Photoacoustic Reconstruction with Progressive Grid Refinement

Gabriel Bordovsky

Faculty of Information Technology, Brno University of Technology, Centre of Excellence IT4Innovations, CZ

BRNO UNIVERSITY OF TECHNOLOGY FACULTY OF INFORMATION TECHNOLOGY

Overview

The photoacoustic tomography (PAT) is based on the fact that chromophones in a specific tissue, such as hemoglobin in veins and tumors, are able to absorb light emitted from the outside of the tissue. The absorbed light is transformed into heat and causes a thermoelastic expansion. This generates broadband ultrasonic (US) waves in the tissue that can be captured by ultrasound sensors.

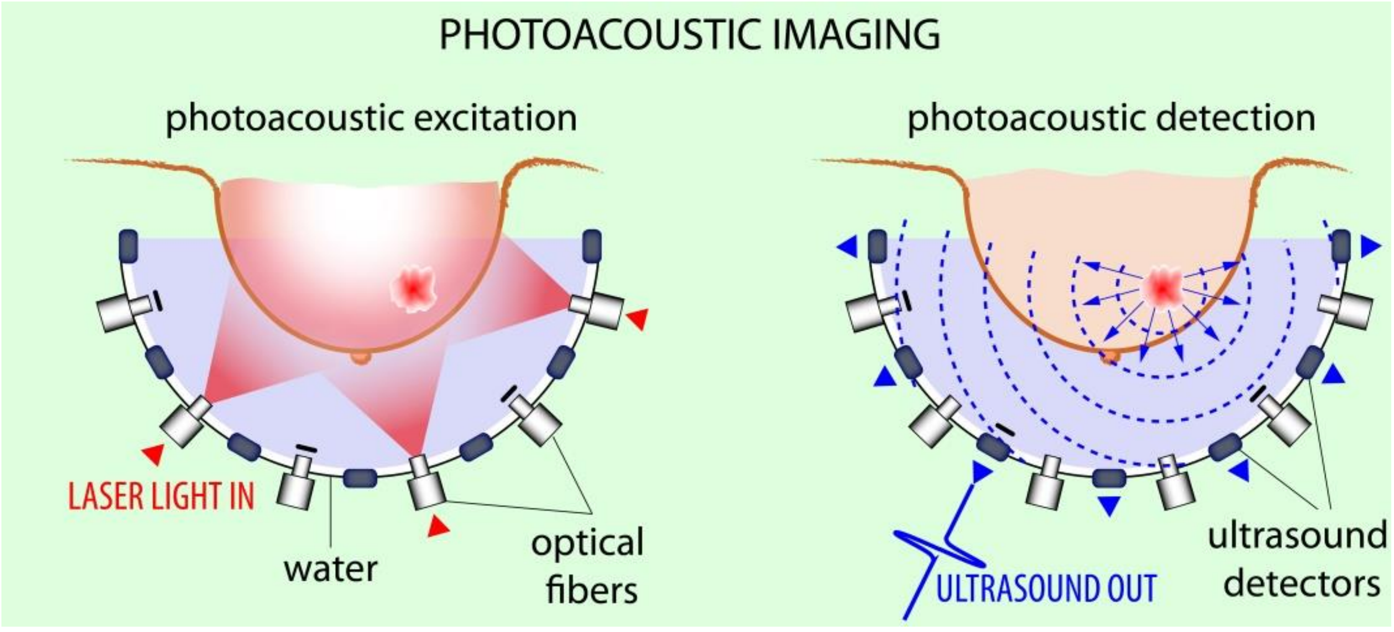

The goal of the PAT is to create an image showing the amount of absorbed light as origin of detected US waves. The model–based iterative approach simulating the US propagation is considered as the least restrictive method. The simulation is computationally demanding.

In this study we investigate an impact of sensor signal downsampling on a residuum. Three pressure distribution approximations are computed from downsampled signal using grids with spacing 0.2 mm (no grid refinement), 0.4 mm and 0.8 mm supporting frequencies up to 2.4 MHz, 1.2 MHz and 0.6 MHz respectively. Each distribution is then used as an initial acoustic pressure for simulation on 0.2 mm grid. Produced signals are then compared with the input signal. Distributions producing lower residuum are closer to the desired solution.

