Computers and Mathematics with Applications 144 (2023) 290–305

Contents lists available at ScienceDirect

Computers and Mathematics with Applications

journal homepage: www.elsevier.com/locate/camwa

Automatic coarsening in Algebraic Multigrid utilizing quality measures for matching-based aggregations

Pasqua D'Ambra [a], Fabio Durastante [c,a,*], Salvatore Filippone [b,a], Ludmil Zikatanov [d]

[a] Institute for Applied Computing "Mauro Picone" (IAC), Consiglio Nazionale delle Ricerche, Via Pietro Castellino 111, Naples, 80131, NA, Italy
[b] Department of Civil and Computer Engineering, University of Rome "Tor Vergata", Via Politecnico 1, Rome, 00133, RM, Italy
[c] Department of Mathematics, University of Pisa, Largo Bruno Pontecorvo, 5, Pisa, 56127, PI, Italy
[d] Department of Mathematics, The Pennsylvania State University, McAllister Building, Pollock Rd, State College, 16802, PA, USA

ARTICLE INFO

Keywords:
AMG
Convergence
Graph matching
Aggregation
Compatible relaxation

ABSTRACT

In this paper, we discuss the convergence of an Algebraic MultiGrid (AMG) method for general symmetric positive-definite matrices. The method relies on an aggregation algorithm, named *coarsening based on compatible weighted matching*, which exploits the interplay between the principle of compatible relaxation and the maximum product matching in undirected weighted graphs. The results are based on a general convergence analysis theory applied to the class of AMG methods employing unsmoothed aggregation and identifying a quality measure for the coarsening; similar quality measures were originally introduced and applied to other methods as tools to obtain good quality aggregates leading to optimal convergence for M-matrices. The analysis, as well as the coarsening procedure, is purely algebraic and, in our case, allows an *a posteriori* evaluation of the quality of the aggregation procedure which we apply to analyze the impact of approximate algorithms for matching computation and the definition of graph edge weights. We also explore the connection between the choice of the aggregates and the compatible relaxation convergence, confirming the consistency between theories for designing coarsening procedures in purely algebraic multigrid methods and the effectiveness of the coarsening based on compatible weighted matching. We discuss various completely automatic algorithmic approaches to obtain aggregates for which good convergence properties are achieved on various test cases.

1. Introduction

We assess here the convergence of a MultiGrid method (MG) for the solution of linear systems of the form

$$A\mathbf{u} = \mathbf{f}, \tag{1}$$

on the finite-dimensional linear vector space V equipped with an inner product $(\cdot,\cdot)$, where $A : V \to V'$ is symmetric positive definite (SPD), $\mathbf{f} \in V'$ and V' is the dual of V; by the Riesz representation theorem V' can be identified with V. More specifically, we focus on a recently proposed method belonging to the class of Algebraic MultiGrid Methods (AMG) with *unsmoothed aggregation* (UA-AMG) or *plain aggregation* [1,2]. These can be seen as particular instances of a general stationary linear iterative method for solving (1)

$$\mathbf{u}^m = \mathbf{u}^{m-1} + B(\mathbf{f} - A\mathbf{u}^{m-1}), \quad m = 1, 2, \dots; \quad \text{given } \mathbf{u}^0 \in V, \tag{2}$$

where $B : V' \to V$ is a linear operator which can be interpreted as an approximate inverse of A. An AMG method, or indeed any MG, is based on the recursive use of a two-grid scheme combining the action of a *smoother*, i.e., a convergent iterative method, and a *coarse-grid correction*, which corresponds to the solution of the residual equations on a coarser grid. In completely general terms, the guiding design principle of an AMG is the optimization of the choice of coarse space for a given smoother. The most commonly used smoothers are the splitting-based methods, such as the Gauss–Seidel method and the (modified or scaled) Jacobi method.

As usual in the MG context, the final objective of any analysis is to achieve uniform convergence with respect to the problem size (optimal convergence). Unfortunately, this is a property that can normally be established only for the two-level AMG (TL-AMG); it is very rarely extended to the multilevel case when no "geometric" information on the matrix A is available. Our task is then to ensure the selection of an appropriate set of aggregates, i.e., the disjoint sets of fine grid un-

* Corresponding author.
E-mail address: fabio.durastante@unipi.it (F. Durastante).

https://doi.org/10.1016/j.camwa.2023.06.026
Received 21 November 2022; Received in revised form 6 June 2023; Accepted 18 June 2023
Available online 14 July 2023

knowns to which the coarse grid unknowns are associated, to guarantee a *fast* convergence at a reasonable cost per iteration. Of the many possible ways of achieving such a result, we narrow down our investigation to the case of UA-AMG; see [3,4] for the first works in this direction. Within this framework, we are going to exploit the unifying theory outlined in the review [2] to assess convergence and to investigate and characterize the quality of the coarse spaces generated by means of the aggregation procedure introduced in [5,6]. The latter is a technique based on the use of matching algorithms for edge-weighted graphs [7–9] that aims to achieve a purely algebraic and automatic approach for the solution of (1), with no further assumption on the SPD system matrix, and independently of any user defined parameter. Indeed, this approach fits within a trend of similar algebraic techniques, e.g., those based on path-covering algorithms [10], or on the use of matching to generate multilevel hierarchies for graph Laplacians relative to coarse subspaces in finite elements applications [11], striving for purely algebraic aggregation procedures that are adaptive in nature and allow for an *a posteriori* analysis of the quality of the generated coarse spaces.

We observe that, as reported in [2, Section 8.5, Section 9.5], the general convergence theory we specialized in this paper for the aggregation based on matching in weighted graphs, was originally designed for the AGMG method in [1,12] and extended in [13,14], to obtain AMG methods based on unsmoothed aggregation with a user-defined bound on the convergence rate. In [13] the authors show that, for the class of nonsingular symmetric M-matrices with nonnegative row sum, if the aggregates can be built in such a way that a meaningful local bound is fulfilled, the resulting multilevel methods employing an appropriate AMLI cycle [15] shows an optimal convergence with a guaranteed convergence rate. The theory is extended to nonsymmetric M-matrices for a TL-AMG in [14]. In [2] the theory is again extended to more general SPD matrices and formalized as an abstract framework for the setup of coarsening methods.

We finally note that the need to define local measures to assess the quality of a coarse space also led to the introduction of the notion of *compatible relaxation*. Compatible relaxation, first defined by Brandt in [16], as *a modified relaxation scheme that keeps the coarse-level variables invariant*, was originally based on the idea to use a smoother to detect slowly converging components. This principle has been widely applied to define a general procedure for coarsening, both for selecting coarse variables and to adapt the prolongators in adaptive AMG (see, e.g., [17–20]). It was a basic guideline for the formulation of our coarsening and of its application in a bootstrap AMG based on composition of multiple AMG hierarchies [5,6]. In our coarsening method, since we explicitly define the complementary space to the coarse space, we can apply a smoother to the only-fine variables and infer the quality of the coarse space by an estimate of the corresponding convergence rate. Our experiments show the coherency between the aggregation quality measure based on the general theory in [2], which has the advantage to be independent of the smoother and only depends on the way we build aggregates, and the quality measure based on the compatible relaxation.

The main contributions of this paper can be summarized as follows.

- We prove that the automatic aggregation-based coarsening, relying on maximum weight matching in graphs equipped with a suitable choice of edge weights, fulfills all the conditions to have a bounded convergence rate of the corresponding TL-AMG for any SPD matrix.
- We show how the resulting quality measure for the aggregation can be used to drive the choice of different (approximate) matching algorithms and of the edge weights in the adjacency graph of the system matrix, without resorting to heuristics and *a priori* information on the near kernel of the matrix.
- We emphasize the connection between the choice of the aggregates and the compatible relaxation principle for the new coarsening, confirming the consistency between the currently available theories for general coarsening in AMG.

The remainder of this paper is organized as follows: to begin with, in Section 2 we introduce a quality measure for a general UA-AMG in terms of the unifying theory from [2]. Then, in Section 3 we reintroduce the UA-AMG from [5,6] and specialize the convergence theory and the quality measure for the aggregates from the previous section to this case. Section 4 is entirely devoted to the application of the theory to some standard benchmarks; specifically, we investigate how the various matching algorithms applied for obtaining the aggregates influence their quality. Section 5 shows the coherency between the quality of aggregates and the convergence ratio of a convergent smoother applied to the effective smoother space, i.e., to the complementary space to the coarse space. Section 6 summarizes conclusions.

2. Convergence theory for TL-AMG algorithms and quality measure for aggregates

The measure of the quality of the aggregates, and thus of the coarse space, for a given TL-AMG algorithm we are interested in depends both on the convergence ratio achieved by the resulting method and on the cost needed for defining and applying the multigrid hierarchy. To set the notation, and the context in which we are performing our analysis, let us briefly recall the components of a TL-AMG method, i.e.:

- a *convergent smoother*, $R : V' \to V$;
- a *coarse space* V_c; this is either a subspace of V or more generally a space with a smaller dimension than V. It is always linked to V via a prolongation operator $P : V_c \to V$;
- a *coarse space solver*, $B_c : V_c' \to V_c$;

and how these components are related to its convergence properties. We follow the approach discussed in [2] that permits to analyze the convergence properties of a multigrid algorithm in a general way. To this end, we need to introduce the inner product

$$(\mathbf{u},\mathbf{v})_{\overline{R}^{-1}} = (\overline{T}^{-1}\mathbf{u},\mathbf{v})_A = (\overline{R}^{-1}\mathbf{u},\mathbf{v}),\ \overline{T} = \overline{R}A,\ \text{and}\ \overline{R} = R' + R - R'AR,$$

together with the accompanying norm $\|\cdot\|_{\overline{R}^{-1}}$, where R' is the adjoint operator of R and $\overline{R}$ is called the *symmetrized* operator of R. We assume, moreover, that $\overline{R}$ is SPD, which implies that the smoother R is always convergent and such that

$$\|\mathbf{v}\|_A^2 \le \|\mathbf{v}\|_{\overline{R}^{-1}}^2.$$

The restriction of (1) to the coarse space is then expressed as

$$A_c\mathbf{u}_c = \mathbf{f}_c$$

where

$$A_c = P'AP,\ \ \mathbf{f}_c = P'\mathbf{f},\ \ \text{with } P' \text{ adjoint operator of } P.$$

For the sake of the analysis, the coarse space solver B_c is often chosen to be the exact solver, namely $B_c = A_c^{-1}$, however, we should distinguish between an exact TL-AMG and an inexact TL-AMG when B_c is only an approximation of A_c^{-1}. Given $\mathbf{g} \in V'$, a TL-AMG operator B, defined by the above components is described in Algorithm 1. The corresponding error propagation operator $E = I - BA$ is $E = (I - RA)(I - \Pi_c)$, where $\Pi_c = PA_c^{-1}P'A$ is the orthogonal projection on V_c.

Algorithm 1: Two-level post-smoothed MG.

Data A: matrix, R: convergent smoother, P: prolongator, B_c: coarse solver, $\mathbf{g}$: arbitrary vector in V'
Result $B\mathbf{g}$: preconditioned vector
Coarse grid correction: $\mathbf{w} := PB_cP'\mathbf{g}$
Post-smoothing: $B\mathbf{g} := \mathbf{w} + R(\mathbf{g} - A\mathbf{w})$

We can now explore the connection between the TL-AMG convergence rate and the selection of the coarse spaces. Let us consider the

prolongation operator P, used in representing the operator Π_c; in our case, P will be a piecewise constant prolongation, a very common choice. This means that the coarse grid correction computed on the residual equation will be transferred back to the fine grid by assigning the same value to all fine grid variables associated with a given coarse variable.

A common alternative to this choice is to smooth out the prolongator P by means of a number of smoothing iterations applied to a piecewise constant *tentative prolongator*; this choice gives rise to the popular class of AMG algorithms with *smoothed* aggregation [21,15,2], but they are out of the scope of the present analysis.

We assume now that there exists a sequence of spaces $V_1, V_2, \dots, V_J$, which are not necessarily subspaces of the vector space V, and that each of them is related to the original space V by a linear operator

$$\Pi_j : V_j \to V. \tag{3}$$

We are moreover assuming that V can be written as a sum of subspaces as

$$V = \sum_{j=1}^{J} \Pi_j V_j.$$

Let $\underline{W} = V_1 \times V_2 \times \dots \times V_J$, with the inner product

$$(\underline{\mathbf{u}}, \underline{\mathbf{v}}) = \sum_{j=1}^{J} (\mathbf{u}_j, \mathbf{v}_j),$$

with $\underline{\mathbf{u}} = (\mathbf{u}_1, \dots, \mathbf{u}_J)^T$ and $\underline{\mathbf{v}} = (\mathbf{v}_1, \dots, \mathbf{v}_J)^T$. Let also $\Pi_W : \underline{W} \to V$ be the operator:

$$\Pi_W \underline{\mathbf{u}} = \sum_{j=1}^{J} \Pi_j \mathbf{u}_j, \quad \forall \underline{\mathbf{u}} \in \underline{W}. \tag{4}$$

We can then write

$$\Pi_W = (\Pi_1, \dots, \Pi_J) \text{ and } \Pi'_W = (\Pi'_1, \dots, \Pi'_J)^T.$$

We assume that for each j there is an operator $A_j : V_j \to V'_j$ which is symmetric positive semi-definite, and we define $\underline{A}_W : \underline{W} \to \underline{W}'$ as follows:

$$\underline{A}_W = \operatorname{diag}(A_1, A_2, \dots, A_J).$$

We also assume that for each j there is a SPD operator $D_j : V_j \to V'_j$, and define $\underline{D} : \underline{W} \to \underline{W}'$ as follows:

$$\underline{D} = \operatorname{diag}(D_1, D_2, \dots, D_J).$$

We associate a coarse space with each V_j: $V_j^c \subset V_j$, and consider the corresponding orthogonal projection $Q_j : V_j \to V_j^c$ with respect to $(\cdot,\cdot)_{D_j}$. We define $\underline{Q} : \underline{W} \to \underline{W}'$ by $\underline{Q} = \operatorname{diag}(Q_1, \dots, Q_J)$.

