NUMERICAL LINEAR ALGEBRA WITH APPLICATIONS
Numer. Linear Algebra Appl. 2010; **17**:557–576
Published online 9 February 2010 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/nla.707

New multigrid smoothers for the Oseen problem

Steven Hamilton*, †, Michele Benzi and Eldad Haber

Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, U.S.A.

SUMMARY

We investigate the performance of smoothers based on the Hermitian/skew-Hermitian (HSS) and augmented Lagrangian (AL) splittings applied to the Marker-and-Cell (MAC) discretization 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. Copyright © 2010 John Wiley & Sons, Ltd.

Received 27 May 2009; Revised 14 December 2009; Accepted 14 December 2009

KEY WORDS: multigrid; smoothing iterations; generalized Stokes and Oseen problems; incompressible Navier–Stokes equations

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

$$\frac{\partial \mathbf{u}}{\partial t} - \nu \Delta \mathbf{u} + (\mathbf{u} \cdot \nabla)\mathbf{u} + \nabla p = \mathbf{f} \quad \text{on } \Omega \times (0, T] \tag{1}$$

$$\operatorname{div} \mathbf{u} = 0 \quad \text{on } \Omega \times [0, T] \tag{2}$$

*Correspondence to: Steven Hamilton, Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, U.S.A.
†E-mail: sphamil@emory.edu

Contract/grant sponsor: U.S. Department of Energy Computational Science Graduate Fellowship
Contract/grant sponsor: National Science Foundation; contract/grant number: DMS-0511336
Contract/grant sponsor: National Science Foundation; contract/grant number: DMS-0724759
Contract/grant sponsor: U.S. Department of Energy; contract/grant number: DMS 0724759

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega \times [0, T] \tag{3}$$

$$\mathbf{u}(\mathbf{x}, 0) = \mathbf{u}_0(\mathbf{x}) \quad \text{on } \Omega \tag{4}$$

where v 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

$$\sigma\mathbf{u} - v\Delta\mathbf{u} + (\mathbf{v}\cdot\nabla)\mathbf{u} + \nabla p = \mathbf{f} \quad \text{in } \Omega \tag{5}$$

$$\operatorname{div}\mathbf{u} = 0 \quad \text{in } \Omega \tag{6}$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega \tag{7}$$

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

$$\begin{bmatrix} A & B^{\mathrm{T}} \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} f \\ g \end{bmatrix} \tag{8}$$

where u represents the discrete velocity, p now represents the discrete pressure, 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. [1].

The efficient solution of (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. [1–7]. Coupled multigrid methods have also been developed; see, e.g. [8–10] 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 v (equivalently, of the Reynolds number Re). As mentioned in [10], 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 v; see also [11] 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 such as Vanka's method [12] and segregated, distributive relaxation schemes such as SIMPLE and related approaches (see, e.g. [9, 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. [3, 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 (8). Such splittings have been intensively studied in recent years in the context of developing preconditioners for Krylov subspace methods,

DOI: 10.1002/nla

with very good results; see, e.g. [2, 3, 13–15]. Their use as smoothers for multigrid has not been previously investigated.

2. HSS SMOOTHING

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

$$\begin{aligned}\mathscr{A} &= (H+\alpha I)-(\alpha I-S)\\ \mathscr{A} &= (S+\alpha I)-(\alpha I-H).\end{aligned} \tag{9}$$

From these splittings, the HSS iteration for solving $\mathscr{A}x=b$ was defined in [13] by:

$$\begin{aligned}(H+\alpha I)x^{(k+\frac{1}{2})} &= (\alpha I-S)x^{(k)}+b\\ (S+\alpha I)x^{(k+1)} &= (\alpha I-H)x^{\left(k+\frac{1}{2}\right)}+b\end{aligned} \tag{10}$$

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 [2, 14, 16]. The optimal selection of the parameter α (as a preconditioner) has been studied in, e.g. [15, 16]. In practice, when carrying out (10) 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 (8) in which a negative sign is placed before the (2,1) block, resulting in the equivalent system

$$\begin{bmatrix} A & B^{\mathrm{T}} \\ -B & 0 \end{bmatrix}\begin{bmatrix} u \\ p \end{bmatrix}=\begin{bmatrix} f \\ -g \end{bmatrix} \quad \text{or } \mathscr{A}x=b. \tag{11}$$

Although the coefficient matrix in (8) is indefinite (its eigenvalues fall on both sides of the imaginary axis), the one in (11) has its spectrum entirely contained in the open-right complex half-plane [2].

For many discretizations (e.g. the Marker-and-Cell (MAC), discretization [17]), the HSS splitting almost exactly corresponds to the natural (i.e. physical) splitting of the relevant differential operators:

$$H=\begin{bmatrix} \sigma I+L & 0 \\ 0 & 0 \end{bmatrix} \tag{12}$$

$$S=\begin{bmatrix} K & B^{\mathrm{T}} \\ -B & 0 \end{bmatrix} \tag{13}$$

DOI: 10.1002/nla

where L represents the discretization of the negative vector Laplacian multiplied by the viscosity, given by

$$L = \begin{bmatrix} -\nu\Delta_h & 0 \\ 0 & -\nu\Delta_h \end{bmatrix}, \tag{14}$$

