

A short note on a fast and high-order Hybridizable Discontinuous Galerkin solver for the 2D high-frequency Helmholtz equation

Matthias Taus^{1(*)}, Leonardo Zepeda-Núñez², and Laurent Demanet¹

¹Dept. of Mathematics and Earth Resources Lab, Massachusetts Institute of Technology;

²Dept. of Mathematics, University of California, Irvine.

SUMMARY

The method of polarized traces provides the first documented algorithm with truly scalable complexity for the high-frequency Helmholtz equation, i.e., with a runtime sublinear in the number of volume unknowns in a parallel environment. However, previous versions of this method were either restricted to a low order of accuracy, or suffered from computationally unfavorable boundary reduction to $\mathcal{O}(p)$ interfaces in the p -th order case. In this note we rectify this issue by proposing a high-order method of polarized traces with compact reduction to two, rather than $\mathcal{O}(p)$, interfaces. This method is based on a primal Hybridizable Discontinuous Galerkin (HDG) discretization in a domain decomposition setting. In addition, HDG is a welcome upgrade for the method of polarized traces, since it can be made to work with flexible meshes that align with discontinuous coefficients, and it allows for adaptive refinement in h and p . High order of accuracy is very important for attenuation of the pollution error, even in settings when the medium is not smooth. We provide some examples to corroborate the convergence and complexity claims.

INTRODUCTION

Solving the time-harmonic wave equation for heterogeneous wave-speeds in the high-frequency regime is a ubiquitous problem in geophysical exploration. The problem is still open in the context of numerical analysis, both from the points of view of optimal complexity and accuracy.

Recent progress has been made on two fronts: (i) new formulations that allow one to obtain more accurate approximations such as the ones presented in Hiptmair et al. (2015); Gittelsohn et al. (2009); Imbert-Gérard (2015), and (ii) new algorithms to solve the linear system in complexity linear in the number of volume unknowns, such as Engquist and Ying (2011a,b); Liu and Ying (2015a,b); de Hoop et al. (2011); Chen and Xiang (2013a,b); Leng (2015). It is difficult, however, to formulate a method that is both fast and of high order of accuracy. The direct methods typically used for solving the systems arising from highly accurate discretizations (cf. Davis (2004); Amestoy et al. (2001)) result in suboptimal complexities, and the state-of-the-art algorithms with optimal complexity tend to use low-order discretizations. There are some notable recent exceptions, such as the methods in Tsuji et al. (2014); Stolk (2015), which have both a high order of accuracy and optimal complexity; however, they use structured meshes and smooth wave-speeds, which reduce their applicability to problems in the context of geophysical exploration.

In this note we aim to bridge accuracy and efficiency. We present a versatile and fast algorithm to solve the constant density acoustic Helmholtz equation in the high-frequency regime using a high-order Hybridizable Discontinuous Galerkin (HDG) discretization (see Nguyen et al. (2011)). The novelty is twofold: first, the method of polarized traces introduced in Zepeda-Núñez and Demanet (2016) is extended to HDG discretizations of arbitrary order in a domain decomposition setting, obtaining a fast and yet high-order method; and second, using static condensation at the interfaces between subdomains, we attenuate the extra cost associated with the application of the method of polarized traces to high-order discretizations presented in Zepeda-Núñez and Demanet (2016).

A HDG discretization has several advantages: it is a high-order method, in which the number of globally coupled degrees of freedom are much smaller than for other Discontinuous or Continuous Galerkin formulations of the same order; local post-processing techniques can be performed in a parallel fashion to increase the accuracy of the solution; it accepts general triangular meshes, which can be used to appropriately align the mesh with possible discontinuities and topographic features; it accepts hp-adaptivity, which allows one to use different mesh- and polynomial-order-refinement, thus reducing the number of degrees of freedom to appropriately approximate the solution.

