Hitachi Previews First PA-RISC Processor

HARP-1 Provides Good Performance But May Not be Marketed

By Linley Gwennap

Three years after licensing the PA-RISC architecture from Hewlett-Packard, Hitachi unveiled the first PA-RISC processor of its own design at the recent Hot Chips conference. The company believes its new chip, called HARP-1 (for Hitachi Advanced RISC Processor), should come close to the performance of HP's own processors. The Japanese firm, however, has no plans to broadly market its new device.

The Hitachi processor is designed to run at 120 MHz. It achieves this clock rate using 0.5-micron (effective) BiCMOS with four metal layers. The Hitachi design is two-way superscalar with dual integer units. At the conference, the company said its chip would exceed 70 SPECint92 and 110 SPECfp92; final measurements will be made after compiler tuning is complete.

Dual Integer Units

HARP-1 includes two complete integer ALUs, each with its own shifter, and a separate branch adder. Unlike SuperSPARC, which has the advantage of being able to issue two dependent integer operations per cycle, on HARP-1, the two operations must have no register dependencies. The integer ALUs also handle address calculations.

The floating-point unit has a pipelined adder and a pipelined multiplier. Both have a latency of four cycles. A third unit handles divide instructions using an iterative algorithm, taking 19 cycles for a double-precision calculation. Square roots are emulated in software.

The dispatcher can issue two FP operations each cycle as long as they use different units (for example, a multiply and an add). The PA-RISC FMPYADD instruction requires both the multiplier and the adder to be available; it can be used to issue two FP operations while still leaving one instruction slot for an integer operation. As with integer instructions, two FP instructions must have no data dependencies to be issued on the same cycle.

HARP-1 can perform a single load or store on each cycle. Adding a second load pipeline, as in TFP (see page 9) would have required dual-porting both the data cache and TLB, significantly increasing die size. As shown in Figure 1, either ALU can generate a 48-bit virtual address for the data cache. The integer unit can load or store one 32-bit value at a time. (PA-RISC registers are 32 bits.) The floating-point unit supports single-, double-, and quad-word loads and stores, requiring a 128-bit path to the data cache.

The chip can process one branch per cycle. If the branch is the first instruction in a pair, it cannot be issued with the following (delay slot) instruction, although a branch can be paired with its predecessor. This simplifies the control logic; only one of the two instruction buses is routed to the branch unit. Compound PA-RISC instructions such as compare-and-branch and add-and-branch follow the same issue rules as simple branch instructions.

Strict Alignment Rules

Although the instruction dispatch is quite flexible, strict alignment rules apply. Instructions are read from the instruction cache two at a time; if they cannot be issued together, they are issued one at a time. Most other superscalar processors use an instruction buffer that allows the second instruction to be paired with its successor if the first must be issued alone.

HARP-1 is more restrictive because it performs the dependency check as instructions are loaded into the on-chip cache. The instruction cache contains an extra bit for each pair of even and odd instructions; the bit indicates whether the pair can be issued together. As shown in Figure 2, this allows instructions to be dispatched during the "IF" stage rather than adding a separate pipeline stage, as is done in the Alpha CPU, for example.

The execution pipeline is quite similar to the PA7100 and other traditional five-stage RISC pipelines. HARP-1, however, adds the "N" stage because, at 120 MHz, there is not enough time to complete the data-cache access and check the cache tags (and TLB) in the

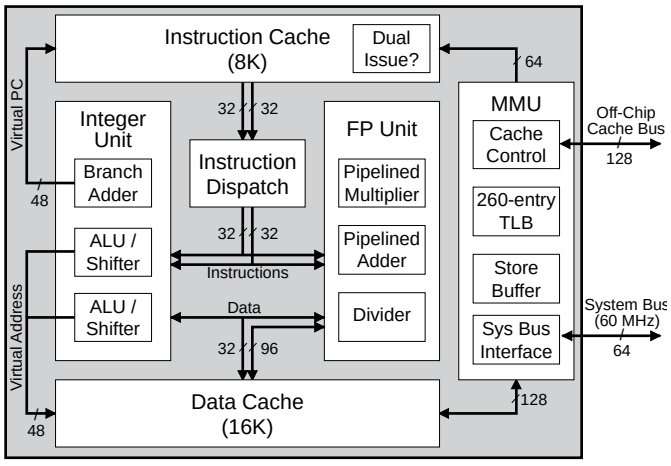

Figure 1. HARP-1 includes dual integer units, one load/store unit, and separate pipelined floating-point units.

