.
.

RISC Family
Performance
Summary

Digital Equipment Corporation

Table of Contents

Table of Contents	Page
Introduction	**3**
Benchmarks by Application Environment	**3**
Methodology	**4**
Summary of Findings	**5**
SPEC Positioning	**6**
SPEC Performance	**7**
Mechanical Design: ANSYS Finite Element Analysis	**9**
Scientific Computing: NCAR	**11**
Scientific Computing: DR Labs CPU2	**13**
Chemical Engineering: MDATOM	**15**
Geo/Process Engineering: TOP	**17**
Geo/Process Engineering: UTCHEM	**19**
Floating Point Benchmarks	**21**
Linpack	**21**
Whetstone	**24**
Lawrence Livermore FORTRAN Kernels	**26**
Integer Benchmarks	**29**
Dhrystone	**29**
Test Configurations	**31**
Raw Data Not Previously Supplied	**33**
References	**35**

Introduction

The purpose of this document is:

1. To describe the performance of the Digital RISC family of systems for Engineering and Scientific applications, including the SPEC[1] benchmarks.
2. To position the SPEC benchmarks relative to other benchmarks.

This document does not cover the performance of multi-user, client/server or commercial (e.g. Transaction Processing) environments.

Benchmarks by Application Environment

Results from a large selection of benchmarks are needed to understand system performance. Understanding the nature of your application environment and matching this information with benchmark results is crucial for accurate performance prediction. The table below categorizes the benchmarks used in this document by market segment and application type.

Application Environment	Benchmark	Application Type
COMPUTER AIDED SOFTWARE ENGINEERING	GCC Li	GNU "C" Compiler Lisp Interpreter
ELECTRONIC DESIGN	Espresso Eqntott Spice 2g6	PLA Simulator Boolean Logic Simulation Analog Circuit Simulation
MECHANICAL DESIGN	ANSYS	Finite Element Analysis
GEO/PROCESS ENGINEERING	TOP UTCHEM	Gridding and Contour Mapping Reservoir Simulation
SCIENCE	DR Labs CPU2 Doduc Fpppp Matrix300 MDATOM Nasa7 NCAR Tomcatv	Scientific Kernels High Energy Physics Quantum Chemistry Matrix Multiplication Chemical Engineering Aerodynamic Simulation Atmospheric Research Fluid Dynamics

[1] The SPEC Release 1.0 Benchmark Suite consists of 10 applications - oriented CPU intensive benchmarks.

Methodology

Wherever possible, we have run the same benchmark codes on all systems. However, due to equipment availability of competitive systems, some data is reported based on published vendor claims and has not been independently verified.

All hardware and software is standard. Tests such as Linpack, Whetstone and Dhrystone were obtained from public sources such as Argonne National Laboratory and Lawrence Livermore National Laboratory. The Digital Review CPU2 suite was obtained from Digital Review magazine.

Accurately Reflecting Comparative Performance

When making comparisons between systems on the same benchmark, it is standard practice to normalize all results to that of a certain system. In this report, we have normalized many of the graphs to the DECstation 3100. This system is then defined to represent 1.0 on the scale, and all other systems fall higher or lower than that system. The number to the right of each bar indicates the raw data for that benchmark.

Results of individual benchmarks can be changed dramatically by the choice of operating system version, compiler version, level of optimization used and memory size. In particular, because some systems have more than one compiler available from the vendor, using different compilers can have a significant impact on benchmark performance.

A variety of systems have been tested and compared in this document. Configurations tested include workstations (with a variety of graphics controllers) as well as servers, both large and small. Our tests are primarily CPU intensive and the reader is encouraged to carefully consider the appropriateness of mapping these results into their own environments.

Specific configuration details and benchmark raw data is presented in the final sections of this document.

Summary of Findings

Digital RISC Family provides leadership performance.

For Engineering and Scientific applications, including the SPEC benchmarks:

1. The DECstation 5000 Model 200 outperforms Sun's highest performing system (the SPARCserver 490) and is 1.5 to 2.0 times faster than the Sun SPARCstation 330.

2. The DECstation 3100 offers similar or better performance than the Sun SPARCstation 330.

3. The DECstation 2100 delivers similar or better performance than the Sun SPARCstation 1.

4. The DECsystem 5400 and DECsystem 5810 provide similar or better performance than the Sun SPARCserver 390.

SPEC Positioning

SPEC (System Performance Evaluation Cooperative) has been formed to identify and create an objective set of applications-oriented tests, which can serve as common reference points and be used to evaluate computer systems performance. Digital is a member of SPEC and endorses its goals. Several factors differentiate the SPEC benchmarks from the other application benchmarks included in this document:

1) SPEC includes only applications which are in the public domain.

2) SPEC includes applications which are easily ported to numerous platforms with minimal changes to the source code.

3) SPEC consists of a small but growing number of benchmarks. It does not cover all market segments and application types at this time.

4) SPEC contains a mix of integer and floating point applications, whereas the other applications described here are primarily floating point intensive.

SPEC does not include commercially available applications due to their proprietary nature and difficulty in porting. However, many commercially available applications are widely used and therefore of interest to users. For this reason, we have included benchmark results for some additional applications in this document.

SPEC Performance

For the SPEC Benchmarks:

- The DECstation 5000 Model 200 outperforms Sun's highest performing system, the SPARCserver 490. The DECstation 5000 Model 200 is also 50% faster than the Sun SPARCstation 330.

- DECstation 3100 performance is within 5% that of the Sun SPARCstation 330.

- The DECstation 2100 provides similar performance to the Sun SPARCstation 1.

System	SPECmark *
DECstation 5000 Model 200	18.5
SPARCserver 490	17.6
DECsystem 5400	11.8
SPARCstation 330	11.8
DECsystem 5810	11.3
DECstation 3100	11.3
SPARCstation 1	8.4
DECstation 2100	8.3

* Sun results reported in [SPEC]. DEC results were the best obtainable at the time of publication.

SPEC Performance

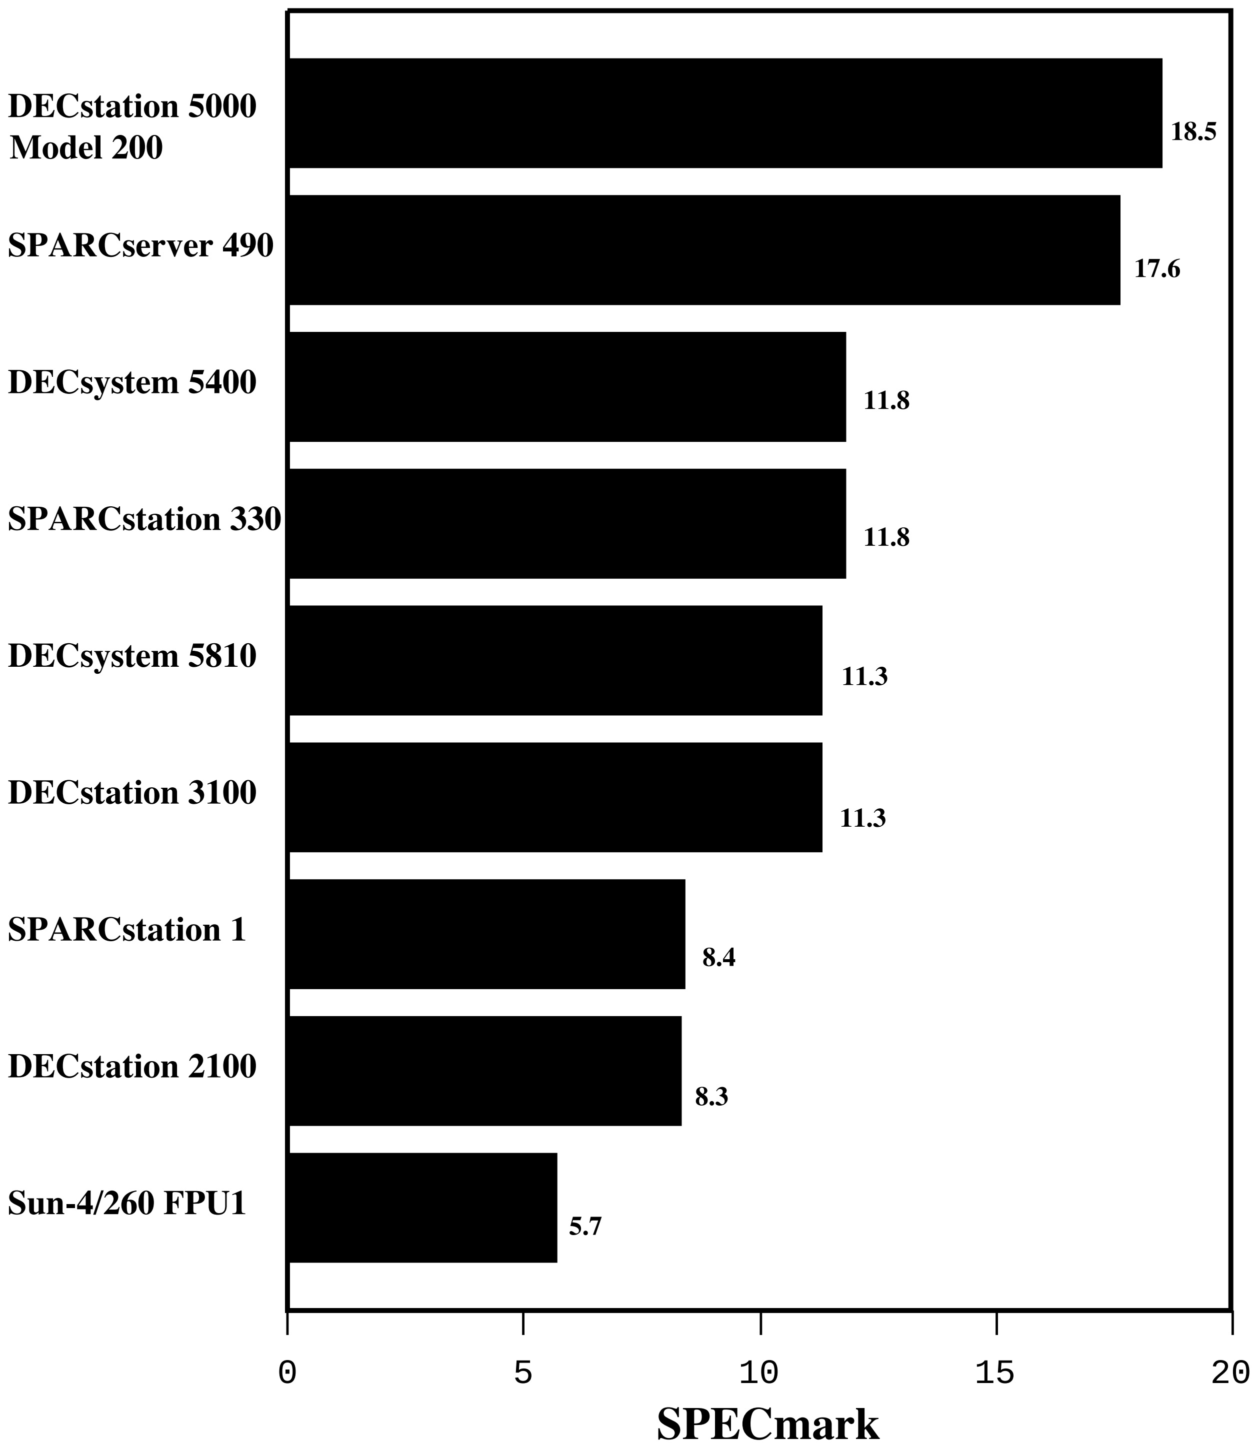

A measure of system speed based upon the SPEC Benchmark Suite.

Mechanical Design: ANSYS Finite Element Analysis

ANSYS is a proprietary finite element analysis system developed and marketed by Swanson Analysis Systems, Inc. (SASI). It is widely used for the design, evaluation and optimization of mechanical structures subject to stress, vibration and field effects.

The ANSYS application is compute intensive, although a moderate amount of I/O is performed. It relies primarily upon double precision arithmetic. The performance metric used for comparison is the CPU time required to complete each job.

SASI has developed a set of standard ANSYS benchmarks, and we have chosen the SP3 benchmark for our performance comparisons. SP3 is a moderately sized 3-D static stress problem with an RMS wave front of 124.

For the SP3 benchmark:

- The DECstation 5000 Model 200 provides comparable performance to the SGI IRIS 4D/240 and is 30% faster than the Sun SPARCserver 490.

- The DECstation 3100 outperforms the SPARCstation 330.

- The DECstation 2100 provides similar performance to the SPARCstation 1.

- Both the DECsystem 5400 and DECstation 3100 deliver superior performance compared to the HP9000/835.

