Johann Radon Institute for
Computational and Applied Mathematics
Austrian Academy of Sciences (ÖAW)

Computing interpolation weights in AMG based on multilevel Schur complements

J.K. Kraus

RICAM-Report 2003-02

COMPUTING INTERPOLATION WEIGHTS IN AMG BASED ON MULTILEVEL SCHUR COMPLEMENTS

JOHANNES K. KRAUS

Abstract. This paper presents a particular construction of neighborhood matrices to be used in the computation of the interpolation weights in AMG (algebraic multigrid). The method utilizes the existence of simple interpolation matrices (piecewise constant for example) on a hierarchy of coarse spaces (grids). Then one constructs by algebraic means graded away coarse spaces for any given fine-grid neighborhood. Next, the corresponding stiffness matrix is computed on this graded away mesh, and the actual neighborhood matrix is obtained by computing the multilevel Schur complement of this matrix where degrees of freedom outside the neighborhood have to be eliminated. The paper presents algorithmic details, provides model complexity analysis as well as some comparative tests of the quality of the resulting interpolation based on the multilevel Schur complements versus element interpolation based on the true element matrices.

1. Introduction

Algebraic Multigrid (AMG) was first introduced in the early 80-ies [1, 2, 3, 4, 5] and immediately attracted substantial interest [12, 13, 14]. Recently, there has been a major resurgence of interest in the field of classical AMG [8], element-based AMG [6, 9, 10], spectral AMG [7], as well as AMG based on energy-minimizing interpolation and smoothed aggregation [11, 15, 16, 17, 18].

Element-based algebraic multigrid methods in general exploit additional information on the dicretization level. The main idea is to measure the quality of interpolation in a way that takes into account a heuristic saying that *interpolation should approximate eigenvectors with error proportional to the size of the associated eigenvalues.* This is a reasonable aim since common relaxation techniques reduce error components associated with small eigenvalues slowly. Hence, assuming that these "low energy" error components can be reduced effectively on coarser *grids*, the prolongation (in the coarse-grid correction step) has to be most accurate for these components.

AMG based on element interpolation (AMGe) [6, 10] has been a significant progress in extending algebraic multigrid, which was originally designed having M-matrices in mind, to non-M matrices, arising from elasticity problems, for instance. This is due to the fact that the knowledge of the individual element matrices (used to create assembled neighborhood matrices) allows a localization of the interpolation quality

Date: December 19, 2003.

1991 *Mathematics Subject Classification.* 65F10, 65N20, 65N30.

Key words and phrases. algebraic multigrid, multilevel Schur complements, element–free interpolation, graded away coarsening.

This work was supported by the Austrian Academy of Sciences.

measure, which then results in a simple rule for computing an efficient prolongation mapping. Thus, these small matrices carry with them implicitly the correct assignment and treatment of "weak" and "strong" connection.

Element-free AMGe [9] tries to accomplish the superior prolongation without the knowledge of the element matrices. The method uses an extension mapping to provide boundary values outside a neighborhood. In essence, this captures information that could be obtained from individual finite element matrices if they were available. However, the most efficient extension mappings assume information on the null-space of the global operator, e.g., the constant vector being in the null-space (for scalar second-order elliptic problems).

Spectral AMGe [7] was designed to apply the AMGe concept to even broader classes of problems. It uses the spectral decomposition of neighborhood matrices (small collections of element stiffness matrices) in order to determine local representations of algebraically smooth error components. Contrary to classical AMG, spectral AMGe avoids any assumptions on the nature of "smooth" error.

In this paper we focus on an alternative approach for constructing neighborhood matrices that can be used in the process of building the interpolation in element-based AMG. However, we neither assume access to individual element matrices nor do we make any assumptions on the nullspace of the global stiffness matrix in the present paper.

The remainder of the paper is organized as follows: In section 2 we describe a method for the construction of two–level neighborhood matrices. Then, in section 3 a multilevel algorithm that recursively applies the two–level method is proposed. Section 4 provides algorithmic details of an efficient implementation as well as a complexity analysis. We summarize some spectral properties of the neigborhood matrices in section 5. Finally, we present some numerical results in section 6.

2. Two–level neighborhood matrices

Let D be the set of fine grid degrees of freedom (dofs) and let P be an interpolation matrix that maps the set of coarse dofs D_c onto D. Let Ω be a fine–grid neighborhood, a small subset of D. Given a sparse matrix A one can use the following partitioning

$$A = \begin{bmatrix} A_{\text{int, int}} & A_{\text{int, ext}} \\ A_{\text{ext, int}} & A_{\text{ext, ext}} \end{bmatrix} \begin{matrix} \} \text{ dofs in } \Omega \\ \} \text{ dofs outside } \Omega. \end{matrix}$$

Accordingly, P admits similar block structure, one first orders the rows of P that correspond to Ω and then the rows that correspond to the exterior of Ω,

$$P = \begin{bmatrix} P_{\text{int}} \\ P_{\text{ext}} \end{bmatrix}.$$

We first build a new interpolation matrix $\pi = \pi_\Omega$, which coarsens only the exterior of Ω, that is, we keep the old coarse degrees of freedom and use them to interpolate the dofs outside Ω based on P_{ext} whereas the dofs in Ω are treated as new (richer) coarse dofs and the interpolation is the identity on Ω; that is, one has

$$\pi = \begin{bmatrix} I & 0 \\ 0 & P_{\text{ext}} \end{bmatrix}.$$

Note that the number of columns of π is equal to the total number of coarse dofs plus the number of fine dofs in Ω. Based on π we compute the new coarse matrix $\widehat{A}_\Omega = \pi^T A \pi$,

$$\widehat{A}_\Omega = \begin{bmatrix} \widehat{A}_{\text{int, int}} & \widehat{A}_{\text{int, ext}} \\ \widehat{A}_{\text{ext, int}} & \widehat{A}_{\text{ext, ext}} \end{bmatrix} = \begin{bmatrix} A_{\text{int, int}} & A_{\text{int, ext}} P_{\text{ext}} \\ P_{\text{ext}}^T A_{\text{ext, int}} & P_{\text{ext}}^T A_{\text{ext, ext}} P_{\text{ext}} \end{bmatrix}.$$

Since Ω will run over a set of fine–grid neighborhoods that provide an overlapping partitioning of D, the direct computation of all $\widehat{A}_\Omega$, especially the blocks $P_{\text{ext}}^T A_{\text{ext, ext}} P_{\text{ext}}$, may appear too costly. That is why we propose the following more computationally feasible algorithm.

One first computes the global coarse matrix $A_c = P^T A P$ once. Then for each neighborhood Ω one computes the small sized matrices

$$A_{c,\text{ int, int}} = P_{\text{int}}^T A_{\text{int, int}} P_{\text{int}} \tag{2.1}$$

$$\widehat{A}_{\text{int, ext}} = A_{\text{int, ext}} P_{\text{ext}} \tag{2.2}$$

$$\widehat{A}_{\text{ext, int}} = P_{\text{ext}}^T A_{\text{ext, int}}. \tag{2.3}$$

Then based on the identity

$$\begin{aligned} A_c &= P_{\text{ext}}^T A_{\text{ext, ext}} P_{\text{ext}} + P_{\text{int}}^T A_{\text{int, int}} P_{\text{int}} + P_{\text{int}}^T A_{\text{int, ext}} P_{\text{ext}} + P_{\text{ext}}^T A_{\text{ext, int}} P_{\text{int}} \\ &= P_{\text{ext}}^T A_{\text{ext, ext}} P_{\text{ext}} + A_{c,\text{ int, int}} + P_{\text{int}}^T \widehat{A}_{\text{int, ext}} + \widehat{A}_{\text{ext, int}} P_{\text{int}}, \end{aligned}$$

one computes

$$\widehat{A}_{\text{ext, ext}} = P_{\text{ext}}^T A_{\text{ext, ext}} P_{\text{ext}} = A_c - A_{c,\text{ int, int}} - P_{\text{int}}^T \widehat{A}_{\text{int, ext}} - \widehat{A}_{\text{ext, int}} P_{\text{int}}. \tag{2.4}$$

This implementation avoids the use of the typically large blocks $A_{\text{ext, ext}}$ in the computation.

