

Technical Report

TR-2009-001

New Multigrid Smoothers for the Oseen Problem

by

S. Hamilton, M. Benzi, E. Haber

MATHEMATICS AND COMPUTER SCIENCE

EMORY UNIVERSITY

NEW MULTIGRID SMOOTHERS FOR THE OSEEN PROBLEM

STEVEN HAMILTON*, MICHELE BENZI†, AND ELDAD HABER‡

Abstract. We investigate the performance of smoothers based on the Hermitian/skew-Hermitian (HSS) and augmented Lagrangian (AL) splittings applied to MAC discretizations of the Oseen problem. Both steady and unsteady flows are considered. Local Fourier analysis and numerical experiments on a 2-D lid-driven cavity problem indicate that the proposed smoothers result in h -independent convergence and are fairly robust with respect to the Reynolds number.

Key words. multigrid, smoothing iterations, generalized Stokes and Oseen problems, incompressible Navier–Stokes equations

AMS subject classifications. Primary 65F10, 65N22, 65F50. Secondary 76M.

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

$$\frac{\partial \mathbf{u}}{\partial t} - \nu \Delta \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p = \mathbf{f} \quad \text{on } \Omega \times (0, T] \quad (1.1)$$

$$\operatorname{div} \mathbf{u} = 0 \quad \text{on } \Omega \times [0, T] \quad (1.2)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega \times [0, T] \quad (1.3)$$

$$\mathbf{u}(\mathbf{x}, 0) = \mathbf{u}_0(\mathbf{x}) \quad \text{on } \Omega \quad (1.4)$$

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

$$\sigma \mathbf{u} - \nu \Delta \mathbf{u} + (\mathbf{v} \cdot \nabla) \mathbf{u} + \nabla p = \mathbf{f} \quad \text{in } \Omega \quad (1.5)$$

$$\operatorname{div} \mathbf{u} = 0 \quad \text{in } \Omega \quad (1.6)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega \quad (1.7)$$

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

Spatial discretization of the preceding equations using finite differences or finite elements results in a large, sparse saddle point system of the form

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} f \\ g \end{bmatrix} \quad (1.8)$$

*Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA (sphamil@emory.edu). The work of this author was supported in part by a U. S. Department of Energy Computational Science Graduate Fellowship.

†Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA (benzi@mathcs.emory.edu). The work of this author was supported in part by the National Science Foundation Grant DMS-0511336.

‡Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA (haber@mathcs.emory.edu). The work of this author was supported in part by the National Science Foundation Grant DMS-0724759 and by the U. S. Department of Energy Grant DMS 0724759.

where u and p represent the discrete velocity and pressure, respectively, A is the discretization of the diffusion, convection, and time-dependent terms, B^T is the discrete gradient, B the (negative) discrete divergence, and f and g contain forcing and boundary terms. Here we assume that the discretization satisfies the LBB (‘inf-sup’) stability condition, so that no pressure stabilization is required; see, e.g., [10].

The efficient solution of (1.8) calls for rapidly convergent iterative methods. Much work has been done in developing efficient preconditioners for Krylov subspace methods applied to this problem; see, e.g., [3, 5, 8, 9, 10, 11, 14]. Coupled multigrid methods have also been developed; see, e.g., [20, 22, 23] and the references therein. The ultimate goal is to develop robust solvers with optimal complexity. In particular, the rate of convergence should be (asymptotically) independent of the mesh size h and of the kinematic viscosity ν (equivalently, of the Reynolds number Re). As mentioned in [23], one of the main challenges in incompressible CFD is the construction of smoothers that are robust over a wide range of values of the viscosity, in particular for small values of ν ; see also [15] where the difficulty of smoothing for the Navier–Stokes equations with low viscosity is pointed out.

There are two main classes of smoothers for incompressible flow problems: fully coupled, ‘box’ smoothers like Vanka’s method [21] and segregated, distributive relaxation schemes like SIMPLE and related approaches (see, e.g., [22, Chapter 7.6] for an overview). Vanka’s method is often found to be superior to other smoothers, but it sometimes fails to deliver h -independent convergence for hard problems and small values of the viscosity; see, e.g., [5, Table 6.4]. Hence, there is a strong interest in developing smoothers that exhibit good robustness over a wide range of problem parameters.

