NAS Parallel Benchmark Results

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

RNR Technical Report RNR-92-002
January 10, 1993

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 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.

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 T045-1, Moffett Field, CA 94035.

1 Introduction

The Numerical Aerodynamic Simulation (NAS) Program, which is based at NASA Ames Research Center, is dedicated to advance 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.

At the present time, there are several commercial highly parallel systems available with computing power roughly competitive with conventional supercomputers (even greater on some special problems). Unfortunately, there is little reliable data on the performance of such systems on state-of-the-art computational aerophysics problems. In general, the science of performance evaluation has not kept pace with advances in parallel computer hardware and architecture. There is not even a generally accepted benchmark strategy for highly parallel supercomputers.

In our view, the best benchmarking approach for highly parallel supercomputers is the "paper and pencil" benchmark. The idea is to specify a set of problems only algorithmically. Even the input data must be specified only on paper. Naturally, the problem has to be specified in sufficient detail that a unique solution exists, and the required output has to be brief yet detailed enough to certify that the problem has been solved correctly. But the details of the implementation should be left to the programmer as far as possible.

To this end, we have devised the NAS Parallel Benchmarks (NPB). These 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 [3]. The five kernels are relatively compact problems, each of which emphasizes 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 [2]. The full specification of the benchmarks is given in [3].

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 from the following address: NAS Systems Division, Mail Stop 258-8, NASA Ames Research Center, Moffett Field, CA 94035, attn: NAS Parallel Benchmark Codes. 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.

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 memory requirements, run times and performance ratios. The performance ratios compare individual timings with the current best time on that benchmark achieved on one processor of a Cray Y-MP. The run times in each case are elapsed time of day figures, measured in accordance with the specifications given in [3]. Memory requirements are currently available for only some of these implementations. We hope to have complete information for these columns in future editions of this paper.

Note that performances rates are not cited in millions of floating point operations per second (megaflops) in these tables. We suggest instead that the actual run times (or, equivalently, the performance ratios) be examined when comparing different systems and implementations. For those who wish to compute megaflops figures for the NAS Parallel Benchmarks on any system, we insist that they be computed using the standard floating point operation (flop) counts given in Table 1. Table 1 also contains megaflops rates calculated in this manner for the current fastest implementation on one processor of the Cray Y-MP.

With the exception of the Integer Sort benchmark, these standard flop counts were determined by using the hardware performance monitor on a Cray Y-MP, 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 megaflops rates. We reserve the right to change these standard flop counts in the future if deemed necessary.

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

Benchmark Name	Abbrev-iation	Operation Count	Megaflops on Y-MP/1
Embarrassingly Parallel	EP	2.668×10^{10}	211
Multigrid	MG	3.905×10^{09}	176
Conjugate Gradient	CG	1.508×10^{09}	127
3-D FFT PDE	FT	5.631×10^{09}	196
Integer Sort	IS	7.812×10^{08}	68
LU Simulated CFD Application	LU	6.457×10^{10}	194
SP Simulated CFD Application	SP	1.020×10^{11}	216
BT Simulated CFD Application	BT	1.813×10^{11}	229

Table 1: Standard Operation Counts and Current Y-MP/1 Megaflops Rates

division of the Research Institute for Advanced Computer Science, which is located at NASA Ames; BBN denotes Bolt, Beranek and Newman; Boeing denotes Boeing Computer Services; CRI denotes Cray Research, Inc.; KSR denotes Kendall Square Research Corp., Intel denotes the Supercomputer Systems Division of Intel Corp.; MasPar denotes MasPar Computer Corp.; Meiko denotes Meiko Scientific Corp.; and TMC denotes Thinking Machines, Inc. Where no individual citation is made for a specific model, the results are due to vendor staff.

Unfortunately, the limited space in this report does not permit discussion of the methods used in any of these implementations. However, we have included references to technical papers describing these methods whenever such papers are available. Readers are referred to these documents for full details.

This report includes a number of new results not previously published. Some of the MasPar MP-1 and MP-2 results, the Kendall Square KSR-1 results and the TMC CM-5 results in particular have not previously been disclosed. In quite a few other instances, results are improved from previous listings, reflecting 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 The Embarrassingly Parallel 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.

