

Pollution-free and fast hybridizable discontinuous Galerkin solvers for the high-frequency Helmholtz equation

Matthias Taus^{1(*)}, Leonardo Zepeda-Núñez², Russell J. Hewett³, and Laurent Demanet¹

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

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

³Total E&P Research and Technology USA.

SUMMARY

In this work we propose a hybridizable discontinuous Galerkin (hdG) discretization of the high-frequency Helmholtz equation in the presence of point sources and highly heterogeneous and discontinuous wave speed models. We show that it delivers solutions that are provably second-order accurate and do not suffer from the pollution error, as long as a slightly higher order hdG method is used where the polynomial degree is chosen such that $p = O(\log \omega)$. These results hold even if the discontinuities in the wave speed are not resolved by the hdG mesh, as long as the integration procedure used in the assembly of the stiffness matrix respects the discontinuities. Further, we show that the associated linear systems can be solved using a modification of the method of polarized traces resulting in a method with linear complexity up to a poly-logarithmic factor, or sub-linear complexity in a parallel environment.

To our knowledge and surprise, this note contains the first instance of a numerical method that is at the same time fast ($O(N)$ runtime) and accurate (second-order, pollution-free) in the context of models of geophysical interest.

INTRODUCTION

The high-frequency Helmholtz equation defined over heterogeneous wave speeds is a ubiquitous problem in science and engineering. In particular, in geophysics, it often serves as the forward model for inverse problems, for example in oil exploration (see Chen (1997), Pratt (1999), Virieux and Operto (2009)). Within this context, it is therefore crucial to solve the Helmholtz equation in the high-frequency regime not only accurately but also efficiently.

In view of accuracy, the main difficulty of the problem is rooted in the pollution effect, whose primary consequence is a large shift in the phase of propagating waves, even if the Shannon-Nyquist sampling rate (see Shannon (1998)) is observed. In the context of inverse problems, these phase errors dramatically decrease the quality of the inversion. Indeed, when computing the imaging condition, these wavefields with large phase error may correlate in a different position, or not at all. It has been shown that the pollution error cannot be avoided for low-order discretizations (see Babuska and Sauter (1997)) unless the discretization is refined. This refinement either involves oversampling of the wavefield, resulting in a suboptimal number of degrees of freedom, or an increase of the order of the polynomial approximation.

A first class of methods proposed to deal with pollution errors

involves the choice of non-polynomial basis functions (see Gittelsohn et al. (2009), Moiola et al. (2011), Turkel et al. (2013), Hiptmair et al. (2015), Imbert-Gérard (2015), Nguyen et al. (2015), Stolk (2015), Peterseim (2016)). However, point sources may be problematic, these methods usually require piecewise constant wave speed, and their associated linear systems cannot be solved fast in general. A second class of methods involves polynomial basis functions and deals with the pollution error by increasing the polynomial degree. However, pollution-free convergence of these polynomial methods has only been established for constant wave speeds (Melenk and Sauter (2011)). There seems to be no published result for problems involving discontinuous wave speeds or point sources, which can induce extra errors polluting the solution in the full domain. These errors can be effectively handled using non-uniform meshes such as meshes aligned with discontinuities in the wave speed or adaptive meshes in the vicinity of point sources. However, meshes that are well-suited for numerical methods and aligned with discontinuities in the wave speed can be very hard to obtain, and adaptive meshes depend on the location of the point source. Consequently, for any new distribution of point sources a new mesh has to be generated.

In this note we aim to solve the deficiencies in accuracy described above. We introduce an 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 method is based on polynomial basis functions and uniform meshes, independently of the wave speed and point source distribution. Even though the mesh is not adapted to the discontinuities in the wave speed or the point sources, the proposed method is rigorously shown to be second-order accurate and to not be subject to the pollution error. This is achieved by utilizing a correct integration rule to handle discontinuities in the wave speed, a singularity removal technique leveraging the Green's function of Laplace's equation to handle point sources, and by increasing the degree of the polynomial degree as $p = O(\log \omega)$ to eliminate pollution error.