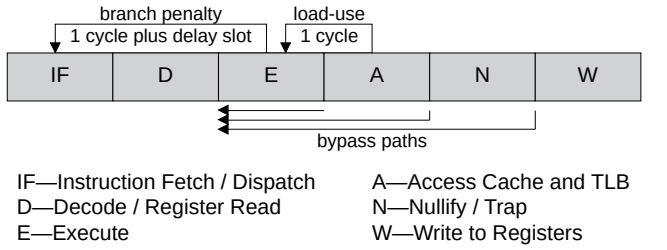

Figure 2. HARP-1 uses a six-stage pipeline that is similar to traditional five-stage pipelines except for the "N" stage.

"A" stage. The added stage provides an extra cycle before the results are committed to the register file. Since this stage is near the end of the pipeline, it does not impact the load-use or branch penalties, but it does increase the complexity of the bypass logic.

Split-Level Cache New for PA-RISC

The on-chip caches are a significant departure from HP's PA-RISC design strategy. HP processors have always relied on large primary caches implemented using external SRAM. Using advanced design techniques at both the chip and board level, the PA7100 is able to achieve single-cycle access to 512K of cache at 99 MHz using 9-ns SRAMs. This strategy has limited the clock speed of HP's chips (although not as much as some would think) and requires the use of expensive cache memory.

Hitachi licensed HP's architecture but chose a different cache design. To reach 120 MHz, it moved the primary caches on-chip to eliminate chip-to-chip delays. Even with an advanced IC process, there is room for only 8K of instruction cache and 16K of data cache on the chip. Both caches are direct-mapped and parity-protected; the instruction cache has a line size of 32 bytes, while the data cache uses a 16-byte line. Both caches are virtually indexed and physically tagged. The data cache uses a write-through protocol to simplify coherency with the external cache.

Like the on-chip cache, the on-chip TLB is split between instructions and data. Each of the two TLBs contains 128 page entries and two block entries that can map large portions of memory. Most recent processors implement fully associative TLBs, but Hitachi has chosen a simpler direct-mapped scheme. The large size of the TLBs compensates for the lower hit rate caused by the lack of associativity.

HARP-1 connects directly to an external, second-level cache, providing all of the necessary control signals. The external cache is unified (instruction and data), direct-mapped, and ECC-protected. It is physically indexed and physically tagged, with a line size of 32 bytes. A write-back protocol reduces traffic on the system bus. The chip supports cache sizes from 256K to 4M. The cache does not include multiprocessor support.

The external cache can be accessed in two processor cycles. A data-cache miss takes only a single access to refill a cache line, while an instruction-cache miss requires two accesses, or four cycles. Hitachi did not specify the SRAM access time, but with a 17-ns cache access cycle, it should be possible to use 12-ns memory chips.

One unusual feature is a second-level TLB that is combined with the external cache tags in the same SRAMs. This TLB is unified and can hold 2K to 32K entries, depending on the size of the cache. This two-way associative TLB is accessed directly by hardware on a first-level TLB miss. The miss penalty is 14 cycles if the translation is found in the first set or 18 cycles if it is in the second. If the translation misses the second-level TLB, the CPU generates a TLB-miss interrupt to request software intervention.

A separate, 64-bit bus connects the processor to main memory and the rest of the system. This bus operates at one-half of the processor frequency. HARP-1 implements a six-entry store buffer to avoid delays while writing data from the external cache to main memory. Hitachi's presentation did not discuss the specifications for the system bus.

An Expensive Chip

The separate, wide buses for cache and main memory data forced Hitachi to select a 595-pin ceramic PGA for HARP-1. This package is larger (and probably more expensive) than the packaging for any other announced microprocessor. The die itself measures 16.2×16.5 mm; at 267 mm^2, it is larger than even SuperSPARC, although somewhat smaller than Pentium or TFP . The design packs 2.8 million transistors into this space using an advanced BiCMOS process with four layers of metal. Hitachi quotes an effective transistor length of 0.5-micron, which is roughly equivalent to a 0.65-micron (drawn) process. The gate oxide thickness is just 90 Å, which should create very fast transistors.

