

## A SCALABLE AND ROBUST VERTEX-STAR RELAXATION FOR HIGH-ORDER FEM\*

PABLO D. BRUBECK<sup>†</sup> AND PATRICK E. FARRELL<sup>†</sup>

**Abstract.** Pavarino proved that the additive Schwarz method with vertex patches and a low-order coarse space gives a  $p$ -robust solver for symmetric and coercive problems [Pavanno, *Numer. Math.*, 66 (1993), pp. 493–515]. However, for very high polynomial degree it is not feasible to assemble or factorize the matrices for each patch. In this work we introduce a direct solver for separable patch problems that scales to very high polynomial degree on tensor product cells. The solver constructs a tensor product basis that diagonalizes the blocks in the stiffness matrix for the internal degrees of freedom of each individual cell. As a result, the nonzero structure of the cell matrices is that of the graph connecting internal degrees of freedom to their projection onto the facets. In the new basis, the patch problem is as sparse as a low-order finite difference discretization, while having a sparser Cholesky factorization. We can thus afford to assemble and factorize the matrices for the vertex-patch problems, even for very high polynomial degree. In the nonseparable case, the method can be applied as a preconditioner by approximating the problem with a separable surrogate. We demonstrate the approach by solving the Poisson equation and a  $H(\text{div})$ -conforming interior penalty discretization of linear elasticity in three dimensions at  $p = 15$ .

**Key words.** preconditioning, high-order, tensor product, additive Schwarz, sparse Cholesky

**MSC codes.** 65F08, 65N35, 65N55

**DOI.** 10.1137/21M1444187

**1. Introduction.** For problems with smooth solutions, high-order finite element methods offer very good convergence properties, and in some cases they do not exhibit locking phenomena found in low-order methods. Moreover, there exist optimal matrix-free algorithms for operator evaluation with high arithmetic intensity, arising from data locality, that make efficient use of modern parallel hardware architectures. Unfortunately, the conditioning of the stiffness matrix is severely affected by the polynomial degree of the approximation. In order to obtain practical iterative solvers, we require good preconditioners.

Optimal solvers are often obtained from a multiplicative multigrid V-cycle where the smoother consists of a domain decomposition method, such as additive Schwarz with a particular space decomposition. The multigrid algorithm is then accelerated by a Krylov subspace method, such as the preconditioned conjugate gradient (PCG) method. The choice of space decomposition in the relaxation is crucial for robustness with respect to the cell size  $h$ , the polynomial degree  $p$ , and parameters in the equation.

One of the cheapest relaxations, with  $\mathcal{O}(p^d)$  computational cost, is diagonal scaling, also known as point-Jacobi. The diagonally preconditioned Laplacian has a condition number of  $\mathcal{O}(p^{2d-2})$  [36]. This implies that the number of PCG iterations, and therefore the number of residual evaluations, is  $\mathcal{O}(p^{d-1})$ , incurring a total cost of

\*Submitted to the journal's Methods and Algorithms for Scientific Computing section September 3, 2021; accepted for publication (in revised form) June 27, 2022; published electronically September 20, 2022.

<https://doi.org/10.1137/21M1444187>

**Funding:** The work of the first author was supported by the University of Oxford Mathematical Institute Graduate Scholarship. The work of the second author was supported by the EPSRC grants EP/V001493/1 and EP/R029423/1.

<sup>†</sup>Mathematical Institute, University of Oxford, Oxford OX1 2JD, UK (brubeckmarti@maths.ox.ac.uk, patrick.farrell@maths.ox.ac.uk).

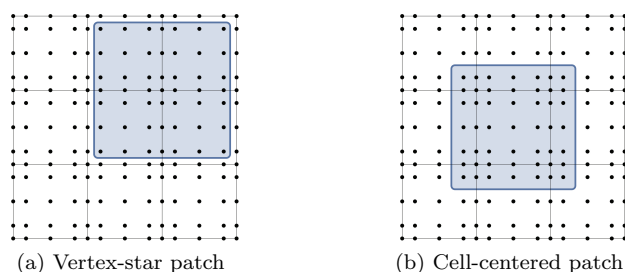


FIG. 1. Subdomains for the additive Schwarz method on a regular mesh ( $p = 4$ ). In combination with a low-order coarse grid, the vertex-star patch gives a  $p$ -robust method for symmetric and coercive problems, while the cell-centered patch does not.

$\mathcal{O}(p^{2d})$ . In order to minimize the time to solution, it is reasonable to consider more expensive relaxation methods that converge in fewer iterations. Ideally, we wish to balance the cost of applying the relaxation with that of updating the residual. On tensor product elements, the latter operation can be done quickly in  $\mathcal{O}(p^{d+1})$  operations via the sum-factorization. Sum-factorization breaks down the residual evaluation into products of one-dimensional operators and diagonal scalings [40].

In 1993, Pavarino proved that the additive Schwarz method with a vertex-centered space decomposition and an additive coarse space of lowest-order ( $p = 1$ ) gives a robust solver with respect to  $h$  and  $p$  for symmetric and coercive problems [42]. This type of space decomposition is often referred to as generous overlap and is illustrated in Figure 1. We use the terminology of [19] and refer to the subdomains in this space decomposition as vertex-star patches, as they are constructed by taking all the degrees of freedom (DOFs) on the topological entities in the *star* of each vertex.

The most straightforward implementation of a vertex-star solver involves the assembly and direct factorization of the  $\mathcal{O}(p^d) \times \mathcal{O}(p^d)$  patch matrices (which are dense for Lagrange shape functions). This becomes prohibitively expensive at very high polynomial degrees, with the Cholesky factorization of such a matrix requiring  $\mathcal{O}(p^{3d})$  operations. However, there exist bases for which the element matrices are sparse on affine cells, such as the hierarchical Lobatto shape functions [53]. In this basis, the stiffness matrix has a 5-point stencil in 2D, and a much larger 13-point stencil in 3D.

Efficient relaxation methods that are  $p$ -robust may arise from the discretization of an auxiliary problem for which fast inversion techniques are available. For a more general approach to auxiliary space techniques, we refer to the work of Xu [56]. In our context, the underlying PDE and/or the domain can be replaced by those of a problem which is solvable by the method of separation of variables. The fast diagonalization method (FDM) [34] is a  $\mathcal{O}(p^{d+1})$  direct factorization that breaks the problem down into a sequence of one-dimensional subproblems.

For the Poisson equation discretized on meshes with all cells Cartesian (all internal angles are right angles), the vertex-star problems can be solved directly with the FDM [55]. Huismann, Stiller, and J. Fröhlich [25] introduced a remarkably fast solver with  $\mathcal{O}(p^d)$  scaling on such meshes. The linear system is statically condensed by elimination of the cell DOFs, and the reduced system on the interface is solved with  $p$ -multigrid and a restricted variant of the FDM onto the interface DOFs of a vertex-star. Since the statically condensed operator requires the exact inversion of the cell matrices, their approach has no obvious extension to the unstructured, non-Cartesian case.

The FDM can be applied as a relaxation by means of an auxiliary problem that






FIG. 2. The FDM may be applied as a relaxation only on vertex-star patches that are structured, i.e., where the cells are laid out in a tensor product grid.

is separable, but this requires a tensor product grid discretization of the patch, which is only possible when the cells are laid out in a tensor product grid [55]. When the cells are not Cartesian, the method of Witte, Arndt, and Kanschä [55] approximates the whole patch as a single Cartesian domain and converges slowly even when the cells are slightly distorted. On general meshes, the patches may not have this structure, thus the FDM cannot be directly applied on such patches. An example of a vertex-star patch to which the FDM cannot be applied as a relaxation is shown in Figure 2.

A popular alternative in the literature has been to use cell-centered patches with minimal overlap by including a few layers of DOFs of the neighboring cells [20, 48, 52]. This can be done in such a way that every patch remains structured. This kind of space decomposition is more amenable to fast implementation but does not give a  $p$ -robust solver. If the number of layers is fixed, then the measure of the overlap region will decrease as  $p$  is increased. Pavarino also proved that when the overlap is not generous, the rate of convergence of the additive Schwarz method will depend inversely on the overlap size [43]. To overcome this, Fischer and Lottes [33] applied a hybrid  $p$ -multigrid/Schwarz method in the context of a Poisson problem. They implemented several levels of  $p$ -multigrid to overcome the lack of  $p$ -robustness of the cell-centered patches with minimal overlap. The use of multiple levels increases the overlap at the coarser levels with a minor impact on the overall computational cost. Cell-centered patches without overlap have also been employed for nonsymmetric problems [45, 46, 17].

Instead of replacing the vertex-star patches with cell-centered ones, the alternative finite element method-spectral element method (FEM-SEM) preconditioner [40] rediscretizes the problem on each vertex-star patch with  $p = 1$  on a Gauss-Lobatto-Legendre (GLL) grid, a mesh with vertices at the DOFs of the high-order space. The theory behind this guarantees the spectral equivalence between the differential operator discretized on the two spaces [13]. Since low-order methods are naturally sparse, this approach is not constrained to Cartesian cells and can deal properly with mixed first derivatives that the FDM cannot handle. A downside of this approach is that the Cholesky factors of the patch matrices are quite dense, limiting its scalability to very high polynomial degree. Computationally advantageous approaches involve incomplete factorizations of the patch FEM-SEM matrices [44] or the use of algebraic multigrid on the global FEM-SEM operator [9].

In this work we develop a solver for vertex-star patches that scales to very high polynomial degree. Our approach does not rely on a particular structure of the patch. In particular, it applies to the patches shown in Figure 2. The key idea is to use the

FDM to numerically construct a basis of functions on an interval that diagonalizes the interior blocks of *both* the stiffness and mass matrices in one dimension. When the problem is assembled with respect to the associated tensor product basis, the resulting stiffness and mass matrices are sparse in the Cartesian case. In particular, the total number of nonzeros is the same as that of a low-order finite difference discretization of the Laplacian. Moreover, fill-in in the Cholesky factorization is only introduced for the interface DOFs, resulting in very sparse Cholesky factors. The factorization requires  $\mathcal{O}(p^{3(d-1)})$  operations, while forward- and back-substitution steps have a cost of  $\mathcal{O}(p^{2(d-1)})$  operations that is optimal for  $d \in \{2, 3\}$ , in contrast with the  $\mathcal{O}(p^{3d})$  and  $\mathcal{O}(p^{2d})$  costs of the naïve approach. A disadvantage of the approach is that the memory required scales like  $\mathcal{O}(p^{2(d-1)})$ , instead of the optimal  $\mathcal{O}(p^d)$  required for storing the solution. In the non-Cartesian case, we approximate the form with one that is separable in the reference problem. Robustness with respect to  $h$  and  $p$  should follow from the spectral equivalence between the forms, and numerical experiments indicate that the approach is effective when the cells are moderately deformed.

We demonstrate the effectiveness of our approach by applying it to a  $H(\text{div}) \times L^2$  conforming discretization of a mixed formulation of incompressible linear elasticity. We present a sequence of problems of increasing complexity building up to this. In section 2 we present the standard  $hp$ -FEM formulation for the Poisson problem and construct a solver based on a sparse discretization of an auxiliary locally separable PDE that employs the numerically computed FDM basis. In section 3 we consider the application of our solver to linear elasticity. In the primal formulation, although our approach can be applied to patch problems for the individual components of displacement, we explain why it cannot be applied to the coupled vector-valued problem, which is necessary for parameter-robustness in the incompressible regime. We therefore consider a mixed formulation instead. Developing a  $p$ -robust solver requires both a  $p$ -robust preconditioner and a  $p$ -robust discretization, and for the latter we choose a  $H(\text{div}) \times L^2$  conforming approach. We then extend the method to symmetric interior-penalty discontinuous Galerkin discretizations, which are required for the displacement block of the mixed problem. We apply our relaxation to the displacement block of the incompressible elasticity system in conjunction with block-preconditioned Krylov methods. We end with conclusions in section 4.

**2. Sparse Poisson solver.** We will first describe a solver for the Poisson equation on Cartesian cells, which will subsequently be extended as a preconditioner for more general symmetric coercive problems on non-Cartesian cells.