In addition, the proposed method can be solved in quasi-linear runtime, or sub-linear runtime in a parallel environment. This is because it can be placed in the framework of the method of polarized traces (see Zepeda-Núñez and Demanet (2016), Taus et al. (2016)). This is an exception to other published fast iterative methods (see Engquist and Ying (2011), Calandra et al. (2013), Chen and Xiang (2013), Stolk (2013), Vion and Geuzaine (2014), Zepeda-Núñez and Demanet (2016)) or direct methods (see Wang et al. (2013), Gillman et al. (2014)), which either rely on low-order discretizations or do not scale optimally in the high-frequency regime. Finally, fast solvers

Pollution-free and fast hdG solvers for the high-frequency Helmholtz equation

were introduced in (Zepeda-Núñez and Zhao (2016)) achieving optimal complexity with respect to the frequency but under much stricter assumptions on the wave speed and source term.

The resulting method is second-order accurate, pollution-free and can be solved with the method of polarized traces in an asymptotic complexity of $\mathcal{O}(\omega^2 \log^6 \omega)$ for the setup, or off-line stage, and $\mathcal{O}(\omega^2 \log^4 \omega)$ for the solve, or on-line stage (These complexity figures are linear rather than sub-linear, 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.) Further, the empirically most expensive operation in the on-line stage, which is solving the global system, can be performed in $\mathcal{O}(\omega^2 \log^2 \omega)$ complexity.

The main advantage of this method is that given the amount of globally coupled degrees of freedom 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, but with without pollution effect. 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

$$-\Delta u(\mathbf{x}) - \omega^2 m(\mathbf{x})u(\mathbf{x}) = f(\mathbf{x}) \quad \text{for } \mathbf{x} \in \Omega, \quad (1)$$

with absorbing boundary conditions (ABCs) on $\partial\Omega$, where ω is the frequency, $m(\mathbf{x})$ is the squared slowness, and $f(\mathbf{x})$ is the source term.

The ABCs are realized via perfectly matched layers (PMLs) (see Bérenger (1994), Bermúdez et al. (2007), Johnson (2010)) which allow one to rewrite the problem as a boundary value problem defined on an extension of Ω :

$$-\operatorname{div} \left[\tilde{\Lambda}(\mathbf{x}) \nabla u(\mathbf{x}) \right] - \omega^2 \tilde{m}(\mathbf{x})u(\mathbf{x}) = \tilde{f}(\mathbf{x}) \quad \text{for } \mathbf{x} \in \tilde{\Omega}, \quad (2)$$

$$u(\mathbf{x}) = 0 \quad \text{for } \mathbf{x} \in \partial\tilde{\Omega}, \quad (3)$$

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

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

and the right-hand side is defined as

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

where

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

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

Singularity Removal

We focus our attention to the case when f is modeled as a point source located at \mathbf{x}_0 , i.e., $f(\mathbf{x}) = \delta(\mathbf{x} - \mathbf{x}_0)$. In such cases, the right-hand side of Eq. 1 is a distribution, and there is little theory available to assess the accuracy of the solution. One alternative to dealing with a point source is to refine the mesh and to mollify the source; however, this approach requires to use a discretization adapted to the source, making the linear system source dependent. We use a different strategy that allows us to recover second-order accuracy and still retain a matrix that is independent of the source distribution.

We suppose that the solution, u , to the Helmholtz equation can be written as

$$u = u_s + v \quad (4)$$

where v captures the singular part of the solution.

