i960® MICROPROCESSOR BENCHMARK REPORT

Order Number: 272950-001

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained from:

Intel Corporation
P.O. Box 7641
Mt. Prospect IL 60056-764
or call 1-800-548-4725

i960® MICROPROCESSOR BENCHMARK REPORT

Contents

FIGURES

TABLES

1.0 INTRODUCTION

This report describes the results of performance benchmarks for the Intel i960® microprocessor family.

The microprocessors covered in this report are:

- 80960SA, 80960SB, 80960KA, 80960KB, 80960CA, 80960CF, 80960JA, 80960JF, 80960JD, 80960HA, 80960HD and 80960HT.

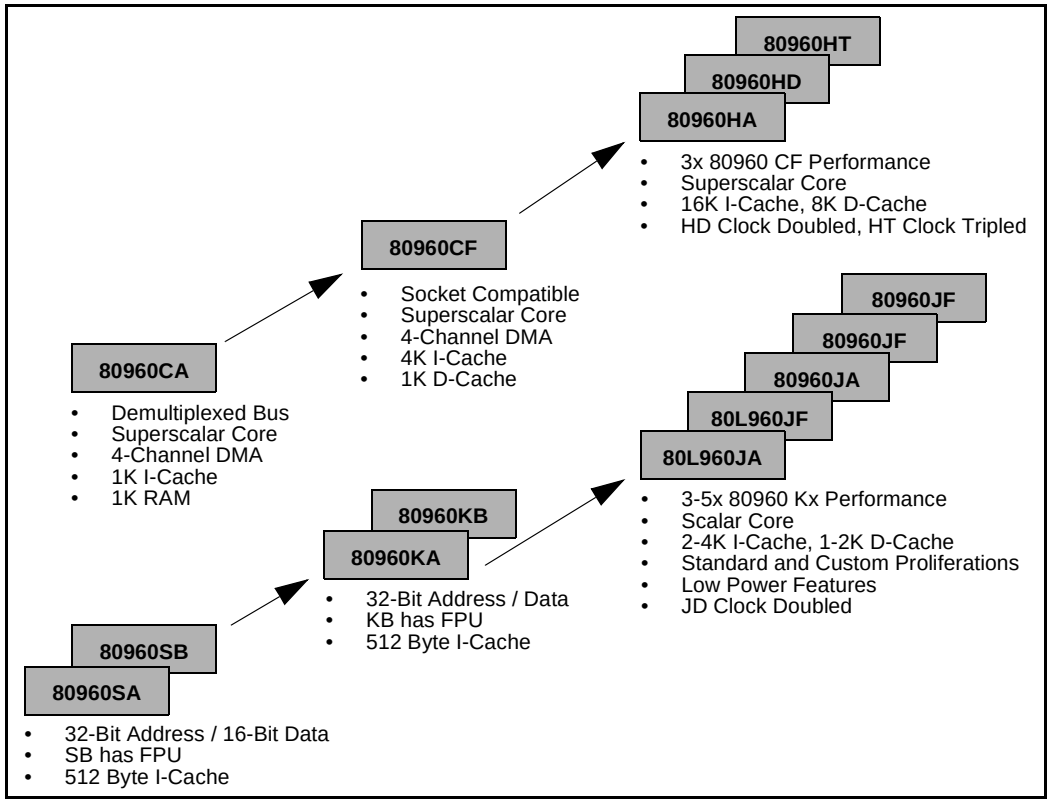

Figure 1. i960® Family Features Comparison

The basis of the comparison is a set of synthetic benchmark programs executed on all i960 microprocessors. This report makes every attempt to:

1. Use accepted practices in running the benchmarks.

2. Provide full and unambiguous disclosure of all factors that significantly affect the results.

Why synthetic benchmarks? The results of synthetic benchmark performance should not be the sole factor in selecting a microprocessor. The best indicator of performance is a customer's own benchmark program or — in the absence of that — one resembling the actual application. Developing a benchmark, however, costs time and money and may not be feasible for all customers. The customer may not have an application program developed when considering microprocessor performance. These constraints cause developers to turn to synthetic benchmarks for an indication of microprocessor performance.

Benchmark data in this report is presented using simple bar charts. Figure 2 is an example of such bar charts; it shows the relative performance of the present line of i960 microprocessors.

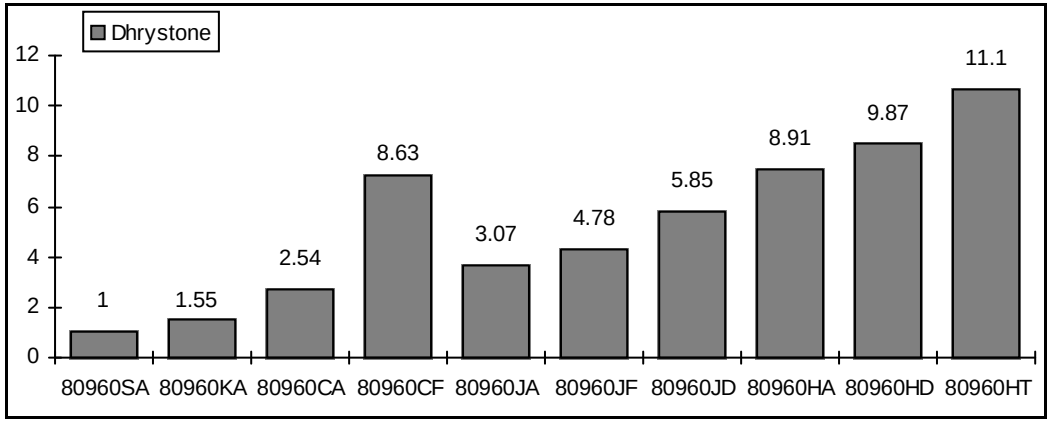

Figure 2. Relative Performance of i960® Microprocessors (not frequency normalized)

The raw results of each run are presented in tables appended to this document. These tables contain all data obtained from the benchmarking process. The exception to this is the Whetstone test. The Whetstone results were not included in the floating point average.

The magnitude of the 80960 Whetstone result may be misleading. The high value is a result of the aggressive optimization that is possible on some types of program structures with the Intel CTOOLS960 compiler.

Selection of programs used in this report was based on general acceptance, well-known behavior and the fact that each is written in the C programming language. Each makes an attempt at a relative measurement of the performance of the microprocessor/compiler combination.

The benchmark programs used include Dhrystone, MIPS, Whetstone, Stanford integer and floating point sections. For today's embedded applications two more meaningful benchmarks are also included: Ghostscript* provided by Aladdin Enterprises* and NET which is composed of three networking benchmarks.