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. The memory requirement for this benchmark was minimal on all systems.

Intel iPSC/860 results are due to J. Baugh of Intel. CM-2, CM-200 and CM-5 results are due to J. Richardson of TMC. MasPar results are due to J. McDonald of MasPar.

3 The Multigrid 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, although the communication patterns are highly structured (as opposed to the conjugate gradient benchmark).

Results for this benchmark are shown in Table 3. Intel results are due to BCS. CM-2 and CM-200 results are due to J. Richardson at TMC.

4 The Conjugate Gradient 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.

The irregular communication requirement of this benchmark is evidently a challenge for all systems. Results are shown in Table 4. Intel results are due to BCS. CM-2 results are due to J. Richardson of TMC. nCUBE-2 results are by B. Hendrickson, R. Leland, and S. Plimpton of Sandia National Laboratory.

5 The 3-D FFT PDE 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.

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 results are due to E. Kushner of Intel. CM-2 and CM-200 results are due to J. Richardson of TMC.

Computer System	Problem Size	No. Proc.	Memory (mwords)	Time (sec.)	Ratio to Y-MP/1
Y-MP	2^{28}	1	4.9	126.2	1.00
		8	4.9	15.9	7.95
Y-MP EL		1	4.9	550.5	0.23
		4	4.9	141.2	0.89
C-90		1	4.9	47.6	2.65
		4	4.9	12.4	10.20
		16	4.9	3.2	39.56
TC2000		64	1	284.0	0.44
iPSC/860		32	1	102.7	1.23
		64	1	51.4	2.46
		128	1	25.7	4.91
CM-2		8K	1	126.6	1.00
		16K	1	63.9	1.97
		32K	1	33.7	3.74
		64K	1	18.8	6.71
CM-200		8K	1	76.9	1.64
		16K	1	39.2	3.22
		32K	1	20.7	6.10
		64K	1	10.9	11.58
CM-5		16	1	42.4	2.98
		32	1	21.5	5.88
		64	1	10.9	11.62
		128	1	5.4	23.49
		256	1	2.7	46.84
		512	1	1.4	90.47
CS-1		16		116.8	1.08
MP-1		4K		248.0	0.51
		16K		69.3	1.82
MP-2		16K		22.4	5.63
KSR-1		32		69.8	1.81
		64		34.9	3.62
		96		23.4	5.39
		128		18.1	6.97

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

Computer System	Problem Size	No. Proc.	Memory (mwords)	Time (sec.)	Ratio to Y-MP/1
Y-MP	256^3	1	56.7	22.22	1.00
		8	56.7	2.96	7.51
Y-MP EL		1	56.7	89.19	0.25
		4	56.7	32.11	0.69
C-90		1	56.7	8.65	2.57
		4	56.7	2.42	9.18
		16	56.7	0.96	23.14
iPSC/860		128		8.6	2.58
CM-2		16K		45.8	0.49
		32K		26.0	0.85
		64K		14.1	1.58
CM-200		16K		30.2	0.74
		32K		17.2	1.29
CS-1		16		42.8	0.52
MP-1		16K		12.0	1.85
MP-2		16K		4.36	5.10
KSR-1		32		20.6	1.08

Table 3: Results of the Multigrid (MG) Benchmark

Computer System	Problem Size	No. Proc.	Memory (mwords)	Time (sec.)	Ratio to Y-MP/1
Y-MP	2.0×10^6	1	10.4	11.92	1.00
		8	10.4	2.38	5.01
Y-MP EL		1	10.4	65.4	0.18
		4	10.4	23.9	0.50
C-90		1	10.4	4.56	2.61
		4	10.4	1.51	7.89
		16	10.4	0.58	20.55
TC2000		40		51.4	0.23
iPSC/860		128		8.6	1.38
CM-2		8K		25.6	0.47
		16K		14.1	0.85
		32K		8.8	1.35
CM-200		8K		15.0	0.79
CS-1		16		67.5	0.18
MP-1		4K		64.5	0.18
		16K		14.6	0.82
MP-2		16K		11.0	1.08
KSR-1		32		21.7	0.55
nCUBE-2		1024		6.1	1.96

