European Conference on Computational Fluid Dynamics
ECCOMAS CFD 2006
P. Wesseling, E. Oñate and J. Périaux (Eds)

ROBUST AGGREGATION-BASED COARSENING FOR ADDITIVE SCHWARZ IN THE CASE OF HIGHLY VARIABLE COEFFICIENTS

R. Scheichl* **and E. Vainikko†**

*Bath Institute for Complex Systems and Department of Mathematical Sciences,
University of Bath, Claverton Down, Bath BA2 7AY, UK.
Email: masrs@maths.bath.ac.uk

†Institute of Computer Science, University of Tartu,
Liivi 2, Tartu 50409, Estonia.
Email: eero.vainikko@ut.ee

Key words: Second-order Elliptic Problems, Heterogeneous Media, Smoothed Aggregation, Two-level Additive Schwarz, Algebraic Multigrid

Abstract. *We study two–level overlapping domain decomposition preconditioners with coarse spaces obtained by smoothed aggregation in iterative solvers for finite element discretisations of second-order elliptic problems. We are particularly interested in the situation where the diffusion coefficient (or the permeability) α is highly variable throughout the domain. Our motivating example is Monte-Carlo simulation for flow in rock with permeability modelled by log-normal random fields. By using the concept of strongly-connected graph r-neighbourhoods (suitably adapted from the algebraic multigrid context) we design a two–level additive Schwarz preconditioner that is robust to strong variations in α as well as to mesh refinement. We give upper bounds on the condition number of the preconditioned system which do not depend on the size of the subdomains (not available previously in the literature) and make explicit the interplay between the coefficient function and the coarse space basis functions in this bound. In particular, we are able to show that the condition number can be bounded independent of the ratio of the two values of α in a binary medium even when the discontinuities in the coefficient function are not resolved by the coarse mesh. Our numerical results show that the bounds with respect to the mesh parameters are sharp and that the method is indeed robust to strong variations in α. We compare the method to other preconditioners (aggregation-type AMG and classical additive Schwarz) as well as to a sparse direct solver, and show its superiority over those methods for highly variable coefficient functions α.*

1 INTRODUCTION

This paper extends work by Brezina & Vanek[3], Jenkins et al.[14], Lasser & Toselli[16], Sala[17], and Sala et al.[18] on two-level additive Schwarz preconditioners based on smoothed

aggregation techniques first introduced by Vanek, Mandel & Brezina[21] in the context of algebraic multigrid methods. We consider the iterative solution of linear systems of equations resulting from discretisations of boundary-value problems for the model elliptic problem

$$-\nabla \cdot (\alpha \nabla u) \;=\; f \;, \tag{1}$$

in a bounded 2D or 3D domain Ω, subject to homogeneous Dirichlet boundary conditions. We are particularly interested in the case where α is highly variable throughout the domain, as for example in the simulation of flow through heterogeneous porous media.

Let us consider the discretisation of (1) using continuous piecewise linear finite elements $\mathcal{V}^h$ on a triangulation $\mathcal{T}^h$ of Ω of mesh width h. Then the condition number of the resulting stiffness matrix A will grow like $O(h^{-2})$ as the mesh is refined. In addition the condition number will also depend on $\sup_{x,y\in\Omega}(\alpha(x)/\alpha(y))$. To improve the conditioning, suppose that Ω is covered by a set $\{\Omega_i : i = 1,\ldots,s\}$ of overlapping subdomains such that $\bigcup_{i=1}^s \Omega_i = \Omega$ and such that $\operatorname{diam}\Omega_i \le H^{sub}$. Secondly, suppose that we also have a coarse space $\mathcal{V}_0 \subset \mathcal{V}^h$. In our case the coarse space $\mathcal{V}_0 := \operatorname{span}\{\phi_j : j = 1,\ldots,N\}$ will be obtained by smoothed aggregation, i.e. the coarse space basis functions ϕ_j are obtained by grouping together fine grid nodes into aggregates of diameter $\le H$, by summing the associated fine grid basis functions and by "smoothing" the result[21]. Let M_{AS}^{-1} denote the classical two-level additive Schwarz preconditioner, obtained by solving discretisations of (1) on each of the overlapping subdomains as well as on the coarse space $\mathcal{V}_0$ (see e.g. Toselli & Widlund[20]).

The main theoretical result of this paper is to improve the bounds for $\kappa(M_{AS}^{-1}A)$ in the literature[16,17,18]. Note that as we will see in the numerical experiments, for efficiency reasons it is of interest to choose $H^{sub} \gg H$. However, the estimates in the literature[16,17,18] all involve H^{sub} as well as H and assume $\alpha \equiv 1$. We extend their results to the case $\alpha \not\equiv 1$ and prove a sharper bound that makes explicit the dependency on α and on the mesh parameters. In particular, we are able to show that $\kappa(M_{AS}^{-1}A)$ is independent of H^{sub}, the size of the subdomains, and only depends linearly on the local ratio of the size of the coarse space aggregates and the size of their overlap. Our numerical experiments show that this bound is sharp.

The dependency of the condition number $\kappa(M_{AS}^{-1}A)$ on α is reduced to the quantity $\gamma(\alpha) := \max_j \{\delta_j^2 \, \||\alpha|\nabla\phi_j|^2\|_{L_\infty(\Omega)}\}$ where δ_j is the size of the overlap of the support of ϕ_j and its neighbours, i.e provided $\nabla\phi_j(x)$ is small wherever $\alpha(x)$ is large, then $\kappa(M_{AS}^{-1}A)$ can be bounded independently of α. We will see that for certain choices of the coefficient function α, smoothed aggregation techniques produce coarse space basis functions such that $\gamma(\alpha)$ remains bounded even when $\sup_{x,y\in\Omega}\frac{\alpha(x)}{\alpha(y)}$ goes to infinity. For highly variable α the strongest results in the domain decomposition literature are for the "structured" case with standard linear coarse space, in which the coarse mesh is constructed to resolve discontinuities in α. In such cases it is possible to bound the condition number independent of α but at the expense of a stronger dependency on the mesh parameters. An excellent survey of such results can be found in Chan & Mathew[6]. Another class of results[10,22,11,4]

applies when the number of discontinuities in α which are not resolved is small. Then it can be shown that domain decomposition preconditioners produce a highly clustered spectrum with relatively few near-zero eigenvalues – which is an advantageous situation for Krylov subspace methods like CG. We are not aware of any theoretical results in the algebraic multigrid literature which make explicit the dependency on α for highly variable coefficient functions. Some other related results about robust coarsening can be found in the domain decomposition and multigrid literature[23,5,9,13].

