

PARALLEL ENERGY-MINIMIZATION PROLONGATION FOR ALGEBRAIC MULTIGRID*

CARLO JANNA[†], ANDREA FRANCESCHINI[‡], JACOB B. SCHRODER[§], AND LUKE OLSON[¶]

Abstract. Algebraic multigrid (AMG) is one of the most widely used solution techniques for linear systems of equations arising from discretized partial differential equations. The popularity of AMG stems from its potential to solve linear systems in almost linear time, that is with an $O(n)$ complexity, where n is the problem size. This capability is crucial at the present, where the increasing availability of massive HPC platforms pushes for the solution of very large problems. The key for a rapidly converging AMG method is a good interplay between the smoother and the coarse-grid correction, which in turn requires the use of an effective prolongation. From a theoretical viewpoint, the prolongation must accurately represent near kernel components and, at the same time, be bounded in the energy norm. For challenging problems, however, ensuring both these requirements is not easy and is exactly the goal of this work. We propose a constrained minimization procedure aimed at reducing prolongation energy while preserving the near kernel components in the span of interpolation. The proposed algorithm is based on previous energy minimization approaches utilizing a preconditioned restricted conjugate gradients method, but has new features and a specific focus on parallel performance and implementation. It is shown that the resulting solver, when used for large real-world problems from various application fields, exhibits excellent convergence rates and scalability and outperforms at least some more traditional AMG approaches.

Key words. algebraic multigrid, AMG, preconditioning, energy minimization, prolongation

MSC codes. 65N55, 68W10, 65F08

DOI. 10.1137/22M1513794

1. Introduction. With the increasing availability of powerful computational resources, scientific and engineering applications are becoming more demanding in terms of both memory and CPU time. For common methods used in the numerical approximation to partial differential equations (e.g., finite difference, finite volume, or finite element), the resulting approximation can easily grow to several millions or even billions of unknowns. The efficient solution to the associated sparse linear system of equations

$$(1.1) \quad \mathbf{Ax} = \mathbf{b},$$

either as a stand-alone system or as part of a nonlinear solve process, often represents a significant computational expense in the numerical application. Thus, research on sparse linear solvers continues to be a key topic for efficient simulation at large scales. One of the most popular sparse linear solvers is algebraic multigrid (AMG) [6, 7, 29]

*Submitted to the journal's Methods and Algorithms for Scientific Computing section August 4, 2022; accepted for publication (in revised form) May 18, 2023; published electronically September 26, 2023.

<https://doi.org/10.1137/22M1513794>

Funding: The work of the third author was supported by NSF grant DMS-2110917.

[†]M³E S.r.l., via Giambellino 7, 35129, Padova, Italy, and Department ICEA, University of Padova, via Marzolo 9, 35131, Padova, Italy (c.janna@m3eweb.it).

[‡]Corresponding author. Department ICEA, University of Padova, via Marzolo 9, 35131, Padova, Italy (andrea.franceschini@unipd.it).

[§]Department of Mathematics and Statistics, University of New Mexico, Albuquerque, NM 87131 USA (jbschroder@unm.edu).

[¶]Siebel Center for Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (lukeo@illinois.edu).

because of its potential for $O(n)$ computational cost for problems with n degrees-of-freedom.

A fast converging AMG method relies on the complementary action of relaxation (e.g., with weighted Jacobi) and coarse grid correction, which is a projection step focused on eliminating the error that is not reduced by relaxation. Even in a purely algebraic setting, the main algorithmic decisions in multigrid are often based on heuristics for elliptic problems. As a result, for more complex applications, traditional methods often break down, requiring additional techniques to improve accuracy with a careful eye on overall computational complexity.

Even with advanced AMG methods, robustness remains an open problem for a variety of applications, especially in parallel. Yet, there have been several advances in recent years that have significantly improved convergence in a range of settings. Adaptive AMG [21] and adaptive smoothed aggregation [10] are among early attempts to assess the quality of the AMG setup phase *during* the setup process, with the ability to adaptively improve the interpolation operators. Later works focus on extending the adaptive ideas to more general settings [22], and in particular, Bootstrap AMG [4] further develops the idea of adaptive interpolation with least-squares interpolation coupled with locally relaxed vectors and multilevel eigenmodes. Other advanced approaches have a focus on specific AMG components, such as energy minimization of the interpolation operator [23, 35, 30, 28, 25], generalizing the strength of connection procedure [27, 5], or by considering the nonsymmetric nature of the problem directly [26, 24].

While AMG robustness and overall convergence has improved with combinations of the advances above, the overarching challenge of controlling cost is persistent. In this paper, we make a number of related contributions with a focus on AMG effectiveness and efficiency at large scale. Our key contributions are as follows:

- The quality and sparsity of tentative interpolation is improved through a novel utilization of sparse QR and a new process for sparsity pattern expansion that targets locally full-rank matrices for improved mode interpolation constraints.
- We accompany the energy minimization construction of interpolation with new energy and convergence monitoring, thus limiting the total cost.
- We apply a new preconditioning technique for the energy minimization process based on Gauss–Seidel applied to the blocks.
- We present the nontrivial and efficient parallel implementation in detail; and
- we demonstrate improved convergence and computational complexity with several large scale experiments.

The remainder of this paper is as follows. We begin with the basics of AMG in section 2. In section 3, we derive the energy minimization process based on QR factorizations and introduce a method for monitoring reduction of energy in practice. Finally, we conclude with several numerical experiments in section 5 along with a discussion on performance.

2. Introduction to classical AMG. The effectiveness of AMG as a solver depends on the complementary relationship between relaxation and coarse-grid correction, where the error not reduced by relaxation on the fine grid (e.g., with weighted-Jacobi or Gauss–Seidel) is accurately represented on the coarse grid, where a complementary error correction is computed. For a more in-depth introduction to AMG; see the works [11, 31]. Here, we focus our description of AMG on the coarse grid and interpolation setup, which are most relevant to the rest of this paper.

Constructing the AMG coarse grid begins with a partition of the n unknowns of A into a C-F partition of n_f fine nodes and n_c coarse nodes: $\{0, \dots, n-1\} = \mathcal{C} \cup \mathcal{F}$. From this, we assume an ordering of A by F-points followed by C-points:

$$(2.1) \quad A = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{fc}^T & A_{cc} \end{bmatrix},$$

where, for example, A_{ff} corresponds to entries in A between two F-points. We also assume A is SPD so that $A_{cf} = A_{fc}^T$. In classical AMG, prolongation takes the form

$$(2.2) \quad P = \begin{bmatrix} W \\ I \end{bmatrix},$$

where W must be sparse (for efficiency) and represents interpolation from the coarse grid to fine grid F-points.

In constructing prolongation of the form (2.2), there are two widely accepted guidelines, the so-called *ideal* [9, 37] and *optimal* [36, 8] forms of prolongation. Although both of these are not feasible in practical applications, leading to very expensive and dense prolongation operators, the concepts behind their definition are valuable guides for constructing effective P .

Ideal prolongation is constructed by starting with the above C-F partition and constructing P_{id} as

$$(2.3) \quad P_{id} = \begin{bmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{bmatrix}.$$

Making P_{id} the goal for interpolation is motivated by Corollary 3.4 from the theoretical work [14]. Here, the main assumption is a classical AMG framework where P is of the form in (2.2).¹ In this setting, the choice of $W = -A_{ff}^{-1}A_{fc}$ minimizes the two-grid convergence of AMG relative to the choice of P , i.e., relaxation is fixed. Motivating our later energy minimization approach, P_{id} can be viewed as having zero energy rows, as AP_{id} is zero at all F-rows.

With *optimal* interpolation, the goal of interpolation is to capture the algebraically smoothest modes in $\text{span}(P)$, i.e., the modes left behind by relaxation. More specifically following [8], let v_1, v_2, \dots, v_n be the ordered eigenvectors of the generalized eigenvalue problem $Ax = \lambda Mx$, where M is the symmetrized relaxation matrix. (See [8] for more details.) Then, the two-grid convergence of AMG is minimized if

$$(2.4) \quad \text{span}(P) = \text{range}(v_1, v_2, \dots, v_{n_c}).$$

Note, that no assumptions on the structure of P are made, as in (2.2). Motivating our later energy minimization approach, (2.4) indicates that $\text{span}(P)$ should capture low-energy modes relative to relaxation, which our Jacobi or Gauss-Seidel preconditioned energy minimization approach will explicitly target. Moreover, our energy minimization approach will incorporate constraints which explicitly force certain vectors, corresponding to the smallest eigenvalues, to be in $\text{span}(P)$.

The idea of energy minimization AMG with constraints has been exploited for both symmetric and nonsymmetric operators in several works [23, 35, 30, 28, 25], and, though requiring more computational effort than classical interpolation formulas, often provides improved preconditioners that balance the extra cost.

¹The other assumptions are specific choices for the map to F-points $S = [I, 0]$, for the map to C-points $R = [0, I]$, and for relaxation $X = \|A\|I$.

TABLE 3.1
Symbols used in this section.

Symbol	Description
A	matrix
$a^{(r)}$	column vector containing the nonzeros of A ordered row-wise
$a^{(c)}$	column vector containing the nonzeros of A ordered column-wise
\mathcal{A}	nonzero pattern of A
\bar{A}	binary matrix obtained replacing the nonzeros of A with ones
v	vector
\bar{v}	vector containing only the nonzeros of v

3. Energy minimization prolongation. To help the reader with the heavy notation needed for this section, Table 3.1 provides a reference for the main symbols.

The energy minimization process combines the key aspects of ideal and optimal prolongation. To define this, we first introduce V , a basis for the *near kernel* of A or the *lowest energy* modes of A . Next, the optimization problem is as follows: find

$$(3.1) \quad P = \underset{P \in \mathcal{P}}{\operatorname{argmin}} (\operatorname{tr}(P^T A P)) \rightarrow \text{objective function,}$$

subject to the constraint

$$(3.2) \quad V \subseteq \operatorname{range}(P).$$

Here, the space \mathcal{P} denotes all possible sparse matrices with nonzero entries only at specified locations defined by the nonzeros of a predetermined binary matrix \bar{P}_k .

