A SCALABLE PRECONDITIONER FOR A PRIMAL DPG METHOD*

A. T. BARKER†, V. DOBREV†, J. GOPALAKRISHNAN‡, AND T. KOLEV†

Abstract. We show how a scalable preconditioner for the primal discontinuous Petrov-Galerkin (DPG) method can be developed using existing algebraic multigrid (AMG) preconditioning techniques. The stability of the DPG method gives a norm equivalence which allows us to exploit existing AMG algorithms and software. We show how these algebraic preconditioners can be applied directly to a Schur complement system arising from the DPG method. One of our intermediate results shows that a generic stable decomposition implies a stable decomposition for the Schur complement. This justifies the application of algebraic solvers directly to the interface degrees of freedom. Combining such results, we obtain the first massively scalable algebraic preconditioner for the DPG system.

Key words. algebraic multigrid, BoomerAMG, ADS, Schur complement, Discontinuous Petrov-Galerkin.

AMS subject classifications. 65F10, 65M55, 65N30.

1. Introduction.

Discontinuous Petrov-Galerkin (DPG) methods, introduced in [12, 13], constructed test spaces that guarantee stability. Today these methods are known to be simultaneously viewable as Galerkin mixed methods, as least-squares methods in nonstandard norms, or as Petrov-Galerkin methods using discontinuous functions [16]. DPG methods have a great deal of flexibility, allowing them to be applied to a wide variety of problems [10, 14, 17], and their convergence theory has now matured [9].

However, there is a lack of fast scalable solvers for the DPG method. In [3], an overlapping Schwarz preconditioner is analyzed: because it has no coarse level, the preconditioner expectedly deteriorates as overlap size become small. A coarse level was added for improved scalability in [25], where the authors analyzed a two-level additive Schwarz preconditioner for an ultraweak DPG method applied to the Poisson equation with Robin boundary condition. Going beyond the Poisson problem to the harmonic wave equation, there are numerical reports of good performance of certain preconditioning strategies [20].

It was clear from the inception of the DPG method that certain norm equivalences stemming from its stability give a natural avenue for the design of preconditioners, as was presented abstractly in [29]. We pursue the same avenue, but investigate all the details necessary to prove optimality. For example, the norm in which the DPG method is stable involves an infimum over an infinite dimensional space. Its equivalence with a finite dimensional infimum arising from a discrete Schur complement is an essential ingredient in our preconditioner analysis, an equivalence not proved in [29]. Another example of an observation immediate from our analysis (but not clear from [29]) is that the interface variable in the DPG method admits a preconditioner which uses only the interface degrees of freedom. In the recent work of [28] geometric multilevel strategies are investigated numerically but without theoretical analysis. In this paper we show how existing algebraic multilevel preconditioners can

*Performed under the auspices of the U.S. Department of Energy under Contract DE-AC52-07NA27344 (LLNL-JRNL-710378), supported in part by AFOSR grant FA9550-17-1-0090, and facilitated by equipment acquired using ARO grant W911NF-16-1-0307.

† Center for Applied Scientific Computing (CASC), Lawrence Livermore National Laboratory, Livermore, CA 94550, atb@llnl.gov, dobrev1@llnl.gov, tzanio@llnl.gov

‡ Portland State University, PO Box 751 (MTH), Portland, OR 97207-0751, gjay@pdx.edu

be effectively combined to precondition a DPG system, including at very large scale on parallel supercomputers.

The particular DPG method we consider is the so-called primal DPG method [15], reviewed in the next section. After describing a basic norm equivalence associated with the method, we proceed to analyze one of the component norms in Section 3. We show that this interface norm, obtained as an infimum over an infinite-dimensional space, is equivalent to an infimum over a finite-dimensional space. An auxiliary algebraic Schur complement result is presented in Section 4. Section 5 identifies the finite-dimensional infimum as a Schur complement norm and proceeds to analyze an auxiliary-space preconditioner for the Schur complement. The preconditioner and the main result are summarized in Section 6. Section 7 reports results from numerical studies of the proposed preconditioner. We conclude by summarizing the main results from the paper in Section 8.

2. The primal DPG method. For completeness and consistency of notation, we recall some definitions and results from the work introducing the primal DPG system [15]. The model problem we consider is the Poisson problem of finding $u \in H^1_0(\Omega)$ such that

$$\int_\Omega \kappa \nabla u \cdot \nabla v \, dx = \int_\Omega f v \, dx \tag{2.1}$$

for all $v \in H^1_0(\Omega)$, where Ω is a polygonal (in $\mathbb{R}^2$) or polyhedral domain (in $\mathbb{R}^3$) with Lipschitz boundary and $\kappa > 0$ is a piecewise constant coefficient. Zero Dirichlet boundary conditions on $\partial\Omega$ are essentially imposed in (2.1). In practice the method can also handle more general problems with varying coefficients and different boundary conditions, but we choose this setting for simplicity.

Even before discretization, the DPG formulation uses a mesh-dependent weak form. We assume that Ω is given together with a mesh Ω_h that partitions Ω into elements of varying shapes. For example, in the $\mathbb{R}^2$ case, the mesh elements $K \in \Omega_h$ may be triangles or quadrilaterals, while in $\mathbb{R}^3$ the elements may be tetrahedra, prisms, hexahedra, etc. Precise assumptions on the mesh and element shapes will be specified later, but for now we only require that the boundary of each element be Lipschitz, so that traces of Sobolev space functions on the element boundaries are well-defined. Specifically, we require a well-defined normal trace operator

$$\mathrm{trc}_n : H(\mathrm{div}, \Omega) \to \prod_{K \in \Omega_h} H^{-1/2}(\partial K),$$

which maps to element-wise traces $(\mathrm{trc}_n\, q)|_{\partial K} = q \cdot n|_{\partial K}$. One can also define $\mathrm{trc}_n\, q$ as a single valued function on the mesh facets: Indeed, if each interface facet $\gamma = \bar{K}_1 \cap \bar{K}_2$ between two elements $K_1, K_2 \in \Omega_h$, as well as boundary facets $\gamma = \bar{K}_1 \cap \partial\Omega$ are Lipschitz, then we may fix a continuous unit normal vector function n_γ on γ and define $(\mathrm{trc}_n\, q)|_\gamma = n_\gamma \cdot q|_\gamma$. This is a well-defined function in the dual space of $\mathring{H}^{1/2}(\gamma)$ (denoted also by $H^{1/2}_{00}(\gamma)$), whenever $q \in H(\mathrm{div}, \Omega)$.

The numerical fluxes of the DPG method lie in the range of trc_n, i.e., in the space $Q = \mathrm{ran}(\mathrm{trc}_n)$ with norm given by

$$\|q\|_Q = \inf_{\tau \in \mathrm{trc}_n^{-1}\{q\}} \|\tau\|_{H(\mathrm{div},\Omega)}. \tag{2.2}$$

Here, as usual, $\mathrm{trc}_n^{-1}\{q\}$ denotes the pre-image of the singleton $\{q\}$. It is standard to prove that the *minimal extension operator* $E : Q \to H(\mathrm{div}, \Omega)$, defined by $\mathrm{trc}_n(Eq) =$

q and $(Eq, v)_{H(\mathrm{div},K)} = 0$ for all $v \in \mathring{H}(\mathrm{div}, K)$ and $K \in \Omega_h$, attains the infimum at (2.2), i.e.,

$$\|q\|_Q = \|Eq\|_{H(\mathrm{div},\Omega)}. \tag{2.3}$$

Here and throughout, for any inner product space W, we use $\|\cdot\|_W$ and $(\cdot,\cdot)_W$ to denote its norm and inner product, respectively. When the space is uniquely understood from the argument or other context, we will drop the subscript. Let

$$H^1(\Omega_h) = \prod_{K\in\Omega_h} H^1(K).$$

This product space is endowed with the standard Cartesian product norm and inner product. For brevity, put $X = H_0^1(\Omega) \times Q$ and $Y = H^1(\Omega_h)$. Define the bilinear form $b : X \times Y \to \mathbb{R}$ by

$$b((w,r),v) = \sum_{K\in\Omega_h} \left(\int_K \kappa \nabla w \cdot \nabla v \, dx + \langle r, v\rangle_{H^{-1/2}(\partial K)} \right)$$

where $\langle r, v\rangle_{H^{-1/2}(\partial K)}$ denotes the duality pairing between $H^{-1/2}(\partial K)$ and $H^{1/2}(\partial K)$. The Dirichlet problem (2.1) can then be reformulated [15] as the problem of finding a pair $(u, q) \in X$ satisfying

$$b((u,q),v) = F(v), \qquad \forall v \in Y, \tag{2.4}$$

with $F(v) = (f, v)_{L^2(\Omega)}$. It is proved in [15, Lemma 3.4] that the problem (2.4) is uniquely solvable for $(u, q) \in X$ given any F in the dual space of Y (see also [9, Example 3.6] for a simplified analysis).

