Optimization of Ultrasound Simulations on Multi-GPU Servers

Filip Vaverka[1], Matej Spetko[1], Bradley E. Treeby[2] and Jiri Jaros[1]

[1]Faculty of Information Technology, Brno University of Technology, Centre of Excellence IT4Innovations, CZ
[2]Department of Medial Physics and Biomedical Engineering, University College London, UK

BRNO UNIVERSITY OF TECHNOLOGY FACULTY OF INFORMATION TECHNOLOGY

UCL

Overview

Photoacoustic tomography is a rapidly emerging imaging modality based on the generation of ultrasound waves using pulsed laser light. To reconstruct a photoacoustic image, an inverse problem is solved to estimate the laser induced initial pressure distribution in the tissue from the ultrasound signals recorded at the tissue surface. Our approach to this problem employs iterative solvers which require both forward and backward ultrasound wave propagation simulation in each iteration. This creates demand for high performance US propagation solvers in moderately sized domains ($\sim 10^9$ unknowns).

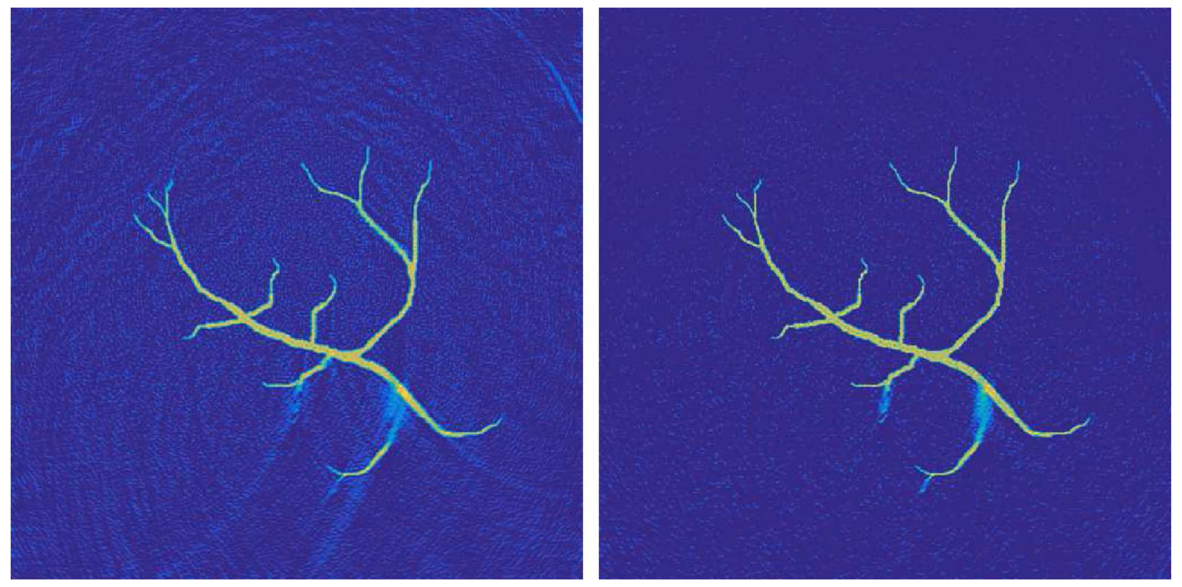

Ultrasound Wave Propagation in Tissue

The governing equations model ultrasound wave propagation in tissue, which is a heterogeneous and absorbing medium. Accounting for accurate acoustic absorption and reflection is crucial for photoacoustic imaging. The governing acoustic equations can be written as:

$$\frac{\partial \mathbf{u}}{\partial t} = -\frac{1}{\rho_0}\nabla p + \mathbf{S}_F \qquad \text{(momentum conservation)}$$

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

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

These equations are discretized using local fourier basis which allows efficient decomposition across multiple GPUs by partitioning the domain into a grid of local subdomains where gradients are calculated locally and the global communication is replaced by nearest-neighbor overlap exchange. The gradient calculation on each sub-domain then reads as follows, where b is a bell function smoothing the subdomain interface:

$$\frac{\partial p}{\partial x} = \mathscr{F}^{-1}\left\{ik\mathscr{F}\left\{b \circ p\right\}\right\}$$

