

Oden Institute REPORT 20-17

October 2020

An Adaptive Multigrid Solver for DPG Methods with Applications in Linear Acoustics and Electromagnetics

by

Socratis Petrides and Leszek Demkowicz



Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin
Austin, Texas 78712

Reference: Socratis Petrides and Leszek Demkowicz, "An Adaptive Multigrid Solver for DPG Methods with Applications in Linear Acoustics and Electromagnetics," Oden Institute REPORT 20-17, Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, October 2020.

AN ADAPTIVE MULTIGRID SOLVER FOR DPG METHODS WITH APPLICATIONS IN LINEAR ACOUSTICS AND ELECTROMAGNETICS

SOCRATIS PETRIDES¹ AND LESZEK DEMKOWICZ²

¹*Center for Applied Scientific Computing, Lawrence Livermore National Laboratory*

²*Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin*

ABSTRACT. We propose an adaptive multigrid preconditioning technology for solving linear systems arising from Discontinuous Petrov-Galerkin (DPG) discretizations. Unlike standard multigrid techniques, this preconditioner involves only trace spaces defined on the mesh skeleton, and it is suitable for adaptive hp-meshes. The key point of the construction is the integration of the iterative solver with a fully automatic and reliable mesh refinement process provided by the DPG technology. The efficacy of the solution technique is showcased with numerous examples of linear acoustics and electromagnetic simulations, including simulations in the high-frequency regime, problems which otherwise would be intractable. Finally, we analyze the one-level preconditioner (smoother) for uniform meshes and we demonstrate that theoretical estimates of the condition number of the preconditioned linear system can be derived based on well established theory for self-adjoint positive definite operators.

1. INTRODUCTION

1.1. Background. Numerical simulation of wave propagation phenomena has been a very challenging problem in the computational science community for many years. The main difficulty is the following. Most commonly used numerical methods for PDEs, such as the Finite Element or the Finite Difference methods lose their stability, unless the wave frequency is “resolved”. That is, the underlying discretization mesh has to be “fine enough” in order to satisfy the *Nyquist criterion* and overcome the so called *pollution effect*. The Nyquist criterion requires a minimum number of finite elements or discretization points per wavelength for the method to remain stable. Moreover, as the frequency increases, even if the Nyquist criterion is met, the solution gets “polluted”, i.e., becomes out of phase or very diffusive. Pollution can be minimized by using high order discretizations, but it can not be eliminated. To the best of our knowledge, to this day there is no *pollution-free* method for simulations in two or three space dimensions [4].

Fine meshes and high order discretizations result in very large linear systems and therefore extremely hard to solve. State of the art direct solvers based on LU decompositions for sparse matrices fail to scale efficiently and they usually run out of memory. While for two dimensional problems multi-frontal direct solvers for sparse matrices can be optimal, i.e., for a linear system of size N the complexity and memory estimates are $\mathcal{O}(N^{3/2})$ and $\mathcal{O}(N \log N)$ respectively, in three space dimensions this is not the case. For a 3D problem on a uniform mesh, the computational cost is estimated to be $\mathcal{O}(N^2)$ and the required memory is $\mathcal{O}(N^{4/3})$ [39].

An alternative approach is to employ iterative solution techniques. This kind of solvers do not need to perform any kind of factorization. On the contrary, they are based on recursive matrix–vector multiplications, operations that have both complexity and memory requirements of $\mathcal{O}(N)$. Usually, the optimal strategy is to construct a “good” approximation of the inverse of the system matrix and use it as a preconditioner to a Krylov subspace method such as the Generalized Minimum Residual (GMRES) or the Conjugate Gradient (CG) method. Unfortunately, constructing an optimal preconditioner for time–harmonic wave simulations is still a notoriously hard problem. The time–harmonic wave operator (acoustics or Maxwell) is highly indefinite and cutting edge

E-mail address: petrides1@llnl.gov, leszek@oden.utexas.edu.

preconditioners for elliptic problems (such as multilevel, multigrid, DDM) fail to work. A common strategy for preconditioning indefinite operators involve the construction of a preconditioner for GMRES which shifts the spectrum to the positive half plane. The resulting solver is not always reliable, especially when the grid is very coarse. In addition, multigrid techniques lose their efficiency, since in order to converge, they usually require expensive coarse grid solves because the coarse mesh has to be fine enough to maintain discrete stability.

Several recent approaches show more promising results [25]. In particular well known attempts include but not limited to, shifted Laplacian techniques [20, 46, 24], domain decomposition methods [36, 27, 8, 47], multigrid methods [48, 28, 9] and stabilized FEM techniques based on artificial absorption [7, 23]. A lot of attention has been paid to the sweeping preconditioner [19, 51] and some of its variants such as the polarized traces [55], the L-sweeps method [49] and the Source transfer method [12, 37]. These are domain decomposition methods which introduce absorbing boundary conditions in the interior of the domain (often a PML) to transfer the information among the subdomains. Although it is hard to exploit parallelization, they are provably very optimal for problems in unbounded domains. However, their efficiency can significantly deteriorate for quasi-resonant problems.

1.2. Preconditioning the DPG system. There are several works on preconditioning DPG systems for elliptic problems. To the best of our knowledge the first one chronologically, was done by Barker et al. in [5]. In this work the authors analyzed the one level additive Schwarz preconditioner for the Poisson problem and showcased, both theoretically and numerically, convergence of the preconditioned CG solver, independent of the mesh size h . A few years later, Barker et al. in [6] constructed a scalable preconditioner for the primal DPG method again for the Poisson problem. The key point of the implementation was the norm equivalence of the DPG bilinear form with the standard H^1 and $H(\text{div})$ norms. Naturally, existing H^1 and $H(\text{div})$ algebraic multigrid (AMG) technologies could be utilized to precondition the DPG system. Following a similar approach in [38], the authors extended theoretical analysis for the Poisson problem to the two level setting. Their work was based on well known results on preconditioning the H^1 , $H(\text{div})$ and $H(\text{curl})$ spaces [2, 18, 33]. Finally, a geometric multigrid preconditioner for the Poisson and Stokes problem is presented in [45].

Designing a preconditioner for wave problems is much more challenging. While, it is relatively straight forward to construct and analyze preconditioners for elliptic problems using the idea of norm equivalence, for wave problems the equivalence constants are frequency dependent. As far as we know there are two works so far attempting to construct robust and efficient preconditioners for wave problems but none of them provide any theoretical results. In [38] the authors present a numerical study on a one level additive Schwarz preconditioner when applied to the Helmholtz problem. Their results indicate uniform convergence of the solver with respect to the frequency ω and the mesh size h . The convergence is however sensitive to the overlap and the size of the Schwarz patches. The second attempt on preconditioning the Helmholtz problem is described in [30]. This work introduces a multiplicative Schwarz (overlapping block Gauss–Seidel) preconditioner for the primal DPG formulation. The preconditioner is shown to converge at a rate independent of the polynomial order p and the frequency ω , but no dependence study on the mesh size h is provided.

For our construction we take a slightly different path than most of the works described above. The preconditioner is constructed directly from the underlying bilinear form of the wave problem. Using the properties of the ultraweak DPG method we invoke a numerical experiment to examine the dependence of the condition number of the preconditioned system with respect to the polynomial order of approximation, the frequency, the Schwarz patch size and the mesh size. Even though the analysis is done only for a one level additive smoother, it gives us useful insights for the construction of a multilevel preconditioner. We note that our results are consistent with the results of related work described above.

1.3. Outline. The paper is organized as follows. First, we give a brief summary of the DPG methods, using as a model problem the linear acoustics equations. Moving on, in Section 3, we provide some theoretical results on the one-level preconditioner (smoother) for these equations. The analysis is complemented with a numerical experiment that is described in detail in Appendix A. In Section 4, we present the multigrid algorithm and in Section 5 we showcase with numerous numerical results the efficacy of this method for both acoustics and electromagnetic simulations. Lastly, in Section 6, we summarize our conclusions and outline future work.

2. THE DPG METHOD AND LINEAR ACOUSTICS

The DPG method for the time-harmonic linear acoustics equations was analyzed in depth in [56, 16]. In this section we briefly mention the DPG setup and provide some background for the analysis of the one-level additive Schwarz smoother for preconditioning the ultraweak DPG formation of the time-harmonic acoustics equations.

2.1. The ideal Petrov–Galerkin method: a brief outline. We consider the following abstract setting. Given a continuous bilinear (sesquilinear) form $b(\cdot, \cdot)$ defined on the product $U \times V$ of Hilbert spaces U (trial space) and V (test space), and a continuous linear (anti-linear) form $l(\cdot)$ defined on V , we want to find the solution to the problem:

$$\begin{cases} u \in U \\ b(u, v) = l(v), \quad v \in V. \end{cases}$$

We assume that the above continuous problem is well posed, i.e., the continuous inf–sup condition [3] for the bilinear form $b(\cdot, \cdot)$ holds. Given now a finite-dimensional trial space $U_h \subset U$, we can recast the above linear problem as a minimization problem in the dual norm $\|\cdot\|_{V'}$, i.e., we now seek the solution to

$$u_h = \arg \min_{w_h \in U_h} \frac{1}{2} \|l - Bw_h\|_{V'}^2,$$

where $B : U \rightarrow V'$ is defined as

$$\langle Bu, v \rangle_{V' \times V} = b(u, v)$$

and $\langle \cdot, \cdot \rangle_{V' \times V}$ denotes the duality pairing on $V' \times V$. By introducing the *Riesz operator* for the test space, the dual norm $\|\cdot\|_{V'}$ can be replaced with $\|R_V^{-1}(\cdot)\|_V$ and arrive at a simpler (to compute with) minimization problem:

$$(2.1) \quad u_h = \arg \min_{w_h \in U_h} \frac{1}{2} \|R_V^{-1}(l - Bw_h)\|_V^2.$$

The Riesz operator $R_V : V \rightarrow V'$ is an isometric isomorphism defined by

$$(R_V y)(v) = \langle R_V y, v \rangle_{V' \times V} = (y, v)_V, \quad y, v \in V.$$

Equation (2.1) leads to two additional characterizations for the DPG method. Taking the Gâteaux derivative, leads to the following linear problem:

$$(2.2) \quad \begin{cases} u_h \in U_h \subset U \\ (R_V^{-1}(l - Bu_h), R_V^{-1}B\delta u_h) = 0, \quad \delta u_h \in U_h \end{cases}$$

We introduce the *trial-to-test operator* $T : U_h \rightarrow V$, defined by

$$(2.3) \quad (T\delta u_h, \delta v)_V = b(\delta u_h, \delta v) \quad \delta u_h \in U_h, \delta v \in V, \quad \text{i.e., } T = R_V^{-1}B$$

and we define the *optimal test space* as $V_h^{\text{opt}} := T(U_h)$, i.e., the image of the trial to test operator T acting on the trial space U_h . Therefore, problem (2.2) leads to the following Petrov–Galerkin scheme with *optimal test functions*.

$$\begin{cases} u_h \in U_h \subset U \\ b(u_h, v_h) = l(v_h), \quad v_h \in V_h^{\text{opt}} \end{cases}$$

Clearly, the above construction of the test space always delivers a symmetric positive definite system and it is *optimal* in then sense that it guarantees uniform and unconditional discrete stability [15]. Indeed, for $v_h \in V_h^{\text{opt}}$ we have:

$$\frac{b(u_h, v_h)}{\|v_h\|_V} = \frac{b(u_h, Tu_h)}{\|Tu_h\|_V} = \|Tu_h\|_V = \sup_{v \in V} \frac{|(Tu_h, v)_V|}{\|v\|_V} = \sup_{v \in V} \frac{|b(u_h, v)|}{\|v\|_V}.$$

Therefore the *discrete inf–sup* stability condition [3], is trivially satisfied.

