Supercomputer Technology for Ultrasound Tomographic Image Reconstruction: Mathematical Methods and Experimental Results

Alexander Goncharsky and Sergey Seryozhnikov(✉)

Lomonosov Moscow State University, Moscow, Russia.
gonchar@srcc.msu.ru
s2110sj@gmail.com

Abstract. This paper is concerned with layer-by-layer ultrasound tomographic imaging methods for differential diagnosis of breast cancer. The inverse problem of ultrasound tomography is formulated as a coefficient inverse problem for a hyperbolic differential equation. The scalar mathematical model takes into account wave phenomena, such as diffraction, refraction, multiple scattering, and absorption of ultrasound. The algorithms were tested on real data obtained in experiments on a test bench for ultrasound tomography studies. Low-frequency ultrasound in the 100 – 500 kHz band was used for sounding. An important result of this study is an experimental confirmation of the adequacy of the underlying mathematical model. The ultrasound tomographic imaging methods developed have a spatial resolution of 2 mm, which is acceptable for medical diagnostics. The experiments were carried out using phantoms with parameters close to the acoustical properties of human soft tissues. The image reconstruction algorithms are designed for graphics processors. Architecture of the GPU cluster for ultrasound tomographic imaging is proposed, which can be employed as a computing device in a tomographic complex.

Keywords: Ultrasound Tomography · Coefficient Inverse Problem · Medical Imaging · GPU Cluster.

1 Introduction

Supercomputer technologies open up new possibilities in medical diagnostics. One such example is ultrasound tomography. This technology is developed extensively in the USA, Germany, and Russia [1–4]. Medical imaging, especially the differential diagnosis of breast cancer, is a promising application of ultrasound tomography. Wave tomography methods can also be applied to nondestructive imaging of solids [5].

The inverse problem of wave tomography is nonlinear. Breakthrough results in the field of solving inverse problems of wave tomography have been obtained in recent years. Explicit formulas for the gradient of the residual functional between the measured and the numerically simulated wave fields in various formulations were derived in the works [6–9]. In our earlier papers [10–12] we derived a representation for the gradient of the residual functional in terms of a scalar mathematical model that accounts for diffraction and absorption effects.

Iterative supercomputer-oriented algorithms for reconstructing ultrasound tomographic images have been developed whose efficiency was confirmed via solving numerous model problems [13]. Testing algorithms by applying them to model problems is necessary for design of ultrasound tomography devices. However, numerical simulations can not assess the adequacy of the underlying mathematical model and only physical experiments can answer this question.

In this study, algorithms for solving the inverse problems of ultrasound tomography in the layer-by-layer formulation are tested on experimental data. For this purpose, a test bench for ultrasound tomography studies was constructed. One of the key results of this study is that it demonstrates that the mathematical model agrees well with real physical processes. Solving inverse problems of wave tomography using experimental data showed that tomographic methods can reconstruct inhomogeneities approximately 2 mm in size. Such spatial resolution is quite acceptable for medical imaging. The parameters of the phantoms used in experiments were close to the acoustic properties of human soft tissues.

The inverse problem of wave tomography is a nonlinear inverse problem with millions of unknowns. The algorithms developed are designed for GPU clusters and we used GPU nodes of the "Lomonosov-2" supercomputer at the Moscow State University [14] to process experimental data. In this paper, we discuss the architecture of a GPU cluster which can be used as a computing device in a tomographic complex.

2 Formulation of the Inverse Problem and its Solution Methods

Figure 1a shows the scheme of a layer-by-layer tomographic examination used in this study. Object G is submerged in water L with known acoustical properties. Acoustic sounding pulses are emitted by transducers at positions 1. Figure 1b shows a typical waveform of a sounding pulse. The scattered ultrasound waves are measured at positions 2. At each position, the signal is recorded as a function of time $u(t)$. Figure 1c shows a typical received signal. This scheme of the experiment allows us to measure both transmitted and reflected waves. The measurement points lie at surface Γ, which in this case is circular.

The objective is to reconstruct the internal structure of the object using these measurements. This formulation of the problem is called the time-domain formulation. The measurements are repeated for multiple horizontal planes, thus obtaining a layered representation of a 3D object. The inverse problem is solved independently for each of the horizontal planes $z = const$.

