Development of General Purpose Numerical Software Infrastructure for Large Scale Scientific Computing

Project Representative
Akira Nishida Research Institute for Information Technology, Kyushu University

Author
Akira Nishida Research Institute for Information Technology, Kyushu University

The Scalable Software Infrastructure Project was initiated as a national project in Japan, for the purpose of constructing a scalable software infrastructure for scientific computing. The project covered three areas: iterative solvers for linear systems, fast integral transforms, and their portable implementation.

Modular programming was adopted to enable users to write their codes by combining elementary mathematical operations. Implemented algorithms were selected from the viewpoint of scalability on massively parallel computing environments. Since the first release in September 2005, the codes have been used by thousands of research projects around the world.

Keywords: high performance computing, parallel algorithms, modular programming

1. Overview

Construction of a software Infrastructure for highly parallel computing environments requires precisely prediction of future hardware technologies, and design of scalable and portable software for these technologies.

The Scalable Software Infrastructure (SSI) Project was initiated in November 2002, as a national project in Japan, for the purpose of constructing a scalable software infrastructure [1], [2], [3]. Based on the policies, we have used various types of parallel computers, and carefully designed our libraries on them, to maintain portability and usability. The architectures include shared memory parallel computers, distributed-memory parallel computers, Linux-based PC clusters, and vector supercomputers. In 2003, we signed a contract with the IBM Watson Research Center on the joint study of the library implementation on massively parallel environments with tens of thousands of processors. Since 2006, the SSI project has been selected for a joint research with the Earth Simulator Center to port our libraries on massively parallel vector supercomputing environments. The results of the SSI project will be evaluated on larger supercomputers in the near future.

In the SSI project, we have studied object-oriented implementation of libraries, autotuning mechanisms, and scripting languages for the implemented libraries. The results were applied to a modular iterative solver library Lis and a fast Fourier transform library FFTSS. The libraries were written in C, and equipped with Fortran interfaces. We have also developed SILC, a simple interface for library collections, with an extension to the scripting language.

2. Lis: a Library of Iterative Solvers for Linear Systems

In the fields such as fluid dynamics and structural analysis, we must solve large-scale systems of linear equations to compute high resolution numerical solutions of partial differential equations. The subgroup released Lis, a library of iterative solvers and preconditioners for linear systems, with various sparse matrix storage formats. Supported solvers, preconditioners, and matrix storage formats are listed in Table 1-4. We present an example of the program using Lis in Fig. 1.

There are a variety of portable software packages that are applicable to the iterative solver of sparse linear systems. SPARSKIT is a toolkit for sparse matrix computations written in Fortran. PETSc is a C library for the numerical solution of partial differential equations and related problems, which is to be used in application programs written in C, C++, and Fortran. PETSc includes parallel implementations of iterative solvers and preconditioners based on MPI. Aztec is another library of parallel iterative solvers and preconditioners written in C. The library is fully parallelized using MPI. From the viewpoint of functionality, our library and all three of the libraries mentioned above support different sets of matrix storage formats, iterative solvers, and preconditioners. In addition, our library is parallelized using OpenMP to take multicore architectures into consideration. Feedbacks from the users have been applied to Lis, and Lis has been tested on various platforms from small PC clusters to massively parallel computers, including NEC's SX, IBM's Blue Gene, and Cray's XT series. The code of Lis has attained the vectorization ratio of 99.1% and the parallelization ratio of 99.99%. We show a comparison of the MPI version

Table 1 Solvers for linear equations

1.0.x	CG	Added in 1.1.x	CR
	BiCG		BiCR
	CGS		CRS
	BiCGSTAB		BiCRSTAB
	BiCGSTAB(l)		GPBiCR
	GPBiCG		BiCRSafe
	Orthomin(m)		FGMRES(m)
	GMRES(m)		IDR(s)
	TFQMR		MINRES
	Jacobi		
	Gauss-Seidel		
	SOR		

Table 2 Solvers for eigenproblems

Added in 1.2.x	Power Iteration
	Inverse Iteration
	Approximate Inverse Iteration
	Rayleigh Quotient Iteration
	Lanczos Iteration
	Subspace Iteration
	Conjugate Gradient
	Conjugate Residual

Table 3 Preconditioners

1.0.x	Jacobi	Added in 1.1.0	Crout ILU
	ILU(k)		ILUT
	SSOR		Additive Schwarz
	Hybrid		User defined preconditioner
	I+S		
	SA-AMG		
	SAINV		