Substituting Eq. 4 in Eq. 1 we obtain

$$-(\Delta + m\omega^2)u_s = \delta(\mathbf{x} - \mathbf{x}_0) + (\Delta + m\omega^2)v. \quad (5)$$

We suppose that v has the form

$$v(\mathbf{x}) = G_0(\mathbf{x}, \mathbf{x}_0)\chi_\varepsilon(\mathbf{x}, \mathbf{x}_0), \quad (6)$$

where $G_0(\mathbf{x}, \mathbf{x}_0)$ is the Green's function of Laplace's equation and χ_ε is a smooth cut-off function such that

$$\chi_\varepsilon(\mathbf{x}, \mathbf{x}_0) = \begin{cases} 1 & \text{if } x \in [a_1 + \varepsilon, b_1 - \varepsilon] \times [a_2 + \varepsilon, b_2 - \varepsilon] \\ 0 & \text{if } x \in \tilde{\Omega} \setminus \Omega \end{cases}, \quad (7)$$

where $\varepsilon > 0$ is frequency independent. Note that the support of the cut-off function does not overlap with the PML, and therefore the boundary conditions of Eq. 1 are preserved.

Following a standard computation and using the definition of $G_0(\mathbf{x}, \mathbf{x}_0)$, it can be shown that the Dirac delta on the right-hand side of Eq. 5 cancels and we get an equation for u_s ,

$$-(\Delta + m\omega^2)u_s = m\omega^2 G_0(\mathbf{x}, \mathbf{x}_0)\chi_\varepsilon + 2\nabla G_0(\mathbf{x}, \mathbf{x}_0) \cdot \nabla \chi_\varepsilon + G_0(\mathbf{x}, \mathbf{x}_0)\Delta \chi_\varepsilon, \quad (8)$$

whose right-hand side is in $L^2(\Omega)$. Using standard regularity theory, we will see that this is sufficient to show that u_s can be approximated with second-order accuracy.

Hybridizable discontinuous Galerkin Methods

After the removal of the singularity, we discretize the boundary-value problem in Eqs. 2 and 3 using a hdG method introduced in (Cockburn et al. (2009)) for elliptic problems and extended to the Helmholtz equation in (Griesmaier and Monk (2011)). To this end, we introduce a partitioning \mathcal{T}_h with mesh-size h of $\tilde{\Omega}$ into squares and its corresponding set of faces \mathcal{E}_h . Furthermore, we introduce spaces of polynomials V_h , W_h and M_h , where V_h and W_h are defined over \mathcal{T}_h and M_h over \mathcal{E}_h . These spaces are used to approximate $q := \tilde{\Lambda} \nabla u$, u , and $\hat{u} := u|_{\mathcal{E}_h}$, respectively. The mesh-size is defined as $h = \sqrt{|T_{\max}|/|\tilde{\Omega}|}$ where $|T_{\max}|$ is the area of the largest square in \mathcal{T}_h .

Pollution-free and fast hdG solvers for the high-frequency Helmholtz equation

An approximate solution (q_h, u_h, \hat{u}_h) of (q, u, \hat{u}) is then found such that

$$\begin{aligned} (i\omega\tilde{\Lambda}^{-1}q_h, \overline{r_h})_{\mathcal{T}_h} - (u_h, \overline{\text{div } r_h})_{\mathcal{T}_h} + \langle \hat{u}_h, \overline{r_h \cdot n} \rangle_{\partial\mathcal{T}_h} &= 0, \\ (i\omega u_h, \overline{w_h})_{\mathcal{T}_h} - (q_h, \overline{\nabla w_h})_{\mathcal{T}_h} + \langle \hat{q}_h \cdot n, \overline{w_h} \rangle_{\partial\mathcal{T}_h} &= (\tilde{f}, \overline{w})_{\mathcal{T}_h}, \\ \langle \hat{u}_h, \overline{\mu_h} \rangle_{\partial\tilde{\Omega}} &= 0, \\ \langle \hat{q}_h \cdot n, \overline{\mu_h} \rangle_{\partial\mathcal{T}_h \setminus \partial\tilde{\Omega}} &= 0, \end{aligned}$$

for all (r_h, w_h, μ_h) in $V_h \times W_h \times M_h$. Here, n is the unit vector and $\hat{q}_h := q_h + \tau(u_h - \hat{u}_h)$, where $\tau > 0$ is a stabilization parameter. Further, we used the common notation

$$(f, g)_{\mathcal{T}_h} = \sum_{T \in \mathcal{T}_h} \int_T f g dx, \quad \langle \hat{f}, \hat{g} \rangle_{\partial\mathcal{T}_h} := \sum_{T \in \mathcal{T}_h} \int_{\partial T} \hat{f} \hat{g} ds$$

for functions f, g, \hat{f} and \hat{g} .