In this study, we use a scalar wave model, which accounts for all the wave phenomena, such as diffraction, refraction, multiple scattering and absorption

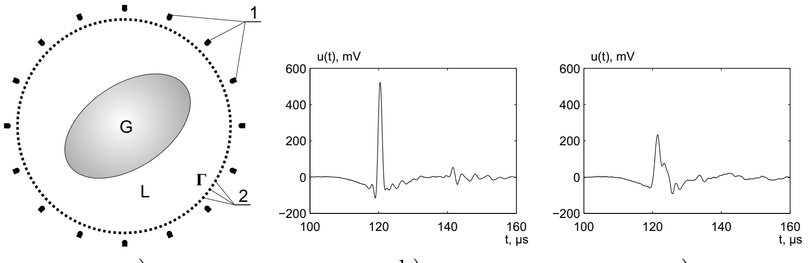

a) b) c)

Fig. 1. Layer-by-layer tomographic examination: a) placement of emitters and detectors; b) waveform of a sounding pulse; c) waveform of a received pulse

of ultrasound. In the scalar model, the acoustic pressure $u(\boldsymbol{r},t)$ is described by a hyperbolic equation. The inverse problem of ultrasound tomography is a coefficient inverse problem, in which we use the measurements of the wave field taken at some surface $\varGamma$ to obtain the coefficients $c(\boldsymbol{r})$ and $a(\boldsymbol{r})$ of the wave equation:

$$c(\boldsymbol{r})u_{tt}(\boldsymbol{r},t) + a(\boldsymbol{r})u_t(\boldsymbol{r},t) - \varDelta u(\boldsymbol{r},t) = 0; \tag{1}$$

$$u(\boldsymbol{r},t)|_{t=0} = F_0(\boldsymbol{r}), \quad u_t(\boldsymbol{r},t)|_{t=0} = F_1(\boldsymbol{r}). \tag{2}$$

Here, $c(\boldsymbol{r}) = 1/v^2(\boldsymbol{r})$, where $v(\boldsymbol{r})$ is the speed of sound; $a(\boldsymbol{r})$ is the absorption factor in the medium; $\boldsymbol{r} = \{x,y\}$ is the position of the point inside the reconstructed 2D plane, and $\varDelta$ is Laplacian with respect to $\boldsymbol{r}$.

This formulation assumes a transparent (non-reflecting) boundary of the computational domain. A non-reflecting boundary condition [15] in the form $\partial u/\partial \boldsymbol{n} = -c^{-0.5}\partial u/\partial t$ is applied at the boundary. Here, $\boldsymbol{n}$ is a vector pointing towards the ultrasound emitter. The initial conditions $F_0(\boldsymbol{r})$ and $F_1(\boldsymbol{r})$ represent the wave field radiated from the emitter at the initial time of the numerical simulation. We obtain them via the time-reversal method [16] by solving equation (1) in reverse time with boundary condition $u(\boldsymbol{s},t)|_{s\in\varGamma} = U_0(\boldsymbol{s},t)$, where $U_0(\boldsymbol{s},t)$ are the data measured at surface $\varGamma$ in a homogeneous medium without the object.

The coefficient inverse problem considered is ill-posed. The methods for solving such problems were developed in [17–19]. We formulate the inverse problem as that of minimizing the residual functional

$$\varPhi(u(c,a)) = \frac{1}{2}\int\limits_0^T \int\limits_S (u(\boldsymbol{s},t) - U(\boldsymbol{s},t))^2 \, \mathrm{d}\boldsymbol{s}\, \mathrm{d}t \tag{3}$$

for its argument (c,a). Here $U(\boldsymbol{s},t)$ are the data measured at surface $\varGamma$ for the time period $(0,T)$, $u(\boldsymbol{s},t)$ is the solution of the direct problem (1)–(2) for the given $c(\boldsymbol{r}) = 1/v^2(\boldsymbol{r})$ and $a(\boldsymbol{r})$. The residual functional is the sum of the residuals (3) obtained for each position of the emitter.

Representations for the gradient of the residual functional in various formulations were obtained by the authors in [9, 13]. The gradient $\Phi'(u(c,a)) = \{\Phi'_c(u), \Phi'_a(u)\}$ of the functional (3) with respect to the variation of the sound speed and absorption factor $\{dc, da\}$ has the form:

$$\Phi'_c(u(c)) = \int\limits_0^T w_t(\boldsymbol{r},t)u_t(\boldsymbol{r},t)\,\mathrm{d}t, \quad \Phi'_a(u(a)) = \int\limits_0^T w_t(\boldsymbol{r},t)u(\boldsymbol{r},t)\,\mathrm{d}t. \tag{4}$$