Indeed, for practical use in AMG, the prolongation operator must be sparse, therefore construction begins by defining a sparse nonzero pattern for P . Assume that a strength of connection (SoC) matrix S is provided where nonzero entries denote a strong coupling between degrees-of-freedom in the discretization matrix [29, 27, 5]. Next, let T be a *tentative* prolongation operator with nonzero pattern \mathcal{T} , to be used as an initial guess for P .² We next obtain an enlarged sparsity pattern \mathcal{P} by growing \mathcal{T} to include all strongly connected neighbors up to distance k . Denoting with \bar{P} the binary matrix obtained from P by replacing its nonzeros with unitary entries, this is equivalent to

$$(3.3) \quad \bar{P}_k = \operatorname{spones}(S^k \bar{T}),$$

where \mathcal{P} is defined by the pattern of \bar{P}_k (see [28]) and $\operatorname{spones}(\cdot)$ transforms the argument into a corresponding binary matrix, i.e., with the same pattern.

For a constraint condition satisfying (3.2), we start by splitting the near kernel basis V with the same C-F splitting as in (2.1):

$$(3.4) \quad V = \begin{bmatrix} V_f \\ V_c \end{bmatrix}.$$

Recalling the form of interpolation (2.2), the near kernel requirement for P becomes

$$(3.5) \quad W V_c = V_f,$$

²See subsection 3.2 for our contributions regarding the construction of T , which is based on the adaptive algorithm [22, 17]. T can be defined similarly to the tentative prolongation from smoothed aggregation AMG [33] in that T interpolates the basis V , but needs further improvement, for example, with energy minimization.

which is a set of n_f conditions on the rows of W . By denoting w_i^T as the i th row of W (and P) and v_i^T as the i th row V , condition (3.5) is then exploited row-wise as

$$(3.6) \quad V_c^T w_i = v_i \quad \forall i \in \mathcal{F}.$$

Using the sparsity pattern \mathcal{P} , we rewrite (3.6) for only the nonzeros in each row w_i . Letting the index set \mathcal{J}_i be the nonzero column indices in the i th row of P , this becomes

$$(3.7) \quad V_c(\mathcal{J}_i, :)^T \bar{w}_i = v_i \quad \forall i \in \mathcal{F},$$

where $\bar{w}_i = w_i(\mathcal{J}_i)$ collects only the nonzeros of w_i . It is important to note that for each of the n_f fine points, the constraints (3.7), are independent of each other, because each entry of W appears in only one constraint. Next, we write the constraints in the following matrix form:

$$(3.8) \quad \tilde{B}^T p^{(r)} = g^{(r)},$$

where $p^{(r)}$ is the row-wise vectorization of P , $g^{(r)}$ collects the v_i in (3.7) into a single vector, and where \tilde{B}^T is a block diagonal matrix composed of $V_c(\mathcal{J}_i, :)^T$ due to the independence of constraints.

Lastly, we describe the reduced linear system framework for approximating (3.1). Minimizing (3.1) is equivalent to minimizing the energy of the individual columns of P on the prescribed nonzero pattern:

$$(3.9) \quad p_i = \operatorname{argmin}_{p_i \in \mathcal{P}_i} p_i^T A p_i,$$

where p_i is the i th column of P and \mathcal{P}_i is the sparsity pattern of column i . The generic p_i represents one term in the trace of $P^T A P$ in (3.1).

Next, let \mathcal{I}_i be the set of nonzero row indices of the i th column of W , and let \bar{h}_i be the vector collecting the nonzero entries of the i th column of W . Then the minimization in (3.9) defines \bar{h}_i with

$$(3.10) \quad A(\mathcal{I}_i, \mathcal{I}_i) \bar{h}_i = -A(\mathcal{I}_i, i) \quad \forall i \in \mathcal{C},$$

where $A(\mathcal{I}_i, \mathcal{I}_i)$ is a square, relatively dense, submatrix of A corresponding to the allowed nonzero indices \mathcal{I}_i and $A(\mathcal{I}_i, i)$ is a vector corresponding to the i th column of A at the allowed nonzero indices. Also in this case, each column of P satisfies (3.9) independently. Thus, denoting by $f^{(c)}$ the vector collecting each $-A(\mathcal{I}_i, i)$, the minimization (3.9) is recast as

$$(3.11) \quad K p^{(c)} = f^{(c)},$$

where K is block diagonal with $A(\mathcal{I}_i, \mathcal{I}_i)$ as the i th block and $p^{(c)}$ is the column ordered version of $p^{(r)}$.

The two conditions, one on the range of P in (3.8) and one on the minimality of the objective function in (3.11), together form a constrained minimization problem, whose solution is the desired energy minimal prolongation. Casting this problem using Lagrange multipliers results in the saddle point system

$$(3.12) \quad \begin{bmatrix} K & B \\ B^T & 0 \end{bmatrix} \begin{bmatrix} p^{(c)} \\ \lambda \end{bmatrix} = \begin{bmatrix} f^{(c)} \\ g^{(c)} \end{bmatrix}.$$

The elements in (3.12) are the same as those defined in (3.8) and (3.11), with the exception that B is \tilde{B} reordered to operate on the column-ordered vector $p^{(c)}$, and that λ is the vector of Lagrange multipliers whose values are not needed for the purpose of setting up the prolongation. We emphasize that, with the entries of $p^{(c)}$ enumerated columnwise with respect to P , K is block diagonal. Likewise, if the vectorization of P is obtained following the rows of P , then B becomes block diagonal. Unfortunately, there is no sorting able to make both K and B block diagonal at the same time. Nevertheless, it is possible to take advantage of this underlying structure in the numerical implementation, as will be shown later. Leveraging the block structure of B is also important, because, as we will see in section 3.1, our algorithm to minimize energy requires several applications of the orthogonal projector given as

$$(3.13) \quad \Pi_B = I - B(B^T B)^{-1} B^T.$$

The system (3.12) follows closely the method from [28]. In sections 3.2–3.4, we outline our proposed improvements to energy minimization.

3.1. Minimization through Krylov subspace methods. Following [28], energy minimization proceeds by starting with a tentative prolongation, T , that satisfies the near kernel constraints (see (3.5)). From here on, we will vectorize matrices only *columnwise*, unless clearly stated, thus the superscript (c) will be ignored.

Denoting by t the tentative prolongation in vector form these constraints read

$$(3.14) \quad B^T t = g.$$

Defining the final prolongation as the tentative t plus a correction δ_p gives

$$(3.15) \quad p = t + \delta_p.$$

Then, the problem is recast as finding the optimal correction ΔP^* :

$$(3.16) \quad \Delta P^* = \operatorname{argmin}_{\Delta P \in \mathcal{P}} (\operatorname{tr}((T + \Delta P)^T A (T + \Delta P))),$$

subject to the constraint $B^T \delta_p = 0$, where δ_p is the vector form of ΔP . By recalling the definition of T —i.e., $T = [W_0^T, I]^T$ —and that ΔP has nonzero components only in ΔW —i.e., $\Delta P = [\Delta W^T, 0]^T$ —the C-F partition (2.1) of $\operatorname{tr}(P^T A P)$ reads

$$(3.17) \quad \begin{aligned} \operatorname{tr}((T + \Delta P)^T A (T + \Delta P)) &= \underbrace{\operatorname{tr}(\Delta W^T A_{ff} \Delta W) + 2 \operatorname{tr}(W_0^T A_{ff} \Delta W) + 2 \operatorname{tr}(A_{fc}^T \Delta W)}_{\text{dependent on } \Delta P} \\ &\quad + \underbrace{\operatorname{tr}(W_0^T A_{ff} W_0) + 2 \operatorname{tr}(W_0^T A_{fc}) + \operatorname{tr}(A_{cc})}_{\text{independent of } \Delta P}. \end{aligned}$$

Using the preset nonzero pattern for ΔW , the problem is rewritten in vector form to minimize *only* the terms depending on ΔP :

$$(3.18) \quad \operatorname{argmin}_{\delta w} (\delta w^T K \delta w + 2w_0^T K \delta w + 2f^T \delta w),$$

subject to the constraint

$$(3.19) \quad B^T \delta w = 0,$$

where K and f are defined as in (3.12) and δw and w_0 are the vector forms of ΔW and W_0 , respectively. Note that the matrix form of δw is nonzero only at fine nodes and satisfies the sparsity pattern of \mathcal{P} at fine nodes.

This minimization can be performed using (preconditioned) conjugate gradients by ensuring that both the initial solution and the search direction satisfy the constraint [28]. To do this, return to the orthogonal projector

$$(3.20) \quad \Pi_B = I - B(B^T B)^{-1} B^T,$$

and apply conjugate gradients to the singular system

$$(3.21) \quad \Pi_B K \Pi_B \delta w = -\Pi_B (f + K w_0),$$

starting from $\delta w = 0$. Due to its block diagonal structure, it is straightforward to find a QR decomposition of $B = QR$, and the projection simply becomes

$$(3.22) \quad \Pi_B = I - QQ^T.$$

Finally, introducing $K_\Pi = \Pi_B K \Pi_B$ and $\bar{f} = f + K w_0$, the Krylov subspace built by conjugate gradients is

$$(3.23) \quad \mathcal{K}_m = \text{span}\{\Pi_B \bar{f}, K_\Pi \bar{f}, K_\Pi^2 \bar{f}, \dots, K_\Pi^m \bar{f}\}.$$

This is equivalent to applying the nullspace method [3] to the saddle-point system (3.12).

3.2. Improved tentative interpolation T and orthogonal projection Π_B with sparsity pattern expansion and sparse QR . A crucial point for energy minimization interpolation is the availability of a tentative prolongation T that satisfies the near kernel representability constraint (3.5). While this is relatively straightforward for scalar diffusion equations where V has only one column and T corresponds to nonoverlapping basis functions, it is not trivial for vector-valued PDEs such as elasticity. One specific difficulty is that while forming the i th constraint equation (3.7), we must ensure that $V_c(\mathcal{J}_i, :)$ is full rank. If it is not full rank, then no prolongation operator is able to satisfy (3.5) and the solution to (3.12) is not possible, in general. We consider two possible remedies:

1. Add strongly connected neighbors to the pattern of the i th row of P to enlarge \mathcal{J}_i until $V_c(\mathcal{J}_i, :)$ is full-rank (i.e., sparsity pattern expansion) or
2. compute the least square solution of (3.7) as is done in [28].

