

Burak Aksoylu · Ivan G. Graham · Hector Klie · Robert Scheichl

Towards a rigorously justified algebraic preconditioner for high-contrast diffusion problems

Dedicated to Prof. Dr. Wolfgang Hackbusch on the occasion of his 60th birthday.

Abstract In this paper we present a new preconditioner suitable for solving linear systems arising from finite element approximations of elliptic PDEs with high-contrast coefficients. The construction of the preconditioner consists of two phases. The first phase is an algebraic one which partitions the degrees of freedom into “high” and “low” permeability regions which may be of arbitrary geometry. This partition yields a corresponding blocking of the stiffness matrix and hence a formula for the action of its inverse involving the inverses of both the high permeability block and its Schur complement in the original matrix. The structure of the required sub-block inverses in the high contrast case is revealed by a singular perturbation analysis (with the contrast playing the role of a large parameter). This shows that for high enough contrast each of the sub-block inverses can be approximated well by solving only systems with constant coefficients. The second phase of the algorithm involves the approximation of these constant coefficient systems using multigrid methods. The result is a general method of algebraic character which (under suitable hypotheses) can be proved to be robust with respect to both the contrast and the mesh size. While a similar performance is also achieved in practice by algebraic multigrid (AMG) methods, this performance is still without theoretical justification. Since the first phase of our method is comparable to the process of identifying weak and strong connections in conventional algebraic multigrid algorithms, our theory provides to some extent a theoretical justification for these successful algebraic procedures. We demonstrate the advantageous properties of our preconditioner using ex-

periments on model problems. Our numerical experiments show that for sufficiently high contrast the performance of our new preconditioner is almost identical to that of the Ruge and Stüben AMG preconditioner, both in terms of iteration count and CPU-time.

Keywords Diffusion problem, high-contrast coefficients, finite element approximation, algebraic preconditioner, Schur complement, multigrid

Mathematics Subject Classification (2000) 15A12, 65F10, 65F35, 65N22, 65N55

Acknowledgements I.G. Graham and R. Scheichl would like to thank the Center for Subsurface Modeling, ICES, University of Texas at Austin and the RICAM institute, Austrian Academy of Sciences, Linz, for financial support. B. Aksoylu would like to thank H.R. Beyer of the Center for Computation and Technology at Louisiana State University for many enlightening discussions.

1 Introduction

Problems with high-contrast coefficients are ubiquitous in porous media flow applications; e.g., [20, 25, 24]. Consequently, development of efficient solvers for high-contrast heterogeneous media has been an active area of research, specifically in the setting of multiscale solvers [2, 10, 11, 22]. In this paper, we are particularly concerned with the convergence of a family of algebraic preconditioners that exploit the binary character of high-contrast coefficients related to those recently proposed by Aksoylu and Klie [3].

We consider preconditioners for piecewise linear finite element discretisations of boundary-value problems for the model elliptic problem

$$-\nabla \cdot (\alpha \nabla u) = f, \quad (1)$$

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

B. Aksoylu
Department of Mathematics and Center for Computation and Technology, Louisiana State University.

I.G. Graham · R. Scheichl
Department of Mathematical Sciences and Bath Institute for Complex Systems, University of Bath, Bath BA2 7AY, UK.
E-mail: I.G.Graham@bath.ac.uk

H. Klie
Center for Subsurface Modeling, Institute for Computational Science and Engineering, The University of Texas at Austin.

this kind arise in groundwater flow and oil reservoir simulation; see for example the comprehensive overviews [1, 17, 7, 15]. For the theoretical statements in this paper, we will assume for convenience that α in (1) is a scalar function. However all the same results hold when $\alpha(x)$ is replaced by a symmetric positive definite matrix $\mathcal{A}(x)$, with spectrum lying in the range $[C^{-1}\alpha(x), C\alpha(x)]$ where C is moderate in size, and where the scalar function $\alpha(x)$ has the properties which we assume below. The case when C is very large (the anisotropic case) presents additional difficulties and should be the subject of future analysis.

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 (1) in this space yields the linear system:

$$A\mathbf{u} = \mathbf{f}, \quad (2)$$

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 (2) which are robust to the heterogeneity as well as to the mesh width h .

In the literature there are many papers devoted to the efficient solution of this problem and provide a rigorous justification when discontinuities in α are simple interfaces which can be resolved by a coarse mesh (see e.g. [4, 13] and the references therein for papers on domain decomposition methods and [26] for results on multigrid methods).

Even if suitable coefficient-resolving coarse meshes are not available, good performance of Krylov-based methods can still be achieved by standard preconditioners when there is a small number of unresolved interfaces. This is because the preconditioning produces a highly clustered spectrum with correspondingly few near-zero eigenvalues [9, 10, 23].

For more general complicated heterogeneous high-contrast media, recent progress was made in [11] where a characterisation of domain decomposition methods which are robust with respect to both contrast and mesh parameters was presented. This analysis indicated explicitly how subdomains and coarse spaces should be designed in order to achieve robustness also with respect to extreme heterogeneities, even inside coarse mesh elements. This approach was further extended in [22] to give a justification of the robustness of smoothed aggregation type domain decomposition methods for problems of this type.

At the same time it is well-known that algebraic multigrid procedures also produce optimal robust solvers for such heterogeneous problems, but so far theoretical justification of this is lacking. In this paper we describe a preconditioner which involves both an algebraic phase (similar to that used in algebraic multigrid) coupled with an application of standard multigrid ([12]) and we prove its robustness and demonstrate this on a sequence of model problems.

The preconditioner which we shall describe is an enhancement of an original method proposed in [14] for solving pressure-saturation coupled systems and recently ap-

plied in [3] for the setting of highly heterogeneous media. In [3], the coupling in the pressure system was interpreted as the interaction of degrees of freedom with different physical properties (as explained later). Moreover, when the underlying physics is not fully captured algebraically by the block partitioning – especially in the case of complex geometry – a deflation strategy was employed to enhance the preconditioner.

To give some more details, the first algebraic phase of our family of preconditioners involves partitioning of the degrees of freedom (subsequently referred to as “DOFs”) into a set corresponding to a “high-permeability” region and a “low-permeability” region. DOFs that lie on the interface between the two regions are (always) included in the high-permeability region. Note that in the context of standard FE matrices and, for high enough contrast, this can easily be obtained by examining the diagonal entries of the matrix A , or by using a strong-connection criterion similar to that used in algebraic multigrid algorithms. Thus any vector $\mathbf{u} \in \mathbb{R}^n$ can be decomposed into $\mathbf{u} = (\mathbf{u}_H^T, \mathbf{u}_L^T)^T$ and the stiffness matrix A in (2) can be partitioned

$$A = \begin{bmatrix} A_{HH} & A_{HL} \\ A_{LH} & A_{LL} \end{bmatrix}. \quad (3)$$

After a little algebra, the exact inverse of A can be written:

$$A^{-1} = \begin{bmatrix} I_{HH} & -A_{HH}^{-1}A_{HL} \\ 0 & I_{LL} \end{bmatrix} \begin{bmatrix} A_{HH}^{-1} & 0 \\ 0 & S^{-1} \end{bmatrix} \begin{bmatrix} I_{HH} & 0 \\ -A_{LH}A_{HH}^{-1} & I_{LL} \end{bmatrix} \quad (4)$$

where $S = A_{LL} - A_{LH}A_{HH}^{-1}A_{HL}$ is the Schur complement of A_{HH} in A and I_{HH} and I_{LL} denote the identity matrices of the appropriate dimension.

A singular perturbation analysis can now be devised to explain the properties of the subblocks in (4). Arguments of this type were first used in the context of condition number analysis for additive Schwarz methods in [9, 10]. More recently this approach was refined to treat the more complicated problem of analysing multigrid preconditioners in [26]. Here we use the singular perturbation-type analysis in a different context.

Suppose for simplicity that Ω_H has coefficient $\alpha = \hat{\alpha} \gg 1$ and that $\Omega_L := \text{int}(\Omega \setminus \Omega_H)$ has coefficient $\alpha = 1$. (Note however, that our method and our analysis are not restricted to this piecewise constant model situation.) It is clear that

$$\hat{\alpha}^{-1}A_{HH} = \mathcal{N}_{HH} + \mathcal{O}(\hat{\alpha}^{-1}), \quad \text{as } \hat{\alpha} \rightarrow \infty, \quad (5)$$

where \mathcal{N}_{HH} is the matrix corresponding to the pure Neumann problem for the Laplace operator on Ω_H . This shows that (after scaling by $\hat{\alpha}^{-1}$), A_{HH} can be preconditioned robustly and efficiently by standard multilevel methods, such as geometric multigrid, with a performance independent of h and $\hat{\alpha}$. The possible benefits of using multigrid as a preconditioner for the conjugate gradient method are well-known and were pointed out, for example, by W. Hackbusch in [12].