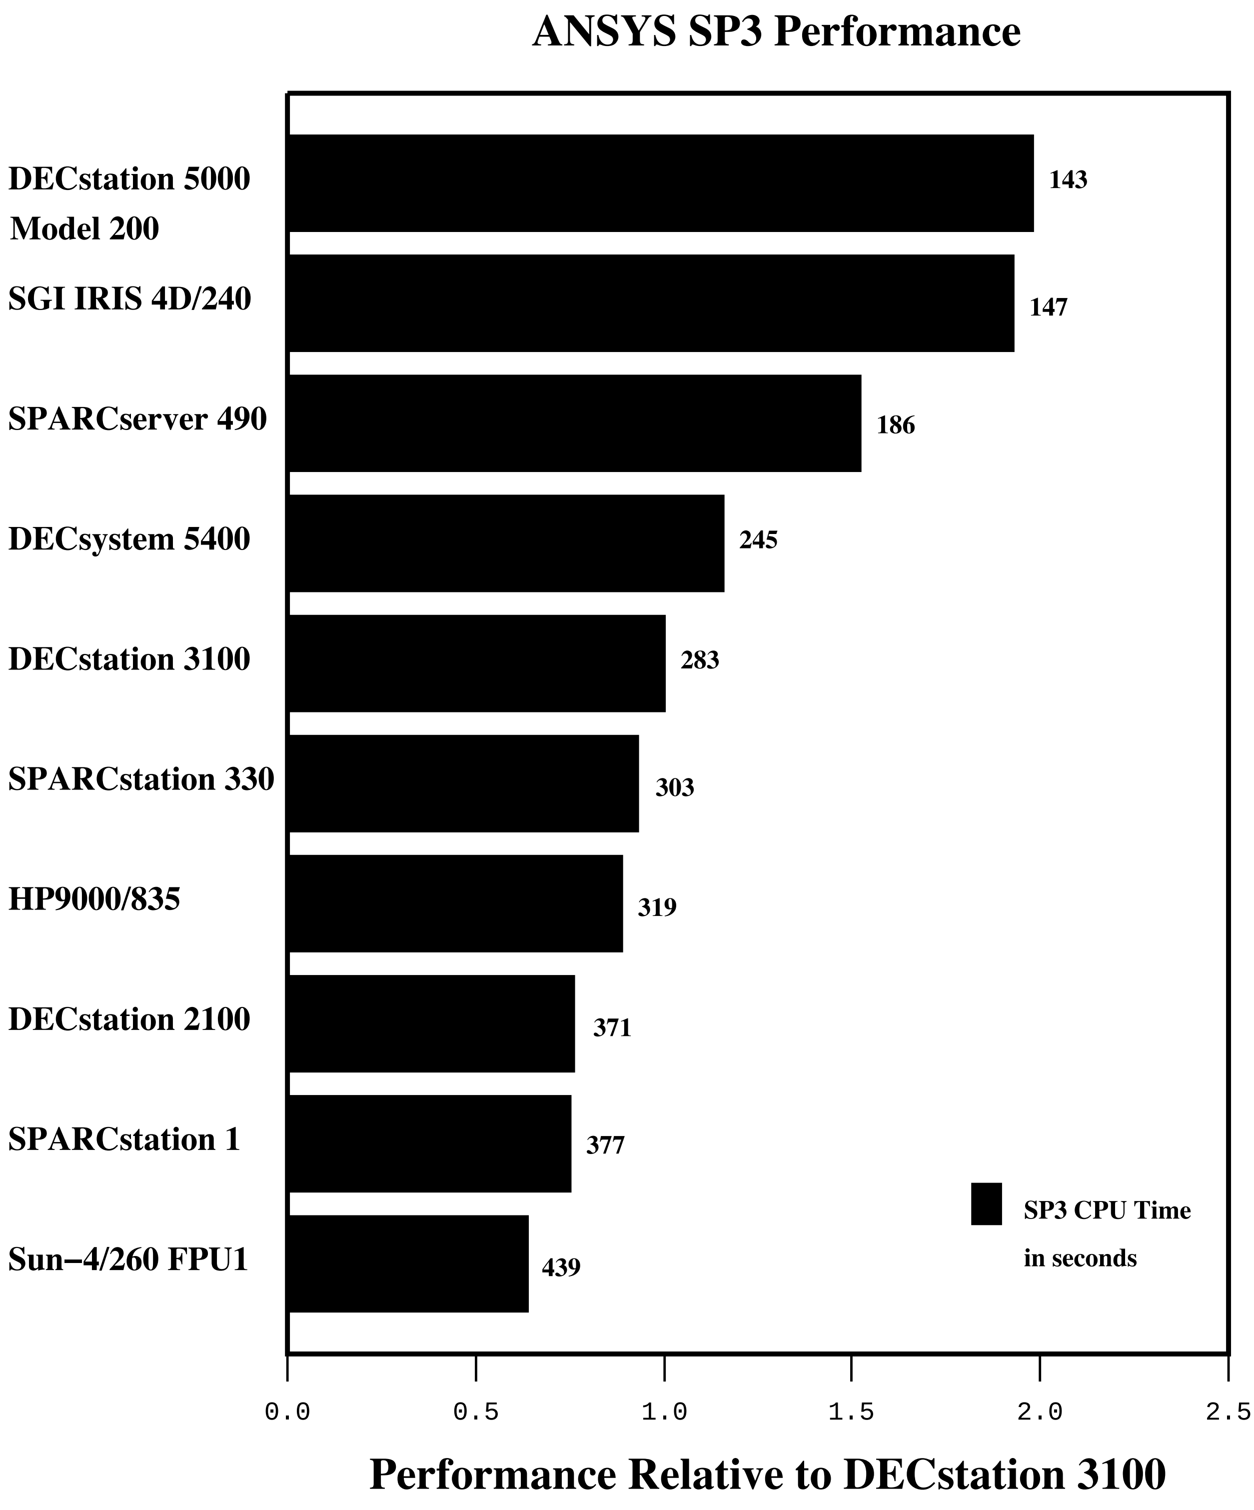

ANSYS performance for vendor benchmark SP3. CPU times for competitive systems reported by Swanson Analysis Systems, Inc. [SASI - 90]. All results are "single stream".

Scientific Computing: NCAR

The Scientific Computing Division of the National Center for Atmospheric Research (NCAR) has developed a shallow water model benchmark [NCAR-89] which it uses to characterize the performance of various computer systems ranging from workstations through supercomputers.

The benchmark is written in FORTRAN, is floating point oriented and does little I/O. The version of the benchmark that we tested solves the shallow water equations, a set of time dependent partial differential equations, on a 256 by 256 planar grid.

The benchmark is typical of applications that perform finite difference approximations. Most of the calculations are performed in three doubly nested DO loops.

The metric of comparison is the measured MFLOPS for each system. The Sun results were obtained from NCAR and the other results were measured by Digital.

For the NCAR shallow water model benchmark:

- DECstation 5000 Model 200 performance was double that of the SPARCstation 330.
- The DECstation 3100 performed 15% faster than the SPARCstation 330.
- The DECstation 2100 achieved a higher MFLOP rating than the SPARCstation 1.

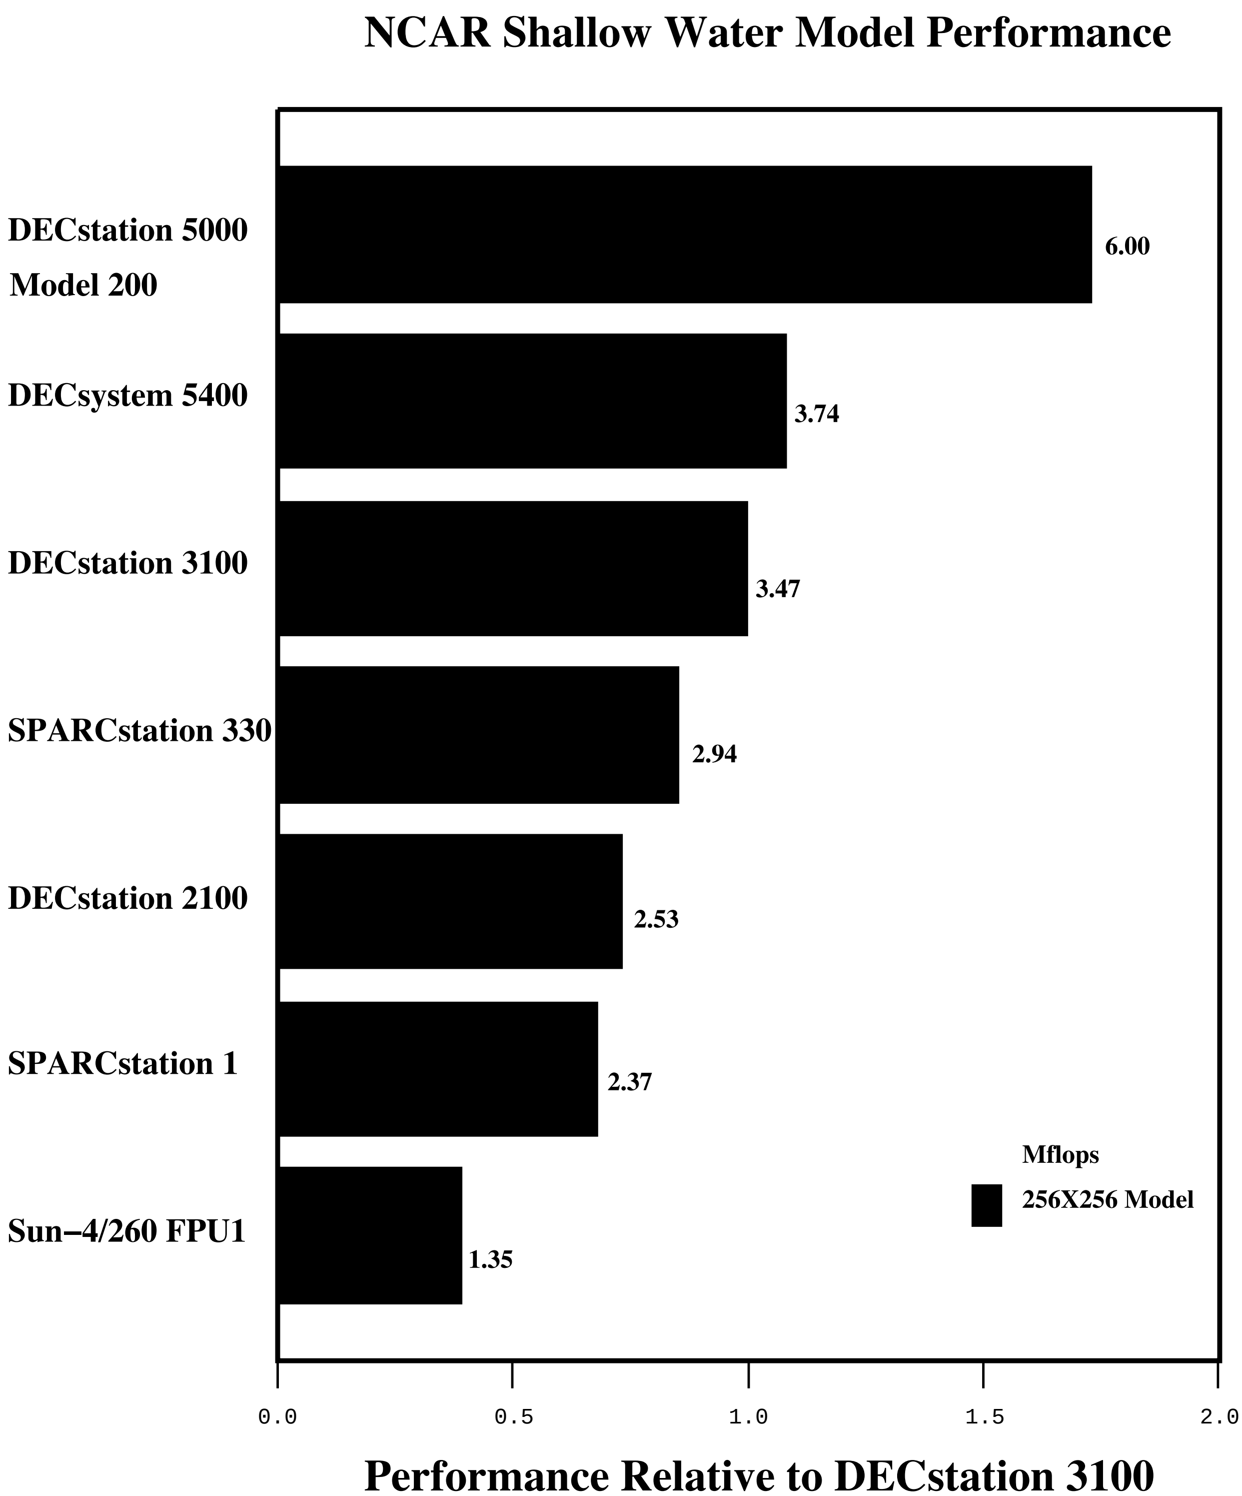

Shallow Water Model performance in Mflops. Benchmark is typical of finite difference approximations and is floating point intensive.