Here $u(\boldsymbol{r},t)$ is the solution of the direct problem (1)–(2), and $w(\boldsymbol{r},t)$ is the solution of the "conjugate" problem with the given $c(\boldsymbol{r})$, $a(\boldsymbol{r})$, and $u(\boldsymbol{r},t)$:

$$c(\boldsymbol{r})w_{tt}(\boldsymbol{r},t) - a(\boldsymbol{r})w_t(\boldsymbol{r},t) - \Delta w(\boldsymbol{r},t) = E(\boldsymbol{r},t); \tag{5}$$

$$w(\boldsymbol{r},t=T) = 0, \quad w_t(\boldsymbol{r},t=T) = 0; \tag{6}$$

$$E(\boldsymbol{r},t) = \begin{cases} u(\boldsymbol{r},t) - U(\boldsymbol{r},t), \text{ where } \boldsymbol{r} \in \Gamma \text{ and } U(\boldsymbol{r},t) \text{ is known;} \\ 0, \text{ otherwise.} \end{cases} \tag{7}$$

Non-reflecting boundary condition [15] is applied at the boundary of the computational domain. Thus, to compute the gradient (4) we need to solve the direct problem (1)–(2) and the "conjugate" problem (5)–(7).

Given a formula for the gradient (4), we can construct various iterative algorithms that minimize the residual functional, such as the steepest descent method. Let us assume that the coefficients $c^{(m)}$ and $a^{(m)}$ for m-th iteration have been determined. To construct the next iterative approximation, we compute the gradient $\{\Phi'_c(u), \Phi'_a(u)\}$ at point $\{c^{(m)}, a^{(m)}\}$ and solve the problem of minimizing a one-dimensional functional along the direction of the gradient. As the next iterative approximation, we choose the point $\{c^{(m+1)}, a^{(m+1)}\} = \arg\min\limits_{\alpha>0} \Phi(c^{(m)} - \alpha\Phi'_c, a^{(m)} - \alpha\Phi'_a)$, and so on.

Minimization methods based on the explicit formula for the gradient allow us to propose efficient numerical algorithms for the approximate solution of inverse problems of wave tomography. The gradient method has regularizing properties and stops when the value of the residual functional becomes equal to the error of the input data [18].

To solve equations (1)–(2) and (5)–(7), we use the finite-difference time-domain method (FDTD) on uniform grids. We introduce a uniform discrete grid with a space step of h and a time step of τ:

$$x_i = ih,\ 0 \le i < n; \quad y_j = jh,\ 0 \le j < n; \quad t_k = k\tau,\ 0 \le k < m.$$

The parameters h and τ are related by the Courant stability condition $\sqrt{2}c^{-0.5}\tau < h$, where $c^{-0.5} = v$ is the speed of sound. To approximate the equation (1) we use the following second-order finite difference scheme:

$$c_{ij}\frac{u_{ij}^{k+1} - 2u_{ij}^k + u_{ij}^{k-1}}{\tau^2} + a_{ij}\frac{u_{ij}^{k+1} - u_{ij}^{k-1}}{\tau} - \frac{\Delta u_{ij}^k}{h^2} = 0. \tag{8}$$

Here $u_{ij}^k = u(x_i, y_j, t_k)$ are the values of $u(\boldsymbol{r}, t)$ at point (i, j) at the time step k; c_{ij} and a_{ij} are the values of $c(\boldsymbol{r})$ and $a(\boldsymbol{r})$ at point (i, j). The first term approximates $c(\boldsymbol{r})u_{tt}(\boldsymbol{r}, t)$, the second — $a(\boldsymbol{r})u_t(\boldsymbol{r}, t)$. Discrete Laplacian Δ is computed using a fourth-order optimized finite difference scheme [20]. Collecting the terms with u_{ij}^{k+1} for $(k+1)$-th time step, we obtain the explicit formula for simulating the ultrasound wave sequentially in time. A similar scheme is used to solve the equations (5)–(7) for $w(\boldsymbol{r}, t)$ in reverse time.

The gradient of the residual functional is computed by formula (4). An explicit formula allows us to use gradient-based mininization methods, such as the steepest descent method [21]

3 GPU Implementation of the Layer-by-layer Tomographic Image Reconstruction Algorithm

The numerical method was implemented on graphics processors using OpenCL™ technology. An iterative gradient method is used for image reconstruction. At the first iteration, we use the initial approximation for the unknown coefficients $c^{(0)}(\boldsymbol{r}) = 1500\,\mathrm{m{\cdot}s^{-1}}$, $a^{(0)}(\boldsymbol{r}) = 0$. The following steps are performed at each gradient descent iteration (m):

