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 two-dimensional lid-driven cavity problem indicate that the proposed smoothers result in h-independent convergence and are fairly robust with respect to the Reynolds number. A direct comparison shows that the new smoothers compare favorably to coupled smoothers of Braess–Sarazin type, especially in terms of scaling for increasing 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

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

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

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

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

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

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

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

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

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

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

*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., [11].

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, 9, 10, 11, 12, 16]. Coupled multigrid methods have also been developed; see, e.g., [22, 24, 25] 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 [25], 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 [17] 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 [23] and segregated, distributive relaxation schemes like SIMPLE and related approaches (see, e.g., [24, 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, 21]. 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 $k = 0, 1, \ldots$, 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, 21]. 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

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

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 [14]), the HSS splitting almost exactly corresponds to the natural (i.e., physical) splitting of the relevant differential operators:

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

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 is not 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 $T = (S + \alpha I)^{-1}(\alpha I - H)(H + \alpha I)^{-1}(\alpha I - S)$. For a two-dimensional (2-D) problem, we denote the symbol of the discrete Laplacian by $\tilde{\Delta}_h$, the symbol of the convective term by $\widetilde{(v \cdot \nabla)}_h$, and the symbols of the x- and y- derivatives by $\tilde{\delta}_x$ and $\tilde{\delta}_y$, respectively. For the MAC discretization, these symbols equate to $\tilde{\Delta}_h = \frac{1}{h^2}(4 - 2\cos\theta_x - 2\cos\theta_y)$, $\widetilde{(v \cdot \nabla)}_h = \mathrm{i}\,(v_x \sin\theta_x/2 + v_y \sin\theta_y/2)$, and $\tilde{\delta}_{x,y} = \mathrm{i}\,\sin\theta_{x,y}$, where $\mathrm{i} = \sqrt{-1}$. The LFA smoothing rate is defined as $L_h = \rho(\tilde{T})$, where ρ indicates the spectral radius. For the HSS iteration, one finds that the dominant eigenvalue of the symbol of the iteration matrix can be written as

$$\lambda = \frac{(\alpha - \tilde{\Delta}_h)(\alpha - \widetilde{v \cdot \nabla})}{(\alpha + \tilde{\Delta}_h)(\alpha + \widetilde{v \cdot \nabla})} \tag{2.5}$$

Because α is a real number and $\widetilde{(v \cdot \nabla)}_h$ is pure imaginary, it is easily seen that

$$\left| \frac{\alpha - \widetilde{(v \cdot \nabla)}_h}{\alpha + \widetilde{(v \cdot \nabla)}_h} \right| = 1 \tag{2.6}$$

and thus

$$L_h(\theta_x, \theta_y) = \left| \frac{\alpha - \tilde{\Delta}_h}{\alpha + \tilde{\Delta}_h} \right| \tag{2.7}$$

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 steady 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 the quantity $\mu = \sup(L_h(\Theta)),\ \Theta \in [-\pi, \pi] \backslash [-\pi/2, \pi/2]$ which is referred to as the smoothing factor. Taking $\alpha = \frac{4\nu}{h^2}$, we find that the smoothing factor is minimized and takes a value of $\mu = \frac{1}{3}$ (independent of ν!). For an unsteady problem, the optimal parameter becomes $\alpha = \frac{4\nu}{h^2} + \sigma$, where σ is the constant related to the time step. A smoothing factor of $\frac{1}{3}$ indicates that all high frequency components of the error are attenuated by a factor of 3 for each iteration of HSS. This value is comparable to many commonly used smoothers [26], 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 [19] 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, 21], 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 $h \to 0$. 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 $h \to 0$.

	$Re \equiv \nu^{-1}$			
$1/h$	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 $S + \alpha I$

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

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

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 [13]. 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 [17] (though the poroelasticity equations lack a convective term).

We will denote the coefficient matrix in the preceeding equation by $\mathcal{A}$. Now we consider a preconditioner of the form

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

where the application of $\hat{A}_\gamma^{-1}$ 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 $\gamma \to \infty$ (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].

