Efficient Helmholtz Equation Solver for Frequency Domain Waveform Inversion Based on the Decomposition into One-Way Wave Equations

Rehman Ali*, Feiyu Wang**, Trevor Mitcham*, and Nebojsa Duric*

*Department of Imaging Sciences, URMC, Rochester, NY, 14620 USA
**Department of Mechanical Engineering, University of Rochester, Rochester, NY, 14620 USA

ABSTRACT

When using a ring array to perform ultrasound tomography, the most computationally intensive component of frequency-domain full waveform inversion (FWI) is the Helmholtz equation solver. The Helmholtz equation is an elliptic partial differential equation (PDE) whose discretization leads to a large system of equations; in many cases, the solution of this large system is itself the inverse problem and requires an iterative method. Our current solution relies on discretizing the 2D Helmholtz equation based on a 9-point stencil and using the resulting block tridiagonal structure to efficiently compute a block LU factorization. Conceptually, the L and U systems are equivalent to a forward and backward wave propagation along one of the spatial dimensions of the grid, resulting in a direct non-iterative solution to the Helmholtz equation based on a single forward and backward sweep. Based on this observation, the PDE representation of the Helmholtz equation is split into two one-way wave equations prior to discretization. The numerical implementations of these one-way wave equations are highly parallelizable and lend themselves favorably to accelerated GPU implementations. We consider the Fourier split-step and phase-shift-plus-interpolation (PSPI) methods from seismic imaging as numerical solutions to the one-way wave equations. We examine how each scheme affects the numerical accuracy of the final Helmholtz equation solution and present its impact on FWI with breast imaging examples.

Keywords: Nonlinear inverse problems, waveform inversion, ultrasound, tomography

1. INTRODUCTION

Ultrasound tomography (UST) uses the transmission of ultrasound through tissue to reconstruct high-resolution images of tissue properties such as sound speed and attenuation. The most common clinical use of UST is breast imaging[1–4] where sound speed and attenuation serve as biomarkers for cancer detection. The full waveform inversion (FWI) algorithm responsible for high-resolution UST imaging relies on the accurate simulation of pressure fields in a heterogeneous medium. Although time-domain FWI[5,6] can accommodate a larger variety of inversion strategies, frequency-domain FWI is often preferred in clinical implementations due to its reduced memory and computational costs. Currently, there are two competing geometries for clinical UST imaging: 1) ring array imaging,[7] and 2) linear-arrays in a pitch-catch setup that are rotated around the tissue-of-interest.[4,8] Both of these imaging approaches rely on frequency-domain FWI to reconstruct the image. The obvious benefit of the ring array system over the rotating pitch-catch setup is that transducer motion is not necessary to obtain the 360 degree view around the tissue, thereby avoiding artifacts of tissue motion and long scan times.[7]

On the other hand, the pitch-catch setup with the linear array transducers allows for the use of the Fourier split-step method to simulate the pressure field for frequency-domain waveform inversion.[4,8] This is a significant computational advantage over ring array imaging which must rely on solutions to the full Helmholtz equation, an elliptic partial differential equation whose solution often depends on solving a large system of equations. The Fourier split-step method reduces the full Helmholtz equation into a one-way wave equation that replaces the system of equations with a fast and explicit depth-advancement scheme for simulating the pressure field. The main caveat of the Fourier split-step method is the limited angular range of wave propagation that can

Further author information: (Send correspondence to Rehman Ali)
Rehman Ali: E-mail: rali8@stanford.edu