Scientific Computing: DR Labs CPU2

The Digital Review magazine DR Labs CPU2 benchmark is a moderately sized FORTRAN program that includes 34 separate tests, many of which have been derived from application codes and benchmarks developed at the National Laboratories. The benchmark is most relevant in predicting the performance of engineering/scientific applications which are primarily single precision floating point oriented.

We follow Digital Review in comparing the geometric mean of the CPU times for the 34 kernels tested. For consistency in the way that these results are reported by Digital Review, we express the performance of each system as a multiple of MicroVAX Units of Performance (MVUPS).

For the DR Labs CPU2 benchmarks:

- The DECstation 5000 Model 200 offers 1.8 times the performance of the SPARCstation 330.
- The DECstation 3100 is over 10% faster than the SPARCstation 330.
- The DECstation 2100 is 25% faster than the SPARCstation 1, and provides comparable performance to the HP9000/835 workstation.

DR Labs CPU2 Performance

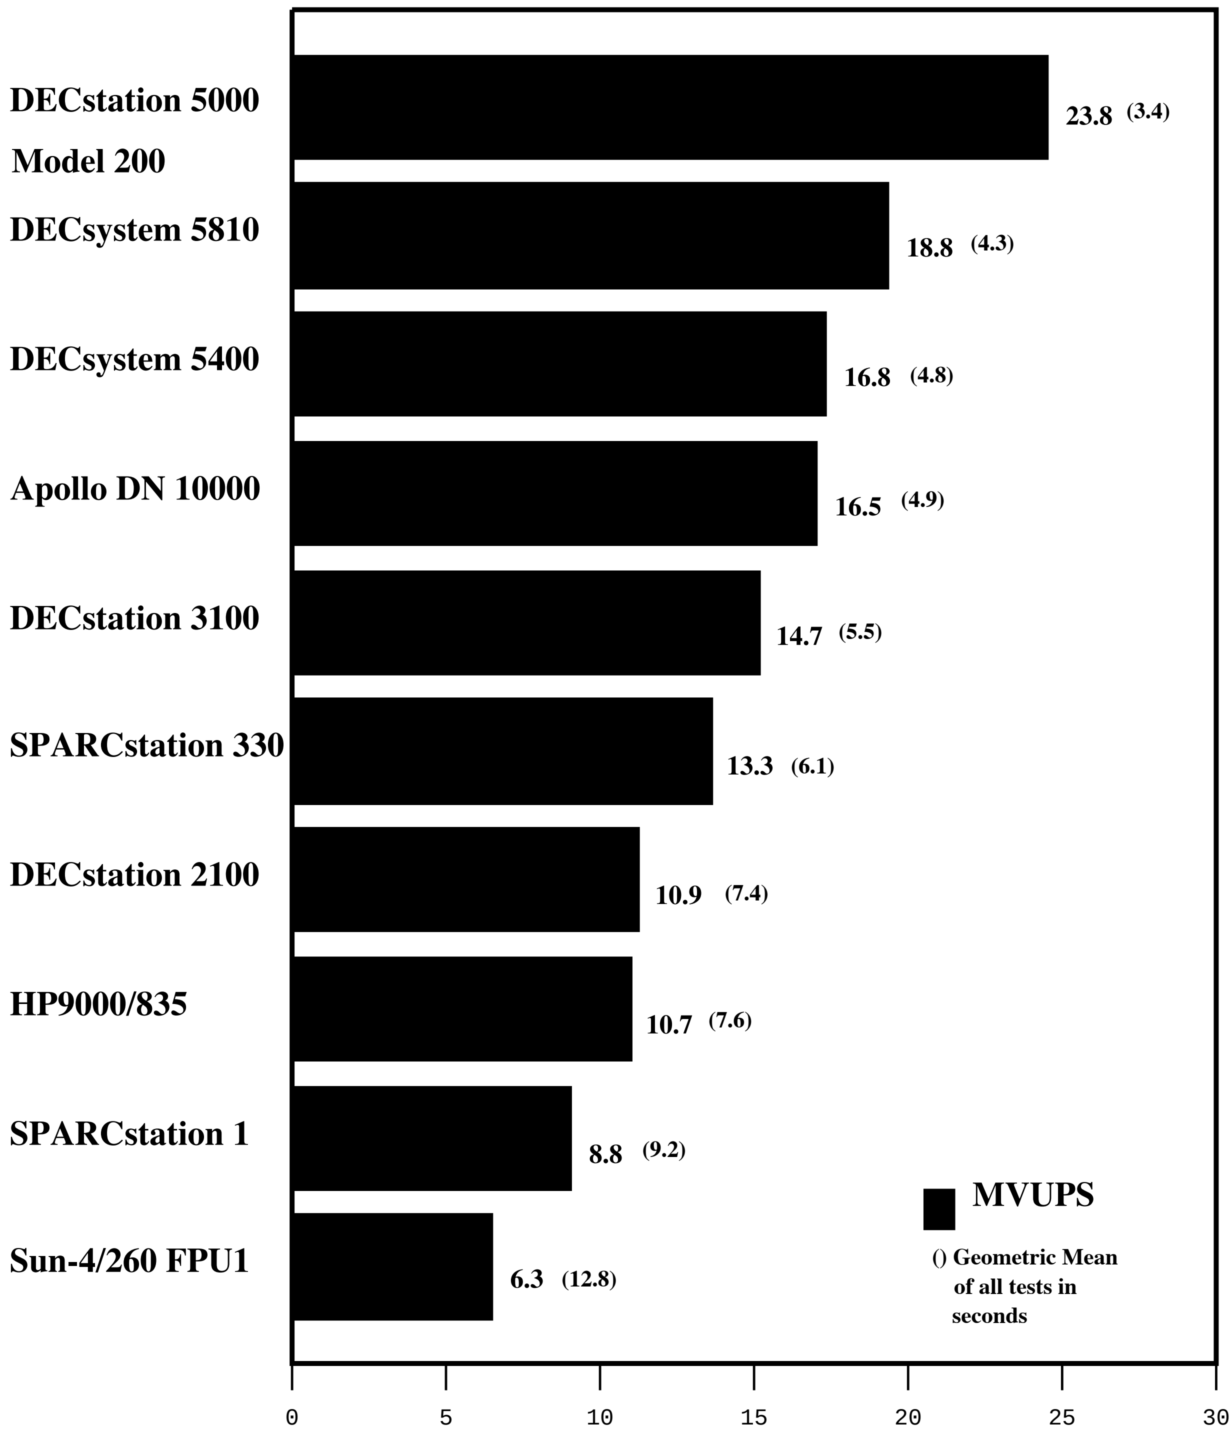

Performance Relative to MicroVAX II (MVUPS)

Performance is normalized to MicroVAX II reference system. The MVUP Rating was calculated by taking the ratios of the listed geometric means to 80.96 (MicroVAX II geo mean).

Chemical Engineering: MDATOM

MDATOM is a FORTRAN Molecular Dynamics simulation that is typical of Computational Chemistry application workloads [MDATOM]. The benchmark tests single precision floating point operations with integer loop counter overhead, while performing almost no I/O.

It was developed by Cornell University's Chemical Engineering Department to model the liquid state by solving the coupled newtonian equations of motion for several hundred model particles in a cell of specified geometry.

The performance metric is the CPU time required to evaluate 1500 time steps for 108 Lennard-Jones atoms. All results are reported for optimization level -O, unless otherwise specified.

For the MDATOM benchmark:

- The DECstation 3100 is 50% faster than the SPARCstation 330.

- The DECstation 2100 is 50% faster than the SPARCstation 1, and also outperforms the SPARCstation 330.

MDATOM Performance

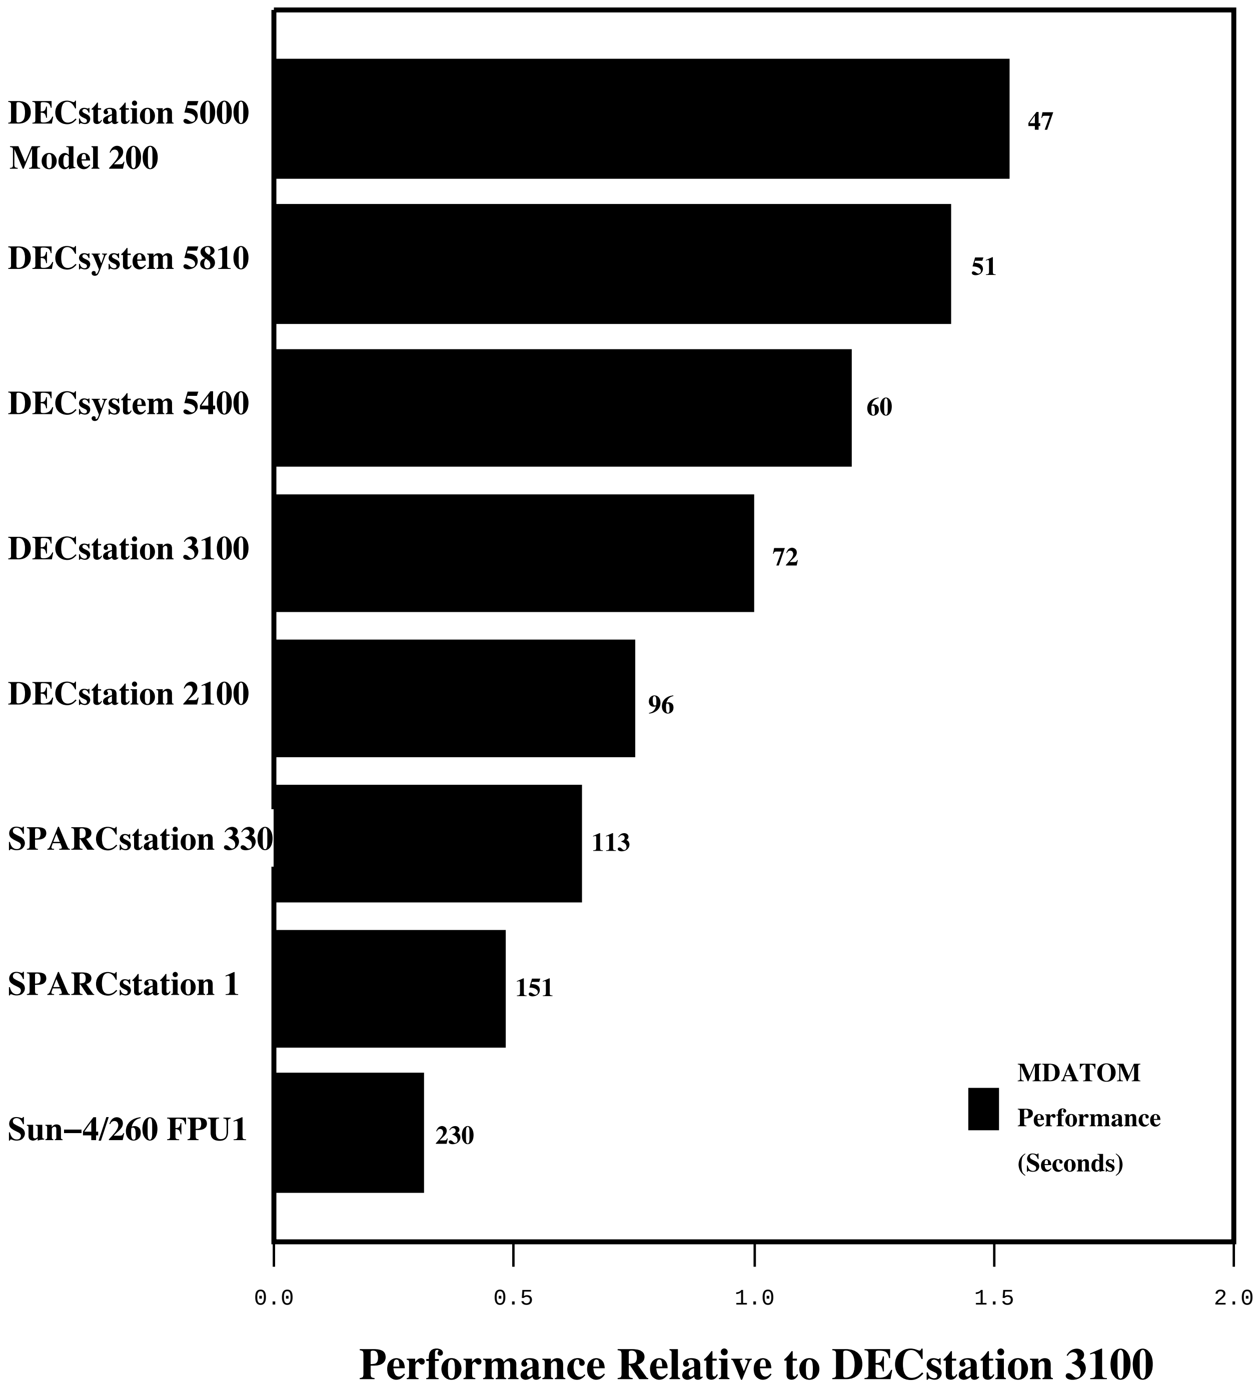

MDATOM Benchmark performs single precision floating point operations and is typical of Computational Chemistry application workloads.

