

ALGEBRAIC MULTIGRID PRECONDITIONING OF HIGH-ORDER SPECTRAL ELEMENTS FOR ELLIPTIC PROBLEMS ON A SIMPLICIAL MESH*

LUKE OLSON†

Abstract. Algebraic multigrid is investigated as a solver for linear systems that arise from high-order spectral element discretizations. An algorithm is introduced that utilizes the efficiency of low-order finite elements to precondition the high-order method in a multilevel setting. In particular, the efficacy of this approach is highlighted on simplexes in two and three dimensions with nodal spectral elements up to order $n = 11$. Additionally, a hybrid preconditioner is also developed for use with discontinuous spectral element methods. The latter approach is verified for the discontinuous Galerkin method on elliptic problems.

Key words. spectral element, discontinuous Galerkin, algebraic multigrid

AMS subject classifications. 65N15, 65N30, 65N55

DOI. 10.1137/060663465

1. Introduction. In recent years spectral element methods have increased in popularity [18, 9]. Their success is often attributed to the ability to approximate solutions to partial differential equations (PDEs) efficiently by requiring fewer degrees-of-freedom than low-order approaches (e.g., bilinear finite elements). Moreover, high-order methods yield more accurate solutions in problems with long time integration [14] and are also beneficial in multiscale simulation [17]. Combined with a compatible discretization method such as the discontinuous Galerkin (DG) method, the method is used successfully in a variety of applications [20, 15].

The overall efficiency of high-order methods, however, remains an open question, particularly when a solution to a system of equations is needed. While fewer degrees-of-freedom are required to achieve a certain accuracy in the solution compared with low-order methods, the degrees-of-freedom are more intimately coupled, and the resulting matrix is more dense. The reduced locality impacts both the amount of storage and the complexity of the solver.

The use of high-order elements also changes many matrix properties, impacting the success of an iterative method. The condition of the matrix is high compared to that of low-order methods resulting in slower convergence (e.g., $\mathcal{O}(N^4/h^{-2})$ for two dimensional (2D) elliptic problems [26]), and the systems are not M-matrices, a valuable theoretical property for algebraic multigrid.

The multigrid method has found success for a variety of applications [32], particularly in the scope of elliptic problems. Algebraic-based multigrid methods [30] in particular have emerged as robust solvers, relying only on the graph of the matrix to achieve a scalable algorithm. The primary goal of this paper is to investigate the use of algebraic multigrid (AMG) with low-order finite elements as a preconditioner of high-order spectral element discretizations of elliptic problems.

*Received by the editors June 21, 2006; accepted for publication (in revised form) June 26, 2007; published electronically September 28, 2007. This work was partially supported by the NSF under award DMS-0612448.

<http://www.siam.org/journals/sisc/29-5/66346.html>

†Department of Computer Science, 201 North Goodwin Avenue, University of Illinois at Urbana-Champaign, Urbana, IL 61801 (lukeo@uiuc.edu).

The contributions of this paper are twofold. First, an AMG algorithm for solving systems arising from nodal spectral elements on simplexes in two and three dimensions is highlighted. The solution technique utilizes a low-order finite element preconditioner based on a local Delaunay triangulation of the interpolatory nodes. Scalable performance of the multigrid preconditioner is verified. Second, a new hybrid preconditioner is identified for use in a nonconforming setting. In particular, we extend the successful AMG preconditioner in the continuous setting to an effective solution method in the discontinuous framework and highlight the necessary components for achieving a scalable algorithm.

The remainder of the paper is organized as follows. Sections 1.1 and 1.2 briefly detail the spectral element method and the common AMG framework used throughout the paper. In section 1.3, previous results are summarized which provide a basis for the algorithmic details and performance outlined in section 2. The preconditioner for the discontinuous method is introduced in section 3, and we summarize the scope of this research in section 4.

1.1. Discretization. The focus of this paper is on spectral elements of nodal type: basis functions defined by Lagrange interpolating polynomials over a set of nodes. While this framework has many benefits, including direct representation of the solution as coefficients and a decoupling of degrees-of-freedom in the case of inter-element connectivity, it does not possess a direct hierarchical representation. Spectral methods of modal type often do form a hierarchical set and are more directly amenable to multilevel solution techniques. The so-called spectral multigrid method developed computationally and theoretically in [22, 27, 23] has proven to be efficient, yielding bounded convergence factors in one dimension and only marginal growth in two dimensions. The results in [12] highlight the versatility of the p -multigrid approach. As discussed in section 1.3, we follow an approach that also utilizes a low-order basis, but we do so through well-established theory on the spectral equivalence of the low-order finite element and high-order spectral element operators.

Let $\Omega \in \mathbb{R}^d$ denote the domain for $d = 2, 3$. Initially we focus on the elliptic PDE

$$(1.1a) \quad -\nabla \cdot \nabla u(\mathbf{x}) = f(\mathbf{x}) \quad \text{in } \Omega,$$

$$(1.1b) \quad u(\mathbf{x}) = g(\mathbf{x}) \quad \text{on } \partial\Omega,$$

where AMG has been successful for low-order finite element methods. The associated weak problem for the Galerkin method becomes the following: Find $u_{h,n} \in V$ so that

$$(1.2) \quad a(u_{h,n}, v_{h,n}) = \ell(v_{h,n}) \quad \forall v_{h,n} \in V,$$

where

$$(1.3) \quad a(u, v) = \int_{\Omega} \nabla u \cdot \nabla v \, d\mathbf{x}, \quad \ell(v) = \int_{\Omega} f(\mathbf{x})v \, d\mathbf{x},$$

and V denotes the spectral element basis. The system of equations $Au = b$ that results from (1.3) is the focus of the first portion of this paper.

As explained above, the basis functions V are nodal. As a reference element, we use the d -simplex $\hat{\kappa}_d = \{(\xi_1, \dots, \xi_d) \in \mathbb{R}^d : \xi_i \in [-1, 1], \sum_{i=1}^d \xi_i < 0\}$ and denote the complete space of n -degree polynomials on $\hat{\kappa}_d$ with minimal dimension as P_n^d . The dimension of this space [13] is $\dim P_n^d \equiv N_n^d = \binom{d+n}{d}$. Over $\hat{\kappa}$ we define $\phi_{n,i}(\mathbf{x}) \in \hat{V}$ as a Lagrange interpolating polynomial with the property $\phi_{n,i}(\mathbf{x}_j) = \delta_{ij}$, where δ_{ij} is the Kronecker delta function, \mathbf{x}_i are nodal locations, and i ranges from 1 to $N = N_n^d$.




FIG. 1.1. Nodal sets with $n = 10$ (1.1(a)) and $n = 6$ (1.1(b)) on the reference elements.

