Performance and accuracy analysis of nonlinear k-Wave simulations using local domain decomposition with an 8-GPU server

Bradley Treeby, Filip Vaverka, and Jiri Jaros

Citation: Proc. Mtgs. Acoust. **34**, 022002 (2018); doi: 10.1121/2.0000883
View online: https://doi.org/10.1121/2.0000883
View Table of Contents: http://asa.scitation.org/toc/pma/34/1
Published by the Acoustical Society of America

Proceedings of Meetings on Acoustics

Volume 34 http://acousticalsociety.org/

21st International Symposium on Nonlinear Acoustics

21st ISNA

Santa Fe, New Mexico, USA July 9-13, 2018, 50th Anniversary

Computational Acoustics: S3-5

Performance and accuracy analysis of nonlinear k-Wave simulations using local domain decomposition with an 8-GPU server

Bradley Treeby
Department of Medical Physics and Biomedical Engineering, University College London, London, WC1E6BT,UNITED KINGDOM; b.treeby@ucl.ac.uk
Filip Vaverka and Jiri Jaros
Centre of Excellence IT4Innovation, Faculty of Information Technology, Brno University of Technology, Brno, South Moravian Region, CZECH REPUBLIC; ivaverka@fit.vutbr.cz, jarosjir@fit.vutbr.cz

Large-scale nonlinear ultrasound simulations using the open-source k-Wave toolbox are now routinely performed using the MPI version of k-Wave running on traditional CPU-based clusters. However, the all-to-all communications required by the 3D fast Fourier transform (FFT) severely impact performance when scaling to large numbers of compute cores. This can be overcome by using a domain decomposition strategy based on a local Fourier basis. In this work, we analyze the performance and accuracy of using local domain decomposition for running a high-intensity focused ultrasound (HIFU) simulation in the kidney on a single server containing eight NVIDIA P40 graphical processing units (GPUs). Different decompositions and overlap sizes are investigated and compared to a global MPI simulation running on a CPU-based supercomputer using 1280 cores. For a grid size of 960 by 960 by1280 grid points and an overlap size of 4 grid points, the error in the simulation using local domain decomposition is on the order of 0.1% compared to the global simulation, which is sufficient for most applications. The financial cost for running the simulation is also reduced by more than an order of magnitude.

Published by the Acoustical Society of America

 https://doi.org/10.1121/2.0000883

1. INTRODUCTION

Large-scale nonlinear ultrasound simulations in heterogeneous media are particularly important for high-intensity focused ultrasound (HIFU), for example, to aid with hardware design, patient selection, and treatment planning.[1] However, these are computationally demanding simulations due to the large size of the area of interest compared to the acoustic wavelength. Such simulations can be performed using the MPI version of k-Wave running on traditional CPU-based clusters.[2,3] However, the all-to-all communications required by the 3D fast Fourier transform (FFT) severely impact performance when scaling to large numbers of compute cores. This can be overcome by using a domain decomposition strategy based on a local Fourier basis as described in Ref. 4. In the current work, the performance and accuracy of this local domain decomposition approach are investigated for a real-world problem.

2. METHODS

Simulations were performed for a HIFU sonication of the kidney using a single element bowl transducer. The transducer had a radius of curvature of 140 mm, circular aperture diameter of 120 mm, and frequency of 1 MHz. These values are within the range of parameters for commercially available focused ultrasound systems used for abdominal treatments. The transducer geometry was modelled using a simply-connected discrete bowl,[5] with a surface intensity of 1 W/cm^2 (175 kPa) and an acoustic power of 119 W.