Medical Imaging 2024: Ultrasonic Imaging and Tomography, edited by Christian Boehm, Nick Bottenus,
Proc. of SPIE Vol. 12932, 1293216 · © 2024 SPIE · 1605-7422
doi: 10.1117/12.3006288

be simulated through the tissue. This limitation is imposed by the strong directionality of the one-way wave equation. Fortunately, for plane waves emitted and received by linear transducer arrays in a pitch-catch setup,[4,8] this caveat of the Fourier split-step method is practically negligible. However, in ring array imaging, the larger angular range of wave propagation that must be simulated from each transmitting element cannot be adequately simulated by the Fourier split step method; the full omnidirectional solution to the Helmholtz equation becomes necessary. The computational advantage of the pitch-catch setup is so significant that it readily enables 3D imaging based on the Fourier split-step method.[9,10] A similar extension to 3D would be far more difficult to accomplish with the full Helmholtz equation as it would require solving a much larger system of equations. In these situations, time-domain solutions become much more attractive.[11]

Therefore, the practical advantages of the ring-array setup (e.g., less motion and reduced scan times) are offset by computational costs. The goal of this work is to disrupt this apparent tradeoff by bringing the advantages of the Fourier split-step method to ring array imaging without its current drawbacks.

2. BACKGROUND

For ring-array ultrasound tomography, the most computationally demanding aspect of frequency-domain FWI is the Helmholtz equation, an elliptic partial differential equation (PDE) that produces a large system of equations when discretized over a computational grid:

$$A(\omega, \mathbf{s})\mathbf{u}(\omega, \mathbf{s}) = \delta(\omega), \tag{1}$$

where $\omega = 2\pi f$ is the frequency, $\mathbf{s}$ is slowness, $\mathbf{u}$ is the pressure field for a given source δ, and A represents the discretization of the Helmholtz equation operator $(\nabla^2 + \omega^2\mathbf{s}^2)$ over the computational grid. For an $N \times N$ cartesian grid, A is an $N^2 \times N^2$ system of equations. Fortunately, A is often sparse, and depending on the numerical scheme, this sparsity may be structured favorably. For example, when a 9-point stencil[12] is used, A has a block tridiagonal structure[13] with $N \times N$ blocks that are themselves tridiagonal:

$$A = \begin{bmatrix} D_1 & U_1 & & & \\ L_1 & D_2 & U_2 & & \\ & \ddots & \ddots & \ddots & \\ & & L_{N-2} & D_{N-1} & U_{N-1} \\ & & & L_{N-1} & D_N \end{bmatrix}, \tag{2}$$

operating on the corresponding partitions of $\mathbf{u}$ and δ, respectively (see Figure 1):

$$\mathbf{u} = \begin{bmatrix} \vec{u}_1 \\ \vec{u}_2 \\ \vdots \\ \vec{u}_{N-1} \\ \vec{u}_N \end{bmatrix} \quad \delta = \begin{bmatrix} \vec{\delta}_1 \\ \vec{\delta}_2 \\ \vdots \\ \vec{\delta}_{N-1} \\ \vec{\delta}_N \end{bmatrix}. \tag{3}$$

This tridiagonal structure leads to an efficient solution via the block LU decomposition[13] $A = LU$:

$$L = \begin{bmatrix} T_1 & & & & \\ L_1 & T_2 & & & \\ & \ddots & \ddots & & \\ & & L_{N-2} & T_{N-1} & \\ & & & L_{N-1} & T_N \end{bmatrix}, \tag{4}$$

$$U = \begin{bmatrix} I & T_1^{-1}U_1 & & & \\ & I & T_2^{-1}U_2 & & \\ & & \ddots & \ddots & \\ & & & I & T_{N-1}^{-1}U_{N-1} \\ & & & & I \end{bmatrix}, \tag{5}$$

