NAS PARALLEL BENCHMARK RESULTS 10-94

David H. Bailey, Eric Barszcz, Leonardo Dagum and Horst D. Simon[1]

NAS Technical Report NAS-94-001
October 1994

Abstract

The NAS Parallel Benchmarks have been developed at NASA Ames Research Center to study the performance of parallel supercomputers. The eight benchmark problems are specified in a "pencil and paper" fashion. In other words, the complete details of the problem to be solved are given in a technical document, and except for a few restrictions, benchmarkers are mostly free to select the language constructs and implementation techniques best suited for a particular system.

This paper presents performance results of various systems using the NAS Parallel Benchmarks. These results represent the best results that have been reported to us for the specific systems listed. Some changes and clarifications to the benchmark rules are also described.

[1] Bailey and Barszcz are employees of NASA Ames Research Center. Dagum and Simon are employees of Computer Science Corp., and their work was funded by the NASA Ames Research Center under contract NAS2-12961. Mailing address for all authors: NASA Ames Research Center, Mail Stop T27A-1, Moffett Field, CA 94035-1000. The latest NPB results are available electronically on the WWW beginning at URL address: `http://www.nas.nasa.gov/RNR/Parallel/NPB/NPBindex.html` or by sending an email request to `leo@nas.nasa.gov`

1 Introduction

The Numerical Aerodynamic Simulation (NAS) Program, located at NASA Ames Research Center, is dedicated to advancing the science of computational aerodynamics. One key goal of the NAS organization is to demonstrate by the year 2000 an operational computing system capable of simulating an entire aerospace vehicle system within a computing time of one to several hours. It is currently projected that the solution of this grand challenge problem will require a computer system that can perform scientific computations at a sustained rate approximately one thousand times faster than 1990 generation supercomputers. Most likely such a computer system will employ hundreds or even thousands of processors operating in parallel.

In order to objectively measure the performance of various highly parallel computer systems and to compare them with conventional supercomputers, we along with other scientists in our organization have devised the NAS Parallel Benchmarks (NPB). Note that the NPB are distinct from the High Speed Processor (HSP) benchmarks and procurements. The HSP benchmarks are used for evaluating production supercomputers for procurement, whereas the NPB are for studying massively parallel processor (MPP) systems not necessarily tied to a procurement.

The NPB are a set of eight benchmark problems, each of which focuses on some important aspect of highly parallel supercomputing for aerophysics applications. Some extension of Fortran or C is required for implementations, and reasonable limits are placed on the usage of assembly code and the like, but otherwise programmers are free to utilize language constructs that give the best performance possible on the particular system being studied. The choice of data structures, processor allocation and memory usage are generally left open to the discretion of the implementer.

The eight problems consist of five "kernels" and three "simulated computational fluid dynamics (CFD) applications". Each of these is defined fully in [2]. The five kernels are relatively compact problems, each emphasizing a particular type of numerical computation. Compared with the simulated CFD applications, they can be implemented fairly readily and provide insight as to the general levels of performance that can be expected on these specific types of numerical computations.

The simulated CFD applications, on the other hand, usually require more effort to implement, but they are more indicative of the types of actual data movement and computation required in state-of-the-art CFD application codes. For example, in an isolated kernel a certain data structure may be very efficient on a certain system, and yet this data structure would be inappropriate if incorporated into a larger application. By comparison, the simulated CFD applications require data structures and implementation techniques that are more typical of real CFD applications.

Space does not permit a complete description of these benchmark problems. A more detailed description of these benchmarks, together with the rules and restrictions associated with the benchmarks, may be found in [1]. The full specification of the benchmarks is given in [2].

Sample Fortran programs implementing the NPB on a single processor system are available as an aid to implementors. These programs, as well as the benchmark document itself, are available through the World Wide Web (WWW) at URL address:

`http://www.nas.nasa.gov/RNR/Parallel/NPB/NPBindex.html`

or through postal mail from the following address: NAS Systems Division, Mail Stop 258-6, NASA Ames Research Center, Moffett Field, CA 94035, attn: NAS Parallel Benchmark Codes or by sending an email request to: bm-codes@nas.nasa.gov. The sample codes are provided on Macintosh floppy disks and contain the Fortran source codes, "README" files, input data files, and reference output data files for correct implementations of the benchmark problems. These codes have been validated on a number of computer systems ranging from conventional workstations to supercomputers.

There are now two standard sizes for the NAS Parallel Benchmarks; these will be referred to as the Class A and Class B size problems. The nominal benchmark sizes for the Class A and Class B are listed in Tables 1a and 1b respectively. These tables also give the standard floating point operation (flop) counts for the two classes of problems. Note that in the case of MG the grid size is unchanged, but a greater flop count results from changes in the inner loop iterations. We insist that those wishing to compute performance rates in millions of floating point operations per second (Mflop/s) use these standard flop counts. The tables contain Mflop/s rates calculated in this manner for the (frozen) 1992 implementation on one processor of the Cray Y-MP for Class A and the current fastest implementation on one processor of the Cray C90 for Class B. Note, however, that in Tables 2 through 9, performance rates are *not* cited in Mflop/s; we present instead the actual run times (and, equivalently, the performance ratios). We suggest that these, and not Mflop/s, be examined when comparing different systems and implementations.

Benchmark Name	Abbrev-iation	Nominal Size	Operation Count ($\times 10^9$)	Mflop/s on Y-MP/1
Embarrassingly Parallel	EP	2^{28}	26.68	211
Multigrid	MG	256^3	3.905	176
Conjugate Gradient	CG	14,000	1.508	127
3-D FFT PDE	FT	$256^2 \times 128$	5.631	196
Integer Sort	IS	$2^{23} \times 2^{19}$	0.7812	68
LU Simulated CFD Application	LU	64^3	64.57	194
SP Simulated CFD Application	SP	64^3	102.0	216
BT Simulated CFD Application	BT	64^3	181.3	229

Table 1a: Standard Operation Counts and YMP/1 Mflop/s for Class A Size Problems

Benchmark Name	Abbrev-iation	Nominal Size	Operation Count ($\times 10^9$)	Mflop/s on C90
Embarrassingly Parallel	EP	2^{30}	1008.8	543
Multigrid	MG	256^3	18.81	498
Conjugate Gradient	CG	75,000	54.89	447
3-D FFT PDE	FT	512×256^2	71.37	560
Integer Sort	IS	$2^{25} \times 2^{21}$	3.150	244
LU Simulated CFD Application	LU	102^3	319.6	493
SP Simulated CFD Application	SP	102^3	447.1	627
BT Simulated CFD Application	BT	102^3	721.5	572

Table 1b: Standard Operation Counts and C90 Mflop/s for Class B Size Problems

In the following, each of the eight benchmarks will be briefly described, and then the best performance results we have received to date for each computer system will be given in Tables 2 through 9. These tables include run times and performance ratios. The performance ratios compare individual timings with the current best time on that benchmark achieved on one processor of either a Cray Y-MP (for Class A) or a Cray C90 (for Class B). The run times in each case are elapsed time of day figures, measured in accordance with the specifications given in [2].

With the exception of the Integer Sort benchmark, these standard flop counts were determined by using the hardware performance monitor on either the Cray Y-MP or the Cray C90, and we believe that they are close to the minimal counts required for these problems. In the case of the Integer Sort benchmark, which does not involve floating-point operations, we selected a value approximately equal to the number of integer operations required, in order to permit the computation of performance rates analogous to Mflop/s rates. We reserve the right to change these standard flop counts in the future if deemed necessary.

The NAS organization reserves the right to verify any NPB results that are submitted to us. We may, for example, attempt to run the submitter's code on another system of the same configuration as that used by the submitter. In those instances where we are unable to reproduce the submitter's supplied results (allowing a 5% tolerance) our policy is to alert the submitter of the discrepancy and allow him or her until the next release of this report to resolve the discrepancy. If the discrepancy is not resolved to our satisfaction, then our own observed results, and not the submitter's results, will be reported. This policy will apply to all results we receive and publish.

Whenever possible, we have tried to credit the actual individuals and organizations who have contributed the performance results cited in the tables. In these citations, NAS denotes the NAS Applied Research Branch at NASA Ames (including both NASA civil servants and Computer Science Corp. contractors); RIACS denotes the parallel systems division of the Research Institute for Advanced Computer Science, which is located at NASA Ames; BBN denotes Bolt, Beranek and Newman; BCS denotes Boeing Computer Services; CRI denotes Cray Research, Inc.; Fujitsu denotes Fujitsu America, Inc.; KSR denotes Kendall Square Research Corp.; IBM denotes International Business Machines, Inc.; Intel denotes the Supercomputer Systems Division of Intel Corp.; MasPar denotes MasPar Computer Corp.; Meiko denotes Meiko Scientific

Corp.; NEC denotes HNSX Supercomputers Inc.; and TMC denotes Thinking Machines, Inc. Where no individual citation is made for a specific model, the results are due to vendor staff.

