

Portland State University

PDXScholar

---

Mathematics and Statistics Faculty  
Publications and Presentations

Fariborz Maseeh Department of Mathematics  
and Statistics

---

8-25-2024

# Matrix-Free High-Performance Saddle-Point Solvers for High-Order Problems in $H(\text{div})$

Will Pazner

*Portland State University, pazner@pdx.edu*

Tzanio Kolev

*Lawrence Livermore National Laboratory*

Panayot s. Vassilevski

*Portland State University*

Follow this and additional works at: [https://pdxscholar.library.pdx.edu/mth\\_fac](https://pdxscholar.library.pdx.edu/mth_fac)



Part of the [Physical Sciences and Mathematics Commons](#)

Let us know how access to this document benefits you.

---

## Citation Details

Pazner, W., Kolev, T., & Vassilevski, P. S. (2024). Matrix-Free High-Performance Saddle-Point Solvers for High-Order Problems in  $H(\text{div})$ . *SIAM Journal on Scientific Computing*, 46(3), B179–B204. <https://doi.org/10.1137/23m1568806>

This Pre-Print is brought to you for free and open access. It has been accepted for inclusion in Mathematics and Statistics Faculty Publications and Presentations by an authorized administrator of PDXScholar. Please contact us if we can make this document more accessible: [pdxscholar@pdx.edu](mailto:pdxscholar@pdx.edu).

# MATRIX-FREE HIGH-PERFORMANCE SADDLE-POINT SOLVERS FOR HIGH-ORDER PROBLEMS IN $\mathbf{H}(\text{div})$

WILL PAZNER<sup>1</sup>, TZANIO KOLEV<sup>2</sup>, AND PANAYOT VASSILEVSKI<sup>1</sup>

**ABSTRACT.** This work describes the development of matrix-free GPU-accelerated solvers for high-order finite element problems in  $\mathbf{H}(\text{div})$ . The solvers are applicable to grad-div and Darcy problems in saddle-point formulation, and have applications in radiation diffusion and porous media flow problems, among others. Using the interpolation–histopolation basis (cf. [44]), efficient matrix-free preconditioners can be constructed for the  $(1, 1)$ -block and Schur complement of the block system. With these approximations, block-preconditioned MINRES converges in a number of iterations that is independent of the mesh size and polynomial degree. The approximate Schur complement takes the form of an M-matrix graph Laplacian, and therefore can be well-preconditioned by highly scalable algebraic multigrid methods. High-performance GPU-accelerated algorithms for all components of the solution algorithm are developed, discussed, and benchmarked. Numerical results are presented on a number of challenging test cases, including the “crooked pipe” grad-div problem, the SPE10 reservoir modeling benchmark problem, and a nonlinear radiation diffusion test case.

## 1. INTRODUCTION

This paper is concerned with the efficient solution of high-order finite element problems posed in the space  $\mathbf{H}(\text{div})$ ; in particular, we are interested in the grad-div problem

$$(1) \quad -\nabla(\alpha \nabla \cdot \mathbf{u}) + \beta \mathbf{u} = \mathbf{f},$$

and the closely related Darcy-type problem

$$(2) \quad \begin{cases} \mathbf{u} + \varepsilon \nabla q = \mathbf{f}, \\ \nabla \cdot \mathbf{u} + \gamma q = g, \end{cases}$$

where  $\mathbf{u}: \Omega \rightarrow \mathbb{R}^d$  is a vector-valued unknown,  $q: \Omega \rightarrow \mathbb{R}$  is a scalar unknown ( $\Omega \subseteq \mathbb{R}^d$ ,  $d \in \{2, 3\}$ ), and  $\alpha, \beta, \gamma$  and  $\varepsilon$  are potentially heterogeneous scalar coefficients. Such problems arise in numerous applications, including, for example, radiation diffusion and porous media flow, among others [9, 13]. We consider finite element discretizations for (1) and (2), where  $\mathbf{u}$  is approximated in the Raviart–Thomas  $\mathbf{H}(\text{div})$  finite element space, and  $p$  is approximated in the discontinuous  $L^2$  finite element space.

The solution of such problems via high-order finite element methods presents several advantageous features, as well as some challenges. High-order methods promise high accuracy per degree of freedom, and have been demonstrated to achieve high performance on modern supercomputing architectures, in particular those based on graphics processing units (GPUs) and other accelerators [1, 31, 43]. However, the solution of the resulting large linear systems remains challenging for a number of reasons. With increasing polynomial degree, the systems become both more ill-conditioned, and less sparse: the number of nonzeros in the matrix scales like  $\mathcal{O}(p^{2d})$  and the condition number typically scales like  $\mathcal{O}(p^4/h^2)$ , where  $p$  is the polynomial degree and  $d$  is the spatial dimension [40]. Consequently, it is necessary to develop effective matrix-free preconditioners that give rise to well-conditioned systems without having access to an assembled matrix. There are a number of existing approaches for matrix-free preconditioning elliptic equations; among these are  $p$ -multigrid, low-order-refined (also known as FEM–SEM) preconditioning, and hybrid Schwarz preconditioning [35, 38, 42]. However, the extension of these approaches to problems in  $\mathbf{H}(\text{div})$  is not straightforward; for example, multigrid-type methods with simple point smoothers are known not to perform well for such problems [5].

Effective preconditioners for low-order  $\mathbf{H}(\text{div})$  discretizations include multigrid with Schwarz smoothers [6], auxiliary space algebraic multigrid preconditioning (e.g. the auxiliary divergence solver, ADS) [32], and algebraic hybridization [16]. We briefly describe some existing approaches for extending these approaches to high-order discretizations. In [44], a low-order-refined approach for preconditioning grad-div problems in

---

<sup>1</sup>FARIBORZ MASEEH DEPARTMENT OF MATHEMATICS AND STATISTICS, PORTLAND STATE UNIVERSITY, PORTLAND, OR

<sup>2</sup>CENTER FOR APPLIED SCIENTIFIC COMPUTING, LAWRENCE LIVERMORE NATIONAL LABORATORY, LIVERMORE, CA

$\mathbf{H}(\text{div})$  was proposed. In that work, the high-order operator discretized using the so-called interpolation–histopolation basis was shown to be spectrally equivalent to the lowest-order discretization posed on a refined mesh. Then, an effective preconditioner for the low-order discretization (for example, the AMG-based ADS solver) can be used to precondition the high-order operator. In [11], a vertex-star domain decomposition was combined with the use of a finite element basis with favorable sparsity properties to develop optimal-complexity multigrid solvers.

In the present work, we develop efficient matrix-free preconditioners for high-order  $\mathbf{H}(\text{div})$  problems by considering the associated  $2 \times 2$  saddle-point system. Such saddle-point systems can be preconditioned effectively using block preconditioners [8, 46, 51]. The main challenge associated with developing such methods is the construction of spectrally equivalent approximations to the inverse of the Schur complement. Using favorable properties of the interpolation–histopolation basis (cf. [44]) it is possible to both enhance the sparsity of the block system, and to form sparse M-matrix approximations to the Schur complement that are effectively preconditioned with standard algebraic multigrid methods. Although the block system is indefinite and contains more unknowns than the positive-definite system obtained by discretizing the grad–div problem directly, it is much less computationally expensive to precondition, resulting in significantly decreased cost per iteration, and faster overall solve time compared with alternative approaches. Additionally, the described methods have low memory requirements, are highly scalable, and are amenable to GPU acceleration.

The structure of this paper is as follows. In Section 2 we describe the high-order finite element discretization and the associated linear systems. This section also describes the construction of the block preconditioners, and includes results concerning the estimation of the condition number of the preconditioned system. Section 3 describes the algorithmic details and GPU acceleration of the proposed methods. Numerical examples are presented in Section 4; these include several challenging grad–div and Darcy benchmark problems, as well as a fully nonlinear time-dependent radiation-diffusion example. We end with conclusions in Section 5.

## 2. DISCRETIZATION AND LINEAR SYSTEM

In the discrete setting, we consider a mesh  $\mathcal{T}$  of the domain  $\Omega$ , where each element  $\kappa \in \mathcal{T}$  is the image of the unit cube  $[0, 1]^d$  under a (typically isoparametric) transformation  $T_\kappa$  (i.e.,  $\kappa = T_\kappa([0, 1]^d)$ ). In other words, we consider quadrilateral meshes in two spatial dimensions, and hexahedral meshes in three dimensions. The variational problem associated with (1) is given by

$$(3) \quad (\alpha \nabla \cdot \mathbf{u}, \nabla \cdot \mathbf{v}) + (\beta \mathbf{u}, \mathbf{v}) = (\mathbf{f}, \mathbf{v}).$$

To discretize (3), we take  $\mathbf{u}, \mathbf{v} \in \mathbf{V}_h$ , where  $\mathbf{V}_h$  is the degree- $p$  Raviart–Thomas finite element space,

$$\mathbf{V}_h = \{ \mathbf{v} \in \mathbf{H}(\text{div}) : \mathbf{v}|_\kappa \circ T_\kappa \in \mathbf{V}_h(\kappa) \}.$$

The local space  $\mathbf{V}_h(\kappa)$  is the image of the reference space  $\mathbf{V}_h(\hat{\kappa})$  under the  $\mathbf{H}(\text{div})$  Piola transformation  $\mathbf{V}_h(\kappa) = \det(J_\kappa)^{-1} J_\kappa \mathbf{V}_h(\hat{\kappa})$ , where  $J_\kappa$  is the Jacobian of the element transformation  $T_\kappa$ . In three dimensions, the reference space is given by  $\mathbf{V}_h(\hat{\kappa}) = \mathcal{Q}_{p,p-1,p-1} \times \mathcal{Q}_{p-1,p,p-1} \times \mathcal{Q}_{p-1,p-1,p}$ , where  $\mathcal{Q}_{p_1,p_2,p_3}$  is the space of trivariate polynomials of degree at most  $p_i$  in the variable  $x_i$ .  $\mathbf{V}_h(\hat{\kappa})$  is defined analogously in two dimensions. Functions  $\mathbf{v} \in \mathbf{V}_h$  have continuous normal components across element interfaces; no tangential continuity is enforced.

After discretization, the linear system of equations

$$(4) \quad \mathbf{A} \mathbf{u} = \mathbf{f}$$

is obtained, where  $\mathbf{A}$  is the matrix given by the sum of the  $\mathbf{H}(\text{div})$  stiffness and mass matrices,  $\mathbf{u}$  is a vector of degrees of freedom representing  $\mathbf{u}$ , and  $\mathbf{f}$  is a dual vector representing  $\mathbf{f}$ .

The grad–div problem (1) can be equivalently recast as the first-order system

$$(5) \quad \begin{cases} \beta \mathbf{u} - \nabla(\alpha q) = \mathbf{f}, \\ \nabla \cdot \mathbf{u} - q = 0, \end{cases}$$

leading to the symmetric variational problem

$$(6) \quad \begin{aligned} (\beta \mathbf{u}, \boldsymbol{\sigma}) + (\alpha q, \nabla \cdot \boldsymbol{\sigma}) &= (\mathbf{f}, \boldsymbol{\sigma}), \\ (\alpha \nabla \cdot \mathbf{u}, \tau) - (\alpha q, \tau) &= 0, \end{aligned}$$

where  $\mathbf{u}, \boldsymbol{\sigma} \in \mathbf{V}_h$  and  $q, \tau \in W_h$ , where  $W_h$  is the degree- $(p-1)$   $L^2$  finite element space

$$W_h = \{ w \in L^2(\Omega) : w|_{\kappa} \circ T_{\kappa} \in W_h(\kappa) \},$$

The local space  $W_h(\kappa)$  is the image of the reference space  $W_h(\hat{\kappa}) = \mathcal{Q}_{p-1}$  under the integral-preserving mapping,  $W_h(\kappa) = \det(J_{\kappa})^{-1} W_h(\hat{\kappa})$ . Note that the space  $W_h$  is equal to the range of divergence operator acting on  $\mathbf{V}_h$

The variational system (6) can be written in matrix form as

$$(7) \quad \begin{pmatrix} \mathbf{M}_{\beta} & \mathbf{B}_{\alpha}^{\top} \\ \mathbf{B}_{\alpha} & -\mathbf{W}_{\alpha} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{q} \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ 0 \end{pmatrix},$$

where the matrices  $\mathbf{M}$ ,  $\mathbf{B}$ , and  $\mathbf{W}$  are defined by

$$(8) \quad \begin{aligned} \mathbf{v}^{\top} \mathbf{M}_{\beta} \mathbf{u} &= (\beta \mathbf{u}, \mathbf{v}), \\ \mathbf{q}^{\top} \mathbf{B}_{\alpha} \mathbf{u} &= (\alpha \nabla \cdot \mathbf{u}, q), \\ \mathbf{r}^{\top} \mathbf{W}_{\alpha} \mathbf{q} &= (\alpha q, r). \end{aligned}$$

It is straightforward to see that eliminating the scalar unknown  $q$  from (5) results in the original grad-div system (1). Likewise, the Schur complement of the saddle-point matrix with respect to the  $(2, 2)$ -block  $\mathbf{M}_{\beta} + \mathbf{B}_{\alpha}^{\top} \mathbf{W}_{\alpha}^{-1} \mathbf{B}_{\alpha}$  is equal to the original  $\mathbf{H}(\text{div})$  matrix  $\mathbf{A}$ .

