Digital's 21164PC Aimed at PC Market

New Chip Should Outrun Pentium II in $2,500 Windows NT Systems

by Linley Gwennap

Digital Semiconductor has performed an impressive magic trick: instead of sawing a woman in half, the company cut its 21164 chip in two without damaging its performance. Compared with its predecessor, the vendor's new 21164PC trims 34% off the die and 86 pins off the package, yet it achieves nearly the same SPEC ratings. As a result, the chip is suitable for systems that sell for as little as $2,500 and deliver better integer performance than any non-Alpha system shipping today. The new device is currently sampling at speeds up to 533 MHz and should reach volume production by 3Q97.

As the name indicates, the 21164PC is a cornerstone of Digital's strategy to push Alpha into the PC market. Alpha is the only RISC architecture that Microsoft has committed to support in future versions of Windows NT, but it has had little success competing against Intel in this market, despite its outstanding performance. The 21164PC addresses the issue of system cost, while the vendor's FX!32 emulation/translation software *(see* **1014MSB.PDF***)* bridges the gap to x86 applications. Digital's agreements with Mitsubishi and Samsung ensure multiple sources for Alpha processors.

All that's missing are major system vendors other than Digital itself. Digital hopes the 21164PC, by bringing Alpha performance to lower price points, will convince a few such companies to sign on. But by the time the 21164PC reaches the market, Intel's Pentium II should come within 10–15% of the Alpha chip's performance on NT business applications. This situation will leave Alpha mainly in the NT workstation market, where FP applications predominate; the new chip is unlikely to attract the support of the top-tier system makers Digital needs to penetrate the mainstream PC market.

Performance Gain Without Pain

Conceptually, Digital introduces a new Alpha CPU core at the high end of its line and then, within two years, puts that core into a low-cost processor. Following this plan, the original 21064 CPU begat the low-end 21066, which sank without a trace due to a poor price/performance ratio. Thus, the goal for the 21164PC was to lower the cost of a 21164-based system without sacrificing performance.

The cost reductions targeted four areas. First, the cost of the processor itself had to be reduced. Second, the large and expensive external cache of the 21164 needed to be replaced with PC-standard parts. Third, the cost of the system logic must be comparable to corresponding PC components. Finally, all other parts of the system—main memory, graphics, peripherals, power supply, etc.—must use high-volume PC technology.

Trimming the Fat From the 21164

The 21164 processor consumes 209 mm^2 in Digital's 0.35-micron four-layer-metal CMOS-5 process. As Figure 1 shows, there are only two options for significant die-size reduction: reducing the size of the CPU core or of the 96K level-two cache. Making the CPU core smaller would have involved significant redesign, probably resulting in a severe loss of performance, so Digital chose instead to completely eliminate the L2 cache. As the figure shows, this change cuts out about a third of the die.

To help compensate, the 21164PC has a 16K instruction cache, twice the size of the 21164's; the data cache remains 8K in size. Although doubling the data cache would have also improved performance, it would have been more costly than adding instruction cache, because the 21164's dual-ported data cache is implemented as two duplicate arrays, doubling its physical size. Also, the 21164 CPU performs some load reordering, helping to hide the latency of data accesses.

Removing 80% of the on-chip memory causes a relatively modest performance reduction of about 13% on SPECint95. Floating-point applications tend to access memory more often, however, so slashing the cache causes a 22% loss on SPECfp95. In both cases, the performance loss of omitting a traditional L2 cache would have been greater, but the 21164's 96K cache is much smaller than typical external level-two caches.

For the 21164PC, moving the L2 off-chip increases the latency to read the first word from 8 cycles to 20 and reduces bandwidth by a factor of four. These changes cause the performance reduction noted above. Digital's performance estimates, however, say this penalty is modest at the application level and is partially mitigated by doubling the instruction-cache size and expanding the L2 cache to 1M.