Moreover the analysis of A_{HH} as $\hat{\alpha} \rightarrow \infty$ has important implications for the behaviour of S . In §2, 3 we show that in this case

$$S = S(\infty) + \mathcal{O}(\hat{\alpha}^{-1}), \quad (6)$$

where $S(\infty)$ is a low rank perturbation of A_{LL} . The rank of the perturbation depends on the number of disconnected components in Ω_H . This special limiting form of S allows us to build a robust approximation of S^{-1} , for example combining solves with A_{LL} (again available robustly using standard multilevel methods) with the Sherman-Morrison-Woodbury formula.

There are a number of further approximations of (4) which can be envisaged. In fact, the simplest version of the Aksoylu-Klie preconditioner [3] is

$$B_{AK0} = \begin{bmatrix} A_{HH}^{-1} & 0 \\ 0 & I_{LL} \end{bmatrix} \begin{bmatrix} I_{HH} & 0 \\ -A_{LH}A_{HH}^{-1} & I_{LL} \end{bmatrix}. \quad (7)$$

As we show in §2, this preconditioner will perform reasonably provided the number of DOFs in Ω_L is not significantly large. To obtain better behaviour with respect to the number of DOFs in Ω_L a suitable modification to the Aksoylu-Klie preconditioner would be

$$B_{AK1} = \begin{bmatrix} A_{HH}^{-1} & 0 \\ 0 & S^{-1} \end{bmatrix} \begin{bmatrix} I_{HH} & 0 \\ -A_{LH}A_{HH}^{-1} & I_{LL} \end{bmatrix}. \quad (8)$$

As a simple consequence of (4) we have $\sigma(B_{AK1}A) = \{1\}$. However, a practical application of this preconditioner requires again robust and efficient approximations of A_{HH}^{-1} and S^{-1} , and so B_{AK1} is in fact nothing else but a nonsymmetric version of the preconditioner which we shall present below. We will only focus on the symmetric version in this paper.

The paper is structured as follows. In the next section we explain the basic idea for a simple model problem of a two scale medium with a simply connected high permeability region inside the domain. This leads to a suggested preconditioner which we show robust as $\hat{\alpha} \rightarrow \infty$. In §3 we extend the perturbation analysis to several high permeability regions and more general coefficients. In §4 we compare the performance of the proposed preconditioners numerically on some model problems. We also include performance comparisons with geometric and algebraic multigrid methods.

2 The one island case

2.1 Singular perturbation analysis

Let Ω be decomposed with respect to permeability value as

$$\Omega = \Omega_H \cup \Omega_L, \quad (9)$$

where Ω_H and Ω_L denote the high and low permeability regions respectively. Note that this is available algebraically, either via inspection of the diagonal entries in A or using the very common notion from AMG of strong and weak




Fig. 1 $\Omega = \Omega_H \cup \Omega_L$ where Ω_H and Ω_L are high and low permeable regions, respectively.

connections in A . Let Γ be the interface between Ω_H and Ω_L ; $\Gamma = \overline{\Omega}_H \cap \overline{\Omega}_L$.

We shall describe the basic idea by assuming first of all that Ω_H is connected and $\overline{\Omega}_H \cap \partial\Omega = \emptyset$, and that pure Dirichlet boundary conditions are enforced on all of $\partial\Omega$. (See Figure 1.) Moreover let $\alpha|_{\Omega_H} = \hat{\alpha} \gg 1$ and $\alpha|_{\Omega_L} = 1$. We will come back to the more general situation in the next section.

From the decomposition (9), we obtain a blocking for A as in (3) where only the block $A_{HH} = A_{HH}(\hat{\alpha})$ depends on $\hat{\alpha}$ and the Schur complement is

$$S(\hat{\alpha}) := A_{LL} - A_{LH}A_{HH}(\hat{\alpha})^{-1}A_{HL}. \quad (10)$$

To analyse the $\hat{\alpha}$ -robustness of preconditioners based on (4), we need to analyse the asymptotic behaviour of the block components $A_{HH}(\hat{\alpha})^{-1}$, $S(\hat{\alpha})^{-1}$ and $A_{LH}A_{HH}(\hat{\alpha})^{-1}$ as $\hat{\alpha} \rightarrow \infty$. This is the purpose of Lemma 1 below. To prepare for this, we further decompose the set of DOFs associated with $\overline{\Omega}_H$ into a set of interior DOFs associated with index I and boundary DOFs with index Γ . This leads to the following further block representation of

$$A_{HH}(\hat{\alpha}) = \begin{bmatrix} A_{II}(\hat{\alpha}) & A_{I\Gamma}(\hat{\alpha}) \\ A_{\Gamma I}(\hat{\alpha}) & A_{\Gamma\Gamma}(\hat{\alpha}) \end{bmatrix}. \quad (11)$$

The entries in the block $A_{\Gamma\Gamma}(\hat{\alpha})$ are assembled from contributions both from finite elements in Ω_H and Ω_L , i.e. $A_{\Gamma\Gamma}(\hat{\alpha}) = A_{\Gamma\Gamma}^{(H)}(\hat{\alpha}) + A_{\Gamma\Gamma}^{(L)}$ and so, inserting this into (11), we obtain

$$A_{HH}(\hat{\alpha}) = \hat{\alpha}\mathcal{N}_{HH} + \Delta, \quad \text{where} \quad \Delta = \begin{bmatrix} 0 & 0 \\ 0 & A_{\Gamma\Gamma}^{(L)} \end{bmatrix}, \quad (12)$$

and where \mathcal{N}_{HH} is the Neumann matrix on $\overline{\Omega}_H$, as described in (5). This is a symmetric positive semidefinite matrix with a simple zero eigenvalue and associated constant eigenvector. If n_H denotes the number of degrees of freedom in $\overline{\Omega}_H$, a suitable normalised eigenvector is the constant vector with entries $n_H^{-1/2}$, which we denote by \mathbf{e}_H . We further write in block form as $\mathbf{e}_H = (\mathbf{e}_I^T, \mathbf{e}_\Gamma^T)^T$.

Finally we note that the off-diagonal blocks in (3) (which are independent of $\hat{\alpha}$) have the decomposition:

$$A_{LH} = \begin{bmatrix} 0 & A_{LR} \end{bmatrix} = A_{HL}^T. \quad (13)$$

The following result describes the asymptotic behaviour of the sub-blocks in (4).

Lemma 1 *Let $\eta := \mathbf{e}_r^T A_{rr}^{(L)} \mathbf{e}_r = \mathbf{e}_h^T \Delta \mathbf{e}_h$. Then*

- (i) $A_{HH}(\hat{\alpha})^{-1} = \mathbf{e}_h \eta^{-1} \mathbf{e}_h^T + \mathcal{O}(\hat{\alpha}^{-1})$
- (ii) $S(\hat{\alpha}) = A_{LL} - (A_{LR} \mathbf{e}_r) \eta^{-1} (\mathbf{e}_r^T A_{rL}) + \mathcal{O}(\hat{\alpha}^{-1})$
- (iii) $A_{LH} A_{HH}(\hat{\alpha})^{-1} = (A_{LR} \mathbf{e}_r) \eta^{-1} \mathbf{e}_h^T + \mathcal{O}(\hat{\alpha}^{-1})$

Proof Since \mathcal{N}_{HH} is symmetric positive semidefinite we have the eigenvalue decomposition:

$$Z^T \mathcal{N}_{HH} Z = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_{n_H-1}, 0), \quad (14)$$

where $\{\lambda_i : i = 1, \dots, n_H\}$ is a non-increasing sequence of eigenvalues of \mathcal{N}_{HH} and Z is orthogonal. Because the eigenvector corresponding to the zero eigenvalue is constant, we may write $Z = [\tilde{Z} \mid \mathbf{e}_h]$ and so, using (12), we have

$$\begin{aligned} Z^T A_{HH} Z &= \begin{bmatrix} \hat{\alpha} \text{diag}(\lambda_1, \dots, \lambda_{n_H-1}) + \tilde{Z}^T \Delta \tilde{Z} & \tilde{Z}^T \Delta \mathbf{e}_h \\ \mathbf{e}_h^T \Delta \tilde{Z} & \mathbf{e}_h^T \Delta \mathbf{e}_h \end{bmatrix} \\ &=: \begin{bmatrix} \tilde{A}(\hat{\alpha}) & \tilde{\delta} \\ \tilde{\delta}^T & \eta \end{bmatrix}, \end{aligned} \quad (15)$$

where $\eta > 0$ (independent of $\hat{\alpha}$), since $\eta = \mathbf{e}_h^T A_{HH}(\hat{\alpha}) \mathbf{e}_h$ and $A_{HH}(\hat{\alpha})$ as a diagonal subblock of $A(\hat{\alpha})$ is SPD. To find the limiting form of $A_{HH}(\hat{\alpha})^{-1}$ note that

$$\begin{aligned} \tilde{A}(\hat{\alpha}) &= \hat{\alpha} \text{diag}(\lambda_1, \dots, \lambda_{n_H-1}) + \tilde{Z}^T \Delta \tilde{Z} \\ &= \hat{\alpha} \text{diag}(\lambda_1, \dots, \lambda_{n_H-1}) * \\ &\quad (I + \hat{\alpha}^{-1} \text{diag}(\lambda_1^{-1}, \dots, \lambda_{n_H-1}^{-1}) \tilde{Z}^T \Delta \tilde{Z}). \end{aligned}$$