Defining $T = I - \mathcal{P}^{-1}\mathcal{A}$ (a Richardson iteration on the preconditioned system), we consider the use of AL as a smoother rather than as a preconditioner. As previously described for the HSS iteration,we can perform a local Fourier analysis for this iteration matrix. For the same MAC-discretized steady problem as before, we find the dominant eigenvalue (two eigenvalues are identically zero) of the symbol of the iteration matrix is given by

$$\lambda = \left(\frac{\nu^2 (\tilde{\Delta}_h)^2 + (\widetilde{v \cdot \nabla})_h^2}{(\nu + \gamma)^2 \tilde{\Delta}_h^2 + (\widetilde{v \cdot \nabla})_h^2} \right)^{1/2} \tag{3.3}$$

and the smoothing rate is simply the modulus of this eigenvalue. Though this expression does not reduce to a simple expression as in the case of the HSS analysis, we can identify a few pertinent features of this quantity. First of all, the smoothing rate is strictly less than unity for any positive values of ν and γ, indicating that this iteration does in fact possess smoothing properties. Secondly, the smoothing rate tends to zero as γ tends to infinity. As mentioned earlier, large values of γ present difficulties when attempting to invert $\mathcal{P}$ inexactly and thus we anticipate that selecting a moderate value of γ will likely be necessary in a practical implementation.

It can be shown that the AL smoothing iteration is a non-standard form of distributed relaxation. The crux to establishing an efficient AL smoother lies in the definition of $\hat{A}_\gamma$. One possibility is to implicitly define $\hat{A}_\gamma^{-1}$ in terms of a multigrid cycle, however efficient smoothing is difficult due to the aforementioned null space of the matrix B (and thus $B^T W^{-1} B$). In [5], a highly effective and robust geometric multigrid solver was developed to address such difficulties, based on the one presented in [20]. Unfortunately, implementation of this multigrid iteration is less than straightforward, particularly on unstructured meshes. As simpler alternatives, we currently consider taking $\hat{A}_\gamma$ 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 $\hat{A}_\gamma$ 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 [11]. Homogeneous Dirichlet boundary conditions are prescribed for all velocity components with the exception of a 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

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

For the solver, we use FGMRES [18] 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 [22] 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 pressure prolongation operator is given by bilinear interpolation

	$Re \equiv \nu^{-1}$			
$1/h$	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

and the pressure restriction is the scaled transpose of the prolongation. The velocity restriction operators are chosen to be a 1-D full weighting in the direction of the velocity component and averaging of nearest neighbors in the orthogonal direction, giving the following stencils for the restriction operators

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

These grid transfer operators are consistent with those suggested in [22] for staggered-grid discretizations.

As a point of comparison for the proposed smoothers to one which has been studied in literature, we introduce the method of Braess and Sarazin (B-S), originally proposed for the Stokes problem [8] but subsequently studied as an Oseen smoother [15]. In this method, a preconditioner of the form

$$\mathcal{P} = \begin{bmatrix} \alpha C & B^T \\ B & 0 \end{bmatrix} \tag{4.3}$$

is considered, where C is some approximation to A, the $(1,1)$ block of (1.8), and α is a damping parameter. In the original study, C was taken to be simply the diagonal of the respective matrix. For problems with convection, however, this approximation degrades very rapidly as $\nu \to 0$ and a better approximation is necessary. In [15], it was proposed to use a zero fill-in incomplete LU factorization of A and this is the selection we make for this study. It should be noted that this approximation, while more scalable in terms of iteration counts, significantly increases the cost of applying $\mathcal{P}^{-1}$. Indeed, one smoothing iteration requires solving two linear systems, one with matrix C and one with the Schur complement matrix $BC^{-1}B^T$. The first one is inexpensive: since $C = \bar{L}\bar{U}$, where $\bar{L}$ and $\bar{U}$ are the computed ILU factors of A, only two sparse triangular solves are required at each smoothing step. The second system, however, is problematic since the Schur complement matrix $B(\bar{L}\bar{U})^{-1}B^T$ cannot be explicitly formed. An inner iteration must be used to approximately solve the Schur complement systems, but convergence can be slow and it is difficult to find effective preconditioners, especially in the steady case. To present a fair comparison, in this study $\mathcal{P}^{-1}$ is applied using a direct solver, though clearly this would not be a feasible approach in a practical application.

Concerning the choice of α, we found that simply taking $\alpha = 1$ appears to produce optimal results, therefore this choice is used for all test cases in this study. The behavior of the method does not appear to be highly sensitive to the choice of α.