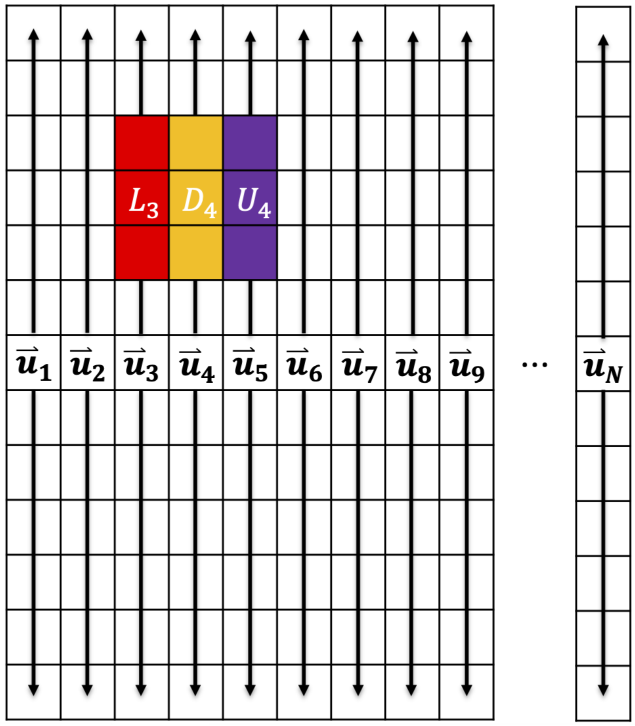

Figure 1. Visualization of the Computational Grid with the 9-Point Stencil Discretization. The full pressure field $\mathbf{u}$ is split into columns $\vec{u}_1$, $\vec{u}_2$, $\ldots$, $\vec{u}_N$ on the computational grid. The L_i (red), D_i (orange), and U_i (purple) blocks of the block-tridiagonal A matrix result from the 9-point stencil ($L_i/D_i/U_i$ partitions shown).

where $T_1 = D_1$ and $T_n = D_n - L_{n-1}T_{n-1}^{-1}U_{n-1}$ for $n = 2, ..., N$. Note that this recurrence relation results in fully dense T_n matrices, so the storage requirements for all the T_n blocks is N^3. This block LU decomposition is then used to solve equation (1) $[A\mathbf{u} = (LU)\mathbf{u} = L(U\mathbf{u}) = \delta]$ in two steps: 1) $L\mathbf{v} = \delta$, and 2) $U\mathbf{u} = \mathbf{v}$, where $\mathbf{v}$ is partitioned similar to $\mathbf{u}$ and δ:

$$\mathbf{v} = \begin{bmatrix} \vec{v}_1 \\ \vec{v}_2 \\ \vdots \\ \vec{v}_{N-1} \\ \vec{v}_N \end{bmatrix}. \tag{6}$$

The L-system is solved using the forward sweep: $\vec{v}_1 = T_1^{-1}\vec{\delta}_1$ and $\vec{v}_n = T_n^{-1}(\vec{\delta}_n - L_{n-1}\vec{v}_{n-1})$ for $n = 2, ..., N$. The U-system is solved using the backward sweep: $\vec{u}_N = \vec{v}_N$ and $\vec{u}_n = \vec{v}_n - T_n^{-1}U_n\vec{u}_{n+1}$ for $n = 2, ..., N$.

3. INNOVATION

The forward and backward sweeps created by the L- and U-systems are one-way wave propagations in opposite directions along the same axis (see Figure 2 for an example). Rather than perform this block LU decomposition after discretizing the Helmholtz equation, we can perform an analogous splitting of the Helmholtz equation in its PDE-form directly:

$$(\nabla^2 + \omega^2\mathbf{s}^2)\mathbf{u} = \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial z^2} + \omega^2\mathbf{s}^2\right)\mathbf{u} = \left(\frac{\partial}{\partial x} + i\sqrt{\omega^2\mathbf{s}^2 + \frac{\partial^2}{\partial z^2}}\right)\left(\frac{\partial}{\partial x} - i\sqrt{\omega^2\mathbf{s}^2 + \frac{\partial^2}{\partial z^2}}\right)\mathbf{u} = \delta, \tag{7}$$

with one-way wave equations $\left(\frac{\partial}{\partial x} + i\sqrt{\omega^2\mathbf{s}^2 + \frac{\partial^2}{\partial z^2}}\right)\mathbf{v} = \delta$ and $\left(\frac{\partial}{\partial x} - i\sqrt{\omega^2\mathbf{s}^2 + \frac{\partial^2}{\partial z^2}}\right)\mathbf{u} = \mathbf{v}$, respectively.

