Large-scale Ultrasound Simulations with Local Fourier Basis Decomposition

Jiri Jaros[1], Matej Dohnal[1], Bradley E. Treeby[2]

[1]Faculty of Information Technology, Brno University of Technology, CZ
[2]Department of Medical Physics and Bioengineering, University College London, UK

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

1 Overview

High-intensity focused ultrasound (HIFU) is an emerging non-invasive cancer therapy that uses tightly focused ultrasound waves to destroy tissue cells through localised heating. The treatment planning goal is to select the best transducer position and transmit parameters to accurately target the tumour. The path of the ultrasound waves can be predicted by solving acoustic equations based on mass, momentum, and energy conservation. However, this is a computationally difficult problem because the domain size is very large compared to the acoustic wavelength.

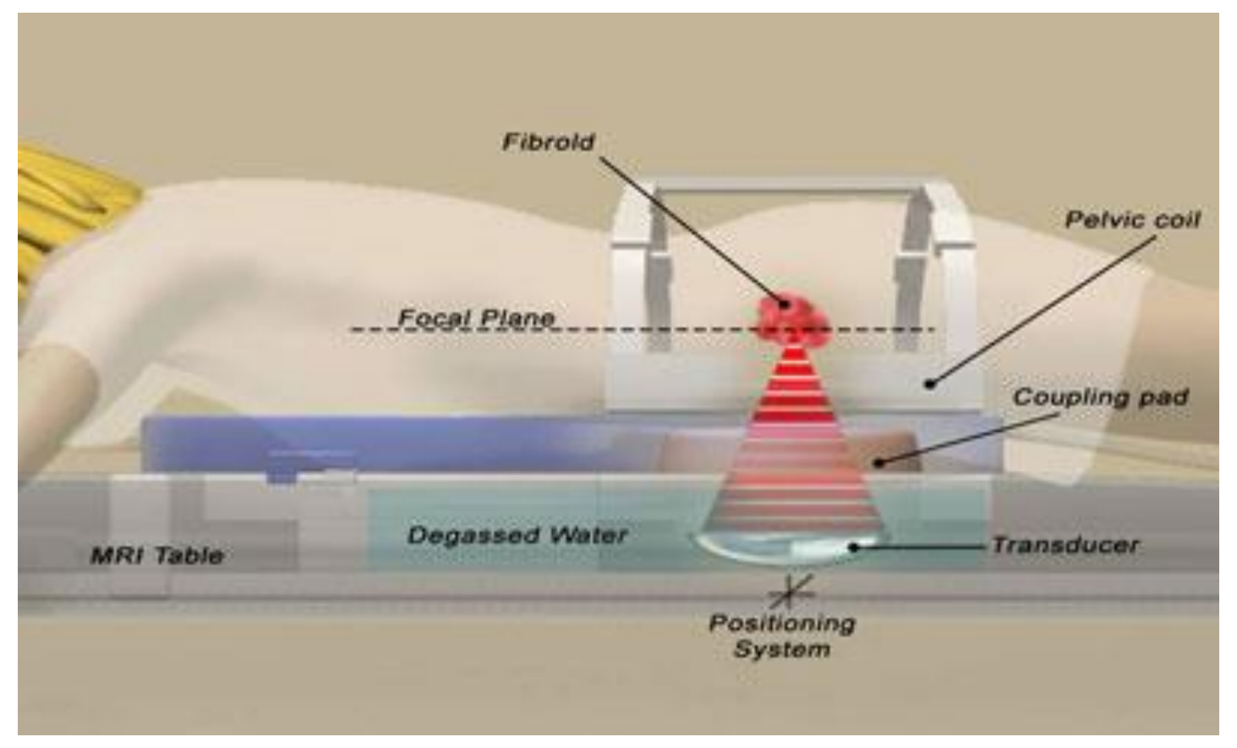

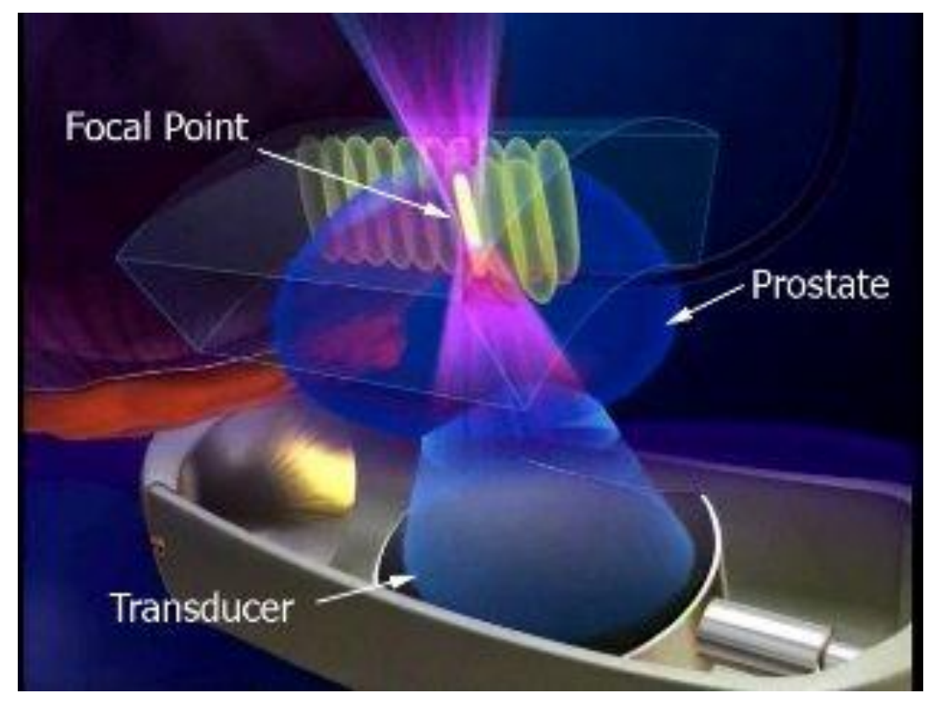

2 Nonlinear Ultrasound Wave Propagation in Tissue

The governing equations must account for the nonlinear propagation of ultrasound waves in tissue, which is a heterogeneous and absorbing medium. Accurately accounting for acoustic absorption is critical for predicting ultrasound dose under different conditions.