In one dimension, the Gauss–Legendre–Lobatto (GLL) nodes offer a highly stable and complete set interpolation points for a Lagrange basis. For higher dimensions, a natural set of nodal locations is not available. Much progress has been made in defining stable and accurate nodal sets [31, 33]; we employ nodal locations based on electrostatics [13] due to their availability and near-optimal numerical properties. The results in this paper were largely invariant with respect to small variations in the nodal locations, namely using the Fekete nodal set [31] or Warp&Blend nodes [33]. In two dimensions, the electrostatic nodal set we use enforces GLL nodes on each edge of $\hat{\kappa}_2$, letting the remaining $N_n^2 - 3N_n^1 - 3$ interior nodes be determined by symmetry and electrostatic distribution. Similarly, in three dimensions, we impose the 2D distribution on each face of the tetrahedron and allow the remaining nodal locations in the interior to be determined also from electrostatics. The values are determined in preprocessing or from a library and only for the reference element $\hat{\kappa}_d$. An example is illustrated in Figure 1.1. As with GLL nodes in one dimension, the electrostatic nodes (equal to GLL nodes in one dimension) are clustered in the corners.

1.2. Algebraic multigrid. We consider the solution of high-order $A_{sp}u = b$ using algebraic-based multigrid methods together with standard Krylov acceleration methods such as the conjugate gradient (CG). The AMG algorithm we adhere to is the classical Ruge–Stüben (RS) [30] algorithm, with a few modifications, which are mentioned throughout this paper. The MG cycling we use is the familiar V-cycle summarized for two levels by the following:

1. Presmooth ν_1 times on the fine grid Ω^0 $u^0 \leftarrow S(u^0; \nu_1)$.
2. Restrict residual $r^0 = b^0 - Au^0$ to the coarse grid Ω^1 $r^1 = Rr^0$.
3. Solve the coarse grid problem $A^1e^1 = r^1$ on Ω^1 $e^1 \approx (A^1)^{-1}r^1$.
4. Interpolate the coarse-grid error to Ω^0 $e^0 = Pe^1$.
5. Correct the previous approximation on Ω^0 $u^0 \leftarrow u^0 + e^0$.
6. Postsmooth ν_2 on the fine grid Ω^0 $u^0 \leftarrow S(u^0; \nu_1)$.

The setup phase is more tedious and impacts on the application to high-order discretizations. Unknown u_i is connected to u_j if $a_{ij} \neq 0$. The size of a_{ij} indicates the impact of this connection, relating how much influence the error at j has on the error at i in the relaxation process. We say that u_i is strongly connected [30] to u_j if

$$(1.4) \quad -a_{ij} \geq \theta \max_{k \neq i} -a_{ik},$$

where a threshold of $\theta = 0.25$ is used in our computations, a typical selection for

elliptic problems. We refer to [30] for details on the selection of coarse grids.

For a given C/F splitting of Ω^0 , the next step in the setup phase is to define interpolation. The θ parameter in the definition of the strength-of-connection controls the density of coarse grids. Likewise, in the definition of interpolation, we need to ensure that corresponding coarse-grid operators A^k are sufficiently sparse. In section 2, long-range connections are not used in interpolation: only direct connections are used in the definition. However, it does not yield effective interpolation for all problems, as highlighted in section 3.

Although AMG is well suited as a solution technique, it is also used for many problems as an effective preconditioner. This process can be viewed as acceleration since Krylov steps can often reduce error components that the multigrid cycle is unable to annihilate. Furthermore, each step of the Krylov method is relatively inexpensive, requiring on the order of one matrix-vector multiply.

ALGORITHM 1. AMG-PCG ALGORITHM.

```

1   $r_{new} = r_0 = b - Ax_0$ 
2   $MG_{setup}(\hat{A})$ 
3   $p = z = MG_{cycle}(\hat{A}, r_0)$ 
4  while  $\|r_{new}\|/\|r_0\| > \delta$ 
5     $\alpha := (r, z)/(Ap, p)$ 
6     $x_{new} := x + \alpha p$ 
7     $r_{new} := r - \alpha Ap$ 
8     $z_{new} := MG_{cycle}(\hat{A}, r_{new})$ 
9     $\beta := (r_{new}, z_{new})/(r, z)$ 
10    $p_{new} := z_{new} + \beta p$ 
11   update  $r, z, p$ 

```

The construction of \hat{A} in Algorithm 1 is central to the algorithm introduced in sections 2 and 3. Specifically, steps 2, 3, and 8 require the formulation of \hat{A} along with the definition of the associated multigrid process. The reference to MG_{cycle} in Algorithm 1 consists of one V-cycle for the computations reported in this paper, but the use of multiple cycles or W-cycles may be easily incorporated.

1.3. AMG vis-à-vis high order. The strategy proposed in this paper relies on the ability of low-order finite element operators to precondition the high-order spectral element discretization. Let A_s be the matrix from a spectral element discretization, and let A_f be the resulting matrix from a low-order method defined over a tessellation of the nodal set in the spectral element discretization. The concept is well developed with early contributions in [24]. In [24], it is shown that using a finite-difference operator to precondition a matrix from a *spectral* method is effective in one dimension. The use of finite elements is detailed in [8] for one dimension, and the effectiveness for two dimensions is detailed in [10]. Norm dependent bounds on $A_f^{-1}A_s$ are further developed in [25], the impact of advection and reaction terms are studied in [25, 11], and the 2D and 3D Helmholtz problem is studied in [4]. The condition number in higher dimensions, however, does not in general remain bounded, and the effectiveness of preconditioning is decreased when using simplexes [10]. Still, the condition of the system is improved and is suitable for Krylov iterative methods. An extensive study focusing on Schwarz methods in this context can be found in [21] and has motivated much the effort here.

As an example of the low-order preconditioner, consider the 3×3 spectral element mesh of order $n = 10$ displayed in Figure 1.2(a). The nodal set used in each element is a tensor set of the one dimensional (1D) GLL nodes. An effective preconditioner




FIG. 1.2. Meshes corresponding to the spectral element operator A_s and to the finite element preconditioner A_f , respectively.

utilizes the nodal locations, forming a bilinear finite element mesh resulting in A_f . The effect is illustrated in Figure 1.2(b). The performance using low-order finite element preconditioning with AMG is studied for quadrilaterals in [16]. One AMG V-cycle is used to approximate the action of A_f^{-1} to precondition the spectral element matrix.

Two aspects are highlighted in [16] which motivate the algorithms developed in this paper. First, for both Poisson's equations and Stoke's flow (H^1 -elliptic least-squares form), the use of a bilinear finite element preconditioner with AMG as illustrated in Figure 1.2 scales with spectral element order. The preconditioner is known to be effective; however, the inexact and inexpensive preconditioning of AMG is sufficient. More precisely, CG iterations exhibit a linear increase with spectral order n , while AMG preconditioned CG (AMG-PCG) iterations remain invariant, needing only 12–14 iterations in the cases tested. Furthermore, the average convergence factor of the AMG-PCG iterations also scaled with n , remaining constant at around 0.34. Also notable regarding this simulation is that the AMG complexity (detailed below in section 2) remains bounded, although somewhat high.