The actual neighborhood matrix that we associate with the neighborhood Ω is the Schur complement

$$A_\Omega = A_{\text{int, int}} - \widehat{A}_{\text{int, ext}} \widehat{A}_{\text{ext, ext}}^{-1} \widehat{A}_{\text{ext, int}}. \tag{2.5}$$

The motivation to consider the Schur complement A_Ω is as follows.

Lemma 2.1. *Consider the vector space*

$$V_\Omega^c = \left\{ \begin{bmatrix} \mathbf{v}_{int} \\ P_{ext} \mathbf{v}_c \end{bmatrix} \right\}. \tag{2.6}$$

Then for a given local fine–grid vector $\mathbf{v}_{int}$ defined on Ω the vector

$$E\mathbf{v}_{int} = \begin{bmatrix} \mathbf{v}_{int} \\ -P_{ext} \left(P_{ext}^T A_{ext,\ ext} P_{ext} \right)^{-1} P_{ext}^T A_{ext,\ int} \mathbf{v}_{int} \end{bmatrix} \in V_\Omega^c$$

provides the minimum energy extension of $\mathbf{v}_{int}$ in the space V_Ω^c. Here we assume that A is symmetric positive (semi–)definite (and $P_{ext}^T A_{ext,\ ext} P_{ext}$ is invertible).

Proof. This is seen from the identity

$$\begin{aligned}\inf_{\mathbf{v}_c}\begin{bmatrix}\mathbf{v}_{\text{int}}\\ P_{\text{ext}}\mathbf{v}_c\end{bmatrix}^T A\begin{bmatrix}\mathbf{v}_{\text{int}}\\ P_{\text{ext}}\mathbf{v}_c\end{bmatrix} &= \inf_{\mathbf{v}_c}\begin{bmatrix}\mathbf{v}_{\text{int}}\\ \mathbf{v}_c\end{bmatrix}^T \pi^T A\pi\begin{bmatrix}\mathbf{v}_{\text{int}}\\ \mathbf{v}_c\end{bmatrix}\\ &= \inf_{\mathbf{v}_c}\begin{bmatrix}\mathbf{v}_{\text{int}}\\ \mathbf{v}_c\end{bmatrix}^T \widehat{A}_\Omega\begin{bmatrix}\mathbf{v}_{\text{int}}\\ \mathbf{v}_c\end{bmatrix}\\ &= \mathbf{v}_{\text{int}}^T A_\Omega \mathbf{v}_{\text{int}}\\ &= (E\mathbf{v}_{\text{int}})^T A(E\mathbf{v}_{\text{int}}).\end{aligned}$$

In fact, the minimum is achieved for $\mathbf{v}_c = -\widehat{A}^{-1}_{\text{ext, ext}}\widehat{A}_{\text{ext, int}}\mathbf{v}_{\text{int}}$. The last equality holds since for any coarse vector $\mathbf{w}_c$ one has

$$\begin{aligned}\begin{bmatrix}\mathbf{v}_{\text{int}}\\ P_{\text{ext}}\mathbf{w}_c\end{bmatrix}^T AE\mathbf{v}_{\text{int}} &= \begin{bmatrix}\mathbf{v}_{\text{int}}\\ P_{\text{ext}}\mathbf{w}_c\end{bmatrix}^T\\ &\quad\cdot\begin{bmatrix}A_\Omega\mathbf{v}_{\text{int}}\\ \left(I - A_{\text{ext, ext}}P_{\text{ext}}\left(P_{\text{ext}}^T A_{\text{ext, ext}}P_{\text{ext}}\right)^{-1}P_{\text{ext}}^T\right)A_{\text{ext, int}}\mathbf{v}_{\text{int}}\end{bmatrix}\\ &= \begin{bmatrix}\mathbf{v}_{\text{int}}\\ \mathbf{w}_c\end{bmatrix}^T\begin{bmatrix}A_\Omega\mathbf{v}_{\text{int}}\\ 0\end{bmatrix}\\ &= (\mathbf{v}_{\text{int}})^T A_\Omega\mathbf{v}_{\text{int}}\end{aligned}$$

□

3. The multilevel algorithm

In this section we will describe the procedure for the computation of the multilevel Schur complements. Assume that we have started with an Ω, a large set of fine grid dofs, and have computed a corresponding neighborhood matrix $A = A_\Omega$. Let G be a smaller set, $G \subset \Omega$, and assume that an interpolation mapping $P = P_\Omega$ which interpolates the dofs in Ω is given. Applying the two–level procedure from the previous section now to the matrices $A{:=}\, A_\Omega$, $P{:=}\, P_\Omega$ one computes the two–level Schur complement A_G. Thus we have now a smaller neighborhood matrix and the process can be repeated recursively.

To be precise, consider a sequence of nested neighborhood sets $\mathcal{T}_0 \subset \mathcal{T}_1 \subset \cdots \subset \mathcal{T}_\ell$. The notation $\mathcal{T}_{k-1} \subset \mathcal{T}_k$ means that for any fine grid neighborhood $G \in \mathcal{T}_{k-1}$ there is a coarse–level neighborhood $\Omega \in \mathcal{T}_k$ such that $G \subset \Omega$. Moreover, every neighborhood set $\mathcal{T}_k$ should provide an overlapping partition of the fine grid D_0. For instance, such nested neighborhood sets can be obtained via the element agglomeration algorithm proposed in [10]. This algorithm, which is based on the face–face graph of (agglomerated) elements, works purely algebraically. It starts from a set of fine grid elements. In a first step of agglomeration it determines a set of fine grid neighborhoods (agglomerates) $\mathcal{T}_0$. Then, the coarse–level neighborhood sets $\mathcal{T}_1, \mathcal{T}_2, \ldots \mathcal{T}_\ell$ are computed successively from the neighborhood set of the previous level. An example

is illustrated in Figure 1: Here the fine grid comprises 1600 triangles; the finest level neighborhood set $\mathcal{T}_0$ contains 382 neighborhoods (upper left picture); the coarse–level neighborhood sets $\mathcal{T}_1$ to $\mathcal{T}_5$ contain 93, 33, 15, 7, respectively 3 neighborhoods (upper right to lower right picture). Finally, the coarsest level neighborhood set $\mathcal{T}_\ell$, in this example $\mathcal{T}_6$ (without picture), is defined by one single element, i.e., the neighborhood that covers the whole domain.

Assume that we have interpolation matrices $P_1, \ldots, P_\ell$. Those can be global interpolation matrices that interpolate from a kth level grid D_k directly to the finest grid $D = D_0$. It makes sense to assume that D_k is a coarser set than D_{k-1}. To build a neighborhood matrix for any $G \in \mathcal{T}_{k-1}$, $G \subset \Omega$, $\Omega \in \mathcal{T}_k$, we will only need the rows of P_k which correspond to the dofs in Ω. That is, we set $Q_\Omega := P_k|_\Omega$, the matrix obtained from P_k by deleting all its rows corresponding to the dofs outside Ω. Then we apply the two level algorithm to the matrices A_Ω and Q_Ω in order to construct the neighborhood matrix A_G.

