Preconditioning and fast solvers for incompressible flow

A. J. Wathen
Oxford University Computing Laboratory

This short conference paper is a brief description with references of work that has been going on for about the last decade on fast solution methods for linearised systems arising from discretisation of the incompresible Navier-Stokes equations

$$\begin{aligned} -\nu\nabla^2\mathbf{u} + \mathbf{u}.\nabla\mathbf{u} + \nabla \mathrm{p} &= f \\ \nabla\cdot\mathbf{u} &= 0 \end{aligned}$$

which are taken to hold in some domain $\Omega \subset \mathbb{R}^d$ together with appropriate boundary conditions. We concentrate on the steady state equations; time-dependent problems are essentially easier in the context of iterative methods. We will make some comment on this later.

The numerical results we present here come from mixed finite element approximation though the techniques are not restricted by approximation method. Whichever approximation techinique is employed the discrete equations arise in the form

$$\begin{bmatrix} F(u) & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} f \\ 0 \end{bmatrix}$$

where u is a vector of velocity coefficients and p is a vector of pressure coefficients. The operators B^T and B are the discrete gradient and discrete negative divergence respectively which naturally arise as adjoints, hence the linear algebraic type notation using transpose. $F(u) = \nu A + N(u)$ is the discrete advection diffusion operator where A represents the discrete (vector) Laplacian and N the discrete advection.

In this paper we consider three different linearisations of these equations, namely

- Slow flow which leads via dropping of the quadratic inertial term to the Stokes problem

$$\begin{bmatrix} \nu A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} f \\ 0 \end{bmatrix}$$

- Picard (simple fixed point) linearisation which leads to a sequence of Oseen problems for increasing k of the form

$$\begin{bmatrix} \nu A + N(u^k) & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} u^{(k+1)} \\ p^{(k+1)} \end{bmatrix} = \begin{bmatrix} f \\ 0 \end{bmatrix}$$

- and Newton linearisation which gives a sequence of linear problems of the form

$$\begin{bmatrix} F(u^k) + M(u^k) & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \delta u^{(k+1)} \\ \delta p^{(k+1)} \end{bmatrix} = \text{residual}$$

  where $M(u) = F_u(u).u$ is a zeroth order derivative term.

Notably all of these linearised systems have a zero block in the $2, 2$ position since the continuity equation does not involve the pressure. We will refer to them collectively as discrete saddle-point systems.

Our purpose is to describe fast solution algorithms for these linearised systems. Possible approaches include

- Sparse direct (elimination) methods: these are generally excellent for discrete systems of dimension $\leq 10^4, 10^5$, but require too much memory/storage and too much computation for really large problems.
- Multigrid: many good solvers for partial differential equation problems are based in some way on multigrid techniques. The choice for these saddle-point systems is whether to use a more complicated algorithm involving for example transformations and non-standard smoothing or to incorporate simple and fast multigrid techniques as part of a solution strategy.
- Krylov subspace methods (Conjugate Gradients, MINRES , GMRES ,...): these general methods tend to work well *only if* good preconditioning is used. A particular feature is that problems with a small number of distinct eigenvalues (or at least a minimum polynomial of low degree) are very rapidly solved.

The approaches described here combine the best of the above: we use simple multigrid solvers as part of block preconditioners employed within appropriate Krylov subspace iteration. Sparse direct methods are usually competitive whenever a discrete problem is of small enough dimension.

In order to motivate preconditioning approaches for our saddle-point problems, we make use of the following theoretical observation (Murphy, Golub, Wathen (2000)). Here $S = BH^{-1}B^T$ is called the *Schur Complement.*

If the saddle-point system

$$\begin{bmatrix} H & B^T \\ B & 0 \end{bmatrix}$$

is preconditioned by the block diagonal matrix

$$\begin{bmatrix} H & 0 \\ 0 & S \end{bmatrix}$$

then the resulting preconditioned system has only the *three* distinct eigenvalues $1, \frac{1}{2} \pm \frac{\sqrt{5}}{2}$. If the same system is preconditioned by the block triangular matrix

$$\begin{bmatrix} H & B^T \\ 0 & S \end{bmatrix}$$

then the resulting presonditioned system has only the *two* distinct eigenvalues ± 1. In both cases the preconditioned matrices are diagonalisable, so that an appropriate Krylov subspace method such as MINRES or GMRES in the symmetric and nonsymmetric case respectively will *terminate* with the correct solution after respectively three or two iterations.

