

ALGEBRAIC MULTIGRID FOR LINEAR SYSTEMS OBTAINED BY EXPLICIT ELEMENT REDUCTION*

THOMAS A. BRUNNER[†] AND TZANIO V. KOLEV[‡]

Abstract. We consider sparse linear systems, where a set of “interior” degrees of freedom have been eliminated in order to reduce the problem size. This elimination is assumed to be local, so the “interior” principal sub-matrix is block-diagonal, and the resulting Schur complement is still sparse. For it to be beneficial, the elimination process should lead to reduced memory requirements, but equally importantly, it should also result in an algebraic problem that can be solved efficiently. In this paper we propose a general element reduction approach, and show how the elimination process can exploit a particular “sub-zonal” discretization to maintain the sparsity of the Schur complement. We also investigate Algebraic Multigrid (AMG) solution algorithms applied to the reduced problem, and we discuss the influence of the local elimination on solver-related properties of the matrix, such as near-nullspace preservation and the availability of stable subspace decompositions. We focus on BoomerAMG, a parallel variant of classical Ruge-Stüben AMG, applied to scalar diffusion problems [18], and the Auxiliary-space Maxwell Solver (AMS) for electromagnetic diffusion applications [22]. In the electromagnetic case, we establish algebraically a reduced version of the HX decomposition from [19], and consider a modification of the reduction process that targets the singular problems arising in simulations with pure void (zero conductivity) regions. For scalar diffusion problems, our particular stencil analysis shows that the reduction has a positive effect on meshes with stretched elements. We present a number of 2D, 3D and axi-symmetric numerical experiments, which demonstrate that the combination of an appropriately chosen local elimination with the use of the BoomerAMG and AMS solvers can lead to significant improvements in the overall solution time.

Key words. algebraic multigrid; static condensation; auxiliary space solvers; electromagnetic problems; Schur complement reduction

AMS subject classifications. 65F08, 65N30, 65N55

1. Introduction. In this paper we consider the algebraic solution of the linear system

$$(1.1) \quad Ax = b,$$

where A is a sparse matrix arising in the finite element discretization of a given partial differential equation (PDE). We assume that the matrix A is symmetric and positive definite (SPD), with the exception of Section 4.3.1, where A is only semi-definite.

Suppose that A is very large and instead of solving (1.1) directly, we are interested in reducing the problem size by first eliminating some of the unknowns in the linear system. Specifically, consider the block splitting of A

$$A = \begin{pmatrix} A_{ii} & A_{ir} \\ A_{ri} & A_{rr} \end{pmatrix}$$

where i stands for the “interior” unknowns, the ones we plan to eliminate, while r corresponds to the “reduced” (remaining) unknowns, i.e. the ones that will remain in the matrix after the reduction is complete. A major assumption in our approach is that the elimination is local, in other words *the interior principal sub-matrix A_{ii} is block-diagonal.*

*This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, LLNL-JRNL-440891

[†]Weapons and Complex Integration, Lawrence Livermore National Laboratory, brunner6@llnl.gov,

[‡]Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, tzanio@llnl.gov

Leaving the semi-definite case for Section 4.3.1, we have that the matrix A_{ii} is SPD, so we can define the harmonic extension interpolation matrix

$$(1.2) \quad P = \begin{pmatrix} -A_{ii}^{-1}A_{ir} \\ I \end{pmatrix},$$

and the Schur complement

$$(1.3) \quad S = A_{rr} - A_{ri}A_{ii}^{-1}A_{ir} = P^tAP.$$

In this notation, the reduced linear system reads

$$(1.4) \quad Sx_r = b_r,$$

where $x = Px_r$ and $b_r = P^tb$. An important consequence of the locality of the elimination process is that the reduced matrix S is sparse. In contrast to other Schur complement approaches, where only the action of S is implemented, in this paper *we actually assemble S as a sparse matrix*. In fact, having access to the individual entries of S will play a critical role in our solution process.

Specifically, we investigate the reduction from (1.1) to (1.4) in the context of Algebraic Multigrid (AMG) for matrices A arising in scalar and electromagnetic diffusion applications. The central question of our inquiry is:

Can we solve larger problems faster by applying AMG to (1.4) instead of (1.1)?

Since the elimination process will generally create new connections between the reduced degrees of freedom, the matrix S may actually require more memory for storage than A , even though its number of unknowns is smaller. This can be especially problematic in 3D, where unlike in 2D, straightforward reduction approaches may lead to a significant increase in the number of nonzeros of the Schur complement. We will address this issue by considering a general element reduction approach in the settings of a particular ‘‘sub-zonal’’ discretization of interest in the following Section 2. We will show that exploiting the properties of this ‘‘sub-zonal’’ discretization is crucial to maintain the sparsity of the Schur complement.

More importantly, we need to investigate the influence of the elimination process on the AMG-related properties of the matrix, some of which are discussed in the overview of algebraic multigrid-type solvers given in Section 3. In particular, we will show in Section 4, that for the considered applications, *if AMG works for A , it will also work for S* . This result is especially interesting in the case of electromagnetic problems, where it is derived based on the proof, in Section 4.3, of a new, *reduced version of the Hiptmair-Xu (HX) decomposition* from [19]. We also show in Section 4.2.1 that, somewhat surprisingly, the elimination process can produce discretizations with improved multigrid properties for 2D problems on stretched grids.

The above considerations will allow us to ultimately give a positive answer to the central question of our inquiry. This will also be demonstrated by the performance of the solvers on 2D, 3D and axi-symmetric problems presented in the numerical experiments in Section 5.

1.1. Related Approaches. Our reduction approach is closely related to the method of static condensation and other elimination algorithms corresponding to different choices of P in (1.2). We review some of these related approaches below, skipping less relevant methods, such as non-overlapping domain decomposition [31, 26], where the Schur complement is dense and is not assembled. We remark that there

have been other algebraic solvers for Schur complement matrices, such as the AMLI method [39], which however require access to the original matrix A .

The classical static condensation algorithm of Wilson [41] was introduced in 1974 to “eliminate the internal degrees of freedom in a quadrilateral finite element constructed from four triangles”. Nowadays, static condensation is frequently used to eliminate the geometrically interior degrees of freedom in high-order finite element discretizations, see e.g., [36, §3.5.9] [27]. A critical property of static condensation is that the reduced matrix S has the same sparsity pattern as the principal sub-matrix A_{rr} , and therefore the solution of (1.4) always requires less memory than that of (1.1). Since the number of nonzeros of the matrix S will increase in the applications that we consider, we prefer to think of our elimination scheme as an element reduction approach (see Section 2.1) instead of a generalized type of static condensation.

Schur complement reduction, especially in the form of static condensation, has been applied in many areas of scientific computing, including mimetic discretizations for diffusion problems on generalized polyhedral meshes [8], hp -adaptive refinement for elliptic equations [27], discontinuous Galerkin methods [17] and physics-based preconditioning in magnetohydrodynamics (MHD) [15]. In a number of cases, AMG was used for solving the condensed system, leading to the central problem of interest in this paper. Frequently, AMG was reported to work well in practice, but we are not aware of any published attempts for analysis of its performance on the Schur complement, especially in the case of electromagnetic problems. We undertake this task in the main portion of the paper presented in Section 4.

Another vein of research has been exploring approximate Schur complements, which have been used e.g., in reservoir simulations, and in fact AMG itself has been described as “an approximate Schur complement approach” [37]. Over the years, AMG researchers have also employed approximate local elimination of degrees of freedom, where P was not the harmonic extension, but instead was defined by the local AMG interpolation operator [34]. One example in this direction is [23], where the “ideal” harmonic extension prolongator is approximated with a sparse matrix derived by replacing A_{ii} with a diagonal approximation. This idea arises also in energy-minimization AMG [25] and in multigrid upscaling [6, 24], where the emphasis is on capturing the homogenized properties of the problem in the interpolation P . In this paper we focus on the exact Schur complement reduction, and leave other choices of P for future work.

Finally, we note the cyclic reduction/multigrid research from [16], see also [12], where it is shown that multigrid converges faster on S for a specific algebraic elimination on a checkerboard colored structured grid. We view the current paper as an extension of this result to general AMG settings, including electromagnetic problems. The work from [16] was continued in [42], where a detailed analysis of the multigrid smoothing efficiency was performed. This is complimentary to the classical AMG topics we consider in the present paper, which are mostly concerned with the interpolation operator and the preservation of the near-nullspace.

2. Memory Considerations. In this section we give a detailed description of our view of the elimination process as a general finite element reduction approach. We also investigate when this approach results in a matrix S having a smaller number of nonzero entries than the original matrix A . The number of nonzero entries is proportional to the memory required for storing the matrix, as well as the number of floating point operations needed for a matrix-vector multiply.




FIG. 2.1. In XY geometry, the full triangular grid has nodal degrees of freedom. The reduced elements are simply the quadrilateral elements, with the interior degrees of freedom identified at zone centers. The final reduced mesh has the same number of degrees of freedom and sparsity pattern as the original quadrilateral grid.

2.1. An Element Reduction Approach. A finite element discretization can be specified by a set of elements $\{e\}$, with corresponding degrees of freedom (dofs) $\{x_e\}$, and local element matrices $\{A_e\}$, which are assembled to form A . With that definition in mind, we consider the following 4-step elimination process of local degrees of freedom, see Figure 2.1 and Figure 2.3 for illustration:

1. *Choose a set of reduced elements.* Each reduced element E is simply a union of several neighboring original elements, similar to the agglomerated elements in the AMGe framework [7, 21]. The set of all reduced elements, $\{E\}$, should provide a non-overlapping decomposition of the original set of elements into small subdomains.
2. *Determine the interior degrees of freedom.* For a given reduced element, we define its interior dofs to be those which do not interact (through the sparsity pattern of A) with dofs outside of the reduced element.
3. *Define the reduced degrees of freedom.* The reduced degrees of freedom are the degrees of freedom on the boundary of the reduced elements.
4. *Define the reduced element matrices.* The reduced element matrices are the local Schur complements

$$S_E = A_{rr,E} - A_{ri,E} A_{ii,E}^{-1} A_{ir,E} = P_E^t A_E P_E,$$

where the subscript E denotes a restriction to the degrees of freedom in the particular reduced element. These are indeed element matrices, since S can be assembled from $\{S_E\}$.

The set of reduced elements is the only input to the above algorithm, so we can specify the elimination process by simply describing the reduced elements. More importantly, the output of the algorithm is an induced finite element discretization on the reduced grid with elements $\{E\}$, degrees of freedom $\{x_E\}$, and element matrices $\{S_E\}$. *We emphasize that this provides a consistent way of introducing a finite element discretization on any set of reduced elements, including geometries without a natural reference element.* An example in this direction will be provided in Section 2.2.2.

