Intel 840 Brings RDRAM to Workstations

New Chip Set Expands on 800-Series Architecture With More Memory and I/O

by Peter N. Glaskowsky

With twice the memory capacity and bandwidth of Intel's 820 chip set (see MPR 10/6/99, p. 30) and an optional second PCI interface, the new 840 chip set is Intel's first workstation core-logic product to offer clear advantages over its desktop-PC cousin. These advantages—and the premium price of Direct RDRAM—make the 840 the best choice for workstation buyers who value performance above all else. The new chip set will also find use in a few servers, but workstations are its primary market.

As Figure 1 shows, the 840's memory controller hub (MCH) includes two interfaces not found on the 820—an extra RDRAM channel and a port that connects to the new P64H PCI controller. The P64H supports a 64-bit PCI bus, marking the (long overdue) first appearance of 64-bit PCI on an Intel workstation chip set. This bus can be configured for 33-MHz or 66-MHz operation. At the latter clock rate, only two slots are allowed. At 33 MHz, four slots may be used.

Memory Capacity, Speed Doubled

Though the 840's second Direct RDRAM interface doubles the theoretical bandwidth of the 820's memory subsystem, its effect on memory capacity is more significant. The 32-device capacity of each RDRAM channel limits Intel's 820 to a maximum 512M using today's 128-Mbit RDRAM chips. Many users need more than 512M, especially those in the publishing, video-production, and computer-aided design industries. The 840's dual channels support 1G of 128-Mbit RDRAMs, and using the 256-Mbit chips expected to arrive in 1H00, the 840 can handle up to 2G—the same as Intel's earlier 440GX (see MPR 7/13/98, p. 11) and enough for almost all workstation applications. (Intel says the RDRAM-related signal-integrity problem that delayed the introduction of the 820 chip set does not affect the 840.)

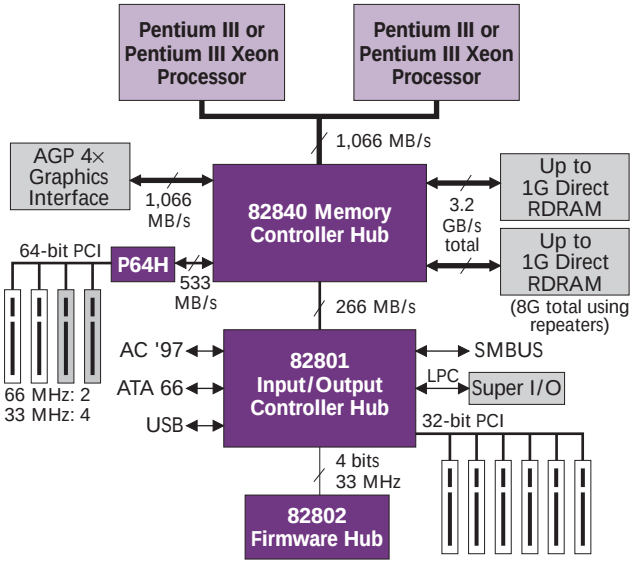

Figure 1. The Intel 840 workstation chip set supports two Direct RDRAM channels and two PCI buses, doubling the memory and I/O capabilities of Intel's 820 desktop chip set.

The optional 82803 RDRAM memory repeater hub (MRH-R) converts each RDRAM channel into two channels, for a total of four channels and 4G of RDRAM. Intel will also offer the 82804 SDRAM repeater hub (MRH-S), which supports up to 4G of PC100 SDRAM with a peak bandwidth of 1.6 GBytes/s. The repeaters may be used to create memory expansion cards equipped with either RDRAM or SDRAM plus the appropriate type of repeater.

The 840's memory channels are run in parallel and must be configured identically. The parallel operation reduces access latency, since an entire cache line may be moved across the bus in 8 cycles, versus 16 for the single-channel 820. This 10-ns reduction should have a significant impact on application-level performance.

The second channel may also bring down the effective cost of 840-based systems. Rather than equipping the 840 with the fastest available 400-MHz RDRAM, OEMs may choose to use slower and less-expensive RDRAM devices, or even cheaper 100-MHz SDRAM. Using 300-MHz RDRAMs instead of 400-MHz chips to fill out a 1G system could save hundreds of dollars. Two channels of 300-MHz RDRAM offer about 2.2 GBytes/s of sustained throughput, more than enough for most users' needs.

The SDRAM option would save even more money, perhaps as much as $2,000 in component costs for a 1G system. Such a configuration would offer about 1.0 GBytes/s, less than can be achieved with RDRAM but still twice the 500-MBytes/s memory bandwidth of the 440GX.

The money saved by the move to slower DRAM would more than pay for the 840's price premium over the 820. The 840's superior I/O would come free. From this perspective, the 840, at just $60, may be a more economical way than the 820 to achieve any desired level of performance.

AGP 4× Adapts to Workstation Needs

Though the 840's 4× AGP interface is identical to that of the 820, many workstations designed with the 840 will implement the optional AGP Pro form-factor specification, developed to support multichip professional 3D cards.