The requirement to achieve these precise conditions is that a solution of a linear system with the preconditioner as coefficient matrix is required at each iteration, hence an exact inversion of H and formation and exact inversion of the Schur complement S would be needed. This would be prohibitive and we do not suggest it, however this result does define a paradigm for formation of a preconditioner: if approximations to H and S are employed then there will be respectively three or two clusters of eigenvalues and Krylov subspace iterative convergence will still be very rapid. Note that the usual paradigm for construction of preconditioners is to in some sence approximate the whole matrix or its inverse: this is not what is suggested here for these saddle-point systems.

Thus, we require approximate solvers/preconditioners $\widehat{H}$ for H and $\widehat{S}$ for S. Since H represents a particular discrete differential operator, it is relatively easy to identify many good choices for $\widehat{H}$. For the sequence of problems described above, we suggest

- for Stokes: $H = \nu A$ is just the discrete vector Laplacian, so use a simple multigrid cycle for $\widehat{H}$
- for Oseen: $H = F(u^k) = \nu A + N(u^k)$ is a discrete advection-diffusion operator, so use an appropriate multigrid cycle also here. This is not so easy as for the self-adjoint Laplacian, but good methods for the advection-diffusion problem nevertheless exist (see for example Ramage (1999)).
- for Newton: $H = F(u^k) + M(u^k)$ is discrete 2nd order operator so an appropriate multigrid cycle may also be applicable in this case, though we will only show results where H is used. This is the subject of current research.

The apparently more difficult issue is how to approximate the Schur complement S without constructing it in the first place. (It is usually a dense matrix and requires H^{-1}!). We describe here some really effective approaches which lead to optimal solvers for the considered saddle-point systems.

For the Stokes problem we directly use div-stabilty (or inf-sup or Babuska-Brezzi stability) which will be a property of any discretization we employ if we wish to get numerical results which converge to the exact solution under mesh refinement. (So-called stabilization methods are also covered by these essential techniques, but for simplicity we consider only the inf-sup stable case here: see references). Boundedness is easily shown and leads to the existence of a constant $\Gamma \leq d$ for $\Omega \subset \mathbb{R}^d$ satisfying the below. Thus for a div-stable approximation we will have the existence of a positive constant γ independent of mesh size satisfying

$$\gamma \|\mathrm{p}\| \leq \sup_{\mathbf{u}} \frac{(\mathrm{p}\nabla \cdot \mathbf{u})}{\|\nabla \mathbf{u}\|} \leq \Gamma \|\mathrm{p}\|$$

which in discrete (matrix) form is preciseley

$$\gamma (p^T Q p)^{1/2} \quad \leq \quad \max_u \frac{p^T B u}{(u^T A u)^{1/2}}$$

$$= \max_{w=A^{1/2}u} \frac{p^T B A^{-1/2} w}{(w^T w)^{1/2}}$$
$$= (p^T B A^{-1} B^T p)^{1/2} \le \Gamma (p^T Q p)^{1/2}.$$

Here Q is the pressure mass matrix which is thus shown to be (spectrally) equivalent to the Schur complement $BA^{-1}B^T$ for the Stokes problem. Since Q is a well-conditioned matrix, there are very fast iterative solution methods for systems involving Q, so a natural choice for the Schur complement approximation $\widehat{S}$ is $\widehat{S} = Q$ or $\widehat{S} = diag(Q)$ (see Wathen (1987)). In fact in the numerical results presented here we use exactly four iterations of diagonally scaled Conjugate Gradients for Q as $\widehat{S}$.