This paper reports benchmark results on the following systems: TC2000 by Bolt, Beranek and Newman (BBN); YMP, EL, C90, and T3D by Cray Research Inc. (CRI); Paragon and iPSC/860 by Intel; SP-1, SP-2 (wide node) and RS6000-590 by International Business Machines (IBM); VPP500 by Fujitsu; KSR1 and KSR2 by Kendall Square Research; ADENART by Kyoto University and Matsushita Electric Industrial Co.; MP-1 and MP-2 by MasPar Computer Corp.; CS-1 and CS-2 by Meiko Scientific; nCUBE-2S by nCUBE; SX-3 by NEC; Power Challenge XL and Power Indigo 2 by Silicon Graphics Inc. (SGI); CM-2, CM-200, CM-5, and CM-5E by Thinking Machines Corp. (TMC); and clusters of distributed workstations including Sparcstation's by Sun; RS6000's by IBM; and 4D25's by SGI. Entries in the tables are ordered alphabetically by vendor, except for distributed workstation results which appear last.

Unfortunately, the limited space in this report does not permit discussion of the methods used in any of these implementations. However, references to technical papers describing these methods have been included whenever such papers are available. In particular, details of the implementation of these benchmarks on the TC2000, the CM2, the CM200, the SP-1 and the IBM Cluster may be found in [5, 6, 11, 13]. General discussion on architectural requirements for the benchmarks may be found in [8]. Readers are referred to these documents for full details.

This report includes a number of new and/or improved results on the Cray C90 and T3D, the Fujitsu VPP500, the Intel Paragon (with OSF1.2 and with SunMos), the IBM wide node SP-2, the Kendall Square KSR2, the nCUBE-2S, the NEC SX-3, and the SGI Power Challenge and Power Indigo systems. The "Parasoft IBM (token)" results were run on a cluster of nine IBM RS6000-320H workstations with 25 MHz clock rate, 16 MB memory and a token ring interconnect capable of 16 MBits/sec transfer rates. The performance improvements observed on some of these systems reflect improvements both in compilers and implementations. Efforts are currently underway to port the NAS Parallel Benchmarks on other systems, and we hope to have more results in the future.

2 Benchmark Changes

Because the benchmarks are specified in only a pencil and paper fashion, it is inevitable that loopholes develop whereby the benchmark rules are not violated but the benchmark intent is defeated. This section addresses changes to be made in the Embarrassingly Parallel (EP) and Conjugate Gradient (CG) benchmark specification in order to close some loopholes that have developed with these kernels.

Eventually we hope that parallel computing technology will advance to the point where we will be able to measure performance by providing source code, rather than pencil and paper, benchmark descriptions. However, the current lack of a common parallel language or architectural paradigm prohibits our movement in this direction.

2.1 Changes to EP

The intent of the EP benchmark is to provide an accuracy and performance check on the Fortran LOG and SQRT intrinsics and to act as an easy kernel which vendors can readily implement on prototype systems. There are two possible loopholes in its implementation which are here disallowed. Results employing these loopholes will not be reported in future releases of this report.

The first loophole involves using a table lookup scheme to compute the SQRT and LOG functions used to generate Gaussian pseudorandom numbers. When the resulting numbers are close to the histogram boundaries in the verification test, a full precision evaluation of these intrinsics is employed. Thus the scheme passes all the verification tests yet defeats the intent of this benchmark.

The second loophole involves replacing calls to the SQRT and LOG intrinsics by a single call to a Fortran coded function that returns the SQRT(-LOG(X)). Again this scheme will pass the verification test yet does not satisfy the intent since the Fortran intrinsic functions have not been employed in the implementation.

Two changes are here made to the benchmark specification. First, two checksums are now required as part of the verification test. Second, only Fortran intrinsic functions (or equivalent calls to the standard C math library) may be used for SQRT and LOG.

2.2 Changes to CG

The intent of the CG benchmark is to test the performance of the system for unstructured grid computations which by their nature require irregular long distance communication or memory access. The benchmark essentially requires computing a sparse matrix-vector product. Rather than distribute a multi-Mbyte file for the matrix, the compact subroutine **makea** is supplied to generate a random sparse matrix. The **makea** procedure generates a sparse matrix by summing outer products of random sparse vectors. This construction is intended to preclude the clever use of *a priori* knowledge of the matrix structure to reduce the communication requirement.

Nonetheless, by saving the random vectors used in **makea**, it is possible to reformulate the sparse matrix-vector multiply and its associated irregular communication in a way such that communication is substantially reduced, and only a few dense vectors are communicated. All sparse operations can be kept local to the processing nodes.

Although this scheme of matrix-vector multiplication may be considered to satisfy the the rules of the CG benchmark, it defeats its intended purpose of measuring random communication performance. Therefore this scheme is no longer allowed and results employing this loophole will not be reported in future releases of this report. A strict interpretation of the benchmark specification [2] precludes this scheme since it is clearly stated that the conjugate gradient method will be used to compute the solution z to $Az = x$, and as part of this method the vector q must be computed via the product $q = Ap$. This means the matrix A must be used, not the vectors employed in its construction.

3 Kernel Results

3.1 Embarrassingly Parallel (EP) Benchmark

The first of the five kernel benchmarks is an "embarrassingly parallel" problem. In this benchmark, two-dimensional statistics are accumulated from a large number of Gaussian pseudorandom numbers, which are generated according to a particular scheme that is well-suited for parallel computation. This problem is typical of many "Monte-Carlo" applications. Since it requires almost no communication, in some sense this benchmark provides an estimate of the upper achievable limits for floating point performance on a particular system. Discussion on the parallel implementation of this benchmark may be found in [3].

Results for the embarrassingly parallel benchmark are shown in Table 2. Not all systems exhibit high rates on this problem. This appears to stem from the fact that this benchmark requires references to several mathematical intrinsic functions, such as the Fortran routines AINT, SQRT, and LOG, and evidently these functions are not highly optimized on some systems.

Results which have employed the reduced precision table lookup scheme described in Section 2.1 are unacceptable and not listed in the tables. The SunMos-turbo operating system for the Paragon allows both i860 processors on the node to be used for computation (in regular SunMos and OSF the second processor is used purely for communication).

Intel Paragon results are due to S. Gupta and T. Phung of Intel. CM-2, CM-200 and CM-5 results are due to J. Richardson of TMC. SP-2 results are due to R. Agarwal, F. Gustavson and M. Zubair of IBM. KSR1 and KSR2 results are due to S. Breit (KSR), J. Singer (U. Houston), and G. Shah (Georgia Tech). VPP500 results are due to B. Elton of Fujitsu. CS-2 results are due to J. Cownie and K. Pickard of Meiko and L. Meadows of Portland Group. SX-3 results are due to G.M. Sastri of NEC. Power Challenge and Power Indigo results are due to J. Richardson of SGI. Distributed workstation results are due to S. White of Emory University [14] except for the SGI results which are due to D. Browning of the NAS System Development branch. The "Mixed-A" computer system consisted of 16 Sun Sparc 1's, one Sun IPC, one Sun Sparc2, 11 Sun SLC's, three IBM RS6000 model 550's, one IBM RS6000 model 530, and one NeXT machine. The listed PVM results used PVM 2.4 and Ethernet.

3.2 Multigrid (MG) Benchmark

The second kernel benchmark is a simplified multigrid kernel, which solves a 3-D Poisson PDE. This problem is simplified in the sense that it has constant rather than variable coefficients as in a more realistic application. This code is a good test of both short and long distance highly structured communication. The Class B problem uses the same size grid but a greater number of outer loop iterations.

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
BBN TC2000	Dec 91	64	284.0	0.44
Convex SPP1000	June 94	1	376.8	0.33
		4	96.0	1.31
		8	48.1	2.62
	Oct 94	16	24.3	5.19
Cray Y-MP	Aug 92	1	126.2	1.00
		8	15.9	7.95
Cray C-90	Oct 94	1	46.31	2.72
		4	11.59	10.89
		8	5.84	21.60
		16	2.95	42.81
Cray T3D	Oct 94	16	35.04	3.60
		32	17.52	7.20
		64	8.76	14.41
		128	4.38	28.81
		256	2.19	57.63
		512	1.09	115.78
		1024	0.55	229.45
Fujitsu VPP500	Aug 94	1	44.25	2.85
		4	11.24	11.21
		8	5.67	22.22
		16	2.87	43.97
		32	1.46	86.44
		64	0.75	167.92
IBM SP-2	Aug 94	8	44.26	2.85
		16	22.15	5.70
		32	11.08	11.38
		64	5.53	22.82
Intel iPSC/860	May 92	32	102.7	1.23
		64	51.4	2.46
		128	25.7	4.91
Intel Paragon (OSF1.2)	Mar 94	64	10.45	12.1
		128	5.24	24.1
		256	2.66	47.4
		512	1.38	91.4
Intel Paragon (SunMos turbo)	Mar 94	64	5.27	23.9
		128	2.76	45.7
		256	1.46	86.4
Kendall Square KSR1	Oct 93	16	101.9	1.2
		32	51.4	2.5
		64	26.0	4.9
		128	12.8	9.9
Kendall Square KSR2	Feb 94	32	24.8	5.1
	May 94	64	13.0	9.7
Kyoto/Matsushita ADENART	Feb 94	256	32.9	3.8
MasPar MP-1	Aug 92	4K	248.0	0.51
		16K	69.3	1.82
MasPar MP-2	Nov 92	16K	22.4	5.63

