Domain Decomposition and Multigrid Algorithms
for Elliptic Problems on Unstructured Meshes

TONY F. CHAN AND BARRY F. SMITH

ABSTRACT. Multigrid and domain decomposition methods have proven to be versatile methods for the iterative solution of linear and nonlinear systems of equations arising from the discretization of PDEs. The efficiency of these methods derives from the use of a grid hierarchy. In some applications to problems on unstructured grids, however, no natural multilevel structure of the grid is available and thus must be generated as part of the solution procedure.

In this paper, we consider the problem of generating a multilevel grid hierarchy when only a fine, unstructured grid is given. We restrict attention to problems in two dimensions. Our techniques generate a sequence of coarser grids by first forming a maximal independent set of the graph of the grid or its dual and then applying a Cavendish type algorithm to form the coarser triangulation. Iterates on the different levels are combined using standard interpolation and restriction operators. Numerical tests indicate that convergence using this approach can be as fast as standard multigrid and domain decomposition methods on a structured mesh.

1. Introduction

Recently, unstructured meshes have become quite popular in large scale scientific computing [2], [10]. They have the advantage over structured meshes of the extra flexibility in adapting efficiently to complicated geometries and to rapid changes in the solution. However, this flexibility may come with a price.

Traditional solvers which exploit the regularity of the mesh may become less effi-
cient on an unstructured mesh. Moreover, vectorization and parallelization may
become more problematic. Thus, there is a need to adapt and modify current
solution techniques for structured meshes so that they can run as efficiently on
unstructured meshes.

In this paper, we will present some domain decomposition (DD) and multigrid
(MG) methods for solving elliptic problems on unstructured triangular meshes
in two space dimensions. These are among the most efficient algorithms for
solving elliptic problems. The application of multigrid methods to unstructured
grid problems have received some attention; see for example [10] and references
therein. There has been relatively little work on domain decomposition methods
for unstructured grid problems. Cai and Saad [4] considered overlapping domain
decomposition methods for general sparse matrices which in principle can be
applied to the stiffness matrices arising from discretizations of elliptic problems
on unstructured grids. However, if a coarse grid is to be used (often necessary
for fast convergence), it cannot be deduced from the algebraic structure of the
sparse matrix alone and geometric information about the coarse grid and the
associated interpolation operators must be supplied.

For multigrid and domain decomposition algorithms, a hierarchy of grids,
together with the associated interpolation and restriction operators, is needed.
For structured meshes, this grid hierarchy is naturally available and is indeed
exploited in these algorithms. For an unstructured mesh, however, the coarser
grids may not be given. Thus, a procedure is needed that generates this grid
hierarchy, as well as the associated interpolation and restriction operators. One
approach is to generate the coarser meshes independently, using a mesh gen-
erator, possibly the one which generated the fine mesh in the first place. This
approach has been used by Mavriplis [10], who constructed multigrid algorithms
for the Navier-Stokes equations on unstructured meshes in two and three space
dimensions. Another approach is to generate the grid hierarchy automatically
and directly from the given unstructured fine grid. This approach requires less
from the user because only the fine grid, on which the solution is sought, is
needed.

In this paper, we will follow the second approach. Our techniques generate
a sequence of coarser grids by first forming two maximal independent sets, one
for the interior vertices and the other for the boundary vertices, and then ap-
plying Cavendish's algorithm [6] to form the coarser triangulation. Thus, in this
approach, the coarse mesh vertices form a subset of the fine mesh vertices. We
also consider a variant in which this nested property of the vertices does not
hold. Iterates on the different levels are combined using standard finite element
interpolation and restriction operators. The mesh can be multiply-connected.
Numerical tests indicate that convergence using both coarsening approaches can
be as fast as standard multigrid and domain decomposition methods on a struc-
tured mesh.

For a recently developed convergence theory for two level Schwarz domain decomposition methods using nonnested coarse grids see Cai [5] and Chan, Smith and Zou [8].

2. Domain Decomposition and Multigrid Algorithms

Thus, we are interested in solving the following elliptic problem:

$$(2.1) \qquad -\nabla \cdot \alpha(x,y)\nabla u = f(x,y), \qquad \alpha(x,y) > 0,$$

on a 2D (not necessarily simply connected) region Ω with appropriate boundary conditions. We assume that Ω is triangulated into a fine grid, which can generally be unstructured and non-quasi-uniform, and a finite element or finite difference method is applied resulting in the algebraic system $Au = b$. The DD and MG algorithms we shall construct are used as preconditioners for A and are used in conjunction with a preconditioned Krylov subspace method.

