

Application benchmark results for Big Red, an IBM e1350 BladeCenter Cluster

Craig A. Stewart¹, Matthew Link², D. Scott McCaulay², Greg Rodgers³, George Turner², David Hancock², Peng Wang⁴, Faisal Saied⁵, Marlon Pierce¹, Ross Aiken³, Matthias S. Mueller⁶, Matthias Jurenz⁶, Matthias Lieber⁶, Jenett Tillotson², and Beth A. Plale¹

¹ Pervasive Technology Institute, Indiana University

² University Information Technology Services, Indiana University

³ IBM Corporation

⁴ University Information Technology Services, Indiana University – Purdue University Indianapolis

⁵ Rosen Center for Advanced Computing, Purdue University

⁶ Center for Information Services and High Performance Computing (ZIH) Technische Universität Dresden

1. Introduction	2
2. Hardware	2
2.a. Basic building blocks and system components.....	2
2.b. Initial 20.48 TFLOPS configuration	2
2.c. Expanded 30.72 TFLOPS configuration.....	3
3. System software	3
4. System architecture benchmarks	3
4.a. System architecture benchmarks utilizing 1024 cores.....	4
4.b. System architecture benchmarks utilizing 2048 cores.....	5
5. Application benchmarks.....	7
6. Architecture benchmark output files	9
6.a. HPCC benchmark data for initial 1024 core configuration	9
6.b. HPCC benchmark data for initial 2048 core configuration	20
6.b.1. HPCC benchmark data for N=70000.....	20
6.b.2. HPCC benchmark data for N=10000.....	28
6.b.3. HPCC benchmark data for N=600000.....	36
6.c. Comparison of application benchmarks	45
6.c.1. WRF.....	45
6.c.2. OOCORE1	45
6.c.3. GAMESS1	46
6.c.4. MILC2	47
6.c.5. PARATEC2.....	48
6.c.6. HOMME3.....	48
7. Acknowledgements.....	49
8. References	49

1. Introduction

The purpose of this report is to present the results of benchmark tests with Big Red, an IBM e1350 BladeCenter Cluster. This report is particularly focused on providing details of system architecture and test run results in detail to allow for analysis in other reports and comparison with other systems, rather than presenting such analysis here.

Big Red is based on the open blade cluster architecture [1], [2], employing the IBM JS21 blade and PowerPC 970MP processor. Having decided in early 2006 to make a major investment in an HPC system, Indiana University (IU) issued a call for proposals for a large supercomputer cluster on 3 March 2006. The system, purchased on 13 April 2006, was selected based on analysis of the proposals submitted by supercomputer vendors. After the initial 20.4 TFLOPS implementation, Big Red was upgraded to 30.72 TFLOPS in spring 2007. This upgrade was part of an economic development initiative in the state of Indiana (US), led jointly by Indiana University, Purdue University, the Indiana Economic Development Corporation (IEDC), and IBM, Inc. [3].

2. Hardware

2.a. Basic building blocks and system components

Big Red is based on the IBM JS21 Bladeserver node (used as user login nodes and compute nodes), configured as follows:

- 2 x 2.5GHz dual-core PowerPC 970MP processors
- 8GB 533MHz Double Data Rate, version 2 (DDR2) Synchronous Dynamic Random Access Memory (SDRAM)
- 73GB Serial Attached Small Computer System Interface (SAS) disk (67GB in /scratch)
- 1 x Myricom M3S-PCIXD-2-I (Lanai XP)

In addition to the JS21 Bladeserver node, there are storage nodes based on the IBM 505 pSeries 505 configured as follows:

- 2 x 1.65GHz dual-core Power5+ processor
- 8 GB 533MHz DDR2 SDRAM
- 73GB SAS disk
- 2 x Emulex LP10000 PCI-X/133MHz FC adapters

2.b. Initial 20.48 TFLOPS configuration

Big Red's initial 20.48 TFLOPS configuration was a distributed shared-memory cluster, consisting of 512 IBM JS21 blades, each with two dual-core PowerPC 970 MP processors, 8GB of memory, a 73B SAS disk, and a PCI-X Myrinet 2000 adapter for high-bandwidth, low-latency Message Passing Interface (MPI) applications. In addition to local scratch disks, the Big Red compute nodes are connected via gigabit Ethernet to a 266TB GPFS file system, hosted on 16 IBM p505 Power5 systems. Two 256-port Myricom M3-CLOS-ENCL switches provide a 2+2Gb/s low-latency interconnect for the 512 compute nodes. The allocation of nodes is described in Table 1 below.

Use	Node type	#
User login	JS21	4
Compute	JS21	512
Storage	pSeries 505	16

Table 1.

The disk for Big Red's 266TB GPFS volume is hosted on 10 DataDirect Networks S2A 9500 storage controllers, each dual-pathed to 5 Service Availability Forum (standard) (SAF) 4248 chassis. Physical disks are aggregated in an 8+2 Redundant Array of Independent (or Inexpensive) Disks (RAID) configuration; combined with the dual-pathed controllers and active/active GPFS storage hosts, this provides multiple levels of redundancy for Big Red's storage space.

Home directories are available via NFSv3 over Gigabit Ethernet. They are housed in an IBM N5500 A20. They consist of two nodes. Each node contains:

- 2 Intel 2.8 GHz Xeon processors,
- 4 GB of Error Correction Codes (ECC) memory
- 512 MB of non-volatile random access memory (NVRAM)
- 4 integrated 10/100/1000 Ethernet ports
- 4 integrated 2 Gbps Fibre Channel ports
- 2 redundant hot-plug integrated power supplies and cooling fans
- 3 PCI-X expansion slots for additional FC HBAs or Ethernet NICs
- Infiniband (IB) cluster interconnect between the two processing nodes

2.c. Expanded 30.72 TFLOPS configuration

Big Red's 30.72 TFLOPS configuration is a distributed memory cluster, consisting of 768 IBM JS21 blades, each with two dual-core PowerPC 970 MP processors, 8GB of memory, a 73B SAS disk, and a PCI-X Myrinet 2000 adapter for high-bandwidth, low-latency MPI applications. In addition to local scratch disks, the Big Red compute nodes are connected via gigabit Ethernet to a 266TB GPFS file system, hosted on 16 IBM p505 Power5 systems. Four 256-port Myricom M3-CLOS-ENCL switches provide a 2+2Gb/s low-latency interconnect for the 768 compute nodes. The allocation of nodes is described in Table 2 below.

Use	Node type	#
User Login	JS21	4
Compute	JS21	768
Storage	pSeries 505	16

Table 2.

3. System software

The Big Red cluster runs the SuSE Linux Enterprise Server (SLES) operating system. Batch jobs are managed with IBM's LoadLeveler and the Moab Workload Manager. Big Red uses LoadLeveler to submit and monitor jobs. LoadLeveler relies on the Moab scheduler software for job scheduling, incorporating a fair share mechanism based on research system time used by each user trying to run a job.

Compilers on Big Red:

- GCC version 3.3.3
- IBM compilers:
 - C compiler: xlc version 8.0
 - C++ compiler xlc version 8.0
 - Fortran compiler: xlf/xlf90 version 10.1

Message passing libraries:

```
MPICH Version: 1.2.7
MPICH Release date: 2005/11/08 03:56:31
MPICH Patches applied: none
MPICH configure: --with-device=ch_mx -prefix=/N/soft/linux-sles9-ppc64/mpich-mx-ibm-64 -opt= -O -
lib= -WI,-rpath,/opt/mx/lib64,-rpath,/opt/mx/lib -L/opt/mx/lib64 -L/opt/mx/lib/ -lmyriexpress
MPICH Device: ch_mx
```

4. System architecture benchmarks

We ran the HPC Challenge Benchmarks [4] as specified in the National Science Foundation (NSF) Benchmark guidance document NSF 06-05 [5].

4.a. System architecture benchmarks utilizing 1024 cores

System Architecture benchmarks run on Big Red with its initial 20.48 TFLOPS configuration were run using SUSE Linux SLES 9, update 3.

The system benchmarks performed are summarized here. Detailed benchmark results are available in section 6. Benchmarks were run without modifications to the run and make files provided as part of the NSF benchmark suite.

- High Performance Linpack (HPL) – the Linpack TPP benchmark, which measures the floating point rate of execution for solving a linear system of equations.
 - Time to completion: 199.22 seconds
 - Speed: 3.346e+03 GFLOPS
- Double precision, General Equation, Matrix Multiply (DGEMM)M – measures the floating point rate of execution of double precision real matrix-matrix multiplication.
 - StarDGEMM
 - Minimum GFLOPS: 4.538
 - Average GFLOPS: 7.009
 - Maximum GFLOPS: 7.552
 - SingleDGEMM
 - Single DGEMM GFLOPS: 7.874
- STREAM – a simple synthetic benchmark program that measures sustainable memory bandwidth (in GB/s) and the corresponding computation rate for simple vector kernel.

- StarSTREAM

Function	Rate (GB/s)
Copy	2.986
Scale	3.023
Add	3.145
Triad	3.244

Table 3.

- SingleSTREAM

Function	Rate (GB/s)	Avg time	Min time	Max time
Copy	1.0220	0.0176	0.0175	0.0177
Scale	1.0678	0.0169	0.0168	0.0170
Add	1.1617	1.2032	1.0231	1.0233
Triad	1.2247	0.0235	0.0220	0.0240

Table 4.

- PTRANS (parallel matrix transpose) – exercises the communications where pairs of processors communicate with each other simultaneously. It is a useful test of the total communications capacity of the network.
 - Average CPU time to completion: 23.997 seconds
 - Average wall clock time to completion: 26.687 seconds
 - Average speed: 25.112 GB/s
- RandomAccess – measures the rate of integer random updates of memory (GUPS).
 - MPIRandomAccess
 - CPU time used: 62.250 seconds
 - Wall clock time used: 62.290 seconds
 - Updates per second [GUP/s]: 0.1799E+09
 - Updates/PE per second [GUP/s]: 0.00018E+09
 - StarRandomAccess
 - CPU time used: 11.830 seconds

- Wall clock time used: 11.836 seconds
- Average GUP/s: 0.0028 billion (10^9)
- SingleRandomAccess
 - Single GUP/s 0.004045
- FFTE – measures the floating-point rate of execution of double precision complex one-dimensional Discrete Fourier Transform (DFT).
 - MPIFFT
 - Generation time: 0.178 seconds
 - Tuning: 0.164 seconds
 - Computing: 3.567 seconds
 - Inverse FFT: 3.469 seconds
 - StarFFT
 - Generation time: 0.357 seconds
 - Tuning: 0.001 seconds
 - Computing: 0.481 seconds
 - Inverse FFT: 3.649 seconds
 - Average GFLOPS: 0.451
 - SingleFFT
 - GFLOP/S: 0.718
- Communication bandwidth and latency – a set of tests to measure latency and bandwidth of a number of simultaneous communications patterns; based on b_eff (effective bandwidth benchmark).
 - Max Ping Pong Latency: 0.005811 msec
 - Randomly Ordered Ring Latency: 0.017495 msec
 - Min Ping Pong Bandwidth: 195.131555 MB/s
 - Naturally Ordered Ring Bandwidth: 135.169320 MB/s
 - Randomly Ordered Ring Bandwidth: 23.885499 MB/s

4.b. System architecture benchmarks utilizing 2048 cores

System Architecture Benchmarks run on Big Red with its 30.72 TFLOPS configuration were run using the SUSE Linux SLES 9, update 3.

A subset of the previously described benchmarks was performed on the upgraded system and is summarized here. Detailed benchmark results are available in section 6. Benchmarks were run without modifications to the run and make files provided as part of the NSF benchmark suite.

- HPL – the Linpack TPP benchmark which measures the floating point rate of execution for solving a linear system of equations.
 - Time to completion: 10641.34 seconds
 - Speed: $1.353e+04$ GFLOPS
- DGEMM – measures the floating point rate of execution of double precision real matrix-matrix multiplication.
 - StarDGEMM
 - Minimum GFLOPS: 6.793
 - Average GFLOPS: 8.269
 - Maximum GFLOPS: 8.320
 - SingleDGEMM
 - Single DGEMM GFLOPS: 8.494

Function	Rate (GB/s)
Copy	2.959

Scale	3.0509
Add	3.318
Triad	3.312

Table 5.

- STREAM – a simple synthetic benchmark program that measures sustainable memory bandwidth (in GB/s) and the corresponding computation rate for simple vector kernel.
 - StarSTREAM
 - Not performed.
 - SingleSTREAM
 - Not performed.
- PTRANS (parallel matrix transpose) – exercises the communications where pairs of processors communicate with each other simultaneously. It is a useful test of the total communications capacity of the network.
 - Average wall clock time to completion: 26.95 seconds
 - Average CPU time to completion: 26.71 seconds
 - Average speed: 28.955 GB/s
- RandomAccess – measures the rate of integer random updates of memory (GUPS).
 - MPIRandomAccess
 - CPU time used: 61.37 seconds
 - Wall clock time used: 61.41 seconds
 - Updates per second [GUP/s]: 0.2497 billion (10⁹)
 - Updates/PE per second [GUP/s]: 0.00012 billion (10⁹)
 - StarRandomAccess
 - CPU time used: 211.78 seconds
 - Wall clock time used: 211.81 seconds
 - Average GUP/s: 0.0025 billion (10⁹)
 - SingleRandomAccess
 - GUP/s: 0.0035 billion (10⁹)
- FFTE – measures the floating-point rate of execution of double precision complex one-dimensional Discrete Fourier Transform (DFT).

Function	Rate (GB/s)	Avg time	Min time	Max time
Copy	1.0377	0.0176	0.0175	0.0177
Scale	1.0497	0.0169	0.0168	0.0170
Add	1.1782	1.2032	1.0231	1.0233
Triad	1.941	0.0235	0.0220	0.0240

Table 6.