The NET benchmark is proprietary code, but listed to show relative performance of the 80960 family in a networking application.

---

* Third-party brands and names are the property of their respective owners.

2.0 METHODOLOGY

There are few choices in selection of memory technology when designing a microprocessor subsystem. DRAM is the predominant read/write memory technology. ROM, Flash, and EPROM are predominant technologies for read-only (or read-mostly) memory subsystems. SRAM subsystems are often expensive beyond practicality, as code and data size of applications continue to grow beyond 1 Mbyte.

The system designer must work within the bounds of these available memory technologies. For example, if a system designer must use 70 ns DRAM technology, the relevant question of performance is — How fast can the microprocessor execute from 70 ns DRAM?

A 32-bit microprocessor memory subsystem typically resembles one of the three subsystems described in Table 1.

Table 1. Typical 32-Bit Microprocessor Memory Systems

Cost	Performance	Code Size	Description
High	High	Small to Moderate	Code and data are located in fast SRAM. SRAM provides access times on the order of 10 - 35 ns. Not a common design, but simple to implement. Practical for applications which have little code and are not cost-sensitive.
Moderate	Moderate to High	Moderate to Large	Code and data are located in DRAM. The code and initialized data are loaded from a backplane bus or inexpensive ROM at initialization. Mainstream, inexpensive DRAM technology typically provides 60 to 70 ns access time and fast page mode access capability. The system designer can trade off interface cost and performance using different degrees of complexity such as burst mode support and interleaving. Performance may also be enhanced in these systems with a small, fast SRAM dedicated to frequently accessed data.
Low	Low to Moderate	Large	Code is executed from ROM. Data is located in DRAM. The system designer can increase performance by interleaving the ROM subsystem. Since this design is driven by low cost, a DRAM subsystem is typically implemented at the lowest possible cost.

Microprocessor performance is influenced by several elements; categorized below as either intrinsic or extrinsic:

- Intrinsic elements generally are not or cannot be varied for performance analysis purposes. These properties are inherent to the microprocessor:
  - — Architecture and internal implementation (e.g., cache size, instruction set, registers)
  - — Efficiency of the external memory interface (e.g., instruction fetch bandwidth)
  - — Compiler efficiency (assuming that the best compiler is selected with the highest optimization turned on)
- Extrinsic elements can be varied depending on application requirements, performance and cost constraints:
  - — Clock speed
  - — Bus bandwidth (memory wait states)

Since many factors influence microprocessor performance, it is necessary to choose an equitable reference or baseline for a fair performance comparison. For this performance report, memory technology and clock speed have been chosen as the common reference for measuring performance of industry standard benchmark programs.

3.0 BENCHMARK PROGRAMS

Table 2. Benchmark Program Descriptions

Program	Description	Units of Measure
Dhrystone	Tests integer performance. String manipulation is a common action in this program. Version 2.1 is used here.	Dhrystones/second
MIPS	Measures instructions per second, based on Dhrystone v 2.1	Millions of instructions/second
Whetstone	Tests floating point performance	Millions of Whetstones/second
Stanford	Contains both integer and floating point sections. Uses a suite of well-known problems such as the towers of Hanoi and sorting algorithms	Stanford integer composite
NET	Composite of networking applications	Stanford floating point composite (Smaller is better)
Ghostscript*	PostScript interpreter using 'Scoop' postscript test page.	Seconds (Smaller is better)

4.0 DEVELOPMENT WORKSTATION/OPERATING SYSTEM

Host: IBM* RS/6000 (AIX* v 3.1)

The 80960 benchmarking environment includes an IBM RS/6000 (AIX v 3.1) machine with several 80960 Cyclone* boards attached.

4.1 Compiler/Assembler/Linker

The software used for this report is Intel GNU960 v5.0.

Compiler optimizations (of the compiler options available, for this report) used:

- For non-profiling, the -O4 option (the most "aggressive" option).
- For profiling, the two pass compile option (provides program-wide optimization).

Refer to the iC-960 Compiler User's Guide (651230) for further details on optimization options.

4.2 Target Hardware

All performance numbers were obtained using a Cyclone Evaluation Platform (Cyclone EP) with 8 Mbyte of 70 nS interleaved DRAM

The Cyclone EP is a stand-alone general purpose evaluation and development tool for Intel's family of i960 embedded processors. The main board provides the capability to install one of several i960 microprocessor modules. Using the different CPU modules allows one to evaluate the various i960 microprocessors in one system environment. This type of environment is most desirable when evaluating CPU performance.

Some of the Cyclone EP advantages are:

- Interchangeable i960 microprocessor modules (referred to as CPU modules)
- SIMM DRAM support up to 32 Mbyte
- Three 16-bit counter/timers
- Selectable processor clock frequency
- DRAM controller automatically optimizes wait states based on processor clock speed and memory speed

4.3 Memory Interface Configurations (specific to the 80960 Cyclone evaluation board)

Table 3. DRAM Access Times

Frequency MHz	Operation	DRAM Speed ns	Clock Cycles	Wait States x1,x2,x3,x4	Sustained Bandwidth[1] Mbytes/sec
16	Read	60, 70	3,1,1,1	1,0,0,0	36
20	Read	60, 70	3,1,1,1	1,0,0,0	45
25	Read	60	3,1,1,1	1,0,0,0	66
25	Read	70	4,1,1,1,1[2]	2,0,0,0	50
33	Read	60, 70	4,1,1,1,1[2]	2,0,0,0	66
40	Read	60	4,1,1,1,1[2]	2,0,0,0	80
40	Read	70	5,2,2,2,1[2]	3,1,1,1	53
16	Write	60, 70	3,2,2,2	1,1,1,1	25.6
20	Write	60, 70	3,2,2,2	1,1,1,1	32
25	Write	60	3,2,2,2	1,1,1,1	44.5
25	Write	70	4,2,2,2,1[2]	2,1,1,1	36
33	Write	60, 70	4,2,2,2,1[2]	2,1,1,1	48
40	Write	60	4,2,2,2,1[2]	2,1,1,1	53
40	Write	70	4,2,2,2,1[2]	2,1,1,1	58

1. Bandwidths stated are sustained bandwidths, not peak.
2. The extra cycle is the overhead of DRAM precharge. DRAM precharge time only impacts back-to-back cycles.

5.0 RELATIVE PERFORMANCE OF THE 80960 FAMILY (not frequency normalized)

The following graphs show the relative performance of the i960 microprocessor family. On one occasion performance numbers are not indicated due to timer issues specific to the platform and/or architecture used during the benchmark run.

