Computing (2007) 80: 319–343
DOI 10.1007/s00607-007-0237-z
Printed in The Netherlands

Computing

Additive Schwarz with aggregation-based coarsening for elliptic problems with highly variable coefficients

R. Scheichl[1] and **E. Vainikko**[2]

[1]Bath Institute for Complex Systems, Department of Mathematical Sciences, University of Bath, Bath, UK

[2]Institute of Computer Science, University of Tartu, Tartu, Estonia

Received 30 May 2006; Accepted 8 June 2007; Published online 30 July 2007

Summary

We develop a new coefficient-explicit theory for two-level overlapping domain decomposition preconditioners with non-standard coarse spaces in iterative solvers for finite element discretisations of second-order elliptic problems. We apply the theory to the case of smoothed aggregation coarse spaces introduced by Vanek, Mandel and Brezina in the context of algebraic multigrid (AMG) and 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 strong connections (suitably adapted from the AMG 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 and make explicit the interplay between the coefficient function and the coarse space basis functions. 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 and to a sparse direct solver.

AMS Subject Classifications: 65F10; 65N22; 65N55.

Keywords: second-order elliptic problems; heterogeneous media; smoothed aggregation; two-level additive Schwarz; algebraic multigrid.

1. Introduction

This paper extends work by Brezina and Vanek [6], Lasser and Toselli [24], and Sala [27] on the theory of two-level additive Schwarz preconditioners with non-standard coarse spaces, in particular those based on smoothed aggregation techniques first introduced by Vanek et al. [31], [32] in the context of algebraic multigrid methods.

Correspondence: R. Scheichl, Bath Institute for Complex Systems, Department of Mathematical Sciences, University of Bath, Claverton Down, Bath BA2 7AY, UK (E-mail: r.scheichl@maths.bath.ac.uk)

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 with $f \in L_2(\Omega)$ and 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 Eq. (1) using continuous piecewise linear finite elements $\mathcal{V}^h$ on a shape-regular 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 (see e.g., [5]). Moreover, as noted in [16], the condition number will also depend on $\sup_{x,y\in\Omega} \frac{\alpha(x)}{\alpha(y)}$. To improve the conditioning, suppose that Ω is covered by a set $\{\Omega_i : i = 1, \ldots, s\}$ of overlapping subdomains such that $\bigcup_{i=1}^{s} \bar{\Omega}_i = \bar{\Omega}$ and such that diam $\Omega_i \leq H^{\text{sub}}$. Secondly, suppose that we also have a coarse space $\mathcal{V}_0 \subset \mathcal{V}^h$. In our case the coarse space $\mathcal{V}_0 := \text{span}\{\Phi_j : j = 1, \ldots, N\}$ can be quite general for theoretical purposes, but in practice it 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 $\leq H$, by summing the associated fine grid basis functions and by "smoothing" the result [31]. Let M_{AS}^{-1} denote the classical two-level additive Schwarz preconditioner, obtained by solving discretisations of Eq. (1) on each of the overlapping subdomains as well as on the coarse space $\mathcal{V}_0$ (see e.g., [29]).

The main theoretical result of this paper is to improve the bounds for $\kappa(M_{\text{AS}}^{-1}A)$ in [24], [27], [28]. In particular, we are able to show that $\kappa(M_{\text{AS}}^{-1}A)$ is independent of H^{sub}, the size of the subdomains, and only depends linearly on the local ratio of the diameter of the supports of the coarse space basis functions and of the size of their overlap. Our numerical experiments with $\alpha \equiv 1$ (i.e., the Laplacian) show that this bound is sharp in the case of smoothed aggregation. Note that as we will see in the numerical experiments, for efficiency reasons it is of interest to choose $H^{\text{sub}} \gg H$. The estimates in [24], [27], [28] 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. The dependency of the condition number $\kappa(M_{\text{AS}}^{-1}A)$ 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 that of its neighbours, i.e., roughly speaking, provided $\nabla\Phi_j(x)$ is small wherever $\alpha(x)$ is large, then $\kappa(M_{\text{AS}}^{-1}A)$ 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 (both overlapping and non-overlapping) 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 [9]. Another class of results (i.e., [7], [14], [15], [33]) 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. Other related results from the domain decomposition and multigrid literature can be found in [34], [22], [8], [13], [19].

Note that the theoretical results in this paper are strongly related to the results in the recent paper [16], but that here we assume no underlying coarse grid. The resulting bounds and the ways in which they are proved are different to the theory in [16]. For a comparison see [18]. Extensions to hybrid additive-multiplicative Schwarz and to deflation methods are discussed in [17].

To test numerically the resilience of our new method (available at [30]) 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. Here we restrict to non-smoothed aggregation, and the method proves indeed to be robust with respect to variations in α. As predicted by our theory (and by the theory in [16]), two-level overlapping additive Schwarz with standard piecewise linear coarsening is not robust on the same test problem; neither seems to be aggregation-based algebraic multigrid (AMG) with Gauss-Seidel smoothing. However, since we only tested one implementation of aggregation-based AMG due to Bastian [2]–[4], it remains to be seen whether different smoothers (such as those employed in the `ML` library [12]) might lead to a multigrid method based on aggregation that is robust on our test problems. We note that the standard `AMG1R5` by Ruge and Stüben [26] and the related `BoomerAMG` from the `HYPRE` library [20] are robust on our test problems (even with standard Gauss–Seidel smoothing) and more efficient than our method. However, this class of AMG methods uses a different (heuristic) strategy to design coarse levels and unfortunately no theory exists yet that explains their robustness to strong variations in α. We hope that our theoretical work will ultimately provide some important insight into this robustness.

We start in the next section by defining the basic notation and tools. Section 3 contains the main theoretical results. In Sect. 4, we construct via smoothed aggregation coarse space basis functions which satisfy the assumptions made in Sect. 3, and finish in Sect. 5 with numerical results to verify the sharpness of our theoretical results.

2. Preliminaries

Let Ω be a bounded, open, polygonal (polyhedral) domain in $\mathbb{R}^d$, with $d = 2$ or 3, with boundary $\partial\Omega$, and let $\mathcal{T}^h$ be a family of conforming meshes $\mathcal{T}^h$ (triangles in 2D, tetrahedra in 3D), which are shape-regular as the mesh diameter $h \to 0$. A typical element of $\mathcal{T}^h$ is $\tau \in \mathcal{T}^h$ (a closed subset of $\bar{\Omega}$). If W is any subset of $\bar{\Omega}$ then $\mathcal{N}^h(W)$ will denote the set of nodes of $\mathcal{T}^h$ which also lie in W and, using a suitable index set $\mathcal{I}^h(W)$, we write this as

$$\mathcal{N}^h(W) = \{x_p : p \in \mathcal{I}^h(W)\}. \tag{2}$$

In particular, $\mathcal{N}^h(\bar{\Omega})$ is the set of all nodes of the mesh, including boundary nodes, and $\mathcal{N}^h(\Omega)$ is the set of all interior nodes.

Let $H^1(\Omega)$ and $H^1_0(\Omega)$ denote the usual Sobolev spaces and let $S^h(\Omega)$ denote the subspace of $H^1(\Omega) \cap C(\bar{\Omega})$, consisting of continuous piecewise linear functions with respect to $\mathcal{T}^h$. Let $\{\varphi_p : p \in \mathcal{I}^h(\bar{\Omega})\}$ denote the set of hat functions corresponding to the nodes $\mathcal{N}^h(\bar{\Omega})$ and set $S^h_0(\Omega) := S^h(\Omega) \cap H^1_0(\Omega)$. Suppose D is a polygonal (polyhedral) subdomain of Ω, such that $\bar{D}$ is a union of elements from $\mathcal{T}^h$. Then we let $|D|$ denote the volume of D and $S^h_0(D)$ the subset of $S^h_0(\Omega)$ consisting of functions whose support is contained in $\bar{D}$.

We consider the bilinear form arising from Eq. (1):

$$a(u, v) := \int_\Omega \alpha \nabla u \cdot \nabla v \, \mathrm{d}x, \quad u, v \in H^1_0(\Omega), \tag{3}$$

and its Galerkin approximation in the n-dimensional space $\mathcal{V}^h := S^h_0(\Omega)$. Let A be the corresponding $n \times n$ stiffness matrix, i.e.,

$$A_{pq} := \int_\Omega \alpha \nabla \varphi_p \cdot \nabla \varphi_q \, \mathrm{d}x, \quad p, q \in \mathcal{I}^h(\Omega). \tag{4}$$

The finite element solution of Eq. (1) is then equivalent to solving the linear system

$$A\mathbf{U} = \mathbf{b} \quad \text{with } b_p = \int_\Omega f \varphi_p \, \mathrm{d}x, \tag{5}$$

where $\mathbf{U}$ is the coefficient vector for the finite element solution u_h of Eq. (1) in the basis $\{\varphi_p\}$.

We are interested in iterative methods for solving Eq. (5) and hence in preconditioners for A which remove the ill-conditioning due to both the non-smoothness of α and the smallness of h. We will be concerned with preconditioners based on domain decomposition. These will be defined using solves on local subdomains and in a global coarse space as follows.

