Robust Domain Decomposition Algorithms for Multiscale PDEs

I.G. Graham and R. Scheichl

Bath Institute For Complex Systems

Preprint 14/06 (2006)

http://www.bath.ac.uk/math-sci/BICS

Robust Domain Decomposition Algorithms for Multiscale PDEs

I.G.Graham and R. Scheichl
Department of Mathematical Sciences ,
University of Bath ,
Bath BA2 7AY , U.K.
email: I.G.Graham@bath.ac.uk , R.Scheichl@bath.ac.uk

In this paper we describe a new class of domain deomposition preconditioners suitable for solving elliptic PDEs in highly fractured or heterogeneous media, such as arise in groundwater flow or oil recovery applications. Our methods employ novel coarsening operators which are adapted to the heterogeneity of the media. In contrast to standard methods (based on piecewise polynomial coarsening), the new methods can achieve robustness with respect to coefficient discontinuities even when these are not resolved by a coarse mesh. This situation arises often in practical flow computation, in both the deterministic and (Monte-Carlo simulated) stochastic cases. An example of a suitable coarsener is provided by multiscale finite elements. In this paper we explore the linear algebraic aspects of the multiscale algorithm, showing that it involves a blend of both classical overlapping Schwarz methods and non-overlapping Schur methods. We also extend the algorithm and the theory from its additive variant to obtain new hybrid and deflation variants. Finally we give extensive numerical experiments on a range of heterogeneous media problems illustrating the properties of the methods.

I. INTRODUCTION

In this paper we discuss new domain decomposition preconditioners for piecewise linear finite element discretisations of boundary-value problems for the model elliptic problem

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

in a bounded polygonal or polyhedral domain $\Omega \subset \mathbb{R}^d$, $d = 2$ or 3 with suitable boundary conditions on the boundary $\partial\Omega$. The coefficient $\alpha(x)$ may vary over many orders of magnitude in an unstructured way on Ω. Many examples of this kind arise in groundwater flow and oil reservoir modelling.

Let $\mathcal{T}^h$ be a conforming shape-regular simplicial mesh on Ω and let $\mathcal{V}^h$ denote the space of continuous piecewise linear finite elements on $\mathcal{T}^h$ which vanish on essential boundaries. The finite element discretisation of (I.1) in this space yields the linear system:

$$A\mathbf{u} = \mathbf{f} \;, \tag{I.2}$$

and it is well-known that the conditioning of A worsens when $\mathcal{T}^h$ is refined or when the heterogeneity (characterised by the range of α) becomes large. It is of interest to find solvers for (I.2) which are robust to changes in the mesh width h as well as in the heterogeneity.

While there are many papers which solve this problem for "layered media" in which discontinuities in α are simple interfaces which can be resolved by a coarse mesh (see e.g. [4] and the references therein), until recently there was no rigorously justified method for general heterogeneous media. The recent paper [9] presented a new analysis of domain decomposition methods for (I.2) (which have inherent robustness with respect to h) and this analysis indicated explicitly how subdomains and coarse spaces should be designed in order to achieve robustness also with respect to heterogeneities. More precisely this analysis introduced new "robustness indicators" (which depend on the choice of subdomains and coarse space and in particular depend on the energy of the coarse space basis functions) and proved that, if these indicators are controlled, then the preconditioner will be robust. Paper [9] then goes on to consider the use of multiscale finite elements to build coarse spaces for domain decomposition and proves a number of results which indicate their robustness in cases where standard coarsening methods fail to be robust. Crucially our methods do not require resolution of the coefficient jumps by the coarse mesh to achieve robustness.

The coarse spaces proposed in [9] yield coefficient-dependent prolongation operators, similar to those which have been tested empirically in the context of (Schur complement based) domain decomposition methods in [3, 6]. The concept of energy-minimising coarse spaces also appears in several papers on the construction of algebraic multigrid methods. For example [23] proposes to compute low energy coarse spaces by solving a global constrained minimisation problem, before attempting the solution of (I.2). This is in principle expensive, but an approximate solution (based on a few PCG iterations to the minimization problem) yields some good emprirical results. In [11] a recursive algorithm to approximately solve the constrained minimisation problem of [23] is proposed, but the behaviour of this in the presence of heterogeneity is not analysed. The use of multiscale finite elements as coarseners was also proposed in [1], but again this was in the Schur-complement context and the analysis depended on classical periodic homogenisation theory. The analysis in our recent paper [9] does not require periodicity and does not appeal to homogenisation theory.

In [9] we presented rigorous error bounds for generalised domain decomposition methods and proved that they can work well in the presence of heterogeneity. The analysis in [9] is presented in terms of the abstract theory of Schwarz methods. However that analysis is not so revealing about the linear algebra aspects, which determine the computational work involved. Thus in this paper we give a linear algebraic presentation of the multiscale preconditioners. In particular we show that the multiscale preconditioner, in its simplest form, is a variant of the two-level overlapping additive Schwarz method with a non-standard coarse solve. The coarse solve proposed in [9] could also be obtained by (i) reduction of A to a Schur complement defined on the "skeleton nodes" (i.e. on all fine freedoms that lie on coarse element boundaries) and (ii) subsequent further coarsening (via interpolation) to obtain a matrix defined on coarse freedoms only. In this sense the method combines aspects of both the classical overlapping and the non-overlapping (Schur complement or iterative substructuring) methods.

We also go on to consider hybrid and deflation variants of our preconditioner, where the coarse solve is either combined in a multiplicative way with the subdomain solves

(see [13, 14, 21]), or used to deflate the system (see [16, 17]). We show that with an appropriate initial guess both these preconditioners lead to the same preconditioned conjugate gradient iteration and perform no worse than the additive variant. Finally we present a range of numerical experiments illustrating the results.

To describe our preconditioners, we restrict to the case of homogeneous Dirichlet boundary conditions for simplicity. For any subdomain D of Ω which consists of a union of elements of $\mathcal{T}^h$, let $\mathcal{S}^h(D)$, denote the space of continuous piecewise linear functions on D and let $\mathcal{S}^h_0(D) = \mathcal{S}^h(D) \cap H^1_0(D)$. In particular, $\mathcal{V}^h = \mathcal{S}^h_0(\Omega)$. Moreover, if W is any subset of $\overline{\Omega}$, then $\mathcal{N}^h(W) := \{x_i \ : i \in \mathcal{I}^h(W)\}$ is the set of all nodes of the mesh $\mathcal{T}^h$ which lie in W (where $\mathcal{I}^h(W)$ is a suitable index set).

Domain decomposition preconditioners for (I.2) are defined by first intoducing an overlapping open covering $\{\Omega_i : i = 1, \ldots, N\}$ of Ω, with each $\overline{\Omega}_i$ assumed to consist of a union of elements from $\mathcal{T}^h$. Then we introduce $\mathcal{V}_i := \{v_h \in \mathcal{V}^h : \ \mathrm{supp}(v_h) \subset \overline{\Omega}_i\}$ and, for $j \in \mathcal{I}^h(\Omega_i)$ and $j' \in \mathcal{I}^h(\Omega)$, we define the restriction matrix $(R_i)_{j,j'} := \delta_{j,j'}$. The matrix $A_i := R_i A R_i^T$, is then just the minor of A corresponding to rows and columns taken from $\mathcal{I}^h(\Omega_i)$. One-level domain decomposition methods are constructed from the inverses A_i^{-1}.

Two-level methods are obtained by adding a global coarse solve. Let $\mathcal{T}^H$ be a shape-regular mesh of coarse simplices on Ω with a typical element of $\mathcal{T}^H$ being the (closed) set K, which again we assume to consist of the union of a set of fine grid elements $\tau \in \mathcal{T}^h$. Also, let $\mathcal{F}^H$ denote the set of all (closed) faces of elements in $\mathcal{T}^H$. (In the 2D case "faces" should be interpreted to mean "edges".) The set of coarse mesh nodes on any subset W of $\overline{\Omega}$ is denoted by $\mathcal{N}^H(W) := \{x^H_p : p \in \mathcal{I}^H(W)\}$ with a suitably chosen index set $\mathcal{I}^H(W)$. The coarse space basis functions Φ_p are required to satisfy (for $p, p' \in \mathcal{I}^H(\overline{\Omega})$) the assumptions:

(C1) $\Phi_p \in S^h(\Omega) \ , \quad \Phi_p(x^H_{p'}) \ = \ \delta_{p,p'} \ ;$

(C2) $\mathrm{supp}(\Phi_p) \ \subset \ \omega_p \ := \ \cup\{K : p \in \mathcal{I}^H(K)\} \ ;$

(C3) $\sum_{p \in \mathcal{I}^H(\overline{\Omega})} \Phi_p(x) \ = \ 1, \quad x \in \overline{\Omega};$

(C4) $\|\Phi_p\|_{L_\infty(\Omega)} \ \lesssim \ 1 \ .$

From these functions we define the coarse space $\mathcal{V}_0 \ := \ \mathrm{span}\{\Phi_p \ : \ p \in \mathcal{I}^H(\Omega)\}$, which, by (C1) and (C2), is the span of all Φ_p that vanish on the boundary $\partial\Omega$, and is thus a subspace of $\mathcal{V}^h$. (A special case of basis functions satisfying (C1)–(C4) is given by the standard continuous piecewise linear functions on the coarse mesh $\mathcal{T}^H$.) Now, if we introduce the restriction matrix

$$(R_0)_{pj} \ := \ \Phi_p(x^h_j) \ , \quad j \in \mathcal{I}^h(\Omega), \quad p \in \mathcal{I}^H(\Omega), \tag{I.3}$$

then the matrix $A_0 := R_0 A R_0^T$ is the stiffness matrix for problem (I.1) discretised in $\mathcal{V}_0$ using the basis $\{\Phi_p : p \in \mathcal{I}^H(\Omega)\}$.