Table 4: Results of the Conjugate Gradient (CG) Benchmark

Computer System	Problem Size	No. Proc.	Memory (mwords)	Time (sec.)	Ratio to Y-MP/1
Y-MP	$256^2 \times 128$	1	42.9	28.77	1.00
		8	42.9	4.19	6.87
Y-MP EL		1	42.9	122.6	0.23
		4	42.9	34.9	0.82
C-90		1	42.9	10.28	2.80
		4	42.9	2.58	11.2
		16	42.9	0.91	31.6
iPSC/860		64		20.9	1.37
		128		9.7	2.96
CM-2		16K		37.0	0.78
		32K		18.2	1.58
		64K		11.4	2.52
CM-200		8K		45.6	0.63
CS-1		16		170.0	0.17
MP-1		16K		18.3	1.57
MP-2		16K		8.0	3.60
KSR-1		32		13.6	2.12
		64		8.4	3.43

Table 5: Results of the 3-D FFT PDE (FT) Benchmark

Computer System	Problem Size	No. Proc.	Memory (mwords)	Time (sec.)	Ratio to Y-MP/1
Y-MP	2^{23}	1	31.1	11.46	1.00
		8	31.1	1.85	6.19
Y-MP EL		1	31.1	153.9	0.07
		4	31.1	41.5	0.28
C-90		1	31.1	5.20	2.20
		4	31.1	1.42	8.07
		16	31.1	0.57	20.10
iPSC/860		32		25.7	0.45
		64		17.3	0.66
		128		13.6	0.84
CM-2		8K		215.1	0.05
		16K		111.5	0.10
		32K		56.0	0.20
CS-1		16		62.7	0.18
MP-1		16K		23.6	0.49
MP-2		16K		18.4	0.62
KSR-1		32		40.2	0.29

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

6 The Integer Sort 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.

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 results are due to to E. Kushner of Intel. CM-2 results are due to L. Dagum of NAS.

7 The Three 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.

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 the third simulated CFD application (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.

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.

Intel and CM-2 results are due to S. Weeratunga, R. Fatoohi, E. Barszcz and V. Venkatakrishnan of NAS, except that BT and SP results on the Intel are due to BCS. CM-5 results are due to J. Richardson of TMC.

8 Other Results

As far as we have been able to determine, the timings presented above all represent runs that fully comply with the rules and restrictions stated in the benchmark document [3]. Two of these rules are that assembly language may not be used, and that assembly-coded library routines may only be used for a restricted set of operations. The allowable exceptions include vendor-supported routines to synchronize processors, communicate data, perform array transpositions, evaluate Fortran intrinsic functions, perform dense matrix multiplication and compute fast Fourier transforms.

There are several reasons for these restrictions on assembly code. For one thing, without restrictions of some sort, an entire benchmark might be implemented in assembly-level code. While such performance results might be interesting, they would hardly be indicative of the performance that a scientist could reasonably expect on a full-scale application program. In other words, the tuning rules for the NPB reflect our expectation (and experience) that real scientific applications consist largely of Fortran or C code, and that usage of library routines is restricted to a handful of widely available functions.

Nonetheless, some scientists have attempted implementations of the NPB using library routines that do not comply with the official rules. In particular, Thinking Machines, Inc. has obtained performance results using assembly-coded library routines to perform key computations in several of the NPB. These results are shown in Table 10 [4]. Their

Computer System	Problem Size	No. Proc.	Memory (mwords)	Time (sec.)	Ratio to Y-MP/1
Y-MP	64^3	1	32.3	333.5	1.00
		8	32.3	49.5	6.74
Y-MP EL		1	32.3	1449.0	0.23
		4	32.3	522.3	0.64
C-90		1	32.3	157.6	2.12
		4	32.3	43.9	7.59
		16	32.3	17.6	18.93
TC2000		62		3032.0	0.11
iPSC/860		64	12	690.8	0.48
		128	16	442.5	0.75
CM-2		8K	14	1307.0	0.26
		16K	14	850.0	0.39
		32K	14	546.0	0.61
CM-5		64		336.0	0.99
CS-1		16		2937.0	0.11
MP-1		4K		1785.0	0.19
MP-2		4K		562.0	0.59
KSR-1		32		1041.3	0.32

