

Parallel Implementation of an Optimal Two Level Additive Schwarz Preconditioner for the 3-D Finite Element Solution of Elliptic Partial Differential Equations

S.A. Nadeem*and P.K. Jimack†
 Computational PDE Unit, School of Computing,
 University of Leeds, Leeds LS2 9JT, UK.

Abstract

This paper presents a description of the extension and parallel implementation of a new two level additive Schwarz (AS) preconditioner for the solution of 3-d elliptic partial differential equations (PDEs). This preconditioner, introduced in [4], is based upon the use of a novel form of overlap between the subdomains which makes use of a hierarchy of meshes: with just a single layer of overlapping elements at each level of the hierarchy. The generalization considered here is based upon the restricted AS approach reported in [5] and the parallel implementation is an extension of work in two dimensions, [3].

1 Introduction

To begin a brief description of the optimal two level AS preconditioner introduced in [4] is provided. The construction of this preconditioner is based upon the existence of a nested sequence of meshes on each subdomain, each with a single layer of overlap onto the neighbouring subdomains.

Let \mathcal{T}_0 be a coarse triangulation of the problem domain, Ω say, consisting of N_0 tetrahedral elements, $\tau_j^{(0)}$, such that $\tau_j^{(0)} = \bar{\tau}_j^{(0)}$,

$$\bar{\Omega} = \bigcup_{j=1}^{N_0} \tau_j^{(0)} \quad \text{and} \quad \mathcal{T}_0 = \{\tau_j^{(0)}\}_{j=1}^{N_0}. \quad (1.1)$$

Also let $\text{diameter}(\tau_j^{(0)}) = O(H)$, and divide Ω into p *non-overlapping* subdomains Ω_i . These subdomains should be such that:

$$\bar{\Omega} = \bigcup_{i=1}^p \bar{\Omega}_i, \quad (1.2)$$

$$\Omega_i \cap \Omega_j = \emptyset \quad (i \neq j), \quad (1.3)$$

$$\bar{\Omega}_i = \bigcup_{j \in I_i} \tau_j^{(0)} \quad \text{where } I_i \subset \{1, \dots, N_0\} \quad (I_i \neq \emptyset). \quad (1.4)$$

The coarse triangulation \mathcal{T}_0 may be refined several times, to produce a family, $\mathcal{T}_0, \dots, \mathcal{T}_J$, where each triangulation in this family, \mathcal{T}_k , consists of N_k tetrahedral elements, $\tau_j^{(k)}$, such that

$$\bar{\Omega} = \bigcup_{j=1}^{N_k} \tau_j^{(k)} \quad \text{and} \quad \mathcal{T}_k = \{\tau_j^{(k)}\}_{j=1}^{N_k}. \quad (1.5)$$

*Funded by the Government of Pakistan through a Quaid-e-Azam scholarship.

†Corresponding author: pkj@comp.leeds.ac.uk

The successive mesh refinements that define this sequence of triangulations need not be global and may be non-conforming, however it is necessary that they satisfy a number of conditions:

1. $\tau \in \mathcal{T}_{k+1}$ implies that either
 - (a) $\tau \in \mathcal{T}_k$, or
 - (b) τ has been generated as a refinement of an element of \mathcal{T}_k into eight regular children (by bisecting each edge of this parent of τ),
2. the level of any tetrahedra which share a common point can differ by at most one,
3. only tetrahedra at level k may be refined in the transition from \mathcal{T}_k to \mathcal{T}_{k+1} .

(Here the level of a tetrahedron is defined to be the least value of k for which that tetrahedron is an element of \mathcal{T}_k .) In addition to the above it is also necessary that:

4. in the final mesh, \mathcal{T}_J , all sets of tetrahedra which share an edge which lies on the interface of any subdomain with any other subdomain have the same level as each other (i.e. the mesh is conforming along subdomain interfaces).

