NUMERICAL LINEAR ALGEBRA WITH APPLICATIONS
Numer. Linear Algebra Appl. 2008; **0**:1–24 *Prepared using* nlaauth.cls *[Version: 2002/09/18 v1.02]*

Energy minimizing coarse spaces for two-level Schwarz methods for multiscale PDEs

Jan Van lent[1,*], Robert Scheichl[1] and Ivan G. Graham[1]

Department of Mathematical Sciences, University of Bath, Claverton Down, Bath, BA2 7AY, U.K.

SUMMARY

Two-level overlapping Schwarz methods for elliptic partial differential equations combine local solves on overlapping domains with a global solve of a coarse approximation of the original problem. To obtain robust methods for equations with highly varying coefficients, it is important to carefully choose the coarse approximation. Recent theoretical results by the authors have shown that bases for such robust coarse spaces should be constructed such that the energy of the basis functions is minimized. We give a simple derivation of a method that finds such a minimum energy basis using one local solve per coarse space basis function and one global solve to enforce a partition of unity constraint. Although this global solve may seem prohibitively expensive, we demonstrate that a one-level overlapping Schwarz method is an effective and scalable preconditioner and we show that such a preconditioner can be implemented efficiently using the Sherman-Morrison-Woodbury formula. The result is an elegant, scalable, algebraic method for constructing a robust coarse space given only the supports of the coarse space basis functions. Numerical experiments on a simple two-dimensional model problem with a variety of binary and multiscale coefficients confirm this. Numerical experiments also show that, when used in a two-level preconditioner, the energy minimizing coarse space gives better results than other coarse space constructions, such as the multiscale finite element approach. Copyright © 2008 John Wiley & Sons, Ltd.

KEY WORDS: overlapping additive Schwarz method; coarse space; constrained energy minimization; domain decomposition; preconditioning

1. Introduction

In this paper we study the construction of scalable and robust domain decomposition methods for elliptic partial differential equations with highly variable coefficients. In particular, we want the work per iteration to be proportional to the number of unknowns. Furthermore, the number of iterations should be independent of the mesh size, the number of subdomains and, in this paper especially, the coefficients of the partial differential equation. Ideally, we would like the cost of our methods to be similar to the cost of solving the Laplace equation, i.e., a problem with constant coefficients, using an optimal solver, such as geometric multigrid.

We consider a scalar elliptic equation of the form

$$-\nabla \cdot (\mathcal{A}\nabla u) = f, \tag{1}$$

*Correspondence to: Department of Mathematical Sciences, University of Bath, Claverton Down, Bath, BA2 7AY, U.K., j.van.lent@maths.bath.ac.uk

on a bounded open domain $\Omega \subset \mathbb{R}^d$, $d = 2, 3$, with Dirichlet or mixed boundary conditions. The symmetric positive definite coefficient tensor $\mathcal{A} : \Omega \to \mathbb{R}^{d\times d}$ can be highly varying. Standard finite element discretization on a fine mesh leads to a system of equations

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

where A is a large, sparse, symmetric positive definite matrix. This problem can be solved iteratively using a preconditioner which solves subproblems corresponding to discretizations of (1) on a set of overlapping subdomains which cover Ω. For each iteration, the problems on the subdomains can be solved independently. This so-called one-level overlapping Schwarz method is not robust since the number of iterations depends on the number of subdomains. The method can be made robust by solving a coarse approximation to the original problem in addition to the subproblems on the subdomains [1, 2]. A standard two-level Schwarz preconditioner uses a coarse approximation based on a finite element discretization on a coarse mesh. This results in a method that is robust with respect to the mesh size and the number of subdomains, but, in general, not the coefficient variation [1, 2, 3]. For a completely robust method, the construction of the coarse level should take into account the coefficient variation.

In Section 5 we recall the convergence theory for overlapping additive Schwarz methods from [4, 5, 6, 7]. This theory refines the standard Schwarz theory [1, 8, 2] and gives coefficient-dependent criteria for choosing a coarse space which motivates a construction based on energy minimization. More precisely our coarse space will be spanned by a partition of unity which is piecewise linear on the fine mesh and which satisfies a certain energy minimizing criterion. Let $\bar{A}$ be the stiffness matrix corresponding to the finite element approximation of (1) on the given fine mesh with natural boundary conditions. Then $\bar{A}$ has dimension N equal to the number of nodes in $\bar{\Omega} = \Omega \cup \partial\Omega$. The matrix $\bar{A}$ is symmetric positive semidefinite with a single zero eigenvalue with eigenvector $\mathbf{1} = \begin{bmatrix} 1 & \cdots & 1 \end{bmatrix}^T$. The energy minimization problem which we are concerned with consists in finding vectors $\bar{\Phi}_j \in \mathbb{R}^N$, that minimize the energy functional

$$\sum\nolimits_j \bar{\Phi}_j^T \bar{A} \bar{\Phi}_j, \tag{3}$$

subject to the condition

$$\sum\nolimits_j \bar{\Phi}_j = \mathbf{1}. \tag{4}$$

In addition we require that the i-th entry of $\bar{\Phi}_j$ is nonzero only for indices $i \in \omega_j^h$. The set $\{\omega_j^h\}$ is a (usually overlapping) partitioning of indices to be chosen. This final requirement is needed to ensure that the resulting coarse matrix is still sparse.

As shown in [9], the minimization problem (3)–(4) has a unique solution. The vectors $\bar{\Phi}_j$ can be found by solving local systems corresponding to subblocks of $\bar{A}$, as well as a global system with a matrix of the same size as $\bar{A}$. We assume that the subblocks can be solved efficiently and concentrate on finding efficient ways to solve the global system, which we call the *Lagrange multiplier system*.

The energy minimizing coarse space construction was proposed and studied in the (related but different) context of multigrid methods in [9, 10, 11]. The construction requires the solution of a Lagrange multiplier system of the same size as the original problem, but with a special structure. In Section 6 we extend ideas from [11], to propose an efficient preconditioner for this system in the domain decomposition case. The main point in [11] is that, because of the special structure of the Lagrange multiplier system, it suffices to use a one-level preconditioner to solve it efficiently. Since the matrix of the system is dense, constructing and applying the preconditioner is not straightforward. We demonstrate in Section 6 how to obtain an efficient and scalable implementation. The numerical results in Section 7 show that the energy minimizing coarse space does indeed give better results than other

coarse space constructions in all test cases. More importantly the results show that the preconditioner for the Lagrange multiplier system is completely robust for all the examples considered, which means the coarse space construction is robust and its cost grows with the problem size only at the same rate as the construction of other coarse spaces. For related work on coarse space constructions based on energy minimization see [12].

There are many approaches for the construction of coarse spaces, or equivalently interpolation operators, that have been studied in the context of multigrid and domain decomposition methods. Carefully chosen coarse grids and interpolation operators form the basis of the success of algebraic multigrid methods [13, 14, 15, 16, 17, 18, 19, 20, 21, 22]. In the context of domain decomposition methods, the multiscale finite element approach, originally introduced as a numerical homogenization tool [23, 24], was used for the construction of robust coarse spaces in [4]. Other related coarse spaces are the partition of unity coarse spaces of [3, 25, 26] and the coarse spaces described in [27]. In [28, 5] the ideas from aggregation type multigrid methods [14, 18] are applied in the context of overlapping domain decomposition methods.

For expositional purposes, we consider only the scalar elliptic equation (1) with continuous piecewise linear finite elements and homogeneous Dirichlet boundary conditions, and set $\mathcal{A}(x) = \alpha(x)I$ in the numerical tests. However, the method itself is much more generally applicable. The energy minimization approach can also be used to set up coarse spaces for linear elasticity problems [10] and edge-element discretizations of curl-curl type problems [29]. Higher-order finite elements or mixed boundary conditions pose no additional difficulties. It is even possible to construct efficient preconditioners for anisotropic problems with this approach provided the overlapping sets ω_j^h are chosen appropriately (see [9]). The numerical results in this paper show in fact very clearly that (even in the isotropic case) when the coefficients are highly varying, in a complicated (possibly random) way, the crucial ingredient for a robust and efficient two-level Schwarz method is a careful choice of the sets ω_j^h. Some ideas on how to choose these sets, inspired by aggregation type algebraic multigrid techniques [14] are provided in [5], but this requires still further investigation. These techniques are related to methods for choosing coarse grid points in other algebraic multigrid approaches [21, 15, 19, 22].

Note that any techniques to find coefficient robust coarse spaces for multigrid or domain decomposition methods are interesting in the context of numerical homogenization and upscaling as well [23, 30].

In what follows, we use the notation $x \lesssim y$ to stand for $x \leq Cy$ for some constant C. It is standard to assume that the constant C is independent of the mesh size and the parameters of the method, such as the number of subdomains, the size of the subdomains and the size of the overlap. Here, in addition we also assume that C is independent of the coefficients $\mathcal{A}$ of the PDE (1). The notation $x \sim y$ means $x \lesssim y$ and $y \lesssim x$.

