
Hybrid Schwarz-Multigrid Methods for the Spectral Element Method: Extensions to Navier-Stokes

Paul F. Fischer¹ and James W. Lottes²

¹ Argonne National Laboratory, Mathematics and Computer Science Division
(<http://www.mcs.anl.gov/~fischer/>)

² University of Illinois, Dept. of Theoretical and Applied Mechanics

Summary. The performance of multigrid methods for the standard Poisson problem and for the consistent Poisson problem arising in spectral element discretizations of the Navier-Stokes equations is investigated. It is demonstrated that overlapping additive Schwarz methods are effective smoothers, provided that the solution in the overlap region is weighted by the inverse counting matrix. It is also shown that spectral element based smoothers are superior to those based upon finite element discretizations. Results for several large 3D Navier-Stokes applications are presented.

1 Introduction

The spectral element method (SEM) is a high-order weighted residual technique that combines the geometric flexibility of finite elements with the rapid convergence properties and tensor-product efficiencies of global spectral methods. Globally, elements are coupled in an unstructured framework with interelement coupling enforced through standard matching of nodal interface values. Locally, functions are represented as tensor products of stable N th-order Lagrangian interpolants based on Gauss-Lobatto (GL) or Gauss (G) quadrature points. For problems having smooth solutions, such as the incompressible Navier-Stokes equations, the SEM converges exponentially fast with the local approximation order N . Because of its minimal numerical dissipation and dispersion, the SEM is particularly well suited for the simulation of flows at transitional Reynolds numbers, where physical dissipation is small and turbulence-model dissipation is absent.

The two-level hierarchy of the spectral element discretization provides a natural route to domain decomposition with several benefits. The loose C^0 interelement coupling implies that the stencil depth does not increase with approximation order, so that interprocessor communication is minimal. The local tensor-product structure allows matrix-vector products to be recast as cache-efficient *matrix-matrix* products and also allows local subdomain problems to

be solved efficiently with fast tensor-product solvers. Finally, the high-order polynomial expansions provide a readily available sequence of nested grids (obtained through successive reductions in polynomial degree) for use in multilevel solvers.

This paper presents recent developments in spectral element multigrid (SEMG) methods. Our point of departure is the original work of Rønquist and Patera [1987] and Maday and Muñoz [1988], who developed variational SEMG for the two-dimensional Poisson problem using intra-element prolongation/restriction operators coupled with Jacobi smoothing. The high-aspect-ratio cells present in the tensor-product GL grid are a well-known source of difficulty in spectral multigrid methods and have drawn much attention over the past decade. We have developed multigrid smoothers in Lottes and Fischer [2004] based on the overlapping additive Schwarz method of Dryja and Widlund [1987] and Fischer et al. [2000]. We bypass the high-aspect-ratio cell difficulty by solving the local problems directly using fast tensor-product solvers; this approach ensures that the smoother cost does not exceed the cost of residual evaluation. Here, we extend our SEMG approach from the two-dimensional Laplacian to the more difficult consistent Poisson operator that governs the pressure in the mixed $\mathbb{P}_{N-1} - \mathbb{P}_{N-2}$ spectral element formulation of Maday and Patera [1989].

In the next section, we introduce the SE discretization for a model Poisson problem. The basic elements of our multilevel iterative procedures are presented in Section 3, along with results for the Poisson problem. Extensions to unsteady Navier-Stokes applications are described in Section 4.

2 Discretization of the Poisson Problem

The spectral element discretization of the Poisson problem in \mathbb{R}^d is based on the weighted residual formulation: *Find* $u \in X_N$ *such that*

$$(\nabla v, \nabla u)_{GL} = (v, f)_{GL} \quad \forall v \in X_N. \quad (1)$$

The inner product $(\cdot, \cdot)_{GL}$ refers to the Gauss-Lobatto-Legendre (GL) quadrature associated with the space $X_N := [Z_N \cap H_0^1(\Omega)]$, where $Z_N := \{v \in L^2(\Omega) | v|_{\Omega^e} \in \mathbb{P}_N(\Omega^e)\}$. Here, L^2 is the space of square integrable functions on Ω ; H_0^1 is the space of functions in L^2 that vanish on the boundary and whose first derivative is also in L^2 ; and $\mathbb{P}_N(\Omega^e)$ is the space of functions on Ω^e whose image is a tensor-product polynomial of degree $\leq N$ in the reference domain, $\hat{\Omega} := [-1, 1]^d$. For $d = 2$, a typical element in X_N is written

$$u(\mathbf{x}^e(r, s))|_{\Omega^e} = \sum_{i=0}^N \sum_{j=0}^N u_{ij}^e h_i^N(r) h_j^N(s), \quad (2)$$

where u_{ij}^e is the nodal basis coefficient; $h_i^N \in \mathbb{P}_N$ is the Lagrange polynomial satisfying $h_i^N(\xi_j) = \delta_{ij}$, where ξ_j , $j = 0, \dots, N$ are the the GL quadrature




Fig. 1. Spectral element configuration ($E = 9$, $N = 8$) showing Lagrange interpolation points for functions in X^N (left) and Y^N (right). The shaded regions illustrate the “minimal overlap” domain extension for the overlapping Schwarz smoothers.