The suggested preconditioner for the Stokes problem is therefore the block diagonal operator

$$\begin{bmatrix} A_{MG} & 0 \\ 0 & Q \end{bmatrix}$$

where A_{MG} is any simple multigrid cycle for the Laplacian. In the numerical results presented here we use one simple V-cycle with standard grid transfer operators and relaxed Jacobi smoothing with one pre- and one post-smoothing iteration.

Briefly the theory (which is given in Silvester and Wathen (1994)) proves that for this preconditioner convergence (in the right norm) is independent of mesh size, so that a constant number of iterations is expected for convergence to any given tolerance on any mesh.

We illustrate this is practice by showing the number of MINRES iterations for 10^{-6} residual reduction and the required CPU time in MATLAB on the same desktop computer in the table below. As mentioned above, $\widehat{H} = A_{MG}$ is one simple multigrid V-cycle and the Schur complement approximation $\widehat{S}$ is four diagonally scaled Conjugate Gradient iterations for the mass matrix Q. The problem is the standard two-dimensional driven cavity flow and the results are shown for various div-stable mixed finite elements as well as the locally stabilized $\boldsymbol{Q}_1$–$\boldsymbol{P}_0$ element. The independence of required iterations with respect to mesh size is apparent as is the applicability across many different discretizations. The approximate quadrupling of the CPU time as the number of degrees of freedom increases by a factor of four under uniform mesh refinement indicates the optimality of this solution method.

I am grateful to Rene Schneider (University of Leeds, UK) for allowing me to present also some of his results for the Stokes problem. These used

Grid	Mixed Element				
	$\boldsymbol{Q}_1$–$\boldsymbol{P}_0$	$\boldsymbol{Q}_2$–$\boldsymbol{Q}_1$	$\boldsymbol{Q}_2$–$\boldsymbol{P}_1$	$\boldsymbol{Q}_2$–$\boldsymbol{P}_0$	direct
16×16	36 (6)	31 (5)	29 (5)	25 (5)	(.3)
32×32	38 (8)	33 (10)	31 (7)	25 (6)	(3)
64×64	38 (21)	31 (21)	31 (19)	27 (16)	(31)
128×128	37 (76)	31 (74)	29 (69)	27 (59)	(221)
256×256	36 (313)	29 (309)	29 (305)	27 (267)	(8961)

a $\boldsymbol{P}_2$–$\boldsymbol{P}_1$ div-stable element on one processor of a Sun Fire 6800 cluster with UltraSPARC II Cu 900MHz processors:

degrees of freedom	iterations	CPU time for solution	for setup
659	23	2.3e-2	1.5e-1
2467	25	6.4e-1	5.0e-2
9539	25	2.0	1.5e-1
37507	25	1.2e+1	5.9e-1
148739	22	6.9e+1	2.5
592387	24	3.5e+2	1.0e+1
2364419	23	1.5e+3	4.2e+1
9447427	24	6.7e+3	1.7e+2
37769219	24	2.7e+4	6.8e+2

The optimality/scalability is even more glaringly apparent in Schneider's results as it takes 24 itertions regardless of whether a grid with 659 degrees of freedom is used or a grid of over 37 million degrees of freedom! The required CPU time also clearly scales with the numbers of degrees of freedom.