**2.1. Continuous Galerkin formulation.** We start from the standard weak formulation of the Poisson equation. Consider a bounded domain  $\Omega \subset \mathbb{R}^d$ ,  $d \in \{1, 2, 3\}$ , and let  $\Gamma_D \subseteq \partial\Omega$  be the part of the boundary where the Dirichlet boundary condition  $u|_{\Gamma_D} = u_0$  is prescribed. The problem is to find  $u - u_0$  in  $V := H_0^1(\Omega) = \{v \in H^1(\Omega), v|_{\Gamma_D} = 0\}$  such that

$$(2.1) \quad a(v, u) = L(v) \quad \forall v \in V,$$

where

$$(2.2) \quad a(v, u) := \int_{\Omega} \nabla v \cdot \nabla u \, d\mathbf{x}, \quad L(v) := \int_{\Omega} v f \, d\mathbf{x}.$$

The standard FEM discretization employs a mesh  $\mathcal{T}_h = \{K\}$  of  $\Omega$ . In this work we consider quadrilateral and hexahedral cells, so that each cell  $K$  can be mapped

with a diffeomorphism  $F_K : \hat{K} \rightarrow K$  from the reference hypercube  $\hat{K} = \hat{\mathcal{I}}^d$ , where  $\hat{\mathcal{I}} = [-1, 1]$  is the reference interval. The approximate solution  $u_h \in V_h$  is sought in the space of piecewise continuous tensor product polynomials on each cell, i.e.,  $V_h := \mathbf{Q}_p(\Omega) \subset V$ . We first define the space of shape functions on  $\hat{K}$

$$(2.3) \quad \mathbf{Q}_p(\hat{K}) := \bigotimes_{j=1}^d \mathbf{P}_p(\hat{\mathcal{I}}), \quad \mathbf{P}_p(\hat{\mathcal{I}}) := \text{span} \{ \hat{x}^j, 0 \leq j \leq p \},$$

and via composition with  $F_K^{-1}$ , we define

$$(2.4) \quad \mathbf{Q}_p(\Omega) := \left\{ v \in H_0^1(\Omega) : \forall K \in \mathcal{T}_h \exists \hat{v} \in \mathbf{Q}_p(\hat{K}) \text{ s.t. } v|_K = \hat{v} \circ F_K^{-1} \right\}.$$

Once we fix a basis  $\{\phi_j\}_{j=1}^n$  for  $V_h$ , the approximate solution is expanded as  $u_h = \sum_{j=1}^n u_j \phi_j$ . The resulting  $n \times n$  system of linear equations is

$$(2.5) \quad A\mathbf{u} = \mathbf{f},$$

where  $[A]_{ij} = a(\phi_i, \phi_j)$  is the stiffness matrix,  $\mathbf{u} = (u_1, \dots, u_n)^\top$  is the vector of coefficients, and  $\mathbf{f} = (L(\phi_1), \dots, L(\phi_n))^\top$  is the load vector.

We recall the standard construction of the basis  $\{\phi_j\}$  [26]. The basis is defined in terms of shape functions  $\{\hat{\phi}_j\}$  on  $\hat{K}$ . Given shape functions  $\{\hat{\phi}_j^{1D}\}_{j=0}^p$  for  $\mathbf{P}_p(\hat{\mathcal{I}})$ , a tensor product basis  $\{\hat{\phi}_j\}$  for  $\mathbf{Q}_p(\hat{K})$  can be constructed as

$$(2.6) \quad \hat{\phi}_j(\hat{\mathbf{x}}) = \prod_{k=1}^d \hat{\phi}_{j_k}^{1D}(\hat{x}_k),$$

where we have expanded  $j = (j_1, \dots, j_d) \in [0, p]^d$  as a multi-index.

The interval shape functions are decomposed into interface and interior modes. The interface modes have nonzero support on either endpoint of  $\hat{\mathcal{I}}$ , while the interior modes vanish at the boundary of  $\hat{\mathcal{I}}$ . In multiple dimensions, the shape functions decompose into interior, facet, edge, and vertex modes, depending on how many one-dimensional interface functions are multiplied together. To generate a  $C^0$  basis, we simply match the shape of individual interface modes. Hence,  $A$  will be block sparse, since  $[A]_{ij} = 0$  when  $i$  and  $j$  correspond to interior modes supported on different cells.

For the interval shape functions, one standard choice is the set of Lagrange polynomials on the GLL nodes  $\{\hat{\xi}_i\}_{i=0}^p \subset [-1, 1]$ . These nodes are the roots of  $(1 - \hat{\xi}^2)P'_p(\hat{\xi})$ , where  $P_k(\hat{\xi})$  is the Legendre polynomial of degree  $k$ . The Lagrange polynomials  $\{\ell_j(\hat{x})\}$  satisfy  $\ell_j(\hat{\xi}_i) = \delta_{ij}$  by construction,

$$(2.7) \quad \ell_j(\hat{x}) = \prod_{k=0, k \neq j}^p \frac{\hat{x} - \hat{\xi}_k}{\hat{\xi}_j - \hat{\xi}_k}, \quad j = 0, \dots, p.$$

Another useful basis is formed by the hierarchical Lobatto shape functions  $\{l_j\}$ , which are constructed by augmenting the so-called bubble functions (integrated Legendre polynomials) with linear Lagrange functions,

$$(2.8) \quad l_j(\hat{x}) = \begin{cases} (1 - \hat{x})/2 & \text{for } j = 0, \\ (1 + \hat{x})/2 & \text{for } j = p, \\ \int_{-1}^{\hat{x}} P_j(z) \, dz & \text{for } j = 1, \dots, p-1. \end{cases}$$

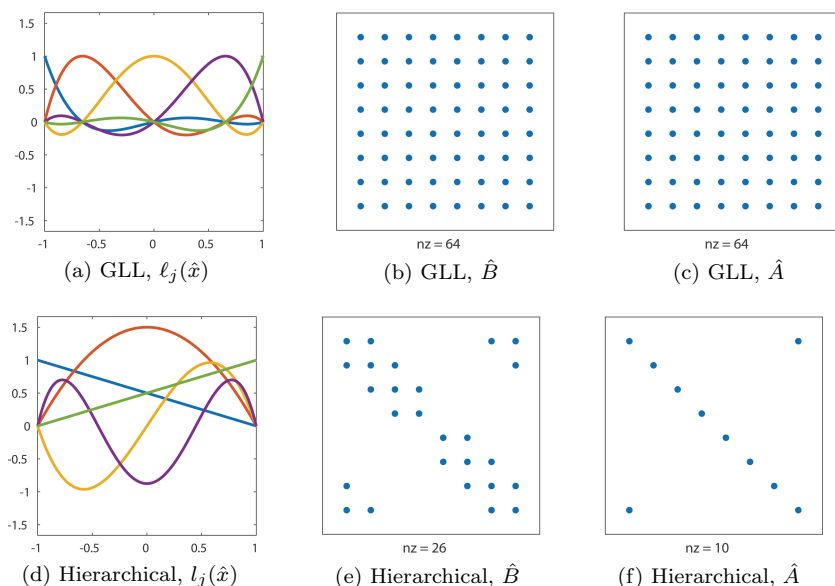


FIG. 3. Plots of the interval shape functions ( $p = 4$ ) and nonzero structure of the mass and stiffness matrices on the reference interval ( $p = 7$ ).

These two choices of shape functions are plotted in Figure 3.

The assembly of the stiffness matrix  $A$  is described as follows. On each cell  $K \in \mathcal{T}_h$  we define the cell stiffness matrix  $A^K \in \mathbb{R}^{(p+1)^d \times (p+1)^d}$  in terms of the basis functions  $\{\phi_j^K\}$  that are supported on  $K$ , which are obtained from the reference shape functions  $\{\hat{\phi}_j\}$  via  $\phi_j^K = \hat{\phi}_j \circ F_K^{-1}$ . Then, the cell stiffness matrices are

$$(2.9) \quad [A^K]_{ij} = \int_K \nabla \phi_i^K \cdot \nabla \phi_j^K \, dx.$$

The global stiffness matrix is then assembled via direct stiffness summation:

$$(2.10) \quad A = \sum_{K \in \mathcal{T}_h} R_K^\top A^K R_K,$$

where  $R_K \in \mathbb{R}^{(p+1)^d \times n}$  is the Boolean restriction matrix from the global DOFs to those local to cell  $K$ .

**2.2. Tensor product structure on Cartesian cells.** If  $d = 1$  and  $F_K$  is an affine mapping, then the cell stiffness matrices are

$$(2.11) \quad [A^K]_{ij} = \frac{1}{L^K} \int_{\hat{I}} \hat{\phi}_i' \hat{\phi}_j' \, d\hat{x} = \frac{1}{L^K} [\hat{A}]_{ij}.$$

Here  $\hat{A} \in \mathbb{R}^{(p+1) \times (p+1)}$  is the stiffness matrix on the reference interval, and  $L^K = |K|/|\hat{I}|$ , where  $|K|$  denotes the measure of the cell  $K$ .

For  $d = \{2, 3\}$ , we will first consider the case where  $\Omega$  can be tessellated with a mesh  $\mathcal{T}_h$  consisting of Cartesian cells; i.e., the cells are rectangular quadrilaterals

or hexahedra (all internal angles are right angles). In this idealized setting, the cell stiffness matrices are separable in the reference coordinates

$$(2.12) \quad A^K = \begin{cases} \mu_1^K \hat{B} \otimes \hat{A} + \mu_2^K \hat{A} \otimes \hat{B} & \text{if } d = 2, \\ \mu_1^K \hat{B} \otimes \hat{B} \otimes \hat{A} + \mu_2^K \hat{B} \otimes \hat{A} \otimes \hat{B} + \mu_3^K \hat{A} \otimes \hat{B} \otimes \hat{B} & \text{if } d = 3, \end{cases}$$

where

$$(2.13) \quad [\hat{B}]_{ij} = \int_{\hat{T}} \hat{\phi}_i \hat{\phi}_j \, d\hat{x}$$

is the mass matrix on the reference interval,  $\mu_j^K = (L_j^K)^{-2} \prod_{i=1}^d L_i^K$ , and  $L_j^K$  is the length of  $K$  along the  $j$ th axis divided by  $|\hat{T}|$ . The symbol  $\otimes$  denotes the Kronecker product, which for matrices  $A \in \mathbb{R}^{m \times n}$ ,  $B \in \mathbb{R}^{r \times s}$ , is defined as the block matrix

$$(2.14) \quad A \otimes B = \begin{bmatrix} a_{11}B & \cdots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \cdots & a_{mn}B \end{bmatrix} \in \mathbb{R}^{rm \times sn}.$$

It follows that if  $A$  and  $B$  are sparse, then  $A \otimes B$  is also sparse.

For the GLL basis  $\{\ell_j\}$ , both  $\hat{A}$  and  $\hat{B}$  are dense, but these are sparse in the hierarchical basis  $\{l_j\}$ . This is illustrated in Figure 3. On affine cells, the hierarchical basis yields a sparse stiffness matrix. The bubble functions  $\{l_j\}_{j=1}^{p-1}$  satisfy  $l'_j(\hat{x}) = P_j(\hat{x})$ , and due to the orthogonality of the Legendre polynomials, the interior block of  $\hat{A}$  is diagonal. The only off-diagonal nonzeros in  $\hat{A}$  are due to the coupling between the interface modes  $l_0, l_p$ . Nevertheless, in order for this sparsity to propagate to higher dimensions, we would additionally wish that  $\hat{B}$  is also as sparse as possible. This is not quite the case for  $\{l_j\}$ , as  $\hat{B}$  has two interior blocks with tridiagonal structure in the even-odd decomposition. Therefore, on a typical row,  $A$  will have the structure of the 5-point stencil for  $d = 2$  and that of a 13-point stencil for  $d = 3$ .

**2.3. The fast diagonalization method.** Linear systems involving structured matrices such as that defined in (2.12) can be solved efficiently using the FDM [34]. This method reduces the computation into a sequence of eigenvalue problems on the interval in a similar fashion as the method of separation of variables. It requires a separable PDE and a tensor product basis; therefore it can only be applied on meshes or mesh patches with tensor product structure.

To illustrate the FDM, we consider solving a problem on the interior of a Cartesian cell,  $A_{II}^K u_I^K = \underline{r}_I^K$ , where  $A_{II}^K = R_I^K A R_I^{K\top}$  and  $R_I^K \in \mathbb{R}^{(p-1)^d \times n}$  is the Boolean restriction matrix onto the interior DOFs of  $K$ . We may first solve the generalized eigenvalue problem on the interior of the reference interval

$$(2.15) \quad \hat{A}_{II} \hat{S}_{II} = \hat{B}_{II} \hat{S}_{II} \hat{\Lambda}_{II}$$

in conjunction with the normalization condition  $\hat{S}_{II}^\top \hat{B}_{II} \hat{S}_{II} = \mathbf{1}$ . Here  $\hat{A}_{II}, \hat{B}_{II} \in \mathbb{R}^{(p-1) \times (p-1)}$  are the interior blocks of  $\hat{A}$  and  $\hat{B}$ , respectively,  $\hat{\Lambda}_{II} \in \mathbb{R}^{(p-1) \times (p-1)}$  is the diagonal matrix of eigenvalues, and  $\hat{S}_{II} \in \mathbb{R}^{(p-1) \times (p-1)}$  is the matrix of eigenvectors. The generalized eigenproblem (2.15) may be equivalently rewritten as

$$(2.16) \quad \hat{S}_{II}^\top \hat{A}_{II} \hat{S}_{II} = \hat{\Lambda}_{II}, \quad \hat{S}_{II}^\top \hat{B}_{II} \hat{S}_{II} = \mathbf{1}.$$

The corresponding continuous problem is to find  $s_j(\hat{x})$  and  $\lambda_j$ ,  $j = 1, \dots, p-1$ , such that

$$(2.17) \quad \int_{\hat{\mathcal{T}}} s'_i(\hat{x}) s'_j(\hat{x}) \, d\hat{x} = \lambda_j \delta_{ij}, \quad \int_{\hat{\mathcal{T}}} s_i(\hat{x}) s_j(\hat{x}) \, d\hat{x} = \delta_{ij}, \quad s_j(\pm 1) = 0,$$

with solution  $s_j(\hat{x}) = \sin(j\pi(1 + \hat{x})/2)$ ,  $\lambda_j = j^2\pi^2/4$ . Thus, when  $\hat{A}$  and  $\hat{B}$  are discretized with the GLL basis,  $\hat{S}_{II}$  approximates a discrete sine transform on the interior GLL nodes.