To test the resilience of our method to strong variations in the coefficient function α we study problem (1) on the unit square with coefficient function α chosen as a realisation of a log-normal random field or of a "clipped" log-normal field with variance σ^2 and correlation length scale λ. The method proves indeed to be extremely robust with respect to variations in α and outperforms AMG and standard two-level additive Schwarz with linear coarse space.

For more numerical results and detailed proofs see Scheichl & Vainikko[19]. Note also that the there is a strong link between the results in this paper and the results in Graham, Lechner & Scheichl[12].

Throughout the paper, the notation $C \lesssim D$ (for two quantities C, D) means that C/D is bounded above independently of the *mesh parameters* and of the *coefficient function* α.

2 GENERAL THEORY

In this section we provide a general framework for the analysis of domain decomposition preconditioners for (1) in which the dependence of the condition number of the preconditioned stiffness matrix on α, as well as on the mesh parameters, is made precise. For details and proofs see Scheichl & Vainikko[19].

Let $\{\mathcal{T}^h\}$ be a shape-regular family of triangulations of Ω of mesh width h, and let $\mathcal{S}^h(\Omega)$ denote the subspace of $H^1(\Omega) \cap C(\bar{\Omega})$, consisting of continuous piecewise linear functions with respect to $\mathcal{T}^h$. We consider the bilinear form

$$a(u,v) \;:=\; \int_\Omega \alpha \nabla u \cdot \nabla v, \qquad u, v \in H^1_0(\Omega)\ , \tag{2}$$

and its Galerkin approximation in the n–dimensional space $\mathcal{V}^h := \mathcal{S}^h(\Omega) \cap H^1_0(\Omega)$. Let A be the corresponding $n \times n$ stiffness matrix using the standard nodal basis for $\mathcal{V}^h$.

We are interested in iterative methods for solving (1) and hence in preconditioners for A which remove the ill-conditioning due to both the non-smoothness of α and the smallness of the mesh width h. We will be concerned with preconditioners based on domain decomposition methods. Let $\{\Omega_i : i = 1, \ldots, s\}$ be an overlapping open covering of Ω. Note that we will neither make any assumptions on the shape of the subdomains Ω_i nor on the way they overlap, except that they should be resolved by the fine mesh. However the conditions on our coarse space basis functions below will implicitly induce some assumptions on the subdomains. Let $\mathcal{V}_i := \mathcal{S}^h_0(\Omega_i)$ and let R_i denote the restriction

matrix which takes degrees of freedom in Ω to degrees of freedom in Ω_i and set $A_i := R_i A R_i^T$. Hence, A_i is just the minor of A corresponding to the degrees of freedom in Ω_i.

To obtain scalability with respect to the number of subdomains, one normally introduces an additional coarse space. We will define a coarse space in a quite general way by defining a set of basis functions which satisfy certain assumptions. In Section 3 we will then describe an aggregation technique to construct a set of functions which satisfy these assumptions. Let $\{\phi_j : j = 1, \ldots, N_H\} \subset S^h(\Omega)$ be a linearly independent set of finite element functions and let

$$\omega_j := \text{interior}\left(\text{supp}\{\phi_j\}\right).$$

(Note that (C1) below guarantees that $\{\omega_j\}$ is a covering of Ω.) Define $H_j := \text{diam}\{\omega_j\}$ and set $H := \max_{j=1}^{N_H} H_j$. We need to assume (for theoretical purposes only) that the covering $\{\omega_j\}$ is shape regular and that the overlap between any support ω_j and its neighbours is uniform of size δ_j (for details see Scheichl et al.[19]). Also, set $\delta := \min_{j=1}^{N_H} \delta_j$. We make the following assumptions on the the functions $\{\phi_j : j = 1, \ldots, N_H\}$:

(C1) $$\sum_{j=1}^{N_H} \phi_j(x) \;=\; 1, \;\text{ for all }\; x \in \bar{\Omega}.$$

(C2) For all $j = 1, \ldots, N_H$ there is a unique $i_j \in \{1, \ldots, s\}$ such that $\omega_j \;\subset\; \Omega_{i_j}$.

(C3) $$\|\phi_j\|_{L_\infty(\Omega)} \;\lesssim\; 1 \,.$$

To simplify our notation we assume that the functions ϕ_j are numbered in such a way that $\phi_j \in \mathcal{V}^h$ for all $j \le N$ and $\phi_j \notin \mathcal{V}^h$ for all $j > N$, with $N < N_H$, i.e. for all $j \le N$ we have $\phi_j|_{\partial\Omega} = 0$. We can then define the coarse space as follows:

$$\mathcal{V}_0 \;=\; \text{span}\{\phi_j : j = 1, \ldots, N\}$$

and we have $\mathcal{V}_0 \subset \mathcal{V}^h$. Now, if we introduce the restriction matrix

$$(R_0)_{j,p} \;=\; \phi_j(x_p)\,, \quad p = 1, \ldots, n, \quad j = 1, \ldots, N,$$

where x_p, $p = 1, \ldots, n$, are the interior nodes of the fine mesh $\mathcal{T}^h$, then the matrix $A_0 := R_0 A R_0^T$ is the stiffness matrix for the bilinear form $a(\cdot,\cdot)$ discretised in $\mathcal{V}_0$ using the basis $\{\phi_j : j = 1, \ldots, N\}$. The corresponding two-level additive Schwarz preconditioner, based on combining coarse and subdomain solves is

$$M_{AS}^{-1} \;=\; \sum_{i=0}^{s} R_i A_i^{-1} R_i^T \,. \tag{3}$$

Note that although we have not directly made any assumptions on the overlap between the subdomains Ω_i, Assumption (C2) implies that the minimum overlap between two of

the subdomains $\{\Omega_i\}$ is in fact bounded from below by δ, the minimum overlap of the supports $\{\omega_j\}$.

It has already been noted in the literature[16,17,18] that in order to bound $\kappa(M_{AS}^{-1}A)$, the condition number of the preconditioned stiffness matrix in the case of the two-level additive Schwarz preconditioner, we need a further assumption on the gradient of the coarse space basis functions. The following two assumptions have been used previously:

(C4a) $$|\phi_j|^2_{H^1(\Omega)} \lesssim \frac{H_j^{d-1}}{\delta_j} \quad j = 1, \ldots, N_H \ ,$$

(C4b) $$\|\nabla\phi_j\|^2_{L_\infty(\Omega)} \lesssim \delta_j^{-2} \quad j = 1, \ldots, N_H \ .$$

Note that (C4b) implies (C4a) and is therefore a stronger assumption.

