INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS
Int. J. Numer. Meth. Fluids 2008; **00**:1–20 *Prepared using* fldauth.cls *[Version: 2002/09/18 v1.01]*

Analysis of preconditioned iterative solvers for incompressible flow problems

S. A. Melchior[1,*,†], V. Legat[1,‡], P. Van Dooren [1,§] and A. J. Wathen[2,¶]

[1] *CESAME, Batiment Euler, avenue Georges Lemaître, 4, Louvain-la-Neuve, B-1348, Belgium*
[2] *Computing Laboratory, Wolfson Building, Parks Road, Oxford OX1 3QD*

SUMMARY

Solving efficiently the incompressible Navier–Stokes equations is a major challenge, especially in the three-dimensional case. The approach investigated by Elman, Silvester and Wathen [*Finite Elements and Fast Iterative Solvers*. Oxford University Press, 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 wedge elements dedicated to stratified oceanic flows. 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.

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 incompressible Navier–Stokes equations

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

*Correspondence to: CESAME, Batiment Euler, avenue Georges Lemaître, 4, Louvain-la-Neuve, B-1348, Belgium
†E-mail: samuel.melchior@uclouvain.be
‡E-mail: vincent.legat@uclouvain.be
§E-mail: paul.vandooren@uclouvain.be
¶E-mail: wathen@maths.ox.ac.uk

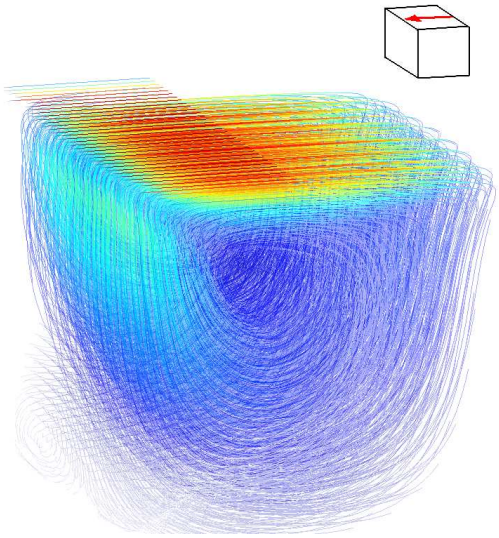

Figure 1. The trajectories of particles of the cubic driven cavity problem. The norm of the velocity ranging from 0 to 1 is given by the color.

where the computational domain Ω is a subset of $\mathbb{R}^3$. The viscosity, the mass density of the fluid and the external forces are denoted by μ, ρ and f, respectively. The velocity field $\underline{u}$ satisfies suitable boundary conditions on $\partial\Omega$, and the pressure field is given by p.

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 [20, 22, 26]. A extensive overview of these methods is given by Benzi et al. [4], Elman et al. [12] and Turek [34]. Some recent works on numerical solution of saddle-point problems include [5, 11, 24, 27].

We investigate the approach of Elman, Silvester and Wathen [12] in order to solve the Stokes and Navier–Stokes equations. This analysis is developed in the framework of the development of new unstructured grid oceanic models [3, 10, 14, 40] . 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 (u and v) are distinguished from the out-of-plane velocity, w. Unstructured grids for oceanic flows are attractive in view of the geometrical flexibility to represent coastlines. However, the stratified vertical property of the flows is better addressed by a regular structured grid in the vertical direction. As the problem in oceanography are typically large scale, we investigate the use of preconditioned Krylov solvers for three-dimensional unstructured grids, with wedge elements. In fact, the meshes used in the oceanic model is obtained from a triangular grid.

We decide to restrict ourselves to a simple three-dimensional square cavity driven problem [9, 25]. 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

Table I. Characteristics of 4 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**	**Iterative solver**
M1	192	3216	0.126286	1 $[s]$	4 $[s]$
M2	588	9053	0.124565	18 $[s]$	10 $[s]$
M3	1998	28717	0.123363	210 $[s]$	42 $[s]$
M4	4602	64497	0.121416	898 $[s]$	114 $[s]$

preconditioned solvers. The motivation of 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 consist of the 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 [8, 13, 17, 30]. 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 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 1980's by Brandt [6] and Ruge & Stüben [28]. 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 grid. Instead of using coarse grids, the AMG variant automatically builds a series of coarser operator by only using the matrix given by the discretization on one fine grid [15]. 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 wedges extruded in 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 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, the numerical estimation of the LBB constant γ is also provided [2, 7, 19].