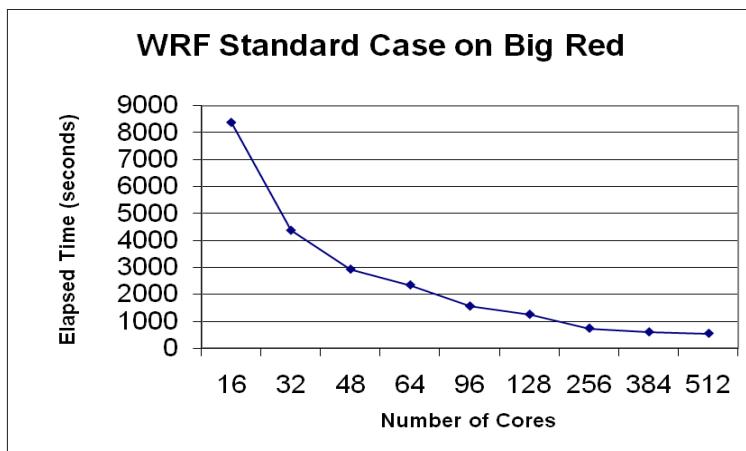
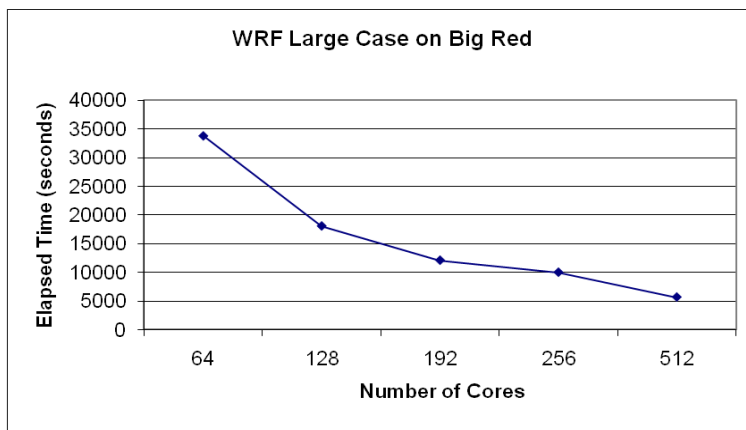
- MPIFFT
 - Generation time: 2.846 seconds
 - Tuning: 2.592 seconds
 - Computing: 43.377 seconds
 - Inverse FFT: 44.124 seconds
- StarFFT
 - Generation time: 5.694 seconds
 - Tuning: 0.003 seconds
 - Computing: 10.507 seconds
 - Inverse FFT: 11.355 seconds
 - Average GFLOPS: 0.399

- SingleFFT
 - GFLOPS: 0.436
- Communication bandwidth and latency – a set of tests to measure latency and bandwidth of a number of simultaneous communications patterns; based on b_eff (effective bandwidth benchmark).
 - Max Ping Pong Latency: 0.006124 msec
 - Randomly Ordered Ring Latency: 0.017735 msec
 - Min Ping Pong Bandwidth: 195.954309 MB/s
 - Naturally Ordered Ring Bandwidth: 92.852883 MB/s
 - Randomly Ordered Ring Bandwidth: 21.201995 MB/s

5. Application benchmarks

Use of Big Red is aimed at a large user community with diverse computational needs. Performance of several applications was measured to provide information about expected performance under real-use conditions.

- WRF – A next-generation mesoscale numerical weather prediction system designed to serve both operational forecasting and atmospheric research needs [6]. Expanded charted data are in section 6.c.1.



- OOCORE1 – Out-of-core solver software code developed by the SCALAPACK group at the University of Tennessee at Knoxville [7].
- GAMESS1 – General Atomic and Molecular Electronic Structure System. a general ab initio quantum chemistry package [8]. Expanded data are in section 6.c.3.

# Processors	# Nodes	# Cores	Time				Linear Scaling				Scaling Efficiency			
			NSF Small	NSF Medium	NSF Large	NSF X-Large	NSF Small	NSF Medium	NSF Large	NSF X-Large	NSF Small	NSF Medium	NSF Large	NSF X-Large
1	1	1												
1	1	2												
2	1	4												
4	2	8												
8	4	16		14508				1				1		
16	8	32		7917				7254				0.91625616		
32	16	64		4626				3627				0.78404669		
64	32	128		3090	23459			1813.5	1			0.5868932	1	
96	48	192		2843	15757			1209	15639.3333			0.42525501	0.99253242	
128	64	256		2837	12103			906.75	11729.5			0.31961579	0.96913988	
160	80	320			9976				9383.6				0.94061748	
192	96	384			8593				7819.66667				0.91000427	
224	112	448			7545				6702.57143				0.91000427	
256	128	512			6869				5864.75				0.85379968	
384	192	768			5283				3909.83333				0.74007824	
448	224	896			4923				3351.28571				0.68074055	
512	256	1024			4641				2932.375				0.6318412	
1024	512	2048												
1536	768	3072												
2048	1024	4096												

Table 7. GAMESS benchmarks.

- MILC2 – Large-scale numerical simulations software to study quantum chromodynamics (QCD), the theory of the strong interactions of subatomic physics [9]. Expanded charted data are in section 6.c.4.

		Core Count								
		1	16	32	64	128	256	512	1024	2048
Test Case	small	92.581								
	medium		1683.432	865.036	447.845	235.026	122.015			
	large				11316.27	5781.716	2964.596	1537.605	817	452
	x-large								6554	3695
	small Target	91								
	medium Target		1680	855	443	242	143			
	large Target				11821	6282	3499	2102	1371	

Table 8.

- PARATEC2 – Parallel Total Energy Code. A package designed primarily for a massively parallel computing platform. The code performs ab-initio quantum-mechanical total energy calculations using pseudopotentials and a plane wave basis set [10]. Expanded charted data are in section 6.c.5.

		Core Count					
		4	32	64	128	256	512
Test Case	small	4					
	medium		1323	737	518		
	large			3794	2415	2068	
	IU case1						
	IU case2						
	IU case3						

Table 9. Tuning level.

- HOMME3 – High Order Methods Modeling Environment, tools to create a high-performance scalable global atmospheric model [11].

6. Architecture benchmark output files

6.a. HPCCL benchmark data for initial 1024 core configuration

```

Hostname: 's1c1b1'
#####
=====
HPLinpack 1.0a -- High-Performance Linpack benchmark -- January 20, 2004
Written by A. Petit et and R. Clint Whaley, Innovative Computing Labs., UTK
=====

```

An explanation of the input/output parameters follows:

```

T/V      : Wall time / encoded variant.
N        : The order of the coefficient matrix A.
NB       : The partitioning blocking factor.
P        : The number of process rows.
Q        : The number of process columns.
Time     : Time in seconds to solve the linear system.
Gflops   : Rate of execution for solving the linear system.

```

The following parameter values will be used:

```

N      : 100000
NB     : 200
PMAP   : Row-major process mapping
P      : 31
Q      : 96

```

```

PFACT : Left
NBMIN : 4
NDIV : 2
RFACT : Right
BCAST : 2ringM
DEPTH : 0
SWAP : Mix (threshold = 64)
L1 : transposed form
U : transposed form
EQUIL : yes
ALIGN : 16 double precision words

```

```

-----
- The matrix A is randomly generated for each test.
- The following scaled residual checks will be computed:
  1) ||Ax-b||_oo / ( eps * ||A||_1 * N )
  2) ||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 )
  3) ||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo )
- The relative machine precision (eps) is taken to be 1.110223e-16
- Computational tests pass if scaled residuals are less than 16.0

```

Begin of PTRANS section.

```

M: 50000
N: 50000
MB: 200 100 200
NB: 200 100 200
P: 31
Q: 96

```

	M	N	MB	NB	P	Q	TIME	CHECK	GB/s	RESID
WALL	50000	50000	200	200	31	96	49.60	PASSED	0.403	0.00
CPU	50000	50000	200	200	31	96	49.62	PASSED	0.403	0.00
WALL	50000	50000	100	100	31	96	12.97	PASSED	1.542	0.00
CPU	50000	50000	100	100	31	96	13.00	PASSED	1.538	0.00
WALL	50000	50000	200	200	31	96	49.52	PASSED	0.404	0.00
CPU	50000	50000	200	200	31	96	49.55	PASSED	0.404	0.00

```

Finished 3 tests, with the following results:
  3 tests completed and passed residual checks.
  0 tests completed and failed residual checks.
  0 tests skipped because of illegal input values.

```

```

END OF TESTS.
Current time (1182228795) is Tue Jun 19 00:53:15 2007

```

```

End of PTRANS section.
Begin of HPL section.

```

```

=====
HPLinpack 1.0a -- High-Performance Linpack benchmark -- January 20, 2004
Written by A. Petitet and R. Clint Whaley, Innovative Computing Labs., UTK
=====

```

```

An explanation of the input/output parameters follows:
T/V : Wall time / encoded variant.
N : The order of the coefficient matrix A.
NB : The partitioning blocking factor.
P : The number of process rows.
Q : The number of process columns.
Time : Time in seconds to solve the linear system.
Gflops : Rate of execution for solving the linear system.

```

The following parameter values will be used:

```

N : 100000
NB : 200
PMAP : Row-major process mapping

```

```

P      :      31
Q      :      96
PFACT  :      Left
NBMIN  :      4
NDIV   :      2
RFACT  :      Right
BCAST  :      2ringM
DEPTH  :      0
SWAP   :      Mix (threshold = 64)
L1     :      transposed form
U      :      transposed form
EQUIL  :      yes
ALIGN  :      16 double precision words

```

```

-----
- The matrix A is randomly generated for each test.
- The following scaled residual checks will be computed:
  1) ||Ax-b||_oo / ( eps * ||A||_1 * N )
  2) ||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 )
  3) ||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo )
- The relative machine precision (eps) is taken to be      1.110223e-16
- Computational tests pass if scaled residuals are less than      16.0

```

```

=====
T/V          N      NB      P      Q          Time          Gflops
-----
WR03R2L4    100000  200    31    96          111.25          5.993e+03

```

```

=====
||Ax-b||_oo / ( eps * ||A||_1 * N ) =      0.0197702 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) =      0.0132985 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) =      0.0023911 ..... PASSED
=====

```

```

Finished      1 tests with the following results:
              1 tests completed and passed residual checks,
              0 tests completed and failed residual checks,
              0 tests skipped because of illegal input values.

```

```

-----
End of Tests.

```

```

=====
Current time (1182228906) is Tue Jun 19 00:55:06 2007

```

```

End of HPL section.
Begin of StarDGEMM section.
Scaled residual: 0.00618034
Node(s) with error 0
Minimum Gflop/s 4.538357
Average Gflop/s 7.008804
Maximum Gflop/s 7.551706
Current time (1182228907) is Tue Jun 19 00:55:07 2007

```

```

End of StarDGEMM section.
Begin of SingleDGEMM section.
Node(s) with error 0
Node selected 2000
Single DGEMM Gflop/s 7.873863
Current time (1182228908) is Tue Jun 19 00:55:08 2007

```

```

End of SingleDGEMM section.
Begin of StarSTREAM section.

```

```

-----
This system uses 8 bytes per DOUBLE PRECISION word.

```

```

-----
Array size = 1120071, Offset = 0
Total memory required = 0.0250 GB.

```

Each test is run 10 times, but only the *best* time for each is used.

Your clock granularity/precision appears to be 1 microseconds.
Each test below will take on the order of 11660 microseconds.
(= 11660 clock ticks)
Increase the size of the arrays if this shows that
you are not getting at least 20 clock ticks per test.

WARNING -- The above is only a rough guideline.
For best results, please be sure you know the
precision of your system timer.

Function Rate (GB/s) Avg time Min time Max time
Copy: 1.0220 0.0176 0.0175 0.0177
Scale: 1.0678 0.0169 0.0168 0.0170
Add: 1.1617 0.0232 0.0231 0.0233
Triad: 1.2247 0.0235 0.0220 0.0240

Results Comparison:
 Expected : 1291778759355468800.000000 258355751871093760.000000
344474335828124992.000000
 Observed : 1291778759353353728.000000 258355751869478400.000000
344474335828124992.000000
Solution Validates

Node(s) with error 0
Minimum Copy GB/s 0.972553
Average Copy GB/s 1.036761
Maximum Copy GB/s 1.496659
Minimum Scale GB/s 0.990273
Average Scale GB/s 1.040788
Maximum Scale GB/s 1.562295
Minimum Add GB/s 1.042977
Average Add GB/s 1.140035
Maximum Add GB/s 1.529104
Minimum Triad GB/s 1.075315
Average Triad GB/s 1.173333
Maximum Triad GB/s 2.186392
Current time (1182228909) is Tue Jun 19 00:55:09 2007

End of StarSTREAM section.
Begin of SingleSTREAM section.
Node(s) with error 0
Node selected 1143
Single STREAM Copy GB/s 2.985886
Single STREAM Scale GB/s 3.023113
Single STREAM Add GB/s 3.144787
Single STREAM Triad GB/s 3.243859
Current time (1182228909) is Tue Jun 19 00:55:09 2007

End of SingleSTREAM section.
Begin of MPIRandomAccess section.

This is the DARPA/DOE HPC Challenge Benchmark version 1.0.0 October 2003
Produced by Jack Dongarra and Piotr Luszczek
Innovative Computing Laboratory
University of Tennessee Knoxville and Oak Ridge National Laboratory

See the source files for authors of specific codes.
Compiled on Jun 18 2007 at 11:44:38
Current time (1182262355) is Tue Jun 19 10:12:35 2007

Hostname: 's1clb1'

=====

HPLinpack 1.0a -- High-Performance Linpack benchmark -- January 20, 2004

An explanation of the input/output parameters follows:

T/V : Wall time / encoded variant.
 N : The order of the coefficient matrix A.
 NB : The partitioning blocking factor.
 P : The number of process rows.
 Q : The number of process columns.
 Time : Time in seconds to solve the linear system.
 Gflops : Rate of execution for solving the linear system.

The following parameter values will be used:

N : 100000
 NB : 200
 PMAP : Row-major process mapping
 P : 32
 Q : 32
 PFACT : Left
 NBMIN : 4
 NDIV : 2
 RFACT : Right
 BCAST : 2ringM
 DEPTH : 0
 SWAP : Mix (threshold = 64)
 L1 : transposed form
 U : transposed form
 EQUIL : yes
 ALIGN : 16 double precision words

-
- The matrix A is randomly generated for each test.
 - The following scaled residual checks will be computed:
 - 1) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * N)$
 - 2) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$
 - 3) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$
 - The relative machine precision (eps) is taken to be 1.110223e-16
 - Computational tests pass if scaled residuals are less than 16.0

Begin of PTRANS section.

