A convergent Born series for solving the inhomogeneous Helmholtz equation in arbitrarily large media

Gerwin Osnabrugge, Saroch Leedumrongwatthanakun, Ivo M. Vellekoop*

Biomedical Photonic Imaging Group, MIRA Institute for Biomedical Technology & Technical Medicine, University of Twente, PO Box 217, 7500 AE Enschede, The Netherlands

A R T I C L E   I N F O

A B S T R A C T

We present a fast method for numerically solving the inhomogeneous Helmholtz equation. Our iterative method is based on the Born series, which we modified to achieve convergence for scattering media of arbitrary size and scattering strength. Compared to pseudospectral time-domain simulations, our modified Born approach is two orders of magnitude faster and nine orders of magnitude more accurate in benchmark tests in 1, 2, and 3-dimensional systems.

© 2016 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

1. Introduction

The Helmholtz equation, the time-independent form of the scalar wave equation, appears in many fields of physics ranging from electron scattering to seismology. Additionally, this equation describes electromagnetic wave propagation in 2-dimensional systems, and it is often used as a scalar approximation to light propagation in 3-dimensional scattering media [1].

A great variety of approaches are available for solving the Helmholtz equation in inhomogeneous media [2]. For most numerical methods, the equation is discretized using the finite difference [3] or the finite element [4,5] approximation. The corresponding system of linear equations can be solved using direct matrix inversion methods [6,7], but these methods are ineffective for solving large inhomogeneous systems. More efficient methods exist, such as the category of Krylov subspace methods [8] extended with multigrid [9,10] or domain decomposition methods [11,12]. However, solving the inhomogeneous Helmholtz equation for large systems remains a computationally challenging task [13].

For the simulation of light propagation, the finite-difference time-domain methods (FDTD) [14,15] and the pseudospectral time domain methods (PSTD) [16–18] are favored the most [19]. Using these time-domain methods, the steady-state solution is found by replacing the source term with a periodically oscillating source and running the simulations until steady state is reached. However, the finite difference approximations required for these methods are only exact in the limit of an infinitely small step size, resulting in a trade-off between accuracy and speed.

A completely different class of methods is formed by volume integral methods utilizing the Green's function [20–23]. From Green's function theorem, the well-known Born series can be derived [1]. The Born series has long been used to solve

---

* Corresponding author.
  E-mail address: i.m.vellekoop@utwente.nl (I.M. Vellekoop).

the scattering problem for small particles with a low scattering potential. However, the series diverges for larger structures, or structures with a high scattering contrast [24].

Here we present a modified version of the Born series, that unconditionally converges to the exact solution of the Helmholtz equation for arbitrarily large structures with an arbitrarily high scattering potential. Our method converges rapidly and monotonically to the exact solution, and is several orders of magnitude faster than PSTD.

First, we will introduce our modified Born series, and discuss the convergence criteria and other characteristics. Then, we compare the accuracy and run-time of our method to that of PSTD for 1, 2, and 3-dimensional homogeneous and inhomogeneous media. A proof of convergence is given in Appendix A and the boundary conditions for our experiments are presented in Appendix B.

2. Traditional Born series

First, the inhomogeneous Helmholtz equation is considered

$$\nabla^2 \psi(\mathbf{r}) + k(\mathbf{r})^2 \psi(\mathbf{r}) = -S(\mathbf{r}), \tag{1}$$

where $\psi(\mathbf{r})$ represents the field at position $\mathbf{r} \in \mathbb{R}^D$, with D representing the dimensionality of the problem. $S(\mathbf{r})$ is the source term and $k(\mathbf{r})$ is the wavenumber. The scattering potential is defined as $V(\mathbf{r}) \equiv k(\mathbf{r})^2 - k_0^2 - i\epsilon$ to find

$$\nabla^2 \psi(\mathbf{r}) + (k_0^2 + i\epsilon)\psi(\mathbf{r}) = -V(\mathbf{r})\psi(\mathbf{r}) - S(\mathbf{r}). \tag{2}$$

Typically, one chooses k_0 as the background potential and ϵ infinitesimally small [20]. However, the solution of Eq. (2) does not depend on the choice for k_0 and ϵ at all. Using the Green's function theorem, the solution can be written as

$$\psi(\mathbf{r}) = \int g_0(\mathbf{r} - \mathbf{r}')[V(\mathbf{r}')\psi(\mathbf{r}') + S(\mathbf{r}')]d\mathbf{r}', \tag{3}$$

where the Green's function g_0 is defined as the solution to

$$\nabla^2 g_0(\mathbf{r}) + (k_0^2 + i\epsilon)g_0(\mathbf{r}) = -\delta(\mathbf{r}). \tag{4}$$

The Green's function can easily be found in Fourier transformed coordinates $\mathbf{p}$, where $\tilde{g}_0(\mathbf{p}) = 1/(|\mathbf{p}|^2 - k_0^2 - i\epsilon)$. Eq. (3) can now be simplified to the matrix form

$$\psi = GV\psi + GS, \tag{5}$$

where the convolution with g_0 is replaced by a single operator $G \equiv F^{-1}\tilde{g}_0(\mathbf{p})F$, with F and F^{-1} the forward and inverse Fourier transforms, respectively. V is a diagonal matrix, and S and ψ are vectors. Eq. (5) can be expanded recursively to arrive at the traditional Born series

$$\psi_{\text{Born}} = [1 + GV + GVGV + \ldots]GS. \tag{6}$$

Note that this series only converges to ψ if the spectral radius ρ of operator GV is less than unity. The Born series has proven to be successful for solving the Helmholtz equation for small scattering structures with weak scattering potential [25]. However, it is well known that there are serious constraints on the convergence of the series for larger and stronger scattering structures, limiting the usefulness of this approach [24].

3. Modified Born series

3.1. Convergent Born series

Here we show that the series *can* be made convergent for any arbitrary scattering potential by choosing a suitable preconditioner γ. Applying this preconditioner to Eq. (5) gives

$$\gamma\psi = \gamma GV\psi + \gamma GS, \tag{7}$$

$$\psi = M\psi + \gamma GS, \tag{8}$$

with

$$M \equiv \gamma GV - \gamma + 1. \tag{9}$$

We now choose

$$\gamma(\mathbf{r}) = \frac{i}{\epsilon}V(\mathbf{r}), \tag{10}$$

$$\epsilon \geq \max_{\mathbf{r}}|k(\mathbf{r})^2 - k_0^2|. \tag{11}$$

As we show in Appendix A, this combination of γ and ϵ ensures that $\rho(M) < 1$, meaning that Eq. (8) can be expanded recursively to arrive at our modified Born series

$$\psi = [1 + M + M^2 + M^3 + \ldots]\gamma G S, \tag{12}$$

or, in iterative form

$$\psi_{k+1} = M\psi_k + \gamma G S, \tag{13}$$

with $\psi_0 = \gamma G S$. This iteration converges to the exact solution of the Helmholtz equation for arbitrarily large media.