In this paper we investigate two types of multigrid smoothers, one based on the Hermitian and skew-Hermitian (HSS) splitting and the other a block triangular smoother based on the augmented Lagrangian (AL) formulation of the saddle point problem (1.8). Such splittings have been intensively studied in recent years in the context of developing preconditioners for Krylov subspace methods, with very good results; see, e.g., [1, 3, 4, 5, 19]. Their use as smoothers for multigrid has not been previously investigated.

2. HSS smoothing. Any matrix can be uniquely written as the sum of its Hermitian and skew-Hermitian (symmetric and skew-symmetric for real-valued matrices) components, $A = \frac{1}{2}(A + A^T) + \frac{1}{2}(A - A^T) = H + S$. From this decomposition, we can define two different splittings of the matrix A by shifting the symmetric and skew-symmetric component by some parameter α :

$$\begin{aligned} A &= (H + \alpha I) - (\alpha I - S) \\ A &= (S + \alpha I) - (\alpha I - H). \end{aligned} \tag{2.1}$$

From these splittings, the HSS iteration was defined in [1] by:

$$\begin{aligned} (H + \alpha I)x^{(k+\frac{1}{2})} &= (\alpha I - S)x^{(k)} + b \\ (S + \alpha I)x^{(k+1)} &= (\alpha I - H)x^{(k+\frac{1}{2})} + b \end{aligned} \tag{2.2}$$

for $k = 0, 1, \dots$, with $x^{(0)}$ arbitrary. Here α is a positive shift parameter. Recently, it has been shown that the corresponding operator $P = (H + \alpha I)(S + \alpha I)$ can be an effective preconditioner for Krylov methods for saddle point problems [3, 4]. The optimal selection of the parameter α (as a preconditioner) has been studied in, e.g.,

[2, 19]. In practice, when carrying out (2.2) inexact solves are sufficient to achieve good convergence rates, making the overall approach practically feasible.

For the solution of the Oseen problem, we consider a slight modification to equation (1.8) in which a negative sign is placed before the (2,1) block, resulting in the equivalent system

$$\begin{bmatrix} A & B^T \\ -B & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} f \\ -g \end{bmatrix}. \quad (2.3)$$

While the coefficient matrix in (1.8) is indefinite (its eigenvalues fall on both sides of the imaginary axis), the one in (2.3) has its spectrum entirely contained in the right half complex plane [3].

For many discretizations (e.g., the Marker-and-Cell, or MAC, discretization [13]), the HSS splitting almost exactly corresponds to the natural (i.e., physical) splitting of the relevant differential operators:

$$\begin{aligned} H &= \begin{bmatrix} \sigma I + L & 0 \\ 0 & 0 \end{bmatrix} \\ S &= \begin{bmatrix} K & B^T \\ -B & 0 \end{bmatrix} \end{aligned} \quad (2.4)$$

where L represents the discretization of the Laplacian, σI corresponds to the time derivative (for finite elements the identity is replaced by the velocity mass matrix) and K is the discretization of the convective term; B and B^T retain their previous meaning. Matrix H is clearly symmetric. For a constant coefficients problem, the convective term is truly skew-symmetric up to boundary conditions. For a general non-constant coefficients problem, however, the convective term isn't exactly skew-symmetric as entries in structurally symmetric positions involve the wind function evaluated at neighboring grid points. Thus, for a continuous wind function the convective term approaches a skew-symmetric term as the mesh is refined. In practice, one can either split the original coefficient matrix algebraically so that the resulting terms are strictly symmetric and skew-symmetric, or it can be split based on the physical terms as in equations (2.4). Though the theoretical analyses on the HSS approach only strictly hold for the algebraic splitting, experimental observations indicate that there is very little difference in the behavior between the two approaches.