Various two-level domain decomposition methods can be defined from these ingredients. Our first methods are the classical one- and two-level Additive Schwarz methods, $M^{-1}_{AS,1}$ and $M^{-1}_{AS,2}$, given by

$$M^{-1}_{AS,1} \ = \ \sum_{i=1}^{N} R_i A_i^{-1} R_i^T \quad \text{and} \quad M^{-1}_{AS,2} \ = \ R_0^T A_0^{-1} R_0 \ + \ M^{-1}_{AS,1} \ . \tag{I.4}$$

In the following section we give a short summary of some theoretical results from [9] which describe the behaviour of the preconditioners in (I.4). Before proceeding, we remark that the results in [9] have also been extended to coarse spaces constructed by smoothed aggregation in [19, 20] .

II. SUMMARY OF THE THEORY

For the purposes of exposition we will describe the theory for scalar α in (I.1) (and all our computations will be for that case), but we remark at the end of this section how the theory extends to the tensor case. Throughout the paper, the notation $C \lesssim D$ (for two quantities C, D) means that C/D is bounded above independently, not only of the *mesh parameter* h, the *domain decomposition parameters* δ_i, ρ_i and H_K (introduced below), but also of the *average values* of α on each $\tau \in \mathcal{T}^h$. Moreover $C \sim D$ means that $C \lesssim D$ and $D \lesssim C$. For theoretical purposes, we shall assume that $\alpha \geq 1$. This is no loss of generality, since problem (I.2) can be scaled by $(\min_x \alpha(x))^{-1}$ without changing its conditioning.

The theory in [9] contains a number of technical assumptions which allow quite general unstructured overlapping subdomains and also allow coarse meshes which are unrelated to the subdomains. Here we give a descriptive treatment of these assumptions and refer the reader to [9] for full details.

First we need some assumptions on the subdomains $\{\Omega_i\}$. For each i, we require that there is an **overlap parameter** $\delta_i > 0$, such that the part of Ω_i which is overlapped by the others is uniformly of order $\mathcal{O}(\delta_i)$. We can state this more precisely by introducing for $\mu > 0$, the "near boundary subset":

$$\Omega_{i,\mu} := \{x \in \Omega_i \ : \mathrm{dist}(x, \Gamma_i) < \mu\} \ , \tag{II.1}$$

where Γ_i is the boundary of Ω_i. The **overlap assumption** then requires that the part of Ω_i which is overlapped by other subdomains should be contained in Ω_{i,δ_i} and moreover, should also contain $\Omega_{i,c\delta_i}$, for some absolute constant $0 < c < 1$.

We also allow the subdomains Ω_i to have quite general shapes, but their "thinness" places a constraint on the design of the coarse mesh (see below). For each i, we define the **shape parameter** ρ_i as the largest μ so that the near-boundary set $\Omega_{i,\mu}$ can be finitely covered by Lipschitz polyhedra, each of which has: (i) closure intersecting Γ_i in a set of measure $\sim \mu^{d-1}$; (ii) diameter $\sim \mu$; (iii) length of edges $\sim \mu$; and (iv) volume $\sim \mu^d$. Since the part of Ω_i which is overlapped by other subdomains consists of a union of (shape–regular) elements from $\mathcal{T}^h$, the overlap assumption implies that $\rho_i \gtrsim \delta_i > 0$.

To see how the shape parameter measures the "thinness" of a subdomain, consider in $3D$, either a rectangular slab-shaped hexahedron Ω_1 with dimensions $a \times a \times b$, or a rectangular rod-shaped hexahedron Ω_2 with dimensions $a \times b \times b$, where, in both cases, $b \ll a$. Then clearly both these subdomains have shape parameter $\rho_i \sim b$, $i = 1, 2$.

Having introduced the subdomains Ω_i which cover Ω, now let $\Pi(\{\Omega_i\})$ denote the set of all partitions of unity $\{\chi_i\} \subset W^1_\infty(\Omega)$ subordinate to this cover. (This is a standard concept - see, e.g., [21].) Then we define our first robustness indicator which will appear in our condition number estimates below:

Definition. (Partitioning robustness indicator). *For a particular partition of unity* $\{\chi_i\}$*, define*

$$\pi(\alpha, \{\chi_i\}) = \max_{i=1}^{N} \left\{ \delta_i^2 \, \left\| \alpha |\nabla \chi_i|^2 \right\|_{L_\infty(\Omega)} \right\} .$$

Then the partition robustness indicator is defined by

$$\pi(\alpha) \;=\; \inf_{\{\chi_i\} \in \Pi(\{\Omega_i\})} \pi(\alpha, \{\chi_i\}) \; .$$

The quantity $\pi(\alpha)$ appears in our estimates for the one–level and the two–level preconditioner below. Note that, roughly speaking, this robustness indicator is well-behaved if there is a partition of unity whose functions have small gradient wherever α is large. So in a certain sense $\pi(\alpha)$ measures the ability of the subdomains Ω_i to handle the coefficient heterogeneity. The weight δ_i^2 is chosen to make $\pi(\alpha) \lesssim 1$ when $\alpha = 1$.

Turning now to the two-level variant, recall that the coarse space basis functions $\{\Phi_p\}$ are required to satisfy assumptions (C1)–(C4) on a coarse simplicial mesh $\mathcal{T}^H$. The coarse mesh will be assumed shape regular with the diameter of a typical element K denoted by H_K. A mild connection is needed between the local coarse and subdomain sizes which is described by introducing the concept of the **local coarse mesh diameter** defined by:

$$H_i \;:=\; \max_{K \in \mathcal{T}^H(\Omega_i)} H_K \; , \quad \text{where} \quad \mathcal{T}^H(\Omega_i) \;:=\; \{K \in \mathcal{T}^H : K \cap \overline{\Omega}_i \neq \emptyset\} \; . \tag{II.2}$$

Our final assumption is then:

(C5) $$H_i \;\lesssim\; \rho_i, \quad i = 1, \ldots, N.$$

This says that a coarse mesh element should not be large in comparison to the shape parameters of the subdomains which it intersects and is a generalisation of [21, Assumption 3.5]. Note that (C5) does not impose any direct structural relation between coarse mesh and subdomains.

Analogous to $\pi(\alpha)$ we now introduce a quantity which reflects how the coarse space handles the coefficient heterogeneity.

Definition. (Coarse space robustness indicator).

$$\gamma(\alpha) \;:=\; \max_{p \in \mathcal{I}^H(\Omega)} \left\{ H_p^{2-d} \, |\Phi_p|^2_{H^1(\Omega),\alpha} \right\} \qquad \textit{where} \qquad H_p \;:=\; \operatorname{diam}(\omega_p) \; .$$

Note also that for the classical case when Φ_p are the nodal basis for the continuous piecewise linear functions on $\mathcal{T}^H$, we have, via standard estimates, $\gamma(\alpha) \;\lesssim\; \max_{\tau \in \mathcal{T}^h} \alpha_\tau$, and so $\gamma(\alpha) \lesssim 1$ when $\alpha \sim 1$. When α varies more rapidly, our framework leaves open the possibility of choosing the Φ_p to depend on α in such a way that $\gamma(\alpha)$ is still well-behaved.

We can now state one of the main results from [9], giving condition number estimates for the one- and two- level additive Schwarz preconditioners.

Theorem 2.1. *The condition numbers of the preconditioned stiffness matrices using the one-level and the two-level additive Schwarz preconditioners satisfy*

$$\kappa \left(M_{AS,1}^{-1} A \right) \;\lesssim\; \pi(\alpha) \max_{i=1}^{N} \left\{ \frac{1}{\rho_i \delta_i} \right\} \; ,$$

$$\kappa \left(M_{AS,2}^{-1} A \right) \;\lesssim\; \pi(\alpha)\, \gamma(1) \max_{i=1}^{N} \left(1 + \frac{H_i}{\delta_i} \right) \;+\; \gamma(\alpha) \; .$$

It is important to emphasise that the quantity H_i introduced in (II.2) is **not** the subdomain diameter. Since it is a measure of the size of coarse mesh elements which intersect Ω_i, it can be much smaller than the diameter of Ω_i.

We note that Theorem 2.1 is a very general result in which the robustness indicators $\pi(\alpha)$ and $\gamma(\alpha)$ both play a rôle. However the following example shows that, in a certain sense the choice of the coarse basis functions Φ_p are the primary issue in determining robustness of the two-level method.

Example. As above, choose a coarse grid $\mathcal{T}^H$ and basis functions $\{\Phi_p : p \in \mathcal{I}^H(\overline{\Omega})\}$. (Note that we include basis functions corresponding to nodes on the boundary $\partial\Omega$.) Index the subdomains by $p \in \mathcal{I}^H(\overline{\Omega})$ and choose the pth subdomain Ω_p to be ω_p (see (C2)). It follows that the overlap parameter can be chosen as $\delta_p \sim H_p$. Now, the Φ_p form a partition of unity subordinate to the covering $\{\Omega_p : p \in \mathcal{I}^H(\overline{\Omega})\}$, and so by applying a trivial bound to $\gamma(\alpha)$, we have

$$\max\{\pi(\alpha), \gamma(\alpha)\} \;\leq\; \pi\left(\alpha, \{\Phi_p\}\right) \;=\; \max_{p \in \mathcal{I}^H(\overline{\Omega})} \left\{\delta_p^2 \|\alpha |\nabla \Phi_p|^2\|_{L_\infty(\Omega)}\right\} \;.$$

Hence Theorem 2.1 implies

$$\kappa\left(M_{AS,2}^{-1} A\right) \;\lesssim\; \gamma(1)\, \pi\left(\alpha, \{\Phi_p\}\right)$$

and so robustness with respect to α is achieved in this case simply by ensuring Φ_p has a small gradient wherever α is large.

