

33-64
N92-11731
4/6/00

Domain-Decomposed Preconditionings for
Transport Operators†

p-19

Tony F. Chan* - CD146017
William D. Gropp** - AX337235
and
David E. Keyes*** - YE661058

Abstract

We test the performance of five different interface preconditionings for domain-decomposed convection-diffusion problems, including a novel one known as the spectral probe, while varying mesh parameter, Reynolds number, ratio of subdomain diffusion coefficients, and domain aspect ratio. The preconditioners are representative of the range of practically computable possibilities that have appeared in the domain decomposition literature for the treatment of nonoverlapping subdomains. We demonstrate through a large number of numerical examples that no single preconditioner can be considered uniformly superior or uniformly inferior to the rest, but that knowledge of particulars, including the shape and strength of the convection, is important in selecting among them in a given problem.

1. Introduction

The solution of linearized convection-diffusion equations of the form

$$\vec{c} \cdot \nabla \phi - \nabla \cdot \epsilon \nabla \phi = f, \quad (1.1)$$

where ϕ is a conserved quantity (energy, mass fraction, momentum component, etc.) transported under the influence of velocity field \vec{c} and diffusivity ϵ is required throughout computational physics. Discretization by finite differences or finite elements results in a large sparse system of algebraic equations whose solution can be demanding in computational resources and is one of the many driving forces for parallel computation. Because the strength of coupling between a pair of discrete unknowns governed by an equation like (1.1) decays with physical separation (more or less isotropically depending upon \vec{c}), it is natural to partition the problem spatially when looking for concurrency in the solution algorithm. Parallelism is, in fact, only one of several compelling reasons for the recent surge of research on domain decomposition algorithms exemplified by the volumes [9, 10, 17]. Others include a powerful theory describing optimal or near-optimal algebraic convergence rates for hierarchical preconditioners of domain-decomposed type, the convenience of composite array data structures for describing complex shapes, a desire to employ solution techniques and quality

†This report is an augmentation of a report entitled *Interface Preconditioning for Domain-Decomposed Convection-Diffusion Operators* by the first and third authors that appeared in "Proceedings of the Third International Symposium on Domain Decomposition Methods," T. F. Chan, R. Glowinski, J. Periaux, and O. B. Widlund, eds., pp. 245-262, SIAM, 1990.

*Department of Mathematics, UCLA, Los Angeles, CA 90024. The work of this author was supported in part by NSF ASC 9003002, DOE DE-FG03-87ER25037, and ARO DAAL03-88-K-0085.

**Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439. The work of this author was supported in part by the Applied Mathematical Sciences subprogram of the Office of Energy Research, U.S. Department of Energy, under Contract W-31-109-ENG-38.

***Department of Mechanical Engineering, Yale University, New Haven, CT 06520. The work of this author was supported in part by NSF ECS-8957475 and by the National Aeronautics and Space Administration under NASA Contract NAS1-18605, while the author was in residence at the Institute for Computer Applications in Science and Engineering, NASA-Langley Research Center, Hampton, VA 23665.

software restricted to problems with various local uniformity requirements (which are subproblems with regard to (1.1)), and sheer problem size, which can ultimately push numerical ill-conditioning and serial memory traffic beyond acceptable limits.

Preconditionings for interfacial degrees of freedom have been the focus of much attention during the development of domain decomposition methods, and deservedly so, since interfaces of lower dimension than the original domain of definition of the partial differential equation are created by a predominant form of nonoverlapping decomposition related to nested dissection of the underlying finite difference or finite element matrix operator. We refer generically to such forms of domain decomposition as Schur iteration, since elimination of the subdomain interiors leaves a Schur complement system for the separator unknowns. Additional interest in interface preconditioning comes from the fact that the classical Schwarz iteration, the prototype for overlapping decompositions, has been placed into correspondence with a stationary iteration having as unknowns the interfacial degrees of freedom of a nonoverlapping decomposition [6, 11]. This correspondence between Schwarz and Schur methods enriches the study of domain decomposition algorithms in general, because properties which are more easily analyzed in one framework may be extended to the other.

The present contribution focuses on the performance of a variety of easily computed Schur complement preconditioners in a rather special context: a single interface dividing a rectangle into two subrectangles in which the capability of performing exact solves is presumed. We consider a scalar convection-diffusion operator under a uniform or "terraced" diffusion coefficient and a variety of continuity-satisfying flow fields chosen to exhibit the relative advantages and disadvantages of the preconditioners. The pristine nature of the problem class allows focusing on the quality of the interfacial preconditioning, alone, in four different limits: large discrete problem size, large Reynolds (or Peclet) number, large diffusion coefficient ratio, and large aspect ratio. (The Reynolds number is the dimensionless ratio $\epsilon l / \bar{\epsilon}$, where $\bar{\epsilon}$ is a characteristic velocity, l a characteristic length, and $\bar{\epsilon}$ a characteristic diffusivity. Large values characterize strongly nonsymmetric, convectively-dominated systems.) Any or all of these limits could be important in a production engineering code whose parallelization might be sought through domain decomposition. We show that no single interface preconditioner is best in all limits, and therefore seek to qualitatively rank their sensitivities to these limits and identify realms of superiority.

Several different coefficient fields $\bar{\epsilon}$ and ϵ are studied because the performance of all of the preconditioners are sensitive to them and unjustified optimism or pessimism can result from too narrow a study. Two of our five preconditioners have been amply studied previously in the symmetric positive definite context of pure diffusion. There have been very few studies of any of them in the convection-diffusion context, and since this case is also relatively untouched by theoretical approaches, apart from spatially invariant velocity distributions, numerical studies are continuing to yield interesting information.

We comment briefly on a few other issues which bear on our choice of scope. It is possible to set up an alternative framework for nonoverlapping decompositions in which interfacial coupling is simply discarded, or partially accounted for in ways that do not require special treatment of a separator set; see, *e.g.*, [1] and [26]. In so doing one obtains the advantages of greatly simplified coding and less inter-domain data traffic per iteration. Problems dominated by local interactions can be handled quite acceptably by decoupling; see *e.g.*, [23]. However, in problems which are diffusively dominated (more fundamentally, problems whose Green's functions have support which is not substantially confined within artificial subdomain boundaries), such approaches have limited applicability to large numbers of gridpoints and/or subdomains.

The special case of a single interface obviates discussion of preconditioning the set of vertices where multiple interfaces intersect. Vertex preconditioning is very important but also more readily prescribable, at least in two dimensions. A coarse grid problem for the vertices having the same

structure as the undecomposed original problem can be derived directly from the differential operator by employing a hierarchical basis discretization. The interface system, on the other hand, corresponds to a pseudo-differential operator, the numerical analysis of which is relatively less well developed in the presence of convective terms. In a preconditioner consisting of component blocks corresponding to subdomains, vertices, and interfacial edges (and interfacial planes in three dimensions), any one block can limit the overall performance. A study of interface preconditioning is thus necessary, but not sufficient, for guiding the construction of complete preconditioners.

Finally, as to the relevance of our scope, we note that practical problems often involve several simultaneous convection-diffusion operators linked through coefficients or source terms. Continued study of the scalar case is, however well motivated by techniques such as the alternating block factorization [4] which successfully employ scalar preconditioners *inside* of a change of dependent variables which partially decouples the original system.

