

February 23, 1995

Revision 1.3

i960[®] JX MICROPROCESSOR BENCHMARKS

This technical bulletin compares the benchmark performance of Intel's i960 JA, JF and JD microprocessors with that of the i960 CA and CF processors.

Device Features

The 80960JA features a 2 Kbyte two-way set associative instruction cache and a 1 Kbyte direct-mapped data cache.

The 80960JF features a 4 Kbyte two-way set associative instruction cache and a 2 Kbyte direct-mapped data cache.

The 80960JD features a 4 Kbyte two-way set associative instruction cache and a 2 Kbyte direct-mapped data cache. The 80960JD core operates at a frequency which is double that of the bus.

The 80960CA features a 1 Kbyte two-way set associative instruction cache. The CA possesses a superscalar core design that enables it to issue multiple instructions in a single clock.

The 80960CF features a 4 Kbyte, two-way set associative instruction cache and a 1 Kbyte direct-mapped data cache. The CF possesses a superscalar core design that enables it to execute multiple instructions in a single clock.

Software Tools

Compiler/Assembler/Linker

Intel ic960 v4.5; Host: IBM* RS/6000 (AIX* v 2.5)

Compiler optimizations used: Of the compiler options available, this benchmark used:

- for non-profiling, the -O2 option (the most "aggressive" option). Used for the net test.
- for profiling, the -qp2 option (provides program-wide optimization). Used for the remainder of the tests.

Refer to the *iC-960 Compiler User's Guide* (484287) for further details on optimization options.

Source Level Debugger

Intel gdb960 v2.4; Host: IBM RS/6000

* Other brands and names are the property of their respective owners.

Benchmarking Hardware

The hardware used in performing the benchmarks is listed in the following table:

Table 1. Benchmarking Hardware

| Processor | Eval. Platform | Clock Rate | Wait States | RAM |
|--------------------------|---|------------|-------------|-------------------------|
| 80960JA/JF (DRAM) | Cyclone 80960 Evaluation Platform w/ 80960JX CPU Module | 33 MHz | 20-1-21 | 2 Mbytes of 70 ns DRAM |
| 80960JD (DRAM) | Cyclone 80960 Evaluation Platform w/ 80960JX CPU Module | 50/25 MHz | 10-1-11 | 2 Mbytes of 70 ns DRAM |
| 80960JA/JF (SRAM) | EV80960JX Evaluation Board | 20 MHz | 00-1-00 | 128 Kbytes of 8 ns SRAM |
| 80960JD (SRAM) | EV80960JX Evaluation Board | 40/20 MHz | 00-1-00 | 128 Kbytes of 8 ns SRAM |
| 80960CA/CF (DRAM & SRAM) | TomCA Evaluation Board | 33 MHz | 20-1-21 | 4 Mbytes of 20 ns SRAM |

Benchmarks

Table 1 contains a synopsis of each of the benchmarks used.

Table 2. Benchmark Program Descriptions

| Program | Description | Units of Measure |
|-----------|--|--|
| Dhrystone | Tests integer performance. String manipulation is a common action in this program. The version used here is 2.1. | Dhrystones/Second. A larger result indicates better performance. |
| Net | A composite of three benchmarks from our customers in the networking business. | Elapsed time, in seconds. A smaller result indicates better performance. |
| 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; Stanford floating point composite; These are weighted averages of each piece of the program. A smaller result indicates better performance. |

Results

The Table 3 shows the results for each benchmark test.

Table 3. Benchmark Results - DRAM

| Benchmark Program | 80960JA | 80960CA | 80960JF | 80960JD | 80960CF | Notes |
|-------------------------|---------|---------|---------|---------|---------|-------------------|
| Dhrystone | 47410 | 36504 | 51775 | 73569 | 77281 | Bigger is better |
| Net | 1065 | 1631 | 983 | 600 | 755 | Smaller is better |
| Stanford Integer | 85 | 73 | 84 | 58 | 57 | Smaller is better |
| Stanford Floating Point | 397 | 312 | 386 | 265 | 272 | Smaller is better |

Memory Impact on Performance

To indicate the impact of memory access performance on benchmark results the Dhrystone program was run at zero wait states for each of the processors. These zero wait state numbers are not included in the bar chart below.

Table 4. Benchmark Results - SRAM


| Benchmark Program | 80960JA | 80960CA | 80960JF | 80960JD | 80960CF | Notes |
|-------------------|--------------------|---------|--------------------|--------------------|---------|------------------|
| Dhrystone | 65327 ¹ | 57591 | 68927 ¹ | 88489 ² | 87998 | Bigger is better |

¹ Benchmark was run at 20MHz. Result was scaled upward to a value corresponding to 33MHz.

² Benchmark was run at 20MHz. Result was scaled upward to a value corresponding to 25MHz

Result Graph

The following graph reflects the DRAM performance of the 80960JA, JF and JD processors relative to the 80960CA and CF.



Summary

The four individual tests can be characterized by their sensitivity to the presence of a data cache in the processor being tested. In the cases of Net and Dhrystone, the 80960JF's 2 Kbyte data cache — in conjunction with a larger instruction cache — clearly distinguish it from the 80960CA. In the Stanford test suite, core performance is the issue, and the superscalar design of the 80960CA is a significant factor.