The Buses and Networks of the Pentium 4

Figure 8.11 shows the I/O system of a PC based on the Pentium 4. The processor connects to peripherals via two main chips. The chip next to the processor is the memory controller hub, commonly called the *north bridge*, and the one connected to it is the I/O controller hub, called the *south bridge*.

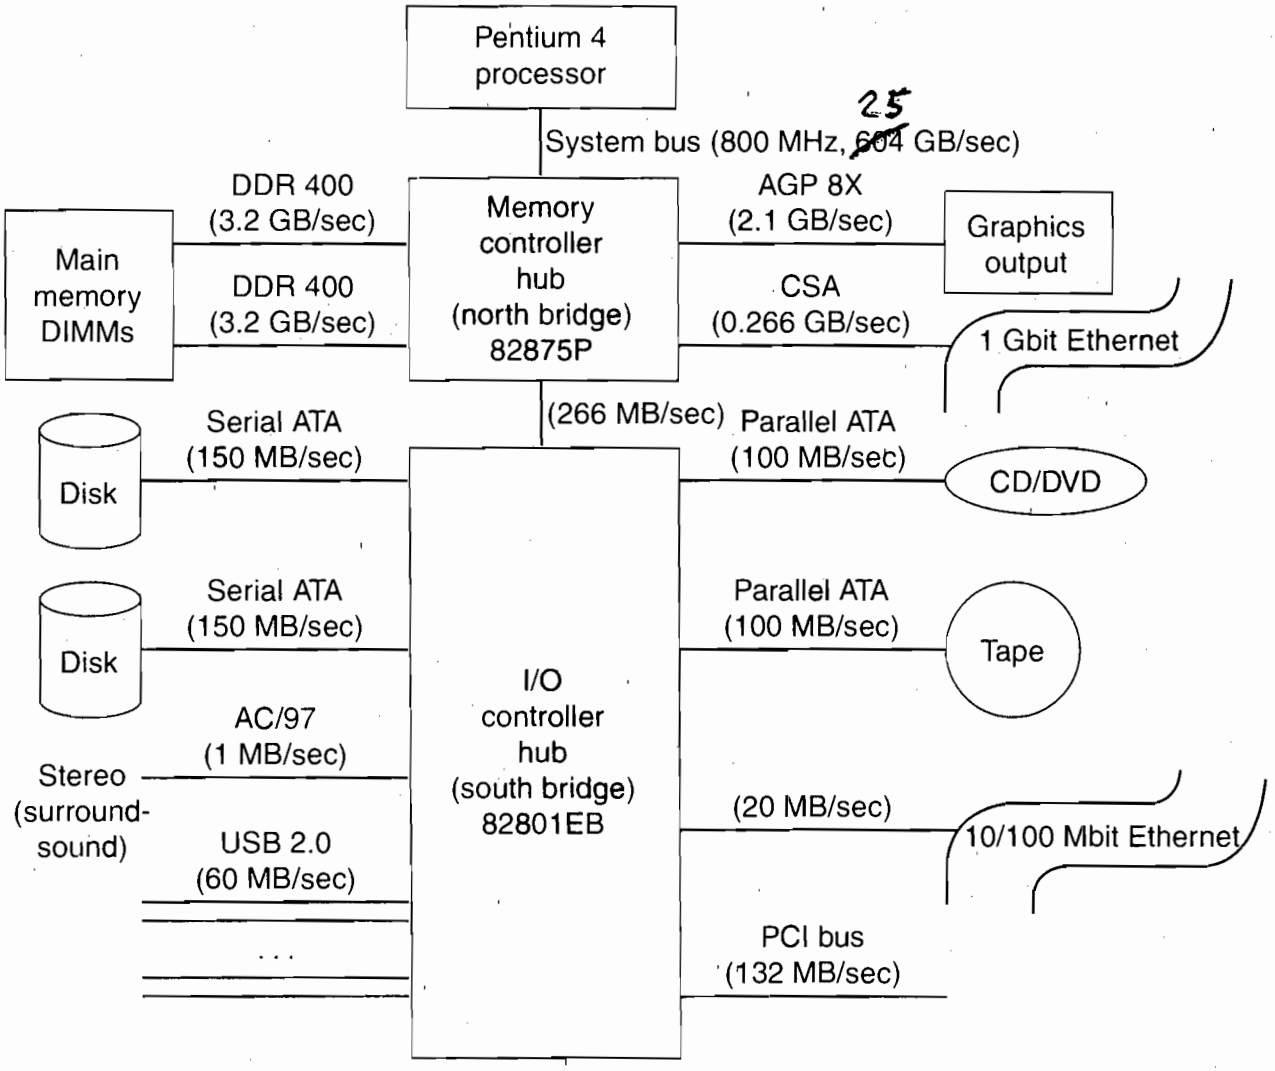

FIGURE 8.11 Organization of the I/O system on a Pentium 4 PC using the Intel 875 chip set. Note that the maximum transfer rate between the north bridge (memory hub) and south bridge (I/O hub) is 266 MB/sec, which is why Intel put the AGP bus and Gigabit Ethernet on the north bridge.

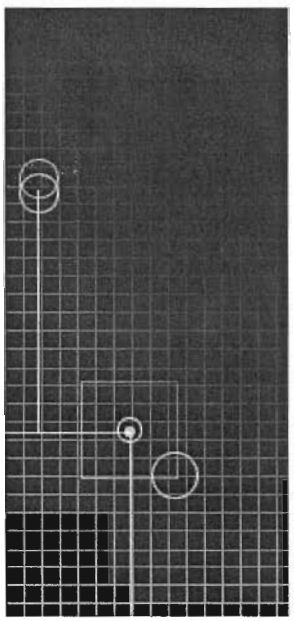

General Information
February 6, 2004 • Vol.26 Issue 6
Page(s) 18 in print issue

What's The Difference?

PCI Extended vs. PCI Express

PCI. PCI-X 1.0. PCI-X 2.0. PCMCIA. PCI Express. Do you know the difference?

The truth is, most of us don't, despite the fact that the local bus is one of the most important parts of modern computing, crucial to the speed of client/server systems. And worse, nearly all of us confuse PCI-X, short for PCI Extended, with PCI Express, the newest implementation of the PCI bus and a true revolution in the way data is moved from peripheral to processor.

■ In The Beginning Was PCI

To understand PCI-X and PCI Express, you need to understand plain old PCI, or Peripheral Component Interconnect. It's the dominant local bus used in servers, clients, communications, and embedded applications, not to mention newer versions of the Macintosh. And it's the base that PCI-X and PCI Express are built on.

Developed by Intel in 1992 to improve on the ISA bus, PCI is a 64-bit technology that's usually implemented as 32-bit. It can run at clock speeds of 33MHz or 66MHz. At 33MHz, 32 bits translates into a throughput of 133Mbps.