If  $A^K$  is given by (2.12), then its inverse has the following diagonal factorization

$$(2.18) \quad (A_{II}^K)^{-1} = \left\{ \bigotimes_{k=1}^d \hat{S}_{II} \right\} (\Lambda_{II}^K)^{-1} \left\{ \bigotimes_{k=1}^d \hat{S}_{II}^\top \right\},$$

where

$$(2.19) \quad \Lambda_{II}^K = \begin{cases} \mu_1^K \mathbb{1} \otimes \hat{\Lambda}_{II} + \mu_2^K \hat{\Lambda}_{II} \otimes \mathbb{1} & \text{if } d = 2, \\ \mu_1^K \mathbb{1} \otimes \mathbb{1} \otimes \hat{\Lambda}_{II} + \mu_2^K \mathbb{1} \otimes \hat{\Lambda}_{II} \otimes \mathbb{1} + \mu_3^K \hat{\Lambda}_{II} \otimes \mathbb{1} \otimes \mathbb{1} & \text{if } d = 3. \end{cases}$$

Therefore, the solution of a system  $A_{II}^K \underline{u}_I^K = \underline{r}_I^K$  can be obtained with  $\mathcal{O}(p^{d+1})$  computational work.

The main limitation of this approach is that it does not generalize to terms that contain first derivatives, ruling out the possible extension to advection problems. Mixed first derivative terms are also very common in symmetric coercive problems, for instance, when the cells have nonorthotropic deformations, or for vector-valued operators that mix first derivatives of distinct vector components, such as  $\nabla(\nabla \cdot \mathbf{u})$ .

**2.4. Sparse FDM discretization.** Our key idea is to construct a new finite element basis on the interval, inspired by the FDM, which yields a sparse stiffness matrix. To construct this new basis, we solve (2.16) with  $\hat{A}$  and  $\hat{B}$  in the GLL basis. Then, we interpolate the eigenvectors of the FDM with polynomials  $\{\hat{s}_j\}_{j=1}^{p-1} \subset P_p(\hat{\mathcal{T}})$  that satisfy  $\hat{s}_j(\pm 1) = 0$ ,  $\hat{s}_j(\hat{\xi}_i) = [\hat{S}_{II}]_{ij}$  for  $i, j \in [1, p-1]$ . The unisolvent dual basis of our proposed FDM discretization of  $P_p(\hat{\mathcal{T}})$  consists of point evaluation at the vertices and integral moments against the orthogonal interior shape functions  $\hat{s}_j$ ,  $j = 1, \dots, p-1$ . This construction ensures that  $\hat{A}_{II}$  and  $\hat{B}_{II}$  become diagonal under this basis. As an additional consequence, the one-dimensional interface shape functions are also orthogonal to the interior ones but not to each other. Hence  $\hat{B}$  becomes as sparse as possible in this new FDM basis. Figure 4 shows the shape functions and the nonzero structure of  $\hat{A}$  and  $\hat{B}$  for the FDM basis.

The stiffness matrix for a Cartesian cell  $A^K$  in this FDM basis is also given by (2.12) and has structured sparsity, since  $A \otimes B$  is sparse when  $A$  and  $B$  are sparse. Consequently, the global matrix  $A$  is also sparse. We can thus apply sparse direct factorization methods or other preconditioners such as the additive Schwarz method.

For problems that are separable in Cartesian coordinates, we obtain a solver that is reminiscent of the FDM. The multiplication times the Kronecker product of the matrices of eigenvectors is incorporated in the definition of the FDM shape functions. If for some reason the problem needs to be solved with any other basis, one can transform the linear system back and forth to the FDM basis via interpolation and restriction at each application of the method. We also replace the diagonal matrix of the traditional FDM with a sparse matrix, which can be assembled even




FIG. 4. Plots of the interval FDM shape functions ( $p = 4$ ) and nonzero structure of the mass and stiffness matrices in the FDM basis on the reference interval ( $p = 7$ ).

on three-dimensional meshes with a very large number of cells and for very high degree. This approach removes the requirement of a global tensor product grid while exploiting the separability of the PDE and the local structure of  $\mathbf{Q}_p(\Omega)$  at the expense of replacing the diagonal factor by a sparse matrix.

**2.5. Hybrid  $p$ -multigrid/Schwarz method.** The solver of Pavarino is fully additive, across both the coarse grid and the vertex-star patches. In our work we consider a small variation of this, with the solver multiplicative between the two levels while remaining additive among the vertex-star patches. This improves the convergence at essentially no cost. The method can be interpreted as a hybrid multiplicative two-level  $V(1, 1)$ -cycle with the additive Schwarz method [18] with vertex-star patches as the fine grid relaxation and the lowest-order discretization on the same mesh as the coarse space. The sparse matrix for the coarse space may be assembled and factorized, or other preconditioners such as geometric or algebraic multigrid may be applied instead. The vertex-star patch  $V_j$  includes the DOFs associated with vertex  $\mathbf{v}_j$  of  $\mathcal{T}_h$  and all cells, facets, and edges adjacent to  $\mathbf{v}_j$  (the topological entities in the *star* of the vertex, a standard concept in algebraic topology [38, section 2]).

We may write the multigrid relaxation as

$$(2.20) \quad P_{\text{ASM}}^{-1} = \sum_{j=1}^J \bar{R}_j^\top A_j^{-1} \bar{R}_j,$$

with  $\bar{R}_j$  the Boolean restriction matrix onto  $V_j$ , and  $A_j = \bar{R}_j A \bar{R}_j^\top$  are the sparse patch matrices for which we may explicitly compute a Cholesky decomposition. The relaxation is scaled by the damping coefficient

$$(2.21) \quad \omega = 2[(1 + \alpha)\lambda_{\max} + (1 - \alpha)\lambda_{\min}]^{-1},$$

where  $\lambda_{\min}, \lambda_{\max}$  are the extremal eigenvalues of  $P_{\text{ASM}}^{-1}A$  estimated via the CG-Lanczos procedure [30], and  $\alpha = 0.25$  is chosen to tackle the high frequency error, also ensuring that the error iteration matrix  $\mathbb{I} - \omega P_{\text{ASM}}^{-1}A$  is contractive.

To illustrate the direct solver on the Cartesian vertex-star patch shown in Figure 1(a), we show in Figure 5 the nonzero structure for the patch matrix  $A_j$  and its Cholesky factor. The sparsity pattern of the global matrix  $A$  connects the interior DOFs to their projections onto the facets; hence a typical interior row of  $A$  will have  $2d+1$  nonzeros. For the patch matrix  $A_j$ , an interior row will only have  $d+1$  nonzeros, as the patch only includes one facet per dimension on each cell. Moreover, the total

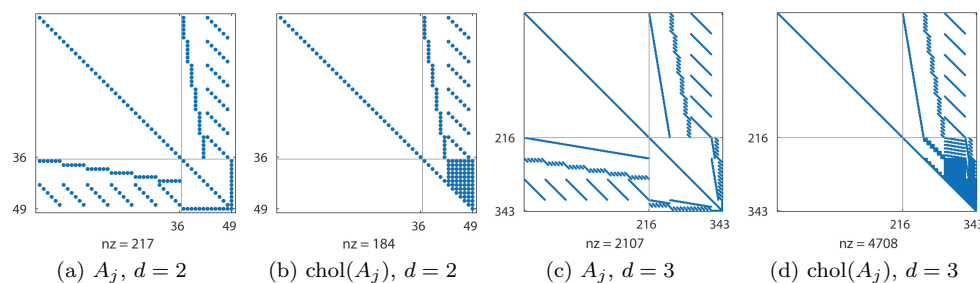


FIG. 5. Nonzero structure of the stiffness matrix in the FDM basis  $A_j = L_j L_j^\top$  and its upper Cholesky factor  $L_j^\top$  for the Poisson problem on a Cartesian vertex-star patch with  $p = 4$ . With the nested dissection reordering, the factor matrix has minimal fill-in, occurring only on the bottom-right block.

number of nonzeros of  $A_j$  is the same as that of a low-order finite difference or finite element discretization with the 5-point or 7-point stencil on the same grid.

**2.6. Computational complexity.** Here we discuss the computational cost of the solution of the patch problem using the Cholesky factorization. Once the factorization has been computed, it may be applied in  $\mathcal{O}(p^{2(d-1)})$  cost, which is optimal for  $d \in \{2, 3\}$ . Unfortunately, the factorization phase is suboptimal, requiring  $\mathcal{O}(p^{3(d-1)})$  operations to compute. The memory required to store the Cholesky factor is  $\mathcal{O}(p^{2(d-1)})$ , which for  $d = 3$  is one factor of  $p$  higher than that required to store the solution.

Consider a stiffness matrix  $A$  discretized with the FDM basis on any mesh with all cells Cartesian. The number of floating point operations (flops) needed to solve a linear system using a sparse Cholesky factorization  $A = LL^\top$  is roughly four times the number of nonzero entries in  $L$  [14]. To maximize the sparsity in  $L$ , it is crucial to reorder the DOFs, such that interior DOFs are followed by the interface DOFs. This ensures that the fill-in is introduced only on the bottom-left block. To analyze the cost of factorization, we first introduce the block  $LDL^\top$  decomposition

$$(2.22) \quad A = \begin{bmatrix} A_{II} & A_{I\Gamma} \\ A_{\Gamma I} & A_{\Gamma\Gamma} \end{bmatrix} = \begin{bmatrix} \mathbf{1} & 0 \\ A_{\Gamma I} A_{II}^{-1} & \mathbf{1} \end{bmatrix} \begin{bmatrix} A_{II} & 0 \\ 0 & S_\Gamma \end{bmatrix} \begin{bmatrix} \mathbf{1} & A_{II}^{-1} A_{I\Gamma} \\ 0 & \mathbf{1} \end{bmatrix},$$

where  $S_\Gamma = A_{\Gamma\Gamma} - A_{\Gamma I} A_{II}^{-1} A_{I\Gamma}$  is the interface Schur complement. By construction, the top-left block  $A_{II}$  is diagonal with positive entries, with Cholesky factor  $A_{II}^{1/2}$ . If we decompose  $A_{II}$  and  $S_\Gamma$  in the second matrix on the right-hand side of (2.22) into their Cholesky factors and distribute each factor onto the other two matrices, we obtain the Cholesky decomposition of  $A$ :

$$(2.23) \quad A = LL^\top = \begin{bmatrix} A_{II}^{1/2} & 0 \\ A_{\Gamma I} A_{II}^{-1/2} & L_\Gamma \end{bmatrix} \begin{bmatrix} A_{II}^{1/2} & A_{II}^{-1/2} A_{I\Gamma} \\ 0 & L_\Gamma^\top \end{bmatrix},$$

where the Schur complement is factorized as  $S_\Gamma = L_\Gamma L_\Gamma^\top$ . Since  $A_{II}$  is diagonal, the off-diagonal block  $A_{\Gamma I} A_{II}^{-1/2}$  will preserve the nonzero structure of  $A_{\Gamma I}$  and similarly for its transpose. Thus, fill-in is only introduced on the interface block through  $L_\Gamma$ .

An ordering strategy that minimizes fill-in consists of applying nested dissection [21] on the adjacency graph that connects topological entities. Each node in this






FIG. 6. Relative number of non-zero entries in the Cholesky factors of the stiffness matrix with the FDM approach and the FEM-SEM preconditioner, for a Cartesian vertex-star patch of  $2^d$  cells. The FDM approach is sparser, with substantial gains at higher degrees.

graph represents a cell, face, edge, or vertex. The ordering of the entities is then used to permute the corresponding blocks in  $A$ .

Assuming for the worst case that  $L_\Gamma^\top$  is dense, the memory required to store the Cholesky factor is  $\mathcal{O}(p^{2(d-1)})$ . This represents a significant increase from the traditional FDM, which is kept at the optimal  $\mathcal{O}(p^d)$ . However, the DOF ordering does lead to some structured sparsity in  $L_\Gamma^\top$ , as can be seen in Figure 5. Nevertheless, we still observe dense  $\mathcal{O}(p^{d-1}) \times \mathcal{O}(p^{d-1})$  blocks. The fact that  $L_\Gamma^\top$  contains these dense blocks indicates that  $\mathcal{O}(p^{3(d-1)})$  operations are required in the factorization phase. However, the forward- and back-substitution steps have a computational cost of  $\mathcal{O}(p^{2(d-1)})$  operations, which is optimal for  $d \leq 3$ .

Compared to the FEM-SEM approach, our approach with the FDM basis has a sparser Cholesky factorization. In Figure 6 we present the ratio of the number of nonzeros in the Cholesky factors of our approach and the FEM-SEM preconditioner ordered with nested dissection for  $d \in \{2, 3\}$ . The fact that the ratio is always below 1 confirms that our approach is sparser, with a substantial gain at higher degrees. Practical FEM-SEM solvers use algebraic multigrid [9] or patchwise multigrid with incomplete LU smoothers [44] to avoid the cost of the Cholesky factorizations of the patch matrices.

**2.7. Extension to non-Cartesian cells.** For arbitrarily deformed cells, the local stiffness matrices  $A^K$  cannot be expressed in terms of tensor products of  $\hat{A}$  and  $\hat{B}$  as in (2.12), and  $A^K$  is not sparse in the FDM basis. The preconditioning techniques found in [16, 20, 55] introduce an auxiliary Cartesian domain to construct a separable problem for which the FDM is a direct solver. The method described by Fischer [20] constructs a preconditioner by replacing  $K$  with its nearest rectangular approximation, whose dimensions are computed as the mean separation between the mapped GLL nodes from opposite facets of  $K$ . Witte, Arndt, and Kanschat [55] obtain the lengths from the average arclength of opposite sides of  $K$ , but it is not clear how this extends to the three-dimensional case. To the best of our knowledge, no theory underpins these choices.