This paper is organized as follows. The Stokes equations, which model creeping flows, are first investigated as tight theoretical bounds can be derived in that case. Then, several choices in the extension of this preconditioned 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

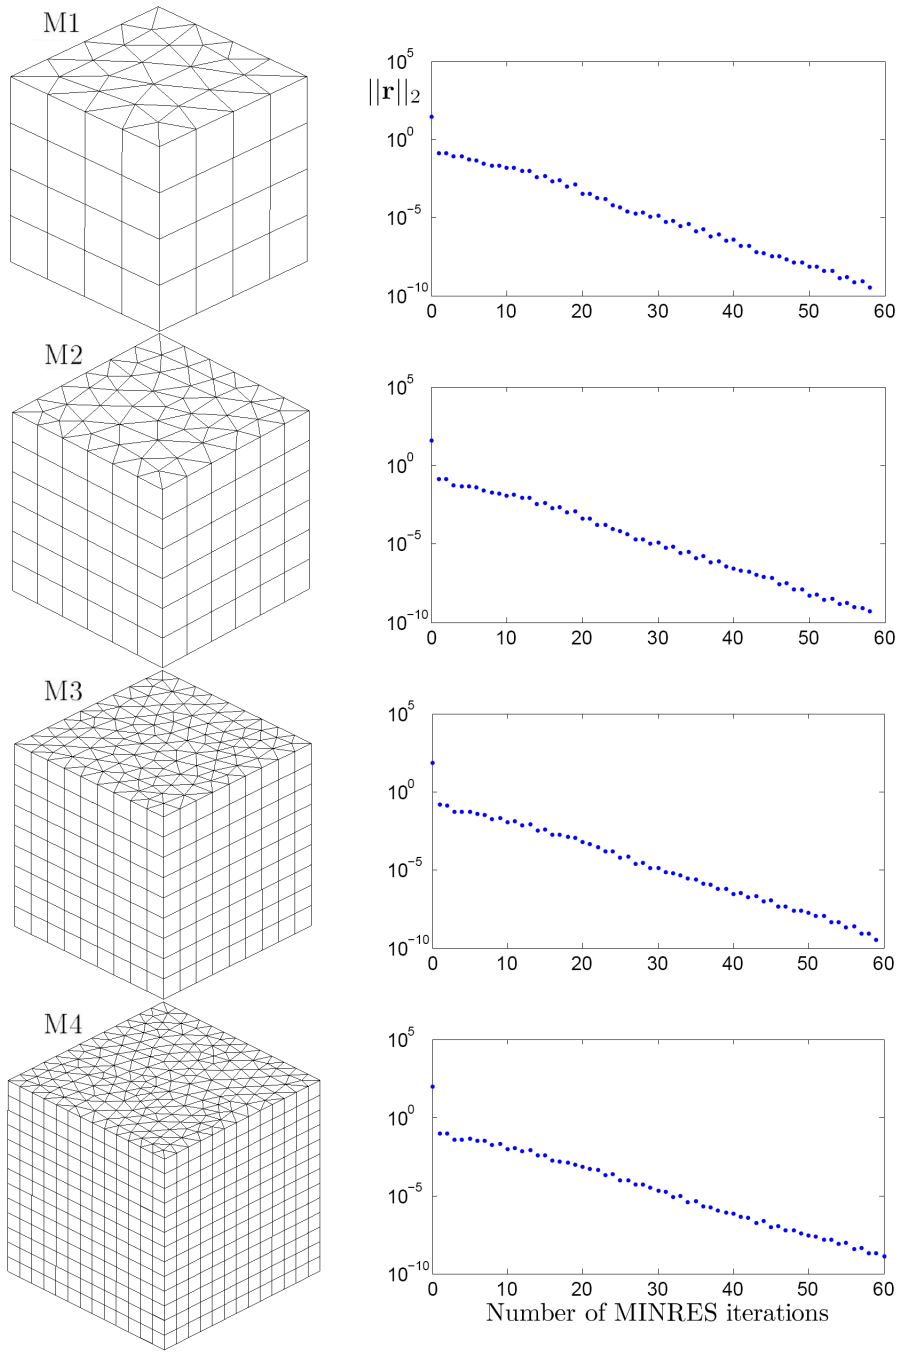

Figure 2. Convergence of the MINRES solver preconditioned with one AMG V-cycle and the mass matrix, for the Stokes equations for meshes M1, M2, M3 and M4, respectively. The rate of convergence is 0.73 on each mesh.

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

where A is the stiffness matrix, i.e. a discrete vectorial Laplacian operator; B and B^T are discrete negative divergence and gradient operator, 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 [23] 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

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

As shown by Murphy, Golub & Wathen [21], the minimum 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 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 [1], 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:

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

In the special case where M_A is a preconditioner that yields a convergence of the Richardson iteration at a rate independent of h, such inequalities are satisfied. (see [12] 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, since as shown for example in [12] $\forall \mathbf{q} \notin \ker(B^T)$:

$$\gamma^2 \leq \frac{\left\langle BA^{-1}B^T\mathbf{q}, \mathbf{q}\right\rangle}{\langle Q\mathbf{q}, \mathbf{q}\rangle} \leq \Gamma^2, \tag{4}$$

where $\gamma > 0$ is the LBB constant; values of γ^2 on several meshes are given in Table I. On the other side, 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, $\Gamma < \sqrt{d}$ 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 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:

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

Following the approach developed in [12], we firstly decouple this system by substituting one equation into the other:

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

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

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:

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

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:

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

The latter gives after using the relations (6) and (4):

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

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

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

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:

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

In the case $\lambda > 0$, the bounds are obtained similarly; it follows from the equation (7) and the relations (3) that

