

Scalable DPG multigrid solver for Helmholtz problems: A study on convergence

Jacob Badger^{a,*}, Stefan Henneking^a, Socratis Petrides^b, Leszek Demkowicz^a

^a Oden Institute, The University of Texas at Austin, United States of America

^b Lawrence Livermore National Laboratory, United States of America

ARTICLE INFO

Dedicated to Leszek Demkowicz on the occasion of his 70th birthday

Keywords:

Multigrid
Discontinuous Petrov-Galerkin
Helmholtz
High-frequency wave propagation
hp-adaptivity

ABSTRACT

This paper presents a scalable multigrid preconditioner targeting large-scale systems arising from discontinuous Petrov–Galerkin (DPG) discretizations of high-frequency wave operators. This work is built on previously developed multigrid preconditioning techniques of Petrides and Demkowicz (Comput. Math. Appl. 87 (2021) pp. 12–26) and extends the convergence results from $\mathcal{O}(10^7)$ degrees of freedom (DOFs) to $\mathcal{O}(10^9)$ DOFs using a new scalable parallel MPI/OpenMP implementation. Novel contributions of this paper include an alternative definition of coarse-grid systems based on restriction of fine-grid operators, yielding superior convergence results. In the uniform refinement setting, a detailed convergence study is provided, demonstrating h and p robust convergence and linear scaling with respect to the wave frequency. The paper concludes with numerical results on *hp*-adaptive simulations including a large-scale seismic modeling benchmark problem with high material contrast.

Acknowledgments

We thank Matteo Croci for helpful discussions. J. Badger, S. Henneking, and L. Demkowicz were supported with AFOSR grant FA9550-19-1-0237 and NSF award 2103524. 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-847116.

1. Introduction

1.1. Background

Wave propagation problems arise in a number of contexts including natural resource exploration, medical imaging, and nuclear fusion research, to name a few. However, developing accurate and efficient numerical algorithms for the solution of time-harmonic wave propagation problems is a notoriously difficult problem. While traditional finite element methods (FEM) can deliver high-accuracy and optimal discretizations, their efficacy for wave operators deteriorates for two main reasons. First, they suffer from stability issues unless very fine meshes are used to resolve the propagating wave. In the high-frequency regime this results in prohibitively expensive problems. The lack of

preasymptotic discrete stability also makes mesh adaptivity techniques unreliable and inefficient. Second, the linear system is highly indefinite and, consequently, standard iterative solution schemes break down [26]. Current leading-edge preconditioning techniques for wave operators, such as multigrid methods [43,56,34], domain decomposition methods with special transmission conditions [7,31,55,44], stabilized methods based on artificial absorption [27,6], shifted Laplacian [53] and sweeping preconditioners [23,24,11,46,57] are promising but they lose their efficiency in heterogeneous media and can be difficult to apply on complex geometries [29,25].

An alternative approach instead employs *minimum-residual* discretization methodologies which, by construction, produce positive-definite discrete systems and may therefore be amenable to more standard preconditioning techniques [33,49,50,48]. Indeed, popularization of the first-order system least-squares methodology (FOSLS) [8,45], and other least-squares methodologies [12], was driven by the applicability of geometric and algebraic multigrid methods to otherwise indefinite problems. However, for wave propagation problems, FOSLS is known to be highly dissipative [30] and thus not competitive in the high-frequency regime. This work discusses a multigrid solver based on a minimum-residual discretization obtained by the discontinuous Petrov–

* Corresponding author.

E-mail address: jcbadger@utexas.edu (J. Badger).

<https://doi.org/10.1016/j.camwa.2023.07.006>

Galerkin (DPG) method with Optimal Test Functions [18] applied to the ultraweak variational formulation.

1.2. DPG-MG solver

The DPG FE methodology of Demkowicz and Gopalakrishnan [17, 16,19] is a non-standard least-squares method with several attractive properties: mesh-independent stability, a built-in error indicator, and applicability to a number of variational formulations with different functional settings. A special case of the DPG method is the well-established FOSLS method in which the residual is minimized in the L^2 test norm. As mentioned previously, however, other formulations are preferable in the context of wave propagation. Among the various DPG formulations, the so-called *ultraweak* variational formulation has proved to be superior: it is less dissipative than other DPG formulations [48], with dispersion error roughly commensurate to Galerkin discretizations [30], and it has been shown to solve problems with many wavelengths accurately by countering the pollution error through a modest increase in the order of discretization [37]. These properties were leveraged by Petrides and Demkowicz in [50] to define an hp -adaptive multilevel preconditioner for DPG wave propagation problems discretized with conforming elements of the exact-sequence energy spaces [15].

Similar to hybridizable methods, the DPG methodology introduces additional trace degrees of freedom (DOFs) on the mesh skeleton resulting from testing with larger discontinuous¹ (“broken”) test spaces [10]. In the case of high-order discretizations, statically condensing all interior DOFs onto the mesh skeleton results in a smaller global system and enables more coherent implementation for field and trace variables. The DPG multigrid solver (DPG-MG) is defined on this condensed global system of trace degrees of freedom. Constructing suitable prolongation operators for the condensed system is complicated by the fact that fine-grid DOFs resulting from h -refinement have no natural coarse-grid representatives; this is a challenge shared by hybridizable methods [52,51]. The construction of a stable prolongation operator between such non-nested condensed systems for general DPG problems² was one of the major contributions in the original DPG-MG work by Petrides and Demkowicz [49,48] and it will be outlined later in Section 3.

1.3. Direction and outline

Based on the initial implementation by Petrides and Demkowicz [48], we have developed a scalable hybrid MPI/OpenMP implementation of the DPG-MG solver. The parallel implementation and its scaling characteristics will be detailed in a forthcoming publication; the present work instead leverages our performant implementation to study the convergence properties of the DPG-MG solver under uniform h , uniform p , and hp -adaptive refinements. This work is intended to elucidate scaling characteristics of the DPG-MG solver and identify aspects of the current construction which may be improved.

The remainder of the paper is structured as follows: Section 2 defines the ultraweak acoustics model and DPG discretization used throughout this work. Section 3 reviews the construction of the DPG-MG solver. In Section 4, a number of convergence studies are performed for a model problem with manufactured solution. The applicability of the solver to state-of-the-art computational challenges is demonstrated in Section 5, using the GO_3D_OBS model [35], a challenging benchmark problem for evaluating next-generation algorithms in seismic modeling. We conclude in Section 6 with a discussion of findings and future work.

