13

Additive Schwarz, CG and Discontinuous Coefficients

I.G. Graham and M.J. Hagger

1 Introduction

This paper is concerned with the performance of the conjugate gradient(CG) method with additive Schwarz preconditioner for computing unstructured finite element approximations to the elliptic problem

$$\nabla . a \nabla u = f, \text{ on } \Omega, \quad u = g \text{ on } \partial\Omega_D, \quad \frac{\partial u}{\partial n} = \tilde{g} \text{ on } \partial\Omega_N. \tag{1}$$

Here $\Omega \subset \mathbb{R}^3$ is a polyhedral domain with boundary $\partial\Omega$ partitioned into disjoint subsets $\partial\Omega_D \neq \emptyset$ and $\partial\Omega_N$, each of which is composed of unions of polygons, and f, g and $\tilde{g}$ are suitably smooth given data. (Analogous results also hold in 2D.) We also assume that a is piecewise constant on each of d open disjoint polyhedral regions Λ_k, such that $\cup_{k=1}^{d} \bar{\Lambda}_k = \bar{\Omega}$, and we write $a|_{\Lambda_k} = a_k$ where each $a_k \in \mathbb{R}_+ := (0, \infty)$ is constant. We have in mind that the regions Λ_k of different material properties are fixed but may have complicated geometry and that the overall mesh used to compute u accurately will be finer than the geometry of the Λ_k. There are many applications of this type of problem, for example in groundwater flow and in electromagnetics.

After discretisation with linear finite elements on a triangulation $\mathcal{T}$ of Ω, (1) reduces to the SPD system

$$K(\mathbf{a})\mathbf{x} = \mathbf{b}(\mathbf{a}), \tag{2}$$

where the stiffness matrix and load vector depend continuously on $\mathbf{a} \in \mathbb{R}_+^{\mathrm{d}}$. Let h denote the diameter of $\mathcal{T}$ and $\mathcal{J} = \max_{k,l}\{a_k/a_l\}$. It is a standard result that $K(\mathbf{a})$ is ill-conditioned in the sense that (under suitable assumptions) $\kappa(K(\mathbf{a})) = O(h^{-2})$ as $h \to 0$ for fixed $\mathbf{a}$ and $\kappa(K(\mathbf{a})) = O(\mathcal{J})$ as $\mathcal{J} \to \infty$ for fixed h. (Here κ denotes the 2-norm condition number.) One of the striking successes of domain decomposition methods has been the construction of preconditioners for which the condition number of the preconditioned problem is bounded independently of both h and $\mathbf{a}$. We refer to these two properties as "*h-optimality*" and "*$\mathbf{a}$-optimality*" respectively. For a review

Ninth International Conference on Domain Decomposition Methods
Editor Petter E. Bjørstad, Magne S. Espedal and David E. Keyes

of the many papers on this subject, see [CM94] or [DSW96]. As far as we are aware all of these results assume meshes are obtained by structured refinement from a coarse grid (or substructures) which resolve the coefficient jumps.

In many large-scale computations, unstructured meshes are obtained by mesh generators, subdomains are obtained from mesh partitioning codes, and coarse grids are obtained by some coarsening strategy. In this context it may be difficult to implement the preconditioners covered by this theory. Recently the theory has been substantially extended to unstructured meshes - [CSZ96, CGZ96] and the references therein, and results about the h-optimality (but not the $\mathbf{a}$-optimality) of the corresponding preconditioners have been obtained. Indeed it is possible to construct some counter examples to $\mathbf{a}$-optimality when the coarse grid does not resolve the discontinuity in $\mathbf{a}$ ([GH96]). Since it seems rather unnatural to consider *unstructured* grids which are restrained to *resolve* the discontinuity in $\mathbf{a}$, this appears to suggest that unstructured grids may be bad for this type of problem. However, there is much empirical evidence to suggest the CG method remains robust to discontinuities in $\mathbf{a}$ even when they are not resolved by the discretisation and preconditioning process. In this work we prove a result which explains this phenomenon. It is obtained not by examining the condition number of the preconditioned matrix (which may be very bad) but by obtaining bounds on its eigenvalues (which, except for a small number of outliers, turn out to be very well behaved). We only have room here for a statement of our results and an idea of the proof. The necessary details are in [GH96].

