Large-scale Ultrasound Simulations Using the Hybrid OpenMP/MPI Decomposition

Jiri Jaros*, Vojtech Nikl*, Bradley E. Treeby†

*Department of Compute Systems, Brno University of Technology

† Department of Medical Physics, University College London

Outline

- **Ultrasound simulations in soft tissues**
  - What is ultrasound
  - Why do we need ultrasound simulations
  - Factors for the ultrasound simulations
  - What is the challenge
- **k-Wave toolbox**
  - Acoustic model
  - Spectral methods
- **Large-scale ultrasound simulations**
  - 1D domain decomposition with pure-MPI
  - 2D hybrid decomposition with OpenMP/MPI
- **Achieved results**
  - FFTW scaling
  - Strong scaling
- **Conclusions and open questions**

What Is Ultrasound?

Longitudinal (compressional) acoustic waves …

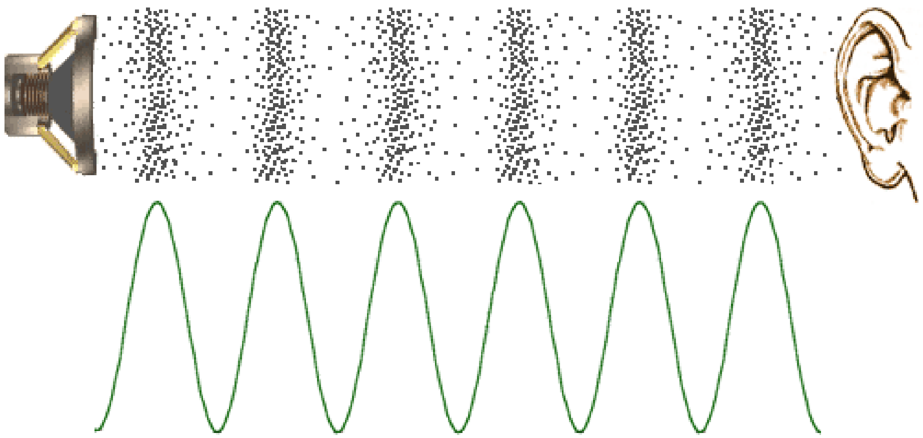

… with a frequency above 20 kHz

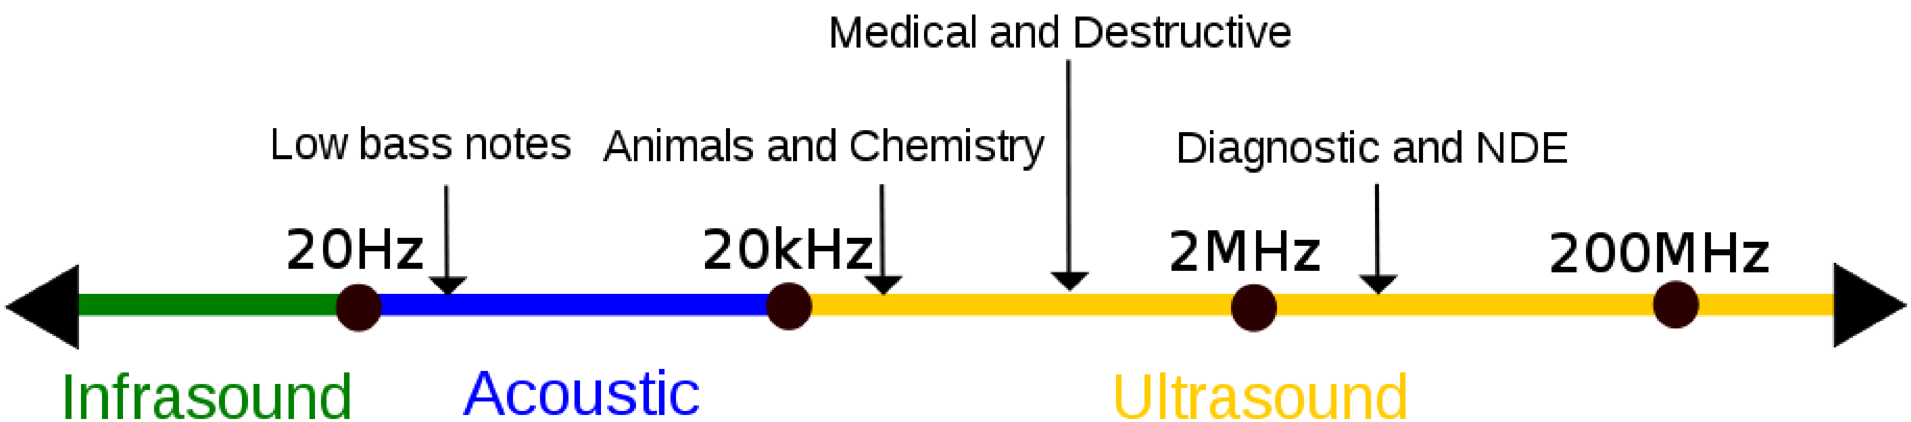

Ultrasound Simulation

- **Photoacoustic imaging**
- **Aberration correction**
- **Training ultrasonographers**
- **Ultrasound transducer design**
- **Treatment planning (HIFU)**

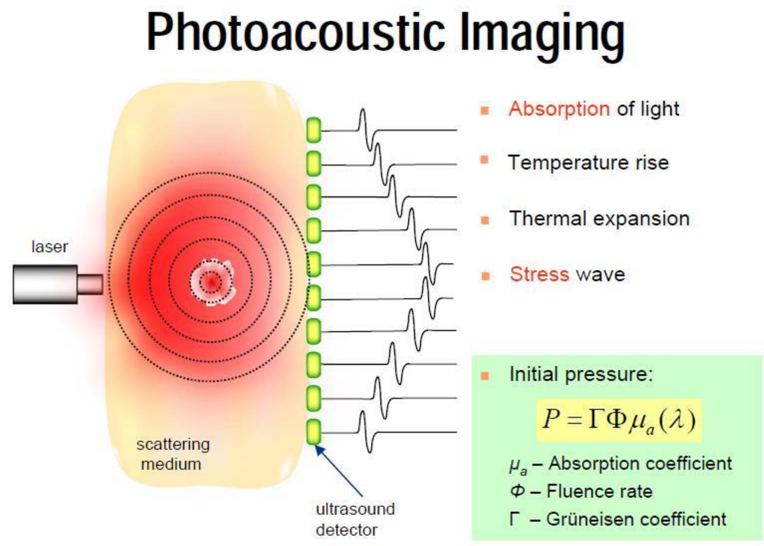

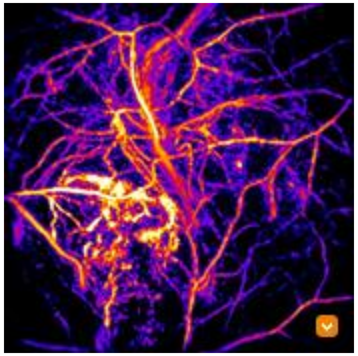

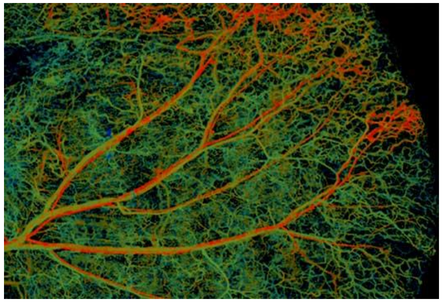

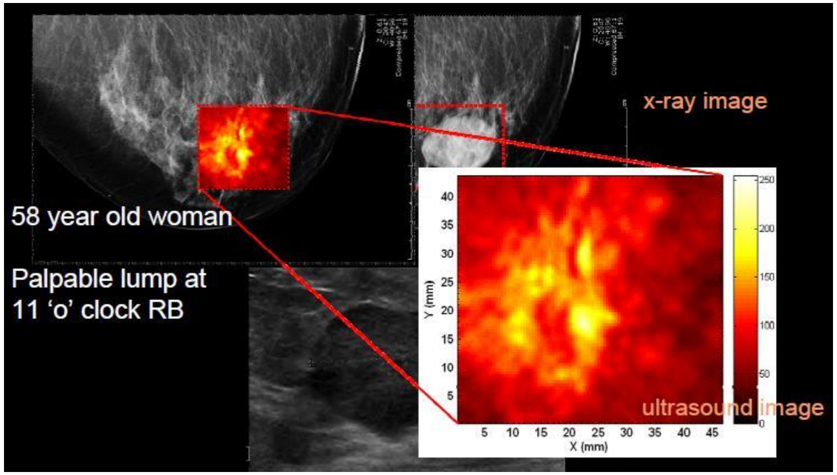

HIFU Treatment Planning

Scan
- CT or MR scan of a patient

Parameter setting
- Scan segmentation (bones, fat, skin, …)
- Medium parameters (density, sound speed)