Note that A depends on α only through the quantities $\int_\tau \alpha(x) \, \mathrm{d}x$, $\tau \in \mathcal{T}^h$, so henceforth we shall assume that α has constant value α_τ on each τ. Throughout the paper, the notation $X \lesssim Y$ (for two quantities X, Y) means that X/Y is bounded above independently, not only of the *mesh parameter* h and of the domain decomposition parameters (i.e., s, N, H_j, δ_j, diam Ω_i, etc.) introduced below, but also of the *coefficient values* $\{\alpha_\tau : \tau \in \mathcal{T}^h\}$. Moreover $X \sim Y$ means that $X \lesssim Y$ and $Y \lesssim X$.

2.1 Two-level overlapping additive Schwarz preconditioners

Let $\{\Omega_i : i = 1, \ldots, s\}$ be an overlapping, open covering of Ω. Each $\bar{\Omega}_i$ is assumed to consist of a union of elements from $\mathcal{T}^h$, and each point $x \in \Omega$ is assumed to

be covered by a number of subdomains which is independent of the mesh size h. Note that we will neither make any assumptions on the shape of the subdomains Ω_i nor on the way they overlap. However the conditions on our coarse space basis functions below will implicitly induce some assumptions on the subdomains. Having introduced the subdomains, we also introduce, for each Ω_i, the local subspace $\mathcal{V}_i := S_0^h(\Omega_i)$ of $\mathcal{V}^h$. Then, for $p \in \mathcal{I}^h(\Omega_i)$ and $q \in \mathcal{I}^h(\Omega)$, we define the matrix $(R_i)_{p,q} = \delta_{p,q}$ and set $A_i := R_i A R_i^{\mathrm{T}}$, which is just the minor of A corresponding to rows and columns taken from $\mathcal{I}^h(\Omega_i)$.

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 Sect. 4 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

$$\omega_j := \text{interior}\big(\text{supp}\{\Phi_j\}\big).$$

We make the following assumptions:

(C1) For all $j = 1, \ldots, N_H$ there exists an $i \in \{1, \ldots, s\}$ such that $\omega_j \subset \Omega_i$.
(C2) $\sum_{j=1}^{N_H} \Phi_j(x) = 1$, for all $x \in \bar{\Omega}$.
(C3) $\|\Phi_j\|_{L_\infty(\Omega)} \lesssim 1$.

Note that Assumption (C2) guarantees that $\{\omega_j\}$ is a covering of Ω. Also, since Φ_j is a finite element function, ω_j will consist of the union of a set of fine grid elements $\tau \in \mathcal{T}^h$.

For theoretical purposes only, we also need to make a series of further assumptions on the sets ω_j. Let $H_j := \text{diam}\{\omega_j\}$ and $H := \max_{j=1,\ldots,N_H} H_j$. First of all we assume that each of the sets ω_j is simply connected. We assume further that the covering $\{\omega_j\}$ is shape regular as $H \to 0$, i.e., for all $j = 1, \ldots, N_H$: $H_j^d \lesssim |\omega_j|$ and $H_j \sim H_{j'}$ if $\omega_j \cap \omega_{j'} \neq \emptyset$ (*shape regularity*). We also need to make an assumption on the overlap between the ω_j. Let $\gamma_j := \partial\omega_j \cap \Omega$, i.e., the "interior boundary" of ω_j, and let

$$\mathring{\omega}_j = \{x \in \omega_j \ : x \notin \bar{\omega}_{j'} \text{ for any } \ j' \neq j\}, \tag{6}$$

be the subset of ω_j which is not overlapped by any of the other supports $\omega_{j'}$. We assume that the overlap between the supports ω_j of the coarse space basis functions is uniform, i.e., for each $j = 1, \ldots, N_H$ there exists a parameter δ_j, the "overlap parameter", such that

$$\text{dist}(x, \mathring{\omega}_j) \ \sim \ \delta_j \quad \text{for all } x \in \gamma_j \tag{7}$$

(*uniform overlap*). If $\mathring{\omega}_j = \emptyset$, we set $\delta_j = H_j$. Finally, we assume that each point $x \in \Omega$ is covered by a number of supports ω_j which is independent of the mesh size h (*finite covering*).

To simplify our notation we assume that the functions Φ_j are numbered in such a way that $\Phi_j|_{\partial\Omega} = 0$ for all $j \leq N$ and $\Phi_j|_{\partial\Omega} \neq 0$ for all $j > N$, with $N \leq N_H$, i.e., for all $j \leq N$ we have $\Phi_j \in \mathcal{V}^h$. 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

$$(R_0)_{j,p} = \Phi_j(x_p), \quad p \in \mathcal{I}^h(\Omega), \quad j = 1, \ldots, N,$$

then the matrix $A_0 := R_0 A R_0^{\mathrm{T}}$ 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

$$M_{\mathrm{AS}}^{-1} = \sum_{i=0}^{s} R_i^{\mathrm{T}} A_i^{-1} R_i. \tag{8}$$

Note that although we have not directly made any assumptions on the subdomains $\{\Omega_i\}$, Assumption (C1) implies that $H^{\mathrm{sub}} := \max_{i=1}^{s} \mathrm{diam}(\Omega_i) \geq H$ and that the minimum overlap δ^{sub} between any two of the subdomains is in fact bounded from below by the minimum overlap of the supports $\{\omega_j\}$.

Before studying in Sect. 3 the dependency of the condition number of $M_{\mathrm{AS}}^{-1}A$ on the mesh parameters and on the coefficient function α, we need to recall two well-known classical results which are central in the analysis of domain decomposition preconditioners of Schwarz type.

2.2 Basic properties of Schwarz preconditioners

For any vectors $\mathbf{V}, \mathbf{W} \in \mathbb{R}^n$, let $\langle \mathbf{V}, \mathbf{W} \rangle_A = \mathbf{V}^{\mathrm{T}} A \mathbf{W}$ denote the inner product induced by A. For any $u \in \mathcal{V}^h$, let $\mathbf{U} \in \mathbb{R}^n$ denote its corresponding vector of coefficients with respect to the nodal basis $\{\varphi_p\}$. Then it is easily shown that the matrices $R_i^{\mathrm{T}} A_i^{-1} R_i A$ are symmetric and positive definite with respect to the inner product $\langle \cdot, \cdot \rangle_A$. Moreover it is a standard observation that

$$\left\langle R_i^{\mathrm{T}} A_i^{-1} R_i A \mathbf{U}, \mathbf{U} \right\rangle_A = a(P_i u, u), \tag{9}$$

where P_i denotes the orthogonal projection onto $\mathcal{V}_i$ with respect to $a(\cdot,\cdot)$. From this, one obtains the following classical results (see [29]) which relate the properties of the subspaces $\mathcal{V}_i$ to the properties of the two-level additive Schwarz preconditioner M_{AS}^{-1}. For any symmetric positive definite matrix B, let $\lambda_{\max}(B)$ and $\lambda_{\min}(B)$ denote its maximum and minimum eigenvalues, respectively.

Theorem 2.1 (Colouring argument): *The collection of subspaces $\{\mathcal{V}_i : i = 1, \ldots, s\}$ can be coloured by N_c different colours so that when $\mathcal{V}_i$ and $\mathcal{V}_{i'}$ have the same colour, we necessarily have $\mathcal{V}_i$ and $\mathcal{V}_{i'}$ mutually orthogonal in the inner product induced by a, and*

$$\lambda_{\max}(M_{\mathrm{AS}}^{-1}A) \leq N_c + 1.$$

Theorem 2.2 (Stable splitting): *Suppose that there exists a constant C_0, such that every $u_h \in \mathcal{V}^h$ admits a decomposition*

$$u_h = \sum_{i=0}^{s} u_i, \quad \text{with } u_i \in \mathcal{V}_i, \ \ i = 0, \ldots, s \qquad \text{and} \qquad \sum_{i=0}^{s} a(u_i, u_i) \leq C_0^2 a(u_h, u_h).$$

Then

$$\lambda_{\min}(M_{\mathrm{AS}}^{-1} A) \geq C_0^{-2}.$$

3. General framework for analysis

In this section, we provide a general framework for the analysis of two-level additive Schwarz preconditioners for Eq. (5) in which the dependency of the condition number on α as well as on the mesh parameters is made precise. Section 3.1 aims to clarify the dependency on the mesh parameters and gives only a very crude estimate with respect to the coefficient function $\alpha(x)$. In particular it is shown there that the condition number is independent of the diameter of the subdomains. Section 3.2 aims to give a sharper bound of the condition number with respect to possible variations in $\alpha(x)$ and makes explicit the interplay between the coefficient function and the coarse space basis functions.

3.1 Improving the dependency on the mesh parameters

We have seen at the end of the previous section that (provided each $x \in \Omega$ is covered by finitely many subdomains Ω_i) all that is needed to prove a bound for the condition number of the preconditioned stiffness matrix in the case of the two-level additive Schwarz preconditioner is a stable splitting for any $u_h \in \mathcal{V}^h$ in the sense of Theorem 2.2. It has already been noted in [24], [27], [28] that the size of the stability constant C_0 in Theorem 2.2 depends on the assumptions which one makes on the gradient of the coarse space basis functions. Let us define the following two assumptions:

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

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

Note that (C4b) implies (C4a) and is therefore a stronger assumption (see [29, sect. 3.10.1]).