points (the zeros of $(1 - \xi^2)L'_N(\xi)$, where L_N is the Legendre polynomial of degree N) and δ_{ij} is the Kronecker delta function; and $\mathbf{x}^e(r, s)$ is the coordinate mapping from $\widehat{\Omega}$ to Ω^e . We assume $\Omega = \cup_{e=1}^E \Omega^e$ and that the intersection of two subdomains (spectral elements) is an entire edge, a single vertex, or void. Function continuity ($u \in H^1$) is enforced by ensuring that nodal values on element boundaries coincide with those on adjacent elements. Figure 1 illustrates a spectral element decomposition of the square using $E = 9$ elements. The Gauss-Lobatto-based mesh on the left shows the nodal distribution for X_N . The Gauss-based mesh on the right is used for functions in Y_N , which will be introduced in the context of the Stokes discretization in Section 4.

Computational Preliminaries. Because we employ iterative solvers, we need an efficient procedure for evaluating matrix-vector products associated with the bilinear forms in (1). As noted by Orszag [1980], tensor-product bases play a key role in this respect, particularly for large N (i.e., $N \geq 8$). Here, we introduce several points that are central to our element-based solution strategy.

As with standard finite element methods, we assume availability of both local element-based and global mesh-based node numberings, with the local-to-global map given by $q(i_1, \dots, i_d, e) \in \{1, \dots, \bar{n}\}$, for $i_k \in \{0, \dots, N\}$, $k \in \{1, \dots, d\}$, and $e \in \{1, \dots, E\}$, where \bar{n} is the number of distinct global nodes. Let Q^T be the $\bar{n} \times E(N+1)^d$ matrix with columns $\hat{\underline{e}}_{q(i_1, \dots, i_d, e)}$, where $\hat{\underline{e}}_q$ denotes the q th column of the $\bar{n} \times \bar{n}$ identity matrix. Then the matrix-vector product $\underline{u}_L = Q\underline{u}$ represents a global-to-local mapping for any function $u(\mathbf{x}) \in X^N$, and the bilinear form on the left of (1) can be written

$$(\nabla v, \nabla u) = \underline{v}^T Q^T A_L Q \underline{u}, \quad (3)$$

where $A_L = \text{block-diag}(A^e)_{e=1}^E$ is the unassembled stiffness matrix comprising the local stiffness matrices, A^e , and Q^T and Q correspond to respective gather and scatter operations. In practice the global stiffness matrix, $A := Q^T A_L Q$, is never formed. One simply effects the *action* of A by applying each matrix to a vector through appropriate subroutine calls.

In the SEM, computational efficiency dictates that *local* stiffness matrices should also be applied in matrix-free form. The local contributions to (3) are

$$(\nabla v, \nabla u)_{GL}^e = (\underline{v}^e)^T A^e \underline{u}^e = (\underline{v}^e)^T \begin{pmatrix} D_1 \\ D_2 \end{pmatrix}^T \begin{pmatrix} G_{11}^e & G_{12}^e \\ G_{12}^e & G_{22}^e \end{pmatrix} \begin{pmatrix} D_1 \\ D_2 \end{pmatrix} \underline{u}^e, \quad (4)$$

with respective geometric factors and derivative operators,

$$G_{ij}^e := (\widehat{B} \otimes \widehat{B}) \left[\sum_{k=1}^d \frac{\partial r_i}{\partial x_k} \frac{\partial r_j}{\partial x_k} \right]^e J^e, \quad D_1 := (I \otimes \widehat{D}), \quad D_2 := (\widehat{D} \otimes I). \quad (5)$$

Here, \underline{v}^e and \underline{u}^e are vectors containing the lexicographically ordered nodal basis coefficients $\{v_{ij}^e\}$ and $\{u_{ij}^e\}$, respectively; $\widehat{B} = \text{diag}(\rho_k)_{k=0}^N$ is the one-dimensional mass matrix composed of the GL quadrature weights; and \widehat{D} is the one-dimensional derivative matrix with entries

$$\widehat{D}_{ij} = \left. \frac{dh_j}{dr} \right|_{\xi_i}, \quad i, j \in \{0, \dots, N\}^2.$$

The Jacobian, J^e , and metric terms (in brackets in (5)) are evaluated pointwise at each GL quadrature point, (ξ_p, ξ_q) , so that each of the composite geometric matrices, G_{ij}^e , is diagonal.

The presence of the cross terms, G_{12}^e , implies that A^e is full and requires storage of $(N+1)^4$ nonzeros for *each* spectral element if explicitly formed. In the spectral element method, this excessive storage and work overhead is avoided by retaining the factored form (5), which requires (to leading order) storage of only $3E(N+1)^2$ nonzeros and work of $\approx 8E(N+1)^3$ per matrix-vector product. The savings is more significant in 3D, where the respective storage and work complexities are $6E(N+1)^3$ and $\approx 12E(N+1)^4$ for the factored form, versus $O(EN^6)$ if A is explicitly formed. Moreover, the leading order work terms for the factored form can be cast as efficient matrix-matrix products, as discussed in detail by Deville et al. [2002]. These complexity savings can be extended to all system matrices and are the basis for efficient realizations of high-order weighted residual techniques.