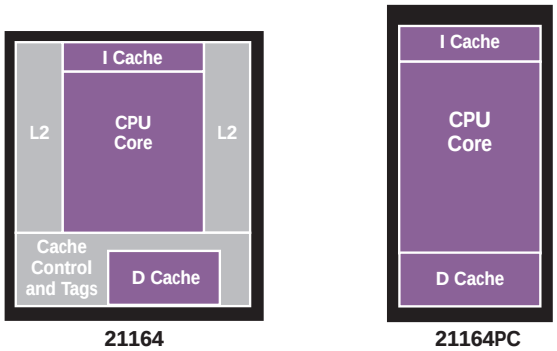

Figure 1. A floor plan of the 21164 shows the 96K L2 cache consumes 34% of the 209-mm^2 die. After removing this cache and making minor adjustments, the 21164PC consumes only 140 mm^2.

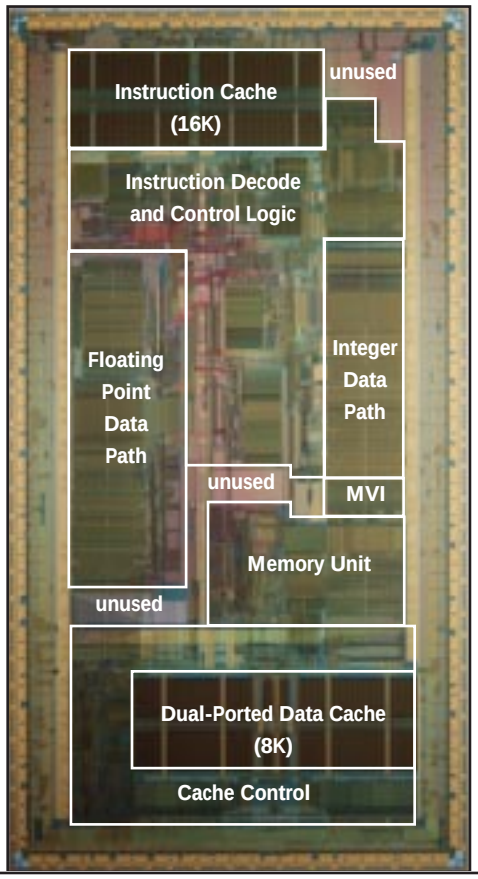

Figure 2. With 3.5 million transistors, Digital's 21164PC measures 8.5 × 16.2 mm in a 0.35-micron four-layer-metal CMOS process.

Leveraging the 21164 Core

In most other respects, the new chip is very similar to the 21164. It borrows the same four-way superscalar CPU core and always executes instructions in order, with reordering of memory references *(see* **081201.PDF***)*. The only change to the core is the addition of Digital's motion video instructions (MVI), which increase the performance of video encoding *(see* **101505.PDF***)*. The MVI circuitry consumes less than 1 mm^2, but as Figure 2 shows, Digital simply added 0.6 mm to the height of the die, avoiding the need to relayout the CPU core; this decision resulted in a 5-mm^2 increase but helped keep the design time to just nine months.

Like other Alpha processors for NT, this chip has a restricted MMU that doesn't support Digital's Unix or VMS operating systems. This restriction allows Digital to maintain higher margins on processors for these systems than on its NT boxes.

Eliminating unnecessary pins trimmed the package cost by about $10 and also reduced the size of the pad ring. Most of the excised pins were unneeded power and grounds dating back to when the 21164 chip was a 3.3-V, 50-W fire breather. Using a 2.0-V core, the 21164PC dissipates a relatively modest 24 W (maximum) at 533 MHz.

Other pins were saved by simplifying the maximum system configuration. For example, the 21164 supports multiple processors; the 21164PC does not. The new chip allows "only" 4M of L2 cache, compared with a maximum of 64M in the 21164. In total, these changes allow the 21164PC to fit into a 413-pin ceramic PGA.

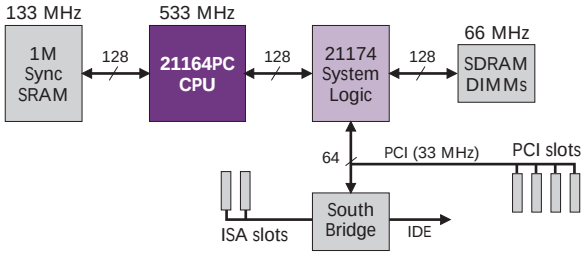

Figure 3. The 21164PC's L2 cache RAMs are similar to Pentium II's, and the Alpha chip takes advantage of standard PC memory and system interfaces via the 21174 system-logic chip.