Note that the block system (7) is also closely related to the discretized system corresponding to the Darcy problem (2). In the case of the mixed finite element problem, the resulting system takes the form

$$(9) \quad \begin{pmatrix} \mathbf{M}_{1/\varepsilon} & \mathbf{B}^{\top} \\ \mathbf{B} & -\mathbf{W}_{\gamma} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{q} \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ \mathbf{g} \end{pmatrix},$$

where  $\mathbf{M}_{1/\varepsilon}$  is the  $\mathbf{H}(\text{div})$  mass matrix weighted by  $\varepsilon^{-1}$ ,  $\mathbf{B}$  corresponds to the unweighted divergence form, and  $\mathbf{W}_{\gamma}$  is the  $L^2$  mass matrix weighted by  $\gamma$  (cf. (8)). In contrast to the variational form corresponding to the grad-div problem, in this context the coefficient  $\gamma$  may be zero (corresponding to the Poisson problem with no reaction term), leading to the saddle-point system with zero  $(2, 2)$ -block

$$(10) \quad \begin{pmatrix} \mathbf{M}_{1/\varepsilon} & \mathbf{B}^{\top} \\ \mathbf{B} & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{q} \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ \mathbf{g} \end{pmatrix}.$$

In what follows, we focus on the solution of the  $2 \times 2$  saddle-point systems (7) and (9). While the original grad-div system is symmetric and positive-definite, and so may be solved using the conjugate gradient method, the saddle-point system is symmetric and indefinite, and so we use the MINRES Krylov method (or GMRES, when using a non-symmetric preconditioner). To precondition this system, we make use of block preconditioners developed in Section 2.2. The preconditioners for the grad-div saddle-point system (7) apply equally well to the Darcy block systems (9) and (10), including the case of zero  $(2, 2)$ -block.

**2.1. Interpolation-histopolation basis.** In order to construct efficient preconditioners for the block systems, we will use properties of a particular choice of bases for the high-order Raviart–Thomas space  $\mathbf{V}_h$  and  $L^2$  space  $W_h$ . These spaces will use the “interpolation-histopolation” basis (see [44] and also [18, 33, 43]). To define this basis, the  $d$ -dimensional reference element  $[0, 1]^d$  is subdivided into  $p^d$  subelements (where  $p$  is the polynomial degree) whose vertices are given by the  $d$ -fold Cartesian product of the  $p+1$  Gauss–Lobatto points (see Figure 1). Then, the basis is induced by the following choice of degrees of freedom:

- The degrees of freedom of the Raviart–Thomas space  $\mathbf{V}_h$  are given by integrals of the normal component over each face of the subelement mesh. The basis functions are induced by the tensor-product of one-dimensional interpolation and histopolation operators.
- The degrees of freedom of the  $L^2$  space  $W_h$  are given by integrals over each subelement volume. The basis functions are induced by the tensor-product of one-dimensional histopolation operators.

Given the above definitions, there is a natural identification of Raviart–Thomas degrees of freedom with subelement faces, and  $L^2$  degrees of freedom with subelement volumes. For our purposes, this choice of basis results in several important properties, the proofs of which can be found in [44].

- The Raviart–Thomas mass matrix  $\mathbf{M}$  defined on  $\mathbf{V}_h$  is spectrally equivalent to its diagonal, independent of mesh size and polynomial degree.

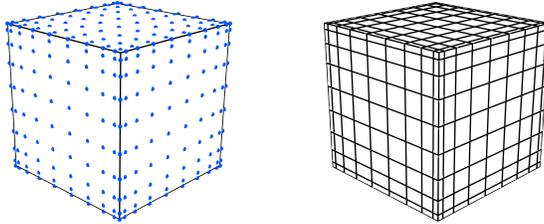


FIGURE 1. Left: high-order ( $p = 9$ ) hexahedral element with 10 Gauss-Lobatto nodes in each dimension. Right: subelement mesh with vertices at the Gauss-Lobatto points.

- The  $L^2$  mass matrix  $W$  defined on  $W_h$  is spectrally to its diagonal, independent of mesh size and polynomial degree.
- The discrete divergence matrix  $D : V_h \rightarrow W_h$  is sparse, with two nonzeros per column (taking values 1 and  $-1$ ), regardless of polynomial degree. The matrix is entirely determined by the mesh topology (its entries are independent of mesh geometry or distortion), and it coincides exactly with the lowest-order discrete divergence matrix defined on the refined mesh (see the right panel of Figure 1).

Because of the sparsity afforded by the discrete divergence matrix, it is advantageous to consider the transformed saddle-point matrix

$$(11) \quad \mathcal{A} := \begin{pmatrix} M_\beta & D^T \\ D & -W_\alpha^{-1} \end{pmatrix} = \begin{pmatrix} I & 0 \\ 0 & W_\alpha^{-1} \end{pmatrix} \begin{pmatrix} M_\beta & B_\alpha^T \\ B_\alpha & -W_\alpha \end{pmatrix} \begin{pmatrix} I & 0 \\ 0 & W_\alpha^{-1} \end{pmatrix}$$

using that  $D = W_\alpha^{-1}B_\alpha$  by the definitions in (8). In contrast to the untransformed saddle-point system introduced in (7), the transformed system  $\mathcal{A}$  has much sparser off-diagonal blocks (whose sparsity, critically, is independent of the polynomial degree). However, a downside of this transformation is that the matrix  $\mathcal{A}$  involves the *inverse* of the  $L^2$  mass matrix  $W_\alpha$ . In the context of an iterative solver such as MINRES, this means that the action of  $W_\alpha^{-1}$  will need to be computed at every iteration. Because  $W_h$  is a discontinuous space,  $W_\alpha$  is block-diagonal, and, at least for low orders, the blocks can be factorized efficiently element-by-element. For higher orders, it is not feasible to factorize (or even assemble) the blocks of  $W_\alpha$ , and so we turn to iterative matrix-free methods to apply the action of  $W_\alpha^{-1}$ . These methods are described in greater detail in Section 2.5, and their performance is studied in the numerical results in Section 3.1. The sparsity of the off-diagonal blocks and the diagonal spectral equivalence properties enumerated above allow for the sparse approximation of the Schur complement and efficient algebraic multigrid preconditioning, enabling the block preconditioning techniques described in the following section.

**2.2. Block preconditioners for saddle-point systems.** At this point, we make a brief digression to state some results concerning block preconditioning of  $2 \times 2$  saddle-point systems. We do not claim novelty of these results, however we collect them here in a form that will be convenient to reference later (see, e.g., [8, 46, 51] for related results concerning block preconditioning of saddle-point systems). Consider the saddle-point system

$$(12) \quad \mathcal{A} = \begin{pmatrix} A & B^T \\ B & -C \end{pmatrix}.$$

While the particular saddle-point system of interest is that given by (11), in this section we only make the assumption that the blocks  $A$  and  $C$  in (12) are symmetric and positive-definite (but make no other assumptions about the particular form of the blocks). The convergence of MINRES applied to this system can be estimated using bounds on the positive and negative parts of the spectrum,

$$\sigma(\mathcal{A}) \subseteq [\mu_1, \mu_2] \cup [\nu_1, \nu_2] \quad \text{with} \quad \mu_1, \mu_2 < 0, \quad \nu_1, \nu_2 > 0.$$

In particular, we are interested in the ratio  $C/c$ , where

$$C = |\lambda|_{\max} = \max\{|\lambda| : \lambda \in \sigma(\mathcal{A})\}, \quad c = |\lambda|_{\min} = \min\{|\lambda| : \lambda \in \sigma(\mathcal{A})\},$$

resulting in

$$c(x, x) \leq |(\mathcal{A}x, x)| \leq C(x, x).$$

Given a symmetric positive-definite preconditioner  $\mathcal{B}$ , the relevant bounds are those that satisfy

$$(13) \quad c(\mathcal{B}x, x) \leq |(\mathcal{A}x, x)| \leq C(\mathcal{B}x, x),$$

which can then be used to estimate the condition number,  $\kappa(\mathcal{B}^{-1}\mathcal{A}) \leq C/c$ , cf. [8, Section 10].

**Proposition 1.** *Let  $\mathcal{B}$  be the block-diagonal preconditioner*

$$\mathcal{B} = \begin{pmatrix} \mathbf{A} & 0 \\ 0 & \mathbf{S} \end{pmatrix},$$

where  $\mathbf{S} = \mathbf{C} + \mathbf{D}\mathbf{A}^{-1}\mathbf{D}^\top$  is the (negative) Schur complement of  $\mathcal{A}$  with respect to the (1,1)-block. This preconditioner is optimal in the sense that

$$\sigma(\mathcal{B}^{-1}\mathcal{A}) \subseteq \left[-1, (1 - \sqrt{5})/2\right) \cup \left[1, (1 + \sqrt{5})/2\right) \quad \text{and} \quad \kappa(\mathcal{B}^{-1}\mathcal{A}) \leq \frac{\sqrt{5} + 1}{\sqrt{5} - 1} \approx 2.618\dots$$

The condition number in Proposition 1 can be improved slightly by using a different scaling of the blocks (see, for example, [8, Theorem 3.8]).

**Proposition 2.** *The block-diagonal preconditioner*

$$\mathcal{B} = \begin{pmatrix} 2\mathbf{A} & 0 \\ 0 & \mathbf{S} \end{pmatrix},$$

is optimal in the sense that

$$\sigma(\mathcal{B}^{-1}\mathcal{A}) \subseteq [-1, -1/2) \cup [1/2, 1) \quad \text{and} \quad \kappa(\mathcal{B}^{-1}\mathcal{A}) \leq 2.$$

The constant 2 is the optimal relative scaling of  $\mathbf{A}$  in terms of minimizing the bound on  $\kappa(\mathcal{B}^{-1}\mathcal{A})$ .

In the above, we are assuming that the applying the preconditioner  $\mathcal{B}$  involves exactly inverting the matrix  $\mathbf{A}$  and Schur complement  $\mathbf{S}$ . In practice this is typically infeasible, and so the inverses of these blocks are themselves approximated by spectrally equivalent preconditioners. The following proposition shows that this does not impact the optimality of the preconditioners.

**Proposition 3.** *Given a saddle-point system  $\mathcal{A}$  and preconditioner  $\mathcal{B}$ , suppose that the inequalities (13) hold with constants  $c$  and  $C$ . Suppose further that  $\tilde{\mathcal{B}}$  is another preconditioner, spectrally equivalent to  $\mathcal{B}$ , i.e.*

$$(14) \quad \tilde{c}(\tilde{\mathcal{B}}x, x) \leq (\mathcal{B}x, x) \leq \tilde{C}(\tilde{\mathcal{B}}x, x).$$

Then,  $\tilde{\mathcal{B}}^{-1}\mathcal{A}$  is uniformly well-conditioned, with condition number bound

$$\kappa(\tilde{\mathcal{B}}^{-1}\mathcal{A}) \leq C\tilde{C}/(c\tilde{c}).$$

*Proof.* This follows immediately by combining (13) and (14). □

*Remark 1.* As a consequence of the above results, a block-diagonal preconditioner of the form

$$\mathcal{B} = \begin{pmatrix} \tau\tilde{\mathbf{M}} & 0 \\ 0 & \tilde{\mathbf{S}} \end{pmatrix},$$

where  $\tilde{\mathbf{M}}$  is spectrally equivalent to  $\mathbf{M}$  and  $\tilde{\mathbf{S}}$  is spectrally equivalent to the Schur complement  $\mathbf{S}$ , will result optimal (i.e. discretization parameter independent) convergence for MINRES when applied to the saddle-point system  $\mathcal{A}$ . The relative scaling of the blocks determined by the coefficient  $\tau$  should be chosen in accordance with Proposition 2. In the numerical examples presented in Section 4, we use the simple choice of  $\tau = 1$ .