If Ω^e is a regular parallelepiped, the local stiffness matrix simplifies to a separable form. For example, for an $L_x^e \times L_y^e$ rectangle, one would have

$$A^e = \frac{L_y^e}{L_x^e} \widehat{B} \otimes \widehat{A} + \frac{L_y^e}{L_x^e} \widehat{A} \otimes \widehat{B}, \quad \widehat{A} := \widehat{D}^T \widehat{B} \widehat{D}. \quad (6)$$

This form has a readily computable (pseudo-) inverse given by the fast diagonalization method (FDM) of Lynch et al. [1964],

$$A_e^{-1} = (S \otimes S) \left[\frac{L_y^e}{L_x^e} I \otimes \Lambda + \frac{L_y^e}{L_x^e} \Lambda \otimes I \right]^{-1} (S^T \otimes S^T), \quad (7)$$

where S is the matrix of eigenvectors and Λ the matrix of eigenvalues satisfying $\widehat{A}S = \widehat{B}S\Lambda$ and $S^T\widehat{B}S = I$. The bracketed term in (7) is diagonal, and its pseudo-inverse is computed by inverting nonzero elements and retaining zeros elsewhere. For arbitrarily deformed elements, the discrete Laplacian cannot be expressed in the tensor-product form (6), and the FDM cannot be used. For the purposes of a preconditioner, however, it suffices to apply the FDM to a regular parallelepiped of equivalent size, as demonstrated in Couzy [1995] and Fischer et al. [2000]. Similar strategies for the case of nonconstant coefficients are discussed by Shen [1996].

3 Multilevel Solvers

We are interested in methods for solving the global system $A\underline{u} = \underline{g}$. To introduce notation, we consider the two-level multigrid sweep.

Procedure Two-Level: (8)

- i) $\underline{u}^{k+1} = \underline{u}^k + \sigma M(\underline{g} - A\underline{u}^k), \quad k = 0, \dots, m_d - 1$
- ii) $\underline{r} = \underline{g} - A\underline{u}^{m_d}$
- iii) $\underline{\tilde{e}} = \sigma_C P A_C^{-1} P^T \underline{r}$
- iv) $\underline{\tilde{u}}^0 = \underline{u}^{m_d} + \underline{\tilde{e}}$
- v) $\underline{\tilde{u}}^{k+1} = \underline{\tilde{u}}^k + \sigma M(\underline{g} - A\underline{\tilde{u}}^k), \quad k = 0, \dots, m_u - 1$
- vi) If $\|A\underline{\tilde{u}}^{m_u} - \underline{g}\| < \text{tol}$, set $\underline{u} := \underline{\tilde{u}}^{m_u}$, *quit*.
Else, $\underline{u}^0 := \underline{\tilde{u}}^{m_u}$, go to (i).

Here M is the smoother, σ and σ_C are relaxation parameters, and m_d and m_u are the number of smoothing steps on the downward and upward legs of the cycle, respectively. Steps (i) and (v) are designed to eliminate high-frequency error components that cannot be represented on the coarse grid. The idea is that the error after (ii), $\underline{e} := A^{-1}\underline{r}$, should be well approximated by $\underline{\tilde{e}}$, which lies in the coarse-grid space represented by the columns of P . The coarse-grid problem, A_C^{-1} , is solved directly, if A_C is sufficiently sparse, or approximated by recursively applying the two-level procedure to the A_C system, giving rise to the multigrid “V” cycle. The prolongation matrix P interpolates from the coarse space to the fine nodes using the local tensor-product basis functions for the coarse space.

If the two-level procedure is used as a preconditioner, we take $\underline{u}^0 = 0$, $m_d = 1$, and $m_u = 0$, and the procedure simplifies to the following.

Procedure Preconditioner: (9)

- i)* $\underline{u}^1 = \sigma M \underline{g}$,
- ii)* $\underline{r} = \underline{g} - A \underline{u}^1$
- iii)* $\underline{\tilde{e}} = \sigma_C P A_C^{-1} P^T \underline{r}$
- iv)* $\underline{u} = \underline{u}^1 + \underline{\tilde{e}}$, return.

The preconditioner can be viewed either as an application of the multigrid V-cycle or as a two-level multiplicative Schwarz method (Smith et al. [1996]). By simply replacing *(ii)* with $\underline{r} = \underline{g}$, we obtain a two-level additive Schwarz method, which has the advantage of avoiding an additional multiplication by A . This savings is important in the Navier-Stokes applications that we consider in Section 4.

Smoothers for the Poisson Problem. Here, we review the SEMG smoothing strategies considered for the Poisson problem in Lottes and Fischer [2004]. Our original intent was to base the smoother, M , on the additive overlapping Schwarz method of Dryja and Widlund [1987], with local subdomain problems discretized by finite elements (FEs) having nodes coincident with the GL nodes, as considered by Casarin [1997], Fischer [1997], and Pahl [1993]. By using the fast diagonalization method to solve the local problems, however, we are freed from the constraint of using FE-based preconditioners because the cost depends only on the use of tensor-product forms and not on the sparsity of the originating operator. Hence, we are able to consider subdomain problems derived as restrictions of the originating spectral element matrix, A , as first suggested by Casarin [1997].