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

From the first inequality, it can be derived that:

$$\lambda \geq \delta = c > 0, \tag{11}$$

as Q and A and are positive definite. From the second inequality, we take advantage of $\left\langle B^T Q^{-1} B\mathbf{u}, \mathbf{u}\right\rangle \leq \Gamma^2 \langle A\mathbf{u}, \mathbf{u}\rangle$, 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:

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

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

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

Those bounds are tight. For problems with 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 [12]. If we restrict ourselves to 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 $\mathcal{T}_*$ which implies the three-iterations 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) to 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

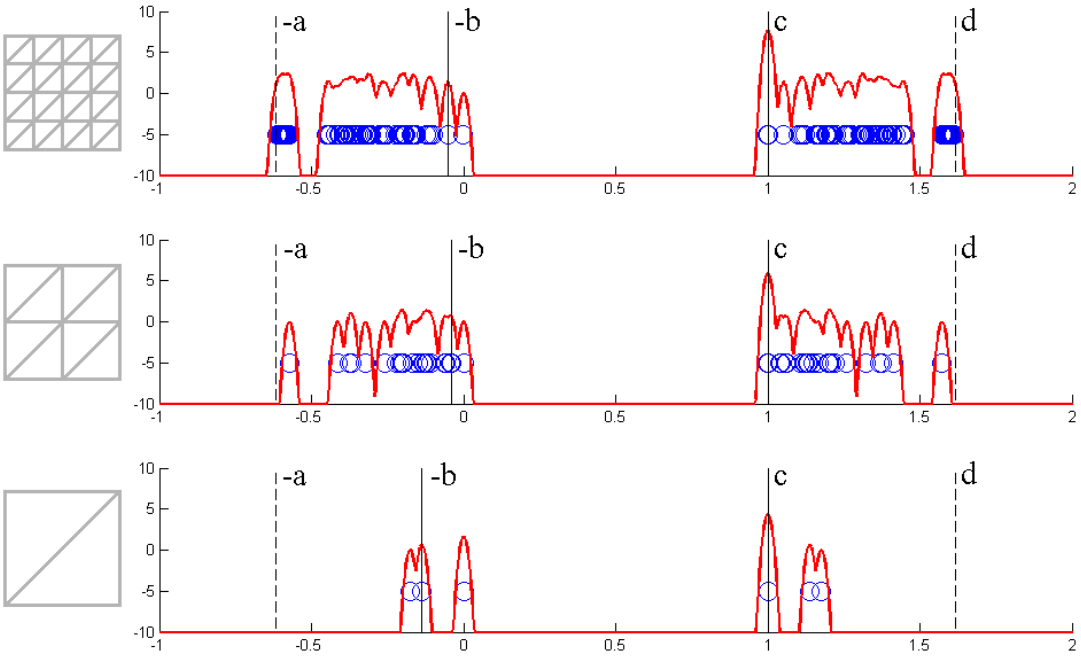

Figure 3. Spectrum of the preconditioned matrix $\mathcal{T} = \mathcal{M}^{-1}\mathcal{A}$ with $M_A = A$ and $M_{BA^{-1}B^T} = Q$ for a series of coarse regular cubic meshes with 1, 2 and 4 layers of wedges. The theoretical bounds are indicated with vertical lines.

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 [12].

In Figure 4, the distributions of the eigenvalues for the mesh M1 and for an equivalent regular mesh are compared. For the regular mesh, we observe the occurrence of two distinct eigenvalues twice as close to 1. This can be viewed as a direct consequence of a smaller LBB constant. As we will observe, it implies that the irregular mesh yields a 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 a better presentation of the eigenvalue bounds presented in Elman, Silvester and Wathen [12]. Such bounds were established by Rusten and Winther [29] and independently for the Stokes problem by Wathen and Silvester [32, 39].

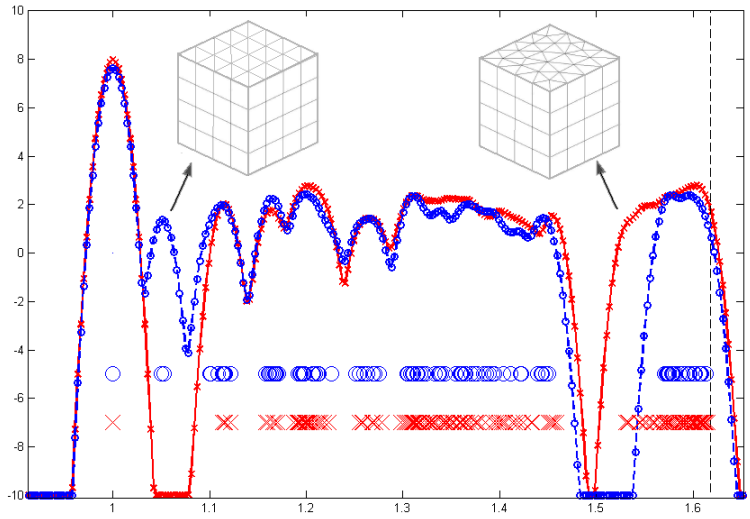

Figure 4. Close-up view of the positive eigenvalues of the preconditioned matrix $\mathcal{T}$ for the regular mesh with 4 layers of wedges and for the irregular mesh M1. Both meshes have approximately the same number of degrees of freedom.

