New Chip Sets Embellish Xeon

The 450NX PCIset and 440GX AGPset Target Servers and Workstations

by Keith Diefendorff

In concert with the Xeon announcement (see MPR 7/13/98, p. 1), Intel has announced two new chip sets to support that processor in the server and workstation markets. The 450NX PCIset is a complete redesign of Pentium Pro's 450GX for Xeon-based servers up to four-way SMP. The 440GX AGPset is a minor tweak to the 440BX (see MPR 4/20/98, p. 18) for Xeon-based workstations. Future eight-way SMP systems will use the Corollary ProFusion chip set.

The 450NX Gives Xeon Four-Way SMP

The 450NX improves on the 450GX in a number of ways, including support for the Pentium II's 100-MHz system bus, support for larger memory, and support for 64-bit and multiple PCI buses. Like the 450GX, the NX forgoes AGP, as high-speed graphics is not a requirement for servers. Instead, the NX provides multiple PCI channels to better support the heavy disk-I/O bandwidth needed in servers.

The NX chip set comprises four different types of chips, as Figure 1 shows. The memory and I/O controller (MIOC) attaches to the system bus and provides the memory interface port and two 64-bit ports for two PCI exchange bridge (PXB) chips. Each PXB can provide either two 32-bit PCI buses or one 64-bit PCI bus. The PCI buses operate at 33 MHz. The 450NX supports Xeon's AGTL+ system-bus signaling.

Included in the NX chip set are three components for implementing the memory card: the RCG (RAS/CAS generator) and two data multiplexer chips (MUXs).

A significant new feature of the NX is its support for 8G of memory, whereas the 450GX supports only 4G. The NX's memory controller can drive two memory cards, each of which can hold 16 EDO DIMMs. The DIMMs operate at 100 MHz and are interleaved four ways.

The choice of EDO over SDRAM for a new chip set is somewhat surprising, but the 450NX's interleaved EDO design actually provides 25% more bandwidth than the 800 Mbytes/s provided by the noninterleaved open-page SDRAM architecture in 440GX. Bandwidth being the most important resource in a server environment, interleaved EDO is a good choice. In addition, large capacity EDO DIMMs are less expensive and more readily available than SDRAM—both important attributes when building very large memory systems.

The NX's MIOC and PXBs are provided in 540-pin PLGA packages. The complete NX chip set—including MIOC, two PXBs, and RCG/MUXs for two memory cards—lists for $409 in 1,000-unit quantities.

The 450NX will also be delivered in a stripped-down "basic" version that lists for $232. The basic version uses the same MIOC but comes with only a single PXB and RCG/MUXs for one memory card.

Two Ways to Go Eight-Way

The 450NX supports SMP with up to four processors. The NX also supports multiprocessing above four processors using a clustered-NUMA (nonuniform memory access) approach through the 450GX-style cluster connector.

The NUMA architecture uses bus bridges to connect multiple SMP clusters, each of which has its own local memory. This approach puts each cluster in a separate address space, with processors seeing different access times for local and remote memory—hence the name nonuniform memory access. In clustered-NUMA systems, cache coherence across the cluster boundary is not enforced by hardware.

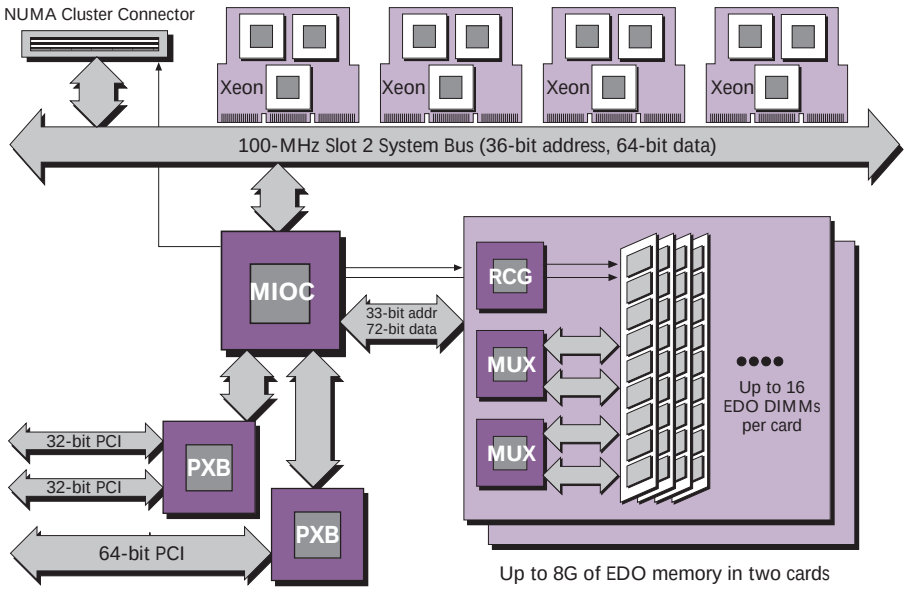

Figure 1. The 450NX attaches to the Xeon system bus and provides a port to main memory and two ports to PCI exchange bridges (PXBs). The system bus operates at 100 MHz for four-way SMP, but when a cluster bridge is used, the system bus is limited to 90 MHz. The four-way interleaved EDO memory is implemented on two cards, each with up to 16 DIMMs. The RCG and MUX chips are supplied with the chip set.