The algorithmic framework of our experiments is described in Section 2, followed by introduction of the five interface preconditioners and a brief discussion of their properties in Section 3. Section 4 contains performance measurements in the form of iteration counts along several axes of problem parameter space. We draw some conclusions and recommendations in Section 5.

2. Schur Domain Decomposition Methods

We take as our starting point the matrix equation $Ax = b$ arising from a finite difference discretization of (1.1). The domain decomposition method we employ is an iterative substructuring method consisting of three elements: (1) the operator A whose inverse action we would like to compute with an accuracy commensurate with the discretization, (2) an approximation B to A , whose inverse action is computationally convenient to compute, and (3) an acceleration scheme for the preconditioned system which requires only the ability to form the actions of A and B^{-1} on a vector. In all cases reported herein, A is derived from a second-order central differencing of the diffusion term and a first-order upwind differencing of the convection term. Extensions to second-order upwind differencing have been carried out in, for instance, [27]. We use right-preconditioned GMRES [30] as our iterative acceleration scheme, that is, we solve $AB^{-1}y = b$ by applying the standard GMRES algorithm to (AB^{-1}) then recover x through the solution of $Bx = y$.

GMRES is guaranteed to converge in a finite number of steps for nonsingular AB^{-1} even in the presence of nonsymmetry or indefiniteness, assuming exact arithmetic. The maximum number of steps required is the number of distinct eigenvalues of the preconditioned operator. This convergence result depends upon dynamically storing a complete basis for the Krylov space built from powers of AB^{-1} acting on the initial residual vector. For large problems, this much memory can easily become excessive, and GMRES is often truncated or restarted [30] in cases where it does not converge within a predetermined number of steps. However, we allow full GMRES iteration in our experiments, up to some maximum number of steps (set at 30 herein) which is sufficient in all but two cases. Because fewer than 30 steps are almost always sufficient, we effectively suppress from consideration the restart or truncation parameter. This parameter can be important in a "production" setting.

The substructuring enters through the manipulation of A and B into forms which possess large block zeros, for the sake of concurrency or for some of the other reasons noted in the introduction. For elliptic operators such as (1.1), A is irreducible; hence there are no block triangular permutations. However, if the domain is cut by the removal of a swath of gridpoints as wide as the semi-bandwidth of the stencil, two large subproblems are created whose only coupling is through the small removed set. For five-point stencils on logically tensor product grids, we may choose a single row or column of unknowns. (A two-point-wide generalization has been studied for the thirteen-point biharmonic stencil in [8].) Ordering the separators last, we obtain

$$Ax \equiv \begin{pmatrix} A_{11} & 0 & A_{13} \\ 0 & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} \equiv b. \quad (2.1)$$

Here, A_{11} and A_{22} are five-point operators with bandwidth no larger than that of the naturally ordered original system, but A_{33} , which renders the coupling between the points on the interface itself, is tridiagonal. The other blocks contain the coupling of the separator unknowns to the subdomains, and vice versa. From the point of view of the continuous operator they represent derivatives in directions normal to the interface.

Block Gaussian elimination of the unknowns x_1 and x_2 would yield the Schur complement system

$$Cx_3 = d \quad (2.2)$$

for x_3 , where

$$C = A_{33} - A_{31}A_{11}^{-1}A_{13} - A_{32}A_{22}^{-1}A_{23} \quad (2.3)$$

and

$$d = b_3 - A_{31}A_{11}^{-1}b_1 - A_{32}A_{22}^{-1}b_2. \quad (2.4)$$

If x_3 can be found, the subdomain problems are decoupled. However, direct computation of the generally dense C in order to solve (2.2) requires as many pairs of exact subdomain solves as there are degrees of freedom in x_3 , which is generally prohibitive. It is also unnecessary inasmuch as iterative techniques have been devised which require many fewer iterations than the dimension of x_3 , and which furthermore require only approximate subdomain solves in each iteration. As mentioned already, we shall ignore the option of inexact subdomain solves in the sequel, effectively reducing the iterations to the interface, but we nevertheless make use of a general purpose code which retains the interior degrees of freedom in carrying out the numerical experiments.

We consider two families of preconditioners B , the structurally symmetric

$$\begin{aligned} B_1 &= \begin{pmatrix} A_{11} & 0 & 0 \\ 0 & A_{22} & 0 \\ A_{31} & A_{32} & M \end{pmatrix} \begin{pmatrix} I & 0 & A_{11}^{-1}A_{13} \\ 0 & I & A_{22}^{-1}A_{23} \\ 0 & 0 & I \end{pmatrix} \\ &= \begin{pmatrix} A_{11} & 0 & A_{13} \\ 0 & A_{22} & A_{23} \\ A_{31} & A_{32} & M + A_{31}A_{11}^{-1}A_{13} + A_{32}A_{22}^{-1}A_{23} \end{pmatrix}, \end{aligned}$$

where M approximates the Schur complement C (2.3) of A_{11} and A_{22} in A , and the simpler block triangular

$$B_2 = \begin{pmatrix} A_{11} & 0 & A_{13} \\ 0 & A_{22} & A_{23} \\ 0 & 0 & M \end{pmatrix}.$$

The factorized form of B_1 above shows that the cost of applying the inverse of B_1 is one solve with M and two solves each with A_{11} and A_{22} . There is an inherent sequentiality to the subdomain solves, however, since the system involving M in the left factor requires data from the first set of subdomain solves. The inverse of B_2 can be applied to a vector at the cost of solving one system each with M , A_{11} , and A_{22} . The system for M is solved first, followed by independent solves in the subdomains which use the interface values as boundary conditions.

We assume throughout that the A_{ii} are invertible. (This is certainly a reasonable requirement for a discrete convective-diffusive operator and is guaranteed herein for all Reynolds numbers by upwind differencing.) Under this assumption, C is also invertible [15].

For matrices arising from standard quasi-uniform finite element discretizations of elliptic partial differential equations, A has a condition number of $\mathcal{O}(h^{-2})$, whereas C has a condition number of $\mathcal{O}(h^{-1})$ [5, 29]. The equivalence of conjugate gradient iterations on the Schur complement system with preconditioner M and on the full *substructured* matrix A with preconditioner B_1 was shown in [24].

For reference in Section 4, it is interesting to note the forms of the preconditioned operators AB_1^{-1} , and AB_2^{-1} . In order to make the formulae more readable, we combine the independent subdomain solves into a block matrix A_Ω , and denote the separator block by A_Γ , to re-express the above matrices as

$$A = \begin{pmatrix} A_\Omega & A_{\Omega\Gamma} \\ A_{\Gamma\Omega} & A_\Gamma \end{pmatrix}, \quad B_1 = \begin{pmatrix} A_\Omega & A_{\Omega\Gamma} \\ A_{\Gamma\Omega} & M + A_{\Gamma\Omega}A_\Omega^{-1}A_{\Omega\Gamma} \end{pmatrix}, \quad B_2 = \begin{pmatrix} A_\Omega & A_{\Omega\Gamma} \\ 0 & M \end{pmatrix},$$

whence

$$B_1^{-1} = \begin{pmatrix} A_\Omega^{-1} + A_\Omega^{-1}A_{\Omega\Gamma}M^{-1}A_{\Gamma\Omega}A_\Omega^{-1} & -A_\Omega^{-1}A_{\Omega\Gamma}M^{-1} \\ -M^{-1}A_{\Gamma\Omega}A_\Omega^{-1} & M^{-1} \end{pmatrix}$$

and

$$B_2^{-1} = \begin{pmatrix} A_\Omega^{-1} & -A_\Omega^{-1}A_{\Omega\Gamma}M^{-1} \\ 0 & M^{-1} \end{pmatrix}.$$

