EXPLORING PROGRAMMING MODELS FOR ACCELERATING SCIENTIFIC APPLICATIONS ON HYBRID CPU-MIC PLATFORMS

Roman Wyrzykowski, Lukasz Szustak, Kamil Halbiniak
Faculty of Mechanical Engineering and Computer Science,
Czestochowa University of Technology, Poland
roman@icis.pcz.pl

Modern heterogeneous computing platforms have become powerful HPC solutions, which could be applied to a wide range of real-life applications. In particular, the hybrid platforms equipped with Intel Xeon Phi coprocessors offer the advantages of massively parallel computing, while supporting practically the same (or similar) parallel programming model as conventional homogeneous solutions. However, there is still an open issue as to how scientific applications can efficiently utilize hybrid platforms with Intel MIC coprocessors.

In paper [1], we proposed a method for porting a real-life scientific application to computing platforms with Intel MICs. We focused on the parallel implementation of a numerical model of alloy solidification, which is based on the generalized finite difference method. We developed a sequence of steps that are necessary for porting this application to platforms with accelerators, assuming no significant modifications of the code. The proposed method considers not only overlapping computations with data movements, but also takes into account an adequate utilization of cores/threads and vector units. Using parallel resources of one Intel Xeon Phi coprocessor (KNC architecture), the developed approach allowed us to execute the whole application 3.45 times faster than the original parallel version running on two CPUs.

In this work, we focus on studying various heterogeneous programming models for accelerating the solidification application on hybrid CPU-MIC platforms. We focus on two models: OpenMP 4.0 Accelerator Model and Hetero Streams Library (hStreams in short) [2]. Now the main challenge for achieving a desired high performance of computations is to take advantage of CPUs and coprocessors to work together, when all the available threads of CPUs and Intel MICs are utilized coherently to solve the modelling problem.

In the paper, we present the performance comparison of the above-mentioned models for various configurations of computing resources. In particular, using the hStreams library, our approach allows us parallelize efficiently the solidification application on hybrid platforms with two CPUs and two MICs, and accelerate computations about 10.5 times in comparison with the basic version for two CPUs. We also conclude that while OpenMP provides an unified directive-based programming model, the current stable version of this standard is not efficient in multi-device heterogeneous platforms. That is why, we plan to investigate new features available in version 4.5 of OpenMP, such as asynchronous offload mechanism.

References

[1] L. Szustak et al. Toward parallel modeling of solidification based on the generalized finite difference method using Intel Xeon Phi. *PPAM 2015, Part I. LNCS*, 9573:411–412, 2016.

[2] Ch. J. Newburn et al. Heterogeneous streaming. *IPDPSW, AsHES*, 2016.