2 Theoretical Results

We present here our results for the unstructured multilevel additive Schwarz preconditioner proposed in [CGZ96]. This is more general than the result in [GH96] which is about the two-level variant ([CSZ96]), but the proof of both results is identical. Let $\{\mathcal{T}^l\}_{l=0}^{Q}$ be a shape-regular sequence of triangulations of Ω with diameters $h^0 > h^1 > ... > h^Q = h$, where $\mathcal{T}^Q = \mathcal{T}$ is the fine mesh on which (1) is discretised. We assume in this section that this fine mesh resolves the interfaces between the regions Λ_k. We remove this restriction in §3. Assume that for each l, Ω is partitioned into non-overlapping subdomains $\tilde{\Omega}_j^l, j = 1, ..., s_l$ which are then extended to overlapping subdomains Ω_j^l with $\delta^l := \min_j \operatorname{dist}(\partial\Omega_j^l, \partial\tilde{\Omega}_j^l) > 0$, and are such that $\partial\Omega_j^l \cap \bar{\Omega}$ contains only boundaries of tetrahedra of $\mathcal{T}^l$. More general meshes are permissible - see [CGZ96] for technical details. Let $\mathcal{N}^l$ denote the degrees of freedom in $\mathcal{T}^l$ and set $\mathcal{N}_j^l = \mathcal{N}^l \cap \Omega_j^l$. For any set of nodes $\mathcal{S}$ let $[\mathcal{S}]$ denote the space of nodal vectors on $\mathcal{S}$. Then (2) is to be solved for $\mathbf{x} \in [\mathcal{N}^Q]$. For each l and j we introduce a prolongation ${R_j^l}^T : [\mathcal{N}_j^l] \to [\mathcal{N}^Q]$ as follows: For $\mathbf{x} \in [\mathcal{N}_j^l]$, we form the piecewise linear function with value x_p at nodes $p \in \mathcal{N}_j^l$ and 0 elsewhere. Then form the piecewise linear interpolant of this with respect to $\mathcal{T}^Q$. The nodal values of this on $\mathcal{N}^Q$ are called ${R_j^l}^T\mathbf{x}$. The operator $R_j^l : [\mathcal{N}^Q] \to [\mathcal{N}_j^l]$ is defined to be the adjoint of ${R_j^l}^T$. Note that in the case $l = Q$, ${R_j^Q}^T$ is just the straightforward extension by 0 from $[\mathcal{N}_j^Q]$ to $[\mathcal{N}^Q]$. The multilevel preconditioner $M(\mathbf{a})$ is then defined by the action of its inverse:

$$M(\mathbf{a})^{-1} = \sum_{l=0}^{Q} \sum_{j=1}^{s_l} {R_j^l}^T (K(\mathbf{a})_j^l)^{-1} R_j^l, \qquad (3)$$

where $K(\mathbf{a})_j^l := R_j^l K(\mathbf{a}) {R_j^l}^T$. In [CGZ96] it is proved that

$$\kappa(M(\mathbf{a})^{-1}K(\mathbf{a})) \leq C(\mathbf{a})Q^2 \max_{1 \leq l \leq Q} \{(h^l + h^{l-1})/\delta^l\}^2, \qquad (4)$$

and so, with the appropriate choice of overlap, δ^l, we have h-optimality. The **a**-optimality is an open question since the behaviour of $C(\mathbf{a})$ is unknown.

Rather than trying to analyse $C(\mathbf{a})$ which (by counterexamples in [GH96]) cannot be expected to be bounded in **a** without further assumptions, we instead obtain a bound on the number of preconditioned CG iterations needed to solve (2) as **a** varies. To do this we characterise the number of CG iterations in terms of the asymptotics of sequences of coefficients $\{\mathbf{a}^{(m)}\}_{m=1}^{\infty} \subset \mathbb{R}_+^{\mathrm{d}}$. To avoid uninteresting pathologies these are required to satisfy two mild assumptions: Firstly we require that for all k, l, $a_l^{(1)} \geq a_k^{(1)}$ implies $a_l^{(m)} \geq a_k^{(m)}$ for all $m \geq 1$; Secondly for all k, l such that $\bar{\Lambda}_k \cap \bar{\Lambda}_l \neq \emptyset$ we require that the ratio $a_k^{(m)}/a_l^{(m)}$ either approaches 0, ∞ or remains in a compact subset of $\mathbb{R}_+$ as $m \to \infty$. These assumptions are consistent with the typical applications of (1) - see [GH96].