Simulation
- Ultrasound propagation simulation
- Dosage, focus position, aberration correction

Operation
- Application of the ultrasound treatment

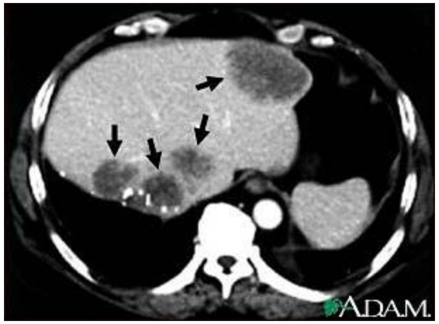

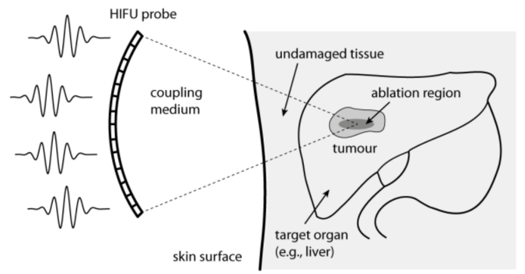

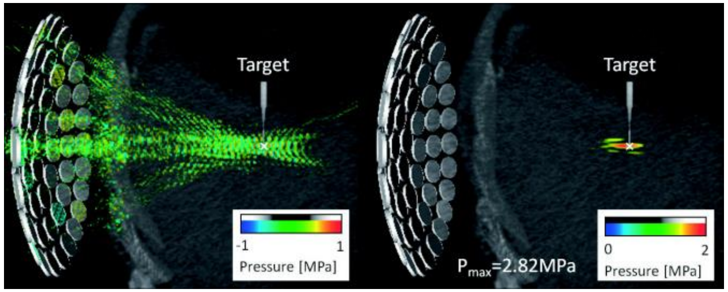

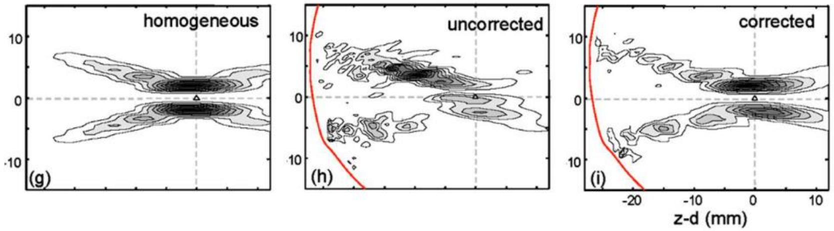

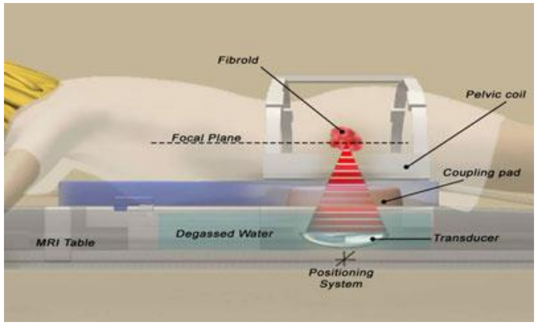

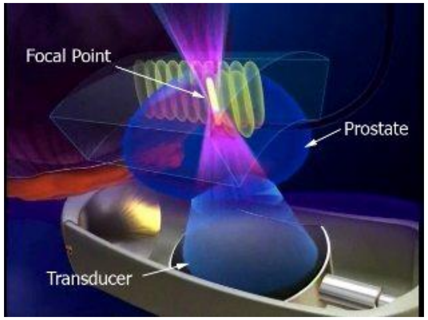

Factors for Ultrasound Simulation

- **Nonlinear wave propagation**
  - Production of harmonics
  - Energy dependent

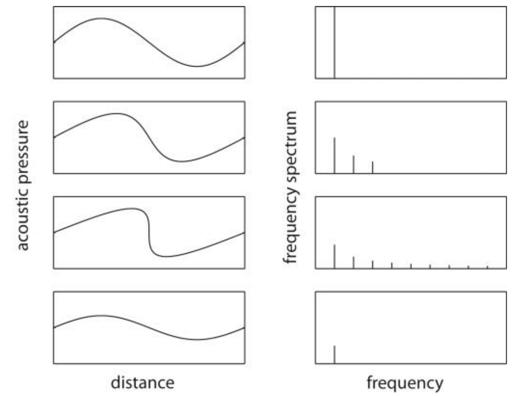

- **Heterogeneous medium**
  - Dispersion
  - Reflection

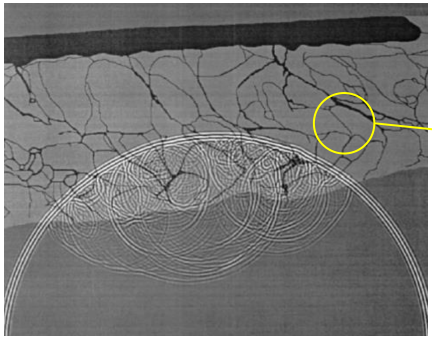

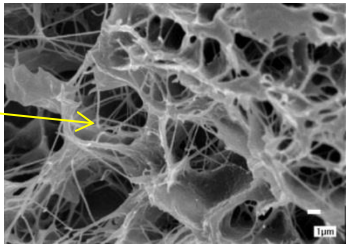

- **Absorbing medium**
  - Frequency dependent
  - Medium dependent

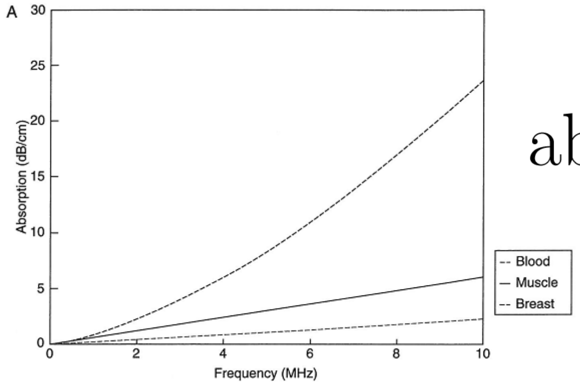

$$\mathrm{abs} = \alpha_0 f^y$$

How Big Simulations Do We Need?

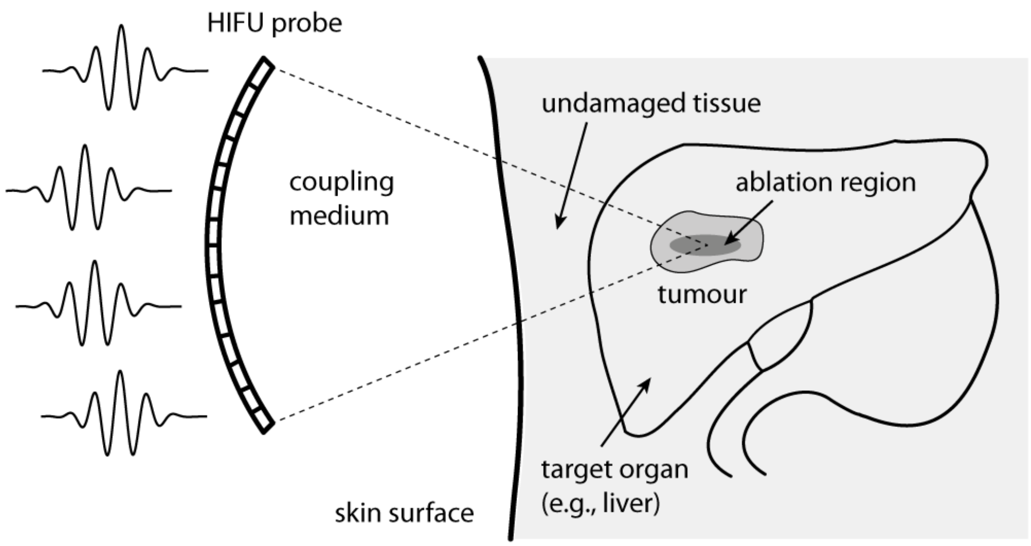

Speed of sound in water ≈ 1500 m/s

At 1 MHz, 20 cm ≈ 133 λ
At 10 MHz, 20 cm ≈ 1333 λ

At 15 grid points per wavelength, each matrix is **30 TB**!