3.2. Interpretation

By introducing ϵ and the preconditioner γ, we are able to ensure the convergence of the Born series. In order to gain an intuitive understanding of our algorithm, we first analyze the role of ϵ. As can be seen from Eq. (4), ϵ introduces an imaginary component to the wavevector of the background medium. As a result, the Green's function in a homogeneous medium decays exponentially with distance. For instance, in three dimensions the solution to Eq. (4) is

$$g_0(\mathbf{r}) = \frac{e^{i|\mathbf{r}|\sqrt{k_0^2 + i\epsilon}}}{4\pi|\mathbf{r}|}. \tag{14}$$

The exponential decay makes that the total amount of energy represented by the Green's function is finite and localized.

The term $i\epsilon$ in the background potential is compensated exactly by an imaginary term in the scattering potential V, such that the final solution of the Helmholtz equation remains the same. A physical interpretation of this construct is that the background medium is lossy, and the scattering potential compensates this loss by an equal amount of gain. As a result, even a homogeneous medium has a non-zero scattering potential.

Figs. 1a–c show the solution for a homogeneous medium after 20, 40, and 60 iterations respectively. It can be seen that the solution expands with each iteration. This expansion should not be confused with time-resolved propagation, since the algorithm is only solving a time-independent wave equation. Still, it is useful to analyze this 'pseudo-propagation' effect since it gives an indication of how many iterations are required to cover a homogeneous medium of a given size. We define the pseudo-propagation distance as the distance traveled by a 1-dimensional wave packet $\Psi_0(x) \equiv A(x)\exp(ik_0 x)$, with x as the spatial dimension and A as a positive-valued envelope function. The center of the wave packet is defined as

$$\langle x_0 \rangle \equiv \frac{\int x A(x)dx}{\int A(x)dx} = \frac{\int x \Psi_0(x)e^{-ik_0 x}dx}{\int \Psi_0(x)e^{-ik_0 x}dx} = \left.\frac{d\tilde{\Psi}_0(p_x)}{\tilde{\Psi}_0(p_x)idp_x}\right|_{p_x=k_0} \tag{15}$$

with $\tilde{\Psi}_0$ denoting the Fourier transform of Ψ_0. After a single iteration the packet has transformed into

$$\Psi_1(x) = [\gamma G V - \gamma + 1]\Psi_0(x). \tag{16}$$

For a homogeneous medium ($V = -i\epsilon$, and $\gamma = iV/\epsilon = 1$)

$$\tilde{\Psi}_1(p_x) = \left[\frac{-i\epsilon}{p_x^2 - k_0^2 - i\epsilon}\right]\tilde{\Psi}_0(p_x) \tag{17}$$

The center of the shifted wave packet Ψ_1 is at

$$\langle x_1 \rangle = \left.\frac{d\tilde{\Psi}_1(p_x)}{\tilde{\Psi}_1(p_x)idp_x}\right|_{p_x=k_0} = 2k_0/\epsilon + \langle x_0 \rangle, \tag{18}$$

with $\langle x_0 \rangle$ the center of the original wave packet. In words, each iteration moves the wave packet over a distance $2k_0/\epsilon$. This behavior is illustrated in Fig. 1d. The plot shows the differential contributions of the 20th, 40th, and 60th iteration. Clearly, each iteration only modifies the field in a finite volume, and the solution expands at a speed of $2k_0/\epsilon$ per iteration. Eq. (18) shows that ϵ should be chosen as low as possible to ensure rapid convergence. Since the minimum value of ϵ is given by Eq. (11), our method will be faster for structures with a lower contrast. Therefore, we let $k_0^2 = \left(\min_{\mathbf{r}} \text{Re}\{k^2(\mathbf{r})\} + \max_{\mathbf{r}} \text{Re}\{k^2(\mathbf{r})\}\right)/2$ in order to minimize ϵ.

The pseudo-propagation speed can be compared to the actual propagation speed in other methods such as PSTD and FDTD. Table 1 summarizes the key characteristics of our method compared to PSTD and FDTD. Because of the required finite difference approximations, the accuracy of FDTD is proportional to both Δt^2 and Δx^2, and the accuracy of PSTD is proportional to Δt^2, whereas the modified Born method is only restricted by machine precision. By comparing the (pseudo-)propagation speed, it can be seen that if $2k_0/\epsilon > c\Delta t$ or

$$\left(\frac{k_m}{k_0} - \frac{k_0}{k_m}\right) < \sqrt{D}, \tag{19}$$
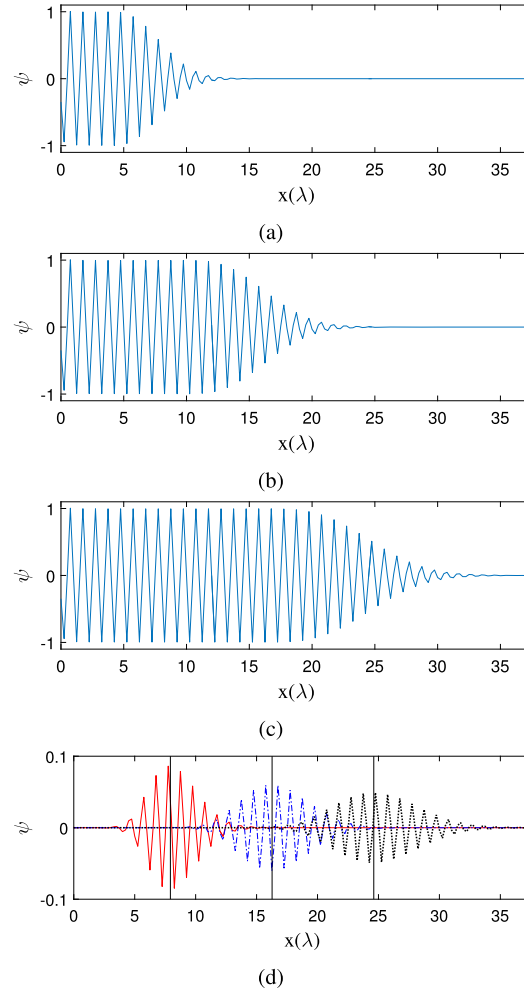
Fig. 1. Solution for wave propagation in a homogeneous medium after a) 20, b) 40, and c) 60 iterations respectively. d) The field added during iteration 20 (solid red), iteration 40 (dashed blue), and iteration 60 (dotted black). The vertical lines show the theoretical positions of the differential contributions, calculated using Eq. (18). In all simulations $\epsilon = 0.8k_0^2$. (Color version online.)

Table 1
Key parameters for our modified Born approach, the PSTD and FDTD numerical methods. Here the largest wavenumber $k_m \equiv \max_{\mathbf{r}} |k(\mathbf{r})|$, the shortest wavelength $\lambda_m \equiv 2\pi/k_m$ and L is the length of the system. The scattering contrast $\nu \equiv \left(\frac{k_m}{k_0} - \frac{k_0}{k_m}\right)$.

