Multistage Iterative Method to Tackle Inverse Problems of Wave Tomography

Alexander V. Goncharsky[1,2], *Sergey Y. Romanov*[1,2], *Sergey Y. Seryozhnikov*[1,2]

This paper is concerned with developing the methods for solving inverse problems of low-frequency ultrasound tomography under scalar wave models using supercomputer technologies. Unlike X-ray tomography, the inverse problem considered is posed as a problem of minimizing a non-convex residual functional. The multistage iterative method (MSM) is proposed as a method for obtaining an approximate solution to the inverse problem. Convergence of the method to the exact solution is achieved via the use of low-frequency sounding signals at the initial stages of the iterative method. The method is illustrated on model problems focused on ultrasound tomographic diagnostics of soft tissues in medicine. Finite-difference time-domain method is used to solve the wave equation, which accounts for most of the computational complexity of the method. The multistage method reduces the computation time, since the initial stages use low-resolution finite difference grids. The effectiveness of the MSM method is investigated on GPU and SIMD-capable CPU computing platforms. Numerical simulations showed that modern processors equipped with AVX-512 FPUs are capable of solving small-scale problems of wave tomography. For large-scale tasks, GPUs equipped with fast on-board memory are preferred. The numerical algorithm is data-parallel and well-suited for GPU architecture. The proposed method can be used in medical imaging and nondestructive testing applications.

Keywords: ultrasound tomography, coefficient inverse problem, gradient method, numerical simulation.

Introduction

At present, it is hard to imagine a field of science or technology where tomographic imaging is not used. First tomographs that appeared in the middle of the latest century used X-ray radiation. However, the history of tomography can well begin at the beginning of the latest century, when X-rays have been discovered (Nobel Prize of the year). At the same time, Radon's solution of the inverse problem of reconstructing a function of two variables given its linear functionals has been published [1]. This result essentially solved the mathematical problems of tomography in a linear model. However, it took humanity half a century of intensive scientific research to develop first X-ray tomographs.

Currently, various tomographs (X-ray, MRI, positron emission tomography) are widely used in medicine, science and technology. All these technologies are united by the fact that, from a mathematical point of view, the problems of interpreting the data of tomographic experiments are inverse problems that can be solved within the framework of linear mathematical models. Solving such problems does not presently pose serious mathematical concerns. Personal computers are sufficient for data interpretation in a linear problem.

At the end of the latest century, it became possible to research a very interesting field of wave tomography, which employs ultrasonic, electromagnetic, seismic or optical radiation for sounding. For all these problems, it is necessary to use nonlinear mathematical models to interpret the data of tomographic experiments.

The progress in wave tomography developments has been driven by several factors. The first is the development of modern methods for solving inverse problems. First results obtained

---

[1]Lomonosov Moscow State University, Moscow, Russia
[2]Moscow Center of Fundamental and Applied Mathematics, Moscow, Russia

by Academician Tikhonov in the 60s [2, 3] were continued in the works of his students and followers [4–6]. By the end of the nineties, exhaustive results had been obtained in the field of solving ill-posed linear and nonlinear problems. Tikhonov brought the concept of a regularizing algorithm as a method for the approximate solution of the inverse problem. Within the framework of this concept, effective numerical methods have been developed for solving a wide range of problems in mathematical physics [7–11]. Important results were obtained in the field of using iterative schemes for the approximate solution of nonlinear inverse problems [5]. Intensive research was carried out in the field of solving coefficient inverse problems of mathematical physics [12, 13] Fast-growing supercomputer technology was another factor that contributed to the development of wave tomography. The solution of inverse problems under the wave model requires ample computational resources due to large problem dimensions and its nonlinearity. It is impossible to solve inverse problems of wave tomography without the use of supercomputers [14, 15].

In short, the situation with wave tomography at the moment can be characterized as follows. Most works on wave tomography consider scalar wave models that take into account the effects of diffraction, refraction, and even multiple scattering. The inverse problem in this case can be posed as a problem of minimizing the residual functional between the experimental data (measured wave field at the detectors) and the wave field computed using the mathematical model of wave propagation. The most important recent result in wave tomography is the ability to calculate the gradient of the residual functional explicitly [16–18]. This result makes it possible to use gradient-based methods for minimizing the residual functional to obtain an approximate solution to the inverse problem. Since the functional is not convex and has local minima, the problem of finding the global minimum of the functional arises. Despite the large number of works on this topic [19–21] this problem is unsolvable in a general case. In this paper, it is proposed to narrow the search area using additional information in order to find the global minimum of the residual functional.

Using additional information in solving inverse problems is not a new approach. In [4, 5] it was proposed to use such information about the sought-for functions as their monotonicity or convexity for constructing approximate solutions to ill-posed problems. In a finite-difference approximation, the problem reduces to minimizing a functional on a convex polyhedron with known vertices. Effective numerical algorithms have been developed for this approach.

To solve the problems of wave tomography, the authors propose a multistage iterative method (MSM) that uses additional prior information specific to these inverse problems. As a possible application of wave tomography, medical tomographic imaging for differential diagnosis of breast diseases is considered. Several groups of researchers are intensively working in this field [22–24]. These developments are currently at the stage of prototypes. At the moment, the main problem is constructing effective algorithms for data interpretation. This paper demonstrates that MSM can effectively find an approximate solution to inverse problems of tomographic image reconstruction in application to medical ultrasound imaging of soft tissues. The question of choosing the optimal computing platform for the proposed method is discussed.

Developing ultrasound tomography devices is a challenging task. A natural question arises: what advantages ultrasound tomography can provide in comparison with existing diagnostic methods. Let us try to give an answer to this question using the example of medical diagnostics, where ultrasound devices have been successfully used for a long time. To discuss the problem more specifically, we will narrow the field of medical diagnostics to soft tissue imaging. Unlike conventional ultrasound instruments, ultrasound tomographs can characterize the inspected tis-

sues. Just as in X-ray tomography, the doctor can obtain the value of the sound wave velocity at any point of the image. This result opens up the possibility of classifying neoplasms by the speed of sound in them. Neural networks can automate the tissue classification process [25]. Why cannot this be done using standard ultrasound diagnostic devices? Both X-ray and ultrasound tomography employ sounding waves transmitted through the object. Conventional ultrasound instruments detect only reflected waves, and this information is principally insufficient for tissue characterization. Finally, unlike X-ray tomography, ultrasound tomography is completely safe and therefore can be used for regular screening.

The article is organized as follows. Section 1 describes the inverse problem of wave tomography and the solution method. In Section 2 we introduce the proposed multistage iterative method. Section 3 describes the finite difference numerical method employed in the solution algorithm, Section 4 describes the parallel implementation of the solution algorithm. Section 5 compares the performance of the algorithm on various CPU and GPU computing systems. Section 6 presents model problems to demonstrate the proposed multistage method. Conclusion summarizes the study and points directions for further work.

1. Formulation of the Inverse Problem of Wave Tomography and its Solution Method

In this study, we consider the waves described by the scalar wave equation. In the scalar model, the scalar wave field $u(\boldsymbol{r}, \boldsymbol{q}, t)$, which represents the acoustic pressure, can be computed from the given initial data using the equation

$$c(\boldsymbol{r})u_{tt}(\boldsymbol{r}, \boldsymbol{q}, t) - \Delta u(\boldsymbol{r}, \boldsymbol{q}, t) = \delta(\boldsymbol{r} - \boldsymbol{q})g(t) \tag{1}$$

$$u(\boldsymbol{r}, \boldsymbol{q}, t = 0) = u_t(\boldsymbol{r}, \boldsymbol{q}, t = 0) = 0. \tag{2}$$

Here, $c^{-0.5}(\boldsymbol{r}) = v(r)$ is the speed of sound in the medium, $\boldsymbol{r} \in \mathbb{R}^2$, Δ is the Laplacian operator with respect to $\boldsymbol{r}$, δ is the Dirac delta function, which defines a point source at $\boldsymbol{q}$. The sounding pulse emitted by the source is described by function $g(t)$. Short broadband sounding pulses with a useable frequency range of 100–600 kHz and a duration up to 50 μs can be used for sounding in medical ultrasound tomography. The sounding pulses are further discussed in Section 6.
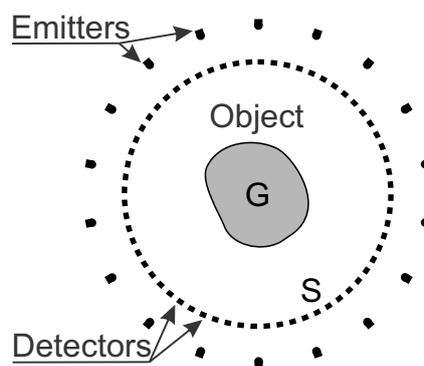

Figure 1. Ultrasound tomographic imaging scheme

The inverse problem of ultrasound tomography can be formulated as follows. Figure 1 shows the scheme of ultrasound tomographic imaging. The object with an unknown speed of sound $v(\boldsymbol{r})$ occupies region G. The object G, emitters and detectors of ultrasound are placed in a homogeneous medium with a known speed of sound $v_0 = c_0^{-0.5} = const$. The emitters are located at coordinates $\boldsymbol{q}_j$. A total of M emitter positions $j = 1, \ldots, M$ are located around region G. Measurements of the wave field $u(\boldsymbol{r}, \boldsymbol{q}, t)$ are taken on a circle S surrounding region G.