To start, we set A_Ω for $\Omega \in \mathcal{T}_\ell$ to be the fine grid matrix, that is, $\Omega = D$, (and in practice we assume that no essential boundary conditions have been imposed yet). Generally, we assume that the neighborhoods on coarser levels k contain a small, i.e. $\mathcal{O}(1)$ number of neighborhoods of the previous level $k-1$. Therefore, once having computed a neighborhood matrix A_Ω for a kth level neighborhood Ω, one can use the two-level method to compute all fine–grid ($(k-1)$th level) neighborhood matrices A_G for neighborhoods G that are contained in Ω.

Denoting the set of kth level neighborhood matrices by $\mathcal{A}_k = \{A_{\Omega_k} : \Omega_k \in \mathcal{T}_k\}$ the multilevel algorithm summarized below ends up with a set of fine–grid neighborhood matrices $\mathcal{A}_0$.

Algorithm 3.1. (Multilevel Schur complements)
Initiate: *set* $\mathcal{A}_\ell = \{A_{\Omega_\ell}\} = \{A\}$.
Recursively apply the two-level algorithm on all levels $k = \ell, \ell-1, \ldots, 1$:

- $\mathcal{A}_{k-1} := \emptyset$
- *For all neighborhoods* $\Omega_k \in \mathcal{T}_k$ *construct* $Q_{\Omega_k} := P_k|_{\Omega_k}$ *and then perform the following steps for all subneighborhoods* $\Omega_{k-1} \subset \Omega_k$ *where* $\Omega_{k-1} \in \mathcal{T}_{k-1}$*:*
  (1) *Build* $\pi = \begin{bmatrix} I & 0 \\ 0 & Q_{ext} \end{bmatrix}$ *where* $Q_{ext} := P_k|_{\Omega_k \setminus \Omega_{k-1}}$.
  (2) *Compute the two-level neighborhood matrix* $\widehat{A}_{\Omega_{k-1}} = \pi^T A_{\Omega_k} \pi$.
  (3) *Compute the Schur complement* $A_{\Omega_{k-1}} = A_{int,\, int} - \widehat{A}_{int,\, ext} \widehat{A}^{-1}_{ext,\, ext} \widehat{A}_{ext,\, int}$.
  (4) *Update* $\mathcal{A}_{k-1} := \mathcal{A}_{k-1} \cup \{A_{\Omega_{k-1}}\}$.

Figure 2 illustrates a multilevel neighborhhood that is involved in the computation of a particular neighborhood matrix $A_{\Omega_0} \in \mathcal{A}_0$. The set of degrees of freedom of the corresponding graded away mesh $\bar{D}_{\Omega_0}$ is the union of the kth level degrees of freedom in Ω_k over all $k = 0, 1, 2, \ldots, \ell$, where $\Omega_0 \subset \Omega_1 \subset \ldots \subset \Omega_\ell$, and $\Omega_k \in \mathcal{T}_k$, that is, $\bar{D}_{\Omega_0} = (D_0 \cap \Omega_0) \cup (D_1 \cap \Omega_1) \cup \ldots \cup (D_\ell \cap \Omega_\ell)$. The left-hand side picture shows the nested coarse–level neighborhoods where the coloring indicates the partitioning $\{\Omega_0, \Omega_1 \setminus \Omega_0, \ldots, \Omega_\ell \setminus \Omega_{\ell-1}\}$. The right-hand side picture indicates the subneighborhoods $G \in \mathcal{T}_{k-1}$ that are contained in $\Omega_k \setminus \Omega_{k-1}$.