The primal DPG method uses the formulation (2.4) and finite element subspaces $U_h \subset H_0^1(\Omega)$, $Q_h \subset Q$, and $Y_h \subset H^1(\Omega_h)$. Here $h = \max_{K\in\Omega_h} \mathrm{diam}(K)$ denotes the mesh size of Ω_h. To describe a computational version of the method, let $\{u_i\}, \{q_j\}, \{v_k\}$ denote a finite element basis for U_h, Q_h and Y_h, respectively. Define the matrices

$$[\mathtt{B}_0]_{ki} = b((u_i, 0), v_k), \qquad [\mathtt{B}_1]_{kj} = b((0, q_j), v_k),$$
$$[\mathtt{M}]_{kl} = (v_l, v_k)_Y \equiv (v_l, v_k)_{L^2(\Omega)} + (\kappa\nabla v_l, \nabla v_k)_{[L^2(\Omega)]^d}$$

and set

$$\mathtt{B} = \begin{bmatrix} \mathtt{B}_0 & \mathtt{B}_1 \end{bmatrix}, \qquad \mathtt{A} = \mathtt{B}^T\mathtt{M}^{-1}\mathtt{B} = \begin{bmatrix} \mathtt{B}_0^T\mathtt{M}^{-1}\mathtt{B}_0 & \mathtt{B}_0^T\mathtt{M}^{-1}\mathtt{B}_1 \\ \mathtt{B}_1^T\mathtt{M}^{-1}\mathtt{B}_0 & \mathtt{B}_1^T\mathtt{M}^{-1}\mathtt{B}_1 \end{bmatrix}.$$

Let $\mathtt{u}$ denote the vector in $\mathbb{R}^{\dim(U_h)}$ representing a function $u \in U_h$ by the basis expansion formula $u = \sum_i [\mathtt{u}]_i u_i$. The vectors $\mathtt{q}$ in $\mathbb{R}^{\dim(Q_h)}$ and $\mathtt{v}$ in $\mathbb{R}^{\dim(Y_h)}$ are similarly defined. Restricting (2.4) to the finite dimensional spaces formally gives

$$\mathtt{B}\begin{bmatrix} \mathtt{u} \\ \mathtt{q} \end{bmatrix} = \mathtt{F}$$

where $[\mathtt{F}]_k = F(v_k)$. The DPG discretization of (2.4) solves instead the following symmetric and positive definite problem for $\mathtt{u}$ and $\mathtt{q}$:

$$\mathtt{A}\begin{bmatrix} \mathtt{u} \\ \mathtt{q} \end{bmatrix} = \mathtt{g} \tag{2.5}$$

where $\mathtt{g} = \mathtt{B}^T\mathtt{M}^{-1}\mathtt{F}$. Note that $\mathtt{M}$ is the Gram matrix in the (weighted) "broken" $H^1(\Omega_h)$-inner product and so is block diagonal (one block per element). Thus, $\mathtt{M}^{-1}$ can be evaluated fast locally.

The DPG method admits three well-known interpretations. The early papers on the DPG method used the concept of *optimal test functions* [12]. Its interpretation as a least-squares method in a nonstandard inner product was pointed out in [13, p. 6]. Its interpretation as a mixed method is now well known (see e.g., [6, Theorem 2.4]). It is easy to see that all these three interpretations, in practice, yield the same matrix system (2.5) when the same spaces and bases are used.

The starting point of our analysis is the stability of the DPG method (2.5). Let $X_h = U_h \times Q_h$ and let $\mathtt{x}$ and $\mathtt{z}$ be vectors representing two functions x and z in X_h, respectively. Per the above-mentioned notational conventions, $(\mathtt{Ax}, \mathtt{z})$ denotes the Euclidean inner product $\mathtt{z}^T\mathtt{Ax}$. It defines a bilinear form in the function space X_h, namely $a(x, z) = (\mathtt{Ax}, \mathtt{z})$. Note that both X_h and Y_h are used in the definition of $\mathtt{A}$. Throughout this paper we assume that the mesh Ω_h and the spaces X_h and Y_h are such that there exist mesh-independent constants c_1 and c_2 satisfying

$$c_1\|x\|_X^2 \le a(x, x) \le c_2\|x\|_X^2, \tag{2.6}$$

for all $x \in X_h$. The connection between (2.6) and the stability of the method is described next.

PROPOSITION 2.1. *Assumption* (2.6) *holds if and only if*

$$c_1\|x\|_X \le \sup_{0\neq v\in Y_h} \frac{|b(x, v)|}{\|v\|_Y} \le c_2\|x\|_X \tag{2.7}$$

for all x in X_h.

Proof. Define $T_h : X_h \to Y_h$ by $(T_h x, y)_Y = b(x, y)$, for all $x \in X_h$ and $y \in Y_h$. Then, for any $x \in X_h$,

$$\|T_h x\|_Y^2 = \sup_{0\neq v\in Y_h} \frac{(T_h x, v)_Y}{\|v\|_Y} = \sup_{0\neq v\in Y_h} \frac{b(x, v)}{\|v\|_Y}.$$

Letting $\mathtt{x}$ and $\mathtt{v}$ denote the vector representations of $x \in X_h$ and $v = T_h x \in Y_h$, respectively, it is easy to see that $\mathtt{v} = \mathtt{M}^{-1}\mathtt{Bx}$. Hence the result follows from $a(x, x) = (\mathtt{MM}^{-1}\mathtt{Bx}, \mathtt{M}^{-1}\mathtt{Bx}) = \|T_h x\|_Y^2$. ∎

Clearly, the upper inequality of (2.7) follows from the continuity of the bilinear form $b(\cdot, \cdot)$, and therefore holds independently of the choice of the discrete spaces. The lower inequality of (2.7) is an inf-sup condition. It follows from the fact that (2.4) is well-posed whenever the discrete spaces are chosen so that a Fortin operator [19] can be constructed. Here are a few known examples of cases where a c_1 independent of h can be obtained for the Dirichlet problem under consideration:

1. Suppose the mesh Ω_h is a quasiuniform tetrahedral geometrically conforming mesh, U_h is the Lagrange finite element space of degree p, $Q_h = \{q : q|_\gamma$ is a polynomial of degree at most $p-1$ on each mesh facet $\gamma\}$, and $Y_h = \{v : v|_K$ is a polynomial of degree at most $p+2\}$. Then a Fortin operator provided in [19] yields a c_1 independent of h, as proved in [15].
2. When Ω_h is a uniform mesh of rectangular elements, $U_h = \{w \in H_0^1(\Omega) : w|_K$ is in the tensor product space of polynomials of degree at most p in each coordinate direction, for all elements $K \in \Omega_h\}$, $Q_h = \{q : q|_\gamma \in P_p(\gamma)$ on

each mesh facet γ}, and $Y_h = \{v : v|_K$ is a polynomial of degree at most $p+3$ in each coordinate direction$\}$, a Fortin operator in [8] gives a mesh-independent c_1.

In the remainder of this paper, we examine an important implication of (2.6). Namely, in order to precondition the large Hermitian positive definite DPG system (2.5), it suffices to obtain a preconditioner for the $\|\cdot\|_X$ norm. In our model problem, this norm is

$$\|x\|_X^2 = \|u\|_{H^1(\Omega)}^2 + \|q\|_Q^2.$$

for any $x = (u, q) \in U_h \times Q_h$, so it suffices to combine preconditioners for the $H^1(\Omega)$ and Q norms. Since the former is standard, we focus on the latter in the next section.

3. Characterizing the Q-norm. The Q-norm (2.2) is defined through a minimization over an infinite dimensional space (the minimal extension E in (2.3) is not computable). In this section, we relate this norm to a minimum over a finite dimensional subspace.

3.1. Tetrahedral case. To present the idea transparently, we first detail the case when Ω_h is a geometrically conforming mesh of tetrahedral elements. For any tetrahedron K, let $R_p(K) = P_p(K)^3 + xP_p(K)$, where x is the coordinate vector and $P_p(K)$ denotes the set of all polynomials of total degree at most $p \geq 1$. The Raviart-Thomas finite element space is $R_h = \{r \in H(\mathrm{div}, \Omega) : r|_K \in R_p(K)$ for all $K \in \Omega_h\}$. Let $Q_h = \mathrm{trc}_n(R_h)$. Clearly Q_h is a finite dimensional subspace of Q. Define $E_h : Q_h \to R_h$ by $\mathrm{trc}_n(E_h q) = q$, and

$$(E_h q, v)_{H(\mathrm{div},K)} = 0, \qquad \forall v \in R_p(K) \cap \mathring{H}(\mathrm{div}, K),$$

for all $K \in \Omega_h$. This computable approximation of the minimal extension operator defines a new norm on Q_h,

$$\|q\|_{Q_h} = \|E_h q\|_{H(\mathrm{div},\Omega)}.$$

We now proceed to prove the equivalence of this norm with the Q-norm (in Theorem 3.3 below). Throughout this section, let c denote a generic positive constant whose value might change from one occurrence to another, but will remain *independent of h and p*. Let $\hat{K}$ denote the unit tetrahedron, $\hat{n}$ denote its outward unit normal on $\partial\hat{K}$, and $\hat{\sigma}_n = \hat{\sigma} \cdot \hat{n}|_{\partial\hat{K}}$. Let $Q_p(\partial\hat{K}) = \{\hat{\sigma}_n : \hat{\sigma} \in R_p(\hat{K})\}$. For any $\hat{\sigma} \in R_p(\hat{K})$ define the constant function

$$\bar{\sigma}_n = \frac{1}{|\partial\hat{K}|} \int_{\partial\hat{K}} \hat{\sigma}_n \, ds \,,$$

where $|\partial\hat{K}|$ denotes the surface area of $\partial\hat{K}$.