From these expressions it can easily be verified that

$$AB_1^{-1} = \begin{pmatrix} I & 0 \\ (I - CM^{-1})A_{\Gamma\Omega}A_\Omega^{-1} & CM^{-1} \end{pmatrix} \quad \text{and} \quad AB_2^{-1} = \begin{pmatrix} I & 0 \\ A_{\Gamma\Omega}A_\Omega^{-1} & CM^{-1} \end{pmatrix}. \quad (2.5)$$

It is evident that if C is exactly represented by M , then AB_1^{-1} reduces to the identity, and an iteration involving AB_1^{-1} will converge in one step requiring two sets of subdomain solves. An iteration involving AB_2^{-1} , on the other hand, will converge in two steps (from an arbitrary initial guess) if $M = C$, but each step requires only one set of subdomain solves. (These iteration counts do not include the final solve with either B_1 or B_2 which is required to unwind the right-preconditioning.) More generally, if M is sufficiently close to C in the sense that the lower-left block of the structurally symmetric system is small, $\|(I - CM^{-1})\| \ll 1$, we expect that an iteration based on B_2 will require an extra iteration relative to an iteration based on B_1 . Conversely, if M is a poor preconditioner for C , so that the lower left block becomes large, the use of the structurally symmetric system could require more iterations than the use of the block triangular system. Both behaviors are illustrated in Section 4.

Note from (2.5) that AB_1^{-1} and AB_2^{-1} have identical spectra, as Arnoldi estimates for the eigenvalues obtained as a by-product of the GMRES iterations also show. However, Krylov sequences based on the respective operators will in general differ, and there is little that can be said about which method will lead to faster convergence for general \vec{c} if M and C are not sufficiently close.

For some of the preconditioner components M we consider, the overall preconditioning process is numerical unstable, as will be seen in Section 4. Even though the iterations involving AB^{-1} may converge, the final result after unwinding the preconditioning may have few or even no significant figures. For this reason, we always check the actual residual $\|f - Ax\|$ at the end of each calculation.

3. Schur Interface Preconditioners

We proceed to delineate five alternatives for the matrix M .

3.1. Interface Probe Preconditioner

Interface probe preconditioning is a family of methods for approximating the true Schur complement C defined in (2.3) by low bandwidth matrices. We use the nomenclature $IP(k)$ to denote

the approximation sequence $M = A_\Gamma - E_k$, $k = 0, 1, 2, \dots$, where E_k is a matrix of semi-bandwidth k which produces the same action as $A_{\Gamma\Omega}A_\Omega^{-1}A_{\Omega\Gamma}$ on a set of $2k + 1$ test vectors. Note that when A arises from a five-point finite-difference discretization both the IP(0) and IP(1) preconditioners are tridiagonal because $A_{\Omega\Gamma}$ is. As k increases beyond 1, M acquires additional diagonals. Selection of test vectors of appropriate sparsity patterns enables the coefficients of E_k to be read directly off of the product involving $A_{\Gamma\Omega}A_\Omega^{-1}A_{\Omega\Gamma}$, hence the term “probe.” We report only on the row-sum conserving IP(0) herein. IP(1) is only rarely more cost effective than IP(0) over the range of non-symmetric scalar five-point stencil problems studied herein, and a law of diminishing returns sets in as k is increased.

IP(0) was invented independently by Chan and Eisenstat in 1985, immediately generalized to IP(k) in [14], and adopted for variable coefficient symmetric problems in [24] (where it was called the “modified Schur complement” method) and for nonsymmetric problems in [25, 26]. Symmetric versions of IP(0) and IP(1) have also been employed in [2, 3]. Many algebraic and spectral properties of banded and circulant probe preconditioners are derived in [13]. The interface probe technique has the advantage of being purely algebraic in character, and hence capable of being defined for arbitrary operators. It is aesthetically pleasing that the tunable parameter k may be taken from the crude approximation of 0 all the way to the full bandwidth exact solution. It has also been generalized in a straightforward way to multicomponent systems [26]. However, IP(k) for low k is not expected to be particularly useful for *arbitrary* matrices. The low k limit is motivated by the observation that the elements of C decay rapidly away from the diagonal for elliptic problems. In sufficiently simple elliptic problems (*e.g.*, those possessing constant coefficients) preconditioners described below taking better advantage of this structure are also possible, leaving IP(k) large but not unlimited regions of problem parameter space in which to exercise. Interface probing has the advantage of being automatically adaptive to spatial variation in the coefficients but the disadvantage of not possessing the property of spectral equivalence, a consequence of which is that it degrades as the mesh is refined. Experimentally [24], the condition number of the preconditioned Schur complement system for the Laplacian goes like $h^{-1/2}$, and this bound is conjectured to be the best attainable for any tridiagonal matrix based on experiments with an optimization code in [20]. An $h^{-1/2}$ bound is proved for a circulant probe-preconditioned system with periodic boundary conditions on the boundaries normal to the interface in [13].

3.2. Spectral Preconditioner

The spectral preconditioner is an exact eigendecomposition of a single interface, rectangular domain, constant coefficient convection-diffusion operator described in [12] as a generalization of [7]. We consider only the Dirichlet case herein, but generalizations to Neumann boundary conditions are straightforward. Let an interface of n interior nodes (*i.e.*, $h^{-1} = n + 1$) separate two subdomains of the same discrete length, and discrete heights m_1 and m_2 , respectively, over all of which is satisfied the difference equation

$$ax_{i-1,j} + bx_{i,j} + cx_{i+1,j} + dx_{i,j+1} + ex_{i,j-1} = f_{i,j}, \quad (3.1)$$

where i denotes the free index *along* the interface. We may write $M = DW\Lambda W^{-1}D^{-1}$, where W is the discrete sine transform of length n with matrix representation

$$[W]_{ij} = \sqrt{2h} \sin ij\pi h, \quad (3.2)$$

D is the diagonal matrix with elements

$$[D]_i = \left(\frac{a}{c}\right)^{(i-1)/2}, \quad (3.3)$$

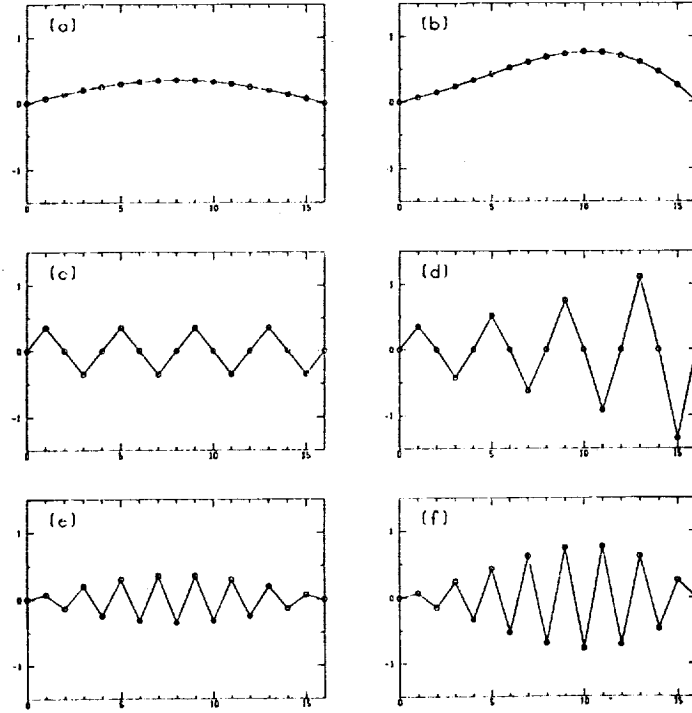


Figure 1: Modes of the Dirichlet problem (3.1) for $n = 15$. (a) $|a/c| = 1$, $j = 1$; (b) $|a/c| = 1.21$, $j = 1$; (c) $|a/c| = 1$, $j = 8$; (d) $|a/c| = 1.21$, $j = 8$; (e) $|a/c| = 1$, $j = 15$; (f) $|a/c| = 1.21$, $j = 15$. (The left-hand column of modes are for the case of no tangential convection.)

