

From: TLE::KISWA::HEINEN "10-May-1988 1205" 10-MAY-1988 12:03
To: MCGREGOR,STAR::CARDOZA,TLE::GROVE,NYLANDER
Subj: FYI - Supnik on Prism

From: ROCK::SUPNIK "New location HLO2-3/C11 09-May-1988 2148" 9-MAY-1988 21:
To: SHARE::PALMER,CHIPS::DICKHUT,@[SUPNIK]STF,@[SUPNIK]SEGSTAFF,@[SUPNIK]AFL
CC:
Subj: uPrism Performance

+-----+
| d | i | g | i | t | a | l |
+-----+

INTEROFFICE MEMORANDUM

To:	Ken Olsen	Date:	9-May-88
	Jack Smith	From:	Bob Supnik
	Bob Palmer	Dept:	SEG/AFL
	Bill Strecker	Ext:	225-4947
cc:	Dave Cutler	Loc:	HLO2-3/C11
	Prism Team	Enet:	ROCK::SUPNIK
	STF		

Subj: What Is uPrism's Performance?

To date, uPrism's performance has been quoted in ultra-conservative terms. As was pointed out at the State of the Company meeting, this can lead to unexpected reactions by customers, and it has certainly led to reactions inside the company. This memo discusses uPrism's performance in the same terms of reference used by RISC chip vendors. In those terms, uPrism is a >30 mips RISC processor.


1 UPRISM'S FREQUENCY


One source of discrepancies is that uPrism's usually quoted operating frequency (40Mhz) is for WORST CASE process parameters, while other RISC chips' quoted frequency is for TYPICAL process parameters. The difference is as follows:

1. In a WORST CASE design, 99% of all chips that function will run at the stated frequency or HIGHER, and essentially none will run slower.
2. In a TYPICAL design, 50% of all chips that function will run at the stated frequency for faster, and 50% will run slower.

Or graphically (imagine a bell instead of peak):

percentage
of chips at
frequency
^





DEC microprocessor chips have always been quoted at worst case frequency, because until CVAX, there was only one system customer for each chip, and hence no users for slower parts. However, all new VAX and Prism chips planned for multiple systems; thus, it is now feasible to plan on "fast" or "typical" parts for the higher priced systems, and "worst case" parts for the lowest priced system. With this scenario:

system	range	uPrism frequency	% in bucket
Personal Prism	worst case	40Mhz	50%
Shrike	typical	50Mhz	30%
Osprey and Moraine	binned	58Mhz	20%

This reflects industry practice:

chip	range	operating frequency
Cypress SPARC	std	25Mhz
	binned	33Mhz
LSI Logic R3000	std	16.67Mhz
	binned	25Mhz

Thus, for purpose of comparisons with 33Mhz SPARC chips and 25Mhz R3000 chips, both of which represent chips at TYPICAL process parameters, the operating frequency of uPrism is 50Mhz.

2 UPRISM BENCHMARKS

The uPrism performance appraisal done at the time of the Crystal-Aquarius study was known to be conservative because:

1. The SIL compiler (a prototype) did not optimize its output code.
2. The SIL compiler did not schedule its code for uPrism.
3. The calling standard was still evolving and contained several features with highly negative performance impact.

More recently, Rich Witek went back and looked at two specific benchmarks, Dhrystone and Linpack. In addition to the problems listed above, he found that:

1. The SIL compiler generated many duplicate instructions.
2. The SIL compiler generated some pessimum code sequences for uPrism.

Rich took the assembly output of Dhrystone and made only the following changes:

1. Removal of duplicate instructions.
2. Simple scheduling for the uPrism pipeline.
3. Replacement of malformed code sequences.

In short, nothing beyond what the production Prism compilers will do. The results were dramatic. Instructions executed dropped by 24%, and benchmark performance improved by 80%. At typical operating frequency, uPrism's simulated performance was >76,000 Dhrystones.

For Linpack, having less time, Rich made only one change: he unrolled the DAXPY loop 4 times, and then let the compiler do its worst. Again, loop unrolling will be standard in the Prism compilers. At typical operating frequency, uPrism's simulated performance was 4.2 mflops. With macro BLAS, and code generator and calling sequence improvements, this number will improve.

3 PERFORMANCE ASSESSMENT

Since Dhrystone is a notoriously pessimal benchmark for VAX systems, I will compare uPrism to the Mips M/1000, said to be an "honest" 10 mips RISC system:

benchmark	uPrism	M/1000	ratio
Dhrystone	76k	22k-25k	3-3.4X
Linpack	4.2mflops	1.2-1.5mflops	2.8-3.5X

(The range of M/1000 numbers reflects differing levels of optimization in the M/1000 benchmarks.) Thus, I conclude that, with reference to other RISC chips, uPrism will deliver >30 RISC mips.

The key to better performance from Prism is better compilers (and, to a lesser extent, optimization of the calling standard). The techniques used above -- scheduling, global optimization, loop unrolling -- are standard in today's RISC compilers. Even further performance improvements are possible. For example, Prism is ideally suited to global allocation of registers at link time (as developed at WRL, performance improvements of up to 20% reported). Prism has a decided cycle time (peak mips) advantage over competitive RISC designs. With proper software support, this theoretical hardware advantage can be translated into superior performance for the customer.

/Bob Supnik