Having defined a decomposition of Ω into subdomains and a nested sequence of triangulations of Ω the restrictions of each of these triangulations onto each subdomain may be defined by

$$\Omega_{i,k} = \{\tau_j^{(k)} : \tau_j^{(k)} \subset \bar{\Omega}_i\}. \quad (1.6)$$

In order to introduce a certain amount of overlap between neighbouring subdomains

$$\tilde{\Omega}_{i,k} = \{\tau_j^{(k)} : \tau_j^{(k)} \text{ has a common point with } \bar{\Omega}_i\} \quad (1.7)$$

is also defined. Following this the finite element (FE) spaces associated with these local triangulations may be introduced. Let G be some triangulation and denote by $\mathcal{S}(G)$ the space of continuous piecewise linear functions on G . Then the following definitions are made:

$$\mathcal{W} = \mathcal{S}(\mathcal{T}_J) \quad (1.8)$$

$$\mathcal{W}_0 = \mathcal{S}(\mathcal{T}_0) \quad (1.9)$$

$$\mathcal{W}_{i,k} = \mathcal{S}(\Omega_{i,k}) \quad (1.10)$$

$$\tilde{\mathcal{W}}_{i,k} = \mathcal{S}(\tilde{\Omega}_{i,k}) \quad (1.11)$$

$$\tilde{\mathcal{W}}_i = \tilde{\mathcal{W}}_{i,0} + \dots + \tilde{\mathcal{W}}_{i,J}. \quad (1.12)$$

It is evident that

$$\mathcal{W} = \mathcal{W}_0 + \tilde{\mathcal{W}}_1 + \dots + \tilde{\mathcal{W}}_p \quad (1.13)$$

where, as in (1.2), p is the number of subdomains. This is the decomposition that forms the basis of the two level additive Schwarz preconditioner (see [13] for example), M say, in [4]. Hence, for a global FE stiffness matrix A , this preconditioner takes the form:

$$M^{-1} = \sum_{i=0}^p R_i^T A_i^{-1} R_i \quad (1.14)$$

where (using the usual local bases for \mathcal{W} , \mathcal{W}_0 and $\tilde{\mathcal{W}}_i$) R_i is the rectangular matrix representing the L^2 projection from \mathcal{W} to \mathcal{W}_0 (when $i = 0$) or $\tilde{\mathcal{W}}_i$ (when $i > 0$), and A_i is the FE stiffness matrix corresponding to the subspace \mathcal{W}_0 (when $i = 0$) or $\tilde{\mathcal{W}}_i$ (when $i > 0$).

2 A Nonsymmetric Generalization of the Preconditioner

It is proved in [4] that, for a symmetric positive definite (SPD) stiffness matrix A , (1.14) is an optimal two level preconditioner. That is, the condition number of $M^{-1}A$ is bounded independently of h , H and p , where the tetrahedral elements at level J have a diameter of $O(h)$. Following the ideas of [5], for example, one may also consider a *restricted* version of this preconditioner that is not symmetric even when each A_i is. This is obtained by replacing the R_i^T terms in (1.14) with simpler alternatives which significantly reduce the amount of computation and communication required in a parallel implementation. The form of this restricted preconditioner is given by

$$\tilde{M}^{-1} = \sum_{i=0}^p \Pi_i^T A_i^{-1} R_i, \quad (2.1)$$

where Π_i^T is a simple prolongation matrix, all entries of which are zero with the following exceptions. For $i > 0$, if the node at which the n^{th} local basis function of $\tilde{\mathcal{W}}_i$ is non-zero is inside Ω_i then the entry in position (m, n) of Π_i^T is 1, where m is the number of this node in \mathcal{T}_J . Also, if the node at which the n^{th} local basis function of $\tilde{\mathcal{W}}_i$ is non-zero is on the interface of Ω_i then the entry in position (m, n) of Π_i^T is $\frac{1}{q}$, where m is as before and q is the number of neighbouring subdomains that share this interface node. For $i = 0$, entry (m, n) of Π_0^T is equal to 1 where m is the number of the node in \mathcal{T}_J which corresponds to each node n in \mathcal{T}_0 .