Geo/Process Engineering: Dynamic Graphics Inc. SGL Benchmark TOP

SGL is a gridding and contouring ANSI FORTRAN 77 subroutine library. It is used in conjunction with the companion DGI library SDL for digital terrain modeling (DTM) and mapping. SGL is also an integral component of the DGI product ISM (Interactive Surface Modeling).

The benchmark TOP is an industry standard gridding and contouring benchmark for a large faulted data set. TOP is based on an actual Geo/Process industry data set and was a test component of the August 1987 Denver CEED/II (Comparison, Evaluation, Exhibition and Demonstration). This conference was sponsored by the Denver Geophysical Society. The conference posed technical problems and test data cases to a wide variety of third party Geo/Process software packages and interpretive workstations. TOP consists of 17,729 scatter data points and 1000 faults. A 256 x 256 grid was used to create the model surface. Actual performance is noted in CPU seconds.

Based upon the TOP benchmark, the DECstations offer superior performance:

- The DECstation 5000 Model 200 offers more than double the performance of the SPARCstation 330.

- The DECstation 3100 provides 1.33 times the performance of the SPARCstation 330.

- The DECstation 2100 achieves similar performance to the SPARCstation 330 and performs 1.36 times faster than the SPARCstation 1.

- Both the DECstation 3100 and 2100 outperform the Apollo DN10000.

TOP Performance

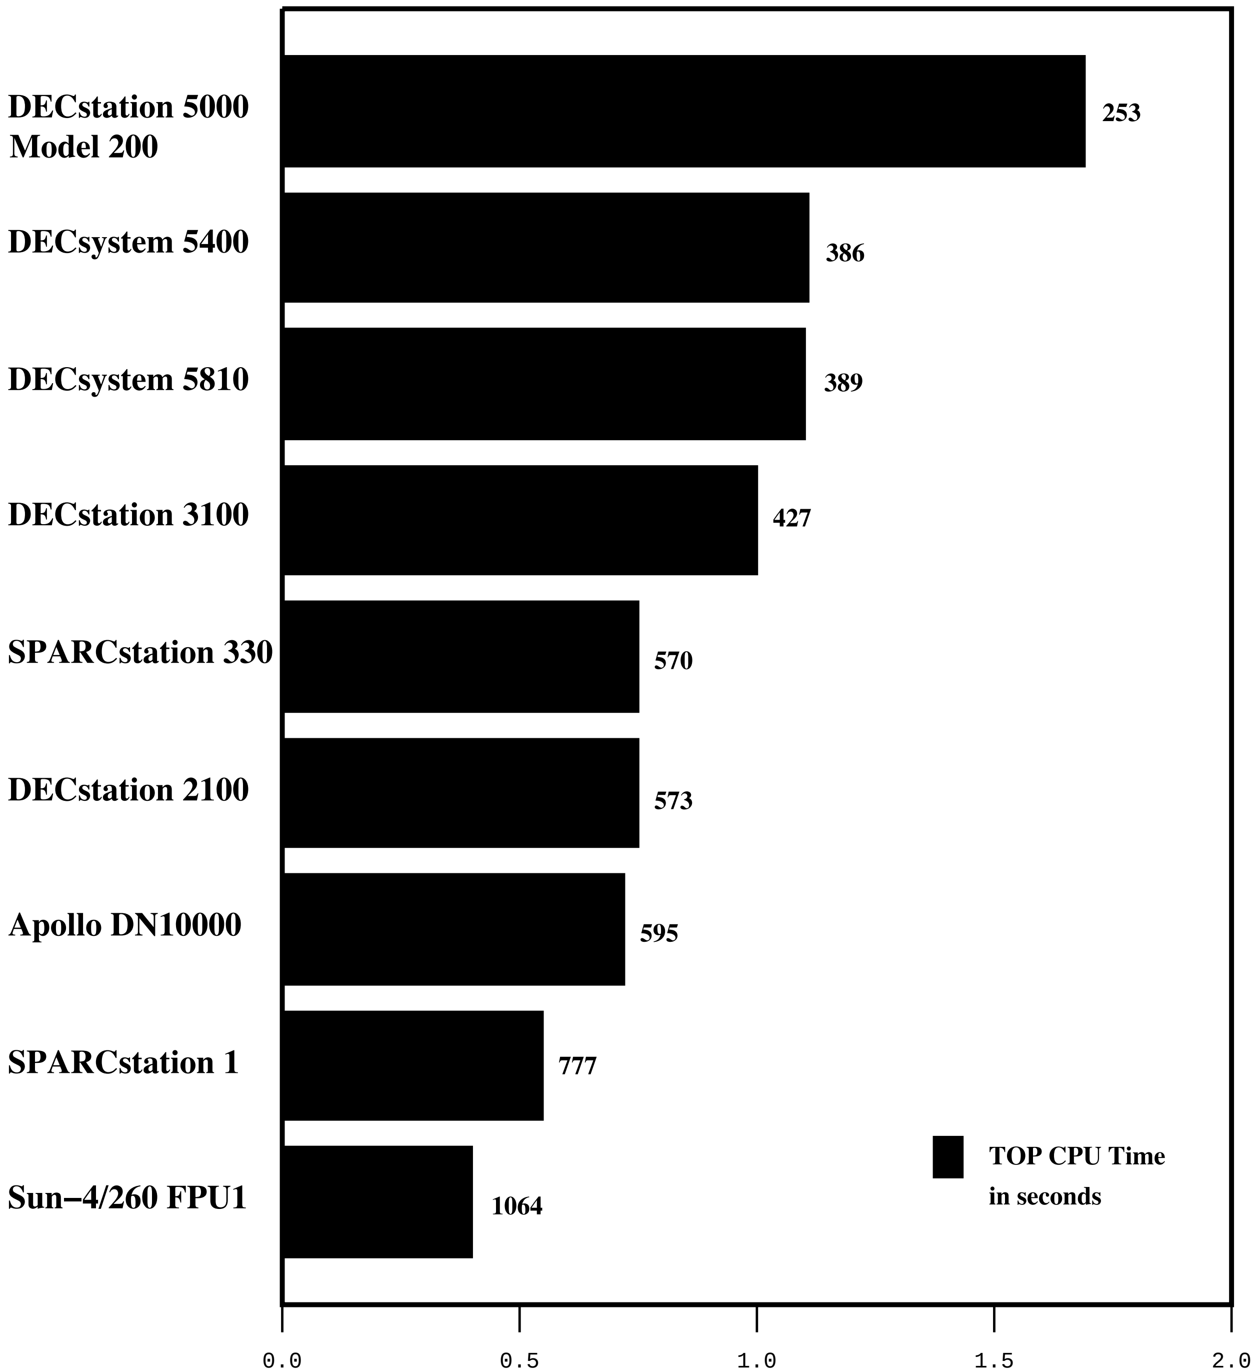

Performance Relative to DECstation 3100

DGI TOP gridding and contouring benchmark. Apollo results were reported by DGI.

Geo/Process Engineering: UTCHEM

UTCHEM [Pashapour] simulates the chemical flooding of an oil reservoir. The application was developed at the University of Texas at Austin.

The program uses a three dimensional finite difference algorithm. The simulator solves the material balance equations for up to nineteen components, each of which may form up to three phases. The governing component conservation and overall material balance equations are solved by a finite-difference approximation of the spatial derivatives and a forward difference approximation of the time derivatives.

UTCHEM is a large FORTRAN program that measures both floating point and integer performance. We have tested UTCHEM with a reservoir size of 22x22x4.

The performance metric for this benchmark is elapsed time in seconds.

For the UTCHEM 22x22x4 benchmark:

- DECstation 3100 performance is within 5% that of the SPARCstation 330.

- The DECstation 2100 provides similar performance to the SPARCstation 1

UTCHEM Performance

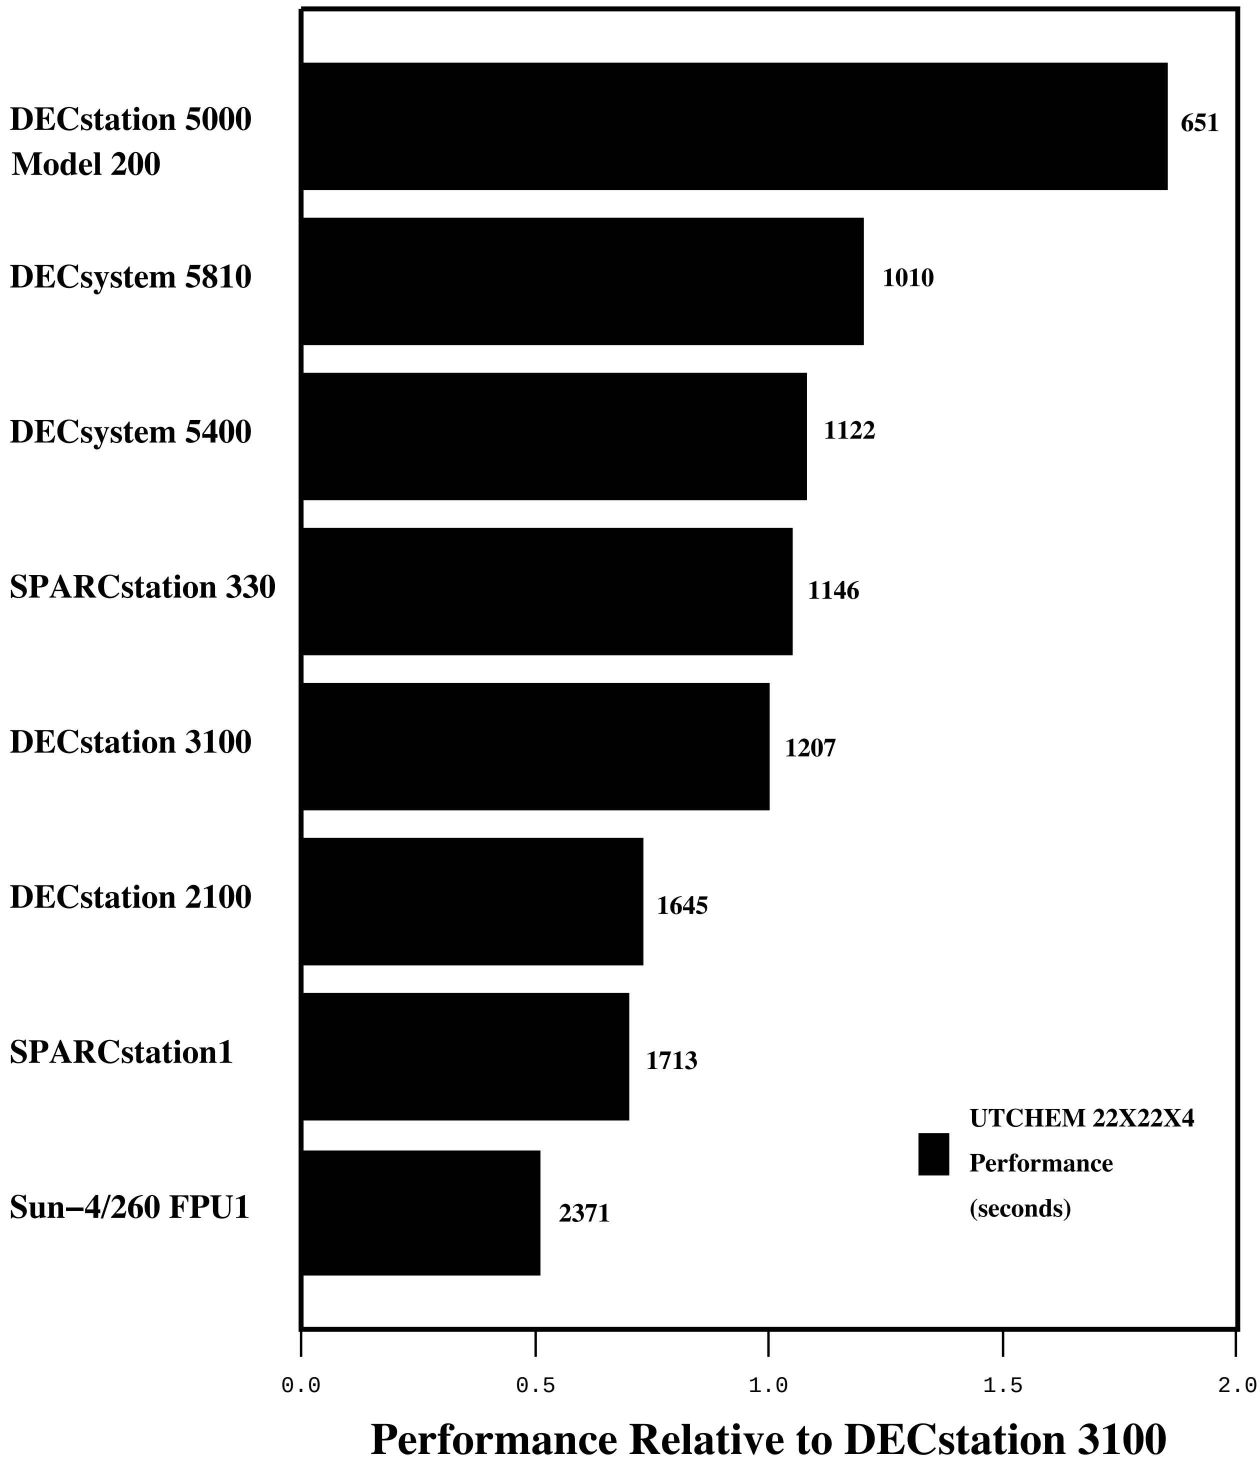

UTCHEM oil reservoir simulation benchmark provided by the University of Texas at Austin.

