

LEGIBILITY NOTICE

A major purpose of the Technical Information Center is to provide the broadest dissemination possible of information contained in DOE's Research and Development Reports to business, industry, the academic community, and federal, state and local governments.

Although a small portion of this report is not reproducible, it is being made available to expedite the availability of information on the research discussed herein.

Handwritten signature or initials

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

TITLE: A BENCHMARK OF THE SCS-40 COMPUTER: A MINI SUPERCOMPUTER
COMPATIBLE WITH THE CRAY X-MP/24

LA-UR--87-659

AUTHOR(S): Harvey J. Wasserman,
Margaret L. Simmons
Ann H. Hayes

DE88 009158

SUBMITTED TO Vector and Parallel Processors in Computational Science Conference
Liverpool, England
August 25-28, 1987

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

MASTER

DISTRIBUTION

By acceptance of this article the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution or to allow others to do so for U.S. Government purposes.

The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

Los Alamos Los Alamos National Laboratory
Los Alamos, New Mexico 87545



(REVISED 5-1-87)

A BENCHMARK OF THE SCS-40 COMPUTER: A MINI SUPERCOMPUTER COMPATIBLE WITH THE CRAY X-MP/24

*Harvey J. Wasserman
Margaret L. Simmons
Ann H. Hayes*

Computer and Communications Division
Los Alamos National Laboratory
Los Alamos, New Mexico 87545

ABSTRACT

An accurate benchmark of the SCS-40 mini supercomputer manufactured by Scientific Computer Systems Corporation has been carried out. A new, revised set of standard ANSI77 Fortran benchmark codes were run on the SCS-40 in a dedicated environment, using Version 1.13 of the CFT compiler. The results are compared with those obtained on one processor of a CRAY X-MP/24 computer using the Cray Research Inc. version of the same compiler. The results suggest that for a typical Los Alamos National Laboratory computational workload, the SCS-40 is equivalent to one-quarter to one-third of a single processor of the CRAY X-MP/24.

1. INTRODUCTION

Recently there has been much interest in computers of the "mini-super" class.¹⁻³ Machines in this category are intended to "bridge the gap" between supercomputers such as those made by Cray Research Inc., and the slower and less expensive minicomputers, such as the Digital Equipment Corporation (DEC) VAXs. The Scientific Computer Systems Corporation (SCS) has announced a mini supercomputer, the SCS-40, that uses an instruction set that is essentially a duplicate of that found in the CRAY X-MP/2 series. Yet, by using "off the shelf" technology, SCS is able to price their machine at considerably less than a CRAY X-MP.

We have already reported a detailed performance evaluation of the CRAY X-MP/24.⁴ We have now benchmarked the SCS-40 to determine precisely what level of performance can be expected from this machine. The benchmark was carried out in the usual Los Alamos manner:^{5, 10} a set of portable Fortran codes, representing exclusively the LANL computational workload, was run on the SCS-40 in a dedicated environment. No throughput or I/O measurements were made. The results of the benchmark appear below.

2. SCS-40 ARCHITECTURE

The system architecture of the SCS-40¹¹ mimics one processor of a CRAY X-MP/2.¹² The central processor, built with commercially available emitter coupled logic technology, operates with a clock period (CP) of 45 nanoseconds (ns). The control section of the processor consists of five microcoded pipelined segments, including a 256-word instruction buffer with a 10-ns access time. The instruction buffer can be filled at the rate of one word per

cycle. The full complement of X-MP registers, namely the A, S, B, T, vector, vector merge, and vector length, has been implemented, as well as three sets of registers used in the X-MP for sharing among processors. The A and B registers do 24-bit address calculations. Functional unit organization is the same as on the X-MP, except that only a single shared shift unit is provided rather than separate vector and scalar shift units. Flexible hardware chaining within the vector registers is possible. Also, SCS has implemented the bidirectional memory feature; this allows block loads or stores between the B, T, or V registers and memory to begin before the completion of a prior block load or store. (The X-MP/2 only allows overlapping of block loads with block stores.) Block loads from memory to the vector registers may be "strided" with use of an offset kept in an A register. Table I compares some representative operation times for the SCS-40 and the CRAY X-MP.