Table 2a: Results of the Class A Embarrassingly Parallel (EP) Benchmark (cont'd)

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
Meiko CS-1	Aug 92	16	116.8	1.08
Meiko CS-2	Oct 94	16	39.39	3.20
		32	20.45	6.16
		64	11.00	11.46
		96	7.84	16.07
		128	6.29	20.06
nCUBE-2S	Mar 94	64	83.8	1.51
		128	41.93	3.01
		256	20.97	6.02
		512	10.50	12.02
		1024	5.25	24.03
NEC SX-3	Oct 94	1	21.27	5.93
Silicon Graphics	Oct 94	1	242.95	0.52
Power Challenge XL		4	61.44	2.05
		8	30.77	4.10
		16	15.48	8.15
Silicon Graphics Power Indigo	Oct 94	1	244.18	0.52
Thinking Machines CM-2	Oct 91	8K	126.6	1.00
		16K	63.9	1.97
		32K	33.7	3.74
		64K	18.8	6.71
Thinking Machines CM-200	Oct 91	8K	76.9	1.64
		16K	39.2	3.22
		32K	20.7	6.10
		64K	10.9	11.58
Thinking Machines CM-5	Nov 92	16	42.4	2.98
		32	21.5	5.88
		64	10.9	11.62
		128	5.4	23.49
		256	2.7	46.84
		512	1.4	90.47
Thinking Machines CM-5E	Feb 94	32	11.5	11.0
		64	5.7	22.1
		128	3.0	42.1
PVM Sparcs (Ethernet)	Sep 93	16	1670.0	0.08
PVM RS6000-550 (Ethernet)	Sep 93	4	890.0	0.14
PVM Mixed-A (Ethernet)	Sep 93	34	494.0	0.26
PVM SGI 4D25 (Ethernet)	Sep 93	4	2536.4	0.05
Parasoft IBM (token)	Jan 94	9	589.0	0.2

Table 2a: (cont'd) Results of the Class A Embarrassingly Parallel (EP) Benchmark

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to C90/1
Convex SPP1000	Aug 94	8	191.01	0.97
		16	96.55	1.92
Cray C90	Oct 94	1	185.26	1.00
		4	46.58	3.98
		8	23.20	7.98
		16	11.80	15.70
Cray T3D	Oct 94	16	140.13	1.32
		32	70.06	2.64
		64	35.03	5.29
		128	17.51	10.58
		256	8.76	21.15
		512	4.38	42.30
		1024	2.19	84.59
Fujitsu VPP500	Aug 94	1	176.64	1.05
		4	44.52	4.16
		8	22.36	8.29
		16	11.26	16.45
		32	5.68	32.62
		64	2.88	64.33
IBM SP-2	Aug 94	8	153.90	1.20
		16	77.08	2.40
		32	38.28	4.84
		64	19.42	9.54
	Oct 94	128	9.60	19.30
Intel Paragon (OSF1.2)	Mar 94	64	41.74	4.44
		128	20.86	8.88
		256	10.47	17.69
		512	5.26	35.22
Intel Paragon (SunMos turbo)	Mar 94	64	21.18	8.75
		128	10.49	17.66
		256	5.41	34.24
Kendall Square KSR2	May 94	64	46.6	3.98
Meiko CS-2	Aug 94	16	152.81	1.21
		32	77.20	2.40
		64	39.48	4.69
		96	26.84	6.90
		128	21.16	8.76
nCUBE-2S	Mar 94	64	336.3	0.55
		128	168.2	1.10
		256	84.1	2.20
		512	42.1	4.40
		1024	21.0	8.82
NEC SX-3	Oct 94	1	81.58	2.27
Silicon Graphics	Oct 94	1	973.62	0.19
Power Challenge XL		4	245.74	0.75
		8	122.98	1.51
		16	61.79	3.00
Thinking Machines CM-5E	Feb 94	32	46.9	3.95
		64	23.6	7.85
		128	11.6	15.97

Table 2b: Results of the Class B Embarrassingly Parallel (EP) Benchmark

Results for this benchmark are shown in Table 3. Intel Paragon results are due to J. Patterson of BCS and E. Kushner of Intel. CM-2 and CM-200 results are due to J. Richardson at TMC. RS6000-590 results are due to L.J. Shieh of IBM. SP-1 and SP-2 results are due to R. Lawrence and C. Douglas of IBM. KSR1 and KSR2 results are due to G. Montry of Southwest Software. VPP500 results are due to J.C.H. Wang of Fujitsu. CS-2 results are due to J. Cownie of Meiko. SX-3 results are due to G.M. Sastri of NEC. Distributed workstation results are due to S. White of Emory University [14] using PVM 2.4 and Ethernet except where noted otherwise.

3.3 Conjugate Gradient (CG) Benchmark

In this benchmark, a conjugate gradient method is used to compute an approximation to the smallest eigenvalue of a large, sparse, symmetric positive definite matrix. This kernel is typical of unstructured grid computations in that it tests irregular long distance communication and employs sparse matrix vector multiplication.

An unfortunate inconsistency has developed in the specification of the Class A size CG benchmark. The original benchmark description (as written in RNR Technical Report RNR-91-002) specified 15 iterations, however subsequent publications (specifically [2]) specify 25 iterations. For historical consistency we continue to report timings for 15 iterations, and results we have received based on 25 iterations have been scaled by 15/25. (The benchmark time scales linearly with number of iterations.)

Results which have circumvented the sparse matrix-vector multiplication by retaining elements of the matrix construction, as described in Section 2.2, are unacceptable and not listed in the tables.

The irregular communication requirement of this benchmark is evidently a challenge for all systems. Results are shown in Table 4. CM-2 results are due to J. Richardson of TMC. Intel iPSC/860 and nCUBE-2 results are by B. Hendrickson, R. Leland, and S. Plimpton of Sandia National Laboratory[9]. Paragon results are due to S. Gupta of Intel, R. van de Geijn of U.T. Austin and John Lewis of BCS[10]. Cray EL and C90 results are due to M. Zagha of Carnegie Mellon University. VPP500 results are due to J. Wang of Fujitsu. SP-1 results are due to D. Klepacki of IBM. SP-2 results are due to D. Klepacki, B. Alpern and L. Carter of IBM. KSR1 and KSR2 results are due to S. Breit and J. Middlecoff of KSR. CS-2 results are due to D. Daniel of Meiko. Power Challenge and Power Indigo results are due to F. Shakib of SGI. Distributed workstation results are due to S. White of Emory University [14] using PVM 2.4 and Ethernet except where noted otherwise.

3.4 3-D FFT PDE (FT) Benchmark

In this benchmark a 3-D partial differential equation is solved using FFTs. This kernel performs the essence of many "spectral" codes. It is a good test of long-distance communication performance. Discussion on the parallel implementation of this benchmark may be found in [3].

The rules of the NAS Parallel Benchmarks specify that assembly-coded, library routines may be used to perform matrix multiplication and one-dimensional, two-dimensional or three-dimensional FFTs. Thus this benchmark is somewhat unique in that computational library routines may be legally employed.

Results are shown in Table 5. Intel Paragon results are due to E. Kushner and T. Phung of Intel. VPP500 results are due to S. Zarantonello of Fujitsu. CM-2 and CM-200 results are due to J. Richardson of TMC. RS6000-590, SP-1 and SP-2 results are due to F. Gustavson, M. Zubair and R. Agarwal of IBM. KSR1 and KSR2 results are due to N. Camp of KSR. MP-1 and MP-2 results are due to J. Fier of MasPar. CS-2 results are due D. Daniel of Meiko. SX-3 results are due to G.M. Sastri of NEC. Power Challenge and Power Indigo results are due to J. Fier of SGI.

3.5 Integer Sort (IS) Benchmark

This benchmark tests a sorting operation that is important in "particle method" codes. This type of application is similar to "particle in cell" applications of physics, wherein particles are assigned to cells and may drift out. The sorting operation is used to reassign particles to the appropriate cells. This benchmark tests both integer computation speed and communication performance. For discussion on general parallel algorithms for this benchmark see [7].

This problem is unique in that floating point arithmetic is not involved. Significant data communication, however, is required. Results are shown in Table 6. Intel Paragon results are due to to S. Gupta and B.

