WAVEHOLTZ: ITERATIVE SOLUTION OF THE HELMHOLTZ EQUATION VIA THE WAVE EQUATION

DANIEL APPELÖ*, FORTINO GARCIA[†], AND OLOF RUNBORG[‡]

Abstract. A new iterative method, the WaveHoltz iteration, for solution of the Helmholtz equation is presented. WaveHoltz is a fixed-point iteration that filters the solution to the solution of wave equation with time-periodic forcing and boundary data. The WaveHoltz iteration corresponds to a linear and coercive operator which, after discretization, can be recast as a positive definite linear system of equations. The solution to this system of equations approximates the Helmholtz solution and can be accelerated by Krylov subspace techniques. Analysis of the continuous and the discrete cases are presented as are numerical experiments.

Key words. Wave equation, Helmholtz equation

AMS subject classifications. 65M22, 65M12

1. Introduction. The defining feature of waves are their ability to propagate over large distances without changing their shape. It is this property that allows them to carry information which underpins all communication, be it through speech or electronic transmission of data. Waves can also be used to probe the interior of the earth, the human body or engineering structures like buildings or bridges. This probing can be turned into images of the interior by the means of solving inverse problems. Harnessing the nature of waves requires high-order accurate and efficient numerical methods that are able to simulate wave propagation in three dimensions and over long distances. For cutting edge problems in scientific and engineering research such simulations must be carried out on parallel high-performance computing platforms and thus the numerical methods must scale while being easy to implement and generally applicable.

In this paper we focus on approximating solutions to the scalar wave equation in the frequency domain, i.e. the Helmholtz equation

(1.1)
$$\nabla \cdot (c^2(x)\nabla u) + \omega^2 u = f(x).$$

However, to to obtain such solutions we will use time domain discretizations of the wave equation. The motivation for developing high order accurate and scalable Helmholtz solvers comes from both mathematics and applications. On the mathematics side the recent results by Engquist and Zhao [14] give sharp lower bounds on the number of terms in a separated representation approximation of the Green's function of the Helmholtz equation as a function of the frequency (wavenumber). These bounds limit the applicability of the state of the art sweeping preconditioners in the high frequency regime and, for example, for interior and wave guide problems. Motivation also comes from applications in seismology, optics and acoustics. For example in full waveform inversion the problems are very large and the robustness of the inversion process can be enhanced by combining frequency and time domain inversion in a multi-scale fashion to avoid getting trapped in local minima.

Designing efficient iterative solvers for the Helmholtz equation (1.1) is notoriously difficult and has been the subject of much research (for detailed reviews see Ernst and Gander, [17], Gander and Zhang [19], and Erlangga, [15]). The main two difficulties in solving the Helmholtz equation are the resolution requirements and the highly indefinite character of the discretized system of equations.

Assuming that (1.1) has been scaled so that the mean of c(x) is about 1 then the typical wavelength is $\lambda = 2\pi/\omega$ and the typical wavenumber is $\omega/2\pi$. In order to numerically propagate solutions to the time dependent wave equation corresponding to (1.1) with small errors it is crucial to control the dispersion by

Funding: Supported in part by STINT initiation grant IB2019-8154.

^{*}Department of Applied Mathematics, University of Colorado, Boulder, CO 80309-0526 (Daniel.Appelo@colorado.edu)

Funding: Supported in part by STINT initiation grant IB2019-8154 and NSF Grant DMS-1913076. Any conclusions or recommendations expressed in this paper are those of the author and do not necessarily reflect the views of the NSF.

[†]Department of Applied Mathematics, University of Colorado, Boulder, CO 80309-0526 (Fortino.Garcia@colorado.edu)

Funding: Supported in part by STINT initiation grant IB2019–8154 and NSF Grant DMS-1913076. Any conclusions or recommendations expressed in this paper are those of the author and do not necessarily reflect the views of the NSF

[‡] Department of Mathematics, KTH, 100 44, Stockholm (olofr@kth.se)

using high order methods. The basic estimate by Kreiss and Oliger [29] shows that in order to propagate a wave over J wavelengths with a pth order finite difference method and with an error no greater than ϵ one must choose the number of points per wavelength PPW(J, p) as

$$PPW(J, p) \ge C(p, \epsilon)J^{\frac{1}{p}}$$
.

Here $C(p,\epsilon)$ depends on the tolerance ϵ but decreases with increasing order of accuracy p. Consequently, for a problem in d-dimensions and with fixed physical size the number of wavelengths in the domain will scale as ω^d and to maintain a fixed tolerance the total number of degrees of freedom needed, $N_p(\omega) = \mathcal{O}(\omega^{d(1+\frac{1}{p})})$, is very large for high frequencies.

The dependence on p and ω in $N_p(\omega)$ immediately reveals two fundamental criteria for designing high frequency Helmholtz solvers:

- 1. The solvers must be **parallel**, **memory lean** and they must **scale well**. In 3D the number of degrees of freedom representing the solution cannot be stored on a single computer, and even on a parallel computer it is important to preserve the sparsity of the discrete version of (1.1).
- 2. The underlying discretizations must be **high order accurate**. At high frequencies and in 3D the extra penalty due to pollution / dispersion errors becomes prohibitive.

Further, the linear system matrix, A, resulting from direct discretization of (1.1) is indefinite so that the robust and easy to implement preconditioned conjugate gradient (PCG) method cannot be used. Instead the method of necessity becomes the preconditioned generalized minimal residual method (GMRES). To efficiently precondition GMRES one must exploit the intrinsic properties of the wave equation. The oscillatory nature of the Helmholtz Green's function and its discrete counterpart A^{-1} can only be well approximated if the (unconditioned) Krylov subspace is allowed to grow quite large (with "large" scaling adversely with the frequency ω , [17]). The slow growth of the "spanning power" of the Krylov vectors is due to the underlying local connectivity of the discretization, preventing information to propagate rapidly. Efficient preconditioners must thus accelerate the propagation of information or reduce the cost of each iteration. Without preconditioners the iteration typically stagnates.

Perhaps the first contribution that aimed to improve the propagation of information was the Analytic Incomplete LU preconditioner (AILU) by Gander and Nataf [18]. The AILU preconditioner finds an LDL^T factorization from an approximation of the same pseudodifferential operators that are used to construct non-reflecting boundary conditions [11, 1, 24] and sweeps forward then backward along one of the coordinate directions in a structured grid.

The pioneering works on sweeping preconditioners by Engquist and Ying [12, 13] were major break-throughs in the solution of the Helmholtz equation. Similar to the AILU, the preconditioners in [12, 13] use a LDL^T decomposition but exploit the low rank properties of off-diagonal blocks together with perfectly matched layers to obtain solvers that converge in a small number of GMRES iterations. The papers [12, 13] were the two first instances of iterative Helmholtz solvers that converge in a small number of iterations that is almost independent of frequency.

Once it had been established that low rank approximations, combined with clever use of sweeping and perfectly matched layers (PML), could be used to find Helmholtz solvers with linear scaling then many extensions and specializations were constructed. For example, in [34] Stolk introduced a domain decomposition method with transmission conditions based on the perfectly matched layer (PML) that is able to achieve near linear scaling. Chen and Xiang, [10], and Vion and Geuzaine, [36], also considered sweeping domain decomposition method combined with PML and showed that their methods could be used as efficient preconditioners for the Helmholtz equation. The method of polarized traces by Zepeda-Núñez, Demanet and co-authors, [41, 40, 39], is a two step sweeping preconditioner that compresses the traces of the Greens function in an offline computation and utilizes incomplete Green's formulas to propagate the interface data. See also the recent review by Gander and Zhang [19] for connections between sweeping methods.

Alongside iterative methods there are some attractive direct and multigrid methods. Examples from the class of direct methods are the Hierarchically Semi-Separable (HSS) parallel multifrontal sparse solver by deHoop and co-authors, [37], the spectral collocation solver by Gillman, Barnett and Martinsson, [20], and the p-FEM approach of Bériot, Prinn and Gabard, [6], which utilizes an $a\ priori$ error indicator to choose the polynomial order of each element . Notable examples of multigrid methods are the Wave-ray method by Brandt and Livshits [8, 31] and the shifted Laplacian preconditioner with multigrid by Erlangga et al. [16].

As mentioned previously, the invention of sweeping preconditioners was a breakthrough and it is likely that they will have lasting and continuing impacts for the solution of the Helmholtz equation in various settings. There are, however, some limitations. First, in the recent paper [14], Engquist and Zhao provide precise lower bounds on how the number of terms that are needed to approximate the Helmholtz Green's function depends on the frequency. In particular, for the high frequency regime they show that for interior problems and waveguides the rank of the off-diagonal elements grows fast, rendering sweeping preconditioners less efficient. They also show that the situation is, in general, worse in 3D than in 2D. This lack of compressibility may, in cases of practical importance, increase the cost of both the factorization and compression as well as the application of the compressed preconditioner. We note that this loss of compressibility at high frequency will also prevent direct methods such as [37, 20, 6] from reaching their most efficient regimes. An additional drawback of direct methods is their memory consumption for 3D problems.

Another potential drawback with the sweeping methods is the long setup times before the solve. Of course all of the algorithms above do not suffer from this deficiency but many of them do. This may not be problematic when considering a background velocity that does not change but this is not the case, for example, when inverting for material parameters. In this case the velocity model will change constantly, necessitating a costly factorization in each update.

Finally, the two criterions 1.) and 2.) above are not so easy to meet for sweeping preconditioners. The sweep itself is intrinsically sequential and although there have been at least partially successful attempts to parallelize the sweeping methods it is hard to say that they are easy to parallelize in a scalable way. In a similar vein most of the methods use (and some rely on) low order discretizations. Although it is possible to use higher order accurate discretizations together with sweeping preconditioners, their scarcity in the literature is noticeable.

Another approach that is somewhat popular in the engineering literature is to simply run the wave equation for a long time to get a Helmholtz solution, see e.g. [27]. The theoretical underpinning of this approach is the *limiting amplitude principle* which says that every solution to the wave equation with an oscillatory forcing, in the exterior of a domain with reflecting boundary conditions tends to the Helmholtz solution. However, since the limiting amplitude principle only holds for exterior problems this approach does not work for interior problems and becomes very slow for problems with trapping waves. See e.g. the articles by Ladyzhenskaya [30], Morawetz [33] and Vainberg [35].

An alternative approach, the so called Controllability Method (CM), was originally proposed by Bristeau et al. [9]. In the CM the solution to the Helmholtz equation is found by solving a convex constrained least-squares minimization problem where the deviation from time-periodicity is minimized in the classic wave equation energy. The basic ingredients in an iteration step in CM are: a.) the solution of a forward wave and a backward wave equation over one time-period, and b.) the solution of a symmetric coercive elliptic (and wave number independent) problem.

In [9] and the later spectral element implementations of CM by Heikkola et al. [26, 25] only sound-soft scatterers were considered. For more general boundary conditions the minimizer of the cost functional of [9] is not unique but alternative cost functionals that does guarantee uniqueness (and thus convergence to the Helmholtz solution) were recently proposed by Grote and Tang in [22]. We also note that if the wave equation is formulated as a first order system it is possible to avoid solving the elliptic problem [21, 23].

In what follows we will present an alternative to the controllability method. Our method, which we call the WaveHoltz Iteration method (WHI), only requires a single forward wave equation solve and no elliptic solves but produces a positive definite (and sometimes symmetric) iteration that can be accelerated by, e.g. the conjugate gradient method or other Krylov subspace methods. As the WaveHoltz iteration is built from a time domain wave equation solver we claim and hope to demonstrate that it meets both criterion 1. and 2. above.

The rest of the paper is organized as follows. In Section 2 we present and analyze our method and its extensions, in Section 3 we briefly outline the numerical methods we use to solve the wave equation, in Section 4 we present numerical experiments, and in Section 5 we summarize and conclude.

Before proceeding we would like to acknowledge that although our method is distinct from the controllability method, it was the work by Grote and Tang, [22], that introduced us to CM and inspired us to derive the method discussed below.

2. WaveHoltz: A New Method for Designing Scalable Parallel Helmholtz Solvers. We consider the Helmholtz equation in a bounded open smooth domain Ω ,

(2.1)
$$\nabla \cdot (c^2(x)\nabla u) + \omega^2 u = f(x), \qquad x \in \Omega,$$

with boundary condtions of the type

(2.2)
$$i\alpha\omega u + \beta(c^2(x)\vec{n}\cdot\nabla u) = 0, \quad \alpha^2 + \beta^2 = 1, \qquad x \in \partial\Omega.$$

We assume $f \in L^2(\Omega)$ and that $c \in L^{\infty}(\Omega)$ with the bounds $0 < c_{\min} \le c(x) \le c_{\max} < \infty$ a.e. in Ω . Away from resonances, this ensures that there is a unique weak solution $u \in H^1(\Omega)$ to (2.1). Due to the boundary conditions u is in general complex valued.