1. The direct problem (1)–(2) is solved and the wave field at the boundary $u(\boldsymbol{s}, t)$, $\boldsymbol{s} \in \Gamma$ is stored in memory.
2. The "conjugate" problem (5)–(7) is solved in reverse time. To compute $u(\boldsymbol{r}, t)$ we use the boundary data $u(\boldsymbol{s}, t)$ previously stored, and to compute $w(\boldsymbol{r}, t)$ we use the measured data $U(\boldsymbol{s}, t), \boldsymbol{s} \in \Gamma$. The gradient Φ'_c, Φ'_a is computed by formula (4), summed over all the emitter positions and accumulated over time.
3. The gradient descent step γ is determined. Its value is adjusted automatically: if the residual decreases ($\Phi^{(m)} < \Phi^{(m-1)}$), then γ is increased, and vice versa.
4. The current approximation is updated: $c^{(m+1)} = c^{(m)} + \gamma\Phi'_c$, $a^{(m+1)} = a^{(m)} + \gamma\Phi'_a$.

The most computationally expensive tasks are solving the "conjugate" problem (5)–(7) and computing the gradient. Figure 2 illustrates the GPU implementation of the algorithm. To adapt the algorithm for GPU, the task is split into independent thread blocks. The gradient is summed over all the emitter positions, and therefore an efficient solution is to compute the gradient sequentially for each emitter position, accumulating the partial sum in the shared memory of the GPU.

Each thread block processes an $LX \times LY$ area of the $L \times L$ image, where LX is equal to the number of threads in a block, and LY is determined by the amount of available shared memory: $LY = SharedMemorySize/(LX \times 8)$. The number of threads per block LX is chosen to optimize the GPU workload. For tested devices, the optimum value of LX was found to be equal to the maximum supported number of threads (512). Therefore, $LY \leq 12$ for a typical NVidia

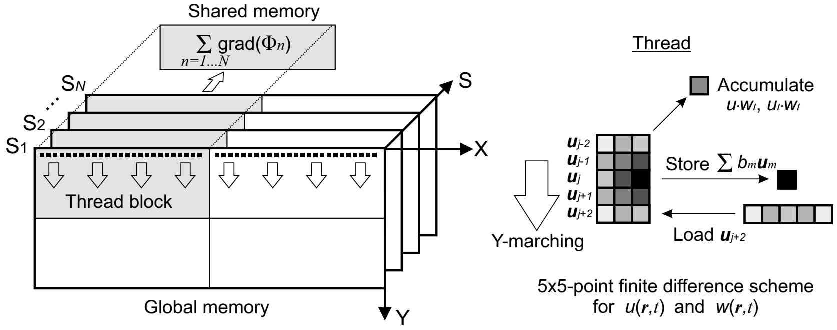

Fig. 2. GPU-optimized algorithm of computing the gradient of the residual functional

device with 48 KB of shared memory. We used $LX = L$ for image sizes $L \leq 512$ and $LX = L/2$ for image sizes $512 < L \leq 1024$.

Inside each $LX \times LY$ data block the computations are carried out sequentially along the Y-axis (Y-marching method) and sequentially for each emitter position $\mathrm{S}_n,\ n = 1, ..., N$. The Y-marching method ensures high automatic caching efficiency due to the sequential memory access pattern and allows increasing the stencil size with no significant performance loss. At each step along the Y-axis, the value of u_{ij}^{k-1} and a vector $\boldsymbol{u}_{j+2}$ containing the acoustic pressure values for 5 consecutive points along the X-axis are loaded into the registers. Since the discrete Laplacian coefficient matrix is symmetric [20], only 3 components per row are stored in the registers for a 5×5-point finite difference scheme.

Acoustic pressure u_{ij}^{k+1} at the next time step is computed by formula (8), which can be expressed as a scalar product with some coefficients b_m. The coefficients are constant for all points of the stencil except the center point (i, j). The wave field w_{ij}^{k+1} is computed using a similar formula. The gradient $\{u \cdot w_t, u_t \cdot w_t\}$ is then computed and added to the partial sum over emitter positions, which is accumulated in the shared memory. The process is repeated for all emitter positions.

The data are stored in a 3D array X×Y×S, where S corresponds to the emitter position. A 2D subarray for one position S_n contains the values of $u(\boldsymbol{r}, t)$ and $w(\boldsymbol{r}, t)$ for three time steps. Computations with experimental data showed that 32-bit floating point type is sufficient for data representation. Thus, the amount of memory required to store all the data is $24 \times L^2 \times N$ bytes, where L is the grid size along one dimension, and N is the number of emitter positions. For 24 positions and $L = 768$, the volume of data amounts to 350 MB. Additionally, $32 \times N \times T \approx 100$ MB of global memory are used to store the values of $u(\boldsymbol{r}, t)$ and $w(\boldsymbol{r}, t)$ at the measurement surface $\mathit{\Gamma}$. Here T is the number of time steps in the simulation, which is on the order of L. Thus, modern graphics cards with at least 1 GB of memory can host all the data in the global memory, eliminating the need to access the system RAM throughout the simulation.