Within the HDG framework, the solution inside each triangle depends on the values of the solution on the triangle edges only. The degrees of freedom associated with the edges of the triangles are globally coupled, constituting the global system. Once a solution is computed on all edges of the mesh, the solution inside each triangle can be very efficiently computed in an embarrassingly parallel fashion, which can be seamlessly implemented using accelerators. In order to obtain a solution of the global system, we use the following approach: Within the domain decomposition framework, the global system is decomposed into layers following the inherent geometry. The interconnectivity of the global system between the layers is then reduced by applying a Schur complement. The resulting linear system can be re-written as a surface integral equation (SIE) posed on the boundaries of the layers, which is efficiently solved using the method of polarized traces.

The resulting algorithm has an asymptotic complexity of $\mathcal{O}(p^4N)$ for the setup, or off-line stage, and $\mathcal{O}(p^2N)$ for the solve, or on-line stage. (These complexity figures are linear rather than sublinear, because the fast algorithm component of the method of polarized traces is absent in this note – it would be a simple matter to restore it.) The empirically most expensive operation in the on-line stage, which is solving the global system, can be performed in $\mathcal{O}(N)$ complexity. The main advantage is that, given the amount of globally coupled degrees of freedom

A short note on a fast and high-order HDG solver for the 2D high-frequency Helmholtz equation

and the use of a Schur complement to decrease the connectivity between subdomains, the cost of solving the global linear system issued from a high-order HDG method is comparable to solving a linear system issued from a second order, finite difference discretization with the same number of degrees of freedom. Finally, we point out that the ideas showcased in this abstract can be easily extended to the 3D case.

METHOD

Model Problem and Absorbing Boundary Conditions

Let $\Omega = [a_1, b_1] \times [a_2, b_2] \subset \mathbb{R}^2$ be a rectangular domain of interest. We solve the constant density acoustic Helmholtz equation given by

$$-\Delta u(x) - \omega^2 m(x)u(x) = f(x) \quad \text{for } x \in \Omega,$$

plus absorbing boundary conditions (ABCs) on $\partial\Omega$; here ω is the frequency, $m(x)$ is the squared slowness, and $f(x)$ is the source term.

The ABCs are realized via perfectly matched layers (PMLs) (cf., Bérenger (1994)) which allow one to rewrite the problem as a boundary value problem defined on an extension of Ω :

$$-\operatorname{div} \left[\tilde{\Lambda}(x) \nabla u(x) \right] - \omega \tilde{m}(x) u(x) = \tilde{f}(x) \quad \text{for } x \in \tilde{\Omega}, \quad (1)$$

$$u(x) = 0 \quad \text{for } x \in \partial\tilde{\Omega}, \quad (2)$$

where $\tilde{\Omega} = [a_1 - \delta_1, b_1 + \delta_1] \times [a_2 - \delta_2, b_2 + \delta_2]$ for $\delta_1, \delta_2 > 0$; the coefficients are defined as

$$\tilde{\Lambda}(x) = \begin{pmatrix} \frac{\alpha_1(x)}{\alpha_2(x)} & 0 \\ 0 & \frac{\alpha_2(x)}{\alpha_1(x)} \end{pmatrix}, \quad \tilde{m}(x) = \frac{m(x)}{\alpha_1(x)\alpha_2(x)};$$

and the right-hand side is defined as

$$\tilde{f}(x) = \frac{f(x)}{\alpha_1(x)\alpha_2(x)}, \quad \text{with} \quad \alpha_i(x) = \frac{1}{1 + i \frac{\sigma_i(x)}{\omega}},$$

where

$$\sigma_i(x) = \begin{cases} \frac{C}{\delta_i} \left(\frac{a_i - x_i}{\delta_i} \right)^2, & \text{for } x_i \in (a_i - \delta_i, a_i) \\ 0, & \text{for } x_i \in (a_i, b_i) \\ \frac{C}{\delta_i} \left(\frac{x_i - b_i}{\delta_i} \right)^2, & \text{for } x_i \in (b_i, b_i + \delta_i) \end{cases}$$

with an appropriately chosen absorption constant $C > 0$.