The material properties were derived from the open-source AustinWoman voxel model,[6] which provides a label-based segmentation of the digital image datasets from the Visible Human Project run by the U.S. National Library of Medicine. Book values for the material properties were assigned based on the IT'IS database.[7] The geometric focus of the transducer was positioned slightly above the inferior pole of the left kidney, with the focal position approximately 90 mm from the skin surface. The domain was discretised using a grid point spacing of 172 μm (giving a maximum supported frequency of 4.4 MHz), and a grid size of 960 $\times$ 960 $\times$ 1280 grid points (corresponding to a domain size of 16.5 $\times$ 16.5 $\times$ 22 cm). The simulation was run until steady state using a time step of 11.4 ns (giving a CFL number of 0.1 in the background medium) and 25407 time steps.

k-Wave simulations using local-domain decomposition were run as described in Ref. 4. The bell-shape was chosen using numerical optimisation by minimising the L_∞ error for a broadband plane-wave pulse in a homogeneous medium propagating between two neighbouring sub-domains. The simulations were run using a single compute server (PNY 2U 8 GPU Barebone Server) with 8 $\times$ NVIDIA Tesla P40 Pascal GPUs each with 3840 CUDA cores and 24 GB of memory. Combined, the server has a total GPU memory of 192 GB, and a theoretical single-precision performance of 96 Tflops.

The global domain (960 $\times$ 960 $\times$ 1280 grid points) was divided into 8 sub-domains, with each sub-domain assigned to a single GPU. Three different decompositions were investigated, using either (1, 1, 8), (1, 2, 4), or (2, 2, 2) sub-domains in the (x, y, z) directions. For each decomposition, four overlap sizes were investigated ranging from 4 to 16 grid points. As the computational performance depends heavily on the FFT, the local domains were spatially zero-padded to give local domain sizes with small prime factors (within the constraints of the 24 GB of memory available on each GPU). The local domain sizes, largest prime factor, and GPU memory usage for each case are given in Table 1.

For reference, simulations were also performed using a global domain with the MPI version of k-Wave.[3] The simulations were run on the Anselm supercomputer operated by the IT4Innovations National Supercomputing Center in Ostrava, Czech Republic. Anselm is an Intel-infiniband cluster based on Sandy Bridge processors (2$\times$8 core Intel E5-2665 at 2.4 GHz and 64 GB RAM per node) interconnected by a 40 Gb Fat-tree infiniband interconnection. To give approximately similar floating point performance to the GPU server, the reference simulation was run using 80 nodes (1280 cores).

Table 1: Compute times and errors for running a 960 × 960 × 1280 simulation using local domain decomposition with different arrangements of sub-domains and overlap sizes. The local domains are padded to give small prime factors, and the largest prime factor for each dimension is reported. The memory usage corresponds to the memory usage per GPU.

Decomp (x, y, z)	**Overlap (grid points)**	**Local Domain (grid points)**	**Factors (x, y, z)**	**Memory (GB)**	**Compute Time (hr:min)**	L_∞ **error (%)**
(1, 1, 8)	4	972×972×192	(3, 3, 3)	21.8	4:55	0.20
	8	972×972×192	(3, 3, 3)	22.2	6:45	0.039
	12	972×972×192	(3, 3, 3)	22.6	8:21	0.020
	16	972×972×192	(3, 3, 3)	23.0	9:50	0.014
(1, 2, 4)	4	972×512×384	(3, 2, 3)	23.1	5:17	0.091
	8	972×512×384	(3, 2, 3)	23.4	7:23	0.020
	12	972×512×384	(3, 2, 3)	23.8	9:28	0.0072
	16	960×512×360	(5, 2, 5)	22.8	11:20	0.0043
(2, 2, 2)	4	512×512×648	(2, 2, 3)	21.0	4:52	0.12
	8	512×512×672	(2, 2, 7)	22.1	6:48	0.030
	12	512×512×672	(2, 2, 7)	22.4	9:13	0.012
	16	512×512×672	(2, 2, 7)	22.8	11:10	0.0064

3. RESULTS

The compute times and the L_∞ error in the peak positive pressure in steady state compared to the global simulation are shown in Table 1 and Fig. 1. The error percentages are calculated relative to the spatial peak positive pressure in the global simulation. For an overlap of four grid points, the maximum error is on the order of 0.2%, which is sufficient for many applications. This decreases with an increasing overlap size, at the expense of an increase in compute time. There is some variation in error between the three decompositions. Figure 2 shows error plots for the three decompositions for an overlap of four grid points. Small reflections from the sub-domain interfaces can be seen in the error plots. The reduced error for the (1, 2, 4) decomposition is most likely due to the position of the domain cuts relative to the spatial distribution of the acoustic pressure (there are no x or z cuts directly through the focal region).

