IDT Spins R4600 for Consumer Market

R4650 Adds Simple DSP Function to MIPS Core

by James L. Turley

Adding to their existing array of RISC processors, MIPS licensees IDT and NKK will begin producing a new device aimed at high-end games, imaging, and other compute-intensive embedded applications. Announced at last month's Microprocessor Forum, the R4650 is a spin-off of the R4600 Orion processor *(see **0714MSB.PDF**)* currently offered by IDT, NKK, and Toshiba.

The new device sacrifices double-precision floating-point and most of the MMU to reduce the die size and cost but improves integer multiply performance considerably. The R4650 is also the second MIPS implementation so far to include a multiply/add instruction, a shot at displacing separate DSP chips.

Like the R4600, the new processor was created by Quantum Effect Design (QED), a small design firm founded by former MIPS employees. The R4650's development was funded by IDT and NKK, both of which will sell the devices. Unlike the R4600, Toshiba did not participate in the design. Tom Riordan, presenting for QED, said the design work is complete; both vendors expect to be sampling parts in 1Q95.

CLARENCE TOWERS

QED's Tom Riordan presenting the details of the new R4650 and R4700 MIPS processors.

The modifications tailor the R4650 for applications where the R4600 has already had some success, namely network routers and page rendering, and for new applications in video compression/decompression and high-end games.

The R4650 improves integer multiplication latency by about a factor of four compared with the R4400 and R4600. Multiply latency is only three clocks for 32-bit operands or two clocks if both operands are 16 bits. In other MIPS processors, multiplying integers is effectively a two-step process. The result of a MULT instruction is placed in a pair of special-purpose registers; two additional instructions then transfer the contents of these registers into the CPU's general registers. For the R4600, this effective multiply time is 10 clock cycles for 32-bit quantities.

The R4650 supports the MULT instruction for software compatibility but adds a new MUL instruction that deposits a sign-extended 32-bit result directly into a general register. This process saves the extra step of copying the result but truncates any value beyond 32 bits with no warning. For well-controlled data sets, the MUL instruction can save considerable time. Most fixed-point DSP chips—which the R4650 is intended to replace—work on 8- to 12-bit operands, so the restriction to 32-bit results should not be a problem for target applications.

The second significant enhancement is the integer MAD (multiply/add) instruction, similar to those found on DSP chips. The MAD instruction performs an integer $A \times B + C$ function in three cycles, assuming all three operands are available. The last cycle can be overlapped with the first cycle of another multiply operation, for an effective rate of one MAD every two cycles. In practice, a LOAD/MAD instruction pair would be required to load one operand from external memory. Assuming no data cache misses, the R4650 can stream LOAD/MAD pairs at one-half the processor clock rate, or 67 million operations per second at 133 MHz.

The MAD instruction is not part of the MIPS-3 standard architecture. QED developed this enhancement for IDT, borrowing an unused opcode with the understanding of MIPS Technologies. IDT has coordinated with MIPS Technologies to reserve more unused opcodes for future enhancements such as this one.

On the floating-point side, single-precision performance is identical to that of the R4600. Double-precision support was removed entirely. Unimplemented double-precision instructions from the R4600 are trapped, permitting software emulation of the missing operations. Creating a narrower data path to the FPU reduces critical clock-line loading in addition to reducing chip area.

Modified Cache, Stripped MMU

Separate data and instruction caches are 8K each, only half the size of those in the R4600. As in its predecessor, each cache is two-way set associative, arranged as 256 lines of 32 bytes. Unlike the R4600, the R4650 has the ability to lock one set from either cache via software. Locking one cache set leaves the other half operating normally, though a 4K cache is at the lower limit of usability for such a fast processor. Locking both cache sets at once is not possible, but the same effect can be achieved by defining the entire virtual address space as noncachable.

Unlike the R4600, both caches are physically indexed and physically tagged. No hardware cache snooping is provided, so software must flush the caches if necessary. The lack of hardware snooping is typical for embedded CPUs, as multiprocessing embedded systems are not yet the norm. Even those systems with DMA can generally confine transfers to noncachable areas of memory, so snooping is unnecessary.

Abandoning the typical MIPS-compliant MMU as too complex and costly for embedded use, QED included a set of "base-bounds" registers for rudimentary memory protection. The new MMU has two upper bounds registers, one for instructions and one for data. If a virtual address is greater than or equal to the value in the bounds register, a fault is generated. Otherwise, a programmable offset is added to the virtual address, producing the physical address. Different tasks can use different base/offset register values, mapping the same virtual address to different physical addresses.