*Remark 2.* Block-triangular preconditioners of the form

$$\mathcal{B} = \begin{pmatrix} \mathbf{M} & \mathbf{B}^\top \\ 0 & \mathbf{S} \end{pmatrix}$$

result in a preconditioned system with exactly one eigenvalue,  $\sigma(\mathcal{B}^{-1}\mathcal{A}) = \{1\}$ , and whose minimal polynomial has degree two, guaranteeing convergence of GMRES in at most two iterations [8]. In practice, the diagonal blocks must be replaced with spectrally equivalent approximations  $\tilde{\mathbf{M}}$  and  $\tilde{\mathbf{S}}$ . In this setting, convergence estimates for GMRES may be obtained using a field of values analysis, where the resulting bounds

depend on the inf-sup constant of the corresponding finite element problem [7]. However, because the triangular preconditioners are non-symmetric, GMRES instead of MINRES must be used, increasing the memory requirements. Whether the reduced number of iterations warrants sacrificing the short-term recurrence is typically problem-specific; in the numerical results we consider only the block-diagonal preconditioner.

*Remark 3.* Using the so-called Bramble–Pasciak (BP) transformation, it is possible to use a block-triangular preconditioner with the conjugate gradient method in a transformed inner product [10]. This inner product requires a scaled approximation of the  $(1, 1)$ -block of the matrix, typically necessitating the solution of an eigenvalue problem with the matrix  $\tilde{M}^{-1}M$ . The total work per iteration of the BP–CG method is similar to that of block-diagonally preconditioned MINRES. Preliminary investigations into the performance of the BP–CG method applied to the problems considered presently indicate largely similar performance to that of MINRES with block-diagonal preconditioning (but with the added cost of estimating the extremal eigenvalue of  $\tilde{M}^{-1}M$ ).

**2.3. Application to the grad-div system.** We now apply the block preconditioning techniques developed in Section 2.2 to the mixed  $(\mathbf{H}(\text{div})$  and  $L^2$ ) saddle-point system  $\mathcal{A}$  in the form given by (11). Using the properties of the interpolation-histopolation basis (see Section 2.1), the  $\mathbf{H}(\text{div})$  mass matrix  $M_\beta$  is spectrally equivalent to its diagonal, and so we can take  $\tilde{M} = \text{diag}(M_\beta)$ . We now turn to approximating the Schur complement  $S = W_\alpha + DM_\beta^{-1}D^\top$ . Using the spectral equivalence of  $W_\alpha$  to its diagonal, we take  $\tilde{W} = \text{diag}(W_\alpha)$ , and then define

$$(15) \quad \tilde{S} = \tilde{W}^{-1} + D\tilde{M}^{-1}D^\top.$$

Because of the properties of the basis, the matrix  $\tilde{S}$  takes a particularly simple structure. Each degree of freedom in the space  $\mathbf{V}_h$  corresponds to a *subelement face*, and each degree of freedom in  $W_h$  corresponds to a *subelement volume*. The connectivity pattern of the  $\tilde{S}$  is determined by the topological connectivity of the subelement mesh. In particular, let  $\mathcal{F}(i)$  denote the set of Raviart–Thomas degrees of freedom  $k$  corresponding to subelement faces belonging to the subelement volume  $i$ . Given two subelement volumes (i.e., two  $L^2$  degrees of freedom) sharing a common face, let  $k_{ij}$  denote the degree of freedom corresponding to the shared face. Then, the entries of  $\tilde{S}$  are given by

$$(16) \quad \tilde{S}_{ij} = \begin{cases} \tilde{W}_{ii}^{-1} + \sum_{k \in \mathcal{F}(i)} \tilde{M}_{kk}^{-1} & \text{if } i = j, \\ -\tilde{M}_{kk}^{-1} & \text{if } i \text{ and } j \text{ share common face } k = k_{ij}, \\ 0 & \text{otherwise.} \end{cases}$$

Consequently, the matrix  $\tilde{S}$  is sparse, since the number of off-diagonal entries is bounded by the number of faces per element (4 for a quadrilateral mesh in 2D, 6 for a hexahedral mesh in 3D), independent of the polynomial degree of the finite element spaces. Additionally,  $D\tilde{M}^{-1}D^\top$  can be viewed as a simple discretization of the Laplacian, and the following arguments show that  $S$  is well-preconditioned by methods such as algebraic multigrid.

**Proposition 4.** *The approximate Schur complement  $\tilde{S}$  is an  $M$ -matrix.*

*Proof.* This follows from the explicit form (16) using that  $W_\alpha$  and  $M_\beta$  are positive-definite, and so  $\tilde{W}_{ii}^{-1} > 0$  and  $\tilde{M}_{kk}^{-1} > 0$ .  $\square$

Furthermore,  $\tilde{S}$  can be interpreted as a diagonal (approximate) mass matrix  $\tilde{W}$  plus a *weighted graph Laplacian* term  $D\tilde{M}^{-1}D^\top$ . The graph Laplacian  $L$  of an edge-weighted graph is defined by

$$(17) \quad L_{ij} = \begin{cases} w_i & \text{if } i = j, \\ -w_{ij} & \text{if } i \sim j, \\ 0 & \text{otherwise,} \end{cases}$$

where we write  $i \sim j$  if the graph vertices  $i$  and  $j$  are connected by an edge (with weight  $w_{ij}$ ). The vertex weight  $w_i$  is defined by  $w_i = \sum_{i \sim j} w_{ij}$ . Comparing (16) to (17), it is straightforward to see that  $D\tilde{M}^{-1}D^\top$  is the graph Laplacian corresponding to the weighted face-neighbor connectivity graph of the subelement mesh, with weights corresponding to the reciprocal of the diagonal entries of  $M_\beta$  (the diagonal entries of  $M_\beta$

scale like area of the corresponding subelement face). Note that this system is very closely related to the lowest-order discontinuous Galerkin discretization described in [44]. Algebraic multigrid preconditioners for linear systems of this form are well-studied, and can be expected to result uniform convergence independent of problem size. The following proposition summarizes the application of block-diagonal preconditioning to the grad-div system.

**Proposition 5.** *Let  $\tilde{M}$  denote the diagonal of  $M_\beta$  and let  $\tilde{S}$  denote a matrix spectrally equivalent to  $S$  (for example,  $\tilde{S}^{-1}$  could be given by one AMG V-cycle applied to  $\tilde{S}$ ). Then,  $B^{-1}\mathcal{A}$  is uniformly well conditioned, where  $B$  is the block diagonal preconditioner*

$$B = \begin{pmatrix} \tilde{M} & 0 \\ 0 & \tilde{S} \end{pmatrix}.$$

*Remark 4* (PDE-based Schur complement approaches). We note that instead of the algebraic approach described above, it is also possible to consider a PDE-based Schur complement, as in the framework of operator preconditioning [39]. In this context, the Schur complement associated with the block system (5) is a Laplacian operator defined on the  $L^2$  space. Preconditioners based on discretizing this operator using an interior penalty method have been considered in [48]. While this approach is robust with respect to mesh size, a standard symmetric interior penalty DG discretization for the Schur complement does not result in a uniform preconditioner with respect to the polynomial degree; the condition number of the preconditioned block system is observed to grow linearly with the polynomial degree. This is in contrast to the algebraic approach described above, which results in condition numbers that are asymptotically independent of  $p$ .

**2.4. Application to the Darcy system.** The treatment of the Darcy system (9) is largely similar to that of the grad-div system (7). The relevant difference is that in the case of the Darcy system, the off-diagonal blocks  $B$  and  $B^T$  are not weighted by the coefficient  $\gamma$ . As a result, the transformed system takes the slightly modified form

$$(18) \quad \mathcal{A}' = \begin{pmatrix} M_{1/\varepsilon} & D^T \\ D & -W^{-1}W_\gamma W^{-1} \end{pmatrix} = \begin{pmatrix} I & 0 \\ 0 & W^{-1} \end{pmatrix} \begin{pmatrix} M_{1/\varepsilon} & B^T \\ B & -W_\gamma \end{pmatrix} \begin{pmatrix} I & 0 \\ 0 & W^{-1} \end{pmatrix}$$

In general, the (2,2)-block of  $\mathcal{A}'$  requires two applications of  $W^{-1}$ . However, in the special case that  $\gamma v_h \in W_h$  for all  $v_h \in W_h$ , this block can be simplified to  $-W_{1/\gamma}^{-1}$ . This can be seen by noting that  $W^{-1}W_\gamma w$  (where the vector  $w$  corresponds to  $w_h \in W_h$ ) is the unique element  $u_h \in W_h$  satisfying

$$(u_h, v_h) = (\gamma w_h, v_h)$$

for all  $v_h \in W_h$ . Likewise,  $W_{1/\gamma}^{-1}Ww$  is the unique element  $u'_h \in W_h$  satisfying

$$(\gamma^{-1}u'_h, v_h) = (w_h, v_h)$$

for all  $v_h \in W_h$ . When  $\gamma v_h \in W_h$ , we have

$$(u_h, v_h) = (\gamma w_h, v_h) = (w_h, \gamma v_h) = (w_h, \gamma v_h) = (\gamma^{-1}u'_h, \gamma v_h) = (u'_h, v_h),$$

and hence  $W^{-1}W_\gamma = W_{1/\gamma}^{-1}W$ , and so  $W^{-1}W_\gamma W^{-1} = W_{1/\gamma}^{-1}WW^{-1} = W_{1/\gamma}^{-1}$ . In particular, this holds when the coefficient  $\gamma$  is piecewise constant (i.e.  $\gamma|_\kappa$  is constant for each  $\kappa \in \mathcal{T}$ ). In the general case, this simplification does not hold, and the full form of the (2,2)-block is used. In the case of the Darcy problem with zero (2,2)-block (corresponding to the Poisson problem  $-\nabla \cdot (\varepsilon \nabla p) = g$ ), the matrix  $W^{-1}$  is not required to compute the action of  $\mathcal{A}'$ , reducing the computational cost of the MINRES iterations.

As before, the  $\mathbf{H}(\text{div})$  mass matrix  $M_{1/\varepsilon}$  can be approximated by its diagonal, and  $W^{-1}W_\gamma W^{-1}$  can be approximated by the product of (the reciprocals of) the diagonals of the terms in the product. The approximate Schur complement  $\tilde{S}$  takes the same form (16) as in the preceding section, and can be effectively preconditioned with algebraic multigrid.

**2.5.  $L^2$  mass inverse.** Solving the transformed saddle-point system  $\mathcal{A}$  given by (11) by a Krylov subspace method such as MINRES (or GMRES in the case of a non-symmetric preconditioner) requires, at every iteration, the action of the operator  $W^{-1}$ . The matrix  $W_\alpha$  corresponds to the  $\alpha$ -weighted  $L^2$  inner product on the space  $W_h$ , i.e.,

$$r^T W_\alpha q = (\alpha q, r), \quad q, r \in W_h.$$

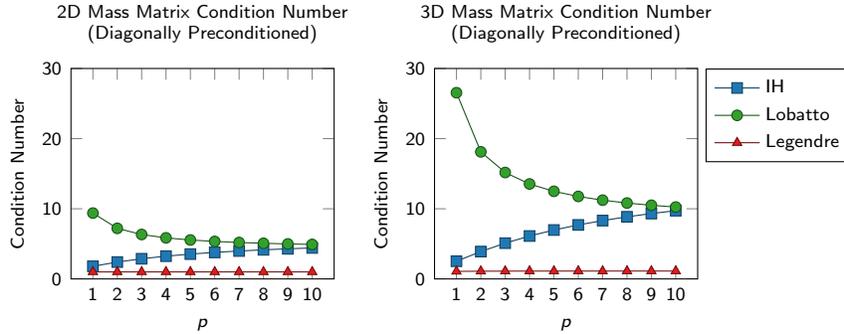


FIGURE 2. Condition number of the  $L^2$  mass matrix preconditioned by diagonal scaling for different choices of basis. Left: 2D case on a skewed quadrilateral. Right: 3D case on a skewed hexahedron.