The required acoustic equations can be written as:

$$\frac{\partial \mathbf{u}}{\partial t} = -\frac{1}{\rho_0}\nabla p \qquad \textit{momentum conservation}$$

$$\frac{\partial \rho}{\partial t} = -\left(2\rho + \rho_0\right)\nabla\cdot\mathbf{u} - \mathbf{u}\cdot\nabla\rho_0 \qquad \textit{mass conservation}$$

$$p = c_0^2\left(\rho + \mathbf{d}\cdot\nabla\rho_0 + \frac{B}{2A}\frac{\rho^2}{\rho_0} - \Pi\rho\right) \qquad \textit{pressure-density relation}$$

These equations are discretised using the *k*-space pseudo-spectral method and solved iteratively. This reduces the number of required grid points per wavelength by an order of magnitude compared to finite element or finite difference methods. For uniform Cartesian grids, the gradients can be calculated using the fast Fourier transform (FFT), e.g.

$$\frac{\partial p}{\partial x} = \mathbb{F}^{-1}\left\{ik_x\mathbb{F}\left\{p\right\}\right\}$$

3 Global Domain Decomposition

The 1D global domain decomposition is the natural way of partitioning simulation domains while using spectral methods. This decomposition offers great accuracy, however, suffers from all-to-all communication bottleneck and limited scaling given by the longest domain size.