Additionally, if we define $\psi := R_V^{-1}(l - Bu_h)$ to be the *error representation function*, (2.2) leads to the following mixed problem:

$$\begin{cases} u_h \in U_h, \psi \in V, \\ (\psi, v)_V + b(u_h, v) = l(v), & v \in V, \\ b(w_h, \psi) = 0, & w_h \in U_h \end{cases}$$

where now we solve simultaneously for the original unknown u_h and the error representation function ψ . Note that the norm of ψ offers a built-in a-posteriori error indicator. Indeed, if we define the *energy* norm as

$$\|u\|_E := \|Bu\|_{V'} = \sup_{v \in V} \frac{|b(u, v)|}{\|v\|_V}$$

then we can compute the Galerkin error measured in the *energy* norm, by computing the norm of the error representation function (residual).

$$\|u_h - u\|_E = \|B(u_h - u)\|_{V'} = \|Bu_h - l\|_{V'} = \|R_V^{-1}(Bu_h - l)\|_V = \|\psi\|_V.$$

2.2. The ultraweak formulation – linear acoustics. Consider the *operator form* of the time-harmonic linear acoustics equations with impedance boundary conditions in a simply connected domain Ω :

$$\begin{cases} i\omega p + \operatorname{div} u = f_1, & \text{in } \Omega \\ i\omega u + \nabla p = f_2, & \text{in } \Omega \\ p - u \cdot n = 0, & \text{on } \partial\Omega \end{cases}$$

Here, p and u denote the pressure and the velocity of the wave respectively. Define the group variable $\mathbf{u} := (p, u)$, the group load $\mathbf{f} := (f_1, f_2) \in L^2(\Omega) \times (L^2(\Omega))^d =: \mathbf{L}^2(\Omega)$ and the acoustics operator $A((p, u)) = (i\omega p + \operatorname{div} u, i\omega u + \nabla p)$. Then the problem can be compactly written as:

$$(2.4) \quad \begin{cases} \mathbf{u} \in D(A) \\ A\mathbf{u} = \mathbf{f} \end{cases}$$

where $D(A)$ denotes the domain of A :

$$D(A) := \{\mathbf{u} \in \mathbf{L}^2(\Omega) : A\mathbf{u} \in \mathbf{L}^2(\Omega), p - u \cdot n = 0 \text{ on } \partial\Omega\}.$$

By testing (2.4) with the group variable $\mathbf{v} := (q, v)$, and using integration by parts, the regularity is passed from the trial to the test space. The resulting variational formulation is called the *ultraweak formulation*:

$$\begin{cases} \mathbf{u} \in \mathbf{L}^2(\Omega) \\ (\mathbf{u}, A^*\mathbf{v}) = (\mathbf{f}, \mathbf{v}), & \mathbf{v} \in D(A^*), \end{cases}$$

where $D(A^*)$ is equipped with the adjoint graph norm,

$$\|\mathbf{v}\|_V^2 := \|A^*\mathbf{v}\|^2 + \alpha^2 \|\mathbf{v}\|^2,$$

with a scaling parameter $\alpha = \mathcal{O}(1)$. Observe that for this specific problem, the operator is (formally) *skew-adjoint*, $A^* = -A$, and

$$D(A^*) = \{\mathbf{v} \in \mathbf{L}^2(\Omega) : A^*\mathbf{v} \in \mathbf{L}^2(\Omega) : q + v \cdot n = 0 \text{ on } \partial\Omega\}.$$

Note that with the above choice of norm on the test space and $\alpha = 0$ the ultraweak formulation delivers L^2 -projection. In practice, when the space is broken, α has to be non-zero in order for the norm to be *localizable*. A common choice is $\alpha = 1$. It is easy to see that the two norms are robustly equivalent. This in turn gives a pollution free method for one-space dimension problems (see [44, 16, 56, 26] for additional details). Additionally, we emphasize that other equivalent (in terms of stability) formulations can be explored [35, 14, 21, 34], however, especially for wave operators, the ultraweak formulation is provably superior because of its approximability properties [56, 16, 43, 40, 32, 31].

The boundary conditions on the test functions (built into the definition of $D(A^*)$), can be eliminated at the expense of introducing extra unknowns $\hat{u} := (\hat{p}, \hat{u}_n)$. We arrive then at the linear problem

$$\begin{cases} u \in L^2(\Omega), \hat{u} \in \hat{U} \\ (u, A^*v) + \langle \hat{u}, v \rangle = (f, v), \quad v \in H_{A^*}(\Omega), \end{cases}$$

where

$$H_{A^*}(\Omega) := \{v \in L^2(\Omega) : A^*v \in L^2(\Omega)\}$$

is equipped with the same graph norm as above. The flux \hat{u} is defined as the trace of a function that lives in the original energy graph space,

$$\hat{U} := \text{tr } D(A).$$

More precisely, the trace space is defined by:

$$(2.5) \quad \begin{aligned} \hat{U} &= \{(\hat{p}, \hat{u}_n) : \text{exists } (p, u) \in D(A) \text{ such that } (\hat{p}, \hat{u}_n) = \text{tr}(p, u)\} \\ &= \{(\hat{p}, \hat{u}_n) \in H^{\frac{1}{2}}(\partial\Omega) \times H^{-\frac{1}{2}}(\partial\Omega) : \hat{p} - \hat{u}_n = 0 \text{ on } \partial\Omega\}. \end{aligned}$$

and it is equipped with the minimum energy extension norm,

$$\|\hat{u}\|_{\hat{U}} := \inf_{\substack{u \in D(A) \\ \text{tr } u = \hat{u}}} \|u\|_{H_A(\Omega)}$$

where

$$(2.6) \quad \|u\|_{H_A(\Omega)}^2 := \|Au\|^2 + \|u\|^2.$$

2.3. The Discontinuous Petrov–Galerkin method: “breaking” the test space. Following the same procedure as above, but with broken (discontinuous) test functions instead, we arrive at the *ultaweak variational formulation with broken test spaces*:

$$\begin{cases} u \in L^2(\Omega), \hat{u} \in \hat{U}_h \\ \underbrace{(u, A^*v) + \langle \hat{u}, v \rangle}_{=: b((u, \hat{u}), v)} = (f, v), \quad v \in V(\Omega_h). \end{cases}$$

The broken test space is defined as:

$$V(\Omega_h) := \{v \in L^2(\Omega) : A_h^*v \in L^2(\Omega)\}$$

and it is equipped with the adjoint graph norm,

$$\|v\|_{V(\Omega_h)}^2 := \sum_K \underbrace{(\|A_h^*v\|^2 + \|v\|^2)}_{=: \|v\|_{V(K)}^2}.$$

Here, the symbol h indicates that the operator is understood *element-wise*. Fluxes and traces are defined now on the whole mesh skeleton Γ_h and equipped with the minimum energy extension norm. The trace space is given by

$$\hat{U}_h = \{(\hat{p}, \hat{u}_n) \in H^{\frac{1}{2}}(\Gamma_h) \times H^{-\frac{1}{2}}(\Gamma_h) : \hat{p} - \hat{u}_n = 0 \text{ on } \partial\Omega\}.$$

Additionally, we define the *energy norm* by

$$(2.7) \quad \|(u, \hat{u})\|_E^2 := \left(\sup_{v \in V(\Omega_h)} \frac{b((u, \hat{u}), v)}{\|v\|_{V(\Omega_h)}} \right)^2 = \sum_K \underbrace{\left(\sup_{v \in V(K)} \frac{|(u, A_h^*v) + \langle \hat{u}, v \rangle|}{\|v\|_{V(K)}} \right)^2}_{=: \|(u, \hat{u})\|_{E,K}^2}$$

The above energy norm satisfies the same property as standard Sobolev energy norms, i.e., the global norm squared ($\|(u, \hat{u})\|_E^2$) equals the sum of element norms squared ($\|(u, \hat{u})\|_{E,K}^2$) [10]. This result is very crucial for the analysis of the preconditioner since it allows us to consider a single element, when proving norm equivalence.

2.4. The practical DPG method. We note that for practical computations one has to derive the so called *practical DPG method* [29]. That is, the inversion of the Riesz operator can only be approximated, since computing it exactly would require the solution of an infinite dimensional boundary value problem. The procedure is then to consider a truncated finite dimensional test space $V^r \subset V$ with $\dim(V^r) \gg \dim(U_h)$, the so called *enriched test space*. Consequently, both the trial-to-test operator T and the error representation function ψ are also approximated. Various studies exist that examine the effect of this approximation to the discrete stability with the overall conclusion that stability is not significantly affected. We refer the reader to [41, 10, 29, 11] for further reading.

3. ANALYSIS OF THE ONE-LEVEL PRECONDITIONER

We are now ready to provide a theoretical convergence analysis for the one level additive Schwarz smoother. The analysis is based on the *subspace correction theory* [54, 53, 52] for self-adjoint positive definite operators. We start by stating the subspace correction theorem and we proceed by verifying its assumptions for the DPG ultraweak formulation for the time-harmonic acoustics problem presented in the previous section.

3.1. Additive Schwarz preconditioner and the subspace correction theory. Let B be self-adjoint and positive definite defined on the vector space U . Suppose that $U_i, i = 1, \dots, J$, are closed subspaces of the Hilbert space $(U, (\cdot, \cdot))$. Additionally, let B_i be self-adjoint and positive definite operator defined on U_i by

$$(B_i u_i, v_i) = (B u_i, v_i), \quad \forall u_i, v_i \in U_i,$$

and let $Q_i : U \rightarrow U_i$ denote the (\cdot, \cdot) -orthogonal projection onto U_i . Then, the operator

$$A = \sum_{i=1}^J B_i^{-1} Q_i$$

is called the *additive preconditioner* based on subspaces U_i and operators B_i . The additive Schwarz algorithm is given by:

Algorithm 1 Additive Schwarz/parallel subspace correction

- | | |
|--|---|
| 1: procedure PSC(u_k, u_{k+1}) | ▷ Given u_k compute u_{k+1} |
| 2: $r = l - B u_k$ | ▷ Compute initial residual |
| 3: $r_i = Q_i r$ | ▷ Project the residual on to U_i |
| 4: $B_i z_i = r_i$ | ▷ Solve on the subspace the local problem |
| 5: $u_{k+1} = u_k + \theta \sum_{i=1}^J z_i$ | ▷ Correct u_k on each subspace, $\theta \in (0, 1]$ |
-

Theorem 1 (*Subspace correction* [53, 54, 52]). *For the above setting, assume the following two statements hold:*

- (*strengthened Cauchy-Schwarz inequality*) *there exists $\beta > 0$ such that for all $u_i, v_i \in U_i$*

$$\sum_{i=1}^J \sum_{j=1}^J |(u_i, v_j)_B| \leq \beta \left(\sum_{i=1}^J \|u_i\|_B^2 \right)^{1/2} \left(\sum_{j=1}^J \|v_j\|_B^2 \right)^{1/2}$$

- (*stable decomposition*) *there exists $\alpha > 0$ such that $\forall u \in U$, there exists a decomposition $u = \sum_{i=1}^J u_i$, with $u_i \in U_i$, that satisfies*

$$\sum_{i=1}^J \|u_i\|_B^2 \leq \alpha^{-1} \|u\|_B^2.$$

Then, the following equivalence relation is true:

$$(3.1) \quad \alpha(u, u)_B \leq (Pu, u)_B \leq \beta(u, u)_B$$

where $P = \sum_{i=1}^J P_i$ and $P_i : U \rightarrow U_i$ is the $(\cdot, \cdot)_B$ -orthogonal projector, i.e.,

$$(P_i u, v_i)_B = (u, v_i)_B, \quad \forall v_i \in U_i.$$

Note that $B_i P_i = Q_i B$ and therefore $P = AB$ is the preconditioned matrix. Consequently, Equation (3.1) gives an estimate of the condition number $\kappa(AB) = \beta/\alpha$. A detailed proof of this theorem can be found in [52].

3.2. Verification of the theorem assumptions in the DPG setting.

3.2.1. Setup. Let $u = (u, \hat{u})$ denote the group variable including the field and trace variables of the ultraweak formulation. We want to precondition the energy norm $\|\cdot\|_E$ given by Equation (2.7). Let $b_E(\cdot, \cdot)$ be the Hermitian and coercive form corresponding to the energy norm, i.e.,

$$b_E(u, v) = b(u, Tv) = (Tu, Tv)_V$$

where T is the trial-to-test operator for the broken formulation defined according to Equation (2.3). We introduce $\{\Omega_i\}_{i=1}^J$ a finite cover of Ω such that each $\bar{\Omega}_i$ is the support of a vertex shape function. We denote the size of a vertex patch by δ . In addition, we assume that the cover satisfies the *finite overlap property*, i.e., there exists an integer r such that each point of Ω is contained in at most r of the sets Ω_i . Equivalently, let χ_i be the characteristic function of Ω_i . Then

$$(3.2) \quad \sum_{i=1}^J \|\chi_i\|_{L^\infty}^2 = \sum_{i=1}^J \|\chi_i\|_{L^\infty} \leq r$$

Finally, the local energy subspaces corresponding to the partition are given by:

$$U_i = L^2(\Omega_i) \times \hat{U}_i,$$

where $\hat{U}_i := \{\hat{u} \in \hat{U}_h : \hat{u} = 0 \text{ on } \Gamma_h - \Omega_i\}$, and Γ_h is the mesh skeleton.

