OEMs Brace for Slot 2 Impact

Massive Module Creates Concerns for Workstation and Server Makers

by Peter N. Glaskowsky

The industry got its first look at the Slot 2 module at the Intel Developer Forum (IDF) last month. Intel has not yet announced the new name (see MPR 3/30/98, p. 1) for processors based on the Deschutes core and Slot 2 module; as the new module is twice the weight and height of a Slot 1 module and slightly wider (about 6" × 6" × 1.5" thick with heat sink), we would like to suggest “Heftium.” The most common server and workstation systems equipped with the new processor will use two Slot 2 processors in a dual retention module (DRM). Complete with heat sinks and fans, these subassemblies will weigh about four pounds.

Figure 1 shows a picture of the new module with a representative heat sink (unlike Slot 1 processors, Slot 2 modules will not be shipped with a heat sink) next to the current Slot 1 module, while Figure 2 shows a prototype DRM developed by Intel’s Workstation Product Division (WPD).

WPD’s retention module design will be used by some midrange server vendors, while other OEMs—and other groups within Intel—are developing retention modules to suit other applications. Intel’s Enterprise Server Group, for example, is working on a quad retention module, two of which will be needed in Intel’s forthcoming eight-way server design, code-named Saber, that is based on Corollary’s Profusion core-logic chip set (see MPR 9/16/96, p. 9).

Designers Must Handle Weighty Problems

Supporting 16 pounds of aluminum and plastic (plus a few grams of silicon) in an eight-way server will require OEMs to pay careful attention to mechanical engineering. The approach used with Slot 1, where the processor module is attached to the motherboard through its card guide/retention module and the motherboard is separately attached to the chassis, cannot be used with Slot 2. The bumps and shocks experienced during routine shipping and handling would apply tremendous torque to the motherboard, likely cracking it.

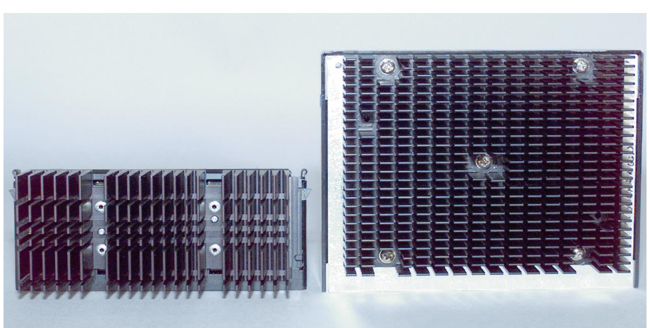

Figure 1. The Slot 2 processor module is about twice the height of the Slot 1 module and slightly wider. To permit greater OEM flexibility, Slot 2 modules will not come with a preassembled heat sink.

Instead, the hefty Slot 2 retention modules must be attached through the motherboard to the chassis itself. WPD’s solution uses four #6 screws, capable of withstanding the brief but intense (up to 50-g) shocks that systems can experience during rough handling. The chassis itself must be unusually sturdy; Intel’s tests show that mechanical stresses caused by such shocks can exceed 800 psi.

To the designers of today’s relatively svelte desktop PCs, these may be daunting problems, but high-end workstation and server manufacturers are already accustomed to dealing with heavy machines, giving them a leg up in supporting the weight of Slot 2. Some deskside systems already tip the scales at more than 200 pounds, making even quad-processor Slot 2 retention modules seem petite by comparison.

Thermal Management a Hot Topic

The heavy aluminum heat sinks on Slot 2 modules are there for a good reason: between the 450-MHz Deschutes processor core and four custom full-speed SRAMs (see MPR 3/9/98, p. 9), each Slot 2 processor can produce up to 37 W of heat, and future Slot 2 CPUs will run still hotter. Even a two-pound aluminum heat sink can’t dissipate this much thermal energy without a stiff wind blowing: at 45° C, Slot 2 systems must provide about 450 linear feet per minute (LFM) of air, roughly nine times the airflow in today’s mainstream PCs.

Slot 2 systems may need special ducting within the case to bring this much air to the processors. Ducting alone won’t solve the problem, however. To ensure enough airflow across the heat sink of the second processor, Intel’s own DRM is designed to support up to four 40-mm fans to force cooling air between the two modules.

The thermal challenges don’t stop with the processors. The forthcoming 450NX chip set, expected to ship in 3Q98 with support for four-way multiprocessing and up to 8G of SDRAM, will itself consume about 10 W. The 450NX north bridge will require its own heat sink and 125 LFM of cooling air. Intel’s recommended motherboard layout places this chip under the primary CPU’s heat sink, necessitating a low-profile heat sink and careful airflow management.

With the CPU core running at 400 MHz or more—not to mention even higher Merced and Direct RDRAM frequencies in the future—keeping EMI inside the system enclosure is becoming more difficult than ever. In the past, designers have been able to vent enclosures by using simple perforations in the sheet metal, but the multi-GHz harmonics of these higher operating frequencies can escape through such openings. Vents in new systems may use mesh screens to prevent the escape of EMI, but these screens will interfere with airflow and may mandate the use of more powerful fans to overcome the higher back pressure.