FIGURE 1. Nested neighborhood sets: unstructured mesh

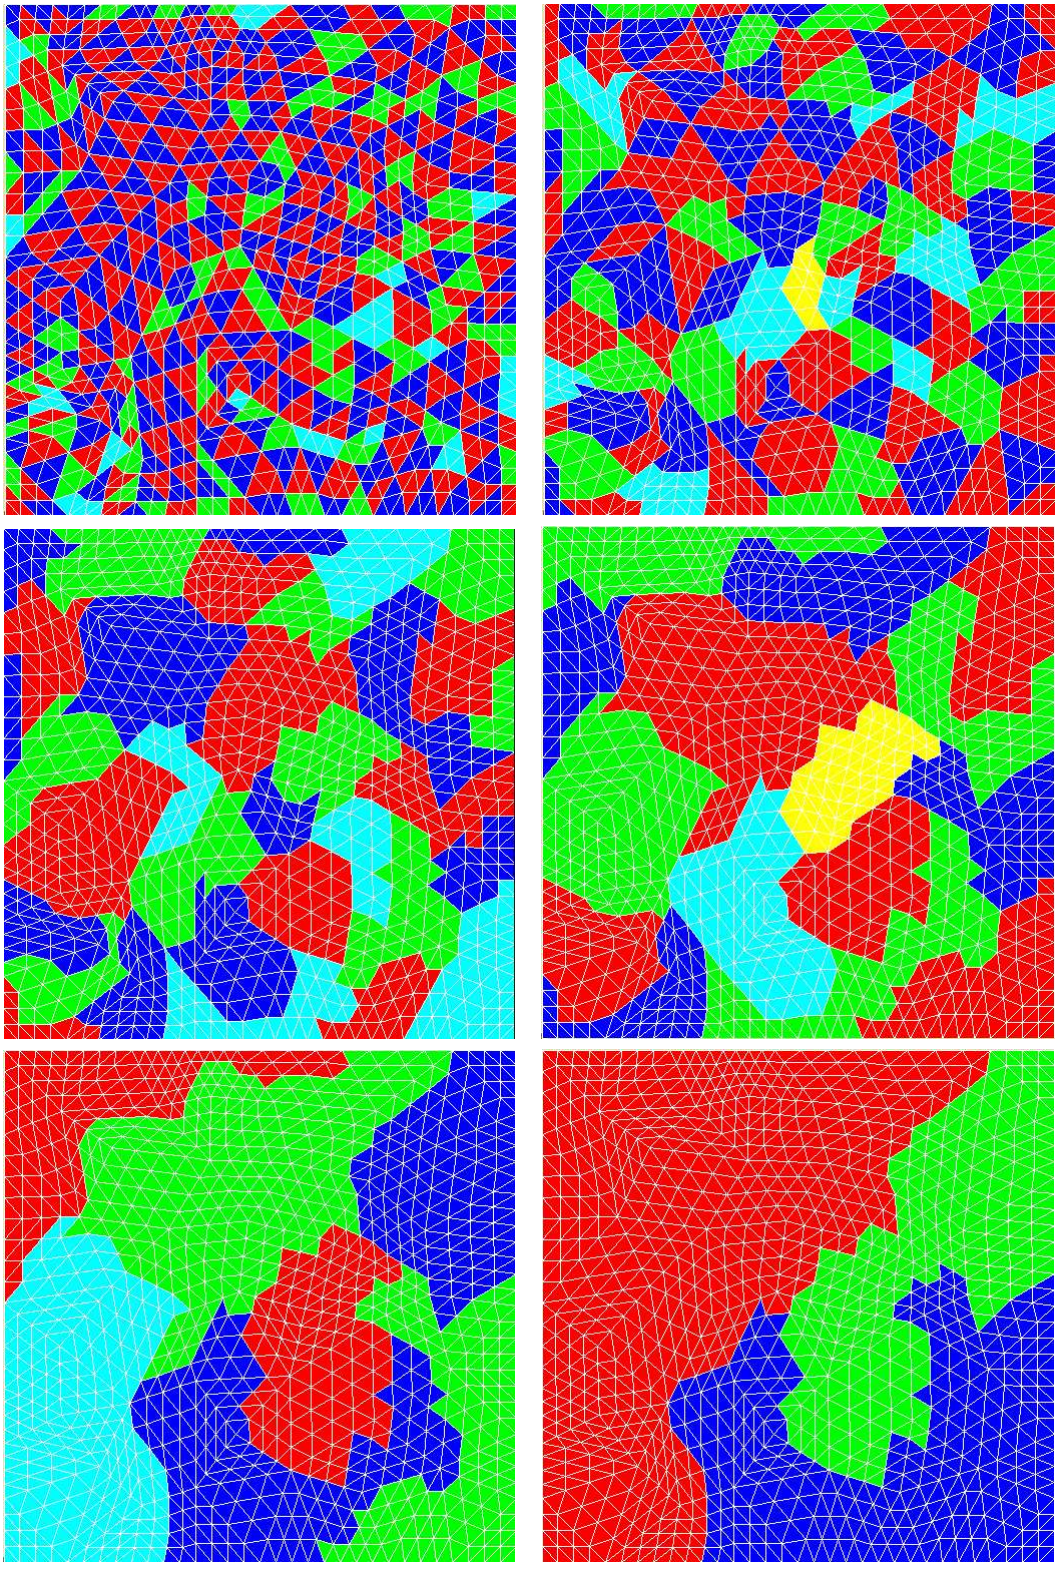

FIGURE 2. Multilevel neighborhood

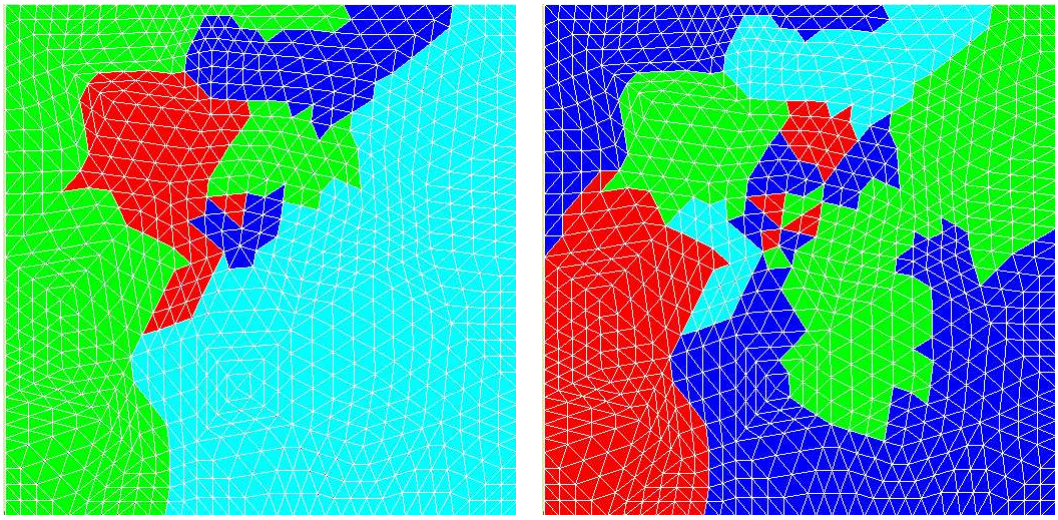

4. Computational complexity

The only complexity issue that may arise is that the matrices A_Ω get dense since they are Schur complements of coarser level neighborhood matrices. However, their size gets smaller; it equals the number of the degrees of freedom in Ω. The potential fill–in problem can be controlled by limiting the number of agglomeration levels ℓ and the size of the coarsest grid D_ℓ. The best one could achieve is a cost of order $\mathcal{O}(\ell N_0)$, where N_0 is the size of the fine grid. This is seen from the fact that at every level the best we could do is to perform a bounded number of operations for every dof in a given neighborhood Ω and since the neighborhoods at a given level cover the whole fine grid $D = D_0$ the minimal cost per level would be $\mathcal{O}(N_0)$.

A more refined analysis involves a proper storage of the Schur complements A_Ω in the form of low–rank updates to a sparse matrix B_Ω, namely, one may assume–and this will become clear from what follows–that

$$A_\Omega = B_\Omega - \sum_j A_{\Omega,\,c,\,j} A_{c,\,j}^{-1} A_{c,\,j,\,\Omega}.$$

Here, $A_{c,\,j}$ are coarse matrices of bounded, i.e., $(\mathcal{O}(1))$ size. For any set M by $\mathcal{N}(M)$ we will denote a quantity proportional to the size of M. In the following we will assume that the global interpolation operators can be localized (at an acceptable cost); a more precise definition of this property is:

Definition 4.1. (Localization property of interpolation) *A set* $\{P_k : D_k \to D_0, 1 \leq k \leq \ell\}$ *of global interpolation matrices is called to be lacalizable with respect to the sequence of nested neighborhood sets* $\mathcal{T}_0 \subset \mathcal{T}_1 \subset \cdots \subset \mathcal{T}_\ell$ *if the following is true: For any* k*th level neighborhood* $\Omega \in \mathcal{T}_k$, $k \in \{1, 2, \ldots, \ell\}$, *and all its* $(k-1)$*th level sub-neighborhoods* $G \subset \Omega$, $G \in \mathcal{T}_{k-1}$, *the* k*th level interpolation matrix* $P = P_k$ *restricted to* Ω *respectively* G, *denoted by* $Q = P|_\Omega$ *respectively* $Q_{int} = Q|_G$, *is such that the*

products $Q^T B_\Omega Q$ and $Q_{int}^T B_G Q_{int}$ can be computed in a number of operations proportional to $\mathcal{N}(\Omega)$ respectively $\mathcal{N}(G)$. Moreover, all fine-grid dofs from Ω (and hence from $\Omega_{ext} = \Omega \setminus G$) are interpolated from $\mathcal{O}(1)$ coarse dofs in $D_k \cap \Omega$.

Consider now any arbitrary but fixed neighborhood $\Omega \in \mathcal{T}_k$. Then the computational work that is required in order to compute the new two-level Schur complements A_G for all $G \subset \Omega$, $G \in \mathcal{T}_{k-1}$ arises from the following computations:

- In a first step we compute the *local* coarse matrix
$$A_{c,\,\Omega} = Q^T A_\Omega Q = Q^T B_\Omega Q - \sum_{j\geq k}^{\ell-1} (Q^T A_{\Omega,\,c,\,j}) A_{c,\,j}^{-1} (A_{c,\,j,\,\Omega} Q),$$
which requires order $\mathcal{N}(\Omega)\left(1 + \sum_{j\geq k}^{\ell-1} \mathcal{O}(1)\right)$ operations.
- Next, for every $G \subset \Omega$ we consider the products $A_{G,\,c,\,k-1} := \widehat{A}_{\text{int},\,\text{ext}} = A_{\text{int},\,\text{ext}} Q_{\text{ext}}$ which are equal to
$$B_{\text{int},\,\text{ext}} Q_{\text{ext}} - \sum_{j\geq k}^{\ell-1} A_{\text{int},\,c,\,j} (A_{c,\,j})^{-1} A_{c,\,j,\,\text{ext}} Q_{\text{ext}}.$$
Here "ext" stands for the dofs in $\Omega \setminus G$ and "int" stands for the dofs in G, respectively. They can be computed in a number of operations of order
$$\mathcal{N}(G)\left(\mathcal{O}(1) + \sum_{j\geq k}^{\ell-1} \mathcal{O}(1)\right).$$
The blocks $A_{c,\,k-1,\,G} := \widehat{A}_{\text{ext},\,\text{int}} = Q_{\text{ext}}^T A_{\text{ext},\,\text{int}}$ which equal
$$Q_{\text{ext}}^T B_{\text{ext},\,\text{int}} - \sum_{j\geq k}^{\ell-1} Q_{\text{ext}}^T A_{\text{ext},\,c,\,j} (A_{c,\,j})^{-1} A_{c,\,j,\,\text{int}},$$
are similarly computed in
$$\mathcal{N}(G)\left(\mathcal{O}(1) + \sum_{j\geq k}^{\ell-1} \mathcal{O}(1)\right)$$
operations. Also, if the numbers of subneighborhoods G contained in a coarse neighborhood Ω stays bounded at all levels k, $1 \leq k \leq \ell$, then the total cost of this step is readily estimated by
$$\mathcal{N}(\Omega)\left(\sum_{j\geq k-1}^{\ell-1} \mathcal{O}(1)\right).$$
- The most expensive block of $\widehat{A}_G$ is $A_{c,\,k-1} := \widehat{A}_{\text{ext},\,\text{ext}} = Q_{\text{ext}}^T A_{\text{ext},\,\text{ext}} Q_{\text{ext}}$, which we compute as described previously (cf. 2.4) by first computing
$$A_{c,\,k-1,\,\text{int},\,\text{int}} = Q_{\text{int}}^T \left(B_{\text{int},\,\text{int}} - \sum_{j\geq k}^{\ell-1} A_{G,\,c,\,j} (A_{c,\,j})^{-1} A_{c,\,j,\,G}\right) Q_{\text{int}},$$

which cost is readily estimated by

$$\mathcal{N}(G) \sum_{j \geq k-1}^{\ell-1} \mathcal{O}(1).$$

Then, using the already computed quantities $A_{G,\,c,\,k-1}$ and $A_{c,\,k-1,\,G}$ (cf. 2.2 and 2.3) the calculation of $Q_{\rm int}^T A_{G,\,c,\,k-1}$ and $A_{c,\,k-1,\,G} Q_{\rm int}$ can be performed in $\mathcal{N}(G) \times \mathcal{O}(1)$ operations. That is, summing over all Gs contained in Ω the actual cost in this step is proportional to

$$\mathcal{N}(\Omega) \left(\sum_{j \geq k-1}^{\ell-1} \mathcal{O}(1) \right).$$

- Finally, the form of the new Schur complements A_G in terms of a sparse matrix B_G plus low rank updates, requires one additional inversion of a coarse matrix $A_{c,\,k-1}$ which has size of order $\mathcal{O}(1)$. Thus, one gets that the new Schur complement A_G has the same form as A_Ω now with one extra low rank update,

$$\begin{aligned} A_G &= A_{\Omega,\,\mathrm{int},\,\mathrm{int}} - A_{G,\,c,\,k-1}(A_{c,\,k-1})^{-1} A_{c,\,k-1,\,G} \\ &= B_{\mathrm{int},\,\mathrm{int}} - \sum_{j \geq k}^{\ell-1} A_{G,\,c,\,j}(A_{c,\,j})^{-1} A_{c,\,j,\,G} - A_{G,\,c,\,k-1}(A_{c,\,k-1})^{-1} A_{c,\,k-1,\,G} \\ &= B_G - \sum_{j \geq k-1}^{\ell-1} A_{G,\,c,\,j}(A_{c,\,j})^{-1} A_{c,\,j,\,G}, \end{aligned}$$

where

$$B_G = B_\Omega|_G\,,\ A_{G,\,c,\,j} = A_{\mathrm{int},\,c,\,j},\ A_{c,\,j,\,G} = A_{\,c,\,j,\,\mathrm{int}},\ j \geq k.$$

We summarize the above estimates in the following theorem.

Theorem 4.1. *If the set $\{P_k : D_k \to D_0, 1 \leq k \leq \ell\}$ of global interpolation matrices satisfies the localization property of Definition 4.1 all Schur complements at levels $k = \ell - 1, \ldots, 0$ can be computed in a total cost of order*

$$\mathcal{O}(N_0 \ell^2), \quad \textit{where } N_0 = \mathcal{N}(D_0).$$

Remark 4.1. *If one assumes more regular structure of the neighborhoods, in particular, that each Ω can be split in a boundary $\partial\Omega$ and an interior set, such that all low rank updates have actually the size of the boundary of Ω then the cost and storage of the above algorithm can substantially be reduced. For example, the low rank updates can be stored in only one block of size $\mathcal{N}(\partial\Omega) \times \mathcal{N}(\partial\Omega)$ and correspondingly the cost of computing the next fine–grid Schur complements A_G for all $G \subset \Omega$ will be reduced to $\mathcal{N}(\Omega) + (\mathcal{N}(\partial\Omega))^2$. For two–dimensional neighborhoods, it is reasonable to assume that $(\mathcal{N}(\partial\Omega))^2 \simeq \mathcal{N}(\Omega)$. Then the cost at every level $k:\ 0 < k \leq \ell$, stays proportional to $\sum_{\Omega \in \mathcal{T}_k} \mathcal{N}(\Omega) = \mathcal{O}(N_0)$ and therefore the total cost reduces to $\mathcal{O}(\ell N_0)$. The storage is readily seen to be $\mathcal{O}(N_0)$.*

5. Spectral properties

We will now study, for the case of a symmetric positive (semi–) definite fine grid matrix A, arising from some Ritz–Galerkin type finite element discretization, the relationship between the multilevel Schur complements A_Ω computed via Algorithm 3.1 and the neighborhood matrices A_Ω^N assembled from the corresponding *true* element matrices, i.e.,

$$A_\Omega^N := \sum_{e\subset\Omega} A_e^N. \tag{5.1}$$

That is, we consider a set of elements $E = \{e\}$, e.g., triangles, and a set of element matrices $\{A_e^N : e \in E\}$. Note that in general E will provide an overlapping partition of the set of fine grid dofs D_0. However, for neighborhoods considered as domains, E provides a non–overlapping partition, which means that no shared elements are allowed for any two neighborhoods in a neighborhood set.

Moreover, let $\mathcal{T}_0, \mathcal{T}_1, \ldots, \mathcal{T}_\ell$ be a sequence of nested neighborhood sets where $\Omega_k \in \mathcal{T}_k$ denotes an arbitrary neighborhood at some level k, $0 \le k \le \ell$. Then, throughout the rest of this section, by $\Omega_j \in \mathcal{T}_j$ we will denote the coarse–level neighborhood that contains Ω_k at level $j > k$, i.e., $\Omega_k \subset \Omega_{k+1} \subset \ldots \subset \Omega_\ell$.