3.2.2. Strengthened Cauchy–Schwarz inequality. Verifying the first assumption, of Theorem 1 is straight forward. Let $u_i \in U_i$ and $v_j \in U_j$. Then, for any inner product $(\cdot, \cdot)_B$ the inequality

$$|(u_i, v_j)_B| \leq \varepsilon_{ij} \|u_i\|_B \|v_j\|_B$$

is true for $\varepsilon_{ij} \leq 1$. Indeed, from Cauchy–Schwarz inequality, $\varepsilon_{ij} = 1$ when $(u_i, v_j)_B \neq 0$, but it can be chosen to be 0 when $(u_i, v_j)_B = 0$. Taking the sum over i and j we have:

$$\begin{aligned} \sum_{i=1}^J \sum_{j=1}^J |(u_i, v_j)_B| &\leq \sum_{i=1}^J \sum_{j=1}^J \varepsilon_{ij} \|u_i\|_B \|v_j\|_B \\ &\leq \rho(\mathcal{E}) \left(\sum_{i=1}^J \|u_i\|_B^2 \right)^{1/2} \left(\sum_{j=1}^J \|v_j\|_B^2 \right)^{1/2} \end{aligned}$$

where $\rho(\mathcal{E}) = \|\mathcal{E}\|_2$ is the spectral radius of the matrix ε_{ij} . Note that the entries of the matrix ε_{ij} are zero for non overlapping subdomains. Indeed, suppose that $\text{supp}\{u_i\} \subseteq \Omega_i$. Then $\text{supp}\{Tu_i\} \subseteq \Omega_i$. This is a direct consequence of the use of broken test spaces, i.e., T , the DPG *trial-to-test operator* is local and therefore Tu_i is discontinuous. Therefore

$$b_E(u_i, v_j) = b(u_i, Tv_j) = (Tu_i, Tv_j)_V = 0,$$

for all $v_j \in \Omega_j$ such that $\text{supp}\{v_j\} \cap \Omega_i = \emptyset$, and consequently

$$b_E(u_i, v_j) = 0,$$

for Ω_i, Ω_j disjoint and $\text{supp}\{u_i\} \subseteq \Omega_i, \text{supp}\{v_j\} \subseteq \Omega_j$. Finally, by the finite overlap assumption, we obtain an upper bound for the spectral radius:

$$\rho(\mathcal{E}) \leq r$$

and the result reads:

$$\sum_{i=1}^J \sum_{j=1}^J |b_E(u_i, v_j)| \leq r \left(\sum_{i=1}^J \|u_i\|_E^2 \right)^{1/2} \left(\sum_{j=1}^J \|v_j\|_E^2 \right)^{1/2}$$

3.2.3. Stable Decomposition. Proving the second assumption of Theorem 1 is a bit more involved. This is exactly the well-known *stable decomposition* assumption, which in simple terms it says that if u is decomposed into patch contributions, then the sum of energies stored in the patches must be controlled by the energy of u . We start with a DPG global stability result [10, 16],

$$\|u\|^2 + \|\hat{u}\|_{\hat{U}}^2 \leq \left[\frac{1}{\gamma^2} + \left(1 + \frac{1}{\gamma}\right)^2 \right] \|(u, \hat{u})\|_E^2.$$

Above γ denotes the global boundedness below constant for operator A . For the acoustics operator and the case of impedance BC, constant γ is independent of frequency ω [16]. Combining the stability estimate with continuity we derive the equivalence relation:

$$\gamma_1^2 (\|u\|^2 + \|\hat{u}\|_{\hat{U}}^2) \leq \|(u, \hat{u})\|_E^2 \leq M_1^2 (\|u\|^2 + \|\hat{u}\|_{\hat{U}}^2)$$

where $\gamma_1^{-2} = \frac{1}{\gamma^2} + \left(1 + \frac{1}{\gamma}\right)^2$ and M_1 is a continuity constant. The above relation allows us to construct a stable decomposition for the original trial norm instead of the energy norm. We will prove the result by considering separate cases for the L^2 space and the trace space \hat{U} . For simplicity we consider a the two-dimensional case for the rest of the proof. The three-dimensional case is fully analogous.

Stable decomposition for the L^2 space. Let $Y_p = \mathcal{Q}^{(p,q)} = \mathcal{P}^p \otimes \mathcal{P}^q$ denote the space of polynomials of order less or equal p, q with respect to x, y respectively. These spaces are described in detail in [43, Ch.2]. Additionally, let $\{\Phi_j\}_{j=1}^J$ be a partition of unity subordinate to the covering Ω_j of Ω , so that

$$\sum_{j=1}^J \Phi_j = 1, \quad 0 \leq \Phi_j \leq 1, \quad \text{supp}(\Phi_j) \subset \Omega_j$$

Given $u \in Y_p$ we define the following decomposition

$$u_j = P \Phi_j u$$

where $P : \mathcal{Q}^{(p+1,q+1)} \rightarrow \mathcal{Q}^{(p,q)}$ is the L^2 -orthogonal projection. Obviously

$$\sum_{i=1}^J u_j = \sum_{j=1}^J P \Phi_j u = P \underbrace{\sum_{i=1}^J \Phi_j}_{=1} u_j = P u = u$$

It easy to see that the decomposition is stable. Indeed

$$\|u_j\| = \|P \Phi_j u\| = \|P \Phi_j u\|_{L^2(\Omega_j)} \leq \|P\| \|\Phi_j\|_{L^\infty} \|u\|_{L^2(\Omega_j)} \leq \|u\|_{L^2(\Omega_j)}$$

Summing up for all subdomains, and by the finite overlap property (3.2) we have

$$\sum_{j=1}^J \|u_j\|^2 \leq \sum_{j=1}^J \|u\|_{L^2(\Omega_j)}^2 \leq r \|u\|^2$$

Stable decomposition for the trace space \hat{U} . We shall consider a single element $K \in \Omega_j$. We denote by $H_A(K)$ the graph energy space of functions defined on element K . Recall that $\hat{\Omega}_j$ is defined by the support of a vertex shape function Φ_j of size δ . Then, $\{\Phi_j\}_{i=1}^J$ is a partition of unity,

$$\sum_{j=1}^J \Phi_j = 1, \quad 0 \leq \Phi_j \leq 1, \quad \text{supp}(\Phi_j) \subset \Omega_j \quad \text{and} \quad \|\nabla \Phi_j\|_{L^\infty} \leq C \delta^{-1}$$

Let $H_A(K)$ be the graph energy space of functions defined on element K , we denote by $H_A^j(K)$ the subspace of functions vanishing on $K - \text{supp}(\Phi_j)$.

Operator dependent projection based interpolant. Let \hat{u} be a sufficiently regular function defined on the skeleton. Additionally, let's assume the existence of a projection-based interpolant $\hat{\Pi}\hat{u}$ (which will be defined explicitly later on) taking values on the mesh skeleton such that

$$(3.3) \quad \|\hat{\Pi}\Phi_j\hat{u}\|_{\hat{U}} \leq \underbrace{\|\hat{\Pi}\|_{\hat{U}}}_{=:C_{\hat{\Pi}}} \|\Phi_j\hat{u}\|_{\hat{U}}$$

Construction of such an interpolant is operator-dependent, and we expect constant $C_{\hat{\Pi}}$ to grow mildly with the frequency ω for the acoustics operator, and be independent of the polynomial order p and the fine mesh size h . We investigate the dependence on ω , h , and p numerically by constructing such an interpolation operator. The construction and the numerical experiments are described in the next section. Consider now a function on the skeleton \hat{u} . We define the corresponding decomposition by

$$\hat{u}_j = \hat{\Pi}\Phi_j\hat{u}.$$

Clearly,

$$\sum_{j=1}^J \hat{u}_j = \sum_{j=1}^J \hat{\Pi}\Phi_j\hat{u} = \hat{\Pi}\left(\underbrace{\sum_{j=1}^J \Phi_j}_{=1}\hat{u}\right) = \hat{\Pi}\hat{u} = \hat{u}$$

and by the postulated property (3.3)

$$\|\hat{u}_j\|_{\hat{U}} = \|\hat{\Pi}\Phi_j\hat{u}\|_{\hat{U}} \leq C_{\hat{\Pi}}\|\Phi_j\hat{u}\|_{\hat{U}} \leq C_{\hat{\Pi}}\|\Phi_j u\|_{H_A}$$

where u is an extension of \hat{u} to the element such that $tru = \hat{u}$. The last term bounds nicely for the acoustics operator,

$$\|\Phi_j u\|_{H_A}^2 = \|A(\Phi_j u)\|^2 + \|(\Phi_j u)\|^2$$

and

$$\begin{aligned} \|A(\Phi_j(u, p))\|^2 &= \|i\omega\Phi_j p + \text{div}(\Phi_j u)\|^2 + \|i\omega u + \nabla(\Phi_j p)\|^2 \\ &\leq 2\left(\|\Phi_j(i\omega p + \text{div} u)\|^2 + \|\Phi_j(i\omega u + \nabla p)\|^2 + \|\nabla\Phi_j \cdot u\|^2 + \|\nabla\Phi_j p\|^2\right) \\ &\leq 2\left(\|A(u, p)\|^2 + \delta^{-2}(\|u\|^2 + \|p\|^2)\right) \end{aligned}$$

Therefore

$$\|\hat{u}_j\|_{\hat{U}} \leq C_{\hat{\Pi}}C_{\delta}\|u\|_{H_A}, \quad \text{where} \quad C_{\delta} = \mathcal{O}(\delta^{-1})$$

Note that the constant C_{δ} is independent of frequency. The above inequality is true for an arbitrary extension u of \hat{u} . Taking the infimum on the right hand side with respect to extensions and summing over all subdomains we obtain the stable decomposition for the trace space. Combining everything together we have the following lemma.

Lemma 1. *For every function $u = (u, \hat{u}) \in L^2(\Omega) \times \hat{U}$ there exists a stable decomposition*

$$u = \sum_{j=1}^J u_j, \quad u_j \in L^2(\Omega_j) \times \hat{U}_j,$$

such that

$$\sum_{j=1}^J \|u_j\|_E^2 \leq C\|u\|_E^2.$$

where $C = \mathcal{O}(rC_{\hat{\Pi}}^2C_{\delta}^2)$. Here r comes from the finite overlap property and $C_{\delta} = \mathcal{O}(\delta^{-1})$. This in turn gives an upper bound for the condition number of the preconditioned DPG system

$$\kappa(\text{AB}) \leq C \frac{r^2}{\delta^2}, \quad \text{where } C = \mathcal{O}(C_{\hat{\Pi}}^2).$$

It remains to study the dependence of the interpolation norm $C_{\hat{\Pi}}$ on the discretization size h , polynomial order p and the frequency ω . For this, we design a numerical experiment to compute the interpolation norm. We present the results below. Further details describing the design of the experiment are given in Appendix A.

3.3. Numerical study of $\|\hat{\Pi}\|_{\hat{\mathcal{U}}}$. We examine the dependence of the norm of the interpolation operator described above, on the polynomial order p , the discretization size h , and the frequency ω . We present the results in the tables below for polynomial orders $p = 2, 4, 6$. For each case of p we also present convergence of the preconditioned CG solver. The subdomains are defined to be the support of a vertex function defined on a quadrilateral uniform mesh of size $h = 1/2$ (i.e, 9 subdomains with a fixed overlap size $\delta = 1/2$).

$h \setminus \omega$	π	2π	4π	8π	16π
1/2	15	16	17	11	8
1/4	16	18	20	14	9
1/8	16	18	22	23	10
1/16	16	18	23	24	25
1/32	16	19	23	25	25

A Iteration count

$h \setminus \omega$	π	2π	4π	8π	16π
1/2	2.553	4.327	6.088	3.768	2.268
1/4	2.540	4.427	7.598	6.320	4.328
1/8	2.510	4.415	6.972	8.247	6.061
1/16	2.500	4.362	6.841	8.379	9.555
1/32	2.497	4.343	6.785	8.651	9.621

B Value of the constant $\|\hat{\Pi}\|_{\hat{\mathcal{U}}}$

TABLE 1. Polynomial order $p = 2$. Left: iteration count for CG preconditioned with additive Schwarz smoother with fixed $\delta = 1/2$. Right: The value of the interpolation norm.

$h \setminus \omega$	π	2π	4π	8π	16π
1/2	16	18	22	17	9
1/4	16	18	22	22	11
1/8	16	18	22	25	23
1/16	17	19	23	25	25
1/32	17	19	23	25	26

A Iteration count

$h \setminus \omega$	π	2π	4π	8π	16π
1/2	1.516	1.976	3.222	3.168	1.871
1/4	1.497	2.021	3.328	5.262	2.582
1/8	1.491	2.046	3.432	5.122	5.891
1/16	1.490	2.055	3.492	5.260	6.478
1/32	1.489	2.057	3.512	5.352	6.556

B Value of the constant $\|\hat{\Pi}\|_{\hat{\mathcal{U}}}$

TABLE 2. Polynomial order $p = 4$. Left: iteration count for CG preconditioned with additive Schwarz smoother with fixed $\delta = 1/2$. Right: The value of the interpolation norm.