Since the space  $W_h$  is discontinuous,  $W$  is block-diagonal, with blocks  $W_i$  corresponding to the  $i$ th element of the mesh  $\kappa_i$ . The block  $W_i$  represents the  $L^2$  inner product over  $\kappa_i$ ,

$$\mathbf{r}_i^\top W_i \mathbf{q}_i = (\alpha q, r)_{\kappa_i},$$

where  $\mathbf{q}_i$  and  $\mathbf{r}_i$  are subvectors of  $\mathbf{q}$  and  $\mathbf{r}$  consisting of those degrees of freedom belonging to the element  $\kappa_i$ . Hence, each block  $W_i$  is symmetric and positive-definite, and has size  $p^d \times p^d$ .

A significant amount of prior work on the topic of efficiently inverting the  $L^2$  mass matrix has been performed [21, 23, 34, 45]. For modest values of  $p$ , assembling and factorizing each of the blocks  $W_i$  (e.g. using Cholesky factorization or spectral decomposition) is an appropriate strategy, allowing for the application of  $W_\alpha^{-1}$  with roughly the same cost as standard operator evaluations or matrix-vector products. However, for higher-order problems, it becomes infeasible to assemble, store, and factorize the blocks  $W_i$ . Each block requires  $\mathcal{O}(p^{2d})$  storage, and the assembly and factorization require  $\mathcal{O}(p^{3d})$  operations. In this case, it is preferable to compute the action of  $W_\alpha^{-1}$  matrix-free using an iterative method.

In the matrix-free setting, the action of  $W_\alpha^{-1}$  is computed block-by-block using a local conjugate gradient (CG) iteration. Because of properties of the basis for the space  $W_h$ , each block  $W_i$  is spectrally equivalent to its diagonal, and so using  $\text{diag}(W_i)$  as a preconditioner will result in a number of CG iterations that is independent of the polynomial degree  $p$ . Furthermore, the action of  $W_i$  can be computed efficiently with matrix assembly using a sum factorization approach [41, 45]. Although this method is asymptotically optimal in terms of iterations (the cost of computing  $W_\alpha^{-1}$  is proportional to the cost of evaluating  $W_\alpha$ , independent of discretization parameters), the constants in the condition number estimates can be improved by using a change of basis.

In Figure 2 we show the condition number of the 2D and 3D  $L^2$  mass matrix on a single highly skewed element. In each case, we consider the diagonally scaled mass matrix using a different basis. In all cases, the condition number is bounded with respect to  $p$ . As predicted by analytical estimates (cf. [29, 49]) the condition number of the Gauss–Lobatto mass matrix *improves* with increasing  $p$ . The condition number obtained using the interpolation–histopolation basis displays a preasymptotic increase in condition number, and for larger  $p$  results in largely similar condition numbers as with the Gauss–Lobatto basis. The Gauss–Legendre nodal basis results in the best conditioned system, with maximal condition number of approximately 1.12 in these tests. Note that on elements whose transformation has constant Jacobian determinant (i.e. parallelepipeds) the Gauss–Legendre mass matrix is diagonal (and so in this special case, diagonal preconditioning is exact). However, in the cases we consider here, the elements are skewed, resulting in non-constant Jacobian determinants, and so the mass matrix is not diagonal with any of the bases considered. Even in this case, the diagonally preconditioned matrix remains extremely well conditioned, with condition number only slightly larger than unity.

Motivated by these results, faster CG convergence can be obtained by first performing a change of basis, and then iteratively solving the system in the Gauss–Legendre basis with diagonal preconditioning. Performing the change of basis necessitates the transformation of the right-hand side and the obtained solution. If a non-zero initial guess is used, one additional transformation must also be performed. Each of these

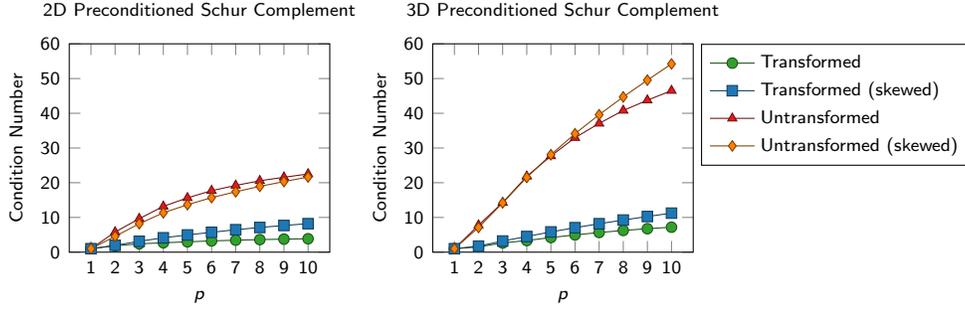


FIGURE 3. Computed condition numbers of preconditioner Schur complement. The transformed system is given by  $\tilde{S}^{-1}S$ , and the untransformed system is given by  $\tilde{S}'^{-1}S'$ .

transformations is generally less expensive than one operator evaluation, and so this will result in an overall efficiency gain if the total iteration count is reduced by at least two or three. Numerical experiments studying these considerations are presented in Section 3.1.

**2.6. Notes on the untransformed system.** Instead of considering the transformed system  $\mathcal{A}$  defined by (11), it is also possible to apply block preconditioners directly to the saddle-point system (7). This has the advantage that each MINRES iteration would require computing the action of  $W_\alpha$  rather than  $W_\alpha^{-1}$ , potentially leading to computational savings. However, this approach has two drawbacks. Firstly, the off-diagonal blocks  $B_\alpha$  and  $B_\alpha^\top$  do not enjoy the same favorable sparsity as  $D$  and  $D^\top$ , and therefore are more expensive to compute. Secondly, and more importantly, the Schur complement of the original system takes the form

$$(19) \quad S' = W_\alpha + B_\alpha M_\beta^{-1} B_\alpha^\top = W_\alpha + W_\alpha D M_\beta^{-1} D^\top W_\alpha.$$

In order to construct block preconditioners, approximations of the inverse of  $S'$  are required. As in the construction of the approximate Schur complement  $\tilde{S}$  in (15), it is possible to replace  $W_\alpha$  and  $M_\beta$  by their diagonals, to obtain a matrix  $\tilde{S}'$  spectrally equivalent to  $S'$ . This approximation of the second term on the right-hand side of (19) is significantly worse than the approximation of the transformed Schur complement  $S$  by (15).

In Figure 3, we display computed condition numbers of the transformed preconditioned Schur complement  $\tilde{S}^{-1}S$  and the untransformed version  $\tilde{S}'^{-1}S'$  for increasing polynomial degree. We consider both the unit cube  $[0, 1]^d$ , and a skewed single-element mesh (the same mesh used in the results shown in Figure 2). The empirical results show that the transformed Schur complement results in significantly smaller condition number, in accordance with arguments described above. These results justify the use of the transformed system (11), despite the increased cost per iteration incurred by the inverse  $L^2$  mass matrix.

### 3. ALGORITHMIC DETAILS AND GPU ACCELERATION

In this section, we give a brief overview of the solution algorithm and describe the strategies used for GPU acceleration. To solve the variational problem (6), we solve the saddle-point system

$$(20) \quad \begin{pmatrix} M_\beta & D^\top \\ D & -W_\alpha^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \tilde{\mathbf{q}} \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ 0 \end{pmatrix},$$

such that  $(\mathbf{u}, \mathbf{q})$  for  $\mathbf{q} = W_\alpha^{-1} \tilde{\mathbf{q}}$  gives the solution to the untransformed system (7). A schematic of the solution procedure is shown in Figure 4. This system is solved using MINRES (with the block-diagonal preconditioner) or GMRES (with block-triangular preconditioners). Each iteration requires application of the operators  $M_\beta$ ,  $D$ ,  $D^\top$  and  $W_\alpha^{-1}$ , and application of the preconditioners  $\tilde{M}^{-1}$  and  $\tilde{S}^{-1}$ . The high-order mass matrix  $M_\beta$  can be applied efficiently using matrix-free algorithms with sum factorization (cf. [1, 22]). The discrete divergence operator  $D$  can be represented explicitly as a sparse matrix (where the number of nonzeros per row remains bounded, independent of the polynomial degree). The efficient application of  $W_\alpha^{-1}$  with GPU acceleration is discussed in Section 3.1. The preconditioner  $\tilde{M}^{-1}$  for the (1,1)-block is simply

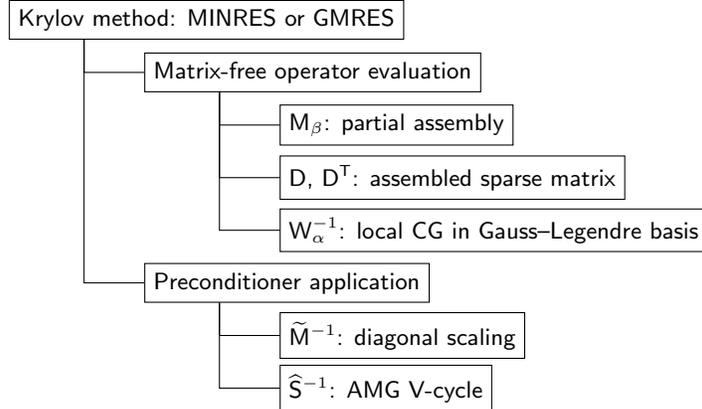


FIGURE 4. Solver diagram for block preconditioners applied to the transformed saddle-point system (20).

diagonal scaling, and is trivially parallelizable. In this work, we use algebraic multigrid for the approximate Schur complement preconditioner; details are discussed in Section 3.2.

**3.1. GPU acceleration of the DG mass inverse.** Because each iteration involving the system (20) requires computing the action of  $W_\alpha^{-1}$ , it is imperative to develop fast algorithms for the action of the inverse of the DG mass matrix. As described in Section 2.5, we consider several approaches, depending on the polynomial degree and spatial dimension. The methods considered here are fairly standard and well-known [21, 23, 34]. The empirical results of this section allow for the comparison of these methods on modern GPU-accelerated hardware, using high-performance implementations. For relatively small polynomial degrees, it is feasible to assemble and store the  $L^2$  mass matrix, allowing for the use of direct methods; for higher polynomial degrees, the memory requirements and computational time required to assemble the matrix render this option infeasible, necessitating the use of iterative solvers.

Since this matrix is block-diagonal, in the case where the matrix may be assembled, each block is stored as a dense matrix, which can then be factorized. These operations are particularly well-suited for use with batched linear algebra libraries [19]. The performance of batched LU solvers and matrix-matrix multiplication operations on GPUs has been well-studied [2, 50]. Presently, we consider two approaches. The first approach is to perform a batched Cholesky factorization (`potrf`) during the setup, and then each application of the mass inverse will perform a triangular solve (`potrs`). Batched versions of these routines are available on Nvidia and AMD GPU hardware through the `cuSOLVER` and `rocSOLVER` libraries. The second approach computes the explicit inverse of each block (`getrf` followed by `getri`), and then each application is computed using a batched matrix-vector product (`gemv`). While the computation of the inverse matrix necessitates additional overhead and is less numerically stable than solving the system using the Cholesky factorization, it has the advantage that the application of the inverse is computed using a matrix-vector product rather than a triangular solve, which is better-suited for fine-grained parallelism. For symmetric positive-definite matrices of small sizes, matrix inversion typically possesses satisfactory stability bounds [15]. Since direct methods are only relevant for small-sized matrices in this application, the numerical stability of explicit matrix inverses is not a primary concern.

