Computing* 2006; **28**(6):2095–2113.
8. de Niet AC, Wubs FW. Two preconditioners for saddle point problems in fluid flows. *International Journal for Numerical Methods in Fluids* 2007; **54**:355–377.
9. Peters J, Reichelt V, Reusken A. Fast iterative solvers for discrete Stokes equations. *SIAM Journal on Scientific Computing* 2005; **27**(2):646–666.
10. Rehman M, Vuik C, Segal G. A comparison of preconditioners for incompressible Navier–Stokes solvers. *International Journal for Numerical Methods in Fluids* 2008; **57**(12):1731–1751.
11. Balay S, Gropp WD, McInnes LC, Smith BF. Efficient management of parallelism in object oriented numerical software libraries. In *Modern Software Tools in Scientific Computing*, Arge E, Bruaset AM, Langtangen HP (eds). Birkhauser: Basel, 1997; 163–202.
12. Comblen R, Legrand S, Deleersnijder E, Legat V. A finite element method for solving the shallow water equations on the sphere. *Ocean Modelling* 2009; **28**:12–23.
13. Hanert E, Le Roux DY, Legat V, Deleersnijder E. An efficient Eulerian finite element for the shallow water equations. *Ocean Modelling* 2005; **10**:115–136.
14. White L, Legat V, Deleersnijder E. Tracer conservation for three-dimensional, finite-element, free-surface, ocean modeling on moving prismatic meshes. *Monthly Weather Review* 2008; **136**:420–442.
15. Chiang TP, Sheu WH, Hwang RR. Effect of Reynolds number on the eddy structure in a lid-driven cavity. *International Journal for Numerical Methods in Fluids* 1998; **26**:557–579.
16. Povitsky A. Three-dimensional flow in cavity at yaw. *Nonlinear Analysis* 2005; **63**:e1573–e1584.
17. Cahouet J, Chabard JP. Some fast 3D finite element solvers for generalized Stokes problem. *International Journal for Numerical Methods in Fluids* 1988; **8**:869–895.
18. Gresho PM, Sani RL. *Incompressible Flow and the Finite Element Method, Isothermal Laminar Flow*, vol. 2. Wiley: Chichester, 1998.
19. Kay DA, Gresho PM, Griffiths DF, Silvester DJ. Adaptive time-stepping for incompressible flow. Part II: Navier–Stokes equations. *SIAM Journal on Scientific Computing* 2010; **32**(1):111–128.
20. Simo JC, Armero F. Unconditional stability and long-term behaviour of transient algorithms for the incompressible Navier–Stokes and Euler equations. *Computer Methods in Applied Mechanics and Engineering* 1994; **111**: 111–154.
21. Paige C, Saunders M. Solution of sparse indefinite systems of linear equations. *SIAM Journal on Numerical Analysis* 1975; **13**:617–629.
22. Brandt A. Algebraic multigrid theory: the symmetric case. *Applied Mathematics and Computation* 1986; **19**:23–56.
23. Ruge JW, Stüben K. Algebraic Multigrid (AMG). In *Multigrid Methods, Frontiers in Applied Mathematics*, McCormick SF (ed.). SIAM: Philadelphia, 1987; 73–130.
24. Henson VE, Yang UM. BoomerAMG: a parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics* 2002; **41**:155–177.
25. Babuška I. The finite element method with Lagrangian multipliers. *Numerische Mathematik* 1973; **20**:179–192.
26. Brezzi F. On the existence, uniqueness and approximation of saddle point problems arising from Lagrangian multipliers. *RAIRO Analyse Numerique* 1974; **8**:129–154.
27. Ladyzhenskaya O. *The Mathematical Theory of Viscous Incompressible Flow*. Gordon and Breach: New York, 1969.
28. Murphy MF, Golub GH, Wathen AJ. A note on preconditioning for indefinite linear systems. *SIAM Journal on Scientific Computing* 2000; **21**:1969–1972.
29. Axelsson O, Barker VA. *Finite Element Solution of Boundary Value Problems*. Academic Press: Orlando, 1984.
30. Rusten T, Winther R. A preconditioned iterative method for saddle-point problems. *SIAM Journal on Matrix Analysis and Applications* 1992; **13**:887–904.
31. Silvester D, Wathen A. Fast iterative solutions of stabilised Stokes systems. Part II: using general block preconditioners. *SIAM Journal on Numerical Analysis* 1994; **31**:1352–1367.
32. Wathen A, Silvester D. Fast iterative solution of stabilised Stokes systems. Part I: using simple diagonal preconditioners. *SIAM Journal on Numerical Analysis* 1993; **30**:630–649.

33. Hestenes MR, Stiefel E. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards* 1952; **49**:409–435.
34. Saad Y, Schultz MH. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing* 1986; **7**:856–869.
35. Trefethen LN, Embree M. *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton University Press: NJ, 2005.
36. Wathen A. Preconditioning and convergence in the right norm. *International Journal of Computer Mathematics* 2007; **84**(8):1199–1209.
37. Kay D, Loghin D, Wathen A. A preconditioner for the steady-state Navier–Stokes equations. *SIAM Journal on Scientific Computing* 2002; **24**(1):237–256.
38. Elman HC, Tuminaro RS. Boundary conditions in approximate commutator preconditioners for the Navier–Stokes equations. *Electronic Transactions on Numerical Analysis* 2009; **35**:257–280.
39. Varga RS. *Matrix Iterative Analysis*. Prentice-Hall: Englewood Cliffs, NJ, 1962. Second edition, Springer: New York, 2000.
40. Wathen A, Rees T. Chebyshev semi-iteration in preconditioning for problems including the mass matrix. *Electronic Transactions on Numerical Analysis* 2009; **34**:125–135.
41. Wathen A. Realistic eigenvalue bounds for the Galerkin mass matrix. *IMA Journal of Numerical Analysis* 1987; **7**:449–457.