

Two-level Overlapping Schwarz Preconditioning of Nodal Discontinuous Galerkin Discretizations of the Indefinite Helmholtz Equation ¹

J. S. Hesthaven, L. N. Olson ^{*}, L. C. Wilcox

Brown University, Division of Applied Mathematics, Box F, Providence, RI 02915

Abstract

A preconditioned two-level overlapping Schwarz method for solving unstructured nodal discontinuous Galerkin discretizations of the indefinite Helmholtz problem is studied. We employ triangles in two dimensions and in a *local* discontinuous Galerkin (LDG) variational setting. We highlight the necessary components of the algorithm needed to achieve efficient results in the context of high-order elements and indefinite algebraic systems. More specifically, we demonstrate the importance of not only coarse-grid solution sweeps, but also for increased overlap in the subdomain solves as the order of the elements increases. In this paper, we detail the discretization strategy and offer an effective approach to solving the resulting system of equations, with numerical evidence in support.

Key words: Discontinuous Galerkin, Nodal Elements, Spectral Element Method, Domain Decomposition, Preconditioned Additive Schwarz, Krylov Subspace Methods

PACS: 02.70.Dh, 02.60.Dc, 02.60.Cb

1991 MSC: 65N30, 65N55, 65F10, 65Y20, 65N22

^{*} Corresponding author.

Email addresses: Jan.Hesthaven@brown.edu (J. S. Hesthaven),
Luke.Olson@brown.edu (L. N. Olson), lucasw@dam.brown.edu (L. C. Wilcox).

URLs: <http://www.cfm.brown.edu/people/jansh> (J. S. Hesthaven),
<http://dam.brown.edu/people/lolson> (L. N. Olson),
<http://dam.brown.edu/people/lucasw> (L. C. Wilcox).

¹ This work was sponsored by the National Science Foundation Division of Mathematical Sciences under grant number DMS-9977371. The work of JSH was partly supported by NSF Career Award DMS-0132967, by ARO under contract DAAD19-01-1-0631 and by the Alfred P. Sloan Foundation through a Sloan Research Fellowship. The work of LCW was partly supported by NSF Career Award DMS-0132967.

1 Introduction

We consider the Helmholtz equation

$$-\nabla \cdot \nabla u(\mathbf{x}) - \omega^2 u(\mathbf{x}) = f(\mathbf{x}) \quad \text{in } \Omega, \quad (1a)$$

$$u(\mathbf{x}) = g(\mathbf{x}) \quad \text{on } \Gamma, \quad (1b)$$

where $\Omega \subset \mathbb{R}^2$ is a bounded, connected domain and $\Gamma = \partial\Omega$. We can interpret u as an electromagnetic quantity—e.g. a normal component of the electric field—and ω as the wave number. For simplicity, we limit our presentation in this paper to Dirichlet boundary conditions and consider smooth solutions $u \in H^s(\Omega)$ for $s \geq 1$, where $H^s(\Omega)$ is the classic Sobolev space of order s [3].

The Helmholtz equation is often considered a highly simplified form of a general wave equation, making it a convenient test problem for detailing many numerical algorithms. Although the form presented in (1) is evidently straightforward, it does still expose a number of difficulties that we discuss in this paper. The problem turns cumbersome quickly as the wave number increases since the resulting system of equations from the discretization scheme becomes indefinite. The Helmholtz equation is also of interest due to its underlying wavelike nature. The equation can quickly be derived from the hyperbolic wave equation by considering time harmonic solutions. Identifying the key components to efficiently solving this wave problem will carry over into more complicated situations, such as Maxwell’s equations.

From a discretization standpoint, research in numerical approximations to wave type equations has been quite successful in recent years. In particular, for Maxwell’s equations, spectral and *hp* element methods have increased in popularity [19,35]. Combined with the robust discontinuous Galerkin methodology [9] this approach yields accurate approximations for unstructured tessellations of complex domains [20,27] and has also been used to solve Maxwell’s equations [21,34]. We refer to [11,25] and the references therein for an extensive overview of spectral/*hp* techniques for a variety of applications.

Recent developments in preconditioning strategies have motivated the work presented in this paper. Compelling arguments are made in the the work of Fischer and Lottes [31] in the viability of using multilevel algorithms. A number of iterative solution strategies are discussed in [31], leading to the conclusion that a *hybrid* form of an overlapping Schwarz method is most robust. The work there focuses on elliptic, positive definite PDEs discretized using a high-order tensor basis on quadrilaterals. Warburton and Pavarino [33] have implemented a one-level robust Overlapping Schwarz method able to handle unstructured spectral elements, while Toselli and Lasser [28,29,30], Feng and Karakashian [14], and Kanschat [24] conduct a study of Preconditioning Schwarz methods for Discontinuous Galerkin methods. Indefiniteness has been addressed

for low-order discretizations on structured grids in a number of publications [7,5,6], again using overlapping additive Schwarz and Schwarz substructuring preconditioners.

Multigrid methods, closely related to domain decomposition (DD) methods, have also been attempted for high-order discretizations with some success. The results from [31] indicate that multigrid methods perform quite well in certain situations. One approach is to derive coarser grids by reducing the polynomial order until the order is one, then invoking a standard multigrid solver. This approach yields optimal results in 1-D, but quickly lose scalability in higher dimensions [32,36]. Similarly, the so-called *multi-p* method constructs V-cycles based on the *modal* basis, reducing the order, to define the coarser grids [23]. Again, the algorithm is not scalable, but does offer a promising direction. Other multigrid attempts have also produced encouraging results. Instead of solving the high-order system directly, a low-order preconditioner is constructed using bilinear finite elements on the same degrees of freedom. A bilinear finite element mesh is superimposed on the degrees of freedom of the high-order nodal element discretization. The size of the system is not decreased, only the sparsity of the matrix. The preconditioner is then formulated using multigrid. For elliptic, positive definite problems, the performance is near optimal [2] mainly due to boundedness of the condition number of $A_1^{-1}A_p$ [15], where A_1 is the first-order basis and A_p results from a high-order spectral element method. Likewise, the related algebraic based multigrid methods (AMG) quickly lose scalability [37], unless a low-order (bilinear finite element) preconditioner is used in a similar manner [22]. Finally, multigrid has also addressed indefiniteness and a variety of coarsening schedules are detailed in [13]. Here, successful error correction is achieved by using Krylov smoothing to properly treat the difficult, 'negative' eigenmodes.

The multilevel approach taken in this paper is a Schwarz-type method. We consider unstructured meshes and complicated domains, thus limiting a multigrid approach to the less accessible and parallelizable AMG. The method presented in this paper attempts to combine the advantageous features of the Schwarz methods reviewed above. We study the performance of a *hybrid* preconditioned additive Schwarz method that utilizes nested/non-nested coarse grid solves and element overlap to maintain efficient performance as the order of the discontinuous spectral element method increases and indefiniteness becomes more prominent.

The remainder of this paper is organized as follows. In Section 2, the discontinuous spectral element formulation is described. We detail the local discontinuous Galerkin (LDG) method for this problem and outline the triangular nodal elements employed in the discretization. A brief description of the preconditioned additive Schwarz method is presented in Section 3 and the efficacy of this approach over a range of parameters is illustrated in Section 4. We

conclude Section 4 by discussing the effect of overlap versus coarse grid solution with respect to problem complexity. In Section 5, we finish with some concluding remarks.