and so, for sufficiently large $\hat{\alpha}$, we have:

$$\|\tilde{A}(\hat{\alpha})^{-1}\|_2 \leq \frac{\hat{\alpha}^{-1} \max_{i < n_H} \lambda_i^{-1}}{1 - \hat{\alpha}^{-1} \max_{i < n_H} \lambda_i^{-1} \|\tilde{Z}^T \Delta \tilde{Z}\|_2} \rightarrow 0$$

as $\hat{\alpha} \rightarrow \infty$. Hence we may write, for $\hat{\alpha}$ sufficiently large,

$$\begin{bmatrix} \tilde{A}(\hat{\alpha}) & \tilde{\delta} \\ \tilde{\delta}^T & \eta \end{bmatrix}^{-1} = \begin{bmatrix} I & -\tilde{A}(\hat{\alpha})^{-1} \tilde{\delta} \\ \mathbf{0}^T & 1 \end{bmatrix} Y(\hat{\alpha}) \begin{bmatrix} I & \mathbf{0} \\ -\tilde{\delta}^T \tilde{A}(\hat{\alpha})^{-1} & 1 \end{bmatrix}$$

$$\text{with } Y(\hat{\alpha}) := \begin{bmatrix} \tilde{A}(\hat{\alpha})^{-1} & \mathbf{0} \\ \mathbf{0}^T & (\eta - \tilde{\delta}^T \tilde{A}(\hat{\alpha})^{-1} \tilde{\delta})^{-1} \end{bmatrix}. \quad (16)$$

which implies

$$\begin{bmatrix} \tilde{A}(\hat{\alpha}) & \tilde{\delta} \\ \tilde{\delta}^T & \eta \end{bmatrix}^{-1} = \begin{bmatrix} O & \mathbf{0} \\ \mathbf{0}^T & \eta^{-1} \end{bmatrix} + \mathcal{O}(\hat{\alpha}^{-1}), \quad (17)$$

and, by (15), we have

$$\begin{aligned} A_{HH}(\hat{\alpha})^{-1} &= Z \begin{bmatrix} O & \mathbf{0} \\ \mathbf{0}^T & \eta^{-1} \end{bmatrix} Z^T + \mathcal{O}(\hat{\alpha}^{-1}) \\ &= \mathbf{e}_h (\mathbf{e}_r^T A_{rr}^{(L)} \mathbf{e}_r)^{-1} \mathbf{e}_h^T + \mathcal{O}(\hat{\alpha}^{-1}), \end{aligned} \quad (18)$$

which proves part (i) of the Lemma.

Parts (ii) and (iii) follow from simple substitution, using (10) and (13). \square

To understand this lemma a bit better, we define the limiting forms:

$$\begin{aligned} A_{HH}(\infty)^{-1} &:= \mathbf{e}_h (\mathbf{e}_r^T A_{rr}^{(L)} \mathbf{e}_r)^{-1} \mathbf{e}_h^T, \\ S(\infty) &:= A_{LL} - A_{LH} A_{HH}(\infty)^{-1} A_{HL} \\ &= A_{LL} - (A_{LR} \mathbf{e}_r) (\mathbf{e}_r^T A_{rr}^{(L)} \mathbf{e}_r)^{-1} (\mathbf{e}_r^T A_{rL}), \\ P_{LH}(\infty) &:= A_{LH} A_{HH}(\infty)^{-1} \\ &= (A_{LR} \mathbf{e}_r) (\mathbf{e}_r^T A_{rr}^{(L)} \mathbf{e}_r)^{-1} \mathbf{e}_h^T. \end{aligned}$$

Note that $S(\infty)$ can also be interpreted as the Schur complement of $c^2 \mathbf{e}_r^T A_{rr}^{(L)} \mathbf{e}_r$ in the matrix

$$A_{LL}^\infty = \begin{bmatrix} c^2 \mathbf{e}_r^T A_{rr}^{(L)} \mathbf{e}_r & c \mathbf{e}_r^T A_{rL} \\ c A_{LR} \mathbf{e}_r & A_{LL} \end{bmatrix},$$

for any nonzero value of c . In particular, if we choose $c := n_H^{1/2}$, then $c \mathbf{e}_r = \mathbf{1}_r$, the vector of all ones on r and, using also (12), we have

$$A_{LL}^\infty := \begin{bmatrix} \mathbf{1}_r^T A_{rr}^{(L)} \mathbf{1}_r & \mathbf{1}_r^T A_{rL} \\ A_{LR} \mathbf{1}_r & A_{LL} \end{bmatrix} = \begin{bmatrix} \mathbf{1}_H^T A_{HH}(1) \mathbf{1}_H & \mathbf{1}_H^T A_{HL} \\ A_{LH} \mathbf{1}_H & A_{LL} \end{bmatrix}. \quad (19)$$

This is the stiffness matrix for a pure Dirichlet problem for the Laplacian on all of Ω with the additional constraint that the solution is constant on Ω_H . See Figure 2.




Fig. 2 The matrix in (19) corresponds to a homogeneous Dirichlet problem for the Laplacian on Ω under the constraint that the solution is constant on Ω_H .

Thus, when $\hat{\alpha} \gg 1$, the original problem decouples almost entirely into a (regularised) Neumann problem (i.e. $A_{HH}(\hat{\alpha})$ for the Laplacian on Ω_H (scaled by $\hat{\alpha}$) and a Dirichlet problem (i.e. A_{LL}^∞) for the Laplacian on all of Ω , but under the additional constraint that the solution is constant on Ω_H . The coupling of the two problems (i.e. $A_{LH} A_{HH}(\hat{\alpha})^{-1}$) reduces to a transfer of the average of the

solution over Ω_H to Ω_L . Efficient and robust multilevel preconditioners exist (with theory) for the two subproblems and we will exploit exactly this fact to construct preconditioners that we can prove are robust with respect to mesh size and coefficient variations. We shall now explain this in the context of the model problem considered in this section.

2.2 A suitable preconditioner

Based on the above perturbation analysis we propose the following preconditioner:

$$B(\hat{\alpha}) := \begin{bmatrix} I_{HH} & -P_{LH}(\infty)^T \\ 0 & I_{LL} \end{bmatrix} \begin{bmatrix} A_{HH}(\hat{\alpha}) & 0 \\ 0 & S(\infty) \end{bmatrix}^{-1} \begin{bmatrix} I_{HH} & 0 \\ -P_{LH}(\infty) & I_{LL} \end{bmatrix}. \quad (20)$$

The following theorem shows that B is an effective preconditioner for $\hat{\alpha} \gg 1$.

Theorem 1 *For $\hat{\alpha}$ sufficiently large we have*

$$\sigma(B(\hat{\alpha})A(\hat{\alpha})) \subset [1 - c\hat{\alpha}^{-1/2}, 1 + c\hat{\alpha}^{-1/2}]$$

for some constant c independent of $\hat{\alpha}$, and therefore

$$\kappa(B(\hat{\alpha})A(\hat{\alpha})) = 1 + \mathcal{O}(\hat{\alpha}^{-1/2}).$$

Remark 1 It is possible to carry out a more detailed perturbation analysis of $A_{HH}(\hat{\alpha})^{-1}$ and $S(\hat{\alpha})$, and to quantify the constant c in the above theorem. It turns out that $c \leq \kappa_{\text{eff}}(\mathcal{N}_{HH})^{1/2}$, where $\kappa_{\text{eff}}(\mathcal{N}_{HH}) = \lambda_{\max}(\mathcal{N}_{HH})/\lambda_2(\mathcal{N}_{HH})$ is the effective condition number of \mathcal{N}_{HH} . In the case of a quasi-uniform mesh $\kappa_{\text{eff}}(\mathcal{N}_{HH})^{1/2} = \mathcal{O}(h^{-1}) = \mathcal{O}(n_H^{1/2})$, where n_H is the number of nodes in $\overline{\Omega}_H$. Therefore provided $\hat{\alpha} \gg n_H$, the preconditioned matrix $B(\hat{\alpha})A(\hat{\alpha})$ is well conditioned, i.e. $\kappa(B(\hat{\alpha})A(\hat{\alpha})) = 1 + \mathcal{O}((n_H/\hat{\alpha})^{1/2})$. The proof of this requires substantial further analysis. Details will be given in the forthcoming work [21].

