平成25年度 情報工学コース卒業研究報告要旨

石井 研究室	氏　　名	高　雨 蘇
卒業研究題目	Research on Parallel Computing for Matrix Multiplication by Using Intel®Xeon Phi™	

In recent years, the demand for increasing compute power is growing rapidly in many fields of research, especially in the field of large scale numeric computation and simulation. Accelerators, like GPUs, are one way to fulfill these requirements, but they often require a laborious rewrite of the application using special programming models like CUDA or OpenCL. In contrast, the Intel®Xeon Phi™coprocessor is based on the Intel®Many Integrated Core Architecture and can be programmed with standard techniques like OpenMP, POSIX threads, or MPI. It is applied by many of the top 500 supercomputers in the world, including the present No.1, Tianhe II. It will provide high performance and low power consumption without the immediate need to rewrite an application.

Nagoya University Information Technology Center has introduced a new Fujitsu CX400 supercomputer system which includes an Intel®Xeon Phi™3100 family coporcessor on board. This accelerator system is expected to make a peak performance of 1.003 Tera FLOPS per node with total 184 nodes.

However, despite the reduction for rewriting and programming effort, supporting shared memory applications with only minimal changes does not necessarily mean that these applications are fully optimized to exert all the power of the accelerators. Based on the data from Fujitsu, the new introduced accelerator system gained a benchmark measurement result of 68.2%.

The objective of this paper is to evaluate the parallel performance for the basic problems. A survey is employed to illustrate how parallel computations for matrix multiplication, which are the most heavy part of most numeric computation programs, are accelerated by using the Intel®Xeon Phi™coprocessors. All programs are written in C programming language with OpenMP parallel programming library and the Intel®Language Extensions for Offload(LEO), compiled by Intel Compiler XE (version sp1.0.080), and executed on Nagoya University large scale symmetric multi-processing computing server Fujitsu Supercomputer CX 400/270 system. Each test program is repeated 10 times to take an average result.

In this work, a simple test program using directive offload pragmas was exploited to demonstrate how much percentage of the peak performance could be achieved, and a result of approximate 96% is observed. A comparison between the performance of parallel computation matrix multiplication executed on host Intel IvyBridge processors and fully offloaded to Intel Xeon Phi coprocessors was made and the coprocessor system performs an obviously higher calculation speed than the host processors. Furthermore, matrix multiplications computed in divided blocks are executed and the results illustrate a requirement of dividing the matrix reasonably as well as assigning threads appropriately.