SAND2022-12164R

SAND2022-XXXXR
LDRD PROJECT NUMBER: 226833

LDRD PROJECT TITLE: Composing preconditioners for multiphysics PDE systems with applications to Generalized MHD

PROJECT TEAM MEMBERS: Raymond S. Tuminaro, Michael M. Crockatt, Allen C. Robinson

ABSTRACT:

New patch smoothers or relaxation techniques are developed for solving linear matrix equations coming from systems of discretized partial differential equations (PDEs). One key linear solver challenge for many PDE systems arises when the resulting discretization matrix has a near null space that has a large dimension, which can occur in generalized magnetohydrodynamic (GMHD) systems. Patch-based relaxation is highly effective for problems when the null space can be spanned by a basis of locally supported vectors. The patch-based relaxation methods that we develop can be used either within an algebraic multigrid (AMG) hierarchy or as stand-alone preconditioners. These patch-based relaxation techniques are a form of well-known overlapping Schwarz methods where the computational domain is covered with a series of overlapping sub-domains (or patches). Patch relaxation then corresponds to solving a set of independent linear systems associated with each patch. In the context of GMHD, we also reformulate the underlying discrete representation used to generate a suitable set of matrix equations. In general, deriving a discretization that accurately approximates the curl operator and the Hall term while also producing linear systems with physically meaningful near null space properties can be challenging. Unfortunately, many natural discretization choices lead to a near null space that includes non-physical oscillatory modes and where it is not possible to span the near null space with a minimal set of locally supported basis vectors. Further discretization research is needed to understand the resulting trade-offs between accuracy, stability, and ease in solving the associated linear systems.

INTRODUCTION AND EXECUTIVE SUMMARY OF RESULTS:

The presence of a large near null space presents unique challenges when solving linear systems via iterative methods. This is primarily due to the fact that the norm of the residual associated with an approximate solution might be small even though the norm of the error in the approximate solution is large. The development of the smoothed aggregation AMG method was one of the earliest linear solver approaches to highlight the challenges of the near null space [16, 20, 21]. Specifically, a

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

technique was presented for developing interpolation operators that accurately represent the near null space. This early work emphasized solving linear elasticity systems where the near null space dimension is six. Unfortunately, these techniques are not well-suited for problems with particularly large near null space dimension such as the eddy current formulation of Maxwell's equations. Here, several different techniques have been proposed including [9, 10, 18, 15, 7, 2, 3, 5, 11, 12, 13, 14, 6] where [2] is most closely related to the direction pursued in the research presented here. In particular, [2] proposed a type of patch method where the set of patches takes advantage of the structure of the near null space for the curl operator. The key is that even though the null space dimension is large, the null space can be represented by a basis of locally supported vectors. The work of this LDRD has been to develop general user-defined patch preconditioning capabilities that allow for patches to span multiple MPI ranks and to adapt this patch relaxation technique to more general PDE systems that are far more complicated than the eddy current equations.

Our primary initial target for the new techniques is a generalized or extended magnetohydrodynamic (GMHD) PDE system. The specific equations are outlined in the next section. One main new ingredient is that this system actually has two different PDE expressions and each of these can generate a large near null space. One expression contains the curl operator while the other is a Hall term. It is possible for one of these terms to be dominant within some sub-regions of the domain while the other term is dominant within other sub-regions. It is also possible that both terms are dominant within the same sub-region or that neither term is dominant in other regions. Further, the character of the near null space is quite different depending on which term is responsible. For the curl operator, the null space corresponds to the space of gradients of certain nodal basis functions while for the Hall term it is a restriction of the magnetic field locally.

One major finding of this LDRD is that the curl operator and the Hall term also complicate the discretization process. The basic issue is that some discretizations are advantageous (e.g., preserve certain null space properties) when considering a subset of the PDE terms but not desirable when considering all PDE terms in the system. For example, we used the ALEGRA multiphysics code base for our research [19]. The original GMHD discretization that had been developed and existed in ALEGRA was found to lead to a discrete null space that included non-physical oscillatory modes. Given this deficiency, we devised and analyzed several other discretizations with different trade-offs in terms of accuracy, stability, and solvability. From these we identified one discretization scheme that seems to be a better candidate for future GMHD work.