2. Model Problem

Let Ω be a bounded, open, polygonal (polyhedral) domain in $\mathbb{R}^2$ ($\mathbb{R}^3$). As our model problem we consider the scalar elliptic equation

$$-\nabla \cdot (\mathcal{A}\nabla u) = f, \quad \text{on } \Omega, \tag{5}$$

with homogeneous Dirichlet boundary conditions

$$u = 0, \quad \text{on } \partial\Omega. \tag{6}$$

The unknown function u, the right hand side f and the coefficient tensor $\mathcal{A}$ are functions defined on Ω. We assume that $\mathcal{A}$ is uniformly bounded and symmetric positive definite, but it may otherwise be highly varying throughout Ω. The model problem can be formulated in weak form as finding $u \in H_0^1(\Omega)$ such that

$$a(u, v) = (f, v), \tag{7}$$

for all test functions $v \in H_0^1(\Omega)$, where $H_0^1(\Omega)$ is the usual Sobolev space on Ω. The bilinear form $a(\cdot, \cdot)$ and the inner product $(\cdot, \cdot)$ are defined as

$$a(u, v) = \int_\Omega \mathcal{A}\nabla u \cdot \nabla v \, dx, \qquad (f, v) = \int_\Omega f \, v \, dx. \tag{8}$$

The continuous problem (7) is discretized using a standard finite element approximation with continuous piecewise linear functions. Let $\mathcal{T}^h$ be a family of conforming meshes of triangles (tetrahedra), which are shape-regular as the mesh diameter h goes to zero. We denote by Ω^h and $\bar{\Omega}^h$ the sets of nodes of $\mathcal{T}^h$ contained in Ω and $\bar{\Omega} = \Omega \cup \partial\Omega$, respectively. Note that throughout we will use "barred" symbols to distinguish between objects related to Ω and $\bar{\Omega}$ (i.e., including the boundary). The space of continuous functions on $\bar{\Omega}$ which are piecewise linear with respect to $\mathcal{T}^h$, is denoted by $\bar{\mathcal{V}}^h$. Using the nodal hat functions ϕ_p, defined such that[†]

$$\phi_p \in \bar{\mathcal{V}}^h, \qquad \phi_p(q) = \delta_{pq}, \qquad p, q \in \bar{\Omega}^h, \tag{9}$$

where $\delta_{pq} = 1$ if $p = q$ and 0 otherwise, the finite element space can be written as

$$\bar{\mathcal{V}}^h = \text{span}\{\phi_p : p \in \bar{\Omega}^h\}. \tag{10}$$

The space of finite element functions that in addition satisfy the homogeneous Dirichlet boundary conditions and are therefore zero on $\partial\Omega$ is given by

$$\mathcal{V}^h = \text{span}\{\phi_p : p \in \Omega^h\}. \tag{11}$$

The Galerkin method replaces the space $H_0^1(\Omega)$ in (7) by the finite dimensional space $\mathcal{V}^h \subset H_0^1(\Omega)$. That is, we look for an approximate solution $u_h \in \mathcal{V}^h$ such that (7) holds for all $v_h \in \mathcal{V}^h$. Using the vector of coefficients $\mathbf{u} = \begin{bmatrix} \cdots & \mathrm{u}_q & \cdots \end{bmatrix}^T$, where $\mathrm{u}_q = u_h(q)$, $q \in \Omega^h$, we can write the approximate solution as $u_h = \sum_{q \in \Omega^h} \mathrm{u}_q \phi_q$. Using the test functions $v_h = \phi_p$, $p \in \Omega^h$, results in the system of linear equations

$$A\mathbf{u} = \mathbf{f}, \tag{12}$$

where the elements of the matrix and of the right hand side vector are given by

$$A_{pq} = a(\phi_q, \phi_p), \qquad \mathrm{f}_p = (f, \phi_p), \qquad p, q \in \Omega^h. \tag{13}$$

The matrix A is symmetric positive definite and typically large, but sparse.

If instead of $\mathcal{V}^h$ we use $\bar{\mathcal{V}}^h$ in the Galerkin method, we obtain a matrix $\bar{A}$ which does not explicitly incorporate boundary conditions (we have so-called natural boundary conditions). The matrix $\bar{A}$ is introduced since it is used in the coarse space construction. It is symmetric positive semidefinite and the kernel of $\bar{A}$ consists of the constant vectors, i.e., all row sums of $\bar{A}$ are zero. The matrix A is a submatrix of $\bar{A}$.

[†]We use the nodes of the finite element triangulation as indices for matrices and vectors. For each vector or matrix the set(s) of indices will be specified. In an actual implementation it may be necessary to choose a numbering for each set of indices.

3. Two-Level Schwarz Method

We now describe a two-level additive Schwarz method for the system of equations (12). Additive Schwarz methods are domain decomposition methods. Thus, let Ω_i be open subsets of the domain Ω. We assume that each subdomain Ω_i is the interior of a union of triangles of the mesh $\mathcal{T}^h$. We denote by Ω_i^h the nodes of $\mathcal{T}^h$ contained in Ω_i, i.e.,

$$\Omega_i^h = \{p \in \Omega^h : p \in \Omega_i\}. \tag{14}$$

For each subdomain Ω_i, we specify a subproblem with homogeneous Dirichlet boundary conditions on $\partial\Omega_i$. We define a restriction matrix R_i as

$$[R_i]_{pq} = \delta_{pq}, \qquad p \in \Omega_i^h, \qquad q \in \Omega^h. \tag{15}$$

The system matrix of the homogeneous Dirichlet problem on Ω_i is then given by the Galerkin product

$$A_i = R_i A R_i^T. \tag{16}$$

For a restriction matrix of the form (15), A_i is a submatrix of A. Note, however, that in general R_i could also be chosen differently. The (one-level) overlapping additive Schwarz preconditioner is given by

$$B = \sum\nolimits_i R_i^T A_i^{-1} R_i \tag{17}$$

Applying B to a vector involves solving subsystems with sparse, symmetric positive definite matrices A_i. The subproblems can be solved in parallel. Note that B is completely determined by the matrix A and the index sets Ω_i^h.

This one-level preconditioner is in general not scalable since the number of iterations depends on the number of subdomains. This problem can be overcome by solving an approximation of the original system in a suitable coarse space [1, 2]. Let Φ be a matrix whose columns span the coarse space. The coarse approximation of A is formed using the Galerkin product

$$A_H = \Phi^T A \Phi. \tag{18}$$

The two-level additive Schwarz preconditioner $\tilde{B}$ is then given by

$$\tilde{B} = \Phi A_H^{-1} \Phi^T + \sum\nolimits_i R_i^T A_i^{-1} R_i. \tag{19}$$

The crucial thing to make this preconditioner robust for highly varying coefficients is a careful choice of Φ. In the next section we present a method for constructing Φ that takes into account coefficient variation.

4. Coarse Space Construction

In this section we describe a practical and efficient coarse space construction based on local solves on overlapping domains. We construct a coarse space that is a subspace of the finite element space $\bar{\mathcal{V}}^h$. This means that (as usual) the coarse space basis functions can be represented by their values at the nodes of the finite element mesh. Let $\{\Psi_j \in \bar{\mathcal{V}}^h\}$ be the set of coarse space basis functions. We denote the vector of nodal values of the j-th coarse space basis function Ψ_j by $\bar{\Phi}_j$. Consequently, the matrix $\bar{\Phi}$ with j-th column equal to $\bar{\Phi}_j$ represents an interpolation from the coarse space to the fine space.

The matrix $\bar{\Phi}^T$ represents a restriction from the fine space to the coarse space. Note that we use the terms fine and coarse space to refer to the spaces of finite element functions as well as the spaces of vectors representing them. Once $\bar{\Phi}$ has been constructed, the coarse space for the Dirichlet problem is obtained by dropping those basis functions that are non-zero on the boundary. The matrix Φ which interpolates from the Dirichlet coarse space to the Dirichlet fine space is a submatrix of $\bar{\Phi}$ with rows restricted to the nodes in the interior of the domain Ω and columns restricted to coarse basis functions which are zero on the boundary $\partial\Omega$. This can be written as

$$\Phi_{pj} = \bar{\Phi}_{pj}, \qquad p \in \Omega^h, \qquad j \in \{j : \bar{\Phi}_{pj} = 0 \text{ for all } p \in \partial\Omega\}. \tag{20}$$

The number of columns in Φ determines the dimension of the coarse space, and therefore the dimension of the coarse matrix A_H. The sparsity pattern of Φ determines the sparsity pattern of A_H. To obtain an efficient and practical method we need the dimension of A_H to be much smaller than the dimension of A, and we want A_H to have a similar sparsity pattern as A, i.e., a small, bounded number of nonzeros per row independent of the dimension of A or of A_H. Therefore, we prescribe the support ω_j for each of the coarse space basis functions Ψ_j, or equivalently a set of nodes $\omega_j^h = \omega_j \cap \bar{\Omega}^h$ such that the i-th element of $\bar{\Phi}_j$ can only be nonzero if $i \in \omega_j^h$. In matrix notation we write

$$\bar{\Phi}_j = \bar{R}_j^T \mathbf{q}_j, \tag{21}$$

where $\mathbf{q}_j$ is a vector containing the non-zero coefficients of the coarse basis function and $\bar{R}_j^T$ is an extension matrix from ω_j^h to $\bar{\Omega}^h$. The corresponding restriction matrix $\bar{R}_j$ from $\bar{\Omega}^h$ to ω_j^h is a zero-one matrix constructed in the same way as the restriction matrix R_i from $\bar{\Omega}^h$ to Ω_i^h for the subdomain problems (see (15) in Section 3).

To obtain a sparse coarse matrix A_H, we assume that each support ω_j overlaps with at most η other supports, where η is independent of the problem size and of the number of subdomains. This implies that each point of Ω is contained in at most $\eta+1$ supports ω_j and also that each node in $\bar{\Omega}^h$ is contained in at most $\eta + 1$ sets ω_j^h. The latter, together with the assumption that each node is contained in at least one set ω_j^h, implies, in matrix notation

$$I \leq \sum\nolimits_j \bar{R}_j^T \bar{R}_j \leq (\eta + 1)\, I, \tag{22}$$

where the inequalities are taken componentwise.

With these preliminaries out of the way we can describe the coarse space construction. We will show in the next section (Theorem 1) that the following construction solves the energy minimization problem (3)–(4).

We propose that $\mathbf{q}_j$ be found by solving a local system

$$A_j \mathbf{q}_j = \bar{\mathbf{g}}_j, \tag{23}$$

where $\bar{\mathbf{g}}_j$ is a "well-chosen" right hand side and $A_j = \bar{R}_j \bar{A} \bar{R}_j^T$, i.e., the system matrix of a local problem on ω_j. The matrix A_j is a principal submatrix of $\bar{A}$.

If we assume that the right hand sides $\bar{\mathbf{g}}_j$ are restrictions of one global vector $\bar{\mathbf{g}}$, i.e.,

$$\bar{\mathbf{g}}_j = \bar{R}_j \bar{\mathbf{g}}, \tag{24}$$

then, combining (21), (23) and (24), we see that

$$\bar{\Phi}_j = \bar{R}_j^T A_j^{-1} \bar{R}_j \bar{\mathbf{g}}. \tag{25}$$

It is important to make sure that the kernel of the matrix $\bar{A}$, which contains all constant vectors, can be represented by the coarse space. This is achieved by enforcing the partition of unity constraint $\sum_j \bar{\Phi}_j = \mathbf{1}$. We see that $\bar{\mathbf{g}}$ is uniquely defined as the solution of the system

$$\bar{B}\bar{\mathbf{g}} = \mathbf{1}. \tag{26}$$

where

$$\bar{B} = \sum\nolimits_j \bar{R}_j^T A_j^{-1} \bar{R}_j \,. \tag{27}$$

In other words to construct the coarse space, we need to solve a system with $\bar{B}$, the one-level overlapping additive Schwarz preconditioner for $\bar{A}$ with respect to the covering $\{\omega_j\}$, as the system matrix.