In the case of $N_H = s$ (i.e., one coarse space basis function per subdomain) and $\alpha \equiv 1$ the existence of a stable splitting has been proved in [24]. It can also be found in [29, theorems 3.17 and 3.19]:

Lemma 3.1: *Assume that* (C1)–(C3) *hold true and that in addition $N_H = s$. Then, for every $u_h \in \mathcal{V}^h$, there exists a decomposition*

$$u_h = \sum_{i=0}^{s} u_i, \quad \text{with } u_i \in \mathcal{V}_i, \ \ i = 0, \ldots, s$$

such that

$$\sum_{i=0}^{s} |u_i|^2_{H^1(\Omega)} \lesssim \left(1 + \max_{j=1,\dots,N_H} \frac{H_j}{\delta_j}\right)^{\beta} |u_h|^2_{H^1(\Omega)},$$

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

The following theorem is a direct consequence of Theorems 2.1 and 2.2 and Lemma 3.1.

Theorem 3.2: *Assume that* (C1)–(C3) *hold true and that in addition* $N_H = s$ *and* $\alpha \equiv 1$. *Then*

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

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 $N_H \gg s$, or in other words subdomains of much larger diameter than H. Let $\delta := \min_{j=1,\dots,N_H} \delta_j$ and suppose that (as above) $H^{\rm sub}$ and $\delta^{\rm sub}$ denote the maximum diameter of any subdomain Ω_i and the minimum overlap between any two subdomains, respectively. Then it has been shown in [27] that for $\alpha \equiv 1$ and under the weaker Assumption (C4a)

$$\kappa(M_{\rm AS}^{-1}A) \lesssim \left(1 + \frac{H^{\rm sub}}{\delta^{\rm sub}}\right)\left(1 + \frac{H}{\delta}\right). \tag{10}$$

This result has been improved in [28] using the stronger Assumption (C4b) to give

$$\kappa(M_{\rm AS}^{-1}A) \lesssim \left(1 + \frac{H^{\rm sub}}{\delta^{\rm sub}} + \frac{H}{\delta}\right). \tag{11}$$

However, both these results are not sharp for $H^{\rm sub} \gg H$ as our numerical results in Sect. 5 will show. Indeed, our first main result in the following theorem, which can be obtained from Lemma 3.1 (even for $N_H \gg s$) using a colouring argument, shows that under Assumptions (C1)–(C3) the diameter of the subdomains does not affect $\kappa(M_{\rm AS}^{-1}A)$. In addition it also makes explicit the dependency of $\kappa(M_{\rm AS}^{-1}A)$ on α when $\alpha \not\equiv 1$.

Theorem 3.3: *Assume that* (C1)–(C3) *hold true. Then*

$$\kappa\left(M_{\rm AS}^{-1}A\right) \lesssim \max_{\tau,\tau' \in \mathcal{T}^h} \frac{\alpha_\tau}{\alpha_{\tau'}} \left(1 + \max_{j=1,\dots,N_H} \frac{H_j}{\delta_j}\right)^{\beta}, \tag{12}$$

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

Proof: This follows directly from Theorems 2.1 and 2.2, if we can find a stable splitting for each $u_h \in \mathcal{V}^h$ into elements $u_i \in \mathcal{V}_i$.

Let $u_h \in \mathcal{V}^h$. We know from Lemma 3.1 that there are functions $v_0 \in \mathcal{V}_0$ and $v_j \in \mathcal{S}_0^h(\omega_j)$, $j = 1, \ldots, N_H$, such that $u_h = \sum_{j=0}^{N_H} v_j$ and

$$\sum_{j=0}^{N_H} |v_j|^2_{H^1(\Omega)} \lesssim \left(1 + \max_{j=1,\ldots,N_H} \frac{H_j}{\delta_j}\right)^{\beta} |u_h|^2_{H^1(\Omega)}, \tag{13}$$

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

Let $\{\mathcal{I}_i : i = 1, \ldots, s\}$ be a non-overlapping partitioning of $\{1, \ldots, N_H\}$ such that $\omega_j \subset \Omega_i$ for all $j \in \mathcal{I}_i$ (i.e., $\cup\{\mathcal{I}_i : i = 1, \ldots, s\} = \{1, \ldots, N_H\}$ and $\mathcal{I}_i \cap \mathcal{I}_{i'} = \emptyset$ for $i \neq i'$). Assumption (C1) guarantees that such a partitioning exists. Note that $\mathcal{I}_i$ is the index set of all coarse space functions whose support is entirely contained in Ω_i. If ω_j is contained in more than one subdomain it is only assigned to one index set $\mathcal{I}_i$. Set

$$u_i = \sum_{j \in \mathcal{I}_i} v_j \quad \text{and} \quad u_0 = v_0. \tag{14}$$

Then $u_h = \sum_{i=0}^{s} u_i$ and Assumption (C1) implies that $u_i \in \mathcal{V}_i$ for all $i = 1, \ldots, s$.

To show that it is a stable splitting we use a colouring argument. We assumed in Sect. 2.1 that each point $x \in \Omega$ is covered by only finitely many supports ω_j. Therefore, we have for each i that the collection of subspaces $\{\mathcal{S}_0^h(\omega_j) : j \in \mathcal{I}_i\}$ can be coloured by N_c different colours so that when $\mathcal{S}_0^h(\omega_j)$ and $\mathcal{S}_0^h(\omega_{j'})$ have the same colour, we necessarily have that $\mathcal{S}_0^h(\omega_j)$ and $\mathcal{S}_0^h(\omega_{j'})$ are mutually orthogonal in the inner product induced by a. Then with u_i defined as in Eq. (14),

$$a(u_i, u_i) = \sum_{j, j' \in \mathcal{I}_i} a(v_j, v_{j'}) \leq N_c \sum_{j \in \mathcal{I}_i} a(v_j, v_j)$$

and hence it follows from Eq. (13) that

$$\begin{aligned} \sum_{i=1}^{s} a(u_i, u_i) &\leq N_c \max_{\tau \in \mathcal{T}^h} \alpha_\tau \sum_{j=1}^{N_H} |v_j|^2_{H^1(\Omega)} \lesssim \max_{\tau \in \mathcal{T}^h} \alpha_\tau \left(1 + \max_{j=1,\ldots,N_H} \frac{H_j}{\delta_j}\right)^{\beta} |u_h|^2_{H^1(\Omega)} \\ &\lesssim \max_{\tau, \tau' \in \mathcal{T}^h} \frac{\alpha_\tau}{\alpha_{\tau'}} \left(1 + \max_{j=1,\ldots,N_H} \frac{H_j}{\delta_j}\right)^{\beta} a(u_h, u_h). \end{aligned}$$

□

We will see below that the bound in Theorem 3.3 can still be improved with respect to the dependency on the coefficient function α. However, Theorem 3.3 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 is independent of the size of the subdomains and that (ii)

it only depends on local ratios of the sizes of the supports of the coarse space basis functions and of their overlap. The numerical results in Sect. 5 will show that this accurately reflects the dependency of the condition number on the mesh parameters in the case of aggregation-type coarse spaces.

3.2 Improving the dependency on variations in α

We have seen in the previous section that the size of the stability constant C_0 (and thus of the condition number bound) 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 $\alpha(x)$ over Ω:

Definition 3.4 (coarse space robustness indicator):

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

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 $\min_{\tau \in \mathcal{T}^h} \alpha_\tau$ without changing the condition number of the resulting discrete problem. For measurable $D \subset \Omega$, we define the weighted H^1-seminorm by

$$|f|^2_{H^1(D),\alpha} = \int_D \alpha |\nabla f|^2 \, \mathrm{d}x.$$

Note that $|f|^2_{H^1(\Omega),\alpha} = a(f,f)$. It follows trivially from the above assumption that

$$|f|^2_{H^1(D)} \leq |f|^2_{H^1(D),\alpha}, \quad \text{for all } \ f \in H^1(D). \tag{15}$$

The first result in this section examines the properties of quasi–interpolation on the abstract coarse space $\mathcal{V}_0$ and makes use of the coarse space robustness indicator.

Lemma 3.5: *There exists a linear operator $\tilde{I}_0 : H^1_0(\Omega) \to \mathcal{V}_0$ such that*

$$|\tilde{I}_0 u|^2_{H^1(\Omega),\alpha} \lesssim \gamma(\alpha) \max_{j=1,\dots,N_H} \frac{H_j}{\delta_j} \, |u|^2_{H^1(\Omega),\alpha} \quad \textit{for all } u \in H^1_0(\Omega).$$

Proof: The proof is obtained using the standard quasi-interpolant:

$$\tilde{I}_0 u := \sum_{j=1}^{N} \bar{u}_j \, \Phi_j, \qquad \text{where } \bar{u}_j := |\omega_j|^{-1} \int_{\omega_j} u \, \mathrm{d}x.$$

First note that for any $j = 1, \dots, N_H$ we have

$$|\bar{u}_j| \leq |\omega_j|^{-1} |\omega_j|^{1/2} \|u\|_{L_2(\omega_j)} = |\omega_j|^{-1/2} \|u\|_{L_2(\omega_j)}. \tag{16}$$

Also note that it follows from the definition of $\mathring{\omega}_j$ in Eq. (6) and from Assumption (C2) that

$$|\Phi_j|^2_{H^1(\omega_j),\alpha} = \int_{\omega_j \setminus \mathring{\omega}_j} \alpha |\nabla \Phi_j|^2 \, \mathrm{d}x \leq \| \alpha |\nabla \Phi_j|^2 \|_{L_\infty(\Omega)} \, |\omega_j \setminus \mathring{\omega}_j|$$

and so using Definition 3.4 together with the shape regularity and the overlap Condition (7), we have

$$|\Phi_j|^2_{H^1(\omega_j),\alpha} \lesssim \| \alpha |\nabla \Phi_j|^2 \|_{L_\infty(\Omega)} H_j^{d-1} \delta_j \leq \gamma(\alpha) \, \frac{H_j^{d-1}}{\delta_j}. \tag{17}$$

Now introduce the set

$$\tilde{\omega}_j := \cup \{ \omega_k : \omega_j \cap \omega_k \neq \emptyset \},$$

i.e., the union of the supports of the coarse space basis functions that intersect ω_j.