2. Ultraweak DPG for Helmholtz

The DPG method constructs automatically stable discretizations of well-posed variational formulations, inheriting stability from the con-

tinuous problem. It achieves this by computing optimal test functions [16] that realize the supremum in the discrete inf-sup condition [3]. The unique space of these specially-selected test functions is called the optimal test space. In practice, this space is approximated by inverting the global Riesz map over an enriched, discontinuous test space [32]. The numerical computations in this paper employ a uniform increase of the polynomial order by 1 for the enrichment of the test space. The discontinuous (broken) nature of the test space enables the element-local computation of optimal test functions. However, this breaking of the test space results in additional (trace) unknowns defined on the mesh skeleton [10].

Notation and energy spaces We briefly introduce some notation and the energy spaces used throughout this work. Consider a bounded domain $\Omega \subset \mathbb{R}^3$ with Lipschitz boundary $\Gamma \equiv \partial\Omega$. The L^2 -inner product over Ω is denoted by $(\cdot, \cdot)$ and the L^2 -norm by $\|\cdot\|$. We define the standard energy spaces

$$\begin{aligned} L^2(\Omega) &= \{y : \Omega \rightarrow \mathbb{C} : \|y\| < \infty\}, \\ H^1(\Omega) &= \{w : \Omega \rightarrow \mathbb{C} : w \in L^2(\Omega), \nabla w \in (L^2(\Omega))^3\}, \\ H(\text{div}, \Omega) &= \{v : \Omega \rightarrow \mathbb{C}^3 : v \in (L^2(\Omega))^3, \nabla \cdot v \in L^2(\Omega)\}. \end{aligned} \quad (2.1)$$

In the DPG method, we use corresponding broken energy spaces for test functions which are defined as product-spaces over elements $\{K\}_{K \in \Omega_h}$ of the finite element mesh Ω_h :

$$\begin{aligned} H^1(\Omega_h) &:= \{w : \Omega \rightarrow \mathbb{C} : w|_K \in H^1(K) \forall K \in \Omega_h\}, \\ H(\text{div}, \Omega_h) &:= \{v : \Omega \rightarrow \mathbb{C}^3 : v|_K \in H(\text{div}, K) \forall K \in \Omega_h\}. \end{aligned} \quad (2.2)$$

Lastly, the breaking of test functions [10] leads to introducing trace unknowns on the mesh skeleton $\Gamma_h := \{\partial K\}_{K \in \Omega_h}$. The trace spaces are understood as element-wise traces of globally conforming functions:

$$\begin{aligned} H^{1/2}(\Gamma_h) &:= \left\{ \prod_{K \in \Omega_h} \gamma^K(w|_K) : w \in H^1(\Omega) \right\}, \\ H^{-1/2}(\Gamma_h) &:= \left\{ \prod_{K \in \Omega_h} \gamma_n^K(v|_K) : v \in H(\text{div}, \Omega) \right\}, \end{aligned} \quad (2.3)$$

where γ^K and γ_n^K are element-wise continuous and normal trace operators [18,14].

Helmholtz problem This work considers the first-order mixed form of time-harmonic linear acoustics with inhomogeneous impedance boundary condition (BC). In operator form, the equations are given by

$$\begin{aligned} i\omega p + \nabla \cdot u &= 0 & \text{in } \Omega, \\ i\omega u + \nabla p &= 0 & \text{in } \Omega, \\ Z^{-1}p - u_n &= u_0 & \text{on } \Gamma, \end{aligned} \quad (2.4)$$

where p is pressure, u is velocity, ω is the angular wave frequency, and $i = \sqrt{-1}$; in the impedance BC, Z is acoustic wave impedance, $u_n := u \cdot n$ is the flux in outward normal direction n , and u_0 is an impedance load.

Broken ultraweak formulation Let $(U, \hat{U})$ and $\mathcal{V}$ be the trial and test space, respectively, and let $\mathcal{V}'$ be the space of antilinear functionals on $\mathcal{V}$. The DPG formulation of the Helmholtz problem is defined by a variational formulation of the form: Given $l \in \mathcal{V}'$, find $u \in U$ and $\hat{u} \in \hat{U}$ that satisfy

$$b(u, v) + \hat{b}(\hat{u}, v) = l(v), \quad v \in \mathcal{V}, \quad (2.5)$$

where b and $\hat{b}$ are sesquilinear forms on $U \times \mathcal{V}$ and $\hat{U} \times \mathcal{V}$, respectively.

We refer to [48,15] for a thorough derivation of the ultraweak Helmholtz formulation. The broken ultraweak formulation, given by (2.5), is defined by the following group variables and forms:

¹ The ‘D’ in the DPG name.

² discretized with exact-sequence energy spaces.

$$\begin{aligned}
\mathbf{u} &= (u, p) \in (L^2(\Omega))^3 \times L^2(\Omega), \\
\hat{\mathbf{u}} &= (\hat{u}_n, \hat{p}) \in H^{-\frac{1}{2}}(\Gamma_h) \times H^{\frac{1}{2}}(\Gamma_h) : Z^{-1}\hat{p} - \hat{u}_n = 0 \text{ on } \Gamma, \\
\mathbf{v} &= (q, v) \in H^1(\Omega_h) \times H(\text{div}, \Omega_h) : Z^{-1}q + v_n = 0 \text{ on } \Gamma, \\
b(\mathbf{u}, \mathbf{v}) &= (i\omega p, q) - (u, \nabla_h q) + (i\omega v, v) - (p, \nabla_h \cdot v) \\
\hat{b}(\hat{\mathbf{u}}, \hat{\mathbf{v}}) &= \langle \hat{u}_n, q \rangle_{\Gamma_h} + \langle \hat{p}, v_n \rangle_{\Gamma_h}, \\
l(\mathbf{v}) &= \langle u_0, q \rangle_{\Gamma}.
\end{aligned} \tag{2.6}$$

Note that the test functions are assumed to satisfy $Z^{-1}q + v_n = 0$ on Γ in order to build in the impedance BC, and that the impedance BC implicitly implies additional regularity of the velocities on boundary Γ . For the load to be well-defined, we can assume $u_0 \in H^{-1/2}(\Gamma)$ in which case $\langle u_0, q \rangle_{\Gamma}$ can be understood in the sense of duality pairing; another option is to assume $u_0 \in L^2(\Gamma)$ interpreting $\langle u_0, q \rangle_{\Gamma}$ in the L^2 -sense; see [15, Section 5.5] for further discussion on the regularity issue for impedance BCs.

The additional unknowns $\hat{u}_n, \hat{p}$ describe the normal velocity (i.e., flux) and the pressure on element boundaries on the mesh skeleton Γ_h ; $\hat{u}_n$ is discretized as the normal trace of $H(\text{div})$ -conforming elements, and $\hat{p}$ as the continuous trace of H^1 -conforming elements. The broken test space is equipped with the adjoint graph norm [50,18]:

$$\|\mathbf{v}\|_{\mathbf{v}}^2 := \|A_h^* \mathbf{v}\|^2 + \alpha \|\mathbf{v}\|^2, \tag{2.7}$$

where $A_h^* \mathbf{v} = -(i\omega q + \nabla_h \cdot v, \nabla_h q + i\omega v)$, and α is a scaling constant. Throughout this work, numerical results are computed with $\alpha = 1$.