$h \setminus \omega$	π	2π	4π	8π	16π
1/2	17	19	23	24	12
1/4	16	18	22	23	22
1/8	17	18	22	24	24
1/16	17	18	22	25	25
1/32	17	18	23	25	26

A Iteration count

$h \setminus \omega$	π	2π	4π	8π	16π
1/2	1.271	1.522	1.874	3.514	2.391
1/4	1.268	1.512	1.753	3.093	3.026
1/8	1.267	1.510	1.734	2.750	3.583
1/16	1.267	1.510	1.730	2.685	3.724
1/32	1.266	1.510	1.730	2.671	3.962

B Value of the constant $\|\hat{\Pi}\|_{\hat{\mathcal{U}}}$

TABLE 3. Polynomial order $p = 6$. Left: iteration count for CG preconditioned with additive Schwarz smoother with fixed $\delta = 1/2$. Right: The value of the interpolation norm.

For these experiments we simulate a plane wave in the square domain $\Omega = (0, 1)^2$, using impedance boundary conditions on the entire boundary of the domain. We run our simulations for five different frequencies and we perform successive uniform h -refinements, starting with a mesh of size $h = 1/2$. The CG solver is terminated when the l_2 -norm of the residual drops below 10^{-6} . The numbers in red color denote that the error is above 90%, i.e., the mesh is not fine enough to resolve the wave.

The following interesting observations can be derived from the results. First, the solver shows convergence, independent of the polynomial order. This, can also be verified in the value of the constant $\hat{\Pi}_{\hat{V}}$. Gopalakrishnan and Schöberl observed the same behavior in their construction of a multiplicative Schwarz preconditioner in [30]. The second observation involves the dependence on the mesh size h . As we can clearly see, both the constant and the CG convergence are independent of h .

Lastly, the number of iterations of the solver grows mildly with respect to the frequency, an observation that is also reflected in the value of the constant. However, in cases where the mesh is too coarse to resolve a high frequency wave, the solver is more efficient with respect to the number of iterations (see numbers in red). This behavior, might not seem to be of great value for uniform meshes, but in case of adaptive refinements, this could be very beneficial. As we have seen in the previous chapter, the DPG method being unconditionally stable, allows for adaptive refinements starting from very coarse meshes. Therefore, integrating an iterative solver within the adaptive process, even in the pre-asymptotic region, seems to be a promising direction to follow.

Overall, the computed interpolation norm gives a very useful insight on the effectiveness of the preconditioner. We emphasize that, there is no reason to believe that the bound provided by the interpolation norm is sharp. However, the overall trend of the value of the norm, is consistent with the convergence behavior of the solver.

Note that in practical computations the interior degrees of freedom (L^2 field variables) can be eliminated in an element-wise fashion. The size of the final system is then significantly reduced, since the condensed system involves only the trace unknowns. A well known Schur complement result guarantees that the upper bound of the condition number of the condensed and original operators are the same [43, Appendix C]. This allows us to analyze the original system even if in actual computations we use the statically condensed system.

4. EXTENSION TO MULTIGRID

We would like to extend our construction to the two- (multi-) level setting, i.e., accelerate the preconditioner by coupling (in the multiplicative way) the additive Schwarz smoother with a coarse grid solve. The main reason is the unavoidable dependence of the one-level preconditioner on the number of subdomains and the size of overlap. In order to, keep the number of iterations of the solver under control, both the number of subdomains and the size of the overlap have to remain constant and unfortunately this adds significant work on the local solves within each subdomain as h decreases. On the other hand, if the number of subdomains increases, the cost of each local solve remains constant, but the overall number of iterations grows. A preliminary implementation for a two-level setting for acoustics for two dimensional problems was introduced in [44]. Here, we extend that work to three dimensional simulations and to the multi-level setting.

4.1. Discussion on implementation. The major components of the multigrid preconditioner are: a) inter-grid transfer operators, b) smoother and c) coarse grid solver. We describe the construction of each component for the ultraweak formulation for the acoustics problem presented in Section 2.3.

Inter-grid transfer operators. Recall that the ultraweak formulation involves two sets of variables, i.e., the field variables $(u, p) \in U = \mathbf{L}^2(\Omega)$ and the trace variables $(\hat{p}, \hat{u}_n) \in \hat{U}$ given by Equation (2.5). Since the field variables are only L^2 -conforming, they can be element-wise eliminated from the global system. The resulting linear system therefore involves only the trace unknowns. Constructing, a prolongation operator from a coarse trace space to an h -refined fine trace space is not so straight forward. The reason is simply because new edges (faces in 3D) which are created after an h -refinement hold new interface variables which have no ancestors. Therefore the usual prolongation operator (natural inclusion) based on constrained approximation [17] is not well defined. Following our work in [44], we overcome this difficulty by considering a two-step procedure for the prolongation. First, we construct the so called *macro-grid*. This is the mesh that is constructed by eliminating all

Coarse-grid solve. For the coarse grid correction we use a direct solver, specifically the Multi-frontal Sparse Direct Solver MUMPS [39, 1]. Note that the coarse grid has the field variables eliminated, and its size is usually significantly smaller than the size of the fine-grid system.

5. NUMERICAL RESULTS

This chapter is devoted to numerical results for relatively large simulations, using the Conjugate Gradient (CG) solver preconditioned with the multigrid technology. Our interest is in computationally challenging acoustics and electromagnetic problems. We present various numerical experiments in acoustics and Maxwell equations in both uniform and adaptive refinement setting. All the simulations were performed on a single node with 24 cores and 256G of RAM memory, using the Finite Element Library *hp3D* [13, 17, 22].

5.1. Time harmonic Maxwell equations. We first consider the time harmonic form of the Maxwell equations given by

$$\begin{cases} i\omega\mu H + \nabla \times E = 0 \\ -i\omega\epsilon E + \nabla \times H = J \end{cases}$$

where E is the electric field, H is the magnetic field, ω is the angular frequency, ϵ is the permittivity and μ and permeability of the material.

5.1.1. Comparison with standard multigrid methods. For our first Maxwell experiment we compare our DPG multigrid technology with the multigrid preconditioner for the standard Galerkin method described in [28]. Following the experiment in [28], we consider the computational domain $\Omega = (0, 1)^3$ and a *perfect electric conductor* (PEC) material, i.e, the boundary condition is $n \times E = 0$ on $\partial\Omega$. For the discretization we use a uniform hexahedral mesh of the lowest order. Define the space $H_0^{-\frac{1}{2}}(\text{curl}, \Gamma_h)$ on the mesh skeleton as

$$H_0^{-\frac{1}{2}}(\text{curl}, \Gamma_h) := \{E \in H^{-\frac{1}{2}}(\text{curl}, \Gamma_h) : n \times E = 0 \text{ on } \partial\Omega\}$$

where $H^{-\frac{1}{2}}(\text{curl}, \Gamma_h)$ is defined as [43, Ch. 2]. Then the ultraweak DPG formulation is

$$\begin{cases} E, H \in (L^2(\Omega))^3, \\ \hat{E} \in H_0^{-\frac{1}{2}}(\text{curl}, \Gamma_h), \hat{H} \in H^{-\frac{1}{2}}(\text{curl}, \Gamma_h) \\ i\omega\mu(H, F) + (E, \nabla_h \times F) + \langle n \times \hat{E}, F \rangle_{\Gamma_h} = 0, & F \in H(\text{curl}, \Omega_h), \\ -i\omega\epsilon(E, G) + (H, \nabla_h \times G) + \langle n \times \hat{H}, G \rangle_{\Gamma_h} = (J, G), & G \in H(\text{curl}, \Omega_h) \end{cases}$$

where the letter h denotes element-wise operations. The simulation is driven by the right hand side which is chosen such that the exact solution for the electric field is given by the finite element lift of the function

$$E_{\text{ex}}(x, y, z) = [y(1-y)z(1-z), yx(1-x)z(1-z), x(1-x)y(1-y)].$$

We choose $\epsilon = \mu = 1$. The initial guess for the CG solver is set to zero and the iterations are terminated when the norm of the residual is reduced by a factor of 10^{-6} . In Tables 4a and 5a we present the iteration count of the CG solver when preconditioned with our multigrid technology for $\omega = 1$ and 10. Likewise, Tables 4b and 5b, retrieved from [28], show the iteration count of the GMRES solver preconditioned with the standard multigrid technology. We note that the smoother used in [28], is of multiplicative type (block Gauss-Seidel). In the tables, h and H denote the fine and the coarse grid discretization size respectively. For the DPG multigrid implementation, On each level we perform one pre- and one post- smoothing step, so that we keep the symmetry. At the coarsest level the problem is solved exactly with a direct solver.

$h \setminus H$	$1/2$	$1/4$	$1/8$	$1/16$
$1/4$	7			
$1/8$	7	7		
$1/16$	7	7	7	
$1/32$	6	7	6	7
$1/64$	6	6	6	6

A CG preconditioned with MG for DPG

$h \setminus H$	$1/2$	$1/4$	$1/8$	$1/16$
$1/4$	6			
$1/8$	7	7		
$1/16$	9	10	8	
$1/32$	10	10	9	7
$1/64$	11	11	9	8

B GMRES preconditioned with MG for FEM

TABLE 4. Iteration count for $\omega = 1$. Observe the uniform convergence with respect to h and H .

$h \setminus H$	$1/2$	$1/4$	$1/8$	$1/16$
$1/4$	9			
$1/8$	11	10		
$1/16$	14	12	9	
$1/32$	14	14	12	10
$1/64$	15	15	12	12

A CG preconditioned with MG for DPG

$h \setminus H$	$1/2$	$1/4$	$1/8$	$1/16$
$1/4$	3*			
$1/8$	2*	37		
$1/16$	3*	48	18	
$1/32$	2*	78*	22	16
$1/64$	2*	78*	21	17

B GMRES preconditioned with MG for FEM

TABLE 5. Iteration count for $\omega = 10$. The number of iterations for the DPG method grows mildly with the frequency but always converges to the true solution. Uniform convergence is achieved when a fine enough coarse grid is used. On the contrary the GMRES method fails to deliver reliable solutions when the coarse grid is in the pre-asymptotic region.

The purpose of this comparison is not to compare number to number the iteration count for each multi-grid technology, but rather to observe the general convergence trend of our preconditioner compared to the preconditioner for the standard Galerkin method. The entries n^* in the Table 5b denote that even though the GMRES solver did converge, the final iterate differed from the true solution by more than 10^{-3} (measured in the appropriate norm). This is a well known flaw of the GMRES solver, i.e., the residual of the GMRES algorithm might be small enough and the stopping criterion is met, but the output solution is far from the true solution. This happens when the coarse grid is not fine enough for the corresponding frequency and stability is lost. On the contrary, this undesirable convergence behavior is not happening for the Conjugate Gradient solver. The discrete pre-asymptotic stability of the DPG method, along with the theory of self-adjoint preconditioners, ensure that the CG solver always converge to the true solution. However, the convergence does depended on how “good” the coarse grid is with respect to the frequency. A similar dependence on the frequency is observed in our 2D simulations [44]. For both preconditioners, uniform convergence with respect to the frequency is recovered when the coarse grid is fine enough.

5.1.2. Fichera “oven” problem. For our second Maxwell experiment we solve the Fichera “oven” problem which was first presented in [10]. The adaptive nature of the DPG method makes it suitable for this problem because the solution is expected to be singular. The set up is as follows. For the construction of the domain we start with the cube $(0, 2)^3$ which is uniformly refined into eight cubes and then one is removed creating the Fichera corner. Then, an infinite waveguide is attached at the top and it’s truncated at a unit distance from the Fichera corner (see Figure 5.1). We choose $\epsilon = \mu = 1$ and $\omega = 5$. The simulation is driven by a non-homogeneous electric boundary condition on the waveguide and a homogeneous electric boundary condition

elsewhere. That is,

$$n \times E = \begin{cases} n \times E_d & \text{across the waveguide section} \\ 0 & \text{elsewhere} \end{cases}$$

where $E_d = (\sin \pi x_2, 0, 0)$ is the first propagating mode.






FIGURE 5.1. Fichera corner with a truncated infinite waveguide attached at the top (retrieved, September 10, 2020 from [10])

We start the simulation with a mesh of eight cubes with a uniform order of approximation $p = 3$ and perform successive h -adaptive refinements. The adaptive refinements are driven by the built-in DPG error indicator (the norm of the error representation function ψ), and terminated when the norm of the residual decreases by one order of magnitude. Note that an exact solution for this problem is not known.