A novel aspect of this work is our pursuit of sparsity pattern expansion. We find that this careful construction of the sparsity pattern, which guarantees that each constraint is exactly satisfied as $V_c(\mathcal{J}_i, :)$ is always full rank, greatly improves performance on some problems.

To accomplish this task, we adopt a dynamic-pattern, least-squares fit (LSF) procedure that satisfies (3.5) or, equivalently, (3.14). For each row of W (corresponding to a fine node i), this is equivalent to satisfying the local dense system (3.7). For simplicity, we rewrite (3.7) by dropping the row subscript i , with $w = \bar{w}_i$ and $v = v_i$, yielding

$$(3.24) \quad \mathbb{B}w = v,$$

where $\mathbb{B} = V_c(\mathcal{J}_i, :)^T$ corresponds to a diagonal block of B^T in (3.12), when the nonzero entries of P are enumerated row-wise.

Algorithm 3.1. Tentative prolongation set-up.

```

1: procedure PTENT_SETUP( $S, V, l_{\max}$ )
2:   input:  $S$  – strength of connection matrix
3:    $V$  – near kernel modes
4:    $l_{\max}$  – maximum interpolation distance
5:   output:  $\hat{t}$  – initial tentative prolongation
6:   for all fine nodes  $i$  do
7:     Set  $l = 0$ ;
8:     while  $l < l_{\max}$  do
9:       Set  $l = l + 1$ ;
10:      Form  $\mathcal{N}_i$  with the strong neighbors of  $i$  up to distance  $l$ ;
11:      Select the best columns of  $V_c(\mathcal{N}_i, :)^T$  using max vol to form  $\mathbb{B}$ ;
12:      Compute the least-squares solution to  $\mathbb{B}\mathbf{w} = \mathbf{v}$ ;
13:      if  $\|\mathbf{v} - \mathbb{B}\mathbf{w}\|_2 = 0$  then
14:        Assign  $\mathbf{w}^T$  as  $i$ -th row of  $\hat{t}$ ;
15:        break;
16:      end if
17:    end while
18:  end for
19: end procedure

```

Considering this generic fine node represented by (3.24), if there are a sufficient number of coarse node neighbors, then \mathbb{B} has more columns than rows. Hence if we assume a full rank \mathbb{B} , then (3.24) is an underdetermined system and can be solved in several ways. In order to have a sparse solution \mathbf{w} , we choose a minimal set of columns of \mathbb{B} using the *max vol* algorithm [19, 16] to have the best basis. Here, we satisfy (3.24) exactly. We note that a related *max vol* approach to computing C-F splittings is used in [8].

Remark 3.1. We adopt this form of a *QR* factorization with *max vol*, that is as sparse as possible, in order to improve the complexity of our algorithm and quality of T . While it is a relatively minor change to the algorithm’s structure, we count it as a useful novelty of our efficient implementation.

If, on the contrary, the number of neighboring coarse nodes is not sufficient ($\mathbf{v} \notin \text{span}(\mathbb{B})$), then (3.24) cannot be satisfied because it is overdetermined. This may occur not only when \mathbb{B} is *skinny*, i.e., the number of columns is smaller than the number of rows, but more often because \mathbb{B} is rank deficient even if it has a larger number of columns, i.e., it is *wide*. In elasticity problems and, in particular, with shell finite elements, this issue arises often in practice with standard distance one coarsening, where during the coarsening process, some fine nodes may occasionally remain isolated. Our strategy is to gradually increase the interpolation distance for violating nodes where $\mathbf{v} \notin \text{span}(\mathbb{B})$, thus widening the interpolatory set (i.e., adding columns to \mathbb{B}) where it is necessary. Algorithm 3.1 describes how to set-up \hat{t} , the vector form of initial tentative prolongation \hat{T} . (Algorithm 3.2, described later, will further process \hat{T} for the final tentative prolongation T).

Remark 3.2. Avoiding a *skinny* or rank deficient local \mathbb{B} block is also important for the construction of the orthogonal projection $\Pi_B = I - B(B^T B)^{-1} B^T$ that maps vectors of \mathbb{R}^n to $\text{Ker}(B^T)$. Note that Π_B is used not only to correct the conjugate

gradients search direction, but also to ensure the initial prolongation satisfies the near kernel constraint. In general, the size of \mathbb{B} during energy minimization is larger than when constructing tentative prolongation because of the additional sparsity pattern expansion in (3.3). Thus, this “rank-deficiency” issue is ameliorated, but there are pathological cases where it has been observed in practice.

We now describe our procedure for computing local blocks of Π_B , called $\Pi_{\mathbb{B}}$, and a single row of the final tentative prolongation T , called w_0^T . The global procedure to build T and Π_B is then obtained by repeating this local algorithm for each fine row. Denote by \widehat{w}_0^T the *starting* tentative prolongation row. We use the word *starting*, because, in general, we can receive a tentative prolongation that does not satisfy the constraint. That is, we may have

$$(3.25) \quad \mathbb{B}\widehat{w}_0 \neq v.$$

Denote by n_l and m_l the dimensions of the local system, so that $\mathbb{B} \in \mathbb{R}^{n_l \times m_l}$. To fulfill condition (3.24), we must find a correction to \widehat{w}_0 , say δ , such that

$$(3.26) \quad \mathbb{B}\delta = v - \mathbb{B}\widehat{w}_0 = r$$

and then set

$$(3.27) \quad w_0 = \widehat{w}_0 + \delta.$$

To enforce condition (3.19) efficiently, we construct an orthonormal basis of $\text{range}(\mathbb{B})$, say Q , that gives rise to the desired local orthogonal projector:

$$(3.28) \quad \Pi_{\mathbb{B}} = I - QQ^T.$$

If the prolongation pattern is large enough, the vast majority of the above local problems will be such that $n_l \leq m_l$ with \mathbb{B} being also full-rank, i.e., $\text{rank}(\mathbb{B}) = n_l$. Thus, an *economy-size* QR decomposition is first performed on \mathbb{B} , $\mathbb{B} = QR$ with $Q \in \mathbb{R}^{n_l \times m_l}$ and $R \in \mathbb{R}^{m_l \times m_l}$, and Q is used to form the local projector. Then, through the same QR decomposition, we compute δ as the least norm solution of the underdetermined system (3.26):

$$(3.29) \quad \delta = (\mathbb{B}^T \mathbb{B})^{-1} \mathbb{B}^T r = (R^T Q^T Q R)^{-1} R^T Q^T r = R^{-1} Q^T r.$$

Note that any solution to (3.26) would be equivalent, as the optimal choice in terms of global prolongation energy is later computed by the restricted CG algorithm. Since the initial tentative prolongation arises from the LSF set-up, it should already fulfill (3.26) and $r \equiv 0$. However, in the most difficult cases even extending the interpolatory set with large distance l_{\max} is not sufficient to guarantee an exact interpolation of the near kernel for all the fine nodes. For these fine nodes— i.e., when $n_l > m_l$ or \mathbb{B} is not full rank—we compute an SVD decomposition of \mathbb{B} :

$$(3.30) \quad \mathbb{B} = U \Sigma V^T.$$

From the diagonal of Σ , we determine the rank of \mathbb{B} , say k_l , and use the first k_l columns of U to form Q and thus $\Pi_{\mathbb{B}}$. From a numerical viewpoint, the rank is the number of singular values larger than $1/\text{condmax}$ times the largest singular value. In our numerical experiments, we used $\text{condmax} = 10^{10}$. The value of condmax has a negligible impact on the overall AMG convergence. It is enough to avoid too small

Algorithm 3.2. Energy minimization set-up.

```

1: procedure EMIN_SETUP( $B, g, \hat{t}$ )
2:   input:  $B$  – block diagonal constraint matrix, as in (3.14)
3:            $g$  – constraint right-hand side
4:            $\hat{t}$  – initial tentative prolongation
5:   output:  $\Pi_B$  – projection matrix, constructed blockwise
6:            $t$  – final tentative prolongation
7:   for all fine nodes  $i$  do
8:     Gather  $\mathbb{B}$ ,  $\mathbf{v}$  and  $\widehat{\mathbf{w}}_0$  for row  $i$ , as in (3.25);
9:     Compute  $\mathbf{r} = \mathbf{v} - \mathbb{B}\widehat{\mathbf{w}}_0$ ;
10:    FAIL_QR = false;
11:    if  $n_l \geq m_l$  then
12:      Compute economy size QR of  $\mathbb{B}$ :  $\mathbb{B} = QR$ ;
13:      if  $\text{rank}(R) < m_l$  then
14:        Set FAIL_QR = true;
15:      else
16:        Compute  $\delta = R^{-1}Q^T\mathbf{r}$ ;
17:      end if
18:    end if
19:    if  $n_l < m_l$  or FAIL_QR then
20:      Compute SVD of  $\mathbb{B}$ :  $\mathbb{B} = U\Sigma V^T$ ;
21:      Determine  $k = \text{rank}(\mathbb{B})$  using the diagonal of  $\Sigma$ ;
22:      Set  $Q = U(:, 1:k)$ ;
23:      Compute  $\delta = V\Sigma^\dagger U^T\mathbf{r}$ ;
24:    end if
25:    Compute  $\mathbf{w}_0 = \widehat{\mathbf{w}}_0 + \delta$  and insert  $\mathbf{w}_0^T$  as row  $i$  of  $t$ ;
26:    Set  $i$ th block of  $\Pi_B$  as  $I - QQ^T$ ;
27:  end for
28: end procedure

```

values (e.g., $\leq 10^4$), which prevent the correct representation of the kernel, or too large values (e.g., $\geq 10^{14}$), which may cause too large weights in T . Finally, since in this case it could be impossible to satisfy the constraint because the system is overdetermined, we use the least square solution to (3.26) to compute the correction for $\widehat{\mathbf{w}}_0$:

$$(3.31) \quad \delta = V\Sigma^\dagger U^T\mathbf{r}.$$

It is important to recognize that for these specific fine nodes, energy minimization cannot reduce the energy, because the constraint does not leave any degree of freedom. Consequently, this situation should be avoided when selecting the coarse variables and the prolongation pattern, because a prolongation violating the near kernel constraint will likely fail in representing certain smooth modes. The pseudocode to set-up Π_B and correct the initial tentative prolongation T (after Algorithm 3.1) is provided in Algorithm 3.2.

