

1 A NEW SEMI-STRUCTURED ALGEBRAIC MULTIGRID METHOD *

2 VICTOR A. P. MAGRI[†], ROBERT D. FALGOUT , AND ULRIKE M. YANG

3 **Abstract.** Multigrid methods are well suited to large massively parallel computer architectures
4 because they are mathematically optimal and display good parallelization properties. Since current
5 architecture trends are favoring regular compute patterns to achieve high performance, the ability
6 to express structure has become much more important. The *hypre* software library provides high-
7 performance multigrid preconditioners and solvers through conceptual interfaces, including a semi-
8 structured interface that describes matrices primarily in terms of stencils and logically structured
9 grids. This paper presents a new semi-structured algebraic multigrid (SSAMG) method built on this
10 interface. The numerical convergence and performance of a CPU implementation of this method are
11 evaluated for a set of semi-structured problems. SSAMG achieves significantly better setup times
12 than *hypre*'s unstructured AMG solvers and comparable convergence. In addition, the new method
13 is capable of solving more complex problems than *hypre*'s structured solvers.

14 **Key words.** algebraic multigrid, semi-structured multigrid, semi-structured grids, structured
15 adaptive mesh refinement

16 **AMS subject classifications.** 65F08, 65F10, 65N55

17 **1. Introduction.** The solution of partial differential equations (PDEs) often
18 involves solving linear systems of equations

$$19 \quad (1.1) \quad \mathbf{Ax} = \mathbf{b},$$

20 where $A \in \mathbb{R}^{N \times N}$ is a sparse matrix; $\mathbf{b} \in \mathbb{R}^N$ is the right-hand side vector, and
21 $\mathbf{x} \in \mathbb{R}^N$ is the solution vector. In modern simulations of physical problems, the
22 number of unknowns N can be huge, e.g., on the order of a few billion. Thus, fast
23 solution methods must be used for Equation (1.1).

24 Multigrid methods acting as preconditioners to Krylov-based iterative solvers are
25 among the most common choices for fast linear solvers. In these methods, a multilevel
26 hierarchy of decreasingly smaller linear problems is used to target the reduction of
27 error components with distinct frequencies and solve (1.1) with $O(N)$ computations
28 in a scalable fashion. There are two basic types of multigrid methods [7]. Geometric
29 multigrid employs rediscrretization on coarse grids, which needs to be defined explicitly
30 by the user. A less invasive and less problem-dependent approach is algebraic multi-
31 grid (AMG) [27], which uses information coming from the assembled fine level matrix
32 A to compute a multilevel hierarchy. The *hypre* software library [21, 15] provides
33 high-performance preconditioners and solvers for the solution of large sparse linear
34 systems on massively parallel computers with a focus on AMG methods. It features
35 three different interfaces, a structured, a semi-structured, and a linear-algebraic inter-
36 face. Its most used AMG method, BoomerAMG [19], is a fully unstructured method,
37 built on compressed sparse row matrices (CSR). The lack of structure presents seri-
38 ous challenges to achieve high performance on GPU architectures. The most efficient
39 solver in *hypre* is PFMG [2], which is available through the structured interface. It
40 is well suited for implementation on accelerators, since its data structure is built on
41 grids and stencils, and achieves significantly better performance than BoomerAMG
42 when solving the same problems [4, 14]; however, it is applicable to only a subset of
43 the problems that BoomerAMG can solve. This work presents a new semi-structured

*This work was performed under the auspices of the U.S. Department of Energy by Lawrence
Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-JRNL-834288.

[†]Lawrence Livermore National Laboratory ({paludettomag1, falgout2, yang11}@llnl.gov).

44 algebraic multigrid (SSAMG) preconditioner, built on the semi-structured interface,
 45 consisting of mostly structured parts and a small unstructured component. It has
 46 the potential to achieve similar performance as PFMG with the ability to solve more
 47 complex problems.

48 There have been other efforts to develop semi-structured multigrid methods. For
 49 example, multigrid solvers for hierarchical hybrid grids (HHG) have shown to be highly
 50 efficient [6, 5, 17, 18, 22]. These grids are created by regularly refining an initial,
 51 potentially unstructured grid. Geometric multigrid methods for semi-structured tri-
 52 angular grids that use a similar approach have also been proposed [25]. More recently,
 53 the HHG approach has been generalized to a semi-structured multigrid method [24].
 54 Regarding applications, there are many examples employing semi-structured meshes
 55 which can benefit from new semi-structured algorithms, e.g., petroleum reservoir sim-
 56 ulation [16], marine ice sheets modeling [9], next-generation weather and climate
 57 models [1], and solid mechanics simulators [26], to name a few. In addition, software
 58 frameworks that support the development of block-structured AMR applications such
 59 as AMReX [29, 30] and SAMRAI [20] can benefit from the development of solvers for
 60 semi-structured problems.

61 This paper is organized as follows. Section 2 reviews the semi-structured concep-
 62 tual interface of *hypre*, which enables the description of matrices and vectors that
 63 incorporate information about the problem’s structure. Section 3 describes the new
 64 semi-structured algorithm in detail. In section 4, we evaluate SSAMG’s performance
 65 and robustness for a set of test cases featuring distinct characteristics and make com-
 66 parisons to other solver options available in *hypre*. Finally, in section 5, we list
 67 conclusions and future work.

68 **2. Semi-structured interface in *hypre*.** The *hypre* library provides three
 69 conceptual interfaces by which the user can define and solve a linear system of equa-
 70 tions: a structured (**Struct**), a semi-structured (**SStruct**) and a linear algebraic (IJ)
 71 interface. They range from highly specialized descriptions using structured grids and
 72 stencils in the case of **Struct** to the most generic case where sparse matrices are stored
 73 in a parallel compressed row storage format (**ParCSR**) [12, 13]. In this paper, we focus
 74 on the **SStruct** interface [12, 13], which combines features of the **Struct** and the IJ in-
 75 terfaces and targets applications with meshes composed of a set of structured subgrids,
 76 e.g, block-structured, overset, and structured adaptive mesh refinement grids. The
 77 **SStruct** interface also supports multi-variable PDEs with degrees of freedom lying in
 78 the center, corners, edges or faces of cells composing logically rectangular boxes. From
 79 a computational perspective, these variable types are associated with boxes that are
 80 shifted by different offset values. In this work, we consider only cell-centered problems
 81 for ease of exposition. The current CPU implementation of SSAMG cannot deal with
 82 problems involving multiple variable types yet; however, the mathematical algorithm
 83 of SSAMG expands to such general cases.

84 There are five fundamental components required to define a linear system in the
 85 **SStruct** interface: a grid, stencils, a graph, a matrix, and a vector. The grid is
 86 composed of n_p structured parts with independent index spaces and grid spacing.
 87 Each part is formed topologically by a group of boxes, which are a collection of cell-
 88 centered indices, described by their “lower” and “upper” corners. Figure 1 shows an
 89 example of a problem geometry that can be represented by this interface. Stencils
 90 are used to define connections between neighboring grid cells of the same part, e.g., a
 91 typical five-point stencil would connect a generic grid cell to itself and its immediate
 92 neighbors to the west, east, south, and north. The graph describes how individual

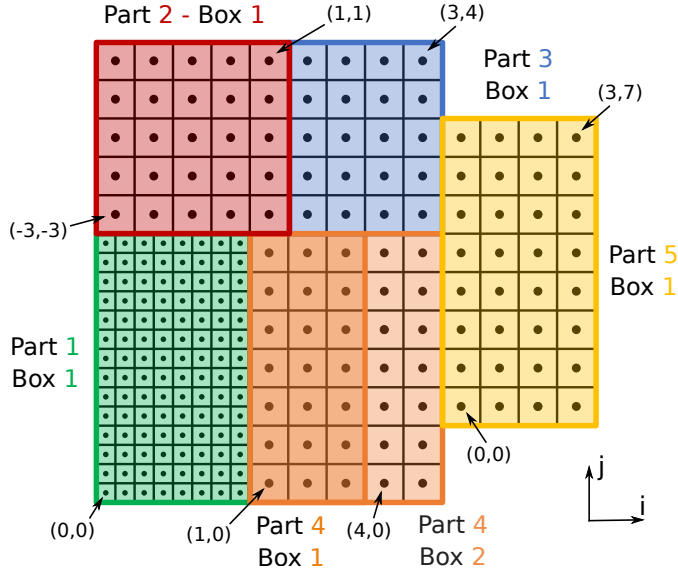


FIGURE 1. A semi-structured grid composed of five parts. Part 4 (orange) consists of two boxes, while the others consist of just a single box. Furthermore, Part 1 (green) has a refinement factor of two with respect to the other parts. The pairs (x, y) denote cell coordinates in the i and j topological directions, respectively. Note that the indices of lower-left cells for each part are independent, since the grid parts live in different index spaces.

93 parts are connected, see Figure 3 for an example. We have now the components
 94 to define a semi-structured matrix $A = S + U$, which consists of structured and
 95 unstructured components, respectively. S contains coefficients that are associated
 96 with stencil entries. These can be variable coefficients for each stencil entry in each
 97 cell within a part or can be set to just a single value if the stencil entry is constant
 98 across the part. U is stored in ParCSR format and contains the connections between
 99 parts. Since this matrix is unstructured and allows for any kind of connection between
 100 two different nodes, we do not restrict the way that two semi-structured parts interact,
 101 such as in the case of some octree-type implementations that require a 2 : 1 balance.
 102 Finally, a semi-structured vector describes an array of values associated with the cells
 103 of a semi-structured grid.

104 **3. Semi-structured algebraic multigrid (SSAMG).** In the *hypre* package,
 105 there is currently a single native preconditioner for solving problems with multiple
 106 parts through the `SStruct` interface, which is a block Jacobi method named `Split`. It
 107 uses one V-cycle of a structured multigrid solver as an approximation to the inverse
 108 of the structured part of A . This method has limited robustness since it consid-
 109 ers only structured intra-grid couplings in a part to build an approximation of A^{-1} .
 110 In this paper, we present a new solver option for the `SStruct` interface that com-
 111 puts a multigrid hierarchy taking into account inter-part couplings. This method
 112 is called SSAMG (Semi-Structured Algebraic MultiGrid). It is currently available
 113 in the `recmat` branch of *hypre*. This section defines coarsening, interpolation, and
 114 relaxation for SSAMG (subsections 3.1, 3.2, and 3.4, respectively). It also describes
 115 how coarse level operators are constructed (subsection 3.3) and discusses a strategy
 116 for improving the method's efficiency at coarse levels (subsection 3.5).