The main memory of the SCS-40 is capable of containing either one, two, or four million 64-bit words. The memory consists of 16 interleaved tiers, has an access time of five CP (225 ns), and operates on a five CP cycle time. In comparison, the CRAY X-MP/24 is interleaved 32 ways, has an access time of 14 CP (133 ns), and has a 4-CP (38-ns) cycle time. The SCS-40 memory is implemented on 256K CMOS chips and employs a SECDED error correction/detection scheme on eight check bits. The machine we measured in this benchmark was configured with four million words of memory.

In the SCS-40, main memory, the register, functional unit, and control sections of the CPU, and the I/O subsystem are interconnected via two sets of 64-bit-wide busses. The functional unit and memory data busses operate at a virtual rate equal to one-half the processor CP, or 22.5 ns. There are three bidirectional data busses connecting main memory to the register and I/O sections with a total maximum capacity of 133 Mword/s (six words per clock). However, actual transfer rates are limited by the four memory control busses, each of which is capable of conveying one write or read request per CP. Thus, the average number of words actually transferred to or from memory per CP is four. The three memory data busses are allocated such that one is dedicated to vector read/writes, one is a shared vector-I/O bus, and the third is reserved for A, S, B, and T register loads.

The functional units and registers are joined by two pairs of busses. Each pair can transmit an operand from a V register to the functional units in one-half of a CP. Additionally, one bus in each pair can transmit a result back from the functional units to the registers in another one-half of a CP. Both pairs of physical busses support the logical operand and result busses, thus enabling the vector functional units to produce one result from two vector operands every CP.

The SCS-40 I/O subsystem may be configured with from one to four I/O modules, each of which contains two I/O channels. The maximum throughput can reach 22.2 Mword/s. The I/O processors, one per channel, communicate with external devices using a 16-bit channel word and are firmware programmed for specific interfaces. Either a DEC VAX-11/750 or 11/780 is used as a front-end. The machine we benchmarked was connected to four DD-680 disk units.

3. SCS-40 SOFTWARE

SCS software products have relied heavily on two important factors: (1) Cray Research placed its Version 1.13 software in the public domain, and (2) much software development for the Cray computers has been carried out at Department of Energy (DOE) institutions and the resulting software is also in the public domain.

	Floating Point Add	Floating Point Multiply	Reciprocal Approximation	Vector Population	Address Multiply
SCS-40	3 CP 135 ns	3 CP 135 ns	6 CP 270 ns	2 CP 90 ns	2 CP 90 ns
X-MP/24	6 CP 57 ns	7 CP 66.5 ns	14 CP 133 ns	5 CP 47.5 ns	4 CP 38 ns

The SCS-40 currently runs the Cray Timesharing System (CTSS) as its operating system. CTSS was developed at the Lawrence Livermore National Laboratory (LLNL) and is based on a similar system used on the Control Data Corporation (CDC) 7600 computers. CTSS is the primary operating system on Cray machines at such DOE sites as LANL, LLNL, and the Livermore National Magnetic Fusion Energy Computational Center. Two of the most important features of CTSS from the user standpoint are the notions of drop files and suffixes, both of which are supported in full on the SCS-40.

SCS currently supports Version 1.13 of the Cray Fortran Compiler (CFT), as well as its associated assembler, CAL. The command line for these products is the same as it is on the LANL CTSS system (the command line is slightly different under Cray's COS operating system). Both the Baselib/Fortlib and CFTLIB Fortran libraries are available; all of the results below were obtained using CFTLIB. There were two reasons for choosing CFTLIB: (1) the syntax for disk file I/O initialization conforms more with the ANSI 77 standard ('OPEN' vs 'CALL LINK' in Fortlib), and (2) SCS personnel informed us that codes loaded with CFTLIB seemed to run about 10% faster than those loaded with Fortlib/Baselib. The only change our codes required was the addition of a statement to create a drop file. This is due to a temporary bug in SCS CFTLIB that prevents automatic interpretation of the PROGRAM statement.

Other important software we used included the debugger DDT; the editors TRIXGL and TEDI; the COSMOS job controller, MOVE, a utility to ship files from the YAX front-end to the SCS-40 worker; and other utilities and control key entries. The control key feature has been dramatically improved over that available at LANL; some 58 control key commands are supported, including several that allow for a split screen.