Modeling Scenario	Source Freq [MHz]	Source Type	Nonlinear Harmonics	Max Freq [MHz]	Domain Size [mm]			Domain Size [Wavelengths]		
					X	Y	Z	X	Y	Z
Diagnostic Ultrasound: Abdominal Curvilinear Transducer	3	Tone Burst	5	18	150	80	25	1800	960	300
Diagnostic Ultrasound: Linear Transducer	10	Tone Burst	5	60	50	80	30	2000	3200	1200
Transrectal Prostate HIFU Minimal Cavitation	4	CW	15	64	80	60	20	3413	2560	853
MR-Guided HIFU Minimal Cavitation	1.5	CW	10	15	250	250	150	2500	2500	1500
Histotripsy Intense Cavitation	1	CW	50	50	250	250	150	8333	8333	5000

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

Acoustic Model for Soft Tissues

- **k-Wave Toolbox** (http://www.k-wave.org)
  - 3,385 registered users
- **Full-wave 3D acoustic model**
  - including nonlinearity
  - heterogeneities
  - power law absorption
- **Solves coupled first-order equations**

$$\left\{\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} &= -\frac{1}{\rho_0}\nabla p \\ \frac{\partial \rho}{\partial t} &= -\left(2\rho + \rho_0\right)\nabla\cdot\mathbf{u} \\ p &= c_0^2\left(\rho + \frac{B}{2A}\frac{\rho^2}{\rho_0} - \mathrm{L}\rho\right) \end{aligned}\right.$$

momentum conservation

mass conservation

pressure-density relation

$$\mathrm{L} = \tau\frac{\partial}{\partial t}\left(-\nabla^2\right)^{\frac{y}{2}-1} + \eta\left(-\nabla^2\right)^{\frac{y+1}{2}-1}$$

absorption term

k-space Pseudospectral Method in C++

- **Technique**
  - Medium properties generated by Matlab scripts from a medical scan.
  - Input signal is injected by a transducer.
  - Sensor data is collected in the form of raw time series or aggregated acoustics values.
  - Post processing and visualization handled by Matlab.

- **Operations executed in every time step**
  - 6 forward 3D FFTs
  - 8 inverse 3D FFTs
  - *3+3 forward and inverse 1D FFTs in the case of non-staggered velocity*
  - About 100 element wise matrix operations (multiplication, addition,…)

- **Global data set**
  - 14 +3 (scratch) + *3 (unstaggering)* real 3D matrices
  - 3+3 complex 3D matrices
  - 6 real 1D vectors
  - 6 complex 1D vectors
  - *Sensor mask, source mask, source input*
  - *<0 , 20> real buffers for aggregated quantities (max, min, rms, max_all, min_all)*

- **Implementation language**
  - C/C++ and MPI parallelization
  - MPI-FFTW library– efficient way to calculate distributed 3D FFTs
  - HDF5 library – hierarchical data format for parallel I/O

- **Data decomposition**
  - Data decomposed along the Z dimension
  - Data distributed when read using parallel I/O
  - Frequency domain operations work on transposed data to reduce the number of global communications (3D transpositions).

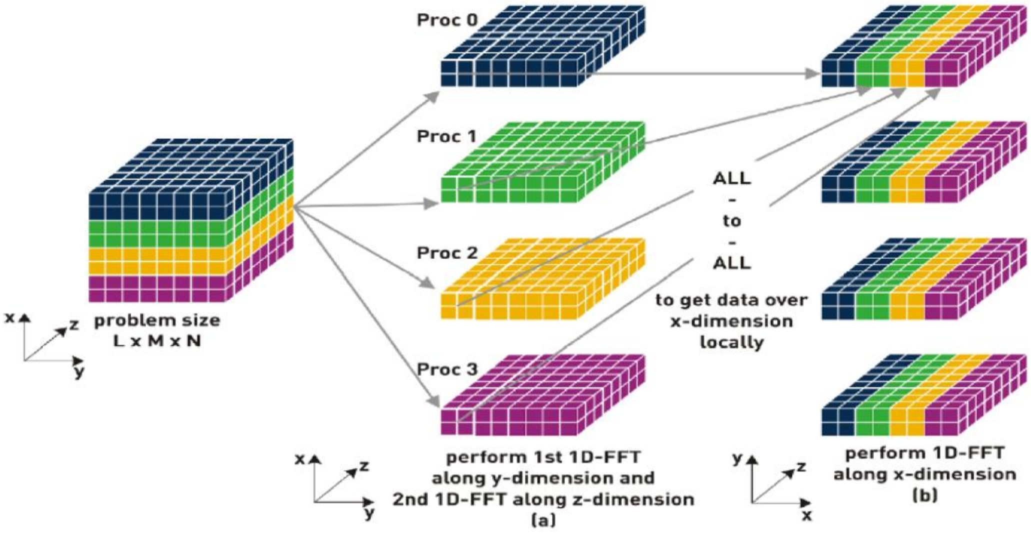

K-Wave++ Toolbox
Strong Scaling (1D decomposition)