In the case of $s = N_H$ (i.e. one coarse space basis function per subdomain) and $\alpha \equiv 1$ the following bound has been proved in Lasser & Toselli[16]. It can also be found in Toselli & Widlund[20].

Theorem 2.1 *Assume that (C1)-(C3) hold true and that in addition $s = N_H$ and $\alpha \equiv 1$. Then*

$$\kappa(M_{AS}^{-1}A) \lesssim \left(1 + \max_{j=1}^{N_H} \frac{H_j}{\delta_j}\right)^\beta$$

where $\beta = 2$ if assumption (C4a) holds and $\beta = 1$ if assumption (C4b) holds.

As we will see, in practice it is often more efficient to choose $s \ll N_H$, or in other words subdomains of much larger diameter than H. Let us denote $H^{sub} := \max_{i=1}^s \operatorname{diam}\Omega_i$ and let δ^{sub} be the minimum overlap between any two subdomains. It has been shown in Sala[17] that for $\alpha \equiv 1$ and under the weaker assumption (C4a)

$$\kappa(M_{AS}^{-1}A) \lesssim \left(1 + \frac{H^{sub}}{\delta^{sub}}\right)\left(1 + \frac{H}{\delta}\right) \tag{4}$$

This result has been improved in Sala et al.[18] using the stronger assumption (C4b) to give

$$\kappa(M_{AS}^{-1}A) \lesssim \left(1 + \frac{H^{sub}}{\delta^{sub}} + \frac{H}{\delta}\right). \tag{5}$$

However, both these results are not sharp for $H^{sub} \gg H$ as our numerical results in Section 4 will show. Indeed it is possible to prove the following result (even for $s \ll N_H$) using a simple colouring argument.

Theorem 2.2 *Assume that (C1)-(C3) hold true. Then*

$$\kappa\left(M_{AS}^{-1}A\right) \lesssim \sup_{x,y\in\Omega} \frac{\alpha(x)}{\alpha(y)} \left(1 + \max_{j=1}^{N_H} \frac{H_j}{\delta_j}\right)^\beta$$

where $\beta = 2$ if assumption (C4a) holds and $\beta = 1$ if assumption (C4b) holds.

Proof. See Scheichl & Vainikko[19]. □

We will see below that the bound in Theorem 2.2 can still be improved with respect to the dependency on the coefficient function $\alpha(x)$. However, Theorem 2.2 constitutes already a new result in its own right since it provides a sharper bound with respect to the mesh parameters than previously available in the literature. The two main points which should be highlighted are that our new result shows that (i) the condition number of the preconditioned stiffness matrix using a two-level additive Schwarz preconditioner with an aggregation-type coarse space is independent of the size of the subdomains and that (ii) it only depends on local ratios of the size of the supports of the coarse space basis functions and their overlap. The numerical results in Section 5 will show that this accurately reflects the dependency of the condition number on the mesh parameters.

We have seen above that the size of the bound on the condition number $\kappa\left(M_{AS,2}^{-1}A\right)$ depends on the assumptions which are made on the gradient of the coarse space basis functions. We now make the dependency even more explicit by introducing the following quantity which measures the robustness of the coarse space $\mathcal{V}_0$ with respect to variations in α:

Definition 2.3 (Coarse space robustness indicator).

$$\gamma(\alpha) \;:=\; \max_{j=1}^{N_H} \left\{ \delta_j^2 \, \| \alpha |\nabla \phi_j|^2 \|_{L_\infty(\Omega)} \right\} \;.$$

For the remainder of this section we assume that $\alpha \geq 1$. This is no loss of generality, since otherwise the problem (1) can be rescaled by dividing through by $\inf_{x\in\Omega} \alpha(x)$ without changing the condition number of the resulting discrete problem.

The following result which constitutes the main theoretical contribution of this paper relies on an improved bound for the stability constant C_0 for the decomposition of an arbitrary $u_h \in \mathcal{V}^h$ into elements $u_i \in \mathcal{V}_i$. The proof of this result is quite technical and is omitted. It can again be found in Scheichl and Vainikko[19].

Theorem 2.4 *Assume that (C1)-(C3) hold true. Then*

$$\kappa\left(M_{AS}^{-1}A\right) \;\lesssim\; \gamma(\alpha) \left(1 + \max_{j=1}^{N_H} \frac{H_j}{\delta_j}\right) \;.$$

In principle, for an arbitrary coefficient function $\alpha(x)$ and for an arbitrary set of coarse space basis functions that satisfy (C4b), the quantity $\gamma(\alpha)$ can become as bad as $\sup_{x,y\in\Omega} \frac{\alpha(x)}{\alpha(y)}$ which is the quantity that appears in Theorem 2.2. However, the huge improvement in Theorem 2.4 lies in the fact that $\gamma(\alpha)$ accurately reflects the interplay between coefficient function and coarse space basis functions. In fact, we will see in Section 3 that for a range of coefficient functions $\alpha(x)$, smoothed aggregation techniques produce coarse space basis functions such that $\gamma(\alpha)$ remains bounded even when $\sup_{x,y\in\Omega} \frac{\alpha(x)}{\alpha(y)}$ goes to infinity.

3 SMOOTHED AGGREGATION COARSE SPACES

Smoothed aggregation techniques have been introduced first in the context of algebraic multigrid methods in Vanek, Mandel & Brezina[21] and further investigated by Brezina & Vanek[3], Jenkins et al.[14], Lasser & Toselli[16], Sala[17], and Sala et al.[18] in the context of Schwarz methods. However, surprisingly, all of the latter papers only concentrate on the case $\alpha \equiv 1$ (or $\alpha \sim 1$) and do not use the concept of strongly-connected neighbourhoods of nodes which plays such a key rôle in the context of the coarse grid construction in algebraic multigrid[21].

To describe the smoothed aggregation algorithm that we use to construct a set of coarse space basis functions $\{\phi_j : j = 1, \ldots, N_H\}$ which satisfy the assumptions in Section 2, we first need to define *strongly-connected graph r–neighbourhoods*. Let $\mathcal{N} := \{x_p : p = 1, \ldots, n_h\}$ be the set of all nodes of $\mathcal{T}^h$ including the boundary nodes (so that $n_h > n$), and let $\mathcal{A}$ be the $n_h \times n_h$ stiffness matrix corresponding to a discretisation of $a(\cdot, \cdot)$ in $S^h(\Omega)$, i.e. including the degrees of freedom on the boundary.

