

Hybrid Multigrid/Schwarz Algorithms for the Spectral Element Method

James W. Lottes* and Paul F. Fischer†

May 30, 2003

Abstract

We study the performance of the multigrid method applied to spectral element (SE) discretizations of the Poisson equation. Smoothers based on finite element (FE) discretizations, overlapping Schwarz methods, and point-Jacobi are considered in conjunction with conjugate gradient and GMRES acceleration techniques. It is found that Schwarz methods based on restrictions of the originating SE matrices converge faster than FE-based methods and that weighting the Schwarz matrices by the inverse of the diagonal counting matrix is essential to effective Schwarz smoothing. Several of the methods considered achieve convergence rates comparable to those attained by classic multigrid on regular grids.

1 Introduction

The availability of fast elliptic solvers is essential to many areas of scientific computing. For unstructured discretizations in three dimensions, iterative solvers are generally optimal from both a work and storage standpoint. Ideally, one would like to have computational complexity that scales as $O(n)$ for an n -point grid problem in \mathbb{R}^d , implying that the iteration count should be bounded as the mesh is refined. Modern iterative methods such as multigrid and Schwarz-based domain decomposition achieve bounded iteration counts through the introduction of multiple representations of the solution (or the residual) that allow efficient elimination of the error at each scale. The theory for these methods is well established for classical finite difference (FD) and finite element (FE) discretizations, and order-independent convergence rates are often attained in practice.

For spectral element (SE) methods, there has been significant work on the development of Schwarz-based methods that employ a combination of local subdomain solves and sparse global solves to precondition conjugate gradient iteration. Theoretical work, due to Pavarino and Widlund [31], Mandel [23], Casarin [2, 3, 4], and others, and practical experience [13, 14, 15, 28, 34] indicate that these methods achieve order-independent convergence, or at least convergence rates that are independent of the number of subdomains (elements). The iteration counts, however, are generally not as low as observed for classical FD- or FE-based multigrid schemes. It is natural to consider exploiting the nested spaces that are intrinsic to

*Dept. of Theoretical and Applied Mechanics, University of Illinois, Urbana, IL 61801

†Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439

the spectral element method, constructed through successive reductions in polynomial order, to develop an efficient spectral element multigrid (SEMG) procedure. Early work in this direction was undertaken by Rønquist and Patera [35], who employed Jacobi smoothing in an intra-element SEMG implementation. Extensive numerical experiments with this approach are presented in the thesis of Rønquist [32], and theoretical analysis is presented by Maday and Muñoz [21] and Maday et al. [22]. They showed that a multigrid convergence factor of $\rho = 0.75$ was attained for $d = 1$, independent of the number of elements E and polynomial degree N . For $d = 2$, however, they found $\rho \sim 1 - c/N$ for standard Jacobi-based smoothing and $\rho \sim 1 - c/\sqrt{N}$ when Jacobi-based smoothing is coupled with Chebyshev acceleration. Several authors have attributed the performance degradation for $d > 1$ to the high-aspect-ratio cells that are present in the SE grids, which are based on tensor-products of Gauss-Lobatto nodal bases. Spectral multigrid schemes incorporating semi-coarsening and/or line relaxation have been proposed to address this problem in [18, 19, 37, 1]. (High-aspect-ratio *subdomains* in the Schwarz context have been addressed in [24, 13, 14].)

An alternative approach to preconditioning the spectral element method is to exploit the equivalence between the spectral operators and FD or FE discretizations based on the *same* set of Gauss-Lobatto nodal points, as first suggested by Orszag [27]. If A_s is the SE stiffness matrix and A_f is the FE stiffness matrix, then $A_f^{-1}A_s$ has a bounded condition number (typically $\kappa \sim \pi^2/4$ [9, 10, 29]), and conjugate gradient iteration will converge in a fixed number of iterations. This equivalence effectively reduces the SE problem to an FE problem, ostensibly solvable by classical multigrid. As we will demonstrate, however, the FE problem inherits the high-aspect ratio difficulties associated with the originating SE problem, and one must resort to more powerful techniques, such as semi-coarsening or line relaxation, to achieve order-independent convergence. A pioneering effort in this direction is the work of Zang et al. [41, 42], who tackled the FE preconditioning step with multigrid that employed incomplete Cholesky/conjugate gradient smoothing. Drawbacks of this approach in the SE context are the lack of parallelism and the need to establish a global ordering for the ICCG sweeps, which moves away from the element-centric paradigm that characterizes the SE method. The SE-FE equivalence has also played a significant role in advancing the development of domain decomposition approaches to solving the spectral element problem. Pahl [28] showed that the FE discretizations provide an excellent basis for Schwarz preconditioning of the SE problem, a result borne out in the theory developed by Casarin [3] and in numerous Navier-Stokes applications [13, 15, 16]. Casarin [3] also shows that the local problems in the overlapping Schwarz method can be based on restrictions of the original SE matrix, a point that we revisit in this paper. We note that Pavarino and Warburton have shown that the FE/Schwarz-based preconditioning strategy also works well in the context of triangular elements using high-order nodal bases [30].

An important development in SE/domain decomposition strategies, originally due to Couzy and Deville [7], is to exploit the local tensor-product structure within each (quadrilateral or hexahedral) spectral element by employing the fast-diagonalization method (FDM) of Lynch et al. [20] for the solution of the local problems. The storage and work complexities of the FDM are $O(EN^{d+1})$ and $O(EN^d)$, respectively, with constants essentially identical to those for matrix-vector products in A_s . Thus, more sophisticated local-solve

strategies (e.g., multigrid) would reduce the global solution time by at most a factor of two in comparison with the FDM approach. In fact, because the $O(EN^{d+1})$ work term arises from cache-efficient matrix-matrix products, the FDM is faster than any other approach for the moderate values of N typically used in practice; any iterative approach would need to converge in fewer than two iterations to compete. The FDM/Schwarz approach has been successfully applied in two- and three-dimensional incompressible Navier-Stokes simulations [5, 15] and in the solution of the shallow-water equations on the sphere [40].

In this paper, we combine the intra-element SEMG technique of Rønquist and Patera [35] and Maday et al. [22] with the more recent developments in spectral element/Schwarz methods. The main idea is to employ the Schwarz overlapping method [11, 13], coupled with subdomain solves based on the FDM, to provide a multigrid smoothing step. Because the subdomain problems are solved exactly (and efficiently), the ill-conditioning attributed to high-aspect ratio cells within the subdomains is avoided. For completeness, we present several competitive methods and compare with performance results presented elsewhere, when available.

In the next section, we introduce the SE discretization for a model Poisson problem. The basic elements of the iterative procedures are described in Section 3 and, several smoothing strategies are presented in Section 4. The error and smoothing behavior for a few of the methods is discussed in Section 5. Numerical experiments demonstrating the convergence behavior of the methods are presented in Section 6. Additional discussion and concluding remarks are given in Section 7.

2 Discretization

We take as our model problem the Poisson equation in the square $\Omega := [-1, 1]^2$,

$$-\nabla^2 u = f \text{ in } \Omega, \quad u = 0 \text{ on } \partial\Omega, \quad (1)$$

where $\partial\Omega := \bar{\Omega} \setminus \Omega$. To begin, however, we consider the one-dimensional problem, $-u_{xx} = f$, $u(\pm 1) = 0$.

2.1 One-Dimensional Case

The starting point for the spectral element discretization is the weak formulation of the one-dimensional model problem, which reads: Find $u \in H_0^1(\omega)$ such that

$$(u_x, v_x) = (f, v) \quad \forall v \in H_0^1(\omega), \quad (2)$$

where

$$\forall u, v \in L^2, \quad (u, v) = \int_{\omega} uv \, dx. \quad (3)$$

Here, $\omega := [-1, 1]$, $L^2 = \{v : \int_{\omega} v^2 \, dx < \infty\}$, $H^1 = \{v : v \in L^2(\omega), v_x \in L^2(\omega)\}$, and $H_0^1 = \{v : v(\pm 1) = 0, v \in H^1(\omega)\}$.