We first note that the function $w(t,x) := u(x) \exp(i\omega t)$ is a $T = 2\pi/\omega$ -periodic (in time) solution to the forced scalar wave equation

$$w_{tt} = \nabla \cdot (c^2(x)\nabla w) - f(x)e^{i\omega t}, \quad x \in \Omega, \quad 0 \le t \le T,$$

$$w(0,x) = v_0(x), \quad w_t(0,x) = v_1(x),$$

$$\alpha w_t + \beta(c^2(x)\vec{n}\cdot\nabla w) = 0, \quad x \in \partial\Omega,$$

$$(2.3)$$

where $v_0 = u$ and $v_1 = i\omega u$. Based on this observation, our approach is to find this w instead of u. We could thus look for initial data v_0 and v_1 such that w is a T-periodic solution to (2.3). However, there may be several such w, see [22], and we therefore impose the alternative constraint that a certain time-average of w should equal the initial data. More precisely, we introduce the following operator acting on the initial data $v_0 \in H^1(\Omega)$, $v_1 \in L^2(\Omega)$,

$$\Pi \left[\begin{array}{c} v_0 \\ v_1 \end{array} \right] = \frac{2}{T} \int_0^T \left(\cos(\omega t) - \frac{1}{4} \right) \left[\begin{array}{c} w(t,x) \\ w_t(t,x) \end{array} \right] dt, \quad T = \frac{2\pi}{\omega},$$

where w(t, x) and its time derivative $w_t(t, x)$ satisfies the wave equation (2.3) with initial data v_0 and v_1 . The result of $\Pi[v_0, v_1]^T$ can thus be seen as a filtering in time of $w(\cdot, x)$ around the ω -frequency. We will further motivate the choice of time averaging in the analysis below. By construction, the solution u of Helmholtz now satisfies the equation

The WaveHoltz method then amounts to solving this equation with the fixed point iteration

(2.6)
$$\left[\begin{bmatrix} v \\ v' \end{bmatrix}^{(n+1)} = \Pi \begin{bmatrix} v \\ v' \end{bmatrix}^{(n)}, \quad \left[\begin{matrix} v \\ v' \end{matrix} \right]^{(0)} \equiv 0. \right]$$

Provided this iteration converges and the solution to (2.5) is unique, we obtain the Helmholtz solution as $u = \lim_{n \to \infty} v^n$.

Remark 1. Note that each iteration is inexpensive and that T is reduced by the reciprocal of ω as ω grows. If we assume that the number of degrees of freedom in each dimension scales with ω and that we evolve the wave equation with an explicit method this means that the number of timesteps per iteration is independent of ω . Also note that the iteration is trivial to implement (in parallel or serial) if there is already a time domain wave equation solver in place. The integral in the filtering is carried out independently for each degree of freedom and simply amounts to adding up a weighted sum (e.g. a trapezoidal sum) of the solution one timestep at a time. Finally, note that WHI allows all the advanced techniques that have been developed for wave equations (e.g. local timestepping, non-conforming discontinuous Galerkin finite elements h- and p-adaptivity etc.) can be transferred to the Helmholtz equation and other time harmonic problems.

2.1. Iteration for the Energy Conserving Case. Here we consider boundary conditions of either Dirichlet ($\beta=0$) or Neumann ($\alpha=0$) type. This is typically the most difficult case for iterative Helmholtz solvers when Ω is bounded. The wave energy is preserved in time and certain ω -frequencies in Helmholtz are resonant, meaning they equal an eigenvalue of the operator $-\nabla \cdot (c^2(x)\nabla)$. Moreover, the limiting amplitude principle does not hold, and one can thus not obtain the Helmholtz solution by solving the wave equation over a long time interval.

We start by introducing a simplified iteration for this case. With the given boundary conditions the solution to Helmholtz will be real valued, since f is a real valued function. Without loss of generality, we may then take $w_t(0,x) = 0$ and $w(t,x) = u(x)\cos(\omega t)$, since for a T-periodic real valued solution there is a time when $w_t(0,x) = 0$. We choose that time as the initial time so that (2.3) becomes

$$w_{tt} = \nabla \cdot (c(x)^2 \nabla w) - f(x) \cos(\omega t), \quad x \in \Omega, \quad 0 \le t \le T,$$

$$w(0, x) = v(x), \quad w_t(0, x) \equiv 0,$$

$$\alpha w_t + \beta(c^2(x)\vec{n} \cdot \nabla w) = 0, \quad x \in \partial\Omega.$$

The simplified iteration is then defined as

$$(2.8) v^{n+1} = \Pi v^n, v^0 \equiv 0,$$

where

(2.9)
$$\Pi v = \frac{2}{T} \int_0^T \left(\cos(\omega t) - \frac{1}{4} \right) w(t, x) dt, \qquad T = \frac{2\pi}{\omega},$$

with w(t,x) solving the wave equation (2.7) with initial data $v=v^n\in H^1(\Omega)$. We now analyze this iteration. By the choice of boundary conditions the operator $-\nabla\cdot(c^2(x)\nabla)$ has a point spectrum with non-negative eigenvalues with corresponding eigenfunctions that form an orthonormal basis of $L^2(\Omega)$. Denote those eigenmodes $(\lambda_j^2,\phi_j(x))$, with $\|\phi_j\|_{L^2(\Omega)}=1$. We assume that the angular frequency ω is not a resonance, i.e. $\omega^2\neq\lambda_j^2$ for all j. The Helmholtz equation (2.1) is then wellposed.

We recall that for any $q \in L^2(\Omega)$ we can expand

$$q(x) = \sum_{j=0}^{\infty} \hat{q}_j \phi_j(x),$$

for some coefficients \hat{q}_i and

$$||q||_{L^2(\Omega)}^2 = \sum_{j=0}^{\infty} |\hat{q}_j|^2, \qquad c_{\min}^2 ||\nabla q||_{L^2(\Omega)}^2 \le \sum_{j=0}^{\infty} \lambda_j^2 |\hat{q}_j|^2 \le c_{\max}^2 ||\nabla q||_{L^2(\Omega)}^2.$$

We start by expanding the Helmholtz solution u, the initial data v to the wave equation (2.7), and the forcing f in this way,

$$u(x) = \sum_{j=0}^{\infty} \hat{u}_j \phi_j(x), \quad v(x) = \sum_{j=0}^{\infty} \hat{v}_j \phi_j(x), \quad f(x) = \sum_{j=0}^{\infty} \hat{f}_j \phi_j(x).$$

Then,

$$-\lambda_j^2 \hat{u}_j + \omega^2 \hat{u}_j = \hat{f}_j \quad \Rightarrow \quad \hat{u}_j = \frac{\hat{f}_j}{\omega^2 - \lambda_i^2}.$$

For the wave equation solution w(t,x) with initial data w=v and $w_t=0$ we have

(2.10)
$$w(t,x) = \sum_{j=0}^{\infty} \hat{w}_j(t)\phi_j(x), \qquad \hat{w}_j(t) = \hat{u}_j\left(\cos(\omega t) - \cos(\lambda_j t)\right) + \hat{v}_j\cos(\lambda_j t).$$




Fig. 2.1: The filter transfer function β for $\omega = 10$.

The filtering step (2.9) then gives

$$\Pi v = \sum_{j=0}^{\infty} \bar{v}_j \phi_j(x), \qquad \bar{v}_j = \hat{u}_j (1 - \beta(\lambda_j)) + \hat{v}_j \beta(\lambda_j),$$

where

$$\beta(\lambda) := \frac{2}{T} \int_0^T \left(\cos(\omega t) - \frac{1}{4} \right) \cos(\lambda t) dt.$$

We introduce the linear operator $S: L^2(\Omega) \to L^2(\Omega)$,

(2.11)
$$\mathcal{S}\sum_{j=0}^{\infty} \hat{u}_j \phi_j(x) := \sum_{j=0}^{\infty} \beta(\lambda_j) \hat{u}_j \phi_j(x),$$

which gives the filtered solution of the wave equation with f = 0, when applied to the initial data v. We can then write the iteration as

(2.12)
$$v^{n+1} = \Pi v^n = \mathcal{S}(v^n - u) + u.$$

The operator \mathcal{S} is self-adjoint and has the same eigenfunctions $\phi_j(x)$ as $-\nabla \cdot (c^2(x)\nabla)$ but with the (real) eigenvalues $\beta(\lambda_j)$. The convergence properties of the iteration depend on these eigenvalues and it is therefore of interest to study the range of the filter transfer function β . Figure 2.1 shows a plot of β which indicates that the eigenvalues of \mathcal{S} are inside the unit interval, with a few of them being close to 1 (when $\lambda_j \approx \omega$), and most of them being close to zero (when $\lambda_j \gg \omega$). In the appendix we show the following lemma about β .

LEMMA 2.1. The filter transfer function β satisfies $\beta(\omega) = 1$ and

$$0 \le \beta(\lambda) \le 1 - \frac{1}{2} \left(\frac{\lambda - \omega}{\omega} \right)^2, \qquad \text{when } \left| \frac{\lambda - \omega}{\omega} \right| \le \frac{1}{2},$$
$$|\beta(\lambda)| \le \frac{1}{2}, \qquad \text{when } \left| \frac{\lambda - \omega}{\omega} \right| \ge \frac{1}{2},$$
$$|\beta(\lambda)| \le b_0 \frac{\omega}{\lambda - \omega}, \qquad \text{when } \lambda > \omega.$$

where $b_0 = \frac{3}{4\pi}$. Moreover, close to ω we have the local expansion

(2.13)
$$\beta(\omega + r) = 1 - b_1 \left(\frac{r}{\omega}\right)^2 + R(r/\omega) \left(\frac{r}{\omega}\right)^3, \qquad b_1 = \frac{2\pi^2}{3} - \frac{1}{4} \approx 6.33, \quad ||R||_{\infty} \le \frac{5\pi^3}{6}.$$

REMARK 2. It is easy to see that $\beta(\omega) = 1$ for any constant besides 1/4. The particular choice 1/4 is made to ensure that $\beta'(\omega) = 0$, which is necessary to keep $\beta \leq 1$ in a neighborhood of ω . We explore other possibilities in Section 2.3.

From this lemma we can derive some results for the operator S. To do this we first quantify the non-resonance condition. We let

 $\delta_j = \frac{\lambda_j - \omega}{\omega},$

be the relative size of the gap between λ_j and the Helmholtz frequency, and then denote the smallest gap (in magnitude) by δ ,

$$\delta = \delta_{j^*}, \quad j^* = \operatorname{argmin}_{j} |\delta_{j}|.$$

Then we have

LEMMA 2.2. Suppose $\delta > 0$. The spectral radius ρ of S is strictly less than one, and for small δ ,

(2.14)
$$\rho = 1 - b_1 \delta^2 + O(\delta^3),$$

with b_1 as in Lemma 2.1. Moreover, S is a bounded linear map from $L^2(\Omega)$ to $H^1(\Omega)$.

Proof. From Lemma 2.1 we get

$$\rho = \sup_{j} |\beta(\lambda_j)| \le \sup_{j} \max\left(1 - \frac{1}{2}\delta_j^2, \frac{1}{2}\right) \le \max\left(1 - \frac{1}{2}\delta^2, \frac{1}{2}\right) < 1.$$

For the more precise estimate when δ is small we will use (2.13). Since $1 > \rho \ge \beta(\omega + \omega\delta) \to 1$ as $\delta \to 0$, we can assume that $\rho > 1 - \eta^2/2$, with $\eta := b_1/2||R||_{\infty}$, for small enough δ . Then, since $|\beta(\omega + \omega\delta_j)| \le 1 - \eta^2/2$ for $|\delta_j| > \eta$ by Lemma 2.1, we have

$$\rho = \sup_{|\delta_j| \le \eta} \beta(\omega + \omega \delta_j) = \beta(\omega + \omega \delta_{k^*}),$$

for some k^* with $|\delta_{k^*}| \leq \eta$. If $\delta_{k^*} = \delta_{j^*}$ (where $\delta = |\delta_{j^*}|$) then (2.13) gives (2.14). If not, we have $\eta \geq |\delta_{k^*}| \geq \delta$ and by Lemma 2.1

$$0 \le \beta(\omega + \omega \delta_{k^*}) - \beta(\omega + \omega \delta_{j^*}) = -b_1(\delta_{k^*}^2 - \delta^2) + R(\delta_{k^*})\delta_{k^*}^3 - R(\delta_{j^*})\delta_{j^*}^3 \le -b_1(\delta_{k^*}^2 - \delta^2) + \frac{b_1}{2}(\delta_{k^*}^2 + \delta^2),$$

which implies that $\delta_{k^*}^2 \leq 3\delta^2$ and that

$$0 \le b_1(\delta_{k^*}^2 - \delta^2) \le R(\delta_{k^*})\delta_{k^*}^3 - R(\delta_{j^*})\delta_{j^*}^3 \le ||R||_{\infty}(1 + 3\sqrt{3})\delta^3.$$

Therefore,

$$\rho = 1 - b_1 \delta_{k^*}^2 + O(\delta_{k^*}^3) = 1 - b_1 \delta^2 + b_1 (\delta^2 - \delta_{k^*}^2) + O(\delta_{k^*}^3) = 1 - b_1 \delta^2 + O(\delta_{k^*}^3 + \delta^3) = 1 - b_1 \delta^2 + O(\delta^3).$$

This shows (2.14). For the second statement, we note first that by (2.1),

$$|\lambda_j \beta(\lambda_j)| \le \omega \begin{cases} 1, & \lambda_j \le \omega, \\ \frac{b_0 \lambda_j}{\lambda_j - \omega}, & \lambda_j > \omega, \end{cases} = \omega \begin{cases} 1, & \lambda_j \le \omega, \\ b_0 (1 + 1/\delta_j), & \lambda_j > \omega, \end{cases} \le \omega \min(1, b_0 (1 + 1/|\delta|)) =: D.$$

Suppose now that $g \in L^2(\Omega)$ and

$$g(x) = \sum_{j=0}^{\infty} \hat{g}_j \phi_j(x).$$

Then

$$||\mathcal{S}g||_{H^1(\Omega)}^2 \leq \sum_{j=0}^{\infty} |\beta(\lambda_j)|^2 |\hat{g}_j|^2 + \sum_{j=0}^{\infty} \frac{\lambda_j^2}{c_{\min}^2} |\beta(\lambda_j)|^2 |\hat{g}_j|^2 \leq \left(1 + \frac{D^2}{c_{\min}^2}\right) \sum_{j=0}^{\infty} |\hat{g}_j|^2 = \left(1 + \frac{D^2}{c_{\min}^2}\right) ||g||_{L^2(\Omega)}^2.$$

This proves the lemma.