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 given by

$$K = \begin{bmatrix} (\mathbf{v}\cdot\nabla)_h & 0 \\ 0 & (\mathbf{v}\cdot\nabla)_h \end{bmatrix} \tag{15}$$

where $\mathbf{v} = (v_x, v_y)^{\mathrm{T}}$ is the vector containing the discretized wind function. B and B^{T} retain their previous meaning. The 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 Equation (12). Though the theoretical analyses on the HSS approach only strictly hold for the algebraic splitting, experimental observations indicate that there is a 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) [8, 18] on the HSS iteration for the steady Oseen problem with the MAC discretization to determine its potential as a multigrid smoother and aid in the selection of the free parameter α. The MAC discretization is a staggered grid finite-difference discretization consisting of a standard five-point discretization of the Laplacian, a centered short-difference discretization of the gradient/divergence terms, and a centered long-difference treatment of the convective term. The iteration matrix describing Equations (10) 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 scalar Laplacian by $\tilde{\Delta}_h$, the symbol of the convective term by $\widetilde{(\mathbf{v}\cdot\nabla)}_h$, and the symbols of the x- and y- derivatives by $\tilde{\partial}_x$ and $\tilde{\partial}_y$, respectively. For the MAC discretization, these symbols equate to $\tilde{\Delta}_h = -(1/h^2)(4 - 2\cos\theta_x - 2\cos\theta_y)$, $\widetilde{(\mathbf{v}\cdot\nabla)}_h = (\mathrm{i}/h)\,(v_x \sin\theta_x + v_y \sin\theta_y)$, and $\tilde{\partial}_{x,y} = (2\mathrm{i}/h)\,\sin(\theta_{x,y}/2)$, where $\mathrm{i} = \sqrt{-1}$. Using these symbols, the relevant matrices can be written as

$$\tilde{H} = \begin{bmatrix} -\nu\tilde{\Delta}_h & 0 & 0 \\ 0 & -\nu\tilde{\Delta}_h & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \tilde{S} = \begin{bmatrix} \widetilde{(\mathbf{v}\cdot\nabla)}_h & 0 & -\tilde{\partial}_x \\ 0 & \widetilde{(\mathbf{v}\cdot\nabla)}_h & -\tilde{\partial}_y \\ -\tilde{\partial}_x & -\tilde{\partial}_y & 0 \end{bmatrix}, \tag{16}$$

where the negative signs appear in the $(1,3)$ and $(2,3)$ positions of $\tilde{S}$ because the adjoints of $\tilde{\partial}_x$ and $\tilde{\partial}_y$ are $-\tilde{\partial}_x$ and $-\tilde{\partial}_y$, respectively. We use the MATLAB [19] symbolic algebra package

DOI: 10.1002/nla

to compute the symbol of the iteration matrix $\tilde{T}=(\tilde{S}+\alpha I)^{-1}(\alpha I-\tilde{H})(\tilde{H}+\alpha I)^{-1}(\alpha I-\tilde{S})$ and its eigenvalues. These eigenvalues are given by

$$\lambda_1=\left(\frac{\alpha+\nu\tilde{\Delta}_h}{\alpha-\nu\tilde{\Delta}_h}\right)\left(\frac{\alpha-(\widetilde{\mathbf{v}\cdot\nabla})_h}{\alpha+(\widetilde{\mathbf{v}\cdot\nabla})_h}\right) \tag{17}$$

$$\lambda_{2,3}=\frac{[\alpha^3+\alpha k-\alpha\nu\tilde{\Delta}_h(\widetilde{\mathbf{v}\cdot\nabla})_h]\pm([\alpha^2(\widetilde{\mathbf{v}\cdot\nabla})_h-\nu\tilde{\Delta}_h(\alpha^2-k)]^2+4\alpha^4k)^{1/2}}{[\alpha^3-\alpha k-\alpha\nu\tilde{\Delta}_h(\widetilde{\mathbf{v}\cdot\nabla})_h]+[\alpha^2(\widetilde{\mathbf{v}\cdot\nabla})_h-\nu\tilde{\Delta}_h(\alpha^2-k)]} \tag{18}$$

where $k=(\tilde{\partial}_x^2+\tilde{\partial}_y^2)$. Because α is taken to be a real number and $(\widetilde{\mathbf{v}\cdot\nabla})_h$ is imaginary, the second term in Equation (17) has modulus 1, thus

$$|\lambda_1|=\left|\frac{\alpha+\nu\tilde{\Delta}_h}{\alpha-\nu\tilde{\Delta}_h}\right|=\left|\frac{\alpha-\dfrac{\nu}{h^2}(4-2\cos\theta_x-2\cos\theta_y)}{\alpha+\dfrac{\nu}{h^2}(4-2\cos\theta_x-2\cos\theta_y)}\right| \tag{19}$$

independent of the convective term. By selecting $\alpha=4\nu/h^2$ in the steady case, this expression becomes

$$|\lambda_1|=\left|\frac{\cos\theta_x+\cos\theta_y}{4-\cos\theta_x-\cos\theta_y}\right| \tag{20}$$

which is independent of both the mesh size and the viscosity. In fact, over the high-frequency portion of the frequency domain ($\mathbf{\Theta}_{\text{high}}=[-\pi,\pi]\backslash(-\pi/2,\pi/2)$) this quantity is bounded above by $\frac{1}{3}$. The LFA frequency-dependent amplification factor is defined as $S_h(\boldsymbol{\theta})=\rho(\tilde{T}(\boldsymbol{\theta}))$, where ρ indicates the spectral radius and $\boldsymbol{\theta}=(\theta_x,\theta_y)^{\mathrm{T}}$. Because in a multigrid method it is assumed that low-frequency errors will be handled on coarser meshes, the rate of smoothing is dictated by the worst-error amplification that occurs within the high-frequency portion of the domain. Thus, we define an LFA smoothing factor as $\mu_{\text{loc}}=\sup\{S_h(\boldsymbol{\theta})|\boldsymbol{\theta}\in\mathbf{\Theta}_{\text{high}}\}$. Figure 1 shows the behavior of the amplification factor for two different instances of the Oseen problem and identifies the corresponding smoothing factors. Here the wind function is taken to have norm 1 by selecting $v_x=\cos\beta$ and $v_y=\sin\beta$ for a parameter β. In this example we choose $\beta=\pi/4$, though the results are highly insensitive to this selection.

Figure 2 shows the progression of the velocity error for a sample problem consisting of flow in the unit square with homogeneous Dirichlet boundary conditions (thus, the exact solution is the all zero vector) and a wind function given by

$$\mathbf{v}(x,y)=\begin{bmatrix}8x(x-1)(1-2y)\\ 8(2x-1)y(y-1)\end{bmatrix} \tag{21}$$

The initial guess is taken to be a random vector with magnitude 1. Figure 3 shows the error in the velocity for the same test problem after 20 iterations for a range of α values. This illustrates the favorable smoothing properties of the HSS iteration when α is selected to be near the optimal value predicted by the LFA.