Unlike for (1.14) a proof of the optimality of the preconditioner (2.1) is not offered, however empirical evidence suggests that it may be optimal and that it performs better than (1.14) for typical test problems. For example, Table 1 shows the number of iterations required to solve a simple 3-d potential flow problem of the form

$$-\Delta\phi = f \quad \text{in } \Omega = (0, 2) \times (0, 1) \times (0, 1) \quad (2.2)$$

subject to Dirichlet boundary conditions throughout $\partial\Omega$. A Galerkin FE discretization is used on a sequence of uniformly refined meshes and the 2-norm of the residual is reduced by a factor of 10^5 in each case. Since (2.2) is a self-adjoint problem the discretization yields a SPD stiffness matrix A . In this case the preconditioner M , given by (1.14), is also SPD however \tilde{M} , given by (2.1), is not. Hence, although the conjugate gradient method ([1, 6]) may be used with M as the preconditioner, a more general iterative technique, such as GMRES ([11]), must be used in the restricted case. Nevertheless, Table 1 clearly demonstrates that this minor disadvantage of (2.1) over (1.14) is more than compensated for by the reduction in the number of iterations required. Furthermore, the computational (and, in the parallel implementation, communication) cost of applying \tilde{M} at each iteration is less than that of applying M .

It should be noted that the results presented in Table 1, and all other results presented in this paper, are for one particular set of partitions of the domain $\Omega = (0, 2) \times (0, 1) \times (0, 1)$. These are obtained by using a very simple recursive coordinate bisection algorithm (e.g. [12]) and, in this case, provide compact subdomains with a relatively small surface-area to volume ratio, as illustrated in Figure 1. It is shown in [9] that the shape of the subdomains can have some bearing on the precise performance of the preconditioner (2.1) however this issue will not be addressed here.

A further advantage of the preconditioner \tilde{M} over M is that it may be applied quite naturally to the solution of nonsymmetric or indefinite linear systems arising from the discretization of non-self-adjoint PDEs. For example, Table 2 shows the number of preconditioned GMRES iterations required to solve two convection-diffusion problems of the form

$$-\varepsilon\Delta u + \underline{b} \cdot \underline{\nabla}u = f \quad \text{in } \Omega = (0, 2) \times (0, 1) \times (0, 1) \quad (2.3)$$

Elements/Procs.	Preconditioner (1.14)				Preconditioner (2.1)			
	2	4	8	16	2	4	8	16
6144	6	11	16	19	3	3	4	4
49152	7	12	17	19	3	5	5	6
393216	8	13	19	20	4	6	7	8
3145728	9	13	19	20	4	8	9	9

Table 1: The performance of the symmetric AS preconditioner (1.14) and the reduced AS preconditioner (2.1) on the Galerkin FE discretization of (2.2): figures quoted represent the number of iterations required to reduce the initial residual by a factor of 10^5 on a sequence of uniform refinements of an initial mesh of 768 tetrahedral elements.