3. DPG-MG solver

DPG-MG is a multigrid-preconditioned conjugate gradient solver, with the multilevel preconditioner defined on a hierarchy of meshes produced through refinement as defined in [50]. Mesh-independent stability of the DPG methodology implies that the DPG-MG solver can be initialized on arbitrarily coarse initial meshes. Once the solution is obtained to a sufficient accuracy on the current mesh, the DPG error indicator is used to produce a set of refinements to define the next mesh. Because the solution on intermediate meshes is needed only to sufficient accuracy to produce the next mesh, optimal hp -adaptive meshes can be produced with relatively few iterations and at little cost.

Prolongation As indicated in Section 1, the DPG-MG solver is defined on trace DOFs located on the mesh skeleton; h -refinements produce fine-grid edges and faces that do not coincide with the previous-grid skeleton and thus have no natural representatives on the previous mesh. To ameliorate this, Petrides and Demkowicz introduced a two-stage prolongation [49]. In the first stage of the restriction, fine-grid DOFs not supported on the previous mesh skeleton are statically condensed; the resulting mesh is called the *macro* grid.

The second stage of the restriction operator is defined as the transpose of the natural inclusion operator, which expresses previous-grid basis functions as a linear combination of macro-grid basis functions. The natural inclusion operator can be constructed by solving a series of projection problems, projecting basis functions on macro-grid edges and faces onto basis functions defined on the previous-grid mesh skeleton. Efficient implementation of these projection problems can be achieved via constrained approximation [20]. Note that in the two-grid setting of the original work [49], the fine grid was constructed using multiple refinements of the coarse grid; thus prolongation was defined via recursive application of constrained approximation. In the multigrid setting, we assume that elements are h -refined at most one time between grid levels. Certain complex cases, including anisotropic h -refinements, can still require the recursive definition but are not considered in this work.

V-cycle and smoother We briefly illustrate the DPG-MG V-cycle and introduce terminology used to refer to the various meshes involved in the algorithm. First, we depart somewhat from convention and refrain

from designating the *initial* mesh the coarse grid (we refer to it simply as the initial mesh). Therefore, terms like *coarse-grid* vertex patches and *coarse-grid* stiffness matrices are not necessarily related to the initial mesh. Instead, we reserve the terms *coarse* and *fine* grid to indicate the coarser and finer meshes, respectively, in a pair of consecutive meshes. In particular, `fine(i)` refers to the grid produced by the i -th refinement of the initial mesh with `coarse(i)` being its coarse counterpart. This naming convention is illustrated along with the V-cycle in Fig. 1. Note that in the MPI-distributed implementation, discussed below, `coarse(i)` and `fine(i-1)` correspond to two different partitions of the same mesh.

The DPG-MG solver applies both conjugate gradient iterations and smoothing on the macro grid. Smoothing is performed using an additive Schwarz (overlapping block-Jacobi) smoother with blocks defined as macro-grid DOFs supported on coarse-grid vertex patches. Smoothing patches on `fine(i)` are constructed using `coarse(i)` vertex patches. Use of vertex patches can lead to large smoothing blocks, especially in the case of high-order discretizations, but avoids additional complexities for vector-valued variables set in $H(\text{curl})$ and $H(\text{div})$ [2]. Thus, the DPG-MG solver is applicable, without modification, to any well-posed DPG problem discretized with exact-sequence energy spaces. Alternative definitions of smoothing patches lead to smaller patch sizes for vector-valued variables [42,43].

Alternative construction of coarse-grid operators Coarse-grid operators can be constructed either by direct matrix assembly, or by restricting fine-grid operators. Note that because of the use of trace spaces and the *on-the-fly* computation of the DPG optimal test space, these two approaches are not equivalent (indeed, in a non-nested multigrid iteration, as is the case here, these approaches are generally not equivalent). The DPG-MG solver assembles and solves the system on the current mesh, then refines to define the next mesh. Thus, for linear problems, direct assembly of coarse-grid operators can be accomplished by simply storing the current-grid system before refinement; this was the approach taken in the original DPG-MG implementation [50]. However, we observed that for high-frequency problems, convergence rapidly deteriorated with increasing frequency. As will be demonstrated in Section 4, computing coarse-grid stiffness matrices as restrictions of fine-grid stiffness matrices restores the expected convergence. We are working to develop a rigorous understanding of how these two coarse-grid approaches differ, and thus defer an analysis of this phenomenon to a later publication. For now it will suffice to say that, for large frequencies (and fixed α), the spectrum of element Gram matrices with the adjoint graph test norm (2.7) changes dramatically as a mesh transitions from preasymptotic to asymptotic regime; this in turn may impart vastly different scales to fine- and coarse-grid systems and, when not accounted for in prolongation, may cause the coarse-grid correction to become unstable. Finally, we note that computing coarse-grid operators as restrictions of fine-grid operators is relatively inexpensive compared to assembly of the fine-grid DPG system, typically requiring between 1% and 10% of the cost of fine-grid assembly, even when sum-factorization [47,4] is used to accelerate element assembly.

Scalable MPI/OpenMP implementation in hp3D The DPG-MG solver is implemented in *hp3D*, a scalable finite element software for analysis and discretization of complex three-dimensional multiphysics applications [39,38]. *hp3D* supports a number of advanced FE technologies including exact-sequence conforming discretizations, fully anisotropic hp -adaptivity, and hybrid meshes with elements of “all shapes” (tetrahedra, hexahedra, prisms, pyramids). The code leverages hybrid MPI/OpenMP parallelism and interfaces with various scientific libraries including PETSc [5], MUMPS [1], and Zoltan [21]. *hp3D* is available as an open-source code under a BSD-3 license.³

³ <https://github.com/Oden-EAG/hp3d>.




Fig. 1. DPG-MG V-cycle. The prolongation operator is defined by first condensing all fine-grid degrees of freedom onto the coarse-grid skeleton; then applying the natural inclusion operator on macro-grid edges and, in 3D, faces.