	Symbol	Modified Born	PSTD [16]	FDTD [14,15]
Grid spacing	Δx	$\leq \lambda_m/2$	$\leq \lambda_m/2$	$\leq \lambda_m/8$
Accuracy		machine precision	$O(\Delta t^2)$	$O(\Delta t^2 \Delta x^2)$
Total grid points	N_{tot}	$\geq (k_m L/\pi)^D$	$\geq (k_m L/\pi)^D$	$\geq (4k_m L/\pi)^D$
FLOPs/iterations	F	$O(N_{tot} \log N_{tot})$	$O(N_{tot} \log N_{tot})$	$O(N_{tot})$
		$O(k_m^D \log k_m^D)$	$O(k_m^D \log k_m^D)$	$O(k_m^D)$
Time step	Δt	n/a	$\leq \frac{2\Delta x}{\pi c\sqrt{D}}$	$\leq \frac{\Delta x}{c\sqrt{D}}$
(Pseudo-)propagation	Δr	$2k_0/\epsilon \leq \frac{2}{k_m \nu}$	$c\Delta t \leq \frac{2}{k_m \sqrt{D}}$	$c\Delta t \leq \frac{\pi}{4k_m \sqrt{D}}$
Total FLOPs		$O(\nu k_m^{D+1} \log k_m^D)$	$O(\sqrt{D} k_m^{D+1} \log k_m^D)$	$O(\sqrt{D} k_m^{D+1})$

our method is always faster than PSTD. For $D = 3$, this condition corresponds with $k_m/k_0 < 2.2$. For example, if the scalar wave approximation is used to simulate light propagation in biological tissue, with a refractive index between 1.33 and 1.45, our exact method 'propagates' at a speed of 2.7 wavelengths per iteration, which is a factor 21 faster than a PSTD simulation ran at the lowest accuracy possible.

4. Implementation

We implemented the modified Born algorithm in Matlab. The field ψ, and potential map V are discretized on a regular 2-dimensional grid with a grid spacing of $\Delta x = \lambda/4$, with λ an arbitrarily chosen wavelength of 1 distance unit. The field ψ, the potential V, and the pre-computed $\tilde{g}_0$ are stored in 2-dimensional arrays. Our iterative algorithm is implemented as

$$\psi_{k+1}(\mathbf{r}) = \psi_k(\mathbf{r}) - \frac{i}{\epsilon} V(\mathbf{r}) \left(\psi_k(\mathbf{r}) - \text{ifft}_2 \left[\tilde{g}_0(\mathbf{p}) \, \text{fft}_2 \left[V(\mathbf{r})\psi_k(\mathbf{r}) + S(\mathbf{r}) \right] \right] \right), \tag{20}$$

where fft$_2$ and ifft$_2$ are the forward and inverse 2-dimensional fast Fourier transform operators, index k is the iteration number, and S is the source. All multiplications are point-wise multiplications.

In order to compare the accuracy and efficiency of our approach with that of PSTD, we used a PSTD algorithm to solve the time-dependent scalar wave equation for a time-harmonic source [26]

$$\nabla^2 \psi(\mathbf{r}, t) - \frac{1}{c^2(\mathbf{r})} \frac{\partial^2 \psi(\mathbf{r}, t)}{\partial t^2} - \frac{\sigma(\mathbf{r})}{c^2(\mathbf{r})} \frac{\partial \psi(\mathbf{r}, t)}{\partial t} = -S(\mathbf{r})e^{-i\omega t}, \tag{21}$$

with ω the angular frequency of the source, $c(\mathbf{r})$ the wave velocity, and $\sigma(\mathbf{r})$ the damping rate. We choose the time units to have unit wave velocity, such that $\omega = 2\pi/\lambda$. By Fourier transforming Eq. (21) with respect to time, we find back the Helmholtz equation

$$\nabla^2 \psi(\mathbf{r}) + \frac{\omega^2}{c^2(\mathbf{r})} \psi(\mathbf{r}) + \frac{i\omega\sigma(\mathbf{r})}{c^2(\mathbf{r})} \psi(\mathbf{r}) = -S(\mathbf{r}). \tag{22}$$

This result shows that we can solve the Helmholtz equation using PSTD by choosing $c^2(\mathbf{r}) = \omega^2/\text{Re}\{k^2(\mathbf{r})\}$ and $\sigma(\mathbf{r}) = \omega\text{Im}\{k^2(\mathbf{r})\}/\text{Re}\{k^2(\mathbf{r})\}$, and propagating the simulated wave until a steady state is reached.

We use the following PSTD algorithm [26] to solve the time-dependent scalar wave equation in presence of attenuation

$$\psi_{k+1}(\mathbf{r}) = C^{(1)}(\mathbf{r})\psi_{k-1}(\mathbf{r}) + C^{(2)}(\mathbf{r})\psi_k(\mathbf{r}) + C^{(3)}(\mathbf{r}) \left(\text{ifft}_2 \left[-|\mathbf{p}|^2 \, \text{fft}_2 \left[\psi_k(\mathbf{r}) \right] \right] + S(\mathbf{r}) \right), \tag{23}$$

where coefficient matrices $C^{(1)}$, $C^{(2)}$ and $C^{(3)}$ are pre-calculated:

$$C^{(1)}(\mathbf{r}) = \frac{\sigma(\mathbf{r})\Delta t - 2}{\sigma(\mathbf{r})\Delta t + 2}, \qquad C^{(2)}(\mathbf{r}) = \frac{4}{\sigma(\mathbf{r})\Delta t + 2}, \qquad C^{(3)}(\mathbf{r}) = \frac{2c^2(\mathbf{r})\Delta t^2}{\sigma(\mathbf{r})\Delta t + 2}. \tag{24}$$

Since PSTD uses a finite-difference scheme to approximate the time derivative, the accuracy of the method will depend on Δt. In the limit of $\Delta t \to 0$, PSTD should converge to the exact solution that is found with our modified Born approach.

We have the option to either use absorbing boundaries or periodic boundary conditions. Choosing a proper absorbing boundary condition is not trivial as any residual reflection or transmission will affect the accuracy of the simulation. We found that for our high-accuracy simulations a good trade-off between layer thickness and accuracy was achieved with the boundaries described in Appendix B.

5. Results

5.1. Homogeneous medium

In the first numerical experiment, we simulated wave propagation through a homogeneous medium with $k(\mathbf{r}) = 2\pi/\lambda$. The simulation grid had a size of 1 by 200 pixels, with a resolution of $\Delta x = \lambda/4$. A source with unit amplitude was placed at the first pixel of the homogeneous medium. We used periodic boundary conditions in the vertical direction, and appended a 25-λ-thick absorbing layer (of type $N = 4$, see Appendix B) on both horizontal sides of the medium. The resulting structure of 1 by 400 pixels was zero-padded to a 1 by 512 pixel grid to increase the performance of the fast Fourier transforms. We simulated wave propagation through this medium. First using our modified Born approach, and then using PSTD with various values for Δt. For both methods, we chose the total number of iterations such that a distance of 100 wavelengths was traversed, i.e. the complete medium twice.