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, like the *conjugate gradient method* (CG) [16] and the *generalized minimal residual method* (GMRES) [31]. 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 a 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 [15], 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 the same as 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 a algebraic multigrid iteration is thus an efficient strategy. The staircase effect is weaker because the spectrum is no longer symmetric in that 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 which gives the same length to both intervals, especially in the case with the AMG V-cycle.

3. PRECONDITIONED GMRES SOLVER FOR THE NAVIER–STOKES EQUATIONS

From the Picard linearization of the discrete Navier–Stokes equations, we obtain:

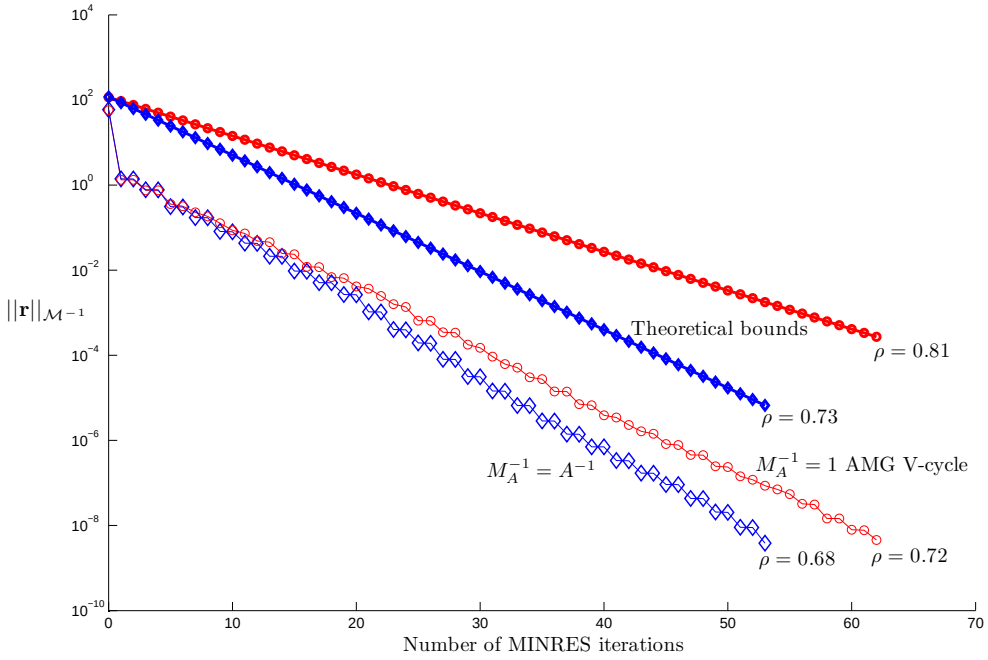

Figure 5. Convergence of the preconditioned MINRES solver and the corresponding theoretical bounds for two different M_A, for the Stokes lid driven cavity with the mesh M4. The preconditioned residual norm $\|\mathbf{r}\|_{\mathcal{M}^{-1}}$ is shown. The Schur complement is approximated by the mass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

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

where F is the linear combination $\mu A + \rho N$ of the Laplacian operator and the linearized convection term. Since 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 in the nonlinear iterations. Each of the latter requires to solve a linear system. Therefore, the inner loop is the preconditioned Krylov subspace method which computes its solution. Applying the inverse of the preconditioner requires to apply 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, we used a symmetric and positive definite block diagonal preconditioner. Now, it is no longer required for the preconditioner to be

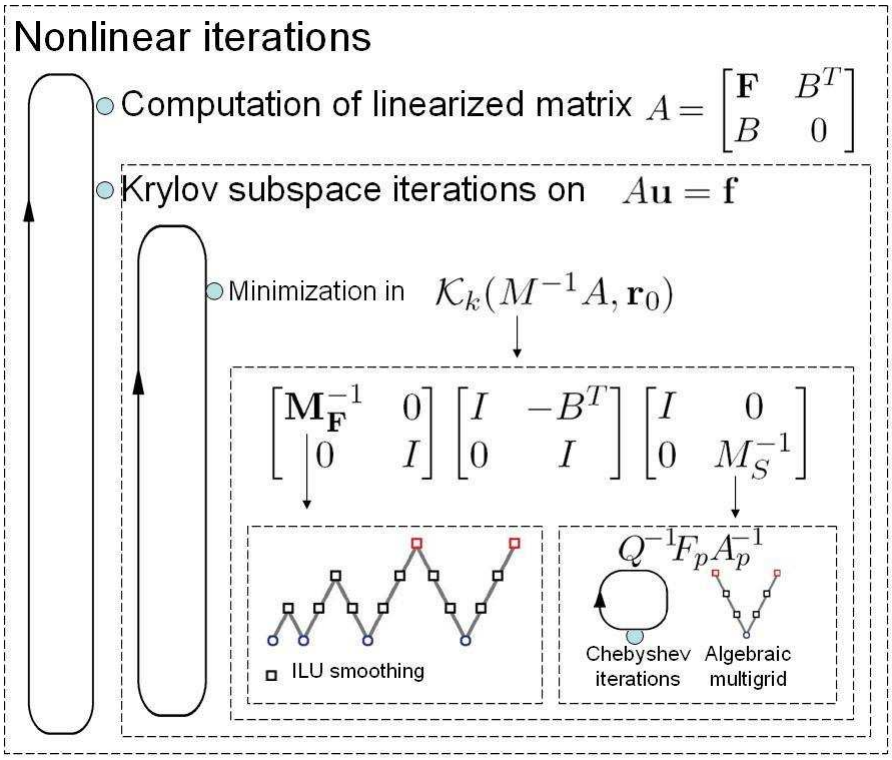

Figure 6. Structure of the nested iterations in the solver of the Navier–Stokes equations.

symmetric and positive definite with the GMRES method applied to the Navier–Stokes equations.