The subtle nature of complexities is the basis of the second observation in [16]. While AMG-PCG applied directly to the spectral operator A (i.e., $\hat{A} = A$) yields attractive convergence factors of around 0.1 for all spectral orders, the cost of each AMG cycle heavily depends on n . The overall cost associated in preconditioning with $\hat{A} = A$ scales as $\mathcal{O}(n^4)$, whereas preconditioning with $\hat{A} = A_f$ scales only as $\mathcal{O}(n^2)$. The underlying statement of [16] is that AMG-PCG iterations can be performed directly on the spectral element operator when $n < 4$ and that $\hat{A} = A_f$ preconditioning is beneficial if $n \geq 4$. The later case is even more remarkable since explicit construction of the spectral element operator A is not needed, saving memory and reducing the $\mathcal{O}(n^4)$ construction cost.

The algorithms proposed in this paper follow a similar process; however, we extend the results, implementation, and analysis in a number of directions. The methodology is extended to simplexes and for nontensor nodal sets—based on electrostatics introduced in section 1.1—requiring an unstructured local tessellation utilizing a reference Delaunay triangulation. Additionally, the quality of the meshes is considered and simulations are performed both in two and in three dimensions. A new complexity measure is also introduced to account for the use of a low-complexity, low-order preconditioner in the context of a high-order solver. The new measure highlights the efficacy of low-order finite element AMG preconditioning. Discontinuous spectral elements are also considered, and a novel multilevel preconditioner is introduced.




FIG. 2.1. Local Delaunay tessellation of the reference element for $n = 10$.

2. Continuous methods. In this section we investigate the use of Algorithm 1 for high-order spectral elements on triangular and tetrahedral meshes. A continuous Galerkin method is used (1.3) to discretize the model problem (1.1).

Central to the algorithm is the ability to use a low-order finite element operator as a preconditioner: $\hat{A} = A_f$. For triangles and for tetrahedrons, the nontensor nodal locations (see Figure 1.1) require additional attention, whereas in the case of quadrilaterals, a straightforward structured mesh is used to locally tessellate the nodal set (see Figure 1.2). Here, we use a Delaunay triangulation of the nodal set on the reference element, mapping the local tessellation to each element in the spectral element mesh to construct the global finite element mesh. The Delaunay triangulation is used since it maximizes the minimum angle of the local elements as depicted in Figure 2.1. Moreover, since the Delaunay operation is executed in preprocessing over just the reference element, it does not add to the complexity of the algorithm. The nonuniqueness of tetrahedral tessellations did not result in problematic elements for the nodal sets tested, although postprocessing could easily be included to ensure the exclusion of so-called sliver elements. Since the quality of the tessellation is enforced only on the reference element, the area ratios and minimum angles are not preserved over the global grid. Yet, given a reasonable initial spectral element mesh, we can assume for d -simplexes that the mesh qualities are preserved since the map Jacobian is constant. On the other hand, as with quadrilaterals, the aspect ratio of the mesh can become skewed particularly at high order due to the grouping of nodes in corners. An example is illustrated in Figure 2.2.

2.1. AMG measures. A classical Ruge–Stüben implementation of AMG is used (see section 1.2). We assume a C/F splitting and define interpolation by considering $Ae \approx 0$ for smooth error. That is,

$$(2.1) \quad a_{ii}e_i \approx - \sum_{j \in C_i^s} a_{ij}e_j - \sum_{j \in F_i^s} a_{ij}e_j - \sum_{j \in N_i^w} a_{ij}e_j,$$

where C_i^s are C -points that strongly influence F -point i , F_i^s are F -points that strongly influence i , and N_i^w are points that weakly influence i . First, weak connections N_i^w

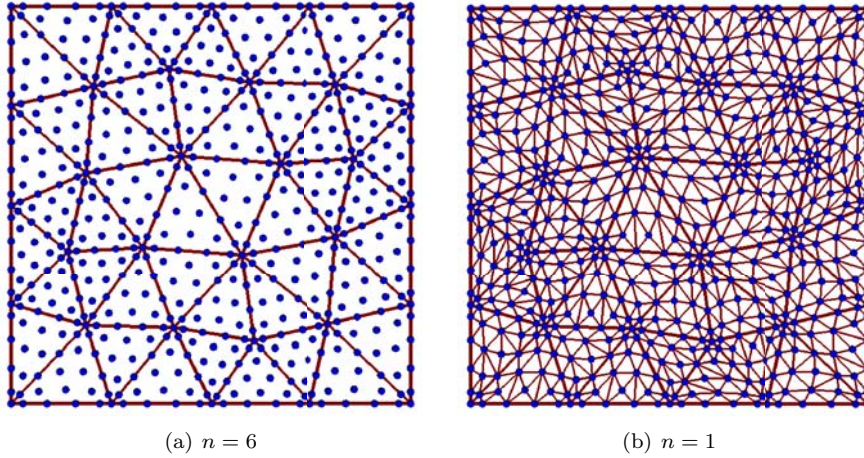


FIG. 2.2. Meshes corresponding to the spectral element operator A and to the finite element preconditioner A_f , respectively.

are distributed to the diagonal as in the original RS description [30] (cf. [3]):

$$(2.2) \quad \left(a_{ii} + \sum_{j \in N_i^w} a_{ij} \right) e_i \approx - \sum_{j \in C_i^s} a_{ij} e_j - \sum_{j \in F_i^s} a_{ij} e_j.$$

Intuitively this is compatible with the high-order discretization as weak connections may be large in magnitude but of the same sign as the diagonal, thus preserving the dominance of the F -point i in the interpolation process. Numerically the performance is only slightly improved.

The remaining term in (2.2) is perhaps more important than the weak connections: F -points strongly influencing i . We also adhere to the original algorithm and distribute direct strong connections along all nodes in C_i^s to ensure that interpolation is reasonable in the event that a strength of an edge has been misidentified for a positive off-diagonal entry.

The remaining parameters in the AMG algorithm are standard and largely invariant with spectral order for the problems tested. We run V(1,1)-cycles in all tests performed, using pointwise Gauss-Seidel over C -points followed by F -points in pre-relaxation and the reverse ordering for postrelaxation in order to preserve symmetry of the AMG preconditioner, a necessary attribute for PCG.