The algorithm developed was tested on NVidia GeForce GTX Titan and Tesla K40s devices. The program profiling statistics are listed in Table 1. We used OpenCL profiling events to obtain start and end times of each computational kernel and data transfer operation. The statistics were accumulated for 10 gradient descent iterations on a GTX Titan device.

Table 1. Program profiling statistics (GTX Titan, 768×768 grid, 10 iterations)

Runs	**Time, μs**	**Min, μs**	**Max, μs**	**Total**	**%**	**Function**
5040	1924	1866	2425	9699.17ms	21.16%	`ForwardWave`
5020	6340	6303	7441	31.83s	69.44%	`BackwardWave`
10060	271	181	766	2733.53ms	5.96%	`BoundaryCond`
5020	31	25	41	156.72ms	0.34%	`LoadBound`
5040	32	24	678	161.89ms	0.35%	`SaveBound`
5020	57	53	70	287.69ms	0.63%	`LoadExData`
5020	36	35	86	183.14ms	0.40%	`SaveExData`
264	2313	1	29224	610.87ms	1.33%	*Data transfers*

The most time-consuming functions are wave simulations that solve the direct (`ForwardWave`) and the conjugate (`BackwardWave`) problems. The time spent in other functions amounted to less than 10%, which means near-linear scaling with respect to the problem size.

Table 2 lists the memory requirements for different problem sizes and the computation times for tested devices. The number of emitter positions was $N = 24$, and 40 gradient descent iterations were performed in each test run. These parameters correspond to the values used to reconstruct the images from experimental data.

Table 2. Memory requirements and computational times for different problem sizes

Grid size (L)	**512**	**768**	**1024**
GPU memory used, MB	220	450	780
Time, GTX Titan	114s	258s	477s
Time, Tesla K40s	130s	294s	562s

The total number of operations per problem is proportional to $L^3 \cdot N$, and the amount of memory is proportional to $L^2 \cdot N$. The experimental reconstructions were performed using a grid size of $L = 768$. For this problem size, reconstructing one horizontal layer using one GPU device takes approximately 5 minutes.

One of the aims of this study is to assess the parameters of a GPU cluster that can be used as a computing device for medical imaging. The benchmarking results showed that to reconstruct the images for multiple layers in practically feasible time, which is 15 – 30 minutes, the number of graphics processors in the cluster should amount to 1/3 – 1/4 of the number of layers. The computations for each layer are independent. To obtain the images for 30 – 40 layers, a GPU cluster containing 8 – 16 devices would be required. Since the memory requirements for layer-by-layer image reconstruction are quite low, the cluster can include 4 – 8 dual-GPU boards such as GeForce GTX 690. Devices equipped with High Bandwidth Memory (HBM) have approximately 3 times higher performance than GDDR5-equipped devices, so a cluster of 4 such devices can perform the image reconstruction in a reasonable time. Such GPU-cluster would be compact and inexpensive, and can be incorporated as a computing device in a diagnostic facility.

4 The Test Bench for Ultrasound Tomography Studies

The experimental measurements were carried out using a custom-made test bench for ultrasound tomography studies, shown in Fig.3. The test bench is designed for ultrasound tomographic imaging of 3D objects with parameters close to the acoustical properties of soft tissues. The ultrasound emitter 1 and the receiver 2 can be rotated using the motors 3 around the object 4 and moved in the vertical direction. The maximum size of the inspected object is approximately 12×12 cm.

A broadband ultrasound transducer with a central frequency of 400 kHz was used as an emitter. The frequency spectrum of the sounding pulses covers the 100 – 500 kHz band. The duration of the sounding pulses was approximately 2 μs. Teledyne Reson TC4038 hydrophone was used as a receiver.

Figure 4 shows the functional scheme of the test bench. Electrical pulses are generated by the arbitrary waveform generator and amplified by the power amplifier. The amplitude of the pulses at the transducer is ±80V. Acoustic signals are received by the hydrophone, amplified by the preamplifier and digitized using the ADC module. The digitized data are collected by a PC. The motors, the waveform generator, and the ADC module are synchronized by a controller unit.