The spectral element discretization is based on a restriction of u and v to a finite-dimensional subspace. We introduce $X^N := Y^N \cap H^1$ and $X_0^N := Y^N \cap H_0^1$, where Y^N is




Figure 1: Local (top) and global (bottom) SE node numbering for $(E_x, N)=(3,4)$.

the space of piecewise discontinuous polynomials defined on elements, or subdomains, whose union constitutes a nonoverlapping partition of ω . For simplicity, we consider a uniform partition consisting of E_x elements, $\omega^e := [x^{e-1}, x^e]$, $e = 1, \dots, E_x$, with $x^e := eL_x - 1$ and $L_x := 2/E_x$. Functions in Y^N are represented in terms of Lagrange polynomials on each element

$$u(x)|_{\omega^e} = \sum_{j=0}^N u_j^e h_j(r^e), \quad (4)$$

where u_j^e are the nodal basis coefficients associated with ω^e and $r^e = 1 + 2(x - x^e)/L_x$ is an affine mapping from ω^e to the reference element $\hat{\omega} := [-1, 1]$. The Lagrangian basis functions satisfy $h_j(r) \in \mathbb{P}_N(\hat{\omega})$, $h_j(\xi_i) = \delta_{ij}$, where $\mathbb{P}_N(\hat{\omega})$ is the set of all polynomials of degree $\leq N$; $\xi_i \in [-1, 1]$, $i = 0, \dots, N$, correspond to the Gauss-Lobatto-Legendre (GLL) quadrature points; and δ_{ij} is the Kronecker delta. Derivatives in this basis are given by

$$\left. \frac{du}{dx} \right|_{x_i^e} = \frac{2}{L_x} \sum_{j=0}^N D_{ij} u_j^e, \quad (5)$$

where

$$D_{ij} = \left. \frac{dh_j}{dr} \right|_{\xi_i} \quad (6)$$

defines the elemental derivative matrix for $\xi_i \in \hat{\omega}$.

In addition to the choice of bases for X^N , the spectral element method is characterized by the substitution of quadrature for the inner products (\cdot, \cdot) in (2). For any $u, v \in Y^N$, we define

$$(u, v)_N := \frac{L_x}{2} \sum_{e=1}^{E_x} \sum_{j=0}^N \rho_j u_j^e v_j^e, \quad (7)$$

where ρ_j is the GLL quadrature weight. With these definitions, the spectral element discretization of (2) reads: Find $u \in X_0^N$ such that

$$(u_x, v_x)_N = (f, v)_N \quad \forall v \in X_0^N. \quad (8)$$

To generate the matrix form of (8), we introduce the local and global node numberings shown in Fig. 1. Sets of variables with both sub- and superscripts indicate local variables and are denoted by vectors with subscript L . Every function in X^N has an equivalent global node numbering, denoted by vector elements with a single subscript. For example, $\underline{u}_L = (u_0^1, u_1^1, \dots, u_N^{E_x})^T$ is the locally indexed set of nodal values, while $\underline{u} = (u_1, u_2, \dots, u_{n_x})^T$ is the globally indexed set. For the one-dimensional domain, the local to global mapping is given by

$$u_{q(i,e)} = u_i^e, \quad q(i,e) = 1 + i + N(e-1), \quad i = 0, \dots, N, \quad e = 1, \dots, E_x. \quad (9)$$

Note that, because of function continuity ($u \in X^N$), global indices have two local representations on the interfaces, that is, $q(N,e) = q(0,e+1)$. We introduce the Boolean matrix

$$Q_x^T := (\hat{e}_{q(i,e)}), \quad i = 0, \dots, N, \quad e = 1, \dots, E_x, \quad (10)$$

where \hat{e}_j is the j th column of the $n_x \times n_x$ identity matrix. Then the local to global mapping can be expressed in matrix form as $\underline{u}_L = Q_x \underline{u}$. Boundary conditions ($u \in X_0^N$) are enforced by introducing a restriction operator R_x , with $R_x^T := (\hat{e}_j)$, $j = 2, \dots, n_x - 1$. For any $\underline{u} \in \mathbb{R}^{n_x-2}$, there is a corresponding function $u \in X_0^N$ with local basis coefficients given by $\underline{u}_L = Q_x R_x^T \underline{u}$. With these definitions, the matrix form of (8) reads: Find $\underline{u} \in \mathbb{R}^{n_x-2}$ such that

$$\underline{v}^T R_x Q_x^T A_{x,L} Q_x R_x^T \underline{u} = \underline{v}^T R_x Q_x^T B_{x,L} \underline{f}_L, \quad \forall \underline{v} \in \mathbb{R}^{n_x-2}. \quad (11)$$

Here, the block-diagonal matrices $A_{x,L} = \text{diag}(A^e)$ and $B_{x,L} = \text{diag}(B^e)$ are respectively the unassembled stiffness and mass matrices, with entries

$$A_{ij}^e = \frac{2}{L_x} \sum_{p=0}^N D_{ip}^T \rho_p D_{pj}, \quad B_{ij}^e = \frac{L_x}{2} \rho_i \delta_{ij}, \quad i, j \in [0, \dots, N]^2, \quad e = 1, \dots, E_x. \quad (12)$$

We define the respective assembled stiffness and mass matrices as

$$A_x := R_x Q_x^T A_{x,L} Q_x R_x^T, \quad B_x := R_x Q_x^T B_{x,L} Q_x R_x^T. \quad (13)$$

Since A_x is symmetric positive definite (SPD), (12) is equivalent to solving the system

$$A_x \underline{u} = \underline{g},$$

where $\underline{g} := R_x Q_x^T B_{x,L} \underline{f}_L$.

In addition to the spectral element matrices (13), we construct finite element matrices based upon the space of piecewise linear Lagrangian basis functions, $\phi_i(x) \in H_0^1$,

$$\phi_i(x) = \begin{cases} \frac{x - x_{i-1}}{x_i - x_{i-1}} & x_{i-1} \leq x < x_i \\ \frac{x - x_{i+1}}{x_i - x_{i+1}} & x_i \leq x < x_{i+1} \\ 0 & \text{otherwise.} \end{cases} \quad i \in \{1, \dots, n_x - 1\} \quad (14)$$

The corresponding (tridiagonal) stiffness and (diagonal) mass matrices for the homogeneous Dirichlet problem are given by

$$\tilde{A}_{ij} = \int_{-1}^1 \frac{d\phi_i}{dx} \frac{d\phi_j}{dx} dx, \quad \tilde{B}_{ij} = \delta_{ij} \int_{-1}^1 \phi_j(x) dx, \quad i, j \in \{1, \dots, n_x - 1\}^2. \quad (15)$$

2.2 Two-Dimensional Case

The usual approach to generation of the governing SE system matrices in multiple space dimensions is to decompose the domain Ω into nonoverlapping quadrilateral or hexahedral elements and perform the necessary quadrature and differentiation on an element by element basis to yield a block-diagonal matrix comprising the element stiffness matrices, which are then assembled through application of a Boolean matrix Q defined by function continuity requirements, as in the preceding section. (See, e.g., [8, 13].) Here, we take a different approach in order to illustrate key features concerning the fast diagonalization method and the construction of tensor-product-based subdomain operators that are also applicable in more general situations (i.e., to problems that are not based on tensor-product arrays of elements).