FLOATING POINT BENCHMARKS

Linpack

A number of benchmarks that solve dense systems of linear equations using Linpack subroutines have been developed by Jack Dongarra of Oak Ridge National Laboratory. Linpack programs can be characterized as having a high percentage of floating point additions and multiplications and most of the time is consumed in a set of subroutines called the Basic Linear Algebra Subprograms (BLAS), which are called repeatedly throughout the benchmark.

The Linpack benchmarks are compared based upon the execution rate as measured in MFlops. The most popular variants solve a 100 x 100 system of equations, either in single or double precision, and have become one of the most widely used benchmarks, to gauge engineering/scientific applications performance. For example, many finite element, finite difference, simulation and regression analysis applications exploit Linpack like equation solvers.

Results can be generated either by using standard FORTRAN versions of the BLAS or by using "hand-coded" variants. The FORTRAN BLAS are typically available in two forms - the standard version which has its vector-scalar multiply-add routine unrolled four times and a version where the unrolling is undone (annotated "Rolled BLAS"). Both results are presented in the following table:

Linpack Floating Point Performance in Mflops
Using Both Standard and Rolled BLAS

	"Standard BLAS"		"Rolled BLAS"	
System	S. Precision	D. Precision	S. Precision	D. Precision
DECstation 5000 Model 200	6.4	3.7	6.5	3.6
DECsystem 5810	4.3	1.8		
DECsystem 5400	4.0	1.6	4.0	1.6
DECstation 3100	4.0	1.6	3.9	1.6
DECstation 2100	2.8	1.2	2.9	1.2
HP9000/835	2.3	1.8	2.1	1.6
SPARCstation 330	---	---	3.6	2.7
SPARCstation 1	2.0	1.1	2.3	1.4
Sun-4/260 FPU2*	---	---	2.2	1.7
Sun-4/260 FPU1	1.3	.89	1.7	1.1

*Measured by Sun [SUN89]

MFlops

Linpack Performance (Single Precision)

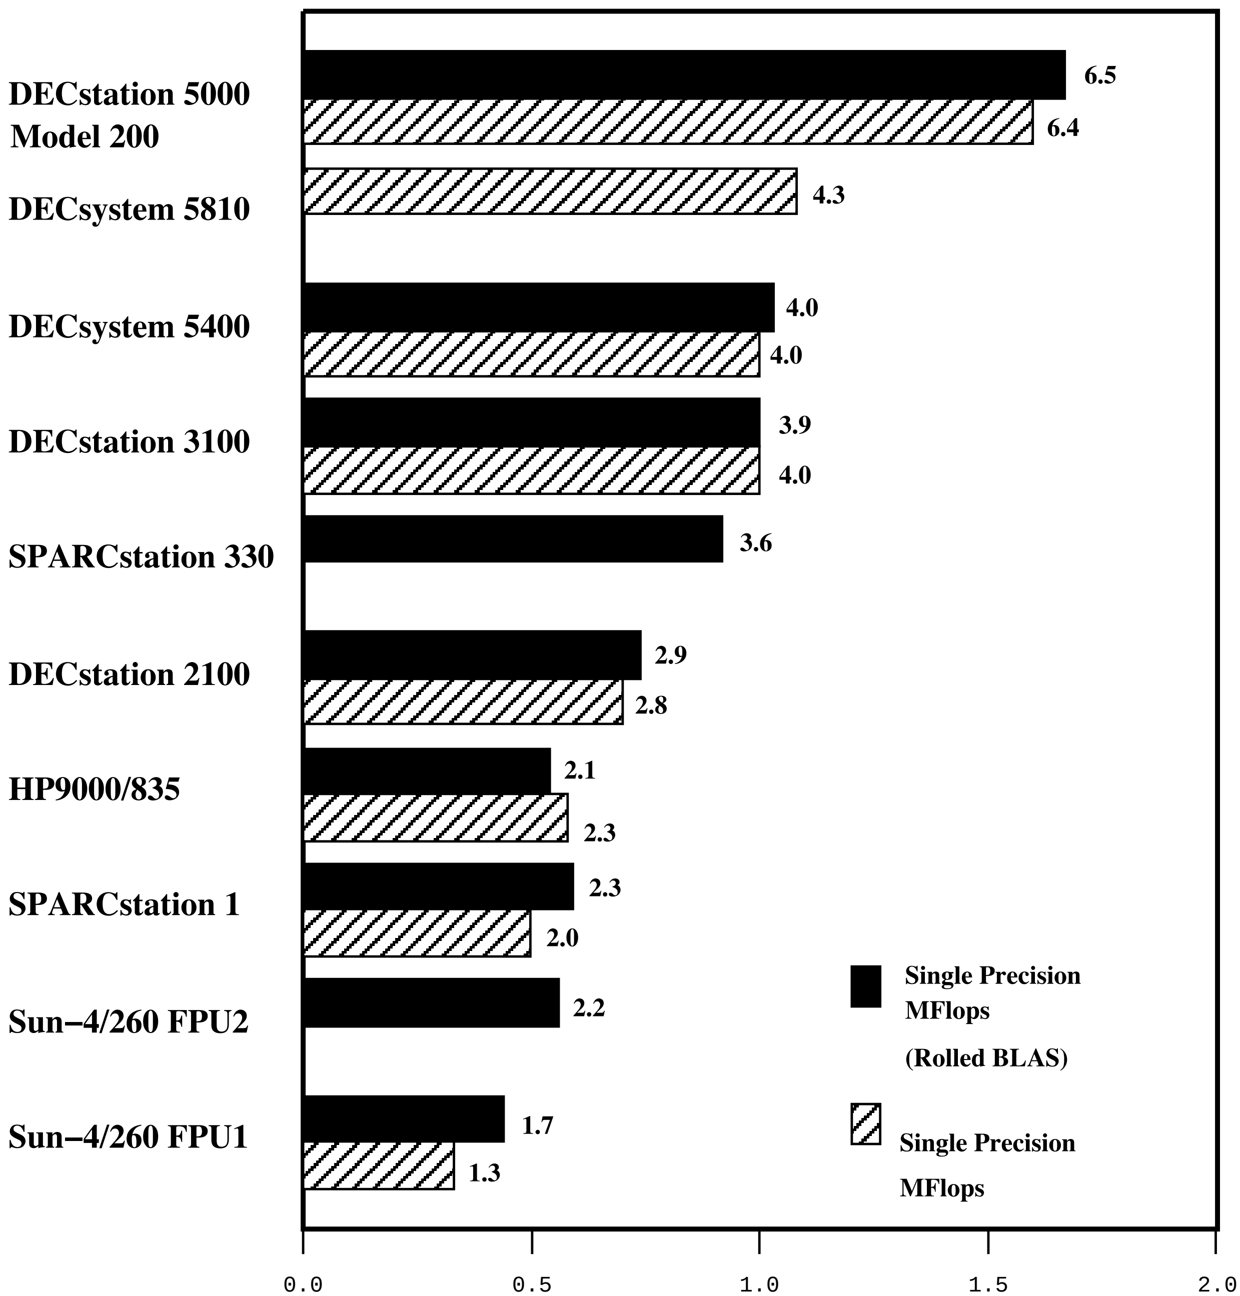

Performance Relative to DECstation 3100

Linpack 100 x 100 Performance in MFlops using both versions of FORTRAN BLAS.

Linpack Performance (Double Precision)

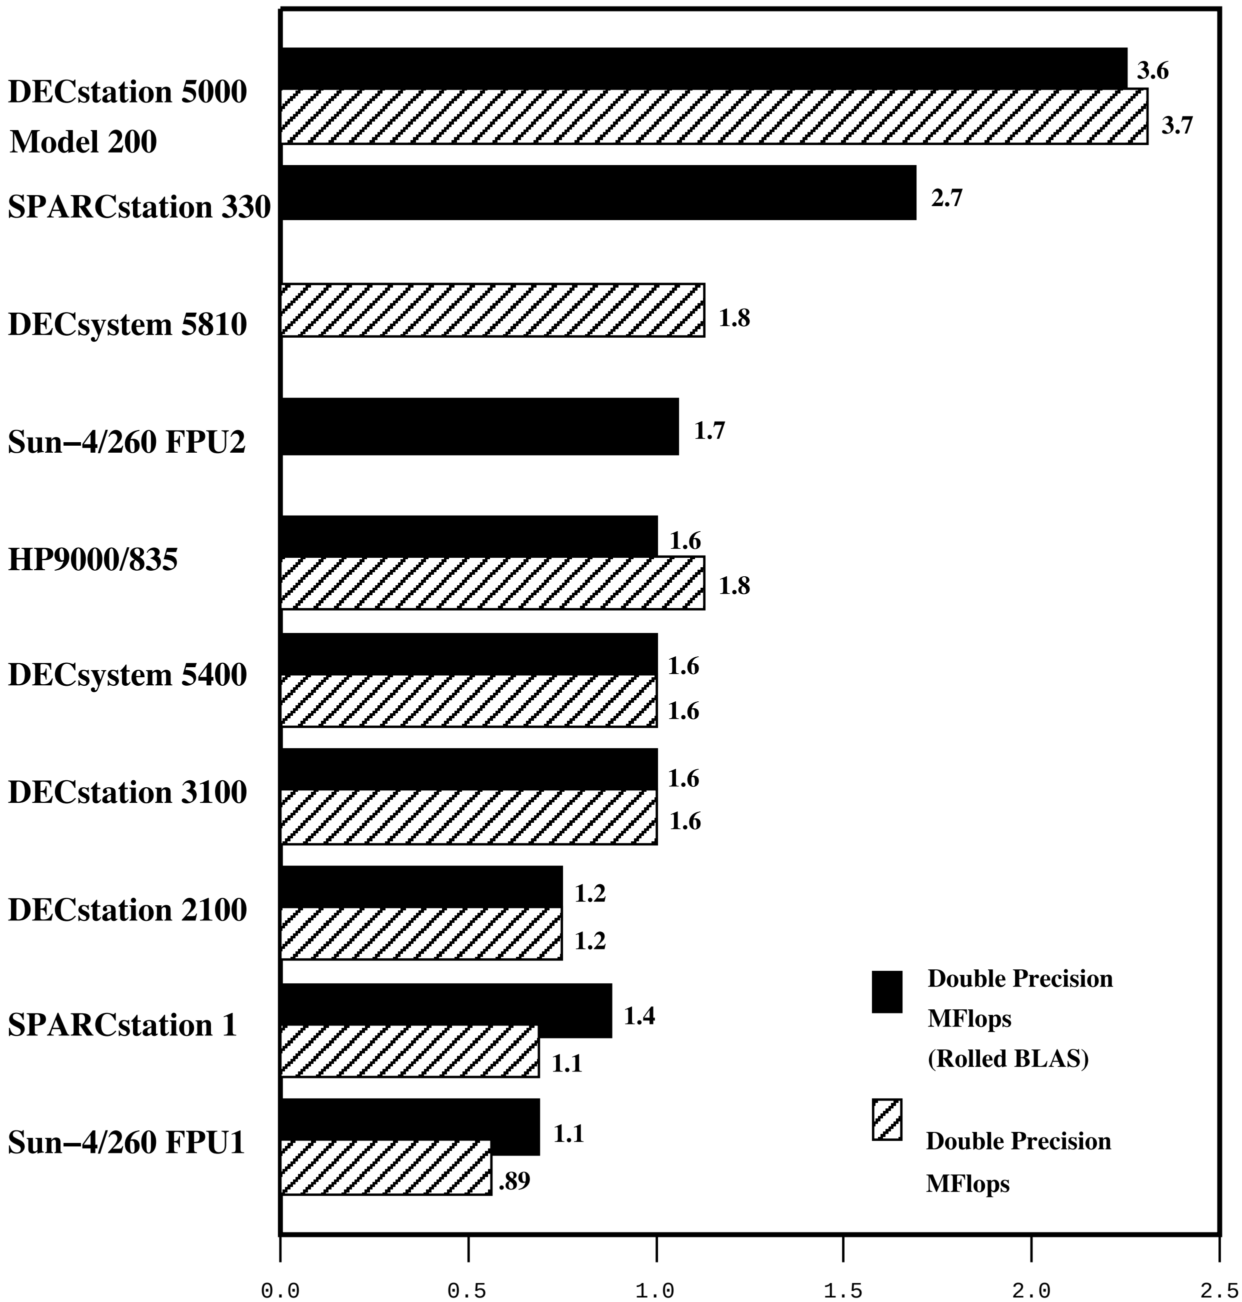

Performance Relative to DECstation 3100

Linpack 100 x 100 Performance in MFlops using both versions of FORTRAN BLAS.