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,

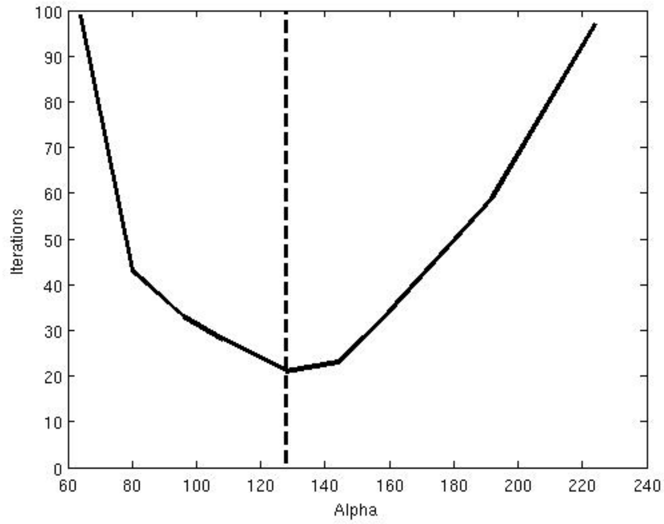

FIG. 4.1. *Sensitivity of HSS multigrid performance with respect to parameter α. Dashed line indicates location of optimal α predicted by LFA.*

	$Re \equiv \nu^{-1}$			
$1/h$	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

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 [10, 16], 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 coresponding to the finest mesh are meaningful.

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. Figure 4.1 illustrates the sensitivity of the convergence with repect to α and also displays the accuracy with which the LFA predicts the true optimal α.

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 $\hat{A}_\gamma$ 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

	$Re \equiv \nu^{-1}$			
$1/h$	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

	$Re \equiv \nu^{-1}$			
$1/h$	16	32	64	128
32	15	20	46	259
64	18	20	35	231
128	19	21	36	112
256	20	23	38	64

TABLE 4.4
Iteration count for B-S multigrid on steady Oseen problem

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.

Table 4.4 illustrates the behavior of the Braess–Sarazin smoother for the same steady 2-D lid-driven cavity problem. Within the range of problem parameters shown, the convergence rate appears to be independent of the mesh size, though the dependence on the viscosity is very pronounced. In particular, note the rather low Reynolds number at which the degradation begins to occur. The lack of robustness with respect to viscosity appears to demonstrate a clear advantage of either the HSS or AL smoothers when compared to the B-S smoothing for steady problems.

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. For the first test we take $\sigma = h^{-1}$, representative of the matrices that would result from a first order implicit discretization of the time derivative. The results for HSS, (inexact) AL multigrid ($\gamma = 1$), and B-S are shown in Tables 4.5, 4.6, and 4.7, respectively. For the HSS smoother, the viscosity dependence is almost completely eliminated, though there is some growth in iteration counts with respect to decreasing mesh size. The inexact AL multigrid

	$Re \equiv \nu^{-1}$				
$1/h$	256	512	1024	2048	4096
32	14	14	13	13	13
64	21	21	20	20	20
128	33	34	34	34	34
256	41	49	56	58	59

TABLE 4.5
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 4.6
Iteration count for inexact (block upper triangular/upper triangular) AL multigrid on unsteady Oseen problem, $\sigma = h^{-1}$

results of Table 4.6 are 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. The Braess–Sarazin smoother displays the poorest performance of the three. As mesh size and viscosity are both decreased, there is a noticeable degradation of performance.

Next we consider the case where h^{-2} multiplied by the identity matrix is added to the $(1, 1)$ block of the original coefficient matrix, simulating a second order discretization of the time derivative (e.g., with Crank-Nicolson). Tables 4.8, 4.9 and 4.10 show the results for the HSS, AL, and B-S smoothers, respectively. The HSS results are quite similar to those for $\sigma = h^{-1}$, exhibiting viscosity-independent convergence but still experiencing moderate growth with respect to decreasing mesh size. The inexact AL smoothers are extremely robust with respect to decreasing viscosity in this case and the iteration counts are quite small on the finest mesh. For this problem the Braess–Sarazin smoother displays the greatest insensitivity to problem parameters, showing convergence nearly independent of both mesh size and Reynolds number. It should be noted, however, that due to the ease of applying the upper triangular preconditioner, the inexact AL smoother actually requires much less computational effort.

Finally, we consider a case where the time step is given a fixed value across all mesh sizes (here we take $\sigma = 10$). This selection would be representative of problems with periodic or quasi-stationary flows. These results are displayed in tables 4.11, 4.12, and 4.13. The HSS smoother displays strong performance on the coarser grids, though the iteration counts increase rapidly for small viscosity on the finer meshes. The inexact AL approaches (as in the steady case, we select $\gamma = 0.1$) suffer from rather high iteration counts on the coarse meshes, though these counts are reduced

	$Re \equiv \nu^{-1}$			
$1/h$	512	1024	2048	4096
32	13	15	16	17
64	16	20	24	63
128	14	24	33	114
256	10	16	36	58

TABLE 4.7

Iteration count for B-S multigrid on unsteady Oseen problem ($\sigma = h^{-1}$)