Letting $e^n := u - v^n$ we can rearrange (2.12) and obtain

$$e^{n+1} = \mathcal{S}e^n \quad \Rightarrow \quad ||e^{n+1}||_{L^2(\Omega)} \le \rho ||e^n||_{L^2(\Omega)} \quad \Rightarrow \quad ||e^n||_{L^2(\Omega)} \le \rho^n ||e^0||_{L^2(\Omega)} \to 0,$$

which shows that v^n converges to u in L^2 . By Lemma 2.2 all iterates $v^n \in H^1(\Omega)$ since $v^0 = 0$. We can therefore also get convergence in H^1 . Let

$$e^{n}(x) = \sum_{j=0}^{\infty} \hat{e}_{j}^{n} \phi_{j}(x),$$

and consider similarly

$$\begin{split} &\sum_{j=0}^{\infty} |\hat{e}_{j}^{n+1}|^{2} \lambda_{j}^{2} = \sum_{j=0}^{\infty} \beta(\lambda_{j})^{2} |\hat{e}_{j}^{n}|^{2} \lambda_{j}^{2} \leq \rho^{2} \sum_{j=0}^{\infty} |\hat{e}_{j}^{n}|^{2} \lambda_{j}^{2} \quad \Rightarrow \\ &||\nabla e^{n}||_{L^{2}(\Omega)}^{2} \leq \frac{1}{c_{\min}^{2}} \sum_{j=0}^{\infty} |\hat{e}_{j}^{n}|^{2} \lambda_{j}^{2} \leq \frac{\rho^{2n}}{c_{\min}^{2}} \sum_{j=0}^{\infty} |\hat{e}_{j}^{0}|^{2} \lambda_{j}^{2} \leq \rho^{2n} \frac{c_{\max}^{2}}{c_{\min}^{2}} ||\nabla e^{0}||_{L^{2}(\Omega)}^{2} \rightarrow 0. \end{split}$$

We conclude that the iteration converges in H^1 with convergence rate ρ . By Lemma 2.1 we have $\rho \sim 1-6.33\delta^2$ and, not surprisingly, the smallest gap, δ , determines the convergence factor. We have thus showed

THEOREM 2.3. The iteration in (2.8) and (2.9) converges in $H^1(\Omega)$ for the Dirichlet and Neumann problems away from resonances to the solution of the Helmholtz equation (2.1). The convergence rate is $1 - O(\delta^2)$, where δ is the minimum gap between ω and the eigenvalues of $-\nabla \cdot (c^2(x)\nabla)$.

As discussed in the introduction, the dependence of the convergence rate on ω is often of interest. For the energy conserving case, however, this question is ambiguous as the problem is not well-defined for all ω . As soon as $\omega = \lambda_j$ there are either no or an infinite number of solutions. In higher dimensions, the eigenvalues λ_j get denser as j increases, meaning that in general the problem will be closer and closer to resonance as ω grows. Therefore, solving the interior undamped Helmholtz equation for high frequencies, with pure Dirichlet or Neumann boundary conditions, may not be of great practical interest.

Nevertheless, we can make the following analysis. By the work of Weyl [38] we know that the eigenvalues grow asymptotically as $\lambda_j \sim j^{1/d}$ in d dimensions. The average minimum gap δ when $\omega \approx \lambda_j$ is then

$$\delta \approx \frac{1}{\lambda_{j+1} - \lambda_j} \int_{\lambda_j}^{\lambda_{j+1}} \frac{\min(\lambda - \lambda_j, \lambda_{j+1} - \lambda)}{\omega} d\lambda = \frac{\lambda_{j+1} - \lambda_j}{4\omega} \sim \frac{(j+1)^{1/d} - j^{1/d}}{\omega} \approx \frac{j^{1/d-1}}{d\omega} \sim \frac{\omega^{1-d}}{\omega} \sim \omega^{-d}.$$

When the convergence rate is $1 - O(\delta^2)$, the number iterations to achieve a fixed accuracy grows as $O(1/\delta^2)$. This shows that the number of iterations would grow at the unacceptable rate ω^{2d} for the iteration.

Fortunately, one can accelerate the convergence by using the conjugate gradient method in the energy conserving case and with any other Krylov method in the general case. The linear system that we actually want to solve is

$$(I - S)v =: Av = b := \Pi 0.$$

Moreover, with $b = \Pi 0$ pre-computed we can easily evaluate the action of \mathcal{A} at the cost of a single wave solve. Precisely, since $\mathcal{A}v = v - \Pi v + b$ we simply carry out the evaluation of $\mathcal{A}v$ by evolving the wave equation for one period in time with v as the initial data and then subtract the filtered solution from the sum of the initial data and the right hand side b.

The operator \mathcal{A} is self adjoint and positive, since $-1/2 < \beta(\lambda_j) < 1$, which implies that the eigenvalues of \mathcal{A} lie in the interval (0,3/2). The condition number of \mathcal{A} is of the same order as $1-\rho$, where ρ is the spectral radius of \mathcal{S} , i.e. by the simple analysis above, $\operatorname{cond}(\mathcal{A}) \sim \omega^{2d}$. If this system is solved using the (unconditioned) conjugate gradient method the convergence rate is $1-1/\sqrt{\operatorname{cond}(\mathcal{A})} \sim 1-1/\omega^d$, [7]. Thus, then the method just requires $\sim \omega^d$ iterations for fixed accuracy.

Remark 3. The operator A is self-adjoint and coercive when $\delta > 0$ since

$$\langle \mathcal{A}u, u \rangle = \langle (I - \mathcal{S})u, u \rangle = \sum_{j=0}^{\infty} (1 - \beta(\lambda_j)) |\hat{u}_j|^2 \ge (1 - \rho) ||u||^2.$$

This should be contrasted with the original indefinite Helmholtz problem, which is not coercive. In fact, the eigenvalues satisfy the simple relation $\lambda_{WHI} = 1 - \beta(\lambda_{Helmholtz} + \omega) > 0$. The two formulations are however mathematically equivalent for the interior Dirichlet and Neumann problems away from resonances, as the analysis above shows.

The coercivity also implies that the solution to (2.5) for the simplified iteration is unique since $w = \Pi w$ is equivalent to A(w - u) = 0.

Remark 4. A discretization would have approximately a fixed number of grid points per wavelength, leading to a (sparse) matrix of size $N \times N$ with $N \sim \omega^d$. Hence, the number of iterations for WHI would be $O(N^2)$ and the total cost $O(N^3)$ since each iteration costs O(N). This should be compared with a direct solution method which is better than $O(N^3)$ when the matrix is sparse.

Remark 5. In the Krylov accelerated case this analysis suggests that the number of iterations would now be O(N) and the total cost $O(N^2)$. However, in the experiments below we observe slightly better complexity for interior problems and significantly better complexity for open problems. In fact, for the open problems we find that, in both two and three dimensions, the number of iterations scale as $\sim \omega$ which is the required number of iterations for the information to travel through the domain.

2.2. Analysis of the Discrete Iteration. To better understand the effects of discretizations we consider the following discrete version of the algorithm for the energy conserving case described above in Section 2.1. We introduce the temporal grid points $t_n = n\Delta t$ and a spatial grid with N points together with the vector $w^n \in \mathbb{R}^N$ containing the grid function values of the approximation at $t = t_n$. We also let $f \in \mathbb{R}^N$ hold the corresponding values of the right hand side. The discretization of the continuous spatial operator $-\nabla \cdot (c^2(x)\nabla)$, including the boundary conditions, is denoted L_h and it can be represented as an $N \times N$ matrix. The values $-\nabla \cdot (c^2(x)\nabla w)$ are then approximated by $L_h w^n$. As in the continuous case, we assume L_h has the eigenmodes (λ_j^2, ϕ_j) , such that $L_h \phi_j = \lambda_j^2 \phi_j$ for $j = 1, \ldots, N$, where all λ_j are strictly positive and ordered as $0 \le \lambda_1 \le \ldots \le \lambda_N$.

We let the Helmholtz solution u be given

$$-L_h u + \omega^2 u = f.$$

The numerical approximation of the iteration operator is denoted Π_h , and it is implemented as follows. Given $v \in \mathbb{R}^N$, we use the leap frog method to solve the wave equation as

(2.15)
$$w^{n+1} = 2w^n - w^{n-1} - \Delta t^2 L_h w^n - \Delta t^2 f \cos(\omega t_n),$$

with time step $\Delta t = T/M$ for some integer M, and initial data

$$w^{0} = v,$$
 $w^{-1} = v - \frac{\Delta t^{2}}{2}(L_{h}v + f).$

The trapezoidal rule is then used to compute $\Pi_h v$,

(2.16)
$$\Pi_h v = \frac{2\Delta t}{T} \sum_{n=0}^M \eta_n \left(\cos(\omega t_n) - \frac{1}{4} \right) w^n, \qquad \eta_n = \begin{cases} \frac{1}{2}, & n = 0 \text{ or } n = M, \\ 1, & 0 < n < M. \end{cases}$$

With these definitions we can prove

THEOREM 2.4. Suppose there are no resonances, such that $\delta_h = \min_j |\lambda_j - \omega|/\omega > 0$. Moreover, assume that Δt satisfies the stability and accuracy requirements

(2.17)
$$\Delta t < \frac{2}{\lambda_N + 2\omega/\pi}, \qquad \Delta t\omega \le \min(\delta_h, 1).$$

Then the fixed point iteration $v^{(k+1)} = \Pi_h v^{(k)}$ with $v^{(0)} = 0$ converges to v^{∞} which is a solution to the discretized Helmholtz equation with the modified frequency $\tilde{\omega}$,

$$-L_h v^{\infty} + \tilde{\omega}^2 v^{\infty} = f, \qquad \tilde{\omega} = 2 \frac{\sin(\Delta t \omega/2)}{\Delta t}.$$

The convergence rate is at least $\rho_h = \max(1 - 0.3\delta_h^2, 0.6)$.

Proof. We expand all functions in eigenmodes of L_h ,

$$w^n = \sum_{j=1}^N \hat{w}_j^n \phi_j, \qquad f = \sum_{j=1}^N \hat{f}_j \phi_j, \qquad u = \sum_{j=1}^N \hat{u}_j \phi_j, \qquad v = \sum_{j=1}^N \hat{v}_j \phi_j, \qquad v^\infty = \sum_{j=1}^N \hat{v}_j^\infty \phi_j.$$

Then the Helmholtz eigenmodes of u and v^{∞} satisfy

$$\hat{u}_j = \frac{\hat{f}_j}{\omega^2 - \lambda_j^2}, \qquad \hat{v}_j^\infty = \frac{\hat{f}_j}{\tilde{\omega}^2 - \lambda_j^2}.$$

We note that $\tilde{\omega}$ is not resonant and \hat{v}_i^{∞} is well-defined for all j, since by (C.2) and (2.17)

$$|\tilde{\omega} - \lambda_j| \ge |\omega - \lambda_j| - |\tilde{\omega} - \omega| \ge \omega \delta_h - \frac{\Delta t^2 \omega^3}{24} \ge \omega \left(\delta_h - \frac{1}{24} \min(\delta_h, 1)^2 \right) > 0.$$

The wave solution eigenmodes are given by the difference equation

(2.18)
$$\hat{w}_j^{n+1} - 2\hat{w}_j^n + \hat{w}_j^{n-1} + \Delta t^2 \lambda_j^2 \hat{w}_j^n = -\Delta t^2 \hat{f}_j \cos(\omega t_n).$$

with initial data

$$\hat{w}_{j}^{0} = \hat{v}_{j}, \qquad \hat{w}_{j}^{-1} = \hat{v}_{j} \left(1 - \frac{1}{2} \Delta t^{2} \lambda_{j}^{2} \right) - \frac{1}{2} \Delta t^{2} \hat{f}_{j}.$$

By (2.17)

$$|2 - \Delta t^2 \lambda_i^2| < 2$$

and the characteristic polynomial for the equation, $r^2 + (\Delta t^2 \lambda_j^2 - 2)r + 1$, then has two roots on the boundary of the unit circle. The solution is therefore stable and is given by (the verification of which is found in Appendix B)