Leveraging the the new discretization scheme, we then adapted a type of patch preconditioner for the linear systems on two different problems (a shock tube system and a current sheet problem). The degrees-of-freedom defining each patch lie within a group of $2 \times 2 \times 2$ neighboring elements. With the previous discretization and several previous linear solvers (including AMG for curl operators employing a variety of different smoothers), the linear solver failed to converge or converged extremely poorly for some problem/parameter regimes. We demonstrate that the newly developed patch solver can address these highly ill-conditioned linear systems when the reformulated discretization is employed. In particular, most of these difficult problems can be solved in less than 50 iterations and the overall convergence rate is not greatly affected by the mesh resolution. Further research is needed to extend the patch methodology to more complex cases and to further evaluate the associated convergence behavior to see if further adaptions might be necessary.

DETAILED DESCRIPTION OF RESEARCH AND DEVELOPMENT AND METHODOLOGY:

To solve PDE systems such as GMHD, we focus on complementing existing AMG procedures with patch-based relaxation schemes. These patch-based relaxation schemes correspond to an overlapping Schwarz or equivalently an overlapping block Jacobi procedure. Figure 1 depicts a sample mesh covered by a set of patches. Theoretically, the sub-domains (or the Jacobi blocks) should be chosen so that there is at least one sub-domain associated with each vector of a minimal basis that spans the PDE system's near null space. The size of the sub-domain should be sufficiently large to completely encompass the local non-zero support of its associated basis vector. Thus, the first part of the project focuses on understanding the discrete near null space of the PDE system that one is interested in solving while the second major part of the project develops a general framework for creating patch-based preconditioners.

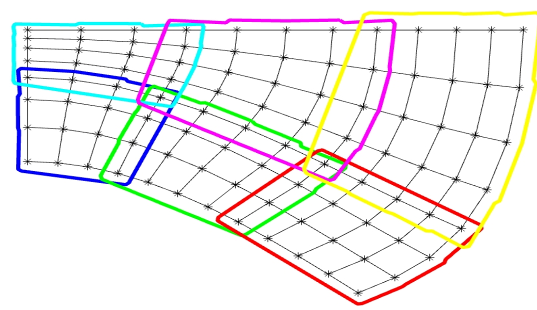

Figure 1: Sample mesh covered with 6 overlapping patches.

Discretization and understanding the near null space. Our initial target application comes from generalized MHD. The equations of interest are

$$\partial_t \mathbf{D} - \operatorname{curl} \mathbf{H} + \mathbf{J} = 0, \qquad \partial_t \mathbf{B} + \operatorname{curl} \mathbf{E} = 0, \qquad \mathbf{D} = \epsilon \mathbf{E}, \qquad \mathbf{B} = \mu \mathbf{H}, \tag{1}$$

and

$$\tau \partial_t \mathbf{J} + \mathbf{J} - \boldsymbol{\beta} \times \mathbf{J} - \sigma \mathbf{E} = 0 \tag{2}$$

with

$$\boldsymbol{\beta} = \frac{\mathrm{e}\tau}{m_e}\mathbf{B}, \qquad \sigma = \frac{\mathrm{e}^2 n_e}{m_e}\tau. \tag{3}$$

Here, $\mathbf{E}$ is the electric field, $\mathbf{H}$ is the magnetic field, $\mathbf{D}$ is the displacement field, $\mathbf{J}$ is the current density, $\mathbf{B}$ is the magnetic flux density (magnetic induction), $|\boldsymbol{\beta}|$ is the so-called Hall parameter, σ is the electrical conductivity, μ is the magnetic permeability, ϵ is the electric permittivity, e denotes the unit electric charge, n_e the electron number density, m_e the electron mass, and τ the ion-electron relaxation time. We note that (2) is a generalized Ohm's law which replaces the classical Ohm's law $\mathbf{J} = \sigma \mathbf{E}$. The presence of (2) gives rise to generalized MHD (as opposed to more traditional ideal or resistive MHD equations) and leads to additional numerical complications. The Hall term $\boldsymbol{\beta} \times \mathbf{J}$ can be problematic due to the fact that the $\boldsymbol{\beta}\times$ operator has a large null space. Recalling the vector identity $\mathbf{v} \times \mathbf{v} = 0$, it follows that for a given $\boldsymbol{\beta}$, $[\beta_1 \delta(x,y,z), \beta_2 \delta(x,y,z), \beta_3 \delta(x,y,z)]$ lies in the null space of $\boldsymbol{\beta}\times$ for any spatial location given by (x,y,z). Here, $\delta(x,y,z)$ is the Dirac delta function centered at (x,y,z) and β_k refers to the k^{th} component of $\boldsymbol{\beta}$. The introduction of $\delta(x,y,z)$ follows from the fact that $\boldsymbol{\beta}\times$ includes no spatial derivatives or that the vector identity holds at any point within the domain. When the magnitude of $\boldsymbol{\beta}$ is large relative to other PDE terms, the $\boldsymbol{\beta}\times$ term leads to a large near null space of the generalized Ohm's law PDE.

