

Overview of the HPC Challenge Benchmark Suite

Jack J. Dongarra
Oak Ridge National Laboratory and
University of Tennessee Knoxville

Piotr Luszczek
University of Tennessee Knoxville

Abstract—The HPC Challenge¹ benchmark suite has been released by the DARPA HPCS program to help define the performance boundaries of future Petascale computing systems. HPC Challenge is a suite of tests that examine the performance of HPC architectures using kernels with memory access patterns more challenging than those of the High Performance Linpack (HPL) benchmark used in the Top500 list. Thus, the suite is designed to augment the Top500 list, providing benchmarks that bound the performance of many real applications as a function of memory access characteristics e.g., spatial and temporal locality, and providing a framework for including additional tests. In particular, the suite is composed of several well known computational kernels (STREAM, HPL, matrix multiply – DGEMM, parallel matrix transpose – PTRANS, FFT, RandomAccess, and bandwidth/latency tests – b_eff) that attempt to span high and low spatial and temporal locality space. By design, the HPC Challenge tests are scalable with the size of data sets being a function of the largest HPL matrix for the tested system.

I. HIGH PRODUCTIVITY COMPUTING SYSTEMS

The DARPA High Productivity Computing Systems (HPCS) [1] is focused on providing a new generation of economically viable high productivity computing systems for national security and for the industrial user community. HPCS program researchers have initiated a fundamental reassessment of how we define and measure performance, programmability, portability, robustness and ultimately, productivity in the High End Computing (HEC) domain.

The HPCS program seeks to create trans-petaflops systems of significant value to the Government HPC community. Such value will be determined by assessing many additional factors beyond just theoretical peak flops (floating-point operations). Ultimately, the goal is to decrease the time-to-solution, which means decreasing both the execution time and development time of an application on a particular system. Evaluating the capabilities of a system with respect to these goals requires a different assessment process. The goal of the HPCS assessment activity is to prototype and baseline a process that can be transitioned to the acquisition community for 2010 procurements. As part of this effort we are developing a scalable benchmark for the HPCS systems.

The basic goal of performance modeling is to measure, predict, and understand the performance of a computer program or set of programs on a computer system. The applications of performance modeling are numerous, including evaluation of algorithms, optimization of code implementations, parallel

library development, and comparison of system architectures, parallel system design, and procurement of new systems.

This paper is organized as follows: sections II and III give motivation and overview of HPC Challenge while section IV provides more detailed description of the HPC Challenge tests and section V talks briefly about the scalability of the tests; sections VI, VII, and VIII describe the rules, the software installation process and some of the current results, respectively. Finally, section IX concludes the paper.

II. MOTIVATION

The DARPA High Productivity Computing Systems (HPCS) program has initiated a fundamental reassessment of how we define and measure performance, programmability, portability, robustness and, ultimately, productivity in the HPC domain. With this in mind, a set of computational kernels was needed to test and rate a system. The HPC Challenge suite of benchmarks consists of four local (matrix-matrix multiply, STREAM, RandomAccess and FFT) and four global (High Performance Linpack – HPL, parallel matrix transpose – PTRANS, RandomAccess and FFT) kernel benchmarks. HPC Challenge is designed to approximately bound computations of high and low spatial and temporal locality (see Fig. 1). In addition, because HPC Challenge kernels consist of simple mathematical operations, this provides a unique opportunity to look at language and parallel programming model issues. In the end, the benchmark is to serve both the system user and designer communities [2].

III. THE BENCHMARK TESTS

This first phase of the project have developed, hardened, and reported on a number of benchmarks. The collection of tests includes tests on a single processor (local) and tests over the complete system (global). In particular, to characterize the architecture of the system we consider three testing scenarios:

- 1) Local – only a single processor is performing computations.
- 2) Embarrassingly Parallel – each processor in the entire system is performing computations but they do not communicate with each other explicitly.
- 3) Global – all processors in the system are performing computations and they explicitly communicate with each other.

The HPC Challenge benchmark consists at this time of 7 performance tests: HPL [3], STREAM [4], RandomAccess, PTRANS, FFT (implemented using FFTE [5]), DGEMM [6],

¹This work was supported in part by the DARPA, NSF, and DOE through the DARPA HPCS program under grant FA8750-04-1-0219.