To state our theorem, let n denote the dimension of $K(\mathbf{a})$ and let $\lambda_1^{(m)} \leq ... \leq \lambda_n^{(m)}$ denote the eigenvalues of the preconditioned matrix $M(\mathbf{a}^{(\mathbf{m})})^{-1}K(\mathbf{a}^{(\mathbf{m})})$. For any integer $0 \leq L \leq n-1$, set $\kappa_{L+1}^{(m)} = \lambda_n^{(m)}/\lambda_{L+1}^{(m)}$ and $\mathcal{J}^{(m)} = \max_{k,l}\{a_k^{(m)}/a_l^{(m)}\}$. Let $\mathbf{x}^j$ be the jth preconditioned CG iterate for (2), and let $\|\,.\,\|_m$ denote the energy norm induced by $K(\mathbf{a}^{(m)})$.

Theorem 1. There is an integer L and a constant C which are independent of m, h^l and δ^l such that for each $\epsilon > 0$

$$\frac{\| \mathbf{x} - \mathbf{x}^j \|_m}{\| \mathbf{x} - \mathbf{x}^0 \|_m} \leq \epsilon, \text{ provided } j \geq L + \sqrt{\kappa_{L+1}^{(m)}} \left\{ \log \frac{2}{\epsilon} + L \log \frac{C}{\lambda_1^{(m)}} \right\}. \qquad (5)$$

In addition $(\lambda_1^{(m)})^{-1} = O(\mathcal{J}^{(m)})$ and $\kappa_{L+1}^{(m)}$ is bounded as $\mathcal{J}^{(m)} \to \infty$, for fixed h.

Thus, apart from an additional L iterates, the number of iterates to obtain a fixed tolerance grows only logarithmically in $\mathcal{J}^{(m)}$ as $m \to \infty$. This is to be compared with the $O(\mathcal{J}^{(m)})$ growth which the condition number $\kappa_1^{(m)} = \kappa(M(\mathbf{a}^{(m)})^{-1}K(\mathbf{a}^{(m)}))$ can experience when the coarse mesh does not resolve the discontinuity in **a** [GH96].

A key question is the size of L: The answer from [GH96] is that L can never be larger than the maximal number of components of sets formed by taking unions of the Λ_k. But in many cases L is known to be smaller. The size of L depends on the *limiting form* of $\mathbf{a}^{(m)}$ as $m \to \infty$ (but not on m). A rigorous definition of L is given in [GH96]. From this it follows, for example, that if each of the Λ_k touches $\partial\Omega_D$ then $L = 0$, $\kappa_1^{(m)}$ is bounded independently of m, and the preconditioning is **a**-optimal. More dramatically, suppose Ω is a square divided into a chequer-board of any number of square regions Λ_k, which are coloured alternately red and black. Suppose that

$a_k^{(m)} \to \infty$ on red squares while $a_k^{(m)} \to 0$ on black squares, then $L = 1$ at most. If any of the red squares touches $\partial\Omega_D$ then $L = 0$ and the preconditioner is again **a**-optimal. We emphasise that these results do *not* require that the subdomains Ω_j^l on any of the levels have any relationship to the regions Λ_k on which a is constant.

As an example of the worst case, if $a_k^{(m)} \to \infty$ on L_0 of the regions Λ_k which do not touch $\partial\Omega_D$ and do not touch each other, and $a_k^{(m)}$ is bounded on the other regions then $L = L_0$ in the theorem.