Our approach to construct the separable surrogate is based on the theory of equivalent operator preconditioning [6]. We work with the bilinear form  $a(\cdot, \cdot)$  in terms of the reference coordinates. We discard the mixed derivative terms that prevent separability, and we replace the coefficients with piecewise constants in the reference coordinates.<sup>1</sup> We will prove that this choice yields a spectrally equivalent operator.

<sup>1</sup>Recall that piecewise constant coefficients in the physical coordinates will not yield piecewise constant coefficients in the reference coordinates.

The bilinear form  $a(\cdot, \cdot)$  can be expressed as a sum of cell contributions  $a_K(\cdot, \cdot)$ , where integration and differentiation are with respect to  $\hat{\mathbf{x}}$ . The measure  $d\mathbf{x}$  is replaced by  $|\det(DF_K)|d\hat{\mathbf{x}}$ , and the gradient is computed via the chain rule, since the arguments of the form become functions of  $\hat{\mathbf{x}}$  after being composed with  $F_K$ . Hence,

$$(2.24) \quad a(v, u) = \sum_{K \in \mathcal{T}_h} a_K(v, u) = \sum_{K \in \mathcal{T}_h} \int_{\hat{K}} \hat{\nabla} v \circ F_K \cdot \hat{G}^K \hat{\nabla} u \circ F_K \, d\hat{\mathbf{x}},$$

where  $\hat{\nabla}$  is the gradient with respect to  $\hat{\mathbf{x}}$ , and  $\hat{G}^K : \hat{K} \rightarrow \mathbb{R}^{d \times d}$  is the inverse metric of the coordinate transformation weighted by the Jacobian determinant,

$$(2.25) \quad \hat{G}^K = |\det(DF_K)| DF_K^{-1} DF_K^{-\top}.$$

This tensor encapsulates all of the geometry-dependent information in the form; it is spatially dependent for generally deformed elements and constant in the case of affine transformations. For a separable geometry,  $\hat{G}^K$  is diagonal, and thus for a Cartesian cell it is diagonal and constant. To construct an auxiliary problem that is separable by the FDM in the reference coordinates, we replace  $\hat{G}^K$  in  $a_K(\cdot, \cdot)$  with a constant diagonal approximation  $\text{diag}(\mu_j^K)$ . Each  $\mu_j^K$  is given by the cellwise average of the diagonal entry  $\hat{G}_{jj}^K$ ,

$$(2.26) \quad \mu_j^K := \frac{1}{|\hat{K}|} \int_{\hat{K}} \hat{G}_{jj}^K \, d\hat{\mathbf{x}},$$

where summation over the index  $j$  is not implied. As the approximation is local to each cell, it is still possible to assemble a sparse stiffness matrix discretizing the auxiliary problem on meshes where cells are not structured in a tensor product grid.

We now establish the spectral equivalence between the original bilinear form and the auxiliary separable one.

**THEOREM 2.1.** *Let  $\hat{\mu}_K := \text{diag}(\mu_j^K)$  be the constant diagonal approximation of  $\hat{G}^K$ , and define the auxiliary bilinear form*

$$(2.27) \quad \tilde{a}(v, u) := \sum_{K \in \mathcal{T}_h} \tilde{a}_K(v, u) := \sum_{K \in \mathcal{T}_h} \int_{\hat{K}} \hat{\nabla} v \circ F_K \cdot \hat{\mu}_K \hat{\nabla} u \circ F_K \, d\hat{\mathbf{x}}.$$

*Then, there exist  $p$ -independent constants  $c, C > 0$  that depend on  $\mathcal{T}_h$  through  $\hat{G}^K$  such that*

$$(2.28) \quad c \leq \frac{a(v, v)}{\tilde{a}(v, v)} \leq C \quad \forall v \in V \setminus \{0\}.$$

*Proof.* Let  $c_K, C_K$  be lower and upper bounds for the spectrum of the diagonally scaled metric, so that  $\sigma(\hat{\mu}_K^{-1/2} \hat{G}^K \hat{\mu}_K^{-1/2}) \in [c_K, C_K] \, \forall \hat{\mathbf{x}} \in \hat{K}$ . We claim that

$$(2.29) \quad c_K \leq \frac{a_K(v, v)}{\tilde{a}_K(v, v)} \leq C_K \quad \forall v \in \{v \in V : v|_K \neq 0\}.$$

This result is obtained by first rewriting  $a_K(v, v)$  with  $\hat{\mu}_K^{1/2} \hat{\mu}_K^{-1/2} \hat{G}^K \hat{\mu}_K^{-1/2} \hat{\mu}_K^{1/2}$  instead of  $\hat{G}^K$  and then replacing  $\hat{\mu}_K^{-1/2} \hat{G}^K \hat{\mu}_K^{-1/2}$  with  $c_K \mathbb{1}$  or  $C_K \mathbb{1}$  to find the lower or upper bounds, respectively. It then follows that

$$(2.30) \quad c := \min_{K \in \mathcal{T}_h} c_K \leq \frac{a(v, v)}{\tilde{a}(v, v)} \leq \max_{K \in \mathcal{T}_h} C_K =: C \quad \forall v \in V \setminus \{0\}. \quad \square$$

Let  $\tilde{A}$  be the stiffness matrix associated with the auxiliary form  $\tilde{a}(\cdot, \cdot)$ . By Theorem 2.1, the condition number  $\kappa(\tilde{A}^{-1}A)$  is bounded by  $C/c$  independently of  $p$ . Numerical experiments also indicate that  $\kappa(\tilde{A}^{-1}A)$  is independent of  $h$  under uniform refinements. Now consider a preconditioner  $P$  where the auxiliary form  $\tilde{a}(\cdot, \cdot)$  is used additively in both the coarse solve and the vertex-star patches. In this case, Theorem 1 of [42] guarantees that  $\kappa(P^{-1}\tilde{A})$  is bounded from above independently of  $h$  and  $p$ . Hence we may conclude that  $\kappa(P^{-1}A) \leq \kappa(P^{-1}\tilde{A})\kappa(\tilde{A}^{-1}A)$  is bounded independently of  $h$  and  $p$ . In practice, we expect that using multiplicative coarse grid correction with the original form  $a(\cdot, \cdot)$  can only improve the preconditioner.

To gain useful insight, we consider the case where  $d = 2$  and  $F_K$  is an affine transformation, that is, when  $K$  is a parallelogram. Without loss of generality, suppose that one of the sides of  $K$  has length  $2L_1$  and is aligned with the first reference coordinate axis, and the other side of length  $2L_2$  is at an angle  $\theta$  with respect to the same axis. The Jacobian of the coordinate transformation is

$$(2.31) \quad DF_K = \begin{bmatrix} L_1 & L_2 \cos \theta \\ 0 & L_2 \sin \theta \end{bmatrix},$$

to which corresponds the Jacobian-weighted inverse metric

$$(2.32) \quad \hat{G}^K = \frac{1}{L_1 L_2 |\sin \theta|} \begin{bmatrix} L_2^2 & -L_1 L_2 \cos \theta \\ -L_1 L_2 \cos \theta & L_1^2 \end{bmatrix}.$$

Since  $\hat{G}^K$  is constant,  $\hat{\mu}_K$  is simply the diagonal part of  $\hat{G}^K$ . The spectrum of the diagonally scaled metric will be independent of  $L_1$  and  $L_2$  but still depend on  $\theta$ ,

$$(2.33) \quad \sigma(\hat{\mu}_K^{-1/2} \hat{G}^K \hat{\mu}_K^{-1/2}) = [1 - |\cos \theta|, 1 + |\cos \theta|].$$

This spectrum is desirable because it is centered at 1 and bounded above for all  $\theta$ . If we follow the geometric approaches of [20, 55], we would have to choose a rectangle with side lengths  $2L_1$  and  $2L_2$  as the auxiliary domain for the Poisson problem. Then, the previous bounds (2.33) would become scaled by  $|\sin \theta|^{-1}$ . In this case, the spectrum is unbounded from above in the limit  $\theta \rightarrow 0$ .

**2.8. Numerical experiments.** We provide an implementation of the  $P_p$  element with the FDM shape functions on the interval in the FIAT [28] package. The extension to quadrilaterals and hexahedra is achieved by taking tensor products of the one-dimensional element with FInAT [23]. Code for the sum-factorized evaluation of the residual is automatically generated by Firedrake [47, 24], implementing a Gauss–Lobatto quadrature rule with  $3(p+1)/2$  points along each direction. The sparse preconditioner discretizing the auxiliary form is implemented in `firedrake.FDMPC` using PETSc [8]. The Cholesky factorization of the patch matrices is computed using CHOLMOD [14]. Most of our computations were performed using an Intel Xeon CPU E5-4627 v2 @ 3.30 GHz with 32 cores and 67.6 GB of RAM storage.

The hybrid  $p$ -multigrid/Schwarz solver employing the FDM/sparse relaxation is illustrated in Figure 7. To achieve scalability with respect to the mesh parameter  $h$ , on the  $p$ -coarse problem we employ geometric multigrid with damped point-Jacobi relaxation and a Cholesky factorization on the coarsest level using MUMPS [2]. We test the effectiveness of this approach on a hierarchy of meshes obtained by  $l \geq 0$  uniform refinements of the base meshes shown in Figure 8.

We present results for the Poisson equation in  $\Omega = (0, 1)^d$  discretized on the three hierarchies of Cartesian, unstructured, and structured but deformed (Kershaw)

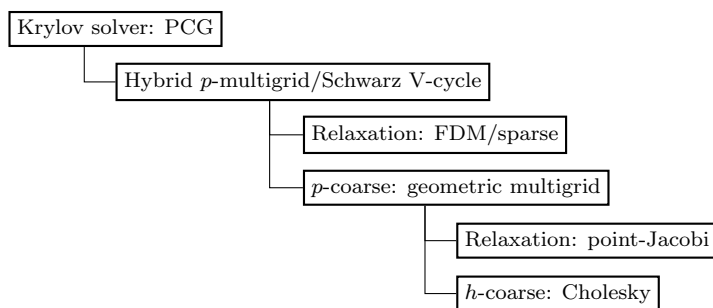


FIG. 7. Solver diagram for the Poisson problem.

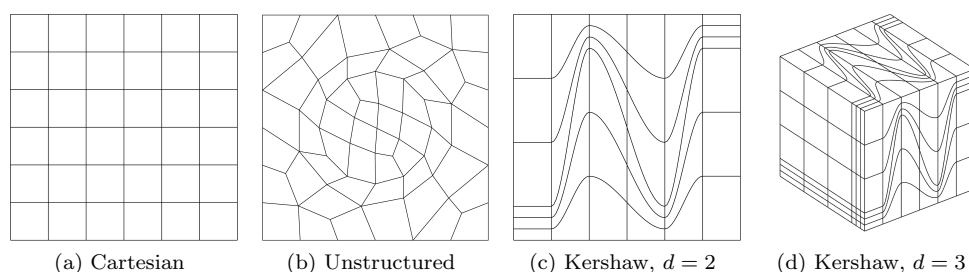


FIG. 8. Base meshes for the Poisson problem. The Cartesian and unstructured base meshes used for  $d = 3$  are the extrusion with six layers of the two-dimensional meshes shown here.

TABLE 1  
PCG iteration counts for the hybrid  $p$ -multigrid/Schwarz solver with the FDM/sparse relaxation. The patch problems are solved exactly on the Cartesian mesh.

$d$	$p \setminus l$	Cartesian			Unstructured			Kershaw		
		0	1	2	0	1	2	0	1	2
2	3	7	8	9	12	13	14	27	35	54
	7	8	8	9	16	16	17	44	56	78
	15	8	9	9	19	19	19	58	69	90
	31	8	9	9	21	20	21	67	80	97
3	3	12	12	12	17	17	18	54	66	102
	7	12	12	12	22	21	21	98	106	158
	15	12	13		25	24		131	132	

[27] meshes. The coordinate field of the Kershaw mesh is in  $[Q_3(\Omega)]^d \cap C^1(\Omega)$ , with a cell aspect ratio of  $\varepsilon_y = \varepsilon_z = 0.3$  near the corners of the domain. We impose homogeneous Dirichlet BCs on  $\Gamma_D = \partial\Omega$  and a constant forcing  $f = 1$ . In Table 1 we present PCG iteration counts required to reduce the Euclidean norm of the residual by a factor of  $10^8$  starting from a zero initial guess. In Table 2 we show the condition number  $\kappa(P^{-1}A)$  estimated by CG-Lanczos. The results show almost complete  $p$ - and  $h$ -robustness in the Cartesian case, and very slow growth of iteration counts in the unstructured case. Given the lack of shape regularity, the Kershaw mesh is significantly more challenging; even with exact patch solvers, we do not expect  $h$ - or  $p$ -robustness.

To assess the computational performance of our approach, we solve the three-dimensional Poisson equation on a Cartesian mesh with  $3 \times 3 \times 3$  cells with a single

TABLE 2

Estimated condition numbers for the preconditioned operator  $\kappa(P^{-1}A)$  using the hybrid  $p$ -multigrid/Schwarz solver with the FDM/sparse relaxation.

