

Digital Reveals PCI Chip Sets For Alpha

Two System-Logic Chip Sets Will Support 21064 Processor

By Linley Gwennap



Taking another step in opening its Alpha architecture, DEC disclosed at the MicroSystems Forum that it is developing two PCI-based system-logic chip sets for its 21064 microprocessor. One is a high-performance design that the company believes will provide workstation performance; the other is a high-integration solution similar to the Intel PCIset (see [070403.PDF](#)).

Current Alpha systems use a large number of PALs and discrete logic chips to interface the 21064 to memory and expansion buses. By reducing both the manufacturing cost and the design time of such interfaces, Digital hopes that the new chip sets will increase the number of Alpha system vendors.

Although DEC is designing the new chips, the company expects to license a third party to actually market them, but no firm arrangements have been made. According to Aaron Bauch, who made the presentation, the company expects that both chip sets will sample this fall and start production shipments by the end of the year, although neither design has yet been fabricated.

Pricing will be announced once a vendor has been identified; DEC is positioning its high-end design against Vitesse's \$150 Pentium cache controller (see [070602.PDF](#)), while the low-cost version will be matched against the PCIset, which is listed at \$84.

Six-Chip and Three-Chip Designs

As shown in Figure 1, the high-performance chip set will use six chips for the basic cache, main memory, and PCI interfaces: one cache/memory controller (CMC), one PCI bus interface (PBI), and four data-path (DP) chips. The six-chip configuration uses a 128-bit path to main memory, just as in Digital's Alpha workstations. The chip set can also be configured with two DP chips, yielding a 64-bit DRAM interface. The 128-bit interface allows for error correction, but only parity protection can be used with the 64-bit memory system.

The CMC incorporates some cache-control logic and a complete main-memory interface. It generates addresses and control signals for the DRAMs. It also controls the flow of data through the data-path chips. The DP chips actually move the data between the DRAM, the processor, and the PCI bus. They consist mainly of various buffers that help reduce delays when reading and writing data.

The PBI implements the PCI interface, taking data from the PCI bus and sending it to the data path using the EBI bus (see Figure 1). Because PCI devices can use virtual addressing, some contiguous virtual addresses may be scattered across multiple physical pages. The PBI contains a "scatter/gather" unit that combines such data into a single burst transaction on the PCI bus. An 8-entry TLB reduces the number of page-table accesses.

The CMC can also drive a single bank of VRAM to create a graphics frame buffer of 1M–16M. DEC believes

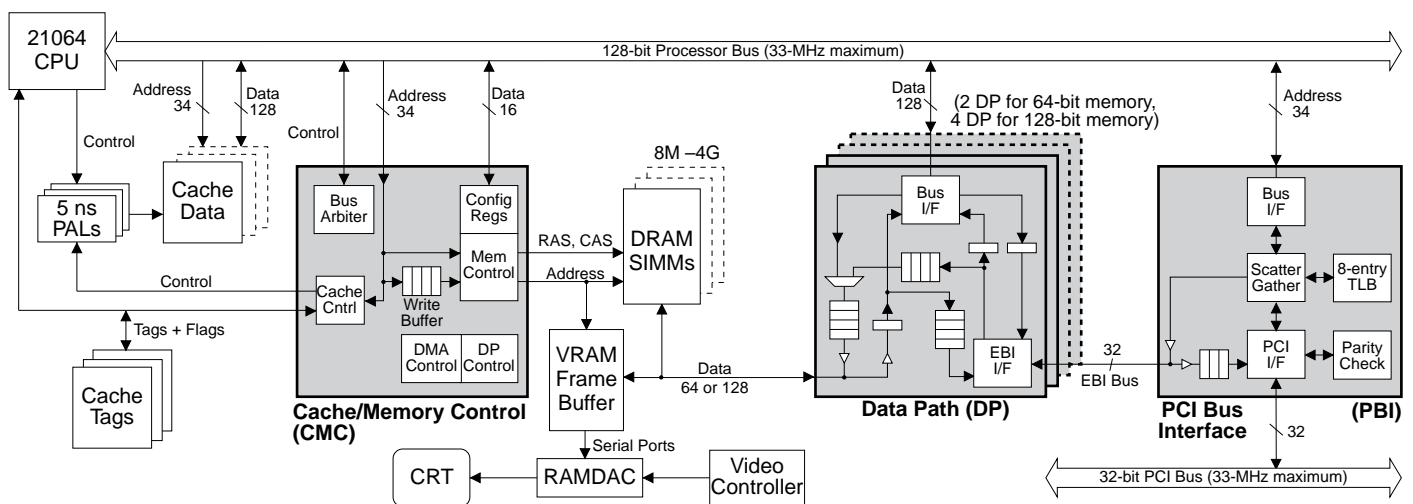


Figure 1. DEC's high-performance PCI chip set provides second-level cache and DRAM control for a 21064 CPU, along with a PCI bus interface and an optional graphics frame buffer.