5 Ultrasound Tomographic Image Reconstruction Using Experimental Data

For ultrasound tomographic image reconstruction experiments we used a cylindrical phantom made of soft silicone, 56 mm in diameter and 130 mm in height. Figure 5 shows the 3D model of the phantom. Phantom 1 contains a tilted 10 mm cylindrical hole 2, and a 15 mm vertical cylindrical hole 3. The phantom is affixed to the mount with three metal pins 1 mm in diameter.

The speed of sound in the silicone is $\approx 1400\,\mathrm{m\cdot s^{-1}}$; in water outside of the phantom and inside the holes — $1500\,\mathrm{m\cdot s^{-1}}$. Figure 5b shows the horizontal sound speed cross-section of the phantom in the z=30 mm plane. Figure 6 shows the reconstructed sound speed images of the phantom in different horizontal

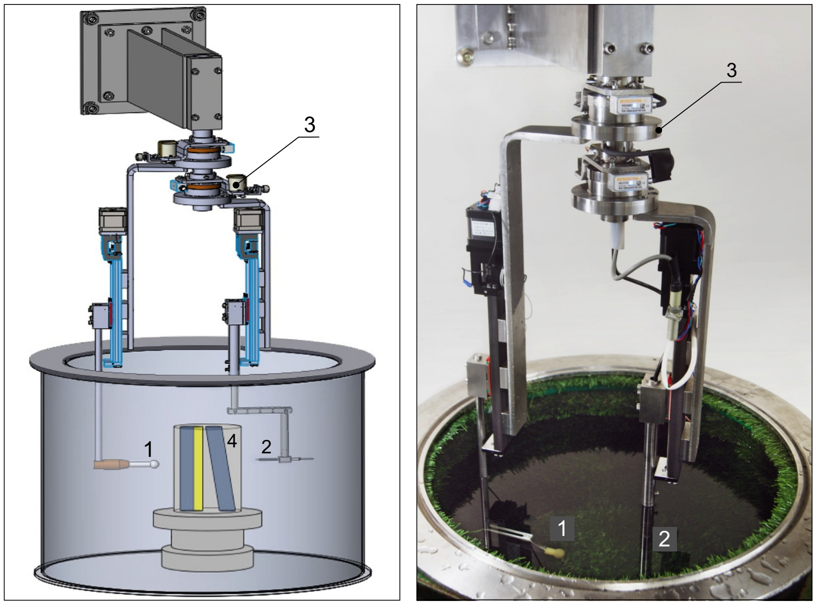

(a) (b)

Fig. 3. The test bench for ultrasound tomography studies: a) 3D model, b) photo of the mechanical assembly

cross-sections. The reconstructed speed of sound inside the holes is equal to the speed of sound in water (1500 $\mathrm{m{\cdot}s^{-1}}$). The 7 mm rod with a higher speed of sound is clearly resolved in all cross-sections. The reconstructed speed of sound inside the rod is 1800 $\mathrm{m{\cdot}s^{-1}}$. The tilted hole changes its position from image to image, as the vertical position of the emitter and the receiver changes. The 1 mm metal pins are clearly visible in all cross-sections.

As is evident from the reconstructed images in Fig.6, the spatial resolution is approximately 2 mm, while the sounding wavelength is approximately 4 mm. The sound speed difference between the silicone material and water amounted to $\approx 5\%$. Yet, the algorithms developed reconstruct the sound speed image with high precision even for a low-contrast object.

The size of the data obtained in a tomographic examination amounted to 35 MB per horizontal layer. The data for 24 emitter positions and 500 receiver positions recorded over a time period of 200μs with 14-bit precision at a sampling rate of 5 MSPS were used to reconstruct the image in each layer. The grid step was 0.4 mm and the number of gradient descent iterations required to obtain the images amounted to $30-40$. Image reconstruction was performed on GPU nodes of the "Lomonosov-2" supercomputer at the Moscow State University.

One of the drawbacks of layer-by-layer ultrasound tomography is that the spatial resolution along the Z-axis is lower than the spatial resolution in X–Y