GHOSTSCRIPT Postscript Interpreter:

Optimization:	Level 04	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Seconds	Smaller is Better

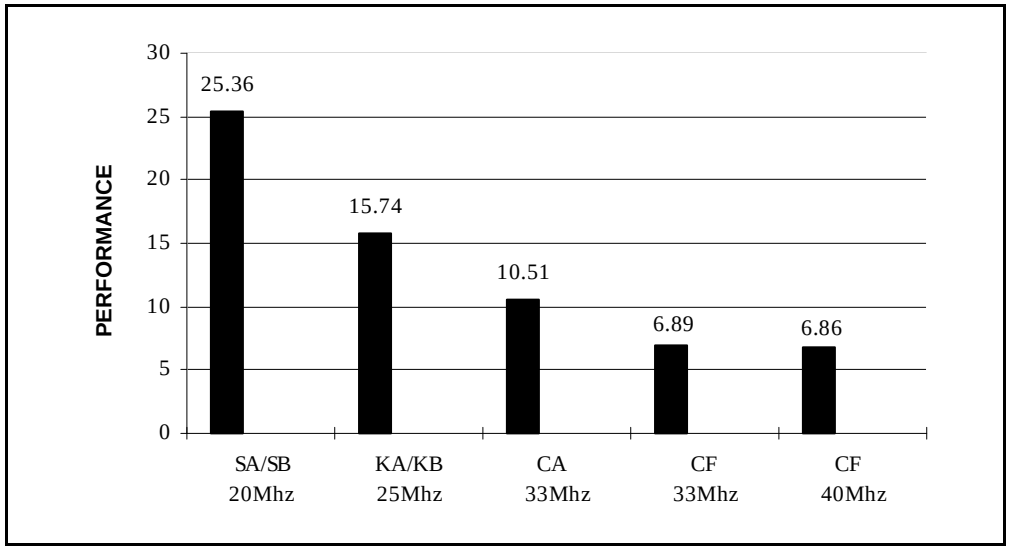

Figure 3. Ghostscript Relative Performance (Performance numbers based on Ventura Scoop PostScript file)

GHOSTSCRIPT Postscript Interpreter cont:

Optimization:	Level 04	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Seconds	Smaller is Better

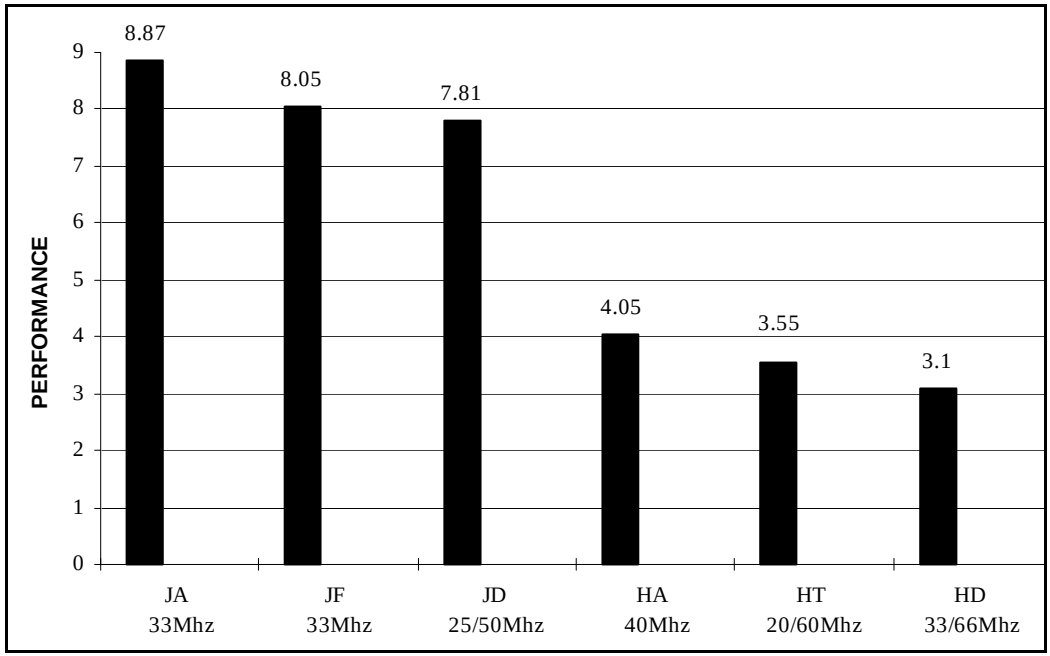

Figure 4. Ghostscript Relative Performance (cont)
(Performance numbers based on Ventura Scoop PostScript file)

5.1 NET:

Composite:	Three networking Benchmarks	
Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Seconds	Smaller is Better

NOTE: Sx note shown due to timer issue.

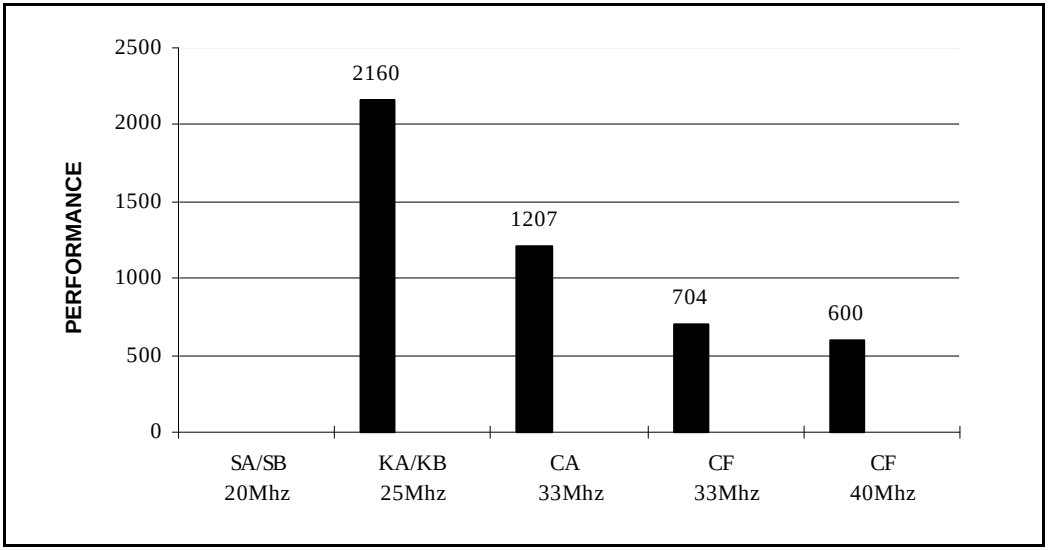

Figure 5. NET Relative Performance

NET cont:

Composite:	Three networking Benchmarks	
Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Seconds	Smaller is Better

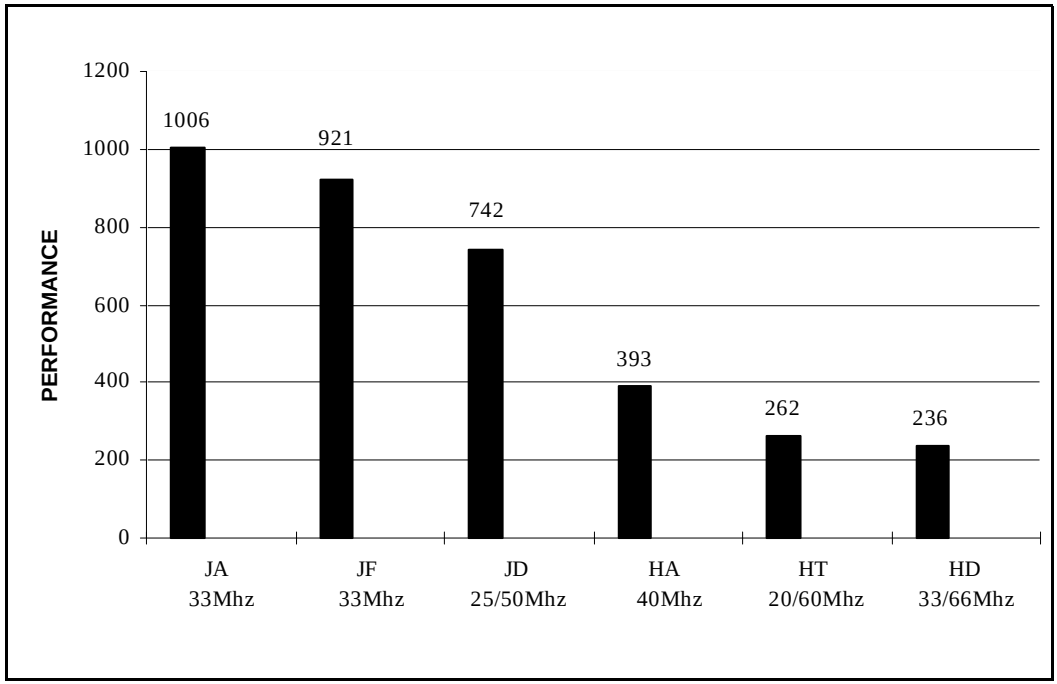

Figure 6. NET Relative Performance (cont)

5.2 DHRYSTONE Version 2.1:

Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Thousand/Second	Larger is Better

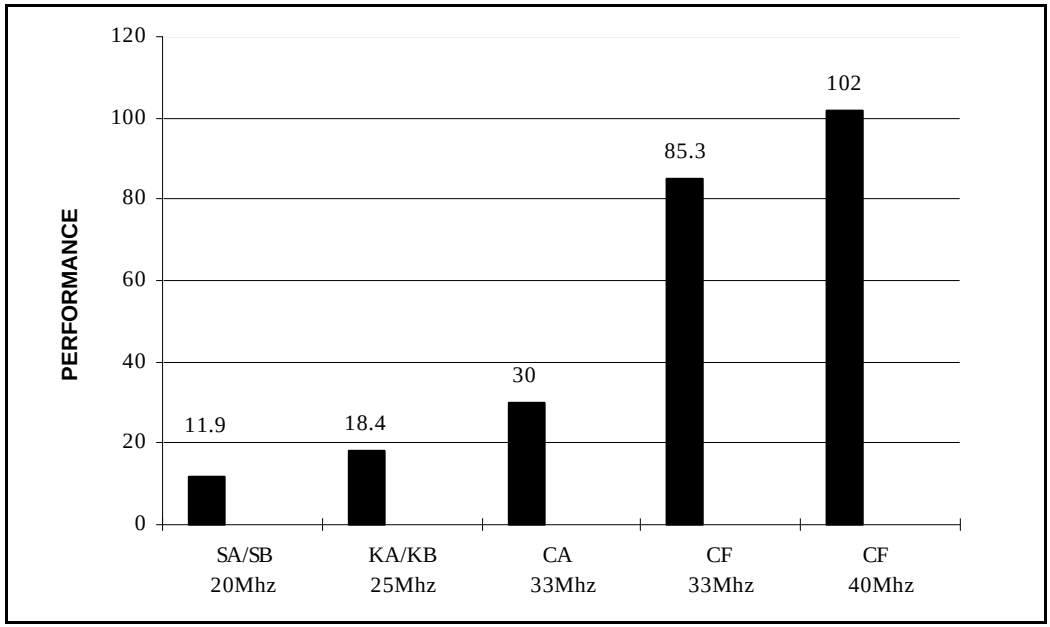

Figure 7. Dhrystone V2.1 Relative Performance

DHRYSTONE Version 2.1 cont:

Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Thousand/Second	Larger is Better

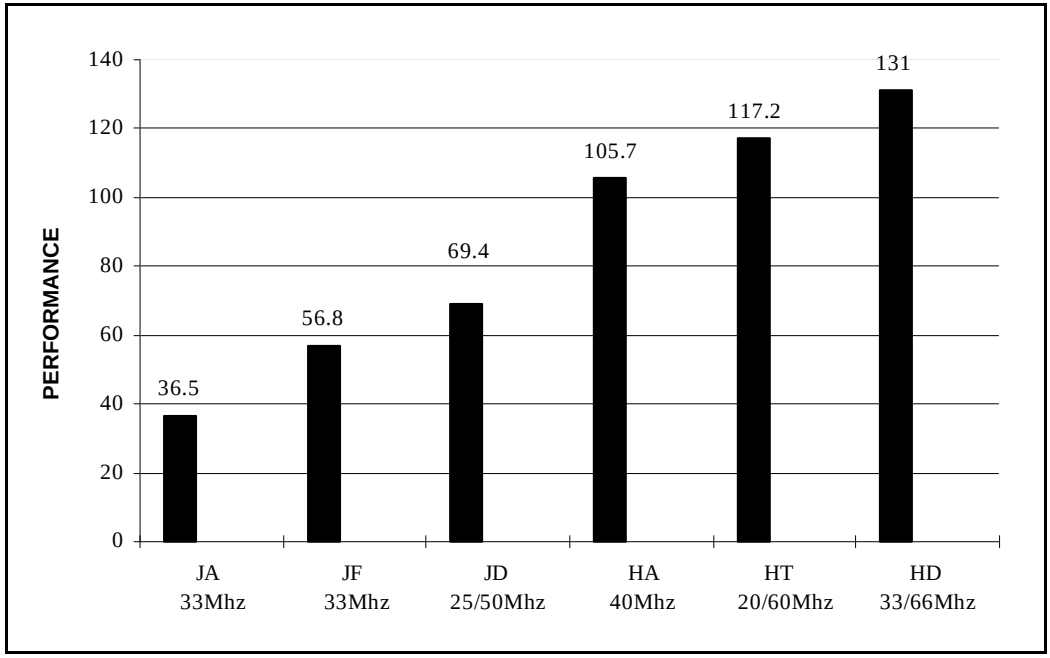

Figure 8. Dhrystone V2.1 Relative Performance (cont)