For an overlap of 4 or 8 grid points, the local domain simulations on a single GPU node are faster than the global domain code running on 80 CPU nodes, where both systems have approximately the same floating point performance. Perhaps more significantly, for all local domain simulations, the financial cost of running the simulation is significantly reduced. For example, using the West Europe prices for the Microsoft Azure cluster, a GPU node with 4 × P40 GPUs is $12.76 per hour, making the simulation cost for 2 nodes for 5 hours $127. In comparison, a high-performance compute node with 2 × 8 cores (similar to Anselm) is $3.41 per hour, making the simulation cost for 80 nodes for 8 hours $2182, a factor of 17 higher.

For comparison, the compute times and errors for the (2, 2, 2) decomposition *without* local subdomain padding are shown in Fig. 1. In this case, the maximum prime factors in (x, y, z) are (61, 61, 3), (31, 31, 41), (7, 7, 93), and (2, 2, 7) for overlap sizes of 4, 8, 12, and 16, respectively. The large prime factors have a significant effect on the compute time, in some cases causing the simulation time to more than double. In contrast, there is almost no change in the accuracy of the simulation. Consequently, it is concluded that local sub-domain padding should always be used to improve the computational performance of the FFT.

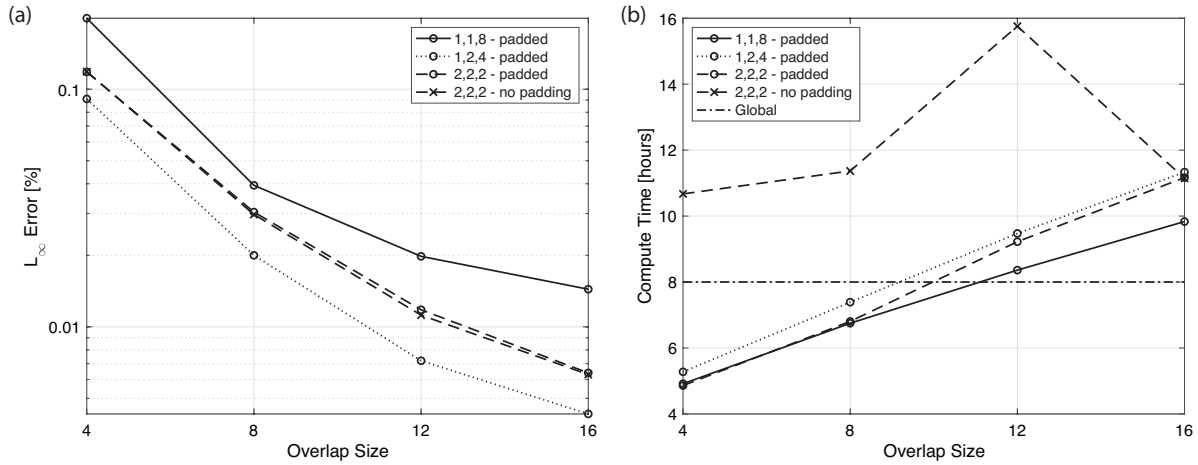

Figure 1: Change in the (a) L_∞ error and (b) compute time with the size of the overlap between subdomains for different domain decompositions.

4. CONCLUSION

k-Wave simulations using local domain decomposition with a local Fourier basis have been performed using an 8-GPU server. For the investigated HIFU scenario, errors on the order of 0.1% can be achieved using an overlap of 4 grid points. Compared to simulations running on a traditional CPU cluster, the financial cost is decreased by more than an order of magnitude.