Computer System	Date Received	No. Proc.	Time (sec)	Ratio to Y-MP/1
Convex SPP1000	Jun 94	1	208.0	0.11
		4	54.9	0.40
		8	30.9	0.72
Cray Y-MP	Aug 92	1	22.22	1.00
		8	2.96	7.51
Cray EL	Aug 92	1	89.19	0.25
		4	27.94	0.80
		8	22.30	0.95
Cray C-90	Dec 93	1	8.15	2.7
		4	2.19	10.1
	Aug 92	16	0.96	23.14
Cray T3D	Oct 94	16	14.15	1.57
		32	6.48	3.43
		64	2.69	8.26
		128	1.40	15.87
		256	0.76	29.24
		512	0.41	54.20
		1024	0.25	88.88
Fujitsu VPP500	Aug 94	4	1.58	14.06
		8	0.86	25.84
		16	0.49	45.35
		32	0.33	67.33
IBM RS6000-590	Mar 94	1	41.78	0.53
IBM SP-1	Mar 94	8	17.50	1.27
		16	9.49	2.34
		32	5.10	4.36
		64	2.89	7.69
IBM SP-2	Aug 94	8	6.36	3.49
		16	3.32	6.69
		32	1.81	12.28
		64	1.00	22.22
Intel iPSC/860	Aug 92	128	8.6	2.58
Intel Paragon (OSF1.2)	Mar 94	64	8.4	2.6
		128	4.5	4.9
		256	3.0	7.4
Intel Paragon (SunMos)	Feb 94	64	9.76	2.3
		128	5.10	4.4
		256	3.48	6.4
Kendall Square KSR1	Feb 94	32	19.7	1.1
		64	10.3	2.2
		128	5.6	4.0
Kendall Square KSR2	Feb 94	32	10.3	2.20
	May 94	64	5.7	3.90

Table 3a: Results of the Class A Multigrid (MG) Benchmark (cont'd)

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
Kyoto/Matsushita ADENART	Feb 94	256	21.4	1.0
MasPar MP-1	Aug 92	16K	12.0	1.9
MasPar MP-2	Nov 92	16K	4.36	5.1
Meiko CS-1	Aug 92	16	42.8	0.5
Meiko CS-2	Oct 94	16	7.60	2.93
		64	2.35	9.83
		128	1.43	15.54
nCUBE-2S	Mar 94	64	37.6	0.6
		128	19.2	1.2
		512	5.3	4.2
		1024	2.8	7.9
NEC SX-3	Oct 94	1	2.80	7.94
Thinking Machines CM-2	Dec 91	16K	45.8	0.5
		32K	26.0	0.9
		64K	14.1	1.6
Thinking Machines CM-200	Dec 91	16K	30.2	0.7
		32K	17.2	1.3
Thinking Machines CM-5	Aug 93	32	19.5	1.1
		64	10.9	2.0
		128	6.1	3.6
Thinking Machines CM-5E	Feb 94	32	3.9	5.7
		64	2.3	9.9
		128	1.3	16.6
PVM RS6000-550 (Ethernet)	Sep 93	4	293.0	0.1
PVM RS6000-560 (FDDI)	Sep 93	4	184.0	0.1
	Sep 93	8	110.4	0.2

Table 3a: (cont'd) Results of the Class A Multigrid (MG) Benchmark

Computer System	Date Received	No. Proc.	Time (sec)	Ratio to C90/1
Cray C90	Dec 93	1	37.77	1.0
		4	9.71	3.9
		16	3.97	9.5
Cray T3D	Oct 94	16	66.58	0.57
		32	30.42	1.24
		64	12.56	3.01
		128	6.57	5.75
		256	3.60	10.49
		512	1.88	20.09
		1024	1.15	32.84
Fujitsu VPP500	Oct 94	4	7.53	5.02
		8	4.07	9.28
		16	2.35	16.07
		32	1.56	24.21
IBM RS6000-590	Mar 94	1	184.92	0.2
IBM SP-1	Mar 94	8	82.03	0.46
		16	44.57	0.85
		32	24.37	1.55
		64	13.86	2.73
IBM SP-2	Aug 94	8	28.77	1.31
		16	15.09	2.50
		32	8.21	4.60
		64	4.53	8.34
	Oct 94	128	2.63	14.36
Intel Paragon (OSF1.2)	Mar 94	64	39.8	0.9
		128	21.3	1.8
		256	13.7	2.8
Intel Paragon (SunMos)	Feb 94	64	43.02	0.9
		128	24.15	1.6
		256	16.74	2.3
Kendall Square KSR2	May 94	64	26.1	1.45
Meiko CS-2	Oct 94	16	35.46	1.07
		64	10.76	3.51
		128	6.55	5.77
NEC SX-3	Oct 94	1	13.16	2.87
Thinking Machines CM-5E	Feb 94	32	20.9	1.8
		64	11.3	3.3
		128	6.7	5.6

Table 3b: Results of the Class B Multigrid (MG) Benchmark

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
BBN TC2000	Dec 91	40	51.4	0.23
Convex SPP1000	Jun 94	1	202.9	0.06
		4	49.7	0.24
		8	23.9	0.50
	Oct 94	16	12.0	0.99
Cray Y-MP	Aug 92	1	11.92	1.00
		8	2.38	5.01
Cray EL	Sep 93	1	45.24	0.26
		4	14.29	0.83
		8	10.14	1.18
Cray C-90	Sep 93	1	3.55	3.36
		4	0.96	12.42
		16	0.34	35.06
Cray T3D	Jul 94	16	14.98	0.80
		32	7.46	1.60
		64	4.20	2.84
		128	2.23	5.35
		256	1.30	9.17
	Oct 94	512	0.81	14.72
		1024	0.58	20.55
Fujitsu VPP500	Mar 94	1	5.68	2.10
		2	3.06	3.90
		4	1.72	6.93
		8	1.04	11.46
	Aug 94	16	0.80	14.90
IBM SP-1	Feb 94	8	21.37	0.6
		16	12.82	0.9
		32	7.98	1.5
		64	4.72	2.5
IBM SP-2	Aug 94	8	4.91	2.43
		16	3.15	3.78
		32	2.45	4.86
		64	1.81	6.58
Intel iPSC/860	Sep 93	128	7.0	1.71
Intel Paragon (OSF1.2)	Mar 94	64	4.10	2.9
		128	3.30	3.6
		256	2.83	4.2
Intel Paragon (SunMos)	Nov 93	64	12.6	1.0
Kendall Square KSR1	Feb 94	32	19.0	0.6
		64	13.4	0.9
Kendall Square KSR2	Feb 94	32	9.8	1.2
	May 94	64	6.1	1.95

Table 4a: Results of the Class A Conjugate Gradient (CG) Benchmark (cont'd)

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
Kyoto/Matsushita ADENART	Feb 94	256	10.8	1.1
MasPar MP-1	Aug 92	4K	64.5	0.18
		16K	14.6	0.82
MasPar MP-2	Nov 92	16K	11.0	1.08
Meiko CS-1	Aug 92	16	67.5	0.18
Meiko CS-2	Oct 94	16	7.18	1.66
		32	5.60	2.10
nCUBE-2S	Mar 94	64	29.6	0.4
		128	16.9	0.7
		256	9.6	1.3
		512	6.2	1.9
		1024	4.1	2.9
Silicon Graphics	Oct 94	1	39.0	0.31
Power Challenge XL		2	16.9	0.71
		4	7.2	1.66
		8	4.5	2.65
		16	3.5	3.41
Silicon Graphics Power Indigo	Oct 94	1	52.27	0.22
Thinking Machines CM-2	Mar 92	8K	25.6	0.47
		16K	14.1	0.85
		32K	8.8	1.35
Thinking Machines CM-200	Mar 92	8K	15.0	0.79
Thinking Machines CM-5	Aug 93	32	20.7	0.58
		64	10.6	1.12
		128	6.2	1.92
PVM RS6000-550 (Ethernet)	Sep 93	4	203.2	0.06
PVM RS6000-560 (FDDI)	Sep 93	4	81.5	0.15
Parasoft IBM (token)	Jan 94	9	277	0.04

Table 4a: (cont'd) Results of the Class A Conjugate Gradient (CG) Benchmark

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to C90/1
Cray C90	Dec 93	1	122.90	1.00
		4	33.19	3.7
		16	10.61	11.6
Cray T3D	Jul 94	16	582.05	0.21
		32	298.62	0.41
		64	166.57	0.74
		128	85.51	1.44
		256	50.18	2.45
	Oct 94	512	27.34	4.50
		1024	16.58	7.41
Fujitsu VPP500	Apr 94	2	104.51	1.18
		4	55.40	2.22
		8	31.80	3.86
	Aug 94	15	20.85	5.89
		30	15.21	8.08
IBM RS6000-590	Mar 94	1	429.0*	0.3*
IBM SP-1	Mar 94	16	638.2	0.2
		32	362.9	0.3
		64	193.4	0.6
IBM SP-2	Aug 94	8	165.70	0.74
		16	93.72	1.31
		32	64.21	1.91
		64	42.68	2.88
	Oct 94	128	26.79	4.59
Intel Paragon (OSF1.2)	Mar 94	128	132.5	0.9
	Jul 94	256	70.0	1.76
		512	47.6	2.58
Kendall Square KSR2	May 94	64	182.0	0.68
Meiko CS-2	Oct 94	16	248.30	0.49
		32	156.50	0.78
Thinking Machines CM-5E	Feb 94	32	449.0*	0.3*
		64	199.0*	0.6*
		128	92.0*	1.3*

Table 4b: Results of the Class B Conjugate Gradient (CG) Benchmark (* indicates result used matrix construction to circumvent sparse matrix-vector multiplication)