5.3 DHRYSTONE MIPS:

Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	MIPS	Larger is Better

NOTE: MIPS performance numbers are extrapolated from Dhrystone v 2.1 performance numbers and supplied for indication only.

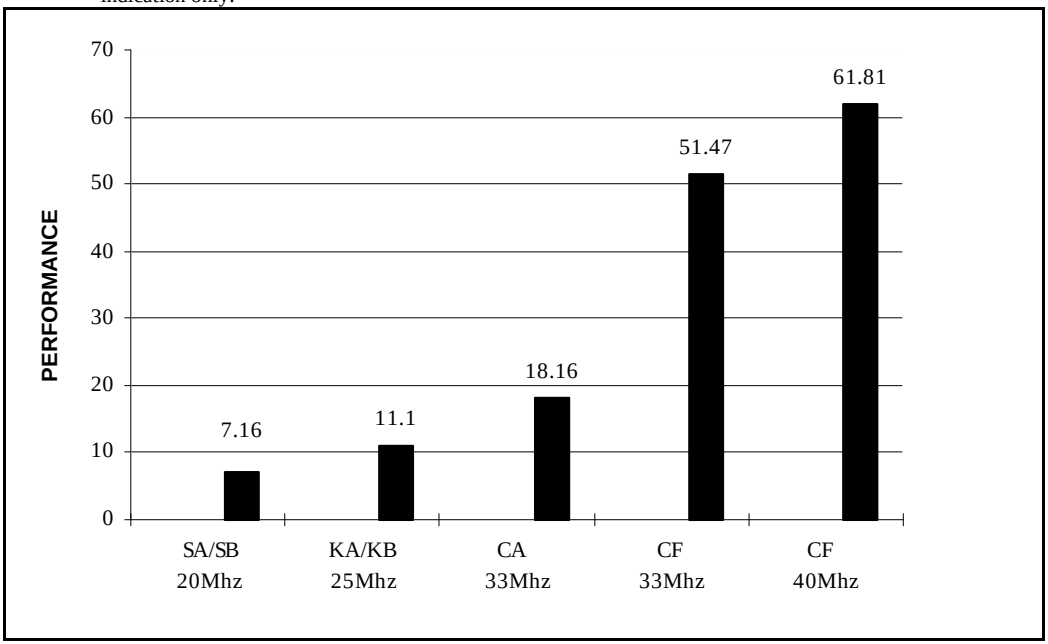

Figure 9. Dhrystone MIPS Relative Performance

DHRYSTONE MIPS cont:

Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	MIPS	Larger is Better

NOTE: MIPS performance numbers are extrapolated from Dhrystone v 2.1 performance numbers and supplied for indication only

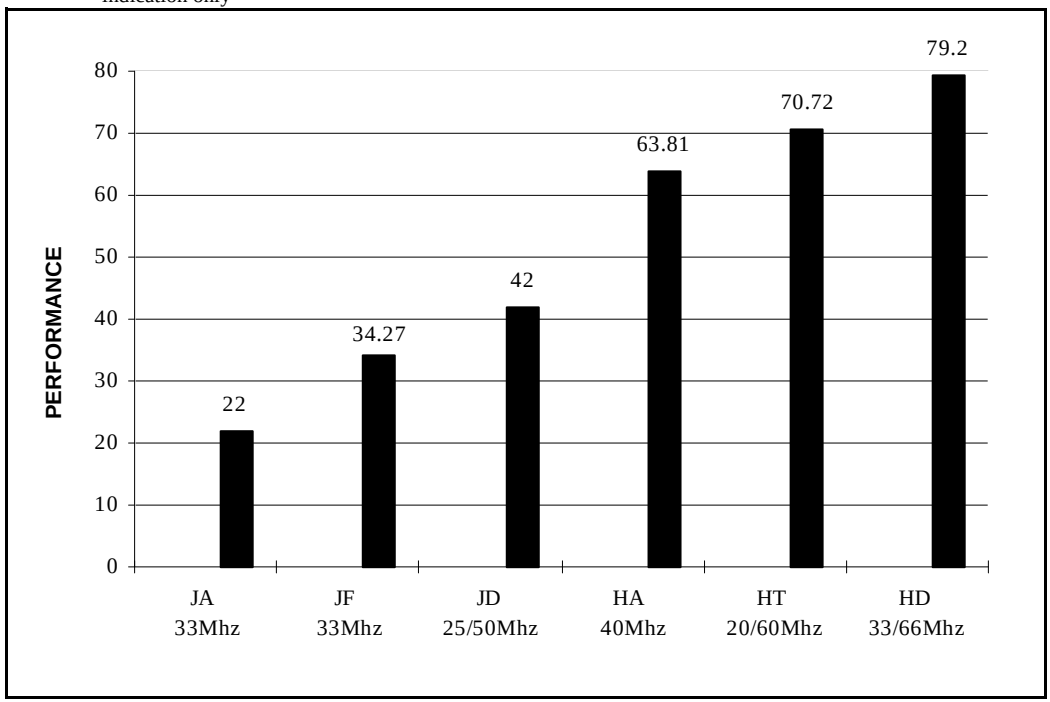

Figure 10. Dhrystone MIPS Relative Performance (cont)

5.4 STANFORD Non-Floating Point Composite:

Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Weighted Averages	Smaller is Better