While the continuous near null space due to the Hall ($\boldsymbol{\beta}\times$) term is easy to identify, it is important to actually understand the *discrete* near null space for the resulting discrete linear system in order to design patches for a patch preconditioner. To begin, we must first define a suitable discretization. In particular, consider the 2×2 E-J block system

$$\begin{bmatrix} \mathbb{A}_{11} & \mathbb{A}_{12} \\ \mathbb{A}_{21} & \mathbb{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{E}^h \\ \mathbf{J}^h \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{bmatrix}, \tag{4}$$

where $\mathbf{E}^h$ and $\mathbf{J}^h$ are the basis coefficients for the electric field and current density, respectively, at time t^{n+1}.

Let $\widehat{\mathbf{E}}_i$ denote an arbitrary edge basis function for the discretization of $\mathbf{E}$, and $\widehat{\mathbf{J}}_i$ denote an arbitrary edge basis function for the discretization of $\mathbf{J}$. $\mathbb{A}_{11}$ is defined by the bilinear form

$$\mathbb{A}_{11}(\widehat{\mathbf{E}}_i, \widehat{\mathbf{E}}_j) = \frac{\epsilon}{\Delta t}\int_\Omega \widehat{\mathbf{E}}_i \cdot \widehat{\mathbf{E}}_j \, d\Omega + \frac{\Delta t}{\mu}\int_\Omega \operatorname{curl}\widehat{\mathbf{E}}_i \cdot \operatorname{curl}\widehat{\mathbf{E}}_j \, d\Omega. \tag{5a}$$

$\mathbb{A}_{12}$ is defined by the bilinear form

$$\mathbb{A}_{12}(\widehat{\mathbf{E}}_i, \widehat{\mathbf{J}}_j) = \int_\Omega \widehat{\mathbf{E}}_i \cdot \widehat{\mathbf{J}}_j \, d\Omega. \tag{5b}$$

The second block row of the discrete system represents the generalized Ohm's law relationship between current and electric field. Originally, this Ohm's law relationship was discretized in ALEGRA by leveraging an element projection operator. Among several advantageous features, this projection operator allowed one to define a smaller linear system involving only the electric field (and not the current) via an analytic Schur complement elimination approach. Unfortunately, one of the main discoveries of this project is that the resulting discrete matrix system is problematic from a linear solver perspective. Specifically, while the edge elements nicely conserve null space properties of the curl operator, the combination of edge elements and the element projection operator introduces non-physical oscillatory modes into the near null space associated with the Hall term. More broadly, our investigation highlights an unanticipated GMHD discretization challenge that centers on trying to preserve important properties of two null spaces, each of which has a very different character.

To avoid some of these oscillatory near null space difficulties, we opted to eliminate the use of the special element projection operator and instead solve the large 2×2 block system. Specifically, the lower block row of the matrix equation is defined by

$$\mathbb{A}_{21}(\widehat{\mathbf{J}}_i, \widehat{\mathbf{E}}_j) = -\int_\Omega \sigma^{n+1} \left(\widehat{\mathbf{J}}_i \cdot \widehat{\mathbf{E}}_j \right) \, d\Omega, \tag{5c}$$

and

$$\mathbb{A}_{22}(\widehat{\mathbf{J}}_i, \widehat{\mathbf{J}}_j) = \int_\Omega \left(\left(1 + \frac{\tau}{\Delta t}\right) \left(\widehat{\mathbf{J}}_i \cdot \widehat{\mathbf{J}}_j \right) - \frac{\mathrm{e}\tau}{m_e} \left[\widehat{\mathbf{J}}_i \cdot \left(\mathbf{B}^n \times \widehat{\mathbf{J}}_j \right) \right] \right) \, d\Omega. \tag{5d}$$

While a complete analysis of this system will be performed in future work, the resulting discrete linear system appears to be more suitable for solving linear systems via iterative methods.