2 Discontinuous Formulation

In this section we detail the use of high-order spectral elements in a *Local* Discontinuous Galerkin DG (LDG) setting. The DG framework has been particularly successful in solving the time dependent Maxwell’s equations [27,20,19], the time harmonic Maxwell’s Equations [21,34], and hyperbolic PDEs in general [9].

The LDG formulation we follow yields several advantageous properties in the resulting linear system of equations. The global mass matrix is block diagonal, allowing cheap inversion, while symmetry is preserved in the global discretization matrix.

We begin by considering an admissible, shape regular triangulation \mathcal{K} of $\Omega \in \mathbb{R}^2$ [3]. Let $h_\kappa = 1/2 \cdot \text{diam}(\kappa)$, for $\kappa \in \mathcal{K}$ and write $\mathcal{K}_h = \mathcal{K}$ when $h_\kappa \leq 2h$ for all $\kappa \in \mathcal{K}$. The tessellation is assumed to be conforming to $\partial\Omega$, the boundary of Ω . We also assume $\mathcal{K}_h = \cup_k \kappa_k$ is a covering of K nonoverlapping triangles κ_k . For each element $\kappa \in \mathcal{K}$, there then exists an isoparametric mapping to a reference element $\hat{\kappa} : \{(r, s) \in [-1, 1]^2; r + s \leq 0\}$.

The numerical approximation u_h on element $\kappa \in \mathcal{K}_h$ are composed of Lagrange interpolating polynomials $L_j(\mathbf{x})$ at selected degrees of freedom \mathbf{x}_j within κ . In 1-D, we describe these locations as the Gauss-Lobatto-Legendre (GLL) quadrature points. Similarly, for our 2-D reference triangle, $\hat{\kappa}$, we choose a distribution of nodes governed by electrostatics [18]. The resulting collection of nodes on an element are arranged on each edge corresponding to the 1-D GLL locations, while the interior nodes are conveniently clustered so that the Lagrange interpolation is stable to high-order. Letting n be the order of the local polynomial approximation on element κ , the number of unknowns required to ensure this level of resolution is

$$N_\kappa = \frac{(n+1)(n+2)}{2}. \tag{2}$$

We limit the focus of this paper to approximations with constant polynomial order in each element—i.e. $N_\kappa = N$. Figure 1 shows an example on the reference element and a few selected (conforming) elements. Finally, we define $\mathcal{P}_n(\kappa)$, the local spectral element space where we seek an approximation, to

be

$$\mathcal{P}_n(\kappa) = \{v_n \in L^2(\kappa) : v_n(\mathbf{x}) = \sum_{j=1}^N \tilde{v}_j L_j(\mathbf{x})\}. \quad (3)$$

Here, $\tilde{v}_j = v(\mathbf{x}_j)$ are the coefficients on the Lagrange interpolating polynomial $L_j(\mathbf{x})$ and consequently the numerical solution at the N degrees of freedom $\mathbf{x}_j \in \kappa$.

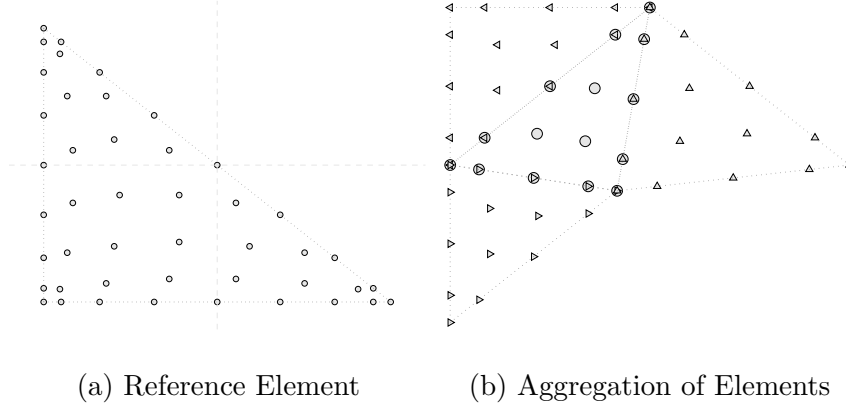


Fig. 1. Node distribution on the reference element for $n = 8$ and on a cluster of elements for $n = 4$.

The standard LDG formulation [1,9] is described first by introducing a slack variable $\mathbf{q} = \nabla u$. The Helmholtz Equation (1) on an arbitrary element κ is transformed into a first-order system

$$-\nabla \cdot \mathbf{q} - \omega^2 u = f \quad \text{in } \kappa, \quad (4a)$$

$$\mathbf{q} - \nabla u = 0 \quad \text{in } \kappa. \quad (4b)$$

Multiplying each equation by scalar and vector test functions $\phi(\mathbf{x})$ and $\vec{\psi}(\mathbf{x})$, respectively, and integrating by parts yields the weak formulation

$$\int_{\kappa} \mathbf{q} \cdot \vec{\psi} \, d\mathbf{x} + \int_{\kappa} u \nabla \cdot \vec{\psi} \, d\mathbf{x} = \int_{\partial\kappa} u \mathbf{n}_k \cdot \vec{\psi} \, d\mathbf{x}, \quad (5a)$$

$$\int_{\kappa} \mathbf{q} \cdot \nabla \phi \, d\mathbf{x} - \omega^2 \int_{\kappa} u \phi \, d\mathbf{x} = \int_{\kappa} f \phi \, d\mathbf{x} + \int_{\partial\kappa} \mathbf{n}_k \cdot \mathbf{q} \phi \, d\mathbf{x}. \quad (5b)$$

For (u, \mathbf{q}) and $(\phi, \vec{\psi})$ in $\Phi \times \Psi$, where

$$\Phi = \{\phi \in L^2(\Omega) : \phi|_{\kappa} \in H^1(\kappa) \forall \kappa \in \mathcal{K}_h\}, \quad (6)$$

$$\Psi = \{\vec{\psi} \in (L^2(\Omega))^2 : \vec{\psi}|_{\kappa} \in (H^1(\kappa))^2 \forall \kappa \in \mathcal{K}_h\}, \quad (7)$$

the weak formulation is well defined [1]. We thus seek an approximation $(u_{h,n}, \mathbf{q}_{h,n})$ to (u, \mathbf{q}) in the discrete spectral element subspace $\Phi_{h,n} \times \Psi_{h,n} \subset$

$\Phi \times \Psi$, where

$$\Phi_{h,n} = \{\phi \in L^2(\Omega) : \phi|_\kappa \in \mathcal{P}_n(\kappa), \forall \kappa \in \mathcal{K}_h\}, \quad (8)$$

$$\Psi_{h,n} = \{\vec{\psi} \in (L^2(\Omega))^2 : \vec{\psi}|_\kappa \in (\mathcal{P}_n(\kappa))^2, \forall \kappa \in \mathcal{K}_h\}. \quad (9)$$