For each simulation, we calculated the relative error E with respect to the analytic solution

$$E \equiv \frac{\langle |\psi - \psi_a|^2 \rangle}{\langle |\psi_a|^2 \rangle}, \tag{25}$$

where ψ_a is the exact analytic solution and $\langle \cdot \rangle$ denotes averaging over all pixels in the medium, excluding the boundaries. To calculate the ψ_a, we represent the single-pixel source as a sinc function, and convolve this source with the Green's function for a homogeneous medium. For simplicity, we only evaluate the convolution for $x > 0$
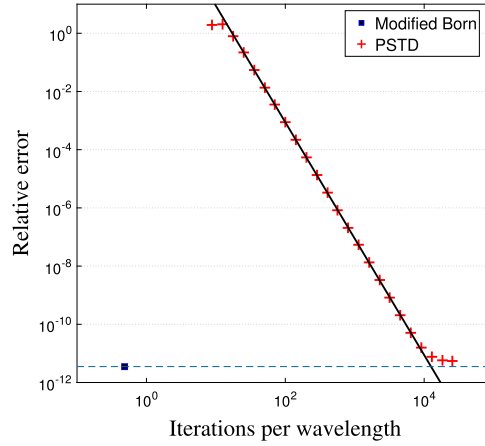
Fig. 2. Relative error for wave propagation in a 1-D homogeneous medium as function of the number of iterations per wave cycle. Plus signs denote the errors for PSTD with varying time step Δt. The square denotes the error for our modified Born approach. The solid line represents the expected error of PSTD, which is proportional to Δt^4. The dashed line at $3.5 \cdot 10^{-12}$ corresponds to the accuracy achieved with our modified Born method.

$$\psi_a(x) = \frac{\sin(\pi x/\Delta x)}{\pi x/\Delta x} * \frac{e^{ik|x|}}{-2ik} \tag{26}$$

$$= \frac{\Delta x}{4k\pi} e^{-ikx} \left[2i\pi + \text{Ei}(ik^- x) - \text{Ei}(ik^+ x) \right] + \frac{\Delta x}{4k\pi} e^{ikx} \left[\text{Ei}(-ik^- x) - \text{Ei}(-ik^+ x) \right], \tag{27}$$

with $k^- \equiv k - \pi/\Delta x$, $k^+ \equiv k + \pi/\Delta x$, and Ei as the exponential integral function. The solution rapidly converges to $\Delta x \exp(ikx)/(-2ik)$ with increasing distance from the source. Still, with the accuracy that our method achieves, the contribution of the exponential integrals (0.06% of the total energy in the simulation) cannot be neglected.

The result of the simulations are shown in Fig. 2. The figure shows that our method achieves a relative error E of less than 10^{-11} at an exceptionally low computational effort of only 0.5 iteration per propagated wavelength. For PSTD, we varied Δt from the largest possible value allowed by the convergence criterion (see Table 1), decreasing the time step until the relative error stabilized. The results of PSTD for large time steps are completely off ($> 100\%$ relative error). To achieve an accuracy of 1% over the full medium, at least 55 iterations per propagated wavelength are needed. The execution speed of both our method and PSTD are limited by the two fft$_2$'s per iteration, so the x-axis of Fig. 2 is directly proportional to simulation time.

This simple test shows that our method is two orders of magnitude faster, and at least nine orders of magnitude more accurate than a PSTD simulation that is designed to achieve a 1% accuracy. As can be seen in the figure, the accuracy of PSTD can be increased by increasing the number of steps per wave cycle. Even though the accuracy is proportional to Δt^4, it takes an extreme value of well over 10^4 iterations per wave cycle before the accuracy of PSTD matches that of our method. The last data point in Fig. 2 took 1.5 days to calculate, compared to 0.3 seconds for our modified Born series.

The main source of error in PSTD simulations is an accumulating error in the phase of the propagating wave due to an error in the calculated wavevector [16]. Under some circumstances (e.g. near-paraxial wave propagation) this error may not be a cause of concern. However, in most interesting scenarios (such as resonators, scattering media, integrated optics devices, etc.) the phase needs to be calculated with a high accuracy in order to correctly simulate multi-path wave interference. Therefore, this simulation, where only 50 wave cycles were simulated, may even give an optimistic picture of the accuracy of PSTD. The relative error of 10^{-11} that is found for our method is caused by residual reflection and leakage of the absorbing boundaries, not by the algorithm itself. By choosing boundaries that are thicker and smoother, an error of less than 10^{-17} was reached. However, since such extreme accuracies are rarely required, we used thinner boundaries, resulting in a smaller medium and shorter run times.

5.2. Disordered inhomogeneous medium

In order to demonstrate the convergence of the modified Born approach for inhomogeneous media, we simulated wave propagation through a 2-dimensional inhomogeneous medium with a random complex potential. The relative scattering potential $k^2(\mathbf{r})/k_0^2 - 1$ is normally distributed with mean $(1.30 + 0.05i)^2$ and variance $(0.10 + 0.02i)^2$. The distribution is low-pass filtered with a cut-off at $1.0\lambda^{-1}$. The simulated medium has a size of 256 by 256 pixels with a resolution of $\Delta x = \min\{\lambda/4\}$ in both spatial dimensions x and y. Periodic boundaries were used and as a source term we used a point source in the center of medium at $x = y = 0$.

First, we simulated wave propagation in the inhomogeneous medium using our method. The solution is presented in Fig. 3. Then, the numerical experiment was performed using PSTD with various numbers of iterations per wave cycle. Similar to in Eq. 25, we define a relative difference:
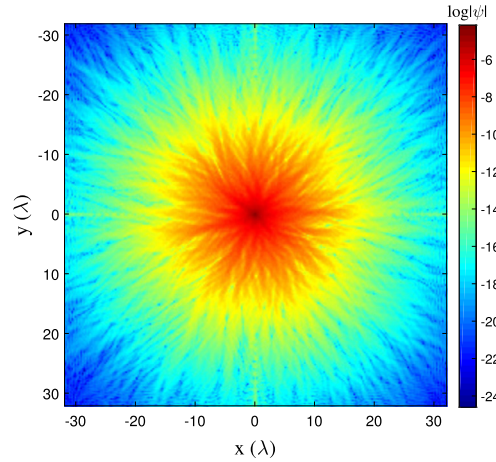
Fig. 3. Numerical experiment simulating an oscillating point source at the center of an inhomogeneous medium with a random potential distribution. The steady-state solution is found using our modified Born approach. The field amplitude is represented on a logarithmic scale. (Color version online.)

$$E_{\text{diff}} \equiv \frac{\langle |\psi_{\text{PSTD}} - \psi_{\text{Born}}|^2 \rangle}{\langle |\psi_{\text{Born}}|^2 \rangle}, \tag{28}$$

where ψ_{PSTD} and ψ_{Born} are the solutions found using PSTD and our method respectively.

The resulting PSTD solution found using a large number of iterations per wave cycle indicates that the solution found by our method was identical to the solution found by PSTD in the limit $\Delta t \rightarrow 0$, to an accuracy of at least 10^{-13}. One should note that the numerical experiment was completed in 8 seconds by the modified Born series method, whereas the PSTD method with 10^4 iterations per wavelength required in order to reach $E_{\text{diff}} < 10^{-13}$ took 20 hours to finish.