CPU Clock Frequency	System Clock		SRAM Access Time	Cache Read Access	
	Divisor	Frequency		Divisor	Cycle Time
150 MHz	5	30 MHz	15 ns	5	33 ns
			12 ns	4	27 ns
			10 ns	4	27 ns
200 MHz	8	25 MHz	15 ns	6	30 ns
			12 ns	6	30 ns
	6	33 MHz	10 ns	5	25 ns

Table 1. The cache read time and system clock rate must be integer divisors of the CPU clock. The minimum cache read time is 14 ns longer than the SRAM access time; the maximum system clock frequency is 33 MHz.

that the Alpha CPU has enough power to eliminate the need for a separate graphics processor and estimates that this chip set will deliver 25–35 WinMarks. This is a particularly rough estimate, since WinBench is not supported on Alpha platforms, but is intended to give a feel for the comparative performance.

The CMC updates the frame buffer and also maintains a refresh pointer that loads the VRAM serial buffers. An external video controller is required to shift the video data out of the VRAMs and to control the RAM-DAC, as shown in Figure 1. If desired, a standard PCI graphics accelerator can be used instead.

For lower-cost system designs, Digital is developing a more integrated version of the chip set that combines the CMC and PBI into a single chip. This version supports only a 64-bit main memory, so the system logic can be implemented in three chips: one control plus two data-path chips. The company has not disclosed complete technical details on this more integrated version, but it appears to be very similar to the high-performance version except for the chip partitioning and 64-bit memory limitation.

Asynchronous Cache

Both chip sets connect to the 21064 processor bus, which is 128 bits wide—twice the width of the Pentium bus—and can be clocked at any integer divisor of the CPU clock. The PCI chip sets support a maximum system clock rate of 33 MHz, and the PCI clock is always the same speed as the system clock. Table 1 shows various possible combinations of CPU and system clock rates.

The second-level cache (L2 cache) is not restricted to the speed of the system clock during a cache hit. In this case, the L2-cache control is performed by the 21064 itself (see *060301.PDF*), which takes data from the processor bus based on programmable cache timing. Table 1 shows that a system with 12-ns SRAMs can read data every 27 ns, which would be 4 cycles at 150 MHz. Using the 128-bit bus, it takes only two transfers to refill the 32-byte lines of the on-chip caches.

On an L2-cache miss or DMA transfer, the CMC becomes involved. In this case, the CMC drives the address

For More Information

Price and availability for Digital's PCI chip sets is not yet available. The company expects the chip sets to sample in the fall and ship by the end of the year. For more information, contact your local Digital sales office or call the DECchip Info Line at 508/568-6868.

to the main memory and also controls writing the new data into the external cache. Memory data is returned with the requested word in the first access, so the processor can restart as soon as possible. The 21064 is capable of "streaming," executing using incoming data at the same time that it is being written into the external cache.

Using 70-ns page-mode DRAMs, it takes 5 system clock cycles (30 ns each) for the first data access and 2 cycles for subsequent accesses. With the full 128-bit memory interface, it takes only two accesses to satisfy a cache miss (5-2 pattern). If a 64-bit interface is used, the pattern is 5-2-2-2. Using 60-ns DRAMs yields a 5-1 (or 5-1-1-1) access pattern. These system clock cycles must be multiplied by the system clock divisor to calculate the processor's cache miss penalty.

External Tag RAMs and Logic

The L2 cache can be 128K to 16M in size, but a typical implementation would use seventeen 32K × 8 SRAMs for a 512K cache (including parity). Two more SRAMs are used for tags and a third for the "valid" and "dirty" flags. These chips are typically 32K × 8 for 512K or 1M caches. SRAMs for both the tag and data must have the same access times; some possible configurations are shown in Table 1. Note that 15-ns parts cannot be used with a 33-MHz system bus as they would not be able to keep up during cache-miss processing.

In the high-performance chip set, three PALs combine the cache control signals from the CPU and CMC, since the CPU controls the cache RAMs during cache hits and the CMC takes over for cache misses. This function is in the critical path for write misses and external cache probes, requiring fast (5-ns) parts. This logic is integrated into the low-cost chip set, adding an extra cycle for these activities and slightly decreasing performance.