Measuring the performance of Algorithm 1 is not straightforward. Typical tools in analyzing the performance of multigrid include the convergence factor, ρ , which is computed as the average reduction in the residual in each iteration. In each test, we drive Algorithm 1 until the relative residual is reduced by a factor of 10^{-6} . One component of a scalable algorithm is to maintain a constant convergence factor (as low as 0.05 for some problems). The other key element is the complexity of each step in the process. For AMG we consider the *operator complexity* for a matrix A^0 :

$$(2.3) \quad \chi_{\text{op}} = \frac{\sum_{\ell=0}^{\ell_{\max}} \aleph(A^\ell)}{\aleph(A^0)},$$

where $\aleph(A)$ is the number of nonzeros in matrix A and ℓ_{\max} is the number of levels to be visited in the AMG cycle. The operator complexity is more instructive than

the *grid complexity*, which similarly measures the number of variables on all levels compared with the finest. Additionally, we define the *cycle complexity* as

$$(2.4) \quad \chi_{\text{cyc}} = \frac{\sum_{\ell=0}^{\ell_{\text{max}}} \nu_{\ell} \cdot \aleph(A^{\ell})}{\aleph(A^0)},$$

where ν_{ℓ} is the number of relaxation sweeps on level ℓ . If we consider one relaxation sweep on the fine grid as one work unit, then χ_{cyc} approximates the amount of work per cycle; for a V(1,1)-cycle, χ_{op} and χ_{cyc} differ by a factor of approximately 2. As we intend to iterate on both spectral and finite element matrices, we denote by χ_{cyc}^s and χ_{cyc}^f the respective cycle complexities.

The results listed in [16] indicate that the cycle complexities χ_{cyc}^f using finite element preconditioning are much larger than the complexities of AMG applied directly to the spectral element operator χ_{cyc}^s . Moreover, χ_{cyc}^s scales with polynomial order, while χ_{cyc}^f grows slightly. If we consider the complexity of the overall algorithm, the results are even more compelling. The fundamental difficulty with χ_{cyc}^f is that $A_f u = b$ is not the target system and does not dominate the process. Rather, the matrix-vector multiply in the CG wrapper involves spectral element operator A_s and comprises much of the cost. This identifies the suboptimal performance indicated by χ_{cyc}^f . Additionally, due to the extreme density of the high-order operator on the finest level, the spectral complexity χ_{cyc}^s is deceptively small. To extend these definitions, we propose the following hybrid operator complexity which accounts for the finite element operator within the spectral PCG wrapper (Algorithm 1):

$$(2.5) \quad \chi_{\text{op}}^* = \frac{\aleph(A^0) + \sum_{\ell=0}^{\ell_{\text{max}}} \aleph(A_f^{\ell})}{\aleph(A^0)}$$

$$(2.6) \quad = \frac{\aleph(A^0) + \chi_{\text{op}}^f \cdot \aleph(A_f^0)}{\aleph(A^0)}.$$

A similar definition for χ_{cyc}^* follows.

Combining the work per cycle (2.6) with the convergence factor ρ , we define the overall work as the number of work units to reduce the residual by a factor of 10:

$$(2.7) \quad W = -\frac{\chi_{\text{cyc}}}{\log_{10} \rho}$$

or

$$(2.8) \quad W^* = -\frac{\chi_{\text{cyc}}^*}{\log_{10} \rho}.$$

2.2. Numerical tests. In this section we present numerical results validating the scalability of Algorithm 1 with the measures introduced in the previous section. In particular, we refer to four situations listed in Table 2.1.

The performance results in Figures 2.3(c) and 2.4 are for a spectral element unstructured triangulation of the unit square with 574 elements. The mesh quality is high, with a normalized minimum angle of 0.7 (over a scale of 0.0 to 1.0) and an area ratio of approximately 0.5. The solver performance was consistent for both coarser and finer spectral element meshes, with nearly constant element diameters across the tessellation; for brevity we present only this particular case.

TABLE 2.1
List of methods.

CG- A_s	The classical CG method applied to the target system of equations $A_s u = b$.
AMG- A_s	AMG V(1,1)-cycles as a solver for $A_s u = b$, rather than used as a preconditioner.
PCG-AMG- A_s	AMG- A_s with acceleration. Algorithm 1 with $\hat{A} = A$.
PCG-AMG- A_f	AMG preconditioned CG using low-order finite element preconditioning. Algorithm 1 with $\hat{A} = A_f$.


The methods in Table 2.1 are compared as the order of the spectral element is increased from linear ($n = 1$) to very high-order ($n = 11$). Figures 2.3(a)–2.3(c) highlight the efficiency of low-order preconditioning with AMG. As a benchmark, pure CG iterations (CG- A_s) are also listed in the first figure. From Figure 2.3(b) it is evident that AMG alone is unable to handle orders beyond $n = 6$. This is consistent with previous attempts where direct application of AMG was found to become ineffective at $n \approx 4$ [29]. Furthermore, both Figures 2.3(a) and 2.3(b) confirm the scalability (in spectral order) for low-order preconditioning. Krylov acceleration alone is unable to effectively capture problematic eigenmodes in the faulty AMG- A_s . The level curve for PCG-AMG- A_f compared with the nontrivial slope in the PCG-AMG- A_s approach confirms that the low-order operator A_f combined with AMG is accurately preconditioning the spectral operator in the CG algorithm.

Figures 2.3(a) and 2.3(b) do not fully address the complexity of the algorithm. The complexity of steps 2 and 8 in Algorithm 1 will vary depending on \hat{A} . For both AMG- A_s and PCG-AMG- A_s , we base the work (2.7) on complexity χ_{cyc}^s as in (2.4). The PCG-AMG- A_f algorithm with $\hat{A} = A_f$ requires the hybrid complexity χ_{cyc}^* defined by χ_{op}^* in (2.6), and the associated work is defined in (2.8). The work is presented in Figure 2.3(c). Since AMG is able to maintain a modest level of complexity for the low-order operator as n increases, the overall complexity remains close to 1.0, while the work increases significantly when AMG is applied to the spectral operator A_s .


The low complexity associated with AMG applied to A_f is predictable, as the density of the operator is much less than that of the spectral operator A_s (cf. Figure 2.4): the number of connections per node in the graph of A_f was on average 6, and the maximum number of connections increases from 8 to 12 with n . The notable AMG advantage is the ability to handle large variations in the size of the elements. The area ratios of the finite element mesh decrease with $\mathcal{O}(n^4)$ (cf. Figure 2.4) as the polynomial order increases, while the minimum angle degrades only slightly, being preserved by the local Delaunay process. This view captivates the necessity of using an algebraic-based solver in the context of low-order preconditioning on unstructured grids.

We extend this approach to three dimensions in a similar manner over a spectral element unstructured mesh of 188 elements. The performance for the methods listed in Table 2.1 are presented in Figure 2.5 and reveal a trend similar to that of two dimensions. AMG combined with CG acceleration effectively preconditions the spectral element operator A_s when combined with A_f , while use of AMG with A_s yields both poor convergence and high complexities.


Similar to the case of quadrilaterals [16], AMG should be used directly with A_s when the spectral order is low $n \leq 4$, while low-order preconditioning is beneficial for high-order $n > 4$. When matrix-free application of the spectral operator A_s is available, then $\hat{A} = A_f$ gains even more appeal, although construction and management



(a) Iterations



(b) Convergence Factors



(c) Total Work

FIG. 2.3. Performance measures in two dimensions: The convergence factors (ρ), iterations to convergence, and work per digit of accuracy are listed for the methods in Table 2.1.