LEMMA 3.1. *There is a $c > 0$ and a $\hat{\mathcal{G}} : \mathbb{R} \to R_p(\hat{K})$ such that for any $\hat{\sigma}$ in $H(\mathrm{div}, \hat{K})$ with $\hat{\sigma}_n \in Q_p(\partial\hat{K})$, we have $\hat{n} \cdot (\hat{\mathcal{G}}\bar{\sigma}_n)|_{\partial\hat{K}} = \bar{\sigma}_n$,*

$$\|\hat{\mathcal{G}}\bar{\sigma}_n\|_{L^2(\hat{K})} \leq c\|\hat{\sigma}\|_{H(\mathrm{div},\hat{K})}, \qquad \textit{and} \qquad \|\mathrm{div}\,\hat{\mathcal{G}}\bar{\sigma}_n\|_{L^2(\hat{K})} \leq c\|\mathrm{div}\,\hat{\sigma}\|_{\hat{K}}.$$

Proof. We show that the following stronger inequality holds on the unit tetrahedron:

$$\|\hat{\mathcal{G}}\bar{\sigma}_n\|_{H(\mathrm{div},\hat{K})} \le c\|\operatorname{div}\hat{\sigma}\|_{\hat{K}}.$$

Let $\hat{x}_I$ denote the incenter of the unit tetrahedron $\hat{K}$. Then $(\hat{x}-\hat{x}_I)\cdot\hat{n} = d$ is constant for any $\hat{x}\in\partial\hat{K}$ ($d = 3|\hat{K}|/|\partial\hat{K}|$ is the radius of the insphere). Define

$$\hat{\mathcal{G}}\bar{\sigma}_n = \frac{\bar{\sigma}_n}{d}(\hat{x}-\hat{x}_I).$$

Then, setting $\hat{c} = \|\hat{\mathcal{G}}1\|_{H(\mathrm{div},\hat{K})}/|\partial\hat{K}|$ and $c = \hat{c}|\hat{K}|^{\frac{1}{2}}$, we have

$$\|\hat{\mathcal{G}}\bar{\sigma}_n\|^2_{H(\mathrm{div},\hat{K})} = \hat{c}^2\left|\int_{\partial\hat{K}}\hat{\sigma}_n\,ds\right|^2 = \hat{c}^2\left|\int_{\hat{K}}\operatorname{div}\hat{\sigma}\,dx\right|^2 \le c^2\|\operatorname{div}\hat{\sigma}\|^2_{\hat{K}} \tag{3.1}$$

and $\hat{\mathcal{G}}\bar{\sigma}_n = \bar{\sigma}_n(\hat{x}-\hat{x}_I)\cdot\hat{n}/d = \bar{\sigma}_n$, for all $\hat{x}\in\partial\hat{K}$. ∎

LEMMA 3.2. *There is a $c>0$ and a $\hat{\mathcal{E}}: Q_p(\partial\hat{K})\to R_p(\hat{K})$ such that for any $\hat{\sigma}$ in $H(\mathrm{div},\hat{K})$ with $\hat{\sigma}_n\in Q_p(\partial\hat{K})$, we have $\hat{n}\cdot(\hat{\mathcal{E}}\hat{\sigma}_n)|_{\partial\hat{K}} = \hat{\sigma}_n$,*

$$\|\hat{\mathcal{E}}(\hat{\sigma}_n-\bar{\sigma}_n)\|_{L^2(\hat{K})} \le c\|\hat{\sigma}\|_{H(\mathrm{div},\hat{K})}, \qquad \textit{and} \qquad \operatorname{div}(\hat{\mathcal{E}}(\hat{\sigma}_n-\bar{\sigma}_n)) = 0.$$

Proof. We use the polynomial extension operator $\mathcal{E}^{\mathrm{div}}$ from [18, Theorem 7.1]: Accordingly *(a)* if $\hat{q}\in Q_p(\partial\hat{K})$, then $\mathcal{E}^{\mathrm{div}}\hat{q}$ is in $R_p(\hat{K})$, *(b)* if $\hat{q}$ has zero mean, then $\operatorname{div}(\mathcal{E}^{\mathrm{div}}\hat{q}) = 0$, and *(c)* if $\hat{\tau}$ is any extension of $\hat{q}$ (i.e., $\hat{\tau}$ is a function in $H(\mathrm{div},\hat{K})$ satisfying $\hat{n}\cdot\hat{\tau}|_{\partial\hat{K}} = \hat{q}$), then

$$\|\mathcal{E}^{\mathrm{div}}\hat{q}\|_{H(\mathrm{div},\hat{K})} \le c\|\hat{\tau}\|_{H(\mathrm{div},\hat{K})}.$$

Since $\hat{\sigma}$ is an extension of $\hat{\sigma}_n$ and $\hat{\mathcal{G}}\bar{\sigma}_n$ is an extension of $\bar{\sigma}_n$,

$$\begin{aligned}\|\mathcal{E}^{\mathrm{div}}(\hat{\sigma}_n-\bar{\sigma}_n)\|_{H(\mathrm{div},\hat{K})} &\le \|\mathcal{E}^{\mathrm{div}}\hat{\sigma}_n\|_{H(\mathrm{div},\hat{K})} + \|\mathcal{E}^{\mathrm{div}}\bar{\sigma}_n\|_{H(\mathrm{div},\hat{K})}\\ &\le c\left(\|\hat{\sigma}\|_{H(\mathrm{div},\hat{K})} + \|\hat{\mathcal{G}}\bar{\sigma}_n\|_{H(\mathrm{div},\hat{K})}\right) \le c\|\hat{\sigma}\|_{H(\mathrm{div},\hat{K})}.\end{aligned}$$

Finally, since $\hat{q} = \hat{\sigma}_n - \bar{\sigma}_n$ has zero mean, we have $\operatorname{div}(\mathcal{E}^{\mathrm{div}}(\hat{\sigma}_n-\bar{\sigma}_n)) = 0$. ∎

Before the next result, recall that for any tetrahedral element K, there is an affine homeomorphism $\Phi:\hat{K}\to K$. Let $[D\Phi]$ denote the Jacobian matrix of its derivatives and let $J = \det[D\Phi]$. Define the Piola maps

$$\Phi_*\hat{\sigma} = J^{-1}[D\Phi_K]\hat{\sigma}\circ\Phi^{-1}, \qquad \Phi^*\sigma = J[D\Phi_K]^{-1}\sigma\circ\Phi.$$

Clearly Φ_* maps functions on $\hat{K}$ to K, while Φ^* maps in the opposite direction, from K to $\hat{K}$. Letting $|K|$ denote the volume of the K and $h_K = \operatorname{diam}(K)$, we recall the following standard estimates [5] for affine Φ: There is a $c>0$, depending only on shape regularity of K, such that

$$\|\Phi_*\hat{\sigma}\|^2_{L^2(K)} \le c\,\frac{h_K^2}{|K|}\|\hat{\sigma}\|^2_{L^2(\hat{K})}, \qquad \|\operatorname{div}(\Phi_*\hat{\sigma})\|^2_{L^2(K)} \le \frac{c}{|K|}\|\operatorname{div}\hat{\sigma}\|^2_{L^2(\hat{K})}, \tag{3.2a}$$

$$\|\Phi^*\sigma\|^2_{L^2(\hat{K})} \le c\,\frac{|K|}{h_K^2}\|\sigma\|^2_{L^2(K)}, \qquad \|\operatorname{div}(\Phi^*\sigma)\|^2_{L^2(\hat{K})} \le c\,|K|\,\|\operatorname{div}\sigma\|^2_{L^2(K)}, \tag{3.2b}$$

for all $\hat{\sigma} \in H(\mathrm{div}, \hat{K})$ and $\sigma \in H(\mathrm{div}, K)$.

We are now in a position to put everything together and prove the main norm equivalence result of this section.

THEOREM 3.3. *If Ω_h is shape regular, then there is a $c_3 > 0$ independent of h and p (and depending only on the shape regularity) such that*

$$\|q\|_Q \le \|q\|_{Q_h} \le c_3 \|q\|_Q \tag{3.3}$$

for all $q \in Q_h$.