Proof Letting $M^{1/2}$ denote the square root of any symmetric positive definite matrix M , we write $B = L^T L$ with

$$L := \begin{bmatrix} A_{HH}^{-1/2} & 0 \\ -S(\infty)^{-1/2}P_{LH}(\infty) & S(\infty)^{-1/2} \end{bmatrix}.$$

(Note that for notational convenience we do not explicitly state which terms depend on $\hat{\alpha}$ everywhere in this proof.) A straightforward calculation shows that

$$\sigma(BA) = \sigma(LAL^T) = \sigma(I + R), \quad (21)$$

with

$$R := \begin{bmatrix} 0 & R_{HL} \\ R_{HL}^T & 0 \end{bmatrix} \quad \text{and} \quad R_{HL} := A_{HH}^{-1/2}(A_{HL} - A_{HH}P_{LH}^T(\infty))S(\infty)^{-1/2}.$$

As an example of the computation leading to (21), note that the bottom right-hand entry of the product LAL^T reads:

$$S(\infty)^{-1/2} [P_{LH}(\infty)A_{HH}P_{LH}(\infty)^T - P_{LH}(\infty)A_{HL} - A_{LH}P_{LH}(\infty)^T + A_{LL}]S(\infty)^{-1/2} = I,$$

since, by definition of $P_{LH}(\infty)$ and of η , we have

$$-A_{LH}P_{LH}(\infty)^T + A_{LL} = S(\infty)$$

and

$$P_{LH}(\infty)A_{HH}P_{LH}(\infty)^T - P_{LH}(\infty)A_{HL} - A_{LH}(\mathbf{e}_H\eta^{-1}\mathbf{e}_H^T A_{HH}\mathbf{e}_H\eta^{-1}\mathbf{e}_H^T - \mathbf{e}_H\eta^{-1}\mathbf{e}_H^T)A_{HL} = 0.$$

To finish the proof we shall show that, for $\hat{\alpha}$ sufficiently large,

$$A_{HH}^{-1/2} = \mathbf{e}_H\eta^{-1/2}\mathbf{e}_H^T + \mathcal{O}(\hat{\alpha}^{-1/2}). \quad (22)$$

On the assumption that (22) holds, we have

$$R_{LH} = S(\infty)^{-1/2}A_{LH}(I_{HH} - \mathbf{e}_H\eta^{-1}\mathbf{e}_H^T A_{HH}\mathbf{e}_H\eta^{-1}\mathbf{e}_H^T + \mathcal{O}(\hat{\alpha}^{-1/2})) = \mathcal{O}(\hat{\alpha}^{-1/2}) \quad (23)$$

and so the spectral radius $\rho(R)$ of R is $\mathcal{O}(\hat{\alpha}^{-1/2})$, which together with (21) completes the proof.

To prove (22), let us write down the eigenvalue decomposition of $A_{HH}(\hat{\alpha})$

$$Q(\hat{\alpha})^T A_{HH}(\hat{\alpha}) Q(\hat{\alpha}) = \text{diag}(\mu_1(\hat{\alpha}), \dots, \mu_{n_H}(\hat{\alpha})) \quad (24)$$

where $\{\mu_i(\hat{\alpha}) : i = 1, \dots, n_H\}$ denotes any non-increasing ordering of the eigenvalues of $A_{HH}(\hat{\alpha})$. Since $A_{HH}(\hat{\alpha})$ is SPD (see the discussion following (15)), we have $\mu_i(\hat{\alpha}) > 0$ for all $i \leq n_H$. Moreover, the μ_i are continuous functions of $\hat{\alpha}$, with

$$\hat{\alpha}^{-1}\mu_i(\hat{\alpha}) = \lambda_i + \mathcal{O}(\hat{\alpha}^{-1}), \quad (25)$$

as $\hat{\alpha} \rightarrow \infty$, where the λ_i are as defined in the proof of Lemma 1 and we have used (15). However, we also know from (18) that (for $\hat{\alpha}$ sufficiently large) the largest eigenvalue of $A_{HH}(\hat{\alpha})^{-1}$ is given by

$$\mu_{n_H}(\hat{\alpha})^{-1} = \eta^{-1} + \mathcal{O}(\hat{\alpha}^{-1}). \quad (26)$$

Therefore, using (24), (25) and (26), we have

$$Q(\hat{\alpha})^T A_{HH}(\hat{\alpha})^{-1/2} Q(\hat{\alpha}) = \text{diag}(0, \dots, 0, \eta^{-1/2}) + \mathcal{O}(\hat{\alpha}^{-1/2}).$$

The required estimate (22) follows by noting that the last column of $Q(\hat{\alpha})$ approaches \mathbf{e}_H with $\mathcal{O}(\hat{\alpha}^{-1/2})$ as $\hat{\alpha} \rightarrow \infty$. \square

Remark 2 Applying the original Aksoylu-Klie preconditioner B_{AK0} ([3]) defined in (7) to $A(\hat{\alpha})$ we get

$$B_{AK0}A(\hat{\alpha}) = \begin{bmatrix} I_{HH} & A_{HH}(\hat{\alpha})^{-1}A_{HL} \\ 0 & S(\hat{\alpha}) \end{bmatrix}.$$

Thus as shown in [3],

$$\sigma(B_{AK0}A) = \{1\} \cup \sigma(S(\hat{\alpha})).$$

We see from Lemma 1 that, as $\hat{\alpha} \rightarrow \infty$, $S(\hat{\alpha})$ converges to a rank 1 perturbation of A_{LL} . By standard theory for such perturbations (e.g. [8, Theorem 8.1.5]), for large $\hat{\alpha}$ and small h , the condition number of $S(\hat{\alpha})$ will be close to the condition number of A_{LL} , which grows with h^{-2} (assuming the area of the domain Ω_L is of fixed size). Therefore B_{AK0} will lose robustness as $h \rightarrow 0$ (even if $\hat{\alpha} \gg h^2$).

To illustrate the sharpness of the estimates in Theorem 1, we computed the eigenvalues of $B^{-1}A$ for the geometry illustrated in Figure 3. This was done by exact computation of the blocks in B , and so is restricted to moderate numbers of degrees of freedom. The domain Ω is a unit square and the dark interior square is a high-permeability region (denoted Ω_H) which is placed at the centre of Ω , and has side length ρ . (See Figure 3.) Then $\Omega_L = \Omega \setminus \overline{\Omega}_H$. The unit square is covered with a uniform triangular mesh with mesh diameter h , which resolves the boundary of Ω_H . Dirichlet conditions are applied on the boundary of Ω and n_H is the number of degrees of freedom in $\overline{\Omega}_H$. We study results for two different choices of ρ , namely $\rho = 1/2$ fixed and $\rho = 4h$, decreasing as the mesh is refined. In the tables, β denotes the quantity $\kappa_{\text{eff}}(\mathcal{N}_{HH})^{1/2}\hat{\alpha}^{-1/2}$ which is an upper bound for the quantity $c\hat{\alpha}^{-1/2}$ in Theorem 1, so that Theorem 1 predicts that the spectrum of $B^{-1}A$ is bracketed in the interval $[1 - \beta, 1 + \beta]$. We also display the exact smallest and largest eigenvalues of $B^{-1}A$ which are denoted λ_1 and λ_n respectively.




Fig. 3 The unit square domain Ω and the high permeability region Ω_H , which is a square of side length ρ centred at the centre of Ω .

From Table 1 we see first of all that the spectrum of $B^{-1}A$ lies within the interval $[1 - \beta, 1 + \beta]$, as predicted.

Observing that λ_1 and λ_n are symmetrically located with respect to 1 let us look more carefully at the behaviour of λ_n . Reading across the rows in Table 1, we see that $\lambda_n - 1$ decreases very clearly with $\mathcal{O}(\hat{\alpha}^{-1/2})$. Reading down the columns we see that $\lambda_n - 1$ increases with a rate bounded by $\mathcal{O}(n_H^{1/2})$, which is what is predicted by Theorem 1 and Remark 1.

Table 2 illustrates the case when $\rho = 4h$, so that the number of degrees of freedom n_H remains fixed at 25 as $h \rightarrow 0$. Reading down the columns we see that for fixed $\hat{\alpha}$, β remains fixed as h decreases and λ_1 and λ_n remain within fixed distances from 1. Again reading across the rows we see that the distance of λ_1 and λ_n from 1 decreases clearly with $\mathcal{O}(\hat{\alpha}^{-1/2})$.

2.3 Practical implementation of the preconditioner

The preconditioner $B(\hat{\alpha})$ still involves inverses of the blocks $A_{HH}(\hat{\alpha})$ and $S(\infty)$. Factorising these will be prohibitively large in realistic applications. However, as discussed above, in the limit as $\hat{\alpha} \rightarrow \infty$ both these blocks contain no more coefficient variation and can be efficiently preconditioned via multilevel preconditioners (with theoretical foundation).

Recall first that $\hat{\alpha}^{-1}A_{HH}(\hat{\alpha}) = \mathcal{N}_{HH} + \hat{\alpha}^{-1}\Delta$ where \mathcal{N}_{HH} is the Neumann problem for the Laplacian on Ω_H . Therefore $A_{HH}(\hat{\alpha})$ can be efficiently preconditioned via a standard multigrid V-cycle B_{HH} (with either geometric or algebraic coarsening strategy). This preconditioner is (justifiably) h -robust on the subspace orthogonal to \mathbf{e}_H , and there exists a well-documented theory explaining this (see e.g. [12] for geometric multigrid and [18] for algebraic multigrid).