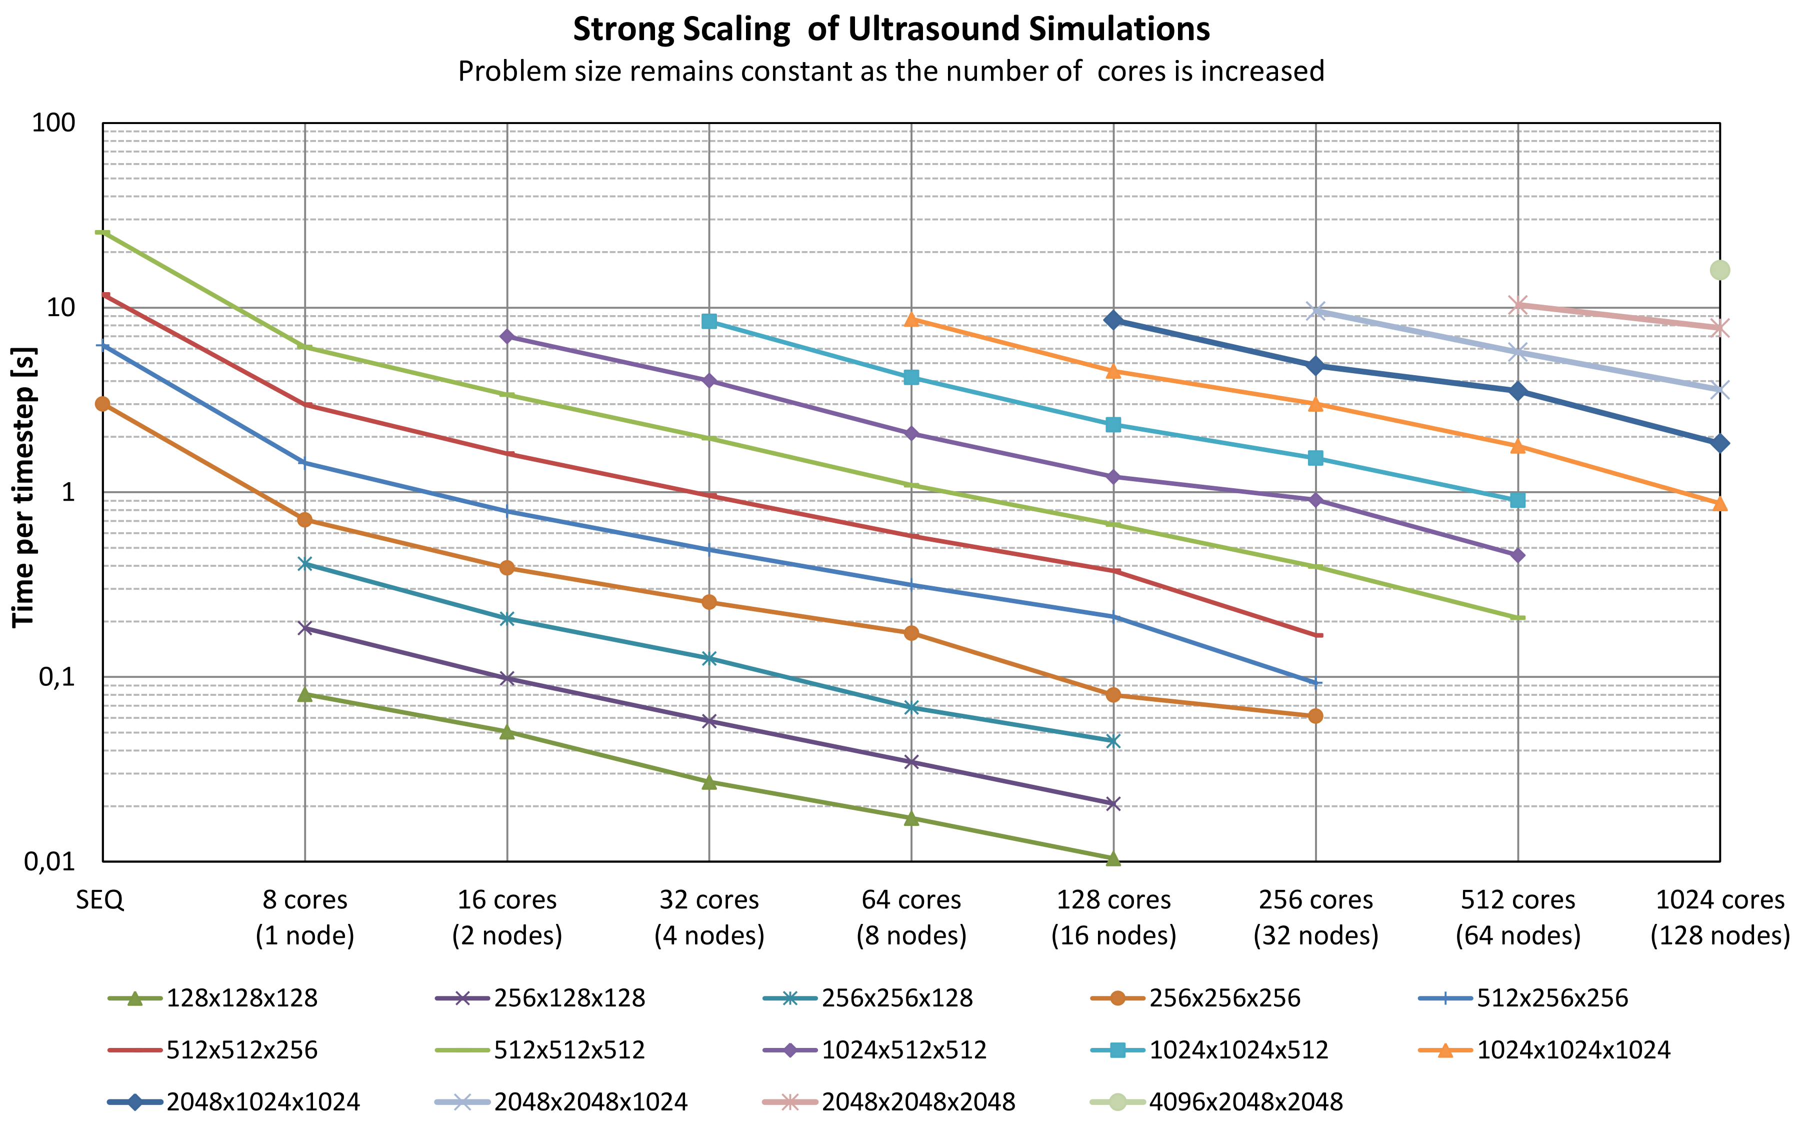

Scalability Problem

- **1D decomposition**
  – The number of cores is limited by the largest dimension
  – It makes some simulation run for too long
  – It makes some simulation not fit into memory
  + It requires less communication
- **2D decomposition**
  + The number of codes is limited by a product of two largest dimensions
  + It is enough for anything we could think of running
  – It requires more communication and smaller messages

- **Example**
  – 4096^3 matrix -> ~256GB in single precision
  – Say 32 matrices -> 8TB RAM in total
  – Max 4096 cores -> 2GB per core (Anselm – 2GB, Fermi – 1GB)

2D Hybrid Domain Decompostion

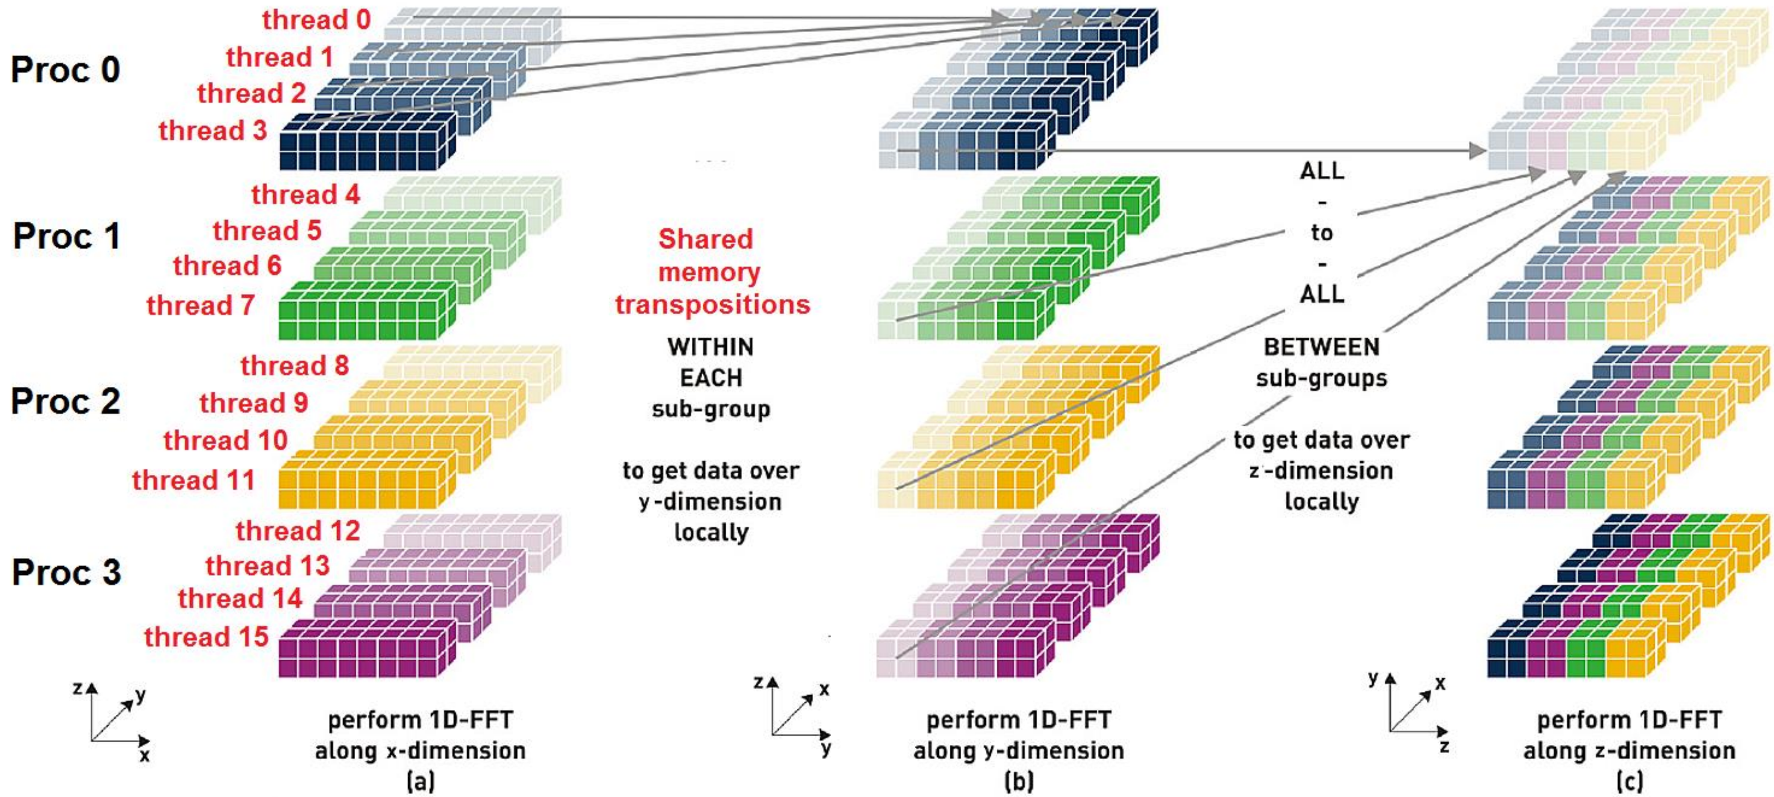

- + high core count limit
- + only 1 MPI global transposition
- + lower number of larger MPI messages

FFT libraries strong scaling Anselm supercomputer (1024^3)