and Λ is a diagonal matrix with elements

$$[\Lambda]_i = \frac{1}{2} \left(\frac{1 + \gamma_i^{m_1+1}}{1 - \gamma_i^{m_1+1}} + \frac{1 + \gamma_i^{m_2+1}}{1 - \gamma_i^{m_2+1}} \right) \sqrt{[b + \sqrt{ac}(2 - \sigma_i)]^2 - 4de}, \quad (3.4)$$

where, in turn,

$$\gamma_i = \frac{1}{4dc} \left[b + \sqrt{ac}(2 - \sigma_i) + \sqrt{[b + \sqrt{ac}(2 - \sigma_i)]^2 - 4de} \right]^2 \quad (3.5)$$

and

$$\sigma_i = 4 \sin^2 \left(\frac{i\pi}{2(n+1)} \right). \quad (3.6)$$

The derivation of these formulae (see [12] for full details) begins with the observation that the columns of the matrix (DW) are the eigenvectors of the tridiagonal matrix formed by the coefficients along the interface, *viz.*, $\text{tridiag}(a, b, c)$. Sample such modes are plotted in Figure 1 for two different values of the ratio $|a/c|$ corresponding to zero and constant nonzero tangential components of the convection. The nonvanishing first-derivative convection term has the effect of multiplying the sinusoids by an exponential.

The philosophy of using the spectral preconditioner for arbitrary interfacial systems is that of solving an approximate (constant coefficient) problem exactly, rather than an exact (general

coefficient) problem approximately. One of its advantages is that it can be defined without requiring the ability to solve problems in adjacent subdomains, as required by the interface probe technique. All that is needed is some averaging rule to obtain the coefficients a through ϵ from the data of the associated regions. All our tests herein employ a simple average of the coefficients along the interface alone. Another advantage is its automatic adaptivity to domain aspect ratio, since the boundary conditions are built into the derivation. We note that application of M^{-1} is inexpensive: two one-dimensional FFTs sandwiched between three diagonal matrix multiplications.

3.3. Spectral Probe Preconditioner

The spectral probe preconditioner, introduced here for the first time, is conceptually a hybrid of the interface probe and the spectral preconditioners. Spectral probing assumes a form for the eigenvectors of C like that derived for the constant coefficient operator of the previous subsection (again based on spatially averaged coefficients), but then populates the diagonal matrix Λ by probing the true Schur complement, so that some spatial adaptivity is accommodated within a spectrally equivalent framework.

We set $M = DW\Lambda W^{-1}D^{-1}$ where W and D are defined as above (or where D is alternatively simply set to the identity matrix, corresponding to $a = c$, for reasons which will become clear in Section 4). Λ is then determined by probing with the interface vector of all 1's. This is the same as the standard test vector for IP(0). To be explicit, we read off the elements of Λ from the equation

$$W^{-1}D^{-1}CDW * \mathbf{1} = \Lambda * \mathbf{1}.$$

The action of C is computed by means of a pair of subdomain solves using $DW * \mathbf{1}$ as the interface boundary condition. Note that the spectral probe preconditioner reduces to the spectral preconditioner in the constant coefficient Dirichlet case, because then C is exactly diagonalized by the given similarity transform.

3.4. Laplacian Square-root Preconditioner

As a base-line reference, and because it appears throughout the literature, we include tests with a method based on the square-root of the one-dimensional Laplacian operator, easily written as:

$$M = W\Lambda W^{-1},$$

where Λ is now the diagonal matrix with elements $[\Lambda]_i = 2\sqrt{\sigma_i}$. We sometimes denote this operator as Dryja's preconditioner because of its popularization in this context in [16]. More general discrete antecedents were considered in [18]. It is difficult to pinpoint the discovery of the spectral equivalence of this operator to the Schur complement of the Laplacian, since the continuous analog of this equivalence has been known for some time. This preconditioner is expected to be good in diffusion-dominated problems, or in the discrete limit $h \rightarrow 0$.

Note that this preconditioner is distinct from the spectral preconditioner (§3.2) for the Laplacian. Dryja's preconditioner achieves a constant bound on the number of iterations as the mesh is refined, but the constant is generally higher than that achievable with the coefficient and aspect ratio adaptability of the previous two techniques. The literature also records two important preconditioners intermediate between the Dryja and spectral techniques, namely [19] and [5]. The latter, the Neumann-Dirichlet preconditioner, contains some of the adaptive capabilities of the spectral preconditioner since it relies on subdomain solves in its construction and hence contains much coefficient information. It is similar to probing techniques in this regard. In fact, the Neumann-Dirichlet preconditioner is exact in problems possessing symmetry across the separator set. All four of the techniques of [5, 7, 16, 19] were tested in [24], but for brevity we test only the extremes here.

3.5. Tangential Preconditioner

Finally, we consider a simple preconditioner possessing partial adaptivity, a lower-dimensional restriction of the operator to the interface created by setting all of the normal derivative terms in

the operator to zero and retaining just the remainder in M . For (1.1) these are just the tangential derivative terms. The obvious motivation for this technique is that it is simple and is expected to work well in the limit of strong convection along the interface, a limit which turns out to be troublesome for the spectral and spectral probe preconditioners. In addition, its very satisfactory behavior in the multidomain experiments in [21] recommend it. For reasons not yet theoretically explained, it performs very well in conjunction with the block triangular form of the the overall preconditioner described in Section 2. A minor disadvantage is its requirement of partial knowledge of the differential operator, rather than simply the elements of the discrete operator A . To be specific, it is necessary to store separately the contributions to A arising from the normal derivative terms, and all other terms.

4. Numerical Experiments

All of the experiments to follow except for those of Table 14 are posed on the unit square ($l = 1$ in the definition of the Reynolds number, Re) with homogeneous Dirichlet boundary conditions. The five different continuity-satisfying flow fields tested are shown in Figure 2(b)-(f), along with a purely diffusive baseline case (Figure 2(a)). When Reynolds numbers are reported below for the variable coefficient cases, they are always based on the maximum velocity in the region. (See [28] for details on the jet and cell flows and other experiments on this particular problem set.) The interface divides the rectangle into equal upper and lower portions, as marked on the figure in the dashed line. In addition to cases with constant diffusion, we study in §4.3 a convectionless case with piecewise constant, but disparate, diffusivities on either side of the interface.