Construction of the one-dimensional problem is extended to \mathbb{R}^d by taking the tensor product of the bases and operators described in Section 2.1. If the nodes are numbered lexicographically, then the stiffness matrix for the two-dimensional Laplacian on $\Omega := \omega_x \times \omega_y$ can be written as the Kronecker product:

$$A_s = B_y \otimes A_x + A_y \otimes B_x. \quad (16)$$

Here, the subscript y indicates a discretization on E_y elements in the y -direction similar to that derived for the x -direction. The finite element discretization has a similar form

$$A_f = \tilde{B}_y \otimes \tilde{A}_x + \tilde{A}_y \otimes \tilde{B}_x. \quad (17)$$

It was observed in [13] that, because A_s is based on diagonal mass matrices, A_f must be similarly constructed in order to be an effective preconditioner of A_s . As a result, A_f has a 5-point, rather than 9-point, stencil.

Matrices of the form (16) and (17) have a particularly simple inverse based upon the fast diagonalization method (FDM). For any $m \times m$ matrices A_* and B_* , with A_* symmetric and B_* SPD, the following similarity transformation holds:

$$S_*^T A_* S_* = \Lambda_*, \quad S_*^T B_* S_* = I, \quad (18)$$

where $\Lambda_* = \text{diag}(\lambda_1, \dots, \lambda_m)$ is the matrix of eigenvalues and $S_* = (\underline{s}_1, \dots, \underline{s}_m)$ is the matrix of eigenvectors associated with the generalized eigenvalue problem $A_* \underline{s} = \lambda B_* \underline{s}$. As a result, A_s is readily diagonalized by taking $*$ = x and y in (18), and its inverse is given by

$$A_s^{-1} = (S_y \otimes S_x) (I \otimes \Lambda_x + \Lambda_y \otimes I)^{-1} (S_y^T \otimes S_x^T). \quad (19)$$

The finite element form is similar, and the extension to three dimensions is straightforward. (See, e.g., [8].) This solution method was introduced by Lynch, Rice, and Thomas [20] and used in spectral element preconditioning applications by Couzy and Deville [6], Couzy [5], and Fischer et al. [15].

We make a few comments regarding the tensor product forms (16) and (17). Here, and in the sequel, we assume that $A \in \mathbb{R}^{n \times n}$ represents either A_s or A_f . Let $R_{i,x}^T$ be a rectangular matrix comprising a contiguous subset of the columns of the $(n_x - 2) \times (n_x - 2)$

identity matrix, and let $R_{j,y}^T$ be defined similarly in the y -direction. Let $R_k = R_{j,y} \otimes R_{i,x}$. Then $\underline{u}_k = R_k^T R_k \underline{u}$ defines a vector consisting of the nodal values of \underline{u} in the rectangular region $\tilde{\Omega}^k$ defined by the restriction matrices $(R_{i,x}, R_{j,y})$, and zero elsewhere. Moreover, $A_k := R_k A R_k^T$ is SPD because A is. By construction, A_k can be inverted by the FDM, and $\underline{u}_k = R_k^T A_k^{-1} R_k \underline{g}$ corresponds to the solution of the discretized Poisson problem, constrained to have zero nodal values outside of $\tilde{\Omega}^k$.

We note that the use of tensor-product forms allows matrix-vector products to be recast as *matrix-matrix* products, which are particularly efficient on modern vector and cache-based processors. For example, if $\underline{u} = u_{ij}$, is the vector of nodal basis coefficients on Ω , then

$$(S_y^T \otimes S_x^T) \underline{u} = S_x^T U S_y,$$

where U is simply \underline{u} viewed as the matrix having entries u_{ij} . For large systems, one might want to consider discretizations that allow fast (i.e., $O(n \log n)$) application of S_x and S_y , as considered by Shen [38]. However, for spectral-element-based subdomain problems, which typically range in size from 5×5 to 16×16 , matrix-matrix approaches are usually optimal.

In general, the FDM cannot be used for arbitrarily deformed subdomains because the discrete Laplacian cannot always be expressed in the tensor product form (16). For the purposes of a preconditioner, however, all one requires is an approximation to A (or A_k). As suggested by Couzy [5], a reasonable approach is to base the FDM on a discretization of Poisson's equation defined on a rectangle having the correct average dimensions in each coordinate direction in the reference domain. This approach was employed in three-dimensional Navier-Stokes applications in [15]. A similar preconditioning strategy for the case of spectral methods with nonconstant coefficients is discussed by Shen [38].

3 Iterative Solution Methods

We are interested in methods to solve the linear equation $A \underline{u} = \underline{g}$. To introduce notation, we consider the Richardson iteration

$$\underline{u}^{k+1} = \underline{u}^k + \sigma M (\underline{g} - A \underline{u}^k), \quad (20)$$

with preconditioner M and relaxation parameter σ . The iteration (20) can be combined with multigrid, in which case it serves the role of a smoother, or accelerated by Krylov-based projection methods such as conjugate gradients (if M is SPD) or GMRES [36]. One also can use *multigrid* as the preconditioner for conjugate gradients (CG) or GMRES. Several combinations of these approaches are considered in Section 6. In this section, we briefly describe our multigrid scheme, as well as the deflation method considered in [12, 33], and a new Schwarz method based on local, coarse, and strip solves.

3.1 Multigrid

The two-level multigrid method is defined as follows:

$$\begin{aligned}
i) \quad & \underline{u}^{k+1} = \underline{u}^k + \sigma_u M_d (\underline{g} - A \underline{u}^k), \quad k = 0, \dots, m_d - 1 \\
ii) \quad & \underline{r} = \underline{g} - A \underline{u}^{m_d} \\
iii) \quad & \underline{\tilde{e}} = J A_C^{-1} J^T \underline{r} \\
iv) \quad & \underline{\tilde{u}}^0 = \underline{u}^{m_d} + \underline{\tilde{e}} \\
v) \quad & \underline{\tilde{u}}^{k+1} = \underline{\tilde{u}}^k + \sigma_u M_u (\underline{g} - A \underline{\tilde{u}}^k), \quad k = 0, \dots, m_u - 1 \\
vi) \quad & \text{If } ||A \underline{\tilde{u}}^{m_u} - \underline{g}|| < \text{tol, set } \underline{u} := \underline{\tilde{u}}^{m_u}, \text{quit.} \\
& \text{Else, } \underline{u}^0 := \underline{\tilde{u}}^{m_u}, \text{ go to (i).}
\end{aligned} \tag{21}$$

The smoothing 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 J . Here, m_d and m_u are the number of smoothing steps on the downward and upward legs of the “V” cycle. In our study, we take $m_d = m_u = 1$, unless otherwise noted. Also, we take $M_d = M_u = M$ and $\sigma_d = \sigma_u = \sigma$, except in the local-coarse-strip (LCS) algorithm described below. We define σ to be either unity or, following Rønquist and Patera [35], the inverse of the maximum eigenvalue of MA (corresponding to maximum damping of the highest wavenumber), which is readily found by a power iteration. The coarse grid correction (Step *iii*) is characterized by the prolongation and restriction operators J and J^T , as well as the coarse-grid solve A_C^{-1} , which are defined below.

For select cases, we have also considered replacing A_C^{-1} with a call to a multigrid V-cycle on the coarser grid, to verify that the algorithms of interest can be nested. This approach is essential to the efficiency of the multigrid method as it ensures that the work on the coarse grid is bounded by the work on the fine grid.