Generation of Source Signal

Our reconstruction study uses a sensor bowl with a radius of 10 cm covered with broadband ultrasound sensors with a central frequency of 0.5 MHz. The simulation domain is represented by a rectangular cuboid with dimensions of 20 cm × 20 cm × 13 cm.

To obtain the signal that would be captured by the ultrasound sensors we simulate the wave propagation from the veins using information obtained by MRI scans. Based on the grid resolution, the simulation is able to support only a limited range of frequencies. Bandwidth of the signals produced by the PA effect on an anatomically correct breast phantoms (with/without a tumor) is in the picture below together with an ideal frequency sensitivity of (0.5 MHz) and 1 MHz sensors.

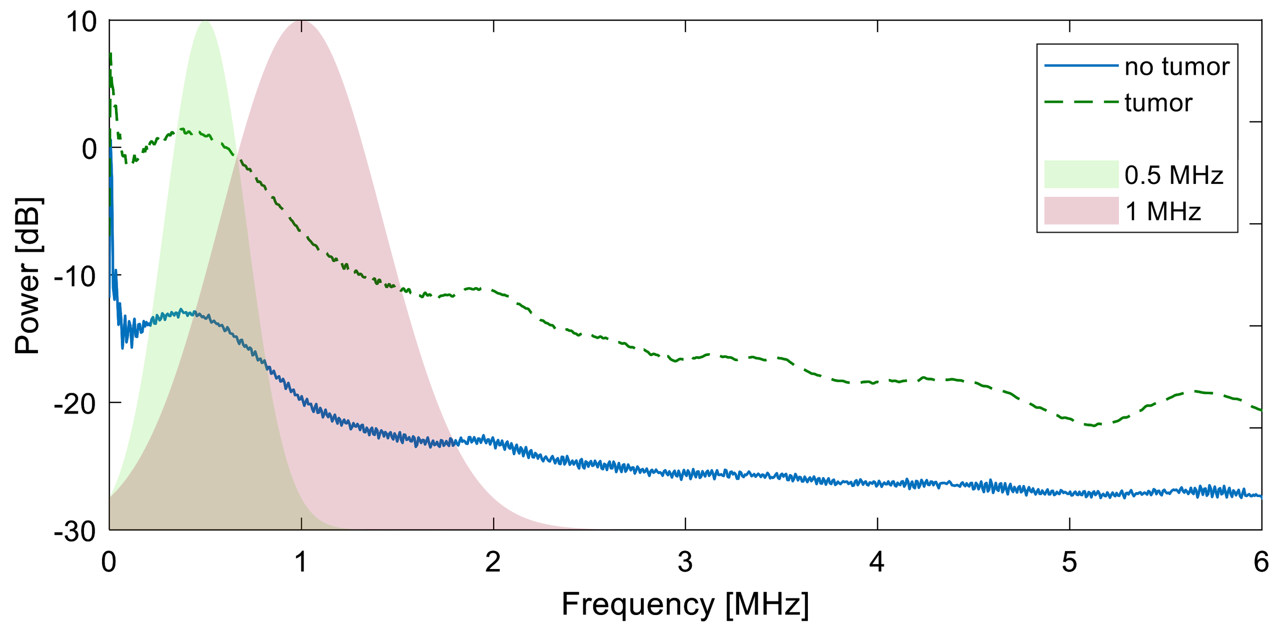

Memory Requirements

Parallel solver using a single MPI process per socket was used for the simulations on the Salomon cluster. The memory required for the biggest used matrix containing the source signal is displayed in the following table.

d_x	size	req. MPI	used MPI	nodes
0.2 mm	36.6 GB	19	32	16
0.4 mm	7.6 GB	4	16	8
0.8 mm	1.1 GB	1	8	4

Since MPI I/O is only able to manage up to 2 GB of data per process in a single dataset, the size of the biggest matrix dictates the minimal number of MPI processes required and consequentially the number of used Salomon nodes.

Reconstruction Time

The following table shows a difference between number of gridpoints (N_x, N_y and N_z), required number of simulation timesteps N_t and wall-time required for each simulation t_{sim}.

d_x	$N_{x,y}$	N_z	N_t	t_{sim}
0.2 mm	1024	672	5220	4h 34min
0.4 mm	528	350	2616	1h 22min
0.8 mm	280	192	1320	9 min

As expected, grids with bigger spacing d_x require less computational resources, fewer simulation steps and produce the result faster for a price of neglecting the higher frequencies. The images below show a frontal (left) and lateral (right) view of the pressure distribution in a breast, computed on 0.2 mm grid (top) and 0.8 mm (bottom).