Given $\bar{\mathbf{g}}$, each column $\bar{\Phi}_j$ of the matrix $\bar{\Phi}$ can then be found by solving the local problem (23) on the interior nodes of the support ω_j.

In Section 6 we investigate how to solve the system (26) efficiently. In the next section we show that the above construction gives the unique solution to a constrained energy minimization problem and that this minimization problem arises naturally in the coefficient-explicit convergence theory for two-level Schwarz methods.

5. Energy Minimization

Consider the following constrained optimization problem for the vectors $\bar{\Phi}_j$, given $\bar{A}$ and the sets $\{\omega_j^h\}$:

$$\min \sum\nolimits_j \bar{\Phi}_j^T \bar{A} \bar{\Phi}_j, \qquad \text{such that} \qquad \sum\nolimits_j \bar{\Phi}_j = \mathbf{1}, \tag{28}$$

and such that the i-th entry of $\bar{\Phi}_j$ is non-zero only for $i \in \omega_j^h$. We show that the construction outlined in Section 4 is equivalent to solving the constrained minimization problem (28).

Theorem 1. *The set of vectors* $\{\bar{\Phi}_j\}$ *defined in* (25)–(27) *is the unique solution to the constrained minimization problem* (28)*.*

Proof. The optimization problem (28) has a quadratic objective function and linear constraints. Such problems can be solved by solving a linear system of equations. We define the Lagrangian

$$L = \frac{1}{2} \sum\nolimits_j \bar{\Phi}_j^T \bar{A} \bar{\Phi}_j - (\sum\nolimits_j \bar{\Phi}_j - \mathbf{1})^T \bar{\lambda}, \tag{29}$$

where we introduced a vector of Lagrange multipliers $\bar{\lambda}$. The factor $\frac{1}{2}$ is introduced for convenience. We can incorporate the sparsity constraints by substituting $\bar{\Phi}_j = \bar{R}_j^T \mathbf{q}_j$

$$L = \frac{1}{2} \sum\nolimits_j \mathbf{q}_j^T \bar{R}_j \bar{A} \bar{R}_j^T \mathbf{q}_j - (\sum\nolimits_j \bar{R}_j^T \mathbf{q}_j - \mathbf{1})^T \bar{\lambda} \tag{30}$$

To find the stationary points of the Lagrangian we set $\nabla_{\bar{\lambda}} L = 0$ and $\nabla_{\mathbf{q}_j} L = 0$, which results in the equations

$$\sum\nolimits_j \bar{R}_j^T \mathbf{q}_j - \mathbf{1} = 0 \quad \text{and} \quad \bar{R}_j \bar{A} \bar{R}_j^T \mathbf{q}_j - \bar{R}_j \bar{\lambda} = 0. \tag{31}$$

Solving the second equation and using the definition of A_j gives

$$\mathbf{q}_j = A_j^{-1} \bar{R}_j \bar{\lambda}. \tag{32}$$

Substituting this in the first equation we find

$$\sum_j \bar{R}_j^T A_j^{-1} \bar{R}_j \bar{\lambda} - \mathbf{1} = 0 \tag{33}$$

or $\bar{B}\bar{\lambda} = \mathbf{1}$. By comparing with (26)–(27), we see that the Lagrange multipliers $\bar{\lambda}$ correspond to the vector $\bar{\mathbf{g}}$ that provides the right hand sides (24) for the local solves (23). □

The energy minimization approach was introduced for the construction of coarse spaces for multigrid methods in [9, 10]. There, the approach was motivated by results for the one-dimensional problem. In the one-dimensional case the optimal coarse space can be worked out explicitly by solving local one-dimensional problems. The supports can be chosen such that the resulting method gives a direct solver for (12) and the corresponding basis functions are in fact the unique solution to the constrained minimization problem (28). This was the motivation in [9, 10] to use this minimization problem and the construction in the previous section also in the two and three dimensional case.

We now present a new way of motivating the energy minimization approach by showing that it follows naturally from convergence theory for two-level Schwarz methods. For the theory we assume that the triangulation $\mathcal{T}^h$ is quasi-uniform and that the eigenvalues of $\mathcal{A}(x)$ satisfy $\lambda_i(\mathcal{A}(x)) \sim \alpha(x)$ for some scalar function $\alpha : \Omega \to [1, \infty)$ that is piecewise constant with respect to $\mathcal{T}^h$. This means we only consider the (nearly) isotropic case here and we possibly need to scale our system, but we still allow $\mathcal{A}$ to be highly variable throughout the domain Ω.

It is well known that for the preconditioned conjugate gradient method applied to (12) with preconditioner P, the number of iterations guaranteeing a required accuracy can be bounded in terms of the square root of the condition number $\kappa(PA) := \lambda_{\max}(PA)/\lambda_{\min}(PA)$ of the preconditioned matrix PA.

A typical two-level Schwarz method uses a discretization of (5) on a coarse mesh $\mathcal{T}^H$ with mesh width H (see Section 7), and extends the coarse elements $K \in \mathcal{T}^H$ to obtain subdomains with overlap width δ. Applying the standard theory for two-level Schwarz methods [1], but taking into account the coefficient variation gives the following bound (see e.g. [4])

$$\kappa(\tilde{B}A) \lesssim \max_{K \in \mathcal{T}^H} \sup_{x,y \in \omega_K} \frac{\alpha(x)}{\alpha(y)} \left(1 + \frac{H}{\delta}\right), \tag{34}$$

where ω_K is the union of all elements in $\mathcal{T}^H$ having a non-empty intersection (i.e., sharing an edge or a vertex) with the coarse element K. This means that as long as the coefficient jumps are small (locally) and the overlap width δ is sufficiently large, standard two-level Schwarz methods are robust. A similar result can be found in [26]. For non-overlapping methods, robustness can usually be shown as long as the coefficients are constant or vary only slightly on each subdomain and provided the coarse mesh is aligned with jumps in the coefficients [2]. In [4] and [5] this standard theory was extended to the case of arbitrary coefficient variation. We summarize here the coefficient explicit theory taken from [5] (see also [7]).

Let $\{\Psi_j \in \bar{\mathcal{V}}^h\}$ be a set of functions spanning the coarse space. Denote the support and the diameter of Ψ_j by

$$\omega_j = \operatorname{supp} \Psi_j \qquad \text{and} \qquad H_j = \operatorname{diam} \omega_j \,, \tag{35}$$

respectively. We assume that the supports ω_j are shape-regular and form a finite cover of $\bar{\Omega}$. Furthermore, we assume that the overlap between a support ω_j and its neighbours is uniform, with $\delta_j > 0$ the width of the overlap region. The parameter δ_j is defined to be the smallest value such that each point in ω_j at a distance less than δ_j from the boundary, is also contained in some ω_k, $k \neq j$. For more technical details see [4, 5]. The following theorem is [5, Theorem 3.8].

Theorem 2. *Assume that*

- $\sum_j \Psi_j(x) = 1$, $x \in \bar{\Omega}$,
- $\|\Psi_j\|_{L_\infty(\Omega)} \lesssim 1$, *and*
- *each* Ω_i *is large enough such that each* ω_j *is contained in at least one* Ω_i.

Then the condition number of the preconditioned matrix can be bounded as follows

$$\kappa(\tilde{B}A) \lesssim \gamma(\alpha)\left(1 + \max_j \frac{H_j}{\delta_j}\right). \tag{36}$$

where

$$\gamma(\alpha) = \max_j \ \delta_j^2 \|\alpha|\nabla\Psi_j|^2\|_{L_\infty(\Omega)} \tag{37}$$

is the coarse space robustness indicator.

The proof of Theorem 2 is similar to the standard proof for two-level additive Schwarz methods with general partition of unity coarse spaces [1, §3.10], but working directly in the energy norm, rather than the H^1-seminorm. If the third assumption is violated, it is still possible to prove the robustness of the method at the expense of some other (technical) assumptions on the basis functions $\{\Psi_j\}$. However, in that case two robustness indicators, one for the subdomain partitioning $\{\Omega_i\}$ and one for the coarse space span$\{\Psi_j\}$, are necessary [4].

If we assume that $\delta_j \sim H_j$, then Theorem 2 shows that we should strive to make $\gamma(\alpha)$ as small as possible to obtain a robust and efficient preconditioner. This can be achieved by a good choice of coarse basis functions. For $\mathcal{A}(x) = I$ and $\{\Psi_j\}$ piecewise linear with respect to some coarse mesh $\mathcal{T}^H$, we have $\gamma(\alpha) = \mathcal{O}(1)$. In this article we try to find basis functions with prescribed supports, that are optimal (in a certain sense) for more general coefficient distributions $\mathcal{A}$.