Such fans must produce enough airflow to remove more than 400 W of heat from the enclosure on dual-processor workstations equipped with 2G of SDRAM, multiple high-speed SCSI hard disks, and several power-hungry AGP and PCI expansion cards. Saber-based servers—with up to eight CPUs and more memory, hard disks, and other peripherals—may produce more than 1.2 kW of heat and will need a proportionate increase in airflow.

Vent screens and stronger fans will also make systems audibly noisier. This isn't a problem for servers that are usually consigned to some out-of-the-way closet, but the 50-dB noise levels of today's typical desktops will seem tomb-like compared with the 60–70 dB we expect will be emitted by Slot 2-based workstations. Systems intended for deskside use, especially in multimedia content creation, will require substantial acoustic baffling to reduce acoustic noise output.

Cool Graphics, Hot Graphics Cards

Yet another source of heat in high-power Slot 2 workstations will be new AGP Pro graphics cards. While AGP Pro uses the same 2× or 4× AGP electrical interface as mainstream AGP cards, it uses a new AGP connector capable of delivering a staggering 110 W of electrical power to AGP Pro boards—more than six times the power consumption of mainstream 3D-graphics cards.

Intel plans to release the AGP Pro specification in 2Q98, defining a new form factor as well as layout and cooling guidelines. AGP Pro cards will support multichip 3D accelerators for CAD and visualization applications as well as multimonitor solutions, designs that are simply too large to fit on cards designed to the current AGP specification.

AGP is facing other difficulties. Although Intel has described a way to provide upgradable 2× AGP graphics on Pentium II motherboards, this three-load AGP solution—connecting the chip set, one AGP device on the motherboard, and one AGP slot—is not compatible with the 4× AGP specification. The faster mode requires the use of 1.5-V signaling levels and can tolerate only two electrical loads; OEMs must choose between integrated graphics and an AGP slot. If three-load 2× AGP becomes part of the mainstream PC feature set, end users are likely to be disappointed by the lack of this capability in the 4× AGP standard.

PCI Facing Performance, Power Problems

Even more serious problems are looming for the venerable PCI expansion bus, which was impressive when it arrived five years ago but is now showing its age. Apart from some minor protocol enhancements, PCI hasn't changed much over the years. Today's 333-MHz Pentium II PCs use the same 5-V, 32-bit, 33-MHz PCI specification found in the first 66-MHz Pentium systems.

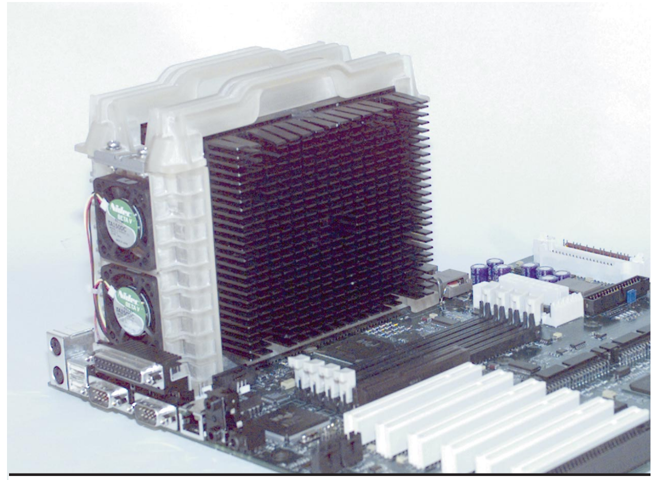

Figure 2. This prototype dual-processor Slot 2 retention module from Intel's Workstation Product Division weighs about four pounds with CPUs and fans installed.

The limited bandwidth and protocol overhead of PCI prompted Intel to develop the AGP interface, but AGP can't be used for mass-storage and network-interface peripherals. The current PCI bus simply isn't fast enough for high-end SCSI RAID controllers or Gigabit Ethernet adapters, peripherals that are already required by high-end workstations and servers and will eventually be needed for mainstream systems as well. Some vendors offer 64-bit PCI slots, others offer 66-MHz PCI, and a few will combine both of these advanced capabilities, but these solutions are basically stopgaps that don't help mainstream users at all, and they aren't really adequate for the high-end market either.

The 5-V signaling on mainstream PCI peripherals is also causing problems for ASIC vendors, which are finding it more difficult to provide 5-V tolerant I/O on new 0.35-micron chips. This problem will only get worse as vendors move to 0.25-micron technology for new 3D accelerators and other advanced peripheral chips.

IDF would have been the ideal venue for Intel to announce a third-generation PCI bus, perhaps with faster clock rates and 1.5-V signaling of AGP but preserving the 10-load capability of 33-MHz PCI, but Intel did not. Sources indicate that the company has begun development of an advanced PCI replacement, though its features and performance remain unknown. At this point, it is unlikely that a new bus could be ready in time for 1999 systems. This means the current PCI bus, with all of its shortcomings, will remain with us for at least the next two years.

Instead, Intel's focus for 1998 and 1999 is on technology for high-end systems. The new Slot 2 modules, chip sets, AGP Pro graphics cards, and related thermal and mechanical improvements are worth all the grief they are sure to cause OEMs. They will finally allow makers of x86-based Windows NT systems to build the "big iron" they need to compete on an equal footing with today's Unix-based workstation and server systems. M