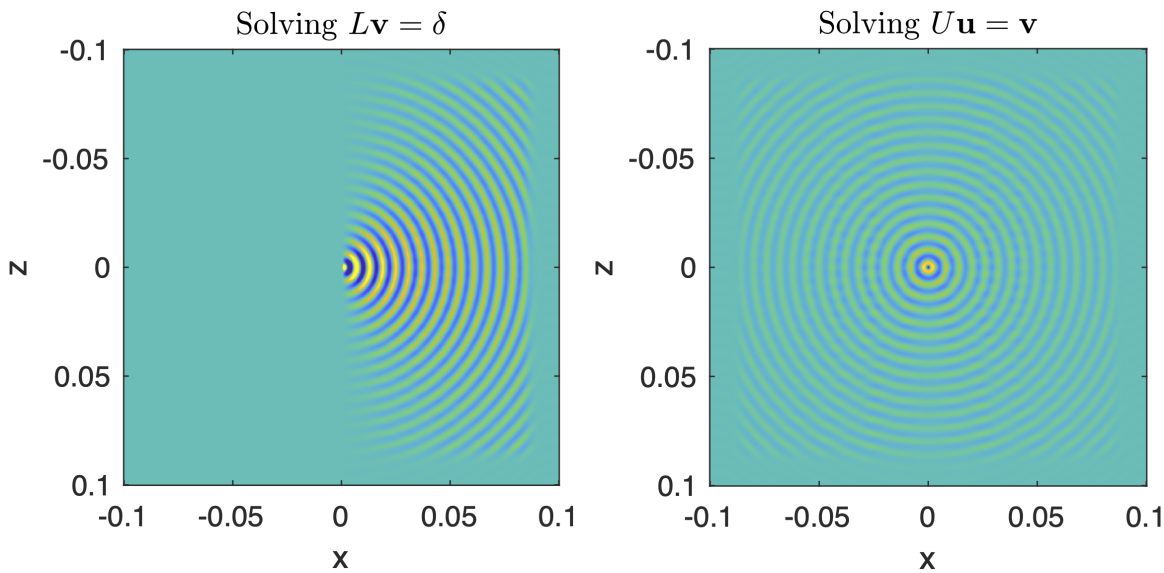

Figure 2. Visual demonstration of the Block LU Helmholtz Equation Solver: Forward sweep via $L\mathbf{v} = \delta$ and backward sweep via $U\mathbf{u} = \mathbf{v}$.

The two most common methods for implementing the one-way wave equations are the Fourier split-step method[4,9,14] and phase-shift-plus-interpolation (PSPI).[15] These mixed-domain methods rely on the Fourier transform and operate in both wavenumber and space, making them highly parallelizable for a fast and efficient GPU implementation. Another motivation behind this work is to find tractable solutions to the 3D Helmholtz equation. The block LU decomposition used in 2D is no longer viable in 3D because the T_n matrices would be fully dense $N^2 \times N^2$ matrices (assuming an $N \times N \times N$ grid). The one-way wave equation solutions remain tractable. Therefore, the purpose of this work is to demonstrate one-way wave equation solutions to the Helmholtz equation that can be incorporated into existing waveform inversion algorithms.

4. METHODS

4.1 Verification of the LU-like Decomposition Based on One-Way Wave Equations

For a given sound speed and attenuation map (real and imaginary part of slowness, respectively), we simulate the pressure field for a given source based on: 1) the Block LU decomposition[13] based on the 9-point stencil discretization,[12] 2) the one-way wave equation decomposition based on the Fourier split-step method, and 3) the one-way wave equation decomposition based on PSPI. Method (1) is a verified solution to the Helmholtz equation. Therefore, methods (2) and (3) will be tested against method (1) to assess accuracy. The normalized root-mean-square error (NRMSE) between (1) & (2) and (1) & (3), will be used as a metric to assess the relative accuracy of methods (2) and (3).