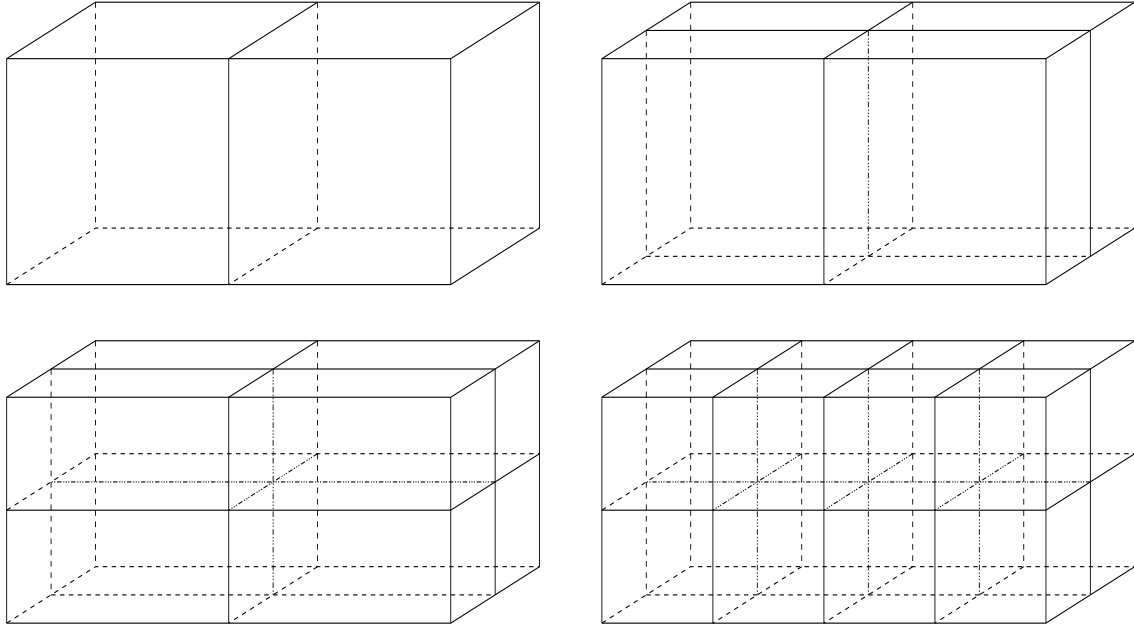


Figure 1: An illustration of the partitioning strategy, based upon recursive coordinate bisection, used to obtain 2, 4, 8 and 16 subdomains for our test problems where $\Omega = (0, 2) \times (0, 1) \times (0, 1)$.

on the same sequence of meshes as used in Table 1. In each case $\underline{b} = (1, 0, 0)^T$, Dirichlet boundary conditions have been applied throughout $\partial\Omega$ and a stabilized (streamline diffusion, [8]) FE discretization has been used. Again only a very slow growth in the number of iterations is observed as h is decreased or p is increased.

3 Parallel Implementation Issues

The parallel implementation of the preconditioned GMRES solver using \tilde{M} as the preconditioner may be undertaken in a number of ways. In particular, the coarse grid solve, $A_0^{-1}\underline{z}_0 = \underline{r}_0$ say, that must be applied at each iteration can be undertaken either in parallel or sequentially. For the two level algorithm described here it is expected that this coarse grid problem will be relatively small and so the overhead of a parallel solve is likely to be prohibitive. Hence this problem is solved sequentially and, following the philosophy of [2], it is repeated on all processors rather than being

Elements/Procs.	$\varepsilon = 10^{-2}$				$\varepsilon = 10^{-3}$			
	2	4	8	16	2	4	8	16
6144	3	4	4	6	5	5	5	7
49152	3	4	4	6	4	5	5	7
393216	3	4	5	7	4	5	5	6
3145728	3	5	6	8	3	4	5	7

Table 2: The performance of the reduced AS preconditioner (2.1) on the streamline diffusion FE discretization of (2.3) for two choices of ε : figures quoted represent the number of iterations required to reduce the initial residual by a factor of 10^5 on a sequence of uniform refinements of an initial mesh of 768 tetrahedral elements.

undertaken on a single processor and the result broadcast, as in [7] for example.

In fact, for the parallel implementation used in this work it has been assumed that each subdomain is assigned to a single processor, so that the coarse grid solve may be combined with each of the local subdomain solves as follows. Each processor, i say, takes its own copy of the partitioned coarse mesh \mathcal{T}_0 but only refines this mesh inside $\tilde{\Omega}_{i,k-1}$ at step k of the refinement process (i.e. from \mathcal{T}_{k-1} to \mathcal{T}_k). The continuous piecewise linear FE spaces on the resulting meshes are then

$$\mathcal{U}_i = \mathcal{W}_0 \cup \tilde{\mathcal{W}}_i \quad (3.1)$$

for $i = 1, \dots, p$. Corresponding meshes are illustrated for a simple 2-d example in Figure 2. Note that in this figure there are a number of “slave” nodes in each processor’s mesh which cause these meshes to be non-conforming. The solution values at these nodes are not free: they are determined by the nodal values at the ends of the edges on which the slave nodes lie. For a practical implementation it turns out to be simpler to allow the solution values at these nodes to be free by performing an interior refinement of those elements on the unrefined sides of the edges that have “hanging” nodes on them. In the 2-d example of Figure 2 this simply involves bisecting all triangular elements containing a hanging node however for 3-d problems, such as those considered in this paper, this intermediate refinement is a little more complicated (see [14] for details of the “green” refinement stencils used here).