The traces of u and \mathbf{q} are replaced by approximations—*numerical fluxes*— u^* and \mathbf{q}^* , respectively, and are described in more detail below. With this substitution the associated weak problem is: Find $(u_{h,n}, \mathbf{q}_{h,n}) \in \Phi_{h,n} \times \Psi_{h,n}$ such that

$$\int_\kappa \mathbf{q}_{h,n} \cdot \vec{\psi}_n \, d\mathbf{x} + \int_\kappa u_{h,n} \nabla \cdot \vec{\psi}_n \, d\mathbf{x} = \int_{\partial\kappa} u^* \mathbf{n}_k \cdot \vec{\psi} \, d\mathbf{x}, \quad (10a)$$

$$\int_\kappa \mathbf{q}_{h,n} \cdot \nabla \phi_n \, d\mathbf{x} - \omega^2 \int_\kappa u_{h,n} \phi_n \, d\mathbf{x} = \int_\kappa f_{h,n} \phi_n \, d\mathbf{x} + \int_{\partial\kappa} \mathbf{n}_k \cdot \mathbf{q}^* \phi_n \, d\mathbf{x}, \quad (10b)$$

for all $\kappa \in \mathcal{K}_h$ and $(\phi_n, \vec{\psi}_n) \in \mathcal{P}(\kappa) \times (\mathcal{P}(\kappa))^2$.

From the weak problem (10) we can further simplify the resulting global system of equations by considering a slightly stronger formulation. Using integration by parts again results in the form that we will use for the remainder of the paper: Find $(u_{h,n}, \mathbf{q}_{h,n}) \in \Phi_{h,n} \times \Psi_{h,n}$ such that

$$- \int_\kappa \nabla \cdot \mathbf{q}_{h,n} \phi_n \, d\mathbf{x} - \omega^2 \int_\kappa u_{h,n} \phi_n \, d\mathbf{x} = \int_\kappa f_{h,n} \phi_n \, d\mathbf{x} + \int_{\partial\kappa} \mathbf{n}_k \cdot (\mathbf{q}^* - \mathbf{q}_{h,n}) \phi_n \, d\mathbf{x}, \quad (11a)$$

$$\int_\kappa \mathbf{q}_{h,n} \cdot \vec{\psi}_n \, d\mathbf{x} - \int_\kappa \nabla u_{h,n} \cdot \vec{\psi}_n \, d\mathbf{x} = \int_{\partial\kappa} (u^* - u_{h,n}) \mathbf{n}_k \cdot \vec{\psi} \, d\mathbf{x}, \quad (11b)$$

for all $\kappa \in \mathcal{K}_h$ and $(\phi_n, \vec{\psi}_n) \in \mathcal{P}(\kappa) \times (\mathcal{P}(\kappa))^2$.

Remark 1 *The computational advantages and shortcomings of using the primal weak form (11) are not well understood. However results indicate that the numerical approximations obtained from (10) and (11) are very close. For smooth solutions, approximations from the primal form are slightly more accurate. Likewise, approximating discontinuities with the typical, dual form of the weak problem is likely beneficial since the differentiation operates on the test function as in (10), allowing more freedom in the approximation.*

Defining the numerical flux is what separates different discontinuous Galerkin approaches [1] and is the most distinguishing feature of a formulation since the interelement connectivity is solely defined by the representation of the numerical flux on each edge. This choice directly impacts the approximation properties as well as the stability of the method. Moreover, the resulting (global) linear system of equations will perhaps exhibit symmetry and varying sparsity patterns depending on how the trace is approximated along each edge of each element in the tessellation. For a given 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 , the *jump* and the

average between neighboring elements are respectively defined as

$$\llbracket u \rrbracket = u^- \mathbf{n}^- + u^+ \mathbf{n}^+ \quad \{\!\!\{ u \}\!\!\} = \frac{1}{2}(u^- + u^+). \quad (12)$$

Similarly, for vector valued functions \mathbf{q} we define

$$\llbracket \mathbf{q} \rrbracket = \mathbf{q}^- \cdot \mathbf{n}_{k^-} + \mathbf{q}^+ \cdot \mathbf{n}_{k^+} \quad \{\!\!\{ \mathbf{q} \}\!\!\} = \frac{1}{2}(\mathbf{q}^- + \mathbf{q}^+). \quad (13)$$

Furthermore, we define the following collections of edges:

$$\Gamma = \partial\Omega \quad \Gamma_h = \cup_{\kappa \in \mathcal{K}} \partial\kappa \quad (14)$$

$$\Gamma_{\text{int}} = \Gamma_h \setminus \Gamma \quad \Gamma_{\text{bdy}} = \Gamma_h \cap \Gamma. \quad (15)$$

For $\kappa \in \mathcal{K}$ with $\partial\kappa \in \Gamma_{\text{bdy}}$, we define the following special cases:

$$\{\!\!\{ u \}\!\!\} = u^- \quad \llbracket u \rrbracket = u^- \mathbf{n}^- \quad (16)$$

$$\{\!\!\{ \mathbf{q} \}\!\!\} = \mathbf{q}^- \quad \llbracket \mathbf{q} \rrbracket = \mathbf{q}^- \cdot \mathbf{n}^- \quad (17)$$

The *local* nature of the LDG formulation will now be exposed. By defining the numerical fluxes u^* and \mathbf{q}^* independent of ∇u , we will be able to formulate the weak problem (11) independent of the slack variable $\mathbf{q}(\mathbf{x})$. In general, the numerical fluxes for the LDG method are defined as [1]

$$u^* = \{\!\!\{ u_{n,h} \}\!\!\} + \beta \cdot \llbracket u_{n,h} \rrbracket, \quad (18)$$

$$\mathbf{q}^* = \{\!\!\{ \mathbf{q}_{n,h} \}\!\!\} - \beta \llbracket \mathbf{q}_{n,h} \rrbracket - \eta_k \llbracket u_{n,h} \rrbracket. \quad (19)$$

The sign on β is specifically opposite to ensure symmetry of the associated stiffness matrix [10]. Adhering to this form of a numerical flux is beneficial since the method is consistent and locally conservative. Further, if $\eta_k > 0$ the method is considered stable [1]. For the remainder of the paper, we choose η_k to be the same on each element edge, $\eta_k = \tau$, and set $\beta = 0$. The selection described yields a *central* flux for u^* and a *stabilized central* flux for \mathbf{q}^* :

$$u^* = \{\!\!\{ u_{n,h} \}\!\!\}, \quad (20)$$

$$\mathbf{q}^* = \{\!\!\{ \mathbf{q}_{n,h} \}\!\!\} - \tau \llbracket u_{n,h} \rrbracket. \quad (21)$$

Remark 2 Forcing the β terms to vanish has an impact on the sparsity of the resulting linear system. Using the central flux (20) and (21) is more correctly identified as the Brezzi method [4]. Due to the ease of implementation, this formulation has grown in popularity, also benefiting from a slightly improved conditioning over a bona fide LDG method where $\beta \neq 0$. Unfortunately, if $\beta = 0$, the data from elements κ^+ is needed to describe equations (11) in element κ^- as well as data from the neighbors of κ^+ , which we label κ^{++} . Thus the influence on one element extends two layers beyond a given element. The

noncompact stencil is also prevalent for $\beta \neq 0$, unless $\beta = 0.5\mathbf{n}^-$, which corresponds to upwind flux. This is considered the LDG method since fortuitous cancellation of the terms eliminates the extension to neighboring elements, resulting in a stencil width of only one layer. Figure 2 articulates this effect. A more detailed explanation of the effects on discretization error and the eigen-spectrum can be found in [26], although convergence of the iterative solution process is not well understood and is a subject of current research.