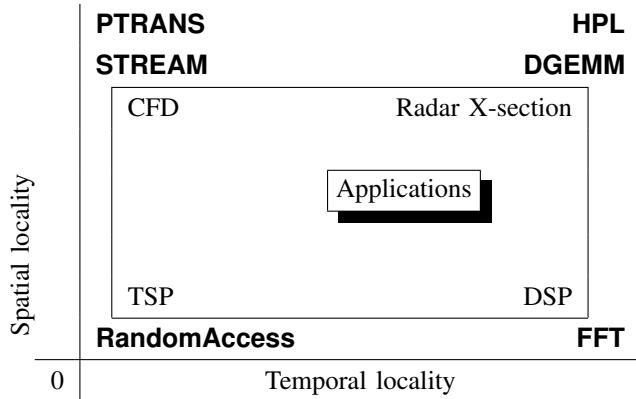


Fig. 1. Targeted application areas in the memory access locality space.

[7] and `b_eff` (MPI latency/bandwidth test) [8], [9], [10]. HPL is the Linpack TPP (toward peak performance) benchmark. The test stresses the floating point performance of a system. STREAM is a benchmark that measures sustainable memory bandwidth (in Gbyte/s), RandomAccess measures the rate of random updates of memory. PTRANS measures the rate of transfer for large arrays of data from multiprocessor’s memory. Latency/Bandwidth measures (as the name suggests) latency and bandwidth of communication patterns of increasing complexity between as many nodes as is time-wise feasible.

Many of the aforementioned tests were widely used before HPC Challenge was created. At first, this may seemingly make our benchmark merely a packaging effort. However, almost all components of HPC Challenge were augmented from their original form to provide consistent verification and reporting scheme. We should also stress the importance of running these very tests on a single machine and have the results available at once. The tests were useful separately for the HPC community before and with the unified HPC Challenge framework they create an unprecedented view of performance characterization of a system – a comprehensive view that captures the data under the same conditions and allows for variety of analysis depending on end user needs.

Each of the included tests examines system performance for various points of the conceptual spatial and temporal locality space shown in Fig. 1. The rationale for such selection of tests is to measure performance bounds on metrics important to HPC applications. The expected behavior of the applications is to go through various locality space points during runtime. Consequently, an application may be represented as a point in the locality space being an average (possibly time-weighted) of its various locality behaviors. Alternatively, a decomposition can be made into time-disjoint periods in which the application exhibits a single locality characteristic. The application’s performance is then obtained by combining the partial results from each period.

Another aspect of performance assessment addressed by HPC Challenge is ability to optimize benchmark code. For that we allow two different runs to be reported:

- Base run done with provided reference implementation.

- Optimized run that uses architecture specific optimizations.

The base run, in a sense, represents behavior of legacy code because it is conservatively written using only widely available programming languages and libraries. It reflects a commonly used approach to parallel processing sometimes referred to as hierarchical parallelism that combines Message Passing Interface (MPI) with threading from OpenMP. At the same time we recognize the limitations of the base run and hence we allow (or even encourage) optimized runs to be made. The optimizations may include alternative implementations in different programming languages using parallel environments available specifically on the tested system. To stress the productivity aspect of the HPC Challenge benchmark, we require that the information about the changes made to the original code be submitted together with the benchmark results. While we understand that full disclosure of optimization techniques may sometimes be impossible to obtain (due to for example trade secrets) we ask at least for some guidance for the users that would like to use similar optimizations in their applications.

IV. BENCHMARK DETAILS

Almost all tests included in our suite operate on either matrices or vectors. The size of the former we will denote below as n and the latter as m . The following holds throughout the tests:

$$n^2 \simeq m \simeq \text{Available Memory}$$

Or in other words, the data for each test is scaled so that the matrices or vectors are large enough to fill almost all available memory.

