第 21 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

2015 年 12 月 2 日（水）13：00〜18:05、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 21 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。本研究会は、科学技術振興機構戦略的創造推進事業 CREST「自動チューニング機構を有するアプリケーション開発・実行環境」（研究代表：中島研吾（東京大学情報基盤センター））との共催です。

大学・研究機関からの参加者 19 名の合計 19 名の参加がありました。活発な議論がなされました。

第 21 回となる今回は、国外から 3 人、国内から 1 人の講演者を招待し、数値アルゴリズム、自動チューニングなど HPC（High-Performance Computing）に関する様々な分野について議論しました。また、若手・女性利用者推薦制度 H27 年度採択課題の採択者の発表も行われました。

まず海外からの講演者として、ドイツの German Aerospace Center（DLR)から、Achim Basermann 博士により、「The Highly Scalable Iterative Solver Library PHIST」と題した講演がなされました。この講演は、German Research Foundation（DFG)による、ESSEX (Equipping Sparse Solvers for Exascale)プロジェクトに関するもので、エクサスケール環境に向けた数値計算ソルバーに関する講演でした。

次に、米国ジョージア工科大学の Edmond Chow 准教授による「Preconditioning for Parallel Computers」と題した講演がありました。本公演では、並列反復解法の前処理に関するものです。厳密な前処理に基づく実装では並列性能が無く、近年のマルチコア計算機で性能向上が望めないという欠点が知られています。そこで Chow 准教授は、並列性能を高め、かつ非同期に計算可能な、新しい前処理方式の提案を行っています。実際の並列計算機を用いてその有効性を検証しています。

次に、国立台湾大学の王偉仲教授により「High-Performance Linear System and Eigenvalue Solvers for Frequency Domain Photonic Device Simulations」と題した講演がなされました。本講演では、光デバイスのシミュレーションにおいて現れる数値計算手法の話題で、特に固有値ソルバーにおける効率の良い計算手法と、近年のよく使われている GPU における効率的な実装に関する講演でした。

一方、若手の招待講演者として、東京女子大学の博士課程学生の小林由佳氏による「Accurate and efficient algorithms for solving ill-conditioned linear systems」と題した講演がありました。この講演では、解きにくい連立一次方程式の行列（悪条件行列）に対する、演算精度を保証する数値手法の提案です。実際に精度保証の数値計算手法を実装するためには、効率の良い実装が必要です。そこで、BLAS と呼ばれる線形代数副プログラム集を用いた実装法を示し、かつ、性能評価をおこなった発表がなされました。

一方、若手・女性利用者推薦制度 H27 年度採択課題の採択者による招待講演では、まず工学院大学の修士課程の学生である野村直也氏による「Performance analysis of SA-AMG method by setting near-kernel vectors」の講演がありました。幾何マルチグリッド法による効率の良い並列実装方式の開発の研究で、問題特性であるメッシュ情報を利用し、収束性のよい実装方式の提案が行われました。

最後に同制度の採択者で、鳥取大学の博士課程学生の井町宏人氏により、「Quantum Wave-packet Dynamics Simulation Solvers and Their Performance on Oakleaf-FX10」と題する講演がなされました。この講演では、大規模分子構造計算においては効率の良い固有値ソルバーが必要になりますが、固有値ソルバーの性能評価を FX10 スーパーコンピュータシステムで行った結果の紹介と、また、反復解法を用いる実装手法についての検討を行いました。

以上のように多種多様な最新成果について発表がなされました。また、若手・助成制度において英語による発表もなされましたが内容については、その他の講演者と比べても興味深い成果発表でした。活発な質疑応答がなされ、盛会のうちに終わりました。

当日のプログラムを以下に載せます。

21th Advanced Supercomputing Environment (ASE) Seminar
December 2nd (Wednesday), 2015, 13:00 - 18:05
4F Remote Conference Room, Information Technology Center (ITC), The University of Tokyo

- Program

13:00 - 13:05 Welcome Address

13:05 - 13:50 Achim Basermann (German Aerospace Center (DLR), Germany)

Title: The Highly Scalable Iterative Solver Library PHIST

Abstract: As modern supercomputers approach the exascale, many numerical libraries face scalability issues due to the massive increase in CPU cores compared to memory bandwidth. Sparse matrix algorithms, e.g. iterative linear and eigenvalue solvers, are particularly affected by the relatively slow memory subsystem.