$d$	$p \setminus l$	Cartesian			Unstructured			Kershaw		
		0	1	2	0	1	2	0	1	2
2	3	1.44	1.49	1.50	2.14	2.37	2.81	9.34	15.6	34.5
	7	1.48	1.48	1.50	3.23	3.27	3.79	19.6	30.3	57.6
	15	1.51	1.51	1.52	4.06	3.78	4.13	30.5	45.8	69.0
	31	1.54	1.52	1.52	4.45	4.06	4.36	40.4	57.1	73.3
3	3	2.87	2.49	2.45	4.16	4.21	4.55	34.8	46.1	117
	7	2.79	2.70	2.67	5.88	5.54	5.47	100	110	266
	15	2.83	2.79		7.12	6.44		165	151	




FIG. 9. Runtimes, flop counts, and achieved arithmetic performance for the Cholesky factorization (symbolic and numeric), solution of the patch problems, and residual evaluation for a Cartesian mesh with  $3 \times 3 \times 3$  cells on a single CPU core. We observe that the factorization runtime scales better than expected, close to the optimal  $\mathcal{O}(p^{d+1})$  complexity.

core of an Intel Core i7-10875H CPU @ 2.30 GHz. We plot in Figure 9 the runtimes, flop counts, and achieved arithmetic performance for the Cholesky factorization of the patch matrices, the solution of the patch problems using this factorization (per application of the relaxation), and the matrix-free evaluation of the residual (per Krylov iteration, excluding the application of the global to local map) as functions of  $p$ . The dotted lines are to indicate powers of  $2p - 1$ , which is the number of DOFs along each side of typical vertex-star patch not intersecting the mesh boundary.

Despite the  $\mathcal{O}(p^{3(d-1)})$  computational cost of the Cholesky factorization, these results show  $\mathcal{O}(p^{2(d-1)})$  scaling for runtime up to  $p = 15$ . This speedup can be explained mainly by data locality. The sparse Cholesky factorization is obtained by recursively applying the block  $LDL^T$  decomposition up to the point where the Schur complement is sufficiently dense. The computation of this Schur complement via dense matrix-matrix multiplication (BLAS-3 `dgemm`) dominates the computational cost. As  $p$  is increased, the utilization of arithmetic units increases in proportion to the dimension of the Schur complement, which explains the  $\mathcal{O}(p^{d-1})$  scaling of the achieved arithmetic performance. As the arithmetic capabilities become saturated for  $p > 15$ , the  $\mathcal{O}(p^{3(d-1)})$  scaling in the factorization runtime should become apparent.

Most of the time in the relaxation step is spent in accessing the factor matrix from memory, given the  $\mathcal{O}(p^{2(d-1)})$  suboptimal storage per patch. The relaxation is therefore limited by memory bandwidth and not arithmetically intense, which explains the poor arithmetic performance. This is in contrast to the sum-factorized residual

evaluation, which has a  $\mathcal{O}(p^d)$  memory footprint and presents better arithmetic utilization [29]. Nevertheless, the results indicate that the runtime for the solution of the patch problems with the sparse Cholesky factorization remains very close to that of the matrix-free residual evaluation for moderate  $p$ , being slightly faster for  $p \leq 7$ , mainly due to lower operation count.

### 3. Application to linear elasticity problems.

**3.1. Primal formulation of linear elasticity.** We now consider how these ideas may be applied in the more complex setting of a nonseparable, vector-valued PDE. The equations of linear elasticity describe the displacement  $\mathbf{u} : \Omega \rightarrow \mathbb{R}^d$  of a solid body with a reference configuration  $\Omega \subset \mathbb{R}^d$ . The primal formulation is to find  $\mathbf{u} - \mathbf{u}_0 \in V := [H_0^1(\Omega)]^d$  such that

$$(3.1) \quad a(\mathbf{v}, \mathbf{u}) = L(\mathbf{v}) \quad \forall \mathbf{v} \in V,$$

where

$$(3.2) \quad a(\mathbf{v}, \mathbf{u}) = \int_{\Omega} 2\mu \varepsilon(\mathbf{v}) : \varepsilon(\mathbf{u}) + \lambda \nabla \cdot \mathbf{v} \nabla \cdot \mathbf{u} \, d\mathbf{x}, \quad L(\mathbf{v}) = \int_{\Omega} \mathbf{v} \cdot \mathbf{B} \, d\mathbf{x}.$$

Here, we assume that the material is homogeneous and isotropic and can thus be described by the Lamé parameters  $\mu, \lambda > 0$ ;  $\varepsilon(\mathbf{u}) = (\nabla \mathbf{u} + \nabla \mathbf{u}^T)/2$  is the linearized strain tensor;  $\mathbf{u}_0$  is Dirichlet data prescribed on  $\Gamma_D \subseteq \partial\Omega$ ; and  $\mathbf{B} \in [L^2(\Omega)]^d$  is a body force. The Poisson ratio  $\nu = \lambda/(2\mu + \lambda)$  measures the compressibility of the material. In the incompressible limit  $\lambda \rightarrow \infty$  (i.e.,  $\nu \rightarrow 1/2$ ), the problem becomes ill-conditioned, as  $a(\cdot, \cdot)$  becomes insensitive to divergence-free perturbations in the arguments.

Consider the partitioning of the stiffness matrix  $A$  into blocks that act on each displacement component,

$$(3.3) \quad A = \begin{bmatrix} A_{11} & \cdots & A_{1d} \\ \vdots & \ddots & \vdots \\ A_{d1} & \cdots & A_{dd} \end{bmatrix}.$$

The diagonal block  $A_{jj}$  discretizes the bilinear form

$$(3.4) \quad \int_{\Omega} \mu \nabla v_j \cdot \nabla u_j + (\mu + \lambda) \frac{\partial v_j}{\partial x_j} \frac{\partial u_j}{\partial x_j} \, d\mathbf{x},$$

where summation is not implied, and  $u_j$  and  $v_j$  are components of  $\mathbf{u}$  and  $\mathbf{v}$ , respectively. The off-diagonal blocks  $A_{ij}$ ,  $i \neq j$ , discretize

$$(3.5) \quad \int_{\Omega} \mu \frac{\partial v_i}{\partial x_j} \frac{\partial u_j}{\partial x_i} + \lambda \frac{\partial v_i}{\partial x_i} \frac{\partial u_j}{\partial x_j} \, d\mathbf{x}.$$

The diagonal blocks can be diagonalized by the FDM on the interior of a Cartesian cell when the reference axes are aligned with the physical coordinates. The same statement does not hold true for the off-diagonal blocks, as they couple together different displacement components. This is because they discretize products of different first derivatives on the different components and hence are not separable.

The separate displacement components (SDC) preconditioner [11, 22] is defined as the block diagonal matrix  $A_{\text{SDC}} = \text{diag}(A_{11}, \dots, A_{dd})$ . In other words, this approach is also described as block-Jacobi in the displacement components. The SDC

TABLE 3

PCG iteration counts for the primal formulation of the linear elasticity problem using the SDC/FDM/sparse relaxation.

$d$	$p \setminus \lambda$	0	$10^0$	$10^1$	$10^2$	$10^3$
2	3	13	14	24	70	199
	7	17	17	28	76	236
	15	18	19	30	81	249
	31	20	20	32	84	258
3	3	20	22	39	114	362
	7	25	28	48	123	381
	15	27	29	51	125	373

preconditioner discretized with the FDM basis is sparse for Cartesian cells aligned with the coordinate axes. On arbitrary cells, for each separate component, we obtain an auxiliary form that is separable in the reference coordinates by selecting constant diagonal coefficients  $\hat{\mu}_K$ .

It is shown in [11] that for a homogeneous isotropic material with principal axes parallel to the axes of the reference coordinate system, the condition number of the preconditioned matrix will depend on the Poisson ratio:

$$(3.6) \quad \kappa(A_{\text{SDC}}^{-1}A) \leq \frac{d-1}{\gamma} \frac{1-\nu}{1-2\nu},$$

where  $\gamma$  is the constant appearing in Korn's inequality,

$$(3.7) \quad \|\mathbf{u}\|_{H^1(\Omega)^d}^2 \leq \gamma \int_{\Omega} \mathbf{u} \cdot \mathbf{u} + \varepsilon(\mathbf{u}) : \varepsilon(\mathbf{u}) \, d\mathbf{x} \quad \forall \mathbf{u} \in V.$$

Thus, the convergence rate of the SDC preconditioner will deteriorate for  $\nu$  approaching  $1/2$ , the so-called nearly incompressible case.

We consider the reference configuration  $\Omega = (0, 1)^d$  discretized on a Cartesian mesh with 8 cells along each direction. We specify  $\mu = 1$ , a uniform downwards body force  $\mathbf{B} = -0.02\mathbf{e}_2$ , and homogeneous Dirichlet BCs on  $\Gamma_D = \{\mathbf{x} \in \partial\Omega, x_1 = 0\}$ . In Table 3 we present the PCG iteration counts required to reduce the Euclidean norm of the residual by a factor of  $10^8$  starting from a zero initial guess. As the preconditioner, we employ the hybrid  $p$ -multigrid/Schwarz method with vertex-star patches and the SDC/FDM/sparse relaxation and a coarse space with  $p = 1$ . As expected from (3.6), the results confirm that the approach is reasonably  $p$ -robust, but that robustness with respect to  $\nu$  cannot be achieved with SDC relaxation on vertex-star patches.

**3.2. Mixed FEM formulations of linear elasticity.** In order to avoid locking in nearly incompressible continua, or impose the incompressibility constraint, the standard approach is to introduce a pressure-like variable and discretize with a mixed FEM. This is expressed by the weak formulation: find  $(\mathbf{u} - \mathbf{u}_0, p) \in V \times Q$  such that

$$(3.8) \quad a(\mathbf{v}, \mathbf{u}) + b(p, \mathbf{v}) = L(\mathbf{v}) \quad \forall \mathbf{v} \in V,$$

$$(3.9) \quad b(q, \mathbf{u}) - c(q, p) = 0 \quad \forall q \in Q,$$

where

$$(3.10) \quad a(\mathbf{v}, \mathbf{u}) = \int_{\Omega} 2\mu \varepsilon(\mathbf{v}) : \varepsilon(\mathbf{u}) \, d\mathbf{x}, \quad b(q, \mathbf{u}) = \int_{\Omega} q \operatorname{div}(\mathbf{u}) \, d\mathbf{x}, \quad c(q, p) = \int_{\Omega} \lambda^{-1} qp \, d\mathbf{x},$$

and  $Q = L_0^2(\Omega)$  for  $\lambda = \infty$  and  $\Gamma_D = \partial\Omega$ , or  $Q = L^2(\Omega)$  otherwise.

In order for this problem to have a unique solution, we require the well-known Brezzi conditions: the solution for  $\mathbf{u}$  is unique if  $a(\cdot, \cdot)$  is coercive on the kernel of  $b(\cdot, \cdot)$ , and the solution for  $p$  is unique if there exists a right inverse for  $b(\cdot, \cdot)$ . This is expressed in the so-called inf-sup condition or LBB condition [7, 12]: there exists  $\beta$ , which might depend on  $\Omega$ , such that

$$(3.11) \quad 0 < \beta := \inf_{q \in Q} \sup_{\mathbf{v} \in V} \frac{b(q, \mathbf{v})}{a(\mathbf{v}, \mathbf{v})^{1/2} \|q\|_Q}.$$

After selecting suitable finite dimensional subspaces  $V_h \subset V$ ,  $Q_h \subset Q$ , we obtain a system of linear equations with the saddle point structure

$$(3.12) \quad \begin{bmatrix} A & B^\top \\ B & -C \end{bmatrix} \begin{bmatrix} \underline{\mathbf{u}} \\ p \end{bmatrix} = \begin{bmatrix} \underline{\mathbf{f}} \\ g \end{bmatrix}.$$

We require the analogous Brezzi conditions for the discrete problem: that  $a(\cdot, \cdot)$  is coercive on the discrete kernel of  $b(\cdot, \cdot)$ , and that there exists a discrete inf-sup constant  $\tilde{\beta}$  independent of the mesh but possibly depending on  $p$  such that

$$(3.13) \quad 0 < \tilde{\beta} := \inf_{q_h \in Q_h} \sup_{\mathbf{v}_h \in V_h} \frac{b(q_h, \mathbf{v}_h)}{a(\mathbf{v}_h, \mathbf{v}_h)^{1/2} \|q_h\|_{Q_h}}.$$

The discretization  $V_h \times Q_h$  must be chosen carefully to satisfy these conditions; the discrete inf-sup condition will not be satisfied by arbitrary discretizations. If they are, we have the well-known error estimates

$$(3.14a) \quad \|\mathbf{u}_h - \mathbf{u}\|_V \leq C_1 \left\{ \inf_{\mathbf{v}_h \in V_h} \|\mathbf{v}_h - \mathbf{u}\|_V + \inf_{q_h \in Q_h} \|q_h - p\|_Q \right\},$$

$$(3.14b) \quad \|p_h - p\|_Q \leq \tilde{\beta}^{-1} C_2 \left\{ \inf_{\mathbf{v}_h \in V_h} \|\mathbf{v}_h - \mathbf{u}\|_V + \inf_{q_h \in Q_h} \|q_h - p\|_Q \right\},$$

where  $C_1, C_2 > 0$  are generic constants independent of the mesh parameter  $h$ . For the use of high-order discretizations, it is desirable to choose an element pair where  $\tilde{\beta}$  does not decrease as the polynomial degree of the approximation is increased. Such a discretization is referred to as  $p$ -stable.

