Oki, Hitachi Prepare Embedded PA-RISCs

First Chips Using HP's Architecture to Hit Merchant Market

by Curtis P. Feigel

For the first time, processors built on HP's PA-RISC architecture may be available on the merchant market. Two new flavors of PA-RISC, from Oki and Hitachi, were revealed at last month's Microprocessor Forum. Clearly targeted at the lower end of the PA-RISC market, the new processors are more integrated than HP's offerings, filling a low-cost niche in the PA-RISC line. Both devices achieve performance befitting a high-end embedded controller, or even an entry-level personal workstation.

The Oki processor is an elegantly simple integer unit with some added interface features to ease its integration into a system. The Hitachi device uses a similar core but with added floating-point and memory-management units. Both use the traditional five-stage RISC pipeline and incorporate on-chip caches for data and instructions—a departure from HP's versions that rely on external primary caches. Integer applications may be hard-pressed to distinguish the two devices.

Oki's Integer RISC For Embedded Systems

The OP32/50N's purpose in life is high-end embedded systems, so Oki concentrated on integer capabilities. As shown in Figure 1, the design does not include an FPU. The processor implements the PA-RISC 1.1 architecture Level-0, which is an absolute memory system, providing 32 bits of physical address space but neither address translation nor memory management.

The Oki processor has a feature designed to enhance interrupt handling: A cache-freezing scheme keeps critical code and data resident in cache, regardless of cache misses. Software may freeze the instruction cache in 2K partitions, and the data cache in 1K partitions. Both caches are two-way set associative, but a frozen partition lies completely within one or the other of the cache's sets. The corresponding partition in the other set continues to operate, but as a direct-mapped cache. This capability (first seen in Fujitsu's SPARClite) is important in real-time situations where guaranteed response time is imperative.

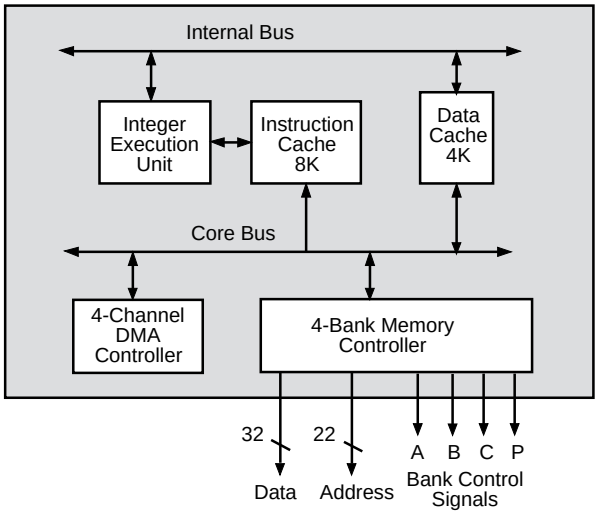

Figure 1. The OP32/50N includes a 4-channel DMA controller, and four memory-bank controllers. Each bank has a programmable data-bus size and a glueless interface to most kinds of memories and peripherals found in embedded systems.

Glueless Memory Interface

The processor segments its 4-G address space into sixteen regions of 256M. Using a control register, software may assign each region to one of four external bank controllers, called banks A, B, C, and P. The region at the top of the address space, Region 15, is permanently assigned to bank P.

The chip provides 22 external address lines. Addresses are generated internally as 32-bit values, then