Digital is one of the few high-end CPU vendors that has not converted to BGA packages, which are considerably less expensive than PGAs for the same number of pins. Sometimes there are advantages to being a technology laggard; because the PGA-to-BGA transition has occurred faster than many package vendors anticipated, there is currently an oversupply of PGAs available, driving prices down to near BGA levels. This imbalance is likely to evaporate over the next year, forcing Digital to update its packaging strategy.

The MDR Cost Model estimates the 21164PC costs about $90 to manufacture, a third less than the 21164. Pentium II, in contrast, is estimated to cost $80, despite having a larger die and more pins. This difference is mainly due to the greater economies of scale achieved by Intel. The Cost Model also penalizes the 21164PC for its PGA package, as the model doesn't take into account short-term supply imbalances. If the market price of PGA packages stays low, the 21164PC could cost slightly less to build than Pentium II.

Cache Interface Improved

Because the 21164PC relies more heavily than its predecessor on its external cache, Digital revised the cache interface to improve performance. As Figure 3 shows, the new chip uses a 128-bit interface, as does the 21164, but it typically accesses the cache at one-fourth of the CPU clock rate: 133 MHz for the 533-MHz CPU. In contrast, the 21164's external cache operates at just 66 MHz. Thus, the new chip has twice the cache bandwidth of its predecessor.

The 21164PC achieves this speed using burst synchronous SRAMs, essentially the same chips used by Pentium II. It's no coincidence that Digital's 466- and 533-MHz processor speeds are exactly twice those of 233- and 266-MHz Pentium II processors; this ratio allows Digital to leverage the Pentium II cache chips. The 21164 relies on older asynchronous SRAMs, reducing the performance of its external cache.

Even though the new cache is faster, it isn't necessarily more expensive. With Pentium II reaching volume production soon, these fast burst SRAMs will carry little, if any, premium over slower asynchronous parts. The main cost of the

cache relates to the number of parts; Digital expects most 21164PC systems to use 1M of cache, whereas most 21164 workstations use 2–4M. The company's performance estimates for the new chip all assume only 1M of cache.

In addition to doubling the cache bandwidth, Digital made a few minor tweaks to the processor's interfaces to improve performance. The cache timing is fully programmable to support a variety of SRAMs. The latency to the external cache is reduced slightly, from 42 ns to 37 ns. The 21164 can have only two pending loads from main memory on the system bus; the 21164PC extends this to three. Although this difference seems small, the change allows the 21164PC to fully saturate the system bus doing reads from memory; in the older design, this bus bandwidth was often wasted.

System Logic Lags Intel's

Because its system bus is different from Pentium II's, the 21164PC can't use standard PC system logic. Digital has developed a new system-logic chip, the 21174, to bridge this gap. This device connects the processor to a 128-bit DRAM bus and a 64-bit PCI bus. The memory bus supports SDRAM (but not EDO), a major step forward for Digital, as its previous chip sets allowed only fast-page-mode DRAM. Up to 512M of DRAM is supported, with parity or ECC. The PCI interface supports standard 32-bit devices as well as emerging 64-bit devices. Intel chip sets support only 32-bit PCI, so there are few 64-bit devices available today.

A 383-pin BGA holds the system-logic chip, which dissipates 2–3 W under normal operation. The 21174 requires a 3.3-V supply, matching the I/O voltage of the CPU, while allowing 5-V DRAM and PCI.

This feature set is serviceable, but it lags that of the latest PC chip sets, such as Intel's 430TX *(see* **1102MSB.PDF***)*. The 430TX supports the ACPI power-management specification; the 21174 has limited power-management capability that isn't used in the initial 21164PC systems. The 430TX supports UltraDMA, increasing I/O performance, and SMBUS for supportability; Digital claims the 21174 has similar, but not compatible, capabilities. The 430TX dissipates about half a watt, far less than Digital's chip set.