In fact,  $p$ -stability is important for solvers also. Approaches based on block-Gaussian elimination, such as the Uzawa algorithm [5] and block-preconditioned MINRES [41], require preconditioners for the negative pressure Schur complement  $S = C + BA^{-1}B^\top$ . It is well known that for the Stokes system, the continuous analogue of  $S$ ,  $\nabla \cdot (-\nabla^2)^{-1} \nabla$ , is well approximated by the identity operator [51]. It follows that  $S$  is spectrally equivalent to the pressure mass matrix,  $M_p$ ,

$$(3.15) \quad \beta_0^2 \leq \frac{\underline{q}^\top S \underline{q}}{\underline{q}^\top M_p \underline{q}} \leq \beta_1^2 \quad \forall \underline{q} \in \mathbb{R}^{\dim(Q_h)} \setminus \{0\}.$$



The rate of convergence of block-preconditioned MINRES will be determined by the ratio  $\beta_1/\beta_0$ . For Stokes flows with pure Dirichlet BCs,  $\beta_1 = 1$ , and  $\beta_1 = \sqrt{d}$  otherwise. In general, we have  $\beta_0 = \tilde{\beta}$ . Since  $A$  is spectrally equivalent to the vector Laplacian, these results also hold for linear elasticity. We may expect solvers based on such techniques to degrade with  $p$ -refinement if the discretization is not  $p$ -stable.

If we choose to work with the  $[H^1(\Omega)]^d$ -conforming space  $V_h = [Q_p(\Omega)]^d$ , some standard inf-sup stable choices for  $Q_h$  are  $Q_{p-1}(\Omega)$ ,  $DQ_{p-2}(\Omega)$  and  $DP_{p-1}(\Omega)$ .  $DQ_{p-2}$  denotes discontinuous piecewise polynomials of degree at most  $p-2$  in each direction, while  $DP_{p-1}$  denotes discontinuous piecewise polynomials of total degree at most  $p-1$ . The choice  $Q_h = Q_{p-1}$  gives rise to the high-order generalization of the Taylor–Hood mixed element [54]. Here,  $M_p$  will not be block diagonal, and hence more expensive preconditioning techniques will be required. Moreover, it is shown numerically in [1] that the Taylor–Hood element is not  $p$ -stable. The choice  $Q_h = DQ_{p-2}$  exhibits an asymptotic decay of  $\tilde{\beta} \leq Cp^{(1-d)/2}$  as  $p \rightarrow \infty$  [10] and thus is not  $p$ -stable. In practice, it is observed that this is quite a pessimistic bound for moderate  $p$  [35]. The choice  $Q_h = DP_{p-1}$  is  $p$ -stable, but numerical experiments reveal that the stability is severely affected by the cell aspect ratio, unlike the previous two choices [50]. Moreover this last space does not have tensor product shape functions, so its efficient implementation becomes challenging.

To construct  $p$ -stable discretizations that are also robust to cell aspect ratio, we turn to nonconforming schemes with  $V_h \subset H(\text{div}, \Omega)$  [15, 31]. In particular we consider the use of Raviart–Thomas elements [4] of degree  $p$  for  $V_h$  for the displacement, paired with  $Q_h = DQ_{p-1}$ . This pair satisfies  $\text{div}(V_h) = Q_h$ , which enforces the incompressibility constraint (3.9) exactly in the numerical approximation for  $\lambda = \infty$ . The Raviart–Thomas elements are defined on the reference quadrilateral as

$$(3.16) \quad \text{RT}_p(\hat{K}) = P_p(\hat{\mathcal{I}}) \otimes DP_{p-1}(\hat{\mathcal{I}}) \oplus DP_{p-1}(\hat{\mathcal{I}}) \otimes P_p(\hat{\mathcal{I}}).$$

The analogous element in 3D is referred to as the Nédélec face element [39]. The definition can be extended to curvilinear cells via the contravariant Piola transform: for a function  $\hat{\mathbf{u}} : \hat{K} \rightarrow \mathbb{R}^d$ , we define  $\mathbf{u} : K \rightarrow \mathbb{R}^d$  as

$$(3.17) \quad \mathbf{u} = \mathcal{P}_K(\hat{\mathbf{u}}) := \frac{1}{|DF_K|} DF_K (\hat{\mathbf{u}} \circ F_K^{-1})$$

and set

$$(3.18) \quad \text{RT}_p(K) = \mathcal{P}_K \left( \text{RT}_p(\hat{K}) \right).$$

These elements have superb properties, but their nonconforming nature must be suitably addressed in the discretization. They only impose continuity of the normal components of  $\mathbf{u}$  across cell facets, and we therefore weakly enforce the tangential continuity via the symmetric interior penalty Galerkin (SIPG) method [3]. The use of SIPG for the displacement requires further extension of the FDM/sparse relaxation; in particular, we must consider the additional facet integrals arising in the method and show that the stiffness remains sparse.

**3.3. Extension to interior penalty DG methods.** Interior penalty discontinuous Galerkin (IP-DG) methods relax the continuity requirement of the discretization space. For instance, instead of  $[H^1(\Omega)]^d$ , we consider a larger function space with weaker continuity, such as  $[L^2(\Omega)]^d$  or  $H(\text{div}, \Omega)$ . As previously mentioned, in order to deal with the nonconformity,  $C^0$ -continuity is weakly enforced via the introduction

of a penalty term on the set of interior facets  $\Gamma_I$  of the mesh  $\mathcal{T}_h$  that vanishes for  $C^0$ -continuous functions. Similarly, the weak prescription of the Dirichlet BC  $\mathbf{u} = \mathbf{u}_0$  on  $\Gamma_D$  is achieved by introducing a penalty term on  $\Gamma_D$ .

We consider the following SIPG formulation:

$$(3.19) \quad \begin{aligned} a(\mathbf{v}, \mathbf{u}) = & \sum_{K \in \mathcal{T}_h} \int_K \nabla \mathbf{v} : \mathcal{F}^v(\nabla \mathbf{u}) \, d\mathbf{x} \\ & + \sum_{e \in \Gamma_I \cup \Gamma_D} \int_e \eta h_e^{-1} \{G^\top\} [\![\mathbf{v}]\!] : [\![\mathbf{u}]\!] - [\![\mathbf{v}]\!] : \{\mathcal{F}^v(\nabla \mathbf{u})\} - \{G^\top \nabla \mathbf{v}\} : [\![\mathbf{u}]\!] \, ds, \end{aligned}$$

$$(3.20) \quad L(\mathbf{v}) = \int_\Omega \mathbf{v} \cdot \mathbf{B} \, d\mathbf{x} + \int_{\Gamma_D} \eta h_e^{-1} G^\top (\mathbf{v} \otimes \mathbf{n}) : (\mathbf{u}_0 \otimes \mathbf{n}) - G^\top \nabla \mathbf{v} : (\mathbf{u}_0 \otimes \mathbf{n}) \, ds.$$

Here  $\mathcal{F}^v(\nabla \mathbf{u})$  is a linear viscous flux. For the vector Poisson equation, the viscous flux is given by  $\mathcal{F}^v(\nabla \mathbf{u}) = \nabla \mathbf{u}$ . For the primal formulation of linear elasticity, the viscous flux corresponds to the stress tensor  $\mathcal{F}^v(\nabla \mathbf{u}) = \mu(\nabla \mathbf{u} + \nabla \mathbf{u}^\top) + \lambda \nabla \cdot \mathbf{u} \mathbf{1}$ . For the mixed formulation of linear elasticity, the (1,1)-block of the system has viscous flux  $\mathcal{F}^v(\nabla \mathbf{u}) = \mu(\nabla \mathbf{u} + \nabla \mathbf{u}^\top)$ .

From left to right, the terms in the surface integral in (3.19) are referred to as the penalty, consistency, and adjoint consistency terms. The quantity  $G$  is known as the homogeneity tensor,

$$(3.21) \quad G_{ijkl} = \frac{\partial}{\partial u_{k,l}} [\mathcal{F}^v(\nabla \mathbf{u})]_{ij},$$

for which we define the adjoint product with  $\nabla \mathbf{v}$ :

$$(3.22) \quad [G^\top \nabla \mathbf{v}]_{kl} = G_{ijkl} v_{i,j}.$$

The average  $\{\cdot\}$  and jump  $[\![\cdot]\!]$  operators are defined for scalar, vector, and tensor arguments as follows. Let  $e$  be a facet of the mesh. For an interior facet, let  $K^-$  and  $K^+$  be the two mesh cells that share it, and let  $\mathbf{w}^-$  and  $\mathbf{w}^+$  be the traces of a function  $\mathbf{w}$  on  $e$  from  $K^-$  and  $K^+$ , respectively. On each facet we define

$$(3.23) \quad \{\mathbf{w}\}_e = \begin{cases} \frac{1}{2}(\mathbf{w}^- + \mathbf{w}^+), & e \in \Gamma_I, \\ \mathbf{w}, & \text{otherwise,} \end{cases} \quad [\![\mathbf{w}]\!]_e = \begin{cases} \mathbf{w}^- \otimes \mathbf{n}^- + \mathbf{w}^+ \otimes \mathbf{n}^+, & e \in \Gamma_I, \\ \mathbf{w} \otimes \mathbf{n}, & \text{otherwise.} \end{cases}$$

In order to ensure coercivity of  $a(\cdot, \cdot)$  as we do  $h$  or  $p$  refinement, the penalty term must be sufficiently large. The penalty coefficient  $\eta h_e^{-1}$  must be chosen inversely proportional to the normal spacing of GLL nodes near the facet, i.e.,  $\eta = \mathcal{O}(p(p+1))$  [37]. For the reciprocal length scale in the direction normal to facet  $e$  we use

$$(3.24) \quad h_e^{-1} := |e| \left\{ |K|^{-1} \right\}_e.$$

The stiffness matrix that corresponds to  $a(\cdot, \cdot)$  in the SIPG formulation is obtained via direct stiffness summation over the cells and facets:

$$(3.25) \quad A = \sum_{K \in \mathcal{T}_h} R_K^\top A^K R_K + \sum_{e \in \Gamma_I \cup \Gamma_D} R_e^\top A^e R_e,$$

where  $A^K$  is the cell matrix discretizing the volume integral in  $K$ ,  $R_K$  is the Boolean restriction onto the DOFs of  $K$ ,  $A^e$  is the facet matrix discretizing the surface integral on  $e$ , and  $R_e$  is Boolean restriction onto the DOFs of the cells sharing facet  $e$ .

To illustrate the extension of our approach to the SIPG discretization, we consider again the scalar Poisson equation. The discrete problem is to find  $u_h \in V_h = \text{DQ}_p(\Omega) \subset L^2(\Omega)$ . On Cartesian cells, both  $A^K$  and  $A^e$  have a tensor product structure of the form (2.12), with matrices of operators on the interval that can be sparsified by the FDM. To illustrate this, suppose that, for  $e \in \Gamma_I$ ,  $R_e$  reorders the DOFs such that the cells  $K^-$  and  $K^+$  share  $e$  along the  $d$ th reference coordinate axis, while leaving the other axes consistently oriented on both cells. The facet matrices are

$$(3.26) \quad A^e = \begin{cases} E^e \otimes \hat{B} & \text{if } d = 2, \\ E^e \otimes \hat{B} \otimes \hat{B} & \text{if } d = 3, \end{cases}$$

where the interval facet matrix  $E^e$  is defined in terms of the coefficients  $\mu_j^K$  appearing in (2.12), the one-dimensional shape functions  $\{\hat{\phi}_j\}$ , and their normal derivatives  $\frac{\partial}{\partial n} \hat{\phi}_j$  on  $\partial \hat{\mathcal{I}}$  (the usual derivative with a sign). When  $e \in \Gamma_D$ ,  $E^e \in \mathbb{R}^{(p+1) \times (p+1)}$  is given by

$$(3.27) \quad [E^e]_{ij} = \mu^e \left( \eta \hat{\phi}_i(\hat{x}^e) \hat{\phi}_j(\hat{x}^e) - \hat{\phi}_i(\hat{x}^e) \frac{\partial}{\partial n} \hat{\phi}_j(\hat{x}^e) - \frac{\partial}{\partial n} \hat{\phi}_i(\hat{x}^e) \hat{\phi}_j(\hat{x}^e) \right).$$

Here  $\mu^e = \mu_l^K$ , where  $\hat{x}_l$  is the reference coordinate normal to  $e$ , and  $\hat{x}^e \in \partial \hat{\mathcal{I}}$  describes the facet  $e$  as the image of the plane  $\hat{x}_l = \hat{x}^e$  under  $F_K$ . When  $e \in \Gamma_I$ ,  $E^e$  is a  $2 \times 2$  block matrix with blocks  $E_{rs}^e \in \mathbb{R}^{(p+1) \times (p+1)}$ ,  $r, s \in \{0, 1\}$ , given by

$$(3.28) \quad [E_{rs}^e]_{ij} = \frac{(-1)^{r-s}}{2} \left( \eta(\mu_0^e + \mu_1^e) \hat{\phi}_i(\hat{x}_r^e) \hat{\phi}_j(\hat{x}_s^e) - \mu_s^e \hat{\phi}_i(\hat{x}_r^e) \frac{\partial}{\partial n} \hat{\phi}_j(\hat{x}_s^e) - \mu_r^e \frac{\partial}{\partial n} \hat{\phi}_i(\hat{x}_r^e) \hat{\phi}_j(\hat{x}_s^e) \right).$$

Here  $\mu_0^e = \mu_l^{K^-}$ ,  $\mu_1^e = \mu_m^{K^+}$ , where  $\hat{x}_l$  and  $\hat{x}_m$  are the reference directions normal to  $e$  on  $K^-$  and  $K^+$ , respectively. Similarly, the facet  $e$  is the image of  $\hat{x}_l = \hat{x}_0^e$  under  $F_{K^-}$  and that of  $\hat{x}_m = \hat{x}_1^e$  under  $F_{K^+}$ .