Under the assumptions of constant coefficients and periodic boundary conditions, we can perform a local Fourier analysis (LFA) on the HSS iteration with the marker-and-cell discretization to determine its potential as a multigrid smoother and aid in the selection of the free parameter α . The iteration matrix describing equations (2.2) is given by $T = (S + \alpha I)^{-1}(H - \alpha I)(H + \alpha I)^{-1}(S - \alpha I)$. Applying the symbols of the 2-D operators, one finds that the frequency-dependent smoothing rate is given by

$$L_h(\theta_x, \theta_y) = \left(\frac{\alpha + \sigma - \tilde{\Delta}_h}{\alpha + \sigma + \tilde{\Delta}_h} \right)^2 \quad (2.5)$$

where $\tilde{\Delta}_h = \frac{\nu}{h^2}(4 - 2\cos(\theta_x) - 2\cos(\theta_y))$ is the symbol of the discrete 2-D Laplacian (scaled by the viscosity ν). Note that this smoothing rate depends only on the diffusive and not the convective term, although it does depend on the viscosity ν . We note in passing that for the Stokes problem ($\mathbf{v} = \mathbf{0}$ in the Oseen problem) one can take $\nu = 1$.

In selecting a multigrid smoother, we wish to choose one which damps errors the most uniformly over high frequency regions. To accomplish this we wish to minimize

	$Re \equiv \nu^{-1}$			
$1/h$	256	512	1024	2048
64	2.7	4.1	6.1	9.1
128	2.1	2.8	4.2	6.5
256	1.6	2.1	2.8	4.3
512	1.4	1.6	2.1	2.8

TABLE 2.1

Ratio of nonzeros in ILU factors over nonzeros in $S + \alpha I$

the quantity $\mu = \sup(L_h(\Theta))$, $\Theta \in [-\pi, \pi] \setminus [-\pi/2, \pi/2]$ which is referred to as the smoothing factor. Taking $\alpha + \sigma = \frac{4\nu}{h^2}$, we find that the smoothing factor is minimized and takes a value of $\mu = \frac{1}{9}$ (independent of ν). For a steady problem ($\sigma = 0$) this reduces to simply $\alpha = \frac{4\nu}{h^2}$ whereas for unsteady problems the optimal parameter is slightly reduced (typically $\sigma = O(h^{-1})$). A smoothing factor of $\frac{1}{9}$ indicates that all high frequency components of the error are attenuated by a factor of 9 for each iteration of HSS. This value is comparable to or better than many commonly used smoothers [24], raising the hope that HSS smoothing might result in a competitive multigrid solver.

In order for HSS to become a feasible method it is necessary to apply an inexact variant with a small computational cost which does not significantly degrade the properties of the exact operator. A single iteration of HSS requires (approximately) solving linear systems with both $H + \alpha I$ and $S + \alpha I$. The matrix $H + \alpha I$ is symmetric positive definite and extremely well-conditioned. In fact, the 2-norm condition number is less than 3 for all mesh sizes and viscosities when α is taken to be the value predicted by Fourier analysis, $\alpha = \frac{4\nu}{h^2}$. A single iteration of conjugate gradients with a zero fill-in incomplete Cholesky preconditioner is sufficient to maintain the effectiveness of the method. The shifted skew-symmetric system, on the other hand, poses a more significant problem. We use a fixed number (in this study, 5) of preconditioned GMRES iterations [17] with a thresholded incomplete LU factorization as the preconditioner. We have found that reordering the original matrix (e.g., with reverse Cuthill-McKee) is necessary to maintain robustness with respect to the mesh size and viscosity, an observation consistent with the findings of [6]. With such reordering, the same value $\tau = 0.01$ of the ILU drop tolerance was used in all cases. Table 2.1 illustrates the storage required for the incomplete factors over a range of problem parameters. A moderate increase in the storage requirement is seen as the viscosity is reduced, but there is actually a decrease in the level of fill-in as the mesh size is reduced. Thus, for a fixed viscosity the total cost per (inexact) HSS smoothing step is linear in the number of unknowns.

It is important to point out a major difference between the use of HSS as a smoother for multigrid and that as a preconditioner for a Krylov subspace method. As shown in [2, 3, 19], the use of HSS as a preconditioner requires that the parameter α should be chosen small; for many problems, the optimal α goes to zero as $h \rightarrow 0$. In contrast, as we just saw, when HSS is used as a smoother for the Oseen (or Stokes) problem the optimal value of α grows like $O(h^{-2})$ as $h \rightarrow 0$.