We remark that Theorem 2.1 is a genuine extension of existing results on this topic. The best general result (for linear coarsening, and written for the case $\alpha \geq 1$ which we consider here) is:

$$\kappa(M_{AS,2}^{-1} A) \;\lesssim\; \left[\max_p \sup_{x,y \in \omega_p} \alpha(x)\right] \left(1 + \frac{H^{\rm sub}}{\delta}\right) \;, \tag{II.3}$$

where $H^{\rm sub}$ is the maximum subdomain diameter and δ is the global minimum of all the overlap parameters δ_i. This guarantees "robustness" with respect to α only when the coarse mesh resolves α sufficiently well. Related results (but not special cases of (II.3)) in fact show robustness with respect to large jumps in α, provided these jumps are resolved by the coarse mesh (see, for example [4] and many references therein). On the other hand if we consider a "binary medium" of two materials, characterised by $\alpha_1 = 1$ and $\alpha_2 = \hat{\alpha} \to \infty$, and we put some of each material into at least one element of the coarse mesh, then (II.3) allows the condition number to grow with $O(\hat{\alpha})$ as $\hat{\alpha} \to \infty$. In contrast the estimates in Theorem 2.1 show how the coarse mesh (and subdomains) should be designed to restore robustness in cases where (II.3) fails to ensure it.

Our methods produce well-conditioned matrices even in the presence of a large number of unresolved interfaces and arbitrary contrast of material properties (see §V.) and so are suitable for random media. We remark that there are also a number of results on the spectral clustering properties of preconditioners in the case of relatively few unresolved interfaces which motivate deflation methods [7, 8, 2, 22]. We return to these in §IV.

Finally we remark that all the theory described above is extended in [9] to the case where α in (I.1) is replaced by an isotropic and symmetric positive definite tensor $\mathcal{A}$ with the property that

$$\underline{C}\; \alpha(x)\; |\xi|^2 \;\leq\; \xi^T \mathcal{A}(x) \xi \;\leq\; \overline{C}\; \alpha(x)\; |\xi|^2 \;, \quad \xi \in \mathbb{R}^d \;,$$

where $\underline{C}, \overline{C}$ are constants and all the above estimates remain true.

III. MULTISCALE COARSENING

In this section we describe the particular case when the coarse basis functions Φ_p are obtained using "multiscale finite elements" ([10]). In this case the Φ_p are obtained by extending predetermined boundary data into the interior of each element K by discrete harmonic extension with respect to the original elliptic operator (I.1). Our exposition here is slightly different to that in [9], but the method is the same.

Recalling the definition of $\mathcal{F}^H$, we can introduce the **skeleton**:

$$\Gamma := \cup\{f : f \in \mathcal{F}^H\} ,$$

i.e. the set of all faces of the mesh, including those belonging to the outer boundary $\partial\Omega$. To introduce boundary data for the multiscale coarsening, for each $p \in \mathcal{I}^H(\overline{\Omega})$, we introduce functions $\psi_p : \Gamma \to \mathbb{R}$ which are required to be piecewise linear (with respect to the mesh $\mathcal{T}^h$ on Γ) and are required also to satisfy the assumptions:

(**M1**) $\psi_p(x^H_{p'}) \;=\; \delta_{p,p'} \;, \quad p, p' \in \mathcal{I}^H(\overline{\Omega}),$

(**M2**) $0 \;\le\; \psi_p(x) \;\le\; 1 \;, \quad \text{and} \quad \sum_{p\in\mathcal{I}^H(\overline{\Omega})} \psi_p(x) \;=\; 1 \;, \quad \text{for all} \;\; x \in \Gamma \;,$

(**M3**) $\psi_p \;\equiv\; 0 \quad$ on all faces $f \in \mathcal{F}^H$ such that $x^H_p \notin f$.

Using ψ_p as boundary data, for each $p \in \mathcal{I}^H(\overline{\Omega})$, the basis functions $\Phi_p \in \mathcal{S}^h(\Omega)$, are then defined by discrete $\alpha-$harmonic extension of ψ_p into the interior of each $K \in \mathcal{T}^H$. That is, for each $K \in \mathcal{T}^H$, $\Phi_p|_K \in \{v_h \in \mathcal{S}^h(K) : \Phi_p|_{\partial K} = \psi_p|_{\partial K}\}$ is such that

$$\int_K \alpha\nabla(\Phi_p|_K)\cdot\nabla v_h \;=\; 0 \quad \text{for all} \quad v_h \in \mathcal{S}^h_0(K)\,. \tag{III.1}$$

Note that since ψ_p vanishes on the boundaries of all coarse elements K which do not contain x^H_p, the number of solves in (III.1) remains $\mathcal{O}(1)$ for each p when $H \to 0$ (since the coarse mesh $\mathcal{T}^H$ is assumed shape regular). The obvious example of boundary data ψ_p satisfying (M1)–(M3) are the standard hat functions restricted to the faces (edges) of the tetrahedron (triangle) K.

However, these are not so appropriate if α varies strongly near the boundary ∂K. The **oscillatory** boundary conditions suggested in [10] are more useful in this case: Let e be an arbitrary edge of the coarse mesh $\mathcal{T}^H$ with end points x^H_p and x^H_q, say, and let α^e denote a (suitable) piecewise constant restriction of α to e (defined for example by averaging). We introduce the solution (with respect to $\mathcal{T}^h$ restricted to e) of the two–point boundary value problem $-(\alpha^e(\psi^e_p)')' = 0$ with boundary conditions chosen to be 1 at x^H_p and 0 at x^H_q. This problem is easily seen to have an explicit continuous piecewise linear solution. Then the "oscillatory boundary data" on a 2D element K is defined by setting $\psi_p|_e = \psi^e_p$ on each edge e of K containing x^H_p, and $\psi_p|_e = 0$ on the other edges. This can be extended to 3D by an obvious bootstrapping procedure (see, e.g. [9]), and it is easy to see (by uniqueness and the maximum principle) that the resulting Ψ_p satisfies the assumptions (M1)–(M3). This recipe specifies basis functions Φ_p which can immediately be seen to satisfy the assumptions (C1) - (C4) (see e.g. [9]).

Matrix form of preconditioner

In order to write the two-level preconditioner in (I.4) in matrix form it is convenient to introduce the cardinalities:

$$n = \#\mathcal{N}^h(\Omega) \ , \quad N = \#\mathcal{N}^H(\Omega) \ , \quad \text{and} \quad n_\Gamma = \#\{\mathcal{N}^h(\Gamma \cap \Omega)\} \ .$$

That is n is the number of fine grid freedoms in the open set Ω, N is the number of coarse grid freedoms in Ω, and n_Γ is the number of fine grid freedoms which lie on the skeleton Γ, but not on $\partial\Omega$. Thus any vector $\mathbf{f} \in \mathbb{R}^n$ can be decomposed into

$$\mathbf{f} = \begin{bmatrix} \mathbf{f}_I \\ \mathbf{f}_B \end{bmatrix}$$

where $\mathbf{f}_I \in \mathbb{R}^{n-n_\Gamma}$ are its values at nodes in $\Omega\backslash\Gamma$ and $\mathbf{f}_B \in n_\Gamma$ are its values at nodes in $\Gamma \cap \Omega$. Correspondingly (following common procedure), the stiffness matrix A in (I.2) can be partitioned

$$A = \begin{bmatrix} A_{II} & A_{IB} \\ A_{BI} & A_{BB} \end{bmatrix} .$$

Note that A_{II} is a block diagonal matrix, with each block corresponding to freedoms in the interior of a particular coarse mesh element. After a little algebra, the exact inverse of A can be written:

$$A^{-1} \ = \ \begin{bmatrix} -A_{II}^{-1}A_{IB} \\ I_{BB} \end{bmatrix} \left[S^{-1}\right] \begin{bmatrix} -A_{BI}A_{II}^{-1} & I_{BB} \end{bmatrix} \ + \ \begin{bmatrix} A_{II}^{-1} & 0 \\ 0 & 0 \end{bmatrix} \tag{III.2}$$

where $S = A_{BB} - A_{BI}A_{II}^{-1}A_{IB}$ is the Schur complement of A_{II} in A and I_{BB} denotes the $n_\Gamma \times n_\Gamma$ identity matrix. (See also [21, p.95].)

The following lemma shows that the coarse grid component in the preconditioner (I.4) in the case of multiscale coarsening is a natural approximation of the first term on the right-hand side of (III.2). To state and prove it we need the following notation. For $p \in \mathcal{I}^H(\Omega)$, the coarse basis function Φ_p is identified with its vector $\boldsymbol{\Phi}_p \in \mathbb{R}^n$ of nodal values and correspondingly, ψ_p is identified with its vector $\boldsymbol{\psi}_p \in \mathbb{R}^{n_\Gamma}$ of nodal values.

Lemma 3.1.

$$R_0^T A_0^{-1} R_0 = \begin{bmatrix} -A_{II}^{-1}A_{IB} \\ I_{BB} \end{bmatrix} \left[Q^T (QSQ^T)^{-1} Q\right] \begin{bmatrix} -A_{BI}A_{II}^{-1} & I_{BB} \end{bmatrix} \quad ,$$

where Q^T is the $N \times n_\Gamma$ matrix whose p-th column is $\boldsymbol{\psi}_p$.