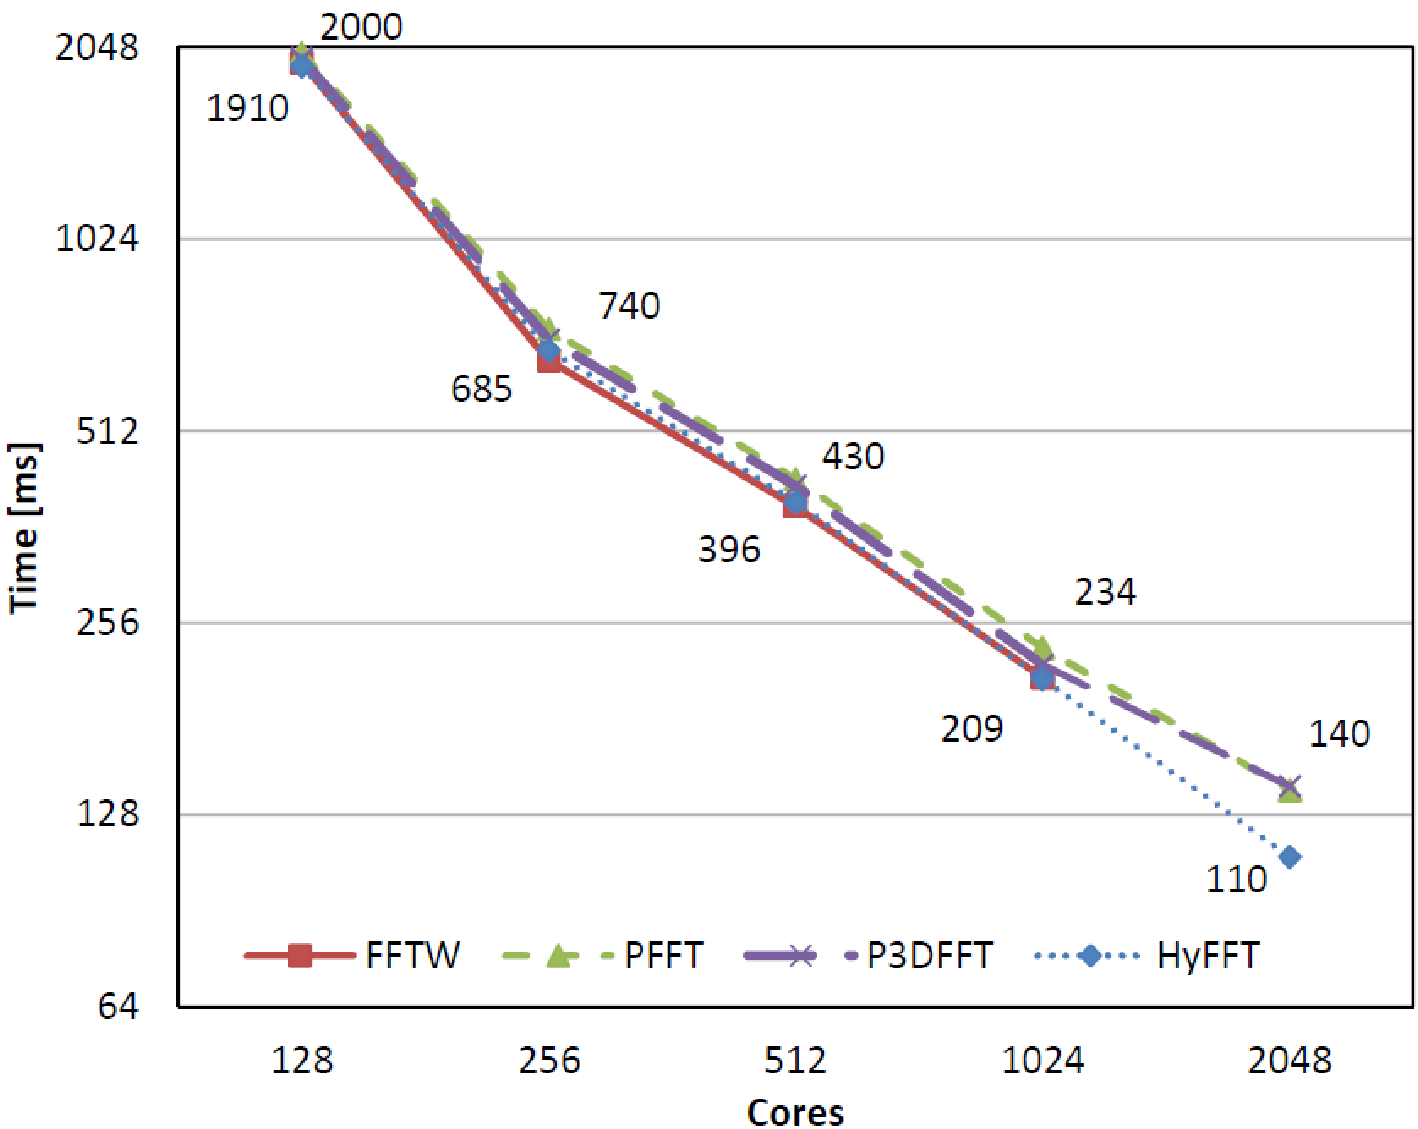

FFT libraries strong scaling
Fermi supercomputer (1024^3)

