Review of a Doctoral Thesis at FIT BUT

Doctoral thesis

LARGE-SCALE ULTRASOUND SIMULATIONS USING ACCELERATED CLUSTERS

Name of the doctoral student

Ing. FILIP VAVERKA

Name and institution of the reviewer

doc. Ing. Lubomír Říha, Ph.D.,

VSB - Technical University of Ostrava, IT4Innovations National Supercomputing Center, Czech Republic

---

I. Thesis

<u>Appropriateness and relevance</u>

The work focuses on effective numerical modeling and computer simulation of wave physical phenomena, specifically in the context of medical ultrasound. This is a relevant and important study area, as advances in medical ultrasound technologies have led to a growing need for more efficient and accurate simulations.

The thesis explains why the spectral methods are the optimal ones due to the significant reduction in spatial resolution when compared to finite-difference time-domain methods. These methods still need distributed memory processing due to the large grid size for practical applications, and therefore high-performance computing is essential for final implementation.

<u>A summary of the contributions of the thesis</u>

The thesis addresses several critical aspects in the chosen field, including:

- the influence of modern accelerated and non-accelerated supercomputer architectures on the design of efficient numerical methods for the simulation of wave phenomena,
- overview of numerical methods typically used for the discretization of PDE-based models describing wave phenomena with a focus on systems of conservation law,
- adaptation of existing Fourier numerical schemes for the environment of modern accelerated clusters, illustrated through adapting the k-Wave scheme for modeling the nonlinear propagation of ultrasonic waves in heterogeneous and absorbing environments,
- parallel implementation for heterogeneous (CPUs + accelerators) HPC architectures that is efficiently applicable to modern HPC clusters currently deployed.

Review of a Doctoral Thesis at FIT BUT

Novelty and significance:

The thesis presents a novel domain decomposition method that, based on the presented performance evaluation:

- extends the scalability of the original global decomposition method due to different decomposition method,
- exhibits significantly better both strong and weak scalability of the CPU version when compared to the original GDD implementation in the k-Wave,
- enables the GPU acceleration in the k-Wave.

The strong part of the thesis is also a very thorough evaluation of the state-of-the-art of related numerical methods and a detailed description of how the new method was developed and validated in terms of numerical properties and stability. The novelty and significance of the thesis are sufficient for a Ph.D. degree.

Question: Can you explain the reasons why, in Figure 6.5 left, the scalability of the GDD method is not measured up to 512 sockets?

Evaluation of the formal aspects of the thesis:

From the formal aspect, the thesis is well-written with an adequate language level. I have no comments in this regard.

Quality of publications

The work from this thesis is published in:

- two journal articles
    - one in 33 average centile on Scopus (Supercomputing Frontiers and Innovations in 2016),
    - one is an extended (5-page) abstract indexed on Scopus as Jsc,
- 6 international conferences proceedings (all as the first author)
    - out of which one is CORE A and one CORE B,
- 4 posters
    - out of which 3 at the prestigious Supercomputing (SC) conference.

This is quite a high publication activity for a Ph.D. study. However, none of the journal papers are indexed on WoS and are without IF.

II. Student's overall achievements

Overall R&D activities evaluation:

From the HPC point of view, the main focus of the thesis is developing a method that also enables execution on large-scale GPU-accelerated clusters.

Review of a Doctoral Thesis at FIT BUT

Can you please address the following comments:

- There is an inconsistency in the results for the Salomon cluster presented in Figures 6.18 and 6.19. For instance, for domain size 256x512x512 and one socket, Figure 6.18 shows a runtime of approx. 1024 seconds, while Figure 6.19 shows a runtime of approx. 512 seconds. The results for Piz Daint are consistent. Can you clarify this?

- The thesis presents both MPI + OpenMP and MPI+CUDA parallelization. However, there is no mention of the MPI + OpenMP + CUDA combination in the thesis. Can you explain what are the challenges?

- The Piz Daint is a Cray XC50 machine with 1 CPU and 1 GPU compute node. On the same compute node footprint, one can have the dual-socket CPU node. Based on Fig. 6.18:
    - for problem size 2048^3 reported speedup is approx. 5x for chip-to-chip comparison,
    - which means that for node-to-node, speedup would be approx. 2.5x (estimated from values for 512 CPU sockets and 256 GPUs),

  Considering my first comment (inconsistency between Figures 6.18 and 6.19), the real speedup could be even lower. Can you check the result and clarify?

  My previous comment is for this rather old Cray CX platform. When one evaluates the architecture of the current HPC clusters, where CPU to GPU ratio is 1:4 (for instance, LUMI, ORNL Frontier) and the CPU to GPU performance ratio is approximately 1:100 (CPU has 2TF and 4 GPUs have 200TF of peak performance in DP64), the method will utilize a significant part of the peak performance of the node. As such, I consider the method more even useful for new machines.

- Comparing the communication overhead presented in Figures 6.20 and 6.21. The Salomon cluster spends a significantly higher part in communication when compared to Piz Daint. There are two points that can potentially cause this overhead:
    - on Salomon, two CPUs share a single NIC, which reduces the bandwidth per CPU, and the second CPU needs to go over the QPI bus to reach the NIC connected to the first CPUs
    - on Piz Daint there is a single GPU per node (based on the text, I expect the CPU is not used for computation but only for communication)

  Have you investigated this, and can you discuss this in more detail?

  I would suggest evaluating the scalability of the CPU and GPU code on the same machine with identical hardware and software stack. In this particular case, to also run the MPI + OpenMP version on Piz Daint.

<u>Assessment of other characteristics</u> (optional):

The student worked for most of his Ph.D. studies in an international environment. Since 2015 he has collaborated with a research group at UCL, which also resulted in several joint publications.

He has also worked on the EU project PAMMOTH and received Excellent Paper Award at PAD 2019 conference.

III. Conclusion

The thesis presents a significant and novel contribution by the development and efficient implementation of scalable domain decomposition method in the ultrasound simulation package k-Wave. In my opinion, it contains all the key features required to fulfill a Ph.D. degree, which are:

- development of a new method based on a thorough investigation of state of the art, which enables to solve of large problems faster due to improved scalability,
- efficient parallel implementation of the method into the existing framework k-Wave,
- development of support for new hardware architectures (Xeon Phi and GPUs).

In my opinion, the thesis and the student´s achievements meet the requirements for the award of an academic Ph.D. degree.

Ostrava 10.5.2023

doc. Ing. Lubomír Říha, Ph.D.,