Let us assume the following hold:

- For all $\underline{\mathbf{w}} \in \underline{W}$:

  $$\|\Pi_W \underline{\mathbf{w}}\|_D^2 \le C_{p,2} \|\underline{\mathbf{w}}\|_{\underline{D}}^2 \tag{5}$$

  for some positive constant $C_{p,2}$ independent of $\mathbf{w}$;
- For each $\mathbf{w} \in V$, there exists $\underline{\mathbf{w}} \in \underline{W}$ such that $\mathbf{w} = \Pi_W \underline{\mathbf{w}}$ and

  $$\|\underline{\mathbf{w}}\|_{\underline{A}_W}^2 \le C_{p,1} \|\mathbf{w}\|_A^2, \tag{6}$$

  for some positive constant $C_{p,1}$ independent of $\mathbf{w}$;
- For all j

  $$N(A_j) \subset V_j^c, \tag{7}$$

  where $N(A_j)$ is the kernel of A_j.

The above assumptions imply that if $\mathbf{w} \in N(A)$, then $\underline{\mathbf{w}} \in N(A_1) \times \dots \times N(A_J)$. We define the global coarse space V_c by

$$V_c = \sum_{j=1}^{J} \Pi_j V_j^c. \tag{8}$$

Furthermore, for each coarse space V_j^c, we define:

$$\mu_j^{-1}(V_j^c) = \max_{\mathbf{v}_j \in V_j} \min_{\mathbf{v}_j^c \in V_j^c} \frac{\|\mathbf{v}_j - \mathbf{v}_j^c\|_{D_j}^2}{\|\mathbf{v}_j\|_{A_j}^2}. \tag{9}$$

In the context of linear algebraic problems arising from finite elements methods, these are usually named the local Poincaré constants (see, e.g., [22, Section 1.5]); finally we define

$$\mu_c = \min_{1 \le j \le J} \mu_j(V_j^c), \tag{10}$$

which is finite thanks to assumption (7).

By TL-AMG convergence theory, if D_j provides a convergent smoother, then $(1 - \mu_j^{-1}(V_j^c))$ is an upper bound of the convergence rate for TL-AMG for V_j with coarse space V_j^c and the following theorem holds:

Theorem 1. *If all the previous assumptions hold, then for each $\mathbf{v} \in V$ we have the error estimate:*

$$\min_{\mathbf{v}^c \in V_c} \|\mathbf{v} - \mathbf{v}_c\|_D^2 \le C_{p,1} C_{p,2} \mu_c^{-1} \|\mathbf{v}\|_A^2.$$

Then the TL-AMG with coarse space defined in (8) converges with a rate:

$$\|E\|_A \le 1 - \frac{\mu_c}{C_{p,1} C_{p,2} c^D} \tag{11}$$

with c^D depending on the convergent smoother, i.e.,

$$c_D \|\mathbf{v}\|_D^2 \le \|\mathbf{v}\|_{\overline{R}^{-1}}^2 \le c^D \|\mathbf{v}\|_D^2. \tag{12}$$

From the above result it is clear why the constant μ_c in (10) represents the convergence quality measure for the aggregates that we were looking for. We will use it in Section 3, to infer the convergence of the TL-AMG based on coarsening relying on weighted matching described in [5,6], as well as to evaluate the quality of the aggregates. Let us also underline that many of the convergence results for TL-AMG methods can be described by means of this set of tools; see, e.g., [2, sections 12.4 and 13.1] for the application to the classical AMG and aggregation-based AMG.

3. Generating aggregates from matching in weighted graphs

We now adopt the theory discussed in the previous section to analyze the construction of the coarse space by means of the *coarsening based on compatible weighted matching* as in [5,6]. We note that, as described in the original papers, our aggregation approach is driven by the idea to generate aggregates automatically, with no use of heuristics nor a priori information on the near kernel of the linear system; however, after generating non-overlapped aggregates by applying maximum weight matching, the setup of the prolongator operator is based on a projection of an arbitrary vector (hopefully a sample of slow-convergent error components, see Section 3.1 for discussion) on the aggregates, in a way similar to the well-known approaches of AMG based on smoothed aggregation [21].

We look at the graph $G = (\mathcal{V}, \mathcal{E})$ associated with the sparse matrix[1] A, also known as the adjacency graph of A. This is the graph G whose set of nodes $\mathcal{V}$ corresponds to the row/column indices $\mathcal{I} = \{1, \dots, n\}$ of A, and whose set of edges $e_{i \mapsto j} = (i, j) \in \mathcal{E}$ is induced by the sparsity

[1] For the sake of simplicity, we are using the same notation for representing linear operators and their corresponding matrices with the only change being the substitution of the adjoint operator with the transpose.

pattern of A. To this graph we associate an edge weight matrix $\hat{A}$ with the following entries:

$$(\hat{A})_{i,j} = \hat{a}_{i,j} = 1 - \frac{2a_{i,j}w_i w_j}{a_{i,i}w_i^2 + a_{j,j}w_j^2},$$

where $a_{i,j}$ are the entries of A and $\mathbf{w} = (w_i)_{i=1}^n$ is a given vector. For such a graph, a *matching* $\mathcal{M}$ is a set of pairwise non-adjacent edges, containing no loops, i.e., no two edges share a common vertex. We call $\mathcal{M}$ a *maximum product matching* if it maximizes the product of the weights of the edges $e_{i\mapsto j}$ belonging to it, i.e., if it maximizes the product of the entries of $\hat{A}$ associated to the matched indices. We stress that for sub-optimal matching algorithms, as discussed in Section 3.2, there may be nodes which are not endpoints of any of the matched edges: we call such nodes *unmatched*. By the above procedure we are choosing as $V_1, \dots, V_J$ the spaces defined by the aggregates $\{\mathcal{A}_j\}_{j=1}^J$ for the row/column indices $\mathcal{I}$ denoting the matrix entries; equivalently, we are decomposing the index set as

$$\mathcal{I} = \bigcup_{i=1}^{J} \mathcal{A}_j,\ \mathcal{A}_i \cap \mathcal{A}_j = \emptyset \text{ if } i \neq j. \tag{13}$$

More generally, to further reduce the dimension of the coarse space, we can perform subsequent pairwise matching steps, i.e., we can iterate ℓ times the matching procedure, acting each time on the graph G' obtained by collapsing together the matched nodes from the previous step.

Let us consider the case in which a single step of pairwise aggregation is performed. We can identify two types of aggregates $\mathcal{A}_j$: those corresponding to pairs of matched nodes, for which $V_j = \mathbb{R}^2$, and those corresponding to the unmatched nodes, for which $V_j = \mathbb{R}$.

The next step in the construction is the definition of the global prolongation matrix P by means of the operators $\Pi_j : V_j \to V$, for $j = 1, \dots, J$, in (3). Let us denote by $n_p = |\mathcal{M}|$ the cardinality of the graph matching $\mathcal{M}$, i.e., the number of matched nodes, and by n_s the number of unmatched nodes. We identify for each edge $e_{j_1\mapsto j_2} \in \mathcal{M}$ the vectors

$$\mathbf{w}_{e_{j_1\mapsto j_2}} = \frac{1}{\sqrt{w_{j_1}^2 + w_{j_2}^2}} \begin{bmatrix} w_{j_1} \\ w_{j_2} \end{bmatrix}, \quad \mathbf{w}^{\perp}_{e_{j_1\mapsto j_2}} = \frac{1}{\sqrt{\frac{w_{j_1}^2}{a_{j_2,j_2}^2} + \frac{w_{j_2}^2}{a_{j_1,j_1}^2}}} \begin{bmatrix} -\frac{w_{j_1}}{a_{j_2,j_2}} \\ \frac{w_{j_2}}{a_{j_1,j_1}} \end{bmatrix}. \tag{14}$$

To build the local prolongator Π_j we introduce the family of maps $\{\eta_j\}_{j=1}^J$ for

$$\begin{aligned} &\eta_j : \{j_1, j_{n_j}\} \to \{1, 2, \dots, n\} \\ &\eta_j(j_p) = i \iff \mathcal{A}_j = \{j_1, j_{n_j}\}, \quad \text{and} \quad i = j_p, \end{aligned} \tag{15}$$

where we assume that in the case of an unmatched node, i.e., when $n_j = 1$, then $\mathcal{A}_j = \{j_1\}$. Thus we have defined the correspondence relation between the indices in the local numbering on the aggregates and the numbering in the global space, that is

$$\{j_1, j_{n_j}\} = \left\{\eta_j(1), \eta_j(j_{n_j})\right\}. \tag{16}$$

Let now $\{\delta_i\}_{i=1}^n$ and $\{e_{j,j_p}\}_{j_p=1}^{n_j}$ be the basis of V and V_j respectively

$$V = \operatorname{span}\{\delta_i\}_{i=1}^n, \quad V_j = \operatorname{span}\{e_{j,j_p}\}_{p=1}^{n_j}.$$

We introduce the operator $\hat{\Pi}_j$ and its dual with respect to the Euclidean/ℓ^2 inner product $\hat{\Pi}'_j$, respectively, as

$$\forall s \in V_j,\ s = \sum_{p=1}^{n_j} s_p e_{j,j_p}, \quad \text{then} \quad \hat{\Pi}_j s = \sum_{p=1}^{n_j} s_p \delta_{\eta_j(j_p)},$$

and

$$\forall w \in V,\ \hat{\Pi}'_j w \in V_j, \qquad \hat{\Pi}'_j w = \sum_{p=1}^{n_j} w_{\eta_j(j_p)} e_{j,j_p},$$

that has been obtained by direct computation of its ℓ^2 inner product. Finally, we define the Π_j associated with the aggregates as

$$\Pi_j = \hat{\Pi}_j \hat{\Pi}'_j, \quad j = 1, \dots, J, \quad \Pi_j = \Pi'_j,\ \Pi_j \Pi_k = 0, \text{ whenever } j \neq k. \tag{17}$$

Then $\underline{A}_W = \operatorname{diag}(A_1, A_2, \dots, A_J) = \operatorname{diag}(\Pi'_1 A \Pi_1, \Pi'_2 A \Pi_2, \dots, \Pi'_J A \Pi_J)$ is the block-diagonal operator corresponding to the restriction of A to the unknowns belonging to the j-th aggregate, and the corresponding columns of the projection matrix are given by

$$\tilde{P} = [\mathbf{p}_1, \dots, \mathbf{p}_{n_p}] \text{ for } \mathbf{p}_j = \Pi_j \mathbf{w}_{e_{i\mapsto j}}.$$

Remark 1. The vectors in (14) are by construction D–orthogonal with respect to the local matrix

$$D_{e_{i\mapsto j}} = \operatorname{diag}([a_{i,i}, a_{j,j}]), \text{ i.e., } \mathbf{w}^T_{e_{i\mapsto j}} D_{e_{i\mapsto j}} \mathbf{w}^{\perp}_{e_{i\mapsto j}} = 0.$$

To complete the construction of the prolongation matrix, we also need to fix an ordering for the unmatched $n_s = n_c - n_p = J - n_p$ nodes, where $n_c = J$ finally denotes the dimension of the coarse space. The local projector Π_j is again the one in (17), but we apply it to the scalars $w_k/|w_k|$, $k = 1, \dots, n_s$, thus obtaining the remaining columns of the prolongation matrix

$$W = [\mathbf{p}_{n_p+1}, \dots, \mathbf{p}_{n_p+n_s}] = [\mathbf{p}_{n_p+1}, \dots, \mathbf{p}_J] \text{ for } \mathbf{p}_k = \Pi_k \frac{\mathbf{w}_k}{|\mathbf{w}_k|}.$$

In an expanded form, the resulting prolongation matrix can then be expressed as

$$\begin{aligned} P &= \underbrace{\begin{bmatrix} \underbrace{\begin{bmatrix} w_{e_1} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & w_{e_{n_p}} \end{bmatrix}}_{n_p} 2n_p & \mathbf{0} \\ \mathbf{0} & \underbrace{\begin{bmatrix} w_1/|w_1| & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & w_{n_s}/|w_{n_s}| \end{bmatrix}}_{n_s} n_s \end{bmatrix}}_{n_c = n_p + n_s = J} n = 2n_p + n_s \\ &= \begin{bmatrix} \tilde{P} & W \end{bmatrix} = [\mathbf{p}_1, \dots, \mathbf{p}_J], \end{aligned} \tag{18}$$

which also allows to express the global coarse space as the space generated by the columns of P, i.e., $V_c = \operatorname{span}\{\mathbf{p}_1, \dots, \mathbf{p}_J\}$. The matrix P we have just built represents a piecewise constant interpolation operator.

3.1. Selecting the weight vector

We can now use again the general theory for the convergence of a multigrid algorithm to discuss what is the optimal choice for the weight vector $\mathbf{w}$, and therefore identify the optimal prolongator operator P. To this aim we recall the following well known result [23,2,24].

Theorem 2. *Let $\{\lambda_j, \mathbf{\Phi}_j\}_{j=1}^n$ be the eigenpairs of $\overline{T} = \overline{R}A$ with $0 < \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$. Let us also assume that $\mathbf{\Phi}_j$ are orthogonal w.r.t. $(\cdot,\cdot)_{\overline{R}^{-1}}$. The convergence rate $\|E(P)\|_A$ is minimal for P such that*

$$\operatorname{range}(P) = \operatorname{range}(P^{opt}), \quad \textit{where} \quad P^{opt} = \left[\mathbf{\Phi}_1, \dots, \mathbf{\Phi}_{n_c}\right].$$

In this case,

$$\|E\|_A^2 = 1 - \lambda_{n_c+1}.$$

Therefore, a sensible choice would be to include in the range of P at least the first eigenvector $\mathbf{\Phi}_1$; this would be sufficient to enforce convergence, albeit possibly with a poor convergence ratio.