HPL (High Performance Linpack) is an implementation of the Linpack TPP (Toward Peak Performance) variant of the original Linpack benchmark which measures the floating point rate of execution for solving a linear system of equations. HPL solves a linear system of equations of order n :

$$Ax = b; \quad A \in \mathbf{R}^{n \times n}; \quad x, b \in \mathbf{R}^n$$

by first computing LU factorization with row partial pivoting of the n by $n + 1$ coefficient matrix:

$$P[A, b] = [[L, U], y].$$

Since the row pivoting (represented by the permutation matrix P) and the lower triangular factor L are applied to b as the factorization progresses, the solution x is obtained in one step by solving the upper triangular system:

$$Ux = y.$$

The lower triangular matrix L is left unpivoted and the array of pivots is not returned. The operation count for the factorization phase is $\frac{2}{3}n^3 - \frac{1}{2}n^2$ and $2n^2$ for the solve phase. Correctness of

the solution is ascertained by calculating the scaled residuals:

$$\frac{\|Ax - b\|_\infty}{\varepsilon \|A\|_1 n},$$

$$\frac{\|Ax - b\|_\infty}{\varepsilon \|A\|_1 \|x\|_1}, \quad \text{and}$$

$$\frac{\|Ax - b\|_\infty}{\varepsilon \|A\|_\infty \|x\|_\infty n},$$

where ε is machine precision for 64-bit floating-point values and n is the size of the problem.

DGEMM measures the floating point rate of execution of double precision real matrix-matrix multiplication. The exact operation performed is:

$$C \leftarrow \beta C + \alpha AB$$

where:

$$A, B, C \in \mathbf{R}^{n \times n}; \quad \alpha, \beta \in \mathbf{R}.$$

The operation count for the multiply is $2n^3$ and correctness of the operation is ascertained by calculating the scaled residual: $\|C - \hat{C}\| / (\varepsilon n \|C\|_F)$ (\hat{C} is a result of reference implementation of the multiplication).

STREAM is a simple benchmark program that measures sustainable memory bandwidth (in Gbyte/s) and the corresponding computation rate for four simple vector kernels:

$$\begin{aligned} \text{COPY:} \quad & c \leftarrow a \\ \text{SCALE:} \quad & b \leftarrow \alpha c \\ \text{ADD:} \quad & c \leftarrow a + b \\ \text{TRIAD:} \quad & a \leftarrow b + \alpha c \end{aligned}$$

where:

$$a, b, c \in \mathbf{R}^m; \quad \alpha \in \mathbf{R}.$$

As mentioned earlier, we try to operate on large data objects. The size of these objects is determined at runtime which contrasts with the original version of the STREAM benchmark which uses static storage (determined at compile time) and size. The original benchmark gives the compiler more information (and control) over data alignment, loop trip counts, etc. The benchmark measures Gbyte/s and the number of items transferred is either $2m$ or $3m$ depending on the operation. The norm of the difference between reference and computed vectors is used to verify the result: $\|x - \hat{x}\|$.

PTRANS (parallel matrix transpose) exercises the communications where pairs of processors exchange large messages simultaneously. It is a useful test of the total communications capacity of the system interconnect. The performed operation sets a random n by n matrix to a sum of its transpose with another random matrix:

$$A \leftarrow A^T + B$$

where:

$$A, B \in \mathbf{R}^{n \times n}.$$

The data transfer rate (in Gbyte/s) is calculated by dividing the size of n^2 matrix entries by the time it took to perform

the transpose. The scaled residual of the form $\|A - \hat{A}\| / (\varepsilon n)$ verifies the calculation.

RandomAccess measures the rate of integer updates to random memory locations (GUPS). The operation being performed on an integer array of size m is:

$$x \leftarrow f(x)$$

$$f: x \mapsto (x \oplus a_i); \quad a_i - \text{pseudo-random sequence}$$

where:

$$f: \mathbf{Z}^m \rightarrow \mathbf{Z}^m; \quad x \in \mathbf{Z}^m.$$

The operation count is $4m$ and since all the operations are in integral values over GF(2) field they can be checked exactly with a reference implementation. The verification procedure allows 1% of the operations to be incorrect (skipped) which allows loosening concurrent memory update semantics on shared memory architectures.