Also shown in Figure 2 is the so-called Interior Penalty method (IP). Here, a local gradient is used in the definition of the flux, which also results in a compact stencil. The IP method offers a straightforward implementation, however the poor conditioning of this approach requires careful attention. Table 1 illustrates a typical situation. The results are presented for the definite case ($\omega = 0.0$) on a grid with $h \approx 1/8$. A single level additive Schwarz scheme, which is presented in more detail in Section 3, is used to precondition the GMRES acceleration. The first column reiterates the fact that the Brezzi approach ($\beta = 0.0$) has slightly better conditioning than the LDG implementation ($\beta = 0.5\mathbf{n}^-$), while the IP system suffers from a very poor spectrum. Column 2 also provides insight, showing that while the LDG scheme is slightly more ill-conditioned, the local type preconditioning scheme is more effective due to the compact stencil. The Brezzi operator responds similarly under preconditioning, but due to the wide stencil, the relative improvement is not as drastic. The preconditioning also has significant influence on the IP method, but due to the poor conditioning, it is difficult to draw a complete conclusion about the routine. We will follow the Brezzi method throughout the paper since it is a widely used formulation of DG and since we expect the preconditioning results to be on the pessimistic side. A more comprehensive study of the various DG methods and preconditioning, similar to Table 1, is an ongoing research effort.




Fig. 2. Stencil width relative to element κ^- .

	N	Without PAS	With PAS
Brezzi	2	73	21
	4	167	28
	6	316	30
	8	534	38
LDG	2	121	21
	4	252	29
	6	456	32
	8	713	36
IP	2	355	57
	4	1291	151
	6	> 2000	294
	8	> 2000	568

Table 1

GMRES iterations (no restart) for Brezzi, LDG, and IP formulations with and without preconditioning.

With the numerical fluxes defined, we can outline in more detail the construction of the global algebraic system. The numerical flux u^* is independent of $\mathbf{q}_{h,n}$ allowing us to write the discrete system completely independent of the slack variable \mathbf{q} —c.f. lifting operators in [1]. As we sum the weak problem over all elements $\kappa \in \mathcal{K}$ we will need the following definitions for global matrices S^x , S^y , and M , which are block diagonal, and for matrices $F_{u^*}^{x,y}$, and $F_{q^*}^{x,y}$, which couple nodes in adjacent elements. The local stiffness and mass matrices, $S^{x,y}(\kappa)$ and $M(\kappa)$ are defined as

$$S_{ij}^x(\kappa) = \left(\frac{\partial L_i(\mathbf{x}^\kappa)}{\partial x}, L_j(\mathbf{x}^\kappa) \right)_\kappa, \quad (22)$$

$$S_{ij}^y(\kappa) = \left(\frac{\partial L_i(\mathbf{x}^\kappa)}{\partial y}, L_j(\mathbf{x}^\kappa) \right)_\kappa, \quad (23)$$

$$M_{ij}(\kappa) = (L_i(\mathbf{x}^\kappa), L_j(\mathbf{x}^\kappa))_\kappa, \quad (24)$$

where $\mathbf{x}^\kappa \in \kappa$. The global matrices are of the form $A = \{A_{k,k}\}_{k=1}^K$, where $A_{k,k} = A(\kappa^k)$ for each of the local definitions in (22-24). The operators associated with

the flux term in (11b) are

$$F_{u^*}^x(\kappa) = \sum_{\substack{\kappa^- \in \mathcal{K} \\ \kappa^- \cap \kappa^+ \neq \emptyset}} \int_{\partial\kappa} n_x \frac{L^+(\mathbf{x}^\kappa) - L^-(\mathbf{x}^\kappa)}{2} d\mathbf{x}, \quad (25)$$

$$F_{u^*}^y(\kappa) = \sum_{\substack{\kappa^- \in \mathcal{K} \\ \kappa^- \cap \kappa^+ \neq \emptyset}} \int_{\partial\kappa} n_y \frac{L^+(\mathbf{x}^\kappa) - L^-(\mathbf{x}^\kappa)}{2} d\mathbf{x}, \quad (26)$$

where L^- and L^+ represent the traces of the Lagrange interpolating polynomials on elements κ^- and κ^+ , respectively. Likewise, the flux term in (11a) is written in matrix form as

$$F_{q^*}^x(\kappa) = \sum_{\substack{\kappa^- \in \mathcal{K} \\ \kappa^- \cap \kappa^+ \neq \emptyset}} \int_{\partial\kappa} n_x \frac{L^+(\mathbf{x}^\kappa) - L^-(\mathbf{x}^\kappa)}{2} d\mathbf{x}, \quad (27)$$

$$F_{q^*}^y(\kappa) = \sum_{\substack{\kappa^- \in \mathcal{K} \\ \kappa^- \cap \kappa^+ \neq \emptyset}} \int_{\partial\kappa} n_y \frac{L^+(\mathbf{x}^\kappa) - L^-(\mathbf{x}^\kappa)}{2} d\mathbf{x}, \quad (28)$$

$$F_{q^*}^\tau = \sum_{\substack{\kappa^- \in \mathcal{K} \\ \kappa^- \cap \kappa^+ \neq \emptyset}} \int_{\partial\kappa} \mathbf{n}^- L^-(\mathbf{x}^\kappa) - \mathbf{n}^+ L^+(\mathbf{x}^\kappa) d\mathbf{x}. \quad (29)$$

Introducing global data vectors $\tilde{\mathbf{q}}^x$, $\tilde{\mathbf{q}}^y$, and $\tilde{\mathbf{u}}$ and summing the weak problem (11) over all elements $\kappa \in \mathcal{K}$, we arrive at the following

$$-S^x \tilde{\mathbf{q}}^x - S^y \tilde{\mathbf{q}}^y - \omega^2 M \tilde{\mathbf{u}} = M \mathbf{f} + F_{q^*}^x \tilde{\mathbf{q}}^x + F_{q^*}^y \tilde{\mathbf{q}}^y - \tau F_{q^*}^\tau \tilde{\mathbf{u}}, \quad (30)$$

$$M \tilde{\mathbf{q}}^x - S^x \tilde{\mathbf{u}} = F_{u^*}^x \tilde{\mathbf{u}}, \quad (31)$$

$$M \tilde{\mathbf{q}}^y - S^y \tilde{\mathbf{u}} = F_{u^*}^y \tilde{\mathbf{u}}. \quad (32)$$

Solving for the slack variable $\tilde{\mathbf{q}}^{x,y}$ in equations (31) and (32), and substituting into (30) eliminates the dependence on $\tilde{\mathbf{q}}$. The system, written in compact form is then

$$(-S + F - \omega^2 M) \tilde{\mathbf{u}} = M \mathbf{f}, \quad (33)$$

where

$$S = S^x M^{-1} S^x + S^y M^{-1} S^y \quad (34)$$

$$F = F_{q^*}^x M^{-1} S^x + F_{q^*}^x M^{-1} F_{u^*}^x + F_{q^*}^y M^{-1} S^y + F_{q^*}^y M^{-1} F_{u^*}^y - \tau F_{q^*}^\tau. \quad (35)$$

The operator S is clearly negative semi-definite, while for $\tau > 0.0$, the composite operator $S - F$ is strictly negative definite. A full eigenspectrum analysis and the impact on the the preconditioner is unknown. However, it suffices to say that for small ω we should expect a well-conditioned system, while for moderate ω , indefinite and near singular matrices should be expected.