Proposition 3. *Using the same notation of Theorem 2, if the weight vector* $\mathbf{w}$ *used to define the prolongator matrix* P *in* (18) *is the* $\mathbf{\Phi}_1$ *eigenvector of* $\overline{T} = \overline{R}A$ *then the* A*–norm of the error propagation matrix* $\|E\|_A^2$ *is less or equal than*

$$\|E\|_A^2 \leq 1 - \lambda_2.$$

Proof. The range of the prolongation matrix P in (18) includes the original vector of the weights $\mathbf{w}$, i.e., there exists $\mathbf{h} \in \mathbb{R}^{n_c}$ such that $P\mathbf{h} = \mathbf{w}$. The conclusion follows immediately by a straightforward application of Theorem 2. □

Unfortunately, this is not an optimal choice from a computational point of view; if we did possess some *a priori* information on the eigenvector, then using this information could improve the quality of the aggregates, and thus the convergence of the method.

In the case where we do not possess information on the eigenvector(s), selecting the appropriate vector $\mathbf{w}$ may not be an easy task. To obtain a good candidate in a completely black–box manner we could exploit the smoother $\overline{R}$ to select as a weight vector an ϵ–smooth algebraic vector in the sense of the following [2]:

Definition 1. Let $R : V \to V$ be a smoothing operator such that its symmetrization $\overline{R}$ is positive definite. Given $\epsilon \in (0, 1)$, we say that the vector v is algebraically ϵ-smooth with respect to A if

$$\|\mathbf{v}\|_A^2 \leq \epsilon \|\mathbf{v}\|_{\overline{R}^{-1}}^2.$$

Such a vector can be obtained by performing a few iterations of the smoother on either a random choice or on the initial theoretical guess.

The last possible adaptive refinement that we are going to consider is the application of a bootstrap iteration exploiting the multigrid hierarchy itself as in [6]. A whole hierarchy B_0 associated with an initial guess $\mathbf{w}_0$, again either a random or user-defined guess, is built in the first step of the bootstrap procedure. Then the hierarchy is used to refine the choice of vectors $\mathbf{w}$ by means of the iteration (2) for the homogeneous linear system, i.e.,

$$\text{Given } \mathbf{w}_0 \text{ compute} \begin{cases} \mathbf{w}^{(0)} = \mathbf{w}_{r-1}, & r = 1, \dots, k-1, \\ \mathbf{w}^{(j)} = \prod_{p=0}^{r-1} (I - B_p^{-1} A)\mathbf{w}^{(j-1)}, & j = 1, \dots, m, \\ \mathbf{w}_{r+1} = \mathbf{w}^{(m)}. \end{cases} \tag{19}$$

To build the multigrid hierarchies B_p for the bootstrap iteration (19) we exploit now the vectors $\mathbf{w}_r$ available at each rth step.

We stress that, from an operational point of view, this means that if one knows at least one ϵ-smooth vector $\mathbf{w}$ to be used as $\mathbf{w}_0$, then it is possible to use it to launch the bootstrap iteration (19) and obtain hierarchies $B_0, B_1, \dots, B_{r-1}$, each satisfying the convergence result in Theorem 1, and generating, when accumulated all–together, an algorithm with improved convergence rate. Moreover, if the bootstrap iteration is launched with a random vector then the TL–AMG algorithm with the bootstrap procedure can still obtain an acceptable convergence rate (see [6,25]).

3.2. Selecting the matching algorithm

One of the main costs in the construction of the multigrid hierarchy is represented by the computation of the maximum product matching needed to identify the aggregates. It is useful to distinguish here between two different approaches. The first approach is to compute an *exact* matching, i.e., a matching that achieves exactly the optimum value for the product. The second approach computes a matching whose product value is not optimal, but is guaranteed to be greater or equal to 1/2 of the maximum; this is called a $\frac{1}{2}$*-approximate* maximum product matching. Relaxing the requirement to obtain the exact optimum allows the achievement of both a reduction of the construction time, as well as the possibility to perform the building phase in a parallel context with a limited amount of data exchange. For the details regarding these computational complexity aspects, we refer to the discussion in [6]; here we focus on the quality of the aggregates obtained by using the different matching algorithms.

For the class of exact algorithms, we employ the algorithm in [9] that is implemented in the `HSL_MC64` routine [26], while for the approximate class we refer to the $\frac{1}{2}$–approximation algorithm by Preis [8], a parallel distributed-memory version of which is employed in [27], the auction type algorithm from [28], and the suitor algorithm in [7], which is what we applied in a parallel Graphics Processing Unit (GPU) setting (see [29]).

3.3. Computing the μ_c constant

First we focus on the task of computing exactly the μ_c constant in Theorem 1. Thus we first need to prove that the Assumptions in (5), (6) and (7) hold for the construction discussed in Section 3.

Lemma 1. *Let the two–grid hierarchy be constructed with the prolongator* P *in* (18). *Then assumptions* (5), *and* (6) *hold true with* $C_{p,1} = 1$, *and* $C_{p,2} = 1$. *Moreover, if* A *is SPD, assumption* (7) *holds since* $N(A_j) = \{\mathbf{0}\}$ *for every* j.

Proof. To prove (5) we observe that by (4) and (17) we have that for all $\underline{\mathbf{v}} \in W$

$$\|\Pi_W \underline{\mathbf{v}}\|_D^2 = \sum_{j=1}^{J} \|\Pi_j \underline{\mathbf{v}}\|_{D_j}^2 = \sum_{j=1}^{J} \left\| \begin{bmatrix} \underline{v}_{j_1} \\ \underline{v}_{j_2} \end{bmatrix} \right\|_{D_j}^2 = \|\underline{\mathbf{v}}\|_{\underline{D}}^2, \Rightarrow C_{p,2} = 1.$$

To prove (5) we use the "local-to-global" maps in (16) to have the index correspondence between the aggregates and the global matrix. Then, noting that $\Pi_j^2 = \Pi_j$, $\Pi_k \Pi_j = 0$ for $k \neq j$, and that $\Pi_j = \Pi_j'$, for $j = 1, \dots, J$, by a direct computation, we find that

$$\begin{aligned} \|\underline{\mathbf{w}}\|_{\underline{A}_W}^2 &= \langle \underline{A}_W \underline{\mathbf{w}}, \underline{\mathbf{w}} \rangle_{\ell^2} = \left\langle \sum_{j=1}^{J} \Pi_j A \Pi_j \mathbf{w}, \sum_{k=1}^{J} \Pi_k \mathbf{w} \right\rangle_{\ell^2} \\ &= \sum_{k=1}^{J} \sum_{j=1}^{J} \langle \Pi_k \Pi_j A \Pi_j \mathbf{w}, \mathbf{w} \rangle_{\ell^2} = \sum_{j=1}^{J} \langle \Pi_j A \Pi_j \mathbf{w}, \mathbf{w} \rangle_{\ell^2} \\ &= \sum_{j=1}^{J} \langle A \Pi_j \mathbf{w}_j, \Pi_j \mathbf{w}_j \rangle_{\ell^2} = \|\mathbf{w}\|_A^2. \end{aligned}$$

The kernel of the projected matrices A_j is reduced to the zero vector since the projector has orthogonal columns, and thus the projected matrices on $\underline{W}$ are SPD. □

The above assumptions practically depend on the fact that independently from the number of aggregation sweeps we collect together, we are decomposing the index set $\mathcal{I}$ as a direct sum of non-overlapping indices as in (13).

This means that we can compute the global constant μ_c in (10) *a posteriori* by solving the generalized eigenvalue problem

$$D(I - Q)\mathbf{x} = \mu_c^{-1} A\mathbf{x}, \tag{20}$$

where Q has been built from the D_j–orthogonal projectors $Q_j : V_j \to V_j^c$, which in our case have the following representation matrices:

$$Q_j = \begin{cases} \mathbf{w}_j (\mathbf{w}_j^T D_j \mathbf{w}_j)^{-1} \mathbf{w}_j^T D_j, & j = 1, \dots, n_p \\ 1, & j = n_p + 1, \dots, n_p + n_s = J \end{cases}$$

and in aggregate form as the D–orthogonal projector represented by:

$$Q = P(P^T D P)^{-1} P^T D = \operatorname{diag}(Q_1, \dots, Q_J). \tag{21}$$

3.4. Estimating the μ_c constant

The general theory for an aggregation-based multigrid, as formalized in [2] and specialized in the previous Section 3.3 for our method, was originally applied in [1,12] for the case of disjoint aggregates with piecewise constant prolongators having unit coefficients; refer also to the *bibliographical notes* in [2, Section 8.5]. An additional tool provided by the discussion in [12] is the possibility of carrying out a purely *local* analysis by looking only at the restriction on the aggregates of the operators $\underline{A}_W$, and $\underline{D}$ under stricter hypothesis on the matrix A of the system and on possible aggregates.

Specifically, to adopt the general strategy introduced in [12], we identify these operators as the restriction of the operator A to the aggregates obtained through the matching algorithm, i.e.,

$$\underline{A}_W = (A_1, \dots, A_J), \quad A_k = A|_{V_k}, \quad \underline{D} = (\underline{D}_1, \dots, \underline{D}_J), \quad D_k = D|_{V_k}. \tag{22}$$

We can then write the complete matrix A as the sum of the block diagonal matrix $\underline{A}_W$ and a remainder A_R containing all the parts we have discarded. Under the stricter hypothesis on A discussed in [12] it is possible to find symmetric and non-negative definite $\underline{A}_W$ and A_R. This allows us to apply [12, Theorem 3.4] and obtain the 'local' bound to the global μ_c constant in Theorem 1. We simply restate the result here in the notation from [2] and the construction from Section 3.

Theorem 4 *(Restatement of [12, Theorem 3.4]). Let $\underline{A}_W = (A_1, \dots, A_J)$ and $\underline{D} = (\underline{D}_1, \dots, \underline{D}_J)$ satisfy the splitting condition $A = \underline{A}_W + A_R$, with $\underline{A}_W$ and A_R both symmetric and non-negative definite, that is, every $\{A_j\}_{j=1}^J$ is non-zero symmetric non-negative definite and $\underline{D}$ is symmetric positive definite. Let $\mathbf{p}$ be one of the columns of P in (18), i.e., $\mathbf{p} = \mathbf{w}_{e_{i\to j}}$ for the indices (i, j) relative to the given aggregate.*

Then μ_c is defined as in (10), and the $\mu_j^{-1}(V_j^c)$ are such that

$$\lambda_2^{-1}(D_j^{-1}A_j) \le \mu_j^{-1}(V_j^c) \le \lambda_1^{-1}(D_j^{-1}A_j).$$

Moreover, if either $(\mathbf{w}_{e_{i\to j}}, \lambda_1(D_j^{-1}A_j))$, or $(\mathbf{w}^{\perp}_{e_{i\to j}}, \lambda_2(D_j^{-1}A_j))$ are eigencouples of the matrix $D_j^{-1}A_j$, then

$$\mu_j^{-1}(V_j^c) = \lambda_2^{-1}(D_j^{-1}A_j).$$

We stress that while in general it is always possible to compute the quantity μ_c in (10) by solving the eigenvalue problem in (20), and thus estimate the overall quality of the matching procedure, application of Theorem 4 to obtain the bound by using only local information requires the stricter hypotheses on the splitting of A.

4. Numerical experiments

To highlight the results of Theorem 4 we consider the case study of the 2D Laplace equation with variable coefficients on the unit square $\Omega = [0, 1]^2$, discretized with 5–point finite differences, i.e. the equation

$$\begin{cases} -\nabla \cdot (a(x, y)\nabla u(x, y)) = f(x, y), & (x, y) \in \Omega, \\ u(x, y) = 0, & (x, y) \in \partial\Omega, \end{cases} \tag{23}$$

and discretized by Lagrangian P1 elements on an unstructured triangular grid. We focus on a 2D example so that we can graphically represent the different aggregates. We concentrate first on the computation of the bounds discussed in Theorem 4 and on the analysis of the different bounds obtained for the different choices of the matching algorithm in Section 3.2 while keeping fixed the choice of the weight vector $\mathbf{w}$. Then, in the second part of the numerical examples, we devote our attention to the analysis of the quality of the aggregates for different choices of the weight vectors $\mathbf{w}$, while considering also the different refinement strategies discussed in Section 3.1.

The version of the BootCMatch algorithm [6] we use here for the tests is available on the repository https://github.com/bootcmatch/BootCMatch. All the plots and the eigenvalues/eigenvectors computations are then performed in Matlab v. 9.6.0.1072779 (R2019a) on the matrices exported in Matrix Market format.

4.1. Computing the μ_c constants

To confirm the applicability of the theory developed in Section 3 we compute both the "true" μ_c constants by solving the generalized eigenvalue problem with the D–orthogonal projector Q in (21), and the estimate obtained by means of Theorem 4, when the splitting for the matrices $\underline{A}_W$ is available, for three different prototypical problems obtained from different choices of the diffusion coefficient in (23). For each of these cases we consider the various matching algorithms discussed in Section 3.2 and the application of $\ell = 1, 2$ steps of pairwise matching, i.e., we consider aggregates made by at most two or four fine variables. In all cases, we consider the weight vector $\mathbf{w} = (1, 1, \dots, 1)^T$, which is suggested by the structure of the matrix. We stress that all the results obtained in the following subsections can be read alongside the numerical experiments in [6] since they complement and further explains the convergence behavior of the method discussed there. To present a wider array of tests, we have given other examples in the Supplementary materials A.1.