Note that

$$A = A^N = \sum_{\Omega_k\in\mathcal{T}_k} A_{\Omega_k}^N$$

for all $k = 0, 1, \ldots, \ell$ (where no essential boundary conditions have been imposed yet); in particular, we have

$$A = A_{\Omega_\ell} = A_{\Omega_\ell}^N \tag{5.2}$$

(by definition). Assume now that the relation

$$\mathbf{v}_{\Omega_j}^T A_{\Omega_j}\mathbf{v}_{\Omega_j} \ge \mathbf{v}_{\Omega_j}^T A_{\Omega_j}^N \mathbf{v}_{\Omega_j} \quad \forall \mathbf{v}_{\Omega_j} := \mathbf{v}|_{\Omega_j} \tag{5.3}$$

holds for some neighborhood $\Omega_j \in \mathcal{T}_j$ at level j and this is true for $j = \ell$ because of 5.2. Then, for any subneighborhood $\Omega_{j-1} \subset \Omega_j$, $\Omega_{j-1} \in \mathcal{T}_{j-1}$ it follows from Lemma 2.1 that

$$\begin{aligned}
\mathbf{v}_{\Omega_{j-1}}^T A_{\Omega_{j-1}}\mathbf{v}_{\Omega_{j-1}} &= \inf_{\mathbf{v}_c} \begin{bmatrix} \mathbf{v}_{\Omega_{j-1}} \\ P_{\text{ext}}\mathbf{v}_c \end{bmatrix}^T A_{\Omega_j} \begin{bmatrix} \mathbf{v}_{\Omega_{j-1}} \\ P_{\text{ext}}\mathbf{v}_c \end{bmatrix} \\
&\ge \inf_{\mathbf{v}_c} \begin{bmatrix} \mathbf{v}_{\Omega_{j-1}} \\ P_{\text{ext}}\mathbf{v}_c \end{bmatrix}^T A_{\Omega_j}^N \begin{bmatrix} \mathbf{v}_{\Omega_{j-1}} \\ P_{\text{ext}}\mathbf{v}_c \end{bmatrix} \\
&\ge \inf_{\mathbf{v}_c} \begin{bmatrix} \mathbf{v}_{\Omega_{j-1}} \\ P_{\text{ext}}\mathbf{v}_c \end{bmatrix}^T \sum_{e\subset\Omega_{j-1}} A_e^N \begin{bmatrix} \mathbf{v}_{\Omega_{j-1}} \\ P_{\text{ext}}\mathbf{v}_c \end{bmatrix} \\
&= \mathbf{v}_{\Omega_{j-1}}^T A_{\Omega_{j-1}}^N \mathbf{v}_{\Omega_{j-1}} \quad \forall \mathbf{v}_{\Omega_{j-1}} := \mathbf{v}|_{\Omega_{j-1}}
\end{aligned}$$

where the last inequality holds since the individual element matrices are symmetric positive semi–definite (for the considered class of problems). Applying this estimate recursively, we conclude

$$\mathbf{v}_{\Omega_k}^T A_{\Omega_k}\mathbf{v}_{\Omega_k} \ge \mathbf{v}_{\Omega_k}^T A_{\Omega_k}^N \mathbf{v}_{\Omega_k} \quad \forall \Omega_k \in \mathcal{T}_k, \quad k \in \{0, 1, \ldots, \ell\}. \tag{5.4}$$

However, a uniform upper bound on the quadratic form associated with any multilevel Schur complement A_{Ω_k} in general involves assumptions on the multilevel interpolation matrices P_j, $j > k$. It makes sense to consider the multilevel version of the space defined by 2.6, that is, we define the vector space

$$V^c_{\Omega_k} = \left\{ \begin{bmatrix} \mathbf{v}_{\Omega_k} \\ P_{\Omega_{k+1},\,\mathrm{ext}}\mathbf{v}_{\Omega_{k+1},\,c} \\ \vdots \\ P_{\Omega_\ell,\,\mathrm{ext}}\mathbf{v}_{\Omega_\ell,\,c} \end{bmatrix} \right\}, \tag{5.5}$$

where

$$P_{\Omega_j,\,\mathrm{ext}} := P_j|_{\Omega_j\setminus\Omega_{j-1}}$$

and $\mathbf{v}_{\Omega_j,\,c}$ is an arbitrary coarse vector defined on the jth level coarse dofs of Ω_j. Then, analogous to the proof of Lemma 2.1 one can show that the multilevel Schur complement A_{Ω_k} gives rise to the minimal energy extension in the space 5.5. That is,

$$\mathbf{v}^T_{\Omega_k} A_{\Omega_k}\mathbf{v}_{\Omega_k} = \inf_{\substack{\mathbf{w}\,\in\, V^c_{\Omega_k} \\ \mathbf{w}|_{\Omega_k} = \mathbf{v}_{\Omega_k}}} \mathbf{w}^T A_{\Omega_\ell}\mathbf{w}. \tag{5.6}$$

Hence, using 5.2, and splitting the quatratic form into the part corresponding to elements contained in Ω_k and the rest, we obtain

$$\mathbf{v}^T_{\Omega_k} A_{\Omega_k}\mathbf{v}_{\Omega_k} = \mathbf{v}^T_{\Omega_k} A^N_{\Omega_k}\mathbf{v}_{\Omega_k} + \inf_{\substack{\mathbf{w}\,\in\, V^c_{\Omega_k} \\ \mathbf{w}|_{\Omega_k} = \mathbf{v}_{\Omega_k}}} \mathbf{w}^T A^N_{\Omega_\ell\setminus\Omega_k}\mathbf{w} \quad \text{for all} \quad \mathbf{v}_{\Omega_k} := \mathbf{v}|_{\Omega_k}, \tag{5.7}$$

where $A^N_{\Omega_\ell\setminus\Omega_k} := A^N_{\Omega_\ell} - A^N_{\Omega_k}$. We summarize the above estimates in the following theorem.

Theorem 5.1. *For any sequence of nested neighborhood sets $\mathcal{T}_0, \mathcal{T}_1, \ldots, \mathcal{T}_\ell = \{\Omega_\ell\}$ and any symmetric positive semi–definite finite element stiffness matrix $A = A_{\Omega_\ell} := \sum_{e\subset\Omega_\ell} A^N_e$ we have that the multilevel Schur complements A_{Ω_k} satisfy*

$$\mathbf{v}^T_{\Omega_k} A_{\Omega_k}\mathbf{v}_{\Omega_k} \geq \mathbf{v}^T_{\Omega_k} A^N_{\Omega_k}\mathbf{v}_{\Omega_k} \quad \textit{for all} \quad \mathbf{v}_{\Omega_k} := \mathbf{v}|_{\Omega_k},$$

where $\Omega_k \in \mathcal{T}_k$ is an arbitrary neighborhood at some level $k \in \{0, 1, \ldots, \ell\}$ and $A^N_{\Omega_k}$ is defined by 5.1.

Moreover, if for any vector $\mathbf{v}_{\Omega_k}$ there exists a vector $\mathbf{w}$, $\mathbf{w}|_{\Omega_k} = \mathbf{v}_{\Omega_k}$, in the space $V^c_{\Omega_k}$ defined by 5.5 such that

$$\mathbf{w}^T A^N_{\Omega_\ell\setminus\Omega_k}\mathbf{w} \leq c \cdot \mathbf{v}^T_{\Omega_k} A^N_{\Omega_k}\mathbf{v}_{\Omega_k}$$

then the upper bound

$$\mathbf{v}^T_{\Omega_k} A_{\Omega_k}\mathbf{v}_{\Omega_k} \leq (1 + c) \cdot \mathbf{v}^T_{\Omega_k} A^N_{\Omega_k}\mathbf{v}_{\Omega_k} \quad \forall \mathbf{v}_{\Omega_k} := \mathbf{v}|_{\Omega_k} \tag{5.8}$$

holds as well.

We close this section by considering an example for which we prove a simple spectral equivalence result.