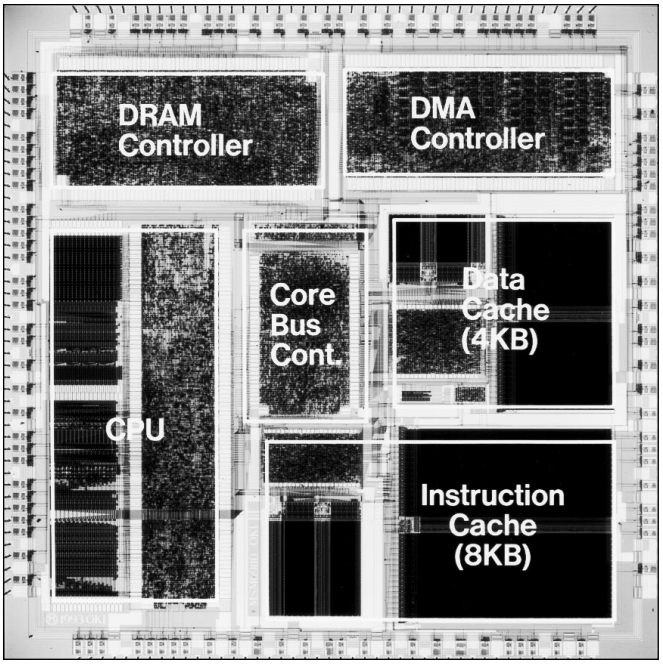

Figure 2. Oki's OP32/50N uses 1.1 million transistors on a 14.3 x 14.3 mm die in a 0.8 µm CMOS process.

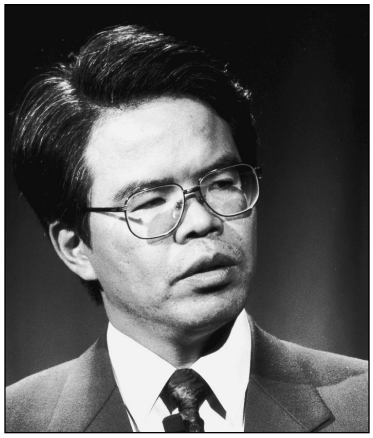

MICHAEL MUSTACCHI

Noriyoshi Ito shows off Oki's handiwork in the OP32/50N at the Microprocessor Forum.

multiplexed onto the external bus. The designer can set each bank controller so the addresses are multiplexed to directly handle certain kinds of devices. Bank A connects to ROM or SRAM. Banks B and C handle SRAM or DRAM. Bank P is reserved for peripherals and I/O.

For example, Bank B and Bank C can each be configured to provide multiplexed addresses for DRAMs from 256K to 16M in size. These banks generate CAS and RAS signals to control the DRAM and handle refreshing. Standard memory accesses take place at up to 66 Mbytes/s (four bytes in two cycles at 33 MHz); bursts access two words in three cycles at speeds up to 88 Mbytes/s (eight bytes in three cycles at 33 MHz).

The data width is programmable, giving the processor access to bytes, halfwords, and words. This not only simplifies connection to a byte-wide boot ROM in Bank A, but lets the designer trade performance for lower cost by using DRAM banks that are 16 bits wide (or smaller) and 8-bit-wide peripherals.

The OP32/50N also has four independent DMA channels, each of which can operate in either single- or multiple-address mode and can be programmed to handle 8-, 16-, or 32-bit-wide transfers. DMA transfers may involve bursts, as well as cycle stealing, for which the DMA channel uses the external bus while the processor core executes using the internal caches. Channel priorities may be fixed or run in round-robin fashion.

The DMA controller implements two particularly useful features. Multiple-block transfers allow the software to set up the next block while transfer of one block is underway. The next transfer is then begun immediately as the preceding block completes. The second feature, called data-match completion, allows the device to transfer data until a specified data value is detected.

The processor decodes 32 interrupt levels from five input pins. It also has a 32-bit interval timer (which runs at the processor's clock speed) and a 32-bit recovery counter, used as watchdog timer. Branch Trap and Break instructions are provided to aid debugging.

Hitachi Offers Workstation Power

Hitachi's PA/50L processor goes one step beyond Oki's to implement Level-1 of the PA-RISC 1.1 architecture, which supports 48-bit virtual addresses. This capability, combined with the device's integer and floating-point execution units, means the PA/50L packs enough power to support simulation and engineer-ing applications common on personal workstations. As shown in Figure 3, the Hitachi chip includes primary caches but lacks the controllers and peripherals that would be useful in embedded applications.