Similarly we can build robust preconditioners for $S(\infty)$ via standard multigrid methods. Recall that $S(\infty) = A_{LL} - \mathbf{v}\eta^{-1}\mathbf{v}^T$ where $\mathbf{v} := A_{LH}\mathbf{e}_H$, $\eta := \mathbf{e}_H^T A_{HH} \mathbf{e}_H$, and A_{LL} is the Dirichlet problem for the Laplacian on Ω_L . If B_{LL} denotes a standard multigrid V-cycle for A_{LL} , we can construct an efficient and robust preconditioner \tilde{S}^{-1} for $S(\infty)$ using the Sherman-Morrison formula, i.e.

$$\tilde{S}^{-1} := B_{LL} + B_{LL}\mathbf{v}(1 - \eta)^{-1}\mathbf{v}^T B_{LL}.$$

Again it follows from standard multigrid theory that this preconditioner is h -robust. Note also that in practice we can precompute and store $B_{LL}\mathbf{v}$ during the setup phase. This means we only need to apply the multigrid V-cycle B_{LL} once per iteration.

Alternatively, we can also obtain an efficient preconditioner for $S(\infty)$ by constructing a standard multigrid V-cycle for A_{LL}^∞ . In this way one can avoid the application of the Sherman-Morrison formula (which may become prohibitively expensive in the case of multiple islands). However, a proof of the h -robustness of this approach can not be directly deduced from existing literature. In the numerical results in §4 we have used this approach. However, in the following (in a slight abuse of notation) we will refer to this preconditioner also by \tilde{S}^{-1} .

h^{-1}	n_H	$\hat{\alpha} = 10^2$			$\hat{\alpha} = 10^4$			$\hat{\alpha} = 10^6$		
		$1 - \beta$	λ_1	λ_n	$1 - \beta$	λ_1	λ_n	$1 - \beta$	λ_1	λ_n
8	25	0.511	0.869	1.311	0.951	0.987	1.013	0.995	0.9987	1.0013
16	81	0.146	0.789	1.211	0.915	0.978	1.022	0.991	0.9978	1.0022
32	289				0.842	0.967	1.033	0.984	0.9967	1.0033
64	1089				0.698	0.953	1.047	0.970	0.9953	1.0047

Table 1 Case 1, i.e. $\rho = 1/2$

h^{-1}	n_H	$\hat{\alpha} = 10^2$			$\hat{\alpha} = 10^4$			$\hat{\alpha} = 10^6$		
		$1 - \beta$	λ_1	λ_n	$1 - \beta$	λ_1	λ_n	$1 - \beta$	λ_1	λ_n
8	25	0.5111	0.8687	1.1313	0.9511	0.9866	1.0134	0.9511	0.9987	1.0013
16	25	0.5111	0.8382	1.1618	0.9511	0.9834	1.0166	0.9951	0.9983	1.0017
32	25				0.9511	0.9829	1.0171	0.9951	0.9983	1.0017
64	25				0.9511	0.9828	1.0171	0.9951	0.9983	1.0017

Table 2 Case 2, i.e. $\rho = 4h$

Therefore, the final (practical) preconditioner which we propose and use is

$$\tilde{B} := \begin{bmatrix} I_{HH} & -P_{LH}(\infty)^T \\ 0 & I_{LL} \end{bmatrix} \begin{bmatrix} B_{HH} & 0 \\ 0 & \tilde{S}^{-1} \end{bmatrix} \begin{bmatrix} I_{HH} & 0 \\ -P_{LH}(\infty) & I_{LL} \end{bmatrix}. \quad (27)$$

This preconditioner is robust on the subspace orthogonal to $\mathbf{e} := [\mathbf{e}_H^T, \mathbf{0}^T]^T$. To deal with the component of the solution in the direction of \mathbf{e} we use deflation techniques, i.e. we apply our preconditioner within a conjugate gradient algorithm for the deflated system

$$\mathcal{P}A\mathbf{u}^\perp = \mathcal{P}\mathbf{f} \quad (28)$$

where \mathcal{P}^T is the A -orthogonal projection into the subspace orthogonal to \mathbf{e} , that means

$$\mathcal{P}^T = I - \mathbf{e}\eta^{-1}\mathbf{e}^T A.$$

and $\mathbf{u}^\perp := \mathcal{P}^T \mathbf{u}$ is the projected solution. The component of the solution \mathbf{u} in the direction of \mathbf{e} is then simply given by

$$\mathbf{u} - \mathbf{u}^\perp = (I - \mathcal{P}^T)\mathbf{u} = \mathbf{e}\eta^{-1}\mathbf{e}^T \mathbf{f}.$$

This approach has a significant additional advantage. Working only on the deflated system (28) (where each CG iteration requires an application of the projections \mathcal{P} and \mathcal{P}^T to project the right hand side into the range of $\mathcal{P}A$ and to project the current iterate into the subspace orthogonal to \mathbf{e} , respectively), we do not require the application of the block $\begin{bmatrix} I_{HH} & 0 \\ -P_{LH}(\infty) & I_{LL} \end{bmatrix}$ and of its transpose. This is a simple consequence of the fact that

$$P_{LH}(\infty)A_{HH}\mathbf{e}_H\eta^{-1}\mathbf{e}_H^T = P_{LH}(\infty)$$

and thus

$$\begin{bmatrix} I_{HH} & 0 \\ -P_{LH}(\infty) & I_{LL} \end{bmatrix} \mathcal{P} = \begin{bmatrix} I_{HH} & 0 \\ -P_{LH}(\infty) & I_{LL} \end{bmatrix} \begin{bmatrix} I_{HH} - A_{HH}\mathbf{e}_H\eta^{-1}\mathbf{e}_H^T & 0 \\ -P_{LH}(\infty) & I_{LL} \end{bmatrix} = \mathcal{P}$$

This implies that within a CG algorithm for the deflated system (28) our preconditioner \tilde{B} decouples entirely into a preconditioner B_{HH} for $A_{HH}(\hat{\alpha})$ (on the subspace A_{HH} -orthogonal to \mathbf{e}_H) and into a preconditioner \tilde{S}^{-1} for $S(\hat{\alpha})$.

Remark 3 The fact that for the model problem considered here deflation against the vector \mathbf{e} leads to a huge improvement in the robustness of the preconditioned conjugate gradient method for any preconditioner, has been observed numerically (e.g. in [16]). The analysis in [16] is for deflation against exact eigenvectors, or small perturbations of these, while here we provide theory for deflation against the “physical” vector \mathbf{e} , which is not a small perturbation of a true eigenvector.

Remark 4 Note that based on the above perturbation analysis it is possible to construct simpler preconditioners which do not require an approximation of $S(\infty)^{-1}$, but only an approximation for A_{LL}^{-1} . It can be shown that (for $\hat{\alpha}$ sufficiently large again) these preconditioners lead to a clustering of the spectrum with a small number (1 or 2) of outliers. The position of these outliers on the real line is independent of the value of $\hat{\alpha}$, but it depends on the shape of the island Ω_H . More details will be in [21].

3 More general coefficients and boundary conditions

3.1 Disconnected high permeability regions

In this section we extend the results of the previous section to the case where the boundary conditions on $\partial\Omega$ may be of mixed type and where Ω_H may consist of multiple disconnected components. In addition the coefficient function α is no longer required to be constant on each of the regions Ω_H and Ω_L .

Specifically, we suppose that $\partial\Omega$ is partitioned into a Neumann part Γ_N and a Dirichlet part Γ_D and that the high-permeability region Ω_H may be written as $\Omega_H = \bigcup_{i=0}^k \Omega_i$

with $\overline{\Omega}_i \cap \overline{\Omega}_j = \emptyset$, for all $i, j = 0, \dots, k$, with $i \neq j$. Moreover, we assume that, for $i = 1, \dots, k$, each of the regions Ω_i is itself connected and that $\Omega_i \cap \Gamma_D = \emptyset$. The remaining component Ω_0 of Ω_H consists of the union of all remaining (disconnected) components of Ω_H that touch the Dirichlet boundary, i.e. for each $x \in \Omega_0$ there exists a continuous path in Ω_0 to Γ_D . Then the low permeability region is defined to be $\Omega_L := \Omega \setminus \overline{\Omega}_H$. Finally Γ denotes the portion of the boundary of Ω_H which does not coincide with $\partial\Omega$ and $\Gamma_i = \Gamma \cap \partial\Omega_i$.

Note that the stiffness matrix A can be scaled globally by any fixed parameter, without changing its condition number, and so we can assume, without loss of generality,

$$\inf_{x \in \Omega} \alpha(x) = 1. \quad (29)$$

Below we perform a singular perturbation analysis analogous to that in §2. The hidden constants in our estimates will depend on the coefficient variation in each of the high permeability and low-permeability regions which we encapsulate in a single parameter μ , chosen so that

$$\mu^{-1} \leq \frac{\alpha(x)}{\alpha(y)} \leq \mu, \quad x, y \in \Omega_i, \quad i = L, 0, \dots, k.$$

Our asymptotic analysis will be valid on the assumption that the minimum value of the coefficient on all the Ω_i is large compared to the maximum value on Ω_L . Accordingly we introduce local and global parameters:

$$\hat{\alpha}_i = \frac{\inf_{x \in \Omega_i} \alpha(x)}{\sup_{x \in \Omega_L} \alpha(x)} \quad \text{and} \quad \hat{\alpha} = \min\{\hat{\alpha}_i : i = 0, \dots, k\} \quad (30)$$

Note that this reduces to the same definition of $\hat{\alpha}$ as in the special case described in §2. Moreover, if $\hat{\alpha} \rightarrow \infty$ then $\hat{\alpha}_i \rightarrow \infty$, for each $i = 0, \dots, k$. Implicitly our estimates below will be valid for fixed μ as $\hat{\alpha} \rightarrow \infty$, i.e. under a kind of “scale separation” assumption. In general there may be considerable freedom in choosing the interface between the high- and low-permeability regions. The results below hold for all choices of boundary so that the scale separation assumption holds, and therefore for an “optimal” choice of interface. The optimal choice may not be easily identifiable in the case of general variable coefficients, but “reasonable” choices could be computed for example by identifying weak and strong connections as in AMG.