For larger polynomial degrees, direct methods become impractical or infeasible. In this case, we perform conjugate gradient iterations locally on each element in parallel. Because the system is block-diagonal, each conjugate gradient solver is independent, and no global reductions are required. In our implementation, the entire CG iteration is performed in a single (fused) kernel, with one block of threads per element. The number of threads per block is chosen to facilitate the application of the local mass operator. This action corresponds to the “*CEED BP1*” benchmark problem, and highly optimized GPU implementations for this operation are available [2, 17]. This kernel also performs the initial change of basis operations, allowing the CG iteration to solve the better-conditioned Gauss-Legendre mass matrix, regardless of the user’s choice basis for the discretization (see Figure 2). The vector operations (`axpy`) in the CG method are local to

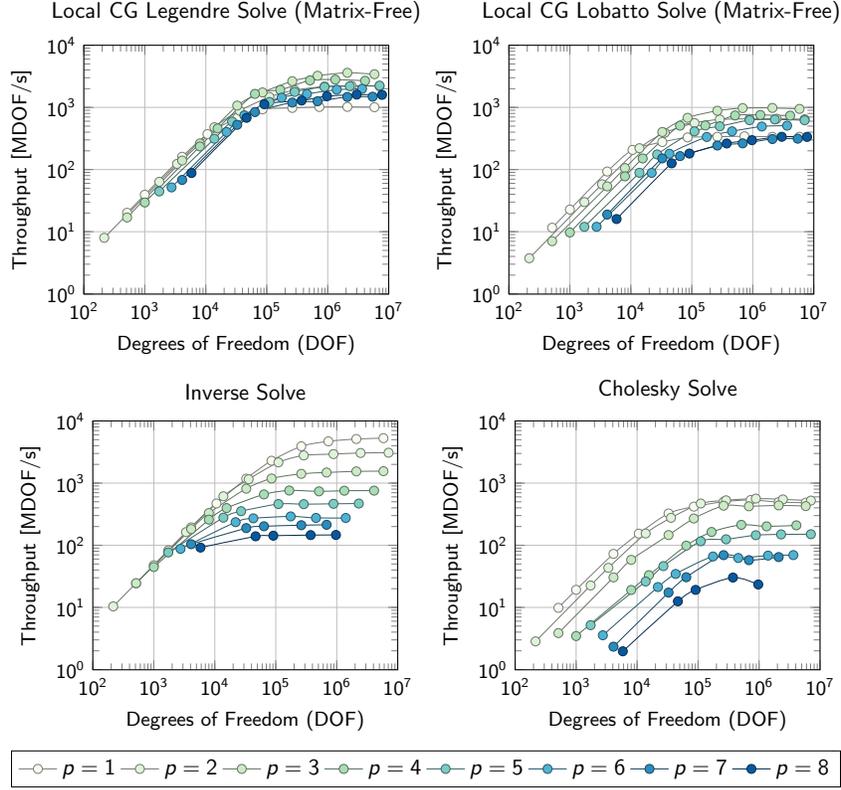


FIGURE 5. GPU DG mass inverse solve phase.

each element, and are trivially parallelizable. The dot products required by CG are performed using parallel reductions in each block.

In the context of the iterative solvers considered presently, the setup time is of lesser concern than the operator application time: the setup is performed once, while the action of  $W^{-1}$  must be computed at least once every iteration (and twice every iteration in the case of variable-coefficient Darcy problems, see Section 2.4). Figure 5 displays GPU kernel throughput for computing the action of  $W^{-1}$  using the methods described in this section. These results make clear that significantly higher performance is achieved in the CG solvers by first changing basis to the Gauss–Legendre basis, and then solving the better-conditioned system (cf. Figure 2). The matrix-free methods using the Gauss–Legendre basis consistently achieve over  $2 \times 10^9$  degrees of freedom per second, with a peak performance of  $3.6 \times 10^9$  at  $p = 3$ . In contrast, a standard CG iteration with global reductions and without fused kernels achieves peak performance about 2–3 orders of magnitude slower than the local CG approach described here.

Comparing the direct solvers, it is clear that using the precomputed explicit inverse matrix results in significantly faster performance than the Cholesky factorization. This is likely attributable to the increased fine-grain parallelism possible when performing batched matrix-vector products compared with triangular solves. For  $p = 1$  and  $p = 2$ , the explicit inverse results in the fastest solve time of all methods considered; the highest throughput observed was  $5.2 \times 10^9$  degrees of freedom per second, achieved for the largest problem size with  $p = 1$ . For  $p = 3$ , the explicit inverse has throughput roughly equal to that of the Gauss–Legendre CG iteration (with significantly more expensive setup costs), and for  $p \geq 4$ , the Gauss–Legendre CG iteration is the most performant method considered.

We conclude with a performance comparison on a series of problems with roughly  $1.7 \times 10^6$  DOFs, using polynomial degree  $p = 1, 2, \dots, 8$ . The mesh considered is Kershaw’s “Z” mesh with distortion parameter  $\epsilon = 0.5$ , cf. [28]. The mesh elements are non-affine, and the mass matrix in the Gauss–Legendre basis is not diagonal. In order to benchmark the solvers in a situation relevant to the saddle-point solvers considered in this paper, on each problem we apply the inverse mass matrix 100 times. In Table 1, we present the

TABLE 1. DG mass inverse setup and solve run times. Solve time is reported for 100 applications of the inverse. The fastest method for each polynomial degree is shown in bold. For  $p = 8$ , there was insufficient GPU memory to compute the explicit inverse, indicated by “—”.

		$p = 1$	$p = 2$	$p = 3$	$p = 4$	$p = 5$	$p = 6$	$p = 7$	$p = 8$
CG	Setup	<b>0.0019</b>	<b>0.0009</b>	<b>0.0007</b>	<b>0.0006</b>	<b>0.0005</b>	<b>0.0005</b>	<b>0.0006</b>	<b>0.0004</b>
CG (Legendre)	Its.	1	2	2	2	3	2	3	3
	Solve $\times 100$	0.2897	0.1333	<b>0.1148</b>	<b>0.1554</b>	<b>0.1320</b>	<b>0.1142</b>	<b>0.1659</b>	<b>0.1353</b>
	Total	0.2918	0.1342	<b>0.1155</b>	<b>0.1560</b>	<b>0.1326</b>	<b>0.1147</b>	<b>0.1664</b>	<b>0.1358</b>
CG (Lobatto)	Its.	5	7	9	10	12	13	13	11
	Solve $\times 100$	0.4401	0.2184	0.1929	0.3260	0.2682	0.2915	0.3375	0.3378
	Total	0.4421	0.2193	0.1936	0.3266	0.2687	0.2921	0.3381	0.3382
Inverse	Setup	0.0039	0.0180	0.2276	1.0510	0.4662	1.2576	1.9993	—
	Solve $\times 100$	<b>0.0380</b>	<b>0.0602</b>	0.1154	0.2299	0.3688	0.6093	0.8066	—
	Total	<b>0.0420</b>	<b>0.0781</b>	0.3430	1.2809	0.8350	1.8669	2.8059	—
Cholesky	Setup	0.0046	0.0110	0.2034	0.9056	0.1459	0.3543	0.7344	12.8996
	Solve $\times 100$	0.3482	0.3188	0.4001	0.8125	1.1600	2.3327	2.5880	5.9370
	Total	0.3528	0.3299	0.6035	1.7181	1.3059	2.6870	3.3224	18.8365

time required to construct the inverse (or set up the matrix-free CG), the time required for 100 inverse applications, and the total time. The shortest runtime for each polynomial degree is shown in bold. The matrix-free CG method has the fastest setup time among the methods considered, with roughly constant time required for all polynomial degrees. The explicit inverse resulted in the fastest overall runtimes for  $p = 1$  ( $4.3\times$  faster than CG) and  $p = 2$  ( $1.1\times$  faster than CG). For polynomial degrees  $p \geq 3$ , the matrix-free CG (changing to Legendre basis) resulted in the fastest runtime; the observed speed-up compared with direct methods increases dramatically with the polynomial degree. The average number of iterations required per element-local CG iteration is listed for Legendre and Lobatto CG. For  $p = 8$ , there was insufficient memory to compute the explicit inverse. For this test case, the CG setup time was  $30,000\times$  faster than the Cholesky factorization, the solve time was  $70\times$  faster than the triangular solve, resulting in an overall speed-up of  $216\times$ . The Cholesky factorization was the slowest of the three methods considered; the slightly faster factorization time compared with the explicit inverse was outweighed by the slower inverse applications.

Based on these results, in the contexts where the inverse of a fixed mass matrix will be applied many times (such as the iterative solvers considered presently, or in the case of time stepping on a fixed mesh), it would be efficient to use explicit inverses for  $p \leq 2$ , and matrix-free CG for  $p \geq 3$ . If, on the other hand, the inverse of the mass matrix will be applied only once (for example, in situations involving a moving mesh, necessitating the recomputation of the mass matrix at every time step), the faster setup time suggests that the matrix-free CG iteration should be used even for  $p \leq 2$ .

**3.2. Schur complement preconditioning.** We now consider GPU-accelerated preconditioning of the approximate Schur complement  $\tilde{\mathbf{S}} = \tilde{\mathbf{W}}^{-1} + \mathbf{D}\tilde{\mathbf{M}}^{-1}\mathbf{D}^\top$ , where we recall that  $\tilde{\mathbf{W}}$  is the diagonal of the  $L^2$  mass matrix,  $\tilde{\mathbf{M}}$  is the diagonal of the  $\mathbf{H}(\text{div})$  mass matrix, and  $\mathbf{D}$  is the discrete divergence matrix. The approach taken presently is algebraic multigrid preconditioning, and so the first step is the assembly of the matrix. The diagonals  $\tilde{\mathbf{W}}$  and  $\tilde{\mathbf{M}}$  can be assembled efficiently using sum factorization [47]. Because of the properties of the interpolation–histopolation basis, the discrete divergence matrix  $\mathbf{D}$  has the same structure as the lowest-order divergence matrix on a refined mesh. In particular, the entries  $\mathbf{D}_{ij}$  are given by

$$\mathbf{D}_{ij} = \begin{cases} \sigma_{ij} & \text{if } j \in \mathcal{F}(i), \\ 0 & \text{otherwise,} \end{cases}$$

where  $\mathcal{F}(i)$  denotes the set of subelement faces adjacent to the subelement volume  $i$  (cf. Section 2.3), and  $\sigma_{ij} = \pm 1$  denotes the orientation of the Raviart-Thomas basis function on subelement face  $j$  relative to

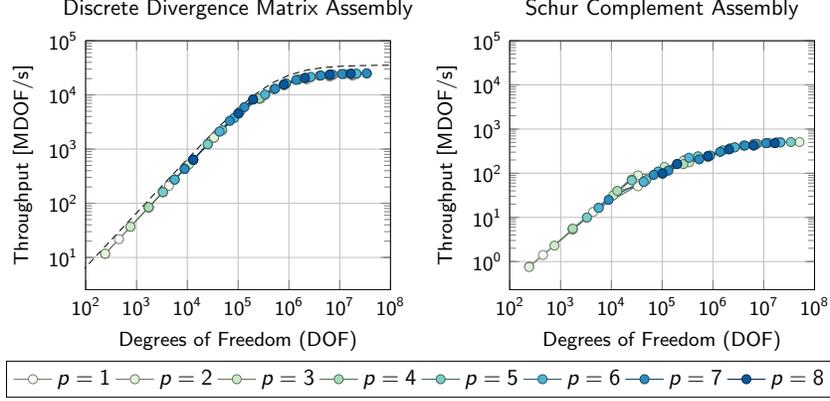


FIGURE 6. Throughput for discrete divergence and approximate Schur complement assembly. The dashed line on the left plot indicates the theoretical peak throughput  $n/t(n)$  based on the performance model (21).

subelement volume  $i$ . This matrix can be constructed directly in CSR format efficiently and in parallel on the GPU; one thread is used per nonzero in the matrix. The algorithm used to construct this matrix is given in Algorithm 1.

The kernel throughput of this algorithm is shown in the left panel of Figure 6. We construct a simple performance model to compare the achieved throughput with the theoretical maximum. For a problem of size  $N$ , the best theoretical runtime based solely on memory transfer can be modeled by

$$(21) \quad t(N) = \frac{DN}{B} + L,$$

where  $D$  is the data transfer required per DOF,  $B$  is the memory bandwidth, and  $L$  is the kernel launch latency. In our model, we choose these parameters based on the specifications of the V100 GPU. The memory bandwidth is given by  $B = 898$  GB/s, and we estimate the kernel launch latency to be  $L = 15$  ns. A CSR-format matrix with the given sparsity pattern requires approximately  $D = 25\frac{1}{3}$  bytes per  $\mathbf{H}(\text{div})$  DOF (note that the algorithm presented in Algorithm 1 requires an additional four-byte global read per DOF to obtain the local DOF index and orientation). Comparing the achieved throughput to the model given by (21), we see that this kernel achieves close to the theoretical peak performance for all considered problem sizes.