M: 50000
 N: 50000
 MB: 200 100 200
 NB: 200 100 200
 P: 32
 Q: 32

TIME	M	N	MB	NB	P	Q	TIME	CHECK	GB/s	RESID
WALL	50000	50000	200	200	32	32	1.26	PASSED	15.859	0.00
CPU	50000	50000	200	200	32	32	1.27	PASSED	15.748	0.00
WALL	50000	50000	100	100	32	32	1.26	PASSED	15.931	0.00
CPU	50000	50000	100	100	32	32	1.27	PASSED	15.748	0.00
WALL	50000	50000	200	200	32	32	1.26	PASSED	15.873	0.00
CPU	50000	50000	200	200	32	32	1.26	PASSED	15.873	0.00

Finished 3 tests, with the following results:
 3 tests completed and passed residual checks.
 0 tests completed and failed residual checks.
 0 tests skipped because of illegal input values.

END OF TESTS.

Current time (1182262368) is Tue Jun 19 10:12:48 2007

End of PTRANS section.

Begin of HPL section.

```

=====
HPLinpack 1.0a -- High-Performance Linpack benchmark -- January 20, 2004
Written by A. Petitet and R. Clint Whaley, Innovative Computing Labs., UTK
=====

```

An explanation of the input/output parameters follows:

```

T/V      : Wall time / encoded variant.
N        : The order of the coefficient matrix A.
NB       : The partitioning blocking factor.
P        : The number of process rows.
Q        : The number of process columns.
Time     : Time in seconds to solve the linear system.
Gflops   : Rate of execution for solving the linear system.

```

The following parameter values will be used:

```

N        : 100000
NB       : 200
PMAP     : Row-major process mapping
P        : 32
Q        : 32
PFACT   : Left
NBMIN    : 4
NDIV     : 2
RFACT   : Right
BCAST   : 2ringM
DEPTH   : 0
SWAP    : Mix (threshold = 64)
L1      : transposed form
U       : transposed form
EQUIL   : yes
ALIGN   : 16 double precision words

```

```

-----
- The matrix A is randomly generated for each test.
- The following scaled residual checks will be computed:
  1) ||Ax-b||_oo / ( eps * ||A||_1 * N )
  2) ||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 )
  3) ||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo )
- The relative machine precision (eps) is taken to be 1.110223e-16
- Computational tests pass if scaled residuals are less than 16.0

```

```

=====
T/V          N      NB      P      Q          Time          Gflops
-----
WR03R2L4    100000  200    32    32          199.22          3.346e+03

```

```

-----
||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0146222 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0098357 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0017685 ..... PASSED
-----

```

```

Finished      1 tests with the following results:
              1 tests completed and passed residual checks,
              0 tests completed and failed residual checks,
              0 tests skipped because of illegal input values.

```

End of Tests.

```

=====
Current time (1182262568) is Tue Jun 19 10:16:08 2007

```

```

End of HPL section.
Begin of StarDGEMM section.
Scaled residual: 0.00510712
Node(s) with error 0

```

Minimum Gflop/s 7.262569
Average Gflop/s 7.564677
Maximum Gflop/s 7.745041
Current time (1182262570) is Tue Jun 19 10:16:10 2007

End of StarDGEMM section.
Begin of SingleDGEMM section.
Node(s) with error 0
Node selected 823
Single DGEMM Gflop/s 8.091957
Current time (1182262572) is Tue Jun 19 10:16:12 2007

End of SingleDGEMM section.
Begin of StarSTREAM section.

This system uses 8 bytes per DOUBLE PRECISION word.

Array size = 3255208, Offset = 0
Total memory required = 0.0728 GB.
Each test is run 10 times, but only
the *best* time for each is used.

Your clock granularity/precision appears to be 1 microseconds.
Each test below will take on the order of 33359 microseconds.
(= 33359 clock ticks)
Increase the size of the arrays if this shows that
you are not getting at least 20 clock ticks per test.

WARNING -- The above is only a rough guideline.
For best results, please be sure you know the
precision of your system timer.

Function	Rate (GB/s)	Avg time	Min time	Max time
Copy:	1.3551	0.0463	0.0384	0.0497
Scale:	1.3465	0.0476	0.0387	0.0504
Add:	1.5210	0.0644	0.0514	0.0683
Triad:	1.4984	0.0640	0.0521	0.0682

Results Comparison:
Expected : 3754233929531250176.000000 750846785906249984.000000
1001129047875000064.000000
Observed : 3754233929605240320.000000 750846785916409344.000000
1001129047886046464.000000
Solution Validates

Node(s) with error 0
Minimum Copy GB/s 1.032456
Average Copy GB/s 1.783751
Maximum Copy GB/s 3.059742
Minimum Scale GB/s 1.033463
Average Scale GB/s 1.745111
Maximum Scale GB/s 2.983031
Minimum Add GB/s 1.173118
Average Add GB/s 1.975090
Maximum Add GB/s 3.382328
Minimum Triad GB/s 1.176597
Average Triad GB/s 1.976104
Maximum Triad GB/s 3.383656
Current time (1182262575) is Tue Jun 19 10:16:15 2007

End of StarSTREAM section.
Begin of SingleSTREAM section.
Node(s) with error 0
Node selected 96
Single STREAM Copy GB/s 3.016144
Single STREAM Scale GB/s 2.940390
Single STREAM Add GB/s 3.319387

Single STREAM Triad GB/s 3.318245
Current time (1182262576) is Tue Jun 19 10:16:16 2007

End of SingleSTREAM section.
Begin of MPIRandomAccess section.
Running on 1024 processors (PowerofTwo)
Total Main table size = $2^{33} = 8589934592$ words
PE Main table size = $2^{23} = 8388608$ words/PE
Default number of updates (RECOMMENDED) = 34359738368
Number of updates EXECUTED = 11208680448 (for a TIME BOUND of 60.00 secs)
CPU time used = 62.250000 seconds
Real time used = 62.290993 seconds
0.179940629 Billion(10^9) Updates per second [GUP/s]
0.000175723 Billion(10^9) Updates/PE per second [GUP/s]
Verification: CPU time used = 25.810000 seconds
Verification: Real time used = 27.247278 seconds
Found 0 errors in 8589934592 locations (passed).
Current time (1182262667) is Tue Jun 19 10:17:47 2007

End of MPIRandomAccess section.
Begin of StarRandomAccess section.
Main table size = $2^{23} = 8388608$ words
Number of updates = 33554432
CPU time used = 11.830000 seconds
Real time used = 11.835544 seconds
0.002835056 Billion(10^9) Updates per second [GUP/s]
Found 0 errors in 8388608 locations (passed).
Node(s) with error 0
Minimum GUP/s 0.002813
Average GUP/s 0.002834
Maximum GUP/s 0.002875
Current time (1182262691) is Tue Jun 19 10:18:11 2007

End of StarRandomAccess section.
Begin of SingleRandomAccess section.
Node(s) with error 0
Node selected 686
Single GUP/s 0.004045
Current time (1182262708) is Tue Jun 19 10:18:28 2007

End of SingleRandomAccess section.
Begin of MPIFFT section.
Number of nodes: 1024
Vector size: 1073741824
Generation time: 0.178
Tuning: 0.164
Computing: 3.567
Inverse FFT: 3.649
 $\max(|x-x_0|)$: 1.788e-15
Current time (1182262716) is Tue Jun 19 10:18:36 2007

End of MPIFFT section.
Begin of StarFFT section.
Vector size: 2097152
Generation time: 0.357
Tuning: 0.001
Computing: 0.481
Inverse FFT: 0.526
 $\max(|x-x_0|)$: 1.617e-15
Node(s) with error 0
Minimum Gflop/s 0.415081
Average Gflop/s 0.450625
Maximum Gflop/s 0.480677
Current time (1182262718) is Tue Jun 19 10:18:38 2007

End of StarFFT section.
Begin of SingleFFT section.

Node(s) with error 0
Node selected 350
Single FFT Gflop/s 0.717585
Current time (1182262719) is Tue Jun 19 10:18:39 2007

End of SingleFFT section.
Begin of LatencyBandwidth section.

Latency-Bandwidth-Benchmark R1.5.1 (c) HLRS, University of Stuttgart
Written by Rolf Rabenseifner, Gerrit Schulz, and Michael Speck, Germany

Details - level 2

MPI_Wtime granularity.
Max. MPI_Wtick is 0.000001 sec
wtick is set to 0.000001 sec

Message Length: 8
Latency min / avg / max: 0.005677 / 0.005677 / 0.005677 msecs
Bandwidth min / avg / max: 1.409 / 1.409 / 1.409 MByte/s

MPI_Wtime granularity is ok.

message size: 8
max time : 10.000000 secs
latency for msg: 0.005677 msecs
estimation for ping pong: 0.510961 msecs
max number of ping pong pairs = 19570
max client pings = max server pongs = 139
stride for latency = 11
Message Length: 8
Latency min / avg / max: 0.004366 / 0.004993 / 0.005811 msecs
Bandwidth min / avg / max: 1.377 / 1.605 / 1.832 MByte/s

Message Length: 2000000
Latency min / avg / max: 10.192990 / 10.192990 / 10.192990 msecs
Bandwidth min / avg / max: 196.213 / 196.213 / 196.213 MByte/s

MPI_Wtime granularity is ok.

message size: 2000000
max time : 30.000000 secs
latency for msg: 10.192990 msecs
estimation for ping pong: 81.543922 msecs
max number of ping pong pairs = 367
max client pings = max server pongs = 19
stride for latency = 59
Message Length: 2000000
Latency min / avg / max: 10.137081 / 10.178978 / 10.249496 msecs
Bandwidth min / avg / max: 195.132 / 196.484 / 197.295 MByte/s

Message Size: 8 Byte
Natural Order Latency: 0.006104 msec
Natural Order Bandwidth: 1.310720 MB/s
Avg Random Order Latency: 0.017494 msec
Avg Random Order Bandwidth: 0.457310 MB/s

Message Size: 2000000 Byte
Natural Order Latency: 14.796257 msec
Natural Order Bandwidth: 135.169320 MB/s
Avg Random Order Latency: 83.732812 msec
Avg Random Order Bandwidth: 23.885499 MB/s

Execution time (wall clock) = 63.332 sec on 1024 processes
- for cross ping_pong latency = 4.310 sec
- for cross ping_pong bandwidth = 26.381 sec
- for ring latency = 0.443 sec

- for ring bandwidth = 32.198 sec

Latency-Bandwidth-Benchmark R1.5.1 (c) HLRS, University of Stuttgart
Written by Rolf Rabenseifner, Gerrit Schulz, and Michael Speck, Germany

Major Benchmark results:

Max Ping Pong Latency: 0.005811 msec
Randomly Ordered Ring Latency: 0.017494 msec
Min Ping Pong Bandwidth: 195.131555 MB/s
Naturally Ordered Ring Bandwidth: 135.169320 MB/s
Randomly Ordered Ring Bandwidth: 23.885499 MB/s

Detailed benchmark results:

Ping Pong:

Latency min / avg / max: 0.004366 / 0.004993 / 0.005811 msec
Bandwidth min / avg / max: 195.132 / 196.484 / 197.295 MByte/s

Ring:

On naturally ordered ring: latency= 0.006104 msec, bandwidth= 135.169320 MB/s
On randomly ordered ring: latency= 0.017494 msec, bandwidth= 23.885499 MB/s

Benchmark conditions:

The latency measurements were done with 8 bytes
The bandwidth measurements were done with 2000000 bytes
The ring communication was done in both directions on 1024 processes
The Ping Pong measurements were done on
- 8742 pairs of processes for latency benchmarking, and
- 324 pairs of processes for bandwidth benchmarking,
out of $1024 \cdot (1024 - 1) = 1047552$ possible combinations on 1024 processes.
(1 MB/s = 10^6 byte/sec)

Current time (1182262782) is Tue Jun 19 10:19:42 2007

End of LatencyBandwidth section.

Begin of Summary section.

VersionMajor=1
VersionMinor=0
VersionMicro=0
VersionRelease=b
LANG=C
Success=1
sizeof_char=1
sizeof_short=2
sizeof_int=4
sizeof_long=8
sizeof_void_ptr=8
sizeof_size_t=8
sizeof_float=4
sizeof_double=8
sizeof_s64Int=8
sizeof_u64Int=8
CommWorldProcs=1024
MPI_Wtick=9.536743e-07
HPL_Tflops=3.34646
HPL_time=199.22
HPL_eps=1.11022e-16
HPL_RnormI=4.0893e-09
HPL_Anorm1=25189.8
HPL_AnormI=25157.5
HPL_Xnorm1=148665
HPL_XnormI=8.2789

