11th World Congress on Computational Mechanics (WCCM XI)
5th European Conference on Computational Mechanics (ECCM V)
6th European Conference on Computational Fluid Dynamics (ECFD VI)
July 20–25, 2014, Barcelona, Spain

PARALLELIZING THE FAST MULTIPOLE METHOD USING A TASK-BASED RUNTIME FOR HETEROGENEOUS ARCHITECTURES

Emmanuel Agullo[1], Bérenger Bramas[1], Olivier Coulaud[1], Eric Darve[2⋆], Matthias Messner[2], Toru Takahashi[3]

[1] Inria, Hiepacs Project, Talence, France, Olivier.Coulaud@inria.fr
[2] Mechanical Engineering Department and Institute for Computational and Mathematical Engineering, Stanford University, USA, darve@stanford.edu
[3] Department of Mechanical Science and Engineering, Nagoya University, Japan, ttaka@nuem.nagoya-u.ac.jp

Key words: *Parallel computing, heterogeneous platforms, GPU, fast multipole method*

A high performance implementation of the fast multipole method (FMM) is crucial of the numerical simulation of many physical problems [4]. In a previous study [3, 2, 1], we have shown that a task-based FMM formulation (i.e., using a direct acyclic graph representation of the computation) provides the flexibility required to process a wide spectrum of particle distributions efficiently on multicore architectures. In this talk, we will show how such an approach can be extended to fully exploit heterogeneous platforms. For that, we designed highly tuned GPU versions of the two dominant operators (P2P and M2L) as well as a scheduling strategy that *dynamically* (i.e., at runtime) decides which tasks get processed on regular CPU cores and on GPU accelerators. We assessed our method with the StarPU runtime system for executing the resulting task flow. The platform is an Intel X5650 Nehalem multicore processor possibly enhanced with one, two or three Nvidia Fermi M2070 or M2090 GPUs. A detailed experimental study on two 30 million particle distributions (a cube and an ellipsoid) show that the resulting software consistently achieves high performance across various parallel processor architectures.

REFERENCES

[1] Emmanuel Agullo, George Bosilca, Bérenger Bramas, Cedric Castagnede, Olivier Coulaud, Eric Darve, Jack Dongarra, Mathieu Faverge, Nathalie Furmento, Luc Giraud, et al. Poster presentation: Matrices over runtime systems at exascale. In *High Performance Computing, Networking, Storage and Analysis, SC'12*, pages 1332–1332. IEEE, 2012.

[2] Emmanuel Agullo, Bérenger Bramas, Olivier Coulaud, Eric Darve, Matthias Messner, and Takahashi Toru. Pipelining the fast multipole method over a runtime system. *arXiv preprint arXiv:1206.0115*, 2012.

[3] Emmanuel Agullo, Bérenger Bramas, Olivier Coulaud, Eric Darve, Matthias Messner, Toru Takahashi, et al. Task-based parallelization of the fast multipole method on NVIDIA GPUs and multicore processors. In *GPU Technology Conference*, 2013.

[4] Matthias Messner, Bérenger Bramas, Olivier Coulaud, and Eric Darve. Optimized M2L kernels for the Chebyshev interpolation based fast multipole method. *arXiv preprint arXiv:1210.7292*, 2012.