Patch preconditioner When solving a linear system

$$Au = f,$$

a single iteration of a patch preconditioner generally has the form

$$M_P^{-1} = \sum_{k=1}^{n_p} V_k^T W_k (V_k A V_k^T)^{-1} V_k \tag{6}$$

where n_p is the number of patches, W_k is a diagonal scaling or weight matrix, and V_k is a rectangular binary restriction matrix. Specifically, V_k selects those entries in the solution vector u that appear within patch k. That is, the $(i,j)^{th}$ element of V_k is nonzero and equal to one if and only if the i^{th} degree-of-freedom within the k^{th} patch corresponds to the j^{th} degree-of-freedom in the global solution vector u. The $(i,i)^{th}$ entry of the weight matrix W_k normally corresponds to the reciprocal of the number of patches that include the associated solution vector degree-of-freedom. Thus, $V_k A V_k^T$ projects the matrix system on to a single patch. Overall, the entire patch preconditioner corresponds to solving n_p independent small linear systems and averaging the result within overlapped regions. In the example depicted in Figure 1, $n_p = 6$.

The Ifpack2 package [17] within the Trilinos framework [8] already included some patch preconditioning capabilities. As these capabilities did not directly suit our needs, we generalized them and adapted them so that they are now much broader and useful when solving PDE systems. Specifically, it is now possible to conveniently specify any patch by providing a set of global ids. Previously, it was necessary to define patches by providing processor-local ids. This is often quite awkward within many applications and effectively made it impossible to define a general patch which could span multiple MPI ranks. The adapted patch methods now allow for overlapping patches to be defined where individual patches might cross several MPI ranks when the new capability is combined with Ifpack2's overlapping Schwarz functionality. A second major addition to Ifpack2 includes the definition of the diagonal scaling matrix W_k. Previously, Ifpack2 only provided a symmetric version of scaling and the definition of the weights did not correspond to averaging in overlap regions, which is the most natural choice for patch preconditioners. Further, it was not possible to properly compute weights when some patches were fully defined on one MPI-rank while other patches spanned multiple MPI-ranks. A new Ifpack2 option now allows one to define W_k so that proper averaging occurs in overlap regions regardless of whether each patch reside on one MPI-rank or across several MPI-ranks. A couple of other minor Ifpack2 adaptations were also added to fix some algorithmic shortcomings in terms of how singleton blocks are defined.

RESULTS AND DISCUSSION:

While patch-based methods have been used for some MHD systems [1] and have also been used for incompressible Navier-Stokes systems [22], we are unaware of any prior work that employs patch preconditioners for GMHD. As already discussed, a range of GMHD discretizations were investigated and devised. We applied the 2×2 discretization (without the special element projection operator) along with a specific set of patches to the initial time step of a shock tube problem. Figure 2 displays the initial geometry of the shock tube. In particular, a $n \times 2 \times 2$ element mesh is used for varying values of n. Dirichlet boundary conditions are employed in the long dimension

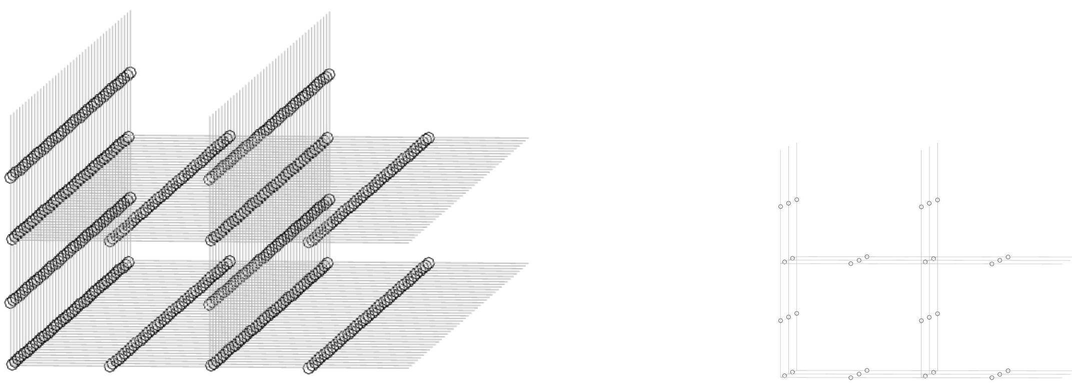

Figure 2: Left: Edges of shock tube (circles denote edge center point). Right: edges associated with one sample patch.