Whetstone

A Whetstone is a synthetic mix of floating point and integer arithmetic, function calls, array indexing, conditional jumps and transcendental functions that was derived from an analysis of one thousand ALGOL programs. The Whetstone benchmark has been implemented in a single precision and double precision FORTRAN program, each carefully arranged to defeat most compiler optimizations. The results of the Whetstone benchmarks are measured in KWIPS (Thousands of Whetstone Instructions Per Second).

There are many permutations of the Whetstone benchmark - so it is important to ensure that comparisons across various systems utilized the same source code and that the internal loop counter is defined large enough to reduce timing variability.

Despite its synthetic mix of operations, Whetstone is generally considered a floating point benchmark and is most representative of small engineering/scientific applications that are able to fit into cache memory.

The results graphed on the following page are for compiler optimization level 3. However, compiler optimization level 4 significantly improves the Whetstone ratings on the Digital RISC platforms. The results for both optimization levels are shown in the table below:

System	Double Precision KWIPS		Single Precision KWIPS	
	-03	-04	-03	-04
DECstation 5000 Model 200	14271	16588	18183	21092
DECsystem 5810	12594	15576	17123	21142
DECsystem 5400	10604	12771	14451	16667
DECstation 3100	9137	10700	12071	13772
DECstation 2100	6798	7983	8996	10278

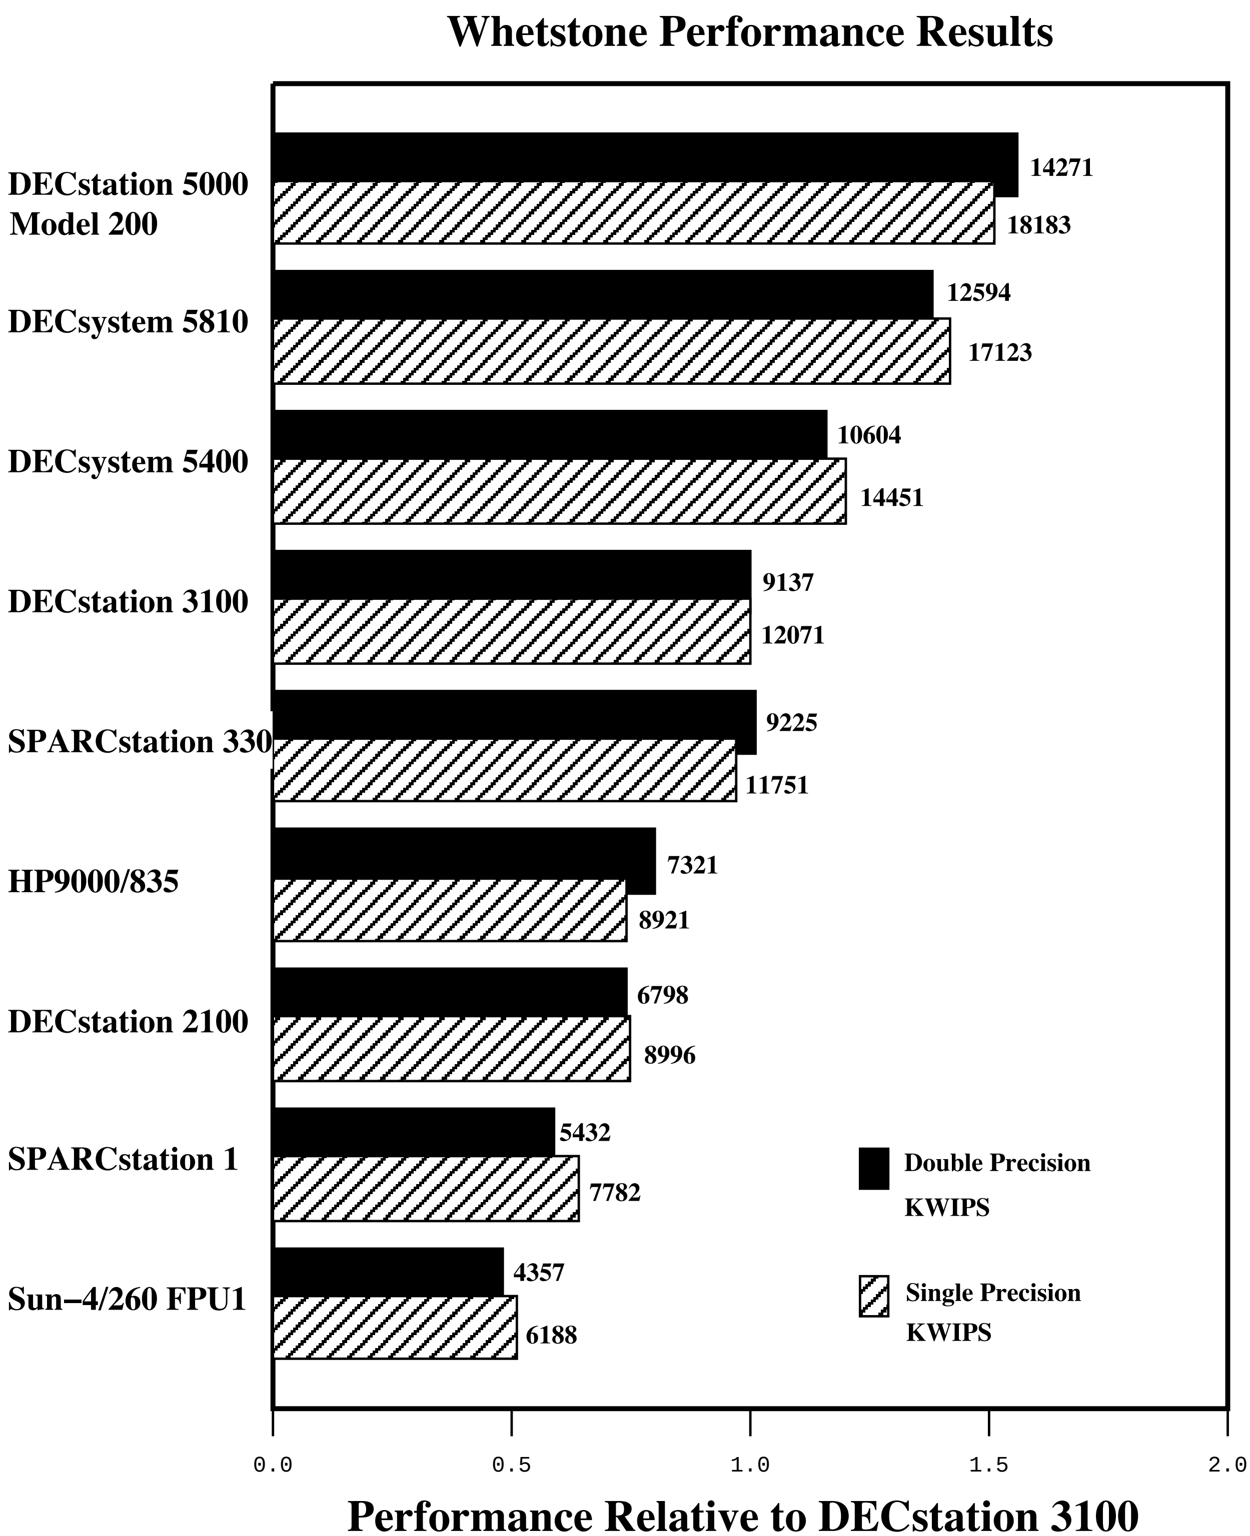

Whetstone Performance, compiler opt=O3

Lawrence Livermore FORTRAN Kernels

Lawrence Livermore National Laboratories' workload is a set of 24 separate tests which are dominated by large scientific calculations which are largely vectorizable. They have been abstracted from the applications run at Lawrence Livermore National Laboratories and run widely on systems from supercomputers to PCs [McMahon 86].

The results of this suite of benchmarks are rather complex to interpret because there is no attempt to distill the results down to a single number. The results, given in Mflops (Millions of Floating Point Operations per second), are reported for minimum, maximum and 3 means: arithmetic, geometric and harmonic.

According to McMahon, the arithmetic mean corresponds to 90%+ vectorization, the geometric mean to approximately 70% vectorization, and the harmonic mean to approximately 40% vectorization. According to John Mashey [MIPS 90], the suite can be considered as different benchmarks which each characterize different applications based on the amount of vectorizable code. For example, the geometric mean offers a good characterization of LINPACK performance; the harmonic mean to SPICE performance.

The graph on the next page shows the geometric and harmonic mean for a loop length equal to 167.