DOI: 10.1002/nla

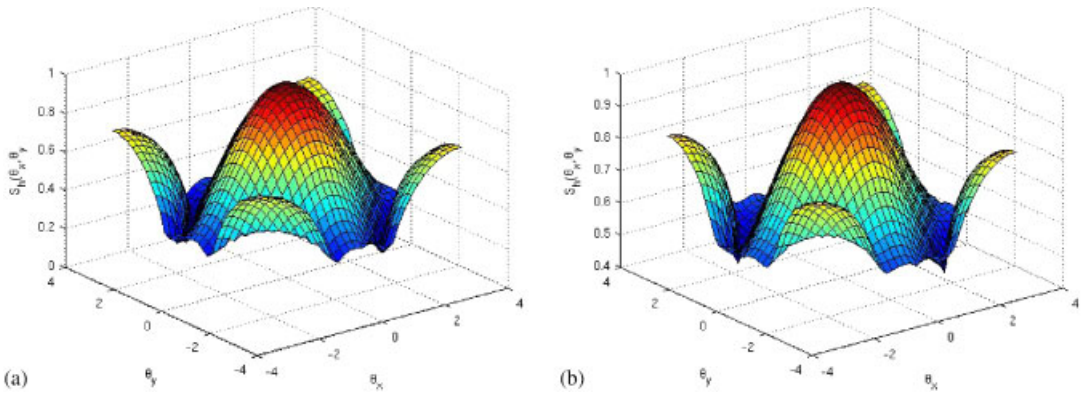

Figure 1. LFA frequency dependent amplification factor with HSS smoother for different parameter combinations: (a) $h=\frac{1}{32}$, $v=0.01$, $\alpha=50.2$, $\mu_{\text{loc}}=0.731$ and (b) $h=\frac{1}{128}$, $v=0.001$, $\alpha=121$, $\mu_{\text{loc}}=0.817$.

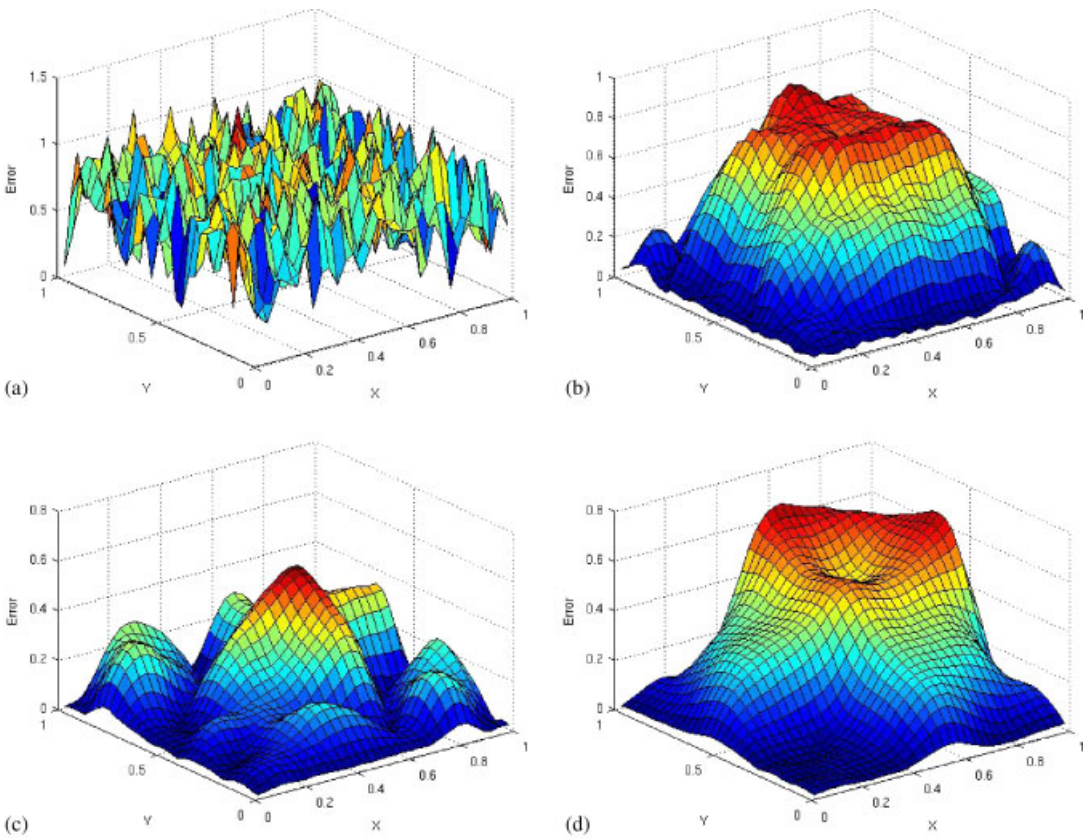

Figure 2. Velocity error versus iteration number ($h=\frac{1}{32}$, $v=0.01$, $\alpha=4v/h^2$): (a) initial error; (b) after 5 iterations; (c) after 10 iterations; and (d) after 20 iterations.