4.2 Impact on FWI with Experimental Data

Please see our prior work[16] for full details on how the Helmholtz equation is used within FWI. We test three different numerical solvers for the Helmholtz equation within an FWI code: 1) block LU decomposition[13] with 9-point stencil,[12] 2) one-way wave equation decomposition with the Fourier split-step method,[14] and 3) one-way wave equation decomposition with PSPI.[15] Raw ultrasound channel data from a whole-breast ultrasound tomography system[16,17] (22-cm diameter ring; 1024 elements) was used to reconstruct sound speed and attenuation in the breast using FWI. The FWI reconstructions for each choice of Helmholtz equation solver are compared. The FWI reconstructions were done in two passes from 0.3 to 1.25 MHz in 50 kHz steps: on the first pass, 3 iterations of conjugate gradient (CG) were used to reconstruct sound speed at each frequency; on the second pass, 3 iterations of CG were used to reconstruct both sound speed and attenuation at each frequency.

5. RESULTS AND DISCUSSION

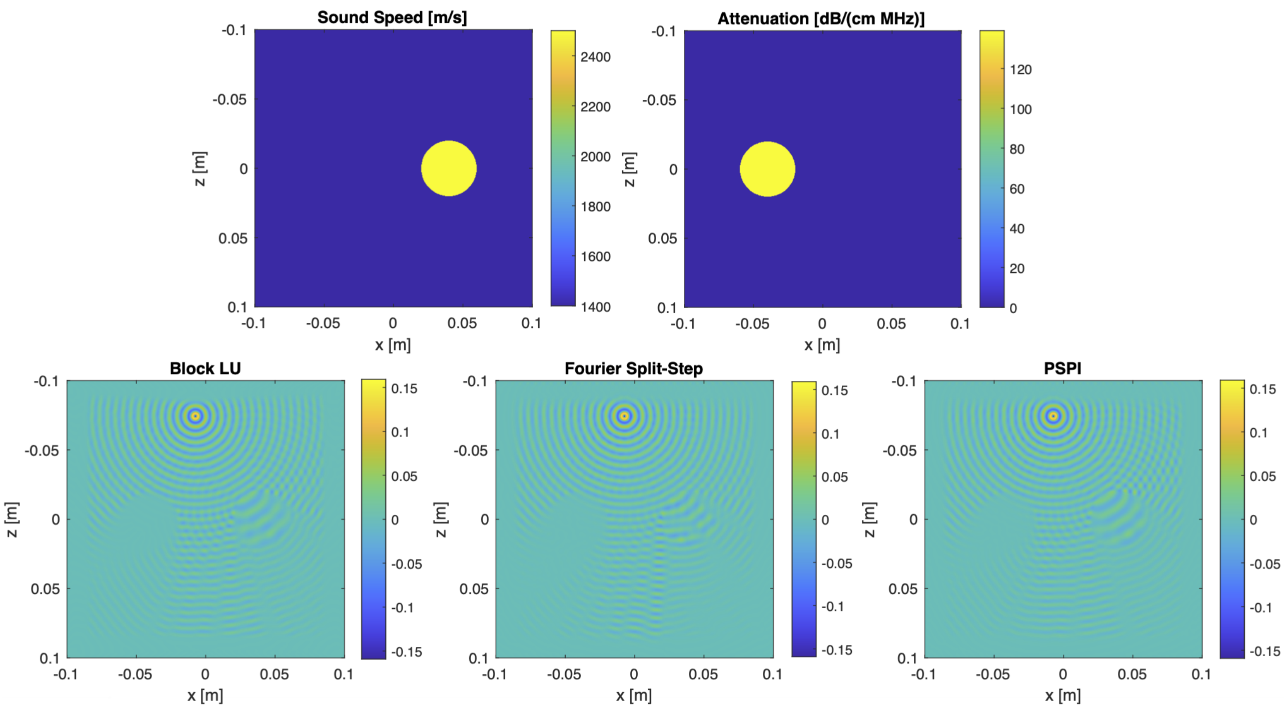

Figure 3. Comparison of One-Way Wave Equation Decomposition (via Fourier split-step method and PSPI) to the Block LU Decomposition for the Helmholtz Equation.

Figure 3 show the simulated pressure fields in a medium with circular 2500 m/s sound speed inclusion in a 1400 m/s background and a 138.974 dB/(cm MHz) attenuation inclusion in a zero attenuation background. Both inclusions were 4 cm in diameter and spaced 8 cm apart center-to-center. The pressure fields were simulated based on (1) the Block LU decomposition,[12,13] (2) the one-way wave equation decomposition with a Fourier split-step implementation, and (3) the one-way wave equation decomposition with a PSPI implementation. The NRMSE between (1) & (2) was 16.46 % while the NRMSE between (1) & (3) was only 5.78 %. Visually, we see that