3 Preconditioning

As detailed in the previous section, while the structure of the algebraic system has many valuable properties such as modest sparsity, symmetry, and being nearly block structured, there are several less advantageous aspects that pose a challenge to common solution techniques. For a fixed problem— i.e. selecting a frequency ω —the resulting system, based on realistic n and h that adequately resolve the solution, is inevitably not definite. The number of eigenvalues of opposite sign, of course, heavily depend on the parameters of the problem. Extensive work by Cai et al. [7,5,6] and Elman [13] conclude that standard Conjugate Gradient based iterative methods handle a moderate number of flipped eigenvalues quite well. We will also use this class of methods and will, in particular, choose the Generalized Minimum Residual method (GMRES) due to it’s robustness and wide popularity [17]. The GMRES method can be applied to indefinite systems and, more importantly, the preconditioned implementation permits indefinite preconditioning matrices. This will be beneficial in the case of the preconditioned additive Schwarz (PAS) method.

Looking more closely at the eigenspectrum exposes a potentially more significant quandary. Recalling the form of the algebraic system (33)

$$\underbrace{-S + F}_{\text{positive}} - \underbrace{\omega^2 M}_{\text{positive}} \tag{36}$$

illustrates that the smallest eigenvalue in magnitude can be arbitrarily close to or equal to zero. Avoiding the special case (and unlucky selection of n and h) when the system is near singular, we are nonetheless left with a possibly large condition number. In this respect, we are forced to consider a range of values in order to avoid anomalies in the approximation parameters. This will be discussed in more detail in the next section.

In the remainder of this section we introduce a *hybrid* additive Schwarz type preconditioner with overlap. Our implementation is a culmination of approaches, which includes a coarse grid solution phase with the ability to handle non-nested coarse grids. This approach fits particularly well with our goal of a fully adaptive *hp* implementation and since nested coarse grids are not always accessible, this adds a layer of flexibility in coarsening tessellations of complicated geometries.

We now turn to the PAS method for

$$A\mathbf{u} = \mathbf{f}. \tag{37}$$

In short, our algorithm calculates solutions to successive ‘subproblems’ and the composite approximation is then realized after a full sweep. The underlying strategy is reminiscent of a block Jacobi sweep, although additive Schwarz

incorporates a number of additional features. We opt for the additive approach as parallelization is generally straightforward. The criteria for selecting subdomains is dependent on the problem and while many applications need only a few levels of h -type refinement, we will focus on situation where complicated domains call for large numbers of elements in mesh. Moreover, we will also focus on regimes where the spectral order is 8 or less since this is a typical range for this solution method in wave-type problems. With these two heuristics in mind, we will describe our subdomains Ω_s as the union of (non-overlapping) $\kappa \in \mathcal{K}$. Computationally, we will further restrict Ω_s to be a connected set, although this is not necessary in general. We define the restriction operator R_s that maps global degrees of freedom in Ω to Ω_s by way of injection. Similarly, the prolongation operator, P_s , is defined to be the transpose: $P_s = R_s^T$. These transfer operators are very sparse having only $|\Omega_s|$ nonzero entries. Using the so-called Galerkin Condition

$$A_s = R_s A R_s^T \quad (38)$$

to define the matrix problem on a given subdomain, we arrive at the classic Additive Schwarz iteration: Given an approximation \mathbf{u}^{old} , a correction can be computed on subdomain Ω_s using the error equation $Ae = r$, where $r = b - Au^{\text{old}}$ is the residual. Sweeping through all subdomains Ω_s for $s = 1 \dots S$, we have

$$\mathbf{u}^{\text{new}} = \mathbf{u}^{\text{old}} + \sum_{s=1}^S R_s^T A_s^{-1} R_s r^{\text{old}}. \quad (39)$$

The preconditioning matrix, M , can be realized if we view (39) as a fixed-point iteration for the preconditioned matrix problem $M^{-1}A\mathbf{u} = M^{-1}\mathbf{f}$ [38]. Writing (39) as

$$\mathbf{u}^{\text{new}} = \left(I - \sum_{s=1}^S R_s^T A_s^{-1} R_s A \right) \mathbf{u}^{\text{old}} + \sum_{s=1}^S R_s^T A_s^{-1} R_s \mathbf{f}. \quad (40)$$

Since $I - M^{-1}A$ is the corresponding iteration matrix in the fixed-point iteration, we have that

$$M^{-1} = \sum_{s=1}^S R_s^T A_s^{-1} R_s. \quad (41)$$

Along with the basic PAS algorithm, presented in (41) in matrix form, we introduce two additional steps that will improve the convergence of the overall (hybrid) method: coarse level solves using non-nested coarse grids and element overlap into neighboring subdomains.

For coarse level solves, we utilize a two-level approach popularized by Dryja and Widlund [12] and with features that are inherent in multigrid algorithms [39]. The current level of mesh refinement, Ω^h , may not have an obvious or natural coarse grid, particularly after several refinement/derefinement levels

have been processed. However, using an approach similar to the work of Chan et al. [8], we can define transfer operators suitable for *any* coarse grid that falls under the same restrictions as the initial tessellation, outlined at the beginning of Section 2. Given a coarse grid tessellation, Ω^H , and subdomain $\Omega_s^h \subset \Omega^h$, we define the restriction operator based on standard finite element interpolation as

$$R_{0_{ij}} = \phi_i(\mathbf{x}_j). \quad (42)$$

Here, $\phi_i(\mathbf{x})$ is a coarse grid finite element bilinear basis function and \mathbf{x}_j is a node in Ω_s on the fine grid. $R_{0_{ij}} = 0$ if \mathbf{x}_j is not in the underlying footprint of ϕ_i and is thus still sparse, although not in comparison to the injection operators above. In order to ensure proper interpolation of constant solutions, we incorporate a row equilibration technique, by rescaling each row of R_0 by the row sum:

$$R_{0_{ij}} \leftarrow \frac{1}{\sum_j R_{0_{ij}}} R_{0_{ij}}. \quad (43)$$

This improves convergence results drastically and is intuitively sensible for non-nested grids since interpolation from a coarse degree of freedom in a coarse grid element κ^H to a fine degree of freedom in a fine grid element κ^h is given proper weighting that depends on the overlap of κ^H and κ^h . Small overlap will have a correspondingly small impact while larger overlap will also be scaled appropriately. As with the subdomain solves, our results also indicate that the Galerkin form of the coarse grid operator A_O is indeed necessary to achieve an efficient algorithm. We define A_0 by

$$A_0 = R_0 A R_T, \quad (44)$$

instead of building the operator A^H directly on Ω^H . The composite preconditioning matrix is then defined to be

$$M^{-1} = R_0^T A_0^{-1} R_0 + \sum_{s=1}^S R_s^T A_s^{-1} R_s. \quad (45)$$

We also present numerical results in the next section for the case when the coarse grid mesh and the fine grid mesh are equal: $\Omega^H = \Omega^h$. Here, bilinear finite elements for the coarse problem are still used, reminiscent of the low-order finite element preconditioning outlined in [15]. It is important to note that a weighted interpolation strategy is needed in this case—e.g. standard finite element interpolation—even though an injection operator is more straightforward. Computation shows that injected coarse grid solves lead to little improvement in the approximation and can pollute the system even for the most basic problems. This indicates that too much information is lost in the injection approach and for this reason, the results present all use bilinear finite element interpolation/restriction operators.