FFT measures the floating point rate of execution of double precision complex one-dimensional Discrete Fourier Transform (DFT) of size m :

$$Z_k \leftarrow \sum_j^m z_j e^{-2\pi i \frac{jk}{m}}; \quad 1 \leq k \leq m$$

where:

$$z, Z \in \mathbf{C}^m.$$

The operation count is taken to be $5m \log_2 m$ for the calculation of the computational rate (in Gflop/s). Verification is done with a residual $\|x - \hat{x}\| / (\varepsilon \log m)$ where \hat{x} is the result of applying a reference implementation of inverse transform to the outcome of the benchmarked code (in infinite-precision arithmetic the residual should be zero).

Communication bandwidth and latency is a set of tests to measure latency and bandwidth of a number of simultaneous communication patterns. The patterns are based on `b_eff` (effective bandwidth benchmark) – they are slightly different from the original `b_eff`. The operation count is linearly dependent on the number of processors in the tested system and the time the tests take depends on the parameters of the tested network. The checks are built into the benchmark code by checking data after it has been received.

V. SCALABILITY ISSUES FOR BENCHMARKS WITH SCALABLE INPUT DATA

A. Notation

The following symbols are used in the formulae:

- P – number of CPUs
- M – total size of system memory
- r – rate of execution (unit: Gflop/s, GUPS, etc.)
- t – time
- N – size of global matrix for HPL and PTRANS
- V – size of global vector for RandomAccess and FFT

B. Assumptions

Memory size per CPU is constant as the system grows. This assumption based on architectural design of almost all systems in existence. Hence, the total amount of memory available in the entire system is linearly proportional to the number of processors.

For HPL the dominant cost is CPU-related because computation has higher complexity order than communication: $\mathcal{O}(n^3)$ versus $\mathcal{O}(n^2)$.

C. Scalability of Time to Solution

Time complexity for HPL is $\mathcal{O}(n^3)$ (the hidden constant is 2/3) so the time to solution is:

$$t_{\text{HPL}} \propto \frac{N^3}{r_{\text{HPL}}} \quad (1)$$

Since N is a size of a square matrix then we need to take square root of available memory:

$$N \propto \sqrt{M} \propto \sqrt{P}$$

The rate of execution for HPL is determined by the number of processors since computations (rather than communication) dominates in terms of complexity:

$$r_{\text{HPL}} \propto P$$

This leads to:

$$t_{\text{HPL}} \propto \sqrt{P}$$

Time complexity for RandomAccess is $\mathcal{O}(n)$ so the time is:

$$t_{\text{RandomAccess}} \propto \frac{V}{r_{\text{RandomAccess}}} \quad (2)$$

The main table for RandomAccess should be as big as half of the total memory, so we get:

$$V \propto M \propto P$$

The rate of execution for RandomAccess can be argued to have various forms:

- If we assume that the interconnect scales with the number of processors then the rate would also scale:

$$r_{\text{RandomAccess}} \propto P$$

- Real-life experience tells that $r_{\text{RandomAccess}}$ is dependent on the interconnect and independent of the number of processors due to interconnect inefficiency:

$$r_{\text{RandomAccess}} \propto \mathbf{1}$$

Even worse, it is also conceivable that the rate is a decreasing function of processors:

$$r_{\text{RandomAccess}} \propto \frac{1}{P}$$

Conservatively assuming that $r_{\text{RandomAccess}} \propto \mathbf{1}$ it follows that:

$$t_{\text{RandomAccess}} \propto P$$

TABLE I

SUBMISSION DATA FOR NUMBER ONE ENTRIES ON THE FIRST AND THE LATEST EDITIONS OF THE TOP500 LISTS.

Date	R_{max}	N_{max}	P	t	\sqrt{P}
June 1993	59.7	52224	1024	1591	32
November 2005	280600	1769471	131072	13163	362

TABLE II

COMPUTER SYSTEMS USED IN TESTS. THE RANDOMACCESS NUMBER WAS OBTAINED DURING EARLY TESTING OF THE CURRENT CODE.