Since α is piecewise constant with respect to $\mathcal{T}^h$ and Ψ_j is piecewise linear, the coarse space robustness indicator can be written as

$$\gamma(\alpha) = \max_j \ \delta_j^2 \max_\tau \left[\alpha|\nabla\Psi_j|^2\right]_\tau, \tag{38}$$

where $\left[\alpha|\nabla\Psi_j|^2\right]_\tau$ denotes the constant value of $\alpha|\nabla\Psi_j|^2$ for the triangle $\tau \in \mathcal{T}^h$. If we define Y to be the matrix with entries

$$Y_{\tau j} = \left(\int_\tau \mathcal{A}|\nabla\Psi_j|^2\right)^{1/2},$$

then by the quasi-uniformity of $\mathcal{T}^h$ and the assumptions made on the eigenvalues of $\mathcal{A}(x)$ we have

$$\gamma(\alpha) \sim h^{-d} \max_j \ \delta_j^2 \max_\tau Y_{\tau j}^2 = h^{-d}\|Y\mathcal{D}\|_{\max}^2 =: \gamma^h(\mathcal{A}), \tag{39}$$

where the matrix max-norm is defined as $\|A\|_{\max} = \max_{i,j}|A_{ij}|$, and $\mathcal{D}$ is the diagonal matrix with entries $\mathcal{D}_{jj} = \delta_j$. The notation "$\sim$" absorbs the quasi-uniformity constants. Note that

$$Y^T Y = \bar{\Phi}^T \bar{A}\bar{\Phi}. \tag{40}$$

Recall that the columns $\bar{\Phi}_j$ of $\bar{\Phi}$ contain the nodal values representing the coarse space basis functions $\Psi_j \in \bar{\mathcal{V}}^h$ (see Section 4). To make the problem of minimizing $\gamma^h(\mathcal{A})$ more tractable, we replace the max-norm $\|\cdot\|_{\max}$ in (39) by the scaled Frobenius norm $h^{d/2}\|\cdot\|_F$ and define a new robustness indicator

$$\gamma_F^h(\mathcal{A}) = \|Y\mathcal{D}\|_F^2 = \mathrm{tr}(\mathcal{D}\bar{\Phi}^T\bar{A}\bar{\Phi}\mathcal{D}) = \sum\nolimits_j \delta_j^2 \bar{\Phi}_j^T \bar{A}\bar{\Phi}_j$$

where we have used the identity $\|A\|_F^2 = \mathrm{tr}(A^T A)$. Note that $\gamma^h(\mathcal{A}) \leq h^d \gamma_F^h(\mathcal{A})$ and therefore if $\gamma_F^h(\mathcal{A})$ is bounded independent of the coefficients than so is $\gamma^h(\mathcal{A})$. However, this leaves open the possibility that $\gamma^h(\mathcal{A})$ is bounded independent of the coefficients for some basis functions $\{\Psi_j\}$, but that no such functions exist for $\gamma_F^h(\mathcal{A})$. In practice, little difference has been found between the different matrix norms in the context of algebraic multigrid methods [10].

We see that provided the overlap width δ_j is of the same order for each support ω_j, the constrained optimization problem (28) corresponds exactly to the minimization of the Frobenius norm robustness indicator $\gamma_F^h(\mathcal{A})$, subject to the partition of unity constraint in Theorem 2. The Frobenius norm is scaled such that for $\mathcal{A}(x) = I$ and piecewise linear $\{\Psi_j\}$, we have $\gamma_F^h(\mathcal{A}) = \mathcal{O}(1)$ again.

6. Solving the Lagrange Multiplier System

The local problems (23) in the coarse space construction in Section 4 are similar to the ones which have to be solved anyway when applying the overlapping Schwarz preconditioner for A. We assume that this can be done efficiently. The hard part in the construction of the coarse space is solving the system (26) for the Lagrange multipliers. Since this system is of the same size as the original problem, we will solve it iteratively as well. Unfortunately, in general $\bar{B}$ is ill-conditioned and so we need to find a good preconditioner for $\bar{B}$.

In [9] it was proposed to use $\bar{A}$ as a cheap preconditioner for $\bar{B}$. However, since

$$\kappa(\bar{A}\bar{B}) = \kappa(\bar{B}\bar{A}), \tag{41}$$

it is clear that the performance of this preconditioner will be only as good as the performance of the one-level preconditioner $\bar{B}$ applied to the matrix $\bar{A}$. In other words, $\bar{A}$ is a cheap preconditioner for $\bar{B}$, but it does not scale well. Note that, since $\bar{A}$ is singular, the condition numbers in (41) should be interpreted as effective condition numbers. Furthermore, since $\bar{A}$ is singular, it is proposed in [9] to use $\bar{A} + \zeta I$, which introduces a parameter ζ that needs to be chosen. These shortcomings were already pointed out in [9].

In [11] it was observed that $\bar{B}$ is a "local" operator. The authors mainly considered constructing coarse spaces for multigrid methods, i.e., for a large number of coarse space basis functions with small supports. For this case, where the size of the supports is of the same order as the mesh size (i.e. $H \lesssim h$), they prove that the condition number $\kappa(\bar{B})$ is uniformly bounded with respect to h and they show numerically that using $\bar{D}$, the inverse of the diagonal of $\bar{B}$, as a preconditioner is sufficient to guarantee robustness also with respect to coefficient jumps. For two-level Schwarz methods, the size of the coarse space is relatively small (compared to the size of the original space). This means that the supports are typically large and the observations from [11] do not apply any longer. The numerical results in Section 7, in particular Tables VI and VIII, illustrate that in this case robustness with respect to the mesh size is lost and moreover, diagonal preconditioning is no longer robust with respect to coefficient variation. It was suggested in [11] that, because of its special structure, $\bar{B}$ could also be preconditioned by a one-level domain decomposition method, that is, without a coarse space. We now construct such a one-level preconditioner $\bar{C}$ and show that it can be implemented efficiently.

Just as the local matrices A_j are constructed based on $\bar{A}$ we can construct local matrices

$$B_j = \bar{R}_j \bar{B} \bar{R}_j^T, \tag{42}$$

based on $\bar{B}$. A one-level overlapping Schwarz preconditioner $\bar{C}$ for $\bar{B}$ is then given by

$$\bar{C} = \sum\nolimits_j \bar{R}_j^T B_j^{-1} \bar{R}_j. \tag{43}$$

Note that the one-level preconditioner $\bar{B}$ for $\bar{A}$ can be applied efficiently since the local matrices A_j, which have to be inverted, are sparse. In contrast, the matrix $\bar{B}$ and therefore the local matrices B_j are in general dense. As we show now, it is nevertheless still possible to apply $\bar{C}$ efficiently.

We illustrate this using an example with a support overlapping with only two others. Consider a support ω_j with 2 neighbours ω_k and ω_l. We have

$$\bar{R}_j \bar{R}_j^T = I_j, \qquad \bar{R}_j \bar{R}_k^T = \hat{I}_{jk} \neq 0, \qquad \bar{R}_j \bar{R}_l^T = \hat{I}_{jl} \neq 0\,,$$

where each row of $\hat{I}_{kj}$ and of $\hat{I}_{lj}$ is zero except for at most one entry which is one. All other products are zero matrices. Thus the matrix B_j can be written as

$$B_j = \bar{R}_j(\textstyle\sum_i \bar{R}_i^T A_i^{-1} \bar{R}_i)\bar{R}_j^T = A_j^{-1} + \hat{I}_{jk} A_k^{-1} \hat{I}_{kj} + \hat{I}_{jl} A_l^{-1} \hat{I}_{lj}.$$

It is well known that this is a dense matrix which means we do not want to explicitly construct it when implementing the method. Instead, the question is how we can efficiently apply

$$B_j^{-1} = (A_j^{-1} + \hat{I}_{jk} A_k^{-1} \hat{I}_{kj} + \hat{I}_{jl} A_l^{-1} \hat{I}_{lj})^{-1}. \tag{44}$$

We use the following form of the Sherman-Morisson-Woodbury formula [31, §2.1.3]

$$(X^{-1} + U\Sigma^{-1}V^T)^{-1} = X - XU(\Sigma + V^T XU)^{-1} V^T X. \tag{45}$$

Using the substitutions

$$X \leftarrow A_j, \quad U = V \leftarrow \begin{bmatrix} \hat{I}_{jk} & \hat{I}_{jl} \end{bmatrix}, \quad \Sigma \leftarrow \begin{bmatrix} A_k & \\ & A_l \end{bmatrix}, \tag{46}$$

we get

$$B_j^{-1} = A_j - A_j \begin{bmatrix} \hat{I}_{jk} & \hat{I}_{jl} \end{bmatrix} G_j^{-1} \begin{bmatrix} \hat{I}_{kj} \\ \hat{I}_{lj} \end{bmatrix} A_j,$$

where

$$G_j = \begin{bmatrix} A_k & \\ & A_l \end{bmatrix} + \begin{bmatrix} \hat{I}_{kj} \\ \hat{I}_{lj} \end{bmatrix} A_j \begin{bmatrix} \hat{I}_{jk} & \hat{I}_{jl} \end{bmatrix}.$$

This shows that B_j^{-1} can be applied to a vector by performing sparse matrix-vector multiplications and by solving a system involving the sparse matrix G_j.

In general let $k_1, \ldots, k_q$ be such that $\omega_{k_i} \cap \omega_j \neq \emptyset$, for all $i = 1, \ldots, q$ (with $q \leq \eta$). Define

$$\hat{R}_j^T = \begin{bmatrix} \hat{I}_{jk_1} & \cdots & \hat{I}_{jk_q} \end{bmatrix}^T, \tag{47}$$

where each row of $\hat{R}_j^T$ is again zero except for at most one entry which is one. This operator essentially duplicates each degree of freedom in the overlap of ω_j as many times as it appears in other supports ω_{k_i}, $i = 1, \ldots, q$, and extends the vector to the remainder of each of the ω_{k_i}. If we also define

$$G_j = \operatorname{diag}(A_{k_1}, \ldots, A_{k_q}) + \hat{R}_j^T A_j \hat{R}_j\,, \tag{48}$$

then, in the general case,

$$B_j^{-1} = A_j - A_j \hat{R}_j G_j^{-1} \hat{R}_j^T A_j. \tag{49}$$

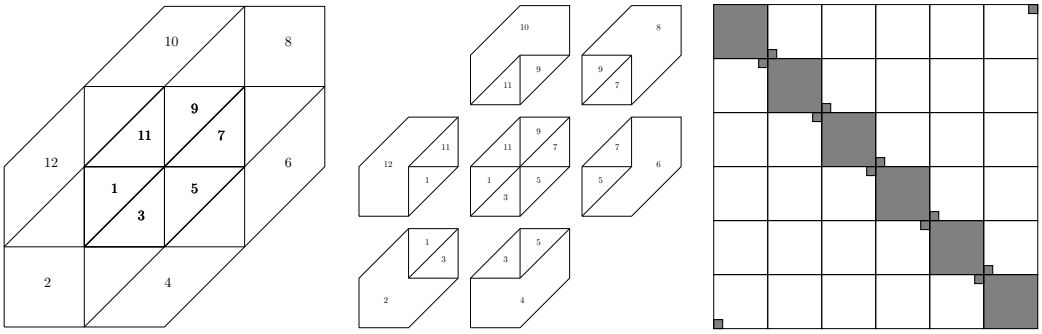

(a) Overlapping supports (b) Exploded view (c) Block matrix G_j

Figure 1. Example of overlap supports and block matrix structure of G_j.