Table 4 Matrix storage formats

Point	Compressed Row Storage
	Compressed Column Storage
	Modified Compressed Sparse Row
	Diagonal
	Ellpack-Itpack generalized diagonal
	Jagged Diagonal Storage
	Dense
	Coordinate
Block	Block Sparse Row
	Block Sparse Column
	Variable Block Row

```
LIS_MATRIX          A;
LIS_VECTOR          b,x;
LIS_SOLVER          solver;
int                 iter;
double              times,itimes,ptimes;

lis_initialize(&argc, &argv);
lis_matrix_create(LIS_COMM_WORLD,&A);
lis_vector_create(LIS_COMM_WORLD,&b);
lis_vector_create(LIS_COMM_WORLD,&x);
lis_solver_create(&solver);
lis_input(A,b,x,argv[1]);
lis_vector_set_all(1.0,b);
lis_solver_set_optionC(solver);
lis_solve(A,b,x,solver);
lis_solver_get_iters(solver,&iter);
lis_solver_get_times(solver,&times, &itimes,&ptimes);
printf("iter = %d time = %e (p=%e i=%e)\n",iter,times, ptimes, itimes);
lis_finalize();
```

Fig. 1 Example of the C program using Lis.

of Lis and PETSc in Fig. 2, for solving a three-dimensional Poisson equation (size: one million, number of nonzero entries: 26 million) on an SGI Altix 3700 with 32 processors.

In our project, we have designed and implemented scalable and robust algorithms of iterative solvers for linear equations and their preconditioning, derived from physical applications. In recent years, multilevel algorithms for large-scale linear equations, such as the algebraic multigrid (AMG), have been investigated by many researches. In most cases, multigrid methods show linear scalability, and the number of iteration counts is O(n) for a problem of size n. The algebraic multigrid method is based on a principle similar to the geometric multigrid, which utilizes the spatial information on physical problems, but this method differs from the geometric multigrid by considering the coefficient as a vertex-edge incidence matrix. In addition, by using the information on the elements and their relations, this method generates coarser level matrices without higher frequency errors. The complexity of the algebraic multigrid is equivalent to the geometric multigrid and can be applied to irregular or anisotropic problems. We proposed an efficient parallel implementation of the algebraic multigrid preconditioned conjugate gradient method based on the smoothed aggregation (SAAMGCG) and found that the proposed implementation provides the best performance as the problem size grows [38]. Currently, the algebraic multigrid is the most effective algorithm for the general-purpose preconditioning, and its scalability is also remarkable. We have implemented the algebraic multigrid in Lis, and have tested it in massively parallel environments. We present the weak scaling results for a two dimensional Poisson equation of dimension 49 million on 1,024 nodes of a Blue Gene system in Fig. 3.

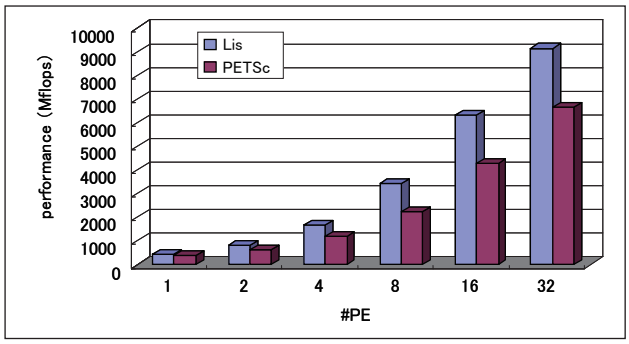

Fig. 2 Comparison of the MPI version of Lis and PETSc.

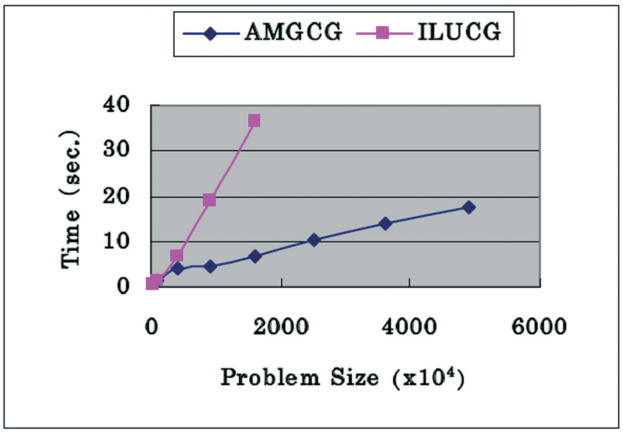

Fig. 3 Comparison of AMGCG and ILUCG.