5.3. Phase conjugation numerical experiments in adipose tissue

In the third numerical experiment, we simulate a phase conjugation experiment inside a two-dimensional adipose tissue sample of 725 by 725 μm in size. Previously, media of such sizes were too time-consuming to simulate. For instance, simulating phase conjugation in a 240 μm thick medium took one week [28], whereas our method took 15 minutes to simulate a 725 μm thick medium.

We rewrite the wavenumber $k(\mathbf{r}) = n(\mathbf{r})2\pi/\lambda_0$, with $n(\mathbf{r})$ being the refractive index of the medium and λ_0 being the wavelength in vacuum. The refractive index distribution of the adipose tissue model is created using a gray scale microscopy image (Young et al. [27]). Based on the optical properties presented in the paper by Jacques [29], the refractive index of fatty cells and the extracellular fluids were estimated at 1.44 and 1.36 respectively. Given this range of refractive indices, the gray values in the microscopy image are linearly converted to a refractive index map shown in Fig. 4a. We chose $\lambda_0 = 1.0$ μm and we scaled the microscope image such that the pixel size $\Delta x = \lambda_0/4$. The resulting medium size is 2900 by 2900 pixels, to which we added a 25-λ-thick absorbing boundary layer and zero padded the medium to a size of 4096 by 4096 pixels. A point source is defined at a depth of $y = 48\lambda$ inside the tissue model and a phase conjugating mirror is placed at the top boundary of the simulation grid.

A phase conjugation experiment consist of two phases: the recording phase and the digital playback phase. In the recording phase, we simulate the light propagation from a point source through the adipose tissue medium, which yields the steady-state solution as shown in Fig. 4b. Afterwards, we take the phase conjugate of the resulting field at the phase conjugating mirror at the top boundary of the simulation grid. This conjugated field is used as the source in the playback phase, which results in the steady-state solution shown in Fig. 4c. As can be seen in this figure, the light is still scattered as it propagates deeper inside the tissue, however, now it focuses at the location of the point source of the recording phase. Phase conjugation is an important technique for high-resolution imaging inside scattering media [30] and is one the many wave phenomena that can now be studied in large scattering structures using our modified Born series approach.

5.4. 3-Dimensional media

Our method can easily be extended to solve the 3-dimensional Helmholtz equation. We simply replaced fft_2 and ifft_2 in Eq. (20) by 3-dimensional fast Fourier transforms. To demonstrate the efficiency and accuracy for a 3-dimensional problem, we first simulated wave propagation in a homogeneous medium of $8 \times 8 \times 200$ pixels. This medium is equivalent to the medium used in Section 5.1, except for the extended size in the lateral dimensions. As can be seen in Fig. 5, the increase in speed and accuracy for the 3-D problem are identical to that of the 1-D problem (Fig. 2). In addition, it can be seen that the PSTD simulations become less accurate at very small time steps ($> 10^4$ iterations per wave cycle). When so many
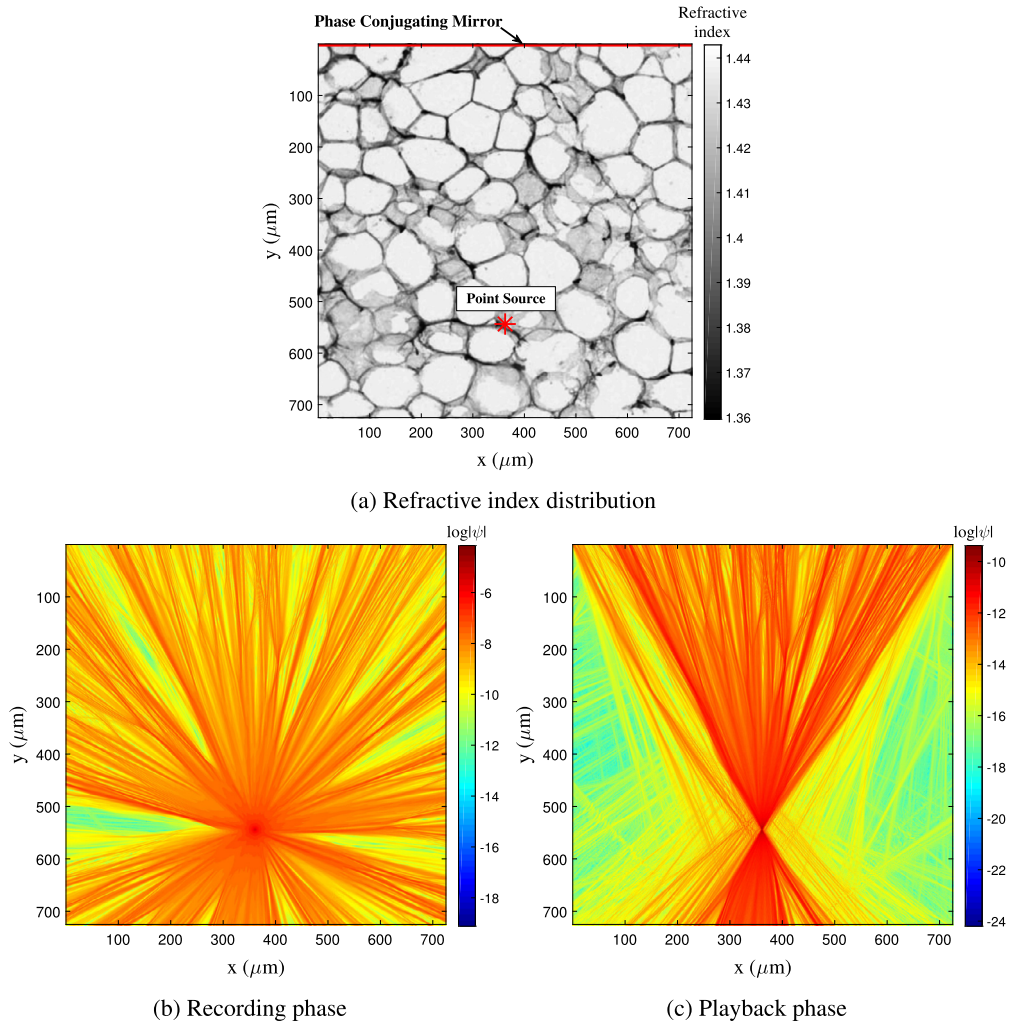
(a) Refractive index distribution

(b) Recording phase   (c) Playback phase

Fig. 4. Numerical experiment simulating phase conjugation of light coming from a point source inside a two-dimensional adipose tissue model. a) The refractive index distribution is based on a gray scale microscopy image of adipose tissue [27]. The red star marks the location of the point source and the line at the top boundary marks the location of the phase conjugating mirror. b) The steady state solution of the recording phase, where light is propagated from a point source through the adipose tissue medium. c) The steady state solution of the playback phase, where the field at the phase conjugating mirror is conjugated and propagated back into the medium. The simulation demonstrates that the light is being focused back at the original location of the point source. The field amplitudes are shown on a logarithmic scale. (Color version online.)

steps are used the rounding errors from the double precision floating point numbers start to accumulate, thereby limiting the accuracy of the result. Since the rounding errors accumulate with every iteration, further decreasing the time step only reduces the accuracy of PSTD.