4.1.1. The constant coefficient diffusion

The first case is the Laplacian with homogeneous coefficients, i.e., $a(x, y) = 1$, on a uniform $n \times n$ grid. This gives rise to the matrix

$$A_{n^2} = I_n \otimes T_n + T_n \otimes I_n, \quad T_n = \mathrm{tridiag}(-1, 2, -1),$$

scaled in such a way that its coefficients are independent from the dimension n^2 of the problem. We first visualize the different aggregates generated by the various matching algorithms in Fig. 1. In this case the aggregation based on the maximum product matching `HSL_MC64` produces the same aggregates that can be obtained by using the standard C\F–splitting. Moreover, by (18) it is straightforward to observe that P is a scalar multiple of the one obtained by choosing $\mathbf{w}_{e_i}$ equal to the vector of all ones; hence, the methods produce exactly the same Q of the classical aggregation, and therefore the same bounds obtained for it in [12, Theorem 3.4]. The aggregates also match the quality of the aggregates in [30], in which the matching strategy for the identification of the aggregates is applied directly to A and coupled with the prolongator P whose nonzero entries are all 1; see the results in Table 1.

Concerning the usage of alternative matching methods, we see that the `HSL_MC64` and the `SUITOR` algorithms do produce the same μ_c constants and bounds, even if `SUITOR` is only guaranteed to reach a value of the objective function one half away from the optimal one. In general, we can observe that in the cases $\ell = 1$ the same constants are reached for different aggregates. This suggests that reaching the maximum weight is not mandatory and that different configurations can yield the same results in terms of the overall quality of the aggregates. To achieve the upper bound from Theorem 4, we use the auxiliary splitting obtained by decreasing the diagonal blocks on the various aggregates by a correction of the form $\pm\delta_j I$ where each δ_j is computed heuristically to enforce the hypotheses. In these cases, for all the matching algorithms when we employ a single sweep, we use $\delta_j = 1/3 \min(A_j\mathbf{1})$, that is $1/3$ of the minimum row sum of the projection of A on the aggregate. When two sweeps are employed, we use instead $\delta_j = \min(A_j\mathbf{1})$ for all the matching but the Auction case in which we employ $\delta_j = 1/2 \min(A_j\mathbf{1})$. We stress that it is difficult to prescribe a formula to achieve the splitting and the local bound without looking into the matrices obtained from the matching procedure, since in general, this may not exist; see, e.g., the next example in which we encounter such a case for one of the matching algorithms.

4.1.2. Diffusion with axial anisotropies

As the second test case we consider having a simple spatial anisotropy oriented with the y–grid lines, i.e.,

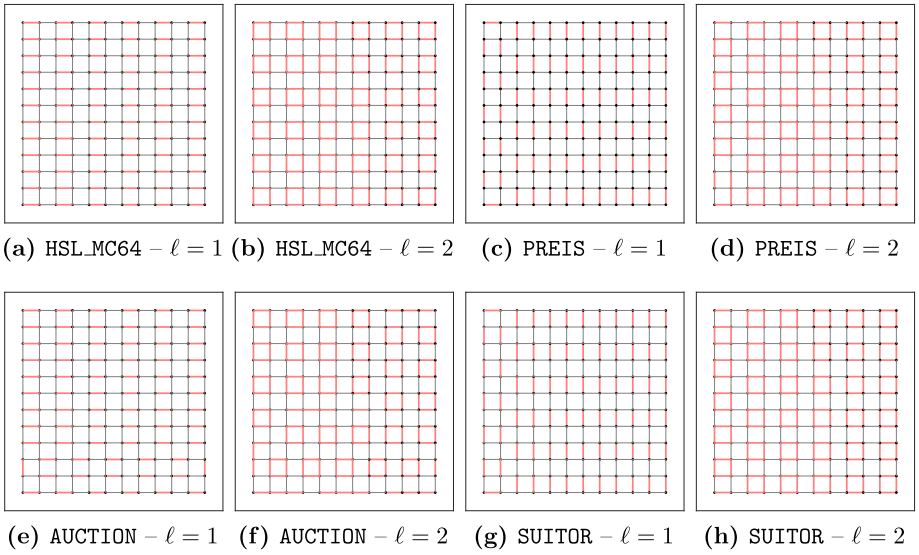

(a) HSL_MC64 – $\ell = 1$ (b) HSL_MC64 – $\ell = 2$ (c) PREIS – $\ell = 1$ (d) PREIS – $\ell = 2$

(e) AUCTION – $\ell = 1$ (f) AUCTION – $\ell = 2$ (g) SUITOR – $\ell = 1$ (h) SUITOR – $\ell = 2$

Fig. 1. Constant coefficient diffusion problem. Aggregates obtained with the weight vector $\mathbf{w} = (1, 1, \ldots, 1)^T$, and the different matching algorithms for $\ell = 1, 2$ pairwise matching steps.

Table 1
Constant coefficient diffusion problem. Comparison of the bound in Theorem 4 with true value of μ_c in (10). Aggregates obtained with the weight vector $\mathbf{w} = (1, 1, \ldots, 1)^T$, and the different matching algorithms for $\ell = 1, 2$ pairwise matching steps.

n	$\ell = 1$		$\ell = 2$	
	bound	μ_c^{-1}	bound	μ_c^{-1}
12	2.000	1.940	2.000	1.959
24	2.000	1.984	2.000	1.989
48	2.000	1.996	2.000	1.997
96	2.000	1.999	2.000	1.999

(a) HSL_MC64 – exact matching

n	$\ell = 1$		$\ell = 2$	
	bound	μ_c^{-1}	bound	μ_c^{-1}
12	2.000	1.923	2.062	2.046
24	2.000	1.982	2.062	2.052
48	2.000	1.996	2.062	2.052
96	2.000	1.999	2.062	2.052

(b) PREIS – $\frac{1}{2}$–approximate matching

n	$\ell = 1$		$\ell = 2$	
	bound	μ_c^{-1}	bound	μ_c^{-1}
12	2.000	1.908	2.667	2.544
24	2.000	1.980	2.894	2.964
48	2.000	1.995	2.667	2.166
96	2.000	1.999	2.667	2.173

(c) AUCTION – $\frac{1}{2}$–approximate matching

n	$\ell = 1$		$\ell = 2$	
	bound	μ_c^{-1}	bound	μ_c^{-1}
12	2.000	1.923	2.000	1.954
24	2.000	1.982	2.000	1.988
48	2.000	1.996	2.000	1.997
96	2.000	1.999	2.000	1.999

(d) SUITOR – $\frac{1}{2}$–approximate matching

$$A_{n^2} = \varepsilon(I_n \otimes T_n) + T_n \otimes I_n, \quad T_n = \text{tridiag}(-1, 2, -1), \quad \varepsilon = 100,$$

in which we are again using a scaling that makes the matrix coefficients independent of the problem size. Intuitively, in this case, we would expect the aggregates to be oriented with the anisotropy, i.e., along the y–axis. If we look at the aggregates we obtain in Fig. 2 we observe that the matching algorithms produce aggregates corresponding to our intuition, with the exception of the PREIS algorithm that for $\ell = 2$ produces some aggregates that do not seem feasible.

Indeed, if we look also at the constants μ_c, and their estimates reported in Table 2 we observe that, excluding the case of the PREIS algorithm, the μ_c constant behaves consistently. The failure in obtaining a bound in the case of the PREIS algorithm is due to the inability of finding a suitable splitting for the aggregates generated by this matching. Indeed, the existence of such splitting is a stricter hypothesis, and cannot be guaranteed in general. We refer back to the discussion in [12] where the original strategy for obtaining the local bound was devised. It is interesting to compare the value of the constant for $\ell = 1$ step of matching for this case with the one obtained for the case with constant coefficients in Table 1: observe in particular that the strong directionality of the diffusion makes the pairwise aggregates much more effective. On the other hand, we observe also that switching to larger aggregates leads to a worse quality of the aggregates than in the case of an isotropic problem.

4.1.3. *Diffusion on an unstructured mesh*

As a final test case, we consider again the Poisson problem with a constant diffusion coefficient but on an unstructured triangular mesh obtained via a Delaunay-based algorithm for which we report the subsequent refinements in Fig. 3.

The aggregates obtained for this test problem are depicted in Fig. 4, whereas the constants and bounds for $\ell = 1$ step of matching are shown in Table 3. Again for this case we could not find an appropriate splitting to produce the local bound of Theorem 4 when $\ell = 2$ steps of pairwise matching were used. If we compare the results in Table 3 with the ones in Table 1, then we observe that the quality of the aggregates, in this case, is analogous to the structured homogeneous case. We also observe that, again, the AUCTION algorithm manages to obtain aggregates with better quality than the ones obtained by all other algorithms, including the ones obtained by the exact matching algorithm. This is in agreement with the computational results discussed in [6].

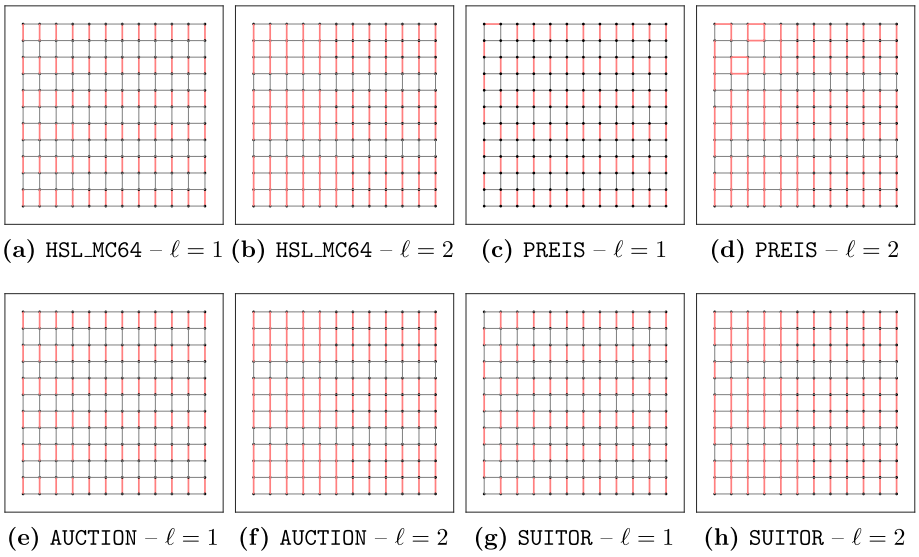

(a) HSL_MC64 – $\ell = 1$ (b) HSL_MC64 – $\ell = 2$ (c) PREIS – $\ell = 1$ (d) PREIS – $\ell = 2$

(e) AUCTION – $\ell = 1$ (f) AUCTION – $\ell = 2$ (g) SUITOR – $\ell = 1$ (h) SUITOR – $\ell = 2$

Fig. 2. Diffusion problem with y–axis oriented anisotropy $\epsilon = 100$. Aggregates obtained with the weight vector $\mathbf{w} = (1, 1, \ldots, 1)^T$, and the different matching algorithms for $\ell = 1, 2$ pairwise matching steps.

Table 2
Diffusion problem with y–axis oriented anisotropy $\epsilon = 100$. Comparison of the bound in Theorem 4 with true value of μ_c in (10) for $\ell = 1, 2$ pairwise aggregation steps, while using the various matching algorithm with weight vector $\mathbf{w} = (1, 1, \ldots, 1)^T$. The † represents a case in which we could not find the splitting needed to apply Theorem 4.

n	$\ell = 1$		$\ell = 2$	
	bound	μ_c^{-1}	bound	μ_c^{-1}
12	1.980	1.010	5.025	3.443
24	1.980	1.010	5.025	3.447
48	1.980	1.010	5.025	3.448
96	1.980	1.010	5.025	3.448

(a) HSL_MC64 – exact matching

n	$\ell = 1$		$\ell = 2$	
	bound	μ_c^{-1}	bound	μ_c^{-1}
12	1.765	1.741	†	8.580
24	1.765	1.745	†	8.725
48	1.765	1.745	†	8.730
96	1.765	1.745	†	8.730

(b) PREIS – $\frac{1}{2}$–approximate matching

n	$\ell = 1$		$\ell = 2$	
	bound	μ_c^{-1}	bound	μ_c^{-1}
12	1.980	1.010	5.025	3.443
24	1.980	1.010	5.025	3.447
48	1.980	1.010	5.025	3.448
96	1.980	1.010	5.025	3.448

(c) AUCTION – $\frac{1}{2}$–approximate matching

n	$\ell = 1$		$\ell = 2$	
	bound	μ_c^{-1}	bound	μ_c^{-1}
12	1.111	1.010	3.448	3.442
24	1.111	1.010	3.448	3.447
48	1.111	1.010	3.448	3.448
96	1.111	1.010	3.448	3.448

(d) SUITOR – $\frac{1}{2}$–approximate matching

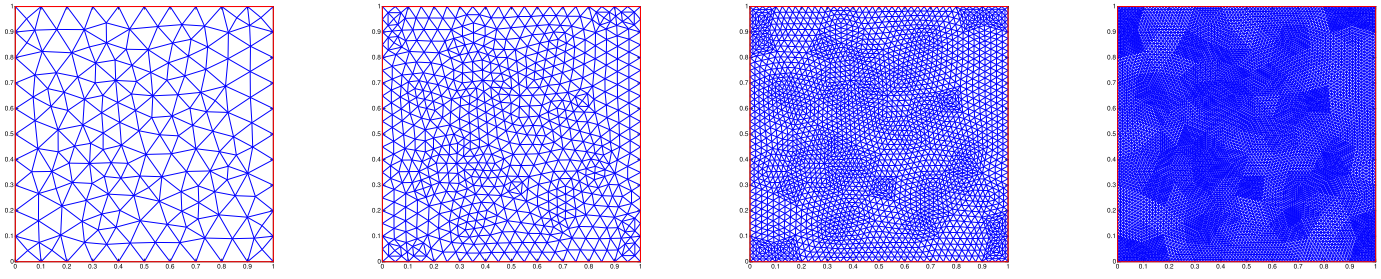

Fig. 3. Unstructured meshes for the Poisson problem, four levels of refinement using a Delaunay-based algorithm.