HPL_N=100000
HPL_NB=200
HPL_nprow=32
HPL_npcol=32
HPL_depth=0
HPL_nbdiv=2
HPL_nbmin=4
HPL_cpfact=L
HPL_crfact=R
HPL_ctop=3
HPL_order=R
HPL_dMACH_EPS=1.110223e-16
HPL_dMACH_SFMIN=2.225074e-308
HPL_dMACH_BASE=2.000000e+00
HPL_dMACH_PREC=2.220446e-16
HPL_dMACH_MLEN=5.300000e+01
HPL_dMACH_RND=1.000000e+00
HPL_dMACH_EMIN=-1.021000e+03
HPL_dMACH_RMIN=2.225074e-308
HPL_dMACH_EMAX=1.024000e+03
HPL_dMACH_RMAX=1.797693e+308
HPL_sMACH_EPS=5.960464e-08
HPL_sMACH_SFMIN=1.175494e-38
HPL_sMACH_BASE=2.000000e+00
HPL_sMACH_PREC=1.192093e-07
HPL_sMACH_MLEN=2.400000e+01
HPL_sMACH_RND=1.000000e+00
HPL_sMACH_EMIN=-1.250000e+02
HPL_sMACH_RMIN=1.175494e-38
HPL_sMACH_EMAX=1.280000e+02
HPL_sMACH_RMAX=3.402823e+38
dweeps=1.110223e-16
sweps=5.960464e-08
HPLMaxProcs=1024
HPLMinProcs=1024
DGEMM_N=1562
StarDGEMM_Gflops=7.56468
SingleDGEMM_Gflops=8.09196
PTRANS_GBs=15.931
PTRANS_time=1.25541
PTRANS_residual=0
PTRANS_n=50000
PTRANS_nb=100
PTRANS_nprow=32
PTRANS_npcol=32
MPIRandomAccess_N=8589934592
MPIRandomAccess_time=62.291
MPIRandomAccess_CheckTime=27.2473
MPIRandomAccess_Errors=0
MPIRandomAccess_ErrorsFraction=0
MPIRandomAccess_ExeUpdates=11208680448
MPIRandomAccess_GUPs=0.179941
MPIRandomAccess_TimeBound=60
RandomAccess_N=8388608
StarRandomAccess_GUPs=0.00283447
SingleRandomAccess_GUPs=0.00404454
STREAM_VectorSize=3255208
STREAM_Threads=1
StarSTREAM_Copy=1.78375
StarSTREAM_Scale=1.74511
StarSTREAM_Add=1.97509
StarSTREAM_Triad=1.9761
SingleSTREAM_Copy=3.01614
SingleSTREAM_Scale=2.94039
SingleSTREAM_Add=3.31939
SingleSTREAM_Triad=3.31824
FFT_N=2097152

```

StarFFT_Gflops=0.450625
SingleFFT_Gflops=0.717585
MPIFFT_N=1073741824
MPIFFT_Gflops=45.1478
MPIFFT_maxErr=1.78836e-15
MaxPingPongLatency_usec=5.81145
RandomlyOrderedRingLatency_usec=17.4936
MinPingPongBandwidth_GBytes=0.195132
NaturallyOrderedRingBandwidth_GBytes=0.135169
RandomlyOrderedRingBandwidth_GBytes=0.0238855
MinPingPongLatency_usec=4.36604
AvgPingPongLatency_usec=4.99336
MaxPingPongBandwidth_GBytes=0.197295
AvgPingPongBandwidth_GBytes=0.196484
NaturallyOrderedRingLatency_usec=6.10352
FFTEblk=16
FFTEnp=8
FFTEl2size=1048576
M_OPENMP=-1
omp_get_num_threads=0
omp_get_max_threads=0
omp_get_num_procs=0
MemProc=-1
MemSpec=-1
MemVal=-1
MPIFFT_time0=3.09944e-06
MPIFFT_time1=0.915237
MPIFFT_time2=0.377161
MPIFFT_time3=0.862327
MPIFFT_time4=0.47812
MPIFFT_time5=0.916633
MPIFFT_time6=3.09944e-06
End of Summary section.
#####
End of HPC Challenge tests.
Current time (1182262782) is Tue Jun 19 10:19:42 2007

```

```
#####
```

6.b. **HPCC benchmark data for initial 2048 core configuration**

6.b.1. *HPCC benchmark data for N=70000*

```
#####
This is the DARPA/DOE HPC Challenge Benchmark version 1.0.0 October 2003
Produced by Jack Dongarra and Piotr Luszczek
Innovative Computing Laboratory
University of Tennessee Knoxville and Oak Ridge National Laboratory

```

```

See the source files for authors of specific codes.
Compiled on Jun 18 2007 at 11:44:38
Current time (1182266522) is Tue Jun 19 11:22:02 2007

```

```

Hostname: 's15c3b3.dim'
#####
=====
HPLinpack 1.0a -- High-Performance Linpack benchmark -- January 20, 2004
Written by A. Petitet and R. Clint Whaley, Innovative Computing Labs., UTK
=====

```

```

An explanation of the input/output parameters follows:
T/V      : Wall time / encoded variant.
N        : The order of the coefficient matrix A.
NB       : The partitioning blocking factor.
P        : The number of process rows.
Q        : The number of process columns.
Time     : Time in seconds to solve the linear system.

```

Gflops : Rate of execution for solving the linear system.

The following parameter values will be used:

N : 70000
NB : 200
PMAP : Row-major process mapping
P : 16
Q : 32
PFACT : Left
NBMIN : 4
NDIV : 2
RFACT : Right
BCAST : 2ringM
DEPTH : 0
SWAP : Mix (threshold = 64)
L1 : transposed form
U : transposed form
EQUIL : yes
ALIGN : 16 double precision words

- The matrix A is randomly generated for each test.
- The following scaled residual checks will be computed:
1) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * N)$
2) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$
3) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$
- The relative machine precision (eps) is taken to be 1.110223e-16
- Computational tests pass if scaled residuals are less than 16.0

Begin of PTRANS section.

M: 35000
N: 35000
MB: 200 100 200
NB: 200 100 200
P: 16
Q: 32

	TIME	M	N	MB	NB	P	Q	TIME	CHECK	GB/s	RESID
WALL	35000	35000	200	200	16	32	0.81	PASSED	12.064	0.00	
CPU	35000	35000	200	200	16	32	0.82	PASSED	11.951	0.00	
WALL	35000	35000	100	100	16	32	0.71	PASSED	13.726	0.00	
CPU	35000	35000	100	100	16	32	0.73	PASSED	13.425	0.00	
WALL	35000	35000	200	200	16	32	0.82	PASSED	12.018	0.00	
CPU	35000	35000	200	200	16	32	0.82	PASSED	11.951	0.00	

Finished 3 tests, with the following results:
3 tests completed and passed residual checks.
0 tests completed and failed residual checks.
0 tests skipped because of illegal input values.

END OF TESTS.

Current time (1182266528) is Tue Jun 19 11:22:08 2007

End of PTRANS section.

Begin of HPL section.

=====
HPLinpack 1.0a -- High-Performance Linpack benchmark -- January 20, 2004
Written by A. Petitet and R. Clint Whaley, Innovative Computing Labs., UTK
=====

An explanation of the input/output parameters follows:

T/V : Wall time / encoded variant.
N : The order of the coefficient matrix A.
NB : The partitioning blocking factor.
P : The number of process rows.

Q : The number of process columns.
 Time : Time in seconds to solve the linear system.
 Gflops : Rate of execution for solving the linear system.

The following parameter values will be used:

N : 70000
 NB : 200
 PMAP : Row-major process mapping
 P : 16
 Q : 32
 PFACT : Left
 NBMIN : 4
 NDIV : 2
 RFACT : Right
 BCAST : 2ringM
 DEPTH : 0
 SWAP : Mix (threshold = 64)
 L1 : transposed form
 U : transposed form
 EQUIL : yes
 ALIGN : 16 double precision words

 The matrix A is randomly generated for each test.
 -- The following scaled residual checks will be computed:
 1) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * N)$
 2) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$
 3) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$
 - The relative machine precision (eps) is taken to be 1.110223e-16
 - Computational tests pass if scaled residuals are less than 16.0

T/V	N	NB	P	Q	Time	Gflops
WR03R2L4	70000	200	16	32	122.87	1.861e+03

					$\frac{\ Ax-b\ _{\infty}}{\ A\ _{\infty}} / (\text{eps} * \ A\ _{\infty} * N)$	0.0074097 PASSED
					$\frac{\ Ax-b\ _{\infty}}{\ A\ _{\infty}} / (\text{eps} * \ A\ _{\infty} * \ x\ _{\infty})$	0.0100375 PASSED
					$\frac{\ Ax-b\ _{\infty}}{\ A\ _{\infty}} / (\text{eps} * \ A\ _{\infty} * \ x\ _{\infty})$	0.0016702 PASSED
=====						

Finished 1 tests with the following results:
 1 tests completed and passed residual checks,
 0 tests completed and failed residual checks,
 0 tests skipped because of illegal input values.

 End of Tests.

=====

Current time (1182266652) is Tue Jun 19 11:24:12 2007
 End of HPL section.
 Begin of StarDGEMM section.
 Scaled residual: 0.00505038
 Node(s) with error 0
 Minimum Gflop/s 5.596380
 Average Gflop/s 7.287612
 Maximum Gflop/s 7.635199
 Current time (1182266655) is Tue Jun 19 11:24:15 2007

End of StarDGEMM section.
 Begin of SingleDGEMM section.
 Node(s) with error 0
 Node selected 402
 Single DGEMM Gflop/s 8.067628
 Current time (1182266657) is Tue Jun 19 11:24:17 2007

End of SingleDGEMM section.
Begin of StarSTREAM section.

This system uses 8 bytes per DOUBLE PRECISION word.

Array size = 3190104, Offset = 0
Total memory required = 0.0713 GB.
Each test is run 10 times, but only
the *best* time for each is used.

Your clock granularity/precision appears to be 1 microseconds.
Each test below will take on the order of 34368 microseconds.
(= 34368 clock ticks)

Increase the size of the arrays if this shows that
you are not getting at least 20 clock ticks per test.

WARNING -- The above is only a rough guideline.
For best results, please be sure you know the
precision of your system timer.

Function Rate (GB/s) Avg time Min time Max time
Copy: 1.0233 0.0500 0.0499 0.0501
Scale: 1.0230 0.0500 0.0499 0.0503
Add: 1.1475 0.0670 0.0667 0.0673
Triad: 1.1906 0.0671 0.0643 0.0675

Results Comparison:

Expected : 3679149435468749824.000000 735829887093750016.000000
981106516124999936.000000

Observed : 3679149435538703872.000000 735829887103518720.000000
981106516135525632.000000

Solution Validates

Node(s) with error 0

Minimum Copy GB/s 0.992160
Average Copy GB/s 1.029710
Maximum Copy GB/s 1.086132
Minimum Scale GB/s 0.988011
Average Scale GB/s 1.030611
Maximum Scale GB/s 1.078901
Minimum Add GB/s 1.073509
Average Add GB/s 1.162837
Maximum Add GB/s 1.200585
Minimum Triad GB/s 1.075087
Average Triad GB/s 1.168802
Maximum Triad GB/s 1.456752
Current time (1182266659) is Tue Jun 19 11:24:19 2007

End of StarSTREAM section.
Begin of SingleSTREAM section.

Node(s) with error 0
Node selected 368
Single STREAM Copy GB/s 3.020930
Single STREAM Scale GB/s 2.946100
Single STREAM Add GB/s 3.314271
Single STREAM Triad GB/s 3.318587
Current time (1182266660) is Tue Jun 19 11:24:20 2007

End of SingleSTREAM section.
Begin of MPIRandomAccess section.
Running on 512 processors (PowerofTwo)
Total Main table size = 2^{32} = 4294967296 words
PE Main table size = 2^{23} = 8388608 words/PE
Default number of updates (RECOMMENDED) = 17179869184
Number of updates EXECUTED = 9200124928 (for a TIME BOUND of 60.00 secs)
CPU time used = 57.540000 seconds

Real time used = 57.549086 seconds
0.159865701 Billion(10^9) Updates per second [GUP/s]
0.000312238 Billion(10^9) Updates/PE per second [GUP/s]
Verification: CPU time used = 23.180000 seconds
Verification: Real time used = 23.601585 seconds
Found 0 errors in 4294967296 locations (passed).
Current time (1182266743) is Tue Jun 19 11:25:43 2007

End of MPIRandomAccess section.
Begin of StarRandomAccess section.
Main table size = 2^{23} = 8388608 words
Number of updates = 33554432
CPU time used = 11.820000 seconds
Real time used = 11.830897 seconds
0.002836170 Billion(10^9) Updates per second [GUP/s]
Found 0 errors in 8388608 locations (passed).
Node(s) with error 0
Minimum GUP/s 0.002805
Average GUP/s 0.002836
Maximum GUP/s 0.002885
Current time (1182266767) is Tue Jun 19 11:26:07 2007

End of StarRandomAccess section.
Begin of SingleRandomAccess section.
Node(s) with error 0
Node selected 475
Single GUP/s 0.004057
Current time (1182266783) is Tue Jun 19 11:26:23 2007

End of SingleRandomAccess section.
Begin of MPIFFT section.
Number of nodes: 512
Vector size: 536870912
Generation time: 0.178
Tuning: 0.163
Computing: 2.916
Inverse FFT: 2.921
 $\max(|x-x_0|)$: $1.717e-15$
Current time (1182266790) is Tue Jun 19 11:26:30 2007

End of MPIFFT section.
Begin of StarFFT section.
Vector size: 2097152
Generation time: 0.357
Tuning: 0.001
Computing: 0.483
Inverse FFT: 0.539
 $\max(|x-x_0|)$: $1.617e-15$
Node(s) with error 0
Minimum Gflop/s 0.412050
Average Gflop/s 0.450655
Maximum Gflop/s 0.466624
Current time (1182266792) is Tue Jun 19 11:26:32 2007

End of StarFFT section.
Begin of SingleFFT section.
Node(s) with error 0
Node selected 453
Single FFT Gflop/s 0.707194
Current time (1182266793) is Tue Jun 19 11:26:33 2007

End of SingleFFT section.
Begin of LatencyBandwidth section.