117 **3.1. Coarsening.** As in PFMG [2], we employ semi-coarsening in SSAMG. The
 118 coarsening directions are determined independently for each part of the `SStructGrid`
 119 to allow better treatment of problems with different anisotropies among the parts.
 120 The idea of semi-coarsening is to coarsen in a single direction of strong coupling such
 121 that every other perpendicular line/plane (2D/3D) forms the new coarse level. For
 122 an illustration, see Figure 3, where coarse points are depicted as solid disks.

123 In the original PFMG algorithm, the coarsening direction was chosen to be the
 124 dimension with smallest grid spacing. This option is still available in *hypre* by allowing
 125 users to provide an initial n_d -dimensional array of “representative grid spacings” that
 126 are only used for coarsening. However, both PFMG and SSAMG can also compute
 127 such an array directly from the matrix coefficients. In SSAMG, this is done separately
 128 for each part, leading to a matrix $W \in \mathbb{R}^{n_p \times n_d}$, where n_p and n_d denote the number
 129 of parts and problem dimensions. Here, element W_{pd} is heuristically thought of as
 130 a grid spacing for dimension d of part p , and hence a small value indicates strong
 131 coupling.

132 To describe the computation of W in part p , consider the two-dimensional nine-
 133 point stencil in Figure 2c and assume that $A_C > 0$ (simple sign adjustments can be
 134 made if $A_C < 0$). The algorithm extends naturally to three dimensions. Note also that
 135 both PFMG and SSAMG are currently restricted to stencils that are contained within
 136 this nine-point stencil (27-point in 3D). The algorithm proceeds by first reducing
 137 the nine-point matrix to a single five-point stencil through an averaging process,
 138 then computing the (negative) sum of the resulting off-diagonal coefficients in each
 139 dimension. That is, for the i -direction ($d = 1$), we compute

$$140 \quad (3.1) \quad c_1 = \sum_{(i,j)} -(A_{SW} + A_W + A_{NW}) - (A_{SE} + A_E + A_{NE}),$$

141 where the stencil coefficients are understood to vary at each point (i, j) in the grid.
 142 Here the left and right parenthetical sums contribute to the “west” and “east” co-
 143 efficients of the five-point stencil. The computation is analogous for the j -direction.
 144 From this, we define

$$145 \quad (3.2) \quad W_{pd} = \sqrt{\frac{\max_{1 \leq i \leq n_d} c_i}{c_d}},$$

146 based on the heuristic that the five-point stencil coefficients are inversely proportional
 147 to the square of the grid spacing.

148 With W in hand, the semi-coarsening directions for each level and part are com-
 149 puted as described in Algorithm 3.1. The algorithm starts by computing a bounding
 150 box¹ around the grid in each part, then loops through the grid levels from finest
 151 (level 1) to coarsest (level n_l). For a given grid level l and part p , the coarsening
 152 direction d^* is set to be the one with minimum² value in W_p (line 8). Then, the
 153 bounding box for part p is coarsened by a factor of two in direction d^* (line 9) and
 154 W_{p,d^*} is updated to reflect the coarser “grid spacing” on the next grid level (line 10).
 155 If the bounding box is too small, no coarsening is done (line 7) and that part becomes
 156 inactive. The coarsest grid level n_l is the first level with total semi-structured grid

¹Given a set of boxes, a bounding box is defined by the cells with minimum index (lower corner) and maximum index (upper corner) over the entire set.

²In the case of two or more directions sharing the same value of W_{pd} , as in an isotropic scenario, we set d^* to the one with smallest index.

157 size less than a given maximum size s_{max} , unless this exceeds the specified maximum
 158 number of levels l_{max} .

Algorithm 3.1 SSAMG coarsening

```

1: procedure SSAMGCOARSEN( $W$ )
2:   for  $p = 1, n_p$  do
3:     Compute part bounding boxes  $bbox_p$ 
4:   end for
5:   for  $l = 1, n_l$  do
6:     for  $p = 1, n_p$  do
7:       if  $\text{volume}\{bbox_p\} > 1$  then
8:          $d^* = \arg \min_d \{W_{pd}\}$ 
9:         Coarsen  $bbox_p$  in direction  $d^*$  by a factor of 2
10:         $W_{pd^*} = 2 * W_{pd^*}$ 
11:       end if
12:     end for
13:   end for
14: end procedure

```

159 **3.2. Interpolation.** A key ingredient in multigrid methods is the interpolation
 160 (or prolongation) operator P , the matrix that transfers information from a coarse
 161 level in the grid hierarchy to the next finer grid. The restriction operator R moves
 162 information from a given level to the next coarser grid. For a numerically scalable
 163 method, error modes that are not efficiently reduced by relaxation should be captured
 164 in the range of P , so they can be reduced on coarser levels [7].

165 In SSAMG, we employ a structured operator-based method for constructing pro-
 166 longation similar to the method used in [2]. It is “structured” because P is composed
 167 of only a structured component; interpolation is only done within a part, not between
 168 them. It is “operator-based” because the coefficients are algebraically computed from
 169 S and are able to capture heterogeneity and anisotropy. In *hypre*, P is a rectangular
 170 matrix defined by two grids (domain and range), a stencil, and corresponding stencil
 171 coefficients. In the case of P , the domain grid is the coarse grid and the range grid
 172 is the fine grid. Since SSAMG uses semi-coarsening, the stencil for interpolation con-
 173 sists of three coefficients that are computed by collapsing the stencil of A , a common
 174 procedure for defining interpolation in algebraic multigrid methods.

175 To exemplify how P is computed, consider the solution of the Poisson equation on
 176 a cell-centered grid (Figure 2a) formed by a single part and box. Dirichlet boundary
 177 conditions are used and discretization is performed via the finite difference method
 178 with a nine-point stencil (Figure 2c). Assume that coarsening is in the i -direction
 179 by selecting fine grid cells with even i -coordinate index (depicted in darker red) and
 180 renumbering them on the coarse grid as shown in Figure 2b. The prolongation oper-
 181 ator connects fine grid cells to their neighboring coarse grid cells with the following
 182 stencil (see [11] for more discussion of stencil notation)

$$183 \quad P \sim [P_W \quad 1 \quad P_E]_c = [P_W \quad * \quad P_E]_c^{r_1} \oplus [* \quad 1 \quad *]_c^{r_2},$$

184 where

$$185 \quad (3.3) \quad P_W = \frac{A_{SW} + A_W + A_{NW}}{A_S + A_C + A_N}, \text{ and } P_E = \frac{A_{SE} + A_E + A_{NE}}{A_S + A_C + A_N}.$$

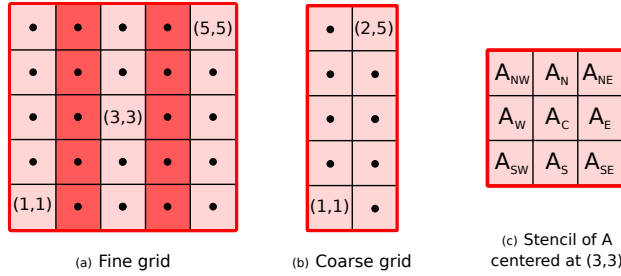


FIGURE 2. (a) and (b) show one example of fine and coarse grids, respectively, also known as range and domain grids for the purpose of prolongation. Coarsening is done in the i -direction, as depicted by the darker cells in the fine grid. (c) shows the stencil coefficients of A relative to the grid point $(3,3)$ from the fine grid. Stencil coefficients for a given grid point can be viewed as the nonzero coefficients of its respective row in a sparse matrix.

186 Here, r_1 denotes the range subgrid given by the light-colored cells in Figure 2a, and r_2
 187 denotes the subgrid given by the dark-colored cells. For a fine-grid cell such as $(3,3)$
 188 in Figure 2, interpolation applies the weights P_W and P_E to the coarse-grid unknowns
 189 associated with cells $(2,3)$ and $(4,3)$ in the fine-grid indexing, or $(1,3)$ and $(2,3)$
 190 in the coarse-grid indexing. For a fine-grid cell such as $(2,3)$, interpolation applies weight 1
 191 to the corresponding coarse-grid unknown.

192 When one of the stencil entries crosses a part boundary that is not a physical
 193 boundary, we set the coefficient associated with it to zero and update the coefficient
 194 for the opposite stencil entry so that the vector of ones is contained in the range
 195 of the prolongation operator. Although this gives a lower order interpolation along
 196 part boundaries, it limits stencil growth and makes the computation of coarse level
 197 matrices cheaper, see section 3.3. It also assures that the near kernel of A is well
 198 interpolated between subsequent levels.

199 Another component needed in a multigrid method is the restriction operator,
 200 which maps information from fine to coarse levels. SSAMG follows the Galerkin
 201 approach, where restriction is defined as the transpose of prolongation ($R = P^T$).

202 **3.3. Coarse level operator.** The coarse level operator A_c in SSAMG is com-
 203 puted via the Galerkin product $P^T A P$. Since the prolongation matrix consists only
 204 of the structured component, the triple-matrix product can be rewritten as

$$205 \quad (3.4) \quad A_c = P^T S P + P^T U P,$$

206 where the first term on the right-hand side is the structured component of A_c , and
 207 the second its unstructured component. Note that the last term involves the mul-
 208 tiplication of matrices of different types, which we resolve by converting one matrix
 209 type to the other. Since it is generally not possible to represent a `ParCSR` matrix in
 210 structured format, we convert the structured matrix P to the `ParCSR` format. How-
 211 ever, we consider only the entries of P that are actually involved in the triple-matrix
 212 multiplication $P^T U P$ to decrease the computational cost of the conversion process.

213 If we examine the new stencil size for A_c , we note that the use of the two-point
 214 interpolation operator limits stencil growth. For example, in the case of a 2D five-
 215 point stencil at the finest level, the maximum stencil size on coarse levels is nine, and
 216 for a 3D seven-point stencil at the finest level, the maximum stencil size on coarse
 217 levels is 27.

218 We prove here that under certain conditions, the unstructured portion of the

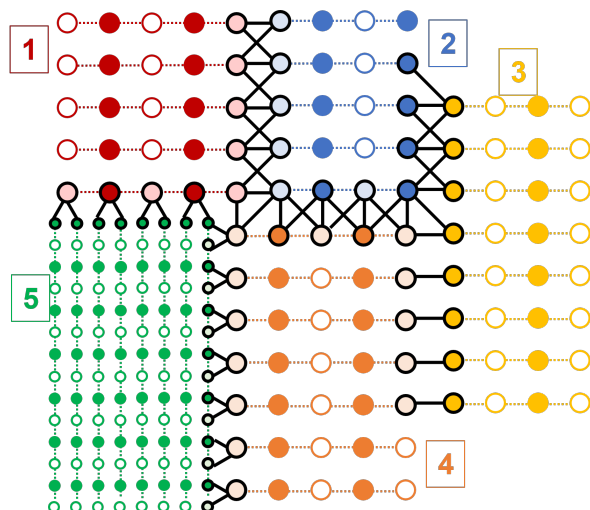


FIGURE 3. Example of a graph of the matrix U and graph of matrix P derived from the semi-structured grid shown in Figure 1. The graph of U is depicted by black-solid edges. Note that these connections are determined by the stencils of each semi-structured part. In this example, all parts have a five-point stencil, except for part 2, which has a nine-point stencil. This explains why parts 2, 3, and 4 are connected diagonally in U . The graph of P consists of five unconnected subgraphs illustrated by the dotted multicolored lines. Lastly, the boundary points are depicted by black-rimmed circles.