Processor	Interconnect	GUPS [GUPS]	Peak [Gflop/s]	Mem [GiB]	SMP [CPU]
Power4	Federation	0.00262	4*1.7	1	32
Cray X1	2D torus	0.14544	16*0.8	4	4
Opteron	Myrinet 2000	0.00319	2*1.4	1	2
Itanium 2	NUMalink	0.00305	4*1.5	8	4
Xeon	InfiniBand	0.00065	2*2.4	1	2

Peak (per processor) = flops per cycle * frequency

Mem – memory per processor

SMP – number of processors in each SMP node

D. Scalability Tests

It is easy to verify the theoretical framework from the previous section for the HPL test. Table I shows the relevant data from the first and the latest editions of the TOP500 list. The time to solution grew over 8 times while the square root of processor count grew over 11 times. This amounts to about 27% error which is relatively good achievement considering how much has changed during those 12 years including progress in software (algorithmic improvements) and hardware (chip architecture, materials, and fabrication).

The computer systems that were used in RandomAccess initial tests are listed in Table II. The table also shows the initially obtained RandomAccess number for each system. This number was used to gauge running time of RandomAccess in the final version of the code.

Table III shows estimates of time it takes to run HPL and RandomAccess using the analysis from the previous section. For a 256-CPU system RandomAccess takes much longer to run than HPL. The only exception is Cray X1: it has large amount of memory per CPU which makes for

TABLE III

ESTIMATED TIME IN MINUTES TO PERFORM FULL SYSTEM TEST OF HPL AND RANDOMACCESS.

Manufacturer	System	64 CPUs		256 CPUs	
		HPL	R.A.	HPL	R.A.
IBM	Power4	29.08	108.9	58.1	435.7
Cray	X1	123.6	7.8	247.2	31.4
Atipa	Opteron	70.6	89.6	141.2	358.4
SGI	Itanium 2	745.9	750.0	1491.9	3000.2
Voltaire	Xeon	41.2	440.2	82.4	1760.8

R.A.=RandomAccess

longer HPL run and large GUPS number which makes for short `RandomAccess` run. One of HPC Challenge' guiding principles was not to exceed running time beyond the twice the HPL's time. To keep the code in line with this principle meant reducing the time required by `RandomAccess` by artificially terminating the original algorithm and reflecting this change in the verification procedure. Both of these enhancements have been implemented in the current version of the code.

VI. RULES FOR RUNNING THE BENCHMARK

There must be one baseline run submitted for each computer system entered in the archive. There may also exist an optimized run for each computer system.

1) *Baseline Runs*

Optimizations as described below are allowed.

a) *Compile and load options*

Compiler or loader flags which are supported and documented by the supplier are allowed. These include porting, optimization, and preprocessor invocation.

b) *Libraries*

Linking to optimized versions of the following libraries is allowed:

- BLAS
- MPI

Acceptable use of such libraries is subject to the following rules:

- All libraries used shall be disclosed with the results submission. Each library shall be identified by library name, revision, and source (supplier). Libraries which are not generally available are not permitted unless they are made available by the reporting organization within 6 months.
- Calls to library subroutines should have equivalent functionality to that in the released benchmark code. Code modifications to accommodate various library call formats are not allowed.
- Only complete benchmark output may be submitted – partial results will not be accepted.

2) *Optimized Runs*

a) *Code modification*

Provided that the input and output specification is preserved, the following routines may be substituted:

- In HPL: `HPL_pdgesv()`, `HPL_pdt_rsv()` (factorization and substitution functions)
- no changes are allowed in the DGEMM component
- In PTRANS: `pdtrans()`
- In STREAM: `tuned_STREAM_Copy()`, `tuned_STREAM_Scale()`, `tuned_STREAM_Add()`, `tuned_STREAM_Triad()`

- In `RandomAccess`:

`MPIRandomAccessUpdate()` and `RandomAccessUpdate()`

- In FFT:

`fftw_malloc()`, `fftw_free()`, `fftw_one()`, `fftw_mpi()`, `fftw_create_plan()`, `fftw_destroy_plan()`, `fftw_mpi_create_plan()`, `fftw_mpi_local_sizes()`, `fftw_mpi_destroy_plan()` (all of these functions are compatible with FFTW 2.1.5 [11], [12])

- In `b_eff` component alternative MPI routines might be used for communication. But only standard MPI calls are to be performed and only to the MPI library that is widely available on the tested system.