We first discuss overlapping domain decomposition algorithms (a recent survey can be found in [7].) The fine grid Ω is decomposed into p overlapping subdomains $\Omega_i, i = 1, \cdots, p$, either as specified by the user or automatically determined by a mesh partitioning algorithm (in this paper, we will use exclusively the *recursive spectral bisection* (RSB) method of Pothen, Simon and Liou [11]). Associated with each Ω_i are restriction and extension operators R_i and R_i^T ($R_i u$ extracts the components of u corresponding to Ω_i and $R_i^T u_i$ is the zero-extension of an iterate u_i on Ω_i to Ω) and the local stiffness matrix $A_i \equiv R_i A R_i^T$. In order to achieve a good convergence rate, we will also use a coarse grid Ω_0. In this paper, we shall assume that Ω_0 is a proper triangular mesh itself which is not necessarily nested to Ω. We construct the associated interpolation operator R_H^T, which maps an iterate u_0 on Ω_0 to Ω, as follows. If a fine grid node lies within a triangle of Ω_0, we use linear interpolation to obtain its value, otherwise we set its value to zero. Once R_H^T is defined, the restriction operator R_H is defined to be its transpose. Finally, we compute the coarse grid stiffness matrix A_H by applying a piecewise linear finite element method to (2.1) on Ω_0. Note that in general $A_H \neq R_H A R_H^T$ due to the non-nestedness of the grids.

With these operators defined, we can now define the additive Schwarz preconditioner (which corresponds to a generalized block Jacobi method) as follows:

$$M_{as}^{-1} = R_H^T A_H^{-1} R_H + \sum_{i=1}^{p} R_i^T A_i^{-1} R_i.$$

Thus, each application of the preconditioner involves restricting the residual vector to each subdomain and performing a subdomain solve. In addition, a weighted restriction of the residual vector is computed on the coarse grid and inverted by a coarse grid solve. These local and coarse solutions are then mapped

back onto the fine mesh and added together to obtain the desired result. Multiplicative versions (i.e. Gauss-Seidel) can also be defined analogously given an ordering of the subdomains.

Multilevel preconditioners (including classical multigrid methods) are closely related to domain decomposition methods and their implementations can be treated in the same framework. A grid hierarchy is needed and the associated interpolation and restriction operators can be defined in an analogous way. For example, let the fine grid be level 1 and the coarsest grid level l. Let R_i denote the restriction operator from level 1 to level i and the transpose R_i^T the corresponding interpolation operator. Then an additive multilevel preconditioner can be written in the following form:

$$M_{aml}^{-1} = \sum_{i=1}^{l} R_i^T S_i R_i,$$

where S_i is a "smoother" on level i. For instance, for multi-level diagonal scaling, S_i is simply the inverse of the diagonal of the stiffness matrix on level i. Classical V-cycle MG methods can be viewed as symmetrized multiplicative versions of the above preconditioner. Note that in practice the action of R_i and R_i^T are computed via a recursion using mappings between *adjacent* grid levels.

In our implementation of the domain decomposition algorithms, the coarse grid Ω_0 is obtained by a sequence of recursively applied coarsening steps (see next section), and hence the grid hierarchy is naturally defined for performing the multigrid iteration as well. Of course, this is not the only way to construct the coarse grid for a DD method. For example, one can use the subdomains to directly construct a coarse grid without going through a grid hierarchy. We shall not pursue these other possibilities in this paper.

3. Construction of the Grid Hierarchy

In this section, we will describe our techniques for constructing the coarse grid hierarchy, as well as the associated interpolation and restriction operators, directly from the given unstructured fine mesh. It suffices to describe this for one coarse level because the procedure can be recursively applied to obtain all the coarse meshes.

We shall need the notion of a *maximal independent set* of the vertices of a graph. A subset of vertices V of a graph G is said to be *independent* if no two vertices of V are connected by an edge. V is said to be *maximally independent* if adding any additional vertex to it makes it dependent. Note that maximal independent sets of vertices of a graph are generally not unique.