219 coarse grid operator stays restricted to the part boundaries and does not grow into
 220 the interior of the parts. Note that we define a part boundary $\delta\Omega_i$ here as the set of
 221 points in a part Ω_i that are connected to neighboring parts in the graph of the matrix
 222 U . See the black-rimmed points in Figure 3 for an illustration. Figure 3 shows also
 223 the graph of P for the semi-structured grid in Figure 1 and an example of a graph for
 224 the unstructured matrix U .

225 **Theorem 1** We make the following assumptions:

- 226 • The grid Ω consists of k parts: $\Omega = \Omega_1 \cup \dots \cup \Omega_k$, where $\Omega_i \cap \Omega_j = \emptyset$.
- 227 • The grid has been coarsened using semi-coarsening.
- 228 • The operator P interpolates fine points using at most two adjacent coarse
 229 points aligned with the fine points and maps coarse points onto themselves.
- 230 • The graph of the unstructured matrix U contains only connections between
 231 boundary points, i.e., $u_{i,j} = 0$ if $i \in \Omega_m \setminus \delta\Omega_m, m = 1, \dots, k$, or $j \in \Omega_n \setminus$
 232 $\delta\Omega_n, n = 1, \dots, k$, and there are no connections within a part, i.e., $u_{i,j} = 0$ for
 233 $i, j \in \Omega_m, m = 1, \dots, k$.

234 Then the graph of the unstructured part $U_c = P^T U P$ also contains only connections
 235 between boundary points, i.e., $u_{i,j}^c = 0$ if $i \in \Omega_m^c \setminus \delta\Omega_m^c, m = 1, \dots, k$, or $j \in \Omega_n^c \setminus$
 236 $\delta\Omega_n^c, n = 1, \dots, k$, and there are no connections within a part, i.e., $u_{i,j}^c = 0$ for $i, j \in$
 237 $\Omega_m^c, m = 1, \dots, k$.

238 **Proof:** Since we want to examine how boundary parts are handled, we reorder
 239 the interpolation matrix P and the unstructured part U , so that all interior points
 240 are first followed by all boundary points. The matrices P and U are then defined as
 241 follows:

$$242 \quad (3.5) \quad P = \begin{pmatrix} P^I & P^{IB} \\ P^{BI} & P^B \end{pmatrix}, \quad U = \begin{pmatrix} 0 & 0 \\ 0 & U^B \end{pmatrix}.$$

243 Note that while U^B maps $\delta\Omega$ onto $\delta\Omega$, P^B maps $\delta\Omega_c$ onto $\delta\Omega$. Thus, in the extreme
 244 case that all boundary points are fine points, P^{BI} and P^B do not exist. The coarse
 245 unstructured part is given as follows:

$$246 \quad (3.6) \quad U_c = P^T U P = \begin{pmatrix} (P^{BI})^T U^B P^{BI} & (P^{BI})^T U^B P^B \\ (P^B)^T U^B P^{BI} & (P^B)^T U^B P^B \end{pmatrix}.$$

247 It is clear already that there is no longer a connection to P^I and P^{IB} , eliminating
 248 many potential connections to interior points; however, we still need to investigate
 249 further the influence of P^{BI} and P^B .

250 Since P^{BI} , P^B , and U^B are still very complex due to their dependence on k
 251 parts, we further rewrite them as follows using the fact that P is defined only on the
 252 structured parts and U only connects boundary points of neighboring parts.

$$253 \quad (3.7) \quad P^x = \begin{pmatrix} P_1^x & & & \\ & P_2^x & & \\ & & \ddots & \\ & & & P_k^x \end{pmatrix}, \quad U^B = \begin{pmatrix} 0 & U_{1,2}^B & \dots & U_{1,k}^B \\ U_{2,1}^B & 0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & U_{k-1,k}^B \\ U_{k,1}^B & \dots & U_{k,k-1}^B & 0 \end{pmatrix}.$$

254 Note that while U_{ij}^B maps $\delta\Omega_i$ to $\delta\Omega_j$, only the coefficients corresponding to edges in
 255 the graph of U that connect points in $\delta\Omega_i$ to $\delta\Omega_j$ are nonzero, all other coefficients
 256 are zero. Then, $(P^x)^T U^B P^y$, where “ x ” and “ y ” can stand for “ BI ” as well as “ B ”,
 257 is given by

$$258 \quad (3.8) \quad \begin{pmatrix} 0 & (P_1^x)^T U_{1,2}^B P_2^y & \dots & (P_1^x)^T U_{1,k}^B P_k^y \\ (P_2^x)^T U_{2,1}^B P_1^y & 0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & (P_{k-1}^x)^T U_{k-1,k}^B P_k^y \\ (P_k^x)^T U_{k,1}^B P_1^y & \dots & (P_k^x)^T U_{k,k-1}^B P_{k-1}^y & 0 \end{pmatrix}.$$

259 This allows us to just focus on the submatrices $(P_i^x)^T U_{ij}^B P_j^y$. Let us define $P_i^x|_{\delta\Omega_{ij}}$
 260 as the matrix that consists of the rows of P_i^x that correspond to boundary points in
 261 $\delta\Omega_i$ that are connected to boundary points in $\delta\Omega_j$. Note that $\delta\Omega_{ij}$ can still be fairly
 262 complex and consist of several sides, e.g., if one part is embedded in another part.
 263 In that situation, we will divide $P_i^x|_{\delta\Omega_{ij}}$ into independent submatrices that belong to
 264 just one side of the part boundary and examine them individually. Boundary points
 265 that can be associated with several sides will be assigned to the side that connects
 266 with the other part. Since coarsening only occurs in one direction, this assignment is
 267 unambiguous. For simplicity, we will assume that $\delta\Omega_{ij}$ is just a single line for now.

268 There are only three potential scenarios that can occur due to our use of semi-
 269 coarsening and a simple two-point interpolation (Figure 3):

- 270 • all boundary points are coarse points as shown at the right boundary of part
 271 2 and the left boundary of part 3;
- 272 • all boundary points are fine points as at the right boundary of part 1 and 4;
- 273 • the boundary points are alternating coarse and fine points as illustrated at
 274 the right boundary of part 5.

275 If all points are coarse points, $P_i^B|_{\delta\Omega_{ij}} = I$ and $P_i^{BI}|_{\delta\Omega_{ij}} = 0$, since there are no
 276 connections from the boundary to the interior for $P_i^{BI}|_{\delta\Omega_{ij}}$. If all points are fine points,
 277 $P_i^B|_{\delta\Omega_{ij}}$ does not exist, and $P_i^{BI}|_{\delta\Omega_{ij}}$ is a matrix with at most one nonzero element
 278 per row in the column corresponding to the interior coarse point connected to the fine

279 boundary point, or it does not exist, if there are no interior points. Coarse points in
 280 Ω_i adjacent to the fine boundary points in $\delta\Omega_i$ become boundary points of Ω_i^c , e.g.,
 281 see right boundary of part 1 or left and right boundaries of part 4. Consequently,
 282 all nonzero elements in $P_i^{BI}|_{\delta\Omega_{ij}}$ are associated with a column belonging to $\delta\Omega_i$.
 283 In the case of alternating fine and coarse points, $P_i^{BI}|_{\delta\Omega_{ij}} = 0$, since there are no
 284 connections from the boundary to the interior, and $P_i^B|_{\delta\Omega_{ij}}$ is a matrix with at most
 285 two nonzeros in the j -th and k -th columns, where j and k are elements of $\delta\Omega_i^c$. Recall
 286 that all columns in U_{ij} belonging to points outside of $\delta\Omega_j$ and all rows belonging
 287 to points outside of $\delta\Omega_i$ are zero. Based on this and the previous observations it is
 288 clear that if all points are coarse or we are dealing with alternating fine and coarse
 289 points, the submatrices in (3.8) that involve P_i^{BI} will be 0, since $P_i^{BI}|_{\delta\Omega_{ij}} = 0$ and
 290 $P_j^{BI}|_{\delta\Omega_{ji}} = 0$. Any additional nonzero coefficients in P_i^{BI} or P_j^{BI} due to boundary
 291 points next to other parts will be canceled out in the matrix product. It is also clear,
 292 since the columns of P_i^B pertain only to points in $\delta\Omega_i^c$, that the graph of the product
 293 $(P_i^B)^T U_{ij} P_j^B$ only contains onnections of points of $\delta\Omega_i^c$ to points of $\delta\Omega_j^c$ and none to
 294 he interior or to itself.

295 Let us further investigate the case where all boundary points are fine points. We
 296 first consider $(P_i^{BI})^T U_{ij} P_j^{BI}$. Since we have already shown that $P_i^{BI}|_{\delta\Omega_{ij}} = 0$ for
 297 boundaries with coarse or alternating points leading to zero triple products in (3.8),
 298 we can ignore these scenarios and assume that for both P_i^{BI} and P_j^{BI} the boundary
 299 points adjacent to each other are fine points. Each row of $P_i^{BI}|_{\delta\Omega_{ij}}$ has at most one
 300 nonzero element in the column corresponding to the interior coarse point connected
 301 to the fine boundary point. This interior point is also an element in $\delta\Omega_i^c$. Therefore
 302 the graph of the product $(P_i^{BI})^T U_{ij} P_j^{BI}$ only contains connections of points of $\delta\Omega_i^c$ to
 303 points of $\delta\Omega_j^c$ and none to the interior or to itself. Finally, this statement also holds
 304 for the triple products $(P_i^B)^T U_{ij} P_j^{BI}$ and $(P_i^{BI})^T U_{ij} P_j^B$ using the same arguments
 305 as above. Note that the number of nonzero coefficients in U_c can still be larger than
 306 those in U , however the growth only occurs along part boundaries, and we have not
 307 observed unlimited growth in our numerical experiments.