There is a constant source term of unit strength in the interior. Although it is special, a zero initial guess for the solution vector is employed throughout, since this will usually be the natural choice when (1.1) arises for a Newton increment, as part of an outer nonlinear iteration. The performance of the preconditioners is measured by the number of iterations required to reduce the initial residual by a factor of 10^{-5} , regardless of the mesh resolution. The tables of iteration counts are grouped by subsections into four sets of experiments.

4.1. Sensitivity to Mesh Refinement

Tables 1 through 6 examine a constant Re situation as the (uniform) mesh is refined by three successive powers of 2. Of course, the discrete diffusion term, the Laplacian, becomes more and more dominant with each refinement of the grid, since it scales as h^{-2} as compared with the h^{-1} scaling of the convection term. This is the asymptotic limit for which D and its relatives S and SP are designed. In the first table, the Laplacian is studied in isolation ($Re = 0$). In the next five convective cases, $Re = 16$. For the coarsest mesh ($h^{-1} = 8$), the contributions to the diagonal of the discrete operator from the two terms are equal at this Reynolds number (the *cell* Reynolds number, $Re_c \equiv ch/\epsilon$, is 2).

h^{-1}	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
8	4	1	1	5	5	5	2	2	4	4
16	6	1	1	5	7	7	2	2	5	4
32	9	1	1	5	9	9	2	2	5	4
64	11	1	1	4	11	12	2	2	5	4

Table 1: Iteration counts for the pure diffusion problem as a function of mesh parameter for two different preconditioner structures and five different interface blocks.

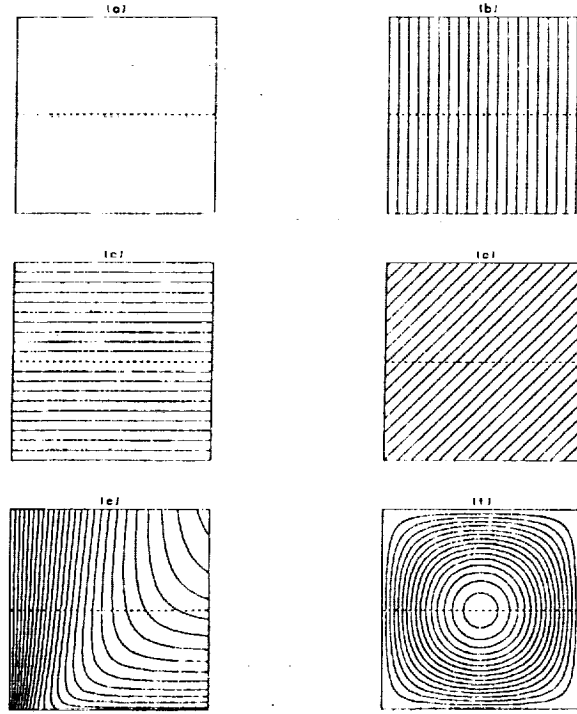


Figure 2: Streamfunction contour plots of the two-dimensional flow fields represented by \vec{c} in the numerical experiments. (a) Pure Diffusion; (b) Normal Convection; (c) Tangential Convection; (d) Skew Convection; (e) Jet Convection (the domain is the right half of a symmetric flow field); (f) Cell Convection.

The S (spectral), SP (spectral probe), and D (Dryja) columns of Table 1 reveal their exactness or spectral equivalence, respectively. Because iteration count is a threshold measurement, most of the data is subject to ± 1 perturbation upon modest adjustment of the convergence tolerances, but the S and SP residuals at convergence are zero to machine precision. The deterioration of IP like some negative power of h is evident on both the structurally symmetric (B_1) and block triangular (B_2) sides of the table. The tangential preconditioner is the only one with markedly different performance depending upon the structure of B . Here, as below, it is excellent in conjunction with the block triangular form.

h^{-1}	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
8	3	1	1	4	5	4	2	2	4	5
16	5	1	1	5	7	6	2	2	5	5
32	6	1	1	5	9	7	2	2	5	5
64	7	1	1	5	12	8	2	2	5	5

Table 2: Same as Table 1, except for normal convection at a Reynolds number of 16.

Table 2, for a normal convection problem, is similar to Table 1 except that IP improves slightly

as each of the terms $A_{31}A_{11}^{-1}A_{13}$ and $A_{32}A_{22}^{-1}A_{23}$ being approximated by a diagonal matrix becomes less important relative to A_{33} because one of the coupling matrices is small. For instance, if the convection is from subdomain 1 into subdomain 2, A_{13} and A_{32} are weak.

h^{-1}	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
8	5	1	7	8	6	5	2	8	7	4
16	6	1	10	10	9	7	2	11	10	5
32	8	1	11	11	11	9	2	12	11	5
64	11	1	12	11	15	12	2	13	11	5

Table 3: Same as Table 1, except for tangential convection at a Reynolds number of 16.

The importance of the D matrix in the spectral preconditioner is evident in Table 3 in which a tangential convection problem is considered. The version of SP employed in this study approximates the D in its definition as the identity; using the true D here would reproduce the spectral results in this constant coefficient case, just as in the previous two tables in which $D = I$ anyway. Though SP and D are spectrally equivalent, they require an order of magnitude more iterations than S, and are surpassed by IP in the smaller problem range on the structurally symmetric side, and by the tangential preconditioner on the block triangular side.

h^{-1}	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
8	4	1	7	8	7	5	2	8	8	6
16	5	1	9	10	9	6	2	10	9	6
32	7	1	10	10	12	8	2	11	10	6
64	9	1	10	10	15	10	2	11	11	7

Table 4: Same as Table 1, except for skew convection at a Reynolds number of 16.

Table 4 contains the last of the constant coefficient test examples, featuring a skew convection (inclined at 45 degrees relative to the interface). Overall results are not very different from the purely tangential convection case. Most of the entries in the skew table lie at or between corresponding entries in the normal and tangential tables. Those of the tangential preconditioner, however, are often worse in the intermediate skew convection case than at either extreme.

h^{-1}	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
8	5	4	5	6	6	5	5	6	6	5
16	6	4	6	6	8	7	5	7	7	6
32	9	4	6	6	11	10	5	7	7	7
64	11	4	6	6	14	12	5	7	7	8

Table 5: Same as Table 1, except for jet convection at a Reynolds number of 16.

The jet case recorded in Table 5 tends to level the preconditioner landscape because the constant coefficient approximation of S is no longer exact. S remains the best technique as h^{-1} increases, but its margin of superiority over other spectrally equivalent techniques (with D as a nonadaptive extreme) is small.

h^{-1}	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
8	4	5	4	5	5	5	5	5	5	5
16	6	5	5	6	7	7	6	6	6	6
32	9	5	5	5	9	10	6	6	6	6
64	11	5	5	5	11	12	6	6	6	6

Table 6: Same as Table 1, except for cell convection at a Reynolds number of 16.

The cell case of Table 6 is the greatest equalizer among the test cases, because the interface cuts a zone of recirculation, *i.e.*, there is normal flow across it in both directions, and none of the methods holds an edge of superiority. Performance under the block tridiagonal preconditioner is notably uniform for the last four methods.

4.2. Sensitivity to Reynolds Number

Tables 7 through 11 examine the influence of increasing Reynolds number. Values of Re of 0, 4, 16, 64, 256, and 1024 are considered at $h^{-1} = 64$. Thus, the third row of each table in this subsection is the same as the last row of the table of corresponding flow type in the first set, and the first row of each table is the same as the last row of the pure diffusion case in Table 1. Progressing down the rows of the table the nonsymmetry of the operator increases. Between rows four and five the convection terms begin to contribute more heavily than the diffusion terms to the diagonal elements of A .