Here, the spaces $V_h, W_h,$ and M_h are spaces of (discontinuous) piecewise polynomials of degree p . However, in order to ensure stability on square elements, V_h has to be enriched by three polynomials of higher degree (Cockburn et al. (2012)). Independently of the order p , this enrichment only involves three basis functions and therefore does not affect the efficiency of the method.

Error Analysis

The hdG method considered in this abstract has been analyzed in (Griesmaier and Monk (2011)) for triangular elements and constant functions m . Using results from (Cockburn et al. (2012)), the techniques developed in (Griesmaier and Monk (2011)) can be easily extended to basis functions defined on squares and to bounded functions m . These error estimates depend on the smoothness of the solution in each element. To this end, we categorize the elements $T \in \mathcal{T}$ in two groups: the set \mathcal{T}_h^s containing the elements of \mathcal{T}_h on which u is smooth, and the set $\mathcal{T}_h^{\text{ns}}$ containing all other elements. It can be easily argued that the sets \mathcal{T}_h^s and $\mathcal{T}_h^{\text{ns}}$ can be determined a-priori as the sets where m is smooth and non-smooth respectively. Finally, assuming that

$$\left\| \frac{\partial u}{\partial x_i} \right\|_{L^2(T)} \leq C\omega \|u\|_{L^2(T)}$$

for all $T \in \mathcal{T}_h$, the error estimate is

$$\|u - u_h\|_{L^2(\tilde{\Omega})} \leq C(\omega)h^2 \left[\sum_{T \in \mathcal{T}_h^s} \left(\frac{\omega h}{p} \right)^{p+1} + \sum_{T \in \mathcal{T}_h^{\text{ns}}} \left(\frac{\omega h}{p} \right)^2 \right] \quad (9)$$

for a constant $C(\omega) > 0$ independent of p and h . Note that this can be bounded by

$$\|u - u_h\|_{L^2(\tilde{\Omega})} \leq C(\omega) \left(\frac{\omega h}{p} \right)^2$$

and therefore we conclude that our hdG method is second-order accurate.

Pollution Error

It has already been argued in (Griesmaier and Monk (2011))

that keeping the polynomial degree p and the number elements within each wavelength fixed causes the considered hdG method to suffer from pollution effects. This can be easily seen from the error estimate in Eq. 9. Indeed, keeping the number of elements per wavelength fixed means that ωh can be bounded by a constant and therefore the error is bounded by $C(\omega)$. Furthermore, we can show that in our context, $C(\omega)$ depends linearly on ω , i.e. $C(\omega) = C\omega$ for a constant $C > 0$ independent of ω . Therefore, if we keep the polynomial degree fixed at $p = 1$ and the numbers of degrees of freedom in each wavelength constant (ωh is constant), with increasing frequency we get that

$$\|u - u_h\|_{L^2(\tilde{\Omega})} \leq C\omega. \quad (10)$$

Numerical examples, presented in the sequel, show that this linear dependency on ω is a sharp estimate and therefore characterizes the pollution error of the proposed hdG method.

If we assume that m is piecewise smooth, there are $O(1/h)$ elements in $\mathcal{T}_h^{\text{ns}}$ since the discontinuities are all along lines. Consequently, the error can be written as

$$\|u - u_h\|_{L^2(\tilde{\Omega})} \leq C \left[\omega h \left(\frac{\omega h}{p} \right)^2 + \left(\omega^{\frac{1}{p+1}} \frac{\omega h}{p} \right)^{p+1} \right].$$

Therefore, if we choose h small enough such that $\omega h < 1$ and p large enough such that $\omega^{\frac{1}{p+1}}/p \leq 1$, the error is bounded and therefore the pollution error is eliminated. Note that this restriction on p is equivalent to $p = O(\log \omega)$.