FIG. 2.4. Mesh diagnostics: The number of connections in both A_s (▽) and in A_f (△) along with the normalized minimum angle (□) and the area ratio (○) in the finite element mesh.


of the low-order finite element operator should be considered. From Figure 2.3(c), we notice a similar separation at $n = 4$, with PCG-AMG- A_f remaining relatively inexpensive. Moreover, the 3D simulations show this threshold to be at slightly lower orders—at $n \approx 2$ or 3—as illustrated in Figure 2.5(c).

The deterioration in convergence for classical AMG applied directly to A_s is largely due to the underlying M-matrix assumptions of the method. Both strength-of-connection and interpolation construction in RS-AMG are motivated through principles based on an M-matrix. Indeed, much of the early convergence theory on AMG relies on the system being an M-matrix [30], and the high-order matrices A_s considered here do not have this property: $a_{ij} > 0$ for some (i, j) (apart from the linear case). A well-known theoretical tool to generalize the theory is the concept of an *essentially positive* matrix [2]. For positive off-diagonal connections a_{ij} , if there exists a d -path with sufficiently large negative weight, then the AMG convergence results still hold [32]. Computationally, however, the adjustments for essentially positive matrices (i.e., basing coarse-grid selection and interpolation on negative weights only) have little impact. The results vary little whether positive connections are ignored or collapsed.


The notable change in performance for AMG at certain polynomial orders (e.g., $n = 4$ in two dimensions) is that the matrices are no longer essentially positive (with a path length of two). As an example, in the 2D case with a effective grid size of $h \approx 1/4$, the matrix is heavily populated with off-diagonals from $n = 2$, but the essentially positive property is not lost until $n = 5$ (see Figure 2.6). The number of positive edges in the graph of A_s grows with $\mathcal{O}(n^4)$, making classical AMG strength measures increasingly less effective. The $n = 4$ case coincides also with the deterioration of AMG convergence, thus confirming these M-matrix diagnostics. One common problem in utilizing very long-range connections in interpolation is that the computation is overwhelmed by the operator complexities.



(a) Iterations



(b) Convergence Factors



(c) Total Work

FIG. 2.5. Performance measures in three dimensions: The convergence factors (ρ), iterations to convergence, and work per digit of accuracy are listed for the methods in Table 2.1.




FIG. 2.6. The increase in positive off-diagonal connections that do not have a sufficiently large negatively weighted path of length two (solid line). The dark dashed line is the $\mathcal{O}(n^4)$ reference, and the dashed gray line indicates the percentage of off-diagonals that are positive.

3. Discontinuous methods. The DG method was initially introduced as a solution method for solving hyperbolic and convection dominated problems [6]. More recently, the DG methodology has successfully been extended to a variety of problems, becoming a robust and popular framework. In this section we describe the DG approach followed throughout this paper.

A variety of DG methods have been proposed in recent years for various problems [1]; we follow a *local* DG method (LDG) [7], which is described as follows. Recalling the model elliptic problem (1.1), let $\mathbf{q} = \nabla u$ and rewrite the problem as a first-order system:

$$\begin{aligned}
 (3.1a) \quad & \mathbf{q} - \nabla u = 0 \quad \text{in } \Omega, \\
 (3.1b) \quad & -\nabla \cdot \mathbf{q} = f \quad \text{in } \Omega, \\
 (3.1c) \quad & u(\mathbf{x}) = g \quad \text{on } \partial\Omega.
 \end{aligned}$$

Central to the DG methodology is the discontinuity in the approximation between elements; the DG varieties are distinguished by the way interelement connectivity is imposed. Let \mathcal{K}^h be a tessellation of Ω , and consider the weak form of (3.1) over each $\kappa \in \mathcal{K}^h$:

$$(3.2a) \quad \int_{\kappa} \mathbf{q} \cdot \boldsymbol{\psi} \, d\mathbf{x} + \int_{\kappa} u \nabla \cdot \boldsymbol{\psi} \, d\mathbf{x} = \oint_{\partial\kappa} u \boldsymbol{\psi} \cdot \mathbf{n}_{\kappa} \, d\mathbf{x},$$

$$(3.2b) \quad \int_{\kappa} \mathbf{q} \cdot \nabla \phi = \oint_{\partial\kappa} \mathbf{q} \cdot \mathbf{n}_{\kappa} \, d\mathbf{x} + \int_{\kappa} f \phi \, d\mathbf{x},$$

where $\phi(\mathbf{x})$ and $\boldsymbol{\psi}(\mathbf{x})$ are scalar and vector test functions, respectively, and \mathbf{n}_{κ} is the outward normal on κ . The interelement connectivity is now determined by definition of numerical flux on each edge. On element κ , define u^- to be the value of u interior to the element, and define u^+ to be the value of u in the adjacent neighboring element. For a scalar function u and vector function \mathbf{q} , the *jump* and the *average* between neighboring elements are respectively defined as $[[u]] = u^- \mathbf{n}^- + u^+ \mathbf{n}^+$, $\{u\} = \frac{1}{2}(u^- + u^+)$, $[[\mathbf{q}]] = \mathbf{q}^- \cdot \mathbf{n}_{\kappa^-} + \mathbf{q}^+ \cdot \mathbf{n}_{\kappa^+}$, $\{\mathbf{q}\} = \frac{1}{2}(\mathbf{q}^- + \mathbf{q}^+)$. We define the numerical fluxes u^* and \mathbf{q}^* independently of ∇u , allowing the formulation of the weak problem (3.2) independently of the flux $\mathbf{q}(\mathbf{x})$. In general, the numerical fluxes for the




FIG. 3.1. Stencil width for the LDG method compared with Brezzi and related DG methods. The width is one for the case of LDG, while dependence extends to two layers (++) for other common DG methods.

LDG method are defined as [1]

$$(3.3) \quad u^* = \{u\} + \beta \cdot \llbracket u \rrbracket, \quad \mathbf{q}^* = \{\mathbf{q}\} - \beta \llbracket \mathbf{q} \rrbracket - \eta_\kappa \llbracket u \rrbracket,$$

which replace u and \mathbf{q} in the elemental boundary terms of (3.2). With this, we strengthen the weak problem slightly using an additional integration-by-parts to arrive at

$$(3.4a) \quad \int_\kappa \mathbf{q} \cdot \boldsymbol{\psi} \, d\mathbf{x} - \int_\kappa \nabla u \cdot \boldsymbol{\psi} \, d\mathbf{x} = \oint_{\partial\kappa} (u^* - u) \boldsymbol{\psi} \cdot \mathbf{n}_\kappa \, d\mathbf{x},$$

$$(3.4b) \quad - \int_\kappa \nabla \cdot \mathbf{q} \phi = \oint_{\partial\kappa} (\mathbf{q}^* - \mathbf{q}) \cdot \mathbf{n}_\kappa \, d\mathbf{x} + \int_\kappa f \phi \, d\mathbf{x}$$

for all $\kappa \in \mathcal{K}$. The construction leads to the linear system $A^{dg}u = b$, which is the focus of the section.