FIGURE 5.2. Evolution of the mesh and the numerical solution for the real part of the x-component of the electric field. These results are for the meshes 1,3,5,7,9,11,13 and 15.

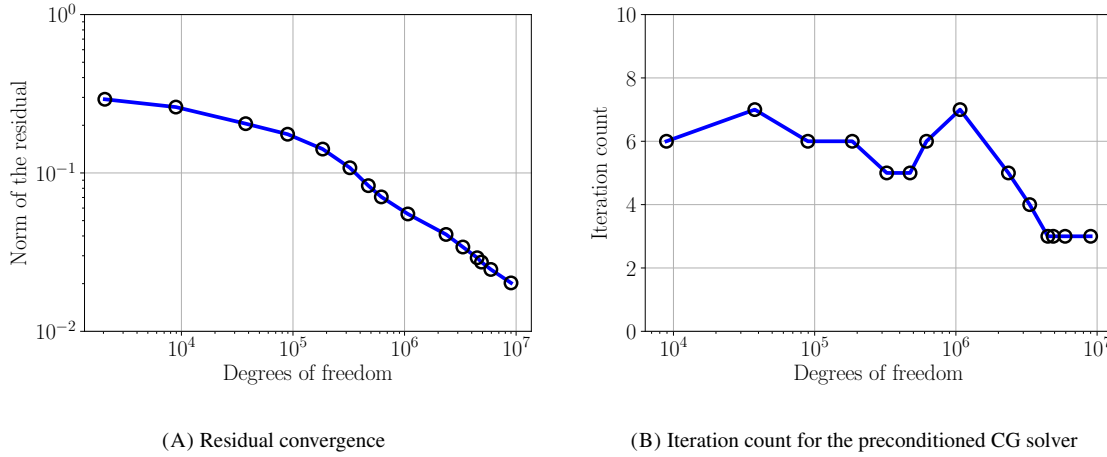


FIGURE 5.3. Fichera problem: convergence of the residual (left) and the CG solver(right)

The multigrid preconditioner setting is as follows. Starting from the finest mesh, several coarser adaptive meshes, which belong to the history of refinements, are selected in the multigrid cycle. An exact solve is performed at a coarse grid, where the size of the system is significantly smaller than the fine grid system and small enough to be solve efficiently with a direct solver. The selection is made at run time and depends on the computer architecture and the current memory available.

In Figure 5.2 we show the evolution of the mesh along with the numerical solution for the real part of the x -component of the electric field. As we can see, the ultraweak DPG method, captures the singularities at the reentrant corner and edges very fast. Observe how the solution changes qualitatively and quantitatively with the resolution of the singularities. The convergence of the residual is displayed in Figure 5.3 (A). Since there is no known exact solution to this problem, the only way to quantify the convergence is through the DPG residual. Clearly the norm of the residual converges to zero as we proceed with refinements. In Figure 5.3 (B) we illustrate the convergence behavior of the CG solver preconditioned with multigrid. For this simulation, at most five levels were used in the multigrid cycles. At all intermediate multigrid levels, we perform 10 smoothing steps with a relaxation parameter selected according to the finite overlap property (here $\theta = 0.2$). The CG iterations are initiated with a zero initial guess and they are terminated when the norm of the discrete residual is reduced below 10^{-6} . As we can observe, the solver shows uniform convergence with respect to the discretization size throughout the adaptive refinement process. We emphasize that the problem was too large to be solved with a direct solver.

5.1.3. Scattering of a Gaussian beam from a cube. Our last Maxwell example involves the simulation of a high-frequency Gaussian beam scattering from a rigid cube. We consider the domain $(0, 1)^3$ where a cube of side length $a = 1/7$ centered at $(0.5, 0.5, 0.5)$ is removed ($\Omega = (0, 1)^3 \setminus (\frac{3}{7}, \frac{4}{7})^3$). The simulation is driven by an impedance boundary condition on the outer cube and homogeneous electric boundary condition on the inner cube. The impedance data around the origin correspond to a high-frequency Gaussian beam propagating inside the domain at a specific angle. Away from the source the impedance data smoothly decay to zero in order to simulate absorbing boundary conditions. The DPG ultraweak formulation for this setting is given by:

$$\begin{cases} E, H \in (L^2(\Omega))^3, \\ \hat{E} \in H_{\Gamma_2}^{-\frac{1}{2}}(\text{curl}, \Gamma_h), \hat{H} \in H^{-\frac{1}{2}}(\text{curl}, \Gamma_h) \\ n \times (n \times \hat{E}) - n \times \hat{H} = \hat{g}, \quad \text{on } \Gamma_1, \\ i\omega\mu(H, F) + (E, \nabla_h \times F) + \langle n \times \hat{E}, F \rangle_{\Gamma_h} = 0, \quad F \in H(\text{curl}, \Omega_h), \\ -i\omega\epsilon(E, G) + (H, \nabla_h \times G) + \langle n \times \hat{H}, G \rangle_{\Gamma_h} = 0, \quad G \in H(\text{curl}, \Omega_h) \end{cases}$$

where, Γ_1, Γ_2 are the boundaries of the outer and the inner cube respectively and

$$H_{\Gamma_2}^{-\frac{1}{2}}(\text{curl}, \Gamma_h) := \{\hat{E} \in H^{-\frac{1}{2}}(\text{curl}, \Gamma_h) : n \times \hat{E} = 0 \text{ on } \Gamma_2\}.$$

We start the simulation with a uniform mesh of 342 cubes of order $p = 3$ and we initiate adaptive hp -refinements. We follow an ad-hoc refinement strategy; an element marked for refinement is h -refined unless its size is less than half a wavelength, in which case it's p -refined (two elements per wavelength). In order to resolve the anticipated singularities on the scatterer, the elements adjacent to its corners and edges are marked and forced to be h -refined when needed.

Finally, when an adaptive p -refinement is performed, we follow the following rule (minimum rule). When an element is marked for a p -refinement we first find the neighboring elements with respect to its faces. The order of a neighboring element is then increased by one if it is less than the intended order of the marked element. The last step is to assign orders to edges and faces. The order of a face is defined to be the minimum of the orders of its neighboring elements. After all faces are assigned their order, the edge orders take the value of the minimum of the orders of the faces they belong to. With this rather complicated rule, we ensure that a p -unrefinement is not possible (the sequence of meshes remains nested), and the constrained approximation technology for handling hanging nodes [17] is well defined. We mention that a maximum rule was also tested successfully. However, in order to maintain stability the enriched order of the test space had to be increased, and this often led to significantly increased element computation times.

We run the simulation for $\epsilon = \mu = 1$ and $\omega = 50\pi$. This frequency corresponds to approximately 40 wavelengths inside the computational domain. The multigrid setting is the same as in the previous example. For this experiment, at most eight multigrid levels were used. Adaptivity is guided by the norm of the residual

which is shown to be driven to zero in Figure 5.8 (A). The iteration count of the CG solver preconditioned with our multigrid technology is presented in Figure 5.8 (B). Note that the number of iterations of the CG solver remains under control throughout the adaptive refinement process. Lastly, the evolution of the mesh and the numerical solution of the real part of E_x are shown in Figures 5.5 and 5.7 respectively. Note that in these figures we show only the part of the domain below the plane defined by the point $(0.5, 0.5, 0.5)$ and the normal vector $(-0.5, -0.5, 1)$.

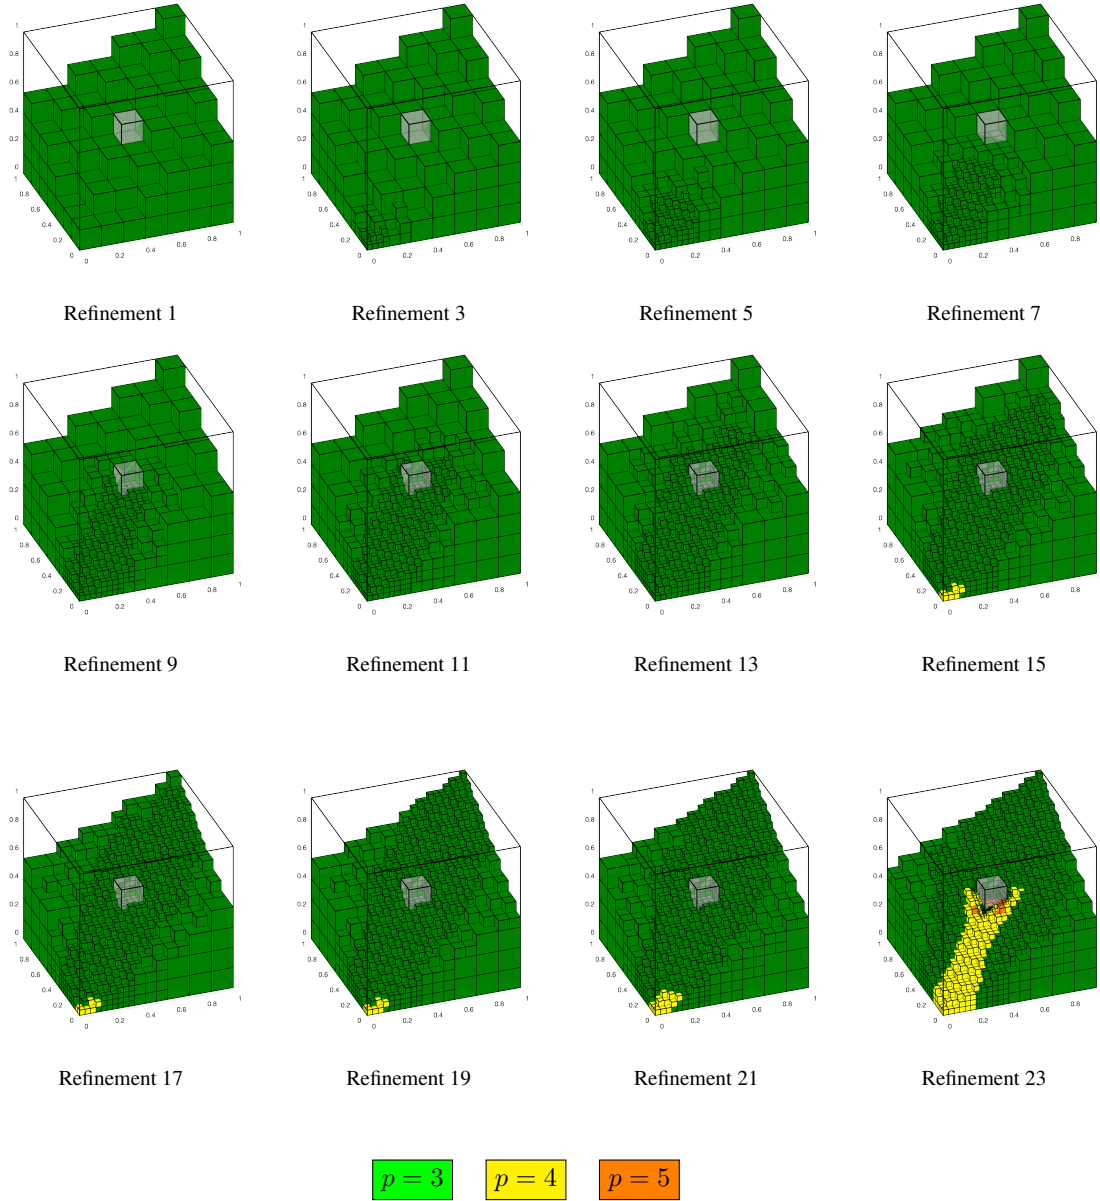


FIGURE 5.5. Evolution of the hp-adaptive meshes.




FIGURE 5.7. Real part of the numerical solution of the x-component of the electric field. Notice how the DPG adaptive technology refines only in regions of the domain where there is wave activity.