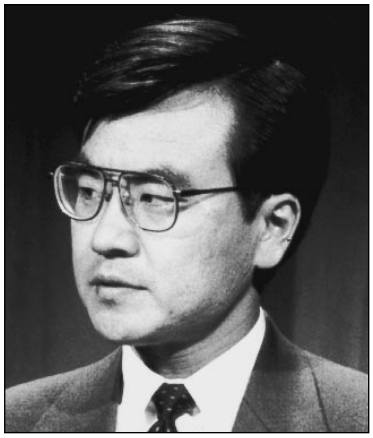

MICHAEL MUSTACCHI

Kunio Uchiyama, at the Microprocessor Forum, reveals the inner workings of Hitachi's PA/50L.

Like Oki's device, the PA/50L is not superscalar; it fetches one instruction per clock, which is then dispatched to the proper unit for decode and execution. Its integer unit has an ALU, an SMU (shift/merge unit), and a set of 32 general-purpose, 32-bit registers that, with the decode and control logic, form a five-stage pipeline. The floating-point unit has separate adder and multiplier sections, and a set of 28 FP registers, each 64-bits wide.

The instruction cache is 8K, while the data cache is 4K and uses a write-back protocol. Both caches are two-way set associative. Software running on the PA/50L can take advantage of a prefetch instruction that can load data into the cache before it's needed. Each cache has a 32-byte incoming block buffer, arranged as four lines of eight bytes each. Here, external 32-bit-wide accesses are converted to match the internal 64-bit width. If the block

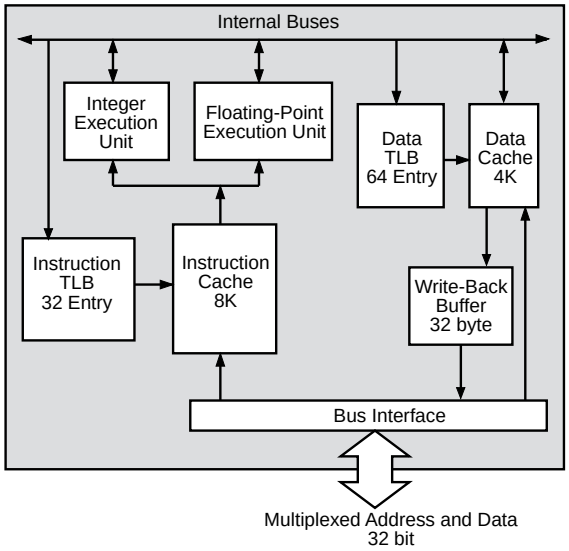

Figure 3. Hitachi's PA/50L has built-in support for synchronous DRAM, big-endian/little-endian code and data, and a 1/8-speed power-save mode.

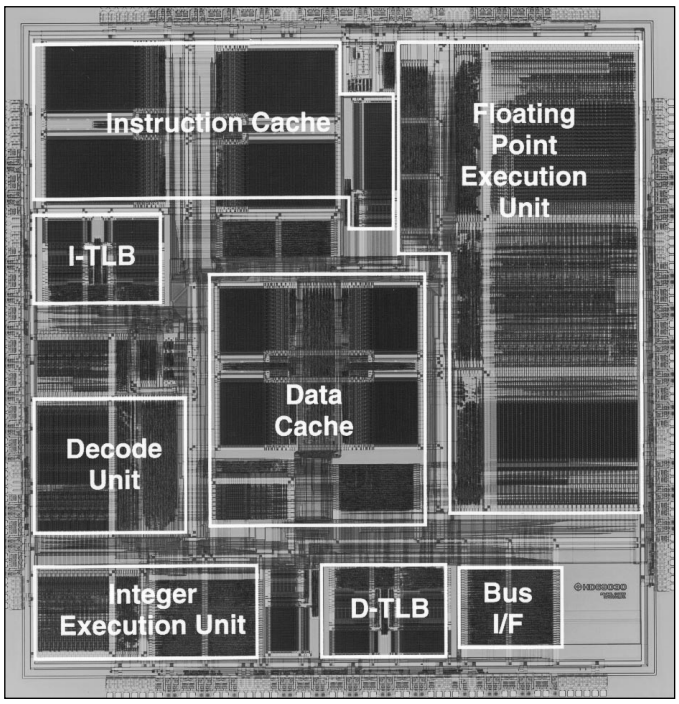

Figure 4. Hitachi's PA/50L measures 11.5 x 12.0 mm. The 1.28 million transistors are fabricated using a 0.6 µm CMOS process.

buffer is empty when data arrives, special bypass buses engage to provide a short path around the buffers, eliminating unnecessary latency.