2.2. A Sub-zonal Discretization. In this section we apply the element reduction approach to a particular “sub-zonal” finite element discretization of the second-




FIG. 2.2. The sub-zonal mesh in 2D (left) and 3D (right). In 3D, 24 tetrahedral sub-elements are added to each hexahedron as show in the center.

order definite Maxwell bilinear form

$$(2.1) \quad a_{\text{curl}}(\mathbf{u}_h, \mathbf{v}_h) = \left(\frac{\Delta t}{\mu} \nabla \times \mathbf{u}_h, \nabla \times \mathbf{v}_h \right) + (\sigma \mathbf{u}_h, \mathbf{v}_h)$$

on a given polyhedral domain Ω . Here, Δt is the simulation time-step, while $\sigma \geq 0$ and $\mu > 0$ are the conductivity and the magnetic permeability coefficients, respectively. We restrict ourselves to finite element functions \mathbf{u}_h and \mathbf{v}_h belonging to \mathbf{V}_h – the space of lowest order Nédélec (edge) finite elements [29] [28, §5–§8]. Recall that the degrees of freedom in \mathbf{V}_h are the tangential moments associated with the edges of the mesh. For simplicity, we assume homogeneous Dirichlet boundary conditions in \mathbf{V}_h , representing a medium surrounded by a perfect conductor. The subscript h is a reminder that the above quantities are defined on a discrete computational mesh. We will also use h as a constant, which denotes the (globally quasi-uniform) mesh size, in Section 3.2.

For generality, we also consider the case of “pure void zones” in (2.1), where $\sigma = 0$ in the elements comprising the non-conductive part of the domain. This implies that the corresponding matrix \mathbf{A} will be singular as discussed in Section 4.3.1. A popular alternative is to set σ to a small positive number in the non-conductive region relative to its value in the conductor, $\sigma_{nc}/\sigma_c \approx 0$. This choice produces a significant jump in the coefficients of $a_{\text{curl}}(\cdot, \cdot)$ and generates a large near-nullspace, which makes \mathbf{A} very challenging for AMG.

We investigate three different PDE models for the above Maxwell problem, which we denote with XY , RZ and $3D$. We will refer to these as modes, or geometries. For example, the $3D$ mode corresponds to the regular discretization of $a_{\text{curl}}(\cdot, \cdot)$ in 3D. The XY and RZ geometries are described below. We sometimes refer collectively to them as 2D problems.

The XY mode corresponds to a problem with data and solution of the form $\mathbf{u}_h(x, y, z) = (0, 0, p_h(x, y))$. The problem can then be reduced to a 2D nodal discretization of the scalar diffusion bilinear form

$$(2.2) \quad a_{\text{grad}}(p_h, q_h) = \left(\frac{\Delta t}{\mu} \nabla p_h, \nabla q_h \right) + (\sigma p_h, q_h),$$

where p_h and q_h belong to S_h – the space of linear finite elements on the mesh in the 2D plane orthogonal to the solution. The degrees of freedom are the values at the vertices of that mesh.

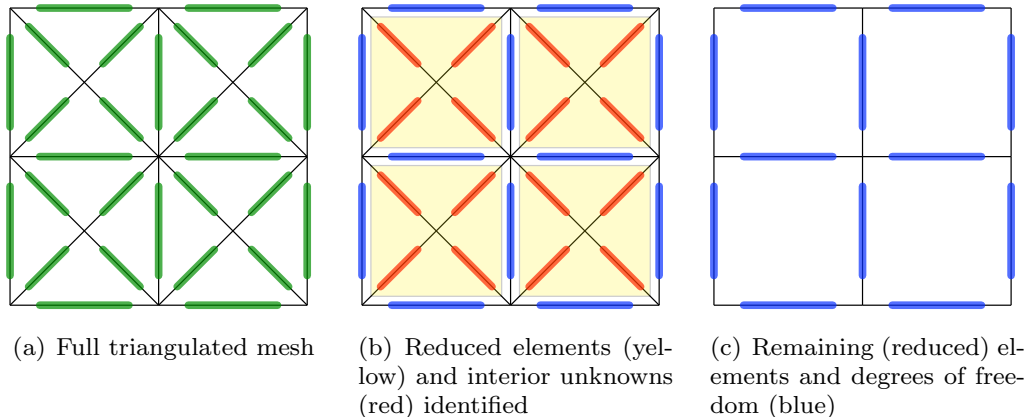


FIG. 2.3. In RZ geometry, the triangular sub-zonal grid has edge degrees of freedom. The reduced elements are simply the quadrilateral elements. The final reduced mesh has the same number of degrees of freedom and sparsity pattern as the original quadrilateral grid.

The RZ geometry corresponds to an axi-symmetric problem reduced to a 2D meridian plane. In this case we get a discretization of the 2D Maxwell bilinear form

$$(2.3) \quad a_{rz}(\mathbf{u}_h, \mathbf{v}_h) = \left(r \frac{\Delta t}{\mu} \nabla \times \mathbf{u}_h, \nabla \times \mathbf{v}_h \right) + (r\sigma \mathbf{u}_h, \mathbf{v}_h),$$

where r is the distance to the axis of rotation, $\nabla \times \mathbf{u}_h = -\partial_r \mathbf{u}_h + \partial_z \mathbf{u}_h$ and \mathbf{V}_h is the 2D Nédélec space on the mesh in the meridian plane. As in 3D, the degrees of freedom are associated with the edges of the mesh.

In the remainder of the paper, we consider a particular radiation hydrodynamics application [32], which simulates inertial confinement fusion experiments. The hydrodynamics discretization uses a compatible algorithm [10, 5], based on a sub-zonal mesh, obtained by splitting the initial arbitrary polygonal or polyhedral grid into triangular or tetrahedral elements. Recently a resistive MHD package [9] was added to the code in order to support simulations of Z-pinch experiments at Sandia National Laboratories. We focus on the electromagnetic linear systems that are solved in this package in XY , RZ and $3D$ modes.

In order to be consistent with the hydrodynamics discretization, (2.2), (2.3) and (2.1) are also discretized on the sub-zonal mesh. In 2D, each sub-triangle is defined by two consecutive vertices and the element center, while in 3D the tetrahedrons are defined by connecting the center of each polyhedral element with the center of a face and two consecutive vertices on that face. This splitting is illustrated in Figure 2.2 for a quadrilateral and hexahedron. While we have the capability to run on an arbitrary mesh, all the tests in this paper are run with initial quadrilateral or hexahedral meshes.

2.2.1. Element Reduction in 2D. We first address the elimination process in the XY and RZ modes of the target application. A natural choice in this case is to pick the reduced elements to be the original quadrilaterals, as shown in Figure 2.1 and Figure 2.3. We point out that even though in both cases we recover the associated quadrilateral mesh, *the reduced finite element discretizations are not the ones based on quadrilateral finite elements*, see Section 4.2.1.

In order to measure the reduction in memory, we introduce the notation $n_{\text{rows}}(\mathbf{A})$ and $n_{\text{nz}}(\mathbf{A})$ for the number of rows and nonzero entries in a matrix. We also denote




FIG. 2.4. Hexahedral (left) and octahedral (right) reduced elements in 3D. The octahedral element spans part of two of the original hexahedrons, and is formed by gluing the eight tetrahedrons adjacent to the original hexahedral face (yellow). The interior degrees of freedom, which will be eliminated, are highlighted in red. The remaining reduced degrees of freedom are colored in blue.

with \mathcal{N}_z , \mathcal{N}_v and \mathcal{N}_e the number of quadrilateral elements (a.k.a. zones), vertices and edges, respectively.

Simple counting of the local balance of nonzero entries in each reduced element shows that we get memory reduction in both cases: $n_{\text{nz}}(\mathbf{S}_{XY}) = n_{\text{nz}}(\mathbf{A}_{XY}) - 5\mathcal{N}_z$ and $n_{\text{nz}}(\mathbf{S}_{RZ}) = n_{\text{nz}}(\mathbf{A}_{RZ}) - 16\mathcal{N}_z$.

TABLE 2.1

Asymptotic estimates for the XY and RZ modes in 2D. The n_{nz} estimates in the third and fourth columns assume an original structured quadrilateral grid. The last column shows the memory reduction on general grids.

matrix	n_{rows} (reduction)	n_{nz} (reduction)	$n_{\text{nz}}/n_{\text{rows}}$	$n_{\text{nz}}(\mathbf{S}) - n_{\text{nz}}(\mathbf{A})$
\mathbf{A}_{XY}	$2 \mathcal{N}_z$	$14 \mathcal{N}_z$	7	
\mathbf{S}_{XY}	$\mathcal{N}_z (\times 2)$	$9 \mathcal{N}_z (\times 1.6)$	9	$-5 \mathcal{N}_z$
\mathbf{A}_{RZ}	$6 \mathcal{N}_z$	$30 \mathcal{N}_z$	5	
\mathbf{S}_{RZ}	$2 \mathcal{N}_z (\times 3)$	$14 \mathcal{N}_z (\times 2.1)$	7	$-16 \mathcal{N}_z$

To get more qualitative estimates, we assume structured grid and use the relationships $\mathcal{N}_z \sim \mathcal{N}_v \sim \frac{1}{2}\mathcal{N}_e$, which follow from asymptotic analysis on quadrilateral meshes. The results are presented in Table 2.1 and show good reduction in the memory requirements both in the XY and RZ cases. The savings are especially significant on RZ structured grids.

2.2.2. Element Reduction in 3D. Next, we consider the elimination process in the 3D case, where there are a number of possibilities for defining the reduced elements. We concentrate on two particular choices: the matrix \mathbf{S}_H , where the reduced elements are the original hexahedrons, and the matrix \mathbf{S}_O , which is based on the octahedral reduced elements that are composed of the eight tetrahedrons adjacent to each face in the original hexahedral mesh, see Figure 2.4. The octahedral elements provide an example of a geometry where the direct application of the finite element method is not straightforward.

Let \mathcal{N}_z , \mathcal{N}_v , \mathcal{N}_e and \mathcal{N}_f denote the number of hexahedral elements, vertices, edges and faces, respectively. Asymptotically, on hexahedral grids we have $\mathcal{N}_z \sim \mathcal{N}_v \sim \frac{1}{3}\mathcal{N}_e \sim \frac{1}{3}\mathcal{N}_f$.

TABLE 2.2

Asymptotic estimates for the two element reduction options in 3D. The n_{nz} estimates in the third and fourth columns assume an original structured hexahedral grid. The last column shows the memory reduction on general grids.

matrix	n_{rows} (reduction)	n_{nz} (reduction)	$n_{\text{nz}}/n_{\text{rows}}$	$n_{\text{nz}}(\text{S}) - n_{\text{nz}}(\text{A})$
A	$29 \mathcal{N}_z$	$461 \mathcal{N}_z$	16	
S_H	$15 \mathcal{N}_z (\times 1.9)$	$1107 \mathcal{N}_z (\times 0.4)$	74	$+646 \mathcal{N}_z$
S_O	$11 \mathcal{N}_z (\times 2.6)$	$335 \mathcal{N}_z (\times 1.4)$	30	$-126 \mathcal{N}_z$

Performing some straightforward calculations one can derive the estimates for the original and reduced matrices presented in Table 2.2. The hexahedral-based reduction leads to a significant increase in the number of nonzeros, and unlike in the 2D case, S_H requires more memory than the original matrix. In contrast, the octahedral-based elimination always results in memory reduction, which asymptotically is comparable with the XY case on structured grids.