Proof. Let $p \in \mathcal{I}^H(\overline{\Omega})$. Note first that each of the problems in (III.1) constitutes a FE discretisation of a local Dirichlet problem with Dirichlet data $\psi_p|_{\partial K}$ on the fine mesh involving only freedoms of Φ_p in K. Since ψ_p was assumed piecewise linear, in matrix notation this is equivalent to

$$\begin{bmatrix} A_{II} & 0 \\ 0 & I_{BB} \end{bmatrix} \begin{bmatrix} \boldsymbol{\Phi}_p \end{bmatrix} = \begin{bmatrix} -A_{IB} \\ I_{BB} \end{bmatrix} \begin{bmatrix} \boldsymbol{\psi}_p \end{bmatrix}.$$

Hence, by definition (I.3),

$$R_0^T = [\boldsymbol{\Phi}_1, \boldsymbol{\Phi}_2, \ldots, \boldsymbol{\Phi}_N] \ = \ \begin{bmatrix} -A_{II}^{-1}A_{IB} \\ I_{BB} \end{bmatrix} \begin{bmatrix} \boldsymbol{\psi}_1, \boldsymbol{\psi}_2, \ldots, \boldsymbol{\psi}_N \end{bmatrix} = \begin{bmatrix} -A_{II}^{-1}A_{IB} \\ I_{BB} \end{bmatrix} Q^T.$$

Thus an easy calculation shows $A_0 = R_0 A R_0^T = QSQ^T$, and the result follows directly. ∎

Thus from Lemma 3.1 we see immediately that the preconditioner $M_{AS,2}^{-1}$ defined in (I.4) can be written

$$M_{AS,2}^{-1} \;=\; \begin{bmatrix} -A_{II}^{-1}A_{IB} \\ I_{BB} \end{bmatrix} \left[Q^T(QSQ^T)^{-1}Q\right] \begin{bmatrix} -A_{BI}A_{II}^{-1} & I_{BB} \end{bmatrix} \;+\; \sum_{i=1}^{N} R_i A_i^{-1} R_i^T \;. \tag{III.3}$$

Comparing (III.3) with (III.2), we notice two differences: firstly the exact inverse S^{-1} of the Schur complement in (III.2) (defined on all fine freedoms on the skeleton) has been replaced by an approximation $Q^T(QSQ^T)^{-1}Q$, where the solve $(QSQ^T)^{-1}$ has to be carried out only on the coarse mesh freedoms.

Secondly, while the second term on the right-hand side of (III.2) consists of the sum of inverses of A on interior freedoms of each coarse element, the corresponding term in (III.3) consists of a sum of inverses of A on subdomains. Because of assumption (C5), the subdomain solves will always be over domains which are at least as big (asymptotically) as the coarse elements. In this sense the second term on the right hand side of (III.3) provides a better preconditioning because the local solves are now over bigger regions which overlap.

Continuing the comparison further we notice that, in the special case of 1D, the skeleton is the whole of the coarse grid and so the interpolation matrix Q is always the identity. Thus the first terms on the right-hand side of (III.2) and (III.3) coincide. This ensures that in 1D the multiscale preconditioner has condition number bounded independently of mesh parameters and of the coefficient α.

To justify this assertion, consider the simplest case where each subdomain Ω_i is the extension of a coarse grid element $K \in \mathcal{T}^H$. Denoting the diagonal blocks in the last term in (III.2) by A_K we have, for any vector $\mathbf{u} \in \mathbb{R}^n$, from (III.2) and Lemma 3.1,

$$\mathbf{u} \;=\; A^{-1}A\mathbf{u} \;=\; R_0^T A_0^{-1} R_0 A\mathbf{u} \;+\; \sum_{K\in\mathcal{T}^H} R_K^T A_K^{-1} R_K A\mathbf{u} \;, \tag{III.4}$$

where $R_K\mathbf{u}$ is the restriction of a vector $\mathbf{u}$ to its values at nodes in the interior of K. Also note that $R_0^T A_0^{-1} R_0 A$ and $R_K A_K^{-1} R_K A$ are the orthogonal projections, with respect to the inner product induced by A, onto range(R_0^T) and range(R_K^T), respectively. Therefore it follows from (III.4) that

$$\|\mathbf{u}\|_A^2 \;=\; \|R_0^T A_0^{-1} R_0 A\mathbf{u}\|_A^2 \;+\; \sum_{K\in\mathcal{T}^H} \|R_K^T A_K^{-1} R_K A\mathbf{u}\|_A^2 \;, \tag{III.5}$$

where $\|\cdot\|_A$ is the norm induced by A. Now, let $\Omega_{i(K)}$ be the subdomain associated with $K \in \mathcal{T}^H$. Then range$(R_K^T) \subset$ range$(R_{i(K)}^T)$, and so (III.4) represents a decomposition of $\mathbf{u}$ into components in range(R_i^T), for $i = 0, \ldots, N$. But $M_{AS,2}^{-1}A$ is simply an analogous sum of orthogonal projections onto range(R_i^T), $i = 0, \ldots, N$. Hence (III.5) and Lions' Lemma (see [21, Lemma 2.5]) ensure that the minimum eigenvalue of $M_{AS,2}^{-1}A$ is no smaller than 1. The maximum eigenvalue is bounded above by a colouring argument (independently of mesh parameters and the coefficient α) and the claim about the 1D case follows.

IV. OTHER VARIANTS

In this section we describe two variants of the additive Schwarz preconditioner (I.4) and give condition number estimates analogous to those in Theorem 2.1 .

Hybrid additive-multiplicative Schwarz

This preconditioner has been first introduced and analysed in [13, 14]. It is obtained by combining the one-level additive method (consisting only of the local subdomain solves) in a multiplicative way with the coarse solve. In order to ensure that the resulting preconditioner is symmetric we need to apply the coarse solve twice. We formulate this preconditioner by first defining the matrix

$$Q_0 = I - AR_0^T\, A_0^{-1}\, R_0 \ . \tag{IV.1}$$

(Note that $I - Q_0^T = R_0^T A_0^{-1} R_0 A$, which is the orthogonal projection discussed in the previous section. We work with Q_0 in this section since it relates best to the notation in the existing literature and simplifies the comparison with deflation below.)

A property of Q_0 which we will use frequently is that it "commutes" with A, i.e.

$$Q_0 A \ = \ A - AR_0^T\, A_0^{-1}\, R_0 A \ = \ AQ_0^T \ . \tag{IV.2}$$

The hybrid preconditioner is now defined as

$$M_{\rm Hyb}^{-1} = R_0^T\, A_0^{-1}\, R_0 + Q_0^T\, M_{\rm AS,1}^{-1} Q_0 \ . \tag{IV.3}$$

The following theorem can be found in [13, Lemma 3.2] (see also [21, Lemma 2.15]).

Theorem 4.1.

$$\kappa(M_{\rm Hyb}^{-1}A) \ \leq \ \kappa(M_{\rm AS,2}^{-1}A) \ .$$

Combining this theorem with Theorem 2.1, we have the following corollary for the hybrid preconditioner.

Corollary.

$$\kappa\left(M_{\rm Hyb}^{-1}A\right) \ \lesssim \ \pi(\alpha)\,\gamma(1)\max_{i=1}^{N}\left(1+\frac{H_i}{\delta_i}\right) \ + \ \gamma(\alpha) \ .$$

Thus the hybrid preconditioner is also robust provided $\gamma(\alpha)$ and $\pi(\alpha)$ are bounded. The inequality in Theorem 4.1 is sharp and we will see in Section V. that in practice the introduction of the hybrid preconditioner instead of the additive one does not lead to a significant improvement in the number of CG iterations in many test cases – in particular when $M_{\rm AS,2}^{-1}$ is already an effective preconditioner. However, the introduction of the hybrid preconditioner does seem to lead to a significant improvement in the case of random coefficients where the additive preconditioner is less robust.

Although it may seem that the application of the hybrid preconditioner requires more work than the application of the standard additive Schwarz method, the extra work can be minimised after making the following observations.

Note first that, provided we choose an initial guess for the CG iteration of the form: $\mathbf{u}^{(0)} = R_0^T\, A_0^{-1}\, R_0\mathbf{f} + \mathbf{w}$, with $\mathbf{w} \in \text{range}(Q_0^T)$, subsequent CG iterates satisfy

$$\mathbf{u}^{(k)} - \mathbf{u}^{(0)} \ \in \ \text{range}(Q_0^T) \qquad \text{for all} \ \ k \geq 0 \tag{IV.4}$$

(see [21, Lemma 2.11]). Now recall that in each CG iteration the preconditioner M_{Hyb}^{-1} is applied to the current residual $\mathbf{r}^{(k)} = \mathbf{f} - A\mathbf{u}^{(k)}$. However, by (IV.2) and the assumption on the initial guess,

$$\mathbf{r}^{(0)} \;=\; \mathbf{f} - A\mathbf{u}^{(0)} \;=\; Q_0\mathbf{f} - A\mathbf{w} \;\in\; \text{range}(Q_0 A)\ .$$

Together with (IV.2) and (IV.4) this implies

$$\mathbf{r}^{(k)} \;\in\; \text{range}(Q_0 A) \qquad \text{for all}\ \ k \geq 0. \tag{IV.5}$$

and so the action of Q_0 does not need to be calculated in practice. Since we also have $R_0^T\, A_0^{-1}\, R_0\, Q_0 = 0$, the additional work required to apply the hybrid preconditioner (as opposed to the additive two-level preconditioner) is one matrix-vector product with A per CG iteration.

