Contribution ID: **183** Type: **Sectional reports**

Population annealing method and hybrid supercomputer architecture

Thursday, 8 July 2021 13:30 (15 minutes)

A population annealing method is a promising approach for large-scale simulations because it is potentially scalable on any parallel architecture. We report an implementation of the algorithm on a hybrid program architecture combining CUDA and MPI [1]. The problem is to keep all general-purpose graphics processing unit devices as busy as possible by efficiently redistributing replicas. We provide testing details on hardware-based Intel Skylake/Nvidia V100, running more than two million replicas of the Ising model samples in parallel. As the complexity of the simulated system increases, the acceleration grows toward perfect scalability.

This work was initiated under Grant No. 14-21-00158 and finished under Grant No. 19-11-00286 from the Russian Science Foundation. We also acknowledge the support within the scientific program of the Landau Institute for Theoretical Physics. We used the Manticore cluster of ANR laboratory at the Science Center in Chernogolovka for the small-scale testing and the supercomputing facility of the National Research University Higher School of Economics for the large-scale testing [2].

[1] A. Russkov, R. Chulkevich, L. Shchur, Computer Physics Communications, 261, 107786 (2021)
[2] P. S. Kostenetskiy, R. A. Chulkevich, and V. I. Kozyrev, J. Phys. Conf. Ser. 1740, 012050 (2021)

Summary

Primary authors: SHCHUR, Lev (leading reseacher, Landau Institute for Theoretical Physics); RUSSKOV, Alexander (Scientific Center in Chernogolovka); CHULKEVICH, R. (National Research University Higher School of Economics)

Presenter: SHCHUR, Lev (leading reseacher, Landau Institute for Theoretical Physics)

Session Classification: HPC

Track Classification: 5. High Performance Computing