3. AL smoothing. We begin the discussion of the augmented Lagrangian formulation by replacing the original system (1.8) with the following equivalent system:

$$\begin{bmatrix} A + \gamma B^T W^{-1} B & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} \hat{f} \\ g \end{bmatrix} \quad (3.1)$$

where $\hat{f} = \gamma B^T W^{-1} g$; here γ is a parameter and W is a positive definite matrix, frequently taken to be the pressure mass matrix or a diagonal approximation thereof, see [12]. It is of interest to note that the (1,1) block of the augmented system (3.1) resembles that of the poroelasticity equations described in [15] (though the poroelasticity equations lack a convective term). We will denote the coefficient matrix in the preceding equation by \mathcal{A} . Now we consider a preconditioner of the form

$$\mathcal{P} = \begin{bmatrix} \hat{A}_\gamma & B^T \\ 0 & -\frac{1}{\gamma} W \end{bmatrix} \quad (3.2)$$

where the application of \hat{A}_γ^{-1} involves the (inexact) inversion of $A + \gamma B^T W^{-1} B$. It was shown in [5] that the eigenvalues of $\mathcal{P}^{-1} \mathcal{A}$ all tend to 1 as $\gamma \rightarrow \infty$ (uniformly in h). However, since B (and thus $B^T W^{-1} B$) has a significant null space, $A + \gamma B^T W^{-1} B$ becomes very ill-conditioned for large γ and thus finding an effective approximation to it is problematic. Thus taking γ to be a moderate value, say $O(1)$, is frequently a better strategy, see [5].

Here we consider the use of \mathcal{P} as a smoother rather than as a preconditioner. It can be shown that this is a non-standard form of distributed relaxation. The crux to establishing an efficient AL smoother lies in the definition of \hat{A}_γ . One possibility is to implicitly define \hat{A}_γ^{-1} in terms of a multigrid cycle, however efficient smoothing is difficult due to the aforementioned null space of the matrix B . In [5], a highly effective and robust geometric multigrid solver was developed to address such difficulties, based on the one presented in [18]. Unfortunately, implementation of this multigrid iteration is less than straightforward, particularly on unstructured meshes. As simpler alternatives, we currently consider taking \hat{A}_γ to be either the block upper triangular or upper triangular portion of $A + \gamma B^T W^{-1} B$. Note that A , and hence $A + \gamma B^T W^{-1} B$ has a natural 2-by-2 block structure, and that inversion of the block upper triangular matrix \hat{A}_γ requires the (approximate) solution of two scalar anisotropic convection-diffusion equations with anisotropy ratio $1 + \frac{\gamma}{\nu}$. In the experiments below we solved these ‘exactly’ by a direct method, but in practice an approximate iterative solver could be used. Efficient iterative solvers for such problems can be found in literature, see for instance [7]. In the upper triangular case, when W is a diagonal matrix then the preconditioner \mathcal{P} as a whole is upper triangular and solving systems involving the preconditioner becomes trivial. The asymptotic cost per iteration for the inexact AL smoother is thus linear in the number of unknowns for the triangular case and though our implementation of the block upper triangular case is not $O(n)$ at the current time, such an implementation is in principle possible in both the steady and unsteady case. Below we investigate the use of \mathcal{P} or one of these approximations as a smoother for a coupled multigrid method for the discrete Oseen problem.