Finally, we repeat the simulations of a disordered medium, this time in three spatial dimensions x, y and z. We used the same parameters as in Section 5.2, with a medium size of $128 \times 128 \times 128$ pixels (equivalent to $32 \times 32 \times 32$ times the wavelength in vacuum), and placed the source at $x = y = z = 0\lambda$. Fig. 6a–c show cross sections of the simulation results at $x = 0\lambda$, $x = 8\lambda$, and $x = 16\lambda$, respectively. Again, we compared the results of the modified Born series with PSTD and find that PSTD converges to the solution of our method in the limit $\Delta t \to 0$ to an accuracy of at least 10^{-9}. The simulations took 14.6 seconds for our method, whereas the most accurate PSTD simulation using 10^3 iterations per propagated wavelength took 19 hours.

6. Conclusion

In conclusion, we presented a fast iterative method for solving the Helmholtz equation. Our method is a modified version of the Born series. In contrast to the original Born series, our method converges for arbitrarily large structures with an arbitrarily high scattering potential. This guaranteed convergence is achieved by introducing the preconditioner γ and the non-vanishing ϵ, which localizes the energy of the Green's function without affecting the solution of the Helmholtz equation.
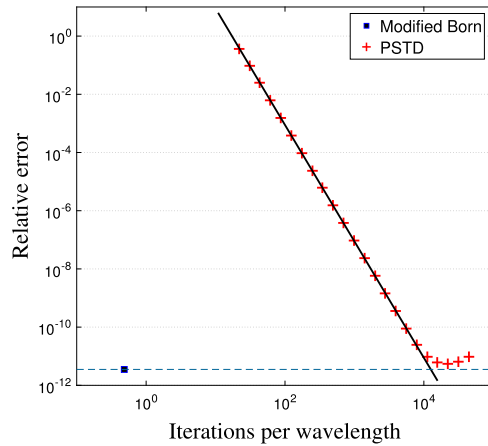
Fig. 5. Relative error for a wave propagating in a 3-D homogeneous medium as a function of a number of iterations per wave cycle for our modified Born approach and PSTD with varying time-step. The solid line is proportional to Δt^4, with Δt the PSTD time step. The dashed line at $3.5 \cdot 10^{-12}$ corresponds to the accuracy achieved with our modified Born method.
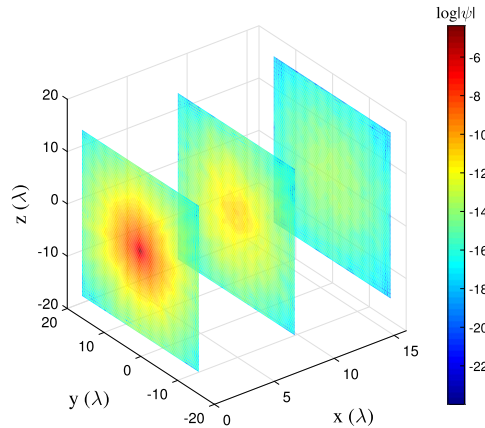

Fig. 6. Simulation results for a point source at $x = y = z = 0\lambda$ in a 3-D disordered medium of $128 \times 128 \times 128$ pixels. The 3D-solution is represented by three cross-sectional slices at $x = 0\lambda$, $x = 8\lambda$, and $x = 16\lambda$, where the field amplitude is shown on a logarithmic scale. The interference fringes at the boundaries of the slices are the result of the periodic boundary conditions. (Color version online.)

Our method is several orders of magnitude faster in finding a steady-state solution to the inhomogeneous wave equation than the commonly used PSTD and FDTD methods, especially for media with a low scattering contrast. In addition, our method converges to the exact solution, whereas finite difference methods are, by definition, limited in accuracy.

We demonstrated our method for 1, 2, and 3-dimensional media. Extensions to scalar waves in higher dimensions are trivial. If a time-dependent solution is required, our modified Born series method could be applied once for each frequency of interest, and the time-dependent solution could be calculated simply with a Fourier transform. Such an approach would still be more accurate, and potentially faster than PSTD, depending on the simulated bandwidth.

As with all grid-based numerical methods, the scattering potential has to be discretized, which imposes a limit on the spatial frequencies of the solution. This limitation will reduce the accuracy for media with very sharp interfaces. Furthermore, the speed of the modified Born series depends on the scattering contrast. Since most optical systems have refractive index typically in the range of 1 and 2, scattering contrast in optical systems is relatively small, which makes our method suitable for simulations of light propagation. However, in acoustic wave simulations, for instance, the scattering contrast can become much larger, drastically reducing the speed of our numerical method.

Future research may be directed towards finding combination of γ and ϵ for which the iterations converge more rapidly, or towards using more advanced iterative methods to speed up convergence. Additionally, due to the localized nature of our Green's function, a large medium may be split into separate domains to accommodate parallel processing or multi-domain calculations.

Appendix A. Proof of convergence

In this appendix, we will show that the series in Eq. (12) always converges for the choice of γ and ϵ as given by Eqs. (10) and (11). A sufficient condition for convergence is $\rho(M) < 1$. We will first prove that $\rho(M) \leq 1$ and consider the eigenvalues with value 1 later. Recognizing that

$$\frac{1}{|\mathbf{p}|^2 - k_0^2 - i\epsilon} = \frac{i}{2\epsilon}\left(1 - \frac{|\mathbf{p}|^2 - k_0^2 + i\epsilon}{|\mathbf{p}|^2 - k_0^2 - i\epsilon}\right),$$

(A.1)

and substituting Eq. (10) we can write operator M as

$$M = \frac{-V}{2\epsilon^2}\left[1 - F^{-1}\frac{|\mathbf{p}|^2 - k_0^2 + i\epsilon}{|\mathbf{p}|^2 - k_0^2 - i\epsilon}F\right]V - \frac{iV}{\epsilon} + 1,$$

(A.2)

$$= \frac{1}{2\epsilon^2}\left[-V^2 + VUV - 2i\epsilon V + 2\epsilon^2\right],$$

(A.3)

where

$$U \equiv F^{-1}\frac{|\mathbf{p}|^2 - k_0^2 + i\epsilon}{|\mathbf{p}|^2 - k_0^2 - i\epsilon}F$$