Solving the Linear System and Complexity

Since the support of every basis function in V_h and W_h is restricted to one single element in \mathcal{T}_h , all degrees of freedom corresponding to u_h and q_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 \hat{u}_h only. The global linear system is then solved using a matrix-free version of the method of polarized traces (see Zepeda-Núñez and Demanet (2016)) specially adapted for this discretization in (Taus et al. (2016)).

The method of polarized traces consists of two stages: the off-line and the on-line stage. The off-line stage has to be computed only once for each wave speed distribution and mesh. Then, in the on-line stage, the system can be solved for any right-hand side. It was shown in (Taus et al. (2016)) that the off-line and on-line stage can be computed in $O(p^4N)$ and $O(p^2N)$ time respectively. It has also been shown that the empirically most expensive part of the on-line stage, which is solving the global system, can be performed in $\mathcal{O}(N)$ complexity. Here, the degrees of freedom can be computed as $N = p^2h^{-2}$. Therefore, for a given frequency ω , the off-line and on-line stages can be performed in $O(\omega^2 \log^6 \omega)$ and $O(\omega^2 \log^4 \omega)$ time, and the global system can be solved in $O(\omega^2 \log^2 \omega)$ complexity.

NUMERICAL EXPERIMENTS

The method presented in this abstract was implemented in Julia v0.5 (see Bezanson et al. (2012)). The numerical experi-

Pollution-free and fast hdG solvers for the high-frequency Helmholtz equation

ments were performed in a dual socket server with two Xeon E5-2780 CPU and 384 GB of RAM.

We performed two sets of numerical experiments: one to demonstrate the second-order accuracy of the method, and the other to verify the elimination of the pollution error. In both problems, the wave speed model is part of the 2004 BP model (see Billette and Brandsberg-Dahl (2005)) and the mesh does not align with the discontinuities of the medium (Figure 1).




Figure 1: 2004 BP model (left) and wavefield generated by a point source (right).

To demonstrate the accuracy of the solver, we fixed the frequency at $\omega = 12\pi$, placed a point source at $(0.8, 0.8)$, and computed a reference solution u_h^{ref} on an overly refined uniform mesh. We then fixed the polynomial degree at $p = 1$ and computed solutions u_h on coarser uniform meshes. To assess the convergence of the method we computed the L^2 -error

$$E(h) := \sqrt{\int_{\Omega} |u_h^{\text{ref}} - u_h|^2 dx}.$$

Figure 2 clearly shows that $E(h) = O(h^2)$, which implies that the solution of the proposed method delivers a solution that is second-order accurate, even if the problem involves point sources and the uniform meshes are neither aligned with the discontinuities in the wave speed nor adaptively refined in the vicinity of the point source.




Figure 2: Error $E(h)$ for a solution computed for $\omega = 12\pi$ using the 2004 BP model and a point source.

To demonstrate that the proposed method does not suffer from pollution errors, we computed reference solutions $u_h^{\text{ref}}(\omega)$ for different values of ω . For these reference solutions a uniform mesh with mesh-size h is chosen so that there are 7 elements within the shortest wavelength and the polynomial degree is fixed at $p = 6$. For each of those meshes using their corresponding frequencies ω , we computed two more solutions:

one using polynomial degree $p = 1$, and another where the polynomial degree p is chosen so that $\omega^{\frac{1}{p+1}}/p \leq 1$. For each of those solutions we computed the L^2 -error $E(h, \omega)$ using the reference solution $u_h^{\text{ref}}(\omega)$ similarly to $E(h)$. We then estimated the constant $C(\omega)$ introduced in Eq. 9 as

$$C_{\text{est}}(\omega) = E(h, \omega) \left(\frac{h}{p}\right)^{-2}.$$

Figure 3 shows that if the polynomial degree is kept constant at $p = 1$, $C(\omega)$ grows as $O(\omega)$ as the frequency is increased. This shows two aspects: (i) if the polynomial degree in the hdG method is kept constant, the solution suffers from pollution errors, and (ii) the estimate in Eq. 10 is sharp. Moreover, Figure 3 also shows that if p is chosen so that $\omega^{\frac{1}{p+1}}/p \leq 1$, the pollution error can be eliminated, even in the presence of point sources and discontinuous wave speeds. In particular, we emphasize that for these examples the discontinuities in the wave speed model were not resolved by the mesh.




