第16回先進スーパーコンピューティング環境研究会（ASE研究会）実施報告

東京大学情報基盤センター　准教授　片桐孝洋

2013年12月4日（水）13:25〜17:30、東京大学情報基盤センター（浅野地区）4階遠隔会議室にて、第16回先進スーパーコンピューティング環境研究会（ASE研究会）が開催されました。

大学・研究機関からの参加者12名、企業からの参加者3名、無所属1名の合計16名の参加がありました。活発な議論がなされました。

今回は、招待講演として海外から4名の研究者の方を招待しました。

まず初めに、米国ローレンスバークレー国立研究所から Francois-Henry Rouet 博士をお呼びし、連立一次方程式の直接解法であるマルチフロンタル法に関する講演を行いました。次にフランスのボルドー大学から Samuel Thibault 博士をお呼びし、非均質計算環境向きのタスクスケジューラである StarPU に関する講演を行いました。次に、米国ローレンスバークレー国立研究所から Osni Marques 博士をお呼びし、行列計算と CG 法における並列言語 UPC での性能に関する講演を行いました。最後に、米国サンディエゴ・スーパーコンピュータセンターから Yifeng Cui 博士をお呼びし、ペタスケールに向けた地震シミュレーションの講演を行いました。

以上の招待講演者の講演内容は、数値アルゴリズム、システムソフトウェア、コンパイラ、アプリケーションと多岐にわたり、HPC分野の観点から、たいへん興味深い講演内容でした。会場からも、活発な質疑がなされました。

当日のプログラムを以下に載せます。

Program

- 13:25 – Opening
- 13:30 - 14:20 Invited Talk（1）
  - Dr. Francois-Henry Rouet (Lawrence Berkeley National Laboratory, USA)
  - Title： A parallel multifrontal solver that exploits hierarchically semi-separable representations
  - Co-authors： Xiaoye S. Li (LBNL), Artem Napov (University of Brussels), Jianlin Xia (Purdue), Maarten V. de Hoop (Purdue)
  - Abstract： Low-rank approximation techniques are an increasingly popular way of speeding-up sparse matrix algorithms.
    We focus on direct methods for the solution of sparse linear systems, and we show how to embed hierarchically semi-separable (HSS) representations into a multifrontal solver.
    These techniques allow to decrease both the operation complexity and the

memory footprint of the solver for many real-life applications.
We focus on a new parallel geometric code that aims at solving discretized Helmholtz equations on regular meshes.
We present experimental results on very large domains (200+ million unknowns on 3D problems) and up to 16,000 cores.
We also present some new research directions, such as the use of randomized sampling.

- 14:20 - 14:30 Break
- 14:30 - 15:20 Invited Talk (2)
  - Dr. Samuel Thibault (University of Bordeaux, France)
  - Title: StarPU: leveraging clusters of heterogeneous machines through dynamic task scheduling.
  - Abstract: Heterogeneous accelerator-based architectures are more and more seen in production HPC clusters, featuring multicore CPUs, GPU accelerators, and now even Xeon Phi accelerators, thus providing an unprecedented amount of processing power per node.
    It has thus become one of the biggest challenges in HPC to deal with such machines which expose such a highly unbalanced computing power.
    To fully tap into the potential of these heterogeneous machines, pure offloading approaches, that consist in running an application on regular processors while offloading part of the code on accelerators, are not sufficient.
    This talk will present the latest advances in the StarPU project, which aims at providing portable optimized performance on clusters of heterogeneous multicore+accelerator machines to task-based applications.
    The goal is to relieve the programmer from the technical aspects of data management and task scheduling, while applying theoretical task scheduling algorithms on actual application execution to improve performance.
    It also provides performance feedback through task profiling and trace analysis.
- 15:20 - 15:30 Break
- 15:30 - 16:20 Invited Talk (3)
  - Dr. Osni Marques (Lawrence Berkeley National Laboratory, USA)
  - Title: Performance of UPC Implementations of Matrix Computations and CG
  - Abstract: Partitioned Global Address Space (PGAS) languages offer programmers a shared memory view that can increase productivity and allow

locality exploitation to obtain good performance on large-scale distributed memory systems.

It is expected that the PGAS model, as provided by Co-array Fortran, Chapel, X10, XcalableMP and Unified Parallel C (UPC), will play an increasing role in upcoming systems as it circumvents memory limitations that may be found in MPI.

In this presentation we show results for a parallel numerical library for dense matrix computations using UPC, and for linear solvers based on Cholesky and LU factorizations that are built upon that library.

We also show results for four different implementation of CG in UPC, using 1D and 2D distributions of the (sparse) matrix involved in CG computations, and evaluate those implementations on two different platforms by means of a set of matrices that exhibit distinct sparse patterns.

- 16:20 - 16:30 Break
- 16:30 - 17:20 Invited Talk (4)
  - Dr. Yifeng Cui (San Diego Supercomputer Center, USA)
  - Title: Earthquake Ground Motion Modeling: Petascale and Beyond
  - Abstract: The future exascale systems will exhibit a substantially different balance among processor speed, system interconnect and memory bandwidth.

    Scientific applications must be radically re-designed and re-implemented for new programming models, new algorithms, and system tools to ensure a continued growth in sustained performance and productivity.

    In this presentation, I will introduce a finite difference earthquake forward code to illustrate the challenges and opportunities on large-scale heterogeneous systems.

    This GPU-based application implements high throughput, memory locality, communication reduction and communication/computation overlap.

    These performance tuning made possible to achieve 2.3 Pflop/s sustained performance on 16,384 OLCF Titan GPUs, while simulating realistic 0-10 Hz earthquake ground motions relevant to building engineering design.

    The simulations used a mesh comprising 443-billion elements in a calculation that includes both small-scale fault geometry and media complexity. Ground motion synthetics were computed using dynamic rupture propagation along a rough fault embedded in a 3D velocity structure with small-scale heterogeneities described by a statistical model.

- 17：30 – Closing
- 18：00 – Reception

図：当日の様子

ASE 研究会の開催情報はメーリングリストで発信をしております。研究会メーリングリストに参加ご希望の方は、ASE 研究会幹事の片桐（katagiri@cc.u-tokyo.ac.jp）までお知らせください。

以上