Let us first look at the case when $\tilde{\omega}_j$ does not touch $\partial\Omega$. In this case it follows from (C2) that

$$\sum_{j=1}^{N} \Phi_j(x) = \sum_{j=1}^{N_H} \Phi_j(x) = 1, \quad \text{for all } x \in \omega_j,$$

which implies that $(\tilde{I}_0 1)(x) = 1$ for all $x \in \omega_j$. Let $\hat{u} := u - |\tilde{\omega}_j|^{-1} \int_{\tilde{\omega}_j} u \, \mathrm{d}x$ and then use Eqs. (16), (17), and the shape regularity of $\{\omega_j\}$ to obtain

$$\begin{aligned} |\tilde{I}_0 u|^2_{H^1(\omega_j),\alpha} = |\tilde{I}_0 \hat{u}|^2_{H^1(\omega_j),\alpha} &\leq \max_{\{k:\omega_j \cap \omega_k \neq \emptyset\}} \left(|\omega_k|^{-1} \|\hat{u}\|^2_{L_2(\omega_k)} \right) |\Phi_k|^2_{H^1(\omega_k),\alpha} \\ &\lesssim \|\hat{u}\|^2_{L_2(\tilde{\omega}_j)} \max_{\{k:\omega_j \cap \omega_k \neq \emptyset\}} H_k^{-d} \, \gamma(\alpha) \, \frac{H_j^{d-1}}{\delta_j} \\ &\lesssim \gamma(\alpha) \, \frac{1}{H_j \delta_j} \|\hat{u}\|^2_{L_2(\tilde{\omega}_j)}. \end{aligned}$$

Thus, since $\hat{u}$ has zero mean on $\tilde{\omega}_j$, Poincaré's inequality (c.f. [29, corollary A.15]) implies

$$|\tilde{I}_0 u|^2_{H^1(\omega_j),\alpha} \lesssim \gamma(\alpha) \, \frac{H_j}{\delta_j} |\hat{u}|^2_{H^1(\tilde{\omega}_j)} = \gamma(\alpha) \, \frac{H_j}{\delta_j} |u|^2_{H^1(\tilde{\omega}_j)}. \tag{18}$$

Next, consider the case when $|\tilde{\omega}_j \cap \partial\Omega| \sim H_j^{d-1}$. Then, since $u = 0$ on $\partial\Omega$ we can apply Friedrich's inequality (c.f. [29, corollary A.14]) together with Eqs. (16), (17), and the shape regularity of $\{\omega_j\}$ to obtain

$$|\tilde{I}_0 u|^2_{H^1(\omega_j),\alpha} \lesssim \gamma(\alpha) \, \frac{1}{H_j \delta_j} \|u\|^2_{L_2(\tilde{\omega}_j)} \lesssim \gamma(\alpha) \, \frac{H_j}{\delta_j} |u|^2_{H^1(\tilde{\omega}_j)}. \tag{19}$$

Finally, the case of $\tilde{\omega}_j$ touching $\partial\Omega$ such that $|\tilde{\omega}_j \cap \partial\Omega| \ll H_j^{d-1}$ can be reduced to the latter case, by adding an additional set ω_k to $\tilde{\omega}_j$ such that $|\tilde{\omega}_j \cap \partial\Omega| \sim H_j^{d-1}$.

The result then follows from Eqs. (18) and (19) by summing over $j = 1, \ldots, N_H$ and by using Eq. (15) together with the fact that each point $x \in \Omega$ was assumed to lie in only finitely many of the supports ω_j. □

To prove the main theorem in this section we need in addition the following two technical lemmas. Let I_h denote the piecewise linear nodal interpolant on the fine mesh.

Lemma 3.6: *Let $v_h \in \mathcal{V}^h$. Then for all $j = 1, \ldots, N_H$,*

$$|I_h(\Phi_j v_h)|^2_{H^1(\omega_j),\alpha} \lesssim \|\alpha|\nabla\Phi_j|^2\|_{L_\infty(\omega_j)} \|v_h\|^2_{L_2(\omega_j\setminus\mathring{\omega}_j)} + |v_h|^2_{H^1(\omega_j),\alpha}.$$

Proof: See [16, lemma 3.5].[1] □

Lemma 3.7: *For all $j = 1, \ldots, N_H$ and for all $u \in H^1(\omega_j)$*

$$\|u\|^2_{L_2(\omega_j\setminus\mathring{\omega}_j)} \lesssim \delta_j^2 \left(\left(1 + \frac{H_j}{\delta_j}\right) |u|^2_{H^1(\omega_j)} + \frac{1}{H_j\delta_j} \|u\|^2_{L_2(\omega_j)} \right).$$

Proof: See [29, lemma 3.10].[2] □

Using these three lemmata together with Theorems 2.1 and 2.2 we can now prove the main result in this paper. It relies on an improved bound for the stability constant C_0 in Theorem 2.2.

Theorem 3.8: *Assume that* (C1)–(C3) *hold true. Then*

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

Proof: We proceed as in the proof to Theorem 3.3 by first constructing for each $u_h \in \mathcal{V}^h$ a stable splitting into elements v_j with $v_0 \in \mathcal{V}_0$ and $v_j \in \mathcal{S}_0^h(\omega_j)$, for $j = 1, \ldots, N_H$.

Let

$$v_0 := \tilde{I}_0 u_h = \sum_{j=1}^{N} \bar{u}_j \Phi_j \quad \text{where } \bar{u}_j := |\omega_j|^{-1} \int_{\omega_j} u_h \, \mathrm{d}x$$

and

$$v_j := \begin{cases} I_h\left(\Phi_j(u_h - \bar{u}_j)\right) & \text{for } j = 1, \ldots, N, \\ I_h\left(\Phi_j u_h\right) & \text{for } j = N+1, \ldots, N_H\,. \end{cases}$$

[1] Note that this is where Assumption (C3) is needed.

[2] Note that this is where the Assumption (7) of uniform overlap is needed again.

Then using (C2)

$$\sum_{j=0}^{N_H} v_j = \sum_{j=0}^{N} \left(\bar{u}_j \Phi_j + I_h(\Phi_j u_h) - \bar{u}_j \Phi_j\right) + \sum_{j=N+1}^{N_H} I_h(\Phi_j u_h)$$

$$= \sum_{j=1}^{N_H} I_h(\Phi_j u_h) = I_h\left(u_h \sum_{j=1}^{N_H} \Phi_j\right) = u_h.$$

Furthermore, by Lemma 3.5 we have

$$|v_0|^2_{H^1(\Omega),\alpha} \lesssim \gamma(\alpha) \max_{j=1,\dots,N_H} \frac{H_j}{\delta_j} |u_h|^2_{H^1(\Omega),\alpha}. \tag{20}$$

For all $j = 1, \dots, N$, Lemma 3.6 gives

$$|v_j|^2_{H^1(\omega_j),\alpha} \lesssim \|\alpha |\nabla \Phi_j|^2\|_{L_\infty(\omega_j)} \|u_h - \bar{u}_j\|^2_{L_2(\omega_j \setminus \mathring{\omega}_j)} + |u_h - \bar{u}_j|^2_{H^1(\omega_j),\alpha} \tag{21}$$

and it follows from Lemma 3.7 that

$$\|u_h - \bar{u}_j\|^2_{L_2(\omega_j \setminus \mathring{\omega}_j)} \lesssim \delta_j^2 \left(\left(1 + \frac{H_j}{\delta_j}\right) |u_h - \bar{u}_j|^2_{H^1(\omega_j)} + \frac{1}{H_j \delta_j} \|u_h - \bar{u}_j\|^2_{L_2(\omega_j)}\right).$$

Since $(u_h - \bar{u}_j)$ has zero mean over ω_j, we can apply Poincaré's inequality [29, corollary A.15] to the last term and use the resulting bound in Eq. (21) to obtain

$$|v_j|^2_{H^1(\omega_j),\alpha} \lesssim \delta_j^2 \|\alpha |\nabla \Phi_j|^2\|_{L_\infty(\omega_j)} \left(1 + \frac{H_j}{\delta_j}\right) |u_h|^2_{H^1(\omega_j)} + |u_h|^2_{H^1(\omega_j),\alpha}. \tag{22}$$

Similarly, for all $j = N+1, \dots, N_H$, Lemma 3.6 gives

$$|v_j|^2_{H^1(\omega_j),\alpha} \lesssim \|\alpha |\nabla \Phi_j|^2\|_{L_\infty(\omega_j)} \|u_h\|^2_{L_2(\omega_j \setminus \mathring{\omega}_j)} + |u_h|^2_{H^1(\omega_j),\alpha} \tag{23}$$

and it follows again from Lemma 3.7 that

$$\|u_h\|^2_{L_2(\omega_j \setminus \mathring{\omega}_j)} \lesssim \delta_j^2 \left(\left(1 + \frac{H_j}{\delta_j}\right) |u_h|^2_{H^1(\omega_j)} + \frac{1}{H_j \delta_j} \|u_h\|^2_{L_2(\omega_j)}\right).$$

Here we can use Friedrich's inequality [29, corollary A.14] to bound the last term, since $\partial\omega_j$ has non-trivial intersection with $\partial\Omega$ and $u_h|_{\partial\Omega} = 0$. We use the resulting bound in Eq. (23) to obtain Eq. (22) as before.