The use of Schwarz-based smoothing, which is arguably more expensive than traditional smoothers, is motivated by several factors. First, it is not practical to apply Gauss-Seidel smoothing in the SEM because the matrix entries are not available (see (4)). The alternative of pointwise-Jacobi smoothing was shown by Rønquist [1988] and Maday et al. [1992] not to scale for $d > 1$. Specifically, the authors demonstrated a convergence factor of $\rho = 0.75$ for $d = 1$, but only $\rho = 1 - c/N \log N$ for $d = 2$. Second, while one can exploit the SE-FE spectral equivalence established by Orszag [1980] to ostensibly convert the SE problem into a FE problem and then apply traditional multigrid, the FE problem inherits the difficulties of its SE counterpart, namely, the high-aspect ratio cells that arise from the tensor-product of the one-dimensional Gauss-Lobatto grids. Moreover, even if the GL-based FE problem could be solved with low work, the iteration count would still be higher than what is observed for the Schwarz-based approach. Third, to minimize cost, it is reasonable to have a smoother whose cost is on par with that of residual evaluation if it can substantially reduce the iteration count.

We illustrate the problem of high-aspect ratio cells by considering application of the two-level procedure (8) to the model Poisson problem (1) discretized on the unit square with an $8N \times 8N$ array of bilinear finite elements. Iteration counts for four different smoothing strategies are shown in Table 1. Jacobi implies $M^{-1} := \text{diag}(A)$; GSRB is a Gauss-Seidel sweep with the nodes

Table 1. MG method on FE problem.

FE Spacing	No.	Smoother/ Preconditioner	Coarse Space	Iterations, 10^{-11} Reduction			
				$N = 4$	$N = 8$	$N = 12$	$N = 16$
Uniform	a	Jacobi	$N/2$	39	38	38	38
	b	GSRB	$N/2$	9	9	9	9
	c	H Schwarz	$N/2$	40	41	42	42
	d	H Schwarz (W)	$N/2$	7	7	6	6
SE	e	Jacobi	$N/2$	41	84	148	219
	f	GSRB	$N/2$	11	28	46	65
	g	H Schwarz	$N/2$	40	43	47	52
	h	H Schwarz (W)	$N/2$	6	7	7	9

ordered into two maximally independent (“red-back”) subsets; and H Schwarz and H Schwarz(W) correspond to the hybrid Schwarz-based smoothers introduced below. In all cases, σ is chosen such that the maximum eigenvalue of σMA is unity, and $\sigma_C = 1$. The coarse system is solved directly and is based on the same FE discretization, save that, in each direction, every other nodal point is discarded. The first set of results is for uniformly sized elements of length $1/8N$ on each side. Resolution-independent convergence is obtained for each of the smoothing strategies, with GSRB and H Schwarz(W) being the most competitive. Although H Schwarz(W) has a lower iteration count, GSRB requires less work per iteration, and the two are roughly equal in computational cost. The second set of results is for an $8N \times 8N$ array of bilinear elements whose vertices coincide with the GL node spacing associated with an 8×8 array of spectral elements of order N . In this case, the pointwise Jacobi and GSRB smoothers break down as N is increased. Only H Schwarz(W) retains performance comparable to the uniform grid case. We note that line-based relaxation strategies proposed by Shen et al. [2000] and Beuchler [2002] also compensate for the high-aspect-ratio cell difficulty. For the values of N considered here, however, the hybrid Schwarz approach is likely to be faster, at least on cache-based architectures, where the matrix-matrix product-oriented fast-diagonalization method is very effective.

Our hybrid Schwarz strategy is based on a multiplicative combination of an additive Schwarz smoother at the fine scale and a coarse-grid correction. The smoother, originally due to Dryja and Widlund [1987], is written as

$$M := \sum_{e=1}^E R_e^T A_e^{-1} R_e. \quad (10)$$

Here, R_e is the standard Boolean restriction matrix that extracts from a global nodal vector those values associated with the *interior* of the extended subdomain $\bar{\Omega}^e$. In all cases, $\bar{\Omega}^e$ is an extension of Ω^e that includes a single row (or plane, in 3D) of nodal values in each of $2d$ directions, as illustrated in Fig. 1 (left). R_e^T extends by zero the vector of nodal values interior to $\bar{\Omega}^e$

