Accelerating Hybrid Local Domain Decomposition for the k-Wave Toolbox on Multi-GPU Systems

Oliver Kuník and Jiri Jaros

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

Overview

The *k-Wave* toolbox is designed for simulating the time-domain propagation of acoustic waves, utilizing the Fourier collocation method to calculate spatial derivatives. While this method is highly accurate, it incurs significant communication overhead when implemented on multi-CPU or multi-GPU systems. To address this issue, we divided the simulation domain into multiple subdomains, allowing spatial derivatives to be computed locally within each subdomain.

Hybrid Local Domain Decomposition

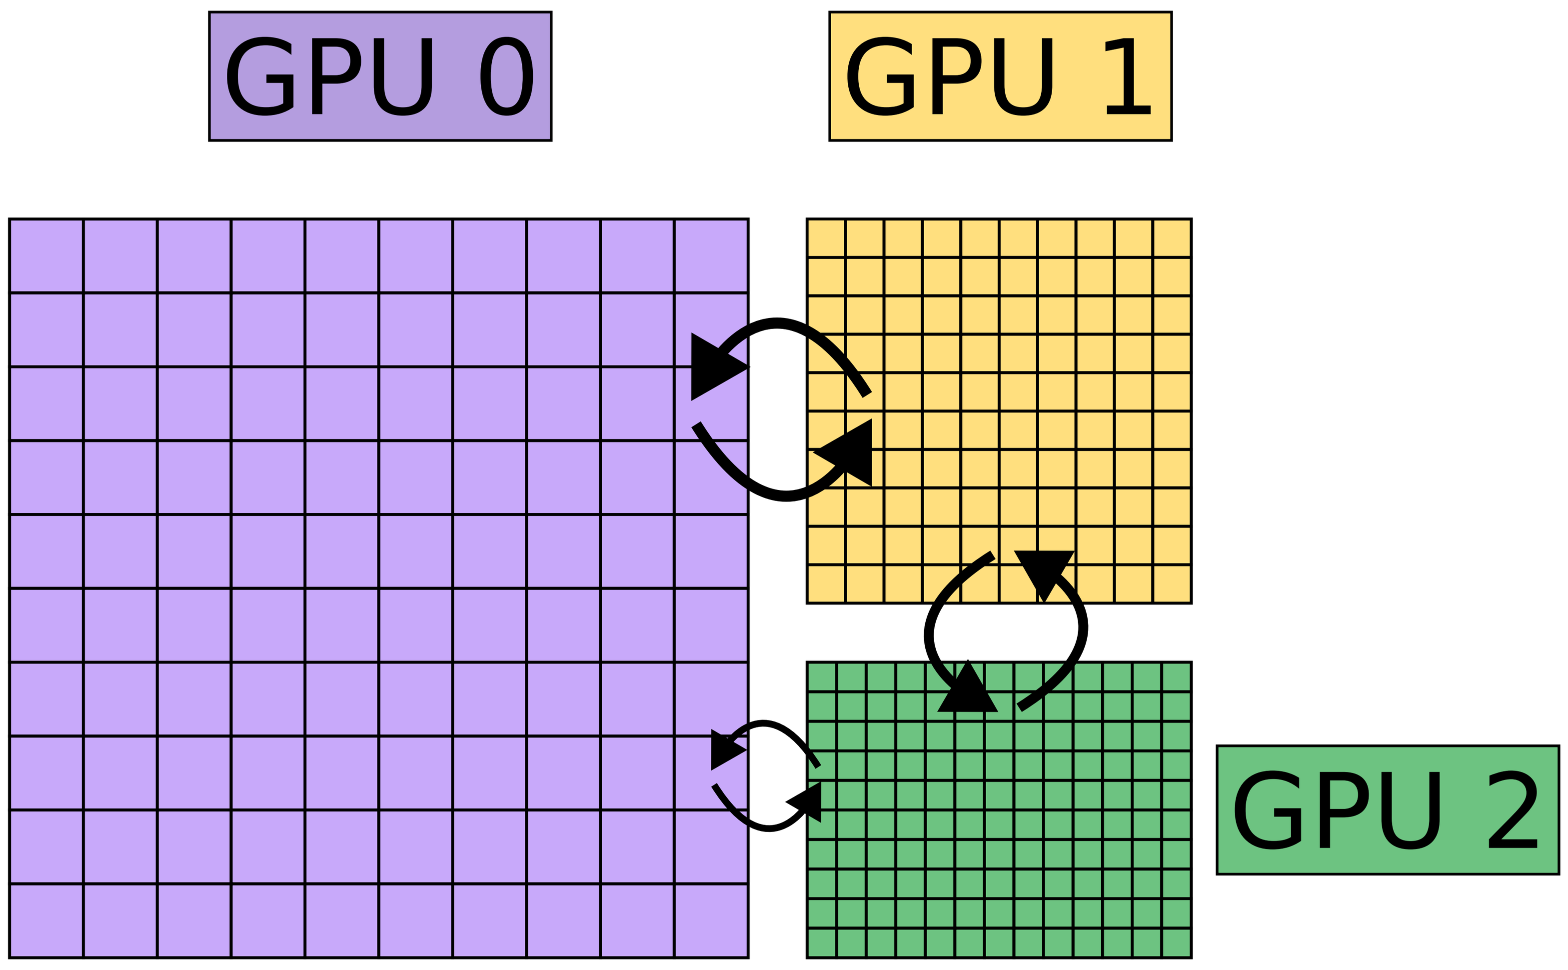

We run the simulation across multiple GPUs, allowing users to specify how the domain is divided and each subdomain assigned to a specific GPU. Each subdomain can operate with its own spatial and temporal resolution.

To enable acoustic wave propagation between subdomains, the halo regions are periodically resampled and exchanged.