Overlap is also introduced in our algorithm. This increases communication, but, as we present in the next section, is an essential component particularly

for high-order approximations and as the matrix increases in indefiniteness and size. We define $\delta = 0$ to be the case where $\Omega_{s_1}^\circ \cap \Omega_{s_2} = \emptyset$ when $\bar{\Omega}_{s_1} \cap \Omega_{s_2} \neq \emptyset$. Even so, due to the nature of the discontinuous discretization, where degrees of freedom in neighboring elements may share a geometric location, there is a slight inheritance of overlap. The process of avoiding overlap in a DG formulation is discussed in more detail in [30] for low order elements. By increasing δ , we simply mean that each subdomain is padded by δ layers of elements. An example case with three subdomains is presented in Figure 3. At first glance, this may seem extreme, since Fischer and Lottes [31] extend only by strips of nodes into the adjacent elements. However, the class of problems we address is altogether different, requiring large numbers of elements, and requiring only moderate polynomial degrees, making overlap overhead costs small as the mesh is further refined. Moreover, layers of nodes within an electrostatic distribution (Figure 1) are not readily available either in the element itself or in the reference element, where they have a straightforward formation in the case of tensor-based element.




Fig. 3. 3 subdomains Ω_s with and without overlap.

4 Numerical Preliminaries

We now explore the numerical efficiency of the preconditioned additive Schwarz Algorithm described in the previous section. In order for the method to be of practical value, it should exhibit robustness with respect to a variety of challenging properties. Namely, indefiniteness of the algebraic system, varying stencil sizes, and a range of grid sizes. We expose these difficulties by varying the frequency, ω , in the Helmholtz equation (1), the polynomial order, n , of the spectral elements, and the minimum diameter, h , of an element in the tessellation \mathcal{K}_h .

Our test problem is basic, yet still presents a level of generality. It is defined as follows:

Example 3

$$-\nabla \cdot \nabla u(\mathbf{x}) - \omega^2 u(\mathbf{x}) = f(\mathbf{x}) \quad \text{in } \Omega, \tag{46a}$$

$$u(\mathbf{x}) = g(\mathbf{x}) \quad \text{on } \Gamma, \tag{46b}$$

$$f(\mathbf{x}) = \left(2(2\pi)^2 - \omega^2\right) \sin(2\pi x) \sin(2\pi y) \tag{46c}$$

$$g(\mathbf{x}) = 0, \tag{46d}$$

$$u_{\text{exact}} = \sin(2\pi x) \sin(2\pi y) \tag{46e}$$

$$\Omega = [0, 1] \times [0, 1] \tag{46f}$$

Remark 4 Notice that $f(\mathbf{x}) \in C^\infty(\Omega) \subset L^2(\Omega)$ and that $u_{\text{exact}} \in C_0^\infty(\Omega) \subset H^2(\Omega)$.

We begin by looking at well conditioned example with only slight indefiniteness—i.e. only a few eigenvalues have flipped sign. Comparing the iterations in Table 2 indicates that a coarse grid is beneficial for high-order discretizations. Here, h_f is the average diameter of elements on the fine grid, while h_c represents the average diameter of the elements on the coarse grid. Dryja and Widlund [12] introduced a coarse grid correction as a means of reducing global error components and the approach parallels the strategy found in classical multigrid schemes [39]. Generalized Minimum Residual (GMRES) iterations [38] (with no restarts) are reduced for each polynomial order when using a richer coarse grid. This trend also continues for larger ω when the system becomes highly indefinite as in Figure 4. The growth in iterations can be attributed to an inaccurate coarse grid problem. As the frequency ω increases, more degrees of freedom are needed to fully resolve the solution. When the problem is viewed on a coarser grid, the discretization lacks resolution and the solution found on the coarse grid no longer resembles an accurate approximation to the fine grid solution. Thus the two-level error correction becomes ineffective and possibly pollutes the fine grid solution. For smaller ω , the coarse grid problem is still adequately resolved and, as Figure 4 confirms, the coarse grid correction component improves the preconditioning for all n .

	order n										
	1	2	3	4	5	6	7	8	9	10	11
$h_c \approx 0$	26	38	49	60	71	82	93	105	116	128	140
$h_c \approx 1/4$	22	32	39	50	58	67	72	81	88	100	108
$h_c \approx 1/8$	14	25	30	36	43	47	55	60	66	73	79

Table 2
GMRES iterations with $h_f \approx 1/8$, $\omega = 1.0$, and $\delta = 0.0$




Fig. 4. GMRES iterations for $\omega = 0, 1, 10$, and 50 (top curve) with no overlap. Left figure: $h_c \approx 1/4$. Right figure: $h_c \approx 1/8$.

Overlap has a larger impact on the convergence of the preconditioned iterative method as indicated in Table 3. Iteration counts are greatly reduced when one layer of information is shared between subdomains. The dependence on n is also improved as the iterations increase only slightly.

	order n										
	1	2	3	4	5	6	7	8	9	10	11
$\delta = 0$	22	32	39	50	58	67	72	81	88	100	108
$\delta = 1$	22	22	23	24	24	25	25	26	26	27	28

Table 3
GMRES iterations for $h_f \approx 1/8$, $h_c \approx 1/4$, and $\omega = 1.0$

As ω increases, the system becomes more indefinite. Figure 5 shows that the iteration counts remain bounded as the polynomial order is increased for each the selected ω . The iterations increase as the frequency is increased, but this is expected as more low eigenvalues are shifted to the positive half-plane. It is also apparent that a level of overlap has more affect than simply enriching the global coarse grid solution. The iterations in the top two plots are only slightly improved in the bottom two plots when a richer coarse grid is used. However, there is significant improvement as we introduce overlap, particularly for the case of the highly indefinite problem, $\omega = 50$.

A more reasonable and methodical test is to consider problem parameters that are typically encountered computationally. In other words, a more definitive exposé is to test problems where the discretization is not under or over resolved. Referring to dispersion analysis, using around 10 degrees of freedom per wavelength (in 1-D) is generally considered well resolved. Thus, we would




Fig. 5. GMRES iterations versus polynomial order: Comparing overlap with and without a global coarse grid solve.

need to choose a mesh and polynomial order according to

$$\frac{2\pi(n+1)}{\omega h} > 10 \quad (47)$$

$$\Rightarrow \frac{5\omega}{\pi} < \frac{(n+1)}{h}. \quad (48)$$

A summary is given in Table 4.