4. Results. We consider the marker-and-cell (staggered-grid finite difference) discretized Oseen problem on the unit square. As a test problem, we take the standard leaky-lid driven cavity problem described, for instance, in [10]. Homogeneous Dirichlet boundary conditions are prescribed for all velocity components with the exception of a

	$Re \equiv \nu^{-1}$			
$1/h$	256	512	1024	2048
64	15	48	63	72
128	15	23	72	104
256	13	25	38	151
512	10	18	37	51

TABLE 4.1

Iteration count for HSS multigrid on steady Oseen problem

positive unit horizontal velocity along the top edge. To approximate the solution of a single Picard iteration, we take the wind function to be the rotating vortex described by

$$\mathbf{v}(x, y) = \begin{bmatrix} 8x(x-1)(1-2y) \\ 8(2x-1)y(y-1) \end{bmatrix}. \quad (4.1)$$

For the solver, we use FGMRES [16] preconditioned with one multigrid cycle. In our experiments, we found that FGMRES acceleration resulted in a more robust solver than using the multigrid iteration alone. In all cases we used the zero vector as the initial guess and a reduction of the 2-norm of the initial residual by four orders of magnitude as the stopping criterion. The cycle is chosen to be a V-cycle [20] with one pre-smoothing and one post-smoothing step. For HSS, a smoothing step is simply one full HSS iteration. For the augmented Lagrangian approach, a smoothing step is a single Richardson iteration on the preconditioned system, i.e. $x^{k+1} = x^k + \mathcal{P}^{-1}r^k$ where $r^k = b - \mathcal{A}x^k$ is the residual and \mathcal{P} is as defined earlier. In all cases, the mesh is refined using a standard coarsening in which the mesh size is doubled in both the x and y directions. The coarse mesh problem is obtained by re-discretizing the underlying problem on the coarser grid. A series of successively coarser grids is used with the coarsest grid being that for which $h = \frac{1}{2}$. An exact solver is used on the coarsest grid. The velocity restriction operators are chosen to be a 1-D full weighting in the direction of the velocity component and linear interpolation in the orthogonal direction. The pressure restriction operator is given by bilinear interpolation. These restriction operators are described by the following stencils:

$$R_{u_x} = \frac{1}{8} \begin{bmatrix} 1 & 2 & 1 \\ 1 & 2 & 1 \end{bmatrix} \quad R_{u_y} = \frac{1}{8} \begin{bmatrix} 1 & 1 \\ 2 & 2 \\ 1 & 1 \end{bmatrix} \quad R_p = \frac{1}{4} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \quad (4.2)$$

The prolongation operators are selected to be the scaled transposes of the corresponding restriction operators. These grid transfer operators are consistent with those suggested in [20] for staggered-grid discretizations.

Table 4.1 shows the results for the steady Oseen problem with HSS smoothing. The convergence behavior is extremely robust with respect to decreasing mesh size, even showing a decreasing trend for most viscosities. A noticeable degradation in performance is observed with respect to decreasing viscosity, which is not surprising. The degradation appears to be similar to that experienced by other solvers in literature [9, 14], and at the finest mesh size the degradation is actually quite manageable. It should be mentioned that since no velocity stabilization is being attempted, for small ν only numerical solutions corresponding to the finest mesh are meaningful.

	$Re \equiv \nu^{-1}$			
$1/h$	256	512	1024	2048
64	16	23	37	46
128	12	17	27	45
256	8	10	16	29
512	5	7	9	15

TABLE 4.2

Iteration count for exact AL multigrid on steady Oseen problem

	$Re \equiv \nu^{-1}$			
$1/h$	256	512	1024	2048
64	38	77	170	413
128	22	43	100	227
256	13	20	37	88
512	8	12	19	32

TABLE 4.3

Iteration count for inexact (block upper triangular) AL multigrid on steady Oseen problem

For all parameter combinations, the parameter α is selected to be the value predicted by the Fourier analysis described above. Since α depends on h , it is necessary to redefine α on each grid level to achieve optimal smoothing for that mesh size. Experimenting with different values of the free parameter α indicates that the prediction from the LFA is the optimal choice (or very nearly so) across all problem parameters.