3.2 Deflation

The deflation method was introduced by Nicolaides [26] and also studied by Mansfield [25] in the context of FE discretizations. Spectral element implementations of the deflation method in Navier-Stokes applications are presented in Rønquist [33] and Fischer [12]. Deflation employs an explicit projection of the solution onto coarse and fine subspaces, denoted by subscript C and F , respectively. The system $A \underline{u} = \underline{g}$ is split into the two equations,

$$A_F \underline{u}_F = (I - A J A_C^{-1} J^T) \underline{g} \tag{22}$$

$$A_C \underline{u}_C = J^T (\underline{g} - A \underline{u}_F), \tag{23}$$

such that $\underline{u} = \underline{u}_C + \underline{u}_F$, where $A_F = A - A J A_C^{-1} J^T A$. The column space of the prolongation operator J is intended to approximate the span of the low eigenmodes of the A system and is readily formed through intra-element interpolation as in the case of SEMG. The fine-space matrix A_F has a null space corresponding to the coarse space, so one has the additional constraint $J^T \underline{u}_F = 0$, which is readily enforced through local operations provided J is based on intra-element interpolation [12]. The fine-space equation is solved by using

preconditioned CG. Every calculation of a matrix-vector product involving A_F (one per CG iteration) requires two matrix-vector products in A and a coarse grid solve. Once \underline{u}_F is known, an additional coarse grid problem (23) is solved to compute \underline{u}_C . Because of the reliance on an exact decomposition of the coarse and fine spaces, it is imperative that $A_C = J^T A J$ and that A_C^{-1} be computed exactly. Deflation is the only method in which we use this form of the coarse-grid operator. In all other cases A_C is built by constructing a (FE or SE) discretization on the mesh of order $N_C < N$.

3.3 Local-Coarse-Strip Algorithm

The main difficulty of the methods presented thus far is in addressing the error on and near the interfaces. In the multigrid schemes, for instance, the high-aspect ratios of the cells near the element interfaces prevent the smoothers from being effective near and at the interfaces. The result, which is seen, in fact, in all methods, is that the error near the interfaces decays slowly. The local-coarse-strip (LCS) method was specifically devised to combat this problem. This method is a multiplicative combination of the preconditioners M_L , M_C , and $W M_{\text{strip}}$, which correspond to the execution of local solves, a coarse grid solve, and (weighted) strip solves, respectively. These will be discussed in detail later. The LCS scheme corresponds to the multigrid method (21) with $M_d = M_L$, $M_u = W M_{\text{strip}}$, and $m_d = m_u = 1$. In a more compact form, the LCS sweep can be written as

$$M_{\text{LCS}} = M_L + [M_C + W M_{\text{strip}}(I - A M_C)](I - A M_L). \quad (24)$$

Although M_L , M_C , and $W M_{\text{strip}}$ are themselves symmetric (at least for $W = I$), M_{LCS} is not and therefore cannot be used as a preconditioner for the conjugate gradient method. It can be used, however, in a scheme such as GMRES.

4 Preconditioners/Smoothers

In this section, we present several multigrid smoothing strategies that are used in the numerical tests of Section 6. Our original intent was to base the multigrid smoother, M , on the additive overlapping Schwarz method of Dryja and Widlund [11], with local subdomain problems discretized by finite elements having nodes coincident with the GLL nodes, as considered in [3, 13, 28]. Following Couzy and Deville [6] and Fischer et al. [15], the subdomain problems can be solved by using a fast diagonalization method similar to (19), guaranteeing that the cost of M is essentially equivalent to the cost of applying A . However, the use of the FDM frees one from the constraint of using FE-based preconditioners, since the FDM solution cost is dependent only on the use of the tensor-product form and not on the sparsity of the originating operator. Hence, one also can consider subdomain problems derived as restrictions of the original SE matrix, A_s , as first suggested by Casarin [3].

To establish a means for comparing the computational complexities, we estimate the cost of each step in terms of the cost of a matrix-vector products in A_s on the fine grid, which we define to have unit cost. For example, for the element-based Schwarz methods, we associate a cost of unity for M , whereas we attribute zero to the Jacobi-based preconditioner. Note that these costs do not account for communication overhead and would need to be

reevaluated in a parallel computing context. The storage complexity of the methods is essentially a small multiple of the number of grid points, so we assume parity in this measure.

Below, we briefly describe the preconditioning strategies considered in the next section. Each definition is based on a generic matrix A , which denotes either A_f or A_s , both of which are considered.

Jacobi

The Jacobi preconditioner is simply the diagonal matrix

$$(M_J)_{ij} = \delta_{ij} \frac{1}{A_{ij}}. \quad (25)$$

Cost: 0.

GSRB


We define the Gauss-Seidel red-black preconditioner as follows:

$$M_r = R_r^T (R_r A_f R_r^T)^{-1} R_r, \quad (26a)$$

$$M_b = R_b^T (R_b A_f R_b^T)^{-1} R_b, \quad (26b)$$

$$M_{\text{GSRB}} = M_r + M_b (I - A M_r), \quad (26c)$$

where R_r and R_b restrict to red and black nodes, respectively. Nodes are designated as red or black in a checkerboard pattern. No red node is the neighbor of another red node. M_{GSRB} has the same sparsity pattern as A_f . Note that this preconditioner is used only in the FE case. Cost: 1.



Local

$$M_{\text{local}} = \sum_{e=1}^E R_e^T (R_e A R_e^T)^{-1} R_e, \quad (27)$$

11

be used in conjunction with some other step, such as M_{strip} , which updates the interface nodes. Cost: 1.

Coarse

The coarse-solve system is given as

$$M_C = JA_C^{-1}J^T. \quad (28)$$

For all cases except deflation, A_C is computed as a spectral or finite element discretization, following (16) or (17), but using GLL points of order $N_C = 1$ or $N_C = N/2$. J interpolates from the coarse to the fine space, using either spectral or piecewise bilinear interpolation, depending on whether A_C is based on the spectral or finite element discretization. For deflation, we take $A_C = J^T A J$. Cost: 1 for $N_C = N/2$, 0 for $N_C = 1$.

Schwarz

The additive Schwarz preconditioner due to Dryja and Widlund [11] is

$$M_{\text{Schwarz}} = \sum_{e=1}^E \bar{R}_e^T (\bar{R}_e A \bar{R}_e^T)^{-1} \bar{R}_e. \quad (29)$$

Here, \bar{R}_e is a tensor-product-based restriction operator that restricts to those nodes that are less than l nodes away from element e , as illustrated in Fig. 2(b) for the case $l = 2$.

For the case of many elements, it is important to augment (29) with a coarse-grid problem [39]. This can be done either in an additive way, in which case one obtains the two-level additive Schwarz preconditioner (which we refer to as the additive Schwarz method)

$$M_{\text{A-Schwarz}} = M_C + M_{\text{Schwarz}},$$

or in a multiplicative way, in which case one obtains a hybrid Schwarz method given by the multigrid algorithm of Section 3.1, with $M = M_{\text{Schwarz}}$, $m_d = 1$ and $m_u = 0$ or 1. Cost of M_{Schwarz} : 1.

Strip

We define the strip smoother as

$$M_{\text{strip}} = \sum_i \tilde{R}_i^T (\tilde{R}_i A \tilde{R}_i^T)^{-1} \tilde{R}_i \quad (30)$$

where \tilde{R}_i restricts to the nodes of strip i . We let a strip be the collection of nodes on an interface together with nodes in either direction of the interface such that the strip width is $2l - 1$ nodes, as illustrated in Fig. 2(c). Cost: 1.

Table 1: Iterative methods and smoothers. Note that a semicolon in the preconditioner indicates a multiplicative combination. That is, “ P ; Q ” indicates $M = P + Q(I - AP)$. Only the weighted cases are shown. The nonweighted cases are obtained by setting $W = I$.