Proof. The lower inequality follows from

$$\|q\|_Q = \inf_{\tau \in \mathrm{trc}_n^{-1}\{q\}} \|\tau\|_{H(\mathrm{div},\Omega)} \le \inf_{\tau_h \in R_h \cap \mathrm{trc}_n^{-1}\{q\}} \|\tau_h\|_{H(\mathrm{div},\Omega)} = \|q\|_{Q_h}.$$

To prove the upper inequality, pick any $K \in \Omega_h$, set

$$\sigma = (Eq)|_K, \quad \hat{\sigma} = \Phi^* \sigma, \quad \mathcal{F}_K q = \Phi_* \hat{\mathcal{F}} \hat{\sigma}_n,$$

where E is the minimal extension in (2.3),

$$\hat{\mathcal{F}} \hat{\sigma}_n = \hat{\mathcal{G}} \bar{\sigma}_n + \hat{\mathcal{E}}(\hat{\sigma}_n - \bar{\sigma}_n),$$

and $\hat{\mathcal{G}}$ and $\hat{\mathcal{E}}$ are as given by Lemmas 3.1 and 3.2. Clearly, $n \cdot (\mathcal{F}_K q)|_{\partial K} = q$ and the function $\mathcal{F}q$, defined by $(\mathcal{F}q)|_K = \mathcal{F}_K q$ for all $K \in \Omega_h$, is in $H(\mathrm{div}, \Omega)$. Moreover the estimates of Lemmas 3.1 and 3.2, together with (3.2), imply

$$\begin{aligned}
\|\mathcal{F}q\|^2_{L^2(K)} &\le c \frac{h_K^2}{|K|} \left(\|\hat{\mathcal{G}} \bar{\sigma}_n\|^2_{L^2(\hat{K})} + \|\hat{\mathcal{E}}(\hat{\sigma}_n - \bar{\sigma}_n)\|^2_{L^2(\hat{K})} \right) \le c \frac{h_K^2}{|K|} \|\hat{\sigma}\|^2_{H(\mathrm{div},\hat{K})} \\
&= c \frac{h_K^2}{|K|} \|\Phi^* \sigma\|^2_{H(\mathrm{div},\hat{K})} \le c \left(\|\sigma\|^2_{L^2(K)} + h_K^2 \|\mathrm{div}\sigma\|^2_{L^2(K)} \right), \\
\|\mathrm{div}\mathcal{F}q\|^2_{L^2(K)} &\le \frac{c}{|K|} \|\mathrm{div}(\hat{\mathcal{G}} \bar{\sigma}_n)\|^2_{L^2(\hat{K})} \le \frac{c}{|K|} \|\mathrm{div}\hat{\sigma}\|^2_{L^2(\hat{K})} = \frac{c}{|K|} \|\mathrm{div}(\Phi^* \sigma)\|^2_{L^2(\hat{K})} \\
&\le c \|\mathrm{div}\sigma\|^2_{L^2(K)}.
\end{aligned}$$

We have thus obtained, for any $q \in Q_h$, an extension $\mathcal{F}q \in R_h$ satisfying

$$\|\mathcal{F}q\|_{H(\mathrm{div},\Omega)} \le c \|\sigma\|_{H(\mathrm{div},\Omega)} = c \|Eq\|_{H(\mathrm{div},\Omega)} = c \|q\|_Q.$$

Since $\|q\|_{Q_h}$ is the infimum of $\|\tau_h\|_{H(\mathrm{div},\Omega)}$ over all $\tau_h \in R_h$ satisfying $\mathrm{trc}_n \tau_h = q$, the inequality $\|q\|_{Q_h} \le \|\mathcal{F}q\|_{H(\mathrm{div},\Omega)}$ holds and completes the proof. ∎

3.2. General meshes. We now briefly remark on how the norm equivalence of Theorem 3.3 may be extended to more general elements and meshes. While a general theorem for all element shapes is beyond the scope of this paper, we wish to provide pointers on what arguments need extension. The proof of Theorem 3.3 depends on three ingredients: *(a)* Lemma 3.1, *(b)* Lemma 3.2, and *(c)* the scaling estimates (3.2). Moving from tetrahedral to other element shapes, we must first obtain generalizations of the extension operators of Lemmas 3.1 and 3.2 on the reference element $\hat{K}$ for the new shapes. We show how this can be done for two other element shapes, one in two dimensions and another in three dimensions.

Triangles: The extension $\hat{\mathcal{G}}$ constructed in the proof of Lemma 3.1 continues to work for the unit triangle if we set $\hat{x}_I$ to be the center of the inscribed circle of the

triangle. As for Lemma 3.2, if $\hat{K}$ is a triangle, then the extension of [1, Corollary 2.2] has all the properties stated in the lemma.

Cubes: To obtain the result of Lemma 3.1 when $\hat{K}$ is the unit cube, we set $\hat{x}_I = (1/2, 1/2, 1/2)$ and $\hat{\mathcal{G}}\bar{\sigma}_n = 2(\hat{x} - \hat{x}_I)\bar{\sigma}_n$. Then proceeding as in (3.1), we obtain the result. The extension operators constructed in [11] for each p provide the required $\hat{\mathcal{E}}$ in Lemma 3.2 when $\hat{K}$ is a cube.

The scaling estimates (3.2) are valid for affine mappings Φ. We next comment on meshes with curved elements, which are images of reference elements under a possibly nonlinear Φ. If Φ is such that the estimates of (3.2) with a properly (re)defined h_K and $|K|$ for curvilinear elements K hold, then the proof of Theorem 3.3 can be generalized. Examples of nonlinear Φ where such geometrical quantities can be identified can be found in [4, 5].

4. An algebraic Schur complement result. The purpose of this section, which can be read independently of the rest of the paper, is to present a simple matrix result, whose relevance to our problem will be clear in the next section. The result is a generalization of [7, Lemma 4.2]. Suppose $i \cup f = \{1, 2, \ldots, m\}$ and $j \cup e = \{1, 2, \ldots, l\}$ are disjoint partitions of two index sets. Let $\mathtt{D}$ be an $m \times m$ symmetric positive definite matrix and $\mathtt{H}$ be an $m \times l$ matrix (both with real entries). We use standard block notations, e.g., $\mathtt{x}_f$ denotes the restriction of a vector $\mathtt{x}$ to f-indices, and the matrices have block forms

$$\mathtt{D} = \begin{bmatrix} \mathtt{D}_{ii} & \mathtt{D}_{if} \\ \mathtt{D}_{fi} & \mathtt{D}_{ff} \end{bmatrix}, \quad \mathtt{H} = \begin{bmatrix} \mathtt{H}_{ij} & \mathtt{H}_{ie} \\ \mathtt{H}_{fj} & \mathtt{H}_{fe} \end{bmatrix}. \tag{4.1}$$

Define $\mathtt{S}$ to be the Schur complement $\mathtt{S} = \mathtt{D}_{ff} - \mathtt{D}_{fi}\mathtt{D}_{ii}^{-1}\mathtt{D}_{if}$. Let $\mathtt{diag(D)}$ denote the diagonal matrix formed from the diagonal of $\mathtt{D}$.

LEMMA 4.1. *Suppose there is a $c_4 > 0$ such that every $\mathtt{u} \in \mathbb{R}^m$ can be decomposed as $\mathtt{u} = \mathtt{v} + \mathtt{Hr}$, for some $\mathtt{v} \in \mathbb{R}^m$ and $\mathtt{r} \in \mathbb{R}^l$, such that*

$$(\mathtt{diag(D)v}, \mathtt{v}) + (\mathtt{DHr}, \mathtt{Hr}) \le c_4(\mathtt{Du}, \mathtt{u}). \tag{4.2}$$

Then for any $\mathtt{u} \in \mathbb{R}^m$ there exist $\mathtt{v} \in \mathbb{R}^m$ and $\mathtt{r} \in \mathbb{R}^l$ (not necessarily the same as in the assumption), depending only on $\mathtt{u}_f$, such that the decomposition $\mathtt{u}_f = \mathtt{v}_f + [\mathtt{Hr}]_f$ holds and satisfies

$$(\mathtt{diag(S)v}_f, \mathtt{v}_f) + (\mathtt{S}[\mathtt{Hr}]_f, [\mathtt{Hr}]_f) \le c_4(\mathtt{Su}_f, \mathtt{u}_f).$$

Proof. Let $\mathtt{E}$ be the matrix representation of the extension operator E_h

$$\mathtt{E} = \begin{bmatrix} -\mathtt{D}_{ii}^{-1}\mathtt{D}_{if} \\ \mathtt{I}_{ff} \end{bmatrix}.$$

Since $\mathtt{S} = \mathtt{E}^T\mathtt{DE}$, from the well-known properties of Schur complements

$$(\mathtt{Sx}_f, \mathtt{x}_f) = (\mathtt{DEx}_f, \mathtt{Ex}_f) = \inf_{\{\mathtt{y} \in \mathbb{R}^m : \mathtt{y}_f = \mathtt{x}_f\}} (\mathtt{Dy}, \mathtt{y}) \le (\mathtt{Dx}, \mathtt{x}), \qquad \forall \mathtt{x} \in \mathbb{R}^m. \tag{4.3}$$

Now, given any $\mathtt{u} \in \mathbb{R}^m$, let us set $\mathtt{w} = \mathtt{Eu}_f$ and let $\mathtt{v} \in \mathbb{R}^m$, $\mathtt{r} \in \mathbb{R}^l$ be such that $\mathtt{w} = \mathtt{v} + \mathtt{Hr}$ (and in particular $\mathtt{u}_f = \mathtt{w}_f = \mathtt{v}_f + [\mathtt{Hr}]_f$) and

$$(\mathtt{diag(D)v}, \mathtt{v}) + (\mathtt{DHr}, \mathtt{Hr}) \le c_4(\mathtt{Dw}, \mathtt{w}). \tag{4.4}$$