4. BENCHMARK RESULTS

Members of the LANL Computer Research and Applications Benchmark Team tested the SCS-40 on December 17, 1986. New versions of the standard Los Alamos benchmark set were run. A detailed description of the codes appears in Appendix A. In addition to this standard set of benchmarks, another set of codes, representing critical computational work currently performed at the Laboratory, and presently being developed into benchmark codes, was also run.

All the codes measured CPU times with a call to the CFTLIB routine SECOND. This routine returns the CTSS charge time for the run (in this case 95% of actual CPU time); a correction factor of 1.05 has thus been applied to all the raw data.

4.1. Standard Benchmarks

Timing data (in seconds) for the standard benchmark set are listed in Table II. For comparative purposes, timings from a single processor of a CRAY X-MP/24 (using CFT 1.13) are also given. The column labeled "options" includes results obtained with the CFT compiler option "opt = btreg" as well as a call "CALL EBM" to enable bidirectional access to memory. The column labeled "no options" actually means that compilation was performed using those options that are included by default with the CFT compiler.

The data show that for a typical Los Alamos workload, the SCS-40 is equivalent to about 20-30% of a single processor of the CRAY X-MP/24. The results do not suggest any particular pattern with respect to the level of vectorization of the benchmarks. For example, the codes LSS and BMK11A both perform at about the same level on the SCS-40 (about 30% of the X-MP/24 (single processor)), although LSS is nearly 100% vectorizable and BMK11A is only about 50% vectorizable. Note that BMK11A would benefit significantly from scatter/gather operations (neither the SCS-40 nor the CRAY X-MP/24 provide this feature); BMK11A runs for 5.0 s on the CRAY X-MP/48 using CTSS and CFT1.14.

Two codes in our benchmark suite attempt to measure performance of basic vector operations as a function of vector length under a variety of memory access conditions such as contiguous load/stores, constant strides, and random gather/scatters. Table III reports the results of program VECOPS, which measures MFLOP rates for contiguously stored vectors. For all the operations, the performance of the SCS-40 for vector length 1000 is about 20% of the performance we measure on one processor of the X-MP/24. An exception to this occurs for the operation $V = V + S * V$, the fifth entry in Table III. For this operation, at vector length 1000, the SCS-40 provides nearly 50% of the performance of the X-MP/24 (single processor). We believe that the reason for this poor performance on the X-MP/24 (single processor) is the result of a compiler bug--extremely inefficient code is produced by Cray's CFT 1.13. The bug has apparently been fixed in the SCS version of the same compiler. The same level of performance relative to the X-MP/24 (single processor) is observed for short vectors on the SCS-40. Use of the BTREG option along with bidirectional memory increases the MFLOP rate for vector length 1000 by almost a factor of 2 on the SCS-40 (not shown). The gather/scatter operations, the last four shown in Tables III and IV, do not vectorize on the SCS-40 or on the X-MP/2, and so very slow rates are observed; note that these operations do vectorize on the X-MP/4 series. MFLOP rates for vectors accessed with a variety of strides are given in Table IV. As with the X-MP/24, there is no significant degradation in performance when using strides rather than contiguous vector accesses.

Table II. Benchmark Execution Times (in Seconds) for the Standard Los Alamos Benchmarks

Program Name	SCS-40		CRAY X-MP/24 (Single Processor)	RATIO (X-MP/SCS-40)
	No Options	Options		
FFT	23.3	21.5	5.2	0.22
MATRIX	207.4	167.4	60.0	0.29
LSS	35.3	28.5	10.6	0.30
GAMTEB	24.2	21.4	6.9	0.29
SCALGAM	410.8	580.4	116.0	0.28
INTMC	227.1	219.6	56.1	0.25
BMK11A	42.8	39.4	12.8	0.30
BMK11B	12.3	12.1	2.7	0.22
VECOPS	a	a	.	.
VECSKIP	b	b	.	.

^aSee Table III.

^bSee Table IV.