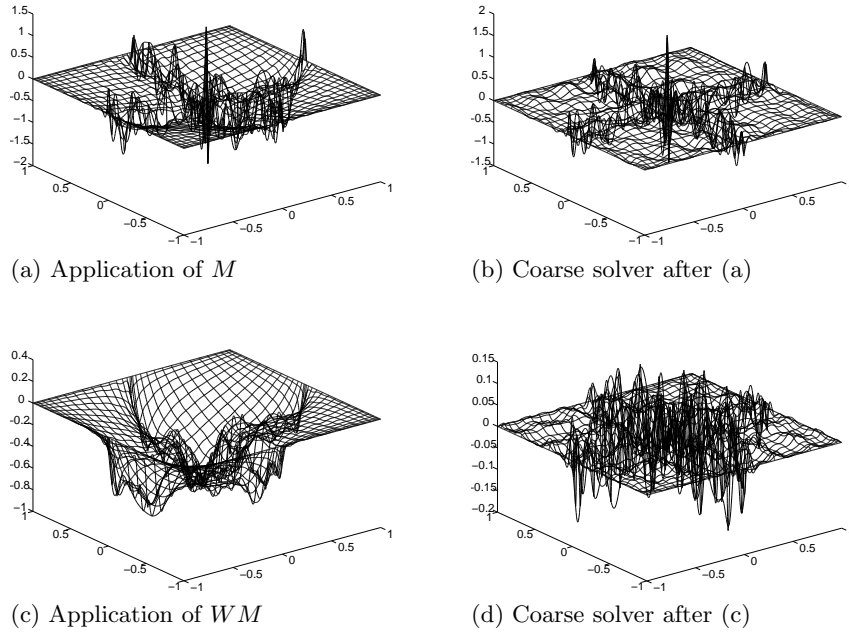


Fig. 2. Error plots for the hybrid Schwarz preconditioner and coarse solve, with $N_C = N/2$ and $(E, N) = (4, 16)$, applied to a random initial guess.

to a full length vector. Multiplication by A_e^{-1} is effected by using the fast diagonalization method similar to (7). In a preprocessing step, one assembles *one-dimensional* stiffness and mass matrices, A_* and B_* ($*$ = x , y or z), for each space dimension, $1, \dots, d$; restricts these to the relevant ranges using a one-dimensional restriction matrix R_*^e ; and solves an eigenvalue problem of the form $((R_*^e)^T A_* R_*^e) S_*^e = ((R_*^e)^T B_* R_*^e) S_*^e \Lambda_*^e$ to obtain the requisite eigenpairs (S_*^e, Λ_*^e) . Because the spectral elements are compactly supported, the preprocessing step requires knowledge only of the size of the elements on either side of Ω^e , in each of the d directions. For subdomains that are not rectilinear, A_e is based on average lengths in each direction.

We have found it important to weight the solutions in the overlap region by the inverse of the diagonal counting matrix

$$C := \sum_{e=1}^E R_e^T R_e. \quad (11)$$

The entries of C enumerate the number of subdomains that share a particular vertex. Setting $W = C^{-1}$ gives rise to the weighted overlapping Schwarz smoother $M_W := WM$. Although convergence theory for the weighted Schwarz method is yet to be developed, the methodology of Frommer and

Szyld [2001] should be applicable to this setting as well. In addition to reducing the maximum eigenvalue of MA (which, by simple counting arguments, is $\max C_{ii}$; see Smith et al. [1996]), multiplication by W significantly improves the smoothing performance of the additive Schwarz step. This latter point is illustrated in Fig. 2, which shows the error when the two-level preconditioner (9) is applied to random right-hand-side vector for a 2×2 array of spectral elements with $N=16$. Figure 2(a) shows the error after a single application of the additive Schwarz smoother (10), with $\sigma=1$. While the solution is smooth in the interior, there is significant undamped error along the interface, particularly at the cross point. As noted by Lottes and Fischer [2004], the error along the interface can be reduced by choosing $\sigma = 1/4$, but the overall error is no longer smooth. In either case, the subsequent coarse-grid correction does not yield a significant error reduction. By contrast, the error after application of M_W , seen in Fig. 2(c), is relatively smooth, and the coarse-grid correction is very effective. Comparing the magnitudes in Figs. 2(b) and 2(d), one sees a tenfold reduction in the error through the introduction of W .

Table 2 presents convergence results for the Poisson problem on the square discretized with an 8×8 array of spectral elements. Case 2(a) shows results for the unweighted additive Schwarz preconditioner using an FE-based smoother. This scheme is the Poisson equivalent to the method developed by Fischer et al. [2000] for the pressure subproblem considered in the next section. For all the other cases, A_e is based on a restriction of A rather than on an FE discretization. Case 2(b) shows that this simple change considerably reduces the iteration count. Enriching the coarse space from $N_C = 1$ to $N/2$ and incorporating the weight matrix W yields further reductions in iteration count and work. (Because of symmetry requirements, W is applied as a pre- and postmultiplication by $W^{1/2}$ for the preconditioned conjugate gradient, PCG, cases). The work shown in the last column of Table 2 is an estimate of the number of equivalent matrix-vector products required to reduce the error by 10^{-11} . Rather than attempting to symmetrize the hybrid Schwarz method (9), we simply switched to GMRES, which allowed W to be applied directly during the summation of the overlapping solutions. Comparison of cases 2(f) and 2(h) underscores the importance of weighting.

4 Extension to Navier-Stokes

Efficient solution of the incompressible Navier-Stokes equations in complex domains depends on the availability of fast solvers for sparse linear systems. For unsteady flows, the pressure operator is the leading contributor to stiffness, as the characteristic propagation speed is infinite. Our pressure solution procedure involves two stages. First, we exploit the fact that we solve similar problems from one step to the next by projecting the current solution onto a subspace of previous solutions to generate a high-quality initial approximation, as outlined in Fischer [1998]. We then compute the correction to this

Table 2. Iteration count for E=8×8 SE problem.