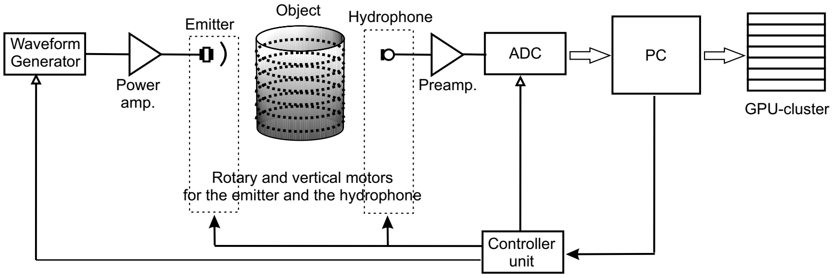

Fig. 4. The functional scheme of the test bench

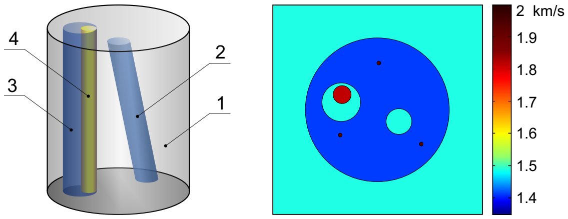

a) b)

Fig. 5. The phantom: a) 3D model, b) exact horizontal sound speed cross-section at z=30 mm

planes. This is due to the fact that in the layered model only waves propagating along the reconstructed plane are taken into account. In order to increase the resolution, out-of-plane wave sources must be used as well. This scheme of tomographic examination requires solving a three-dimensional inverse problem, which is much more computationally expensive than a layer-by-layer problem. The numerical simulations presented in [13, 22] showed that the 3D scheme can be implemented in practice and can provide higher image reconstruction quality than the layered scheme. However, a 3D setup requires much more complex measuring equipment.

6 Conclusion

The main result of this study is a successful application of the ultrasound tomography algorithms developed to real experimental data. Numerous experiments performed on a test bench for experimental tomographic studies showed that the scalar wave model based on a second-order hyperbolic equation can be used to reconstruct images of real objects. The tomographic methods developed have high resolution, which is quite acceptable for medical imaging. The experiments were carried out using phantoms with parameters close to the acoustic properties of human soft tissues.

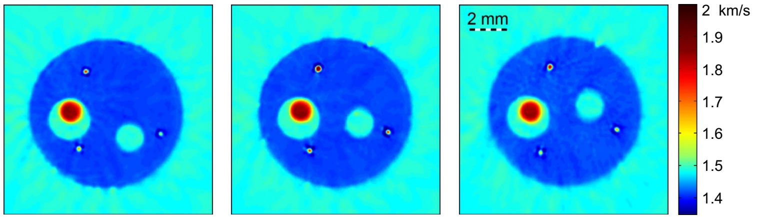

z=30 mm z=60 mm z=90 mm

Fig. 6. Reconstructed sound speed cross-sections of the phantom

The use of GPU clusters is the most practically feasible option for implementing the iterative solution algorithms for 2D inverse problems. A finite-difference method can be efficiently parallelized using data-parallel architectures such as SIMD and GPU. Experiments showed that single-precision floating-point arithmetic is sufficient for solving direct and inverse problems. Thus, widely available GPU devices can be used for image reconstruction.

The algorithms and the scheme of experiments are designed for implementation on a dedicated GPU cluster The experimental setup assumes detailed recording of the signal on a cylindrical surface for multiple emitter positions. The number of measurement points in each horizontal layer amounts to 500. At the same time, the number of emitter positions is small — approximately 20 in each layer. Thus, an image of a single layer can be reconstructed by a single GPU device in a short time. We propose the architecture of a GPU cluster with a relatively small number of devices proportional to the number of reconstructed layers of the 3D image. The properties of such a GPU cluster, such as the size, power consumption, and cost, allow it to be used in the new medical ultrasound tomographic facilities being developed.

Acknowledgements

This work was supported by Russian Science Foundation [grant number 17-11-01065]. The research is carried out at Lomonosov Moscow State University. The research is carried out using the equipment of the shared research facilities of HPC computing resources at Lomonosov Moscow State University.

References

1. M. Sak, N. Duric, P. Littrup, L. Bey-Knight, H. Ali, P. Vallieres, M. E. Sherman, G. L. Gierach, Using speed of sound imaging to characterize breast density, Ultrasound Med. Biol. 43 (1) (2017) 91–103. doi:10.1016/j.ultrasmedbio.2016.08.021.
2. R. Jifik, I. Peterlik, N. Ruiter, J. Fousek, R. Dapp, M. Zapf, J. Jan, Sound-speed image reconstruction in sparse-aperture 3-d ultrasound transmission tomography, IEEE T. Ultrason. Ferr. 59 (2) (2012) 254–264. doi:10.1109/tuffc.2012.2185.