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

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

where $S = BF^{-1}B^T$ is the negative Schur complement of the linearized matrix. It would achieve an 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 $\begin{bmatrix} F^{-1}B^T - I \end{bmatrix}^T \begin{bmatrix} S^{-1}B & I \end{bmatrix}$

and it follows from the definition of S that

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

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

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

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, we observe almost the same convergence behavior for both preconditioners. Intuitively, this might be understood as we apply the same operations but in a different order. Moreover, $\mathcal{M}_L^T = \mathcal{M}_U$ in the Stokes case.

The spectrum of $\mathcal{T}_L$ and $\mathcal{T}_U$ only contain eigenvalues $\lambda_j = 1$ with geometrical multiplicity of one or two. If the sign in front of the matrix S is changed, the spectrum now consists in a set of eigenvalues -1 and 1, with geometrical multiplicity 1. For any selection of sign, convergence would be achieved in two iterations for both upper and lower ideal preconditioner. 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 equation (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 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 $\mathcal{H}^{-1}\mathcal{A}\mathcal{H}^{-T}(\mathcal{H}^T\mathbf{x}) = \mathcal{H}^{-1}\mathbf{b}$, 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 since 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. Since 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.

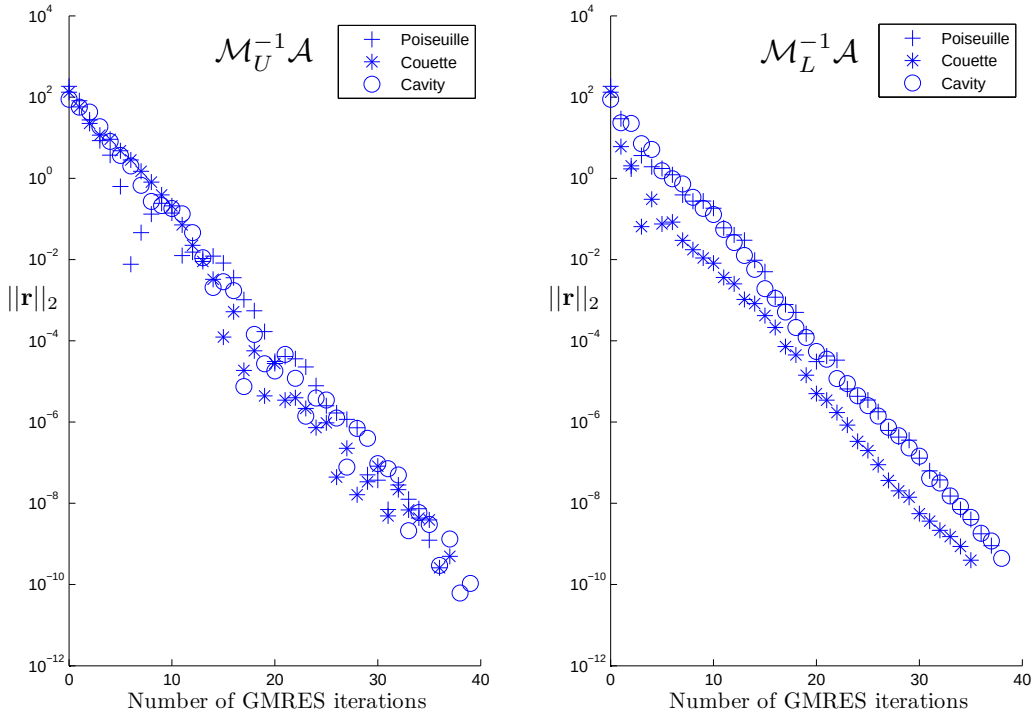

Figure 7. Convergence of the GMRES solver, with upper and lower triangular preconditioners, for several Stokes problems on the mesh M4. F and S are approximated in the preconditioner by 1 AMG V-cycle and the mass matrix, respectively. The true residual norm $\|\mathbf{r}\|_2$ is shown.

One way to compare both possibilities would be to evaluate the non-normality of the preconditioned matrix, since 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 [33]. Another possible measure is $\|\mathcal{N}(\mathcal{T})\|_F/\|\mathcal{T}^2\|_F$. This quantity is also scale-invariant and goes to infinity if $\nu(A) \to \infty$, unlike $\|\mathcal{N}(\mathcal{T})\|_F/\|\mathcal{T}\|_F^2$, which is ≤ 2.