Choosing $\beta = \mathbf{n}_\kappa/2$ ensures symmetry and locality of the associated stiffness matrix [1] and is a standard choice which leads to the *local* DG method.

3.1. AMG modification. While the LDG method does exhibit many advantageous properties in terms of the approximation both mathematically and computationally, it introduces additional complexities in the algebraic system. The issues concerning AMG are illustrated in this section, including modifications to the basic RS algorithm.

A key component of the LDG formulation is the so-called penalty term, identified by the coefficient η_κ in (3.3). From an approximation standpoint, η_κ offers stability control as it moderates the effect of the jump in u , $\llbracket u \rrbracket$. Moreover, it is well known that the spectrum of the linear system is influenced by the η_κ parameter, impacting many iterative solution methods. For elliptic problems, we can expect a relationship of $\text{cond}(A) = \mathcal{O}(\eta)$ between the condition number and η_κ , with $\eta_\kappa = \eta$ for all $\kappa \in \mathcal{K}$ [5]. With a well-balanced penalty term of $\eta_\kappa = \frac{\eta}{h^2}$ (cf. [34]), the preconditioner proposed below performed well and exhibited no dependence on slight variations in η_κ . This may be partly due to the results reported in [5] indicating less variability with respect to the penalty parameter in the condition number at higher orders.

Our choice in DG schemes was partly based on the fact that LDG limits the breadth of dependence. Figure 3.1 illustrates this effect. The result is that (1) the locality of AMG interpolation will be more effective for tighter stencils and (2) a low-order preconditioner will be more easily identified since construction can be arguably constrained in the first level to interelement connectivity only.

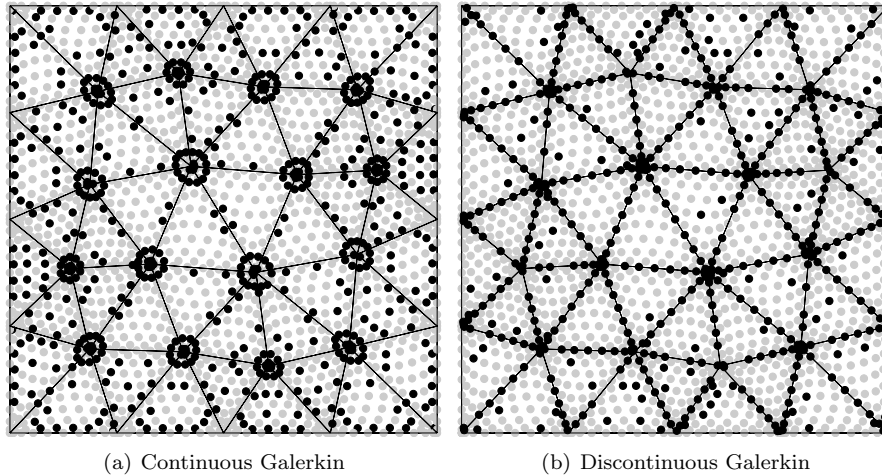


FIG. 3.2. AMG C/F splitting for $n = 11$. The C -points are represented by gray circles and the F points by black circles.

Still, central to DG is the concept of additional degrees-of-freedom along element boundaries. AMG applied directly to the spectral operator in the continuous Galerkin method yields adequate results; we cannot expect such straightforward application in the context of DG methods, as the connections between degrees of freedom along a boundary play a much different role than connectivity with interior nodes. Indeed, Figure 3.2(a) identifies this subtlety for the case of $n = 11$. Figure 3.2(a) is the C/F splitting for the continuous methods analyzed in the previous section, while 3.2(b) depicts the C/F splitting for the first-level. The latter case results in a nonconvergent method and, more notably, very high complexities.

The solution is to extend the definition of “strong connections” in the AMG setup phase. In the classical RS algorithm presented earlier, only *direct* connections in the graph of the matrix are considered as possible interpolation nodes. In the DG framework, this results in element boundary nodes being interpolated only from nodes interior to the element and from nodes on adjacent element boundaries. Conversely, in the continuous Galerkin setting, the scope of interpolation extends into both elements for points along an edge. We accomplish this extension by considering long-range connections in the graph for interpolation (see Figure 3.3).

To consider long-range influence, we use an approach based on relaxation in defining the strong connections since using relaxation circumvents the need for an M-matrix. The implementation was introduced in [28] but has not been used extensively in practice. For node i , we seek a measure of influence from some node j (not necessarily connected to i). There are three cases in which we would like to account for long-range influence; three examples are illustrated in Figure 3.4. In Figure 3.4(a), we account for the strength of the direct connection i to j . In Figure 3.4(b), we see that i will benefit from having two connections to a pair $\{j, k\}$ with a strong connection. The influence of k on j will either strengthen the existing connection to i or possibly balance the connection, leaving the influence of j on i as weak. Finally, Figure 3.4(b) identifies the most common example of long-range influence: no connection between i and j , but a common connection to node k .

To identify the influence in the cases of Figure 3.4 we employ the following algo-




FIG. 3.3. Long-range connections: Node i is directly connected to j indices, while having long-range connections to k indices. The direct connections are solid black, and the long-range (or direct-direct) connections are dashed black.




FIG. 3.4. Three cases for long-range influence: only direct, common direct, and indirect.

rithm:

1. Initialize vector v to be 1 at node j and zero otherwise.
2. Relax at i .
3. Relax at each k .
4. Relax again at i .

Since we consider smooth error to have little variation in the direction of a strong connection, a large value of v_i will indicate a strong connection between i and j . Classical RS interpolation is then implemented.

3.2. Hybrid preconditioner. The construction of a low-order preconditioner resembles the low-order finite element preconditioner studies in section 2; however, we introduce several modifications to account for the discontinuity. Coupled with the modified AMG algorithm in section 3.1, the process yields a scalable method.

The preconditioner we present is motivated by three components:

1. Low-order preconditioning is effective in the simplex (cf. section 2).
2. Continuous, linear finite elements are effective coarse spaces in domain decomposition for low-order DG methods [19].
3. Interelement connectivity in the spectral element DG method is closely associated with the spectrum of the matrix (cf. (3.4)).