3. J. Wiskin, D. Borup, M. Andre, S. Johnson, J. Greenleaf, Y. Parisky, J. Klock, Three-dimensional nonlinear inverse scattering: Quantitative transmission algorithms, refraction corrected reflection, scanner design, and clinical results, J. Acoust. Soc. Am. 133 (5) (2013) 3229–3229. doi:10.1121/1.4805138.
4. V. A. Burov, D. I. Zotov, O. D. Rumyantseva, Reconstruction of the sound velocity and absorption spatial distributions in soft biological tissue phantoms from experimental ultrasound tomography data, Acoust. Phys. 61 (2) (2015) 231–248. doi:10.1134/s1063771015020013.
5. E. G. Bazulin, A. V. Goncharsky, S. Y. Romanov, S. Y. Seryozhnikov, Parallel CPU- and GPU-algorithms for inverse problems in nondestructive testing, Lobachevskii Journal of Mathematics 39 (4) (2018) 486–493. doi:10.1134/S1995080218040030.
6. F. Natterer, Sonic imaging, in: Handbook of Mathematical Methods in Imaging, Springer Nature, 2014, pp. 1–23. doi:10.1007/978-3-642-27795-5_37-2.
7. M. V. Klibanov, A. A. Timonov, Carleman Estimates for Coefficient Inverse Problems and Numerical Applications, Walter de Gruyter GmbH, 2004. doi:10.1515/9783110915549.
8. A. Goncharsky, S. Romanov, S. Seryozhnikov, Inverse problems of 3D ultrasonic tomography with complete and incomplete range data, Wave Motion 51 (3) (2014) 389–404. doi:10.1016/j.wavemoti.2013.10.001.
9. A. V. Goncharsky, S. Y. Romanov, Supercomputer technologies in inverse problems of ultrasound tomography, Inverse Probl. 29 (7) (2013) 075004. doi:10.1088/0266-5611/29/7/075004.
10. A. V. Goncharsky, S. Y. Romanov, Iterative methods for solving coefficient inverse problems of wave tomography in models with attenuation, Inverse Probl. 33 (2) (2017) 025003. doi:10.1088/1361-6420/33/2/025003.
11. A. V. Goncharsky, S. Y. Romanov, S. Y. Seryozhnikov, Low-frequency three-dimensional ultrasonic tomography, Doklady Physics 468 (3) (2016) 268–271. doi:10.7868/S0869565216150093.
12. A. V. Goncharsky, S. Y. Romanov, Inverse problems of ultrasound tomography in models with attenuation, Phys. Med. Biol. 59 (8) (2014) 1979–2004. doi:10.1088/0031-9155/59/8/1979.
13. A. Goncharsky, S. Romanov, S. Seryozhnikov, A computer simulation study of soft tissue characterization using low-frequency ultrasonic tomography, Ultrasonics 67 (2016) 136–150. doi:10.1016/j.ultras.2016.01.008.
14. V. Voevodin, S. A. Zhumatiy, S. Sobolev, A. Antonov, P. Bryzgalov, D. A. Nikitenko, K. S. Stefanov, V. V. Voevodin, Practice of "Lomonosov "supercomputer, Open Systems J. (7) (2012) 36–39.
URL https://www.osp.ru/os/2012/07/13017641/
15. B. Engquist, A. Majda, Absorbing boundary conditions for the numerical simulation of waves, Math. Comput. 31 (139) (1977) 629–629. doi:10.1090/s0025-5718-1977-0436612-4.
16. M. Fink, Time reversal in acoustics, Contemporary Physics 37 (2) (1996) 95–109. doi:10.1080/00107519608230338.
17. A. N. Tikhonov, Solution of incorrectly formulated problems and the regularization method, Soviet Math. Dokl. 4 (1963) 1035–1038.
18. A. Bakushinsky, A. Goncharsky, Ill-Posed Problems: Theory and Applications, Springer Nature, 1994. doi:10.1007/978-94-011-1026-6.
19. A. N. Tikhonov, A. V. Goncharsky, V. V. Stepanov, A. G. Yagola, Numerical Methods for the Solution of Ill-Posed Problems, Springer Netherlands, 1995. doi:10.1007/978-94-015-8480-7.

20. B. Hamilton, S. Bilbao, 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), Springer, 2013, pp. 363–395.
21. A. Bakushinsky, A. Goncharsky, Iterative methods for solving ill-posed problems, Nauka, Moscow, 1989.
22. A. V. Goncharsky, S. Y. Seryozhnikov, The architecture of specialized GPU clusters used for solving the inverse problems of 3d low-frequency ultrasonic tomography, in: V. Voevodin, S. Sobolev (Eds.), Supercomputing. RuSCDays 2017. Communications in Computer and Information Science, Vol. 793, Springer, Cham., 2017, pp. 363–395.