Table 7: Results for the LU Simulated CFD Application

Computer System	Problem Size	No. Proc.	Memory (mwords)	Time (sec.)	Ratio to Y-MP/1
Y-MP	64^3	1	9.2	471.5	1.00
		8	9.2	64.6	7.30
Y-MP EL		1	9.2	2026.0	0.23
		4	9.2	601.9	0.78
C-90		1	9.2	184.70	2.55
		4	9.2	49.74	9.48
		16	9.2	13.06	36.10
TC2000		112		880.0	0.54
iPSC/860		64		667.3	0.71
		128		449.5	1.05
CM-2		8K		3900	0.12
		16K		2104	0.22
		32K		1080	0.44
CM-5		64		280.0	1.68
CS-1		16		2975	0.16
MP-1		4K		1772	0.27
MP-2		4K		657	0.72
KSR-1		32		377.7	1.25
		64		228.8	2.06
		96		170.2	2.77
		128		150.0	3.14

Table 8: Results for the SP Simulated CFD Application

Computer System	Problem Size	No. Proc.	Memory (mwords)	Time (sec.)	Ratio to Y-MP/1
Y-MP	64^3	1	42.3	792.4	1.00
		8	42.3	114.0	6.95
Y-MP EL		1	42.3	4033	0.20
		4	42.3	1208	0.66
C-90		1	42.3	356.9	2.22
		4	42.3	96.10	8.25
		16	42.3	28.39	27.91
TC2000		112		1378	0.58
iPSC/860		64		714.7	1.11
		128		414.3	1.91
CM-2		16K		3328	0.24
		32K		1914	0.41
CM-5		64		240.0	3.30
CS-1		16		2984	0.27
MP-1		4K		2396	0.33
MP-2		4K		803	0.99
KSR-1		32		439.0	1.81
		64		239.4	3.31
		96		167.9	4.72
		128		134.5	5.89

Table 9: Results for the BT Simulated CFD Application

Benchmark	Computer System	No. Proc.	Time (sec.)	Ratio to Y-MP/1	Ratio to CM-x without library
IS	CM-2	16K	35.8	0.32	3.11
		32K	21.0	0.55	2.67
		64K	14.9	0.77	
	CM-200	64K	5.7	2.01	
SP	CM-2	16K	1444	0.33	1.46
		32K	917.0	0.51	1.18
		64K	640.0	0.74	
	CM-5	64	180.0	2.62	1.56
BT	CM-2	16K	1118	0.71	2.98
		32K	634.0	1.25	3.01
		64K	370.0	2.14	
	CM-200	16K	832.0	0.95	
		32K	601.0	1.32	
	CM-5	64	176.0	4.50	1.36

Table 10: TMC Results Using Library Routines

implementation of the IS benchmark on the CM-2, for example, runs more than twice as fast as reported in Table 6, and their CM-2 rates for the BT benchmark are nearly three times as fast as reported in Table 9. The CM-5 results in this table are based on implementations that utilize conventional CM Fortran routines for computation, but which utilize a low-level array transposition routine not yet officially supported by TMC. As Table 10 indicates, using this routine resulted in roughly a 1.5 speedup in BT and SP over straight CM Fortran.

In addition to library results from TMC, other results include a comparison of the relative performance between the Intel Gamma iPSC/860 and the Intel Delta prototype. The Delta machine employs the same i860 microprocessor as the Gamma but has a two dimensional grid topology and a faster router. Since the two dimensional topology results in a lower connectivity than is available with the hypercube topology of the Gamma, it is of interest to measure the extent to which the faster router can make up for the additional distance messages may have to travel. Some of the benchmark codes written by the NAS Applied Research branch were ported directly to the Delta machine. Table 11 presents the relative performance between the two machines. Note that we do not present the absolute performance on the Delta, because the codes used are not the fully optimized codes used for the results in Tables 2-8. Furthermore, note that these codes were not tuned for the mesh topology of the Delta machine, and there may be a significant performance improvement if this were done.

Benchmark	No. Proc.	ratio to Gamma
EP	64	1.03
IS	64	1.36
	128	1.09
MG	32	0.70
CG	128	1.05
LU	16	0.96

Table 11: Performance of Intel Delta Compared to Gamma Machine (dated 5/92).