The procedure has four steps:

    (i) Form a maximally independent set of the boundary vertices and from these construct a set of coarse boundary edges,

(ii) Form a maximally independent set of the interior vertices,

(iii) Apply a Cavendish type algorithm [6] to triangulate the resulting collection of coarse boundary edges and coarse interior vertices,

(iv) Construct the interpolation and restriction operators.

Step (i) is fairly straightforward. For each disjoint boundary segment, the boundary vertices are ordered say in a clockwise direction, starting with a random vertex. Then every other vertex is thrown out and the remaining ones are connected with new coarse boundary edges. This forms a coarse representation of the boundary segment. After several coarsenings, one may find that the boundary is no longer qualitatively similar to the original boundary. This may be prevented by simply retaining some of the vertices in the coarse grid boundary that would normally be dropped.

Step (ii) uses a greedy wavefront type algorithm. A random interior vertex is selected for inclusion in the maximally independent set. Then every interior vertex connected to it is eliminated from consideration for inclusion in the maximally independent set. Next, one of the interior vertices connected to the newly eliminated vertices is selected for inclusion and the procedure repeats until all interior vertices have been considered. An algorithm similar to this has been used by Barnard and Simon [1] in designing graph partitioning algorithms. This procedure can be implemented in linear time, i.e. proportional to the total number of interior vertices.

The input to Step (iii) is thus a collection of coarse boundary edges and coarse interior vertices. A version of Cavendish's algorithm [6] is then applied to triangulate this collection. This algorithm is an advancing front technique and "grows" new triangles from those already built by selecting an interior vertex to be "mated" to an existing edge. In doing so, it tries to optimize the aspect ratio of the new triangle formed, preferring those that are close to being equilateral. It is possible to implement this algorithm in linear time, i.e. proportional to the number of interior vertices, but our current implementation is not optimal.

Finally, in Step (iv), the interpolation operator is constructed in the form of a sparse matrix and stored. To determine the entries of this interpolation matrix, the coarse triangles are taken in sequence and the entries corresponding to all the fine grid vertices within the coarse triangle are then computed using the standard piecewise linear interpolation. This procedure can also be implemented in linear time because the fine grid triangles close to the vertices of the coarse triangle (which are also fine grid vertices as well) can be found by a local search. We emphasize that this is not possible if the coarser grids are generated completely independently. The restriction matrix is then just the transpose of the interpolation matrix. Clearly higher order finite elements may also be used; then the interpolation would be piecewise polynomial and could still be calculated in a local manner and hence remain a linear time algorithm.

In the alternative variant only Step (ii) is changed. We consider a candidate coarse vertex at the center of each element. Those that are adjacent to

the selected boundary nodes are then eliminated. We then construct a maximal independent set of the remaining candidate vertices using the same approach as indicated above. This procedure is equivalent to calculating a maximal independent set of the *dual* graph of the mesh. It is important to note that the coarse grid vertices generated in this manner will not lie in the same location as any fine grid nodes. In our experiments on a few sample grids, the number of coarse grid nodes generated using this alternative approach is slightly more than that generated using the first approach outlined.

4. Numerical Results

All the numerical experiments were performed using the Portable, Extensible Toolkit for Scientific Computation (PETSc) of Gropp and Smith, [9] running on a Sun SPARC 10.

We will report numerical results for solving the Poisson equation on three different unstructured triangular meshes. All meshes are enclosed in the unit square. Two kinds of boundary conditions are used: (1) homogeneous Dirichlet, or (2) a *mixed* condition: if $x > .2$, a homogeneous Neumann boundary condition is imposed, otherwise homogeneous Dirichlet is imposed. We use piecewise linear finite elements for the discretization and we solve the resulting systems of linear equations by either a V-cycle multigrid method (with a pointwise Gauss Seidel smoother, using 2 pre and 2 post smoothing sweeps per level) or an overlapping Schwarz domain decomposition method. In all cases, the discrete right hand side is chosen to be a vector of all 1's and the initial iterate set to zero. Both the MG and DD methods are used as preconditioners, with full GMRES [12] as an outer accelerator. The iteration is stopped when the l_2 norm of the residual has been reduced by a factor of 10^{-6}. Our goal is to compare the performance of our versions of domain decomposition and multigrid algorithms on unstructured meshes to that of the same algorithms on similar *structured* meshes. For this purpose, we also use two structured meshes in our experiments (a uniform mesh on a square and on an annulus).