b) *Limitations of Optimization*

i) *Code with limited calculation accuracy*

The calculation should be carried out in full precision (64-bit or the equivalent). However the substitution of algorithms is allowed (see next).

ii) *Exchange of the used mathematical algorithm*

Any change of algorithms must be fully disclosed and is subject to review by the HPC Challenge Committee. Passing the verification test is a necessary condition for such an approval. The substituted algorithm must be as robust as the baseline algorithm. For the matrix multiply in the HPL benchmark, Strassen Algorithm may not be used as it changes the operation count of the algorithm.

iii) *Using the knowledge of the solution*

Any modification of the code or input data sets, which uses the knowledge of the solution or of the verification test, is not permitted.

iv) *Code to circumvent the actual computation*

Any modification of the code to circumvent the actual computation is not permitted.

VII. SOFTWARE DOWNLOAD, INSTALLATION, AND USAGE

The reference implementation of the benchmark may be obtained free of charge at the benchmark's web site: <http://icl.cs.utk.edu/hpcc/>. The reference implementation should be used for the base run. The installation of the software requires creating a script file for Unix's `make(1)` utility. The distribution archive comes with script files for many common computer architectures. Usually, few changes to one of these files will produce the script file for a given platform.

After, a successful compilation the benchmark is ready to run. However, it is recommended that changes are made to the benchmark's input file such that the sizes of data to use during the run are appropriate for the tested system. The sizes

should reflect the available memory on the system and number of processors available for computations.

We have collected a comprehensive set of notes on the HPC Challenge benchmark. They can be found at <http://icl.cs.utk.edu/hpcc/faq/>.

VIII. EXAMPLE RESULTS

Fig. 2 shows a sample rendering of the results web page: http://icl.cs.utk.edu/hpcc/hpcc_results.cgi. It is impossible to show here all of the results for nearly 60 systems submitted so far to the web site. The results database is publicly available at the aforementioned address and can be exported to Excel spreadsheet or an XML file. Fig. 3 shows a sample kiviati diagram generated using the benchmark results. Kiviati diagrams can be generated at the website and allow easy comparative analysis for multi-dimensional results from the HPC Challenge database.

IX. CONCLUSIONS

No single test can accurately compare the performance of HPC systems. The HPC Challenge benchmark test suite stresses not only the processors, but the memory system and the interconnect. It is a better indicator of how an HPC system will perform across a spectrum of real-world applications. Now that the more comprehensive, informative HPC Challenge benchmark suite is available, it can be used in preference to comparisons and rankings based on single tests. The real utility of the HPC Challenge benchmarks are that architectures can be described with a wider range of metrics than just flop/s from HPL. When looking only at HPL performance and the Top500 List, inexpensive build-your-own clusters appear to be much more cost effective than more sophisticated HPC architectures. Even a small percentage of random memory accesses in real applications can significantly affect the overall performance of that application on architectures not designed to minimize or hide memory latency. HPC Challenge benchmarks provide users with additional information to justify policy and purchasing decisions. We expect to expand and perhaps remove some existing benchmark components as we learn more about the collection.

REFERENCES

- [1] "High Productivity Computer Systems," (<http://www.highproductivity.org/>).
- [2] W. Kahan, "The baleful effect of computer benchmarks upon applied mathematics, physics and chemistry," 1997, the John von Neumann Lecture at the 45th Annual Meeting of SIAM, Stanford University.
- [3] J. J. Dongarra, P. Luszczek, and A. Petitet, "The LINPACK benchmark: Past, present, and future," *Concurrency and Computation: Practice and Experience*, vol. 15, pp. 1–18, 2003.
- [4] J. McCalpin, "STREAM: Sustainable Memory Bandwidth in High Performance Computers," (<http://www.cs.virginia.edu/stream/>).
- [5] D. Takahashi and Y. Kanada, "High-performance radix-2, 3 and 5 parallel 1-D complex FFT algorithms for distributed-memory parallel computers," *The Journal of Supercomputing*, vol. 15, no. 2, pp. 207–228, 2000.
- [6] J. J. Dongarra, J. D. Croz, I. S. Duff, and S. Hammarling, "Algorithm 679: A set of Level 3 Basic Linear Algebra Subprograms," *ACM Transactions on Mathematical Software*, vol. 16, pp. 1–17, March 1990.
- [7] —, "A set of Level 3 Basic Linear Algebra Subprograms," *ACM Transactions on Mathematical Software*, vol. 16, pp. 18–28, March 1990.