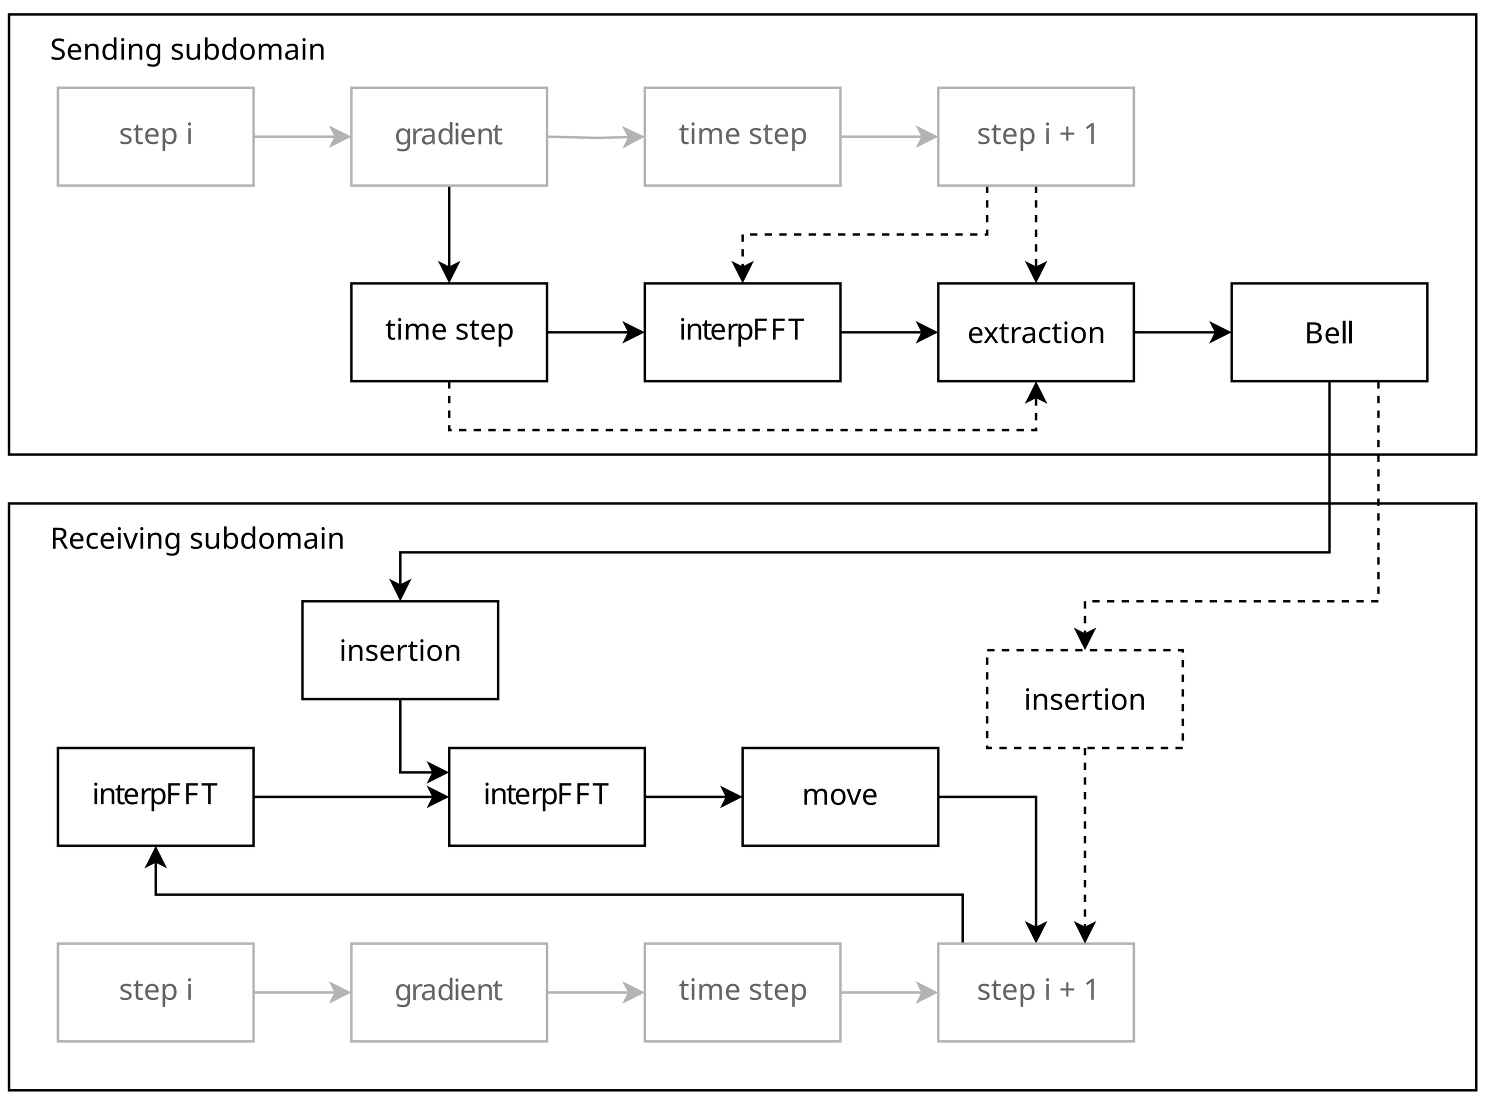

The simulation is implemented in CUDA, utilizing the cuFFT library to perform Fourier transformations. The exchange of halo regions is accelerated with cuFFT callback functions and transmitted via NVLinks.

Results

A maximum error of 0.8% was achieved when decomposing the domain into 8 subdomains without any change in resolution. Higher error levels occur with multiple resolutions due to sound reflections from subdomain interfaces. The error depends on the halo region width and the type of Bell function used, both of which are user-defined.