while the other two dimensions enforce periodic boundary conditions. The degrees-of-freedom are the electric field and current basis coefficients for the edge elements. Figure 2 depicts edges corresponding to the electric field and current basis coefficients that are being determined by solving a linear system. Edges on the far right and top of the $n \times 2 \times 2$ element mesh do not appear explicitly in the linear system due to the periodic boundary conditions. The various model parameters are taken to be $\sigma = 10^6$, $m_e = 10^{-4}/(1 + 10^{-4})$, $\mathrm{e} = 10^4/(1 + 10^{-4})$, $\mu = 1$, and $\epsilon = 10^{-4}$. The initial conditions for the electron density and magnetic field are given by

$$n_e = \begin{cases} 1.0, & 0.0 \leq x < 0.5, \\ 0.125, & 0.5 < x \leq 1.0, \end{cases} \qquad \mathbf{B} = \begin{cases} (0.75, +1.0, 0.0), & 0.0 \leq x < 0.5, \\ (0.75, -1.0, 0.0), & 0.5 < x \leq 1.0, \end{cases} \tag{7}$$

It should be noted that in this parameter regime the generalized Ohm's law term is dominant and that the curl term does not contribute significantly to the ill-conditioning of the matrix. The Ohm's law Hall term does not include any spatial derivatives and so it resembles a type of mass matrix whose sole source of ill-conditioning is due to the near null space aspects of the $\boldsymbol{\beta}\times$ operator. Despite the relatively modest size of this problem, the condition number is over 10^7 primarily due to these near null space effects. For the patch preconditioner a set of $n - 1$ patches are defined. In particular, each patch corresponds to all degrees-of-freedom that lie within a $2 \times 2 \times 2$ set of adjacent elements that effectively form a two-element wide slice of the tube. Figure 3 shows the convergence profile when the patch method is used as a preconditioner for GMRES on the shock tube problem. Here, one can see that convergence rates are quite rapid and nearly constant

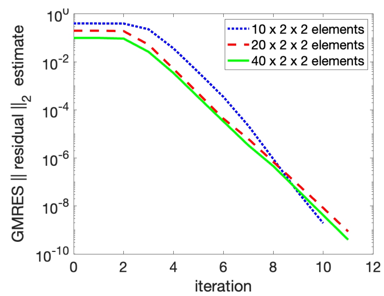

Figure 3: Residual histories using GMRES with patch preconditioner on shock tube problem.

regardless of the number of elements in the long direction of the shock tube. It should be noted that we were unable to satisfactorily solve this problem with our traditional algebraic multigrid solvers. In this particular case, the patch preconditioner is used without multigrid, which is not needed due to the lack of dominant spatial derivatives in this parameter regime. That is, a local overlapping Jacobi-style method can completely address the ill-conditioning due to the near null space of the Hall term. When we use the patch method as a smoother within a two level geometric multigrid scheme, the convergence is almost identical to that obtained without multigrid as expected. We anticipate that in other parameter regimes multigrid will be needed and so the patch method is best employed as a multigrid smoother (perhaps introducing a scalar damping parameter before the summation in (6)).

As a second test example, an $2n \times n \times 2$ element mesh is considered for a perturbed current sheet problem (cf. [4] for details of the problem setup). Dirichlet boundary conditions are enforced in the y dimension while the other two dimensions enforce periodic boundary conditions. The various model parameters are taken to be $\sigma = 10^6$, $m_e = 4.16 \times 10^{-2}/(1 + 4 \times 10^{-2})$, e $= 1$, $\mu = 1$, $\epsilon = 1/c^2$, where $c = \delta m_i^{-1/2}$, $m_i = 1.04/(1 + 4 \times 10^{-2})$, and δ is taken to be either 10^2 or 10^6. The matrix condition number for the $64 \times 32 \times 2$ element mesh is 4.7×10^{10} when $\delta = 10^2$ while the condition number is 7.4×10^{11} when $\delta = 10^6$. It should be noted that for this problem the curl term now also contributes to the condition number of the matrix proportionally to the value of δ. While the near null space ill-conditioning is problematic, we found that large variations in the magnitude of matrix entries coming from the four sub-blocks also contribute somewhat to the ill-conditioning and additionally lead to significant rounding error effects within GMRES. Specifically, the preconditioned GMRES internal estimate of the residual's 2-norm becomes highly inaccurate after a 10^{-4} reduction, ultimately leading to convergence stagnation. Fortunately, a simple symmetric scaling transformation helps reduce rounding error effects on convergence. In particular, we found that scaling $\mathbb{A}_{11}$ by 10^4, $\mathbb{A}_{12}$ by 10^2, $\mathbb{A}_{21}$ by 10^2 and not scaling $\mathbb{A}_{22}$ approximately equilibrates the Frobenius norm of the four sub-matrices leading to an equivalent linear system where the GMRES internal estimate of the residual's 2-norm remains accurate even when the residual is reduced by nearly 8 orders of magnitude. Further, the condition number is reduced by about 2 orders of magnitude with this scaling. To maintain consistency, the first component of the linear system right hand side b_1 must be scaled by 100 and the first component of the final GMRES solution must also be scaled by 100 to arrive at the final approximate solution. In this case, the patch preconditioner is applied to the scaled system and each patch again corresponds to a $2 \times 2 \times 2$ set of adjacent elements. The total number of patches is 465 ($= 31 \times 15 \times 1$) and 1953 ($= 63 \times 31 \times 1$) for the small and large meshes respectively. Figure 4 shows the convergence profile when the patch method is used as a preconditioner for GMRES on the scaled current sheet problem. Here, one can see that