The original DPG-MG implementation [50] employed shared-memory parallelism via OpenMP threading for single-node computation. Memory limitations of typical compute node configurations limited the scalability of the adaptive solver to $\mathcal{O}(10^7)$ DOFs. We have extended the original solver implementation to support scalable distributed-memory computation with MPI. The approach is based on distributing solution and geometry DOFs on subdomains [38,36]; however, the DPG-MG solver employs unique data structures and algorithms that extend *hp*3D's data structures to allow for asymmetric inclusion of ghost elements and enable efficient asynchronous communication with neighboring subdomains. To maintain satisfactory parallel efficiency on *hp*-adaptive meshes, dynamic load balancing is performed at each new grid level during the refinement process. For a fixed frequency, the distributed DPG-MG solver implementation has been shown to scale with near-linear parallel efficiency to $\mathcal{O}(10^9)$ DOFs on $\mathcal{O}(100)$ compute nodes. The details of the solver's parallel data structures, algorithms, and its scaling characteristics are not the focus of this paper and will instead be discussed in a future publication.

4. Convergence studies

In this section, we perform a number of convergence studies to investigate the robustness of the DPG-MG solver with respect to element size h , polynomial order p , and angular frequency ω . Previous expositions of the DPG-MG solver considered a variety of physical problems, smoothing steps, and tolerances, illustrating the versatility of the DPG-MG solver but somewhat confounding the scaling behavior. Instead, to elucidate the convergence characteristics of the DPG-MG solver, we fix the following parameters:

- Conjugate gradient iterations are terminated when the relative ℓ^2 -norm of the discrete residual has been reduced by a factor of 10^7 .
- After each refinement, the (initial) solution is reset to zero; in other words, solutions from previous grids are *not* used to generate initial guesses for following grids.

- A single pre- and post-smoothing step is performed on each grid level (V(1,1)-cycle), except in one case in Section 4.2 in which both *one* and *five* smoothing steps are employed to aid in comparison; this case will be noted.
- The initial mesh is a single element of order $p = 2$; however, iterations are not reported for the initial mesh which is solved using the MUMPS direct solver [1].
- No initial-grid solver is employed during the iteration. We have observed no effect on convergence when the initial grid is far from resolving the wave (as is the case throughout this work).


All experiments in this section were performed on *Frontera's* Cascade Lake (CLX) nodes at the Texas Advanced Computing Center [54]. Timing statistics are neglected in this section, but will be provided in Section 5.

4.1. Problem setup


Throughout this section we consider propagation of a Gaussian beam with waist-radius 0.1 and direction given by spherical angles $(\theta, \phi) = (45^\circ, 55^\circ)$ in a homogeneous unit cube domain $[0, 1]^3$. Homogeneous impedance boundary conditions are imposed on all surfaces except near the origin, where the Gaussian beam is injected through a manufactured impedance load. More precisely, let $g(\mathbf{x})$ denote the prescribed Gaussian beam solution; the impedance load was defined on boundary Γ as:

$$u_0(\mathbf{x}) = e^{-1000r^6(\mathbf{x})} (Z^{-1}g + (i\omega)^{-1}\partial_n g)$$

where the exponential term corresponds to a fast-decaying window function that localizes the load near the origin, $r(\mathbf{x}) := \sqrt{x^2 + y^2 + z^2}$ is the radial coordinate of $\mathbf{x}$, and ∂_n denotes the outward normal derivative. Equations (2.4) were solved in a non-dimensional form with (non-dimensional) acoustic impedance $Z = 1$ used throughout this work. The solution is depicted in Fig. 5.



(a) Coarse-grid system restricted from fine-grid system



(b) Coarse-grid system stored from previous meshes

Fig. 2. Convergence of the DPG-MG solver with *one smoothing step* applied to uniformly h -refined meshes. The solver convergence is h -robust and the iteration numbers are lower when using fine-grid restriction. The iterations until convergence depend linearly on the frequency ω .

The ultraweak formulation for the Helmholtz problem employed here has been added to the public *hp3D* GitHub repository⁴ under the problem directory `HELMHOLTZ/ULTRAWEAK_DPG/`; additional information on the implementation will be provided in an updated version of the user manual [39]. This precise problem setup can be run by setting the following parameters in the application script `run.sh`: `cube geometry prob=0`, impedance boundary condition `ibc=3`, and Gaussian beam exact solution `exact=5`. A public version of the scalable DPG-MG solver is planned for release to the *hp3D* repository by June 2024 but is not yet available at the time of publication.

4.2. Direct assembly vs. fine-grid restriction for coarse-grid operators

As indicated in Section 3, coarse-grid systems can be either directly assembled (or stored from previous meshes) or computed from fine-grid systems by applying the restriction operator. The two approaches are referred to as *store* and *restrict*, respectively. As will be demonstrated, the construction of coarse-grid systems has significant implications for the convergence of the DPG-MG solver.

Uniform h -refinements We begin by studying convergence of the DPG-MG solver under uniform h -refinements; i.e. each subsequent grid is produced by a uniform h -refinement of the previous grid. The number of iterations required for convergence under each of the approaches, for a variety of frequencies, is reported in Fig. 2. Examining the results in Fig. 2a (*restrict*), it can be seen that the number of iterations increases roughly linearly with frequency but demonstrates clear h -robustness in the asymptotic regime. The increase in number of iterations with frequency is expected: meshes that cannot resolve the wave do not contribute to preconditioning the operator. Note that unlike multigrid preconditioners for the standard Galerkin method, which would diverge in this setting due to lack of discrete stability on the coarse grid, the DPG solver remains stable. Next, comparing Fig. 2a (*restrict*) and Fig. 2b (*store*), it can be seen that storing the coarse-grid system consistently resulted in a larger number of iterations than restricting; additionally, storing the coarse-grid system does not demonstrate h -robustness.

Uniform h -refinements; five smoothing steps In the original implementation of the DPG-MG solver, a relatively large number of smoothing iterations (typically between 5 and 10) were used in numerical experiments. For comparison, we repeat the previous study using five smoothing steps per iteration (V(5,5)-cycle); the results are depicted

in Fig. 3. Using a large number of smoothing steps tends to restore the h -robust convergence when the coarse-grid system is stored (Fig. 3b); however, the number of smoothing steps needed to attain h -robustness tends to increase with frequency. In particular, note that for the higher-frequency cases, the number of iterations until convergence is in fact lower when using one smoothing step with restriction from fine-grid systems (Fig. 2a) than when using five smoothing steps with coarse-grid operators stored from previous meshes (Fig. 3b). Comparing Fig. 2a and Fig. 3a, it can be seen that when the coarse-grid systems are defined via restriction, increasing the number of smoothing steps by a factor of five results in a decrease of the number of smoothing iterations by only a factor of two. We neglect an explicit study of convergence in terms of the number of smoothing steps but qualitatively report that a single smoothing step per grid level is optimal for convergence in all of our numerical experiments to date.

Uniform p -refinements To investigate convergence of the DPG-MG solver under p -refinements, we first perform h -refinements until there are at least two elements per wavelength (satisfying the Nyquist criterion); then, the polynomial order of discretization p is incremented on each subsequent grid. The results of this study are shown in Fig. 4, where it can be seen that both storing and restricting leads to p -robust convergence; however, restricting again requires far fewer iterations.