Table III. SCS-40 MFLOP Rates as a Function of Vector Length for a Series of One Million Vector Operations Performed on Vectors Stored in Contiguous Locations (Program VECOPS)

Vector Length	10	25	50	100	200	500	1000
Operation							
$V = V + S$	3.6	9.0	15.3	17.3	17.8	19.9	20.2
$V = S * V$	3.7	9.2	15.3	17.3	17.8	19.9	20.2
$V = V + V$	3.5	8.7	13.8	15.5	15.9	18.2	18.2
$V = V * V$	3.5	8.7	13.8	15.5	16.0	18.2	18.2
$V = V + S * V$	6.6	16.4	27.5	30.7	31.9	36.4	36.4
$V = V * V + S$	6.5	16.4	27.7	20.5	31.8	36.4	36.7
$V = V * V + V$	6.1	12.9	17.7	20.3	22.4	25.7	26.4
$V = S * V + S * V$	8.8	8.8	27.3	29.0	29.6	31.5	31.7
$V = V * V + V * V$	8.4	7.7	24.6	25.8	26.2	28.0	27.3
$V = V(I) + S$	0.9	1.0	1.0	1.0	0.9	1.0	1.1
$V(I) = V * V$	1.0	1.1	1.2	1.2	1.2	1.2	1.2
$V(I) = V(I) + V * V$	1.9	2.1	2.1	2.2	2.2	2.2	2.2

Table IV. SCS-40 MFLOP Rates as a Function of Vector Length for a Series of One Million Vector Operations Performed on Vectors Accessed with a Stride of 23 (Program VECSKIP)

Vector Length	10	25	50	100	200	500	1000
Operation							
$V = V + S$	2.6	6.4	10.9	13.8	15.3	18.2	18.6
$V = S * V$	2.6	6.5	12.4	14.6	16.1	18.2	18.7
$V = V + V$	2.5	6.3	10.9	13.7	14.6	17.3	17.4
$V = V * V$	2.5	5.2	11.9	14.6	14.6	17.7	17.4
$V = V + S * V$	4.8	12.0	20.4	25.6	29.1	34.3	35.0
$V = V * V + S$	4.8	12.0	20.0	25.5	29.1	34.3	35.1
$V = V * V + V$	4.7	10.2	14.6	18.1	20.3	24.3	24.9
$V = S * V + S * V$	7.0	14.0	21.4	25.0	27.2	30.1	30.7
$V = V * V + V * V$	6.6	14.2	20.6	23.1	24.6	27.3	27.3
$V = V(I) + S$	0.8	0.9	1.0	1.0	1.0	1.0	1.0
$V(I) = V * V$	1.0	1.2	1.2	1.3	1.3	1.3	1.3
$V(I) = V(I) + V * V$	1.6	1.9	2.0	2.0	2.1	2.1	2.1
$V = V + V * V(I)$	1.7	2.2	2.3	2.4	2.4	2.5	2.5

We also ran three of the benchmarks, FFT, MATRIX, and BMK11A, on both the CRAY X-MP/24 (single processor) and the SCS-40 in scalar mode, using the "off=v" compiler option (these results are not shown in the tables). The results of this test suggest that the improvement of vector over scalar performance on the SCS-40 is about the same as it is on the X-MP/24.

It is also instructive to compare the performance of the SCS-40 with another computer marketed in the mini-super class: the Convex C-1.¹³ Consider, for example, the code BMK11A. The run time for this code on the C-1 is 103.1 s, compared with 39.4 s on the SCS-40. (The Convex C-1 benchmark was carried out in July, 1985 and most likely does not represent current performance of this machine.) The ratio of the cycle times of the two machines is 2.2 (the C-1 cycle time is 100 ns), however, the SCS-40 performs about 20% better on BMK11A than the ratio of

cycle times would indicate. The difference is probably because 64-bit computations are more efficient on the SCS-40 than they are on the C-1; the corresponding time for BMK11A in 32-bit mode on the C-1 is 77.2 s. Regrettably, further comparison between the two machines cannot be made at this time, because the same set of benchmarks was not run on both machines. Another benchmark of the Convex C-1 is planned.

0.1. Miscellaneous Benchmarks

Table V presents the execution times for some additional codes not presently included in the standard benchmark set.