Function $g(t)$, which describes the sounding pulse, is known. In the inverse problem, the objective is to determine an unknown function $v(\boldsymbol{r})$ at $\boldsymbol{r} \in G$, using experimental data $U(s, \boldsymbol{q}_j, t)$ obtained at the boundary S ($\boldsymbol{s} \in S$) for emitter positions $\boldsymbol{q}_j, j = 1, \ldots, M$. Thus, the wave field $u(s, \boldsymbol{q}_j, t)$ from equations (1)–(2) satisfies the following equation for all emitter positions $\boldsymbol{q}_j$:

$$u(s, \boldsymbol{q}_j, t)|_{s \in S} = U(s, \boldsymbol{q}_j, t). \tag{3}$$

The system of equations (1)–(3) defines the inverse problem. Thus, solving inverse problem of ultrasound tomography in the scalar model involves reconstructing the unknown wave velocity $v(\boldsymbol{r})$ in region G according to equations (1)–(3).

The residual functional $\Phi(c)$ of the argument $c(\boldsymbol{r})$ is the difference between the experimental data and the data computed from equations (1)–(2). The residual functional for a computed wave field on the boundary S determined by a given speed of sound $c(\boldsymbol{r})$ can be written as

$$\Phi(c) = \sum_{j=1}^{M} \frac{1}{2} \int_0^T \int_S \left(u(s, \boldsymbol{q}_j, t) - U(s, \boldsymbol{q}_j, t) \right)^2 \, \mathrm{d}s \, \mathrm{d}t. \tag{4}$$

Here $U(s, \boldsymbol{q}_j, t)$ are the experimental data on the boundary S for the time interval $(0, T)$, and $u(s, \boldsymbol{q}_j, t)$ is the wave field obtained via solving the direct problem (1)–(2), which depends on the specified coefficient $c(\boldsymbol{r})$. For multiple sounding wave sources, the residual functional is the sum over $j = 1, \ldots, M$ of the residual values obtained for each source. For each fixed source j, the integral is summed over time $(0, T)$ and over the boundary S – for all the detectors receiving the signal from the selected source. Mathematically, the inverse problem is posed as a problem of finding a function $\bar{c}(\boldsymbol{r})$ that minimizes the residual functional (4) $\bar{c}(\boldsymbol{r}) : \min_{c(\boldsymbol{r})} \Phi(c) = \Phi(\bar{c})$ The $\bar{c}(\boldsymbol{r})$ function is taken as an approximate solution to the inverse problem. Gradient methods have proven effective for minimizing the residual functional $\Phi(c)$. A rigorous mathematical formulation for the gradient of the residual functional has been obtained in [17, 18, 26]. The gradient of the functional (4) has the form

$$\Phi'(c) = \sum_{j=1}^{M} \int_0^T \frac{1}{2} w_t(s, \boldsymbol{q}_j, t) u_t(s, \boldsymbol{q}_j, t) \, \mathrm{d}t, \tag{5}$$

where $u(s, \boldsymbol{q}_j, t)$ is the solution of the main problem (1)–(2) and $w(s, \boldsymbol{q}_j, t)$ is the solution of the conjugate problem (6)–(7). Both solutions depend on $c(\boldsymbol{r})$ coefficients [14, 18].

$$c(\boldsymbol{r}) w_t t(\boldsymbol{r}, \boldsymbol{q}_j, t) - \Delta w(\boldsymbol{r}, \boldsymbol{q}_j, t) = u(s, \boldsymbol{q}_j, t) - U(s, \boldsymbol{q}_j, t)|_{s \in S}, \tag{6}$$

$$w(\boldsymbol{r}, \boldsymbol{q}_j, t = T) = w_t(\boldsymbol{r}, \boldsymbol{q}_j, t = T) = 0. \tag{7}$$

The inverse problem of wave tomography in the considered formulation is a nonlinear coefficient inverse problem. In nonlinear problems, typically, the residual functional (4) is not convex,

which means that the functional may have local minima. As a consequence, gradient methods for minimizing the residual functional from an arbitrary initial approximation may converge to a local minimum, but not to the global one.

2. The Main Idea of the Multistage Method for Obtaining Approximate Solutions to Nonlinear Inverse Problems of Ultrasound Tomography

As shown in Section 1, inverse problem of ultrasound tomography can be solved by minimizing the residual functional (4), which may have local minima. There are many works concerned with finding global minima of functionals. However, this problem has no solution in a general case. In problems of wave tomography, an important prior information is present, which is that the area of convergence of iterative processes of minimizing the residual functional strongly depends on the wavelength of the sounding radiation. If the central frequency of the sounding pulses tends to 0, then passing to the limit in equation (1) reduces it to a linear integral equation with respect to the unknown function $c(\boldsymbol{r})$. This idea of using linear models in problems of wave diagnostics is actively discussed in [7, 27–30].

It would seem that this result opens up wide possibilities for solving inverse problems of nonlinear wave tomography, but this is actually not the case for several reasons. The first reason is that in real ultrasound imaging, the center frequency of the sounding wave determines the resolution, and generally all medical ultrasound devices have center frequencies of 1 MHz or higher. It is not possible to obtain experimental data in the frequency range close to zero.

The second reason is that the experimental data are measured with some error. In order to use the linear approximation, it is necessary to calculate the second derivative of measured waveforms, which is poorly conditioned. Nevertheless, the idea of using low frequencies in problems of ultrasound diagnostics is fruitful, and this idea is used in the proposed multistage iterative method (MSM) for obtaining an approximate solution to the inverse problem of ultrasound tomography.

In this article, the capabilities of the MSM method are illustrated using the problems of ultrasound tomographic imaging of soft tissues in medicine, namely, breast imaging. The problem of early-stage breast cancer diagnosis is one of the most important problems of modern medicine. A characteristic feature of soft tissue imaging is that the difference between the speed of sound in soft tissues and the speed of sound in surrounding water v_0 is quite low and does not exceed 15%. It seems natural to use the known speed of sound v_0 in a homogeneous medium surrounding an object in region G as an initial approximation in iterative processes. However, as it will be shown via numerical simulations, the choice of an initial approximation in the form of a constant v_0 in nonlinear problems does not guarantee convergence to the global minimum of the residual functional.

Multistage iterative method (MSM) proposed in this study for solving coefficient inverse problems of ultrasound tomography ensures convergence of the gradient descent algorithm to the global minimum. By increasing the mean wavelength of the sounding signal we expand the range of initial approximations from which the gradient descent method of minimizing the residual functional converges to the global minimum. This idea is at the heart of the MSM method, and it is demonstrated in this study on a large number of model problems that simulate the problem of

ultrasound tomographic imaging for breast cancer diagnosis for various configurations of objects and sounding pulses.

In a practical implementation of the MSM method, experimental data for two (or more) frequency bands with different central frequencies f_1 and f_2, $f_1 < f_2$ are used. First, the inverse problem is solved via an iterative gradient descent method for minimizing the residual functional using the lowest frequency band f_1. The initial approximation is chosen as a constant equal to the speed of sound v_0 in the environment. Then, at the second stage, the inverse problem is solved via an iterative method using a higher frequency band f_2. The result of solving the inverse problem at the first stage at a lower frequency f_1 is used as an initial approximation of the speed of sound for the gradient descent method at the second stage.

The described scheme of MSM method application consists of two stages. In a real situation, it may be necessary to apply several successive stages for three or more frequency bands. The number of stages depends on the task of ultrasound tomography considered. It is important that at each stage the initial approximation is close enough to the point of the global minimum of the residual functional for the frequency band used at that stage. Numerous model calculations have shown that a five-stage method is optimal for tomographic imaging of soft tissues with low-frequency ultrasound in the $100 - 600$ kHz band. To ensure the convergence of the MSM method to the exact solution of the inverse problem, the first stage should be carried out using a wavelength $\lambda = 12$ mm (the central frequency of the pulse is 125 kHz). Subsequent stages were carried out using wavelengths of 8, 6, 4 mm, and the last stage uses $\lambda = 3$ mm (500 kHz). The wavelength of 3 mm provides a high spatial resolution of the method – 1–1.5 mm in application to soft tissue imaging.

The following model problem illustrates the capabilities of the MSM method. The parameters of model problems correspond to breast ultrasound tomography. The variation of the speed of sound in model problems does not exceed 15%. In order to ensure the convergence of the iterative process, the developed methods of ultrasound diagnostics use significantly lower sounding frequencies compared to conventional medical ultrasound devices.

The following parameter values were used in the numerical simulations. The speed of sound in the medium surrounding the object is $v_0 = 1.5$ mm/μs (water), the speed of sound in the object varies from 1.4 mm/μs to 1.7 mm/μs. The calculations were carried out for five frequency bands with central frequencies of 125, 188, 250, 375 and 500 kHz (mean wavelengths of sounding pulses equal to 12, 8, 6, 4 and 3 mm, respectively). The size of the two-dimensional computational domain was 25×25 cm, the size of the finite difference grid for the 500 kHz band was 1600×1600 points, for the 125 kHz band – 480×480 points. In the calculations, the sources and receivers were located around the object. There were 24 source positions in total, located on a circle 200 mm in diameter. The receivers were located with a step of ≈ 1 mm on a circle 165 mm in diameter.