Hybridizable Discontinuous Galerkin Methods

We discretize the boundary-value problem in Eqs. 1 and 2 using a (HDG) method based on the primal weak formulation. A similar HDG method has been previously considered for Poisson's equation (Waluga and Egger (2012)). Following this approach, we introduce a triangulation \mathcal{T}_h of $\tilde{\Omega}$ with a mesh-size h , the space V_h^p of piecewise discontinuous polynomials of degree p defined on each triangle $T \in \mathcal{T}_h$, and the space W_h^p of piecewise discontinuous polynomials of degree p defined on each edge $E \in \mathcal{T}_h$. The mesh-size is defined as $h = \sqrt{|T_{\max}|/|\tilde{\Omega}|}$ where $|T_{\max}|$ is the area of the largest triangle in \mathcal{T}_h and $|\tilde{\Omega}|$ is the area of $\tilde{\Omega}$.

An approximate solution u_h of u is then found such that

$$\sum_{T \in \mathcal{T}_h} a_T(u_h, \lambda_h; v_h, \mu_h) = \sum_{T \in \mathcal{T}_h} F_T(v_h, \mu_h)$$

for all $(v_h, \mu_h) \in V_h^p \times W_h^p$ where

$$\begin{aligned} a_T(u_h, \lambda_h; v_h, \mu_h) &:= \int_T \left[\tilde{\Lambda} \nabla u_h \cdot \nabla v_h - \omega \tilde{m} u_h v_h \right] dx \\ &\quad - \int_{\partial T} [n \cdot \nabla u_h (v_h - \mu_h) ds_x + (u_h - \lambda_h) n \cdot \nabla v_h] ds_x \\ &\quad + \tau \frac{p^2}{h} \int_{\partial T} (u_h - \lambda_h) (v_h - \mu_h) ds_x \end{aligned}$$

and

$$F_T(v_h, \mu_h) := \int_T \tilde{f} v_h dx.$$

Here, $\tau > 0$ is a stabilization constant that has to be chosen large enough to ensure the invertibility of the resulting linear system. Note that since the support of every basis function in V_h^p is restricted to one single element in \mathcal{T}_h , all degrees of freedom corresponding to u_h can be locally eliminated by static condensation on the edges of the element, resulting in a global system constituted by the degrees of freedom corresponding to λ_h only. To this end, in what follows, we focus on the efficient solution of this reduced global system. Once λ_h is obtained, u_h can be reconstructed locally on each element.

The Method of Polarized Traces

The method of polarized traces has been established as a method to efficiently invert system matrices originating from Finite Difference and low-order Finite Element Methods. In this work, we extend these results to HDG Methods of arbitrary order. The method of polarized traces consists of four main steps:

1. Divide the computational domain into layers.
2. Reduce the problem to the degrees of freedom $\bar{\lambda}_h$ corresponding to the boundaries of the layers, resulting in a surface integral equation (SIE) with sparsity pattern depicted in Fig. 1 (left).
3. Double the degrees of freedom on the boundaries of the layers by polarizing $\bar{\lambda}_h = \bar{\lambda}_h^\uparrow + \bar{\lambda}_h^\downarrow$, and reformulate an equivalent polarized SIE, such that the two diagonal blocks of the resulting system for $\bar{\lambda}_h^\uparrow$ and $\bar{\lambda}_h^\downarrow$ are an upper- and a lower-triangular matrix as depicted in Fig. 1 (right).
4. Efficiently solve the resulting system of linear algebraic equations using the upper- and lower-triangular matrix as a block-diagonal preconditioner.

This allows one to find $\bar{\lambda}_h$ efficiently from which λ_h can be reconstructed on each layer in an embarrassingly parallel way.

In our approach, we assume that the layers are chosen so that their boundaries cut through edges of the triangulation, see Figure 2, and eliminate all degrees of freedom corresponding to cut faces. In contrast to existing approaches using the method of polarized traces, this elimination allows one to reduce the problem to the degrees of freedom corresponding to

A short note on a fast and high-order HDG solver for the 2D high-frequency Helmholtz equation




Figure 1: Sparsity pattern of the SIE matrix (left), and the polarized SIE matrix (right).

the boundaries of the layers only, independently of the polynomial degree, see Fig. 2. This is a major advantage of our approach, as it allows for a significantly smaller linear system to obtain $\bar{\lambda}_h^\uparrow$ and $\bar{\lambda}_h^\downarrow$, especially for high polynomial degrees p .