	$Re \equiv \nu^{-1}$			
$1/h$	512	1024	2048	4096
32	35	34	33	32
64	42	41	40	39
128	43	42	40	39
256	50	48	46	44

TABLE 4.8

Iteration count for HSS multigrid on unsteady Oseen problem ($\sigma = h^{-2}$)

	$Re \equiv \nu^{-1}$			
$1/h$	512	1024	2048	4096
32	20/22	20/22	20/22	20/22
64	29/31	29/31	29/31	29/31
128	28/31	29/31	29/31	29/31
256	13/13	13/13	13/13	13/13

TABLE 4.9

Iteration count for inexact AL multigrid (block/lower triangular) on unsteady Oseen problem ($\sigma = h^{-2}$)

	$Re \equiv \nu^{-1}$			
$1/h$	512	1024	2048	4096
32	10	10	10	10
64	11	11	11	11
128	12	11	11	11
256	13	13	12	12

TABLE 4.10

Iteration count for B-S multigrid on unsteady Oseen problem ($\sigma = h^{-2}$)

	$Re \equiv \nu^{-1}$			
$1/h$	512	1024	2048	4096
32	8	8	8	8
64	9	11	12	10
128	10	19	57	66
256	12	12	91	132

TABLE 4.11

Iteration count for HSS multigrid on unsteady Oseen problem ($\sigma = 10$)

	$Re \equiv \nu^{-1}$			
$1/h$	512	1024	2048	4096
32	38/110	56/162	68/204	74/239
64	35/137	65/276	113/476	184/NC
128	20/77	40/186	74/486	159/NC
256	12/36	18/74	29/158	46/325

TABLE 4.12

Iteration count for inexact AL multigrid (block/lower triangular) on unsteady Oseen problem ($\sigma = 10$). NC indicates that convergence was not achieved within 500 iterations.

	$Re \equiv \nu^{-1}$			
$1/h$	32	64	128	256
32	26	55	305	NC
64	24	71	NC	NC
128	28	43	490	NC
256	41	42	104	NC

TABLE 4.13

Iteration count for B-S multigrid on unsteady Oseen problem ($\sigma = 10$). NC indicates that convergence was not achieved within 500 iterations

to more manageable levels on the finest mesh. As expected, the viscosity scaling is more favorable than in the steady case and in fact even the triangular approximation (which was unacceptable for steady flow) shows potential, considering the low cost per iteration. The behavior of the Braess–Sarazin smoother is again quite similar to the steady case, showing rapid deterioration in performance for even rather large values of the viscosity.

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. Both approaches appear to be superior to a coupled smoothing iteration of Braess–Sarazin type from the point of view of robustness and overall computational cost.

Future work includes an 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. An investigation into the relationship between the augmented Lagrangian formulation and the poroelasticity equations of [17] 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] D. Braess and R. Sarazin, *An efficient smoother for the Stokes problem*, Appl. Numer. Math., 23 (1997), pp. 3–19.

[9] H. C. Elman, *Preconditioners for saddle point problems arising in computational fluid dynamics*, Appl. Numer. Math., 43 (2002), pp. 75–89.

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

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] V. John and L. Tobiska, *Numerical performance of smoothers in coupled multigrid methods for parallel solution of the incompressible Navier–Stokes equations*, Int. J. Numer. Meth. Fluids, 33 (2000), pp. 453–473.

[16] 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.

[17] C. W. Oosterlee and F. J. Gaspar, *Multigrid relaxation methods for systems of saddle point type*, Appl. Numer. Math., 58 (2008), pp. 1933–1950.

[18] Y. Saad, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM J. Sci. Comput., 14 (1993), pp. 461–469.

[19] 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.

[20] J. Schöberl, *Multigrid methods for a parameter dependent problem in primal variables*, Numer. Math., 84 (1999), pp. 97–119.

[21] 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.

[22] U. Trottenberg, C. Oosterlee, and A. Schuller, *Multigrid*, Academic Press, San Diego, 2001.

[23] S. P. Vanka, *Block-implicit multigrid solution of Navier–Stokes equations in primitive variables*, J. Comput. Phys., 65 (1986), pp. 138–158.

[24] P. Wesseling, *Principles of Computational Fluid Dynamics*, Springer Series in Computational Mathematics 29, Springer, New York, 2001.

[25] P. Wesseling and C. W. Oosterlee, *Geometric multigrid with applications to computational fluid dynamics*, J. Comput. Appl. Math., 128 (2001), pp. 311–334.

[26] R. Wienands and W. Joppich, *Practical Fourier Analysis for Multigrid Methods*, Chapman & Hall, New York, 2005.