Latency-Bandwidth-Benchmark R1.5.1 (c) HLRS, University of Stuttgart
Written by Rolf Rabenseifner, Gerrit Schulz, and Michael Speck, Germany

Details - level 2

MPI_Wtime granularity.
Max. MPI_Wtick is 0.000001 sec
wtick is set to 0.000001 sec

Message Length: 8
Latency min / avg / max: 0.005260 / 0.005260 / 0.005260 msecs
Bandwidth min / avg / max: 1.521 / 1.521 / 1.521 MByte/s

MPI_Wtime granularity is ok.
message size: 8
max time : 10.000000 secs
latency for msg: 0.005260 msecs
estimation for ping pong: 0.473410 msecs
max number of ping pong pairs = 21123
max client pings = max server pongs = 145
stride for latency = 5
Message Length: 8
Latency min / avg / max: 0.000998 / 0.005111 / 0.005379 msecs
Bandwidth min / avg / max: 1.487 / 1.600 / 8.013 MByte/s

Message Length: 2000000
Latency min / avg / max: 10.195971 / 10.195971 / 10.195971 msecs
Bandwidth min / avg / max: 196.156 / 196.156 / 196.156 MByte/s

MPI_Wtime granularity is ok.
message size: 2000000
max time : 30.000000 secs
latency for msg: 10.195971 msecs
estimation for ping pong: 81.567764 msecs
max number of ping pong pairs = 367
max client pings = max server pongs = 19
stride for latency = 29
Message Length: 2000000
Latency min / avg / max: 10.139465 / 10.179064 / 10.231495 msecs
Bandwidth min / avg / max: 195.475 / 196.482 / 197.249 MByte/s

Message Size: 8 Byte
Natural Order Latency: 0.005603 msec
Natural Order Bandwidth: 1.427848 MB/s
Avg Random Order Latency: 0.017202 msec
Avg Random Order Bandwidth: 0.465049 MB/s

Message Size: 2000000 Byte
Natural Order Latency: 17.032743 msec
Natural Order Bandwidth: 117.420903 MB/s
Avg Random Order Latency: 66.029662 msec
Avg Random Order Bandwidth: 30.289417 MB/s

Execution time (wall clock) = 56.667 sec on 512 processes
- for cross ping_pong latency = 5.185 sec
- for cross ping_pong bandwidth = 25.942 sec
- for ring latency = 0.276 sec
- for ring bandwidth = 25.264 sec

Latency-Bandwidth-Benchmark R1.5.1 (c) HLRS, University of Stuttgart
Written by Rolf Rabenseifner, Gerrit Schulz, and Michael Speck, Germany

Major Benchmark results:

Max Ping Pong Latency: 0.005379 msecs
Randomly Ordered Ring Latency: 0.017202 msecs
Min Ping Pong Bandwidth: 195.474857 MB/s

Naturally Ordered Ring Bandwidth: 117.420903 MB/s
Randomly Ordered Ring Bandwidth: 30.289417 MB/s

Detailed benchmark results:

Ping Pong:

Latency min / avg / max: 0.000998 / 0.005111 / 0.005379 msec
Bandwidth min / avg / max: 195.475 / 196.482 / 197.249 MByte/s

Ring:

On naturally ordered ring: latency= 0.005603 msec, bandwidth= 117.420903 MB/s
On randomly ordered ring: latency= 0.017202 msec, bandwidth= 30.289417 MB/s

Benchmark conditions:

The latency measurements were done with 8 bytes
The bandwidth measurements were done with 2000000 bytes
The ring communication was done in both directions on 512 processes
The Ping Pong measurements were done on
- 10609 pairs of processes for latency benchmarking, and
- 324 pairs of processes for bandwidth benchmarking,
out of $512 \times (512 - 1) = 261632$ possible combinations on 512 processes.
(1 MB/s = 10^6 byte/sec)

Current time (1182266850) is Tue Jun 19 11:27:30 2007

End of LatencyBandwidth section.

Begin of Summary section.

VersionMajor=1
VersionMinor=0
VersionMicro=0
VersionRelease=b
LANG=C
Success=1
sizeof_char=1
sizeof_short=2
sizeof_int=4
sizeof_long=8
sizeof_void_ptr=8
sizeof_size_t=8
sizeof_float=4
sizeof_double=8
sizeof_s64Int=8
sizeof_u64Int=8
CommWorldProcs=512
MPI_Wtick=9.536743e-07
HPL_Tflops=1.86107
HPL_time=122.873
HPL_eps=1.11022e-16
HPL_RnormI=1.01772e-09
HPL_Anorm1=17673.5
HPL_AnormI=17648.8
HPL_Xnorm1=51674.1
HPL_XnormI=4.44249
HPL_N=70000
HPL_NB=200
HPL_nprow=16
HPL_npcol=32
HPL_depth=0
HPL_nbddiv=2
HPL_nbmin=4
HPL_cpfact=L
HPL_crfact=R
HPL_ctop=3
HPL_order=R

HPL_dMACH_EPS=1.110223e-16
HPL_dMACH_SFMIN=2.225074e-308
HPL_dMACH_BASE=2.000000e+00
HPL_dMACH_PREC=2.220446e-16
HPL_dMACH_MLEN=5.300000e+01
HPL_dMACH_RND=1.000000e+00
HPL_dMACH_EMIN=-1.021000e+03
HPL_dMACH_RMIN=2.225074e-308
HPL_dMACH_EMAX=1.024000e+03
HPL_dMACH_RMAX=1.797693e+308
HPL_sMACH_EPS=5.960464e-08
HPL_sMACH_SFMIN=1.175494e-38
HPL_sMACH_BASE=2.000000e+00
HPL_sMACH_PREC=1.192093e-07
HPL_sMACH_MLEN=2.400000e+01
HPL_sMACH_RND=1.000000e+00
HPL_sMACH_EMIN=-1.250000e+02
HPL_sMACH_RMIN=1.175494e-38
HPL_sMACH_EMAX=1.280000e+02
HPL_sMACH_RMAX=3.402823e+38
dweeps=1.110223e-16
sweps=5.960464e-08
HPLMaxProcs=512
HPLMinProcs=512
DGEMM_N=1546
StarDGEMM_Gflops=7.28761
SingleDGEMM_Gflops=8.06763
PTRANS_GBs=13.7264
PTRANS_time=0.713952
PTRANS_residual=0
PTRANS_n=35000
PTRANS_nb=100
PTRANS_nproW=16
PTRANS_npcol=32
MPIRandomAccess_N=4294967296
MPIRandomAccess_time=57.5491
MPIRandomAccess_CheckTime=23.6016
MPIRandomAccess_Errors=0
MPIRandomAccess_ErrorsFraction=0
MPIRandomAccess_ExeUpdates=9200124928
MPIRandomAccess_GUPs=0.159866
MPIRandomAccess_TimeBound=60
RandomAccess_N=8388608
StarRandomAccess_GUPs=0.00283563
SingleRandomAccess_GUPs=0.00405748
STREAM_VectorSize=3190104
STREAM_Threads=1
StarSTREAM_Copy=1.02971
StarSTREAM_Scale=1.03061
StarSTREAM_Add=1.16284
StarSTREAM_Triad=1.1688
SingleSTREAM_Copy=3.02093
SingleSTREAM_Scale=2.9461
SingleSTREAM_Add=3.31427
SingleSTREAM_Triad=3.31859
FFT_N=2097152
StarFFT_Gflops=0.450655
SingleFFT_Gflops=0.707194
MPIFFT_N=536870912
MPIFFT_Gflops=26.6959
MPIFFT_maxErr=1.71659e-15
MaxPingPongLatency_usec=5.37932
RandomlyOrderedRingLatency_usec=17.2025
MinPingPongBandwidth_GBytes=0.195475
NaturallyOrderedRingBandwidth_GBytes=0.117421
RandomlyOrderedRingBandwidth_GBytes=0.0302894
MinPingPongLatency_usec=0.998378

```

AvgPingPongLatency_usec=5.1106
MaxPingPongBandwidth_GBytes=0.197249
AvgPingPongBandwidth_GBytes=0.196482
NaturallyOrderedRingLatency_usec=5.60284
FFTEblk=16
FFTEnp=8
FFTEl2size=1048576
M_OPENMP=-1
omp_get_num_threads=0
omp_get_max_threads=0
omp_get_num_procs=0
MemProc=-1
MemSpec=-1
MemVal=-1
MPIFFT_time0=4.05312e-06
MPIFFT_time1=0.747922
MPIFFT_time2=0.259638
MPIFFT_time3=0.683326
MPIFFT_time4=0.474093
MPIFFT_time5=0.725643
MPIFFT_time6=3.09944e-06
End of Summary section.
#####
End of HPC Challenge tests.
Current time (1182266850) is Tue Jun 19 11:27:30 2007

```

6.b.2. HPC benchmark data for N=10000

```

#####
This is the DARPA/DOE HPC Challenge Benchmark version 1.0.0 October 2003
Produced by Jack Dongarra and Piotr Luszczek
Innovative Computing Laboratory
University of Tennessee Knoxville and Oak Ridge National Laboratory

See the source files for authors of specific codes.
Compiled on Jun 18 2007 at 11:44:38
Current time (1182266961) is Tue Jun 19 11:29:21 2007

```

```

Hostname: 'slc1b1.dim'
#####
=====
HPLinpack 1.0a -- High-Performance Linpack benchmark -- January 20, 2004
Written by A. Petitet and R. Clint Whaley, Innovative Computing Labs., UTK
=====

```

An explanation of the input/output parameters follows:

```

T/V      : Wall time / encoded variant.
N        : The order of the coefficient matrix A.
NB       : The partitioning blocking factor.
P        : The number of process rows.
Q        : The number of process columns.
Time     : Time in seconds to solve the linear system.
Gflops   : Rate of execution for solving the linear system.

```

The following parameter values will be used:

```

N       : 100000
NB      : 200
PMAP    : Row-major process mapping
P       : 32
Q       : 64
PFACT   : Left
NBMIN   : 4
NDIV    : 2
RFACT   : Right
BCAST   : 2ringM

```

DEPTH : 0
 SWAP : Mix (threshold = 64)
 L1 : transposed form
 U : transposed form
 EQUIL : yes
 ALIGN : 16 double precision words

 - The matrix A is randomly generated for each test.
 - The following scaled residual checks will be computed:
 1) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * N)$
 2) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$
 3) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$
 - The relative machine precision (eps) is taken to be 1.110223e-16
 - Computational tests pass if scaled residuals are less than 16.0

Begin of PTRANS section.

M: 50000
 N: 50000
 MB: 200 100 200
 NB: 200 100 200
 P: 32
 Q: 64

TIME	M	N	MB	NB	P	Q	TIME	CHECK	GB/s	RESID
WALL	50000	50000	200	200	32	64	0.53	PASSED	37.462	0.00
CPU	50000	50000	200	200	32	64	0.54	PASSED	37.037	0.00
WALL	50000	50000	100	100	32	64	0.54	PASSED	36.716	0.00
CPU	50000	50000	100	100	32	64	0.55	PASSED	36.364	0.00
WALL	50000	50000	200	200	32	64	0.53	PASSED	37.493	0.00
CPU	50000	50000	200	200	32	64	0.55	PASSED	36.364	0.00

Finished 3 tests, with the following results:
 3 tests completed and passed residual checks.
 0 tests completed and failed residual checks.
 0 tests skipped because of illegal input values.

END OF TESTS.
 Current time (1182266975) is Tue Jun 19 11:29:35 2007

End of PTRANS section.

Begin of HPL section.

=====

HPLinpack 1.0a -- High-Performance Linpack benchmark -- January 20, 2004
 Written by A. Petitet and R. Clint Whaley, Innovative Computing Labs., UTK
 =====

An explanation of the input/output parameters follows:

T/V : Wall time / encoded variant.
 N : The order of the coefficient matrix A.
 NB : The partitioning blocking factor.
 P : The number of process rows.
 Q : The number of process columns.
 Time : Time in seconds to solve the linear system.
 Gflops : Rate of execution for solving the linear system.

The following parameter values will be used:

N : 100000
 NB : 200
 PMAP : Row-major process mapping
 P : 32
 Q : 64
 PFACT : Left
 NBMIN : 4
 NDIV : 2

```

RFACT : Right
BCAST : 2ringM
DEPTH : 0
SWAP : Mix (threshold = 64)
L1 : transposed form
U : transposed form
EQUIL : yes
ALIGN : 16 double precision words

```

```

-----
- The matrix A is randomly generated for each test.
- The following scaled residual checks will be computed:
  1)  $\frac{\|Ax-b\|_{\infty}}{\epsilon * \|A\|_1 * N}$ 
  2)  $\frac{\|Ax-b\|_{\infty}}{\epsilon * \|A\|_1 * \|x\|_1}$ 
  3)  $\frac{\|Ax-b\|_{\infty}}{\epsilon * \|A\|_{\infty} * \|x\|_{\infty}}$ 
- The relative machine precision (eps) is taken to be 1- Computational tests
pass if scaled residuals are less than 16.0

```

```

=====
T/V                N      NB      P      Q                Time                Gflops
-----
WR03R2L4          100000    200     32     64                129.38                5.153e+03
=====
 $\frac{\|Ax-b\|_{\infty}}{\epsilon * \|A\|_1 * N} = 0.0162180$  ..... PASSED
 $\frac{\|Ax-b\|_{\infty}}{\epsilon * \|A\|_1 * \|x\|_1} = 0.0109091$  ..... PASSED
 $\frac{\|Ax-b\|_{\infty}}{\epsilon * \|A\|_{\infty} * \|x\|_{\infty}} = 0.0019615$  ..... PASSED
=====

```

```

Finished      1 tests with the following results:
              1 tests completed and passed residual checks,
              0 tests completed and failed residual checks,
              0 tests skipped because of illegal input values.

```

```

-----
End of Tests.

```

```

=====
Current time (1182267105) is Tue Jun 19 11:31:45 2007

```

```

End of HPL section.
Begin of StarDGEMM section.
Scaled residual: 0.00593874
Node(s) with error 0
Minimum Gflop/s 4.424290
Average Gflop/s 7.079863
Maximum Gflop/s 7.642500
Current time (1182267106) is Tue Jun 19 11:31:46 2007

```

```

End of StarDGEMM section.
Begin of SingleDGEMM section.
Node(s) with error 0
Node selected 1587
Single DGEMM Gflop/s 7.912201
Current time (1182267107) is Tue Jun 19 11:31:47 2007

```

```

End of SingleDGEMM section.
Begin of StarSTREAM section.

```

```

-----
This system uses 8 bytes per DOUBLE PRECISION word.

```

```

-----
Array size = 1627604, Offset = 0
Total memory required = 0.0364 GB.
Each test is run 10 times, but only
the *best* time for each is used.

```

```

-----
Your clock granularity/precision appears to be 1 microseconds.
Each test below will take on the order of 17235 microseconds.

```

(= 17235 clock ticks)

Increase the size of the arrays if this shows that you are not getting at least 20 clock ticks per test.

WARNING -- The above is only a rough guideline.
For best results, please be sure you know the precision of your system timer.