(2.19)
$$\hat{w}_j^n = (\hat{v}_j - \hat{v}_j^\infty)\cos(\tilde{\lambda}_j t_n) + \hat{v}_j^\infty\cos(\omega t_n),$$

where λ_j is well-defined by the relation

$$2\frac{\sin(\Delta t \tilde{\lambda}_j/2)}{\Delta t} = \lambda_j.$$

Now, let

$$\Pi_h v = \sum_{j=1}^{\infty} \bar{v}_j \phi_j.$$

Then the numerical integration gives

$$\bar{v}_j = \frac{2\Delta t}{T} \sum_{n=0}^{M} \eta_n \left(\cos(\omega t_n) - \frac{1}{4} \right) \left((\hat{v}_j - \hat{v}_j^{\infty}) \cos(\tilde{\lambda}_j t_n) + \hat{v}_j^{\infty} \cos(\omega t_n) \right)$$
$$= (\hat{v}_j - \hat{v}_j^{\infty}) \beta_h(\tilde{\lambda}_j) + \hat{v}_j^{\infty} \beta_h(\omega) = \hat{v}_j \beta_h(\tilde{\lambda}_j) + (1 - \beta_h(\tilde{\lambda}_j)) \hat{v}_j^{\infty},$$

where

$$\beta_h(\lambda) = \frac{2\Delta t}{T} \sum_{n=0}^{M} \eta_n \cos(\lambda t_n) \left(\cos(\omega t_n) - \frac{1}{4} \right),$$

and we used the fact that the trapezoidal rule is exact, and equal to one, when $\lambda = \omega$. (Recall that for periodic functions the trapezoidal rule is exact for all pure trigonometric functions of order less than the number of grid points.) Hence, if $|\beta_h(\tilde{\lambda}_j)| < 1$ the j-th mode in the fixed point iteration converges to \hat{v}_j^{∞} . This is ensured by the following lemma, the proof of which is found in Appendix C.

Lemma 2.5. Under the assumptions of Theorem 2.4,

(2.20)
$$\max_{1 \le j \le N} |\beta_h(\tilde{\lambda}_j)| \le \rho_h =: \max(1 - 0.3\delta_h^2, 0.63).$$

Since the bound $|\beta_h(\tilde{\lambda}_j)| \le \rho_h < 1$ in the lemma is uniform for all j the convergence $v^{(k)} \to v^{\infty}$ with rate at least ρ_h follows. This concludes the proof of the theorem.

Remark 6. The discretization above is used as an example to illustrate the impact of going from the continuous to the discrete iteration. For a particular discretization we can improve the iteration further by using the knowledge of how it approximates ω and the eigenvalues of the continuous operator. Indeed, for the discretization above, let us define $\bar{\omega}$ by the relation

$$\omega = 2 \frac{\sin(\Delta t \bar{\omega}/2)}{\Delta t}.$$

Then if we use $f\cos(\bar{\omega}t_n)$ instead of $f\cos(\omega t_n)$ in the time stepping (2.15), the limit will be precisely the Helmholtz solution, $v^{\infty} = u$. Furthermore, the condition $\Delta t\omega \leq \min(\delta_h, 1)$ can be quite restrictive for problems close to resonance. It is only important to ensure convergence of the iterations. Another way to do that is to slightly change the discrete filter by replacing the constant 1/4 in (2.16) by a Δt -dependent number such that $|\beta_h(\lambda)| < 1$ for $\lambda \neq \omega$. Another option is to use a higher order quadrature rule, which would mitigate the restriction on Δt .

2.3. Tunable Filters. In Lemma 2.1 we saw that the filter transfer function satisfies $\beta(\omega) = 1$ and $-1/2 < \beta(r) < 1$ when $r \neq \omega$ and that these conditions guaranteed convergence of the WaveHoltz iteration. To improve convergence when $r \approx \omega$ we now consider a more general filter transfer function

(2.21)
$$\bar{\beta}(\lambda) = \frac{2}{T} \int_0^T (\cos(\omega t) + \alpha(t)) \cos(\lambda t) dt, \quad \alpha(t) = a_0 + \sum_{n=1}^\infty a_n \sin(n\omega t),$$

where we refer to $\alpha(t)$ as a time-dependent shift. As before, necessary conditions for convergence are $\bar{\beta}(\omega) = 1$, $\bar{\beta}'(\omega) = 0$. Straightforward calculations reveal that these conditions require that the two first coefficients must satisfy

$$a_1 = \frac{1}{2\pi}(1 + 4a_0).$$

The remaining terms in the sum are orthogonal to $\cos(\lambda t)$ when $\lambda = \omega$. Carrying out the integration in full for each term yields the general form

$$\bar{\beta}(\lambda) = \frac{\lambda \omega \sin(\lambda T)}{\pi (\lambda^2 - \omega^2)} + a_0 \frac{\omega \sin(\lambda T)}{\pi \lambda} + \sum_{n=1}^{\infty} a_n \frac{n\omega^2}{\pi (\lambda^2 - n^2 \omega^2)} \left(\cos(\lambda T) - 1\right),$$

from which it follows that another necessary condition is $|a_0| < 1/2$ since $|\bar{\beta}(r)| < 1$ and

$$\bar{\beta}(0) = a_0 \lim_{\lambda \to 0} \frac{\omega \sin(2\pi\lambda/\omega)}{\pi\lambda} = 2a_0.$$

We note that the standard filter, where $a_0 = -1/4$ and $a_1 = 0$, satisfies the necessary conditions.

REMARK 7. For the remaining coefficients a_n we only need to ensure that $|\bar{\beta}(r)| < 1$ which leaves large freedom to design $\bar{\beta}$. For example we may try to maximize $|\bar{\beta}''(\omega)|$ (minimize $\bar{\beta}''(\omega)$) so that $\bar{\beta}(r)$ is sharply peaked around $r = \omega$. We do not pursue a systematic study of this here but illustrate the utility of the added flexibility of (2.21) with numerical experiments below in Section 4.

2.4. Multiple Frequencies in One Solve. We can use the WaveHoltz algorithm to solve for multiple frequencies at once. Suppose we look for the solutions u_i of

$$\nabla \cdot (c^2(x)\nabla u_i) + \omega_i^2 u_i = f_i(x), \qquad i = 1, \dots, N,$$

with the same c and boundary condition for all i. To find those solutions we include all frequencies in the wave equation part of the iteration (2.3), and solve

(2.22)
$$w_{tt} = \nabla \cdot (c(x)^2 \nabla w) - \sum_{i=1}^{N} f_i(x) \cos(\omega_i t).$$

We then seek a decomposition

(2.23)
$$w(x,t) \equiv \sum_{i=1}^{N} u_i(x) \cos(\omega_i t),$$

of the solution. The filtering part of the WaveHoltz iteration is also updated to reflect the multiple frequencies

$$v_{n+1} = \frac{2}{T} \int_0^T \left(\sum_{i=1}^N \cos(\omega_i t) - \frac{1}{4} \right) w(x, t) dt.$$

As before we take $v_0 = 0$ when we deal with energy conserving boundary conditions. To this end we assume that the frequencies are related by an integer multiple in a way so that the period T can be chosen based on the lowest frequency.

The different $u_i(x)$ in (2.23) can be found as follows. Once we have found the time periodic solution to (2.22) evolve one more period and sample w(x,t) at N distinct times t_j , j = 1, ..., N. We then have

$$u_i(x) = \sum_{j=1}^{N} \beta_{ij} w(x, t_j),$$

where the coefficients β_{ij} are the elements of A^{-1} with the elements of A being $a_{ij} = \cos(\omega_j t_i)$.

- **3.** Wave Equation Solvers. In this section we briefly outline the numerical methods we use in the experimental section below. We consider both discontinuous Galerkin finite element solvers and finite difference solvers. In all the experiments we always use the trapezoidal rule to compute the integral in the WaveHoltz iteration.
- **3.1.** The Energy Based Discontinuous Galerkin Method. Our spatial discretization is a direct application of the formulation described for general second order wave equations in [3, 4]. Here we outline the spatial discretization for the special case of the scalar wave equation in one dimension and refer the reader to [3] for the general case.

The energy of the scalar wave equation is

$$H(t) = \int_D \frac{v^2}{2} + G(x, w_x) dx,$$

where

$$G(x, w_x) = \frac{c^2(x)w_x^2}{2},$$

is the potential energy density, v is the velocity (not to be confused with the iterates v^n above) or the time derivative of the displacement, $v = w_t$. The wave equation, written as a second order equation in space and first order in time then takes the form

$$w_t = v,$$

 $v_t = -\delta G - f(x)\cos(\omega t),$

where δG is the variational derivative of the potential energy

$$\delta G = -(G_{w_x})_x = -(c^2(x)w_x)_x.$$

For the continuous problem the change in energy is

(3.1)
$$\frac{dH(t)}{dt} = \int_{D} vv_{t} + w_{t}(c^{2}(x)w_{x})_{x} dx = -\int_{D} vf(x)\cos(\omega t)dx + [w_{t}(c^{2}(x)w_{x})]_{\partial D},$$

where the last equality follows from integration by parts together with the wave equation. Now, a variational formulation that mimics the above energy identity can be obtained if the equation $v-w_t=0$ is tested with the variational derivative of the potential energy. Let Ω_j be an element and $\Pi^s(\Omega_j)$ be the space of polynomials of degree s, then the variational formulation on that element is:

PROBLEM 1. Find $v^h \in \Pi^s(\Omega_j)$, $w^h \in \Pi^r(\Omega_j)$ such that for all $\psi \in \Pi^s(\Omega_j)$, $\phi \in \Pi^r(\Omega_j)$

(3.2)
$$\int_{\Omega_j} c^2 \phi_x \left(\frac{\partial w_x^h}{\partial t} - v_x^h \right) dx = \left[c^2 \phi_x \cdot n \left(v^* - v^h \right) \right]_{\partial \Omega_j},$$

(3.3)
$$\int_{\Omega_j} \psi \frac{\partial v^h}{\partial t} + c^2 \psi_x \cdot w_x^h + \psi f(x) \cos(\omega t) dx = [\psi (c^2 w_x)^*]_{\partial \Omega_j}.$$

Let $[[\zeta]]$ and $\{\zeta\}$ denote the jump and average of a quantity ζ at the interface between two elements, then, choosing the numerical fluxes as

$$v^* = \{v\} - \tau_1[[c^2 w_x]]$$
$$(c^2 w_x)^* = \{c^2 w_x\} - \tau_2[[v]],$$

will yields a contribution $-\tau_1([[c^2 w_x]])^2 - \tau_2([[v]])^2$ from each element face. To this end we choose $\tau_i > 0$ (so called upwind or Sommerfeld fluxes) which together with the choice that the approximation spaces be of the same degree r=s result in methods that are r+1 order accurate in space and measured in the L_2 norm. We note that even in the case of energy conserving numerical fluxes the formulation does not lead to a *symmetric* matrix for the WaveHoltz iteration (it is of course positive definite though).

Physical boundary conditions can also be handled by appropriate specification of the numerical fluxes, see [3] for details. The above variational formulation and choice of numerical fluxes results in an energy identity similar to (3.1). However, as the energy is invariant to certain transformations the variational problem does not fully determine the time derivatives of w^h on each element and independent equations must be introduced. In this case there is one invariant and an independent equation is $\int_{\Omega_j} \left(\frac{\partial w^h}{\partial t} - v^h \right) = 0$.

Denoting the degrees of freedom on element Ω_j by v_j and w_j the semi-discretization according to (3.2)-(3.3) on element Ω_j can be written

(3.4)
$$S(\frac{\partial w_j}{\partial t} - v_j) = L_1(v_{j-1}, v_j, v_{j+1}, w_{j-1}, w_j, w_{j+1}),$$

(3.5)
$$M\frac{\partial v_j}{\partial t} + Sw_j + f_j \cos(\omega t) = L_2(v_{j-1}, v_j, v_{j+1}, w_{j-1}, w_j, w_{j+1}),$$

where the elements of the element matrices M and S are $M_{kl} = \int_{\Omega_j} \phi_x \phi_l dx$ and $S_{kl} = \int_{\Omega_j} c^2(\phi_k)_x (\phi_l)_x dx$ respectively and the lift operators L_1 and L_2 represents the numerical fluxes. Note that a convenient way to directly enforce the independent equation is to compute the time derivatives of w_i according to

$$\frac{\partial w_j}{\partial t} = v_j = S^{\dagger} L_1(v_{j-1}, v_j, v_{j+1}, w_{j-1}, w_j, w_{j+1}),$$

where S^{\dagger} is the pseudo inverse of S.

3.2. Finite Difference Discretizations. For the finite difference examples we exclusively consider Cartesian domains $(x, y, z) \in [L_x, R_x] \times [L_y, R_y] \times [L_z, R_z]$ discretized by uniform grids $(x_i, y_j, z_k) = (L_x + L_y, R_y)$ $ih_x, L_y + jh_y, L_z + kh_z$, with $i = 0, \dots, n_x$ and $h_x = (R_x - L_x)/n_x$, etc.