Simplified Access to SDRAM

The PA/50L multiplexes address and data onto a single 32-bit bus, and provides support for synchronous DRAM (SDRAM). By comparing the row and bank of the current address with those of the previous address, the processor can control what is basically a pipelined interface into the SDRAM *(see **070205.PDF**)*.

Memory management uses a 32-entry TLB for instructions and a 64-entry TLB for data. Both are two-way set associative, with each entry covering a 4K page. Each also has two additional entries that translate larger blocks of addresses, adjustable from 256K to 32M. The TLBs can mark memory pages as uncachable. The 32-byte store buffer holds pending data until the external bus is available. A control bit selects between big-endian and little-endian mode for both code and data.

	Oki OP32/50N	Hitachi PA/50L	Intel i960CF
Clock	33 MHz	33 MHz	33 MHz
Dhrystones	39 MIPS	55 MIPS	51 MIPS
Caches			
Instruction	8K, 2-way	8K, 2-way	4K, 2-way
Data	4K, 2-way	4K, 2-way	1K, direct
Registers			
Integer	32 × 32 bits	32 × 32 bits	32 × 32 bits
FP	None	28 × 64 bits	None
Die Area	204 mm^2	138 mm^2	125 mm^2
Transistors	1.1M	1.28M	900K
Process	CMOS	CMOS	CMOS
Size	0.8 micron	0.6 micron	0.8 micron
Type	3-layer metal	3-layer metal	2-layer metal
Supply	5 V	3.3 V	5 V
Power	4 W	1.3 W	4.8 W
Package	208-pin PGA	160-pin PQFP	196-pin PQFP
Mfg. Cost (est.)	$120	$95	$55

Table 1. The PA/50L delivers higher performance with lower power than Oki's OP32/50N, but i960 offers lower costs.

The PA/50L also provides an opportunity to optimize block moves through a software "hint." This is a flag in an instruction that prepares the processor for the upcoming block move. In normal operation, moving a block of data would cause the following sequence of events:

- read data from a source block.
- attempt to write data to its destination block (and in doing so, cause a cache miss).
- cache miss causes a readahead of the destination block (which brings unwanted data into the cache).
- newly-read data is overwritten by destination data.
- cache is flushed to write destination data to memory.

By providing the "block copy" hint in software, the processor skips the readahead of the destination block, which dramatically reduces the number of cycles.

Like the Oki processor, the PA/50L decodes 32 interrupt levels that come in on five pins. For fast interrupt handling, the processor saves the state of seven of its registers into "shadow" registers. After the interrupt handler completes, the CPU reloads the shadowed contents to the original registers.

Large Die Size Indicates Higher Cost

It's difficult to directly compare Oki's processor with Hitachi's. The PA/50L would require a companion chip set to place it on an even footing with the OP32/50N. Even so, the price/performance ratio of the two is similar.

Oki builds its OP32/50N in 0.8-micron CMOS process with three metal layers. The device comprises 1.1 million transistors and is packaged in a 208-pin PGA. At 204 mm^2, Oki's large die approaches the size of a Pentium. While the company has yet to set a price, the sheer acreage of silicon will force it to the high end. In fact, the MPR Cost Model *(see **071004.PDF**)* predicts a manufacturing cost of $120 for the OP32/50N.

Using a 0.6-micron CMOS process with three metal layers, Hitachi puts 1.28 million transistors onto a 138 mm^2 die, which is packaged in a 160-pin PQFP. Despite the fact that it is only two-thirds the size of the Oki chip, the PA/50L's predicted manufacturing cost is $95, due to its more expensive process.

The OP32 and the PA/50L will be pitted against other high-performance embedded processors, such as Intel's i960CF—stiff competition considering the maturity of that part. As shown in Table 1, these three parts have comparable performance, but the i960CF costs half as much: While the Intel device is priced well above $100, its estimated manufacturing cost is $55.

Clearly, the price/performance ratio of existing products is going to be a barrier to the new PA-RISC processors. To add further perspective, at an estimated manufacturing cost of $150, HP's PA7100 offers two to three times the performance at less than double the cost.