Re	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
0	11	1	1	4	11	12	2	2	5	4
4	10	1	1	4	11	11	2	2	5	4
16	7	1	1	5	12	8	2	2	5	5
64	5	1	1	7	14	6	2	2	7	6
256	3	1	1	11	17	4	2	2	10	9
1024	2	1	1	15	22	3	2	2	12	16

Table 7: Iteration counts for the normal convection problem at a constant mesh parameter of $1/64$ as a function of Reynolds number for two different preconditioner structures and five different interface blocks.

Table 7 shows that in the presence of constant normal convection, techniques S and SP remain exact at any Re , while IP catches up at high Re , and D and T successively worsen. (For D , this is the drawback of finite h in a method derived for $h \rightarrow 0$.) Qualitatively, the trends are the same for either preconditioner structure, although the tangential method continues to be much better in the block triangular formulation.

The Achilles' heel of the spectral technique appears when there is strong convection tangential to the interface, as seen in Table 8. In this limit in which $|a/c|$ differs sufficiently from unity the

Re	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
0	11	1	1	4	11	12	2	2	5	4
4	12	1	7	7	14	13	2	8	8	5
16	11	1	12	11	15	12	2	13	11	5
64	8	1	20	15	14	9	2	21	15	3
256	7	-	>	20	12	8	-	>	19	1
1024	5	-	>	26	8	6	-	>	24	1

Table 8: Same as Table 7, except for tangential convection. The hyphen denotes loss of precision, and the “>” more than 30 iterations.

latter terms in D , which have this ratio raised to as much as $(n - 1)^{\text{st}}$ power, can approach the machine epsilon (or its reciprocal, depending upon the direction of the convection). In reference to (3.3), we note that for $n = 64$, $(a/c)^{(n-1)/2} \approx \epsilon_{\text{mach}} \approx 10^{-16}$ when $(a/c) \approx 10^{(-32/63)} \approx 0.31$. Under upwind differencing, it only takes a cell Reynolds number of 2 to produce a ratio of 3 between the upwind and downwind stencil coefficients a and c . Thus, $\text{Re}_c = 2$ is the borderline of stability for the spectral preconditioner with respect to tangential convection. In the tables, the $\text{Re} = 64$ row corresponds to a cell Reynolds number of unity, and the $\text{Re} = 256$ row to a cell Reynolds number of 4; thus, they straddle the transition. GMRES iterations based on the spectral interface preconditioner do terminate for the hyphenated entries, but the residuals based on the resulting x vectors shows complete loss of precision.

The spectral probe technique does not lose precision, because of the assumption that $D = I$; however, a diagonal approximation for $W^{-1}CW$ is poor, and it simply takes too long to converge. The Dryja preconditioner, which makes no attempt to adapt to the strong directionality of the problem also deteriorates with Re . At low Re where M and C are close, the B_1 iteration count is lower by one; but at high Re , where isotropic M is very different from unidirectional C , the B_2 iteration count is better. The interface probe technique meanwhile improves with Re as it captures more of the increasingly diagonally dominant problem within its own sparsity structure. Finally, the tangential technique is excellent for a tangentially dominated operator. The cross-diffusion which it neglects becomes of negligible importance as the problem effectively decouples into independent problems for A_{11} , A_{22} , and A_{33} in which the upstream boundary condition is all important.

Re	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
0	11	1	1	4	11	12	2	2	5	4
4	11	1	7	7	14	12	2	8	8	6
16	9	1	10	10	15	10	2	11	11	7
64	7	1	13	14	17	7	2	14	14	8
256	4	-	13	18	19	5	-	14	16	9
1024	3	-	13	20	19	4	-	15	17	10

Table 9: Same as Table 7, except for skew convection.

Most of the observations of the high Re tangential flow also apply to the skew flow in Table 9, however, the latter differentiates between IP, which repeats its tendency to improve as Re grows, and T, which no longer matches the physics of the problem, and is worse than IP, although it is still superior to the Fourier-based methods as a module of the block triangular preconditioner.

We also note that the spectral probe technique captures a significant part of the physics in this case, adapting to the co-dominant normal convection and tending to a constant iteration count without breaking down. It thus rescues the spectral technique and indicates how the robustness of the spectral preconditioner can be maintained with a compromise in efficiency. In a general purpose code, the elements of the D matrix could differ from unity but be bounded artificially, allowing partial tangential adaptivity with full normal adaptivity.

Re	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
0	11	1	1	4	11	12	2	2	5	4
4	11	3	4	5	14	12	4	5	6	6
16	11	4	6	6	14	12	5	7	7	8
64	10	6	9	9	14	11	7	10	10	9
256	9	7	11	14	13	10	8	12	14	10
1024	9	-	20	19	14	10	-	21	18	11

Table 10: Same as Table 7, except for jet convection.

As in the spectral equivalence tests in Table 5, the jet case recorded in Table 10 tends to diminish the extremes of preconditioner behavior relative to the uniform skew convection case. The IP and SP results worsen while the D and T methods nearly hold their own relative to Table 9. The pure spectral method survives at a higher overall Reynolds number because the tangential velocities at the interface are lower than the maximum core velocity of the jet, to which Re is scaled.

Re	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
0	11	1	1	4	11	12	2	2	5	4
4	11	3	3	4	11	12	4	4	5	5
16	11	5	5	5	11	12	6	6	6	6
64	13	8	7	8	13	14	9	8	9	8
256	18	14	11	14	16	19	15	12	14	15
1024	25	23	16	23	26	26	23	17	23	24

Table 11: Same as Table 7, except for cell convection.

Again, the cell case of Table 11 is an equalizer for most of the methods; however, the performance of the spectral probe technique is singularly good. This is easily explained as follows: the average tangential velocity along the interface is zero because of the symmetry of Figure 2(f), so $D = I$ for both S and SP. However, S also employs an zero average for the normal velocity, whereas SP adapts to strong inflow and outflow along opposite halves of the interface. The performance of IP, which improves with Re in all previous tables, deteriorates in this table because no pair of coupling matrices is weak (see comments on Table 2) under recirculation. Performances under the structurally symmetric and block tridiagonal schemes of the tangential preconditioner become similar at high Re. The recirculating cell flow is in some sense a worst case for a single interface. If the domain of Figure 2(f) is further decomposed by a vertical interface, putting a vertex in the center, all four interfaces would be free of two-signed velocity components, and easier to precondition. (The cross-point preconditioning then becomes an important subject.)