It is possible to reduce this extra cost further by precalculating AR_0^T and by storing it in place of R_0^T. This can be achieved with one sparse matrix-matrix product at the cost of approximately three matrix-vector products with A. The additional storage is also equivalent to a few vectors.

Deflation

Another preconditioning strategy that has proven successful when there are a few isolated near-zero eigenvalues is *deflation* (see [18, 16, 17] and the references therein). Here, given an $n \times n$ matrix A we choose a set of vectors $\mathbf{z}_1, \ldots, \mathbf{z}_m$ that span an m-dimensional subspace of $\mathbb{R}^n$ which we aim to project out of the iterates, with $m \ll n$. Following the notation in the deflation literature [16, 17], let Z be the matrix with columns $\mathbf{z}_1, \ldots, \mathbf{z}_m$ and let

$$P = I - AZ\, E^{-1}\, Z^T \qquad \text{where} \qquad E = Z^T AZ\ .$$

(Note that, anologously to (IV.1), $(I - P^T)$ is the orthogonal projection of $\mathbb{R}^n$ onto range(Z) with respect to the inner product induced by A and that again $AP^T = PA$.)

Now, to solve the system (I.2), we write

$$\mathbf{u} = P^T\mathbf{u} + (I - P^T)\mathbf{u}, \tag{IV.6}$$

so that $AP^T\mathbf{u} \;+\; A(I - P^T)\mathbf{u} \;=\; \mathbf{f}$. Then observe that, by applying P to each side of this equation and using $PA(I - P^T) = P(I - P)A = 0$, we obtain a "deflated equation" for the projection $\tilde{\mathbf{u}} := P^T\mathbf{u}$, which reads simply

$$PA\tilde{\mathbf{u}} = P\mathbf{f}\ . \tag{IV.7}$$

The matrix PA in (IV.7) is positive semidefinite (with respect to the inner product induced by A), and has rank $n - m$. To obtain the complimentary part of $\mathbf{u}$, namely $(I - P^T)\mathbf{u}$, note that $(I - P^T)\mathbf{u} = ZE^{-1}Z^T\mathbf{f}$ which can be found by solving directly a small m-dimensional system with coefficient matrix E.

The large system (IV.7) can still be solved iteratively by preconditioned CG. (This is because the right hand side of (IV.7) lies in the range of PAP^T, so this system is consistent and its *effective condition number* is positive.) The convergence rate depends on the square root of the effective condition number:

$$\kappa_{\text{eff}}(M^{-1}PA) \;=\; \frac{\lambda_n}{\lambda_{m+1}} \tag{IV.8}$$

where λ_i is the ith-largest eigenvalue of $M^{-1}PA$ and M is an SPD preconditioner [18, 16, 17].

This method is particularly useful if the m smallest eigenvalues of A are isolated from the rest of the spectrum and if $\mathbf{z}_1, \ldots, \mathbf{z}_m$ are corresponding eigenvectors. In this case the deflated matrix PA has spectrum $\{0, \ldots, 0, \lambda_{m+1}(A), \ldots, \lambda_n(A)\}$ and the performance of preconditioned CG depends only on the $n-m$ largest eigenvalues of A. This is of particular relevance here, since standard one-level preconditioners for (I.1) often produce highly clustered spectra with relatively few near–zero eigenvalues [7, 8, 2, 22]. However, usually the eigenvectors corresponding to the m smallest eigenvalues of A are not available in practice, and deflation would be too expensive were they computed for this purpose only. For this reason deflation preconditioners for (I.1) usually employ so-called subdomain deflation (see [18, 16, 17]). The success of subdomain deflation is analysed in [16, 17] by comparing it to additive and multiplicative coarse grid correction in Schwarz methods. Since we have already established in Theorem 4.1 that the multiplicative coarse grid correction is no worse than the additive one in Schwarz methods, we only quote the result from [17] comparing deflation to multiplicative coarse grid correction.

Theorem 4.2. *For any symmetric positive definite M,*

$$\kappa_{\text{eff}}(M^{-1}PA) \leq \kappa((ZE^{-1}Z^T + P^T M^{-1} P)A) \ .$$

The result in Theorem 4.2 is true for any choice of Z and M. In particular we may choose $Z = R_0^T$ (so that in this case $P = Q_0$) and $M = M_{\text{AS},1}$, and in conjunction with the corollary to Theorem 4.1 we have the following corollary.

Corollary.

$$\kappa_{\text{eff}}(M_{\text{AS},1}^{-1}PA) \ \leq \ \kappa(M_{\text{Hyb}}^{-1}A) \ \lesssim \ \pi(\alpha)\,\gamma(1)\,\max_{i=1}^{N}\left(1+\frac{H_i}{\delta_i}\right)+\gamma(\alpha) \ .$$

This shows that subdomain deflation with $Z = [\boldsymbol{\Phi}_1, \boldsymbol{\Phi}_2, \ldots, \boldsymbol{\Phi}_N]$ also leads to a robust preconditioner provided $\gamma(\alpha)$ and $\pi(\alpha)$ are bounded. Here the deflation vectors are the nodal values of the coarse space basis functions $\boldsymbol{\Phi}_p$.

Note that in [18, 16, 17], given a (non-overlapping) partitioning of the set of freedoms $\mathcal{N}^h(\Omega)$ into m "aggregates" $\{W_j : j = 1, \ldots, m\}$ (usually related to physical subdomains of Ω), the (subdomain) deflation vectors $\{\mathbf{z}_j\}$ are chosen such that

$$(z_j)_p \ = \ \begin{cases} 1 & \text{if } \ p \in W_j, \\ 0 & \text{otherwise,} \end{cases}$$

i.e. subdomain deflation with piecewise constant interpolation onto the subdomains. The robustness of this type of (aggregation-based) coarse space is analysed in [19]. It is very similar to the theory presented in §II, and in particular, it is shown there that

$$\kappa(M_{AS,2}^{-1}A) \ \lesssim \ \max_{j=1}^{m}\left\{\delta_j^2\|\alpha|\nabla\zeta_j|^2\|_{L_\infty(\Omega)}\right\}\max_{j=1}^{m}\left(1+\frac{H_j}{\delta_j}\right) \qquad \text{(IV.9)}$$

where ζ_j is the finite element function associated with the nodal vector $\mathbf{z}_j$ and H_j is the diameter of the support of the function ζ_j in Ω. The overlap parameter is $\delta_j = h$ in this case. Theorems 4.1 and 4.2 imply that the same bound as in (IV.9) also holds for $\kappa_{\text{eff}}(M_{\text{AS},1}^{-1}PA)$. This means that unless care is taken that $\nabla\zeta_j$ is small wherever α is

large the original subdomain deflation preconditioners in [18, 16, 17] are not robust with respect to heterogeneities in α. In particular, if $\alpha \to \infty$ in an element $\tau \in \mathcal{T}^h$ in the overlap between two aggregates then $\kappa_{\text{eff}}(M_{\text{AS},1}^{-1}PA)$ may grow unboundedly. Techniques which avoid this and ensure coarse space robustness also in the case of aggregation are described in [19]. These techniques use ideas from algebraic multigrid to construct the aggregates W_j based on strong connections in A.

We finish this section with a result that links the hybrid preconditioner and deflation even further. Provided the appropriate initial guess is chosen in each case, the deflation preconditioner is in fact simply a different implementation of the hybrid preconditioner. We state and prove the theorem for $Z = R_0^T$ and $M = M_{\text{AS},1}$, but the result is true for arbitrary (full-rank) Z and (symmetric positive definite) M and the proof is identical.

Theorem 4.3. *Let $\mathbf{u}^{(k)}$ denote the kth CG iterate for solving (I.2) with preconditioner M_{Hyb}^{-1} and let $\tilde{\mathbf{u}}^{(k)}$ denote the kth CG iterate for solving the deflated system (IV.7) with $Z = R_0^T$ and preconditioner $M_{\text{AS},1}^{-1}$. If we choose the initial guesses for (IV.7) and (I.2) such that $\tilde{\mathbf{u}}^{(0)} \in \text{range}(Q_0^T)$ and $\mathbf{u}^{(0)} = R_0^T A_0^{-1} R_0 \mathbf{f} + \tilde{\mathbf{u}}^{(0)}$, respectively, then in exact arithmetic*

$$\mathbf{u}^{(k)} \;=\; R_0^T A_0^{-1} R_0 \mathbf{f} \;+\; Q_0^T \tilde{\mathbf{u}}^{(k)} \qquad \textit{for all} \;\; k \geq 0.$$

Proof. This result was first observed by Reinhard Nabben [15], but we know of no proof in the literature. We prove by induction that

$$\mathbf{r}^{(k-1)} \;=\; \tilde{\mathbf{r}}^{(k-1)} \qquad \text{and} \qquad \mathbf{p}^{(k)} \;=\; Q_0^T \tilde{\mathbf{p}}^{(k)} \qquad \text{for all} \;\; k \geq 1, \tag{IV.10}$$

where $\mathbf{p}^{(k)}$ and $\tilde{\mathbf{p}}^{(k)}$ are the search directions in the kth iteration of CG in the hybrid and in the deflation case, respectively; $\mathbf{r}^{(k)}$ is the kth residual in the hybrid case and $\tilde{\mathbf{r}}^{(k)} = Q_0(\mathbf{f} - A\tilde{\mathbf{u}}^{(k)})$ is the (projected) residual in the deflation case (see e.g. [21, Appendix C] for a description of the CG method). The result then follows easily from (IV.10) together with (IV.2), i.e.