Figure 2 presents the results of solving an inverse problem. Figure 2a shows the original image of the object (simulated phantom). Figure 2b plots the residual functional (4) for three frequency bands with mean wavelengths of $\lambda = 3$ mm, 6 mm and 12 mm. The abscissa corresponds to the parameter α in the interval (–0.1, 1.4), which determines the value of the function $c(\boldsymbol{r})$ by the following formula $c(\boldsymbol{r}; \alpha) = (1 - \alpha)c_0 + \alpha \bar{c}(\boldsymbol{r})$. Here, $c^{-0.5}(\boldsymbol{r}) = v(\boldsymbol{r})$ is the wave velocity in the medium, $c_0^{-0.5} = v_0 = const$. For $\alpha = 1$, we get $c(\boldsymbol{r}) = \bar{c}(\boldsymbol{r})$, which is the exact solution and the residual functional equals to 0. For $\alpha = 0$, $c(\boldsymbol{r}) = c_0$, which corresponds to the initial approximation of the iterative process. For short wavelengths $\lambda = 3$ mm and 6 mm, the value of the residual functional in the interval $0 < \alpha < 1$ first increases with α, and only then decreases to 0.
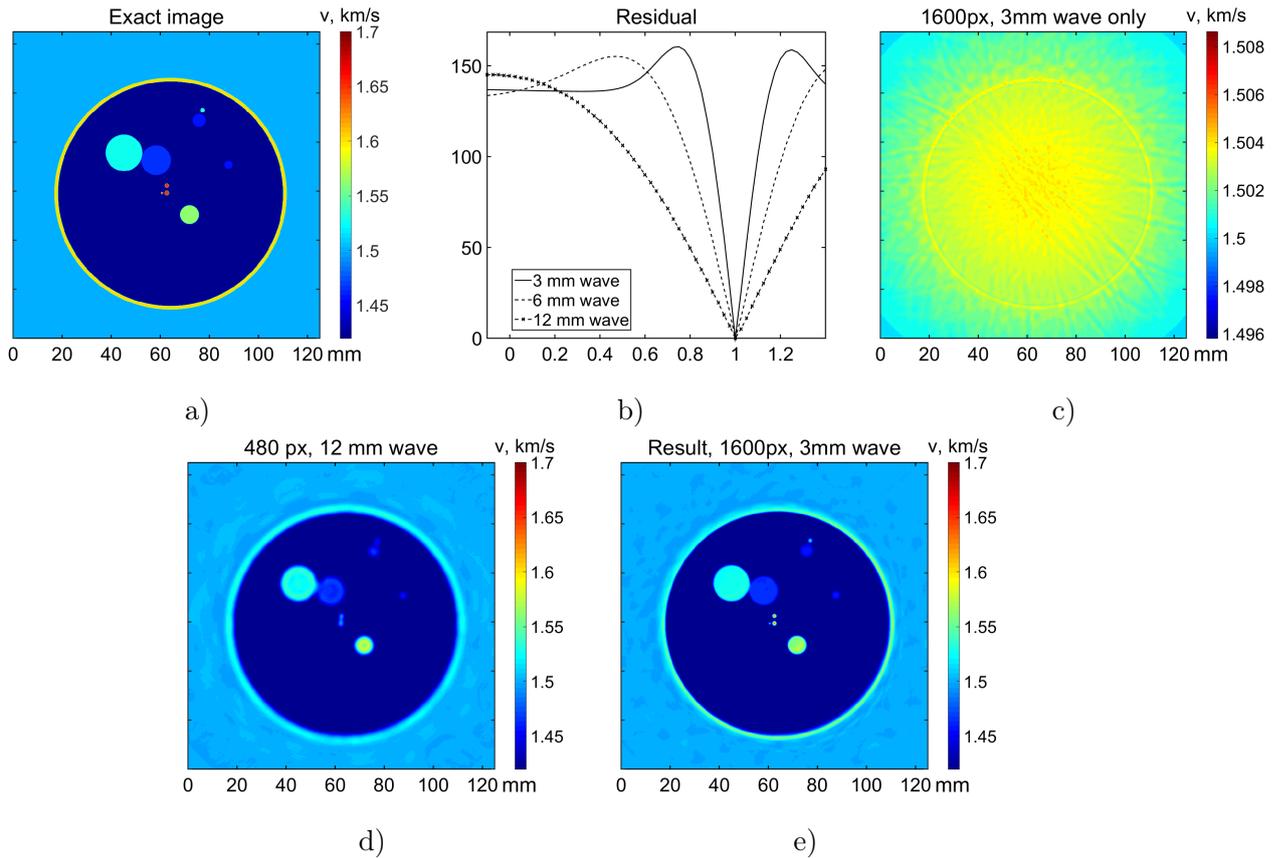
Figure 2. Numerical simulation: a – exact image (phantom); b – plots of the residual functional (4) for wavelengths $\lambda = 3$ mm, 6 mm and 12 mm; c – an image reconstructed using a wavelength $\lambda = 3$ mm from an initial approximation $c_0 = const$, d – an image reconstructed using a wavelength $\lambda = 12$ mm from an initial approximation $c_0 = const$, e – an image reconstructed via the multistage method

This illustrates the idea that the iterative gradient descent process for these wavelengths stops at a local minimum of the residual functional if the iterative process is started from an initial approximation of $c(\boldsymbol{r}) = c_0$. For the wavelength $\lambda = 12$ mm, the residual functional decreases monotonically to 0 as the parameter α changes from 0 to 1. This result allows us to assume that the initial approximation $c(\boldsymbol{r}) = c_0$ lies in the vicinity of the global minimum of the residual functional if the mean wavelength equals to 12 mm or longer. The global minimum is reachable from the initial approximation via the gradient descent method in this case.

To ensure the convergence of the iterative gradient descent process, a five-stage method has been used. For $\lambda = 12$ mm, the iterative process was started from an initial approximation of $c(\boldsymbol{r}) = c_0$. The resulting approximate solution was used as an initial approximation for the second stage with a mean wavelength $\lambda = 8$ mm, and so on. The result of the iterative process for $\lambda = 4$ mm was used as an initial approximation for the last stage with $\lambda = 3$ mm. The difference between wavelengths should be sufficiently small for the global minimum to be reachable from the initial approximation via the gradient descent method at each stage.

Figure 2c shows the image reconstruction results for a sounding pulse wavelength of $\lambda = 3$ mm, where $c_0 = const$ is chosen as the initial approximation for the iterative process. In this case, the iterative process stops at a local minimum, and the resulting solution is very different from the original image. Figure 2d shows the image reconstruction results for a sounding

pulse wavelength of $\lambda = 12$ mm; $c_0 = const$ is chosen as the initial approximation. In this case, the initial approximation lies in the vicinity of the global minimum, and an approximate solution to the inverse problem is obtained. However, due to the large wavelength, the spatial resolution of the resulting image is rather low.

Figure 2e shows the image reconstruction results using the multistage method with 5 stages. The central wavelength at the last stage is $\lambda = 3$ mm. The image obtained at the first stage for $\lambda = 12$ mm (shown in Fig. 2d) was used as an initial approximation for the iterative process using a wavelength $\lambda = 8$ mm at the second stage, and so on for wavelengths of 6 mm, 4 mm and 3 mm. The multistage method made it possible to avoid the iterative process stopping at a local minimum and to obtain the resulting high-quality image.

It turns out that the MSM method not only provides convergence to an approximate solution to the problem, but also significantly reduces the computation time. This issue will be discussed in more detail in Section 6.

3. Numerical Approximation of the Wave Equation

Finite-difference time-domain method (FDTD) was employed to solve equations $(1)-(2)$. We define a uniform rectangular finite difference grid: $x_i = ih, y_j = jh, t_k = k\tau; i, j = 1, ..., N, k = 1, ..., M$, where h is the spatial discretization step, and τ is the time step. A second-order finite difference scheme approximates equation (1):

$$c_{ij} \frac{u_{ij}^{k+1} - 2u_{ij}^k + u_{ij}^{k-1}}{\tau^2} - \frac{\boldsymbol{L}_{ij}^k}{h^2} = 0.$$

Here, $u_{ij}^k = u(x_i, y_j, t_k)$ are the values of $u(\boldsymbol{r}, \boldsymbol{q}, t)$ at point (i, j) at the time step k for a fixed $\boldsymbol{q}$; c_{ij} and a_{ij} are the values of $c(\boldsymbol{r})$ and $a(\boldsymbol{r})$ at point (i, j). The first term approximates $c(\boldsymbol{r})u_{tt}(\boldsymbol{r}, \boldsymbol{q}, t)$, the second term approximates $a(\boldsymbol{r})u_t(\boldsymbol{r}, \boldsymbol{q}, t)$. The discrete Laplacian is denoted by $\boldsymbol{L}_{ij}^k$. A fourth-order numerical approximation [31] on a 5×5-point stencil is used for the discrete Laplacian:

$$\boldsymbol{L}_{ij}^k = \sum_{m=i-2}^{i+2} \sum_{n=j-2}^{j+2} v_{mn} u_{mn}^k. \tag{8}$$

The parameters h and τ are related by the Courant stability condition $c^{-0.5}\tau < h/\sqrt{2}$. For the problem considered, we used a time step equal to $\tau = 0.3c_0^{0.5}h$, which ensured the stability of the finite difference method. The number of operations required to compute a wave propagation simulation is proportional to $O(N^3)$, where N is the number of grid points along spatial dimensions. The number of points N is chosen so that the precision of the wave simulation for the selected wavelength range is sufficient. Thus, computational complexity of the numerical method scales as a third power of wave frequency and spatial image resolution.