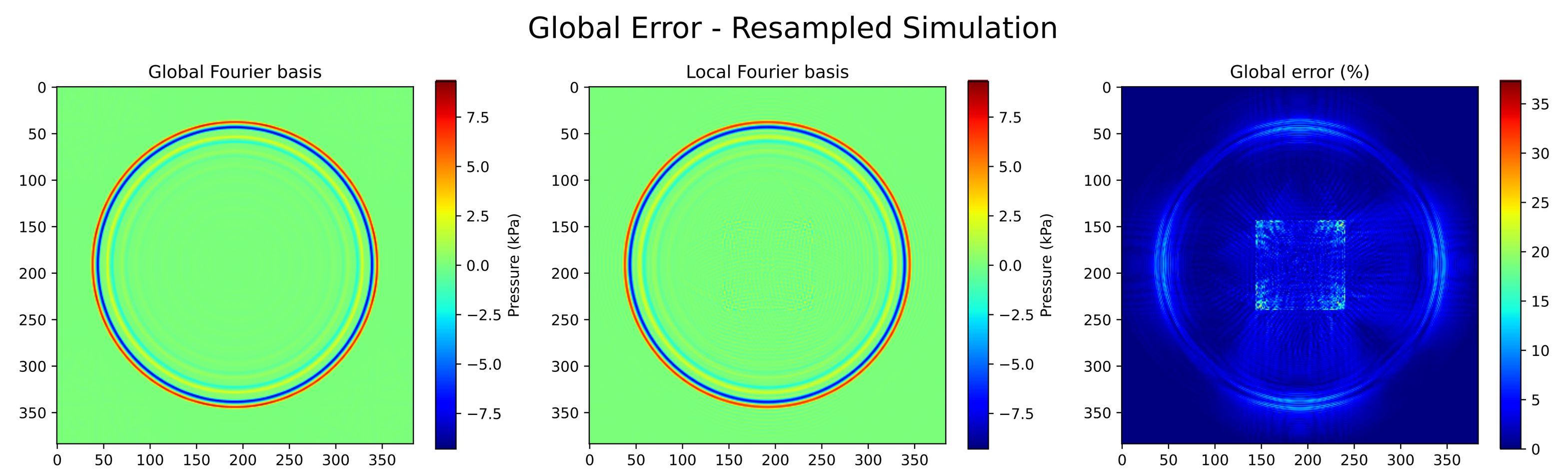

The simulation speed was tested on domains ranging in size from 128^3 to 4096^3 grid points, using a uniform 3D decomposition across up to 256 GPUs on the Karolina supercomputer, with all subdomains maintaining the same resolution. The strong scaling graph indicates substantial speedup for sufficiently large domains.