Sketch of Proof. The proof is obtained in three stages. First it is shown that, for each k, the kth smallest eigenvalue of $M(\mathbf{a}^{(m)})^{-1}K(\mathbf{a}^{(m)})$ can be bounded below in terms of the kth smallest eigenvalue of the diagonally scaled matrix $(\text{diag}K(\mathbf{a}^{(m)}))^{-1}K(\mathbf{a}^{(m)})$ (or, its equivalent symmetric version $S(\mathbf{a}^{(m)}) := (\text{diag}K(\mathbf{a}^{(m)}))^{-1/2}K(\mathbf{a}^{(m)})(\text{diag}K(\mathbf{a}^{(m)}))^{-1/2}$). This is done by a routine application of the fact that these preconditioned matrices can be written as sums of orthogonal projections onto certain subspaces.

The second and substantial part of the proof is a characterisation of the spectrum of $S(\mathbf{a}^{(m)})$ in terms of the asymptotics of the maximum jumps of $\mathbf{a}^{(m)}$. The result is that only a fixed number L of eigenvalues of $S(\mathbf{a}^{(m)})$ may approach zero as these jumps worsen. This number depends on the geometry of the Λ_k but not on the mesh or the values of the coefficients. Some examples of the size of L have already been given above. The rest of the eigenvalues are bounded above and below by positive numbers independent of the size of the jumps. The same statement then holds for the reduced condition number $\kappa_{L+1}^{(m)}$. The third and final stage is to use a well-known extension of the convergence theory of the conjugate gradient method for the case of outlying clusters of bad eigenvalues. Full details are in [GH96].

Remark 1. An analogous estimate to (5) holds true with $\| \,.\, \|_m$ replaced by the Euclidean norm $\| \,.\, \|_2$, but inside the braces on the right-hand side of (5) we must add the term $\frac{1}{2}\log(\kappa\mathcal{J}^{(m)})$ where κ is the condition number of $K(\mathbf{1})$. This is of practical interest since the coefficient-dependent energy norm, $\| \,.\, \|_m$, is not a good place to measure the error if $a_k^{(m)} \to 0$ on some Λ_k.

Remark 2. An analogous clustering effect of preconditioners for problems of type (1) (for a more restrictive class of coefficient variations) is obtained [CNT96] and also leads to logarithmic estimates for the growth in the number of conjugate gradient iterates as the discontinuity worsens. However, the preconditioner proposed there essentially requires the solution of the global Laplace operator $a \equiv 1$ on Ω, which is implemented, for example, by embedding Ω into a rectangular grid and using fast Poisson solvers there. By contrast our present results concern standard additive Schwarz preconditioners widely used in domain decomposition methods.

3 An extension of the theory

The theory described above does not require that any of the subdomains (either at the finest level or any of the coarser levels) resolve the discontinuity in **a**. However, it *does* assume that the fine mesh itself should resolve this discontinuity. This assumption simplifies the analysis of the diagonally scaled matrix, on which the theory depends, but we shall show here that it is not necessary. This generalisation has obvious practical importance since in the case of very irregular regions Λ_k, it may not even be reasonable to expect the fine grid to resolve the discontinuities.