Figure 2: The layers of the computational domain. The degrees of freedom corresponding to faces shown in gray are eliminated. In step 2, the system is reduced to the degrees of freedom corresponding to faces shown in red.

Once the degrees of freedom between any two layers are eliminated, steps 2-4 of the method of polarized traces can be applied on a purely algebraic level in exactly the same way as described in Zepeda-Núñez and Demanet (2016).

COMPLEXITY

We suppose that $N = \mathcal{O}(p^2 h^{-2})$, and that the number of subdomains, L , scales linearly with h^{-1} ; moreover, we suppose that each layer is $q = h^{-1}/L = \mathcal{O}(1)$ triangles wide. The method of polarized traces is decomposed in two stages: (i) the off-line stage performed once per linear system; and (ii) the on-line stage performed for each right-hand side.

In the off-line stage the domain is decomposed in L subdomains, and the matrices are assembled and factorized. For each subdomain, the assembly of the local matrices and the static condensation of the degrees of freedom at the interior of the triangles can be performed in $\mathcal{O}(p^6 h^{-1} q)$ complexity; this operation is embarrassingly parallel at the level of each triangle and can be performed using accelerators. Given that each layer is q triangles wide, the problem inside each layer is a quasi 1D

problem (cf. Engquist and Ying (2011b)), and consequently the LU factorization of the linear system can be performed in $\mathcal{O}(p^3 h^{-1} q^3)$ complexity using multi-frontal methods (see George (1973); Duff and Reid (1983)). This needs to be performed for L different layers resulting in an overall $\mathcal{O}(p^4 N)$ complexity for the off-line stage.

For the on-line stage, the right-hand side is generated via static condensation, which can be done in $\mathcal{O}(p^4 h^{-2})$ complexity. This procedure is embarrassingly parallel at the level of each triangle. Then the elimination of degrees of freedom at the interfaces between subdomains is efficiently performed in $\mathcal{O}(p^2 h^{-1})$ complexity via multi-frontal methods. Finally, the global system can be solved using the method of polarized traces in $\mathcal{O}(p^2 h^{-2}) = \mathcal{O}(N)$ complexity. The overall complexity of the on-line stage is $\mathcal{O}(p^2 N)$.

We point out that the operations with the highest complexities, which are the inversion and the solve of all the local systems within the triangles, can be performed in an embarrassingly parallel fashion. Further, given that the local systems fit in cache, the runtime of inverting each local system is dominated by the memory latency and bandwidth. Thus, the on-line runtime of the algorithm is $\mathcal{O}(N)$ for the range of examples considered in this abstract. Moreover, we point out that if each triangle contains $\mathcal{O}(1)$ wavelengths, then each local problem is essentially elliptic. In such cases, it is possible to further accelerate the algorithm using \mathcal{H} -matrices (cf., Bebendorf (2008)).

NUMERICAL EXPERIMENTS

The method presented in this abstract was implemented in Julia (Bezanson et al. (2012)), and distMesh (Persson and Strang (2004)) was used to generate the meshes. The numerical experiments were performed in a dual socket server with two Xeon E5-2780 CPU and 386 Gigs of RAM.

We performed two sets of numerical experiments: one to demonstrate the accuracy of the method, and one to corroborate the efficiency of the solver. To demonstrate the accuracy of the solver, we fix the frequency, and for different polynomial degrees p we test the convergence rate of the solver as the mesh is uniformly refined. To corroborate the efficiency, we time the execution time of the solver for h -refinements in the high-frequency regime, i.e., increasing the frequency as $\omega \sim N^{1/2}$.

To test the accuracy, we used a fault model shown in Fig. 3 (left) with a mesh that is aligned with the discontinuities. Given the discontinuities in the model, it is known that high-order finite differences will converge sub-optimally (see Larson (1999) and references therein). To test the convergence of our HDG method with respect to h , we fixed the frequency ω at 10 wave lengths within the side length of the square Ω and estimated the order of convergence (*roc*) using the L^2 -error corresponding to two consecutive refinement levels. A reference solution of this problem (Fig. 3) is computed on the finest mesh using $p = 4$. Table 1 shows the convergence of the HDG method for uniformly refined meshes. It can be seen that, in contrast to Finite Difference methods, the HDG method converges at the expected rates for arbitrary polynomial degrees.