When we have impedance boundary conditions on the form $w_t \pm \vec{n} \cdot \nabla w = 0$ we evolve the wave equation as a first order system in time according to the semi-discrete approximation

(3.6)
$$\frac{dv_{ijk}(t)}{dt} = (D_{+}^{x}D_{-}^{x} + D_{+}^{y}D_{-}^{y} + D_{+}^{z}D_{-}^{z})w_{ijk},$$

$$\frac{dw_{ijk}(t)}{dt} = v_{ijk},$$

$$\frac{dw_{ijk}(t)}{dt} = v_{ijk}$$

for all grid points that do not correspond to Dirichlet boundary conditions. On boundaries with impedance conditions we find the ghost point values by enforcing (here illustrated on the top of the domain)

$$(3.8) v_{ijn_z} - D_0^z w_{ijn_z} = 0.$$

Here we have used the standard forward, backward and centered finite difference operators, for example $h_x D_+^x w_{i,j,k} = w_{i+1,j,k} - w_{i,j,k}$ etc. For problems with variable coefficients the above discretization is generally alized as in [5].

We note that in some of the examples where we require high order accuracy we use the summation by parts discretization for variable coefficients developed by Mattson in [32] and described in detail there.

3.3. Time Discretization. In most of the numerical examples we use either an explicit second order accurate centered discretization of w_{tt} (for finite differences with energy conserving boundary conditions we eliminate v and time discretize w_{tt} directly as in the analysis in Section 2.2) or the classic fourth order accurate explicit Runge-Kutta method.

For some of the DG discretizations we employ Taylor series time-stepping in order to match the order of accuracy in space and time. Assuming that all the degrees of freedom have been assembled into a vector \mathbf{w} we can write the semi-discrete method as $\mathbf{w}_t = Q\mathbf{w}$ with Q being a matrix representing the spatial discretization. Assuming we know the discrete solution at the time t_n we can advance it to the next time step $t_{n+1} = t_n + \Delta t$ by the simple formula

$$\mathbf{w}(t_n + \Delta t) = \mathbf{w}(t_n) + \Delta t \mathbf{w}_t(t_n) + \frac{(\Delta t)^2}{2!} \mathbf{w}_{tt}(t_n) \dots$$
$$= \mathbf{w}(t_n) + \Delta t Q \mathbf{w}(t_n) + \frac{(\Delta t)^2}{2!} Q^2 \mathbf{w}(t_n) \dots$$

The stability domain of the Taylor series which truncates at time derivative number $N_{\rm T}$ includes part of the imaginary axis if $mod(N_T, 4) = 3$ or $mod(N_T, 4) = 0$ (see e.g. [28]). However as we use a slightly dissipative spatial discretization the spectrum of our discrete operator will be contained in the stability domain of all sufficiently large choices of $N_{\rm T}$ (i.e. the $N_{\rm T}$ should not be smaller than the spatial order of approximation). Note also that the stability domain grows linearly with the number of terms.

- 4. Numerical Examples. In this section we illustrate the properties of the proposed iteration and its Krylov accelerated version by a sequence of numerical experiments in one, two and three dimensions.
- **4.1. Examples in One Dimension.** We begin by presenting some very basic numerical experiments in one dimension.
- **4.1.1.** Convergence of Different Iterations / Solvers at a Fixed Frequency. We start by repeating the example described in Section 3.5 in [23]. This example is used in [23] to illustrate that the original cost functional from [9] (denoted J in [23]) does not yield the correct solution due to the existence of multiple minimizers.

The example solves the Helmholtz equation with c=1 and with the exact solution

$$u(x) = 16x^2(x-1)^2, \quad 0 \le x \le 1.$$

| Method / b. c. | WHI | LSQR | QMR | $\overline{\text{CG}}$ | GMRES |
|----------------|-----------|------------|------------|------------------------|------------|
| D-D | 94.5(-15) | 76.1(-15) | 75.9(-15) | 151.5(-15) | 97.9(-15) |
| N-N | 49.2(-15) | 142.8(-15) | 144.4(-15) | 158.5(-15) | 144.1(-15) |
| D-N | 28.3(-15) | 55.4(-15) | 81.9(-15) | 272.3(-15) | 67.0(-15) |

Table 4.1: Maximum error for various combinations of boundary conditions and methods.




Fig. 4.1: Convergence of the residual for the plain WaveHoltz iteration and its accelerated versions using LSQR, QMR, CG and GMRES. The titles of the figures indicate the boundary conditions used to the left and right, e.g. D-N means Dirichlet on the left and Neumann on the right.

Here both u (and w_t for the time-dependent problem) and u_x vanish at the endpoints so any boundary condition of the form

$$\alpha w_t + \beta(\vec{n} \cdot w_x) = 0, \quad \alpha^2 + \beta^2 = 1,$$

will be satisfied. Dirichlet boundary conditions correspond to $\alpha = 1$ and Neumann boundary conditions correspond to $\alpha = 0$, all other values will be an impedance boundary condition. Here, as in [23], we take the frequency to be $\omega = \pi/4$.

We discretize using the energy based DG method discussed above and use upwind fluxes which adds a small amount of dissipation. For this experiment we use 5 elements with degree q=7 polynomials and we use an 8th order accurate Taylor series method in time. We set Δt so that $n_t \Delta t = T = 2\pi/\omega$ while making the inequality $\Delta t \leq C_{\rm CFL} \Delta x/(q+1)$ as sharp as possible (in this experiment we fix $C_{\rm CFL} = 1/2$). With this resolution in space and time the truncation errors are negligible and we expect that the observed convergence properties should match those of the continuous analysis.

As mentioned above we expect that our method works best when combined with a classical iterative Krylov subspace method. The energy based DG method will produce a matrix A with real eigenvalues in (0,3/2) but it will not yield a symmetric matrix A. We present results for the WaveHoltz iteration (denoted WHI in figures and tables), and its acceleration with Matlab implementations of LSQR, QMR, CG and GMRES (we use the default unconditioned settings with a tolerance of 10^{-13}). In Figure 4.1 we display the convergence histories for various combinations of boundary conditions. The residuals for the Krylov accelerated iterations are the ones returned by the Matlab functions and the residual for the WaveHoltz iteration is simply the L_2 norm of the difference between two subsequent iterations. As can be seen the convergence behavior for QMR and GMRES are uniformly the fastest and appears to be insensitive to the type of boundary condition used. Note that the numerical method used here does not yield a symmetric matrix and CG is not guaranteed to work. Evidence of this loss or stagnation of convergence can be found in the cases D-D and D-N in Figure 4.1.

The actual errors in the converged solutions can be found in Table 4.1, where it can be seen that the error for all of the iteration methods are close to the residual tolerance.

4.1.2. Convergence with Increasing Frequency. To study how the number of iterations scale with the Helmholtz frequency ω we solve the wave equation on the domain $x \in [-6, 6]$ with constant wave speed




Fig. 4.2: Left: Number of iterations divided by ω as a function of ω for different boundary conditions. Middle and right: Zoom in around a resonance for the Dirichlet problem when using Krylov acceleration (middle) and when using WHI (right).

$$c^2(x) = 1$$
 and with a forcing

$$f(x) = \omega^2 e^{-(\omega x)^2},$$

that results in the solution being $\mathcal{O}(1)$ for all ω . The solver is the same as in the previous example. We keep the number of degrees of freedom per wave length fixed by letting the number of elements be $5\lceil\omega\rceil$. We always take the polynomial degree to be 7 and the number of Taylor series terms in the timestepping to be 8. As we now also consider impedance boundary conditions, with $\alpha = 1/2$, we use WHI accelerated by GMRES.

We report the number of iterations it takes to reach a GMRES residual smaller than 10^{-10} for the six possible combinations of Dirichlet, Neumann and impedance boundary conditions for 50 frequencies distributed evenly between 1 and 100. The results are displayed to the left in Figure 4.2 where we plot the number of iterations divided by ω as a function of ω . It is clear that the asymptotic scaling is linear with growing frequency. Interestingly all the combinations of boundary conditions collapse to two different curves with the Dirichlet-Dirichlet and impedance-impedance conditions converging the fastest.

We know from the analysis in Section 2.1 that the rate of convergence of the WaveHoltz iteration deteriorates near resonant frequencies (for non-impedance problems) but from Figure 4.2 it appears that all frequencies converge more or less the same rate. To study the behavior of the accelerated algorithm for homogenous Dirichlet boundary conditions we zoom in around $\omega \approx 19.114$ where the continuous problem has a resonance. In the middle graph in Figure 4.2 we display the required number of iterations around the resonant frequency. As can be seen there is some deterioration but only in very narrow band around a frequency that is slightly less than 19.114 and probably is the modified resonant frequency discussed in Section 2.2. This behavior can be contrasted to the growth of the number of iterations for the WaveHoltz iteration without GMRES acceleration, see the right figure in Figure 4.2. Clearly the acceleration of the WaveHoltz iteration by GMRES improves the robustness of the method near resonances.

4.1.3. Multiple Frequencies in One Solve. Here we illustrate the technique described in Section 2.4 for finding solutions of multiple frequencies at once. We set $\omega=1$ and $\omega_j=2^{j-1}\omega$ for $j=1,\ldots,4$, and consider the domain $x\in[0,1]$. We use the finite difference discretization discussed in Section 3.2 with Dirichlet boundary conditions. The time evolution is done by a second order centered discretization of w_{tt} , as was done in the discrete analysis in Section 2.2. The problem is forced by a point source centered at x=1/2 for $j=1,\ldots,4$, and we consider a constant wave speed $c^2(x)=1$. We display the convergence with decreasing h in Figure 4.3 on the right, where it can be seen that each solution u_j converges at a rate of h^2 .

4.1.4. Tunable Filters. Here we consider solving a Helmholtz problem in the domain $x \in [0, 1]$ with Dirichlet boundary conditions and constant wave speed $c^2 = 1$. The discretization is the same as in the previous experiment and we use a point source centered at x = 1/2. A straightforward calculation shows that the resonant frequencies of the problem are integer multiples of π and we specifically consider solving the Helmholtz problem with frequency $\omega = 4.1\pi$, which has a minimum relative gap to resonance of $\delta = 1$




Fig. 4.3: Left: Convergence history of the near resonant frequency 4.1π for the WaveHoltz filter and a tunable filter, and that of the frequency 1.5π for reference. Middle: The error between successive WaveHoltz iterates with the usual WaveHoltz filter. Right: Convergence of the solution for the CG accelerated WaveHoltz iteration with a point forcing.




Fig. 4.4: (Left) The usual WaveHoltz filter (in blue) and updated tunable filter (in red). (Right) Closeup of both the usual WaveHoltz filter and the updated tunable filter near the resonant frequency 4π .

 $1/41 \approx 0.024$. As discussed previously, we expect that the convergence rate of the WaveHoltz iteration will stagnate since ω is close to resonance. We compare the convergence against the problem with frequency $\omega = 1.5\pi$ which has a minimum relative gap to resonance of $\delta = 1/3$. The iteration history is displayed in Figure 4.3 on the left. It can be seen that the usual WaveHoltz iteration converges rapidly for the frequency 1.5π but that of 4.1π stagnates considerably. In the middle of Figure 4.3 we display the difference between successive WaveHoltz iterates for the Helmholtz problem with frequency ω , from which it is clear that the residual is a scaling of the resonant mode $\sin(4\pi x)$.

To improve the rate of convergence close to resonance we leverage a tunable filter as mentioned in Section 2.3. To obtain this filter, we consider the filter transfer function (2.21) and truncate the sin expansion of the time-dependent shift $\alpha(t)$ to 12 terms such that $a_n = 0$ for n > 11. In this example we take the usual choice of $a_0 = -1/4$ for the constant term in the filter transfer function (2.21) which, as discussed in Section 2.3, requires $a_1 = 0$. We then perform a minimization over a discrete set of 3000 equispaced points $r_j \in [0, 16\pi]$ of the empirically constructed functional

(4.1)
$$J(a_2, a_3, \dots, a_{11}) = 10.6 \bar{\beta}''(\omega) + 0.1 \sum_{|r_j - \omega| > 0.1} |\bar{\beta}(r_j)|^{20},$$

via 100 steepest descent iterations. The first term in the functional (4.1) minimizes the second derivative at the peak $\omega = 4.1\pi$, while the second weakly enforces that $|\beta(r)| \le 1$ for all r > 0 to ensure convergence of the fixed point iteration.

In Figure 4.4 on the right we see that the updated filter is steeper near $\omega = 4.1\pi$ so that repeated application of the updated filter will more quickly remove the resonant mode with frequency 4π and we




Fig. 4.5: Typical solutions computed with the GMRES accelerated WHI at $\omega = 77.5$. The thick lines indicate Dirichlet boundary conditions.

thus expect faster convergence. This is confirmed in the resulting iteration history of the updated filter, shown in Figure 4.3 on the left. The cost of improving convergence behavior near resonance, however, is a larger value of $\bar{\beta}$ for many other modes as shown in Figure 4.4 on the left. A more careful investigation of optimized filters is left for the future.