	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$	$n = 6$	$n = 7$
h	$\leq \omega \leq$	$\leq \omega \leq$	$\leq \omega \leq$	$\leq \omega \leq$	$\leq \omega \leq$	$\leq \omega \leq$	$\leq \omega \leq$
$\frac{1}{2}$	0 4	3 5	4 7	6 8	7 10	9 12	11 13
$\frac{1}{4}$	0 7	6 10	9 13	12 16	15 20	19 23	22 26
$\frac{1}{8}$	0 13	12 20	19 26	25 32	31 39	38 45	44 51
$\frac{1}{16}$	0 26	25 39	38 51	50 64	63 77	76 90	89 102
$\frac{1}{32}$	0 51	50 77	76 102	101 128	127 153	152 179	178 204

Table 4

Limits of realistic ω for select mesh sizes and polynomial orders.

In Table 5 we present GMRES iterations with no preconditioning. The work involved is excessive and increases as the range of ω is increased, since the system becomes more indefinite. As we vary ω by one, the iterations vary, but remain within a certain range, indicating that we have not selected a fortuitous set of parameters nor are we considering a *near* singular system. Comparing these values with Table 6, we see that preconditioning improves the iteration counts, particularly when one layer of overlap is introduced ($\delta = 1$). The introduction of coarse grids, however, does not play a significant role until the coarse grid is fine enough. Tables 7 and 8 show slight improvement when a coarse solution is included and the density is increased for problems with low wave numbers. This is illustrated by the first three rows of Tables 6, 7, and 8, where the system includes only a few negative eigenvalues and where the coarse grid problem is able to adequately resolve the solution. Difficulties arise for higher ω when a higher percentage of eigenvalues have changed sign and the coarse grid problem does not have enough degrees of freedom for meaningful resolution of the high frequency waves. Comparing the last three or four rows of Tables 6, 7, and 8, we see no improvement and some pollution when increasing the size of the coarse grid correction.

Recalling coarse-grid correction from a multigrid viewpoint, we introduce coarse levels in order to eliminate global modes not reduced by the local subdomain solves on the fine-grid. As we transfer the problem to a coarse grid, the coarse grid problem will have an even higher percentage of indefinite eigenmodes, making effective coarse-grid correction more elusive. This is precisely what we observe numerically. As the wave number increases, a *fixed* coarse grid mesh is not rich enough to resolve the wave problem on the coarse grid. As the coarse grid is enriched from Table 7 to Table 8, we are able to effectively precondition wider range of wave numbers.

ω	iterations								avg.
0 ... 7	42	42	43	44	49	52	52	63	48.4
6 ... 10				89	100	106	112	125	106.4
9 ... 13				155	176	156	166	197	170.0
12 ... 16				234	260	286	253	321	270.8
15 ... 20			333	416	378	393	395	439	392.3
19 ... 23				490	532	512	564	571	533.8
22 ... 26				676	690	697	683	780	705.2

Table 5
GMRES iterations: $h_f \approx 1/8$, no preconditioning

ω	iterations								avg.	
$\delta = 0$	0 ... 7	20	20	21	22	23	25	26	29	23.2
	6 ... 10				37	47	39	43	48	42.8
	9 ... 13				55	57	53	55	66	57.2
	12 ... 16				66	78	74	57	84	71.8
	15 ... 20		97		117	102	98	106	113	105.5
	19 ... 23				142	152	148	154	159	151.0
	22 ... 26				191	194	192	190	200	193.4
ω	iterations								avg.	
$\delta = 1$	0 ... 7	17	17	17	17	17	18	19	24	18.2
	6 ... 10				20	28	25	27	33	26.6
	9 ... 13				28	32	27	30	35	30.4
	12 ... 16				30	36	38	32	43	35.8
	15 ... 20		40		49	51	45	47	56	48.0
	19 ... 23				62	66	71	67	70	67.2
	22 ... 26				64	70	70	76	79	71.8

Table 6
GMRES iterations: $h_f \approx 1/8$, no coarse grid

5 Concluding Remarks

In this paper, we have highlighted a number of important algorithmic details for effectively preconditioning high-order nodal discontinuous Galerkin discretizations of the indefinite Helmholtz operator. The approach utilizes the robustness of an overlapping Schwarz-based preconditioner to handle the localized nature of the problem and increases in polynomial order while relying on the proven success of Krylov acceleration methods to efficiently condition indefinite modes. The results emphasize the importance of sufficient overlap and the delicate issue of global coarse grid correction. If wave numbers are high and the fine grid problem are fully resolved, then adequate resolution on the coarse grid is necessary for improved preconditioning and to avoid pollution from the two-level approach.

This research has also introduced a number of possibilities for further development. The various DG formulations give rise to matrix problems with a wide range of condition numbers and stencil width. This is currently not accounted for in most preconditioning strategies. More proficient algorithms are presumably achievable if the stability parameters and element boundary transmission

ω	iterations								avg.
0 ... 7	21	21	21	22	23	24	24	29	23.1
6 ... 10				36	46	41	44	51	43.6
9 ... 13				56	62	57	59	71	61.0
$\delta = 0$ 12 ... 16				70	81	78	72	93	78.8
15 ... 20			95	113	104	102	107	116	106.2
19 ... 23				142	152	144	151	158	149.4
22 ... 26				185	191	190	190	199	191.0
ω	iterations								avg.
0 ... 7	20	20	20	20	21	22	22	31	22.0
6 ... 10				26	34	30	31	37	31.6
9 ... 13				32	42	33	36	43	37.2
$\delta = 1$ 12 ... 16				36	42	45	42	54	43.8
15 ... 20			46	56	58	52	56	67	55.8
19 ... 23				70	79	76	75	79	75.8
22 ... 26				69	75	73	78	82	75.4

Table 7

GMRES iterations: $h_f \approx 1/8$, $h_c \approx 1/4$

conditions in these methods are addressed more directly. Finally, the encouraging results for the Helmholtz problem motivates the study of more complex wave problems such as Maxwell's equation.

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 (2001/02), pp. 1749–1779 (electronic).
- [2] S. BEUCLER, *Multigrid solver for the inner problem in domain decomposition methods for P-FEM*, SIAM J. Numer. Anal., 40 (2002), pp. 928–944 (electronic).
- [3] D. BRAESS, *Finite elements*, Cambridge University Press, Cambridge, second ed., 2001. Theory, fast solvers, and applications in solid mechanics, Translated from the 1992 German edition by Larry L. Schumaker.
- [4] F. BREZZI, G. MANZINI, D. MARINI, P. PIETRA, AND A. RUSSO, *Discontinuous Galerkin approximations for elliptic problems*, Numer. Methods Partial Differential Equations, 16 (2000), pp. 365–378.