Price & Availability

The 450NX PCIset, including one MIOC, two PXBs, two RCGs, and four data MUXs, has a list price of $409 in 1,000-piece quantities.

The 450NX Basic PCIset, including one MIOC, one PXB, one RCG, and two data MUXs, lists for $232 in 1,000-piece quantities.

The 440GX AGPset comes with a PIIX4E south bridge and has a 1,000-piece list price of $63.

The 440GX is available now for volume delivery. The 450NX and 450NX Basic will be available in July.

For more information, access *http://developer.intel.com/design/pcisets*.

The NUMA architecture divides the system bus into independent segments that together provide more aggregate bandwidth than a single shared bus. This organization reduces the bus utilization that is seen by each processor, which, according to queuing theory, can dramatically improve performance. NUMA architecture, however, forces software to use an explicit message-passing system to communicate between CPUs in remote clusters. Such protocols can have high overheads and can force extra memory-copy operations that waste time and horsepower.

NUMA organizations are effective for applications where the parallelism is coarse-grained and where little sharing or communication needs to occur between processors. NUMA is less effective in applications where processors need to be tightly coupled, for example, where all the processors must cooperate on a single data set. In general, the asymmetry of the NUMA hardware makes it difficult for an operating system to schedule computing resources efficiently.

In symmetric multiprocessing, on the other hand, all CPUs share the same address space, and all see uniform memory access times. Cache coherence across all CPUs is enforced by hardware. These factors allow software to use a simpler shared-memory interprocessor communication protocol, which can have performance advantages over the message-passing protocol in many applications. The hardware symmetry also makes resource scheduling and load balancing easier for the operating system.

The 440NX does not support SMP above four-way. Eight-way SMP systems, however, are already available from several RISC vendors, and this technology will be important if Xeon is to compete with them in the high-end server markets. To get eight-way SMP technology, Intel last year bought a company named Corollary (see MPR 9/16/96, p. 9) for its Profusion architecture. Intel has not said when it will introduce the Profusion technology, but we expect it to be sometime late this year. Sources tell us, however, that Profusion will be offered only as a full motherboard solution, not as a chip set. Intel may also license the board design to OEMs, but Intel will not support changes to the design.

Chip Set for Workstations Is Less Compelling

The 440GX is a more modest enhancement to 440BX. In fact, the 440GX is built on the 440BX core, and the only change is support for larger memory (2G vs. 1G). The GX retains the open-page SDRAM architecture of the BX, which provides better average memory latency under typical workstation and PC workloads than does the 450NX's interleaved EDO.

The 440GX comes in the same 492-pin BGA package as the 440BX and uses the same PIIX4E south bridge. The 440GX lists for $63 in 1,000-piece quantities, 21% more than the $52 Intel is charging for the nearly identical 440BX. As Table 1 shows, there is no obvious reason, other than their intended markets, that these parts should have any price difference at all.

The value of 440GX is unclear. The only advantage over the 440BX appears to be its support for 2G of memory. While a few workstation vendors may want this, for Intel to come out with a new chip set just for this feature seems questionable. A much more aggressive part was in order.

On the other hand, the 450NX is a clear upgrade from Pentium Pro's 450GX for the server market. The 450NX provides four-way SMP, larger memory, and significantly more I/O and memory bandwidth—all critical to Xeon's success. Its $409 price seems high but it is appropriate for systems carrying several thousand dollars worth of CPUs. M

Feature	Servers			Workstations	
	450GX	450NX Basic	450NX	440BX	440GX
Max CPUs	4 CPUs	4 CPUs	4 CPUs	2 CPUs	2 CPUs
System Bus	66 MHz	100 MHz	100 MHz	100 MHz	100 MHz
Mem Speed	66 MHz	100 MHz	100 MHz	100 MHz	100 MHz
Mem Type	FPM	EDO	EDO	SDRAM	SDRAM
Mem Arch.	4:1 Interleave	4:1 Interleave	4:1 Interleave	Open-page	Open-page
Mem B/W	533 Mbytes/s	1 Gbytes/s	1 Gbytes/s	800 Mbytes/s	800 Mbytes/s
Max Mem	4G	4G	8G	1G	2G
DIMMs	16 DIMM	16 DIMM	32 DIMM	4 DIMM	8 DIMM
ECC	Yes	Yes	Yes	Yes	Yes
Scrubbing	None	Scrub-on-error	Scrub-on-error	Scrub-on-error	Scrub-on-error
32-bit PCI	1	2 or 0	4 or 2 or 0	1	1
64-bit PCI	0	0 or 1	0 or 1 or 2	0	0
AGP	No	No	No	2×	2×
Clusters	NUMA	No	NUMA	No	No
Signals	GTL+	AGTL+	AGTL+	GTL+	GTL+
Package	352 BGA	540 PLGA	540 PLGA	492 BGA	492 BGA
USB	No	Yes	Yes	Yes	Yes
List Price	$391	$232	$409	$52	$63

Table 1. The 450NX PCIset replaces the 450GX used with the Pentium Pro to serve the high-end server market. The 440GX AGPset is a minor enhancement of the 440BX to support larger memory in high-end workstations. (Source: Intel)