One final shortcoming: Digital is pricing the 21174 at $142 in quantities of 1,000; the 430TX lists for $33. Intel's price includes a "south bridge" PCI-to-ISA chip; Digital's does not. The 430TX is designed for Pentium, not Pentium II, but we expect Intel's forthcoming 440LX chip set to include similar features at a similar price. The 440LX is slated for volume production at about the same time as the 21174.

The 440LX will also add a new feature: AGP support. AGP is Intel's initiative to increase bandwidth to the graphics accelerator, particularly for 3D chips *(see* **100803.PDF***)*. We expect Intel to exert its marketing muscle to push AGP into the market quickly, although it won't become widespread until 1998. Digital plans to release a 21174 follow-on with AGP support; this version is due to ship sometime in 1H98.

Mitsubishi Supplies 21164PC

Long-suffering Alpha partner Mitsubishi will be a second source for the 21164PC processor. The Japanese vendor signed up for Alpha exactly four years ago *(see* **0704MSB.PDF***)* but has yet to ship an Alpha processor for revenue. Instead of abandoning Alpha, Mitsubishi joined the development team, sending about a dozen engineers to Massachusetts to help Digital design the 21164PC.

Mitsubishi is slightly behind Digital in bringing the chip to market, with samples planned for June and volume production for September. The company has not finalized its pricing, but it appears significantly higher than Digital's: about $600 for the entry-level 400-MHz part, 50% more than Digital's price. Mitsubishi plans to reduce prices "quickly," moving below $300 early in 1998.

At first glance, two Alpha processor suppliers appear to be one too many. Unlike Digital Semiconductor, however, Mitsubishi has strong sales channels into many PC makers, due to its SRAM and DRAM products. The company may also differentiate its Alpha products from Digital's in the future. One possibility is combining an Alpha CPU with embedded DRAM, a technology Mitsubishi pioneered with its M32R/D chip *(see* **100702.PDF***)*.

Although Samsung did not initially have rights to the 21164PC design as part of its Alpha license, the company recently renegotiated its deal and now plans to sell the part. Samsung supplies memory chips to more PC makers than Mitsubishi does, and the Korean vendor is likely to offer very aggressive pricing. This combination could spell trouble for Mitsubishi once Samsung gets the Alpha chip into production.

Motherboard Provides PC Compatibility

By providing standard PC interfaces such as SDRAM and PCI, the 21174 allows Digital to leverage the PC infrastructure. A complete example is provided by Digital's 164SX motherboard, which supports the 21164PC and the 21174 system-logic chip within a standard PC system. It uses the same ATX form factor as Intel's motherboards, so no mechanical redesign is needed.

The motherboard uses a Cypress 82C693U PCI-to-ISA bridge that provides a real-time clock, keyboard controller, and IDE interface. The initial motherboard doesn't include a USB serial port, but a future version of the Cypress chip will support USB, and it could be brought out on future versions of the board. An SMC 37C669 provides standard serial and parallel ports.

The 164SX motherboard includes four slots for 168-pin SDRAM DIMMs, four PCI slots (two for 64-bit cards), and two ISA slots. The motherboard includes no graphics or networking support; these functions must be added via PCI cards. Digital plans to sell the board for less than $500,

Price & Availability

In quantities of 1,000, Digital is selling the 21164PC for $295, $395, and $495 in speed grades of 400, 466, and 533 MHz, respectively. The chips are sampling now, with production slated for Q3. For more information, access the Web at *http://ftp.digital.com/pub/Digital/info/semiconductor/literature/dsc-library.html* or contact Digital Semiconductor at 1.800.332.2717.

Mitsubishi will also sell the 21164PC in speed grades of 400 and 533 MHz, with volume production scheduled for September. The company has not yet announced pricing. Contact your local Mitsubishi sales office.

including the system logic, 1M of cache, and Windows NT BIOS but no CPU or DRAM. The company will also supply schematics free of charge to companies interested in building and selling the motherboard.

Digital has also signed up VLSI Technology to deploy a 21164PC system-logic chip. The VLSI chip set will better match the features of Intel's chip sets, although it won't include ACPI. The company hasn't announced price or availability; we expect the chip set to appear by the end of the year at a price between Intel's and Digital's. By early 1998, motherboards based on the VLSI chip set should be available for about $300, improving the competitiveness of the 21164PC.