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
Convex SPP1000	Aug 94	1	178.57*	0.16
		4	46.78*	0.62
		8	25.54*	1.13
Cray Y-MP	Aug 92	1	28.77*	1.00
		8	4.19*	6.87
Cray EL	May 93	1	105.1*	0.27
		4	27.9*	1.03
		8	18.5*	1.56
Cray C-90	Aug 92	1	10.28*	2.80
		4	2.58*	11.20
		16	0.91*	31.60
Cray T3D	Oct 94	16	11.86*	2.43
		32	6.00*	4.80
		64	3.07*	9.37
	Jul 94	128	1.57*	18.32
		256	0.80*	35.96
	Oct 94	512	0.54*	53.28
		1024	0.32*	89.91
Fujitsu VPP500	Aug 94	4	2.93	9.82
		8	1.45	19.81
		16	0.75	38.51
		32	0.40	72.47
		64	0.24	121.91
IBM RS6000-590	Feb 94	1	61.01*	0.5
IBM SP-1	Feb 94	8	43.68*	0.7
		16	22.86*	1.3
		32	12.08*	2.4
		64	6.46*	4.5
IBM SP-2	Aug 94	8	14.59*	1.97
		16	7.79*	3.70
		32	4.87*	5.91
		64	2.42*	11.89
Intel iPSC/860	Dec 91	64	20.9*	1.37
	Apr 92	128	9.7*	2.96
Intel Paragon (OSF1.2)	Mar 94	64	9.1*	3.2
		128	4.9*	5.9
		256	3.6*	8.0
Intel Paragon (SunMos)	Mar 94	64	7.2*	4.0
		128	3.9*	7.4
		256	3.0*	9.7
Kendall Square KSR1	Feb 94	32	16.2*	1.8
		64	9.2*	3.1
Kendall Square KSR2	Feb 94	32	9.0*	3.2
	May 94	64	6.5*	4.43

Table 5a: Results of the Class A 3-D FFT PDE (FT) Benchmark (cont'd) (* indicates library result).

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
Kyoto/Matsushita ADENART	Feb 94	256	72.7	0.4
MasPar MP-1	Aug 92	16K	18.3*	1.57
MasPar MP-2	Nov 92	16K	8.0*	3.60
Meiko CS-1	Aug 92	16	170.0*	0.17
Meiko CS-2	Oct 94	16	12.67	2.27
		32	7.17	4.01
		64	4.53	6.35
nCUBE-2S	Mar 94	64	62.8*	0.5
		128	32.9*	0.9
		256	16.0*	1.8
		512	8.4*	3.4
		1024	4.1*	7.0
NEC SX-3	Oct 94	1	2.79*	10.31
Silicon Graphics	Oct 94	1	61.17*	0.47
Power Challenge XL		2	35.53*	0.81
		4	19.98*	1.44
		8	12.57*	2.29
		16	11.18*	2.57
Thinking Machines CM-2	Dec 91	16K	37.0*	0.78
		32K	18.2*	1.58
		64K	11.4*	2.52
Thinking Machines CM-200	Dec 91	8K	45.6*	0.63
Thinking Machines CM-5	Aug 93	32	14.9*	1.93
		64	7.9*	3.64
		128	6.6*	4.36
Thinking Machines CM-5E	Feb 94	32	7.4*	3.9
		64	3.9*	7.4
		128	2.9*	9.9

Table 5a: (cont'd) Results of the Class A 3-D FFT PDE (FT) Benchmark (* indicates library result).

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to C90/1
Convex SPP1000	Aug 94	8	375.43*	0.34
Cray C90	Dec 93	1	127.44*	1.00
		2	63.74*	2.0
		16	8.43*	15.1
Cray T3D	Oct 94	64	40.80*	3.12
		128	20.96*	6.08
	July 94	256	10.89*	11.70
	Oct 94	512	6.73*	18.94
		1024	3.76*	33.89
Fujitsu VPP500	Aug 94	16	7.95	16.03
		32	4.07	31.33
		64	2.18	58.54
IBM RS6000-590	Mar 94	1	856.3*	0.1
IBM SP-1	Mar 94	16	286.5*	0.4
		32	143.2*	0.9
		64	74.5*	1.7
IBM SP-2	Aug 94	16	96.02*	1.33
		32	52.98*	2.40
		64	28.50*	4.47
	Oct 94	128	14.57*	8.75
Kendall Square KSR2	May 94	64	124.0*	1.03
Intel Paragon (OSF1.2)	Mar 94	128	56.5*	2.3
		256	30.6*	4.2
Intel Paragon (SunMos)	Feb 94	256	25.1*	5.1
Meiko CS-2	Oct 94	32	82.71	1.54
		64	48.04	2.65
NEC SX-3	Oct 94	1	37.52*	3.40
Silicon Graphics	Oct 94	1	761.67*	0.17
Power Challenge XL		2	414.52*	0.31
		4	223.97*	0.57
		8	130.15*	0.98
		16	110.37*	1.15
Thinking Machines CM-5E	Feb 94	32	89.0*	1.4
		64	46.0*	2.8
		128	34.0*	3.7

Table 5b: Results of the Class B 3-D FFT PDE (FT) Benchmark (* indicates library result).

Greer of Intel. CM-2, CM-200 and MasPar results use a library sorting routine. Cray Y-MP results are due to CRI. Cray C-90 and EL results are due to M. Zagha of Carnegie Mellon University using a radix sort optimized for interleaved memories [16]. VPP500 results are due to B. Elton of Fujitsu. RS6000-590, SP-1 and SP-2 results are due to F. Gustavson, M. Zubair and R. Agarwal of IBM. KSR1 and KSR2 results are due to C. Nowacki of KSR.

4 Simulated CFD Application Benchmarks

The three simulated CFD application benchmarks are intended to accurately represent the principal computational and data movement requirements of modern CFD applications.

The first of these is the called the lower-upper diagonal (LU) benchmark. It does not perform a LU factorization but instead employs a symmetric successive over-relaxation (SSOR) numerical scheme to solve a regular-sparse, block (5×5) lower and upper triangular system. This problem represents the computations associated with a newer class of implicit CFD algorithms, typified at NASA Ames by the code "INS3D-LU". This problem exhibits a somewhat limited amount of parallelism compared to the next two. Discussion of the serial algorithm underlying this benchmark may be found in [15]. Discussion of the parallel algorithms may be found in [4].

The second simulated CFD application is called the scalar pentadiagonal (SP) benchmark. In this benchmark, multiple independent systems of non-diagonally dominant, scalar pentadiagonal equations are solved. The third simulated CFD application is called the block tridiagonal (BT) benchmark. In this benchmark, multiple independent systems of non-diagonally dominant, block tridiagonal equations with a 5×5 block size are solved.

SP and BT are representative of computations associated with the implicit operators of CFD codes such as "ARC3D" at NASA Ames. SP and BT are similar in many respects, but there is a fundamental difference with respect to the communication to computation ratio. Discussion of the serial algorithm underlying this benchmark may be found in [12].

Performance figures for the three simulated CFD applications are shown in Tables 7, 8 and 9. Timings are cited as complete run times, in seconds, as with the other benchmarks. A complete solution of the LU benchmark requires 250 iterations. For the SP benchmark, 400 iterations are required. For the BT benchmark, 200 iterations are required.

For LU, credits are as follows: iPSC/860 and CM-2 results are due to S. Weeratunga, R. Fatoohi, E. Barszcz and V. Venkatakrishnan of NAS; VPP500 results are due to C. Chen of Fujitsu; CM-5 results are due to J. Richardson and D. Sandee of TMC; MP-1 and MP-2 results are due to J. McDonald of MasPar; Intel Paragon results are due to T. Phung and E. Kushner of Intel; KSR1 and KSR2 results are due to S. Breit of KSR; RS600-590 results are due to L.E. Hannon of IBM; SP-1 and SP-2 results are due to V. Naik of IBM; nCUBE-2S results are due to E. Schulman of nCUBE; Power Challenge and Power Indigo results are due to J. McDonald of SGI.

For SP, credits are as follows: CM-2 results employ a library scalar pentadiagonal solver; CM-5 results are due to J. Richardson and D. Sandee of TMC; iPSC/860 results are due to J. Patterson of BCS; Paragon results are due to T. Phung of Intel for tranpose algorithm, and R. van de Wijngaar of MCAT for multipartition method; MP-1 and MP-2 results are due to J. McDonald of MasPar; KSR1 and KSR2 results are due to S. Breit of KSR and G. Shah of Georgia Tech; RS600-590 results are due to L.J. Shieh of IBM; SP-1 and SP-2 results are due to V. Naik of IBM; VPP500 results are due to S. Gavali of Fujitsu; SX-3 results are due to G.M. Sastri of NEC; nCUBE-2S results are due to E. Schulman of nCUBE; Power Challenge and Power Indigo results are due to J. Bannning of SGI.