A short note on a fast and high-order HDG solver for the 2D high-frequency Helmholtz equation




Figure 3: The wave speed for the fault model (*left*) and the corresponding reference solution (*right*).

h	$p = 1$		$p = 2$		$p = 3$	
	Error	eoc	Error	eoc	Error	eoc
7.1E-2	7.5E-1		1.4E-2		1.0E-3	
3.5E-2	2.2E-1	1.7	1.1E-3	3.7	7.0E-5	3.9
1.8E-2	5.8E-2	1.9	1.0E-5	3.4	4.7E-6	3.9

Table 1: The convergence of the HDG method in the normalized L^2 -error with respect to uniform mesh-refinement.

To test the efficiency, we used the fault and the BP 2004 model from Billette and Brandsberg-Dahl (2005) (shown in Fig. 4). Given the complex structure of the salt body in the BP 2004 model, the mesh is not aligned with the discontinuities. However, an adaptive integration strategy similar to Zepeda-Núñez and Demanet (2015) was employed in order to achieve accurate solutions in this case. For each model, we generated 10 right-hand sides by randomly putting a point source inside the domain Ω , and we timed the runtime of the algorithm.

Table 2 shows the timings of the off-line and the on-line stage for the fault model. We can observe that the total assembly time scales like $\mathcal{O}(N)$. For the on-line stage, we can observe that the number of iterations is weakly dependent on h and ω , resulting in an overall $\mathcal{O}(N)$ scaling. Table 3 shows the timings of the on-line stage for the BP 2004 model. We do not show timings for the off-line-stage of the BP 2004 model because the adaptive integration technique affects the assembly times so that the times of the off-line stage are not representative. We can observe that for the on-line stage, the time to apply the SIE and its preconditioner scales as $\mathcal{O}(N)$. The number of iterations needed to solve the resulting linear system depends weakly on h and ω , resulting in the advertised $\mathcal{O}(N)$ complexity.

DISCUSSION

We presented an extension of the method of polarized traces (Zepeda-Núñez and Demanet (2016)) to HDG discretizations of arbitrary order. The resulting method reduces the global system to a SIE defined only on the interfaces between subdomains, which can be efficiently solved using the method of polarized traces.




Figure 4: BP 2004 model (*left*) and wavefield generated by a point source (*right*).

h	$\frac{\omega}{2\pi}$	L	off-line stage	on-line stage
2.4E-2	5.21	3	41.19	1.43 (3.5)
1.2E-2	10.42	7	160.08	10.16 (2.5)
6.0E-3	20.84	15	701.85	47.73 (2.4)

Table 2: Time (in seconds) to perform the off-line stage for the fault model, and the average number of GMRES iterations (*bold*) required to reduce the relative residual to 10^{-5} along with the average execution time (in seconds) of one GMRES iteration for $p = 2$.

h	$\frac{\omega}{2\pi}$	L	on-line stage
1.2E-2	7.38	3	5.73 (3)
6.0E-3	14.75	7	18.76 (3.8)
3.0E-3	29.50	15	52.30 (4.9)

Table 3: Average number of GMRES iterations (*bold*) required to reduce the relative residual to 10^{-5} along with the average execution time (in seconds) of one GMRES iteration for the BP 2004 model for $p = 2$.

The presented method is particularly interesting for problems involving discontinuous coefficients. We have shown that if the discontinuities are resolved in the mesh, HDG methods converge optimally. If the discontinuities in the coefficients are not resolved in the mesh, the method still can be applied efficiently, as long as an appropriate integration is performed.

We point out that this approach can be further optimized by a parallel implementation taking advantage of all embarrassingly parallelizable stages of the algorithm. Moreover, this approach can be easily extended to 3D problems.

ACKNOWLEDGMENTS

The authors thank Total SA for support. LD is also supported by AFOSR, ONR, and NSF.