By summing over j and using the bounds (20) and (22), together with Eq. (15) and the assumption of a finite covering, we get

$$\sum_{j=0}^{N_H} |v_j|^2_{H^1(\Omega),\alpha} \lesssim \gamma(\alpha) \left(1 + \max_{j=1,\dots,N_H} \frac{H_j}{\delta_j}\right) |u_h|^2_{H^1(\Omega),\alpha}.$$

The proof is complete by using a colouring argument as in the proof of Theorem 3.3. Note that here we have already got a stable splitting $\{v_j\}$ in the weighted norm $|\cdot|_{H^1(\Omega),\alpha}$. Therefore no extra $\max_{\tau,\tau' \in \mathcal{T}^h} \frac{\alpha_\tau}{\alpha_{\tau'}}$ term appears. □

In principle, for an arbitrary coefficient function α and an arbitrary set of coarse space basis functions that satisfy (C4b), the quantity $\gamma(\alpha)$ can become as bad as $\max_{\tau,\tau'\in\mathcal{T}^h} \frac{\alpha_\tau}{\alpha_{\tau'}}$ which is the quantity that appears in Theorem 3.3. However, the huge improvement in Theorem 3.8 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 Sect. 4 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 $\max_{\tau,\tau'\in\mathcal{T}^h} \frac{\alpha_\tau}{\alpha_{\tau'}}$ goes to infinity.

4. Smoothed aggregation coarse spaces

Smoothed aggregation techniques have been introduced first in the context of algebraic multigrid methods by Vanek et al. [31], [32] and further investigated by Brezina and Vanek [6], Jenkins et al. [21], Lasser and Toselli [24], Sala [27], and Sala et al. [28] 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 [26], [31], [32].

To describe the smoothed aggregation algorithm that we use to construct a set of coarse space basis functions $\{\Phi_j : j = 1, \ldots, N_H\}$ we first 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 4.1: *(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*

$$|\tilde{\mathcal{A}}_{pq}| \geq \varepsilon \max_{k\neq p} |\tilde{\mathcal{A}}_{pk}|, \tag{24}$$

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

(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 is strongly connected *to x_q 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 the node x_p and all nodes x_q that are strongly connected to x_p with threshold ε and shortest path length $\ell_{pq} \leq r$.*

To our knowledge strongly-connected graph r-neighbourhoods have not yet been used in the context of domain decomposition methods. We have also not seen the

criterion (24) anywhere in the literature. It stems from the algebraic multigrid code of Bastian [2] and is a modified version of the criterion in [31]. Note that in contrast to the criterion in [31], criterion (24) is "directed", i.e., a node x_q may not be strongly connected to a node x_p, even if x_p is strongly connected to x_q.

The construction of the coarse space basis functions is now very similar to the algorithm described in [6]. 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, and (c) to minimise the number of nonzeros in the coarse matrix are different. They are inspired by Bastian [2], [3] and Raw [25]. The algorithm consists of two main steps: aggregation and smoothing. For both of those steps we make use of the so-called *filtered matrix* $\mathcal{A}^\varepsilon$ (as introduced originally in [31]) with entries

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

We first create a set of aggregates $\{W_j : j = 1, \ldots, N_H\}$ such that

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

(i.e., a non-overlapping partition of $\mathcal{N}$). Given an aggregation "radius" $r \in \mathbb{N}$ and a threshold $\varepsilon \in [0, 1]$, roughly speaking each of the sets W_j will be calculated by finding the strongly-connected graph r-neighbourhood $S_{r,\varepsilon}(x_j^H)$ of a suitably chosen seed node $x_j^H \in \mathcal{N}$. The aggregation procedure which is used in our code is outlined in the algorithm in Fig. 1. It uses an advancing front in the graph $\mathcal{G}$ to choose good seed nodes and to create the aggregates W_j.

To advance the front in each step of the algorithm we make use of the filtered matrix $\mathcal{A}^\varepsilon$. Note that given a node $x_p \in \mathcal{N}$, the set $S_\varepsilon(x_p)$ consists of the nodes x_q corresponding to the nonzero entries in the p-th row of $\mathcal{A}^\varepsilon$. This is readily available, if (the sparse matrix) $\mathcal{A}^\varepsilon$ is stored in compressed row storage format. To reduce the number of connections between neighbouring aggregates, and thus the number of nonzero elements in the coarse matrix A_0, the strongly-connected graph r-neighbourhood $S_{r,\varepsilon}(x_j^H)$ for a given seed node x_j^H is "rounded off" to produce the final aggregate W_j, by slightly enlarging the set with all nodes $x_q \notin S_{r,\varepsilon}(x_j^H)$ that are strongly connected to at least d nodes in $S_{r,\varepsilon}(x_j^H)$ (see step 4b in Fig. 1). This idea stems from Raw [25]. The seed nodes are chosen from a front (approximately) r nodes ahead of the "boundary" of the current set of aggregates, to give the new aggregates a better chance to grow and to decrease the chance of leaving a gap between new and old aggregates. The final step in our aggregation algorithm (i.e., step 7 in Fig. 1) aims to obtain aggregates of similar size and to avoid very small aggregates. (Note that this may not be possible, i.e., there may be an aggregate with less than $a_{\min}$ nodes that is not strongly connected to any other aggregate. In this case we keep the small aggregate.)

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

Input: Matrix A, set of nodes $\mathcal{N}$, threshold $\varepsilon \in [0,1]$, aggregation radius $r \in \mathbb{N}$, minimum aggregate size a_{min}, maximum aggregate size a_{max}.

Output: Set of aggregates $\{W_j : j = 1, \ldots, N_H\}$.

1. Calculate the filtered matrix $\mathcal{A}^{\varepsilon}$ and set $\mathcal{S} \leftarrow \emptyset$ and $j \leftarrow 0$.
2. Set $j \leftarrow j+1$ and choose a seed node x_j^H from $\mathcal{S}$ (if $\mathcal{S} = \emptyset$ choose an arbitrary node from $\mathcal{N}$).
3. Set layer $\mathcal{L}(0) \leftarrow \{x_j^H\}$, set $\mathcal{N} \leftarrow \mathcal{N} \backslash \mathcal{L}(0)$ and set $W_j \leftarrow \mathcal{L}(0)$.
4. For $k = 1$ to $2r+1$
   (a) Set layer $\mathcal{L}(i) \leftarrow \bigcup_{x \in \mathcal{L}(i-1)} (S_\varepsilon(x) \cap \mathcal{N})$ (i.e., all free nodes that are strongly connected to $\mathcal{L}(i-1)$).
   (b) If $i \leq r$, add to $\mathcal{L}(i)$ all $x \in \mathcal{N}$ that are strongly connected to at least d nodes in $\mathcal{L}(i)$, set $\mathcal{N} \leftarrow \mathcal{N} \backslash \mathcal{L}(i)$ and set $W_j \leftarrow W_j \cup \mathcal{L}(i)$.
5. Find $i_{max} \leftarrow \operatorname{argmax} |\mathcal{L}(i)|$ over the set $\{r+1, \ldots, 2r+1\}$ (i.e., the largest layer) and add to $\mathcal{S}$ all $x \in \mathcal{L}(i_{max})$ of shortest path length from x_j^H.
6. If $\mathcal{N} \neq \emptyset$, goto Step 2; else set $N_H \leftarrow j$.
7. Merge any aggregate W_j that is too small (i.e., $|W_j| < a_{min}$) with a strongly connected neighbouring aggregate W_k (subject to the requirement $|W_j \cup W_k| \leq a_{max}$; it may be necessary to split up W_j to achieve this) and set $N_H \leftarrow N_H - 1$. (If no such strongly connected neighbour exists, W_j is kept unchanged.)

Fig. 1. Aggregation algorithm

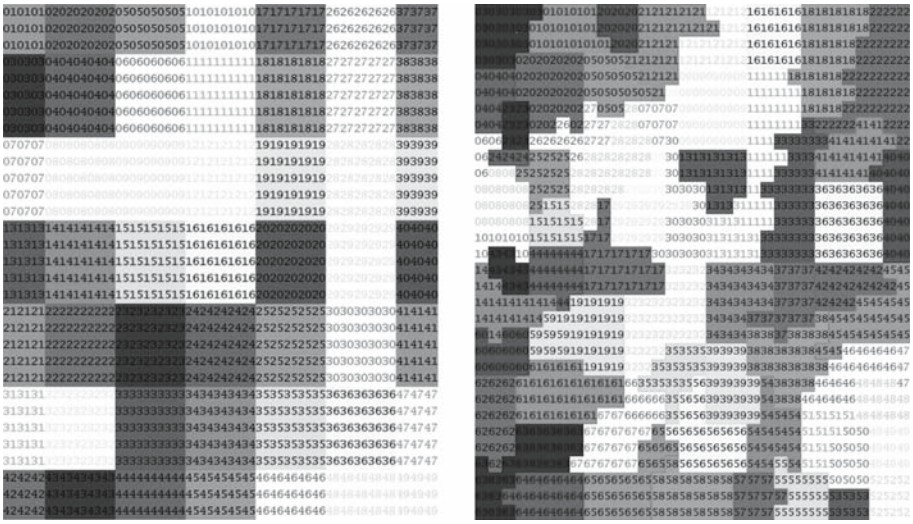

Fig. 2. Aggregates W_j for $\alpha \equiv 1$ (*left*) and α strongly-varying (*right*) ($r = 2$, $\varepsilon = \frac{2}{3}$)

where all connections in A are strong (see Fig. 2 (left) for the case $\alpha \equiv 1$ and $r = 2$ on a uniform mesh). For an arbitrary coefficient function α the shapes and sizes of the aggregates W_j may vary strongly (see Fig. 2 (right) for an example), and thus the constant in the shape-regularity assumption for $\{\omega_j\}$ in Sect. 2.1 may become large.