For BT, credits are as follows: CM-2 and CM-200 results employ a library block tridiagonal solver; CM-5 results are due to J. Richardson and D. Sandee of TMC; iPSC/860 results are due to J. Patterson of BCS; Paragon results are due to T. Phung of Intel; MP-1 and MP-2 results are due to J. McDonald of MasPar; KSR1 and KSR2 results are due to S. Breit of KSR; RS600-590 results are due to L.J. Shieh of IBM; SP-1 and SP-2 results are due to V. Naik of IBM; VPP500 results are due to H. Lai of Fujitsu and staff of Fujitsu Limited; CS-2 results are due G. Montry of Southwest Software. SX-3 results are due to G.M. Sastri of NEC; nCUBE-2S results are due to E. Schulman of nCUBE; Power Challenge and Power Indigo results are due to J. McDonald of SGI.

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
Convex SPP1000	Jun 94	1	76.8	0.15
		4	24.5	0.47
		8	14.6	0.78
Cray Y-MP	Aug 92	1	11.46	1.00
		8	1.85	6.19
Cray EL	Sep 93	1	43.76	0.26
		4	12.99	0.88
		8	8.45	1.35
Cray C-90	Sep 93	1	3.33	3.44
		4	0.85	13.46
		16	0.27	42.38
Cray T3D	Oct 94	16	11.86	0.97
		32	5.87	1.95
		64	2.89	3.97
		128	1.49	7.69
		256	0.81	14.15
		512	0.54	21.22
		1024	0.44	26.05
Fujitsu VPP500	Apr 94	1	2.189	5.24
		2	1.574	7.28
		4	1.098	10.44
		8	0.917	12.50
IBM RS6000-590	Feb 94	1	21.73	0.5
IBM SP-1	Feb 94	8	16.81	0.7
		16	8.85	1.3
		32	5.04	2.3
		64	3.06	3.7
IBM SP-2	Aug 94	8	5.00	2.29
		16	2.79	4.11
		32	1.77	6.47
		64	0.93	12.32
Intel iPSC/860	May 92	32	25.7	0.45
		64	17.3	0.66
		128	13.6	0.84
Intel Paragon (OSF1.2)	Mar 94	32	7.81	1.5
		64	4.34	2.6
		128	2.41	4.8
Intel Paragon (SunMos)	Mar 94	32	5.48	2.1
		64	3.77	3.0
Kendall Square KSR1	Feb 94	32	10.8	1.1
		64	6.6	1.7
Kendall Square KSR2	Feb 94	32	7.0	1.6
	May 94	64	3.9	2.94

Table 6a: Results of the Class A Integer Sort (IS) Benchmark (cont'd) (* indicates library result).

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
Kyoto/Matsushita ADENART	Feb 94	256	46.6	0.3
MasPar MP-1	Jan 93	16K	11.5*	1.00
MasPar MP-2	Jan 93	16K	7.7*	1.49
Meiko CS-1	Aug 92	16	62.7	0.18
nCUBE-2S	Mar 94	64	23.2	0.5
		128	12.0	1.0
		256	6.1	1.9
		512	3.2	3.6
		1024	1.7	6.8
Thinking Machines CM-2	Dec 91	16K	35.8*	0.32
		32K	21.0*	0.55
		64K	14.9*	0.77
Thinking Machines CM-200	Dec 91	64K	5.7*	2.01
Thinking Machines CM-5	Aug 93	32	43.1	0.27
		64	24.2	0.47
		128	12.0	0.96
Thinking Machines CM-5E	Feb 94	32	6.3	1.8
		64	3.1	3.7
		128	1.66	6.9

Table 6a: (cont'd) Results of the Class A Integer Sort (IS) Benchmark (* indicates library result).

5 Sustained Performance Per Dollar

One aspect of the relative performance of these systems has not been addressed so far, namely the differences in price between these systems. One way to compensate for these price differences is to compute sustained performance per million dollars, i.e. the performance ratio figures shown in Tables 2 through 9 divided by the list price in millions. Some figures of this type are shown in Table 11 for two of the benchmarks (the Class B size MG and SP benchmarks) for the most recent of the systems tested. The table includes the list price of the minimal system (in terms of memory per node, disk space, etc.) required to run the full Class B size NPB as implemented by the vendor. These prices were provided by the vendors and include any associated software costs (i.e. operating system, compilers, scientific libraries as required, etc.) but do not include maintenance. Hardware configurations for the various systems as tested by the vendors and associated list prices are provided in Table 10. Be aware that list prices are similar to peak performance in that they are guaranteed not to be exceeded.

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to C90/1
Cray C90	Dec 93	1	12.92	1.00
		4	3.30	3.9
		16	0.98	13.7
Cray T3D	Oct 94	32	25.46	0.51
		64	12.88	1.00
		128	6.57	1.97
		256	3.27	3.95
		512	1.94	6.66
		1024	1.22	10.59
Fujitsu VPP500	Apr 94	4	3.70	3.49
		8	3.03	4.26
IBM RS6000-590	Mar 94	1	91.6	0.1
IBM SP-1	Mar 94	16	37.3	0.3
		32	20.1	0.6
		64	11.2	1.2
IBM SP-2	Aug 94	8	19.98	0.65
		16	11.04	1.17
		32	6.88	1.88
		64	3.55	3.64
	Oct 94	128	1.99	6.49
Intel Paragon (OSF1.2)	Mar 94	64	17.33	0.7
		128	9.52	1.4
		256	5.94	2.2
		512	4.69	2.8
Intel Paragon (SunMos)	Mar 94	64	11.98	1.1
		128	7.22	1.8
Kendall Square KSR2	May 94	64	20.3	0.64
nCUBE-2S	Mar 94	128	47.5	0.3
		512	12.5	1.0
		1024	6.5	2.0
Thinking Machines CM-5E	Feb 94	32	32.0	0.4
		64	16.4	0.8
		128	8.4	1.5

Table 6b: Results of the Class B Integer Sort (IS) Benchmark

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
BBN TC2000	Dec 91	62	3032.0	0.11
Convex SPP1000	Oct 94	1	2668.0	0.13
		4	597.0	0.56
		8	331.0	1.01
		16	209.0	1.60
Cray Y-MP	Aug 92	1	333.5	1.00
		8	49.5	6.74
Cray EL	Aug 92	1	1449.0	0.23
		4	522.3	0.64
		8	351.6	0.95
Cray C-90	Aug 92	1	157.6	2.12
		4	43.9	7.59
		16	17.6	18.93
Cray T3D	Oct 94	16	214.23	1.56
		32	111.82	2.98
		64	58.69	5.69
		128	30.38	10.98
		256	16.99	19.63
		512	9.59	34.78
		1024	7.09	47.04
Fujitsu VPP500	Aug 94	1	146.89	2.27
IBM RS6000-590	Mar 94	1	645.2	0.5
IBM SP-1	Feb 94	8	291.4	1.1
		16	172.9	1.9
		32	101.8	3.3
		64	63.2	5.3
IBM SP-2	Aug 94	8	116.23	2.87
		16	69.09	4.83
		32	38.90	8.57
		64	24.94	13.37
Intel iPSC/860	Mar 91	64	690.8	0.48
		128	442.5	0.75
Intel Paragon (OSF1.2)	Jul 94	64	190.0	1.76
		128	118.0	2.83
		256	75.0	4.45
Kendall Square KSR1	Feb 94	32	341.0	1.0
		64	199.0	1.7
		128	155.0	2.2
Kendall Square KSR2	Feb 94	32	172.0	1.9
	May 94	64	102.0	3.27

Table 7a: Results for the Class A LU Simulated CFD Application (cont'd)

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
Kyoto/Matsushita ADENART	Feb 94	256	327.5	1.0
MasPar MP-1	Aug 92	4K	1580.0	0.2
MasPar MP-2	Nov 92	4K	463.5	0.7
Meiko CS-1	Aug 92	16	2937.0	0.1
nCUBE-2S	Mar 94	64	1322.0	0.3
		128	712.5	0.5
		256	389.1	0.9
		512	226.1	1.5
		1024	134.1	2.5
Silicon Graphics	Jul 94	1	604.0	0.55
Power Challenge XL		4	231.8	1.44
		8	111.7	2.99
		16	65.3	5.11
Silicon Graphics Power Indigo	Oct 94	1	716.5	0.47
Thinking Machines CM-2	Mar 91	8K	1307.0	0.26
		16K	850.0	0.39
		32K	546.0	0.61
Thinking Machines CM-5	Aug 93	32	418.0	0.80
		64	272.0	1.23
		128	171.0	1.95
Thinking Machines CM-5E	Feb 94	32	152.0	2.2
		64	97.0	3.4
		128	65.0	5.1

Table 7a: (cont'd) Results for the Class A LU Simulated CFD Application

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to C90/1
Cray C90	Dec 93	1	648.5	1.00
		4	166.1	3.9
		16	51.6	12.6
Cray T3D	Oct 94	16	875.49	0.74
		32	470.82	1.38
		64	241.14	2.69
		128	124.48	5.21
		256	66.03	9.82
		512	36.39	17.82
		1024	20.77	31.22
Fujitsu VPP500	Aug 94	1	591.05	1.10
IBM RS6000-590	Mar 94	1	2694.6	0.2
IBM SP-1	Feb 94	16	604.8	1.1
		32	348.1	1.9
		64	207.5	3.1
IBM SP-2	Aug 94	8	434.59	1.49
		16	238.72	2.72
		32	135.92	4.77
		64	79.64	8.14
		128	49.82	13.02
Kendall Square KSR2	May 94	64	424.0	1.53
Intel Paragon (OSF1.2)	Jul 94	64	675.0	0.96
		128	406.0	1.60
		256	254.0	2.55
		512	175.0	3.71
Thinking Machines CM-5E	Feb 94	32	595.0	1.1
		64	367.0	1.8
		128	318.0	2.0
Silicon Graphics	Jul 94	1	2617.9	0.25
Power Challenge XL		4	1010.5	0.64
		8	550.2	1.18
		16	308.1	2.10