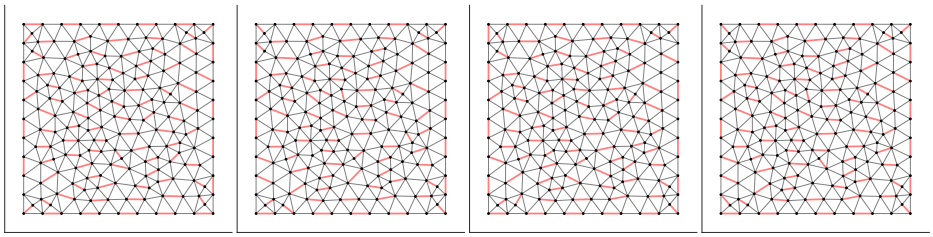

(a) HSL_MC64 – $\ell = 1$ (b) PREIS – $\ell = 1$ (c) AUCTION – $\ell = 1$ (d) SUITOR – $\ell = 1$

Fig. 4. Diffusion problem with constant coefficients on an unstructured grid. Aggregates obtained with the weight vector $\mathbf{w} = (1, 1, \ldots, 1)^T$, and the different matching algorithms for $\ell = 1$ pairwise matching steps.

Table 3
Diffusion problem with constant coefficients on an unstructured grid. Comparison of the bound in Theorem 4 with true value of μ_c in (10). Aggregates obtained with the weight vector $\mathbf{w} = (1, 1, \ldots, 1)^T$, and the different matching algorithms for $\ell = 1$ pairwise matching steps.

dofs	bound	μ_c^{-1}
185	3.000	1.613
697	3.000	1.562
2705	3.000	1.639
10657	3.000	1.897

(a) HSL_MC64 – exact matching

dofs	bound	μ_c^{-1}
185	2.396	1.830
697	2.306	1.667
2705	2.258	2.157
10657	2.249	2.001

(b) PREIS – $\frac{1}{2}$–approximate matching

dofs	bound	μ_c^{-1}
185	3.000	1.583
697	3.000	1.596
2705	2.103	1.794
10657	2.106	1.759

(c) AUCTION – $\frac{1}{2}$– approximate matching

dofs	bound	μ_c^{-1}
185	2.695	1.686
697	2.484	1.645
2705	2.258	1.690
10657	2.249	1.893

(d) SUITOR – $\frac{1}{2}$–approximate matching

4.2. Selecting the weight vector

We consider here the same test problems of the previous section, in which all the aggregates were computed by using the weight vector $\mathbf{w} = (1, 1, \ldots, 1)^T$, and compare them with the possible different choices for the weight vector discussed in Section 3.1. In every case we compare the aggregates obtained by using as weight vector $\mathbf{w}$ either:

1. a random initial guess, refined by some smoother iterations,
2. the vector $\mathbf{w} = (1, 1, \ldots, 1)^T$, refined by some smoother iterations,
3. the eigenvector associated with the smallest eigenvalue.

Information on using the bootstrap procedure is contained in the Appendix Section Supplementary materials A.2.1.

4.2.1. Random weight

We start considering the choice of an initial random weight vector $\mathbf{w}$ for all the test problems in Section 3.1, and consider using as smoother for its refinement the ℓ_1–Jacobi method [31]; each refinement step, in this case, has a cost that is dominated by a diagonal scaling. We test the procedure for all the matching algorithms discussed in Section 3.2, but we visualize the attained aggregates only for SUITOR. From what we have seen in the previous section, the SUITOR matching algorithm consistently gives good results for all the problems, and is, from a computational point of view, the best candidate when looking for the parallel applicability of the AMG algorithms [6]. In Fig. 5 we report the results obtained; as we can observe, a random initial guess without any refinement is a very poor choice, and we need several refinement steps to obtain constants μ_c that are comparable with the ones we have seen in Section 4.1. However, we can still go below the results obtained with the theoretical guess given by the constant weight vector $\mathbf{w} = (1, 1, \ldots, 1)^T$, at the cost of performing many refinement iterations. Note also that the aggregates for which these results are obtained would have been difficult to guess.

We consider for this case also a Poisson problem with an axially rotated anisotropy of angle θ and modulus ε on the same unstructured grid from Fig. 3, that is, we consider the discretization of

$$\begin{cases} -\nabla \cdot (\mathbf{A}\nabla u) = f, & (x, y) \in \Omega, \\ u = 0, & (x, y) \in \partial\Omega, \end{cases} \qquad \mathbf{A} \in \mathbb{R}^{2\times 2}. \tag{24}$$

Results for this test case are given in Fig. 6.

If we compare these results with the one in Fig. 5c, we observe that there is a moderate increase in the convergence constant for all combinations of rotation angle and modulus. Moreover, we can observe that over-refinement of the weight vector does not improve the overall quality of the aggregation procedure.

4.2.2. Refined uniform weight

As we have seen from the previous set of examples, a sufficient number of refinement steps on a random weight vector $\mathbf{w}$ already improves the quality of the aggregates obtained through the matching algorithms. Therefore, we expect to obtain a similar result when we start from a more reasonable guess for the weight vector. We consider the same experimental setting and only change the initial guess from a random $\mathbf{w}$ to the uniform vector $\mathbf{w} = (1, 1, \ldots, 1)^T$. For this case, we plot in Fig. 7 the aggregates obtained with the AUCTION algorithm, which attains the best constants. What is interesting to notice in this case is that very few iterations of the smoother coupled with the AUCTION algorithm generate aggregates that are better than the ones obtained by the complete

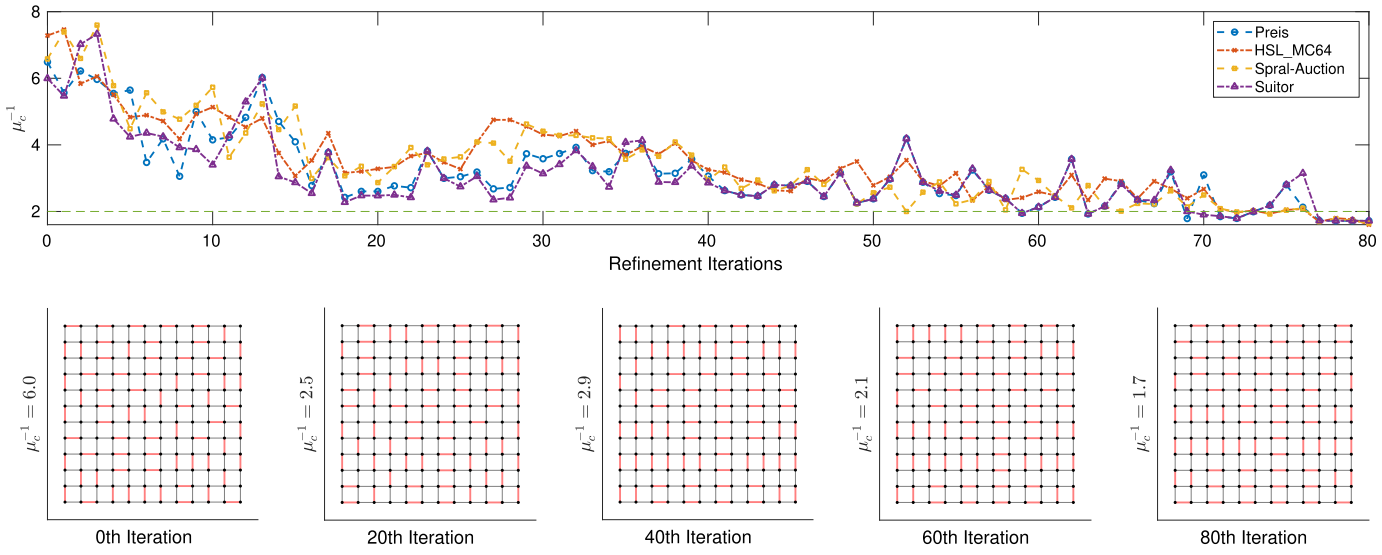

(a) Constant coefficient diffusion problem

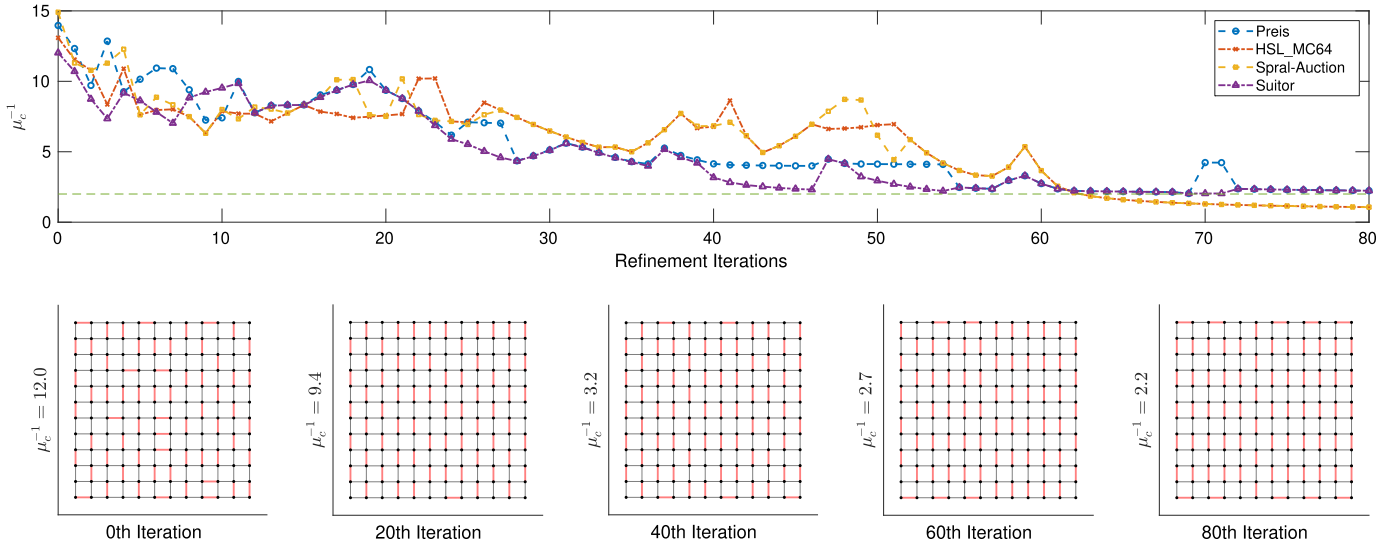

(b) Diffusion problem with y–axis oriented anisotropy $\varepsilon = 100$

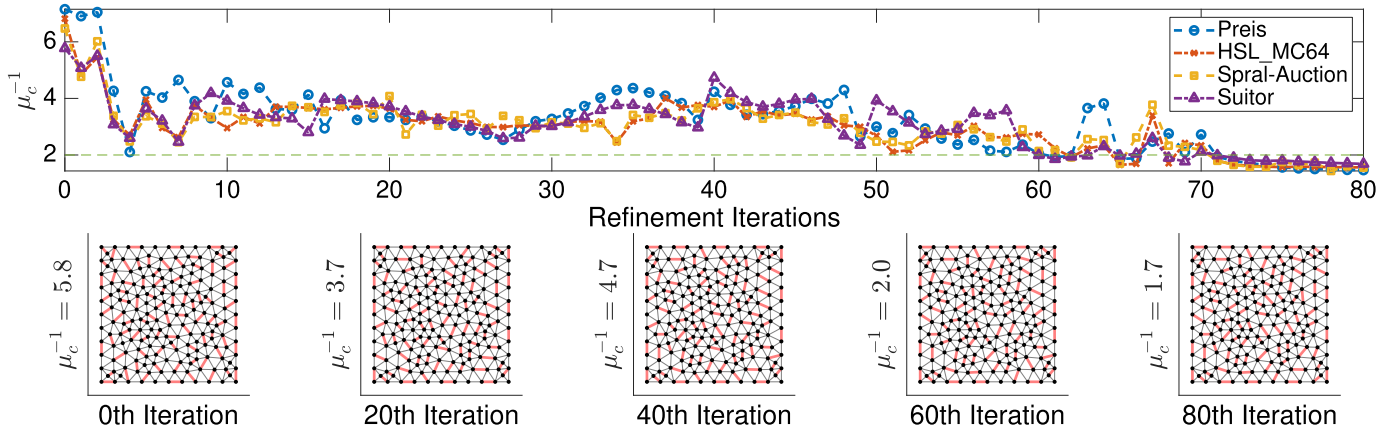

(c) Constant coefficient diffusion problem on an unstructured grid

Fig. 5. Refinement of the weight vector starting from a random guess, and using the ℓ_1–Jacobi smoother. We report a graph containing the μ_c^{-1} constant up to 80 refinement steps for a single sweep of pairwise aggregation. The depicted aggregates are the ones obtained with the SUITOR algorithm.