Preconditioner	Symmetric	Overall	Method
σM_J ; M_C	no	Richardson	MG, Jacobi, 1 smoothing
σM_{GSRB} ; M_C	no	Richardson	MG, GSRB, 1 smoothing
$\sigma M_{\text{GSRB}} + \alpha W M_{\text{strip}}$; M_C	no	Richardson	MG, GSRB w/ strip 1 smoothing
$W M_{\text{Schwarz}}$; M_C	no	Richardson	MG, H Schwarz, 1 smoothing
		GMRES	MG/GMRES, H Schwarz
M_J	yes	CG	PCG, Jacobi
		Deflation	Deflation, Jacobi
$M_{\text{GSRB}} + (M_{\text{GSRB}})^T$	yes	CG	PCG, GSRB
		Deflation	Deflation, GSRB
$M_{\text{GSRB}} + (M_{\text{GSRB}})^T + \alpha \sqrt{W} M_{\text{strip}} \sqrt{W}$	yes	CG	PCG, GSRB w/ strip
		Deflation	Deflation, GSRB w/ strip
$\sqrt{W} M_{\text{Schwarz}} \sqrt{W}$	yes	Deflation	Deflation, A Schwarz
$\sigma M_J + M_C$	yes	CG	PCG, Jacobi
$\frac{\sigma}{2} [M_{\text{GSRB}} + (M_{\text{GSRB}})^T] + M_C$	yes	CG	PCG, GSRB
$\frac{\sigma}{2} [M_{\text{GSRB}} + (M_{\text{GSRB}})^T] + \alpha \sqrt{W} M_{\text{strip}} \sqrt{W} + M_C$	yes	CG	PCG, GSRB w/ strip
$\sqrt{W} M_{\text{Schwarz}} \sqrt{W} + M_C$	yes	CG	PCG, A Schwarz
σM_J ; M_C ; σM_J	yes	Richardson	MG, Jacobi, 2 smoothings
		CG	MCG, Jacobi
$\frac{\sigma}{2} [M_{\text{GSRB}} + (M_{\text{GSRB}})^T]$; M_C ; $\frac{\sigma}{2} [M_{\text{GSRB}} + (M_{\text{GSRB}})^T]$	yes	CG	MCG, GSRB
	yes	CG	MCG, GSRB w/ strip
$\frac{\sigma}{2} [M_{\text{GSRB}} + (M_{\text{GSRB}})^T] + \alpha \sqrt{W} M_{\text{strip}} \sqrt{W}$	yes	CG	MCG, GSRB w/ strip
$\sqrt{W} M_{\text{Schwarz}} \sqrt{W}$; M_C ; $\sqrt{W} M_{\text{Schwarz}} \sqrt{W}$	yes	CG	MCG, H Schwarz
M_{local} ; M_C ; $W M_{\text{strip}}$	no	Richardson	LCS
		GMRES	LCS/GMRES
σM_{GSRB} ; M_C ; σM_{GSRB}	no	Richardson	MG, GSRB, 2 smoothings
$\sigma M_{\text{GSRB}} + \alpha W M_{\text{strip}}$; M_C ; $\sigma M_{\text{GSRB}} + \alpha W M_{\text{strip}}$	no	Richardson	MG, GSRB w/ strip 2 smoothings
$W M_{\text{Schwarz}}$; M_C ; $W M_{\text{Schwarz}}$	no	Richardson	MG, H Schwarz, 2 smoothings

Table 2: Per iteration work estimate.

Method	Total work	Mat-vec products	Smoother apps	Coarse grid grid work
MG, 2 smoothings	6	3	2	1
MG, 1 smoothing	4	2	1	1
PCG, $N_C = 0, 1$	2	1	1	0
PCG, $N_C = N/2$	3	1	1	1
Deflation	4	2	1	1
MGCG	6	3	2	1
LCS	6	3	2	1
LCS/GMRES	6	3	2	1

Diagonal weight matrix W

Both the strip and Schwarz preconditioners involve a sum of overlapping local solves. The diagonal counting matrix C given by

$$C = \sum_i R_i^T R_i \quad (31)$$

indicates the number of subdomains sharing a given vertex. In several cases considered, the additive overlapping solves are weighted with the diagonal weight matrix given by

$$W_{ii} = \begin{cases} 0 & \text{if } C_{ii} = 0 \\ \frac{1}{C_{ii}} & \text{otherwise.} \end{cases} \quad (32)$$


For CG applications, where M must be symmetric, we pre- and postmultiply M by \sqrt{W} . The distribution of the weights (for either M_{strip} or M_{Schwarz}) is illustrated in Fig. 2(d).

Table 4 summarizes the combinations of iterative methods and smoothers considered in Section 6. Additive coupling of the local and coarse solves is denoted by “+”. Multiplicative coupling is denoted by “;”. Table 2 gives the work-per-iteration estimates used in evaluating each of the methods. The estimates provided in the table are roughly the number of matrix-vector products in A_s that would correspond to an equivalent amount of work.


5 Error Behavior for Hybrid Schwarz and LCS

Figure 3 illustrates the error behavior for three variants of the hybrid Schwarz scheme for a 2×2 array of elements with $N = 16$.


Figure 3(a) shows the error after a single application of the smoothing step (20) with the (additive) Schwarz preconditioner (29) with minimal overlap as depicted in Fig. 2(b). While the error is clearly smooth in the element interiors, it is not smooth along the interfaces, and there is a definite spike at the cross-point. As seen in Fig. 3(b), subsequent application of the coarse-grid solve does little to remove the error along the interface—the rapid decay of the error in the interface-normal direction prevents it from being effectively represented on the coarse grid, and the overall convergence is slow.




(a) Application of M_{Schwarz}




(b) Coarse solve after (a)




(c) Application of $\sigma M_{\text{Schwarz}}$, $\sigma = 1/4$



(d) Coarse solve after (c)



(e) Application of WM_{Schwarz}



(f) Coarse solve after (e)

Figure 3: 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.

Figures 3(c) and (d) show the error when the relaxation parameter $\sigma = 1/4$ is introduced. (The maximum eigenvalue of

$$M_{\text{Schwarz}}A = \sum_{e=1}^E \bar{R}_e^T (\bar{R}_e A \bar{R}_e^T)^{-1} \bar{R}_e A =: \sum_{e=1}^E P_e$$


is bounded by E , since each $P_e := \bar{R}_e^T (\bar{R}_e A \bar{R}_e^T)^{-1} \bar{R}_e A$ is the A -orthogonal projector onto $\mathcal{R}(\bar{R}_e^T)$ and thus has eigenvalues of 0 and 1. By grouping the subdomains into sets such that no two subdomains within a given set overlap, the bound estimate can be improved to \mathcal{C} , where \mathcal{C} is the minimum number of disjoint subdomains required to cover all of Ω . In the present case, $\mathcal{C} = 4$, implying $\sigma = 1/4$.) After the smoothing step, the error in the element interiors (Fig. 3(c)) is not as smooth as in Fig. 3(a). It is, however, improved on the whole, and the coarse-grid step is able to make a more significant reduction, as seen in Fig. 3(d).

Figures 3(e) and (f) show the error when the Schwarz smoother is weighted by W (32). This weighting eliminates much of the high-frequency error on the interface observed in Fig. 3(a) and makes the coarse-grid correction much more effective. Comparing the magnitudes in Figs. 3(b) and 3(f) one sees, for this particular case, a tenfold reduction in the error through the introduction of W .


Figure 4 shows how each step of the LCS scheme affects the error. Starting with random error in a 2×2 element grid of order $N = 16$, applying M_L , the local solves, immediately smoothes the error in the interiors of the elements, as seen in Fig. 4(a). The coarse solve, M_C , then effectively eliminates the interior error, leaving the high-frequency modes on the interfaces, as Fig. 4(b) shows. These modes are significantly reduced by WM_S , the strip solves, as shown by Fig. 4(c). The tenfold reduction in the high-frequency components of the error near the interfaces depends on the presence of W in the strip-solve update. Without it, the strip solves would have to be damped by a relaxation factor σ_u and the error would be at most reduced by a factor of two, or less. This factor would then determine the overall convergence rate.