3.3. Improved stopping criterion and energy monitoring for CG-based energy minimization. Stopping criteria plays an important role in the overall cost and effectiveness of energy minimization. Here, we introduce a measure for monitoring energy and halting in the algorithm.

Since CG often converges quickly for energy minimization, it is common to fix the number of iterations in advance [23, 28]. However, in the case of a more challenging problem, several iterations may be needed, thus requiring an accurate stopping criterion. One immediate option is to use the relative residual, yet this may not be a close indicator of *energy*. In the following, we analyze CG for a generic $Ax = b$; however, our observations extend to PCG as well.

In the CG algorithm, once the search direction p_k is defined at iteration k , the scalar α is computed:

$$(3.32) \quad \alpha_k = \frac{p_k^T r_k}{p_k^T A p_k},$$

in such a way that the new approximation $x_{k+1} = x_k + \alpha_k p_k$ minimizes the square of the energy norm of the error, i.e.,

$$(3.33) \quad E_k = (x_k - h)^T A (x_k - h),$$

where $h = A^{-1}b$ is the true solution. The difference in energy $\Delta E_{k+1} = E_{k+1} - E_k$ between two successive iterations k and $k + 1$ can be computed as

$$(3.34) \quad \begin{aligned} \Delta E_{k+1} &= (x_k + \alpha_k p_k)^T A (x_k + \alpha_k p_k) - 2b^T (x_k + \alpha_k p_k) - x_k^T A x_k + 2b^T x_k \\ &= 2\alpha_k x_k^T A p_k + \alpha_k^2 p_k^T A p_k - 2\alpha_k b^T p_k = -2\alpha_k r_k^T p_k + \alpha_k^2 p_k^T A p_k \\ &= \alpha_k (-2r_k^T p_k + \alpha_k p_k^T A p_k) = -\alpha_k (p_k^T r_k) = -\frac{(p_k^T r_k)^2}{p_k^T A p_k} < 0. \end{aligned}$$

From (3.32) and (3.34), it is possible to measure, with minimal cost, the energy decrease provided by the $(k + 1)$ st iteration. Indeed, by noting that α_k is computed as the ratio between the two values $\alpha_{\text{num}} = p_k^T r_k$ and $\alpha_{\text{den}} = p_k^T A p_k$, the energy decrease reads

$$(3.35) \quad \Delta E_{k+1} = \frac{\alpha_{\text{num}}^2}{\alpha_{\text{den}}}.$$

The relative value of the energy variation with respect to the initial variation (first iteration) is monitored and convergence is achieved when energy is sufficiently reduced:

$$(3.36) \quad \frac{\Delta E_k}{\Delta E_1} \leq \tau$$

for a small user-defined τ .

3.4. Improved preconditioning for CG-based energy minimization. Before introducing PCG, we present some important properties of matrices K and B , that we leverage in the design of effective preconditioners for energy minimization. In particular, we will see that, thanks to the nonzero structures of K and B , pointwise Jacobi or Gauss–Seidel iterations prove particularly effective in preconditioning the projected block $\Pi_B K \Pi_B$.

Let us assume that the vector form of prolongation p has been obtained by collecting the nonzeros of P *row-wise*, so that B is block diagonal. Denoting by Q the matrix collecting an orthonormal basis of $\text{range}(B)$, and Z an orthonormal basis of $\text{ker}(B)$, by construction we have

$$(3.37) \quad [Q \ Z][Q \ Z]^T = [Q \ Z]^T [Q \ Z] = I,$$




FIG. 3.1. Nonzero pattern of the columns of B , Q , and Z corresponding to a specific row of the prolongation P .

i.e., the matrix $[Q \ Z]$ is square and orthogonal. Moreover, since B is block diagonal, both Q and Z are block diagonal and can be easily computed and stored. We note that, by construction, each column of Q and Z refers to a specific row of the prolongation, and, due to the block diagonal pattern chosen for B , also each column of Q and Z is nonzero only in the positions corresponding to the entries of p collecting the nonzeros of a specific row of P , as is schematically shown in Figure 3.1. More precisely, using the same notation as in (3.7), let us define \mathcal{L}_i as the set of indices in p corresponding to the nonzero entries in the i th row of P so that

$$(3.38) \quad p(\mathcal{L}_i) = P(i, \mathcal{J}_i)$$

and define $\mathcal{J}_{B,i}$ and $\mathcal{J}_{Z,i}$ as the set of columns of B and Z , referring to the i th row of P . Then,

$$(3.39) \quad \left. \begin{aligned} B(k, \mathcal{J}_{B,i}) &= 0 \\ Q(k, \mathcal{J}_{B,i}) &= 0 \\ Z(k, \mathcal{J}_{Z,i}) &= 0 \end{aligned} \right\} \text{ if } k \notin \mathcal{L}_i.$$

We are now ready to state two theorems that will be useful in explaining the choice of our preconditioners.

THEOREM 3.1. *The diagonal of the projected matrix $Z^T K Z$ is equal to the projection of the diagonal of K :*

$$(3.40) \quad \text{diag}(Z^T K Z) = \text{diag}(Z^T D_K Z),$$

where D_K is the matrix collecting the diagonal entries of K . Moreover, $Z^T D_K Z$ is a diagonal matrix.

Proof. Let us consider the block of columns of Z relative to row i , that is $Z(:, \mathcal{J}_{Z,i})$, remembering that it is nonzero only for the row indices in \mathcal{L}_i . As a consequence, the square block H_i obtained by pre- and postmultiplying K by $Z(:, \mathcal{J}_{Z,i})$ is computed as

$$(3.41) \quad H_i(j, k) = \sum_{r \in \mathcal{L}_i} \left(\sum_{s \in \mathcal{L}_i} Z(s, k) K(r, s) \right) Z(r, j) \quad \text{for } j, k \in \mathcal{J}_{Z,i}.$$

However, $K(r, s)$ for $r, s \in \mathcal{L}_i$ represents the connection between $P(i, j_r)$ and $P(i, j_s)$, which is nonzero only for $j_r = j_s$, i.e., $r = s$, because in K there is no connection between different columns of P . Moreover, due to (3.10), $K(r, r) = A(i, i)$ for every $r \in \mathcal{L}_i$. As the columns of Z are orthonormal by construction, $Z^T Z = I$, it immediately follows that the square block H_i is diagonal with all its nonzero entry equal to $A(i, i)$.

Algorithm 3.3. Preconditioned conjugate gradients for energy minimization.

```

1: procedure EMIN_PCG(maxit,  $\tau$ ,  $K$ ,  $f$ ,  $\Pi_B$ ,  $M$ ,  $T$ )
2:   input: maxit – maximum iterations
3:            $\tau$  – energy convergence tolerance
4:            $K$  – system matrix (applied matrix-free with  $A$ )
5:            $f$  – right-hand side  $f$  from (3.11)
6:            $\Pi_B$  – projection matrix
7:            $M$  – preconditioner
8:            $T$  – tentative prolongation
9:   output:  $P$  – final prolongation
10:  Extract global weight vector  $w_0$  row-wise from  $T$ 
11:   $\Delta w = 0$ ;
12:   $r = -\Pi_B(f + Kw_0)$ ;
13:  for  $k = 1, \dots, \text{maxit}$  do
14:     $z = \Pi_B M^{-1} r$ ;
15:     $\gamma = r^T z$ ;
16:    if  $i = 1$  then
17:       $y = z$ ;
18:    else
19:       $\beta = \gamma / \gamma_{old}$ ;
20:       $y = z + \beta y$ ;
21:    end if
22:     $\gamma_{old} = \gamma$ ;
23:     $\check{y} = \Pi_B K y$ ;
24:     $\alpha = \gamma / (y^T \check{y})$ ;
25:     $\Delta E_k = \gamma \alpha$ 
26:    if  $\Delta E_k < \tau \Delta E_1$  return
27:     $\Delta w = \Delta w + \alpha y$ ;
28:     $r = r - \alpha \check{y}$ ;
29:  end for
30:   $w = w_0 + \Delta w$ ;
31:  Form final prolongation  $P = [W; I]$  with global weight vector  $w$ 
32: end procedure

```

The fact that $Z^T D_K Z$ is a diagonal matrix follows from the above observation that $K(\mathcal{L}_i, \mathcal{L}_i) = D_K(\mathcal{L}_i, \mathcal{L}_i) = A(i, i)I_m$, with I_m the identity matrix having size m equal to the cardinality of \mathcal{L}_i . \square

COROLLARY 3.2. *The product $Z^T D_K Q$, where D_K is defined as in Theorem 3.1, is equal to the null matrix:*

$$(3.42) \quad Z^T D_K Q = 0.$$

Proof. Due to the block-diagonal structure of Q and Z , all the off-diagonal blocks of $Z^T D_K Q$ are empty. Then, (3.42) follows from $D_K(\mathcal{L}_i, \mathcal{L}_i) = A(i, i)I_m$ and $Z(:, \mathcal{J}_i) \perp Q(:, \mathcal{J}_i)$. \square

3.4.1. Preconditioned CG-based energy minimization. Preconditioning CG can greatly improve convergence, but special care should be taken to maintain the search direction y in the space of vectors satisfying the near kernel constraint. In

other words, y must satisfy $\Pi_B y \equiv y$. In [28], a Jacobi preconditioner is adopted that satisfies this requirement, but, due to the special properties of the matrix K , it is possible to compute a more effective preconditioner. Denoting by M^{-1} any approximation of K^{-1} , we use $\Pi_B M^{-1} \Pi_B$ to precondition $\Pi_B K \Pi_B$ in order to guarantee the constraint. The resulting PCG algorithm is outlined in Algorithm 3.3, where, since Π_B is a projection, we can avoid premultiplying r by Π_B (line 14) as r already satisfies the constraint.