To simplify this discussion, let us assume for the moment that $k = 2$. (The general case is presented in Lemma 2 below.) Then, since the Ω_i are pairwise disjoint, for $i = 0, 1, 2$, we have

$$A_{HH}(\alpha) = \begin{bmatrix} A_{00}(\alpha) & 0 & 0 \\ 0 & A_{11}(\alpha) & 0 \\ 0 & 0 & A_{22}(\alpha) \end{bmatrix}, \quad (31)$$

and, generalising (12), we can write

$$A_{ii}(\alpha) = \hat{\alpha}_i \mathcal{N}_i + \Delta_i. \quad (32)$$

Here \mathcal{N}_i and Δ_i also depend on α . In fact \mathcal{N}_i is the stiffness matrix of a Neumann boundary value problem ($i = 1, 2$), respectively a mixed boundary value problem ($i = 0$) on Ω_i , with coefficient function

$$\hat{\alpha}_i^{-1} \alpha(x) = \left(\frac{\sup_{x \in \Omega_L} \alpha(x)}{\inf_{x \in \Omega_i} \alpha(x)} \right) \alpha(x),$$

which, from the assumptions above, can be easily seen to satisfy

$$1 \leq \hat{\alpha}_i^{-1} \alpha(x) \leq \mu^2, \quad x \in \Omega_i, \quad i = 0, 1, 2.$$

Moreover, analogously to (12),

$$\Delta_i = \begin{bmatrix} 0 & 0 \\ 0 & A_{\Gamma_i, \Gamma_i}^{(L)} \end{bmatrix},$$

where $A_{\Gamma_i, \Gamma_i}^{(L)}$ represents the coupling between nodes on Γ_i coming from the low permeability region, in which (by assumptions above), the coefficient varies between 1 and μ .

As in Section 2 we find that, as $\hat{\alpha} \rightarrow \infty$,

$$A_{ii}(\alpha)^{-1} = \mathbf{e}_i (\mathbf{e}_i^T \Delta_i \mathbf{e}_i)^{-1} \mathbf{e}_i^T + \mathcal{O}(\hat{\alpha}^{-1}) \quad \text{for } i = 1, 2, \quad (33)$$

and, since \mathcal{N}_0 has trivial nullspace, we have

$$A_{00}(\alpha)^{-1} = \mathcal{O}(\hat{\alpha}^{-1}), \quad \hat{\alpha} \rightarrow \infty. \quad (34)$$

Thus, if we denote by $\mathbf{e}_{H,i}$ the n_H -vector which coincides with \mathbf{e}_i on $\overline{\Omega}_i$ and is 0 elsewhere, then

$$A_{HH}(\alpha)^{-1} = \sum_{i=1}^2 \mathbf{e}_{H,i} (\mathbf{e}_i^T \Delta_i \mathbf{e}_i)^{-1} \mathbf{e}_{H,i}^T + \mathcal{O}(\hat{\alpha}^{-1}) \quad (35)$$

and

$$S(\alpha) = A_{LL} - \sum_{i=1}^2 (A_{LH} \mathbf{e}_{H,i}) (\mathbf{e}_i^T \Delta_i \mathbf{e}_i)^{-1} (\mathbf{e}_{H,i}^T A_{HL}) + \mathcal{O}(\hat{\alpha}^{-1}). \quad (36)$$

As in Section 2.1 we define

$$A_{HH}(\infty)^{-1} := \lim_{\hat{\alpha} \rightarrow \infty} A_{HH}(\alpha)^{-1}, \quad S(\infty) := \lim_{\hat{\alpha} \rightarrow \infty} S(\alpha),$$

$$P_{LH}(\infty) := \lim_{\hat{\alpha} \rightarrow \infty} A_{LH} A_{HH}(\alpha)^{-1},$$

and collect the results in a lemma that generalises Lemma 1.

Lemma 2

$$A_{HH}(\infty)^{-1} = \sum_{i=1}^k \mathbf{e}_{H,i} (\mathbf{e}_i^T \Delta_i \mathbf{e}_i)^{-1} \mathbf{e}_{H,i}^T,$$

$$S(\infty) = A_{LL} - \sum_{i=1}^k (A_{LH} \mathbf{e}_{H,i}) (\mathbf{e}_i^T \Delta_i \mathbf{e}_i)^{-1} (\mathbf{e}_{H,i}^T A_{HL}),$$

$$P_{LH}(\infty) = \sum_{i=1}^k (A_{LH} \mathbf{e}_{H,i}) (\mathbf{e}_i^T \Delta_i \mathbf{e}_i)^{-1} \mathbf{e}_{H,i}^T.$$

Thus, in the limit as $\hat{\alpha} \rightarrow \infty$ the Schur complement $S(\alpha)$ is a simple rank- k update of the block A_{LL} . As in the previous section $S(\infty)$ can again be interpreted as the Schur complement of the $k \times k$ diagonal matrix

$$\begin{aligned} E &:= \text{diag}(n_1 \mathbf{e}_1^T \Delta_1 \mathbf{e}_1, \dots, n_k \mathbf{e}_k^T \Delta_k \mathbf{e}_k) \\ &= \text{diag}(\mathbf{1}_{H,1}^T A_{HH} \mathbf{1}_{H,1}, \dots, \mathbf{1}_{H,k}^T A_{HH} \mathbf{1}_{H,k}) \end{aligned}$$

in

$$A_{LL}^\infty := \begin{bmatrix} E & F^T \\ F & A_{LL} \end{bmatrix} \quad \text{with} \quad F := A_{LH} [\mathbf{1}_{H,1}, \dots, \mathbf{1}_{H,k}] \quad (37)$$

and $\mathbf{1}_{H,i} := n_i^{1/2} \mathbf{e}_{H,i}$, i.e. the n_H -vector that is 1 on Ω_i and 0 elsewhere.

Thus, analogously to the discussion at the end of Section 2.1, for large $\hat{\alpha}$, the problem again essentially decouples into Neumann or mixed boundary value problems on each of the Ω_i and a Dirichlet problem on $\Omega \setminus \Omega_0$ with the additional constraint that the solution is constant on each of the Ω_i , $i = 1, \dots, k$. Again (and assuming of course that μ is small compared to $\hat{\alpha}$), there exist efficient and robust multilevel preconditioners for each of these decoupled problems.

With this insight, a suitable preconditioner B can be defined exactly as in (20) and we then have, analogously to Theorem 1, the following result.

Theorem 2 *For $\hat{\alpha}$ sufficiently large we have*

$$\sigma(BA) \subset [1 - c\hat{\alpha}^{-1/2}, 1 + c\hat{\alpha}^{-1/2}]$$

for some constant $c = c(\mu, h)$, which is independent of $\hat{\alpha}$ and in which the dependence on h is understood. (Details will be in [21].)

3.2 Multiscale media

Finally, we also consider multiscale media, i.e. high-permeability regions within regions of intermediate strength permeability. We simply apply the above framework recursively. For simplicity let us consider the following model situation: as in Section 2, we assume that Ω_H is connected and $\overline{\Omega}_H \cap \partial\Omega = \emptyset$, but now we assume further that $\Omega_H = \Omega_1 \cup \Omega_2$ where Ω_1 is again connected and $\overline{\Omega}_1 \cap \partial\Omega_H = \emptyset$, i.e. an island within an island. We will consider the case when α_2 is high relative to α_L and α_1 is high relative to α_2 , so we define

$$\hat{\alpha}_1 = \left(\frac{\inf_{x \in \Omega_1} \alpha(x)}{\sup_{x \in \Omega_2} \alpha(x)} \right) \quad \text{and} \quad \hat{\alpha}_2 = \left(\frac{\inf_{x \in \Omega_2} \alpha(x)}{\sup_{x \in \Omega_L} \alpha(x)} \right),$$

and assume that $\hat{\alpha} := \min\{\hat{\alpha}_1, \hat{\alpha}_2\} \rightarrow \infty$, i.e. we have a three-scale medium with good scale separation. Again assume moderate coefficient variation inside the subdomains, with μ defined as in Section 3.1.