An approximate solution to the inverse problem is obtained via an iterative gradient descent method. Each iteration involves solving direct (1)–(2) and conjugate (6)–(7) problems in order to compute the gradient of the residual functional, which requires simulating the wave propagation process in forward and reverse time.

The numerical method was implemented in software for GPU and SIMD-capable CPU computing platforms. The discrete Laplacian computation (8) is the most compute-intensive operation in this method. The flowchart of the SIMD algorithm for computing the discrete Laplacian

is shown in Fig. 3. The Laplacian is spherically symmetrical, which makes it less computationally expensive than a convolution problem in a general case.
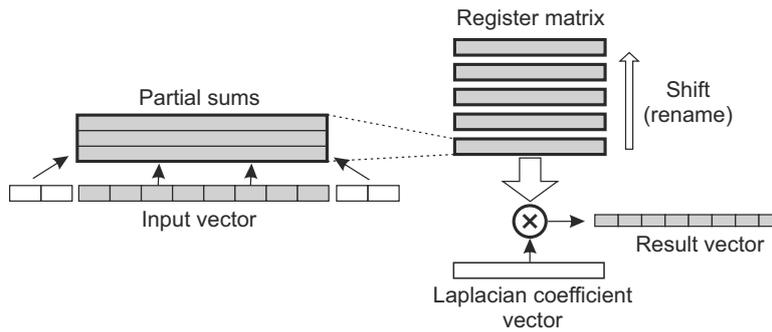

Figure 3. SIMD Discrete Laplacian computation algorithm

Y-marching method was employed to compute the convolution using registers for temporary storage. The results are calculated sequentially in vertical direction. Using the input data vector and horizontally adjacent cells, partial sums are computed and stored in the register matrix, which contains the data for 5 lines of the image. The result vector is computed by multiplying the register matrix by the Laplacian coefficient vector. The algorithm advances to the next line by shifting the lines in the register matrix up and reloading the last line from the input vector. The data is shifted via renaming the registers.

Modern processors (AVX, AVX-512, ARM NEON-class FPUs) typically have 32 SIMD registers, each of which holds a vector of 8 32-bit floating point elements for AVX FPU, 16 elements for AVX-512 FPU and 4 elements for ARM NEON FPU. There are three partial sums per line and five lines in the register matrix; thus, the input vector can be two registers long for the single wave simulation ($u(\boldsymbol{r}, \boldsymbol{q}, t)$ for the direct problem (1)–(2)) and one register long for the dual wave simulation ($u(\boldsymbol{r}, \boldsymbol{q}, t)$ and $w(\boldsymbol{r}, \boldsymbol{q}, t)$ for the conjugate problem (6)–(7)).

The computations on multi-core CPUs were parallelized using OpenMP. MPI interface was used for data exchange between computing nodes (CPU sockets or GPU devices). Figures 4a, b illustrate the order of computations for multi-core CPUs and GPUs, respectively.


a)   b)
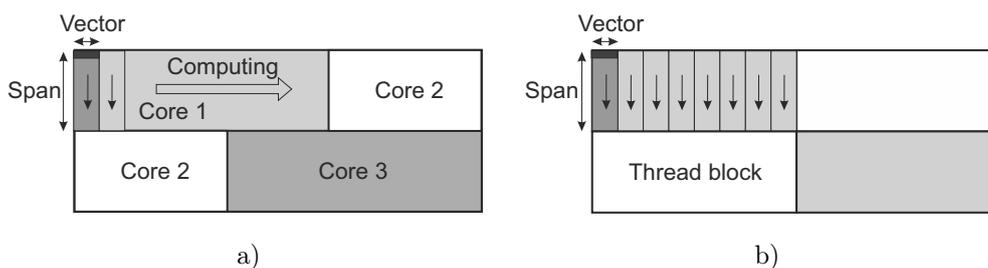
Figure 4. Parallelizing the computations on multi-core processors (a) and GPU (b)

For efficient use of cache memory, the length of the vertical segment ("Span") of the Y-marching method was limited to some specified value. The span length typically ranges from 5 to 40 pixels. On multi-core CPUs, individual spans are computed in sequence along the horizontal dimension. The data of the previous span remains in the cache memory and is used to compute a part of the next span. The optimal span length depends on the image size and CPU cache properties and can be determined for each system via performance tests. The better the data fits into the cache, the larger span lengths are preferred. An equal amount of data is distributed to

each computing core. For GPU, the computations are performed in parallel within each thread block. The thread block size can be adjusted for better performance.

4. Parallel Implementation of the Inverse Problem Solution Algorithm

Figure 5a illustrates the direct problem solution algorithm. The algorithm simulates the wave field propagating through an inhomogeneous medium. A predefined numerical phantom simulates the object being imaged. The phantom specifies the speed of sound $c(\boldsymbol{r})$ in the imaging plane. Ultrasound emitters and detectors are placed in a circular formation around the phantom, as shown in Fig. 1.
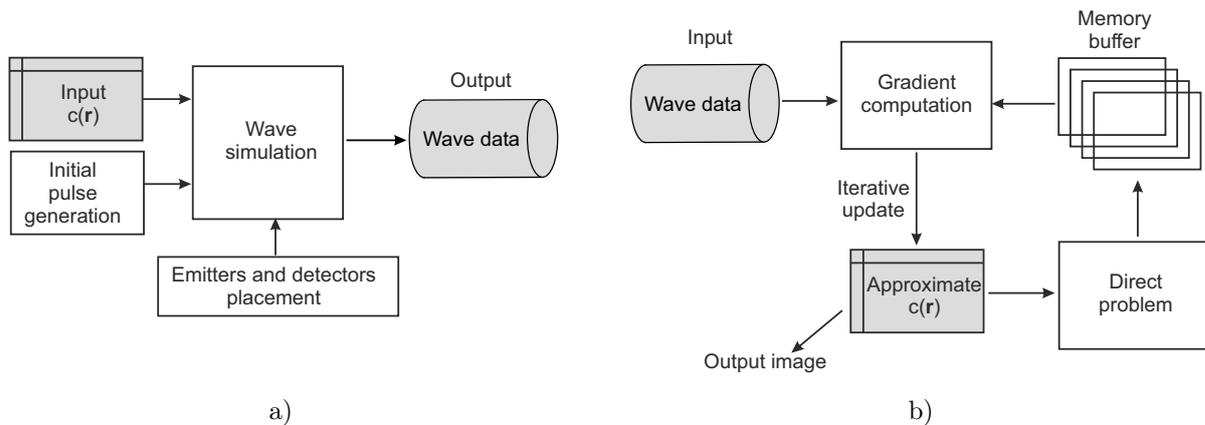
Figure 5. Direct (a) and inverse (b) problem solution algorithms

The wave field is simulated sequentially in time, starting from the initial pulse that is computed as a spherical wave radiating from the emitter position. The resulting wave field at the detector positions for each emitter is recorded in the output data.

Figure 5b illustrates the inverse problem solution algorithm. An approximate solution to the inverse problem of wave tomography is computed via the iterative gradient descent method. An initial approximation of $c(\boldsymbol{r})$ coefficients is set at the beginning of the iterative process. The direct problem is solved for the current approximation of $c(\boldsymbol{r})$.

The boundary values of computed $u(\boldsymbol{r}, \boldsymbol{q}, t)$ wave field are stored in the memory buffer. The buffer is used to reverse the wave propagation direction of $u(\boldsymbol{r}, \boldsymbol{q}, t)$ in formula (5). At the gradient computation stage (Fig. 5b) the data from the buffer are applied to the boundary of the computational domain in reverse time in order to compute $u(\boldsymbol{r}, \boldsymbol{q}, t)$ in reverse time simultaneously with $w(\boldsymbol{r}, \boldsymbol{q}, t)$.

The $w(\boldsymbol{r}, \boldsymbol{q}, t)$ wave is computed from $u(\boldsymbol{r}, \boldsymbol{q}, t)$ and the input wave data, and the gradient is computed using formula (5). The current approximate solution is updated by adding the computed gradient to the $c(\boldsymbol{r})$ coefficient array, and the process is repeated. The iterative process continues until the residual functional ceases to decrease. At the end of the process, the resulting approximate solution is the output of the inverse problem solution algorithm.

The iterative gradient descent method permits parallelizing the computations contained within a single iteration.

Figure 6a illustrates the parallelized computing process. Computation of the gradient of the residual functional can be subdivided into independent sub-tasks for each ultrasound emitter.

a)   b)
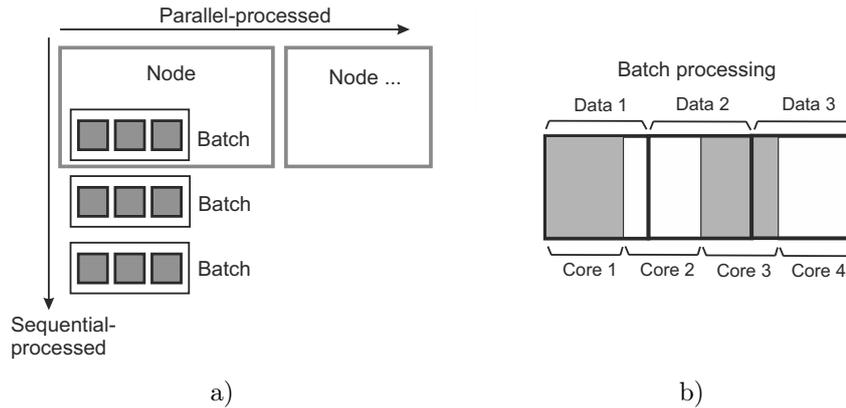
Figure 6. Parallelizing the computations on multi-core processors (a) and GPU (b)

The total number of emitters in wave tomography typically ranges from 10 to 100. The emitters are divided evenly between the computing nodes.