the Fourier split-step method is most inaccurate inside the high sound speed inclusion while PSPI shows much greater agreement with the block LU solution inside the sound speed inclusion. The Fourier split-step method propagates the wave by first operating in the wavenumber domain and then applying a phase screen in the spatial domain to account for lateral heterogeneities. However, implicit in this spatial step is a linearization that is only valid for small lateral heterogeneities. The PSPI method improves on the Fourier split-step concept by applying the method at several different reference velocities (we choose 9 reference velocities spanning the range of sound speeds in each layer) and interpolating between the results to improve accuracy. Based on the results of Figure 3, the additional interpolation step in PSPI greatly improves the accuracy of the one-way wave equation decomposition. These results extend to FWI as well.

Figure 4 shows the reconstructed sound speed and attenuation in the breast using FWI with each of the three Helmholtz equation solver. The first column of Figure 4 shows the result of using the block LU solver, the most accurate but also the most memory and computation intensive of the solvers. The second and third columns of Figures 4 show the two methods for implementing the decomposition into one-way wave equations. The explicit directionality of these solvers introduces errors into the final solution. The numerical inaccuracies of the Fourier split-step method result in the strong artifacts seen in the reconstructed sound speed and attenuation images. Although some artifacts may remain with PSPI, the images produced as a result of PSPI are a significant improvement upon the images resulting from the Fourier split-step method.

We present efficient and parallelable solutions to the Helmholtz equation based on the decomposition into