4.2. Problems in Two Dimensions. In this section we present experiments in two space dimensions.

4.2.1. Convergence in Different Geometries. In this example we solve the Helmholtz equation with a constant wave speed, $c^2 = 1$, in the domain $(x, y) \in [-1, 1]^2$ and with forcing

$$f(x,y) = -\omega^2 e^{-\sigma[(x-0.01)^2 + (y-0.015)^2]},$$

where $\sigma = \max(36, \omega^2)$. We vary the frequency according to $\omega = 1/2 + k$, k = 1, ..., 100, and keep the number of points per wavelength roughly constant by choosing $n_x = n_y = 8\lceil \omega \rceil$. Here we use the finite difference method outlined in Section 3.2 combined with the classic fourth order Runge-Kutta method in time with a timestep $\Delta t = h_x/c$.

For each frequency we solve six different problems consisting of combinations of Dirichlet and impedance boundary conditions with zero to four open sides and with the two open boundary case forking into two cases: (1) the open boundaries are opposite each other, or (2) next to each other forming a corner. In Figure 4.5 we display the real part of the solution for the frequency $\omega = 77.5$ for the six different problems.

In this example, the WaveHoltz iteration is accelerated by GMRES without restarts. Given that the storage requirement for GMRES grows with the number of iterations, it is often beneficial (especially for high frequency problems) to integrate and average over several periods to allow further propagation of information within the domain while mitigating the rapid growth of the Krylov subspace. For this example we thus choose to perform the WaveHoltz iteration with an integration time of 10 periods (i.e. we choose $T=10\frac{2\pi}{\omega}$). In Figure 4.6 we report the number of iterations needed to reduce the relative residual below 10^{-7} . It is clear from the results that the geometries where the waves can get trapped are considerably more difficult and requires more iterations. The computational results appear to indicate that the number of iterations to reach the tolerance scale as $\omega^{1.55}$ for the inner Dirichlet problem and similarly for the waveguide and the case with three Dirichlet boundary conditions. As the frequency increases and the distance between resonant




Fig. 4.6: To the left: number of iterations as a function of frequency to reduce the relative residual below 10^{-7} for problems with no trapped waves. Middle: the same but for problems with trapped waves and for the interior problem. Both are with the GMRES accelerated WHI . To the right: Residuals for the GMRES accelerated WHI , the CG accelerated WHI and for GMRES solution of the directly discretized Helmholtz problem.

frequencies decreases the iteration is not able to reduce the relative residual below the tolerance 10^{-7} within the prescribed maximum 1000 iterations. On the other hand for geometries with no trapped waves we see faster convergence (see the left figure in Figure 4.6) with the number of iterations scaling roughly as $\omega^{9/10}$.

To the right in Figure 4.6 we display the residual as a function of the number of right hand side evaluations (for the wave equation this is equivalent to taking a timestep and for the direct discretization of Helmholtz this is equivalent to one application of the sparse system matrix, the cost of these are roughly equivalent) when $\omega = 51.5$ for the pure Dirichlet boundary condition problem. The three different results are for: 1. the WaveHoltz acceleration, 2. the WaveHoltz iteration accelerated with conjugate gradient and based on the same spatial discretization but with a second order accurate centered discretization of w_{tt} using $\Delta t = 0.7h_x$ and, 3. a direct discretization of the Helmholtz equation (using the spatial discretization described in Section 3.2) combined with GMRES for solving the resulting system of equations. Precisely we use GMRES with restart every 100 iterations. For space reasons we only display this for one frequency but note that although the results may differ a bit between frequencies the trend is similar in the problems we have investigated.

It is clear from the residuals that both the GMRES and conjugate gradient accelerated WaveHoltz iterations are radically faster than applying GMRES to the direct discretization of Helmholtz. As all the methods use the same spatial discretization this is an indication of the importance of changing the problem from an indefinite system of equations to a positive definite and to a symmetric positive definite system.

Remark 8. We note that the problems considered in this experiment can be naturally solved with integral equation techniques since the the wave speed is constant. In addition as the problem is posed in two dimensions and can be stored in memory a good sparse solver will also be a very good alternative. What we want to demonstrate is: 1. The positive definiteness of the accelerated WHI makes it faster than standard iterative techniques for the direct discretization of Helmholtz, 2. The complexity is different for open and closed problems as predicted by the theory in [14].

4.2.2. Smoothly Varying Wave Speed in an Open Domain. In this example we consider a smoothly varying medium in a box $(x, y) \in [-1, 1]^2$. The wave speed is

$$c^{2}(x,y) = 1 - 0.4e^{-\left(\frac{x^{2}+y^{2}}{0.25^{2}}\right)^{4}},$$

and is also depicted in Figure 4.7.

Here we use the energy based DG solver and impose a right going plane wave $e^{i\omega(t-x)}$ through impedance boundary conditions on the left, bottom and top faces of the domain. On the right boundary we impose a zero Dirichlet condition. In all the computations we use degree 5 polynomials and a 6th order Taylor series method. The elements used form a Cartesian structured grid and we scale the number of elements so that we have 8 degrees of freedom per wavelength. The WHI is applied with an integration time of 5 periods and is accelerated by GMRES with a termination tolerance 10^{-7} on the relative residual. We solve the Helmholtz




Fig. 4.7: Left: the speed of sound (squared) used in example 4.2.2. Red indicates a rigid wall and black indicates open walls. Middle: Number of iterations as a function of frequency. Right: Compute time normalized by the frequency times the number of degrees of freedom.




Fig. 4.8: The magnitude of the Helmholtz solution for, from left to right, $\omega = 25\pi, 50\pi$ and 100π .

problem with $\omega = k\pi, k = 3, 4, \dots, 100$ and measure the total time from start to time of solution and we also measure the number of iterations needed to converge. The results, displayed in Figure 4.7, again show that for this type of open problem the iteration appears to require $N_{\text{iter}} \sim \mathcal{O}(\omega^{0.9})$ iterations to converge to a fixed tolerance. In terms of total computational time we observe $T_{\text{Total}} \sim \mathcal{O}(\omega N_{\text{DOF}})$ which is slightly higher than what would be expected from the $\mathcal{O}(\omega^{0.9})$ behavior.

However, as the distance traveled by the wave solution is proportional to $cT = 2\pi c/\omega$ and the information must travel through the domain at least once the time to solution is as good as can be expected. To reduce the computational complexity further we would need to propagate the solution faster than the speed of sound by applying a preconditioner or some type of multi-level strategy. Although we believe this is possible we leave such attempts to future work.

The magnitude of the solutions with $\omega = 25\pi, 50\pi$ and 100π are plotted in Figure 4.8.

4.2.3. Convergence of the Approximation Error and the Residual. As our iteration leads to a linear system of equations (and consequently a different residual) we should check that the residual is still a suitable proxy for the discretization error. Although we have no reason to believe this would not be the case we note that we have not yet performed a detailed analysis and resort to checking this numerically. We consider the same computational domain and method as above but with speed of sound c = 1 and with zero Dirichlet boundary conditions. We set $\omega = 2$ and choose the forcing so that the solution is

$$u = -(x^2 - 1)^2(y^2 - 1)^2,$$

and compute the solution using polynomials of degree three in the energy DG method and a fourth order




Fig. 4.9: The maximum error GMRES residuals as a function of number of iterations for four different mesh sizes. The rates of convergence agree with the order of the method.

accurate Taylor time stepper. In Figure 4.9 we display the maximum errors in u and the residuals for each GMRES iteration for Cartesian grids with grid spacings 1/2, 1/4, 1/8 and 1/16. As can be seen the residuals and the errors track well until the errors saturate. To the right in the figure we also indicate the rates of convergence based on the subsequent grid refinements. As expected they are very close to four.

4.2.4. The Marmousi2 Model. In the last two examples in this section we use the sixth order summation-by-parts finite difference operators developed by Mattson in [32]. Here we use the classic fourth order Runge-Kutta method for timestepping. In this example we simulate the solution caused by a point source placed in a material model where the speed of sound is taken from P-wave velocity in the Marmousi2 model¹. We discretize the full model which consists of 13601×2801 grid points and covers a domain that is roughly 17×3.5 kilometers. On the top surface we prescribe a zero Dirichlet condition and on the remaining three sides we add a 50 grid point wide supergrid layer (see [2]) that is terminated by zero Dirichlet boundary conditions. We accelerate the WHI by the transpose free quasi minimal residual (TFQMR) method and terminate the iteration when the relative residual is below 10^{-5} . We perform each iteration over 8 periods and take 500 timesteps per iteration. The time periodic point forcing is applied near the surface in grid point (6750, 2600) and we perform computations with $\omega = 200,400$ and 800.

As the number of unknowns is relatively large, $\sim 76 \cdot 10^6$, we parallelize the finite difference solver by a straightforward domain decomposition with the communication handled by MPI. The simulations were carried out on Maneframe II at the Center for Scientific Computation at Southern Methodist University using 60 dual Intel Xeon E5-2695v4 2.1 GHz 18-core Broadwell processors with 45 MB of cache each and 256 GB of DDR4-2400 memory. The results displayed in Figure 4.10 and 4.11 illustrate the ability of the method to find solutions to large problems and at high frequencies.

4.2.5. Multiple Frequencies. In this final example in two dimensions we again use the sixth order accurate summation-by-parts discretization from [32] with homogenous Dirichlet boundary conditions on the domain $(x, y) \in [-1, 1]^2$. The spatial discretization size is the same in both coordinates and is taken to be 2/300. The velocity model is taken to be smoothly varying. Precisely we have that

$$c^{2}(x,y) = 1 - 0.9 \left(e^{-\left(\frac{(x^{2} + (y - 0.4)^{2} - 0.4^{2})}{0.2^{2}}\right)^{4}} + e^{-\left(\frac{(x^{2} + (y + 0.4)^{2} - 0.3^{2})}{0.2^{2}}\right)^{4}} \right)$$

¹http://www.agl.uh.edu/downloads/downloads.htm




Fig. 4.10: Displayed is the base 10 logarithm of the magnitude of the Helmholtz solution ($\log_{10} |u|$) caused by a point source near the surface. The results are, from top to botton, for $\omega = 800, 400$ and 200.




Fig. 4.11: Zoom in of the base 10 logarithm of the magnitude of the Helmholtz solution ($\log_{10} |u|$) caused by a point source near the surface. The results are, from left to right, for $\omega = 800,400$ and 200.

see also Figure 4.12. We consider three frequencies, $\omega = 15$, 30, 60, and use the same forcing in Helmholtz for all frequencies,

$$f(x,y) = \frac{\sigma}{\pi} e^{-\sigma(x^2 + y^2)}, \quad \sigma = (4\omega)^2.$$

Here we use the WaveHoltz iteration over three periods of the lowest frequency, accelerated by GMRES (with tolerance 10^{-8}). We time step using a centered second order approximation to w_{tt} with a timestep $\Delta t = 1/600$. Since we solve for three frequencies at once we adjust the filter as described in Section 2.3 and extract all three solutions at once. Those solutions along with the material model are displayed in Figure




Fig. 4.12: Computation of three Helmholtz problems by one solve. The frequencies are $\omega = 15,30$ and 60. The material model is also displayed, red is $c^2 = 1$ and dark blue is $c^2 = 0.1$.




Fig. 4.13: To the left: number of iterations as a function of frequency to reduce the relative residual below $5 \cdot 10^{-5}$ for problems with no trapped waves. Here WHI is accelerated by TFQMR. To the right: the same but for the interior problem. Here WHI is accelerated with either CG or TFQMR.

4.12.

4.3. Problems in Three Dimensions. In this section we present experiments in three dimensions.

4.3.1. Convergence in Different Geometries. We solve the wave equation in a box $(x, y, z) \in [-1, 1]^3$ with the smoothly varying medium


$$c^{2}(x, y, z) = 1 + \frac{1}{10}e^{-(x^{2}+y^{2}+z^{2})}.$$

We use a uniform grid $(x_i, y_j, z_k) = (-1 + ih, -1 + jh, -1 + kh)$ with grid spacing h = 1/n and choose $n = \max(\lceil 10 \omega \rceil, 20)$ to keep the resolution fixed. The Helmholtz problem is forced by

(4.2)
$$F(x,y,z) = \omega^3 e^{-36\omega^2 ((x-x_0)^2 + (y-y_0)^2 + (z-z_0)^2)},$$

where $x_0 = 1/100$, $y_0 = 3/250$, and $z_0 = 1/200$. We impose a mixture of boundary conditions consisting of homogenous Dirichlet and/or impedance boundary conditions: (1) impedance on all sides, (2) Dirichlet at z = 1 and impedance on all other sides, (3) Dirichlet at z = -1, z = 1, and z = 1 with impedance on all other sides, and (4) Dirichlet on all sides. We solve the equations in first order form in time and use the semi-discrete approximation described in Section 3.2.

In this example the WaveHoltz iteration is performed over 5 periods, numerically integrated in time with the classic Runge Kutta method of order four, and accelerated by TFQMR method. For the pure Dirichlet problem we also use CG but note that although the spatial discretization leads to a symmetric WHI matrix when combined with a centered finite difference approximation in time the matrix is only close to symmetric when combined with the slightly dissipative Runge Kutta method. The experiments indicate that this slight non-symmetry does not destroy the convergence iteration of CG in this case.