4.3. hp -adaptive refinements

We now consider hp -adaptive refinements, employing the Dörfler marking strategy [22] to determine elements to be refined. Marked elements are h -refined until the maximum edge-length is less than one-half the wavelength, otherwise they are p -refined. We end refinements one mesh after no additional h -refinements are requested. As shown in Fig. 5, hp -adaptive refinements produce a series of meshes with a “sweeping” structure, i.e., they follow the direction of propagation of the beam.

In the case of uniform h -refinements, the observed linear increase in iterations with frequency is expected and is related to the inadequacy of coarse-space corrections when meshes are not sufficiently fine to resolve the wave. With hp -adaptive refinements, the behavior of the number of iterations until convergence with respect to frequency is less obvious since intermediate meshes are able to *partially* resolve the wave. Indeed, we initially believed the “sweeping” structure of meshes helped to reduce the frequency dependence of convergence. The convergence study in Fig. 6 seems to indicate this is not the case; the number of iterations show a clear linear increase with frequency. However, note the maximum number of iterations required for convergence was consistently higher than for uniform refinements; this is unexpected since the

⁴ <https://github.com/Oden-EAG/hp3d>.




Fig. 3. Convergence of the DPG-MG solver with *five smoothing steps* applied to uniformly h -refined meshes. Doing additional smoothing tends to restore h -robustness in (b) to some extent; in (a), h -robustness is still observed, however using five smoothing steps per level only reduces the number of iterations until convergence by a factor of approximately two when compared to one smoothing step. Iterations depend linearly on the frequency ω .




Fig. 4. Convergence of the DPG-MG solver applied to uniformly hp -refined meshes: grids are uniformly h -refined until two elements per wavelength, then uniformly p -refined. The solver convergence is p -robust and the iteration numbers are lower when using fine-grid restriction. The iterations until convergence depend linearly on the frequency ω .

adaptive case smoothes on each grid level, thus a much larger number of smoothing steps are performed overall.

An initial hypothesis on the cause of the deteriorated convergence of the DPG-MG solver on hp -adaptive meshes implicates a similar phenomenon underlying the deteriorated convergence when coarse-grid systems are stored from previous meshes. In that case, roughly speaking, different scales between fine and coarse systems, imparted by element-wise Gram matrices, were ameliorated by constructing coarse-grid systems as restrictions of fine-grid systems, so that all systems inherit the fine-grid scale. Under hp -adaptive refinements, multiple element sizes—with various scales imparted by element-wise Gram matrices—are simultaneously present. We are undertaking a more rigorous investigation; however, note that in Fig. 6a, the number of iterations for convergence on the final meshes decreases considerably. Returning to Fig. 5, we can see that these final meshes are characterized exclusively by p -refinements with a consistent, uniform element size; from which coarse-grid systems are restricted.

5. Application – seismic modeling

The observed linear increase of iterations with frequency ω implies a suboptimal $\mathcal{O}(N^{4/3})$ computational complexity of the DPG-MG solver in the preasymptotic regime; this is comparable to other methods in-

cluding shifted Laplacian [28,53], multilevel [56,9,34], and domain decomposition [7,44,55] methodologies, but is worse than the logarithmic increase observed for sweeping-type preconditioners [24,59] including source-transfer [46], and others. Despite this linear increase, the DPG-MG solver is competitive for solving large-scale high-frequency wave propagation problems. To illustrate the performance and flexibility of the DPG-MG solver, we consider the GO_3D_OBS benchmark [35] from seismic modeling. The benchmark problem is set in a hexahedral domain with high-contrast heterogeneous structures representing a subduction zone, inspired by the geology of the Nankai Trough. Following [58], we consider a $20 \times 102 \times 28.3$ km³ section of the model. The wavespeed in this section is illustrated in Fig. 7 and varies from 1500 m/s to 8500 m/s. Material data is specified on a uniform grid with spacing 100 m, downsampled from the original 25 m spacing of the GO_3D_OBS model (which is rather large, occupying 132 GB per parameter). The following computations were performed on 512 *Frontiera* Cascade Lake (CLX) compute nodes (28672 cores) at the Texas Advanced Computing Center.

The problem is driven by a point source with a frequency of 3.75 Hz located at (10.0, 12.5, 0.0), implemented as a tight Gaussian with standard deviation $\sigma = 50$ m. For simplicity, we use an initial mesh consisting of $8 \times 42 \times 12$ hexahedral quadratic ($p = 2$) elements with a total of 205757 degrees of freedom. General unstructured meshes, fitted to




Fig. 5. hp -adaptive propagation of a Gaussian beam in a cube domain; crinkle cut of hp -adaptive meshes (top) and surface cut of the real part of the acoustic pressure field (bottom). The hp -adaptive meshes have a “sweeping” structure as refinements first accumulate near the corner and then propagate into the domain; solutions on intermediate meshes are stable and partially resolve the wave.




Fig. 6. Convergence DPG-MG solver applied to hp -adaptive meshes. The number of iterations required for convergence again increases linearly with frequency (ω) and the maximum number of iterations is higher than for uniform refinements (compare Fig. 2a).

high-contrast interfaces or adapted to wavespeed, could be used in conjunction with the DPG-MG solver with great effect; however, we illustrate that adaptive refinements, starting from an arbitrarily coarse initial mesh, can resolve complex problems without the need for hand-tuned or time-intensive meshing.

The Dörfler marking strategy [22] is again used to mark elements for adaptive refinement. We perform seven initial h -adaptive refinements to resolve the region around the point source, followed by hp -adaptive refinements until the DPG residual is reduced by a factor of $3 \cdot 10^2$ (which coincided with the exhaustion of computer memory). The hp -adaptive strategy selects h -refinements when the maximum edge length of an element is less than one-half of the wavelength, except in regions of high contrast—which we define to be a greater than 10% change in wavespeed over the element—where one-quarter of the wavelength is used. Otherwise, p -refinements are selected until a maximum order of $p = 5$, after which h -refinements are again performed. Element

order was limited to $p = 5$ since higher-order elements can result in large smoothing patches that are expensive to store; we are pursuing a number of strategies to reduce patch storage, including a GPU implementation that recomputes smoothing patches during the solution. For applications that solve for many loads simultaneously, the cost of recomputing patches can be amortized over multiple loads, increasing the appeal of this approach.