In the remainder of this section we focus our attention on $Z^T K Z$ instead of $\Pi_B K \Pi_B$, because, as $\Pi_B = I - Q Q^T = Z Z^T$, they have the same spectrum. Our aim is to find a good preconditioner for $Z^T K Z$. Unfortunately, although K is block diagonal and several effective preconditioners can be easily built for it, $Z^T K Z$ is less manageable and further approximations are needed.

By pre- and postmultiplying K by $[Z Q]^T$ and $[Z Q]$, respectively, we can write the following 2×2 block expression:

$$(3.43) \quad [Z Q]^T K [Z Q] = \begin{bmatrix} Z^T K Z & Z^T K Q \\ Q^T K Z & Q^T K Q \end{bmatrix},$$

and, since we are interested in the inverse of $Z^T K Z$, we can express it as the Schur complement of the leading block of the inverse of $[Z Q]^T K [Z Q]$ [34, Chapter 3]:

$$(3.44) \quad ([Z Q]^T K [Z Q])^{-1} = [Z Q]^T K^{-1} [Z Q] = \begin{bmatrix} Z^T K^{-1} Z & Z^T K^{-1} Q \\ Q^T K^{-1} Z & Q^T K^{-1} Q \end{bmatrix},$$

from which it follows that

$$(3.45) \quad (Z^T K Z)^{-1} = Z^T K^{-1} Z - Z^T K^{-1} Q (Q^T K^{-1} Q)^{-1} Q^T K^{-1} Z.$$

When K^{-1} is approximated with the inverse of the diagonal of K , $M_J = \text{diag}(K)$, because of Theorem 3.2, we have that $Z^T M_J^{-1} Q = 0$, and the expression (3.45) becomes

$$(3.46) \quad (Z^T K Z)^{-1} \simeq M_1^{-1} = Z^T M_J^{-1} Z,$$

which corresponds to a Jacobi preconditioning of $Z^T K Z$.

We highlight that only for Jacobi can the postmultiplication by Π_B be neglected in line 14 of Algorithm 3.3, since M_J^{-1} does not introduce components along $\text{range}(Q)$. This is consistent with [28], where the Jacobi preconditioner is used, and no postmultiplication by Π_B is adopted.

If a more accurate preconditioner is needed, K^{-1} can be approximated using a blockwise symmetric Gauss–Seidel (SGS) iteration, that is,

$$(3.47) \quad K^{-1} \simeq M_{SGS}^{-1} = (L + D)^{-T} D (L + D)^{-1},$$

which substituted into equation (3.45) reads

$$(3.48) \quad (Z^T K Z)^{-1} \simeq M_2^{-1} = Z^T M_{SGS}^{-1} Z - Z^T M_{SGS}^{-1} Q (Q^T M_{SGS}^{-1} Q)^{-1} Q^T M_{SGS}^{-1} Z.$$

Since the application of (3.48) is still impractical due to the presence of the term $(Q^T M_{SGS}^{-1} Q)^{-1}$, we neglect the second member of the right-hand side, based on the heuristic that $Z^T M_{SGS}^{-1} Q$ should be small, because $Z^T M_J^{-1} Q = 0$. After this simplification, we obtain the final expression of the projected SGS preconditioner:

$$(3.49) \quad K^{-1} \simeq M_2^{-1} = Z^T M_{SGS}^{-1} Z = Z^T (L + D)^{-T} D (L + D)^{-1} Z.$$

Note that while M_1^{-1} is exactly the Jacobi preconditioner of $Z^T K Z$, M_2^{-1} is only an approximation of the exact SGS preconditioner of $Z^T K Z$. However, we will show in the numerical results that it is able to significantly accelerate convergence.

4. Efficient parallel implementation. Energy minimization has historically been considered computationally expensive for AMG setup and real applications. Nevertheless, a cost effective implementation is still possible, but requires special care with the algorithm parallelization. In this work, we build our AMG preconditioner with Chronos [17, 15], which provides the standard numerical kernels usually required in a distributed memory AMG implementation such as the communication of ghost unknowns, coarse grid construction (e.g., computing the PMIS C-F partition [12]), sparse matrix-vector, and sparse matrix-matrix product, etc. For energy minimization, however, we developed three specific kernels that are not required in other AMG approaches, but are critical for an efficient parallel implementation here, i.e., the sparse product between K and p , application of the projection Π_B , and symmetric Gauss–Seidel with K . We do not list Jacobi preconditioning, because it simply consists of a row-wise scaling of P .

The first issue related to the product by K is that in practical applications K cannot be stored. In fact, if we consider a prolongation P having r nonzeros per row on average, the number of nonzeros per column will be approximately $s \simeq \frac{n_f}{n_c} r$ and the K matrix would be of size $n_{cs} \times n_{cs}$ with about $n_{cs} t \simeq n_f r t$ nonzeros to be stored, where t is the average number of nonzeros per row of A . Often, $r \simeq t$ and storing K becomes several times more expensive than storing A . For instance in practical elasticity problems, K can be up to 20 times larger than A , making it unavoidable to proceed in matrix-free mode for K . Fortunately, the special structure of K allows us to interpret the product Kp as a sparse matrix-matrix (SpGEMM) product between A and P , but with a fixed, prescribed pattern on P . This property can be easily derived from the definition of K (3.11) and the vector form of the prolongation p . One advantage of prescribing a fixed pattern for P is that the amount of data exchanged in SpGEMM is greatly reduced. First, the sparsity pattern adjacency information of P can be communicated only once before entering PCG, then for all the successive AP products only the entries of P are exchanged. Moreover, all the buffers to receive and send messages through the network, can be allocated and set-up only once at the beginning and removed at the end of the minimization process. In practice, we find that these optimizations reduce the cost of SpGEMM by about 50%.

By contrast, the construction and application of Π_B does not require any communication. In fact, we distribute the prolongation row-wise among processes, and since B is block diagonal, with each block corresponding to a row of P as shown in (3.8), we distribute B blocks to the process owning the respective row of P . Following this scheme, each block \mathbb{B} of B is factorized independently to obtain Q which is efficiently processed by the LAPACK routine DGEMV when applying Π_B .

Finally, we emphasize that there is no parallel bottleneck when parallelizing the application of symmetric Gauss–Seidel with K . This is because K is block diagonal and, at least in principle, each block can be assigned to a different process. However, as in the Kp product, the main issue is the large size of K which prevents explicit storage. As above, we rely on the equivalence between the Kp product and the AP product with a prescribed sparsity pattern. The symmetric Gauss–Seidel step is then performed matrix-free, similar to the AP product, by saving again a considerable amount of data exchange because it is not necessary to communicate adjacency information or reallocate communication buffers. As expected, the symmetric

Gauss–Seidel application exhibits a computational cost comparable to that of the AP product.

5. Numerical experiments. In this section, we investigate the improved convergence offered by energy minimization and the resulting computational performance and parallel efficiency for real-world applications. This numerical section is divided into three parts: a detailed analysis on the prolongation energy reduction and related AMG convergence for two medium-size matrices, a weak scalability test for an elasticity problem on a uniformly refined cube, and a comparison study for a set of large real world matrices representing fluid dynamics and mechanical problems.

As a reference for the proposed approach, we compare with the well-known and open source solvers GAMG [32], a smoothed aggregation based AMG code from PETSc, and BoomerAMG [20], a classical AMG based code in *hypre*. In all cases, Chronos, GAMG and BoomerAMG are used with preconditioned conjugate gradients (PCG). Both GAMG and BoomerAMG parameters are set for each problem starting from the default parameters, as suggested by the PETSc user guide [32] and the “best practices” parameters for *hypre* from [13]. For relaxation, we have chosen powerful standard options available in the three packages, that is FSAI with 10 nonzeros per row in Chronos, degree 2-Chebyshev accelerated Jacobi in GAMG, and hybrid symmetric-SOR BoomerAMG. Each time such default parameters are modified, we report the chosen values in the relevant section.

A truly fair comparison between different solver packages is indeed difficult. Thus, our goal in the comparison is to primarily highlight the competitiveness of the proposed energy-minimization approach with existing approaches. Here, aspects affecting the comparison include the different relaxation methods (note, no Chebyshev or symmetric SOR options exist in Chronos and no FSAI option exists in PETSc or *hypre*) and the different strength-of-connection methods. For the filtering of the strength of connection matrix, Chronos retains a fixed number of nonzeros per row (50 by default), while GAMG and BoomerAMG use a threshold for dropping small values (μ). Overall, we choose standard, effective options available in each package. However, in order to encourage a useful comparison, we do the following. Regarding the near kernel modes used by each multigrid solver’s interpolation routines, either the constant vector (scalar problems) or the rigid body modes (systems problems) are used. We note that *hypre* has the capability to account for multiple near kernel modes [1]. Additionally, the input near kernel modes are never later smoothed or updated, since not all packages allow for this option. Last, since all of the experiments are run in parallel, the system matrices are partitioned with ParMETIS [18] before the AMG preconditioner set-up to reduce communication overhead.

All numerical experiments have been run on the Marconi100 supercomputer located in the Italian consortium for supercomputing (CINECA). Marconi100 consists of 980 nodes based on the IBM Power9 architecture, each equipped with two 16-core IBM POWER9 AC922 @3.1 GHz processors. For each test, the number of nodes, N , is selected so that each node has approximately 640,000,000 nonzero entries, and, consequently, the number of nodes is problem dependent. Each node is always fully exploited by using 32 MPI tasks, i.e., each task (core) has an average load of 20,000,000 nonzero entries. The number of cores is denoted n_{cr} . Only during the smaller cases in the weak scalability analysis are nodes partially used (i.e., with less than 32 MPI tasks). Even though Chronos can exploit hybrid MPI-OpenMP parallelism, for the sake of comparison, it is convenient to use just one thread, i.e., pure MPI parallelism. Moreover, for such a high load per core, we do not find that fine-grained OpenMP parallelism is of much help.

The numerical results are presented in terms of total number of computational cores used, n_{cr} , the grid and operator complexities, C_{gd} and C_{op} , respectively, the number of iterations, n_{it} , and the setup, iteration, and total times, T_p , T_s , and $T_t = T_p + T_s$, respectively. For all the test cases, the right-hand side is a unit vector. The linear systems are solved with PCG and a zero initial guess, and convergence is achieved when the ℓ_2 -norm of the iterative residual drops below eight orders of magnitude with respect to the ℓ_2 -norm of the right-hand side.