Corollary 5.1. *Consider the two-dimensional model problem where A corresponds to the Laplace operator discretized by linear finite elements on a quasiuniform triangular mesh. Then, choosing the multilevel interpolation operators P_k ($k \geq 1$) as (piecewise) constant interpolation based on simple averaging, the multilevel Schur complements $A_\Omega = A_{\Omega_0}$ computed via Algorithm 3.1 are spectrally equivalent to the corresponding neighborhood matrices $A_\Omega^N = A_{\Omega_0}^N$ defined via 5.1.The constant c in the upper bound 5.8 does not depend on the number of levels ℓ in this case. Here we assume that the size of the fine grid neighborhoods is uniformly bounded, i.e., $\mathcal{N}(\Omega) = \mathcal{O}(1)$.*

Proof. We note that under the assumption of (piecewise) constant interpolation every vector $\mathbf{v}$ in the space $V_{\Omega_0}^c$ is constant on the sets $\Omega_k \setminus \Omega_{k-1}$, $k \in \{1, 2, \ldots, \ell\}$, i.e., $\mathbf{v}$ has the representation

$$\mathbf{v} = \begin{bmatrix} \mathbf{v}_{\Omega_0} \\ \mathbf{c}_{\Omega_1,\,\mathrm{ext}} \\ \vdots \\ \mathbf{c}_{\Omega_\ell,\,\mathrm{ext}} \end{bmatrix}$$

with some constant vectors $\mathbf{c}_{\Omega_1,\,\mathrm{ext}}, \ldots, \mathbf{c}_{\Omega_\ell,\,\mathrm{ext}}$. Now, for any vector $\mathbf{v}_{\Omega_0}$ we consider the extension

$$\bar{\mathbf{w}} = \bar{E}\mathbf{v}_{\Omega_0} = \begin{bmatrix} \mathbf{v}_{\Omega_0} \\ \bar{\mathbf{v}}_{\Omega_0} \end{bmatrix} \in V_{\Omega_0}^c, \quad \bar{\mathbf{v}}_{\Omega_0} = \begin{bmatrix} \bar{v} \\ \bar{v} \\ \vdots \\ \bar{v} \end{bmatrix}$$

that is constant outside Ω_0 where the constant $\bar{v}$ is chosen to be the average value of $\mathbf{v}_{\Omega_0}$ over Ω_0, i.e., $\bar{v} = (\sum_{i\in\Omega_0} v_i)/\mathcal{N}(\Omega_0)$. Then, with this particular choice of $\bar{\mathbf{w}}$ one gets

$$\begin{aligned}
\inf_{\substack{\mathbf{w}\,\in\, V_{\Omega_0}^c \\ \mathbf{w}|_{\Omega_0} = \mathbf{v}_{\Omega_0}}} \mathbf{w}^T A_{\Omega_\ell\setminus\Omega_0}^N \mathbf{w} \;&\leq\; \bar{\mathbf{w}}^T A_{\Omega_\ell\setminus\Omega_0}^N \bar{\mathbf{w}} \\
&=\; \sum_{e\not\subset\Omega_0, e\cap\Omega_0\neq\emptyset} \int_e \nabla w \cdot \nabla w \\
&\leq\; c_1 \cdot \sum_{i\in\Omega_0} (\bar{v} - v_i)^2 \\
&\leq\; c_1 \cdot \sum_{i,j\in\Omega_0} (v_i - v_j)^2 \\
&\leq\; c_2 \cdot \mathbf{v}_{\Omega_0}^T A_{\Omega_0}^N \mathbf{v}_{\Omega_0}
\end{aligned}$$

where for a quasiuniform triangulation c_1 and c_2 are uniformly bounded constants. □

6. Numerical experiments

In the numerical experiments presented in this section we used the agglomeration algorithm from [10] in order to provide the nested neighborhood sets. The simple initial multilevel interpolation operators $P_k : D_k \to D_0$, involved in the construction of the two-level neighborhood matrices, were taken to be piecewise constant interpolation based on averaging. As already mentioned in section 4 the crucial point for the

complexity issue of our algorithm is that all dofs from any coarse-level neighborhood Ω are interpolated from $\mathcal{O}(1)$ coarse dofs. This assumption (c.f. Theorem 4.1) obviously can be met if the number $\mathcal{N}(\Omega \cap D_k)$ of coarse dofs contained in a coarse-level neighborhood Ω is bounded by some constant. The following Table 1 lists $\mathcal{N}(\mathcal{T}_k)$, the number of coarse-level neighborhoods at level k,

$$n_{\max} := \max_{\Omega \in \mathcal{T}_k} \left(\sum_{G \subset \Omega} 1 \right), \quad \text{and} \quad n_{\text{ave}} := \frac{\sum_{\Omega \in \mathcal{T}_k} \left(\sum_{G \subset \Omega} 1 \right)}{\mathcal{N}(\mathcal{T}_k)},$$

the maximal and average number of fine-level neighborhoods G that are contained a coarse-level neighborhood Ω, as well as,

$$c_{\max} := \max_{\Omega \in \mathcal{T}_k} \left(\mathcal{N}(\Omega \cap D_k) \right), \quad \text{and} \quad c_{\text{ave}} := \frac{\sum_{\Omega \in \mathcal{T}_k} \mathcal{N}(\Omega \cap D_k)}{\mathcal{N}(\mathcal{T}_k)},$$

the maximal and average number of coarse dofs that are contained in a coarse–level neighborhood Ω. Here the fine grid was an unstructured triangular mesh with 25600 elements and 6291 fine-level neighborhoods G at level zero. We observe that the number of coarse dofs that are used for interpolation in any neighborhood Ω is bounded uniformly in the levels.

TABLE 1. Number of fine-level neighborhoods and coarse dofs contained in a coarse-level neighborhood

level k	$\mathcal{N}(\mathcal{T}_k)$	$n_{\max}$	n_{ave}	$c_{\max}$	c_{ave}
1	1571	10	4.0	14	6.4
2	422	8	3.7	17	7.2
3	127	7	3.3	16	8.0
4	49	4	2.6	12	7.6
5	21	4	2.3	16	6.6
6	10	3	2.1	9	5.8
7	5	2	2.0	6	4.2
8	2	3	2.5	3	3.0

Then, it seems likely to study the convergence of the AMGe method, replacing in the interpolation setup the neighborhood matrices assembled from true element stiffness matrices—to which we will refer as natural neighborhood matrices—with the artificial neighborhood matrices computed via multilevel Schur complements.

Note that assembling the artificial neighborhood matrices globally, yields an auxiliary problem that is spectrally equivalent to the original problem under reasonable assumptions. We showed this in the previous section (cf. Corollary 5.1) for the model problem corresponding to the Laplace operator.

In our numerical tests we used the artificial neighborhood matrices as input for the agglomeration-based AMGe method as it is described in [10]. That is, in a preprocessing step we computed a sequence of (improved) single–level interpolation matrices $P_k^\star : D_k \to D_{k-1}, 1 \leq k \leq \ell$ based on the artificial neighborhood matrices. The coarse grid operators A_k were computed via the standard Galerkin approach, i.e.,

$A_k = (P_k^\star)^T A_{k-1} P_k^\star$ where A_0 was the matrix corresponding to the original problem (with essential boundary conditions). The resulting AMG method is denoted by AMGe* in the following.

Our test problem was the second order elliptic PDE

$$-\nabla \cdot \left(\epsilon \cdot I + b \cdot b^T\right) \nabla u = f \quad \text{on } \Omega, \tag{6.1}$$

$$u = 0 \quad \text{on } \partial\Omega, \tag{6.2}$$

$\Omega = (0,1) \times (0,1)$, and $b = (\cos\theta, \sin\theta)^T$.

For discretization we used linear finite elements on a sequence of unstructured triangular meshes where the dimension of the FE space was increased. The results presented below are for $\theta = 30°$, which yields a direction of strongest couplings that is not aligned with the grid points. The convergence rates are almost the same if we have a spatial variability or a jump discontinuity in the angle θ, so we did not list them separately. However, the factor ϵ is a critical problem parameter if the coarse–grid selection is not adapted to the anisotropy.