Definition 3.1 (a) Let x_p and x_q be two neighbouring nodes of $\mathcal{T}^h$, $p \neq q$, i.e. there exists a $\tau \in \mathcal{T}^h$ such that $x_p, x_q \in \tau$. Then node x_q is *strongly connected* to x_p iff

$$|\tilde{A}_{pq}| \;\geq\; \varepsilon \max_{k \neq p} |\tilde{A}_{pk}| \tag{6}$$

where $\tilde{A} := (\operatorname{diag} \mathcal{A})^{-1/2} \mathcal{A} (\operatorname{diag} \mathcal{A})^{-1/2}$ and $\varepsilon \in [0, 1]$ is a pre–determined threshold. Let $S_\varepsilon(x_p)$ denote the set consisting of all nodes x_q that are strongly connected to x_p with threshold ε, as well as the node x_p itself.

(b) Let $\mathcal{G} := (\mathcal{N}, \mathcal{E})$ be the graph induced by the mesh $\mathcal{T}^h$, where $\mathcal{E}$ denotes the set of all edges of $\mathcal{T}^h$. Now, let x_p and x_q be two (arbitrary) nodes of $\mathcal{T}^h$. Then x_p and x_q are *strongly connected* iff there exists a path γ_{pq} in $\mathcal{G}$ with nodes $x_p = x_{p_0}, x_{p_1}, \ldots, x_{p_\ell} = x_q$ such that x_{p_i} is strongly connected to $x_{p_{i-1}}$ for all $i = 1, \ldots, \ell$ (in the sense of (a)). Let ℓ_{pq} be the length of the shortest such path γ_{pq}.

(c) The *strongly–connected graph r–neighbourhood* of a node x_p is the set $S_{r,\varepsilon}(x_p)$ consisting of all nodes x_q that are strongly connected to x_p with $\ell_{pq} \leq r$, as well as the node x_p itself.

To our knowledge the criterion (6) does not appear anywhere in the literature. It stems from the algebraic multigrid code of Bastian[2] and is a modified version of the criterion in Vanek et al.[21]. Strongly–connected graph r-neighbourhoods are to our knowledge also a novel concept in the context of domain decomposition methods.

The construction of the coarse space basis functions is now almost identical to the algorithm described in Brezina & Vanek[3]. However, the heuristics which we use (a) to choose good seed nodes for each aggregate, (b) to ensure that the aggregates are shape-regular where possible, (c) to minimise the number of nonzeros in the coarse matrix, and

(d) to deal with isolated nodes are inspired by Bastian[2]. For details on those heuristics and on the more practical aspects of our implementation see Scheichl et al.[19]. The algorithm consists of two main steps: aggregation and smoothing. For both of those steps we make use of the so–called *filtered matrix* A^ε with entries

$$A^\varepsilon_{pq} := \begin{cases} \mathcal{A}_{pp} + \sum_{x_q \notin S_\varepsilon(x_p)} \mathcal{A}_{pq} & \text{if} \quad p = q, \\ \mathcal{A}_{pq} & \text{if} \quad x_q \in S_\varepsilon(x_p)\backslash\{x_p\}, \\ 0 & \text{otherwise.} \end{cases}$$

We start by creating a set of aggregates $\{W_j : j = 1, \ldots, N_H\}$ such that

$$\mathcal{N} = \bigcup_{j=1,\ldots,N_H} W_j \quad \text{and} \quad W_j \cap W_{j'} = \emptyset \quad \forall j \neq j'$$

(i.e. a non-overlapping partition of $\mathcal{N}$). The sets W_j can be obtained by choosing an aggregation "radius" $r \in \mathbb{N}$, a threshold $\varepsilon \in [0, 1]$ and a set of seed nodes $x^H_j \in \mathcal{N}$, $j = 1, \ldots, N_H$, and by calculating the strongly–connected graph r-neighbourhoods $S_{r,\varepsilon}(x^H_j)$. To calculate $S_{r,\varepsilon}(x^H_j)$ we make use of the filtered matrix A^ε. Note that given a node $x_p \in \mathcal{N}$ we can calculate $S_{r,\varepsilon}(x_p)$ by calculating the vector $\boldsymbol{w}^p = (A^\varepsilon)^r \boldsymbol{e}^p$ where $\boldsymbol{e}^p$ is the vector with entries $e^p_q = \delta_{pq}$. $S_{r,\varepsilon}(x_p)$ is then given by all the nodes x_q such that $w^p_q \neq 0$. It is sufficient to carry out the sparse matrix-vector products needed in the calculation of $\boldsymbol{w}^p$ symbolically.

In practice we choose the seed nodes by sweeping through Ω with an advancing front. This might not lead to a complete partitioning of $\mathcal{N}$ and some heuristic procedures need to be put in place that deal with nodes that are left over and/or isolated (not strongly–connected to any other node). The algorithm aims to produce shape–regular aggregates W_j and is guaranteed to achieve this in the case where $\{\mathcal{T}^h\}$ is a quasi-uniform family of triangulations and all connections are strong (see Figure 1 (left) for the case $\alpha \equiv 1$ and $r = 2$). However, for an arbitrary strongly–varying coefficient function α the aggregates W_j may (and indeed should in many cases) not be shape-regular.

For each $j = 1, \ldots, N_H$ we now define a vector $\boldsymbol{\Psi}^j \in \mathbb{R}^{n_h}$ as follows:

$$\Psi^j_p := \begin{cases} 1 & \text{if } x_p \in W_j \\ 0 & \text{otherwise.} \end{cases}$$

Let $\psi_j \in S^h(\Omega)$ be the finite element function corresponding to the coefficient vector $\boldsymbol{\Psi}^j$. We further smooth these vectors $\boldsymbol{\Psi}^j$ by using a damped Jacobi smoother

$$S := \left(I - \omega \left(\operatorname{diag} A^F\right)^{-1} A^F\right).$$

with damping parameter ω. Let

$$\boldsymbol{\Phi}^j := S^\mu \boldsymbol{\Psi}^j$$

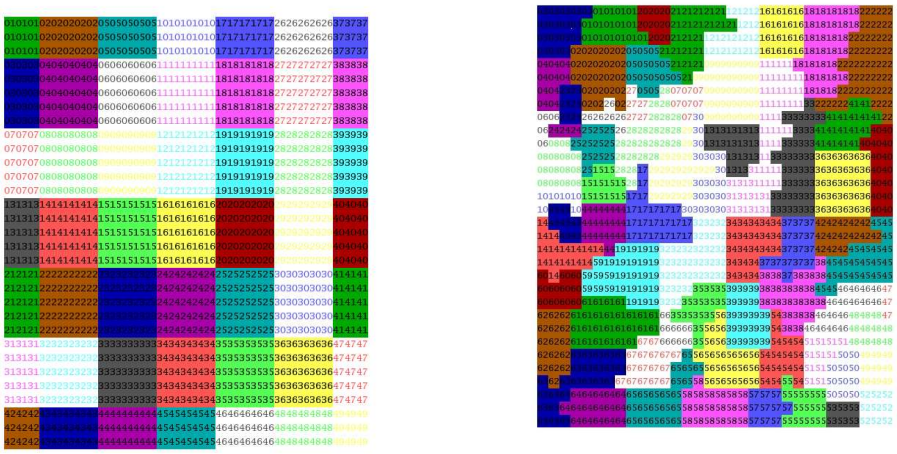

Figure 1: Aggregates W_j for $\alpha \equiv 1$ (left) and α strongly–varying (right) with $r = 2$, $\varepsilon = \frac{2}{3}$.

where μ is the number of smoothing steps. The jth coarse space basis function $\phi_j \in S^h(\Omega)$ is the finite element function corresponding to the coefficient vector $\mathbf{\Phi}^j$.