Table 7b: Results for the Class B LU Simulated CFD Application

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
BBN TC2000	Dec 91	112	880.0	0.54
Convex SPP1000	Oct 94	1	2813.0	0.17
		4	751.0	0.63
		8	379.0	1.24
		16	250.0	1.89
Cray Y-MP	Aug 92	1	471.5	1.00
		8	64.6	7.30
Cray EL	Aug 92	1	2025.7	0.23
		4	601.9	0.78
		8	488.4	0.97
Cray C-90	Aug 92	1	184.70	2.55
		4	49.74	9.48
		16	13.06	36.10
Cray T3D	Jul 94	16	206.08	2.29
		32	107.54	4.38
		64	55.39	8.51
		128	28.58	16.50
		256	15.31	30.80
	Aug 94	512	8.91	52.92
	Oct 94	1024	5.41	87.15
Fujitsu VPP500	Aug 94	1	176.75	2.67
		2	108.85	4.33
		4	57.24	8.24
		8	29.87	15.79
		16	20.99	22.47
IBM RS6000-590	Mar 94	1	993.1	0.5
IBM SP-1	Feb 94	8	441.6	1.1
		16	268.7	1.8
		32	165.0	2.9
		64	100.4	4.7
IBM SP-2	Aug 94	8	177.39	2.65
		16	100.19	4.71
		32	58.00	8.13
		64	34.77	13.56
Intel iPSC/860	Jul 94	64	640.0	0.74
	Aug 92	128	449.5	1.05
Intel Paragon (OSF1.2)	Jul 94	64	226.0	2.09
		128	143.0	3.30
		256	97.0**	4.86**
		324	89.0**	5.30**
Kendall Square KSR1	Feb 94	32	418.0	1.1
		64	257.0	1.8
		128	160.0	2.9
Kendall Square KSR2	Feb 94	32	221.0	2.1
	May 94	64	131.0	3.6

Table 8a: Results for the Class A SP Simulated CFD Application (cont'd) (* indicates library result; ** indicates multipartition algorithm instead of transpose algorithm).

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
Kyoto/Matsushita ADENART	Feb 94	256	209.9	2.3
MasPar MP-1	Aug 92	4K	1772	0.27
MasPar MP-2	Nov 92	4K	615	0.77
Meiko CS-1	Aug 92	16	2975	0.16
nCUBE-2S	Mar 94	64	1243.2	0.4
		128	717.4	0.7
	Jul 94	256	387.3	1.22
		512	208.6	2.26
		1024	120.9	3.90
NEC SX-3	Oct 94	1	75.72	6.23
Silicon Graphics	Jul 94	1	858.3	0.55
Power Challenge XL		4	225.8	2.09
		8	119.5	3.94
		16	67.2	7.01
Silicon Graphics Power Indigo	Oct 94	1	986.8	0.48
Thinking Machines CM-2	Dec 91	16K	1444.0*	0.33
		32K	917.0*	0.51
		64K	640.0*	0.74
Thinking Machines CM-5	May 93	32	289.0	1.63
		64	170.0	2.77
		128	119.0	3.96
Thinking Machines CM-5E	Feb 94	32	169.0	2.8
		64	104.0	4.5
		128	61.0	7.7

Table 8a: (cont'd) Results for the Class A SP Simulated CFD Application (* indicates library result; ** indicates multipartition algorithm instead of transpose algorithm).

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to C90/1
Convex SPP1000	Oct 94	8	1739.0	0.41
Cray C90	Dec 93	1	713.1	1.00
		4	203.1	3.5
		16	80.4	8.9
Cray T3D	Jul 94	16	818.07	0.87
		32	463.62	1.54
		64	242.69	2.94
		128	130.45	5.47
		256	77.29	9.23
	Oct 94	512	42.63	16.73
		1024	25.23	28.26
Fujitsu VPP500	Aug 94	1	664.76	1.07
		2	417.78	1.71
		4	228.37	3.12
		6	143.20	4.98
		8	120.05	5.94
		17	53.12	13.42
	Sep 94	34	39.01	18.28
IBM RS6000-590	Mar 94	1	4047.2	0.2
IBM SP-1	Feb 94	16	941.2	0.8
		32	522.4	1.4
		64	302.3	2.5
IBM SP-2	Aug 94	8	701.66	1.02
		16	368.02	1.94
		32	194.19	3.67
		64	111.19	6.41
	Oct 94	128	63.86	11.17
Intel Paragon (OSF1.2)	Mar 94	64	960.0	0.7
	Jul 94	102	610.0	1.17
		204	387.0	1.84
		256	301.0*	2.37*
		324	262.0*	2.72*
		400	246.0*	2.90*
		484	209.0*	3.41*
Kendall Square KSR2	May 94	64	495.0	1.44
NEC SX-3	Oct 94	1	294.68	2.42
Silicon Graphics	Jul 94	1	3719.5	0.19
Power Challenge XL		4	947.6	0.75
		8	491.4	1.45
		16	313.1	2.28
Thinking Machines CM-5E	Feb 94	32	1014.0	0.7
		64	595.0	1.2
		128	320.0	2.2

Table 8b: Results for the Class B SP Simulated CFD Application (* indicates multipartition algorithm instead of transpose algorithm).

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
BBN TC2000	Dec 91	112	1378.0	0.58
Convex SPP1000	Oct 94	1	2825.0	0.28
		4	732.0	1.08
		8	366.0	2.17
		16	211.0	3.76
Cray Y-MP	Aug 92	1	792.4	1.00
		8	114.0	6.95
Cray EL	May 93	1	3832.8	0.21
		4	1090.2	0.73
		8	764.1	1.04
Cray C-90	Aug 92	1	356.9	2.22
		4	96.1	8.25
		16	28.4	27.91
Cray T3D	Oct 94	16	234.12	3.38
		32	117.69	6.72
		64	60.55	13.09
	Jul 94	128	30.84	25.69
		256	15.89	49.87
	Aug 94	512	8.39	94.45
	Oct 94	1024	4.56	173.77
Fujitsu VPP500	Oct 94	2	75.17	10.54
		4	39.14	20.25
		8	19.82	39.98
		16	9.99	79.32
		32	5.09	155.68
		64	2.66	297.89
IBM RS6000-590	Feb 94	1	1249.4	0.6
IBM SP-1	Aug 94	8	443.9	1.78
		16	249.2	3.18
		32	143.0	5.54
		64	83.1	9.53
IBM SP-2	Aug 94	8	226.23	3.35
		16	128.11	6.18
		32	69.72	11.36
		64	39.87	19.87

Table 9a: Results for the Class A BT Simulated CFD Application (* indicates library result) (cont'd).

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to Y-MP/1
Intel iPSC/860	Aug 92	64	714.7	1.11
		128	414.3	1.91
Intel Paragon (OSF1.2)	Mar 94	64	235.0	3.4
		128	129.0	6.1
		256	83.0	9.5
		512	63.0	12.5
Intel Paragon (SunMos)	Nov 93	64	224.0	3.5
	Mar 94	128	113.0	7.0
Kendall Square KSR1	Feb 94	32	457	1.7
		64	256	3.1
		128	145	5.5
Kendall Square KSR2	Feb 94	32	225	3.5
	May 94	64	130	6.10
Kyoto/Matsushita ADENART	Feb 94	256	314.1	2.5
MasPar MP-1	Aug 92	4K	2396.0	0.33
MasPar MP-2	Nov 92	4K	789.0	1.00
Meiko CS-1	Aug 92	16	2984.0	0.27
Meiko CS-2	Oct 94	8	570.4	1.39
		16	286.6	2.77
		32	149.3	5.31
nCUBE-2S	Mar 94	64	1243.2	0.6
		128	644.7	1.2
	Jul 94	256	336.7	2.35
		512	179.1	4.42
		1024	100.9	7.85
NEC SX-3	Oct 94	1	100.31	7.90
Silicon Graphics	Jul 94	1	1330.3	0.60
Power Challenge XL		4	355.9	2.23
		8	177.0	4.48
		16	91.8	8.63
Silicon Graphics Power Indigo	Oct 94	1	1499.6	0.53
Thinking Machines CM-2	Dec 91	16K	1118.0*	0.71
		32K	634.0*	1.25
		64K	370.0*	2.14
Thinking Machines CM-200	Dec 91	16K	832.0*	0.95
		32K	601.0*	1.32
Thinking Machines CM-5	May 93	32	284.0	2.79
		64	175.0	4.50
		128	119.0	6.66
Thinking Machines CM-5E	Feb 94	32	146.0	5.4
		64	84.0	9.4
		128	48.0	16.5

Table 9a: (cont'd) Results for the Class A BT Simulated CFD Application (* indicates library result).