---

**Algorithm 1** Construction of the discrete divergence matrix  $D$  in CSR format.

---

$N \leftarrow \#$  of  $L^2$  DOFs

**parallel for**  $i \in \{0, \dots, N-1\}$ ,  $k \in \{1, \dots, 2d\}$  **do**  $\triangleright$  For each  $L^2$  DOF (subelement volume)  $i$ , and adjacent subelement face  $k$

$I[i] \leftarrow 2di$   $\triangleright$   $2d$  nonzeros per row

$e \leftarrow \text{element\_map}[i]$   $\triangleright$   $e$  is the mesh element containing DOF  $i$

$i_{loc} \leftarrow \text{l2\_global\_to\_local}[i]$   $\triangleright$   $i_{loc}$  is the local  $L^2$  DOF index of  $i$  within  $e$

$j_{loc} \leftarrow \text{volume\_to\_face}[k, i_{loc}]$   $\triangleright$   $j_{loc}$  is the local RT DOF index of the subelement face  $k$

$\sigma_{loc} \leftarrow \text{local\_orientation}[j_{loc}]$   $\triangleright$   $\sigma_{loc}$  is the orientation of DOF  $j_{loc}$  in the reference element

$j \leftarrow \text{rt\_local\_to\_global}[e, j_{loc}]$   $\triangleright$   $j$  is the global DOF index

$\sigma_{glob} \leftarrow \text{global\_orientation}[j]$   $\triangleright$   $\sigma_{loc}$  is the orientation of the  $j$ th DOF relative to  $j_{loc}$

$J[2di + k] \leftarrow j$

$A[2di + k] \leftarrow \sigma_{loc}\sigma_{glob}$

**end**

$I[N] \leftarrow 2dN$

$\triangleright$  The matrix  $D$  in CSR format is given by the arrays  $(I, J, A)$

---

After constructing the discrete divergence, the approximate Schur complement can be constructed by computing the sparse matrix-matrix product  $\tilde{D}\mathbf{M}^{-1}\mathbf{D}^T$ ; other approaches, including custom assembly kernels,

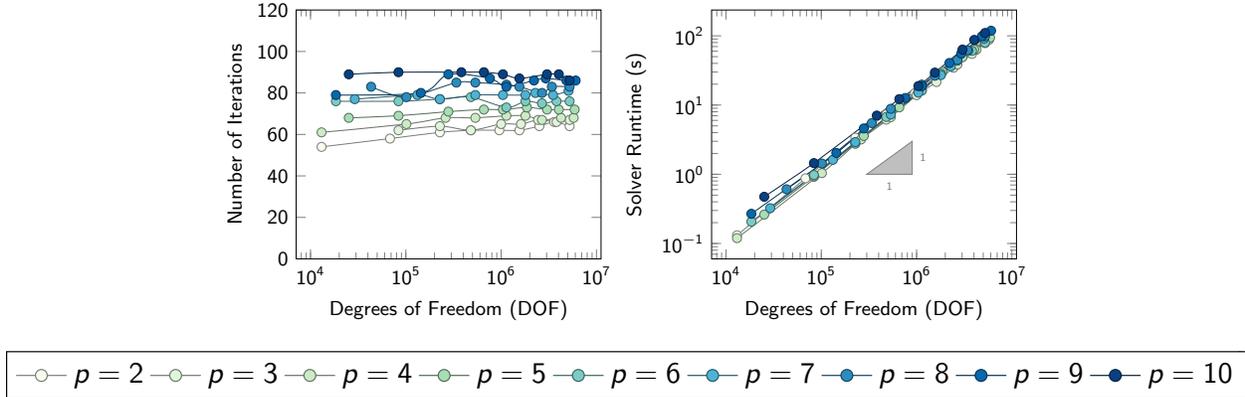


FIGURE 7. Runtime and iteration counts for the Cartesian grid test case under  $h$ - and  $p$ -refinement.

are also possible. In the current implementation, this product is computed using the sparse triple-product algorithms provided by the *hypre* solvers library [20]. The throughput for the triple product operation is shown in the right panel of Figure 6. Given the assembled Schur complement in CSR format, the approximate inverse  $\widehat{S}^{-1}$  is given by one V-cycle of an algebraic multigrid preconditioner; in this work, we use the BoomerAMG preconditioner from *hypre* [26]. Detailed GPU performance studies for AMG preconditioners applied to similar problems are given in [43].

#### 4. NUMERICAL RESULTS

In the following sections, we test the performance of the saddle-point solvers on a number of grad-div and Darcy test problems. GPU results were performed on LLNL’s *Lassen* supercomputer, with four Nvidia V100 GPUs per node. The methods were implemented in the framework of the open-source MFEM finite element library [4], see also <https://mfem.org>. The solvers described in this paper are open-source, and have been released as part of the MFEM library; full source code and examples are available in MFEM’s GitHub repository at <https://github.com/mfem/mfem>. A relative tolerance of  $10^{-12}$  was used as the stopping criterion for all iterative solvers, unless specifically noted otherwise. The BoomerAMG algebraic multigrid preconditioner from the *hypre* library is used to precondition the approximate Schur complement [26], and *hypre*’s ADS preconditioner applied to the low-order-refined system is used as a comparison [32]. The AMG preconditioner for the Schur complement uses parallel maximal independent set (PMIS) coarsening with no aggressive coarsening levels.

**4.1. Cartesian grid.** As a first numerical example, we consider the simplest case of a constant-coefficient Darcy problem on a three-dimensional regular Cartesian grid. The exact solution and right-hand side are chosen according to the solver benchmark test-case proposed in [30]. The saddle-point system is discretized as described above, and solved using MINRES with the block-diagonal preconditioner outlined in Figure 4. A relative tolerance of  $10^{-12}$  was used for this problem. We consider both  $h$ - and  $p$ -refinements, with  $p = 2, 3, \dots, 10$ . The results are presented in Figure 7.

This example illustrates robust convergence typical of multigrid-type methods. The iteration counts are roughly constant under  $h$ -refinement, and display only relatively modest growth with increasing polynomial degree. Furthermore, the total solver runtime (including both the setup and solve phases) is observed to scale as  $\mathcal{O}(N)$ , where  $N$  is the number of degrees of freedom; for a given problem size, only relatively minor increases in runtime were observed for increasing polynomial degree.

We also use this simple test case to measure the runtimes of the individual components of the solver algorithm. In Figure 8, we display the runtimes for each of the components; we run both on one V100 GPU and four Power 9 CPU cores, and report both sets of timings. (Note that these times are not intended to compare GPU vs. CPU performance, but rather to investigate the relative weighting of each of the algorithm components on different architectures). These runtimes were measured for  $2 \leq p \leq 7$ , on the largest mesh considered above (corresponding to approximately  $5 \times 10^6$  degrees of freedom). On both architectures, the relative runtimes are fairly similar; solver and AMG setup represents slightly under half of the overall

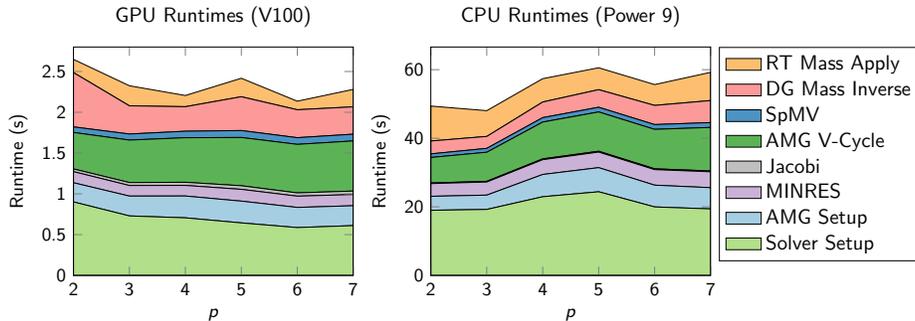


FIGURE 8. Runtimes for solver components (cf. Figure 4) for the Cartesian grid test case. Left panel: V100 GPU. Right panel: Four Power 9 CPU cores.

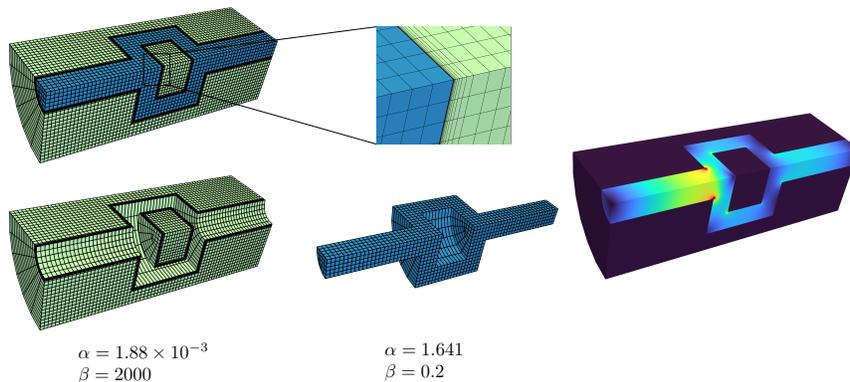


FIGURE 9. Crooked pipe grad-div problem (figure reproduced from [44]). Left panel: mesh with inset showing anisotropic elements and definition of piecewise constant coefficients on subdomains. Right panel: illustration of solution with constant forcing term.

runtime. On CPU and GPU, the AMG V-cycle for the approximate Schur complement represents the largest per-iteration time at high orders.

**4.2. Grad-div problem: crooked pipe.** In this section, we consider the “crooked-pipe” grad-div problem, which is a more realistic and challenging test case. This problem, which serves as a model problem related radiation diffusion simulations, is posed on a cylindrical sector. The elements are divided into two materials; the interface between the materials consists of highly stretched, anisotropic elements, increasing the problem difficulty. We solve the problem

$$\nabla(\alpha \nabla \cdot \mathbf{u}) - \beta \mathbf{u} = \mathbf{f},$$

where the coefficients  $\alpha$  and  $\beta$  are piecewise constants determined by the materials. A schematic of this problem is shown in Figure 9. In the outer region (colored green in Figure 9), we have  $\alpha = 1.88 \times 10^{-3}$  and  $\beta = 2000$ , and in the inner region (colored blue in Figure 9), we have  $\alpha = 1.641$  and  $\beta = 0.2$ . This problem, first proposed in [25], was considered as a benchmark problem using the ADS solver in [32], using algebraic hybridization in [16], and with low-order-refined preconditioning in [44].

We use the same problem setup as in [44], and solve this problem with increasing polynomial degree. We compare the performance of the saddle-point solvers discussed presently to applying the ADS preconditioner directly to the high-order assembled matrix, and to ADS applied to the spectrally equivalent low-order-refined system. We first compare CPU results computed using 144 MPI ranks on four nodes of LLNL’s *Quartz* supercomputer, which is the same configuration considered in [44]. The results are shown in Table 2. Both methods display a pre-asymptotic increase in the number of iterations with increasing polynomial degree. At  $p = 6$ , the saddle-point solver required  $1.9\times$  as many iterations as at  $p = 2$ , and the LOR-ADS solver required  $2.5\times$  as many iterations. Overall, the saddle-point solver required approximately  $1.5\times$  as many iterations as the LOR-ADS solver, however it was faster overall by a factor of  $6\text{--}9\times$ . The reason

TABLE 2. CPU results from the crooked pipe test problem (144 CPU cores on LLNL’s *Quartz*).

$p$	Saddle-Point		LOR-ADS		Speed-up	# DOFs
	Its.	Time (s)	Its.	Time (s)		
2	168	0.26	96	1.73	$6.65\times$	356,500
3	231	0.87	143	5.73	$6.59\times$	1,190,115
4	269	2.11	172	14.91	$7.07\times$	2,805,520
5	299	4.45	206	36.90	$8.29\times$	5,461,375
6	323	8.73	241	80.84	$9.26\times$	9,416,340

TABLE 3. GPU results from the crooked pipe test problem (4 Nvidia V100 GPUs on LLNL’s *Lassen*).

$p$	Saddle-Point		LOR-ADS		Speed-up	# DOFs
	Its.	Time (s)	Its.	Time (s)		
2	193	0.85	143	7.73	$9.09\times$	356,500
3	251	1.52	206	14.97	$9.85\times$	1,190,115
4	298	2.19	269	23.72	$10.83\times$	2,805,520
5	335	3.49	332	37.60	$10.77\times$	5,461,375
6	360	4.01	392	61.94	$15.45\times$	9,416,340

TABLE 4. Matrix-based methods: CPU results from the crooked pipe test problem (144 CPU cores on LLNL’s *Quartz*).