And while PCI-X and PCI Express are both improvements on PCI, don't expect PCI to go anywhere anytime soon. In much the same way that PCI and ISA slots can often be found side by side on the same machine, neither PCI-X nor PCI Express is designed to replace PCI.

■ PCI-X

Often confused for PCI Express, PCI-X stands for PCI Extended, a standard designed jointly by HP, IBM, and Compaq to increase the performance of high-bandwidth devices, such as Gigabit Ethernet and Fibre Channel, as well as processors that are part of a cluster. It's useful to think of it as an enhanced form of the PCI bus.

PCI-X is fully backward compatible with PCI and improves on the throughput of PCI from 133Mbps to 1Gbps. It comes in two flavors: PCI-X 1.0, with a maximum clock speed of 133MHz, and PCI-X 2.0, which tops out at 533MHz. (In fact, PCI-X 2.0 has four "speed grades": PCI-X 66, PCI-X 133, PCI-X 266, and PCI-X 533, which run at 66MHz, 133MHz, 266MHz, and 533MHz, respectively; which grade is used depends

on the application of the bus.)

533MHz
x 32 BITS
17,056 Mbps
÷ 8 BITS PER BYTE
= 2,132 MBps

But 2.0 also has other refinements beyond speed, including ECC (Error Correction Code), to improve the robustness of the interface; 1.5V signaling and source synchronous strobes, which improve performance so it can run at 533MHz; device ID messages, which are designed to enable a new class of peer-to-peer transfer applications; and a 16-bit version for embedded applications.

Servers and workstations are the initial targets for PCI-X 2.0. InfiniBand, 10 Gigabit Ethernet, 10 Gigabit Fibre Channel, and multifunction cards that have large bandwidth requirements will benefit from the PCI-X 266 and PCI-X 533 flavors. The PCI-SIG, the ultimate authority on the use and deployment of the technology, is mum on whether PCI-X 2.0 will make it to the desktop segment of the market. That, it says, will be determined by the needs of individual developers and their products.