308 **3.4. Relaxation.** Relaxation, or smoothing, is an important element of multi-
 309 grid whose task is to eliminate high frequency error components from the solution
 310 vector \mathbf{x} . The relaxation process at step $k > 0$ can be described via the generic
 311 formula:

$$312 \quad (3.9) \quad \mathbf{x}_k = \mathbf{x}_{k-1} + \omega M^{-1} (\mathbf{b} - A\mathbf{x}_{k-1}),$$

313 where M^{-1} is the smoother operator and ω is the relaxation weight. In SSAMG, we
 314 provide two pointwise relaxation schemes. The first one is weighted Jacobi, in which
 315 $M^{-1} = D^{-1}$, with D being the diagonal of A . Moreover, ω varies for each multigrid
 316 level and semi-structured part as a function of the grid-spacing metric W :

$$317 \quad (3.10) \quad \omega_p = \frac{2}{3 - \beta_p/\alpha_p},$$

318 where

$$319 \quad (3.11) \quad \alpha_p = \sum_{d=1}^{n_d} \frac{1}{W_{pd}^2} \quad \text{and} \quad \beta_p = \sum_{\substack{d=1, \\ d \neq p}}^{n_d} \frac{1}{W_{pd}^2}.$$

320 The ratio β_p/α_p adjusts the relaxation weight to more closely approximate the optimal
 321 weight for isotropic problems in different dimensions. To see how this works, consider

322 as an example a highly-anisotropic 3D problem that is nearly decoupled in the k -
 323 direction and isotropic in i and j . Because of the severe anisotropy, the problem is
 324 effectively 2D, so the optimal relaxation weight is $4/5$. Since our coarsening algorithm
 325 will only coarsen in either directions i or j , we get $\beta_p/\alpha_p = 1/2$, and $\omega_p = 4/5$ as
 326 desired.

327 The second relaxation method supported by SSAMG is L1-Jacobi. This method
 328 is similar to the previous one, in the sense that a diagonal matrix is used to construct
 329 the smoother operator; however, here, the i -th diagonal element of M equals the
 330 L1-norm of the i -th row of A :

$$331 \quad M_{ii} = \sum_{j=1}^N |A_{ij}|.$$

332 This form leads to guaranteed convergence when A is positive definite, i.e., the error
 333 propagation operator $E = I - M^{-1}A$ has a spectral radius smaller than one. We refer
 334 to [3] for more details. This option tends to give slower convergence than weighted
 335 Jacobi; however, a user-defined relaxation factor in the range $(1, 2/\lambda_{max}(M^{-1}A))$
 336 (λ_{max} is the maximum eigenvalue) can be used to improve convergence.

337 To reduce the computational cost of a multigrid cycle within SSAMG, we also
 338 provide a way to turn off relaxation on certain multigrid levels in fully and partially
 339 isotropic scenarios. We call this option “skip”, and it has the action of mimicking
 340 full-coarsening. With this option, relaxation levels (and skipped relaxation levels) are
 341 defined in sequence moving from fine to coarse as follows. Let d_l^* be the coarsening
 342 direction on level l and let r be the relaxation level with largest value $r < l$. If $d_l^* = d_p^*$
 343 for some $p \geq r$, then we define level l to be a relaxation level. We also ensure relaxation
 344 is never skipped on the finest and coarsest levels. For example, in an isotropic setting,
 345 the coarsening directions (from fine to coarse) might be 1, 2, 3, 1, 2, 3, ... with relaxation
 346 occuring on levels where $d_l^* = 1$. In an anisotropic setting with strong coupling in
 347 dimension 1, the coarsening directions might be 1, 1, 1, 1, 2, 3, ... with relaxation again
 348 occuring on levels where $d_l^* = 1$.

349 **3.5. Hybrid approach.** Since SSAMG uses semi-coarsening, the ratio between
 350 the number of variables on subsequent grids is equal to two. In classical algebraic
 351 multigrid, this value tends to be larger, especially when aggressive coarsening strate-
 352 gies are applied. This leads to the creation of more levels in the multigrid hierarchy
 353 of SSAMG when compared to BoomerAMG. Since the performance benefits of ex-
 354 ploiting structure decreases on coarser grid levels, we provide an option to transition
 355 to an unstructured multigrid hierarchy at a certain level or coarse problem size cho-
 356 sen by the user. This is done by converting the matrix type from `SStructMatrix` to
 357 `ParCSRMatrix` at the transition level. The rest of the multigrid hierarchy is set up
 358 using BoomerAMG configured with the default options used in *hypre* as of version
 359 2.25.0, i.e., HMIS coarsening, strength threshold of value of 0.25, ext+i interpolation,
 360 and forward/backward L1-Gauss-Seidel relaxation. With a properly chosen transition
 361 level, the hybrid approach can improve convergence and thus solve times while main-
 362 taining a similar overall setup cost for SSAMG. In the non-hybrid case, the coarsest
 363 level problem in SSAMG is solved with a single sweep of the same relaxation method
 364 used in previous levels.

365 **4. Numerical results.** In this section, we investigate convergence and perfor-
 366 mance of SSAMG when used as a preconditioner for the conjugate gradient method
 367 (PCG). We also compare it to three other multigrid schemes in *hypre*, namely PFMG,

368 Split, and BoomerAMG. The first is the flagship multigrid method for structured
 369 problems in *hypre* based on semi-coarsening [2, 8], the second, a inexact block-Jacobi
 370 method built on top of the `SStruct` interface [21], in which blocks are mapped to
 371 semi-structured parts, and the last scheme is *hypre*'s unstructured algebraic multi-
 372 grid method [19]. Each of these preconditioners has multiple setup parameters that
 373 affect its performance. For the comparison made here, we select those leading to the
 374 best solution times on CPU architectures. In addition, we consider four variants of
 375 SSAMG in an incremental setting to demonstrate the effects of different setup options
 376 described in the paper. A complete list of the methods considered here is given below:

- 377 • PFMG: weighted Jacobi smoother and “skip” option, as described in section
 378 3.4 for SSAMG, turned on.
- 379 • Split: inexact block-Jacobi method with one V-cycle of PFMG as the inner
 380 solver for parts.
- 381 • BoomerAMG³: Forward/Backward L1-Gauss-Seidel relaxation [3]; coarsen-
 382 ing via HMIS [10] with a strength threshold value of 0.25; modularized option
 383 for computing the Galerkin product *RAP*; one level (first) of aggressive coarsen-
 384 ing with multi-pass interpolation [28] and, in the following levels, matrix-
 385 based extended+i interpolation [23] truncated to a maximum of four nonzero
 386 coefficients per row.
- 387 • SSAMG-base: baseline configuration of SSAMG employing weighted L1-
 388 Jacobi smoother with relaxation factor equal to 3/2.
- 389 • SSAMG-skip: above configuration plus the “skip” option.
- 390 • SSAMG-hybrid: above configuration plus the “hybrid” option for transition-
 391 ing to BoomerAMG, with the aggressive coarsening option and multipass
 392 interpolation options disabled, as the coarse solver at the 10th level, which
 393 corresponds to three steps of full grid refinement in 3D, i.e., 512 times reduc-
 394 tion on the number of degrees of freedom (DOFs).
- 395 • SSAMG-opt: refers to the best SSAMG configuration and employs the same
 396 parameters as SSAMG-hybrid except for transitioning to BoomerAMG at
 397 the 7th level. This results in six pure SSAMG coarsening levels and reduction
 398 factor of 64 on the number of DOFs.

399 Every multigrid preconditioner listed above is applied to the residual vector via a
 400 single V(1,1)-cycle. The global coarsest grid size is equal to at most eight unknowns
 401 in all cases where BoomerAMG is used, one unknown for PFMG, and the number of
 402 parts for Split, SSAMG-base and SSAMG-skip. The number of parts varies according
 403 to the test case, e.g., four in test cases 1 and 2, three in test case 3, and one in the last
 404 test case. Since we test the solvers for increasing global problem sizes, the number of
 405 levels in the various multigrid hierarchies increases for larger problems.

406 We consider four test cases featuring three-dimensional semi-structured grids,
 407 different part distributions, problem sizes and anisotropy directions. In the coarsest
 408 problem size for a test case, each semi-structured part is owned by a different processor
 409 and formed by a single box containing $m \times m \times m$ cells. In the remaining problem
 410 sizes, each semi-structured part is uniformly refined in all directions by a factor equal
 411 to p and distributed to $p \times p \times p$ unique MPI tasks. This leads to a total of $n_p p^3$
 412 MPI tasks for n_p parts. Note that, in this strategy, the number of unknowns owned
 413 by a processor and the total number of parts in the grid are kept constant, while the
 414 global number of unknowns per part is $(m \times p)^3$. For an example of grid partitioning,

³We tuned the BoomerAMG configuration parameters for performance, i.e. best overall setup and solve times. Note that such a strategy generally does not lead to the fastest convergence.

415 see Figure 4. We are particularly interested in evaluating the weak scalability of the
 416 proposed method for a few tasks up to a range of thousands of MPI tasks. Thus, we
 417 vary the value of p from one to eight with unitary increments.

418 For the results, we report the number of iterations needed for convergence, setup
 419 time of the preconditioner, and solve time of the iterative solver. All experiments
 420 were performed on Lassen, a cluster at LLNL equipped with two IBM POWER9
 421 processors (totaling 44 physical cores) per node. However, we note that 32 cores per
 422 node at most were used in the numerical experiments to reduce the effect of limited
 423 memory bandwidth. Convergence of the iterative solver is achieved when the L2-
 424 norm of the residual vector is less than $10^{-6}\|\mathbf{b}\|_2$. The linear systems were formed
 425 via discretization of the Poisson equation through finite differences via the following
 426 seven-point stencil:

$$427 \quad (4.1) \quad A \sim [-\gamma] \begin{bmatrix} & & -\beta \\ -\alpha & 2(\alpha + \beta + \gamma) & -\alpha \\ & & -\beta \end{bmatrix} [-\gamma]$$

428 where α , β , and γ denote the coefficients in the i , j , and k topological directions. For
 429 the isotropic problems, $\alpha = \beta = \gamma = 1$, for the anisotropic cases we define their values
 430 in section 4.2. Finally, we used a vector of ones for the right hand side and an initial
 431 solution guess composed of random numbers between zero and one.

432 **4.1. Test case 1 - cubes side-by-side.** The first test case is made of an iso-
 433 tropic and block-structured three-dimensional domain composed of four cubes, where
 434 each contains the same number of cells and refers to a different semi-structured part.
 435 Figure 4a shows one particular case with cubes formed by four cells in each direction.
 436 Regarding the solver choices, since PFMG works only for single-part problems, we
 437 translated parts into independent boxes in an equivalent structured grid. Note that
 438 such a transformation is only possible due to the simplicity of the current problem
 439 geometry and it is unattainable in more general cases such as those described later in
 440 sections 4.3 and 4.4.