Combining steps 1–3 into a suitable algorithm provides the basis of our approach. Given the bilinear form $a^{dg}(u, v)$ defined by the LDG method (3.4), we have the

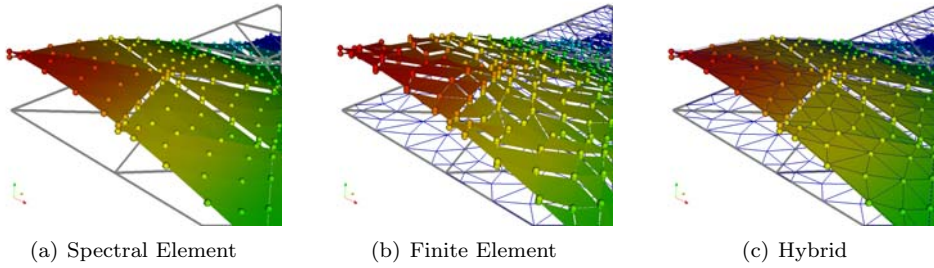


FIG. 3.5. Grid vectors in a spectral element, finite element, and locally projected (hybrid) space.

following weak problem defined: Find $u, v \in V^{dg}$ so that

$$(3.5) \quad a^{dg}(u, v) = \ell(v) \quad \forall v \in V^{dg},$$

where $V^{dg} = V_s^{dg}$ defines the discontinuous nodal spectral element basis—Lagrange interpolating polynomials over electrostatic nodes as defined in section 1.1—over the spectral element mesh $\mathcal{K}_s = \mathcal{K}$. This leads to the matrix equation $A_s^{dg}u = f$. Similarly, we define (3.5) with a discontinuous linear finite element space V_f^{dg} over the locally constructed finite element mesh \mathcal{K}_f of section 2 (cf. Figure 2.2), resulting in a (low-order) matrix equation $A_f^{dg}u = b$.

Motivated by [19] we construct operator C to project the LDG problem into a locally *continuous* finite element space. Denote by V^* a hybrid space of continuous finite elements within each spectral element. V^* is defined as

$$(3.6) \quad V^* = \{u \in V_f^{dg} : u \in H_0^1(\kappa) \forall \kappa \in \mathcal{K}_s\}.$$

Functions in V^* are continuous across local elements but remain discontinuous across spectral elements $\kappa \in \mathcal{K}_s$. Defining $C : V_f^{dg} \rightarrow V^*$ by averaging the jump, $u \leftarrow (u^+ + u^-)/2$, we arrive at the hybrid preconditioner matrix

$$(3.7) \quad A^* = CA_f^{dg}.$$

The preconditioner is illustrated in Figure 3.5 using an example grid vector. The underlying spectral element mesh with for $n = 5$ is apparent in Figure 3.5(a), while the discontinuous finite element problem $A_f^{dg}u = b$ is evident by the jumps and locally refined mesh in Figure 3.5(b). Figure 3.5(c) highlights the hybrid preconditioner, where jumps are maintained between spectral elements and the local finite element space is continuous.

With the hybrid preconditioner defined, the advantage of using long-range connections as proposed in section 3.1 can be seen in Figure 3.6. The connectivity between an element boundary node i and nodes in adjacent elements (e.g., j) is limited in the case of direct connections. However, with long-range connections, the possible influence is more extensive, mimicking the case of continuous elements.

3.3. Numerical tests. We confirm the efficacy of the hybrid preconditioner proposed in the previous section by considering the model problem (1.1) in two dimensions. An LDG method is constructed based on (3.4), and we test the performance as the order of the spectral elements increases by considering the complexity measures presented in section 2.1.




FIG. 3.6. Long-range connections in the hybrid preconditioner: node i and node j with direct connections (solid) and additional long-range connections (dotted).

TABLE 3.1

AMG preconditioning for DG: Polynomial order n is compared with the number of iterations of Algorithm 1, the average convergence factor ρ , the total work per digit accuracy W^* , the number of nonzeros in the matrices, and the total number of degrees-of-freedom in the spectral element problem N .

n	Iterations	ρ	W^*	$\mathfrak{N}(A_s^{dg})$	$\mathfrak{N}(A^*)$	N
1	19	0.46	9.77	3420	3420	138
2	8	0.18	3.33	13680	14076	276
3	9	0.21	3.19	38000	35100	460
4	9	0.20	2.57	85500	59184	690
5	6	0.09	1.61	167580	98244	966
6	9	0.18	1.87	297920	141408	1288
7	11	0.26	2.05	492480	194508	1656
8	12	0.29	2.02	769500	251712	2070
9	12	0.29	1.88	1149500	321372	2530
10	13	0.34	2.05	1655280	399312	3036
11	12	0.30	1.68	2311920	478908	3588

We present in Table 3.1 the performance measures for Algorithm 1 with $\hat{A} = A^*$. As the polynomial order increases, the number of iterations needed to converge along with the average convergence factor remain generally bounded. Furthermore, the complexity measure χ_{op}^* justified in section 2.1 combined with the average convergence factor yields a low, bounded work measure W^* , confirming the scalability of the method with respect to spectral order. Additional simulations on both coarser and finer spectral element grids generated nearly identical trends.

The performance listed in Table 3.1 is encouraging,¹ particularly in view of the number of nonzeros in the matrix. The ratios of nonzeros (columns 5 and 6) to problem size N (column 7) is plotted in Figure 3.7. The complexity of the fine grid preconditioning operator A^* remains low, while the density in the spectral operator A_s^{dg} grows with n , as expected.

Finally, the C/F splitting in AMG for the fine-grid hybrid preconditioner is presented in Figure 3.8. Comparing with Figure 3.2, the long-range connections used for

¹The case of $n = 1$ is trivial. The additional and unnecessary work in constructing the preconditioner is evident in the results, which are provided for completeness.




FIG. 3.7. The ratio of nonzeros in matrix A_s^{dg} (solid) and in matrix A^* (dashed) to the matrix size N .




FIG. 3.8. AMG C/F splitting for $n = 11$ for the matrix of the hybrid discontinuous preconditioner. The C-points are represented by gray circles and the F points by black circles.

the hybrid preconditioner are able to produce a more compatible splitting, resulting in both better convergence and lower complexity.

4. Summary. The proposed computational framework is shown to be an effective approach for solving high-order spectral element methods. The algorithm, based on low-order preconditioning, provides a straightforward way of combining two highly successful methodologies: algebraic multigrid and high-order spectral element discretizations.

The present work contributes to the development of scalable solvers for high-order discretizations in several ways. A low-order preconditioner is extended to d -simplexes using a local Delaunay triangulation, while hybrid measures are proposed to accurately weigh the cost of using a low-order multilevel preconditioner in the context of a high-order spectral element discretization. Both 2D and 3D results highlight the efficacy for continuous methods and provide a motivational basis in moving toward discontinuous discretizations.

A new multilevel preconditioner for discontinuous methods is introduced. The approach utilizes an extension to the classical Ruge–Stüben AMG which accounts for long-range connections in the coarse-grid selection phase and in construction of the intergrid transfer operators. The AMG modifications are used in combination with a local low-order continuous finite element preconditioner to efficiently solve the spectral element problem.

Acknowledgment. The author thanks the developers of Sledge++, a discontinuous spectral element library [35], which provided the basis for the numerical results presented in this paper. Additionally, portions of the RS implementation and modifications were developed by John Ruge; the author is appreciative of the many helpful discussions.