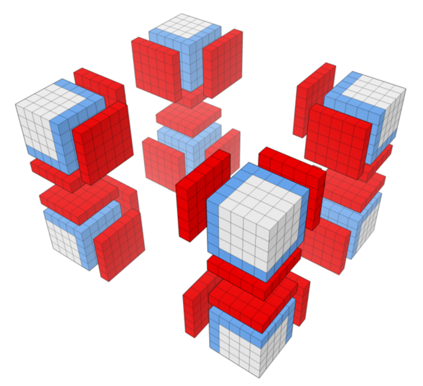

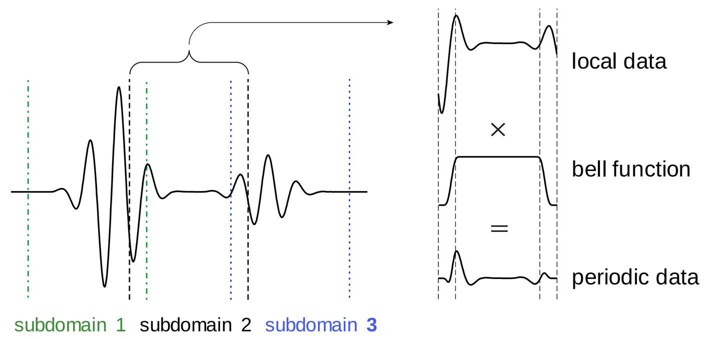

Simulation Accuracy

Since the gradient is not calculated on the whole data, numerical error is introduced. The error level can be controlled by the shape of the bell function and the size of the overlap region.

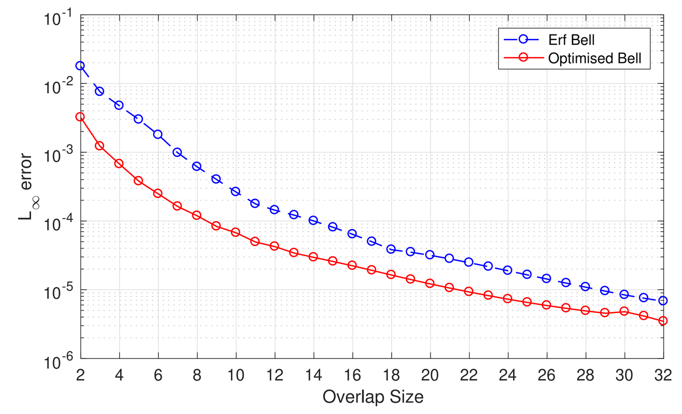

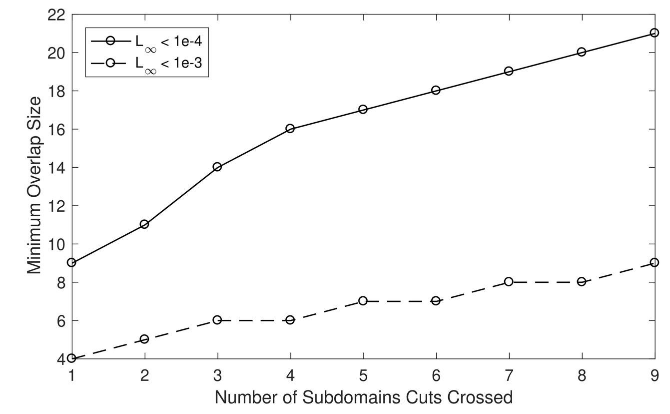

Dense Multi-GPU System

The architecture of an 8-GPU dual socket dense compute node used for our experiments offers up to 96 Tflop of compute power coupled with 192 GB of high-speed memory.