Method	No.	Smoother/ Preconditioner	Coarse Space	Iterations, 10^{-11} Reduction				Work $N = 16$
				$N = 4$	$N = 8$	$N = 12$	$N = 16$	
PCG	a	A Schwarz (FE)	1	28	35	46	58	116
	b	A Schwarz	1	25	27	35	43	86
	c	A Schwarz	$N/2$	26	26	26	27	81
	d	A Schwarz (W)	1	17	24	33	43	86
	e	A Schwarz (W)	$N/2$	16	21	22	24	72
MG/ GMRES	f	H Schwarz	$N/2$	21	23	24	25	100
	g	H Schwarz (W)	1	14	20	29	36	108
	h	H Schwarz (W)	$N/2$	13	12	12	13	52

approximation using a scalable iterative solver. Here, we extend the multigrid methods presented in the preceding sections to computation of the pressure in SE-based simulations of incompressible flows.

To introduce notation, we review the Navier-Stokes discretization presented in detail in Fischer [1997]. The temporal discretization is based on a semi-implicit scheme in which the nonlinear term is treated explicitly and the remaining unsteady Stokes problem is solved implicitly. Our spatial discretization is based on the $\mathbb{P}_N - \mathbb{P}_{N-2}$ spectral element method of Maday and Patera [1989]. Assuming \mathbf{f}^n incorporates all terms explicitly known at time t^n , the $\mathbb{P}_N - \mathbb{P}_{N-2}$ formulation of the Navier-Stokes problem reads: Find $(\mathbf{u}^n, p^n) \in X_N \times Y_N$ such that

$$\begin{aligned} \frac{1}{Re}(\nabla \mathbf{v}, \nabla \mathbf{u}^n)_{GL} + \frac{1}{\Delta t}(\mathbf{v}, \mathbf{u}^n)_{GL} - (\nabla \cdot \mathbf{v}, p^n)_G &= (\mathbf{v}, \mathbf{f}^n)_{GL}, \\ (q, \nabla \cdot \mathbf{u}^n)_G &= 0, \end{aligned} \quad (12)$$

$\forall (\mathbf{v}, q) \in X_N \times Y_N$. The inner products $(\cdot, \cdot)_{GL}$ and $(\cdot, \cdot)_G$ refer to the Gauss-Lobatto-Legendre (GL) and Gauss-Legendre (G) quadratures associated with the spaces $X_N := [Z_N \cap H_0^1(\Omega)]^d$ and $Y_N := Z_{N-2}$, respectively, and Z_N is the space introduced in conjunction with (1). The local velocity basis is given, componentwise, by the form (2). The pressure is similar, save that the nodal interpolants are based on the $N-1$ Gauss points, $\eta_i \in (-1, 1)$, as illustrated in Fig. 1 (right).

Insertion of the SEM bases into (12) yields a discrete Stokes system to be solved at each step:

$$\mathbf{H} \underline{\mathbf{u}}^n - \mathbf{D}^T \underline{p}^n = \mathbf{B} \underline{\mathbf{f}}^n, \quad \mathbf{D} \underline{\mathbf{u}}^n = 0. \quad (13)$$

$\mathbf{H} = \frac{1}{Re} \mathbf{A} + \frac{1}{\Delta t} \mathbf{B}$ is the discrete equivalent of the Helmholtz operator, $(-\frac{1}{Re} \nabla^2 + \frac{1}{\Delta t})$; $-\mathbf{A}$ is the discrete Laplacian; \mathbf{B} is the (diagonal) mass matrix associated with the velocity mesh; \mathbf{D} is the discrete divergence operator, and $\underline{\mathbf{f}}^n$ accounts for the explicit treatment of the nonlinear terms. Note that the Galerkin approach implies that the governing system in (13) is symmetric and

that the matrices \mathbf{H} , \mathbf{A} , and \mathbf{B} are all symmetric positive definite. We have used bold capital letters to indicate matrices that interact with vector fields.

The Stokes system (13) is advanced by using the operator splitting approach presented by Maday et al. [1990] and Perot [1993]. One first solves

$$\mathbf{H} \hat{\mathbf{u}} = \mathbf{B} \mathbf{f}^n + \mathbf{D}^T \underline{p}^{n-1}, \quad (14)$$

which is followed by a pressure correction step

$$E \delta \underline{p}^n = -\frac{1}{\Delta t} \mathbf{D} \hat{\mathbf{u}}, \quad \underline{\mathbf{u}}^n = \hat{\mathbf{u}} + \Delta t \mathbf{B}^{-1} \mathbf{D}^T \delta \underline{p}^n, \quad \underline{p}^n = \underline{p}^{n-1} + \delta \underline{p}^n, \quad (15)$$

where $E := \mathbf{D} \mathbf{B}^{-1} \mathbf{D}^T$ is the Stokes Schur complement governing the pressure in the absence of the viscous term.

E is the consistent Poisson operator for the pressure and is spectrally equivalent to A . Through a series of tests that will be reported elsewhere, we have found the following to be an effective multilevel strategy for solving E . We employ (9) to precondition GMRES with a weighted additive Schwarz smoother. The local problems are based on $E_e := \tilde{R}_e E \tilde{R}_e^T$, where the subdomains defined by the restriction matrices \tilde{R}_e correspond to the minimal-overlap extension illustrated in Fig. 1 (right). The coarse-grid problem, A_C , is based on A with $N_C = N/2$ (typically), which was found not only to be cheaper but also better at removing errors along the element interfaces. At all intermediate levels, A_C^{-1} is approximated with a single V-cycle (8).