In the German Research Foundation (DFG) project ESSEX (Equipping Sparse Solvers for Exascale), we developed the flexible software framework PHIST for implementing iterative methods on HPC systems. PHIST (Pipelined Hybrid Iterative Solver Toolkit) has been developed containing an interface to the existing numerical software framework Trilinos originally. PHIST also includes adapters to basic building block libraries so that high-level algorithm developments can benefit from high-performance

kernel implementations, e.g. sparse matrix-vector multiplication kernels.

Moreover, PHIST provides systematic and continuous testing of all software components and allows us to develop stable implementations of innovative iterative methods in an evolving hard- and software environment.

Using the PHIST framework, we present recent work on sparse eigenvalue solvers, from scalable basic operations to complete algorithms, with a focus on the difficult problem of finding interior eigenvalues. In this context, we also discuss the application of the CARP-CG method as a preconditioner for applications such as Graphene simulation.

13:50 - 14:00 Break

14:00 - 14:45 Edmond Chow (Georgia Institute of Technology, USA)

Title: Preconditioning for Parallel Computers

Abstract: Preconditioners are approximations to a matrix or to the inverse of a matrix, called implicit and explicit preconditioners, respectively.

Using an implicit preconditioner generally requires triangular solves or other operations that are difficult to parallelize. On the other hand, using an explicit preconditioner generally only requires easy-to-parallelize sparse matrix-vector products. However, many explicit preconditioners, such as sparse approximate inverses, are not able to provide the same convergence rates as implicit preconditioners such as incomplete factorizations, due to their local nature. In this talk, we first review recent research on parallel preconditioning and then present a new explicit preconditioner that is designed to propagate data globally at each step. Although it is expensive compute, this new preconditioner can lead to very fast convergence rates while also being very efficient to apply on parallel computers.

14:45 - 14:55 Break

14:55 - 15:40 Weichung Wang (National Taiwan University, Taiwan)

Title: High-Performance Linear System and Eigenvalue Solvers for Frequency Domain Photonic Device Simulations

Abstract: Large-scale linear systems and eigenvalue problems arise in simulations of three-dimensional photonics devices in frequency domain. In particular, we concern frequency domain photonic device simulations involving periodic or homogeneous structures. These three-dimensional devices can be modeled by the Maxwell vector wave

equations with or without absorptive (perfect matching layer) boundary conditions. To efficiently solve these time and memory consuming ill-conditioned problems, our newly developed compressed hierarchical Schur method (ChiS) provides enhanced computation performance.

CHiS takes advantages of the periodic or homogeneous structures to remove the redundant computations and storages. CHiS also uses domain decompositions to achieve coarse-grains parallelism in its elimination tree. The dense BLAS3 operations within the sub-matrices and the Schur complement allow us to achieve fine-grained parallelism on multicore CPU or GPU. Several numerical results in three-dimensional photonics devices including guided wave examples are demonstrated.

15:40 - 15:50 Break

15:50 - 16:20 Satoshi Ohshima (ITC, The University of Tokyo, Japan)

Title: Auto-Tuning for OpenACC: Utilization and Expansion of ppOpen-AT for OpenACC

Abstract: OpenACC makes GPU programming easy. But, in order to obtain high performance, users have to consider various optimization techniques and parameters. In the case of multi-core and many-core environment, we have developed Auto-Tuning language named ppOpen-AT since several years and shown the availability of it. In this study, we investigate the usability of ppOpen-AT for OpenACC. Moreover, we propose to expand ppOpen-AT for further optimization of OpenACC.

16:20 - 16:50 Yuka Kobayashi and Takeshi Ogita (Tokyo Woman's Christian University, Japan)

Title: Accurate and efficient algorithms for solving ill-conditioned linear systems

Abstract: We are concerned with accurate solutions of ill-conditioned linear systems by using floating-point arithmetic.