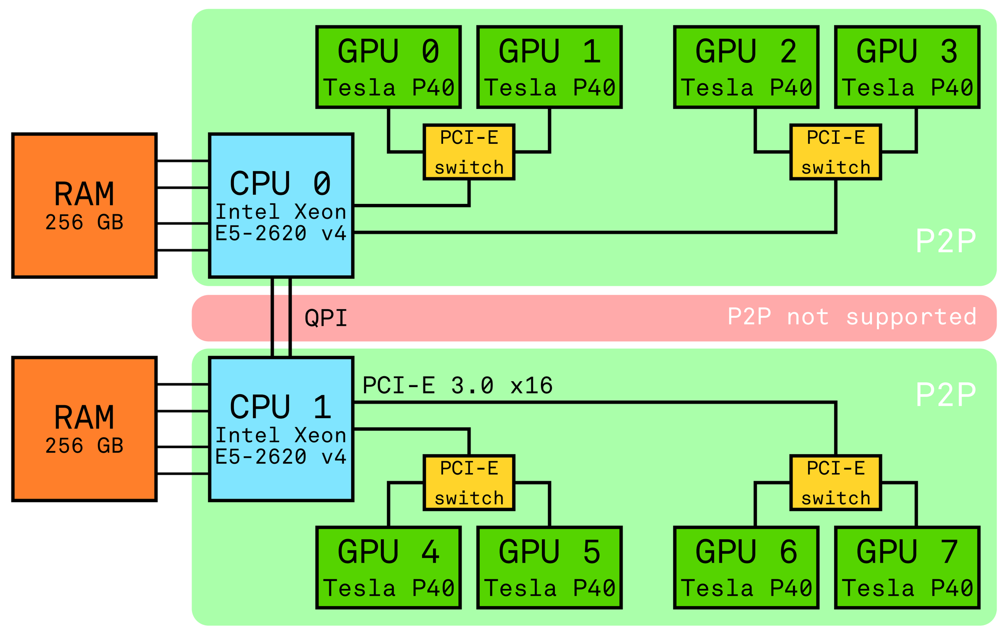

Communication Patterns

Naive combinations of CUDA with MPI result in many temporary intermediate copies of the data being transferred between GPUs. On the top of that all data transfers have to go through system memory and CPU which leads to inneficient use of the total system throughput. In contrast, CUDA-Aware MPI (possibly with explicit P2P) allows to eliminate most of intermediate buffers and copies. The only remaining buffers are those required by CUDA-Aware MPI for transfers over network or QPI interconnect.

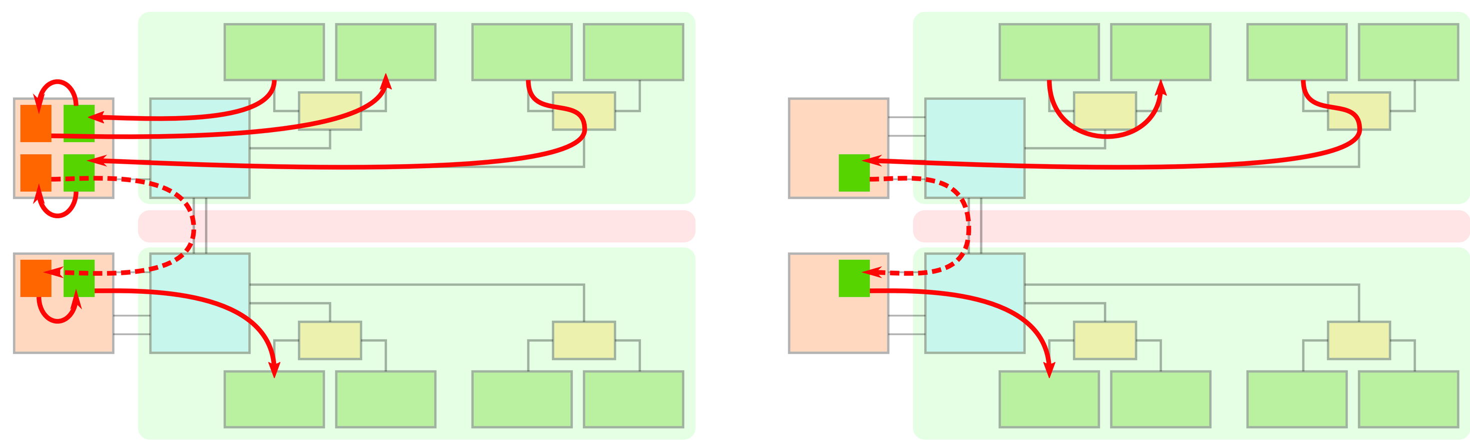

Naive CUDA+MPI Performance

The implementation of GPU-to-GPU data transfers without the native CUDA support in MPI introduces significant overheads (which is much more pronounced in dense multi-GPU systems).
The performance of the simulation is heavily impacted by the overlap size and to lesser degree by a chosen decomposition strategy.