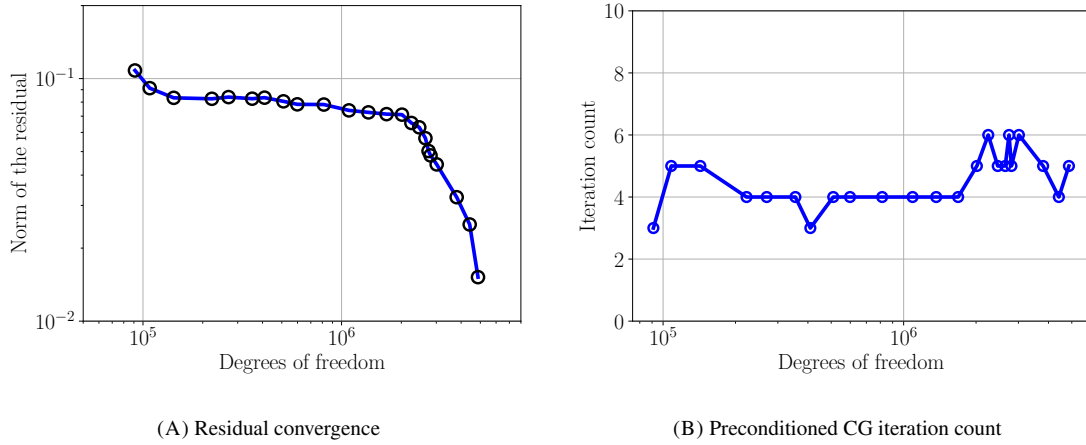


FIGURE 5.8. Convergence of the residual and the preconditioned CG solver

5.2. Linear acoustics equations. We continue with three dimensional examples for the linear acoustics problem. Recall the time harmonic form of the acoustics equations,

$$\begin{cases} i\omega p + \operatorname{div} u = f, \\ i\omega u + \nabla p = 0, \end{cases}$$

where ω is the angular frequency, p is the pressure and u is the velocity.

5.2.1. Scattering of a plane wave from a sphere. For our first example we consider the simulation of the scattered wave when a plane wave hits a rigid sphere (see Figure 5.9 (A)). The domain is meshed¹ by hexahedra as shown in Figure 5.9 (B) and the spherical surfaces are approximated using transfinite interpolation [17]. The incident plane wave is traveling in the direction $(1, 1, 1)$ with frequency $\omega = 35\pi$. This corresponds to approximately 30 wavelengths inside the domain.

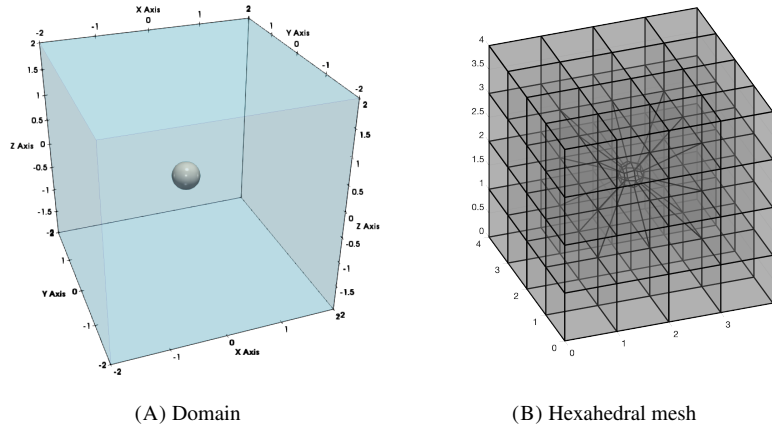


FIGURE 5.9. Computational domain with a spherical scatterer and initial mesh

¹The authors would like to thank Brendan Keith for all his help on constructing the geometry files for this simulation.

The simulation is driven by hard boundary conditions on the spherical boundary (Γ_{sph}),

$$u \cdot n = g = -u_{\text{inc}} \text{ on } \Gamma_{\text{sph}}$$

where u_{inc} is the incident plane wave. The computational domain is truncated by a homogeneous impedance condition on the outer boundary of the cube Γ_{cube} .

$$p = u \cdot n \text{ on } \Gamma_{\text{cube}}$$

We solve the problem using our multigrid technology in the uniform refinement setting. The fine grid consists of 40960 hexahedra of quartic polynomial order. This results in a linear system of approximately 14 million degrees of freedom. The coarse grid is constructed by two h -coarsening steps of the fine grid. It consists of 640 quartic hexahedra and the linear system has size of approximately 200 thousands degrees of freedom. The Conjugate Gradient algorithm starts with zero initial guess and terminates when the residual is less than 10^{-5} . We use total of 10 smoothing iterations at each level and the smoother relaxation parameter is chosen to be $\theta = 0.2$. In this setting the preconditioned Conjugate Gradient algorithm converges in 12 iterations. We emphasize that we attempted to solve the problem with a sparse multi-frontal solver but it was not possible because of high memory requirements ($\approx 350Gb$). The solution is shown in Figure 5.10 below. Note that we show the solution on the part of the domain below the plane defined by the point $(0.5, 0.5, 0.5)$ and the normal vector $(-1, -1, 2)$.




FIGURE 5.10. Scattered wave: real part of pressure

5.2.2. Scattering of a plane wave from a cube. Our second acoustics example involves scattering of a plane wave by a cube. The computational domain is $\Omega = (\frac{1}{7}, \frac{6}{7})^3 \setminus ((\frac{3}{7}, \frac{4}{7})^3)$. The incident wave has low to medium frequency of $\omega = 16\pi$, and has direction of propagation $(1, 1, 0)$. The domain is truncated by a *Perfectly Matched Layer* (PML) region of length $L = \frac{1}{7}$ in each direction. The construction of the PML is based on the work described in [50]. Additional details on how the DPG formulation is modified inside the PML region are given in [43, Appendix D].

We start the simulation with a mesh consisting of 342 cubes of side size $h = \frac{1}{7}$ and polynomial order $p = 3$. Note that the initial mesh is fine enough to control the best approximation and the pollution effect. However, the singularities on the scatterer and the exponential decay of the wave in the PML region are not resolved and therefore the quality of the solution is not good. Starting automatic h -adaptivity, we anticipate that the DPG method will perform extensive refinements in order to resolve the singularities at the corners and edges of the cubic scatterer. Additionally, some refinements are expected to occur at the transition from the computational domain into the PML region.




FIGURE 5.11. Residual and preconditioned CG convergence. Plane wave scattering from a cube. Direction of propagation: $(1,1,0)$.

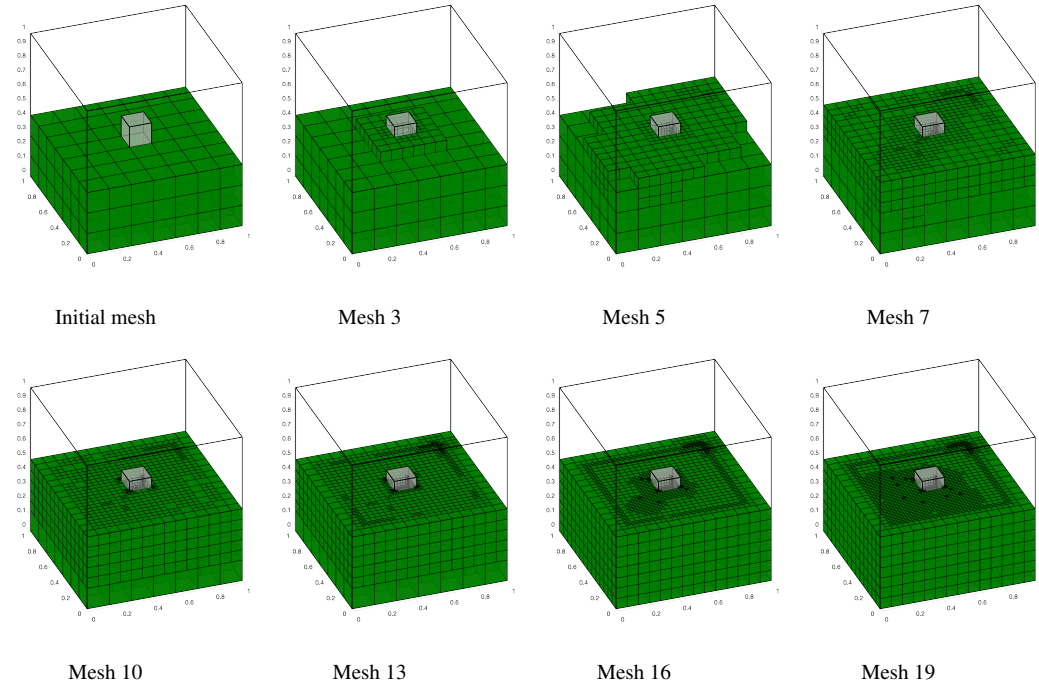


FIGURE 5.12. Wave propagating in the direction $(1,1,0)$. Evolution of the h -adaptive mesh. As expected, a lot of refinements occur in the region close to the scatterer because the singularities have to be resolved. Additional refinements occur in the PML region.

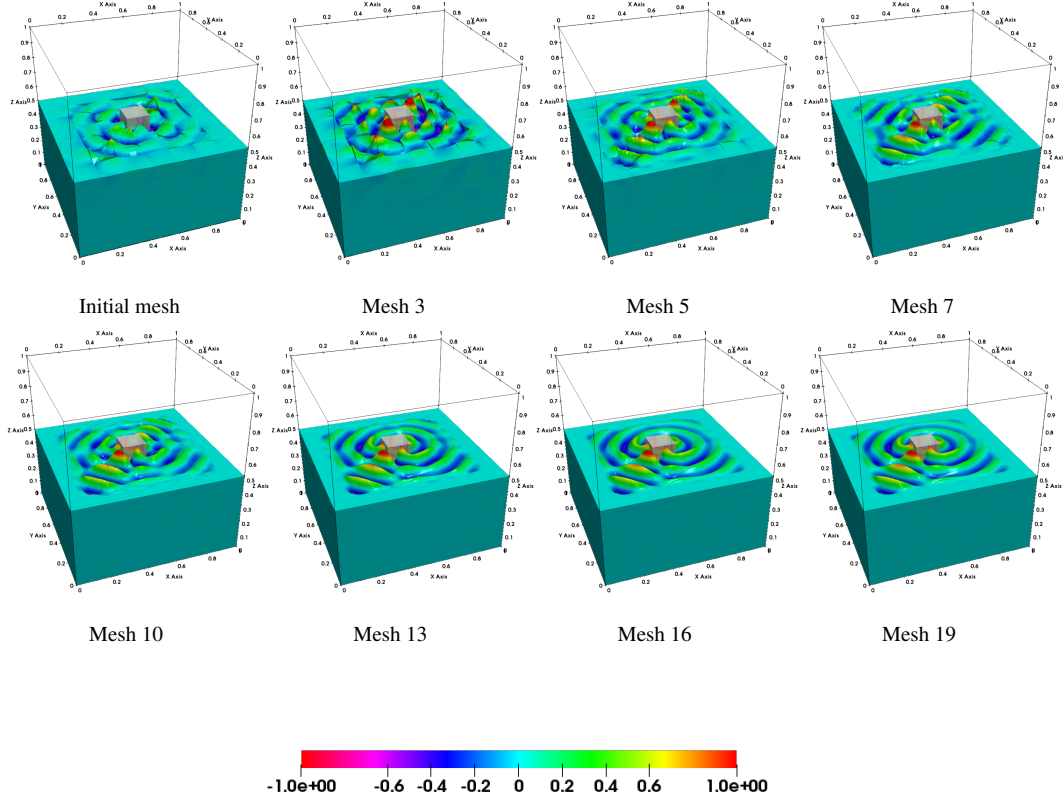


FIGURE 5.11. Wave propagating in the direction $(1,1,0)$. Evolution of the solution. The solution rapidly decays in the PML region. Notice that the quality of the solution is affected by the resolution of the singularities

The multigrid preconditioner setup is the same as in the previous examples. Here, we use at most 6 multigrid levels and the stopping criterion of the CG solver is set to 10^{-5} . The sequence of meshes along with the solutions are shown in Figure 5.12 and Section 5.2.2. Notice that the meshes and the solutions are shown only in the part of the domain where $z \leq 0.5$. As expected, refinements occur close to the singularities and in the PML region. Observe how the solution changes as the singularities are resolved. Convergence results for the DPG residual and iteration counts for the preconditioned CG solver are given in Figure 5.11. As in the previous numerical examples the solver shows robust convergence throughout the adaptive refinement process.

5.3. Scattering of a Gaussian beam from a cube. Our last experiment involves the simulation of a high frequency Gaussian beam scattering from a cube. The computational domain and the PML region are defined as in Section 5.2.2. We run the simulation with angular frequency $\omega = 140\pi$. This corresponds to about 85 wavelengths inside the computational domain. We begin the simulation with a uniform mesh consisting of 342 cubes of polynomial order $p = 3$. We then perform adaptive hp -refinements with the following strategy. An element inside the computational domain is h -refined up to the point where its maximum side size becomes smaller than half a wavelength. An element with size smaller than that is p -refined. An exception is made for the elements adjacent to the corners and edges of the cubic scatterer and the elements inside the PML region, where only h -refinements are allowed. The simulation is terminated when the DPG residual reduces by an order of magnitude.

The iterative solver setup is as follows. The initial iterate consists of the degrees of freedom corresponding to the prolongation of the solution on the previous mesh to the current mesh. We perform 10 smoothing steps at

each multigrid level and we choose the relaxation parameter to be $\theta = 0.2$. Additional computational time is saved by omitting to smooth in areas of the domain where there is no wave activity (the local patch residual is close to zero). Finally, in this simulation we use at most eight multigrid levels and the CG iterations are terminated when the l^2 -norm of the residual becomes less than 10^{-3} .