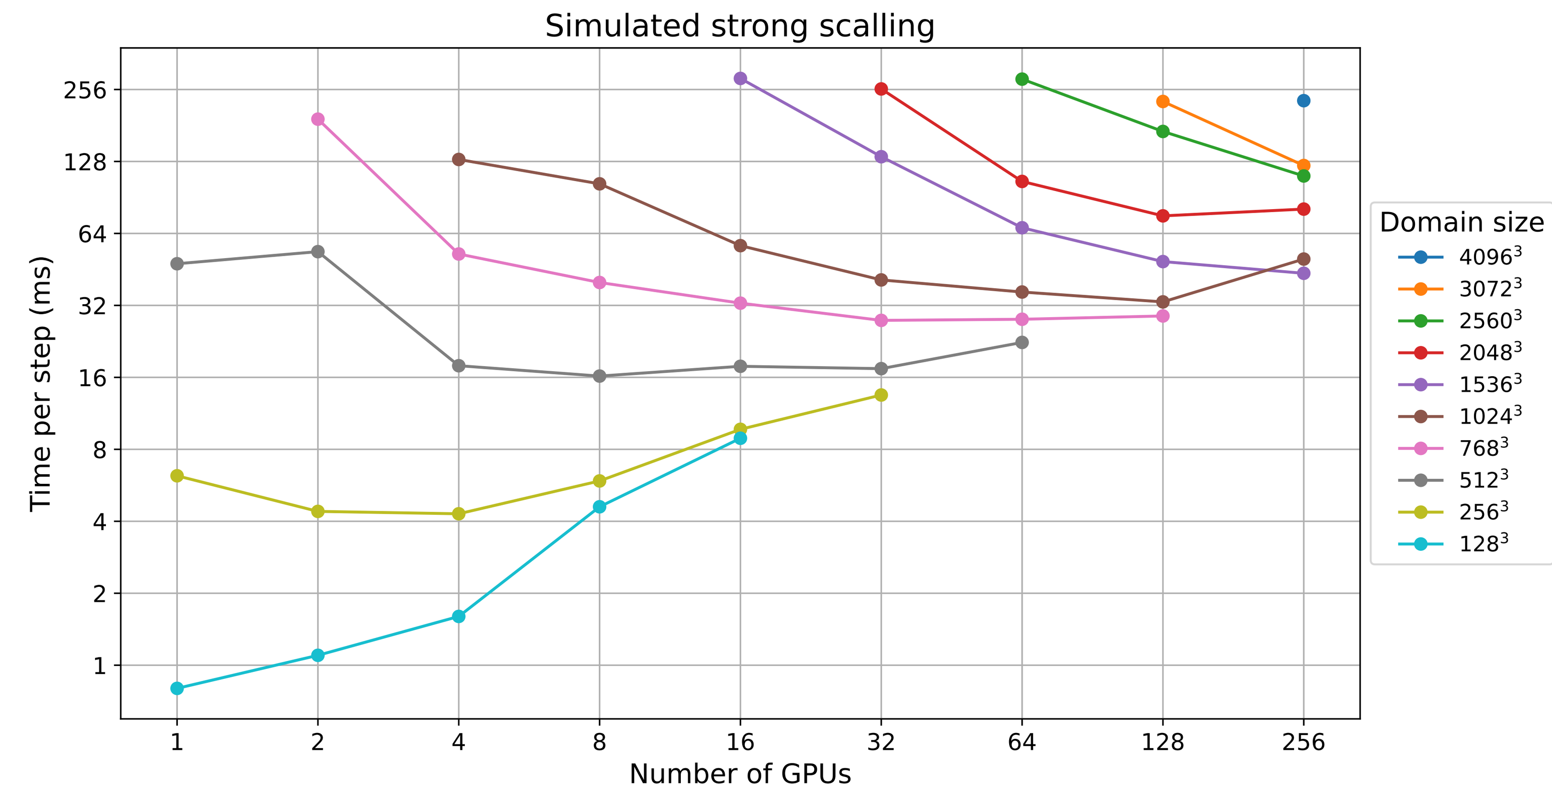

The weak scaling graph shows that the time per iteration slowly rises due to growing rank of the decompostion but stabilizes at 64 GPUs for most domain sizes.