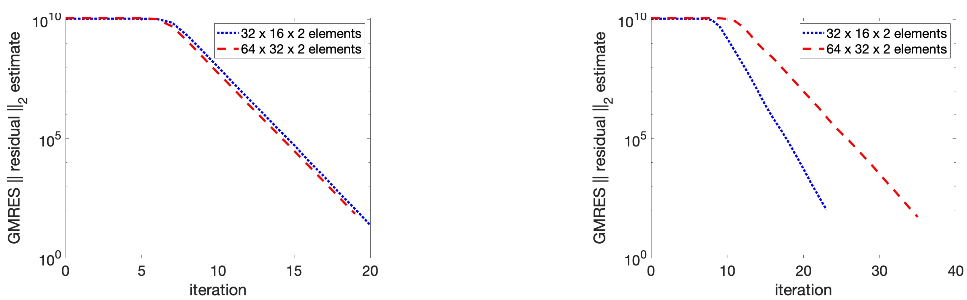

Figure 4: Residual histories using GMRES with patch preconditioner on scaled current sheet problem. Left: Hall term dominates ($\delta = 10^2$). Right: significant curl term ($\delta = 10^6$)

convergence rates are still fairly rapid and do not vary too significantly between the two mesh sizes for the smaller δ case. Once again, we are unable to satisfactorily solve either of these problems with our traditional algebraic multigrid solvers. It is worth noting that incorporating the patch method as a smoother within a 2-level Multigrid cycle does not significantly improve the residual convergence behavior of preconditioned GMRES even in the larger δ scenario where convergence is slower on the refined mesh. Clearly, further investigation is needed to analyze this large δ case and to understand whether convergence rates can be improved for refined meshes. It should be emphasized that these results are extremely preliminary in nature and that more follow-up work is necessary to see if the reformulated discretization can be used in conjunction with the preconditioner to address the Hall term near null space for more complex problems defined with complex domains. In these cases, one would need to adapt the patch definition used here for a structured domain to something more general for complex domains. One possibility might be to define one patch for each internal mesh node such that it includes all degrees-of-freedom defined within all elements adjacent to the mesh node.

ANTICIPATED OUTCOMES AND IMPACTS:

Many multi-physics systems encompass a range of complex interacting effects that significantly complicate the linear solution process. Unfortunately, these solver challenges could severely limit future simulations. This project has helped highlight the importance of discrete representation trade-offs between accuracy, stability, and linear system solvability. A new discrete GMHD representation that we developed shows promise in terms of accuracy, stability, and solvability, though further research is still needed. In particular, we have shown that patch-based methods can be very effective on at least some highly ill-conditioned GMHD matrices. The LDRD reformulation work will spur GMHD Sandia solver developments, and we anticipate that this work will be of interest within the GMHD research community. It also has relevance when considering other extended MHD extensions that incorporate additional physics (Nernst, Ettingshausen, Righi-Leduc). This physics is extremely important, requiring care in developing discrete representations and solvers.

More broadly, patch-based preconditioners and patch-based multigrid smoothers will continue to play an important role for solving PDE systems. Previously, they have demonstrated their utility for some incompressible Navier-Stokes systems and for some MHD systems. As described in this report, patch-based methods also now show potential for solving the reformulated generalized MHD systems described in this manuscript. Currently Sandia and much of the DOE community frequently use ILU-based relaxation schemes for PDE systems, which is generally less-scalable both theoretically and in practice than the patch-based methods. It is also somewhat more challenging to leverage advanced accelerators when employing ILU methods. The heavy use of ILU-based schemes is due in part because of the lack of easy-to-use generalized patch-based software that can automatically address complex domains and function naturally with algebraic multigrid methods. The work that we have done has extended the Trilinos capabilities to construct fairly arbitrary patch-based methods and so this should help promote/facilitate their adoption. Further work, however, is necessary to continue to simplify their use for application developers. Additionally, new tools are needed to help employ patch-based smoothers within algebraic multigrid methods where there is normally no notion of an element on automatically generated coarser levels. Finally, the current implementation of the patch method is most likely not optimal and so some future work will be required to optimize this algorithm, which may involve some non-trivial refactorization.