REFERENCES

- [1] D. N. ARNOLD, F. BREZZI, B. COCKBURN, AND L. D. MARINI, *Unified analysis of discontinuous Galerkin methods for elliptic problems*, SIAM J. Numer. Anal., 39 (2002), pp. 1749–1779.
- [2] A. BRANDT, *Algebraic multigrid theory: The symmetric case*, Appl. Math. Comput., 19 (1986), pp. 23–56.
- [3] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A Multigrid Tutorial*, 2nd ed., SIAM, Philadelphia, 2000.
- [4] C. CANUTO AND P. PIETRA, *Boundary and interface conditions within a finite element preconditioner for spectral methods*, J. Comput. Phys., 91 (1990), pp. 310–343.
- [5] P. CASTILLO, *Performance of discontinuous Galerkin methods for elliptic PDEs*, SIAM J. Sci. Comput., 24 (2002), pp. 524–547.
- [6] B. COCKBURN, G. E. KARNIADAKIS, AND C.-W. SHU, *The development of discontinuous Galerkin methods*, in *Discontinuous Galerkin Methods* (Newport, RI, 1999), Lecture Notes in Comput. Sci. Eng. 11, Springer, Berlin, 2000, pp. 3–50.
- [7] B. COCKBURN AND C.-W. SHU, *The local discontinuous Galerkin method for time-dependent convection-diffusion systems*, SIAM J. Numer. Anal., 35 (1998), pp. 2440–2463.
- [8] M. DEVILLE AND E. MUND, *Chebyshev pseudospectral solution of second-order elliptic equations with finite element preconditioning*, J. Comput. Phys., 60 (1985), pp. 517–533.
- [9] M. O. DEVILLE, P. F. FISCHER, AND E. H. MUND, *High-Order Methods for Incompressible Fluid Flow*, Cambridge Monogr. Appl. Comput. Math. 9, Cambridge University Press, Cambridge, UK, 2002.
- [10] M. O. DEVILLE AND E. H. MUND, *Finite-element preconditioning for pseudospectral solutions of elliptic problems*, SIAM J. Sci. Statist. Comput., 11 (1990), pp. 311–342.
- [11] M. O. DEVILLE AND E. H. MUND, *Finite element preconditioning of collocation schemes for advection-diffusion equations*, in *Iterative Methods in Linear Algebra* (Brussels, 1991), North-Holland, Amsterdam, 1992, pp. 181–189.
- [12] B. T. HELENBROOK, D. MAVRIPLIS, AND H. L. ATKINS, *Analysis of “p”-multigrid for continuous and discontinuous finite element discretizations*, in *Proceedings of the 16th AIAA Computational Fluid Dynamics Conference*, Orlando, FL, 2003.
- [13] J. S. HESTHAVEN, *From electrostatics to almost optimal nodal sets for polynomial interpolation in a simplex*, SIAM J. Numer. Anal., 35 (1998), pp. 655–676.
- [14] J. S. HESTHAVEN, *High-order accurate methods in time-domain computational electromagnetics. A review*, Adv. Imaging Electron Phys., 17 (2003), pp. 59–123.
- [15] J. S. HESTHAVEN AND T. WARBURTON, *Nodal high-order methods on unstructured grids. I. Time-domain solution of Maxwell’s equations*, J. Comput. Phys., 181 (2002), pp. 186–221.
- [16] J. J. HEYS, T. A. MANTEUFFEL, S. F. MCCORMICK, AND L. N. OLSON, *Algebraic multigrid for higher-order finite elements*, J. Comput. Phys., 204 (2005), pp. 520–532.
- [17] M. ISKANDARANI, D. B. HAIDVOGEL, J. C. LEVIN, E. CURCHITSER, AND C. A. EDWARDS, *Multiscale geophysical modeling using the spectral element method*, Computing in Science & Engineering, 4 (2002), pp. 42–48.
- [18] G. E. KARNIADAKIS AND S. J. SHERWIN, *Spectral/hp Element Methods for CFD*, Numer. Math. Sci. Comput., Oxford University Press, New York, 1999.
- [19] C. LASSER AND A. TOSELLI, *An overlapping domain decomposition preconditioner for a class of discontinuous Galerkin approximations of advection-diffusion problems*, Math. Comp., 72 (2003), pp. 1215–1238.

- [20] I. LOMTEV AND G. E. KARNIADAKIS, *A discontinuous Galerkin method for the Navier-Stokes equations*, Internat. J. Numer. Methods Fluids, 29 (1999), pp. 587–603.
- [21] J. W. LOTTES AND P. F. FISCHER, *Hybrid multigrid/Schwarz algorithms for the spectral element method*, J. Sci. Comput., 24 (2005), pp. 45–78.
- [22] Y. MADAY AND R. MUÑOZ, *Spectral element multigrid. II. Theoretical justification*, J. Sci. Comput., 3 (1988), pp. 323–353.
- [23] Y. MADAY, R. MUÑOZ, A. PATERA, AND E. RØNQUIST, *Spectral element multigrid methods*, in Iterative Methods in Linear Algebra (Brussels, 1991), North-Holland, Amsterdam, 1992, pp. 191–201.
- [24] S. A. ORSZAG, *Spectral methods for problems in complex geometries*, J. Comput. Phys., 37 (1980), pp. 70–92.
- [25] S. V. PARTER AND E. E. ROTHMAN, *Preconditioning Legendre spectral collocation approximations to elliptic problems*, SIAM J. Numer. Anal., 32 (1995), pp. 333–385.
- [26] 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.
- [27] E. M. RØNQUIST AND A. T. PATERA, *Spectral element multigrid. I. Formulation and numerical results*, J. Sci. Comput., 2 (1987), pp. 389–406.
- [28] J. RUGE, *AMGS01*, FORTRAN Library, June, 1997.
- [29] J. RUGE, *AMG for higher-order discretizations of second-order elliptic problems*, in Proceedings of the Eleventh Copper Mountain Conference on Multigrid Methods, Copper Mountain, CO, 2003.
- [30] J. W. RUGE AND K. STÜBEN, *Algebraic multigrid*, in Multigrid Methods, S. F. McCormick, ed., Frontiers Appl. Math. 3, SIAM, Philadelphia, 1987, pp. 73–130.
- [31] M. A. TAYLOR, B. A. WINGATE, AND R. E. VINCENT, *An algorithm for computing Fekete points in the triangle*, SIAM J. Numer. Anal., 38 (2000), pp. 1707–1720.
- [32] U. TROTTENBERG, C. W. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Academic Press, San Diego, CA, 2001.
- [33] T. WARBURTON, *An explicit construction for interpolation nodes on the simplex*, J. Engrg. Math., 56 (2006), pp. 247–262.
- [34] T. WARBURTON AND M. EMBREE, *The role of the penalty in the local discontinuous Galerkin method for Maxwell's eigenvalue problem*, Comput. Methods Appl. Mech. Engrg., 195 (2006), pp. 3205–3223.
- [35] T. WARBURTON, J. HESTHAVEN, AND L. WILCOX, *Sledge++ Users' Guide*, <http://www.caam.rice.edu/~timwar/TimWarburton/Sledge++.html>, 2006.