Figure 1(a) shows a typical hexagonal support ω_j, as used in the numerical examples in Section 7, and the six supports $\omega_{k_1}, \ldots, \omega_{k_6}$ it overlaps with. All seven supports are shown separately in Figure 1(b). If the unknowns are grouped and numbered as illustrated in Figures 1(a) and 1(b), then the matrix G_j has the block structure shown in Figure 1(c). The shaded regions indicate the non-zero blocks of G_j which are themselves sparse.

To illustrate that (49) can indeed be applied efficiently to a vector, we estimate the complexity of this operation under some reasonable assumptions: (i) the cost of adding two sparse vectors of size m and the cost of multiplying an $m \times m$ sparse matrix by a vector are both proportional to m; (ii) the cost of solving a sparse $m \times m$ system is $\mathcal{S}(m) = \mathcal{O}(m^\beta)$ for some $\beta \geq 1$ (ideally $\beta = 1$). Now recall that each support overlaps with η other supports and assume further that the number of interior nodes in each support is m. This implies that $A_j \in \mathbb{R}^{m\times m}$, $\hat{R}_j \in \mathbb{R}^{m\times \eta m}$ and $G_j \in \mathbb{R}^{\eta m\times \eta m}$. Applying B_j^{-1} involves vector additions and matrix-vector multiplications with a cost proportional to ηm and a sparse solve with a cost of $\mathcal{S}(\eta m)$. Since the coupling between the $m \times m$ blocks of G_j is sparse (only through the shared overlap of two supports with a third, see Figure 1(c)), we can expect the cost of this sparse solve to be closer to $\eta\mathcal{S}(m)$, the cost of solving η decoupled $m \times m$ systems. Since we assumed η to be uniformly bounded, the cost of applying $\bar{C}$ will be $\mathcal{O}(m^\beta)$ which (asymptotically) is the same as the cost of applying $\bar{B}$. More specifically the cost of applying $\bar{C}$ will be roughly η times the cost of applying $\bar{B}$.

Remark 1. The number of iterations for solving (26) is determined by $\kappa(\bar{C}\bar{B})$. The numerical results in §7.2 indicate that this condition number is bounded by a constant that only depends on η and not on the size of the supports, on the size of the overlap, on the number of coarse basis functions or on the coefficients of the problem. The proof of this result is beyond the scope of the present paper and will be the subject of a forthcoming paper.

7. Numerical Results

As a typical domain for our numerical tests we take the unit square $\Omega = [0, 1]^2$. In all the examples we take $\mathcal{A}(x) = \alpha(x)I$, with $\alpha(x)$ piecewise constant with respect to $\mathcal{T}^h$. To make it easier to compare with

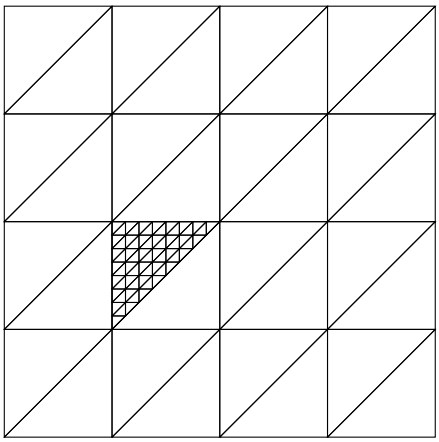

Figure 2. Coarse grid and fine mesh in one coarse element ($n_s = 4, n = 8$).

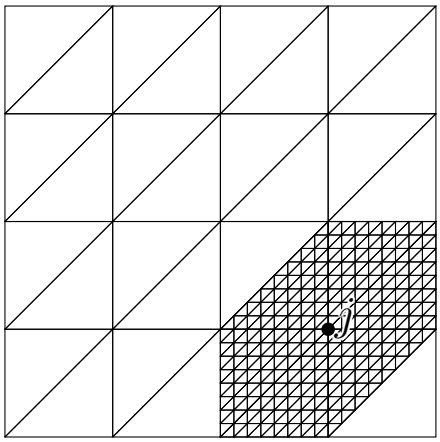

Figure 3. The support for the coarse basis function associated with coarse node j.

other methods, we assume that there is a coarse triangulation $\mathcal{T}^H$ which partitions the fine mesh into non-overlapping sets of triangles. Note, however, that the energy minimizing coarse space construction also works for more general partitionings. For our example, we consider a coarse triangulation obtained by dividing the unit square into $n_s \times n_s$ squares and subdividing each square along its main diagonal. Each coarse square is further subdivided into $n \times n$ small squares which are further subdivided into triangles. Each coarse triangle now contains $n \times n$ fine mesh triangles. For the numerical experiments we use r as a measure of the total number of unknowns where $n_s n = 2^r$. The coarse grid and one coarse triangle with subtriangles are shown in Figure 2 for $n_s = 4$, $n = 8$, $r = 5$.

The subdomains Ω_i for the one-level additive Schwarz method are chosen to consist of the union of all coarse triangles around a node of the coarse mesh. This means there are $(n_s + 1)^2$ subdomains, one for each coarse mesh node (see Figure 3 for a typical subdomain). This choice of subdomains gives $\delta_j \sim H_j$, which is often referred to as *generous overlap* [1].

For our numerical examples we take the supports ω_j of the coarse space basis functions to be the same as the subdomains Ω_i. This means that the coarse space basis functions (in all the studied methods) will have the same supports as the piecewise linear finite element basis functions with respect to the coarse mesh.

We compare several additive Schwarz preconditioners. The first method is the one-level additive Schwarz method based on the subdomains described above, i.e., without a coarse space (`NO`). For the two-level additive Schwarz methods we consider coarse spaces based on standard piecewise linear (`PL`) finite elements (with respect to the coarse mesh) and multiscale finite elements with linear (`LB`) and oscillatory (`OB`) boundary conditions [23, 24, 4], as well as the energy minimizing (`EM`) coarse space described above.

The basis functions in the case of the piecewise linear (`PL`) coarse space are the standard finite element "hat" functions which take the value one at one of the nodes of the coarse mesh and are zero at all other nodes of the coarse mesh.

The multiscale finite element basis functions are constructed by discrete α-harmonic extension of data on the boundary of each coarse triangle to the interior of the triangle. As for the piecewise linear

basis functions, we assume that the function Ψ_j associated with coarse node j satisfies $\Psi_j(j) = 1$ and is zero at all other coarse nodes. Different types of multiscale finite element bases arise from different choices of the boundary data. For the first type of multiscale finite element basis functions, we use linear boundary conditions (LB) on the boundary of each coarse triangle. The discrete α-harmonic extension of this boundary data to the interior of a coarse triangle K is obtained by (numerically) solving the homogeneous problem

$$-\nabla \cdot (\alpha \nabla u) = 0, \quad \text{on } K. \tag{50}$$

In practice this corresponds to solving a linear system with system matrix A_K, the principal subblock of A with row/column indices in $\Omega^h \cap K$. Note that $\Psi_j\big|_K \equiv 0$ for all coarse triangles K such that $j \notin \bar{K}$. Therefore we only need to solve a small number of local subproblems for each Ψ_j.

For the second type of multiscale finite element functions, the boundary data is obtained by first (numerically) solving a one-dimensional problem on each edge E of the coarse mesh. In [23] these are referred to as *oscillatory boundary conditions* (OB). The one-dimensional problem is the restriction of the homogeneous two-dimensional problem (50) to an edge E of the coarse mesh. Since the coefficient function α is a piecewise constant function with respect to $\mathcal{T}^h$, we have to define a value for α on E. Several choices are possible. Here, as suggested in [4], we assign to each fine mesh edge the maximum of the coefficient values of the triangles containing that fine mesh edge, since this choice guarantees that the basis functions are flat where α is large. This is different to the standard choice for discretization where typically a harmonic average is used (e.g., motivated by Darcy's Law). The resulting coefficient function on each coarse mesh edge E is piecewise constant with respect to the fine mesh edges. Again, we only have to solve these one-dimensional problems for a particular basis function Ψ_j if $j \in \bar{E}$. The discrete α-harmonic extension of the boundary data to the interior K of each coarse triangle is obtained, as above, by (numerically) solving (50) subject to the obtained oscillatory boundary conditions. For details see [4].

The construction of the multiscale finite element functions has many similarities with the construction based on energy minimization described in the previous sections. In both cases, local problems are solved to find the coarse space basis functions. The main differences are that for the multiscale finite elements we need a coarse mesh and we have to specify artificial boundary conditions on the triangles of the coarse mesh. For the energy minimizing basis it suffices to specify the overlapping supports of the basis functions. Note that if the supports are the same in both cases and if the boundary data in the multiscale finite element case are chosen to coincide with the values of the energy minimizing basis, then both methods will produce identical coarse spaces. However, since the optimal boundary data are not known a priori, the coarse spaces will normally be different. Intuitively, the energy minimizing coarse space can be expected to be more robust, since the multiscale finite element construction only takes into account coefficients near coarse edges when setting up the boundary conditions on coarse elements. Of course, the energy minimization construction is more expensive, since the local problems are slightly larger. Furthermore, we have to solve a global system (26) for the Lagrange multipliers to satisfy the partition of unity constraint, which is satisfied by construction in the case of the multiscale finite element coarse spaces. However, as explained in Section 6, using the preconditioner proposed here, this additional global solve has a cost that is of the same order as the remainder of the additive Schwarz algorithm. Thus asymptotically the cost of applying the two-level preconditioner $\tilde{B}$ grows at the same rate with respect to the number of unknowns for all four coarse spaces (PL, LB, OB, EM). Of course in absolute terms the energy minimizing coarse space is the most costly. However, we will see below that in some cases this is easily outweighed by the increased robustness of the preconditioner and thus the reduced number of CG iterations.