Following [6] further, we now define for each $j = 1, \ldots, N_H$ a vector $\boldsymbol{\Psi}^j \in \mathbb{R}^{n_h}$ as follows:

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

Let $\Psi_j \in S^h(\Omega)$ be the finite element function corresponding to the coefficient vector $\boldsymbol{\Psi}^j$. In the case of *non-smoothed aggregation* we set $\Phi_j = \Psi_j$ and we are finished.

If we want to use *smoothed aggregation*, we apply in addition a suitable smoother S to the functions Ψ_j, $j = 1, \ldots, N_H$. Several different choices for S have been investigated (see [31], [6], [24]). Here, we use the smoother proposed in [31], i.e.,

$$S := \left(I - \omega \, (\operatorname{diag} \mathcal{A})^{-1} \mathcal{A}^{\varepsilon} \right),$$

with damping parameter ω. (Note that this is similar to a damped Jacobi smoother.) We apply μ smoothing steps to each coefficient vector $\boldsymbol{\Psi}^j$, i.e.,

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

Then, the jth coarse space basis function $\Phi_j \in S^h(\Omega)$ is the finite element function corresponding to the (smoothed) coefficient vector $\boldsymbol{\Phi}^j$.

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 (C1) is satisfied.

Let us now consider whether the functions Φ_j, $j = 1, \ldots, N_H$, satisfy the other assumptions made in Sect. 3. Note that for quasi–uniform $\{\mathcal{T}^h\}$ and $\alpha \sim 1$ all connections in $\mathcal{A}$ are strong (provided ε is not too close to 1), and so this case has already been covered in the literature. See [6] and [24] for details. It is important to note however, that it has so far only been possible to prove the weaker Assumption (C4a), and not (C4b), in the case of smoothed basis functions (including the case of damped Jacobi smoothing used here). In the non-smoothed case, i.e., for $\mu = 0$, (C4a) and (C4b) follow 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 non-smoothed case here, i.e., $\mu = 0$ and so $\Phi_j = \Psi_j$. The case of smoothed aggregation will be covered in a forthcoming paper. In the case of $\mu = 0$, the size of the overlap parameter δ_j is of order $O(h_j)$ where h_j is the diameter of the largest element $\tau \subset \omega_j$ which touches the boundary of ω_j. All the assumptions made in Sect. 3 are satisfied by construction, except possibly the shape regularity of the supports ω_j. This is not guaranteed and depends on the coefficient function α.

However, 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 3.4 is bounded independently of α and the mesh parameters. Take for instance the following example of a binary medium α:

Example 4.2: 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 Fig. 3 for an

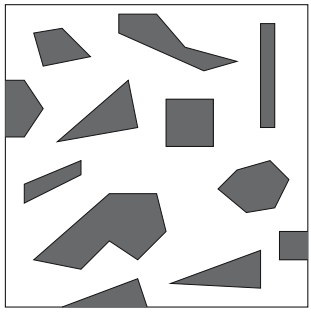

Fig. 3. Typical situation for Example 4.2, i.e., "islands" B_k (in *red*) where α is large

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

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

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 $k = 1, \dots, m$, or (ii) $S_\varepsilon(x_p) \cap B_k = \emptyset$ for all $k = 1, \dots, m$; 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, $j = 1, \dots, N_H$, constructed above satisfy either $W_j \subset B_k$ for some $k = 1, \dots, m$ or $W_j \cap B_k = \emptyset$ for all $k = 1, \dots, m$.

Provided we have a quasi–uniform family of meshes and h is small enough, our aggregation algorithm will therefore produce (non-smoothed) coarse space basis functions Φ_j with shape regular supports ω_j 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 $\nabla \Phi_j(x) \lesssim \delta_j^{-1}$ for all $x \in \Omega$, this implies that the coarse space robustness indicator

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

In the case of a quasi–uniform family of meshes, when $\delta_j \sim h$, and for suitably chosen r, this implies that the bound in Theorem 3.8 reduces to

$$\kappa\left(M_{\mathrm{AS}}^{-1} A\right) \lesssim \left(1 + \max_{j=1,\dots,N_H} \frac{H_j}{\delta_j}\right) \lesssim h^{-1} \max_{k=1,\dots,m} \{\operatorname{diam} B_k\}$$

independent of the size of $\hat{\alpha}$. Thus, if the maximum diameter of the islands (and therefore also $\max_j H_j$) is $O(h)$, the bound is completely independent of α.

5. Numerical results

In the following section we present some 2D experiments to verify the sharpness of our theoretical results in Sect. 3 in the case of aggregation coarse spaces. We only report results with smoothed aggregation in the case of $\alpha \equiv 1$. In the case of variable α we consider only non-smoothed aggregation. Let $\Omega = [0, 1]^2$ and let $\{\mathcal{T}^h\}$ be a family of uniform triangulations of Ω. (However, our implementation of the aggregation procedure does not require a structured triangulation.) We solve the resulting linear equation systems with preconditioned CG and use a tolerance of 10^{-6} for the relative reduction in the residual norm. For the solution of subdomain and coarse grid problems we employ the multifrontal sparse direct solver `UMFPACKv4.4` [11]. All CPU times given below were obtained on a 3 GHz Intel P4 processor with 1 GB RAM and include the setup times for each solver if not stated otherwise. See [30] for details regarding the implementation of our method.

Let us first consider the case $\alpha \equiv 1$, i.e., the Laplacian. In this case it is not crucial how to choose the threshold ε for strong connections, since all the off-diagonal entries in A are the same. We choose $\varepsilon = \frac{2}{3}$. As we have seen in the previous section (cf. Fig. 2), in this case 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 $H^{\text{sub}} = (2r_0+1)H + \delta^{\text{sub}}$, where r_0 is the coarse grid aggregation radius (i.e., aggregating the degrees of freedom in A_0).

In Fig. 4, we plot the condition numbers $\kappa\left(M_{\text{AS}}^{-1}A\right)$ of the preconditioned systems and the CPU times in the case of $n = 1024^2$ (and thus $h = 1/1025$) for various choices of r and r_0. For these experiments we choose the subdomain overlap $\delta^{\text{sub}} = 3h$ and use $\mu = 1$ smoothing steps with $\omega = \frac{2}{3}$. Hence $\delta = \delta^{\text{sub}} = 3h$.

We note that the agreement with the theory is extremely good: the condition number $\kappa\left(M_{\text{AS}}^{-1}A\right) \leq 5\frac{H}{\delta}$, 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 $H^{\text{sub}} \gg H$. We see that the best efficiency of the method is attained for H^{sub} somewhere between $10H$ and $30H$.

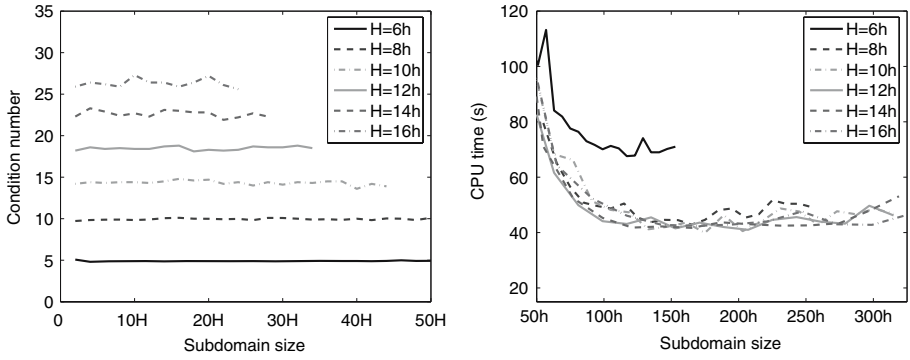

Fig. 4. Condition numbers and CPU times for Laplacian ($n = 1024^2$, $\delta^{\text{sub}} = 3h$, $\mu = 1$)

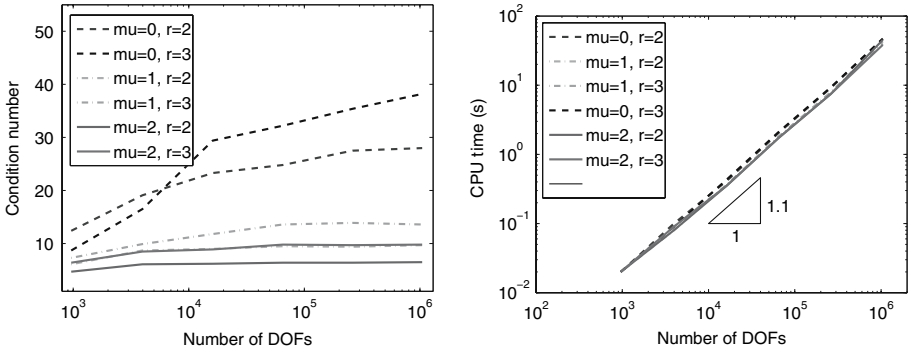

Fig. 5. Condition numbers and CPU times for Laplacian with $\delta^{\text{sub}} = 3h$ (varying n)