9 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. We should not be too surprised that the Cray C-90 system, for example, exhibits superior performance rates on these benchmarks, since its current purchase price is much higher than that of the iPSC/860 and the CM-2.

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 purchase price in millions. Some figures of this type are shown in Table 12 for two of the benchmarks, the EP benchmark and the SP benchmark and for ten different systems. The table includes the nominal cost of the various systems. In arriving at the nominal cost figure we have assumed a complete system with a typical set of peripherals. These figures are very approximate and may not necessarily reflect current prices; typically old system prices drop substantially when a vendor introduces a new model. Therefore, because of the approximate and changeable nature of these prices, and because the memory sizes, disk capacities and I/O performances of these systems are certainly not equivalent, the figures in the last column of Table 12 should be interpreted as only very rough indications of sustained performance per dollar.

10 Conclusions

With some algorithmic experimentation and tuning, respectable NPB performance rates have been achieved on several multiprocessor systems. Except for the EP benchmark, the following conclusions hold: (1) the 16 processor Cray C-90 system is the highest performing system tested; (2) the Intel 128 processor iPSC/860 system, the 32K CM-2 system, the 64 node CM-5 system and the 16K MasPar MP-2 system appear to be equivalent to roughly one (although in some cases significantly more than one) Y-MP processor; and (3) when sustained performance rates are normalized by system prices, the CM-5 and the two MasPar

Benchmark	Computer System	No. Proc.	Ratio to Y-MP/1	Nominal cost ($)	Perf. per million $
EP	C-90	16	39.56	36M	1.10
	Y-MP	8	7.95	15M	0.53
	Y-MP EL	4	0.89	1.5M	0.59
	iPSC/860	128	4.91	3M	1.64
	CM-2	32K	3.74	5M	0.75
	CM-5	512	90.47	15M	6.03
	MP-1	16K	1.82	0.5M	3.64
	MP-2	16K	5.63	1M	5.63
	CS-1	16	1.08	0.3M	3.60
	KSR-1	128	6.97	6M	1.16
SP	C-90	16	36.10	36M	1.00
	Y-MP	8	7.30	15M	0.49
	Y-MP EL	4	0.78	1.5M	0.52
	iPSC/860	128	1.05	3M	0.35
	CM-2	32K	0.44	5M	0.09
	CM-5	64	2.62	2.5M	1.05
	MP-1	4K	0.27	0.2M	1.35
	MP-2	4K	0.72	0.3M	2.40
	CS-1	16	0.16	0.3M	0.53
	KSR-1	128	3.14	6M	0.52

Table 12: Approximate Sustained Performance Per Dollar

systems appear to deliver somewhat more performance per dollar than the C-90.

Of the latest generation of parallel computers, only Kendall Square has supplied us with a complete set of benchmark numbers. Thinking Machines, Inc. and MasPar have supplied us with a limited set of results. These results are rather encouraging, as some systems exhibit performance exceeding that of a Y-MP/1, even on the challenging SP and BT benchmarks. It is also reasonable to expect even better performance from these machines as the compilers mature and the implementations receive further tuning.

Some scientists have suggested that the answer to obtaining high performance rates on highly parallel computers is to substitute alternative algorithms that have lower interprocessor communication requirements. However, it has been the experience of the scientists in our research group that a certain amount of long-distance communication is unavoidable for these types of applications. Alternative algorithms that have higher computation rates usually require more iterations to converge to a solution and thus require more overall run time. Clearly it is pointless to employ numerically inefficient algorithms merely to exhibit artificially high performance rates on a particular parallel architecture [1].

References

[1] D. H. Bailey, "Twelve Ways to Fool the Masses When Giving Performance Results on Parallel Computers", *Supercomputing Review*, August 1991, p. 54 – 55. Also published in *Supercomputer*, September 1991, p. 4 – 7.

[2] 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.

[3] D. Bailey, J. Barton, T. Lasinski, and H. Simon, eds., "The NAS Parallel Benchmarks", Technical Report RNR-91-02, NASA Ames Research Center, Moffett Field, CA 94035, January 1991.

[4] 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.

[5] 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", submitted to *Concurrency*, April 1991.