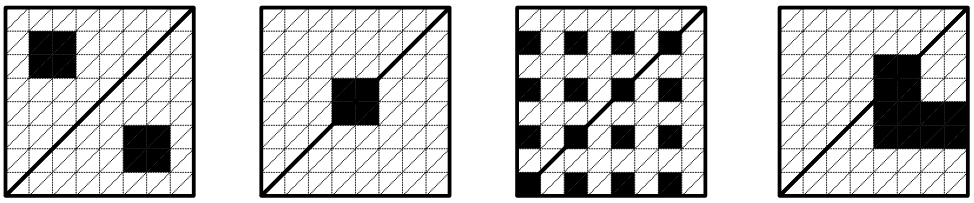

(a) Example 1 (b) Example 2 (c) Example 3 (d) Example 4

Figure 4. Coefficient patches for the binary media examples. Each patch is repeated n_s times in each direction. White indicates $\alpha = 1$ and black indicates $\alpha = \hat{\alpha}$.

For structured grids, robust coarse spaces can also be constructed based on the operator dependent prolongation used in certain geometric multigrid methods such as BoxMG [30]. Our construction is more costly (although again the cost is assymptotically the same), but it allows for arbitrary supports, not just uniform coarsenings of structured grids.

We consider several example problems similar to the ones in [4]. For the first set of so-called binary media examples, the coefficient α takes the value $\hat{\alpha} \geq 1$ in some parts of the domain (resolved by the fine mesh) and the value 1 in the remainder. The four binary media examples are illustrated in Figure 4 with white indicating $\alpha = 1$ and black indicating $\alpha = \hat{\alpha}$. Each figure shows the $n \times n$ small squares making up a representative square, which is repeated n_s times in each direction. We use the side lengths $H = 1/n_s$ and $h = H/n$ of the coarse and fine mesh squares as mesh parameters.

Example 1. For the first example, each coarse triangle has a high coefficient region in its interior. The regions are squares of size $H/4$, located at a distance $H/8$ from the horizontal and vertical edges of the coarse triangles.

Example 2. For the second example there is a high coefficient region on top of the middle of each diagonal edge shared by two coarse triangles. The regions are squares of size $H/4$, located at a distance $3H/8$ from the horizontal and vertical edges of the coarse triangles.

Example 3. The third example is an example of a fine scale binary medium. The high coefficient regions are regularly spaced squares of size h with corners $(2ih, 2jh)$ and $(2ih + h, 2jh + h)$, $(i, j) \in \mathbb{Z}^2$.

Example 4. The fourth example has an L-shaped region with high coefficients (see Figure 4(d)).

All the examples have high coefficient regions. The coefficient explicit convergence theory (see Section 5) shows that for a robust method we need basis functions with low energy. This means that the basis functions should be flat in the high coefficient regions. We therefore expect that the standard piecewise linear coarse space will not perform well as $\hat{\alpha}$ increases. Since Example 1 has no high coefficients on or near the coarse element boundaries, the multiscale finite element coarse space with linear boundary conditions will be robust with respect to coefficient variation. The resulting coarse basis functions are flat in the high coefficient region and they behave like the piecewise linear hat functions near the edges of the coarse mesh. On the other hand, Examples 2, 3 and 4 have high coefficient regions on the boundaries of the coarse elements and therefore linear boundary conditions will not be sufficient for robustness. Because of the linear boundary conditions, the solution has to go from high to low values within a region of high coefficients and therefore it is impossible for the α-harmonic extensions to be flat there. On the other hand, the oscillatory boundary conditions obtained by solving one-dimensional problems on the coarse edges, allow low energy basis functions to be

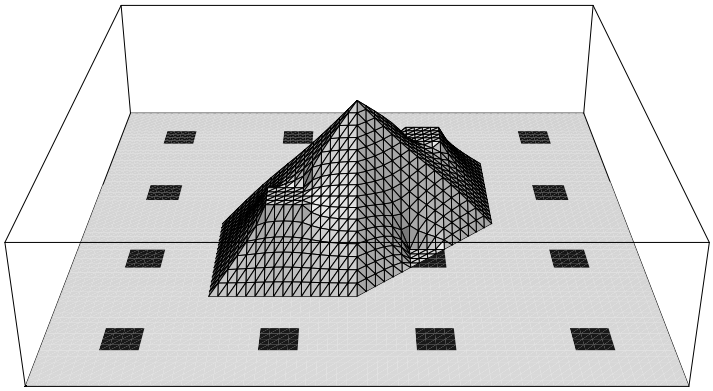

Figure 5. Typical energy minimizing coarse space basis function for Example 2.

found for Examples 2 and 3. Figure 5 shows a typical energy minimizing coarse space basis function for Example 2. Example 4 aims to illustrate that, for the same set of supports, it is possible to have a coefficient function for which the multiscale finite element functions do not work very well and the energy minimizing basis functions do. Since the area where the coefficient is high stretches across a coarse triangle and the (linear or oscillatory) boundary conditions on opposite edges of the coarse element specify both high and low values for this area, the multiscale finite element basis function cannot be flat everywhere. It will therefore have a high energy, which blows up as $\hat{\alpha} \to \infty$. This means that, in this case, the multiscale finite element functions do not provide a coarse space that is robust with respect to the coefficients. Figures 6(a) and 6(b) show surface and contour plots for some of the coarse basis functions for Example 4, both for the multiscale finite elements with oscillatory boundary conditions (`OB`) and the energy minimizing coarse basis functions (`EM`). We see that the basis functions are flat in regions with high coefficients for the `EM` case, but not for the `OB` case.

Finally we consider random media examples. Let $Z(p)$, $p \in \mathbb{R}^2$ be a Gaussian random field with mean $\mu(p) = E(Z(p))$ and covariance function $\Sigma(p,q) = E((Z(p) - \mu(p))(Z(q) - \mu(q)))$. We use the software package Gaussian [32] to construct a Gaussian random field $Z(p)$ on a grid with $n_s n \times n_s n$ squares. We consider homogeneous, stationary, isotropic Gaussian random fields with mean $\mu(p) = \mu = 0$ and an exponential covariance function of the form (see, e.g., eq. (2.2) in [33])

$$\Sigma(p,q) = \sigma^2 \exp(-|p-q|/\lambda), \tag{51}$$

where the parameters are the variance σ^2 and the correlation length λ. The coefficients are then taken to be the exponential of this field to obtain a lognormal random field. This gives a coefficient function with a certain smoothness determined by the correlation length λ. For each random field realization we calculate $\breve{\alpha} = \sup_{x,y\in\Omega} \alpha(x)/\alpha(y)$. Table I shows the minimum, median and maximum of $\breve{\alpha}$ for the 100 realizations used in the experiments. We also consider random binary media examples obtained by taking Gaussian random fields with the same covariance structure and by replacing the 50% lowest

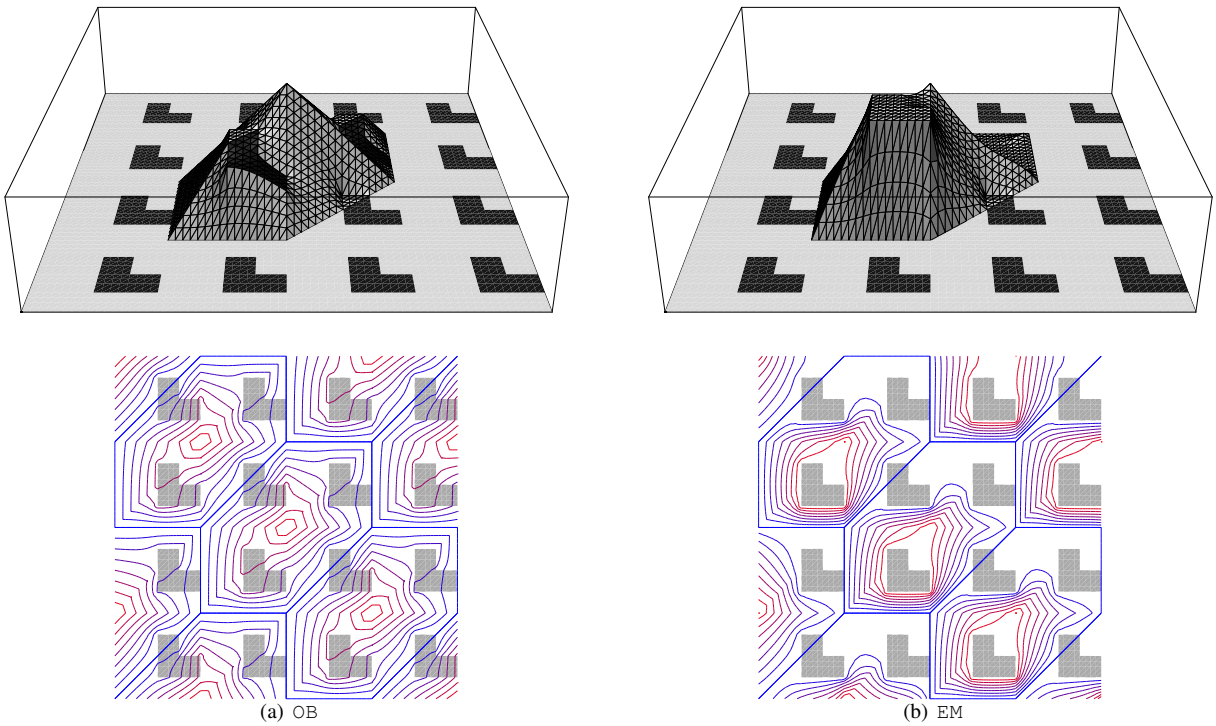

(a) OB (b) EM

Figure 6. Typical coarse space basis functions for Example 4 ($n_s = 4, n = 8, \hat{\alpha} = 10^6$) with a clearly visible non-zero gradient in the high coefficient region for OB and not for EM.

Table I. Statistics for $\check{\alpha} = \sup_{x,y\in\Omega} \alpha(x)/\alpha(y)$ for 100 log-normal random field realizations used in the experiments.

σ^2	$\check{\alpha}_{\text{min}}$	$\check{\alpha}_{\text{med}}$	$\check{\alpha}_{\text{max}}$
0	1.0e+00	1.0e+00	1.0e+00
2	4.8e+04	1.6e+05	6.2e+05
4	4.2e+06	2.2e+07	1.5e+08
6	1.3e+08	9.9e+08	1.1e+10
8	2.4e+09	2.4e+10	3.8e+11

and the 50% highest values of the field by 1 and $\hat{\alpha}$, respectively. A Gaussian random field and the corresponding clipped random field are shown in Figure 7.