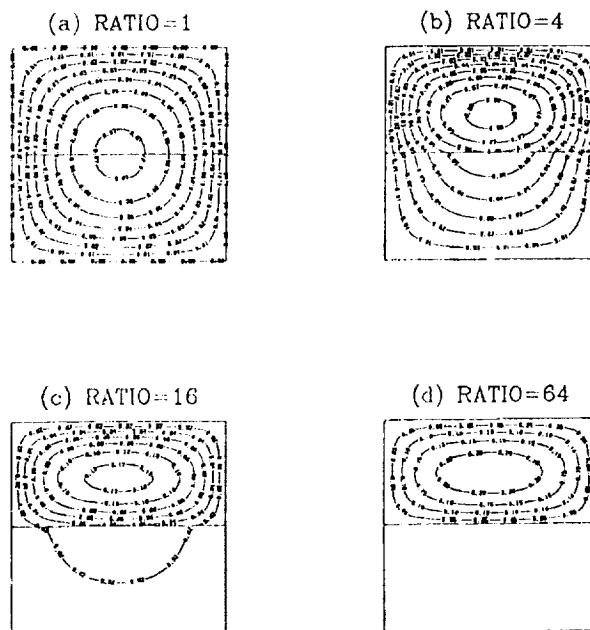


Figure 3: Contour plots of the solution to $\nabla \cdot \epsilon \nabla u = f$ with homogeneous Dirichlet boundary conditions and $f = 1$ for various ratios of the (piecewise constant) diffusivity in the adjoining subdomains. “Ratio” is $\epsilon_{lower}/\epsilon_{upper}$ and the product $\epsilon_{lower}\epsilon_{upper}$ is maintained at unity.

4.3. Piecewise Constant Diffusivity

This subsection focuses on the effectiveness of the preconditioners in heterogeneous problems. With reference to Figure 3, we assign a higher diffusivity in the lower domain than in the upper. We report convergence results in two limits: fixed diffusivity ratio and increasing h^{-1} (Table 12) and fixed h^{-1} and increasing diffusivity ratio (Table 13).

The performance of the two probe and the tangential preconditioners in the terraced diffusivity case in Table 12 is identical to their performance in the uniform case of Table 1. However, the Dryja preconditioner is slightly worse and the spectral preconditioner (based on the arithmetically averaged diffusivity) is much worse than in the uniform case. In fact, S is the worst preconditioner in the terraced diffusion case while SP is the best, the probe supplying crucial information.

As the ratio of diffusivities increases for a fixed problem size, as in Table 13, all methods asymptote monotonically to fixed convergence rates, as shown theoretically in [13]. The physical interpretation of the asymptotically large diffusivity ratio is that virtually all of the variation of solution in response to the (fixed) forcing occurs in the subdomain with the lesser diffusivity. Figure 3 illustrates this phenomenon. The contours are kinked at the interface since it is the normal component of the diffusive flux, $c\nabla u$, not ∇u itself, which must balance on either side. The solution along the interface is asymptotically the boundary value of zero. The spectral probe method is a singularly good performer in this limit.

4.4. Sensitivity to Aspect Ratio

Table 14 examines the constant diffusion case under aspect ratios ranging from $1:\frac{1}{16}$ (a squat

h^{-1}	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
8	4	4	1	5	5	5	5	2	5	4
16	6	6	1	5	7	7	7	2	5	4
32	9	9	1	6	9	9	9	2	5	4
64	11	12	1	6	11	12	13	2	5	4

Table 12: Iteration counts for the step diffusion problem as a function of mesh parameter for two different preconditioner structures and five different interface blocks. The ratio of the diffusion coefficients of the two subdomains is 64.

Ratio	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
1	9	1	1	5	9	9	2	2	5	4
4	9	7	1	5	9	9	8	2	5	4
16	9	8	1	5	9	9	9	2	5	4
64	9	9	1	6	9	9	9	2	5	4
256	9	9	1	6	9	9	9	2	6	4
1024	9	9	1	6	9	9	9	2	6	4

Table 13: Iteration counts for the step diffusion problem at a constant mesh parameter of $1/32$ as a function of diffusivity ratio for two different preconditioner structures and five different interface blocks.

rectangle with subdomains just two cells high) to 1:2 (a tall rectangle composed of two square subdomains). Note that the discrete length of the interface, n , remains constant at 63 in all examples, while heights m_1 and m_2 vary from 1 to 63.

(l_x, l_y)	Structurally Symmetric					Block Triangular				
	IP	S	SP	D	T	IP	S	SP	D	T
$(1, \frac{1}{16})$	4	1	1	8	14	4	2	2	8	13
$(1, \frac{1}{8})$	5	1	1	6	13	6	2	2	6	10
$(1, \frac{1}{4})$	7	1	1	5	12	8	2	2	5	7
$(1, \frac{1}{2})$	9	1	1	5	11	10	2	2	5	5
$(1, 1)$	11	1	1	4	11	12	2	2	5	4
$(1, 2)$	9	1	1	4	8	10	2	2	5	2

Table 14: Iteration counts for the pure diffusion problem at a constant mesh parameter of $1/64$ as a function of aspect ratio for two different preconditioner structures and five different interface blocks.

We confirm that S and SP are completely adaptive to aspect ratio in this constant coefficient problem, which is not true of any other method, including Dryja's. The tangential preconditioner is understandably good when the narrow (more strongly boundary influenced) direction is the tangential one, and poorer when this aspect ratio is reversed. The nonmonotonic character of IP is interesting, showing that it adapts well to either extreme and is poorest in the balanced

intermediate case.

5. Conclusions

We conclude by pulling together some overall assessments of the preconditioners tested in the previous section. The interface probe technique is the most general purpose and robust of the methods. It is always definable and adapts well to high Re and extreme aspect ratio, but is nonoptimal at high h^{-1} and is occasionally the worst method. It does very well in predominantly unidirectional strongly convective flows. Generalizations to multiple interfaces, multiple components, and inexact subdomain solves are straightforward, but not pursued herein (see [26]).

The spectral method is mathematically the method of choice asymptotically in h^{-1} , where physically well resolved problems should end up. However, since it is based on a global constant-coefficient approximation it does not perform well in heterogeneous problems. Furthermore, high cell Reynolds numbers can cause it to go unstable, and flow fields whose actual values along the interface differ greatly from their average values are poorly represented. A stabilizing technique was proposed which could preserve the robustness of the spectral method at high cell Reynolds numbers, namely selection of an exponent base for D in (3.3) closer to unity than the true $|a/c|$.

The spectral probe method is as good as the spectral method when W alone is a good approximation to the eigenvectors of C (i.e., low tangential convection). A D -modulated version of spectral probe can be just as effective (and just as vulnerable to strong tangential convection) as the pure spectral method in a constant coefficient problem. SP adapts better than S to normal convection variation and it adapts perfectly to piecewise constant heterogeneities in the diffusivity. As a “probe” method, it shares the coding disadvantages of IP.

The Dryja preconditioner is never exclusively the best method, but is as good as either S or SP in a variable coefficient problem in the limit $h \rightarrow 0$, where the diffusive contributions to A dominate. However, the extra cost of S is insignificant compared to D, and SP costs only extra subdomain solves in the pre-processing, so these techniques (suitably stabilized for tangential convection) will almost always be preferable in applications.

All of the above techniques are relatively insensitive to the choice of the overall preconditioner structure. The tangential interface preconditioner is an exception, for reasons yet to be explained theoretically. It is much better under the block triangular form of the overall preconditioner, and is very competitive with the other techniques under physically predictable circumstances, namely tangentially dominant convection or narrow aspect ratio. It is also simple to code and has been successfully generalized to multicomponent systems (see [22]).

We note that when exact subdomain solves are performed, the block triangular form of the preconditioner, B_2 , with its one-directional flow of information from the separators to the interiors, is almost always preferable to the structurally symmetric form B_1 in terms of execution time, since a full set of subdomain solves is saved at each iteration and the iteration counts are usually comparable. The structurally symmetric form is also obviously useful when A is itself symmetric, since it then admits preconditioned conjugate gradients rather than GMRES as the iterative solver.