DOI: 10.1002/nla

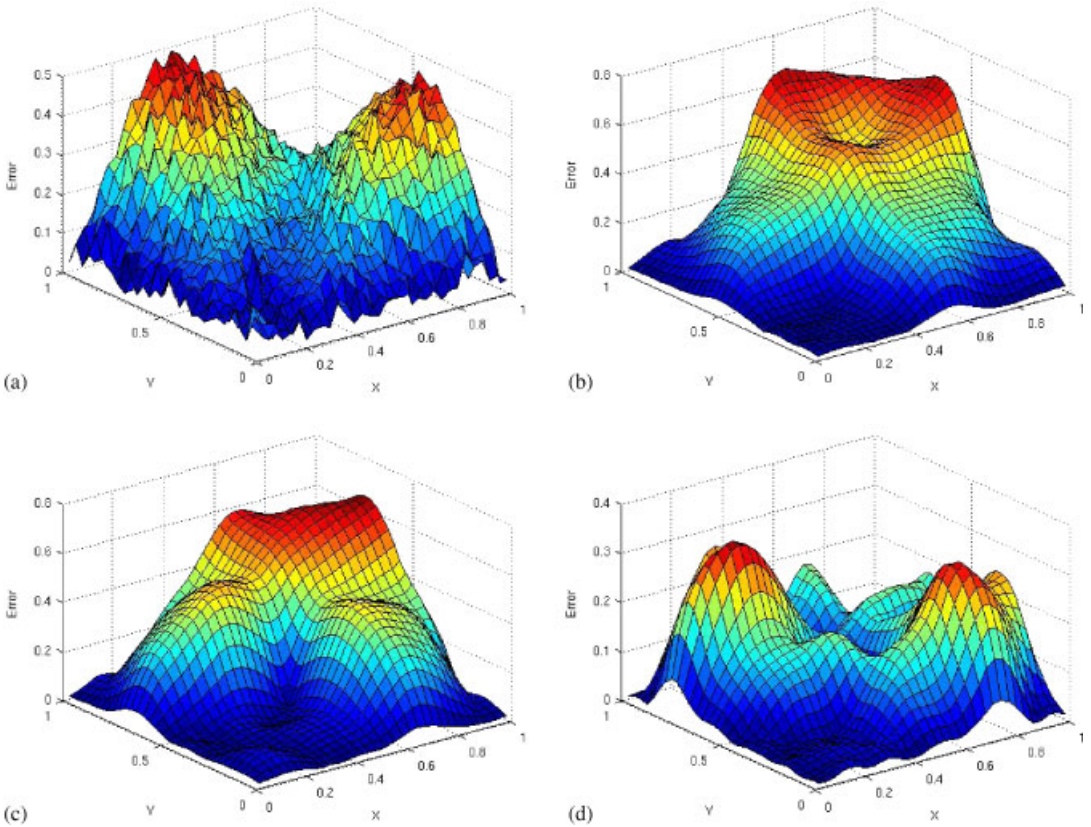

Figure 3. Velocity error after 20 iterations for various α values ($h=\frac{1}{32}$, $\nu=0.01$): (a) $\alpha=20$; (b) $\alpha=\alpha_{\mathrm{LFA}}=50$; (c) $\alpha=60$; and (d) $\alpha=80$.

The preceding analysis can be easily generalized to the unsteady case (i.e. $\sigma\neq 0$ in Equation (12)). The result is that $-\nu\tilde{\Delta}_h$ is replaced by $-\nu\tilde{\Delta}_h+\sigma$ in Equations (17)–(19) and selecting $\alpha=(4\nu/h^2)+\sigma$ results in

$$|\lambda_1|=\left|\frac{\cos\theta_x+\cos\theta_y}{4-\cos\theta_x-\cos\theta_y+\sigma\dfrac{h^2}{\nu}}\right| \tag{22}$$

which is always smaller than $|\lambda_1|$ for the steady case. Table I shows the computed LFA smoothing factor, μ_{loc}, and corresponding optimal α values for several different parameter combinations for both the steady and unsteady problems. The smoothing factor for the unsteady problem is almost always less than that for the steady problem, although in the second to last line of the table we see it is possible for the unsteady case to be worse.

In practice, we have observed that using HSS as a multigrid smoother leads to excellent convergence behavior for certain problems but results in a large sensitivity with respect to problem

DOI: 10.1002/nla

Table I. LFA smoothing factor, μ_{loc} for the HSS iteration applied to the steady and unsteady ($\sigma = 1/h$) Oseen problems. Numbers in parentheses are the optimal α value.

		$Re \equiv \nu^{-1}$	
$1/h$	ν	Steady	Unsteady
32	0.01	0.731 (50.2)	0.562 (74)
32	0.001	0.904 (14.7)	0.544 (44.7)
64	0.01	0.658 (153)	0.604 (193)
64	0.001	0.866 (42)	0.547 (97.4)
128	0.01	0.753 (378)	0.773 (375)
128	0.001	0.817 (121)	0.554 (223)

parameters, namely a significant degradation in the convergence rate as the viscosity is decreased. However, by using the multigrid V-cycle as a preconditioner for GMRES [20] this sensitivity is greatly reduced. In [21] it was observed that utilizing a multigrid method as a preconditioner for a Krylov method typically yielded more favorable results than multigrid used as an iterative solver. Additionally, a technique for performing a Fourier analysis of GMRES(m) preconditioned by multigrid was developed in [22]. Experimentally, it has been observed that the optimal value of α when using HSS multigrid as a preconditioner for GMRES differs from that predicted by LFA. In particular, the selection $\alpha = 4\nu/h^2$ appears to produce optimal (or very nearly so) behavior for the steady case across a range of problem parameters, as demonstrated in Figure 4 for the same test problem as before. Rather than minimizing the eigenvalue of maximum modulus in Equations (17) and (18), this choice of α actually minimizes $|\lambda_1|$ (λ_1 is always the eigenvalue of minimum modulus) as discussed previously. For the unsteady case, the optimal selection appears to be $\alpha = 4\nu/h^2 + \sigma$, which is again the selection that minimizes $|\lambda_1|$.

In order for HSS to become a feasible method it is necessary to apply an inexact variant with a small computational cost that 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 the mesh sizes and viscosities when α is taken to be the value suggested previously for use with GMRES, $\alpha = 4\nu/h^2$. 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 [20] 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 [23]. With such reordering, the same value $\tau = 0.01$ of the ILU drop tolerance was used in all cases. Table II 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 at worst linear in the number of unknowns. We note that when the inner iterations are performed inexactly using a non-stationary method, it becomes necessary to use a flexible variant of GMRES (FGMRES) [24] for the outer solver.