Having obtained a mesh on each processor it is possible for processor i to assemble the local FE problem for its mesh independently of the other processors. For a linear PDE this algebraic system may be expressed as

$$\begin{bmatrix} A_i & 0 & B_i \\ 0 & \bar{A}_i & \bar{B}_i \\ C_1 & \bar{C}_2 & A_{i,s} + \bar{A}_{i,s} \end{bmatrix} \begin{bmatrix} \underline{u}_i \\ \bar{\underline{u}}_i \\ \underline{u}_{i,s} \end{bmatrix} = \begin{bmatrix} \underline{f}_i \\ \bar{\underline{f}}_i \\ \underline{f}_{i,s} + \bar{\underline{f}}_{i,s} \end{bmatrix}. \quad (3.2)$$

Here \underline{u}_i are the unknowns inside Ω_i , $\bar{\underline{u}}_i$ are the unknowns outside $\bar{\Omega}_i$ and $\underline{u}_{i,s}$ are the unknowns on the interface of Ω_i . Correspondingly, $A_i, B_i, C_i, A_{i,s}$ and $\underline{f}_{i,s}$ are matrix and vector blocks assembled over elements inside Ω_i , and $\bar{A}_i, \bar{B}_i, \bar{C}_i, \bar{A}_{i,s}$ and $\bar{\underline{f}}_{i,s}$ are matrix and vector blocks assembled over elements outside Ω_i on processor i .

Note that in addition to the block notation used in (3.2) the stiffness matrix A may also be