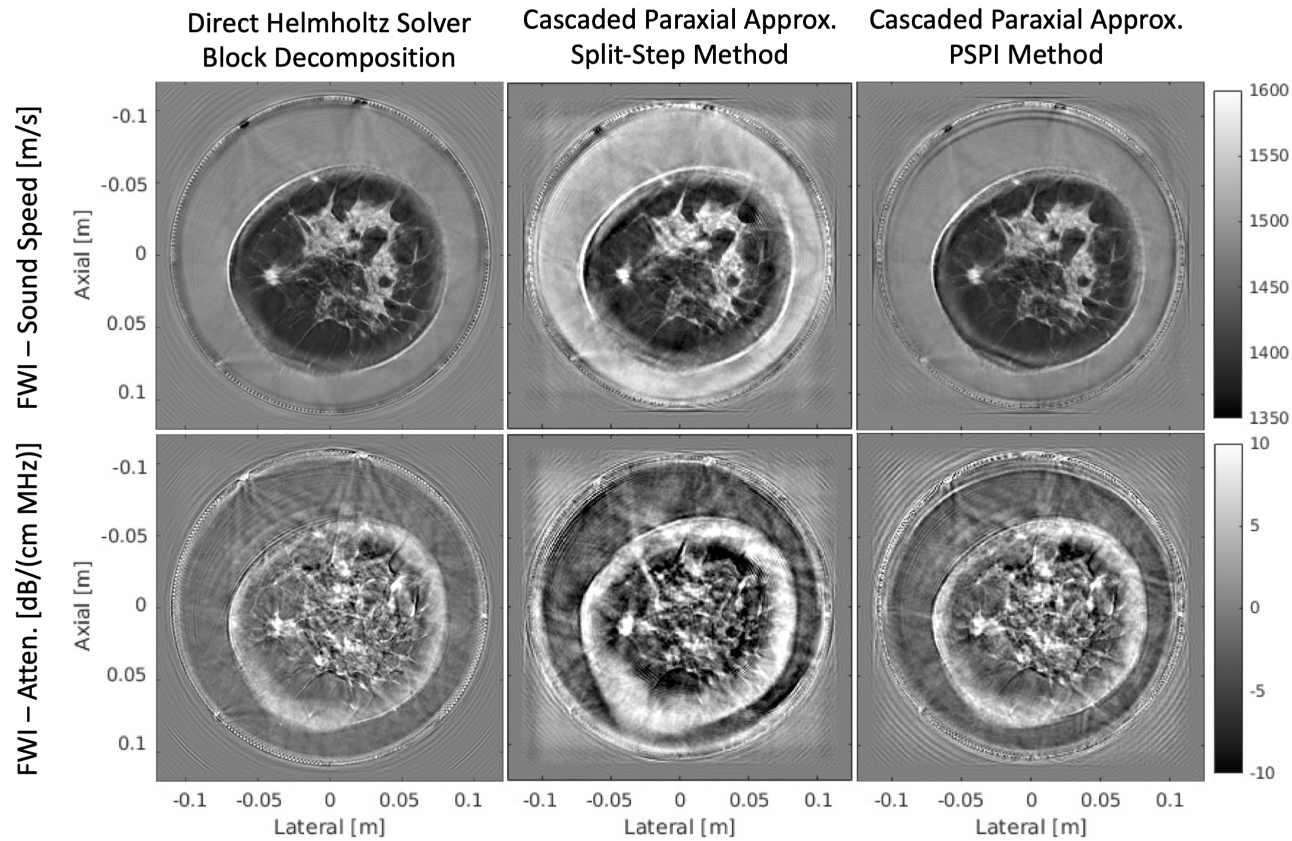

Figure 4. FWI reconstructions of sound speed and attenuation based on three Helmholtz equation solvers: 1) block LU decomposition, 2) Fourier split-step solutions to the one-way wave equation decomposition, and 3) PSPI solutions to the one-way wave equation decomposition.

one-way wave equations. The first of these methods is the Fourier split-step method, whose accuracy decreases for propagation directions orthogonal to the axis along which the decomposition is performed. The second method, known as PSPI, interpolates between different applications of the split-step method to reduce those inaccuracies. These solutions are presented in the context of frequency-domain FWI in ring-array ultrasound tomography. This is the first time these concepts have been utilized outside of a plane-wave imaging setup.[4, 9, 10]

6. CONCLUSIONS

The one-way wave equations lead to computationally efficient and parallelizable solvers for the Helmholtz equation. However, the choice of numerical method impacts the accuracy of the final solution. We show that the PSPI method is the most accurate of these methods and can be used to accurately reconstruct sound speed and attenuation in breast tissue using FWI. These results have positive implications for future work on 3D imaging.

ACKNOWLEDGMENTS

The authors would like to acknowledge Dr. Cuiping Li at Delphinus Medical Technologies for her valuable input and advice on this work. The authors would also like to thank Electrical Engineering PhD student Thurston Brevett from Stanford University for helping the authors to optimize their GPU implementation of the block LU decomposition and waveform inversion routines.