We can now apply the above framework recursively. Provided $\hat{\alpha}_2$ is sufficiently large we obtain the limiting form of

the Schur complement $S(\alpha)$ (of A_{HH} in A) as before. However, there is still strong coefficient variation in the Neumann problem on Ω_H and so we need to apply our analysis again. We write

$$A_{HH} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}$$

and find as above that, since $\hat{\alpha}_1 \rightarrow \infty$, this problem again essentially decouples into a Neumann problem on Ω_1 and a problem on Ω_H with the constraint that the solution is constant on Ω_1 . This latter problem can again be described either through the limiting Schur complement of A_{11} in A_{HH} , i.e.

$$S_2(\infty) := A_{22} - (A_{21} \mathbf{e}_1) (\mathbf{e}_1^T \Delta_1 \mathbf{e}_1)^{-1} (\mathbf{e}_1^T A_{12}),$$

or through the $(n_2 + 1) \times (n_2 + 1)$ matrix

$$A_{22}^\infty := \begin{bmatrix} \mathbf{1}_1^T A_{11} \mathbf{1}_1 & \mathbf{1}_1^T A_{12} \\ A_{21} \mathbf{1}_1 & A_{22} \end{bmatrix},$$

and the preconditioner B can be extended accordingly.

4 Numerical Experiments

In this section we apply the conjugate gradient (CG) algorithm preconditioned using our preconditioner \tilde{B} defined in (27) and we compare this with several other possible choices of preconditioner in the context of several model problems. In the following tables the notation $CG + MG$ means that the CG algorithm is preconditioned using one V -cycle of standard geometric multigrid with SSOR smoother, piecewise linear interpolation and full weighting restriction. The problem on the coarsest grid is solved using the banded Cholesky solver `dpbsv.f` in LAPACK. The notation $CG + AMG$ means that one V -cycle of the Ruge and Stüben AMG algorithm `amg1r5.f` [19] is used as preconditioner for CG. Finally the notation $CG + \tilde{B}$ means that the preconditioner is \tilde{B} , which is constructed using the (above) geometric multigrid V -cycle on subblocks of A as described in §2.3. $\kappa(MG)$, $\kappa(AMG)$ and $\kappa(\tilde{B}A)$ are estimates (based on Ritz values obtained from the CG iteration) of the condition numbers of the respective preconditioned matrices.

The initial guess for the CG algorithm is taken to be zero and the stopping criterion is taken to be a relative residual reduction of $\epsilon = 10^{-8}$ in the Euclidean norm. In all the experiments below the domain is the unit square and the relevant elliptic problem is solved subject to a Dirichlet condition $u(x) = 1 - x_1$ where $x = (x_1, x_2)$, consistent with flow from left to right in the domain.

All computations are done using the GNU `fortran95` compiler `gfortran` on a linux intel core 2 laptop with clock-speed 2GHz, 2Gb of memory and cache size 4Mb.

Example 1 Our first example concerns the geometry in Figure 3, with $\rho = 1/2$ where the central island is given a coefficient $\alpha = \hat{\alpha}$ and the surrounding region a coefficient

h^{-1}	$CG + MG$		$CG + AMG$		$CG + \tilde{B}$		$\kappa(MG)$	$\kappa(AMG)$	$\kappa(\tilde{B}A)$
128	6	(0.11)	7	(0.10)	7	(0.11)	1.19	1.27	1.20
256	6	(0.30)	7	(0.36)	7	(0.40)	1.21	1.29	1.23
512	6	(1.29)	7	(1.48)	7	(1.48)	1.25	1.31	1.24
1024	6	(5.05)	7	(5.81)	7	(5.98)	1.24	1.35	1.26

Table 3 Iteration numbers (cpu times) for the geometry given in Figure 3, with $\rho = 1/2$ and $\hat{\alpha} = 10^6$ and associated condition numbers.

$\hat{\alpha}$	$CG + MG$		$CG + AMG$		$CG + \tilde{B}$		$\kappa(MG)$	$\kappa(AMG)$	$\kappa(\tilde{B}A)$
10^2	6	(5.06)	7	(5.83)	23	(17.7)	1.23	1.36	27.9
10^4	6	(5.04)	7	(5.81)	7	(5.98)	1.23	1.35	1.49
10^6	6	(5.05)	7	(5.81)	7	(5.98)	1.23	1.35	1.26
10^8	6	(5.05)	7	(5.81)	7	(5.98)	1.23	1.35	1.26

Table 4 Iteration numbers (cpu times) for the geometry given in Figure 3, with $\rho = 1/2$ and $h^{-1} = 1024$ and associated condition numbers.

$\alpha = 1$. Table 3 gives the iteration numbers and (in brackets) the cpu times for fixed $\hat{\alpha} = 10^6$ as $h \rightarrow 0$. Table 4 gives the iteration numbers and (in brackets) the cpu times for fixed $h^{-1} = 1024$ as $\hat{\alpha} \rightarrow \infty$.

We observe from these tables that the \tilde{B} preconditioner performs almost identically to the MG and AMG preconditioners for large enough $\hat{\alpha}$. In particular all three (as expected) are h -robust for $\hat{\alpha}$ fixed. In all three cases the condition numbers of the preconditioned matrices are very near to unity. A proof of the robustness of the CG+MG method is given in [26], under the assumption that the coarse grid resolves the discontinuous coefficient (which is the case here). We have given a proof of the robustness of \tilde{B} under quite general assumption on the shape of the interface between low and high permeability regions earlier in this paper. A proof of the $\hat{\alpha}$ -robustness of the AMG preconditioner is still lacking.

Example 2 The second experiment is for the geometry depicted in Figure 4. For this we perform the same experiments as in Example 1. Again the dark shaded regions are given a coefficient $\alpha = \hat{\alpha}$ and the remainder a coefficient $\alpha = 1$. The results are given in Tables 5 and 6. The results lead to similar conclusions as in Example 1.

We also remark that preconditioner \tilde{B} does not work so well as $\hat{\alpha}$ gets smaller (see Table 6). This is to be expected, since this method is based on asymptotic expansions as $\hat{\alpha} \rightarrow \infty$. However as mentioned earlier our analysis demonstrates theoretically the success of algebraic methods based on strong and weak connections in the high contrast case. Also, of course, standard geometric multigrid works well in the low-contrast case.

With respect to Examples 1 and 2, we would like to highlight a point that has already been made many times in the literature, but so far without theoretical justification. Using algebraic procedures to identify strong and weak couplings in the stiffness matrix arising from FE discretisations of high-contrast diffusion problems, it is indeed possible (as we have proved in this paper) to design multigrid meth-




Fig. 4 The unit square domain Ω and the high permeability region Ω_H , which consists of two squares of side length $1/5$ centred at the points $(3/10, 3/10)$ and $(7/10, 7/10)$.


ods that are $\hat{\alpha}$ -robust and have almost the same computational complexity as geometric multigrid methods for diffusion problems with constant coefficients. In fact, applying CG+MG to the Laplace problem $\hat{\alpha} = 1$ with $h^{-1} = 1280$ we require 6 iterations and 7.72 seconds with our code which is only slightly faster than the performance of CG+AMG and CG+ \tilde{B} for the high contrast case.

Example 3 As a final example we use the geometry pictured in Figure 5. This is derived from the Society of Petroleum Engineers benchmark example SPE10 [5]. We took Layer 59 of this benchmark test and created a binary medium by imposing a coarse mesh on the medium and identifying boxes where the permeability was above and below average. Regions of low permeability are shaded dark in Figure 5 in this case. In the case of applications in groundwater flow, to obtain physical flow fields, conservative discretizations like finite volume or mixed finite elements are usually used. When mixed finite elements with Raviart-Thomas velocity elements are applied to the problem (1) it is well-known that a resulting linear system of saddle-point type results. However this system can be reduced to a positive-definite

h^{-1}	$CG + MG$		$CG + AMG$		$CG + \tilde{B}$		$\kappa(MG)$	$\kappa(AMG)$	$\kappa(\tilde{B}A)$
160	6	(0.14)	7	(0.14)	7	(0.16)	1.20	1.39	1.21
320	6	(0.51)	7	(0.57)	7	(0.61)	1.21	1.38	1.23
640	6	(1.93)	7	(2.26)	7	(2.39)	1.22	1.38	1.24
1280	6	(7.84)	7	(8.88)	7	(9.81)	1.24	1.43	1.25

Table 5 Iteration numbers (cpu times) for the geometry given in Figure 4, with $\hat{\alpha} = 10^6$ and various condition numbers.

$\hat{\alpha}$	$CG + MG$		$CG + AMG$		$CG + \tilde{B}$		$\kappa(MG)$	$\kappa(AMG)$	$\kappa(\tilde{B}A)$
10^2	6	(7.82)	8	(9.54)	19	(22.9)	1.23	1.57	11.8
10^4	6	(7.85)	7	(8.90)	7	(9.85)	1.24	1.43	1.28
10^6	6	(7.84)	7	(8.88)	7	(9.81)	1.24	1.43	1.25
10^8	6	(7.79)	7	(8.87)	7	(9.82)	1.24	1.42	1.25

Table 6 Iteration numbers (cpu times) for the geometry given in Figure 4, with $h^{-1} = 1280$ and various condition numbers.**Fig. 5** The geometry of the SPE10-derived problem.

system of form (2) by the employment of a divergence-free basis ([6]). Moreover this reduced positive definite system corresponds (in the case of scalar coefficient) to a standard discretisation of a problem of form (1) with coefficient α^{-1} . Therefore to make our numerical experiments physically relevant, we perform them on the geometry depicted in Figure 5 with the coefficient value $\hat{\alpha}$ on the dark shaded areas taken to be large (corresponding to small permeability in the physical example).