DOI: 10.1002/nla

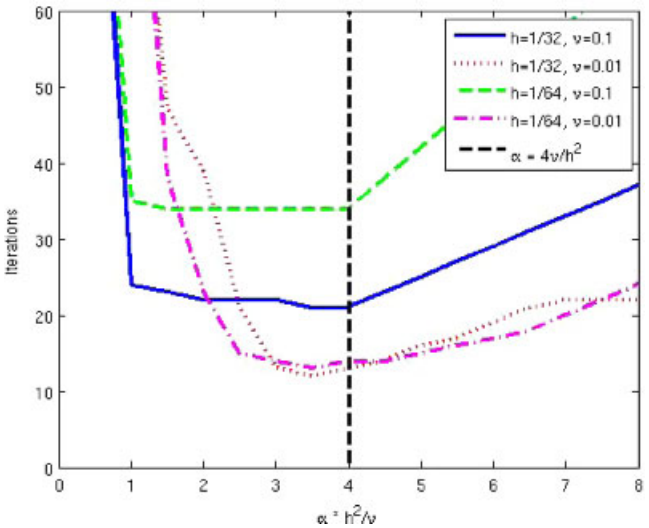

Figure 4. Iterations required for convergence of GMRES preconditioned by an HSS multigrid V-cycle for different parameter combinations. The vertical line indicates $\alpha = 4\nu/h^2$.

Table II. Ratio of non-zeros in ILU factors over non-zeros in $S+\alpha I$.

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

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, 15, 16], 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 $h \to 0$.

3. AL SMOOTHING

We begin the discussion of the AL formulation by replacing the original system (8) with the following equivalent system:

$$\begin{bmatrix} A+\gamma B^{\mathrm{T}} W^{-1} B & B^{\mathrm{T}} \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} \hat{f} \\ g \end{bmatrix} \tag{23}$$

DOI: 10.1002/nla

where $\hat{f} = f + \gamma B^{\mathrm{T}} W^{-1} g$; 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 [25]. It is of interest to note that the (1,1) block of the augmented system (23) resembles that of the poroelasticity equations described in [11] (though the poroelasticity equations lack a convective term).

We let $\mathscr{A}$ now denote the coefficient matrix in the preceding equation. Now we consider a preconditioner of the form

$$\mathscr{P} = \begin{bmatrix} \hat{A}_\gamma & B^{\mathrm{T}} \\ 0 & -\dfrac{1}{\gamma} W \end{bmatrix} \tag{24}$$

where the application of $\hat{A}_\gamma^{-1}$ involves the (inexact) inversion of $A + \gamma B^{\mathrm{T}} W^{-1} B$. It was shown in [3] that the eigenvalues of $\mathscr{P}^{-1}\mathscr{A}$ all tend to 1 as $\gamma \to \infty$ (uniformly in h). However, as B (and thus $B^{\mathrm{T}} W^{-1} B$) has a significant null space, $A + \gamma B^{\mathrm{T}} W^{-1} B$ 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 [3].

Defining $T = I - \mathscr{P}^{-1}\mathscr{A}$ (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 LFA for this iteration matrix again for the MAC-discretized steady problem. Observe that the preconditioned matrix can be written as

$$\mathscr{P}^{-1}\mathscr{A} = \begin{bmatrix} I + \gamma A_\gamma^{-1} B^{\mathrm{T}} B & A_\gamma^{-1} B^{\mathrm{T}} \\ -\gamma B & 0 \end{bmatrix} \tag{25}$$

Thus, the iteration matrix is given by

$$T = I - \mathscr{P}^{-1}\mathscr{A} = \begin{bmatrix} \gamma A_\gamma^{-1} B^{\mathrm{T}} B & -A_\gamma^{-1} B^{\mathrm{T}} \\ \gamma B & I \end{bmatrix} \tag{26}$$

As $A_\gamma = A + \gamma B^{\mathrm{T}} B$ (here we take $W = I$ as the pressure mass matrix is the identity for finite differences), we can write

$$\tilde{A}_\gamma = \begin{bmatrix} -\nu\tilde{\Delta}_h + (\widetilde{\mathbf{v}\cdot\nabla})_h + \gamma\tilde{\partial}_x^2 & \gamma\tilde{\partial}_x\tilde{\partial}_y \\ \gamma\tilde{\partial}_x\tilde{\partial}_y & -\nu\tilde{\Delta}_h + (\widetilde{\mathbf{v}\cdot\nabla})_h + \gamma\tilde{\partial}_y^2 \end{bmatrix} \tag{27}$$

and therefore

$$\tilde{A}_\gamma^{-1} = \frac{1}{c}\begin{bmatrix} -\nu\tilde{\Delta}_h + (\widetilde{\mathbf{v}\cdot\nabla})_h + \gamma\tilde{\partial}_y^2 & -\gamma\tilde{\partial}_x\tilde{\partial}_y \\ -\gamma\tilde{\partial}_x\tilde{\partial}_y & -\nu\tilde{\Delta}_h + (\widetilde{\mathbf{v}\cdot\nabla})_h + \gamma\tilde{\partial}_x^2 \end{bmatrix} \tag{28}$$

DOI: 10.1002/nla

where $c=(-\nu\tilde{\Delta}_h+(\widetilde{\mathbf{v}\cdot\nabla})_h)[-(\nu+\gamma)\tilde{\Delta}_h+(\widetilde{\mathbf{v}\cdot\nabla})_h]$. Multiplying out these terms allows us to write the symbol of the iteration matrix as

$$\tilde{T}=\begin{bmatrix} \frac{\gamma}{c_2}\tilde{\partial}_x^2 & \frac{\gamma}{c_2}\tilde{\partial}_x\tilde{\partial}_y & \frac{1}{c_2}\tilde{\partial}_x \\ \frac{\gamma}{c_2}\tilde{\partial}_x\tilde{\partial}_y & \frac{\gamma}{c_2}\tilde{\partial}_y^2 & \frac{1}{c_2}\tilde{\partial}_y \\ \gamma\tilde{\partial}_x & \gamma\tilde{\partial}_y & 1 \end{bmatrix} \tag{29}$$

where $c_2=-(\nu+\gamma)\tilde{\Delta}_h+(\widetilde{\mathbf{v}\cdot\nabla})_h$. Using the MATLAB symbolic algebra package to find the eigenvalues of this matrix, we find that two eigenvalues are identically zero and the third is given by

$$\lambda=\frac{-\nu\tilde{\Delta}_h+(\widetilde{\mathbf{v}\cdot\nabla})_h}{-(\nu+\gamma)\tilde{\Delta}_h+(\widetilde{\mathbf{v}\cdot\nabla})_h} \tag{30}$$