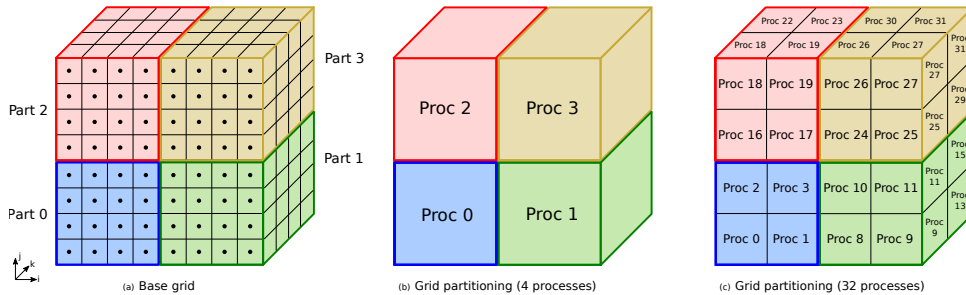


FIGURE 4. (a) Three-dimensional base grid used for test case 4.1. Note that there are no adjacent parts in the k -direction. Colors denote different parts, and the numerical experiments showed in this section are produced by uniformly refining the semi-structured parts composing the base grid in all topological directions. The next two illustrations show how portions of the semi-structured grid are mapped to different MPI tasks for $p = 1$ (b) and $p = 2$ (c). Note that part zero (blue) is entirely owned by processor zero when $p = 1$, and it is distributed among processors zero to seven when $p = 2$ (processors four to seven cannot be seen in the figure), while keeping the same number of unknowns per processor.

441 For the numerical experiments, we consider $m = 128$, which gives a local problem
 442 size per part of 2,097,152 DOFs and a global problem size of 8,388,608 DOFs for 4

443 MPI tasks ($p = 1$). The largest problem we consider here, obtained when $p = 8$, has a
 444 global size of about 4.3 billion DOFs. For a complete list of problem sizes considered
 445 for this test case, see Table 1. In addition, we show the number of levels in the
 446 various multigrid hierarchies. Note that SSAMG-base, SSAMG-skip, and Split have
 447 the same numbers, which increase with the problem sizes since these methods have
 448 a fixed coarsening ratio. The same observation is valid for PFMG, while its number
 449 of levels are shifted by two since the coarsest problem size is set to a single degree of
 450 freedom. BoomerAMG (abbreviated to AMG on the table) has the least number of
 451 levels among all methods, this is due to the use of aggressive coarsening and the fact
 452 that the coarsening ratio varies through the hierarchy. Lastly, SSAMG-hybrid and
 453 SSAMG-opt demonstrates a mixed behavior due to the transition to BoomerAMG at
 454 some point in their multigrid hierarchies.

TABLE 1

Number of levels in the multigrid hierarchy generated by each method (columns) for several problem sizes (rows). Structured multigrid methods have an increasing number of levels with larger problem sizes, while unstructured multigrid (BoomerAMG) shows a less pronounced increase.

p	N_{procs}	DOFs	Number of multigrid levels						
			base	skip	hybrid	opt	AMG	Split	PFMG
1	4	8,388,608	22	22	14	12	8	22	24
2	32	67,108,864	25	25	15	12	9	25	27
3	108	226,492,416	26	26	15	13	9	26	28
4	256	536,870,912	28	28	15	13	10	28	30
5	500	1,048,576,000	28	28	15	13	10	28	30
6	864	1,811,939,328	29	29	15	13	10	29	31
7	1,372	2,877,292,544	30	30	15	13	10	30	32
8	2,048	4,294,967,296	31	31	15	13	10	31	33

455 Figure 5 shows weak scalability results for this test case. Analyzing the iteration
 456 counts, Split is the only method that does not converge in less than the maximum
 457 iteration count of 100 for runs larger than 108 million DOFs ($p = 3$). This lack
 458 of numerical scalability was already expected since couplings among parts are not
 459 captured in Split’s multigrid hierarchy. The best iteration counts are reached by
 460 PFMG, which is natural since this method can take full advantage of the problem’s
 461 geometry. Noticeably, the iteration counts of SSAMG-opt follow PFMG closely, since
 462 part boundaries are no longer considered after transitioning to BoomerAMG on the
 463 coarser levels and the transition level takes place earlier here than in SSAMG-hybrid;
 464 the other SSAMG variants need a higher number of iterations for achieving conver-
 465 gence, since the interpolation is of lower quality along part boundaries. Lastly, the
 466 BoomerAMG preconditioner shows a modest increase in iteration counts for increasing
 467 problem sizes, and this is common in the context of algebraic multigrid.

468 Solve times are directly related to iteration counts. Since Split has a similar
 469 iteration cost to the other methods but takes the largest number of iterations to
 470 converge, it is the slowest option in solution time. For the same reason, the three
 471 SSAMG variants except for SSAMG-opt are slower than the remaining precondition-
 472 ers. Still, SSAMG-skip is faster than SSAMG-base, despite showing more iterations,
 473 because the “skip” option reduces its iteration cost. The optimal variant SSAMG-opt
 474 is able to beat BoomerAMG by a factor of 1.6x for $p = 1$ and 1.8x for $p = 8$. More-

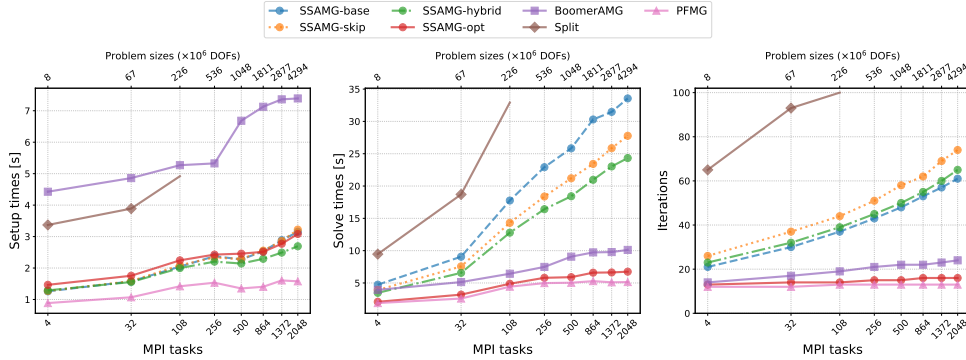


FIGURE 5. Weak scalability results for test case 1. Three metrics are shown in the figure, i.e., setup phase times in seconds (left); solve phase times in seconds (middle), and number of iterations (right). All curves are plotted with respect to the number of MPI tasks, N_{procs} , which varies from 4 ($p = 1$) up to 2048 ($p = 8$).

over, SSAMG-opt shows little performance degradation with respect to the fastest preconditioner (PFMG). Lastly, the jumps in solve times between $N_{\text{procs}} = 32$ and $N_{\text{procs}} = 108$ are mainly due to the higher costs associated with inter-node communication versus intra-node communication. The same observation is valid for the remaining test cases in this section.

BoomerAMG is the slowest option when analyzing setup times. This is a result of multiple reasons, the three most significant ones being:

- BoomerAMG employs more elaborate formulas for computing interpolation, which require more computation time than the simple two-point scheme used by PFMG and SSAMG;
- the triple-matrix product algorithm for computing coarse operators implemented for CSR matrices is less efficient than the specialized algorithm employed by `Struct` and `SStruct` matrices;
- BoomerAMG’s coarsening algorithm involves choosing fine/coarse nodes on the matrix graph besides computing a strength of connection matrix. Those steps are not necessary for PFMG or SSAMG.

This is followed by Split, which should have setup times close to PFMG, but due to a limitation of its parallel implementation, the method does not scale well with an increasing number of parts. On the other hand, all the SSAMG variants show comparable setup times, up to 2.8x faster than BoomerAMG. The first two SSAMG variants share the same setup algorithm, and their lines are superposed. SSAMG-opt has a slightly slower setup for $p \leq 5$ than SSAMG-base, but for $p > 5$ the setup times of these two methods match. The fastest SSAMG variant by a factor of 1.2x with respect to the others is SSAMG-hybrid, and that holds because it generates a multigrid hierarchy with fewer overall levels than the non-hybrid SSAMG variants leading to less communication overhead associated with collective MPI calls. The same argument is true for SSAMG-opt; however, the benefits of having fewer levels is outweighed by the higher cost of converting the `SStructMatrix` to a `ParCSRMatrix` at a finer transition level, which involves more data. Still, SSAMG-opt is 2.9x and 2.5x faster than BoomerAMG for $p = 1$ and $p = 8$, respectively. PFMG yields the best setup times with a speedup of nearly 4.8x with respect to BoomerAMG and up

506 to 1.9x with respect to SSAMG.

507 We note that PFMG is naturally a better preconditioner for this problem than
 508 SSAMG since it does not have the same restrictions as SSAMG for computing inter-
 509 polation coefficients across part boundaries. However, this test case was significant to
 510 show how close the performance of SSAMG can be to PFMG, and we demonstrated
 511 that SSAMG-opt is fairly close to PFMG, besides yielding faster solve and setup times
 512 than BoomerAMG.

513 **4.2. Test case 2 - anisotropic cubes.** This test case has the same problem
 514 geometry and sizes ($m = 128$) as the previous test case; however, it employs different
 515 stencil coefficients (α , β , and γ) for each part of the grid with the aim of evaluating
 516 how anisotropy affects solver performance. Particularly, we consider three different
 517 scenarios (Figure 6) where the coefficients relative to stencil entries belonging to the
 518 direction of strongest anisotropy for a given part are 100 times larger than the re-
 519 maining ones. The directions of prevailing anisotropy for each scenario are listed
 520 below:

- 521 (A) “ i ” (horizontal) for all semi-structured parts.
- 522 (B) “ i ” for parts zero and two; “ j ” (vertical) for parts one and three.
- 523 (C) “ i ” for part zero, “ j ” for part three, and “ k ” (depth) for parts one and two.

524 Regarding the usage of PFMG for this problem, the same transformation mentioned
 525 in section 4.1 applies here as well.

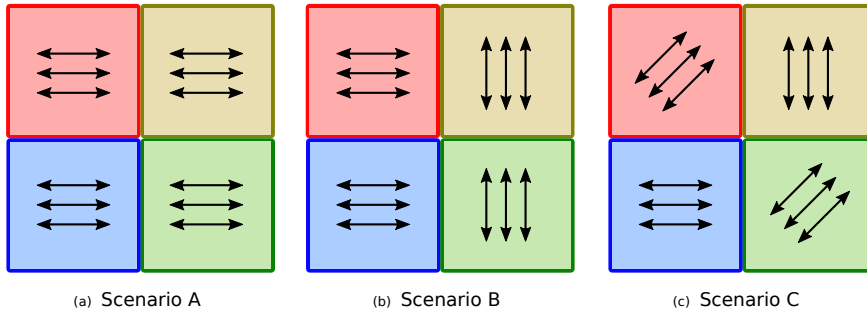


FIGURE 6. XY -plane cut of the three-dimensional grids used in test case 4.2. We consider three anisotropy scenarios. Arrows indicate the direction of prevailing anisotropy in each part of the grid, e.g., i -direction in scenario A. Diagonal arrows in the rightmost case indicate the k -direction.