Given the aggressive process, the die seems unusually large. The die size and transistor count are roughly equivalent to those of SuperSPARC and Pentium, but those two chips use lower-density 0.8-micron BiCMOS processes. IBM's PowerPC 601 uses a 0.5-micron (effective) process with four layers of interconnect to put 2.8 million transistors on a die just 121 mm^2, less than half the size of HARP-1. IBM cheats a bit by using a fifth metal layer to put the pads on top of the die, but this does not account for the huge difference in size.

Cranking these figures through the µPR Cost Model *(see **071004.PDF**)* results in an estimated manufacturing cost of $650 for the Hitachi chip, the highest yet for any current or announced single-chip processor. Only Pentium and the TFP chips are larger, but HARP-1 has a higher estimated wafer cost due to the complexity of its BiCMOS process. The package is also expensive and the projected volumes are low.

With its large transistor count and 120-MHz clock, it isn't surprising that HARP-1 dissipates 20 watts at 3.3V. While high, this is within the range of modern cooling techniques, as demonstrated by the 30-watt Alpha CPU.

Going for the High End

Hitachi's figures show that HARP-1's integer performance is slightly less than the 80.0 SPECint92 posted by HP's PA7100. Hitachi's processor has a 20% frequency advantage and also gains from its dual integer ALUs; the PA7100 can issue only one integer instruction per cycle. The HP chip has a major advantage due to its 512K, single-cycle external cache; HARP-1 has just 24K of on-chip single-cycle cache and takes an extra cycle to access its external cache. Hitachi's cache design offsets the performance gains from its higher clock rate and dual integer units.

Hitachi's figures are also lower than performance ratings for the DEC Alpha and MIPS R4400 processors. The Alpha chip is two-way superscalar and runs at up to 200 MHz. The R4400, at 150 MHz, is also faster than HARP-1 and has slightly more on-chip cache.

The HP, DEC, and MIPS chips are all currently in production. Hitachi has not announced availability for the chip but is currently testing the first silicon; this means that the chip will probably reach production in 1H94. By that time, new chips from these vendors may be available at even higher performance levels.

Hitachi has also not indicated a price for its processor, if indeed it is sold at all. Given our manufacturing cost estimate, it appears that HARP-1 could be sold for around $1000–$2000, which is at the very high end of current microprocessor prices. There are no system logic chips available for this processor, forcing potential customers to design their own.

Hitachi refused to comment on whether it will sell HARP-1, and its American subsidiary says it has no plans to market the chip in the US. Hitachi may choose to use it only in its own PA-RISC systems, reducing its reliance on HP processors. Infocorp estimates that Hitachi sold only a few hundred PA-RISC systems in 1992, making this a very small business to support an entire microprocessor development effort, although Hitachi certainly plans to expand its volumes in the future.

If Hitachi chooses to keep its PA-RISC processors to itself, it would be disappointing. The formation of the Precision RISC Organization last year indicated an intent to open the PA-RISC architecture, a promise that so far lies unfulfilled *(see **0710ED.PDF**)*. HP is not willing to offer its chips on the open market, leaving Hitachi as the best hope for promoting PA-RISC to other system vendors. Perhaps we must wait a bit longer for a high-performance, merchant-market PA-RISC chip.

HARP-1 is an impressive result for Hitachi's first microprocessor using a new architecture. Unfortunately, the circuit design and layout appear inefficient, resulting in a die that is too large to be cost-competitive with other leading microprocessors. The name indicates that the company plans to continue this line of processors; having one good design under its belt, Hitachi could deliver a true industry leader with HARP-2. ♦

More PA-RISC Processors Due

Two new PA-RISC processors will be announced in October but neither will come from HP, the originator of PA-RISC. Unlike HARP-1, both of the new chips are aimed at the low end of the cost scale and are intended for embedded applications.

Hitachi's low-cost controller is named PA/50. The second chip is from Oki Semiconductor, another PA-RISC licensee. Details on these chips will be provided at the upcoming Microprocessor Forum.