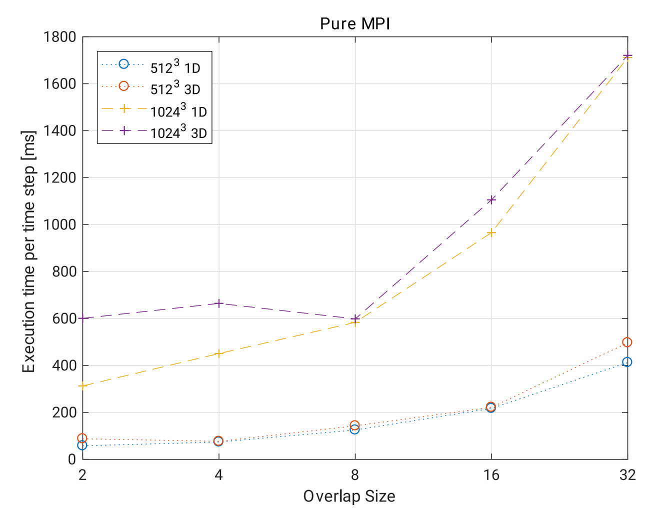

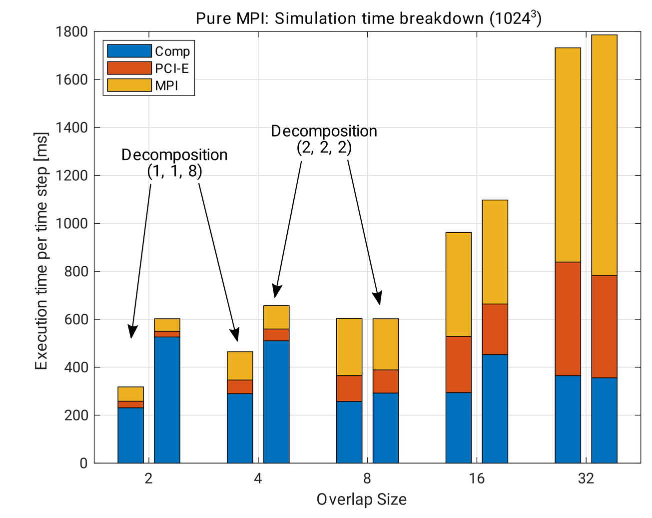

Hybrid CUDA-Aware MPI + P2P

Significant speedups (up to 3.6×) are achieved when the combination of CUDA-Aware MPI and CUDA Peer-to-Peer transfers is employed. It is also necessary to map processes and GPUs optimally so that the communication through QPI is minimized. Only about 6% of measured speedup comes from CUDA Peer-to-Peer transfers.

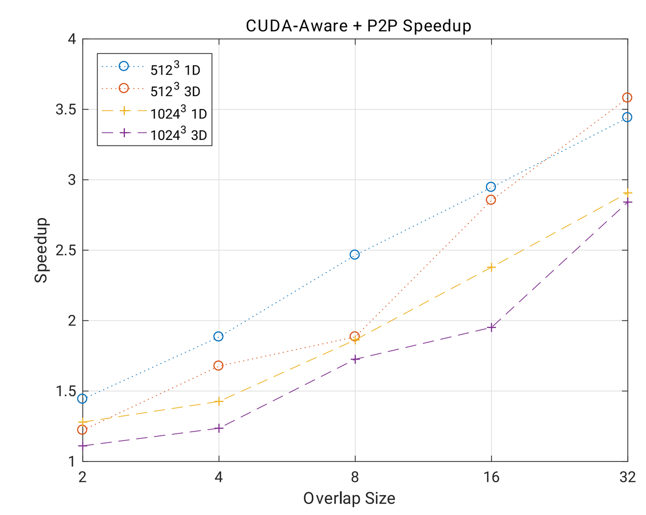

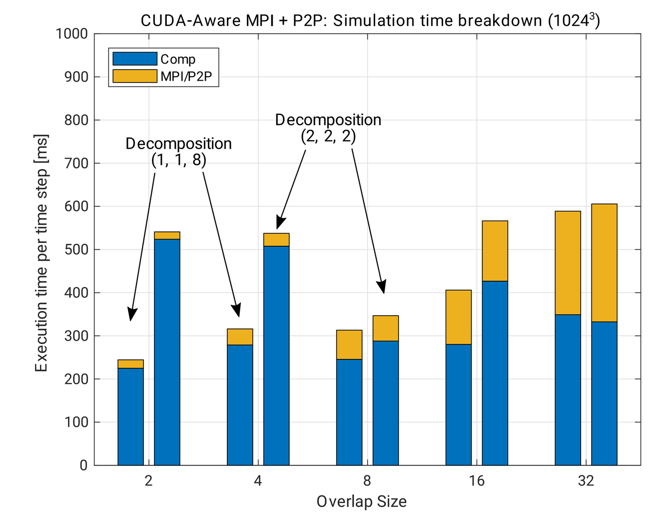

Impact and Outlook

Typical photoacoustic tomography image reconstruction on a domain of 8 dm^3 with a maximum frequency of 2 MHz requires up to 50 simulations, each of which with 5500 time-steps on a domain of 1024^3 grid points.

This translates to about 88 hours of computation on the Salomon cluster using 32 nodes, 24 Haswell cores each. Such an image reconstruction would cost about $3,855.

Using our local Fourier Basis approach optimized for multi-GPU servers, we are able to reduce this time to 35 hours, and more importantly, reduce the cost to $893.

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

PHOTONICS21

IT4Innovations national supercomputing center

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 also supported by the Engineering and Physical Sciences Research Council, UK, grant numbers EP/L020262/1.