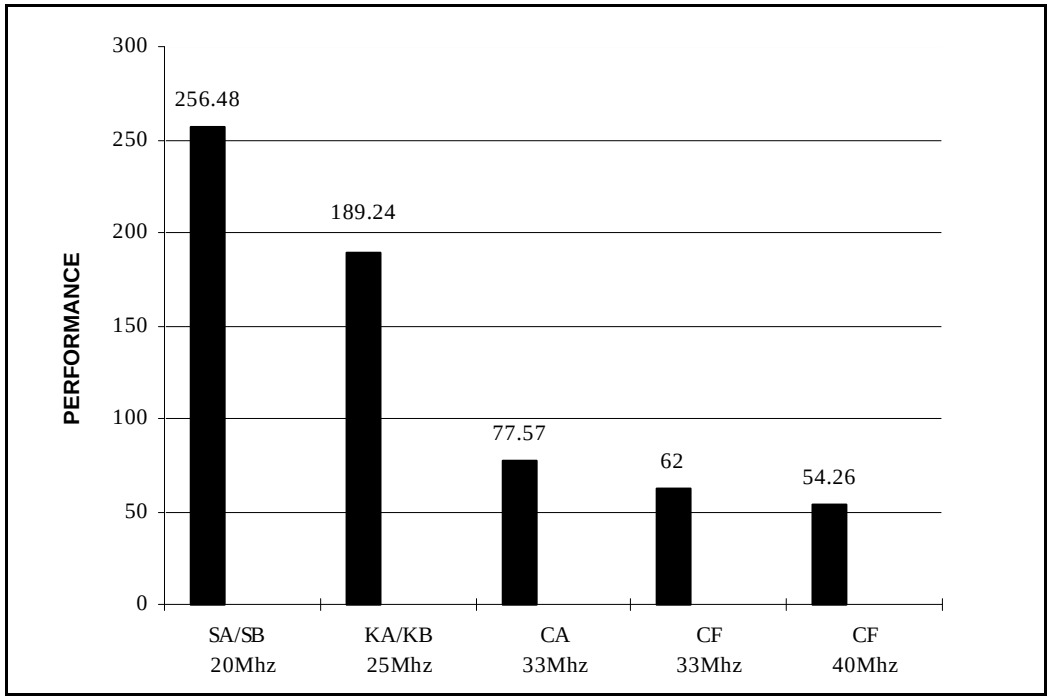

Figure 11. Stanford Non-Floating Point Composite Relative Performance

STANFORD Non-Floating Point Composite cont:

Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Weighted Averages	Smaller is Better

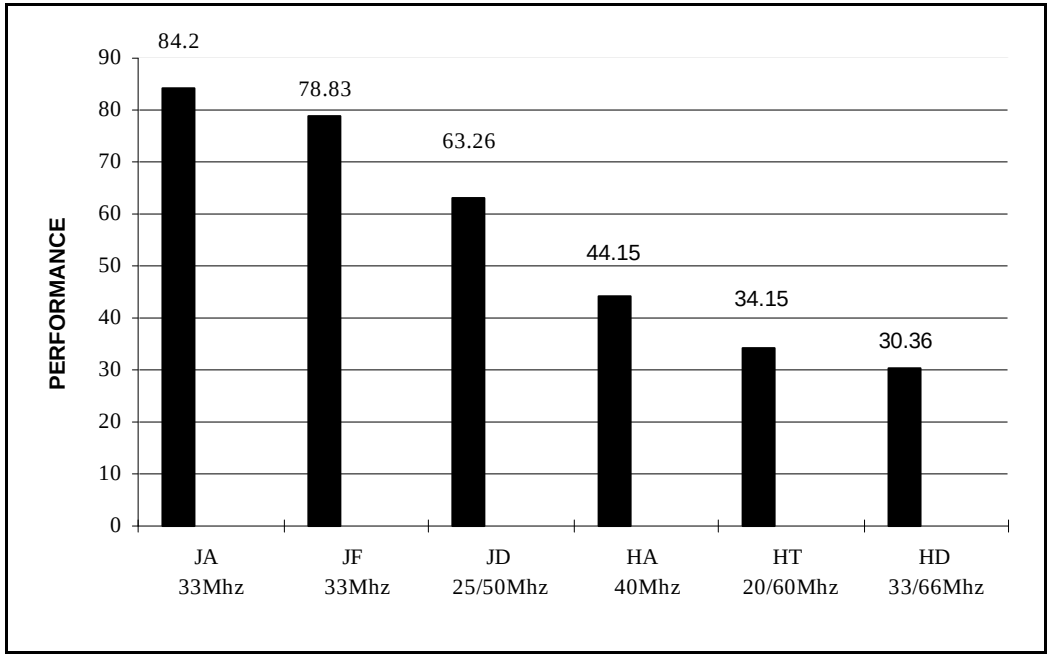

Figure 12. Stanford Non-Floating Point Composite Relative Performance (cont)

5.5 STANFORD Floating Point Composite:

Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Weighted Averages	Smaller is Better

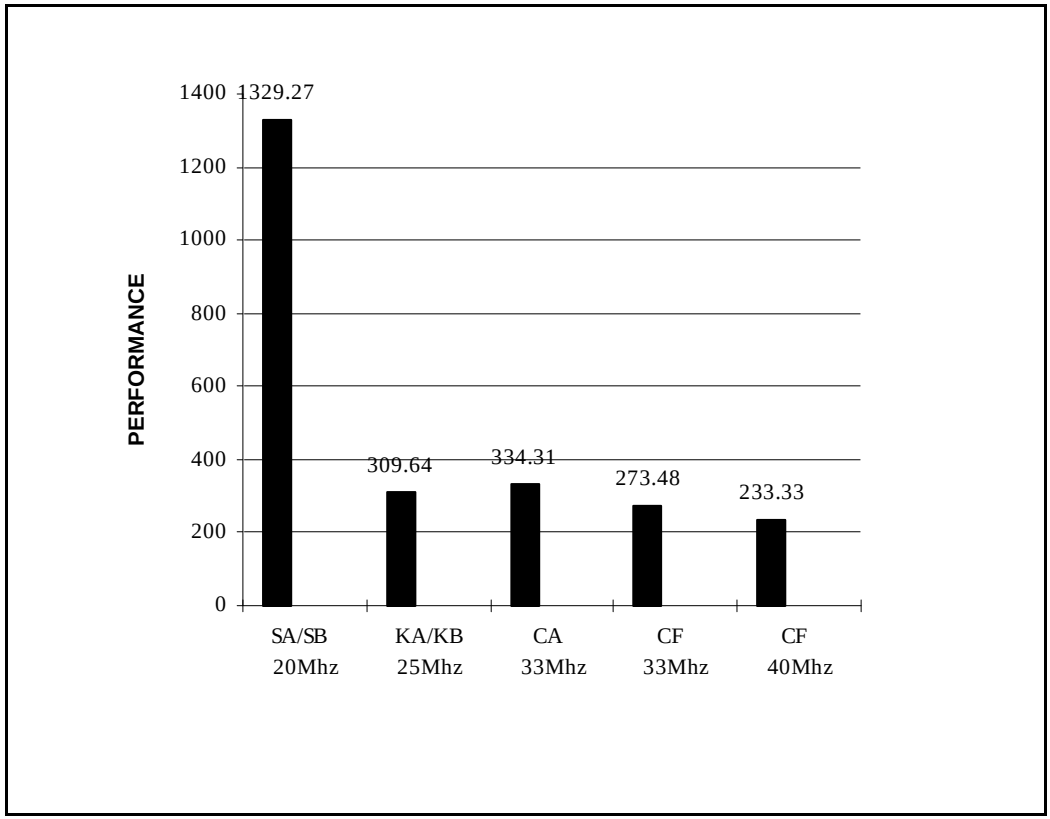

Figure 13. Stanford Floating Point Composite Relative Performance

STANFORD Floating Point Composite cont:

Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Weighted Averages	Smaller is Better

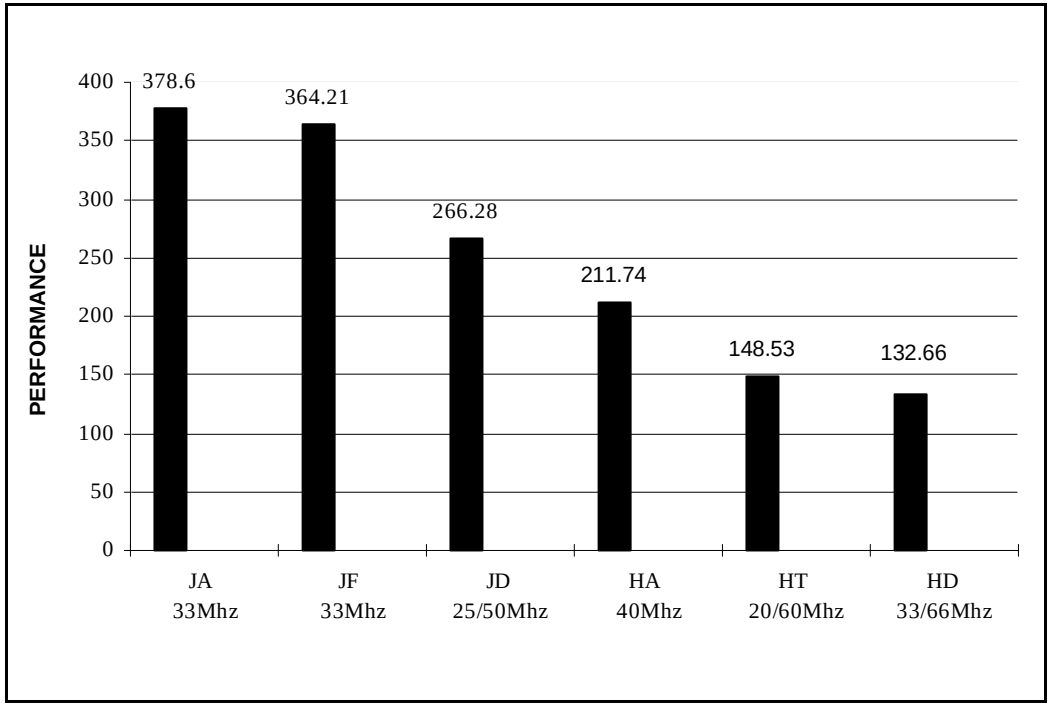

Figure 14. Stanford Floating Point Composite Relative Performance (cont)

5.6 WHETSTONE Single Precision:

Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Million/Second	Larger is Better

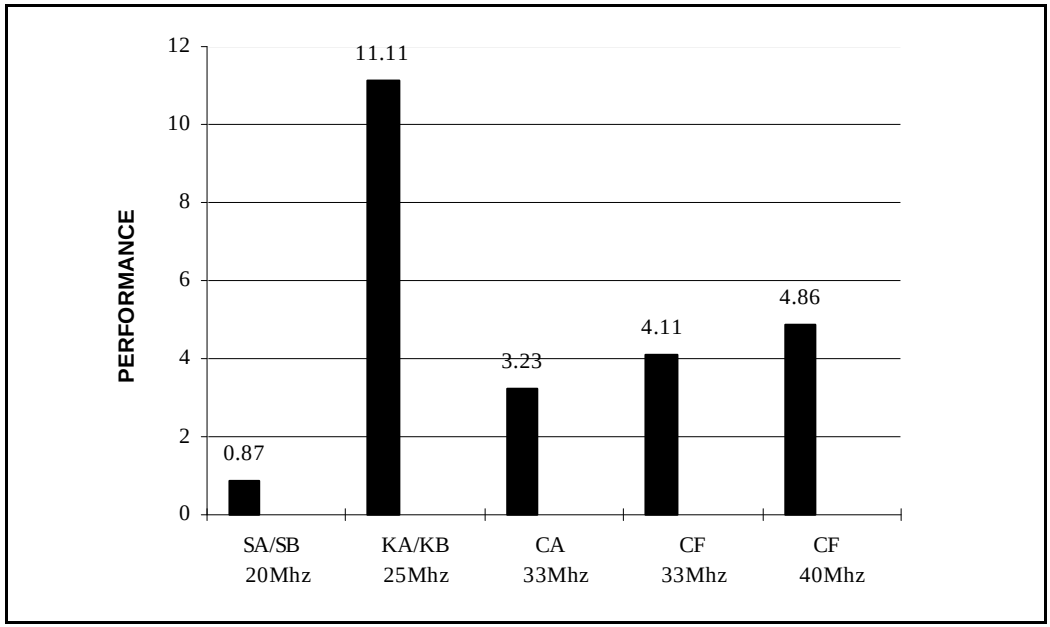

Figure 15. Whetstone Single Precision Relative Performance

WHETSTONE Single Precision cont:

Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Million/Second	Larger is Better