For each node, the computations are grouped into one or more batches executed sequentially. A batch consists of the data for several ultrasound emitters that are processed in parallel. For GPUs, the batch size is determined by the GPU memory capacity. A typical modern GPU can process most ultrasound tomography problems within its on-board memory as a single batch. For CPUs, the batch size is optimized for maximum performance, and typically is chosen close to the CPU last-level cache size. The data for a single emitter can amount to several megabytes; thus, small batch sizes of 1–2 emitters are common for CPUs.

Figure 6b illustrates the order of computations within a batch. The computations performed for each ultrasound emitter are identical and consist of a direct problem solution (wave simulation) and a gradient computation that involves simulating two wave fields $u(\boldsymbol{r}, \boldsymbol{q}, t)$ and $w(\boldsymbol{r}, \boldsymbol{q}, t)$. The only difference between emitters is in the data contents. Thus, wave fields for multiple emitters can be processed as a single data array in order to divide the computations evenly between computing cores.

The result of each iteration is the gradient of the residual functional, which is the sum of partial gradients computed for each ultrasound emitter. The partial gradients are summed up within each node, and then summed up over nodes using MPI interface. Data exchanges between nodes occur only once per iteration and therefore do not incur any noticeable delay.

5. Computing Performance for Different Computing Platforms

The software implementation of the algorithm has been tested on multiple computing platforms: Intel Haswell-EP (14 cores, AVX2 FPU), Intel 6240R (12 cores, AVX-512 FPU), NVidia Tesla P100 and NVidia Tesla V100 GPUs.

The multistage method involves solving inverse problems of wave tomography using data with bandwidth gradually increasing from stage to stage. Since the computation time strongly depends on the finite difference grid size, at each stage of the multistage method the grid size is chosen as the smallest size that still provides sufficient accuracy of the finite difference scheme. The computation time for a multistage task is the sum of the time intervals spent on each stage. To estimate the computation time for multistage tasks, the performance of computing platforms was tested on the computations of separate iterations of the gradient descent method for various grid sizes.

CPU performance significantly depends on cache utilization, as external memory is much slower than cache memory. Thus, for optimal performance we choose the batch size (the amount of data to be processed in parallel by the CPU, Fig. 6a) close to the CPU cache size. In order to optimize the computations, performance tests were conducted to determine the optimal batch size for parallel processing on each target system. The batch sizes determined may differ form the physical CPU cache size due to the use of an additional memory buffer (Fig. 5b).
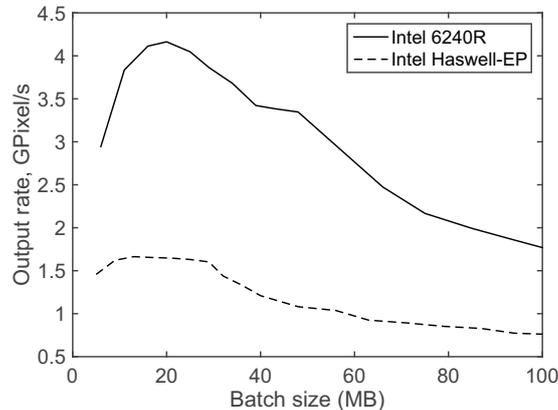

Figure 7. CPU performance depending on batch size

Figure 7 plots CPU performance depending on batch size. Output data rate is measured in the tests as the number of computed gradient pixels per second in gigapixels/s. Each pixel of a reconstructed image uses 32 bytes of data. The tests showed that using too large batch sizes results in a performance decrease due to cache misses. Too small batch sizes are also non-optimal because the task execution time for a small batch becomes very short and the thread synchronization latency becomes noticeable, especially on a fast CPU such as Intel 6240R. As a result, a batch size of 22 Mb was chosen for computations on Intel CPUs.

A single simulation frame uses $N \times N$ pixels of output, where N is the grid size along each dimension. Wave simulation requires approximately N time steps for the computed wave to reach the detectors. The wave simulation is performed for every ultrasound emitter in order to compute the gradient of the residual functional. The multistage method uses multiple grids with different resolutions. To determine the computing time for a task, we determine the computing time for a single wave simulation and multiply that by the number of emitters and the number of gradient descent iterations.

A series of tests were conducted to determine the computing time for each grid resolution. Figure 8a plots the computing devices' performance depending on the finite difference grid size. For Intel CPUs, the performance decreases for larger grids, as such grids do not fit in the CPU cache. For GPUs the performance may slightly increase with increased grid size due to more parallelism being available on larger grids. The NVidia Tesla P100 and NVidia Tesla V100 GPUs tested are equipped with sufficient amount of VRAM to process the whole inverse problem as a single batch in parallel.

Figure 8b plots the computing time per iteration per emitter corresponding to the output rate shown in Fig. 8a. The time scale is logarithmic in this plot. The number of operations to compute a wave simulation scales as a third power of the grid size. Employing smaller grids to complete the first iterations quickly in the multistage iterative method significantly decreases

a)   b)
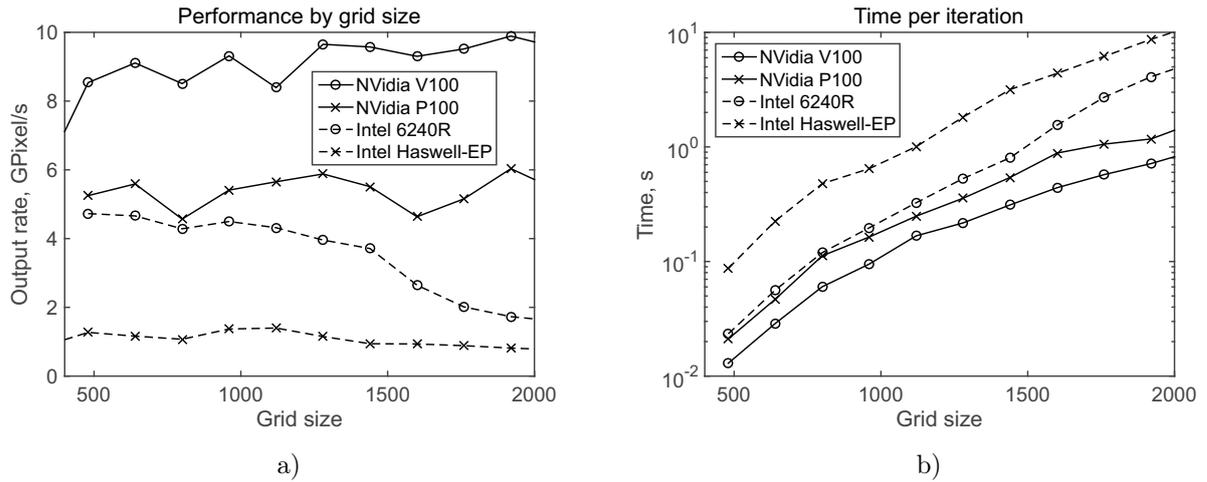
Figure 8. Device performance depending on finite difference grid size (a); corresponding computing time per iteration per emitter (b)

the total computing time compared to performing all the iterations using the highest-resolution grid.

Table 1. Multistage method parameters

Stage	1	2	3	4	5
Resolution, px	480	640	800	1200	1600
Wavelength, mm	12	8	6	4	3
Iterations	30	30	30	30	30
Time, sec	3	7	15	50	120

Numerical simulations showed that on average 30 iterations are sufficient at each stage of the multistage method to obtain an approximate solution suitable for the next stage, or to obtain the final result at the last stage. Multistage method parameters suitable for medical imaging for breast cancer diagnosis are summarized in Tab. 1. The computing time row in Tab. 1 lists the actual times achieved on a computing node of "Lomonosov-2" supercomputer [32] equipped with two NVidia Tesla V100 GPUs working in parallel.
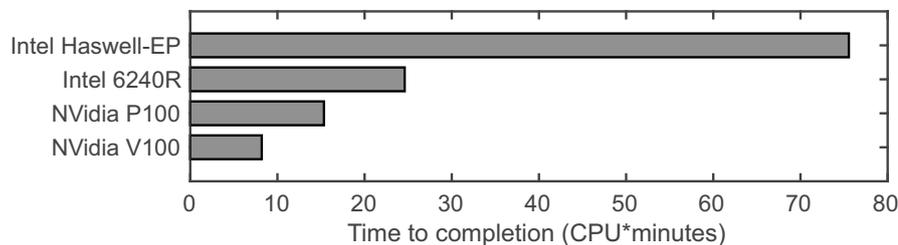
Figure 9. Overall performance of the computing devices

Using these parameters, the total computing time required to complete the inverse problem solution can be estimated for each computing device. Figure 9 shows the computing time for the devices tested. For example, to obtain 30 images per hour in a medical imaging setup (one image is a single cross-section of an object) a computing cluster of either 4 NVidia Tesla V100, 8 NVidia Tesla P100, 12 Intel 6240R or 36 Intel Haswell-EP processors would be needed.