AGP Pro uses the same logical interface, but AGP Pro slots include extra power-supply pins and provide a fixed

physical relationship to two nearby PCI slots. These characteristics allow use of multiboard modules that can consume much more power than a single AGP graphics card. An AGP Pro module is allowed to consume up to 110 W compared with the 25 W of conventional AGP. Standard AGP graphics cards may still be used in AGP Pro slots.

Faster PCI Needed for New I/O Standards

Other workstation chip sets, such as those developed for Digital's (now Compaq's) Alpha systems, have supported 64-bit PCI for years. Intel's previous core-logic products, including the 440GX (see MPR 7/13/98, p. 11), lacked this feature, making them less suitable for use in demanding workgroup environments, where Gigabit Ethernet, SCSI RAID arrays, and other fast I/O interfaces are often found.

Intel could have added a second PCI interface to the existing I/O controller hub (ICH), but this would have overloaded the 266-MByte/s I/O hub interface. Instead, Intel added a second I/O hub interface to the MCH. The new hub interface is twice as wide (16 bits) as the primary I/O hub interface and provides 533 MBytes/s of peak bandwidth.

A new chip, the P64H, translates between the new I/O hub interface and a 64-bit 66-MHz PCI bus. The P64H is an optional part of the 840 chip set. It will be used on most high-end systems but will not be included on some entry-level workstations.

We believe the second I/O hub interface may be used for other purposes in the future, though Intel has not announced any plans to do so. Intel uses a similar concept on its 450NX (see MPR 7/13/98, p. 11) to connect PCI exchange bridge (PXB) chips. Two PXB chips are available—one with a single 64-bit bus and another with two 32-bit buses. Intel could offer a similar selection for the 840 without redesigning the 840 MCH chip.

The new port could also be used to connect other high-performance peripheral interfaces to the 840. The System I/O specification (see MPR 9/13/99, p. 4), now under development, could be adapted to the 840 in this way. System I/O provides about 500 MBytes/s of bandwidth per pair of unidirectional connections, matching the capacity of the 840's P64H interface.

One limitation of the 840's I/O architecture is that peer-to-peer PCI transactions between the two PCI buses are not supported. This type of bridging is not needed for workstations, but it would have improved I/O throughput in 840-based servers using Intelligent I/O (I_2O) controllers.

The ICH and firmware hub (FWH) in the 840 set are the same chips used in the 820 and 810/810E, bringing the same advantages to the 840 platform. Because the MCH and ICH are linked by a 266-MByte/s hub-to-hub link, IDE hard-disk I/O coexists peacefully with the ICH's PCI bus rather than fighting for bandwidth.

Though the ICH indirectly supports ISA-bus slots in 820-based systems, Intel says none of its customers plan to include ISA slots in 840-based machines. Accordingly, Intel will not validate ISA-slot compatibility with the 840. The 840 is thus the first Intel chip set to eliminate ISA support—but not, we trust, the last.

Price & Availability

The 840 chip set is available now and is priced at $60 in quantity. Intel has not released pricing for the optional MRH and P64H chips. More information is available online at *http://developer.intel.com/design/chipsets*.

Server Applications Optional

The 840's impressive memory and I/O bandwidth appear to make it a good choice for dual-processor servers. Unfortunately, the high cost of RDRAM memory will delay acceptance of the new chip set in the server arena. The 840 also lacks some features now considered important for servers, notably hot-plug PCI and peer-to-peer transactions between its PCI buses. Hot-plug support allows PCI cards to be removed and replaced without shutting down the system, a capability that helps improve server uptime but is not especially important for workstations. Finally, Intel's initial Coppermine-based Pentium III Xeon processors (see MPR 10/25/99, p. 1) have just 256K of L2 cache, which is insufficient for some server applications.

Because of these issues, workstations are Intel's initial focus for the 840. For the next year or so, Intel's 450NX (see MPR 7/13/98, p. 11) and Reliance's Champion 3 (see MPR 9/13/99, p. 16) chip sets will manage most new Pentium III two- and four-processor servers. We expect, however, that some OEMs will identify specific target markets within the server space, where the 840's unique features make sense.

The 840 supports Pentium III and Pentium III Xeon processors with bus speeds up to 133 MHz. High-end workstations and most servers will be designed for Slot 2 Xeon modules, which offer slightly better reliability and scalability, due to the improved thermal characteristics of the Xeon module and its use of on-board voltage regulation. Servers will need Slot 2 to support future Xeons with more L2 cache. Less-expensive 840 motherboards will include 370-pin sockets for Pentium III instead.

840 Clearly Differentiates PC Workstations

Intel's previous workstation chip sets were based on very similar desktop PC chip-set designs. The 440GX, for example, was a 440BX with one extra DRAM address bit. This similarity made it difficult for OEMs to establish a clear distinction between PCs and PC workstations.

The 840 solves this problem entirely, enabling a range of systems with real advantages over mainstream desktops. OEMs will have no trouble communicating these advantages to their customers. Buyers may wince at the price tags on 840-based systems, but users will rejoice at the unparalleled performance available from the new chip set. M