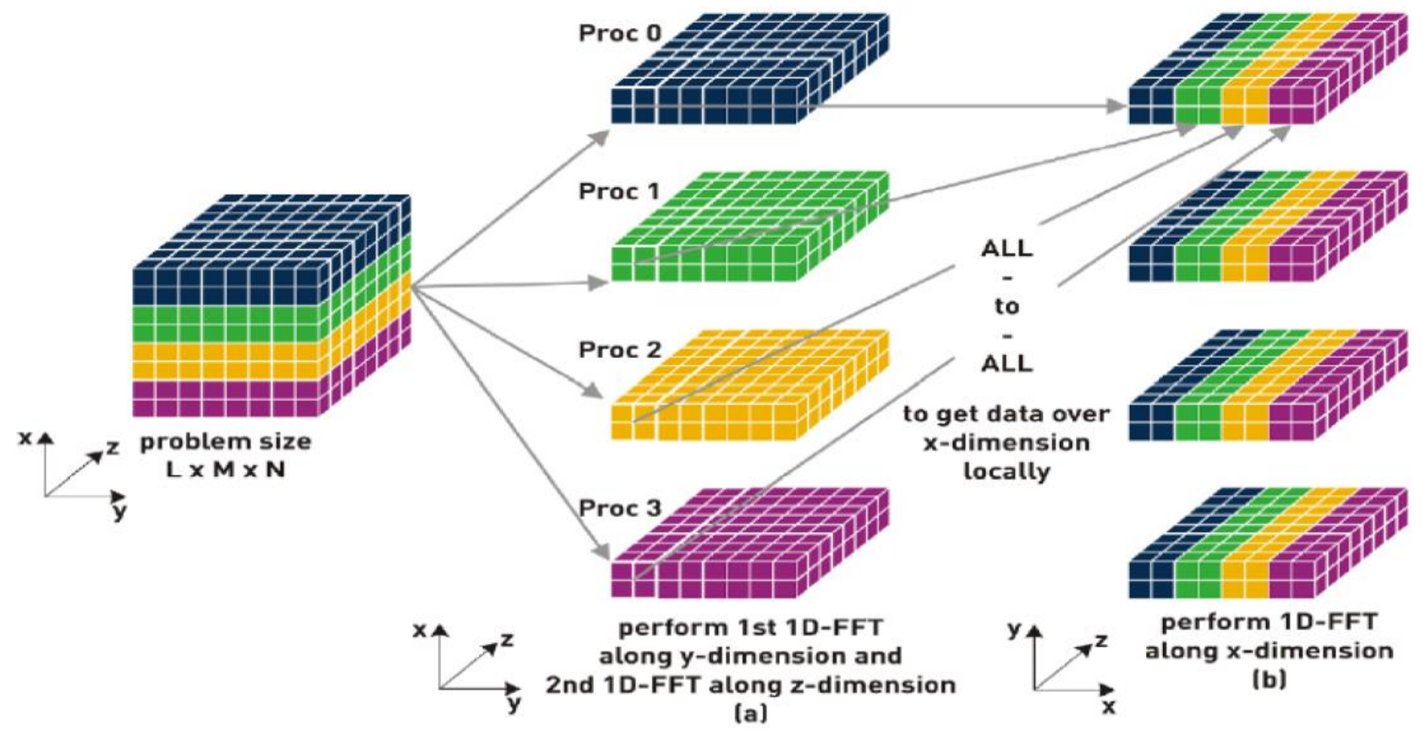

4 Local Fourier Basis Decomposition

Local domain decomposition reduces the communication burden by partitioning the domain into a grid of local subdomains where gradients are calculated locally and the global communication is replaced by the nearest-neighbor halo exchange.

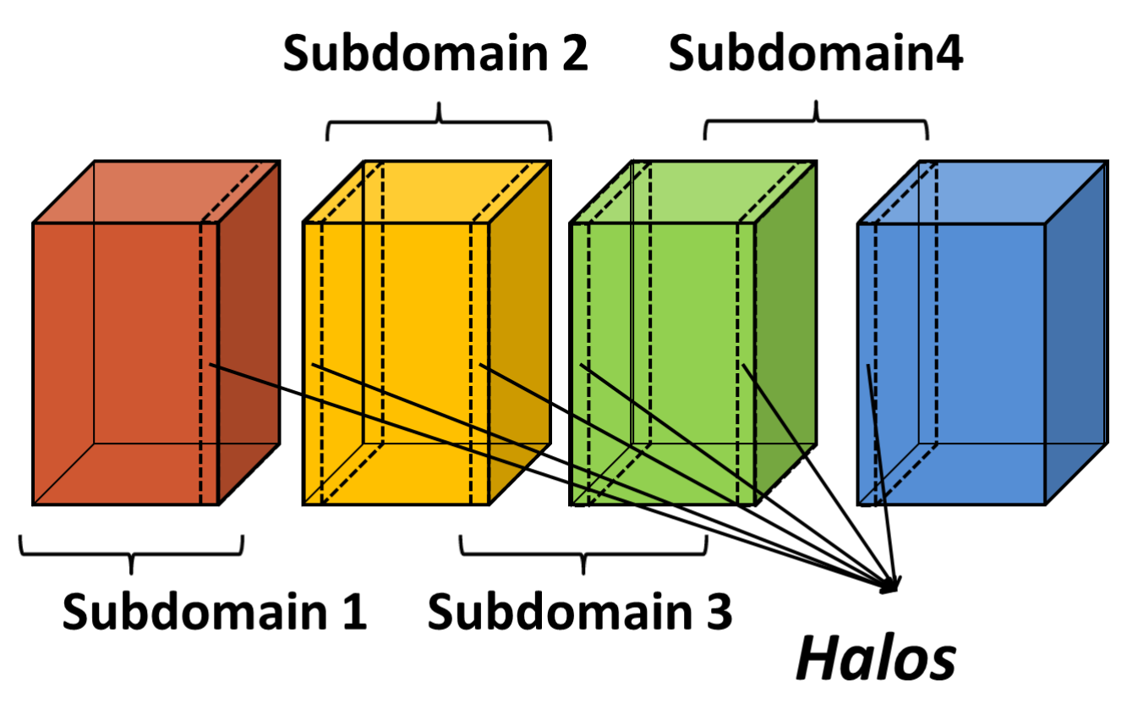

The gradient calculation with the hallo on a *i*-th subdomain reads as follows (*b* is a bell function smoothening the subdomain interface):

$$\frac{\partial p_i}{\partial t} = \mathbb{F}^{-1}\{ik_i\mathbb{F}(b\cdot p_i)\}$$

5 Local Decomposition Accuracy

Since the gradient is not calculated on the whole data, the numeric error is introduced. The level of error can be tuned by the thickness of the halo region.