Provided that almost all the data is cached, Intel 6240R with AVX-512 FPU is only 1.5 times behind NVidia Tesla P100 GPU. Thus, CPU or GPU clusters can be used for image reconstruction using the multistage method with image sizes up to 1600 pixels. Such grid sizes are sufficient for low-frequency wave tomography of relatively small objects, where the object size is on the order of 25–30 wavelengths. For higher frequencies or larger objects, larger grids would be required, putting CPU systems at a significant disadvantage compared to GPUs.

GPUs were found to be the preferred architecture for solving direct and inverse problems of wave tomography, especially for higher-resolution images. GPU performance on a typical FDTD algorithm is approximately proportional to its memory throughput. The numerical algorithm is data-parallel and requires neither synchronized data exchanges between computing cores nor cache coherence. Thus, the algorithm can benefit from the specific structure of graphics processors.

6. Model Problem Examples

The multistage iterative method for solving inverse problems of wave tomography proposed in this study is designed primarily to ensure convergence of the gradient iterative method to the global minimum. Approximate solutions of the inverse problem are computed using gradually increasing sounding pulse bandwidth and image resolution.

It may seem that an experimental setup with 5 different ultrasound emitters is needed to use a 5-stage iterative method in practice. However, the input data for multiple stages can be produced from a single experimental measurement via application of a low-pass filter to the broadband signal. This approach requires the signal to contain sufficiently strong low frequencies. Broadband ultrasound transducers with a usable frequency range of 100 to 600 kHz can be used for this task.

a)   b)   c)
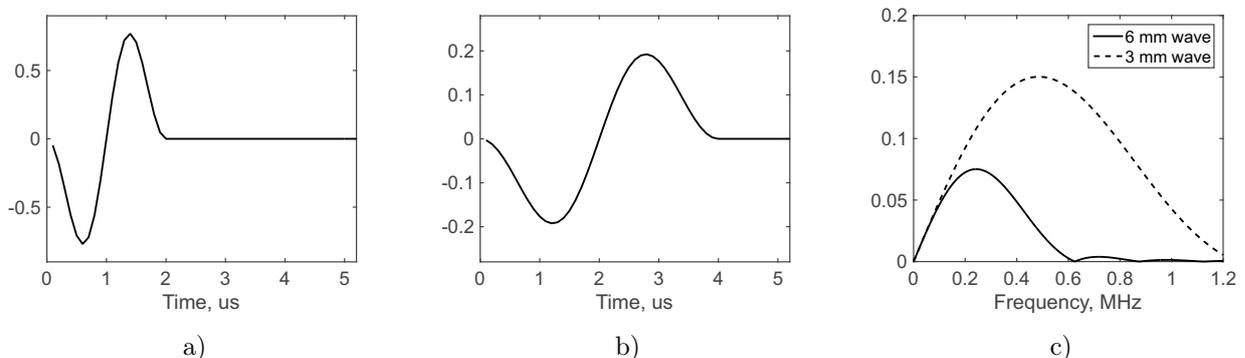
Figure 10. Sounding pulses: a – waveform with a mean wavelength of 3 mm; b – waveform with a mean wavelength of 6 mm; c – frequency spectra of the sounding pulses

Figure 10a shows the waveform of a broadband sounding pulse with a mean wavelength of 3 mm used in the presented numerical simulations at the last stage of the MSM method. Figure 10b shows the sounding pulse with a mean wavelength of 6 mm used for an intermediate stage of the MSM method. Figure 10c shows the frequency spectra of these two sounding pulses. The spectrum of the longer wave with a central frequency of 250 kHz is a part of the spectrum of the shorter wave with a central frequency of 500 kHz. A short pulse contains both low and high frequencies, lower parts of the spectrum can be filtered and used for low-resolution stages.

Lower-resolution approximations can be computed much faster; thus, the multistage iterative method can be used to improve computing time even if multiple stages are not necessary for convergence. Figure 11 presents an example of a low-contrast phantom, the image of which can be reconstructed via the gradient descent method using a wavelength of 3 mm and a constant initial approximation.
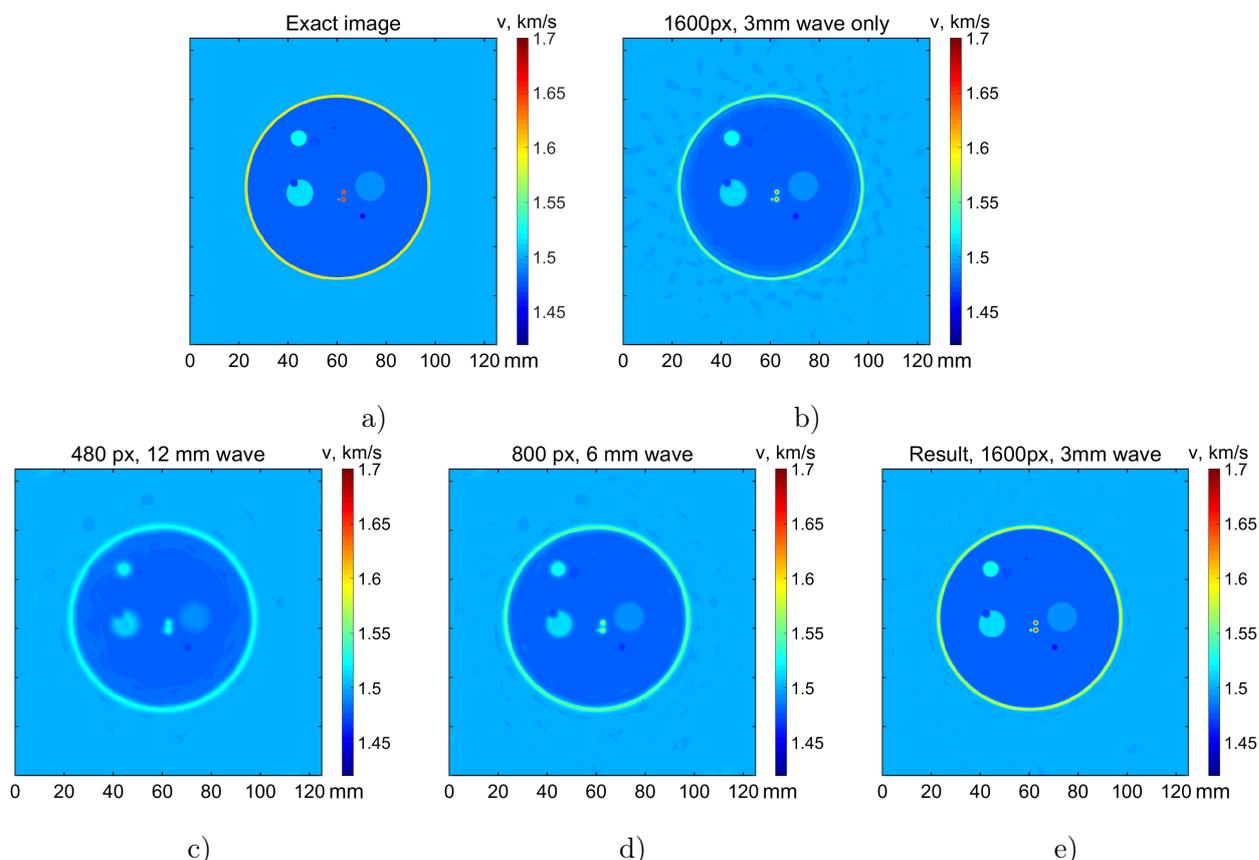
Figure 11. Single-stage and multistage reconstruction: a – exact image, b – image reconstructed at 3 mm wavelength from a constant initial approximation, c – image reconstructed at the first stage from a constant initial approximation, d – image reconstructed at the 3rd stage, e – image reconstructed using at the final stage

Figure 11a shows an exact image of the phantom. Figure 11b shows an image reconstructed using the gradient method from a constant initial approximation and with a pulse wavelength of 3 mm. Even without using the multistage method, the image is reconstructed quite accurately. However, to solve this inverse problem, 100–120 gradient descent iterations are required.

Figures 11c–e show approximate solutions obtained at various stages of the multistage method. Figure 11c shows an image reconstructed at the first stage of the multistage method from a constant initial approximation using a wavelength of 12 mm and a grid size of 480×480 points. Figure 11d shows an image reconstructed at the third stage of the multistage method using a wavelength of 6 mm and a grid size of 800×800 points. Figure 11e shows an image reconstructed at the last stage of the multistage method using a wavelength of 3 mm and a grid size of 1600×1600 points.

Figure 11 demonstrates that the image Fig. 11e obtained via the multistage method is closer to the original than the image in Fig. 11b obtained via the gradient descent method with fixed parameters and a constant initial approximation. Although the parameters of the iterative

method are the same for Fig. 11b and Fig. 11e, a better image quality in Fig. 11e is achieved due to a better initial approximation being used in the last stage of the MSM method. Instead of a constant, an image computed at the previous stage is used, which is much closer to the exact image than a constant.

In this example, the multistage method can reduce the computation time and improve the image quality. The parameters of the multistage method are listed in Tab. 1. The total number of gradient descent iterations at all stages is 150, however, only 30 of them are performed in a high resolution of 1600×1600. These iterations are the most time-consuming. Starting from an initial approximation computed at the previous stage instead of a constant initial approximation, only 30 iterations at the highest resolution are sufficient to complete the process.