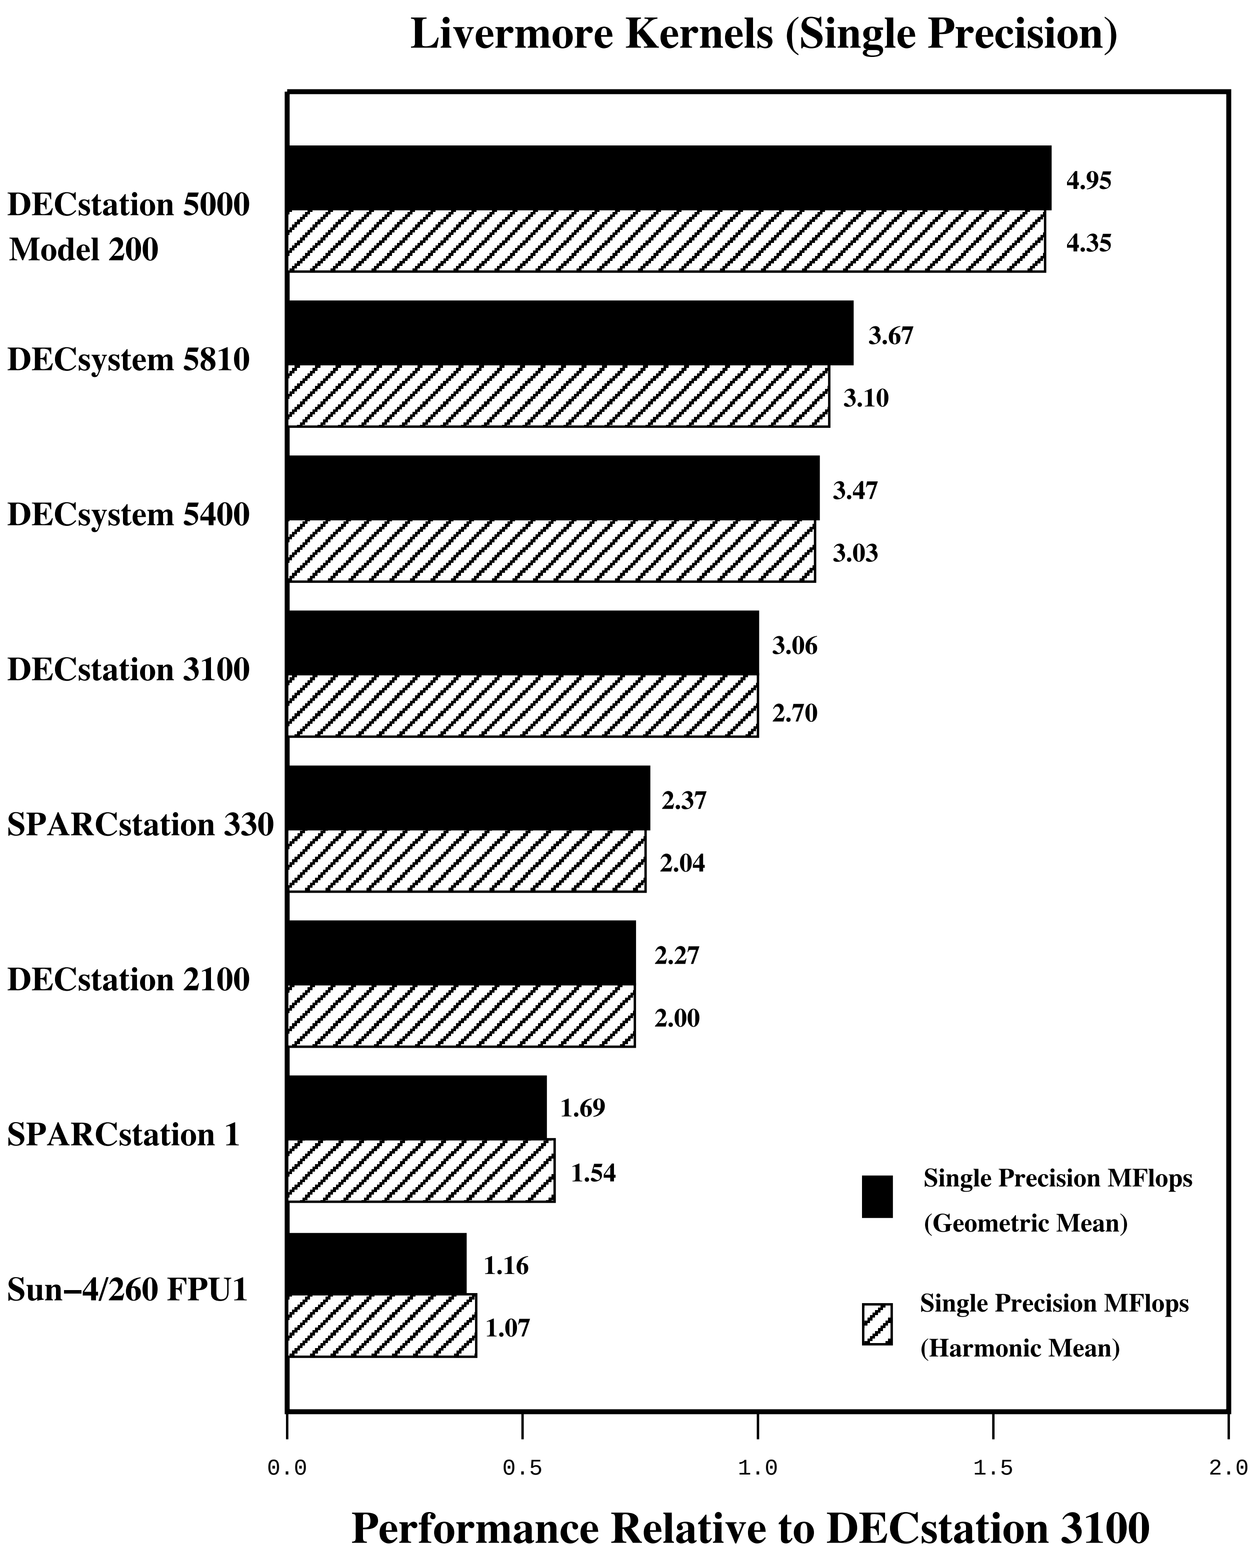

Lawrence Livermore FORTRAN Kernel Performance, Harmonic and Geometric Mean, in Mflops.

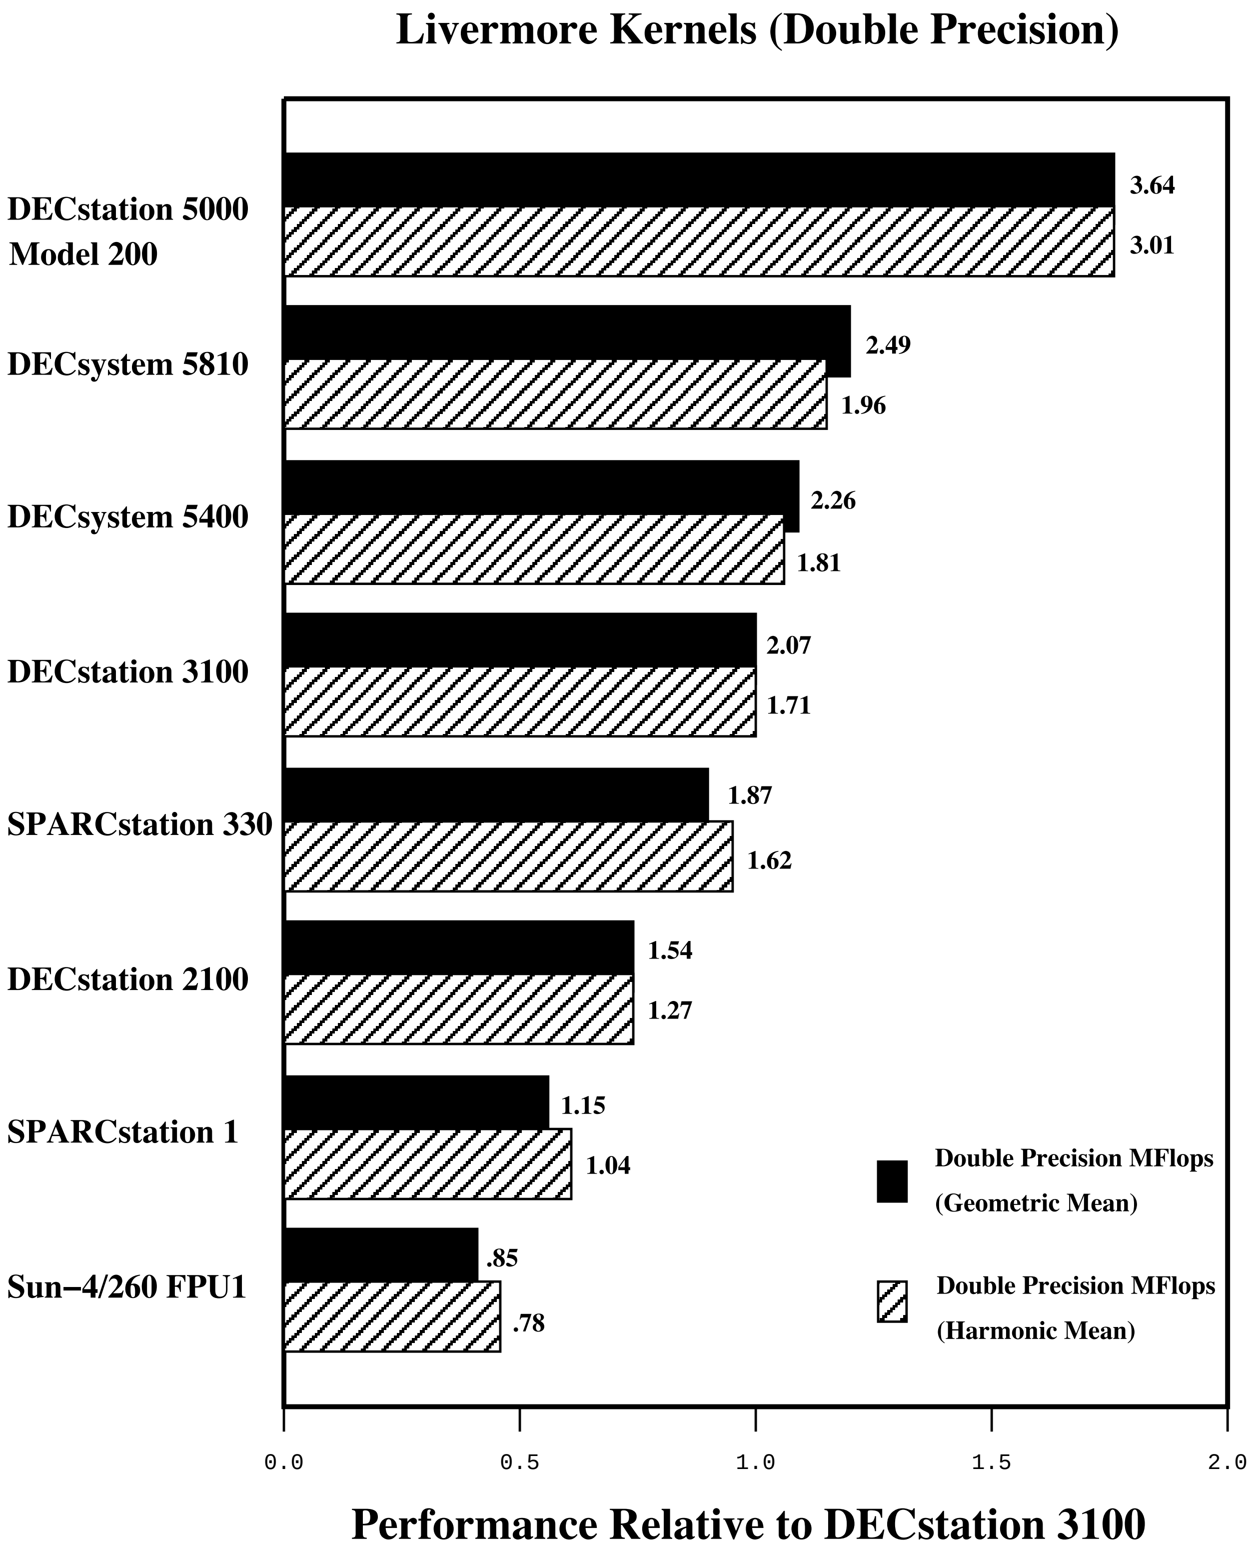

Lawrence Livermore FORTRAN Kernel Performance, Harmonic and Geometric Mean, in Mflops.

INTEGER BENCHMARKS

Dhrystone

The Dhrystone program was developed by Reinhold Weicker in 1984 as a Synthetic Systems Programming Benchmark. Despite the fact that it is distributed as either a "C", "Pascal" or "Ada" program, Dhrystone results are almost always compared using the "C" variant.

Dhrystone is widely available, easy to run and is arguably the industry's most popular Integer benchmark. Unfortunately, the result obtained is difficult to fairly compare amongst differing computing architectures and is almost as sensitive to how the Dhrystone executable image is compiled and linked as it is to the underlying processor speed. The benchmark documentation presents a set of ground rules for building and executing Dhrystone. Today, the accepted practice is to run the benchmark under any environment you wish, as long as the environment is clearly described and procedure inlining compiler optimization is not employed. We increased the loop counter in the program to five million in order to improve the consistency of the benchmark results.

The Dhrystone benchmark has been updated several times, so when comparing results, it is important to reference which version was utilized. In this analysis, we use V1.1 and V2.1 (versions 1.0, and 2.0 also exist). Dhrystone does not seem to be the best indication of application performance and is unusual in the following respects:

- Unusually low dynamic nesting depth of function calls
- Unusually low number of instructions executed per function call
- Large percentage of time spent in "strcpy" and "strcmp" routines, processing unusually large character strings
- Character strings are typically alignable on a word boundary

We include the Dhrystone benchmark in our performance evaluation because of its popularity, but warn against using it as the sole basis of comparing system performance and of accepting results that don't explicitly label how the benchmark was built and what optimizations were exploited.