By (4.3), with $\mathtt{x} = \mathtt{Hr}$

$$(\mathtt{S}[\mathtt{Hr}]_f, [\mathtt{Hr}]_f) \leq (\mathtt{DHr}, \mathtt{Hr}). \tag{4.5}$$

Next, consider the k-th diagonal entry of $\mathtt{S}$ which can be expressed as $\mathtt{S}_{kk} = (\mathtt{Se}_k, \mathtt{e}_k)$ where $\mathtt{e}_k$ is the vector with entries $[\mathtt{e}_k]_s = \delta_{ks}$. Setting $\mathtt{x}^T = [\mathtt{0}^T \ \mathtt{e}_k^T]$ in (4.3), we get

$$\mathtt{S}_{kk} = (\mathtt{Se}_k, \mathtt{e}_k) \leq (\mathtt{D}_{ff}\mathtt{e}_k, \mathtt{e}_k) = [\mathtt{D}_{ff}]_{kk}\,.$$

Since all diagonal entries of $\mathtt{D}$ are positive, we conclude that

$$(\mathtt{diag}(\mathtt{S})\mathtt{v}_f, \mathtt{v}_f) \leq (\mathtt{diag}(\mathtt{D}_{ff})\mathtt{v}_f, \mathtt{v}_f) \leq (\mathtt{diag}(\mathtt{D})\mathtt{v}, \mathtt{v}). \tag{4.6}$$

Adding the estimates (4.5) and (4.6) and then using (4.4) we arrive at

$$(\mathtt{diag}(\mathtt{S})\mathtt{v}_f, \mathtt{v}_f) + (\mathtt{S}[\mathtt{Hr}]_f, [\mathtt{Hr}]_f) \leq (\mathtt{diag}(\mathtt{D})\mathtt{v}, \mathtt{v}) + (\mathtt{DHr}, \mathtt{Hr}) \leq c_4(\mathtt{Dw}, \mathtt{w}).$$

Noting that $(\mathtt{Dw}, \mathtt{w}) = (\mathtt{DEu}_f, \mathtt{Eu}_f) = (\mathtt{Su}_f, \mathtt{u}_f)$ completes the proof. □

The statement of Lemma 4.1 can be easily extended to the case of more than one matrix $\mathtt{H}$: assume that we have a sequence of real matrices $\mathtt{H}_k$ with dimensions $m \times l_k$, $k = 1, \ldots, n$.

COROLLARY 4.2. *Suppose there is $c_4 > 0$ such that for all $\mathtt{u} \in \mathbb{R}^m$ there exist $\mathtt{v} \in \mathbb{R}^m$ and $\mathtt{r}_k \in \mathbb{R}^{l_k}$, $k = 1, \ldots, n$, such that*

$$\mathtt{u} = \mathtt{v} + \sum_{k=1}^{n} \mathtt{H}_k\mathtt{r}_k\,, \qquad \text{and} \qquad (\mathtt{diag}(\mathtt{D})\mathtt{v}, \mathtt{v}) + \sum_{k=1}^{n} (\mathtt{DH}_k\mathtt{r}_k, \mathtt{H}_k\mathtt{r}_k) \leq c_4(\mathtt{Du}, \mathtt{u}).$$

Then for any $\mathtt{u} \in \mathbb{R}^m$ there exist $\mathtt{v} \in \mathbb{R}^m$ and $\mathtt{r}_k \in \mathbb{R}^{l_k}$, $k = 1, \ldots, n$ (not necessarily the same as in the assumption), depending only on $\mathtt{u}_f$, such that

$$\mathtt{u}_f = \mathtt{v}_f + \sum_{k=1}^{n} [\mathtt{H}_k\mathtt{r}_k]_f\,, \quad \text{and} \quad (\mathtt{diag}(\mathtt{S})\mathtt{v}_f, \mathtt{v}_f) + \sum_{k=1}^{n} (\mathtt{S}[\mathtt{H}_k\mathtt{r}_k]_f, [\mathtt{H}_k\mathtt{r}_k]_f) \leq c_4(\mathtt{Su}_f, \mathtt{u}_f).$$

5. Preconditioning the Q_h-norm using an interface decomposition. In Section 3, we reduced the problem of preconditioning $\|\cdot\|_Q^2$ to that of preconditioning $\|\cdot\|_{Q_h}^2$. In this section we propose a scalable method for preconditioning the Q_h-norm, by further reducing the problem to that of preconditioning the Gram matrix of the $H(\mathrm{div}, \Omega)$ inner product. Such matrices can be efficiently handled by recent algebraic multigrid techniques [24], resulting ultimately in a good preconditioner for the DPG matrix $\mathtt{A}$, as shown in the next section. We will also show that only computations involving the interface degrees of freedom are needed to precondition the Q_h-norm.

Let $\{r_m\}$ denote a finite element basis of R_h. Define $\mathtt{D}$ to be the Gram matrix of the $H(\mathrm{div}, \Omega)$ inner product in the $\{r_m\}$ basis. We partition the degrees of freedom of $\{r_m\}$ into those associated with the interior of elements – denoted by i – and those on the element interfaces – denoted by f – and block partition $\mathtt{D}$ as in (4.1). Recall the notational conventions from Section 2 that allow us to move from functions q to their vector representations $\mathtt{q}$ using appropriate basis expansions. As already noted in (4.3), the Schur complement $\mathtt{S} = \mathtt{D}_{ff} - \mathtt{D}_{fi}\mathtt{D}_{ii}^{-1}\mathtt{D}_{if}$ satisfies

$$(\mathtt{Sq}, \mathtt{q}) = \inf_{\{r \in R_h\colon \mathtt{r}_f = \mathtt{q}\}} (\mathtt{Dr}, \mathtt{r}) = \|E_h q\|_{H(\mathrm{div},\Omega)}^2 = \|q\|_{Q_h}^2, \tag{5.1}$$

i.e., to precondition the Q_h-norm we need to construct a good preconditioner for $\mathtt{S}$.

The characterization of the Q_h-norm in terms of an $H(\mathrm{div})$-norm suggests the use of an $H(\mathrm{div})$ preconditioner. Indeed, if $\mathtt{T} = \begin{bmatrix}\mathtt{0}_{fi} & \mathtt{I}_{ff}\end{bmatrix}$ denotes the restriction operator such that $\mathtt{Tr} = \mathtt{r}_f$ for all $r \in R_h$, then it follows from

$$\mathtt{D}^{-1} = \begin{bmatrix}\mathtt{I} & -\mathtt{D}_{ii}^{-1}\mathtt{D}_{if}\\ 0 & \mathtt{I}\end{bmatrix}\begin{bmatrix}\mathtt{D}_{ii}^{-1} & 0\\ 0 & \mathtt{S}^{-1}\end{bmatrix}\begin{bmatrix}\mathtt{I} & 0\\ -\mathtt{D}_{fi}\mathtt{D}_{ii}^{-1} & \mathtt{I}\end{bmatrix}$$

that $\mathtt{S}^{-1} = [\mathtt{D}^{-1}]_{ff} = \mathtt{T}\mathtt{D}^{-1}\mathtt{T}^T$. Thus, replacing $\mathtt{D}^{-1}$ by any spectrally equivalent $H(\mathrm{div},\Omega)$-preconditioner will give us a spectrally equivalent preconditioner for $\mathtt{S}$. In particular, we may use the Auxiliary-space Divergence Solver (ADS) of [24].

It is well known that ADS is a good preconditioner for many problems set in the $H(\mathrm{div},\Omega)$-conforming space R_h. However, we want to precondition the interface operator $\mathtt{S}$ using only the interface degrees of freedom. The ADS preconditioner when applied to R_h uses all degrees of freedom of R_h, and not merely the interface degrees of freedom in Q_h. This can become a significant addition to the cost as the order p increases.

What can we expect when the algebraic ADS is directly applied to the interface space Q_h? To answer this, we examine below the stable decomposition underpinning the theory of ADS and employ Corollary 4.2 to get an analogous stable decomposition restricted to the interface. For simplicity, we now focus on the three-dimensional case. (The two-dimensional case is similar once curl is properly defined.) Let N_h denote the $H(\mathrm{curl})$-conforming Nedelec space of the first kind on the same mesh, which is in correspondence with R_h in the standard finite element exact sequence.

ADS is based on a decomposition of an arbitrary $H(\mathrm{div},\Omega)$ function using an H^1 component, the curl of an $H(\mathrm{curl})$ component, and a component that is "small" in the sense that it can be handled by simple smoothing. This decomposition leads to an additive preconditioner for $\mathtt{D}$ in the form

$$\mathtt{B}^{\mathtt{ADS}} = \mathtt{R} + \mathtt{\Pi}\,\mathtt{B}^{\mathtt{\Pi}}\,\mathtt{\Pi}^T + \mathtt{C}\,\mathtt{B}^{\mathtt{C}}\,\mathtt{C}^T \tag{5.2}$$

where the ingredients are as follows:

1. $\mathtt{R}$ is a simple smoother for the global matrix $\mathtt{D}$, for example, one symmetrized Gauss-Seidel iteration.
2. $\mathtt{\Pi}$ is the matrix representation of the Raviart-Thomas interpolation operator from $U_h \times U_h \times U_h$ (or simply U_h^3) to R_h obtained using a standard basis $\{u_l\}$ of U_h and the basis $\{r_m\}$ of R_h.
3. $\mathtt{C}$ is the matrix representation of curl : $N_h \to R_h$ using a standard basis $\{n_k\}$ of N_h and the basis $\{r_m\}$ of R_h.
4. $\mathtt{B}^{\mathtt{\Pi}}$ is a standard algebraic H^1 solver, for example BoomerAMG from [21, 2], applied to the matrix $\mathtt{\Pi}^T\mathtt{D}\mathtt{\Pi}$.
5. $\mathtt{B}^{\mathtt{C}}$ is an algebraic Maxwell solver, such as the Auxiliary-space Maxwell Solver (AMS) of [23] applied to $\mathtt{C}^T\mathtt{D}\mathtt{C}$.