Figure 3: Estimated values for $C(\omega)$ for increasing ω where for one set of solutions the polynomial degree is kept at $p = 1$ and for another set increased as $O(\log \omega)$.

DISCUSSION

We presented a modification of the hdG methods introduced in (Taus et al. (2016)), which is pollution-free and second-order accurate even in the presence of discontinuous media and distributional (point source) right-hand sides.

The presented method is particularly interesting for problems involving discontinuous coefficients, in which meshing the discontinuities can be problematic. We point out that this approach can be further optimized by a parallel implementation and by pipelining the right-hand sides as explained in (Scheuer et al. (2016)). Moreover, the results can in principle be easily extended to 3D problems.

ACKNOWLEDGMENTS

The authors thank Total SA for support. LD is also supported by AFOSR grants FA9550-12-1-0328 and FA9550-15-1-0078, ONR grant N00014-16-1-2122, and NSF grant DMS-1255203.

Pollution-free and fast hdG solvers for the high-frequency Helmholtz equation

REFERENCES

- Babuska, I. M., and S. A. Sauter, 1997, Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers?: *SIAM Journal on Numerical Analysis*, **34**, 2392–2423.
- Bérenger, J.-P., 1994, A perfectly matched layer for the absorption of electromagnetic waves: *Journal of Computational Physics*, **114**, 185–200.
- Bermúdez, A., L. Hervella-Nieto, A. Prieto, and R. Rodríguez, 2007, An optimal perfectly matched layer with unbounded absorbing function for time-harmonic acoustic scattering problems: *Journal of Computational Physics*, **223**, 469 – 488.
- Bezanson, J., S. Karpinski, V. B. Shah, and A. Edelman, 2012, Julia: A fast dynamic language for technical computing: arXiv preprint arXiv:1209.5145.
- Billette, F., and S. Brandsberg-Dahl, 2005, The 2004 BP velocity benchmark.: EAGE.
- Calandra, H., S. Gratton, X. Pinel, and X. Vasseur, 2013, An improved two-grid preconditioner for the solution of three-dimensional Helmholtz problems in heterogeneous media: *Numerical Linear Algebra with Applications*, **20**, 663–688.
- Chen, Y., 1997, Inverse scattering via Heisenberg’s uncertainty principle: *Inverse Problems*, **13**, 253.
- Chen, Z., and X. Xiang, 2013, A source transfer domain decomposition method for Helmholtz equations in unbounded domain: *SIAM Journal on Numerical Analysis*, **51**, 2331–2356.
- Cockburn, B., J. Gopalakrishnan, and R. Lazarov, 2009, Unified hybridization of discontinuous galerkin, mixed, and continuous galerkin methods for second order elliptic problems: *SIAM Journal on Numerical Analysis*, **47**, 1319–1365.
- Cockburn, B., W. Qiu, and K. Shi, 2012, Conditions for superconvergence of hdg methods for second-order elliptic problems: *Mathematics of Computation*, **81**, 1327–1353.
- Engquist, B., and L. Ying, 2011, Sweeping preconditioner for the Helmholtz equation: moving perfectly matched layers: *Multiscale Modeling & Simulation*, **9**, 686–710.
- Gillman, A., A. Barnett, and P. Martinsson, 2014, A spectrally accurate direct solution technique for frequency-domain scattering problems with variable media: *BIT Numerical Mathematics*, 1–30.
- Gittelsohn, C. J., R. Hiptmair, and I. Perugia, 2009, Plane wave discontinuous Galerkin methods: Analysis of the h-version: *ESAIM: Mathematical Modelling and Numerical Analysis*, **43**, no. 02, 297–331.