5.1. Analysis of the energy minimization process. We use two matrices for studying prolongation energy reduction, `Cube` and `Pflow742` [22]. While the former is quite simple, as it is the fourth refinement level of the linear elasticity cube used in the weak scalability study, the latter arises from a three-dimensional (3D) simulation of the pressure field in a multilayered porous medium discretized by a sufficiently regular Q2-hexahedral finite elements. The main source of ill-conditioning here is the large contrast in the material properties for different layers. The dimensions of `Cube` are 1,778,112 rows and 78,499,998 entries, with 44.15 entries per row on average. The size of `Pflow742` is 742,793 rows and 37,138,461 entries for an average entry-per-row ratio of 50.00.

The overall solver performance is compared against the energy reduction for the fine-level P when using the energy minimization Algorithm 3.3 with either Jacobi and Gauss–Seidel as the preconditioner. Results are analyzed in terms of computational costs and times. The main algorithmic features we want to analyze are

- how the prolongation energy reduction affects AMG convergence; and
- the effectiveness of the preconditioner, i.e., Jacobi or Gauss–Seidel.

The energy minimization iteration count (Algorithm 3.3) is denoted by n_{it}^E . Between brackets, we also report the relative energy reduction that is used to monitor the restricted PCG convergence of Algorithm 3.3, as shown in (3.36). T_i is the time spent to improve the prolongation, with either classical prolongation smoothing with weighted-Jacobi (SMOOTHED) or the energy minimization process. Note that T_i is only part of T_p .

Table 5.1 shows the results for `Cube`. First, it can be seen that the energy minimization algorithm produces prolongation operators which lead to somewhat lower complexities than for the smoothed case. Moreover, energy minimization builds more effective prolongation operators overall, since the global iteration count (n_{it}) is lower. As the energy minimization iteration count increases, we can observe that n_{it} decreases, while the setup time (T_p) increases. For this simple problem, the optimal point in terms of total time (T_t) is reached using two iterations of Jacobi (EMIN-J). Figure 5.1a further shows that two iterations of Jacobi (EMIN-J) already reaches close to the achievable energy minimum. Figure 5.2a compares the cost in wall-clock seconds for each energy minimization iteration using different preconditioners. For this case and implementation, Jacobi is more efficient.

Similar conclusions can be drawn for the `Pflow742` case, as n_{it} monotonically decreases as the energy associated with the prolongation operator is reduced. As reported by Table 5.2, the optimal total time (T_t) is obtained with four Jacobi iterations (EMIN-J). Figure 5.1b shows how the energy of the prolongation operator decreases slower than for `Cube` and that Jacobi converges significantly slower than Gauss–Seidel. However, as reported by Figure 5.2b, the cost of Gauss–Seidel is still more than that of Jacobi, although the performance difference is smaller than for `Cube`.

Regarding algorithmic complexity, each energy minimization iteration with Gauss–Seidel should cost exactly twice an iteration with Jacobi, and from the above tests, Gauss–Seidel is able to reduce the energy more than twice as fast as Jacobi. Thus

TABLE 5.1
Analysis of energy minimization for the *Cube* problem.

Prolongation	n_{it}^E	$\frac{\Delta E_k}{\Delta E_1}$	C_{gr}	C_{op}	n_{it}	T_p [s]	T_s [s]	T_t [s]	T_i [s]
SMOOTHED	—	—	1.075	1.648	58	52.1	13.4	65.5	6.3
EMIN-J	1	10^0	1.075	1.592	54	47.6	11.2	58.8	11.3
EMIN-J	2	$2 \cdot 10^{-1}$	1.075	1.589	32	50.9	6.7	57.6	14.3
EMIN-J	4	10^{-2}	1.075	1.589	27	57.3	5.7	63.0	20.0
EMIN-GS	1	10^0	1.075	1.587	28	55.5	6.0	61.5	19.3
EMIN-GS	2	$2 \cdot 10^{-2}$	1.075	1.588	26	65.6	5.5	71.1	28.8
EMIN-GS	4	$1 \cdot 10^{-4}$	1.075	1.589	26	84.4	5.6	90.0	47.7

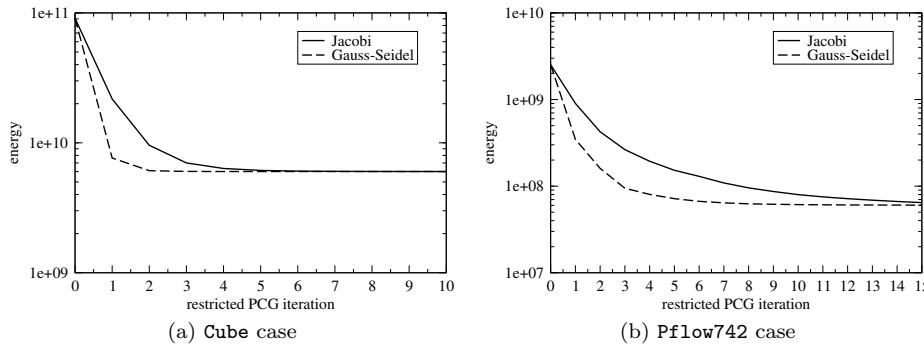


FIG. 5.1. Energy reduction versus energy minimization iterations with Jacobi and Gauss-Seidel.

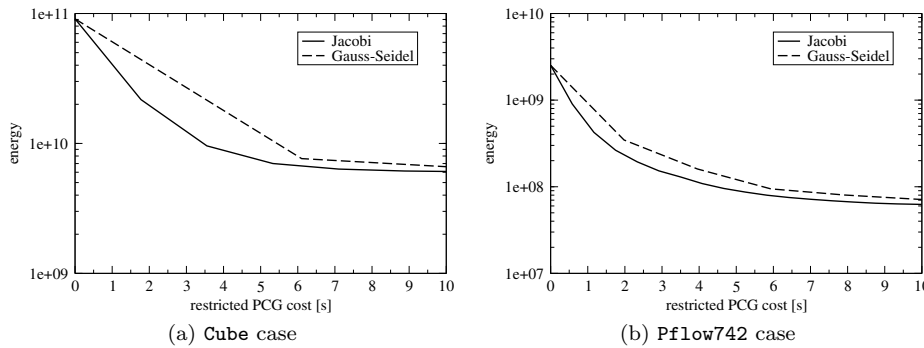


FIG. 5.2. Energy reduction versus computational cost for energy minimization preconditioned with Jacobi and Gauss-Seidel.

TABLE 5.2
Analysis of energy minimization for the *Pflow742* problem.

Prolongation	n_{it}^E	$\frac{\Delta E_k}{\Delta E_1}$	C_{gr}	C_{op}	n_{it}	T_p [s]	T_s [s]	T_t [s]	T_i [s]
SMOOTHED	—	—	1.061	1.465	369	23.0	29.8	52.8	2.3
EMIN-J	1	10^0	1.061	1.339	377	20.4	27.3	47.7	2.9
EMIN-J	2	$3 \cdot 10^{-1}$	1.061	1.344	270	21.5	19.6	41.1	3.7
EMIN-J	4	$4 \cdot 10^{-2}$	1.062	1.352	219	23.3	15.7	39.0	5.4
EMIN-J	8	$8 \cdot 10^{-3}$	1.062	1.360	189	26.2	13.8	40.0	8.1
EMIN-GS	1	10^0	1.061	1.346	276	23.0	19.9	42.9	5.3
EMIN-GS	2	$9 \cdot 10^{-2}$	1.062	1.352	221	26.2	16.0	42.2	8.2
EMIN-GS	4	$6 \cdot 10^{-3}$	1.062	1.363	184	31.7	14.3	46.0	13.8
EMIN-GS	8	$7 \cdot 10^{-4}$	1.063	1.367	183	42.9	13.8	56.7	24.8

Gauss–Seidel should be cheaper than Jacobi. Unfortunately, this is not confirmed by our numerical experiments and this is likely due to a suboptimal parallel implementation of the Gauss–Seidel preconditioner. Although the block diagonal structure of K allows, theoretically, for a perfectly parallel implementation, the Gauss–Seidel implementation requires more communication and synchronization stages than Jacobi, likely leading to our results where the Jacobi preconditioner is faster. A more cost effective implementation of Gauss–Seidel will be the focus of future work. For now, Jacobi is chosen as our default preconditioner and will be used in all subsequent cases.

Finally, we observe that a relative energy reduction of one order of magnitude gives generally the best trade-off between set-up time and AMG convergence. Therefore, we use $\tau = 0.1$ as default.

5.2. Weak scalability test. Here, we carry out a weak scalability study of energy minimization AMG for linear elasticity on a unit cube and different levels of refinement. The unit cube $[0, 1]^3$ is discretized by regular tetrahedral elements, the material is homogeneous, and all displacements are prevented on the square region $[0, 0.125] \times [0, 0.125]$ at $z = 0$. The mesh sizes are chosen such that each subsequent refinement produces about twice the number degrees of freedom with respect to the previous mesh. The problem sizes range from 222k rows to 124M rows, with an average entries-per-row ratio of 44.47.

Two sets of AMG parameters are used with Chronos: the first set targets a constant PCG iteration count, while the second minimizes the total solution time (T_t). The first section of Table 5.3 provides the outcome of the first test. The iteration count increases very slowly from 23 to 33, while the problem size increases by a factor of about 2^9 . However, this nearly optimal scaling comes with relatively large complexities. As a consequence, total times are also relatively large, especially the setup time (T_p). The time for energy minimization (T_i) scales quite well, however, with only a factor of 2 difference between the first and last refinement levels.

Next, to reduce the setup and solution times, we increase the AMG strength threshold [29] to allow for lower complexities. The second section of Table 5.3 collects the new outputs. The relative trends are the same as in the previous tests, but the timings are lower (30% on average). The reduced complexity is thus advantageous by providing faster wall-clock times, even at the expense of more iterations n_{it} .