Programmers familiar with the R4000's peculiar "external write request" method of generating interrupts will celebrate the departure of that feature on the R4650. Where the pin-limited R4000 was forced to multiplex all interrupts onto a single pin, the R4650 has a pin for each interrupt source. The new chip also distinguishes between internally generated faults and hardware interrupts, jumping directly to the user's interrupt handler for the latter, saving the four clock cycles needed for a table lookup.

Bus Options Save Power, Cost

The R4650's external bus interface is identical to the multiplexed 64-bit address/data bus found on the R4600. Normally, the bus operates in a 64-bit mode. A power-up option can configure the device for 32-bit operation. When writing 64-bit operands in this mode, the R4650's bus controller will initiate two separate 32-bit write cycles rather than broadcasting one address followed by two data cycles. Although the R4650 is not plug-compatible with the R4600, the bus protocols are the same, making the R4650 useable with existing PC system-logic chips from Toshiba, Acer, DeskStation, and NEC.

Externally, the R4650 operates like the R4600, but its internal bus is a new design for IDT and NKK. Both suppliers intend to use the chip as a springboard for custom devices, so an intermediate bus between the CPU core, caches, and external bus was developed. The R4650 is the first in the MIPS R4000 family device to be expressly designed for such customization.

Power consumption varies widely depending on internal and external clock speed and whether the processor is operating with the 32-bit or 64-bit external bus. From a low of 0.7 W at 40 MHz in 32-bit mode at 3.3 V, the 5-V version of the R4650 dissipates as much as 10 W in 64-bit mode at 133 MHz.

Price & Availability

Samples of the IDT79R4650 (5 V) and IDT79RV4650 (3.3 V) will be available 1Q95, with production in 2Q95. Three version will be available: 80, 100, and 133 MHz. Prices start at $64 for the 80-MHz part in 10,000-unit quantities; other pricing is not available. For more information, contact Integrated Device Technologies (Santa Clara, Calif.) at 408.727.6116; fax 408.492.8674.

To reduce power usage, it is possible to stop the clock, but a WAIT instruction works better for most applications. After a WAIT instruction is executed, power consumption drops to around 200 mW; the part can be reawakened by an external interrupt. Stopping the clock stops the PLL and drops power to about 30 mW, but restarting from a stopped clock takes much longer than from a WAIT. The R4650 is fully static, so no state information is lost in either case.

The die size is 56 mm^2 in IDT's 0.65-micron three-layer-metal process (vs. 81 mm^2 for the R4600), making it smaller than an Am29040 or i960Hx. According to the MPR Cost Model, the estimated manufacturing cost is $22, or 40% less than that of the R4600.

A standard 208-pin PQFP is large enough for the processor's 64-bit multiplexed address/data bus. A smaller package that does not bond out the upper external bus lines may be offered in the future for 32-bit systems. Both 3.3-V and 5-V versions are planned, with internal processor speeds of 80, 100, and 133 MHz.

IDT is pushing the limits of its low-cost packaging with the R4650. Only the two lower speed grades are available in a plastic QFP. The 133-MHz version dissipates too much heat for a plastic package, pushing IDT to a 208-pin MQUAD package.

64 Bits for 64 Bucks

IDT has set an aggressive price for the R4650—as low as $64 in quantities of 10,000. The company's RISC parts start at about $15 for the R3041, on up to the 150-MHz R4400 at $1,035, so the R4650 is positioned at about the middle of IDT's processor line.

The R4600 has proved popular in high-bandwidth communications and graphics applications. Both Sony and Nintendo have adopted the MIPS architecture for their next-generation high-end video games. The Sony Playstation will use a derivative of the R3000 fabricated by LSI Logic *(see **080902.PDF**)*. Nintendo will get its R4200 processors from NEC for its upcoming Project Reality game player. QED seems to have done a good job preserving the R4600's best assets while increasing its appeal with simple DSP functions and lower cost.

The R4650's quoted performance of 175 Dhrystone MIPS at 133 MHz places it well above Intel's top-of-the-

R4700 Aims for Graphics

At the request of its biggest customer, IDT has created a new version of its R4600 processor that is tuned for workstation applications. The new R4700—which, like the R4650, was announced at last month's Microprocessor Forum—boosts floating-point performance by about 20%, making it ideal for Silicon Graphics' low-end and midrange workstations. SGI is already using the R4600 in its Indy workstations but is eager to offer better FP performance.