- Griesmaier, R., and P. Monk, 2011, Error analysis for a hybridizable discontinuous galerkin method for the helmholtz equation: *Journal of Scientific Computing*, **49**, 291–310.
- Hiptmair, R., A. Moiola, and I. Perugia, 2015, A survey of Trefftz methods for the Helmholtz equation: ArXiv e-prints.
- Imbert-Gérard, L.-M., 2015, Interpolation properties of generalized plane waves: *Numerische Mathematik*, 1–29.
- Johnson, S., 2010, Notes on perfectly matched layers (PMLs).
- Melenk, J. M., and S. Sauter, 2011, Wavenumber explicit convergence analysis for galerkin discretizations of the helmholtz equation: *SIAM Journal on Numerical Analysis*, **49**, 1210–1243.
- Moiola, A., R. Hiptmair, and I. Perugia, 2011, Vekua theory for the Helmholtz operator: *Zeitschrift für angewandte Mathematik und Physik*, **62**, 779–807.
- Nguyen, N., J. Peraire, and B. Cockburn, 2011, High-order implicit hybridizable discontinuous galerkin methods for acoustics and elastodynamics: *Journal of Computational Physics*, **230**, 3695 – 3718.
- Nguyen, N. C., J. Peraire, F. Reitich, and B. Cockburn, 2015, A phase-based hybridizable discontinuous Galerkin method for the numerical solution of the Helmholtz equation: *J. Comput. Physics*, **290**, 318–335.
- Peterseim, D., 2016, Eliminating the pollution effect in Helmholtz problems by local subscale correction: *Math. Comp.* (Also available as INS Preprint No. 1411).
- Pratt, R. G., 1999, Seismic waveform inversion in the frequency domain; part 1: Theory and verification in a physical scale model: *Geophysics*, **64**, 888–901.
- Scheuer, A., L. Zepeda-Núñez, R. Hewett, and L. Demanet, 2016, A short note on a pipelined polarized-trace algorithm for 3D Helmholtz: *SEG Technical Program Expanded Abstracts 2016*, 3819–3824.
- Shannon, C. E., 1998, Communication in the presence of noise: *Proceedings of the IEEE*, **86**, 447–457.
- Stolk, C., 2013, A rapidly converging domain decomposition method for the Helmholtz equation: *Journal of Computational Physics*, **241**, 240–252.
- Stolk, C. C., 2015, A dispersion minimizing scheme for the 3-D Helmholtz equation with applications in multigrid based solvers: ArXiv e-prints.
- Taus, M., L. Demanet, and L. Zepeda-Núñez, 2016, A short note on a fast and high-order hybridizable discontinuous Galerkin solver for the 2D high-frequency Helmholtz equation: *SEG Technical Program Expanded Abstracts 2016*, 3835–3840.
- Turkel, E., D. Gordon, R. Gordon, and S. Tsynkov, 2013, Compact 2D and 3D sixth order schemes for the Helmholtz equation with variable wave number: *Journal of Computational Physics*, **232**, 272 – 287.
- Vion, A., and C. Geuzaine, 2014, Double sweep preconditioner for optimized Schwarz methods applied to the Helmholtz problem: *Journal of Computational Physics*, **266**, 171–190.
- Virieux, J., and S. Operto, 2009, An overview of full-waveform inversion in exploration geophysics: *GEOPHYSICS*, **74**, WCC1–WCC26.
- Wang, S., X. S. Li, X. J., Y. Situ, and M. V. de Hoop, 2013, Efficient scalable algorithms for solving dense linear systems with hierarchically semiseparable structures: *SIAM Journal on Scientific Computing*, **35**, C519–C544.
- Zepeda-Núñez, L., and L. Demanet, 2016, The method of polarized traces for the 2D Helmholtz equation: *Journal of Computational Physics*, **308**, 347 – 388.
- Zepeda-Núñez, L., and H. Zhao, 2016, Fast alternating bidirectional preconditioner for the 2D high-frequency Lippmann–Schwinger equation: *SIAM Journal on Scientific Computing*, **38**, B866–B888.