Table 1. CG-iterations and CPU times for Laplacian ($\delta^{\text{sub}} = 3h$ and $r = 3$)

n	No. of CG-iterations			Setup time (s)			Iteration time (s)		
	$\mu = 0$	$\mu = 1$	$\mu = 2$	$\mu = 0$	$\mu = 1$	$\mu = 2$	$\mu = 0$	$\mu = 1$	$\mu = 2$
64^2	21	16	14	0.05	0.05	0.05	0.04	0.03	0.03
128^2	25	18	16	0.20	0.21	0.21	0.22	0.15	0.14
256^2	27	18	16	1.01	1.02	1.02	1.10	0.78	0.71
512^2	26	18	15	4.53	4.66	4.56	4.79	3.23	2.89
1024^2	27	18	15	26.7	25.9	26.6	20.5	14.0	11.7

The method is also almost independent of h as we see in Fig. 5. Here, using the same 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 $\kappa\left(M_{\text{AS}}^{-1}A\right) \leq 5\frac{H}{\delta}$. (Note that the method does not break down in the case $\mu = 2$ where $\delta = 5h > \delta^{\text{sub}}$ and our theoretical assumption (C1) is violated. This has already been noted in [28].) The CPU times are growing approximately like $O(n^{1.1})$ which is almost optimal. The growth stems from the fact that (for fixed r and r_0) our coarse problem size and the subdomain problem sizes in the tests in Fig. 5 are growing proportionally to n. However, this increase in the CPU time is extremely mild, since `UMFPACK` scales very well (i.e., roughly like $O(n^{1.1})$) for reasonably small problems (up to 10^5 degrees of freedom). We are able to exploit this good performance of the sparse direct solver fully here. Note that for larger problems ($n \approx 10^6$ or bigger) `UMFPACK` and other sparse direct solvers start to slow down dramatically, e.g., for $n \approx 10^6$ `UMFPACK` scales only like $O(n^{1.9})$ on our system, while our method continues to scale like $O(n^{1.1})$. In Table 1, we give a more detailed break-down of the performance of our method in the case $r = 3$ and present number of iterations, setup time and time spent in CG. Note that for $r = 2$ (i.e., a larger coarse space) the number of iterations is consistently reduced by 2–3 iterations, but the total CPU time is almost identical due to slightly higher setup cost.

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

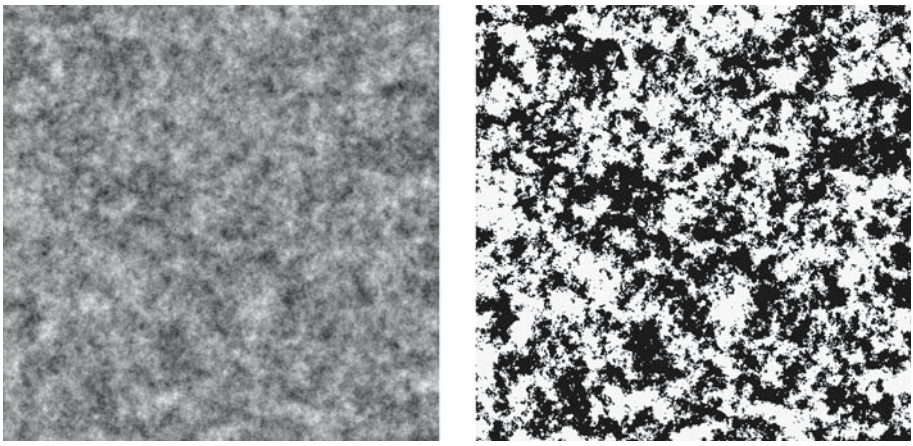

Fig. 6. Log–normal and "clipped" log-normal random fields ($n = 512^2$ and $\lambda = 1/64$)

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., [10]. We use `Gaussian` [23] to create these random fields. See Fig. 6 (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 large and small values of α. For example for the field in Fig. 6 for $\sigma^2 = 8$ we have $\max_{\tau,\tau' \in \mathcal{T}^h} \frac{\alpha_\tau}{\alpha_{\tau'}} \approx 3 \times 10^{10}$.

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 Fig. 6 (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 Fig. 6 for $\sigma^2 = 8$ we have $\max_{\tau,\tau' \in \mathcal{T}^h} \frac{\alpha_\tau}{\alpha_{\tau'}} \approx 5 \times 10^5$. Although the ratio 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 (see, e.g., [1]).

We will now test the robustness of our method (denoted `ADOUG` below) for these clipped log–normal fields. In all the tests below we choose $r = 2$, $\delta^{\text{sub}} = 3h$ and $\mu = 0$, i.e., no smoothing. The threshold for strong connections is chosen to be $\varepsilon = \frac{2}{3}$. See Fig. 2 (right) for a typical set of aggregates. Note that this value of ε was found by experiment. However, the method is not very sensitive to changes in ε. In fact, for a typical realisation of α, the number of iterations only varies by about 1 or 2 when varying ε in the interval [0.6, 0.8]. However, this does not mean the criterion is irrelevant. If the threshold is chosen too large (i.e., $\varepsilon = 0.9$ or larger), then the method is not robust anymore.

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

CG-iterations					CPU-time (s)			
σ^2	$\max_{\tau,\tau'} \frac{\alpha_\tau}{\alpha_{\tau'}}$	ADOUG	AAMG	DOUG	σ^2	ADOUG	AAMG	UMFPACK
2	1.5×10^1	24	14	32	2	2.12	1.35	1.85
4	2.2×10^2	27	27	89	4	2.14	2.27	1.70
6	3.3×10^3	29	40	296	6	2.34	3.31	1.33
8	4.9×10^4	26	77	498	8	2.41	6.23	4.88
10	7.4×10^5	26	27	724	10	2.37	2.39	4.98

In our first set of results in Table 2, we fix the problem size to $n = 256^2$ and the correlation length to $\lambda = 1/64$ and vary σ^2 (and thus the size of the jump in the coefficient function α) for a typical realisation of the clipped log-normal random field. We see that, as suggested by the above theory, 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 significantly with σ^2. Note that the size of the coarse space does also not vary significantly, i.e., $N \approx 5000$ in all cases.

To put this performance in perspective we compare ADOUG with an aggregation-type AMG preconditioner by Bastian [2], [3], [4] (denoted AAMG below) that uses a similar aggregation algorithm (without smoothing) and the same criterion for strong connections as our code, i.e., (24) with $\varepsilon = \frac{2}{3}$. The average coarsening rate in AAMG is about 4 and the setup time for the experiments in Table 2 is about 0.35 s. We observe from the results in Table 2 (columns 4 and 7) that AAMG is not robust to variations in σ^2 for clipped random fields. Although it is clearly superior to ADOUG both in terms of the number of iterations and the CPU-time for small variations in α, the number of CG-iterations is strongly dependent on the size of the jump in α, and in the case of $\sigma^2 = 8$ AAMG is more than 2.5 times slower than ADOUG. Since most components in AAMG are similar to ADOUG this leads to the conjecture that the loss of robustness is related to the smoother in AAMG which is standard Gauss–Seidel. Algebraic multigrid methods rely on the fact that the coarse grid correction and the smoother are "compatible" and this might not be the case here. However the theoretical analysis of this is still an open problem. The performance of AAMG might also be improved, if smoothed aggregation (such as employed in the ML library [12]) was used. However, our experiments with ADOUG suggest that in the case of clipped log–normal random fields the difference between smoothed and non-smoothed aggregation is very small.

We also include in Table 2 (column 5) the iteration count for a "classical" two-level additive Schwarz preconditioner DOUG [14], [15] with standard piecewise linear FE coarse space. This code uses graph partitioning software to construct the subdomains, but as predicted by the theory in [16], [18] applied to standard piecewise linear coarsening, DOUG is not robust when the coefficient jumps are not resolved by the coarse mesh and the number of iterations grows very strongly with the size of the jump. Finally, we also include in Table 2 (column 9) the CPU times for solving each of the problems directly with the sparse direct solver UMFPACKv4.4 [11]. We see that even UMFPACK shows a dependency on the size of the jump. This is due

Table 3. Comparison of solvers for clipped random fields with $n = 256^2$ and $\sigma^2 = 8$

CG-iterations				CPU-time (s)			
λ	ADOUG	AAMG	DOUG	λ	ADOUG	AAMG	UMFPACK
1/17	26	18	355	1/17	2.20	1.67	4.52
1/33	27	64	430	1/33	2.24	5.14	4.77
1/65	26	77	498	1/65	2.41	6.23	4.88
1/129	33	70	655	1/129	2.71	5.77	7.48
1/257	48	166	858	1/257	3.84	13.5	10.2

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

CG-iterations					CPU-time (s)			
h	λ	ADOUG	AAMG	DOUG	h	ADOUG	AAMG	UMFPACK
1/65	1/17	20	12	60	1/65	0.10	0.06	0.05
1/129	1/33	25	35	136	1/129	0.46	0.68	0.52
1/257	1/65	26	77	498	1/257	2.41	6.23	4.88
1/513	1/129	34	100	1111	1/513	16.8	33.8	88.8
1/1025	1/257	74	422	–	1/1025	105.9	540	–

to the extra cost for partial pivoting in the case of largely varying entries in the matrix A.