Some implementations of  $\text{DQ}_p$  do not feature an interior-interface decomposition and use the Gauss–Legendre (GL) nodal shape functions. The GL nodes do not include the endpoints, thus all shape functions have nonzero support at the facets, causing  $E^e$  to be dense. The matrices  $E^e$  are sparse for a basis with an interior-interface decomposition, such as the GLL Lagrange polynomials  $\{\ell_j\}$ , the hierarchical Lobatto polynomials  $\{l_j\}$ , and the FDM polynomials  $\{\hat{s}_j\}$ . Since  $\hat{\phi}_j(\pm 1)$  is nonzero for a single  $j \in \{0, p\}$ , each term in (3.28) and (3.27) corresponds to a nonzero entry, a nonzero row, and a nonzero column of  $E^e$ , respectively, as seen in Figure 10.




FIG. 10. Nonzero structure for the interior facet matrix  $E^e$  on the interval ( $p = 4$ ).

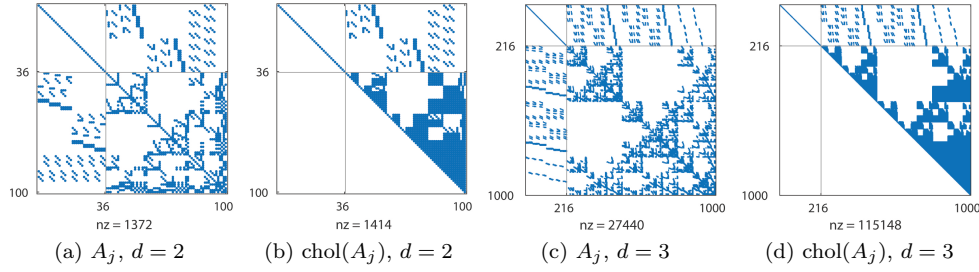


FIG. 11. Nonzero structure of the SIPG stiffness matrix in the FDM basis  $A_j = L_j L_j^\top$  and its upper Cholesky factor  $L_j^\top$  for the Poisson problem on a Cartesian vertex-star patch with  $p = 4$ . Since the space is discontinuous, the number of DOFs in a patch is increased. The number of nonzeros in an interior row is  $3d + 1$ , since the interior DOFs are connected to each of the  $2d$  facets of their corresponding cell, plus  $d$  more facets from the adjacent cells.

Instead of diagonalizing the SIPG patch matrix as in [55], our FDM-based approach produces a sparse matrix with diagonal interior blocks on (possibly) unstructured vertex-star patches. Figure 11 shows the sparsity pattern of the matrix for the SIPG formulation of the Poisson equation on a Cartesian vertex-star patch in the FDM basis, along with its Cholesky factor. Here the matrix size is increased from  $(2p - 1)^d$  DOFs in the CG case to  $(2p + 2)^d$ . At low polynomial degrees, the interface DOFs form a large fraction of the total number, but the proportion decreases as  $p$  increases. The computational complexity analysis of section 2.6 carries over to the SIPG case.

For a general viscous flux, we construct an auxiliary separable form by expressing the cell integrals in terms of the reference coordinates

$$(3.29) \quad a_K(\mathbf{v}, \mathbf{u}) = \int_K G_{ijkl}^K \frac{\partial v_i}{\partial x_j} \frac{\partial u_k}{\partial x_l} d\mathbf{x} = \int_{\hat{K}} \hat{G}_{ijkl}^K \frac{\partial}{\partial \hat{x}_j} (v_i \circ F_K) \frac{\partial}{\partial \hat{x}_l} (u_k \circ F_K) d\hat{\mathbf{x}},$$

where  $\hat{G}^K$  is the homogeneity tensor in the reference coordinates

$$(3.30) \quad \hat{G}_{ijkl}^K = |\det(DF_K)| [DF_K^{-1}]_{jm} [DF_K^{-1}]_{ln} G_{imkn}^K.$$

The auxiliary form  $\tilde{a}(\cdot, \cdot)$  is constructed by approximating  $\hat{G}_{ijkl}^K$  with a piecewise constant tensor that discards the entries where  $i \neq k$  or  $j \neq l$ . The corresponding cell stiffness matrices become sparse in the FDM basis and have a similar form as (2.12), except that the coefficients  $\mu_j^K$  are diagonal matrices that multiply each term through an additional Kronecker product. Hence, we expect that preconditioners based on the auxiliary form to be limited by the coupling between vector components, by the mesh geometry, and by how  $G^K$  varies within  $K$ .

This approach carries over to non-Cartesian cells for DG discretizations of the Poisson equation in the same way as the CG case. Unfortunately, the extension of our approach to vector-valued problems in  $H(\text{div})$  on non-Cartesian cells does not yield a good preconditioner. In this setting, we construct a block diagonal preconditioner separating the components of the DOFs, which are in the reference coordinates. For the vector Poisson problem on non-Cartesian cells, the Piola transform introduces off-diagonal contributions from the volume and surface terms, which does not occur on non-Piola-mapped elements, such as  $[Q_p]^d$  and  $[DQ_p]^d$ . The excluded terms are required for the surface integral terms to vanish for arguments with  $C^0$  continuity, and without them the preconditioner might become indefinite on non-Cartesian cells.

**3.4. Results for mixed formulations of linear elasticity.** We consider the same problem as in Table 3, with both a conforming  $[Q_p]^d \times DQ_{p-2}$  discretization and a nonconforming  $RT_p \times DQ_{p-1}$  discretization. For the  $H(\text{div})$ -conforming discretization, the normal components of the Dirichlet BCs are enforced strongly, while the tangential components of the BCs are weakly enforced with SIPG. Enforcing the normal conditions strongly is crucial for achieving a divergence-free solution in the Stokes limit  $\lambda = \infty$ . For the penalty coefficient, we use  $\eta = p(p+1)$ . We restrict our experiment to Cartesian cells, so that the FDM/sparse relaxation is applicable to the  $H(\text{div})$ -conforming discretization.

We iteratively solve the discrete system (3.12) via MINRES with a symmetric positive definite block diagonal preconditioner

$$(3.31) \quad \mathcal{P}_{\text{diag}} = \begin{bmatrix} P_1 & 0 \\ 0 & P_2 \end{bmatrix}.$$

Here  $P_1$  is a preconditioner for the displacement block  $A$ , and  $P_2$  is a preconditioner for the scaled pressure mass matrix  $(\mu^{-1} + \lambda^{-1})M_p$ . For  $P_1$  we employ the hybrid  $p$ -multigrid/Schwarz method with the SDC/FDM/sparse relaxation and  $[Q_1]^d$  as the coarse space. In our tests, we discretize the pressure space with the GL basis, and employ point-Jacobi on the pressure mass matrix, i.e.,  $P_2 = (\mu^{-1} + \lambda^{-1})\text{diag}(M_p)$ . When  $\mathcal{T}_h$  consists of Cartesian cells,  $M_p = \text{diag}(M_p)$  in the GL basis. The solver is illustrated in Figure 12.

In Table 4 we present MINRES iteration counts for the same configuration considered in Table 3 in section 3.1, using the  $[Q_p]^d \times DQ_{p-2}$  and  $RT_p \times DQ_{p-1}$  elements, respectively. Both discretizations yield robust iteration counts with respect to  $\lambda$ ; the iterations grow with the former discretization much more quickly than the latter, especially in 3D.

The solver configuration shown in Figure 12 is optimized for memory usage, employing a block diagonal preconditioner so that the short-term recurrences of MINRES may be exploited. If one is willing to trade memory for time, one may consider an alternative configuration shown in Figure 13 employing right-preconditioned GMRES [49] with a block upper triangular preconditioner,

$$(3.32) \quad \mathcal{P}_{\text{upper}} = \begin{bmatrix} P_1 & B^\top \\ 0 & -P_2 \end{bmatrix},$$

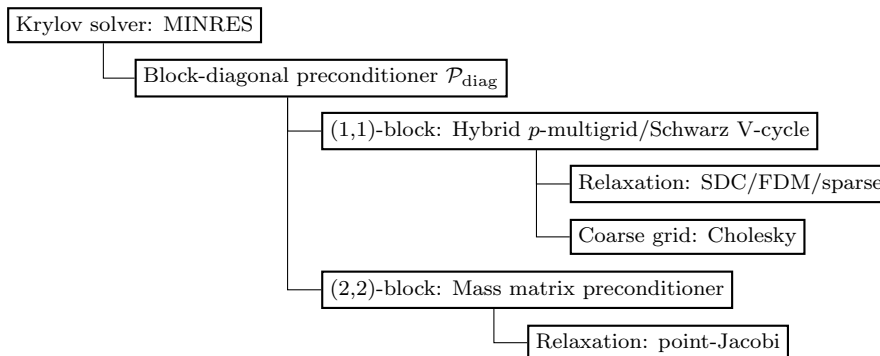


FIG. 12. Solver diagram for the mixed linear elasticity problem.

TABLE 4

MINRES iteration counts for the mixed linear elasticity problem, using the solver in Figure 12.

$d$	$p \setminus \lambda$	$[Q_p]^d \times DQ_{p-2}$					$RT_p \times DQ_{p-1}$				
		$10^0$	$10^1$	$10^2$	$10^3$	$\infty$	$10^0$	$10^1$	$10^2$	$10^3$	$\infty$
2	3	28	40	43	43	43	25	36	39	40	40
	7	31	45	50	51	51	28	40	43	45	45
	15	34	50	57	57	57	30	43	48	48	48
	31	36	53	64	65	65	31	45	51	51	51
3	3	44	67	75	76	76	34	50	55	56	56
	7	50	83	96	97	98	39	58	63	65	65
	15	53	88	111	118	119	41	63	70	70	70

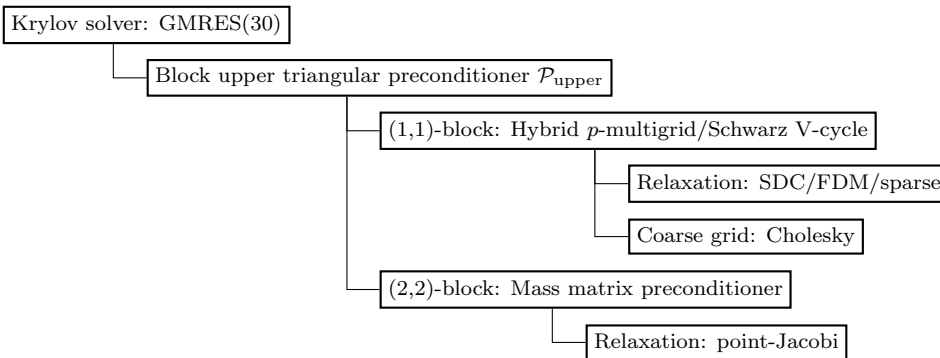


FIG. 13. Solver diagram for the mixed linear elasticity problem that trades memory for iteration counts.

TABLE 5

GMRES iteration counts for the mixed linear elasticity problem, using the solver in Figure 13.

$d$	$p \setminus \lambda$	$[Q_p]^d \times DQ_{p-2}$					$RT_p \times DQ_{p-1}$				
		$10^0$	$10^1$	$10^2$	$10^3$	$\infty$	$10^0$	$10^1$	$10^2$	$10^3$	$\infty$
2	3	17	23	25	26	26	14	20	22	22	22
	7	18	25	27	28	28	16	22	24	24	24
	15	20	27	33	34	34	17	23	26	26	26
	31	21	30	38	38	39	18	24	28	28	28
3	3	24	32	33	34	34	17	23	25	25	25
	7	27	35	38	38	38	21	26	29	29	29
	15	28	38	44	46	46	22	28	31	31	32


which requires a single application each of  $P_1^{-1}$ ,  $P_2^{-1}$ , and  $B^\top$  per GMRES iteration. The GMRES iteration counts are presented in Table 5.

In Table 6 we study the performance of our solver on an unstructured mesh. We consider the  $[Q_p]^d \times DQ_{p-2}$  discretization of incompressible linear elasticity ( $\lambda = \infty$ ). We prescribe  $\mu = 1$ , a uniform downwards body force  $\mathbf{B} = -0.02\mathbf{e}_2$ , and homogeneous Dirichlet BCs on the displacement on the holes of the domain. The three-dimensional mesh is obtained via extrusion by 16 layers of the two-dimensional mesh. The iteration counts follow the same pattern as before for this element: they are not  $p$ -robust as expected, but they remain modest even at very high degrees.

TABLE 6

GMRES iteration counts for the mixed formulation of the incompressible linear elasticity problem on the unstructured mesh shown here, using the solver in Figure 13.