First stages of the multistage method require significantly less computation time than the last. The total computation time in this example is 195 seconds, of which the last stage takes 120 s. In 195 seconds, it is possible to perform 48 high-resolution gradient descent iterations, which is usually not enough to reconstruct an image with good accuracy. Thus, the multistage method allows us to reduce the computation time.
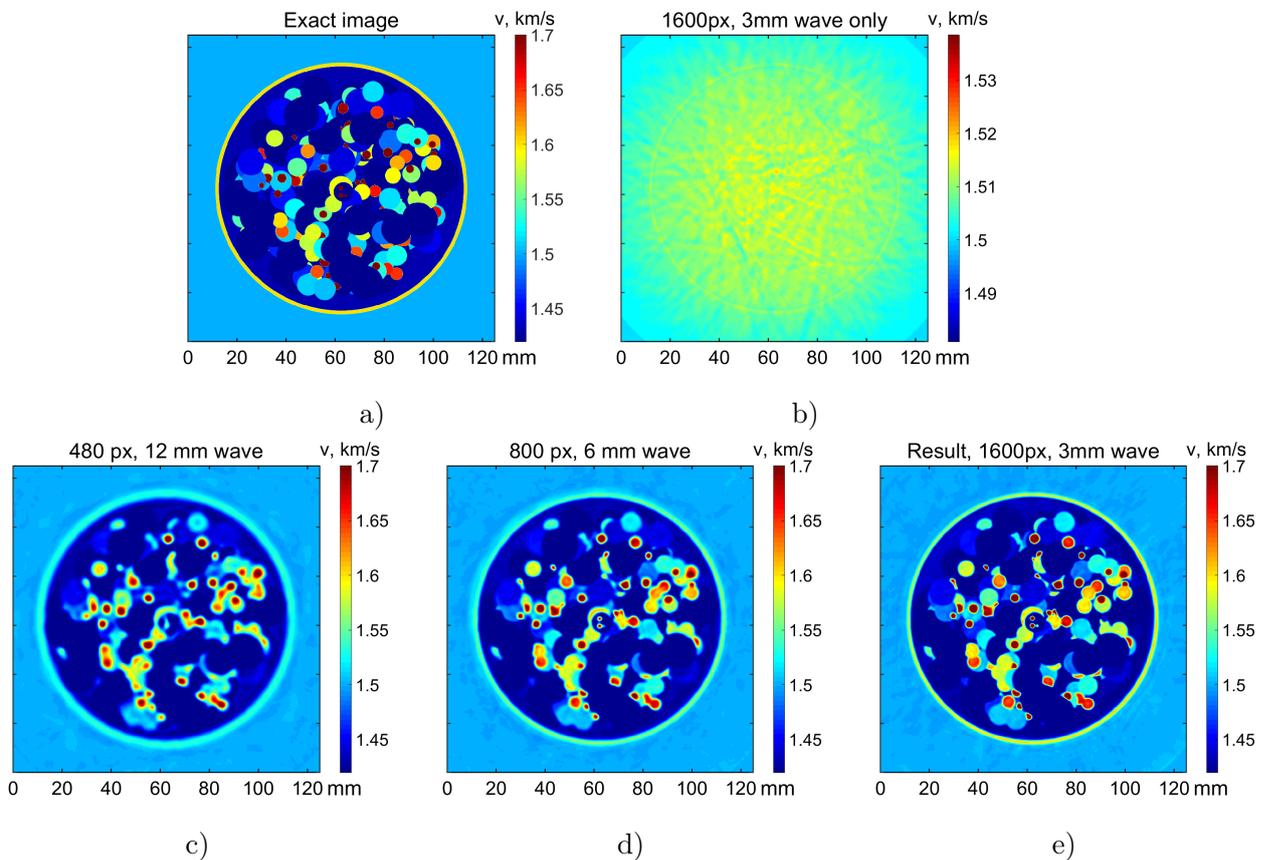
Figure 12. Complex phantom example: a – exact image, b – image reconstructed at 3 mm wavelength from a constant initial approximation, c – image reconstructed using 12 mm wavelength from a constant initial approximation, d – image reconstructed using 6 mm wavelength at the 3rd stage, e – final image after 5 stages

Figure 12 shows an example of reconstructing a complex internal structure of an object. Figure 12a shows the exact image of the model object (phantom) In this example, the gradient descent process does not converge using a constant initial approximation and a short sounding pulse with a wavelength of 3 mm. Figure 12b shows the reconstruction result for this case.

The iterative process stops at a local minimum of the residual functional and the image is not reconstructed. To obtain a tomographic image of such an object, it is necessary to apply the multistage method.

At the first stage, an approximate solution is computed using a constant initial approximation, a central wavelength of 12 mm and a grid size of 480×480 points. The result of the first stage is shown in Fig. 12c. This image has a low resolution, but is sufficiently close to the original. It is used as an initial approximation for the next stage of the multistage method. In total, 5 stages are performed according to Tab. 1. The image quality gradually improves from stage to stage. This means that the images can be analyzed before the whole reconstruction process is completed – as early as the sought-for image features are resolved. This property of the multistage method can improve image analysis time in practical applications. Highest-resolution stages take considerably more time to compute. Figure 12d shows the result of the third stage of the multistage method with an average wavelength of 6 mm and a grid size of 800×800 points. Figure 12e shows the result of the last stage with an average wavelength of 3 mm and a grid dimension of 1600×1600 points. Thus, the MSM method ensures the convergence of the iterative process of gradient-descent minimization of the residual functional and allows for high accuracy tomographic image reconstructions via wave tomography technology.

Conclusion

This article discusses the methods of ultrasound tomography, which can be used to inspect various objects, for example, in nondestructive testing or in medical imaging such as tomographic imaging of soft tissues for early-stage breast cancer diagnosis. A characteristic feature of ultrasound tomography is that it takes into account not only the reflected radiation, but also the radiation transmitted through the object, similarly to X-rays. It is well known that the higher the frequency, the higher the spatial resolution can be achieved in ultrasound imaging. Conventional medical ultrasound diagnostic devices usually employ frequencies above 1 MHz. The tomographic methods considered in this study use low frequencies in the 100–600 kHz band for medical ultrasound tomography. The proposed MSM method uses low frequencies for initial stages of the method to ensure its convergence and high frequencies for final stages to achieve high resolution.

The article discusses a supercomputer implementation of the multistage iterative method (MSM) for solving nonlinear inverse problems of ultrasound tomography. From a mathematical point of view, the problem is posed as a problem of minimizing the residual functional, which is not convex and has local minima. The method is based on a prior information which is typical for most inverse problems of wave tomography. The effectiveness of the MSM method is illustrated on a large number of model problems focused on ultrasound tomographic diagnostics of soft tissues. As shown in the article, the MSM method completely covers the problem of constructing an approximate solution in the problems of medical ultrasound tomographic diagnostics of soft tissues.

In contrast to previous works [33, 34], this article considers the inverse problem of ultrasonic tomography in the wave model without taking into account the absorption of the medium. The numerical algorithms have been adapted for SIMD-capable CPUs. The performance of computing systems is compared for different sizes of the computational grid and for a multistage run. To ensure the convergence of the iterative process of solving the inverse problem, filtering the spectrum of a broadband sounding pulse is proposed. This approach greatly simplifies the exper-

iment compared to using multiple narrow-band emitters. In this study, a significant acceleration of the algorithm has been achieved due to the use of smaller grids for lower frequencies. The method has been tested on phantoms of a complex structure close to reality.

Both CPU- and GPU-based computing clusters can be used to implement the MSM method in practice. It is shown that GPU supercomputers have an advantage, especially for large volumes of data. The numerical algorithm is data-parallel an well-suited for GPU architecture. Modern processors equipped with AVX-512 FPUs are capable of solving small-scale tasks that fit in the CPU cache memory. For large-scale tasks, it is always better to use a GPU equipped with fast on-board memory.

Recently, a new processor architecture has been developed, consisting of several hundred thousand computing cores located on a single silicon wafer. For problems of wave tomography, this line of work is also of great interest. Wafer-scale processors manufactured by Cerebras are designed for machine learning tasks. A specific feature of such tasks is a large number of operations being performed on a fixed data array. Wave tomography problems have the same property – the operation of wave propagation simulation is performed many times on the same data array in the iterative process. On an actual problem of physical process simulation, the wafer-scale system has achieved a performance of 860 TFlops [35]. To date, the wafer-scale architecture has the highest performance-to-memory ratio and is thus the most promising architecture for implementing large-scale wave tomography applications.

To conclude, we note that the use of supercomputer technologies for solving wave tomography problems opens up the possibility of using complex mathematical models describing such physical processes as diffraction, refraction, multiple scattering, and so on. Under such models, the inverse problems are nonlinear and the methods developed in this study can be used to solve such problems. First of all, the developed solution methods focus on the problems of medical ultrasound tomography and inverse problems of electromagnetic sounding.