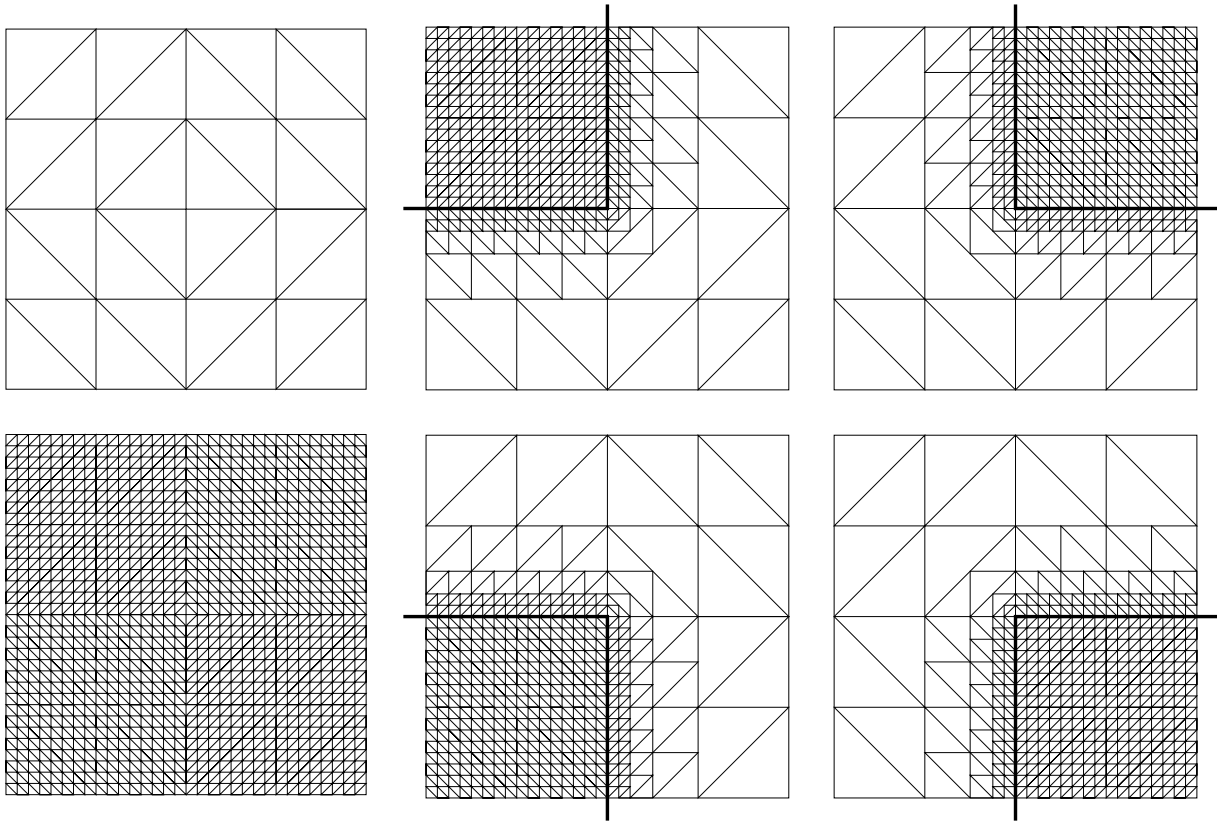


Figure 2: A 2-d illustration of a simple coarse mesh (top left) and a corresponding uniformly refined final mesh (bottom left) along with the actual meshes generated on four different processors (where the coarse mesh has been locally refined only inside $\tilde{\Omega}_{i,k}$ at level k of the refinement process).

written in a global block structure:

$$A = \begin{bmatrix} A_1 & & & B_1 P_1 \\ & \ddots & & \vdots \\ & & A_p & B_p P_p \\ P_1^T C_1 & \cdots & P_p^T C_p & A_I \end{bmatrix}, \quad (3.3)$$

where $A_I = \sum_{i=1}^p P_i^T A_{i,s} P_i$ for some Boolean prolongation matrices P_i^T which are used to extend a vector of nodal values on the interface of Ω_i to a larger vector of all subdomain interface values. Hence a matrix-vector product, $\underline{w} = A\underline{v}$ say, may be calculated in parallel as

$$\underline{w}_i = A_i \underline{v}_i + B_i P_i \underline{v}_I, \quad (3.4)$$

$$\underline{w}_I = \sum_{i=1}^p (P_i^T C_i \underline{v}_i + P_i^T A_{i,s} P_i \underline{v}_I) \quad (3.5)$$

$$= \sum_{i=1}^p P_i^T \underline{w}_{i,s} \quad \text{say.} \quad (3.6)$$

Similarly, a vector inner product may also be calculated efficiently in parallel using the blocks that are stored on each processor.

Hence the only remaining computationally significant task in the preconditioned GMRES algorithm to be considered is the parallel implementation of the preconditioner itself. The overall effect

Processors	$\varepsilon = 10^{-2}$					$\varepsilon = 10^{-3}$				
	1	2	4	8	16	1	2	4	8	16
Solution Time	770.65	484.53	347.61	228.39	136.79	688.12	442.44	277.78	187.16	108.75
Speed-Up	—	1.6	2.2	3.4	5.6	—	1.6	2.5	3.7	6.3

Table 3: Timings for the parallel solution using the stabilized FE discretization of (2.3) for two choices of ε : the solution times are quoted in seconds and the speed-ups are relative to the best sequential solution time.

of the modifications outlined at the start of this section are to yield a preconditioner that now takes the form

$$\hat{M}^{-1} = \sum_{i=1}^p \hat{\Pi}_i^T \hat{A}_i^{-1} \hat{R}_i, \quad (3.7)$$

where \hat{R}_i is the projection matrix from \mathcal{T}_J to the mesh stored on processor i , \hat{A}_i is the matrix on the left-hand side of (3.2) and $\hat{\Pi}_i^T$ is the simple prolongation matrix which combines Π_0^T and Π_i^T in (2.1). There are two main implementation issues associated with solving a preconditioning system of the form $\hat{M}\underline{z} = \underline{r}$ at each iteration in parallel. The first of these is the calculation of the restriction $\hat{R}_i \underline{r}$, which requires an all-to-one communication for *each* processor i . The second main issue that must be addressed is the solution of subproblems of the form $\hat{A}_i \underline{z} = \hat{R}_i \underline{r}$ on each processor. These systems are of the form (3.2) however and may therefore be solved simultaneously on each processor without the need for any further communication. The prolongation operations $\hat{\Pi}_i^T \underline{z}$ are extremely cheap to compute.

In Table 3 some parallel timings are presented for the solution of the convection-diffusion test problem (2.3) on the finest meshes used in Table 2 (i.e. using 3145728 tetrahedral elements). The subproblems of the form (3.2) have been solved sequentially (and only approximately, for improved overall efficiency) using an incomplete LU preconditioned GMRES solver ([10]) and this is also the solver used for the sequential times quoted. All timings are in seconds on a shared memory SG Origin2000 computer.