$d$	$p$	# DOFs	Iter.
2	3	17 466	26
	7	104 250	29
	15	499 002	35
	31	2 173 242	44
3	3	588 927	39
	7	7 876 575	47
	11	31 236 927	51



**4. Conclusion.** We have introduced a fast relaxation method required for unstructured vertex-centered patch problems arising in Pavarino's approach, extending its practicality to much higher polynomial degrees. Our method relies on a spectrally equivalent form, constructed such that it is separable in the reference coordinates. We show promising results for the Poisson equation and mixed formulations of linear elasticity. A downside of the approach is its narrow applicability; it will not be effective on more general problems, especially for those where the dominant terms include mixed derivatives and mixed vector components. In addition, our method relies on having a good quality mesh, with its performance depending on the minimal angle; however, mesh generators with guarantees on the minimal angle are available in two dimensions [32]. So far, we have only considered constant-coefficient problems, but the theory of [6] suggests that our approach would remain effective for spatially varying coefficients. Work in progress shows that the memory bandwidth limitations and the suboptimal complexity of the sparse Cholesky factorization can be overcome by iterative patch solvers, such as incomplete factorizations and algebraic multigrid.

## REFERENCES

- [1] M. AINSWORTH, P. COGGINS, AND B. SENIOR, *Mixed hp-Finite Element Methods for Incompressible Flow*, CRC Research Notes in Mathematics, Chapman and Hall, 2000.
- [2] P. R. AMESTOY, I. S. DUFF, J.-Y. L'EXCELLENT, AND J. KOSTER, *A fully asynchronous multi-frontal solver using distributed dynamic scheduling*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 15–41.
- [3] D. N. ARNOLD, *An interior penalty finite element method with discontinuous elements*, SIAM J. Numer. Anal., 19 (1982), pp. 742–760.
- [4] D. N. ARNOLD, D. BOFFI, AND R. S. FALK, *Quadrilateral  $H(\text{div})$  finite elements*, SIAM J. Numer. Anal., 42 (2005), pp. 2429–2451.
- [5] K. J. ARROW, L. HURWICZ, AND H. UZAWA, *Studies in Non-linear Programming*, Stanford University Press, Redwood City, CA, 1958.
- [6] O. AXELSSON AND J. KARÁTON, *Equivalent operator preconditioning for elliptic problems*, Numer. Algorithms, 50 (2009), pp. 297–380.
- [7] I. BABUŠKA, *The finite element method with Lagrangian multipliers*, Numer. Math., 20 (1973), pp. 179–192.
- [8] S. BALAY, S. ABHYANKAR, M. F. ADAMS, J. BROWN, P. BRUNE, K. BUSCHELMAN, L. DALCIN, V. ELJKHOUT, W. D. GROPP, D. KARPEYEV, D. KAUSHIK, M. G. KNEPLEY, D. A. MAY, L. C. MCINNES, R. T. MILLS, T. MUNSON, K. RUPP, P. SANAN, B. F. SMITH, S. ZAMPINI, H. ZHANG, AND H. ZHANG, *PETSc Users Manual*, Tech. Report ANL-95/11 - Revision 3.11, Argonne National Laboratory, 2019.
- [9] P. D. BELLO-MALDONADO AND P. F. FISCHER, *Scalable low-order finite element preconditioners for high-order spectral element Poisson solvers*, SIAM J. Sci. Comput., 41 (2019), pp. S2–S18.

- [10] C. BERNARDI, Y. MADAY, AND B. MÉTIVET, *Spectral approximation of the periodic-nonperiodic Navier-Stokes equations*, Numer. Math., 51 (1987), pp. 655–700.
- [11] R. BLAHEA, *Displacement decomposition–incomplete factorization preconditioning techniques for linear elasticity problems*, Numer. Linear Algebra Appl., 1 (1994), pp. 107–128.
- [12] F. BREZZI, *On the existence, uniqueness and approximation of saddle-point problems arising from Lagrangian multipliers*, ESAIM Math. Model. Numer., 8 (1974), pp. 129–151.
- [13] C. CANUTO, *Stabilization of spectral methods by finite element bubble functions*, Comput. Methods Appl. Mech. Engrg., 116 (1994), pp. 13–26, [https://doi.org/10.1016/S0045-7825\(94\)80004-9](https://doi.org/10.1016/S0045-7825(94)80004-9).
- [14] Y. CHEN, T. A. DAVIS, W. W. HAGER, AND S. RAJAMANICKAM, *Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate*, ACM Trans. Math. Softw., 35 (2008), 22, <https://doi.org/10.1145/1391989.1391995>.
- [15] B. COCKBURN, G. KANSCHAT, AND D. SCHÖTZAU, *A note on discontinuous Galerkin divergence-free solutions of the Navier–Stokes equations*, J. Sci. Comput., 31 (2007), pp. 61–73.
- [16] W. COUZY AND M. O. DEVILLE, *A fast Schur complement method for the spectral element discretization of the incompressible Navier-Stokes equations*, J. Comp. Phys., 116 (1995), pp. 135–142, <https://doi.org/10.1006/jcph.1995.1011>.
- [17] L. T. DIOSADY AND S. M. MURMAN, *Scalable tensor-product preconditioners for high-order finite-element methods: Scalar equations*, J. Comp. Phys., 394 (2019), pp. 759–776, <https://doi.org/10.1016/j.jcp.2019.04.047>.
- [18] M. DRYJA AND O. WIDLUND, *An Additive Variant of the Schwarz Alternating Method for the Case of Many Subregions*, Technical Report 339, Ultracomputer Note 131, Department of Computer Science, Courant Institute, 1987.
- [19] P. E. FARRELL, M. G. KNEPLEY, L. MITCHELL, AND F. WECHSUNG, *PCPATCH: Software for the Topological Construction of Multigrid Relaxation Methods*, ACM Trans. Math. Softw., 47 (2021), 25, <https://doi.org/10.1145/3445791>.
- [20] P. F. FISCHER, H. M. TUFO, AND N. I. MILLER, *An Overlapping Schwarz Method for Spectral Element Simulation of Three-dimensional Incompressible Flows*, in Parallel Solution of Partial Differential Equations, Springer, New York, 2000, pp. 159–180.
- [21] A. GEORGE, *Nested dissection of a regular finite element mesh*, SIAM J. Numer. Anal., 10 (1973), pp. 345–363.
- [22] I. GUSTAFSSON AND G. LINDSKOG, *On parallel solution of linear elasticity problems: Part I: Theory*, Numer. Linear Algebra Appl., 5 (1998), pp. 123–139.
- [23] M. HOMOLYA, R. C. KIRBY, AND D. A. HAM, *Exposing and Exploiting Structure: Optimal Code Generation for High-order Finite Element Methods*, preprint, arXiv:1711.02473, 2017.
- [24] M. HOMOLYA, L. MITCHELL, F. LUPORINI, AND D. A. HAM, *TSFC: A structure-preserving form compiler*, SIAM J. Sci. Comput., 40 (2018), pp. C401–C428.
- [25] I. HUISMANN, J. STILLER, AND J. FRÖHLICH, *Scaling to the stars—A linearly scaling elliptic solver for  $p$ -multigrid*, J. Comp. Phys., 398 (2019), 108868.
- [26] G. KARNIADAKIS AND S. SHERWIN, *Spectral/hp Element Methods for Computational Fluid Dynamics*, Oxford University Press, Oxford, UK, 2013.
- [27] D. S. KERSHAW, *Differencing of the diffusion equation in Lagrangian hydrodynamic codes*, J. Comp. Phys., 39 (1981), pp. 375–395.
- [28] R. C. KIRBY, *Algorithm 839: FIAT, a new paradigm for computing finite element basis functions*, ACM Trans. Math. Softw., 30 (2004), pp. 502–516, <https://doi.org/10.1145/1039813.1039820>.
- [29] M. KRONBICHLER AND W. A. WALL, *A performance comparison of continuous and discontinuous Galerkin methods with fast multigrid solvers*, SIAM J. Sci. Comput., 40 (2018), pp. A3423–A3448.
- [30] C. LANCZOS, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, J. Res. Natl. Bur. Stand. (U.S.), 45 (1950), 2133.
- [31] P. L. LEDERER AND J. SCHÖBERL, *Polynomial robust stability analysis for  $H(\text{div})$ -conforming finite elements for the Stokes equations*, IMA J. Numer. Anal., 38 (2018), pp. 1832–1860.
- [32] X. LIANG AND Y. ZHANG, *Hexagon-based all-quadrilateral mesh generation with guaranteed angle bounds*, Comput. Methods Appl. Mech. Engrg., 200 (2011), pp. 2005–2020, <https://doi.org/10.1016/j.cma.2011.03.002>.
- [33] J. W. LOTTES AND P. F. FISCHER, *Hybrid multigrid/Schwarz algorithms for the spectral element method*, J. Sci. Comput., 24 (2005), pp. 45–78.
- [34] R. E. LYNCH, J. R. RICE, AND D. H. THOMAS, *Direct solution of partial difference equations by tensor product methods*, Numer. Math., 6 (1964), pp. 185–199.
- [35] Y. MADAY, D. MEIRON, A. T. PATERA, AND E. M. RÖNQVIST, *Analysis of iterative methods for the steady and unsteady Stokes problem: Application to spectral element discretizations*, SIAM J. Sci. Comput., 14 (1993), pp. 310–337.



- [36] J.-F. MAITRE AND O. POURQUIER, *Condition number and diagonal preconditioning: Comparison of the  $p$ -version and the spectral element methods*, Numer. Math., 74 (1996), pp. 69–84.
- [37] J. MANZANERO, A. M. RUEDA-RAMÍREZ, G. RUBIO, AND E. FERRER, *The Bassi Rebay 1 scheme is a special case of the symmetric interior penalty formulation for discontinuous Galerkin discretisations with Gauss–Lobatto points*, J. Comp. Phys., 363 (2018), pp. 1–10.
- [38] J. R. MUNKRES, *Elements of Algebraic Topology*, CRC Press, Boca Raton, FL, 1984.
- [39] J.-C. NÉDÉLEC, *Mixed finite elements in  $\mathbb{R}^3$* , Numer. Math., 35 (1980), pp. 315–341.
- [40] S. A. ORSZAG, *Spectral methods for problems in complex geometries*, J. Comp. Phys., 37 (1980), pp. 70–92, [https://doi.org/10.1016/0021-9991\(80\)90005-4](https://doi.org/10.1016/0021-9991(80)90005-4).
- [41] C. C. PAIGE AND M. A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.
- [42] L. F. PAVARINO, *Additive Schwarz methods for the  $p$ -version finite element method*, Numer. Math., 66 (1993), pp. 493–515.
- [43] L. F. PAVARINO, E. ZAMPIERI, R. PASQUETTI, AND F. RAPETTI, *Overlapping Schwarz methods for Fekete and Gauss–Lobatto spectral elements*, SIAM J. Sci. Comput. 29 (2007), pp. 1073–1092.
- [44] W. PAZNER, *Efficient low-order refined preconditioners for high-order matrix-free continuous and discontinuous Galerkin methods*, SIAM J. Sci. Comput., 42 (2020), pp. A3055–A3083.
- [45] W. PAZNER AND P.-O. PERSSON, *Approximate tensor-product preconditioners for very high order discontinuous Galerkin methods*, J. Comp. Phys., 354 (2018), pp. 344–369.
- [46] W. PAZNER AND P.-O. PERSSON, *Interior penalty tensor-product preconditioners for high-order discontinuous Galerkin discretizations*, in Proceedings of the 2018 AIAA Aerospace Sciences Meeting, 2018, p. 1093.
- [47] F. RATHGEBER, D. A. HAM, L. MITCHELL, M. LANGE, F. LUPORINI, A. T. T. McRAE, G.-T. BERCEA, G. R. MARKALL, AND P. H. J. KELLY, *Firedrake: Automating the finite element method by composing abstractions*, ACM Trans. Math. Softw., 43 (2016), pp. 24:1–24:27.
- [48] J.-F. REMACLE, R. GANDHAM, AND T. WARBURTON, *GPU accelerated spectral finite elements on all-hex meshes*, J. Comp. Phys., 324 (2016), pp. 246–257.
- [49] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869, <https://doi.org/10.1137/0907058>.
- [50] D. SCHÖTZAU AND C. SCHWAB, *Mixed hp-FEM on anisotropic meshes*, Math. Models Methods Appl. Sci., 8 (1998), pp. 787–820.
- [51] D. SILVESTER AND A. WATHEN, *Fast iterative solution of stabilised Stokes systems Part II: Using general block preconditioners*, SIAM J. Numer. Anal., 31 (1994), pp. 1352–1367.
- [52] J. STILLER, *Robust multigrid for Cartesian interior penalty DG formulations of the Poisson equation in 3D*, in Spectral and High Order Methods for Partial Differential Equations ICOSAHOM 2016, M. L. Bittencourt, N. A. Dumont, and J. S. Hesthaven, eds., Springer International Publishing, Cham, 2017, pp. 189–201.
- [53] B. SZABÓ AND I. BABUŠKA, *Finite Element Analysis*, John Wiley & Sons, New York, 1991.
- [54] C. TAYLOR AND P. HOOD, *A numerical solution of the Navier-Stokes equations using the finite element technique*, Comput. Fluids, 1 (1973), pp. 73–100.
- [55] J. WITTE, D. ARNDT, AND G. KANSCHAT, *Fast tensor product Schwarz smoothers for high-order discontinuous Galerkin methods*, Comput. Methods Appl. Math., 21 (2021), pp. 709–728, <https://doi.org/10.1515/cmam-2020-0078>.
- [56] J. XU, *The auxiliary space method and optimal multigrid preconditioning techniques for unstructured grids*, Computing, 56 (1996), pp. 215–235, <https://doi.org/10.1007/bf02238513>.