The other way is to compare the convergence when the approximation proposed for the Stokes problem are used instead of F and S in this preconditioner. This experimentation, shown in Figure 8, is not in agreement with the first observation. Preconditioning on the right-hand side seems to give a faster convergence. This result confirms as well that the minus sign in front of the Schur complement is more efficient, whatever the side of preconditioning.

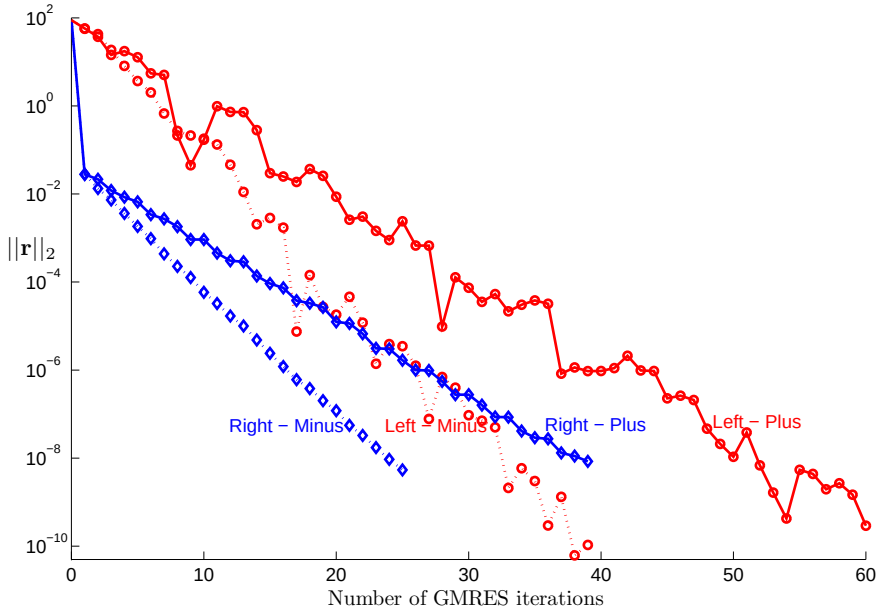

Figure 8. Convergence of the preconditioned residual norm $\|\mathbf{r}\|_{\mathcal{M}^{-1}}$ with the GMRES solver on the Stokes lid driven cavity flow on the mesh M4. F and S are approximated in the preconditioner by 1 AMG V-cycle and the mass matrix, respectively

Whatever side is chosen, the observed rate of convergence is 0.53 in the minus case and 0.67 otherwise. The difference between both preconditioning sides seems to be only due to the gap at the first iteration. The bounds (15) shows that the rate of convergence only depends on the spectrum. We observe that 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 [37].

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 which works only with 3 vectors while 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

One possibility is the idea derived in Kay, Loghin & Wathen [18]:

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

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 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 since only M_S^{-1} has to be applied. On the boundary, no conditions are imposed; this seems reasonable since the mean value of the pressure is free for enclosed flow. Therefore, both A_p and F_p are singular.

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 [35] pp. 149–156), preconditioned by its diagonal D_Q, are completed (see [38]). 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 [36]). The computation of these bounds for the linear wedge elements shows that:

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

Since 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).

As far as the ordering of the matrices in the definition of M_S is concerned, there is no proof that it has to be as in equation (20). We studied the six different possibilities for solving the Stokes equations, the Oseen problem and for the inversion of their Schur complement. In the Stokes case, the preconditioner M_S is not really a scaled mass matrix, because A_p is singular and thus approximated by $\hat{A}_p$ where $\hat{A}_p^{-1}$ is the matrix representation of 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 Table II and Table III.

There are no clear explanations of these different results. Nevertheless, intuition can be

[†]For the matrices $A \in \mathbb{R}^{m\times n}$ and $B \in \mathbb{R}^{p\times q}$, the Kronecker product $A \otimes B \in \mathbb{R}^{mp\times nq}$ is the block matrix $\begin{bmatrix} a_{11}B & \cdots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \cdots & a_{mn}B \end{bmatrix}$

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
$\hat{A}_p^{-1}F_pQ^{-1}$	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
$F_pQ^{-1}\hat{A}_p^{-1}$	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
$\hat{A}_p^{-1}F_pQ^{-1}$	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
$F_pQ^{-1}\hat{A}_p^{-1}$	distorted norm
$\hat{A}_p^{-1}Q^{-1}F_p$	highly distorted norm

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 matrix F_p and A_p cannot be separated since, 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 since, in the Stokes case, one is the transpose of the other.