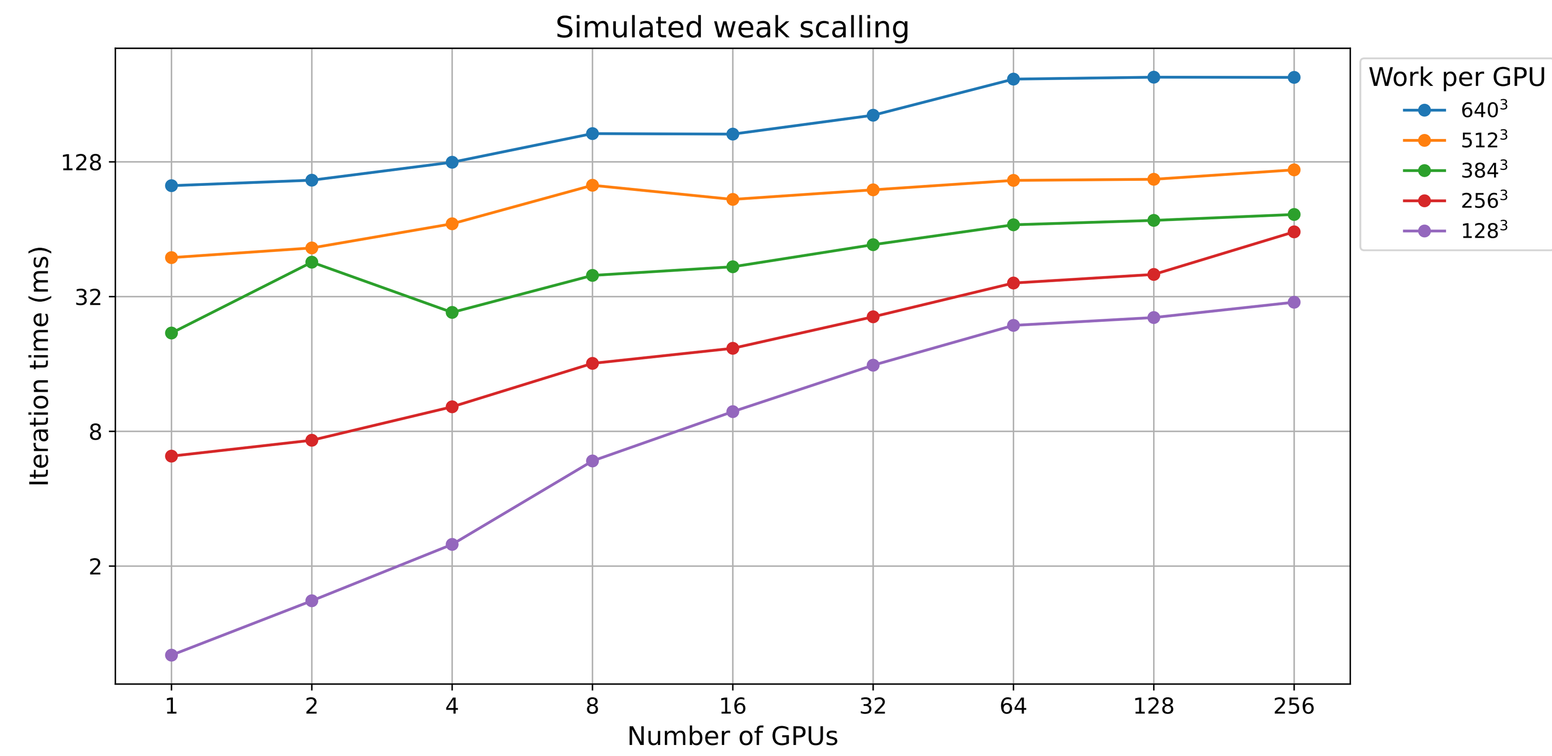

For speed tests with multiple resolutions, a 2D decomposition was used, with sections of the domain downscaled by a factor of four. The domain was divided into eight subdomains and compared against a uniform 3D decomposition without resolution changes, as well as the default single-GPU k-Wave implementation.

The multi-resolution simulation outperforms the single-GPU implementation but is considerably slower than the uniform Multi-GPU resolution simulation. While simulation time is significantly reduced, the interpolation and halo exchanges take substantially longer.