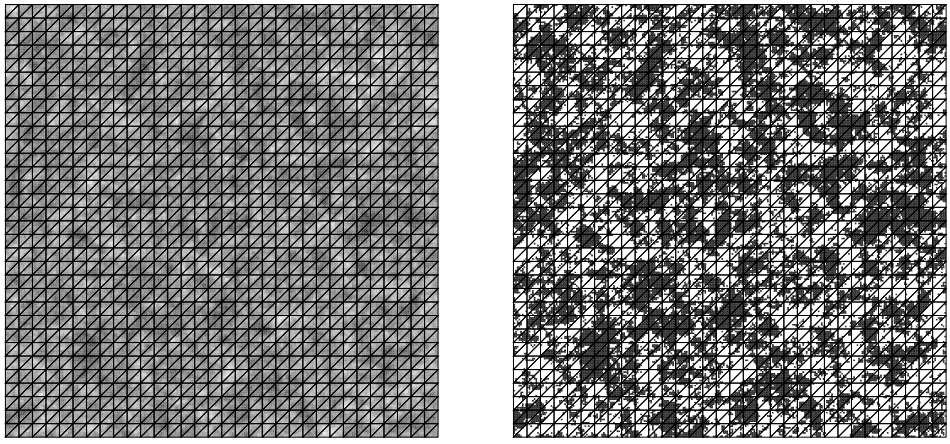

(a) Unclipped (b) Clipped

Figure 7. Gaussian random media examples ($n_s = 32$, $n = 8$, $\lambda = 4h$).

7.1. Scalability and Coefficient Robustness of the Two-Level Preconditioner

We first consider how the different preconditioners perform when using them in a conjugate gradient iteration to solve (12), with right hand side $\mathbf{f} = \mathbf{1}$ and initial guess $\mathbf{u} = 0$. Tables II–V show the number of CG iterations needed to attain a relative reduction in the residual norm by a factor of 10^{-6} for the different methods applied to the binary and random media examples. Tables II–IV show results for the binary media examples, Table V shows results for the random media.

For our choice of subdomains, the one-level method (NO) is robust with respect to coefficient variation in Examples 1–4 (see [4]) and this can be observed in Table II. However, Table III shows that, for the one-level method (NO), the number of iterations increases with the number of subdomains. Adding a standard piecewise linear finite element coarse space (PL) leads to a significant improvement for small values of $\hat{\alpha}$, but as soon as $\hat{\alpha}$ becomes larger the number of iterations increases. For Example 1 the multiscale finite element coarse space with linear boundary conditions (LB) is sufficient to obtain a fully robust method. For Examples 1–3, the multiscale finite element coarse space with oscillatory boundary conditions (OB) and the energy minimizing coarse space (EM) give equally good results. Example 4 illustrates that there are cases where the EM space is robust and the OB space is not.

Table V shows results for the clipped and unclipped random media examples. As σ^2 increases the coefficient variation increases (see Table I). For the clipped (binary) case this is similar to increasing $\hat{\alpha}$ in Examples 1–4. The performance of the multiscale finite element coarse space with oscillatory boundary conditions and the energy minimizing coarse space is very similar here. Both perform better than piecewise linear coarsening but unfortunately both are not fully robust. However, the energy minimizing framework is much more flexible than the multiscale finite element framework. In particular, it allows for more general supports. To achieve a fully robust method it is necessary to adapt not only the values, but also the supports of the coarse space basis functions to the coefficients. This can be done fully algebraically using ideas from Algebraic Multigrid. In [5], numerical experiments with adaptively chosen supports based on an aggregation method are performed. The aggregation method

Table II. Number of CG iterations to solve (12) with different preconditioners for different values of the coefficient $\hat{\alpha}$ ($n_s = 32, n = 8, r = 8$).

	Example 1					Example 2					Example 3					Example 4				
$\hat{\alpha}$	NO	PL	LB	OB	EM	NO	PL	LB	OB	EM	NO	PL	LB	OB	EM	NO	PL	LB	OB	EM
10^0	79	13	13	13	13	79	13	13	13	13	79	13	13	13	13	79	13	13	13	13
10^2	79	48	14	14	15	79	46	37	15	13	80	63	40	14	14	123	64	51	46	29
10^4	79	89	14	14	15	79	104	89	15	14	84	113	110	15	14	131	167	171	157	32
10^6	79	89	14	14	15	79	113	92	15	14	87	115	112	15	14	135	179	183	162	32

Table III. Number of CG iterations to solve (12) with different preconditioners for an increasing number n_s^2 of subdomains ($\hat{\alpha} = 10^6$).

			Example 1					Example 2					Example 3					Example 4				
n_s	n	r	NO	PL	LB	OB	EM	NO	PL	LB	OB	EM	NO	PL	LB	OB	EM	NO	PL	LB	OB	EM
4	8	5	14	18	13	13	13	14	18	17	13	13	15	19	18	13	13	25	29	27	26	21
8	8	6	21	23	14	14	15	21	31	25	15	14	23	33	31	15	14	34	48	49	48	26
16	8	7	42	48	15	15	15	41	58	48	15	14	45	59	57	15	14	74	89	95	88	29
32	8	8	79	89	14	14	15	79	113	92	15	14	87	115	112	15	14	135	179	183	162	32

in [5], without smoothing, results in supports with minimal overlap. For this case, the basis functions are the same as the ones found by the energy minimization construction (since there is only one set of functions that satisfies the partition of unity constraint). The numerical results in [5] show that with an adaptive choice of supports the method is fully robust even in the random coefficient case. For problems of the same size as the ones used here (but with slightly smaller supports) and for values of $\hat{\alpha}$ varying between 10 and 10^6 the experiments in [5, Table 1] show that the number of CG iterations in the clipped case varies only very slightly between 24 and 29 iterations. Similar results are reported for the unclipped case.

7.2. Scalability and Coefficient Robustness of the Coarse Space Construction

We now consider the construction of the coarse spaces based on energy minimization as described in Sections 4–6. The main part of this construction is the solution of the Lagrange multiplier system (26). We solve this system using a preconditioned conjugate gradient method. As preconditioners we consider a shifted version of the semidefinite system matrix $\bar{A} + \zeta I$, $\zeta = 10^{-3}$ (indicated as $\bar{A}$ in the tables) as proposed in [9], the inverse diagonal $\bar{D}$ of $\bar{B}$ as proposed in [11] and the one-level additive Schwarz preconditioner $\bar{C}$ described in Section 6. Another idea that we have tried, is to use a "localized" version of $\bar{A}$, that is

$$\bar{E} = \sum\nolimits_j \bar{R}_j^T A_j \bar{R}_j. \tag{52}$$

Table IV. Number of CG iterations to solve (12) with different preconditioners for growing subdomain problem size ($\hat{\alpha} = 10^6$).

			Example 1					Example 2					Example 3					Example 4				
n_s	n	r	NO	PL	LB	OB	EM	NO	PL	LB	OB	EM	NO	PL	LB	OB	EM	NO	PL	LB	OB	EM
8	4	5	28	37	35	35	26	21	22	14	14	14	24	32	31	14	13	40	48	49	46	24
8	8	6	21	23	14	14	15	21	31	25	15	14	23	33	31	15	14	34	48	49	48	26
8	16	7	22	24	15	15	15	22	32	26	15	14	24	32	31	15	15	37	49	51	48	26
8	32	8	23	24	15	15	15	23	32	27	16	15	25	34	31	15	15	39	50	52	49	27

Table V. Number of CG iterations (median for 100 realizations) to solve (12) for the random media examples for different values of $\hat{\alpha}$ and σ^2 ($n_s = 32$, $n = 8$, $r = 8$, $\lambda = 4h$).

Clipped

$\hat{\alpha}$	NO	PL	LB	OB	EM
10^0	79	13	13	13	13
10^2	169	35	31	28	26
10^4	216	59	52	42	40
10^6	265	93	80	64	60
10^8	322	138	120	94	89

Unclipped

σ^2	NO	PL	LB	OB	EM
0	79	13	13	13	13
2	174	31	28	24	23
4	238	48	41	31	30
6	322	68	56	37	37
8	428	93	74	44	45

This operator can be motivated by considering the case of minimal overlap between the supports, i.e., when the overlap consists of one layer of elements of the fine mesh. For this case no interior nodes are shared between the supports so that $\bar{R}_j \bar{R}_k^T = 0$ for $j \neq k$ and $\bar{R}_j \bar{R}_j^T = I_j$ and therefore $\bar{E}$ is the exact inverse of $\bar{B}$.

To see how each of the different preconditioners performs, we show in Tables VI–IX the number of CG iterations needed for solving the Lagrange multiplier system (26). Again the stopping tolerance is 10^{-6}. We observe that the performance of the one-level additive Schwarz preconditioner $\bar{C}$ is outstanding for all test cases. It gives a coarse space construction that is robust with respect to coefficient variation, number of subdomains and subproblem size. All the other preconditioners break down in one or more cases.

Figure 8 shows typical graphs for the processor time taken to construct the multiscale finite element coarse space with oscillatory boundary conditions (OB) and the energy minimizing coarse space (EM) for increasing subproblem size (see Tables IV and VIII, i.e. $n_s = 8$). In each case the subproblem size is $\mathcal{O}(n^2)$. The results are for Example 4. We see that the more robust EM construction is slightly more expensive, but it still scales similarly to the OB construction. Two lines indicating linear ($\beta = 1$) and quadratic ($\beta = 2$) behaviour are also shown. As explained in Section 6, a good sparse solver should give a computational complexity that scales close to linearly with subproblem size.

Table VI. Number of CG iterations to solve (26) for different values of the coefficient $\hat{\alpha}$ ($n_s = 32, n = 8, r = 8$).

	Example 1				Example 2				Example 3				Example 4			
$\hat{\alpha}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$
10^0	53	43	18	10	53	43	18	10	53	43	18	10	53	43	18	10
10^2	57	80	16	9	55	124	32	10	70	108	57	10	93	139	35	10
10^4	57	46	16	9	55	167	50	10	71	119	135	9	96	200^+	46	10
10^6	57	46	16	9	55	42	67	10	71	37	200^+	9	96	160	53	10

Table VII. Number of CG iterations to solve (26) for an increasing number n_s^2 of subdomains ($\hat{\alpha} = 10^6$).