In total, the mesh is refined 34 times, resulting in a final mesh (illustrated in Fig. 8) with over 6.3 million elements and 1.9 billion degrees of freedom. The upper right-hand corner of the mesh in Fig. 8 is not fully refined and a few more adaptive steps would be needed to further refine the mesh in that region; however, looking ahead to Fig. 9, it can be seen that the solution near this region has a relatively small amplitude. As alluded to previously, the refinement process was terminated due to a lack of memory; this is because the current implementation stores all 34 adaptive meshes and associated coarse-grid operators (which are re-




Fig. 7. Cutaway of acoustic wavespeed for a section of the GO_3D_OBS model, representing a subduction zone.




Fig. 8. Cross-sections of the hexahedral mesh (top); element order (middle); and wavespeed (bottom), at a cross-wise distance of 10 km. h -adaptivity is used to resolve high-contrast interfaces and the point source (enlarged to show detail) while p -adaptivity helps to mitigate the pollution error.




Fig. 9. Acoustic pressure (real) for 3.75 Hz frequency and a point source located at (10, 12.5, 0). The following cross-sections are shown: from the top, 0.5 km depth, 15 km depth, 25 km depth, and 10 km cross-wise (left); and 12.5 km dip-wise (right). The field is amplitude-compensated (scaled by the distance from the source) to enable visualization throughout the domain.

computed as restrictions of fine-grid operators after each refinement). Various amelioration strategies are possible, e.g. refactoring the refinement tree to group refinements of similar depth and reduce the total number of grids.

Cross-sections of the solution are illustrated in Fig. 9. These results are qualitatively similar to those depicted in [58], however we note the location of the point source between the two solutions differs slightly. A quantitative comparison of solution accuracy is deferred for a future work.

Convergence and timings for this example are shown in Fig. 10. The number of iterations required for convergence (again using a tolerance of 10^{-7}) on each grid is depicted in Fig. 10a; however, note that in contrast to the convergence studies in Section 4, here we use the solution on the previous meshes to initialize the solution on subsequent meshes. The effect of initializing with previous solutions becomes apparent in later iterations, when the solution is reasonably well resolved in much of the domain and further refinements result in fairly localized perturbations to the solution. Fig. 10b shows that the DPG residual decreases early in the adaptive process, when compared to the hp -adaptive Gaussian beam problem in Fig. 6b; this is likely because the wave decays relatively quickly away from the point source.

Timings for the assembly and solution phase on each mesh are shown in Fig. 10; the solution on the finest mesh was completed in 210 seconds whereas the total runtime for the job (including assembly and solution on 34 meshes) was 3029 seconds (51 minutes). Because the solution on different grids involve different numbers of iterations, the time *per iteration* is depicted in Fig. 10d to better gauge efficiency of the implementation. After an early preasymptotic regime, it can be seen that the time per iteration scales roughly linearly with respect

to DOFs; however, some super-linearity is observed in the largest instances. Super-linearity of multigrid solvers is expected when the number of DOFs grows sub-geometrically between grid levels (as is the case late in the refinement process). Refactoring the refinement tree to reduce the number of grid levels could help mitigate super-linear scaling and significantly reduce the cost of hp -adaptive solver iterations.

Finally, we note that the variations in time per iteration in Fig. 10d are related to the significant challenge of load balancing on hp -adaptive meshes. Indeed, hp -adaptive refinements produce elements and smoothing patches with *highly* disparate costs that must be accurately predicted then properly partitioned. We neglect definition of our load balancing strategy in this work, but we note that the multilevel approach employed here operates on a large number of relatively small elements and patches; this greatly simplifies the estimation of costs, can provide opportunities for more fine-grain parallelism, and is often more conducive to shared-memory parallelism than other methods including domain decomposition and sweeping-type preconditioners. Scalable adaptivity is a key differentiator of the DPG-MG solver which, to our knowledge, is novel among Helmholtz solvers.

Uniform refinements We conclude this section by considering uniform h - and p -refinements for the GO_3D_OBS benchmark. As remarked earlier, under the current implementation the DPG-MG solver stores and operates on all previous grids, which can become expensive when a small number of elements are refined between meshes. This added expense can be somewhat justified since adaptive meshes often attain a similar accuracy with a small fraction of the number of DOFs. Still, we are working toward reducing the additional expense of applying the solver to adaptively refined meshes. The following uniform refinement




Fig. 10. Convergence and timings for hp -adaptive solution of the GO_3D_OBS model with the DPG-MG solver (28672 cores). The number of iterations is better controlled when solutions on previous grids are used to initialize subsequent grids. For a fixed frequency, the DPG-MG solver scales nearly linearly with respect to degrees of freedom; super-linear scaling is due to sub-geometric growth in DOFs.

example is intended to provide a baseline for potential performance of adaptive refinements and to provide a more direct point of comparison to other Helmholtz solver implementations.

We start from the same $8 \times 42 \times 10$ initial mesh of uniform order $p = 2$ and perform four uniform h -refinements (corresponding to roughly two elements per wavelength in the water), followed by two uniform p -refinements. The final mesh has 16 million elements and 5.6 billion degrees of freedom. Convergence and timing information are provided in Fig. 11. Comparing Fig. 11b to Fig. 10b, it can be seen that the uniform refinement setting requires nearly three times as many DOFs and only reaches a DPG residual ten times larger than for hp -adaptive refinements, illustrating the optimality of meshes produced with hp -adaptive refinements using the DPG error indicator. Still, comparing Fig. 11d and Fig. 10d, it can also be seen that the time per iteration for uniformly refined meshes is three times smaller than for adaptive meshes, or nine times smaller when normalizing for the number of DOFs. A significant benefit may thus be attained by reducing the number of grid levels used during hp -adaptive refinements. Finally, we note that the number of iterations for uniform refinements (Fig. 11a) was higher than for the adaptive case (Fig. 10a); this is contrary to findings in Section 4 but seems to be the case here because adaptive refinements resolve the point source on early grids and thus typically begin iterations from a much lower residual than uniformly refined meshes.

6. Conclusion

The DPG-MG solver leverages the unique properties of the DPG methodology including mesh-independent stability, a built-in error indi-

cator, and Hermitian positive-definite discrete systems to enable robust, adaptable, and scalable solution of high-frequency wave propagation problems. When coarse-grid operators are constructed as restrictions of fine-grid operators, the DPG-MG solver demonstrates clear h - and p -robust convergence and a linear dependence with respect to wave frequency. A similar linear dependence on frequency was observed for hp -adaptive refinements. Despite the linear increase in number of iterations with respect to frequency, a scalable MPI/OpenMP implementation of the DPG-MG solver was demonstrated to be competitive for high-frequency wave propagation problems. In initial, moderate-scale tests, the DPG-MG solver was able to solve a challenging high-contrast seismic modeling benchmark (GO_3D_OBS) with 1.9 billion DOFs on hp -adaptive meshes and 5.6 billion DOFs on uniformly refined meshes; a larger scale than any work we are currently aware of for high-frequency wave propagation in heterogeneous media. A significantly smaller DPG residual was achieved when hp -adaptive meshes were used, however we defer quantitative comparisons of accuracy to a later publication.