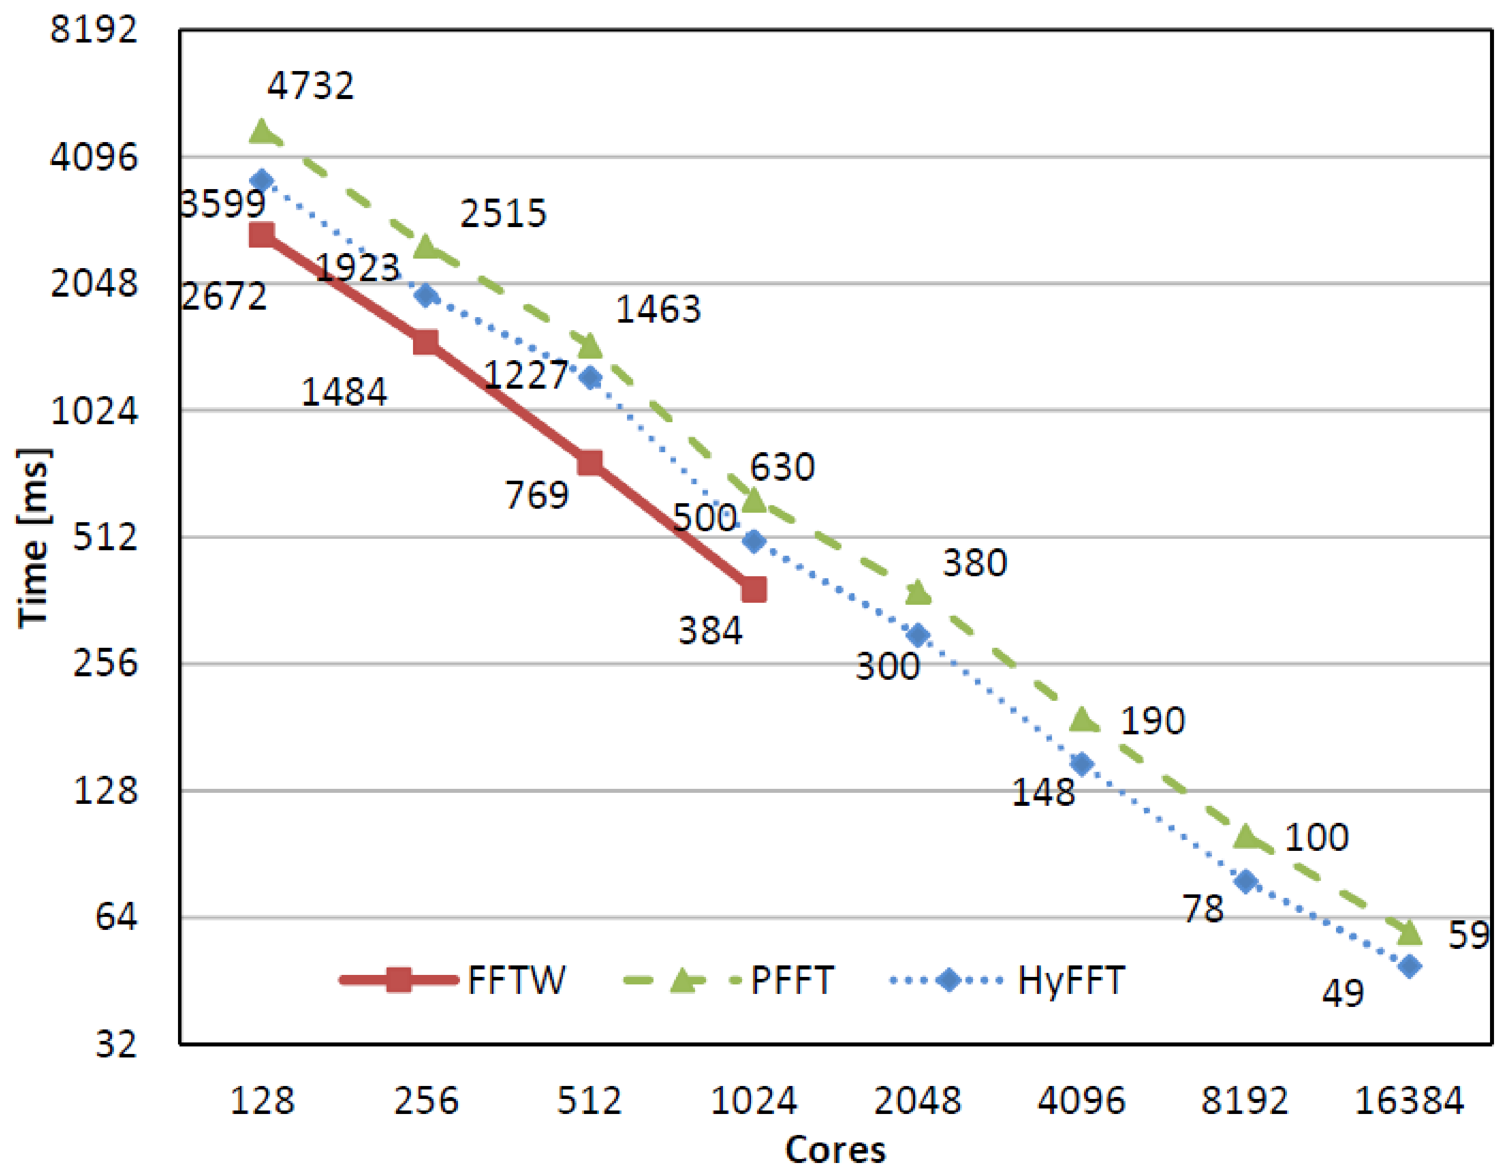

k-Wave
A MATLAB toolbox for the time-domain simulation of acoustic wave fields

Time distribution of hybrid FFT

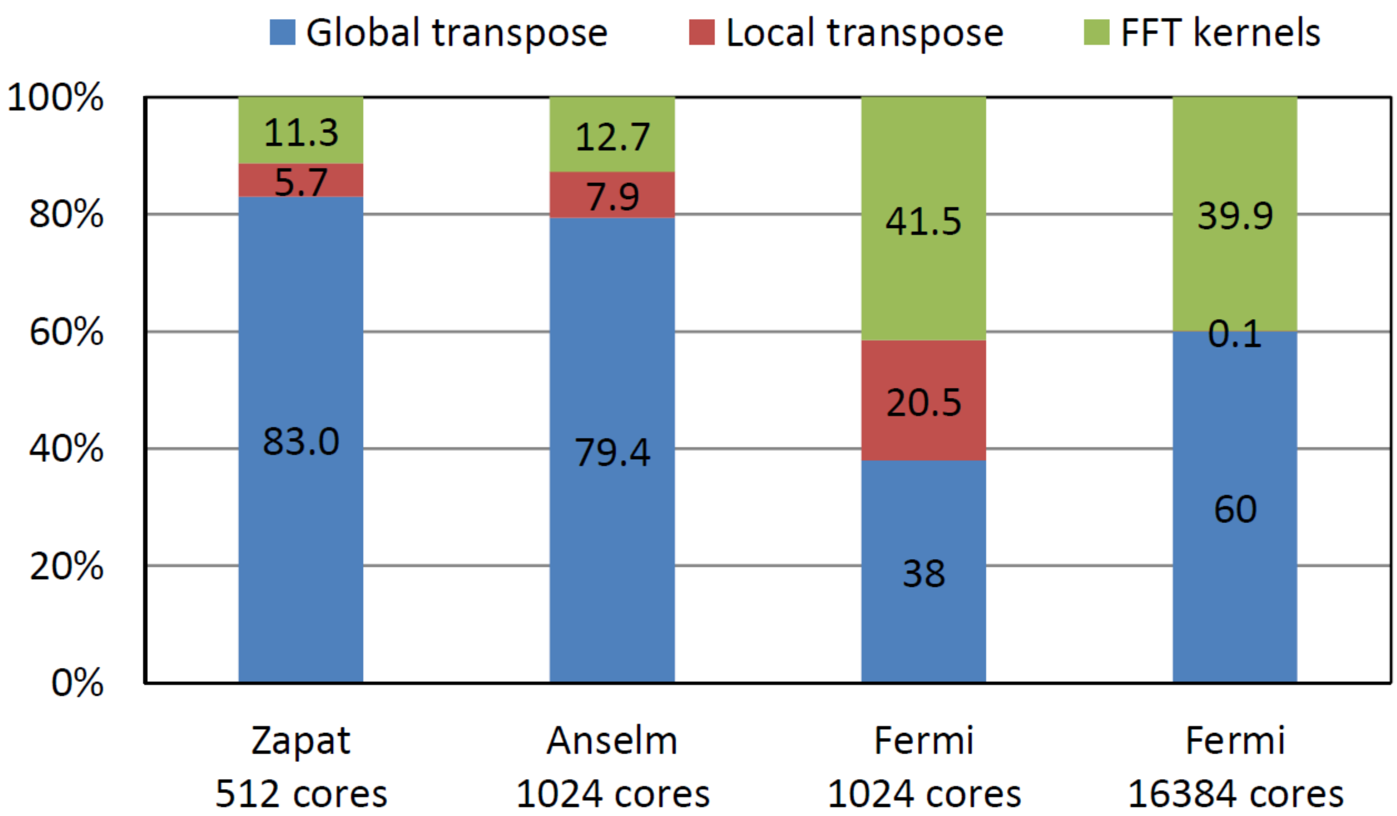

k-Wave
A MATLAB toolbox for the time-domain simulation of acoustic wave fields

Simulation scaling (Anselm)