REFERENCES

[1] Li, C., Duric, N., Littrup, P., and Huang, L., "In-vivo breast sound-speed imaging with ultrasound tomography," *Ultrasound in Medicine and Biology* **35**(10), 1615–1628 (2009).

[2] Duric, N., Littrup, P., Poulo, L., Babkin, A., Pevzner, R., Holsapple, E., Rama, O., and Glide, C., "Detection of breast cancer with ultrasound tomography: First results with the computed ultrasound risk evaluation (CURE) prototype," *Medical physics* **34**(2), 773–785 (2007).

[3] Schreiman, J., Gisvold, J., Greenleaf, J. F., and Bahn, R., "Ultrasound transmission computed tomography of the breast.," *Radiology* **150**(2), 523–530 (1984).

[4] Wiskin, J., Borup, D., Johnson, S., and Berggren, M., "Non-linear inverse scattering: high resolution quantitative breast tissue tomography," *The Journal of the Acoustical Society of America* **131**(5), 3802–3813 (2012).

[5] Warner, M. and Guasch, L., "Adaptive waveform inversion: Theory," *Geophysics* **81**(6), R429–R445 (2016).

[6] Song, C., Wang, Y., Richardson, A., and Liu, C., "Weighted envelope correlation-based waveform inversion using automatic differentiation," *IEEE Transactions on Geoscience and Remote Sensing* (2023).

[7] Sandhu, G., Li, C., Roy, O., Schmidt, S., and Duric, N., "Frequency domain ultrasound waveform tomography: breast imaging using a ring transducer," *Physics in Medicine & Biology* **60**(14), 5381 (2015).

[8] Ali, R., "Open-source full-waveform ultrasound computed tomography based on the angular spectrum method using linear arrays," in [*Medical Imaging 2022: Ultrasonic Imaging and Tomography*], **12038**, 187–205, SPIE (2022).

[9] Wiskin, J., Borup, D., Iuanow, E., Klock, J., and Lenox, M. W., "3-d nonlinear acoustic inverse scattering: Algorithm and quantitative results," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* **64**(8), 1161–1174 (2017).

[10] Wiskin, J., Malik, B., Natesan, R., Borup, D., Pirshafiey, N., Lenox, M., and Klock, J., "Full wave 3d inverse scattering transmission ultrasound tomography: breast and whole body imaging," in [*2019 IEEE International Ultrasonics Symposium (IUS)*], 951–958, IEEE (2019).

[11] Appelo, D., Garcia, F., and Runborg, O., "Waveholtz: Iterative solution of the helmholtz equation via the wave equation," *SIAM Journal on Scientific Computing* **42**(4), A1950–A1983 (2020).

[12] Chen, Z., Cheng, D., Feng, W., and Wu, T., "An optimal 9-point finite difference scheme for the helmholtz equation with pml.," *International Journal of Numerical Analysis & Modeling* **10**(2) (2013).

[13] Gander, M. J. and Zhang, H., "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**(1), 3–76 (2019).

[14] Stoffa, P., Fokkema, J. T., de Luna Freire, R., and Kessinger, W., "Split-step fourier migration," *Geophysics* **55**(4), 410–421 (1990).

[15] Gazdag, J. and Sguazzero, P., "Migration of seismic data by phase shift plus interpolation," *Geophysics* **49**(2), 124–131 (1984).

[16] Ali, R., Mitcham, T., and Duric, N., "Impact of starting model on waveform inversion in ultrasound tomography," in [*Medical Imaging 2023: Ultrasonic Imaging and Tomography*], **12470**, 89–98, SPIE (2023).

[17] Duric, N., Li, C., Roy, O., and Schmidt, S., "Acoustic tomography: promise versus reality," in [*2011 IEEE International Ultrasonics Symposium*], 2033–2041, IEEE (2011).