In Table 3, 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 $\lambda > 2h$. 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. AAMG shows a much stronger dependency on λ and seems to have real problems with short correlation lengths. DOUG and UMFPACK are also affected, but to a lesser extent than AAMG.

In Table 4 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 since the correlation length is linked to the mesh width, the problem actually gets harder the more the mesh is refined (in addition to the growing problem size n). Note that the size of the coarse space is about 50,000 in our largest problem, i.e., for $h = 1/1025$. 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 AAMG. 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 Fig. 5. UMFPACK starts to slow down quite dramatically at $h = 1/512$ and runs out of memory for $h = 1/1024$. The classical additive Schwarz method DOUG with piecewise linear coarse space even diverges for $h = 1/1024$. It would be interesting to investigate whether using smoothed aggregation would improve the results of AAMG in Table 4. However, again our experience with smoothed aggregation within ADOUG suggests that this would lead to no significant improvement over non-smoothed aggregation.

To substantiate our claim that unclipped log–normal random fields are indeed simpler, we include one set of results with an unclipped log–normal random field

Table 5. Comparison for an unclipped log-normal field with $n = 512^2$, $\sigma^2 = 8$ and $\lambda = 8h$

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

in Table 5. Note that although we have $\max_{\tau,\tau' \in \mathcal{T}^h} \frac{\alpha_\tau}{\alpha_{\tau'}} \approx 3 \times 10^{10}$ for the random field in Table 5, the condition number of A is in fact nowhere near as large in this case, since the correlation length of $\lambda = 8h$ ensures that values of α are correlated on two neighbouring elements $\tau, \tau' \in \mathcal{T}^h$. Since only the local ratio of $\frac{\alpha_\tau}{\alpha_{\tau'}}$ plays a rôle in the analysis of all the methods, they all perform better for unclipped fields than for clipped fields, in particular the classical additive Schwarz method DOUG (see Table 5). The aggregation-based Schwarz method ADOUG is about 1.5 times faster than the aggregation-based AMG code AAMG and slightly faster than UMFPACK. However, UMFPACK fails to run for larger problem (e.g., for $n = 1024^2$) due to memory requirements.

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 and Dr. Jan Van lent for many useful discussions regarding the theory and the results.

References

[1] Almond, D. P., Bowen, C. R., Rees, D. A. S.: Composite dielectrics and conductors: simulation, characterization and design. J Phys D Appl Phys **39**, 1295–1304 (2006)
[2] Bastian, P.: Aggregation-based algebraic multigrid code. IWR, University of Heidelberg (1996) (available at http://www.dune-project.org/index.html)
[3] Bastian, P., Helmig, R.: Efficient fully-coupled solution techniques for two-phase flow in porous media: parallel multigrid solution and large scale computations. Adv Water Res **23**, 199–216 (1999)
[4] Blatt, M., Bastian, P.: The iterative solver template library. In: Proc. Workshop on State-of-the-Art in Scientific and Parallel Computing, PARA '06, Umea, June 18–21, 2006. Lecture Notes in Scientific Computing. Springer, Berlin (forthcoming)
[5] Brenner, S. C., Scott, L. R.: The mathematical theory of finite element methods. Springer, Berlin (1994)
[6] Brezina, M., Vanek, P.: A black-box iterative solver based on a two-level Schwarz method. Computing **63**, 233–263 (1999)
[7] Cai, X., Nielsen, B. F., Tveito, A.: An analysis of a preconditioner for the discretized pressure equation arising in reservoir simulation. IMA J Numer Anal **19**, 291–316 (1999)
[8] Carvalho, L. M., Giraud, L., Le Tallec, P.: Algebraic two-level preconditioners for the Schur complement method. SIAM J Sci Comput **22**, 1987–2005 (2001)
[9] Chan, T. F., Mathew, T.: Domain decomposition methods. Acta Numerica 1994. Cambridge University Press (1994)
[10] Cliffe, K. A., Graham, I. G., Scheichl, R., Stals, L.: Parallel computation of flow in heterogeneous media modelled by mixed finite elements. J Comput Phys **164**, 258–282 (2000)
[11] Davies, T. A.: UMFPACKv4.4 multifrontal software, CISE, University of Florida (2005) (available at http://www.cise.ufl.edu/research/sparse/umfpack)
[12] Gee, M. W., Siefert, C. M., Hu, J. J., Tuminaro, R. S., Sala, M. G.: ML 5.0 Smoothed Aggregation User's Guide, Technical Report SAND2006-2649, Sandia National Laboratories (2006)

[13] Giraud, L., Guevara Vasquez, F., Tuminaro, R. S.: Grid transfer operators for highly-variable coefficient problems in two-level non-overlapping domain decomposition methods. Numer Linear Algebra Appl **10**, 467–484 (2003)
[14] Graham, I. G., Hagger, M. J.: Additive Schwarz, CG and discontinuous coefficients. In: (Bjorstad, P., Espedal, M., Keyes, D. E., eds.) Proc. 9th Int. Conf. on Domain Decomposition Methods, Bergen, 1996. Domain Decomposition Press, Bergen (1998)
[15] Graham, I. G., Hagger, M. J.: Unstructured additive Schwarz–CG method for elliptic problems with highly discontinuous coefficients. SIAM J Sci Comput **20**, 2041–2066 (1999)
[16] Graham, I. G., Lechner, P. O., Scheichl, R. O.: Domain decomposition for multiscale PDEs. Numer Math **106**, 589–626 (2007)
[17] Graham, I. G., Scheichl, R.: Robust domain decomposition algorithms for multiscale PDEs. Numer Meth Partial Differ Equ **23**, 859–878 (2007)
[18] Graham, I. G., Scheichl, R.: Coefficient-explicit condition number bounds for overlapping additive Schwarz. In: Proc. 17th Int. Conf. on Domain Decomposition Methods, Strobl, Austria, 3–7 July 2006. Lecture Notes in Computational Science and Engineering. Springer, Berlin (2007) (forthcoming)
[19] Griebel, M., Scherer, K., Schweitzer, M. A.: Robust norm equivalences and optimal preconditioners for diffusion problems. Math Comput **76**, 1141–1161 (2007)
[20] Henson, V. E., Yang, U. M.: BoomerAMG: a parallel algebraic multigrid solver and preconditioner. Appl Numer Math **41**, 155–177 (2002)
[21] Jenkins, E. W., Kelley, C. T., Miller, C. T., Kees, C. E.: An aggregation-based domain decomposition preconditioner for groundwater flow. SIAM J Sci Comput **23**, 430–441 (2001)
[22] Jones, J. E., Vassilevski, P. S.: AMGe-based on element agglomeration. SIAM J Sci Comput **23**, 109–133 (2001)
[23] Kozintsev, B., Kedem, B.: Gaussian package, University of Maryland (1999) (available at http://www.math.umd.edu/ bnk/bak/generate.cgi)
[24] Lasser, C., Toselli, A.: Convergence of some two-level overlapping domain decomposition preconditioners with smoothed aggregation coarse spaces. In: (Pavarino, L., Toselli, A., eds.) Recent developments in domain decomposition methods, pp. 95–117. Lecture Notes in Computational Science and Engineering, vol. 23. Springer, Berlin (2002)
[25] Raw, M.: A coupled algebraic multigrid method for the 3D Navier–Stokes equations. In: Proc. 34th Aerospace Sciences Meeting and Exhibit, Reno, AIAA-1996-297. AIAA Press, Washington (1996)
[26] Ruge, J., Stüben, K.: Efficient solution of finite difference and finite element equations by algebraic multigrid (AMG). In: (Paddon, D. J., Holstein, H., eds.) Multigrid methods for integral and differential equations, pp. 169–212. IMA Conference Series. Clarendon, Oxford (1985)
[27] Sala, M.: Analysis of two-level domain decomposition preconditioners based on aggregation. ESAIM M2AN **38**, 765–780 (2004)
[28] Sala, M., Shadid, J. N., Tuminaro, R. S.: An improved convergence bound for aggregation-based domain decomposition preconditioners. SIAM J Matrix Anal **27**, 744–756 (2006)
[29] Toselli, A., Widlund, O.: Domain decomposition methods – algorithms and theory. Springer, Berlin (2005)
[30] Vainikko, E., Tehver, M., Batrashev, O., Niitsoo, M., Skaburskas, K., Scheichl, R.: ADOUG – aggregation-based domain decomposition on unstructured grids. Universities of Tartu and Bath (2006) (available at http://www.dougdevel.org)
[31] Vanek, P., Mandel, J., Brezina, M.: Algebraic multigrid by smoothed aggregation for 2nd and 4th order elliptic problems. Computing **56**, 179–196 (1996)
[32] Vanek, P., Brezina, M., Mandel, J.: Convergence of algebraic multigrid based on smoothed aggregation. Numer Math **88**, 559–579 (2001)
[33] Vuik, C., Segal, A., Meijerink, J. A.: An efficient preconditioned CG method for the solution of a class of layered problems with extreme contrasts in the coefficients. J Comput Phys **152**, 385–403 (1999)
[34] Wan, W. L., Chan, T. F., Smith, B.: An energy-minimizing interpolation for robust multigrid methods. SIAM J Sci Comput **21**, 1632–1649 (2000)