526 Figure 7 shows the results referent to scenario A. Numerical scalabilities of the
 527 hybrid SSAMG variants look good, and for SSAMG-hybrid particularly, they look
 528 better than in the previous test case since the two-point interpolation strategy is
 529 naturally a good choice for the first few coarsening levels when anisotropy is present in
 530 the same direction as the coarsening one. However, the non-hybrid SSAMG variants
 531 do not show a reasonable scalability, which can be explained by their inability to
 532 interpolate accross part boundaries when the coarse level problems get isotropic. Note
 533 that BoomerAMG does not suffer from this limitation, which helps the hybrid SSAMG
 534 variants to be more scalable. For the same reason as discussed in the previous test
 535 case, Split takes more than 100 iterations to converge, thus, it is not shown in the
 536 figures for this test case. Lastly, PFMG uses the least number of iterations followed
 537 closely by SSAMG-opt and BoomerAMG.

538 Regarding solve times, SSAMG-opt is about 1.3x faster than BoomerAMG for
 539 $p \leq 2$ and $p > 5$. The “skip” option of SSAMG is not beneficial for this case since
 540 the solve times of SSAMG-skip are higher than SSAMG-base. In fact, such an option

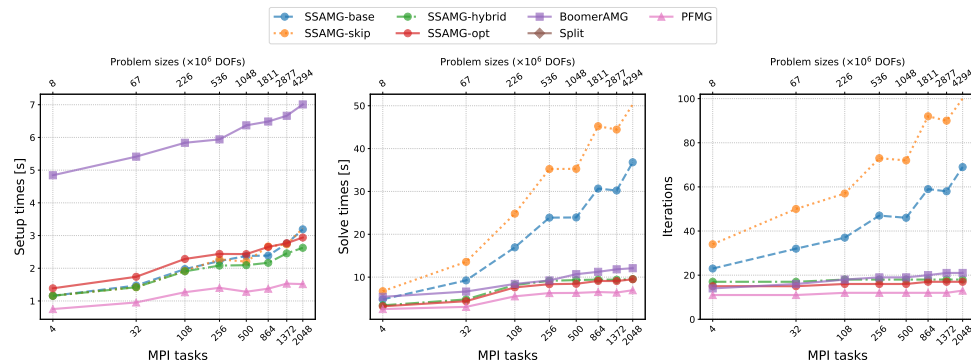


FIGURE 7. Weak scalability results for scenario A of test case 2. Three metrics are shown in the figure, i.e., setup phase times in seconds (left); solve phase times in seconds (middle), and number of iterations (right). All curves are plotted with respect to the number of MPI tasks, N_{procs} , which varies from 4 ($p = 1$) up to 2048 ($p = 8$).

541 does not play a significant role in reducing the solve time compared to isotropic test
 542 cases. This is because coarsening happens in the same direction for the first few levels
 543 in anisotropic test cases, and thus relaxation is skipped only in the later levels of the
 544 multigrid hierarchy where the cost per iteration associated with them is already low
 545 compared to the initial levels. Moreover, the omission of relaxation in coarser levels
 546 of the multigrid hierarchy can be detrimental for convergence in SSAMG, explaining
 547 why SSAMG-skip requires more iterations than SSAMG-base. Following the fact
 548 that PFMG is the method that needs fewer iterations for convergence, it is also the
 549 fastest in terms of solution times. For setup times, the four SSAMG variants show
 550 comparable results, and similar conclusions to test case 1 are valid here. Lastly, the
 551 speedups of SSAMG-opt over BoomerAMG are 3.5x and 2.4x for $p = 1$ and $p = 8$,
 552 respectively.

553 Results for scenario B are shown in Figure 8. The most significant difference
 554 here compared to scenario A are the results for PFMG. Particularly, the number of
 555 iterations for PFMG is above 100 and not shown in the plots. This is caused by the
 556 fact that PFMG employs the same coarsening direction everywhere on the grid, and
 557 thus it cannot recognize the different regions of anisotropy as done by SSAMG. This
 558 is clearly sub-optimal since a good coarsening scheme should adapt to the direction
 559 of largest coupling of the matrix coefficients. The larger number of iterations is also
 560 reflected in the solve times of PFMG, which become less favorable than those by
 561 SSAMG and BoomerAMG. Setup times of PFMG continue to be the fastest ones;
 562 however, this advantage is not sufficient to maintain its position of fastest method
 563 overall. The comments regarding the speedups of SSAMG compared to BoomerAMG
 564 made for scenario A also apply here.

565 We conclude this section by analyzing the results, given in Figure 9, for the last
 566 anisotropy scenario C. Since there is a mixed anisotropy configuration in this case as
 567 in scenario B, PFMG does not show a satisfactory convergence behavior and it is not
 568 shown in the graph. On the other hand, the SSAMG variants show good numerical
 569 and computational scalabilities, and, particularly, SSAMG-opt shows similar speedups
 570 compared to the BoomerAMG variants as discussed in the previous scenarios. When
 571 considering all three scenarios discussed in this section, we note that SSAMG shows

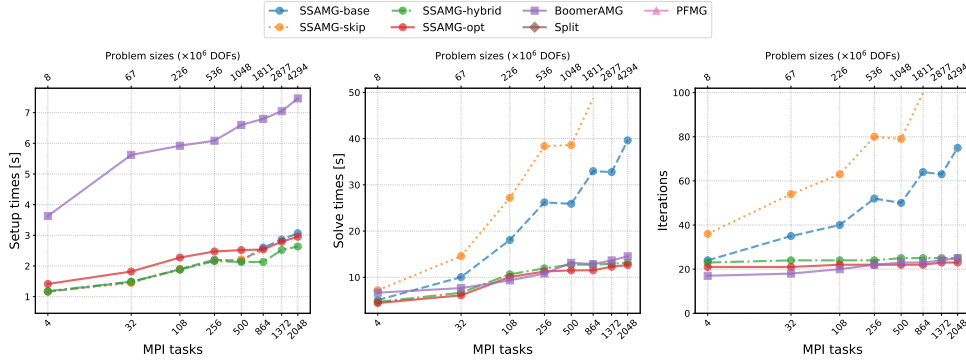


FIGURE 8. Weak scalability results for scenario B of test case 2. Three metrics are shown in the figure, i.e., setup phase times in seconds (left); solve phase times in seconds (middle), and number of iterations (right). All curves are plotted with respect to the number of MPI tasks, N_{procs} , which varies from 4 ($p = 1$) up to 2048 ($p = 8$).

572 good robustness with changes in anisotropy, and this an important advantage over
 573 PFMG.

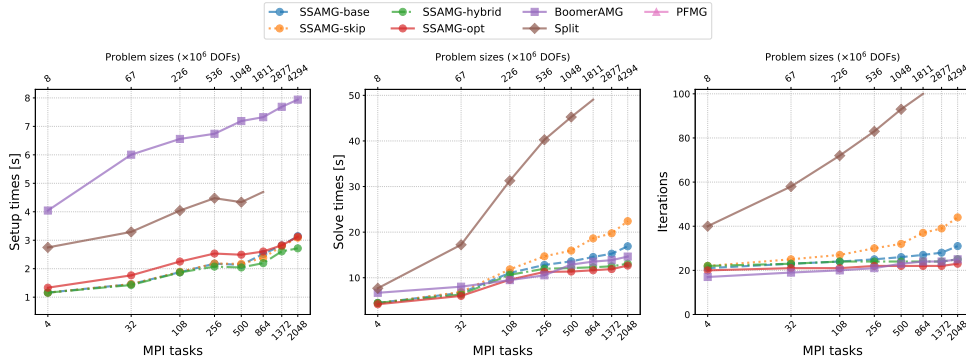


FIGURE 9. Weak scalability results for scenario C of test case 2. Three metrics are shown in the figure, i.e., setup phase times in seconds (left); solve phase times in seconds (middle), and number of iterations (right). All curves are plotted with respect to the number of MPI tasks, N_{procs} , which varies from 4 ($p = 1$) up to 2048 ($p = 8$).

574 **4.3. Test case 3 - three-points intersection.** In this test case, we consider a
 575 grid composed topologically of three semi-structured cubic parts that share a common
 576 intersection edge in the k -direction (Figure 10). Stencil coefficients are isotropic,
 577 but this test case is globally non-Cartesian. In particular, the coordinate system is
 578 different on either side of the boundary between parts 1 and 2. For example, an east
 579 stencil coefficient coupling Part 1 to Part 2 is symmetric to a north coefficient coupling
 580 Part 2 to Part 1.

581 For the numerical experiments of this section, we use $m = 160$, which gives a local
 582 problem size per part of 4,096,000 DOFs, and a global problem size of 12,288,000
 583 DOFs, when $p = 1$, i.e., three parts and MPI tasks. Figure 11 reports weak scalability

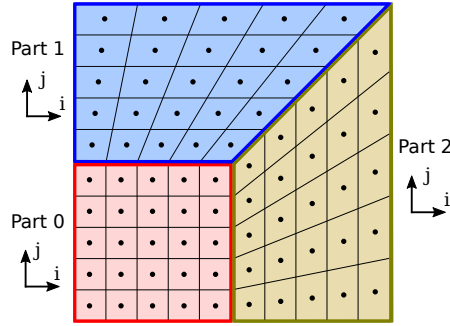


FIGURE 10. ij -plane view of the base geometry for test case 4.3. Uniformly refined instances of this problem in all directions are used for obtaining the results.

584 results for the current test case. As noted in section 4.1, it is not possible to recast
 585 this problem into a single part; thus, we cannot show results for PFMG here.