■ PCI Express

In the past few years, a slew of new technologies, such as 10GHz processors, faster types of RAM, high-performance graphics, multiple-Gbps networking, fast storage, and more have begun to demand a faster interconnect, one with greater bandwidth to service their hungry needs. Enter PCI Express, a high speed, serial, point-to-point, hot-pluggable, and hot-swappable system bus that's fully compatible with PCI, at least on the software level.

2,000 MHz
32 BITS
64,000 Mbps
÷ 8 BITS PER BYTE
= 8,000 MBps
ONE WAY

Each point-to-point interconnect in the PCI Express architecture can have 1, 2, 4, 8, 12, 16, or 32 dual simplex 2.5Gbps lanes (with a 2.0Gbps effective rate), which translates into scalable bandwidth up to 128Gbps (or 16GBps) between nodes.
↑ 8 ONE WAY

PCI Express is also called 3GIO (Third Generation I/O) and occasionally Arapahoe, after the IEEE working group where it was first developed as a joint project with Intel at the head. In early 2002, ownership of the technology was transferred to the Portland-based PCI-SIG where it was renamed PCI Express Architecture. Products with PCI Express are expected sometime in the early part of this year (with initial silicon in the second half of 2003).

PCI Express is not intended to replace PCI or PCI-X. Rather, the PCI Special Interest Group will continue to support both standards because both have attributes that are needed on a per-application basis. As proof of this, consider the fact that PCI/PCI-X and PCI Express slots will coexist in the same machines in the future and be used for different types of interconnects. (PCI Express is, however, intended to replace AGP.)

Some key features of PCI Express include software compatibility with PCI, meaning a smooth transition to new hardware can be made with software that leverages the new, advanced features of PCI Express, and optimum bandwidth per pin, to enable unique and small form factors, simplify the design of motherboards and the way they route signals, and reduce problems with signal integrity. There's also native hot-plug/ hot-swap capacity and native power management capacity.

All in all it's a powerful and promising technology, the exact adoption of which remains to be seen. But along with PCI-X, it's sure to make server rooms and data centers faster the world over. ■

by David Garrett

SATA

Serial ATA

Next Generation Storage Interface

By Mark Kyrnin

Related Resources

- Hard Drive Manufacturers
- Net Reviews of Hard Drives

From Other Guides

- Serial ATA For Your Next Hard Drive

Elsewhere on the Web

- Serial ATA Working Group

Stay Up-to-Date with PC Hardware

Email:

subscribe

Hard drive and optical storage devices for the personal computer market have been using variations of the IDE and ATA standards for over 10 years now. It has undergone a great many revisions over the years and is finally reaching the limits of its capabilities. To address the limitations of the ATA interface, the new Serial ATA interface has been under development for several years and is now finally starting to become available to consumers in the forms of controllers and more importantly hard drives. So lets take a look at this new interface and what its capabilities are compared to the older ATA format.

Serial vs. Parallel

The fundamental difference between the two formats is how the data is transferred between the device and the processors. Traditional ATA devices and controllers use a parallel data transfer mechanism. Parallel processing is a fairly common technique where multiple channels of data are sent simultaneously to try and increase the amount of data transferred in a single clock cycle. In the case of the ATA/100 standards used by today's IDE drives and controllers, they send the data across a 16-bit channel. The problem with this type of mechanism is the number of wire required to transfer that data. This is why the ATA cables are so wide. It is necessary to have the 40 or 80 wires required to transfer the data. The problem with this is the interference caused between these wires. At higher clock speeds necessary for faster speeds, the interference between the wires is too great to allow for reliable transmission.

Over the last couple of years, many advances have been developed in serial transmission techniques. Specifically through the development of the Universal Serial Bus interfaces. Serial transmissions run across a single control channel compared to the multiple channels of a parallel interface. This means that at the same clock speeds, the serial line will carry less data, but because the serial method requires fewer wires, less interference is generated to cause data integrity problems. This allows for serial transmission methods to run at much higher speeds than the equivalent parallel methods. In the case of the first Serial ATA standard, the clock runs at 1500 MHz compared to a clock rate of 50 MHz of the ATA/100 standard.