Function	Rate (GB/s)	Avg time	Min time	Max time
Copy:	1.0152	0.0259	0.0257	0.0261
.110223e-16				
Scale:	1.0667	0.0245	0.0244	0.0247
Add:	1.1655	0.0338	0.0335	0.0341
Triad:	1.2319	0.0338	0.0317	0.0343

Results Comparison:
Expected : 1877116964765625088.000000 375423392953124992.000000
500564523937500032.000000
Observed : 1877116964762494976.000000 375423392953518720.000000
500564523937500032.000000
Solution Validates

Node(s) with error 0
Minimum Copy GB/s 0.953661
Average Copy GB/s 1.020977
Maximum Copy GB/s 1.196598
Minimum Scale GB/s 0.987395
Average Scale GB/s 1.031436
Maximum Scale GB/s 1.385633
Minimum Add GB/s 1.054997
Average Add GB/s 1.143028
Maximum Add GB/s 1.461643
Minimum Triad GB/s 1.074565
Average Triad GB/s 1.153736
Maximum Triad GB/s 2.150413
Current time (1182267109) is Tue Jun 19 11:31:49 2007

End of StarSTREAM section.
Begin of SingleSTREAM section.
Node(s) with error 0
Node selected 2028
Single STREAM Copy GB/s 2.923391
Single STREAM Scale GB/s 3.046514
Single STREAM Add GB/s 3.180433
Single STREAM Triad GB/s 3.170646
Current time (1182267109) is Tue Jun 19 11:31:49 2007

End of SingleSTREAM section.
Begin of MPIRandomAccess section.
Running on 2048 processors (PowerofTwo)
Total Main table size = 2^{33} = 8589934592 words
PE Main table size = 2^{22} = 4194304 words/PE
Default number of updates (RECOMMENDED) = 34359738368
Number of updates EXECUTED = 14134075392 (for a TIME BOUND of 60.00 secs)
CPU time used = 58.090000 seconds
Real time used = 58.085687 seconds
0.243331467 Billion(10^9) Updates per second [GUP/s]
0.000118814 Billion(10^9) Updates/PE per second [GUP/s]
Verification: CPU time used = 32.160000 seconds
Verification: Real time used = 33.033619 seconds
Found 0 errors in 8589934592 locations (passed).
Current time (1182267202) is Tue Jun 19 11:33:22 2007

End of MPIRandomAccess section.
Begin of StarRandomAccess section.
Main table size = 2^{22} = 4194304 words
Number of updates = 16777216

CPU time used = 5.280000 seconds
Real time used = 5.280341 seconds
0.003177298 Billion(10^9) Updates per second [GUP/s]
Found 0 errors in 4194304 locations (passed).
Node(s) with error 0
Minimum GUP/s 0.003133
Average GUP/s 0.003171
Maximum GUP/s 0.003232
Current time (1182267213) is Tue Jun 19 11:33:33 2007

End of StarRandomAccess section.
Begin of SingleRandomAccess section.
Node(s) with error 0
Node selected 1563
Single GUP/s 0.004696
Current time (1182267220) is Tue Jun 19 11:33:40 2007

End of SingleRandomAccess section.
Begin of MPIFFT section.
Number of nodes: 2048
Vector size: 1073741824
Generation time: 0.089
Tuning: 0.081
Computing: 2.119
Inverse FFT: 2.076
max(|x-x0|): 1.755e-15
Current time (1182267224) is Tue Jun 19 11:33:44 2007

End of MPIFFT section.
Begin of StarFFT section.
Vector size: 1048576
Generation time: 0.178
Tuning: 0.001
Computing: 0.218
Inverse FFT: 0.238
max(|x-x0|): 1.590e-15
Node(s) with error 0
Minimum Gflop/s 0.410551
Average Gflop/s 0.458268
Maximum Gflop/s 0.490154
Current time (1182267225) is Tue Jun 19 11:33:45 2007

End of StarFFT section.
Begin of SingleFFT section.
Node(s) with error 0
Node selected 1143
Single FFT Gflop/s 0.722309
Current time (1182267226) is Tue Jun 19 11:33:46 2007

End of SingleFFT section.
Begin of LatencyBandwidth section.

Latency-Bandwidth-Benchmark R1.5.1 (c) HLRS, University of Stuttgart
Written by Rolf Rabenseifner, Gerrit Schulz, and Michael Speck, Germany

Details - level 2

MPI_Wtime granularity.
Max. MPI_Wtick is 0.000001 sec
wtick is set to 0.000001 sec

Message Length: 8
Latency min / avg / max: 0.005439 / 0.005439 / 0.005439 msecs
Bandwidth min / avg / max: 1.471 / 1.471 / 1.471 MByte/s

```

MPI_Wtime granularity is ok.
message size:                               8
max time :                                  10.000000 secs
latency for msg:                             0.005439 msecs
estimation for ping pong:                    0.489503 msecs
max number of ping pong pairs =              20428
max client pings = max server pongs =        142
stride for latency =                          17
Message Length: 8
Latency  min / avg / max:  0.004306 /  0.005327 /  0.005990 msecs
Bandwidth min / avg / max:  1.335 /    1.508 /    1.858 MByte/s

```

```

Message Length: 2000000
Latency  min / avg / max: 10.218978 / 10.218978 / 10.218978 msecs
Bandwidth min / avg / max: 195.714 / 195.714 / 195.714 MByte/s

```

```

MPI_Wtime granularity is ok.
message size:                               2000000
max time :                                  30.000000 secs
latency for msg:                             10.218978 msecs
estimation for ping pong:                    81.751823 msecs
max number of ping pong pairs =              366
max client pings = max server pongs =        19
stride for latency =                          109
Message Length: 2000000
Latency  min / avg / max: 10.151982 / 10.201659 / 10.652542 msecs
Bandwidth min / avg / max: 187.749 / 196.053 / 197.006 MByte/s

```

```

Message Size:                               8 Byte
Natural Order Latency:                       0.006413 msec
Natural Order Bandwidth:                     1.247377 MB/s
Avg Random Order Latency:                    0.018138 msec
Avg Random Order Bandwidth:                  0.441071 MB/s

```

```

Message Size:                               2000000 Byte
Natural Order Latency:                       18.761754 msec
Natural Order Bandwidth:                     106.599841 MB/s
Avg Random Order Latency:                    90.116342 msec
Avg Random Order Bandwidth:                  22.193533 MB/s

```

```

Execution time (wall clock) = 71.540 sec on 2048 processes
- for cross ping_pong latency = 7.573 sec
- for cross ping_pong bandwidth = 29.036 sec
- for ring latency = 1.127 sec
- for ring bandwidth = 33.804 sec

```

```

-----
Latency-Bandwidth-Benchmark R1.5.1 (c) HLRS, University of Stuttgart
Written by Rolf Rabenseifner, Gerrit Schulz, and Michael Speck, Germany
Major Benchmark results:
-----

```

```

Max Ping Pong Latency:                       0.005990 msecs
Randomly Ordered Ring Latency:               0.018138 msecs
Min Ping Pong Bandwidth:                     187.748612 MB/s
Naturally Ordered Ring Bandwidth:           106.599841 MB/s
Randomly Ordered Ring Bandwidth:            22.193533 MB/s

```

```

-----
Detailed benchmark results:
Ping Pong:
Latency  min / avg / max:  0.004306 /  0.005327 /  0.005990 msecs
Bandwidth min / avg / max: 187.749 / 196.053 / 197.006 MByte/s
Ring:
On naturally ordered ring: latency= 0.006413 msec, bandwidth= 106.599841 MB/s
On randomly ordered ring: latency= 0.018138 msec, bandwidth= 22.193533 MB/s

```

Benchmark conditions:

The latency measurements were done with 8 bytes
The bandwidth measurements were done with 2000000 bytes
The ring communication was done in both directions on 2048 processes
The Ping Pong measurements were done on
- 14641 pairs of processes for latency benchmarking, and
- 361 pairs of processes for bandwidth benchmarking,
out of $2048 \times (2048 - 1) = 4192256$ possible combinations on 2048 processes.
(1 MB/s = 10^6 byte/sec)

Current time (1182267297) is Tue Jun 19 11:34:57 2007

End of LatencyBandwidth section.

Begin of Summary section.

VersionMajor=1
VersionMinor=0
VersionMicro=0
VersionRelease=b
LANG=C
Success=1
sizeof_char=1
sizeof_short=2
sizeof_int=4
sizeof_long=8
sizeof_void_ptr=8
sizeof_size_t=8
sizeof_float=4
sizeof_double=8
sizeof_s64Int=8
sizeof_u64Int=8
CommWorldProcs=2048
MPI_Wtick=9.536743e-07
HPL_Tflops=5.153
HPL_time=129.377
HPL_eps=1.11022e-16
HPL_RnormI=4.53557e-09
HPL_Anorm1=25189.8
HPL_AnormI=25157.5
HPL_Xnorm1=148665
HPL_XnormI=8.2789
HPL_N=100000
HPL_NB=200
HPL_nprow=32
HPL_npcol=64
HPL_depth=0
HPL_nbdiv=2
HPL_nbmin=4
HPL_cpfact=L
HPL_crfact=R
HPL_ctop=3
HPL_order=R
HPL_dMACH_EPS=1.110223e-16
HPL_dMACH_SFMIN=2.225074e-308
HPL_dMACH_BASE=2.000000e+00
HPL_dMACH_PREC=2.220446e-16
HPL_dMACH_MLEN=5.300000e+01
HPL_dMACH_RND=1.000000e+00
HPL_dMACH_EMIN=-1.021000e+03
HPL_dMACH_RMIN=2.225074e-308
HPL_dMACH_EMAX=1.024000e+03
HPL_dMACH_RMAX=1.797693e+308
HPL_sMACH_EPS=5.960464e-08
HPL_sMACH_SFMIN=1.175494e-38

HPL_sMACH_BASE=2.000000e+00
HPL_sMACH_PREC=1.192093e-07
HPL_sMACH_MLEN=2.400000e+01
HPL_sMACH_RND=1.000000e+00
HPL_sMACH_EMIN=-1.250000e+02
HPL_sMACH_RMIN=1.175494e-38
HPL_sMACH_EMAX=1.280000e+02
HPL_sMACH_RMAX=3.402823e+38
dweeps=1.110223e-16
sweps=5.960464e-08
HPLMaxProcs=2048
HPLMinProcs=2048
DGEMM_N=1104
StarDGEMM_Gflops=7.07986
SingleDGEMM_Gflops=7.9122
PTRANS_GBs=37.4932
PTRANS_time=0.53343
PTRANS_residual=0
PTRANS_n=50000
PTRANS_nb=200
PTRANS_nprow=32
PTRANS_npcol=64
MPIRandomAccess_N=8589934592
MPIRandomAccess_time=58.0857
MPIRandomAccess_CheckTime=33.0336
MPIRandomAccess_Errors=0
MPIRandomAccess_ErrorsFraction=0
MPIRandomAccess_ExeUpdates=14134075392
MPIRandomAccess_GUPs=0.243331
MPIRandomAccess_TimeBound=60
RandomAccess_N=4194304
StarRandomAccess_GUPs=0.00317136
SingleRandomAccess_GUPs=0.00469598
STREAM_VectorSize=1627604
STREAM_Threads=1
StarSTREAM_Copy=1.02098
StarSTREAM_Scale=1.03144
StarSTREAM_Add=1.14303
StarSTREAM_Triad=1.15374
SingleSTREAM_Copy=2.92339
SingleSTREAM_Scale=3.04651
SingleSTREAM_Add=3.18043
SingleSTREAM_Triad=3.17065
FFT_N=1048576
StarFFT_Gflops=0.458268
SingleFFT_Gflops=0.722309
MPIFFT_N=1073741824
MPIFFT_Gflops=75.9955
MPIFFT_maxErr=1.75542e-15
MaxPingPongLatency_usec=5.99027
RandomlyOrderedRingLatency_usec=18.1377
MinPingPongBandwidth_GBytes=0.187749
NaturallyOrderedRingBandwidth_GBytes=0.1066
RandomlyOrderedRingBandwidth_GBytes=0.0221935
MinPingPongLatency_usec=4.30644
AvgPingPongLatency_usec=5.32659
MaxPingPongBandwidth_GBytes=0.197006
AvgPingPongBandwidth_GBytes=0.196053
NaturallyOrderedRingLatency_usec=6.41346
FFTEblk=16
FFTEnp=8
FFTEl2size=1048576
M_OPENMP=-1
omp_get_num_threads=0
omp_get_max_threads=0
omp_get_num_procs=0
MemProc=-1

```

MemSpec=-1
MemVal=-1
MPIFFT_time0=2.86102e-06
MPIFFT_time1=0.524337
MPIFFT_time2=0.1812
MPIFFT_time3=0.556872
MPIFFT_time4=0.237667
MPIFFT_time5=0.608035
MPIFFT_time6=2.86102e-06
End of Summary section.
#####
End of HPC Challenge tests.
Current time (1182267297) is Tue Jun 19 11:34:57 2007
#####

```

6.b.3. HPC benchmark data for N=600000

```

#####
This is the DARPA/DOE HPC Challenge Benchmark version 1.0.0 October 2003
Produced by Jack Dongarra and Piotr Luszczek
Innovative Computing Laboratory
University of Tennessee Knoxville and Oak Ridge National Laboratory

```

```

See the source files for authors of specific codes.
Compiled on May 10 2007 at 14:49:43
Current time (1180068203) is Fri May 25 00:43:23 2007

```

Hostname: 's11c1b2'

```

#####
=====
HPLinpack 1.0a -- High-Performance Linpack benchmark -- January 20, 2004
Written by A. Petitet and R. Clint Whaley, Innovative Computing Labs., UTK
=====

```

An explanation of the input/output parameters follows:

```

T/V      : Wall time / encoded variant.
N        : The order of the coefficient matrix A.
NB       : The partitioning blocking factor.
P        : The number of process rows.
Q        : The number of process columns.
Time     : Time in seconds to solve the linear system.
Gflops   : Rate of execution for solving the linear system.

```

The following parameter values will be used:

```

N       : 600000
NB      : 200
PMAP    : Row-major process mapping
P       : 30
Q       : 68
PFACT   : Left
NBMIN   : 4
NDIV    : 2
RFACT   : Right
BCAST   : 2ringM
DEPTH   : 0
SWAP    : Mix (threshold = 64)
L1      : transposed form
U       : transposed form
EQUIL   : yes
ALIGN   : 16 double precision words

```

```

-----
- The matrix A is randomly generated for each test.
- The following scaled residual checks will be computed:
  1) ||Ax-b||_oo / ( eps * ||A||_1 * N )

```

- 2) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$
- 3) $\frac{\|Ax-b\|_{\infty}}{\|A\|_{\infty}} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$

- The relative machine precision (eps) is taken to be 1.110223e-16
 - Computational tests pass if scaled residuals are less than 16.0

Begin of PTRANS section.