6 Results


We performed numerical experiments to test the described methods on the model problem using MATLAB version 6.5. The results are tabulated in Tables 3–8. For all runs, the domain was decomposed into $E = E_x \times E_y = 8 \times 8$ elements, which was found to be sufficiently large to exhibit asymptotic behavior. The source function was $2\pi^2 \sin(x) \sin(y)$, and the initial guess was chosen at random with values confined to the interval $[0, 1]$. Iterations continued until the norm of the error, determined by an initial exact solve, dropped below 1×10^{-11} . Total work estimates, based on the per iteration estimates of Table 2, are reported for the $N = 16$ case. Table 3 lists iteration counts for multigrid applied to the FE discretized model problem. Tables 4–8 list iteration and work counts for algorithms that solve the SE discretized model problem. The methods of Table 4 use FE-based preconditioners; those of Tables 5–8 use SE-based preconditioners. In all the tables, the Schwarz methods use a minimal overlap, $l = 2$ (1 node beyond the element interfaces), while the



(a) Local solves



(b) Coarse solve



(c) Strip solves

Figure 4: Error after the various stages of the LCS algorithm for $(E, N) = (4, 16)$.

Table 3: MG method on FE problem.

FE Spacing	#	Smoother/ Preconditioner	Coarse Space	Iterations			
				$N = 4$	$N = 8$	$N = 12$	$N = 16$
Uniform	a	Jacobi	$N/2$	39	38	38	38
	b	GSRB	$N/2$	9	9	9	9
	c	GSRB Strip	$N/2$	40	45	45	44
	d	H Schwarz	$N/2$	40	41	42	42
	e	H Schwarz (W)	$N/2$	7	7	6	6
SE	f	Jacobi	$N/2$	41	84	148	219
	g	GSRB	$N/2$	11	28	46	65
	h	GSRB Strip	$N/2$	34	23	25	29
	i	H Schwarz	$N/2$	40	43	47	52
	j	H Schwarz (W)	$N/2$	6	7	7	9

strip solves use $l = 3$. The strip solves are weighted with $\alpha = 0.2$ relative to the GSRB solves. In each table, the “weighted” designation indicates the use of W (32) or \sqrt{W} , as indicated in Table 4. The unweighted cases imply W is the identity matrix. For clarity, the table and cases throughout the discussion below are identified in parentheses.

The first set of experiments was designed to test the effect of grid spacing on the multi-grid method. Table 3 shows the results of various smoothing strategies applied to the FE problem, $A_f \underline{u} = \underline{g}$, with a uniform grid spacing and with the spacing given by a tensor product of the nodal spacing depicted in Fig. 1. For the uniform spacing, the GSRB algorithm (3b) produces the classic multigrid convergence rate—more than one digit per iteration. This rate is not attained for any of the other smoothers except for the weighted hybrid Schwarz scheme (3e). However, all the methods exhibit mesh resolution independence. For the nonuniform spacing based upon the SE nodal points, only the weighted hybrid Schwarz scheme (3j) is comparable to the best case for the uniform spacing. A slight degradation in performance is observable with increasing N . None of the methods attains mesh resolution independence.

Table 4 shows results for the problem $A_s \underline{u} = \underline{g}$ with preconditioners M and M_C derived from FE discretizations based on the SE nodal points. All of the results in this table are accelerated by CG, as straight multigrid did not converge for the parameters considered. One point of comparison with Table 3 is the weighted hybrid Schwarz method (4w), which has iteration counts ranging from 16 to 22 as N ranges from 4 to 16, in contrast to the values 6 to 9 given in (3j). This increase in iteration count is consistent with a condition number of $\pi^2/4$ for $A_f^{-1} A_s$, which would require approximately 16 CG iterations [17].

Another point of interest in Table 4 is additive Schwarz with bilinear coarse-grid space (4j), which corresponds to the state of the art as employed in [13, 15, 28, 30]. These results are in close agreement with those reported by Pahl [28] (Table 4.4). Comparing (4j) with the equivalent weighted method (4l), which employs \sqrt{W} as a pre- and post-multiplier of M_{Schwarz} , shows that a simple weighting provides roughly a 25 percent reduction in iteration count and work. Enriching the coarse grid to $N_C = N/2$ (4m) reduces the N dependence in the iteration count and yields an additional 18 percent reduction in work for the case $N = 16$, resulting in the lowest overall work for the FE-based preconditioning strategy.

Table 4: FE preconditioners for SE problem.

Method	#	Smoother/ Preconditioner	Coarse Space	Iterations				Work
				$N = 4$	$N = 8$	$N = 12$	$N = 16$	$N = 16$
PCG	a	Jacobi	–	132	> 300	> 300	> 300	> 600
	b	GSRB	–	102	221	> 300	> 300	> 600
	c	GSRB Strip	–	100	198	285	> 300	> 600
	d	Jacobi	1	52	116	187	264	528
	e	GSRB	1	34	73	114	159	318
	f	GSRB Strip	1	33	66	95	119	238
	g	Jacobi	$N/2$	29	40	50	61	183
	h	GSRB	$N/2$	20	26	32	37	117
	i	GSRB Strip	$N/2$	20	23	27	30	90
	j	A Schwarz	1	28	35	46	58	116
	k	A Schwarz	$N/2$	25	28	30	32	96
	l	A Schwarz (W)	1	21	26	36	46	92
	m	A Schwarz (W)	$N/2$	20	24	25	25	75
Deflation	n	Jacobi	$N/2$	23	35	47	57	228
	o	GSRB	$N/2$	20	25	32	37	148
	p	GSRB Strip	$N/2$	18	22	27	31	124
	q	A Schwarz	$N/2$	23	25	26	26	104
	r	A Schwarz (W)	$N/2$	21	22	22	22	88
MGCG	s	Jacobi	$N/2$	21	29	36	44	264
	t	GSRB	$N/2$	17	20	24	28	168
	u	GSRB Strip	$N/2$	23	23	23	24	144
	v	H Schwarz	$N/2$	18	19	19	20	120
	w	H Schwarz (W)	$N/2$	16	19	20	22	132

Table 5: SE preconditioners for SE problem.

Method	#	Smoother/ Preconditioner	Coarse Space	Iterations				Work
				$N = 4$	$N = 8$	$N = 12$	$N = 16$	$N = 16$
MG	a	Jacobi	$N/2$	56	107	183	275	1650
	b	H Schwarz	$N/2$	40	42	44	48	288
	c	H Schwarz (W)	$N/2$	9	15	17	18	108
PCG	d	Jacobi	–	130	> 300	> 300	> 300	> 600
	e	Jacobi	1	53	116	182	255	510
	f	Jacobi	$N/2$	32	45	57	69	207
	g	A Schwarz	1	25	27	35	43	86
	h	A Schwarz	$N/2$	26	26	26	27	81
	i	A Schwarz (W)	1	17	24	33	43	86
	j	A Schwarz (W)	$N/2$	16	21	22	24	72
Deflation	k	Jacobi	$N/2$	23	34	44	55	220
	l	A Schwarz	$N/2$	21	22	22	23	92
	m	A Schwarz (W)	$N/2$	12	12	13	14	56
MGCG	n	Jacobi	$N/2$	19	26	33	41	246
	o	H Schwarz	$N/2$	14	15	16	17	102
	p	H Schwarz (W)	$N/2$	10	14	16	16	96
LCS	q	Local/Strip	$N/2$	11	10	10	10	60
LCS/GMRES	r	Local/Strip	$N/2$	9	9	9	9	54

The weighted hybrid Schwarz (4w) and deflation (4r) schemes also have low iteration counts. It is not realistic to consider $N_C = N/2$ for deflation, however, since this enriched coarse-grid problem must be solved exactly and cannot be nested as would be the case for multigrid (= hybrid Schwarz). Surprising, the unweighted hybrid Schwarz scheme (4v) has slightly lower iteration counts than its weighted counterpart (4w). This is the only occurrence of this phenomenon in the tests considered.