3.3. Results

Figure 10 shows that the number of iterations increases rather linearly with the mass matrix while a curve fitting with a power law gives an experimental exponent of approximately $\frac{3}{5}$. The value for $Re = 500$ is ignored in the curve fitting because the behavior of the iterative solver

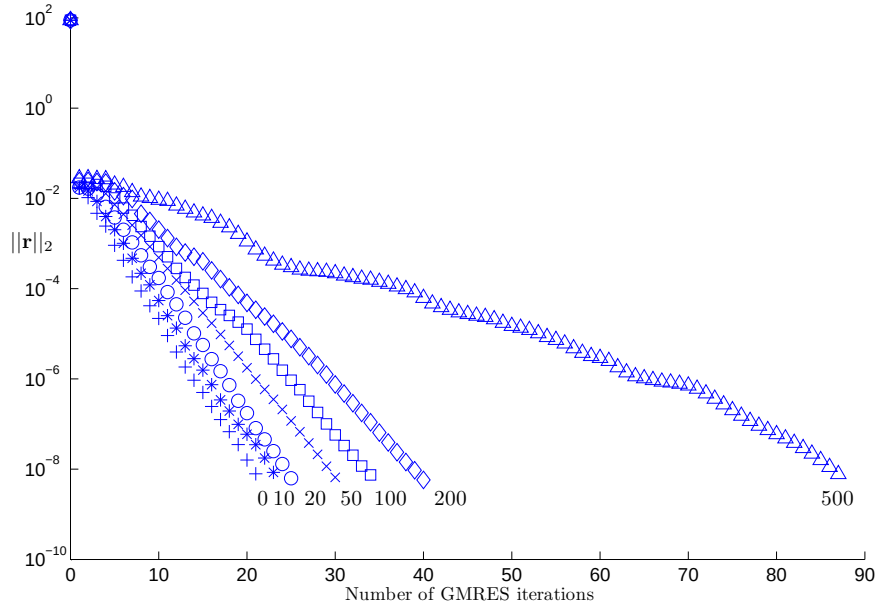

Figure 9. Convergence of the GMRES solver to compute the second Picard iterate, with the preconditioner using one AMG V-cycle for M_F and $M_S = QF_p^{-1}A_p$, for several Reynolds numbers on the mesh M4.

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 that we observe a similar behavior in the iterative scheme and in the physical observations from the transient flows obtained numerically [9]. 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 a 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, a 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 from numerical experiments. Firstly, we substitute the MINRES method by the GMRES scheme suitable for any matrices. Secondly, the symmetric positive definite block

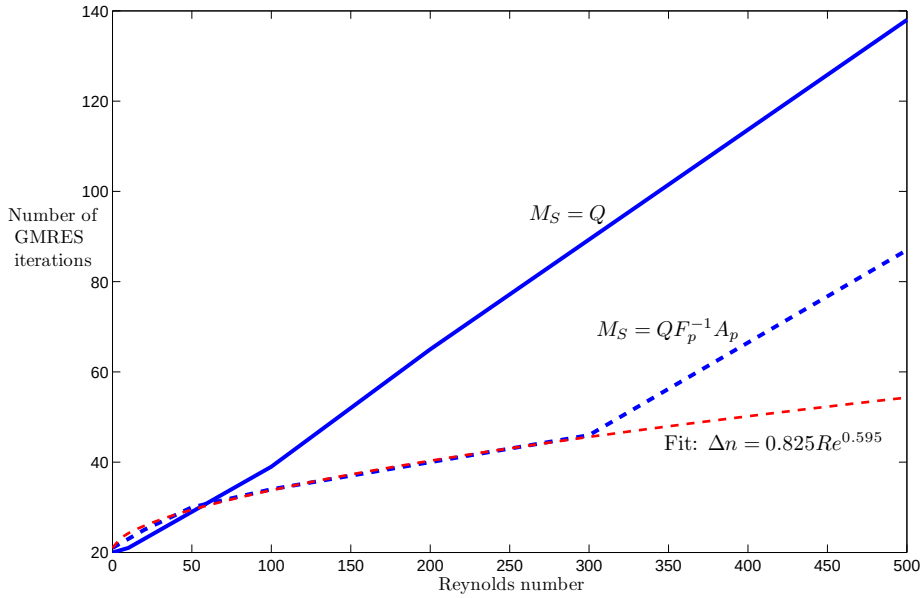

Figure 10. Number of GMRES iterations to reach convergence in function of the Reynolds number for both approximations of the Schur complement, on the mesh M4. A curve fitting of Δn, the increase in the number of iterations compared to the case $Re = 0$, is given for the pressure convection–diffusion preconditioner.

diagonal preconditioner is now replaced by a block triangular one. Therefore, new degrees of freedom appear in the definition of such a preconditioner. In this paper, the following questions are analyzed.

- Firstly, the issue of selecting an upper or an lower triangular preconditioner is investigated. Surprisingly, both structures provide 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 yields 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 yields 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.

ACKNOWLEDGEMENTS

Samuel Melchior is Research fellow with the Belgian National Fund for Scientific Research (FNRS).

The present 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.