For comparison, the same set of problems is next solved with GAMG from PETSc. The default values for all parameters are used, as it turned out they were already the best. The only exception is the threshold for dropping edges in the aggregation graph (μ), whose value is reported alongside the results. The third section of Table 5.3 shows the output for GAMG. The complexities and run-times are higher than for Chronos, especially for the setup stage. The iteration counts are usually between two and three times larger than those required by Chronos. For this test problem, reducing complexity by setting $\mu = 0.0$ is not beneficial as the increase in the iteration count cancels the set-up time reduction. We also note that both the operator complexity and set-up time increase significantly with the refinement level, while energy minimization is able to better limit growth in these quantities. Comparing two codes is fraught with difficulty, but these results do indicate that energy minimization is an efficient approach in parallel for this problem.

5.3. Challenging real-world problems. We now examine the proposed energy minimization approach for a set of challenging real-world problems arising from discretized PDEs in both fluid dynamics and mechanics. The former class consists of problems from the discretization of the Laplace operator, such as underground fluid

TABLE 5.3

Weak scalability results for regular cube. Three sets of results are reported: (i) energy minimization-based AMG via Chronos to produce an almost constant iteration count; (ii) the same tuned to minimize the total time; and (iii) PETSc's GAMG using all default parameters with the exception of μ , which is chosen to reduce the overall solution time.

Solver	nrows	n_{cr}	N	C_{gr}	C_{op}	n_{it}	T_p [s]	T_s [s]	T_t [s]	T_i [s]
Energy minimization minimal iteration count	222k	1	1	1.077	1.540	23	28.7	2.7	31.4	0.12
	447k	2	1	1.076	1.559	24	33.0	2.9	35.9	0.12
	902k	4	1	1.076	1.580	26	36.1	3.4	39.5	0.13
	1,778k	8	1	1.075	1.596	27	39.0	3.6	42.7	0.13
	3,675k	16	1	1.075	1.610	28	43.3	4.1	47.4	0.15
	7,546k	32	1	1.075	1.620	29	53.9	5.3	59.2	0.18
	15,533k	64	2	1.075	1.630	30	61.2	5.9	67.2	0.20
	31,081k	128	4	1.075	1.670	30	74.9	6.5	81.4	0.22
	62,391k	256	8	1.075	1.642	32	72.8	6.8	79.7	0.21
	124,265k	512	16	1.075	1.646	33	88.8	7.8	96.7	0.24
Energy minimization best solution time	222k	1	1	1.047	1.252	31	19.5	3.1	22.6	0.10
	447k	2	1	1.047	1.260	31	22.6	3.3	25.9	0.11
	902k	4	1	1.047	1.269	34	24.4	3.8	28.2	0.11
	1,778k	8	1	1.047	1.274	34	27.4	4.0	31.4	0.12
	3,675k	16	1	1.047	1.279	37	30.6	4.7	35.3	0.13
	7,546k	32	1	1.047	1.283	38	37.8	6.1	43.9	0.16
	15,533k	64	2	1.046	1.286	49	44.3	8.5	52.8	0.17
	31,081k	128	4	1.046	1.289	41	45.0	7.5	52.5	0.18
	62,391k	256	8	1.046	1.291	43	48.4	8.6	57.0	0.20
	124,265k	512	16	1.046	1.292	51	52.5	10.8	63.3	0.21
GAMG ($\mu = 0.01$) best solution time	222k	1	1	N/A	1.479	58	11.19	13.72	24.92	0.24
	447k	2	1	N/A	1.503	69	10.02	17.12	27.14	0.25
	902k	4	1	N/A	1.526	73	11.41	19.19	30.59	0.26
	1,778k	8	1	N/A	1.549	78	12.63	21.05	33.68	0.27
	3,675k	16	1	N/A	1.571	82	14.83	24.86	39.68	0.30
	7,546k	32	1	N/A	1.608	86	23.04	35.52	58.56	0.41
	15,533k	64	2	N/A	1.618	93	27.15	39.33	66.48	0.42
	31,081k	128	4	N/A	1.703	97	37.48	41.58	79.06	0.43
	62,391k	256	8	N/A	1.803	98	58.98	43.37	102.35	0.44
	124,265k	512	16	N/A	1.953	100	119.16	50.29	169.45	0.50

flow, compressible or incompressible airflow around complex geometries, or porous flow. The latter class consists of mechanical applications such as subsidence analysis, hydrocarbon recovery, gas storage (geomechanics), mesoscale simulation of composite materials (mesoscale), mechanical deformation of human tissues or organs subjected to medical interventions (biomedicine), and design and analysis of mechanical elements, e.g., cutters, gears, air-coolers (mechanical). The selected problems are not only large but also characterized by severe ill-conditioning due to mesh distortions, material heterogeneity, and anisotropy. They are listed in Table 5.4 with details about the size, the number of nonzeros, and the application field from which they arise.

Table 5.5 reports the AMG performance on these benchmarks when using the energy minimization procedure, classical prolongation smoothing (one step of weighed-Jacobi on the tentative prolongation), GAMG, and BoomerAMG, respectively. The overall best time is highlighted in boldface. As before, please note that for GAMG and BoomerAMG only the threshold value for dropping edges in the SoC graph (μ) has been tuned. All the other parameters are used with their default value, as they were already the best.

TABLE 5.4
Matrix sizes and number of nonzeros for the real-world problems.

Matrix	nrows	nterms	avg nt/row	Application
guenda11m	11,452,398	512,484,300	44.75	geomechanics
agg14m	14,106,408	633,142,730	44.88	mesoscale
tripod20m	19,798,056	871,317,864	44.01	mechanical
M20	20,056,050	1,634,926,088	81.52	mechanical
wing	33,654,357	2,758,580,899	81.97	mechanical
Pflow73m	73,623,733	2,201,828,891	29.91	reservoir
c4zz134m	134,395,551	10,806,265,323	80.41	biomedicine

TABLE 5.5

Results for the real-world cases using: (i) energy minimization AMG; (ii) classical prolongation smoothing; (iii) PETSc's GAMG, and (iv) hypre's BoomerAMG. Default PETSc GAMG parameters are used. Also BoomerAMG is used with default parameters, listed as "best practices" in [13]. Only μ , i.e., the threshold for dropping edges in aggregation graph, is changed. '*' means that the matrix has not been partitioned with ParMETIS. Please note that only the threshold value for dropping edges in the aggregation graph (μ) has been tuned. All the other parameters are used with their default value, as it turned out they were already the best.

Solver	case	N	μ	C_{gr}	C_{op}	n_{it}	T_p [s]	T_s [s]	T_t [s]
Energy minimization	guenda11m	1	N/A	1.041	1.325	987	314.0	399.0	713.0
	agg14m	1	N/A	1.042	1.322	23	66.7	7.2	74.0
	tripod20m	2	N/A	1.049	1.302	104	40.3	22.2	62.6
	M20	2	N/A	1.055	1.304	111	98.0	40.2	138.0
	wing	8	N/A	1.055	1.297	140	47.2	25.3	72.5
	Pflow73m	4	N/A	1.028	1.101	1169	225.0	424.0	649.0
	c4zz134m	8	N/A	1.029	1.122	154	72.7	48.8	122.0
	smoothed prolongation	guenda11m	1	N/A	1.041	1.378	1771	307.0	750.0
agg14m		1	N/A	1.042	1.371	48	62.5	15.5	78.1
tripod20m		2	N/A	1.048	1.336	212	34.6	47.5	82.2
M20		2	N/A	1.054	1.733	154	167.0	76.6	244.0
wing		8	N/A	1.055	1.697	301	93.5	71.6	165.1
Pflow73m		4	N/A	1.058	1.371	841	441.3	394.5	836.0
c4zz134m		8	N/A	1.028	1.199	277	79.2	98.9	178.0
GAMG		guenda11m	1	0.00	N/A	1.524	2237	22.3	1553.9
	agg14m	1	0.00	N/A	1.557	33	20.6	32.2	52.7
	tripod20m	2	0.01	N/A	1.679	48	32.2	32.2	64.4
	M20	2	0.01	N/A	1.203	60	36.7	62.4	99.1
	wing	8	0.01	N/A	1.204	250	34.0	108.4	142.4
	Pflow73m	4	—	N/A	—	—	—	—	—
	c4zz134m*	8	0.01	N/A	1.233	156	110.9	250.38	361.33
	Boomer AMG	guenda11m	1	0.90	1.772	2.342	—	—	>1800
agg14m		1	0.90	1.538	1.817	85	5.7	77.2	82.9
tripod20m		2	0.90	1.761	2.620	351	5.7	296.6	302.3
M20		2	0.90	1.552	1.589	634	4.3	667.6	671.9
wing		8	0.70	1.574	2.050	590	4.0	316.2	320.2
Pflow73m		4	0.25	1.154	1.302	—	—	—	—
c4zz134m		8	0.90	1.657	1.959	—	—	—	—

With respect to classical prolongation smoothing, the energy minimization procedure is able to reduce the complexities, in particular C_{op} , the setup time T_p , and also the iteration count (with the only exception being Pflow73m). It is the reduced complexities that allow energy minimization to achieve the lower setup time. The overall gain in total time is in the range 5–55% for all test cases.




FIG. 5.3. Comparison in terms of relative total time for the different preconditioners. Darker portions of the bars represent AMG setup, while the lighter segments represent the time spent in the solve phase. The relative baseline is time taken by the energy minimization phase to build prolongation.

Energy minimization also compares favorably with GAMG and BoomerAMG. GAMG provides a faster total time than energy minimization-based AMG on two cases out of seven, `agg14m` and `M20`, and a similar total time on `tripod20m`, while the solution stage in BoomerAMG is only competitive for the `agg14m` case. The situation is reversed on the most challenging examples, where GAMG is significantly slower and is unable to solve `Pflow73m`. We briefly note that, unexpectedly, the ParMETIS partitioning significantly harms GAMG effectiveness on `c4zz134m`. For this case, we report GAMG performance on the matrix with its native ordering and mark this test with a ‘*’. Typically, the GAMG and BoomerAMG set-up times are faster than energy minimization, but energy minimization allows for a preconditioner of higher quality, which significantly reduces the total time in the most difficult cases. Figure 5.3 collects all these results, reporting the relative total times. The setup and solve phases are denoted by different shading.