The frequency-dependent amplification factor is simply the modulus of this eigenvalue (recall that $\tilde{\Delta}_h$ is a real number and $(\widetilde{v\cdot\nabla})_h$ is imaginary), given by

$$\mu_{\text{loc}}=|\lambda|=\left(\frac{\nu^2\tilde{\Delta}_h^2+|(\widetilde{\mathbf{v}\cdot\nabla})_h|^2}{(\nu+\gamma)^2\tilde{\Delta}_h^2+|(\widetilde{\mathbf{v}\cdot\nabla})_h|^2}\right)^{1/2} \tag{31}$$

This expression allows us to identify a few pertinent features of the amplification factor. First of all, the amplification factor is strictly less than unity for any positive values of ν and γ, indicating that this iteration does in fact possess smoothing properties. Second, the amplification factor tends to zero as γ tends to infinity. As mentioned earlier, large values of γ present difficulties when attempting to invert $\mathscr{P}$ inexactly and thus we anticipate that selecting a moderate value of γ will likely be necessary in a practical implementation. A plot of the amplification factor as a function of Fourier frequency is shown in Figure 5 for two different parameter combinations on the same problem as described for Figure 1.

An expression for the smoothing factor for the unsteady Oseen problem can again be obtained by replacing $-\nu\tilde{\Delta}_h$ by $-\nu\tilde{\Delta}_h+\sigma$ in Equations (27)–(30). The resulting non-zero eigenvalue becomes

$$\mu_{\text{loc}}=\left(\frac{(\sigma-\nu\tilde{\Delta}_h)^2+|(\widetilde{\mathbf{v}\cdot\nabla})_h|^2}{[\sigma-(\nu+\gamma)\tilde{\Delta}_h]^2+|(\widetilde{\mathbf{v}\cdot\nabla})_h|^2}\right)^{1/2} \tag{32}$$

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 $B^{\mathrm{T}}W^{-1}B$). In [3], a highly effective and robust geometric multigrid solver was developed to address such difficulties, based on the one presented in [26]. Unfortunately, implementation of this multigrid iteration is less than straightforward, particularly on unstructured meshes. As a simpler alternative, we consider taking $\hat{A}_\gamma$ to be the block upper triangular portion of $A+\gamma B^{\mathrm{T}}W^{-1}B$. Note that A, and hence

DOI: 10.1002/nla

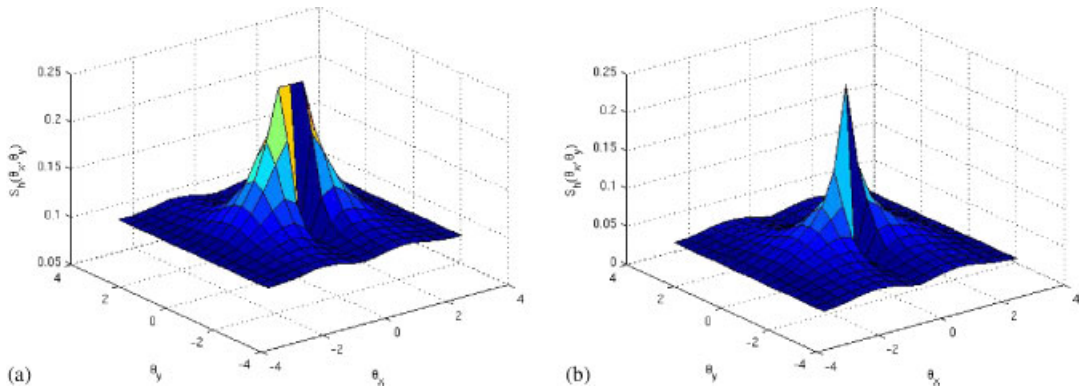

Figure 5. LFA frequency-dependent amplification factor with AL smoother for different parameter combinations. Note that the maximum value shown on z-axis is 0.25: (a) $h=\frac{1}{32}$, $\nu=0.01$, $\gamma=1$, $\mu_{\text{loc}}=0.136$ and (b) $h=\frac{1}{128}$, $\nu=0.001$, $\gamma=1$, $\mu_{\text{loc}}=0.064$.

$A+\gamma B^{\mathrm{T}}W^{-1}B$, has a natural two-by-two block structure so that the upper triangular portion of A_γ can be written as

$$\hat{A}_\gamma=\begin{bmatrix} -\nu\tilde{\Delta}_h+(\widetilde{\mathbf{v}\cdot\nabla})_h+\gamma\tilde{\partial}_x^2 & \gamma\tilde{\partial}_x\tilde{\partial}_y \\ 0 & -\nu\tilde{\Delta}_h+(\widetilde{\mathbf{v}\cdot\nabla})_h+\gamma\tilde{\partial}_y^2 \end{bmatrix}. \tag{33}$$

The inversion of this block upper triangular matrix requires the (approximate) solution of two scalar anisotropic convection–diffusion equations with anisotropy ratio $1+(\gamma/\nu)$. For a detailed discussion of the merits of this modified AL approach as a preconditioner, see [27]. In the following experiments, we solved the convection–diffusion equations 'exactly' using 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 [28]. To simplify matters even further we additionally consider taking $\hat{A}_\gamma$ to be simply the upper triangular portion of A_γ. If W is taken to be a diagonal matrix, then the preconditioner $\mathscr{P}$ as a whole becomes 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. As in the case of HSS, we elect to use the AL multigrid as a preconditioner to GMRES rather than as a standalone solver. Below we investigate the use of $\mathscr{P}$ or one of these approximations as a smoother for a coupled multigrid method for the discrete Oseen problem.