The performance of the augmented Lagrangian smoother on the same problem is shown in Table 4.2 for the exact case and in Table 4.3 for the inexact case. ‘Inexact’ here refers to the selection of \hat{A}_γ as the block upper triangular portion of $A + \gamma B^T W^{-1} B$ as discussed in the previous section. For the exact AL smoother the value $\gamma = 1$ was used to produce the results. In the inexact variant, a slightly smaller value was found to produce better results, and $\gamma = 0.1$ was used instead. In both variants, however, a single value for γ is used for all mesh sizes and viscosities, effectively resulting in a parameter-free smoother; slightly better results can be obtained by fine-tuning the free parameter γ . Results for the upper triangular variant are not shown here as convergence was not achieved in 500 iterations for most of the problem parameters shown. The augmented Lagrangian multigrid displays the same h -independent convergence as that seen with the HSS smoother. The dependence on the viscosity is actually quite weak in the exact case, exhibiting better behavior than the HSS results. For the inexact AL, the convergence is worse than in the exact case when the mesh size is moderate, but at the finer meshes, which are the ones needed to achieve an acceptable resolution of the computed flow, the increase in iteration count as ν decreases is more than compensated by the reduced computational effort required by the inexact smoother. Also, looking at the numbers on the main diagonal of Tables 4.2-4.3 one can see that the iteration count is essentially independent of the ratio ν/h , a highly desirable property.

For the next set of tests, we consider an unsteady Oseen problem. The underlying problem remains the same as before, except now a multiple (σ) of the identity is added to the (1,1) block of the coefficient matrix. Here we take $\sigma = h^{-1}$. The results for HSS and (inexact) AL multigrid are shown in Tables 4.4 and 4.5, respectively. The results for the HSS multigrid are better than in the steady case. The viscosity

	$Re \equiv \nu^{-1}$				
$1/h$	256	512	1024	2048	4096
64	11	19	18	50	51
128	9	15	12	35	70
256	8	12	10	28	26
512	7	9	8	20	16

TABLE 4.4

Iteration count for HSS multigrid on unsteady Oseen problem, $\sigma = h^{-1}$

	$Re \equiv \nu^{-1}$					
$1/h$	256	512	1024	2048	4096	8192
64	14/20	15/21	16/22	17/22	17/23	17/23
128	12/18	13/18	13/19	14/19	15/19	15/19
256	11/14	11/15	11/15	12/15	12/16	12/16
512	11/16	11/17	11/17	12/17	12/17	12/17

TABLE 4.5

Iteration count for inexact (block upper triangular/upper triangular) AL multigrid on unsteady Oseen problem, $\sigma = h^{-1}$

dependence is less pronounced, as is to be expected for the less difficult unsteady problem. The inexact AL multigrid results of Table 4.5 are quite remarkable: the convergence shows no degradation with respect to decreases in either the mesh size or the viscosity—even for viscosities which are quite small. This behavior is observed for both the block upper triangular and the upper triangular approximations to the AL preconditioner. Such robustness with respect to problem parameters combined with the small computational cost per iteration places the inexact AL multigrid solver among the most effective unsteady Oseen solvers in literature.

5. Conclusions and future work. We have investigated some new smoothers for the coupled multigrid solution of the discrete Oseen problem. The smoothers are based on the HSS splitting and on the augmented Lagrangian formulation of the discrete equations, respectively. In practice, the smoothers are applied inexactly so that the cost per smoothing step is $O(n)$, where n is total number of unknowns.

Although still preliminary, our analysis and numerical experiments indicate that the new smoothers are quite promising, showing h -independent behavior (for h sufficiently small) in all cases. The robustness with respect to decreasing viscosity is also good, and indeed excellent for the AL-based smoother. Especially good performance is observed in the unsteady case. For moderate Reynolds numbers and in the limiting case of the Stokes and generalized Stokes problem (not discussed here), both the HSS and the AL-based smoothers are extremely effective.

Future work includes an analysis of the AL-based smoothers, extension to the 3-D case, and implementation and testing for more complicated problems. In particular, the performance of the smoothers for problems with high cell aspect ratios (stretched grids) and for higher order discretizations needs to be investigated. In addition, comparisons with other smoothers that have been proposed in the literature should be carried out. An investigation into the relationship between the augmented Lagrangian formulation and the poroelasticity equations of [15] may also be of interest, especially with regards to the distributed relaxation smoothers described therein. It is also expected that the use of coarse grid velocity stabilization, which was not considered

here, will improve the multigrid performance.