(A.4)

is a unitary operator. In order to prove that $\rho(M) \leq 1$, it suffices to show that $|\langle x, Mx\rangle| \leq \langle x, x\rangle$, for all x, where $\langle \cdot, \cdot \rangle$ is the inner product. We now use the Cauchy–Schwartz inequality to write $|\langle x, VUVx\rangle| = |\langle V^\dagger x, UVx\rangle| \leq \sqrt{\langle UVx, UVx\rangle}\sqrt{\langle V^\dagger x, V^\dagger x\rangle} = \langle Vx, Vx\rangle$, where we used the fact that V is diagonal (and, hence $V^\dagger V = VV^\dagger$) in the final step. With this result, we can eliminate U, resulting in

$$|\langle x, Mx\rangle| \leq \frac{1}{2\epsilon^2}\left|\left\langle x, \left[2\epsilon^2 - 2i\epsilon V - V^2\right]x\right\rangle\right| + \frac{1}{2\epsilon^2}\langle Vx, Vx\rangle.$$

(A.5)

To complete the proof, we now need to demonstrate that the right hand side of this equation is never larger than 1. Since $V = V(\mathbf{r})$, we require

$$\left|2\epsilon^2 - 2i\epsilon V(\mathbf{r}) - V^2(\mathbf{r})\right| + |V(\mathbf{r})|^2 \leq 2\epsilon^2,$$

(A.6)

for all $\mathbf{r}$. To show that this condition is always fulfilled, we define $\Delta \equiv V + i\epsilon = k^2(\mathbf{r}) - k_0^2$ and rewrite Eq. (A.6) as

$$\left|\epsilon^2 - \Delta(\mathbf{r})^2\right| + |\Delta(\mathbf{r}) - i\epsilon|^2 \leq 2\epsilon^2,$$

(A.7)

which can be written as

$$\left|\epsilon^2 - |\Delta(\mathbf{r})|^2 - 2i\Delta(\mathbf{r})\text{Im}\{\Delta(\mathbf{r})\}\right| + |\Delta(\mathbf{r})|^2 + \epsilon^2 - 2\epsilon\text{Im}\{\Delta(\mathbf{r})\} \leq 2\epsilon^2.$$

(A.8)

A slightly stricter criterion follows from triangle inequality

$$\left|\epsilon^2 - |\Delta(\mathbf{r})|^2\right| + 2|\Delta(\mathbf{r})|\text{Im}\{\Delta(\mathbf{r})\} + |\Delta(\mathbf{r})|^2 + \epsilon^2 - 2\epsilon\text{Im}\{\Delta(\mathbf{r})\} \leq 2\epsilon^2,$$

(A.9)

where we require that $\text{Im}\{\Delta(\mathbf{r})\} \geq 0$, which means that the medium cannot have any gain. Since we have $\epsilon \geq |\Delta(\mathbf{r})|$ from Eq. (11), condition (A.9) is always fulfilled, and therefore $\rho(M) \leq 1$.

Eigenvalues of 1 are only possible for an infinite non-absorbing medium. In this case, the solutions to the Helmholtz equation carry infinite total energy, which means that the solution cannot be found using our method. However, if there are absorbing boundaries, or even if there is a single a finite-size volume with non-zero absorption, there will be some points where the wave is absorbed. In this case the left hand side of Eq. (A.9) is strictly less than the right hand side, and convergence of our method is guaranteed.

Appendix B. Boundary conditions

In the implementation of our method, we use a fast Fourier transform to evaluate the convolution with the Green's function. As a result, the system has periodic boundary conditions by default. In order to prevent waves to 'wrap around' the boundaries, we implement absorbing boundary layers. A wide choice of boundary conditions is available for simulating scalar wave propagation in finite size systems (e.g. [31,32]). We choose to use a type of absorbing boundaries with the following properties: firstly, the layer is designed to have zero reflectivity for normal incidence. Secondly, the scattering potential $|k^2(\mathbf{r}) - k_0^2|$ of the layer is bounded to a specified maximum value.

We design the absorbing layers by requiring that the wave has the following form

$$\psi(x) \propto P_N(x)e^{ik_0 x - \alpha x}, \tag{B.1}$$

with P_N an N'th order polynomial inside the absorbing layer ($x > 0$) and equal to 1 for $x \le 0$. Substituting this desired solution into Eq. (1) gives

$$P_N''(x) + 2(ik_0 - \alpha)P_N'(x) + \left[k^2(x) - k_0^2 - 2ik_0\alpha + \alpha^2\right]P_N(x) = 0. \tag{B.2}$$

The simplest solution is found for $N = 1$:

$$P_1(x > 0) = 1 + \alpha x, \tag{B.3}$$

$$k^2(x) - k_0^2 = \frac{\alpha^2(1 - \alpha x + 2ik_0 x)}{1 + \alpha x}. \tag{B.4}$$

Theoretically, this absorbing boundary has zero reflectivity. However, due to the discretization of the field, waves with a high spatial frequency cannot be represented. These frequency components are required to truthfully represent the waves at the transition from the medium to the absorbing boundary. When these components are missing, some residual reflectivity results. We found that this reflectivity can be reduced by smoothing the boundary. This can be achieved by imposing the constraint that the function $k^2(x) - k_0^2$ and its first $N - 2$ derivatives vanish at $x = 0$. A general solution for a polynomial of degree N is given by

$$P_N(x > 0) = \sum_{n=0}^{N} \frac{(\alpha x)^n}{n!} \tag{B.5}$$

$$k^2(x) - k_0^2 = \frac{\alpha^2(N - \alpha x + 2ik_0 x)(\alpha x)^{N-1}}{P_N(x)N!}. \tag{B.6}$$

Choosing a higher value for N results in a smoother boundary with a lower residual reflection. However, it also takes longer for the boundary to reach its maximum absorption coefficient of α. In Eq. (B.5) it can be seen that $P_N(x)$ equals the first N terms of the Taylor expansion of $\exp(\alpha x)$, thereby partially canceling the decay term $\exp(-\alpha x)$ close to the boundary.

The boundary potential saturates at a value of $k^2(x) - k_0^2 = -\alpha^2 + 2ik_0\alpha$ for $x \to \infty$, irrespective of N. Higher values of α result in stronger absorption, at the cost of increasing the scattering potential. Because of condition Eq. (11), a higher scattering potential results in a higher value of ϵ, which gives a lower pseudo-propagation speed (and hence, slower convergence of the method). Throughout this manuscript used layers with a thickness of 25 wavelengths, with $N = 4$ and $\max|k^2(\mathbf{r}) - k_0^2| = 0.2$. These parameters gave a good balance between scattering potential, residual reflectivity, and residual transmission for our high-accuracy simulations.