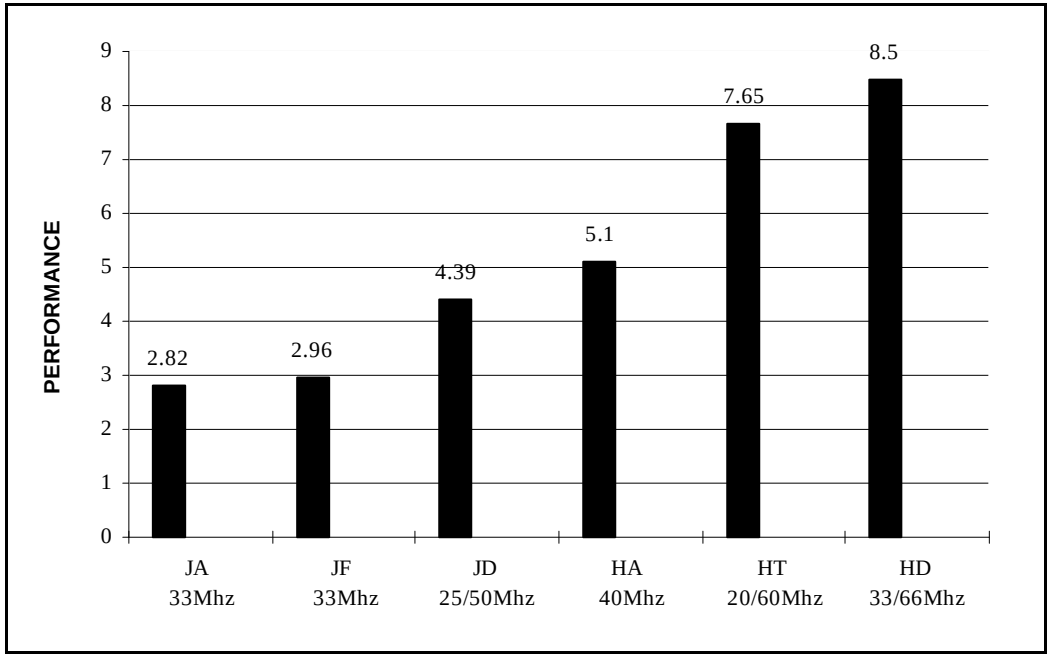

Figure 16. Whetstone Single Precision Relative Performance (cont)

5.7 WHETSTONE Double Precision:

Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Million/Second	Larger is Better

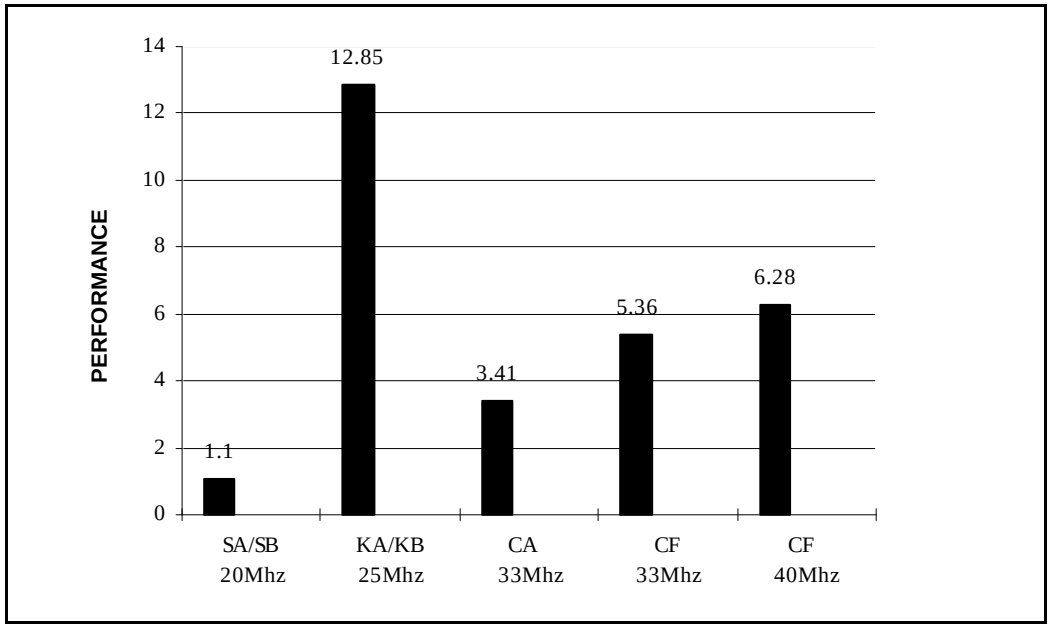

Figure 17. Whetstone Double Precision Relative Performance

WHETSTONE Double Precision cont:

Optimization Level:	Profiled	
Memory:	Interleaved 70 ns DRAM	
Wait State Profile:	20 MHz	10111
	25/33 MHz	20121
	40 MHz	31121
Unit of Measure:	Million/Second	Larger is Better

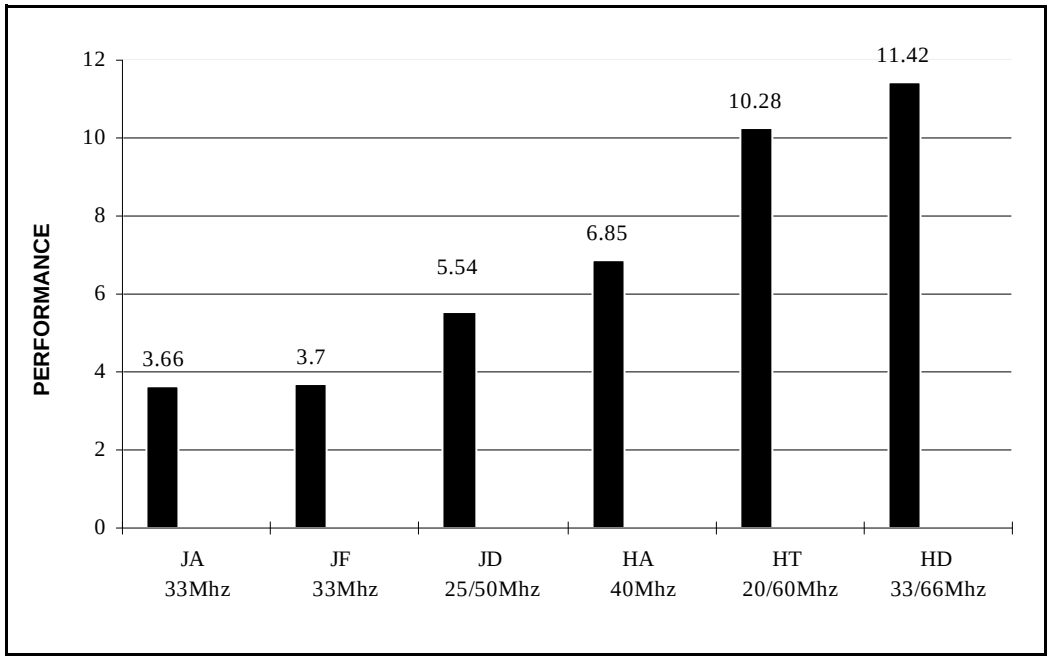

Figure 18. Whetstone Double Precision Relative Performance (cont)