Similar analysis can be performed for all the other possible reduced element choices, by enumerating what type of edges will be eliminated (element-face, element-vertex, face-vertex, edge-vertex). This analysis shows that, in fact, *the octahedral-based reduction is the optimal choice with respect to both the problem size and the memory usage in 3D.*

3. Algebraic Multigrid Solvers. In this section we review some of the properties of AMG methods which will be needed in the following sections. We use the family of algebraic multigrid solvers, since they provide particularly attractive solution techniques for the problems described in the previous section, where A is very large, sparse and badly conditioned due to the jumps in the coefficients and the presence of both high- and low-frequency eigenvectors in the spectrum of the PDE. Below, we consider two AMG-type solvers targeting scalar and electromagnetic diffusion problems, respectively.

3.1. Classical AMG. The classical Ruge-Stüben AMG algorithm [35] has been very successful on matrices arising from nodal discretizations of the scalar diffusion bilinear form $a_{\text{grad}}(\cdot, \cdot)$. Several efficient and robust implementations of AMG are currently available, including the parallel solver BoomerAMG in *hypra* [18, 20], which has demonstrated robustness with respect to jumps in the coefficients and scalability on more than 125 thousand processors [13, 2]. Smoothed aggregation [38, 14] is another viable algebraic solver for (2.2).

A critical component of classical AMG is the knowledge of the *near-nullspace* which consists of normalized vectors e satisfying

$$Ae \approx 0.$$

Typically, a near-nullspace vector e can be approximated locally by a constant, which is reflected in the construction of AMG interpolation operators and the requirement that they have row sums of one. Furthermore, it is assumed that e varies smoothly in the direction of “strong connections”.

The concept of *strength of connection* forms the foundation of the coarsening and the interpolation procedures in classical AMG, and is defined as follows: the dof i

depends strongly on the dof j if

$$-A_{ij} \geq \theta \max_{k \neq i} \{-A_{ik}\}.$$

Here, $0 < \theta \leq 1$ is the strength threshold parameter. Note that there are other definitions for strength of connection, such as the more symmetric version used in the smoothed aggregation algorithm, but we focus on classical AMG and will later refer to the θ parameter given above.

3.2. Auxiliary-space Maxwell Solver (AMS). Traditionally, electromagnetic diffusion problems described by the second-order definite Maxwell forms $a_{\text{curl}}(\cdot, \cdot)$ and $a_{\text{rz}}(\cdot, \cdot)$ have been challenging for classical AMG due to the large number of locally-supported gradients in their near-nullspaces. Even though the research in extending classical AMG ideas to these problems have met with some success [33, 3], we choose to focus on the recently developed Auxiliary-space Maxwell Solver (AMS) [22]. This choice is motivated by the strong theoretical foundation of AMS [19] and its clear connection to classical AMG for nodal problems. A related approach, based on compatible discretization principles, can be found in [4].

The AMS algorithm is an example of an algebraic solver that takes advantage of discretization information. Specifically, AMS targets the lowest order Nédélec finite element methods in the RZ and $3D$ cases that were considered in Section 2.2. The solver is built on the relation between the Nédélec space \mathbf{V}_h and the “auxiliary” nodal linear finite element spaces S_h and $\mathbf{S}_h \equiv S_h^d$, where $d = 2$ for RZ and $d = 3$ for $3D$. In particular, to characterize the near-nullspace of A , AMS employs two natural interpolation operators from the nodal spaces into \mathbf{V}_h . These operators are illustrated in Figure 3.1. The first one is the discrete representation of the mapping

$$p_h \in S_h \mapsto \nabla p_h \in \mathbf{V}_h,$$

which we refer to as the *discrete gradient matrix* G . Note that G is a simple topological table describing the edges of the mesh (including their orientation) in terms of the vertices. The second operator, Π , in Figure 3.1 is the matrix representation of the *Nédélec interpolation operator* Π_h which transfers linear vector fields $\mathbf{z}_h \in \mathbf{S}_h$ into \mathbf{V}_h by computing the corresponding Nédélec degrees of freedom:

$$\Pi_h \mathbf{z}_h = \sum_e \left(\int_e \mathbf{z}_h \cdot \mathbf{t}_e ds \right) \Phi_e.$$

Here Φ_e is the basis functions corresponding to edge e with unit tangent \mathbf{t}_e . As described in [22], Π can be computed from G and the vectors \mathbf{x} , \mathbf{y} and \mathbf{z} providing the coordinates of the vertices of the mesh (in 3D), since $\Pi = [\Pi_x \ \Pi_y \ \Pi_z]$ with $(\Pi_x)_{ij} = |G_{ij}|(G\mathbf{x})_i/2$, $(\Pi_y)_{ij} = |G_{ij}|(G\mathbf{y})_i/2$ and $(\Pi_z)_{ij} = |G_{ij}|(G\mathbf{z})_i/2$.

Given the discrete gradient matrix G and the Nédélec interpolation matrix Π (or the coordinate vectors \mathbf{x} , \mathbf{y} and \mathbf{z}), we can construct various additive and multiplicative AMS methods. For example, AMS can take the following form as an additive preconditioner:

$$(3.1) \quad \mathbf{B} = \mathbf{R} + G\mathbf{B}_G G^T + \Pi \mathbf{B}_\Pi \Pi^T.$$

Here \mathbf{R} is a *point smoother* for A (e.g., a sweep of Gauss-Seidel), while \mathbf{B}_G and \mathbf{B}_Π are *classical AMG V-cycles* for the variationally-constructed scalar and vector Poisson-like coarse grid nodal matrices $G^T A G$ and $\Pi^T A \Pi$, respectively. Note that AMS is not




FIG. 3.1. AMS interpolation operators G and Π relating the Nédélec space V_h (center) with the two nodal auxiliary spaces S_h (left) and S_h (right).

a direct AMG solver for $a_{\text{curl}}(\cdot, \cdot)$, but is instead based on classical AMG sweeps in the nodal auxiliary subspaces (alternatively AMS can be viewed as an AMG method with multiple coarse grids). One can also derive multiplicative solver versions, as well as variants based on related auxiliary spaces. For example, the following 4-subspace version, proposed in [22], works well in practice

$$(3.2) \quad B = R + GB_G G^T + \Pi_x B_{\Pi_x} \Pi_x^T + \Pi_y B_{\Pi_y} \Pi_y^T + \Pi_z B_{\Pi_z} \Pi_z^T.$$

Here B_{Π_x} is a V-cycle for the matrix $\Pi_x^T A \Pi_x$, and similarly for B_{Π_y} and B_{Π_z} . The subspace problems in this case are scalar Poisson-like matrices restricted to subsets of the variables. Indeed, let $z = (z_x, 0, 0)$ correspond to the finite element function $z_h = (z_{h,x}, 0, 0)$. On meshes with simplex elements we have $\nabla \times z_h \in \nabla \times V_h$, so

$$\nabla \times \Pi_h z_h = \nabla \times z_h = (0, \partial_z z_{h,x}, -\partial_y z_{h,x})$$

for $z_h \in S_h$. Therefore,

$$(\Pi_x^T A \Pi_x z_x, z_x) = a_{\text{curl}}(\Pi_h z_h, \Pi_h z_h) = \left(\frac{\Delta t}{\mu} \nabla_{yz} z_{h,x}, \nabla_{yz} z_{h,x} \right) + (\sigma \Pi_h z_h, \Pi_h z_h),$$

so the higher-order term of $\Pi_x^T A \Pi_x$ is similar to the yz -restricted Laplacian operator $\partial_y^2 + \partial_z^2$. Hence, it is reasonable to expect that classical AMG will be a good solver for the subspace problems in (3.2).

In practice, AMS for $a_{\text{curl}}(\cdot, \cdot)$ has shown comparable performance to classical AMG for $a_{\text{grad}}(\cdot, \cdot)$. This robustness is based on the following fundamental theoretical result by Hiptmair and Xu [19], see also [30], which provides a general statement about the structure of the space V_h , similar to the classical Helmholtz decomposition [28, §7.2.1]. We use the notation $f \lesssim g$ and $f \gtrsim g$ to denote that $f \leq Cg$ and $f \geq Cg$, respectively with an absolute constant C .

THEOREM 3.1 (HX decomposition). *Any $u_h \in V_h$ can be split into*

$$(3.3) \quad u_h = v_h + \nabla p_h + \Pi_h z_h$$

where $v_h \in V_h$, $p_h \in S_h$ and $z_h \in S_h$ satisfy

$$h^{-1} \|v_h\|_0 + \|z_h\|_1 \lesssim \|\nabla \times u_h\|_0, \quad \|\nabla p_h\|_0 \lesssim \|u_h\|_0.$$

Without going into details, we note that (3.1) simply addresses each of the terms in (3.3), taking into account that the remainder v_h is small. Since $\|\Pi_h z_h\|_{H(\text{curl})} \lesssim \|z_h\|_1$, Theorem 3.1 can be used to assert the optimality of AMS in the case of

constant coefficients μ and σ (provided the subspace solvers are optimal). However, we are interested in problems with variable coefficients, so we switch to a new, purely matrix perspective on the HX decomposition, which will be the cornerstone of our analysis in Section 4.3.

LEMMA 3.2 (matrix form of HX). *Suppose that the matrix A has the following properties:*

- (A) *The HX stable decomposition holds in the sense that any vector u can be decomposed into*

$$u = v + Gp + \Pi z,$$

such that

$$(D_A v, v) + (AGp, Gp) + (A\Pi z, \Pi z) \lesssim (Au, u).$$

Here D_A is the diagonal of A .

- (B) *Classical AMG is an optimal solver for the nodal subspace matrices $G^T A G$ and $\Pi^T A \Pi$.*

Then AMS is an optimal solver for A .

Proof. Follows from classical Schwarz theory (e.g., applied to (3.1)), and the fact that point smoothers such as Gauss-Seidel are spectrally equivalent to D_A^{-1} , see [40, 44]. An analogous result holds for the decomposition corresponding to (3.2). \square

Note that the matrix form of HX given in condition (A) above corresponds to the following generalization of Theorem 3.1 to problems with variable coefficients:

$$(3.4) \quad h^{-2} \|\mathbf{v}_h\|_{\frac{\Delta t}{\mu}}^2 + \|\mathbf{v}_h\|_{\sigma}^2 + \|\nabla p_h\|_{\sigma}^2 + \|\nabla \mathbf{z}_h\|_{\frac{\Delta t}{\mu}}^2 + \|\mathbf{z}_h\|_{\sigma}^2 \lesssim \|\nabla \times \mathbf{u}_h\|_{\frac{\Delta t}{\mu}}^2 + \|\mathbf{u}_h\|_{\sigma}^2,$$

where the subscripts indicate weighted $L^2(\Omega)$ inner products. For example, if p and u are the coefficient vectors for p_h and \mathbf{u}_h , respectively, then

$$(AGp, Gp) = a_{\text{curl}}(\nabla p_h, \nabla p_h) = \|\nabla p_h\|_{\sigma}^2 \lesssim a_{\text{curl}}(\mathbf{u}_h, \mathbf{u}_h) = (Au, u).$$

A similar relation follows for the term involving \mathbf{v}_h since $\|\nabla \times \Phi_e\|_0^2 \sim h^{-2} \|\Phi_e\|_0^2$ and $\|\mathbf{v}_h\|_{\alpha}^2 \sim h^{2-d} (v, v)_{\alpha}$ imply

$$(D_A v, v) \sim h^{-d} (v, v)_{\frac{\Delta t}{\mu} + h^2 \sigma} \sim h^{-2} \|\mathbf{v}_h\|_{\frac{\Delta t}{\mu}}^2 + \|\mathbf{v}_h\|_{\sigma}^2.$$

Finally, as in AMS, we replace $\|\nabla \mathbf{z}_h\|_{\frac{\Delta t}{\mu}}^2 + \|\mathbf{z}_h\|_{\sigma}^2$ with the variational coarse grid form $\|\nabla \times \Pi_h \mathbf{z}_h\|_{\frac{\Delta t}{\mu}}^2 + \|\Pi_h \mathbf{z}_h\|_{\sigma}^2 = a_{\text{curl}}(\mathbf{z}_h, \mathbf{z}_h)$.