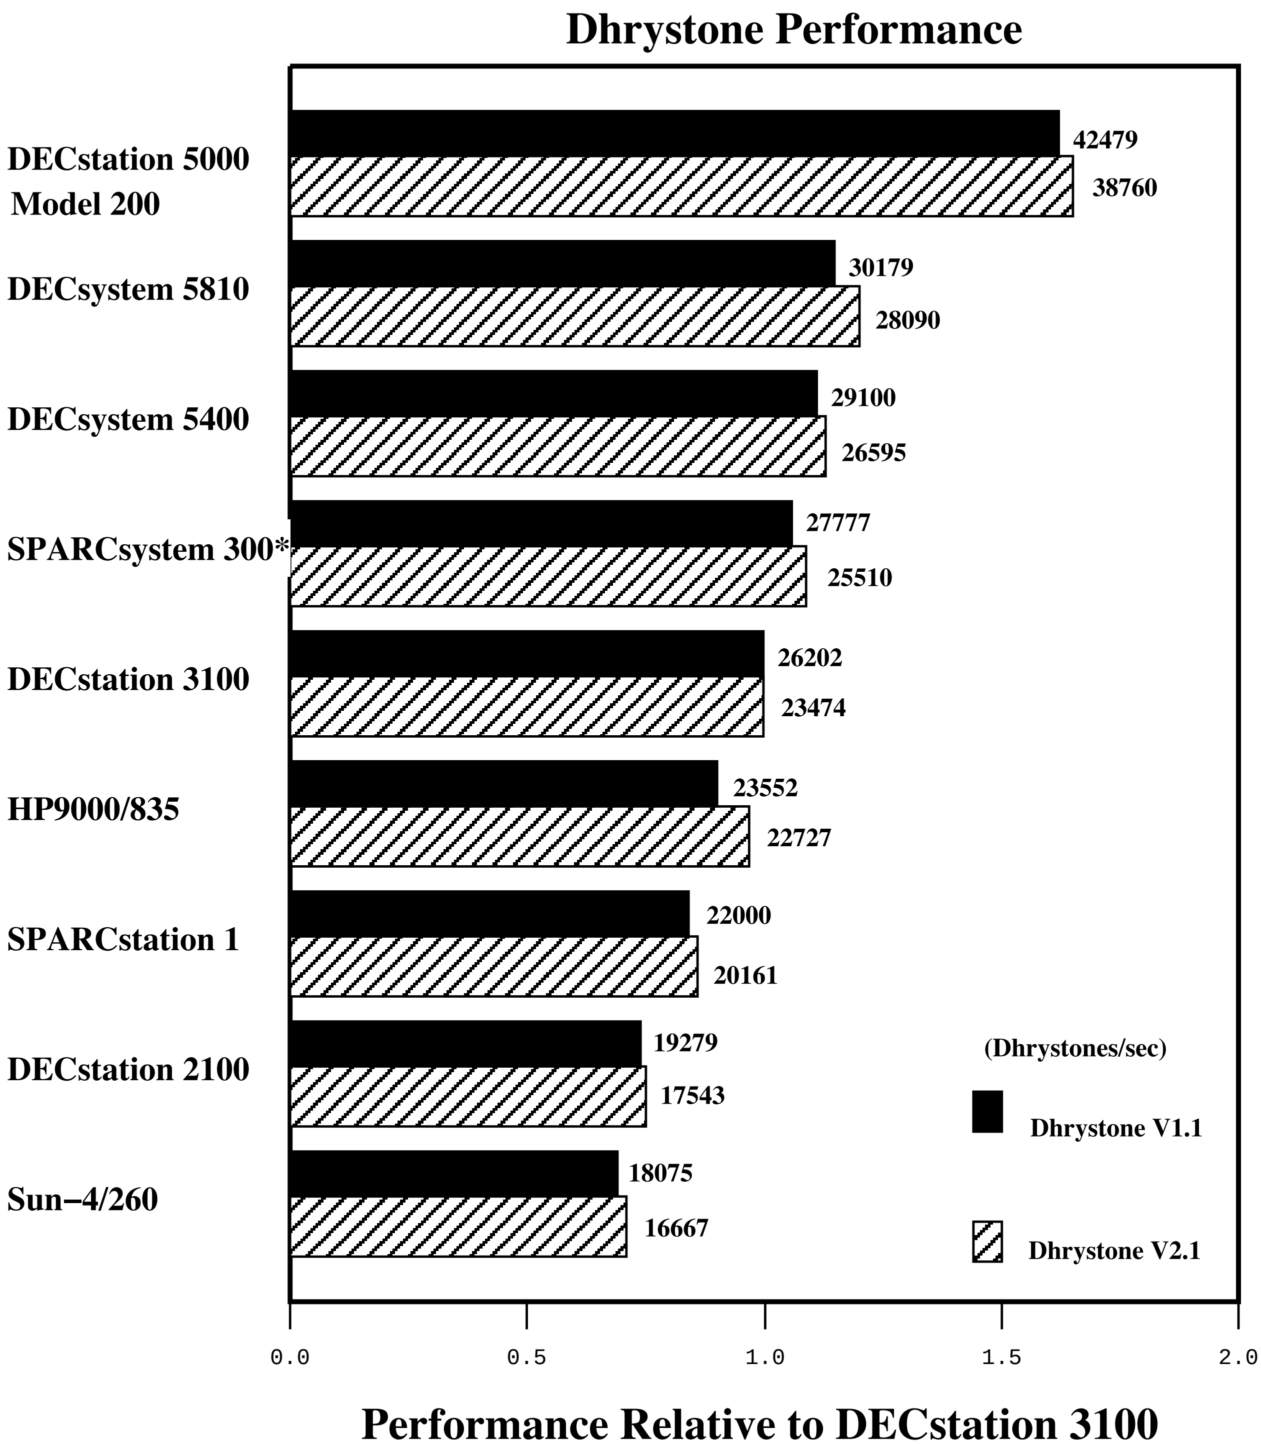

Dhrystone Performance V2.1 and V1.1 using compiler opt=O3, loop counter = five million.

* SPARCsystem 300 results for V1.1 measured by Sun [SUN 89] and are based upon compiler optimization level -04

Test Configurations

The benchmarks were run on systems with the following configurations:

DECstation 2100 Workstation

CPU: MIPS R2000, 12.5 Mhz
FPU: MIPS R2010, 12.5 Mhz
Main Memory: 16 MB
Disk: 332 MB RZ55
Operating System: ULTRIX-32 V3.1
Compilers: cc, f77 V2.10

DECstation 3100 Workstation

CPU: MIPS R2000, 16.67 Mhz.
FPU: MIPS R2010, 16.67 Mhz.
Main Memory: 24 MB
Disk: 332 MB RZ55
Operating System: ULTRIX-32 V3.1
Compilers: cc, f77 V2.10

Note: The DECsystem 3100 is the server version of the DECstation 3100. Both have identical performance.

DECsystem 5400 Server

CPU: MIPS R3000, 20 Mhz.
FPU: MIPS R3010, 20 Mhz.
Main Memory: 64 MB
Disk: 1.2 GB RA90
Operating System: ULTRIX-32 3.1c rev 35
Compilers: cc, f77 V2.10

DECsystem 5810 Server

CPU: MIPS R3000, 25 Mhz.
FPU: MIPS R3010, 25 Mhz.
Main Memory: 64MB
Disk: RA70
Operating System: ULTRIX-32 3.1c rev 35
Compilers: cc, f77 V2.10

DECstation 5000 Model 200 Workstation

CPU: MIPS R3000, 25Mhz
FPU: MIPS R3010, 25Mhz
Main Memory: 16MB
Disk: 665 MB RZ56
Operating System: ULTRIX - 32 V3.1d
Compilers: cc, f77 V2.10

Note: The previous version of the Digital RISC compilers was V1.0 (internal ID was 1.31). The version 2.10 compilers will be available soon and provide enhanced performance over the previous version. This should be considered when comparing results from this summary with those obtained from systems running the V1.0 compilers.

Test Configurations (cont.)

Sun-4/260

CPU: MB86900 (SPARC) IU and MB86910 FPC, 16.67 Mhz.
FPU: Weitek 1164/1165
Main Memory: 8 MB
Disk: 327 MB SCSI
Operating System: SunOS 4.0
Compilers: SunOS 4.0 cc, Sun FORTRAN 12

SPARCstation 1

CPU: SPARC IU, 20 Mhz
FPU: Weitek WTL3170
Main Memory: 8 MB
Disk: 327 MB SCSI
Operating System: SunOS 4.0.3
Compilers: SunOS 4.0 cc, Sun FORTRAN 1.2

SPARCstation 330

CPU: SPARC IU, 25 Mhz
FPU: SPARC FPC and TI FPU, 25 Mhz
Main Memory: 32 MB
Operating System: SunOS 4.0.3
Compilers: SunOS 4.0 cc, Sun FORTRAN 1.2

HP9000/835S

CPU: HP9000/835s Precision Architecture
FPU: HP9000/835s Precision Architecture
Main Memory: 32 MB
Disk: 300 MB HP 7436
Operating System: HP-UX Version A.B3.10
Compilers: HPC HP92453-01A.03.14, FORTRAN 77/UK HP92430A.03.17

Apollo DN10000

Digital did not test this system. Results are from [DR 88] and DGI.

Raw Data Not Previously Supplied

SPEC

Elapsed time in seconds. SPARC results reported by Sun in [SPEC].

Benchmark	DS2100	DS3100	DS5400	DS5810	DS5000 Model 200	Sun-4/260	SPARC Station 1	SPARC Station 330	SPARC Server 490
GCC	196	136	135	153	83	150	139	108	69
Espresso	254	189	161	143	122	292	254	196	138
Spice 2g6	3418	2523	2633	2876	1747	5017	2915	2096	1458
Doduc	224	165	146	128	103	551	369	196	133
Nasa7	2068	1525	1569	2047	889	5878	1968	1439	1012
Li	641	473	510	529	284	683	690	553	318
Eqntott	125	98	81	73	60	133	113	88	63
Matrix 300	627	460	434	545	266	495	410	308	201
Fppp	340	243	227	216	138	1030	389	232	162
Tomcatv	363	268	268	291	153	628	445	324	215
Geometric Mean	467	343	328	342	209	678	459	327	220

DR Labs CPU2

Results are in seconds.

Apollo DN10000 numbers are reported in Digital Review, 12/19/88, "At the Speed of Light through a PRISM".

Benchmark	SUN 4/260	SPARC Station 1	SPARC Station 330	HP9000/ 835	DS2100	DS3100	DS5400	DS5810	DS5000 Model 200	Apollo DN1000
FFTR	12.31	9.03	5.94	6.0	5.36	4.00	3.31	2.68	2.64	2.88
LAHYDR	13.21	8.53	5.58	8.3	6.48	4.86	4.04	3.85	3.20	4.78
LABLE	8.97	5.94	4.34	6.4	5.39	4.03	3.34	2.67	2.68	3.94
LATRD1	12.53	8.48	6.33	7.2	4.74	3.53	3.59	3.52	2.32	5.27
LATRD2	14.21	10.02	7.49	7.4	6.34	4.69	3.91	3.39	3.11	7.87
LAPDE	19.51	12.46	9.24	9.3	6.91	5.16	4.27	3.70	3.42	4.65
LADIFF	13.84	9.25	7.34	18.1	10.64	7.79	8.70	10.14	2.62	8.38
LAFSUM	14.30	8.98	7.95	9.5	7.34	5.40	5.93	7.82	2.25	6.15
LA2DPP	14.96	9.52	8.55	8.7	9.57	7.14	5.94	4.78	4.73	7.50
LA1DPP	14.78	10.12	7.23	6.7	7.71	5.75	4.75	3.87	3.83	5.20
LA2DH	9.08	6.40	3.87	5.2	4.26	3.17	2.63	2.10	2.11	3.17
LAFIMI	16.63	9.47	9.06	10.2	12.26	9.14	7.56	6.07	6.06	11.72
LAIMCO	13.12	8.31	6.07	9.3	6.13	4.58	3.79	4.54	3.04	4.87
LAMAPR	9.29	6.62	4.60	5.7	3.55	2.61	2.79	4.37	1.68	4.58
LINPAC	11.55	9.11	5.58	5.7	8.28	6.09	6.32	5.47	2.96	4.18
AIRREL	5.47	3.42	1.40	4.8	2.36	1.76	1.45	1.16	1.16	5.65
DWHET	12.35	9.76	5.46	6.5	7.53	5.62	5.05	4.20	3.63	3.07
DOUBLE	13.75	9.59	4.35	5.1	5.27	3.93	3.26	2.61	2.60	3.08
EGYPT	15.39	12.23	6.33	6.2	11.84	8.85	7.31	5.87	5.85	5.23
EGYPTD	19.06	14.81	6.08	6.4	11.78	8.79	7.30	5.87	5.84	4.43
GAUSS	13.62	9.64	6.32	8.5	7.46	5.43	4.64	3.86	3.51	2.10
HANOI	13.11	11.74	9.30	10.7	11.62	8.62	7.47	8.29	5.59	7.65
NEFF	16.49	12.60	9.16	7.6	6.13	4.57	3.78	3.03	3.03	2.6
PHILCO	6.96	5.75	3.49	4.5	7.32	5.43	4.83	4.10	3.52	4.87
PRIME	25.87	21.59	15.73	31.0	21.45	16.05	13.30	10.61	10.62	7.52
RR2	16.90	12.15	7.47	8.1	11.93	8.78	7.48	6.37	5.53	4.32
RR3	12.68	8.76	6.60	6.4	8.23	6.21	5.14	4.25	3.95	5.94
RR4	16.18	11.37	8.05	13.0	15.80	11.59	9.93	8.71	7.28	8.88
SINGLE	12.69	8.95	4.38	5.6	6.26	4.67	3.86	3.09	3.09	4.30
SWHET	11.50	9.00	5.84	7.5	7.58	5.65	4.72	3.86	3.74	4.16
ALAM18	9.33	8.00	4.59	7.1	11.43	8.44	6.87	6.31	4.12	6.16
GAMSIM	14.36	10.51	7.40	10.5	6.11	4.55	3.76	3.01	3.01	2.92
LUDD	9.73	7.52	4.76	5.1	5.19	3.83	3.56	3.21	2.34	4.83
LUSD	9.98	7.38	5.03	5.2	4.94	3.69	3.19	2.98	2.36	5.92

References

[DR88]
ETA Explodes CPU Benchmarks", October 10, 1988, first use of CPU2 benchmark suite; "At the Speed of Light Through a PRISM", December 19, 1988, report on Apollo DN10000, Digital Review

[McMahon 86]
"The Livermore Fortran Kernels: A Computer Test of the Numerical Performance Range", December 1986, Lawrence Livermore National Laboratories.

[MDATOM]
Thompson, S., Cornell School of Chemical Engineering, "Molecular Dynamics Benchmarks", USENET, June 13, 1989

[MIPS 90]
"MIPS Performance Brief", Issue 3.9, January 1990

[NCAR-89]
Sato, R, "Performance Comparisons - Shallow Water Model", National Center for Atmospheric Research, Scientific Computing Division, Benchmark Update, 5/23/89

[PASHAPOUR]
Pashapour, Pope, Sepehrnoori and Shiles, "Application of Vectorization and Microtasking for Reservoir Simulation", Parallel Supercomputing: Methods, Algorithims and Applications, John Wiley & Sons, 1989

[SASI-90]
Swanson Analysis Systems, Inc., "ANSYS Benchmark Timing Results - Revision 4.4", Vendor Performance Summary, February 26, 1990

[SGL]
SGL is a proprietary contouring and gridding library developed by Dynamic Graphics, Inc. of Berkeley, CA

[SPEC]
Systems Performance Evaluation Cooperative Newsletter, Volume 2, Issue 1,Winter 1990

[Sun 89]
System Performance Overview, Preliminary Report, Sun MicroSystems, Inc. April 1989

Trademarks