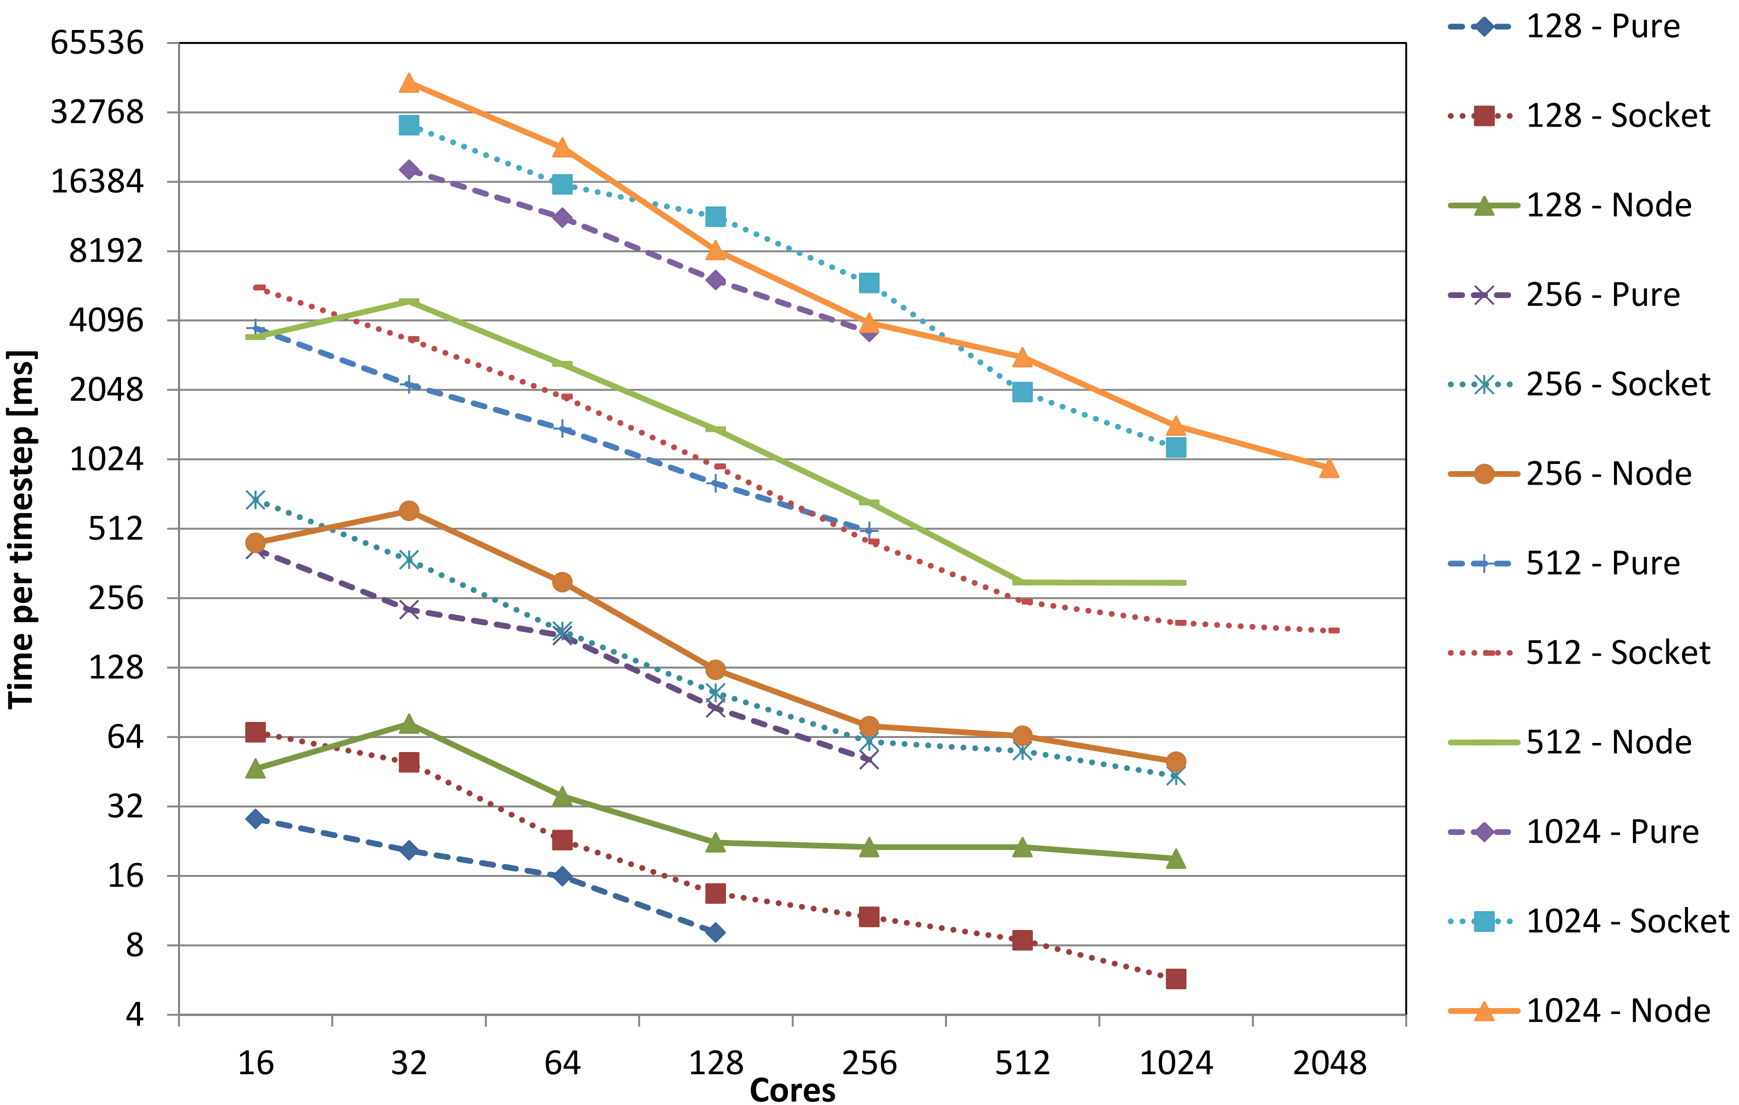

Simulation scaling (SuperMUC)

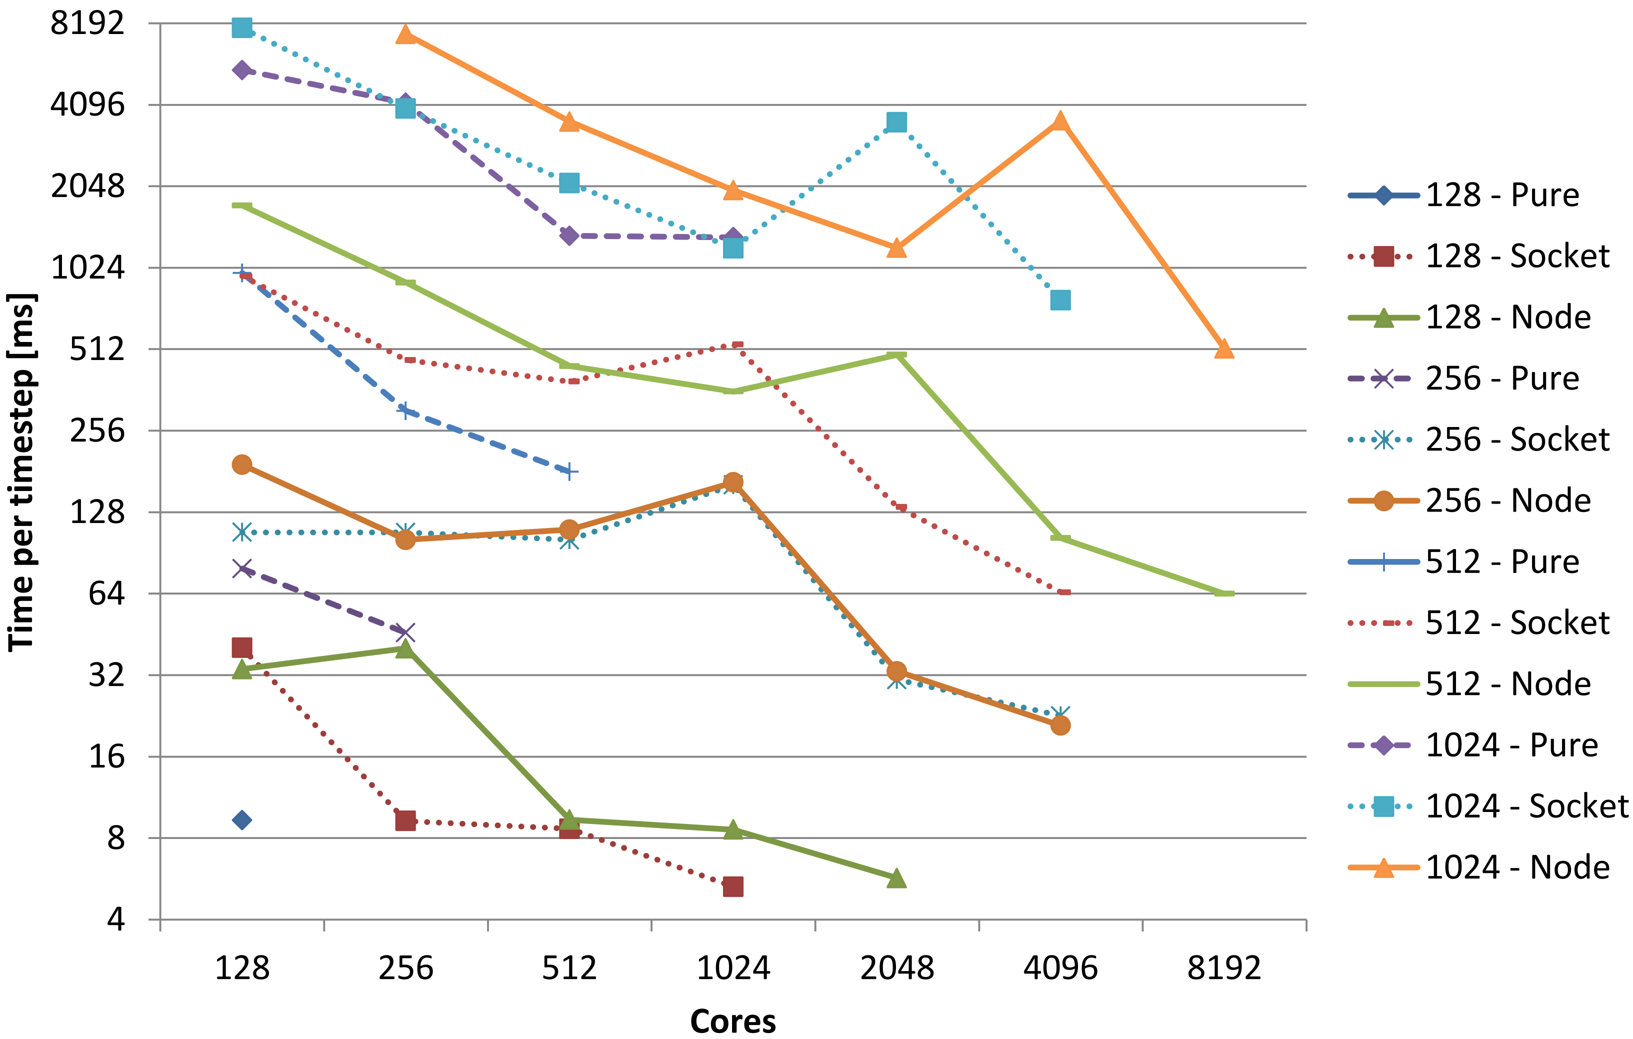

Memory scaling (SuperMUC)