4. AMG Performance on the Reduced Problem. We are now in position to investigate the performance of AMG applied directly to the reduced problem (1.4). The main result of this section can be summarized as follows: *if classical AMG/AMS works for A , then it will also work for S .*

4.1. Schur Complement Properties. The Schur complement inherits many solver-friendly properties from the original matrix, see [43] [40, §3.1] [1, §3.2] [26, §3]. In particular, its condition number is less than or equal to that of A , it can be assembled locally as was pointed out in Section 2.1, and if A is an M-matrix, so is S . The latter is of importance in classical AMG, where theory is typically available only in the M-matrix case. For the definition of M-matrix, see e.g., [1].

We also have the defining energy minimization property

$$(4.1) \quad (\mathbf{S}\mathbf{x}_r, \mathbf{x}_r) = \inf_{\mathbf{x}|_r = \mathbf{x}_r} (\mathbf{A}\mathbf{x}, \mathbf{x}),$$

which implies the diagonal inequality

$$(4.2) \quad D_S \leq D_A$$

where each entry on the left is bounded by the entry on the right corresponding to the same reduced degree of freedom. These properties will be used later in the analysis of AMS in Section 4.3/

Assuming that S requires less memory than A , see Section 2, and based on the above statements, we can generally expect faster matrix-vector multiplies and fewer solver iterations for the reduced problem (1.4) compared to the original problem (1.1). This expectation of course, is provided that AMG is equally suitable for both problems. We investigate this very topic in the following sections.

4.2. Reduced Classical AMG. In this section we prove that the local elimination process based on harmonic extension preserves the near-nullspace of the problem. Since the character of the near-nullspace is central to classical AMG, as discussed in Section 3.1, we conclude that if classical AMG works well on A , it should also work well on S . Therefore, AMG can be applied directly to the reduced problem in the XY case.

LEMMA 4.1. *Consider typical element reduction where the size of the diagonal blocks in A_{ii} is small and independent of the total problem size. Then, the near-nullspace of S is precisely the restriction of the near-nullspace of A to the reduced degrees of freedom. In other words, the reduction process preserves the near-nullspace.*

Proof. First we note that due to the local nature of the elimination and the fact that the reduced elements correspond to small subdomains, the matrix A_{ii}^{-1} is block-diagonal with small local blocks which are spectrally equivalent to their diagonals. For example, these blocks are actually diagonal for the elimination depicted on Figure 2.1. More generally, for elliptic problems in \mathbb{R}^d , a block of A_{ii} can be viewed as the discretization in the interior of the corresponding reduced element with zero Dirichlet boundary conditions. This means that each block is a small version of the original matrix, so if A is $n \times n$ with the typical scaling of $\lambda_{\max}(A) \sim 1$ and $\lambda_{\min}(A) \sim n^{-\frac{2}{d}}$, then for blocks of size m , $\lambda_{\max}(A_{ii}) \sim 1$ and $\lambda_{\min}(A_{ii}) \sim m^{-\frac{2}{d}} \sim 1$.

Therefore,

$$\|A_{ii}^{-1}A_{ir}\| \leq \frac{\lambda_{\max}(A)}{\lambda_{\min}(A_{ii})} \leq C,$$

where C is a constant depending on the particular choice of reduced elements, but independent of the global problem size. This implies that $\|e_r\|$ and $\|Pe_r\|$ are comparable in the sense that


$$\|e_r\| \leq \|Pe_r\| \lesssim \|e_r\|.$$

Now, suppose that e is in the near-nullspace of A , i.e. $\|e\| = 1$ and $Ae \approx 0$. Then,

$$A_{ii}e_i + A_{ir}e_r \approx 0$$

implies $e \approx Pe_r$, so

$$Se_r = P^t A Pe_r \approx P^t Ae \approx 0.$$


 FIG. 4.1. *Stretched elements in 2D (left) and 3D (right).*

Because $\|e\|$ and $\|e_r\|$ are comparable, we conclude that e_r is in the near-nullspace of S .

On the other hand, let e_r be a given vector in the near-nullspace of S . Then, with $e = Pe_r$ we have

$$0 \approx (Se_r, e_r) = (Ae, e).$$

Since A is SPD, this implies $Ae \approx 0$, which together with $\|e\| \approx \|e_r\|$ means that e is in the near-nullspace of A . \square

4.2.1. Meshes with Stretched Elements. Meshes with stretched elements are a common occurrence in the motivating applications. In 2D the reduction process from Section 2.2.1 will eliminate badly shaped triangles. In 3D the improvement is only marginal, see Figure 4.1.

To further investigate the behavior of classical AMG in this case, we perform a stencil analysis for a model 2D Laplacian, where we set $\mu = \Delta t$ and $\sigma = 0$ in (2.2), and we discretize using the P_1 finite element space on a Cartesian grid with initial quadrilateral elements having aspect ratios $1/\varepsilon$. We take the aspect ratio parameter ε to be small.

After carrying through the local elimination on the triangular grid, we get a 9-point reduced stencil at the vertices of the original quadrilaterals which is proportional to

$$\begin{array}{|c|c|c|} \hline -1 & -6 & -1 \\ \hline 2 & 12 & 2 \\ \hline -1 & -6 & -1 \\ \hline \end{array} + \varepsilon^2 \begin{array}{|c|c|c|} \hline -2 & -4 & -2 \\ \hline -4 & 24 & -4 \\ \hline -2 & -4 & -2 \\ \hline \end{array} + \varepsilon^4 \begin{array}{|c|c|c|} \hline -1 & 2 & -1 \\ \hline -6 & 12 & -6 \\ \hline -1 & 2 & -1 \\ \hline \end{array}.$$

Therefore, for large stretching ($\varepsilon \sim 0$), we have

$$\text{asymptotic reduced } P_1 \text{ stencil} \sim \begin{array}{|c|c|c|} \hline -1 & -6 & -1 \\ \hline 2 & 12 & 2 \\ \hline -1 & -6 & -1 \\ \hline \end{array}.$$

This is in contrast to the direct discretization with Q_1 elements, where we get a *different 9-point stencil* at the same points:

$$\text{asymptotic } Q_1 \text{ stencil} \sim \begin{array}{|c|c|c|} \hline -1 & -4 & -1 \\ \hline 2 & 8 & 2 \\ \hline -1 & -4 & -1 \\ \hline \end{array}.$$

Surprisingly, the reduced P_1 stencil on the quadrilateral grid is better for AMG than the direct Q_1 discretization. Indeed, when $\varepsilon \sim 0$, AMG needs to detect strong vertical dependence in order to perform the semi-coarsening needed for the robust solution of this problem, see e.g., [13]. Using the strength threshold parameter θ from Section 3.1, we see that for $\theta = 0.25$, which is the default value in BoomerAMG, the reduced stencil considers only the vertical connections to be strong, while the Q_1 stencil detects also strong dependence in the diagonal directions. To summarize, *the*

artificial introduction and elimination of the interior quadrilateral unknowns leads to a more robust discretization with respect to classical AMG.

In fact, the reduced stencil will produce the correct AMG coarsening for any $\theta > 1/6$, when the center of the stencil is in the interior of the computational domain. Nodes on the Neumann part of the domain boundary have a stencil which leads to additional restrictions on θ . For example, the reduced stencil needs $\theta > 1/3$ in order to produce the correct AMG coarsening near the Neumann boundary. The Q_1 discretization requires $\theta > 1/2$ in this case.

The original XY stencils are also problematic for AMG, because they are asymptotically proportional to

$$\begin{array}{|c|c|c|} \hline -1 & & -1 \\ \hline & 4 & \\ \hline -1 & & -1 \\ \hline \end{array} \quad \text{and} \quad \begin{array}{|c|c|c|c|c|} \hline & & -1 & & \\ \hline & -1 & & -1 & \\ \hline 1 & & 4 & & 1 \\ \hline & -1 & & -1 & \\ \hline & & -1 & & \\ \hline \end{array},$$

where the stencil on the left corresponds to a node in the center of a quadrilateral element, and the right stencil is obtained at the vertices of the mesh. Thus, the Schur complement reduction process has the unexpected benefit of improving the AMG “strength of connection”-related properties of the matrix on stretched grids.

Similar analysis can also be performed for the subspace Poisson problems in RZ , with the difference that θ in this case should stay bounded away from one, in order to preserve the strong vertical dependence near the axis of rotation. Therefore, we can expect an improved AMG performance on 2D reduced stretched grids.

4.3. Reduced AMS. In this section we show how AMS can be applied to the reduced matrices corresponding to (2.3) and (2.1), and prove an appropriate reduced version of the HX decomposition in matrix form. This allows us to conclude that if AMS works well on A , it will also work well on S , and therefore we can apply it directly to the reduced problem in the RZ and $3D$ cases.

First, we note that due to the element reduction nature of the elimination process, the reduction of edge degrees of freedom implies a reduction in the nodal degrees of freedom. Specifically, the reduced vertices are simply the vertices of the reduced edges. With this definition, the discrete gradient matrix can be partitioned as follows:

$$G = \begin{pmatrix} G_{ii} & G_{ir} \\ 0 & G_{rr} \end{pmatrix},$$

where the lower left block is zero, since by definition the reduced edges are not connected to the interior vertices.

Similarly, we have

$$\Pi = \begin{pmatrix} \Pi_{ii} & \Pi_{ir} \\ 0 & \Pi_{rr} \end{pmatrix}$$

with $\Pi_{rr} = [\Pi_{rr,x} \ \Pi_{rr,y} \ \Pi_{rr,z}]$ and $(\Pi_{rr,x})_{ij} = |(G_{rr})_{ij}|(G_{rr,x_r})_i/2$, etc. Therefore, Π_{rr} can still be computed from G_{rr} and the coordinates of the reduced vertices.

By Lemma 4.1, the parts of the near-nullspace of A described by $\text{Ran}(G)$ and $\text{Ran}(\Pi)$ are given by the restriction of the corresponding parts of the near-nullspace of A , which are clearly $\text{Ran}(G_{rr})$ and $\text{Ran}(\Pi_{rr})$. Therefore, G_{rr} and Π_{rr} are the natural reduced versions of the discrete gradient and Nédélec interpolation matrices,

and we can use them, in principle, to apply AMS directly to S . The question is, how well is this going to work? The answer lies in a reduced version of the HX decomposition (3.3), which we consider below.