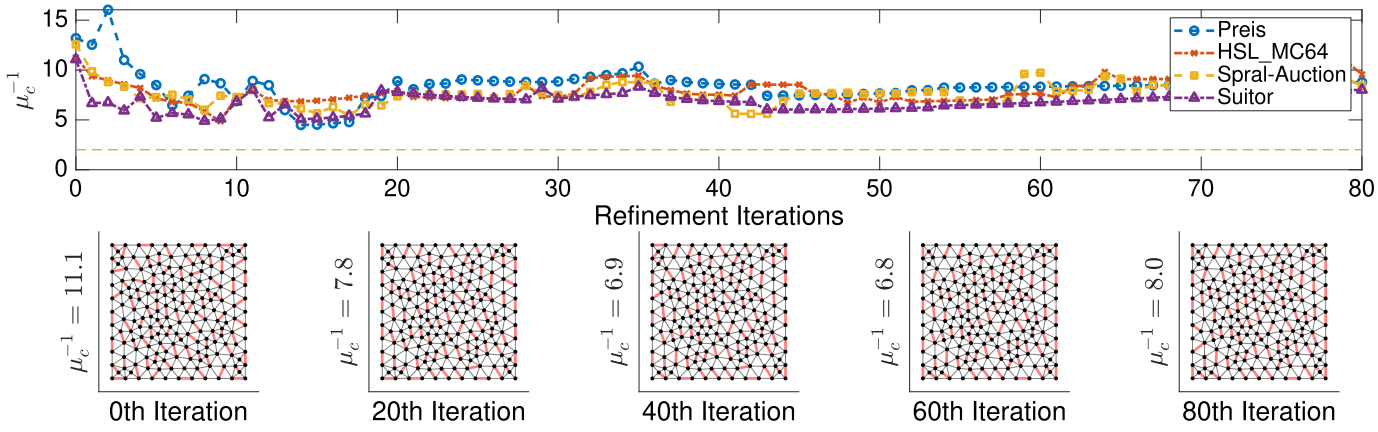

(a) $\theta = \pi/6$, $\varepsilon = 1e-2$

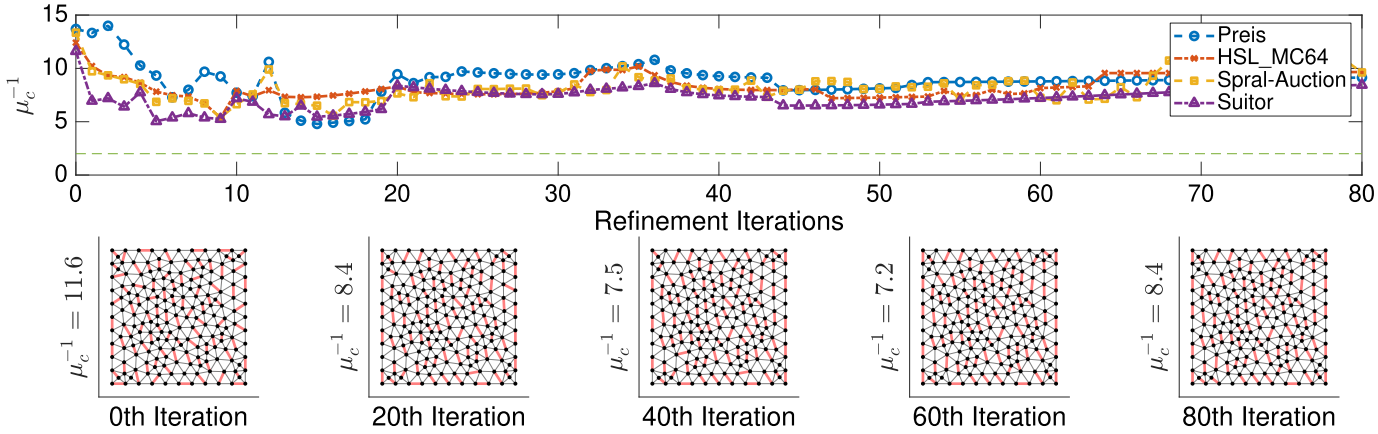

(b) $\theta = \pi/6$, $\varepsilon = 1e-3$

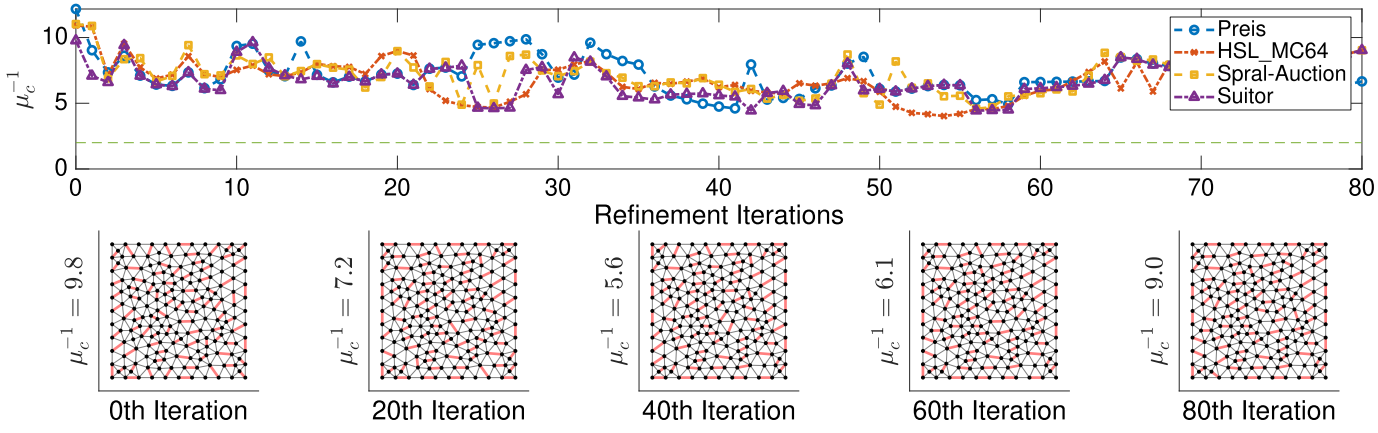

(c) $\theta = \pi/3$, $\varepsilon = 1e-2$

Fig. 6. Poisson problem on an unstructured grid with rotated anisotropy of angle θ, and modulus ε. Refinement of the weight vector starting from a random guess, and using the ℓ_1–Jacobi smoother. We report a graph containing the μ_c^{-1} constant up to 80 refinement steps for a single sweep of pairwise aggregation. The depicted aggregates are the ones obtained with the SUITOR algorithm.

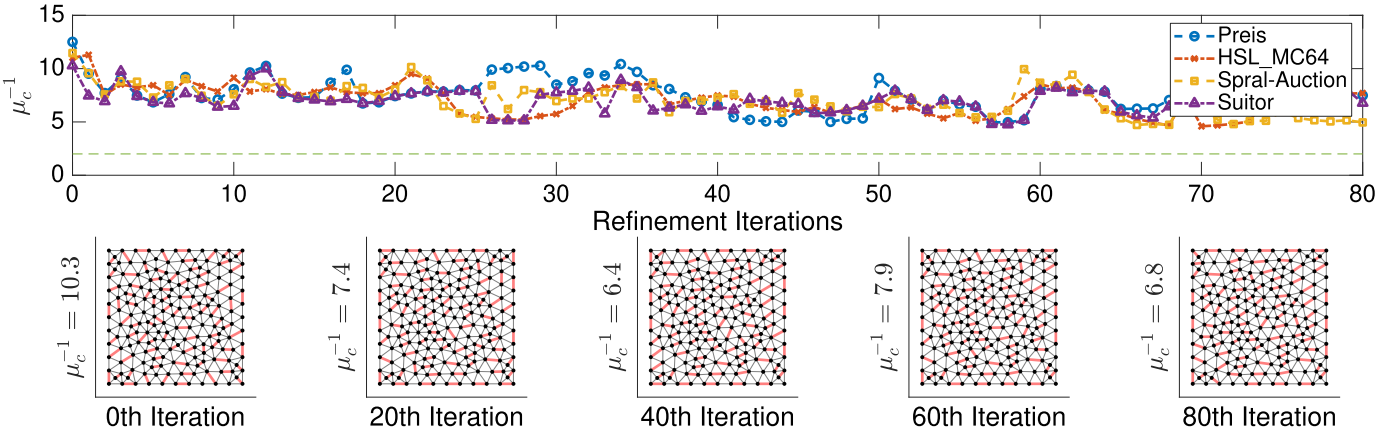

(d) $\theta = \pi/3$, $\varepsilon = 1e-3$

Fig. 6. (*continued*)

Table 4
Constants μ_c^{-1} obtained by using as weight vector $\mathbf{w}$ the eigenvector relative to the smallest eigenvalue as suggested by Proposition 3.

n	Homogeneous		*y*–axis		Homogeneous unstructured		
	$\ell = 1$	$\ell = 2$	$\ell = 1$	$\ell = 2$	dofs	$\ell = 1$	$\ell = 2$
12	1.476	2.336	0.973	2.699	185	1.5076	2.1977
24	1.737	3.826	1.001	3.249	697	1.5184	2.7255
48	1.809	4.274	1.008	3.401	2705	1.6349	3.1663
96	1.808	4.854	1.009	3.437	10657	1.7281	4.0177

matching algorithm `HSL_MC64`. The cases in which directionality in the coefficient is present end up in reproducing the expected aggregates with very few iterations.

As for the previous case, we consider again the Poisson problem on an unstructured mesh with rotated anisotropy from (24). Again, if we compare the results for this case in Fig. 8 with the ones in Fig. 7c we observe that there is a decrease in the performance of the aggregation procedure. Nevertheless, a small number of refinement iterations brings the quality of the aggregates near to the one of the homogeneous case.

4.2.3. *The eigenvector weight*

To complete our analysis we consider the aggregates generated by using as weight vector $\mathbf{w}$ the eigenvector associated with the smallest eigenvalue as in Proposition 3. Since this is a theoretical test, we consider only the application of the full matching algorithm `HSL_MC64`. We report the constants μ_c obtained by this choice in Table 4. If we compare them with the results in Tables 1, 2, and A.6 we observe two different behaviors. In the case of the simpler homogeneous problem selecting the eigenvector makes for worse μ_c constants when $\ell = 2$ steps of pairwise aggregations are used with respect to the case in which the vector $\mathbf{w} = (1, 1, \ldots, 1)^T$ is used in Table 1. If we look at the aggregates obtained by this choice in Fig. 9a and compare them with the one in Fig. 1, we see that the new aggregates are very far from the *box* aggregates obtained in that case, this causes that for certain aggregates we get an M–matrix A_k

$$A_k = \begin{bmatrix} 4 & & -1 & \\ & 4 & -1 & -1 \\ -1 & -1 & 4 & \\ & -1 & & 4 \end{bmatrix}, \qquad D_k = \begin{bmatrix} 4 & & & \\ & 4 & & \\ & & 4 & \\ & & & 4 \end{bmatrix},$$

whose scaled version $D_k^{-1} A_k$ is not a matrix with constant row sum. Therefore the associated $\mathbf{w}_{e_k}$ is not an eigenvector, i.e., we get a μ_k constant that is intermediate between λ_1 and λ_2, as discussed in Theorem 4. On the other hand, the constant vector choice always provides an irreducible and diagonally dominant M–matrix $D_k^{-1} A_k$, hence the vector $\mathbf{w}_{e_k} = (1, 1, 1, 1)^T$ is the unique eigenvector associated with the smallest eigenvalue, thus we obtain a better constant. Focusing now on the other cases in Table 4, whose aggregates are also depicted in Fig. 9, we obtain nearly the same results with the exception of the piecewise regular coefficients in which we are able to improve the attained constants – observe also that they are near the one obtained with the `SUITOR` algorithm and the $\mathbf{w} = (1, \ldots, 1)^T$ vector, even if the aggregates are very different.

What we can conclude from testing the usage of the eigenvector associated with the smallest eigenvalue is that, although guaranteeing the convergence due to Proposition 3, it can generate sub-optimal aggregates. On the other hand, either selecting a vector knowing the structure of the matrices $\{A_k\}_k$, as in the constant coefficient case with the $\mathbf{w} = (1, 1, \ldots, 1)^T$ vector or refining a choice by means of the smoothing procedure, can yield better results as we have seen.

5. Quality of the aggregates and the compatible relaxation principle

As already mentioned in Section 1, the need to measure the quality of a coarse space and to set up a general procedure for coarsening of the widest range of linear systems led to the nice principle of *compatible relaxation*. After its introduction in [16], it has been widely analyzed and related to the general theories for AMG convergence in many papers, starting from [18]. This principle has been applied as a guideline to define the coarsening method described in this paper, as emphasized in the original papers [5,6]. In the following, we show that the results obtained by the quality measure discussed in this paper are in good agreement with a quality measure based on the convergence rate of a compatible relaxation, showing the coherence of the convergence theories. Main advantage in using the constant μ_c in (11) is that it does not depend on a selected smoother and often gives more accurate information on the quality of the coarse space, as also shown in some of our experiments. Furthermore, we observe that the setup of a compatible relaxation scheme requires to build in an explicit way the complementary space to the coarse space, as explained in the following.

