Efficient low cost parallel computing using GPU and computer clusters

C.M.N.A. Pereira[1], A.C.A. Mól[1], M. Waintraub[1], A. Heimlich[1]
e-mail: cmnap@ien.gov.br

[1] *Division of Nuclear Engineering - IEN*

Keywords: parallel programming, GPU, MPI, evolutionary computation, Monte Carlo

Introduction

Due to the high computational overhead imposed a great variety of simulations and optimization problems found in the nuclear engineering field, research and applications of parallel computing have become very attractive to the Artificial Intelligence group of the Applied Artificial Intelligence Laboratory (LIAA/IEN).

In the very beginning, looking for low cost approaches in order to accelerate evolutionary computation applications, the firsts parallel programs have been developed using TCP/IP sockets and C++ programming language [1-3].

In 2009, parallel models of Particle Swarm Optimization (PSO) were investigated by Waintraub [4]. In his research, Message-passing Interface (MPI) and C language have been used for programming a computer cluster.

Following a world tendency for low cost high performance computing, Heimlich [5] proposed the use of graphics processing units (GPU) and CUDA technology in the parallel processing of some nuclear engineering problems.

The last investigations are intended to explore a hybrid approach, using both: GPU/CUDA and Cluster/MPI technologies [6]. This research is under development and is supported by FAPERJ and CNPq.

Parallel computing in a GPU-Cluster using MPI and CUDA

In this work a hybrid parallel Monte Carlo based neutron transport simulation program has been developed using Message-passing Interface (MPI) and Compute Unified Device Architecture (CUDA) technologies. Such program is aimed to run on a GPU-Cluster, that means, a computer cluster in which the nodes are provided with programmable Graphics Processing Units (GPU). A quite simple, but very time consuming Monte Carlo simulation has been considered in order to shown that making use of an uncomplicated and low cost computer architecture (see Figure 1), it is possible to achieve great gains in terms of computational performance.

Results and Discussion

Using the 8-GPU-cluster (4 multi-core PC, with 2 GPU each) several simulations have been done. Speedups over 900 times have been obtained for all cases. In the best case, the parallel simulation was more than 2,000 times faster than the sequential program running on a single processor.

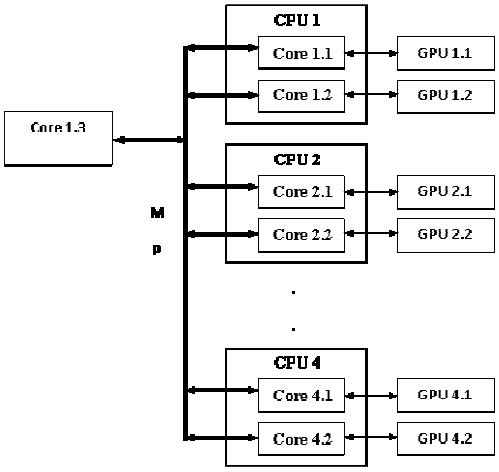

Figure 1. The parallel hardware architecture.

References

[1] Pereira, C.M.N.A. e Lapa, C.M.F. "Coarse-grained parallel genetic algorithm applied to a nuclear reactor core design optimization problem". Annals of Nuclear Energy, v. 30, n.5, p. 555-565, 2003.

[2] Pereira, C.M.N.A. e Lapa, C.M.F. "Parallel island genetic algorithm applied to a nuclear power plant auxiliary feedwater system surveillance tests policy optimization". Annals of Nuclear Energy, v. 30, n.16, p. 1665-1675, 2003.

[3] Pereira, C.M.N.A. e Sacco, W.F. "A parallel genetic algorithm with niching technique applied to a nuclear reactor core design optimization problem". Progress in Nuclear Energy, v. 50, p. 740-746, 2008.

[4] Waintraub, M., Schirru, R. e Pereira, C.M.N.A. "Multiprocessor modeling of parallel Particle Swarm Optimization applied to nuclear engineering problems". Progress in Nuclear Energy, v. 51, p. 680-688, 2009.

[5] Heimlich, A., Mol, A.C.A. e Pereira, C.M.N.A. "GPU-based Monte Carlo simulation in neutron transport and finite differences heat equation evaluation". Progress in Nuclear Energy, 2011.

[6] Pereira, C.M.N.A., Mol, A.C.A., Heimlich, A., Moraes, S.R.S. e Resende, P. "Development and performance analysis of a parallel Monte Carlo neutron transport simulation program for GPU-Cluster using MPI and CUDA technologies". Progress in Nuclear Energy, 2013 (aceito para publicação).