To construct the subdomains Ω_i we apply the aggregation procedure (described above) to A_0. Therefore each subdomain Ω_i will consist of the union of the supports of a set of coarse space basis functions ϕ_j and (C2) is satisfied.

Let us now consider whether the functions ϕ_j, $j = 1, \ldots, N_H$, satisfy the other assumptions made in Section 2. Note that for quasi-uniform $\{\mathcal{T}^h\}$ and $\alpha \sim 1$ all connections in A are strong (provided ε is not too close to 1). This case has already been covered in the literature. See Brezina & Vanek[3] and Lasser & Toselli[16] for details. It is important to note however, that it has so far not been possible to prove Assumption (C4b) in the case of smoothed basis functions (including the case of damped Jacobi smoothing used here). In the unsmoothed case, i.e. for $\mu = 0$, (C4b) follows directly from the construction of the functions ψ_j.

In the case of strongly varying α nothing has been proved so far. We will restrict to the unsmoothed case here, i.e. $\mu = 0$ and so $\phi_j = \psi_j$. (For the case of smoothed aggregation see Scheichl et al.[19].) In the case of $\mu = 0$ all the assumptions made in Section 2 are satisfied by construction, except the shape regularity of the supports ω_j. This is not guaranteed and depends on the coefficient function α. The size of the overlap $\delta_j = O(h_j^{min})$ where h_j^{min} is the diameter of the smallest element $\tau \subset \omega_j$ which touches the boundary of ω_j.

For certain special choices of α it can be shown that the covering $\{\omega_j\}$ is shape regular even when α varies very strongly, and moreover that the coarse space robustness indicator $\gamma(\alpha)$ in Definition 2.3 is bounded independently of α and of the mesh parameters. Take

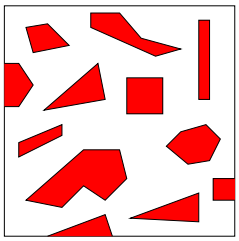

Figure 2: Typical situation for Example 3.2, i.e. "islands" B_k (in red) where α is large.

for instance the following example of a binary medium α:

Example 3.2 Let B_k, $k = 1, \ldots, m$, be closed, simply connected, disjoint, polygonal subsets of Ω ("islands"), i.e. $B_k \cap B_{k'} = \emptyset$ for all $k' \neq k$ (see Figure 2 for an example). Let us assume for simplicity that the distance between two islands B_k and $B_{k'}$ is comparable in size to their diameter. Now let

$$\alpha(x) \;=\; \begin{cases} \hat{\alpha} & \text{if} \quad x \in B_k \quad \text{for some} \quad k = 1, \ldots, m \\ 1 & \text{otherwise} \end{cases}$$

with $\hat{\alpha} \gg 1$. Note first of all that for $\hat{\alpha}$ large enough (and for h small enough) we have for any $x_p \in \mathcal{N}$ either (i) $S_\varepsilon(x_p) \subset B_k$ for some $k = 1, \ldots, m$, or (ii) $S_\varepsilon(x_p) \cap B_k = \emptyset$ for all $k = 1, \ldots, m$; i.e. if two nodes are strongly connected they either both lie in one of the sets B_k or they do not lie in any of the sets at all. Hence, the aggregates W_j, $j = 1, \ldots, N_H$, constructed above satisfy either $W_j \subset B_k$ for some $k = 1, \ldots, m$ or $W_j \cap B_k = \emptyset$ for all $k = 1, \ldots, m$.

A good choice of seed nodes will then ensure that the supports ω_j of the coarse space basis functions ϕ_j are shape regular as $h \to 0$. The shape regularity constant will depend on the original configuration of the sets B_k, e.g. it might be large if one of the sets B_k is very long and thin, or if the gap between two islands is small, but it will not depend on h or any other mesh parameter as $h \to 0$. Moreover, if we choose the aggregation radius r large enough, then each island B_k will contain exactly one aggregate W_j. Since the sets B_k were assumed to be closed, we then have $\alpha|_\tau = 1$ for all elements τ in the overlap of any two supports ω_j and $\omega_{j'}$ with $j' \neq j$. Since $\nabla\phi_j(x) \leq \delta_j^{-1}$ for all $x \in \Omega$, this implies that the coarse space robustness indicator

$$\gamma(\alpha) \leq 1.$$

Therefore in the case of a quasi-uniform family of meshes and for suitably chosen r, the bound in Theorem 2.4 reduces to

$$\kappa\left(M_{AS}^{-1}A\right) \;\lesssim\; h^{-1} \max_{k=1}^{m} \{\operatorname{diam} B_k\}$$

independent of the size of $\hat{\alpha}$. Hence, if the maximum diameter of the islands is $O(h)$, the bound is completely independent of α.