To introduce the measure based on compatible relaxation, we need to define the following 2×2–block factorization

$$\begin{bmatrix} P_f^T \\ P^T \end{bmatrix} A \begin{bmatrix} P_f & P \end{bmatrix} = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix}, \qquad \text{for } P^T D P_f = 0, \tag{25}$$

where $\mathrm{range}(P_f)$ is the space in which the smoother should be effective; this can be used to obtain a decomposition of the whole $\mathbb{R}^n$ since for all $\mathbf{e} \in \mathbb{R}^n$ we have $\mathbf{e} = P_f \mathbf{e}_f + P \mathbf{e}_c$. Exploiting the observation in Remark 1, we can express the matrix P_f through the block factorization (25) in a straightforward way as

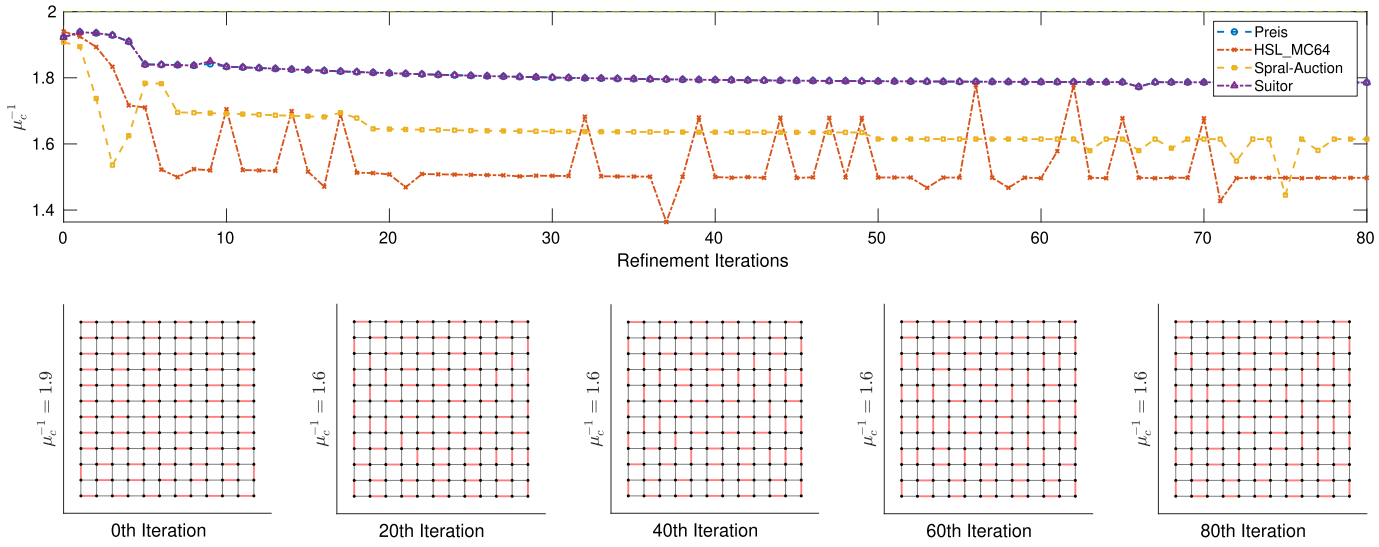

(a) Constant coefficient diffusion problem

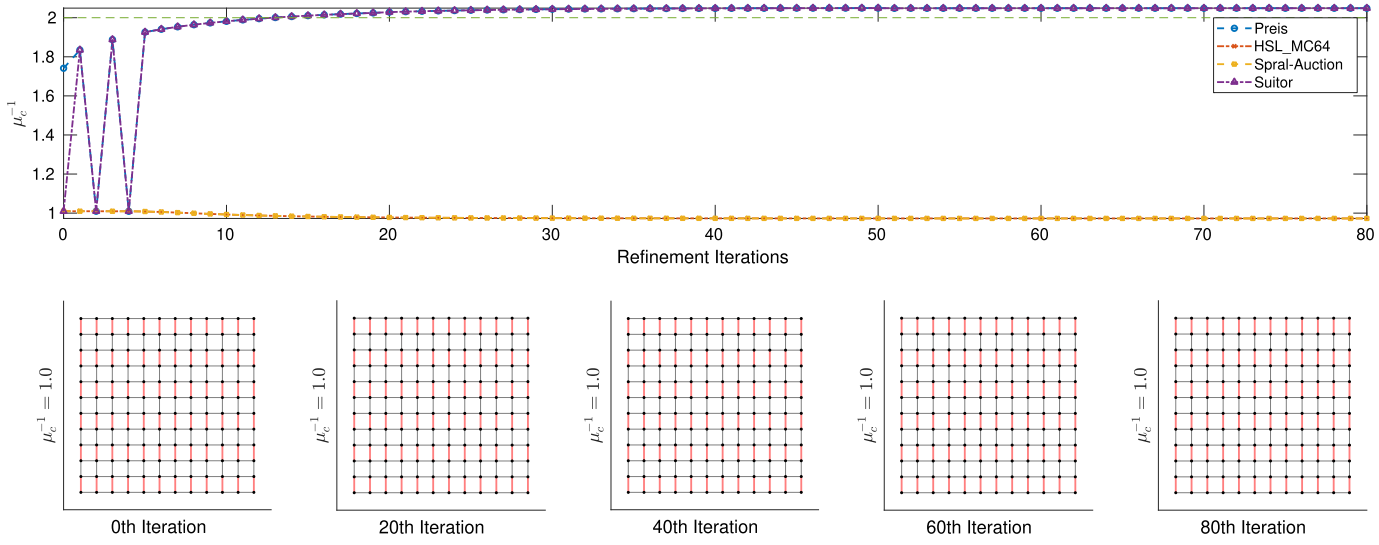

(b) Diffusion problem with y–axis oriented anisotropy $\varepsilon = 100$

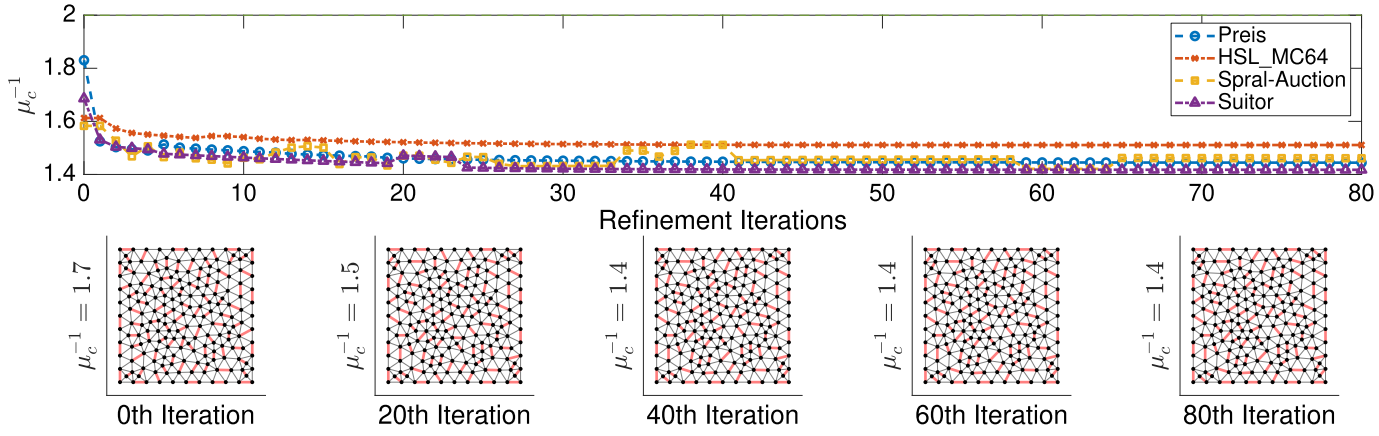

(c) Constant coefficient diffusion problem on an unstructured grid

Fig. 7. Refinement of the weight vector starting from the all one guess, and using the ℓ_1–Jacobi smoother. We report a graph containing the μ_c^{-1} constant up to 80 refinement steps for a single sweep of pairwise aggregation. The depicted aggregates are the ones obtained with the AUCTION algorithm.

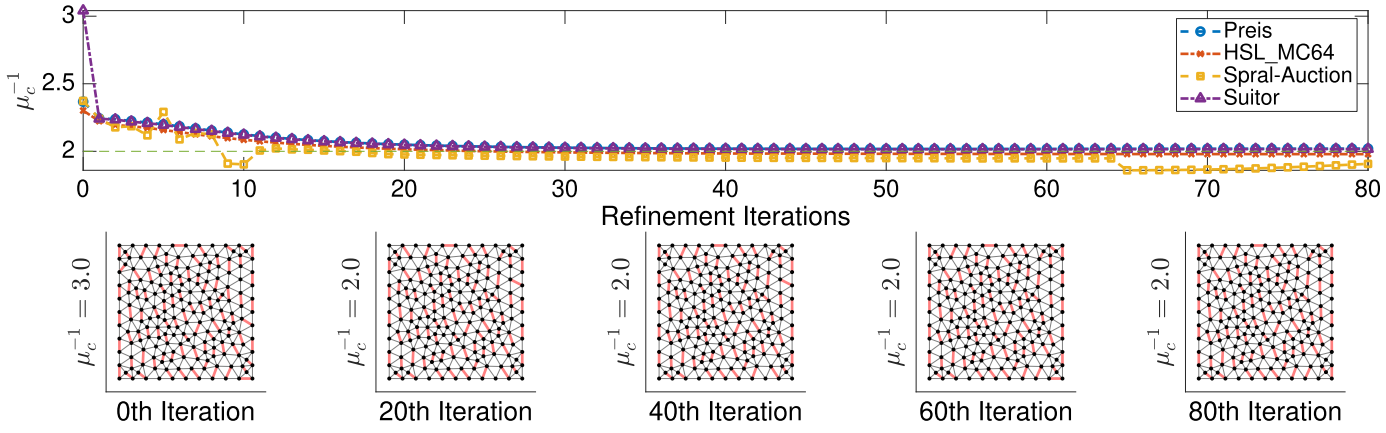

(a) $\theta = \pi/6$, $\varepsilon = 1e-2$

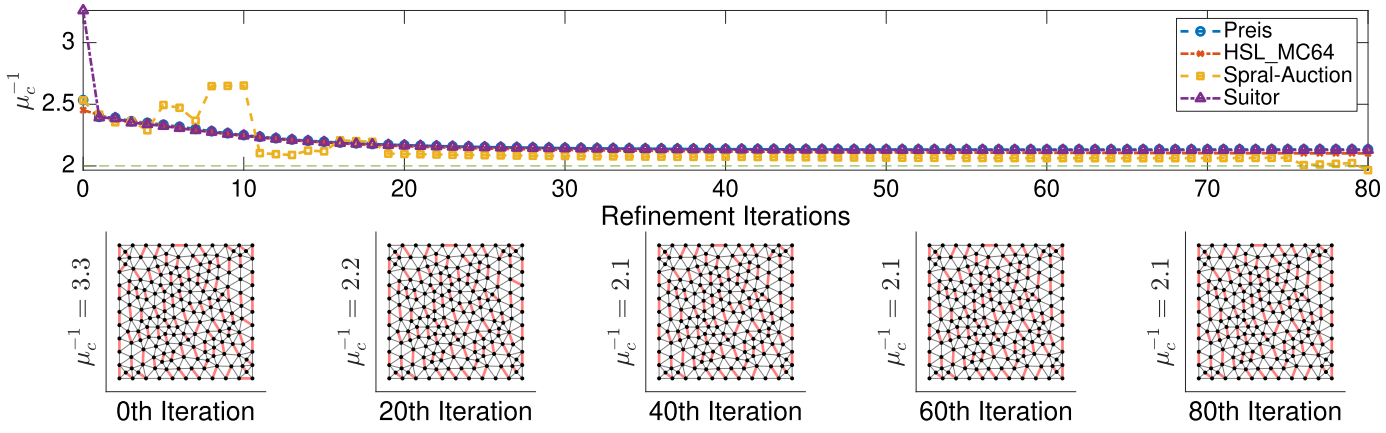

(b) $\theta = \pi/6$, $\varepsilon = 1e-3$

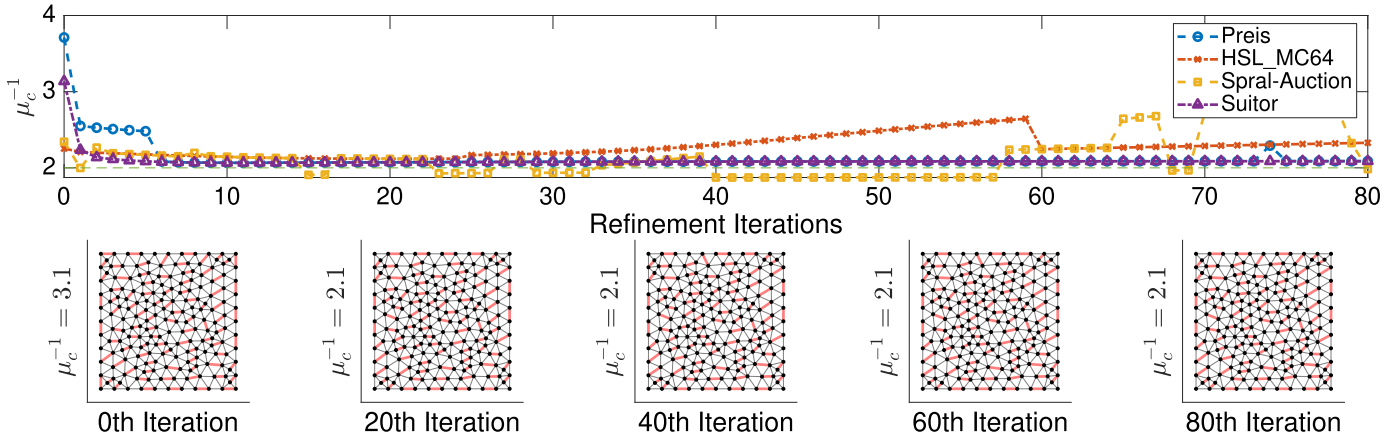

(c) $\theta = \pi/3$, $\varepsilon = 1e-2$

Fig. 8. Poisson problem on an unstructured grid with rotated anisotropy of angle θ, and modulus ε. Refinement of the weight vector starting from all one guess, and using the ℓ_1–Jacobi smoother. We report a graph containing the μ_c^{-1} constant up to 80 refinement steps for a single sweep of pairwise aggregation. The depicted aggregates are the ones obtained with the AUCTION algorithm.