The local problems are solved using the fast diagonalization method, which requires that E_e (and therefore E) be separable. In two dimensions, we need to cast E in the form

$$E = J_y \otimes E_x + E_y \otimes J_x. \quad (16)$$

For simplicity, we assume that we have a single element with $\Omega = \hat{\Omega}$ and ignore the details of boundary conditions. From the definition of E , we have

$$E = D_x B^{-1} D_x^T + D_y B^{-1} D_y^T. \quad (17)$$

The divergence and inverse mass matrices have the tensor-product forms

$$D_x = (\tilde{B} \otimes \tilde{B})(\tilde{J} \otimes \tilde{D}), \quad D_y = (\tilde{B} \otimes \tilde{B})(\tilde{D} \otimes \tilde{J}), \quad B^{-1} = (\hat{B}^{-1} \otimes \hat{B}^{-1}) \quad (18)$$

Here, $\tilde{B} = \text{diag}(\tilde{\rho}_i)_{i=1}^{N-1}$ consists of the Gauss-Legendre quadrature weights, and \tilde{J} and \tilde{D} are respective interpolation and derivative matrices mapping from the GL points to the G points,

$$\tilde{J}_{ij} = h_j^N(\eta_i), \quad \tilde{D}_{ij} = \left. \frac{dh_j^N}{dr} \right|_{r=\eta_i}. \quad (19)$$

Inserting (18) into (17) yields the desired form (16) with




Fig. 3. SE Navier-Stokes examples: (a) $E = 1021$ mesh, inlet profile, and vorticity contours for roughness element; (b) $E = 1536$ mesh and (c) temperature contours for buoyancy driven convection; (d) $E = 2544$ mesh and (e) coherent structures for flow in a diseased carotid artery.

$$J_x = J_y = \tilde{B}\tilde{J}\hat{B}^{-1}\tilde{J}^T\tilde{B}^T, \quad E_x = E_y = \tilde{B}\tilde{D}\hat{B}^{-1}\tilde{D}^T\tilde{B}^T. \quad (20)$$

The extension to multiple elements follows by recognizing that the gather-scatter operator used to assemble the local matrices can be written as $Q = Q_y \otimes Q_x$ for a tensor-product array of elements. Following our element-centric solution strategy, we thus generate E_e by viewing Ω^e as being embedded in a 3^d array of rectilinear elements of known dimensions. Unlike A_e , the entries of E_e are also influenced by the “neighbors of neighbors.” This influence, however, is small and is neglected.

Navier-Stokes Results. We turn now to application of spectral element multigrid (SEMG) to the simulation of unsteady incompressible flows. In nu-




Fig. 4. Iteration histories for FE-based two-level (STD), weighted SE-based two-level (WGT), additive multilevel (ADD), and multiplicative multilevel (HYB) schemes for spectral elements simulations of order $N=9$: (a) hairpin vortex, $E=1021$; (b) hemispherical convection, $E=1536$; and (c) carotid artery simulation, $E=2544$.

merous 2D and 3D Navier-Stokes test problems, we have found the additive variant of the procedure outlined in the preceding section to be roughly two to three times faster than the two-level additive Schwarz method developed in Fischer et al. [2000]. A sample of these results is presented below.

We consider the three test cases shown in Fig. 3. The first case, Fig. 3(a), is boundary-layer flow past a hemispherical roughness element at Reynolds number $Re=1000$ (based on roughness height). The flow generates a pre-transitional chain of hairpin vortices evidenced by the spanwise vorticity contours shown in the symmetry plane. The second example, Fig. 3(b)-(c), is buoyancy-driven convection in a rotating hemispherical shell having inner radius 2.402 and outer radius 3.3. The Rayleigh number (based on shell thickness) is $Ra=20,000$ and the Taylor number is $Ta=160,000$. The third case, Fig. 3(d)-(e), simulates transitional flow in a diseased carotid artery. The severe stenosis in the internal (right) branch results in high flow velocities and, ultimately, transition to turbulence. Figure 3(e) shows the coherent structures that arise just before peak systole.

Figure 4 shows the pressure iteration history for the first 85 timesteps of the three examples, using the initial conditions of Fig. 3. For all cases, $N=9$ and the coarse problem is based on linear elements whose vertices are derived from an oct-refinement of the SE mesh. Four methods are considered: STD refers to the two-level additive Schwarz method of Fischer et al. [2000]; WGT is the same as STD, save that E_e is based on a restriction of E , rather than an FE-based discretization of the Poisson problem, and that the weight matrix W is included; ADD is the same as WGT, save that three levels are employed, with $N_{\text{mid}}=5$; HYB is the same as ADD, save that the multiplicative variant of (9) is used. PCG is used for STD and WGT, whereas GMRES is used for ADD and HYB. Although HYB requires fewer iterations, ADD is the fastest method because it requires only one product in E per iteration. The prominent spikes in Fig. 4(b) result from resetting the projection basis set (Fischer [1998]).