One of the standard methods of solving linear systems accurately is using LU factorization with the iterative refinement method. It is effective if the problem is not ill-conditioned. However, the method cannot work well for ill-conditioned problems. Another possibility is to use multiple-precision arithmetic. Although it can work for ill-conditioned problems, it requires significant computing time.

To remedy the defects, we propose algorithms based on a preconditioned technique using a result of an LU factorization performed in floating-point arithmetic. The proposed algorithms can provide accurate numerical solutions for ill-conditioned problems beyond the limit of the working precision. Moreover, it requires less

computational cost than the previous method using an approximate inverse of the coefficient matrix as a preconditioner.

We conducted numerical experiments using the proposed algorithms on MATLAB. Results of numerical experiments are presented for confirming the effectiveness of the proposed algorithm.

16:50 - 17:00 Break

- Special Session: Initiative on Promotion of Supercomputing for Young or Women Researchers

17:00 - 17:30 Naoya Nomura, Akihiro Fujii, Teruo Tanaka (Kogakuin University), Kengo Nakajima (The University of Tokyo)

Title: Performance analysis of SA-AMG method by setting near-kernel vectors

Abstract: SA-AMG (Smoothed Aggregation - Algebraic Multigrid Method) method is one of the fastest solvers for large scale linear equations. Convergence of SA-AMG method can be improved by setting near-kernel vectors. Our research aims to investigate effectiveness of setting multiple near-kernel vectors and the method to find important near-kernel vectors for fast convergence. At first stage, we are testing 3-dimensional elastic problem with near-kernel vectors that correspond to parallel transition and rotation. Iteration numbers and execution time with different number of kernel vectors are compared. In addition, we investigate a performance relationship between the number of near-kernel vectors and problem size.

17:30 - 18:00 Hiroto Imachi, Takeo Hoshi (Tottori University, Japan, and JST-CREST)

Title: Quantum Wave-packet Dynamics Simulation Solvers and Their Performance on Oakleaf-FX10

Abstract: We develop a parallel quantum wave-packet dynamics simulation program and investigate its parallel efficiency on Oakleaf-FX10. The program is now a dense solver and we also show preliminary results for a sparse solver.

With the help of recent development of supercomputers, large-scale quantum material simulations for next-generation material design are realizing. The authors have developed theory of massively parallel electronic structure calculations and a simulation code named ELSES (Extra-Large-Scale Electronic Structure calculations). We are recently developing a parallel quantum conduction simulation program as an extension of ELSES.

We model quantum conduction as the time-dependent Schrödinger equation (TDSE), i S dx/dt = H x. Here x, H, S are called a wave-packet vector, a Hamiltonian matrix and an overlap matrix, respectively. For real applications, we have to solve the TDSE whose the dimension m = 10^4 - 10^6 over the number of time steps n = 10^4 - 10^5. It involves computational challenges.

One way to decrease the computational cost to solve the TDSE is dimension reduction by eigenvector expansion. We suppose Hamiltonian H(t) changes slowly as H(t) = H_0 + H_1(t) with a relatively small H_1(t). Then using the solution of the generalized eigenvalue problem H_0 V = S V Λ, we can treat effects of the perturbation term H_1(t) directly and reduce the dimension of an approximate TDSE to solve by selecting eigenvectors to expand the wave-packet x. We show parallel performance of the TDSE solver on Oakleaf-FX10 with input from a real-world application.

The number of non-zero elements in the Hamiltonian and overlap matrices is determined by the number of atoms that each atom interact with. Matrices from systems we are interested in have intermediate property between dense and sparse. Therefore we are developing a sparse solver as a complement of the dense solver with eigenvector expansion. The sparse solver is based on sparse linear equations and we investigate which combination of a sparse linear equation solver and a preconditioner is optimal for our input matrices with LIS (Library of Iterative Solvers for linear systems), a software library of parallel iterated solvers for numerical linear algebra problems. We also show preliminary results on the performance of the sparse solver.

This research is partially supported by Initiative on Promotion of Supercomputing for Young or Women Researchers, Supercomputing Division, Information Technology Center, The University of Tokyo.

18:00 - 18:05 Closing Address

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 研究会幹事までお知らせください。

以上