The paper discusses the possibilities of solving inverse problems of wave tomography in the framework of scalar wave models. From a physical point of view, only one type of wave propagates in a scalar medium – longitudinal compression wave. This model is ideal, for example, for ultrasound tomographic screening for breast cancer diagnosis. However, the scalar model is no longer adequate to reality in such cases as, for example, obtaining ultrasound tomographic images of a knee joint. Unlike in soft tissues, multiple types of waves propagate in solid bodies. There are works concerned with solving direct problems of wave propagation in solids. In some works, attempts are made to solve inverse problems of ultrasound imaging in solids [36, 37]. These problems typically arise in the field of nondestructive testing. Solving inverse problems in vector models is a much more complicated task, compared to scalar models. Solving such problems is impossible without the use of supercomputers.

Acknowledgements

References

1. Radon, J.: Uber die Bestimmung von Funktionen durch ihre Integralwerte langs gewisser Mannigfaltigkeiten. Berichte uber die Verhandlungen der Koniglich-Sachsischen Akademie der Wissenschaften zu Leipzig, Mathematisch-Physische Klasse [Reports on the Proceedings of the Royal Saxonian Academy of Sciences at Leipzig, Mathematical and Physical Section], Leipzig: Teubner 69, 262–277 (1917)

2. Tikhonov, A.N.: Solution of incorrectly formulated problems and the regularization method. Soviet. Math. Dokl. 4, 1035–1038 (1963)

3. Tikhonov, A.N.: Regularization of incorrectly posed problems. Soviet Math. Dokl. 4, 1624–1627 (1963)

4. Tikhonov, A.N., Goncharsky, A.V., Stepanov, V.V., Yagola, A.G.: Numerical Methods for the Solution of Ill-Posed Problems. Springer, Dordrecht (1995). `https://doi.org/10.1007/978-94-015-8480-7`

5. Bakushinsky, A., Goncharsky, A.: Ill-Posed Problems: Theory and Applications. Kluwer Academic Publishers. Springer, Dordrecht (1994). `https://doi.org/10.1007/978-94-011-1026-6`

6. Vinokurov, V. A.: Regularizability of Functions. In: Ill-posed problems in the Natural Sciences, pp. 52–70, Mir, Moscow (1987)

7. Lavrentiev, M.M., Romanov, V.G., Shishatskii, S.P.: Ill-Posed Problems of Mathematical Physics and Analysis. American Mathematical Society, Providence, (1986)

8. Ramm, A.G.: Non-uniqueness of the solution to an inverse problem in geophysics. Inverse problems 2, 123–125 (1986)

9. Engl, H.W., Kunish, K., Neubaer, A.: Convergence rate for Tikhonov regularization of nonlinear ill-posed problem. Inverse problems 4, 532–540 (1989)

10. Groetch, C.W.: The theory of Tikhonov regularization for Fredholm equations of the first kind. SIAM Review 28, 116–118 (1986). `https://doi.org/10.1137/1028033`

11. Nashed, M.Z.: Ill-posed Problems: Theory and Practice. Reidel, Dordrecht (1985)

12. Klibanov, M.V., Timonov, A.A.: Carleman Estimates for Coefficient Inverse Problems and Numerical Applications. Walter de Gruyter GmbH (2004). `https://doi.org/10.1515/9783110915549`

13. Romanov, V.G., Kabanikhin, S.I. Inverse Problems for Maxwell's Equations. VSP, Utrecht, (1994)

14. Goncharsky, A.V., Romanov, S.Y., Seryozhnikov, S.Y.: Inverse problems of 3D ultrasonic tomography with complete and incomplete range data. Wave Motion 51(3), 389–404 (2014). https://doi.org/10.1016/j.wavemoti.2013.10.001

15. Goncharsky, A.V., Romanov, S.Y., Seryozhnikov, S.Y.: Comparison of the capabilities of GPU clusters and general-purpose supercomputers for solving 3D inverse problems of ultrasound tomography. Journal of Parallel and Distributed Computing 133, 77–92 (2019). https://doi.org/10.1016/j.jpdc.2019.06.008

16. Natterer, F., Sielschott, H., Dorn, O., et al.: Frechet derivatives for some bilinear inverse problems. SIAM J. Appl. Math. 62, 2092–2113 (2002). https://doi.org/10.1137/s0036139901386375

17. Beilina, L., Klibanov, M.V., Kokurin, M.Y.: Adaptivity with relaxation for ill-posed problems and global convergence for a coefficient inverse problem. J. Math. Sci. 167, 279–325 (2010). https://doi.org/10.1007/s10958-010-9921-1

18. Goncharsky, A.V., Romanov, S.Y.: Iterative Methods for Solving Coefficient Inverse Problems of Wave Tomography in Models with Attenuation. Inverse Probl. 33(2), 025003 (2017). https://doi.org/10.1088/1361-6420/33/2/025003

19. Global Optimization: From Theory to Implementation. Edition by Liberti, L., Maculan, N. Springer (2006)

20. Gel'fand, I.M., Tsetlin, M.L.: Some methods of control for complex systems. Russian Mathematical Surveys 17, 95–117 (1962). https://doi.org/10.1070/rm1962v017n01abeh001124

21. Sulimov, A.V., Zheltkov, D.A., Oferkin, I.V., et al.: Tensor Train Global Optimization: Application to Docking in the Configuration Space with a Large Number of Dimensions. Communications in Computer and Information Science (CCIS), vol. 793, pp. 151–167. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71255-0_12

22. Duric, N., Littrup, P., Li, C., et al.: Breast ultrasound tomography: bridging the gap to clinical practice. Proc. SPIE, 8320, 83200O (2012). https://doi.org/10.1117/12.910988

23. Jirik, R., Peterlik, I., Ruiter, N., et al.: Sound-speed image reconstruction in sparse-aperture 3D ultrasound transmission tomography. IEEE Trans. Ultrason. Ferroelectr. Freq. Control. 59, 254–264 (2012). https://doi.org/10.1109/tuffc.2012.2185

24. Goncharsky, A.V., Romanov, S.Y., Seryozhnikov, S.Y.: A computer simulation study of soft tissue characterization using low-frequency ultrasonic tomography. Ultrasonics 67, 136–150 (2016). https://doi.org/10.1016/j.ultras.2016.01.008

25. Shen, Y., Shamout, F.E., Oliver, J.R., et al.: Artificial intelligence system reduces false-positive findings in the interpretation of breast ultrasound exams. Nat. Commun. 12, 5645 (2021). https://doi.org/10.1038/s41467-021-26023-2

26. Natterer, F.: Sonic imaging. In: Scherzer O. (eds) Handbook of Mathematical Methods in Imaging, pp. 1253–1278. Springer, New York (2015). https://doi.org/10.1007/978-1-4939-0790-8_37

27. Bakushinsky, A., Goncharsky, A., Romanov, S., Seatzu, S.: On the identification of velocity in seismics and in acoustic sounding. Pubblicazioni IAGA, Series "Problemi non ben posti e inversi" 71 (1994)

28. Bakushinskii, A.B., Kozlov, A.I., Kokurin, M.Yu.: One Inverse Problem for a Three-Dimensional Wave Equation. Computational Mathematics and Mathematical Physics 43(8), 1149–1158 (2003)

29. Bakushinskii, A.B., Leonov, A.S.: Low-cost numerical method for solving a coefficient inverse problem for the wave equation in three-dimensional space. Computational Mathematics and Mathematical Physics 58, 548–561 (2018). https://doi.org/10.1134/s0965542518040073

30. Klibanov, M.V., Li, J., Zhang, W.: Linear Lavrent'ev Integral Equation for the Numerical Solution of a Nonlinear Coefficient Inverse Problem. SIAM J. Appl. Math. 81(5), 1954–1978 (2021). https://doi.org/10.1137/20M1376558

31. Hamilton, B., Bilbao, S.: Fourth-order and optimised finite difference schemes for the 2-D wave equation. In: Proc. of the 16th Int. Conference on Digital Audio Effects (DAFx-13). pp. 363–395. Springer (2013)

32. Voevodin, V., Antonov, A., Nikitenko, D., et al.: Supercomputer Lomonosov-2: Large Scale, Deep Monitoring and Fine Analytics for the User Community. Supercomputing Frontiers and Innovations 6(2), 4–11 (2019). https://doi.org/10.14529/jsfi190201

33. Goncharsky, A., Seryozhnikov, S.: Supercomputer technology for ultrasound tomographic image reconstruction: mathematical methods and experiment. In: Voevodin, V., Sobolev, S. (eds) Supercomputing. RuSCDays 2018. Communications in Computer and Information Science, vol. 965, pp. 401–413. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-05807-4_34

34. Romanov, S. Supercomputer simulation study of the convergence of iterative methods for solving inverse problems of 3D acoustic tomography with the data on a cylindrical surface // In: Voevodin, V., Sobolev, S. (eds) Supercomputing. RuSCDays 2018. Communications in Computer and Information Science, vol. 965, pp. 388–400. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-05807-4_33

35. Rocki, K., Essendelft, D.V., Sharapov, I.: Fast stencil-code computation on a wafer-scale processor. In: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (2020)

36. Lechleiter, A., Schlasche, J.W.: Identifying Lame parameters from time-dependent elastic wave. Inverse Problems in Science and Engineering 25(1), 2–26 (2017). https://doi.org/10.1080/17415977.2015.1132713

37. He, J., Rao, J., Fleming, J.D., et al.: Numerical ultrasonic full waveform inversion (FWI) for complex structures in coupled 2D solid/fluid media. Smart Materials and Structures 30, 085044 (2021). https://doi.org/10.1088/1361-665X/ac0f44