Main memory is installed using standard SIMMs. Memory timing can be programmed to within one-half of a system clock cycle (15 ns at 33 MHz). With the wider interface, the minimum memory increment is 16M. Using the narrower interface reduces the minimum increment to 8M. While 16M is a reasonable memory granularity for workstations, Windows NT users may want to add memory in 8M chunks.

Both chip sets are standard 5V CMOS designs. The total power dissipation of the six-chip set is estimated at

about 8.5 W, or 4 W for the three-chip version.

Workstation Performance Achievable

Digital believes that the six-chip configuration will yield performance comparable to its own workstations. For example, the DEC 3000 Model 500 uses a 150-MHz 21064 processor and 512K of cache to achieve 84 SPECint92 and 128 SPECfp92. Although no measurements have been made on the unfabricated chip set, it is designed to provide the same cache access time (using 15-ns SRAMs) and the same cache miss time (using 60-ns DRAMs) as the Model 500. In fact, a configuration using 12-ns SRAMs could outperform the workstation.

A lower-cost design will not reach the same performance. Moving to a 64-bit interface and 70-ns DRAMs increases the cache-refill time from 35 CPU cycles to 60, although the time to return the critical first word remains the same. Combined with the longer write-miss penalty of the high-integration chip set, DEC believes that this design could have 15%–20% less performance than the workstation.

The new chip sets demonstrate Digital's commitment to PCI. They are intended to support all versions of the 21064, including future parts at even higher clock rates. The company hopes that, by using standard PC buses and peripherals, it can provide Alpha performance at PC prices.

DEC's forthcoming 21066 processor goes one step further by building the PCI interface directly onto the processor chip itself. In fact, the 21066 includes cache and DRAM control as well, completely eliminating the need for a system-logic chip set. This processor, due by the end of the year, could help Digital (and possibly others) bring the price of Alpha systems below \$3000. The 21066, combined with the 21064's PCI chip sets, will allow DEC to build a line of PCI-based Alpha systems stretching from this low price point to high-performance servers.

DEC Could Match Intel's Costs

Digital plans to use these PCI systems to attack the nascent Windows NT market. The company has positioned the 150-MHz 21064 against Pentium, and the price of these two CPUs is fairly similar (\$862 for the 21064 versus \$878 for the 60-MHz Pentium, both in thousands). To be competitive at the system level, Digital must also price its PCI chip sets against Pentium chip sets.

DEC claims it will price its low-cost chip set com-

petitively with Intel's PCIset and similar products from other vendors. The PCIset is priced at \$84, including a PCI-to-ISA bridge, in quantities of 1000. The Intel design also includes integrated cache tags. After adding these extra costs, the DEC product would have to be priced under \$50 to achieve the same system cost.

For minimum cost, both the Intel and DEC chip sets can be configured with standard 15-ns cache RAMs. In this setup, a 60-MHz Pentium system could deliver 50–55 SPECint92, while a 150-MHz 21064 could be rated at 75 SPECint92. (No measured benchmarks are available for either chip set.) Thus, the DEC chip set could allow system vendors to offer perhaps 40% more performance with a similar system cost, if the company delivers on its pricing promise.


For maximum performance, the six-chip set will be somewhat more expensive and requires external PALs along with the external tags. The Vitesse Pentium cache controller also requires external tags, however, and needs much faster SRAMs. Neither the Vitesse design nor Intel's zero-wait-state 82496 include a memory controller or PCI interface, which are included in the DEC chip set. Digital says it will match Vitesse's \$150 price tag, yielding about a \$100 cost advantage due to the cheaper tags and inte-

grated system logic.

With 512K of 12-ns SRAMs and a 150-MHz 21064, the Digital design should hit about 85 SPECint92, while the Vitesse cache controller should coax 60 SPECint92 out of a 60-MHz Pentium. In this comparison, the Alpha design produces a 40% performance advantage at a slightly lower system cost. Alternately, using a top-of-the-line 200-MHz 21064 and a 66-MHz Pentium, Digital could offer a 60% advantage with a slightly higher system cost.

DEC is making a smart move by joining Intel in supporting PCI, neutralizing any advantage that x86 platforms might hold in peripheral cost. The new Alpha chip sets will significantly decrease both the design time and system cost of 21064-based products. They should allow both DEC and other vendors to match the cost of Pentium systems while offering superior performance.

To succeed with this strategy, Digital must deliver on its promised chip set pricing and then continue to track Intel's Pentium pricing with its own Alpha CPU prices. Finally, Digital must cross its corporate fingers and hope that a 40% price/performance advantage is enough to convince Windows NT buyers to choose a non-Intel platform. ♦



Digital's Aaron Bauch discusses Alpha technology at the MicroSystems Forum.

MICHAEL MUSTACCH