SCOUT: Scheduling Core Utilization to Optimize the Performance of Scientific Computing Applications on CPU/Coprocessor-based Cluster

M. T. Chung[1], **K. T. Pham**[1], **M. T. Nguyen**[1], **and N. Thoai**[1]

Abstract: Today's scientific computing applications require many different kinds of task and computational resource. The success of scientific computing hinges on the development of High Performance Computing (HPC) system in the role of decreasing execution time. This is easier to be found when we have the support of accelerators like GPU or Intel Xeon Phi coprocessor. At the side of our system - CPU/Coprocessor-based cluster, the speed up increases from 1.4 to over 2 times with scientific applications in the scope of different research groups, for example, environmental issues simulation, material modeling.

One of the benefits of Intel Xeon Phi allows jobs easily to be performed by native mode (running directly on the coprocessor) or offload mode (running a main host program and offloading work to the coprocessor). However, problems related to coprocessor underutilization of offload jobs can cause oversubscription of memory and thread resource. Furthermore, through profiling the runtime behaviors of scientific applications, we recognize that a general offloading job does not occupy the coprocessor during an entire execution time, for typical jobs, it is only 32% the Xeon Phi cores being busy on a compute node. These problems can be solved to improve the performance significantly with a suitable method for scheduling jobs. In this paper, we improve the efficiency of coprocessor use through a scheduling method based on greedy choice. The proposed method enables coprocessor sharing without the overhead and runs as many jobs as possible on each coprocessor. Thus, our solution can improve the throughput and core utilization. Given a set of popular scientific applications, we show this method for increasing performance and throughput of workloads on our system. We implement the entire system as a seamless plugin to PBS Professional, a popular HPC job scheduler and show the efficiency with testbeds in practice.

[1] High Performance Computing Lab, Faculty of Computer Science and Engineering
Ho Chi Minh City University of Technology
268 Ly Thuong Kiet Street, District 10, Ho Chi Minh City, Vietnam
{*ctminh, 51301941, nmthin, namthoai*}*@hcmut.edu.vn*