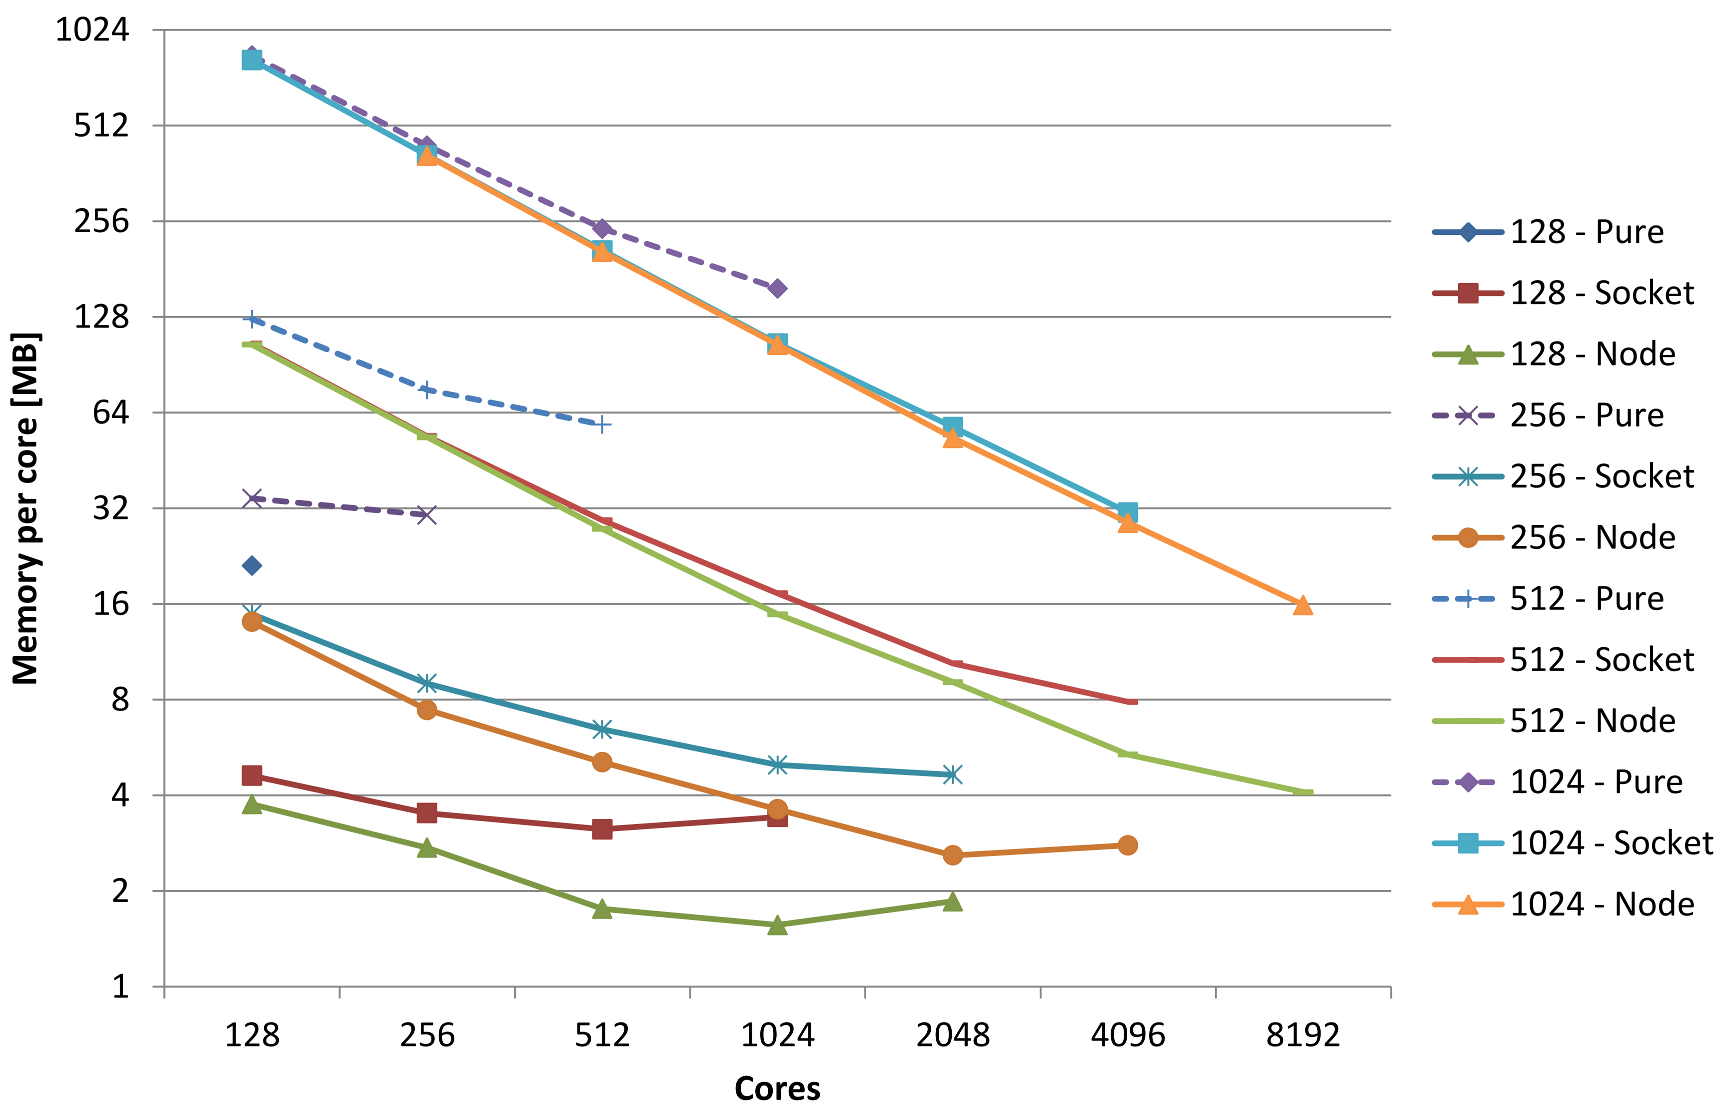

Conclusions

- **Clinical relevant results**
  - To get clinically relevant simulation we need grid sizes of 4096^3 to 8192^3 at least for 50k simulation timesteps
- **Two different Implementations**
  - 1D domain decomposition gives better results for small core counts
  - 2D domain decomposition works well on Anselm, however there is a room for improvement on SuperMUC
  - Memory scaling enables us to run much bigger simulations
- **Future work**
  - Communication and synchronization reduction via overlapping

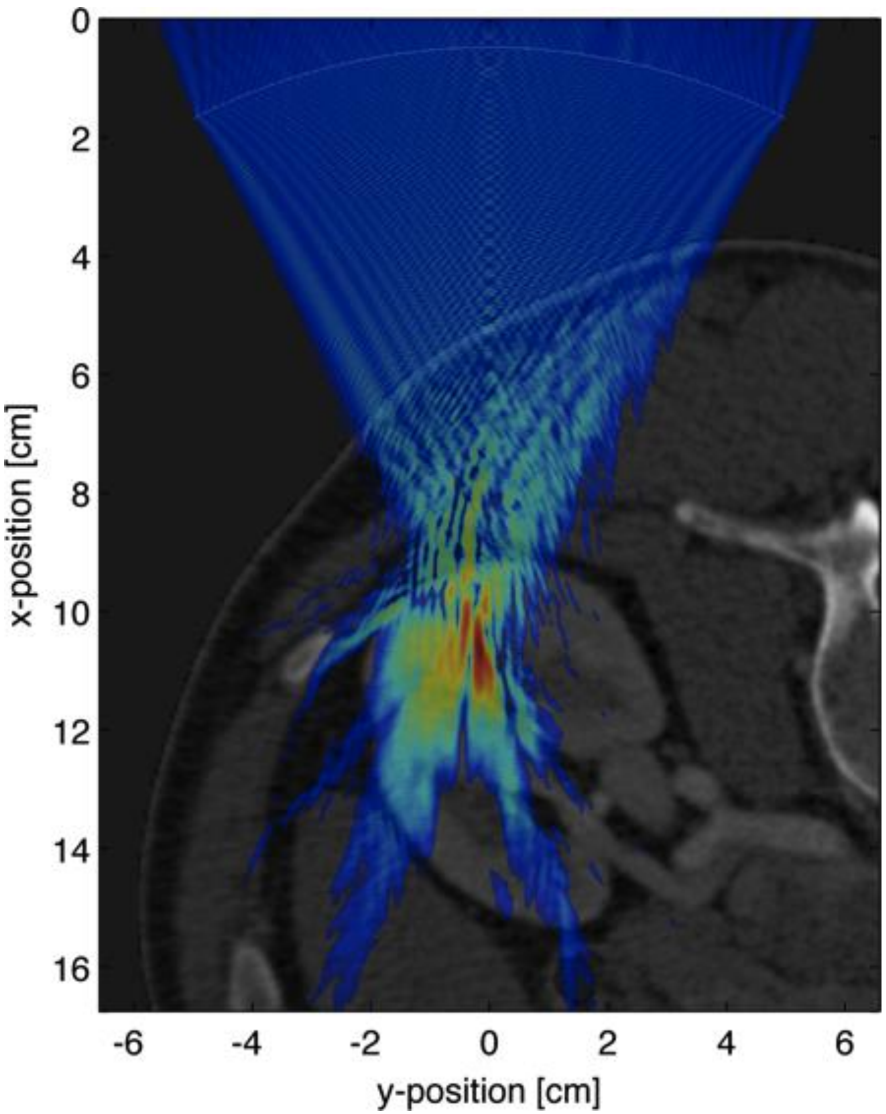

Questions and Comments

Our work has been supported by following institutions and grants

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 also supported by the research project "Architecture of parallel and embedded computer systems", Brno University of Technology, FIT-S-14-2297, 2014-2016.

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).

We acknowledge CINECA and PRACE Summer of HPC project for the availability of high performance computing resources.

The authors gratefully acknowledge 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).