In Figures 5.13 and 5.18 we show the sequence of meshes and the corresponding numerical solutions for the real part of the pressure respectively. For demonstration purposes a partial region of the mesh is shown, constructed by the union of the regions below three planes. These planes are defined by the point $(0.5, 0.5, 0.5)$ and the normal vectors $(-0.5, -0.5, 1)$, $(0.5, -0.5, 1)$ and $(-0.5, 0.5, 1)$. Notice that the singularities at the corners and edges of the cubic scatterer have to be resolved before the wave can propagate. When the wave reaches the PML region, additional refinements occur in order to capture the steep decay. Convergence results are presented in Figure 5.19. Observe in Figure 5.19 (A) that the norm of DPG residual, which drives the adaptive refinements, tends to zero only after the decay in the PML region is resolved. Lastly, iteration counts for the CG solver preconditioned with our multigrid technology are given in Figure 5.19 (B). Notice that the number of iterations remains under control throughout the adaptive process.




FIGURE 5.13. Evolution of the hp -adaptive mesh. Notice that the singularities at the scatterer have to be resolved before the wave can propagate.




FIGURE 5.18. Evolution of the solution. Here the real part of the acoustic pressure is displayed.




FIGURE 5.19. Convergence of the DPG residual and the preconditioned CG solver. Note that the number of iterations of the iterative solver is controlled throughout the adaptive process.

6. CONCLUSION

The main accomplishment of this work is the design and implementation of an efficient and robust preconditioner for linear systems arising from DPG discretizations. The construction is heavily based on the attractive properties of the DPG method, but also on well established theory of Schwarz Domain Decomposition and multigrid methods. Special emphasis was given in the simulation of wave problems, specifically acoustics and electromagnetic simulations. As it was showcased from numerous numerical results the method is stable and reliable and it is suitable for adaptive hp -meshes. Under certain circumstances, where the coarse grid is “good” enough, uniform convergence with respect to the frequency, polynomial order and discretization size can be achieved. Integrating the iterative solver with the adaptive refinement procedure allowed us to solve efficiently, problems in the high-frequency regime, problems which under a uniform-mesh setting would be intractable. Additionally, we presented some theoretical results for the overlapping Schwarz smoother, and demonstrated its dependence on the frequency for the acoustics problem in two-space dimensions.

Our ongoing work focuses on a new implementation of the solver for distributed memory computer architectures. Our aim is to be able to solve high-frequency electromagnetic problems, arising from the simulation of optical laser amplifiers [32, 40], problems which require the sufficient resolution of waves of hundred of thousands of wavelengths. Since direct solvers fail to scale efficiently in modern many-core architectures, an efficient iterative solution scheme seems to be the only promising alternative choice. Finally, our theoretical analysis on the one-level additive preconditioner for the acoustics problem can serve as a stepping stone on a further study in the multilevel setting. While several other theoretical works exist on preconditioning DPG systems, analytical results on preconditioning wave operators are yet to be explored.

ACKNOWLEDGMENTS

This work has been supported with grants by the Air Force Office of Scientific Research (AFOSR grant FA9550-12-1-0484 and FA9550-17-1-0090). The authors would also like to express their gratitude to Dr. Jay Gopalakrishnan for his suggestions and his insightful comments throughout this work.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, LLNL-JRNL-815524. This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

APPENDIX A. COMPUTING THE INTERPOLATION NORM

The continuity constant $C_{\hat{\Pi}} = \|\hat{\Pi}\|_{\hat{U}}$ can be computed using the definition of a norm of an operator, i.e., by solving the maximization problem:

$$\|\hat{\Pi}\|_{\hat{U}} = \max_{\hat{v} \in \hat{U}, \hat{v} \neq 0} \frac{\|\hat{\Pi}\hat{v}\|_{\hat{U}}}{\|\hat{v}\|_{\hat{U}}}$$

This leads to the generalized eigenvalue problem

$$(\hat{\Pi}\hat{v}, \hat{\Pi}\hat{\delta v})_{\hat{U}} = \lambda^2 (\hat{v}, \hat{\delta v})_{\hat{U}}, \quad \hat{\delta v} \in \hat{U}.$$

Consider now a discrete basis $\{\hat{v}_i\}_{i=1}^n$ for the polynomial subspace $\hat{U}_h \subset \hat{U}$. Then, we need to solve

$$\mathbf{P}^* \mathbf{G} \mathbf{P} \mathbf{v} = \lambda^2 \mathbf{G} \mathbf{v}$$

where $\mathbf{v} \in \mathbb{C}^n$, \mathbf{P} is the matrix representation of $\hat{\Pi}$ and $\mathbf{G} = (\hat{v}_i, \hat{v}_j)_{\hat{U}}$ is the Gram matrix corresponding to the inner product $(\cdot, \cdot)_{\hat{U}}$. The continuity constant is therefore given by the square root of the maximum generalized eigenvalue of (\mathbf{A}) .

Computation of matrix \mathbf{G} . Given a basis $\{\hat{v}_i\}_{i=1}^n$ of the polynomial subspace $\hat{U}_h \subset \hat{U}$ we can compute the entries of the Gram matrix by using the $(\cdot, \cdot)_{H_A}$ inner product (see Equation (2.6)). Indeed, using the polarization formula [42, Ch. 6], it is easy to see that

$$(A.1) \quad (\hat{v}_i, \hat{v}_j)_{\hat{U}} = (v_i, v_j)_{H_A}$$

where $v_i, v_j \in H_A(K)$ are the minimum energy extensions of \hat{v}_i, \hat{v}_j respectively. However, in practice we can only approximate the minimum energy extensions using polynomial extensions. We note that for all the numerical results presented in Section 3, polynomials of sufficiently large order² were used.

Approximating the minimum energy extension. Given $\hat{v} \in \hat{U}_h$ we can derive the extension $v \in H_A(K)$ that realizes the minimum energy by:

$$\|\hat{v}\|_{\hat{U}} = \inf_{tr v = \hat{v}} \|v\|_{H_A}$$

Consider the polynomial subspace $H_{A,h}(K) \subset H_A(K)$. Then the above minimization problem leads to the following Dirichlet problem.

$$\begin{cases} \text{Find } v_h \in H_{A,h}(K) \\ (v_h, \delta v_h)_{H_A} = 0, & \delta v_h \text{ in } H_{A,h}(K) \\ v_h = \hat{v}, & \text{on } \partial K \end{cases}$$

or equivalently

$$\begin{cases} \text{Find } v_b \in H_{A,h}^b(K) \text{ (bubble functions)} \\ (Av^b, A\delta v^b) + (v^b, \delta v^b) = -(\hat{v}, \delta v^b)_{H_A}, \quad \forall \delta v^b \text{ in } H_{A,h}^b(K) \end{cases}$$

Definition of the interpolation operator Π - computation of matrix \mathbf{P} . Let $\hat{v} \in \hat{U}_h(K)$. We define the projection based interpolant $\hat{v}^p = \hat{\Pi}\hat{v}$ using the following steps:

- Interpolation at vertices: the interpolant \hat{v}^p matches the function \hat{v} at the vertices

$$\hat{v}^p(a) = \hat{v}(a), \quad \forall \text{ vertex } a.$$

We lift the vertex values using a polynomial extension that lives in the element trace space. This leads to the linear interpolant $\hat{v}_1 \in tr\mathcal{P}^1(K)$.

- Edge projection: We subtract the linear interpolant \hat{v}_1 from the function \hat{v} . Now the difference $\hat{v} - \hat{v}_1$ vanishes at the element vertices. Then we project the difference onto the trace space of edge polynomials of order p_e vanishing at the vertices (i.e, the edge bubbles $\mathcal{P}_0^{p_e}$), i.e.,

$$\begin{cases} \hat{v}_{2,e} \in \mathcal{P}_0^{p_e} \\ \|\hat{v} - \hat{v}_1 - \hat{v}_{2,e}\|_{\hat{U}(e)} \rightarrow \min. \end{cases}$$

Here, the norm $\|\cdot\|_{\hat{U}(e)}$ is defined by the inner product (A.1) on the edge. The edge interpolant is then the sum of the edge projections

$$\hat{v}_2 = \sum_e \hat{v}_{2,e}.$$

The final interpolant is defined by the sum of the vertex and the edge interpolant

$$\hat{v}^p = \hat{v}_1 + \hat{v}_2.$$

We obtain a matrix representation of \mathbf{P} by applying it to a basis of the polynomial space. The polynomial spaces are defined as [43, Ch. 2]. In particular, consider

$$W^r := \mathcal{Q}^{(r,r)} \subset H^1(\Omega), \text{ and } V^r := \mathcal{Q}^{(r,r-1)} \times \mathcal{Q}^{(r-1,r)} \subset H(\text{div}, \Omega).$$

²No notable change in the numerical results could be observed when increasing further the polynomial order of approximation

Then $\hat{\Pi} : W^{r+1} \times V^{r+1} \rightarrow W^r \times V^r$, and its matrix representation P is computed as follows. For $\hat{p} \in tr|_{\partial K} W^{r+1}$, we have vertex shape functions and bubble shape functions. For the vertex shape functions the interpolation operator matches the values at the vertices and so the first 4 columns of matrix P are the orthonormal vectors $\{e_i\}_{i=1}^4$. Note that if \hat{p} is an edge bubble then it vanishes on the vertices. Therefore it is enough to solve the following minimization problem. Let $\hat{p} \in tr|_{\partial K} \mathcal{W}_0^{r+1}$ be an edge bubble vanishing on the boundary. Then we solve

$$\begin{cases} \text{Find } \hat{v}^b \in tr|_{\partial K} \mathcal{W}_0^r \\ \|\hat{p} - \hat{v}^b\|_{\hat{U}} \rightarrow \min \end{cases}$$

or equivalently

$$\begin{cases} \text{Find } \hat{v}^b \in tr|_{\partial K} \mathcal{W}_0^r \\ (\hat{v}^b, \delta \hat{v}^b)_{\hat{U}} = (\hat{p}, \delta \hat{v}^b)_{\hat{U}} \quad \forall \delta \hat{v}^b \in tr|_{\partial K} \mathcal{W}_0^r \end{cases}$$

or

$$(A.2) \quad \begin{cases} \text{Find } v^b \in \mathcal{W}_0^r \\ (v^b, \delta v^b)_{H_A} = (p^b, \delta v^b)_{H_A} \quad \forall \delta v^b \in \mathcal{W}_0^r \end{cases}$$

where $p^b, v^b, \delta v^b$ are the minimum energy polynomial extensions of $\hat{p}, \hat{v}^b, \delta \hat{v}^b$ respectively. Finally the interpolant is defined to be the trace of the solution v^b . Notice that both left and right hand sides of (A.2) can be computed using the Gram matrix G . For the edge bubbles of the variable $\hat{u}_n \in tr|_{\partial K} V^{r+1}$, we follow a similar procedure.

Computation on the master element. We can reduce the computation on the master element using appropriate scalings. Assuming that each element K is obtained by a simple scaling of the master element \bar{K} , and using the same H^1 scaling for both H^1 and $H(\text{div})$ functions we have:

$$\begin{aligned} \|A_\omega(p, u)\|_{L^2(K)}^2 &:= \int_K (|\omega p + \text{div } u|^2 + |\omega u + \nabla p|^2) dK \\ &= \int_K h^2 (|\omega \bar{p} + \frac{1}{h} \text{div } \bar{u}|^2 + |\omega \bar{u} + \frac{1}{h} \nabla \bar{p}|^2) d\bar{K} \\ &= \int_K (|\omega h \bar{p} + \text{div } \bar{u}|^2 + |\omega h \bar{u} + \nabla \bar{p}|^2) d\bar{K} \\ &=: \|A_{\omega h}(p, u)\|_{L^2(\bar{K})}^2 \end{aligned}$$

and

$$\|(p, u)\|_{L^2(K)}^2 = h^2 \|(\bar{p}, \bar{u})\|_{L^2(\bar{K})}^2$$

We can therefore perform all the computations on the master element by using the following norm:

$$\|(p, u)\|_{H_A(\omega h, h)}^2 = \|A_{\omega h}(p, u)\|_{L^2(\bar{K})}^2 + h^2 \|(\bar{p}, \bar{u})\|_{L^2(\bar{K})}^2$$

REFERENCES

- [1] P. R. AMESTOY, I. S. DUFF, J.-Y. L'EXCELLENT, AND J. KOSTER, *A fully asynchronous multifrontal solver using distributed dynamic scheduling*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 15–41 (electronic).
- [2] D. ARNOLD, R. FALK, AND R. WINTHER, *Preconditioning in $H(\text{div})$ and applications*, Math. Comp., 66 (1997), pp. 957–984.
- [3] I. BABUŠKA, *Error-bounds for finite element method*, Numer. Math., 16 (1971), pp. 322–333.
- [4] I. M. BABUŠKA AND S. A. SAUTER, *Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers?*, SIAM J. Numer. Anal., 34 (1997), pp. 2392–2423.
- [5] A. T. BARKER, S. C. BRENNER, E.-H. PARK, AND L.-Y. SUNG, *A one-level additive Schwarz preconditioner for a discontinuous Petrov–Galerkin method*, in Domain Decomposition Methods in Science and Engineering XXI, Springer, 2014, pp. 417–425.
- [6] A. T. BARKER, V. DOBREV, J. GOPALAKRISHNAN, AND T. KOLEV, *A Scalable Preconditioner for a Primal Discontinuous Petrov–Galerkin Method*, SIAM J. Sci. Comput., 40 (2018), pp. A1187–A1203.