Small Performance Advantage

Digital has not published measured SPEC95 scores for the 21164PC, since no systems are yet available. The company estimates the 533-MHz chip will deliver 7% less performance than a 500-MHz 21164 on SPECint95 and 17% less on SPECfp95. At that rate, the 21164PC should reach 11.5 SPECint95 (base) and 15 SPECfp95 (base). (Digital prefers to quote scores on SPEC95 peak, a different and less realistic benchmark.) This integer score is higher than that of any non-Alpha processor, but a few other RISC/Unix chips can match or exceed the FP score.

By the time the 21164PC is in systems, we expect Intel's Pentium II will be in volume production. At 266 MHz, this chip should deliver about 10.5 SPECint95 (base), 10% less than the projected performance of a 21164PC-533.

On SYSmark for NT version 1, the 21164-500 is 22% faster than a Pentium Pro-200. Some of these tests are I/O-bound, however, and do not show off the full capabilities of the Alpha chip. Early testing on SYSmark/NT 4.0 shows the 21164 outperforming Pentium Pro by about 45%. Considering that the 21164PC will be slightly slower than the 21164 and that Pentium II at 266 MHz will do significantly better than Pentium Pro, we estimate the new Alpha chip will outperform Pentium II by 10–15% on native NT business applications. Users of PowerPoint or other non-native applications will, of course, see worse performance.

Digital's pricing for the 21164PC is quite low, particularly compared with that of previous Alpha processors. The company has delivered on its goal of "less than $1 per MHz," with prices ranging from $295 for the 400-MHz version to $495 for the 533-MHz speed grade. This should give Digital a price advantage over the 266-MHz Pentium II, which we expect will debut at about $750, including at least 256K of cache. Adding system logic might push this total to $800. A 533-MHz 21164PC, with 1M of cache and the 21174 system-logic chip, totals less than $700. The price of Pentium II will probably drop below $500 by the end of the year, but Digital says it will keep pace with Intel's price cuts.

Both Intel and Digital will increase the performance of their processors by moving them to 0.25-micron CMOS. Intel's new CMOS process appears to be about six months ahead of Digital's, but Digital may gain a bit more clock speed from the shrink, perhaps getting as high as 733 MHz in 2H98. Intel, however, plans to improve the performance of its chips by releasing two new designs, code-named Katmai and Willamette, in 1998 or early 1999. In short, Digital's PC processor roadmap doesn't appear to offer a compelling integer performance advantage for the next year or two.

Ultimately, Digital is likely to develop a 21264PC, using the 21264 core in a low-cost product. This product should offer a large performance improvement over the 21164PC. Since the original 21264 isn't due until 1H98, however, the low-cost version is unlikely to appear before 2H99.

Tough to Penetrate PC Market

Digital plans to launch an advertising campaign aimed at convincing PC buyers that Alpha is a hot box that will be faster than Intel-based systems while offering full Windows NT compatibility. The basic message, "Alpha is fast," is easy to communicate, but people who buy into this message may be disappointed by reality. The 21164PC appears to have only a 10–15% performance advantage over Pentium II on typical NT business applications. This small advantage will convince few PC buyers to switch from a safe Intel-based system to a radical Alpha product. The 21164PC will offer about twice the floating-point performance of Pentium II, so users with heavy FP demands will find the Alpha chip valuable.

Thus, the 21164PC should assist Digital's efforts in the growing NT workstation market, but it seems unlikely to attract major PC makers or allow Alpha to penetrate the mainstream PC market. Digital's trick of gutting the on-chip cache while salvaging performance is impressive, but the two-year lag between the original 21164 and the low-cost version allowed Intel to nearly catch up in performance.

The 21164PC shows how difficult it is to exceed Intel's performance in a PC infrastructure that Intel itself has defined to precisely meet the needs of its own products and of mainstream PC software. Digital will need a more radical approach to successfully penetrate the mainstream PC market, but in the meantime, Alpha may prosper within the confines of the NT workstation market. M