$$\mathbf{u}^{(k)} \;=\; A^{-1}(\mathbf{f} - \mathbf{r}^{(k)}) \;=\; A^{-1}(\mathbf{f} - Q_0\mathbf{f} + Q_0 A\tilde{\mathbf{u}}^{(k)}) \;=\; R_0^T A_0^{-1} R_0 \mathbf{f} \;+\; Q_0^T \tilde{\mathbf{u}}^{(k)}.$$

Now to prove (IV.10) note first that by (IV.5) and the subsequent discussion we have

$$M_{\text{Hyb}}^{-1}\mathbf{r}^{(k)} \;=\; (R_0^T A_0^{-1} R_0 + Q_0^T M_{\text{AS},1}^{-1} Q_0)\mathbf{r}^{(k)} \;=\; Q_0^T M_{\text{AS},1}^{-1}\mathbf{r}^{(k)}. \tag{IV.11}$$

By (IV.2) and the assumption on the initial guess, we have

$$\mathbf{r}^{(0)} \;=\; \mathbf{f} - AR_0^T A_0^{-1} R_0 \mathbf{f} - A\tilde{\mathbf{u}}^{(0)} \;=\; Q_0(\mathbf{f} - A\tilde{\mathbf{u}}^{(0)}) \;=\; \tilde{\mathbf{r}}^{(0)}$$

Hence, using (IV.11), $\mathbf{p}^{(1)} = M_{\text{Hyb}}^{-1}\mathbf{r}^{(0)} = Q_0^T M_{\text{AS},1}^{-1}\mathbf{r}^{(0)} = Q_0^T\tilde{\mathbf{p}}^{(1)}$.

Now assume that $\mathbf{r}^{(k-1)} = \tilde{\mathbf{r}}^{(k-1)}$ and $\mathbf{p}^{(k)} = Q_0^T\tilde{\mathbf{p}}^{(k)}$ for some $k \geq 1$. This implies

$$(\mathbf{r}^{(k-1)})^T M_{\text{Hyb}}^{-1}\mathbf{r}^{(k-1)} = (\tilde{\mathbf{r}}^{(k-1)})^T M_{\text{AS},1}^{-1}\tilde{\mathbf{r}}^{(k-1)} \quad \text{and} \quad (\mathbf{p}^{(k)})^T A\mathbf{p}^{(k)} = (\tilde{\mathbf{p}}^{(k)})^T Q_0 A\tilde{\mathbf{p}}^{(k)}$$

and so the search parameters in the hybrid and deflation cases coincide, i.e. $\alpha^{(k)} = \tilde{\alpha}^{(k)}$. Similarly, $\beta^{(k)} = \tilde{\beta}^{(k)}$. Moreover, we also have using (IV.2) again that

$$\mathbf{r}^{(k)} \;=\; \mathbf{r}^{(k-1)} - \alpha^{(k)} A\mathbf{p}^{(k)} \;=\; \tilde{\mathbf{r}}^{(k-1)} - \tilde{\alpha}^{(k)} AQ_0^T\tilde{\mathbf{p}}^{(k)} \;=\; \tilde{\mathbf{r}}^{(k)}$$

and $\mathbf{p}^{(k+1)} = M_{\text{Hyb}}^{-1}\mathbf{r}^{(k)} + \beta^{(k)}\mathbf{p}^{(k)} = Q_0^T(M_{\text{AS},1}^{-1}\tilde{\mathbf{r}}^{(k)} + \tilde{\beta}^{(k)}\tilde{\mathbf{p}}^{(k)}) = Q_0^T\tilde{\mathbf{p}}^{(k+1)}$. ∎

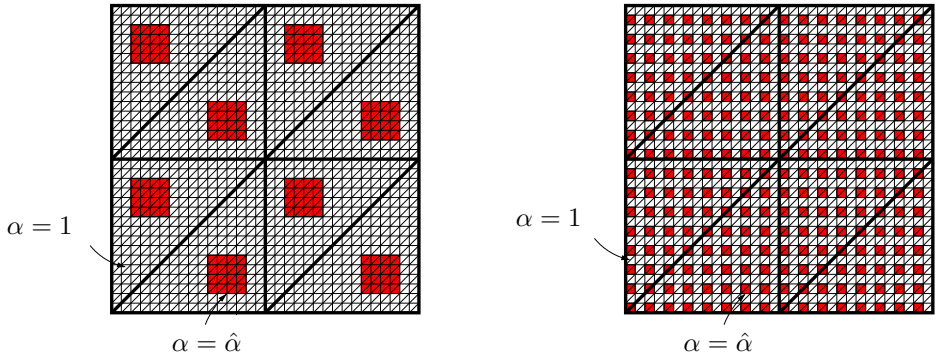

FIG. 1. Examples 1 (left) and 2 (right) for $r = 5$ and $R = 1$.

V. NUMERICAL EXPERIMENTS

In this section, by a series of examples involving "binary" media and random media, we explain how our analysis in Section II yields sharp estimates for standard domain decomposition methods and, moreover leads to new effective robust multiscale preconditioners.

We will work with the unit square $\Omega = [0,1]^2$ and let $\mathcal{T}^h$ be the family of uniform (isosceles) triangulations of Ω obtained by taking a uniform mesh of squares of side $h = 2^{-r}$ (for some $r \in \mathbb{N}$) and subdividing each by a line joining the bottom left to top right corners. Analogously, let $\mathcal{T}^H$ be a uniform family of coarse meshes with mesh width $H = 2^{-R}$, $R < r$, so that each $K \in \mathcal{T}^H$ is a union of a set of fine grid elements as assumed above. For each coarse mesh $\mathcal{T}^H$ we consider two types of coarse space basis functions: standard piecewise linear basis functions and the multiscale basis functions constructed in Section III with oscillatory boundary conditions. They will be denoted by Φ_p^{L} and Φ_p^{M} for each $p \in \mathcal{I}^H(\overline{\Omega})$, respectively. The corresponding two–level additive Schwarz preconditioners will be denoted by $M_{AS,L}^{-1}$ and $M_{AS,M}^{-1}$, respectively, while the one–level method is still denoted by $M_{AS,1}^{-1}$. Moreover, $\gamma^{\mathrm{L}}(\alpha)$ and $\gamma^{\mathrm{M}}(\alpha)$ will denote the corresponding coarse space robustness indicators as defined in Section II. The overlapping open covering $\{\Omega_i : i = 1, \ldots, N\}$ of Ω is obtained from $\mathcal{T}^H$ by extending each element $K \in \mathcal{T}^H$ with β layers of fine grid elements, where $\beta \in \mathbb{N}$ is fixed as $h \to 0$. It follows that the overlap parameter δ_i satisfies $\delta_i \sim \beta h$ for each $i = 1, \ldots, N$.

Let us first demonstrate how our analysis in Section II yields sharp estimates by studying the condition numbers of the preconditioned stiffness matrices in the case of two particular choices of $\alpha(x)$. Here, we restrict to the additive variant (I.4). The hybrid and the deflation preconditioner will be studied later.

Example 1. Here we take $r \geq R + 3$ (i.e. $H \geq 8h$) and let $\alpha(x)$ describe a binary medium where $\alpha(x) = \hat{\alpha}$ on a square island of width $H/4$ in the "centre" of each coarse element $K \in \mathcal{T}^H$. The islands are chosen such that they are located at a distance $H/8$ from the horizontal and vertical edges of K (see FIG. 1 (left)). In the rest of the domain Ω we choose $\alpha(x) = 1$. We study the behaviour of our preconditioners when $\hat{\alpha} \to \infty$.

Note that for this choice of α it can be shown that $\gamma^{\mathrm{M}}(\alpha) \lesssim 1$ and $\pi(\alpha) \sim 1$ (see [9] for a proof). On the other hand, $\gamma^{\mathrm{L}}(\alpha) \to \infty$ as $\hat{\alpha} \to \infty$.

The first set of numerical results in Table 1 shows the loss of robustness of two–level additive Schwarz with linear coarsening in Example 1. It also demonstrates that $\gamma^{\mathrm{L}}(\alpha)$

is a good indicator for this loss of robustness. Moreover, the results show that in the case of linear coarsening the two–level method is performing asymptotically like the one–level method. This is in accordance with the theory since $\kappa(M_{AS,2}^{-1}A) \lesssim \kappa(M_{AS,1}^{-1}A)$ for any choice of the coarse space (cf. [9] for details). In contrast, the last two columns in Table 1 highlight the robustness of multiscale coarsening and show that $\gamma^{\mathrm{M}}(\alpha)$ is also a good indicator for this robustness. The observations confirm the predictions of Theorem 2.1 with respect to variations in α.

TABLE 1. Standard one–level and two–level additive Schwarz preconditioning (with linear and multiscale coarsening) for Example 1 with $h = \frac{1}{256}$, $H = 8h$ and $\delta = 2h$.

$\hat{\alpha}$	$\kappa(M_{AS,1}^{-1}A)$	$\kappa(M_{AS,L}^{-1}A)$	$\gamma^{\mathrm{L}}(\alpha)$	$\kappa(M_{AS,M}^{-1}A)$	$\gamma^{\mathrm{M}}(\alpha)$
10^0	8410	22.0	3.0	22.0	3.0
10^2	6100	111.0	40.1	17.7	4.26
10^4	6040	3870	3.75(+3)	17.6	4.31
10^6	6040	6000	3.75(+5)	17.6	4.31

The results in Table 2 show the sharpness of the $\max_i(1 + H_i/\delta_i)$ term in the bound in Theorem 2.1. Finally, our third set of results in Table 3 explains more the loss of robustness of the standard method as the coarse mesh is refined, i.e. the two–level method with linear coarsening behaves asymptotically like the one–level method and degenerates as the fine mesh is refined while multiscale coarsening leads to a robust preconditioner, with respect to both h and $\hat{\alpha}$.