$p$	Hybridization		ADS		# DOFs
	Its.	Time (s)	Its.	Time (s)	
2	28	0.16	74	2.49	356,500
3	35	0.86	85	19.20	1,190,115
4	36	5.73	92	100.77	2,805,520
5	31	29.82	102	390.32	5,461,375
6	36	245.54	112	1293.12	9,416,340

for this speed-up (despite a larger number of iterations) is that the saddle-point preconditioner is much less computationally expensive than the ADS preconditioner: each application of the saddle-point preconditioner requires just one diagonal scaling and one scalar AMG V-cycle (applied to the smaller  $L^2$  system), whereas each application of ADS requires (depending on the specific solver settings) around 11 AMG V-cycles.

These tests are repeated on the GPU using one node of LLNL’s *Lassen* with 4 Nvidia V100 GPUs. The results are shown in Table 3. The iteration counts for the GPU test are larger than those obtained on the CPU. This is because different BoomerAMG settings are used when *hypre* is run on the CPU or GPU [20]. In particular, hybrid Gauss-Seidel smoothing is used on the CPU, while  $\ell_1$ -Jacobi smoothing is used on the GPU. This appears to have a more significant detrimental effect on the ADS solver; the iteration counts for the saddle-point solver are largely similar to the CPU case. The saddle-point solver consistently achieves 9–15 $\times$  speed-up compared with the LOR-ADS solver, with greater speed-up at higher orders. At  $p = 6$ , the total solve time using the saddle-point solver using the GPU on one node of *Lassen* is over twice as fast as on the CPU using 4 nodes of *Quartz*.

For comparison, we also include CPU results for the matrix-based algebraic hybridization and ADS solvers in Table 4; GPU results are not reported because GPU-accelerated versions of these algorithms are not available. Consistent with the results reported in [16], the algebraic hybridization solver greatly outperforms the ADS solver on all cases considered. Similarly, the LOR-ADS solver outperforms the ADS solver for  $p \geq 2$ , as reported in [44], and consequently the saddle-point solver also compares favorably to matrix-based ADS for these test cases. The comparison with algebraic hybridization is more interesting. Our results are largely consistent with previous comparisons of matrix-based and matrix-free solvers (including hybridizable discontinuous Galerkin methods), see e.g. [36]. The hybridization solver (which performs a conjugate gradient iteration directly on the hybridized Schur complement system) requires significantly fewer iterations than



FIGURE 10. SPE10 problem, logarithmic color scale (cf. [14]). Left: magnitude of inverse permeability coefficient. Right: solution velocity magnitude.

either ADS or the saddle-point solver. For  $p = 2$ , hybridization results in the fastest overall runtime, and for  $p = 3$ , we obtained roughly equal runtimes with the hybridization and saddle-point solvers. However, for  $p > 3$ , the time required to assemble and set up the hybridized solver becomes prohibitive, and the saddle-point solver results in significant speed-ups. For  $p = 4$ , the saddle-point solver is almost twice as fast, and for  $p = 6$ , the saddle-point solver results in an almost  $30\times$  speedup. We conclude that for relatively low orders, when assembling and forming the hybridized Schur complement system is not prohibitively expensive, algebraic hybridization is a highly competitive option. For higher orders, matrix-free methods such as the saddle-point solver discussed presently become necessary, and we note that the benefits of the matrix-free methods, such as reduced memory footprint, are more pronounced on GPU-based architectures.

**4.3. Darcy flow: SPE10 benchmark.** In this section, we consider the SPE10 benchmark problem from the reservoir simulation community [14]. This problem has been widely used as a test case for solvers in the literature, see e.g. [16, 27]. The problem geometry is  $\Omega = [0, 1200] \times [0, 2200] \times [0, 170]$ , represented as a Cartesian mesh with  $60 \times 220 \times 85$  elements. The permeability tensor is defined as a piecewise constant on this mesh. This coefficient is highly heterogeneous, exhibiting contrast of over  $10^7$ ; see Figure 10. On coarser meshes, a coarsened piecewise constant coefficient is obtained by a simple pointwise sampling procedure at element centers. We solve Darcy's equations for the velocity and pressure,

$$\begin{cases} \kappa \mathbf{u} + \nabla p = 0, \\ \nabla \cdot \mathbf{u} = 0, \end{cases} \quad \text{in } \Omega, \\ \mathbf{u} \cdot \mathbf{n} = (1, 0, 0)^\top \cdot \mathbf{n} \quad \text{on } \partial\Omega,$$

where  $\kappa$  is inverse permeability matrix. Since pure Neumann boundary conditions are enforced for this problem, the Schur complement is singular, with a nullspace consisting of all constant functions. To ensure a convergent solver, every application of the approximate inverse  $\widehat{S}^{-1}$  is followed by an orthogonalization step.

We use this problem to perform a weak scaling study using  $p = 4$  and  $p = 6$  elements. With each refinement, the number of degrees of freedom is roughly doubled by solving on a finer mesh, and the problem is solved using twice the number of GPUs. With increasing mesh refinement the iteration counts remain roughly constant. The largest problem has roughly  $7 \times 10^8$  unknowns, and is solved in about 10 seconds on 256 GPUs. For each of the cases considered, the solve time and iterations are similar between the  $p = 4$  and  $p = 6$  cases, illustrating the uniform convergence of the solver with respect to polynomial degree. The parallel scalability of the solver is largely consistent with similar results reported in [31, 43]. For the largest problem, the weak parallel efficiency of the solver is 31%, comparing favorably to the parallel efficiency of 25% for the algebraic hybridization solver and 7.5% for the ADS solver reported for similarly sized SPE10 problems in [16].

TABLE 5. Scalability study for the SPE10 benchmark problem.

$p = 4$				$p = 6$			
# GPUs	# DOFs	Its.	Time (s)	# GPUs	# DOFs	Its.	Time (s)
1	$3.36 \times 10^6$	147	3.25	1	$3.30 \times 10^6$	165	3.14
2	$5.76 \times 10^6$	155	4.31	2	$5.87 \times 10^6$	166	4.36
4	$1.20 \times 10^7$	157	5.15	4	$1.13 \times 10^7$	166	4.82
8	$2.68 \times 10^7$	160	6.45	8	$2.63 \times 10^7$	165	6.41
16	$4.58 \times 10^7$	159	6.35	16	$4.67 \times 10^7$	165	6.19
32	$9.53 \times 10^7$	164	6.94	32	$9.01 \times 10^7$	171	7.20
64	$2.13 \times 10^8$	164	8.62	64	$2.10 \times 10^8$	176	8.97
128	$3.66 \times 10^8$	166	8.55	128	$3.73 \times 10^8$	173	8.79
256	$7.61 \times 10^8$	165	10.44	256	$7.20 \times 10^8$	175	10.53

4.4. **Nonlinear radiation diffusion.** As a final test case, we consider the solution of time-dependent grey (one-group) nonlinear radiation diffusion equations. The governing equations are

$$(22) \quad \rho \frac{\partial e}{\partial t} = -c\sigma(aT_{\text{mat}}^4 - E) + Q,$$

$$(23) \quad \frac{\partial E}{\partial t} + \nabla \cdot F = c\sigma(aT_{\text{mat}}^4 - E) + S,$$

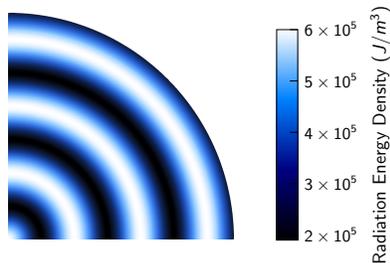
$$(24) \quad F = -\frac{c}{3\sigma} \nabla E,$$

where  $e$  is the material specific internal energy,  $E$  is the radiation energy density, and  $F$  is the radiation flux,  $\sigma$  is the absorption opacity,  $c$  is the speed of light,  $a$  is the black body constant, and  $\rho$  is the material density.  $Q$  and  $S$  are given source terms. The material temperature  $T_{\text{mat}}$  is related to the internal energy  $e$  through the specific heat  $C_v$ . The energy unknowns  $e$  and  $E$  are taken in the  $L^2$  finite element space  $W_h$ , and the radiation flux is taken in the  $\mathbf{H}(\text{div})$  space  $\mathbf{V}_h$ .

For this test case, we use the manufactured solution proposed in [12]. The equations are integrated in time implicitly; we use the three-stage  $L$ -stable singly diagonal implicit Runge–Kutta (SDIRK) method [3]. The resulting nonlinear algebraic equations are solved using an iterative scheme as follows. First, the radiation flux  $F$  is lagged, and a coupled nonlinear system for the material energy  $e$  and radiation energy  $E$  is solved. Then, the material energy  $e$  is fixed, and a Darcy-type linear system is solved for the radiation energy  $E$  and radiation flux  $F$ . This procedure is repeated until the full nonlinear residual has converged to the desired tolerance; for the time-accurate time steps considered presently, this method typically requires only one or two outer iterations to reach a relative tolerance of  $10^{-12}$ .

For fixed radiation flux  $F$ , the (inner) nonlinear equations corresponding to (22) and (23) are solved element-by-element using Newton’s method. The resulting Jacobian matrix takes the form of a  $2 \times 2$  element-wise block diagonal system, which is well-preconditioned by a block Jacobi preconditioner with blocks of size  $2 \times 2$ . After solving for the energies  $e$  and  $E$ , the material energy  $e$  is fixed, and the coupled linear equations for  $E$  and  $F$  corresponding to (23) and (24) are solved; the solution of this system is often a bottleneck in radiation hydrodynamics simulations, emphasizing the importance of the development of efficient preconditioners [24, 37]. These equations take the form of a Darcy-type problem, which can be preconditioned using the block preconditioners described in this work.

The test problem in [12] is posed in spherical coordinates; presently, we solve this problem in Cartesian coordinates on a three-dimensional mesh with 884,736 hexahedral elements using polynomial degree  $p = 4$ . This problem configuration has  $1.7 \times 10^8$  Raviart–Thomas degrees of freedom and  $5.7 \times 10^7$   $L^2$  degrees of freedom, and is solved in parallel with 16 Nvidia V100 GPUs on four nodes of LLNL’s *Lassen*. The equations are integrated until a final time of  $0.1/\tau \approx 4.39 \times 10^{-11}$  s, where  $\tau \approx 2.28 \times 10^9$  s $^{-1}$  is the decay constant of the test problem. The time step is chosen to be  $\Delta t = h/\tau$ , and 10 implicit time steps (each with three Runge–Kutta stages) are taken. Each stage requires an outer nonlinear solve (with relative tolerance  $10^{-12}$ ), an inner nonlinear solve (with relative tolerance  $10^{-8}$ ), and a linear solver (with relative tolerance  $10^{-8}$ ). At the final time, the relative  $L^2$  error was  $7.72 \times 10^{-8}$  in the material energy and  $5.25 \times 10^{-7}$  in the radiation energy. Each Runge–Kutta stage required two outer solver iterations and four inner Newton iterations. In



Alg. Component	Time (s)	%
Nonlinear residual	8.22	4.45%
Inner Newton solver (2 iterations per solve)	7.15	3.87%
Linear solver (84 iterations per solve)	169.39	91.68%

FIGURE 11. Left: snapshot of radiation energy density from radiation diffusion test case. Right: timing results and iteration counts on 16 V100 GPUs.

total, 60 saddle-point linear solves were performed, and on average the linear solver converged in 84 MINRES iterations. The timing results and iteration counts are presented in the right panel of Figure 11. Consistent with results reported in [24], approximately 92% of the total runtime for this problem was spent in the linear solver.

## 5. CONCLUSIONS

In this work, we have developed solvers for grad-div and Darcy-type problems in  $\mathbf{H}(\text{div})$  based on block preconditioning for the associated saddle-point system. Properties of the interpolation–histopolation high-order basis result in off-diagonal blocks with enhanced sparsity. The Schur complement of the block system can be approximated by an M-matrix with sparsity independent of the polynomial degree; such systems are amenable to preconditioning with scalable algebraic multigrid methods. The  $(1, 1)$ -block can be well-approximated by a diagonal matrix. The resulting block preconditioners result in uniform convergence, independent of mesh size and polynomial degree. These solvers are amenable to GPU acceleration; high-throughput algorithms are developed and benchmarked. The performance of the solvers is demonstrated on several benchmark problems, including the “crooked-pipe” grad-div problem, the SPE10 benchmark from reservoir simulation, and nonlinear radiation-diffusion.