Future directions Scalable implementation of the DPG-MG solver has motivated a number of promising research directions. First, as indicated in Section 4, we intend to further investigate the deteriorated convergence rate of the DPG-MG solver when coarse-grid operators are stored and in the case of hp -adaptive meshes. Second, we are working to integrate the solver with automatic, fully-anisotropic hp -adaptivity [13,20], where the DPG error indicator is first used to mark isotropic hp -refinements, an optimal set of anisotropic h - and p -refinements is then extracted from the isotropic refinements, and the remaining refinements are finally discarded. We anticipate the combination of the DPG-MG




Fig. 11. Convergence and timings for uniform h - and p -refinements of the GO_3D_OBS model with the DPG-MG solver (28672 cores). Nearly three times as many DOFs are required and the final DPG residual is ten times larger than the hp -adaptive case (compare Fig. 10b); however, the time per iteration is three times smaller (compare Fig. 10d).

solver with automatic hp -adaptivity will be competitive for complex boundary layer problems and other problems with highly anisotropic features. We are continuously working to improve scaling and performance of the DPG-MG solver; near-term improvements include integration of GPUs for a memory-efficient implementation, refactorization of the refinement tree to reduce the number of grids under adaptive refinements, and implementation of a fully distributed data structure in $hp3D$ (a fairly light-weight but replicated data structure currently limits scalability).

Finally, we are working to apply the scalable DPG-MG solver to challenging problems in science and engineering. For example, we have implemented an ultraweak time-harmonic Maxwell model of a tokamak device on an unstructured tetrahedral mesh with numerous reentrant corners and complex features; hp -adaptivity is expected to be particularly advantageous in this application. Additional applications of interest include extension of an optical fiber amplifier model [40,41] to bent and complex cross-section fibers, and implementation of ultraweak elastic Helmholtz for seismic modeling and, ultimately, seismic inversion.

Data availability

Data will be made available on request.

References

- [1] P. Amestoy, I. Duff, J. L'Excellent, J. Koster, A fully asynchronous multifrontal solver using distributed dynamic scheduling, *SIAM J. Matrix Anal. Appl.* 23 (1) (2001) 15–41.
- [2] D.N. Arnold, R.S. Falk, R. Winther, Multigrid in $H(\text{div})$ and $H(\text{curl})$, *Numer. Math.* 85 (2) (2000) 197–217.
- [3] I.M. Babuška, Error-bounds for finite element method, *Numer. Math.* 16 (4) (1971) 322–333.
- [4] J. Badger, S. Henneking, L. Demkowicz, Sum factorization for fast integration of DPG matrices on prismatic elements, *Finite Elem. Anal. Des.* 172 (2020) 103385.
- [5] S. Balay, et al., PETSc users manual, Tech. Rep. ANL-95/11 - Revision 3.13, Argonne National Laboratory, 2020.
- [6] M. Bonazzoli, et al., Domain decomposition preconditioning for the high-frequency time-harmonic Maxwell equations with absorption, preprint, arXiv:1711.03789, 2017.
- [7] X.-C. Cai, O.B. Widlund, Domain decomposition algorithms for indefinite elliptic problems, *SIAM J. Sci. Stat. Comput.* 13 (1) (1992) 243–258.
- [8] Z. Cai, T.A. Manteuffel, S.F. McCormick, First-order system least squares for the Stokes equations, with application to linear elasticity, *SIAM J. Numer. Anal.* 34 (5) (1997) 1727–1741.
- [9] H. Calandra, S. Gratton, X. Pinel, X. Vasseur, An improved two-grid preconditioner for the solution of three-dimensional Helmholtz problems in heterogeneous media, *Numer. Linear Algebra Appl.* 20 (4) (2013) 663–688.
- [10] C. Carstensen, L. Demkowicz, J. Gopalakrishnan, Breaking spaces and forms for the DPG method and applications including Maxwell equations, *Comput. Math. Appl.* 72 (3) (2016) 494–522.
- [11] Z. Chen, X. Xiang, A source transfer domain decomposition method for Helmholtz equations in unbounded domain, *SIAM J. Numer. Anal.* 51 (4) (2013) 2331–2356.
- [12] H. De Sterck, T.A. Manteuffel, S.F. McCormick, L. Olson, Least-squares finite element methods and algebraic multigrid solvers for linear hyperbolic PDEs, *SIAM J. Sci. Comput.* 26 (1) (2004) 31–54.
- [13] L. Demkowicz, Computing with hp Finite Elements. I. One and Two Dimensional Elliptic and Maxwell Problems, Chapman & Hall/CRC Press, Taylor and Francis, 2006.
- [14] L. Demkowicz, Energy Spaces, Lecture Notes, The University of Texas at Austin, 2018.
- [15] L. Demkowicz, Mathematical Theory of Finite Elements, Lecture Notes, The University of Texas at Austin, 2020.