Given the significant potential of patch-based methods, a follow-on project is planned to consider many of these issues and to further develop tools in the context of GMHD systems. Specifically, the reformulated GMHD system that we developed will serve as a starting point for our future GMHD work at Sandia. While the resulting linear system is larger than the original matrix system formulated within ALEGRA, we believe it is more amenable to a combination of patch-based smoothers and multigrid. We will also investigate generalizations of patch methods to complex domains and with algebraic multigrid as smoothers.

CONCLUSION:

We have studied a class of patch preconditioners for multi-physics PDE systems. These methods correspond to a type of overlapping block Jacobi method and can potentially be used for a wide range of systems. The purely local nature of the patch linear sub-systems make them highly parallelizable, but often they may need to be used in conjunction with a multigrid method in order to achieve rapid linear solver convergence. This project has focused specifically on adapting patch methods to GMHD systems that previously proved problematic (significant convergence limitations or even stagnation) for our existing linear solution techniques. A spectral analysis of our initial set of GMHD test matrices revealed potential discretization concerns from the perspective of solving linear systems when the Hall term plays a dominant role. In particular, the discrete near null space included some non-physical oscillatory modes. As a result of this finding, we investigated alternative discrete representations settling on one that avoids an element projection operator at the expense of now needing to solve a larger linear system. Using this reformulated matrix, we devised a set of patches for a shock tube GMHD simulation and for a current sheet GMHD test problem. In both cases, patches are chosen by taking sets of overlapping $2 \times 2 \times 2$ element cubes. We showed that the corresponding patch preconditioner is able to solve an initial time step matrix system rather easily. For example, the shock tube system requires only 13 iterations, which is independent of the number of elements used to discretize the length of the tube. Previously, our standard algebraic multigrid procedures struggled to solve the original matrix systems. Our preliminary work will lead to a follow-on project where several unanswered research questions need to be addressed before a practical and general-purpose solver can be realized. Additional questions revolve around adaptations to more complex domains and the more general integration of the patch techniques with an algebraic multigrid procedure.

References

[1] J. Adler, T. Benson, E. Cyr, P. Farrell, S. MacLachlan, and R. Tuminaro. Monolithic multigrid methods for magnetohydrodynamics. *accepted for publication in SIAM J. on Sci. Comput.*, 2021. https://arxiv.org/abs/2006.15700.

[2] D. Arnold, R. Falk, and R. Winther. Multigrid in H(div) and H(curl). *Numer. Math.*, 85:197–217, 2000.

[3] R. Beck. Algebraic multigrid by component splitting for edge elements on simplicial triangulations. Technical Report SC 99-40, Zuse Institute Berlin, December 1999.

[4] J. Birn, J. F. Drake, M. A. Shay, B. N. Rogers, R. E. Denton, M. Hesse, M. Kuznetsova, Z. W. Ma, A. Bhattacharjee, A. Otto, and P. L. Pritchett. Geospace environmental modeling (GEM) magnetic reconnection challenge. *Journal of Geophysical Research: Space Physics*, 106(A3):3715–3719, 2001.

[5] P. Bochev, C. Garasi, J. Hu, A. Robinson, and R. Tuminaro. An improved algebraic multigrid method for solving Maxwell's equations. *SIAM J. Sci. Comput.*, 25(2):623–642, 2003.

[6] 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. Comput.*, 31(1):557–583, 2008.

[7] T. Boonen, G. Deliége, and S. Vandewalle. On algebraic multigrid methods derived from partition of unity nodal prolongators. *Numerical Linear Algebra with Applications*, 13(2-3):105–131, March 2006.

[8] M. A. Heroux, R. A. Bartlett, V. E. Howle, R. J. Hoekstra, J. J. Hu, T. G. Kolda, R. B. Lehoucq, K. R. Long, R. P. Pawlowski, E. T. Phipps, A. G. Salinger, H. K. Thornquist, R. S. Tuminaro, J. M. Willenbring, A. Williams, and K. S. Stanley. An overview of the Trilinos project. *ACM Transactions on Mathematical Software*, 31(3):397–423, 2005.