- [7] M. BONAZZOLI, V. DOLEAN, I. G. GRAHAM, E. A. SPENCE, AND P.-H. TOURNIER, *Domain decomposition preconditioning for the high-frequency time-harmonic Maxwell equations with absorption*, arXiv preprint arXiv:1711.03789, (2017).
- [8] X.-C. CAI AND O. B. WIDLUND, *Domain decomposition algorithms for indefinite elliptic problems*, SIAM J. Sci. Statist. Comput., 13 (1992), pp. 243–258.
- [9] H. CALANDRA, S. GRATTON, X. PINEL, AND X. VASSEUR, *An improved two-grid preconditioner for the solution of three-dimensional Helmholtz problems in heterogeneous media*, Numer. Linear Algebra Appl., 20 (2013), pp. 663–688.
- [10] C. CARSTENSEN, L. DEMKOWICZ, AND J. GOPALAKRISHNAN, *Breaking spaces and forms for the DPG method and applications including Maxwell equations*, Comput. Math. Appl., 72 (2016), pp. 494–522.
- [11] C. CARSTENSEN AND F. HELLWIG, *Low-Order Discontinuous Petrov–Galerkin Finite Element Methods for Linear Elasticity*, SIAM J. Numer. Anal., 54 (2016), pp. 3388–3410.
- [12] Z. CHEN AND X. XIANG, *A source transfer domain decomposition method for Helmholtz equations in unbounded domain*, SIAM J. Numer. Anal., 51 (2013), pp. 2331–2356.
- [13] L. DEMKOWICZ, *Computing with hp-adaptive finite elements. Vol. 1*, Chapman & Hall/CRC Applied Mathematics and Nonlinear Science Series, Chapman & Hall/CRC, Boca Raton, FL, 2007. One and two dimensional elliptic and Maxwell problems, With 1 CD-ROM (UNIX).
- [14] L. DEMKOWICZ, *Various Variational Formulations and Closed Range Theorem*, ICES Report, 15-03 (2015).
- [15] L. DEMKOWICZ AND J. GOPALAKRISHNAN, *A class of discontinuous Petrov–Galerkin methods. II. Optimal test functions*, Numer. Methods Partial Differential Equations, 27 (2011), pp. 70–105.
- [16] L. DEMKOWICZ, J. GOPALAKRISHNAN, I. MUGA, AND J. ZITELLI, *Wavenumber explicit analysis of a DPG method for the multidimensional Helmholtz equation*, Comput. Methods Appl. Mech. Engrg., 213/216 (2012), pp. 126–138.
- [17] L. DEMKOWICZ, J. KURTZ, D. PARDO, M. PASZYŃSKI, W. RACHOWICZ, AND A. ZDUNEK, *Computing with hp-adaptive finite elements. Vol. 2*, Chapman & Hall/CRC Applied Mathematics and Nonlinear Science Series, Chapman & Hall/CRC, Boca Raton, FL, 2008. Frontiers: three dimensional elliptic and Maxwell problems with applications.
- [18] M. DRYJA AND O. B. WIDLUND, *Domain decomposition algorithms with small overlap*, SIAM J. Sci. Comput., 15 (1994), pp. 604–620.
- [19] B. ENGQUIST AND L. YING, *Sweeping preconditioner for the Helmholtz equation: moving perfectly matched layers*, Multiscale Model. Simul., 9 (2011), pp. 686–710.
- [20] O. G. ERNST AND M. J. GANDER, *Why it is difficult to solve Helmholtz problems with classical iterative methods*, in Numerical analysis of multiscale problems, vol. 83 of Lect. Notes Comput. Sci. Eng., Springer, Heidelberg, 2012, pp. 325–363.
- [21] F. FUENTES, B. KEITH, L. DEMKOWICZ, AND P. LE TALLEC, *Coupled variational formulations of linear elasticity and the DPG methodology*, J. Comput. Phys., 348 (2017), pp. 715–731.
- [22] F. FUENTES, B. KEITH, L. DEMKOWICZ, AND S. NAGARAJ, *Orientation embedded high order shape functions for the exact sequence elements of all shapes*, Comput. Math. Appl., 70 (2015).
- [23] I. G. GRAHAM, E. A. SPENCE, AND E. VAINIKKO, *Domain decomposition preconditioning for High-Frequency Helmholtz problems with absorption*, Math. Comp., 86 (2015).
- [24] M. J. GANDER, I. G. GRAHAM, AND E. A. SPENCE, *Applying GMRES to the Helmholtz equation with shifted Laplacian preconditioning: what is the largest shift for which wavenumber-independent convergence is guaranteed?*, Numer. Math., 131 (2015), pp. 567–614.
- [25] M. J. GANDER AND H. ZHANG, *A class of iterative solvers for the Helmholtz equation: Factorizations, sweeping preconditioners, source transfer, single layer potentials, polarized traces, and optimized Schwarz methods*, Siam Review, 61 (2019), pp. 3–76.
- [26] J. GOPALAKRISHNAN, I. MUGA, AND N. OLIVARES, *Dispersive and Dissipative Errors in the DPG Method with Scaled Norms for Helmholtz Equation*, SIAM J. Sci. Comput., 36 (2014), pp. A20–A39.
- [27] J. GOPALAKRISHNAN AND J. E. PASCIAC, *Overlapping Schwarz preconditioners for indefinite time harmonic Maxwell equations*, Math. Comp., 72 (2003), pp. 1–15 (electronic).
- [28] J. GOPALAKRISHNAN, J. E. PASCIAC, AND L. F. DEMKOWICZ, *Analysis of a multigrid algorithm for time harmonic Maxwell equations*, SIAM J. Numer. Anal., 42 (2004), pp. 90–108 (electronic).
- [29] J. GOPALAKRISHNAN AND W. QIU, *An analysis of the practical DPG method*, Math. Comp., 83 (2014), pp. 537–552.
- [30] J. GOPALAKRISHNAN AND J. SCHÖBERL, *Degree and wavenumber [in] dependence of Schwarz preconditioner for the DPG method*, in Spectral and High Order Methods for Partial Differential Equations ICOSAHOM 2014, Springer, 2015, pp. 257–265.
- [31] S. HENNEKING AND L. DEMKOWICZ, *A numerical study of the pollution error and DPG adaptivity for long waveguide simulations*, CAMWA, (2020).
- [32] S. HENNEKING, J. GROSEK, AND L. DEMKOWICZ, *A dpg maxwell approach for studying nonlinear thermal effects in active gain fiber amplifiers*, arXiv preprint arXiv:1912.01185, (2019).
- [33] R. HIPTMAIR AND R. H. HOPPE, *Multilevel methods for mixed finite elements in three dimensions*, Numer. Math., 82 (1999), pp. 253–279.
- [34] B. KEITH, F. FUENTES, AND L. DEMKOWICZ, *The DPG methodology applied to different variational formulations of linear elasticity*, Comput. Methods Appl. Mech. Engrg., 309 (2016), pp. 579–609.
- [35] B. KEITH, S. PETRIDES, F. FUENTES, AND L. DEMKOWICZ, *Discrete least-squares finite element methods*, Comput. Methods Appl. Mech. Engrg., 327 (2017), pp. 226–255.

- [36] S. KIM AND H. ZHANG, *Optimized Schwarz method with complete radiation transmission conditions for the Helmholtz equation in waveguides*, SIAM J. Numer. Anal., 53 (2015), pp. 1537–1558.
- [37] W. LENG AND L. JU, *A diagonal sweeping domain decomposition method with source transfer for the Helmholtz equation*, arXiv preprint arXiv:2002.05327, (2020).
- [38] X. LI AND X. XU, *Domain decomposition preconditioners for the discontinuous Petrov-Galerkin method*, ESAIM Math. Model. Numer. Anal., 51 (2017), pp. 1021–1044.
- [39] J. W. H. LIU, *The multifrontal method for sparse matrix solution: Theory and practice*, SIAM Rev., 34 (1992), pp. 82–109.
- [40] S. NAGARAJ, J. GROSEK, S. PETRIDES, L. F. DEMKOWICZ, AND J. MORA, *A 3D DPG Maxwell approach to nonlinear Raman gain in fiber laser amplifiers*, Journal of Computational Physics: X, 2 (2019), p. 100002.
- [41] S. NAGARAJ, S. PETRIDES, AND L. F. DEMKOWICZ, *Construction of DPG Fortin operators for second order problems*, Comput. Math. Appl., 74 (2017), pp. 1964–1980.
- [42] J. ODEN AND L. DEMKOWICZ, *Applied Functional Analysis, Third Edition*, Textbooks in Mathematics, CRC Press, Taylor & Francis Group, 2018.
- [43] S. PETRIDES, *Adaptive multilevel solvers for the discontinuous Petrov-Galerkin method with an emphasis on high-frequency wave propagation problems*, PhD thesis, The University of Texas at Austin, 2019.
- [44] S. PETRIDES AND L. F. DEMKOWICZ, *An adaptive DPG method for high frequency time-harmonic wave propagation problems*, Comput. Math. Appl., 74 (2017), pp. 1999–2017.
- [45] N. V. ROBERTS AND J. CHAN, *A geometric multigrid preconditioning strategy for DPG system matrices*, Comput. Math. Appl., 74 (2017), pp. 2018 – 2043.
- [46] A. H. SHEIKH, D. LAHAYE, AND C. VUIK, *On the convergence of shifted Laplace preconditioner combined with multilevel deflation*, Numer. Linear Algebra Appl., 20 (2013), pp. 645–662.
- [47] C. C. STOLK, *A rapidly converging domain decomposition method for the Helmholtz equation*, J. Comput. Phys., 241 (2013), pp. 240–252.
- [48] C. C. STOLK, M. AHMED, AND S. K. BHOWMIK, *A multigrid method for the Helmholtz equation with optimized coarse grid corrections*, SIAM J. Sci. Comput., 36 (2014), pp. A2819–A2841.
- [49] M. TAUS, L. ZEPEDA-NÚÑEZ, R. J. HEWETT, AND L. DEMANET, *L-Sweeps: A scalable, parallel preconditioner for the high-frequency Helmholtz equation*, J. Comput. Phys., 420 (2020), p. 109706.
- [50] A. VAZIRI ASTANEH, B. KEITH, AND L. DEMKOWICZ, *On perfectly matched layers for discontinuous Petrov-Galerkin methods*, Comput. Mech., (2018).
- [51] A. VION AND C. GEUZAIN, *Double sweep preconditioner for optimized Schwarz methods applied to the Helmholtz problem*, J. Comput. Phys., 266 (2014), pp. 171–190.
- [52] J. XU, *Iterative methods by space decomposition and subspace correction*, SIAM Rev., 34 (1992), pp. 581–613.
- [53] ———, *An introduction to multilevel methods*, in Wavelets, multilevel methods and elliptic PDEs (Leicester, 1996), Numer. Math. Sci. Comput., Oxford Univ. Press, New York, 1997, pp. 213–302.
- [54] ———, *The method of subspace corrections*, J. Comput. Appl. Math., 128 (2001), pp. 335–362. Numerical analysis 2000, Vol. VII, Partial differential equations.
- [55] L. ZEPEDA-NÚÑEZ AND L. DEMANET, *The method of polarized traces for the 2D Helmholtz equation*, J. Comput. Phys., 308 (2016), pp. 347–388.
- [56] J. ZITELLI, I. MUGA, L. DEMKOWICZ, J. GOPALAKRISHNAN, D. PARDO, AND V. M. CALO, *A class of discontinuous Petrov-Galerkin methods. Part IV: the optimal test norm and time-harmonic wave propagation in 1D*, J. Comput. Phys., 230 (2011), pp. 2406–2432.