ACKNOWLEDGMENTS

This work was supported by the Engineering and Physical Sciences Research Council, UK, and the European Union's Horizon 2020 research and innovation programme H2020 ICT 2016-2017 under grant agreement No 732411 (as an initiative of the Photonics Public Private Partnership). This work was also 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.

REFERENCES

[1] K. Okita, R. Narumi, T. Azuma, H. Furusawa, J. Shidooka, and S. Takagi. Effects of breast structure on high-intensity focused ultrasound focal error. *J. Therapeutic Ultrasound*, 6(4):1–12, 2018.

[2] J. Jaros, V. Nikl, and B. E. Treeby. Large-scale Ultrasound Simulations Using the Hybrid OpenMP / MPI Decomposition. In *Proceedings of the 3rd International Conference on Exascale Applications and Software*, pages 115–119, 2015.

[3] J. Jaros, A. P. Rendell, and B. E. Treeby. Full-wave nonlinear ultrasound simulation on distributed clusters with applications in high-intensity focused ultrasound. *Int. J. High Perform. C.*, 30(2):137–155, 2016.

[4] J. Jaros, F. Vaverka, and B. E. Treeby. Spectral Domain Decomposition Using Local Fourier Basis: Application to Ultrasound Simulation on a Cluster of GPUs. *Supercomputing Frontiers and Innovations*, 3(3):39–54, 2016.

[5] E. Martin, Y. T. Ling, and B. E. Treeby. Simulating focused ultrasound transducers using discrete sources on regular cartesian grids. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 63(10):1535–1542, 2016.

[6] J. Massey, C. Geyik, N. Techachainiran, C. Hsu, R. Nguyen, T. Latson, M. Ball, and A. Yilmaz. AustinMan and AustinWoman: High fidelity, reproducible, and open-source electromagnetic voxel models. In *Bioelectromagnetics Society 34th Annual Meeting*, 2012.

[7] P. Hasgall, F. Di Gennaro, C. Baumgartner, E. Neufeld, M. Gosselin, D. Payne, A. Klingenböck, and N. Kuster. IT'IS Database for thermal and electromagnetic parameters of biological tissues (Version 3.0), 2015.