Hydro is a two-dimensional Lagrangian hydrodynamics code representative of codes that are a significant portion of the Laboratory workload. Two problem sizes were run on the SCS-40. The first employed a 6-by-6 grid for 4 timesteps and is called HYDRO6; the second calculated over a 50-by-50 grid for 50 timesteps and is called HYDRO50. When compiling Hydro with the CFT compiler it is necessary to use the MAXBLOCK = 1350 option to ensure maximum vectorization. On the larger Hydro problem the SCS-40 runs at 35% of the CRAY X-MP/24 (single processor) speed, the largest such ratio observed for our benchmarks. This is probably because more time has been spent optimizing Hydro for the Cray computers than has been spent on the other codes we ran.

MCNP500 is a Monte Carlo neutron photon transport code¹⁴ that accounts for a significant amount of the Laboratory's production computing time. The code consists of about 40000 source lines and is essentially nonvectorizable. The current problem involves 500 source particles, a much smaller computation than a typical production run. The SCS-40 ran MCNP500 in 61.8 s; in other words, the SCS-40 achieved 29% of the CRAY X-MP/24 (one processor) performance.

ESN is another almost entirely scalar code, one that simulates deterministic particle transport.¹⁵ On this code the SCS-40 also achieves about 29% of the one processor CRAY X-MP/24 performance.

Table V. Execution Times (in Seconds) for the Miscellaneous Los Alamos Benchmarks

Program Name	SCS-40		CRAY X-MP/24 (Single Processor)	RATIO (X-MP/SCS-40)
	No Options	Options	No Options	
MCNP500	62.5	61.8	19.1	0.29
HYDRO6	0.1	0.1	-	-
HYDRO50	36.7	-	12.7	0.35
ESN	72.8	68.6	21.0	0.29

5. CONCLUSIONS

The SCS-40 provides about one-quarter to one-third of the performance of a single processor of the CRAY X-MP/24, the supercomputer it was designed to emulate. On all of our benchmark codes, the ratio of the SCS-40 execution times to X-MP/24 execution times is greater than the ratio of the CPU cycle times for the two machines. The SCS-40 derives some of this additional speed from more efficient functional units (see Table I above) and possibly from some improvements to the CFT 1.13 compiler.

We have chosen to compare the speed of the SCS-40 with that of the X-MP/24 because of their instruction set compatibility. However, the X-MP/24 is no longer Cray's premier machine, either in terms of hardware features (it lacks hardware scatter/gather) or hardware speed (the new X-MP/416 has a CP of 8.5 ns). Some of our benchmark codes run as much as a factor of 2 times faster on a single processor of a CRAY X-MP/48 than they do on the X-MP/24, largely because of the scatter/gather feature. It would be interesting to observe the performance of these codes on an SCS machine with hardware scatter/gather.

However, the SCS-40 need not be viewed solely in terms of its performance relative to, and compatibility with, the Cray computers. It is a well-designed mini supercomputer in its own right, providing more computational power on our benchmarks than other machines in its class.¹³

Of course, our benchmark codes are intended to represent the computational workload at LANL and caution should be used in comparing these results with those based on other workloads.