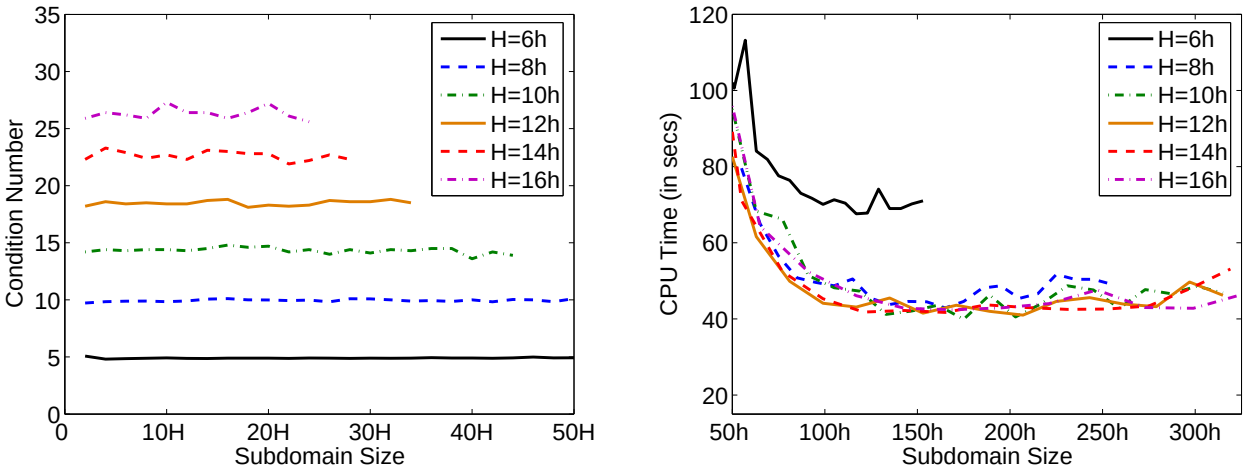

Figure 3: Condition numbers and CPU times for Laplacian with $n = 1024^2$, $\delta^{sub} = 3h$, $\mu = 1$.

4 NUMERICAL RESULTS

In all the numerical experiments below $\Omega = [0,1]^2$ and $\{\mathcal{T}^h\}$ is a family of uniform triangulations of Ω. We solve the resulting linear equation systems with preconditioned CG and tolerance 10^{-6}. We use the sparse direct solver `UMFPACKv4.4`[8] to solve the subdomain and coarse grid problems. All CPU times were obtained on a 3GHz Intel P4 processor with 1GByte RAM.

Let us first consider the case $\alpha \equiv 1$, i.e. the Laplacian. As we have seen in the previous section (cf. Figure 1), in this case the aggregates $\{W_j\}$ and thus the supports $\{\omega_j\}$ are uniformly of size $H = 2(r+\mu+1)h$ and overlap $\delta = (2\mu+1)h$ where r is the aggregation radius in the fine grid aggregation (i.e. aggregating the degrees of freedom in A) and μ is the number of smoothing steps. Similarly, the subdomains Ω_i are uniformly of size $H^{sub} \approx (2r_0+1)H$, where r_0 is the coarse grid aggregation (i.e. aggregating the degrees of freedom in A_0).

In Figure 3 we plot the condition numbers $\kappa\left(M_{AS}^{-1}A\right)$ of the preconditioned systems and the CPU times for various choices of r and r_0. The various parameters in our method are: the problem size $n = 1024^2$ and thus $h = 1/1025$; the subdomain overlap $\delta^{sub} = 3h$; the threshold for strong connections $\varepsilon = \frac{2}{3}$; the damping parameter in the Jacobi smoother $\omega = \frac{2}{3}$ and the number of smoothing steps $\mu = 1$ (hence $\delta = 3h$).

We note that the agreement with the theory is extremely good: the condition number $\kappa\left(M_{AS}^{-1}A\right) \leq 5\frac{H}{\delta}$, for all values of r and r_0, and it is independent of the subdomain size H^{sub}. The CPU times confirm the statement made earlier that it is more efficient computationally to choose $H^{sub} \gg H$. We see that the best efficiency of the method is attained for $H^{sub} \approx 10H - 30H$.

The method is also independent of h as we see in Figure 4. Here, using the same

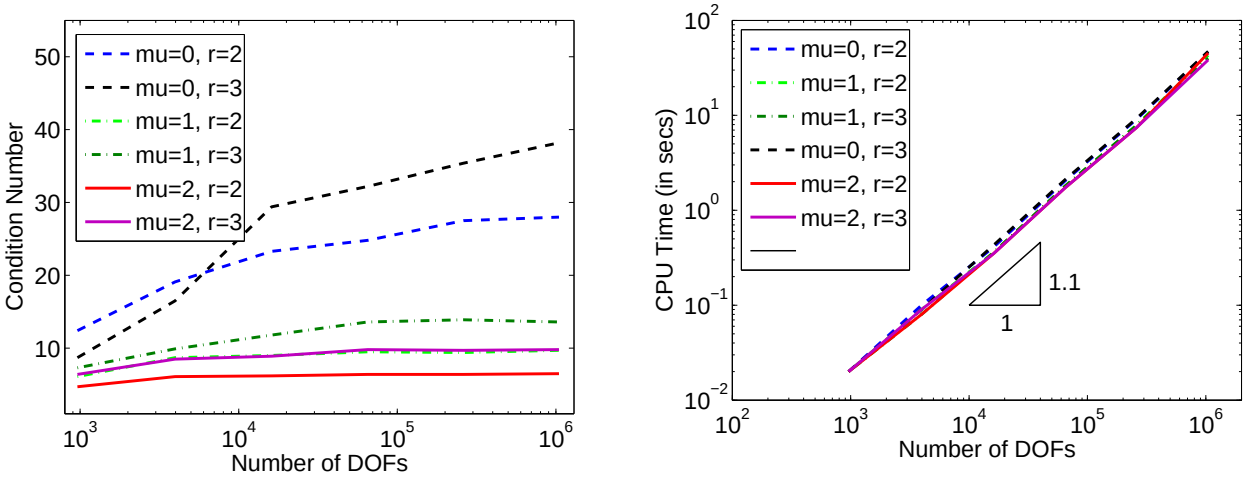

Figure 4: Condition numbers and CPU times for Laplacian with $\delta^{sub} = 3h$ (varying n).

parameters as above but varying the number n of degrees of freedom we see that the condition number is (asymptotically) independent of n (and thus of h) for various choices of the number of smoothing steps μ and of the aggregation radius r. We observe again that $\kappa\left(M_{AS}^{-1}A\right) \leq 5\frac{H}{\delta}$. (Note that the method does not break down in the case $\mu = 2$ where $\delta = 5h > \delta^{sub}$ and our theoretical assumption (C2) is violated.) The CPU times are growing approximately like $O(n^{1.1})$ which is almost optimal. The growth stems from the fact that our coarse problem size and the subdomain problem sizes in the tests in Figure 4 are growing proportionally to n. However, this increase in the CPU time is extremely mild, since `UMFPACK` scales very well for reasonably small problems (up to $n = 10^5$ degrees of freedom), i.e. it scales approximately like $O(n^{1.1})$. We are able to exploit this good performance of sparse direct solvers fully here. Note that for larger problems (say $n \approx 10^6$ or bigger) `UMFPACK` and other sparse direct solvers start to slow down dramatically, e.g. `UMFPACK` scales only like $O(n^{1.9})$ for $n \approx 10^6$ on our system, while our method continues to scale like $O(n^{1.1})$. Note also that there seems to be little dependency of the computational efficiency on small changes in the number of smoothing steps μ or the aggregation radius r.