LEMMA 4.2 (reduced HX decomposition). *Assume that A has a HX decomposition in matrix form as described in Lemma 3.2 (A). Then the same holds for the reduced matrix S . In other words, any reduced vector u_r can be decomposed into*

$$u_r = v_r + G_{rr}p_r + \Pi_{rr}z_r,$$

such that

$$(D_S v_r, v_r) + (S G_{rr} p_r, G_{rr} p_r) + (S \Pi_{rr} z_r, \Pi_{rr} z_r) \lesssim (S u_r, u_r).$$

Proof. Fix the vector u_r and consider its harmonic extension $u = P u_r$. Applying the A -based decomposition from Lemma 3.2 (A), we get

$$u = v + G p + \Pi z,$$

which by restriction gives us the desired decomposition on the reduced degrees of freedom:

$$u_r = v_r + G_{rr} p_r + \Pi_{rr} z_r.$$

To establish the required stability estimates, we use the definition (1.3), the corresponding estimates for the original matrix, and the energy minimization property (4.1):

$$(S u_r, u_r) = (A u, u) \gtrsim (A G p, G p) \geq (S G_{rr} p_r, G_{rr} p_r).$$

Similarly, since Πp is a (non-harmonic) extension of $\Pi_{rr} p_r$ to the interior edges, we get the stability estimate of the Π term from $(A u, u) \gtrsim (A \Pi p, \Pi p) \geq (S \Pi_{rr} p_r, \Pi_{rr} p_r)$. Finally, the remainder v_r is bounded due to the Schur complement diagonal estimate (4.2) from Section 4.1,

$$(S u_r, u_r) \gtrsim (D_A v, v) \geq (D_S v_r, v_r).$$

This completes the proof. \square

We next focus on the second condition in Lemma 3.2, by showing that the subspace problems of the reduced system are close to the Schur complements of the subspace problems in the original system. To this end, we first define a nodal version of the harmonic extension operator and establish a commuting diagram property for certain near-nullspace vectors.

LEMMA 4.3. *Let*

$$P_G = \begin{pmatrix} -(G^T A G)_{ii}^{-1} (G^T A G)_{ir} \\ I \end{pmatrix}$$

be the nodal $G^T A G$ -based harmonic extension operator. Then

$$(4.3) \quad P G_{rr} e_r = G P_G e_r$$

for any vector e which satisfies

$$(A G e)_i = 0.$$

An analogous result holds for P_Π – the nodal $\Pi^T A \Pi$ -based harmonic extension operator.

Proof. The equality (4.3) is trivial on the reduced edge degrees of freedom, so we only need to establish it in the interior edges, where it reads

$$lhs_i := -A_{ii}^{-1} A_{ir} G_{rr} e_r = -G_{ii} (G^T A G)_{ii}^{-1} (G^T A G)_{ir} e_r + G_{ir} e_r =: rhs_i.$$

Since any vector v_i can be uniquely decomposed into gradient and A_{ii} -discretely divergence free parts,

$$v_i = G_{ii} w_i + u_i, \quad \text{where } w_i = (G_{ii}^T A_{ii} G_{ii})^{-1} G_{ii}^T A_{ii} v_i \quad \text{and} \quad G_{ii}^T A_{ii} u_i = 0,$$

then it is enough to show that $(A_{ii}(rhs_i - lhs_i), G_{ii} w_i) = 0$ and $(A_{ii}(rhs_i - lhs_i), u_i) = 0$. The first identity follows from the fact that

$$(G^T A G)_{ii} = G_{ii}^T A_{ii} G_{ii} \quad \text{and} \quad (G^T A G)_{ir} = G_{ii}^T A_{ii} G_{ir} + G_{ii}^T A_{ir} G_{rr},$$

so

$$G_{ii}^T A_{ii} rhs_i = -G_{ii}^T A_{ir} G_{rr} = G_{ii}^T A_{ii} lhs_i.$$

For the equality involving u_i , note that $G_{ii}^T A_{ii} u_i = 0$ and

$$A G = \begin{pmatrix} A_{ii} G_{ii} & A_{ii} G_{ir} + A_{ir} G_{rr} \\ A_{ri} G_{ii} & A_{ri} G_{ir} + A_{rr} G_{rr} \end{pmatrix}$$

imply

$$(A_{ii}(rhs_i - lhs_i), u_i) = ((A G P_{G e_r})_i, u_i) = ((A G e)_i, u_i) = 0.$$

□

We now combine the lemmas into the main result of this section, which ensures that the AMS convergence properties can be transferred to the Schur complement.

THEOREM 4.4. *Suppose that AMS is optimal for A in the sense of Lemma 3.2, i.e. A satisfies conditions (A) and (B) there. Then, the reduced version of AMS based on G_{rr} and Π_{rr} is also optimal for S in the sense of Lemma 3.2.*

Proof. The part regarding condition (A) follows by Lemma 4.2. For condition (B), note that by Lemma 4.3

$$(G_{rr}^T S G_{rr} e_r, e_r) = (G_{rr}^T P^T A P G_{rr} e_r, e_r) = (P_G^T (G^T A G) P_{G e_r}, e_r),$$

so $G_{rr}^T S G_{rr}$ coincides with the (nodal) Schur complement of $G^T A G$, for vectors satisfying $(A G e)_i = 0$. If this is the case for the reduced near-nullspace vectors, e.g., if $G e \approx 0$ or if σ is uniformly small, then we can apply Lemma 4.1 and conclude that classical AMG should work for the above reduced subspace problems. More generally, we can extend the arguments from the Lemma, as shown below.

Let e be in the near-nullspace of $G^T A G$, then by the above considerations, $A G e \approx 0$ implies that e_r is in the near-nullspace of the reduced subspace matrix $G_{rr}^T S G_{rr}$. Conversely, a vector in the near-nullspace of $G_{rr}^T S G_{rr}$ satisfies either $G_{rr} e_r \approx 0$ or $e_{c,r} \approx 0$, where e_c is the restriction to the part of the domain where σ is not small, see Section 4.3.1. In the latter case, zero is extended by zero, so $e_c = (P_{G e_r})_c \approx 0$. If $G_{rr} e_r \approx 0$, then e is approximately a constant on the boundary of each reduced

element, and therefore its nodal harmonic extension will also be close to a constant. We conclude that $\text{GP}_{\text{G}e_r} \approx 0$, so $e = \text{P}_{\text{G}e_r}$ is in the near-nullspace of $\text{G}^T \text{A} \text{G}$.

Similar arguments can be applied to the subspace matrix $\Pi_{rr}^T \text{S} \Pi_{rr}$ which is close to the Schur complement of $\Pi^T \text{A} \Pi$. For example, if we consider the form (3.2), then a vector in the near-nullspace of $\Pi_{rr,x}^T \text{S} \Pi_{rr,x}$ satisfies either $\Pi_{rr,x} e_r \approx 0$ or $e_{c,r} \approx \text{const}$, since $\Pi_x 1 = \text{G}x$, which is the representation of the constant vector $(1, 0, 0)$ in \mathbf{V}_h . The second case can be handled as above, while $\Pi_{rr,x} e_r \approx 0$ means that e_r has a checkerboard pattern on the edges not parallel to the x direction, so $e_r \approx 0$ due to the boundary conditions. On meshes with simplex elements, and on the octahedral grid from Figure 2.4, $\Pi_{rr,x} e_r \approx 0$ directly implies $e_r \approx 0$ locally. \square

4.3.1. Problems with Zero-conductivity Regions. We now consider the case of AMS applied to pure void electromagnetic diffusion problems, where $\sigma = 0$ in the non-conducting part of the domain Ω_{nc} . The original matrix in this case has a non-trivial nullspace

$$\text{Ker}(\text{A}) = \text{Ran}(\text{G}_{nc}),$$

where G_{nc} is the restriction of G to the vertices in the interior of Ω_{nc} . The right-hand side b should satisfy the compatibility condition

$$\text{G}_{nc}^T \text{b} = 0,$$

so that (1.1) has a solution (determined up to a component in the kernel).

One way to stabilize this singular problem is by the introduction of a Lagrange multiplier leading to a saddle-point linear system [11]. We prefer, however, to work with the original system and apply AMS directly to A . This can be done by employing a decomposition in the factor-space $\mathbf{V}_h / \nabla(S_h|_{\Omega_{nc}})$, cf. Corollary 3.2 in [22]. However, when the problem needs to be solved to a very high-accuracy, e.g., relative reduction in the residual norm of order 10^{-14} , we can also employ the following more robust solution procedure:

1. Construct an AMS preconditioner for the non-singular matrix

$$\text{A} + \epsilon \text{G}_{nc} \text{G}_{nc}^T,$$

where $\epsilon > 0$ is a small parameter.

2. Use the above preconditioner in a Krylov iteration for (1.1) with periodic restrictions onto the compatible subspace $\text{Ker}(\text{G}_{nc}^T)$ using the projection operator $\text{I} - \text{G}_{nc}(\text{G}_{nc}^T \text{G}_{nc})^{-1} \text{G}_{nc}^T$.

Note that the above procedure should be used only when necessary, because it is more expensive than regular AMS due to larger sparsity pattern in the non-conducting region, and because it requires additional information from the user, namely a list of the nodes which are interior to the zero-conductivity regions.

Since Step 2 of the solution procedure uses a Krylov iteration for (1.1) we cannot simply apply the element reduction approach only to the matrix $\text{A} + \delta \text{G}_{nc} \text{G}_{nc}^T$ treated by AMS. However, performing the reduction directly on A is also problematic, since A_{ii} may not be invertible. Hence, *we need to modify the elimination process for the “pure void” RZ and 3D cases.* We describe the modified procedure that we use in the numerical experiments below.

Consider a reduced element E in pure void. Then

$$\text{Ker}(\text{A}_{ii,E}) = \text{Ran}(\text{G}_{ii,E}).$$

We therefore replace

$$(4.4) \quad \mathbf{A}_{ii,E} \mathbf{x}_{i,E} + \mathbf{A}_{ir,E} \mathbf{x}_{r,E} = \mathbf{b}_{i,E}$$

with the following problem, which is uniquely solvable for the interior dofs (and allows us to proceed with the local elimination):

$$(4.5) \quad (\mathbf{A}_{ii,E} + \delta \mathbf{G}_{ii,E} (\mathbf{G}_{ii,E})^T) \tilde{\mathbf{x}}_{i,E} + \mathbf{A}_{ir,E} \mathbf{x}_{r,E} = \mathbf{b}_{i,E}.$$

Here, $\delta > 0$ is a small parameter.