4. RESULTS

We consider the MAC discretized Oseen problem on the unit square. As a test problem, we take the standard leaky lid-driven cavity problem described, for instance, in [1]. Homogeneous Dirichlet boundary conditions are prescribed for all velocity components with the exception of a positive

DOI: 10.1002/nla

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 Equation (21).

For the solver, we use FGMRES [24] preconditioned with one multigrid cycle. In all cases we used the zero vector as the initial guess and a reduction of the two-norm of the initial residual by four orders of magnitude as the stopping criterion. The cycle is chosen to be a V-cycle [8] with one pre-smoothing and one post-smoothing step. For HSS, a smoothing step is simply one full HSS iteration. For the AL approach, a smoothing step is a single Richardson iteration on the preconditioned system, i.e. $x^{k+1}=x^k+\mathscr{P}^{-1}r^k$ where $r^k=b-\mathscr{A}x^k$ is the residual and $\mathscr{P}$ 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 and pressure restriction is achieved via the four-point averaging operator. 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:

$$R_{u_x}=\frac{1}{8}\begin{bmatrix}1 & 2 & 1\\ & \cdot & \\ 1 & 2 & 1\end{bmatrix}\quad R_{u_y}=\frac{1}{8}\begin{bmatrix}1 & & 1\\ 2 & \cdot & 2\\ 1 & & 1\end{bmatrix}\quad R_p=\frac{1}{4}\begin{bmatrix}1 & & 1\\ & \cdot & \\ 1 & & 1\end{bmatrix} \tag{34}$$

Bilinear interpolation is used for velocity prolongation. These grid transfer operators are consistent with those suggested in [8] 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 [29] but subsequently studied as an Oseen smoother [30]. In this method, a preconditioner of the form

$$\mathscr{P}=\begin{bmatrix}\alpha C & B^{\mathrm{T}}\\ B & 0\end{bmatrix} \tag{35}$$

is considered, where C is some approximation to A, the $(1,1)$ block of (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 [30], 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 $\mathscr{P}^{-1}$. Indeed, one smoothing iteration requires solving three linear systems, two with the matrix C and one with the Schur complement matrix $BC^{-1}B^{\mathrm{T}}$. The first two are inexpensive: as $C=\bar{L}\bar{U}$, where $\bar{L}$ and $\bar{U}$ are the computed ILU factors of A, only four sparse triangular solves are required at each smoothing step. The second system, however, is problematic as the Schur complement matrix $B(\bar{L}\bar{U})^{-1}B^{\mathrm{T}}$ 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 $\mathscr{P}^{-1}$ is applied using a direct solver, though clearly this would not be a recommended approach in a practical application.

DOI: 10.1002/nla

Table III. Iteration count for HSS multigrid on steady Oseen problem.

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

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 III 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 no worse than that experienced by other solvers in literature [5, 7], and at the finest mesh size the degradation is actually quite manageable. It should be mentioned that as no velocity stabilization is being attempted, for small ν only numerical solutions corresponding to the finest mesh are meaningful.

For all parameter combinations, α is selected to be the value described in Section 2, namely $\alpha = 4\nu/h^2$. As this α depends on h, it is necessary to redefine α on each grid level using the value of h corresponding to the mesh size on that level.

The performance of the AL smoother on the same problem is shown in Table IV for the exact case and in Table V for the inexact case. 'Inexact' here refers to the selection of $\hat{A}_\gamma$ as the block upper triangular portion of $A + \gamma B^{\mathrm{T}} W^{-1} B$ as discussed in Section 3. 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 AL 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 IV and V one can see that the iteration count is essentially independent of the ratio ν/h, a highly desirable property.

Table VI illustrates the behavior of the B–S 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.

DOI: 10.1002/nla

Table IV. Iteration count for exact AL multigrid on steady Oseen problem.

	$Re \equiv \nu^{-1}$			
$1/h$	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 V. Iteration count for inexact (block upper triangular) 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 VI. Iteration count for B–S 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

The lack of robustness with respect to viscosity appears to demonstrate a clear advantage of either the HSS or AL smoothers when compared with 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 ($\alpha = (4\nu/h^2) + \sigma$), (inexact) AL multigrid ($\gamma = 1$), and B–S are shown in Tables VII–IX, 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 results of Table VIII are remarkable: the convergence shows no degradation with respect to decreases in either the mesh size or the viscosity—even for viscosities that 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 B–S smoother displays the most

DOI: 10.1002/nla

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

	$Re \equiv \nu^{-1}$				
$1/h$	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 VIII. Iteration count for inexact (block upper triangular/upper triangular) AL 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 IX. Iteration count for B–S multigrid on unsteady Oseen problem, $\sigma = h^{-1}$.

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

sensitivity with respect to decreasing viscosity, although on the finest mesh the iteration counts are comparable to HSS and B–S may be slightly more robust with respect to mesh refinement. However, due to the previously mentioned difficulties in applying the B–S smoother, significantly more computational effort is required.

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–Nicholson). Tables X–XII show the results for the HSS, AL, and B-S smoothers, respectively. The HSS results are slightly better than for $\sigma = h^{-1}$, still exhibiting viscosity-independent convergence but now showing less degradation with 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 B–S smoother displays the greatest insensitivity to problem parameters, showing convergence nearly independent

DOI: 10.1002/nla

Table X. Iteration count for HSS multigrid on unsteady Oseen problem, $\sigma = h^{-2}$.

	$Re \equiv \nu^{-1}$			
$1/h$	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 XI. Iteration count for inexact AL multigrid (block/upper triangular) 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 XII. Iteration count for B–S multigrid 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

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 XIII–XV. 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 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

DOI: 10.1002/nla

Table XIII. Iteration count for HSS multigrid on unsteady Oseen problem, $\sigma = 10$.

	$Re \equiv \nu^{-1}$			
$1/h$	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 XIV. Iteration count for inexact AL multigrid (block/upper triangular) on the unsteady Oseen problem, $\sigma = 10$. NC indicates that convergence was not achieved within 500 iterations.

	$Re \equiv \nu^{-1}$			
$1/h$	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 XV. Iteration count for B–S multigrid 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

flow) shows potential, considering the low-cost per iteration. The behavior of the B–S 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 AL 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 the total number of unknowns.