An appropriate thickness of 16 grid points was experimentally set as a tradeoff between accuracy and performance

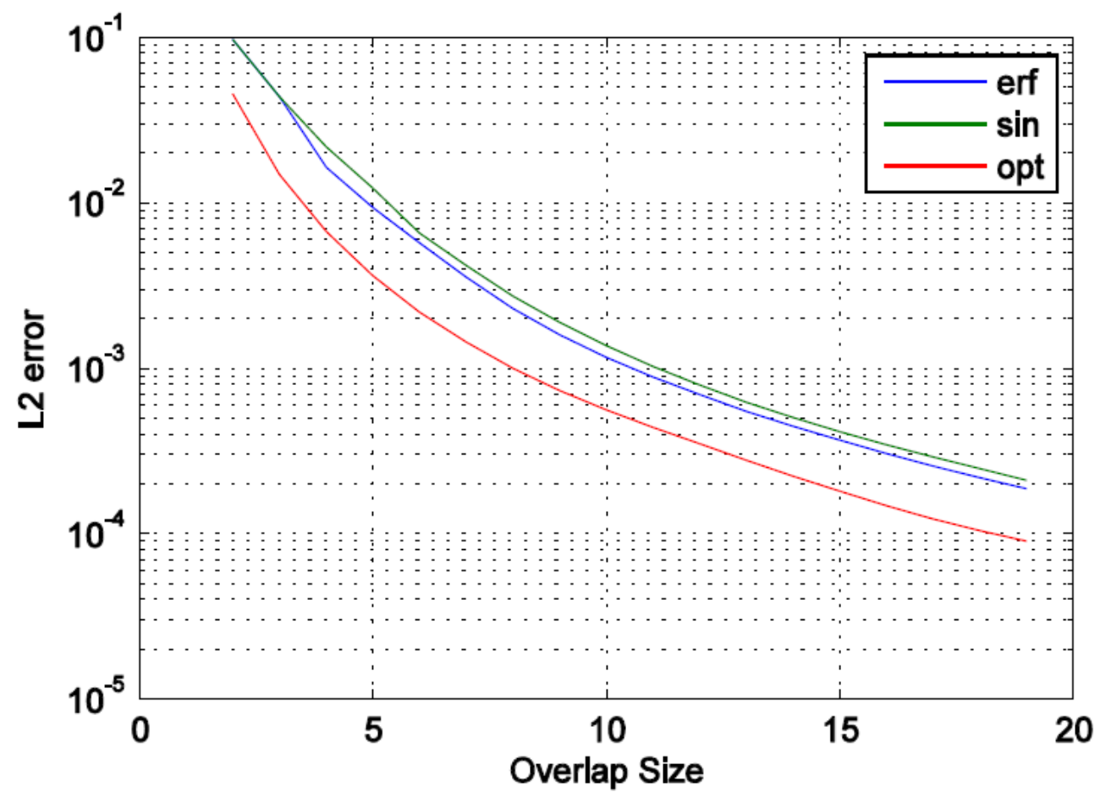

6 Strong Scaling Comparison

Significant improvements in the strong scaling were achieved on the whole simulation size range <512^3, 2048^3> grid points while using 8 to 8192 cores of SuperMUC. The scaling curves all go longer, steeper and smoother.