- [8] A. E. Koniges, R. Rabenseifner, and K. Solchenbach, "Benchmark design for characterization of balanced high-performance architectures," in *Proceedings of the 15th International Parallel and Distributed Processing Symposium (IPDPS'01), Workshop on Massively Parallel Processing (WMPP)*, vol. 3. San Francisco, CA: In IEEE Computer Society Press, April 23-27 2001.
- [9] R. Rabenseifner and A. E. Koniges, "Effective communication and file-i/o bandwidth benchmarks," in *J. Dongarra and Yiannis Cotronis (Eds.), Recent Advances in Parallel Virtual Machine and Message Passing Interface, Proceedings of the 8th European PVM/MPI Users' Group Meeting, EuroPVM/MPI 2001*. Santorini, Greece: LNCS 2131, September 23-26 2001, pp. 24–35.
- [10] R. Rabenseifner, "Hybrid parallel programming on HPC platforms," in *Proceedings of the Fifth European Workshop on OpenMP, EWOMP '03*, Aachen, Germany, September 22-26 2003, pp. 185–194.
- [11] M. Frigo and S. G. Johnson, "FFTW: An adaptive software architecture for the FFT," in *Proc. 1998 IEEE Intl. Conf. Acoustics Speech and Signal Processing*, vol. 3. IEEE, 1998, pp. 1381–1384.
- [12] —, "The design and implementation of FFTW3," *Proceedings of the IEEE*, vol. 93, no. 2, 2005, special issue on "Program Generation, Optimization, and Adaptation".

http://icl.cs.utk.edu/hpcc/hpcc_results.cgi

HPCCHALLENGE HPC

Home Rules News Download FAQ Links Collaborators Sponsors Upload Results

Condensed Results - Base Runs Only - 56 Systems - Generated on Mon Apr 25 12:56:11 2005

System Information				G-HPL	G-PTRANS	G-Random Access	G-FFTE	EP-STREAM Triad	EPDGEMM	RandomRing Bandwidth	RandomRing Latency		
System	Processor	Speed	Count	Threads	Processes	TFlop/s	GB/s	Gup/s	GFlop/s	GB/s	GFlop/s	usec	
MA/PT/PS/PC/TH/PR/CM/CS/JC/IA/SD													
Atipa Conquest cluster	AMD Opteron	1.4GHz	128	1	128	0.2526	3.247			1.629		0.03627	23.68
Cray X1 MSP		0.8GHz	64	1	64	0.5216	3.229			14.990		0.94074	20.34
Cray X1 MSP		0.8GHz	60	1	60	0.5778	30.431			14.974		1.03291	20.83
Cray X1 MSP		0.8GHz	120	1	120	1.0610	2.460			8.496		0.83014	20.12
Cray T3E Alpha	21164	0.6GHz	1024	1	1024	0.0482	10.277			0.517		0.03174	12.09
Cray X1 MSP		0.8GHz	252	1	252	2.3847	97.408			14.914		0.42899	22.27
Cray X1 MSP		0.8GHz	124	1	124	1.2054	39.525			14.973		0.70857	20.15
Cray X1 MSP		0.8GHz	60	1	60	0.5087	1.634	0.003075	3.144	14.902	10.915	1.16779	14.66
Cray T3E Alpha	21164	.675GHz	512	1	512	0.2232	9.774	0.028946	15.477	0.532	0.661	0.03571	8.14
Cray XD1	AMD Opteron	2.2GHz	64	1	64	0.2239	10.592	0.022397	16.361	2.656	4.034	0.22697	1.63
Cray X1 MSP		.8GHz	32	1	32	0.2767	32.661	0.001662	2.965	14.870	8.258	1.41269	14.94

Fig. 2. Sample HPC Challenge results web page.




Fig. 3. Sample kiviati diagram of HPC Challenge results for three clusters with Opteron 2.2 GHz processors and three different interconnects.