TABLE 2. Condition numbers for the two-level additive method with multiscale coarsening for Example 1 with $h = \frac{1}{256}$ and $\hat{\alpha} = 10^6$.

δ	$H = 8h$	$H = 16h$	$H = 32h$	$H = 64h$
$2h$	**17.6**	33.2	62.4	115.4
$4h$	9.9	**17.9**	32.8	59.4
$8h$	6.4	9.9	**17.7**	31.4
$16h$		6.4	9.8	**17.1**

TABLE 3. Condition numbers for the one-level and for the two-level methods with $H = 8h$, $h = 2^{-r}$ and $\delta = 2h$ for $\hat{\alpha} = 10^6$ in Example 1 (as the fine mesh is refined).

r	$1/h$	$\kappa(M_{AS,1}^{-1}A)$	$\kappa(M_{AS,L}^{-1}A)$	$\kappa(M_{AS,M}^{-1}A)$
7	128	1510	1510	17.5
8	256	6040	6000	17.6
9	512	24160	23630	17.7
10	1024	96640	88680	17.7

Example 2. Here, we want to let the areas with large coefficients also touch the edges of our coarse mesh $\mathcal{T}^H$ and investigate the effectiveness of the oscillatory boundary conditions in this case. Let $\alpha(x)$ describe a binary medium, where $\alpha(x) = \hat{\alpha}$ on uniformly placed, square islands of diameter h that are separated by exactly one layer of fine grid elements (see Figure 1 (right)). In the rest of the domain we choose $\alpha(x) = 1$. We study again the behaviour of our preconditioners when $\hat{\alpha} \to \infty$.

It can easily be shown again that $\gamma^{\rm L}(\alpha) \to \infty$ as $\hat{\alpha} \to \infty$. In contrast, we will see in Tables 4 and 5 below that $\gamma^{\rm M}(\alpha)$ is bounded as $\hat{\alpha} \to \infty$.

Our first set of numerical results for Example 2 in Table 4 was obtained with minimal subdomain overlap, i.e. with $\beta = 1$ (i.e. $\delta = 2h$). In this case $\pi(\alpha) \sim \hat{\alpha}$ and thus grows unboundedly as $\hat{\alpha} \to \infty$ (see [9] for details). Therefore none of the preconditioners is robust, even though we see in the last column that in the case of multiscale coarsening (with oscillatory boundary conditions) we have coarse space robustness.

TABLE 4. Condition numbers for Example 2 with $h = \frac{1}{256}$, $H = 8h$ and $\delta = 2h$ (i.e. $\beta = 1$).

$\hat{\alpha}$	$\kappa(M_{AS,1}^{-1}A)$	$\kappa(M_{AS,L}^{-1}A)$	$\kappa(M_{AS,M}^{-1}A)$	$\gamma^{\rm M}(\alpha)$
10^0	8.41(+3)	2.20(+1)	2.20(+1)	3.0
10^2	6.40(+4)	2.36(+2)	2.31(+2)	7.9
10^4	5.11(+6)	2.13(+4)	2.07(+4)	8.0
10^6	$> 10^8$	$> 10^6$	$> 10^6$	8.0

On the other hand, for $\beta \geq 2$, it can be shown that $\pi(\alpha) \lesssim \beta^2$ (see again [9] for details). The results in Table 5 (column 5) confirm this, i.e. as predicted by our theory, multiscale coarsening with oscillatory boundary conditions leads to a robust two–level preconditioner also in Example 2. The coarse space robustness indicator $\gamma^{\rm M}(\alpha)$ is able to predict this behaviour accurately. As in Example 1, linear coarsening leads to a two–level method which performs no better than the one-level method as $\hat{\alpha} \to \infty$ (Table 5, column 3).

TABLE 5. Condition numbers for Example 2 with $h = \frac{1}{256}$, $H = 8h$ and $\delta = 4h$ (i.e. $\beta = 2$).

$\hat{\alpha}$	$\kappa(M_{AS,1}^{-1}A)$	$\kappa(M_{AS,L}^{-1}A)$	$\kappa(M_{\rm Hyb,L}^{-1}A)$	$\kappa(M_{AS,M}^{-1}A)$	$\kappa(M_{\rm Hyb,M}^{-1}A)$	$\gamma^{\rm M}(\alpha)$
10^0	3300	11.9	10.4	11.9	10.4	3.0
10^2	3430	116.0	43.1	12.0	10.4	7.9
10^4	3440	2650	1840	12.0	10.4	8.0
10^6	3440	3430	3410	12.0	10.4	8.0

Table 5 also gives the condition numbers in the case of the hybrid preconditioner defined in Section IV, both for standard piecewise linear and for multiscale coarsening – denoted by $M_{\rm Hyb,L}^{-1}$ and $M_{\rm Hyb,M}^{-1}$, respectively. As predicted (cf. Theorem 4.1) the condition number for the hybrid preconditioner is never any worse than that of the corresponding additive preconditioner. However, more importantly, we see that it is the choice of the coarse space functions which is crucial for the robustness of the preconditioner, and not the way the coarse solve is combined with the local subdomain solves.

We will now explore the efficiency of the new coarsening strategies by using our preconditioners within a preconditioned Conjugate Gradient (CG) method for (I.2) with $\mathbf{f} = \mathbf{1}$. The stopping criterion for CG is a reduction in the Euclidean norm of the residual by a factor $\varepsilon = 10^{-6}$. Apart from the additive preconditioners $M_{AS,1}^{-1}$, $M_{AS,L}^{-1}$ and $M_{AS,M}^{-1}$, already introduced above, we will also include results with the hybrid and the deflation preconditioners defined in Section IV. The initial guess in the CG method is $\mathbf{u}^{(0)} = R_0^T\, A_0^{-1}\, R_0\mathbf{f}$ for all the two-level methods and $\mathbf{u}^{(0)} = \mathbf{0}$ for the one-level method.

In Tables 6 and 7 we compare for varying problem sizes n, the number of CG iterations in the case of $\hat{\alpha} = 10^6$ in Examples 1 and 2, respectively. The results confirm the

statement made in Theorem 4.3, i.e. the number of CG iterations with the hybrid and the deflation preconditioners is identical in each case (up to the effect of rounding errors). The results confirm again the robustness of multiscale coarsening in comparison to linear coarsening and show that the way the coarse solve is combined with the local subdomain solves is not crucial.

TABLE 6. CG iterations for Example 1 with $H = 8h$, $\delta = 2h$ and $\hat{\alpha} = 10^6$.

			Linear Coarsening			Multiscale Coarsening		
r	n	1-Level	Additive	Hybrid	Deflation	Additive	Hybrid	Deflation
7	1.61(+4)	77	79	76	76	22	21	21
8	6.45(+4)	153	150	145	145	22	20	20
9	2.58(+5)	292	287	287	287	20	19	19
10	1.03(+6)	586	574	573	575	21	18	18

TABLE 7. CG iterations for Example 2 with $H = 8h$, $\delta = 4h$ and $\hat{\alpha} = 10^6$.

			Linear Coarsening			Multiscale Coarsening		
r	n	1-Level	Additive	Hybrid	Deflation	Additive	Hybrid	Deflation
7	1.61(+4)	77	100	102	102	22	26	26
8	6.45(+4)	144	185	187	185	22	24	24
9	2.58(+5)	292	355	362	362	22	21	21
10	1.03(+6)	534	681	730	729	21	21	21

In Table 8 we compare the setup time for each of the preconditioners, and the total CPU–time in the case of Example 2. The CPU–times were all obtained on a 3GHz Intel P4 processor. The coarse problem and all the local problems were solved using banded solvers in `LAPACK`. We see that the extra setup cost to construct the coarse space basis functions in the multiscale case is negligible. Moreover, the highly reduced iteration count leads to a speedup on the finest mesh of over 20 in comparison to the standard two-level method with linear coarsening and of over 15 in comparison to the one-level method. We note the slightly worse timings for the hybrid method compared to the additive variant in this example.

TABLE 8. Total CPU–time (in secs) for the problem in Example 2 with $H = 8h$, $\delta = 4h$ and $\hat{\alpha} = 10^6$. The setup times for the preconditioners are given in brackets.

			Linear Coarsening		Multiscale Coarsening	
r	n	1-Level	Additive	Hybrid	Additive	Hybrid
7	1.61(+4)	0.65 (0.06)	0.87 (0.06)	0.95 (0.06)	0.26 (0.07)	0.30 (0.07)
8	6.45(+4)	4.71 (0.21)	6.68 (0.26)	7.16 (0.26)	1.11 (0.31)	1.23 (0.31)
9	2.58(+5)	38.8 (0.87)	52.6 (1.12)	57.1 (1.12)	4.63 (1.31)	4.71 (1.31)
10	1.03(+6)	286 (3.55)	409 (4.97)	470 (4.97)	18.8 (5.72)	19.7 (5.72)

Finally, we illustrate by some numerical experiments that the multiscale method also leads to greatly improved performance over standard preconditioners for random media.