Just as we partitioned the degrees of freedom of R_h into interior (i) and interface (f) ones, we can partition the degrees of freedom of U_h^3 into its interior $\tilde{i}$ and its interface ($\tilde{f}$) degrees of freedom. Similarly the degrees of freedom of N_h are partitioned into sets $\breve{i}$ (interior) and $\breve{f}$ (interface). An important property of the matrices $\mathtt{\Pi}$ and $\mathtt{C}$ is that when we decompose them into the interior and interface degrees of freedom, their block form is

$$\mathtt{\Pi} = \begin{bmatrix}\mathtt{\Pi}_{i\tilde{i}} & \mathtt{\Pi}_{i\tilde{f}}\\ \mathtt{0} & \mathtt{\Pi}_{f\tilde{f}}\end{bmatrix}, \quad \mathtt{C} = \begin{bmatrix}\mathtt{C}_{i\breve{i}} & \mathtt{C}_{i\breve{f}}\\ \mathtt{0} & \mathtt{C}_{f\breve{f}}\end{bmatrix}. \tag{5.3}$$

The fact that $\Pi_{f\tilde{i}}$ and $C_{f\check{i}}$ are zero blocks follows from the definition of the finite element spaces U_h, R_h, N_h and their degrees of freedom, e.g., the R_h degrees of freedom on a face for the curl of a function in N_h depend only on the N_h degrees of freedom associated with that face.

The rationale behind the preconditioner construction in (5.2) comes from the theory of auxiliary space preconditioners [22]. For example, it is possible to prove [24, Section 5.2] under further simplifying assumptions that any $u \in R_h$ can be decomposed into

$$\mathtt{u} = \mathtt{v} + \Pi\mathtt{z} + \mathtt{C}\mathtt{y} \tag{5.4a}$$

with $z \in U_h \times U_h \times U_h$, $y \in N_h$, and $v \in R_h$ such that

$$(\mathtt{diag}(\mathtt{D})\mathtt{v}, \mathtt{v}) + (\mathtt{D}\Pi\mathtt{z}, \Pi\mathtt{z}) + (\mathtt{D}\mathtt{C}\mathtt{y}, \mathtt{C}\mathtt{y}) \le c_5(\mathtt{D}\mathtt{u}, \mathtt{u}) \tag{5.4b}$$

where $c_5 > 0$ is a constant independent of the size of the problem. This is enough to conclude [27] that $\mathtt{B}^{\mathtt{ADS}}$ is a good preconditioner for $\mathtt{D}^{-1}$ (and the "goodness" is measured by c_5 and the condition numbers of $\mathtt{B}^{\Pi}\Pi^T\mathtt{D}\Pi$ and $\mathtt{B}^{\mathtt{C}}\mathtt{C}^T\mathtt{D}\mathtt{C}$). In practice, $\mathtt{B}^{\mathtt{ADS}}$ often serves as a good preconditioner for $\mathtt{D}^{-1}$ even when a rigorous proof of (5.4) is difficult (such as for non-conforming irregular meshes and discontinuous material coefficients). Loosely speaking, (5.4) means that $\mathtt{u}$ can be decomposed into well-behaved components in the ranges of Π and $\mathtt{C}$ with a small remainder $\mathtt{v}$.

When a purely algebraic implementation of ADS is applied to $\mathtt{S}$, it results in the preconditioner

$$\mathtt{B}^f = \mathtt{R}^f + \Pi_{f\tilde{f}}\tilde{\mathtt{B}}^{\Pi}\Pi^T_{f\tilde{f}} + \mathtt{C}_{f\check{f}}\check{\mathtt{B}}^{\mathtt{C}}\mathtt{C}^T_{f\check{f}} \tag{5.5}$$

which *uses only the interface degrees of freedom* of all the spaces involved. Here $\mathtt{R}^f$ is a simple point smoother, like the symmetrized Gauss-Seidel iteration, applied to $\mathtt{S}$, $\tilde{\mathtt{B}}^{\Pi}$ is BoomerAMG solver applied to $\Pi^T_{f\tilde{f}}\mathtt{S}\Pi_{f\tilde{f}}$, and $\check{\mathtt{B}}^{\mathtt{C}}$ is AMS applied to $\mathtt{C}^T_{f\check{f}}\mathtt{S}\mathtt{C}_{f\check{f}}$. Arguments for the effectiveness of the interface preconditioners $\check{\mathtt{B}}^{\mathtt{C}}$ and $\tilde{\mathtt{B}}^{\Pi}$ are presented in [7]. Our goal is to discuss the effectiveness of $\mathtt{B}^f$. Just as (5.4) implies that $\mathtt{B}$ is a good preconditioner for $\mathtt{D}$, a stable interface decomposition is required for $\mathtt{B}^f$ to be a good preconditioner for $\mathtt{S}$. We will now show that the decomposition (5.4) implies a stable interface decomposition.

LEMMA 5.1. *If (5.4) holds, then any $q \in Q_h$ can be decomposed as*

$$\mathtt{q} = \mathtt{v}_f + \Pi_{f\tilde{f}}\mathtt{z}_f + \mathtt{C}_{f\check{f}}\mathtt{y}_f$$

where $v \in R_h$, $z \in U_h^3$ and $y \in N_h$ and their interface degrees of freedom satisfy

$$(\mathtt{diag}(\mathtt{S})\mathtt{v}_f, \mathtt{v}_f) + (\mathtt{S}\Pi_{f\tilde{f}}\mathtt{z}_{\tilde{f}}, \Pi_{f\tilde{f}}\mathtt{z}_{\tilde{f}}) + (\mathtt{S}\mathtt{C}_{f\check{f}}\mathtt{y}_{\check{f}}, \mathtt{C}_{f\check{f}}\mathtt{y}_{\check{f}}) \le c_5(\mathtt{S}\mathtt{q}, \mathtt{q}).$$

In addition, if

$$\kappa\left(\tilde{\mathtt{B}}^{\Pi}\,\Pi^T_{f\tilde{f}}\mathtt{S}\Pi_{f\tilde{f}}\right) \le c_\kappa, \qquad \kappa\left(\check{\mathtt{B}}^{\mathtt{C}}\,\mathtt{C}^T_{f\check{f}}\mathtt{S}\mathtt{C}_{f\check{f}}\right) \le c_\kappa\,,$$

then the condition number $\kappa(\mathtt{B}^f\mathtt{S})$ depends only on c_5 and c_κ.

Proof. Apply Corollary 4.2 with $\mathtt{H}_1 = \Pi$, and $\mathtt{H}_2 = \mathtt{C}$, and observe that $[\Pi\mathtt{z}]_f = \Pi_{f\tilde{f}}\mathtt{z}_{\tilde{f}}$ and $[\mathtt{C}\mathtt{y}]_f = \mathtt{C}_{f\check{f}}\mathtt{y}_{\check{f}}$ due to (5.3). ∎

Informally, the result of the lemma can be stated as follows: if ADS works for the matrix $\mathtt{D}$ (a volumetric discretization of $\|\cdot\|_{H(\mathrm{div},\Omega)}$), it will also work for its Schur complement $\mathtt{S}$ (an interfacial discretization of $\|\cdot\|_{Q_h}$). Since we assume the former, we can conclude that ADS will be an effective preconditioner for $\mathtt{S}$.

6. Scalable preconditioner. We are now ready to put all the pieces together and define a scalable preconditioner for the original DPG matrix $\mathtt{A}$. Our basic premise is that *(i)* the algebraic ADS is a good solver for the Gram matrix of the $H(\mathrm{div},\Omega)$-inner product in R_h, in the sense that (5.4) holds, and *(ii)* the algebraic solver BoomerAMG [21], denoted by $\mathtt{B}^o$, is a good preconditioner for the Gram matrix $\mathtt{G}$ of the $H^1(\Omega)$-inner product on U_h, in the sense that the spectral condition number is bounded independent of discretization size h and polynomial order p, that is,

$$\kappa(\mathtt{B}^o\mathtt{G}) \le c_\kappa. \tag{6.1}$$

Combining this with the $\mathtt{B}^f$ defined in (5.5), we have the following result.

THEOREM 6.1. *Assume that (2.6), (3.3), (5.4) and (6.1) hold. Then the block-diagonal matrix*

$$\begin{bmatrix} \mathtt{B}^o & \\ & \mathtt{B}^f \end{bmatrix} \tag{6.2}$$

is a preconditioner for $\mathtt{A}$ *and the condition number of the preconditioned system depends only on* c_1, c_2, c_3, c_5 *and* c_κ.