Table 5 presents results for the spectral element problem with M and M_C based upon spectral element discretizations. The case (5a) corresponds to the Jacobi-based SEMG method of Maday et al. [22] and Rønquist [32], who found comparable work estimates. The Schwarz- and LCS-based smoothers are constructed by restricting the A_s matrix, as suggested by Casarin [3]. In every case, the iteration count is lower than the corresponding FE-preconditioned case in Table 4. By virtue of the FDM, the costs per iteration are equivalent, so it is clearly preferable to base the local Schwarz problems on the originating SE operators. It is also clear that using the diagonal weight matrix is beneficial, particularly in the unaccelerated multigrid scheme (5c).

The best method in Table 5 is the LCS/GMRES scheme (5r), with a work count of 54 for $N = 16$. Again, the deflation method (5m) has to be discounted because of the requirement for an exact coarse-grid solve of order $N_C = N/2$. Somewhat surprising, the simple additive Schwarz method with a bilinear coarse space (5g,i) is a strong contender, because of its low cost per iteration. In fact, for $N = 8$, the weighted SE-based variant (5i) has the lowest work estimate of all methods considered, 48 versus 54 for the next lowest scheme. The iteration count for $N = 8$ (16) was reduced reduced from 35 (58) for the state-of-the-art unweighted FE-based Schwarz method (4j) to 24 (43) for the weighted SE-based

local problems (5i).

Table 6 presents two-level multigrid results for the case of only one smoothing per iteration ($m_d = 1$, $m_u = 0$). This yields a nonsymmetric preconditioner, so one must switch from CG to GMRES. The best method, weighted hybrid Schwarz (6e), has a slightly higher iteration count than the LCS/GMRES method (5r), but the work estimates are comparable.

Table 7 illustrates the impact of the coarse-grid space on the performance of the hybrid Schwarz and LCS methods. The standard approach with Schwarz methods is to have the coarse-grid mesh size to be roughly the same as the size of the subdomains, corresponding to $N_C = 1$ in the present case. Using $N_C = N/2$ clearly improves the methods and yields nearly order-independent convergence rates. To verify that it is realistic to consider an enriched coarse space for these methods, we have also tested the fully nested multigrid V-cycle. For $N > 2$, the exact coarse-grid solve is replaced by a recursive call to the multigrid algorithm with $N \leftarrow N_C = N/2$. For $N = 2$, the coarse grid problem with $N_C = 1$ is solved exactly. The results are reported in Table 8. In all cases except the unaccelerated hybrid Schwarz method (8a), the iteration count increases by 1 over the corresponding case in Table 7, where exact coarse-grid solves are employed.

7 Conclusions

We have studied several combinations of FE-, Schwarz-, and MG-based preconditioning strategies for spectral element solution of the Poisson equation in \mathbb{R}^2 . All of the methods involved two or more levels of discretization, with the coarser approximation spaces based on the original $E = E_x \times E_y$ spectral element decomposition of Ω , such that intra-element prolongation/restriction could be employed. All methods demonstrated E independent convergence rates for $E_x = E_y \geq 8$, so only this case was considered. Values of N ranged from 4 to 16.

The fastest schemes were found to be the LCS (5q), LCS/GMRES (5r), deflation (5m), and weighted hybrid Schwarz (6e). Practical implementation of deflation is limited because of its need for an exact coarse-grid solve, which is very expensive in the case of $N_C = N/2$. As constructed in this study, the Schwarz scheme (5e) requires a global tensor-product mesh topology because construction of the overlap assumed that elements were (locally) arrayed in a tensor-product fashion. Strategies to circumvent this constraint have been proposed in [15]. The local solves in the LCS scheme involve no overlap and can be done very quickly, particularly in a parallel context where no communication is required. The strip solves do

Table 6: SE preconditioners for SE problem, MG with 1 smoothing.

Method	#	Smoother/ Preconditioner	Coarse Space	Iterations				Work
				$N = 4$	$N = 8$	$N = 12$	$N = 16$	$N = 16$
MG	a	Jacobi	$N/2$	111	215	> 300	> 300	> 1200
	b	H Schwarz	$N/2$	78	83	87	95	380
	c	H Schwarz (W)	$N/2$	16	17	18	19	76
MG/GMRES	d	H Schwarz	$N/2$	21	23	24	25	100
	e	H Schwarz (W)	$N/2$	13	12	12	13	52

Table 7: SE preconditioners for SE problem, MG with 1 smoothing.

Method	#	Smoother/ Preconditioner	Coarse Space	Iterations				Work
				$N = 4$	$N = 8$	$N = 12$	$N = 16$	$N = 16$
MG	a	H Schwarz (W)	1	16	98	249	> 300	> 900
	b	H Schwarz (W)	$N/2$	16	17	18	19	76
MG/GMRES	c	H Schwarz (W)	1	14	20	29	36	108
	d	H Schwarz (W)	$N/2$	13	12	12	13	52
LCS	e	Local/Strip	1	21	27	49	82	410
	f	Local/Strip	$N/2$	11	10	10	10	60
LCS/GMRES	g	Local/Strip	1	13	15	19	24	120
	h	Local/Strip	$N/2$	9	9	9	9	54

require overlap. The important advantage of the strip solves over the overlapping solves of the Schwarz method is that there is no complication with geometries where three or five or more elements, instead of four, join at a vertex. Strip solves rely only on the condition of conformity that two elements join at an interface. The straightforward extension to three dimensions, however, is less clear. For example, one may need an edge solve in addition to the local, coarse, and face solves. One additional feature of the LCS scheme is that it is the only method that proved effective in the straight multigrid V-cycle (8c), without the need for GMRES acceleration. This feature is particularly attractive in large-scale parallel settings because it avoids the communication overhead associated with the inner-product computations required for CG and GMRES.

The additive Schwarz PCG methods had work counts in the range 72–86. With a rich coarse space of $N/2$, the work count was 72; for a coarse space of order 1, the count was 86. Here, one should consider that the gain in going to a richer coarse space is offset by the added complexity and that the iteration count of the rich coarse-space method will degrade, to at least some degree, when the coarse solve is replaced by nesting the method. The improvement in using a richer coarse space may not be dramatic enough to warrant its implementation.

In general, basing the Schwarz operators on restrictions of the originating SE matrices, as suggested by Casarin [3], was found to be superior to using FE-based discretizations. This is made possible through the use of the fast diagonalization method, which can be applied to tensor-product subdomains at a cost that is independent of the sparsity of the governing subdomain system and that is essentially equal to the cost of matrix-vector products in A_s .

We found that weighting the overlapping Schwarz updates by the inverse of the count-

Table 8: Fully nested SE preconditioners for SE problem, MG with 1 smoothing.

Method	#	Smoother/ Preconditioner	Coarse Space	Iterations	Work
				$N = 16$	$N = 16$
MG	a	H Schwarz (W)	$N/2$	32	128
MG/GMRES	b	H Schwarz (W)	$N/2$	14	56
LCS	c	Local/Strip	$N/2$	11	66
LCS/GMRES	d	Local/Strip	$N/2$	10	60

ing matrix C (31) plays a crucial role in obtaining multigrid-like iteration counts.

Finally, extensions to high-aspect-ratio subdomains, three dimensions, and more complex systems, such as the consistent Poisson operator that governs the pressure in the $\mathbb{P}_N - \mathbb{P}_{N-2}$ SE discretization of the Navier-Stokes equations [13, 15], have yet to be considered. We hope that this study provides a useful starting point and some guiding principles for the development of algorithms that address these practical issues.

Acknowledgments

We thank Barry Smith for many useful discussions during the course of this work. This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-Eng-38. We also thank the Division of Educational Programs at Argonne National Laboratory, which was instrumental in enabling the collaboration that resulted in this study.