LEMMA 4.5. *Assume that \mathbf{b} is compatible. Then the problems (4.4) and (4.5) are equivalent in the sense that $\tilde{\mathbf{x}}_{i,E}$ and $\mathbf{x}_{i,E}$ differ by a component in the nullspace of $\mathbf{A}_{ii,E}$.*

Proof. Since $(\mathbf{G}_{ii,E})^T \mathbf{A}_{ii,E} = 0$ and $(\mathbf{G}_{ii,E})^T \mathbf{b}_{i,E} = 0$, equation (4.4) implies

$$(\mathbf{G}_{ii,E})^T \mathbf{A}_{ir,E} \mathbf{x}_{r,E} = 0.$$

Plugging this into (4.5) gives us

$$(\mathbf{G}_{ii,E})^T \mathbf{G}_{ii,E} (\mathbf{G}_{ii,E})^T \tilde{\mathbf{x}}_{i,E} = 0,$$

so $(\mathbf{G}_{ii,E})^T \tilde{\mathbf{x}}_{i,E} = 0$. This means that $\tilde{\mathbf{x}}_{i,E}$ is also locally discretely divergence free and $\mathbf{A}_{ii,E}(\mathbf{x}_{i,E} - \tilde{\mathbf{x}}_{i,E}) = 0$. \square

In practice, we simply perform the local elimination in (4.5) instead of (4.4) and continue with the pure void solution procedure from the beginning of this section. The addition of $\delta \mathbf{G}_{ii,E} (\mathbf{G}_{ii,E})^T$ is easy to implement, because it simply updates each entry of the dense matrix $\mathbf{A}_{ii,E}$ with δ , $-\delta$ or 2δ . With trivial modifications we can also handle reduced elements E which are only partially in Ω_{nc} .

Let $\mathbf{P}_{\mathbf{A} + \delta \mathbf{G}_{ii} \mathbf{G}_{ii}^T}$ be the harmonic extension corresponding to (4.5). Then in the reduced pure void case we apply AMS to the matrix

$$\mathbf{P}_{\mathbf{A} + \delta \mathbf{G}_{ii} \mathbf{G}_{ii}^T}^T \mathbf{A} \mathbf{P}_{\mathbf{A} + \delta \mathbf{G}_{ii} \mathbf{G}_{ii}^T} + \epsilon \mathbf{G}_{nc,rr} \mathbf{G}_{nc,rr}^T.$$

Since for small ϵ and δ this is a small perturbation of the reduced problem in the factor-space, we expect AMS to work well on this problem.

5. Numerical Results. In this section we present numerical results for the XY , RZ and $3D$ modes of the Maxwell problem, that illustrate the theory developed in the previous sections. In particular, we answer the central question of the paper, by demonstrating that in all cases we can solve larger problems faster, with typical speedups of at least a factor of two.

In the following experiments we use the Conjugate Gradient (CG) Krylov solver with the BoomerAMG and AMS preconditioners from the *hypre* library [20] applied in the XY and $RZ/3D$ cases, respectively. The tests were run on the multi-core cluster Hera at Lawrence Livermore National Laboratory (LLNL).

We used a sweep of symmetric hybrid Gauss-Seidel pre- and post-relaxation in BoomerAMG (applied to both the XY and the subspace $RZ/3D$ problems) and its convergent ℓ_1 version in AMS. BoomerAMG also used the low-complexity HMIS coarsening with an extended long-range interpolation, while AMS employed a multiplicative cycle based on the decomposition (3.2). In the singular case, we set the values of the perturbation parameters as follows: $\epsilon = 10^{-8} \max_i \mathbf{A}_{ii}$ and $\delta = 10^{-12} \max_{km} |(\mathbf{A}_{ii,E})_{km}|$.

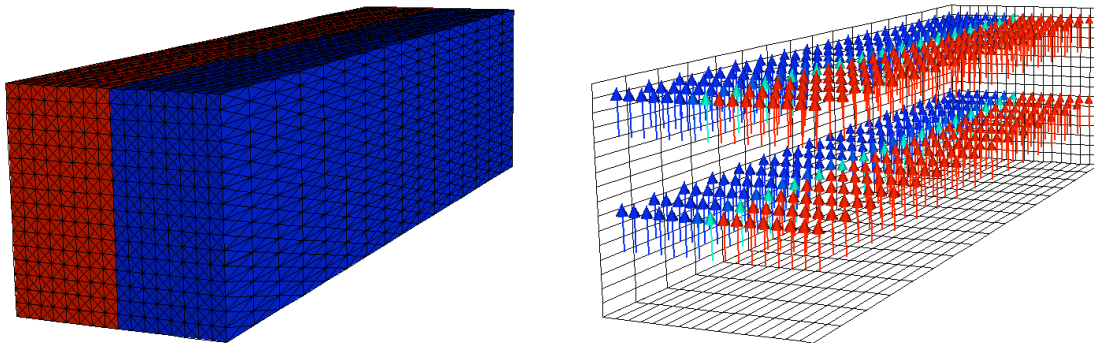


FIG. 5.1. *The Box problem is a simple test problem that we use to test the effects of varying solver parameters, material properties, and aspect ratios. The mesh, with sub-zones, is shown on the left with a high conductivity region shown in red ($\sigma_c = 1$) and a varying low conductivity shown in blue ($0 \leq \sigma_{nc} \leq 1$). The right figure shows the direction of the magnetic field, with the arrows colored by field strength.*

5.1. Parameter Study - The Box Problem. We start with a simple diffusion problem posed on a structured box in $3D$, see Figure 5.1, with the XY and RZ cases corresponding to the front and the top sides, respectively. The box is split in two parts, with the conductivity σ varying between $\sigma_c = 1$ in the material half of the domain and $0 \leq \sigma_{nc} \leq 1$ in the non-conducting (void) half. The mesh is initially uniform, but we stretch it to test the dependence on the aspect ratio. For this problem, we take $\Delta t/\mu = 10^{-3}$ and use convergence tolerance of 10^{-10} in CG.

The properties of the original and the reduced matrices are given in Table 5.1. Note that we get problem size and memory reduction in accordance with the estimates in Table 2.1 and Table 2.2.

TABLE 5.1

Problem size and number of nonzero entries in the original/reduced matrix for the three different modes of the Box problem. The reduction in the number of nonzeros corresponds to the reduction in memory requirements.

mode	n_{rows} (reduction)	n_{nz} (reduction)
XY	33025/16641 ($\times 2.0$)	230145/148225 ($\times 1.6$)
RZ	98560/33024 ($\times 3.0$)	491776/229632 ($\times 2.1$)
$3D$	239260/90460 ($\times 2.6$)	3724060/2658460 ($\times 1.4$)

The Box problem is a simplified serial test to explore the parameter space for the aspect ratio ε and the strength threshold θ . We also performed tests to confirm that, in all cases, the reduced solver is insensitive to varying the conductivity ratio σ_{nc}/σ_c in the range $[0, 1)$. In the interest of brevity, and since the results were uniform, we only report one of these tests in Table 5.7.

We start by running the problems with several different values for the strength threshold parameter θ for both the reduced and full problems and record the number of AMG-CG iterations needed. The results in Table 5.2 confirm the stencil analysis from Section 4.2.1; due to the Neumann boundary conditions in XY , the critical value for the reduced matrix here is $\theta > 1/3$. Note that when θ is above this threshold, convergence becomes stable, and for $\theta = 0.4$ the number of AMG-CG iterations is independent of the element aspect ratios. The critical value for the full problem is somewhat higher.

TABLE 5.2

Solver iteration dependence on the strength threshold θ for the reduced Box problem in XY mode with conductivity ratio $\sigma_{nc}/\sigma_c = 0$ and varying aspect ratio ε is reported in the original/reduced format.

$1/\varepsilon$	1	4	16	64	256	1024	4096
$\theta = 0.33$	7/ 7	12/ 8	16/20	35/42	45/57	46/58	46/58
$\theta = 0.34$	7/ 7	12/ 8	16/12	35/ 9	45/ 7	46/ 7	46/ 7
$\theta = 0.40$	7/ 7	12/ 8	16/ 8	35/ 7	45/ 6	46/ 7	46/ 7
$\theta = 0.50$	10/10	12/ 8	12/ 8	11/ 7	11/ 6	11/ 7	11/ 7

Next we explore the run-time behavior in XY with respect to increasing aspect ratio $1/\varepsilon$ in Table 5.3. Based on Table 5.2, we choose a value of $\theta = 0.34$ for the reduced matrix and $\theta = 0.5$ for the full matrix to minimize run time for each. We set $\sigma_{nc}/\sigma_c = 0$ and report the number of CG iterations (n_{it}) as well as the combined time spent in the solver setup and solution phases (t_{solver}). The remaining time, including the matrix assembly, as well as the elimination and the recovery of the internal degrees of freedom is denoted by ($t_{assemble}$). For all of these quantities we present the data for both problems (1.1) and (1.4) in the format “original/reduced”. Finally, we compute and report the total run time speedup of the computational cycle due to the reduction.

TABLE 5.3

Comparison of overall solution times for the Box problem in XY mode with strength threshold $\theta = 0.34$, conductivity ratio $\sigma_{nc}/\sigma_c = 0$ and varying aspect ratio ε . The number of CG iterations (n_{it}), the solution time (t_{solver}) and the remaining simulation time ($t_{assemble}$) are reported in the original/reduced format.

$1/\varepsilon$	n_{it}	$t_{assemble}$	t_{solver}	speedup
1	10/ 10	0.29/0.23	0.46/0.25	$\times 1.6$
4	12/ 8	0.28/0.21	0.38/0.14	$\times 1.8$
16	12/ 8	0.28/0.22	0.36/0.15	$\times 1.8$
64	11/ 7	0.28/0.22	0.34/0.14	$\times 1.7$
256	11/ 6	0.27/0.22	0.33/0.12	$\times 1.8$
1024	11/ 7	0.28/0.23	0.33/0.15	$\times 1.6$
4096	11/ 7	0.28/0.23	0.33/0.15	$\times 1.6$

From Table 5.3 we see that the AMG-CG solver performs better on the reduced problem, both in terms of number of iterations and time. Even when the number of iterations is the same ($\varepsilon = 1$) there is still a factor of 1.6 speedup (1.8 in the solver). The speedup factor is nearly constant for all aspect ratios. Note also the interesting fact that, even with the extra work of inverting the local A_{ii} and the recovery of x_i , the assemble time in Table 5.3 is always less for the reduced problem. This is a trend in all of our results that we discuss in the following section.

Next, we consider similar tests in the RZ case, where $\sigma_{nc}/\sigma_c = 0$ and we use the the pure void solution procedure from Section 4.3.1. In other words, *we apply AMS directly to the Schur complement of a singular matrix with a large kernel.*

We perform the RZ strength threshold analysis in Table 5.4. Since this problem has only Dirichlet boundary conditions, the critical value in the subspace problems according to Section 4.2.1 is $\theta > 1/6$. This is clearly confirmed by the significant change in the iteration counts shown in the table. Though the convergence is not

uniform, the dependence on ε is significantly reduced when θ is above the critical value. Unlike in the XY case, the iteration counts for the original problem remain high, even for large values of θ . This suggests that the reduced problem has a wider range of good strength threshold values, making it easier to pick a reasonable choice for a wider range of problems. When θ is too large, both methods suffer since strong vertical dependence is lost near the axis of rotation.

TABLE 5.4

Solver iteration dependence on the strength threshold θ for the reduced Box problem in RZ mode with conductivity ratio $\sigma_{nc}/\sigma_c = 0$ and varying aspect ratio ε is reported in the original/reduced format.