[9] R. Hiptmair. Multigrid method for Maxwell's equations. *SIAM J. Numer. Anal.*, 36:204–225, 1998.

[10] R. Hiptmair and J. Xu. Nodal auxiliary space preconditioning in H(curl) and H(div) spaces. *SIAM J. on Numer. Anal.*, 45(6):2483–2509, 2007.

[11] J. Hu, R. Tuminaro, P. Bochev, C. Garasi, and A. Robinson. Toward an h-independent algebraic multigrid method for Maxwell's equations. *SIAM J. Sci. Comput.*, 27(5):1669–1688, 2006.

[12] T. Kolev, J. Pasciak, and P. Vassilevski. H(curl) auxiliary mesh preconditioning. Technical Report UCRL-JRNL-224227, Lawrence Livermore National Labs, June 2006.

[13] T. Kolev and P. Vassilevski. Parallel H1-based auxiliary space AMG solver for H(curl) problems. Technical Report UCRL-TR-222763, Lawrence Livermore National Labs, July 2006.

[14] T. Kolev and P. Vassilevski. Some experience with a H1-based auxiliary space AMG for H(curl) problems. Technical Report UCRL-TR-21841, Lawrence Livermore National Labs, June 2006.

[15] B. Lee and C. Tong. A novel algebraic multigrid-based approach for Maxwell's equations. Technical Report UCRL-JC-218750, Lawrence Livermore National Laboratory, 2006.

[16] J. Mandel, M. Brezina, and P. Vaněk. Energy optimization of algebraic multigrid bases. *Computing*, 62(3):205–228, 1999.

[17] A. Prokopenko, C. M. Siefert, J. J. Hu, M. Hoemmen, and A. Klinvex. Ifpack2 User's Guide 1.0. Technical Report SAND2016-5338, Sandia National Labs, 2016.

[18] S. Reitzinger and J. Schöberl. An algebraic multigrid method for finite element discretizations with edge elements. *Numer. Linear Algebra Appl.*, 9:223–238, 2002.

[19] A. C. Robinson and et. al. ALEGRA: An arbitrary Lagrangian-Eulerian multimaterial, multiphysics code. *Proceedings of AIAA Meeting*, pages 1–39, 2008.

[20] P. Vaněk, J. Mandel, and M. Brezina. Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems. *Computing*, 56:179–196, 1996.

[21] P. Vaněk, J. Mandel, and M. Brezina. Convergence of algebraic multigrid based on smoothed aggregation. *Numerische Mathematik*, 88:559–579, 2001.

[22] S. Vanka. Block-implicit multigrid solution of Navier-Stokes equations in primitive variables. *J. of Comp. Phys.*, 65(1):138 – 158, 1986.

LABORATORY DIRECTED RESEARCH & DEVELOPMENT

WHERE INNOVATION BEGINS

ADDENDUM:

Composing preconditioners for multiphysics PDE systems with applications to Generalized MHD

PI: R. Tuminaro, PM: J. Feddema , Team Members: M. Crockatt & A. Robinson

Project goal

Develop algorithms and software to solve linear systems associated with multi-physics PDEs using generalized magneto-hydrodynamics (GMHD) as an application driver

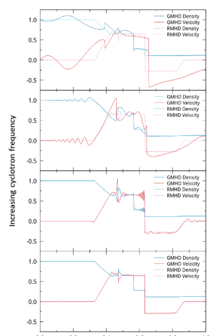

Mission Impact

This capability opens up new multigrid strategies for multi-physics PDEs. It could impact future Sandia Next Generation Pulsed Power Facility objectives and impact future computational models for inertial confinement fusion that include additional physical effects beyond those of traditional resistive MHD.

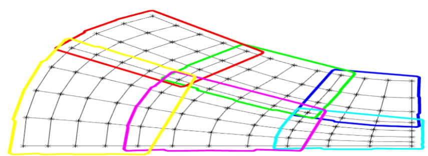

Completed Technical Milestones

Analysis of GMHD discrete linear system near null space	3/2022
Reformulation of discrete GMHD system with improved near null space properties	5/2022
Design new patch solvers (leveraging ASC Trilinos capabilities)	7/2022
Apply patch solve to GMHD shock tube matrix	8/2022

Transition Plan

Algorithms and software developed in LDRD will transition into Trilinos (NNSA NA-10 funded Advanced Simulation and Computing (ASC) program), an FY23 1400/1600 partnership project, and possibly a future DOE Office of Science Advanced Scientific Computing Research (ASCR) project.

Figure 5: Year-ending PowerPoint summary.