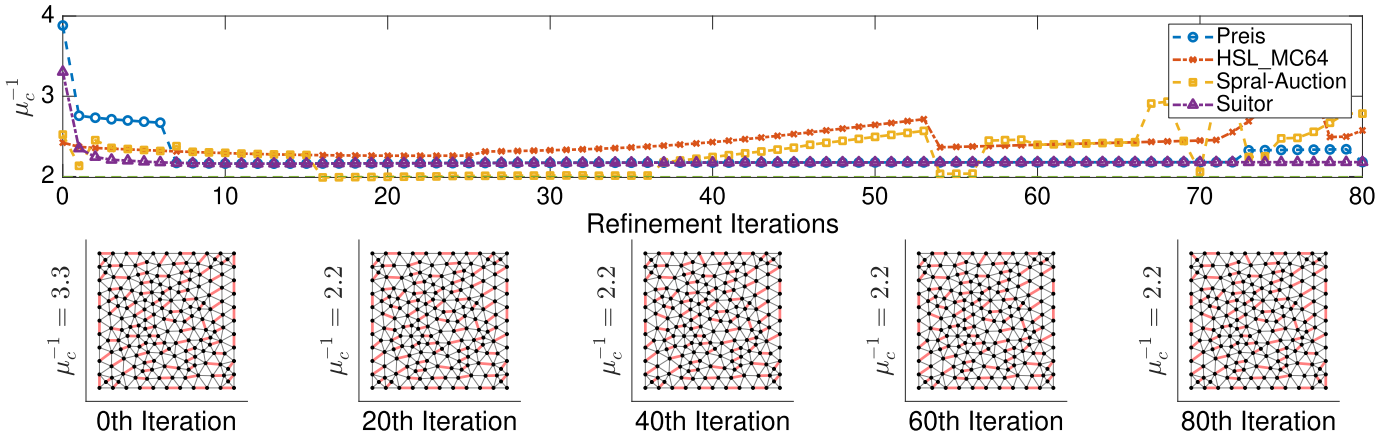

(d) $\theta = \pi/3$, $\varepsilon = 1e-3$

Fig. 8. (*continued*)

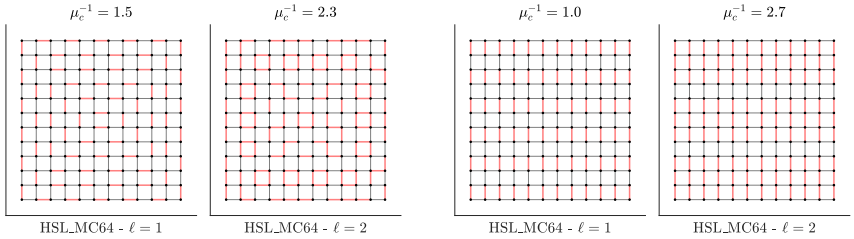

(a) Constant coefficient diffusion problem (b) Diffusion problem with y–axis oriented anisotropy $\varepsilon = 100$

Fig. 9. Aggregates obtained by using as weight vector $\mathbf{w}$ the eigenvector associated with the smallest eigenvalue as suggested by Proposition 3.

$$P_f = \begin{bmatrix} \tilde{P}_f \\ 0 \end{bmatrix} \in \mathbb{R}^{n \times n_p}, \text{ where } \tilde{P}_f = [\mathbf{p}_1^f, \ldots, \mathbf{p}_{n_p}^f] \text{ for } \mathbf{p}_j^f = \Pi_j \mathbf{w}_{e_{i \mapsto j}}^{\perp}.$$

By this construction, each relaxation scheme that is well defined for the block A_{ff} is then a *compatible relaxation*, i.e., a scheme that keeps the values of the coarse variables intact, and therefore makes the smoothing and coarse correction operators work each on the appropriate subspaces.

To validate numerically this claim we then look at the convergence radius $\rho(\cdot)$ of the iterative method induced by the restriction of the ℓ_1–Jacobi global smoother on the matrix A_{ff} in (25), i.e., we look at

$$\rho_f = \rho(I - M_{ff}^{-1} A_{ff}) < 1, \quad M_{ff} = P_f^T M P_f, \qquad A_{ff} = P_f^T A P_f, \tag{26}$$

where M is the iteration matrix of the ℓ_1–Jacobi global smoother for A. In Table 5 we report the value of ρ_f for each combination of test problem and matching algorithm, while setting the weight vector $\mathbf{w} = (1, 1, \ldots, 1)^T$, and the number of matching steps to $\ell = 1$.

If we compare the constants obtained here with the ones in the columns for $\ell = 1$ in the Tables 1, 2 and 3, we observe that the value of the ρ_f constants behaves consistently with quality measure μ_c within the same experiment, while it is harder to use it to compare among the aggregates for different test cases. This is specifically true for the case of the unstructured mesh, where, even if the quality of the aggregates seem to be degraded with respect to the corresponding finite difference case, the convergence ratio of the compatible relaxation is only mildly affected.

6. Conclusions

This paper has presented some theoretical results which complement the available computational evidence on the convergence properties of the *coarsening based on compatible weighted matching*. This is a purely algebraic and automatic procedure, exploiting unsmoothed aggregation for coarsening of general SPD matrices in AMG, introduced in [5,6]. We have shown that the necessary conditions for convergence of AMG, as stated in [2], are satisfied. Furthermore, we used the theory to have a quality measure of aggregates which we used as *a posteriori* guideline to analyze the effectiveness of different edge weights and maximum weight matching algorithms exploited in the coarsening procedure. We have applied the theory to different test cases arising from scalar elliptic PDEs, and we have shown that the good quality of the coarsening procedure is preserved in the case of using sub-optimal algorithms for computing maximum weight matching and that it appears also insensitive to anisotropy and discontinuities in the coefficients of the considered test cases. A possible generalization of our results to the case of smoothed aggregation, i.e., when prolongator operators are defined by $\overline{P} = (I - \omega D^{-1} A)P$ for some suitable ω, can be obtained by using results in [2, Lemma 9.3]. Likely, our final TL-AMG algorithm will have better convergence rate, as also demonstrated by experimental results in [27].

Funding

The research leading to these results received funding from Horizon 2020 Project "Energy oriented Centre of Excellence: toward exascale for energy" (EoCoE–II), Project ID: 824158. The first three authors are members of the INdAM–GNCS research group.

Data availability

The datasets generated during and/or analyzed during the current study are available in the GitHub repository, https://github.com/bootcmatch/BootCMatch.

Table 5
Convergence ratio ρ_f of the compatible relaxation scheme (26) for all the test problems. The coarse space is built from a single step of all the matching algorithms from Section 3.2 with weight vector choice $\mathbf{w} = (1, 1, \ldots, 1)^T$, and no refinement iterations.

n	HSL_MC64	PREIS	AUCTION	SUITOR
12	0.766	0.794	0.755	0.794
24	0.816	0.824	0.805	0.824
48	0.826	0.831	0.829	0.832
96	0.832	0.833	0.832	0.833

(a) Constant coefficient diffusion problem

n	HSL_MC64	PREIS	AUCTION	SUITOR
12	0.969	0.963	0.978	0.963
24	0.986	0.986	0.989	0.986
48	0.993	0.777	0.994	0.871
96	0.996	0.994	0.996	0.994

(b) Diffusion problem with y–axis oriented anisotropy $\epsilon = 100$

dofs	HSL_MC64	PREIS	AUCTION	SUITOR
185	0.803	0.802	0.796	0.799
697	0.831	0.846	0.811	0.843
2705	0.851	0.863	0.862	0.854
10657	0.882	0.917	0.873	0.869

(c) Rotated anisotropy $\theta = \pi/6$ and $\epsilon = 100$ on an unstructured grid

dofs	HSL_MC64	PREIS	AUCTION	SUITOR
185	0.723	0.756	0.729	0.725
697	0.735	0.750	0.754	0.743
2705	0.746	0.788	0.770	0.768
10657	0.785	0.794	0.775	0.800

(d) Constant coefficient problem on an unstructured grid

Acknowledgements

We thank the anonymous reviewers for their helpful comments which allowed us to improve the presentation of the material.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.camwa.2023.06.026.

References

[1] Y. Notay, An aggregation-based algebraic multigrid method, Electron. Trans. Numer. Anal. 37 (2010) 123–146.
[2] J. Xu, L.T. Zikatanov, Algebraic multigrid methods, Acta Numer. 26 (2017) 591–721, https://doi.org/10.1017/S0962492917000083.
[3] I. Marek, Aggregation Methods of Computing Stationary Distributions of Markov Processes, Birkhäuser, Basel, Basel, 1991, pp. 155–169.
[4] R. Blaheta, A multilevel method with correction by aggregation for solving discrete elliptic problems, Apl. Mat. 31 (5) (1986) 365–378.
[5] P. D'Ambra, P.S. Vassilevski, Adaptive AMG with coarsening based on compatible weighted matching, Comput. Vis. Sci. 16 (2) (2013) 59–76, https://doi.org/10.1007/s00791-014-0224-9.
[6] P. D'Ambra, S. Filippone, P.S. Vassilevski, BootCMatch: a software package for bootstrap AMG based on graph weighted matching, ACM Trans. Math. Softw. 44 (4) (2018) 39, https://doi.org/10.1145/3190647.
[7] F. Manne, M. Halappanavar, New effective multithreaded matching algorithms, in: 2014 IEEE 28th International Parallel and Distributed Processing Symposium, 2014, pp. 519–528.
[8] R. Preis, Linear time $\frac{1}{2}$-approximation algorithm for maximum weighted matching in general graphs, in: STACS 99 (Trier), in: Lecture Notes in Comput. Sci., vol. 1563, Springer, Berlin, 1999, pp. 259–269.
[9] I.S. Duff, J. Koster, On algorithms for permuting large entries to the diagonal of a sparse matrix, SIAM J. Matrix Anal. Appl. 22 (4) (2001) 973–996, https://doi.org/10.1137/S0895479899358443.
[10] X. Hu, J. Lin, L.T. Zikatanov, An adaptive multigrid method based on path cover, SIAM J. Sci. Comput. 41 (5) (2019) S220–S241, https://doi.org/10.1137/18M1194493.
[11] W. Xu, L.T. Zikatanov, Adaptive aggregation on graphs, J. Comput. Appl. Math. 340 (2018) 718–730, https://doi.org/10.1016/j.cam.2017.10.032.
[12] A. Napov, Y. Notay, Algebraic analysis of aggregation-based multigrid, Numer. Linear Algebra Appl. 18 (3) (2011) 539–564, https://doi.org/10.1002/nla.741.
[13] A. Napov, Y. Notay, An algebraic multigrid method with guaranteed convergence rate, SIAM J. Sci. Comput. 34 (2) (2012) A1079–A1109, https://doi.org/10.1137/100818509.
[14] Y. Notay, Aggregation-based algebraic multigrid for convection-diffusion equations, SIAM J. Sci. Comput. 34 (4) (2012) A2288–A2316, https://doi.org/10.1137/110835347.
[15] P.S. Vassilevski, Multilevel block factorization preconditioners, in: Matrix-Based Analysis and Algorithms for Solving Finite Element Equations, 1st edition, Springer Science & Business Media, NY, New York, 2008.
[16] A. Brandt, General highly accurate algebraic coarsening, in: Multilevel Methods, Copper Mountain, CO, 1999, Electron. Trans. Numer. Anal. 10 (2000) 1–20.
[17] O.E. Livne, Coarsening by compatible relaxation, Numer. Linear Algebra Appl. 11 (2–3) (2004) 205–227, https://doi.org/10.1002/nla.378.
[18] R.D. Falgout, P.S. Vassilevski, On generalizing the algebraic multigrid framework, SIAM J. Numer. Anal. 42 (4) (2004) 1669–1693, https://doi.org/10.1137/S0036142903429742.
[19] J.J. Brannick, R.D. Falgout, Compatible relaxation and coarsening in algebraic multigrid, SIAM J. Sci. Comput. 32 (3) (2010) 1393–1416, https://doi.org/10.1137/090772216.
[20] A. Brandt, J. Brannick, K. Kahl, I. Livshits, A.M.G. Bootstrap, SIAM J. Sci. Comput. 33 (2) (2011) 612–632, https://doi.org/10.1137/090752973.
[21] P. Vaněk, J. Mandel, M. Brezina, Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems, in: International GAMM-Workshop on Multi-Level Methods, Meisdorf, 1994, Computing 56 (3) (1996) 179–196, https://doi.org/10.1007/BF02238511.
[22] H.C. Elman, D.J. Silvester, A.J. Wathen, Finite Elements and Fast Iterative Solvers: with Applications in Incompressible Fluid Dynamics, 2nd edition, Numerical Mathematics and Scientific Computation, Oxford University Press, Oxford, 2014.
[23] R.D. Falgout, P.S. Vassilevski, L.T. Zikatanov, On two-grid convergence estimate, Numer. Linear Algebra Appl. 12 (2005) 471–494, https://doi.org/10.1002/nla.437.
[24] J. Brannick, F. Cao, K. Kahl, R.D. Falgout, X. Hu, Optimal interpolation and compatible relaxation in classical algebraic multigrid, SIAM J. Sci. Comput. 40 (3) (2018) A1473–A1493, https://doi.org/10.1137/17M1123456.
[25] P. D'Ambra, P.S. Vassilevski, Improving solve time of aggregation-based adaptive AMG, Numer. Linear Algebra Appl. 26 (E2269) (2019) 1–14, https://doi.org/10.1002/nla.2269.
[26] HSL, A collection of Fortran codes for large–scale scientific computation, http://www.hsl.rl.ac.uk/.
[27] P. D'Ambra, F. Durastante, S. Filippone, AMG preconditioners for linear solvers towards extreme scale, SIAM J. Sci. Comput. 43 (5) (2021) S679–S703, https://doi.org/10.1137/20M134914X.
[28] D.P. Bertsekas, Auction algorithms for network flow problems: a tutorial introduction, Comput. Optim. Appl. 1 (1) (1992) 7–66, https://doi.org/10.1007/BF00247653.
[29] M. Bernaschi, P. D'Ambra, D. Pasquini, AMG based on compatible weighted matching on GPUs, Parallel Comput. 29 (2020), https://doi.org/10.1016/j.parco.2019.102599.
[30] H.H. Kim, J. Xu, L.T. Zikatanov, A multigrid method based on graph matching for convection-diffusion equations, Numer. Linear Algebra Appl. 10 (1–2) (2003) 181–195, https://doi.org/10.1002/nla.317, dedicated to the 60th birthday of Raytcho Lazarov.
[31] A.H. Baker, R.D. Falgout, T.V. Kolev, U.M. Yang, Multigrid smoothers for ultraparallel computing, SIAM J. Sci. Comput. 33 (5) (2011) 2864–2887, https://doi.org/10.1137/100798806.