References

[1] M.C.W. van Rossum, T.M. Nieuwenhuizen, Multiple scattering of classical waves: microscopy, mesoscopy and diffusion, Rev. Mod. Phys. 71 (1999) 313–371.
[2] Y.A. Erlangga, Advances in iterative methods and preconditioners for the Helmholtz equation, Arch. Comput. Methods Eng. 15 (1) (2008) 37–66.
[3] M. Nabavi, M.H.K. Siddiqui, J. Dargahi, A new 9-point sixth-order accurate compact finite-difference method for the Helmholtz equation, J. Sound Vib. 307 (3) (2007) 972–982.
[4] J.M. Melenk, S. Sauter, Convergence analysis for finite element discretizations of the Helmholtz equation with Dirichlet-to-Neumann boundary conditions, Math. Comput. 79 (2010) 1871–1914, http://dx.doi.org/10.1090/S0025-5718-10-02362-8.
[5] L.L. Thompson, A review of finite-element methods for time-harmonic acoustics, J. Acoust. Soc. Am. 119 (3) (2006) 1315–1330, http://dx.doi.org/10.1121/1.2164987.
[6] P.G. Martinsson, A fast direct solver for a class of elliptic partial differential equations, J. Sci. Comput. 38 (3) (2009) 316–330, http://dx.doi.org/10.1007/s10915-008-9240-6.
[7] P.G. Schmitz, L. Ying, A fast direct solver for elliptic problems on general meshes in 2d, J. Comput. Phys. 231 (4) (2012) 1314–1338, http://dx.doi.org/10.1016/j.jcp.2011.10.013.
[8] M.J. Gander, I.G. Graham, E.A. Spence, Applying GMRES to the Helmholtz equation with shifted Laplacian preconditioning: what is the largest shift for which wavenumber-independent convergence is guaranteed?, Numer. Math. 131 (3) (2015) 567–614, http://dx.doi.org/10.1007/s00211-015-0700-2.
[9] C.C. Stolk, M. Ahmed, S.K. Bhowmik, A multigrid method for the Helmholtz equation with optimized coarse grid corrections, SIAM J. Sci. Comput. 36 (2014) 2819–2841, http://dx.doi.org/10.1137/13092349X.

[10] Y.A. Erlangga, C.W. Oosterlee, C. Vuik, A novel multigrid based preconditioner for heterogeneous Helmholtz problems, SIAM J. Sci. Comput. 27 (4) (2006) 1471–1492, http://dx.doi.org/10.1137/040615195.

[11] C.C. Stolk, A rapidly converging domain decomposition method for the Helmholtz equation, J. Comput. Phys. 241 (2013) 240–252, http://dx.doi.org/10.1016/j.jcp.2013.01.039.

[12] D. Gordon, R. Gordon, Robust and highly scalable parallel solution of the Helmholtz equation with large wave numbers, J. Comput. Appl. Math. 237 (1) (2013) 182–196, http://dx.doi.org/10.1016/j.cam.2012.07.024.

[13] O.G. Ernst, M.J. Gander, Why it is difficult to solve Helmholtz problems with classical iterative methods, in: Numerical Analysis of Multiscale Problems, vol. 83, Springer, Berlin, Heidelberg, 2012, pp. 325–363.

[14] A. Taflove, S.C. Hagness, Computational Electrodynamics: The Finite-Difference Time-Domain Method, Antennas and Propagation Library, Artech House, 1995.

[15] K. Yee, Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media, IEEE Trans. Antennas Propag. 14 (3) (1966) 302–307, http://dx.doi.org/10.1109/TAP.1966.1138693.

[16] Q.H. Liu, The PSTD algorithm: a time-domain method requiring only two cells per wavelength, Microw. Opt. Technol. Lett. 15 (3) (1997) 158–165, http://dx.doi.org/10.1002/(SICI)1098-2760.

[17] J.L. Vay, I. Haber, B.B. Godfrey, A domain decomposition method for pseudo-spectral electromagnetic simulations of plasmas, J. Comput. Phys. 243 (2013) 260–268, http://dx.doi.org/10.1016/j.jcp.2013.03.010.

[18] S.H. Tseng, A. Taflove, D. Maitland, V. Backman, Pseudospectral time domain simulations of multiple light scattering in three-dimensional macroscopic random media, Radio Sci. 41 (4) (2006), http://dx.doi.org/10.1029/2005RS003408.

[19] B. Gallinet, J. Butet, O.J.F. Martin, Numerical methods for nanophotonics: standard problems and future challenges, Laser Photonics Rev. 9 (6) (2015) 577–603, http://dx.doi.org/10.1002/lpor.201500122.

[20] G. Beylkin, C. Kurcz, L. Monzón, Fast convolution with the free space Helmholtz Green's function, J. Comput. Phys. 228 (8) (2009) 2770–2791, http://dx.doi.org/10.1016/j.jcp.2008.12.027.

[21] M. Wubs, A. Lagendijk, Local optical density of states in finite crystals of plane scatterers, Phys. Rev. E 65 (2002), http://dx.doi.org/10.1103/PhysRevE.65.046612.

[22] P.T. Kristensen, P. Lodahl, J. Mørk, Light propagation in finite-sized photonic crystals: multiple scattering using an electric field integral equation, J. Opt. Soc. Am. B 27 (2) (2010) 228–237, http://dx.doi.org/10.1364/JOSAB.27.000228.

[23] O.J.F. Martin, A. Dereux, C. Girard, Iterative scheme for computing exactly the total field propagating in dielectric structures of arbitrary shape, J. Opt. Soc. Am. A 11 (3) (1994) 1073–1080, http://dx.doi.org/10.1364/JOSAA.11.001073.

[24] R.E. Kleinman, G.F. Roach, P.M. van den Berg, Convergent Born series for large refractive indices, J. Opt. Soc. Am. A 7 (5) (1990) 890–897, http://dx.doi.org/10.1364/JOSAA.7.000890.

[25] S. Trattner, M. Feigin, H. Greenspan, N. Sochen, Validity criterion for the born approximation convergence in microscopy imaging, J. Opt. Soc. Am. A 28 (4) (2011) 665–666, http://dx.doi.org/10.1364/JOSAA.28.000665.

[26] C. Spa, P. Reche-López, E. Hernández, Numerical absorbing boundary conditions based on a damped wave equation for pseudospectral time-domain acoustic simulations, Sci. World J. 2014 (2014) 343–358, http://dx.doi.org/10.1155/2014/285945.

[27] D.A. Young, K.L. Christman, Injectable biomaterials for adipose tissue engineering, Biomed. Mater. 29 (6) (2012) 997–1003, http://dx.doi.org/10.1088/1748-6041/7/2/024104.

[28] S.H. Tseng, PSTD simulation of optical phase conjugation of light propagating long optical paths, Opt. Express 17 (7) (2009) 5490–5495.

[29] S.L. Jacques, Optical properties of biological tissues: a review, Phys. Med. Biol. 58 (11) (2013) R37–R61, http://dx.doi.org/10.1088/0031-9155/58/11/R37.

[30] R. Horstmeyer, H. Ruan, C. Yang, Guidestar-assisted wavefront-shaping methods for focusing light into biological tissue, Nat. Photonics 9 (1) (2015) 564–571, http://dx.doi.org/10.1038/nphoton.2015.140.

[31] J.P. Berenger, A perfectly matched layer for the absorption of electromagnetic waves, J. Comput. Phys. 114 (2) (1994) 185–200, http://dx.doi.org/10.1006/jcph.1994.1159.

[32] Q.H. Liu, The pseudospectral time-domain (PSTD) algorithm for acoustic waves in absorptive media, IEEE Trans. Ultrason. Ferroelectr. Freq. Control 45 (4) (1998) 1044–1055, http://dx.doi.org/10.1109/58.710587.