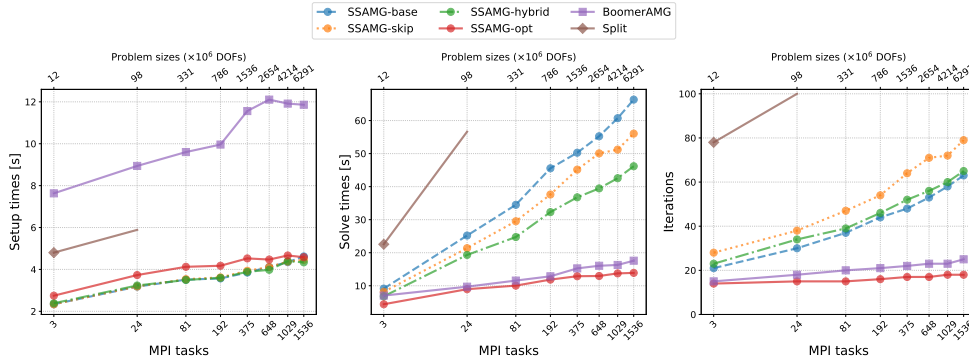


FIGURE 11. Weak scalability results for test case 3. Three metrics are shown in the figure, i.e., setup phase times in seconds (left); solve phase times in seconds (middle), and number of iterations (right). All curves are plotted with respect to the number of MPI tasks, N_{procs} , which varies from 3 ($p = 1$) up to 1536 ($p = 8$).

586 Examining the iteration counts reported in Figure 11, we see that SSAMG-opt
 587 is the fastest converging option with the number of iterations ranging from 17, for
 588 $p = 1$ (3 MPI tasks), to 19, for $p = 8$ (1536 MPI tasks). This is the best numerical
 589 scalability among the other methods, including BoomerAMG. On the other hand, the
 590 remaining SSAMG variants do not show such good scalability as in the previous test
 591 cases. Once again, this is related to how SSAMG computes interpolation weights of
 592 nodes close to part boundaries. In this context, we plan to investigate further how to
 593 improve SSAMG’s interpolation such that the non-hybrid SSAMG variants can have
 594 similar numerical scalability to SSAMG-opt. As in the previous test cases, the Split
 595 method is the least performing method and does not converge within 100 iterations
 596 for $p \geq 2$ ($N_{\text{procs}} \geq 24$).

597 Regarding solve times, SSAMG-opt is the fastest method since it needs the min-
 598 imum amount of iterations to reach convergence. Compared to BoomerAMG, its
 599 speedup is 1.3x for $p = 1$ and $p = 8$. SSAMG-skip shows solution times smaller than
 600 SSAMG-base, and, here, the “skip” option is beneficial to performance. Lastly, look-

601 ing at setup times, all SSAMG variants show very similar timings and the optimal
 602 variant is up to 3.2x faster than BoomerAMG, proving once again the benefits of
 603 exploiting problem structure.

604 **4.4. Test case 4 - structured adaptive mesh refinement (SAMR).** In the
 605 last problem, we consider a three-dimensional SAMR grid consisting of one level of
 606 grid refinement, and thus composed of two semi-structured parts (Figure 12). The
 607 first one, in red, refers to the outer coarse grid, while the second, in blue, refers to
 608 the refined patch (by a factor of two) located in the center of the grid. Each part has
 609 the same number of cells. To construct the linear system matrix for this problem,
 610 we treat coarse grid points living inside of the refined part as ghost unknowns, i.e.,
 611 the diagonal stencil entry for these points is set to one and the remaining off-diagonal
 612 stencil entries are set to zero. Inter-part couplings at fine-coarse interfaces are stored
 613 in the unstructured matrix (U), and the value for the coefficients connecting fine
 614 grid cells with its neighboring coarse grid cells (and vice-versa) is set to $2/3$. This
 615 value was determined by composing a piecewise constant interpolation formula with a
 616 finite volume discretization rule. We refer the reader to the SAMR section of *hypr*'s
 617 documentation [21] for more details.

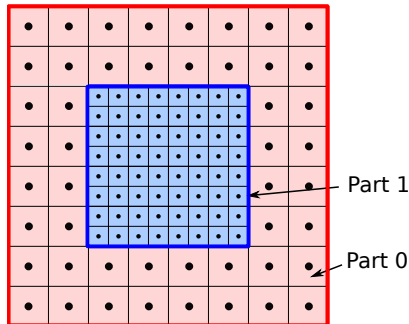


FIGURE 12. XY -plane cut of the three-dimensional semi-structured grid used in test case 4.4 when $m = 8$. The semi-structured parts represent two levels of refinement and contain the same number of cells.

618 The numerical experiments performed in this section used $m = 128$, leading to a
 619 local problem size per part of 2, 097, 152 DOFs, and a global problem size of 4, 194, 304
 620 DOFs, for $p = 1$ ($N_{\text{procs}} = 2$). Figure 13 shows weak scalability results for this test
 621 case. This problem is not suitable for PFMG, thus we do not show results for it.

622 As in the previous test cases, Split does not reach convergence within 100 iterations
 623 when $p \geq 2$. Then, SSAMG-skip is the second least convergent option followed
 624 by SSAMG-base. The best option is again SSAMG-opt with the number of iterations
 625 ranging from 18 ($p = 1$) to 29 ($p = 8$). Furthermore, its iteration counts are
 626 practically constant for the several parallel runs, except for slight jumps located at
 627 $p = 4$ ($N_{\text{procs}} = 128$) and $p = 8$ ($N_{\text{procs}} = 1024$), which are more pronounced for
 628 SSAMG-hybrid.

629 Solve times are generally better when the methods converge faster; however, that
 630 is not always true. In this test case, the iteration costs of SSAMG-base are higher
 631 than SSAMG-skip, due to more time spent on relaxation, and the faster convergence
 632 of the former method is not able to offset the cheaper cost of the latter, leading to
 633 very similar solve times for these methods. Once again, Split is the least performing
 634 option due to its lack of robustness. Lastly, SSAMG-opt and BoomerAMG have

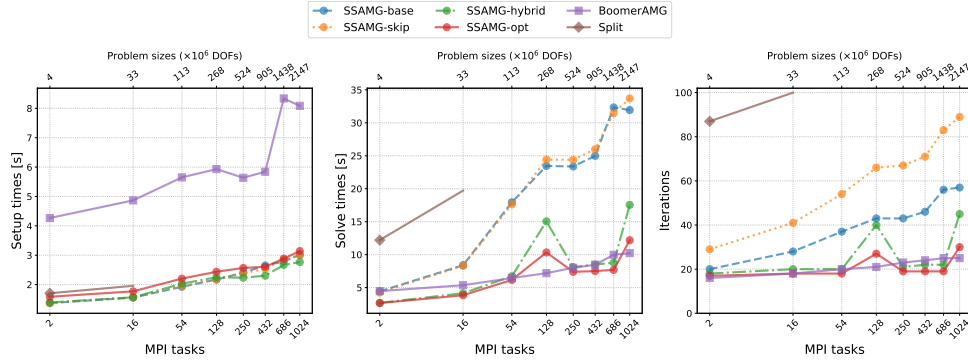


FIGURE 13. *Weak scalability results for test case 4. Three metrics are shown in the figure, i.e., setup phase times in seconds (left); solve phase times in seconds (middle), and number of iterations (right). All curves are plotted with respect to the number of MPI tasks, N_{procs} , which varies from 2 ($p = 1$) up to 1024 ($p = 8$).*

635 similar performance, with SSAMG-opt being slightly better for various cases, but
 636 BoomerAMG showing more consistent performance here.

637 Setup times of the SSAMG variants are very similar. Listing them in decreasing
 638 order of times, SSAMG-base and SSAMG-skip show nearly the same values, followed
 639 by SSAMG-opt, and SSAMG-hybrid is the fastest option. The results for Split are
 640 better here than in the previous test cases, and this is due to the small number of semi-
 641 structured parts involved in this SAMR problem. Still, SSAMG leads to the fastest
 642 options. The slowest method for setup is again BoomerAMG, while SSAMG-opt
 643 shows speedups with respect to the latter method of 2.6x for $p = 1$ and 2.7x for
 644 $p = 8$.

645 **5. Conclusions.** In this paper, we presented a novel algebraic multigrid method,
 646 built on the semi-structured interface in *hypre*, capable of exploiting knowledge about
 647 the problem’s structure and having the potential of being faster than an unstruc-
 648 tured algebraic multigrid method such as BoomerAMG on CPUs and accelerators.
 649 Moreover, SSAMG features a multigrid hierarchy with controlled stencil sizes and
 650 significantly improved setup times.

651 We developed a distributed parallel implementation of SSAMG for CPU archi-
 652 tectures in *hypre*. Furthermore, we tested its performance, when used as a preconditioner
 653 to PCG, for a set of semi-structured problems featuring distinct characteristics
 654 in terms of grid, stencil coefficients, and anisotropy. SSAMG proves to be numerically
 655 scalable for problems having up to a few billion degrees of freedom and its current
 656 implementation achieves speedups with respect to BoomerAMG up to a factor of 3.5
 657 for the setup phase and 1.8 for the solve phase.

658 For future work, we plan to improve different aspects of SSAMG and its implemen-
 659 tation. We will further investigate SSAMG convergence for more complex problems
 660 than have been considered so far. We want to explore adding an unstructured com-
 661 ponent to the prolongation matrix to improve interpolation across part boundaries
 662 and evaluate how this benefits convergence and time to solution. We also plan to add
 663 a non-Galerkin option for computing coarse operators targeting isotropic problems
 664 since this approach applied in PFMG has shown excellent runtime improvements on
 665 both CPU and GPU. Finally, we will develop a GPU implementation for SSAMG.

666 **Acknowledgments.** This material is based upon work supported by the U.S.
 667 Department of Energy, Office of Science, Office of Advanced Scientific Computing
 668 Research, Scientific Discovery through Advanced Computing (SciDAC) program.