Example 3. Here, 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 (as

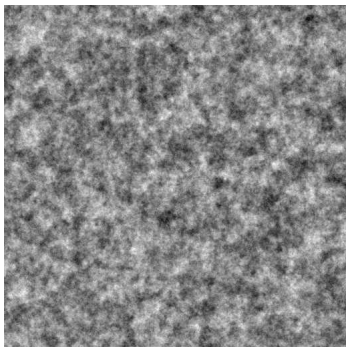

FIG. 2. Typical realisation of a log-normal random field for Example 3 ($h = \frac{1}{512}$ and $\lambda = 4h$). Black areas represent large values of α; white areas represent small values of α.

defined in e.g. Cliffe et al. [5]). This is a commonly studied model for flow in heterogeneous porous media. For more details on the physical background see e.g. [5]. We use `Gaussian` [12] to create realisations of these random fields (see Figure 2 for a grey-scale plot of a typical realisation). 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 contrast, i.e. the ratio of the largest and the smallest values of α. For example, for the field in Figure 2 with $\sigma^2 = 8$, we have $\max_{\tau,\tau' \in \mathcal{T}^h} \frac{\alpha_\tau}{\alpha_{\tau'}} = O(10^{10})$.

TABLE 9. Average number of CG iterations (over 100 realisations) for the problem in Example 3 with $h = \frac{1}{256}$, $H = 8h$, $\delta = 8h$ and $\lambda = 4h$.

			Linear Coarsening		Multiscale Coarsening	
σ^2	max $\frac{\alpha_\tau}{\alpha_{\tau'}}$	1-Level	Additive	Hybrid	Additive	Hybrid
0	1.0	77	18	14	18	14
2	1.9(+5)	180	30	18	23	16
4	3.3(+7)	252	44	25	28	19
8	5.2(+10)	453	85	47	39	25
12	1.6(+13)	730	145	80	51	32
16	2.1(+15)	1021	231	128	64	40
20	1.5(+17)	1345	349	193	79	48

In Table 9 we compare the average number of CG iterations necessary to solve (I.2) with right hand side $\mathbf{f} = \mathbf{1}$ up to a tolerance of $\varepsilon = 10^{-6}$, for 100 different realisations of α for variances between $\sigma^2 = 0$ and 20. We see that for the largest variance $\sigma^2 = 20$ multiscale coarsening performs more than four times faster than standard linear coarsening. Here we see a more substantial superiority of the hybrid method, since it improves on the additive variant by a factor of nearly two in both the linear and multiscale cases. This is also reflected in the average CPU–times. The (small) additional cost for the hybrid preconditioner is outweighed by the benefits achieved: for example in the case $\sigma^2 = 20$, the additive variant takes 7.05 seconds, while the hybrid takes 4.75 seconds. What is more, the latter is only about 2.5 times the time it takes to solve in the case of the Laplacian (i.e. $\sigma^2 = 0$), namely 1.87 seconds. The improvement is of the same order for other choices of h, H, λ and δ.

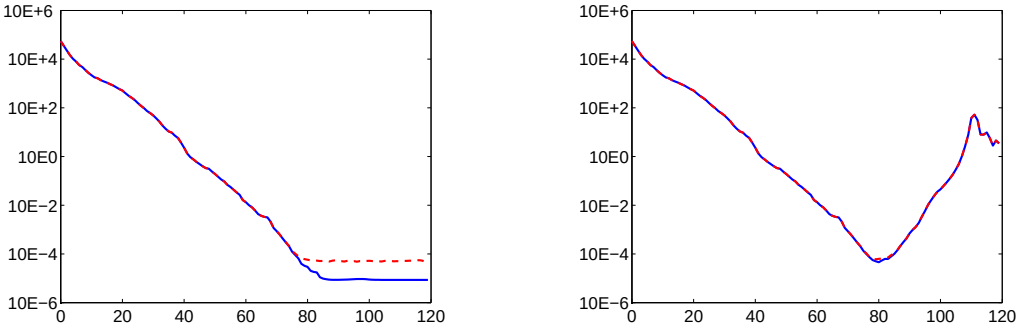

FIG. 3. Exact (dashed) and updated (solid) residual norm for Example 3 with $h = \frac{1}{256}$, $\lambda = 4h$ and $\sigma^2 = 20$, in the case of the hybrid (left) and the deflation preconditioner (right).

Our final set of result in FIG. 3 investigates the robustness of the hybrid and the deflation preconditioner. It has been shown in [16] that both the hybrid and the deflation preconditioner are sensitive to perturbations in A_0^{-1}. Extreme contrasts in the coefficient $\alpha(x)$ can lead to a severe accumulation of round-off error in the LU factors of A_0 (even with partial pivoting) and thus to large perturbations in A_0^{-1}. It was also shown in [16] that the additive two-level method is insensitive to these perturbations. The results in FIG. 3 show that the effect which these perturbations have is more severe in the case of deflation, causing the CG iteration to diverge (since $\tilde{\mathbf{r}}^{(k)}$ is no longer in the range of $Q_0 A Q_0^T$). In contrast, the hybrid preconditioned iterative method produces good solutions to the problem being solved, even in the case of the most extreme contrasts in the coefficients (i.e. for variance $\sigma^2 = 20$). In the graph the dashed line represents the Euclidean norm of the exact residual, while the solid line represents the Euclidean norm of the residual computed by updates inside the CG algorithm. The graphs show that the stopping criterion (based on CG residuals) is producing an actual residual which is within tolerance. Similar behaviour is observed in Examples 1 and 2. In conclusion, the hybrid preconditioner seems to be more robust than deflation, and since both preconditioners lead to the same CG iterates (cf. Theorem 4.3) and to the same computational cost there seems to be no benefit in using deflation. It should be noted though that in practice this issue is not so crucial since it occurs only for very large contrasts in $\alpha(x)$ and for very strict stopping criteria.

For numerical results with other choices of the overlapping covering and the coefficient function α (including multiscale basis functions with linear boundary conditions) see [9].

REFERENCES

1. J. Aarnes and T.Y. Hou, Multiscale domain decomposition methods for elliptic problems with high aspect ratios, *Acta Mathematica Applicatae Sinica*, English Ser. 18 (2002), 63–76.
2. X. Cai, B. F. Nielsen and A. Tveito, An analysis of a preconditioner for the discretized pressure equation arising in reservoir simulation. *IMA J. Numer. Anal.* 19 (1999), 291-316.
3. L.M. Carvalho, L. Giraud, and P. Le Tallec, Algebraic two-level preconditioners for the Schur complement method, *SIAM J. Sci. Comp.* 22 (2001), 1987-2005.

4. T.F. Chan and T. Mathew, Domain Decomposition Methods, *Acta Numerica* 1994 (Cambridge University Press, 1994).
5. 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 (2000), 258-282.
6. 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 (2003), 467-484.
7. 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. Bjørstad, M. Espedal and D.E. Keyes (Domain Decomposition Press, Bergen, 1998).
8. I.G. Graham and M.J. Hagger, Unstructured additive Schwarz-CG method for elliptic problems with highly discontinuous coefficients, *SIAM J Sci. Comp.* 20 (1999), 2041-2066.
9. I.G. Graham, P.O. Lechner and R. Scheichl, Domain Decomposition for Multiscale PDEs, submitted, BICS Preprint 11/06, Bath (2006), available electronically at `http://www.bath.ac.uk/math-sci/BICS/`.
10. T.Y. Hou and X.-H. Wu, A multiscale finite element method for elliptic problems in composite materials and porous media, *J. Comp. Phys.* 134 (1997), 169-189.
11. J.E. Jones and P.S. Vassilevski, AMGe based on element agglomeration, *SIAM J. Sci. Comp.* 23 (2001), 109-133.
12. B. Kozintsev and B. Kedem, `Gaussian` package, University of Maryland, 1999, available at `http://www.math.umd.edu/~bnk/bak/generate.cgi`
13. J. Mandel, Hybrid domain decomposition with unstructured subdomains, *Domain Decomposition Methods in Science and Engineering. 6th International Conference of Domain Decomposition*, Eds: J. Mandel, C. Farhat and X.-C. Cai, AMS Contemporary Mathematics 157 (1994), 103-112.
14. J. Mandel and M. Brezina, Balancing domain decomposition for problems with large jumps in coefficients, *Math. Comp.* 65 (1996), 1387-1401.
15. R. Nabben, Personal Communication, ECCOMAS CFD Conference 2006, Egmond aan Zee, The Netherlands, September 5-8, (2006).
16. R. Nabben and C. Vuik, A comparison of deflation and coarse grid correction applied to porous media flow, *SIAM J. Numer. Anal.* 42 (2004), 1631-1647.
17. R. Nabben and C. Vuik, A comparison of deflation and the balancing preconditioner, *SIAM J. Sci. Comput.* 27 (2006), 1742-1759.
18. A. Nicolaides, Deflation of conjugate gradients with applications to boundary value problems, *SIAM J. Numer. Anal.* 24 (1987), 355-365.
19. R. Scheichl and E. Vainikko, Robust aggregation-based coarsening for additive Schwarz in the case of highly variable coefficients, in *Proc. European Conference on Computational Fluid Dynamics*, ECCOMAS CFD 2006 (P. Wesseling, E. Onate, J. Periaux, eds.), (2006).
20. R. Scheichl and E. Vainikko, Additive Schwarz and aggregation-based coarsening for elliptic problems with highly variable coefficients, submitted, BICS preprint 9/06, Bath (2006), available electronically at `http://www.bath.ac.uk/math-sci/BICS/`.
21. A. Toselli and O. Widlund, *Domain Decomposition Methods Algorithms and Theory*, Springer-Verlag, Berlin-Heidelberg-New York (2005).
22. K. 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.* 21 (2000), 1632 - 1649.
23. W.L. Wan, T.F. Chan and B. Smith, An energy-minimizing interpolation for robust multigrid methods, *SIAM J Sci. Comput.* 21 (2000), 1632 - 1649.