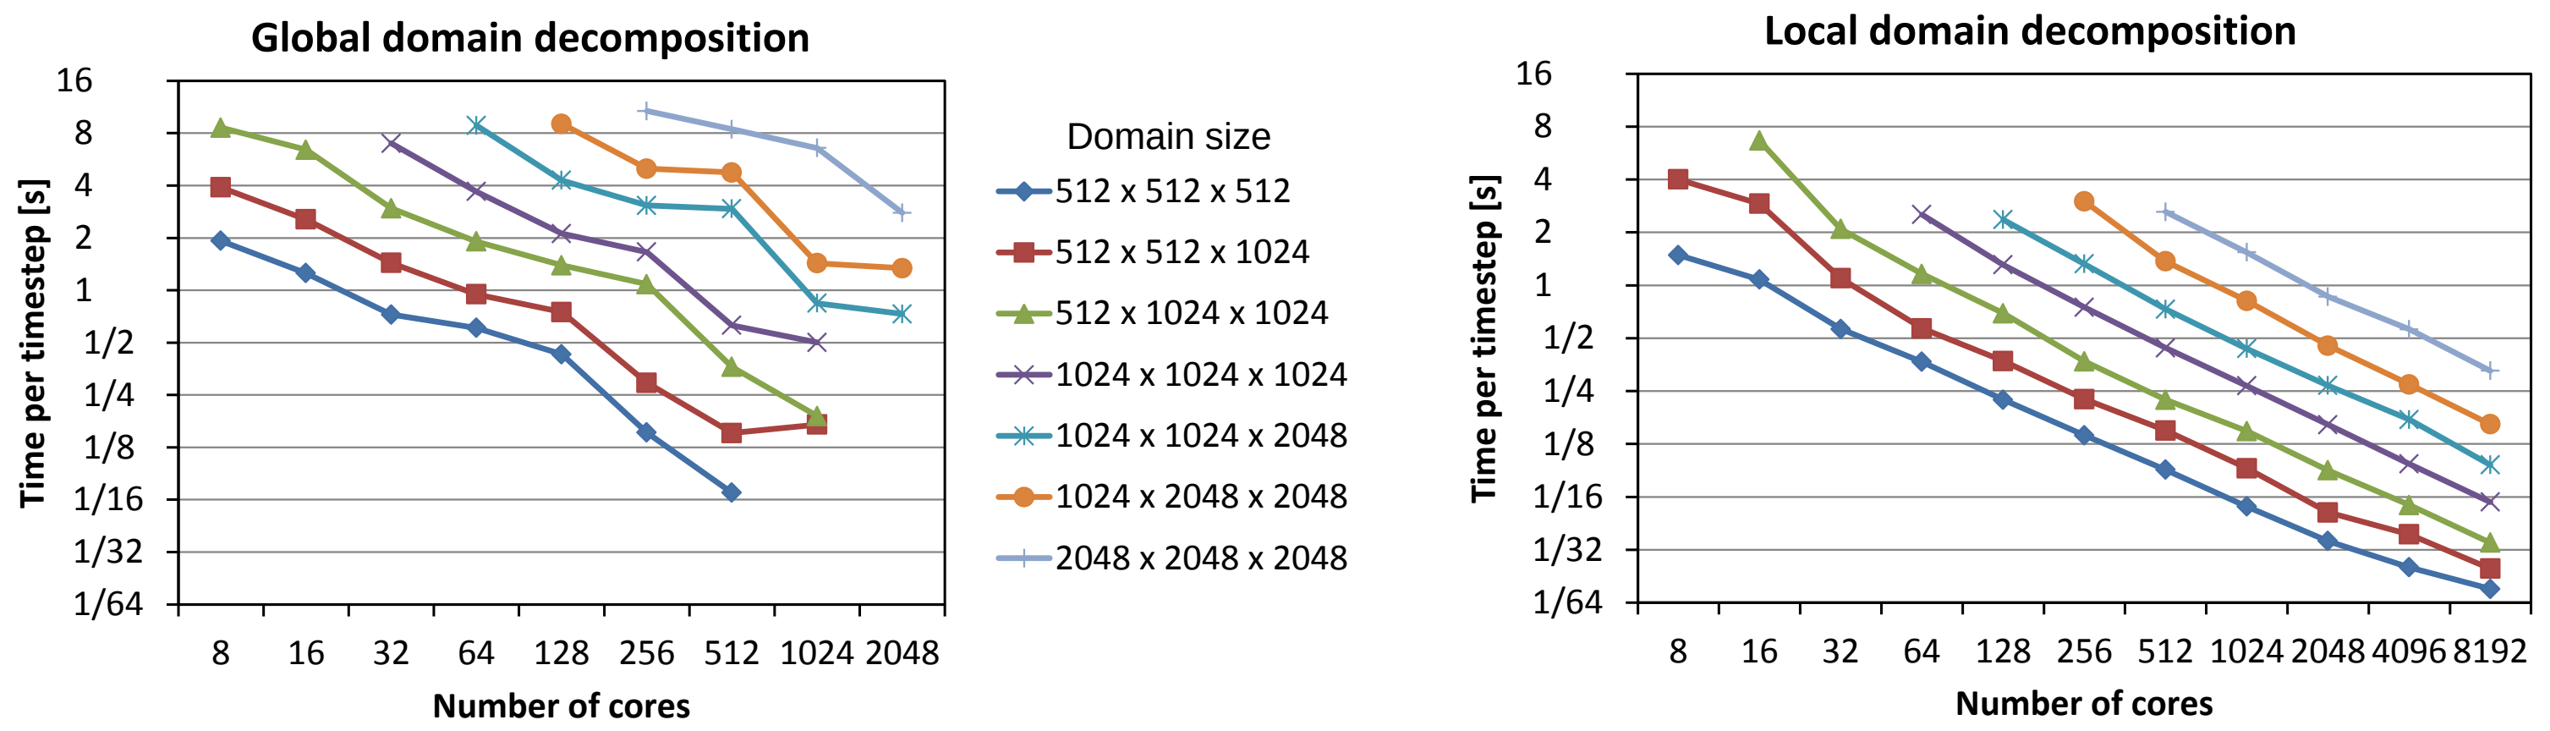

The local Fourier basis decomposition was implemented using both pure-MPI and hybrid OpenMP/MPI paradigms. The benefits mainly show for large domains offering better up to 8.55 times higher performance compared to the global decomposition.