M: 300000
 N: 300000
 MB: 200 100 200
 NB: 200 100 200
 P: 30
 Q: 68

	TIME	M	N	MB	NB	P	Q	TIME	CHECK	GB/s	RESID
WALL	300000	300000	300000	200	200	30	68	31.22	PASSED	23.062	0.00
CPU	300000	300000	300000	200	200	30	68	31.24	PASSED	23.047	0.00
WALL	300000	300000	300000	100	100	30	68	17.66	PASSED	40.765	0.00
CPU	300000	300000	300000	100	100	30	68	17.69	PASSED	40.701	0.00
WALL	300000	300000	300000	200	200	30	68	31.18	PASSED	23.095	0.00
CPU	300000	300000	300000	200	200	30	68	31.22	PASSED	23.062	0.00

Finished 3 tests, with the following results:
 3 tests completed and passed residual checks.
 0 tests completed and failed residual checks.
 0 tests skipped because of illegal input values.

END OF TESTS.
 Current time (1180068420) is Fri May 25 00:47:00 2007
 End of PTRANS section.

Begin of HPL section.

=====
 HPLinpack 1.0a -- High-Performance Linpack benchmark -- January 20, 2004
 Written by A. Petitet and R. Clint Whaley, Innovative Computing Labs., UTK
 =====

An explanation of the input/output parameters follows:
 T/V : Wall time / encoded variant.
 N : The order of the coefficient matrix A.
 NB : The partitioning blocking factor.
 P : The number of process rows.
 Q : The number of process columns.
 Time : Time in seconds to solve the linear system.
 Gflops : Rate of execution for solving the linear system.

The following parameter values will be used:

N : 600000
 NB : 200
 PMAP : Row-major process mapping
 P : 30
 Q : 68
 PFACT : Left
 NBMIN : 4
 NDIV : 2
 RFACT : Right
 BCAST : 2ringM
 DEPTH : 0
 SWAP : Mix (threshold = 64)
 L1 : transposed form
 U : transposed form
 EQUIL : yes
 ALIGN : 16 double precision words

- The matrix A is randomly generated for each test.
- The following scaled residual checks will be computed:

- 1) $\frac{\|Ax-b\|_{\infty}}{\epsilon * \|A\|_1 * N}$
- 2) $\frac{\|Ax-b\|_{\infty}}{\epsilon * \|A\|_1 * \|x\|_1}$
- 3) $\frac{\|Ax-b\|_{\infty}}{\epsilon * \|A\|_{\infty} * \|x\|_{\infty}}$

- The relative machine precision (eps) is taken to be 1.110223e-16
 - Computational tests pass if scaled residuals are less than 16.0

```
=====
```

T/V	N	NB	P	Q	Time	Gflops
WR03R2L4	600000	200	30	68	10641.34	1.353e+04

```
=====
```

$\frac{\ Ax-b\ _{\infty}}{\epsilon * \ A\ _1 * N}$ =	0.0099666	PASSED
$\frac{\ Ax-b\ _{\infty}}{\epsilon * \ A\ _1 * \ x\ _1}$ =	0.0066616	PASSED
$\frac{\ Ax-b\ _{\infty}}{\epsilon * \ A\ _{\infty} * \ x\ _{\infty}}$ =	0.0011291	PASSED

```
=====
```

Finished 1 tests with the following results:
 1 tests completed and passed residual checks,
 0 tests completed and failed residual checks,
 0 tests skipped because of illegal input values.

 End of Tests.

```
=====
```

Current time (1180079078) is Fri May 25 03:44:38 2007
 End of HPL section.

Begin of StarDGEMM section.

Scaled residual: 0.00242414
 Node(s) with error 0
 Minimum Gflop/s 6.792508
 Average Gflop/s 8.268927
 Maximum Gflop/s 8.319971
 Current time (1180079189) is Fri May 25 03:46:29 2007

End of StarDGEMM section.

Begin of SingleDGEMM section.

Node(s) with error 0
 Node selected 719
 Single DGEMM Gflop/s 8.494016
 Current time (1180079284) is Fri May 25 03:48:04 2007

End of SingleDGEMM section.

Begin of StarSTREAM section.

 This system uses 8 bytes per DOUBLE PRECISION word.

Array size = 58823529, Offset = 0
 Total memory required = 1.3148 GB.
 Each test is run 10 times, but only
 the *best* time for each is used.

Your clock granularity/precision appears to be 1 microseconds.
 Each test below will take on the order of 630958 microseconds.
 (= 630958 clock ticks)
 Increase the size of the arrays if this shows that

you are not getting at least 20 clock ticks per test.

WARNING -- The above is only a rough guideline.
For best results, please be sure you know the
precision of your system timer.

Function	Rate (GB/s)	Avg time	Min time	Max time
Copy:	1.0377	0.9083	0.9070	0.9100
Scale:	1.0497	0.9008	0.8966	0.9028
Add:	1.1782	1.2023	1.1982	1.2069
Triad:	1.1941	1.1958	1.1823	1.2006

Results Comparison:

Expected : 67841221951582027776.000000 13568244390316406784.000000
18090992520421875712.000000
Observed : 67841221942089768960.000000 13568244391914422272.000000
18090992517186793472.000000

Solution Validates

Node(s) with error 0
Minimum Copy GB/s 1.002325
Average Copy GB/s 1.039966
Maximum Copy GB/s 1.119272
Minimum Scale GB/s 1.039187
Average Scale GB/s 1.057497
Maximum Scale GB/s 1.113818
Minimum Add GB/s 1.156317
Average Add GB/s 1.176485
Maximum Add GB/s 1.195994
Minimum Triad GB/s 1.169383
Average Triad GB/s 1.210038
Maximum Triad GB/s 1.327515
Current time (1180079330) is Fri May 25 03:48:50 2007

End of StarSTREAM section.

Begin of SingleSTREAM section.

Node(s) with error 0
Node selected 1806
Single STREAM Copy GB/s 2.959842
Single STREAM Scale GB/s 3.050984
Single STREAM Add GB/s 3.318012
Single STREAM Triad GB/s 3.312121
Current time (1180079348) is Fri May 25 03:49:08 2007

End of SingleSTREAM section.

Begin of MPIRandomAccess section.

Running on 2040 processors
Total Main table size = 238 = 274877906944 words
PE Main table size = (238)/2040 = 134744073 words/PE MAX
Default number of updates (RECOMMENDED) = 1099511627776
Number of updates EXECUTED = 15331277280 (for a TIME BOUND of 2660.33 secs)
CPU time used = 61.370000 seconds
Real time used = 61.405342 seconds
0.249673348 Billion(109) Updates per second [GUP/s]

0.000122389 Billion(109) Updates/PE per second [GUP/s]
Verification: CPU time used = 32.340000 seconds
Verification: Real time used = 33.676471 seconds
Found 0 errors in 274877906944 locations (passed).
Current time (1180079488) is Fri May 25 03:51:28 2007

End of MPIRandomAccess section.

Begin of StarRandomAccess section.

Main table size = 227 = 134217728 words
Number of updates = 536870912
CPU time used = 211.780000 seconds
Real time used = 211.805222 seconds
0.002534739 Billion(109) Updates per second [GUP/s]
Found 0 errors in 134217728 locations (passed).
Node(s) with error 0
Minimum GUP/s 0.002520
Average GUP/s 0.002536
Maximum GUP/s 0.002540
Current time (1180079912) is Fri May 25 03:58:32 2007

End of StarRandomAccess section.

Begin of SingleRandomAccess section.

Node(s) with error 0
Node selected 1960
Single GUP/s 0.003517
Current time (1180080218) is Fri May 25 04:03:38 2007

End of SingleRandomAccess section.

Begin of MPIFFT section.

Number of nodes: 1024
Vector size: 17179869184
Generation time: 2.846
Tuning: 2.592
Computing: 43.377
Inverse FFT: 44.124
max(|x-x0|): 2.220e-15
Current time (1180080317) is Fri May 25 04:05:17 2007

End of MPIFFT section.

Begin of StarFFT section.

Vector size: 33554432
Generation time: 5.694
Tuning: 0.003
Computing: 10.507
Inverse FFT: 11.355
max(|x-x0|): 2.095e-15
Node(s) with error 0
Minimum Gflop/s 0.375156
Average Gflop/s 0.399146
Maximum Gflop/s 0.416085
Current time (1180080352) is Fri May 25 04:05:52 2007

End of StarFFT section.

Begin of SingleFFT section.

Node(s) with error 0
Node selected 1334
Single FFT Gflop/s 0.435536
Current time (1180080381) is Fri May 25 04:06:21 2007

End of SingleFFT section.

Begin of LatencyBandwidth section.

Latency-Bandwidth-Benchmark R1.5.1 (c) HLRS, University of Stuttgart
Written by Rolf Rabenseifner, Gerrit Schulz, and Michael Speck, Germany
Details - level 2

MPI_Wtime granularity.
Max. MPI_Wtick is 0.000001 sec
wtick is set to 0.000001 sec

Message Length: 8
Latency min / avg / max: 0.005186 / 0.005186 / 0.005186 msecs
Bandwidth min / avg / max: 1.543 / 1.543 / 1.543 MByte/s

MPI_Wtime granularity is ok.
message size: 8
max time : 10.000000 secs
latency for msg: 0.005186 msecs
estimation for ping pong: 0.466704 msecs
max number of ping pong pairs = 21426
max client pings = max server pongs = 146
stride for latency = 17

Message Length: 8
Latency min / avg / max: 0.000998 / 0.005393 / 0.006124 msecs
Bandwidth min / avg / max: 1.306 / 1.516 / 8.013 MByte/s

Message Length: 2000000
Latency min / avg / max: 10.182977 / 10.182977 / 10.182977 msecs
Bandwidth min / avg / max: 196.406 / 196.406 / 196.406 MByte/s

MPI_Wtime granularity is ok.
message size: 2000000
max time : 30.000000 secs
latency for msg: 10.182977 msecs
estimation for ping pong: 81.463814 msecs
max number of ping pong pairs = 368
max client pings = max server pongs = 19
stride for latency = 109

Message Length: 2000000
Latency min / avg / max: 10.131001 / 10.173062 / 10.206461 msecs
Bandwidth min / avg / max: 195.954 / 196.598 / 197.414 MByte/s

Message Size: 8 Byte
Natural Order Latency: 0.006104 msec
Natural Order Bandwidth: 1.310720 MB/s
Avg Random Order Latency: 0.017735 msec
Avg Random Order Bandwidth: 0.451096 MB/s

Message Size: 2000000 Byte
Natural Order Latency: 21.539450 msec
Natural Order Bandwidth: 92.852883 MB/s
Avg Random Order Latency: 94.330747 msec

Avg Random Order Bandwidth: 21.201995 MB/s

```
Execution time (wall clock)      = 72.957 sec on 2040 processes
- for cross ping_pong latency    = 7.494 sec
- for cross ping_pong bandwidth = 28.893 sec
- for ring latency               = 0.985 sec
- for ring bandwidth             = 35.585 sec
```

Latency-Bandwidth-Benchmark R1.5.1 (c) HLRS, University of Stuttgart
Written by Rolf Rabenseifner, Gerrit Schulz, and Michael Speck, Germany
Major Benchmark results:

Max Ping Pong Latency: 0.006124 msecs
Randomly Ordered Ring Latency: 0.017735 msecs
Min Ping Pong Bandwidth: 195.954309 MB/s
Naturally Ordered Ring Bandwidth: 92.852883 MB/s
Randomly Ordered Ring Bandwidth: 21.201995 MB/s

Detailed benchmark results:

Ping Pong:

Latency min / avg / max: 0.000998 / 0.005393 / 0.006124 msecs
Bandwidth min / avg / max: 195.954 / 196.598 / 197.414 MByte/s

Ring:

On naturally ordered ring: latency= 0.006104 msec, bandwidth= 92.852883 MB/s
On randomly ordered ring: latency= 0.017735 msec, bandwidth= 21.201995 MB/s

Benchmark conditions:

The latency measurements were done with 8 bytes
The bandwidth measurements were done with 2000000 bytes
The ring communication was done in both directions on 2040 processes
The Ping Pong measurements were done on
- 14400 pairs of processes for latency benchmarking, and
- 361 pairs of processes for bandwidth benchmarking,
out of $2040 \cdot (2040 - 1) = 4159560$ possible combinations on 2040 processes.
(1 MB/s = 10^6 byte/sec)

Current time (1180080454) is Fri May 25 04:07:34 2007

End of LatencyBandwidth section.