As we said already earlier, we are also interested in the robustness of our method to large jumps in the value of the coefficient function α. To test this we choose α as a realisation of a log-normal random field, i.e $\log \alpha(x)$ is a realisation of a homogeneous, isotropic Gaussian random field with exponential covariance function, mean 0, variance σ^2 and correlation length scale λ. This is a commonly studied model for flow in heterogeneous porous media. For more details on the physical background see e.g. Cliffe et al.[7]. We use `Gaussian`[16] to create these random fields. See Figure 5 (left) for a grey-scale plot of a typical realisation: Black areas in the plot represent large values of α, white areas represent small values of α. The larger the correlation length λ, the more correlated (and thus smoother) is the field. The larger the variance σ^2, the larger is the difference between

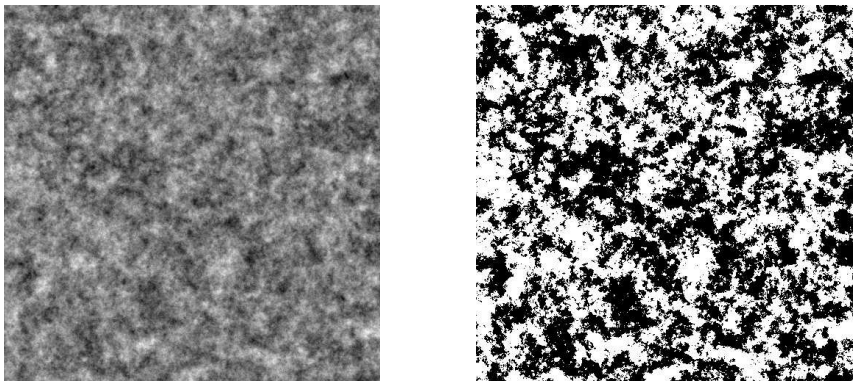

Figure 5: Log-normal and "clipped" log-normal random fields with $n = 512^2$ and $\lambda = 1/64$.

large and small values of α. For example for the field in Figure 5 for $\sigma^2 = 8$ we have $\sup_{x,y\in\Omega} \frac{\alpha(x)}{\alpha(y)} \approx 3 * 10^{10}$.

As an even harder test for our method we use "clipped" log-normal random fields, i.e. the smallest 50% and the largest 50% of the values of $\log \alpha(x)$ are each set to their average value (see Figure 5 (right)). The size and the "roughness" of the areas with small and large coefficients is again related to the correlation length λ. The size of the jump in the value of α is related to the variance σ^2. For the clipped field in Figure 5 for $\sigma^2 = 8$ we have $\sup_{x,y\in\Omega} \frac{\alpha(x)}{\alpha(y)} \approx 5*10^5$. Although this fraction is smaller here, the fact that α changes very rapidly throughout the domain and that the size of the jump at each discontinuity is of the order 10^5 makes it a more challenging problem for the linear solver. Note that clipped random fields play an important rôle in the modelling of *emergent* (electrical, mechanical or thermal) behaviour of micro-structures[1].

We will now test the robustness of our method (denoted `ADOUG` below) in the case of these clipped log-normal fields and compare its performance with the (aggregation-type) AMG code of Bastian[2,a], with the multifrontal sparse direct solver `UMFPACKv4.4`[8], and with the "classical" additive Schwarz solver `DOUG`[10,11,b] which uses graph partitioning software to construct the subdomains and a standard linear FE coarse space. In all the tests below we choose $r = 2$, $\varepsilon = \frac{2}{3}$, $\delta^{sub} = 3h$ and $\mu = 0$, i.e. no smoothing. See Figure 1 (right) for a typical set of aggregates.

We begin in Table 1 by fixing $n = 256^2$ and $\lambda = 1/64$ and by varying σ^2. We see that the new method `ADOUG` is extremely robust with respect to the size of the jumps. Both the number of iterations and the CPU-time do not grow with σ^2. All the other methods show a dependency on the size of the jump, even the direct solver `UMFPACK`. This is due

[a] Note that for efficiency reasons Bastian's AMG code uses unsmoothed piecewise constant prolongation and can therefore not be expected to be independent of the problem size n.

[b] Note that `DOUG` is a parallel code that uses a master-slave model. All timings for `DOUG` are for a parallel run with one processor handling the coarse solve and one doing the rest. Due to the slow network speed on our system the CPU-times are pessimistic.

CG–iterations

σ^2	$\sup_\Omega \frac{\alpha(x)}{\alpha(y)}$	ADOUG	AMG	DOUG
2	$1.5 * 10^1$	24	14	32
4	$2.2 * 10^2$	27	27	89
6	$3.3 * 10^3$	29	40	296
8	$4.9 * 10^4$	26	77	498
10	$7.4 * 10^5$	26	27	724

CPU–time (in secs)

σ^2	ADOUG	AMG	DOUG	UMFPACK
2	2.12	1.35	5.54	1.85
4	2.14	2.27	8.18	1.70
6	2.34	3.31	19.1	1.33
8	2.41	6.23	29.9	4.88
10	2.37	2.39	42.2	4.98

Table 1: Comparison of solvers for clipped random fields with $n = 256^2$ and $\lambda = 1/64$.

to the extra cost for partial pivoting in the case of largely varying diagonal entries in A. The classical additive Schwarz method DOUG does not cope at all with this problem.

In Table 2 we fix $n = 256^2$ and $\sigma^2 = 8$ and study the behaviour of all the methods as we vary the correlation length λ. Again ADOUG is extremely robust and does not vary at all for correlation lengths of size $\lambda \geq 4h$. Only for extremely short correlation lengths (i.e. close to the size of the fine mesh width h) do we start to see any deterioration, and even then the number of iterations does not even double. AMG shows a much stronger dependency on λ and seems to have real problems with short correlation lengths. UMFPACK also is affected strongly by the correlation length.

CG–iterations

λ	ADOUG	AMG	DOUG
1/17	26	18	355
1/33	27	64	430
1/65	26	77	498
1/129	33	70	655
1/257	48	166	858

CPU–time (in secs)

λ	ADOUG	AMG	DOUG	UMFPACK
1/17	2.20	1.67	22.3	4.52
1/33	2.24	5.14	26.3	4.77
1/65	2.41	6.23	29.9	4.88
1/129	2.71	5.77	38.2	7.48
1/257	3.84	13.5	49.5	10.2

Table 2: Comparison of solvers for clipped random fields with $n = 256^2$ and $\sigma^2 = 8$.