669

REFERENCES

- 670 [1] S. ADAMS, R. FORD, M. HAMBLEY, J. HOBSON, I. KAVČIČ, C. MAYNARD, T. MELVIN,
 671 E. MÜLLER, S. MULLERWORTH, A. PORTER, M. REZNY, B. SHIPWAY, AND R. WONG,
 672 *LFRic: Meeting the challenges of scalability and performance portability in weather and*
 673 *climate models*, Journal of Parallel and Distributed Computing, 132 (2019), pp. 383–396,
 674 <https://doi.org/10.1016/j.jpdc.2019.02.007>.
- 675 [2] S. F. ASHBY AND R. D. FALGOUT, *A parallel multigrid preconditioned conjugate gradient al-*
 676 *gorithm for groundwater flow simulations*, Nuclear Science and Engineering, 124 (1996),
 677 pp. 145 – 159, <https://doi.org/10.13182/nse96-a24230>.
- 678 [3] A. H. BAKER, R. D. FALGOUT, T. V. KOLEV, AND U. M. YANG, *Multigrid smoothers for*
 679 *ultraparallel computing*, SIAM Journal on Scientific Computing, 33 (2011), pp. 2864–2887,
 680 <https://doi.org/10.1137/100798806>.
- 681 [4] A. H. BAKER, R. D. FALGOUT, T. V. KOLEV, AND U. M. YANG, *Scaling Hypre's Multigrid*
 682 *Solvers to 100,000 Cores*, Springer London, London, 2012, pp. 261–279, [https://doi.org/](https://doi.org/10.1007/978-1-4471-2437-5_13)
 683 [10.1007/978-1-4471-2437-5_13](https://doi.org/10.1007/978-1-4471-2437-5_13).
- 684 [5] B. BERGEN, G. WELLEN, F. HÜLSEMANN, AND U. RÜDE, *Hierarchical hybrid grids: achieving*
 685 *TERAFLOP performance on large scale finite element simulations*, International Journal of
 686 of Parallel, Emergent and Distributed Systems, 22 (2007), pp. 311–329, [https://doi.org/](https://doi.org/10.1080/17445760701442218)
 687 [10.1080/17445760701442218](https://doi.org/10.1080/17445760701442218).
- 688 [6] B. K. BERGEN AND F. HÜLSEMANN, *Hierarchical hybrid grids: data structures and core algo-*
 689 *rithms for multigrid*, Numerical Linear Algebra with Applications, 11 (2004), pp. 279–291,
 690 <https://doi.org/10.1002/nla.382>.
- 691 [7] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A Multigrid Tutorial, Second Edition*,
 692 Society for Industrial and Applied Mathematics, second ed., 2000, [https://doi.org/10.1137/](https://doi.org/10.1137/1.9780898719505)
 693 [1.9780898719505](https://doi.org/10.1137/1.9780898719505).
- 694 [8] P. N. BROWN, R. D. FALGOUT, AND J. E. JONES, *Semicoarsening multigrid on distributed*
 695 *memory machines*, SIAM Journal on Scientific Computing, 21 (2000), pp. 1823–1834,
 696 <https://doi.org/10.1137/S1064827598339141>.
- 697 [9] S. L. CORNFORD, D. F. MARTIN, D. T. GRAVES, D. F. RANKEN, A. M. LE BROCCQ, R. M.
 698 GLADSTONE, A. J. PAYNE, E. G. NG, AND W. H. LIPSCOMB, *Adaptive mesh, finite volume*
 699 *modeling of marine ice sheets*, Journal of Computational Physics, 232 (2013), pp. 529–549,
 700 <https://doi.org/10.1016/j.jcp.2012.08.037>.
- 701 [10] H. DE STERCK, U. M. YANG, AND J. J. HEYS, *Reducing complexity in parallel algebraic*
 702 *multigrid preconditioners*, SIAM Journal on Matrix Analysis and Applications, 27 (2006),
 703 pp. 1019–1039, <https://doi.org/10.1137/040615729>.
- 704 [11] C. ENGWER, R. D. FALGOUT, AND U. M. YANG, *Stencil computations for PDE-based applica-*
 705 *tions with examples from DUNE and hypre*, Concurrency and Computation: Practice and
 706 Experience, 29 (2017), p. e4097, <https://doi.org/10.1002/cpe.4097>.
- 707 [12] R. D. FALGOUT, J. E. JONES, AND U. M. YANG, *Conceptual interfaces in hypre*, Future Gen-
 708 eration Computer Systems, 22 (2006), pp. 239–251, [https://doi.org/10.1016/j.future.2003.](https://doi.org/10.1016/j.future.2003.09.006)
 709 [09.006](https://doi.org/10.1016/j.future.2003.09.006).
- 710 [13] R. D. FALGOUT, J. E. JONES, AND U. M. YANG, *The design and implementation of hypre,*
 711 *a library of parallel high performance preconditioners*, in Numerical Solution of Par-
 712 tial Differential Equations on Parallel Computers, A. M. Bruaset and A. Tveito, eds.,
 713 Berlin, Heidelberg, 2006, Springer Berlin Heidelberg, pp. 267–294, [https://doi.org/10.](https://doi.org/10.1007/3-540-31619-1_8)
 714 [1007/3-540-31619-1_8](https://doi.org/10.1007/3-540-31619-1_8).
- 715 [14] R. D. FALGOUT, R. LI, B. SJÖGREEN, L. WANG, AND U. M. YANG, *Porting hypre to heteroge-*
 716 *neous computer architectures: Strategies and experiences*, Parallel Computing, 108 (2021),
 717 p. 102840, <https://doi.org/10.1016/j.parco.2021.102840>.
- 718 [15] R. D. FALGOUT AND U. M. YANG, *hypre: A library of high performance preconditioners*, in
 719 Computational Science — ICCS 2002, P. M. A. Sloot, A. G. Hoekstra, C. J. K. Tan, and
 720 J. J. Dongarra, eds., Berlin, Heidelberg, 2002, Springer Berlin Heidelberg, pp. 632–641,
 721 [https://doi.org/10.1007/3-540-47789-6.66](https://doi.org/10.1007/3-540-47789-6_66).
- 722 [16] B. GANIS, G. PENCHEVA, AND M. F. WHEELER, *Adaptive mesh refinement with an en-*
 723 *hanced velocity mixed finite element method on semi-structured grids using a fully cou-*
 724 *pled solver*, Computational Geosciences, 23 (2019), pp. 149–168, <https://doi.org/10.1007/>

- 725 [s10596-018-9789-6](https://doi.org/10.1007/978-3-662-43880-0_50).
- 726 [17] B. GMEINER AND U. RÜDE, *Peta-scale hierarchical hybrid multigrid using hybrid paralleliza-*
727 *tion*, in Large-Scale Scientific Computing, I. Lirkov, S. Margenov, and J. Waśniewski,
728 eds., Berlin, Heidelberg, 2014, Springer Berlin Heidelberg, pp. 439–447, [https://doi.org/](https://doi.org/10.1007/978-3-662-43880-0_50)
729 [10.1007/978-3-662-43880-0_50](https://doi.org/10.1007/978-3-662-43880-0_50).
- 730 [18] B. GMEINER, U. RÜDE, H. STENDEL, C. WALUGA, AND B. WOHLMUTH, *Towards textbook effi-*
731 *ciency for parallel multigrid*, Numerical Mathematics: Theory, Methods and Applications,
732 8 (2015), p. 22–46, <https://doi.org/10.4208/nmtma.2015.w10si>.
- 733 [19] V. E. HENSON AND U. M. YANG, *BoomerAMG: A parallel algebraic multigrid solver and pre-*
734 *conditioner*, Applied Numerical Mathematics, 41 (2002), pp. 155–177, [https://doi.org/10.](https://doi.org/10.1016/S0168-9274(01)00115-5)
735 [1016/S0168-9274\(01\)00115-5](https://doi.org/10.1016/S0168-9274(01)00115-5). Developments and Trends in Iterative Methods for Large
736 Systems of Equations - in memorium Rudiger Weiss.
- 737 [20] R. D. HORNING AND S. R. KOHN, *Managing application complexity in the SAMRAI object-*
738 *oriented framework*, Concurrency and Computation: Practice and Experience, 14 (2002),
739 pp. 347–368, <https://doi.org/10.1002/cpe.652>.
- 740 [21] *hypre: High performance preconditioners*. <http://www.llnl.gov/CASC/hypre/>, [https://github.](https://github.com/hypre-space/hypre)
741 [com/hypre-space/hypre](https://github.com/hypre-space/hypre).
- 742 [22] N. KOHL AND U. RÜDE, *Textbook efficiency: Massively parallel matrix-free multigrid for the*
743 *stokes system*, SIAM Journal on Scientific Computing, 44 (2022), pp. C124–C155, [https:](https://doi.org/10.1137/20M1376005)
744 [//doi.org/10.1137/20M1376005](https://doi.org/10.1137/20M1376005).
- 745 [23] R. LI, B. SJÖGREEN, AND U. M. YANG, *A new class of AMG interpolation methods based on*
746 *matrix-matrix multiplications*, SIAM Journal on Scientific Computing, 43 (2021), pp. S540–
747 S564, <https://doi.org/10.1137/20M134931X>.
- 748 [24] M. MAYR, L. BERGER-VERGIAT, P. OHM, AND R. S. TUMINARO, *Non-invasive multigrid for*
749 *semi-structured grids*, SIAM Journal on Scientific Computing, 44 (2022), pp. A2734–A2764,
750 <https://doi.org/10.1137/20M1375413>.
- 751 [25] C. RODRIGO, F. J. GASPAR, AND F. J. LISBONA, *Multigrid methods on semi-structured grids*,
752 Archives of Computational Methods in Engineering, 19 (2012), pp. 499–538, [https://doi.](https://doi.org/10.1007/s11831-012-9078-9)
753 [org/10.1007/s11831-012-9078-9](https://doi.org/10.1007/s11831-012-9078-9).
- 754 [26] B. RUNNELS, V. AGRAWAL, W. ZHANG, AND A. ALMGREN, *Massively parallel finite differ-*
755 *ence elasticity using block-structured adaptive mesh refinement with a geometric multigrid*
756 *solver*, Journal of Computational Physics, 427 (2021), p. 110065, [https://doi.org/10.1016/](https://doi.org/10.1016/j.jcp.2020.110065)
757 [j.jcp.2020.110065](https://doi.org/10.1016/j.jcp.2020.110065).
- 758 [27] K. STÜBEN, *A review of algebraic multigrid*, Journal of Computational and Applied Mathemat-
759 ics, 128 (2001), pp. 281–309, [https://doi.org/10.1016/S0377-0427\(00\)00516-1](https://doi.org/10.1016/S0377-0427(00)00516-1). Numerical
760 Analysis 2000. Vol. VII: Partial Differential Equations.
- 761 [28] U. M. YANG, *On long-range interpolation operators for aggressive coarsening*, Numerical Linear
762 Algebra with Applications, 17 (2010), pp. 453–472, <https://doi.org/10.1002/nla.689>.
- 763 [29] W. ZHANG, A. ALMGREN, V. BECKNER, J. BELL, J. BLASCHKE, C. CHAN, M. DAY, B. FRIESEN,
764 K. GOTT, D. GRAVES, ET AL., *AMReX: a framework for block-structured adaptive mesh*
765 *refinement*, Journal of Open Source Software, 4 (2019), pp. 1370–1370, [https://doi.org/10.](https://doi.org/10.21105/joss.01370)
766 [21105/joss.01370](https://doi.org/10.21105/joss.01370).
- 767 [30] W. ZHANG, A. MYERS, K. GOTT, A. ALMGREN, AND J. BELL, *AMReX: Block-structured adap-*
768 *tive mesh refinement for multiphysics applications*, The International Journal of High
769 Performance Computing Applications, 35 (2021), pp. 508–526, [https://doi.org/10.1177/](https://doi.org/10.1177/10943420211022811)
770 [10943420211022811](https://doi.org/10.1177/10943420211022811).