Moving on to Schur Complement Approximation for the Oseen problem, the first feature that is apparently different from the Stokes problem is the non-symmetry of $S = BF^{-1}B^T$ as $F = \nu A + N$ is an advection-diffusion operator. We give a motivational arguement for the derivation of our Oseen preconditioner here, for a more rigorous derivation see Kay, Loghin and Wathen (2002). The papers by Elman, Silvester and Wathen (2002) and Loghin (2001) are also relevant. Firstly notice that $BB^T \sim \nabla \cdot \nabla \sim QA_p$ where QA_p is the (correctly scaled) discrete Laplacian on the pressure space. That is, A_p is a pressure stiffness matrix. The key to our Oseen preconditioner is to recognise (and construct!) a matrix F_p which is similarly an advection-diffusion operator on the pressure space. Such a matrix does not arise in the original problem formulation but clearly can be constructed for any space of suitably smooth approximation functions. Considering now the underlying differential operators, we can expect that

$$FB^T \sim B^T F_p$$

which implies that

$$BB^T \sim BF^{-1}B^T F_p = SF_p$$

and so

$$S^{-1} \sim F_p(BB^T)^{-1} \sim F_p A_p^{-1} Q^{-1}$$

The differential operators commute unlike the matrices, so the best order of application of the matrices F_p, A_p^{-1} and Q^{-1} is not so clear from this arguement. Computations indicate that though there is little difference, the smallest number of iterations is achieved with the choice $\widehat{S}^{-1} = Q^{-1}F_pA_p^{-1}$ (Kay, Loghin and Wathen (2002)). The action of A_p^{-1} is approximated by a simple (Laplacian) multigrid cycle and again the action of Q^{-1} is approximated by a few (four in our computations here) diagonally scaled Conjugate Gradient iterations. Note that in the limit of Stokes flow this Schur complement approximation reduces to that given above for the Stokes problem.

Our Oseen preconditioner is therefore

$$\begin{bmatrix} F_{MG} & B^T \\ 0 & \widehat{S} \end{bmatrix}$$

with $\widehat{S}^{-1} = Q^{-1}F_pA_p^{-1}$ - notice that application of the preconditioner therefore requires just multiplication by F_p and not inversion of the pressure advection-diffusion operator and that we include the B^T here as the problem is already non-symmetric.

Briefly, the theory (see Krzyzanowski (2001), Loghin (2001), Elman, Silvester and Wathen (2002)) establishes that the eigenvalues of the preconditioned Oseen matrix are bounded independently of mesh size but have a mild dependence on the viscosity ν. Typically just a few of the eigenvalues grow like $\nu^{-\frac{1}{2}}$ (see Elman, Silvester and Wathen (2002) for a careful analysis of this). Though non-symmetric iterative methods such as the GMRES method which we use have less descriptive convergence bounds than MINRES , it is not unreasonable that the independence of eigenvalues to mesh size and mild dependence on ν should be reflected in convergence rates which similarly reflect this: this is what we see in practice.

To demonstrate this in computations in the table below we present the average number of GMRES iterations for each Picard iteration (each of which is an Oseen solve). The problem is again the common driven cavity flow and we use a $\boldsymbol{P}_2$–$\boldsymbol{P}_1$ div-stable element. The expected trends with mesh size and

	$\nu = 1/Re$			
Grid	1/10	1/40	1/160	1/640
16×16	3.00	5.85	10.66	19.54
32×32	3.14	5.85	10.88	22.27
64×64	2.83	6.16	11.12	22.22

Reynolds number (inverse viscosity) are observed.

We present also some 3-dimensional results: a driven cavity flow in the cubic cavity $[0,1]\times[0,1]\times[0,1]$ with driving top velocity $u_{\text{top}} = (1/\sqrt{3}, 2/\sqrt{3}, 0)$. For this problem we tabulate the maximum number of GMRES iterations for each Picard iteration (i.e. each Oseen solve). Here we use the div-stable $\boldsymbol{Q}_2$–$\boldsymbol{Q}_1$ element. Mesh-independence and mild ν-dependence is again evident.

	$\nu = 1/Re$			
degees of freedom	1/20	1/40	1/80	1/160
6934	31	39	49	58
15468	30	37	51	64
49072	29	36	48	64
112724	28	35	45	61

The same 3-dimensional problem was also solved with a mesh of stretched elements with a discretisation of 49072 degees of freedom. The element aspect ratio is defined as maximum edge length/minimum edge length.