The results are given in Tables 7, 8. Again we observe comparable performance of the three preconditioners. Because of the way we have chosen the permeability field in this case it is easy to create a coarse mesh in geometrical multigrid which resolves the coefficient discontinuity, and so geometrical multigrid works well. For more complicated coefficient fields this would not be the case and in the context of such examples, the comparison between AMG and \tilde{B} is the more realistic one.

Finally we would like to mention a phenomenon which we discovered while implementing these examples. That is that in the case of high contrast media, geometric multigrid is very sensitive to the choice of solver on the coarsest grid. In Tables 9,10 below we repeat the experiments of

Tables 7,8 but this time we use 200 iterations of SSOR for the solution of the problems on the coarsest level in $CG+MG$ and in $CG+\tilde{B}$ (instead of a direct solver). This changes the picture for $CG+MG$ rather strongly and it now no longer appears to be robust. The Ritz values (on the Krylov subspace generated by the CG iteration) suggest that the reason for this lack of robustness is related to a cluster of eigenvalues close to 0 which is not dealt with properly by the coarse solver. (The largest Ritz value is ~ 1.0 in all cases.) This phenomenon is not restricted to the case of the complicated geometry in Example 3. A similar behaviour of $CG+MG$ can be observed in the case of 1 or 2 islands in Examples 1 and 2. Since the smallest eigenvalues are dealt with explicitly in \tilde{B} , $CG+\tilde{B}$ does not suffer from this problem.

References

1. J. E. Aarnes. Modelling of multiscale structures in flow simulations for petroleum reservoirs. In *Geometric Modelling, Numerical Simulation, and Optimization Applied Mathematics at SINTEF*, pages 307–360. Springer Verlag, 2007.
2. J. E. Aarnes and T.Y. Hou. Multiscale domain decomposition methods for elliptic problems with high aspect ratios. *Acta Mathematicae Applicatae Sinica*, 18(1):63–76, 2002.
3. B. Aksoylu and H. Klie. A family of physics-based preconditioners for solving elliptic equations on highly heterogeneous media. Accepted for publication in *Applied Numerical Mathematics*, 2008.
4. T.F. Chan and T. Mathew. Domain decomposition methods. In *Acta Numerica*. Cambridge University Press, 2004.
5. M.A. Christie and M.J. Blunt, Tenth SPE Comparative Solution Project: A Comparison of Upscaling Techniques, *SPE Reservoir Engineering*, 12 (2001),308–317.
6. 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.
7. M.G. Gerritsen and L.J. Durlofsky. Modeling fluid flow in oil reservoirs. *Annu. Rev. Fluid Mech.*, 37:211–238, 2005.
8. G.H. Golub and C.F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 1989.
9. I. G. Graham and M. J. Hagger. Additive Schwarz, CG and discontinuous coefficients. In M. Espedal P. Bjørstad and D.E. Keyes, editors, *Proc. 9th Intern. Confer. on Domain Decomposition Methods*, Bergen, Norway, 1998.

h^{-1}	$CG + MG$		$CG + AMG$		$CG + \tilde{B}$		$\kappa(MG)$	$\kappa(AMG)$	$\kappa(\tilde{B}A)$
160	8	(0.18)	7	(0.18)	8	(0.20)	1.96	1.28	1.90
320	8	(0.64)	7	(0.66)	8	(0.75)	2.56	1.36	2.48
640	8	(2.52)	7	(2.45)	8	(3.17)	3.26	1.45	3.20
1280	8	(10.1)	8	(10.1)	9	(14.4)	4.06	1.62	4.12

Table 7 Iteration numbers (cpu times) for the geometry given in Figure 5, with $\hat{\alpha} = 10^6$ and various condition numbers.

$\hat{\alpha}$	$CG + MG$		$CG + AMG$		$CG + \tilde{B}$		$\kappa(MG)$	$\kappa(AMG)$	$\kappa(\tilde{B}A)$
10^2	10	(12.3)	8	(10.3)	>100	(***)	3.27	1.53	1.2(E+3)
10^4	8	(10.2)	8	(10.1)	17	(24.3)	3.98	1.51	9.71
10^6	8	(10.1)	8	(10.1)	9	(14.4)	4.06	1.62	4.12
10^8	8	(10.2)	7	(9.9)	8	(13.2)	4.06	1.55	4.06

Table 8 Iteration numbers (cpu times) for the geometry given in Figure 5, with $h^{-1} = 1280$ and various condition numbers.

h^{-1}	$CG + MG$		$CG + \tilde{B}$		$\lambda_1(MG)$	$\lambda_7(MG)$	$\lambda_1(\tilde{B}A)$
160	22	(1.34)	11	(0.75)	1.26(E-4)	0.53	0.40
320	25	(2.74)	11	(1.42)	1.25(E-4)	0.41	0.35
640	28	(8.84)	12	(4.84)	1.24(E-4)	0.32	0.28
1280	29	(33.6)	13	(19.8)	1.24(E-4)	0.26	0.23

Table 9 Iteration numbers (cpu times) for the geometry given in Figure 5, with $\hat{\alpha} = 10^6$ and various Ritz values (coarse solver: 200 iterations of SSOR).

$\hat{\alpha}$	$CG + MG$		$CG + \tilde{B}$		$\lambda_1(MG)$	$\lambda_7(MG)$	$\lambda_1(\tilde{B}A)$
10^2	11	(13.8)	>100	(***)	0.31	0.88	1.65(E-2)
10^4	29	(33.6)	18	(26.3)	1.14(E-2)	0.25	0.18
10^6	29	(33.6)	13	(19.8)	1.24(E-4)	0.26	0.23
10^8	11	(13.8)	13	(19.8)	0.24	0.91	0.23

Table 10 iteration numbers (cpu times) for the geometry given in Figure 5, with $h^{-1} = 1280$ and various Ritz values (coarse solver: 200 iterations of SSOR).

10. I. G. Graham and M. J. Hagger. Unstructured additive Schwarz-conjugate gradient method for elliptic problems with highly discontinuous coefficients. *SIAM J. Sci. Comp.*, 20(6):2041–2066, 1999.
11. I.G. Graham, P. Lechner, and R. Scheichl. Domain decomposition for multiscale PDEs. *Numerische Mathematik*, 106:589–626, 2007.
12. W. Hackbusch. *Iterative Solution of Large Sparse Systems of Equations*. Springer-Verlag, New York, 1993, 1993.
13. B.N. Khoromskij and G. Wittum. *Numerical Solution of Elliptic Differential Equations by Reduction to the Interface*. Lecture Notes in Computer Science and Engineering vol 36, Springer Verlag, 2004.
14. H. Klie. *Krylov-Secant methods for solving Large Scale Systems of Coupled Nonlinear Parabolic Equations*. PhD thesis, Dept. of Computational and Applied Mathematics, Rice University, Houston, TX, 1996.
15. P. Ming and X. Ye. Numerical methods for multiscale elliptic problems. *Journal of Computational Physics*, 214:421–445, 2006.
16. R. Nabben and C. Vuik. A comparison of deflation and the balancing preconditioner. *SIAM J. Sci. Comp*, 27:1742–1759, 2006.
17. B. Nötinger, V. Artus, and G. Zargar. The future of stochastic and upscaling methods in hydrogeology. *Hydrogeology Journal*, 13:184–201, 2005.
18. M. Brezina P. Vanek and J. Mandel. Convergence of algebraic multigrid based on smoothed aggregation. *Numer. Math.*, 88:559–579, 2001.
19. J. W. Ruge and K. Stüben. Algebraic multigrid. In *Multigrid methods, Frontiers Appl. Math.*, 3, pages 73–130. SIAM, Philadelphia, 1987.
20. T. Scheibe and S. Yabusaki. Scaling of flow and transport behavior in heterogeneous groundwater systems. *Advances in Water Resources*, 22:223238, 1998.
21. R. Scheichl. A rigorously justified algebraic preconditioner for high-contrast diffusion problems. *in preparation*, 2007.
22. R. Scheichl and E. Vainikko. Additive schwarz and aggregation-based coarsening for elliptic problems with highly variable coefficients. *Computing*, 80(4):319–343, 2007.
23. C. Vuik, A.Segal, and J.A. Meijerink. An efficient preconditioned CG method for the solution of a class of layered problems with extreme contrasts of coefficients. *J. of Comp. Phys.*, 152:385–403, 1999.
24. A. Guadagnini X. Sanchez-Vila and J. Carrera. Representative hydraulic conductivities in saturated groundwater flow. *Reviews of Geophysics*, 44:1–46, 2006.
25. L.J. Durlofsky X.H. Wen and M.G. Edwards. Upscaling of channel systems in two dimensions using flow-based grids. *Transport in Porous Media*, 51:343–366, 2003.
26. J. Xu and Y. Zhu. Uniform convergent multigrid methods for elliptic problems with strongly discontinuous coefficients. Technical Report Technical Report No. AM311, Dept of Mathematics, Penn State, Pennsylvania, USA, 2007.