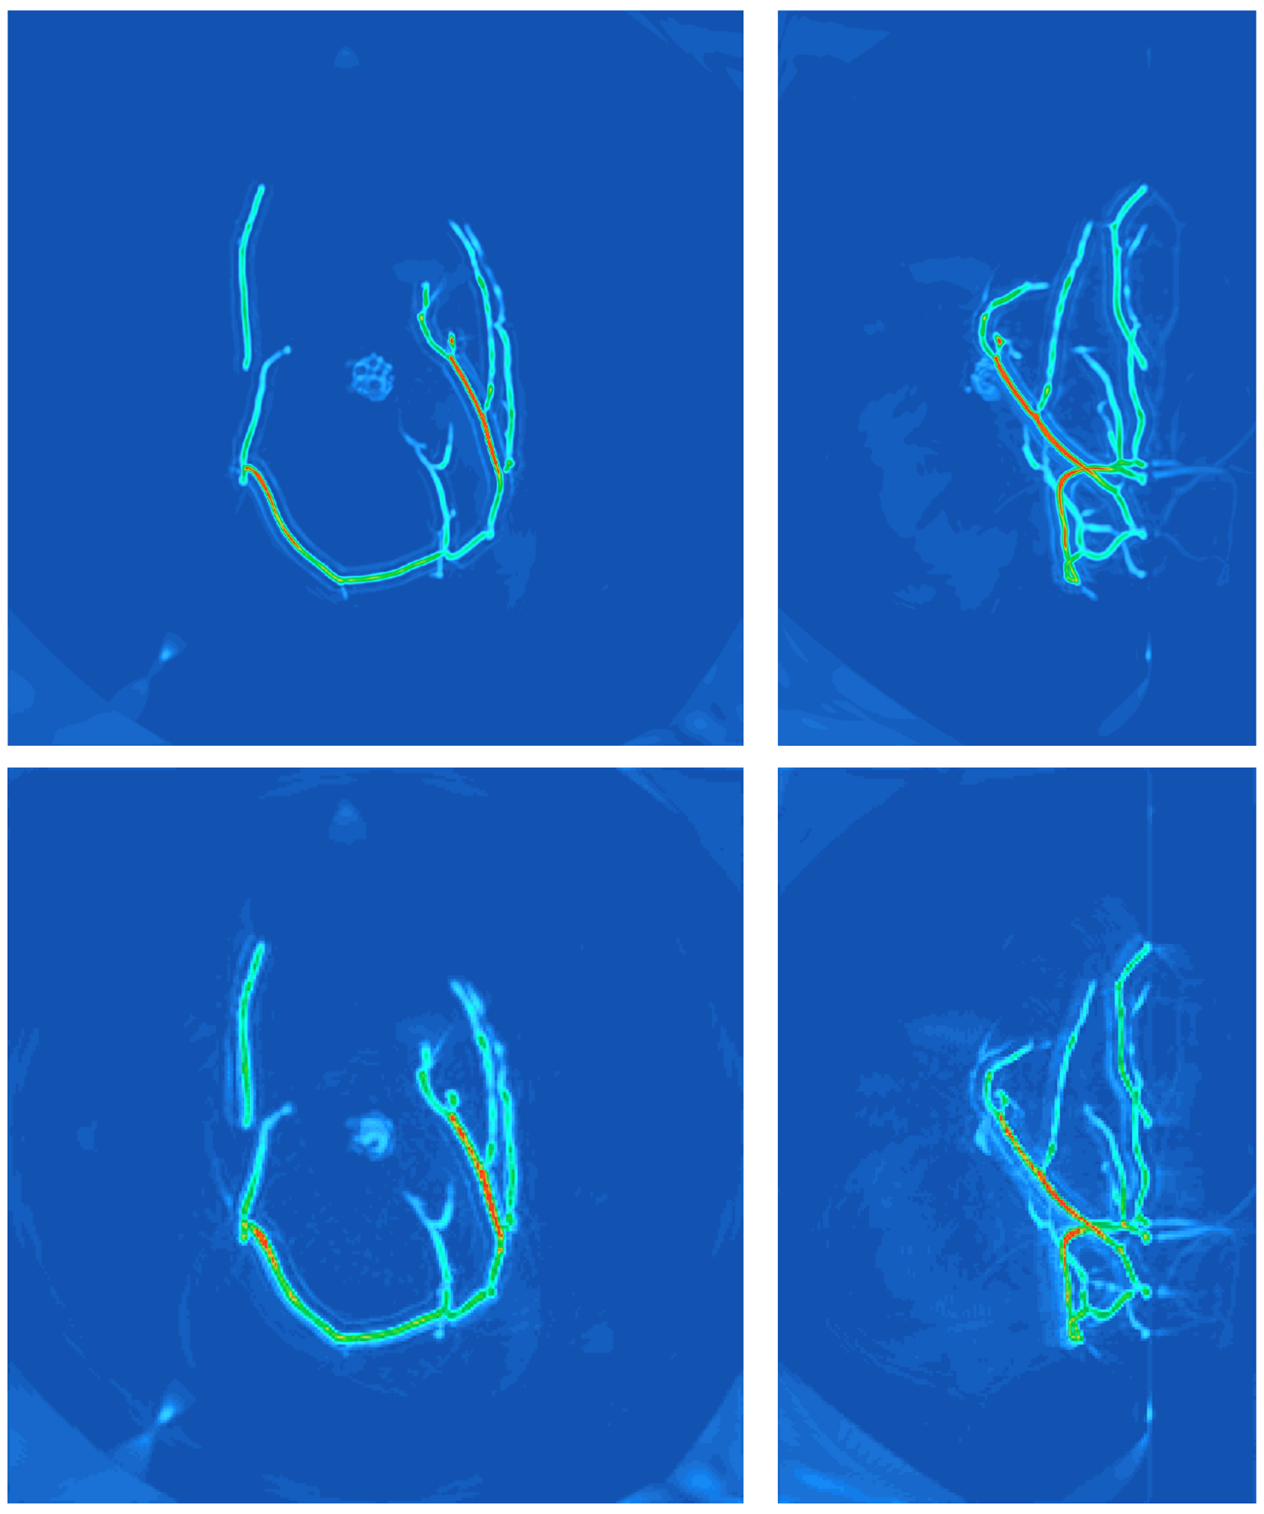

The neglect of the higher frequencies on the 0.8 mm grid causes a noticeable blur.

Results

To compare the results of each variant, we up-sampled the pressure distributions and ran the forward simulations on the 0.2 mm grid. The produced signals are then compared against the input signal obtaining normalized residuum R. The residuum would then be used in iterative framework to correct the approximations. The error in residuum ϵ caused by the down sampling describes the price we had to pay for the speedup S.

d_x	R [%]	ϵ	S
0.2 mm	91.72	-	1
0.4 mm	96.79	5.5%	3.1
0.8 mm	99.81	8.8%	30.4

The quantitative difference 8.8% between residuum produced by reconstruction on 0.2 mm and 0.8 mm grids seems insignificant compared to gain speedup over 30 with 4 times less computational resources.

Conclusion

The whole iterative reconstruction require around 50 simulations to reduce residuum under reasonable amount, which equals to 9 days of computational time. The results indicate that it is possible to create an initial low resolution approximation and speedup the whole process. The high resolution image is still required, but the total time of its construction could be reduced if some part of the process employed the down-sampled grid. This requires further studies.

k-Wave
A MATLAB toolbox for the time-domain simulation of acoustic wave fields

VSB TUO | IT4I

This work was supported by The Ministry of Education, Youth and Sports from the National Programme of Sustainability (NPU II) project "IT4Innovations excellence in science - LQ1602" and by the IT4Innovations infrastructure which is supported from the Large Infrastructures for Research, Experimental Development and Innovations project "IT4Innovations National Supercomputing Center - LM2015070". This project has received funding from the European Union's Horizon 2020 research and innovation programme H2020 ICT 2016-2017 under grant agreement No 732411 and is an initiative of the Photonics Public Private Partnership. This work was supported by the FIT-S-17-3994 Advanced parallel and embedded computer systems project.