Table 1 shows the number of MG iterations for the *Eppstein* mesh, [3], shown in Figure 1, a relative small quasi-uniform unstructured mesh on the unit square. Figures 2 and 3 show the coarser meshes using regular and dual graph coarsening. These results should be compared to those for a uniform square mesh on the unit square in Table 2, because the two meshes are topologically similar. We see that although the performance of our MG algorithm on the unstructured mesh is slightly higher than that on the structured square mesh, its performance is quite satisfactory, for both types of boundary conditions.

Next, we look at a more realistic mesh, the *Airfoil* mesh, (from T. Barth and D. Jesperson of NASA Ames), shown in Figure 4. The coarse meshes are shown in Figures 5, 6 and 7. One may note that several poorly shaped triangles

FIGURE 1. The *Eppstein* mesh: 547 nodes

FIGURE 2. The *Eppstein* mesh: level 2. regular coarsening (left) and dual graph coarsening (right)


FIGURE 3. The *Eppstein* mesh: level 3. regular coarsening (left) and dual graph coarsening (right)

T. F. CHAN AND B. F. SMITH

TABLE 1. MG iterations for the *Eppstein* mesh, 547 nodes

MG Levels	Regular coarsening		Dual graph coarsening	
	Dir. B.C.	Mixed B.C.	Dir. B.C.	Mixed B.C.
2	4	4	3	4
3	4	5	4	6

TABLE 2. MG iterations for the uniform *Square* mesh, 4225 nodes

| MG Levels	Nodes	Dir. B.C.	Mixed B.C.
2	1089	3	3
3	289	3	3
4	81	3	3

are generated on the coarsest grid. These do not seem to seriously effect the convergence rate. In theory, these bad elements could be adjusted during a "cleanup" pass over the mesh, after the Cavendish algorithm was applied. We see that the performance for the Dirichlet boundary condition cases, given in Table 3, are quite comparable to that for the *Eppstein* mesh but is noticeably worse for the mixed boundary conditions. Note there is no difference in the performance for both types of coarsening. The deterioration in performance for the mixed boundary condition case is due to the fact that the coarser grid domains do not completely cover the portion of the fine grid boundary on which Neumann boundary conditions are applied. In Chan, Smith and Zou [8] it is shown that to obtain an optimal convergence rate the coarser grid must completely cover the Neumann part of the fine grid boundary. We also compare the performance to that on a quasi-uniform mesh on an annulus region; see Figure 4 and Table 4. Due to our treatment of the curved boundary, the fine grid boundary is not covered by the coarser grids and hence poor convergence results. Overall, the performance of our MG algorithm on the unstructured *Airfoil* mesh is comparable or better than on the *Annulus* mesh.

In Table 5, we show results for a larger unstructured mesh around an airfoil, namely the *Barth* mesh, (from T. Barth of NASA Ames) shown in Figure 9. The levels 2 and 4 coarse meshes are shown in Figure 10 (not all the grid points are shown). We only include the Dirichlet boundary condition results. We can observe that the MG performance is quite comparable to the other unstructured meshes (i.e. *Airfoil, Eppstein*) and the structured meshes (i.e. *Square, Annulus*).

TABLE 3. MG iterations for the *Airfoil* mesh, 4253 nodes

MG Levels	Regular coarsening			Dual graph coarsening		
	Nodes	Dir. B.C.	Mixed B.C.	Nodes	Dir. B.C.	Mixed B.C.
2	1180	4	8	1507	4	8
3	518	4	9	328	4	9
4	89	4	10	171	5	10

FIGURE 4. The *Airfoil* mesh: 4253 nodes

FIGURE 5. The *Airfoil* mesh: Level 2. Regular (left) and dual graph coarsening (right).

TABLE 4. MG iterations for the *Annulus* mesh, 2176 nodes

| MG Levels	Nodes	Dir. B.C.	Mixed B.C.
2	576	4	18
3	160	5	18
4	48	5	18

FIGURE 6. The *Airfoil* mesh: Level 3. Regular (left) and dual graph coarsening (right).

FIGURE 7. The *Airfoil* mesh: Level 4. Regular (left) and dual graph coarsening (right).