Proof. From (2.6), for any $x = (u, q) \in X$, we have

$$c_1(\|u\|^2_{H^1(\Omega)} + \|q\|^2_Q) \le (\mathtt{Ax}, \mathtt{x}) \le c_2(\|u\|^2_{H^1(\Omega)} + \|q\|^2_Q).$$

Using (3.3),

$$c_1\|u\|^2_{H^1(\Omega)} + c_1c_3^{-2}\|q\|^2_{Q_h} \le (\mathtt{Ax}, \mathtt{x}) \le c_2\|u\|^2_{H^1(\Omega)} + c_2\|q\|^2_{Q_h} \tag{6.3}$$

Hence the result follows from Lemma 5.1 and (6.1). ∎

In practice, the application of $\mathtt{B}^o$ and $\mathtt{B}^f$ requires the availability of the Gram matrices $\mathtt{G}$ and $\mathtt{D}$, which may be inconvenient. What we have in hand is $\mathtt{A}$. Hence instead of the preconditioner in (6.2), we may use the block preconditioner

$$\begin{bmatrix} \mathtt{P}^o & \\ & \mathtt{P}^f \end{bmatrix}$$

where $\mathtt{P}^o$ and $\mathtt{P}^f$ are the algebraic solvers BoomerAMG and ADS applied directly to the principal minors of $\mathtt{A}$ corresponding to U_h and Q_h, namely to $\mathtt{A}_0 = \mathtt{B}_0^T\mathtt{M}^{-1}\mathtt{B}_0$ and $\mathtt{A}_1 = \mathtt{B}_1^T\mathtt{M}^{-1}\mathtt{B}_1$ respectively. The justification for this comes from the observation that by taking $q = 0$ in (6.3), we can conclude that

$$c_1\|u\|^2_{H^1(\Omega)} \le (\mathtt{A}_0\mathtt{u}, \mathtt{u}) \le c_2\|u\|^2_{H^1(\Omega)},$$

i.e., $\mathtt{A}_0$ is spectrally equivalent to $\mathtt{G}$. Similarly, by taking $u = 0$ in (6.3), we have

$$c_1c_3^{-2}\|q\|^2_{Q_h} \le (\mathtt{A}_1\mathtt{q}, \mathtt{q}) \le c_2\|q\|^2_{Q_h}.$$

Thus instead of preconditioning the matrices $\mathtt{G}$ and $\mathtt{S}$, whose quadratic forms give the norms $\|\cdot\|^2_{H^1(\Omega)}$ and $\|\cdot\|^2_{Q_h}$ respectively, we can directly precondition their spectrally equivalent principal minors $\mathtt{A}_0$ and $\mathtt{A}_1$. In our implementation it is in fact straightforward to construct the Gram matrix $\mathtt{G}$, and we do so in order to build the AMG preconditioner $\mathtt{B}^o$, but we use the principal minor $\mathtt{A}_1$ to construct the ADS preconditioner $\mathtt{P}^f$, so that the preconditioner we use in the numerical results below takes the form

$$\begin{bmatrix} \mathtt{B}^o & \\ & \mathtt{P}^f \end{bmatrix}. \tag{6.4}$$

TABLE 7.1
Number of CG iterations and average reduction factors per iteration (in parenthesis) for various h and p refinement levels.

	order (p)				
elements	1	2	4	6	8
64	5 (0.06)	8 (0.14)	12 (0.30)	13 (0.34)	13 (0.34)
512	7 (0.12)	10 (0.23)	12 (0.31)	14 (0.36)	—
4096	8 (0.18)	10 (0.25)	13 (0.33)	—	—
32768	10 (0.22)	10 (0.24)	—	—	—
262144	10 (0.22)	—	—	—	—

7. Numerical results. In this section we report some numerical results with the proposed DPG preconditioner that test its performance with respect to the mesh size h, the polynomial order of the trial space p, as well as the orders of the test and interfacial spaces. We also examine the parallel scalability of the new algorithm and examine its behavior on more challenging problems with unstructured meshes and large coefficient jumps.

We apply a Conjugate Gradients (CG) solver to the problem (2.5) preconditioned with the preconditioner (6.4) where $\mathtt{B}^o$ and $\mathtt{P}^f$ use a single V-cycle of BoomerAMG and ADS respectively. The CG relative tolerance we used was 10^{-6}.

Our implementation is freely available in the MFEM finite element library [26] and we used a slightly modified version of MFEM's parallel Example 8 (version 3.2) to perform the numerical experiments in this section. Specific ADS and BoomerAMG parameters and additional details can be found in the source code of that example.

7.1. Scalability with respect to h, p for structured mesh. Here we solve the test problem (2.1) on the domain $\Omega = (0,1)^3$ with constant coefficient $\kappa = 1$ meshed with a uniform hexahedral grid. The right hand side f is set to the constant one and zero Dirichlet boundary conditions are imposed on all of $\partial\Omega$.

Table 7.1 reports results for experiments with varying mesh size h (reported as number of finite elements) and polynomial orders p. The order p sets the polynomial degree of U_h to p, the order of Q_h to $p-1$, and the order of Y_h to $p+d-1$ where $d = 3$ is the spatial dimension of Ω. As mentioned in Section 2, Assumption (2.6) holds in this setting. The table reports iteration counts as well as the average reduction factors in the PCG iteration. We observe that both of these convergence metrics are quite stable with respect to h and p.

In Table 7.2, we explore the parallel scalability of this algorithm, doing a weak scaling study where the number of elements per processor is kept almost constant as we increase the number of processors. This particular experiment uses a trial space order of $p = 1$ but a test space order of 2 rather than the theoretically necessary 3 (see the remarks in Section 7.2). The test was run on an IBM BlueGene/Q machine, where we use four MPI tasks per node.

While the number of iterations in Table 7.2 exhibits some growth, the overall performance is reasonably scalable, and we are continuing to work on improving the per-iteration run time in our implementation.

7.2. Influence of the order of the test space. Currently known theoretical results on the DPG method requires one to set the test space a few degrees higher than

TABLE 7.2
Weak scaling for the solver with polynomial order fixed at $p = 1$.

processors	elements	iterations	conv. factor	solve time	time/iteration
4	2.62e+5	9	0.21	249.42s	27.7s
32	2.10e+6	11	0.26	473.84s	43.1s
256	1.68e+7	12	0.29	547.95s	45.7s
2048	1.34e+8	13	0.32	665.81s	51.2s
16384	1.07e+9	14	0.37	745.69s	53.3s

TABLE 7.3
Effect of test space order r on scalability with respect to h-refinement for a triangular mesh.

	$p = 4$		$p = 5$	
refine	$r = 4$	$r = 5$	$r = 5$	$r = 6$
1	14	12	18	18
2	20	14	16	16
3	28	15	17	17
4	46	15	16	16
5	50	16	16	16
6	67	17	15	15

the trial space. Higher order test spaces can significantly add to the size of the discrete system (2.5) and the overall computational cost. Our numerical results indicate that test spaces of one degree lower than the theoretical requirement often continue to yield a scalable method. We have observed this for triangles, quadrilaterals, tetrahedra, and hexahedra. In Table 7.3, we present some representative results for the interesting case of triangles in two dimensions, where the scalability depends on the parity of p. For even $p = 4$, scalability requires a test space order one degree higher than for the odd order $p = 5$. The dependence of the error convergence rate on the parity of p was discussed in [6]. It is interesting to observe that our preconditioner also exhibits such dependence.

7.3. Scalability with respect to h on unstructured meshes. Next we consider problems with different meshes, including unstructured triangular, quadrilateral, and tetrahedral meshes, using in particular the meshes shown in Figure 7.1 at various levels of refinement. The problem is the same as in Section 7.1 except for the mesh. We fix $p = 1$ and focus on the scalability with respect to h. The convergence results in Table 7.4 demonstrate that the preconditioner continues to be scalable in these more general settings.

7.4. Behavior of solver with respect to contrast in coefficient. In the following numerical results the coefficient κ in (2.1) is piecewise constant, chosen randomly on each element, so that it is 1 with probability 1/2 and κ_0 with probability 1/2, where κ_0 is a specified constant across the mesh. Here the $\mathtt{B}^o$ component in (6.4) is constructed from an H^1 matrix assembled from the bilinear form in (2.1) using the varying coefficient κ, and $\mathtt{P}^f$ is constructed as usual using the principal minor of $\mathtt{A}$, which also includes the coefficient κ. In Table 7.5 we report the number of iterations and average reduction factor for several refinement levels and choice of contrast κ_0. As

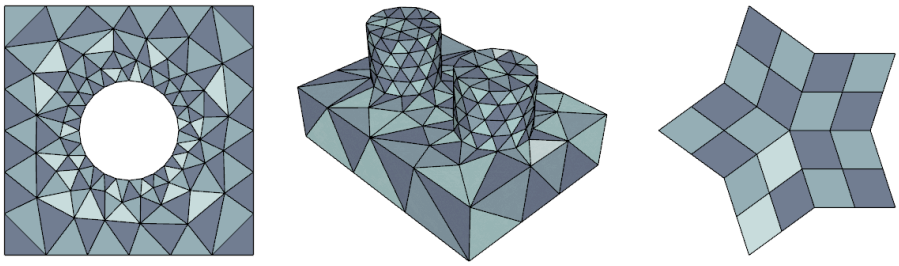

FIG. 7.1. *Three unstructured meshes, using triangles, tetrahedra, and quadrilaterals.*

TABLE 7.4
Number of CG iterations and average reduction factors per iteration (in parenthesis) for several unstructured meshes at various refinement levels.

refine	triangles	tetrahedra	quadrilaterals
0	13 (0.34)	8 (0.17)	9 (0.21)
1	14 (0.37)	11 (0.27)	12 (0.31)
2	14 (0.36)	13 (0.35)	13 (0.32)
3	14 (0.35)	15 (0.39)	13 (0.33)
4	14 (0.36)	16 (0.42)	12 (0.31)
5	14 (0.37)		12 (0.30)
6	15 (0.38)		12 (0.30)
7	15 (0.38)		12 (0.29)
8	15 (0.39)		12 (0.30)

the underlying discretization we use here is not robust to contrast in coefficients, we do not expect the solver to perform perfectly with respect to variation in coefficients. Nonetheless, we see that the preconditioner performs acceptably for a wide range of contrasts and we are able to solve the problem effectively.