Begin of Summary section.

```
VersionMajor=1
VersionMinor=0
VersionMicro=0
VersionRelease=b
LANG=C
Success=1
sizeof_char=1
sizeof_short=2
sizeof_int=4
sizeof_long=8
sizeof_void_ptr=8
sizeof_size_t=8
sizeof_float=4
sizeof_double=8
sizeof_s64Int=8
sizeof_u64Int=8
CommWorldProcs=2040
MPI_Wtick=9.536743e-07
HPL_Tflops=13.5322
```

HPL_time=10641.3
HPL_eps=1.11022e-16
HPL_RnormI=9.99117e-08
HPL_AnormI=150490
HPL_AnormI=150357
HPL_XnormI=897682
HPL_XnormI=8.8349
HPL_N=600000
HPL_NB=200
HPL_nprow=30
HPL_npcol=68
HPL_depth=0
HPL_nbdiv=2
HPL_nbmin=4
HPL_cpfact=L
HPL_crfact=R
HPL_ctop=3
HPL_order=R
HPL_dMACH_EPS=1.110223e-16
HPL_dMACH_SFMIN=2.225074e-308
HPL_dMACH_BASE=2.000000e+00
HPL_dMACH_PREC=2.220446e-16
HPL_dMACH_MLEN=5.300000e+01
HPL_dMACH_RND=1.000000e+00
HPL_dMACH_EMIN=-1.021000e+03
HPL_dMACH_RMIN=2.225074e-308
HPL_dMACH_EMAX=1.024000e+03
HPL_dMACH_RMAX=1.797693e+308
HPL_sMACH_EPS=5.960464e-08
HPL_sMACH_SFMIN=1.175494e-38
HPL_sMACH_BASE=2.000000e+00
HPL_sMACH_PREC=1.192093e-07
HPL_sMACH_MLEN=2.400000e+01
HPL_sMACH_RND=1.000000e+00
HPL_sMACH_EMIN=-1.250000e+02
HPL_sMACH_RMIN=1.175494e-38
HPL_sMACH_EMAX=1.280000e+02
HPL_sMACH_RMAX=3.402823e+38
dweeps=1.110223e-16
sweps=5.960464e-08
HPLMaxProcs=2040
HPLMinProcs=2040
DGEMM_N=6642
StarDGEMM_Gflops=8.26893
SingleDGEMM_Gflops=8.49402
PTRANS_GBs=40.7648
PTRANS_time=17.6623
PTRANS_residual=0
PTRANS_n=300000
PTRANS_nb=100
PTRANS_nprow=30
PTRANS_npcol=68
MPIRandomAccess_N=274877906944
MPIRandomAccess_time=61.4053
MPIRandomAccess_CheckTime=33.6765
MPIRandomAccess_Errors=0
MPIRandomAccess_ErrorsFraction=0
MPIRandomAccess_ExeUpdates=15331277280
MPIRandomAccess_GUPs=0.249673
MPIRandomAccess_TimeBound=2660.33
RandomAccess_N=134217728
StarRandomAccess_GUPs=0.00253592
SingleRandomAccess_GUPs=0.00351655
STREAM_VectorSize=58823529
STREAM_Threads=1
StarSTREAM_Copy=1.03997
StarSTREAM_Scale=1.0575

```

StarSTREAM_Add=1.17648
StarSTREAM_Triad=1.21004
SingleSTREAM_Copy=2.95984
SingleSTREAM_Scale=3.05098
SingleSTREAM_Add=3.31801
SingleSTREAM_Triad=3.31212
FFT_N=33554432
StarFFT_Gflops=0.399146
SingleFFT_Gflops=0.435536
MPIFFT_N=17179869184
MPIFFT_Gflops=67.3302
MPIFFT_maxErr=2.22045e-15
MaxPingPongLatency_usec=6.12438
RandomlyOrderedRingLatency_usec=17.7346
MinPingPongBandwidth_GBytes=0.195954
NaturallyOrderedRingBandwidth_GBytes=0.0928529
RandomlyOrderedRingBandwidth_GBytes=0.021202
MinPingPongLatency_usec=0.998378
AvgPingPongLatency_usec=5.39269
MaxPingPongBandwidth_GBytes=0.197414
AvgPingPongBandwidth_GBytes=0.196598
NaturallyOrderedRingLatency_usec=6.10352
FFTEblk=16
FFTEnp=8
FFTEl2size=1048576
M_OPENMP=-1
omp_get_num_threads=0
omp_get_max_threads=0
omp_get_num_procs=0
MemProc=-1
MemSpec=-1
MemVal=-1
MPIFFT_time0=3.09944e-06
MPIFFT_time1=9.02557
MPIFFT_time2=3.8635
MPIFFT_time3=8.04679
MPIFFT_time4=12.7379
MPIFFT_time5=9.16573
MPIFFT_time6=4.05312e-06
End of Summary section.

#####
End of HPC Challenge tests.
Current time (1180080454) is Fri May 25 04:07:34 2007
#####

```

6.c. Comparison of application benchmarks

6.c.1. WRF

# Processors	# Nodes	# Cores	Time				Linear Scaling				Scaling Efficiency			
			NSF Small	NSF Medium	NSF Large	NSF X-Large	NSF Small	NSF Medium	NSF Large	NSF X-Large	NSF Small	NSF Medium	NSF Large	NSF X-Large
1	1	1												
1	1	2												
2	1	4												
4	2	8												
8	4	16		8379				1				1		
16	8	32		4372				4189.5				0.95826		
32	16	64		2345	33863			2094.75	1			0.89328	1	
64	32	128		1253	18102			1047.38	16931.5			0.83589	0.93534	
96	48	192			12114				11287.7				0.93179	
128	64	256		737	10006			523.688	8465.75			0.71057	0.84607	
160	80	320												
192	96	384		599				349.125				0.58285		
224	112	448												
256	128	512		546	5666			261.844	4232.88			0.47957	0.74707	
384	192	768												
448	224	896												
512	256	1024												
1024	512	2048												
1536	768	3072												
2048	1024	4096												

Table 10.

6.c.2. OOCORE1

Expanded data not available.

6.c.3. GAMESS1

standard.inp												
MPI ver.												
#-way	#-nodes	Total wall time (s)			CPU utilization	Real time				Startup		
			Scalability	Efficiency		(min +	sec)	(s)	(s)	%	Scalability	Efficiency
16	4	14508	1	100.00%	99.17%	242	17.192	14537.19	29.192	0.20%	1	100.00%
32	8	7917	1.832512	91.63%	97.54%	132	33.846	7953.846	36.846	0.46%	1.828	91.38%
48	12	5707	2.542141	84.74%	96.29%	96	15.588	5775.588	68.588	1.19%	2.517	83.90%
64	16	4626	3.136187	78.40%	94.57%	78	11.126	4691.126	65.126	1.39%	3.099	77.47%
96	24	3582	4.050251	67.50%	91.36%	61	46.772	3706.772	124.772	3.37%	3.922	65.36%
128	32	3090	4.695146	58.69%	87.81%	52	51.608	3171.608	81.608	2.57%	4.584	57.29%
192	48	2843	5.10306	42.53%	83.62%	49	43.659	2983.659	140.659	4.71%	4.872	40.60%
256	64	2837	5.113853	31.96%	76.66%	49	5.8	2945.8	108.8	3.69%	4.935	30.84%
socket ver.												
#-way	#-nodes	Total wall time (s)			CPU utilization	Real time				Startup		
			Scalability	Efficiency		(min +	sec)	(s)	(s)	%	Scalability	Efficiency
16	4	14398	1	100.00%	0.9674	240	8.068	14408.07	10.068	0.07%	1	100.00%
32	8	7967	1.807	90.36%	0.9322	133	3.157	7983.157	16.157	0.20%	1.805	90.24%
48	12	5820	2.474	82.46%	0.9108	97	10.373	5830.373	10.373	0.18%	2.471	82.37%
64	16	4774	3.016	75.40%	0.8944	80	0.201	4800.201	26.201	0.55%	3.002	75.04%
96	24	3762	3.827	63.79%	0.8441	63	19.452	3799.452	37.452	0.99%	3.792	63.20%
128	32	3321	4.335	54.19%	0.801	56	13.991	3373.991	52.991	1.57%	4.270	53.38%
192	48	3087	4.664	38.87%	0.7426	52	40.465	3160.465	73.465	2.32%	4.559	37.99%
256	64	3123	4.610	28.81%	0.6816	53	44.793	3224.793	101.793	3.16%	4.468	27.92%
large.inp												
MPI ver.												
#-way	#-nodes	Total wall time (s)			CPU utilization	Real time				Startup		
			Scalability	Efficiency		(min +	sec)	(s)	(s)	%	Scalability	Efficiency
128	32	23459	1	100.00%	92.77%	392	21.645	23541.65	82.645	0.35%	1	100.00%
192	48	15757	1.489	99.25%	91.39%	264	30.689	15870.69	113.689	0.72%	1.483	98.89%
256	64	12103	1.938	96.91%	90.34%	203	45.447	12225.45	122.447	1.00%	1.926	96.28%
320	80	9976	2.352	94.06%	88.38%	169	2.116	10142.12	166.116	1.64%	2.321	92.85%
384	96	8593	2.730	91.00%	88.00%	148	43.477	8923.477	330.477	3.70%	2.638	87.94%
448	112	7545	3.109	88.83%	86.09%	130	22.32	7822.32	277.32	3.55%	3.010	85.99%
512	128	6869	3.415	85.38%	84.85%	117	18.43	7038.43	169.43	2.41%	3.345	83.62%
768	192	5283	4.440	74.01%	77.67%	95	41.957	5741.957	458.957	7.99%	4.100	68.33%
896	224	4923	4.765	68.07%	73.29%	90	18.02	5418.02	495.02	9.14%	4.345	62.07%
1024	256	4641	5.055	63.18%	70.80%	89	13.444	5353.444	712.444	13.31%	4.397	54.97%
socket ver.												

#-way	#-nodes	Total wall time (s)			CPU utilization	Real time			Startup			
			Scalability	Efficiency		(min +	sec)	(s)	(s)	%	Scalability	Efficiency
128	32	23971	1	100.00%	0.9022	400	13.418	24013.42	42.418	0.18%	1	100.00%
192	48	16208	1.478961	98.60%	0.8536	271	6.213	16266.21	58.213	0.36%	1.476276	98.42%
256	64	12287	1.950924	97.55%	0.8503	205	59.374	12359.37	72.374	0.59%	1.942932	97.15%
320	80	10030	2.38993	95.60%	0.8239	167	27.305	10047.31	17.305	0.17%	2.390036	95.60%
384	96	8640	2.774421	92.48%	0.8083	146	1.655	8761.655	121.655	1.39%	2.74074	91.36%
448	112	7645	3.135513	89.59%	0.8046	129	27.663	7767.663	122.663	1.58%	3.09146	88.33%

Table 11. (Compiler flags: "-O3 -q64 -qintsize=8 -qspillsize=1500 -qtbtable=full -qarch=ppc970 -qtune=ppc970 -lessl")

6.c.4. MILC2

# Processors	# Nodes	# Cores	Time				Perfect Scaling				Scaling Efficiency			
			NSF Small	NSF Medium	NSF Large	NSF X-Large	NSF Small	NSF Medium	NSF Large	NSF X-Large	NSF Small	NSF Medium	NSF Large	NSF X-Large
1	1	1	92.581											
1	1	2												
2	1	4												
4	2	8												
8	4	16		1683.432				1				1		
16	8	32		865.036				841.716				0.97304158		
32	16	64		447.845	11316.27			420.858	1			0.93974031	1	
64	32	128		235.026	5781.716			210.429	5658.135			0.89534349	0.97862555	
96	48	192												
128	64	256		122.015	2964.596			105.2145	2829.0675			0.86230791	0.95428433	
160	80	320												
192	96	384												
224	112	448												
256	128	512			1537.605									
384	192	768												
448	224	896												
512	256	1024			817									
1024	512	2048			452	6554								1
1536	768	3072				3695								1.18249887
2048	1024	4096												

Table 12.

6.c.5. *PARATEC2*

# Processors	# Nodes	# Cores	Time				Linear Scaling				Scaling Efficiency			
			NSF Small	NSF Medium	NSF Large	NSF X-Large	NSF Small	NSF Medium	NSF Large	NSF X-Large	NSF Small	NSF Medium	NSF Large	NSF X-Large
1	1	1												
1	1	2												
2	1	4												
4	2	8												
8	4	16												
16	8	32		1323				1				1		
32	16	64		737	3794			661.5	1			0.89755767	1	
64	32	128		518	2415			330.75	1897			0.63851351	0.78550725	
96	48	192												
128	64	256			2068				948.5				0.45865571	
160	80	320												
192	96	384												
224	112	448												
256	128	512												
384	192	768												
448	224	896												
512	256	1024												
1024	512	2048												
1536	768	3072												
2048	1024	4096												

Table 13.

6.c.6. *HOMME3*

Expanded data not available.

7. Acknowledgements

Indiana University's Big Red supercomputer was funded in part by the Indiana METACyt Initiative and is supported in part by a Shared University Research grant from IBM Inc. to Indiana University. The Indiana METACyt Initiative is supported by a grant from the Lilly Endowment, Inc. to Indiana University. Big Red is connected to the NSF-funded TeraGrid. IU's participation in the TeraGrid is funded by National Science Foundation grant numbers 0338618, 0504075, and 0451237. The IU Data Capacitor was funded by grant number CNS-0521433. IU's work with LEAD has been supported by NSF grants ATM 0331480 and EIA-0202048. CICC was supported by P20HG003894-01. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF) or the National Institutes of Health (NIH). The authors wish to thank the people involved in building CTSS software on Big Red: Eric Blau, Charles Bacon, JP Navarro (University of Chicago), and Jason Brechin (National Center for Supercomputer Applications). The MyCluster component of CTSS was built by Evan Turner (Texas Advanced Computing Center) and the SRB client was built by George Kremenek (San Diego Supercomputer Center). Thanks to David Hart and Nancy Wilkins-Diehr (SDSC) and Suresh Marru (IU) for providing information and statistics related to the TeraGrid, LEAD, and Science Gateways generally. We would like to thank the staffs of the Research Technologies Division of UIITS, the Pervasive Technology Institute at IU, and ZIH at Technische Universität Dresden, for their contributions and effort that made the implementation of Big Red and the Data Capacitor, their integration within the TeraGrid, and the performance analysis possible. The authors would particularly like to thank Ray Sheppard and Malinda Lingwall for writing and editing. Gratitude goes particularly to IU President Michael A. McRobbie, whose leadership made Big Red's acquisition and implementation possible.

8. References

- [1] IBM, Inc. IBM BladeCenter JS21: The power of blade innovation. <http://www.redbooks.ibm.com/redbooks/pdfs/sg247273.pdf>
- [2] IBM, Inc. IBM BladeCenter. <http://www-03.ibm.com/systems/bladecenter/>
- [3] Accelerating Growth: Indiana's Strategic Economic Development Plan. 2006. http://www.in.gov/iedc/files/Strategic_Plan.pdf
- [4] University of Tennessee at Knoxville. HPC Challenge. <http://icl.cs.utk.edu/hpcc/>
- [5] National Science Foundation. Benchmarking Information Referenced in the NSF 05-625 "High Performance Computing System Acquisition: Towards a Petascale Computing Environment for Science and Engineering". 2005. <http://www.nsf.gov/pubs/2006/nsf0605/nsf0605.pdf>
- [6] The Weather Research & Forecasting Model (WRF). Home page. <http://www.wrf-model.org/>
- [7] Cable, SB and E D'Azevedo. Benchmarking OOCORE, an Out-of-Core Matrix Solver. 2001. <http://www.ornl.gov/~webworks/cppr/y2001/misc/123564.pdf>
- [8] GAMESS. Home Page. <http://www.msg.ameslab.gov/GAMESS/>
- [9] MIMD Lattice Computation (MILC) Collaboration. Home page. <http://www.physics.indiana.edu/~sg/milc.html>
- [10] National Energy Research Scientific Computing Center. PARAllel Total Energy Code (PARATEC). Home page. <http://www.nersc.gov/projects/paratec/>
- [11] University Corporation for Atmospheric Research. High-Order Methods Modeling Environment (HOMME). Home page. <http://www.image.ucar.edu/CMG/Software/HOMME/>