FIGURE 8. The *Annulus* mesh: 2176 nodes

TABLE 5. MG iterations for the *Barth* mesh, 6691 nodes

MG Levels	Regular coarsening		Dual graph coarsening	
	Nodes	Dir. B.C.	Nodes	Dir. B.C.
2	1614	5	1810	5
3	405	6	574	6
4	112	7	189	6

Again, both coarsening strategies work equally well.

Finally, we show in Table 6 the results for the multiplicative version of our domain decomposition algorithm on the *Airfoil* mesh. The column labeled *overlap* refers to the number of fine grid elements that are extended from each subdomain into the interior of its neighbors. Thus, an overlap of 0 means there is no overlap at all. The column labeled *Level of coarse grid* refers to which level of the grid hierarchy is used as the coarse grid in the DD algorithm. The 16 subdomains computed by the Recursive Spectral Bisection method are shown in Figure 11. We can make several observations from the results. First, the use of a coarse grid reduces the number of iterations significantly. Second, the use of some overlap is very cost effective but the number of iterations levels off quickly as the overlap increases. These results for overlapping Schwarz are also very similar to those obtained for structured grids.

FIGURE 9. The *Barth* mesh: 6691 nodes

FIGURE 10. The *Barth* mesh: Level 2, 1614 nodes (left) Level 4, 112 nodes (right)

TABLE 6. Multiplicative DD iterations for the *Airfoil* mesh. 16 Subdomains

Overlap (no. elements)	Level of coarse grid	Regular coarsening	Dual graph coarsening
0	None	56	56
0	4	21	22
0	3	15	12
1	None	16	16
1	4	10	10
1	3	7	7
2	None	14	14
2	4	8	8
2	3	5	5

FIGURE 11. The *Airfoil* mesh: 16 subdomains computed by RSB

5. Summary

In summary, we have constructed domain decomposition and multigrid algorithms for solving elliptic problems on general unstructured meshes, which in our limited experience perform nearly as well as these algorithms would perform on similar structured meshes. Only the fine mesh is needed and all auxiliary components of the algorithms, such as the coarse grid hierarchy, the interpolation operators, and the domain partitioning, are computed automatically. The algorithms can in principle be extended to three space dimensions and to indefinite, non-self-adjoint and higher order problems.

REFERENCES

1. S.T. Barnard and H.D. Simon. A fast multilevel implementation of recursive spectral bisection. *Proc. 6th SIAM Conf. Parallel Proc. for Sci. Comp.*, pp. 711–718, 1993.
2. H. Deconinck and T.J. Barth. Special course on unstructured grid methods for advection dominated flows. *Lecture Series, Von Karman Inst., Belgium*, AGARD-R-787, March 1992.
3. M. Bern, D. Eppstein and J. Gilbert. Provably good mesh generation. *31st Annual Symp. on Foundations of Computer Science*, IEEE, pp. 231-241, 1990.
4. X.-C. Cai and Y. Saad. Overlapping domain decomposition algorithms for general sparse matrices. Technical Report 93-027, Univ. of Minn., AHPCRC, March 1993.
5. X.-C. Cai. A non-nested coarse space for Schwarz type domain decomposition methods. Technical Report: Department of Computer Science, University of Colorado at Boulder, November 1993.
6. J. C. Cavendish. Automatic triangulation of arbitrary planar domains for the finite element method. *Int'l J. Numer. Meth. Engineering*, 8:679–696, 1974.
7. T. F. Chan and T. Mathew. Domain decomposition algorithms. *Acta Numerica*, pages 1–83, 1994.
8. T. F. Chan and B. F. Smith and J. Zou. Overlapping Schwarz methods on unstructured grids using non-matching coarse grids. Technical Report: CAM 94-08, UCLA Department of Mathematics, Febuary 1994.
9. W. D. Gropp and Barry F. Smith. Portable, Extensible Toolkit for Scientific Computation (PETSc), Available via anonymous ftp at info.mcs.anl.gov in the directory pub/pdetools.
10. D. J. Mavriplis. Unstructured mesh algorithms for aerodynamic calculations. Technical Report 92-35, ICASE, NASA Langley, Virginia, July 1992.
11. A. Pothen, H. Simon, and K. P. Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Mat. Anal. Appl.*, 11:430–452, 1990.
12. Y. Saad and M. H. Schultz. GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Statist. Comput.*, 7:856–869, 1986.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, CA 90024-1555

E-mail address: chan@math.ucla.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, CA 90024-1555

E-mail address: bsmith@math.ucla.edu