$1/\varepsilon$	1	4	16	64	256	1024	4096
$\theta = 0.16$	8/ 8	10/ 8	28/22	75/14	205/24	520/234	835/501
$\theta = 0.17$	8/ 8	10/ 8	28/ 9	74/14	211/24	500/ 30	828/ 33
$\theta = 0.50$	9/ 8	9/ 8	27/ 8	81/ 9	200/15	490/ 26	844/ 27
$\theta = 0.70$	10/10	10/ 8	28/ 7	84/ 7	216/14	594/ 22	694/ 21
$\theta = 0.90$	11/11	10/10	27/10	85/10	246/16	524/ 23	951/ 22

The timing results, presented in Table 5.5, show that the convergence deteriorates due to vanishing coefficients close to the axis of rotation, but overall the reduced AMS solver significantly outperforms the solver applied directly to A. In particular, for the problem with the worse aspect ratio, we get *more than a factor of 45 speedup in the reduced solver leading to more than 39 times total simulation speedup.*

TABLE 5.5

Comparison of overall solution times for the Box problem in RZ mode with strength threshold $\theta = 0.7$, conductivity ratio $\sigma_{nc}/\sigma_c = 0$ and varying aspect ratio ε . The number of CG iterations (n_{it}), the solution time (t_{solver}) and the remaining simulation time ($t_{assemble}$) are reported in the original/reduced format.

$1/\varepsilon$	n_{it}	$t_{assemble}$	t_{solver}	speedup
1	10/10	1.48/0.84	13.5/3.91	$\times 3.2$
4	10/ 8	1.77/0.75	13.7/3.00	$\times 4.1$
16	28/ 7	1.80/0.80	32.4/2.89	$\times 9.2$
64	84/ 7	1.49/0.75	69.5/2.51	$\times 21.7$
256	216/14	1.57/0.76	194./4.43	$\times 37.6$
1024	594/22	1.74/0.74	451./6.25	$\times 64.7$
4096	694/21	0.94/0.75	252./5.61	$\times 39.8$

Finally, we consider the Box problem tests in 3D. In Table 5.6 we investigate both regular AMS for conductivity jump of four orders of magnitude, as well as the robust AMS version for the pure void case. We note that there is a little difference between these cases in terms of solver performance (except that the pure void solver is a bit slower). This trend is typical for all the experiments we have run. Looking at the iteration counts in Table 5.6, we see that the convergence deteriorates significantly on stretched grids. The performance is practically uniform in θ , as the theory from Section 4.2.1 does not apply in this case. Still, there is a significant improvement due to the reduction, with speedup factors between 2 and 4.

In Table 5.7 we investigate the dependence of AMS on the magnitude of the conductivity jump. Note that regular AMS is used in all cases, except $\sigma_{nc} = 0$, when

TABLE 5.6

Comparison of overall solution times for the Box problem in 3D mode with strength threshold $\theta = 0.5$, two sets of conductivity ratios and varying aspect ratio ε . The number of CG iterations (n_{it}), the solution time (t_{solver}) and the remaining simulation time ($t_{assemble}$) are reported in the original/reduced format.

$1/\varepsilon$	n_{it}	$t_{assemble}$	t_{solver}	speedup
$\sigma_{nc}/\sigma_c = 10^{-4}$				
1	9/ 8	6.58/5.04	40.3/17.3	$\times 2.1$
2	9/ 8	7.34/5.14	47.6/16.1	$\times 2.6$
4	16/ 9	7.10/5.07	67.6/16.5	$\times 3.5$
8	29/ 15	7.71/5.15	111./23.8	$\times 4.1$
16	49/ 26	7.40/5.15	178./37.1	$\times 4.4$
32	79/ 42	8.15/5.11	262./55.1	$\times 4.5$
64	121/ 66	7.83/4.95	372./85.1	$\times 4.2$
128	180/107	6.66/5.23	546./138.	$\times 3.8$
$\sigma_{nc}/\sigma_c = 0$				
1	9/ 8	6.30/5.39	53.8/25.1	$\times 2.0$
2	9/ 8	5.88/5.32	51.8/23.4	$\times 2.0$
4	16/ 9	6.15/5.26	72.6/24.8	$\times 2.6$
8	29/ 15	5.95/5.23	117./32.5	$\times 3.3$
16	50/ 26	6.41/5.30	190./54.3	$\times 3.3$
32	79/ 42	6.32/5.31	283./79.2	$\times 3.4$
64	122/ 66	6.02/5.29	440./122.	$\times 3.5$
128	177/103	6.60/5.30	657./187.	$\times 3.4$

we employ the pure void version from Section 4.3.1. The results clearly show little dependence on $\varepsilon \rightarrow 0$ and on $\sigma_{nc}/\sigma_c \rightarrow 0$. Since the robustness with respect to conductivity jumps is a main property of AMS for the original problem, the fact that it gets transferred to the reduced problem is another confirmation of Theorem 4.4.

TABLE 5.7

Solver iteration dependence on the conductivity ratio $\sigma_{nc}/\sigma_c = 0$ for the reduced Box problem in 3D mode with strength threshold $\theta = 0.5$ and varying aspect ratio ε .

$1/\varepsilon$	σ_{nc}/σ_c					
	1	10^{-2}	10^{-4}	10^{-6}	10^{-8}	0
1	8	8	8	8	8	8
2	8	8	8	8	8	8
4	9	9	9	9	9	9
8	15	15	15	15	15	15
16	26	26	26	26	26	26
32	41	42	42	42	42	42
64	59	66	66	66	66	66
128	78	102	107	103	105	103

5.2. Scalability Study - The Coax Problem. We next consider a more realistic problem representing four coaxial cylindrical conductors with varying conductivity

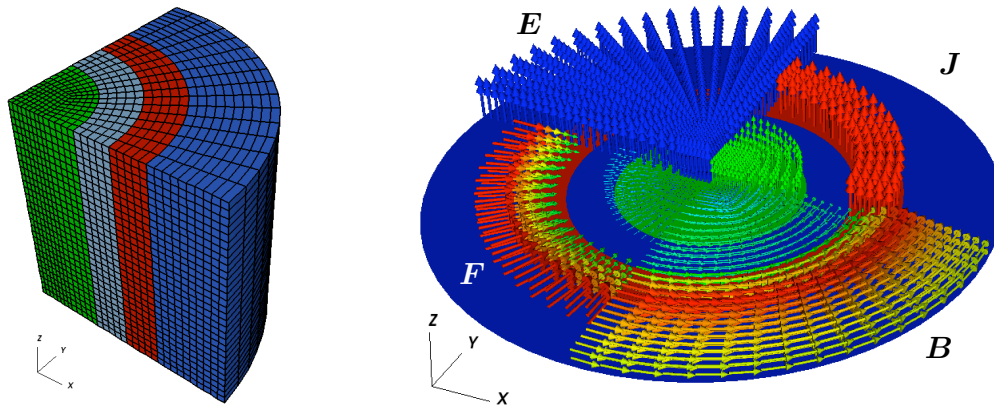


FIG. 5.2. An idealized test problem mocks up the conductivity jumps seen in Z-pinch simulations with four regions of varying conductivity (left). The XY simulations use the top plane, while in RZ the simulation is cut in XZ plane. The right figure shows the magnitude and direction of the electric field \mathbf{E} , the induced magnetic field \mathbf{B} , the current density \mathbf{J} , and the force \mathbf{F} .

that mock up the variations seen in Z-pinch simulations. The domain is a quarter of four concentric cylinders with different conductivities $\sigma = \{10^{-2}, 10^{-8}, 10^{-2}, 0\}$ from the inside out. The mesh and an approximate solution are shown in Figure 5.2. The XY and RZ cases correspond to the top and front sides of the 3D domain. Note that the jumps in σ and the pure void outer region make this problem's (near-)nullspace very challenging, and AMS is required for its robust solution. In these experiments we used the pure void solution procedure from Section 4.3.1, but we emphasize that qualitatively similar results are obtained with regular AMS when the pure void is replaced by a small number.

In fact, the value of supporting pure void regions, with $\sigma = 0$, was highlighted in this problem, since in practice, picking a robust small value for σ is dependent on the mesh resolution. This fact makes it difficult for users of our target application to reliably run problems requiring mesh resolution studies, and therefore despite the extra cost, the pure void solver option is preferable for its robustness.

In this test, we set $\theta = 0.5$, $\varepsilon = 1$, $\Delta t/\mu \sim 10^{-4}$ and use an AMS-CG convergence tolerance of 10^{-10} . We perform a weak scalability test, increasing the mesh refinement by a factor of two each step, and increasing the number of processors proportional to the total number of elements in the problem. Our goal is not to show the full scalability of the methods but rather to demonstrate their relative performance on large problems.

We report the number of processors used (n_p), total problem size (n_{rows}), the CG iteration counts, the assembly time t_{assemble} , the AMS setup time t_{setup} , and the AMS-CG solve time t_{solve} as well as the total simulation speedup. The results for all three geometries (XY, RZ, and 3D) are shown in Table 5.8.

The results for the Coax problem are comparable to the Box problem in the previous section; the problem sizes, number of nonzero matrix entries, iteration counts, times, and overall speedup all follow the same trends. The matrix reduction is robust and remains beneficial even for reasonably large number of processors and problem sizes, with an overall speedup of around a factor of two. For example, the reduced version of the original 3D problem on 512 processors with more than 100 million unknowns required 1.4 times less memory and was solved 2.5 times faster.

Overall, the iteration counts for the reduced and the original problems are comparable, but the reduced solver is always faster, especially in 3D. The solver setup

TABLE 5.8

The matrix sizes, iteration counts, and timings for the Coax problem in all three geometries (XY, RZ, and 3D) show that the benefits of the matrix reduction extend to more realistic problems. The results are reported in the original/reduced format.

n_p	n_{rows}	n_{it}	t_{assemble}	t_{setup}	t_{solve}	speedup
<i>XY results</i>						
1	15013/7589	13/10	0.13/0.11	0.07/0.03	0.21/0.08	$\times 1.8$
4	59721/30025	14/10	0.14/0.13	0.09/0.05	0.23/0.09	$\times 1.6$
16	238225/119441	15/13	0.15/0.12	0.12/0.08	0.36/0.16	$\times 1.7$
64	951585/476449	17/15	0.21/0.14	0.35/0.19	0.79/0.29	$\times 2.1$
<i>RZ results</i>						
1	21720/7320	10/11	0.20/0.12	0.21/0.08	0.57/0.27	$\times 2.1$
4	86640/29040	10/12	0.20/0.11	0.34/0.15	0.84/0.42	$\times 2.0$
16	346080/115680	11/13	0.22/0.13	0.52/0.26	1.42/0.62	$\times 2.1$
64	1383360/461760	12/13	0.23/0.14	1.05/0.63	2.13/1.17	$\times 2.0$
<i>3D results</i>						
1	208370/78774	12/10	3.66/2.59	10.0/3.20	23.1/6.83	$\times 2.9$
8	1640728/621224	13/10	4.08/2.80	32.4/6.95	53.2/10.7	$\times 4.4$
64	13021568/4934592	14/11	4.37/3.00	73.1/16.9	89.7/20.6	$\times 4.1$
512	103756864/39337280	15/13	4.53/3.22	122./41.8	149./66.5	$\times 2.5$

time, which accounts for the calculation of the restriction and prolongation operators and the coarse level matrices in the subspace AMG solvers, is also faster, due to the reduced problem size and number of matrix entries (which also has beneficial effect on the cost of the coarse grid matrix construction). Interestingly, the problem assembly time is also faster, despite the additional computations. This effect is especially striking in *RZ*, and can be possibly explained by better utilization of the multi-core architecture we used to run the test problem.