In the interests of brevity we shall not give here a completely general extension, but instead restrict to the case of a two coefficient problem where $\Lambda_1 \subset \Omega$ and $\Lambda_2 = \Omega \backslash \bar{\Lambda}_1$. Suppose the interface $\Gamma = \bar{\Lambda}_1 \cap \bar{\Lambda}_2$ is continuous and lies entirely inside Ω. As before $a|_{\Lambda_k} = a_k \in \mathbb{R}_+$, $k = 1, 2$. Let $\mathcal{T}$ be a mesh of tetrahedra on Ω which do not need to resolve Γ. Let $\mathcal{N}$ be the nodes of $\mathcal{T}$ which are not on $\partial\Omega_D$. For $p \in \mathcal{N}$, set $\mathcal{T}(p) = \{T \in \mathcal{T} : p \in T\}$. Then for any $\mathbf{a} \in \mathbb{R}_+^2$ and $p, q \in \mathcal{N}$, $K(\mathbf{a})_{pq} = \sum_{T \in \mathcal{T}(p) \cap \mathcal{T}(q)} a_T (K_T)_{pq}$, where K_T is the piecewise linear element stiffness matrix corresponding to the Laplace operator on T and $a_T = (\int_T a)/|T|$, where $|T|$ is the volume of T. Consider the symmetrically diagonally scaled matrix $S(\mathbf{a}) = (\mathrm{diag} K(\mathbf{a}))^{-1/2} K(\mathbf{a}) (\mathrm{diag} K(\mathbf{a}))^{-1/2}$ and a sequence of coefficients $\mathbf{a}^{(m)}$ with constant values on the Λ_k, represented by $\{\mathbf{a}^{(m)}\} \subset \mathbb{R}_+^2$. For $T \in \mathcal{T}$ set $a_T^{(m)} = \int_T a^{(m)}/|T|$. The extreme behaviour of the spectrum of $S(\mathbf{a}^{(m)})$ can be characterised by considering the limit of $S(\mathbf{a}^{(m)})$ as $m \to \infty$. It is easy to see that $S(\mathbf{a}^{(m)})_{pq}$ is independent of m unless $p, q \in T \cap \mathcal{N}$ for some $T \in \mathcal{T}$ with $p \neq q$ and $\mathcal{T}(p) \cup \mathcal{T}(q)$ intersects both Λ_1 and Λ_2 in sets of positive volume. Then, for $T \in \mathcal{T}(p) \cup \mathcal{T}(q)$ we have $a_T^{(m)} = (|T \cap \Lambda_1|/|T|)\, a_1^{(m)} + (|T \cap \Lambda_2|/|T|)\, a_2^{(m)}$. If $\{\mathbf{a}^{(m)}\}$ is constrained to satisfy the mild assumptions of §2 then we need only consider the two cases $\lim_{m\to\infty} \{a_1^{(m)}/a_2^{(m)}\} = \infty$ or 0.

Consider $a_1^{(m)}/a_2^{(m)} \to \infty$. Introduce the slight extension of Λ_1: $\tilde{\Lambda}_1 = \cup\{T : |T \cap \Lambda_1| \neq \emptyset\}$ with extended interface $\tilde{\Gamma} = \partial\tilde{\Lambda}_1$, and the matrix $\tilde{K}_1 = \sum_{T \in \mathcal{T}} (|T \cap \Lambda_1|/|T|) K_T$, which corresponds to a Neumann problem on $\tilde{\Lambda}_1$ for an operator of the form (1) with coefficient $|T \cap \Lambda_k|/|T|$ on each T. Let $\tilde{K}_2$ be the finite element matrix on $\Omega \backslash \tilde{\Lambda}_1$ corresponding to Dirichlet condition on $\tilde{\Gamma}$ and given boundary conditions on $\partial\Omega$. Then $\lim_{m\to\infty} S(\mathbf{a}^{(m)}) \to \tilde{S}$, where $\tilde{S}$ is the diagonally scaled version of the block diagonal matrix $\mathrm{diag}(\tilde{K}_1, \tilde{K}_2)$. $\tilde{S}$ has a single zero eigenvalue with all other eigenvalues positive (but depending on h). So $S(\mathbf{a}^{(m)})$ has a single eigenvalue approaching zero as $m \to \infty$ and Theorem 1 holds with $L = 1$. If $a_1^{(m)}/a_2^{(m)} \to 0$, then $S(\mathbf{a}^{(m)})$ also approaches a diagonally scaled version of a matrix of the general form $\mathrm{diag}(\tilde{K}_1, \tilde{K}_2)$. But here $\tilde{K}_2$ is a stiffness matrix on a slight extension, $\tilde{\Lambda}_2$, of Λ_2 with Neumann condition on $\partial\tilde{\Lambda}_2 \backslash \partial\Omega$ and given mixed conditions on $\partial\Omega$. $\tilde{K}_1$ is a matrix corresponding to a Dirichlet problem on $\Omega \backslash \tilde{\Lambda}_2$. This time $\tilde{S}$ has all positive eigenvalues and Theorem 1 holds with $L = 0$.