References

- [1] S. Beuchler. Multigrid solver for the inner problem in domain decomposition methods for p-fem. *SIAM J. Numer. Anal.*, 40:928–944, 2002.
- [2] M.A. Casarin. Diagonal edge preconditioners in p-version and spectral element methods. *SIAM J. Sci. Comput.*, 18:610–620, 1997.
- [3] M.A. Casarin. Quasi-optimal Schwarz methods for the conforming spectral element discretization. *SIAM J. Numer. Anal.*, 34:2482–2502, 1997.
- [4] M.A. Casarin. Schwarz preconditioners for the spectral element Stokes and Navier-Stokes discretizations. In P. Bjørstad, M. Espedal, and D. Keyes, editors, *Domain Decomposition 9 Proceedings*, pages 72–79, New-York, 1998. J. Wiley.
- [5] 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.
- [6] W. Couzy and M.O. Deville. Spectral-element preconditioners for the Uzawa pressure operator applied to incompressible flows. *J. Sci. Comput.*, 9:107–112, 1994.
- [7] W. Couzy and M.O. Deville. A fast Schur complement method for the spectral element discretization of the incompressible Navier-Stokes equations. *J. Comput. Phys.*, 116:135–142, 1995.
- [8] M.O. Deville, P.F. Fischer, and E.H. Mund. *High-Order methods for incompressible fluid flow*. Cambridge University Press, Cambridge, 2002.

- [9] M.O. Deville and E.H. Mund. Chebyshev pseudospectral solution of second-order elliptic equations with finite element preconditioning. *J. Comp. Phys.*, 60:517–533, 1985.
- [10] M.O. Deville and E.H. Mund. Finite element preconditioning for pseudospectral solutions of elliptic problems. *SIAM J. Sci. Stat. Comput.*, 11:311–342, 1990.
- [11] 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.
- [12] P.F. Fischer. Parallel multi-level solvers for spectral element methods. In A.V. Ilin and L.R. Scott, editors, *Third International Conference on Spectral and High Order Methods*, pages 595–604. Houston Journal of Mathematics, 1996.
- [13] P.F. Fischer. An overlapping Schwarz method for spectral element solution of the incompressible Navier-Stokes equations. *J. Comput. Phys.*, 133:84–101, 1997.
- [14] P.F. Fischer, G.W. Kruse, and F. Loth. Spectral element methods for transitional flows in complex geometries. *J. Sci. Comput.*, 17:81–98, 2002.
- [15] 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.
- [16] P.F. Fischer and E.M. Rønquist. Spectral element methods for large scale parallel Navier-Stokes calculations. *Comput. Methods Appl. Mech. Engrg.*, 116:69–76, 1994.
- [17] G. Golub and C.F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, 1996.
- [18] W. Heinrichs. Line relaxation for spectral multigrid methods. *J. Comput. Phys.*, 77:166–182, 1988.
- [19] W. Heinrichs. Improved condition number for spectral methods. *Math. Comp.*, 53:103–119, 1989.
- [20] 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.
- [21] Y. Maday and R. Muñoz. Numerical analysis of a multigrid method for spectral approximations. In M. Y. Hussaini D. L. Dwoyer and R. G. Voigt, editors, *Lecture Notes in Physics, Volume 323: Proc. of the 11th Int. Conf. on Numerical Methods in Fluid Dynamics*, pages 389–394. Springer, 1989.
- [22] Y. Maday, R. Muñoz, A.T. Patera, and E.M. Rønquist. Spectral element multigrid methods. In P. de Groen and R. Beauwens, editors, *Proceedings of the IMACS International Symposium on Iterative Methods in Linear Algebra, Brussels, 1991*, pages 191–201, Amsterdam, 1992. Elsevier.

- [23] J. Mandel. Two-level domain decomposition preconditioning for the p -version finite element method in three dimensions. *Int. J. Numer. Methods Eng.*, 29:1095–1108, 1990.
- [24] J. Mandel and G.S. Lett. Domain decomposition preconditioning for p -version finite elements with high aspect ratios. *Appl. Numer. Math.*, 8:411–425, 1991.
- [25] L. Mansfield. On the use of deflation to improve the convergence of conjugate gradient iteration. *Comm. in Appl. Numer. Meth.*, 4:151–156, 1988.
- [26] R.A. Nicolaides. Deflation of conjugate gradients with application to boundary value problems. *SIAM J. Numer. Anal.*, 24:355–365, 1987.
- [27] S.A. Orszag. Spectral methods for problems in complex geometry. *J. Comput. Phys.*, 37:70–92, 1980.
- [28] S.S. Pahl. *Schwarz Type Domain Decomposition Methods for Spectral Element Discretizations*. PhD thesis, Univ. of Witwatersrand, Johannesburg, South Africa, 1993. Master’s Thesis, Dept. of Computational and Applied Math.
- [29] S.V. Parter and E.E. Rothman. Preconditioning Legendre spectral collocation approximations to elliptic problems. *SIAM J. Numer. Anal.*, 32:333–385, 1995.
- [30] L.F. Pavarino and T. Warburton. Overlapping Schwarz methods for unstructured spectral elements. *J. Comput. Phys.*, 160:298–317, 2000.
- [31] L.F. Pavarino and O.B. Widlund. A polylogarithmic bound for an iterative substructuring method for spectral elements in three dimensions. *SIAM J. Numer. Anal.*, 33:1303–1335, 1996.
- [32] E. Rønquist. *Optimal Spectral Element Methods for the Unsteady Three-Dimensional Incompressible Navier-Stokes Equations*. PhD thesis, Massachusetts Institute of Technology, 1988. Cambridge, MA.
- [33] E.M. Rønquist. A domain decomposition method for elliptic boundary value problems: Application to unsteady incompressible fluid flow. In D.E. Keyes, T.F. Chan, G. Meurant, J.S. Scroggs, and R.G. Voigt, editors, *Fifth Conference on domain decomposition methods for partial equations*, pages 545–557. SIAM, 1992.
- [34] E.M. Rønquist. A domain decomposition solver for the steady Navier-Stokes equations. In A.V. Ilin and L.R. Scott, editors, *Third International Conference on Spectral and High Order Methods*, pages 469–485. Houston Journal of Mathematics, 1996.
- [35] E.M. Rønquist and A.T. Patera. Spectral element multigrid. I. formulation and numerical results. *J. Sci. Comput.*, 2:389–406, 1987.
- [36] Y. Saad and M.H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7:856–869, 1986.

- [37] J. E. Shen, F. Wang, and J. Xu. A finite element multigrid preconditioner for chebyshev-collocation methods. *Appl. Numer. Math.*, 33:471–477, 2000.
- [38] J.E. Shen. Efficient Chebyshev-Legendre Galerkin methods for elliptic problems. In A.V. Ilin and L.R. Scott, editors, *Third International Conference on Spectral and High Order Methods*, pages 233–239. Houston Journal of Mathematics, 1996.
- [39] B. Smith, P. Bjørstad, and W. Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic PDEs*. Cambridge University Press, Cambridge, 1996.
- [40] S.J. Thomas, J. M. Dennis, H. M. Tufo, and P. F. Fischer. A Schwarz preconditioner for the cubed-sphere. *SIAM J. Sci. Comput.*, 00, (to appear).
- [41] T.A. Zang, Y.S. Wong, and M.Y. Hussaini. Spectral multigrid methods for elliptic equations. *J. Comput. Phys.*, 48:485–501, 1982.
- [42] T.A. Zang, Y.S. Wong, and M.Y. Hussaini. Spectral multigrid methods for elliptic equations II. *J. Comput. Phys.*, 54:489–507, 1984.

<p>The submitted manuscript has been created by the University of Chicago as Operator of Argonne National Laboratory ("Argonne") under Contract No. W-31-109-ENG-38 with the U.S. Department of Energy. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.</p>
--