DOI: 10.1002/nla

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 B–S 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 AL formulation and the poroelasticity equations of [11] may also be of interest, especially with regard 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.

ACKNOWLEDGEMENTS

We thank the anonymous referee for suggestions that led to substantial improvements in the paper. The work of Steven Hamilton was supported in part by a U.S. Department of Energy Computational Science Graduate Fellowship. The work of Michele Benzi was supported in part by the National Science Foundation Grant DMS-0511336. The work of Eldad Haber was supported in part by the National Science Foundation Grant DMS-0724759 and by the U.S. Department of Energy Grant DMS 0724759.

REFERENCES

1. Elman H, Silvester D, Wathen A. *Finite Elements and Fast Iterative Solvers with Applications in Incompressible Fluid Dynamics*. Oxford University Press: Oxford, U.K., 2005.
2. Benzi M, Golub GH. A preconditioner for generalized saddle point problems. *SIAM Journal on Matrix Analysis and Applications* 2004; **26**:20–41.
3. Benzi M, Olshanskii MA. An augmented Lagrangian-based approach to the Oseen problem. *SIAM Journal on Scientific Computing* 2006; **28**:2095–2113.
4. Elman HC. Preconditioners for saddle point problems arising in computational fluid dynamics. *Applied Numerical Mathematics* 2002; **43**:75–89.
5. Elman HC. Preconditioning for the steady-state Navier–Stokes equations with low viscosity. *SIAM Journal on Scientific Computing* 1999; **20**:1299–1316.
6. Elman HC, Silvester DJ, Wathen AJ. Performance and analysis of saddle point preconditioners for the discrete steady-state Navier–Stokes equations. *Numerische Mathematik* 2002; **90**:665–688.
7. Olshanskii MA, Vassilevski YV. Pressure Schur complement preconditioners for the discrete Oseen problem. *SIAM Journal on Scientific Computing* 2007; **29**:2686–2704.
8. Trottenberg U, Oosterlee C, Schuller A. *Multigrid*. Academic Press: San Diego, 2001.
9. Wesseling P. *Principles of Computational Fluid Dynamics*. Springer Series in Computational Mathematics, vol. 29. Springer: New York, 2001.
10. Wesseling P, Oosterlee CW. Geometric multigrid with applications to computational fluid dynamics. *Journal of Computational and Applied Mathematics* 2001; **128**:311–334.
11. Oosterlee CW, Gaspar FJ. Multigrid relaxation methods for systems of saddle point type. *Applied Numerical Mathematics* 2008; **58**:1933–1950.
12. Vanka SP. Block-implicit multigrid solution of Navier–Stokes equations in primitive variables. *Journal of Computational Physics* 1986; **65**:138–158.
13. Bai ZZ, Golub GH, Ng MK. Hermitian and skew-Hermitian splitting methods for non-Hermitian positive definite linear systems. *SIAM Journal on Matrix Analysis and Applications* 2003; **24**:603–626.

DOI: 10.1002/nla

14. Benzi M, Liu J. An efficient solver for the incompressible Navier–Stokes equations in rotation form. *SIAM Journal on Scientific Computing* 2007; **29**:1959–1981.
15. Simoncini V, Benzi M. Spectral properties of the Hermitian and skew-Hermitian splitting preconditioner for saddle point problems. *SIAM Journal on Matrix Analysis and Applications* 2004; **26**:377–389.
16. Benzi M, Gander MJ, Golub GH. Optimization of the Hermitian and skew-Hermitian splitting iteration for saddle-point problems. *BIT Numerical Mathematics* 2003; **43**:881–900.
17. Harlow FH, Welch JE. Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface. *Physics of Fluids* 1965; **8**:2182–2189.
18. Wienands R, Joppich W. *Practical Fourier Analysis for Multigrid Methods*. Chapman & Hall: New York, 2005.
19. MATLAB R2009a. The MathWorks Inc., Natick, MA, 2009.
20. Saad Y, Schultz MH. GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing* 1986; **7**:856–869.
21. Oosterlee CW, Washio T. An evaluation of parallel multigrid as a solver and as a preconditioner for singularly perturbed problems. *SIAM Journal on Scientific Computing* 1998; **19**:87–100.
22. Wienands R, Oosterlee CW. Fourier analysis of GMRES(m) preconditioned by multigrid. *SIAM Journal on Scientific Computing* 2000; **22**:582–603.
23. Benzi M, Szyld DB, van Duin A. Orderings for incomplete factorization preconditioning of nonsymmetric problems. *SIAM Journal on Scientific Computing* 1999; **20**:1652–1670.
24. Saad Y. A flexible inner-outer preconditioned GMRES algorithm. *SIAM Journal on Scientific Computing* 1993; **14**:461–469.
25. Fortin M, Glowinski R. *Augmented Lagrangian Methods*: *Applications to the Numerical Solution of Boundary-Value Problems*. Studies in Applied Mathematics, vol. 15. North-Holland: Amsterdam, New York, Oxford, 1983.
26. Schöberl J. Multigrid methods for a parameter dependent problem in primal variables. *Numerische Mathematik* 1999; **84**:97–119.
27. Benzi M, Olshanskii MA, Wang Z. Modified Augmented Lagrangian preconditioners for the incompressible Navier–Stokes equations. *International Journal for Numerical Methods in Fluids* 2009; to appear.
28. Bertaccini D, Golub GH, Serra-Capizzano S. Spectral analysis of a preconditioned iterative method for the convection-diffusion equation. *SIAM Journal on Matrix Analysis and Applications* 2007; **29**:260–278.
29. Braess D, Sarazin R. An efficient smoother for the Stokes problem. *Applied Numerical Mathematics* 1997; **23**:3–19.
30. John V, Tobiska L. Numerical performance of smoothers in coupled multigrid methods for parallel solution of the incompressible Navier–Stokes equations. *International Journal for Numerical Methods in Fluids* 2000; **33**:453–473.

DOI: 10.1002/nla