6. Conclusions. This work provides evidence of the potential of a novel energy minimization procedure in constructing the AMG prolongation operator. While the theoretical advantages of this idea are well known in the literature, the computational aspects and its practical feasibility in a massively parallel software were still under discussion.

With this contribution, we have highlighted how the energy minimization approach can be effectively implemented in a classical AMG setting, leading to robust and cost-effective prolongation operators when compared to other approaches. It is also shown how, especially in challenging problems, this technique can lead to considerably faster AMG convergence. The prolongation energy can be minimized with several schemes. We have adopted a restricted conjugate gradient, accelerated by suitable preconditioners. We presented and analyzed Jacobi and Gauss–Seidel preconditioners, restricting our attention to these two because of their applicability in matrix-free mode. The experiments show that Gauss–Seidel has the potential to be faster than Jacobi; however, its efficient parallel implementation is not straightforward and needs more attention.

Weak scalability has been assessed on an elasticity model problem by discretizing a cube with regular tetrahedra. The proposed algorithms have been implemented in a hybrid MPI-OpenMP linear solver and its performance has been compared to

other well recognized AMG implementations (GAMG and BoomerAMG) using a set of large and difficult problems arising from real-world applications.

In the future, we plan to further reduce set-up time by investigating other preconditioning techniques as those described in [2], and to extend this approach to nonsymmetric problems as well.

REFERENCES

- [1] A. H. BAKER, T. V. KOLEV, AND U. M. YANG, *Improving algebraic multigrid interpolation operators for linear elasticity problems*, Numer. Linear Algebra Appl., 17 (2010), pp. 495–517.
- [2] M. BENZI AND C. FACCIO, *Solving linear systems of the form $(A + \gamma UU^T)x = b$ by preconditioned iterative methods*, SIAM J. Sci. Comput., 2023, pp. S51–S70, (2023), <https://epubs.siam.org/doi/10.1137/22M1505529>.
- [3] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137.
- [4] A. BRANDT, J. BRANNICK, K. KAHL, AND I. LIVSHITS, *Bootstrap AMG*, SIAM J. Sci. Comput., 33 (2011), pp. 612–632, <https://doi.org/10.1137/090752973>.
- [5] A. BRANDT, J. BRANNICK, K. KAHL, AND I. LIVSHITS, *Algebraic distance for anisotropic diffusion problems: Multilevel results*, Electron. Trans. Numer. Anal., 44 (2015), pp. 472–496.
- [6] A. BRANDT, S. F. MCCORMICK, AND J. W. RUGE, *Algebraic Multigrid (AMG) for Automatic Multigrid Solution with Application to Geodetic Computations*, Technical report, Institute for Computational Studies, Colorado State University, Fort Collins, CO, 1982.
- [7] A. BRANDT, S. F. MCCORMICK, AND J. W. RUGE, *Algebraic multigrid (AMG) for sparse matrix equations*, in Sparsity and Its Applications, D. J. Evans, ed., Cambridge University Press, Cambridge, 1984, pp. 257–284.
- [8] J. BRANNICK, F. CAO, K. KAHL, R. D. FALGOUT, AND X. HU, *Optimal interpolation and compatible relaxation in classical algebraic multigrid*, SIAM J. Sci. Comput., 40 (2018), pp. A1473–A1493, <https://doi.org/10.1137/17M1123456>.
- [9] J. J. BRANNICK AND R. D. FALGOUT, *Compatible Relaxation and Coarsening in Algebraic Multigrid*, SIAM J. Sci. Comput., 32 (2010-01), pp. 1393–1416, <https://doi.org/10.1137/090772216>.
- [10] M. BREZINA, R. FALGOUT, S. MACLACHLAN, T. MANTEUFFEL, S. MCCORMICK, AND J. RUGE, *Adaptive smoothed aggregation (αSA) multigrid*, SIAM Rev., 47 (2005), pp. 317–346, <https://doi.org/10.1137/050626272>.
- [11] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A Multigrid Tutorial*, 2nd ed., SIAM, Philadelphia, 2000, <https://doi.org/10.1137/1.9780898719505>.
- [12] H. DE STERCK, U. M. YANG, AND J. J. HEYS, *Reducing complexity in parallel algebraic multigrid preconditioners*, SIAM J. Matrix Anal. Appl., 27 (2006), pp. 1019–1039, <https://doi.org/10.1137/040615729>.
- [13] R. D. FALGOUT AND J. B. SCHRODER, *Non-Galerkin coarse grids for algebraic multigrid*, SIAM J. Sci. Comput., 36 (2014), pp. C309–C334, <https://doi.org/10.1137/130931539>.
- [14] R. D. FALGOUT AND P. S. VASSILEVSKI, *On generalizing the algebraic multigrid framework*, SIAM J. Numer. Anal., 42 (2004), pp. 1669–1693, <https://doi.org/10.1137/130931539>.
- [15] M. FRIGO, G. ISOTTON, AND C. JANNA, *Chronos web page*, 2021, <https://www.m3eweb.it/chronos>.
- [16] S. A. GOREINOV, I. V. OSELEDETS, D. SAVOSTYANOV, E. E. TYRTYSHNIKOV, AND N. L. ZAMARASHKIN, *How to find a good submatrix*, World Scientific Publishing, 2010, pp. 247–256.
- [17] G. ISOTTON, M. FRIGO, N. SPIEZIA, AND C. JANNA, *Chronos: A general purpose classical AMG solver for high performance computing*, SIAM J. Sci. Comput., 43 (2021), pp. C335–C357, <https://doi.org/10.1137/21M1398586>.
- [18] KARYPIS LAB, *ParMETIS - Parallel Graph Partitioning and Fill-reducing Matrix Ordering*, 2022, <http://glaros.dtc.umn.edu/gkhome/metis/parmetis/overview>.
- [19] D. E. KNUTH, *Semi-optimal bases for linear dependencies*, Linear Multilinear Algebra, 17 (1985), pp. 1–4.
- [20] LAWRENCE LIVERMORE NATIONAL SECURITY, LLC, AND OTHER HYPRE PROJECT DEVELOPERS, *BoomerAMG*, 2023, <https://hypre.readthedocs.io/en/latest/solvers-boomeramg.html>.
- [21] S. MACLACHLAN, T. MANTEUFFEL, AND S. MCCORMICK, *Adaptive reduction-based amg*, Numer. Linear Algebra Appl., 13 (2006), pp. 599–620.
- [22] V. A. P. MAGRI, A. FRANCESCHINI, AND C. JANNA, *A novel algebraic multigrid approach based on adaptive smoothing and prolongation for ill-conditioned systems*, SIAM J. Sci. Comput., 41 (2019), pp. A190–A219, <https://doi.org/10.1137/17M1161178>.

- [23] J. MANDEL, M. BREZINA, AND P. VANĚK, *Energy optimization of algebraic multigrid bases*, Computing, 62 (1999), pp. 205–228.
- [24] T. A. MANTEUFFEL, S. MÜNZENMAIER, J. RUGE, AND B. SOUTHWORTH, *Nonsymmetric reduction-based algebraic multigrid*, SIAM J. Sci. Comput., 41 (2019), pp. S242–S268, <https://doi.org/10.1137/18M1193761>.
- [25] T. A. MANTEUFFEL, L. N. OLSON, J. B. SCHRODER, AND B. S. SOUTHWORTH, *A root-node-based algebraic multigrid method*, SIAM J. Sci. Comput., 39 (2017), pp. S723–S756, <https://doi.org/10.1137/16M1082706>.
- [26] T. A. MANTEUFFEL, J. RUGE, AND B. S. SOUTHWORTH, *Nonsymmetric algebraic multigrid based on local approximate ideal restriction (ℓ AIR)*, SIAM J. Sci. Comput., 40 (2018), pp. A4105–A4130, <https://doi.org/10.1137/17M1144350>.
- [27] L. N. OLSON, J. SCHRODER, AND R. S. TUMINARO, *A new perspective on strength measures in algebraic multigrid*, Numer. Linear Algebra Appl., 17 (2010), pp. 713–733.
- [28] L. N. OLSON, J. B. SCHRODER, AND R. S. TUMINARO, *A general interpolation strategy for algebraic multigrid using energy minimization*, SIAM J. Sci. Comput., 33 (2011), pp. 966–991, <https://doi.org/10.1137/100803031>.
- [29] J. W. RUGE AND K. STÜBEN, *Algebraic multigrid (AMG)*, in Multigrid Methods, Frontiers Appl. Math., S. F. McCormick, ed., SIAM, Philadelphia, 1987, pp. 73–130, <https://doi.org/10.1137/1.9781611971057.ch4>.
- [30] M. SALA AND R. S. TUMINARO, *A new Petrov–Galerkin smoothed aggregation preconditioner for nonsymmetric linear systems*, SIAM J. Sci. Comput., 31 (2008), pp. 143–166, <https://doi.org/10.1137/060659545>.
- [31] U. TROTTEBERG, C. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Academic Press, London, 2001.
- [32] UCHICAGO ARGONNE, LLC, AND THE PETSC DEVELOPMENT TEAM, *PCGAMG*, 2023, <https://petsc.org/main/docs/manualpages/PC/PCGAMG/index.html>.
- [33] P. VANĚK, J. MANDEL, AND M. BREZINA, *Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems*, Computing, 56 (1996), pp. 179–196.
- [34] P. S. VASSILEVSKI, *Multilevel Block Factorization Preconditioners: Matrix-based Analysis and Algorithms for Solving Finite Element Equations*, Springer, New York, 2008.
- [35] W. L. WAN, T. F. CHAN, AND B. SMITH, *An energy-minimizing interpolation for robust multigrid methods*, SIAM J. Sci. Comput., 21 (2000), pp. 1632–1649, <https://doi.org/10.1137/S1064827598334277>.
- [36] J. XU AND L. ZIKATANOV, *Algebraic multigrid methods*, Acta Numer., 26 (2017), pp. 591–721.
- [37] X. XU AND C.-S. ZHANG, *On the ideal interpolation operator in algebraic multigrid methods*, SIAM J. Numer. Anal., 56 (2018), pp. 1693–1710, <https://doi.org/10.1137/17M1162779>.