References

1. S. Ohr, "Computer Scores with Cray Compatibility," *Electronic Design*, pp. 61-62, March 6, 1986.
2. J.R. Lineback, "CMOS Gates Key to 'Affordable' Supercomputer," *Electronics Week*, pp. 17-20, October 29, 1984.
3. S. Lackey, J. Veres, and M. Ziegler, "Supercomputer Expands Parallel Processing Options," *Computer Design*, pp. 76-81, August 15, 1985.
4. O. Lubeck, J. Moore, and R. Mendez, "A Benchmark Comparison of Three Supercomputers: Fujitsu VP-200, Hitachi S810/20, and Cray X-MP/2," *IEEE Computer Magazine*, pp. 10-23, December 1985.
5. M.L. Simmons and H.J. Wasserman, "Los Alamos National Laboratory Computer Benchmarking 1986," *Los Alamos National Laboratory report LA-10898-MS*, January 1987.
6. J.H. Griffin and M.L. Simmons, "Los Alamos National Laboratory Computer Benchmarking 1983," *Los Alamos National Laboratory report LA-10151-MS*, June 1984.
7. J.L. Martin, "Los Alamos National Laboratory Computer Benchmarking 1982," *Los Alamos National Laboratory report LA-9698-MS*, June 1983.
8. J.L. Martin and R.G. Martinez, "Los Alamos National Laboratory Computer Benchmarking 1981," *Los Alamos National Laboratory report LA-9104-MS*, November 1981.
9. I.Y. Bucher and A.H. Hayes, "Los Alamos National Laboratory Computer Benchmarking 1979," *Los Alamos National Laboratory report LA-8689-MS*, February 1981.
10. R.G. Brickner, H.J. Wasserman, A.H. Hayes, and J.W. Moore, "Benchmarking the IBM 3090 with Vector Facility," *Los Alamos National Laboratory unclassified release LA-UR-86-3300*, September 1986.
11. Scientific Computer Systems Corporation, *SCS-40 Architecture Reference Manual*, SCS Document Number 950000001-001A, 1986.
12. Cray Research Incorporated, *CRAY X-MP/48 Mainframe Reference Manual*, Cray Publication Number HR-0097, 1986.
13. M. Simmons and O. Lubeck, "Benchmark of the Convex C-1 Mini Supercomputer," *Los Alamos National Laboratory unclassified release LA-UR-86-2890*, August 1986.
14. J. Briesmeister, Ed., "MCNP: A General Monte Carlo Code for Neutron and Photon Transport," *Los Alamos National Laboratory unclassified release LA-7396-M Rev 2*, September 1986.
15. B.R. Wienke, "ESN: One-Dimensional S₁ Transport Module for Electrons," *Journal of Quantitative Spectroscopy and Radiation Transport*, vol. 28, p. 311, 1982.

ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy. We wish to thank Jim Taft, Paul Robertson, Arlen Petty, and Hanan Potash of Scientific Computer Systems Corporation for their assistance during all phases of the benchmark.

APPENDIX A DESCRIPTION OF THE NEW BENCHMARK PROGRAMS

The Computing and Communications Division at Los Alamos has maintained a set of portable benchmark programs representing characteristic tasks that a new large computer would be required to run at the Laboratory. A database exists containing results of past runs of these programs on a large variety of computers. Recently, a decision was made to update this suite of benchmarks for the following reasons:

1. Some of the codes had problem sizes that were too small; on some supercomputers the timings would not allow meaningful comparisons.
2. Some of the codes implemented algorithms that were no longer deemed to be of importance to the Laboratory.
3. All of the codes were updated to be more in keeping with the ANSI77 Fortran standard.
4. All of the names of the codes were changed from a simple numerical designation to one that more closely represents the type of work the program performs.

Some old benchmarks were eliminated entirely. Table A-1 contains a list of the new benchmarks, a cross-reference to the name of the corresponding old benchmark, a brief description of the code and any changes that were made, and the ratio of the run times of the two versions on the CRAY X-MP/48.

Table A-I. The New Los Alamos Standard Benchmark Set

New Name	Old Name	Description	Ratio(new:old)
INTMC	BMK1	Integer Monte Carlo with almost no floating-point arithmetic; does not vectorize; measures primarily speed of integer arithmetic; no I/O involved--all data internally generated.	1.0
FFT	BMK4A	Highly vectorizable fast Fourier transform (FFT); measures the speed of 512 transforms; involves many short vectors and is therefore sensitive to vector startup times. FFT library routines supplied by many computer manufacturers generally perform multiple FFTs much more efficiently than this code; no I/O.	1.0
MATRIX	BMK14	Basic matrix operations, including multiplication and transpose on matrices of order 100; highly vectorizable but not optimized for vector computers; the problem size has been increased over BMK14.	29.9
LSS	BMK22	Linear system solver from LINPACK for systems of equations of order 100; uses Gaussian elimination; highly vectorizable but not optimized for vector computers.	1.0
GAMTEB	BMK21A	Nonvectorizable Monte Carlo photon transport code.	1.0
SCALGAM	BMK21b	Monte Carlo photon transport code using binary tree random number sequence.	1.0
VECOPS	BMK8A1	Tests rates of elementary vector operations as a function of vector length with vectors stored in contiguous locations; typically one million floating-point operations are timed for all vector lengths.	1.0
VECSKIP	BMK8A2	Performs same operations as VECOPS code but with vectors stored in noncontiguous memory locations; several values of stride are used to determine when memory bank conflicts occur.	1.0