REFERENCES

1. Axelsson O, Barker VA. *Finite element solution of boundary value problems*. Academic Press: Orlando, 1984.
2. Babuška I. The finite element method with Lagrangian multipliers. *Numerische Mathematik* 1973; **20**:179–192.
3. Balay S, Gropp WD, McInnes LC, Smith BF. Efficient Management of Parallelism in Object Oriented Numerical Software Libraries, *Modern Software Tools in Scientific Computing*, Arge E, Bruaset AM, Langtangen HP (eds). Birkhuser Press, 1997; 163–202.
4. Benzi M, Golub GH, Liesen J. Numerical solution of saddle point problems. *Acta Numerica* 2005; Cambridge University Press; 1–137.
5. Benzi M, Olshanskii MA. An augmented Lagrangian-based approach to the Oseen problem. *SIAM Journal on Scientific Computing* 2006; **28**(6):2095–2113.
6. A. Brandt, *Algebraic multigrid theory: The symmetric case*, Appl. Math. Comput., 19 (1986), pp. 23–56.
7. Brezzi F. On the existence, uniqueness and approximation of saddle point problems arising from Lagrangian multipliers. *R.A.I.R.O. Anal.Numr.*, 1974; **8**:129–154.
8. Cahouet J, Chabard JP. Some fast 3 D finite element solvers for generalized Stokes problem, *Int. J . Numer. Meth. Fluids* 1988 ; **8**:869–895.
9. Chiang TP, Sheu WH, Hwang RR. Effect of Reynolds number on the eddy strucutre in a lid-driven cavity. *Int. J. Numer. Meth. Fluids* 1998 ; **26**:557–579.
10. 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.
11. de Niet AC, Wubs FW. Two preconditioners for saddle point problems in fluid flows. *Int. J. Numer. Meth. Fluids* 2007; **54**:355–377
12. Elman HC, Silvester DJ, Wathen AJ. *Finite Elements and Fast Iterative Solvers*. Oxford University Press: Oxford, UK, 2005.
13. Gresho PM, Sani RL. *Incompressible Flow and the Finite Element Method: Volume 2: Isothermal Laminar flow*. John Wiley: Chichester, 1998.
14. 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.
15. Henson VE, Yang UM. BoomerAMG: A parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics* 2002; **41**:155–177.
16. 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.
17. Kay DA, Gresho PM, Griffiths DF, Silvester DJ. Adaptive Time-Stepping for Incompressible Flow Part II: Navier–Stokes Equations. *SIAM J. Sci Comput.* 2010; **32**(1):111–128.
18. Kay D, Loghin D, Wathen A. A preconditioner for the steady-state Navier–Stokes equations. *SIAM J. Sci Comput.* 2002; **24**(1):237–256.
19. Ladyzhenskaya O. *The Mathematical Theory of Viscous Incompressible Flow*. Gordon & Breach: New York, 1969.
20. Lipnikov K, Vassilevski Y. Parallel adaptive solution of the Stokes and Oseen problems on unstructured 3D meshes. *Parallel Computational Fluid Dynamics* 2004;153:161.
21. Murphy MF, Golub GH, Wathen AJ. A note on preconditioning for indefinite linear systems. *SIAM Journal on Scientific Computing* 2000; **21**:1969–1972.
22. 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.
23. Paige C, Saunders M. Solution of sparse indefinite systems of linear equations. *SIAM Journal on Numerical Analysis* 1975; **13**:617–629.
24. Peters J, Reichelt V, Reusken A. Fast iterative solvers for discrete Stokes equations. *SIAM Journal on Scientific Computing* 2005; **27**(2):646–666.
25. Povitsky A. Three-dimensional flow in cavity at yaw. *Nonlinear Analysis* 2005; **63**: e1573–e1584.
26. 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.

27. Rehman M, Vuik C, Segal G. A comparison of preconditioners for incompressible Navier–Stokes solvers. *Int. J. Numer. Meth. Fluids* 2008; **57**(12):1731–1751.
28. Ruge JW, Stüben K. *Algebraic multigrid (AMG)*, in Multigrid Methods, Frontiers in Applied Mathematics, S.F. McCormick, ed., SIAM, Philadelphia, 1987, pp. 73–130.
29. Rusten T, Winther R. A preconditioned iterative method for saddle-point problems. *SIAM J. Matrix Anal. Appl.* 1992; **13**:887–904.
30. Simo JC, Armero F. Unconditional stability and long-term behaviour of transient algorithms for the incompressible Navier–Stokes and Euler equations. *Comput. Methods Appl. Mech. Engrg* 1994; **111**:111–154.
31. 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.
32. Silvester D, Wathen A. Fast iterative solutions of stabilised Stokes systems. Part II: Using general block preconditioners. *SIAM J. Numer. Anal.* 1994; **31**:1352–1367.
33. Trefethen LN, Embree M. *Spectra And Pseudospectra: The Behavior of Nonnormal Matrices And Operators*. Princeton University Press: New Jersey, 2005.
34. Turek S. *Efficient Solvers for Incompressible Flow Problems*. Spinger: Berlin, 1999.
35. Varga RS. *Matrix Iterative Analysis*. Prentice-Hall, Englewood Cliffs: New Jersey, 1962. Second Edition, Springer-Verlag: New York, 2000.
36. Wathen A. Realistic eigenvalue bounds for the Galerkin mass matrix. *IMA J. Numer. Anal.* 1987; **7**:449–457.
37. Wathen A. Preconditioning and convergence in the right norm. *Int. J. Comput. Math.* 2007; **84**(8):1199–1209.
38. Wathen A, Rees T. Chebyshev semi-iteration in preconditioning for problems including the mass matrix. *Electronic Transactions on Numerical Analysis* 2009; **34**:125–135.
39. Wathen A, Silvester D. Fast iterative solution of stabilised Stokes systems. Part I: Using simple diagonal preconditioners. *SIAM J. Numer. Anal.* 1993; **30**:630–649.
40. 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.