In Table 3 we carry out the comparison for fixed $\sigma^2 = 8$ varying the mesh width h, with λ linked to h by $\lambda = 4h$. Note that this means that the problem actually gets harder the more the mesh is refined (not only because of the growing problem size n). As before ADOUG is robust for most of the range but starts to struggle slightly for $h = 1/1025$. However, the growth in the number of iterations is much milder than that of AMG. Moreover, the growth in the CPU-times for ADOUG is even milder, it grows like $O(n^{1.3})$ which is almost as good as in the case of the Laplacian (i.e. $\alpha \equiv 1$) in Figure 4. UMFPACK starts to bail out at $h = 1/512$ and the classical additive Schwarz method DOUG again does not cope either.

To finish we give one set of results with an unclipped log-normal random field in Table

CG–iterations

h	ADOUG	AMG	DOUG
1/65	20	12	60
1/129	25	35	136
1/257	26	77	498
1/513	34	100	1111
1/1025	74	422	****

CPU–time (in secs)

h	ADOUG	AMG	DOUG	UMFPACK
1/65	0.10	0.06	0.89	0.05
1/129	0.46	0.68	2.62	0.52
1/257	2.41	6.23	29.9	4.88
1/513	16.8	33.8	258	88.8
1/1025	105.9	540	****	****

Table 3: Comparison of solvers for clipped random fields with $\sigma^2 = 8$ and $\lambda = 4h$.

4, to show that this case is indeed simpler and that our method also works here. The only new thing to observe is that in this case the performance of the classical additive Schwarz method DOUG is strongly improved.

	ADOUG	AMG	DOUG	UMFPACK
Iterations	19	38	62	
CPU-time	8.3	13.1	29.7	10.3

Table 4: Comparison for an unclipped log-normal field with $n = 512^2$, $\sigma^2 = 8$ and $\lambda = 8h$.

Acknowledgements. We would like to thank Professor Peter Bastian for kindly providing us with his AMG code (including the source code), and we are very grateful to Professor Ivan Graham for many useful discussions regarding the theory.

REFERENCES

[1] D.P. Almond, C.R. Bowen and D.A.S. Rees, Composite dielectrics and conductors: simulation, characterization and design, *J. Phys. D: Appl. Phys.*, **39**, 1295–1304 (2006)

[2] P. Bastian, Source code of aggregation-based algebraic multigrid, IWR, University of Heidelberg, 1996.

[3] M. Brezina and P. Vanek, A black-box iterative solver based on a two-level Schwarz method, *Computing*, **63**, 233–263 (1999).

[4] X. Cai, B.F. Nielsen and A. Tveito, An analysis of a preconditioner for the discretized pressure equation arising in reservoir simulation, *IMA J Num. Anal.*, **19**, 291–316 (1999).

[5] L.M. Carvalho, L.Giraud and P. Le Tallec, Algebraic two-level preconditioners for the Schur complement method, *SIAM J. Sci. Comp.*, **22**, 1987–2005 (2001).

[6] T.F. Chan and T. Mathew, Domain Decomposition Methods, *Acta Numerica* 1994 (Cambridge University Press, 1994).

[7] K.A. Cliffe, I.G. Graham, R. Scheichl and L. Stals, Parallel computation of flow in heterogeneous media modelled by mixed finite elements, *J. Comp. Phys.*, **164**, 258–282 (2000).

[8] T.A. Davies, `UMFPACKv4.4` multifrontal software, CISE, University of Florida, 2005, available at `http://www.cise.ufl.edu/research/sparse/umfpack`

[9] L. Giraud, F. Guevara Vasquez and R.S. Tuminaro, Grid transfer operators for highly-variable coefficient problems in two-level non-overlapping domain decomposition methods, *Numer. Linear Algebra Appl.*, **10**, 467–484 (2003).

[10] I.G. Graham and M.J. Hagger, Additive Schwarz, CG and discontinuous coefficients, in *Proc. Ninth Intern. Confer. on Domain Decomposition Methods*, Bergen, 1996, P. Bjorstad, M. Espedal and D.E. Keyes (Domain Decomposition Press, Bergen, 1998).

[11] I.G. Graham and M.J. Hagger, Unstructured additive Schwarz-CG method for elliptic problems with highly discontinuous coefficients, *SIAM J Sci. Comp.*, **20**, 2041–2066 (1999).

[12] I.G. Graham, P.O. Lechner and R. Scheichl, Domain decomposition for multiscale PDEs, in preparation (2006).

[13] M. Griebel, K. Scherer and M.A. Schweitzer, Robust norm equivalences and optimal preconditioners for diffusion problems, *Math. Comp.*, to appear (2006).

[14] E.W. Jenkins, C.T. Kelley, C.T. Miller and C.E. Kees, An aggregation-based domain decomposition preconditioner for groundwater flow, *SIAM J Sci. Comp.*, **23**, 430–441 (2001).

[15] B. Kozintsev and B. Kedem, `Gaussian` package, University of Maryland, 1999, available at `http://www.math.umd.edu/~bnk/bak/generate.cgi`

[16] C. Lasser and A. Toselli, Convergence of some two-level overlapping domain decomposition preconditioners with smoothed aggregation coarse spaces. In L. Pavarino and A. Toselli (eds.), *Recent Developments in Domain Decomposition Methods*, Lecture Notes in Computational Science and Engineering, Vol. 23, 95–117, Springer (2002).

[17] M. Sala, Analysis of two-level domain decomposition preconditioners based on aggregation, *ESAIM M2AN*, **38**:5, 765–780 (2004).

[18] M. Sala, J.N. Shadid and R.S. Tuminaro, An improved convergence bound for aggregation-based domain decomposition preconditioners, *SIAM J. Matrix Anal.*, **27**, 744–756 (2006).

[19] R. Scheichl and E. Vainikko, Robust coarsening for overlapping domain decomposition preconditioners: Theory and method, in preparation (2006).

[20] A. Toselli and O. Widlund, *Domain Decomposition Methods Algorithms and Theory*, Springer (2005).

[21] P. Vanek, J. Mandel and M. Brezina, Algebraic multigrid by smoothed aggregation for 2nd and 4th order elliptic problems, *Computing*, **56**, 179–196 (1996).

[22] C. Vuik, A. Segal and J.A. Meijerink, An efficient preconditioned CG method for the solution of a class of layered problems with extreme contrasts in the coefficients, *J. Comp. Phys.*, **152**, 385–403 (1999).

[23] W.L. Wan, T.F. Chan and B. Smith, An energy-minimizing interpolation for robust multigrid methods, *SIAM J Sci. Comput.*, **21**, 1632–1649 (2000).