We emphasise that there is a more rigorous derivation (using Greens tensors) of our preconditioner and analysis of ν-dependence (see Kay, Loghin and Wathen (2002), Elman, Silvester and Wathen (2002), Loghin and Wathen (2002)).

element aspect ratio	$\nu = 1/Re$ 1/50	1/100	1/200
1	40	52	67
2	37	49	63
4	38	47	59
8	45	49	61

Also we emphasise that any good components for the various scalar operator approximations can be used: for example these ideas have been employed using Algebraic Multigrid techniques (see Elman et al (2003))

Finally we consider a preconditioner for the Newton linearisation. As $M(u^k)$ is zeroth order operator, we use for preconditioner

$$\begin{bmatrix} F_1(u^k) + M_{1,1}(u^k) & M_{1,2}(u^k) & \multirow{2}{*}{B^T} \\ 0 & F_2(u^k) + M_{2,2}(u^k) & \\ \multicolumn{2}{c}{0} & \widehat{S} \end{bmatrix}$$

with $\widehat{S}^{-1} = Q^{-1} F_p A_p^{-1}$ as before for the Oseen problem.

The limited theory can be found in Elman, Loghin and Wathen (2003). Practical application needs approximations to $F_i(u^k) + M_{i,i}(u^k)$ as well as multigrid cycles for A_p, Conjugate Gradients for Q and construction of F_p which again is just need to perform a matrix vector multiply.

Here for simpliicity of implementation we do direct factorisation of the blocks $F_i(u^k)+M_{i,i}(u^k)$ for a driven cavity problem with $\boldsymbol{Q}_2$–$\boldsymbol{Q}_1$approximation. The average number of GMRES iterations for each Newton iteration is tabulated below.

degees of freedom	$\nu = 1/Re$ 1/10	1/40	1/160	1/640
2467	14.7	19.7	36.4	71.3
9539	12.5	17.7	37.4	67.2
37507	12.7	18.2	34.6	67.5

It would be inappropriate not to mention an alternative purely algebraic preconditioner for the Oseen and Newton problems which was derived by Elman (1999). Instead of

$$FB^T \sim B^T F_p,$$

Elman starts from

$$BFB^T \sim BB^T F_p$$

which leads to

$$(BB^T)^{-1} BFB^T A_p^{-1} Q^{-1} \sim F_p A_p^{-1} Q^{-1}$$

so

$$(BB^T)^{-1} BFB^T (BB^T)^{-1} \sim S^{-1}$$

and $(BB^T)^{-1} \sim (\nabla \cdot \nabla)^{-1}$ so that simple Laplace multigrid can be used for these two actions of the inverse. The resulting Schur complement approximation

$$\widehat{S}^{-1} = (BB^T)^{-1} BFB^T (BB^T)^{-1} \sim S^{-1}$$

is called the 'least squares commutator preconditioner' or BFBt prconditioner: again it only requires a multiply by the advection-diffusion operator F and multigrid cycles for the Laplacian. In practice mild ($O(h^{-1/2})$) mesh-dependence of convergence is observed for this algebraic preconditioner, but ν dependence is good.

Briefly considering various generalisations

- Unsteady flows: ν-dependence disappears for Oseen (Loghin (2001))

  the BFBt preconditioner is better for small time-steps than the Kay and Loghin preconditioner decribed above

  essentially time-dependent problems have an additional velocity mass matrix and are therefore easier problems for iterative methods than the steady problems considered here (Silvester et al (2001)).

- Boussinesq problem are currently under consideration.

- The relationship of the methods discussed here to projection methods have been recently described by Elman (2003).

- Approximate Schur complement idea applies generally, in particular to other saddle-point problems e.g. magnetostatics (Loghin and Wathen (2004)).

- some interesting preconditioning ideas for the Oseen problem involving the advection form have been considered by Olshanskii (1999).