Computer System	Date Received	No. Proc.	Time (sec.)	Ratio to C90/1
Convex SPP1000	Oct 94	8	1606.0	0.79
Cray C90	Dec 93	1	1261.4	1.00
		4	324.9	3.9
		16	96.4	13.1
Cray T3D	Oct 94	16	918.04	1.37
		32	487.33	2.59
		64	254.82	4.95
		128	132.27	9.54
		256	69.39	18.18
	Jul 94	512	38.01	33.19
	Oct 94	1024	20.45	61.68
Fujitsu VPP500	Oct 94	17	37.26	33.85
		34	18.82	67.02
		51	12.61	100.03
IBM RS6000-590	Mar 94	1	5242.4	0.2
IBM SP-1	Apr 94	16	987.4	1.28
		32	511.2	2.47
		64	274.6	4.59
IBM SP-2	May 94	8	976.91	1.29
		16	498.49	2.53
		32	257.52	4.90
		64	135.98	9.28
	Oct 94	128	75.41	16.73
Intel Paragon (OSF1.2)	Mar 94	102	633.0	2.0
		204	359.0	3.5
		306	257.0	4.9
		408	226.0	5.6
		510	196.0	6.4
Intel Paragon (SunMos)	Mar 94	102	598.0	2.1
		204	324.0	3.9
		306	215.0	5.9
Kendall Square KSR2	May 94	64	542.0	2.33
NEC SX-3	Oct 94	1	399.11	3.16
Silicon Graphics	Jul 94	1	5698.7	0.22
Power Challenge XL		4	1450.0	0.87
		8	775.0	1.63
		16	426.0	2.96
Thinking Machines CM-5E	Feb 94	32	806.0	1.6
		64	464.0	2.7
		128	253.0	5.0

Table 9b: Results for the Class B BT Simulated CFD Application

Computer System	No. PE	Memory	Disk	List Price	Date
Convex SPP1000	16	64 MB/PE		0.94M	Jul 94
Cray C90	16	256 MW (total)		30.90M	Oct 93
Cray EL98	8	1 GB (total)	12 GB	1.11M	Oct 93
Cray T3D	256	16 MB/PE		9.25M	Mar 94
Fujitsu VPP500	16	256 MB/PE		17.0M	Mar 94
IBM SP-1	64	64 MB/PE	64 GB	2.66M	Oct 93
IBM SP-2	64	128 MB/PE (wide)	64 GB	5.94M	Oct 94
IBM RS6000-590	1	1 GB		0.25M	Mar 94
Intel Paragon	256	32 MB/PE		7.49M	Mar 94
Kendall Square KSR1	128	32 MB/PE	25 GB	1.7M	Mar 94
Kendall Square KSR2	32	32 MB/PE	25 GB	1.43M	Mar 94
MasPar MP-2	16K	1 GB (total)		1.61M	Oct 93
nCUBE-2S	1024	4 MB/PE		4.0M	Mar 94
NEC SX-3/14R	1	2 GB		4.1M	Oct 94
SGI Power Challenge	16	2 GB (total)	2 GB	1.02M	Jun 94
Thinking Machines CM-5E	128	32 MB/PE		4.0M	Mar 94

Table 10: U.S. List Price for systems configured as tested.

References

[1] D. H. Bailey, E. Barszcz, J. T. Barton, D. S. Browning, R. L. Carter, L. Dagum, R. A. Fatoohi, P. O. Frederickson, T. A. Lasinski, R. S. Schreiber, H. D. Simon, V. Venkatakrishnan, and S. K. Weeratunga, "The NAS Parallel Benchmarks", *Intl. Journal of Supercomputer Applications*, v. 5, no. 3 (Fall 1991), pp. 63 – 73.

[2] D. Bailey, J. Barton, T. Lasinski, and H. Simon, eds., "The NAS Parallel Benchmarks", NASA Technical Memorandum 103863, Ames Research Center, Moffett Field, CA 94035-1000, July 1993.

[3] D.H. Bailey and P.O. Frederickson, "Performance Results for Two of the NAS Parallel Benchmarks", in Proceedings of Supercomputing '91, Albuquerque NM, pp. 166-173, Nov 18-22, 1991.

[4] E. Barszcz, R. Fatoohi, V. Venkatakrishnan, and S. Weeratunga, "Solution of Regular Sparse Triangular Linear Systems on Vector and Distributed Memory Multiprocessors", Tech Report RNR-93-07, NASA Ames Research Center, Moffett Field, CA 94035, April 1993.

[5] G. Bhanot, K. Jordan, J. Kennedy, J. Richardson, D. Sandee and M. Zagha, "Implementing the NAS Parallel Benchmarks on the CM-2 and CM200 Supercomputers", Thinking Machines Corp, Cambridge, MA 02142.

[6] S. Breit, W. Celmaster, W. Coney, R. Foster, B. Gaiman, G. Montry and C. Selvidge, "The Role of Computational Balance in the Implementation of the NAS parallel Benchmarks on the BBN TC2000 Computer", FED-Vol. 156, CFD Algorithms and Applications, ASME, 1993.

[7] L. Dagum, "Parallel Integer Sorting with Medium and Fine-Scale Parallelism", *International Journal of High Speed Computing*, Vol. 5, No. 4, pp. 503–522, 1993.

[8] M. Fillo, *Architectural Support for Scientific Applications on Multicomputers*, Hartung Gorre Verlag, Series in Microelectronics, Volume 27, Konstanz, Germany, 1993.

[9] B. Hendrickson, R. Leland, and S. Plimpton, "An Efficient Parallel Algorithm for Matrix-Vector Multiplication", Sandia Report SAND92-2765, Sandia National Lab, Albuquerque, NM 87185, March 1993.

[10] J. Lewis, and R. van de Geijn, "Distributed Memory Matrix-vector Multiplication and Conjugate Gradient Algorithms", in Proceedings Supercomputing '93, Portland, OR, Nov. 15-19, 1993.

B'mark	Computer System	No. Proc.	Ratio to C90/1	Nominal cost ($)	Date	Perf. per million $
MG-B	Cray C-90	16	9.5	30.90M	Dec 93	0.31
	Cray T3D	256	10.5	9.25M	Oct 94	1.13
	Fujitsu VPP500	16	16.1	17.00M	Oct 94	0.95
	IBM RS6000-590	1	0.2	0.25M	Mar 94	0.82
	IBM SP-1	64	2.7	2.66M	Mar 94	1.03
	IBM SP-2	64	8.3	5.43M	Aug 94	1.54
	Intel Paragon (OSF1.2)	256	2.8	7.49M	Mar 94	0.37
	NEC SX-3	1	2.9	4.10M	Oct 94	0.70
	Thinking Machines CM-5E	128	5.6	4.00M	Feb 94	1.40
SP-B	Convex SPP1000	16	0.4	0.94M	Oct 94	0.44
	Cray C-90	16	8.9	30.90M	Dec 93	0.29
	Cray T3D	256	9.2	9.25M	Jul 94	1.00
	Fujitsu VPP500	17	13.42	18.06M	Aug 94	0.74
	IBM RS6000-590	1	0.18	0.25M	Mar 94	0.70
	IBM SP-1	64	2.5	2.66M	Feb 94	0.94
	IBM SP-2	64	6.4	5.43M	Aug 94	1.18
	Intel Paragon	256	2.4	7.49M	Jul 94	0.32
	NEC SX-3	1	2.4	4.10M	Oct 94	0.59
	SGI Power Challenge	16	2.3	1.02M	Jul 94	2.25
	Thinking Machines CM-5E	128	2.2	4.00M	Feb 94	0.55

Table 11: Approximate Sustained Performance Per Dollar on Two Class B Benchmarks

[11] V. K. Naik, "Performance Issues in Implementing NAS Parallel Benchmark Applications on IBM SP-1", Research Report, T.J. Watson Research Center, IBM, (in preparation) 1993.

[12] T. H. Pulliam. "Efficient Solution Methods for the Navier-Stokes Equations", Lecture Notes for The Von Karman Institute for Fluid Dynamics Lecture Series, Jan. 20 - 24, 1986.

[13] F. Sukup and J. Fritscher, "Efficiency Evaluation of Some Parallelization Tools on a Workstation Cluster Using the NAS Parallel Benchmarks", Computing Center, Vienna University of Technology, Vienna, Austria.

[14] S. White, "NAS Benchmark on Virtual Parallel Machines", Master's thesis, Emory University, 1993.

[15] S. Yoon, D. Kwak, and L. Chang, "LU-SGS Implicit Algorithm for Implicit Three Dimensional Navier-Stokes Equations with Source Term", AIAA Paper 89-1964-CP, American Institute of Aeronautics and Astronautics, Washington, D.C., 1989.

[16] M. Zagha and G.E. Blelloch, "Radix Sort for Vector Multiprocessors", Proceedings of Supercomputing '90, pp. 712-721, New York, NY, Nov. 1991.