The convergence of the Krylov subspace methods are much influenced by the rounding errors. Higher precision operations are effective for the improvement of convergence, although the arithmetic operations are costly. We implemented the quadruple precision operations on Lis, and accelerated them by using architecture specific SIMD instructions, such as Intel's SSE2 and IBM's FNA. To improve their performance, we also applied techniques such as loop unrolling. The computation time of our implementation is only 3.5 times as much as Lis' double precision, and 0.2 times as much as Intel Fortan's REAL*16.

Furthermore, we proposed the DQ-SWITCH algorithm, which efficiently switches the double precision iterations to the quadruple precision to reduce the computation time. The idea of the SIMD accelerated double-double precision operations has also been incorporated into Japan's next generation supercomputer project by RIKEN.

In the fields such as solid-state physics and quantum chemistry, efficient algorithms for eigenproblems for large-scale simulations derived from first principle calculation. There are several methods to compute eigenvalues of large-scale sparse matrices. Based on the observations, we proposed that the scalability of the conjugate gradient method can improve the performance of eigensolvers in parallel environments, where the extreme eigenvalues of a generalized eigenproblem can be solved by reducing these problems to the calculation of the local minimum of the Rayleigh quotient, combined with appropriate preconditioners, such as the algebraic multigrid. We have focused on the implementation of the existing major eigensolvers for sparse matrices on Lis, which was released as version 1.2.

The performance of iterative solvers is affected by the data structure of given matrices, the methodology of their parallelization, and the hierarchy of computer architectures. In the fiscal year 2009, we have studied the validity of the performance optimization of iterative solvers by benchmarking the FLOPS performance of matrix vector product kernels on given computing environments. Figure 4 shows the performance of a kernel spmvtest1, derived from a discretized 1-dimensional Poisson equation, for size up to 1,280,000 on a single node of SX-9 at JAMSTEC, and Fig. 5-7 show the performance for size up to 40,960,000 on the three scalar clusters at Kyushu University.

Although the scalar clusters show performance degradation after they reach their peak performance with the data size of 500kB to 1MB per core, SX-9 shows gradual performance increase until it reaches about 8-9GFLOPS per core (with the diagonal (DIA) format in this case), and keep it as the data size grows.

References

[1] A. Nishida, "SSI: Overview of simulation software infrastructure for large scale scientific applications (in Japanese)," IPSJ, Tech. Rep. 2004-HPC-098, 2004.

[2] A. Nishida, "Experience in Developing an Open Source Scalable Software Infrastructure in Japan," Lecture Notes in Computer Science, vol. 6017, pp. 87-98, 2010.

[3] A. Nishida, R. Suda, H. Hasegawa, K. Nakajima, D. Takahashi, H. Kotakemori, T. Kajiyama, A. Nukada, A. Fujii, Y. Hourai, S. L. Zhang, K. Abe, S. Itoh, and T. Sogabe, the Scalable Software Infrastructure for Scientific Computing Project, CREST, JST, 2009, http://www.ssisc.org/.

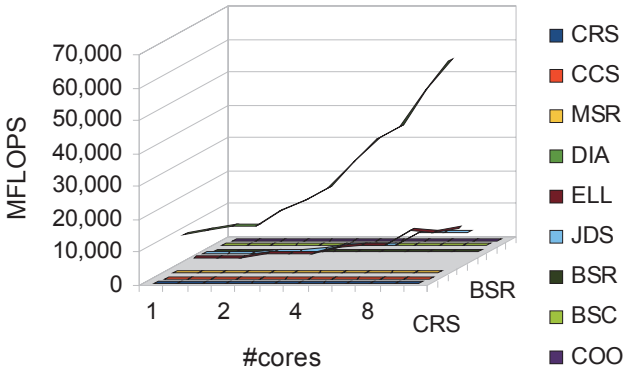

Fig. 4 Performance of spmvtest1 on a single node of the Earth Simulator 2.