Another factor related to die size is power. The PA/50L uses just 1.3 W, but the OP32/50N requires up to 4 W. While not excessive (considering the number of chips dissipating over 10 W), this places some constraints on where the Oki device can be used.

The OP32/50N's power requirements can be reduced in a number of ways. For example, running the processor at 16 MHz cuts power by half. Also, the chip has a power-down mode activated by setting a control-register bit. This halts the processor by preventing clocking. DRAM is not refreshed during power-down, so critical data should be kept in non-volatile memory. (This mode is exited by asserting external interrupts, system reset, or certain other signals.)

Faster Chips in the Works

Hitachi intends to introduce a faster version of its processor, which will be called the PA/50M. This device reportedly uses the same die as the slower PA/50L, but runs at 60 MHz (almost double the PA/50L's speed) because it uses a 5 -V supply. The PA/50M should reach 100 MIPS, but its higher power dissipation will require a different package. It may be introduced as soon as 2Q94.

Oki is also working on a second version of its processor, which will be called the OP32/60N. This device uses the same design as the OP32/50N, but it will be fabricated using 0.5-micron process and run at 3.3 V. These changes will allow the OP32/60N to achieve the same performance level yet use a lower-cost plastic package.

Cooperation Mitigates Competition

Oki and Hitachi, through their association in the Precision RISC Organization (PRO), have managed to build processors different enough from each other that they can hit separate segments of the market. The Oki device appears to be a likely candidate for high-performance embedded control. Its cost is relatively high for embedded applications, and its price/performance ratio somewhat less than competitors', but these factors are easily offset. At the high end of the market, where volumes are lower, development costs and time-to-market can be more critical than cost-of-goods. The OP32/50N's flexible glueless interfaces can dramatically cut the effort required to design and debug a product. Oki's device is probably the easiest way to get PA-RISC compatibility.

Hitachi has been shipping workstations based on HP processor modules, which use a design that relies on external primary cache. The PA/50L brings the cache onboard, reduces the bus to 32 lines, and wraps plenty of processing in a handy package. This simple external interface makes the PA/50L attractive as a high-end embedded controller, but its floating-point and memory-management units—and most of all, its PA-RISC compatibility—make it applicable in an entry-level workstation, too.

Hitachi has a commitment to PA-RISC and has announced that such a workstation is imminent. Other uses are not yet disclosed, but the PA/50L could find its way into Hitachi products such as X-terminals, laser printers, and hardware for networks and communications. The company also indicates that it doesn't use HP's memory controller or other interface chips. Undoubtedly, Hitachi is developing its own companion chipset for the PA/50L; whether the company will make these companion chips available to the outside world is unclear. The company has committed to making the processor itself available to members of the PRO group, but not necessarily to the market at large.

One Small Step Forward

For years, the PA-RISC architecture has been held back by restricted supply to all but close partners, lack of the companion chips required to form a system, and lack of third-party applications. What's more, the embedded market where these two devices are aimed pays close attention to price/performance, absolute price, and especially to development tools such as compilers, debuggers, logic analyzers, and emulators. In none of these categories do the PA-RISC processors stand out.

The advent of these two devices may represent a turning point, but there's still a long way to go. To its credit, Oki's design eliminates the requirement for companion chips, and choosing embedded systems as a target reduces the need for support from third-party applications. Further, the lower cost of both devices (compared with HP's chips) may well translate to higher volume and increased acceptance.

While there is a market for these chips within the PRO group, the world probably won't notice any change in the low-cost desktop market. It's unlikely these two processors will be enough to make the PA-RISC architecture a contender in this arena, especially considering the momentum of established processors. ♦

Price and Availability

Price and availability for Oki's OP32/50 have not yet been determined, but should be announced early in 1994. For more information, contact Rebecca Franks at Oki Semiconductor, 408.737.6378.

Hitachi's PA/50L is scheduled to begin sampling in 4Q93, while PA/50M samples are slated to be available 2Q94. Hitachi has not yet determined a price for the parts. For more information, contact Hitachi America at 415.244.7600.