The problem was solved using a preconditioned conjugate gradient method, where the preconditioning consisted of a single V(1,1)-cycle of AMG, with a Gauß-Seidel smoother. We took a random right-hand side and the iteration was initialized with the zero start vector. The stopping criterion was a reduction of the residual norm by a factor 10^{-6}.

In a first experiment we compared the AMGe method using either the natural or the artificial neighborhood matrices for building the prolongation. Table 2 summarizes the number of preconditioned CG iterations to achieve the desired residual size and ρ, the average convergence factor over the iterations for different problem size. We also report the grid- and operator complexity. We observe that the convergence results are even better for the artificial neighborhood matrices although interpolation is based on the auxiliary problem in this case.

In a second experiment, regarding the same model problem, we determined the two-level convergence rate and discovered the effect of a more sophisticated smoother, i.e., a multiplicative Schwarz method, which we used in the multilevel method. Once more, we compared the performance of AMGe to that of AMGe*. As Table 3 shows, the results for the two-level method with interpolation based on the artificial neighborhood matrices are very close to those for the multilevel method. This indicates that there is not much room to improve this element-free interpolation. The fact that the (average) two-level convergence factors are sometimes even slightly worse than in the multilevel case can be explained by a non-monotonic convergence of the preconditioned conjugate gradient method. Nevertheless, one way to improve the results (for small values of ϵ) would be a problem oriented coarse-grid selection. Another possibility is the choice of a more powerful smoother, for example, the multiplicative Schwarz method. As the results in the two right-most columns of Table 3 indicate, the use of a Schwarz smoother improves AMGe convergence for this type of problem significantly, the use of the artificial neighborhood matrices, however, even gives an additional benefit.

TABLE 2. Convergence results for different problem size

			400 elts 5 levels	1600 elts 6 levels	6400 elts 8 levels	25600 elts 9 levels
$\epsilon = 1$	AMGe*	iterations	6	8	10	11
		ρ	0.087	0.168	0.218	0.274
	AMGe	iterations	6	8	10	11
		ρ	0.082	0.165	0.231	0.276
$\epsilon = 0.1$	AMGe*	iterations	11	14	19	24
		ρ	0.265	0.368	0.478	0.556
	AMGe	iterations	11	15	20	26
		ρ	0.255	0.375	0.488	0.577
$\epsilon = 0.01$	AMGe*	iterations	12	16	25	35
		ρ	0.299	0.413	0.570	0.671
	AMGe	iterations	12	18	29	41
		ρ	0.286	0.445	0.617	0.711
grid complexity			1.43	1.55	1.54	1.42
operator complexity			1.48	1.73	1.71	1.49

TABLE 3. Two-level convergence and multilevel results for Schwarz smoothing

25600 elements 9 levels		2-level: GS smoother AMGe*	2-level: GS smoother AMGe	multilevel: S smoother AMGe*	multilevel: S smoother AMGe
$\epsilon = 1$	iterations	8	7	6	7
	ρ	0.138	0.131	0.079	0.105
$\epsilon = 0.1$	iterations	24	21	13	13
	ρ	0.550	0.509	0.326	0.337
$\epsilon = 0.01$	iterations	37	35	17	18
	ρ	0.688	0.670	0.431	0.463

Acknowledgements. The author is very grateful to Panayot Vassilevski, who initiated the consideration of the presented algorithm, and acknowledges the stimulating discussions on the AMGe topic he also had with Van Emden Henson and other researchers from Lawrence Livermore National Laboratory.

REFERENCES

[1] A. BRANDT, *Algebraic multigrid theory: The symmetric case*, in Preliminary Proceedings for the International Multigrid Conference, Copper Mountain, Colerado, 1983.

[2] ——, *Algebraic multigrid (AMG) for sparse matrix equations*, in Sparsity and Its Applications, D.J. Evans, ed., Cambridge: Cambridge University Press, 1984.

[3] ——, *Algebraic multigrid theory: The symmetric case*, Appl. Math. Comput., 19 (1986), pp. 23–56.

[4] A. BRANDT, S. MCCORMICK, AND J. RUGE, *Algebraic multigrid (AMG) for automatic multigrid solutions with applications to geodetic computations*, Report, Inst. for Computational Studies, Fort Collins, Colorado, 1982.

[5] ——, *Algebraic multigrid (AMG) for sparse matrix equations*, in Sparsity and Its Applications, D.J. Evans, ed., Cambridge University Press, Cambridge, 1984.
[6] M. Brezina, A. Cleary, R. Falgout, V. Henson, J. Jones, T. Manteuffel, S. McCormick, and J. Ruge, *Algebraic multigrid based on element interpolation AMGe*, SIAM J. Sci. Comput., 22 (2000), pp. 1570–1592.
[7] T. Chartier, R. Falgout, V. Henson, J. Jones, T. Manteuffel, S. McCormick, J. Ruge, and P. Vassilevski, *Spectral AMGe (ρAMGe)*, to appear in SIAM J. Sci. Comput.
[8] A. Cleary, R. Falgout, V. Henson, J. Jones, T. Manteuffel, S. McCormick, G. Miranda, and J. Ruge, *Robustness and scalability of algebraic multigrid*, SIAM J. Sci. Stat. Comput., 21 (2000), pp. 1886–1908.
[9] V. Henson and P. Vassilevski, *Element-free AMGe: General algorithms for computing the interpolation weights in AMG*, SIAM J. Sci. Comput., 23 (2001), pp. 629–650.
[10] J. Jones and P. Vassilevski, *AMGe based on element agglomeration*, SIAM J. Sci. Comput., 23 (2001), pp. 109–133.
[11] J. Mandel, M. Brezina, and P. Vaněk, *Energy optimization of algebraic multigrid bases*, Computing, 62 (1999), pp. 205–228.
[12] J. Ruge and K. Stüben, *Efficient solution of finite difference and finite element equations by algebraic multigrid (AMG)*, in Multigrid Methods for Integral and Differential Equations, D.J. Paddon and H. Holstein, eds., The Institute of Mathematics and Its Applications Conference Series, Clarendon Press, Oxford, 1985, pp. 169–212.
[13] ——, *Algebraic multigrid (AMG)*, in Multigrid Methods, S.F. McCormick, ed., vol. 3 of Frontiers in Applied Mathematics, SIAM, Philadelphia, PA, 1987, pp. 73–130.
[14] K. Stüben, *Algebraic multigrid (AMG): Experiences and comparisons*, Appl. Math. Comput., 13 (1983), pp. 419–452.
[15] R. Tuminaro and C. Tong, *Parallel smoothed aggregation multigrid: Aggregation strategies on massively parallel machines*, Report, Sandia National Laboratories, 2000.
[16] P. Vaněk, M. Brezina, and J. Mandel, *Convergence of algebraic multigrid based on smoothed aggregation*, Numer. Math., 88 (2001), pp. 559–579.
[17] P. Vaněk, J. Mandel, and M. Brezina, *Algebraic multigrid based on smoothed aggregation for second and fourth order problems*, Computing, 56 (1996), pp. 179–196.
[18] W. Wan, T. Chan, and B. Smith, *An energy-minimizing interpolation for robust multigrid methods*, Siam J. Sci. Comput., 21 (2000), pp. 1632–1649.

Johann Radon Institute for Computational and Applied Mathematics, Austrian Academy of Sciences, Altenbergerstr. 69, A-4040 Linz, Austria
E-mail address: johannes.kraus@oeaw.ac.at