## 6. ACKNOWLEDGEMENTS

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-JRNL-848035). W. Pazner and Tz. Kolev were partially supported by the LLNL-LDRD Program under Project No. 20-ERD-002. W. Pazner was also partially supported the ORAU Ralph E. Powe Junior Enhancement Award, 2023. The authors thank the anonymous reviewers for their constructive comments.

## REFERENCES

- [1] A. Abdelfattah et al. “GPU algorithms for efficient exascale discretizations”. In: *Parallel Computing* 108 (2021), p. 102841. DOI: 10.1016/j.parco.2021.102841.
- [2] A. Abdelfattah et al. “Performance, Design, and Autotuning of Batched GEMM for GPUs”. In: *Lecture Notes in Computer Science*. Springer International Publishing, 2016, pp. 21–38.
- [3] R. Alexander. “Diagonally Implicit Runge–Kutta Methods for Stiff ODEs”. In: *SIAM Journal on Numerical Analysis* 14.6 (1977), pp. 1006–1021. DOI: 10.1137/0714068.
- [4] R. Anderson et al. “MFEM: a modular finite element methods library”. In: *Computers & Mathematics with Applications* (2020). DOI: 10.1016/j.camwa.2020.06.009.
- [5] D. N. Arnold, R. S. Falk, and R. Winther. “Preconditioning in  $H(\text{div})$  and applications”. In: *Mathematics of Computation* 66.219 (1997), pp. 957–985. DOI: 10.1090/s0025-5718-97-00826-0.
- [6] D. N. Arnold, R. S. Falk, and R. Winther. “Multigrid in  $H(\text{div})$  and  $H(\text{curl})$ ”. In: *Numerische Mathematik* 85.2 (2000), pp. 197–217. DOI: 10.1007/p100005386.
- [7] E. Aulisa, S. Calandrini, and G. Capodaglio. “FOV-equivalent block triangular preconditioners for generalized saddle-point problems”. In: *Applied Mathematics Letters* 75 (2018), pp. 43–49.

- [8] M. Benzi, G. H. Golub, and J. Liesen. “Numerical solution of saddle point problems”. In: *Acta Numerica* 14 (2005), pp. 1–137.
- [9] D. Boffi, F. Brezzi, and M. Fortin. *Mixed Finite Element Methods and Applications*. Vol. 44. Springer Series in Computational Mathematics. Springer Berlin Heidelberg, 2013.
- [10] J. H. Bramble and J. E. Pasciak. “A preconditioning technique for indefinite systems resulting from mixed approximations of elliptic problems”. In: *Mathematics of Computation* 50.181 (1988), pp. 1–17.
- [11] P. D. Brubeck and P. E. Farrell. “Multigrid solvers for the de Rham complex with optimal complexity in polynomial degree”. arXiv eprint 2211.14284. 2022.
- [12] T. A. Brunner. *Development of a grey nonlinear thermal radiation diffusion verification problem*. Tech. rep. SAND2006-4030C. Sandia National Laboratories, 2006.
- [13] T. A. Brunner et al. “Preserving spherical symmetry in axisymmetric coordinates for diffusion problems”. In: *Proceedings of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2013)*. 2013.
- [14] M. A. Christie and M. J. Blunt. “Tenth SPE Comparative Solution Project: A Comparison of Upscaling Techniques”. In: *SPE Reservoir Evaluation & Engineering* 4.04 (2001), pp. 308–317. DOI: 10.2118/72469-pa.
- [15] J. J. D. Croz and N. J. Higham. “Stability of methods for matrix inversion”. In: *IMA Journal of Numerical Analysis* 12.1 (1992), pp. 1–19. DOI: 10.1093/imanum/12.1.1.
- [16] V. Dobrev et al. “Algebraic hybridization and static condensation with application to scalable  $H(\text{div})$  preconditioning”. In: *SIAM Journal on Scientific Computing* 41.3 (2019), B425–B447. DOI: 10.1137/17M1132562.
- [17] V. Dobrev et al. *ECP Milestone Report: Identify initial kernels, bake-off problems (benchmarks) and miniapps (CEED-MS6)*. Tech. rep. LLNL-TR-734169. U.S. Department of Energy, 2017. DOI: 10.5281/zenodo.2542333.
- [18] C. R. Dohrmann. “Spectral equivalence of low-order discretizations for high-order  $H(\text{curl})$  and  $H(\text{div})$  spaces”. In: *SIAM Journal on Scientific Computing* 43.6 (2021), A3992–A4014.
- [19] J. Dongarra et al. “The Design and Performance of Batched BLAS on Modern High-Performance Computing Systems”. In: *Procedia Computer Science* 108 (2017), pp. 495–504. DOI: 10.1016/j.procs.2017.05.138.
- [20] R. D. Falgout et al. “Porting *hypre* to heterogeneous computer architectures: strategies and experiences”. In: *Parallel Computing* 108 (2021), p. 102840. DOI: 10.1016/j.parco.2021.102840.
- [21] N. Fehn, W. A. Wall, and M. Kronbichler. “Efficiency of high-performance discontinuous Galerkin spectral element methods for under-resolved turbulent incompressible flows”. In: *International Journal for Numerical Methods in Fluids* 88.1 (2018), pp. 32–54. DOI: 10.1002/flid.4511.
- [22] P. Fischer et al. “Scalability of high-performance PDE solvers”. In: *The International Journal of High Performance Computing Applications* 34.5 (2020), pp. 562–586. DOI: 10.1177/1094342020915762.
- [23] G. J. Gassner and A. D. Beck. “On the accuracy of high-order discretizations for underresolved turbulence simulations”. In: *Theoretical and Computational Fluid Dynamics* 27.3-4 (2012), pp. 221–237. DOI: 10.1007/s00162-011-0253-7.
- [24] M. González et al. “Multigroup radiation hydrodynamics with flux-limited diffusion and adaptive mesh refinement”. In: *Astronomy & Astrophysics* 578 (2015), A12. DOI: 10.1051/0004-6361/201525971.
- [25] F. Graziani and J. LeBlanc. *The crooked pipe test problem*. Tech. rep. UCRL-MI-143393. Lawrence Livermore National Laboratory, 2000.
- [26] V. E. Henson and U. M. Yang. “BoomerAMG: a parallel algebraic multigrid solver and preconditioner”. In: *Applied Numerical Mathematics* 41.1 (2002), pp. 155–177. DOI: 10.1016/s0168-9274(01)00115-5.
- [27] G. Kanschat, R. Lazarov, and Y. Mao. “Geometric multigrid for Darcy and Brinkman models of flows in highly heterogeneous porous media: a numerical study”. In: *Journal of Computational and Applied Mathematics* 310 (2017), pp. 174–185. DOI: 10.1016/j.cam.2016.05.016.
- [28] D. S. Kershaw. “Differencing of the diffusion equation in Lagrangian hydrodynamic codes”. In: *Journal of Computational Physics* 39.2 (1981), pp. 375–395. DOI: 10.1016/0021-9991(81)90158-3.
- [29] T. Kolev and W. Pazner. “Conservative and accurate solution transfer between high-order and low-order refined finite element spaces”. In: *SIAM Journal on Scientific Computing* 44.1 (2022), A1–A27. DOI: 10.1137/21M1403916.

- [30] T. Kolev et al. *CEED ECP milestone report: high-order algorithmic developments and optimizations for large-scale GPU-accelerated simulations*. Tech. rep. CEED-MS36. U.S. Department of Energy, 2021. DOI: 10.5281/zenodo.4672664.
- [31] T. Kolev et al. “Efficient exascale discretizations: high-order finite element methods”. In: *The International Journal of High Performance Computing Applications* (2021), pp. 527–552. DOI: 10.1177/109434202111020803.
- [32] T. V. Kolev and P. S. Vassilevski. “Parallel auxiliary space AMG solver for  $H(\text{div})$  problems”. In: *SIAM Journal on Scientific Computing* 34.6 (2012), A3079–A3098.
- [33] J. Kreeft, A. Palha, and M. Gerritsma. “Mimetic framework on curvilinear quadrilaterals of arbitrary order”. arXiv:1111.4304. 2011.
- [34] M. Kronbichler et al. “Comparison of implicit and explicit hybridizable discontinuous Galerkin methods for the acoustic wave equation”. In: *International Journal for Numerical Methods in Engineering* 106.9 (2015), pp. 712–739. DOI: 10.1002/nme.5137.
- [35] M. Kronbichler and K. Ljungkvist. “Multigrid for matrix-free high-order finite element computations on graphics processors”. In: *ACM Transactions on Parallel Computing* 6.1 (2019), pp. 1–32. DOI: 10.1145/3322813.
- [36] M. Kronbichler and W. A. Wall. “A performance comparison of continuous and discontinuous Galerkin methods with fast multigrid solvers”. In: *SIAM Journal on Scientific Computing* 40.5 (2018), A3423–A3448. DOI: 10.1137/16m110455x.
- [37] S. H. Langer, I. Karlin, and M. M. Marinak. “Performance Characteristics of HYDRA — A Multiphysics Simulation Code from LLNL”. In: *High Performance Computing for Computational Science — VECPAR 2014*. Lecture Notes in Computer Science. Springer International Publishing, 2015, pp. 173–181.
- [38] J. W. Lottes and P. F. Fischer. “Hybrid multigrid/Schwarz algorithms for the spectral element method”. In: *Journal of Scientific Computing* 24.1 (2005), pp. 45–78. DOI: 10.1007/s10915-004-4787-3.
- [39] K.-A. Mardal and R. Winther. “Preconditioning discretizations of systems of partial differential equations”. In: *Numerical Linear Algebra with Applications* 18.1 (2010), pp. 1–40. DOI: 10.1002/nla.716.
- [40] J. Melenk. “On condition numbers in  $hp$ -FEM with Gauss–Lobatto-based shape functions”. In: *Journal of Computational and Applied Mathematics* 139.1 (2002), pp. 21–48. DOI: 10.1016/s0377-0427(01)00391-0.
- [41] S. A. Orszag. “Spectral methods for problems in complex geometries”. In: *Journal of Computational Physics* 37.1 (1980), pp. 70–92.
- [42] W. Pazner. “Efficient low-order refined preconditioners for high-order matrix-free continuous and discontinuous Galerkin methods”. In: *SIAM Journal on Scientific Computing* 42.5 (2020), A3055–A3083. DOI: 10.1137/19m1282052.
- [43] W. Pazner, T. Kolev, and J.-S. Camier. “End-to-end GPU acceleration of low-order-refined preconditioning for high-order finite element discretizations”. arXiv eprint 2210.12253. (Submitted for publication). 2022.
- [44] W. Pazner, T. Kolev, and C. R. Dohrmann. “Low-order preconditioning for the high-order finite element de Rham complex”. In: *SIAM Journal on Scientific Computing* 45.2 (2023), A675–A702. DOI: 10.1137/22M1486534.
- [45] W. Pazner and P.-O. Persson. “Approximate tensor-product preconditioners for very high order discontinuous Galerkin methods”. In: *Journal of Computational Physics* 354 (2018), pp. 344–369.
- [46] J. Pestana and A. J. Wathen. “Natural preconditioning and iterative methods for saddle point systems”. In: *SIAM Review* 57.1 (2015), pp. 71–91.
- [47] E. M. Rønquist and A. T. Patera. “Spectral element multigrid. I. Formulation and numerical results”. In: *Journal of Scientific Computing* 2.4 (1987), pp. 389–406. DOI: 10.1007/bf01061297.
- [48] T. Rusten, P. S. Vassilevski, and R. Winther. “Interior penalty preconditioners for mixed finite element approximations of elliptic problems”. In: *Mathematics of Computation* 65.214 (1996), pp. 447–467. DOI: 10.1090/s0025-5718-96-00720-x.
- [49] S. A. Teukolsky. “Short note on the mass matrix for Gauss–Lobatto grid points”. In: *Journal of Computational Physics* 283 (2015), pp. 408–413.
- [50] O. Villa et al. “Power/Performance Trade-Offs of Small Batched LU Based Solvers on GPUs”. In: *Euro-Par 2013 Parallel Processing*. Springer Berlin Heidelberg, 2013, pp. 813–825.

- [51] W. Zulehner. “Nonstandard norms and robust estimates for saddle point problems”. In: *SIAM Journal on Matrix Analysis and Applications* 32.2 (2011), pp. 536–560.