REFERENCES

- [1] Z. Z. BAI, G. H. GOLUB, AND M. K. NG, *Hermitian and skew-Hermitian splitting methods for non-Hermitian positive definite linear systems*, SIAM J. Matrix Anal. Appl., 24 (2003), pp. 603–626.
- [2] M. BENZI, M. J. GANDER, AND G. H. GOLUB, *Optimization of the Hermitian and skew-Hermitian splitting iteration for saddle-point problems*, BIT Numerical Mathematics, 43 (2003), pp. 881–900.
- [3] M. BENZI AND G. H. GOLUB, *A preconditioner for generalized saddle point problems*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 20–41.
- [4] M. BENZI AND J. LIU, *An efficient solver for the incompressible Navier–Stokes equations in rotation form*, SIAM J. Sci. Comput., 29 (2007), pp. 1959–1981.
- [5] M. BENZI AND M. A. OLSHANSKII, *An augmented Lagrangian-based approach to the Oseen problem*, SIAM J. Sci. Comput., 28 (2006), pp. 2095–2113.
- [6] M. BENZI, D. B. SZYLD, AND A. VAN DUIN, *Orderings for incomplete factorization preconditioning of nonsymmetric problems*, SIAM J. Sci. Comput., 20 (1999), pp. 1652–1670.
- [7] D. BERTACCINI, G. H. GOLUB, AND S. SERRA-CAPIZZANO, *Spectral analysis of a preconditioned iterative method for the convection-diffusion equation*, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 260–278.
- [8] H. C. ELMAN, *Preconditioners for saddle point problems arising in computational fluid dynamics*, Appl. Numer. Math., 43 (2002), pp. 75–89.
- [9] H. C. ELMAN, *Preconditioning for the steady-state Navier–Stokes equations with low viscosity*, SIAM J. Sci. Computing 20 (1999), pp. 1299–1316.
- [10] H. ELMAN, D. SILVESTER, AND A. WATHEN, *Finite Elements and Fast Iterative Solvers with Applications in Incompressible Fluid Dynamics*, Oxford University Press, Oxford, UK, 2005.
- [11] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Performance and analysis of saddle point preconditioners for the discrete steady-state Navier–Stokes equations*, Numer. Math. 90 (2002), pp. 665–688.
- [12] M. FORTIN AND R. GLOWINSKI, *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems*, Stud. Math. Appl. 15, North-Holland, Amsterdam, New York, Oxford, 1983.
- [13] F. H. HARLOW AND J. E. WELCH, *Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface*, Phys. Fluids, 8 (1965), pp. 2182–2189.
- [14] M. A. OLSHANSKII AND Y. V. VASSILEVSKI, *Pressure Schur complement preconditioners for the discrete Oseen problem*, SIAM J. Sci. Comput., 29 (2007), pp. 2686–2704.
- [15] C. W. OOSTERLEE AND F. J. GASPAS, *Multigrid relaxation methods for systems of saddle point type*, Appl. Numer. Math., 58 (2008), pp. 1933–1950.
- [16] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM J. Sci. Comput., 14 (1993), pp. 461–469.
- [17] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.
- [18] J. SCHÖBERL, *Multigrid methods for a parameter dependent problem in primal variables*, Numer. Math., 84 (1999), pp. 97–119.
- [19] V. SIMONCINI AND M. BENZI, *Spectral properties of the Hermitian and skew-Hermitian splitting preconditioner for saddle point problems*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 377–389.
- [20] U. TROTTEBERG, C. OOSTERLEE, AND A. SCHULLER, *Multigrid*, Academic Press, San Diego, 2001.
- [21] S. P. VANKA, *Block-implicit multigrid solution of Navier–Stokes equations in primitive variables*, J. Comput. Phys., 65 (1986), pp. 138–158.
- [22] P. WESSELING, *Principles of Computational Fluid Dynamics*, Springer Series in Computational Mathematics 29, Springer, New York, 2001.
- [23] P. WESSELING AND C. W. OOSTERLEE, *Geometric multigrid with applications to computational fluid dynamics*, J. Comput. Appl. Math., 128 (2001), pp. 311–334.
- [24] R. WIENANDS AND W. JOPPICH, *Practical Fourier Analysis for Multigrid Methods*, Chapman & Hall, New York, 2005.