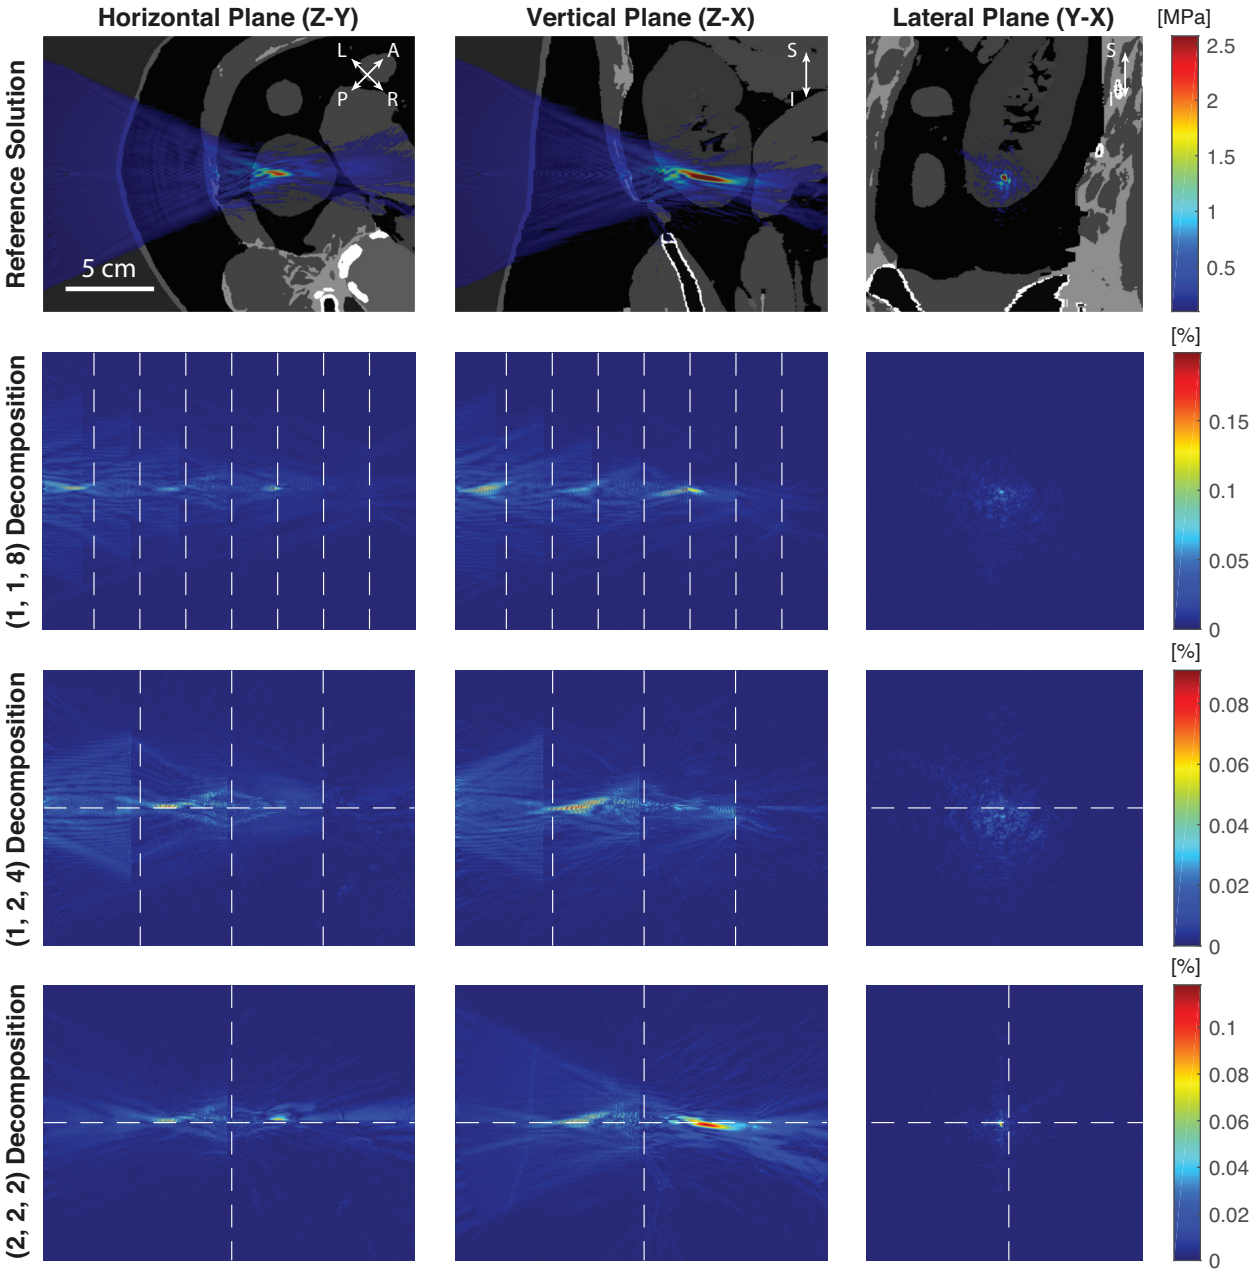

Figure 2: (Top Row) Output from the global domain simulation showing the peak positive pressure in steady state overlaid on the map of the tissue sound speed derived from the AustinWoman model. Three slices through the geometrical focus of the transducer are shown. (Lower Rows) Error plots showing the difference between the local domain and global domain simulations for three different decompositions using an overlap of four grid points. The domain boundaries are shown with dashed lines.