- [16] L. Demkowicz, J. Gopalakrishnan, A class of discontinuous Petrov–Galerkin methods. II. Optimal test functions, *Numer. Methods Partial Differ. Equ.* 27 (1) (2011) 70–105.
- [17] L. Demkowicz, J. Gopalakrishnan, A class of discontinuous Petrov–Galerkin methods. Part I: the transport equation, *Comput. Methods Appl. Mech. Eng.* 199 (23–24) (2010) 1558–1572.
- [18] L. Demkowicz, J. Gopalakrishnan, Discontinuous Petrov–Galerkin (DPG) method, in: *Encyclopedia of Computational Mechanics*, second edition, 2017, pp. 1–15.
- [19] L. Demkowicz, J. Gopalakrishnan, A.H. Niemi, A class of discontinuous Petrov–Galerkin methods. Part III: adaptivity, *Appl. Numer. Math.* 62 (4) (2012) 396–427.
- [20] L. Demkowicz, et al., *Computing with hp Finite Elements. II. Frontiers: Three Dimensional Elliptic and Maxwell Problems with Applications*, Chapman & Hall/CRC, 2007.
- [21] K. Devine, et al., Zoltan data management services for parallel dynamic applications, *Comput. Sci. Eng.* 4 (2) (2002) 90–97.
- [22] W. Dörfler, A convergent adaptive algorithm for Poisson’s equation, *SIAM J. Numer. Anal.* 33 (3) (1996) 1106–1124.
- [23] B. Engquist, L. Ying, Sweeping preconditioner for the Helmholtz equation: hierarchical matrix representation, *Commun. Pure Appl. Math.* 64 (5) (2011) 697–735.
- [24] B. Engquist, L. Ying, Sweeping preconditioner for the Helmholtz equation: moving perfectly matched layers, *Multiscale Model. Simul.* 9 (2) (2011) 686–710.
- [25] Y.A. Erlangga, Advances in iterative methods and preconditioners for the Helmholtz equation, *Arch. Comput. Methods Eng.* 15 (1) (2008) 37–66.
- [26] O. Ernst, M.J. Gander, Why it is difficult to solve Helmholtz problems with classical iterative methods, in: *Numerical Analysis of Multiscale Problems*, in: *Lect. Notes Comput. Sci. Eng.*, vol. 83, Springer, 2012, pp. 325–363.
- [27] I.G. Graham, E.A. Spence, E. Vainikko, Domain decomposition preconditioning for high-frequency Helmholtz problems with absorption, *Math. Comput.* 86 (2015).
- [28] M.J. Gander, I. Graham, E. 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 (3) (2015) 567–614.
- [29] M.J. Gander, 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 Rev.* 61 (1) (2019) 3–76.
- [30] J. Gopalakrishnan, I. Muga, N. Olivares, Dispersive and dissipative errors in the DPG method with scaled norms for Helmholtz equation, *SIAM J. Sci. Comput.* 36 (1) (2014) A20–A39.
- [31] J. Gopalakrishnan, J. Pasciak, Overlapping Schwarz preconditioners for indefinite time harmonic Maxwell equations, *Math. Comput.* 72 (241) (2003) 1–15.
- [32] J. Gopalakrishnan, W. Qiu, An analysis of the practical DPG method, *Math. Comput.* 83 (286) (2014) 537–552.
- [33] J. Gopalakrishnan, 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.
- [34] J. Gopalakrishnan, J.E. Pasciak, L.F. Demkowicz, Analysis of a multigrid algorithm for time harmonic Maxwell equations, *SIAM J. Numer. Anal.* 42 (1) (2004) 90–108.
- [35] A. Górszczyk, S. Operto, GO_3D_OBS: the multi-parameter benchmark geoscientific model for seismic imaging method assessment and next-generation 3D survey design (version 1.0), *Geosci. Model Dev.* 14 (3) (2021) 1773–1799.
- [36] S. Henneking, A scalable hp -adaptive finite element software with applications in fiber optics, PhD thesis, The University of Texas at Austin, 2021.
- [37] S. Henneking, L. Demkowicz, A numerical study of the pollution error and DPG adaptivity for long waveguide simulations, *Comput. Math. Appl.* 95 (2021) 85–100.
- [38] S. Henneking, L. Demkowicz, *Computing with hp Finite Elements. III. Parallel hp 3D Code*, 2023, in preparation.
- [39] S. Henneking, L. Demkowicz, hp 3D user manual, preprint, arXiv:2207.12211, 2022.
- [40] S. Henneking, J. Grosek, L. Demkowicz, Model and computational advancements to full vectorial Maxwell model for studying fiber amplifiers, *Comput. Math. Appl.* 85 (2021) 30–41.
- [41] S. Henneking, J. Grosek, L. Demkowicz, Parallel simulations of high-power optical fiber amplifiers, in: *Spectral and High Order Methods for Partial Differential Equations ICOSAHOM 2020+1*, Springer, 2022, pp. 349–360.
- [42] R. Hiptmair, Multigrid method for $H(\text{div})$ in three dimensions, *Electron. Trans. Numer. Anal.* 6 (1) (1997) 133–152.
- [43] R. Hiptmair, Multigrid method for Maxwell’s equations, *SIAM J. Numer. Anal.* 36 (1) (1998) 204–225.
- [44] S. Kim, H. Zhang, Optimized Schwarz method with complete radiation transmission conditions for the Helmholtz equation in waveguides, *SIAM J. Numer. Anal.* 53 (3) (2015) 1537–1558.
- [45] B. Lee, T.A. Manteuffel, S.F. McCormick, J. Ruge, First-order system least-squares for the Helmholtz equation, *SIAM J. Sci. Comput.* 21 (5) (2000) 1927–1949.
- [46] W. Leng, L. Ju, A diagonal sweeping domain decomposition method with source transfer for the Helmholtz equation, preprint, arXiv:2002.05327, 2020.
- [47] J.D. Mora, L. Demkowicz, Fast integration of DPG matrices based on sum factorization for all the energy spaces, *Comput. Methods Appl. Math.* 19 (3) (2019) 523–555.
- [48] 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.
- [49] S. Petrides, L. Demkowicz, An adaptive DPG method for high frequency time-harmonic wave propagation problems, *Comput. Math. Appl.* 74 (8) (2017) 1999–2017.
- [50] S. Petrides, L. Demkowicz, An adaptive multigrid solver for DPG methods with applications in linear acoustics and electromagnetics, *Comput. Math. Appl.* 87 (2021) 12–26.
- [51] S. Rhebergen, G.N. Wells, Preconditioning for a pressure-robust HDG discretization of the Stokes equations, *SIAM J. Sci. Comput.* 44 (1) (2022) A583–A604.
- [52] S. Rhebergen, G.N. Wells, Preconditioning of a hybridized discontinuous Galerkin finite element method for the Stokes equations, *J. Sci. Comput.* 77 (3) (2018) 1936–1952.
- [53] A.H. Sheikh, D. Lahaye, C. Vuik, On the convergence of shifted Laplace preconditioner combined with multilevel deflation, *Numer. Linear Algebra Appl.* 20 (4) (2013) 645–662.
- [54] D. Stanzione, et al., Frontera: the evolution of leadership computing at the national science foundation, in: *Practice and Experience in Advanced Research Computing*, 2020, pp. 106–111.
- [55] C.C. Stolk, A rapidly converging domain decomposition method for the Helmholtz equation, *J. Comput. Phys.* 241 (2013) 240–252.
- [56] C.C. Stolk, M. Ahmed, S.K. Bhowmik, A multigrid method for the Helmholtz equation with optimized coarse grid corrections, *SIAM J. Sci. Comput.* 36 (6) (2014) A2819–A2841.
- [57] M. Taus, L. Zepeda-Núñez, R.J. Hewett, L. Demanet, L-Sweeps: a scalable, parallel preconditioner for the high-frequency Helmholtz equation, *J. Comput. Phys.* 420 (2020) 109706.
- [58] P.-H. Tournier, et al., 3D finite-difference and finite-element frequency-domain wave simulation with multilevel optimized additive Schwarz domain-decomposition preconditioner: a tool for full-waveform inversion of sparse node data sets, *Geophys.* 87 (5) (2022) T381–T402.
- [59] A. Vion, C. Geuzaine, Double sweep preconditioner for optimized Schwarz methods applied to the Helmholtz problem, *J. Comput. Phys.* 266 (2014) 171–190.