4 Discussion

There are a number of factors that limit the efficiency of the current parallel implementation. In all cases the sequential time of the AS preconditioned solver is greater than that of the sequential solver used (based upon ILU preconditioning, which does not naturally parallelize), so even a perfect parallel implementation would not deliver 100% efficiency. Furthermore, the mesh partitioning strategy used in this work is extremely basic and could certainly be improved to yield faster convergence and, in the case of sixteen or more subdomains, better load-balancing. Also, our decision to solve the coarse grid problem on every processor does result in a greater total amount of communication and, whilst much of this can be overlapped with useful computations, some additional parallel overhead is observed.

Overall however, in this paper the potential of the two level AS preconditioner considered has been demonstrated for the solution of a range of FE problems in three dimensions. Not only does the preconditioner yield optimal convergence rates for symmetric problems, as predicted in [4], but it also appears to perform optimally in its nonsymmetric reduced form and when applied to nonsymmetric model flow problems such as convection-diffusion equations. Furthermore, the practical parallel implementation of this approach has been successfully demonstrated and encouraging results have

been obtained on moderate numbers of processors. Further work is underway to better understand the effects of different mesh partitioning strategies so as to improve the parallel efficiency and to apply the technique as part of an adaptive strategy, as in [2] for example. Extensions to nonlinear equations and systems are also being investigated.

References

- [1] S.F. Ashby, T.A. Manteuffel and P.E. Taylor, “*A Taxonomy for Conjugate Gradient Methods*”, SIAM J. on Numerical Analysis, 27, 1542–1568, 1990.
- [2] R.E. Bank and M. Holst, “*A New Paradigm for Parallel Adaptive Meshing Algorithms*”, SIAM J. on Scientific Computing, 22, 1411-1443, 2000.
- [3] R.E. Bank and P.K. Jimack, “*A New Parallel Domain Decomposition Method for the Adaptive Finite Element Solution of Elliptic Partial Differential Equations*”, Concurrency and Computation: Practice and Experience, 13, 327–350, 2001.
- [4] R.E. Bank, P.K. Jimack, S.A. Nadeem and S.V. Nepomnyaschikh, “*A Weakly Overlapping Domain Decomposition Method for the Finite Element Solution of Elliptic Partial Differential Equations*”, SIAM J. on Scientific Computing, 23, 1818-1842, 2001.
- [5] X.-C. Cai and M. Sarkis, “*An Restricted Additive Schwarz Preconditioner for General Sparse Linear Systems*”, SIAM J. on Scientific Computing, 21, 792–797, 1999.
- [6] G.H. Golub and C.F. van Loan, “*Matrix Computations*”, John Hopkins Press, 3rd edition, 1996.
- [7] D.C. Hodgson and P.K. Jimack, “*A Domain Decomposition Preconditioner for a Parallel Finite Element Solver on Distributed Unstructured Grids*”, Parallel Computing, 23, 1157–1181, 1997.
- [8] C. Johnson “*Numerical Solution of Partial Differential Equations by the Finite Element Method*”, Cambridge University Press, 1987.
- [9] S.A. Nadeem “*Parallel Domain Decomposition Preconditioning for the Adaptive Finite Element Solution of Elliptic Problems in Three Dimensions*”, Ph.D. thesis, University of Leeds, 2001.
- [10] Y. Saad, “*SPARSKIT: A Basic Tool Kit for Sparse Matrix Computations, Version 2*”, Technical Report, Center for Supercomputing Research and Development, University of Illinois at Urbana-Champaign, Urbana, IL, USA, 1994.
- [11] Y. Saad, and M. Schultz, “*GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems*”, SIAM J. on Scientific Computing, 7, 856–869, 1986.
- [12] H.D. Simon, “*Partitioning of Unstructured Problems for Parallel Processing*”, Computing Systems in Engineering, 2, 135–148, 1991.
- [13] B. Smith, P. Bjorstad and W. Gropp, “*Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*”, Cambridge University Press, 1996.
- [14] W. Speares and M. Berzins, “*A 3-d Unstructured Mesh Adaptation Algorithm for Time-Dependent Shock Dominated Problems*”, International Journal for Numerical Methods in Fluids, 25, 81–104, 1997.