			Example 1				Example 2				Example 3				Example 4			
n_s	n	r	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$
4	8	5	14	36	16	10	14	37	54	11	18	33	167	10	27	132	52	11
8	8	6	23	43	17	10	23	42	66	10	31	37	200^+	10	48	149	57	10
16	8	7	39	46	16	9	39	42	71	10	52	38	200^+	10	76	155	55	10
32	8	8	57	46	16	9	55	42	67	10	71	37	200^+	9	96	160	53	10

Table VIII. Number of CG iterations to solve (26) for growing subdomain problem size ($\hat{\alpha} = 10^6$).

			Example 1				Example 2				Example 3				Example 4			
n_s	n	r	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$
8	4	5	32	43	28	8	24	20	46	11	32	45	145	11	49	76	46	11
8	8	6	23	43	17	10	23	42	66	10	31	37	200^+	10	48	149	57	10
8	16	7	22	83	22	10	22	77	83	11	28	60	200^+	10	41	200^+	64	10
8	32	8	18	151	29	11	17	142	110	11	22	103	200^+	10	29	200^+	87	11

Table IX. Number of CG iterations (median for 3 realizations) to solve (26) for the random media examples for different values of $\hat{\alpha}$ and σ^2 ($n_s = 32, n = 8, r = 8, \lambda = 4h$).

Clipped

$\hat{\alpha}$	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$
10^0	53	43	18	10
10^2	189	101	50	14
10^4	200^+	174	88	15
10^6	200^+	200^+	145	16
10^8	200^+	200^+	200^+	16

Unclipped

σ^2	$\bar{A}$	$\bar{D}$	$\bar{E}$	$\bar{C}$
0	53	43	18	10
2	122	115	43	14
4	200^+	200^+	74	15
6	200^+	200^+	106	15
8	200^+	200^+	143	16

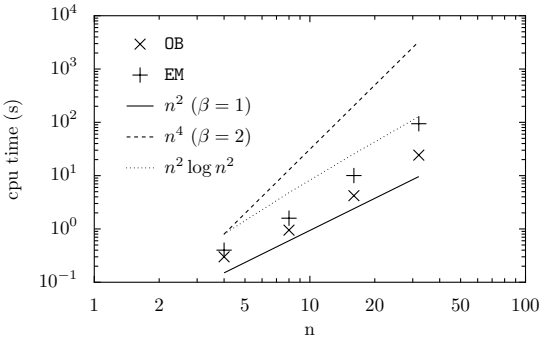

Figure 8. Timings for the OB and EM coarse space constructions for increasing subdomain size (Example 4, $\hat{\alpha} = 10^6$, $n_s = 8$).

8. Concluding Remarks

In this paper we have investigated the construction of coarse spaces for two-level Schwarz type domain decomposition methods. We used a coarse space construction based on energy minimization [9, 11]. The construction of the energy minimizing coarse space involves solving a Lagrange multiplier system of the same size as the original problem. To obtain an efficient solver, we propose a one-level overlapping additive Schwarz preconditioner for this system and we demonstrate how this preconditioner can be implemented effectively. Our numerical results show that, as opposed to previously studied preconditioners, the one-level overlapping additive Schwarz preconditioner results in a coarse space construction that is robust, efficient and scalable.

The numerical results show that the two-level method using the energy minimizing coarse space, performs well for many of the examples considered. However, for a fixed choice of supports for the coarse space basis functions, there are examples for which the method is not robust with respect to coefficient variation. Further investigations will show how to improve the robustness of the method by combining the energy minimizing coarse space construction with techniques for choosing the supports adaptively, depending on the coefficient function (see, e.g., [5]).

We have considered here the use of energy minimizing coarse spaces in the context of two-level Schwarz methods. However, these coarse spaces are also very interesting for numerical homogenization and upscaling.

REFERENCES

1. Toselli A, Widlund O. *Domain Decomposition Methods—Algorithms and Theory*, *Springer Series in Computational Mathematics*, vol. 34. Springer-Verlag: Berlin, 2005.
2. Chan TF, Mathew TP. Domain decomposition algorithms. *Acta Numerica, 1994*. Cambridge Univ. Press: Cambridge, 1994; 61–143.
3. Sarkis M. Schwarz preconditioners for elliptic problems with discontinuous coefficients using conforming and non-conforming elements. PhD Thesis, Courant Institute of Mathematical Sciences, New York University, Department of Computer Science 1994. TR-671.
4. Graham IG, Lechner PO, Scheichl R. Domain decomposition for multiscale PDEs. *Numer. Math.* 2007; **106**(4):589–626.
5. Scheichl R, Vainikko E. Additive Schwarz with aggregation-based coarsening for elliptic problems with highly variable coefficients. *Computing* 2007; **80**(4):319–343.

6. Graham IG, Scheichl R. Robust domain decomposition algorithms for multiscale PDEs. *Numer. Methods Partial Differential Equations* 2007; **23**(4):859–878.
7. Graham IG, Scheichl R. Coefficient-explicit condition number bounds for overlapping additive Schwarz. *Domain Decomposition Methods in Science and Engineering XVII*, *Lect. Notes Comput. Sci. Eng.*, vol. 60. Springer: Berlin, 2008.
8. Xu J. Iterative methods by space decomposition and subspace correction. *SIAM Rev.* 1992; **34**(4):581–613.
9. Wan WL, Chan TF, Smith B. An energy-minimizing interpolation for robust multigrid methods. *SIAM J. Sci. Comput.* 1999; **21**(4):1632–1649.
10. Mandel J, Brezina M, Vaněk P. Energy optimization of algebraic multigrid bases. *Computing* 1999; **62**(3):205–228.
11. Xu J, Zikatanov L. On an energy minimizing basis for algebraic multigrid methods. *Comput. Vis. Sci.* 2004; **7**(3-4):121–127.
12. Brannick J, Zikatanov L. Algebraic multigrid methods based on compatible relaxation and energy minimization. *Domain Decomposition Methods in Science and Engineering XVI*, *Lect. Notes Comput. Sci. Eng.*, vol. 55. Springer: Berlin, 2007; 15–26.
13. Ruge JW, Stüben K. Algebraic multigrid. *Multigrid methods*, *Frontiers Appl. Math.*, vol. 3. SIAM: Philadelphia, 1987; 73–130.
14. Vaněk P, Mandel J, Brezina M. Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems. *Computing* 1996; **56**(3):179–196. International GAMM-Workshop on Multi-level Methods (Meisdorf, 1994).
15. Brandt A. General highly accurate algebraic coarsening. *Electron. Trans. Numer. Anal.* 2000; **10**:1–20. Multilevel methods (Copper Mountain, CO, 1999).
16. Brezina M, Cleary AJ, Falgout RD, Henson VE, Jones JE, Manteuffel TA, McCormick SF, Ruge JW. Algebraic multigrid based on element interpolation (AMGe). *SIAM J. Sci. Comput.* 2000; **22**(5):1570–1592.
17. Jones JE, Vassilevski PS. AMGe based on element agglomeration. *SIAM J. Sci. Comput.* 2001; **23**(1):109–133.
18. Vaněk P, Brezina M, Mandel J. Convergence of algebraic multigrid based on smoothed aggregation. *Numer. Math.* 2001; **88**(3):559–579.
19. Livne OE. Coarsening by compatible relaxation. *Numer. Linear Algebra Appl.* 2004; **11**(2-3):205–227.
20. Brezina M, Falgout R, MacLachlan S, Manteuffel T, McCormick S, Ruge J. Adaptive smoothed aggregation (αSA) multigrid. *SIAM Rev.* 2005; **47**(2):317–346.
21. Brannick J, Brezina M, MacLachlan S, Manteuffel T, McCormick S, Ruge J. An energy-based AMG coarsening strategy. *Numer. Linear Algebra Appl.* 2006; **13**(2-3):133–148.
22. MacLachlan S, Saad Y. A greedy strategy for coarse-grid selection. *SIAM J. Sci. Comput.* 2007; **29**(5):1825–1853.
23. Hou TY, Wu XH. A multiscale finite element method for elliptic problems in composite materials and porous media. *J. Comput. Phys.* 1997; **134**(1):169–189.
24. Hou TY, Wu XH, Cai Z. Convergence of a multiscale finite element method for elliptic problems with rapidly oscillating coefficients. *Math. Comp.* 1999; **68**(227):913–943.
25. Sarkis M. Partition of unity coarse spaces and Schwarz methods with harmonic overlap. *Recent Developments in Domain Decomposition Methods (Zürich, 2001)*, *Lect. Notes Comput. Sci. Eng.*, vol. 23. Springer: Berlin, 2002; 77–94.
26. Sarkis M. Partition of unity coarse spaces: enhanced versions, discontinuous coefficients and applications to elasticity. *Domain Decomposition Methods in Science and Engineering*. Natl. Auton. Univ. Mex., México, 2003; 149–158.
27. Dohrmann C, Klawonn A, Widlund O. A family of energy minimizing coarse spaces for overlapping Schwarz preconditioners. *Domain Decomposition Methods in Science and Engineering XVII*, *Lect. Notes Comput. Sci. Eng.*, vol. 60, Springer, 2008.
28. Brezina M, Vaněk P. A black-box iterative solver based on a two-level Schwarz method. *Computing* 1999; **63**(3):233–263.
29. Musy F, Nicolas L, Perrussel R. Compatible coarse nodal and edge elements through energy functionals. *SIAM J. Sci. Comput.* 2007; **29**(3):1315–1337.
30. MacLachlan SP, Moulton JD. Multilevel upscaling through variational coarsening. *Water Resour. Res* 2006; **42**(2).
31. Golub GH, Van Loan CF. *Matrix Computations*. Third edn., Johns Hopkins University Press: Baltimore, MD, 1996.
32. Kozintsev B. Gaussian user's manual. `http://www.math.umd.edu/~bnk/bak/SOURCE/manual.pdf` 1999.
33. Cliffe KA, Graham IG, Scheichl R, Stals L. Parallel computation of flow in heterogeneous media modelled by mixed finite elements. *J. Comput. Phys.* 2000; **164**(2):258–282.