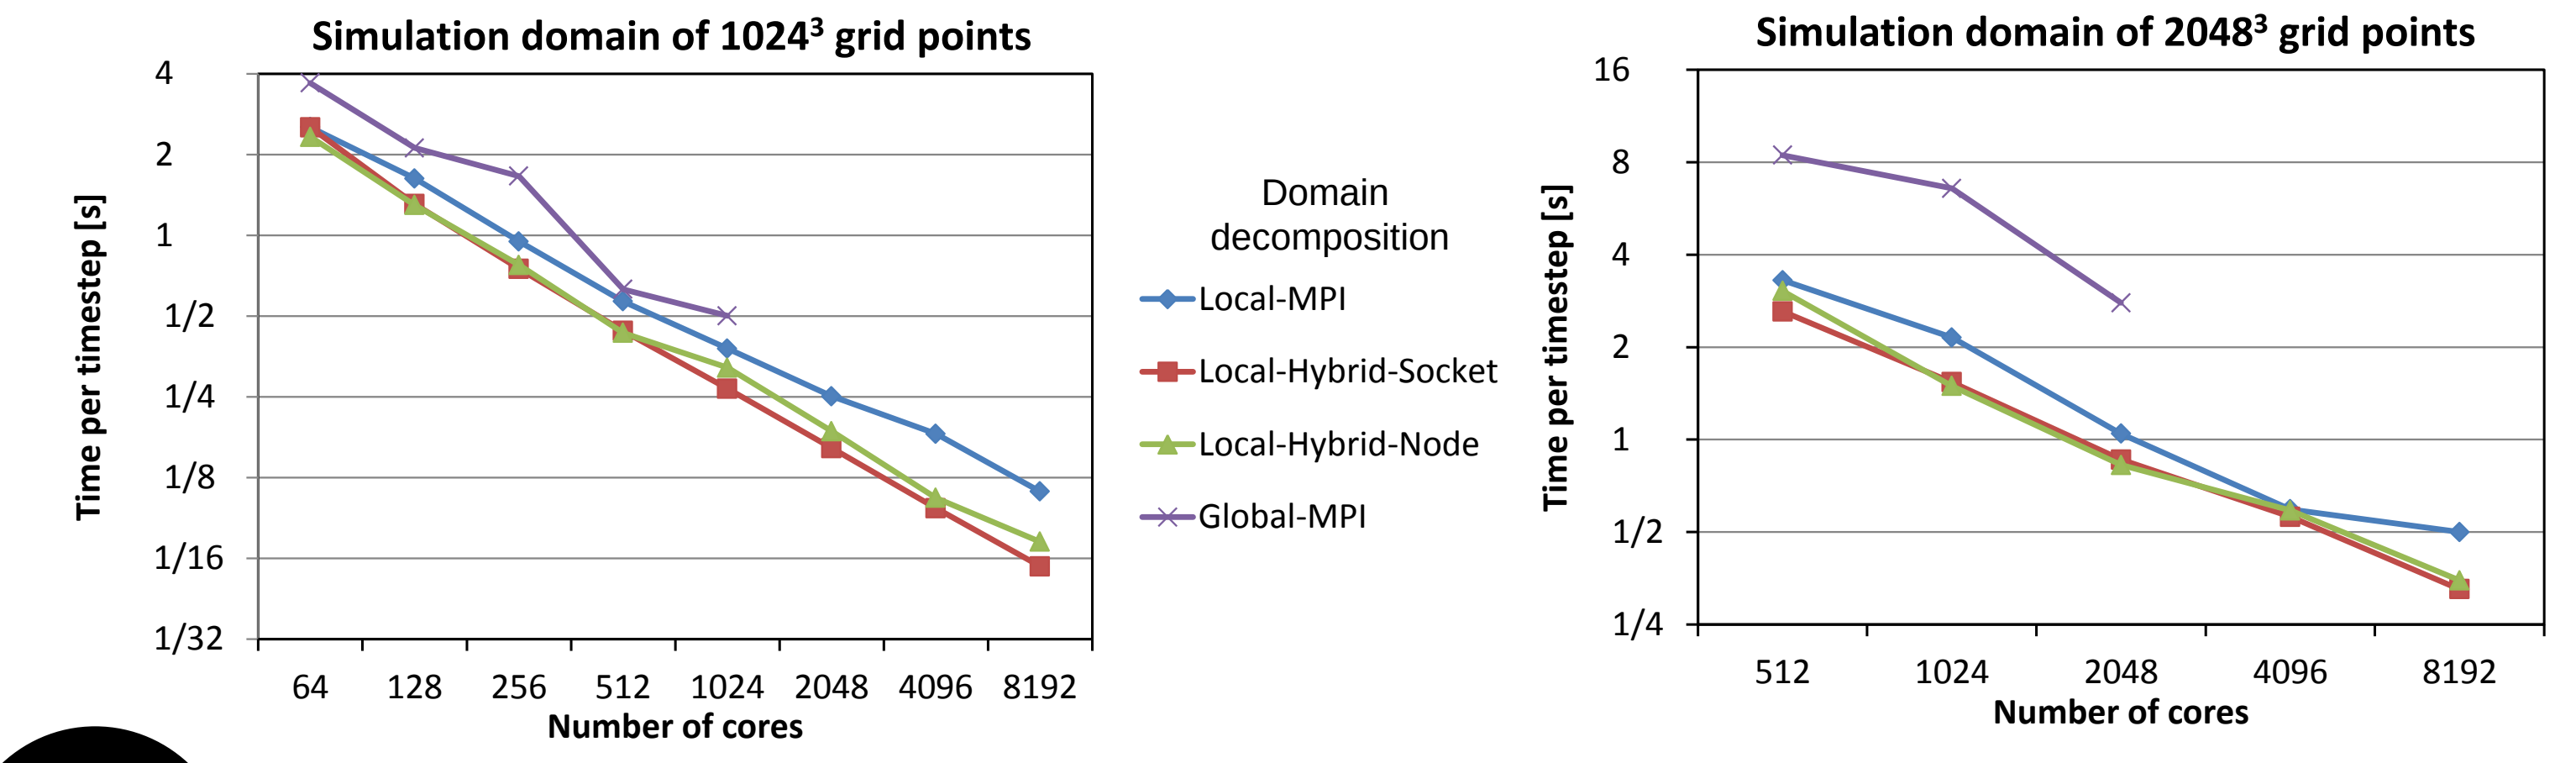

7 Result of Ultrasound Simulation

An example of a realistic HIFU kidney treatment simulation output (distribution of the maximum acoustic pressure).