8. Conclusions. In this paper we presented a scalable preconditioner for the primal DPG formulation of the Poisson problem based on parallel algebraic multigrid techniques. We proved that the preconditioner is optimal under certain assumptions on the mesh and problem coefficients. We also demonstrated that the new algorithm performs well on a wide variety of problems, including some where the theory is not applicable. Due to its algebraic nature, the preconditioner is easy to apply in practice, and has a freely available implementation in the MFEM library.

REFERENCES

[1] M. AINSWORTH AND L. DEMKOWICZ, *Explicit polynomial preserving trace liftings on a triangle*, Math. Nachr., 282 (2009), pp. 640–658. 8

[2] A. BAKER, R. FALGOUT, T. KOLEV, AND U. YANG, *Scaling hypre's multigrid solvers to 100,000 cores*, in High Performance Scientific Computing: Algorithms and Applications, Springer, 2012, pp. 261–279. LLNL-JRNL-479591. 10

[3] A. T. BARKER, S. C. BRENNER, E.-H. PARK, AND L.-Y. SUNG, *A one-level additive Schwarz preconditioner for a discontinuous Petrov-Galerkin method*, in Domain Decomposition Methods in Science and Engineering XXI, vol. 98 of Lecture Notes in Computational Science and Engineering, 2014, pp. 417–425. 1

TABLE 7.5
Number of CG iterations and average reduction factors per iteration (in parenthesis) for various values of the contrasts coefficient κ_0.

	contrast					
elements	1e-06	1e-04	1e-02	1e+00	1e+02	1e+04
64	13 (0.29)	12 (0.31)	10 (0.24)	5 (0.06)	8 (0.15)	12 (0.24)
512	31 (0.64)	29 (0.61)	14 (0.36)	7 (0.10)	11 (0.27)	17 (0.44)
4096	64 (0.80)	49 (0.75)	15 (0.39)	8 (0.16)	13 (0.33)	33 (0.64)
32768	119 (0.89)	73 (0.83)	16 (0.41)	9 (0.20)	14 (0.36)	43 (0.72)

[4] C. BERNARDI, *Optimal finite-element interpolation on curved domains*, SIAM J. Numer. Anal., 26 (1989), pp. 1212–1240. 8

[5] D. BOFFI, F. BREZZI, AND M. FORTIN, *Mixed Finite Element Methods and Applications*, vol. 44 of Springer Series in Computational Mathematics, Springer Berlin Heidelberg, 2013. doi: 10.1007/978-3-642-36519-5. 6, 8

[6] T. BOUMA, J. GOPALAKRISHNAN, AND A. HARB, *Convergence rates of the DPG method with reduced test space degree*, Computers and Mathematics with Applications, 68 (2014), pp. 1550–1561. 4, 14

[7] T. A. BRUNNER AND T. V. KOLEV, *Algebraic multigrid for linear systems obtained by explicit element reduction*, SIAM J. Sci. Comput., 33 (2011), pp. 2706–2731. 8, 11

[8] V. M. CALO, N. O. COLLIER, AND A. H. NIEMI, *Analysis of the discontinuous Petrov-Galerkin method with optimal test functions for the Reissner-Mindlin plate bending model*, Computers and Mathematics with Applications, 66 (2014), pp. 2570–2586. 5

[9] C. CARSTENSEN, L. DEMKOWICZ, AND J. GOPALAKRISHNAN, *Breaking spaces and forms for the dpg method and applications including maxwell equations*, Computers and Mathematics with Applications, 72 (2016), pp. 494–522. 1, 3

[10] J. CHAN, N. HEUER, T. BUI-THANH, AND L. DEMKOWICZ, *A robust DPG method for convection-dominated diffusion problems II: adjoint boundary conditions and mesh-dependent test norms*, Comput. Math. Appl., 67 (2014), pp. 771–795. 1

[11] M. COSTABEL, M. DAUGE, AND L. DEMKOWICZ, *Polynomial extension operators for H^1, H(curl) and H(div)-spaces on a cube*, Math. Comp., 77 (2008), pp. 1967–1999. 8

[12] L. DEMKOWICZ AND J. GOPALAKRISHNAN, *A class of discontinuous Petrov–Galerkin methods. Part I: The transport equation*, Comp. Meth. Appl. Math. Engrg., 199 (2010), pp. 1558–1572. 1, 4

[13] ———, *A class of discontinuous Petrov–Galerkin methods. Part II: Optimal test functions*, Num. Meth. Part. Diff. Eq., 27 (2011), pp. 70–105. 1, 4

[14] ———, *A class of discontinuous Petrov–Galerkin methods. Part IV: The optimal test norm and time–harmonic wave propagation in 1D*, J. Comp. Phys., 230 (2011), pp. 2406–2432. 1

[15] L. DEMKOWICZ AND J. GOPALAKRISHNAN, *A primal DPG method without a first-order reformulation*, Comp. Math. Applic., 66 (2013), pp. 1058–1064. 2, 3, 4

[16] L. DEMKOWICZ AND J. GOPALAKRISHNAN, *Discontinuous Petrov Galerkin (DPG) method*, in Encyclopedia of Computational Mechanics, Wiley Computational Mechanics Online, 2016 (to appear). 1

[17] L. DEMKOWICZ, J. GOPALAKRISHNAN, AND A. H. NIEMI, *A class of discontinuous Petrov–Galerkin methods. Part III: Adaptivity*, Appl. Numer. Math., (2011), p. in press. 1

[18] L. DEMKOWICZ, J. GOPALAKRISHNAN, AND J. SCHÖBERL, *Polynomial extension operators. Part III.*, Math. Comp., 81 (2012), pp. 1289–1326. 6

[19] J. GOPALAKRISHNAN AND W. QIU, *An analysis of the practical DPG method*, Mathematics of Computation, 83 (2014), pp. 537–552. 4

[20] J. GOPALAKRISHNAN AND J. SCHÖBERL, *Degree and wavenumber [in]dependence of Schwarz preconditioner for the DPG method*, in Spectral and High Order Methods for Partial Differential Equations (ICOSAHOM 2014), R. M. Kirby, M. Berzins, and J. S.Hesthaven, eds., no. 106 in Lecture Notes in Computational Science and Engineering, Springer, 2015, pp. 257–265. 1

[21] V. E. HENSON AND U. M. YANG, *BoomerAMG: A parallel algebraic multigrid solver and preconditioner*, Appl. Numer. Math., 41 (2002), pp. 155–177. 10, 12

[22] R. Hiptmair and J. Xu, *Nodal auxiliary space preconditioning in* **H**(**curl**) *and* **H**(div) *spaces*, SIAM J. Numer. Anal., 45 (2007), pp. 2483–2509 (electronic). 11

[23] T. V. Kolev and P. S. Vassilevski, *Parallel auxiliary space AMG solver for H(curl) problems*, J. Comput. Math., 27 (2009), pp. 604–623. 10

[24] ———, *Parallel auxiliary space AMG solver for H(div) problems*, SIAM J. Sci. Comput., 34 (2012), pp. A3079–A3098. 9, 10, 11

[25] X. Li and X. Xu, *Domain decomposition preconditioners for the discontinuous Petrov-Galerkin method*, ESAIM: Mathematical Modelling and Numerical Analysis, in press (2016). 1

[26] *MFEM: Modular finite element methods*. http://mfem.org. 13

[27] S. Nepomnyaschikh, *Domain decomposition methods*, in Lectures on advanced computational methods in mechanics, vol. 1 of Radon Ser. Comput. Appl. Math., Walter de Gruyter, Berlin, 2007, pp. 89–159. 11

[28] N. V. Roberts and J. Chan, *A geometric multigrid preconditioning strategy for DPG system matrices*, ArXiV Preprint: 1608.02567, (2016). 1

[29] C. Wieners and B. Wohlmuth, *Robust operator estimates and the application to substructuring methods for first-order systems*, ESAIM: Mathematical Modelling and Numerical Analysis, 48 (2014), pp. 1473–1494. 1