Clearly, a user who understands his problem physically and is willing to customize can do well by choosing among a variety of interface preconditioner modules, perhaps using different ones on differently resolved grids within the same overall solution procedure. Hopes of finding a single “best” method from among those considered must clearly be dismissed, but the field is still young. It is clear that general purpose adaptivity will require some form of “probing” via subdomain solves; taken to the limit, one obtains C directly. Such probing has an associated cost comparable to an iteration step, and must be undertaken conservatively. Future developments which iteratively improve the interface preconditioning based on accumulated subdomain solve data would be welcome, and so would more hybrids along the lines of the spectral probe which incorporate both analytical and numerical data.

References

- [1] C. Ashcraft, *Domain Decoupled Incomplete Factorizations*, Technical Report ETA-TR-49, Boeing Computer Services, April 1987.
- [2] O. Axelsson & B. Polman, *Block Preconditioning and Domain Decomposition Methods, I*, Technical Report 8735, Dept. of Mathematics, Catholic University, Nijmegen, December 1987.
- [3] ———, *Block Preconditioning and Domain Decomposition Methods, II*, J. Comp. Appl. Math., 24(1988), pp. 55-72.
- [4] R. E. Bank, T. F. Chan, W. M. Coughran, Jr. & R. K. Smith, *The Alternate-Block-Factorization Procedure for Systems of Partial Differential Equations*, Technical Report Numerical Analysis Manuscript 89-5, AT& Bell Laboratories, July 1989.
- [5] P. E. Bjorstad & O. B. Widlund, *Iterative Methods for the Solution of Elliptic Problems on Regions Partitioned into Substructures*, SIAM J. Num. Anal., 23(1986), pp. 1097-1120.
- [6] ———, *To Overlap or Not to Overlap: A Note on a Domain Decomposition Method for Elliptic Problems*, SIAM J. Sci. Stat. Comp., 10(1989), pp. 1053-1061.
- [7] T. F. Chan, *Analysis of Preconditioners for Domain Decomposition*, SIAM J. Num. Anal., 24(1987), pp. 382-390.
- [8] ———, *Boundary Probe Preconditioners for Fourth Order Elliptic Problems*, T. F. Chan, R. Glowinski, J. Periaux & O. Widlund ed., *Second International Symposium on Domain Decomposition Methods*, SIAM, Philadelphia, 1989, pp. 160-167.
- [9] T. F. Chan, R. Glowinski, J. Periaux & O. Widlund, *Second International Symposium on Domain Decomposition Methods*, SIAM, Philadelphia, 1989.
- [10] ———, *Third International Symposium on Domain Decomposition Methods for Partial Differential Equations*, SIAM, Philadelphia, 1990.
- [11] T. F. Chan & D. Goovaerts, *Schwarz = Schur: Overlapping versus Nonoverlapping Domain Decomposition*, Technical Report 88-21, UCLA Comp. and App. Math., August 1988. (To appear in SIAM J. Matrix Anal. and Appl.).
- [12] T. F. Chan & T. Y. Hou, *Domain Decomposition Preconditioners for General Second Order Elliptic Operators*, Technical Report 88-16, UCLA Comp. and App. Math., June 1988.
- [13] T. F. Chan & T. P. Mathew, *The Boundary Probing Technique in Domain Decomposition*, Technical Report 91-02, UCLA Comp. and App. Math., February 1991.
- [14] T. F. Chan & D. Resasco, *A Survey of Preconditioners for Domain Decomposition*, Technical Report 414, Computer Science Dept., Yale University, September 1985. (In Proceedings of the IV Coloquio de Matemáticas del CINVESTAV, Workshop in Numerical Analysis and its applications, Taxco, Mexico, Aug. 18-24, 1985).
- [15] R. W. Cottle, *Manifestations of the Schur Complement*, Lin. Alg. Appl., 8(1974), pp. 189-211.
- [16] M. Dryja, *A Capacitance Matrix Method for Dirichlet Problem on Polygonal Region*, Numer. Math., 39(1982), pp. 51-64.
- [17] R. Glowinski, G. H. Golub, G. A. Meurant & J. Periaux, *First International Symposium on Domain Decomposition Methods for Partial Differential Equations*, SIAM, Philadelphia, 1988.
- [18] R. Glowinski & O. Pironneau, *Numerical Methods for the First Biharmonic Equation and for the Two-Dimensional Stokes Problem*, SIAM Review, 21(1979), pp. 167-212.

- [19] G. H. Golub & D. Mayers, The Use of Preconditioning over Irregular Regions, R. Glowinski & J. L. Lions ed., *Computing Methods in Applied Sciences and Engineering VI*, North Holland, 1984, pp. 3-14.
- [20] A. Greenbaum & G. Rodrigue, *Optimal Preconditioners of a Given Sparsity Pattern*, Technical Report 431, Courant Institute, February 1989.
- [21] W. D. Gropp & D. E. Keyes, *Domain Decomposition with Local Mesh Refinement*, Technical Report 91-20, Inst. for Comput. Appl. Sci. Eng., January 1991. (to appear in SIAM J. Sci. Stat. Comp.).
- [22] ———, *Domain Decomposition Methods in Computational Fluid Dynamics*, Technical Report 91-20, Inst. for Comput. Appl. Sci. Eng., January 1991. (to appear in Int. J. Num. Meths. Fluids).
- [23] D. E. Keyes, *Domain Decomposition Methods for the Parallel Computation of Reacting Flows*, Comput. Phys. Comm., 53(1989), pp. 181-200.
- [24] D. E. Keyes & W. D. Gropp, *A Comparison of Domain Decomposition Techniques for Elliptic Partial Differential Equations and their Parallel Implementation*, SIAM J. Sci. Stat. Comp., 8(1987), pp. s166-s202.
- [25] ———, Domain Decomposition for Nonsymmetric Systems of Equations: Examples from Computational Fluid Dynamics, T. F. Chan, R. Glowinski, J. Periaux & O. Widlund ed., *Second International Symposium on Domain Decomposition Methods*, SIAM, Philadelphia, 1989, pp. 321-339.
- [26] ———, *Domain Decomposition Techniques for the Parallel Solution of Nonsymmetric Systems of Elliptic BVPs*, Appl. Num. Math., 6(1990), pp. 281-301.
- [27] ———, *Domain-Decomposable Preconditioners for Second-order Upwind Discretizations of Multicomponent Systems*, Technical Report MCS-P187-1090, Argonne National Laboratory, November 1990.
- [28] D. E. Keyes, W. D. Gropp & A. Ecker, Domain Decomposition Techniques for Large Sparse Nonsymmetric Systems Arising from Elliptic Problems with First-Order Terms, J. H. Kane & A. D. Carlson ed., *Proceedings of a Symposium on the Solution of Super Large Problems in Computational Mechanics*, Plenum, New York, 1989, pp. 251-266.
- [29] L. Mansfield, *On the Conjugate Gradient Solution of the Schur Complement System Obtained from Domain Decomposition*, SIAM J. Num. Anal., 27(1990), pp. 1612-1620.
- [30] Y. Saad & M. H. Schultz, *GMRES: A Generalized Minimum Residual Algorithm for Solving Nonsymmetric Linear Systems*, SIAM J. Sci. Stat. Comp., 7(1986), pp. 856-869.