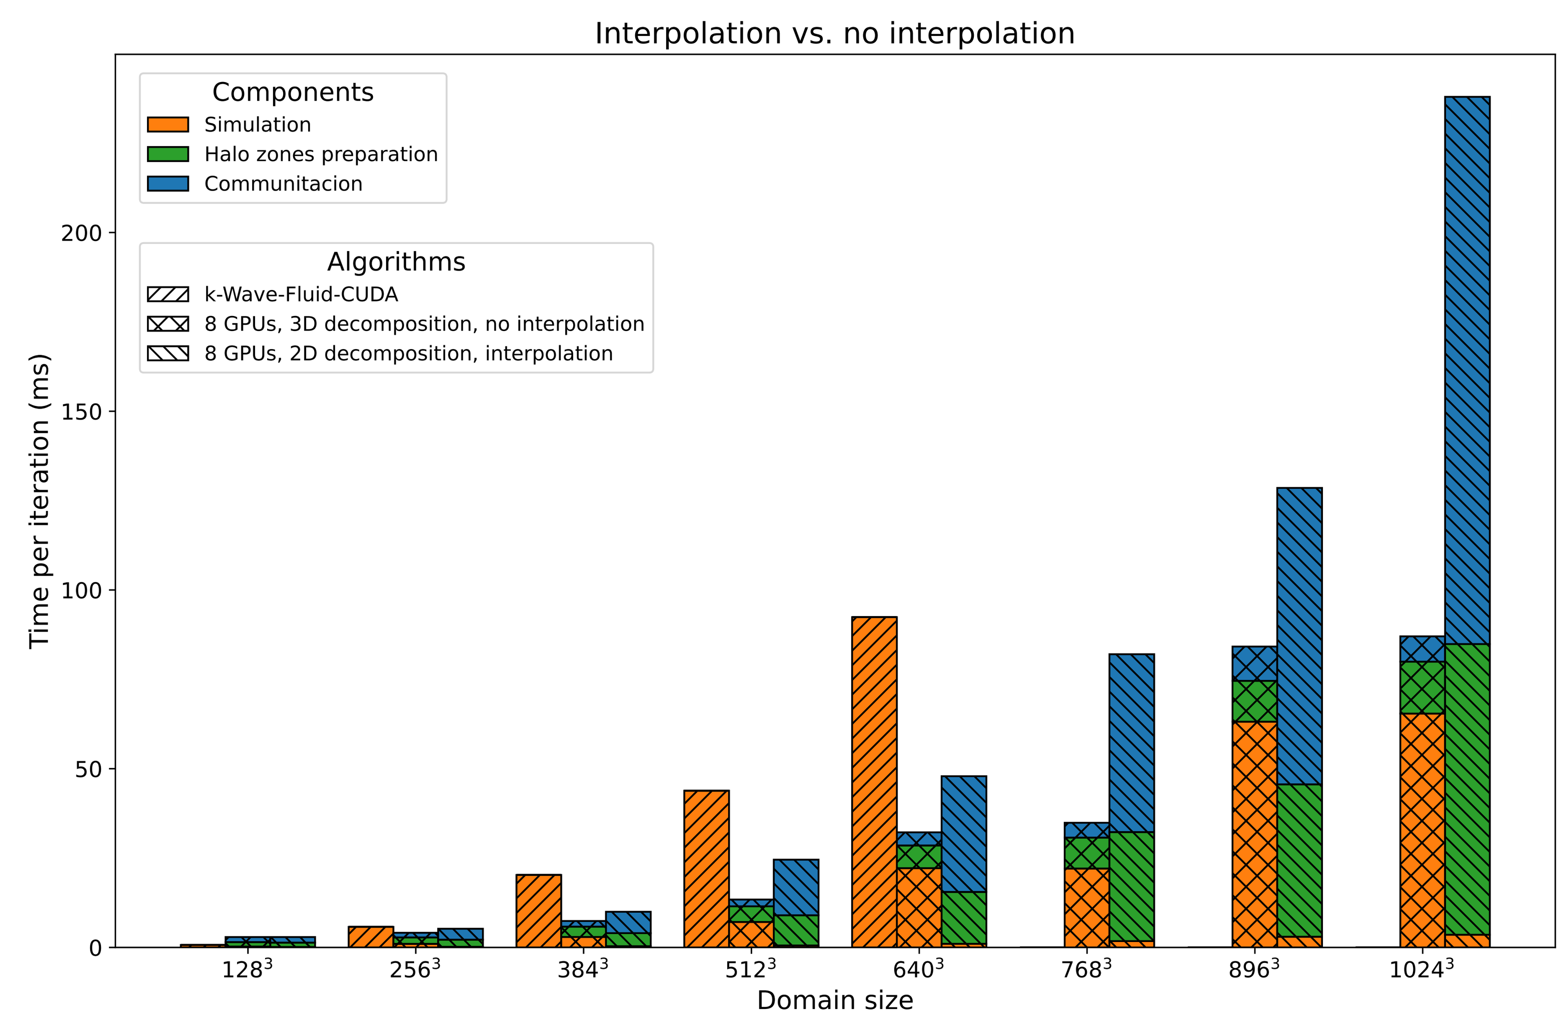

Conclusion

Our approach demonstrates that local decomposition can effectively use the Fourier collocation method for sound propagation simulations. With uniform-resolution subdomains, maximum error remains low, achieving high efficiency on 32-256 GPUs with large domains. This approach also enables simulations on larger domains that exceed single-GPU memory limits.

With multiple resolutions, sound propagates accurately, and with optimized halo regions resampling and exchange parameters, this method shows potential for broader usability.

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

This work was supported by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90254). This project has received funding from the European Unions Horizon Europe research and innovation programme under grant agreement No 101071008. This work was supported by Brno University of Technology under project number FIT-S-23-8141.