The important points to note about the preconditioned iterative solution methods described here are:

- one needs only approximate solvers/preconditioners for scalar Laplacian and advection-diffusion operators
- simple multigrid for such scalar problems can be applied
- they are applicable to any (stable or stabilised) discretisation - this differs from the leading purely multigrid approaches which require approximation methods with specific properties (see Elman (1996)))
- one needs an advection-diffusion operator for pressure space: such an operator does not arise in the original problem. Without this, poor ν-dependence results.

1 Acknowledgements

Most of this work has been done incollaboration with a number of people and I would like to acknowledge Howard Elman (University of Maryland, USA), David Kay (University of Sussex, UK), Daniel Loghin (CERFACS, France), Alison Ramage (University of Strathclyde, UK) and David Silvester (UMIST, UK).

I gratefully acknowledge funding from the EPSRC and British Energy.

References

M.F. Murphy, G.H. Golub and A.J. Wathen, *A note on preconditioning for indefinite linear systems,* SIAM J. Sci. Comput. **21** (2000), pp. 1969-1972.

Ramage, A, *A multigrid preconditioner for stabilised discretisations of advection-diffusion problems,* J. Comput. Appl. Math. **101** (1999), pp. 187-203.

A.J. Wathen, *Realistic eigenvalue bounds for the Galerkin mass matrix,* IMA J. Numer. Anal. **7** (1987), pp. 449-457.

D.J. Silvester and A.J. Wathen, *Fast iterative solution of stabilised Stokes systems Part II: Using general block preconditioners,* SIAM J. Numer Anal. **31** (1994), pp. 1352-1367.

D. Kay, D. Loghin and A.J. Wathen, *A preconditioner for the steady-state Navier-Stokes equations,* SIAM J. Sci. Comput. **24** (2002), pp. 237-256.

H.C. Elman, D.J. Silvester and A.J. Wathen, *Performance and analysis of saddle-point preconditioners for the discrete Navier-Stokes equations,* Numer. Math. **90** (2002), pp. 665-688.

D. Loghin, *Analysis of preconditioned Picard iteration for the Navier-Stokes equations,* submitted to Numer. Math., 2001.

P. Krzyzanowski, *On block preconditioners for nonsymmetric saddle point problems,* SIAM J. Sci. Comput. **23** (2001), pp. 157-169.

D. Loghin and A.J. Wathen, *Schur complement preconditioners for the Navier-Stokes equations,* Int. J. Numer. Meths. Fluids **40** (2002), pp. 403-412.

H.C. Elman, V.E. Howle, J Shadid and R. Tuminaro,‘ *A parallel block multilevel for the 3D incompressible Navier-Stokes equations,* J. Comput. Phys. **187** (2003), pp. 504-523.

H.C. Elman, D. Loghin and A.J. Wathen, *Preconditioning techniques for Newton's method for the incompressible Navier-Stokes equations,* BIT **43** (2003), pp. 961-974.

H.C. Elman, *Preconditioning for the steady-state Navier-Stokes equations with low viscosity,* SIAM J. Sci. Comput. **20** (1999), pp. 1299-1316.

D.J. Silvester, H.C. Elman, D. Kay and A.J. Wathen, *Efficient preconditioning of the linearized Navier-Stokes equations,* J. Comput. Appl. Math. **128** (2001), pp. 261-279.

H.C. Elman, *Preconditining strategies for models of incompressible flow,* University of Maryland, Computer Science report CS-TR-4543, (2003).

D. Loghin and A.J. Wathen, *Analysis of preconditioners for saddle-point problems,* accepted for publication in SIAM J. Sci. Comput. , 2004.

M. Olshanskii, *An iterative solver for the Oseen problem and numerical solution of incompressible Navier-Stokes equations,* Numer. Linear Algebra Appl. **6** (1999), pp. 353-378.

H.C. Elman, *Multigrid and Krylov subspace methods for discrete Stokes equations,* Int. J. Numer. Meths. Fluids **22** (1996), pp. 755-770.