ω	iterations								avg.	
$\delta = 0$	0 ... 7	12	13	13	13	13	13	14	14	13.1
	6 ... 10				26	31	31	38	45	34.2
	9 ... 13				43	57	43	43	61	49.4
	12 ... 16				49	66	68	63	97	68.6
	15 ... 20		83		110	109	104	112	122	106.7
	19 ... 23				140	152	144	156	164	151.2
	22 ... 26				172	186	191	195	208	190.4
ω	iterations								avg.	
$\delta = 1$	0 ... 7	13	13	13	14	14	15	17	17	14.5
	6 ... 10				23	34	24	30	35	29.2
	9 ... 13				33	37	27	35	45	35.4
	12 ... 16				36	44	43	42	54	43.8
	15 ... 20		46		62	60	55	62	66	58.5
	19 ... 23				71	80	86	82	83	80.4
	22 ... 26				77	76	77	84	81	79.0

Table 8

GMRES iterations: $h_f \approx 1/8$, $h_c \approx 1/8$

- [5] X.-C. CAI, *A family of overlapping Schwarz algorithms for nonsymmetric and indefinite elliptic problems*, in Domain-based parallelism and problem decomposition methods in computational science and engineering, SIAM, Philadelphia, PA, 1995, pp. 1–19.
- [6] X.-C. CAI, M. A. CASARIN, F. W. ELLIOTT, JR., AND O. B. WIDLUND, *Overlapping Schwarz algorithms for solving Helmholtz’s equation*, in Domain decomposition methods, 10 (Boulder, CO, 1997), vol. 218 of Contemp. Math., Amer. Math. Soc., Providence, RI, 1998, pp. 391–399.
- [7] X.-C. CAI AND O. B. WIDLUND, *Domain decomposition algorithms for indefinite elliptic problems*, SIAM J. Sci. Statist. Comput., 13 (1992), pp. 243–258.
- [8] T. F. CHAN, B. F. SMITH, AND J. ZOU, *Overlapping Schwarz methods on unstructured meshes using non-matching coarse grids*, Numer. Math., 73 (1996), pp. 149–167.
- [9] B. COCKBURN, G. E. KARNIADAKIS, AND C.-W. SHU, *The development of discontinuous Galerkin methods*, in Discontinuous Galerkin methods (Newport, RI, 1999), vol. 11 of Lect. Notes Comput. Sci. Eng., Springer, Berlin, 2000, pp. 3–50.

- [10] 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 (electronic).
- [11] M. O. DEVILLE, P. F. FISCHER, AND E. H. MUND, *High-order methods for incompressible fluid flow*, vol. 9 of Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, 2002.
- [12] M. DRYJA AND O. B. WIDLUND, *Some domain decomposition algorithms for elliptic problems*, in Iterative methods for large linear systems (Austin, TX, 1988), Academic Press, Boston, MA, 1990, pp. 273–291.
- [13] H. C. ELMAN, O. G. ERNST, AND D. P. O’LEARY, *A multigrid method enhanced by Krylov subspace iteration for discrete Helmholtz equations*, SIAM J. Sci. Comput., 23 (2001), pp. 1291–1315 (electronic).
- [14] X. FENG AND O. A. KARAKASHIAN, *Two-level additive Schwarz methods for a discontinuous Galerkin approximation of second order elliptic problems*, SIAM J. Numer. Anal., 39 (2001), pp. 1343–1365 (electronic).
- [15] P. F. FISCHER, *An overlapping Schwarz method for spectral element solution of the incompressible Navier-Stokes equations*, J. Comput. Phys., 133 (1997), pp. 84–101.
- [16] D. GOTTLIEB AND J. S. HESTHAVEN, *spectral methods for hyperbolic problems*, j. comput. appl. math., 128 (2001), pp. 83–131. numerical analysis 2000, vol. vii, partial differential equations.
- [17] A. GREENBAUM, *Iterative methods for solving linear systems*, vol. 17 of Frontiers in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997.
- [18] 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 (electronic).
- [19] ———, *High-order accurate methods in time-domain computational electrodynamics*, Review 17, Brown University, 2002.
- [20] 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.
- [21] ———, *High-order nodal discontinuous Galerkin methods for the Maxwell eigenvalue problem*, Phil. Trans. R. Soc. Lond., 362 (2004), pp. 493–524.
- [22] J. HEYS, T. MANTEUFFEL, S. MCCORMICK, AND L. OLSON, *Algebraic multigrid (amg) preconditioning for high-order finite elements*, tech. rep., University of Colorado at Boulder, June 2004. submitted.
- [23] N. HU, X.-Z. GUO, AND I. N. KATZ, *Multi-p preconditioners*, SIAM J. Sci. Comput., 18 (1997), pp. 1676–1697.

- [24] G. KANSCHAT, *Preconditioning methods for local discontinuous Galerkin discretizations*, SIAM J. Sci. Comput., 25 (2003), pp. 815–831 (electronic).
- [25] G. E. KARNIADAKIS AND S. J. SHERWIN, *Spectral/hp element methods for CFD*, Numerical Mathematics and Scientific Computation, Oxford University Press, New York, 1999.
- [26] R. M. KIRBY, *Toward dynamic spectral/hp refinement: algorithms and applications to flow-structure interactions*, PhD thesis, Brown University, May 2003.
- [27] D. A. KOPRIVA, S. L. WOODRUFF, AND M. Y. HUSSAINI, *Discontinuous spectral element approximation of Maxwell's equations*, in Discontinuous Galerkin methods (Newport, RI, 1999), vol. 11 of Lect. Notes Comput. Sci. Eng., Springer, Berlin, 2000, pp. 355–361.
- [28] C. LASSER AND A. TOSELLI, *Convergence of some two-level overlapping domain decomposition preconditioners with smoothed aggregation coarse spaces*, in Recent developments in domain decomposition methods (Zürich, 2001), vol. 23 of Lect. Notes Comput. Sci. Eng., Springer, Berlin, 2002, pp. 95–117.
- [29] C. LASSER AND A. TOSELLI, *Overlapping preconditioners for discontinuous Galerkin approximations of second order problems*, in Domain decomposition methods in science and engineering (Lyon, 2000), Theory Eng. Appl. Comput. Methods, Internat. Center Numer. Methods Eng. (CIMNE), Barcelona, 2002, pp. 78–84.
- [30] 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 (electronic).
- [31] J. W. LOTTES AND P. F. FISCHER, *Hybrid multigrid/schwarz algorithms for the spectral element method*, Tech. Rep. ANL/MCS-P1052-0403, Mathematics and Computer Science Division, Argonne National Laboratory, April 2003.
- [32] 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.
- [33] L. F. PAVARINO AND T. WARBURTON, *Overlapping Schwarz methods for unstructured spectral elements*, J. Comput. Phys., 160 (2000), pp. 298–317.
- [34] I. PERUGIA AND D. SCHÖTZAU, *The hp-local discontinuous Galerkin method for low-frequency time-harmonic Maxwell equations*, Math. Comp., 72 (2003), pp. 1179–1214 (electronic).
- [35] G. PIEPER, *New perspectives for spectral and high-order methods*, SIAM Newsletter, 37 (2004).
- [36] E. M. RØNQUIST AND A. T. PATERA, *Spectral element multigrid. I. Formulation and numerical results*, J. Sci. Comput., 2 (1987), pp. 389–406.

- [37] J. RUGE, *Amg for higher-order discretizations of second-order elliptic problems*, in Abstracts of the Eleventh Copper Mountain Conference on Multigrid Methods, 2003.
- [38] Y. SAAD, *Iterative methods for sparse linear systems*, Society for Industrial and Applied Mathematics, Philadelphia, PA, second ed., 2003.
- [39] U. TROTTEBERG, C. W. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Academic Press Inc., San Diego, CA, 2001. With contributions by A. Brandt, P. Oswald and K. Stüben.