Fig. 4.14: Displayed is the base 10 logarithm of the magnitude of the Helmholtz solution ($\log_{10}|u|$) caused by a point source near the surface for $\omega = 200$ (left) and $\omega = 300$ (right) at the slice x = 0.1.

In Figure 4.13 we report the number of iterations needed to reduce the relative residual below $5 \cdot 10^{-5}$. As was seen before in the 2D case, the fully Dirichlet case is notably more difficult and requires more iterations than the other problems considered. The computational results indicate that the number of iterations to reach the tolerance scale as $\omega^{2.57}$ for the inner Dirichlet problem with either CG or TFQMR, with the former taking fewer overall iterations than the latter. By comparison, the set of problems with boundary conditions (1)-(3) listed above appear to converge in a number of iterations that scales as $\omega^{0.97}$, i.e. close to linear in the frequency ω .

4.3.2. Scattering from a Plate. For our final example, we again consider the box $(x, y, z) \in [-1, 1]^3$ with smoothly varying medium

$$c^{2}(x, y, z) = \frac{1}{2} [3 + \sin(16\pi z)\sin(4\pi(x+y))],$$

and impose Dirichlet boundary conditions at z=1 and impedance boundary conditions on all other sides. We use a uniform grid $(x_i, y_j, z_k) = (-1 + ih, -1 + jh, -1 + kh)$ with grid spacing h=1/n where n is the number of gridpoints along a single dimension. The discretization in space and in time is exactly as in the previous example and the problem is forced by F(x, y, z) as in (4.2) with $x_0 = 1/100$, $y_0 = 3/250$, and $z_0 = 4/5$. As in the Marmousi example in the previous section, we parallelize the finite difference solver by a straightforward domain decomposition with the communication handled by MPI. This simulation was carried out on Maneframe II at the Center for Scientific Computation at Southern Methodist University using 64 dual Intel Xeon E5-2695v4 2.1 GHz 18-core Broadwell processors with 45 MB of cache each and 256 GB of DDR4-2400 memory. The magnitude of the solution with $\omega = 200$ and $\omega = 300$ is plotted in Figure 4.14. We use n = 1000 for a total of 10^9 gridpoints in the first case, and n = 1500 for a total of $3.375 \cdot 10^9$ gridpoints in the second for roughly 15-16 points per wavelength.

5. Summary and Future Work. We have presented and analyzed the WaveHoltz iteration, a new iterative method for solving the Helmholtz equation. The iteration results in positive definite and sometimes symmetric matrices that are more amenable for iterative solution by Krylov subspace methods. In choosing a Krylov subspace method we note that CG is the most efficient and memory lean choice when the resulting system is symmetric positive definite, otherwise GMRES generally outperforms other methods such as QMR, LSQR, and TFQMR. As the iteration is based on solving the wave equation it naturally parallelizes and can exploit techniques and spatial discretizations that have been developed for the time dependent problem. Numerical experiments indicate that our iteration appears to converge significantly faster than when the Helmholtz equation is discretized directly and solved iteratively with GMRES.

We believe that the numerical and theoretical results above are promising and note that there are many possible avenues for future exploration. For example we have exclusively used unconditioned Krylov solvers here but the spectral properties of the operator S indicate that preconditioning should be possible. Further,

we have not tried to exploit adaptivity in space or time or any type of sweeping ideas here and we have only briefly touched on the possibilities for more advanced filter design. We hope to study both the numerical and theoretical properties of these in the future.

Finally, here we only analyzed the energy conserving problem. We have carried out a preliminary but (we believe) suboptimal analysis of the impedance problem. We hope to improve our analysis of this case in the near future.

6. Acknowledgement. We would like to acknowledge the careful reading of the manuscript by the reviewers that led to several improvements and clarifications. In particular one of the reviews was extremely careful and thorough and we want to extend an extra thank you for this.

Appendix A. Proof of Lemma 2.1.

Proof. We show the results for the rescaled transfer function

$$\bar{\beta}(r) := \beta(r\omega) = \frac{2}{T} \int_0^T \left(\cos(\omega t) - \frac{1}{4} \right) \cos(r\omega t) dt = \frac{1}{\pi} \int_0^{2\pi} \left(\cos(t) - \frac{1}{4} \right) \cos(rt) dt.$$

By direct integration we get

(A.1)
$$\bar{\beta}(r) = \frac{1}{\pi} \int_0^{2\pi} \frac{1}{2} (\cos((r+1)t) + \cos((r-1)t)) - \frac{1}{4} \cos(rt) dt =$$

(A.2)
$$\frac{1}{2\pi} \left(\frac{\sin(2\pi(r+1))}{r+1} + \frac{\sin(2\pi(r-1))}{r-1} - \frac{1}{2} \frac{\sin(2\pi r)}{r} \right) = \operatorname{sinc}(r+1) + \operatorname{sinc}(r-1) - \frac{1}{2} \operatorname{sinc}(r),$$

where

$$\operatorname{sinc}(r) = \frac{\sin(2\pi r)}{2\pi r}.$$

We use the fact that $\sin(x) \leq x - \tilde{\alpha}x^3$ in the interval $x \in [0, \pi]$ for any $\tilde{\alpha} \in [0, \pi^{-2}]$. This leads to the following estimate for the sinc function

(A.3)
$$0 \le \operatorname{sinc}(r) \le 1 - \alpha r^2, \quad r \in [-0.5, 0.5], \quad \alpha \in [0, 4].$$

We also note that $\operatorname{sinc}(r+n) = \operatorname{sinc}(r)r/(r+n)$ for all integer n.

We now first consider $0 \le r \le 0.5$ and use (A.3) with $\alpha = 4$ and $\alpha = 0$,

$$|\bar{\beta}(r)| = \mathrm{sinc}(r) \left| \frac{r}{r+1} + \frac{r}{r-1} - \frac{1}{2} \right| = \frac{1}{2} \mathrm{sinc}(r) \frac{1+3r^2}{1-r^2} \leq \frac{1}{2} \frac{1-4r^2+3r^2}{1-r^2} = \frac{1}{2}.$$

For $0.5 \le r \le 1.5$ we instead center around r = 1 and get for $|\delta| \le 0.5$,

$$\bar{\beta}(1+\delta) = \text{sinc}(\delta) \frac{3(\delta+1)^2 + 1}{2(2+\delta)(1+\delta)} \ge 0,$$

since $\operatorname{sinc}(\delta) \geq 0$. Moreover, using again (A.3) with $\alpha = 1$,

$$\bar{\beta}(1+\delta) \le (1-\delta^2) \frac{3(\delta+1)^2+1}{2(2+\delta)(1+\delta)} = \frac{4+2\delta-3\delta^2-3\delta^3}{2(2+\delta)} \le \frac{4+2\delta-2\delta^2-\delta^3}{2(2+\delta)} = 1 - \frac{\delta^2}{2}.$$

Finally, for r > 1 we have $1/(r+1) - 1/2r \ge 0$ and therefore

$$|\bar{\beta}(r)| = \frac{|\sin(2\pi r)|}{2\pi r} \left(\frac{r}{r+1} + \frac{r}{r-1} - \frac{1}{2}\right) \le \frac{1}{2\pi} \left(\frac{1}{r+1} + \frac{1}{r-1} - \frac{1}{2r}\right) \le \frac{3}{4\pi} \frac{1}{(r-1)}.$$

We note also that for $r \geq 1.5$ this gives $|\bar{\beta}(r)| \leq 3/2\pi \leq 1/2$.

To prove (2.13) we use the Taylor expansion of $\bar{\beta}$ around r=1,

$$\bar{\beta}(1+\delta) = 1 + \frac{\delta^2}{2}\bar{\beta}''(1) + \frac{\delta^3}{6}\bar{R}(\delta),$$

where $\bar{R}(\delta)$ is the remainder term, which can be bounded as

$$|\bar{R}(\delta)| \le \sup_{r \ge 0} \left| \bar{\beta}^{(3)}(r) \right| \le \frac{1}{\pi} \int_0^{2\pi} t^3 \left(1 + \frac{1}{4} \right) dt = 5\pi^3.$$

Hence, $|R(\delta)| \leq |\bar{R}(\delta)|/6 \leq 5\pi^3/6$. Finally,

$$\beta''(1) = \operatorname{sinc}''(2) + \operatorname{sinc}''(0) - \frac{1}{2}\operatorname{sinc}''(1) = \frac{-1}{2} - \frac{(2\pi)^2}{3} + 1 = -2b_1.$$

This shows (2.13) and concludes the proof of the lemma.

Appendix B. Verification of Discrete Solution. Here we verify that (2.19) is indeed a solution to the difference equation (2.18). Direct substitution yields

$$\hat{w}_{j}^{n+1} - 2\hat{w}_{j}^{n} + \hat{w}_{j}^{n-1} + \Delta t^{2}\lambda_{j}^{2}\hat{w}_{j}^{n} = (\hat{v}_{j} - \hat{v}_{j}^{\infty})\cos(\tilde{\lambda}_{j}t_{n})\left(\cos(\tilde{\lambda}_{j}\Delta t) - 2 + \Delta t^{2}\lambda_{j}^{2} + \cos(\tilde{\lambda}_{j}\Delta t)\right) + \hat{v}_{j}^{\infty}\cos(\omega t_{n})\left(\cos(\omega \Delta t) - 2 + \Delta t^{2}\lambda_{j}^{2} + \cos(\omega \Delta t)\right) = (\hat{v}_{j} - \hat{v}_{j}^{\infty})\cos(\tilde{\lambda}_{j}t_{n})\left(-4\sin^{2}(\tilde{\lambda}_{j}\Delta t/2) + \Delta t^{2}\lambda_{j}^{2}\right) + \hat{v}_{j}^{\infty}\cos(\omega t_{n})\left(-4\sin^{2}(\omega \Delta t/2) + \Delta t^{2}\lambda_{j}^{2}\right) = \hat{v}_{j}^{\infty}\cos(\omega t_{n})\left(-\tilde{\omega}^{2} + \lambda_{j}^{2}\right)\Delta t^{2} = -\Delta t^{2}f_{j}\cos(\omega t_{n}).$$

Second, the initial conditions are satisfied since

$$\hat{w}_{j}^{0} = (\hat{v}_{j} - \hat{v}_{j}^{\infty}) + \hat{v}_{j}^{\infty} = \hat{v}_{j},$$

$$\hat{w}_{j}^{-1} = (\hat{v}_{j} - \hat{v}_{j}^{\infty})\cos(\tilde{\lambda}_{j}\Delta t) + \hat{v}_{j}^{\infty}\cos(\omega\Delta t) = (\hat{v}_{j} - \hat{v}_{j}^{\infty})\left(1 - \frac{1}{2}\Delta t^{2}\lambda_{j}^{2}\right) + \hat{v}_{j}^{\infty}\left(1 - \frac{1}{2}\Delta t^{2}\tilde{\omega}^{2}\right)$$

$$= \hat{v}_{j}\left(1 - \frac{1}{2}\Delta t^{2}\lambda_{j}^{2}\right) + \frac{1}{2}\Delta t^{2}\hat{v}_{j}^{\infty}\left(\lambda_{j}^{2} - \tilde{\omega}^{2}\right) = \hat{v}_{j}\left(1 - \frac{1}{2}\Delta t^{2}\lambda_{j}^{2}\right) - \frac{1}{2}\Delta t^{2}\hat{f}_{j}.$$

This shows that (2.19) solves (2.18).

Appendix C. Proof of Lemma 2.5.

Proof. In general, we introduce the trapezoidal rule applied to $\cos(\alpha t)$ in [0, 1],

$$\mathcal{T}_h(\alpha) := h \sum_{n=0}^{M} \eta_n \cos(\alpha t_n) \approx \int_0^1 \cos(\alpha t) dt = \frac{\sin(\alpha)}{\alpha}, \qquad h = 1/M,$$

from which we attain the following lemma:

LEMMA C.1. The error in $\mathcal{T}_h(\alpha)$ satisfies²

$$\left| \int_0^1 \cos(\alpha t) dt - \mathcal{T}_h(\alpha) \right| \le \frac{h^2 |\alpha|}{\pi^2}, \quad when \ |h\alpha| \le \pi.$$

Proof. A direct calculation shows that

$$\mathcal{T}_h(\alpha) = g(h\alpha) \int_0^1 \cos(\alpha t) dt, \qquad g(x) = \frac{x}{2\tan(x/2)}.$$

The function g(x) can be bounded as $1 - x^2/\pi^2 \le g(x) \le 1$ for $|x| \le \pi$. This gives

$$\left| \int_0^1 \cos(\alpha t) dt - \mathcal{T}_h(\alpha) \right| = |1 - g(h\alpha)| \left| \frac{\sin \alpha}{\alpha} \right| \le \frac{(h\alpha)^2}{\pi^2} \left| \frac{\sin \alpha}{\alpha} \right| \le \frac{h^2 |\alpha|}{\pi^2}.$$

²Note that this estimate is sharper than the standard error estimate for the trapezoidal rule, which would have the factor α^2 from the second derivative of the integrand, not just α .