An extension of this argument to many coefficients will lead to the proof of Theorem 1, in the case when the fine grid does not resolve the discontinuity. The extension to $\|\,.\,\|_2$ mentioned in Remark 1 also holds by the arguments in [GH96].

4 A Numerical Example

To demonstrate our results we consider a two-dimensional problem with geometry motivated by electromagnetic field computations, see [EST94] and the references contained therein. Here there are three interior regions with varying material properties, as in Figure 1(a). The domain is the unit square and the regions with differing material properties are given by $\Lambda_1 = (0.44, 0.56) \times (0.81, 0.94)$, $\Lambda_2 =$

$(0.44, 0.56) \times (0.063, 0.19), \Lambda_3 = \{(x, y) : 0.063 \leq (x-0.5)^2+(y-0.3)^2 \leq 0.14, y \geq 0.3\}$. For the static field case the differential equation is in the form (1), and for this experiment we imposed homogeneous Neumann boundary conditions on the left and right boundaries and Dirichlet boundary conditions of 1 and 0 to the top and bottom boundaries respectively of the unit square.

Figure 1 Geometry of an electromagnetic problem (a), fine mesh (b) and coarse mesh node points (c)

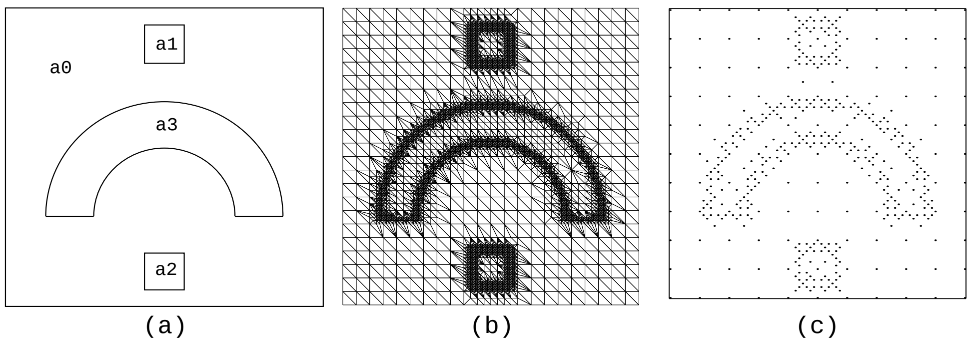

The domain is discretised with a uniform mesh of triangles except for strong refinement near the boundaries of the regions with differing material properties. This results in a mesh of 30856 triangles with 15484 nodes (Figure 1(b)). These do not resolve completely the semicircular geometry of Λ_3.

In this experiment we use the two-level additive Schwarz method, see for example [CSZ96], for which we, in contrast to (3), require the solution of a global coarse problem together with local problems on subdomains of the fine mesh. In principle the coarse mesh is not required to have any direct relation to the fine mesh. However, it may be expected that a coarse mesh which pays no attention to the underlying PDE (e.g. fails to have some refinement where the fine mesh is refined in this example) may not work well. To determine our coarse mesh we first impose a uniform coarse mesh and then perform hierarchical local refinement with "slave nodes" as, e.g., in §7 of [CM94], to increase the density in regions where the fine mesh is dense. The result has 465 nodes and is pictured in Figure 1(c). This coarse mesh is represented by a locally uniform data structure which is completely uniform in large sections of the domain, this allows for a very simple and efficient implementation. More details on the creation and performance of such coarse meshes will be available in a future publication. The partitioning of the fine mesh into the local subdomains is performed using the graph partitioning package METIS[KK95]. Use of this package allows us to produce load balanced connected subdomains, with a single node overlap, based only on the connectivity of the mesh. Hence the geometry of the problem has no direct bearing on the subdomains and neither the fine mesh, coarse mesh nor subdomains resolve the discontinuity in $\mathbf{a}$.