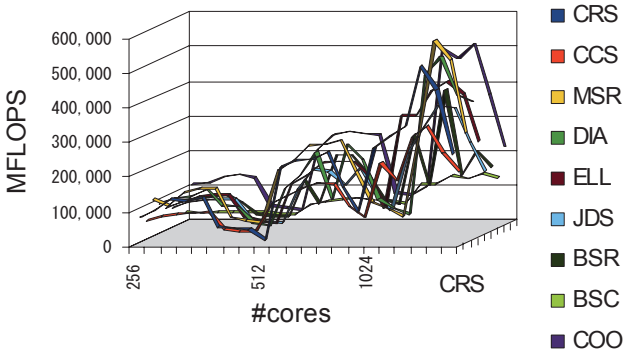

Fig. 5 Performance of spmvtest1 on the Fujitsu PRIMEGY Cluster at Kyushu University.

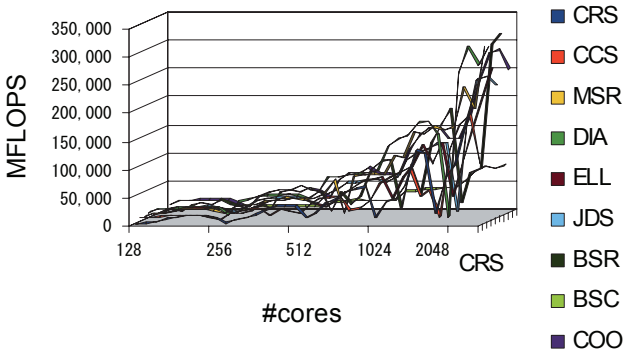

Fig. 6 Performance of spmvtest1 on the Fujitsu PRIMEQUEST Cluster at Kyushu University.

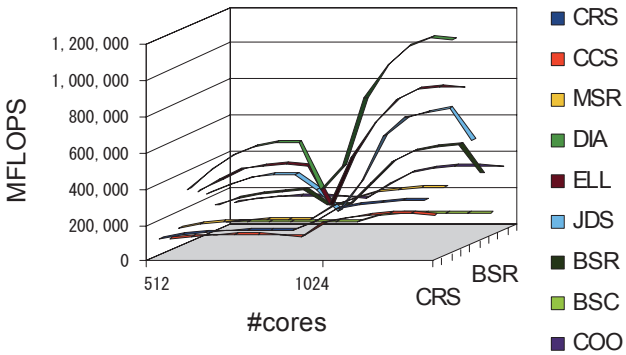

Fig. 7 Performance of spmvtest1 on the Hitachi SR16000 Cluster at Kyushu University.

大規模科学計算向け汎用数値ソフトウェア基盤の開発

プロジェクト責任者
西田　晃　九州大学情報基盤研究開発センター
著者
西田　晃　九州大学情報基盤研究開発センター

本プロジェクトでは、従来それぞれの分野において別個に進められてきた並列アルゴリズムや実装に関する知見をもとに、大規模化が予想される今後の計算環境に対応したスケーラブルなソフトウェア基盤を整備することを目的として、反復解法、高速関数変換、及びその効果的な計算機上への実装手法を中心に、平成 14 年度より科学技術振興機構戦略的創造研究推進事業の一環として、多様な計算機環境を想定した開発を行っている。モジュール化されたインタフェースを採用し、複雑な機能を持つライブラリを容易に構築できるようにするとともに、スケーラビリティの観点から並列化に適したアルゴリズムを開発、実装し、高並列な環境での使用に耐えるライブラリを実現している。本研究の成果はネットワークを通じて広く一般に配布し、フィードバックをもとにより汎用性の高いソフトウェアとしていく方針を採っており、平成 17 年 9 月よりソースコードを無償公開するとともに、ユーザの要望を反映した更新を適宜行なっている。平成 18 年度からは、地球シミュレータセンター共同プロジェクトの一環として、高並列なベクトル計算機環境への最適化を実施し、その成果をライブラリとして公開している。

反復解法においては、疎行列ベクトル積が計算時間の大半を占めることが多く、性能の最適化は最も重要な課題のひとつである。しかしながら、疎行列ベクトル積の性能は、行列の形状、並列化手法、メモリの階層構造等によって大きく性能が変化するため、すべての場合に最適な解法を見出すのは難しい。そこで、本年度は様々なデータ格納形式・並列化手法における疎行列ベクトル積の性能を事前に評価するためのベンチマークプログラムを Lis 上に実装・公開し、その有効性を複数の計算環境を用いて評価した。本ベンチマークは局所的なものであるが、これによって解法の計算性能をある程度正確に見積もることが可能になった。

キーワード : 高性能計算 , 並列アルゴリズム , モジュール化