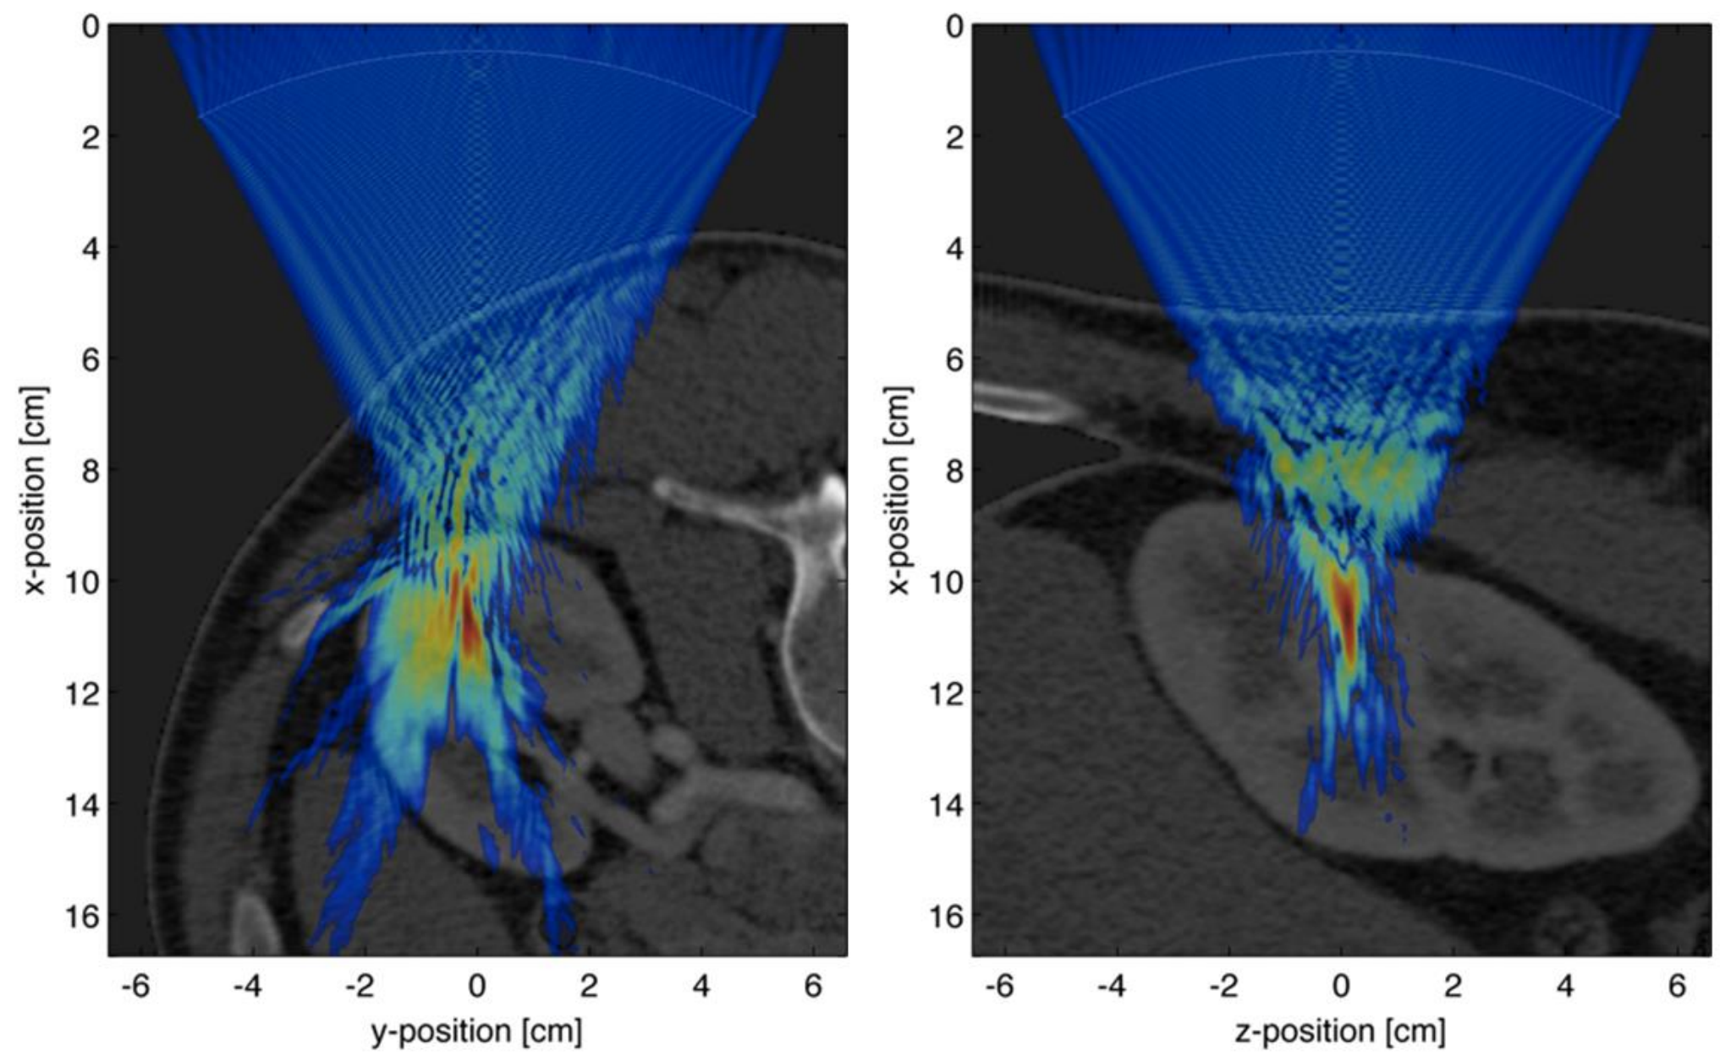

SoMoPro
South Moravian Region
Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities
jcmm
IT4Innovations national supercomputing center
MARIE CURIE ACTIONS

The project is financed from the SoMoPro II programme. The research leading to this invention has acquired a financial grant from the People Programme (Marie Curie action) of the Seventh Framework Programme of EU according to the REA Grant Agreement No. 291782. The research is further co-financed by the South-Moravian Region. This work reflects only the author's view and the European Union is not liable for any use that may be made of the information contained therein.
This work was supported by the IT4Innovations Centre of Excellence project (CZ.1.05/1.1.00/02.0070), funded by the European Regional Development Fund and the national budget of the Czech Republic via the Research and Development for Innovations Operational Programme, as well as Czech Ministry of Education, Youth and Sports via the project Large Research, Development and Innovations Infrastructures (LM2011033). The authors gratefully acknowledge the assistance of Sebastian Lehrack from the LMU Faculty of Physics, and the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputer SuperMUC at Leibniz Supercomputing Centre (LRZ, www.lrz.de).