6. Conclusions. The goal of this paper was to characterize the behavior of algebraic multigrid solvers applied to Schur complements of matrices arising in scalar and electromagnetic diffusion applications.

We showed that in the particular sub-zonal discretization of interest, the element reduction approach leads to a smaller dimensional matrix with fewer nonzeros than the original problem. In the specific application that we considered in Section 2.2, the typical memory reduction factors were: 1.6 for *XY*, 2.1 for *RZ* and 1.4 for *3D*. The corresponding average problem size reduction factors were: 2 for *XY*, 3 for *RZ* and 2.6 for *3D*. The *3D* case was particularly interesting, because it resulted in algebraically constructed octahedral finite elements.

The good performance of the BoomerAMG and AMS solvers on the reduced problem was shown to follow from a nullspace reduction property in the classical AMG case and the reduced HX decomposition proved in Lemma 4.2 in the AMS case. Stencil analysis on stretched grids indicated that improved Q_1 discretization can be obtained through the reduction process. The fact that the already existing AMG and AMS codes can be applied directly to the matrix S is a testament to the power of algebraic solution methods.

The numerical results in Section 5 demonstrated that, even with the elimination

and the recovery overhead, our new approach solves most of the problems at least twice as fast, with no dependence on the jumps in σ . In addition, we showed how the reduction can be easily modified to handle the pure void case. The obtained speedup factors in the considered simulations were 1.6–2.1 for XY , 2.0–39.8 for RZ and 2.0–4.4 for $3D$. These speedups have been observed to hold for other, more realistic user problems as well.

A possible extension of this work is to consider algebraic approaches for determining reduced elements with desirable properties in unstructured settings. Our experience with the sub-zonal discretization suggests that a good starting point for electromagnetic problems is to define the reduced elements as the neighborhoods of the vertices contained in minimal number of edges (cf. Figure 2.3 and Figure 2.4).

Acknowledgments. The authors would like to acknowledge the help of Panayot Vassilevski, Robert Falgout and Tom Gardiner.

REFERENCES

- [1] O. AXELSSON, *Iterative Solution Methods*, Cambridge University Press, 1996.
- [2] A. BAKER, T. GAMBLIN, M. SCHULZ, AND U. M. YANG, *Challenges of algebraic multigrid across multicore architectures*, Tech. Report LLNL-CONF-427763, Lawrence Livermore National Laboratory, 2010.
- [3] P. BOCHEV, C. GARASI, J. HU, A. ROBINSON, AND R. TUMINARO, *An improved algebraic multigrid method for solving Maxwell's equations*, SIAM J. Sci. Comp., 25 (2003), pp. 623–642.
- [4] P. BOCHEV, J. HU, C. SIEFERT, AND R. TUMINARO, *An algebraic multigrid approach based on a compatible gauge reformulation of Maxwell's equations*, SIAM J. Sci. Comp., 31 (2008), pp. 557–583.
- [5] P. BOCHEV AND A. ROBINSON, *Matching algorithms with physics: exact sequences of finite element spaces*, in Preservation of stability under discretization, D. Estep and S. Tavener, eds., Philadelphia, 2001, Colorado State University, SIAM, pp. 145–165.
- [6] A. BRANDT, *Multiscale solvers and systematic upscaling in computational physics*, Comp. Phys. Comm., 169 (2005), pp. 438–441. Proceedings of the Europhysics Conference on Computational Physics 2004.
- [7] M. BREZINA, A. CLEARY, R. FALGOUT, V. HENSON, J. JONES, T. MANTEUFFEL, S. MCCORMICK, AND J. RUGE, *Algebraic multigrid based on element interpolation (AMGe)*, SIAM J. Sci. Comp., 22 (2000), pp. 1570–1592.
- [8] F. BREZZI, K. LIPNIKOV, M. SHASHKOV, AND V. SIMONCINI, *A new discretization methodology for diffusion problems on generalized polyhedral meshes*, Comput. Meth. Appl. Mech. Eng., 196 (2007), pp. 3682–3692. Special Issue Honoring the 80th Birthday of Professor Ivo Babuska.
- [9] T. BRUNNER AND T. GARDINER, *Status of the magnetic physics in KULL*. Presentation given Sandia National Laboratory, July 2009.
- [10] E. CARAMANA, D. BURTON, M. SHASHKOV, AND P. WHALEN, *The construction of compatible hydrodynamics algorithms utilizing conservation of total energy*, Jour. Comp. Phys., 146 (1998), pp. 227–262.
- [11] L. DEMKOWICZ AND L. VARDAPETYAN, *Modeling of electromagnetic absorption/scattering problems using hp-adaptive finite elements*, Comp. Meth. Appl. Mech. Eng., 152 (1998), pp. 103–124.
- [12] H. ELMAN AND G. GOLUB, *Iterative methods for cyclically reduced non-self-adjoint linear systems*, Math. Comp., 54 (1990), pp. 671–700.
- [13] R. FALGOUT, *An introduction to algebraic multigrid*, Comput. Sci. Eng., 8 (2006), pp. 24–33. Special issue on Multigrid Computing.
- [14] M. GEE, C. SIEFERT, J. HU, R. TUMINARO, AND M. SALA, *ML 5.0 smoothed aggregation user's guide*, Tech. Report SAND2006-2649, Sandia National Laboratories, 2006.
- [15] A. GLASSER, *A scalable parallel extended MHD solver: Application of physics-based preconditioning to high-order spectral elements*. Talk at the 11th Copper Mountain Conference on Iterative Methods, 2010.
- [16] G. GOLUB AND R. TUMINARO, *Cyclic reduction/multigrid*, Tech. Report NA-92-14, Stanford University, Stanford, CA, USA, 1992.

- [17] R. HARTMANN, M. LUKÁCOVÁ-MEDVID'OVÁ, AND F. PRILL, *Efficient preconditioning for the discontinuous Galerkin finite element method by low-order elements*, Appl. Numer. Math., 59 (2009), pp. 1737–1753.
- [18] V. HENSON AND U. YANG, *BoomerAMG: A parallel algebraic multigrid solver and preconditioner*, Appl. Numer. Math., 41 (2002), pp. 155–177.
- [19] R. HIPTMAIR AND J. XU, *Nodal auxiliary space preconditioning in $H(\text{curl})$ and $H(\text{div})$ spaces*, SIAM J. Numer. Anal., 45 (2007), pp. 2483–2509.
- [20] *hypre: High performance preconditioners*. <http://www.llnl.gov/CASC/hypre/>.
- [21] J. JONES AND P. VASSILEVSKI, *AMGe based on element agglomeration*, SIAM J. Sci. Comp., 23 (2001), pp. 109–133.
- [22] Tz. KOLEV AND P. VASSILEVSKI, *Parallel auxiliary space AMG for $H(\text{curl})$ problems*, J. Comput. Math., 27 (2009), pp. 604–623. Special issue on Adaptive and Multilevel Methods in Electromagnetics.
- [23] S. MACLACHLAN, T. MANTEUFFEL, AND S. MCCORMICK, *Adaptive reduction-based AMG*, Numer. Linear Algebra Appl., 13 (2006), pp. 599–620.
- [24] S. MACLACHLAN, J. TANG, AND C. VUIK, *Fast and robust solvers for pressure-correction in bubbly flow problems*, J. Comput. Phys., 227 (2008), pp. 9742–9761.
- [25] J. MANDEL, M. BREZINA, AND P. VANĚK, *Energy optimization of algebraic multigrid bases*, Computing, 62 (1999), pp. 205–228.
- [26] T. P. A. MATHEW, *Domain Decomposition Methods for the Numerical Solution of Partial Differential Equations*, vol. 61 of Lecture Notes in Computational Science and Engineering, Springer, 2008.
- [27] W. MITCHELL, *The hp-multigrid method applied to hp-adaptive refinement of triangular grids*, Numer. Linear Algebra Appl., 17 (2010), pp. 211–228.
- [28] P. MONK, *Finite Element Methods for Maxwell's Equations*, Numerical Mathematics and Scientific Computation, Oxford University Press, Oxford, UK, 2003.
- [29] J.-C. NÉDÉLEC, *Mixed finite elements in \mathbb{R}^3* , Numer. Math., 35 (1980), pp. 315–341.
- [30] J. PASCIAK AND J. ZHAO, *Overlapping schwarz methods in $H(\text{curl})$ on nonconvex domains*, East West J. Num. Anal., 10 (2002), pp. 221–234.
- [31] A. QUARTERONI AND A. VALLI, *Domain Decomposition Methods for Partial Differential Equations*, Oxford Science Publications, 1999.
- [32] J. RATHKOPF, D. MILLER, J. OWEN, L. STUART, M. ZIKA, P. ELTGROTH, N. MADSEN, K. MCCANDLESS, P. NOWAK, M. NEMANIC, N. GENTILE, N. KEEN, AND T. PALMER, *KULL: LLNL's ASCI inertial confinement fusion simulation code*, in PHYSOR 2000, ANS International Topical Meeting on the Advances in Reactor Physics and Mathematics and Computation into the Next Millennium, American Nuclear Society, May 2000.
- [33] S. REITZINGER AND J. SCHÖBERL, *An algebraic multigrid method for finite element discretizations with edge elements*, Numer. Linear Algebra Appl., 9 (2002), pp. 223–238.
- [34] J. RUGE. private communications.
- [35] J. RUGE AND K. STÜBEN, *Algebraic multigrid (AMG)*, in Multigrid Methods, S. F. McCormick, ed., vol. 3 of Frontiers in Applied Mathematics, SIAM, Philadelphia, PA, 1987, pp. 73–130.
- [36] P. ŠOLÍN, K. SEGETH, AND I. DOLEŽEL, *Higher-Order Finite Element Methods*, Studies in Advanced mathematics, Chapman & Hall/CRC, 2004.
- [37] K. STÜBEN, *Solving reservoir simulation equations*, (2007). 9th International Forum on Reservoir Simulation, December 9-13, Abu Dhabi, United Arab Emirates.
- [38] 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. International GAMM-Workshop on Multi-level Methods (Meisdorf, 1994).
- [39] P. VASSILEVSKI, *Multilevel preconditioners for elliptic problems by substructuring*, Appl. Math. Comput., 46 (1991), pp. 79–104.
- [40] P. VASSILEVSKI, *Multilevel block factorization preconditioners: Matrix-based analysis and algorithms for solving finite element equations*, Springer, New York, 2008.
- [41] E. WILSON, *The static condensation algorithm*, Int. J. Numer. Methods Eng., 8 (1974), pp. 199–203.
- [42] I. YAVNEH, *Cyclic reduction multigrid revisited*. Talk at the 14th Copper Mountain Conference on Multigrid Methods, 2009.
- [43] F. ZHANG, ed., *The Schur Complement and Its Applications*, vol. 4 of Numerical Methods and Algorithms, Springer, 2005.
- [44] L. ZIKATANOV, *Two-sided bounds on the convergence rate of two-level methods*, Numer. Linear Algebra Appl., 15 (2008), pp. 439–454.