To demonstrate the effect of the discontinuous coefficients on the additive Schwarz preconditioner we use five sets of coefficient values, as specified in Table 1. Each of these problems is also tested with a range of subdomain numbers, from 15 to 120. Three different preconditioners are tested in each case: One with no coarse solve (AS),

Problem	a0	a1	a2	a3
1	1.0	1.0	1.0	1.0
2	1.0	1.0(-1)	1.0(1)	1.0(2)
3	1.0	1.0(-2)	1.0(2)	1.0(4)
4	1.0	1.0(-3)	1.0(3)	1.0(6)
5	1.0	1.0(-4)	1.0(4)	1.0(8)

Table 1 Coefficient values for the test problems

15 Subdomains				30 Subdomains			
Problem	AS	ASC*	ASC	Problem	AS	ASC*	ASC
1	70	25	27	1	87	26	28
2	92	27	26	2	116	31	28
3	109	30	29	3	131	35	33
4	117	31	31	4	149	36	35
5	135	31	31	5	165	36	34
60 Subdomains				**120 Subdomains**			
Problem	AS	ASC*	ASC	Problem	AS	ASC*	ASC
1	123	33	33	1	158	39	34
2	150	36	32	2	186	40	36
3	172	38	37	3	213	44	39
4	196	39	37	4	241	45	43
5	220	39	37	5	272	45	44

Table 2 Results for the 15,30,60 and 120 subdomain case

a second with a coarse solve but based on a purely uniform coarse mesh, with 121 nodes, (ASC*), and the third with coarse mesh pictured in Figure 1(c) (ASC). An initial guess of **0** was used for the CG algorithm.

Table 2 shows the number of iterations of the preconditioned CG method required to satisfy the convergence criterion $\| \mathbf{x}^j - \mathbf{x} \|_2 / \| \mathbf{x} \|_2 < 10^{-5}$, where $\mathbf{x}^j$ is the jth iterate and $\mathbf{x}$ is the true solution computed using an exact factorisation. In all cases convergence to the same tolerance in the energy norm took place within one or two iterates of the results given here. The results for the preconditioner with no coarse solve (AS) show the expected logarithmic growth in the number of CG iterations required, as the jumps in the coefficient worsen. Additionally an increase in the number of subdomains also results in the expected increase in iterations. For the remaining two preconditioners (ASC* and ASC) the growth with respect to the number of subdomains is almost identical, indicating that, for this problem, a simple coarse mesh would be sufficient. The results for ASC* in this case raise the interesting question of whether a uniform coarse mesh would suffice for more general preconditioning tasks. Preliminary experiments indicate that this is not always the case. This will be the subject of future work.

Acknowledgement

This work was supported by UK EPSRC grant GR/J88616.

REFERENCES

[CGZ96] Chan T. F., Go S., and Zou J. (1996) Multilevel domain decomposition and multigrid methods for unstructured meshes: Algorithms and theory. In Glowinski R., Périaux J., Shi Z.-C., and Widlund O. B. (eds) *Proceedings of the Eighth International Conference on Domain Decomposition*. Wiley and Sons, Chichester.

[CM94] Chan T. F. and Mathew T. P. (1994) Domain decomposition algorithms. In Iserles A. (ed) *Acta Numerica*, pages 61–143. Cambridge University Press.

[CNT96] Cai X., Nielsen B. F., and Tveito A. (1996) An analysis of a preconditioner for the discretised pressure equation arising in reservoir simulation. Preprint.

[CSZ96] Chan T. F., Smith B., and Zou J. (1996) Overlapping schwarz methods on unstructured meshes using non-matching coarse grids. *Numer. Math.* 73: 149–167.

[DSW96] Dryja M., Sarkis M. V., and Widlund O. B. (1996) Multilevel Schwarz methods for elliptic problems with discontinuous coefficients in 3 dimensions. *Numer. Math.* 72(3): 313–348.

[EST94] Emson C. R. I., Simkin J., and Trowbridge C. W. (1994) A status report on electromagnetic field computation. *IEEE Trans. Magnetics* 30(4): 1533–1540.

[GH96] Graham I. G. and Hagger M. J. (1996) Unstructured additive Schwarz - CG method for elliptic problems with highly discontinuous coefficients. Submitted to SIAM J. Sci. Comp., June 1996.

[KK95] Karypis G. and Kumar V. (1995) *METIS: Unstructured graph partitioning and sparse matrix ordering system*. Dept. Computer Science, Univ. of Minnesota, Minneapolis.