Since

$$\left(\cos(\omega t) - \frac{1}{4}\right)\cos(\lambda t) = \frac{1}{2}\left(\cos((\omega + \lambda)t) + \cos((\omega - \lambda)t) - \frac{1}{2}\cos(\lambda t)\right),$$

and $h = \Delta t/T$, we can write

$$\beta_h(\lambda) = \frac{\Delta t}{T} \sum_{n=0}^{M} \eta_n \left(\cos((\omega + \lambda)t_n) + \cos((\omega - \lambda)t_n) - \frac{1}{2}\cos(\lambda t_n) \right)$$
$$= \left[\mathcal{T}_h(T(\omega + \lambda)) + \mathcal{T}_h(T(\omega - \lambda)) - \frac{1}{2}\mathcal{T}_h(T\lambda) \right].$$

From Lemma C.1 we then get that

$$|\beta(\tilde{\lambda}_j) - \beta_h(\tilde{\lambda}_j)| \le \frac{h^2}{\pi^2} \left(|T(\omega + \tilde{\lambda}_j)| + |T(\omega - \tilde{\lambda}_j)| + \frac{1}{2} |T\tilde{\lambda}_j| \right) \le \frac{5h^2T}{2\pi^2} (\omega + \tilde{\lambda}_j) = \frac{5\Delta t^2}{2\pi^2T} (\omega + \tilde{\lambda}_j),$$

when

$$\pi \ge hT(\omega + \tilde{\lambda}_j) = \Delta t(\omega + \tilde{\lambda}_j)$$

which is true by (2.17) and the fact that $\arcsin(x) \le \pi x/2$ for $x \in [0, 1]$:

(C.1)
$$\tilde{\lambda}_j = \frac{2}{\Delta t} \arcsin\left(\frac{\Delta t \lambda_j}{2}\right) \le \frac{\pi}{2} \lambda_j.$$

Next, we use the inequality $|x - \sin(x)| \le x^3/6$ for $|x| \le \pi/2$ to show that

(C.2)
$$\left| \frac{\sin(xh)}{h} - x \right| = \frac{1}{h} \left| \sin(xh) - xh \right| \le \frac{(xh)^3}{6h} = \frac{h^2 x^3}{6}, \qquad |hx| \le \frac{\pi}{2}.$$

It gives us an estimate for $\tilde{\lambda}_j - \lambda_j$,

$$|\lambda_j - \tilde{\lambda}_j| = \left| \frac{\sin\left(\Delta t \tilde{\lambda}_j/2\right)}{\Delta t/2} - \tilde{\lambda}_j \right| \le \frac{\Delta t^2}{24} \tilde{\lambda}_j^3,$$

which is valid for all j since $\Delta t \tilde{\lambda}_j/2 \leq \Delta t \pi \lambda_j/4 \leq \Delta t \pi \lambda_N/4 \leq \pi/2$, by (2.17) and (C.1).

By Lemma 2.1

$$|\beta(\omega + r)| \le \begin{cases} 1 - \frac{r^2}{2\omega^2}, & |r/\omega| \le \frac{1}{2}, \\ \frac{1}{2}, & |r/\omega| \ge \frac{1}{2}. \end{cases}$$

We now claim that the statement (2.20) in the lemma holds for all j if $\Delta t\omega \leq \min(\delta_h, 1)$. On the one hand, if $|\omega - \tilde{\lambda}_j| \geq \omega/2$, by (2.17) and (C.1)

$$|\beta_h(\tilde{\lambda}_j)| \leq |\beta(\tilde{\lambda}_j)| + \frac{5\Delta t^2}{2\pi^2 T}(\omega + \tilde{\lambda}_j) \leq \frac{1}{2} + \frac{5\Delta t^2}{2\pi^2 T}(\omega + \tilde{\lambda}_j) = \frac{1}{2} + \frac{5\Delta t\omega}{4\pi^3} \Delta t(\omega + \tilde{\lambda}_j) \leq \frac{1}{2} + \frac{5\Delta t\omega}{4\pi^2} \leq 0.63.$$

On the other hand, if $|\omega - \tilde{\lambda}_j| < \omega/2$,

$$\frac{|\omega - \tilde{\lambda}_j|}{\omega} \ge \frac{|\omega - \lambda_j|}{\omega} - \frac{|\tilde{\lambda}_j - \lambda_j|}{\omega} \ge \delta_h - \frac{\Delta t^2}{24\omega} \tilde{\lambda}_j^3 \ge \delta_h - \frac{\Delta t^2}{24\omega} (\omega + |\tilde{\lambda}_j - \omega|)^3 \ge \delta_h - \frac{(3/2)^3}{24} \Delta t^2 \omega^2 \ge \frac{55}{64} \delta_h,$$

since $\min(\delta_h, 1)^2 \leq \delta_h$. Then

$$|\beta_h(\tilde{\lambda}_j)| \le |\beta(\tilde{\lambda}_j)| + \frac{5\Delta t^2}{48T}(\omega + \tilde{\lambda}_j) \le 1 - \frac{1}{2} \left(\frac{|\omega - \tilde{\lambda}_j|}{\omega}\right)^2 + \frac{5\Delta t^2 \omega}{96\pi}(\omega + \omega + \omega/2)$$

$$\le 1 - \frac{55^2}{2 \cdot 64^2} \delta_h^2 + \Delta t^2 \omega^2 \frac{25}{192\pi} \le 1 - \left(\frac{55^2}{2 \cdot 64^2} - \frac{25}{192\pi}\right) \delta_h^2 \le 1 - 0.3\delta_h^2.$$

This proves the lemma.

REFERENCES

- [1] D. Appelö, Absorbing Layers and Non-Reflecting Boundary Conditions for Wave Propagation Problems, PhD thesis, Royal Institute of Technology, October 2005.
- [2] D. Appelö and T. Colonius, A high-order super-grid-scale absorbing layer and its application to linear hyperbolic systems, Journal of Computational Physics, 228 (2009), pp. 4200–4217.
- [3] D. Appelö and T. Hagstrom, A new discontinuous Galerkin formulation for wave equations in second order form, SIAM Journal On Numerical Analysis, 53 (2015), pp. 2705–2726.
- [4] D. Appelö and T. Hagstrom, An energy-based discontinuous Galerkin discretization of the elastic wave equation in second order form, Comput. Meth. Appl. Mech. Engrg., 338 (2018), pp. 362–391.
- [5] D. Appelö and N. A. Petersson, A stable finite difference method for the elastic wave equation on complex geometries with free surfaces, Communications in Computational Physics, 5 (2009), pp. 84–107.
- [6] HADRIEN BÉRIOT, ALBERT PRINN, AND GWÉNAËL GABARD, Efficient implementation of high-order finite elements for helmholtz problems, International Journal for Numerical Methods in Engineering, 106 (2016), pp. 213–240.
- [7] A. Björck, Numerical methods in matrix computations, vol. 59, Springer, 2015.
- [8] A. Brandt and I. Livshits, Wave-ray multigrid method for standing wave equations, Electron. Trans. Numer. Anal, 6 (1997), p. 91.
- [9] M.O. Bristeau, R. Glowinski, and J. Périaux, Controllability methods for the computation of time-periodic solutions; application to scattering, Journal of Computational Physics, 147 (1998), pp. 265–292.
- [10] Z. Chen and X. Xiang, A source transfer domain decomposition method for Helmholtz equations in unbounded domain, SIAM Journal on Numerical Analysis, 51 (2013), pp. 2331–2356.
- [11] B. Engquist and A. Majda, Absorbing boundary conditions for the numerical simulation of waves, Math. Comp., 31 (1977), p. 629.
- [12] B. Engquist and L. Ying, Sweeping preconditioner for the Helmholtz equation: hierarchical matrix representation, Communications on pure and applied mathematics, 64 (2011), pp. 697-735.
- [13] ——, Sweeping preconditioner for the Helmholtz equation: moving perfectly matched layers, Multiscale Modeling & Simulation, 9 (2011), pp. 686–710.
- [14] B. Engquist and H. Zhao, Approximate separability of the Green's function of the Helmholtz equation in the high frequency limit, Communications on Pure and Applied Mathematics, 71 (2018), pp. 2220–2274.
- [15] Y.A. Erlangga, Advances in iterative methods and preconditioners for the Helmholtz equation, Archives of Computational Methods in Engineering, 15 (2008), pp. 37–66.
- [16] Y. Erlangga, C. Oosterlee, and C. Vuik, A novel multigrid based preconditioner for heterogeneous Helmholtz problems, SIAM Journal on Scientific Computing, 27 (2006), pp. 1471–1492.
- [17] O.G. Ernst and M.J. Gander, Why it is difficult to solve Helmholtz problems with classical iterative methods, in Numerical analysis of multiscale problems, Springer, 2012, pp. 325–363.
- [18] M. GANDER AND F. NATAF, AILU for Helmholtz problems: a new preconditioner based on an analytic factorization, Comptes Rendus de l'Académie des Sciences - Series I - Mathematics, 331 (2000), pp. 261–266.
- [19] M. Gander and H. Zhang, A class of iterative solvers for the Helmholtz equation: Factorizations, sweeping preconditioners, source transfer, single layer potentials, polarized traces, and optimized Schwarz methods, SIAM Review, 61 (2019), pp. 3–76.
- [20] A. GILLMAN, A.H. BARNETT, AND P-G. MARTINSSON, A spectrally accurate direct solution technique for frequencydomain scattering problems with variable media, BIT Numerical Mathematics, 55 (2015), pp. 141–170.
- [21] R. GLOWINSKI AND T. ROSSI, A mixed formulation and exact controllability approach for the computation of the periodic solutions of the scalar wave equation. (i): Controllability problem formulation and related iterative solution, Comptes Rendus Math., 343 (2006), pp. 493–498.
- [22] M.J. Grote and J.H. Tang, On controllability methods for the Helmholtz equation, Journal of Computational and Applied Mathematics, 358 (2019), pp. 306–326.
- [23] M. J. GROTE, F. NATAF, J. H. TANG, AND P.-H. TOURNIER, Parallel Controllability Methods For the Helmholtz Equation, arXiv e-prints, (2019), p. arXiv:1903.12522.
- [24] T. HAGSTROM, Radiation boundary conditions for the numerical simulation of waves, Acta Numerica, 8 (1999), pp. 47– 106.
- [25] E. Heikkola, S. Mönkölä, A. Pennanen, and T. Rossi, Controllability method for acoustic scattering with spectral elements, Journal of Computational and Applied Mathematics, 204 (2007), pp. 344–355.
- [26] ——, Controllability method for the Helmholtz equation with higher-order discretizations, Journal of Computational Physics, 225 (2007), pp. 1553–1576.
- [27] C. V. HILE AND G. A KRIEGSMANN, A hybrid numerical method for loaded highly resonant single mode cavities, Journal of Computational Physics, 142 (1998), pp. 506–520.
- [28] D. KETCHESON, L. LÓCZI, AND T. KOCSIS, On the absolute stability regions corresponding to partial sums of the exponential function, IMA Journal of Numerical Analysis, 35 (2015), pp. 1426-1455.
- [29] H.-O. Kreiss and J. Oliger, Comparison of accurate methods for the integration of hyperbolic equations, Tellus, 24 (1972), pp. 199-215.
- [30] O. A. Ladyzhenskaya, On the limiting-amplitude principle, Uspekhi Mat. Nauk, 12 (1957), pp. 161–164.
- [31] I. Livshits and A. Brandt, Accuracy properties of the Wave-ray multigrid algorithm for Helmholtz equations, SIAM Journal on Scientific Computing, 28 (2006), pp. 1228–24.
- [32] K. Mattsson, Summation by parts operators for finite difference approximations of second-derivatives with variable coefficients, Journal of Scientific Computing, 51 (2012), pp. 650-682.

- [33] C.S. Morawetz, The limiting amplitude principle, Communications on Pure and Applied Mathematics, 15 (1962), pp. 349–361.
- [34] C.C. Stolk, A rapidly converging domain decomposition method for the Helmholtz equation, Journal of Computational Physics, 241 (2013), pp. 240–252.
- [35] B. R. VAINBERG, On short-wave asymptotic behaviour of solutions to steady-state problems and the asymptotic behaviour as t → ∞ of solutions of time-dependent problems, Uspekhi Mat. Nauk, 30 (1975), pp. 1–58.
- [36] A. Vion and C. Geuzaine, Double sweep preconditioner for optimized Schwarz methods applied to the Helmholtz problem, Journal of Computational Physics, 266 (2014), pp. 171–190.
- [37] S. Wang, M.V. de Hoop, and J. Xia, On 3d modeling of seismic wave propagation via a structured parallel multifrontal direct Helmholtz solver, Geophysical Prospecting, 59 (2011), pp. 857–873.
- [38] H. Weyl., Über die asymptotische verteilung der eigenwerte, Nachr. Konigl. Ges. Wiss., (1911), pp. 110-117.
- [39] L. Zepeda-Núñez, A. Scheuer, R. J. Hewett, and L. Demanet, The method of polarized traces for the 3D Helmholtz equation, ArXiv e-prints, (2018).
- [40] L. Zepeda-Núñez and L. Demanet, The method of polarized traces for the 2d Helmholtz equation, Journal of Computational Physics, 308 (2016), pp. 347–388.
- [41] ——, Nested domain decomposition with polarized traces for the 2d Helmholtz equation, SIAM Journal on Scientific Computing, 40 (2018), pp. B942–B981.