Table 3. ADD Avg. Iteration Count for Navier-Stokes Examples

Problem	$N=5$	$N=7$	$N=9$	$N=11$	$N=13$	$N=15$	$N=17$
Hairpin Vortex	9.8	11.1	15.1	17.5	20.4	23.5	26.1
Spherical Conv.	8.2	7.8	8.3	8.9	9.9	10.9	11.6
Carotid Artery	18.5	20.6	23.7	26.0	29.3	32.5	36.0
Carotid (wGT)	16.5	22.2	30.0	39.5	48.4	59.4	65.8

The scalability of the three-level ADD method is illustrated in Table 3, which shows the average iteration count over the last 20 steps for varying N with $N_{\text{mid}}=N/2$. Order-independence is not assured in complex domains, particularly if the mesh contains high aspect-ratio elements (Fischer [1997]). The performance of ADD is nonetheless quite reasonable when one considers that the number of pressure nodes varies by a factor of 64 in moving from $N = 5$ to 17. For purposes of comparison, results for the wGT method are shown for the carotid. The additional level of ADD clearly reduces order dependence.

Acknowledgement. This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract W-31-109-Eng-38.

References

- S. Beuchler. Multigrid solver for the inner problem in domain decomposition methods for p-fem. *SIAM J. Numer. Anal.*, 40:928–944, 2002.
- M. A. Casarin. Quasi-optimal Schwarz methods for the conforming spectral element discretization. *SIAM J. Numer. Anal.*, 34:2482–2502, 1997.
- W. Couzy. *Spectral Element Discretization of the Unsteady Navier-Stokes Equations and its Iterative Solution on Parallel Computers*. PhD thesis, Swiss Federal Institute of Technology-Lausanne, 1995. Thesis nr. 1380.
- M. O. Deville, P. F. Fischer, and E. H. Mund. *High-order methods for incompressible fluid flow*. Cambridge University Press, Cambridge, 2002.
- M. Dryja and O. B. Widlund. An additive variant of the Schwarz alternating method for the case of many subregions. Technical Report TR 339, Courant Inst., NYU, 1987. Dept. Comp. Sci.
- P. F. Fischer. An overlapping Schwarz method for spectral element solution of the incompressible Navier-Stokes equations. *J. Comput. Phys.*, 133:84–101, 1997.
- P. F. Fischer. Projection techniques for iterative solution of $Ax = b$ with successive right-hand sides. *Comput. Methods Appl. Mech. Engrg.*, 163:193–204, 1998.
- P. F. Fischer, N. I. Miller, and H. M. Tufo. An overlapping Schwarz method for spectral element simulation of three-dimensional incompressible flows. In

- P. Bjørstad and M. Luskin, editors, *Parallel Solution of Partial Differential Equations*, pages 158–180, Berlin, 2000. Springer.
- A. Frommer and D. B. Szyld. An algebraic convergence theory for restricted additive Schwarz methods using weighted max norms. *SIAM Journal on Numerical Analysis*, 39:463–479, 2001.
- J. W. Lottes and P. F. Fischer. Hybrid multigrid/Schwarz algorithms for the spectral element method. *J. Sci. Comput.*, (to appear), 2004.
- R. E. Lynch, J. R. Rice, and D. H. Thomas. Direct solution of partial difference equations by tensor product methods. *Numer. Math.*, 6:185–199, 1964.
- Y. Maday and R. Muñoz. Spectral element multigrid: Numerical analysis. *J. Sci. Comput.*, 3:323–354, 1988.
- Y. Maday, R. Muñoz, A. T. Patera, and E. M. Rønquist. Spectral element multigrid methods. In P. de Groen and R. Beauwens, editors, *Proc. of the IMACS Int. Symposium on Iterative Methods in Linear Algebra, Brussels, 1991*, pages 191–201, Amsterdam, 1992. Elsevier.
- Y. Maday and A. T. Patera. Spectral element methods for the Navier-Stokes equations. In A. K. Noor and J. T. Oden, editors, *State-of-the-Art Surveys in Computational Mechanics*, pages 71–143. ASME, New York, 1989.
- Y. Maday, A. T. Patera, and E. M. Rønquist. An operator-integration-factor splitting method for time-dependent problems: Application to incompressible fluid flow. *J. Sci. Comput.*, 5:263–292, 1990.
- S. A. Orszag. Spectral methods for problems in complex geometry. *J. Comput. Phys.*, 37:70–92, 1980.
- S. S. Pahl. Schwarz type domain decomposition methods for spectral element discretizations. Master's thesis, Univ. of Witwatersrand, Johannesburg, South Africa, 1993. Dept. of Computational and Applied Math.
- J. B. Perot. An analysis of the fractional step method. *J. Comp. Phys.*, 108:51–58, 1993.
- E. Rønquist. *Optimal Spectral Element Methods for the Unsteady Three-Dimensional Incompressible Navier-Stokes Equations*. PhD thesis, Massachusetts Institute of Technology, 1988. Cambridge, MA.
- E. M. Rønquist and A. T. Patera. Spectral element multigrid, I.: Formulation and numerical results. *J. Sci. Comput.*, 2:389–406, 1987.
- J. E. Shen. Efficient Chebyshev-Legendre Galerkin methods for elliptic problems. In A. V. Ilin and L. R. Scott, editors, *Third Int. Conference on Spectral and High Order Methods*, pages 233–239. Houston Journal of Mathematics, 1996.
- J. E. Shen, F. Wang, and J. Xu. A finite element multigrid preconditioner for Chebyshev-collocation methods. *Appl. Numer. Math.*, 33:471–477, 2000.
- B. Smith, P. Bjørstad, and W. Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic PDEs*. Cambridge University Press, Cambridge, 1996.