This performance increase comes mainly through an enhanced floating-point multiplier. To conserve space, the original design had a relatively slow multiplier that needed eight cycles. The R4700 uses a larger circuit to complete a single-precision multiply in four cycles and double-precision in five. The issue rate has also been improved from six cycles to four, the same as for adds. Adds and multiplies share a single two-cycle rounding circuit; floating-point operations can be launched on every other cycle by alternating adds and multiplies, delivering a peak rate of 87.5 MFLOPS at 175 MHz.

The R4700 also gains performance through a slight increase in clock speed. QED's Tom Riordan expects the new chip to reach 175 MHz, up from 150 MHz, by moving to a 0.60-micron CMOS process instead of the 0.64-micron process used by the R4600. At that speed, the R4700 is expected to deliver 130 SPECint92 and 100 SPECfp92, approaching the integer performance of a 200-MHz R4400.

IDT rates the part at 3.5 W maximum (at 175 MHz and 3.3 V), but this figure does not include any bus loading, which would push it higher. IDT plans 3.3-V and 5-V versions. Due to the smaller process geometries, the die size of the R4700 is slightly smaller than that of the R4600, even with the larger multiplier array. At just 73 mm^2, the new chip carries an estimated manufacturing cost of $35, about the same as its predecessor.

QED recently taped out the design but has not yet received first silicon. Because of the relatively minor changes, IDT plans to sample using first silicon, starting in 1Q95. The production version is expected in 2Q95. So far, the company has announced speed grades of 100, 133, 150, and 175 MHz but has provided pricing only for the 100-MHz version: $220 in quantities of 10,000 (using an MQUAD package). IDT says that it will provide pricing for the faster versions once it determines the frequency yield of the new process.

Like the R4650, the R4700 was funded by IDT and NKK; these two companies both plan to produce the part, but Toshiba will not. As the initial demand is expected to be almost entirely from one customer (SGI), there may be a limited market for the R4700. The new part offers some advantages over the R4600, but these are due mainly to the process shrink; Toshiba could simply shrink the R4600 masks to achieve the same effect. It may also wait for the second-generation "P4" Orion, due late next year.

—LG

line 75-MHz i960HT *(see **081302.PDF**)*. The R4650 carries a lower price and throws in a floating-point unit to boot. The 960 die is also nearly twice as large, at 100 mm^2, partly because of its 16K data cache. Fabrication processes for both parts are comparable, and they share a 208-pin QFP outline.

Like the Intel device, the R4650 operates with a relatively slow external bus, using clock multiplication to accelerate the CPU core. Whereas the i960Hx parts offer double- and triple-speed internal cores, the MIPS implementation permits any integer multiple from 2× up to 8×. Finally, the MIPS device supports all clock multiplication frequencies on one chip; frequency is selected dynamically during power-up initialization. Intel requires customers to buy a different chip for each multiple of bus speed.

Performance to Spare

All these benefits may not outweigh the perversity of the R4650's underlying Orion bus protocol, however. The entire R4600 family was designed from the start for high-performance desktop workstations. High speed, low latency, and huge bandwidth were favored over moderate cost and ease of use. Any processor design that considers a 32-bit bus to be a low-end option is clearly not aimed at cost-sensitive office automation or mobile computing environments.

The dual out-of-phase clocks, multiplexed parity-protected address and data bus, 9-bit command bus, and six handshake signals all conspire to make the R4650 a nontrivial exercise to interface. A designer cannot simply hook up address, data, and a few control signals and produce a working prototype. Extensive, high-speed control logic is required, and involved state machines are a necessity; ASICs would be better. A high-volume application can amortize ASIC design costs over the life of the product. Low- to medium-volume applications must depend on PLDs or similar logic to work with the bus.

Although both the 960Hx and AMD's new superscalar 29K devices will sit at the top end of their respective families, the R4650 is nowhere near the highest-performing MIPS implementation available. R4650 users will have the advantage of using a processor that leaves plenty of room to maneuver—both up and down the performance path. The new R4700 (see sidebar) and R10000 *(see **081403.PDF**)* prove that there are still performance gains to be had for the architecture.

With its multiply/add function and 64-bit data bus, the R4650 is better suited for data shovel and simple DSP applications than many competitors. While Intel and AMD battle over laser printers, telecom, and RAID controllers, the R4650 is pursuing large network routers as well as the more glamorous—and potentially much more lucrative—set-top box and commercial video game markets.♦

